NetWitness® Platform XDR

Version 12.0.0.0

Release Notes

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

What's New **5**

Upgrade Paths 5

Security Fixes 5

Product Version Life Cycle for NetWitness Platform 5

Enhancements 5

Policy Based Centralized Content Management 6

Springboard 7

Enhanced Springboard to Support New Built-in Panels 8

Create Custom Springboard at the User Level 8

Automated Custom Springboard from Query 9

Respond 9

Incident Workflow Enhancements 10

Incident Details View Enhancements 10

Incident Overview Panel Enhancements 10

Investigation 10

Indicators for Searchable Meta 11

Unified Discovery and Interaction of Events Metadata 11

Enhanced Querying on Events View to Exclude any Specific Meta 13

View Encrypted Data in Decrypted Format 14

Select Custom Date and Time Range in the Events View 14

User Interface 15

NetWitness User Interface Enhancements 15

Detection of removable Storage Devices 16

Block Multiple File Hashes Using an Imported File 16

Support for Arm-based Windows Machines 16

Download MFT from Multiple Hosts in One Step 16

Customizable Maximum File Download Limits 17

Redesigned Alert Details View for Endpoint Alerts in Respond 17

Concentrator, Decoder, and Log Decoder Services 18

Log Parsing Enhancements 18

Enhanced Network Decoder to Decrypt Incoming TLS 1.3 Packets 18

Event Stream Analysis (ESA) 18

Improved ESA Rules Deployment 18

Reports 19

Build Rule View Enhancements 19

Fixed Issues **21**

Administration Fixes 21

Endpoint Fixes 21
Respond Fixes 21
Core Services (Broker, Concentrator, Decoder, Archiver) Fixes 22
Reporting Engine Fixes 22
Log Collection Fixes 22
End of Life Functionality 23
End of Life Functionality and Features in 12.0.0.0 or higher releases 23
Product Documentation 24
Feedback on Product Documentation 24
Getting Help with NetWitness Platform 25
Self-Help Resources 25
Contact NetWitness Support 25
Build Numbers 26
Revision History 28

What's New

The NetWitness 12.0.0.0 release provides new features and enhancements for every role in the Security Operations Center.

Upgrade Paths

The following upgrade paths are supported for NetWitness 12.0.0.0

- NetWitness 11.6.0.0 to 12.0.0.0
- NetWitness 11.6.0.1 to 12.0.0.0
- NetWitness 11.6.1.0 to 12.0.0.0
- NetWitness 11.6.1.1 to 12.0.0.0
- NetWitness 11.6.1.2 to 12.0.0.0
- NetWitness 11.6.1.3 to 12.0.0.0
- NetWitness 11.6.1.4 to 12.0.0.0
- NetWitness 11.7.0.0 to 12.0.0.0
- NetWitness 11.7.0.1 to 12.0.0.0
- NetWitness 11.7.0.2 to 12.0.0.0
- NetWitness 11.7.1.0 to 12.0.0.0
- NetWitness 11.7.1.1 to 12.0.0.0

For more information on upgrading to 12.0.0.0, see Upgrade Guide for NetWitness 12.0.0.0

Security Fixes

For more information on Security Fixes, see Security Advisories.

Product Version Life Cycle for NetWitness Platform

See for Product Version Life Cycle for NetWitness Platform a list of versions that reach End of Primary Support (EOPS).

Enhancements

The following sections are a complete list and description of enhancements to specific capabilities:

- Policy Based Centralized Content Management
- Springboard
- Respond
- Investigation
- User Interface
- Endpoint Investigation
- Concentrator, Decoder, and Log Decoder Services
- Event Stream Analysis (ESA)
- Reports

To locate the documents that are referred to in this section, go to the NetWitness Master Table of Contents.

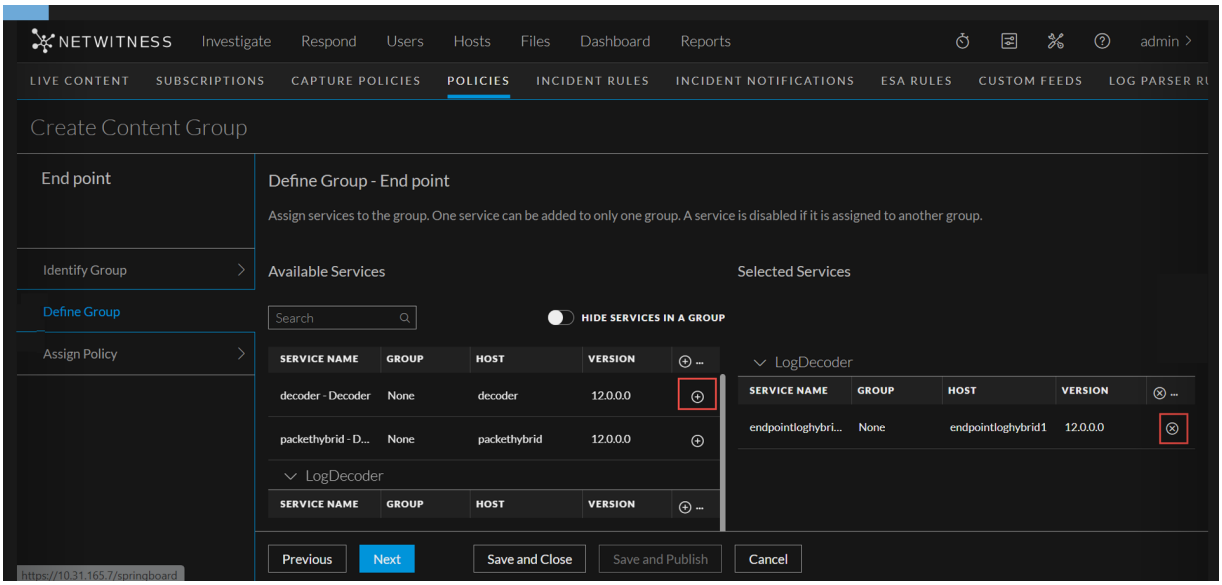
The Product Documentation section has links to the documentation for this release.

Policy Based Centralized Content Management

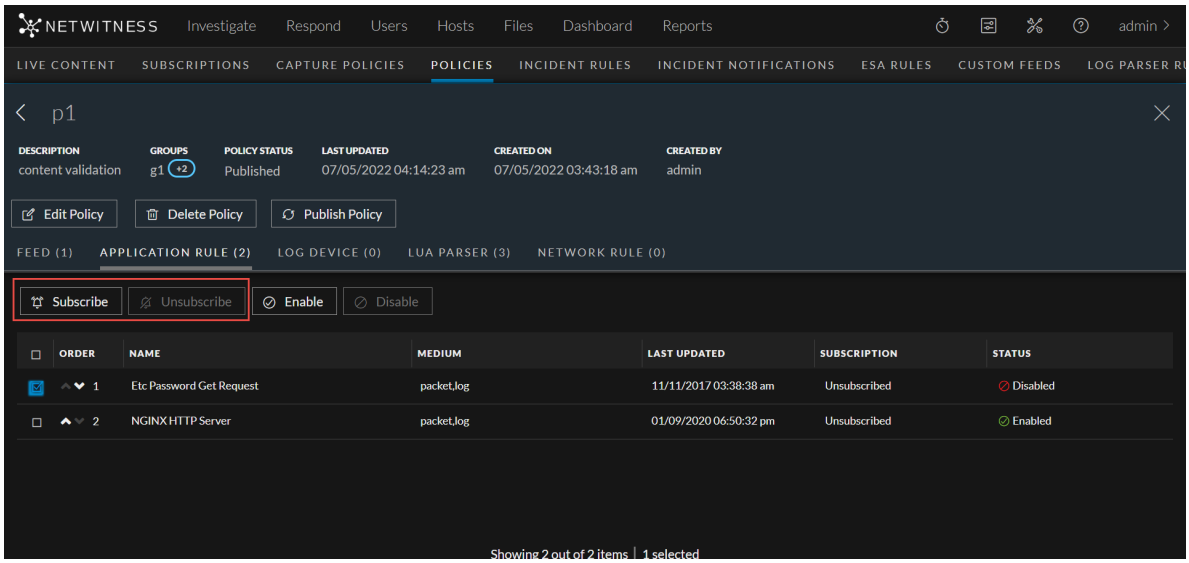
Policy based Centralized Content Management is a unified approach to find, deploy, and manage content through the entire life cycle based on policies that can be assigned to groups of devices. It is a single location to view, modify and manage the content deployed across all services in the environment.

Benefits of Policy based Centralized Content Management:

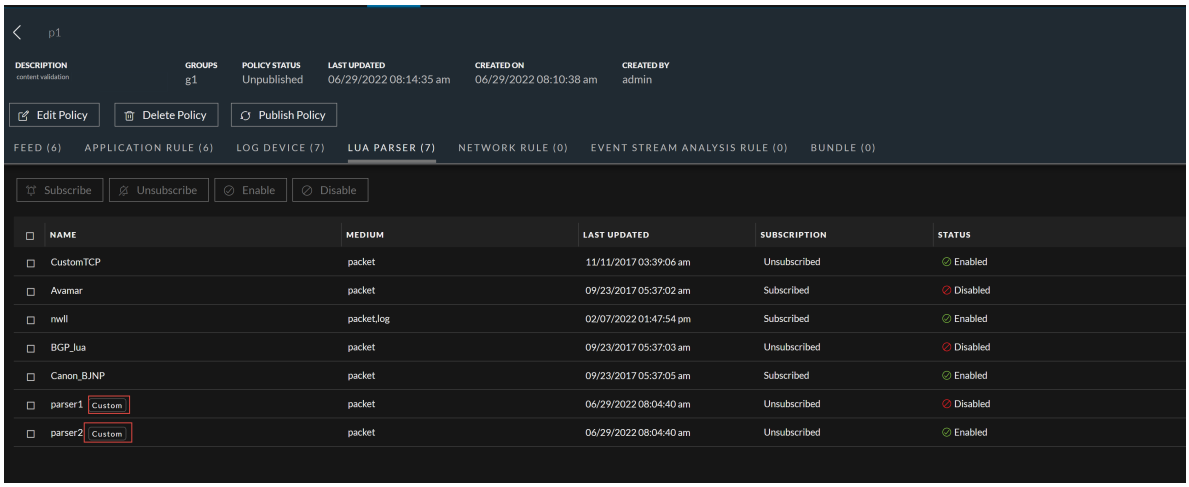
- Add content from RSA Live or add your own custom content.
- Add or remove content without repeating the process on each individual service.
- Add a new service to an existing group to automatically deploy all necessary content.



- One-click management of subscriptions and automatic updates
  - Simply toggle the **Subscribe** button to enable automatic updates of content.



- Provide highly responsive and updated UI for browsing RSA Live content that can help you with the following:
  - View Live and custom content along with your content policies and click  to add content.
  - Seamlessly view Live content along with your own custom content.
  - Centrally import and deploy live and custom content.



For more information, see *Policy based Centralized Content Management* topic in the Live Services Management Guide.

Springboard

The following section describes the new enhancements for the Springboard component:

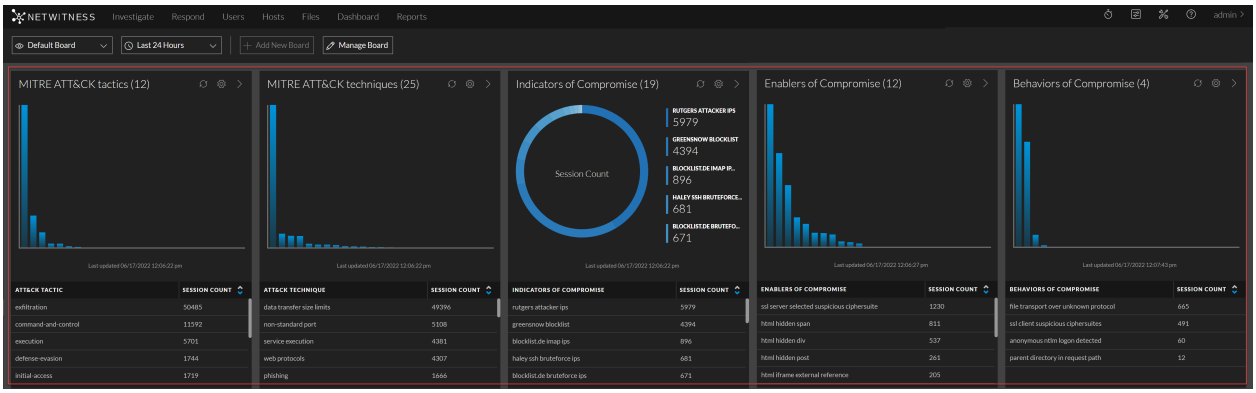
Enhanced Springboard to Support New Built-in Panels

NetWitness Platform Springboard introduces five more out-of-the-box panels based on the events processed and presented on Springboard view. On the Springboard, Administrators and Analysts can now view the following panels of events data which helps in threat hunting and investigation:

- **MITRE ATT&CK tactics**
- **MITRE ATT&CK techniques**
- **Indicators of Compromise**
- **Enablers of Compromise**
- **Behaviors of Compromise**

Administrators can customize these panels to display only the event-focused data for analysts to carry out further investigation.

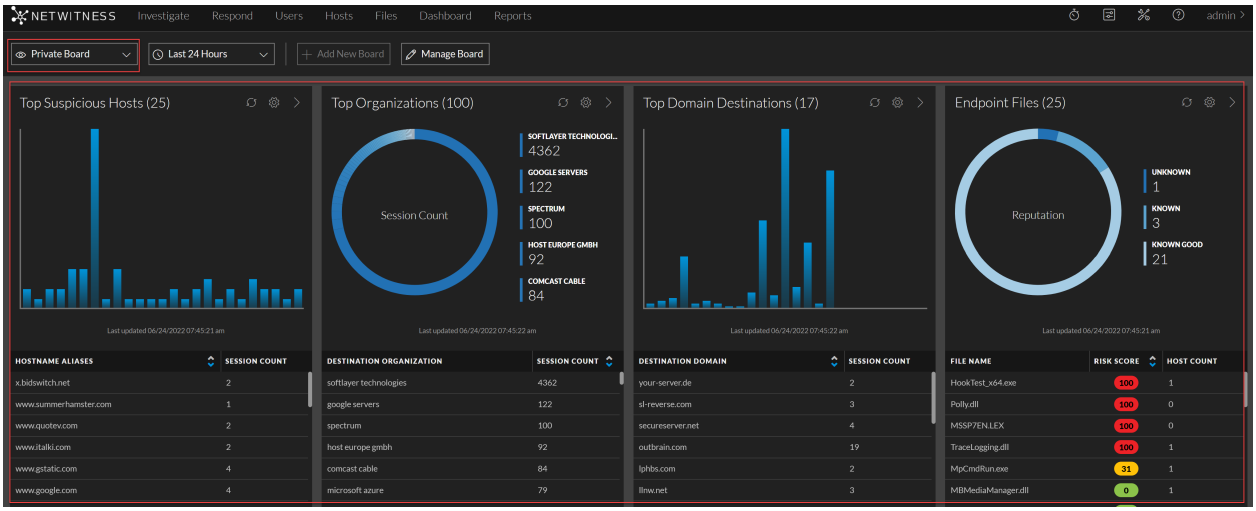
For more information, see *Managing the Springboard* topic in NetWitness Platform Getting Started Guide.



Create Custom Springboard at the User Level

Administrators and Analysts can now add their own custom private board to the NetWitness Platform Springboard and add panels with important system indicators, which helps in threat hunting and investigation. The custom private board is visible only for users who created it. The board allows users to organize and manage information in an easy manner.

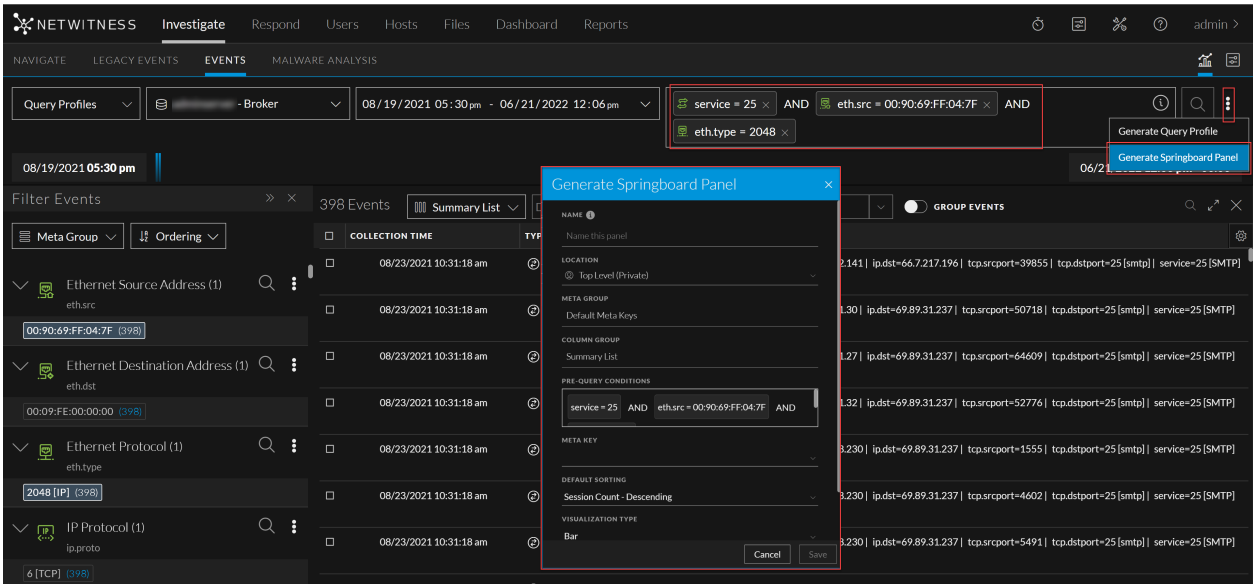
For more information, see *Managing the Springboard* topic in NetWitness Platform Getting Started Guide.



Automated Custom Springboard from Query

During investigation, Administrators and Analysts can add a Springboard panel from the **Investigate** > **Events** view. You can add any number of filters on the query search bar and convert them to Springboard panels for further detection and watch results. The newly added panels will be saved under a custom private board. The board will allow users to organize and manage information in an easy manner.

For more information, see *Add Springboard Panels from Events view* topic in NetWitness Platform Investigate User Guide.



Respond

The Respond view is enhanced to track and capture all the events performed by the users on an incident. The toolbar actions are enhanced to allow users select only the valid priority, status, and assignee for an incident.

Incident Workflow Enhancements

The following changes have been made to the **Change Status** drop-down list in the **Respond** > **Incidents** view:

- Added the new Incident status **Reopen** to help users open the closed incidents.
- Removed **New** and **Assigned** statuses but they are still displayed in the Status column in the **Respond** > **Incidents** > **Incidents List** view.
- Streamlined the incident status change workflow. All the invalid statuses are grayed out, allowing the users to select only the valid status for any incident.

  For more information, see *Escalate or Remediate the Incident* topic in the NetWitness Respond User Guide.

Incident Details View Enhancements

The new **History** Panel is added to display every action performed by the user on an incident. The various actions performed on an incident are as shown below:

- Incident Assignee Change
- Incident Status Change
- Incident Priority Change
- Incident Creation

  For more information, see *Incident Details View* topic in the NetWitness Respond User Guide.

Incident Overview Panel Enhancements

The Incident Overview Panel is enhanced to include the following fields:

- **Time to Acknowledge**(tta): Displays the time taken to assign an Incident after creating it.
- **Time to Detect**(ttd): Displays the time taken for completing the task after the Incident is assigned.
- **Time to Resolve**(ttr): Displays the time taken for closing the task after the Incident is created.
- **External ID**: Allows storing the Incident ID referrals from a different platform.

  For more information, see *Incident Overview Panel* topic in the NetWitness Respond User Guide.

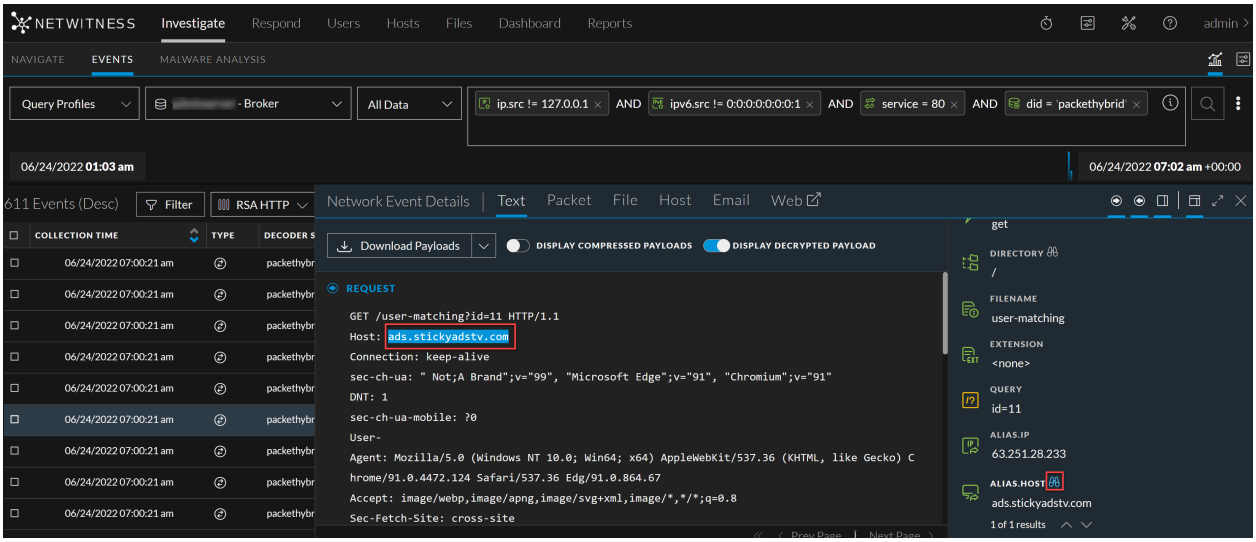
Investigation

The following section describes the new enhancements for the Investigation component:

Indicators for Searchable Meta

The meta key and meta value pairings now display a binocular icon while viewing a text reconstruction in the Event Meta panel, indicating the search option. This enhancement helps the analysts to visually see the indication rather than going through the list of all metadata to figure out which ones may be searched.

For more information, see the NetWitness Platform Investigate User Guide.

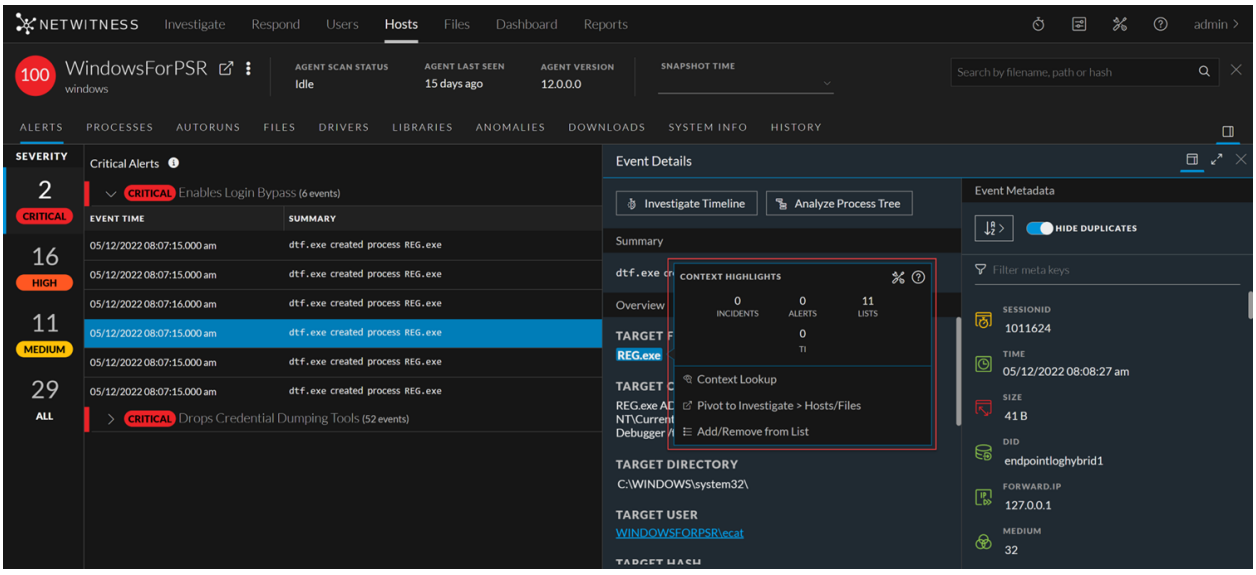


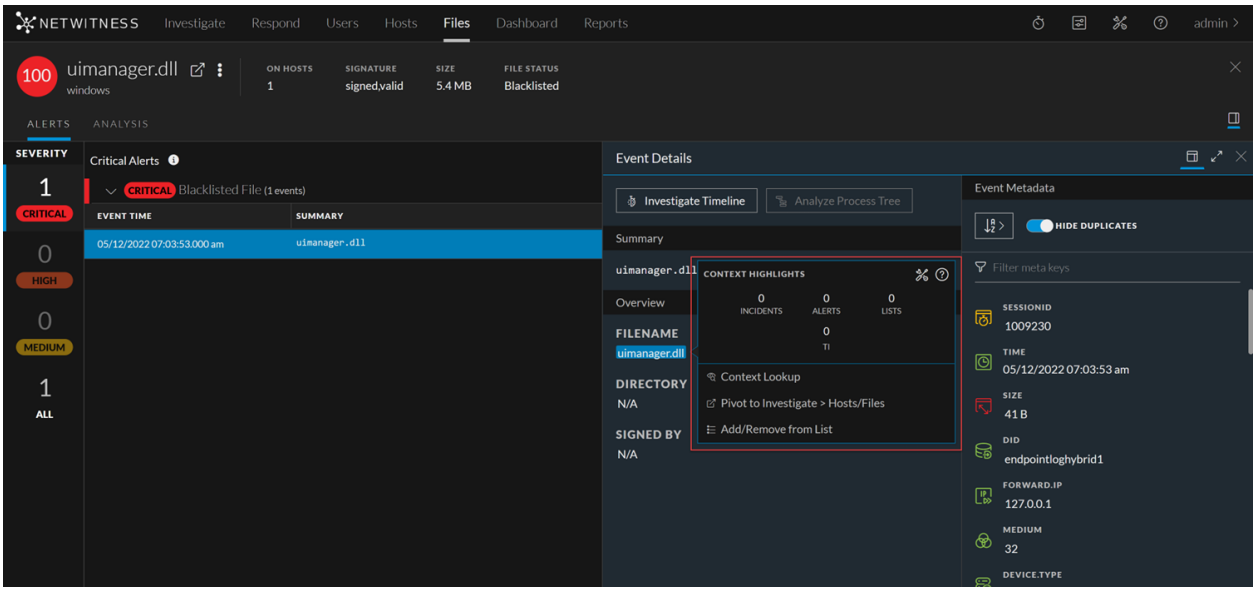
Unified Discovery and Interaction of Events Metadata

Hosts and Files Alerts Details View

Analysts have a unified way to interact with events metadata presented in the Alerts tab of Hosts and Files details view to perform actions or review contextual information. Analysts can use the right and left click options to view the unified panel data.

For more information on Hosts and Files, see *Analyze Hosts Using the Risk Score* and *Analyze Files Using the Risk Score* topics in NetWitness Platform Endpoint User Guide.

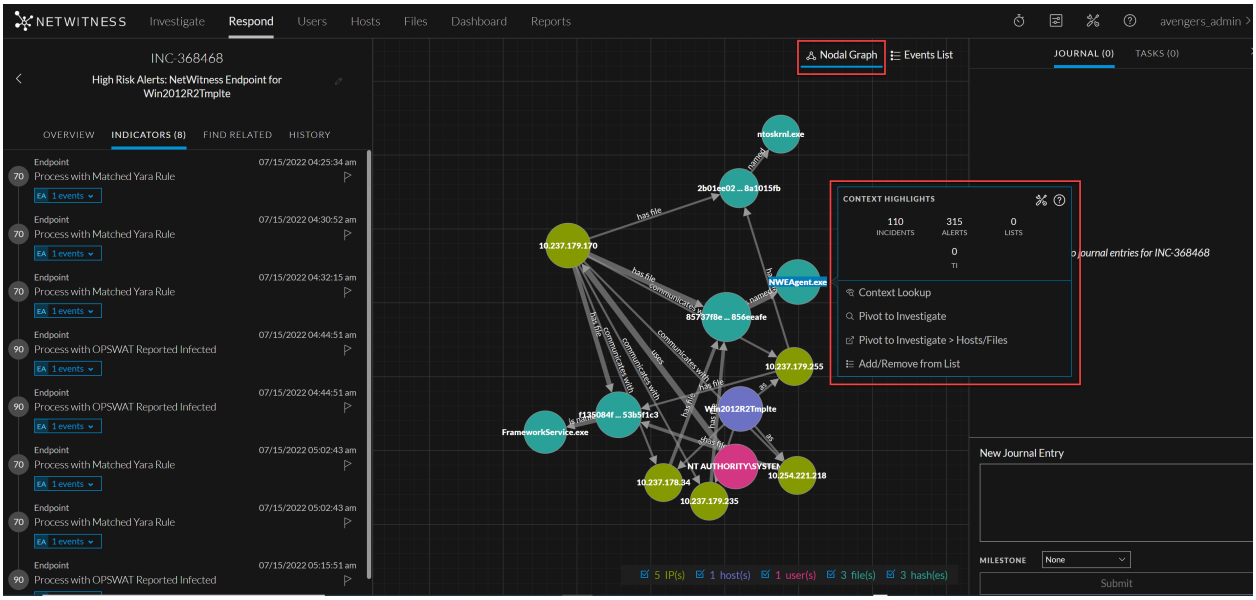


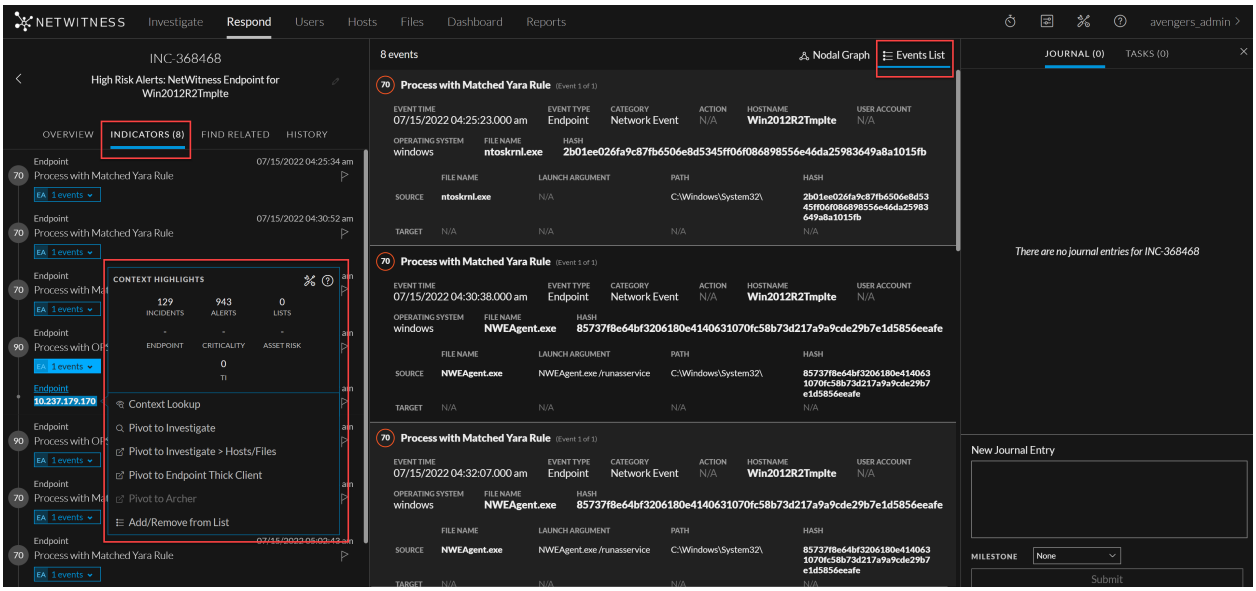


Respond View

Analysts have a unified way to interact with events metadata presented in the Respond view to perform actions or review contextual information.

On the Respond Indicators panel, Nodal Graph, and Events List view, analysts can use the left and right click options to view the unified panel data.



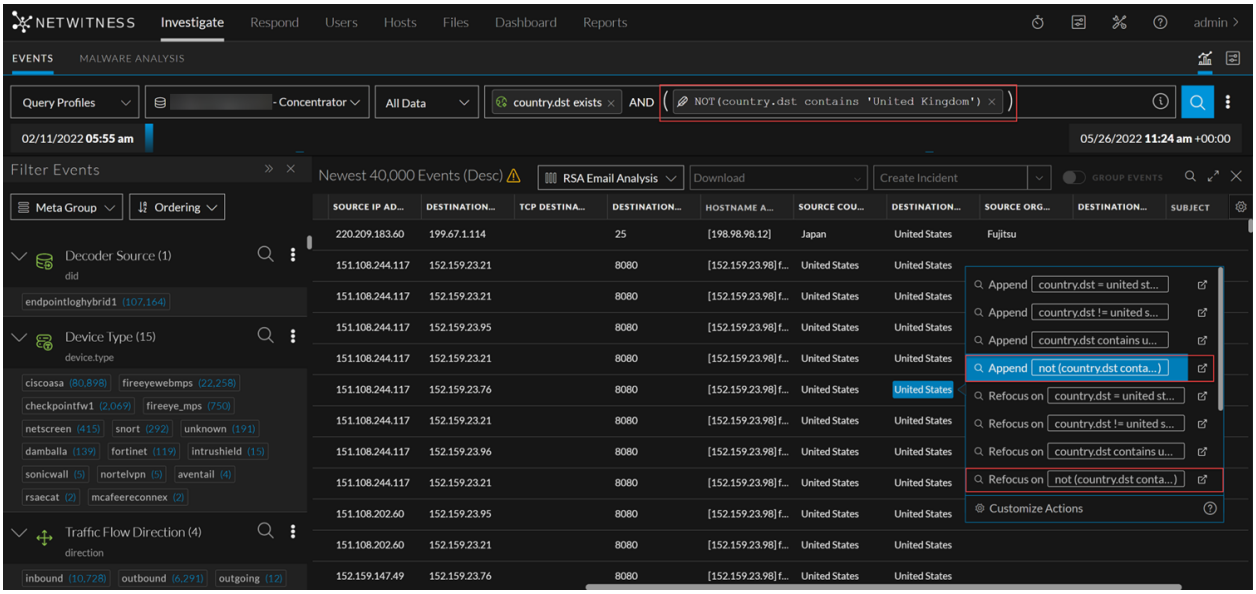


For more information, see NetWitness Platform Respond User Guide.

Enhanced Querying on Events View to Exclude any Specific Meta

Analysts can now exclude particular meta values while querying using the NOT(*meta* contains 'meta value') option available in the investigate unified panel. The specified meta value is removed from the query results when you use NOT(*meta* contains 'meta value') with **Append** or **Refocus** option on a specific meta value. This enhancement helps the analysts to view only the required data results in an optimized manner and conduct further investigation efficiently.

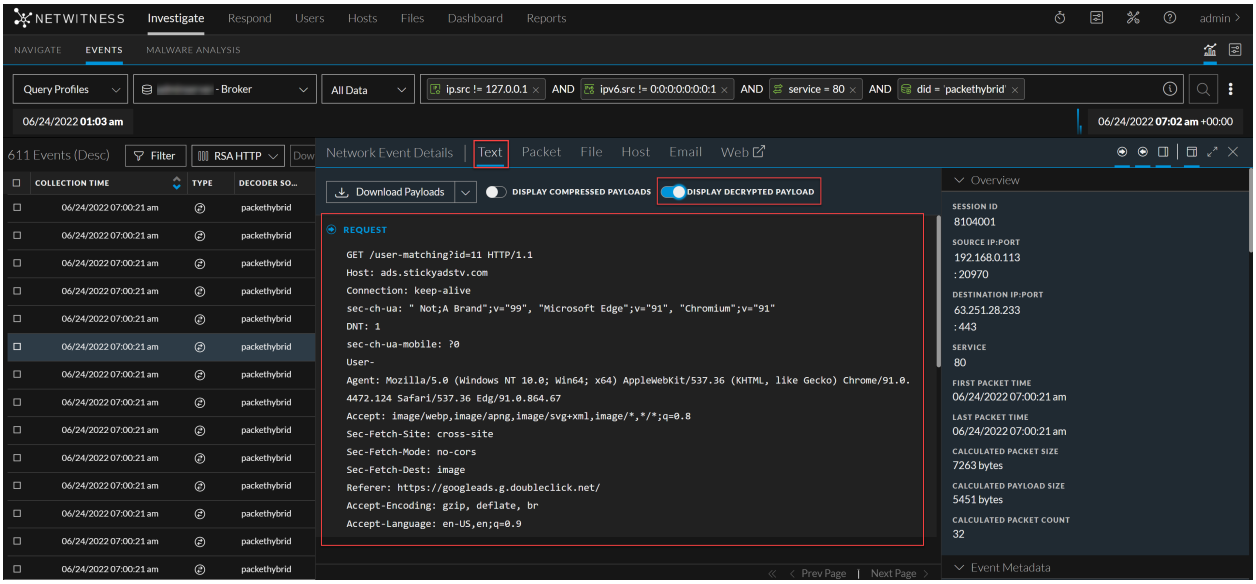
For more information, see the NetWitness Platform Investigate User Guide.



View Encrypted Data in Decrypted Format

Analysts can directly view encrypted data that has been decrypted by the decoder, thereby reducing time and effort in converting data into readable format. The analysts can enable using the **Display Decrypted Payload** toggle option in the **Events** > **Text** view.

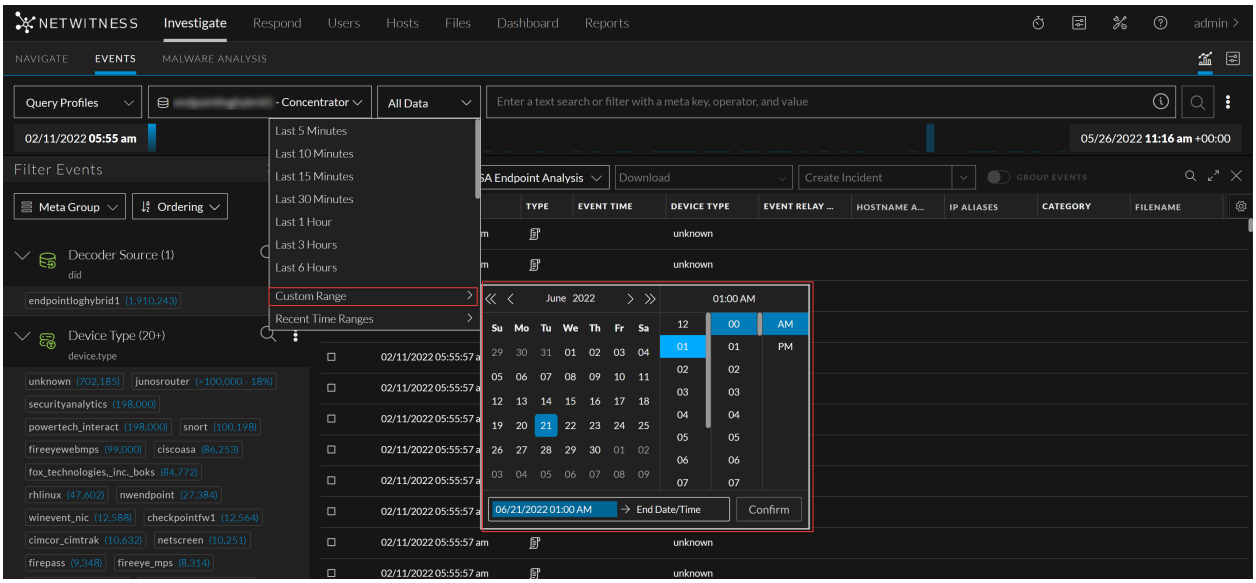
For more information, see the *Text Reconstruction* topic in the NetWitness Platform Investigate User Guide.



Select Custom Date and Time Range in the Events View

Analysts can set a custom range in the **Investigate** > **Events** view to select a specific time, date, month, and year using the calendar view that is displayed on clicking the **Custom Range** option. This enhancement helps the analysts to select date and time quickly and avoid manual intervention therefore avoiding human errors (typos).

For more information, see *Select a Time Range* topic in the NetWitness Platform Investigate User Guide.

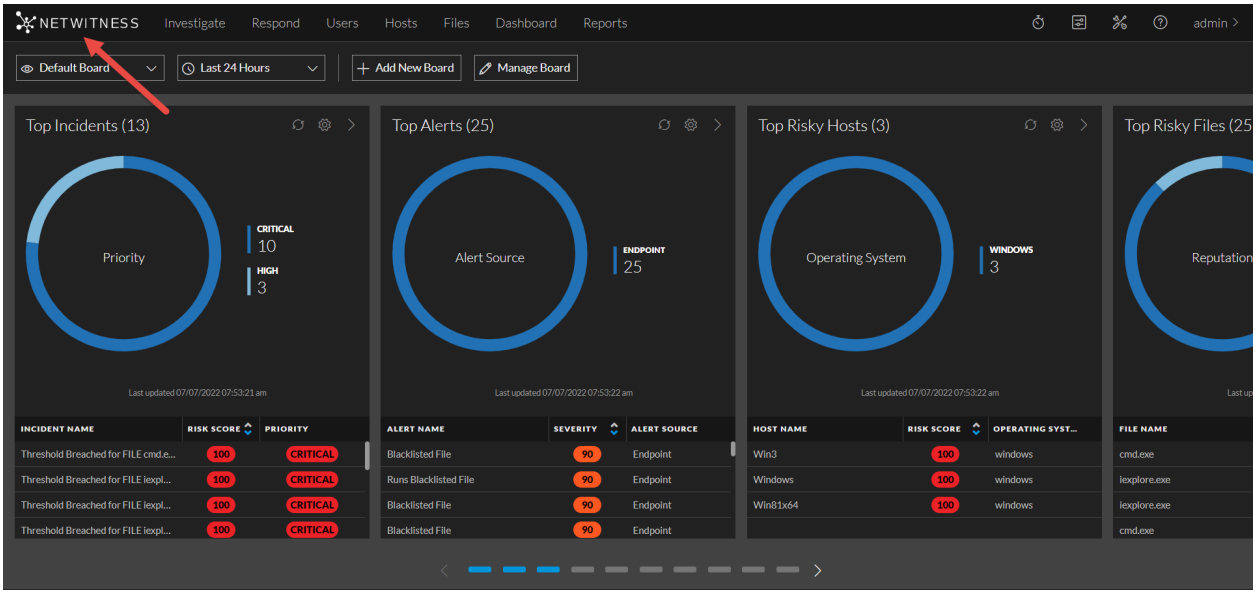


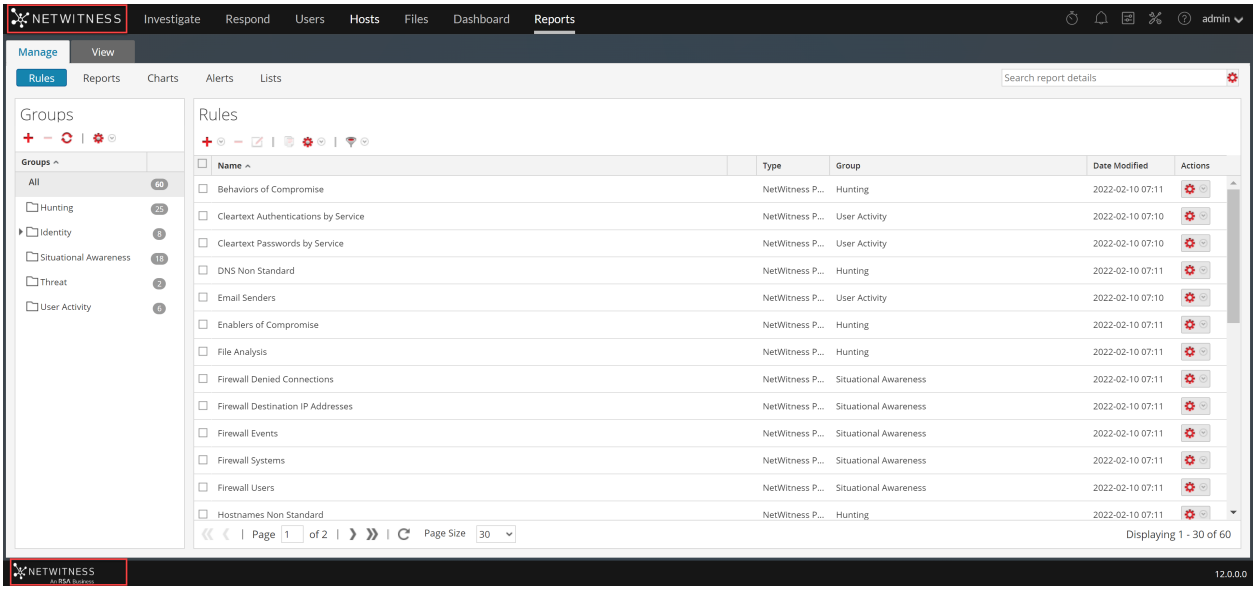
User Interface

The following section describes the new enhancements for the NetWitness user interface:

NetWitness User Interface Enhancements

- The 12.0.0.0 release includes the new NetWitness corporate logo. You can view the new logo in NetWitness Platform, which updates the identity of NetWitness as a trusted brand.
- As part of the repositioning, we are renaming our product as **NetWitness Platform XDR**. This change aims to simplify communications and improve our customers' understanding of how each product secures and protects within the NetWitness portfolio.





Endpoint Investigation

The following section describes the new enhancements for the Endpoint component:

Detection of removable Storage Devices

NetWitness Endpoint Agents are enhanced with the capabilities to detect and report removable storage devices. The Endpoint agents will detect and report when a removable storage device is plugged in or removed. This enhancement provides analysts with extended threat detection capabilities. For more information, see the NetWitness Endpoint User Guide.

Block Multiple File Hashes Using an Imported File

Administrators can import a file with a list of known file hashes that are not present in the environment and block them as soon as they are detected. This enhancement will help analysts to block multiple hashes without manual intervention.

Support for Arm-based Windows Machines

Administrators can install Endpoint agents on Arm-based Windows machines. This enhancement provides analysts with threat detection capabilities on more types of devices.

Download MFT from Multiple Hosts in One Step

Analysts can now download MFT(Master File Table) from multiple hosts on the Hosts list view in one step. This enhancement helps analysts download MFT without opening the Host details view of each host. For more information, See *Download Master File Table* topic on NetWitness Endpoint User Guide.

Customizable Maximum File Download Limits

The limit to the maximum number of file downloads on the Endpoint server is enhanced. On the explore page of an Endpoint server, Administrators can set the limit from 100 to 1000 files. For more information, see *Download Files Using Full Path or Wildcard* on NetWitness Endpoint User Guide.

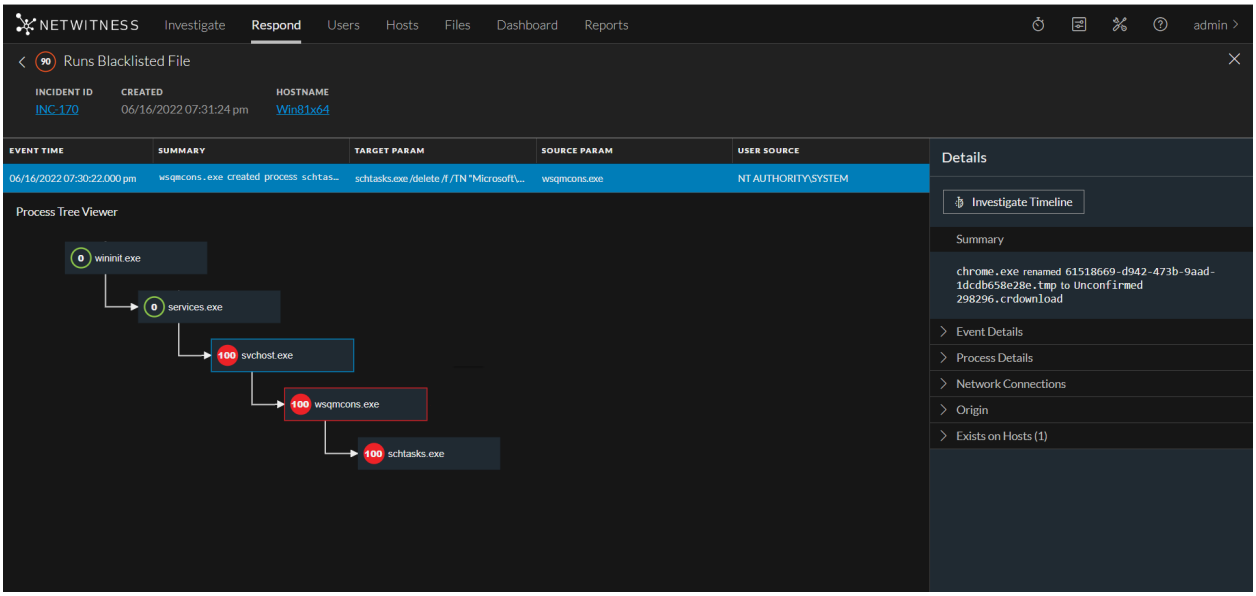
Redesigned Alert Details View for Endpoint Alerts in Respond

In the Respond view, the alert details view for Endpoint alerts shows end-to-end details about an alert. The details are presented in the form of a process tree along with a right panel that provides detailed information about the alert categorized into the following sections:

- **Summary**: A short summary of the alert.
- **Event Details**: Shows the directory, user, hash, signature, risk score, etc.
- **Process Details**: Shows the tactics, techniques, times and details about the targets.
- **Network Connections**: Shows any network connection established ten minutes before and till ten minutes after the alert triggered time.
- **Origin**: Shows how the selected file in the process tree is originated.
- **Exists on Hosts**: The host in which the selected file in the process tree exists.

Besides the above sections, the **Investigate Timeline** takes to the investigate view that has more detailed information.

For more information, see *Review Endpoint Alerts using Process Tree* on the NetWitness Respond User Guide.



Concentrator, Decoder, and Log Decoder Services

The following section describes the new enhancements for the Concentrator, Decoder, and Log Decoder components:

Log Parsing Enhancements

The following log parsing enhancements are made in 12.0.0.0 version. These are new elements that you use in the creation of a log parser:

New Selector Parsing Element Added to Dynamically Map Captured Values to a Meta Key

This will allow the log parser to automatically choose from two or more optional meta keys to assign to a parsed value depending upon the value of another meta key. Consider the following sample log snippet:

```
{ "Direction": "src", "Address": "1.2.3.4" }
{ "Direction": "dest", "Address": "5.6.7.8" }
```

In the above example, if the value of Direction is ”src”, then the preferred meta key to use for the value of Address would likely be **ip.src**. Conversely, if the value for Direction is ”dest”, then the meta key **ip.dst** might be preferred. This can now be achieved with the new **SELECTOR** log parsing element.

Support for Advanced Parsing Elements within CEF Parser and DataType

Support added to CEF parser for VARTYPE, SCANNED, DataType, and Selector parsing elements.

- Allows the CEF parser to take advantage of the fine parsing capabilities found in other parsers.

Dynamic parsing support including PARSERULESCAN added to DataType parsing element.

- Allows nesting of dynamic parsing elements (parse rules) from within an existing DataType.

Enhanced Network Decoder to Decrypt Incoming TLS 1.3 Packets

The enhanced network packet decryption capability helps inspect TLS 1.3 encrypted communications using ephemeral session keys. Administrators can configure Network Decoder to enable decryption of incoming TLS 1.3 network packets.

For more information, see the NetWitness Decoder Configuration Guide.

Event Stream Analysis (ESA)

The Event Stream Analysis is enhanced to reduce the time consumed for new rules deployment.

Improved ESA Rules Deployment

The ESA Rule Deployment has been enhanced with a new option to deploy the rules faster. If you want to push rule-related changes, you can quickly deploy the new rules by clicking the **Fast Deploy** option. For more information, see Alerting with ESA Correlation Rules User Guide.

Reports

The following section describes the new enhancements for the Reports component:

Build Rule View Enhancements

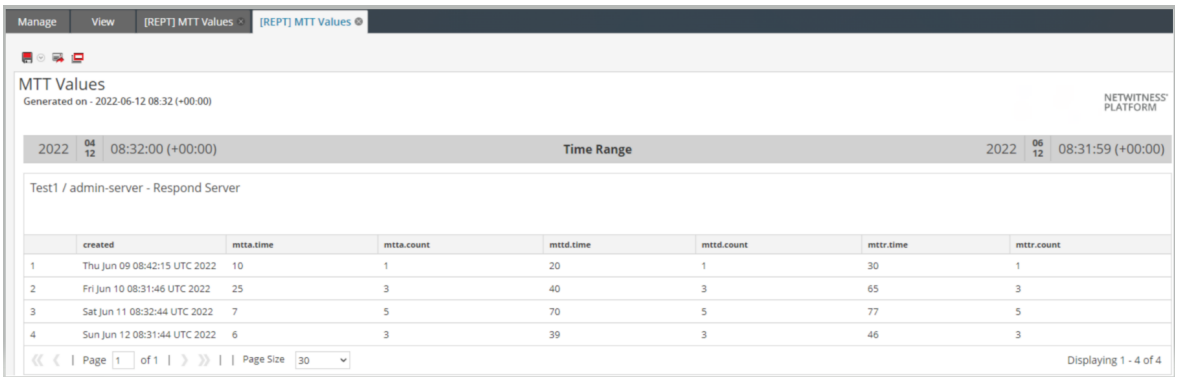
The **Build Rule** view is enhanced to help users view the following information in the report generated:

- The average time taken to assign the incident.
- The average time taken to complete the task.
- The average time taken to close the incident.

The following changes have been made in the **Build Rule** view:

1. Two new options are added in the **From** field:
   - **incidentStats**: The following metas are supported for **incidentStats**:
     a. **created**
     b. **mtta.time**: Displays the average time taken to acknowledge the incidents in a single day.
     c. **mtta.count**: Displays the number of incidents acknowledged in a single day.
     d. **mttd.count**: Displays the number of incidents detected in a single day.
     e. **mttd.time**: Displays the average time taken to detect the incidents in a single day.
     f. **mttr.time**: Displays the average time taken to resolve the incidents in a single day.
     g. **mttr.count**: Displays the number of incidents resolved in a single day.

     These metas are displayed in the report generated. Refer the following figure.

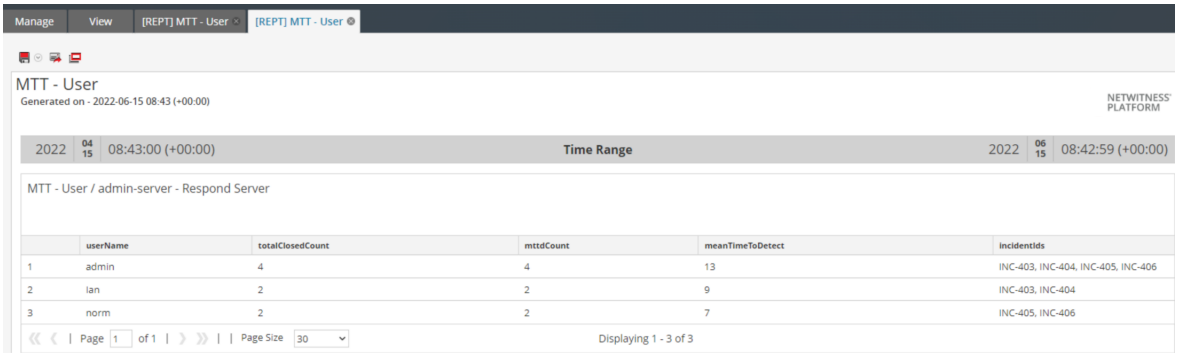


   - **incidentUserStats**: The following metas are supported for **incidentUserStats**:
     a. **userName**: Displays the assignee's or the user's ID for the associated user stats.
     b. **totalClosedCount**: Displays the total number of Incidents closed by the assignee till date.
     c. **meanTimeToDetect**: Displays the average time taken by the user to detect the incidents in the time range selected.

d. **mttdCount**: Displays the count of incidents contributing to the MTTD value computed.

e. **incidentIds**: Displays the list of incident IDs closed by the user during the time range selected.

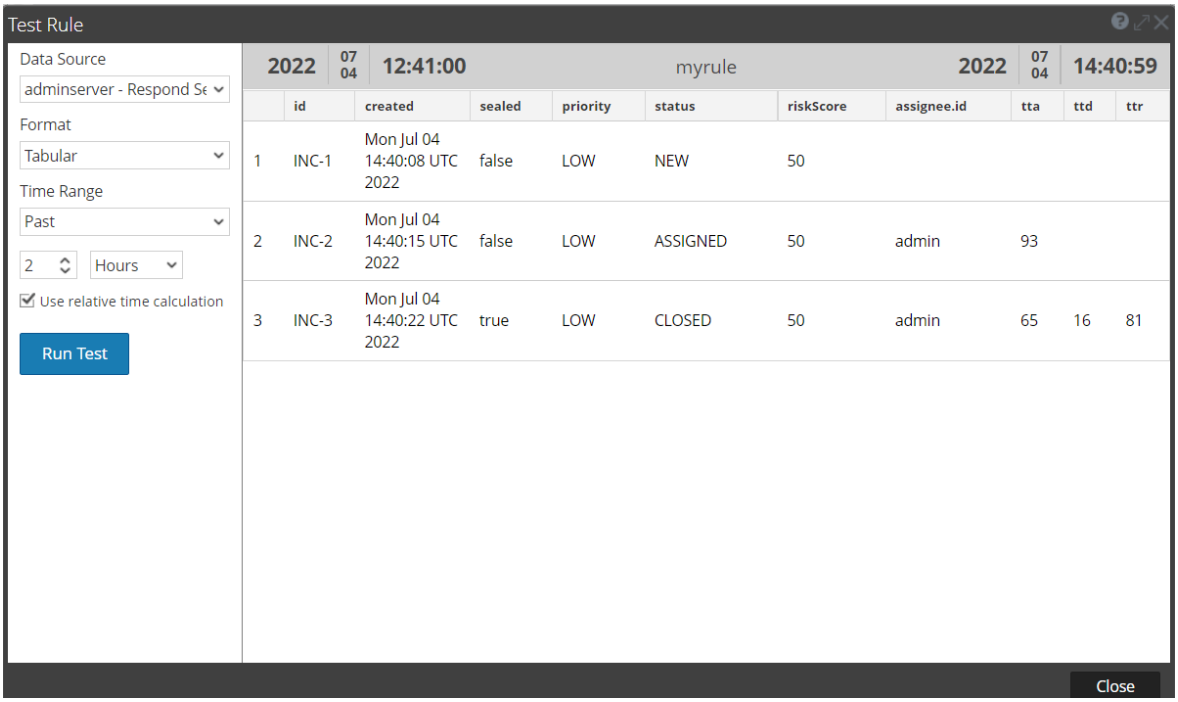
These metas are displayed in the report generated. Refer the following figure.



2. New metas are added for **incident**. The newly added metas are as shown below:

- **assignee.id**
- **tta**(Time to Acknowledge): Displays the time taken to assign an Incident after creating it.
- **ttd**(Time to Detect): Displays the time taken for completing the task after the Incident is assigned.
- **ttr**(Time to Resolve): Displays the time taken for closing the task after the Incident is created.

These metas are populated on the **Test Rule** view. Refer the following figure.



For more information, see the *Create a Rule Using Respond Data Source* topic in the NetWitness Reporting User Guide.

Fixed Issues

This section lists issues fixed after the last major release. For additional information on fixed issues, see the Fixed Version column in the NetWitness® Platform Known Issues list on NetWitness Community Portal.

Administration Fixes

Tracking Number	Description
ASOC-114808	Multiple operations fail after you enable logging for CRUD operation on app rules and use separate API for ADD, DELETE, and UPDATE.
ASOC-113813	The Jobs **(Admin > System > Jobs)** with lengthy queries take a long time to load. As a result, the load time of the Jobs page is affected.

Endpoint Fixes

Tracking Number	Description
SACE-17157	Network adapter in promiscuous state is not identified due to querying wrong wmi namespace by NetWitness agent.
ASOC-113849	Failed icon is displayed beside the agent version (**Hosts> Agent History**) when there is a time difference between agent and server.

Respond Fixes

Tracking Number	Description
ASOC-115196	Aggregation syntax count is not displayed in the report, when you try to create a report with metas such as alert.name, alert.severity,count (alert.numEvents).

Core Services (Broker, Concentrator, Decoder, Archiver) Fixes

Tracking Number	Description
ASOC-114900	Brotil data fails to decompress in Events page (**Investigate> Events**). When exporting data, it is decompressed but on Events page invalid data is displayed.
ASOC -115052	Due to the large size of incoming logs aggregation fails and archiver service crashes. The mem page issue occurs if sessions to be aggregated from log decoder are with large packet andlog sizes.

Reporting Engine Fixes

Tracking Number	Description
ASOC -115232	NetWitness Weekly Scheduled Reports run on a different day after conversion to UTC flows into either the previous or next day for time zones other than the UTC time zone

Log Collection Fixes

Tracking Number	Description
ASOC-120851	After upgrading to 11.7.1.0, WinRM log collection is interrupted due to the conflict in the internal artifact. As a result, logs are not collected from the WinRM server.

End of Life Functionality

The following table provides information on end of life functionality, features and hardware in NetWitness 12.0.0.0 or later releases.

End of Life Functionality and Features in 12.0.0.0 or higher releases

Feature	Notes
Dell Series 4 and Series 4s Hardware	As the Dell S4 and S4s appliances reached the end of life (EOL) in June 2021, we recommend that you not perform any installation or upgrade activities on these and upgrade to new hardware.

Product Documentation

The following documentation is provided with this release.

Documentation	Location URL
NetWitness Platform Master Table of Contents	https://community.netwitness.com/t5/netwitness-platform-online/tkb-p/netwitness-online-documentation
NetWitness Platform 12.0 Product Documentation	https://community.netwitness.com/t5/netwitness-platform-online/tkb-p/netwitness-online-documentation
NetWitness Platform 12.0.0.0 Upgrade Guide	https://community.netwitness.com//t5/rsa-netwitness-platform-staged/upgrade-guide-for-11-7-1/ta-p/658666

Feedback on Product Documentation

You can send an email to nwdocsfeedback@netwitness.com to provide feedback on NetWitness Platform documentation.

Getting Help with NetWitness Platform

Self-Help Resources

There are several options that provide you with help as you need it for installing and using NetWitness:

- See the documentation for all aspects of NetWitness here: https://community.netwitness.com/t5/netwitness-platform/ct-p/netwitness-documentation
- Use the **Search** and **Create a Post** fields in NetWitness Community portal to find specific information here:https://community.netwitness.com/t5/netwitness-discussions/bd-p/netwitness-discussions
- See the NetWitness Knowledge Base: https://community.netwitness.com/t5/netwitness-knowledge-base/tkb-p/netwitness-knowledge-base
- See Troubleshooting section in the guides.
- See also RSA NetWitness® Platform Blog Posts.
- If you need further assistance, contact NetWitness Support.

Contact NetWitness Support

If you contact NetWitness Support, you should be at your computer. Be prepared to provide the following information:

- The version number of the NetWitness Platform product or application you are using.
- The type of hardware you are using.

Use the following contact information if you have any questions or need assistance.

NetWitness Community Portal	https://community.netwitness.com In the main menu, click **Support > Case Portal > View My Cases**.
International Contacts (How to Contact RSA NetWitness Support)	https://community.netwitness.com/t5/support/ct-p/support
Community	https://community.netwitness.com/t5/netwitness-discussions/bd-p/netwitness-discussions
NW Update	update.netwitness.com
LiveUi	live.netwitness.com

Build Numbers

The following table lists the build numbers for various components of NetWitness 12.0.0.0

Component	Version Number
NetWitness Audit Plugins	rsa-audit-plugins-12.0.0.0-4762.5.8853c5bef.el7.noarch.rpm
NetWitness Appliance	rsa-nw-appliance-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness Archiver	rsa-nw-archiver-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness Broker	rsa-nw-broker-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness Concentrator	rsa-nw-concentrator-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness Config Management	rsa-nw-config-management-12.0.0.1-2208052108.5.104cf25.el7.noarch.rpm
NetWitness Config Server	rsa-nw-config-server-12.0.0.0-220531042629.5.b236ead.el7.centos.noarch.rpm
NetWitness Console	rsa-nw-console-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness Content Server	rsa-nw-content-server-12.0.0.0-220426021902.5.dc28ce9.el7.centos.noarch.rpm
NetWitness ContextHub Server	rsa-nw-contexthub-server-12.0.0.0-220613050521.5.b986911.el7.centos.noarch.rpm
NetWitness Correlation Server (ESA)	rsa-nw-correlation-server-12.0.0.0-220610051650.5.57a3c20.el7.centos.noarch.rpm
NetWitness Decoder	rsa-nw-decoder-12.0.0.0-12536.5.ee5a55cf6.el7.x86_64.rpm
NetWitness Deployment Upgrade	rsa-nw-deployment-upgrade-12.0.0.1-2208060358.5.3df31a1.el7.noarch.rpm
NetWitness Endpoint Agents	rsa-nw-endpoint-agents-12.0.0.0-2206211619.5.3fe862e.el7.x86_64.rpm
NetWitness Endpoint Broker Server	rsa-nw-endpoint-broker-server-12.0.0.0-220610063852.5.b91f656.el7.centos.noarch.rpm
NetWitness Endpoint Server	rsa-nw-endpoint-server-12.0.0.0-220620061620.5.ea8ba76.el7.centos.noarch.rpm
NetWitness Integration Server	rsa-nw-integration-server-12.0.0.0-220609130615.5.71ca677.el7.centos.noarch.rpm
NetWitness Investigate Server	rsa-nw-investigate-server-12.0.0.0-220607015824.5.dbfcf19.el7.centos.noarch.rpm
NetWitness Legacy Web Server	rsa-nw-legacy-web-server-12.0.0.1-220806042518.5.2c3da62.el7.centos.noarch.rpm

NetWitness License Server	rsa-nw-license-server-12.0.0.0-220412123008.5.1fb53c8.el7.centos.noarch.rpm
NetWitness Log Collector	rsa-nw-logcollector-12.0.0.0-15032.5.5e731802a.el7.x86_64.rpm
NetWitness Log Decoder	rsa-nw-logdecoder-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness Log Player	rsa-nw-logplayer-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness Malware Analytics Server	rsa-nw-malware-analytics-server-12.0.0.0-220520044551.5.56297e2.el7.centos.x86_64.rpm
NetWitness Metrics Server	rsa-nw-metrics-server-12.0.0.0-220530065358.5.24033a9.el7.centos.noarch.rpm
NetWitness Orchestration Server	rsa-nw-orchestration-server-12.0.0.0-220531093544.5.ed07d6c.el7.centos.noarch.rpm
NetWitness Recovery Tools	rsa-nw-recovery-tool-12.0.0.0-2205190541.5.c20bc1d.el7.noarch.rpm
NetWitness Reporting Engine Server	rsa-nw-re-server-12.0.0.0-5948.5.378a71f6b.el7.x86_64.rpm
NetWitness Respond Server	rsa-nw-respond-server-12.0.0.0-220613050612.5.15a626b.el7.centos.noarch.rpm
NetWitness Root CA Update	rsa-nw-root-ca-update-12.0.0.0-2206021205.5.47aad12.el7.noarch.rpm
NetWitness SDK	rsa-nw-sdk-11.6.0.0-11374.5.fe9457e29.el7.x86_64.rpm
NetWitness Security Server	rsa-nw-security-server-12.0.0.0-220624043254.5.b9a8a8f.el7.centos.noarch.rpm
NetWitness Source Server	rsa-nw-source-server-12.0.0.0-220623003834.5.d4e49fb.el7.centos.noarch.rpm
NetWitness User Interface	rsa-nw-ui-12.0.0.1-220806005807.5.e33efada87.el7.centos.noarch.rpm
NetWitness Workbench	rsa-nw-workbench-12.0.0.0-12540.5.9f874c638.el7.x86_64.rpm
NetWitness SA Tools	rsa-sa-tools-12.0.0.1-2208060720.5.4d4a622.el7.noarch.rpm
NetWitness SMS Runtime	rsa-sms-runtime-rt-12.0.0.0-4762.5.8853c5bef.el7.x86_64.rpm
NetWitness SMS Server	rsa-sms-server-12.0.0.0-4762.5.8853c5bef.el7.x86_64.rpm
NetWitness Windows Legacy Log Collector	NWLegacyWindowsCollector-12.0.0.1-15036.5.056fd2393.exe

Revision History

Date	Description
August 2022	General Availability

NetWitness® Platform XDR

Version 12.0.0.0

Upgrade Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

January, 2023

Contents

Upgrade Overview **6**
- Upgrade Paths 6
- Running in Mixed Mode 6
  - Upgrade Considerations for ESA Hosts 7
- Upgrade Considerations for STIX Custom Feeds 7
- Upgrade or install Legacy Windows Log Collection 7
- Feedback on Product Documentation 7

Getting Help with NetWitness Platform **8**
- Self-Help Resources 8
- Contact NetWitness Support 8

Pre Upgrade Checks **9**

Upgrade Preparation Tasks **11**
- Task 1. (Optional) Remove Legacy Package Repositories 11
- Task 2. Backup and Remove the Rotated RabbitMQ Logs 11
- Task 3. Uninstall the Security Analytics l10n language pack 12

Upgrade Options **13**
- Important Notes - Read This First 13
  - Synchronize Time on Component Hosts with NW Server Host 13
  - Mixed Mode Unsupported for ESA Hosts 13
  - Endpoint Hybrid Systems Not Supported 14
  - Respond Server Service Not Enabled Until NW Server and Primary ESA Host Upgraded to 12.0.0.0 14
  - Deploy_Admin Password Guidelines 14
  - Additional Post Upgrade Steps for 12.0.0.0 Version with Legacy Windows Log Collector 14
- Upgrade Options 14
- Option 1: User Interface Method with Connectivity to the Internet 15
  - Procedure 15
- Option 2: User Interface with No Connectivity to the Internet 17
  - Task 1. Populate Staging Folder (/var/lib/netwitness/common/update-stage/) with Version Upgrade Files 17
  - Task 2. Apply Upgrades from the Staging Area to Each Host 17
- Option 3: Command Line Interface (CLI) with No Connectivity to the Internet 18
- Option 4 (Optional): Pre-Stage Upgrade Repo 18

Post Upgrade Tasks **20**
- General 20
  - (Conditional) Configure NAT-Based IP Addresses 20

(Conditional - For Warm-Standby Hosts Only) Register the Secondary IP Address of Warm-Standby Hosts ... 20
Review Contents of /etc/hosts.user for Obsolete Host Entries ... 21
Jetty Configuration ... 21
Reconfigure DNS Servers ... 21
Make Sure Services Have Restarted and Are Capturing and Aggregating Data ... 21
New Health and Wellness ... 23
Deploy the New Health and Wellness Content from Live ... 23
(Optional) Update UUID of New Health and Wellness Host to Update Service Configuration Documents ... 24
Investigate ... 25
(Conditional - For Custom Roles Only) Adjust investigate-server Permissions for Custom User Roles ... 25
Respond ... 25
(Conditional) Restore Any Respond Service Custom Keys in the Aggregation Rule Schema ... 26
Reference Log Decoder ... 26
Legacy Windows Log Collector ... 26
Update the Legacy Windows Log Collector UUID ... 26
Refresh Legacy Windows Log Collector Certificates with Updated SA Certificates ... 26
User Entity Behavior Analytics ... 27
Endpoint Upgrade Tasks ... 30
Install the 12.0.0.0 Relay Server ... 30
Upgrade Endpoint Agents ... 30
Appendix A. Offline Method (No Connectivity to Live Services) - Command Line Interface ... 31
External Repo Instructions for CLI upgrade ... 32
Appendix B. Set Up External Repo ... 33
Appendix C. Troubleshooting Version Installations and Upgrades ... 35
deploy_admin User Password Has Expired Error ... 37
Downloading Error ... 38
Error Deploying Version <version-number> Missing Update Packages ... 39
Upgrade Failed Error ... 39
External Repo Update Error ... 41
Host Update Failed Error ... 41
Missing Update Packages Error ... 42
OpenSSL 1.1.x ... 42
Patch Update to Non-NW Server Error ... 43
Reboot Host After Update from Command Line Error ... 43
Reporting Engine Restarts After Upgrade ... 43
Log Collector Service (nwlogcollector) ... 45
NW Server ... 47

Orchestration 48
Reporting Engine Service 48
Event Stream Analysis 48
Legacy Windows Log Collector 49
ESA Troubleshooting Information 49
ESA Rules are Not Creating Alerts 49
Endpoint, UEBA, and Live Content Rules are Not Working 50
Example ESA Correlation Server Warning Message for Missing Meta Keys 51

Upgrade Overview

NetWitness 12.0.0.0 provides enhancements and fixes for all products in NetWitness Platform. The instructions in this guide apply to both physical and virtual hosts (including AWS, Azure Public Cloud, and Google Cloud Platform) unless stated to the contrary.

In 12.0.0.0, NetWitness has several new features in the user interface.

Note: In 11.5 and above versions, Administrative tasks are consolidated as icons in the upper right corner to keep administration, configuration, notifications, jobs, and user preferences together.

Upgrade Paths

The following upgrade paths are supported for NetWitness 12.0.0.0:

- NetWitness 11.6.0.0 to 12.0.0.0
- NetWitness 11.6.0.1 to 12.0.0.0
- NetWitness 11.6.1.0 to 12.0.0.0
- NetWitness 11.6.1.1 to 12.0.0.0
- NetWitness 11.6.1.2 to 12.0.0.0
- NetWitness 11.6.1.3 to 12.0.0.0
- NetWitness 11.6.1.4 to 12.0.0.0
- NetWitness 11.7.0.0 to 12.0.0.0
- NetWitness 11.7.0.1 to 12.0.0.0
- NetWitness 11.7.0.2 to 12.0.0.0
- NetWitness 11.7.1.0 to 12.0.0.0
- NetWitness 11.7.1.1 to 12.0.0.0

This guide applies to both physical and virtual hosts (including AWS and Azure Public Cloud).

Note: If you have installed pre-upgrade checks hot fix in any version (11.6.x.x or 11.7.x.x), you must uninstall the hot fix before you upgrade to 12.0.0.0 using the below command:

```
- rpm -qa | grep prechecks-hotfix | xargs -I{} yum remove -y {}
```

Running in Mixed Mode

Running in mixed mode occurs when some services are upgraded to the latest version and some services are on older versions. See "Running in Mixed Mode" in the *NetWitness Platform Hosts and Services Getting Started Guide* for further information.

Note: If you are running Endpoint Log Hybrid in mixed mode, make sure Endpoint Broker is on the same version as one of the Endpoint Servers.

Upgrade Considerations for ESA Hosts

Mixed mode is not supported for ESA hosts in NetWitness version 11.5 and later.

IMPORTANT: The NetWitness server, ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

Upgrade Considerations for STIX Custom Feeds

The custom feeds created before version 12.0.0.0 are processed automatically. On upgrade, the data sources created for ADHOC, REST and TAXII server and the feeds are pulled automatically. See "Create a STIX Custom Feed" in the *NetWitness Platform Live Service Management Guide* and "Configure STIX as a Data Source" in the *NetWitness Platform Context Hub Configuration Guide* for further information.

Upgrade or install Legacy Windows Log Collection

Refer to the Legacy Windows Log Collection Guide for NetWitness.

Note: After you update or install Legacy Windows Log Collection, reboot the system to ensure that Log Collection functions correctly.

Feedback on Product Documentation

You can send an email to nwdocsfeedback@netwitness.com to provide feedback on NetWitness documentation.

Getting Help with NetWitness Platform

Self-Help Resources

There are several options that provide you with help as you need it for installing and using NetWitness:

- See the documentation for all aspects of NetWitness here: https://community.netwitness.com/t5/netwitness-platform/ct-p/netwitness-documentation
- Use the **Search** and **Ask it** fields in NetWitness Community to find specific information here: https://community.netwitness.com/
- See the NetWitness Knowledge Base: https://community.netwitness.com/t5/netwitness-knowledge-base/tkb-p/netwitness-knowledge-base
- See Troubleshooting section in the Guides.
- See also NetWitness® Platform Blog Posts.
- If you need further assistance, contact NetWitness Support.

Contact NetWitness Support

If you contact NetWitness Support, you should be at your computer. Be prepared to provide the following information:

- The version number of the NetWitness Platform product or application you are using.
- The type of hardware you are using.

Use the following contact information if you have any questions or need assistance.

NetWitness Community	https://community.netwitness.com/ In the main menu, click **Support** > **Case Portal** > **View My Cases**.
International Contacts (How to Contact NetWitness Support)	https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897
Community	https://community.netwitness.com/t5/support/ct-p/support
NW Update	update.netwitness.com
LiveUi	live.netwitness.com

Pre Upgrade Checks

You must run the pre-upgrade checks before you upgrade to NetWitness 12.0.0.0 to identify any issues that may result in upgrade failure.

To run the pre-upgrade checks, perform the following:

1. Log in to Admin console.
2. Run the following command:
   `nw-precheck-tool upgrade-checklist`

   The pre-upgrade checks verifies the following:

   - **Security Client File Check** - Ensures `security-client-amqp.yml` file is not present.
   - **Node-0 NW Service-id Status** - Ensures all the service-id is intact with all the different services in Node 0.
   - **Broker Service Trustpeer Symlink** - Ensures Broker symlink file `(/etc/netwitness/ng/broker/trustpeers/)` is not broken.
   - **Node-0 NW Services Status** - Checks the status of all the services in Node 0.
   - **Yum External Repo Check** - Ensures external repos are not available and not enabled.
   - **Node-0 RPM DB Index Check** - Checks if the RPM DB is corrupted or not.
   - **Salt Master Communication** - Verifies the salt communication from Node 0 to all the Nodes.
   - **Node-0 Certificates Check** - Checks if any certificates are missing, expired, or invalid issuer type.
   - **Mongo Authentication** - Validates the `deploy_admin` credentials fetched from `security-cli-client` using Mongo client.
   - **Rabbitmq Authentication** - Validates the `deploy_admin` credentials fetched from `security-cli-client` using RabbitMQ.
   - **(Component Hosts) Node X NW Service Status** - Verifies the status of services (Active or In Active) on all the Node X.
   - **(Component Hosts) Node X Certificates Check** - Checks the certificate expiry, missing, corrupted, and issuer mismatch in all categories of Node X.
   - **Nodes CPU-Memory Info** - Provides CPU and Memory details of all the nodes along with the real-time available memory.
   - **(Admin Server) Node 0 File System Utilization** - Verifies the disk partition utilization of `/var/netwitness/mongo`, `/var/netwitness`, and `root` on Node 0.

- **(Component Hosts) Node X File System Utilization** - Verifies the disk partition utilization of `/var/netwitness/mongo`, `/var/netwitness`, and `root` for ESA Primary and Endpoint Log Hybrid services on Node X.
- **Mongo File (ESAPrimary)** - Checks the ESA Primary node in the system or stack and verifies the permission mode of Mongo file.
- **Orchestration Server Normal Mode** - Checks if the orchestration service is running in normal or safe mode.
- **(Admin Server) Node 0 Init status** - Checks if there are any issues that might fail init process.
- **(Admin Server) Node 0 closed ports** - Checks if the service ports required for NetWitness services are open and listening on Node 0.
- **(Component Hosts) Node X closed ports** - Checks if the service ports required for NetWitness services are open and listening on Node X.
- **Fips Mode Check** - Checks to ensure that the Fips mode is disabled (set to false) before and after upgrade.
- **Node-X RPM DB Index Check** - Checks for the status of RPM DB on Node-X to make sure it is not corrupted.
- **Node-Z Yum Proxy Check** - Checks for the existence of yum.conf file and availability of proxy within the file on Node -Z.
- **Node-X Yum Proxy Check** - Checks for the existence of yum.conf file and availability of proxy within the file on Node -X.
- **Host Info Check Probe** - Checks if the required fields of information of all the hosts in the system (Host IP, Hostname, Installed Services, and Raw Version) are available.
- **Node-Z Cipher Check Probe** - Checks if the required ciphers are available in the location `/etc/rabbitmq/rabbitmq.config` on Node-0.
- **Node-X Cipher Check Probe** - Checks if the required ciphers are available in the location `/etc/rabbitmq/rabbitmq.config` on all Node-X.

Note: We recommend you to disable FIPS before upgrading and re-enable after a successful upgrade, to avoid appliance boot issues. To disable, run the following commands:

```
manage-stig-controls --disable-control-groups 3 --host-all
grub2-mkconfig -o /boot/grub2/grub.cfg
```

Upgrade Preparation Tasks

Complete the following tasks to prepare for the upgrade to NetWitness 12.0.0.0

Warning: The Dell S4 and S4s appliances reached the End of Life (EOL) in June 2021. We recommend that you discontinue installation or upgrade activities on these and upgrade to new hardware.

Task 1. (Optional) Remove Legacy Package Repositories

Perform this task to free up space by removing unused repositories from previous releases from your system.

1. Determine the version of the oldest NetWitness Platform host in your environment by reviewing the host list in the Admin user interface, or by running the following command on the NW Server: `upgrade-cli-client --list`
2. You can safely remove all legacy package repository folders located at `/var/netwitness/common/repo/<version>` on the NW Server for all versions prior the baseline major release version of the oldest active host in the environment.
   - If the oldest host version is 11.7.x.x, you can safely remove 11.0.x.x, 11.1.x.x, 11.2.x.x, 11.3.x.x, 11.4.x.x, 11.5.x.x, and 11.6.x.x repository folders. However, do not remove repository versions greater than or equal to 11.7.0.0.
   - If the oldest host version is 11.3.x.x, you can safely remove 11.0.x.x, 11.1.x.x, and 11.2.x.x repository folders. However, do not remove repository versions greater than or equal to 11.3.0.0.

Task 2. Backup and Remove the Rotated RabbitMQ Logs

Before upgrading to 12.0.0.0, you must remove the old RabbitMQ logs and free up the space in /var/log mount disk. Follow the below procedure to free up the space in /var/log mount disk.

1. Backup the rotated RabbitMQ logs into var/netwitness directory. Do the following.

   ```
   mkdir /var/netwitness/rabbitmq_logsbkp
   scp -r /var/log/rabbitmq/ /var/netwitness/rabbitmq_logsbkp
   ```

2. Remove the rotated RabbitMQ logs from /var/log/rabbitmq pre-upgrade. Do the following.

   ```
   cd /var/log/rabbitmq
   rm -f rabbit\@<sa-uuid>.log.*
   rm -f rabbit\@<sa-uuid>_upgrade.log.*
   rm -f *.gz
   rm -f rabbit@<sa-uuid>.log-*
   ```

Note:
- This procedure must be performed only once before upgrading to 12.0.0.0 Post-upgrade, the RabbitMQ service automatically handles the log rotation.
- The command `rm -f rabbit\@<sa-uuid>.log.*` is used to clean up the old uncompressed logs such as log.1, log.2, and log.3.
- The command `rm -f rabbit\@<sa-uuid>_upgrade.log.*` is used to clean up the old uncompressed upgrade logs.
- The command `rm -f *.gz` is used to clean up the old compressed logs.
- The command `rm -f rabbit@<sa-uuid>.log-*` is used to clean up the old uncompressed logs rotated with logrotate.

Task 3. Uninstall the Security Analytics l10n language pack

Before you upgrade from 11.6.x.x to 11.7.x.x or 12.0.0.0 version, you must uninstall the Security Analytics l10n language pack.

Upgrade Options

Upgrade the systems in your environment in the following order:

1. NW Server hosts
2. Analyst UI hosts
3. ESA Primary hosts
4. ESA Secondary hosts
5. The rest of your component hosts

Note: NW Server, Analyst UI, and ESA Primary and Secondary hosts must all be upgraded on the same day. The rest of your component hosts can be upgraded on the same day or later.

For information about all the host types in NetWitness, see the *Host and Services Getting Started Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Important Notes - Read This First

Synchronize Time on Component Hosts with NW Server Host

Before upgrading your hosts, make sure that the time on each host is synchronized with the time on the NetWitness Server.
To synchronize the time, do one of the following:

- Configure the NTP Server. For more information, see"Configure NTP Servers" in the *System Configuration Guide*.
- Perform the following steps on each host:
  a. SSH to a component host.
  b. Run the following commands.
  ```
  systemctl stop ntpd
  ntpdate nw-node-zero
  systemctl start ntpd
  ```

Mixed Mode Unsupported for ESA Hosts

Mixed mode is not supported for ESA hosts in NetWitness Platform version. The NetWitness server, ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

Endpoint Hybrid Systems Not Supported

For NetWitness Endpoint customers only, Endpoint Hybrid is not supported in 11.3.0.0 and later releases.
If you have deployed an Endpoint Hybrid host in 11.2.x.x and did not install an Endpoint Log Hybrid host in 11.3.x.x or 11.4.x.x, you must install an Endpoint Log Hybrid host in 12.0.0.0. See the *Physical Host Installation Guide for NetWitness Platform* or the *Virtual Host Installation Guide for NetWitness Platform* for instructions on how to install an 12.0.0.0 Endpoint Log Hybrid on a physical host.

Respond Server Service Not Enabled Until NW Server and Primary ESA Host Upgraded to 12.0.0.0

After upgrading the primary NW Server (including the Respond Server service), the Respond Server service is not automatically re-enabled until after the Primary ESA host is also upgraded to 12.0.0.0 The Respond post-upgrade tasks only apply after the Respond Server service is upgraded and is in the enabled state.

Deploy_Admin Password Guidelines

In NetWitness Platform version 11.6 or Later, deployment account password (only on node-zero) must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`.

If `deploy_admin` password is changed on Primary NW Server, it must be changed in the Warm Standby Server if it exists

Additional Post Upgrade Steps for 12.0.0.0 Version with Legacy Windows Log Collector

For 12.0.0.0 version with Legacy Windows Log Collector, you should perform few additional post upgrade tasks. Refer to Legacy Windows Log Collection section in Post Upgrade Tasks for these additional post upgrade tasks.

Upgrade Options

You can choose one of the following upgrade methods based on your Internet connectivity. They are listed in the order recommended by NetWitness.

- Option 1: User Interface Method with Connectivity to the Internet
- Option 2: User Interface with No Connectivity to the Internet
- Option 3: Command Line Interface (CLI) with No Connectivity to the Internet
- Option 4 (Optional): Pre-Stage Upgrade Repo

The following rules apply when you are upgrading hosts for all of these upgrade methods:

- You must upgrade the NW Server host first.
- You can only apply a version that is compatible with the existing host version.
- The NW Server, ESA primary, ESA secondary, and Analyst UI hosts must all be on the same NetWitness Platform version.
- Add Warm Standby (if exists) to the list of hosts and it must be on same version as NetWitness Platform.

Option 1: User Interface Method with Connectivity to the Internet

You can use this method if the NW Server host is connected to Live Services and if you are able to obtain the package.

Prerequisites

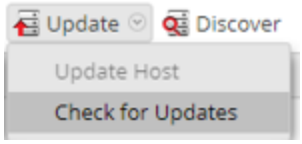
Make sure that:

1. The **Automatically download information about new upgrades every day** option is selected and is applied in **(Admin) > System > Updates**.
2. Updates are available. Go to **(Admin) > Hosts > Update > Check for Updates** to check for updates. The Host view displays the **Update Available** status.
3. 12.0.0.0 is available in the **Update Version** column.

Procedure

If you are upgrading from 11.6.x.x and 11.7.0.x to 12.0.0.0, follow the steps below:

1. Go to **(Admin) > Hosts**.
2. Select the NW Server (`nw-server`) host.
3. Check for the latest updates.



4. Update Available is displayed in the **Status** column if you have a version update in your Local Update Repository for the selected host.
5. Select **12.0.0.0** from the **Update Version** column. If you:
   - Want to view a dialog with the major features in the upgrade and information on the updates, click the information icon ( ) to the right of the upgrade version number.

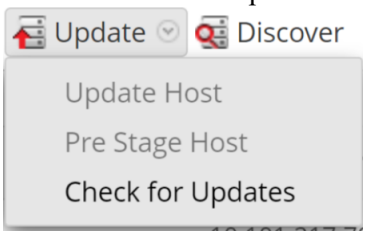
- Cannot find the version you want, select **Update > Check for Updates** to check the repository for any available updates. If an update is available, the message "New updates are available" is displayed and the **Status** column updates automatically to show Update Available. By default, only supported updates for the selected host are displayed.

6. Click **Update** > **Update Host** from the toolbar.
7. Click **Begin Update**.
8. Click **Reboot Host**.
9. Repeat steps 6 to 8 for other hosts.

Note: You can select multiple hosts to upgrade at the same time only after updating and rebooting the NW Server host. All ESA, Endpoint, and Malware Analysis hosts should be upgraded to the same version as that of the NW Server host.

If you are upgrading from 11.7.1.0 and 11.7.1.1 to 12.0.0.0, follow the steps below:

1. Go to **(Admin) > Hosts**.
2. Select the NW Server (`nw-server`) host.
3. Check for the latest updates.



4. Update Available is displayed in the **Status** column if you have a version update in your Local Update Repository for the selected host.
5. Select **12.0.0.0** from the **Update Version** column. If you:
   - Want to view a dialog with the major features in the upgrade and information on the updates, click the information icon ( ) to the right of the upgrade version number.
   - Cannot find the version you want, select **Update > Check for Updates** to check the repository for any available updates. If an update is available, the message "New updates are available" is displayed and the **Status** column updates automatically to show Update Available. By default, only supported updates for the selected host are displayed.
6. Click **Update** > **Update Host** from the toolbar.
7. Click **Begin Update**.
8. Click **Reboot Host**.
9. Repeat steps 6 to 8 for other hosts.

Note: You can select multiple hosts to upgrade at the same time only after updating and rebooting the NW Server host. All ESA, Endpoint, and Malware Analysis hosts should be upgraded to the same version as that of the NW Server host.

Option 2: User Interface with No Connectivity to the Internet

Task 1. Populate Staging Folder (`/var/lib/netwitness/common/update-stage/`) with Version Upgrade Files

1. Download the upgrade package `netwitness-12.0.0.0.zip` from NetWitness Community (https://community.netwitness.com/) > **Downloads** > **NetWitness Platform** > **Version 12.0** to a local directory:
2. SSH to the NW Server host.
3. Copy `netwitness-12.0.0.0.zip` from the local directory to the `/var/lib/netwitness/common/update-stage/` staging folder.
   For example:

```
sudo cp /tmp/netwitness-12.0.0.0.zip /var/lib/netwitness/common/update-stage/
```

If you are logged as a root user you can ignore sudo in the command, for example,

```
cp /tmp/netwitness-12.0.0.0.zip /var/lib/netwitness/common/update-stage/
```

Note: NetWitness Platform unzips the file automatically.

Task 2. Apply Upgrades from the Staging Area to Each Host

Caution: You must upgrade the NW Server host before upgrading any non-NW Server host.

1. Log in to NetWitness.
2. Go to  **(Admin) > Hosts**.

   Note: If you are already on the  **(Admin) > Hosts** page and the **Check for Updates** option (**Update** > **Check for Updates**) is grayed out, refresh the page from the browser to check for the updates.
3. Check for updates and wait for the upgrade packages to be copied, validated, and ready to be initialized.

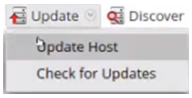
   "Ready to initialize packages" is displayed if:

   - NetWitness Platform can access the upgrade package.
   - The package is complete and has no errors.

   Refer to Troubleshooting Version Installations and Updates for instructions on how to troubleshoot errors (for example, "Error deploying version <version-number>" and "Missing the following update package(s)," are displayed in the **Initiate Update Package for RSA NetWitness Platform** dialog.)
4. Click **Initialize Update**.

It takes some time to initialize the packages because the files are large and need to be unzipped. The time varies depending on how the host is configured.
After the initialization is successful, the **Status** column displays Update Available and you complete the rest of the steps in this procedure to finish the upgrade of the host.

5. Click **Update** > **Update Hosts** from the toolbar.



6. Click **Begin Update** from the Update Available dialog.
   After the host is upgraded, it prompts you to reboot the host.
7. Click **Reboot Host** from the toolbar.

Option 3: Command Line Interface (CLI) with No Connectivity to the Internet

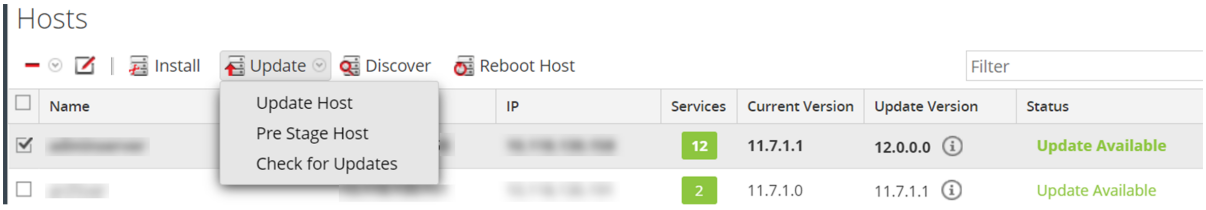
Follow the instructions in Appendix A. Offline Method (No Connectivity to Live Services) - Command Line Interface .

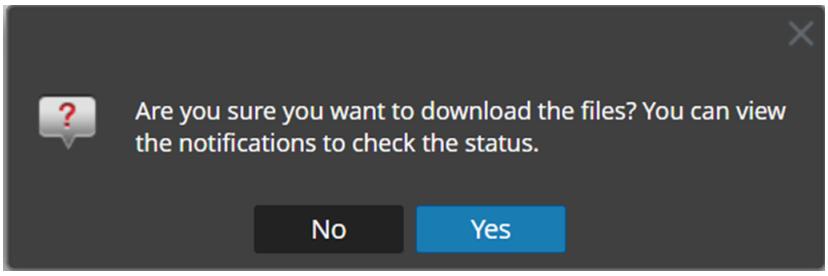
Option 4 (Optional): Pre-Stage Upgrade Repo

You can pre-stage the upgrade repository by downloading the required packages (.zip) without affecting the system. This minimizes the upgrade downtime and ensures the upgrade is completed within the planned time.

Procedure

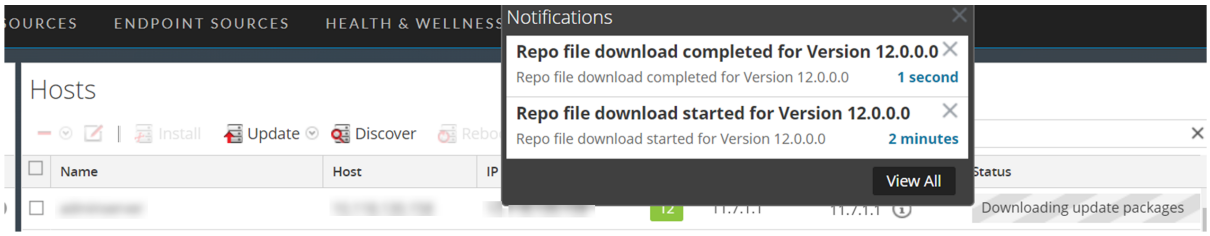
1. Go to  **(Admin) > Hosts**.
2. Click **Update > Check for Updates** from the toolbar.
   All possible update versions will be displayed in the Versions drop-down list.
3. Click **Update** > **Pre Stage Host** and select the version in the update version column.
   A confirmation message for downloading the files is displayed.





4. Click **Yes** to download the upgrade packages to the repo.
5. Verify the status of the download in the notifications tray as shown below.

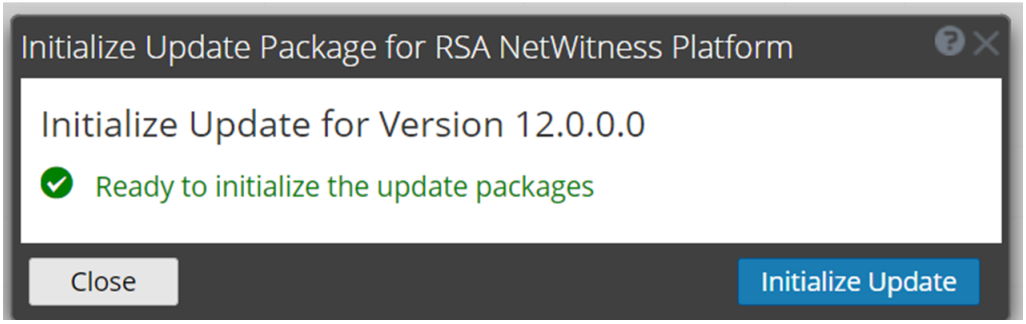
   The **Pre Stage Host** and **Upgrade Host** will be disabled until pre stage is completed.



Note: The current version and the update version in the UI will be the same during the pre stage as it is not the actual update. This is because only the repo files are downloaded and no actual upgrade is done. The version will change only after upgrade.

6. If the download is successful, **Check for Updates** again to start the initialization.
7. Click **Initialize Update**.

   The initialization of the package will take some time as the files are large and will need to be unzipped.



IMPORTANT: Pre Stage Repo preparation steps from 1 to 4 can be performed at any time. However, from steps 5 to 8 the upgrade process begins and you must NOT reboot the host or restart the jetty server during this time as it will corrupt the .ZIP files.

8. Check the status of initialization in the notifications tray.
9. After the initialization is completed successfully, click **Update** > **Update Host**.

   After the host is updated, you will be prompted to reboot the host.
10. Set up the host and reboot the host.

Post Upgrade Tasks

After you upgrade to 12.0.0.0, NetWitness has several new features in the user interface. Complete the tasks that apply to the hosts in your environment.

- General
- Event Stream Analysis (ESA)
- New Health and Wellness
- Investigate
- Respond
- Reference Log Decoder
- Legacy Windows Log Collector
- User Entity Behavior Analytics

General

(Conditional) Configure NAT-Based IP Addresses

If you have a host, such as a VLC, that requires a NAT-based IP address in order to connect to the NW Server host, you must update the host configuration with the following steps.

1. Log in to the host that requires the use of NAT IP addresses, using the console or SSH.
2. Run the following command:
   `nw-manage --enable-nat-usage`
3. To set the NAT address for the NW Server:
   a. Log into the NW Server using the console or SSH.
   b. Run the following command:
   `nw-manage -update-host --host-id <UUID of NW Server> --ipv4-public <NAT IP of NW Server>`

   Note: You can find the UUID and view the current NAT IP address of the host by running `nw-manage --list-hosts`.

(Conditional - For Warm-Standby Hosts Only) Register the Secondary IP Address of Warm-Standby Hosts

The Warm-Standby server must be upgraded to 12.0.0.0 before completing the following steps.

1. Log in to the NW Server using the console or SSH.
2. Run the following command:
   `nw-manage --add-nws-secondary-ip --ipv4 <ip address of Warm/Standby Server>`

   Note: If the Warm-Standby server requires a NAT-based IP address (IPv4-public) for any host to access it during failover, the NAT IP address must also be registered by running the following command: `nw-manage --add-nws-secondary-ip --ipv4 <NAT-based IP address of Warm Standby Server>`
3. Verify the correct Warm Standby host IP address value by running the following command:
   `nw-manage --get-nws-secondary-ip`

Review Contents of /etc/hosts.user for Obsolete Host Entries

After upgrading the NW Server host or a component host, review the contents of the `/etc/hosts.user` file for any obsolete host entries. The `/etc/hosts.user` file contains system and user-generated entries that are not managed by NetWitness Platform. However, entries from `/etc/hosts.user` are merged with NetWitness Platform-generated host mappings to create and update `/etc/hosts`. To avoid conflicts with NetWitness Platform-generated mappings, and to avoid generating connectivity errors resulting from an IP address change, NetWitness recommends that you remove any entries in `/etc/hosts.user` that include a non-loopback IP address of a NetWitness Platform host.

After updating `/etc/hosts.user`, you must refresh the system by running the following command:
`nw-manage --refresh-host --host-key <ID, IP, hostname or display name of host>`

Jetty Configuration

For Jetty Configuration and related information, see **Manage Custom Host Entries** topic in the *System Maintenance Guide*.

Reconfigure DNS Servers

By default, a component host upgraded from 11.4 or earlier is configured with the same system DNS server as the NW Server. If this component host requires a different system DNS address, see "Change Host Network Configuration" in the *System Maintenance Guide* for instructions.

Make Sure Services Have Restarted and Are Capturing and Aggregating Data

Make sure that services have restarted and are capturing data (this depends on whether or not you have auto-start enabled).

If required, restart data capture and aggregation for the following services:

- Decoder
- Log Decoder

- Broker
- Concentrator
- Archiver

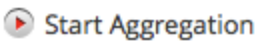
Start Network Capture

1. In the NetWitness Platform menu, go to  (Admin) > **Services**.
   The Services view is displayed.
2. Select each **Decoder** service.
3. Under  (actions), select **View** > **System**.
4. In the toolbar, click 

Start Log Capture

1. In the NetWitness Platform menu, go to  (Admin) > **Services**.
   The Services view is displayed.
2. Select each **Log Decoder** service.
3. Under  (actions), select **View** > **System**.
4. In the toolbar, click 

Start Aggregation

1. In the NetWitness Platform menu, go to  (Admin) > **Services**.
   The Services view is displayed.
2. For each **Concentrator**, **Broker**, and **Archiver** service:
   a. Select the service.
   b. Under  (actions), select **View** > **Config**.
   c. In the toolbar, click 
3. Event Stream Analysis (ESA)

Note: Mixed mode is not supported for ESA hosts in NetWitness Platform version 11.6 and later. The NetWitness server, ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

There are no required post-upgrade tasks for ESA. For ESA troubleshooting, see ESA Troubleshooting Information.

If you want to add support for Endpoint, UEBA, and Live content rules, you must update the `multi-valued` and `single-valued` parameter meta keys on the ESA Correlation service to include all the required meta keys. It is not necessary to make these adjustments during the upgrade; you can make the adjustments later at a convenient time. For detailed information and instructions, see "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*

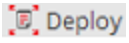
New Health and Wellness

Note: New Health and Wellness in 11.5 replaces Next GEN Health and Wellness (BETA) in 11.4.x.x.

Deploy the New Health and Wellness Content from Live

After you upgrade from version 11.4.x.x to 12.0.0.0, New Health and Wellness content is not updated. To use the latest (default) content, you must deploy the content through NetWitness Live Services.

Note: NetWitness recommends you to take a copy of 11.4.x.x Health and Wellness content before you deploy the content from NetWitness Live Services, as it overwrites the existing content.

1. Log in to NetWitness Platform UI.
2. Click  (**CONFIGURE**) > **LIVE CONTENT**.
3. In the **Search Criteria** panel, select the **Resource Types** as:
   - Health and Wellness Dashboards
   - Health and Wellness Monitors
4. Click **Search**.
5. In the **Matching Resources** view, select the checkbox to the left of the resources that you want to deploy.
6. In the **Matching Resources** toolbar, click 

.
7. In the **Deployment Wizard** > **Resources** tab, click **Next**.
8. In the **Services** tab, select the Metrics Server service.
9. Click **Next**.
10. Click **Deploy**.
    The **Deploy** page is displayed. The Progress bar turns green when you have successfully deployed the resources to the selected services.
11. Click **Close**.

(Optional) Update UUID of New Health and Wellness Host to Update Service Configuration Documents

If you have configured services for New Health and Wellness from `nw-shell` using set-config API and upgrade NetWitness Platform version from 11.4.x.x to 12.0.0.0, you must update IP with UUID for a host on which New Health and Wellness is installed.

1. SSH to Admin Server.
2. Check the UUID of a host on which New Health and Wellness is installed using the command:
   `orchestration-cli-client --list-hosts`

   This lists NetWitness Platform hosts along with the respective UUIDs. Make a note of the UUID of host on which New Health and Wellness is installed.
3. Identify the services on which set-config is invoked using the command:
   `mongo localhost/metrics-server -u deploy_admin -p <deployment_password> --authenticationDatabase admin --eval 'db.metric_config.find({ "createdBy": { $ne: "system" }})'`
   This will list the configuration documents of the services on which set-config is invoked.

Note: If no service documents are listed which means no services are configured before the upgrade, so you can ignore the remaining steps.

4. In the configuration file, update the service document “host” field by replacing IP with the UUID of the host on which New Health and Wellness is installed.
   For example, update the host IP with UUID as below:

```
"host" : "e086511c-121c-4e66-95a3-e87e27b12acb"
```

5. Log in to `nw-shell` using the command:
   `nw-shell`
6. Connect to metrics-server service using the command:
   `connect --service metrics-server`
7. Enter the log in command:
   `login`
8. Enter the admin username and password.
9. Go to `/rsa/metrics/elastic/set-config` and invoke configuration files using the command:
   `invoke --file /<absolute_path_of_service_config_file>`

Investigate

(Conditional - For Custom Roles Only) Adjust investigate-server Permissions for Custom User Roles

After upgrading to Version 12.0.0.0, the built-in user roles for analysts (and others) using Investigate have the `investigate-server.event.filter` permission enabled, but the upgrade process does not enable the permission for custom user roles. Users who are assigned a custom user role that does not have this permission enabled cannot see the Filter Events panel, a new panel in 12.0.0.0 where they can drill into metadata.

Note: The built-in user roles for analysts using Investigate have three additional permissions added in Version 11.4 enabled, but the upgrade process does not enable the permissions for custom user roles. Users who are assigned a custom user role that does not have these permissions cannot see the Navigate view and Legacy Events view in the Investigate menu. The three permissions that need to be enabled for custom user roles are:
`investigate-server.columngroup.read`, `investigate-server.metagroup.read`, and `investigate-server.profile.read`

To enable the permissions for a user role:

1. Go to  (Admin) > **Security** and click the **Roles** tab.
2. Select the custom user role that needs to be edited and click  (edit icon).
3. In the Edit Role dialog, ensure that these four permissions are enabled:
   `investigate-server.event.filter`
   `investigate-server.columngroup.read`
   `investigate-server.metagroup.read`
   `investigate-server.profile.read`
4. Click **Save** to save your changes. When analysts with the custom user role log in to the NetWitness Platform, the changes are in effect.

Respond

The Primary ESA server must be upgraded to 12.0.0.0 before you can complete these tasks.

Note: After upgrading the primary NW Server (including the Respond Server service), the Respond Server service is not automatically re-enabled until after the Primary ESA host is also upgraded to 12.0.0.0 The Respond post-upgrade tasks only apply after the Respond Server service is upgraded and is in the enabled state.

(Conditional) Restore Any Respond Service Custom Keys in the Aggregation Rule Schema

Note: If you did not manually customize the incident aggregation rule schema, you can skip this task.

If you added custom keys in the `var/lib/netwitness/respond-server/data/aggregation_rule_schema.json` file for use in the `groupBy` clause for 12.0.0.0, modify the `/var/lib/netwitness/respond-server/data/aggregation_rule_schema.json` file and add the custom keys from the automatic backup file.

The backup file is located in `/var/lib/netwitness/respond-server/data` and it is in the following format:
`aggregation_rule_schema.json.bak-<time of the backup>`

Reference Log Decoder

For full functionality, make sure your reference Log Decoder is at 11.6 or later. If you never set up a reference Log Decoder, there is no need to take action. For details, see the *Log Parser Customization Guide*.

Legacy Windows Log Collector

Update the Legacy Windows Log Collector UUID

After upgrading to 12.0.0.0, for each Legacy Windows Log Collector configured in your environment, run the following command on the NW Server:

```
wlc-cli-client --update-to-uuid --host <WLC host address>
```

Refresh Legacy Windows Log Collector Certificates with Updated SA Certificates

Post Upgrade Steps:

1. Execute the following command in SA:
   a. `wlc-cli-client --host-display-name` **`hostDisplayName`** `--service-display-name` **`serviceDisplayName`** `--host` **`WLChostIPAddress`** `--port 50101 --use-ssl false`

      Enter following information:

      i. **Legacy Windows Log Collector REST Username and Legacy Windows Log Collector REST Password**: Enter the admin credentials for the Legacy Windows Log Collector.

      ii. **Security Server Username and Security Server Password**: Enter admin credentials for NetWitness.
2. Restart the system.

User Entity Behavior Analytics

IMPORTANT: Before the upgrade, if you encountered and resolved the task failure issues, then after the upgrade, you must replace the `authentication.json` file before you run the post-upgrade tasks. The task failure issues in Airflow and their solutions are described in the 'Troubleshooting' topic of the *UEBA Configuration Guide*.

IMPORTANT: Every UEBA deployment when upgraded requires additional steps to complete the upgrade process. When you upgrade from 11.6.x to 11.6.x.x, you must follow UEBA instructions in the Upgrade Guide for 11.6.x.x, before you upgrade to 11.7.x.

Note: When you upgrade to 11.7.1.0 from 11.4.x.x, you don't need to rerun the UEBA system for the last 28 days, if you don't update the current processing schemas. When you upgrade to 11.7.1.0 from a version prior to 11.4.x, the UEBA system runs a rerun automatically.

1. (For Virtual Machines Only) Update the airflow parallelism on VM.
   If the UEBA system is running on VM, update the airflow parallelism to be 64 by running the following command as root from the UEBA host.

   ```
   sed -i "s|parallelism = 256|parallelism = 64|g"
   /var/netwitness/presidio/airflow/airflow.cfg
   ```

   Note: Copy this command in a single line.

2. Update the UEBA configuration using the following command from the UEBA machine.

   ```
   source /etc/sysconfig/airflow

   source $AIRFLOW_VENV/bin/activate

   OWB_ALLOW_NON_FIPS=on python
   /var/netwitness/presidio/airflow/venv/lib/python2.7/site-packages/presidio_
   workflows-1.0-py2.7.egg/presidio/resources/rerun_ueba_server_config.py
   ```

3. (Optional) Update the UEBA processing schema, if needed.

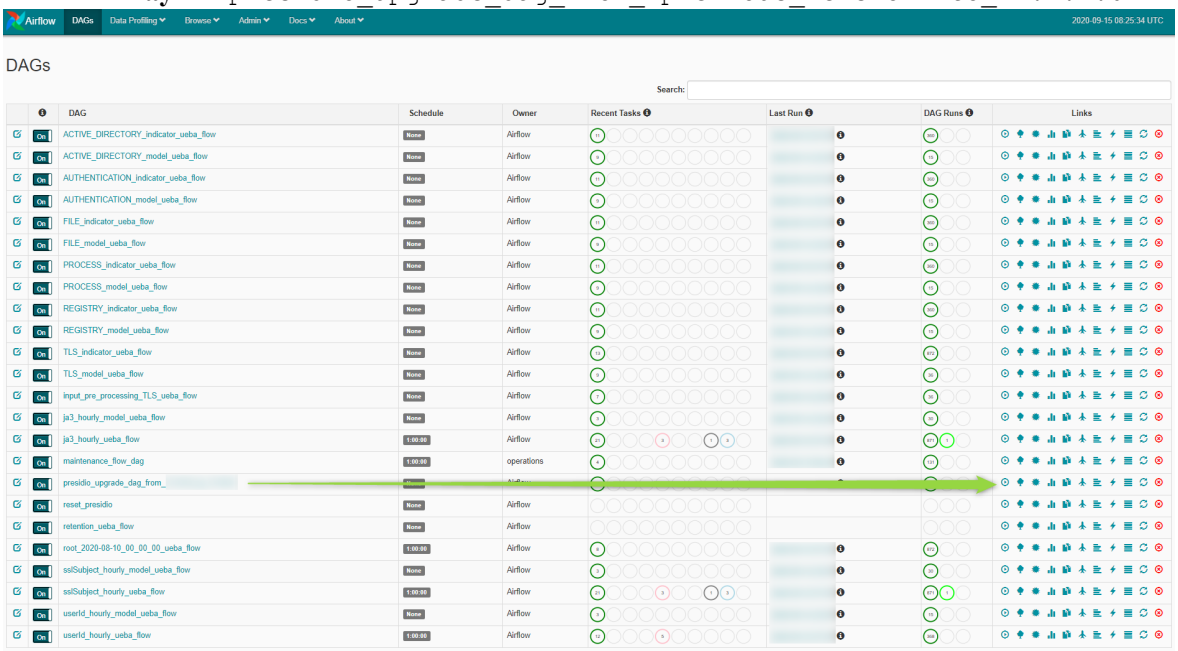
   NetWitness recommends that the UEBA start date is set to 28 days earlier than the current date. For UEBA systems that intend to process TLS data, you must make sure that the start date is set to no later than 14 days earlier than the current date.

   For more information, see the "reset-presidio script" section in the *UEBA Configuration Guide*.

4. Run the airflow upgrade DAG.

   - Go to Airflow main page `https://<UEBA-host-name>/admin`
   - Enter the admin username and password.

- Click the **Play** in `presidio_upgrade_dag_from_<previous_version> to_11.7.1.0`.



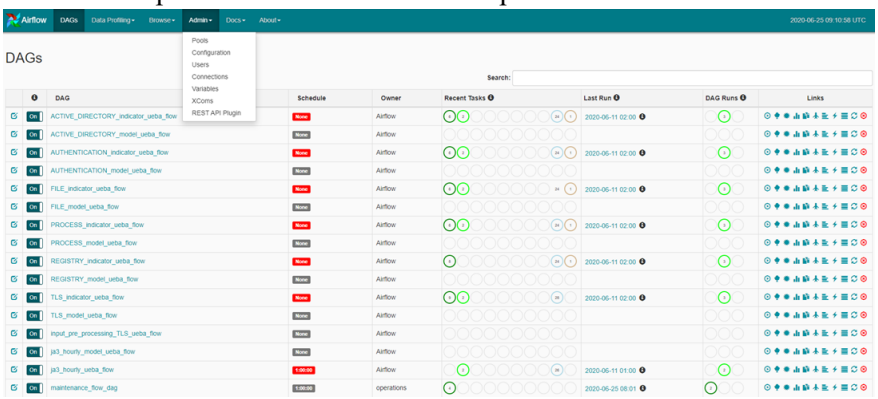
> **Note:** A light green circle will appear next to the upgrade DAG row during the upgrade. If the upgrade process is completed successfully the light green circle changes to green. If the upgrade process fails, the light green circle changes to red.

5. Set the appropriate "Boot Jar Pools" slots:

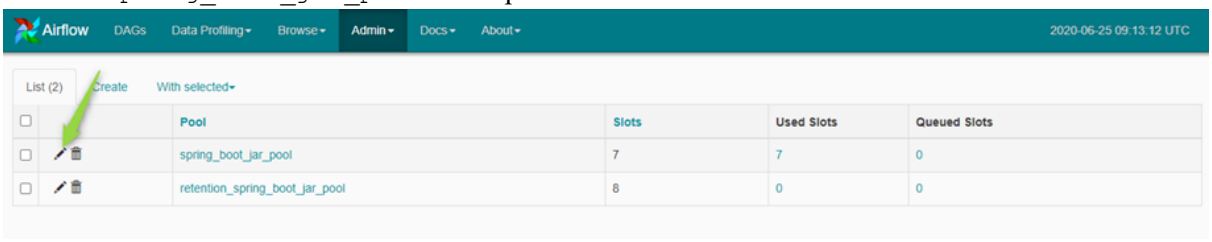
   - **Physical Appliance**: Update the `spring_boot_jar_pool` slot value be 18.
   - **Virtual Appliance**: Update the `spring_boot_jar_pool` slot value to 22.
     To update the **Spring Boot Jar Pools** slots, go to the Airflow main page, tap the **Admin** tab at the top bar and tap **Pools**.

   a. To access the Airflow UI, go to `https://<UEBA_host>/admin` and enter the credentials.
      User: admin
      Password: The environment deploy admin password.

   b. Click on the pencil mark of the Pools to update the slot values.



6. Edit the `spring_boot_jar_pool` and update the slots amount to 5.



Endpoint Upgrade Tasks

Install the 12.0.0.0 Relay Server

If you have configured Relay Server, perform the following:

1. You must upgrade the Relay Server to 12.0.0.0 by downloading the Relay Server installer from the upgraded Endpoint Server. For more information see "(Optional) Installing and Configuring Relay Server" section in the *Endpoint Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
2. Restart the Endpoint Server using the command:

   ```
   systemctl restart rsa-nw-endpoint-server
   ```

Upgrade Endpoint Agents

See "Upgrade Agents" in the *Endpoint Agent Installation Guide for NetWitness Platform 12.0* for instructions on how to upgrade agents.

Appendix A. Offline Method (No Connectivity to Live Services) - Command Line Interface

You can use this method if the NW Server host is not connected to Live Services.

Prerequisites

Make sure that you have downloaded the `netwitness-12.0.0.0.zip` file from NetWitness Community (https://community.netwitness.com/) > **Products > NetWitness Platform** > **Downloads** > **Version 12.0** > **Full Product Downloads** to a local directory:

Procedure

You must perform the upgrade steps for NW Server hosts and for component servers.

Note: If you copy and paste the commands from PDF to Linux SSH terminal, the characters do not work. It is recommended to type the commands.

1. Stage the 12.0.0.0 files to prepare them for the upgrade.
   - Log into the NW Server as `root` and create the following directory:
     `/tmp/upgrade/12.0.0.0`
     and then copy the package zip files to the `/root` directory of the NW Server and extract the package files from `/root` to the appropriate directories using the following command: `unzip netwitness-12.0.0.0.zip -d /tmp/upgrade/12.0.0.0`

   Note: If you copied the .zip file to the created staging directory to unzip, make sure that you delete the initial .zip file that you copied to the staging location after you extract it.
2. Initialize the upgrade, using the following command:
   `upgrade-cli-client --init --version 12.0.0.0 --stage-dir /tmp/upgrade`
3. Upgrade the NW Server host, using the following command:
   `upgrade-cli-client --upgrade --host-key <ID / display name / (hostname/ IP address)> --version 12.0.0.0`
4. When the NW Server host upgrade is successful, reboot the host from NetWitness Platform user interface in the Hosts view.
5. Repeat steps 3 and 4 for other component hosts.

Note: You can check versions of all the hosts, using the command `upgrade-cli-client --list` on the NW Server host. If you want to view the help content of `upgrade-cli-client`, use the command `upgrade-cli-client --help`.

Note: If the following error is displayed during the upgrade process:
`2017-11-02 20:13:26.580 ERROR 7994 — [ 127.0.0.1:5671] o.s.a.r.c.CachingConnectionFactory : Channel shutdown: connection error; protocol method: #method<connection.close>(reply-code=320, reply-text=CONNECTION_FORCED - broker forced connection closure with reason 'shutdown', class-id=0, method-id=0)`
the service pack will install correctly. No action is required. If you encounter additional errors when updating a host to a new version, contact Customer Support for assistance.

External Repo Instructions for CLI upgrade

1. Stage the 12.0.0.0 files to prepare them for the upgrade.
2. Log into the NW Server as `root` and create the following directory: `/tmp/upgrade/12.0.0.0` and then copy the package zip files to the `/root` directory of the NW Server and extract the package files from `/root` to the appropriate directories using the following command: `unzip netwitness-12.0.0.0.zip -d /tmp/upgrade/12.0.0.0`

   Note: If you copied the .zip file to the created staging directory to unzip, make sure that you delete the initial .zip file that you copied to the staging location after you extract it.
3. Initialize the upgrade, using the following command:
   `upgrade-cli-client --init --version 12.0.0.0 --stage-dir /tmp/upgrade`
4. Upgrade the NW Server host using the following command:
   `upgrade-cli-client --upgrade --host-key <ID / display name / (hostname/ IP address)> --version 12.0.0.0`
5. When the NW Server host upgrade is successful, reboot the host from NetWitness UI.
6. Repeat steps 3 and 4 for other component hosts.

Note: You can check versions of all the hosts, using the command `upgrade-cli-client --list` on the NW Server host. If you want to view the help content of `upgrade-cli-client`, use the command `upgrade-cli-client --help`.

Note: If the following error displays during the upgrade process:
`2017-11-02 20:13:26.580 ERROR 7994 — [ 127.0.0.1:5671] o.s.a.r.c.CachingConnectionFactory : Channel shutdown: connection error; protocol method: #method<connection.close>(reply-code=320, reply-text=CONNECTION_FORCED - broker forced connection closure with reason 'shutdown', class-id=0, method-id=0)`
the service pack will install correctly. No action is required. If you encounter additional errors when updating a host to a new version, contact Customer Support for assistance.

Appendix B. Set Up External Repo

Complete the following procedure to set up an external repository (Repo).

Note: 1.) You need an unzip utility installed on the host to complete this procedure. 2.) You must know how to create a web server before you complete the following procedure.

1. (Conditional) Complete this step if you have an external repo and you want to override it.
   - Case 1: You bootstrapped the host from an external repo and you want to upgrade using a local repo on the Admin Server.
     a. Create the `/etc/netwitness/platform/repobase` file.
        `vi /etc/netwitness/platform/repobase`
     b. Edit the `repobase` file so that the only information in the file is the following URL.
        `https://nw-node-zero/nwrpmrepo`
     c. Complete the instructions on how to run the upgrade using the `upgrade-cli-client` tool.
   - Case 2: You bootstrapped the host from local repo on the Admin server (NW Server host) and you want to use an external repo for the upgrade.
     a. Create the `/etc/netwitness/platform/repobase` file.
        `vi /etc/netwitness/platform/repobase`
     b. Edit the `repobase` file so that the only information in the file is the following URL.
        `https://<webserver-ip>/<alias-for-repo>`
     c. Complete the instructions on how to run the upgrade using the `upgrade-cli-client` tool. The instructions are in "Appendix A. Offline Method (No Connectivity to Live Services) - Command Line Interface" in the *Upgrade Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
2. Set up the external repo.
   a. Log in to the web server host.
   b. Create directory to host the NW repository (`netwitness-12.0.0.0.zip`), for example `ziprepo` under `web-root` of the web server. For example, /var/netwitness is the web-root, run the following command string.
      `mkdir -p /var/netwitness/<your-zip-file-repo>`
   c. Create the 12.0.0.0 directory under `/var/netwitness/<your-zip-file-repo>`.
      `mkdir -p /var/netwitness/<your-zip-file-repo>/12.0.0.0`
   d. Create the `OS` and `RSA` directories under `/var/netwitness/<your-zip-file-repo>/12.0.0.0`
      `mkdir -p /var/netwitness/<your-zip-file-repo>/12.0.0.0/OS`
      `mkdir -p /var/netwitness/<your-zip-file-repo>/12.0.0.0/RSA`

e. Unzip the `netwitness-12.0.0.0.zip` file into the `/var/netwitness/<your-zip-file-repo>/12.0.0.0` directory.
```
unzip netwitness-12.0.0.0.zip -d /var/netwitness/<your-zip-file-repo>/12.0.0.0
```
Unzipping `netwitness-12.0.0.0.zip` results in two zip files (`OS-12.0.0.0.zip` and `RSA-12.0.0.0.zip`)
and some other files.

f. Unzip the:
`OS-12.0.0.0.zip` into the `/var/netwitness/<your-zip-file-repo>/12.0.0.0/OS` directory.
```
unzip /var/netwitness/<your-zip-file-repo>/12.0.0.0/OS-12.0.0.0.zip -d /var/netwitness/<your-zip-file-repo>/12.0.0.0/OS
```
The external url for the repo is `http://<web server IP address>/<your-zip-file-repo>`.

g. Unzip the:
`RSA-12.0.0.0.zip` into the `/var/netwitness/<your-zip-file-repo>/12.0.0.0/RSA` directory.
```
unzip /var/netwitness/<your-zip-file-repo>/12.0.0.0/RSA-12.0.0.0.zip -d /var/netwitness/<your-zip-file-repo>/12.0.0.0/RSA
```

h. (Conditional - For Azure) Follow these steps for Azure update.

   i. `mkdir -p /var/netwitness/<your-zip-file-repo>/12.0.0.0/OS/other`

   ii. `unzip nw-azure-12.0-extras.zip -d /var/netwitness/<your-zip-file-repo>/12.0.0.0/OS/other`

   iii. `cd /var/netwitness/<your-zip-file-repo>/12.0.0.0/OS`

   iv. `createrepo`

i. Use the `http://<web server IP address>/<your-zip-file-repo>` in response to **Enter the base URL of the external update repositories** prompt from NW 12.0.0.0 Setup program (`nwsetup-tui`) prompt.

Appendix C. Troubleshooting Version Installations and Upgrades

This section describes the error messages displayed in the Hosts view when it encounters problems updating host versions and installing services on hosts in the Hosts view. If you cannot resolve an update or installation issue using the following troubleshooting solutions, contact Customer Support.

Troubleshooting instructions for the following errors that may occur during the upgrade are described in this section.

- deploy_admin Password Expired Error
- Downloading Error
- Error Deploying Version <version-number> Missing Update Packages
- Upgrade Failed Error
- External Repo Update Error
- Host Update Failed Error
- Missing Update Packages Error
- OpenSSL 1.1.x Error
- Patch Update to Non-NW Server Error
- Reboot Host After Update from Command Line Error
- Reporting Engine Restarts After Upgrade

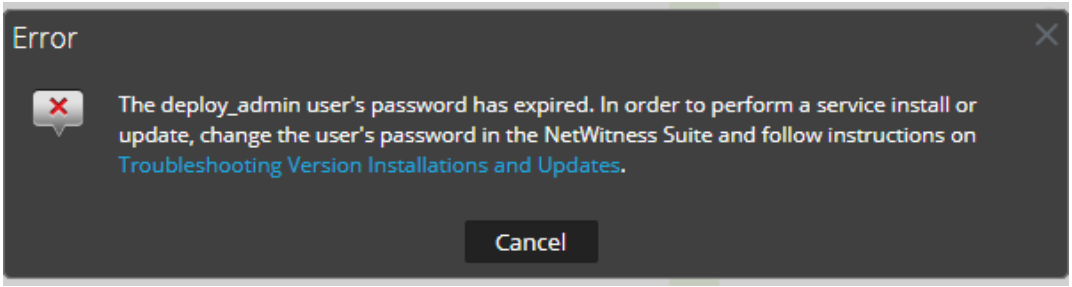
Troubleshooting instructions are also provided for errors for the following hosts and services that may occur during or after an upgrade.

- Log Collector Service
- NW Server
- Orchestration
- Reporting Engine
- Event Stream Analysis
- Legacy Windows Log Collector

Problem	Unable to boot the appliance after upgrading
Wokaround	1. Manually modify the GRUB boot line to `FIPS=0` to get it to boot. 2. From here, disable FIPS using the following command: `manage-stig-controls --disable-control-groups 3 --host-all`

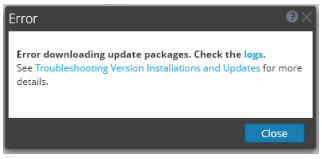
	3. Verify the line `FIPS=1` is removed from `/boot/grub2/grub.cfg`
	• If not, run the following command:
	`grub2-mkconfig -o /boot/grub2/grub.cfg`
	4. Reboot.
	5. Run the following command to enable FIPS:
	`manage-stig-controls --enable-control-groups 3 --host-all`
	6. Reboot again.

deploy_admin User Password Has Expired Error

Error Message	
Cause	The `deploy_admin` user password has expired.

Solution	Reset your `deploy_admin` password password.
	1. On the NW Server host only, run the following command.
	`nw-manage --update-deploy-admin-pw`
	`Please enter the new deploy_admin account password: <new-deploy-admin-password>`
	`Please confirm the new deploy_admin account password: <new-deploy-admin-password>`
	2. Review the output of the `nw-manage --update-deploy-admin-pw` command to verify the `deploy_admin` password was successfully updated on all hosts. If an NW host is down or fails for any reason as displayed by the output of the `nw-manage --update-deploy-admin-pw` command, run `nw-manage --sync-deploy-admin-pw --host-key <host-identifier>` to synchronize the password between the NW Server and the host that failed once the communication failure is resolved.
	3. On the host that failed installation or orchestration, run the `nwsetup-tui` command and use the new **deploy_admin** password in response to the **Deployment Password** prompt.

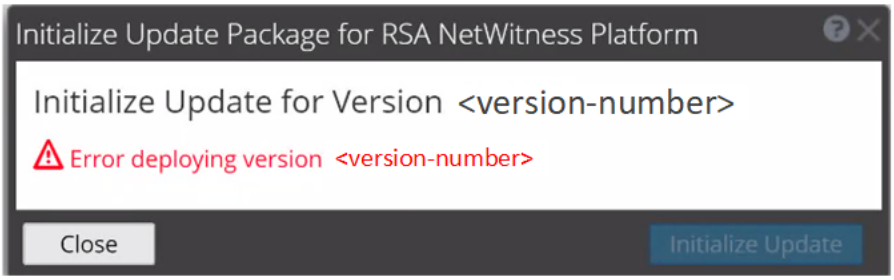
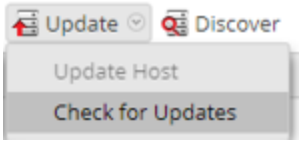
Downloading Error

Error Message	
Problem	When you select an update version and click **Update** >**Update Host**, the download starts but fails to complete.
Cause	Version download files can be large and take a long time to download. If there are communication issues during the download it will fail.
Solution	1. Try to update again. 2. If it fails again with the same error, try to update using the offline methods as described in "Offline Method from Hosts View" or "Offline Method Using Command Line Interface" in the *Upgrade Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues. 3. If you are still not able to update, contact Customer Support.

Error Message	If you are upgrading from NetWitness Platform 11.x.x.x to 11.6.x.x or later, offline UI upgrade fails with the Download error message.

Solution	1. In the Command Line Interface (CLI):
	a. SSH to NW Server.
	b. Run the following command:
	`upgrade-cli-client --upgrade --host-key <ID, IP address, hostname or display name of host> --version <version number>`
	`For example:`
	`upgrade-cli-client --upgrade --host-key <ID, IP address, hostname or display name of host> --version 11.6.0.0`
	2. After the NW Server is successfully updated, log in to the NW Server user interface and go to **(Admin) > Hosts**, where you are prompted to reboot the host.
	3. Click **Reboot Host** from the toolbar.
	You can upgrade all the other hosts directly from the user interface:
	1. Click **Begin Update** from the Update Available dialog.
	After the host is upgraded, it prompts you to reboot the host.
	2. Click **Reboot Host** from the toolbar.

Error Deploying Version <version-number> Missing Update Packages

Error Message	
Problem	Error deploying version <version-number> is displayed in the **Initialize Update Package for NetWitness Platform** dialog after you click on **Initialize Update** if the update package is corrupted.
Solution	1. Click **Close** to close the dialog. 2. Remove the version folder from staging folder. 3. Make sure that the salt-master service is running. 4. Recopy the update package zip file to the staging folder. 5. In the Hosts view toolbar, select Check for Updates again.  6. Click **Initialize Update**. 7. Click **Update** > **Update Hosts** from the toolbar. 8. Click **Begin Update** from the **Update Available** dialog. After the host is updated, it prompts you to reboot the host. 9. Click **Reboot** from the toolbar.

Upgrade Failed Error

Error Message	While updating/installing a device to version 11.2 or above, the following error can occur and be found in /var/log/netwitness/config-management/chef-solo.log: [2019-04-16T20:55:32+00:00] ERROR: Running exception handlers [2019-04-16T20:55:32+00:00] ERROR: Exception handlers complete [2019-04-16T20:55:32+00:00] FATAL: Stacktrace dumped to /var/lib/netwitness/config-management/cache/chef-stacktrace.out [2019-04-16T20:55:32+00:00] FATAL: Please provide the contents of the stacktrace.out file if you file a bug report [2019-04-16T20:55:32+00:00] ERROR: ruby_block[resolve ips] (nw-dns-client::config line 69) had an error: Resolv::ResolvError: no address for 889e5752-6ae3-4286- 33f4ccbc [2019-04-16T20:55:32+00:00] FATAL: Chef::Exceptions::ChildConvergeError: Chef run process exited unsuccessfully (exit code 1)

Cause	The reason can be because the target host is unable to communicate to the Admin Server on port 53 as it is attempting to use the dnsmasq service on the Admin Server to resolve, in this case, 889e5752-6ae3-4286-a944-c182 33f4ccbc. This is the salt minion id of the admin server. You can see this by running "cat /etc/salt/minion" on the Admin Server to compare. Example output:
	`[root@S5-NWAdmin ~]# cat /etc/salt/minion`
	`master: localhost`
	`hash_type: sha256`
	`log_level: info`
	`id: 889e5752-6ae3-4286-a944-c18233f4ccbc`

Solution	If possible, configure any firewalls between the target host and the Admin Server host to be able to communicate on port 53. If this is not possible, the workaround is to include the minion id in the /etc/host file on the component hosts and starting in the 11.4 release, modify the chef recipe not to overwrite this workaround.
Workaround	Refer to Install/Upgrade fails in NetWitness Platform because Resolv::ResolvError: no address for a particular host KB Article.

Error Message	Received an error in the error log similar to the following when trying to update to version 11.6 or later: `FATAL: Chef::Exceptions::Package: yum_package[rsa-protobufs-rt] (rsa-audit::packages line 11) had an error: Chef::Exceptions::Package: No version specified, and no candidate version available for rsa-protobufs-rt`
Cause	Custom builds/rpms installed for certain components installed on hosts, such as in the case of installing Hotfixes.

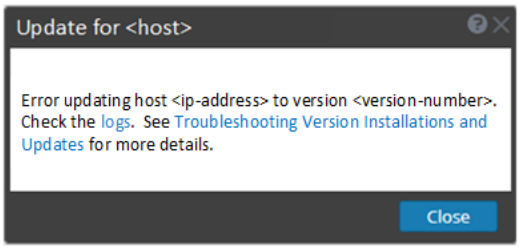
Solution	To resolve the issue, follow the below steps.
	1. SSH to Admin Server.
	2. Locate the component descriptor file by running the following command.
	`cd /etc/netwitness/component-descriptor/`
	3. Open the component descriptor file by running the following command.
	`vi nw-component-descriptor. json`
	4. Search for “packages” section for the component you have custom build/rpm. For example, below shown is the package details for “concentrator” host that has custom build/rpm.
	`“concentrator”: {`
	`“cookbook_name”: “rsa-concentrator”,`
	`“service_names”: [“rsa-nw-concentrator”],`
	`“family”: “launch”,`
	`“default_port”: xxxx, “description”: “Concentrator”,`
	`“packages”:[{ “name”: “rsa-nw-concentrator”,`
	`“version” : “11.6.0.0-2003001075220.5.cecf24b.e.17.centos”`
	`},`
	5. Delete the complete version details including (,) character in the packages section. For

	example, it should look like as shown below after you delete the version details.
	`"packages": [{`
	`"name": "rsa-nw-concentrator"`
	`},`
	Note: You must delete the version details for all the host that has custom builds/rpms in the component descriptor of the admin server.
	6. Run the upgrade process again.

External Repo Update Error

Error Message	Received an error similar to the following error when trying to update to a new version from the : `.Repository 'nw-rsa-base': Error parsing config: Error parsing "baseurl = 'https://nw-node-zero/nwrpmrepo /<version-number>/RSA'": URL must be http, ftp, file or https not ""`
Cause	There is an error the path you specified.
Solution	Make sure that: • the URL does exist on the NW Server host. • you used the correct path and remove any spaces from it.

Host Update Failed Error

Error Message	
Problem	When you select an update version and click **Update** > **Update Host**, the download process is successful, but the update process fails.
Solution	1. Try to apply the version update to the host again. Often this is all you need to do. 2. If you still cannot apply the new version update: Monitor the following logs on NW Server as it progresses (for example, run the `tail -f` command from the command line): `/var/netwitness/uax/logs/sa.log`

	/var/log/netwitness/orchestration-server/orchestration-server.log /var/log/netwitness/deployment-upgrade/chef-solo.log /var/log/netwitness/config-management/chef-solo.log /var/lib/netwitness/config-management/cache/chef-stacktrace.out The error appears in one or more of these logs. 3. If you still cannot apply the update, gather the logs from step 2 and contact Customer Support.

Missing Update Packages Error

Error Message	**Initialize Update for Version xx.x.x.x** Missing the following update package(s) Download Packages from NetWitness Link
Problem	Missing the following update package(s) is displayed in the **Initialize Update Package for NetWitness Platform** dialog when you are updating a host from the **Hosts** view offline and there are packages missing in the staging folder.
Solution	1. Click Download Packages from NetWitness Community in the **Initialize Update Package for NetWitness Platform** dialog. The NetWitness Community page that contains the update files for the selected version is displayed. 2. Select the missing packages from the staging folder. The **Initialize Update Package for NetWitness Platform** dialog is displayed telling you that it is ready to initialize the update packages.

OpenSSL 1.1.x

Error Message	The following example illustrates an ssh error that can occur when the ssh client is run from a host with OpenSSL 1.1.x installed: `$ ssh root@10.1.2.3` **`ssh_dispatch_run_fatal: Connection to 10.1.2.3 port 22: message authentication code incorrect`**

Problem	Advanced users who want to ssh to a NetWitness Platform host from a client that is using OpenSSL 1.1.x encounter this error because of incompatibility between CENTOS 7.x and OpenSSL 1.1.x. For example:
	`$ rpm -q openssl`
	`openssl-1.1.1-8.el8.x86_64`

Solution	Specify the compatible cipher list on the command line. For example:
	`$ ssh -oCiphers=aes128-ctr,aes192-ctr,aes256-ctr root@10.1.2.3`
	`I've read & consent to terms in IS user agreement.`
	`root@10.1.2.3's password:`

	`Last login: Mon Oct 21 19:03:23 2019`

Patch Update to Non-NW Server Error

Error Message	The `/var/log/netwitness/orchestration-server/orchestration-server.log` has an error similar to the following error: **`API|Failure /rsa/orchestration/task/update-config-management [counter=10 reason=IllegalArgumentException::Version '11.x.x.n' is not supported`**
Problem	After you update the NW Server host to a version, you must update all non-NW Server hosts to the same version. For example, if you update the NW Server from 11.4.0.0 to 11.6.0.0 or later, the only update path for the non-NW Server hosts is the same version (that is, 11.6.0.0). If you try to update any non-NW Server host to a different version (for example, from 11.4.0.0 to an 11.4.x.x) you will get this error.
Solution	You have two options: • Update the non-NW Server host to 11.6.0.0 or later, or • Do not update the non-NW Server host (keep it at its current version)

Reboot Host After Update from Command Line Error

Error Message	You receive a message in the User Interface to reboot the host after you update and reboot the host offline. 
Cause	You cannot use CLI to reboot the host. You must use the User Interface.
Solution	Reboot the host in the Host View in the User Interface.

Reporting Engine Restarts After Upgrade

Problem	In some cases, after you upgrade to 11.6 or later from versions of 11.x, such as 11.4, the Reporting Engine service attempts to restart continuously without success.
Cause	The database files for live charts, alert status, or report status may not be loaded successfully as the files may be corrupted.
Solution	To resolve the issue, do the following: 1. Check which database files are corrupted: Navigate to the file located at `/var/netwitness/reserver/rsa/soc/reporting-engine/logs/reporting-engine.log` and check the following blocks: • If the live charts db file is corrupted, the following logs are displayed:

Attempt Failed!!! Clearing pending acquires. While trying to acquire a needed new resource, we failed to succeed more than the maximum number of allowed acquisition attempts (30). Last acquisition attempt exception:

org.h2.jdbc.JdbcSQLException: File corrupted while reading record: null. Possible solution: use the recovery tool [90030-196]

at org.h2.message.DbException.getJdbcSQLException(DbException.java:345)

at org.h2.message.DbException.get(DbException.java:168)

org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'chartSummaryDAOImpl': Invocation of init method failed; nested exception is com.rsa.soc.re.exception.ReportingException: java.sql.SQLException: Connections could not be acquired from the underlying database!

- If the alert status db file is corrupted, the following logs are displayed:

Attempt Failed!!! Clearing pending acquires. While trying to acquire a needed new resource, we failed to succeed more than the maximum number of allowed acquisition attempts (30). Last acquisition attempt exception:

org.h2.jdbc.JdbcSQLException: File corrupted while reading record: null. Possible solution: use the recovery tool [90030-196]

at org.h2.message.DbException.getJdbcSQLException(DbException.java:345)

at org.h2.message.DbException.get(DbException.java:168)

org.springframework.beans.factory.UnsatisfiedDependencyException: Error creating bean with name 'alertStatusHandler': Unsatisfied dependency expressed through field 'alertExecutionStatsDAO'; nested exception is org.springframework.beans.factory.UnsatisfiedDependencyException: Error creating bean with name 'alertExecutionStatsDAOImpl': Unsatisfied dependency expressed through field 'persistedAlertExecutionStatsDAO'; nested exception is org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'persistedAlertExecutionStatsDAOImpl'

- If the report status db file is corrupted, the following logs are displayed:

```
org.h2.jdbc.JdbcSQLException: File corrupted while reading
record: null. Possible solution: use the recovery tool
[90030-196]
```

2. To resolve the live charts database file corruption, perform the following steps:
   a. Stop the Reporting Engine service.
   b. Move the `livechart.mv.db` file from `/var/netwitness/reserver/rsa/soc/reporting-engine/livecharts` folder to a temporary location.
   c. Restart the Reporting Engine service.

   Note: Some live charts data may be lost on performing the above steps.

   To resolve the alert status or report status database file corruption, perform the following steps:

   a. Stop the Reporting Engine service.
   b. Replace the corrupted db file with the latest `alertstatusmanager.mv.db` or `reportstatusmanager.mv.db` file from `/var/netwitness/reserver/rsa/soc/reporting-engine/archives` folder.
   c. Restart the Reporting Engine service.

   For more information, see the Knowledge Base article Reporting Engine restarts After upgrade to NetWitness Platform 11.4.

Problem	After you upgrade to version 11.6 or later, the Reporting Engine service does not restart.
Cause	The Reporting Engine service may not start due to any of the following reasons. - workspace.xml not updated. - Time is not converted properly in livechart h2 database. - JCR (Jackrabbit repository) is corrupted with primary key violation.

Solution	To resolve the issue, run the Reporting Engine Migration Recovery tool (`rsa-nw-re-migration-recovery.sh`) on the Admin Server where the Reporting Engine service is installed.
	Note: You can find the Reporting Engine Migration Recovery tool in the below location.
	`/opt/rsa/soc/reporting-engine-<version number>-<Tag>/nwtools`
	For example:
	`/opt/rsa/soc/reporting-engine-11.6.0.0-<Tag>/nwtools`
	1. SSH to Admin Server.
	2. Untar the RE (Reporting Engine) tool, run the following command.
	`tar -xvf rsa-nw-re-recovery-tool-bundle.tar`
	3. (Optional) If you want to untar the RE tool file in some other directory, you can create a directory and untar the RE tool. Run the following commands.
	`mkdir <NAME OF THE DIRECTORY>`
	`tar -xvf rsa-nw-re-recovery-tool-bundle.tar --directory <PATH OF THE DIRECTORY>`
	4. Run the script, run the following command.
	`./<PATH OF THE DIRECTORY>/rsa-nw-re-recovery-tool.sh`
	For more information, see the Knowledge Base article **Reporting Engine Migration Recovery Tool**.

Log Collector Service (`nwlogcollector`)

Log Collector installation logs posted to `/var/log/install/nwlogcollector_install.log` on the host running the `nwlogcollector` service.

Error Message	`<timestamp>.NwLogCollector_PostInstall: Lockbox Status : Failed to open lockbox: The lockbox stable value threshold was not met because the system fingerprint has changed. To reset the system fingerprint, open the lockbox using the passphrase.`
Cause	The Log Collector Lockbox failed to open after the update.
Solution	Log in to NetWitness and reset the system fingerprint by resetting the stable system value password for the Lockbox as described in the "Reset the Stable System Value" topic under "Configure Lockbox Security Settings" topic in the *Log Collection Configuration Guide*.

Error	`<timestamp> NwLogCollector_PostInstall: Lockbox Status : Not Found`

Message	
Cause	The Log Collector Lockbox is not configured after the update.
Solution	If you use a Log Collector Lockbox, log in to NetWitness and configure the Lockbox as described in the "Configure Lockbox Security Settings" topic in the *Log Collection Configuration Guide*.

Error Message	`<timestamp>: NwLogCollector_PostInstall: Lockbox Status : Lockbox maintenance required: The lockbox stable value threshold requires resetting. To reset the system fingerprint, select Reset Stable System Value on the settings page of the Log Collector.`
Cause	You need to reset the stable value threshold field for the Log Collector Lockbox.
Solution	Log in to NetWitness and reset the stable system value password for the Lockbox as described in "Reset the Stable System Value" topic under "Configure Lockbox Security Settings" topic in the *Log Collection Configuration Guide*.

Error Message	Decoder tries to start capture events but fails. `Aug 27 01:44:41 nwdecoder NwDecoder[8052]: [Decoder] [failure] The PF_RING driver is not installed.`

Solution	To resolve the issue, do the following steps,
	1. SSH to the Decoder host.
	2. Run the following commands.
	`yum reinstall pfring*`
	`systemctl restart nwdecoder`

NW Server

These logs are posted to `/var/netwitness/uax/logs/sa.log` on the NW Server Host.

Problem	After upgrade, you notice that Audit logs are not getting forwarded to the configured Global Audit Setup; or, The following message seen in the `sa.log`. `Syslog Configuration migration failed. Restart jetty service to fix this issue`
Cause	NW Server Global Audit setup migration failed to migrate from 11.4.x.x or 11.5.x.x. to 11.6.0.0 or later.
Solution	1. SSH to the NW Server. 2. Submit the following command. `orchestration-cli-client --update-admin-node`

Orchestration

The orchestration server logs are posted to `/var/log/netwitness/orchestration-server/orchestration-server.log` on the NW Server Host.

Problem	1. Tried to upgrade a non-NW Server host and it failed. 2. Retried the upgrade for this host and it failed again. You will see the following message in the `orchestration-server.log`. `"'file' _virtual_ returned False: cannot import name HASHES""`
Cause	Salt minion may have been upgraded and never restarted on failed non-NW Server host

Solution	1. SSH to the non-NW Server host that failed to upgrade.
	2. Submit the following commands.
	`systemctl unmask salt-minion`
	`systemctl restart salt-minion`
	3. Retry the upgrade of the non-NW Server host.

Reporting Engine Service

Reporting Engine Update logs are posted to to`/var/log/re_install.log` file on the host running the Reporting Engine service.

Error Message	`<timestamp> : Available free space in /var/netwitness/re-server/rsa/soc/reporting-engine [ ><existing-GB ] is less than the required space [ <required-GB> ]`
Cause	Update of the Reporting Engine failed because you do not have enough disk space.
Solution	Free up the disk space to accommodate the required space shown in the log message. See the "Add Additional Space for Large Reports" topic in the *Reporting Engine Configuration Guide* for instructions on how to free up disk space.

Event Stream Analysis

Problem	After upgrading to version 11.6 or later, the ESA correlation server does not aggregate events from the configured data sources.
Error Message	`Invalid username or password at com.rsa.netwitness.streams.base.RecordSourceSubscription.run (RecordSourceSubscription.java:173)`
Solution	To resolve the issue, do the following steps. In the NetWitness user interface,

	1. Go to **(Configure) > ESA Rules**. **ESA Rules** panel is displayed with **Rules** tab open. 2. In the Rules tab options panel, under Deployments, select a deployment. 3. In the **Data Sources** section, select the data source and click in the toolbar. 4. In the **Edit Service** dialog, type the password for that data source. 5. Click the **Test Connection** button to make sure that it can communicate with the ESA service and then click **OK**. **Note:** Do the above procedure for all the configured data sources. 6. After you finish making changes to the deployment, click **Deploy Now** to redeploy the ESA rule deployment.

Legacy Windows Log Collector

Problem	• Legacy Windows Log Collector appears as inactive post upgrade of SA to 12.0.0.0 version and Legacy Windows Log Collector to 11.6.x or 11.7.x versions. • Legacy Windows Log Collector appears as inactive when the stack is upgraded to 12.0.0.0.
Cause	Certificate update in the SA node.
Solution	Refer Legacy Windows Log Collector section in the Post Update Tasks.

ESA Troubleshooting Information

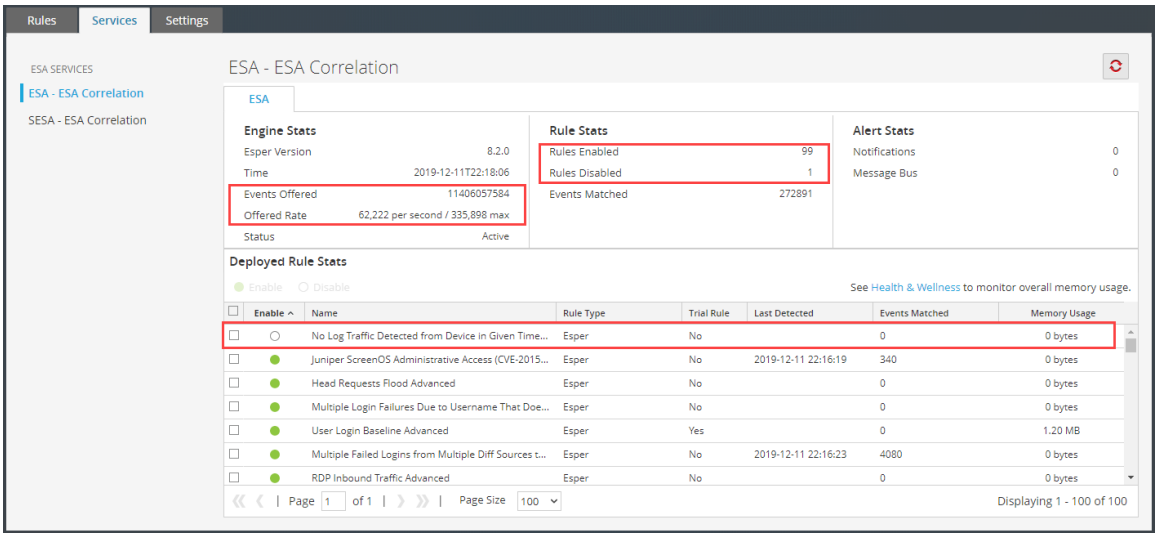
ESA Rules are Not Creating Alerts

If you are not seeing any alerts, check the status of the ESA rule deployments.

1. Go to **(Configure) > ESA Rules > Services** tab.
   The Services view is displayed, which shows the status of your ESA services and deployments.
2. In the options panel on the left, select an ESA service.
3. For each service listed, look at the deployment tabs in the panel on the right. Each tab represents a separate ESA rule deployment.
4. For each ESA rule deployment:
   a. In the **Engine Stats** section, look at the **Events Offered** and the **Offered Rate**. They confirm that the data is being aggregated and analyzed properly. If you see 0 for Events Offered, nothing is

coming in for the deploymIent.

b. In the **Rule Stats** section, look at the **Rules Enabled** and **Rules Disabled**. If there are any disabled rules, look in the **Deployed Rule Stats** section below to view the details of the disabled rules. Disabled rules show a white circle. Enabled rules show a green circle.



5. If you notice any disabled rules that should be enabled:

   a. Go to **(Configure) > ESA Rules > Rules** tab and redeploy the ESA rule deployments that contain disabled rules.

   b. Go back to the **Services** tab and check to see if the rules are still disabled. If the rules are still disabled, check the ESA Correlation service log files, which are located at `/var/log/netwitness/correlation-server/correlation-server.log`.

Note: To avoid unnecessary processing overhead, the Ignore Case option has been removed from the ESA Rule Builder - Build a Statement dialog for meta keys that do not contain text data values. During the upgrade to 11.4 or later, NetWitness Platform does not modify existing rules for the Ignore Case option. If an existing Rule Builder rule has the Ignore Case option selected for a meta key that no longer has the option available, an error occurs if you try to edit the statement and try to save it again without clearing the checkbox.

Endpoint, UEBA, and Live Content Rules are Not Working

To support Endpoint and UEBA content as well as changes to ESA rules from Live, a data change from single-value (string) to multi-value (string array) is required for several meta keys within the ESA Correlation service. In NetWitness Platform 11.4 or later, ESA automatically adjusts the operator in the rule statement when there is a change from string to string array, but you still may need to make manual adjustments to adjust for the string array changes.

To change the string type meta keys to string array type meta keys manually in 11.4 or later, see "Configure Meta Keys as Arrays in ESA Correlation Rule Values" in the *ESA Configuration Guide*.

To use the latest Endpoint, UEBA, and Live content rules, the following default **multi-valued** meta keys are required on the ESA Correlation service in NetWitness Platform version 11.4 or later:

```
action , alert , alert.id , alias.host , alias.ip , alias.ipv6 , analysis.file
, analysis.service , analysis.session , boc , browserprint , cert.thumbprint ,
checksum , checksum.all , checksum.dst , checksum.src , client.all , content ,
context , context.all , context.dst , context.src , dir.path , dir.path.dst ,
dir.path.src , directory , directory.all , directory.dst , directory.src ,
email , email.dst , email.src , eoc , feed.category , feed.desc , feed.name ,
file.cat , file.cat.dst , file.cat.src , filename.dst , filename.src , filter
, function , host.all , host.dst , host.orig , host.src , host.state ,
inv.category , inv.context , ioc , ip.orig , ipv6.orig , netname , OS , param
, param.dst , param.src , registry.key , registry.value , risk , risk.info ,
risk.suspicious , risk.warning , threat.category , threat.desc , threat.source
, user.agent , username
```

The following default **single-valued** meta keys are also required on the ESA Correlation service in NetWitness Platform 11.4 or later:

```
accesses , context.target , file.attributes , logon.type.desc , packets
```

To update your meta keys, see "Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules" in the *ESA Configuration Guide*.

If you used any meta keys in the ESA rule notification templates from the Required String Array or String Meta Keys list, update the templates with the meta key changes. See "Configure Global Notification Templates" in the *System Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Note: Advanced EPL rules may get disabled and are not automatically updated so they must be fixed manually.

For additional troubleshooting information, see “Troubleshoot ESA” in the *Alerting with ESA Correlation Rules User Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Example ESA Correlation Server Warning Message for Missing Meta Keys

If you see a warning message in the ESA Correlation server error logs that means there is a difference between the `default-multi-valued` parameter and `multi-valued parameter` meta key values, the new Endpoint, UEBA, and Live content rules will not work. Completing the "Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules" procedure in the *ESA Configuration Guide* should fix the issue.

Multi-Valued Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[alert, alert_id,
browserprint, cert_thumbprint, checksum, checksum_all, checksum_dst, checksum_
src, client_all, content, context, context_all, context_dst, context_src, dir_
path, dir_path_dst, dir_path_src, directory, directory_all, directory_dst,
directory_src, email_dst, email_src, feed_category, feed_desc, feed_name,
file_cat, file_cat_dst, file_cat_src, filename_dst, filename_src, filter,
function, host_all, host_dst, host_orig, host_src, host_state, ip_orig, ipv6_
orig, OS, param, param_dst, param_src, registry_key, registry_value, risk,
risk_info, risk_suspicious, risk_warning, threat_category, threat_desc,
threat_source, user_agent] are still MISSING from multi-valued
```

Single Value Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[accesses, context_target,
file_attributes, logon_type_desc, packets] are still MISSING from single-
valued
```

NetWitness® Platform XDR

Version 12.0.0.0

API Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

October, 2022

Table of Contents

Overview . . . 2
- Current Version . . . 2
- Schema . . . 2
- HTTP Usage . . . 2
- Case Sensitive . . . 3
- Error Response . . . 3
- Pagination . . . 4

Authentication and Authorization . . . 5
- Obtaining a Token . . . 5
  - Using a Username and Password . . . 5
  - Using a Refresh Token . . . 6
- Authorization . . . 8

Incident APIs . . . 9
- Attributes . . . 9
  - Incident Priority . . . 10
  - Incident Status . . . 11
  - Milestone . . . 11
- Requests . . . 11
  - Get a Single Incident . . . 11
  - Get Incidents by Date Range . . . 13
  - Update an Incident . . . 15
  - Remove an Incident . . . 17
  - Add a Journal Entry . . . 18
  - Get an Incident's Alerts . . . 19
  - Persisting events in an alert . . . 23
  - Suspending persist of events in an alert . . . 24
  - Persisting events in an incident . . . 24
  - Suspending persist of events in an incident . . . 25
  - Get Incident history . . . 26
  - Get Incident Stats . . . 27
  - Get Incident User Stats . . . 29

Services Information . . . 30
- Get Service IDs of all Services . . . 30
  - Sample Request . . . 30
  - Sample Response . . . 30
- Get List of Service IDs by Service Name . . . 31
  - Sample Request . . . 31
  - Sample Response . . . 31

Endpoint APIs . . . 33
Get Hosts . . . 33
Sample Request . . . 34
Sample Response . . . 34
Get Hosts with Filter . . . 35
Sample Request . . . 35
HTTP request . . . 36
Sample Response . . . 36
Get List of Snapshots for Host . . . 38
Sample Request . . . 38
Sample Response . . . 38
Get Snapshot Details for the Host . . . 38
Response Fields . . . 38
Sample Request . . . 46
Sample Response . . . 46
Get Files . . . 48
Sample Request . . . 51
Sample Response . . . 51
Request Scan . . . 52
Path Parameters . . . 52
Request Parameters . . . 53
Sample Request . . . 53
Sample Response . . . 53
Request Scan with given CPU Usage . . . 53
Path Parameters . . . 53
Request Parameters . . . 53
Sample Request . . . 54
Sample Response . . . 54
Request Stop Scan . . . 54
Path Parameters . . . 54
Request Parameters . . . 54
Sample Request . . . 54
Sample Response . . . 54
Get Alerts for a Host . . . 55
Sample Request . . . 55
Sample Response . . . 55
Get Alerts for a File . . . 56
Sample Request . . . 57
Sample Response . . . 57
Request File Download to Server . . . 58
Path Parameters . . . 58

Request Parameters . . . 59
Sample Request . . . 59
HTTP request . . . 59
Sample Response . . . 59
Request Multiple File Downloads to Server . . . 59
Path Parameters . . . 60
Request Parameters . . . 60
Sample Request . . . 60
HTTP request . . . 60
Sample Response . . . 60
Request MFT Download . . . 61
Path Parameters . . . 61
Request Parameters . . . 61
Sample Request . . . 61
Sample Response . . . 61
Request System Dump Download . . . 62
Path Parameters . . . 62
Request Parameters . . . 62
Sample Request . . . 62
Sample Response . . . 62
Request Process Dump Download . . . 63
Path Parameters . . . 63
Sample Request . . . 63
Sample Response . . . 63
Request Network Isolation . . . 63
Path Parameters . . . 64
Request Parameters . . . 64
Sample Request . . . 64
HTTP request . . . 64
Sample Response . . . 64
Update Network Isolation Exclusion List . . . 65
Path Parameters . . . 65
Request Parameters . . . 65
Sample Request . . . 65
HTTP request . . . 65
Sample Response . . . 65
Release from Network Isolation . . . 66
Path Parameters . . . 66
Request Parameters . . . 66
Sample Request . . . 66
HTTP request . . . 66

Sample Response 67

Overview

The NetWitness Platform API can be accessed using the same host and port as the NetWitness user interface.

Current Version

By default, all requests to the REST API will automatically use the latest version of the API available. To provide API stability, clients can specify the API version to use by adding the `NetWitness-Version` HTTP header:

```
NetWitness-Version: 1.0
```

Schema

All data is sent and received as JSON. Any resources containing fields without values will have those fields included with `null` as the value instead of being omitted.

Any fields containing timestamps or dates will be in ISO 8601 format:

```
YYYY-MM-DDTHH:MM:SS.SSSZ
```

HTTP Usage

The RSA NetWitness API tries to adhere as closely as possible to standard HTTP and REST conventions in its use of HTTP verbs and status codes.

HTTP Verbs

Verb	Usage
`GET`	Used to retrieve a resource.
`POST`	Used to create a new resource.
`PATCH`	Used to update an existing resource, including partial updates. Only fields that are modified should be included in the request.
`PUT`	Used to replace an existing resource.
`DELETE`	Used to delete an existing resource.

HTTP Status Codes

Status code	Usage
`200 OK`	The request completed successfully.

`201 Created`	A new resource has been created successfully. The resource's URI is available from the response's `Location` header.
`204 No Content`	An update to an existing resource has been applied successfully.
`400 Bad Request`	The request was malformed. The response body will include an error providing further information. See Error Response.
`401 Unauthorized`	Similar to `403 Forbidden`, but specifically for use when authentication is required and has failed or has not yet been provided. See Authentication and Authorization.
`403 Forbidden`	The request was valid, but the server is refusing the action. The user might not have the necessary permissions for a resource.
`404 Not Found`	The requested resource does not exist.
`500 Internal Server Error`	An unexpected error has occurred. The response body will include a message providing further information.

Case Sensitive

All URLs, request parameters and JSON fields are case sensitive.

Error Response

A common JSON structure is always returned for errors:

Path	Type	Description
`status`	`Number`	The HTTP status code returned.
`timestamp`	`String`	The timestamp of the request.
`errors[]`	`Array`	An array of errors for the given request.
`errors[].message`	`String`	A user-friendly error message explaining what went wrong.
`errors[].field`	`String`	The field or parameter containing the error.

```
{
  "status" : 400,
  "timestamp" : "2022-12-12T09:10:36.6424477Z",
  "errors" : [ {
    "message" : "Value must be less than or equal to \"10\"",
    "field" : "start"
  }, {
    "message" : "Invalid range"
  } ]
}
```

Pagination

A common JSON structure is always used for paginated results:

Path	Type	Description
`items`	`Array`	An array containing the requested resources.
`pageNumber`	`Number`	The requested page number.
`pageSize`	`Number`	The requested number of items to return in a single page.
`totalPages`	`Number`	The total number of pages available.
`totalItems`	`Number`	The total number of items available.
`hasNext`	`Boolean`	Indicates if there is a page containing results after this page.
`hasPrevious`	`Boolean`	Indicates if there is a page containing results before this page.

```
{
  "items" : [ ],
  "pageNumber" : 0,
  "pageSize" : 10,
  "totalPages" : 3,
  "totalItems" : 25,
  "hasNext" : true,
  "hasPrevious" : false
}
```

Authentication and Authorization

All requests must include the `NetWitness-Token` HTTP header containing a valid JSON Web Token (JWT):

```
NetWitness-Token:
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9.eyJleHAiOjE1MTEyNDczODYyNjMsImlzcyI6InNlY3VyaXR5L
XNlcnZlci0zODA1NTA0OS0xZWMyLTQ0MDAtOTUwYS0zZTVkMmJiYTljMjIiLCJpYXQiOjE1MTEyMTEzODYyNjM
sImF1dGhvcml0aWVzIjpbIkFkbWluaXN0cmF0b3JzIl0sInVzZXJfbmFtZSI6ImFkbWluIn0.StBjg9ruIX4Fr
yfCX8qvrSBGZHF8DN3qHZM0Ei9-
thFndm1q_DLP_cnh8Fpm43fdKcs1ErcVRTqhaYvVULYmsF9ShUaSThpLts6zbJVEKlq3ldUGWWCY9bfVGRH3n5
KmWzITPi7xZ-
Rf_Kp2Sj8ecVAip3qDwha7TxYrReXefCnUj0UxgaaXjeZIFjwxFmK6NPZ7TAK9OcvcVhozaR8V92g1kUVP8_54
x7iZ2jL4JvDPaScWBjBTvVEffHNbX9_iLNoRmKqvDELSla6E_trkSREogCt6pZh709Qh70uoC3BsKwNQKbHNEO
U1tRPFaUFfRH7bCdp8v3Aeh3PTaKEuQA
```

The JSON Web Token is defined in RFC-7519. Tokens can be obtained using the methods outlined below.

In the remainder of this document, the token will be truncated to just `eyJ...AT` for brevity.

Obtaining a Token

A JSON Web Token can be obtained using the methods below.

Using a Username and Password

Users can retrieve an access token using their username and password credentials. Since the API gateway is secured using TLS, all credentials will be encrypted in transit.

```
POST /rest/api/auth/userpass
```

Request Parameters

Parameter	Description
`username`	The username of the account to authenticate.
`password`	The password of the account.

Response Fields

Path	Type	Description
`id`	`String`	The account identifier.
`roles`	`Array`	The roles assigned to the user.

Path	Type	Description
accessToken	String	A digitally signed access token that is acceptable as proof of authentication at any Launch service that trusts the public key of this service. The string holds a JSON web-token. See RFC-7519.
refreshToken	String	A digitally signed refresh token that can be used to refresh an expired access token. Refresh tokens have longer expiry periods and can be used by services to re-authenticate users without (storing and) presenting credentials.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/auth/userpass' -i -X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1' \
    -d 'username=ian&password=changeMe'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:14:33 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 106
```

```
{
  "id" : "ian",
  "roles" : [ "Analyst" ],
  "accessToken" : "eyJ...AT",
  "refreshToken" : "eyJ...AT"
}
```

Using a Refresh Token

Users can also retrieve an access token using a refresh token.

```
POST /rest/api/auth/token
```

Request Parameters

Parameter	Description
token	A refresh token.

Response Fields

Path	Type	Description
id	String	The account identifier.
roles	Array	The roles assigned to the user.
accessToken	String	A digitally signed access token that is acceptable as proof of authentication at any Launch service that trusts the public key of this service. The string holds a JSON web-token. See RFC-7519.
refreshToken	String	A digitally signed refresh token that can be used to refresh an expired access token. Refresh tokens have longer expiry periods and can be used by services to re-authenticate users without (storing and) presenting credentials.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/auth/token' -i -X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1' \
    -d
'token=eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9.eyJleHAiOjE2NzM0Mjg0NzMyNzEsImlzcyI6InNlY3
VyaXR5LXNlcnZlci0xZDAzZDVkNi0zYjg3LTQzYjUtOGI4Yi1mZTE0OTRlODE4ZjIiLCJpYXQiOjE2NzA4MzY0
NzMyNzEsInJlZnJlc2giOnRydWUsInVzZXJfbmFtZSI6ImlhbiJ9.oN50Pmw09pw5HTqVZN3FkAkYm6urCzCCP
eHIRDFN2OL1WHUklb_YzkTls9sf_8jleZ8jJblR6avffEBBBthE8livw0Nt-
7iWatG7IHm_oijB2EookKXRbxv9461VssfxWdbY7zs-
s5ltI96atQYbQPFOmPoGogpW_bzZJgFcEAbRBydaVYktKR0t01xIU0PKabBo_GBwutYt0054KljizOhiL_2qAB
m7WrjjluL0nscLCIZyN64DdYtWaE3ONhRyvrjKsnlkCec5ucFXOjoRK03cocqFJuCyoM-
TKSbBpjUDJh2c0jP1wP17dOR-KQYrvbNSaqRef2_KyBbT6Bm0cg'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:14:33 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 106
```

```
{
  "id" : "ian",
  "roles" : [ "Analyst" ],
  "accessToken" : "eyJ...AT",
  "refreshToken" : "eyJ...AT"
}
```

Authorization

In order to make requests through the NetWitness Platform API, users must belong to roles that have the `integration-server.api.access` permission, as well as any underlying permissions required to fulfill the request.

Incident APIs

An Incident is a logically grouped set of alerts created automatically by the Incident Aggregation Engine and grouped by a specific criteria. An Incident, available in the Respond Interface, allows an Analyst to triage, investigate, and remediate these groups of alerts. Incidents can be moved between users, notated, and explored via the nodal graph. Incidents allow users to ensure they understand the full scope of an attack or event in their NW system and then take action.

Attributes

The incident resource is comprised of the following attributes:

Path	Type	Description
`id`	`String`	The unique identifier of the incident.
`title`	`String`	The title of the incident.
`summary`	`String`	The summary of the incident.
`priority`	`String`	The incident priority. See the valid values.
`riskScore`	`Number`	The incident risk score is calculated based on the associated alert's risk score. Risk score ranges from 0 (no risk) to 100 (highest risk).
`status`	`String`	The current status. See the valid values.
`alertCount`	`Number`	The number of alerts associated with an incident.
`averageAlertRiskScore`	`Number`	The average risk score of the alerts associated with the incident. Risk score ranges from 0 (no risk) to 100 (highest risk).
`sealed`	`Boolean`	Indicates if additional alerts can be associated with an incident. A `sealed` incident cannot be associated with additional alerts.
`totalRemediationTaskCount`	`Number`	The number of total remediation tasks for an incident.
`openRemediationTaskCount`	`Number`	The number of open remediation tasks for an incident.
`created`	`String`	The timestamp of when the incident is created.
`lastUpdated`	`String`	The timestamp of when the incident was last updated.
`lastUpdatedBy`	`String`	The NetWitness user identifier of the user who last updated the incident.
`assignee`	`String`	The NetWitness user identifier of the user currently working on the incident.
`sources`	`Array`	Unique set of sources for all of the alerts in an incident.

Path	Type	Description
`ruleId`	`String`	The unique identifier of the rule that created the incident.
`firstAlertTime`	`String`	The timestamp of the earliest occurring Alert in this incident.
`categories`	`Array`	The list of categories this incident is categorized under.
`categories[].id`	`String`	The unique category identifier.
`categories[].parent`	`String`	The parent name of the category.
`categories[].name`	`String`	The friendly name of the category.
`journalEntries`	`Array`	Set of notes about the incident investigation, also known as the JournalEntry.
`journalEntries[].id`	`String`	The unique journal entry identifier.
`journalEntries[].author`	`String`	The author of this entry.
`journalEntries[].notes`	`String`	Notes and observations about the incident.
`journalEntries[].created`	`String`	The timestamp of the journal entry created date.
`journalEntries[].lastUpdated`	`String`	The timestamp of the journal entry last updated date.
`journalEntries[].milestone`	`String`	Incident milestone classifier. See the valid values.
`createdBy`	`String`	The NetWitness user id or name of the rule that created the incident.
`deletedAlertCount`	`Number`	The number of alerts that are deleted from the incident.
`eventCount`	`Number`	The number of events associated with incident.
`alertMeta`	`String`	An object containing unique set of meta values, by type, across all alerts associated with this incident.
`alertMeta.SourceIp`	`Array`	Unique source IP addresses.
`alertMeta.DestinationIp`	`Array`	Unique destination IP addresses.

Incident Priority

The `priority` field can contain these values:

Value	Description
`Low`	Low Priority
`Medium`	Medium Priority
`High`	High Priority
`Critical`	Critical

Incident Status

The `status` field can contain these values:

Value	Description
`New`	New incident.
`Assigned`	Incident is assigned to a user.
`InProgress`	Incident response is in progress.
`RemediationRequested`	Remediation tasks have been requested.
`RemediationComplete`	Remediation tasks are complete.
`Closed`	Incident is closed.
`ClosedFalsePositive`	Incident is closed as it was created due to false positive.

Milestone

Each journal entry can contain a `milestone` consisting of these values:

Value	Description
`Reconnaissance`	Intruder is in the initial phase of the attack where targets and vulnerabilities are identified.
`Delivery`	Intruder transmitted malware to the target.
`Exploitation`	Malware code triggers, which takes action on target network to exploit vulnerability.
`Installation`	Malware weapon installs access point usable by intruder.
`CommandAndControl`	Malware enables intruder to have persistent access to target network.
`ActionOnObjective`	Intruder takes action to achieve their goals, such as data exfiltration, data destruction, or encryption for ransom.
`Containment`	Incident is contained.
`Eradication`	Necessary actions taken to eliminate components of incident and restore the system status.
`Closure`	Incident is addressed.

Requests

Get a Single Incident

A single incident can be retrieved using an incident's unique identifier.

```
GET /rest/api/incidents/{id}
```

Path Parameters

Parameter	Description
`id`	The unique identifier of the incident.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/INC-100' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:50 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 1329
```

```
{
  "id" : "INC-100",
  "title" : "Suspected C&C with suspicious-domain.com",
  "summary" : "Security Analytics detected communications with suspicious-domain.com
that may be command and control malware.",
  "priority" : "Critical",
  "riskScore" : 100,
  "status" : "InProgress",
  "alertCount" : 1,
  "averageAlertRiskScore" : 100,
  "sealed" : true,
  "totalRemediationTaskCount" : 4,
  "openRemediationTaskCount" : 5,
  "created" : "2018-01-01T04:49:27.870Z",
  "lastUpdated" : "2022-12-12T09:15:49.787Z",
  "lastUpdatedBy" : "duke",
  "assignee" : "ian",
  "sources" : [ "Malware Analysis" ],
  "ruleId" : "55e49a79e4b01a1d2be502bc",
  "firstAlertTime" : "2017-08-04T16:49:22Z",
  "categories" : [ {
    "id" : "55e49a79e4b01a1d2be5022e",
    "parent" : "Malware",
    "name" : "Password dumper"
  }, {
    "id" : "55e49a79e4b01a1d2be50228",
    "parent" : "Hacking",
    "name" : "Path traversal"
  } ],
  "journalEntries" : [ {
    "id" : "20",
    "author" : "admin",
    "notes" : "Updated status",
    "created" : "2017-11-15T20:20:54.785Z",
    "lastUpdated" : "2017-11-15T20:20:54.785Z",
    "milestone" : "Containment"
  } ],
  "createdBy" : "norm",
  "deletedAlertCount" : 100,
  "eventCount" : 0,
  "alertMeta" : {
    "SourceIp" : [ "10.11.12.345" ],
    "DestinationIp" : [ "11.11.11.111", "11.22.33.444" ]
  }
}
```


Get Incidents by Date Range

Incidents can be retrieved by the date and time they were created.

```
GET /rest/api/incidents
```

The requested date range can be unbounded, by only supplying either the `since` or `until` parameter, or bounded, by providing both parameters.

Request Parameters

Parameter	Description
`pageNumber`	The requested page number.
`pageSize`	The maximum number of items to return in a single page.
`since`	A timestamp in ISO 8601 format (e.g., `1018-01-01T14:00:00.000Z`). Retrieve incidents created on and after this timestamp.
`until`	A timestamp in ISO 8601 format (e.g., `1018-01-01T14:00:00.000Z`). Retrieve incidents created on and before this timestamp.

All results will be returned using the paginated response payload sorted by the `created` date, in descending order.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents?since=2018-01-
01T04%3A00%3A00.000Z&until=2018-01-01T05%3A00%3A00.000Z&pageSize=100&pageNumber=0' -i
-X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:47 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 1560
```

```
{
  "items" : [ {
    "id" : "INC-100",
    "title" : "Suspected C&C with suspicious-domain.com",
    "summary" : "Security Analytics detected communications with suspicious-domain.com
that may be command and control malware.",
    "priority" : "Critical",
    "riskScore" : 100,
    "status" : "Assigned",
```

```
    "alertCount" : 1,
    "averageAlertRiskScore" : 100,
    "sealed" : true,
    "totalRemediationTaskCount" : 4,
    "openRemediationTaskCount" : 5,
    "created" : "2018-01-01T04:49:27.870Z",
    "lastUpdated" : "2017-08-04T16:49:27.870Z",
    "lastUpdatedBy" : "norm",
    "assignee" : "tony",
    "sources" : [ "Malware Analysis" ],
    "ruleId" : "55e49a79e4b01a1d2be502bc",
    "firstAlertTime" : "2017-08-04T16:49:22Z",
    "categories" : [ {
      "id" : "55e49a79e4b01a1d2be5022e",
      "parent" : "Malware",
      "name" : "Password dumper"
    }, {
      "id" : "55e49a79e4b01a1d2be50228",
      "parent" : "Hacking",
      "name" : "Path traversal"
    } ],
    "journalEntries" : [ {
      "id" : "20",
      "author" : "admin",
      "notes" : "Updated status",
      "created" : "2017-11-15T20:20:54.785Z",
      "lastUpdated" : "2017-11-15T20:20:54.785Z",
      "milestone" : "Containment"
    } ],
    "createdBy" : "norm",
    "deletedAlertCount" : 100,
    "eventCount" : 0,
    "alertMeta" : {
      "SourceIp" : [ "10.11.12.345" ],
      "DestinationIp" : [ "11.11.11.111", "11.22.33.444" ]
    }
  } ],
  "pageNumber" : 0,
  "pageSize" : 100,
  "totalPages" : 1,
  "totalItems" : 1,
  "hasNext" : false,
  "hasPrevious" : false
}
```


Update an Incident

Currently an incident's `status` and `assignee` can be modified using the incidents endpoint.

```
PATCH /rest/api/incidents/{id}
```

The `assignee` field must include the unique identifier for a valid NetWitness user. The list of users can be found in the security section of the administration user interface.

Request Fields

Path	Type	Description
`status`	`String`	The current status. See the valid values.
`assignee`	`String`	The NetWitness user identifier of the user currently working on the incident.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/INC-100' -i -X PATCH \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8' \
    -d '{"status":"InProgress"}'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:49 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 1330
```

```
{
  "id" : "INC-100",
  "title" : "Suspected C&C with suspicious-domain.com",
  "summary" : "Security Analytics detected communications with suspicious-domain.com
that may be command and control malware.",
  "priority" : "Critical",
  "riskScore" : 100,
  "status" : "InProgress",
  "alertCount" : 1,
  "averageAlertRiskScore" : 100,
  "sealed" : true,
  "totalRemediationTaskCount" : 4,
  "openRemediationTaskCount" : 5,
  "created" : "2018-01-01T04:49:27.870Z",
  "lastUpdated" : "2022-12-12T09:15:49.153Z",
  "lastUpdatedBy" : "duke",
  "assignee" : "tony",
  "sources" : [ "Malware Analysis" ],
  "ruleId" : "55e49a79e4b01a1d2be502bc",
  "firstAlertTime" : "2017-08-04T16:49:22Z",
  "categories" : [ {
    "id" : "55e49a79e4b01a1d2be5022e",
    "parent" : "Malware",
    "name" : "Password dumper"
  }, {
    "id" : "55e49a79e4b01a1d2be50228",
    "parent" : "Hacking",
    "name" : "Path traversal"
  } ],
  "journalEntries" : [ {
    "id" : "20",
    "author" : "admin",
    "notes" : "Updated status",
    "created" : "2017-11-15T20:20:54.785Z",
    "lastUpdated" : "2017-11-15T20:20:54.785Z",
    "milestone" : "Containment"
  } ],
  "createdBy" : "norm",
  "deletedAlertCount" : 100,
  "eventCount" : 0,
  "alertMeta" : {
    "SourceIp" : [ "10.11.12.345" ],
    "DestinationIp" : [ "11.11.11.111", "11.22.33.444" ]
  }
}
```


Remove an Incident

A single incident can be removed using the incident's unique identifier.

```
DELETE /rest/api/incidents/{id}
```

Path Parameters

Parameter	Description
id	The unique identifier of the incident.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/INC-100' -i -X DELETE \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Mon, 12 Dec 2022 09:15:51 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```
```

Add a Journal Entry

A journal entry, or note, can be added to an existing incident.

```
POST /rest/api/incidents/{id}/journal
```

Path Parameters

Parameter	Description
id	The unique identifier of the incident.

Request Fields

Path	Type	Description
author	String	The NetWitness user id of the user creating the journal entry.
notes	String	Notes and observations about the incident.
milestone	String	The incident milestone classifier.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/INC-100/journal' -i -X POST \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8' \
    -d '{"author":"duke","notes":"This incident is
contained.","milestone":"Containment"}'
```

Sample Response

```
HTTP/1.1 201 Created
Location: https://api.netwitness.local/rest/api/incidents/INC-100
Date: Mon, 12 Dec 2022 09:15:51 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```
```

Get an Incident's Alerts

All the alerts that are associated with an incident can be retrieved using the incident's unique identifier.

```
GET /rest/api/incidents/{id}/alerts
```

Path Parameters

Parameter	Description
id	The unique identifier of the incident.

Request Parameters

Parameter	Description
pageNumber	The requested page number.
pageSize	The maximum number of items to return in a single page.

Response Fields

Path	Type	Description
items	Array	An array containing the requested resources.
pageNumber	Number	The requested page number.

Path	Type	Description
pageSize	Number	The requested number of items to return in a single page.
totalPages	Number	The total number of pages available.
totalItems	Number	The total number of items available.
hasNext	Boolean	Indicates if there is a page containing results after this page.
hasPrevious	Boolean	Indicates if there is a page containing results before this page.
items[].id	String	The unique alert identifier.
items[].title	String	The title or name of the rule that created the alert.
items[].detail	String	The details of the alert. This can be the module name or meta that the module included.
items[].created	String	The timestamp of the alert created date.
items[].source	String	The source of this alert. For example, "Event Stream Analysis", "Malware Analysis", etc.
items[].riskScore	Number	The risk score of this alert, usually in the range 0 - 100.
items[].type	String	The type of alert, "Network", "Log", etc.
items[].events	Array	The events that make up this alert.
items[].events[].source	Object	The source of the event.
items[].events[].source.device	Object	The device contains the endpoint network information.
items[].events[].source.device.ipAddress	String	The IP address.
items[].events[].source.device.port	Number	The port.
items[].events[].source.device.macAddress	String	The ethernet MAC address.
items[].events[].source.device.dnsHostname	String	The DNS resolved hostname.
items[].events[].source.device.dnsDomain	String	The top-level domain from the DNS resolved hostname.
items[].events[].source.user	Object	The user contains the endpoint user information.
items[].events[].source.user.username	String	The unique username.
items[].events[].source.user.emailAddress	String	An email address.
items[].events[].source.user.adUsername	String	An Active Directory (AD) username.
items[].events[].source.user.adDomain	String	An Active Directory (AD) domain.

Path	Type	Description
items[].events[].destination	Object	The destination of the event.
items[].events[].destination.device	Object	The device contains the endpoint network information.
items[].events[].destination.device.ipAddress	String	The IP address.
items[].events[].destination.device.port	Number	The port.
items[].events[].destination.device.macAddress	String	The ethernet MAC address.
items[].events[].destination.device.dnsHostname	String	The DNS resolved hostname.
items[].events[].destination.device.dnsDomain	String	The top-level domain from the DNS resolved hostname.
items[].events[].destination.user	Object	The user contains the endpoint user information.
items[].events[].destination.user.username	String	The unique username.
items[].events[].destination.user.emailAddress	String	An email address.
items[].events[].destination.user.adUsername	String	An Active Directory (AD) username.
items[].events[].destination.user.adDomain	String	An Active Directory (AD) domain.
items[].events[].domain	String	The top-level domain or Windows domain.
items[].events[].eventSource	String	The source of the event. This may be a fully-qualified hostname with a port, or simple name.
items[].events[].eventSourceId	String	The unique identifier of the event on the source. For Network and Log events, this is the Nextgen Session ID.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/INC-
100/alerts?pageSize=10&pageNumber=0' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:51 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 1301
```


```
{
  "items" : [ {
    "id" : "5a6b81639491573f1e73676c",
    "title" : "LogOn Rule",
    "detail" : "Module_5a5cddb3e4b0ac40016df562_Alert",
    "created" : "2018-01-26T19:28:35Z",
    "source" : "Event Stream Analysis",
    "riskScore" : 90,
    "type" : "Network",
    "events" : [ {
      "source" : {
        "device" : {
          "ipAddress" : "58.229.117.56",
          "port" : 57429,
          "macAddress" : "00:13:c3:3b:c7:00",
          "dnsHostname" : null,
          "dnsDomain" : null
        },
        "user" : {
          "username" : "wwwrun",
          "emailAddress" : null,
          "adUsername" : null,
          "adDomain" : null
        }
      },
      "destination" : {
        "device" : {
          "ipAddress" : "128.164.35.184",
          "port" : 21,
          "macAddress" : "00:17:df:6b:c8:00",
          "dnsHostname" : null,
          "dnsDomain" : null
        },
        "user" : {
          "username" : "wwwrun",
          "emailAddress" : null,
          "adUsername" : null,
          "adDomain" : null
        }
      },
      "domain" : null,
```

```
      "eventSource" : "10.4.61.48:56005",
      "eventSourceId" : "9318"
    } ]
  } ],
  "pageNumber" : 0,
  "pageSize" : 10,
  "totalPages" : 1,
  "totalItems" : 1,
  "hasNext" : false,
  "hasPrevious" : false
}
```


Persisting events in an alert

Persist all the events present in the alert using the alert's unique identifier.

```
POST /rest/api/alerts/persist/{id}
```

Path Parameters

Parameter	Description
id	The unique identifier of the alert.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/alerts/persist/59440ad7e4b0ff674d2cd85e'
-i -X POST \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:54 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 87
```

```
{
  "status" : "SUCCESS",
  "message" : "Persisted alert id 59440ad7e4b0ff674d2cd85e"
}
```

Suspending persist of events in an alert

Suspend all the events present in persist alert using the alert's unique identifier.

```
POST /rest/api/alerts/suspend-persist/{id}
```

Path Parameters

Parameter	Description
id	The unique identifier of the alert.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/alerts/suspend-
persist/59440ad7e4b0ff674d2cd85e' -i -X POST \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:54 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 99
```

```
{
  "status" : "SUCCESS",
  "message" : "Suspended persist for alert id 59440ad7e4b0ff674d2cd85e"
}
```

Persisting events in an incident

Persist all the events present in the incident using the incident's unique identifier.

```
POST /rest/api/incidents/persist/{id}
```

Path Parameters

Parameter	Description
id	The unique identifier of the incident.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/persist/INC-75287' -i -X POST
\
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:54 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 106
```

```
{
  "status" : "SUCCESS",
  "message" : "Task successfully submitted to persist incident id - INC-75287"
}
```

Suspending persist of events in an incident

Suspend all the events present in persist incident using the incident's unique identifier.

```
POST /rest/api/incidents/suspend-persist/{id}
```

Path Parameters

Parameter	Description
id	The unique identifier of the incident.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/suspend-persist/INC-75287' -i
-X POST \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:54 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 87
```

```
{
  "status" : "SUCCESS",
  "message" : "Suspended persist for incident id INC-75287"
}
```

Get Incident history

Fetch the entire history details of the incident using the incident's unique identifier.

```
POST /rest/api/incidents/history/{id}
```

Path Parameters

Parameter	Description
id	The unique identifier of the incident.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/incidents/history/INC-75287' -i -X POST
\
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:55 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 608
```


```
[ {
  "type" : "STATUS",
  "date" : "2022-05-11T09:24:43.025+00:00",
  "changedBy" : "admin",
  "changedFrom" : "IN_PROGRESS",
  "changedTo" : "REMEDIATION_REQUESTED"
}, {
  "type" : "PRIORITY",
  "date" : "2022-05-11T09:25:18.767+00:00",
  "changedBy" : "admin",
  "changedFrom" : "HIGH",
  "changedTo" : "MEDIUM"
}, {
  "type" : "ASSIGNEE",
  "date" : "2022-05-11T09:25:31.157+00:00",
  "changedBy" : "admin",
  "changedFrom" : "rolden",
  "changedTo" : "admin"
}, {
  "type" : "CREATED",
  "date" : "2022-05-11T09:20:31.157+00:00",
  "changedBy" : "admin",
  "changedFrom" : null,
  "changedTo" : null
} ]
```


Get Incident Stats

IncidentStats can be retrieved by the date and time they were created. It contains day wise details of the overall MTTA, MTTD and MTTR as well as the count associated with each metric for that day.

```
GET /rest/api/incidents/stats
```

The requested date range can be unbounded, by only supplying either the `since` or `until` parameter, or bounded, by providing both parameters.

Request Parameters

Parameter	Description
`pageNumber`	The requested page number.
`pageSize`	The maximum number of items to return in a single page.
`since`	A date format (e.g., `2022-01-01`). Retrieve incidents stats created on and after this date.
`until`	A date format (e.g., `2022-01-31`). Retrieve incident stats created on and before this date.

All results will be returned using the paginated response payload sorted by the `created` date, in descending order.

Sample Request

```
$ curl
'https://api.netwitness.local/rest/api/incidents/stats?pageSize=100&pageNumber=0&since
=2022-05-01&until=2022-05-30' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:15:55 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 290
```

```
{
  "items" : [ {
    "date" : "28-May-22",
    "mtta" : 30,
    "mttaCount" : 3,
    "mttd" : 44,
    "mttdCount" : 1,
    "mttr" : 81,
    "mttrCount" : 1
  } ],
  "pageNumber" : 0,
  "pageSize" : 100,
  "totalPages" : 1,
  "totalItems" : 1,
  "hasNext" : false,
  "hasPrevious" : false
}
```


Get Incident User Stats

Incident User Stats can be retrieved by user and date range. It computes the MTTD values for the range of dates mentioned and responds with a consolidated list.

```
GET /rest/api/incidents/user-stats
```

The requested date range can be unbounded, by only supplying either the `since` or `until` parameter, or bounded, by providing both parameters.

Request Parameters

Snippet request-parameters not found for operation::incident-specification-it/get-incident-user-stats

All results will be returned based on the username input given.

Sample Request

Snippet curl-request not found for operation::incident-specification-it/get-incident-user-stats

Sample Response

Snippet http-response not found for operation::incident-specification-it/get-incident-user-stats

Services Information

Get Service IDs of all Services

The following API lists all services with their service IDs.

The response resource is comprised of the following attributes:

Path	Type	Description
`[]`	`Array`	Array of service information.
`[].id`	`String`	Unique identifier of each service installed in the RSA NetWitness suite.
`[].name`	`String`	Name of the service. For example, endpoint-server.
`[].displayName`	`String`	Display name of the service.
`[].host`	`String`	Host details of the service.
`[].version`	`String`	Version of the service.

A list of all services can be retrieved using the following API:

```
GET /rest/api/services
```

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/services' -i -X GET \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:00 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 542
```

```
[ {
  "id" : "5304012d-94c0-42c2-a39f-6720bdab0801",
  "name" : "endpoint-broker-server",
  "displayName" : "Endpoint Broker Server",
  "host" : "endpoint-server",
  "version" : "11.4.0.0"
}, {
  "id" : "5f6e15bb-9384-4666-aa71-e5d134afd7d8",
  "name" : "endpoint-server",
  "displayName" : "Endpoint Server",
  "host" : "endpoint-server",
  "version" : "11.4.0.0"
}, {
  "id" : "f8af0111-f8d4-4fce-b86e-d3efcd073733",
  "name" : "respond-server",
  "displayName" : "Respond Server",
  "host" : "endpoint-server",
  "version" : "11.4.0.0"
} ]
```


Get List of Service IDs by Service Name

```
GET /rest/api/services?name=<service-name>
```

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/services?name=endpoint-broker-server' -i
-X GET \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'Content-Type: application/json;charset=UTF-8'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:00 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 192
```

```
[ {
  "id" : "5304012d-94c0-42c2-a39f-6720bdab0801",
  "name" : "endpoint-broker-server",
  "displayName" : "Endpoint Broker Server",
  "host" : "endpoint-server",
  "version" : "11.4.0.0"
} ]
```

Endpoint APIs

Endpoint Log Hybrid collects and manages endpoint data from hosts. Using the APIs, analyst can:

- View list of host names with agent IDs from one or more Endpoint Servers.
- Retrieve host data, such as drivers, processes, DLLs, files (executables), services, autoruns, security information, anomalies, system configurations, and scripts found on the host.
- Filter hosts on indexed values and sort columns.
- Get list of snapshots and snapshot details for a host.
- Start or stop a scan.
- Get list of files from a specific Endpoint Server.
- Get alerts for a host or file.
- Download file/files to the server.
- Download system and process dump.
- Isolate host from the network.

Note: All endpoint APIs require a service ID to connect to the specific Endpoint Server.

Get Hosts

The Get Hosts API lists all hosts' information from a particular Endpoint Server. It provides a paged response with a standard paged response structure as mentioned in the 'Pagination' section.

The "items" field in paged response consists of individual host information.

Path	Type	Description
`items`	`Array`	An array containing the requested resources.
`pageNumber`	`Number`	The requested page number.
`pageSize`	`Number`	The requested number of items to return in a single page.
`totalPages`	`Number`	The total number of pages available.
`totalItems`	`Number`	The total number of items available.
`hasNext`	`Boolean`	Indicates if there is a page containing results after this page.
`hasPrevious`	`Boolean`	Indicates if there is a page containing results before this page.
`items[].agentId`	`String`	Agent ID of the host.
`items[].hostName`	`String`	Name of the host.
`items[].riskScore`	`Number`	Risk score of the host.
`items[].networkInterfaces`	`Array`	List of network interfaces with details, such as IP address and MAC address.

Path	Type	Description
items[].networkInterfaces[].name	String	Name of the network interface.
items[].networkInterfaces[].macAddress	String	MAC Address of the network interface.
items[].networkInterfaces[].ipv4	Array	List of IPV4 in the network interface.
items[].networkInterfaces[].ipv6	Array	List of IPV6 in the network interface.
items[].networkInterfaces[].networkIdv6	Array	List of network IDV6 in the network interface.
items[].networkInterfaces[].gateway	Array	List of gateway in the network interface.
items[].networkInterfaces[].dns	Array	List of DNS in the network interface.
items[].networkInterfaces[].promiscuous	String	Specifies if the network interface is in the promiscuous mode.
items[].lastSeenTime	String	Agent last seen time.

```
GET /rest/api/hosts?serviceId=<service-id>&&pageNumber=0&pageSize=100
```

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/hosts?serviceId=5f6e15bb-9384-4666-aa71-
e5d134afd7d8&pageNumber=0&pageSize=1' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:01 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 691
```

```
{
  "items" : [ {
    "agentId" : "B27DDED7-6FFA-A9D3-6577-3DDE367B2820",
    "hostName" : "WIN-854PACLCQ07-VC",
    "riskScore" : 0,
    "networkInterfaces" : [ {
      "name" : "Intel(R) 82574L Gigabit Network Connection",
      "macAddress" : "00:50:56:01:1B:31",
      "ipv4" : [ "10.40.15.202" ],
      "ipv6" : [ "fe80::651d:c035:a19b:ea03%13" ],
      "networkIdv6" : [ "fe80::%13" ],
      "gateway" : [ "10.40.12.1" ],
      "dns" : [ "10.31.64.22", "10.31.64.23" ],
      "promiscuous" : false
    } ],
    "lastSeenTime" : "2019-01-28T09:25:02.484Z"
  } ],
  "pageNumber" : 0,
  "pageSize" : 1,
  "totalPages" : 2,
  "totalItems" : 2,
  "hasNext" : true,
  "hasPrevious" : false
}
```


Get Hosts with Filter

The following fields are supported for filtering and sorting on Get Hosts API - 'agentId', 'hostName', 'riskScore' and 'networkInterfaces.ipv4'.

Filter and sort are specified as a part of the request body.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/hosts?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8&pageNumber=0&pageSize=1' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/json' \
    -d '{"criteria":{"criteriaList":[{"criteriaList":[],"expressionList":[{"propertyName":"hostName","restrictionType":"LIKE","propertyValues":[{"value":"WIN-854PACLCQ07-VC","relative":false}]}],"predicateType":"AND"},{"criteriaList":[],"expressionList":[{"propertyName":"riskScore","restrictionType":"BETWEEN","propertyValues":[{"value":0,"relative":false},{"value":100,"relative":false}]}],"predicateType":"OR"}],"expressionList":[],"predicateType":"AND"},"sort":{"keys":["riskScore"],"descending":true}}'
```

HTTP request

```
GET /rest/api/hosts?serviceId=5f6e15bb-9384-4666-aa71-
e5d134afd7d8&pageNumber=0&pageSize=1 HTTP/1.1
Accept: application/json;charset=UTF-8
NetWitness-Token: eyJ...AT
Content-Type: application/json
Host: api.netwitness.local
Content-Length: 498

{"criteria":{"criteriaList":[{"criteriaList":[],"expressionList":[{"propertyName":"hos
tName","restrictionType":"LIKE","propertyValues":[{"value":"WIN-854PACLCQ07-
VC","relative":false}]}],"predicateType":"AND"},{"criteriaList":[],"expressionList":[{
"propertyName":"riskScore","restrictionType":"BETWEEN","propertyValues":[{"value":0,"r
elative":false},{"value":100,"relative":false}]}],"predicateType":"OR"}],"expressionLi
st":[],"predicateType":"AND"},"sort":{"keys":["riskScore"],"descending":true}}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:03 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 692
```

```
{
  "items" : [ {
    "agentId" : "B27DDED7-6FFA-A9D3-6577-3DDE367B2820",
    "hostName" : "WIN-854PACLCQ07-VC",
    "riskScore" : 0,
    "networkInterfaces" : [ {
      "name" : "Intel(R) 82574L Gigabit Network Connection",
      "macAddress" : "00:50:56:01:1B:31",
      "ipv4" : [ "10.40.15.202" ],
      "ipv6" : [ "fe80::651d:c035:a19b:ea03%13" ],
      "networkIdv6" : [ "fe80::%13" ],
      "gateway" : [ "10.40.12.1" ],
      "dns" : [ "10.31.64.22", "10.31.64.23" ],
      "promiscuous" : false
    } ],
    "lastSeenTime" : "2019-01-28T09:25:02.484Z"
  } ],
  "pageNumber" : 0,
  "pageSize" : 1,
  "totalPages" : 1,
  "totalItems" : 1,
  "hasNext" : false,
  "hasPrevious" : false
}
```


The following is an example for filter criteria in the request body:

```
{
  "criteria": {
    "criteriaList": [
      {
        "expressionList": [{ "propertyName": "agentId", "restrictionType":
"EQUAL","propertyValues": [{"value": "2F53FC2C-A737-B34B-6813-12E48379C15D"}]}]
      }
    ]
  }
}
```

The following are the supported 'restrictionType'

- Operators that require no value: IS_NULL, IS_NOT_NULL.
- Operators that require one value: LIKE, NOT_LIKE, EQUAL, NOT_EQUAL, LESS_THAN,LESS_THAN_OR_EQUAL_TO, GREATER_THAN, GREATER_THAN_OR_EQUAL_TO.
- Operators that require two value: BETWEEN, NOT_BETWEEN.
- Operators that uses multiple value: IN, NOT_IN.

The following are the supported 'predicateType' - AND, OR, NOT.

Get List of Snapshots for Host

This API provides a list of snapshots, which are IDs to fetch the snapshot details of the host.

```
GET /rest/api/host/<Host-Agent-Id>/snapshots?serviceId=<service-id>
```

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/snapshots?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:03 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 30
```

```
[ "2017-12-22T14:34:05.985Z" ]
```

Get Snapshot Details for the Host

This API provides snapshot details of the given host for the provided snapshot time.

Response Fields

Path	Type	Description
[]	Array	Array of scan snapshot for the agent.
[].machineOsType	String	Type of operating system (Windows, Mac, Linux).
[].hostName	String	Name of the host.
[].agentId	String	Agent ID of the host.
[].agentVersion	String	Version of the agent.
[].scanStartTime	String	Start time of the scan snapshot.
[].directory	String	Directory of the file.
[].fileName	String	Name of the file.

Path	Type	Description
[].owner.username	String	User name of the owner of the file.
[].owner.groupname	String	Group name of the owner of the file.
[].owner.uid	String	UID of the user name.
[].owner.gid	String	GID of the user name.
[].timeCreated	String	Time when file was created.
[].timeModified	String	Time when file was modified.
[].timeAccessed	String	Time when file was last accessed.
[].attributes	Array	List of file attributes.
[].accessMode	Number	Access mode of the file.
[].sameDirectoryFileCounts.nonExe	Number	Number of non-exe files in the same directory of the file.
[].sameDirectoryFileCounts.exe	Number	Number of exe files in the same directory of the file.
[].sameDirectoryFileCounts.subFolder	Number	Number of sub-folders in the same directory of the file.
[].sameDirectoryFileCounts.exeSameCompany	Number	Number of executables with the same company name in the same directory of the file.
[].sameDirectoryFileCounts.hiddenFiles	Number	Count of hidden files in the same directory of the file.
[].fileContext	Array	List of file context.
[].directoryContext	Array	List of directory context.
[].autorunContext	Array	List of autorun context.
[].networkContext	Array	List of network context.
[].kernelModeContext	Array	List of kernel mode context.
[].userModeContext	Array	List of user mode context.
[].processContext	Array	List of process context.
[].rpm.packageName	String	RPM package name to which the file belongs.
[].rpm.category	String	Category to which the rpm package belongs.
[].windows.processes[].eprocess	String	Identifier of the process.
[].windows.processes[].sessionId	Number	Session ID of the process.
[].windows.processes[].parentPath	String	Directory of the parent process.
[].windows.processes[].imageSize	Number	Size of the process image.
[].windows.processes[].integrityLevel	Number	Integrity level of the process.
[].windows.processes[].context	Array	List of process context.
[].windows.processes[].pid	Number	ID of the process.
[].windows.processes[].parentPid	Number	ID of the parent process.

Path	Type	Description
[].windows.processes[].imageBase	Number	Base address of the process.
[].windows.processes[].createUtcTime	String	Creation time of the process.
[].windows.processes[].owner	String	Name of the user.
[].windows.processes[].launchArguments	String	Launch arguments of the process.
[].windows.processes[].threadCount	Number	Number of threads running in the process.
[].windows.dlls[].createTime	String	Creation time of the process.
[].windows.dlls[].eprocess	String	Identity of the process.
[].windows.dlls[].imageSize	Number	Size of the DLL image in memory.
[].windows.threads[].processName	String	Name of the process.
[].windows.threads[].processTime	String	Creation time of the process.
[].windows.threads[].eprocess	String	Identifier of the process.
[].windows.threads[].pid	Number	PID of the process.
[].windows.threads[].ethread	String	Identifier of the thread.
[].windows.threads[].tid	Number	ID of the thread.
[].windows.threads[].teb	String	Address of thread environment block.
[].windows.threads[].startAddress	String	Start address of the thread in memory.
[].windows.threads[].state	Array	Thread state.
[].windows.threads[].behaviorKey	String	Floating behavior resolution of the thread.
[].windows.drivers[].windows.imageBase	Number	Base address of the driver image.
[].windows.drivers[].windows.imageSize	Number	Size of the driver image.
[].windows.services[].windows.serviceName	String	Service name as identified by the system.
[].windows.services[].displayName	String	Display name for the service.
[].windows.services[].description	String	Description of the service.
[].windows.services[].account	String	Name of the user the service executes as.
[].windows.services[].launchArguments	String	Launch arguments of the service.
[].windows.services[].windows.serviceMain	String	Service’s main.
[].windows.services[].hostingPid	Number	Service’s hosting process ID.

Path	Type	Description
[].windows.services[].state	String	Current state of the service.
[].windows.services[].win32ErrorCode	Number	Last Windows 32 error code from registry.
[].windows.services[].context	Array	List of service context.
[].windows.tasks[].name	String	Name of the task.
[].windows.tasks[].executeUser	String	Name of the user the task executes as.
[].windows.tasks[].creatorUser	String	Name of the user who created the task.
[].windows.tasks[].launchArguments	String	Launch arguments of the task.
[].windows.tasks[].status	Array	Status of the task.
[].windows.tasks[].lastRunTime	String	Time when the task was last run.
[].windows.tasks[].nextRunTime	String	Next scheduled time of the task.
[].windows.tasks[].triggerString	String	Textual trigger string of the task.
[].windows.autoruns[].type	String	Type of the autorun.
[].windows.autoruns[].registryPath	String	Registry path where autorun is located.
[].windows.autoruns[].launchArguments	String	Launch argument of the autorun.
[].windows.imageHooks[].process.pid	String	PID of the process in which hook was detected.
[].windows.imageHooks[].process.fileName	String	Filename of the process in which hook was detected.
[].windows.imageHooks[].process.createUtcTime	String	Creation time of the process in which hook was detected.
[].windows.imageHooks[].hookLocation.section	String	Name of the image section that was modified by the hook.
[].windows.imageHooks[].hookLocation.sectionBase	String	Base of the image section that was modified by the hook.
[].windows.imageHooks[].hookLocation.symbol	String	Closest symbol name to the memory location that was modified.
[].windows.imageHooks[].hookLocation.symbolOffset	Number	Closest symbol +/- offset to the hook location when relevant.
[].windows.imageHooks[].inlinePatch.originalBytes	String	Hexadecimal bytes which were replaced.
[].windows.imageHooks[].inlinePatch.originalAsm	Array	Array of decoded ASM instructions that were replaced.
[].windows.imageHooks[].inlinePatch.currentBytes	String	Hexadecimal bytes which have overwritten the original code.
[].windows.imageHooks[].inlinePatch.currentAsm	Array	Array of decoded ASM instructions that have overwritten the original code.

Path	Type	Description
[].windows.kernelHooks[].hookLocation.objectName	String	Name of the object that was hooked in kernel.
[].windows.kernelHooks[].hookLocation.objectFunction	String	Name of the object function that was hooked in kernel.
[].mac.processes[].priority	Number	Priority of the process.
[].mac.processes[].flags	Number	Process flags.
[].mac.processes[].nice	Number	Nice value of process.
[].mac.processes[].openFilesCount	Number	Number of open files by process at scan time.
[].mac.processes[].context	Array	Process context.
[].mac.processes[].pid	Number	ID of the process.
[].mac.processes[].parentPid	Number	ID of the parent process.
[].mac.processes[].imageBase	Number	Base address of the process.
[].mac.processes[].createUtcTime	String	Creation time of the process.
[].mac.processes[].owner	String	Name of the user.
[].mac.processes[].launchArguments	String	Launch arguments of the process.
[].mac.processes[].threadCount	Number	Number of threads running in the process.
[].mac.dylibs[].pid	Number	Process ID in dylib which is loaded.
[].mac.dylibs[].processName	String	Name of the process.
[].mac.dylibs[].imageBase	String	Base address of image in the process.
[].mac.drivers[].preLinked	Boolean	True if Kext bundle is prelinked.
[].mac.drivers[].numberOfReferences.	Number	Number of references.
[].mac.drivers[].dependencies	Array	List of kexts(name) the driver is linked against.
[].mac.drivers[].imageBase	String	Base address of the driver image.
[].mac.drivers[].imageSize	String	Size of the driver image.
[].mac.daemons[].name	String	Label of the daemon.
[].mac.daemons[].sessionName	String	Name of the session in which daemon runs.
[].mac.daemons[].user	String	Name of the user under which the daemon runs.
[].mac.daemons[].pid	Number	ID of the process.
[].mac.daemons[].onDemand	Boolean	True if daemon is configured to run on demand.
[].mac.daemons[].lastExitCode	Number	Last exit code.
[].mac.daemons[].timeout	Number	Time out value.
[].mac.daemons[].daemons.launchArguments	String	Launch argument of the daemon.
[].mac.daemons[].daemons.configFile	String	Full path of the configuration file used to configure this daemon.

Path	Type	Description
[].mac.tasks[].name	String	Name of the task.
[].mac.tasks[].cronJob	Boolean	True if the task is cron job, else launchd.
[].mac.tasks[].launchArguments	String	Launch argument of the task.
[].mac.tasks[].user	String	Name of the user under which this task will run.
[].mac.tasks[].triggerString	String	Trigger string of the task.
[].mac.tasks[].configFile	String	Full path of the configuration file used to configure this task.
[].mac.autoruns[].type	String	Type of autorun.
[].mac.autoruns[].user	String	Name of the user under which the autorun is run.
[].mac.autoruns[].name	String	Label of the autorun.
[].mac.autoruns[].detail	String	Details of the autorun.
[].linux.processes[].priority	Number	Priority of the process.
[].linux.processes[].uid	Number	UID of the user.
[].linux.processes[].environment	String	Environment variables.
[].linux.processes[].nice	Number	Nice value of the process.
[].linux.processes[].securityContext	String	Security context.
[].linux.processes[].pid	Number	ID of the process.
[].linux.processes[].parentPid	Number	ID of the parent process.
[].linux.processes[].imageBase	Number	Base address of the process.
[].linux.processes[].createUtcTime	String	Time of creation of the process.
[].linux.processes[].owner	String	Name of the user.
[].linux.processes[].launchArguments	String	Launch arguments of the process.
[].linux.processes[].threadCount	Number	Number of threads running in the process.
[].linux.loadedLibraries[].pid	String	Process ID in which library is loaded.
[].linux.loadedLibraries[].processName	String	Name of the process.
[].linux.loadedLibraries[].imageBase	String	Base address of image in the process.
[].linux.drivers[].numberOfInstances	Number	Number of instances loaded in memory.
[].linux.drivers[].loadState	String	Load state of the driver.
[].linux.drivers[].dependencies	Array	Dependent driver names.
[].linux.drivers[].author	String	Name of the author of driver.
[].linux.drivers[].description	String	Description of the driver.

Path	Type	Description
[].linux.drivers[].sourceVersion	String	Source version of the driver.
[].linux.drivers[].versionMagic	String	Version magic of the driver.
[].linux.initds[].initdHashSha256	String	Hash of the init-d script file.
[].linux.initds[].initdPaths	String	Path of the init-d script file.
[].linux.initds[].pid	Number	ID of the process.
[].linux.initds[].description	String	Description of the init-d.
[].linux.initds[].status	String	Status of the init-d.
[].linux.initds[].runLevels	Array	List of run levels in which the init-d is enabled.
[].linux.systemds[].systemdHashSha256	String	Hash value of the systemd script file.
[].linux.systemds[].systemdPaths	String	Path value of the systemd script file.
[].linux.systemds[].name	String	Name of the systemd.
[].linux.systemds[].description	String	Description of the systemd.
[].linux.systemds[].state	String	State of the systemd.
[].linux.systemds[].launchArguments	String	Launch argument of the systemd.
[].linux.systemds[].pid	Number	ID of the process.
[].linux.systemds[].triggeredBy	Array	Triggered by list of the systemd.
[].linux.systemds[].triggerStrings	Array	Trigger strings of the systemd.
[].linux.autoruns[].type	String	Type of autorun.
[].linux.autoruns[].label	String	Label of the autorun.
[].linux.autoruns[].comments	String	Comments of the autorun.
[].linux.crons[].user	String	User account under which cron job was created.
[].linux.crons[].triggerString	String	Trigger string that would launch the cron job.
[].linux.crons[].launchArguments	String	Launch arguments of the cron job.
firstFileName	String	First name of the file sent by the agent.
reputationStatus	String	Reputation status of the file.
globalRiskScore	String	Global risk score.
firstSeenTime	String	Time when the file was first seen by the Endpoint Server.
machineOsType	String	Type of operating system (Windows, Mac, Linux).
signature	Object	Signatory information of the file.
signature.timeStamp	String	Timestamp of the signature.
signature.thumbprint	String	Thumbprint of the certificate.
signature.context	Array	Context information of the certificate.
signature.signer	String	Signer information of the certificate.

Path	Type	Description
size	String	Size of the file.
checksumMd5	String	MD5 of the file.
checksumSha1	String	SHA1 of the file.
checksumSha256	String	SHA256 of the file.
pe	Object	PE information of the file. This is applicable for Windows files.
pe.timeStamp	String	Timestamp of the PE File.
pe.imageSize	String	Image size of the PE file.
pe.numberOfExportedFunctions	String	Number of exported function in the PE file.
pe.numberOfNamesExported	String	Number of names exported in the PE file.
pe.numberOfExecuteWriteSections	String	Number of execute write sections in the PE file.
pe.context	Array	Context information of the PE file.
pe.resources	Object	Resources of the PE file.
pe.resources.originalFileName	String	Original filename as per PE information.
pe.resources.company	String	Company name as per PE information.
pe.resources.description	String	Description of the file as per PE information.
pe.resources.version	String	Version of the file as per PE information.
pe.sectionNames	Array	List of section names in the PE file.
pe.importedLibraries	Array	List of imported libraries in the PE file.
elf	Object	ELF information of the file. This is applicable for Linux files.
elf.classType	String	Class type of the ELF file.
elf.data	String	Data of ELF file.
elf.entryPoint	String	Entry point for the ELF file.
elf.context	Array	Context information of ELF file.
elf.type	String	Type of ELF file.
elf.sectionNames	Array	List of section names in ELF file.
elf.importedLibraries	Array	List of imported libraries in ELF file.
macho	Object	Macho information of the file. This is applicable for Mac files.
macho.uuid	String	UUID of the Macho file.
macho.identifier	String	Identifier of the Macho file.
macho.minOsxVersion	String	Minimum OSx version for the Macho file.
macho.context	Array	Context information of the Macho file.
macho.flags	String	Flags of Macho file.
macho.numberOfLoadCommands	String	Number of load commands for the Macho file.
macho.version	String	Version of the Macho file.

Path	Type	Description
macho.sectionNames	Array	Section names in the Macho file.
macho.importedLibraries	Array	Imported libraries list in the Macho file.
entropy	String	Entropy of the file.
format	String	Format of the file.
fileStatus	String	Status of the file as assigned by the analyst. (Whitelist, Blacklist, Neutral, and Graylist).
remediationAction	String	Remediation action as assigned by the analyst. For example, Blocked.
[].localRiskScore	Number	File's score based on alerts triggered in the given agent.

```
GET /rest/api/host/<Host-Agent-Id>/snapshots/2019-06-
17T04:24:14.608Z?serviceId=<service-id>
```

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/snapshots/2017-12-22T14%3A34%3A05.985Z?serviceId=67a84b72-3d9a-4377-9096-
7e6af9f13306' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 22 Jul 2019 08:34:38 GMT
Content-Length: 6710
Content-Disposition: inline;filename=f.txt
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
```


```
[ {
 "machineOsType" : "windows",
 "hostName" : "HOSTNAME",
 "agentId" : "B27DDED7-6FFA-A9D3-6577-3DDE367B2820",
 "agentVersion" : "11.1.0.0",
 "scanStartTime" : "2017-12-22T14:34:05.985Z",
 "directory" : "F\\SonarQube\\DIRECTCP\\PolicyDefinitions",
 "fileName" : "shtctky.bat",
 "owner" : {
   "username" : "Sinha, Vidya",
   "groupname" : "CORP",
```

```
    "uid" : 9,
    "gid" : 1
  },
  "timeCreated" : "2017-12-22T10:00:15.000Z",
  "timeModified" : "2017-12-22T06:00:15.000Z",
  "timeAccessed" : "2017-12-22T11:00:15.000Z",
  "attributes" : null,
  "accessMode" : 15,
  "sameDirectoryFileCounts" : {
    "nonExe" : 1,
    "exe" : 4,
    "subFolder" : 0,
    "exeSameCompany" : 3,
    "hiddenFiles" : 0
  },
  "fileContext" : [ "ads", "accessDenied", "hiddenDifferentialView", "encrypted" ],
  "directoryContext" : [ "desktop" ],
  "autorunContext" : [ "winlogon" ],
  "networkContext" : [ "accessNetwork", "listen" ],
  "kernelModeContext" : [ "loaded", "hookEat", "hookSsdt", "createThreadNotification",
"imageMismatch", "remoteThreadCreator" ],
  "userModeContext" : [ "loaded", "hookEat", "mapped", "image", "threadFloating",
"remoteMemoryAllocator", "setWindowsHook" ],
  "processContext" : [ "accessDenied", "dyldInserted", "ldPreloaded" ],
  "rpm" : null,
  "windows" : {
    "processes" : [ {
      "pid" : 46270,
      "parentPid" : 4,
      "imageBase" : 55076,
      "createUtcTime" : null,
      "owner" : "Sinha, Vidya",
      "launchArguments" : "/B /nologo %systemroot%\\system32\\calluxxprovider.vbs
$(Arg0) $(Arg1) $(Arg2)",
      "threadCount" : 0,
      "eprocess" : "0x8F76",
      "sessionId" : 1,
      "parentPath" : null,
      "imageSize" : 0,
      "integrityLevel" : 0,
      "context" : [ "UsingNamedPipe" ]
    } ],
    "dlls" : [ {
      "pid" : 46270,
      "processName" : null,
      "imageBase" : 55590,
      "createTime" : "2017-12-22T16:00:15.000+0000",
      "eprocess" : "0x3DFE",
      "imageSize" : 55459
    } ],
    "threads" : [ ],
```

```
    "drivers" : [ ],
    "services" : [ {
      "serviceName" : "wlanext",
      "displayName" : "wlanext",
      "description" : "wlanext description",
      "account" : "Sinha, Vidya",
      "launchArguments" : "-id 1",
      "serviceMain" : "ServiceMain",
      "hostingPid" : 0,
      "state" : null,
      "win32ErrorCode" : 26218,
      "context" : null
    } ],
    "tasks" : [ {
      "name" : "Shaktiman",
      "executeUser" : "Sinha, Vidya",
      "creatorUser" : "Sinha, Vidya",
      "launchArguments" : "-Embedding",
      "status" : null,
      "lastRunTime" : null,
      "nextRunTime" : null,
      "triggerString" : null
    } ],
    "autoruns" : [ ],
    "imageHooks" : [ ],
    "kernelHooks" : [ ]
  },
  "mac" : null,
  "linux" : null,
  "fileProperties" : null,
  "localRiskScore" : 0
} ]
```


To filter the snapshot details based on the category, the following fields can be provided as path parameter - PROCESSES, LOADED_LIBRARIES, SERVICES, AUTORUNS, TASKS, DRIVERS, THREADS, IMAGE_HOOKS, and KERNEL_HOOKS.

```
GET /rest/api/host/<Host-Agent-Id>/snapshots/2019-06-
17T04:24:14.608Z?serviceId=<service-id>&categories=PROCESSES
```

Get Files

This API lists all related information of files from a specific Endpoint Server. These information are specific to the unique file and does not include any host information.

It provides a paged response with a standard paged response structure as mentioned in the 'Pagination' section.

The "items" field in paged response consists of individual file information.

Path	Type	Description
items	Array	An array containing the requested resources.
pageNumber	Number	The requested page number.
pageSize	Number	The requested number of items to return in a single page.
totalPages	Number	The total number of pages available.
totalItems	Number	The total number of items available.
hasNext	Boolean	Indicates if there is a page containing results after this page.
hasPrevious	Boolean	Indicates if there is a page containing results before this page.
items[].firstFileName	String	First name of the file sent by the agent.
items[].reputationStatus	String	Reputation status of the file.
items[].globalRiskScore	String	Global risk score.
items[].firstSeenTime	String	Time when the file was first seen by the Endpoint Server.
items[].machineOsType	String	Type of operating system (Windows, Mac, Linux).
items[].signature	Object	Signatory information of the file.
items[].signature.timeStamp	String	Timestamp of the signature.
items[].signature.thumbprint	String	Thumbprint of the certificate.
items[].signature.context	Array	Context information of the certificate.
items[].signature.signer	String	Signer information of the certificate.
items[].size	String	Size of the file.
items[].checksumMd5	String	MD5 of the file.
items[].checksumSha1	String	SHA1 of the file.
items[].checksumSha256	String	SHA256 of the file.
items[].pe	Object	PE information of the file. This is applicable for Windows files.
items[].pe.timeStamp	String	Timestamp of the PE File.
items[].pe.imageSize	String	Image size of the PE file.
items[].pe.numberOfExportedFunctions	String	Number of exported function in the PE file.
items[].pe.numberOfNamesExported	String	Number of names exported in the PE file.
items[].pe.numberOfExecuteWriteSections	String	Number of execute write sections in the PE file.
items[].pe.context	Array	Context information of the PE file.
items[].pe.resources	Object	Resources of the PE file.
items[].pe.resources.originalFileName	String	Original filename as per PE information.

Path	Type	Description
items[].pe.resources.company	String	Company name as per PE information.
items[].pe.resources.description	String	Description of the file as per PE information.
items[].pe.resources.version	String	Version of the file as per PE information.
items[].pe.sectionNames	Array	List of section names in the PE file.
items[].pe.importedLibraries	Array	List of imported libraries in the PE file.
items[].elf	Object	ELF information of the file. This is applicable for Linux files.
items[].elf.classType	String	Class type of the ELF file.
items[].elf.data	String	Data of ELF file.
items[].elf.entryPoint	String	Entry point for the ELF file.
items[].elf.context	Array	Context information of ELF file.
items[].elf.type	String	Type of ELF file.
items[].elf.sectionNames	Array	List of section names in ELF file.
items[].elf.importedLibraries	Array	List of imported libraries in ELF file.
items[].macho	Object	Macho information of the file. This is applicable for Mac files.
items[].macho.uuid	String	UUID of the Macho file.
items[].macho.identifier	String	Identifier of the Macho file.
items[].macho.minOsxVersion	String	Minimum OSx version for the Macho file.
items[].macho.context	Array	Context information of the Macho file.
items[].macho.flags	String	Flags of Macho file.
items[].macho.numberOfLoadCommands	String	Number of load commands for the Macho file.
items[].macho.version	String	Version of the Macho file.
items[].macho.sectionNames	Array	Section names in the Macho file.
items[].macho.importedLibraries	Array	Imported libraries list in the Macho file.
items[].entropy	String	Entropy of the file.
items[].format	String	Format of the file.
items[].fileStatus	String	Status of the file as assigned by the analyst. (Whitelist, Blacklist, Neutral, and Graylist).
items[].remediationAction	String	Remediation action as assigned by the analyst. For example, Blocked.

Note: The following is a sample response with all fields populated. However, the response for pe, macho, and elf is generated based the operating system type. The fields without values will display as null.

```
GET /rest/api/files?serviceId=<service-id>&&pageNumber=0&pageSize=90
```

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/files?serviceId=5f6e15bb-9384-4666-aa71-
e5d134afd7d8&pageNumber=0&pageSize=1' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:05 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 2054
```


```
{
  "items" : [ {
    "firstFileName" : "cmd.exe",
    "reputationStatus" : "Known",
    "globalRiskScore" : 0,
    "firstSeenTime" : "2019-04-28T05:40:20.000Z",
    "machineOsType" : "windows",
    "signature" : {
      "timeStamp" : "2019-04-28T05:40:20.000Z",
      "thumbprint" : "ae9c1ae54763822eec42474983d8b635116c8452",
      "context" : [ "microsoft", "signed", "valid", "catalog" ],
      "signer" : "Microsoft Windows"
    },
    "size" : 278528,
    "checksumMd5" : "0d088f5bcfa8f086fba163647cd80cab",
    "checksumSha1" : "08cc2e8dca652bdda1acca9c446560d4bc1bcdf9",
    "checksumSha256" :
"9023f8aaeda4a1da45ac477a81b5bbe4128e413f19a0abfa3715465ad66ed5cd",
    "pe" : {
      "timeStamp" : "2019-04-28T05:40:20.000Z",
      "imageSize" : 413696,
      "numberOfExportedFunctions" : 0,
      "numberOfNamesExported" : 0,
      "numberOfExecuteWriteSections" : 0,
      "context" : [ "file.exe" ],
      "resources" : {
        "originalFileName" : "Cmd.Exe",
        "company" : "Microsoft Corporation",
        "description" : "Windows Command Processor",
        "version" : null
```

```
      },
      "sectionNames" : [ ".text" ],
      "importedLibraries" : [ "msvcrt.dll" ]
    },
    "elf" : {
      "classType" : 0,
      "data" : 0,
      "entryPoint" : 0,
      "context" : [ "file.arch64", "file.lkm" ],
      "type" : 1,
      "sectionNames" : [ ".note.gnu.build-id", ".text" ],
      "importedLibraries" : null
    },
    "macho" : {
      "uuid" : "277163DE-842E-390D-A7FF-EC4CF2D211A4",
      "identifier" : "com.apple.geod",
      "minOsxVersion" : "10.11.0",
      "context" : [ "file.arch64" ],
      "flags" : 2097285,
      "numberOfLoadCommands" : 22,
      "version" : "1151.49.1",
      "sectionNames" : [ "__PAGEZERO" ],
      "importedLibraries" : [ "Foundation" ]
    },
    "entropy" : 6.17224886172381,
    "format" : "pe",
    "fileStatus" : "Blacklist",
    "remediationAction" : "Unblock"
  } ],
  "pageNumber" : 0,
  "pageSize" : 1,
  "totalPages" : 2,
  "totalItems" : 2,
  "hasNext" : true,
  "hasPrevious" : false
}
```


Request Scan

This API starts a scan for the host with the specified agent ID.

```
POST /rest/api/host/{agentId}/scan?serviceId=<service-id>&scanType=<scanType>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.
scanType	Type of scan command.
cpuMax	You can use cpuMax to specify the amount of CPU the agent can use to run the scan. You can choose a value from 5 to 100. If you do not specify a value, the agent uses the default 25% CPU for the scan

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/scan?serviceId=5f6e15bb-9384-4666-aa71-
e5d134afd7d8&scanType=QUICK_SCAN&cpuMax=67' -i -X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1'
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:05 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```
```

Request Scan with given CPU Usage

This API starts a scan for the host with the specified agent ID and given CPU usage.

```
POST /rest/api/host/{agentId}/scan?serviceId=<service-
id>&scanType=<scanType>&cpuMax=<cpuMax>
```

Path Parameters

Snippet path-parameters not found for operation::endpoint-api-specification-it/request-scan-for-a-host-with-cpuMax/

Request Parameters

Snippet request-parameters not found for operation::endpoint-api-specification-it/request-scan-for-a-host-with-cpuMax/

Sample Request

Snippet curl-request not found for operation::endpoint-api-specification-it/request-scan-for-a-host-with-cpuMax/

Sample Response

Snippet http-response not found for operation::endpoint-api-specification-it/request-scan-for-a-host-with-cpuMax/

- cpuMax: Minimum value allowed is 5 and Maximum value allowed is 100.

Request Stop Scan

This API stops a scan for the host with the specified agent ID.

```
DELETE /rest/api/host/{agentId}/scan?serviceId=<service-id>&scanType=<scanType>
```

Path Parameters

Parameter	Description
`agentId`	Unique identifier of the host.

Request Parameters

Parameter	Description
`serviceId`	Service ID of the Endpoint Server to be connected.
`scanType`	Type of scan command.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/scan?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8&scanType=CANCEL_SCAN'
-i -X DELETE \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:05 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

Get Alerts for a Host

This API gets all alerts triggered for a given host. Alerts are categorized as 'Critical', 'High', 'Medium' and 'Low'. The response provides category level count of alerts triggered along with list of alerts.

Path	Type	Description
id	String	ID of the entity for which score needs to be queried. Agent ID in case of host and checksum in case of files.
distinctAlertCount	Object	Count of distinct Alert/category for the entity.
distinctAlertCount.critical	String	Number of critical alerts.
distinctAlertCount.high	String	Number of high alerts.
distinctAlertCount.medium	String	Number of medium alerts.
distinctAlertCount.low	String	Number of low alerts.
categorizedAlerts	String	Count of alert and events for a file/host, categorized by severity.

```
GET /rest/api/host/{agentId}/alerts?serviceId=<service-id>
```

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/alerts?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i -X GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:05 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 855
```

```
{
  "id" : "B27DDED7-6FFA-A9D3-6577-3DDE367B2820",
  "distinctAlertCount" : {
    "critical" : 0,
    "high" : 4,
    "medium" : 2,
    "low" : 0
  },
  "categorizedAlerts" : {
    "High" : {
      "Possibly Renamed net.exe Detected" : {
        "alertCount" : 1,
        "eventCount" : 1
      },
      "Powershell Injects Remote Process" : {
        "alertCount" : 3,
        "eventCount" : 3
      },
      "Unexpected taskhostw.exe Parent" : {
        "alertCount" : 1,
        "eventCount" : 1
      },
      "Unexpected runtimebroker.exe Parent" : {
        "alertCount" : 1,
        "eventCount" : 1
      }
    },
    "Medium" : {
      "Performs Scripted File Transfer" : {
        "alertCount" : 1,
        "eventCount" : 1
      },
      "Unexpected Svchost.Exe Parent" : {
        "alertCount" : 1,
        "eventCount" : 1
      }
    }
  }
}
```


To filter alerts based on the category, add the alertCategory path parameter.

```
GET /rest/api/host/{agentId}/alerts?serviceId=<service-id>&alertCategory=Medium
```

Get Alerts for a File

This API gets all alerts triggered for a given file. Alerts are categorized as 'Critical', 'High', 'Medium' and 'Low'. The response provides category level count of alerts triggered along with list of alerts.

Only checksum supported are `sha256` and `md5`.

Path	Type	Description
`id`	`String`	ID of the entity for which score needs to be queried. Agent ID in case of host and checksum in case of files.
`distinctAlertCount`	`Object`	Count of distinct Alert/category for the entity.
`distinctAlertCount.critical`	`String`	Number of critical alerts.
`distinctAlertCount.high`	`String`	Number of high alerts.
`distinctAlertCount.medium`	`String`	Number of medium alerts.
`distinctAlertCount.low`	`String`	Number of low alerts.
`categorizedAlerts`	`String`	Count of alert and events for a file/host, categorized by severity.

```
GET /rest/api/file/{checksum}/alerts?serviceId=<service-id>
```

Sample Request

```
$ curl
'https://api.netwitness.local/rest/api/file/d1c79a36593f0d5f7d07502b963d97acc851dc0291
f4556ce8f110a58a48fda4/alerts?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i -X
GET \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT'
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
Content-Length: 569
```

```
{
  "id" : "d1c79a36593f0d5f7d07502b963d97acc851dc0291f4556ce8f110a58a48fda4",
  "distinctAlertCount" : {
    "critical" : 0,
    "high" : 2,
    "medium" : 1,
    "low" : 0
  },
  "categorizedAlerts" : {
    "High" : {
      "Possibly Renamed net.exe Detected" : {
        "alertCount" : 1,
        "eventCount" : 1
      },
      "Powershell Injects Remote Process" : {
        "alertCount" : 3,
        "eventCount" : 3
      }
    },
    "Medium" : {
      "Performs Scripted File Transfer" : {
        "alertCount" : 1,
        "eventCount" : 1
      }
    }
  }
}
```


To filter alerts based on the category, add the alertCategory path parameter.

```
GET /rest/api/file/{checksum}/alerts?serviceId=<service-id>&alertCategory=Medium
```

Request File Download to Server

This API initiates the file download to the Endpoint Server. The following are the fields in the request body:

- path: Full path where the files may be present.

```
POST /rest/api/host/{agentId}/download/download-file?serviceId=<service-id>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/download/download-file?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i
-X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/json' \
    -d '{"path":"C:\\Users\\sample\\test.exe"}'
```

HTTP request

```
POST /rest/api/host/B27DDED7-6FFA-A9D3-6577-3DDE367B2820/download/download-
file?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8 HTTP/1.1
Accept: application/json;charset=UTF-8
NetWitness-Token: eyJ...AT
Content-Type: application/json
Host: api.netwitness.local
Content-Length: 38

{"path":"C:\\Users\\sample\\test.exe"}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```
```

Request Multiple File Downloads to Server

This API initiates the multiple file downloads to the Endpoint Server matching the wildcard path. The following are the fields in the request body:

- path: Path where the files may be present using wildcard.
- countFiles: Maximum number of files returned by the host matching the wildcard path (default 10).

- maxFileSize: Maximum size of each file (in MB) (default 100 MB).

```
POST /rest/api/host/{agentId}/download/download-files?serviceId=<service-id>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/download/download-files?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8'
-i -X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/json' \
    -d '{"path":"C:\\Users\\sample\\*","countFiles":3,"maxFileSize":150}'
```

HTTP request

```
POST /rest/api/host/B27DDED7-6FFA-A9D3-6577-3DDE367B2820/download/download-
files?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8 HTTP/1.1
Accept: application/json;charset=UTF-8
NetWitness-Token: eyJ...AT
Content-Type: application/json
Host: api.netwitness.local
Content-Length: 64

{"path":"C:\\Users\\sample\\*","countFiles":3,"maxFileSize":150}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

Request MFT Download

This API initiates the MFT download to the Endpoint Server.

```
POST /rest/api/host/{agentId}/download/mft?serviceId=<service-id>&path=<path>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.
path	Drive or NTFS Mount path for which MFT is requested

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/download/mft?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8&path=E%3A%5C'
-i -X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1'
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

Request parameter "path" can be any valid windows drive or windows path. Below are the valid sample values

- C:\
- d:\

- E:\mountPath\NTFSmount
- f:\mount path\NTFS Mount

Note: MFT Download is not supported in Relay Mode

Request System Dump Download

This API initiates the download of the system dump to the Endpoint Server.

```
POST /rest/api/host/{agentId}/download/system-dump?serviceId=<service-id>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/download/system-dump?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i
-X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1'
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```
```

Note: System Dump Download is not supported in Relay Mode

Request Process Dump Download

This API initiates the download of the process dump to the Endpoint Server.

Process information are specified as a part of the request body.

```
POST /rest/api/host/{agentId}/download/process-dump?serviceId=<service-id>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/download/process-dump?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i
-X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/json' \
    -d
'{"processId":5744,"eprocess":"0xFFFFE10DC62C6440","fileName":"spoolsvX.bat","path":"E
\\Windows\\ReportServer\\PolicyDefinitions","hash":"687685b7531648c39fbb24fa81312b7fd2
e3ece1bf1347b386f8725783767e5c","processCreateUtcTime":1670836387232}'
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```
```

Request Network Isolation

This API isolates the host with the specified agent ID from the network.

```
POST /rest/api/host/{agentId}/isolation?serviceId=<service-id>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/isolation?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i -X POST \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/json' \
    -d
'{"allowDnsOnlyBySystem":false,"exclusions":[{"ip":"10.125.0.1","v4":true}],"comment":
"Found malicious"}'
```

HTTP request

```
POST /rest/api/host/B27DDED7-6FFA-A9D3-6577-3DDE367B2820/isolation?serviceId=5f6e15bb-
9384-4666-aa71-e5d134afd7d8 HTTP/1.1
Accept: application/json;charset=UTF-8
NetWitness-Token: eyJ...AT
Content-Type: application/json
Host: api.netwitness.local
Content-Length: 103

{"allowDnsOnlyBySystem":false,"exclusions":[{"ip":"10.125.0.1","v4":true}],"comment":"
Found malicious"}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```

```

Update Network Isolation Exclusion List

This API updates the network isolation exclusion list for the host with the specified agent ID.

```
PATCH /rest/api/host/{agentId}/isolation?serviceId=<service-id>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/isolation?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i -X PATCH \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/json' \
    -d
'{"allowDnsOnlyBySystem":false,"exclusions":[{"ip":"1.2.3.4","v4":true},{"ip":"10.125.
0.1","v4":true}],"comment":"Updating IP exclusions for network isolation"}'
```

HTTP request

```
PATCH /rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/isolation?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8 HTTP/1.1
Accept: application/json;charset=UTF-8
NetWitness-Token: eyJ...AT
Content-Type: application/json
Host: api.netwitness.local
Content-Length: 159

{"allowDnsOnlyBySystem":false,"exclusions":[{"ip":"1.2.3.4","v4":true},{"ip":"10.125.0
.1","v4":true}],"comment":"Updating IP exclusions for network isolation"}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```
```

Release from Network Isolation

This API restores the network connection and removes IP addresses added to the exclusion list for the host with the specified agent ID.

```
DELETE /rest/api/host/{agentId}/isolation?serviceId=<service-id>
```

Path Parameters

Parameter	Description
agentId	Unique identifier of the host.

Request Parameters

Parameter	Description
serviceId	Service ID of the Endpoint Server to be connected.

Sample Request

```
$ curl 'https://api.netwitness.local/rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/isolation?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8' -i -X DELETE \
    -H 'Accept: application/json;charset=UTF-8' \
    -H 'NetWitness-Token: eyJ...AT' \
    -H 'Content-Type: application/json' \
    -d '{"allowDnsOnlyBySystem":false,"exclusions":null,"comment":"Release from
isolation"}'
```

HTTP request

```
DELETE /rest/api/host/B27DDED7-6FFA-A9D3-6577-
3DDE367B2820/isolation?serviceId=5f6e15bb-9384-4666-aa71-e5d134afd7d8 HTTP/1.1
Accept: application/json;charset=UTF-8
NetWitness-Token: eyJ...AT
Content-Type: application/json
Host: api.netwitness.local
Content-Length: 83

{"allowDnsOnlyBySystem":false,"exclusions":null,"comment":"Release from isolation"}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 12 Dec 2022 09:13:06 GMT
Keep-Alive: timeout=60
Connection: keep-alive
```

```

```

NetWitness® Platform XDR

Version 12.0

Archer Integration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Archer Integration **4**

Configure NetWitness to Work With Archer **5**

Create Archer User Accounts for Push and Pull 5

Integrate NetWitness Platform With Archer Cyber Incident & Breach Response 6

Unified Collector Framework 7

Configure Respond for Integration with Archer Cyber Incident & Breach Response 7

Configure Endpoints in Unified Collector Framework 10

Configure Reporting Engine for Integration with Archer Cyber Incident & Breach Response 12

Configure ESA Correlation for Integration with Archer Cyber Incident & Breach Response 14

Archer Feeds 17

Manage Unified Collector Framework **21**

Troubleshoot Archer Integration **22**

Archer Integration

Administrators can integrate NetWitness with Archer Cyber Incident & Breach Response to send alerts and incidents from NetWitness to Archer for incident management and remediation. This guide provides a high-level workflow for configuring this integration.

Note: When you upgrade from Security Analytics 10.6.5 to NetWitness 11.x, the Archer Cyber Incident & Breach Response integration is no valid, and must be re-configured.

The following table list the NetWitness 11.x integration options with Archer Cyber Incident & Breach Response Version 1.3.1.2.

Archer Cyber Incident & Breach ResponseVersion	NetWitness 11.x Integration	Reference
1.3.1.2	ESA Correlation	See "Configure ESA Correlation for Integration with Archer Cyber Incident & Breach Response" section.
1.3.1.2	Reporting Engine (RE)	See "Configure Reporting Engine for Integration with Archer Cyber Incident & Breach Response" section.
1.3.1.2	Respond	See "Configure Respond for Integration with Archer Cyber Incident & Breach Response 1.3.1.2" section.
1.3.1.2	Archer Feeds	See "Archer Feeds" section.

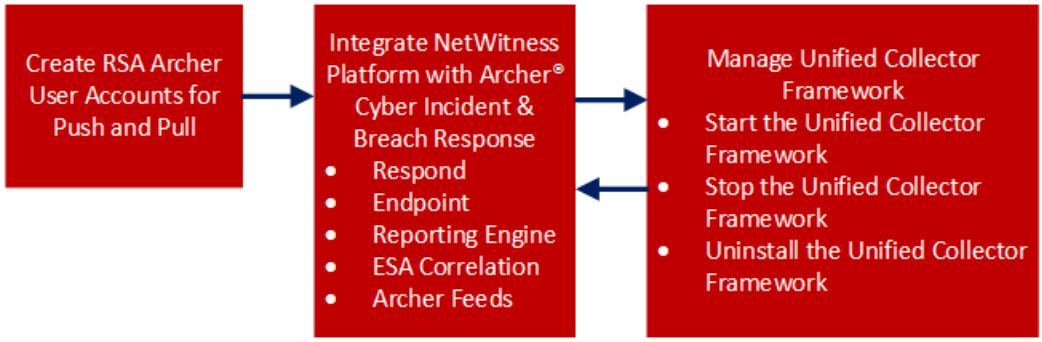
Configure NetWitness to Work With Archer

The Archer Cyber Incident & Breach Response solution enables you to aggregate all actionable security alerts, allowing you to become more effective, proactive, and targeted in your incident response and SOC management. For more information on Archer Cyber Incident & Breach Response capabilities, see Archer documentation on the Archer Community or on the Archer Exchange Community.

This version of Archer determines how NetWitness will be integrated. For supported Archer platforms, see the *SecOps Installation Guide*.

Archer Cyber Incident & Breach Response1.3.1.2 integrates with NetWitness using the UCF (Unified Collector Framework), which comprises of NetWitness Respond integration service and Archer Cyber Incident & Breach Response Watchdog service.

This figure represents the flow of NetWitness 11.x integration with Archer Cyber Incident & Breach Response 1.3.1.2.



Create Archer User Accounts for Push and Pull

You must create a user account for the web service client to transfer data to the Archer GRC Platform.

You require two Archer user accounts to avoid conflicts while sending and receiving data from NetWitness.

To create a user account for push and pull:

1. On the Archer UI, click **Administration** > **Access Control** > **Users** > **Add New**.
2. In the **First Name** and **Last Name** fields, enter a name that indicates that the Unified Collector Framework (UCF) uses this account to push data into Archer GRC. For example, UCF User, Push.

   Note: When configuring the Pull account, enter a name that indicates that the UCF uses this account to pull data from Archer GRC. For example, UCF User, Pull.

3. (Optional) Enter a user name for the new user account.

   Note: If you do not specify a user name, the Archer GRC Platform creates the user name from the first and last name entered when you save the new user account.

4. In the **Contact Information** panel, in the **Email** field, enter an email address to associate with the new user account.

5. In the **Localization** section, change the time zone to (UTC) Coordinated Universal Time.

   Note: The UCF uses UTC time to baseline all the time-related calculations.

6. In the **Account Maintenance** section, enter and confirm a new password for the new user account.

   Note: Make a note of the user name and password for the new user account that you created. You need to enter these credentials when you set up the UCF to communicate with the Archer GRC Platform through the web service client.

7. Clear the Force Password Change On **Next Sign-In** option.
8. In the **Security Parameter** field, select the security parameter that you want to use for this user.

   Note: If you assign a default security parameter with a password change interval of 90 days, you also must update the user account password stored in the SA IM integration service every 90 days. To avoid this, you can optionally create a new security parameter for the SA IM integration service user account, and set the password change interval to the maximum value allowed by your corporate standards.

9. Click the **Groups** tab, and perform the following:
   a. In the **Groups** panel, click **Lookup**.
   b. In the **Available Groups** window, expand Groups.
   c. Scroll down and select **SOC: Solution Administrator and EM: Read Only**.
   d. Click **OK**.
10. Click **Apply** and click **Save**.
11. If the machine language and regional settings of your Archer GRC system are set to anything other than English-US, perform the following:
    a. Open the user account you just created, and in the **Localization** section, in the **Locale** field, select **English (United States)**, and click **Save**.
    b. On the Windows system hosting your Archer GRC Platform, open **Internet Information Services (IIS) Manager**.
    c. Expand your Archer GRC site, click **.Net Globalization**, in both the **Culture** and **UI Culture** fields, select **English (United States)**, and click **Apply**.
    d. Restart your Archer GRC site.
12. Repeat steps 1 – 11 to create a second user account for the UCF to pull data from Archer GRC.

Integrate NetWitness Platform With Archer Cyber Incident & Breach Response

You have to configure the system integration settings to manage incident workflow in Archer Cyber Incident & Breach Response.

For information on how to configure system integration settings, see the "Manage Incidents in Archer Cyber Incident & Breach Response" in the *NetWitness Respond Configuration Guide*.

Unified Collector Framework

NetWitness integrates with Archer Cyber Incident & Breach Response 1.3.1.2 using the UCF. The UCF integrates with all supported SIEM tools and the Archer Cyber Incident & Breach Response solution. After you configure the system integration settings, all incidents are managed in Archer Cyber Incident & Breach Response instead of NetWitness Respond. Incidents created before the integration will not be managed in Archer Cyber Incident & Breach Response.

Note:
• You must configure the same option in both NetWitness and the Unified Collector Framework.
• Integration of the NetWitness Respond module with Reporting Engine or ESA Correlation can result in duplicate events, alerts, and incidents created in Archer Cyber Incident & Breach Response.

UCF supports multiple SIEM tools connections at the same time, such as supporting NetWitness Reporting Engine, HP ArcSight, and NetWitness Respond. However, different instances of the same SIEM tool are not supported, such as two NetWitness servers connected to the same UCF.

Prerequisites

- Install the Archer Cyber Incident & Breach Response package on Archer. See Archer documentation Archer Community or on the Content Tab at https://www.archerirm.community/t5/exchange-overviews/archer-exchange-offering-list/ta-p/672315.
- Install Archer Cyber Incident & Breach Response 1.3.1.2.
- Ensure you have NetWitness 11.1 as it is compatible with Archer Cyber Incident & Breach Response 1.3.1.2.
- Ensure that Respond is configured inNetWitness.

The UCF allows you to integrate your Archer Cyber Incident & Breach Response system with the following:

- NetWitness Respond
- NetWitness Reporting Engine
- NetWitness ESA Correlation.
- Archer Feeds

Configure Respond for Integration with Archer Cyber Incident & Breach Response

Step 1: Select the Mode for NetWitness Respond

1. Go to  **(Admin) > Services**, select the **Respond Server** service, and select  > **View** > **Explore**.

2. Navigate to `respond/integration/export`.
3. Set the `archer-sec-ops-integration-enabled` field to **true**.
4. Restart the Respond service by running the following command:
   `systemctl restart rsa-nw-respond-server`

Step 2: Configure NetWitness Respond to Forward Alerts to UCF

1. Navigate to `C:\Program Files\RSA\SA IM integration service\cert-tool\certs` in the SecOps Middleware box.
2. Copy both `keystore.crt.pem` and `rootcastore.crt.pem` from the certs folder (to the import folder of NetWitess server):
   ```
   cp rootcastore.crt.pem /etc/pki/nw/trust/import
   cp keystore.crt.pem /etc/pki/nw/trust/import
   ```

   Note: Before you copy the files from UCF to NetWitness Admin server, examine the files to remove any blank lines and save them.

3. SSH to NW-server box
   a. Run the update-admin-node command:
      `orchestration-cli-client --update-admin-node`
   b. Restart the RabbitMQ service:
      `systemctl restart rabbitmq-server`
   c. Restart the SMS service:
      `systemctl restart rsa-sms.service`

      Note: This step is mandatory to avoid receiving the "message bus down" error message which indicates that the EventSourceMessagePublisher has failed to reconnect to RabbitMQ on restart. This can cause some features such as deleting event sources to function improperly.

   d. Create user archer and set permissions for the virtual host '`/rsa/system`'
      ```
      rabbitmqctl add_user archer archer
      rabbitmqctl clear_password archer
      rabbitmqctl set_permissions -p /rsa/system archer ".*" ".*" ".*"
      ```

Step 3: Forward Alerts to the NetWitness Respond

- To forward NetWitness ESA Correlation alerts to the NetWitness Respond:
  a. Go to  **(Admin) > Services > ESA** service.
  b. Select an **ESA Correlation** service and select  **> View > Explore**.
  c. Navigate to the Correlation option in the left panel and select the Alert option.
  d. Edit the respond-enabled option to TRUE.
  e. Restart the Correlation server.

- To forward NetWitness Reporting Engine alerts to NetWitness Respond:
  a. Go to **(Admin) > Services > Reporting Engine** service.
  b. Select the **Reporting Engine** service, and then select **> View > Config**.
  c. Click the **General** tab.
  d. In the **System Configuration section**, select the **Forward Alerts to Respond** checkbox and click **Apply**.
- To forward NetWitness Malware Analysis alerts to NetWitness Respond:
  a. Go to **(Admin) > Services > Malware Analysis** service
  b. Select the **Malware Analysis** service, and select > **View** > **Config**.
  c. Click the **Auditing** tab.
  d. In the **Respond Alerting** panel, verify that the **Enabled Config Value** checkbox is selected. If the checkbox is not selected, select the checkbox, and click **Apply**.

Step 4: Forward Endpoint Alerts to the NetWitness Respond

You can forward Endpoint alerts to the Archer GRC through NetWitness Respond. For more information on how to Configure NetWitness Endpoint Alerts via Message Bus, see "Configure NetWitness Endpoint Alerts via Message Bus" in the *NetWitness Endpoint Integration Guide*.

Step 5: Aggregate Alerts into Incidents

The Respond Server service consumes alerts from the message bus and normalizes the data to a common format (while retaining the original data) to enable simpler rule processing. It periodically runs rules to aggregate multiple alerts into an incident and set some attributes of the Incident (for example, severity, category, and so on). For more information on aggregating alerts, see the "Configure Alert Sources to Display Alerts in Respond View" topic in the *NetWitness Respond Configuration Guide*.

To configure alert aggregation:

1. Go to **(Configure) > Incident Rules**.
2. To enable the rules provided out-of- the-box:
   a. Double-click the rule.
   b. Select **Enabled**.
   c. Click **Save**.
   d. Repeat steps a-c for each rule.
3. To add a new rule:
   a. Click **Create Rule**.
   b. Select **Enabled**.
   c. Enter the values in the following fields:

- Rule Name
- Action
- Match Conditions
- Grouping Options
- Incident Options
- Priority
- Notifications

4. Click **Save**.

Configure Endpoints in Unified Collector Framework

Endpoints provide the connection details required for the UCF to reach both your NetWitness and Archer GRC systems.

Note: Some endpoints are necessary to use different integrations. The following list shows the mandatory endpoints.

Mandatory Endpoint Integration

- Archer Push endpoint
- Archer Pull endpoint
- Mode selection: SecOps or Non SecOps mode.

Note:
• If Non SecOps mode is selected, incidents are managed in NetWitness Respond instead of Archer Cyber Incident & Breach Response.
• You must configure the port depending on the protocol (TCP, UDP, or secure TCP).
• Make sure the certificate subject name for your Archer GRC server matches the hostname.

Procedure

1. On the UCF system, open the Connection Manager, as follows:
   a. Open a command prompt.
   b. Change directories to `<install_dir>\SA IM integration service\data-collector`.
   c. Enter `runConnectionManager.bat`.
2. In the **Connection Manager**, enter **1** for Add Endpoint.
3. Add an endpoint for pushing data to Archer Cyber Incident & Breach Response, as follows:
   a. Enter the number for Archer.

      Note: Enable SSL to add the Archer endpoints.

   b. For the endpoint name, enter **push**.

c. Enter the URL of your Archer GRC system.

d. Enter the instance name of your Archer GRC system.

e. Enter the user name of the user account you created to push data into your Archer GRC system.

f. Enter the password for the user account you created to push data into your Archer GRC system, and confirm the password.

g. When prompted if this account is used for pulling data, enter **False**.

4. Add an endpoint for pulling data from Archer Cyber Incident & Breach Response, as follows:

a. Enter the number for Archer.

> **Note:** SSL must be enabled to add the Archer endpoints.

b. For the endpoint name, enter **pull**.

c. Enter the URL of your Archer GRC system.

d. Enter the instance name of your Archer GRC system.

e. Enter the user name of the user account you created to pull data from your Archer GRC system.

f. Enter the password for the user account you created to pull data from your Archer system, and confirm the password.

g. When prompted if this account is used for pulling data, enter **True**.

5. Add an endpoint for NetWitness:

- For Respond

a. Enter the number for NetWitness IM.

b. Enter a name for the endpoint.

c. Enter the SA Host IP address.

d. For SA Messaging Port, enter **5671**.

e. Enter the target queue for remediation tasks. Selecting All processes both the Archer Integration (GRC) and IT Helpdesk (Operations).

f. When prompted to automatically add certificates to the SA trust store, enter **No**.
The certificates are added manually in previous steps.

g. In UCF connection manager, select the mode, as follows:

i. Enter the number for Mode Selection.

ii. Select Manage incident workflow exclusively in Archer Cyber Incident & Breach Response from the drop-down.

> **Note:** Make sure you select the second option as the first option is not supported in NetWitness Platform 11.x release.

- For Reporting Engine and ESA Correlation
  a. To use third-party integrations, add the Syslog Server Endpoint, as follows:
     i. Enter the number for Syslog Server Endpoint.
     ii. Enter the following:
     - User defined name
     - SSL Configured TCP port number

       Note: Defaults to 1515. If you do not want to host the Syslog server in this mode, enter **0**.
     - TCP port number - Enter the TCP port if the Syslog client sends the Syslog message in TCP mode.

       Note: Defaults to 1514. If you do not want to host the Syslog server in this mode, enter **0**.
     - UDP port number - Enter the UDP port if the Syslog client sends the Syslog message in UDP mode.

       Note: Defaults to 514. If you do not want to host the Syslog server in this mode, enter **0**.

     By default, the Syslog server runs in the above three modes, unless it is disabled by entering **0**.
  b. To test the Syslog client, enter the number for Test Syslog Client. Use the Test Syslog client with the files from `<install_dir>\SA IM integration service\config\mapping\test-files\`.

6. In the Connection Manager, enter **5** to test each endpoint.

Configure Reporting Engine for Integration with Archer Cyber Incident & Breach Response

To configure Syslog Output Action for the Reporting Engine:

1. Go to  **(Admin) > Services**.
2. Select the **Reporting Engine** service, and click  **View** > **Config**.
3. Click the **Output Actions** tab.
4. In the **NetWitness Platform Configuration** panel, in the **Host Name** field, enter the host name or IP address of the Reporting Engine server.
5. In the **Syslog Configuration** section, add the Syslog Configuration as follows:

a. In the **Server Name** field, enter the host name of the UCF.

b. In the **Server Port** field, enter the port that you selected in the UCF Syslog configuration.

c. In the **Protocol** field, select the transport protocol.

Note: Configure SSL if you select Secure TCP.

6. Click **Save**.

To configure NetWitness Reporting Engine SSL for Secure Syslog Server:

1. Copy the certificate `keystore.crt.der` from the UCF machine to NetWitness server box at `/usr/lib/jvm/java-1.8.0-openjdk-1.8.0.161-0.b14.el7_4.x86_64/jre/lib/security`.
2. Run the following command:

```
keytool -import -file keystore.crt.der -alias ucf-syslog -keystore /etc/pki/nw/trust/truststore.jks -storepass changeit
```

Note: Do not copy and paste the above command. Type the command to avoid errors.

3. Enable **ServerCertificateValidationEnabled** to **true:**
   - Go to  **(Admin) > Services**.
   - Select  **> View > Explore** of the **Reporting Engine** service .
   - Expand **com.rsa.soc.re** > **Configuration** > **SSLContextConfiguration**.
   - Expand **SSLContextConfiguration** and set **ServerCertificateValidationEnabled** to **true**.
4. Restart the Reporting Engine service by running the following command:
   `service rsasoc_re restart`

To configure rules in NetWitness Platform:

1. Go to **Reports > Manage**.
   The Manage tab is displayed.
2. In the **Rules > Groups** panel, click **+**.
3. Enter a name for the new group.
4. Select the group you created, and in the Rule toolbar, click **+**.
   1. In the **Rule Type** field, select **NetWitness Platform DB**.
5. Enter a name for the rule.
6. Enter values in the **Select** and **Where** fields based on the rule that you want to create.

Note: Add the Syslog configuration with the Syslog name set above.

7. Click **Save**.

Note: To see the same number of alerts in the Reporting Engine and Archer GRC, make sure that you have selected **Once** for execute in both the Syslog and Record tabs.

To add Alert Templates for the Reporting Engine in NetWitness Platform:

The UCF syslog configuration contains out-of-the-box alert templates that you can use to create an alert with a syslog output action. These templates define the criteria used to aggregate alerts into incidents in your Archer GRC Platform.

The sample templates are located in the following location on the UCF system:

`<install_dir>\SA IM integration service\config\mapping\templates\SecOps_SA_Templates`

1. Go to **Reports > Manage > Alerts**.
2. Select the **Template** tab and click +.

Note: After you copy the template in the Create/Modify Template window, make sure to replace `cs25=${sa.host} cs25Label=sahost` to `cs25=${nw.host} cs25Label=nwhost`.

3. In the **Name** field, enter a name for the alert template.
4. In the **Message** field, enter the alert message.
5. Click **Create**.
6. Repeat steps 3 to 6 for each alert template that you want to add.

To Configure Alerts in NetWitness Platform:

In NetWitness Reporting Engine, an alert is a rule that you can schedule to run on a continuous basis and log its findings to several different alerting outputs.

1. Go to **Reports > Manage > Alerts**.
2. Click +, select **Enable**, then Select the rule you created.
3. Select **Push to Decoders**.

   Note: If you do not enter a value in this field, the link in the Archer Security Alerts application to NetWitness does not work.

4. From the Data Sources list, select your data source.
5. In the **Notification** section, select **Syslog**.
6. Click + and complete the Syslog configuration fields.
7. In the **Body Template** field, select the template that you want to use for this Syslog alert.
8. Click **Save**.

Configure ESA Correlation for Integration with Archer Cyber Incident & Breach Response

To configure ESA Correlation Syslog Notification Settings in NetWitness:

1. Go to  **(Admin) > System > Global Notifications**.
2. Click the **Output** tab.
3. Define and enable an ESA Correlation Syslog notification.
4. Click the **Servers** tab.
5. Define and enable a Syslog notification server.
6. In the Syslog Server Configuration section, enter the following:

   Field Description:

   - Name - Specify the custom name.
   - Server IP (Hostname) - Specify the hostname or IP Address of the system on which you installed the UCF.
   - Port - Specify the port number on which you want the UCF to listen.
   - Facility - Specify the Syslog facility.
   - Protocol - Select the protocol.
7. Click **Save**.

To configure NetWitness Platform ESA Correlation SSL for Secure Syslog Server:

If the Syslog server is configured with Secure TCP, configure the SSL.

1. Copy the `keystore.cert.pem` from UCF machine to NetWitness server.
2. Run the following command:

   ```
   security-cli-client --add-trusts --service integration-server --superuser-id admin --superuser-pwd netwitness --chain-file <path to keystore.cert.pem file>
   ```
3. Restart the integration server by running the following command:

   ```
   service rsa-nw-integration-server restart
   ```

To Add ESA Correlation Alert Templates

The UCF syslog configuration contains out-of-the-box alert templates that you can use to create an alert with a syslog output action. These templates define the criteria used to aggregate alerts into incidents in your Archer GRC Platform.

The sample templates are located in the following location on the UCF system:

`<install_dir>\SA IM integration service\config\mapping\templates\SecOps_SA_Templates\SecOps_SA_ESA_templates.txt`

1. Select  **(Admin) > System > Global Notifications**.
2. Click the **Templates** tab and click .
3. In the **Template Type** field, select ESA Correlation.
4. In the **Name** field, enter the name for the template.
5. 6. (Optional) In the **Description** field, enter a brief description for the template.

6. 7. In the **Template** field, enter the alert message.
7. 8. Click **Save**.
8. 9. Repeat steps 3 – 8 for each alert template that you want to add.

To Create ESA Correlation Rules

1. Click **(Configure) > ESA Rules**.
2. In the **Rule Library**, click +.
3. Select **Rule Builder**.
4. In the **Rule Name** field, enter a name for the rule.
5. In the **Description** field, enter a description for the rule.
6. Select the **Severity**.
7. In the **Condition** panel:
   a. Click + to build a statement.
   b. Enter a name, select a condition type, and add meta data/value pairs for the statement.
   c. Click **Save**.
   d. Repeat steps a – c until you have built all the statements for the rule.
8. In the **Notifications** section, select **Syslog**.
9. Select the notification, Syslog server, and template that were created previously.
10. Click **Save** and click **Close**.
11. In the Deployments section, click .
12. Click + and enter a deployment name.
13. Click + for ESA Correlation services section.
14. Select the ESA Correlation Service.
15. In the Data Source section, click + and select a data source.
16. Click **Save**.
17. In the **ESA Correlation Rules** section, click + to select the ESA Correlation Rule that you created.
18. Click **Save**.
19. Click **Deploy Now**.

Archer Feeds

By default, only the IP address and Criticality Rating fields in the Archer Devices application are fed into NetWitness by the SA IM Integration Service. You can customize the Enterprise Management plug-in to include the Business Unit and Facility fields that are cross-referenced in the Devices application in the feed. For more details, see Archer documentation at https://www.archerirm.community/t5/exchange-overviews/archer-exchange-offering-list/ta-p/672315.

Note: If you want to feed Business Unit and Facility information from your Archer GRC Platform into Live, you must also add keys for these fields in the index-concentrator-custom.xml file.

Update the Concentrator and Decoder Services

The SA IM Integration Service in Archer Cyber Incident & Breach Response manages the files for a custom feed and deposits these files in a local folder that you specify when you configure the Enterprise Management Endpoint. The Live module of NetWitness retrieves the feed files from this folder. Live then pushes the feed to the Decoders, which start creating metadata based on the captured network traffic and the feed definition. To enable the Concentrator to detect a new metadata created by the Decoders, make sure to edit the `index-concentrator-custom.xml`, `index-logdecoder-custom.xml`, and `index-decoder-custom.xml` files.

1. Select  **(Admin) > Services**.
2. Select the **Concentrator** and select  > **View > Config**.
3. Click the **Files** tab.
4. From the drop-down list, select the `index-concentrator-custom.xml` file. Do one of the following:
   - If content already exists in the file, add a key for the new metadata element:

     ```
     <key description="Criticality" format="Text" level="IndexValues"
     name="criticality" defaultAction="Open"/>
     ```

     Note: Do not copy and paste above command . Type the command to avoid errors.
   - If the file is blank, add the following content:

     ```
     <?xml version="1.0" encoding="utf-8"?>
     <language level="IndexNone" defaultAction="Auto">
     <key description="Criticality" format="Text" level="IndexValues"
     name="criticality" defaultAction="Open"/>
     </language>
     ```
5. Click **Apply**.
6. To add multiple devices:
   a. Click **Push**.
   b. Select the devices to which you want to push this file.
   c. Click **OK**.

7. Repeat steps 1 to 6 for the Log Decoders and Index Decoders, using index-logdecoder-custom.xml and index-decoder-custom.xml.
8. Restart the Concentrator and Decoder services by running the following commands:
```
service nwdecoder restart
service nwconcentrator restart
```

Add the Archer Enterprise Management Endpoint in UCF

1. In the UCF connection manager, select the mode:
   a. Enter the number for Mode Selection.
   b. Select one of the following options:
      - Manage incident workflow in NetWitness.
      - Manage incident workflow exclusively in Archer Cyber Incident & Breach Response.
2. Add the Archer Enterprise Management Endpoint:
   a. Enter the number for Enterprise Management.
   b. Enter the values in the fields as described in the table below.

Field	Description
Endpoint Name	Optional endpoint name.
Web Server Port	Defaults to 9090. You can configure this to host the web server url by providing the URL with the port number as in the NetWitness live feed: http://hostname:port/archer/sa/feed.
Criticality	Criticality of the assets to be pulled from Archer GRC. If **false**, pull assets with any criticality. If **true**, pull assets with only high criticality. To configure this manually, edit the em.criticality property in the collector-config properties file to provide a comma-separated list of criticalities: LOW, MEDIUM, HIGH.
Feed Directory	Directory where the assets CSV file from Archer GRC are saved. **Note:** The directory path provided must exist.
Web Server Username	Username for authenticating to the EM web server.
Web Server Password	Password for authenticating to the EM web server.

Field	Description
SSL Mode	Defaults to No. If **No**, the URL uses `http mode: http://hostname:port/archer/sa/feed` If you have not updated the host file, see "Update the NetWitness Host File" section. **Note:** NetWitness currently does not support Archer recurring feeds in SSL mode.

Update the NetWitness Platform Host File

1. Edit the host file on the NetWitness server at the following location: `vi /etc/hosts`.
2. Enter the following for the UCF host IP address:

   `<ucf-host-ip> <ucf-host-name>`
3. Restart NetWitness server by running the following command:

   `service jetty restart`
4. While configuring the NetWitness live feed, enter the host name for the URL instead of the IP address and the port number configured for Enterprise Management endpoint in the UCF:

   `http: //<ucf-host-name> : <EM_Port>/archer/sa/feed.`
5. Verify that the connection works.

Create a Recurring Feed Task

For NetWitness to download feed files from the NetWitness Respond Integration Service and push the feeds to Decoders, you must create a recurring feed task and define the feed settings.

Note: For Archer Cyber Incident & Breach Response 1.2: For NetWitness to download feed files from the UCF machine and push the feeds to Decoders, you must create a recurring feed task and define the feed settings. The procedure is similar to Archer Cyber Incident & Breach Response 1.3, with a few exceptions. See documentation on the Archer Exchange Community for details.

1. Select  **(Configure) > Custom Feeds**.
2. In the **Feeds** view, Click ✚.
3. Select **Custom Feed**, and click **Next**.
4. Select **Recurring**.
5. Enter a name for the feed.
6. In the URL field, enter the following:

   `http://ucf_hostname/archer/sa/feed`

   where, *`http :ucf_hostname_or_ip:port`* is the address of the NetWitness Respond Integration Service system. For example: `http://<ucf-host-name>` .
7. Select **Authenticated**.

8. In the **User Name** and **Password** fields, enter the credentials of the user account you created in the Add the Archer Enterprise Management Endpoint in UCF procedure.
9. Define the recurrence interval for the feed.
10. In the **Date Range** panel, define a start and end date for the feed, and click **Next.**
11. Select each Decoder to which you want to push this feed, and click **Next**.
12. In the **Type** field, amke sure that IP is selected.
13. In the **Index Column** field, select 1.
14. In the second column, set the Key value to criticality, and click **Next**.
15. Review your feed configuration details and click **Finish**.

Manage Unified Collector Framework

This section provides additional tasks for configuring and managing the UCF for Archer Cyber Incident & Breach Response 11.3.1.2 Integration.

Start the Unified Collector Framework

1. Click **Control Panel > Administrative Tools > Services**.
2. Select **RSA Unified Collector Framework**.
3. Click **Start**.

Stop the Unified Collector Framework

1. Click **Control Panel > Administrative Tools > Services**.
2. Stop the Archer Cyber Incident & Breach Response WatchDog Service.

   Note: If you do not stop the Watchdog service, the Watchdog service starts the NetWitness Respond Service.

3. Select **RSA Unified Collector Framework**.
4. Click **Stop**.

Note: If the service takes too long to shutdown, use the Task Manager to end the RSASAIMDCService.

Uninstall the Unified Collector Framework

1. Click **Control Panel** > **Programs and Features**.
2. Select **RSA Unified Collector Framework**.
3. Click **Uninstall**.

Troubleshoot Archer Integration

This section provides resolutions to common problems that you may encounter while configuring Archer Cyber Incident & Breach Response 1.3.1.2 with NetWitness Respond.

Problem	Solutions
After adding the endpoint for NetWitness Respond, the Certificate Authority truststore fails to set. **Resolution**	1. Make sure that the SSH credentials for the NetWitness host are valid. 2. If the credentials are correct, but the error still occurs, manually copy certificates.
Remediation Tasks being pushed to the operations queue through the UCF are not appearing in Archer Cyber Incident & Breach Response as Findings.	1. Open the Connection Manager using the command prompt: • `Change directories to <install_dir>\SA IM integration service\data-collector.` • `Type: runConnectionManager.bat` 2. Enter **2** to edit endpoint. 3. Enter **3** to NetWitness Respond. 4. Make sure the Target Queue is set to **All** or **Operations**.
In the `<install_dir>\SA IM integration service\logs\collector.log`, there are SSL errors between NetWitness and RSA Unified Collector Framework.	1. Verify that the SSL certificates are valid. **Note:** NetWitness Respond certificates are valid for two years. 2. If your certificates are expired, regenerate and copy the expired certificates. **To regenerate and copy the certificates:** 1. In the Command Prompt, go to `<install_dir>\SA IM integration service\data-collector.` 2. Enter `runConnectionManager.bat` 3. Enter the number for Regenerate NetWitness RespondIntegration Service Certificate.

Problem	Solutions
NetWitness unable to forward incidents to UCF.	4. In the NetWitness Respond endpoint, in Connection Manager, enter the number for Edit Endpoint. 5. Enter **Yes** to copy the certificates automatically to the NetWitness trust store. **Note:** If certificates fail to copy, manually copy the certificates.
	1. In the collector config (C:\PROGRAM FILES\RSA\SA IM INTEGRATION SERVICE\CONFIG\collector-config), change the following: im.virtualhost=/rsa/im/integration to im.virtualhost=/rsa/system 2. Restart UCF. For more information on restarting UCF, see Start the Unified Collector Framework. 3. In the data collector (C:\PROGRAM FILES\RSA\SA IM INTEGRATION SERVICE\data-collector), double click on the following file to run it. runConnectionmanager.bat

NetWitness® Platform XDR

Version 12.0

Archiver Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Archiver Overview **5**

Configuring an Archiver **6**

Prerequisites 6

Workflow 6

Edit the Archiver Service 8

Add Log Decoder as a Data Source to Archiver 10

Add Log Decoder as a Data Source to Archiver 10

Archiver Meta Settings Considerations 11

(Optional) Configure Meta Filters for Aggregation 11

(Optional) Add Index Entries for Archiver Reporting 13

Configure Archiver Storage and Log Retention 15

Configure Hot, Warm, and Cold Storage 17

Hot, Warm, and Cold Tier Storage Features 20

Collections 21

Collections Features 21

Retention Rules 23

Collection Dialog 25

Rule Definition Dialog 26

Configure Log Storage Collections 28

Prerequisites 28

Configure a Log Storage Collection 28

Define Retention Rules 31

Add Archiver as a Data Source to Reporting Engine 34

Configure Archiver Monitoring 36

Additional Archiver Configuration **37**

Configuring Data Backup and Restore 38

Add Archiver Service 38

Prerequisites 38

Procedure 38

Create Collection 40

Prerequisites 40

Procedure 40

Add Archiver Service as a Data Source to Reporting Engine 42

Prerequisites 42

Procedure 42

Mount Archiver Directories 44
Create a Collection 45
Delete a Collection 46
Example Procedure: How to Restore a Collection for Reporting and Investigation 47
Investigate a Collection 48
View Archiver Collection Statistics 49
View Archiver Logs 49
Add Archiver Service as a Data Source to Broker 50
Prerequisites 50
Procedure 50
Retrieve Hash Information 53
References 58
Archiver Collection Dialog 59
Archiver Services Config View - General Tab 62
Aggregate Services Section 64
System Configuration Section 65
Aggregation Configuration Section 66
Archiver Service Configuration 67
Data Retention Tab - Archiver 69
Total Hot, Warm, and Cold Storage 71
Hot, Warm, and Cold Storage Mount Points Dialogs 71
Services Config View - Archiver 72
General 74
Aggregate Services 75
System Configuration 75
Aggregation Configuration 76
Aggregation Settings 76
Service Heartbeat 77
Files 77

Archiver Overview

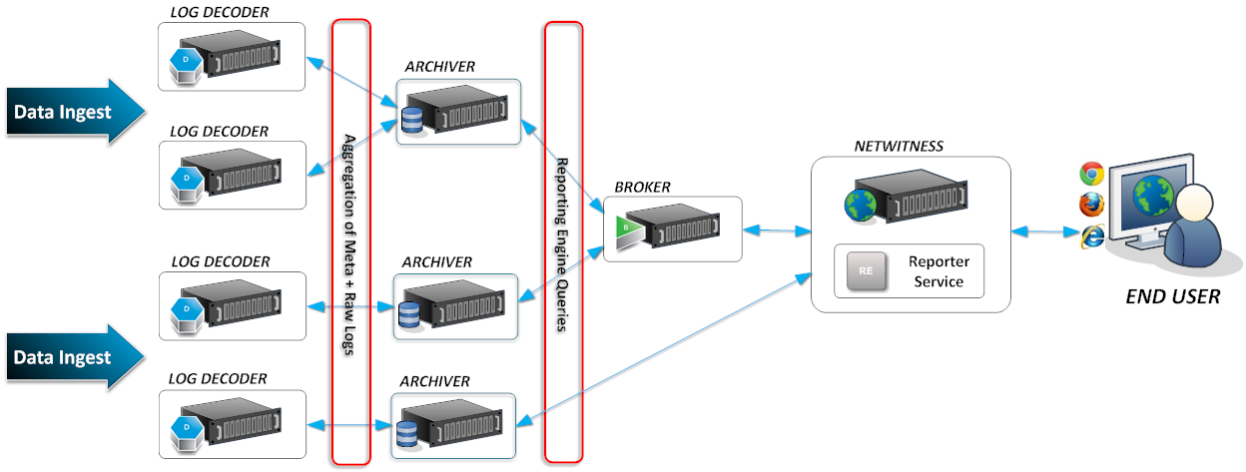
This guide provides detailed instructions on how to configure Archiver in your network, additional procedures that are used at other times, and reference materials that describe the user interface for configuring Archiver in your network.

The NetWitness Archiver is an appliance that enables long-term log archiving by indexing and compressing log data and sending it to Archiving storage. The Archiving storage is then optimized for long-term data retention and compliance reporting.

Archiver stores raw logs and log meta from Log Decoders for long-term retention and it uses Direct-Attached Capacity (DAC) for storage.

Note: Raw packet and packet meta are not stored in the Archiver.

The following figure depicts the architecture of a NetWitness network that implements the Archiver.



Configuring an Archiver

The NetWitness Archiver is an appliance that enables long-term log archiving by indexing and compressing log data and sending it to Archiving storage. The Archiving storage is then optimized for long-term data retention and compliance reporting.

Archiver stores raw logs and log meta from Log Decoders for long-term retention and it uses Direct-Attached Capacity (DAC) for storage.

Note: Raw packet and packet meta are not stored in the Archiver.

Prerequisites

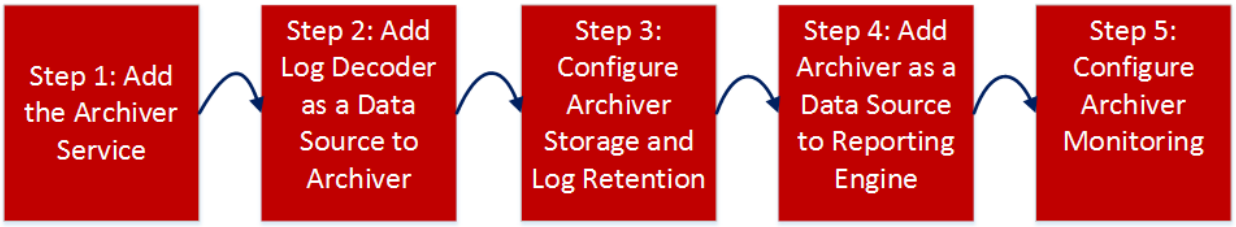
Ensure that you have:

- Installed the Archiver host in your network environment.
- Installed and configured Log Decoder version 11.2 in your network environment.

If you want to configure multiple Archiver or Concentrator services as a group and share the aggregation tasks between them, refer to **Group Aggregation** in the *Deployment Guide*.

Workflow

This workflow illustrates the end-to-end installation and configuration process for an Archiver.



The following table describes the basic steps for configuring an Archiver.The tasks must be completed in the sequence they are given.

Configuration Step	Description
Edit the Archiver Service	Provides information on how to add an Archiver service to the Archiver host and apply a license to it.
Add Log Decoder as a Data Source to Archiver	Provides instructions on how to add a Log Decoder to an Archiver.
Configure Archiver Storage and Log Retention	Provides instructions on how to configure storage and log retention on an Archiver.

Configuration Step	Description
Add Archiver as a Data Source to Reporting Engine	Provides instructions on how to add an Archiver as a data source to Reporting Engine to generate reports for the data collected by an Archiver.
Configure Archiver Monitoring	Provides instructions on how to configure the alert mechanism related to Archiver storage.

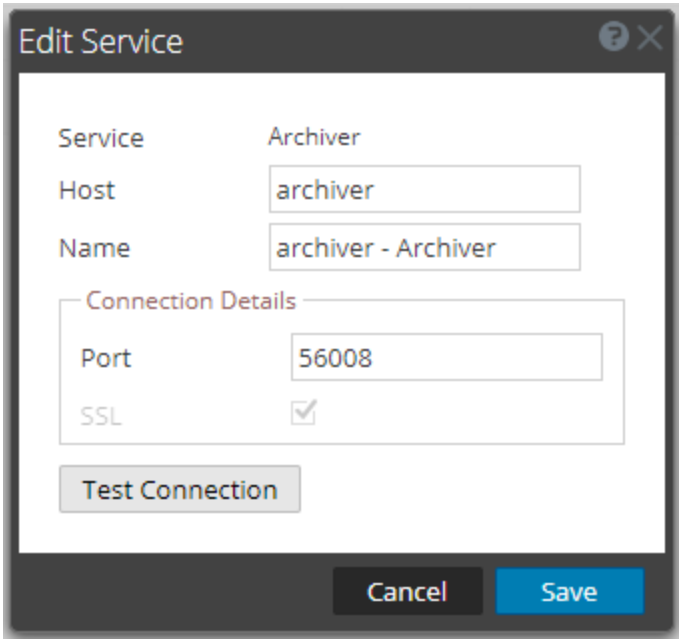
Edit the Archiver Service

To add an Archiver service, ensure that you have installed an Archiver host on which you want to run the Archiver service. See "Step 1: Add or Update Host" topic in the *Host and Services Getting Started Guide* for the procedure that explains how to add a host.

Perform the following steps to edit the Archiver service:

1. Go to **(Admin) > Services**.
   The Services panel is displayed.

2. Select an Archiver service and click in the **Services** panel toolbar.
   The Edit Service dialog is displayed.



3. Provide the following details.

Field	Description
Host	Select a host from the drop-down menu.
Name	Type a name for the service.
Port	Default port is 50008.
SSL	Select **SSL** if you want NetWitness to communicate with the service using SSL. The security of data transmission is managed by encrypting information and providing authentication with SSL certificates. **Note:** If you select SSL, ensure SSL is enabled in the System Configuration panel.

Field	Description
Username	(Optional) Type the username for the service.
Password	(Optional) Type the password for the service.
Entitle Service	Select if you want to apply the entitlements currently configured to this service. For more information, see "Entitlement Capability Implementation" topic in the *Licensing Guide*.

4. Click **Test Connection** to determine if NetWitness connects to the service.
5. When the result is successful, click **Save**.

   The edited service is now displayed in the services panel.

   Note: If the test is unsuccessful, edit the service information and retry.

6. Apply license to the Archiver service.

   Refer to the "Synchronize NetWitness Server**"** topic in the *Licensing Guide* for details on the procedure to activate (apply a license to) the Archiver service.

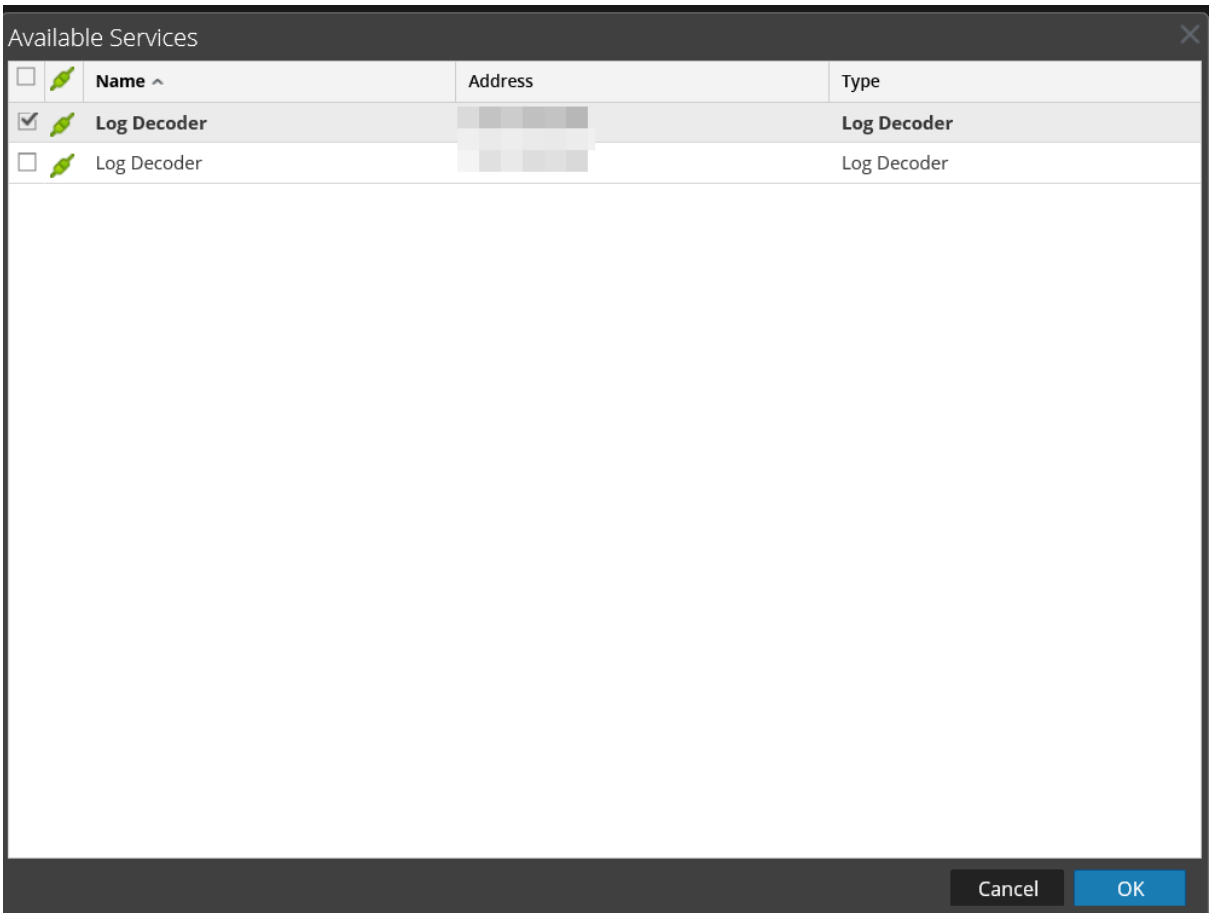
Add Log Decoder as a Data Source to Archiver

To add a Log Decoder as a data source to Archiver, you need to have installed the Archiver host in your network environment, installed and configured a Log Decoder in your network environment, and added the Archiver host to NetWitness and make sure the Archiver service shows as active and licensed.

Add Log Decoder as a Data Source to Archiver

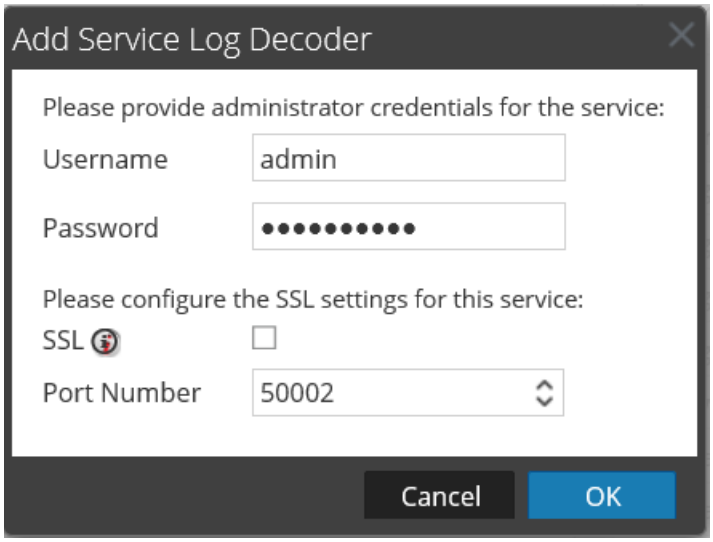
To add a Log Decoder as a data source to an Archiver:

1. Go to **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, click **> View > Config**.
   The Services Config view of Archiver is displayed.
3. On the **General** tab, in the **Aggregate Services** panel, click .
   The Available Services dialog is displayed.



4. Select the Log Decoder service to add as a data source to the Archiver and click **OK**.

5. If the Log Decoder is using the trust model, an Add Service dialog is displayed.



6. Type the username and password for the Log Decoder, and configure the SSL settings.
7. Click **OK**.
   The selected Log Decoder service is listed in the **Aggregate Services** panel.

Archiver Meta Settings Considerations

To maximize retention time, the meta items and index of the Archiver have been reduced (when compared to the Concentrator) to support common reporting needs. This means that, by default, you may not be able to run all of the reports you run on the Concentrator on the Archiver. You can view a list of the current meta and index items used by the Archiver in the following locations:

- **Explorer view**: The **/archiver/devices/<logdecoder>/config/options** path in the **metaInclude** field shows the current list of meta items.
- **Config view > Files tab**: The **index-archiver.xml** shows the default index configuration. The **index-archiver-custom.xml** shows any modifications.

The meta items and index of the Archiver can be customized to support customer specific reporting needs, however this will require additional storage, CPU resources, and Memory resources to support, and may impact retention time. As more meta items are added to the Archiver, the maximum aggregation rate will decrease, and the time to execute reports will increase.

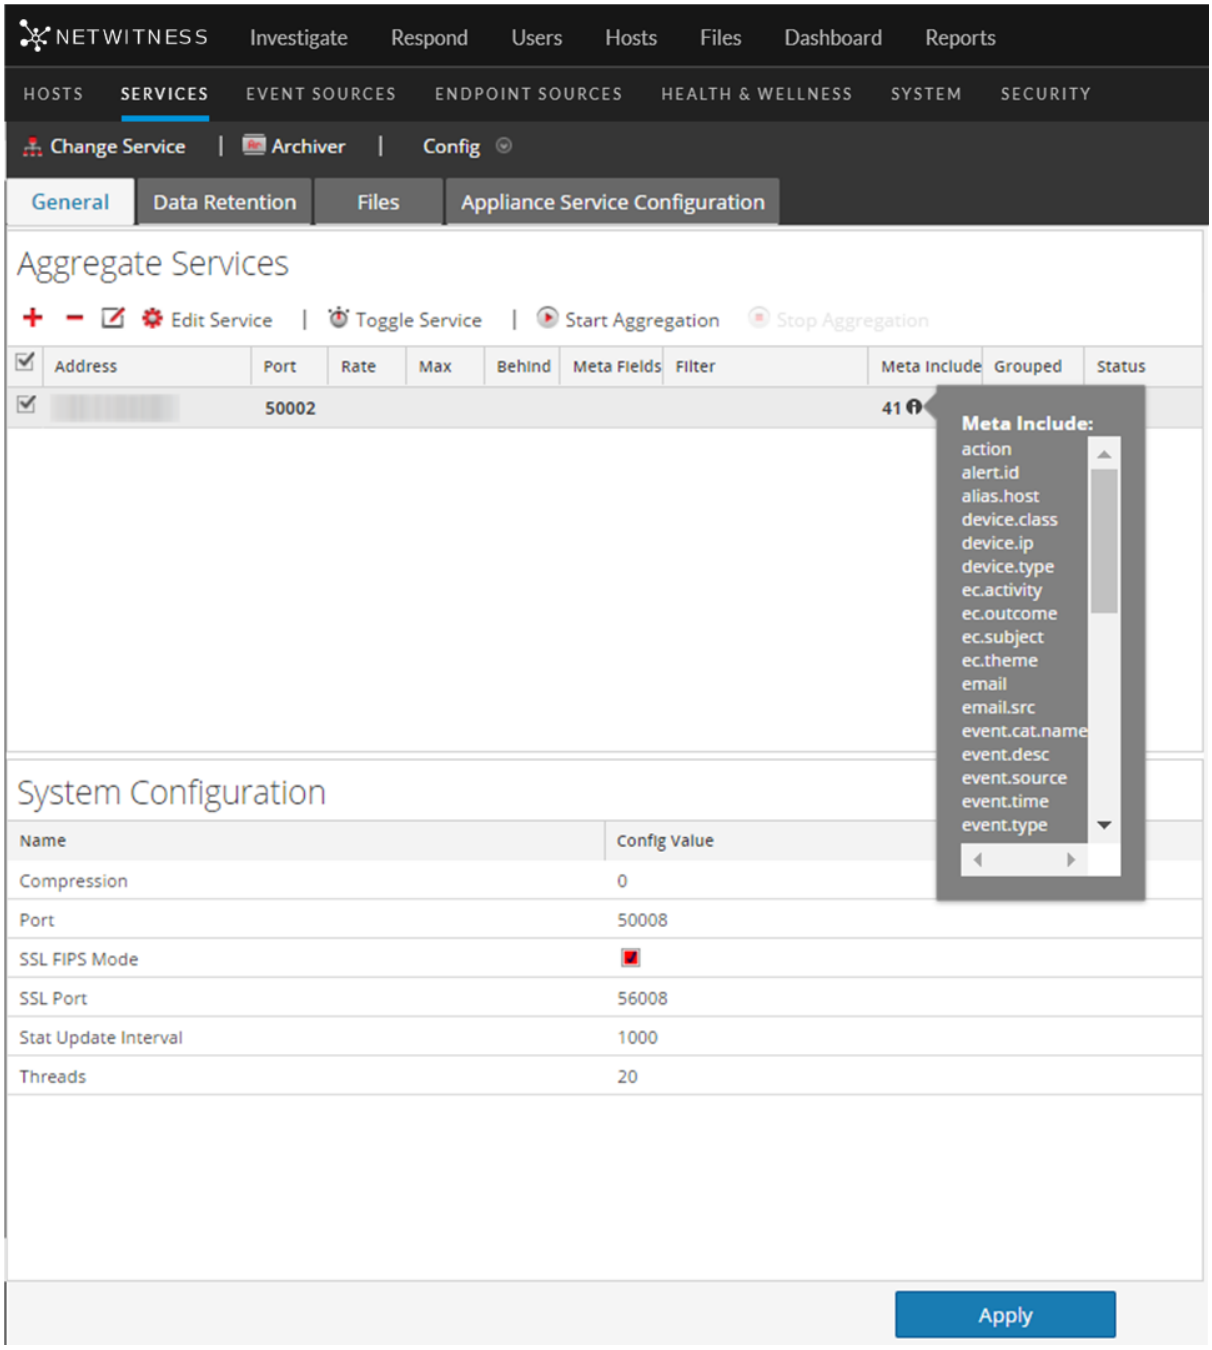
See (Optional) Configure Meta Filters for Aggregation and (Optional) Add Index Entries for Archiver Reporting for additional details.

(Optional) Configure Meta Filters for Aggregation

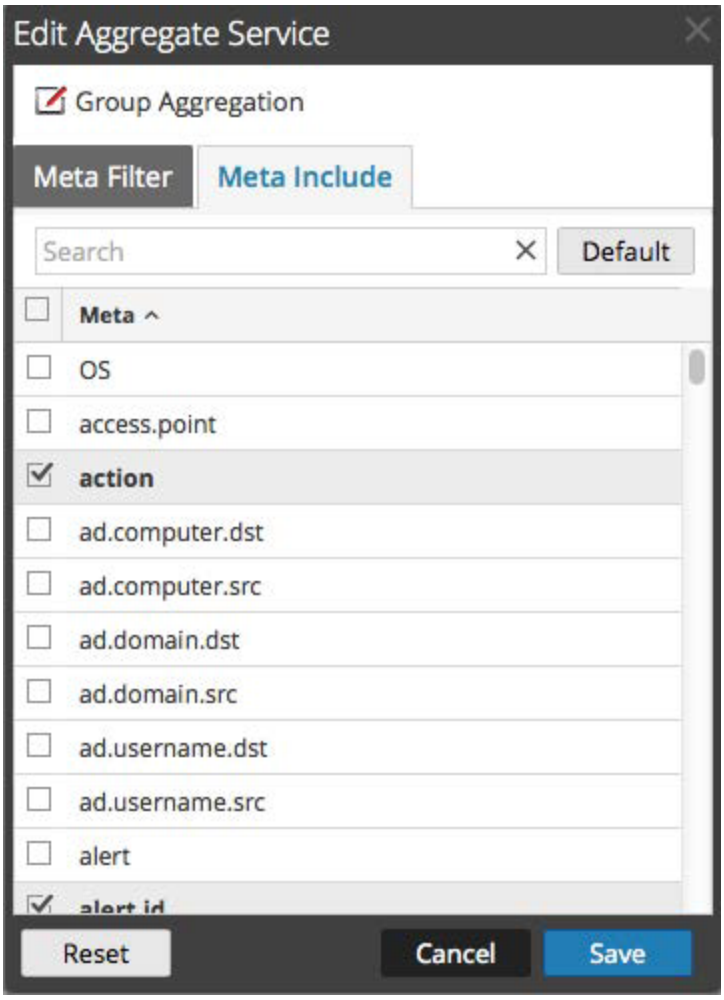
Follow this procedure to view and add additional meta items to the Archiver.

Caution: Adding meta or indexes will require additional storage, CPU resources, and Memory resources to support, and may impact retention time. As more meta items are added to the Archiver, the maximum aggregation rate will decrease, and the time to execute reports will increase.

1. To view the current meta items, in the **Aggregate Services** panel, select the Log Decoder service and click ⓘ in the **Meta Include** field.



2. To add additional meta items, select the Log Decoder service and click ✎.



3. In the Edit Aggregate Service dialog, select the meta items to include in the Meta Include list. For example, you may want to consider including ip.srcport, tcp.srcport, udp.srcport, msg, url, query, bytes, alias.host, ip.dst, ip.dstport, ip.src, tcp.dstport, megabytes, time, event.desc, and word.
4. Click **Save** and then click **Apply**.
5. See (Optional) Add Index Entries for Archiver Reporting below for information on how to index the additional meta keys.

(Optional) Add Index Entries for Archiver Reporting

Caution: Adding meta or indexes will require additional storage, CPU resources, and Memory resources to support, and may impact retention time. As more meta items are added to the Archiver, the maximum aggregation rate will decrease, and the time to execute reports will increase.

The Archiver's default index configuration only includes value indexes for these keys:

- time
- decoder source (did)

- destination user account (user.dst),
- alert ID (alert.id)
- device IP (device.ip)
- source IP address (ip.src)
- destination IP address (ip.dst)
- event description (event.desc)
- device class (device.class)
- medium
- object name (obj.name)
- word

For information on customizing this list, see "Index Customization" in the *Core Database Tuning Guide*.

Configure Archiver Storage and Log Retention

This topic provides instructions for Administrators to configure storage and log retention on an Archiver.

For compliance reasons, it is often necessary to retain some logs longer than other logs. Some logs are legally sensitive and cannot be retained for a long period of time. Other logs have a requirement to be retained for years. In addition to compliance, some logs are useful for historic forensics and other logs have little to no security or operationally relevant value and can be deleted after a short time.

Because business requirements vary, NetWitness enables you to configure Collections, which are log retention sets for storing log data. For each collection, you can specify how much of the total storage space to use and how many days to retain the logs in the collection. To specify the type of logs to put in the collection, you define retention rules to associate with the collections. Retention rules for all of your collections execute sequentially in an order that you define.

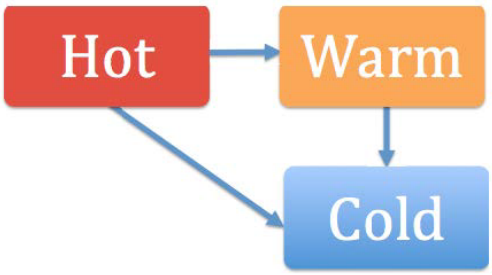
To do this, you must first define the total physical storage space for your collections. NetWitness enables you to define three types of storage:

- **Hot Tier Storage**: This storage contains log data that is in active use as part of the business process. Users can access these logs faster than other types of storage and they can use these logs for reporting and other tasks. Hot storage is usually Direct-Access Capacity (DAC) or SAN storage.
- **Warm Tier Storage**: (Optional) This storage contains older log data aggregated by Archiver. Log data access is slower than hot storage. Users can also use these logs for reporting and other tasks. Warm storage is usually Network Attached Storage (NAS).
- **Cold Tier Storage**: (Optional) This storage contains the oldest log data that is either required for the operation of the business or mandated by regulatory requirements. The logs are offline and Archiver cannot access these logs for reporting or other tasks. However, if you want to access this log data, you can restore it to the collections created on the Archiver service and then use it for reporting. Cold storage is usually offline storage, such as NAS, or temporary storage before archiving to tape. Once data moves to the Cold Tier, that data is no longer managed by Archiver. Once moved, it is incumbent on external processes to back it up or manage that Cold Tier space such that it does not reach 100% capacity. If capacity is reached, this will cause the Archiver to stop aggregation until the problem is fixed.

Archivers are preconfigured to use available hot storage and a default log collection, so you do not have to configure Archiver storage and log retention if you do not have complex log retention requirements.

Logs can move from one type of storage to another in the following ways:

- Hot Storage > Cold Storage
- Hot Storage > Warm Storage > Cold Storage



When a collection reaches its retention limits for hot and warm storage, NetWitness deletes the log data from hot or warm storage. With cold storage configured, a copy goes into cold storage before the logs are deleted from hot or warm storage. For example, if you have a collection with Hot Storage of 1 TB, Warm Storage of 1 TB, and Cold Storage enabled, when the log data reaches 1 TB of hot storage, the oldest log data moves to warm storage. When the log data in warm storage reaches 1 TB, the oldest log data from warm storage is copied to cold storage before it is removed from warm storage.

For Hot and Warm Storage, size and retention period settings for a collection can override each other based on which criterion (size or time) is satisfied first. For example, if you have a collection with Hot Storage of 1 TB, no Warm or Cold Storage, and a Retention period of 20 days, if the Log data exceeds 1 TB after 11 days, the oldest logs over 1 TB are deleted even though the collection has a 20 day retention period.

After you create hot, warm, and cold storage, you configure your log retention storage collections. You can specify the maximum size of the Hot and Warm Storage for the collection, whether to use Cold Storage, the number of days to retain the logs in the collection, the data compression, and whether to use a hash algorithm to be able to verify the data integrity of the files being saved.

After configuring your collections, you define retention rules for your collection. These rules specify the type of logs to be stored in the collection. Each collection must have at least one retention rule associated with it in order to store log data.

Procedure

Perform the following tasks in the order shown to configure storage and log retention.

Task	Reference
1. Configure total hot, warm, and cold storage.	Refer to Configure Hot, Warm, and Cold Storage.
2. Configure log retention storage collections.	Refer to Configure Log Storage Collections.
3. Define retention rules for the collections and determine the order of execution of the overall list of retention rules.	Refer to Define Retention Rules.

Configure Hot, Warm, and Cold Storage

This topic provides instructions for Administrators on how to configure total hot, warm, and cold storage on an Archiver.

An Archiver host has hot storage pre-configured to the defaults. Administrators can configure total hot, warm, and cold storage to meet their specific business requirements. An Archiver must have total hot storage configured, but warm and cold storage configurations are optional. NetWitness does not manage cold storage.

Prerequisites

Ensure that you have:

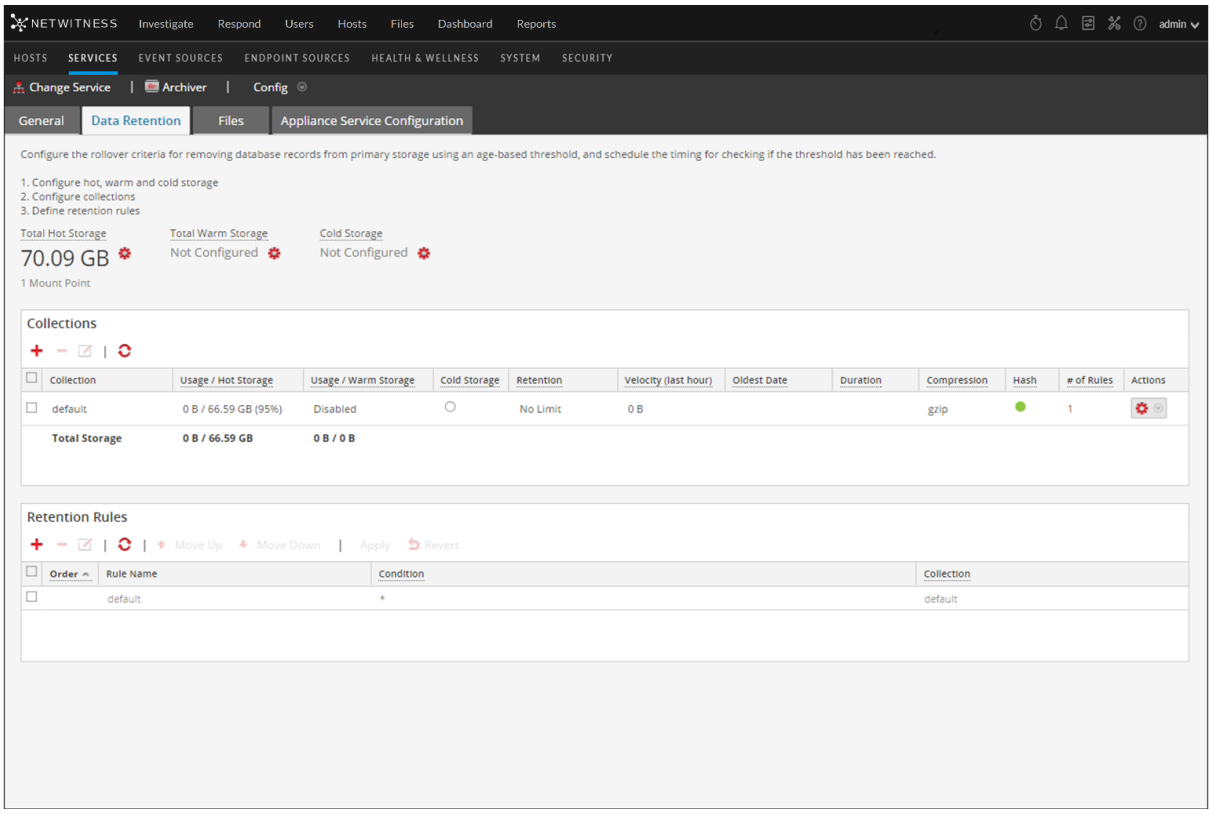
1. Installed the Archiver host in your network environment.
2. Installed and configured Log Decoder in your network environment.
3. Added Archiver as a Core service to your NetWitness deployment.
4. Added Log Decoder services as a data source for Archiver.
5. Installed and configured a DAC or other physical storage in your network environment.
6. Determined your log retention and storage requirements.

Procedures

Configure Total Hot Storage for an Archiver

1. Go to  **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, click  **> View > Config**.

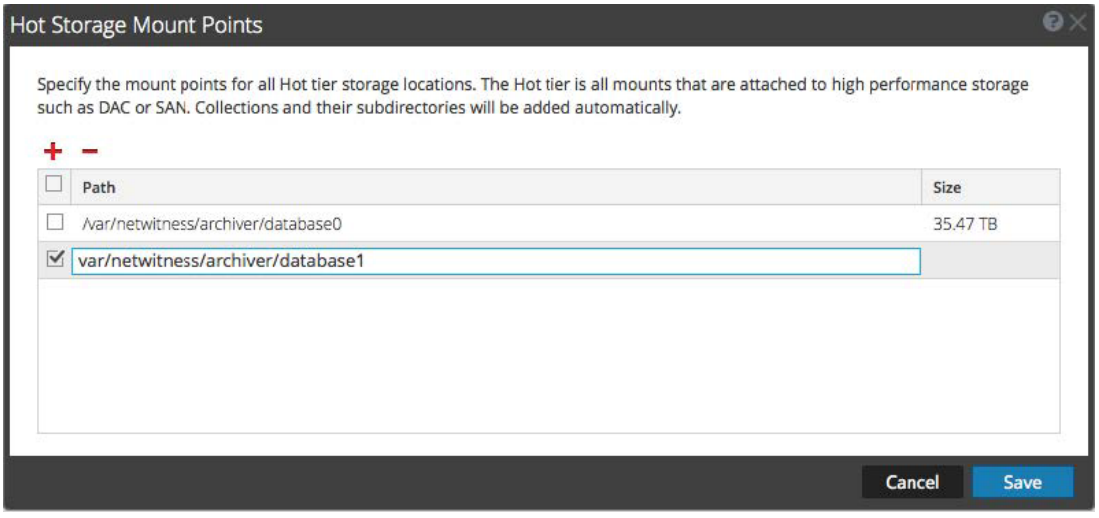
   The Services Config view of Archiver is displayed with General tab open.
3. Click the **Data Retention** tab, in the **Total Hot Storage** section, click  to configure total hot storage.



4. In the **Hot Storage Mount Points** dialog, add the mount points attached to the Archiver host that you want to include in Total Hot Storage.

   These are the paths to high performance storage, such as DAC storage and SAN. Do not add collections or subdirectories to the mount points.

   To add a mount point, click + and type the path to the mount point.



5. Verify that your mount point paths are correct and click **Save**.

NetWitnesswill automatically create metadb, packetdb, sessiondb, and index directories for each collection defined on the Archiver:

```
<storageLocation>/<CollectionName>/metadb
<storageLocation>/<CollectionName>/packetdb
<storageLocation>/<CollectionName>/sessiondb
<storageLocation>/<CollectionName>/index
```

For example, if your mount point is /var/netwitness/archiver, then the following directories will be created for each of your collections:

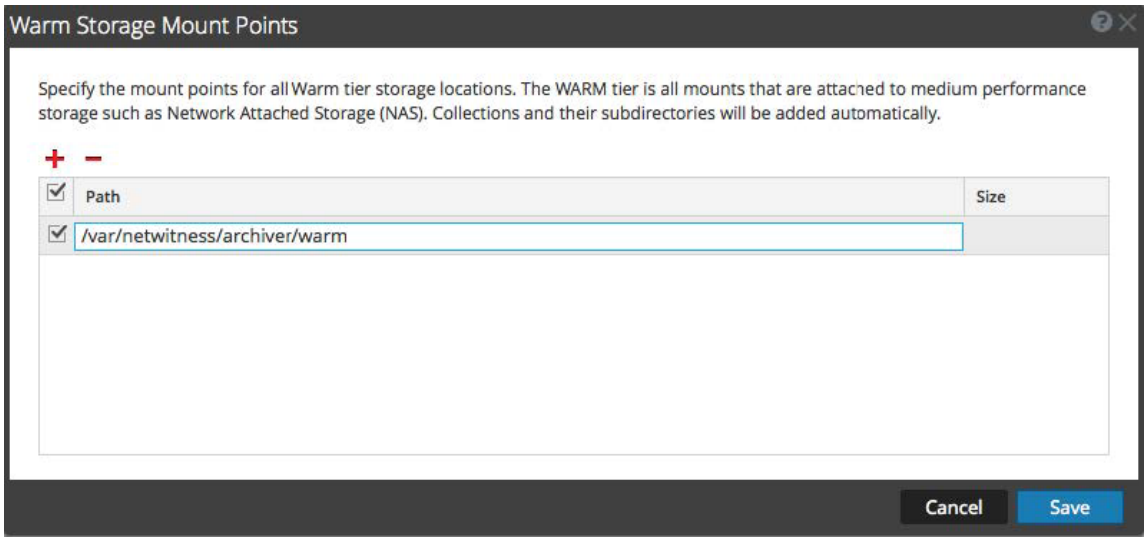
```
/var/netwitness/archiver/<CollectionName>/metadb
/var/netwitness/archiver/<CollectionName>/packetdb
/var/netwitness/archiver/<CollectionName>/sessiondb
/var/netwitness/archiver/<CollectionName>/index
```

After the Archiver service is restarted, data will start being saved to your defined collections. Ensure that your log retention collections are correct before restarting the Archiver service.

Caution: After data has been saved to a mount point, it cannot be removed from the user interface.

Configure Total Warm Storage for an Archiver

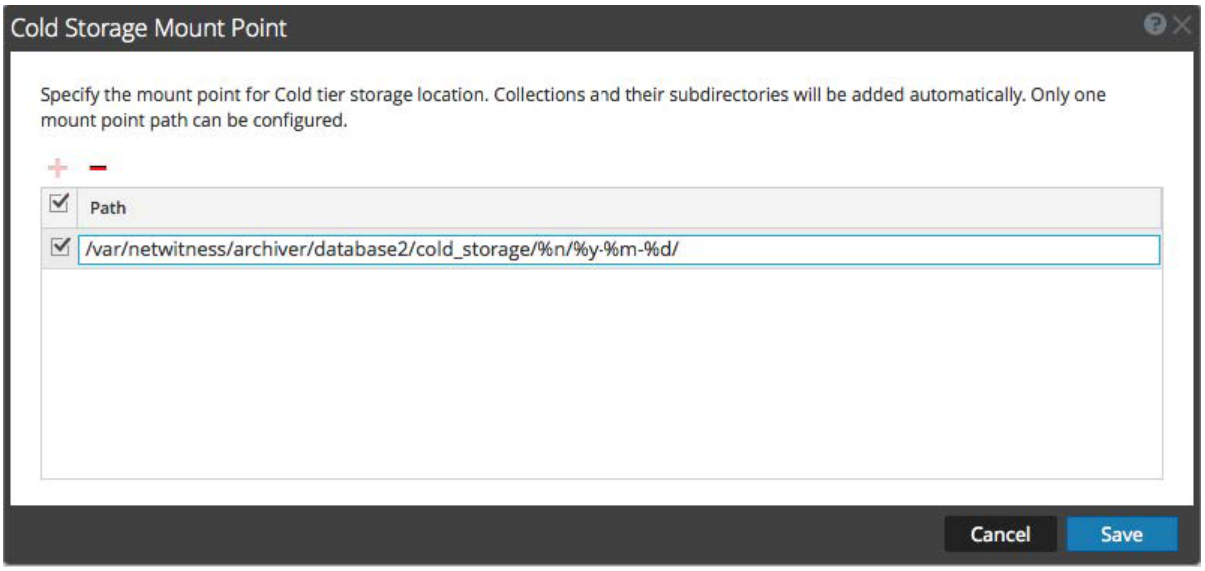
(Optional) The procedure to configure Total Warm Storage for an Archiver is the same as for Total Hot Storage, except that you click in the Total Warm Storage section and add the mount points that you want to use for warm storage, which are the physical paths to warm storage, such as Network Attached Storage (NAS).



Configure Total Cold Storage for an Archiver

(Optional) The procedure to configure Total Cold Storage for an Archiver is the same as for Total Hot Storage, except that you click in the Total Cold Storage section and you add only one mount point for cold storage. NetWitness does not manage cold storage.

You must include the collection name format specifier `%n` somewhere in the cold storage mount point path name to avoid filename collisions between collections.



The following format specifiers are allowed in the path:

Format Specifier	Description
%n	collection name (required)
%y	year the data moved to cold storage
%m	month
%d	day
%h	hour
%##r	block of hours for the current day. For example, if you want three 8 hour blocks, you can set it to %8r. The first 8 hours of the day returns 0, the second 8 hours returns 1, and last 8 hours of the day returns 2.

Changes take effect immediately.

For example, if you have a collection named **compliance** and you create the following cold storage path:

`/sa-cold-storage/%n/%y-%m-%d/`

NetWitness creates a directory each day with the following format:

/sa-cold-storage/compliance/2015-11-20/

Hot, Warm, and Cold Tier Storage Features

The following table describes features of the Hot, Warm, and Cold Tier Storage dialogs.

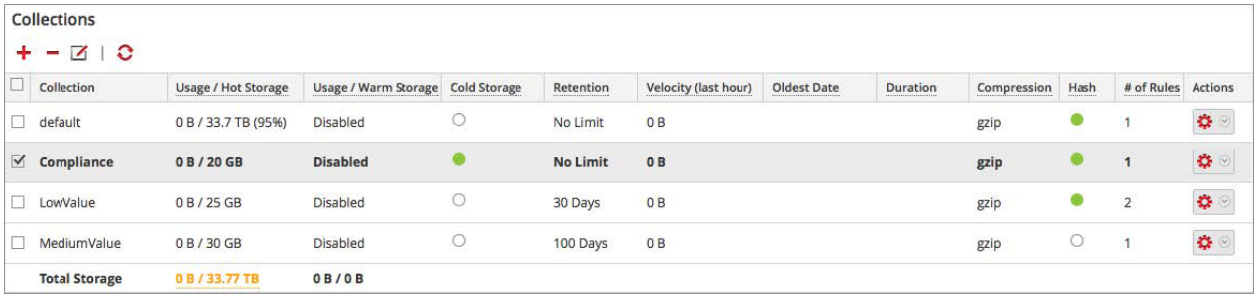
Feature	Description
+	Adds a mount point.
−	Removes a mount point. You cannot delete a mount point that is in use unless you delete the associated collections.
☐	Select the mount points that you want to include for the Total Hot, Warm, and Cold Storage. You can only select one mount point for Total Cold Storage.

Mount Point	Shows the path to the attached physical storage. For example: **/var/netwitness/archiver/database0**, which is the location of the hot storage DAC.
	Do not add collections or subdirectories to the mount points. NetWitness will automatically create metadb, packetdb, sessiondb, and index directories for each collection defined on the Archiver:
	`<storageLocation>/<CollectionName>/metadb`
	`<storageLocation>/<CollectionName>/packetdb`
	`<storageLocation>/<CollectionName>/sessiondb`
	`<storageLocation>/<CollectionName>/index`
	For example, if your hot storage mount point is /var/netwitness/archiver, then the following directories will be created for each of your collections:
	`/var/netwitness/archiver/<CollectionName>/metadb`
	`/var/netwitness/archiver/<CollectionName>/packetdb`
	`/var/netwitness/archiver/<CollectionName>/sessiondb`
	`/var/netwitness/archiver/<CollectionName>/index`
	For Cold Storage, you must include the collection name format specifier `%n` somewhere in the cold storage mount point path name to avoid filename collisions between collections.

Storage Size	Shows the size of the attached storage. The Data Retention tab shows the total amount of storage for your reference.

Collections

The Collections section lists all of your storage collections along with Total Storage for Hot and Warm Storage.



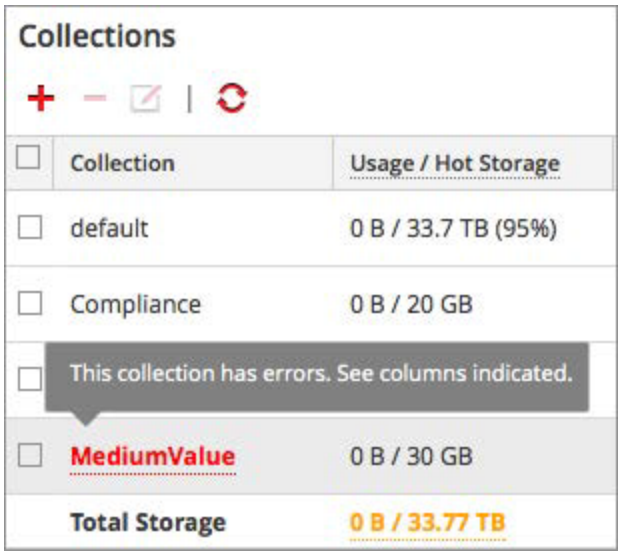
Collections Features

The following table describes the icons and columns of the Collections section. You can hide some of the columns based on your requirements.

Feature	Description
	Opens the Collections dialog, in which you can add a storage collection.
	Removes the selected collection. Deleting the collection permanently removes all stored data from the collection, but the empty data directories remain.
	Opens the Collections dialog, in which you can edit the selected collection.
	Refreshes collection information.
	Selects a collection. For example, you can select a collection for editing or removal.
Collection	Shows the name of your collection, such as Default, Compliance, MediumValue, and LowValue. You can create multiple collections with different criteria for retaining logs. If you do not create any collections, the Default collection is used. If a collection has errors, the collection name and the columns with errors appear in red text.
Usage / Hot Storage	Shows the current hot storage usage and the maximum hot storage for the collection. When the size of the logs reach the maximum hot storage amount, the logs are removed or they roll to the next available storage tier (warm or cold).
Usage / Warm Storage	Shows the current warm storage usage and the maximum warm storage for the collection. When the size of the logs reach the maximum warm storage amount, the logs are removed or they roll to available cold storage.
Cold Storage	Indicates whether cold storage is enabled or disabled. A solid colored green circle indicates that cold storage is enabled ( ). An blank white circle indicates that cold storage is disabled.
Retention	Shows the number of days that logs are retained before being removed or optionally moved to cold storage. No Limit indicates that log retention is not restricted by a specified number of days. For Hot and Warm Storage, size and retention period settings for a collection can override each other based on which criterion (size or time) is satisfied first.
Velocity (last hour)	Shows the number of logs captured over the last hour.
Oldest Date	Shows the date and time of the last log capture.
Duration	Shows how may days ago that the last log was captured. For example: 20 days.
Compression	Shows the compression type used for the meta and raw data in the collection.
Hash	Shows whether hash is enabled or disabled. When enabled, the hash algorithm is used to ensure the data integrity of the files being saved. By default, the only data being hashed is raw logs and the hash files are saved in the same directory as data.

Feature	Description
# of Rules	Shows the number of rules applied to the collection. Define at least one rule for each collection. A collection without any associated rules shows a zero in red text as a warning: The collection name also appears in red text, which indicates an error in the collection. **Caution:** If a collection does not have a rule, no logs will ever go into that collection.
Actions	Enables you to see the rules associated with a collection in the Retention Rule section when you select <actions button> **> Select Rules**. In the Retention Rule section, you can change the overall priority of the collection rules.
Total Storage	Shows the current total hot storage usage and the maximum total hot storage at the bottom of the **Usage / Hot Storage** column. It also shows the current total warm storage usage and the maximum total warm storage at the bottom of the **Usage / Warm Storage** column.

Any errors in the collection appear in red text. A dotted underline indicates that a tooltip is available with information about the error.



Collections that have editing disabled (grayed out) also have tooltips that provide information on the problem.

Retention Rules

The Retention Rules section lists all of the retention rules used for your storage collections listed in the order of rule execution.



The following table describes the features of the Retention Rule section.

Feature	Description
	Opens the Rule Definition dialog, in which you can add a retention rule to use in a storage collection.
	Removes the selected retention rule. In order for your log collections to gather and store log data, you must associate them with at least one retention rule.
	Opens the Rule Definition dialog, in which you can edit the selected retention rule.
	Refreshes retention rule information.
Move Up	Moves the selected retention rule up in the Retention Rule priority list. Retention Rule order is very important. NetWitness evaluates the the retention rules for all of the collections in numerical order by the number listed in the Order column in the Retention Rule section. You can also use drag and drop to reorder retention rules.
Move Down	Moves the selected retention rule down in the Retention Rule priority list. Retention Rule order is very important. NetWitness executes the retention rules for all of the collections in numerical order by the number listed in the Order column in the Retention Rule section.
Apply	Saves the rule order change.
Revert	Reverts the rule order change.
	Selects or shows a selected retention rule.
Order	Shows the order of a rule in the overall list of retention rules.
Rule Name	Shows the name of rule, such as ComplianceDevices and GeneralWindowsLogs.
Condition	Shows the conditions for the rule. These conditions specify the type of logs to include in the collection. Define Retention Rules presents the guidelines for all queries and rule conditions in Core services.
Collection	Shows Collection name and how many days that the collection is retained. For example: MediumValue (30 Days)

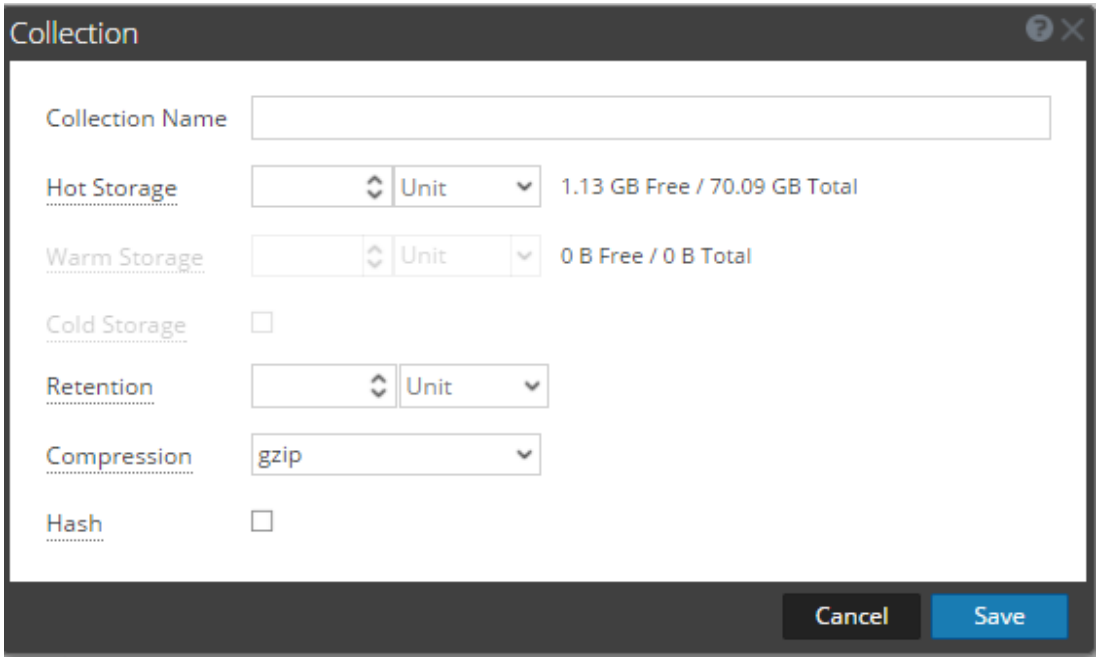
Collection Dialog

On the (Admin) > Services > Config view > Data Retention tab of an Archiver, administrators can define the criteria for log retention and storage. In the Collection dialog, which is accessible from the Collections section, you can define individual storage collections to use for different log types. For example, you may want to create collections for compliance reasons or to selectively retain critical logs.

Procedures related to this dialog box are described in Configure Archiver Storage and Log Retention and Configure Log Storage Collections.

To access the Collection dialog:

1. Go to **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, click **> View > Config**.
   The Services Config view of Archiver is displayed with General tab open.
3. Click the **Data Retention** tab.
4. In the **Collections** section, click to add or edit the rule.
   The Collection dialog is displayed.



The following table describes the fields in the Collection dialog.

Field	Description
Collection Name	Specify a name for your collection, such as Compliance, MediumValue, or LowValue.

Field	Description
Hot Storage	Specify the maximum size or percentage of hot storage to use for this collection. The free space available to use for hot storage and the total hot storage are shown next to this field. When the size of the logs reach the maximum hot storage size, the logs are removed or they roll to the next available storage tier (warm or cold).
Warm Storage	(Optional) Specify the maximum size or percentage of warm storage to use for this collection. The free space available to use for warm storage and the total warm storage are shown next to this field. When the size of the logs reach the maximum warm storage size, the logs are removed or they roll to available cold storage.
Cold Storage	(Optional) Specify whether to use cold storage for this collection. If you use cold storage for the collection, logs outside of the specified size and retention limits roll over to cold storage. If you do not use cold storage, logs outside of the specified size and retention limits are removed.
Retention	(Optional) Specify the number of days that logs are retained before they are removed or rolled over to cold storage. For Hot and Warm Storage, size and retention period settings for a collection can override each other based on which criterion (size or time) is satisfied first.
Compression	Specify the type of compression to use for meta and raw logs in the collection. You can compress the meta and raw logs using GZIP or LZMA to save space. GZIP is very fast at compressing and decompressing, but it does not compress as well as LZMA. LZMA offers better compression at a cost of decompression speed (roughly three times slower than GZIP). Compression ratios are highly dependent on your data. The default compression is GZIP.
Hash	Specify whether to enable or disable hash. When enabled, the hash algorithm is used to verify the data integrity of the files being saved. By default, the only data being hashed is raw logs and the hash files are saved in the same directory as data.

Note: When decreasing collection storage allocations or lowering retention time, it may take several minutes to hours for the data to move and space to become available depending on the amount of moving (rolling) data. The default times are every 20 minutes for a size roll and every six hours for a time roll.

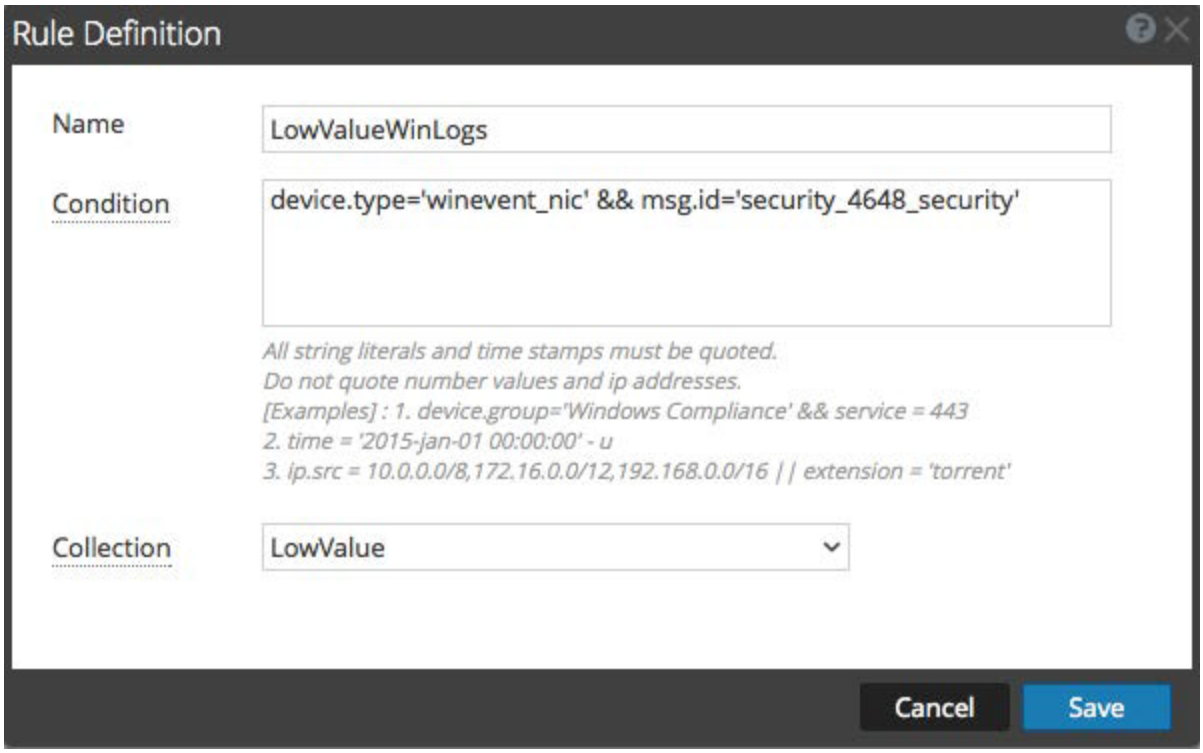
Rule Definition Dialog

In the  (Admin) > Services > Config view > Data Retention tab of an Archiver, administrators can define the criteria for log retention and storage. In the Rule Definition dialog, which is accessible from the Retention Rules section, you can define retention rules to use for your storage collections.

Procedures related to this dialog box are described in Configure Archiver Storage and Log Retention and .Define Retention Rules

To access the Rule Definition dialog:

1. Go to **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, click **> View > Config**.
   The Services Config view of Archiver is displayed with General tab open.
3. Click the **Data Retention** tab.
4. In the **Retention Rule** section, click or .
   The Rule Definition dialog is displayed.



The following table describes fields in the Rule Definition dialog.

Field	Description
Name	Specify a unique name for your retention rule. For example: ComplianceDevices
Condition	Specify the conditions for the type of logs that you want to include in the collection. All sting literals and time stamps must be quoted. Do not quote number values and IP addresses. For example: `device.group='PCI Devices' || device.group='HIPPA Devices'`
Collection	Select the collection on which you want to apply this rule. For example: Compliance

Configure Log Storage Collections

This topic provides instructions for Administrators on how to configure log storage collections on an Archiver.

NetWitness enables you to define individual storage collections for different log types. You can specify the maximum size of the Hot and Warm Storage space used by the collection, whether to use offline storage (Cold Storage), the number of days to retain the logs in the collection, the data compression, and whether to use a hash algorithm to be able to verify the data integrity of the files being saved. You should create collections based on your log retention storage requirements. Each collection that you create must be associated with at least one retention rule.

Prerequisites

Before you configure your log retention storage collections, configure total hot, warm, and cold storage.

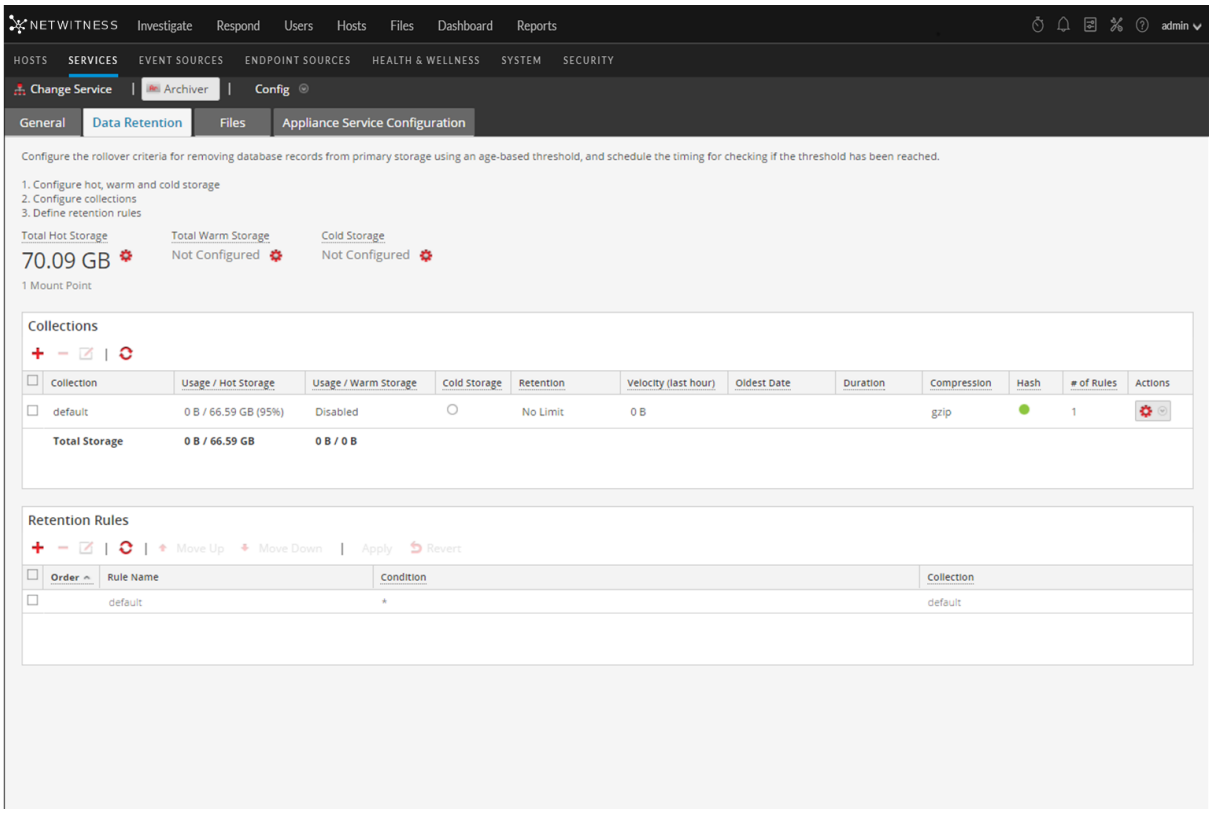
Configure a Log Storage Collection

To configure a log retention storage collection on an Archiver:

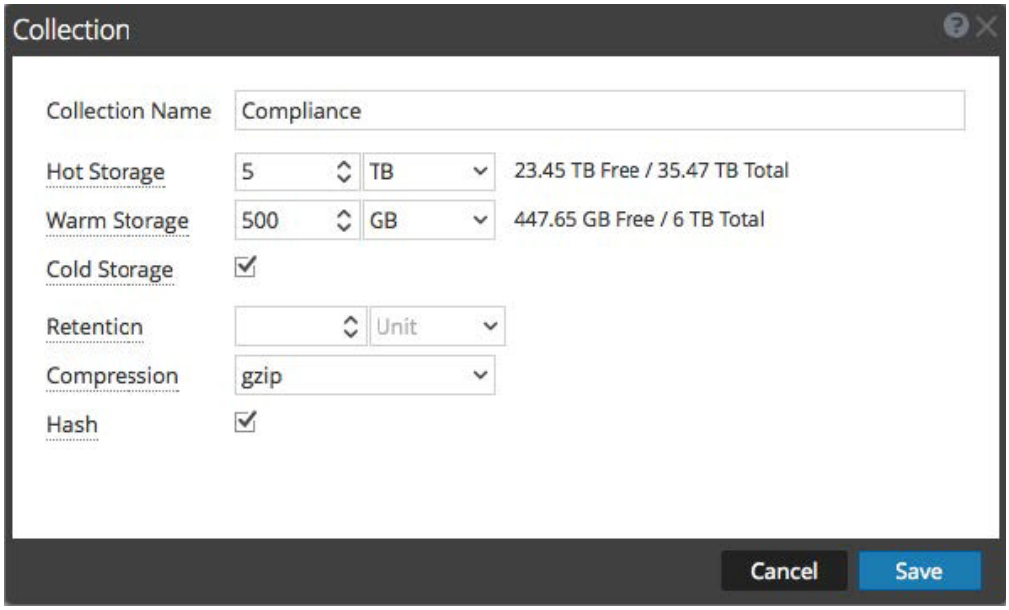
1. Go to  **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, click  **> View > Config**.

   The Services Config view of Archiver is displayed with General tab open.
3. Click the **Data Retention** tab, in the **Collections** section, click  to add a collection.

   (If you decide to make changes to an existing collection, you can select the collection and click  to change the settings.)



The **Collection** dialog is displayed.



4. Configure the collection as described in the following table.

Field	Description
Collection Name	Specify a unique name for your collection, such as Compliance, MediumValue, or LowValue.
Hot Storage	Specify the maximum size or percentage of hot storage to use for this collection. The free space available to use for hot storage and the total hot storage is shown next to this field.
Warm Storage	(Optional) Specify the maximum size or percentage of warm storage to use for this collection. The free space available to use for warm storage and the total warm storage is shown next to this field.
Cold Storage	(Optional) Specify whether to use cold storage for this collection. If you use cold storage for the collection, logs outside the storage limits are copied to cold storage before they are deleted from hot or warm storage.
Retention	(Optional) Specify the number of days that logs are retained before they are removed or rolled over to cold storage. For Hot and Warm Storage, size and retention period settings for a collection can override each other based on which criterion (size or time) is satisfied first.
Compression	Specify the type of compression to use for meta and raw logs in the collection. You can compress the meta and raw logs using GZIP or LZMA to save space. GZIP is very fast at compressing and decompressing, but it does not compress as well as LZMA. LZMA offers better compression at a cost of decompression speed (roughly three times slower than GZIP). Compression ratios are highly dependent on your data. The default compression is GZIP.
Hash	Specify whether to enable or disable hash. When enabled, the hash algorithm is used to verify the data integrity of the files being saved. By default, the only data being hashed is raw logs and the hash files are saved in the same directory as data.

5. Click **Save**.

   Any errors in the collection appear in red text. A dotted underline indicates that a tooltip is available with information about the error. Your collection name appears in red text until at least one retention rule is defined for your collection.
   If you have a collection with editing disabled (grayed out), look at the associated tooltip for more information.

Note: When decreasing collection storage allocations or lowering retention time, it may take several minutes to hours for the data to move and space to become available depending on the amount of moving (rolling) data. The default times are every 20 minutes for a size roll and every six hours for a time roll.

Define Retention Rules

Administrators can define and order retention rules for log storage collections on an Archiver. Retention rules specify the type of logs to be stored in the collection. For your log collections to gather and store log data, you must associate them with at least one retention rule. When you configure a retention rule, you specify a condition and a collection for that rule. The condition (rule definition) determines the type of logs stored in that collection.

For the condition, you can use anything that works in a regular query `where` clause.

For example, to get logs from compliance services, you can use the following condition:

```
device.group='PCI Devices' || device.group='HIPPA Devices'
```

After you define the retention rules for your collections, it is important that you specify the order of your retention rules. NetWitness evaluates the retention rules for all of the collections in numerical order by the number listed in the Order column in the Retention Rule section of the Data Retention tab of the Archiver ((Admin) > Services Config view).



Caution: Rule order is very important. It determines the priority for evaluating the log data for storage retention.

Prerequisites

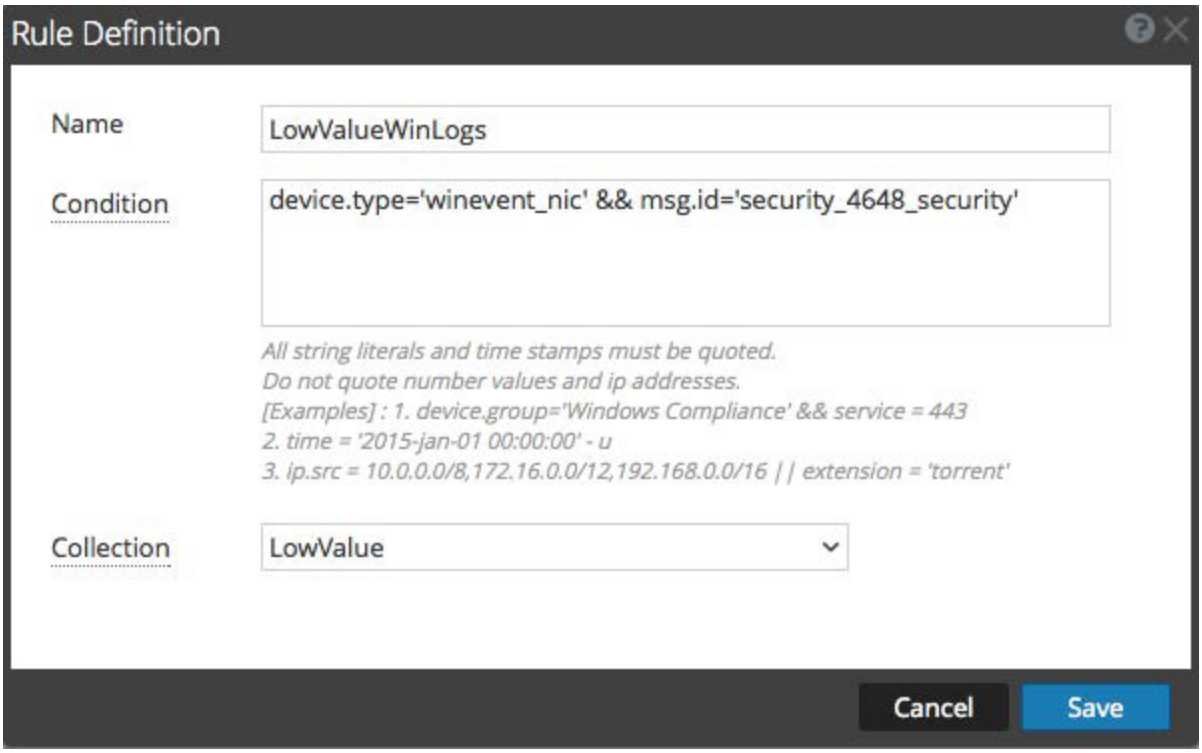
Before you configure your retention rules:

- Configure total hot, warm, and cold storage
- Configure log storage collections

Procedures

Define a Retention Rule for a Collection

1. Go to **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, click **> View > Config**.
   The Services Config view of Archiver is displayed with General tab open.
3. Click the **Data Retention** tab, in the **Retention Rules** section, click .
   The **Rule Definition** dialog is displayed.

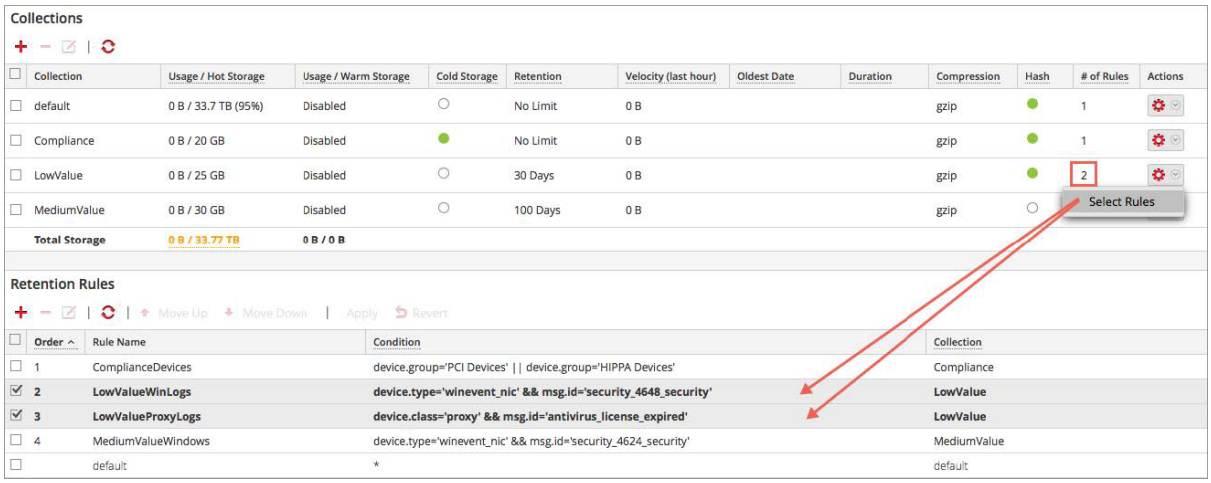


4. Configure the fields in the Rule Definition dialog as described in the following table:

Field	Description
Rule Name	Specify a unique name for your retention rule. It cannot include spaces. For example: LowValueWinLogs
Condition	Specify the conditions for the type of logs that you want to include in the collection. All string literals and time stamps must be quoted. Do not quote number values and IP addresses. For example: `device.type='winevent_nic' && msg.id='security_4648_security'`
Collection	Select the collection on which you want to apply this rule. For example: LowValue.

5. Click **Save**.

   The retention rule that you define becomes associated with the collection you selected. On the **Data Retention** tab, in the **Collections** section, you can click  **> Select Rules** in the **Actions** column for the selected collection to view the retention rules associated with the collection in the **Retention Rule** section.



Specify the Order of your Retention Rules

To prioritize the complete list of all of your retention rules:

1. In the **Retention Rules** section of the **Data Retention** tab, select a retention rule and use drag and drop (or select **Move Up** and **Move Down**) to change its order in the priority list.



2. Click **Apply** to save the order of the retention rules.

Caution: Rule order is very important. It determines the priority for evaluating the log data for storage retention.

Add Archiver as a Data Source to Reporting Engine

This topic provides instructions on how to add Archiver as a data source to Reporting Engine to generate reports for the data collected by Archiver.

Prerequisites

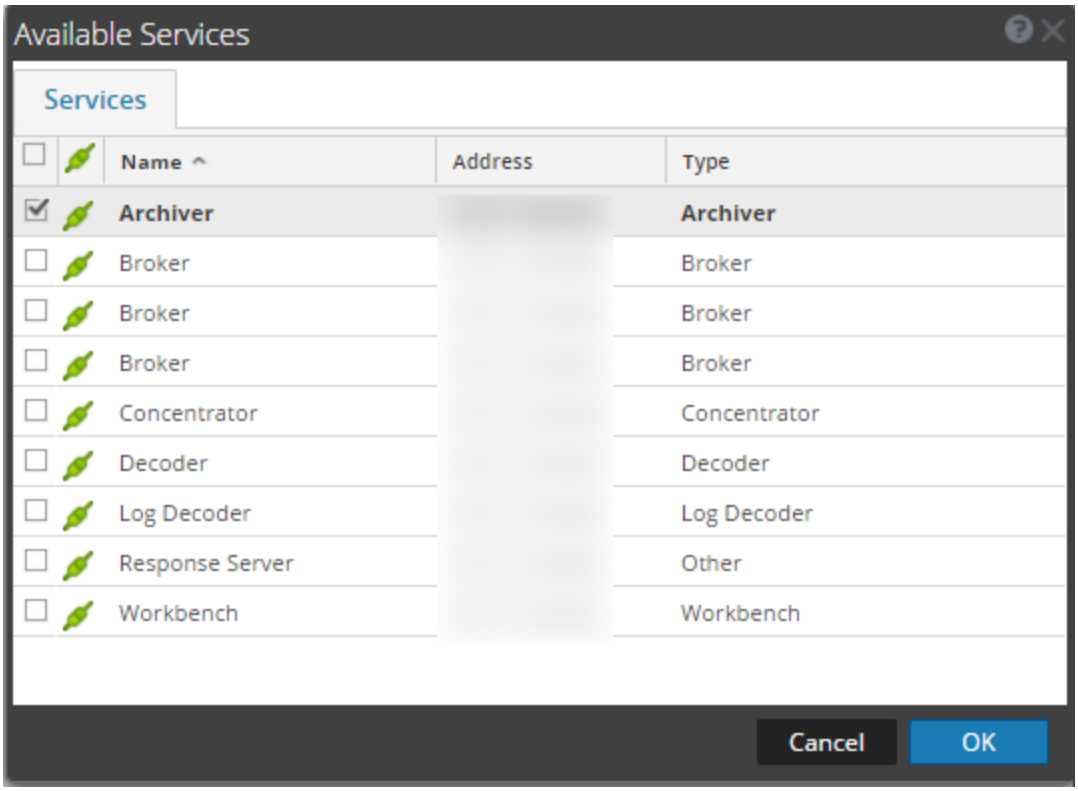
Ensure that you have:

1. Installed the Archiver host in your network environment.
2. Installed and configured a Log Decoder in your network environment.
3. Verified that Reporting Engine and Archiver services are active.

Procedure

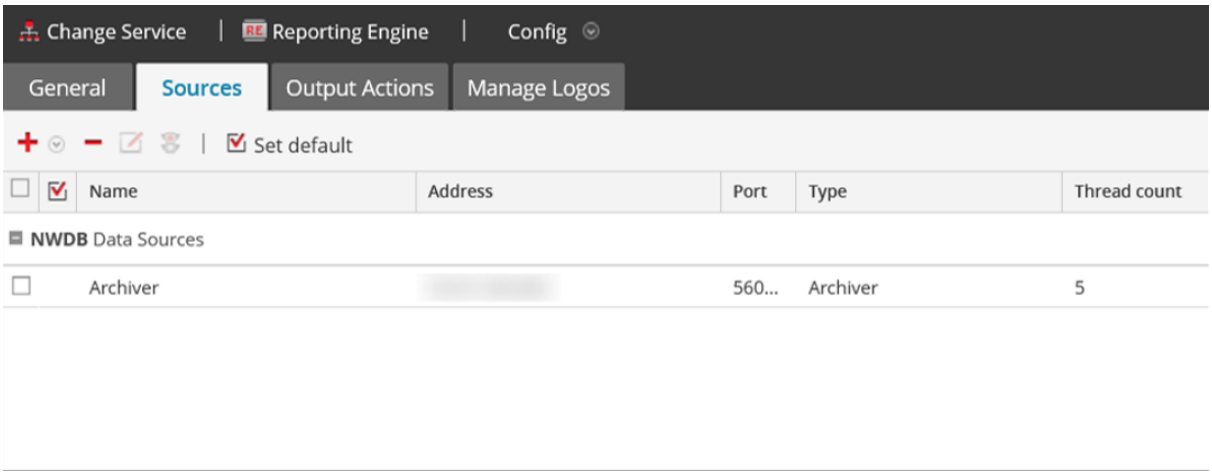
To associate an Archiver data source with Reporting Engine:

1. Go to  **(Admin) > Services**, click , and select **Reporting Engine**.
2. In the **Actions** column, select **View > Config**.
   The Services Config view of Reporting Engine is displayed with General tab open.
3. Select the **Sources** tab.
4. Click  and select **Available Services**.
   The Available Services dialog is displayed.



5. Select the Archiver that you want to add as data source to the Reporting Engine and click **OK**. The Service Information dialog is displayed.
6. Type the username and password for the Archiver.
7. Click **OK**.

   The selected Archiver is listed in the NWDB Data Sources category.

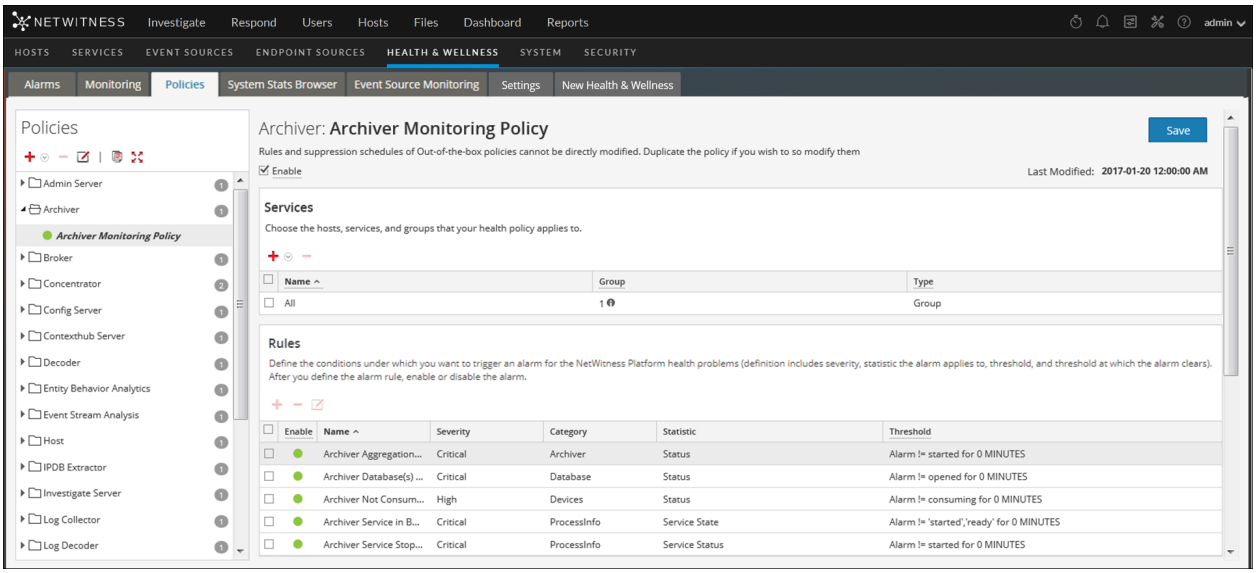


You can now create reports on the data collected by Archiver.

Configure Archiver Monitoring

Health & Wellness enables you to automatically generate notifications when critical thresholds are met.

Review the Health & Wellness policies for Archiver and Host in the Health & Wellness Policies section. Make updates as required.



For detailed information, see "Manage Policies" in the *System Maintenance Guide.*

Additional Archiver Configuration

This topic is a collection of individual procedures, which an administrator may perform at any time and they are not required to complete the initial setup of the Archiver. These procedures are presented in alphabetical order.

Use this section when you are looking for instructions to perform a specific task after the initial setup of an Archiver.

Configuring Data Backup and Restore

This topic provides information on the data backup and restore feature for an Archiver. You can use this feature to back up Archiver data and retrieve the backed up data.

You can back up the data in the following ways:

- Use scripts to copy files from cold storage backup folders onto an offline storage.
- Use backup software to copy files from cold storage backup folders onto an offline storage.
- Run EMC Networker or other backup software on Archiver and have it do daily incremental backup of the database files.

Note: For details on the procedure to back up data using Networker, see the *Administration Guide for Networker*.

Once you have the data backup, you have to perform the following tasks to restore the backed up data that is installed on the Archiver.

Action	Description
1. Restore your data to a location accessible by the Archiver.	Refer to Create Collection
2. Create a collection in Archiver that uses that location.	Refer to "Manage Collections" topic in the *Workbench Configuration Guide*.
3. Add the Archiver service as a data source on Reporting Engine to generate reports for the data restored on the Archiver service.	Refer to Add Archiver as a Data Source to Reporting Engine

Add Archiver Service

The NetWitness Archiver service enables you to create collections with restored data from Archiver offline (cold) storage. This procedure is only required if you do not have the Archiver service installed.

Prerequisites

Make sure you have added an Archiver host and applied a license to it.

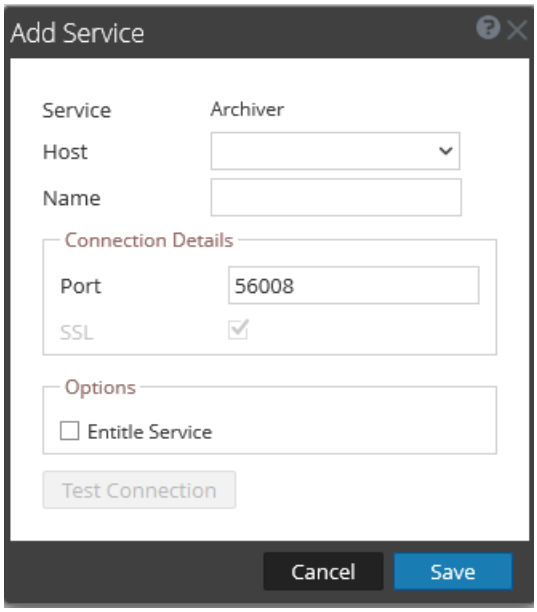
Procedure

Note: This procedure is only required if you do not have Archiver service installed.

Perform the following steps to add the Archiver service:

1. Go to **(Admin) > Services**.
2. In the **Services** panel, select **+>Archiver**.

   The Add Service dialog is displayed, as shown below.



3. Provide the following details.

Field	Description
Host	Select an Archiver host from the drop-down menu.
Name	Type a name for the service.
Port	Default port is 50007.
SSL	Select **SSL** if you want NetWitness to communicate with the service using SSL. The security of data transmission is managed by encrypting information and providing authentication with SSL certificates. **Note:** If you select SSL, ensure SSL is enabled in the System Configuration panel.
Username	(Optional) Type the username for the service.
Password	(Optional) Type the password for the service.

4. Click **Test Connection** to determine if NetWitness connects to the service.
5. When the result is successful, click **Save**.

   The added service is now displayed in the Services panel.

Note: If the test is unsuccessful, edit the service information and retry.

Create Collection

You can create a collection on an Archiver service using data restored from the backed-up data or an existing subset of data. When you recover the backed-up data, you have to place it in the collection folder created on the Archiver service to enable you to generate the required reports for the retrieved data. For example, if you have backed up the data using EMC Networker at <*location*>, you can use the restore options in Networker to restore the backed-up data to the collection folder created on the Archiver service. For restore procedure using EMC Networker, see the *Administration Guide for Networker*.

Prerequisites

Ensure that you have the following:

- Archiver service installed on an Archiver host.
- The Archiverservice has enough space to hold the collection.
- The backed-up data placed in a known location on your local host, if you are creating a collection using the data restored from the backed-up data.

Procedure

The Data Retentions tab enables Administrators to restore and save data that is restored from a backup or from an existing set of data.

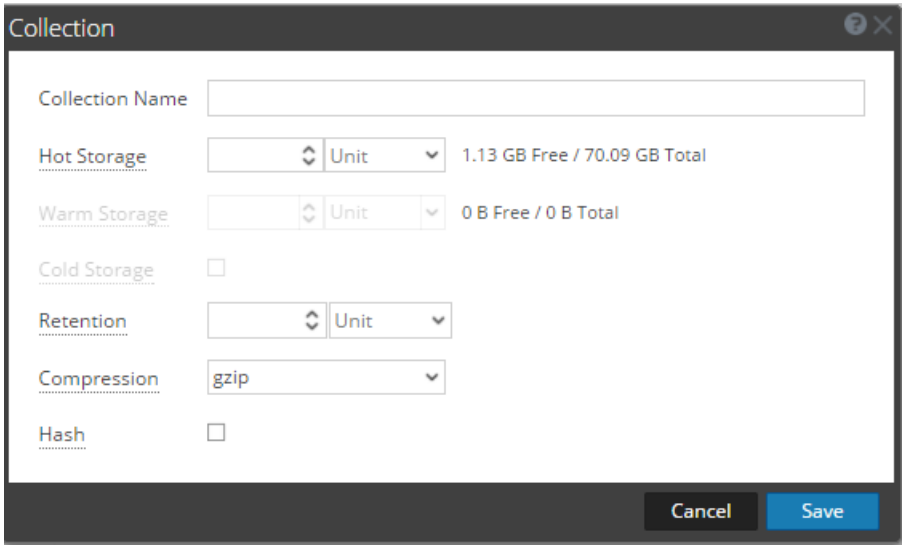
Note: The Administrator can point the source path to the location of the database files and the restore command copies them to the Archiver. The Administrator needs to mount those directories to the Archiver before a restoration collection can be created.

To create a collection using data restored from the backed-up data or existing subset of data:

1. Go to **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, click **> View > Config**.

   The Services Config view of Archiver is displayed with General tab open.
3. Click the **Data Retentions** tab and click in the **Collections** panel to add a collection.

   The **Collection** dialog is displayed.



4. Provide the following information:
   - **Collection Name:** Name of the Archiver collection that you want to restore.
   - **Hot Storage:**Enter the number of Archiver database files and unit size (either Gigabytes or Terabytes) that have been moved from cold storage.
   - **Retention:** Select the number of days or hours that you want to store the collection.
   - **Compression:** Select the compression type for the collection.
5. Click **Save** to restore the collection.

   Note the following:
   - Target is the location where the collection is created.
   - If the source path provided to create the restoration collection does not exist, the following error message is displayed:

     "The source path does not exist '/xxx/xxx/'."

If there is insufficient storage to restore your collection, the following error is displayed:

"Error during disk space checking. Insufficient disk space in location '/xxx/xxx'."

6. The Schedule Job dialog is displayed with the following message:

   "Restoring data into a new collection. Check the jobs page for progress."

7. Click **Jobs**  icon in the top right area of the main menu to expand the list of restoration collection jobs with their current status.

Note: When restoring a collection, the larger the dataset that you have to restore, the longer the restoration will take. If you are restoring a collection containing hundreds of gigabytes or more, restoration may take several hours.

Add Archiver Service as a Data Source to Reporting Engine

This topic provides instructions on how to add the Archiver service as a data source to Reporting Engine to generate reports for the data restored onto the Archiver.

Prerequisites

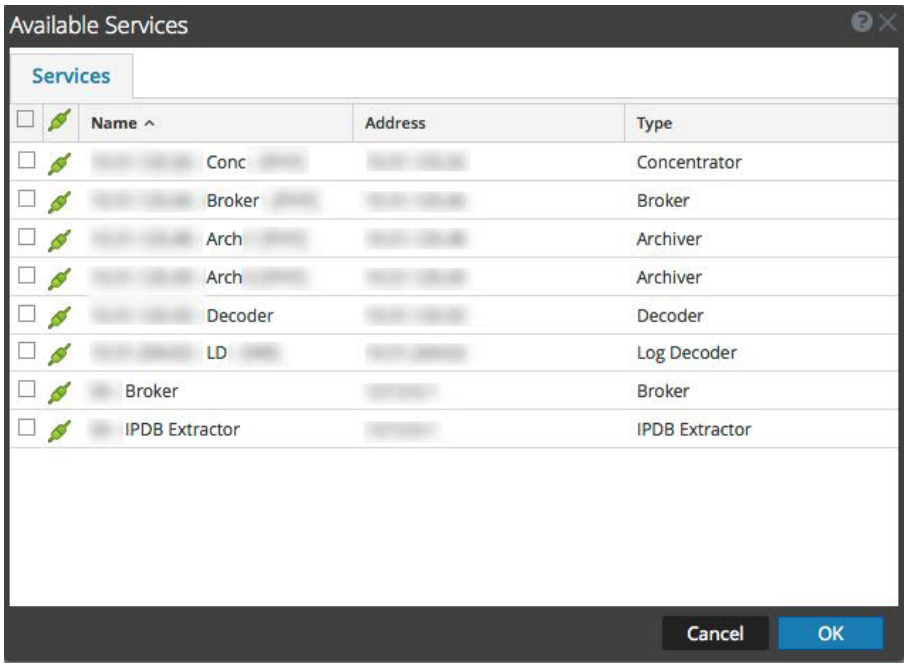
Ensure that you have:

- Installed the Archiver service on the Archiver host.
- Added a collection on the Archiver service.

Procedure

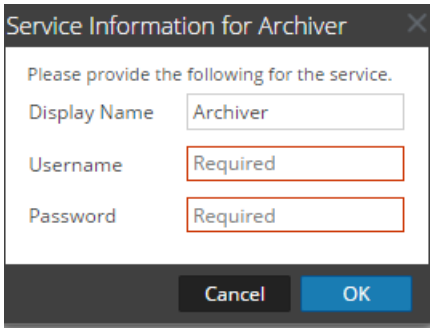
Perform the following steps to add the Archiver service as a data source to Reporting Engine:

1. Go to  **(Admin) > Services**, click  , and select **Reporting Engine** service.
2. In the **Actions** column, click  **> View > Config**.
   The Services Config view of Archiver is displayed with General tab open.
3. Click the **Sources** tab.
4. Click  and select **Available Services**.
   The Available Services dialog is displayed.



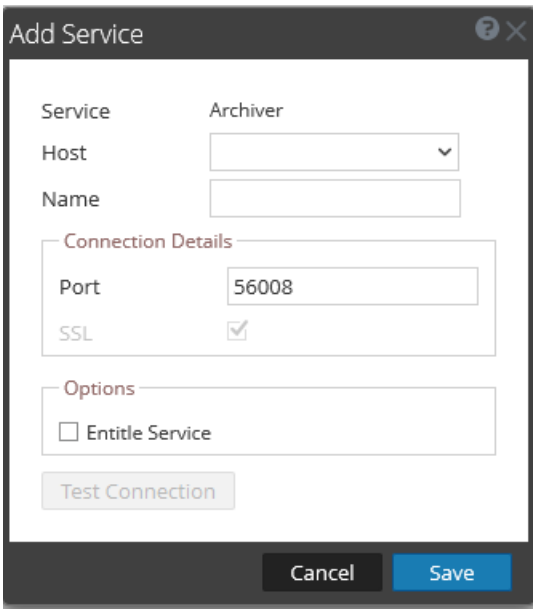
5. Select the Archiver service and click **OK**.

   If the Archiver service is using a Trust Model, the Service Information dialog for the selected service is displayed with the username and password fields required. If the service is not using a Trust Model, these fields will be optional.



6. Type the username and password for admin credentials for the service.
7. Click **OK**.

   The Add Service dialog is displayed.



8. Select a host from the drop-down list and click **Save**.

   The Archiver service is now added as a data source to the Reporting Engine and is listed in the NWDB Data Sources list.

Note: This procedure has to be performed for each collection.

An Administrator can create and delete Workbench collections. and view Workbench statistics and logs. This topic provides all of these procedures and an example procedure for restoring a collection for Reporting and Investigation.

- Mount Archiver Directories
- Create a Collection
- Delete a Collection
- Investigate a Collection
- View Workbench Collection Statistics
- View Workbench Logs

Mount Archiver Directories

If data is in offline storage or cold-tier storage, you need to mount the Archiver directories in order to restore the data for reporting and investigation purposes:

1. Go to  **(Admin) > Services**.
2. Select an **Archiver** from the **Services** panel and select  **> View > Explore**.

   The Explorer view for the Archiver is displayed

3. Right-click on the **Database** node in left-hand tree and select **Database** properties to open them in the right-hand panel.
4. Run the **manifest** command for a time range, for example, 2015-April-01 to 2015-April-10.

   The search returns all files that need to be restored for the selected query.

Create a Collection

Administrators can create collections of restored data from a backup or from an existing set of data.

Note: You can point the source path to the location of the database files and the restore command copies them to the Archiver. You need to mount those directories to the Archiver (where the Workbench is installed) before a restoration collection can be created.

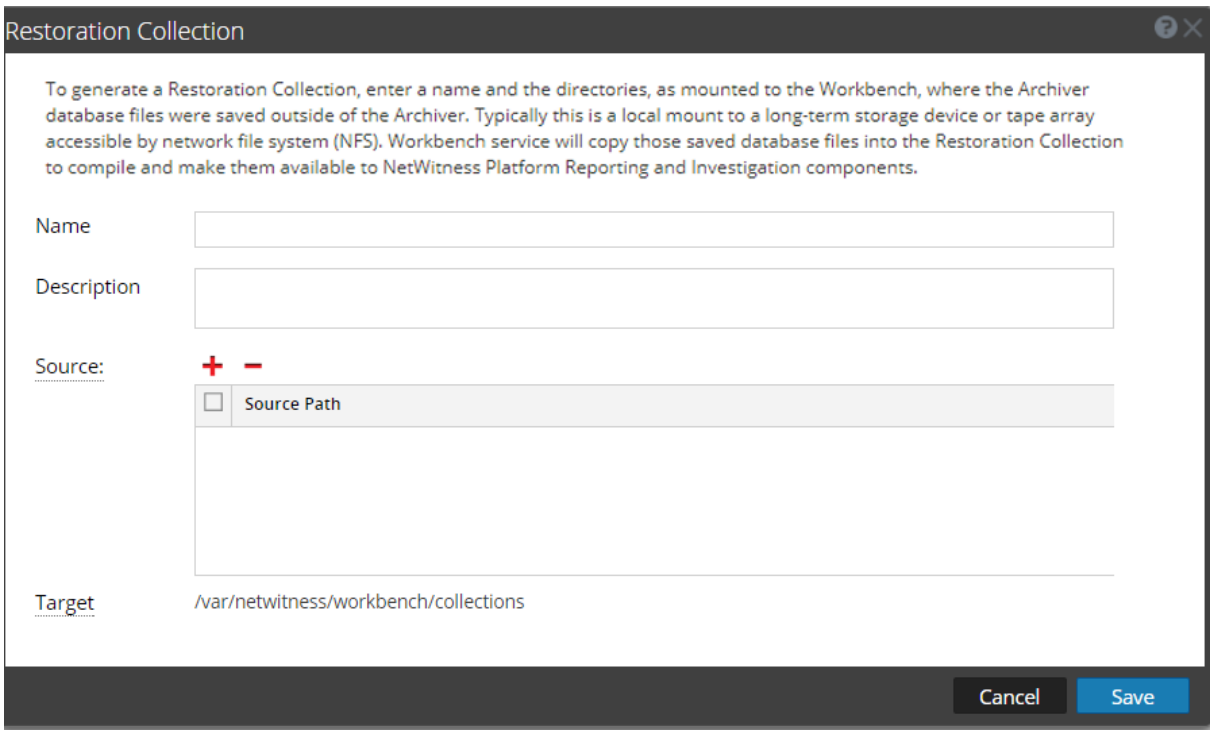
To create a collection using data restored from the backed up data or existing subset of data:

1. Go to **(Admin) > Services**.
2. In the **Services** view, select a **Workbench**, then select **> View > Config**.

   The Services Config view is displayed with the General tab open.
3. Click the **Collections** tab.

   The Collections panel is displayed.
4. Click in the toolbar.

   The Restoration Collection dialog is displayed.



5. Provide the following information:
   - **Name:** Name of the Workbench collection that you want to restore.
   - **Source:** Location where the Archiver database files have been moved from cold storage.

   Note: Target is the location where the collection is created.

6. Click **Save** to restore the collection.

   Note: If the source path provided to create the restoration collection does not exist, the following error message is displayed:
   `The source path does not exist '/xxx/xxx/'.`

   If there is insufficient storage to restore your collection, the following error is displayed:
   `Error during disk space checking. Insufficient disk space in location '/xxx/xxx'.`

   The Schedule Job dialog is displayed with the following message:
   `Restoring data into a new collection. Check the jobs page for progress.`

7. Click the **Jobs** icon  in the NetWitness toolbar to expand the list of restoration collection jobs with their current status.

Note: Restoring a collection that is larger than 550 GB may take several hours to process.

Delete a Collection

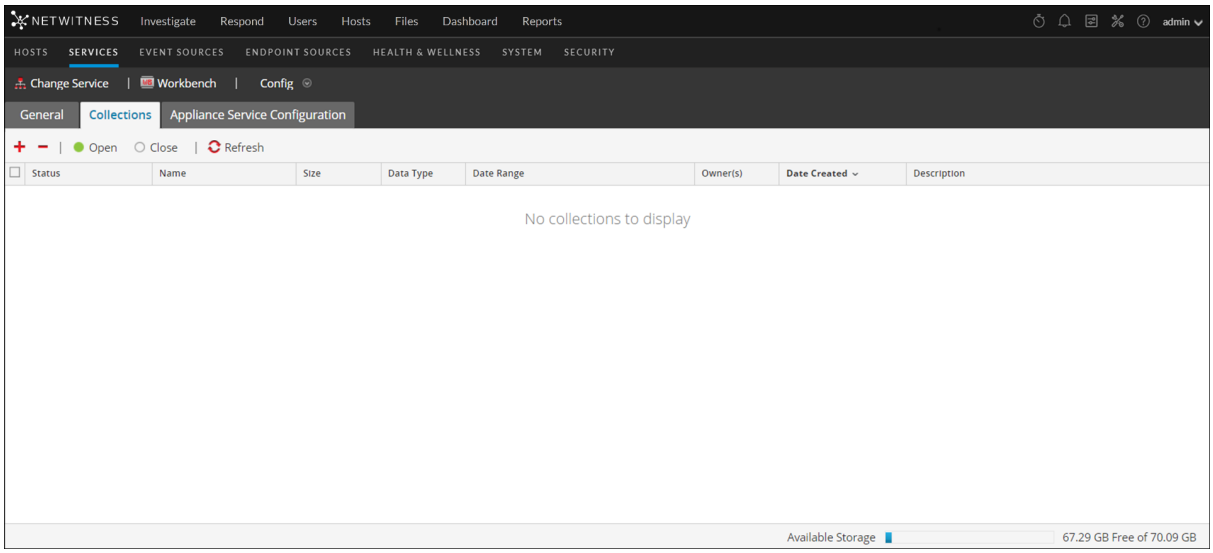
Administrators can delete collections from the Workbench service.

Perform the following steps to delete a collection:

1. Go to  **(Admin) > Services**.
2. In the **Services** view, select a **Workbench**, and click  **> View > Config**.

   The Services Config view opens with the General tab displayed.
3. Select the **Collections** tab.

   The Collections panel is displayed.



4. In the Collections panel, select the collection that you want to delete.
5. Click − from the toolbar.

   A warning dialog requests confirmation.
6. If you want to delete the collection, click **Yes**.

   The collection is removed from the Workbench service.

Example Procedure: How to Restore a Collection for Reporting and Investigation

The following steps illustrate how to restore data for reporting and investigation purposes that is in offline storage or cold-tier storage. In the following example, data is restored for the time range beginning on 2015-April-01 through 2015-April-10.

To restore data for reporting and investigation purposes:

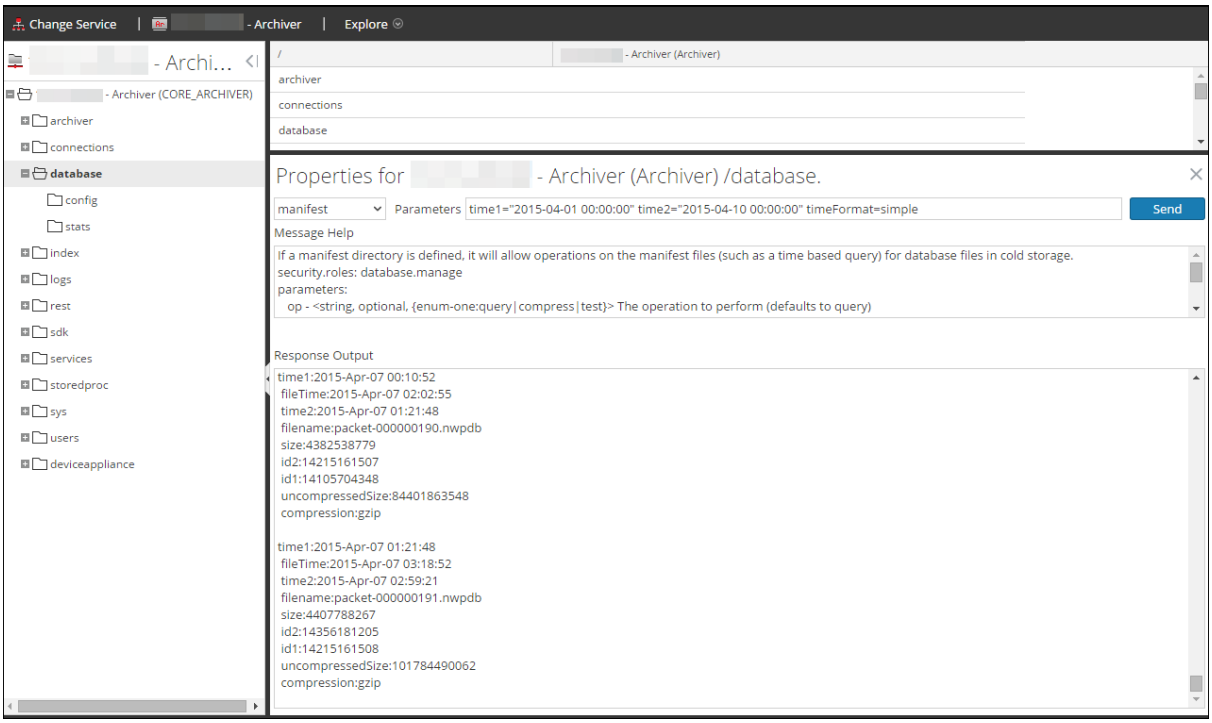
1. Go to **(Admin) > Services**, select the **Archiver**, and click **> View > Explore**.

   The Explorer view for Archiver is displayed.
2. Right-click on **Database** node in left-hand tree and select **Database** properties to open them in the right-hand panel.
3. Run the **manifest** command for the selected time range 2015-April-01 to 2015-April-10.

   The search returns all files that need to be restored for your selected query.

   Example Search:

   ```
   time1="2015-04-01 00:00:00" time2="2015-04-10 00:00:00" timeFormat=simple
   ```



4. Go to **(Admin) > Services**, select an **Archiver**, then select **> View > Config**.
   The Services Config view is displayed with the General tab open
5. Select the **Collections** tab.
6. Create a restoration collection with the source path pointing to files listed in the manifest command output.
7. Save the collection.
   After successfully creating a collection, you can use this collection for reporting and investigation purposes.

Investigate a Collection

To perform an investigation on an Archiver collection:

1. Go to **Investigate**.
   The Investigate dialog is displayed.
2. Click the **Collections** tab in the Investigate dialog.
3. Select an Archiver service in the left panel.
4. Select the collection you want to investigate in the right panel.
5. Click **Navigate**.

The Navigate view is displayed showing data pertaining to the Archiver collection that you selected.

Note: For detailed information about using Investigate, see the *Investigate User Guide.*.

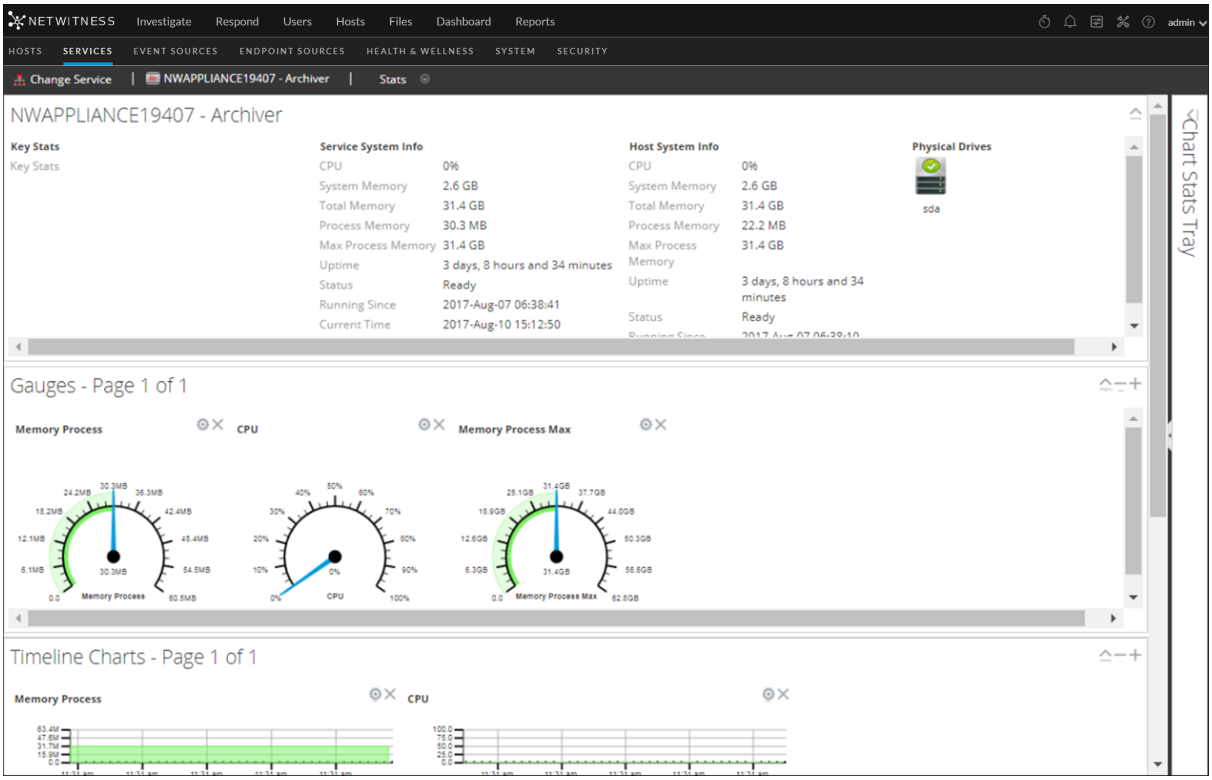
View Archiver Collection Statistics

The same statistics available for other services are provided for the Archiver service. The Services Stats view displays key statistics and system information that pertain to your selected Archiver service. The information is displayed in several different sections within the Stats view: Archiver, Gauges, Timeline Charts and Chart Stats Tray. The Chart Stats Tray lists all available statistics for the Archiver. Any statistic in the Chart Stats Tray can be displayed in a timeline chart.

Perform the following steps to view Archiver statistics:

1. Go to **(Admin) > Services**.
2. In the **Services** view, select an **Archiver**, then select **View > Stats**

   The Services Stats view is displayed.



Note: For more information about Archiver statistics, see the *Host and Services Getting Started Guide*.

View Archiver Logs

Perform the following steps to view logs on an Archiver service:

1. Go to **(Admin) > Services**.
2. In the **Services** view, select a **Archiver**, then select **> View > Logs**

The Services Logs panel is displayed.

Note: For information about viewing and configuring audit logs, see "Configure Global Audit Logging" topic in the *System Configuration Guide.*

Add Archiver Service as a Data Source to Broker

Adding the Archiver service as a data source to Broker is useful when you have more than one collection and you want a report on the archived data. To do this, you can add more than one collection as a downstream service to a Broker and then generate a report on it.

Prerequisites

Ensure that you have:

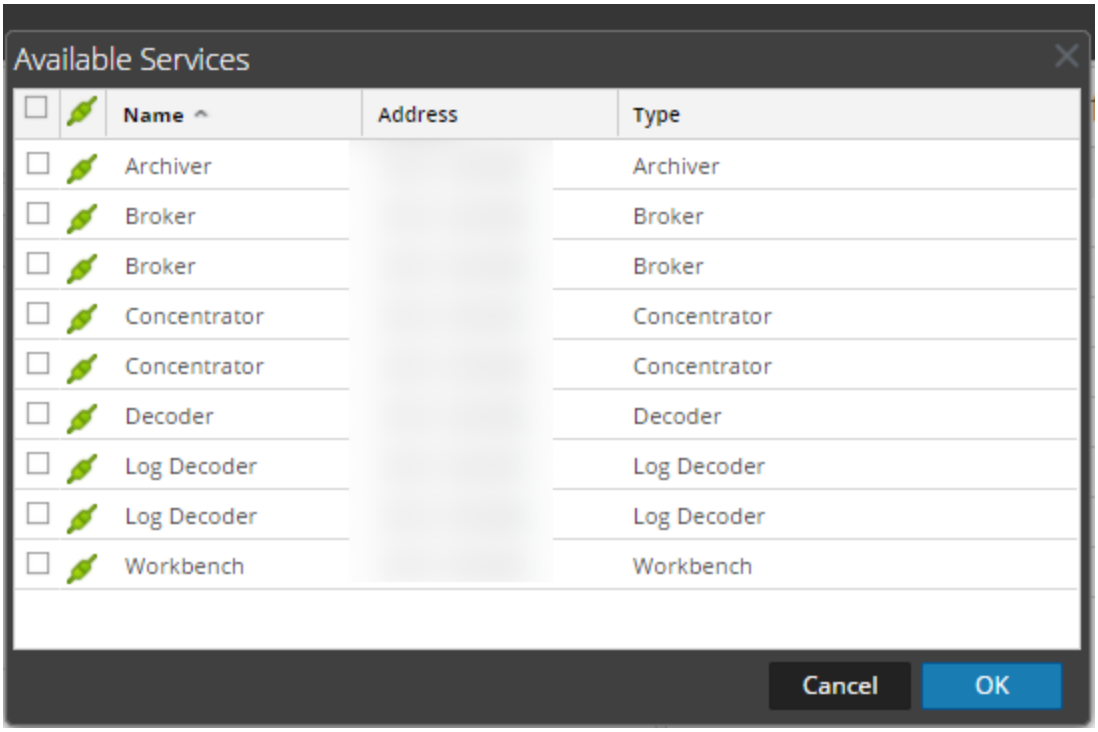
- Installed the Archiver service on the Archiver host.
- Added a collection on the Archiver service.

Procedure

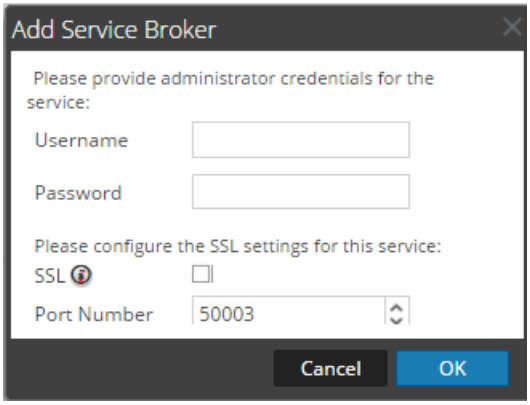
To add an Archiver service as a data source on the Broker:

1. Go to  **(Admin) > Services** and select a **Broker** service.
2. In the **Actions** column, select  **> View > Config**.
   The Config view is displayed with the General tab open.
3. In the **Aggregate Services** section, click .

The Available Services dialog is displayed.

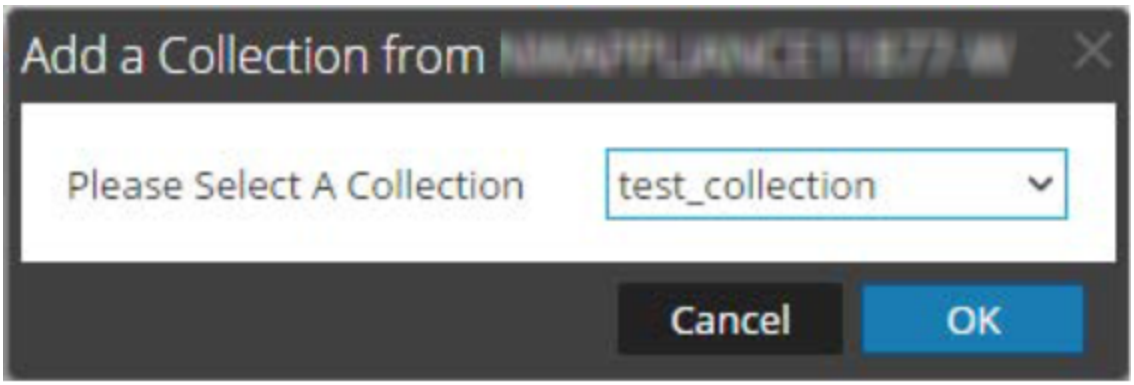


6. Select the Broker service and click **OK**.
7. If the Archiver service is using a Trust Model, a Service Information dialog for the selected service is displayed.



8. Type the username and password for admin credentials for the service.
9. Click **OK**.

   The Add Collection dialog is displayed.



10. Select a collection from the drop-down list and click **OK**.

    The Archiver service is now added as a data source to the Broker.

Note: This procedure has to be performed for each collection.

Retrieve Hash Information

Archiver provides a command, **hashInfo**, which you can use to retrieve the hash information for each session, meta, and packet database that meets the session list or date range criteria. The hash information retrieved is in the form of a list of string parameters, each string parameter corresponding to the hash information for a single database file. You can retrieve the hash information of the database files using the Archiver Service Explore view or REST interface of the Archiver service. The hash information thus retrieved is used to compare the database files in the original location and the exported location to validate data integrity.

The following table lists the criteria that you can use to retrieve the hash files from the database.

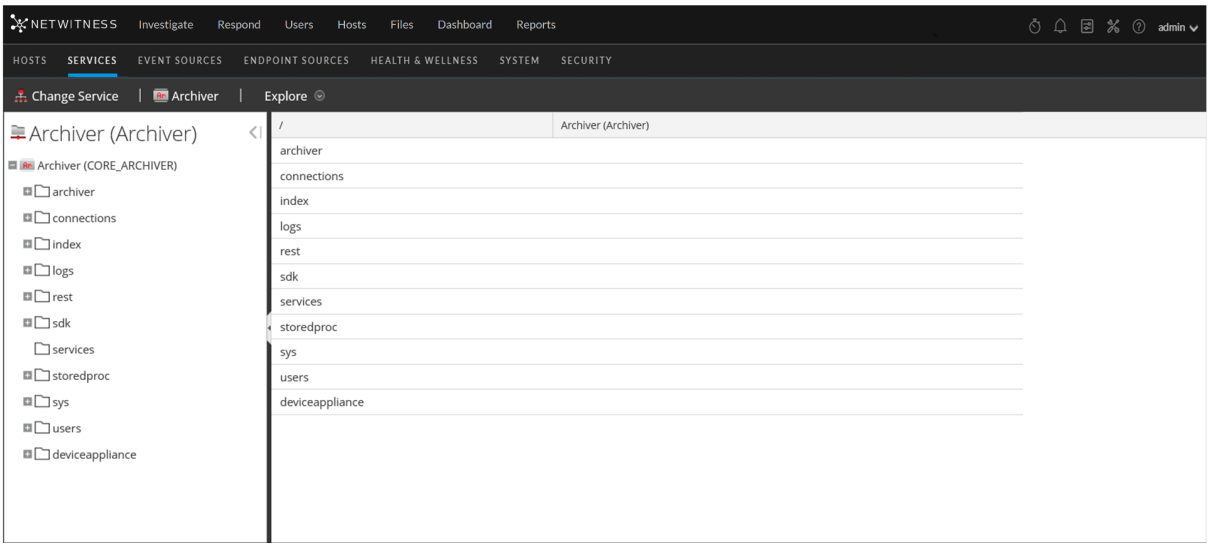
Criteria	Description
sessions	You can retrieve the hash information of the database files by specifying the sessions that exist or read from the session database to determine the associated meta and packet id required to determine which meta and packet database files are needed to retrieve the hash information. **For example:** sessions=100 - Retrieves the hash information of all database files that contain the constituent components(session, meta, content) of session 100. sessions=100,500000 - Retrieves the hash information of all database files that contain the constituent components(session, meta, content) of session 100 and 500000
beginDate	You can specify a begin date as a filter against the database files. This finds the hash information for the files created after the specified date. The begin date specified has to be in the format YYYY-MM-DD HH:MM:SS.
endDate	You can specify an end date as a filter against the database files. This finds the hash information for the files created before the specified date. The end date specified has to be in the format YYYY-MM-DD HH:MM:SS **For example:** beginDate: “2014-Mar-25 05:52:00” endDate=”2014-Mar-27 05:52:00” – Retrieves the hash information of all the database files in between March 25, 2014 and March 27, 2014 in the specified time range on those days.
directories	By default, the hash information files are stored with the database files they were created for. You can also store the hash information file in different location by defining multiple locations in the hash.dir configuration parameter. You can define the location as a filter and retrieve the hash information files for the configured location. **For example:** directories="/home/hash" – Retrieves the hash information of the database files from the location /home/hash

Procedure

To retrieve hash information of the database files:

1. Go to **(Admin) > Services**, click , and select **Archiver**.
2. In the **Actions** column, select **View > Explore**.

   The Explore view of the Archiver service is displayed.



3. In the node tree, right-click on **archiver** and select **Properties**.

   The Properties dialog is displayed.



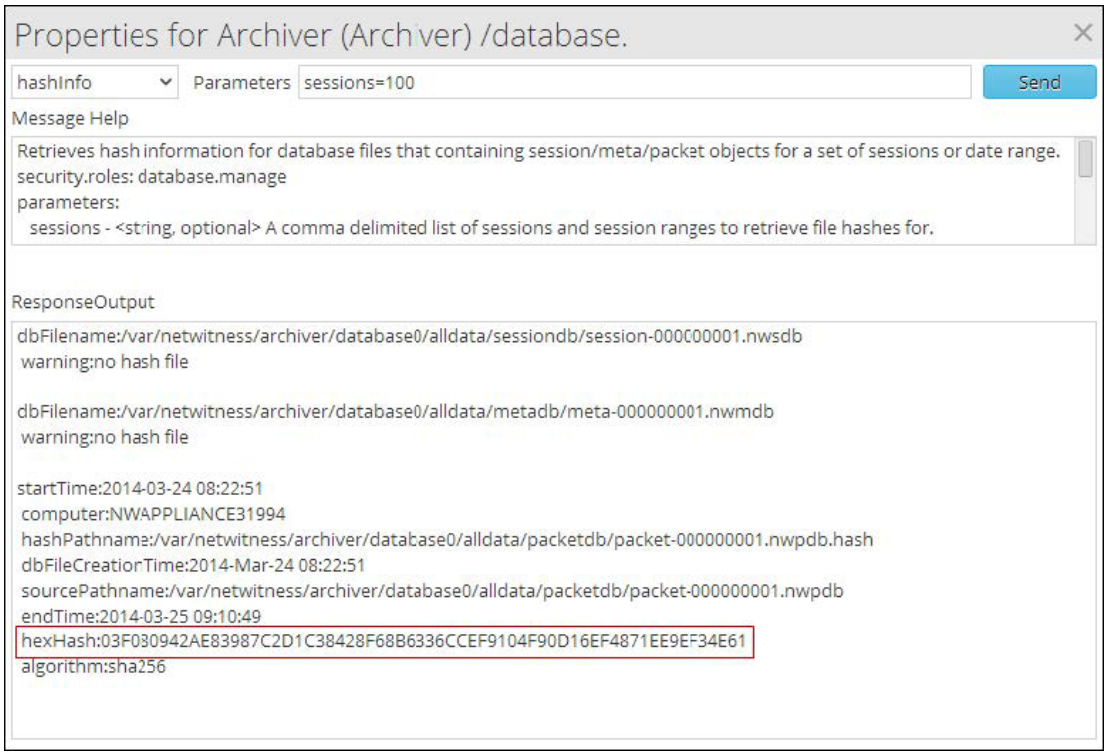
4. In the drop-down menu, select **hashInfo**.
5. In the **Parameters** field, type the criteria that you want to use to retrieve the hash information from the database.
6. Click **Send**.

   The output of the command is displayed in the ReponseOutput textbox. In the output, the hash information is shown in the hexHash parameter. You can use this hash information to verify data integrity manually.

Examples

Retrieve the hash information of the database files for the sessions that exist.
Criteria: sessions=100

Output

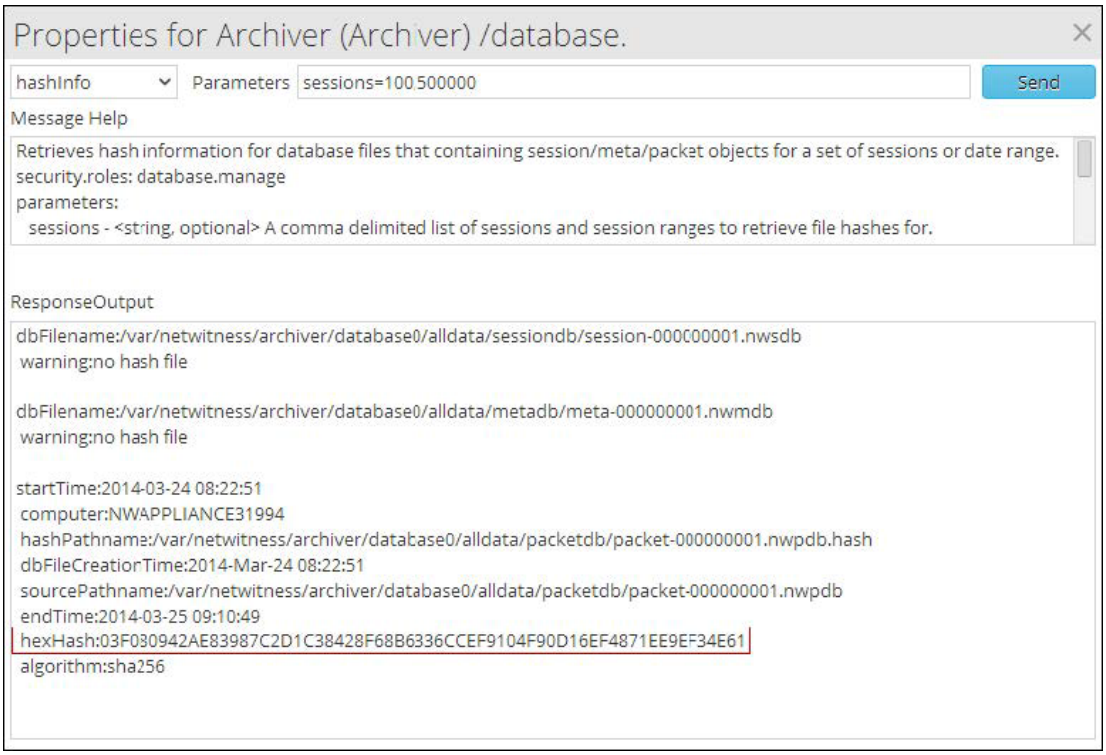


The hash information shown in the hexHash parameter is retrieved and you can use this to verify data integrity manually for session 100.

Retrieve the hash information of the database files for the session ranges that exist.

Criteria: sessions=100,500000

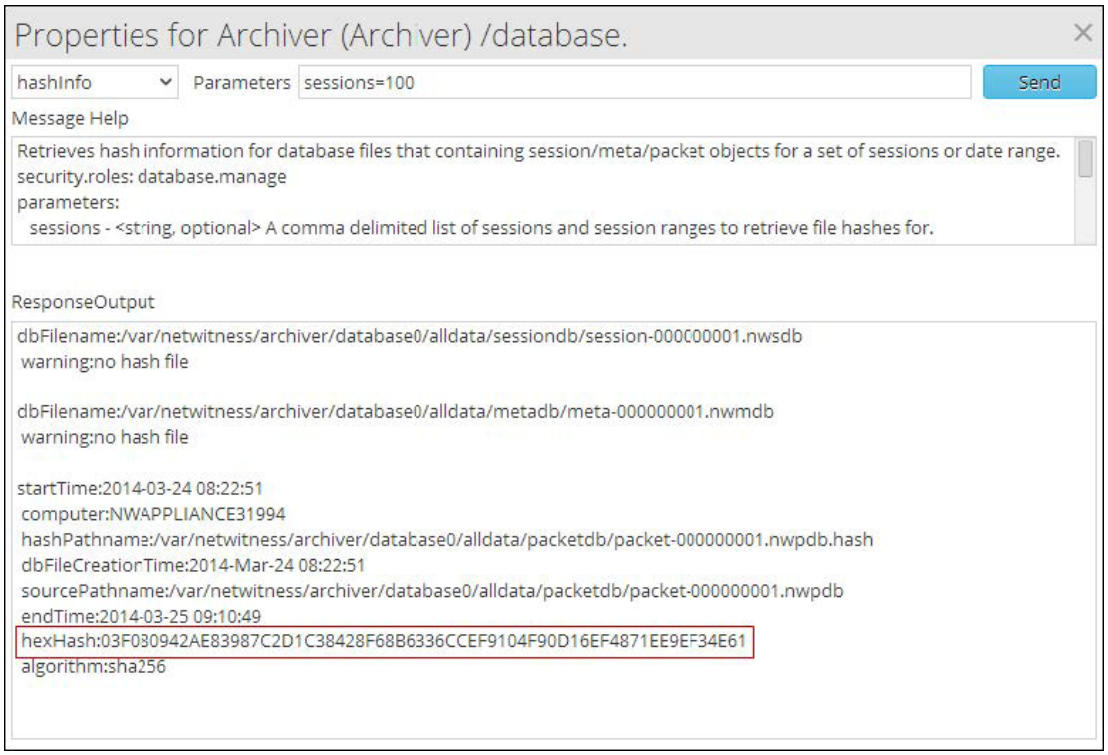
Output



The hash information shown in the hexHash parameter is retrieved and you can use this to verify data integrity manually for session range 100 - 500000

Retrieve the hash information of the database files created in the specified date range
Criteria: beginDate="2017-Mar-25 05:52:15" endDate="2017-Mar-27 05:52:15"

Output



The hash information shown in the hexHash parameter is retrieved and you can use this to verify data integrity manually for the date range specified.

References

This topic is a collection of references, which describe the user interface for Archiver in NetWitness.

Topics

- Archiver Collection Dialog
- Archiver Service Configuration
- Data Retention Tab - Archiver
- Archiver Services Config View - General Tab
- Services Config View - Archiver

Archiver Collection Dialog

On the (Admin) > Services > Config view > Data Retention tab of an Archiver, administrators can define the criteria for log retention and storage. In the Collection dialog, which is accessible from the Collections section, you can define individual storage collections to use for different log types. For example, you may want to create collections for compliance reasons or to selectively retain critical logs.

Workflow

This workflow illustrates the end-to-end installation and configuration process for an Archiver.



What do you want to do?

Role	I want to...	Show me how...
Administrator	Add an Archiver service	Edit the Archiver Service
Administrator	Add a Log Decoderas a source to the Archiver	Add Log Decoder as a Data Source to Archiver
Administrator	*Configure Archiver Storage and Log Retention	Configure Archiver Storage and Log Retention
Administrator	Add an Archiver as a Data Source to the Reporting Engine	Add Archiver as a Data Source to Reporting Engine
Administrator	Configure Archiver Monitoring	Configure Archiver Monitoring

*You can perform this task here.

Related Topics

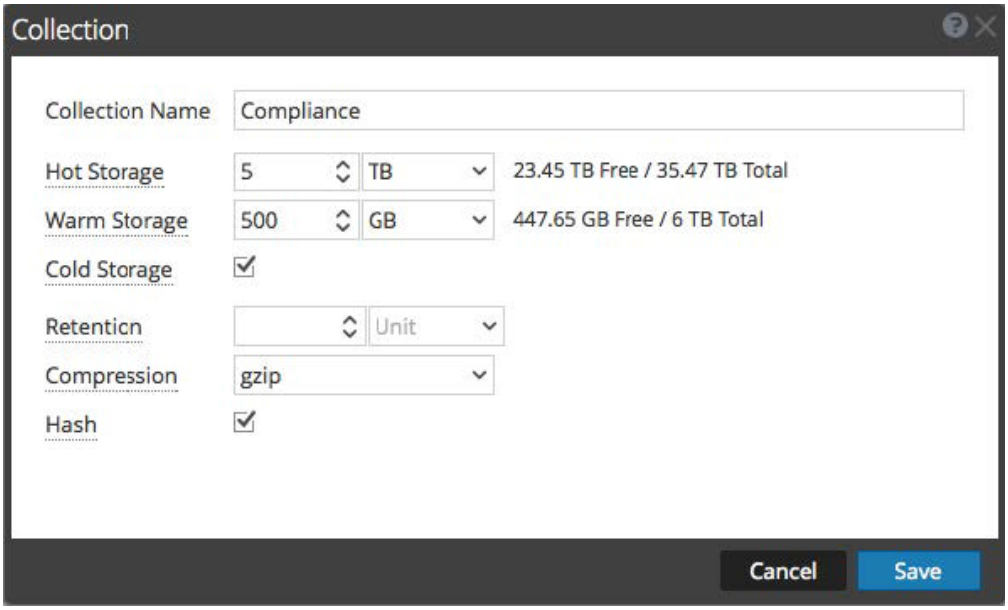
Configuring an Archiver

Quick Look

To access the Collection dialog:

1. Go to **(Admin) > Services**, click , and select **Archiver**.
2. In the Actions column, click **> View > Config**.
   The Services Config view of Archiver is displayed with General tab open.

3. Click the **Data Retention tab**.
4. In the Collections section, click .
   The Collection dialog is displayed.



Note: When decreasing collection storage allocations or lowering retention time, it may take several minutes to hours for the data to move and space to become available depending on the amount of moving (rolling) data. The default times are every 20 minutes for a size roll and every six hours for a time roll.

The following table describes the fields in the Collection Dialog.

Field	Description
Collection Name	Specify a name for your collection, such as Compliance, MediumValue, or LowValue.
Hot Storage	Specify the maximum size or percentage of hot storage to use for this collection. The free space available to use for hot storage and the total hot storage are shown next to this field. When the size of the logs reach the maximum hot storage size, the logs are removed or they roll to the next available storage tier (warm or cold).
Warm Storage	(Optional) Specify the maximum size or percentage of warm storage to use for this collection. The free space available to use for warm storage and the total warm storage are shown next to this field. When the size of the logs reach the maximum warm storage size, the logs are removed or they roll to available cold storage.
Cold Storage	(Optional) Specify whether to use cold storage for this collection. If you use cold storage for the collection, logs outside of the specified size and retention limits roll over to cold storage. If you do not use cold storage, logs outside of the specified size and retention limits are removed.

Field	Description
Retention	(Optional) Specify the number of days that logs are retained before they are removed or rolled over to cold storage. For Hot and Warm Storage, size and retention period settings for a collection can override each other based on which criterion (size or time) is satisfied first.
Compression	Specify the type of compression to use for meta and raw logs in the collection. You can compress the meta and raw logs using GZIP or LZMA to save space. GZIP is very fast at compressing and decompressing, but it does not compress as well as LZMA. LZMA offers better compression at a cost of decompression speed (roughly three times slower than GZIP). Compression ratios are highly dependent on your data. The default compression is GZIP.
Hash	Specify whether to enable or disable hash. When enabled, the hash algorithm is used to verify the data integrity of the files being saved.

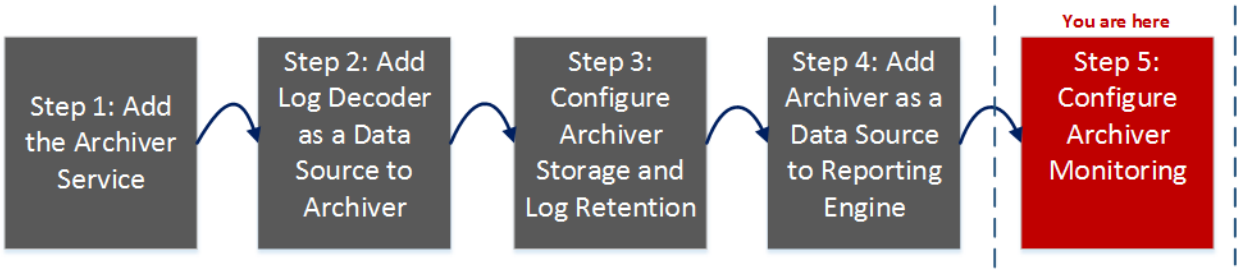
Archiver Services Config View - General Tab

The General tab for an Archiver in the Services Config view helps manage basic service configuration, configure the aggregate service, and configure the aggregation process between an Archiver and the aggregate service.

To access the General tab, go to **(Admin) > Services**, select an **Archiver** service, then select **View > Config**.

Workflow

This workflow illustrates the end-to-end installation and configuration process for an Archiver.



Configuring the aggregate service (whose data is consumed and aggregated) includes:

- Adding, editing, and deleting Archivers as aggregate services
- Toggling an aggregate service online and offline
- Monitoring statistics for aggregate services
- Starting and stopping aggregation

Configuring the aggregation process includes setting:

- Aggregation autostart
- Timing and performance parameters, such as the number of sessions per round of aggregation and time between rounds
- The timing of attempts to restart, reconnect, or take offline a non-responsive aggregate service

Note: Archiver can aggregate from Log Decoder and Archiver (Archiver-Archiver aggregation). Archiver can be aggregated by Archiver and Broker. A Broker can aggregate from either set of concentrators or set of Archivers

What do you want to do?

Role	I want to...	Show me how...
Administrator	Add an Archiver service	Edit the Archiver Service
Administrator	Add a Log Decoder as a Data Source to an Archiver	Add Log Decoder as a Data Source to Archiver

Role	I want to...	Show me how...
Administrator	Configure Archiver Storage and Log Retention	Configure Archiver Storage and Log Retention
Administrator	Add an Archiver as a Data Source to a Reporting Engine	Add Archiver as a Data Source to Reporting Engine
Administrator	*Configure Archiver Monitoring	Configure Archiver Monitoring
Administrator	Start and Stop aggregation Add, edit, delete, and toggle an aggregate service	Aggregate Services Section
Administrator	Manage System Configuration	System Configuration Section

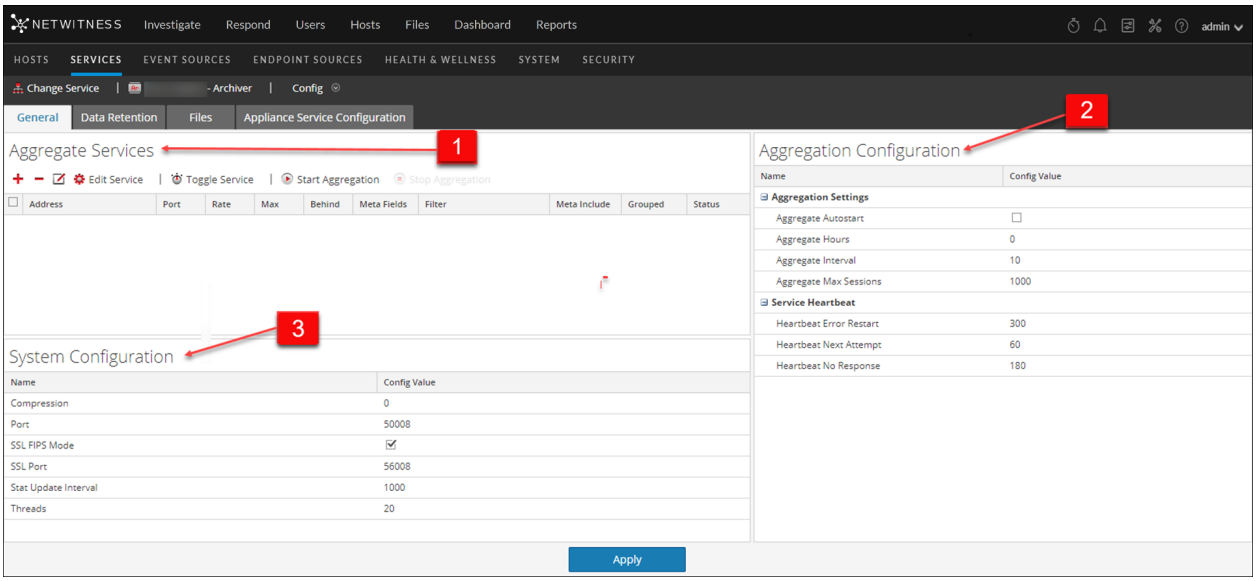
*You can perform this task here.

Related Topics

Configure Log Storage Collections

Quick Look

This is an example of the General tab.

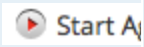
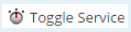


These are the three major sections in the General tab for Archivers:

1	Aggregate Services section provides a way to start and stop aggregation, as well as add, edit, delete, and toggle an aggregate service.
2	System Configuration section manages service configuration for a service.
3	Aggregation Configuration section provides configuration settings that affect various aspects of the aggregation process.

Aggregate Services Section

This is an example of the Aggregate Services section for a Concentrator. The Aggregate Services section toolbar offers these options.

Option	Description
	Opens a dialog in which you can add a Archiver or a Log Decoder.
	Removes the selected aggregate service.
	Opens a dialog to edit **Meta Fields** and **Filter** values.
	When aggregation has been stopped or has not started, starts aggregating data from the online service in the list using the rules defined for the service.
	When aggregation is in progress, stops aggregation on the Broker or Concentrator. This stops all services and flushes the index, which may take several minutes to complete. It is necessary to stop aggregate services in order to perform various administrative procedures.
	Toggles the state of a service between offline and online. Only data from online service is consumed during aggregation.

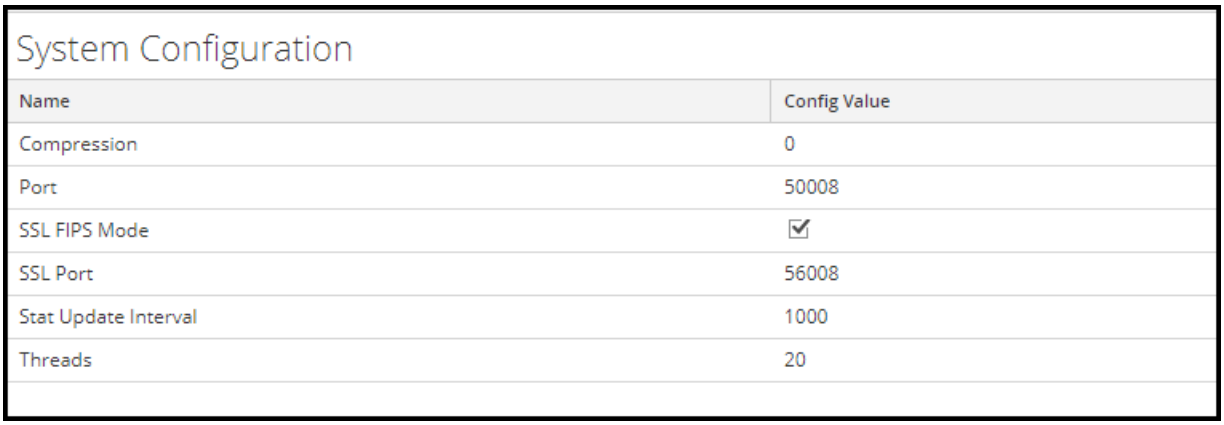
The Aggregate Services section list has these columns.

Column	Description
Address	Lists the address of the service.
Port	Lists the port on which the service listens. The default ports are: • 50001 for Log Collectors • 50002 for Log Decoders • 50003 for Brokers • 50004 for Decoders • 50005 for Concentrators • 50007 for other services
Rate	Lists the number of metadata objects being written to the database per second. Values are rolling average samples over a short time period (10 seconds). After capture stops, the rate is reset to **0**.
Max	Lists the maximum number of metadata objects written to the database per second since capture started. Values are rolling average samples over a short time period (10 seconds). After capture stops, **Max** continues to show the maximum value during capture.
Behind	Lists the number of sessions on the service that need to be aggregated.

Column	Description
Collection	For Brokers only, indicates the collection that was selected when the Analyst Workbench service was added to the Aggregate Services section.
Meta Fields	For Concentrators only, lists the types of metadata being consumed by the aggregate service.
Filter	For Concentrators only, lists any filter being applied to the metadata being consumed by the aggregate service.
Meta Include	For Concentrators only, lists the number of types of meta included in the aggregate service.
Grouped	Whether or not the aggregate service is part of a group.
Status	Lists the current status of the service: • online = available to provide data for consumption by the Broker or Concentrator • offline = not available to provide data for consumption by the Broker or Concentrator • consuming = providing data for consumption by the Broker or Concentrator

System Configuration Section

When a service is first added, default values are in effect. You can edit these values to tune performance.



The System Configuration section has these parameters.

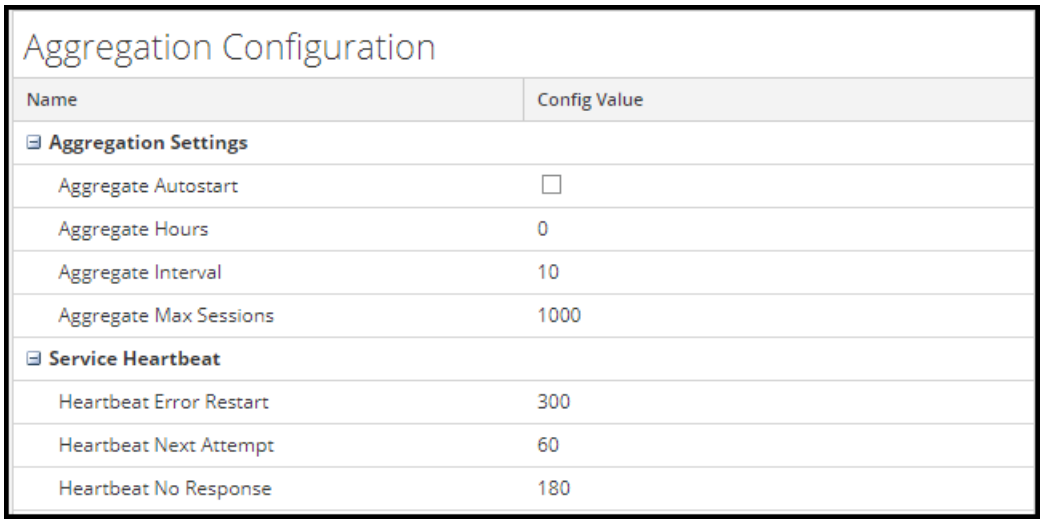
Parameter	Description
Compression	The minimum number of bytes that must be transmitted per response before compression. A setting of **0** disables compression. The default value is **0**. A change in value is effective immediately for all subsequent connections.

Parameter	Description
Port	The port on which the service listens. The default ports are: • 50001 for Log Collectors • 50002 for Log Decoders • 50003 for Brokers • 50004 for Decoders • 50005 for Concentrators • 50007 for other services
SSL FIPS Mode	When enabled (**on**), the security of data transmission is managed by encrypting information and providing authentication with SSL certificates. The default value is **off**.
SSL Port	Indicates the SSL port.
Stat Update Interval	The number of milliseconds between statistic updates on the system. Lower numbers cause more frequent updates and can slow down other processes. The default value is **1000**. A change in value is effective immediately.
Threads	The number of threads in the thread pool to handle incoming requests. A setting of **0** lets the system decide. The default value is **15**. Changes takes effect on service restart.

Aggregation Configuration Section

The Aggregation Configuration Section provides configuration settings for aggregation. When you click **Apply**, the changes are saved; however, not all settings take effect immediately. The tables for Aggregation Settings and Service Heartbeat provide details.

Caution: Do not edit any of these settings without guidance from Customer Support.

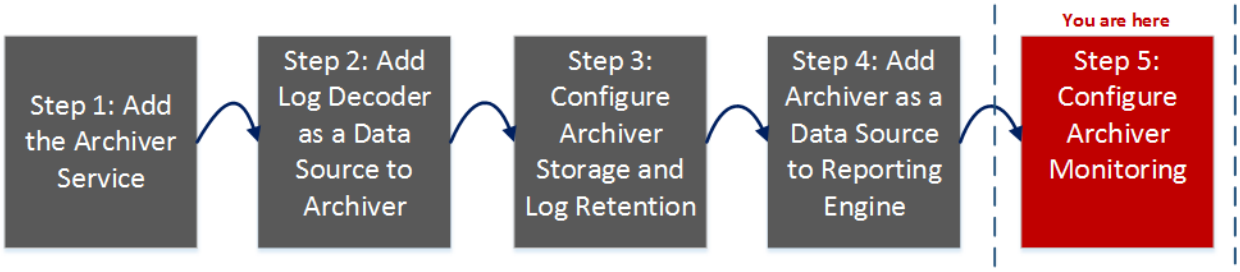


Archiver Service Configuration

This topic lists and describes the available configuration settings for NetWitness Archivers.

Workflow

This workflow illustrates the end-to-end installation and configuration process for an Archiver



What do you want to do?

Role	I want to...	Show me how...
Administrator	Add an Archiver service	Edit the Archiver Service
Administrator	Add a Log Decoder as a Data Source to an Archiver	Add Log Decoder as a Data Source to Archiver
Administrator	Configure Archiver Storage and Log Retention	Configure Archiver Storage and Log Retention
Administrator	Add Archiver as a Data Source to a Reporting Engine	Add Archiver as a Data Source to Reporting Engine
Administrator	*Configure Archiver Monitoring	Configure Archiver Monitoring
Administrator	Configure Archiver settings	/archiver/config
Administrator	Configure Database settings	/database/config
Administrator	Configure Index settings	/index/config
Administrator	Configure Logs settings	/logs/config
Administrator	Configure REST settings	/rest/config
Administrator	Configure SDK settings	/sdk/config
Administrator	Configure Services settings	/services/<service name>/config
Administrator	Configure System settings	/sys/config

*You can perform this task here.

Related Topics

- For more information on configuring Database settings, refer to the "Database Configuration Nodes" topic in the *Core Database Tuning Guide.*
- For more information on configuring Index settings, refer to the "Index Configuration Nodes" topic in the *Database Tuning Guide.*
- For more information on configuring SDK settings, refer to the "SDK Configuration Nodes" topic in the *Database Tuning Guide.*

Data Retention Tab - Archiver

From the (Admin) > Services > Config view > Data Retention tab of an Archiver, administrators can define the criteria for log retention and storage.

Workflow

This workflow illustrates the end-to-end installation and configuration process for an Archiver. From the Data Retention Tab you can configure hot, warm, and cold storage along with configuring multiple storage collections for data retention.



What do you want to do?

Role	I want to...	Show me how...
Administrator	Add the Archiver service.	Edit the Archiver Service
Administrator	Add Log Decoder as a Data Source to an Archiver.	Add Log Decoder as a Data Source to Archiver
Administrator	*Configure Archiver Storage and Log Retention.	Configure Archiver Storage and Log Retention
Administrator	Add an Archiver as a Data Source to Reporting Engine.	Add Archiver as a Data Source to Reporting Engine
Administrator	Configure Archiver Monitoring.	Configure Archiver Monitoring

*You can perform these tasks here.

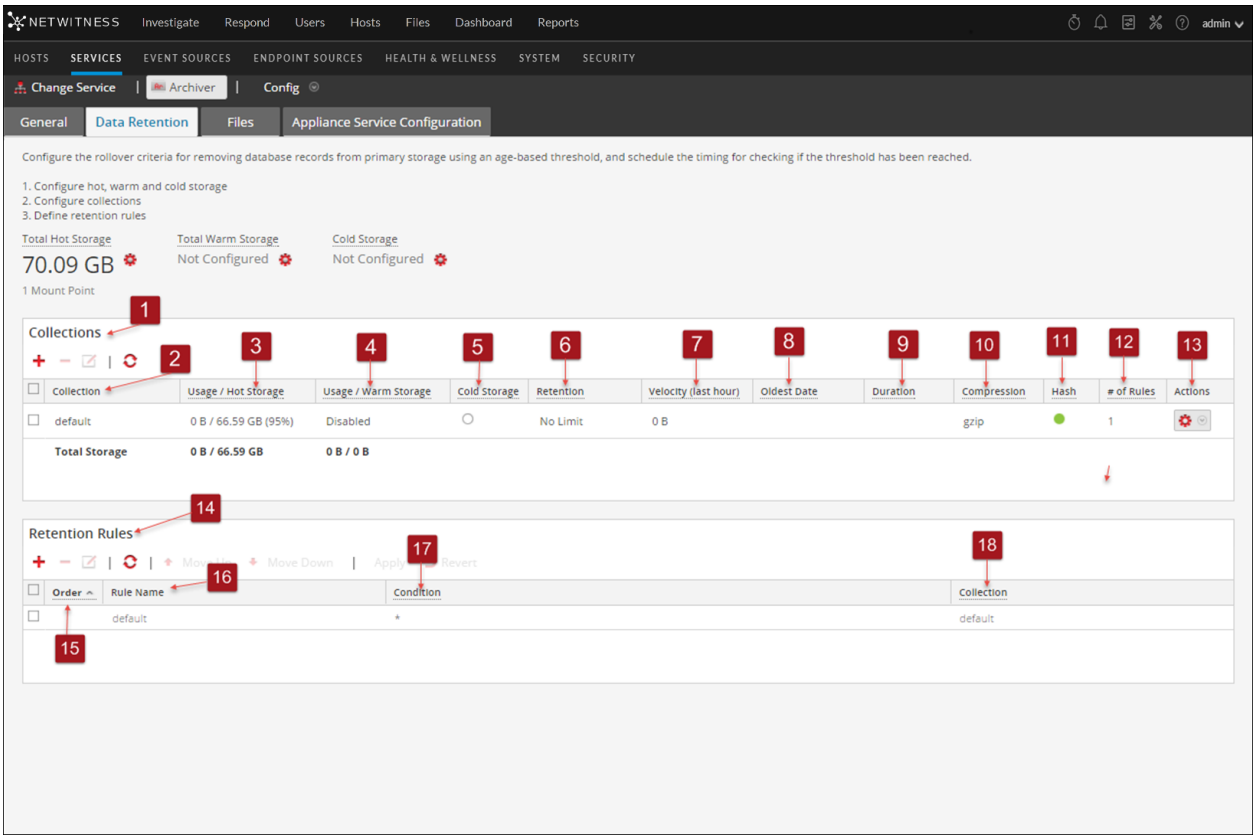
Related Topics

Configuring an Archiver

Configure Hot, Warm, and Cold Storage

Quick Look

As an administrator, you can configure hot, warm, and cold storage as well as multiple storage collections with different locations and criteria for retaining logs. For example, you can create a Compliance collection that stores logs for a specific time period as required by government regulations. You can create another collection that stores low value logs in hot storage with a much shorter retention period. The flexibility of these collections enables you to have significantly less overall storage requirements.



1	Displays the Collections panel with the Data Retention tab open.
2	Allows you to sort the collections in ascending or descending order.
3	Displays the allocated hot storage space for the collection, as well as the approximate current usage.
4	Displays the allocated warm storage space for the collection, as well as the approximate current usage.
5	Displays whether the collection uses cold storage for long-term backup.
6	Displays the time range used to determine when data is moved to cold storage or discarded.
7	Displays the amount of data written to the collection during the past hour.
8	Displays the date of the oldest data stored in the collection.
9	Displays the approximate age of the oldest data stored in the collection.
10	Displays the compression type used in collection storage.
11	Displays whether or not hashes are used when storing data in the collection.
12	Displays the number of retention rules that use this collection for storing data.
13	Displays the Actions drop-down menu.
14	Displays the Retention Rules panel.
15	Displays the order in which Retention Rules are evaluated in the execution chain.
16	Displays the name of the Retention Rule.
17	Data that satisfies this condition is stored in the corresponding collection.
18	Displays the collection used to store the data that satisfies this particular rule condition.

Total Hot, Warm, and Cold Storage

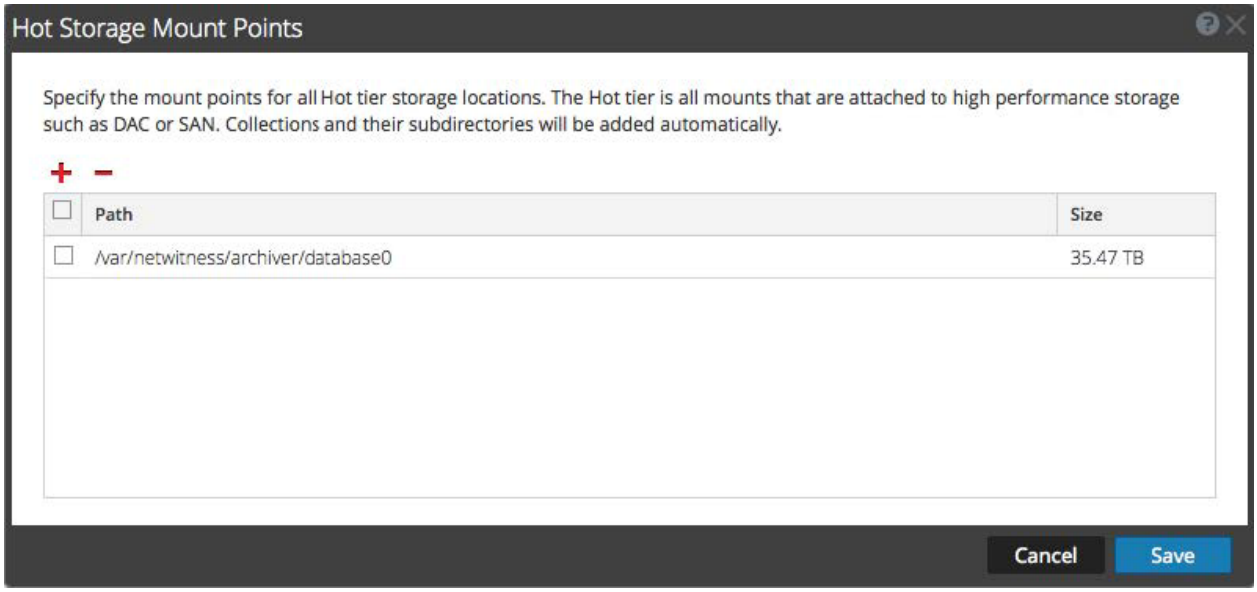
The Total Hot Storage section shows the total amount of Hot storage available and the number of hot storage mount points. The Total Hot Storage section shows the total amount of Warm storage available and the number of warm storage mount points. The Total Cold Storage section shows the total amount of Cold storage and the remaining free space available in Cold storage.



Hot, Warm, and Cold Storage Mount Points Dialogs

In the Hot, Warm, and Cold Storage Mount Points dialogs, you can specify the mount points for your storage locations. You can specify portions of this storage to use for your log storage collections.

To access the Hot, Warm, and Cold Storage Mount Points dialogs, click the  icon near the respective section.

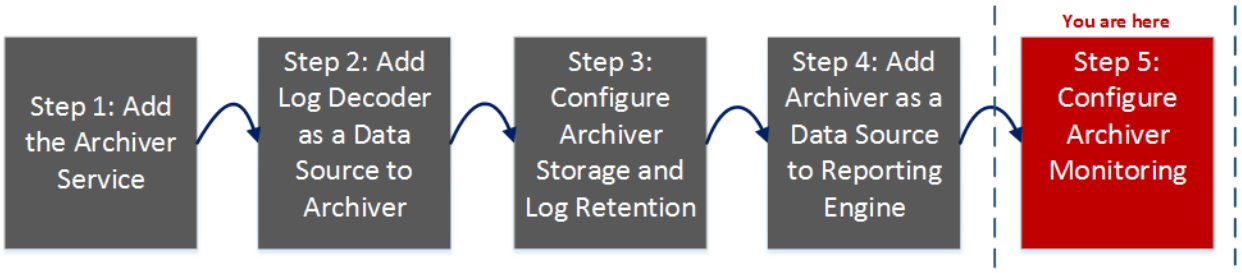


Services Config View - Archiver

The Services Config view ((Admin) > Services > select Archiver service and select >View > Config) provides a way to manage basic service configurations, configure aggregate services, configure log retention and storage, edit service configuration files, and configure the appliance service for an Archiver.

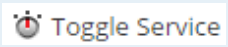
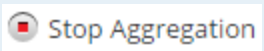
Workflow

This workflow illustrates the end-to-end installation and configuration process for an Archiver.



What do you want to do?

Role	I want to...	Show me how...
Administrator	Add the Archiver service.	Edit the Archiver Service
Administrator	Add a Log Decoder as a Data Source to an Archiver.	Add Log Decoder as a Data Source to Archiver
Administrator	Configure Archiver Storage and Log Retention.	Configure Archiver Storage and Log Retention
Administrator	Add an Archiver as a Data Source to Reporting Engine..	Add Archiver as a Data Source to Reporting Engine
Administrator	Configure Archiver Monitoring.	Configure Archiver Monitoring
Administrator	*Add a Log Decoder as an aggregate service.	Click in the Aggregate Services section.
Administrator	*Remove the selected aggregate service.	Click in the Aggregate Services section.

Role	I want to...	Show me how...
Administrator	*Edit Meta Fields and Filter values of the aggregate service.	Click  in the Aggregate Services section. You can specify the type of metadata that the Archiver consumes from this service. You can also specify a rule to filter data that the Archiver consumes from this service.
Administrator	*Communicate with the Archiver.	Click  in the Aggregate Services section. This enables you to enter the administrator credentials of the selected aggregate service so that it can communicate with the Archiver.
Administrator	*Toggle the state of a service between offline and online.	Click  in the Aggregate Services section.
Administrator	*Aggregate data using the rules defined for the service.	Click  in the Aggregate Services section. Note that it is necessary to start aggregate service after aggregation has been stopped.
Administrator	*Stop aggregation on the Archiver.	Click  in the Aggregate Services section. This stops all services and flushes the index, which may take several minutes to complete. It is necessary to stop aggregate services in order to perform various administrative procedures.

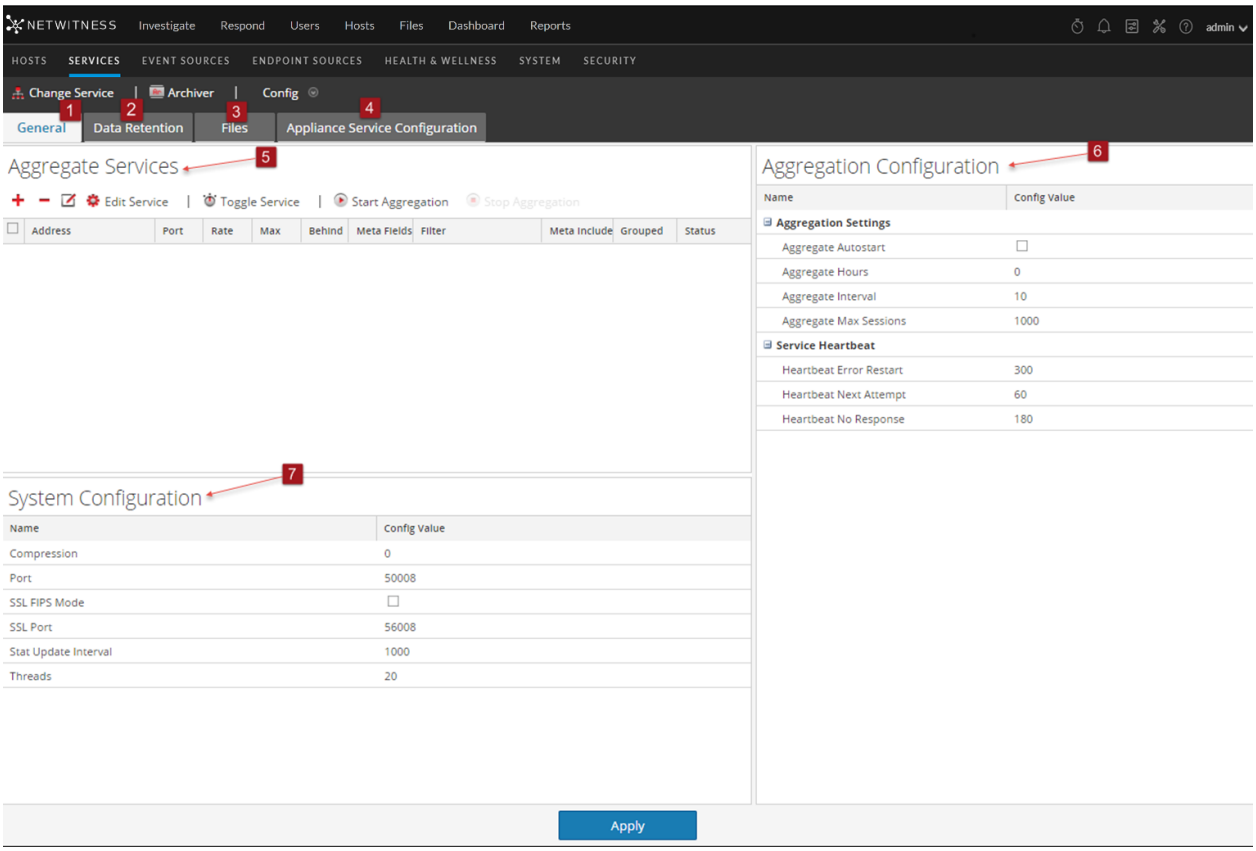
*You can perform this task in the current view.

Related Topics

Configure Log Storage Collections

Quick Look

The Services Config view has four tabs and three panels.



1	General tab provides a way to manage basic Archiver service configuration.
2	Data Retention tab provides a way to view and edit collections and retention rules.
3	Files tab allows you to edit enables you to edit the service configuration files for the Archiver as text files
4	Appliance Service Configuration tab provides a way to configure an Archiver service.
5	Aggregate Services panel provides a way to start and stop aggregation, as well as add, edit, delete, and toggle an aggregate service.
6	Aggregation Configuration panel provides configuration settings that affect various aspects of the aggregation process.
7	System Configuration panel provides a way to manage service configuration for an Archiver service.

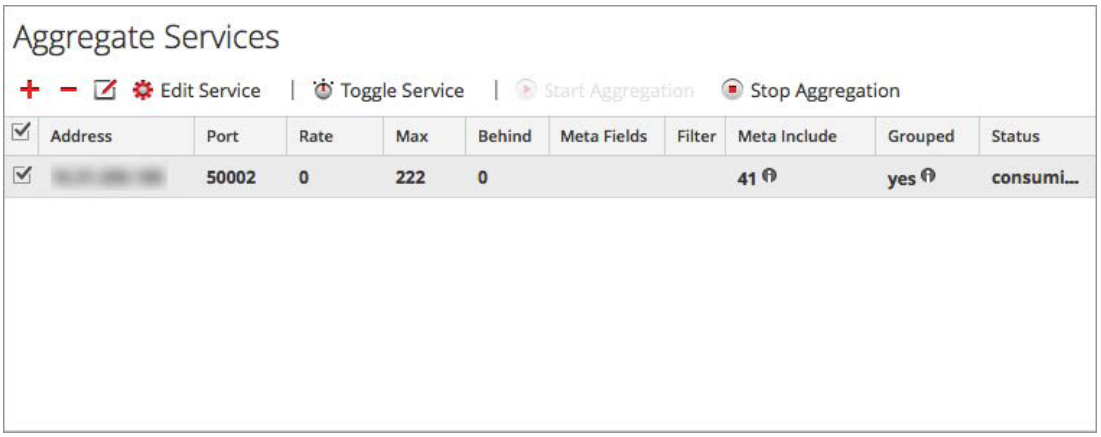
General

The General tab contains the following sections:

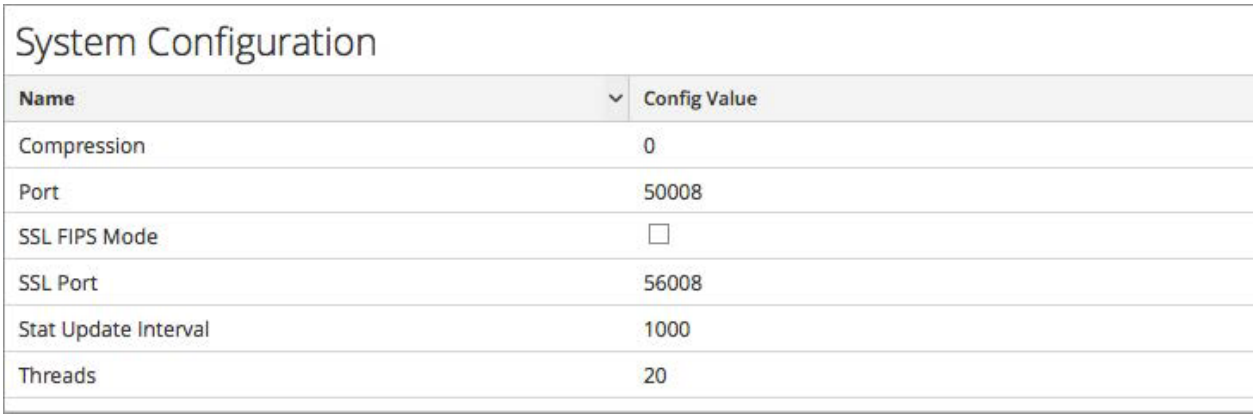
- Aggregate Services
- System Configuration
- Aggregation Configuration

Aggregate Services

The Aggregate Services section provides a way to start and stop aggregation, as well as add, edit, delete, and toggle an aggregate service.



System Configuration

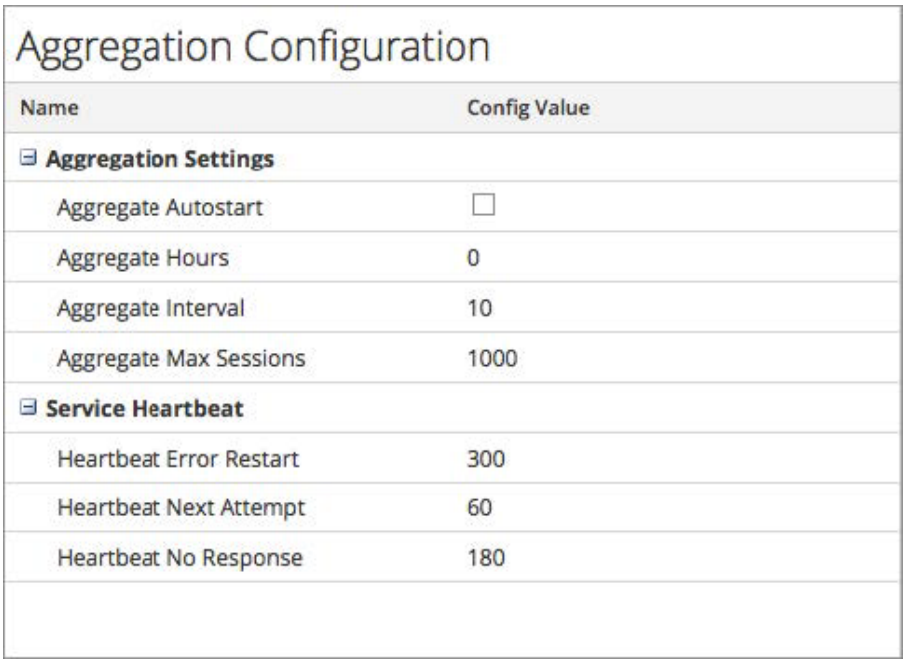


When you add an Archiver service, default values are in effect. NetWitness designed the default values to accommodate most environments and recommends that you do not edit these values because it may adversely affect performance. The following table describes the System Configuration parameters.

Task	Description
Compression	Determines the minimum amount of bytes before a message is compressed. If set to zero, messages are not compressed.
Port	Determines the port used by the service. **Note:** If you change the port number, ensure that you restart the service.
SSL FIPS mode	If enabled, all the data transferred in the network will be encrypted using SSL.
SSL Port	Indicates the port used for encrypting using SSL.

Task	Description
Stat Update Interval	Determines how often (in milliseconds) statistic nodes are updated in the system.
Threads	Determines the number of threads in the thread pool to handle incoming requests.

Aggregation Configuration



The Aggregation Configuration section contains the following sections:

- Aggregation Settings
- Service Heartbeat

Aggregation Settings

The Aggregations Settings section has the following parameters.

Parameter	Description
Aggregate Autostart	If enabled, data aggregation will automatically restart after a service restart.
Aggregate Hours	Determines the maximum number of hours a service is allowed to start aggregation.
Aggregate Interval	Determines the minimum number of milliseconds before another round of aggregation is requested.

Parameter	Description
Aggregate Max Sessions	Determines the number of sessions to aggregate on each round.

Service Heartbeat

The Service Heartbeat section has the following parameters.

Parameters	Description
Heartbeat Error Restart	Determines the number of seconds to wait after a service error before attempting a service reconnect.
Heartbeat Next Attempt	Determines the number of seconds to wait before attempting a service reconnect.
Heartbeat No Response	Determines the number of seconds to wait before taking unresponsive service to offline.

Files

The **Files** tab in the Service Config view enables you to edit the service configuration files for the Archiver as text files. The files available to edit vary depending upon the type of service being configured.

The following files are common to all core services:

- Service index file
- NetWitness file
- Crash reporter file
- Scheduler file
- Feed definitions file

For more information on the **Files** tab, see "Files Tab" in the *Host and Services Getting Started Guide*.

NetWitness® Platform XDR

Version 12.0

AWS Installation Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

AWS Installation Overview **5**

AWS Environment Recommendations **6**

AWS Deployment Scenarios 7

Process 7

NetWitness High-Level Deployment Diagram 8

Prerequisites 8

AWS Deployment **9**

Rules 9

Checklist 9

Establish AWS Environment 9

Find NetWitness Platform AMIs 10

Launch an Instance and Configure a Host 12

Storage Configurations 16

Installation Tasks 16

Task 1 - Install 12.0.0.0 on the NW Server Host and Component Hosts 16

Set Up ESA Hosts 24

Install Component Services on Hosts 25

Complete Licensing Requirements 26

(Optional) Install Warm Standby NW Server 26

(Optional) Boot NetWitness 11.4 EC2 Instance using c5 Instance Types 26

Configure Hosts (Instances) in NetWitness Platform 27

Configure Packet Capture 27

Integrate Gigamon GigaVUE with the Network Decoder 27

Task 1. Integrate the Gigamon Solution 28

Task 2. Configure Tunnel on the Network Decoder 28

Integrate Ixia with the Network Decoder 29

Task 1. Deploy Client Machines 30

Task 2. Create CloudLens Project 30

Task 3. Install Docker Container on Decoder 31

Task 4. Install Docker Container on Clients 31

Task 5. Map the Network Decoder to Ixia Clients 32

Task 6. Validate CloudLens Packets Arriving at Decoder 34

Task 7. Set the Interface in the Network Decoder 35

Integrate f5® BIG-IP with the Network Decoder 36

f5® BIG-IP VE Deployment Information 36

Task 1: Set Up a BIG-IP VE Virtual Server Instance 36
Task 2: Create a Clone Pool 36
Guidelines 37
Troubleshooting Tips 37
Integrate VPC Traffic Mirroring with the Network Decoder 37
Task 1. Configure the Network Decoder as a VPC Traffic Mirroring Destination. 38
Task 2. Configure a VPC Traffic Mirror Filter 38
Task 3. Configure a VPC Traffic Mirror Session 39
Task 4. Set Up a new VXLAN Interface on the Network Decoder 40
Task 5. Validate VPC Traffic Mirroring Packets Arriving at the Network Decoder 41
AWS Instance Configuration Recommendations 43
Archiver 43
Broker 45
Concentrator - Log Stream 46
Packet Stream Solutions 47
Concentrator - Gigamon Solution 47
Concentrator - f5 BIG-IP Solution 47
Decoder - Gigamon Solution 48
Decoder - f5 BIG-IP Solution 48
ESA and Context Hub on Mongo Database 50
Log Collector (Syslog, Netflow, and File Collection Protocols) 51
Log Decoder 52
NW Server, Reporting Engine, Respond and Health & Wellness 53
NetWitness Endpoint Hybrid 54
Appendix A. Silent Installation Using CLI 55

AWS Installation Overview

Before you can deploy NetWitness in the Amazon Web Services (AWS) you need to:

- Review the recommended compute and memory specifications needed for each NetWitness instance.
- Get familiar with the NetWitness Storage Guide to understand the types of drives and volumes needed to support NetWitness instances. For more information, see the Storage Guide for NetWitness® Platform 12.0.
- Make sure that you have a NetWitness Throughput license.

When you are ready to begin deployment, you can purchase any of the following Third-Party solutions for packet capture in AWS. If you engage one of these third-parties, they will assign an account representative and a professional services engineer to you who will work closely with NetWitness Support.

- Gigamon® GigaVUE
- Ixia CloudLens™
- f5® BIG-IP

AWS Environment Recommendations

AWS instances have the same functionality as the NetWitness Azure, virtual, and hardware hosts. NetWitness recommends that you perform the following tasks when you set up your AWS environment.

- Based on the resource requirements of the different components, follow the best practices to use the system and the dedicated storage Elastic Block Store (EBS) Volumes appropriately.
- Make sure that the compute capacity provides a write speed of 10% greater than the required sustained capture and ingest rate for the deployment.
- Build a Concentrator directory for the index database on the Provisioned IOPS SSD.

AWS Deployment Scenarios

Before you can deploy NetWitness you need to:

- Consider the requirements of your enterprise and understand the deployment process.
- Have a high-level picture of the complexity and scope of a NetWitness deployment.

Process

The components and topology of a NetWitness network can vary greatly between installations, and should be carefully planned before the process begins. Initial planning includes:

- Consideration of site requirements and safety requirements.
- Review of the network architecture and port usage.
- Support of group aggregation on Archivers and Concentrators, and virtual hosts.

When updating hosts and services, follow recommended guidelines under the "Running in Mixed Mode" topic in the *NetWitness Host and Services Getting Started Guide*.

You should also become familiar with Hosts, Host Types, and Services as they are used in the context of NetWitness also described in the *NetWitness Host and Services Getting Started Guide*.

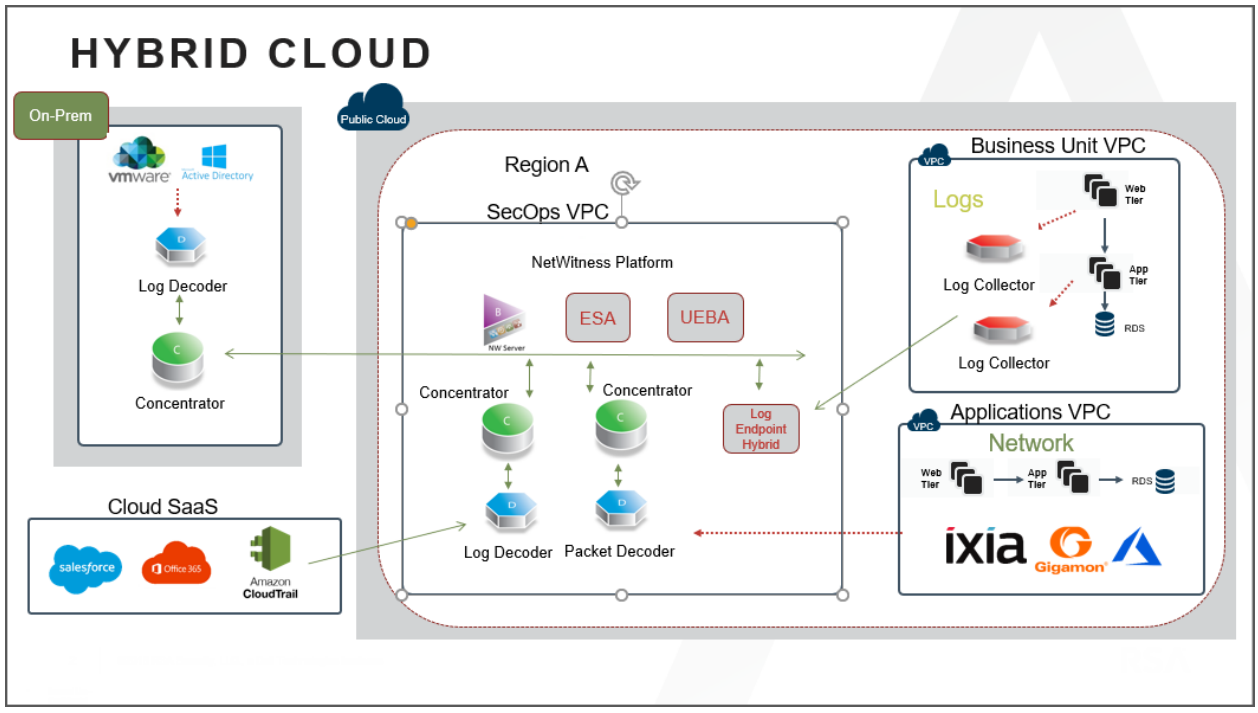
NetWitness High-Level Deployment Diagram

NetWitness is inherently modular. Whether organizations are looking to deploy on-premise or in the cloud, the NetWitness components are decoupled in a way which allows flexible deployment architectures to satisfy a variety of use cases.

The following figure is an example of a hybrid cloud deployment, where the base of the components are residing within the SecOps VPC. Centralizing these components make management easier while keeping network latency to a minimum.

Network, log and endpoint traffic could then be aggregated up to the SecOps VPC. The on-premise location would function just like a normal physical deployment and would be accessible for investigations and analytics.

Cloud SaaS visibility could be captured from a Log Decoder residing in either the cloud or on-premise locations.



Prerequisites

You need the following items before you begin the installation process:

- Ixia account (https://login.ixiacom.com/)
- Access to AWS console
- Network rout-able (and proper AWS Security Groups) for the containers to transfer data to the NetWitness Decoder.

AWS Deployment

This topic contains the rules and high-level tasks that you must follow to deploy NetWitness components in the AWS.

Rules

You must adhere to the following rules when deploying NetWitness in AWS.

- If you reboot the Network Decoder instance, the tunnel is not retained. Create the tunnel on Network Decoder again and restart the Decoder service.
- It is recommended to use private IP addresses when you provision AWS NetWitness instances.

Note: If you assign a public IP to the NW Server Host, update the /etc/nginx/conf.d/nginx.conf configuration file as follows:

```
location /nwrpmrepo
{
alias /var/lib/netwitness/common/repo;
index index.html index.htm;
allow <Subnet-Gateway>/Subnet mask ;
#example
# allow 10.0.0.1/25;
deny all;
autoindex on;
}
```

Checklist

Step	Description	√
1	Establish AWS Environment	
2	Find NetWitness Platform AMIs	
3	Launch an Instance and Configure a Host	
4	Configure Hosts (Instances) in NetWitness Platform	
5	Configure Packet Capture	

Establish AWS Environment

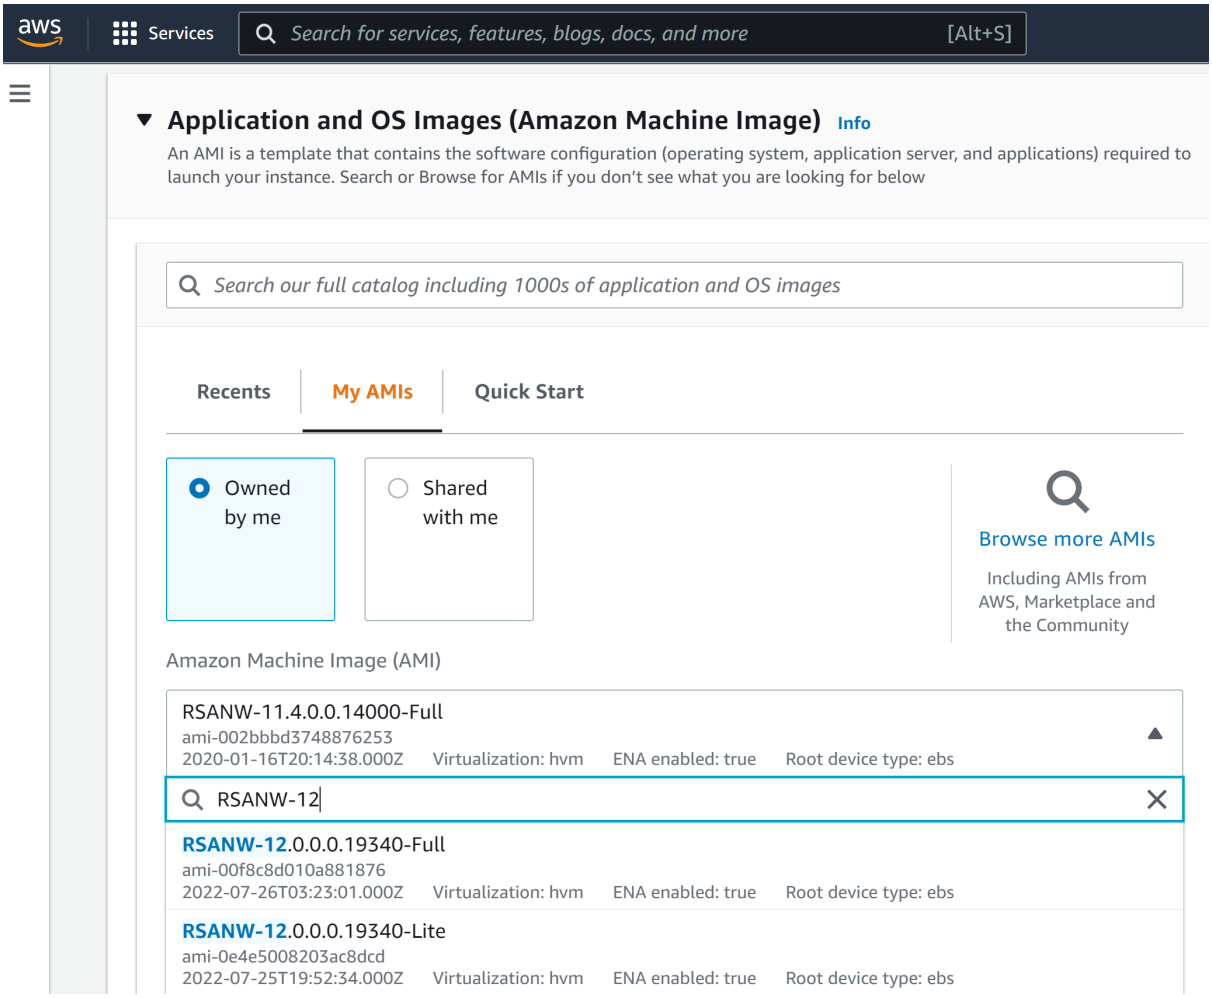
1. Make sure that you have an AWS environment with the capacity to meet or exceed the NetWitness performance guidelines described in AWS Instance Configuration Recommendations.
2. Go to Find NetWitness Platform AMIs.

Find NetWitness Platform AMIs

You can search for NW- AMI files within the Public/Shared/Community repository, using the keyword "RSANW".

Note: For more information, see AWS **Finding Shared AMIs** documentation (http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/usingsharedamis-finding.html).

1. Open the new Amazon EC2 console (New Subscriber Account) at https://console.aws.amazon.com/ec2/.
2. Go to **EC2 Dashboard** and click **Launch Instance**.
   Launch an Instance page is displayed.
3. Under **Application and OS Images (Amazon Machine Image)**, click **My AMIs** and select either **Owned by me** or **Shared with me** option.
4. Search for **RSANW-12** from the **Amazon Machine Image (AMI)** drop-down menu to find the NetWitness AMIs.



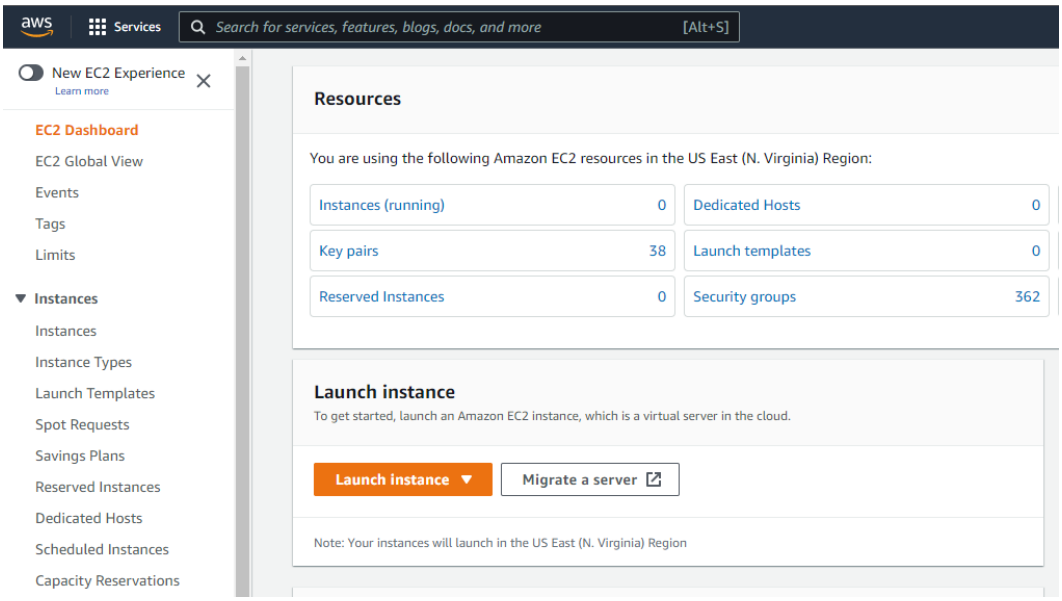
> **Note:** Contact NetWitness Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897) to obtain access to the **RSANW-12.0.0.0.19340-Full**.

5. Go to Launch an Instance and Configure a Host.

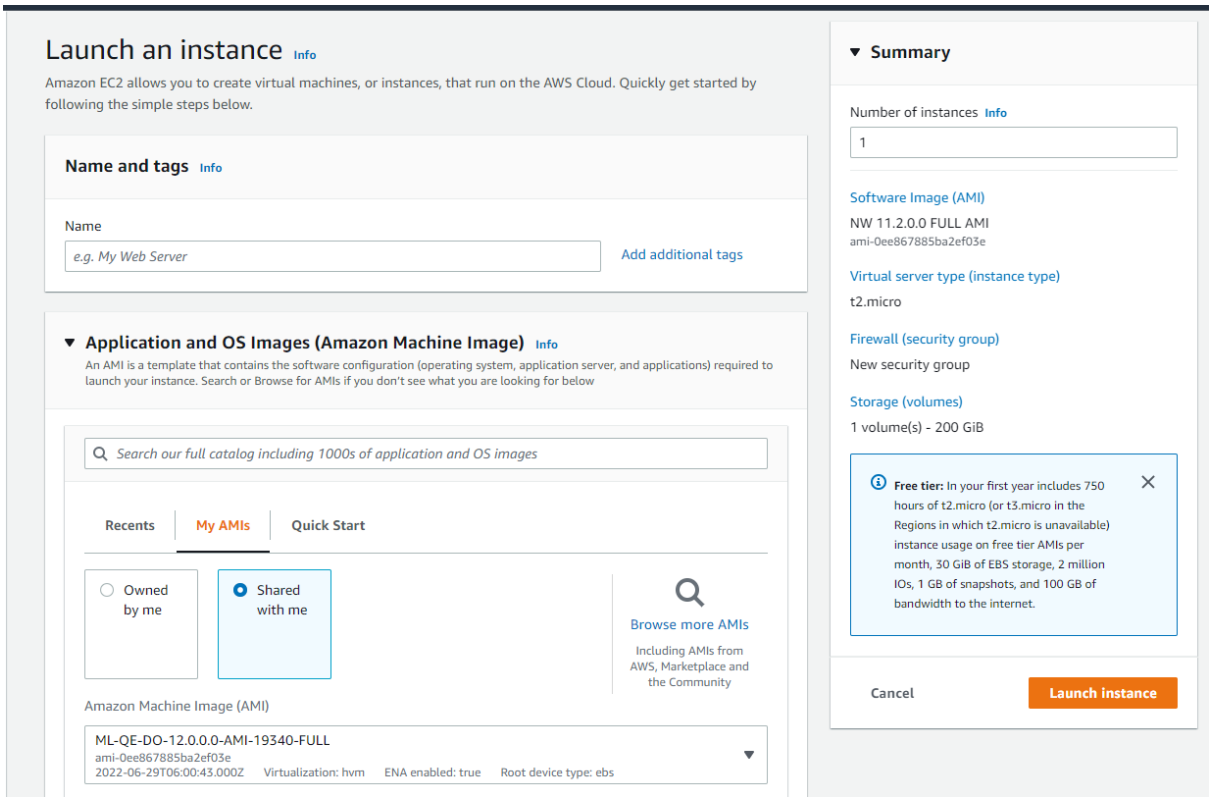
Launch an Instance and Configure a Host

Note: Fore more information, see AWS "Launch an instance using the new launch instance wizard" documentation (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-launch-instance-wizard.html).

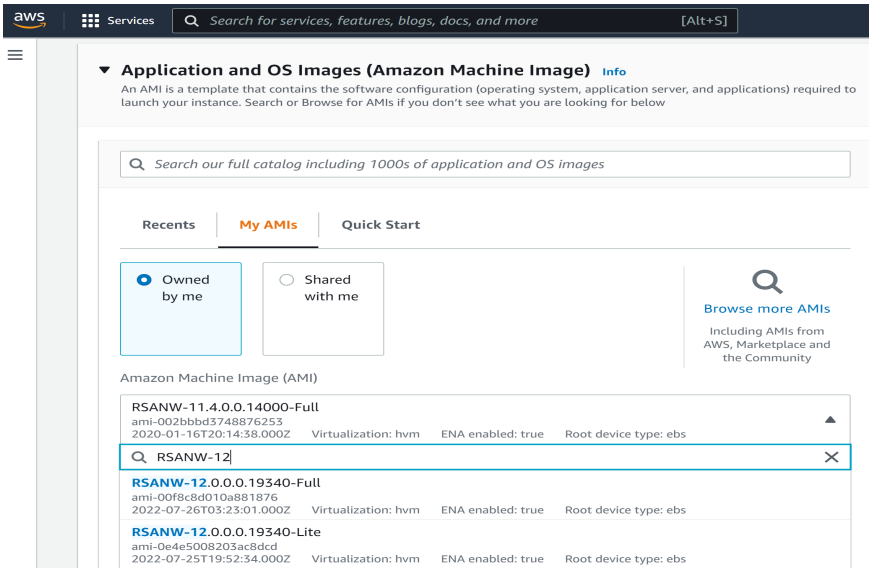
1. Log in to the Amazon EC2 console.
2. Go to **EC2 Dashboard** and click **Launch Instance**.



Launch an Instance page is displayed.

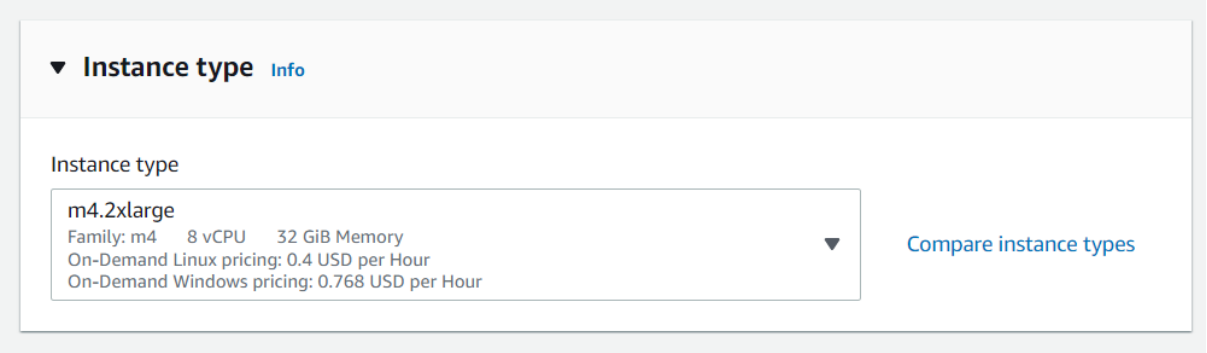


3. (Optional) Under **Name and tags**, enter a name for your instance.
4. Under **Application and OS Images (Amazon Machine Image)**, click **My AMIs** and select either **Owned by me** or **Shared with me** option.
5. Search for **RSANW-12** and select the **RSANW-12.0.0.0.19340-Full** image from the **Amazon Machine Image (AMI)** drop-down menu.

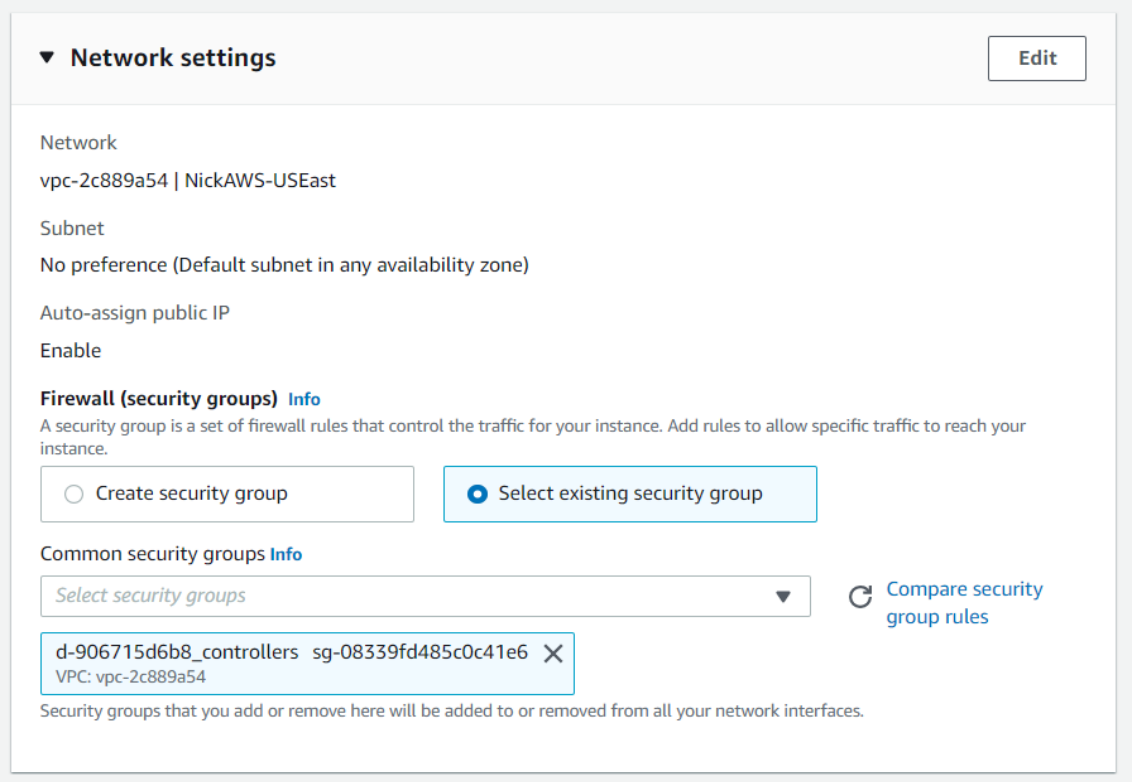


6. Under **Instance type**, Choose the RAM and CPUs by selecting the instance type.
   For more information, see Storage Guide for NetWitness® Platform 12.0 for guidelines on how to

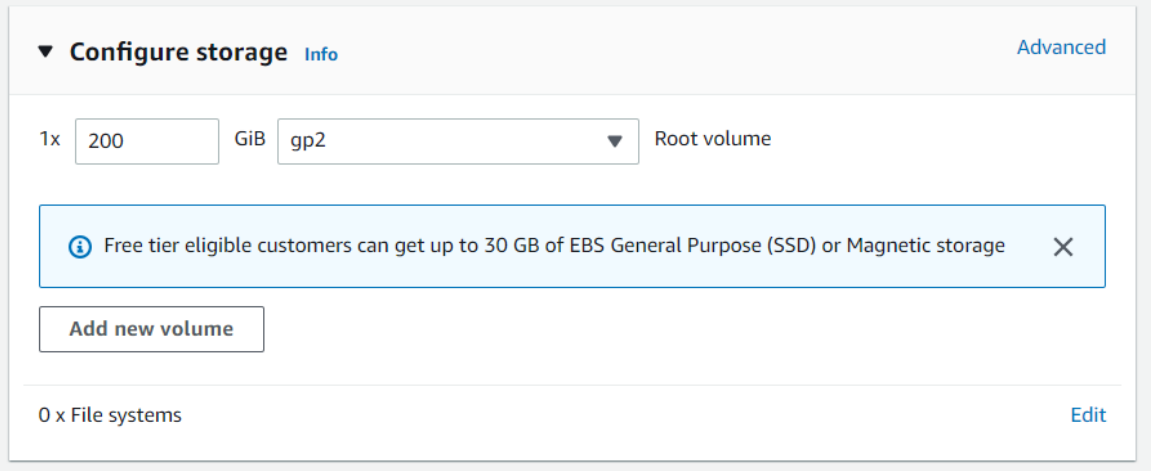
configure the EC2 Instance based on the requirements of the NetWitness component (that is, service) for which you are launching an instance. The following example has the **m4.2xlarge** instance type selected with **8** CPUs and **32** GB of RAM.



7. Under **Network settings**, select **Select existing security grop** option and select your security group from **Common security groups** drop-down menu.

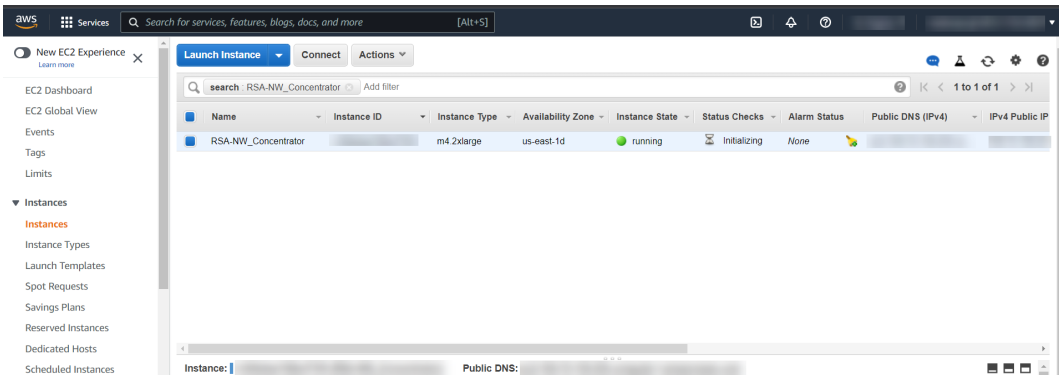


8. Under **Configure Storage** specify the size and type of the volume.

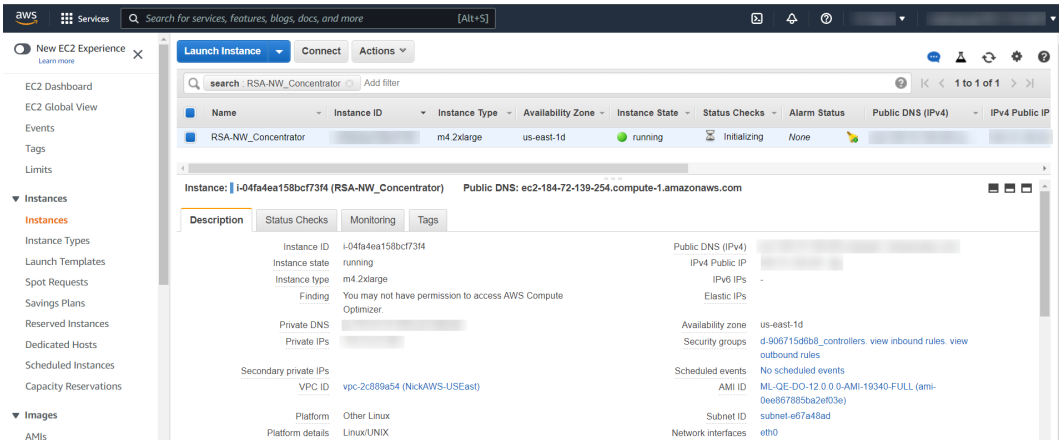


   For more information, see the Storage Guide for NetWitness® Platform 12.0 for guidelines on how to configure storage based on based on the requirements of the NetWitness component (that is, service) for which you are launching an instance.

9. Under **Summary**, you can specify the number of instances required and to review your instance configuration.
10. Click **Launch Instance**.
11. Click **Instances**.
12. Select **Instances** from the left navigation panel to review all instances that AWS is initializing (for example, the **RSA-NW_Concentrator**) .



The IP Address for the new **RSA-NW-Concentrator** host is ***sample-ip-address***.



13. SSH to the newly-created instance using the default NetWitness credentials.
14. Go to Configure Hosts (Instances) in NetWitness Platform .

Storage Configurations

For storage allocations of all host types, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness® Platform 12.0*.

Installation Tasks

Before you begin the installation tasks make sure you open the firewall ports. For more information on the lists of all the ports in a deployment, see the"Network Architecture and Ports" topic in the *Deployment Guide for NetWitness Platform 12.0*.

Caution: Do not proceed with the installation until the ports on your firewall are configured.

Task 1 - Install 12.0.0.0 on the NW Server Host and Component Hosts

Complete the following steps to install 12.0 on NW Server host and other component hosts. Steps that are specific to the NW Server host or to component hosts are noted.

Note: You can perform this task for **RSANW-12.0.0.0.19340-Full** instance.

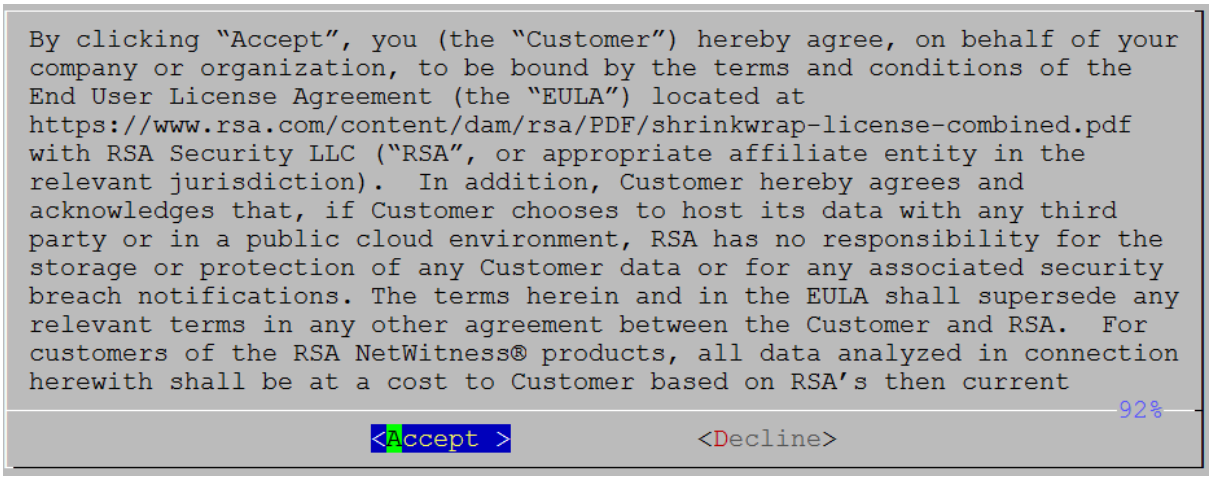
Caution: If you want to install the Endpoint Relay Server, do not run the `nwsetup-tui` script. Follow the instructions in "(Optional) Installing and Configuring Relay Server" in the *NetWitness Endpoint Configuration Guide*.

IMPORTANT: In NetWitness Platform version 11.6 or later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`. If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

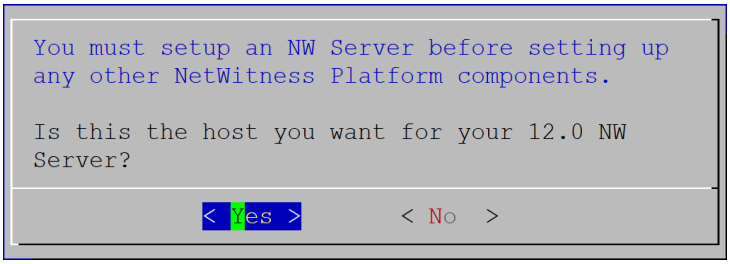
1. Log in to the host with the `root` credentials and run the `nwsetup-tui` command to set up the host.

   This initiates the `nwsetup-tui` (Setup program) and the EULA is displayed.

> **Note:** Use the following options to navigate the Setup prompts.
> 1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, and use the Tab key to move to and from commands (such as **<Yes>**, **<No>**, **<OK>**, and **<Cancel>**). Press **Enter** to register your command response and move to the next prompt.
> 2.) The Setup program adopts the color scheme of the desktop or console you use to access the host.
> 3.) If you specify DNS servers during the Setup program (`nwsetup-tui`) execution, they MUST be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` script to proceed. Any misconfigured DNS servers cause the Setup program to fail. If you need to reach a DNS server after setup that is unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see "Change Host Network Configuration" topic in the System Maintenance Guide.
> If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Update Repository** prompt in step 12 (the DNS servers are not defined so the system cannot access the external repo).



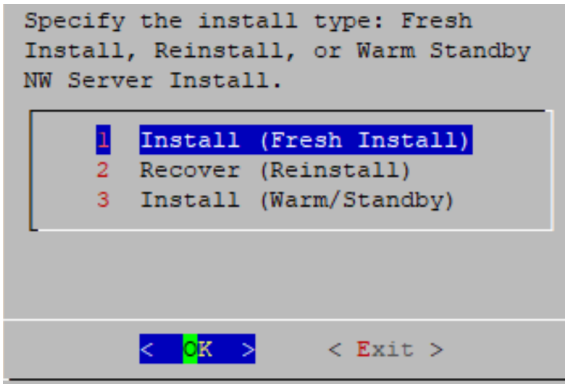
2. Tab to **Accept** and press **Enter**.
   The **Is this the host you want for your 12.0 NW Server** prompt is displayed.



3. Tab to **Yes** and press **Enter** to install 12.0 on the NW Server.
   Tab to **No** and press **Enter** to install 12.0 on other component hosts.

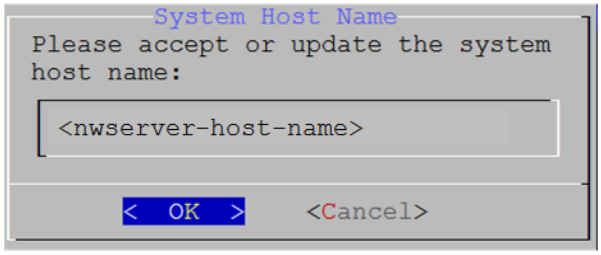
Caution: If you choose the wrong host for the NW Server and complete the Setup, you must restart the Setup Program (step 2) and complete steps all the subsequent steps to correct this error.

4. The **Install** prompt is displayed (**Recover** does not apply to the installation. It is for 12.0 Disaster Recovery.).
   NW Server Host prompt:



   Other Component Hosts, the prompt is the same, but does not include option 3 Install (Warm/Standby)

5. Press **Enter**. **Install (Fresh Install)** is selected by default.
   The **System Host Name** prompt is displayed.
   NW Server prompt:

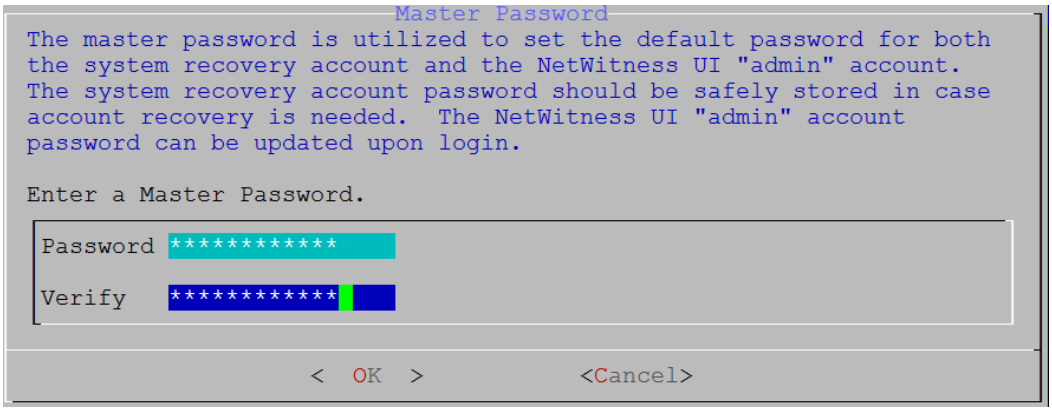


   Other Component Hosts prompt says <non-nwserver-host-name>

   Caution: If you include "." in a host name, the host name must also include a valid domain name.

   Press **Enter** if want to keep this name. If not, edit the host name, tab to **OK**, and press **Enter** to change it.

6. **This step applies only to NW Server hosts**.
   The **Master Password** prompt is displayed.



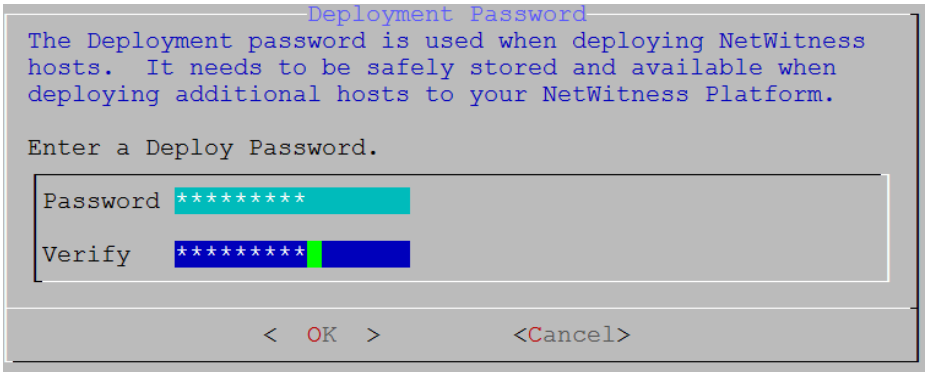
The following list of characters are supported for Master Password and Deployment Password:

- Symbols: **! @ # % ^ +**
- Numbers: **0-9**
- Lowercase Characters: **a-z**
- Uppercase Characters: **A-Z**

No ambiguous characters are supported for Master Password and Deployment Password. For example:
space { } [ ] ( ) / \ ' " ` ~ ; : . < > -

Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

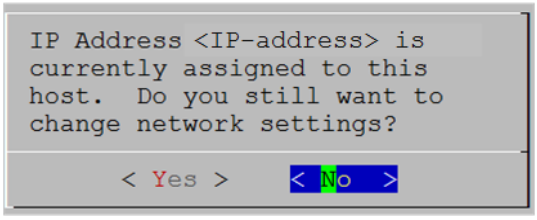
7. **This step applies to both NW Server hosts and component hosts.**
   The **Deployment Password** prompt is displayed.



   Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

8. One of the following conditional prompts is displayed.

   - If the Setup program finds a valid IP address for this host, the following prompt is displayed.

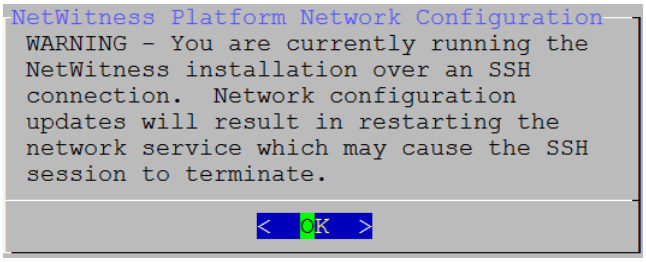


     Press **Enter** if you want to use this IP and avoid changing your network settings. Tab to **Yes** and

press **Enter** if you want to change the IP configuration on the host.

- If you are using an SSH connection, the following warning is displayed.

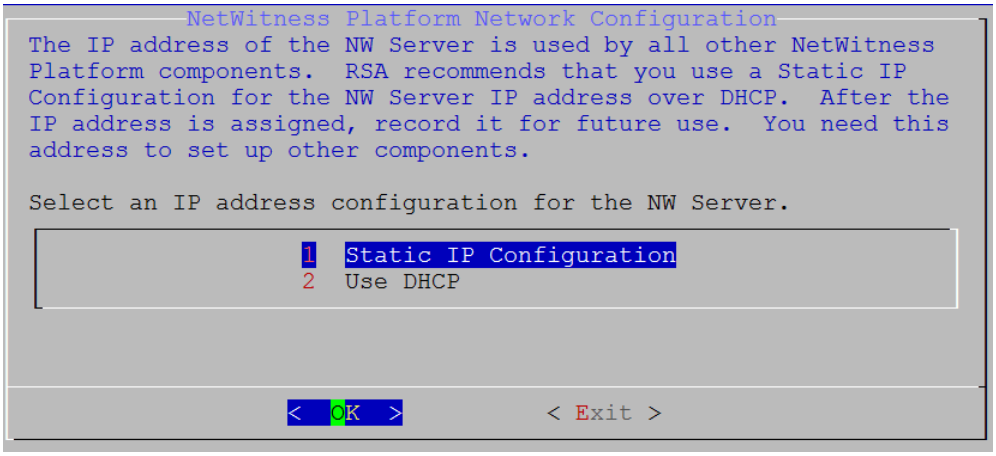
Note: If you connect directly from the host console, the following warning is not displayed.



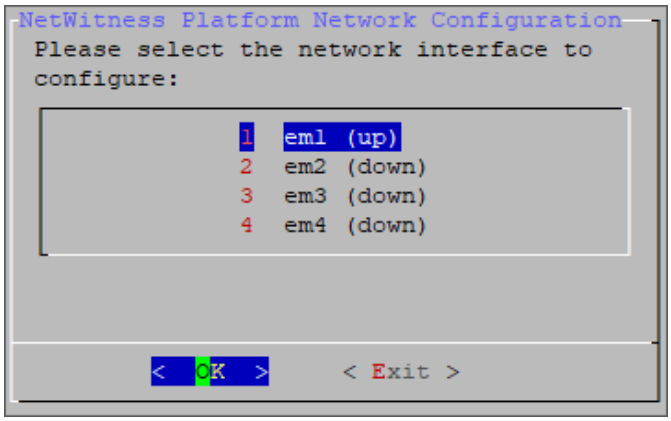
Press **Enter** to close warning prompt.

- If the Setup Program finds an IP configuration and you choose to use it, the **Update Repository** prompt is displayed. Go to step 12 and complete the installation.
- If the Setup Program did not find an IP configuration or if you choose to change the existing IP configuration, the **Network Configuration** prompt is displayed.

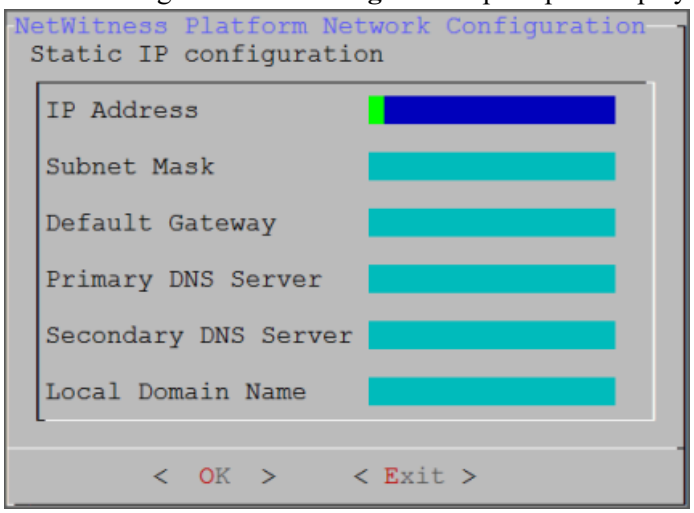
Caution: Only select "**Use DHCP**" as an IP address configuration for the NW Server if DHCP issues static IP addresses.



Tab to **OK** and press **Enter** to use **Static IP**.
If you want to use DHCP, down arrow to **2 Use DHCP** and press **Enter**.
The **Network Configuration** prompt is displayed.



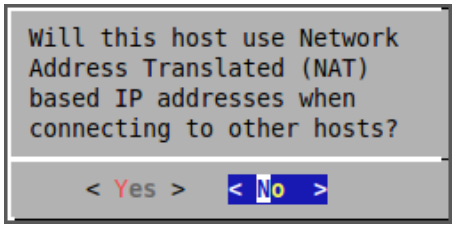
9. Down arrow to the network interface you want, tab to **OK**, and press **Enter**. If you do not want to continue, tab to **Exit.**
   The following **Static IP Configuration** prompt is displayed.



10. Type the configuration values, tab to **OK**, and press **Enter**. If you do not complete all the required fields, an `All fields are required` error message is displayed (**Secondary DNS Server** and **Local Domain Name** fields are not required). If you use the wrong syntax or character length for any of the fields, an `Invalid <field-name>` error message is displayed.

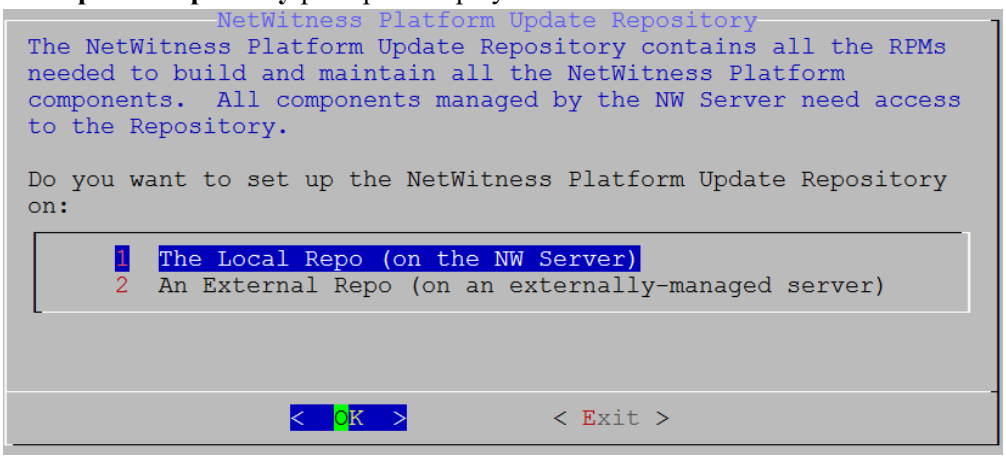
    Caution: If you select **DNS Server**, make sure that the DNS Server is correct and the host can access it before proceeding with the installation.

11. The Use Network Address Translation (NAT) prompt is displayed.



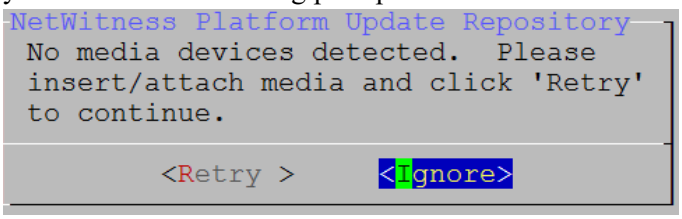
    For the NW Server, tab to **No** and press **Enter**.
    For component hosts, if this host requires the use of NAT-based addresses to communicate with the NW Server, tab to **Yes**. Otherwise, tab to **No** and press **Enter**.

12. The **Update Repository** prompt is displayed.

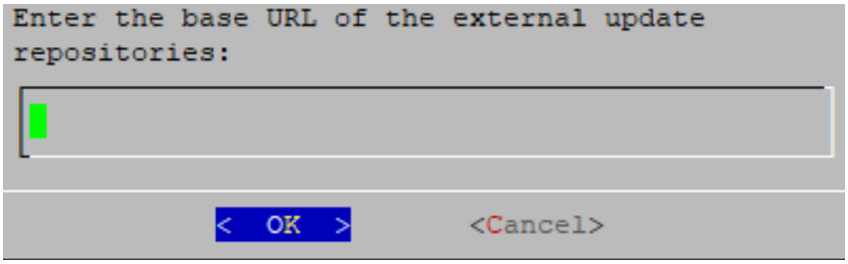


For the NW Server:

- Press **Enter** to choose the **Local Repo**.
- If you want to use an external repo, down arrow to **External Repo**, tab to **OK**, and press **Enter**. If you select **1 The Local Repo (on the NW Server)** in the Setup program, make sure that you have the appropriate media attached to the host (media that contains the ISO file, for example a build stick) from which it can install NetWitness 12.0. If the program cannot find the attached media, you receive the following prompt.



- If you select **2 An External Repo (on an externally-managed server)**, the UI prompts you for a URL. The repositories give you access to RSA updates and CentOS updates. Refer to "Appendix B. Create an External Repo" in this guide for instructions on how to create this repo and its external repo URL so you can enter it in the following prompt.

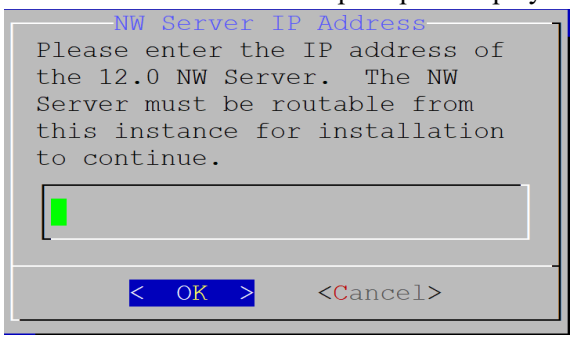


Enter the base URL of the NetWitness external repo and click **OK.** The **Start Install** prompt is displayed.

For component hosts:

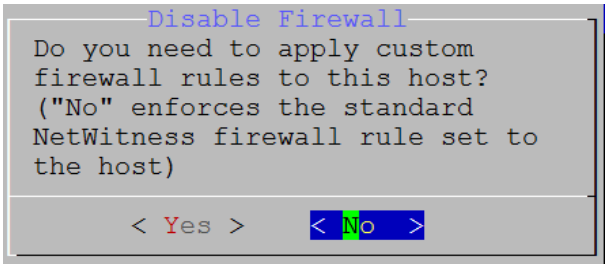
- Select the same repo that you selected when you installed the NW Server host and follow the steps above.

- The NW Server IP Address prompt is displayed.

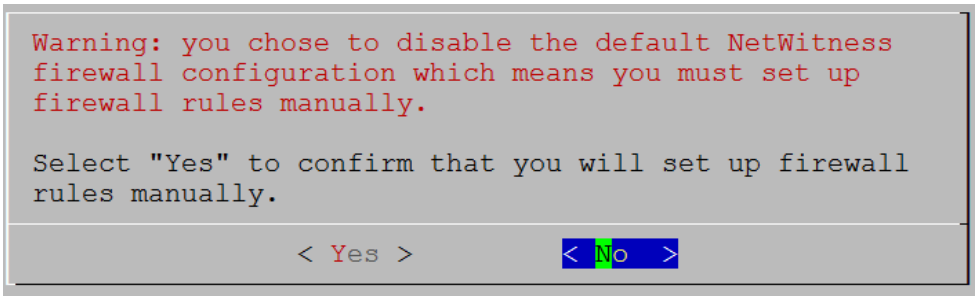


Type the NW Server IP address. Tab to **OK** and press **Enter**.

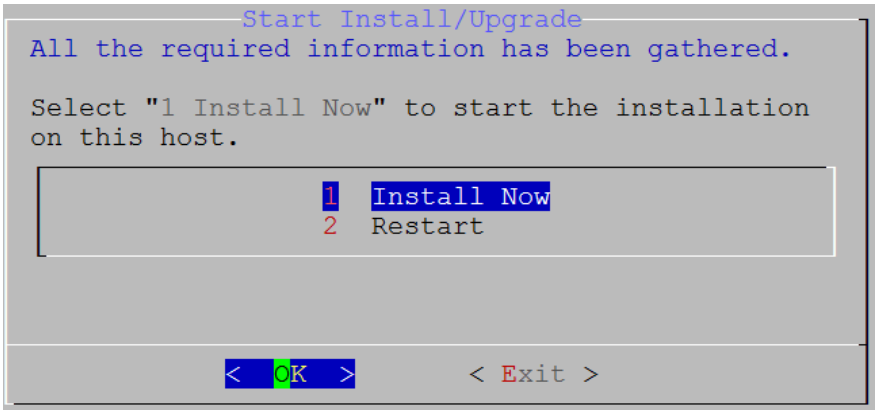
13. The Disable firewall prompt is displayed.



Tab to **No** (default), and press **Enter** to use the standard firewall configuration.
To disable the standard firewall configuration, tab to **Yes**, and press **Enter**.
If you select **Yes**, confirm your selection(select **Yes** again) or select **No** to use the standard firewall configuration.



14. The **Start Install** prompt is displayed.



15. Press **Enter** to install 12.0.
    When **Installation complete** is displayed, you have installed 12.0 on this host.

    Note: Ignore the hash code errors similar to the errors shown in the following figure that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.

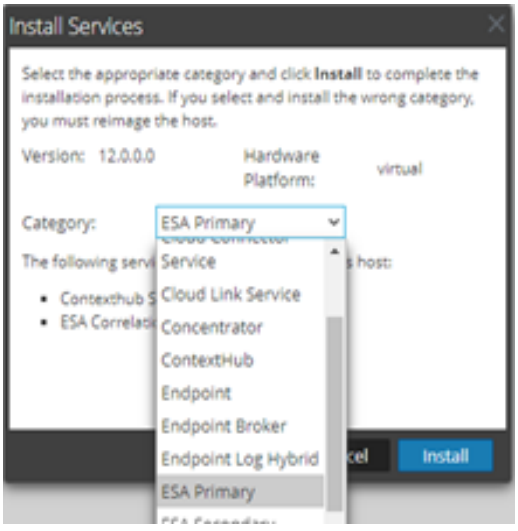
```
ValueError: error:3207A06D:lib(50):B_HASH_init:cr new
Checksum type 'md5' disabled
 (skipped due to only_if)
    * file[/etc/yum.repos.d/CentOS-Base.repo] action delete (up to date)
    * ruby_block[yum-cache-reload-CentOS-Base] action nothing (skipped due to action :nothing)
     (up to date)
  * yum_repository[Remove CentOS-CR repository] action delete
    * execute[yum clean all CentOS-CR] action runERROR:root:code for hash md5 was not found.
Traceback (most recent call last):
  File "/usr/lib64/python2.7/hashlib.py", line 129, in <module>
    globals()[__func_name] = __get_hash(__func_name)
  File "/usr/lib64/python2.7/hashlib.py", line 98, in __get_openssl_constructor
    f(usedforsecurity=False)
```

16. (Optional) If your system configuration requires that a component host must use a NAT IP address to reach the NW Server host, you must configure the NAT IP address of the NW Server by running the following command:
    `nw-manage --update-host --host-id <NW Server Host UUID> --ipv4-public <NAT IP address>`

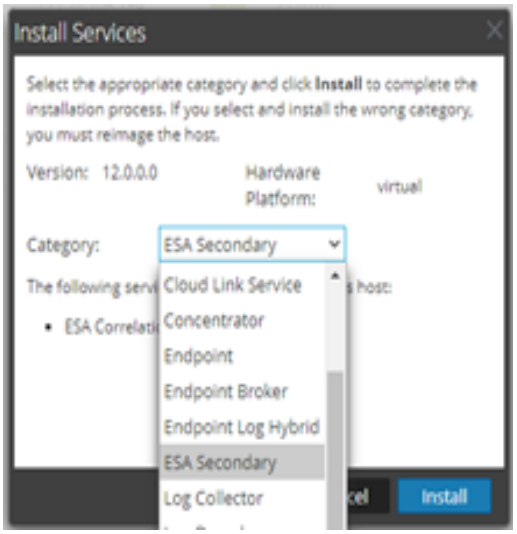
Set Up ESA Hosts

After you install your NW Server and component hosts, follow these steps to set up your ESA hosts.

- Install your primary ESA host following the instructions in "Install 12.0 on the NetWitness Server (NW Server) Host and Other Component Hosts" in this guide, and install the **ESA Primary** service on it after you finish the Set Up program in the UI in (Admin) > **Hosts** > Install:



- (Conditional) If you have a secondary ESA host, install it and install the **ESA Secondary** service on it after you finish the Set Up program in the UI in (Admin) > **Hosts** > Install:



Install Component Services on Hosts

After you have installed NW Server and component hosts, and set up your ESA hosts, follow these steps to install component services, such as Decoders and Concentrators, on your host systems.

1. Install a component service on the host.

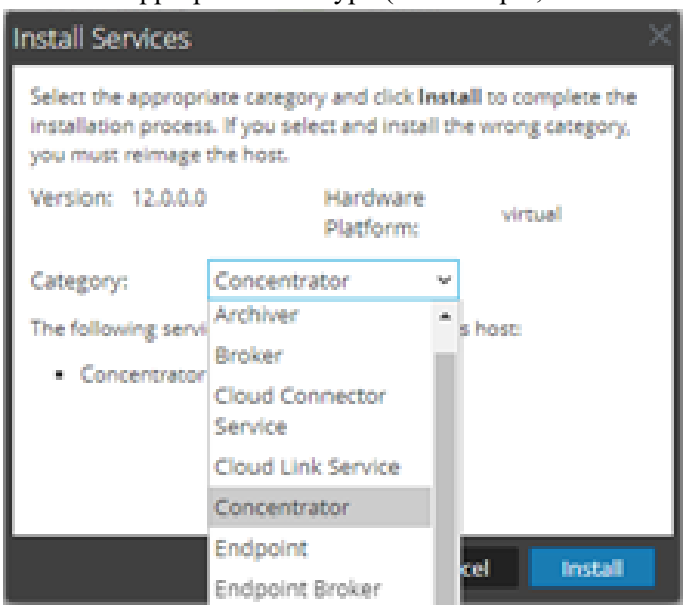
   a. Log into NetWitness and go to (Admin) > **Hosts**.
   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background.

   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the host in the **New Hosts** dialog and click **Enable**.
   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the **Hosts** view and click Install.
   The **Install Services** dialog is displayed.

d. Select the appropriate host type (for example, **Concentrator**) in **Category** and click **Install**.



Complete Licensing Requirements

Complete licensing requirements for installed services. See the *NetWitness Platform 12.0 Licensing Management Guide* for more information. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

(Optional) Install Warm Standby NW Server

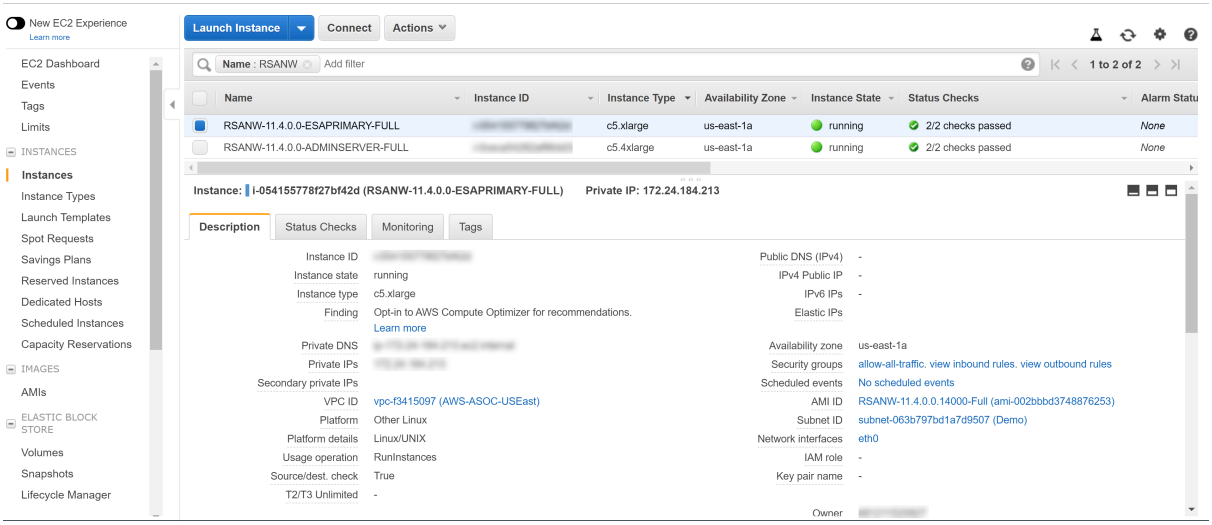
Refer to "Warm Standby NW Server Host" under "Deployment Option Setup Procedures" in the *Deployment Guide for NetWitness Platform 12.0* for instructions on how to set up a Warm Standby NW Server.

(Optional) Boot NetWitness 11.4 EC2 Instance using c5 Instance Types

Following are the steps:

1. Deploy an EC2 instance using NW 11.4 Lite/Full AMI (RSANW-11.4.0.0.14000-Full, ami-002bbbd3748876253).
2. Select a non-nitro based instance type. For example, "_m4.large_" and boot the instance.
3. Install the package "_dracut-config-generic_" and rebuild the initramfs image.
4. For environment without internet connection, copy the provided RPM package to target 11.4 EC2 instance and install it.
5. Shutdown or stop the EC2 instance. Ensure that the instance shutdown behavior is set to "Stop" and not "_Terminate._".

6. Change the instance type of existing EC2 instance to nitro-based c5 types (For example, c5.4xlarge) and start the EC2 instance.
7. The EC2 instance boots successfully now, passing the system and instance status health checks.



Note: Once the instance is moved to c5 instance types and booted successfully, the system will remain fixed in c5 instance state and cannot be reverted back to m4 instance types.

Configure Hosts (Instances) in NetWitness Platform

Configure individual hosts and services as described in NetWitness *Host and Services Getting Started Guide*. This guide also describes the procedures for applying updates and preparing for version upgrades.

Note: After you successfully launch an instance, AWS assigns a default hostname to it. For more information, see "Change Host Network Configuration" in the *System Maintenance Guide* for instructions on changing a hostname.Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Configure Packet Capture

You can integrate any of the following Third-Party solutions with the Network Decoder to capture packets in the AWS cloud:

- Gigamon® GigaVUE
- Ixia CloudLens™
- f5® BIG-IP
- VPC Traffic Mirroring

Integrate Gigamon GigaVUE with the Network Decoder

There are two main tasks to configure the Gigamon® third-party Tap vendor packet capture solution:

Task 1. Integrate the Gigamon® solution.

Task 2. Configure a tunnel on Network Decoder

Task 1. Integrate the Gigamon Solution

Gigamon® Visibility Platform on AWS will be available through the AWS Marketplace and activated by a BYOL license. A thirty-day free trial is also available.

For more information on the Gigamon® solution, see "Gigamon® Visibility Platform for AWS Data Sheet"https://www.gigamon.com/sites/default/files/resources/datasheet/ds-gigamon-visibility-platform-for-aws-4095.pdf.

For more information on the deployment details, see "Gigamon® Visibility Platform for AWS Getting Started Guide" https://www.gigamon.com/sites/default/files/resources/deployment-guide/dg-visibility-platform-for-aws-getting-started-guide-4111.pdf.

After the “Monitoring Session” is deployed within the Gigamon GigaVUE-FM, you can configure the Network Decoder Tunnel.

Task 2. Configure Tunnel on the Network Decoder

1. SSH to the Decoder.
2. Submit the following command strings.

```
$ sudo ip link add tun0 type gretap local any remote <ip_address_of_
VSERIES_NODE_TUNNEL_INTERFACE> ttl 255 key 0

$ sudo ip link set tun0 up mtu <MTU-SIZE>

$ sudo ifconfig (to verify if the tunnel tun0 is being listed in the list
of interfaces)

$ sudo lsmod | grep gre ( to make sure if the below kernel modules are
running:

ip_gre 18245 0

ip_tunnel 25216 1)

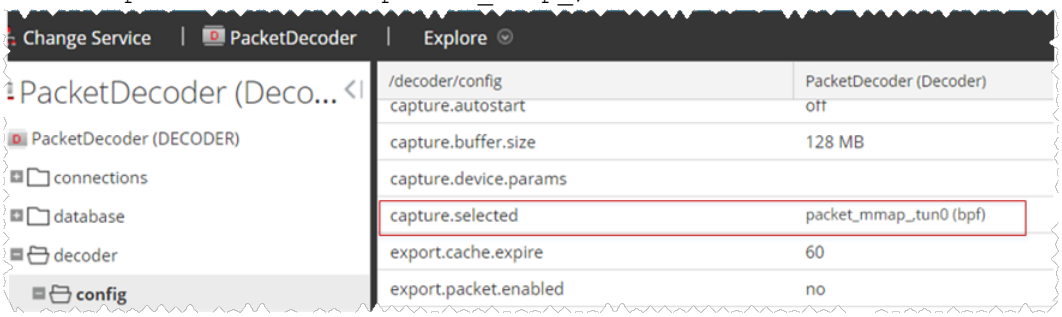
If they are not running then execute the below commands to enable the
modules

$ sudo modprobe act_mirred

$ sudo modprobe ip_gre
```

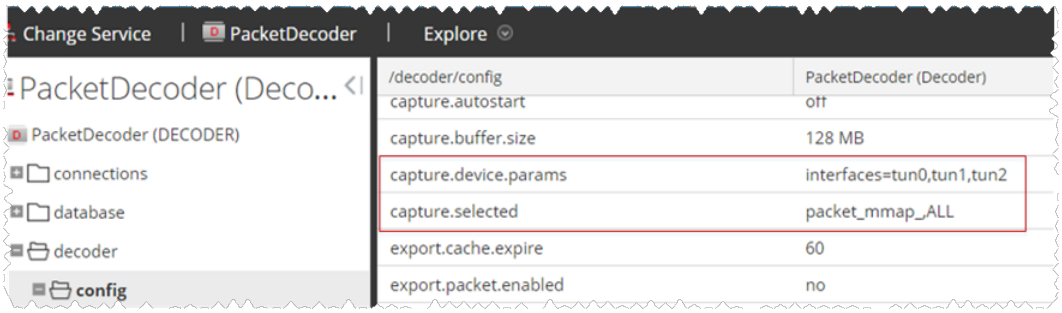
3. Create a firewall rule in the Network Decoder to allow traffic through the tunnel.
   a. Open the `iptables` file.
      `vi /etc/sysconfig/iptables`
   b. Append the line `-A INPUT -p gre -j ACCEPT` before the `commit` statement
   c. Restart `iptables` by executing the following commands.
      `service iptables restart`
      `service ip6tables restart`

4. Set the interface in the Network Decoder.
   a. Log in NetWitness, select the `decoder/config` node in Explorer view for the Network Decoder service.
   b. Set the `capture.selected = packet_mmap_,tun0`.



5. (Conditional) - If you have multiple tunnels on the Network Decoder.
   a. Restart Decoder service after you create the tunnel in Network Decoder.
   b. Log in to NetWitness, select the `decoder/config` node in Explorer view for the Network Decoder service, and set the following parameters.

   ```
   capture.device.params = interfaces=tun0,tun1,tun2
   capture.selected = packet_mmap_,All
   ```



6. Restart decoder service.

   ```
   $ sudo restart nwdecoder
   ```

   The user should be all set to capture the network traffic in Decoder.

Integrate Ixia with the Network Decoder

You must complete the following tasks to integrate the Network Decoder with Ixia CloudLens.

Task 1. Deploy Client Machines

Task 2. Create CloudLens Project

Task 3. Install Docker Container on Decoder

Task 4. Install Docker Container on Clients

Task 5. Map Network Decoder to Ixia Clients

Task 6. Validate CloudLens Packets Arriving at Decoder

Task 7. Set Interface in Network Decoder

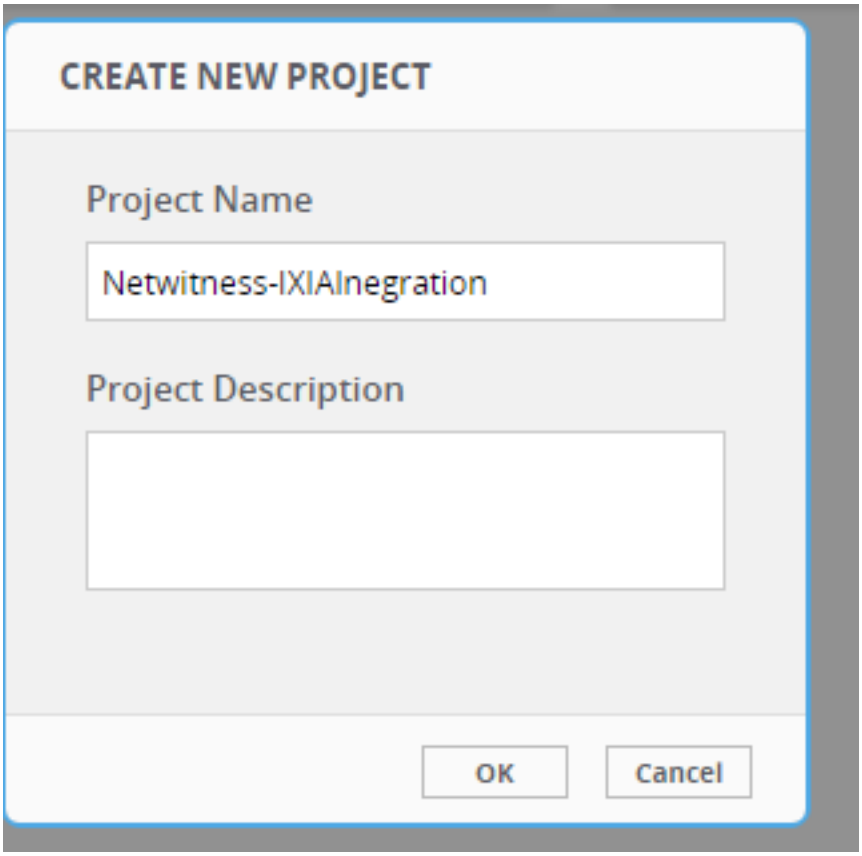
Task 1. Deploy Client Machines

- Deploy client machines onto which you want to route the traffic to the Network Decoder. See the Ixia CloudLens documentation (https://www.ixia.cloud/help/Default.htm) for specifications needed for supported client machines.
- For Client Machines (as well as Decoder machine) the following ports must be opened on AWS Security Group Inbound Rules; UDP 19993 from all, TCP 22 from Admin IP.

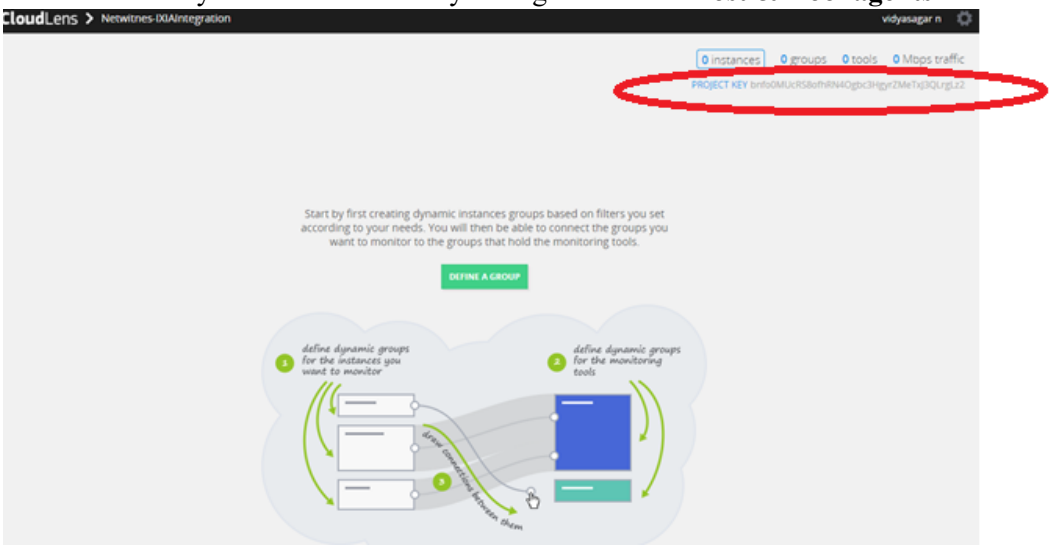
Task 2. Create CloudLens Project

Complete the following steps to create a new project and get your project key.

1. Get Cloudlens login credentials and access to a free trial.
   a. Create an Ixia login account at https://www.ixiacom.com/products/cloudlens-trial-a.
2. Go to the Cloudlens public site (https://www.ixia.cloud).
3. Click + (add) to create a new project with a name of your choosing (for example, **NetWitness-IxiaIntegration**).



4. Click on your newly created project and make note of your Project Key.
   You need the key later for the API key configured on the **Host & Tool agents**.



Task 3. Install Docker Container on Decoder

Complete the following steps to install the Docker container onto the Network Decoder.

1. SSH to the Network Decoder.
2. Enter the following commands to complete the install the Docker service on the Decoder.

```
#yum clean all
# yum -y install docker
```

3. Enter the following command string to start the Docker service.

```
# service docker start
```

4. Enter the following commands to:
   - Access the Ixia repository and obtain the cloudlens-agent container.
   - Replace the **ProjectKeyFromIxiaProjectPortal** variable, which identifies your project key in Ixia portal, with the Project Key you created in Task 2. Create CloudLens Project.

```
sudo docker run \
--name cloudlens \
 -v /:/host \
-v /var/run/docker.sock:/var/run/docker.sock \
-d --restart=always \
--net=host \
--privileged \
ixiacom/cloudlens-agent:latest \
--server agent.ixia.cloud \
--accept_eula y \
--apikey ProjectKeyFromIxiaProjectPortal \
```

Task 4. Install Docker Container on Clients

Complete the follow steps to Y install the Docker Container onto the client machines for which you want to route the traffic to the Network Decoder.

1. SSH to the AWS Client instance.
2. Enable root access to OS CLI (for example `sudo su -`).
3. Enter the following commands to install Docker.
   ```
   # yum -y install docker
   ```
   Caution: The above example of the installed docker engine is for CentOS7. The instructions may vary slightly for different Linux Distributions. For more information, see the Docker docs at https://docs.docker/install.
4. Enter the following commands to start the Docker service.
   ```
   # service docker start
   ```
5. Enter the following commands to:
   - Access the Ixia repository and obtain the **cloudlens-agent** container.
   - Replace the variable **ProjectKeyFromIxiaProjectPortal**, which identifies your project key in Ixia portal, with the Project Key you created in the previous section.
   ```
   sudo docker run \
   --name cloudlens \
    -v /:/host \
   -v /var/run/docker.sock:/var/run/docker.sock \
   -d --restart=always \
   --net=host \
   --privileged \
   ixiacom/cloudlens-agent:latest \
   --server agent.ixia.cloud \
   --accept_eula y \
   --apikey ProjectKeyFromIxiaProjectPortal \
   ```

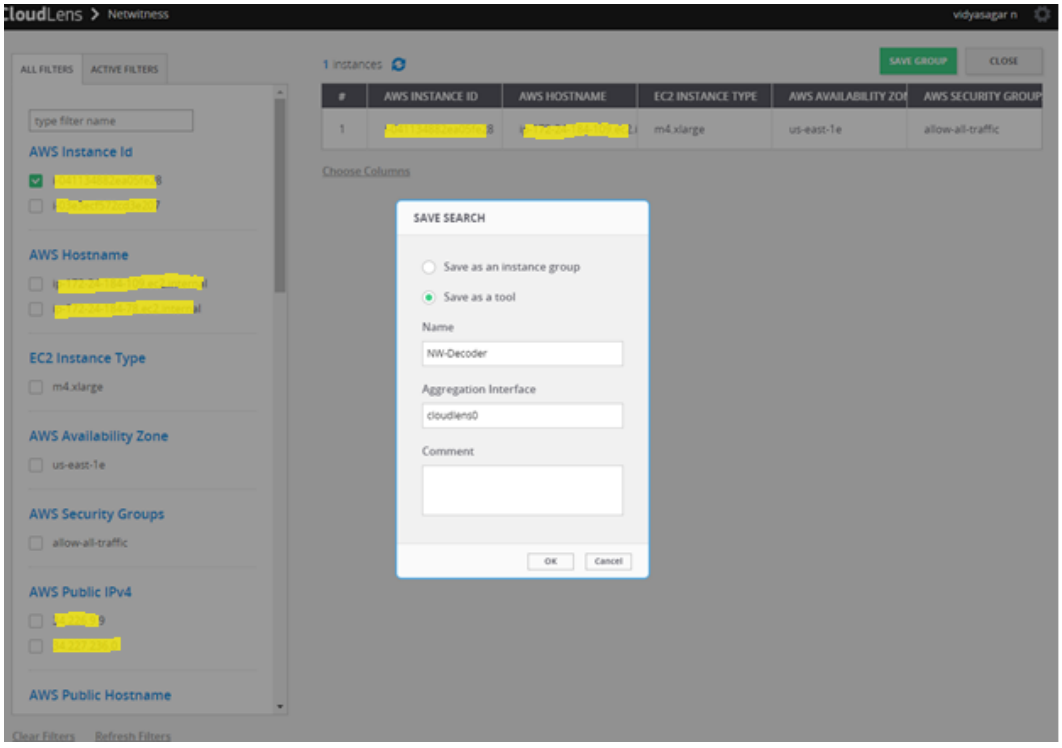
Warning: If you cut and paste commands from a PDF, first paste them into a test editor such as Notepad to confirm the syntax before pasting into the OS CLI. Direct cut and paste between PDF and CLI can contain dashes or other special characters that should not be part of the commands.

Task 5. Map the Network Decoder to Ixia Clients

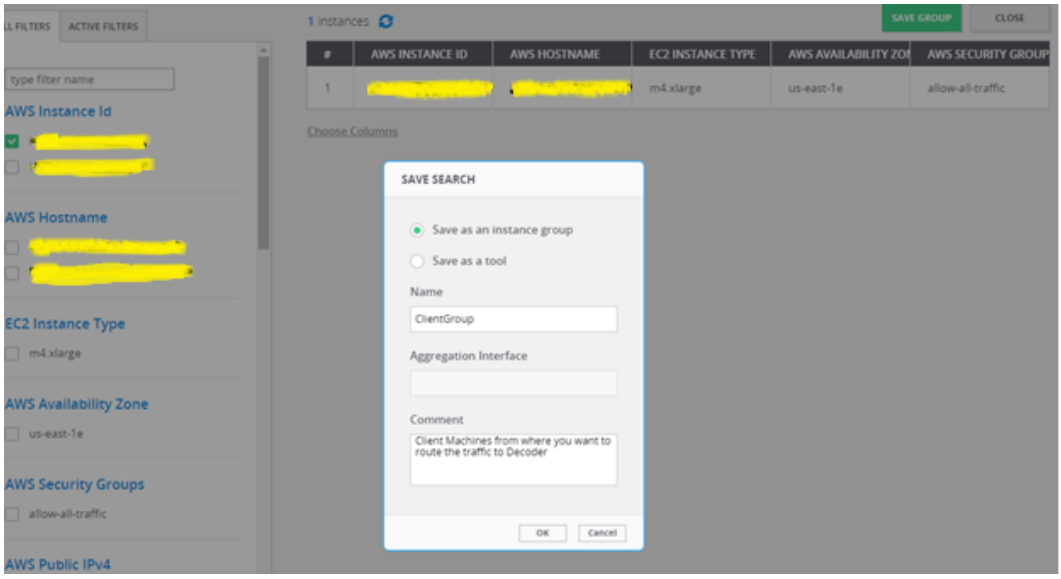
Complete the following steps to map the Network Decoder to the client machines to route the traffic to the Network Decoder.

1. Go to the Cloudlens public site (https://www.ixia.cloud).
2. Double-click on your project to open it.
3. Click the **Define Group** button or the Instances count.
   You should see two instances listed, one for your decoder and the other for the client machines.
4. Filter for the decoder instance and click **Save Search**.
5. Choose **Save as a tool**.
6. Specify a name for the tool, and the **Aggregation Interface**.
   Use a meaningful name for the Aggregation Interface (for example **cloudlens0**. This is a virtual

interface that appears in the OS where your Tool is installed. You need to instruct your tool to 'listen' to that interface in a subsequent step.

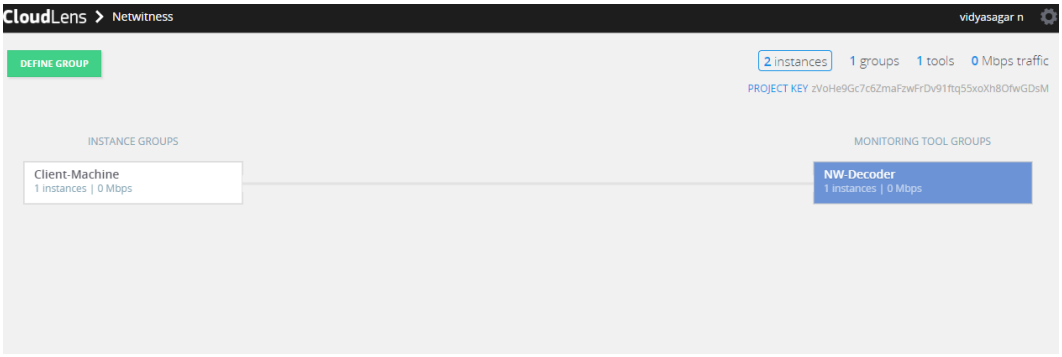


7. Filter the client host instance from the list, and click **Save Search**.



8. Navigate back to the top-level view of the project.
   Your client machine instance and Decoder instance are now displayed.

9. Drag a connection between the your client machine instance and Decoder instance to allow the flow of packets.



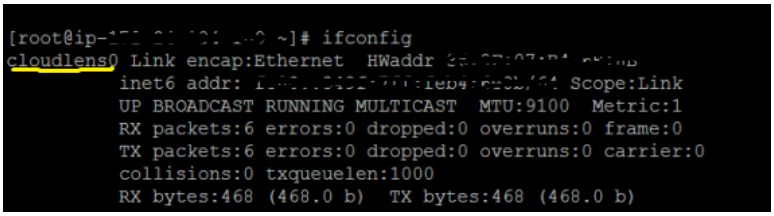
Task 6. Validate CloudLens Packets Arriving at Decoder

Complete the following steps to validate that packets are actually arriving at the Network Decoder.

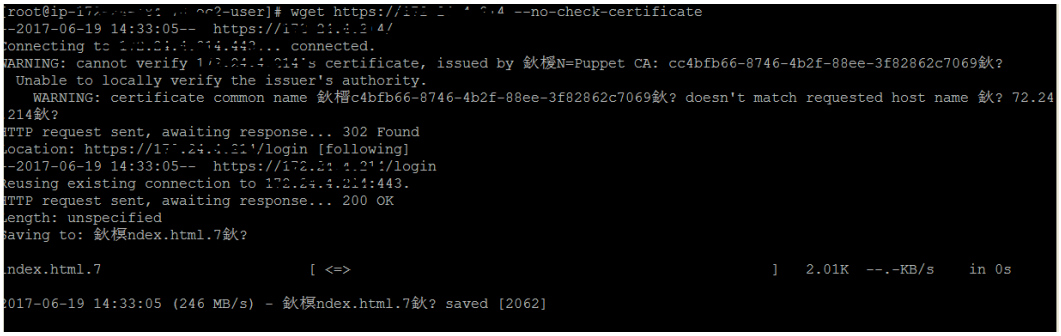
1. SSH to the Network Decoder.
2. Enter the following command.

   `ifconfig`

   The new aggregation interface you created is displayed.

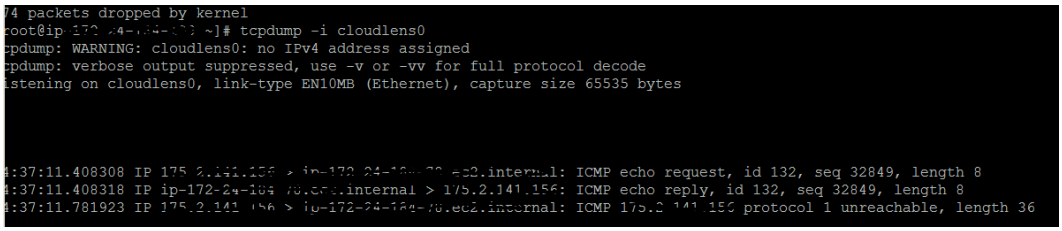


3. Generate traffic from the client OS instance CLI (for example, `wget http://www.google.com/`).



4. SSH to Network Decoder to go to your Network Decoder instance CLI.

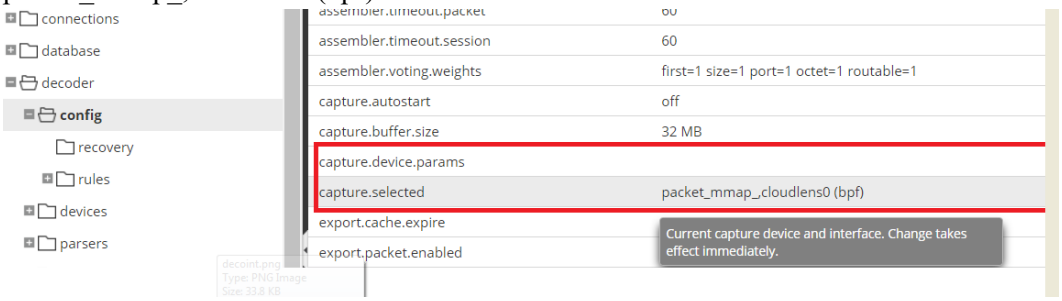
5. Enter the following commands to look for suitable results in the tcpdump.
   `tcpdump -I Cloudlens0`



Task 7. Set the Interface in the Network Decoder

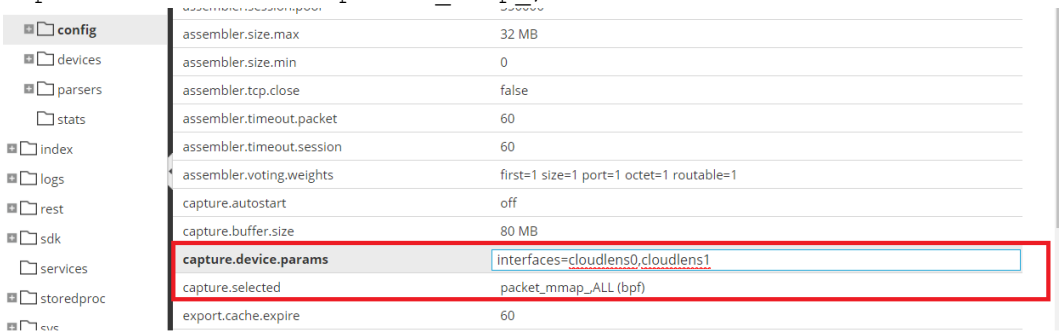
Complete the following steps in the Network Decoder to set the interface to use for the Ixia integration.

1. SSH to the Network Decoder.
2. Enter the following commands to restart decoder service.
   `$ sudo restart nwdecoder`
   The Network Decoder is now set to capture network traffic.
3. Log in to NetWitness and click  **(Admin) > Services**.
4. Select a Decoder service and click  > **View > Explore**.
5. Expand the **decoder** node and click **config** to view the configuration settings.
6. Set the **capture.selected** parameter to the following value.
   packet_mmap_,cloudlens0(bpf)



7. (Conditional) - If you have multiple capture interfaces on the Network Decoder, set the parameters with the following values.
   `capture.device.params --> interfaces=cloudlens0,cloudlens1`

```
capture.selected --> packet_mmap_,All
```



8. Restart the Decoder service after you set the **capture.selected** parameter.

Integrate f5® BIG-IP with the Network Decoder

IG-IP Virtual Edition (VE) is an inline virtual server and load balancer. A common use case would be for the f5® box to be a virtual web server that presents a single IP address / host name that manages requests to a pool of web servers in the cloud.

All traffic to NetWitness flows through the f5® BIG-IP VE virtual server.

The virtual server functions of the BIG-IP clone all traffic to a designated computer by re-writing mac addresses and loading them into a subnet shared with the destination sniffer. This guide describes how to set up the Decoder as the sniffer.

f5® BIG-IP VE Deployment Information

f5® BIG-IP VE on AWS will be available through the AWS Marketplace and activated by a BYOL license. A thirty-day free trial is also available.

For more information on this solution refer to the f5® BIG-IP DNS Data Sheet (https://www.f5.com/pdf/products/big-ip-dns-datasheet.pdf).

Task 1: Set Up a BIG-IP VE Virtual Server Instance

Set up a BIG-IP VE Virtual Server Instance according to the instructions in the "BIG-IP Virtual Edition 12.1.0 and Amazon Web Services: Multi-NIC Manual" (https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-ve-multi-nic-setup-amazon-ec2-12-1-0.html). Complete all the steps through the last steps, "Creating a virtual server."

This virtual server performs packet capture. You may need to create multiple virtual servers to depending on your volume.

As part of creating the virtual server, you must have at least one server in your NetWitness domain to handle the traffic routed by the virtual server (for example, you can create another instance in AWS to host the internal server).

Task 2: Create a Clone Pool

1. Make sure that your Decoder has a network interface on the same subnet as one of the network interfaces on the BIG-IP VE instance.
   The clone pool sends packets to the Decoder by rewriting MAC addresses and sending them out a

network interface. MAC address rewriting can be used to route packets to another subnet.

2. Set up the clone pool within the BIG-IP VE virtual server according to the instructions in "K13392: Configuring the BIG-IP system to send traffic to an intrusion detection system (11.x - 13.x)" article (https://support.f5.com/kb/en-us/solutions/public/13000/300/sol13392.html).
   This document explains how to create the clone pool, and how to make an existing virtual server copy traffic to the clone pool. In this case, we will place the Decoder instance in the clone pool.

Guidelines

The following guidelines will help you to configure packet capture correctly using BIG-IP VE.

- The Decoder instance must have its own IP address on one of the same subnets as BIG-IP VE. BIG-IP uses that IP address to identify the Decoder as being part of the clone pool.
- When adding the Decoder instance to the clone pool, BIG-IP asks for a port number in addition to the IP address. This port number does not matter for the cloned traffic. The Decoder will receive all the cloned traffic, regardless of what port number was used here.
- By default, the AWS subnet shared by the Decoder and BIG-IP VE will not allow the cloned traffic to travel from the BIG-IP VE interface to the Decoder interface. You must disable the **`source/dest. check`** on both the Decoder and BIG-IP VE network interfaces in AWS.
- The Decoder instance must have a single network interface, eth0, by default. The Decoder captures traffic on this interface, but it may also receive administrative traffic on this interface. NetWitness recommends using network rules to filter out ssh and nwdecoder traffic from the capture stream. These are ports 22 (ssh) and 50004/56004 (nwdecoder).

Troubleshooting Tips

There are areas to troubleshoot if packets are not being accepted by the Decoder.

- Make sure that the BIG-IP VE is sending the packets out of the correct interface.
  The BIG-IP VE instance contains tcpdump. Use it to verify the cloned packets are being sent out the expected interface. If they are not, there is a problem in the setup of the clone pool or the virtual server.
- Make sure that the Decoder is receiving packets.
  The Decoder has `tcpdump` installed on it. Use it to verify that the Decoder is receiving packets. If the Decoder is not capturing packets, make sure that
  - The AWS **`source/dest. check`** is turned off.
  - The Decoder is on the same subnet as the interface the BIG-IP VE is using to clone packets.

Integrate VPC Traffic Mirroring with the Network Decoder

VPC Traffic Mirroring allows users to capture and inspect network traffic to analyze packets without using any third-party packet forwarding agents. The solution provides insight and access to network traffic across VPC infrastructure. Users can copy network traffic at any ENI (Elastic Network Interfaces) in VPC, and send it to NetWitness to analyze, monitor, and troubleshoot performance issues.

You must complete the following tasks to integrate the Network Decoder with VPC Traffic Mirroring:

Task 1. Configure the Network Decoder as a VPC Traffic Mirroring Destination

Task 2. Configure a VPC Traffic Mirroring Filter

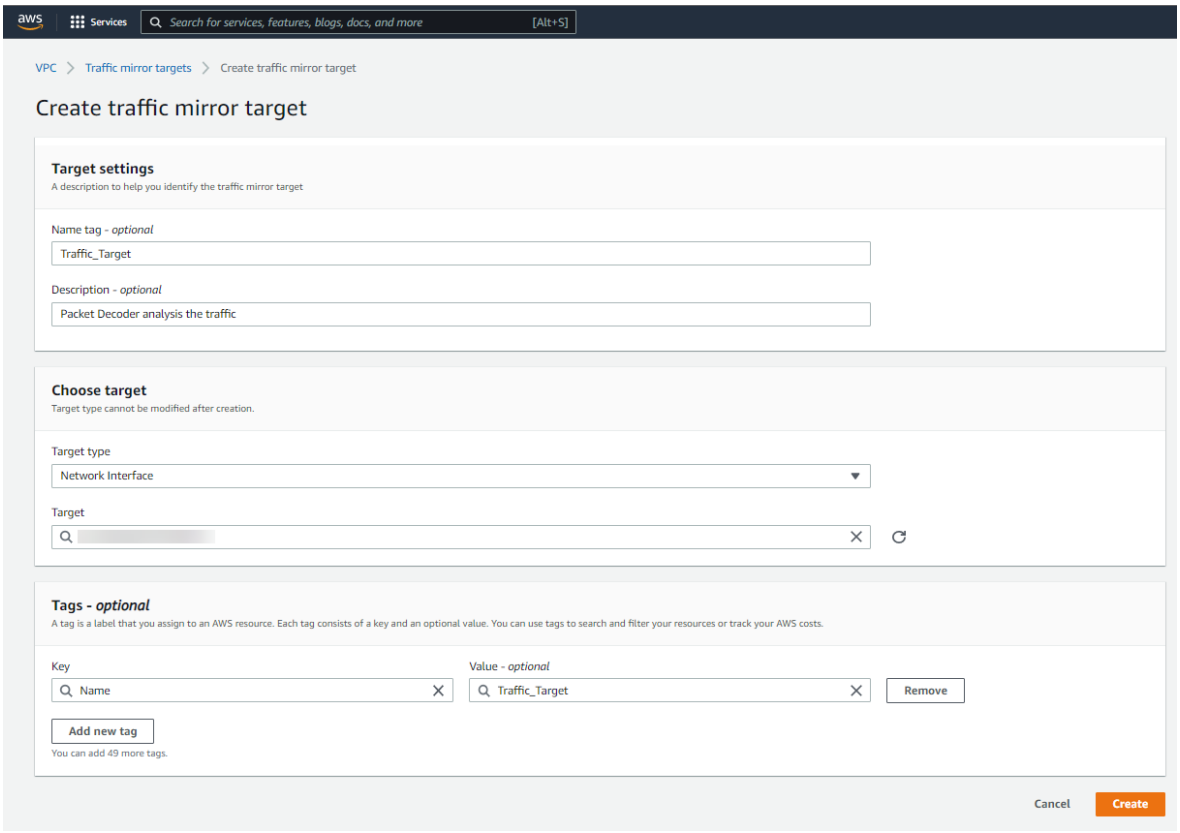
Task 3. Configure a VPC Traffic Mirroring Session

Task 4. Setup a new VXLAN interface on the Network Decoder

Task 5. Validate VPC Traffic Mirroring Packets Arriving at Network Decoder

Task 1. Configure the Network Decoder as a VPC Traffic Mirroring Destination.

1. Open the VPC service console view at https://console.aws.amazon.com/vpc/home.
2. In the navigation panel, select **Traffic mirror targets**.
3. Click **Create traffic mirror target**.



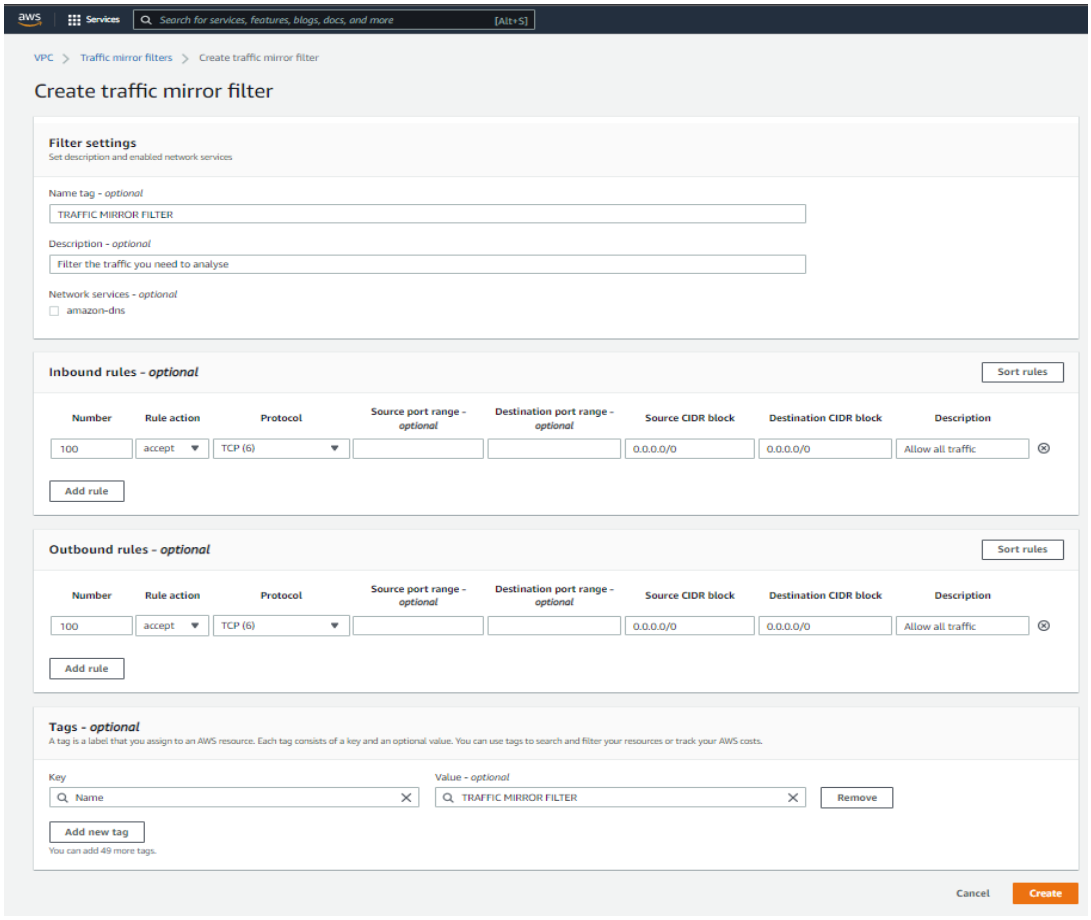
4. In the **Create traffic mirror target** dialog, provide all the required information and click **Create**.

Task 2. Configure a VPC Traffic Mirror Filter

Configure the VPC traffic mirror filter to send only the required packets to the Network Decoder. You can determine if the inbound or outbound traffic needs to be captured or not.

Note: Make sure the UDP port 4789 is open on the AWS instance of Network Decoder.

1. In the navigation panel, select **Traffic mirror filters**.
2. Click **Create traffic mirror filter**.

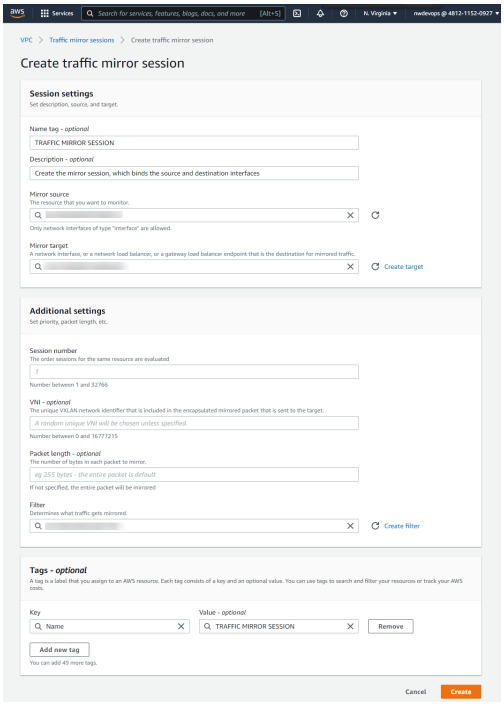


3. In the **Create traffic mirror filter** dialog, provide all the required information and click **Create**.

Task 3. Configure a VPC Traffic Mirror Session

You must configure a VPC Traffic Mirror Session to mirror the traffic by a communication channel between source ENI and destination ENI.

1. In the navigation panel, select **Traffic mirror sessions**.
2. Click **Create traffic mirror session**.



3. In the **Create traffic mirror session** dialog, provide all the required information and click **Create**.

Task 4. Set Up a new VXLAN Interface on the Network Decoder

To capture the UDP enabled traffic you must create an interface and tunnel it to Network Decoder by performing the following steps.

1. SSH to the Decoder.
2. Enter the following commands.

```
sudo ip link add tun0 type vxlan id <VXNLAN ID> local any dev <primary interface ex: eth0> dstport 4789

sudo ip link set tun0 up

ifconfig
```

```
tun0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 8951
        inet6 fe80::bc58:afff:fe06:ec29  prefixlen 64  scopeid 0x20<link>
        ether be:58:af:06:ec:29  txqueuelen 1000  (Ethernet)
        RX packets 989  bytes 74140 (72.4 KiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 0  bytes 0 (0.0 B)
        TX errors 0  dropped 8 overruns 0  carrier 0  collisions 0
```

3. To create a firewall rule in the Network Decoder to allow traffic through the tunnel.

a. Open the IP tables file using the command `vi /etc/sysconfig/iptables`.

b. Append the line `-I INPUT -p udp -m udp --dport 4789 -j ACCEPT`.

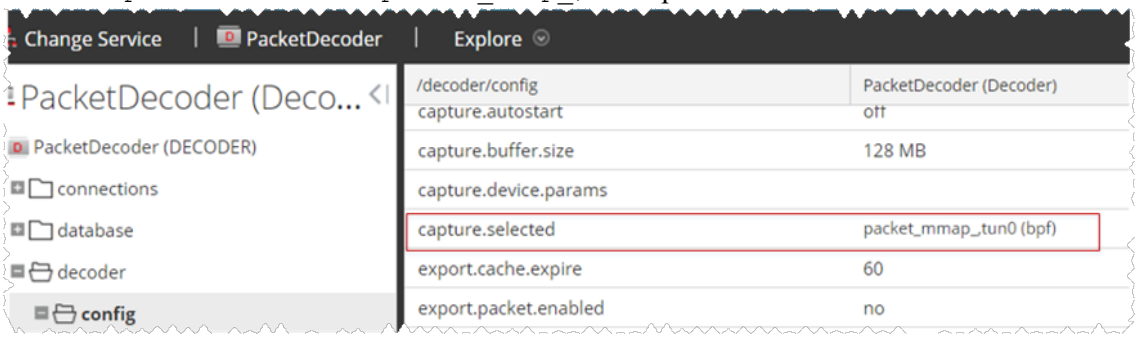
c. Restart IP tables by using the following commands.

```
service iptables restart
service iptables status
```

4. To set the interface in the Network Decoder.

   a. Log in to NetWitness, select the `decoder/config` node in Explorer view of the Network Decoder service.

   b. Set the `capture.selected = packet_mmap_,tun0` parameter.

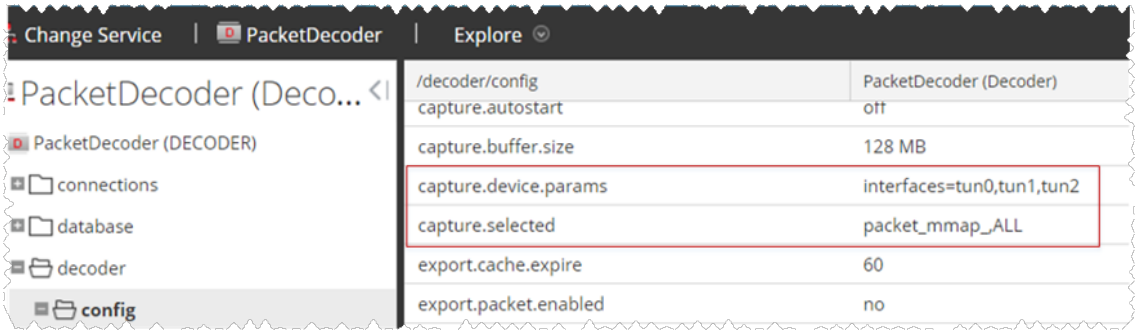


5. (Conditional) If you have multiple tunnels on the Network Decoder.

   a. Restart the Decoder service after you create the tunnel in Network Decoder.

   b. Log in to NetWitness, select the `decoder/config` node in Explorer view of the Network Decoder service, and set the following parameters.

```
capture.device.params = interfaces=tun0,tun1,tun2
capture.selected = packet_mmap_,All
```



6. Restart the Decoder service.

```
$ sudo restart nwdecoder
```

   The user should be all set to capture the network traffic in the Network Decoder.

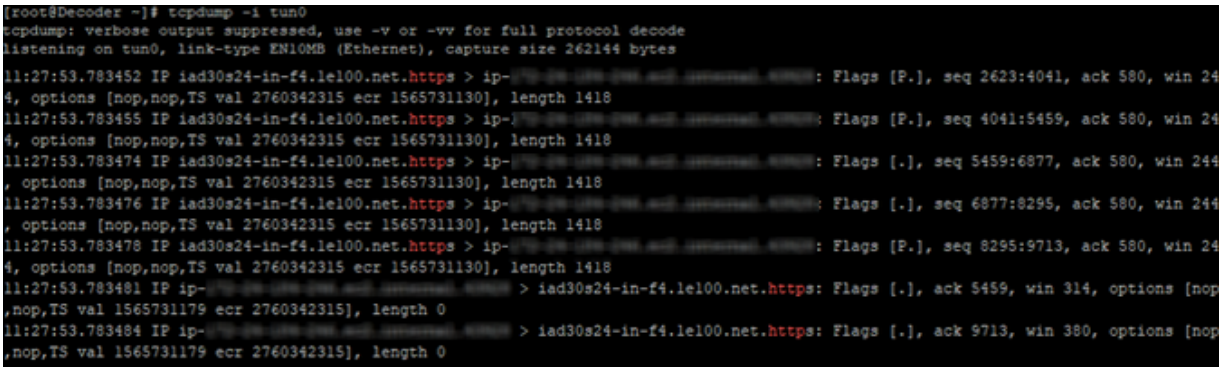
Task 5. Validate VPC Traffic Mirroring Packets Arriving at the Network Decoder

Perform the following steps to validate if the Network Decoder is receiving the network data (packets) successfully.

1. Generate traffic from the client OS instance CLI (for example, `wget http://www.google.com/`).



2. Enter the `tcpdump -i tun0` command to look for suitable results in the tcpdump.



3. The NetWitness reflects meta values as shown below.



Note: You can mirror traffic from an EC2 instance that is supported by the AWS Nitro system (A1, C5, C5d, C5n, I3en, M5, M5a, M5ad, M5d, p3dn.24xlarge, R5, R5a, R5ad, R5d, T3, T3a, and z1d).

Note: For more information, see "New – VPC Traffic Mirroring" documentation at https://aws.amazon.com/blogs/aws/new-vpc-traffic-mirroring/.

AWS Instance Configuration Recommendations

Note: These recommendations can be used as a baseline for 12.0.0.0 and adjusted as needed.

This topic contains the minimum AWS instance configuration settings recommended for the NetWitness virtual stack components.

- EC2 Instance:
  - Instance type adjustments -you must adjust instance types according to your ingestion rate, content and parsers, dashboard reports, scheduled reports, investigations, and active users.
  - Recommended settings - the recommended settings in the NW component instance tables below were calculated under the following conditions.
    - Ingestion rates of 15,000 EPS and 1.5 Gbps were used.
    - All the components were integrated.
    - The Log stream includes a Log Decoder, Concentrator, and Archiver.
    - The Packet stream includes a Network Decoder and Concentrator.
    - The Endpoint Hybrid stream includes a Endpoint Server, Concentrator and Log Decoder.
    - Respond is receiving alerts from the Reporting Engine and Event Stream Analysis.
    - The background load includes reports, charts, alerts, investigation, and respond.
- Block Storage

  For more information on the required volumes and the storage allocations, see the Storage Guide for NetWitness® Platform 12.0.

Archiver

EC2 Instance			
EPS	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
5,000	m4.xlarge No of CPU: 4 Memory: 16 GB	No	Yes
10,000	m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
15,000	m4.4xlarge No of CPU: 16 Memory: 64 GB	No	Yes

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
archiver	/dev/sdg	Throughput Optimized HDD	240 MB/s
workbench	/dev/sdh	Throughput Optimized HDD	N/A

Broker

EC2 Instance		
Instance Type	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
m4.xlarge No of CPU: 4 Memory: 16 GB	No	Yes

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
broker	/dev/sdg	General Purpose SSD	N/A

Concentrator - Log Stream

EC2 Instance			
EPS	**Instance Type**	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
5,000	m4.xlarge No of CPU: 4 Memory: 16 GB	No	Yes
10,000	m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
15,000	m4.4xlarge No of CPU: 16 Memory: 64 GB	No	Yes

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index	/dev/sdg	Provisioned IOPS	10,000
session, metadb	/dev/sdh	Throughput Optimized HDD	240 MB/s

Packet Stream Solutions

Concentrator - Gigamon Solution

EC2 Instance			
Mbps/Gbps	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
500 Mbps	c4.4xlarge No of CPU: 16 Memory: 30 GB	No	Yes
1,000 Mbps	c4.8xlarge No of CPU: 36 Memory: 60 GB	No	Yes
1.5 Gbps	m4.10xlarge No of CPU: 40 Memory: 160 GB	No	Yes

Concentrator - f5 BIG-IP Solution

To be updated when f5 BIG-IP performance testing is complete.

EC2 Instance			
Mbps/Gbps	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
230 Mbps	m4.4xlarge No. of CPU: 16 Memory: 64 GB	No	No

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index	/dev/sdg	Provisioned IOPS	15,000
session, metadb	/dev/sdh	Throughput Optimized HDD	240 MB/s

Decoder - Gigamon Solution

EC2 Instance			
Mbps/Gbps	**Instance Type**	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
500 Mbps	c4.2xlarge No of CPU: 8 Memory: 15 GB	Yes	Yes
1000 Mbps	c4.4xlarge No of CPU: 16 Memory: 30 GB	Yes	Yes
1.5 Gbps	c4.8xlarge No of CPU: 36 Memory: 60 GB	Yes	Yes

Decoder - f5 BIG-IP Solution

To be updated when f5 BIG-IP performance testing is complete.

EC2 Instance			
Mbps/Gbps	**Instance Type**	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
230 Mbps	m4.4xlarge No. of CPU: 16 Memory: 64 GB	No	No

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index,session,meta	/dev/sdg	Throughput Optimized HDD	240 MB/s
packet	/dev/sdh	Throughput Optimized HDD	240 MB/s

ESA and Context Hub on Mongo Database

	EC2 Instance		
EPS	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
9,000	m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
18,000	r4.2xlarge No of CPU: 8 Memory: 61 GB	No	Yes
30,000 Aggregation Rate	r4.4xlarge No of CPU: 16 Memory: 122 GB	No	Yes

Cloud Provider Block Storage			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
apps (/opt/rsa)	/dev/sdg	General Purpose SSD	N/A

Log Collector (Syslog, Netflow, and File Collection Protocols)

EC2 Instance			
EPS	**Instance Type**	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
30,000 NON SSL	c4.2xlarge No of CPU: 8 Memory: 15 GB	No	Yes

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
logcollector	/dev/sdg	General Purpose SSD	N/A

Log Decoder

EC2 Instance			
EPS	**Instance Type**	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
5,000	c4.2xlarge No of CPU: 8 Memory: 15 GB	Yes	Yes
10,000	c4.4xlarge No of CPU: 16 Memory :30 GB	Yes	Yes
15,000	c4.8xlarge No of CPU: 36 Memory: 60GB	Yes	Yes

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index,session,meta	/dev/sdg	Throughput Optimized HDD	240 MB/s
packet	/dev/sdh	Throughput Optimized HDD	240 MB/s

NW Server, Reporting Engine, Respond and Health & Wellness

EC2 Instance		
Instance Type	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
m4.4xlarge No of CPU: 16 Memory: 64 GB	No	Yes

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
uax,ipdb	/dev/sdg	General Purpose SSD	N/A
redb,rehome	/dev/sdh	General Purpose SSD	N/A

NetWitness Endpoint Hybrid

	EC2 Instance		
Agents	**Instance Type**	**Enhanced Networking Enabled**	**Tenancy Type - Dedicated - Run a Dedicated Instance**
15,000 agents	m4.10xlarge No of CPU: 40 Memory: 160 GB RAM	Yes	Yes

Cloud Provider Block Storage			
Volumes	**Device**	**Volume Type**	**IOPS/Baseline Throughput**
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index,session,meta (Log Decoder)	/dev/sdg	Throughput Optimized HDD	240 MB/s
packet (Log Decoder)	/dev/sdh	Throughput Optimized HDD	240 MB/s
index (Concentrator)	/dev/sdi	Provisioned IOPS	10,000
session,meta (Concentrator)	/dev/sdj	Throughput Optimized HDD	240 MB/s
mongoDB	/dev/sdl	Throughput Optimized HDD	240 MB/s

Appendix A. Silent Installation Using CLI

You can use the following Command Line Interface commands to run the installation script (`nwsetup-tui`) without getting prompted for inputs. This enables you to automate the installation of a host by supplying response to the scripts prompts through the command line.

1. After you have created a base image on the host, log in to the host with the `root` credentials.
2. Submit the `nwsetup-tui` script with the `--silent` command and the arguments that you want to apply.

   The following command string is an example of how you would install a basic NW Server host.

   ```
   nwsetup-tui --silent --is-head=true --host-name=new-host --master-
   pass=netwitness --deploy-pass=netwitness --repo-type=1 --custom-
   firewall=false --ip-override=false --eula=true
   ```

Note: In NetWitness Platform version 11.6 or later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`.
If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

3. (Conditional - For Component Hosts Only) Install the appropriate service **Category** on the newly provisioned host in the NetWitness Platform Hosts view.

   a. Log into NetWitness and go to  **(Admin)** > **Hosts**.

   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background

   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the host in the **New Hosts** dialog and click **Enable**.

   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the **Hosts** view (for example, **Event Stream Analysis**) and click 

.

   The **Install Services** dialog is displayed.

   d. Select the appropriate host type in **Category** and click **Install**.

Arguments

Argument	Description
`--help-install-opts`	Display all the arguments in this table.

Argument	Description
`--eula`	Accept or decline the End User License Agreement (EULA). Specify: • `true` (default) to accept the agreement • `false` to decline it and cancel the installation. For example: `--eula=true`
`--is-head`	Designate the host as the NW Server host or a component host. Specify: • `true` for NW Server host. • `false` for Component host. For example: `--is-head=true`
`--host-name`	Specify new hostname. If you do not specify this argument, NetWitness Platform retains the existing hostname. For example: `--host-name=<hostname>`
`--master-pass`	Enter master password. For example: `--master-pass=<password>`
`--deploy-pass`	Enter deployment password. For example: `--deploy-pass=<password>`
`--iface-name`	Specify network interface. For example: `--iface-name=eth0`
`--ip-override`	Accept or override IP address found for this host or change the IP configuration found on the host. Specify: • `true` provide IP address. • `false` use IP address found on the host. For example: `--ip-override=false`
`--ip-type`	Select ip address configuration type. Specify: • `1` Static IP Configuration) • `2` DCHP For example: `--ip-type=1`
`--ip-addr`	For Static IP configuration, enter IP Address for static address. For example: `--ip-addr=<ip-address>`
`--ip-netmask`	For Static IP configuration, enter Subnet Mask for static address. For example: `--ip-gateway=<subnet-mask>`

Argument	Description
`--ip-gateway`	For Static IP configuration, enter default gateway for static address. For example: `--ip-gateway=<default-gateway>`
`--ip-nameserver`	IP address assigned to DNS server. `--ip-nameserver=<ip-address>`
`--ip-nameserver-secondary`	Optional - IP address assigned to a secondary DNS server. For example: `--ip-nameserver-secondary=<ip-address>`
`--ip-domain`	For Static IP configuration, enter Local Domain Name for static address. For example: `--ip-domain=<default-gateway>`
`--repo-type`	Select type of update repository. Specify: • `1` Local repository • `2` External repository For example: `--repo-type=1`
`--repo-url`	For an external update repository, specify the url of the repository. For example: `--repo-url=<url>`
`--head-ip`	For a component host, specify IP Address of the NW Server. For example: `--head-ip=<ip-address>`
`--custom-firewall`	Disable default firewall configuration and use your custom configuration. Specify: • `true` use custom firewall configuration. • `false` use default firewall configuration. For example: `--custom-firewall=true`
`--use-nat`	Configure the host to use Network Address Translation (NAT) based IP addresses: • `true` use NAT IPs to connect to other hosts • `false` do not use NAT IPs to connect to other hosts (default) For example: `--use-nat=false`

NetWitness® Platform XDR

Version 12.0

Azure Installation Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

Azure Installation Overview **4**

Azure Environment Recommendations 4

Azure Deployment Scenarios 5

Process 5

NetWitness High-Level Deployment Diagram 6

Azure Configuration Recommendations **7**

Packet Stream Solutions 8

Decoder - Gigamon Solution 8

Concentrator - Gigamon Solution 9

Updating Partition Size 9

Azure Deployment **11**

Rules 11

Checklist 11

Storage Configurations 12

Enabling Swap Partition in Azure Deployments 12

Deploy NW Server Host 14

Task 1. - Upload NW Server VHDs 14

Task 2. - Create NW Server Image 16

Task 3. Create Virtual Machine (VM) 18

Deploy Component Core Services in Azure 29

Installation Tasks 38

Install 12.0.0.0 on the NetWitness Server (NW Server) and Component Hosts 38

Set Up ESA Hosts 46

Install Component Services on Hosts 46

Complete Licensing Requirements 47

(Optional) Install Warm Standby NW Server 47

NetWitness Azure Storage Allocation Procedure 47

Configure Hosts (Instances) in NetWitness Platform 49

Appendix A. Silent Installation Using CLI **50**

Azure Installation Overview

Azure instances have the same functionality as the NetWitness hardware and virtual hosts. NetWitness recommends that you perform the following tasks when you set up your Azure environment.

Before you can deploy NetWitness in Azure, you need to:

- Review the recommended compute and memory specifications needed for each NetWitness instance.
- Get familiar with the NetWitness Storage Guide to understand the types of drives and volumes needed to support NetWitness instances. For more information, see Storage Guide for NetWitness® Platform 12.0.
- Make sure that you have a NetWitness Throughput license.
- Use Chrome for your browser (Internet Explorer is not supported).

Azure Environment Recommendations

Azure instances have the same functionality as the NetWitness hardware hosts. NetWitness recommends that you perform the following tasks when you set up your Azure environment.

- Based on the resource requirements of the different components, follow best practices to use the system and dedicated storage appropriately.
- Build Concentrator directory for index database and Decoder directory for Packet database on SSD Disks with high IOPS / write throughput.

Azure Deployment Scenarios

Before you can deploy NetWitness you need to:

- Consider the requirements of your enterprise and understand the deployment process.
- Have a high-level picture of the complexity and scope of a NetWitness deployment.

Process

The components and topology of a NetWitness network can vary greatly between installations, and should be carefully planned before the process begins. Initial planning includes:

- Consideration of site requirements and safety requirements.
- Review of the network architecture and port usage.
- Support of group aggregation on Archivers and Concentrators, and virtual hosts.

When updating hosts and services, follow recommended guidelines under the "Running in Mixed Mode" topic in the *NetWitness Host and Services Getting Started Guide*.

You should also become familiar with Hosts, Host Types, and Services as they are used in the context of NetWitness also described in the *NetWitness Host and Services Getting Started Guide*.

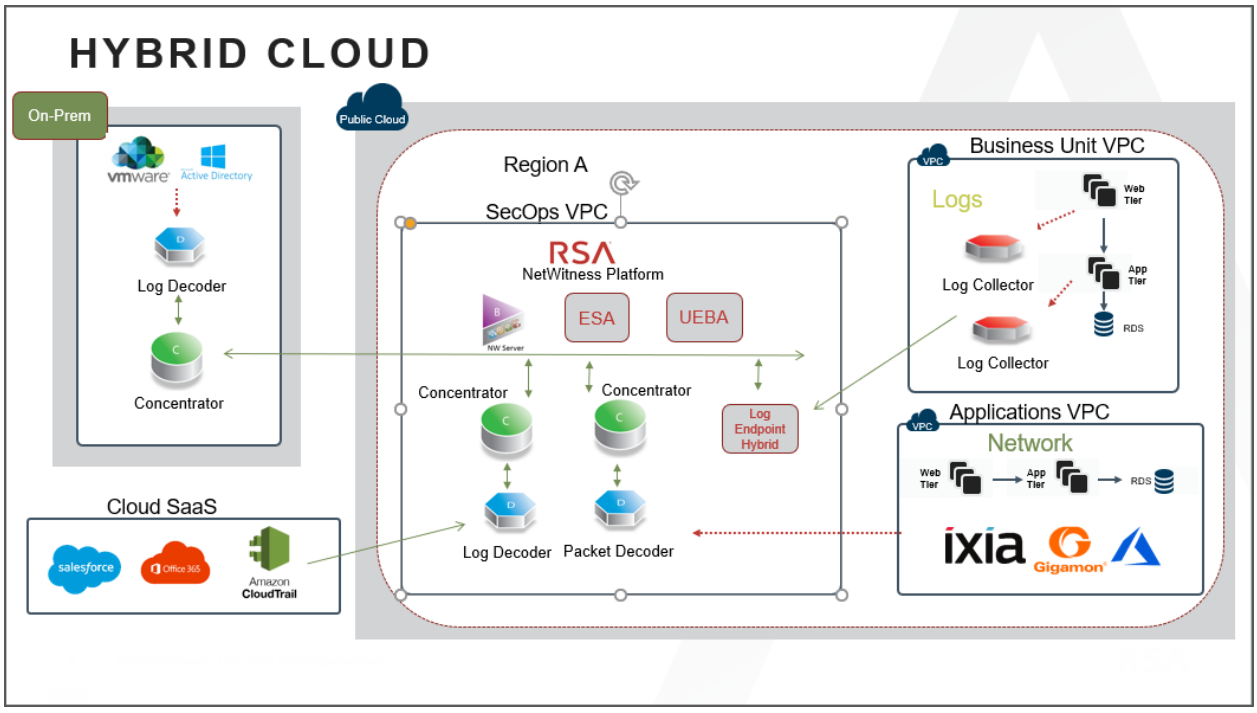
NetWitness High-Level Deployment Diagram

NetWitness is inherently modular. Whether organizations are looking to deploy on-premise or in the cloud, the NetWitness components are decoupled in a way which allows flexible deployment architectures to satisfy a variety of use cases.

The following figure is an example of a hybrid cloud deployment, where the base of the components are residing within the SecOps VPC. Centralizing these components make management easier while keeping network latency to a minimum.

Network, log and endpoint traffic could then be aggregated up to the SecOps VPC. The on-premise location would function just like a normal physical deployment and would be accessible for investigations and analytics.

Cloud SaaS visibility could be captured from a Log Decoder residing in either the cloud or on-premise locations.



Azure Configuration Recommendations

This topic contains the minimum Azure VM configuration settings recommended for the NetWitness (NW) virtual stack components.

- VM:
  - The recommended settings in the NetWitness component VM tables below were calculated under the following conditions.
    - Ingestion rates of 15,000 EPS and 1.5GBps were used.
    - All the components were integrated.
    - The Log stream included a Log Decoder, Concentrator, and Archiver.
    - The Packet stream included a Network Decoder and Concentrator.
    - Incident Management was receiving alerts from the Reporting Engine and Event Stream Analysis.
    - The background load included reports, charts, alerts, investigation, and respond.
    - The default partition size of Azure VM hosts for `/root` is 8GB and for `/var/netwitness` is 15GB. These partitions can be increased to a minimum of 40GB. For more information see, Updating Partition Size.
- VHD (Storage)
  For more information, see Storage Guide for NetWitness® Platform 12.0 on how to increase the number of volumes based on your storage requirements using the NetWitness Sizing & Scoping Calculator.

Azure Instance Recommendations

The following table shows the storage recommendations for NetWitness Azure VMs.

Azure Image Type	Rate (EPS)	CPU (Cores)	RAM (GB)	Instance Type (Azure Name)
NW	Does not apply	16	112	Standard D14_v2
Log Decoder	15,000	32	128	Standard D32s_v3
Log Concentrator	15,000	16	112	Standard DS14_v2
Archiver	15,000	16	112	Standard D14_v2
ESA	15,000	20	140	Standard D15_v2
Log Collector	15,000	8	32	Standard D8s_v3
UEBA*	Does not apply	16	112	Standard D14_v2

Note: *If your log collection volume is low, NetWitness recommends you to deploy UEBA only on a virtual host. If you have a moderate to high log collection volume, NetWitness recommends you to deploy UEBA on the physical host as described under "NetWitness UEBA Host Hardware Specifications" in the *Physical Host Installation Guide*.

Refer to the *Storage Guide for NetWitness Platform* for additional storage information.

Packet Stream Solutions

The following tables show Instance recommendations for Different EPS rates for Packet stream.

Note: NetWitness Decoder is supported with Gigamon Packet broker from version 11.7.x or higher on Azure Cloud environment.

Decoder - Gigamon Solution

Azure Image Type	Rate (Mbps)	CPU (Cores)	RAM (GB)	Instance Type (Azure Name)	Accelerated Networking Enabled
Decoder	500	16	64	Standard D16ds_v4	Yes
Decoder	1000	16	64	Standard D16ds_v4	Yes
Decoder	1500	32	128	Standard D32ds_v4	Yes

Rate (Mbps)	Volumes	Volume Type	IOPS / Baseline Throughput
500	index, session, meta	RAID5 of minimum 3 P15 Premium SSD Disks	80MB/s
500	packet	RAID5 of minimum 3 P15 Premium SSD Disks	80MB/s
1000	index, session, meta	RAID5 of minimum 3 P20 Premium SSD Disks	170MB/s
1000	packet	RAID5 of minimum 3 P30 Premium SSD Disks	170MB/s
1500	index, session, meta	RAID5 of minimum 3 P40 Premium SSD Disks	300MB/s
1500	packet	RAID5 of minimum 3 P40 Premium SSD Disks	300MB/s

Concentrator - Gigamon Solution

Azure Image Type	Rate (Mbps)	CPU (Cores)	RAM (GB)	Instance Type (Azure Name)	Accelerated Networking Enabled
Packet Concentrator	500	16	64	Standard D16ds_v4	No
Packet Concentrator	1000	16	114	Standard DS14_v2	No
Packet Concentrator	1500	16	114	Standard DS14_v2	No

Note: For Packet Concentrator with **500**Mbps rate, if the query load on the environment is on the higher side (max concurrent queries > **5**), it is recommended to use **Standard DS14_v2** Instance.

Rate (Mbps)	Volumes	Volume Type	IOPS / Baseline Throughput
500	index	RAID5 of minimum 3 P30 Premium SSD Disks	10000
500	session, meta	RAID5 of minimum 3 P15 Premium SSD Disks	80MB/s
1000	index	RAID5 of minimum 3 P40 Premium SSD Disks	12000
1000	session, meta	RAID5 of minimum 3 P20 Standard SSD Disks	170MB/s
1500	index	RAID5 of minimum 3 P40 Premium SSD Disks	15000
1500	session, meta	RAID5 of minimum 3 P40 Premium SSD Disks	300MB/s

Updating Partition Size

You can increase the partition size to a minimum of 40GB each.

After adding additional required disk size to the Azure VM, you can extend the partition sizes using the following commands:

1. SSH to the VM, login as a root user and execute the following command to view the existing partitions along with the new partition added.

```
lsblk
```

2. Check the name of the new partition. Eg: `sdc`

```
pvcreate /dev/sdc -y
vgextend netwitness_vg00 /dev/sdc -y
lvextend -L 40G /dev/netwitness_vg00/root -y
xfs_growfs /dev/netwitness_vg00/root
lvextend -L 40G /dev/netwitness_vg00/nwhome -y
xfs_growfs /dev/netwitness_vg00/nwhome
```

These commands are provided assuming that `sdc` is the new disk added and 40GB is the extended partition size for each of the partitions.

Azure Deployment

This topic contains the rules and high-level tasks you must perform to deploy NetWitness components in Azure.

Rules

You must adhere to the following rules:

It is recommended to use private IP addresses when you provision Azure NetWitness VMs.

Checklist

Step	Description	√
1.	Deploy NW Server Host	
2.	Deploy Component Core Services in Azure	

Storage Configurations

This topic contains the recommended Azure storage configurations.

For storage allocations of all host types, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness® Platform 12.0*.

Enabling Swap Partition in Azure Deployments

After completing the Azure deployment, you must enable the swap in your deployment.

To do this, perform the following steps:

1. Modify the default parameters at `/etc/waagent.conf` to

```
ResourceDisk.Format=y
ResourceDisk.Filesystem=ext4
ResourceDisk.MountPoint=/mnt/resource
ResourceDisk.EnableSwap=y
ResourceDisk.SwapSizeMB=4096
```

The following screenshot displays the default parameters.

```
# Format if unformatted. If 'n', resource disk will not be mounted.
ResourceDisk.Format=y

# File system on the resource disk
# Typically ext3 or ext4. FreeBSD images should use 'ufs2' here.
ResourceDisk.Filesystem=ext4

# Mount point for the resource disk
ResourceDisk.MountPoint=/mnt/resource

# Create and use swapfile on resource disk.
ResourceDisk.EnableSwap=n

# Size of the swapfile.
ResourceDisk.SwapSizeMB=0
```

The following screenshot displays the modified parameters.

```
# Format if unformatted. If 'n', resource disk will not be mounted.
ResourceDisk.Format=y

# File system on the resource disk
# Typically ext3 or ext4. FreeBSD images should use 'ufs2' here.
ResourceDisk.Filesystem=ext4

# Mount point for the resource disk
ResourceDisk.MountPoint=/mnt/resource

# Create and use swapfile on resource disk.
ResourceDisk.EnableSwap=y

# Size of the swapfile.
ResourceDisk.SwapSizeMB=4096
```

Note: You can set the `ResourceDisk.SwapSizeMB` parameter based on your requirement.

2. Restart the `waagent.service` using the command: `systemctl restart waagent.service`

Note: To check the status of the swap use the command `swapon --show`.

Deploy NW Server Host

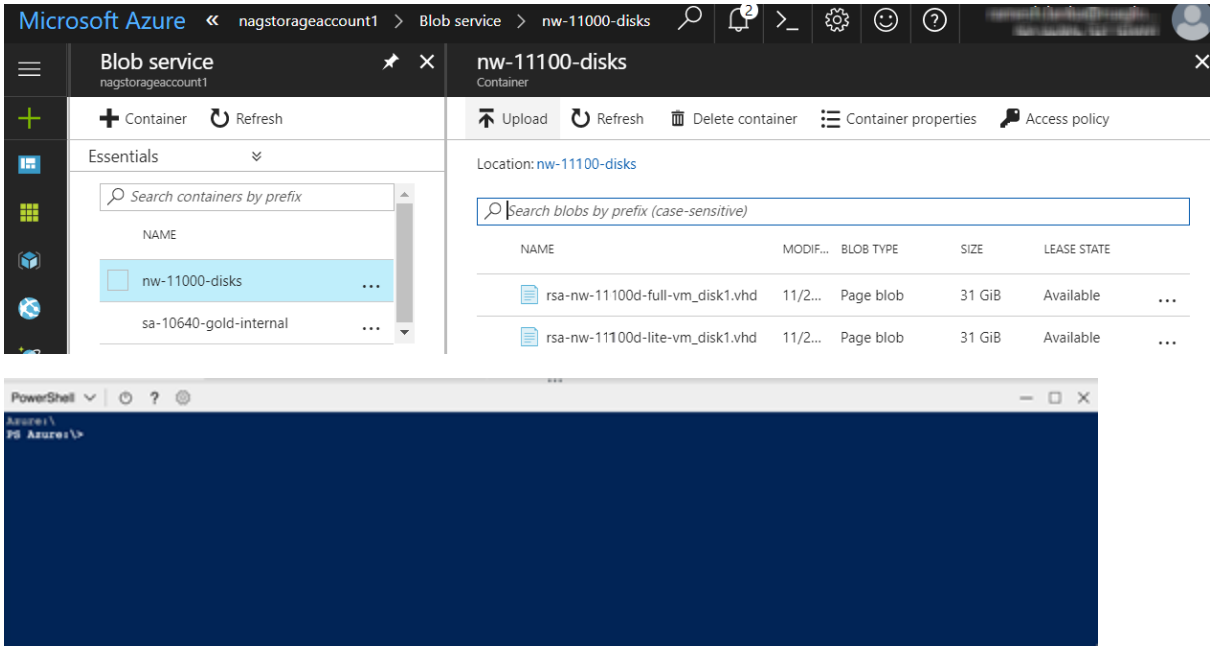
The following tasks must be performed to deploy a NetWitness Server (NW Server) on a virtual machine (VM) in the Azure Cloud environment.

Note: It is not mandatory to deploy the NW Server in the Azure Cloud environment . For more information on how to deploy other components, see Azure Deployment Scenarios.

Task 1. - Upload NW Server VHDs

To upload NW Server VHDs to Azure.

1. Contact Customer Support (https://community.netwitness.com/t5/support/ct-p/support) to open a support case requesting the NW Server VHDs. A valid throughput license is required.
2. Customer Support will update the case with VHD URI's.
3. In the Azure Portal, open the Powershell CLI.



You will need a storage account, blob service and container setup. This is where the VHD's are copied. After these are in place, you can execute the following command within the Azure Portal Powershell CLI. Alternatively, you can also run these commands from the Powershell on your workstation:

a. Run this command from Powershell to install AzureRM: `Install-Module -Name AzureRM -AllowClobber`

b. Execute this command to verify the installation process has been successfully done: `Import-Module -Name AzureRM`

c. If you find any error regarding execution policy, execute this command: - `Set-ExecutionPolicy -ExecutionPolicy RemoteSigned` (then repeat step b)

d. (Optional) If you are running the commands from the Powershell on your workstation, log in to your Azure account using this command: `Login-AzureRmAccount`

e. Select the Subscription: `Select-AzureRmSubscription -SubscriptionId <subscriptionid>`

f. Create a target context: `$targetStorageContext = (Get-AzureRmStorageAccount -ResourceGroupName <resource-group-name> -Name <storage-account-name>).Context`

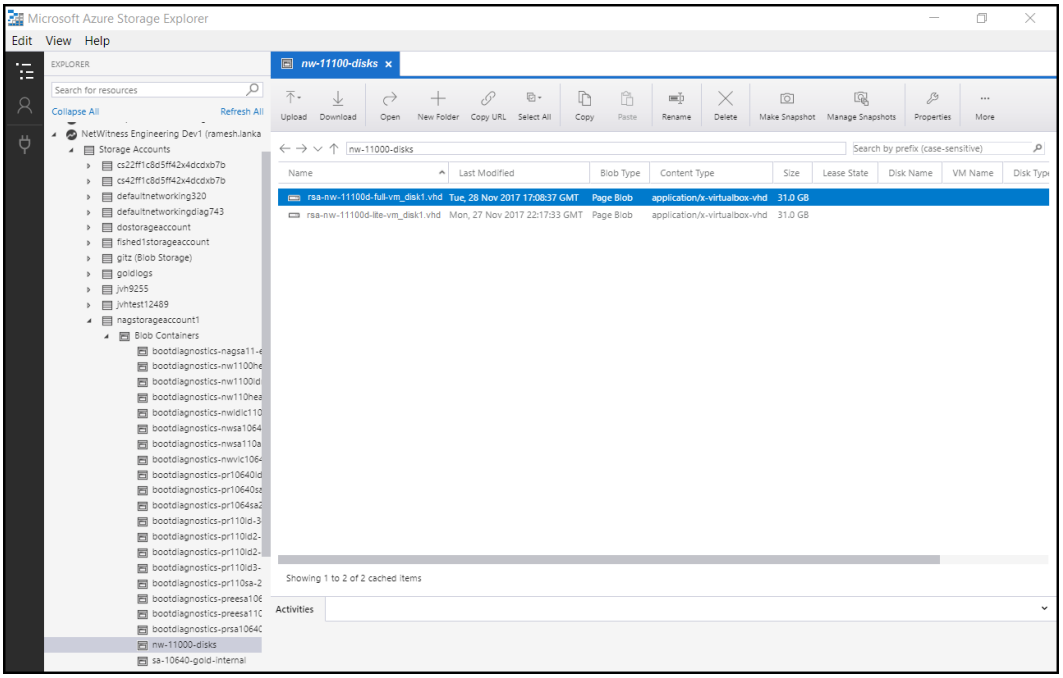
g. Start the copy: `Start-AzureStorageBlobCopy -AbsoluteUri "<SAS-URL>" -DestContainer <container-name> -DestBlob <destination-blob-name> -DestContext $targetStorageContext`

h. Obtain the Blob copy status by using the command: `Get-AzureStorageBlobCopyState -Blob "< destination-blob-name>" -Container "<container-name> " -Context $targetStorageContext`

4. Once the VHD's are successfully copied. You'll must create an image and a VM.

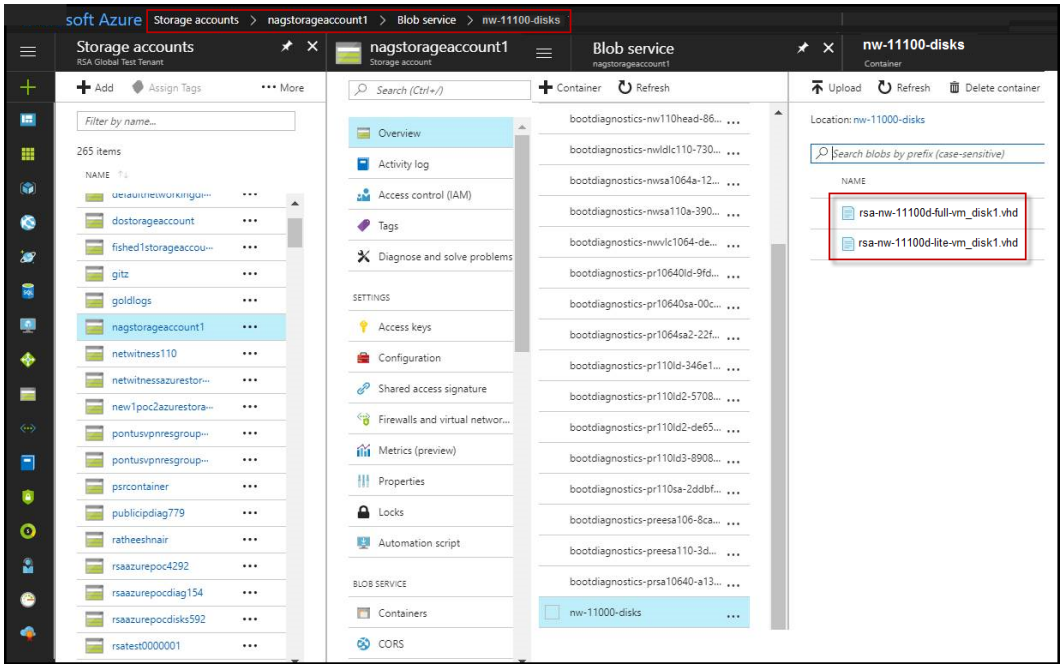
5. Verify if all the NW Server VHDs are uploaded into the Azure Cloud.

Note: Alternatively, you can use the Microsoft Azure Storage Explorer windows utility (http://storageexplorer.com/) to verify that all the VHDs from the following location subscription exist. This utility helps you manage the contents of your storage.



a. Log in to the Azure portal (https://portal.azure.com).

b. From the right panel, click **Storage accounts** > **netwitnessazurestorage1** > **Blob service** > **nwazurevhdstore**.

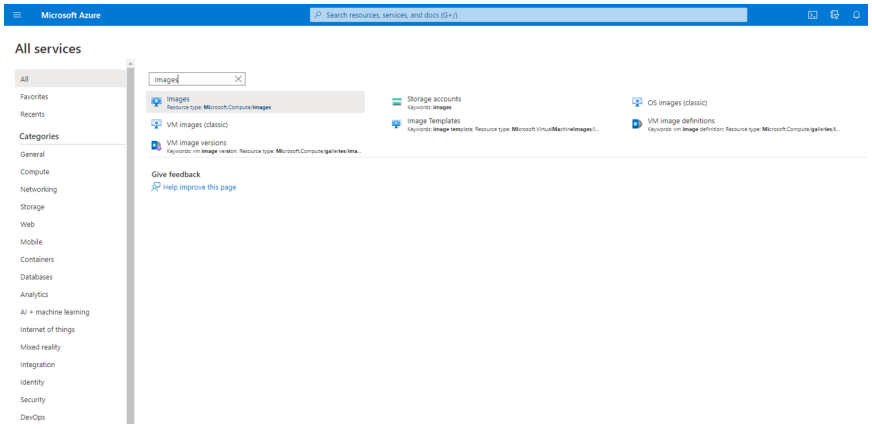


6. (Optional) In the Azure Explorer, go to the **NetWitness** group > **Storage Accounts** > **netwitnessazurestorage1**) > **Blob Containers** > **nwazurevhdstore**).

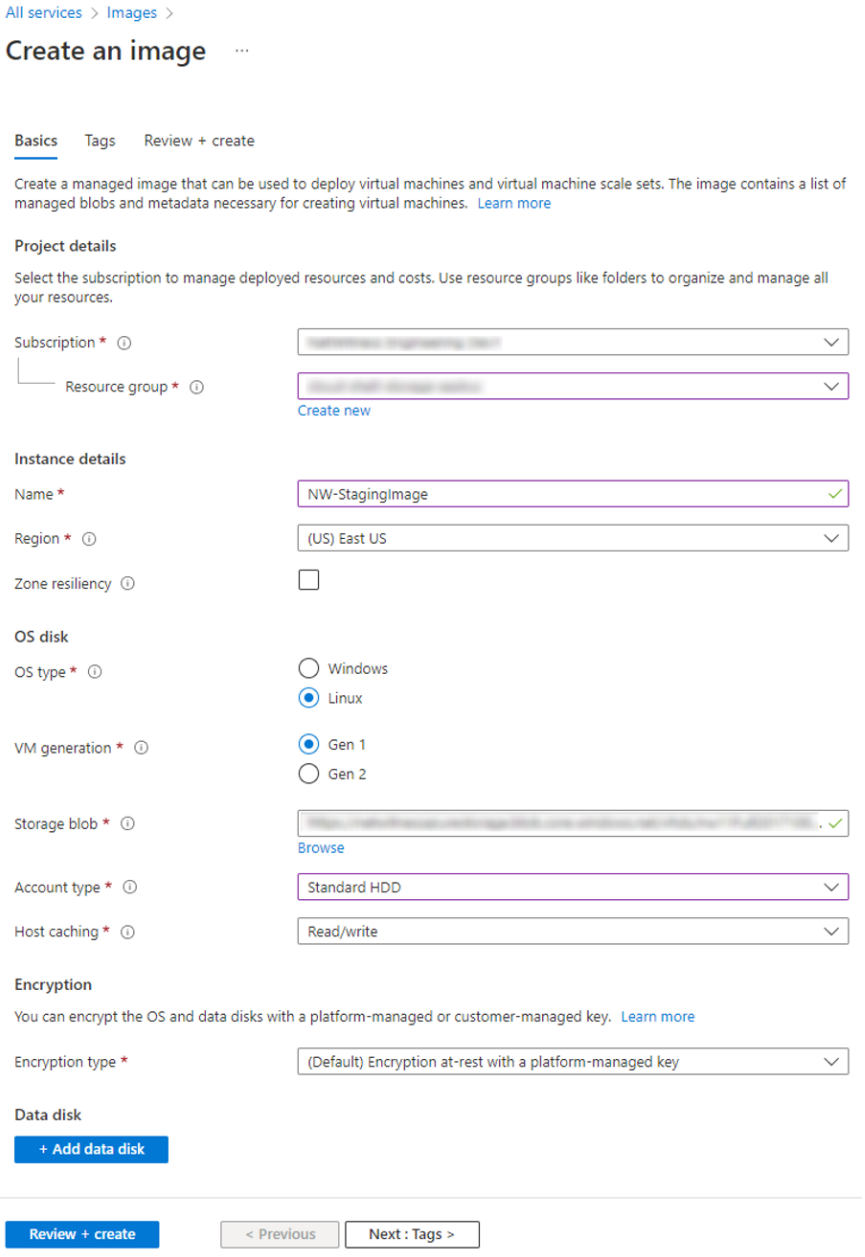
Task 2. - Create NW Server Image

To create a NW Server image in Azure from upload VHDs, perform the following steps:

1. Log in to https://portal.azure.com.
2. From the left panel, click **All Services** and filter by Images.
3. Click **Images**.

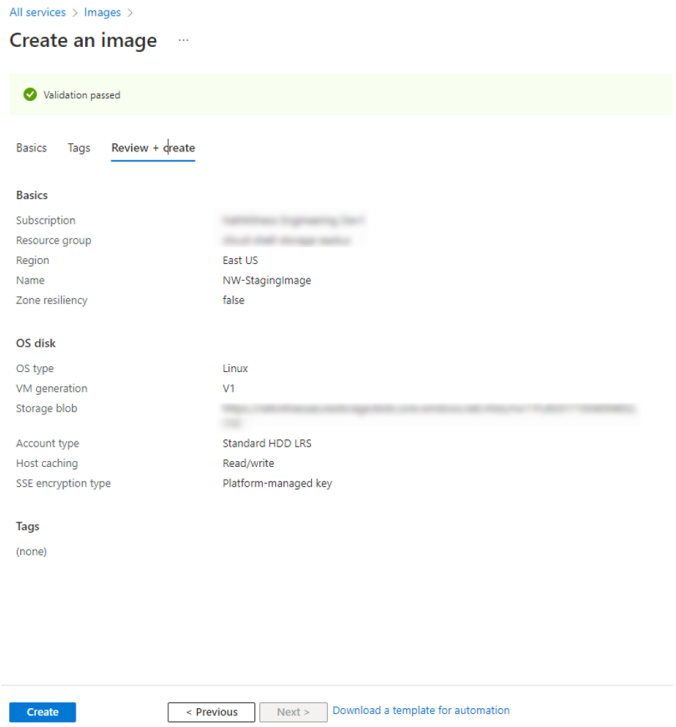


4. To create and configure the Image.
   a. Click **Create**.
   b. Enter an image **Name**, select the correct **Resource Group**, select a valid **Region**, and set the **OS Disk** to Linux.
      In the **Storage blob**, browse to the uploaded location of the VHDs .
   c. Make sure that **Standard (HDD)** is selected for **Account Type**.
      The following screen shot illustrates a completed **Create Image** view.



   d. Click **Next : Tags >** to add the tags for the Image (optional) and then Click **Review + create**.
      Azure does a validation check.

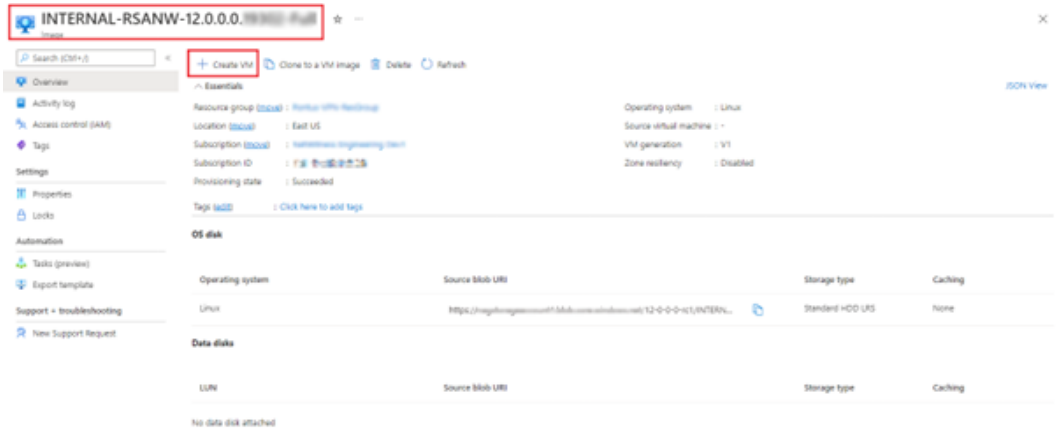
e. Click **Create** to create the image.
Check notifications on top right for the confirmation.



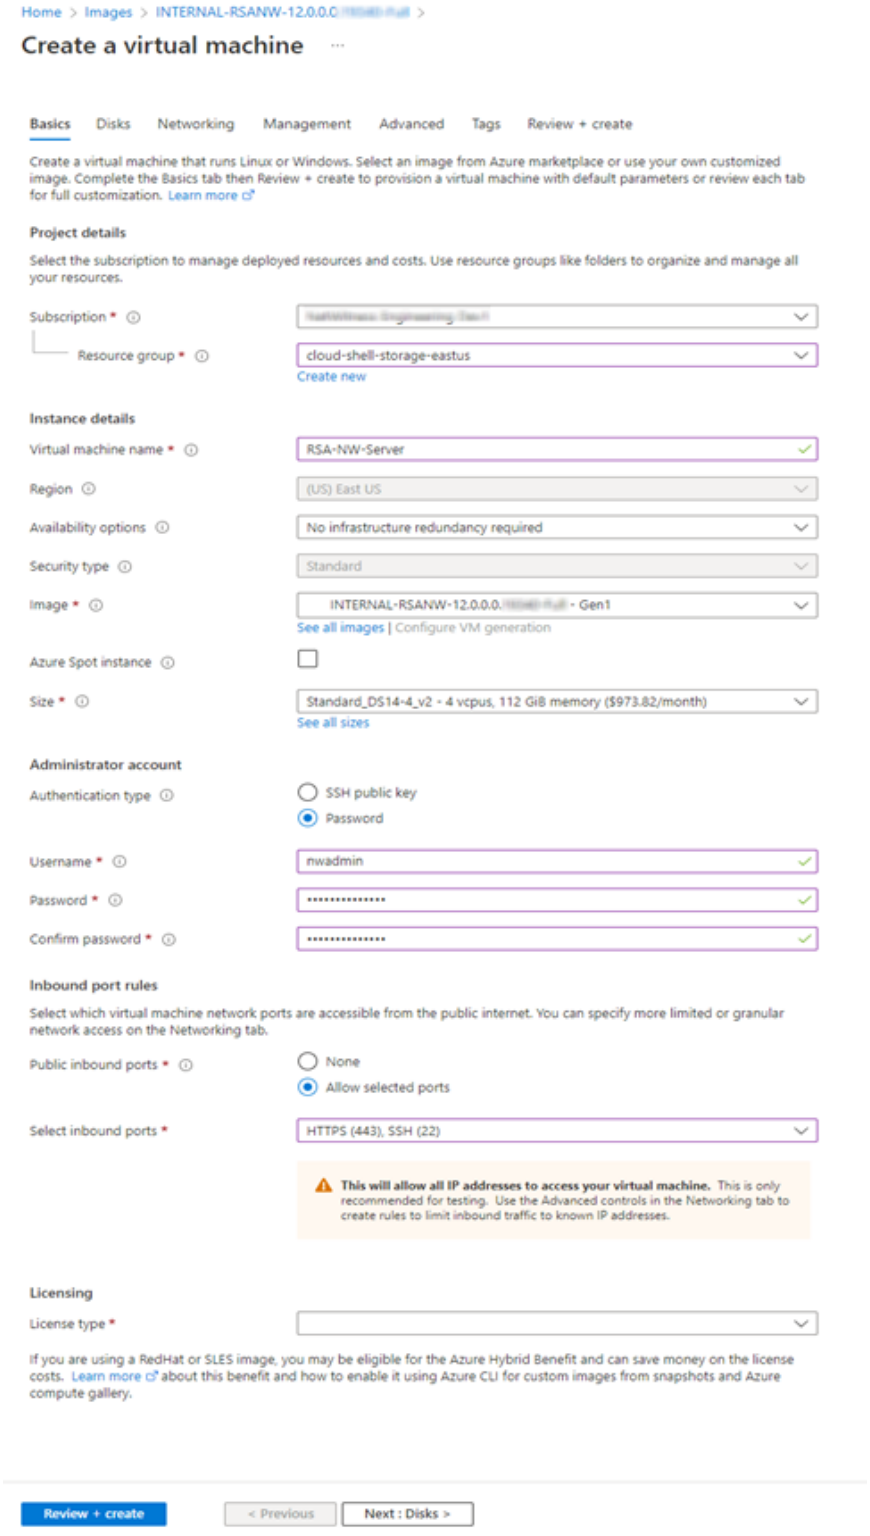
Task 3. Create Virtual Machine (VM)

To create a VM in Azure using the NW Server image:

1. Go to **Images** and click **Create VM**.



The **Basics** tab is displayed.

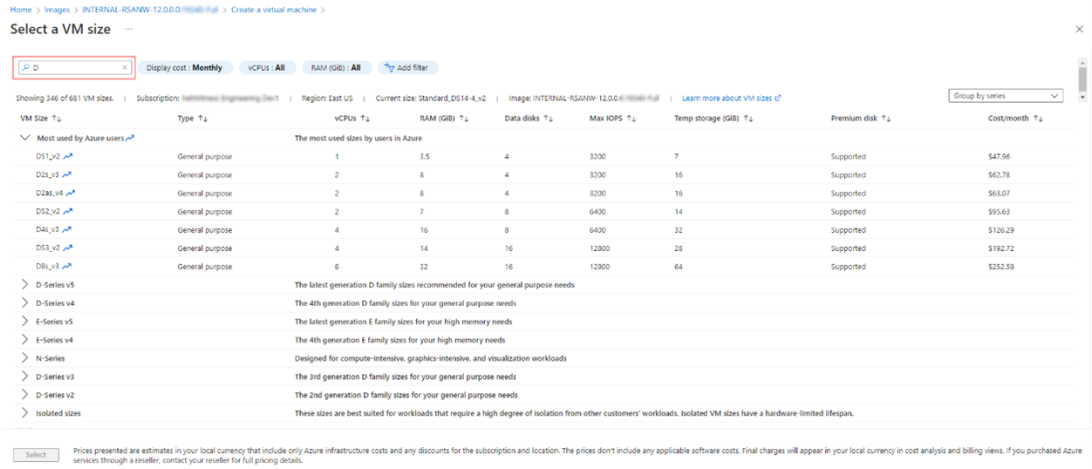


2. Enter the values in following fields.
   a. In the **Name** field, enter a user-defined name (for example, **RSA-NW-Server**).

> **Caution:** The username and password that you define is used to login to the system as a non-administrator user. Do not use the root user (the login does not have superuser permissions). You must change the root password the first time that you log in to the VM by executing the `su passwd root` command. This is a critical step and should not be missed. You cannot use `root` for a username (Azure-specific).

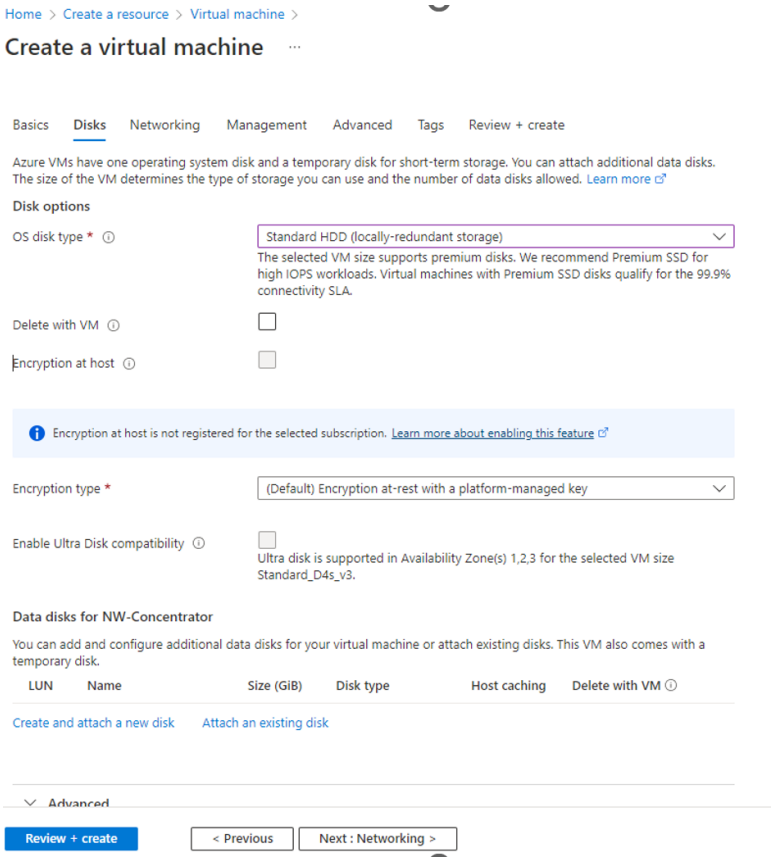
> **Note:** Make sure the values selected in the **Subscription**, **Resource group**, and **Region** fields are correct.

   b. Click **See all sizes** and select appropriate Size and Instance. The recommended instance for Concentrator is **Standard F8**.



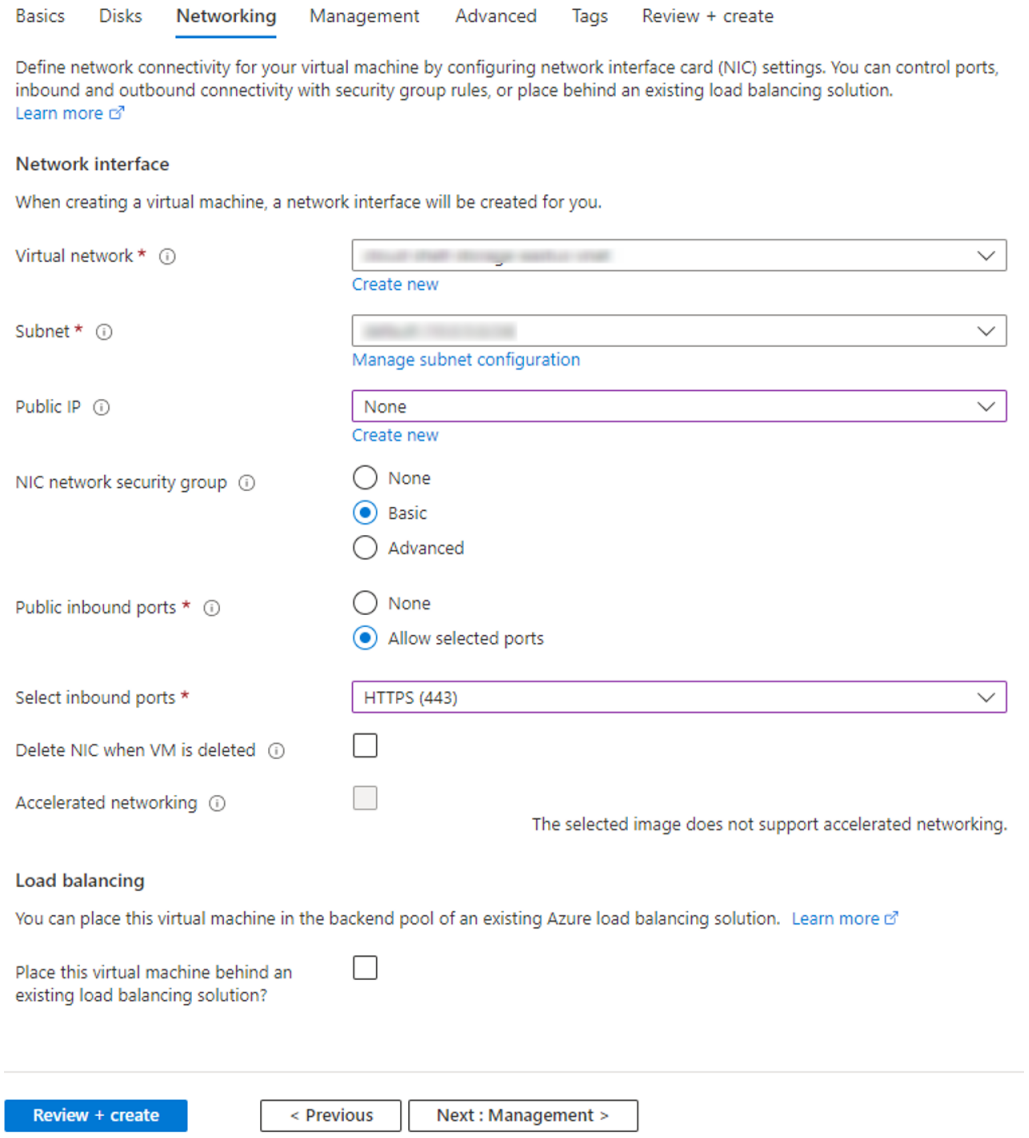
> **Note:** The sizing is based upon the capacity requirements of your enterprise. For more information on NetWitness VM size recommendations based on log capture rates, see Azure Configuration Recommendations. The minimum size NetWitness recommends for the NW Server is **F8 Standard**.

   c. In the **User name** field, enter a valid username.
   d. In the **Authentication type** field, click **Password** and enter a strong password that is a combination of lowercase, uppercase, numeral and a symbol (for example, **Password@123**).
   e. Click **Next** : **Disks >**.
   The **Disks** tab is displayed.
3. In the OS Disk type, select **Standard HDD** from the drop-down list and click **Next : Networking >**.



The **Networking** tab is displayed.

4. Click and define the fields.
   a. In the **Networking** tab, select:
   - A valid **Virtual network** and **Subnet**.

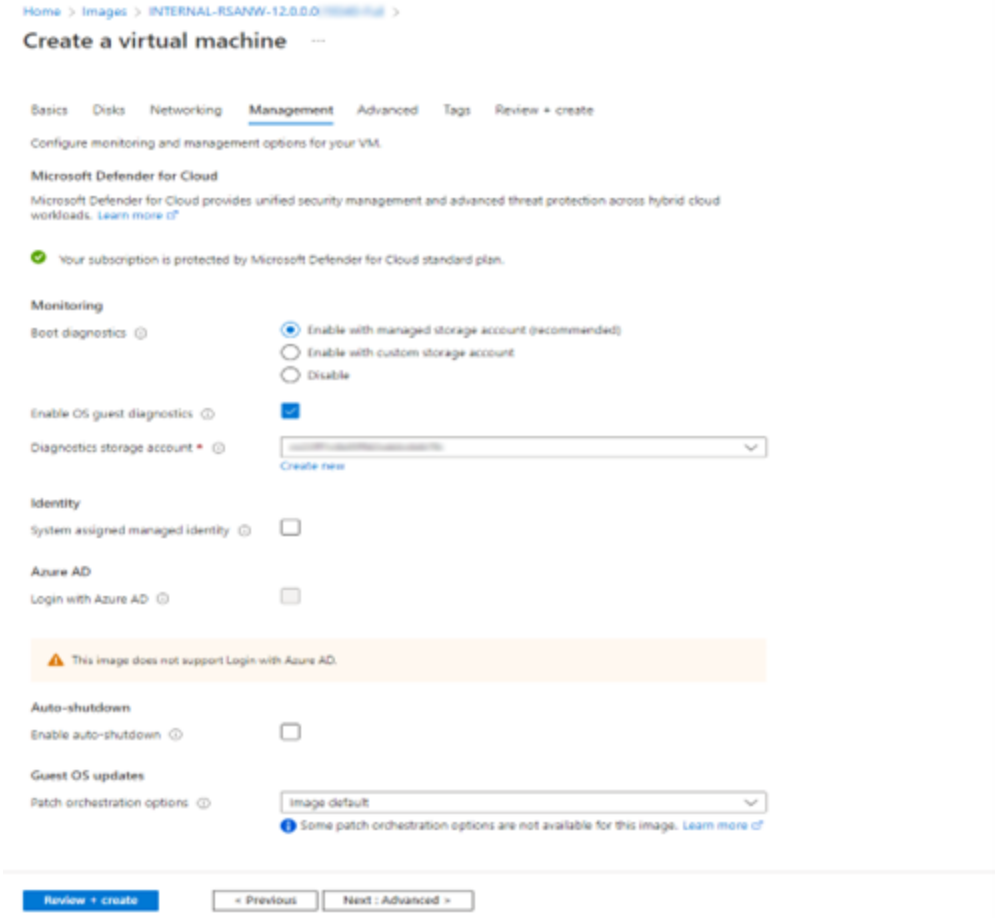


   - **None** for the **Public IP address**.
     NetWitness recommends **None** for the **Public IP address** (this is not mandatory). You can assign a public IP address, but it countermands Best Practices to assign a public IP to something that is based in the Azure Cloud.
   - A valid **Network security group**.
     For information on Network security groups, see the Microsoft Azure documentation (https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-nsg).

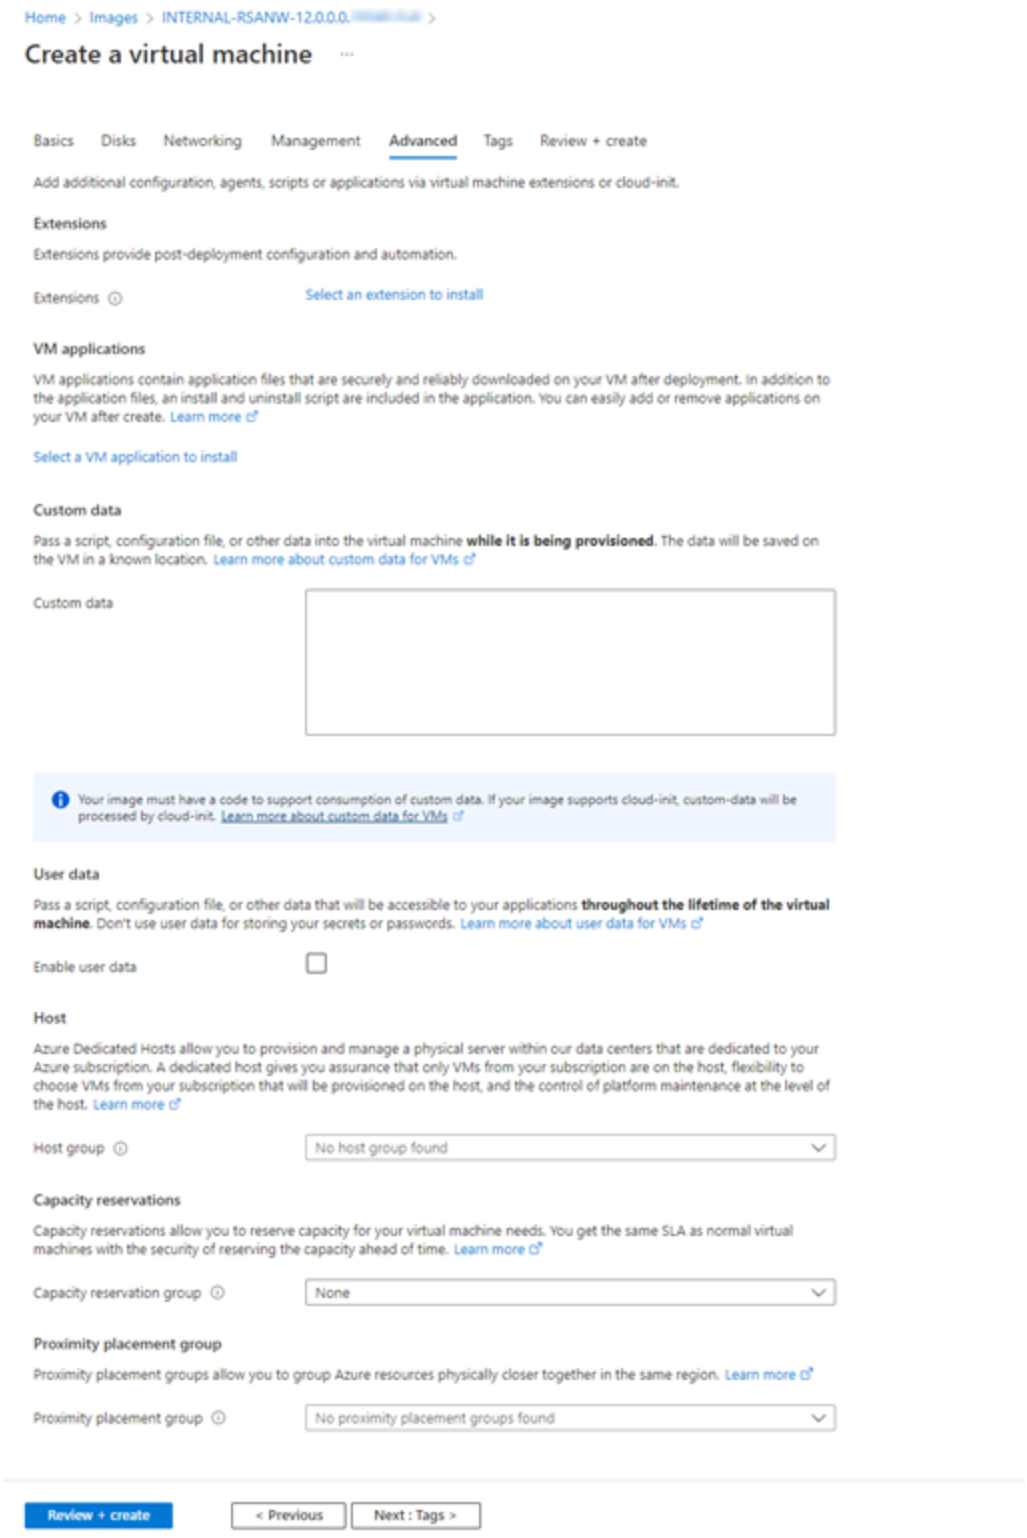
b. In the **Management** tab, select:

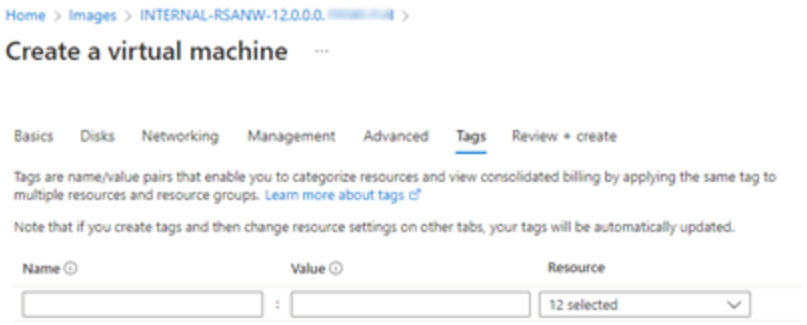
- **On** for **Boot Diagnostics**
- **On** for **Guest OS diagnostics**
- a valid **Diagnostics storage account**

The following figure illustrates a completed Settings panel.



Note: By default, the settings remain unchanged in the **Advanced** and **Tags** tab. Add any name and value pairs for tags based on requirement.


Home > Images > INTERNAL-RSANW-12.0.0.0.
Create a virtual machine
Basics
Disks
Networking
Management
Advanced
Tags
Review + create
Add additional configuration, agents, scripts or applications via virtual machine extensions or cloud-init.
Extensions
Extensions provide post-deployment configuration and automation.
Extensions
Select an extension to install
VM applications
VM applications contain application files that are securely and reliably downloaded on your VM after deployment. In addition to the application files, an install and uninstall script are included in the application. You can easily add or remove applications on your VM after create. Learn more
Select a VM application to install
Custom data
Pass a script, configuration file, or other data into the virtual machine while it is being provisioned. The data will be saved on the VM in a known location. Learn more about custom data for VMs
Custom data
Your image must have a code to support consumption of custom data. If your image supports cloud-init, custom-data will be processed by cloud-init. Learn more about custom data for VMs
User data
Pass a script, configuration file, or other data that will be accessible to your applications throughout the lifetime of the virtual machine. Don't use user data for storing your secrets or passwords. Learn more about user data for VMs
Enable user data
Host
Azure Dedicated Hosts allow you to provision and manage a physical server within our data centers that are dedicated to your Azure subscription. A dedicated host gives you assurance that only VMs from your subscription are on the host, flexibility to choose VMs from your subscription that will be provisioned on the host, and the control of platform maintenance at the level of the host. Learn more
Host group
No host group found
Capacity reservations
Capacity reservations allow you to reserve capacity for your virtual machine needs. You get the same SLA as normal virtual machines with the security of reserving the capacity ahead of time. Learn more
Capacity reservation group
None
Proximity placement group
Proximity placement groups allow you to group Azure resources physically closer together in the same region. Learn more
Proximity placement group
No proximity placement groups found
Review + create
< Previous
Next : Tags >


Home > Images > INTERNAL-RSANW-12.0.0.0.
Create a virtual machine
Basics
Disks
Networking
Management
Advanced
Tags
Review + create
Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. Learn more about tags
Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.
Name
Value
Resource
12 selected

5. Click **Review + Create** after the validation is successful.



Home > Images > INTERNAL-RSANW-12.0.0.0. >

Create a virtual machine ...

Validation passed

Basics Disks Networking Management Advanced Tags Review + create

INTERNAL-RSANW-12.0.0.0.
Image

Standard DS14-4 v2
4 vcpus, 112 GiB memory

Basics

Subscription	
Resource group	
Virtual machine name	RSA-NW-Server
Region	East US
Availability options	No infrastructure redundancy required
Security type	Standard
Image	INTERNAL-RSANW-12.0.0.0.
Size	Standard DS14-4 v2 (4 vcpus, 112 GiB memory)
Authentication type	Password
Username	nwadmin
Public inbound ports	HTTPS
Azure Spot	No

Disks

OS disk type	Standard HDD LRS
Use managed disks	Yes
Delete OS disk with VM	Disabled
Ephemeral OS disk	No

Networking

Virtual network	
Subnet	
Public IP	None
Accelerated networking	Off
Place this virtual machine behind an existing load balancing solution?	No
Delete NIC when VM is deleted	Disabled

Management

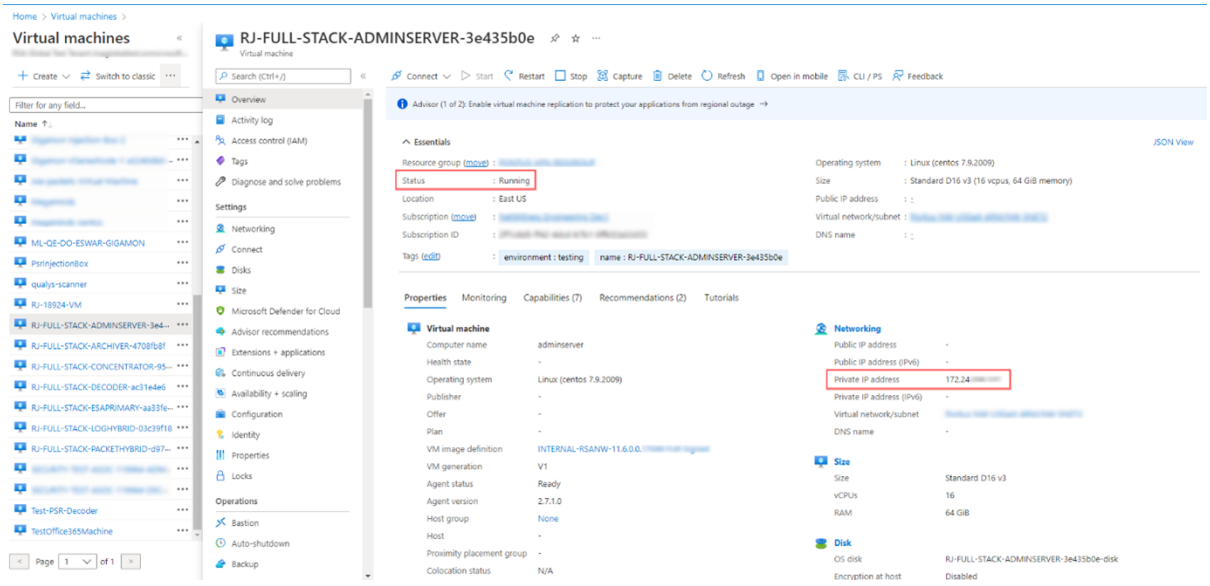
Microsoft Defender for Cloud	Standard
Boot diagnostics	On
Enable OS guest diagnostics	On
Diagnostics storage account	
System assigned managed identity	Off
Login with Azure AD	Off
Auto-shutdown	Off
Enable hotpatch	Off
Patch orchestration options	Image Default

Advanced

Extensions	None

The NW Server VM Deployment is successful when you see the VM status as **Running**.

6. Click **Overview** on the Virtual Machine to view all the required details such as VM status and IP Address.



7. SSH to the VM using the username that you specified in Step 2d of Task 3 and reset the **root** password. Use the `su passwd root` command string to reset the root password.

```
[psradmin@12 ~]$ sudo passwd root
[sudo] password for psradmin:
Changing password for user root.
New password:
BAD PASSWORD: The password contains more than 4 characters of the same class consecutively
Retype new password:
passwd: all authentication tokens updated successfully.
```

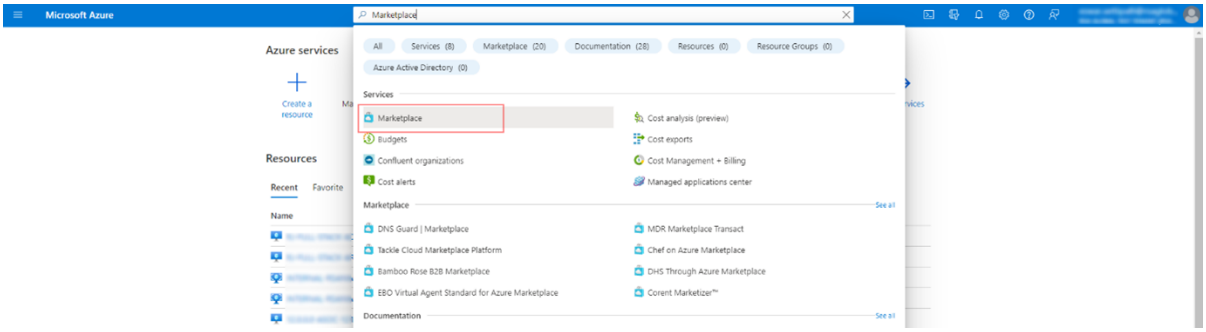
8. Close the current SSH session and open a new SSH session with **root** using the username and the password created in the previous step.

> **Note:** Step 8 is a critical, one-time step for a new deployment. If you do not complete this step, the NetWitness User Interface will not load.

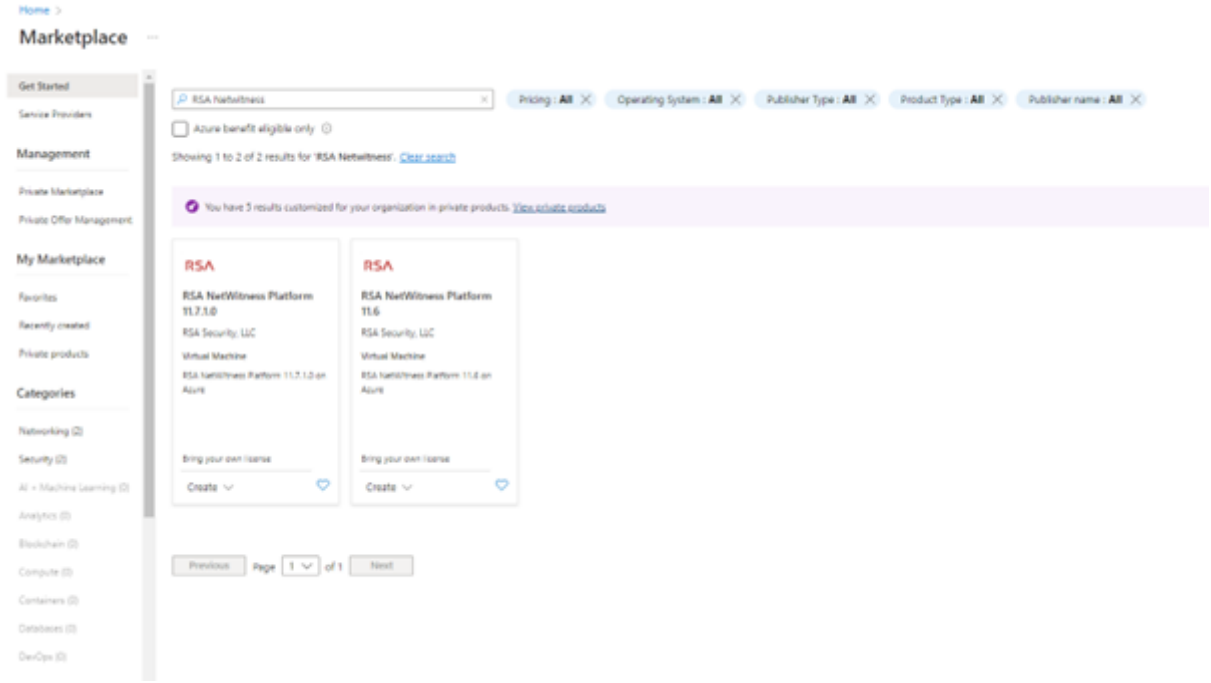
Deploy Component Core Services in Azure

You must perform the following tasks to configure the core NetWitness component services on a virtual machine (VMs) in the Azure Cloud environment.

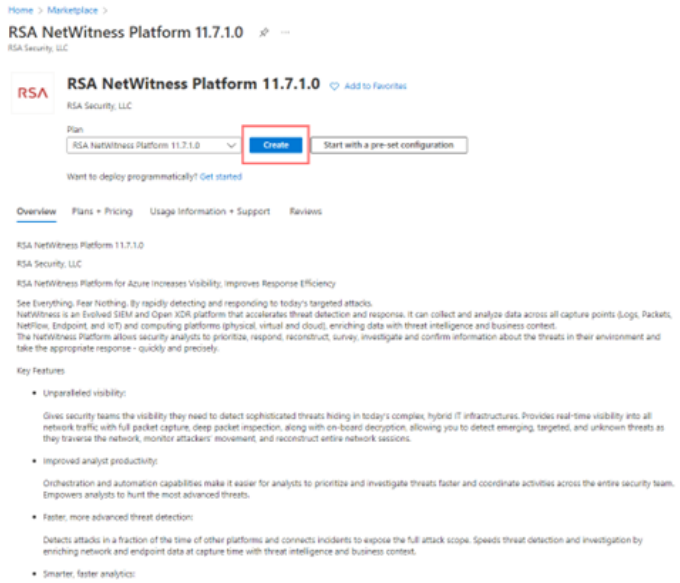
1. Search for Marketplace in Azure portal.



2. Search for **RSA NetWitness**.

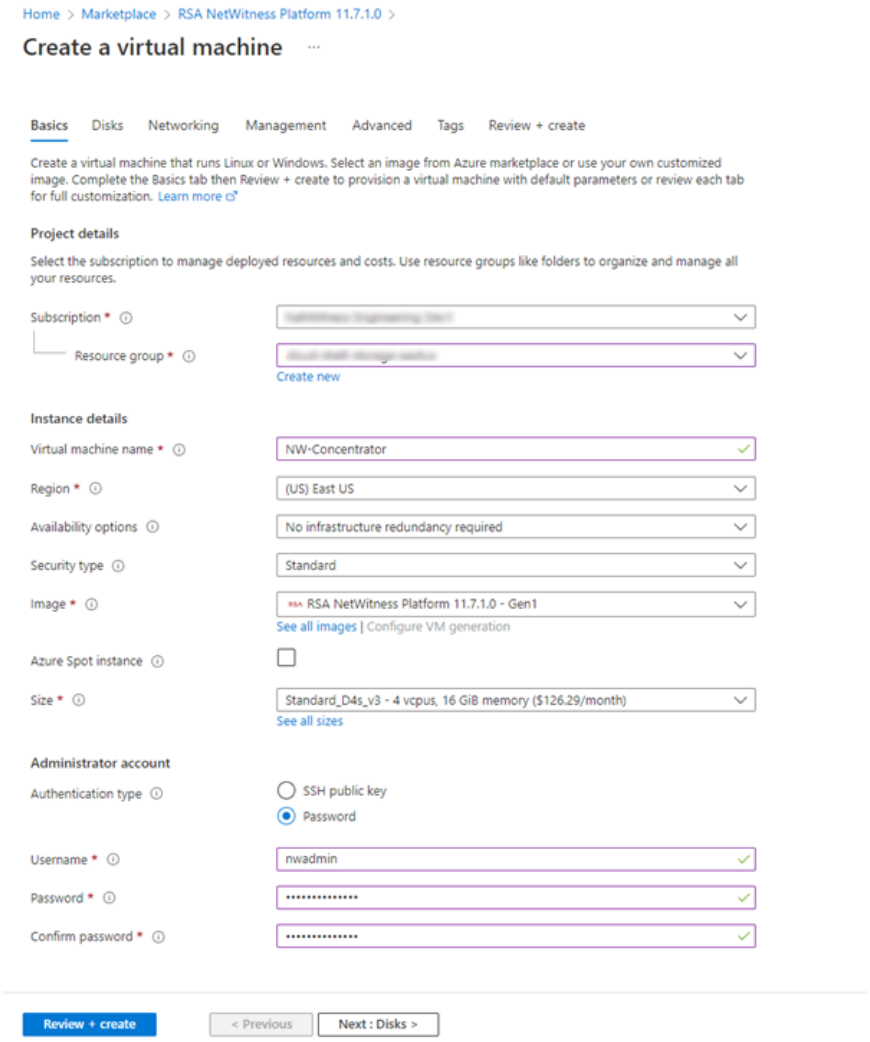


3. Select RSA NetWitness Platform 12.0.0.0 and click **Create**.



The **Create virtual machine** wizard opens and displays the **Basics** tab.

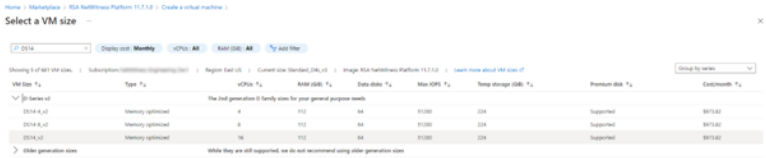
4. Enter the values in the following fields:
   a. Specify a VM **Name** (for example, **NW-Concentrator**).
   b. Select **Password** for **Authentication type**.
   c. Enter your credentials (**User name** and **Password**) and **Confirm Password**.
   d. Click **OK**.



Azure validates the **Basic** specifications and the **2 Size** page is displayed.

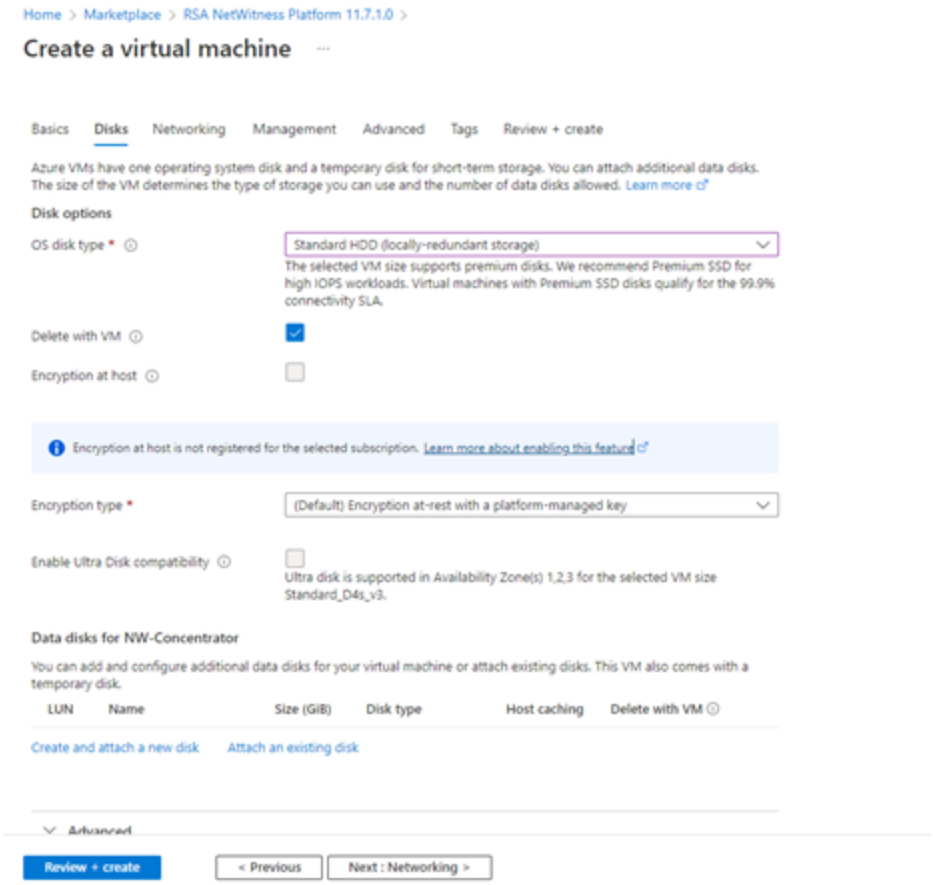
5. Click on the appropriate VM size (for example, **Standard DS14 v2** for the Concentrator) for the service and click **Select** for a VM **Size**.

   For more information on NetWitness recommendations of the VM sizes for each service, see Azure Configuration Recommendations.



   Azure validates the **Size** specifications. Click **Next : Disks >**.

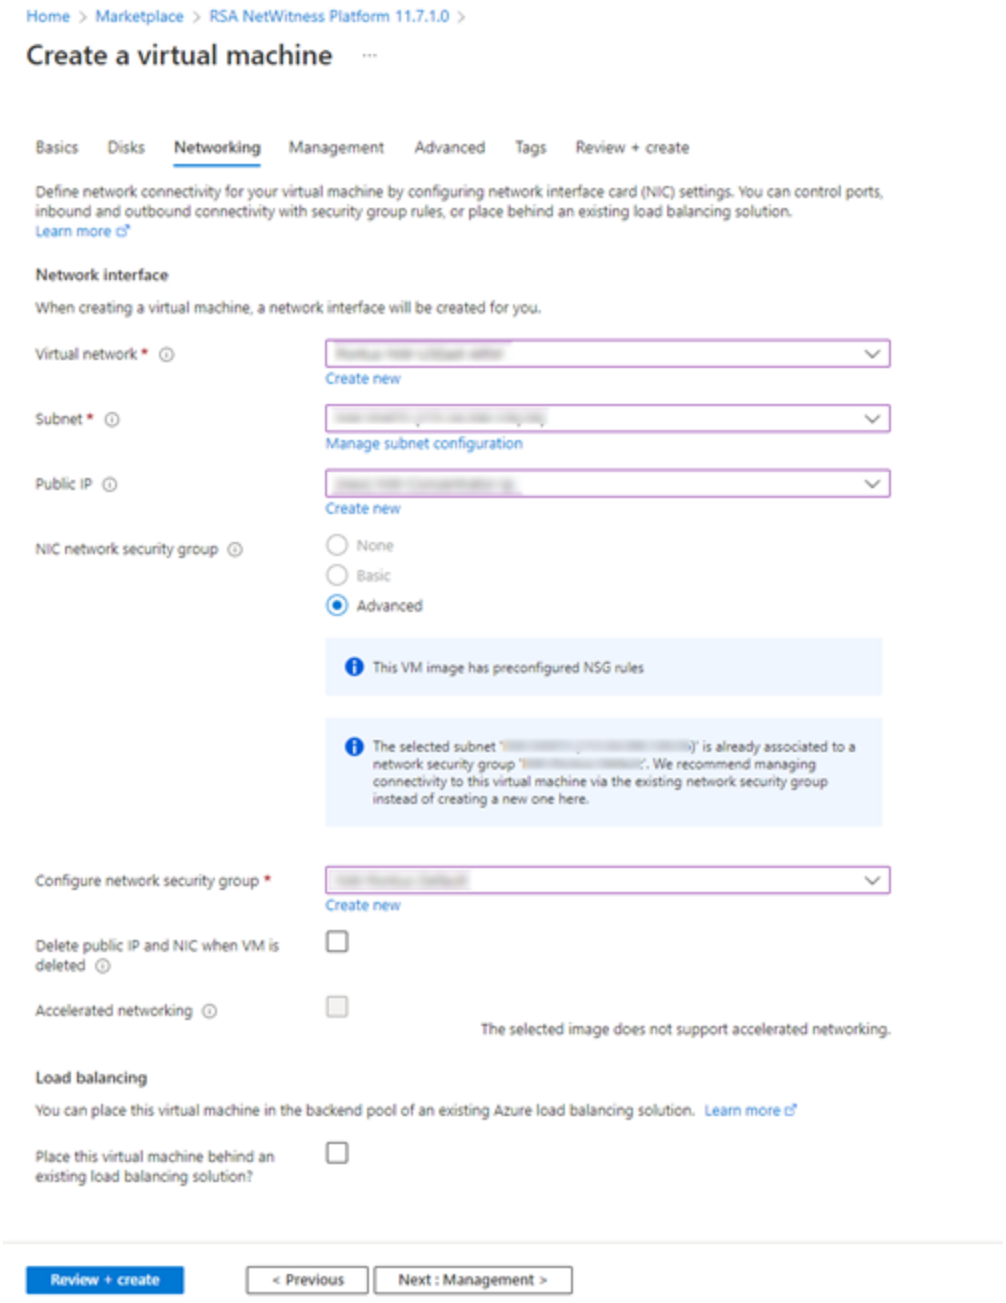
6. Under **Disks** tab, Select **SSD** for the **VM disk type** of the Concentrator or **HDD** for all other components. Solid State Disk (SSD) performs better than a Hard Drive (HDD).



Click **Next : Networking >**.

7. In the **Networking** tab:
   a. Adjust **Virtual network**, **Subnet**, and **Public IP address** according to the requirements of your network.
   b. Enabling **Accelerated Networking** is recommended for Decoder hosts with higher line rates (> **800**Mbps).
   c. Specify a valid **Network Security group**.

> **Note:** For information on Network security groups, see the Microsoft Azure documentation (https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-nsg). Refer to *Deployment: Network Architecture and Ports* (https://community.netwitness.com/t5/netwitness-platform-online/network-architecture-and-ports/ta-p/668996) for a comprehensive list of the ports you must set up for all NetWitness components.

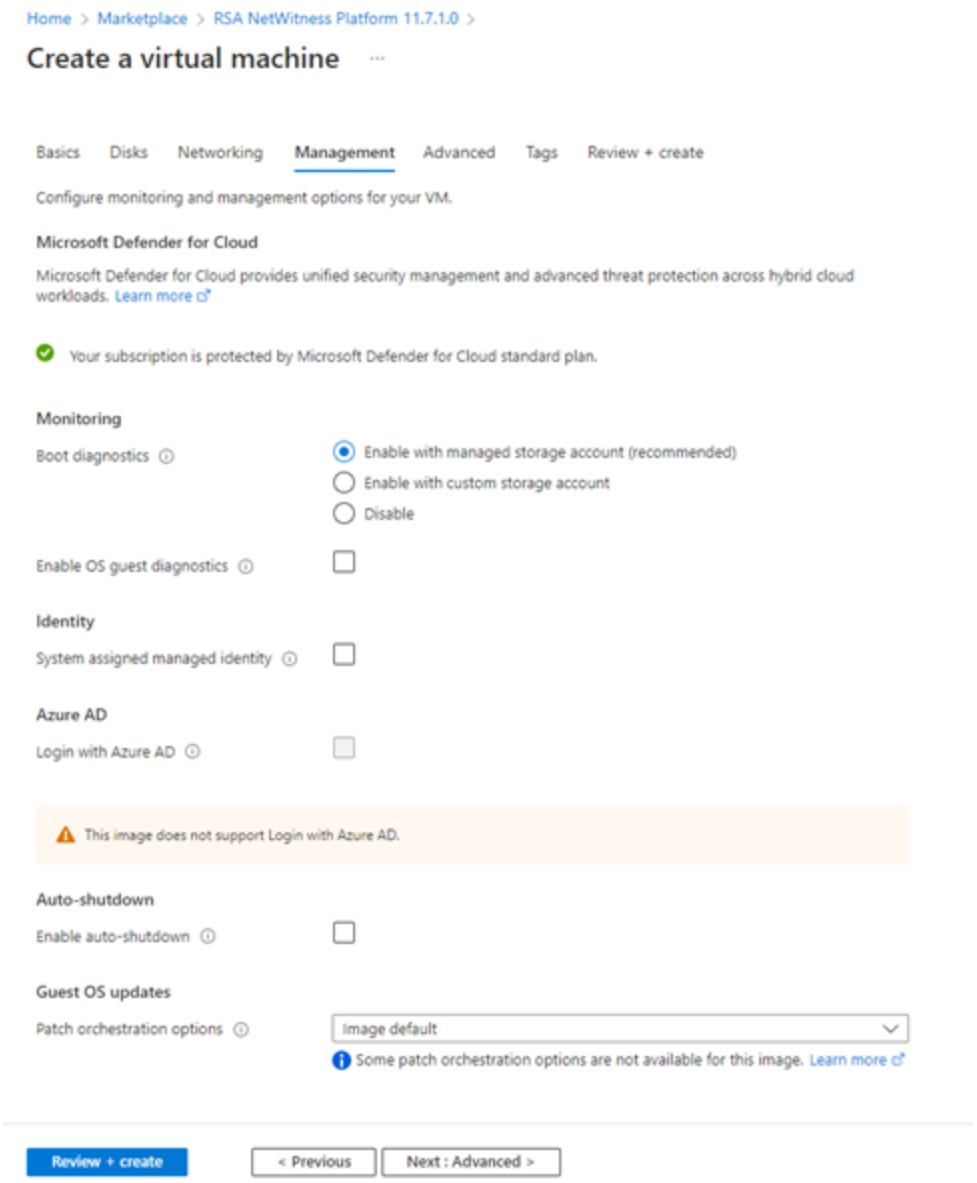


d. After completing the configurations, click **Next : Management >**.

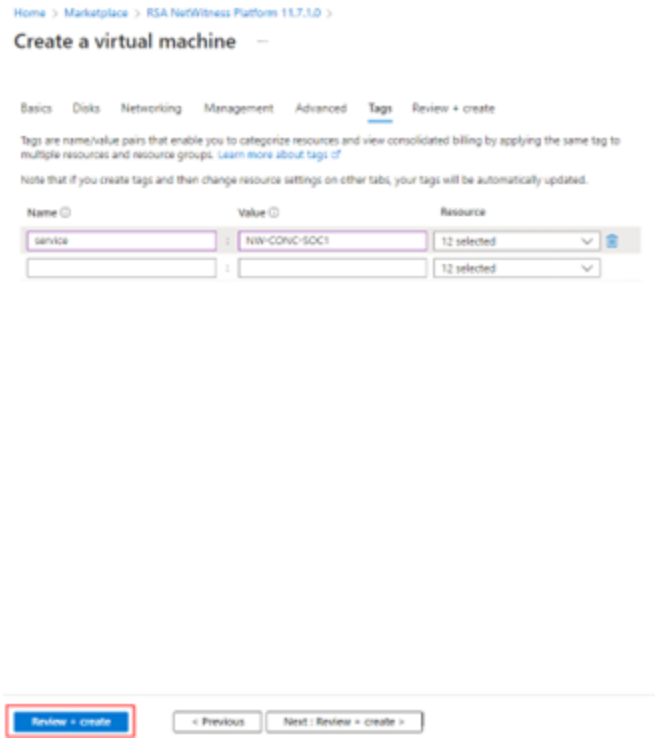
8. In the **Management** Tab:

   a. Enable Boot diagnostics and OS guest diagnostics.

   b. Configure Identity based on your requirements.



9. Click **Next : Advanced >** to enter **Advanced** settings tab. Make the required changes if any and click **Next : Tags >**.
10. In the **Tags** menu, add **Name**, **Value** pairs if any and click **Review + create**.



Azure validates the VM and displays the status check.


Home > Images > INTERNAL-RSANW-12.0.0.0.
Create a virtual machine
Validation passed
Basics Disks Networking Management Advanced Tags Review + create
INTERNAL-RSANW-12.0.0.0.
Image
Standard DS14-4 v2
4 vcpus, 112 GiB memory
Basics
Subscription
Resource group
Virtual machine name RSA-NW-Server
Region East US
Availability options No infrastructure redundancy required
Security type Standard
Image INTERNAL-RSANW-12.0.0.0.
Size Standard DS14-4 v2 (4 vcpus, 112 GiB memory)
Authentication type Password
Username nwadmin
Public inbound ports HTTPS
Azure Spot No
Disks
OS disk type Standard HDD LRS
Use managed disks Yes
Delete OS disk with VM Disabled
Ephemeral OS disk No
Networking
Virtual network
Subnet
Public IP None
Accelerated networking Off
Place this virtual machine behind an existing load balancing solution? No
Delete NIC when VM is deleted Disabled
Management
Microsoft Defender for Cloud Standard
Boot diagnostics On
Enable OS guest diagnostics On
Diagnostics storage account
System assigned managed identity Off
Login with Azure AD Off
Auto-shutdown Off
Enable hotpatch Off
Patch orchestration options Image Default
Advanced
Extensions None
VM applications None
Cloud init No
User data No
Proximity placement group None
Capacity reservation group None
Create
< Previous
Next >
Download a template for automation

11. Click **Create** to deploy the NW-Concentrator VM in Azure.
12. Configure the host VM in NetWitness 12.0.0.0.
13. Repeat steps 1 through 12 inclusive for the rest of the core NetWitness component services.

Installation Tasks

Before you begin the installation tasks make sure you open the firewall ports. For more information on the lists of all the ports in a deployment, see the"Network Architecture and Ports" topic in the *Deployment Guide for NetWitness Platform 12.0.*

Caution: Do not proceed with the installation until the ports on your firewall are configured.

Install 12.0.0.0 on the NetWitness Server (NW Server) and Component Hosts

Note: You can perform this task for **INTERNAL-RSANW-12.0.0.0.19340-Full** signed instance.

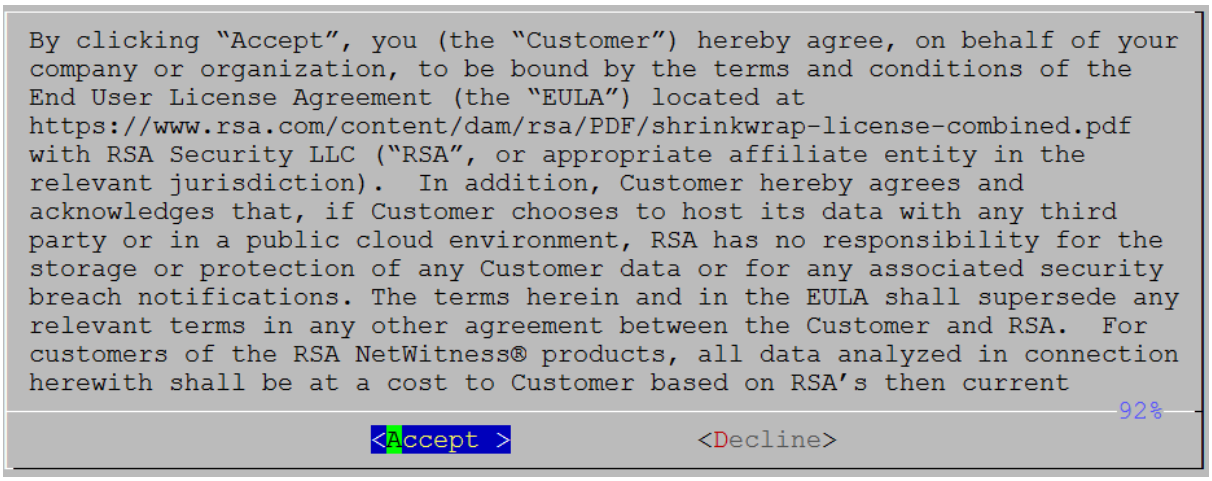
Caution: If you want to install the Endpoint Relay Server, do not run the `nwsetup-tui` script. Follow the instructions in "(Optional) Installing and Configuring Relay Server" in the *NetWitness Endpoint Configuration Guide*.

IMPORTANT: In NetWitness Platform version 11.6 or later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`. If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

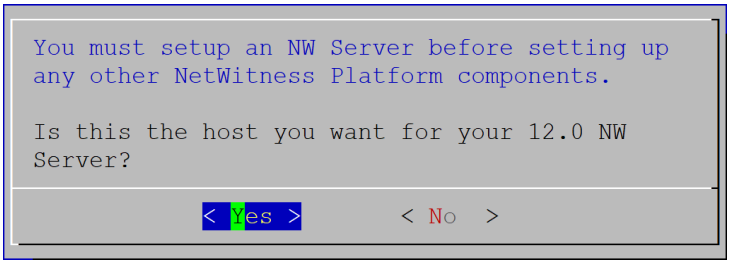
1. Log in to the host with the `root` credentials and run the `nwsetup-tui` command to set up the host.

   This initiates the `nwsetup-tui` (Setup program) and the EULA is displayed.

   Note: Use the following options to navigate the Setup prompts.
   1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, and use the Tab key to move to and from commands (such as **<Yes>**, **<No>**, **<OK>**, and **<Cancel>**). Press **Enter** to register your command response and move to the next prompt.
   2.) The Setup program adopts the color scheme of the desktop or console you use to access the host.
   3.) If you specify DNS servers during the Setup program (`nwsetup-tui`) execution, they MUST be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` script to proceed. Any misconfigured DNS servers cause the Setup program to fail. If you need to reach a DNS server after setup that is unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see "Change Host Network Configuration" topic in the System Maintenance Guide.
   If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Update Repository** prompt in step 12 (the DNS servers are not defined so the system cannot access the external repo).



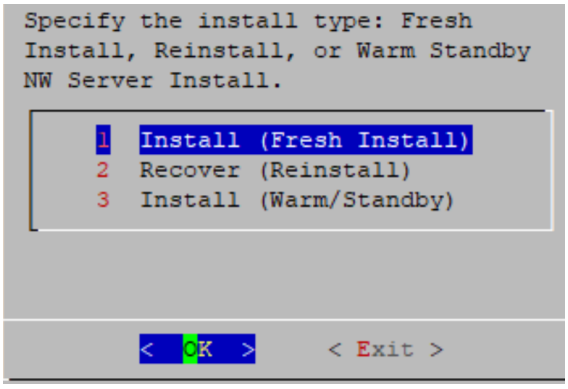
2. Tab to **Accept** and press **Enter**.
   The **Is this the host you want for your 12.0 NW Server** prompt is displayed.



3. Tab to **Yes** and press **Enter** to install 12.0 on the NW Server.
   Tab to **No** and press **Enter** to install 12.0 on other component hosts.

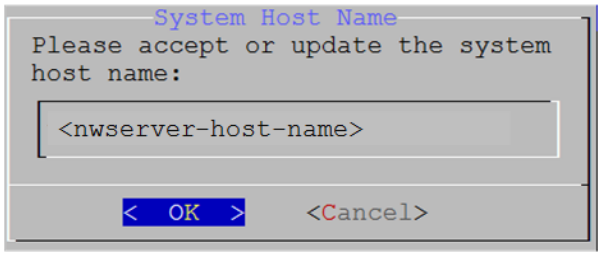
   Caution: If you choose the wrong host for the NW Server and complete the Setup, you must restart the Setup Program (step 2) and complete steps all the subsequent steps to correct this error.

4. The **Install** prompt is displayed (**Recover** does not apply to the installation. It is for 12.0 Disaster Recovery.).
   NW Server Host prompt:



   Other Component Hosts, the prompt is the same, but does not include option 3 Install (Warm/Standby)

5. Press **Enter**. **Install (Fresh Install)** is selected by default.
   The **System Host Name** prompt is displayed.
   NW Server prompt:

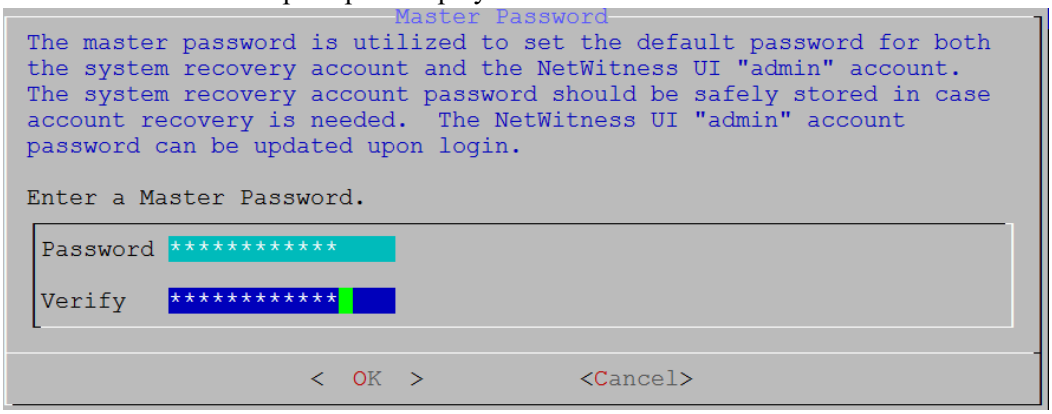


   Other Component Hosts prompt says <non-nwserver-host-name>

   Caution: If you include "." in a host name, the host name must also include a valid domain name.

   Press **Enter** if want to keep this name. If not, edit the host name, tab to **OK**, and press **Enter** to change it.

6. **This step applies only to NW Server hosts**.
   The **Master Password** prompt is displayed.



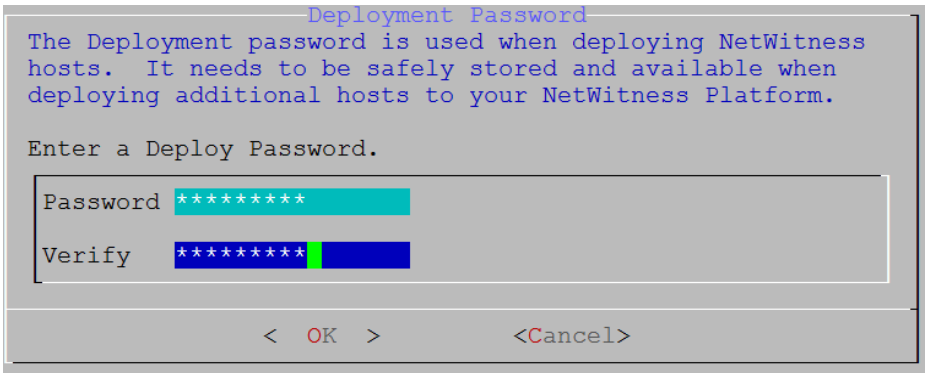
   The following list of characters are supported for Master Password and Deployment Password:

   - Symbols: **! @ # % ^ +**
   - Numbers: **0-9**
   - Lowercase Characters: **a-z**
   - Uppercase Characters: **A-Z**

   No ambiguous characters are supported for Master Password and Deployment Password. For example:
   space { } [ ] ( ) / \ ' " ` ~ ; : . < > -

   Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

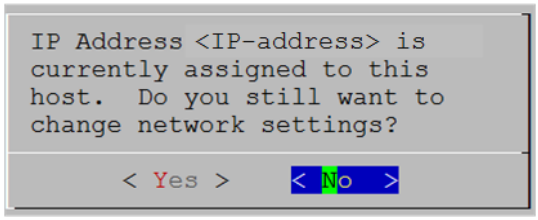
7. **This step applies to both NW Server hosts and component hosts.**
   The **Deployment Password** prompt is displayed.



Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

8. One of the following conditional prompts is displayed.

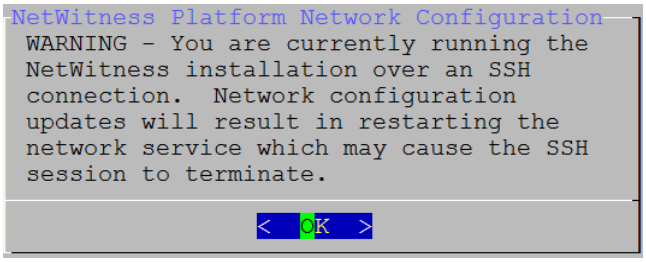
- If the Setup program finds a valid IP address for this host, the following prompt is displayed.



Press **Enter** if you want to use this IP and avoid changing your network settings. Tab to **Yes** and press **Enter** if you want to change the IP configuration on the host.

- If you are using an SSH connection, the following warning is displayed.

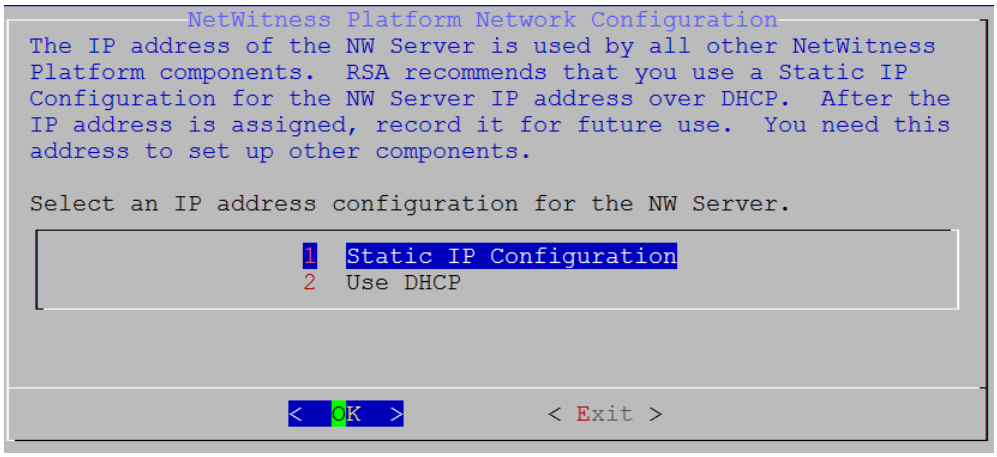
Note: If you connect directly from the host console, the following warning is not displayed.



Press **Enter** to close warning prompt.

- If the Setup Program finds an IP configuration and you choose to use it, the **Update Repository** prompt is displayed. Go to step 12 and complete the installation.
- If the Setup Program did not find an IP configuration or if you choose to change the existing IP configuration, the **Network Configuration** prompt is displayed.

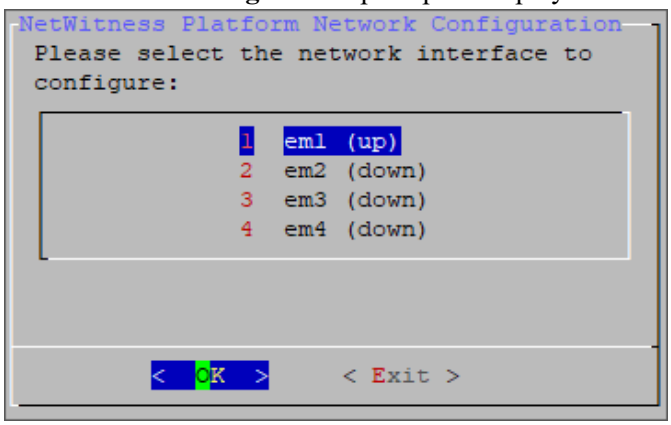
Caution: Only select "**Use DHCP**" as an IP address configuration for the NW Server if DHCP issues static IP addresses.



Tab to **OK** and press **Enter** to use **Static IP**.

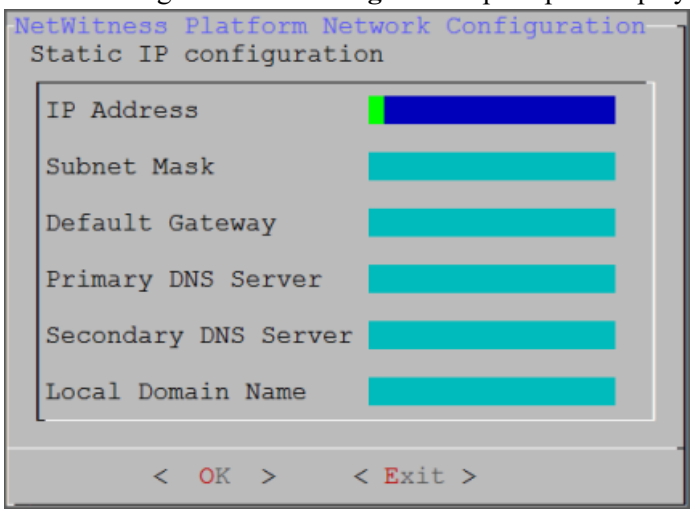
If you want to use DHCP, down arrow to **2 Use DHCP** and press **Enter**.

The **Network Configuration** prompt is displayed.



9. Down arrow to the network interface you want, tab to **OK**, and press **Enter**. If you do not want to continue, tab to **Exit.**

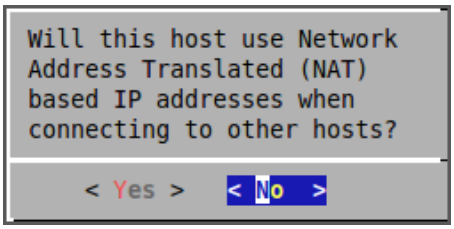
   The following **Static IP Configuration** prompt is displayed.



10. Type the configuration values, tab to **OK**, and press **Enter**. If you do not complete all the required fields, an `All fields are required` error message is displayed (**Secondary DNS Server** and **Local Domain Name** fields are not required). If you use the wrong syntax or character length for any of the fields, an `Invalid <field-name>` error message is displayed.

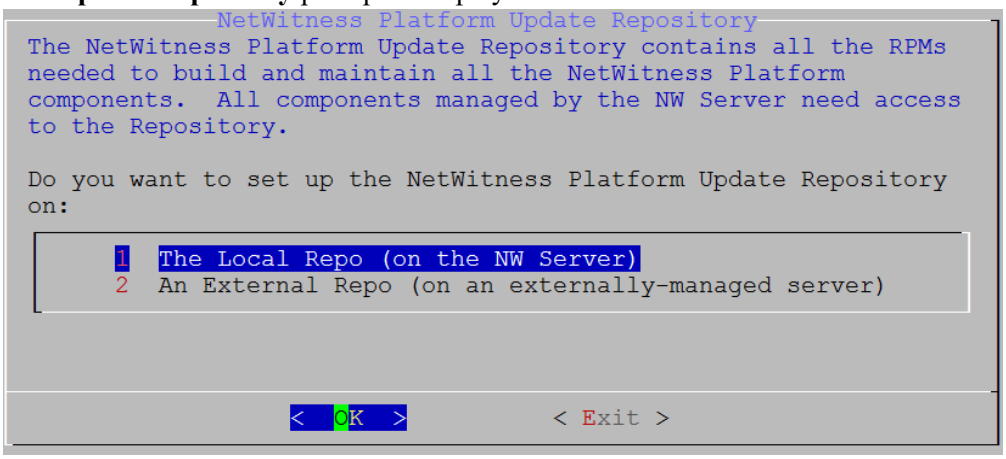
    Caution: If you select **DNS Server**, make sure that the DNS Server is correct and the host can access it before proceeding with the installation.

11. The Use Network Address Translation (NAT) prompt is displayed.



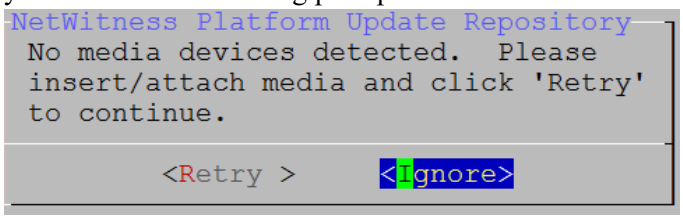
    For the NW Server, tab to **No** and press **Enter**.
    For component hosts, if this host requires the use of NAT-based addresses to communicate with the NW Server, tab to **Yes**. Otherwise, tab to **No** and press **Enter**.

12. The **Update Repository** prompt is displayed.

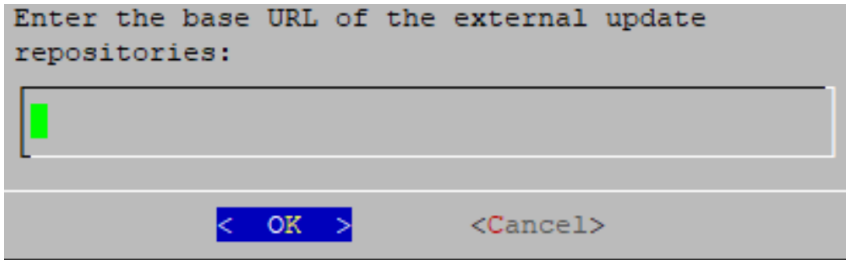


    For the NW Server:

    - Press **Enter** to choose the **Local Repo**.
    - If you want to use an external repo, down arrow to **External Repo**, tab to **OK**, and press **Enter**. If you select **1 The Local Repo (on the NW Server)** in the Setup program, make sure that you have the appropriate media attached to the host (media that contains the ISO file, for example a build stick) from which it can install NetWitness 12.0. If the program cannot find the attached media, you receive the following prompt.



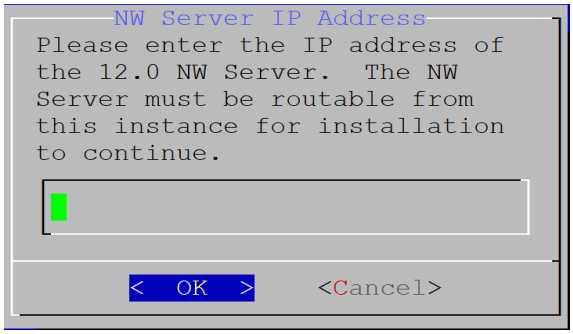
- If you select **2 An External Repo (on an externally-managed server)**, the UI prompts you for a URL. The repositories give you access to RSA updates and CentOS updates. Refer to "Appendix B. Create an External Repo" in this guide for instructions on how to create this repo and its external repo URL so you can enter it in the following prompt.



Enter the base URL of the NetWitness external repo and click **OK.** The **Start Install** prompt is displayed.

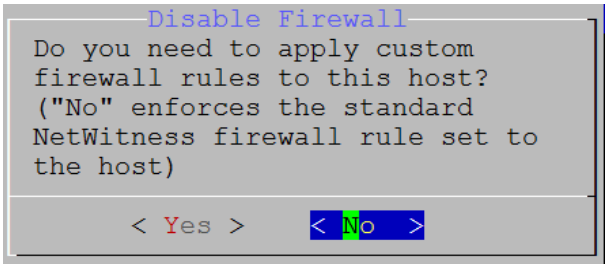
For component hosts:

- Select the same repo that you selected when you installed the NW Server host and follow the steps above.
- The NW Server IP Address prompt is displayed.



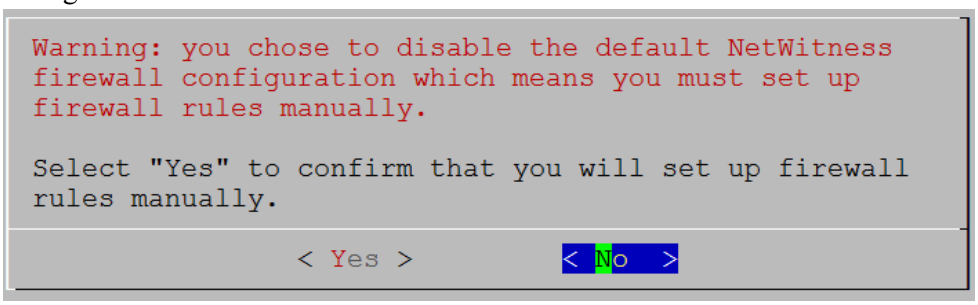
Type the NW Server IP address. Tab to **OK** and press **Enter**.

13. The Disable firewall prompt is displayed.

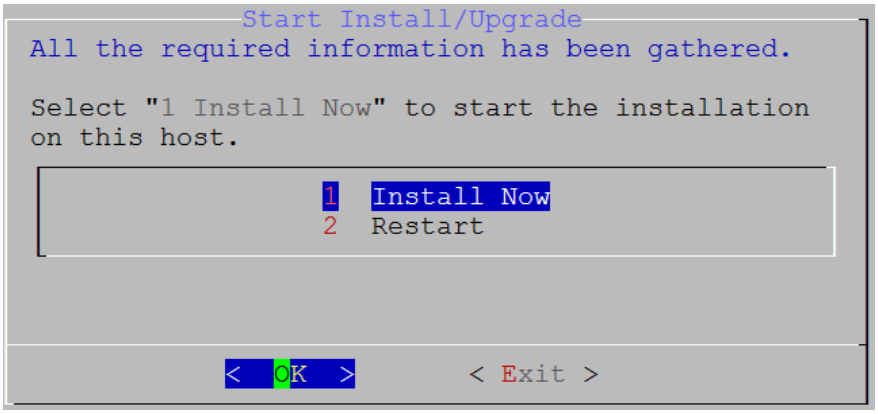


Tab to **No** (default), and press **Enter** to use the standard firewall configuration.
To disable the standard firewall configuration, tab to **Yes**, and press **Enter**.
If you select **Yes**, confirm your selection(select **Yes** again) or select **No** to use the standard firewall

configuration.

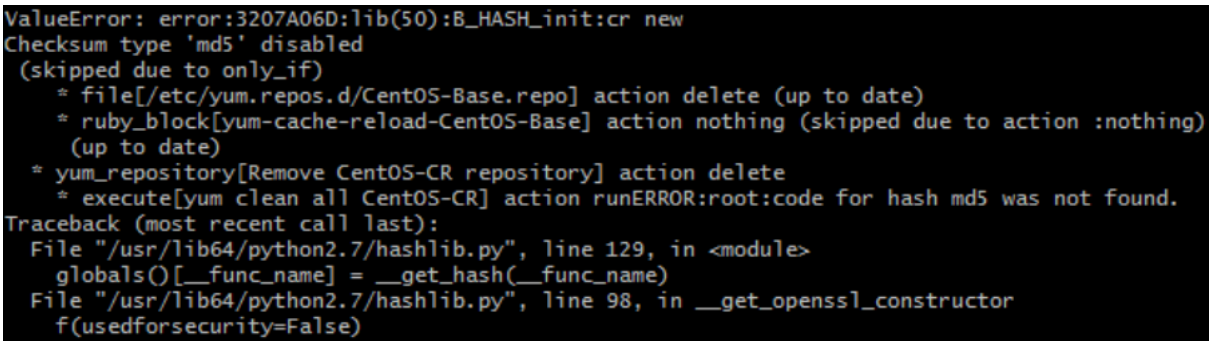


14. The **Start Install** prompt is displayed.



15. Press **Enter** to install 12.0.
When **Installation complete** is displayed, you have installed 12.0 on this host.

Note: Ignore the hash code errors similar to the errors shown in the following figure that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.

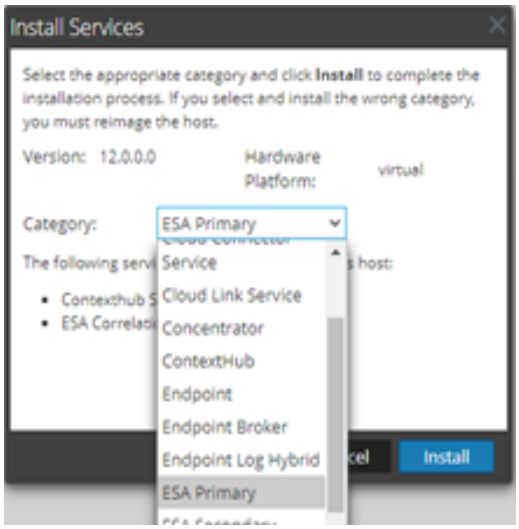


16. (Optional) If your system configuration requires that a component host must use a NAT IP address to reach the NW Server host, you must configure the NAT IP address of the NW Server by running the following command:
`nw-manage --update-host --host-id <NW Server Host UUID> --ipv4-public <NAT IP address>`

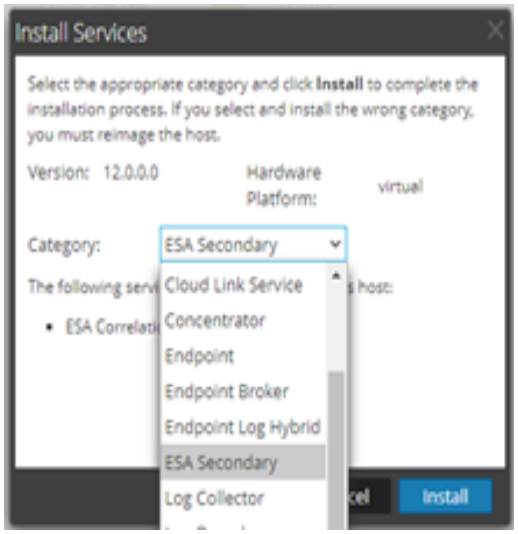
Set Up ESA Hosts

After you install your NW Server and component hosts, follow these steps to set up your ESA hosts.

- Install your primary ESA host following the instructions in "Install 12.0 on the NetWitness Server (NW Server) Host and Other Component Hosts" in this guide, and install the **ESA Primary** service on it after you finish the Set Up program in the UI in (Admin) > **Hosts** > Install:



- (Conditional) If you have a secondary ESA host, install it and install the **ESA Secondary** service on it after you finish the Set Up program in the UI in (Admin) > **Hosts** > Install:



Install Component Services on Hosts

After you have installed NW Server and component hosts, and set up your ESA hosts, follow these steps to install component services, such as Decoders and Concentrators, on your host systems.

1. Install a component service on the host.

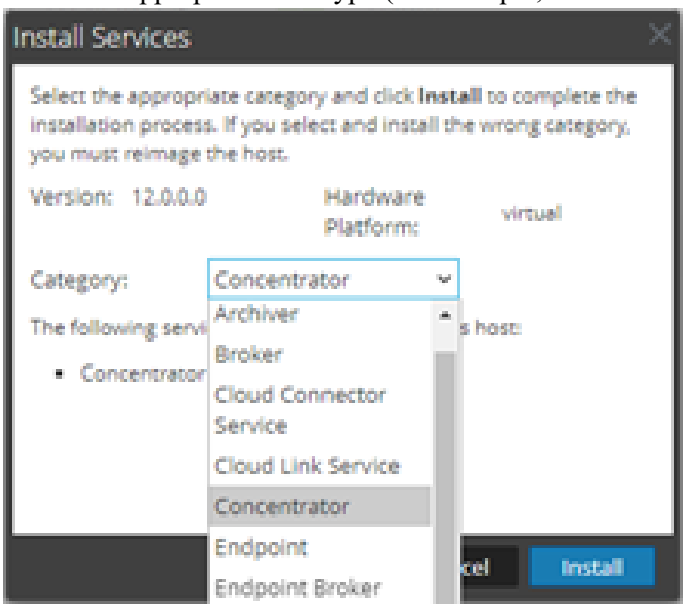
   a. Log into NetWitness and go to (Admin) > **Hosts**.
   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background.

   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the host in the **New Hosts** dialog and click **Enable**.
   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the **Hosts** view and click Install.
   The **Install Services** dialog is displayed.

   d. Select the appropriate host type (for example, **Concentrator**) in **Category** and click **Install**.



Complete Licensing Requirements

Complete licensing requirements for installed services. See the *NetWitness Platform 12.0 Licensing Management Guide* for more information. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

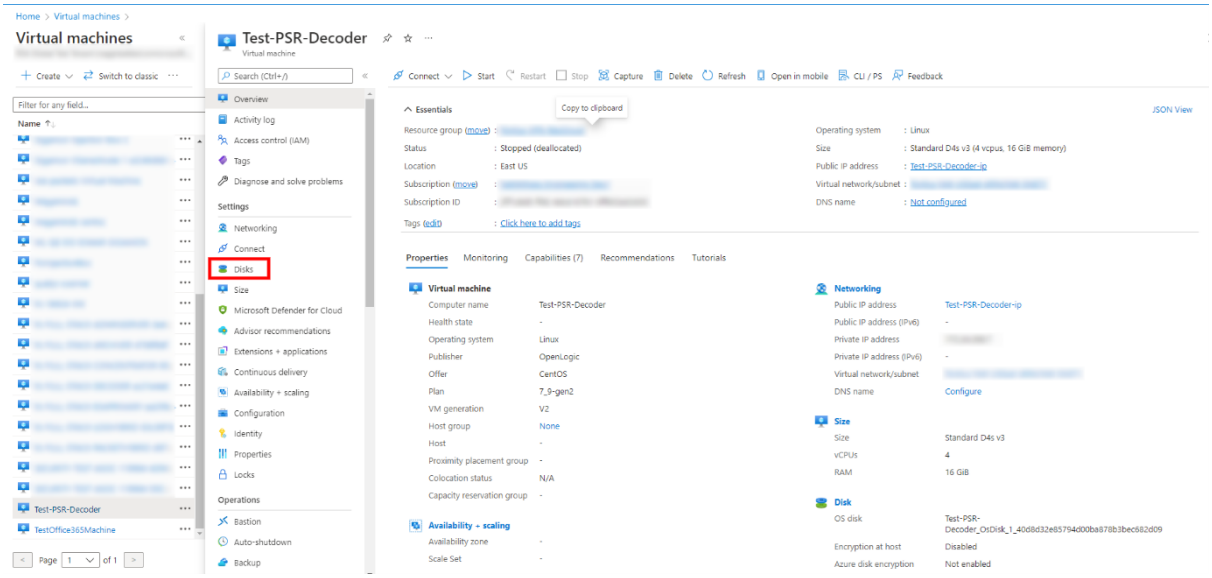
(Optional) Install Warm Standby NW Server

Refer to "Warm Standby NW Server Host" under "Deployment Option Setup Procedures" in the *Deployment Guide for NetWitness Platform 12.0* for instructions on how to set up a Warm Standby NW Server.

NetWitness Azure Storage Allocation Procedure

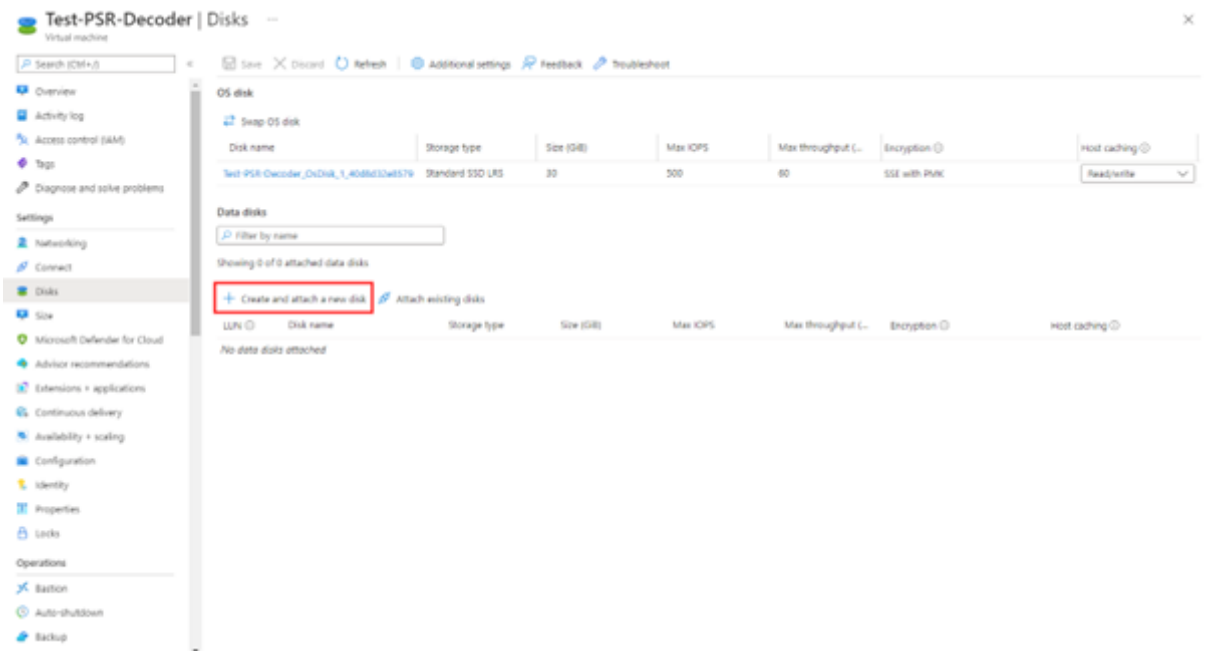
To allocate storage in NetWitness Platform 12.0.0.0, perform the following steps:

1. In Microsoft Azure portal (https://portal.azure.com/), go to **Virtual Machines**.
2. Click on the required VM > **Disks**.



3. Click **Create and attach a new disk**.

> **Note:** You need to add the appropriate amount of disks / IOPS to meet the retention requirements. If you need to add more than a single disk, a RAID configuration is needed.



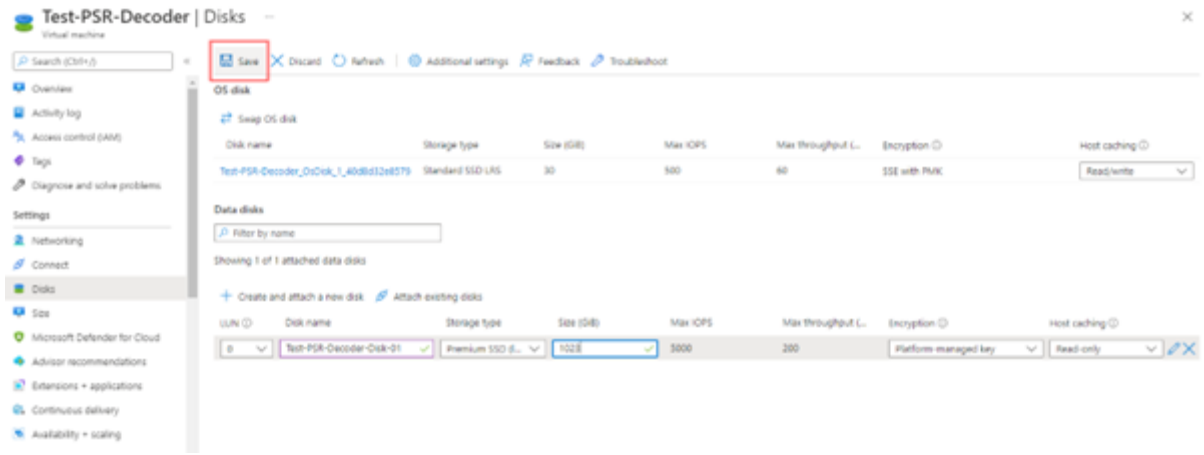
4. In the Disks view,
   a. Enter the Disk name.
   b. Select the storage type of disk.

Note: Premium SSD with high throughput / IOPS is recommended for Concentrator-IndexDB, Decoder-PacketDB.

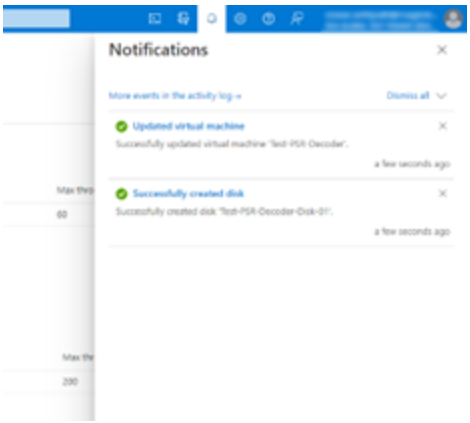
c. Select appropriate disk tier and size based on IOPS / Required retention. For more information, refer Azure Managed Disks documentation (https://docs.microsoft.com/en-us/azure/virtual-machines/disks-types).

d. Enable **Read/write** caching.

5. Click **Save** to finish adding the disk.



6. Once the disk is saved, the success notification messages are displayed in the **Notifications** view.



Configure Hosts (Instances) in NetWitness Platform

Configure individual hosts and services as described in NetWitness *Host and Services Getting Started Guide*. This guide also describes the procedures for applying updates and preparing for version upgrades.

Note: After you successfully launch an instance, Azure assigns a default hostname to it. For more information, see "Change Host Network Configuration" in the *System Maintenance Guide* for instructions on changing a hostname. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Appendix A. Silent Installation Using CLI

You can use the following Command Line Interface commands to run the installation script (`nwsetup-tui`) without getting prompted for inputs. This enables you to automate the installation of a host by supplying response to the scripts prompts through the command line.

1. After you have created a base image on the host, log in to the host with the `root` credentials.
2. Submit the `nwsetup-tui` script with the `--silent` command and the arguments that you want to apply.

   The following command string is an example of how you would install a basic NW Server host.

   ```
   nwsetup-tui --silent --is-head=true --host-name=new-host --master-
   pass=netwitness --deploy-pass=netwitness --repo-type=1 --custom-
   firewall=false --ip-override=false --eula=true
   ```

Note: In NetWitness Platform version 11.6 or later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`.
If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

3. (Conditional - For Component Hosts Only) Install the appropriate service **Category** on the newly provisioned host in the NetWitness Platform Hosts view.

   a. Log into NetWitness and go to  **(Admin)** > **Hosts**.

   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background

   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the host in the **New Hosts** dialog and click **Enable**.

   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the **Hosts** view (for example, **Event Stream Analysis**) and click 

.

   The **Install Services** dialog is displayed.

   d. Select the appropriate host type in **Category** and click **Install**.

Arguments

Argument	Description
`--help-install-opts`	Display all the arguments in this table.

Argument	Description
`--eula`	Accept or decline the End User License Agreement (EULA). Specify: • `true` (default) to accept the agreement • `false` to decline it and cancel the installation. For example: `--eula=true`
`--is-head`	Designate the host as the NW Server host or a component host. Specify: • `true` for NW Server host. • `false` for Component host. For example: `--is-head=true`
`--host-name`	Specify new hostname. If you do not specify this argument, NetWitness Platform retains the existing hostname. For example: `--host-name=<hostname>`
`--master-pass`	Enter master password. For example: `--master-pass=<password>`
`--deploy-pass`	Enter deployment password. For example: `--deploy-pass=<password>`
`--iface-name`	Specify network interface. For example: `--iface-name=eth0`
`--ip-override`	Accept or override IP address found for this host or change the IP configuration found on the host. Specify: • `true` provide IP address. • `false` use IP address found on the host. For example: `--ip-override=false`
`--ip-type`	Select ip address configuration type. Specify: • `1` Static IP Configuration) • `2` DCHP For example: `--ip-type=1`
`--ip-addr`	For Static IP configuration, enter IP Address for static address. For example: `--ip-addr=<ip-address>`
`--ip-netmask`	For Static IP configuration, enter Subnet Mask for static address. For example: `--ip-gateway=<subnet-mask>`

Argument	Description
`--ip-gateway`	For Static IP configuration, enter default gateway for static address. For example: `--ip-gateway=<default-gateway>`
`--ip-nameserver`	IP address assigned to DNS server. `--ip-nameserver=<ip-address>`
`--ip-nameserver-secondary`	Optional - IP address assigned to a secondary DNS server. For example: `--ip-nameserver-secondary=<ip-address>`
`--ip-domain`	For Static IP configuration, enter Local Domain Name for static address. For example: `--ip-domain=<default-gateway>`
`--repo-type`	Select type of update repository. Specify: • `1` Local repository • `2` External repository For example: `--repo-type=1`
`--repo-url`	For an external update repository, specify the url of the repository. For example: `--repo-url=<url>`
`--head-ip`	For a component host, specify IP Address of the NW Server. For example: `--head-ip=<ip-address>`
`--custom-firewall`	Disable default firewall configuration and use your custom configuration. Specify: • `true` use custom firewall configuration. • `false` use default firewall configuration. For example: `--custom-firewall=true`
`--use-nat`	Configure the host to use Network Address Translation (NAT) based IP addresses: • `true` use NAT IPs to connect to other hosts • `false` do not use NAT IPs to connect to other hosts (default) For example: `--use-nat=false`

NetWitness® Platform XDR

Version 12.0

Broker and Concentrator Configuration Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Broker and Concentrator Basics **4**
Overview of Broker and Concentrator **5**
Broker and Concentrator Configuration **6**
Basic Configuration Checklist 6
Step 1. Verify Service System Configuration 6
Step 2. Configure the Aggregation Process 9
Step 3. Configure Aggregate Services 11
Step 4. (Optional) Configuring Group Aggregation 15
RSA Group Aggregation Deployment Recommendations 15
Advantages of Using Group Aggregation 15
Configure Group Aggregation 17
Prerequisites 17
Set up Group Aggregation 18
Step 5. Start and Stop Aggregation 21
Broker and Concentrator Configuration References **24**
Services Config View - Broker or Concentrator General Tab 25
What do you want to do? 25
Related Topics 25
General tab 25
Aggregate Services Section 26
System Configuration Section 28
Aggregation Configuration Section 29
Service Heartbeat 31
Services System View - Broker or Concentrator 32
What do you want to do? 32
Related Topics 32
Services System View 32

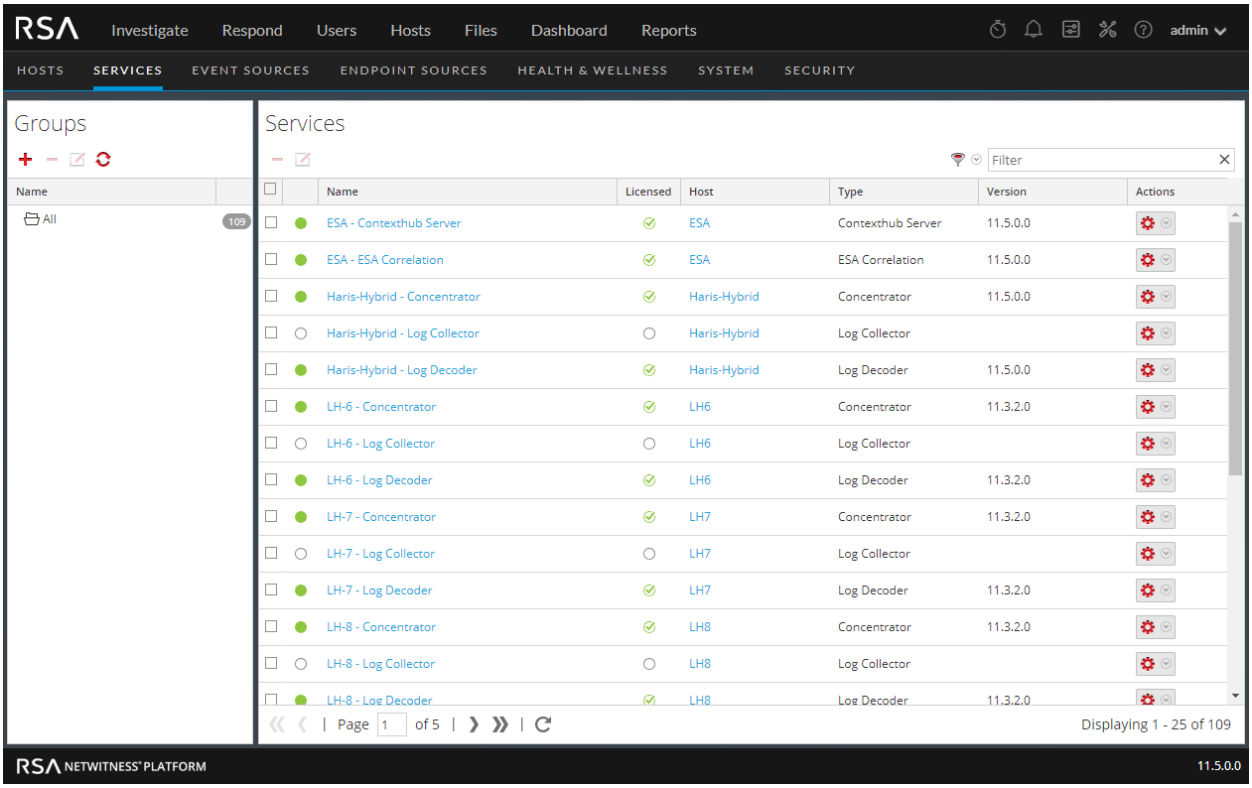
Broker and Concentrator Basics

Concentrators and Brokers aggregate data captured or aggregated by other services unlike Decoders, which capture data.

NetWitness supports the following Broker and Concentrator services:

- Brokers - aggregate data across entire infrastructure from configured Concentrators. You can have multiple concentrators aggregating into one broker. You can also have multiple brokers aggregating into a single broker.
- Concentrators - aggregates and analyzes data across multiple capture locations from decoders, indexes and directs queries.

You can configure various Brokers and Concentrators together under a Broker. Brokers are able to pull in data quickly from the Concentrators because they acquire index information only. This configuration is done using the NetWitness user interface. Most of the configuration is performed in the Administration Services view (**(Admin) > Services**).



You can also configure the aggregate services and perform the whole aggregation process using the Services view. This helps setup aggregation autostart, timing and performance parameters, maximum number of open meta and session files. In addition to this, you can also time the attempts to restart, reconnect, or take a non-responsive aggregate service offline. Configuring Aggregate services includes managing Concentrators and Decoders as aggregate services. You can also limit the data being consumed from an aggregate service using meta fields and filters. The aggregation tasks are performed in the General tab of Administration Services view (**(Admin) > Services**).

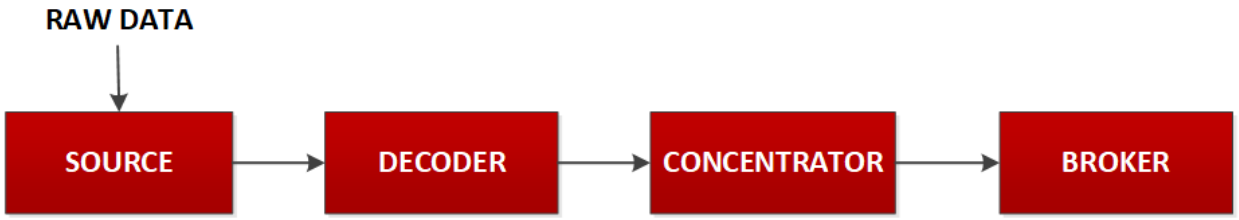
Overview of Broker and Concentrator

Brokers and Concentrators work in conjunction with Decoders and Log Decoders in the NetWitness Platform network. Unlike the two types of Decoders, which capture packets and logs, Concentrators and Brokers aggregate the data captured or aggregated by other services. Brokers aggregate data from configured Concentrators; Concentrators aggregate data from Decoders. A complete overview of the NetWitness Platform is provided in the *NetWitness Platform Getting Started Guide*.

Note: Go to the Master Table of Contents in NetWitness community portal to find and view referenced documents.

As raw data is entered in the system from the source for analysis, it has to be collected and parsed. This raw data is collected, parsed, and stored using a Decoder. The packet data is then indexed, stored, and parsed by the Concentrator. Parsed packet data is also provided as an endpoint for queries. Eventually, the Broker routes queries across multiple Decoder and Concentrator appliances. Here is how information flows to a Concentrator and Broker.

In most cases, the default values for compression, statistics update interval, and number of threads in the thread pool are set at a good point for optimal system performance.

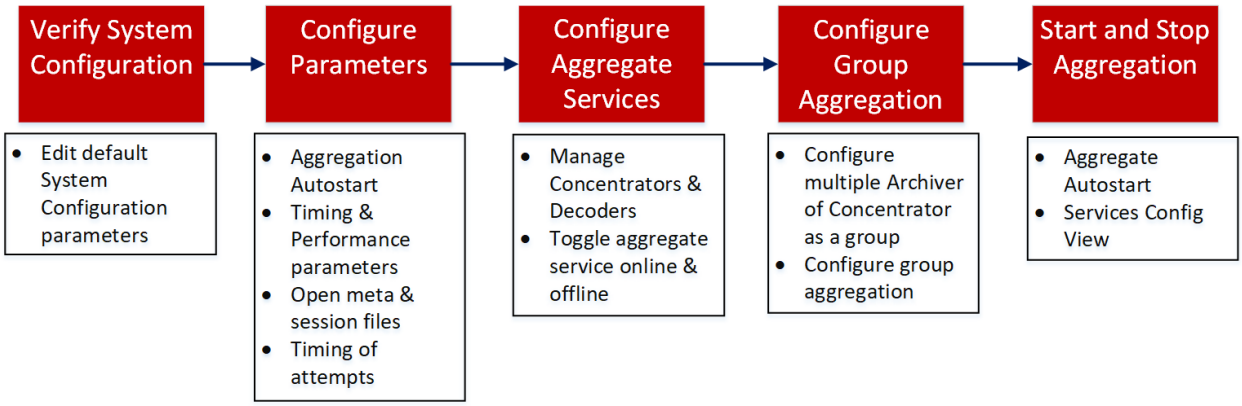


- Concentrator: is required for any large environment to store the Meta data that is generated by the parsers and feeds being triggered by packets and logs ingested into the decoders.
- Broker: The Broker service is similar to the Concentrator service except that it indexes the collected information. It performs virtual mapping of indices on all connected concentrators. Due to the less internal processing performed, the response time is fast. To allow investigation, multiple brokers and/or concentrators report data into a broker.

Broker and Concentrator Configuration

Setting up a Broker or Concentrator involves configuring the basic system parameters, the aggregate services, and the aggregation process between a Broker or Concentrator and the aggregate services.

These are the required configuration steps for a new Broker or Concentrator, and also for changing the configuration of an existing Broker. Perform the steps in the section in the sequence they are given.



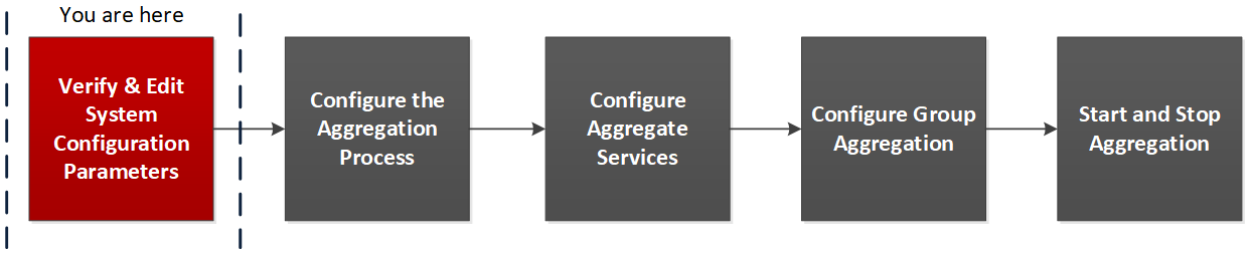
Basic Configuration Checklist

The following checklist provides the sequence for tasks that are required to configure a Broker or Concentrator that has been added to NetWitness in accordance with the *Hosts and Services Guide*.

Configuration Step	Description
Step 1 - Verify System Configuration	Verify system configuration default values for the host and service are appropriate as described in Step 1. Verify Service System Configuration
Step 2 - Configure Parameters	Configure parameters that govern the overall aggregation process as described in Step 2. Configure the Aggregation Process
Step 3 - Configure Aggregate Services	Configure aggregate services as described in Step 3. Configure Aggregate Services
Step 4 - Configure Group Aggregation	(Optional) Configure group aggregation as described in Step 4. (Optional) Configuring Group Aggregation
Step 5 - Start and Stop Aggregation	Start and stop aggregation as described in Step 5. Start and Stop Aggregation

Step 1. Verify Service System Configuration

When a service is first added to NetWitness, default values for the system configuration parameters are in effect. You can edit these values to tune performance.

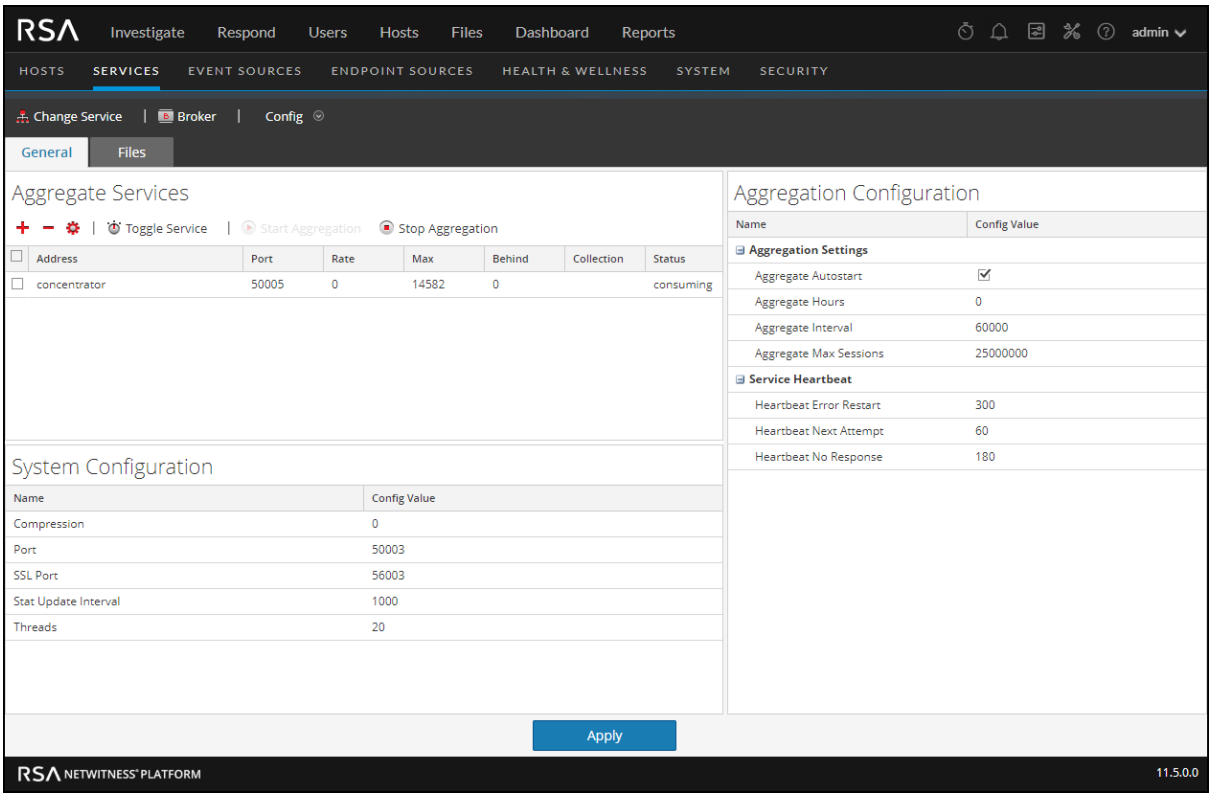


In most cases, the default values for compression, statistics update interval, and number of threads in the thread pool are set at a good point for optimal system performance.

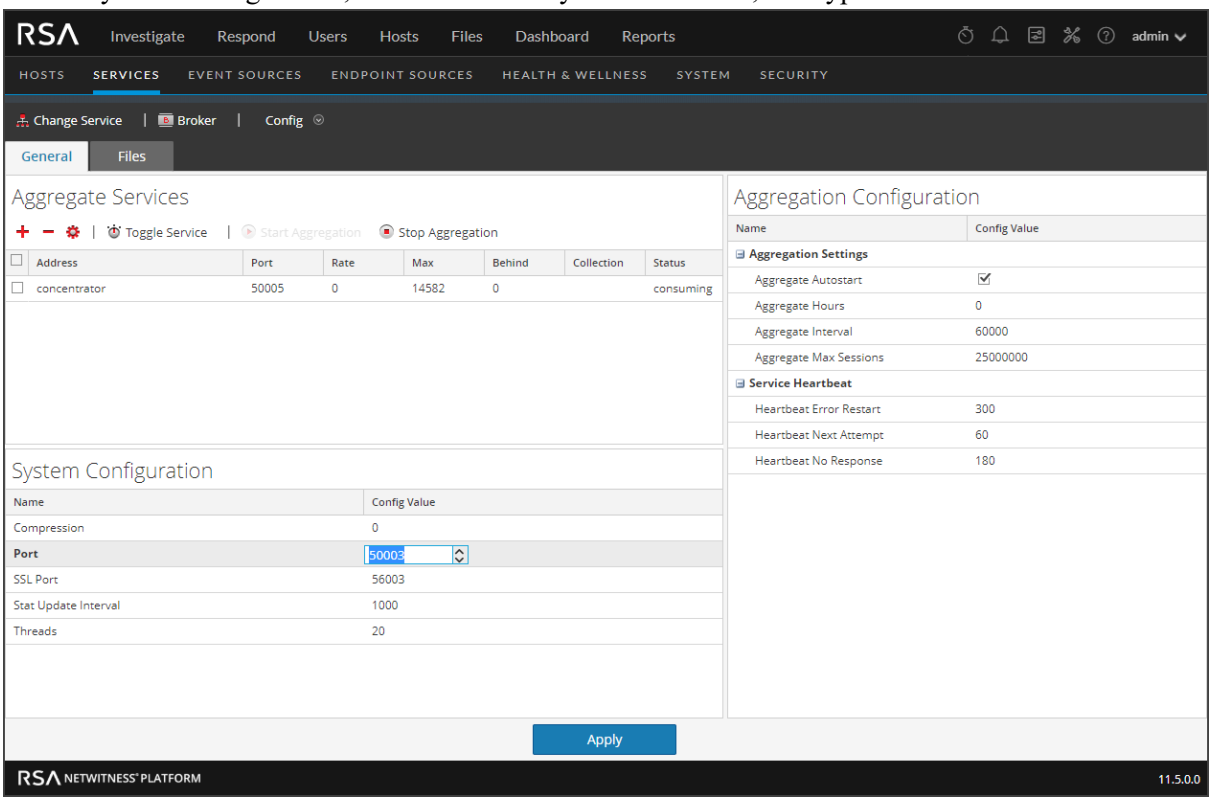
To edit system configuration parameters for a Broker or Concentrator:

1. Go to  **(Admin) > Services**.
2. In the **Services** view, select a Broker or Concentrator, and select  **> View > Config**.

   The Services Config view for the selected service is displayed.



3. Under System Configuration, click a field that you want to edit, and type a new value.

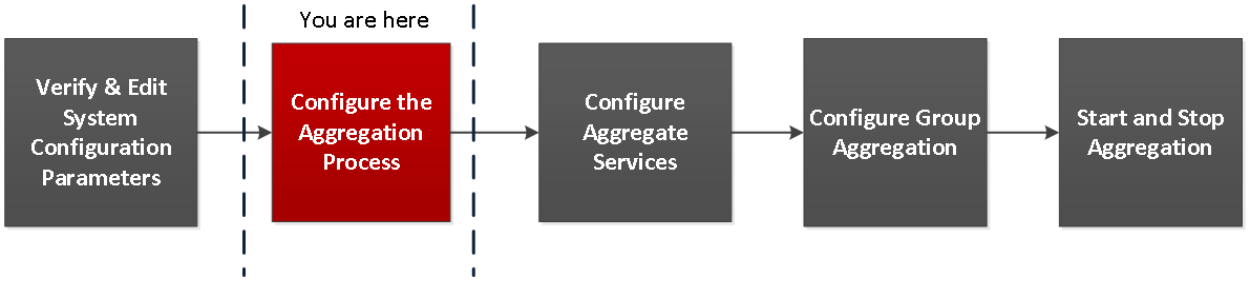


4. When finished editing, click **Apply**.

Step 2. Configure the Aggregation Process

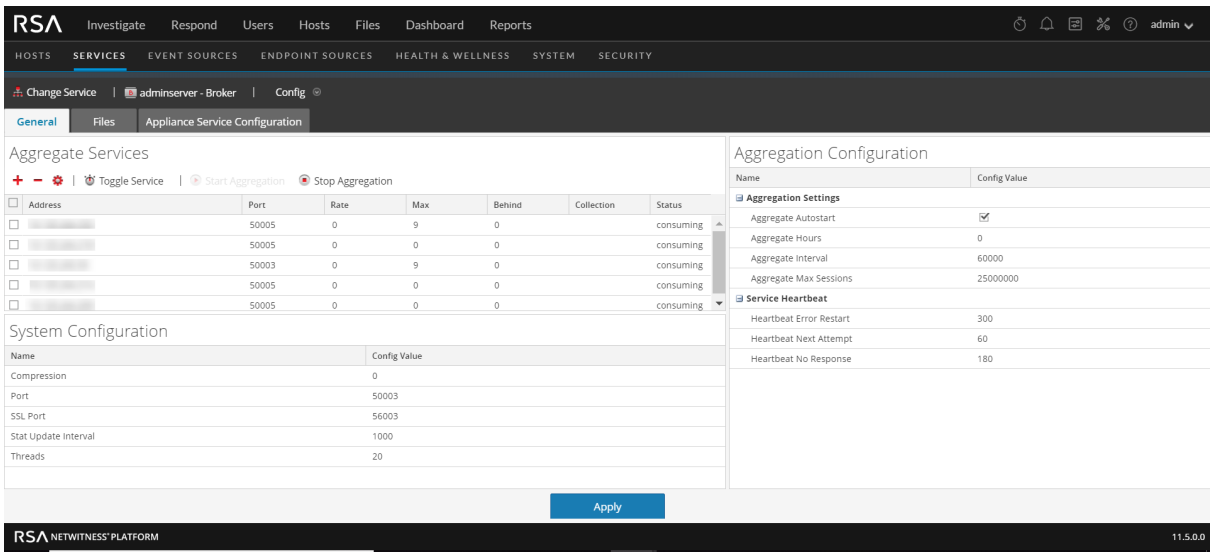
Configuring the aggregation process for a Broker or Concentrator includes setting:

- Aggregation autostart
- Timing and performance parameters, such as the number of sessions per round of aggregation and time between rounds
- Maximum number of open meta and session files
- The timing of attempts to restart, reconnect, or take offline a non-responsive aggregate service

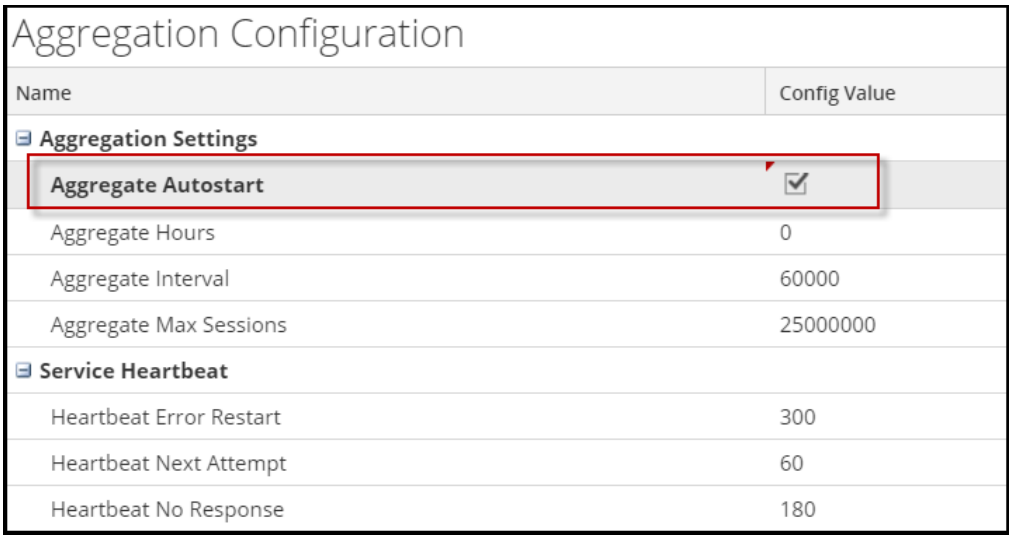


To configure the aggregation process on a Broker or Concentrator:

1. Go to **(Admin) > Services**.
2. In the **Services** view, select a Broker or Concentrator, and select **> View > Config**.
   The Services Config view, which includes the Aggregation Configuration section, is displayed.



3. (Optional) Select **Aggregate Autostart** to enable automatic start of aggregation when a service is online.



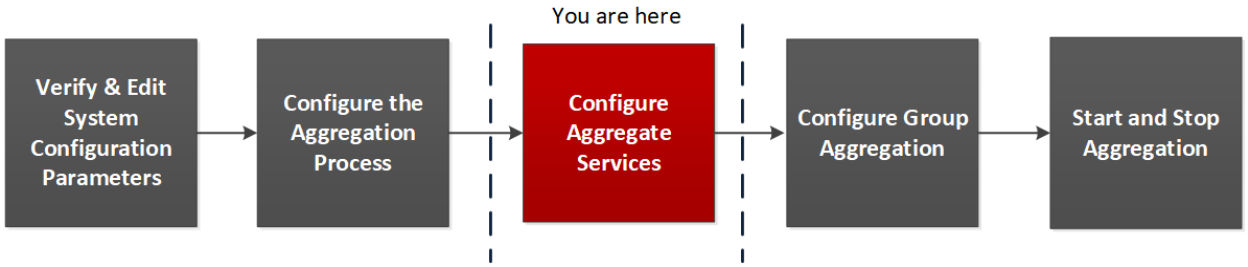
4. (Optional) Edit any of the aggregation settings: the hours back to begin aggregation, the milliseconds between rounds of aggregation, and maximum number of sessions per aggregation round.
5. (Optional) Edit any of the Service Heartbeat settings, which specify the timing of the first attempt to reconnect to a service after an error, the next attempt to reconnect, and taking the service offline after failure to reconnect.
6. When finished editing the settings, click **Apply**.
   The settings become effective immediately.

Step 3. Configure Aggregate Services

This topic introduces basic tasks related to data aggregation on Brokers and Concentrators. For information on the optional setup of group aggregation, see Step 4. (Optional) Configuring Group Aggregation.

Configuring the aggregate services (whose data is consumed and aggregated) includes:

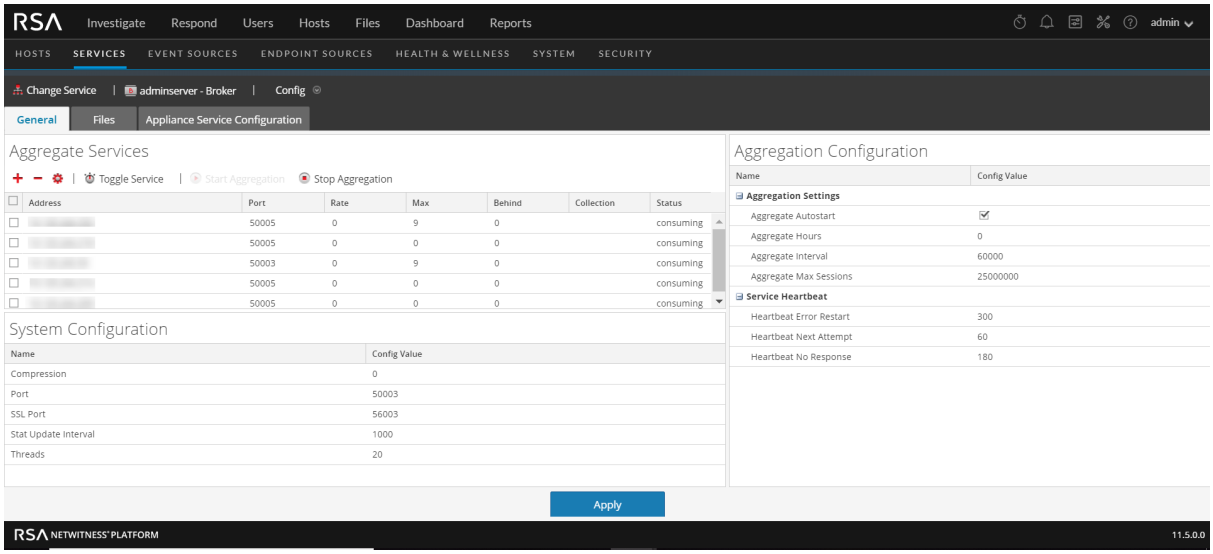
- Adding, editing, and deleting Concentrators and Decoders as aggregate services
- Toggling an aggregate service online and offline



To configure aggregate services to a Broker or Concentrator:

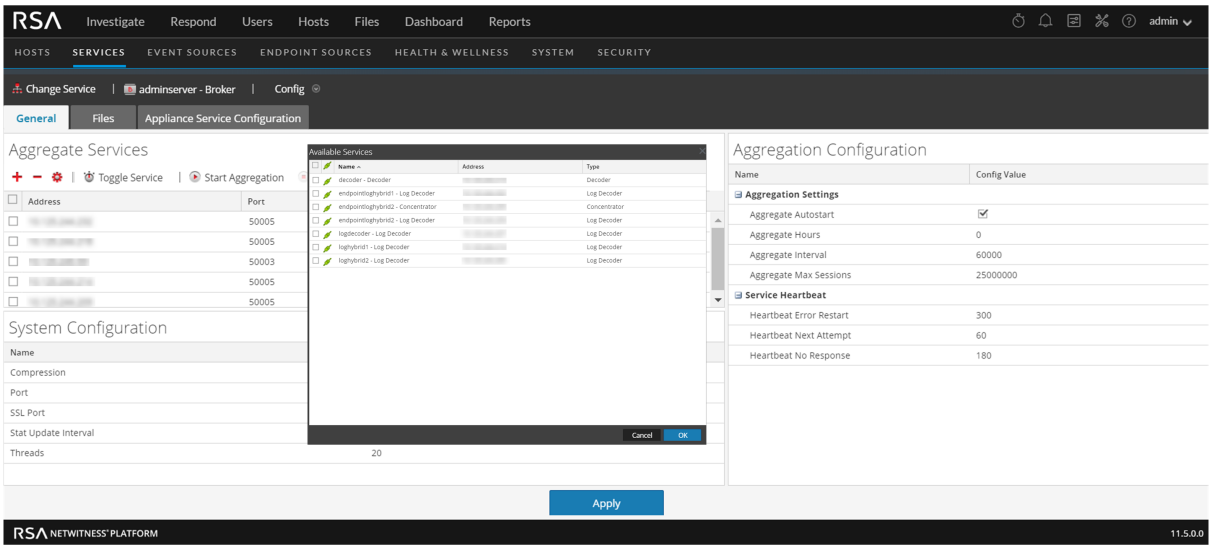
1. Go to  **(Admin) > Services**.
2. In the **Services** view, select a Broker or Concentrator, and select  **> View > Config**.

   The Services Config view for the selected service is displayed.

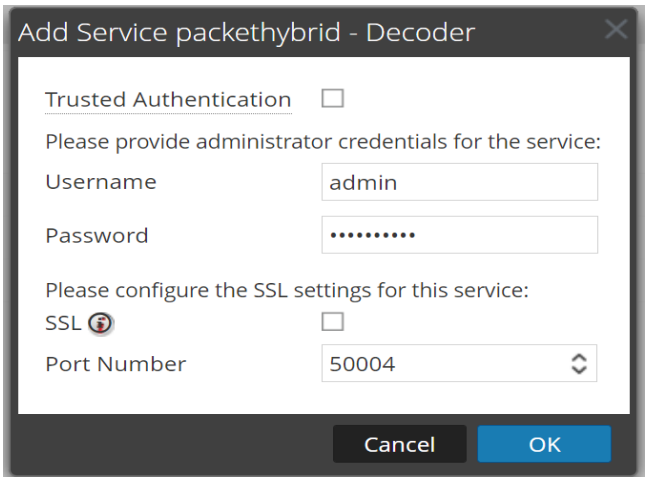


3. Click  in the **Aggregate Services** toolbar.

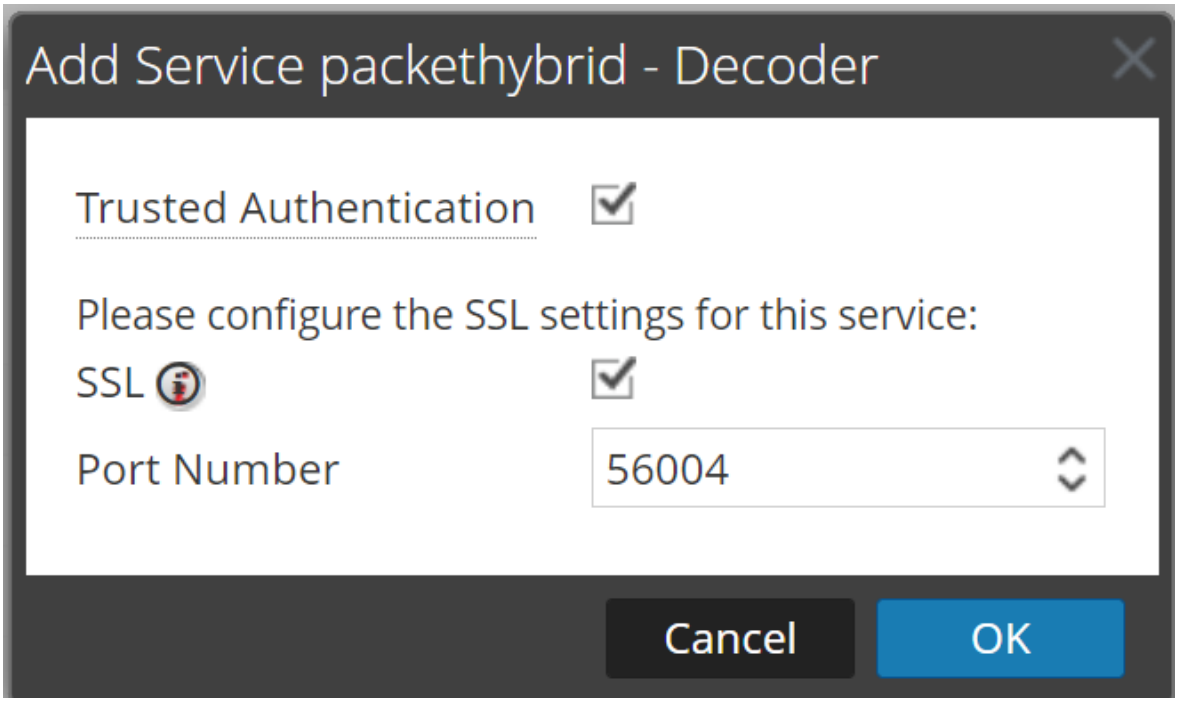
   The Available Services dialog is displayed.



4. Select one or more services to be added and click **OK**.
5. Use one of the following ways to add a service for aggregation:
   - Use Password-Based Authentication. Enter the Administrator username and password. Click **OK**.



- Select **Trusted Authentication**. Click **OK**.



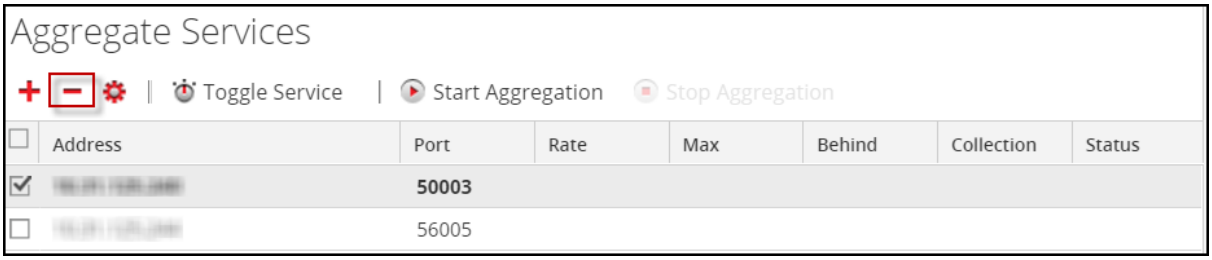
The added services are listed in the Aggregate Services list.

6. To save the changes, click **Apply**.

To remove aggregate services from a Broker or Concentrator:

Note: This option applies only to offline services. If the aggregate service is online, you must first toggle the service offline.

1. In the **Aggregate Services** list, select one or more services.
2. Click − in the toolbar.



The service is removed from Aggregate Services list.

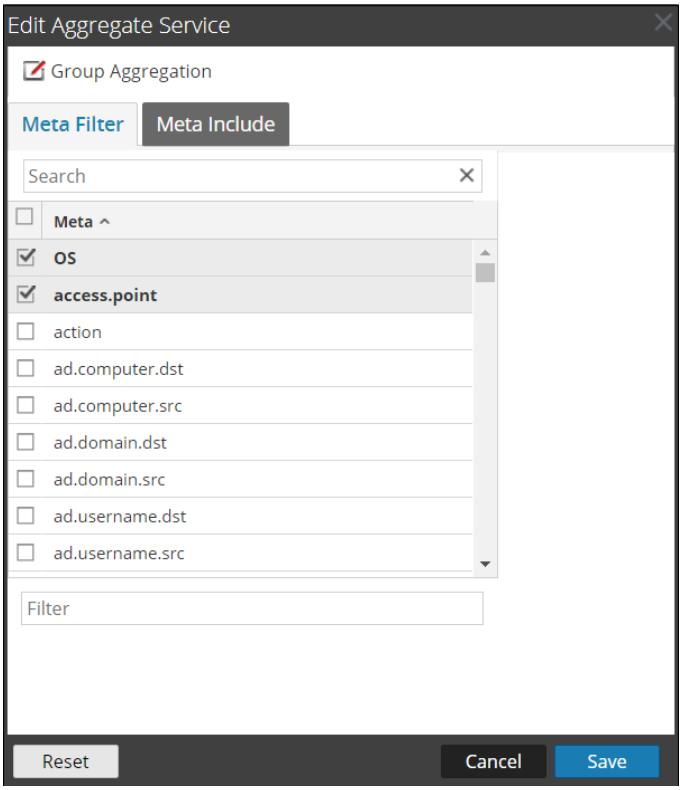
3. To save the change, click **Apply**.

To edit aggregate services on a Concentrator:

Note: This option applies only to offline services. If the aggregate service is online, you must first toggle the service offline. You can edit only one service at a time.

You can limit the data being consumed from an aggregate service using meta fields and filters.

1. Click **Change Service** to change the service to Concentrator.
2. In the **Aggregate Services** list, select one or more services.
3. Click in the toolbar. Enter the authentication information in the pop up dialog box.
   - If the service was added on a different instance of NetWitness, you must add it to this instance of NetWitness in order to edit. A warning dialog allows you to add the service. If you click **Yes**, the Add Service dialog is displayed.
   - If the service is online, a dialog notifies that the service must be offline and requests confirmation that you want to continue. If you click **Yes**, NetWitness takes the service offline and the Edit Aggregate Service dialog is displayed.
   - If the service is offline, the Edit Aggregate Service dialog is displayed with the editable properties for an aggregate service on a Concentrator.
4. Click a type of metadata in the **Meta Include** tab to select the type of metadata for the Concentrator to consume from this service. Click **Save**.



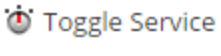
5. To specify a rule to filter data that the Concentrator consumes from this service, compose a rule in the **Meta Filter** tab. Click **Save**.
6. Click **Close**.

   The Edit Aggregate Service dialog closes and the changes are shown in the Aggregate Services list. In this example, two meta were selected on the Meta Include tab. When you click the information icon in the Meta Include field, it shows the selections.
7. To save the changes, click **Apply**.

Toggle a Service

When data aggregation starts, Brokers and Concentrators consume data from aggregate services that are online. When first added to a Broker or Concentrator, aggregate services are offline.

To toggle a service between online and offline:

1. Select a service in the **Aggregate Services** list.
2. Click 

.

   The status is changed.

Step 4. (Optional) Configuring Group Aggregation

You use Group Aggregation to configure multiple Archiver or Concentrator services as a group and share the aggregation tasks between them. You can configure multiple Archiver services or Concentrator services to efficiently aggregate from multiple Log Decoder services to improve query performance on the data:

- Stored in the Archiver.
- Processed through the Concentrator.

RSA Group Aggregation Deployment Recommendations

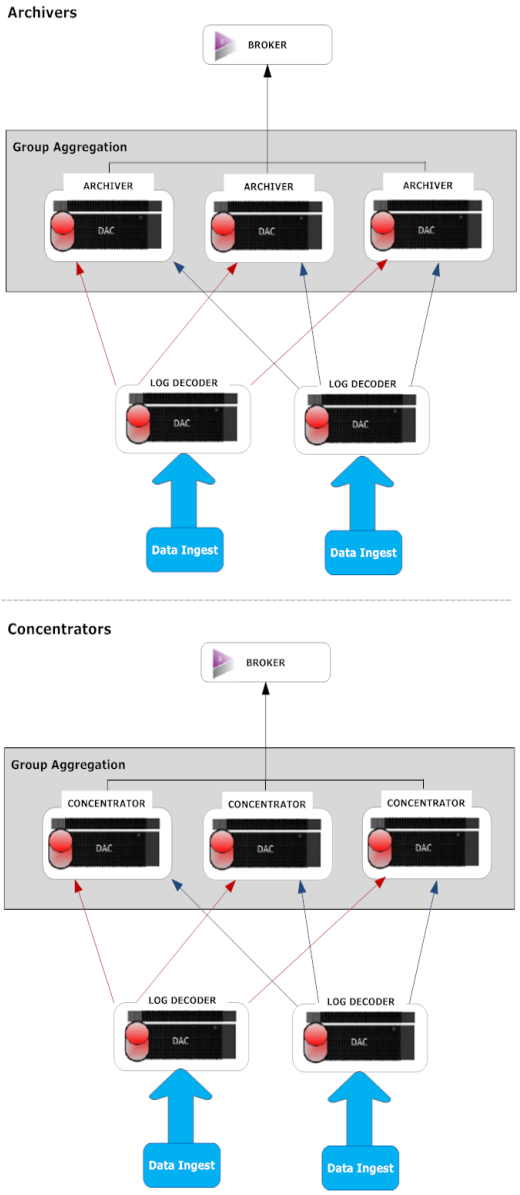
RSA recommends the following deployment for Group Aggregation:

- 1 - 2 Log Decoders
- 3 - 5 Archivers or Concentrators

Advantages of Using Group Aggregation

- Increases the speed of NetWitness queries.
- Improves the performance of aggregate queries (Count and Sum) on the environment.
- Enhances investigation service performance.
- Gives you the option of storing data for a longer duration for investigation purposes.

The following diagram illustrates Group Aggregation.



You can have any number of Archivers or Concentrators grouped together and form an aggregation group. The Archiver or Concentrator services in the group divide all the aggregated sessions between them based on the number of sessions defined in the Aggregate Max Sessions parameter.

For example, in an aggregation group containing two Archiver services or two Concentrator services with the Aggregate Max Sessions parameter set to 10,000, the services would divide the session between themselves as illustrated in the following table.

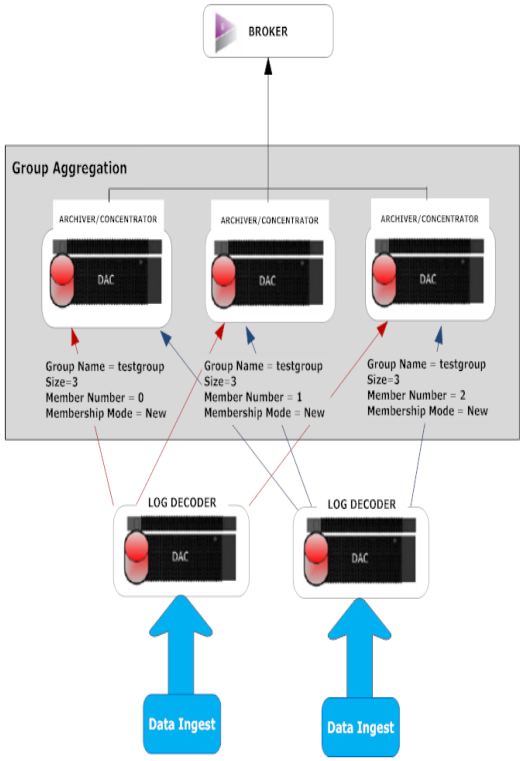
Archiver 0 or Concentrator 0	Archiver 1 or Concentrator 1
1 - 9,999	10,000 - 19,999
20,000 - 29,999	30,000 - 39,999
40,000 - 49,999	50,000 - 59,999

Configure Group Aggregation

Complete this procedure to configure multiple Archiver or Concentrator services as a group and share the aggregation tasks between them.

Prerequisites

Plan the network design for group aggregation. The following figure is an example of a group aggregation setup.



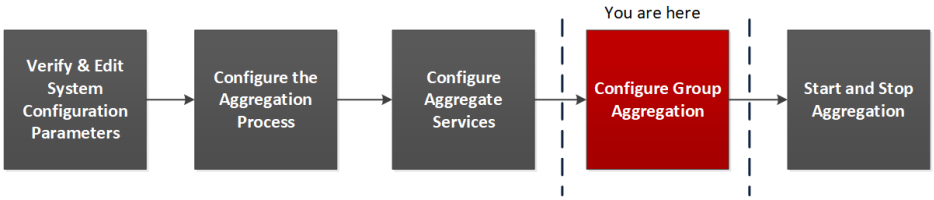
Ensure that you understand the Group aggregation parameters in the following table, and create a group aggregation plan.

Parameter	Description
Group Name	It determines the group to which the Archiver or Concentrator belongs. You can add any number of groups aggregating data from a Log Decoder. The Group Name parameter is used by the Log Decoder to identify which Archiver or Concentrator services are working together. All Archiver or Concentrator services in the group should have the same group name.
Size	It determines the number of Archiver or Concentrator services in the aggregation group.
Member Number	It determines the position of the Archiver or Concentrator in the aggregation group. For a group of size N, member number from 0 to N-1 must be set on each of the Archiver or Concentrators services in the aggregation group. For example: If the size of the aggregation group is 2, the member number of one of the Archiver or Concentrator service should be set to 0 and the member number of the other Archiver or Concentrator should be set to 1.
Membership Mode	There are two membership modes: • **New**: Adding a new Archiver or Concentrator service as a member to the existing aggregation group or creating an aggregation group. The Archiver or Concentrator service does not aggregate any existing sessions from the service as other members of the group would have already aggregated all the sessions on the service. This Archiver or Concentrator service will only aggregate new sessions as they appear on the service. • **Replace**: Replacing an existing aggregation group member. The Archiver or Concentrator will begin aggregation from the oldest session available on the service it is aggregating from.

Note: The Membership Mode parameter has an effect only when no sessions have been aggregated from the service. After some sessions are aggregated, this parameter has no effect.

Set up Group Aggregation

This workflow shows the procedures you complete to configure group aggregation.



Complete the following steps to set up group aggregation.

1. Configure multiple Archiver or Concentrator services in your environment. Make sure that you add the same Log Decoder as data source to all the services.

2. Perform the following on all the Archiver or Concentrator services that you want to be part of aggregation group:

   a. Go to **(Admin) > Services**.

   b. Select the Archiver or Concentrator service, and select > **View** > **Config**.

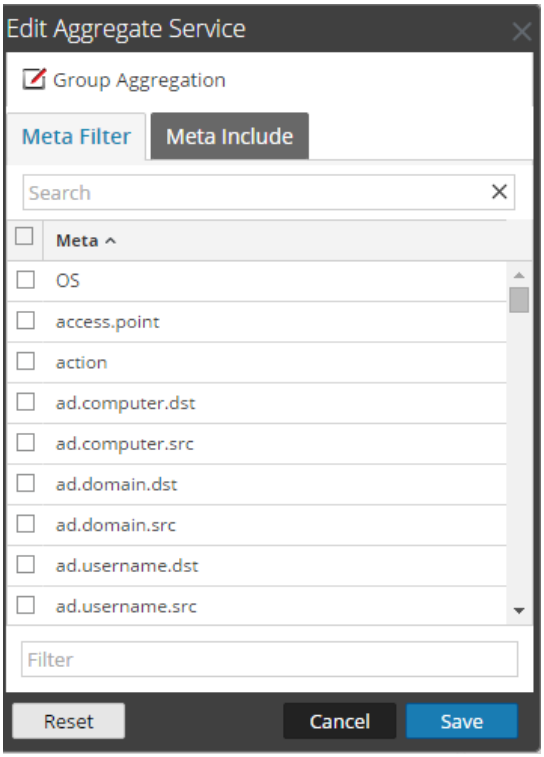
   The Service Config view of the Archiver or Concentrator is displayed.

   c. In the **Aggregate Services** section, select **Log Decoder**.

   d. Click Toggle Service to change the status of the Log Decoder to offline if it is online.

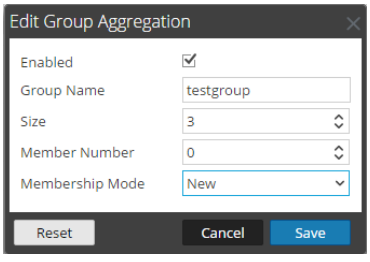
   e. Click .

   The **Edit Aggregate Service** dialog is displayed.



   f. Click Group Aggregation.

   The **Edit Group Aggregation** dialog is displayed.



   g. Select the **Enabled** checkbox and set the following parameters:

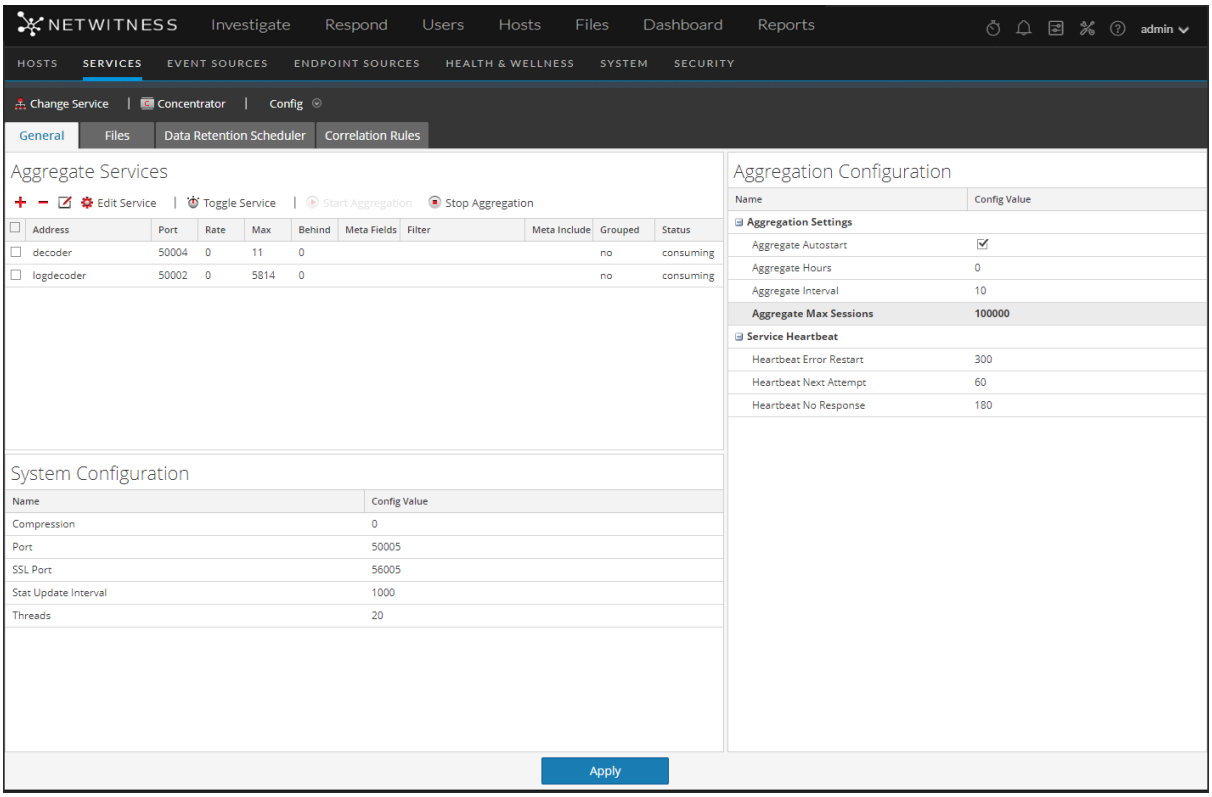
- In the **Group Name** field, type the group name.
- In the **Size** field, select the number of Archiver or Concentrator services in the aggregation group.
- In the **Member Number** field, select the position of the Archiver or Concentrator in the aggregation group.
- In the **Membership Mode** drop-down menu, select the mode.

h. Click **Save**.

i. In the Service Config view, click **Apply**.

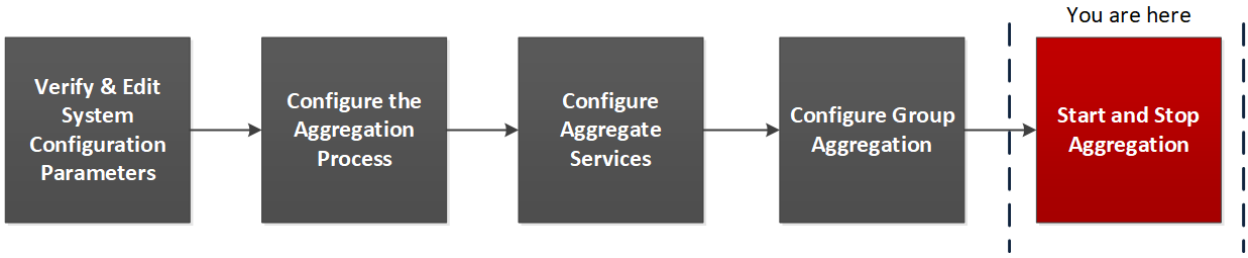
j. Perform **Step b** to **Step i** on all other Archiver or Concentrator services that need to be part of group aggregation.

3. In the **Aggregation Configuration** section, set the **Aggregate Max Sessions** parameter set to **10000**.



Step 5. Start and Stop Aggregation

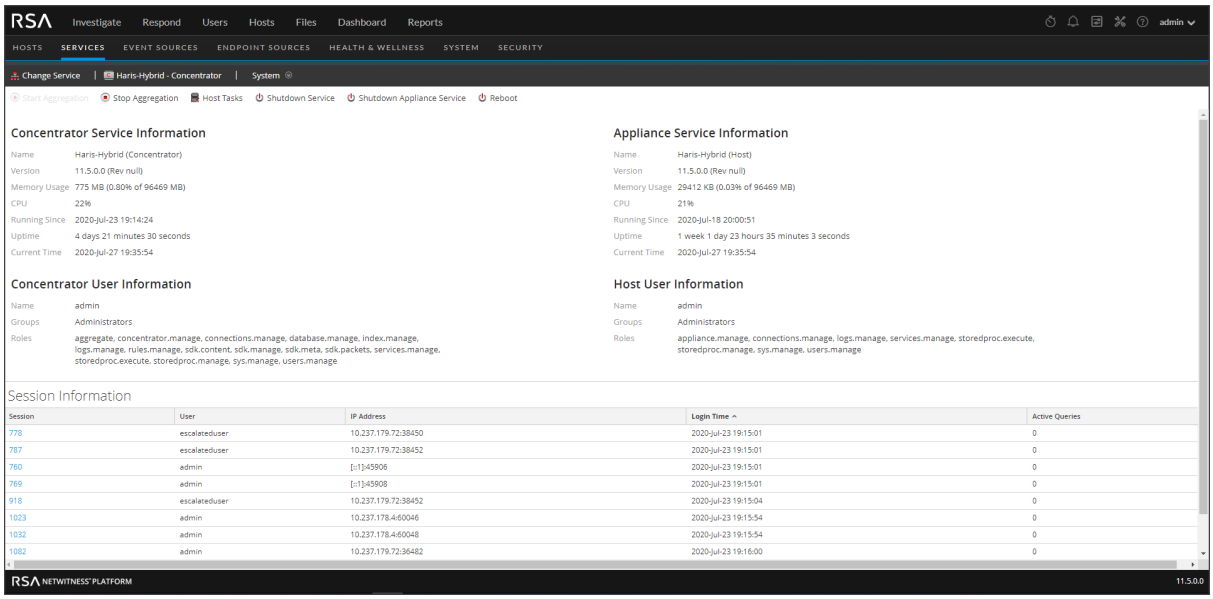
When a Broker or Concentrator starts up, it automatically begins aggregating data if Aggregate Autostart is enabled. When autostart is not enabled, you can start and stop data aggregation manually.



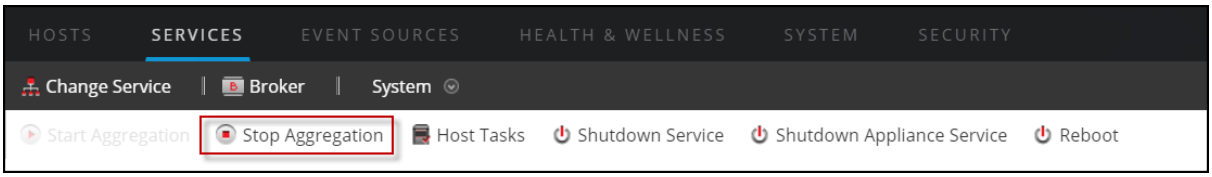
Note: The Aggregate Configuration Settings (in the Services Config View - Broker or Concentrator General Tab) determine whether Aggregate autostart is enabled, as well as the size of a round of aggregation and time between rounds.

To start and stop data aggregation in the services system view:

1. Go to **(Admin) > Services**.
2. In the **Services** view, select a Broker or Concentrator, and select > **View** > **System.**

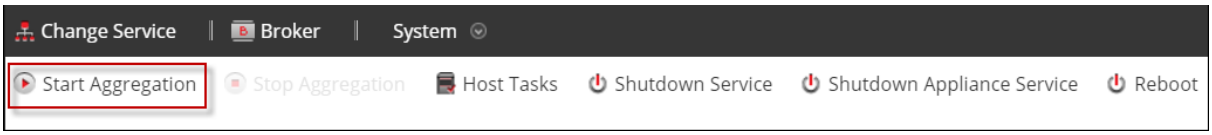


3. To stop a Broker or Concentrator that is capturing data, click **Stop Aggregation** in the toolbar. The service stops aggregating data and the **Stop Aggregation** option in the toolbar is unavailable. The **Start Aggregation** option becomes active.



4. If you want the service to start aggregating data again, click **Start Aggregation**.

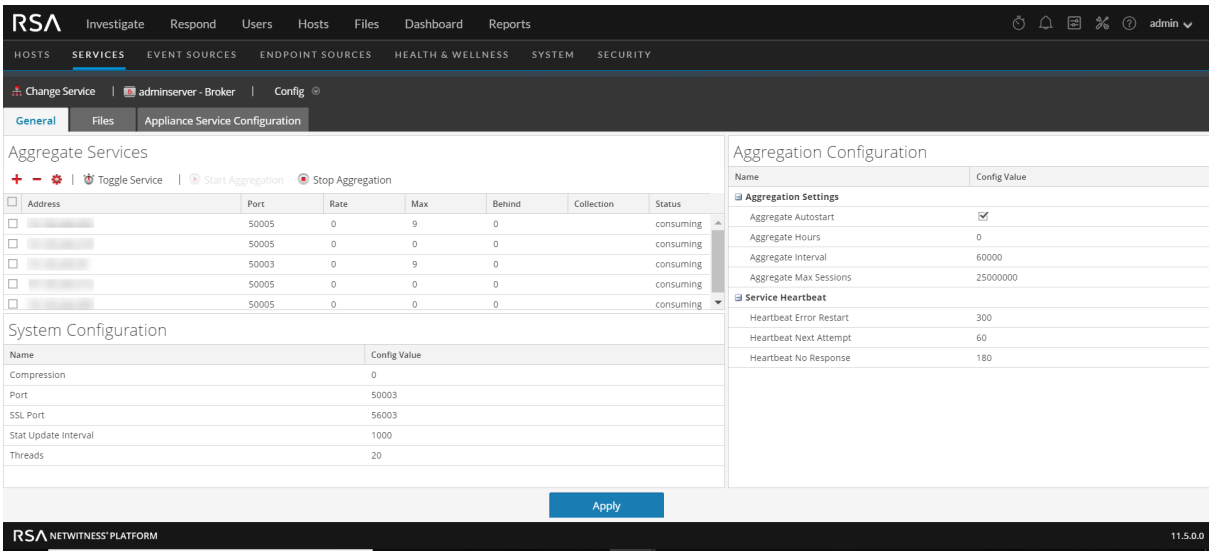
   You can now investigate the captured data in the Investigation module.



To start and stop aggregation in the services config view:

1. Go to **(Admin) > Services**.
2. In the **Services** view, select a Broker or Concentrator, and select > **View** > **Config**.

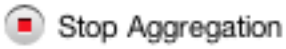
   The Services Config view, which includes the Aggregate Services section, is displayed.



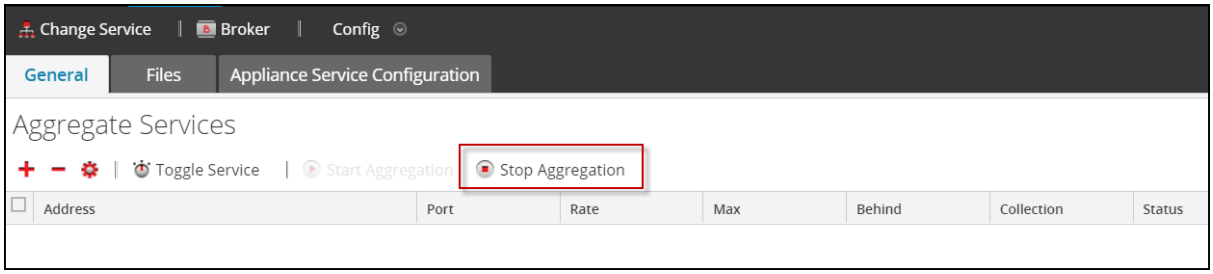
3. To start aggregation on the selected Broker or Concentrator, click Start Aggregation in the **Aggregate Services** toolbar.

   When aggregation starts, the status of all online aggregate services changes to **consuming**. The Start Aggregation button is disabled and the Stop Aggregation button is enabled.



4. To stop aggregation, click  in the **Aggregate Services** toolbar.

When aggregation stops, the status of all consuming aggregate services changes to **online**. The Stop Aggregation button is unavailable and the Start Aggregation button is available.



Broker and Concentrator Configuration References

You can configure Brokers and Concentrators using the NetWitness user interface.

In addition to the views described here, you can view the complete service nodes in a tree form in the Services Explore view, see the "Services Explore View" topic in the *Hosts and Services Getting Started Guide*.

Related Topics

- Services Config View - Broker or Concentrator General Tab
- Services System View - Broker or Concentrator

Services Config View - Broker or Concentrator General Tab

The General tab for a Broker or Concentrator in the Services Config helps manage basic service configuration, configure the aggregate service, and configure the aggregation process between a Broker or Concentrator and the aggregate service.

Configuring the aggregate service (whose data is consumed and aggregated) includes:

- Adding, editing, and deleting Concentrators and Brokers as aggregate services
- Toggling an aggregate service online and offline
- Monitoring statistics for aggregate services
- Starting and stopping aggregation

Configuring the aggregation process includes setting:

- Aggregation autostart
- Timing and performance parameters, such as the number of sessions per round of aggregation and time between rounds
- The timing of attempts to restart, reconnect, or take offline a non-responsive aggregate service

What do you want to do?

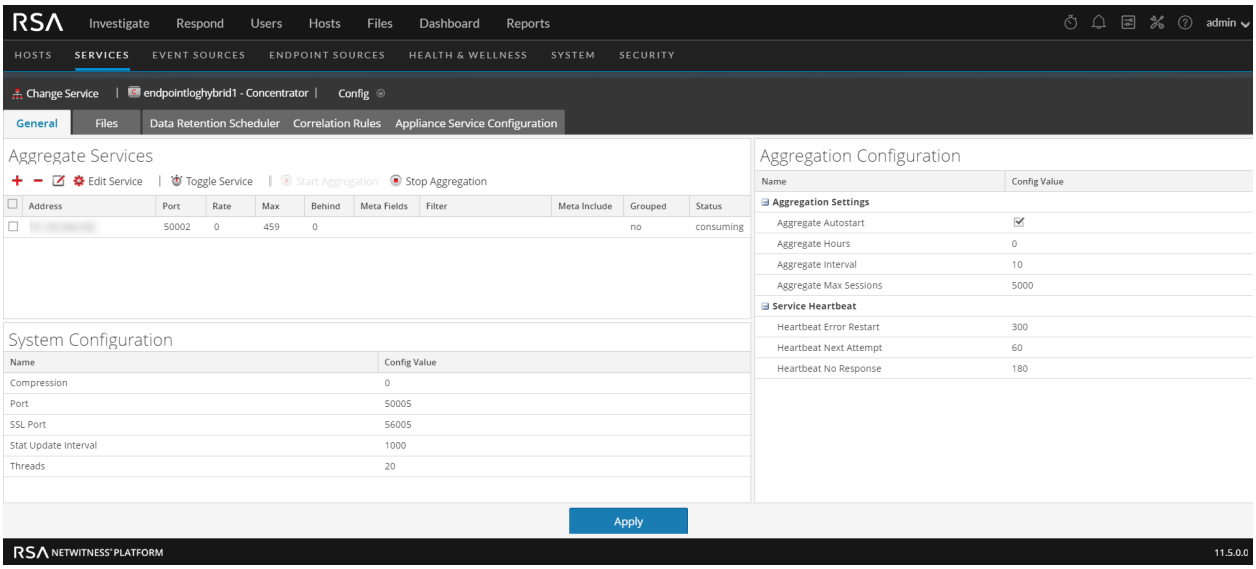
Role	I want to...	Refer to...
Administrator	Start and Stop aggregation Add, edit, delete, and toggle an aggregate service	Aggregate Services Section
Administrator	Manage System Configuration	System Configuration Section

Related Topics

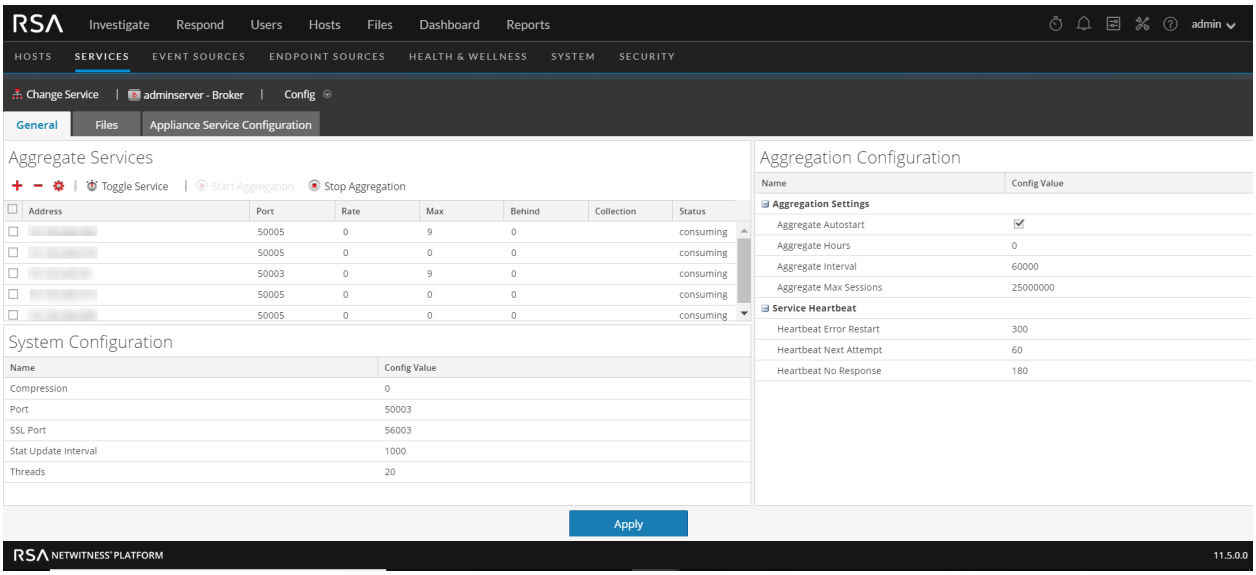
- Broker and Concentrator Basics
- Broker and Concentrator Configuration

General tab

This is an example of the General tab for a Concentrator.



This is an example of the General tab for a Broker.

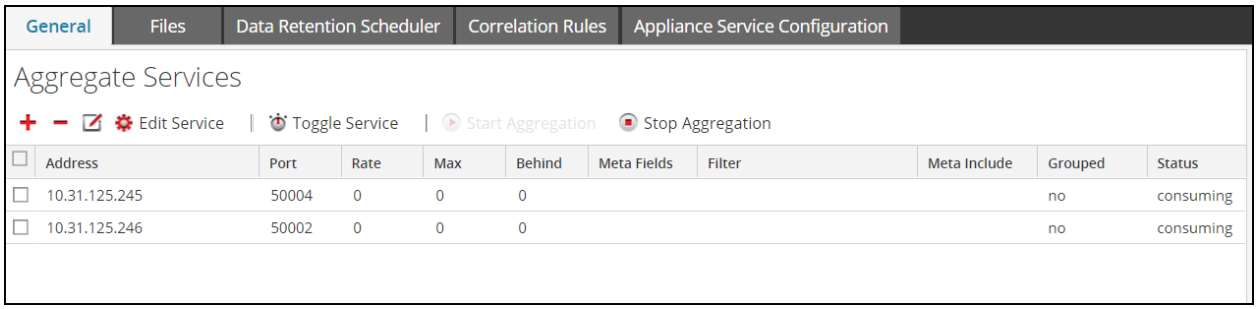


These are the three major sections in the General tab for Brokers and Concentrators:

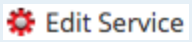
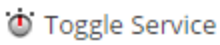
- Aggregate Services
- System Configuration
- Aggregation Configuration

Aggregate Services Section

The Aggregate Services section provides a way to start and stop aggregation, as well as add, edit, delete, and toggle an aggregate service. This is an example of the Aggregate Services section for a Concentrator.



The Aggregate Services section toolbar offers these options.

Option	Description
	Opens a dialog in which you can add a Concentrator, Decoder, or Log Decoder as an aggregate service.
	Removes the selected aggregate service.
	For Concentrators only, opens a dialog to edit **Meta Fields** and **Filter** values for the Concentrator.
	Enables you to enter the administrator credentials of the selected aggregate service so that it can communicate with the Broker or Concentrator.
	When aggregation has been stopped or has not started, starts aggregating data from the online service in the list using the rules defined for the service.
	When aggregation is in progress, stops aggregation on the Broker or Concentrator. This stops all services and flushes the index, which may take several minutes to complete. It is necessary to stop aggregate services in order to perform various administrative procedures.
	Toggles the state of a service between offline and online. Only data from online service is consumed during aggregation.

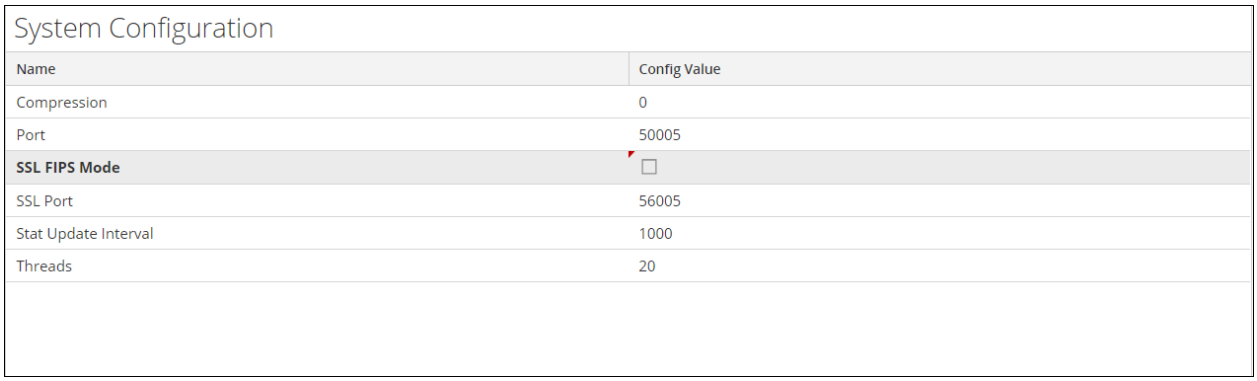
The Aggregate Services section list has these columns.

Column	Description
Address	Lists the address of the service.

Column	Description
Port	Lists the port on which the service listens. The default ports are: • 50001 for Log Collectors • 50002 for Log Decoders • 50003 for Brokers • 50004 for Decoders • 50005 for Concentrators • 50007 for other services
Rate	Lists the number of metadata objects being written to the database per second. Values are rolling average samples over a short time period (10 seconds). After capture stops, the rate is reset to **0**.
Max	Lists the maximum number of metadata objects written to the database per second since capture started. Values are rolling average samples over a short time period (10 seconds). After capture stops, **Max** continues to show the maximum value during capture.
Behind	Lists the number of sessions on the service that need to be aggregated.
Collection	For Brokers only, indicates the collection that was selected when the Analyst Workbench service was added to the Aggregate Services section.
Meta Fields	For Concentrators only, lists the types of metadata being consumed by the aggregate service.
Filter	For Concentrators only, a rule expression (as used in a 'where' clause) can be used to filter the results. You must add a meta key along with an operator and a value, for example `ip.src !=127.0.0.1 && word exists`
Meta Include	For Concentrators only, lists the number of types of meta included in the aggregate service.
Grouped	Whether or not the aggregate service is part of a group.
Status	Lists the current status of the service: • online = available to provide data for consumption by the Broker or Concentrator • offline = not available to provide data for consumption by the Broker or Concentrator • consuming = providing data for consumption by the Broker or Concentrator

System Configuration Section

The System Configuration section manages service configuration for a service. When a service is first added, default values are in effect. You can edit these values to tune performance.



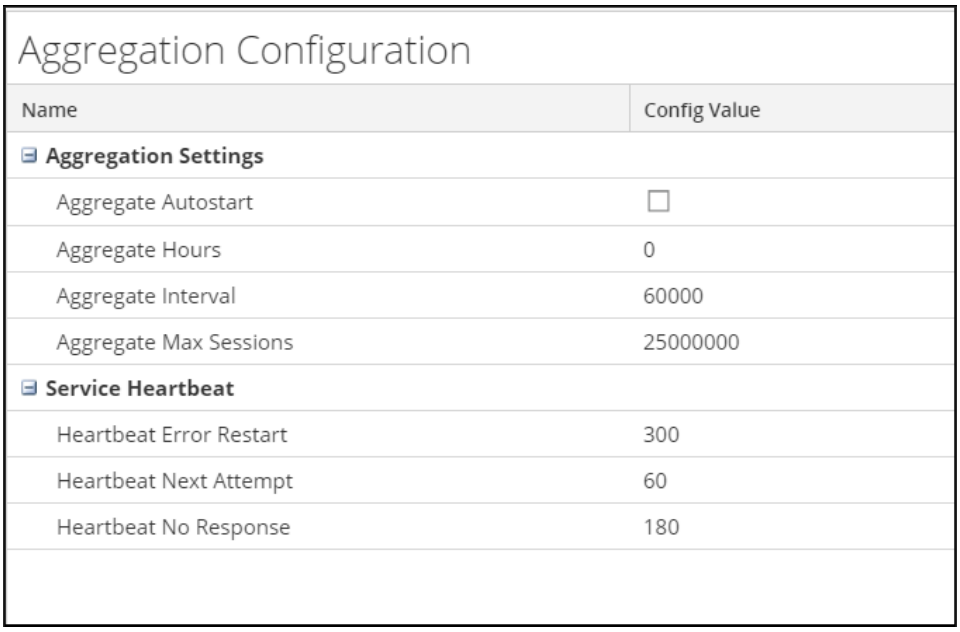
The System Configuration section has these parameters.

Parameter	Description
Compression	The minimum number of bytes that must be transmitted per response before compression. A setting of **0** disables compression. The default value is **0**. A change in value is effective immediately for all subsequent connections.
Port	The port on which the service listens. The default ports are: • 50001 for Log Collectors • 50002 for Log Decoders • 50003 for Brokers • 50004 for Decoders • 50005 for Concentrators • 50007 for other services
SSL FIPS Mode	When enabled (**on**), the security of data transmission is managed by encrypting information and providing authentication with SSL certificates. The default value is **off**.
SSL Port	Indicates the SSL port.
Stat Update Interval	The number of milliseconds between statistic updates on the system. Lower numbers cause more frequent updates and can slow down other processes. The default value is **1000**. A change in value is effective immediately.
Threads	The number of threads in the thread pool to handle incoming requests. A setting of **0** lets the system decide. The default value is **15**. A change takes effect on service restart.

Aggregation Configuration Section

The Aggregation Configuration section provides configuration settings that affect various aspects of the aggregation process. When you click **Apply**, the changes are saved; however, not all settings take effect immediately. The tables for Aggregation Settings and Service Heartbeat provide details.

Caution: Do not change any of these settings unless guided by the Developers or the Customer Support team. Contact the Customer Support for any questions before editing any of these settings.



The following table describes the aggregation settings

Setting	Description
Aggregate Autostart	Option to start aggregation automatically each time the Broker or Concentrator is started. Checked means yes, unchecked means no. This change takes effect immediately.
Aggregate Hours	The number of hours back for each service that the Concentrator or Broker attempts to recover at the beginning of aggregation. This change takes effect immediately. • If the value is set to 0, aggregation for each service starts where it last left off, no matter the number of hours behind. • If the value is any positive integer, the Concentrator or Broker only consumes sessions less than that number of hours back. For example, if a service's most current session is +10 hours from the last session, this is what happens with two different Aggregate Hours values: • With a value of 12, the Concentrator or Broker starts consuming where it left off. • With a value of 4, all sessions between 5 and 10 hours back are skipped and the Concentrator or Broker starts consuming the session that started 4 hours back.
Aggregate Interval	The number of milliseconds between rounds of service aggregation. All services managed by the Broker or Concentrator request additional rounds of session and metadata to be aggregated. If a Broker or Concentrator is still consuming the previous round of data, it cannot request more until it finishes. Change takes effect immediately.
Aggregate Max Sessions	The maximum number of sessions that the Broker or Concentrator requests in a given round of data aggregation. Change takes effect after restart.

Service Heartbeat

In communicating with each aggregate service, Brokers and Concentrators monitor the heartbeat of the service. These parameters specify the timing of the first attempt to reconnect to a service after an error, the next attempt to reconnect, and taking the service offline after failure to reconnect.

Setting	Description
Heartbeat Error Restart	After a heartbeat error is detected on an aggregate service, specifies the number of seconds for a Broker or Concentrator to wait before attempting a service reconnect.
Heartbeat Next Attempt	After a failed attempt to reconnect to an aggregate service, specifies the number of seconds for a Broker or Concentrator to wait before attempting another service reconnect. Change takes effect immediately.
Heartbeat No Response	After failing to reconnect to an unresponsive service, specifies the number of seconds for the Broker or Concentrator to wait before taking the unresponsive service offline. Change takes effect immediately.

When editing parameters in the General tab, you must click **Apply** to save changes.

Services System View - Broker or Concentrator

The Services System view displays information specific to specific to Brokers and Concentrators.

While information displayed in this view is the same for all types of Core services, several options in the toolbar are relevant only for Brokers and Concentrators.

What do you want to do?

Role	I want to...	Refer to...
Administrator	Start and Stop aggregation Add, edit, delete, and toggle an aggregate service	Services System View - Broker or Concentrator
Administrator	Manage System Configuration	Services System View - Broker or Concentrator

Related Topics

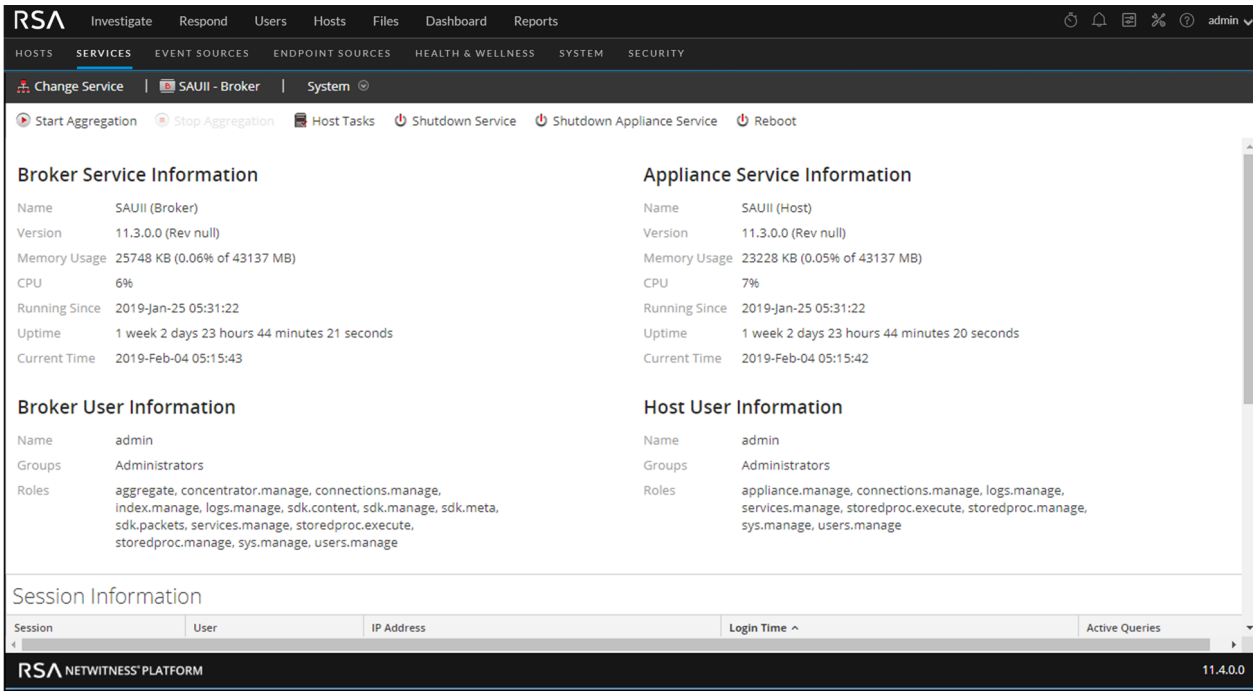
- Broker and Concentrator Basics
- Broker and Concentrator Configuration

Services System View

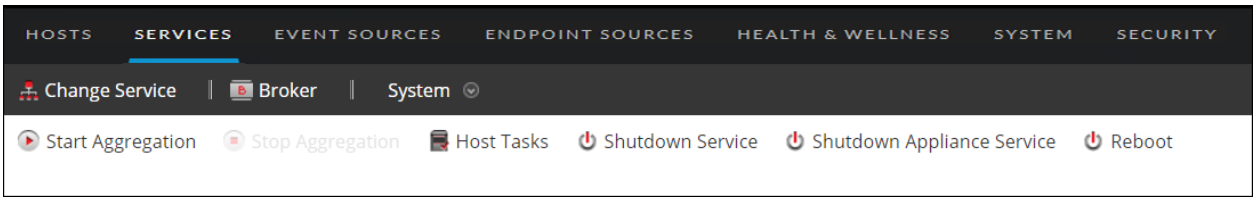
You can access this view by doing the following:

1. Go to  **(Admin)** > **Services.**
2. Select a Concentrator or Broker, and select  > **View** > **System**.

   The System view for the selected Concentrator or Broker is displayed.



The following figure is an example of the toolbar for a Broker or Concentrator.



Host Tasks, Shutdown Service, Shutdown Appliance Service or (Shutdown Appliance), and Reboot are common to all services and are described in the "Services System view" topic in the *Host and Services Getting Started Guide*.

This table describes toolbar options that apply only to a Concentrator or Broker. Both buttons are unavailable until aggregator services are configured and consuming data.

Action	Description
Start Aggregation	Starts aggregation of data being consumed on a Concentrator or Decoder configured as an aggregation service for the selected Broker or Concentrator. The Start Aggregation button is available only when aggregator services are configured and consuming data.
Stop Aggregation	Stops aggregation of data being consumed on a Concentrator or Decoder configured as an aggregation service for the selected Broker or Concentrator. The Stop Aggregation button is available only when aggregation is occurring.

NetWitness® Platform XDR

Version 12.0

Context Hub Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

How Context Hub Works **5**

Overview of Context Hub Configuration **6**

Configure Lists as a Data Source **7**

Add List data source using Local File Store 8

Add List data source using HTTP(S) 10

Next Steps 12

Configure Archer as Data Source **13**

Configure Active Directory as a Data Source **18**

Configure NetWitness Endpoint as a Data Source **22**

Configure Respond as a Data Source **26**

Configure File Reputation Server as a Data Source **28**

Enable or Disable File Reputation Data Source 28

Edit File Reputation Server Data Source Settings 30

Configure STIX as a Data Source **33**

Configure STIX File 33

Configure REST Server 34

Configure TAXII Server 36

Configure REST API as a Data Source **39**

Edit REST API Data Source 41

Delete REST API Data Source 41

Configure Context Hub Data Source Settings **42**

Import or Export Lists for Context Hub **50**

Import a List 50

Import Single-Column List 50

Import Values to an existing List 52

Export List for Context Hub 52

Configure Meta Type Mapping for Context Hub **54**

Context Hub Data Sources Tab **56**

Context Hub Lists Tab **59**

Workflow 59

What do you want to do? 59

Related Topics 60

Quick Look 60

Context Hub STIX Tab **63**

What do you want to do? 63

Related Topics 63
Quick Look 63
Troubleshooting **67**

How Context Hub Works

Context Hub service provides enrichment lookup capability in the Respond and Investigate views. An Administrator can configure the Context Hub service and the data sources to enable an Analyst to perform the context lookup for the required data sources.

By default, the Context Hub service supports enrichment lookups for meta types such as IP address, User, Domain, MAC address, File Name, File Hash, and Host.

The following data sources are supported by NetWitness and provide enriched data when configured.

Lists- Provides contextual information from a list of blacklists, whitelists, or watchlists.

Archer- Provides Criticality information of a device or specific asset based on the IP or Host which needs constant monitoring.

Active Directory - Provides contextual information of a user to help determine if the user is suspicious or not.

RSA NetWitness Endpoint - Provides context information for endpoint module and machine indicators and to help determine if any of the Endpoint devices are compromised.

Respond- Provides contextual information of a specific meta available in respond and enables analyst to respond faster based on context data.

File Reputation Server - Provides contextual information for reputation status of files.

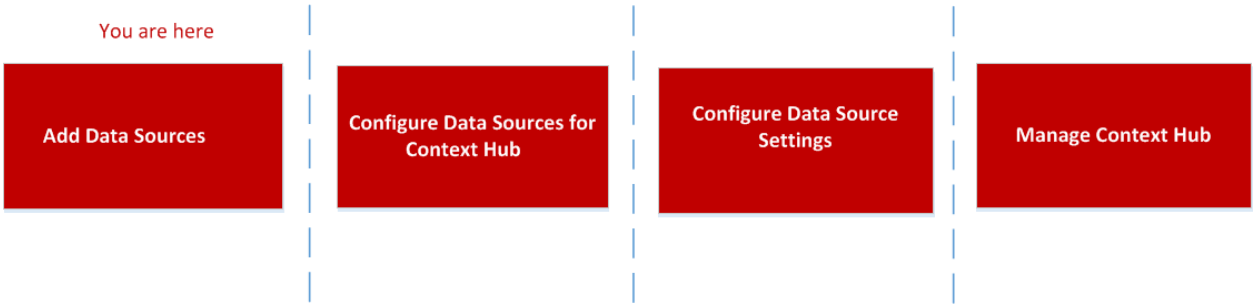
STIX - Provides contextual information on IP address, email address, domain, filename, URL, and file hash from STIX sources.

REST API - Provides contextual information for any accessible web service with an exposed REST API.

Overview of Context Hub Configuration

The Administrator needs to perform each step in the proper sequence to configure the services to perform the context lookup effectively. In the **(Admin) > Services**. Services Config view of Context Hub service, an administrator can configure data sources for Context Hub Service. The administrator can also configure Context Lookups for custom meta keys, if required and also import lists or export lists.

The workflow below describes how the Context Hub service can be configured:



Context Hub service is pre-installed on primary ESA host, and automatically added to the NetWitness.

Note: You can have only one Context Hub service instance enabled in your NetWitness deployment. If there are multiple ESA service in NetWitness, you must choose the appropriate ESA host for Context Hub. A minimum of 8GB space is required to configure Context Hub on ESA host.

Configure Lists as a Data Source

Lists as a Data Source use the Context Hub service to fetch contextual information for meta types that support context lookup. You can create one or more lists and add relevant list values to the list. Make sure that you create meaningful lists such as blacklisted IPs, whitelisted IPs, and so on. The lists can contain supported entities such as IP address, MAC address, User name, Host name, Domain name, File name or File hash. You can import a single-column list or a multi-column list from the Data Source tab. Additionally, all feeds (except STIX feeds) that are created are converted to lists and displayed on the context lookup. If Context Hub is not configured or the service is down, then the feeds will be made available whenever Context Hub is up and running. For more information on creating feeds, see the *Live Services Management Guide*.

Note: When you create a feed, a list is automatically generated with the same name as the feed. If the list name already exists, then the name of the new list is suffixed with the number '2". For example if the existing feed name is test1.csv, then the new list will be named as test2.csv.

List values are in CSV format available in an external location and can be accessed through the following two methods:

- **Local File Store**: You can share a file from a local location.
- **HTTP(S)**: You can share a file using a web server location.

Note: You can also set up recurring job to fetch data on regular intervals by using the Prefetch settings while configuring meta mapping.

Prerequisites

Before you configure Lists data source, ensure that:

- User should have admin permissions.
- Context Hub service is available in  **(Admin) > Services** view of NetWitness.
- If you are using Local File Store or HTTP(S) server, the path mentioned should contain the CSV file In case of remote Local File Store, the file must be mounted or placed on the local drive location `/var/lib/netwitness/contexthub-server/data`.
- The NetWitness user must have read permission to access the file.

Caution: If you are creating a Context Hub list for use as an enrichment source in ESA, the list name cannot include any spaces or special characters, or start with a number. If you do not follow this naming convention, when you attempt to add the list as an enrichment source in ESA, an error message will be displayed and you will not be allowed to add the list.

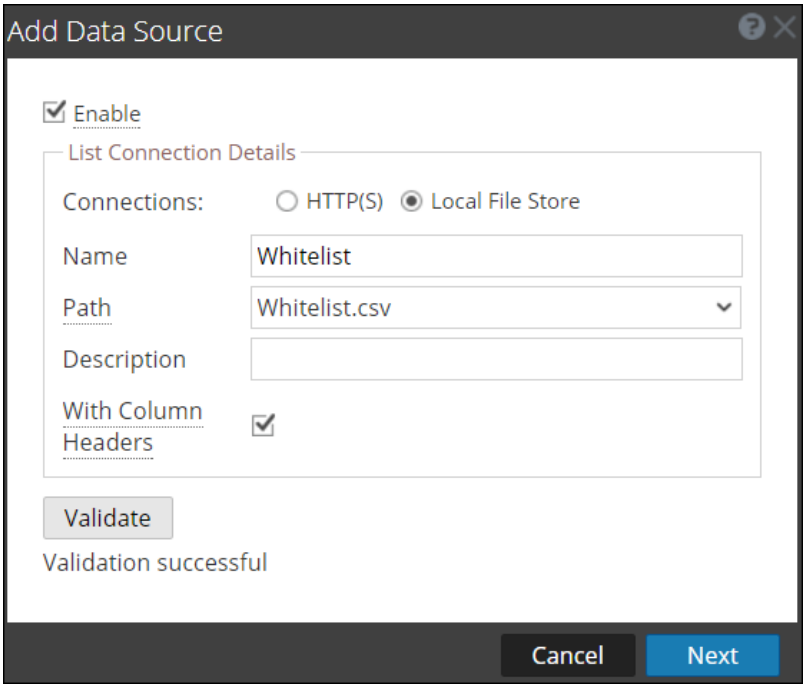
Add List data source using Local File Store

To add a List as a data source:

1. Go to **(Admin) > Services**.
   The services view is displayed.
2. Select the Context Hub service and click > **View** > **Config**.
   The Services Config View of Context Hub is displayed.
3. In the **Data Sources** tab, click > **LIST**.
   The **Add Data Source** dialog is displayed.
4. By default, the **Enable** checkbox is selected. If this option is unchecked, the Next button is disabled, you cannot add the data source, view the list in the list tab and view the contextual information.
5. By default, the Context Highlighting checkbox is enabled. This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub.

Note: You can disable the context highlighting globally in the Context Hub explorer view. After you disable this option, the entity values for all the data sources configured will not be highlighted if there are any contextual information.

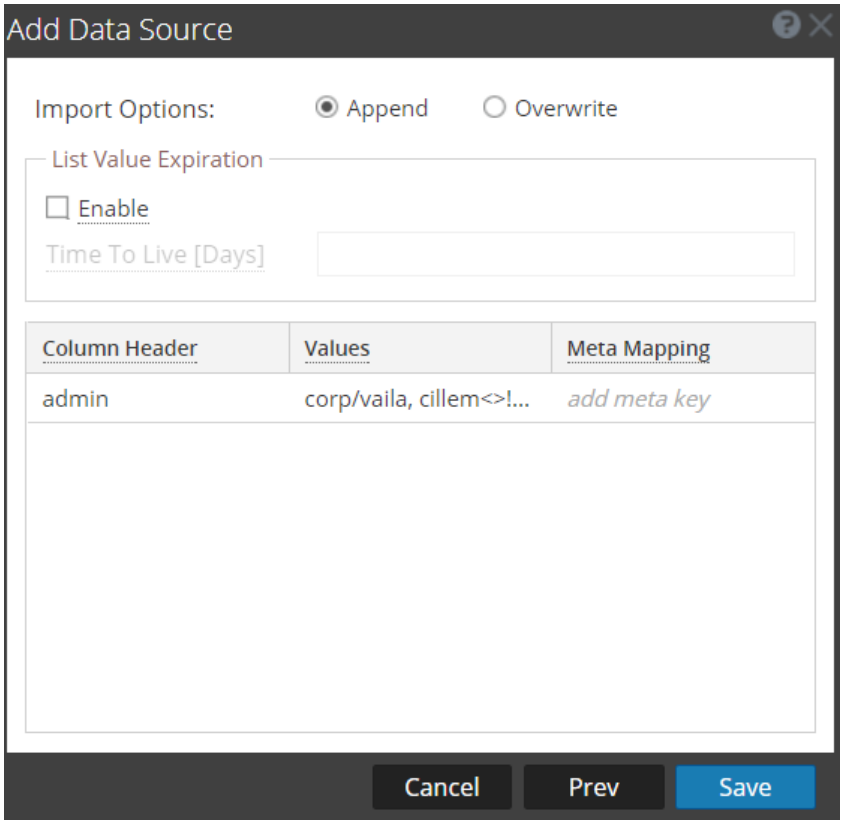
6. Select the **Local File Store** Connection Type.



7. Provide the following database connection details. Enter the following fields for Local File Store Connection Type:
   - **Name:** Provide a name for the list data source.
   - **Path:** This field displays all the data files available in the data folder `/var/lib/netwitness/contexthub-server/data`, where context hub service is running. Select the file name from the drop-down.
     A maximum of 32 columns of CSV file are supported that adhere to the RFC1480 standards.
   - (Optional) **Description**: Add a description for the selected file.
   - **With Column Headers**: Select this option to consider the first row as column headers from the CSV file. If you don't select this option, you need to enter the column headers in the next screen.
8. Click **Validate**.

   If the validation fails, you cannot add the data source.
9. Click **Next**.

   The next dialog is displayed.



10. Select any one of the following options:
    - **Append** - Select this option to add the imported values to an existing list.
    - **Overwrite** - Select this option to replace the values in an existing list with the imported values.
11. In the **List Value Expiration** section, the **Enable** option is unchecked, by default. If you want to store the looked up list values in the cache for a specified number of days then select the **Enable**

checkbox and enter the number of days in the **Time to Live (days)** field for the list values to be retained.

12. In the next screen, map at least one meta key with one or more meta types by mapping a column header with a meta. The description for each field is as follows:
    - **Column Header**: Display headers of the CSV file which must be mapped to a meta type.
    - **Meta Mapping**: Maps a column header field to a meta type.
    - **Values**: Displays the first three values from the imported list.
13. Click **Save**.

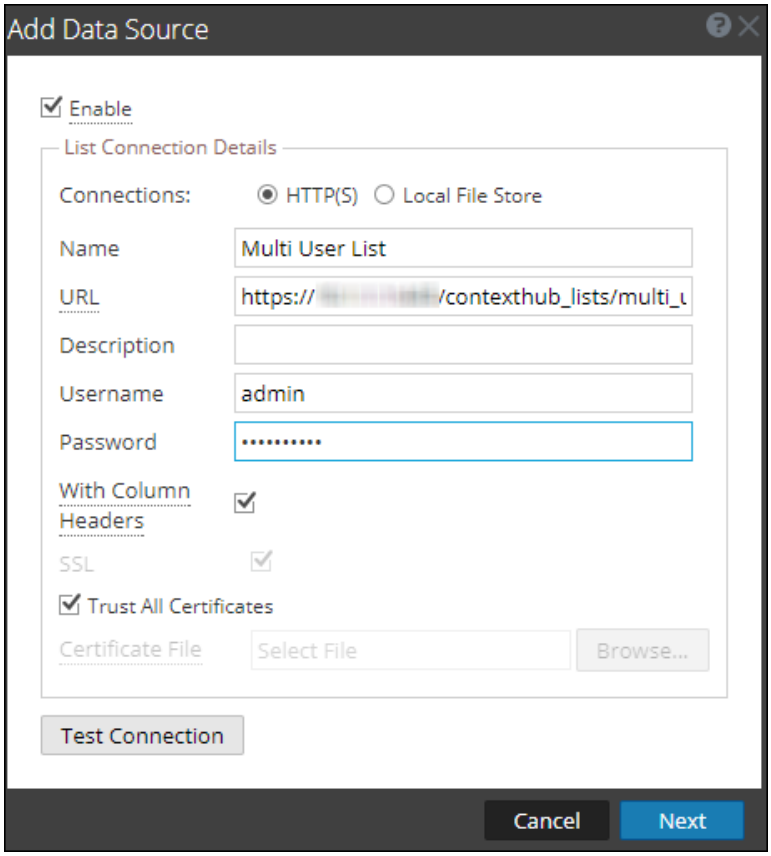
Add List data source using HTTP(S)

To add List as a data source:

1. Select  **(Admin) > Services**.
   The services view is displayed.
2. Select the Context Hub service and click  > **View** > **Config**.
   The Services Config View of Context Hub is displayed.
3. In the **Data Sources** tab, click  > **LIST**.
   The **Add Data Source** dialog is displayed.
4. By default, the **Enable** checkbox is selected. If this option is unchecked, the Next button is disabled, you cannot add the data source, view the list in the list tab and view the contextual information.
5. By default, the Context Highlighting checkbox is enabled. This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled.

Note: You can disable the context highlighting globally in the Context Hub explorer view. After you disable this option, the entity values for all the data sources configured will not be highlighted if there are any contextual information.

6. Select the HTTP(S) Connection Type.

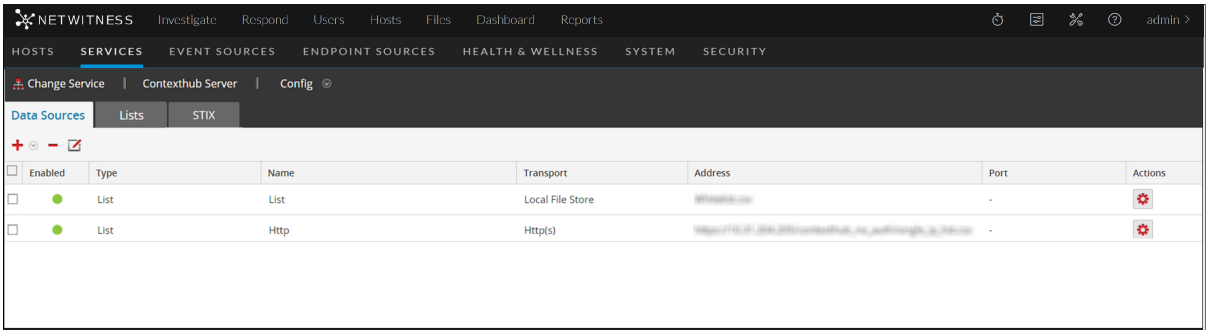


- Enter the following fields for HTTP(S) Connection Type:
  - **Name:** Provide a name for the list data source.
  - **URL**: Enter the path of the CSV file available on the HTTP(S) location along with the host name or IP address of the remote machine where the list is stored. The URL must be of the format: `https://<Hostname or IP-address of the HTTP(S)server>:<Port on which the HTTP(S) server is hosted>/<Absolute path of CSV file>`. For example, `https://10.1.1.1:443/contexthub_lists/multi_user_list.csv`
  - (Optional) **Description**: Add a description for the selected file.
  - (Optional) **Username:** Enter the username to connect to the HTTP(S) server requires basic authentication.
  - (Optional) **Password:** Enter the password to connect to the HTTP(S) server requires basic authentication.
  - **With Column Headers**: Select this option if you want to import a CSV file with headers. If this option is selected and you import the CSV without headers, the first row will be considered as a header which can be edited.
  - **SSL:** If you enter a URL with HTTPS in this field, then this is selected automatically. If you enter a URL with HTTP, then this checkbox is unselected.

- **Trust All Certificates**: Select this checkbox to add the data source without validating the certificate. If you uncheck this option, you need to upload a valid .cer or .crt format HTTP (S)server certificate for the connection to be successful.

7. Click **Test Connection** to test the connection between Context Hub and the data source.
8. Click **Save** to save the settings.

   List is added as a data source for the configured Context Hub and is displayed in the **Data Sources** tab.



Next Steps

- Add, edit, or remove values from a specific list.
- Configure the data source settings to determine the data source fields to be displayed in the Context panel. For instructions, see Configure Context Hub Data Source Settings .
- Import and export a list. For more information, see Import or Export Lists for Context Hub.
- View the contextual data in the Context Summary Panel of the Respond view or Investigate view. For more information, see the *NetWitness Respond User Guide* and *NetWitness Investigate User Guide*.

Configure Archer as Data Source

You can configure Archer as a data source for Context Hub and use the Context Hub service to fetch contextual information from Archer. Use the procedures in this topic to add Archer as a data source for Context Hub service and configure the settings (if required) for Archer.

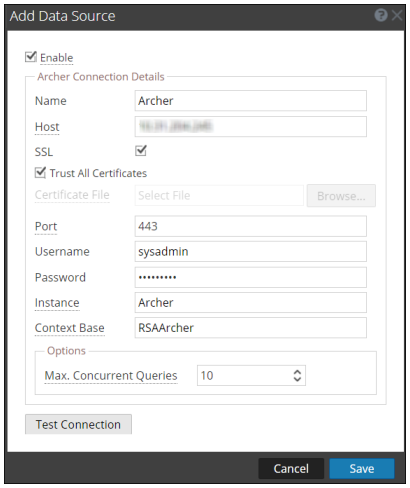
Prerequisites

Before you configure Archer data source, ensure that:

- Context Hub service is available in **(Admin) > Services** view of NetWitness.
- Archer is installed with Licensed Devices application.

To add Archer as a data source for Context Hub:

1. Go to **(Admin) > Services**.
   The Services view is displayed.
2. Select the Context Hub service, and click > **View** > **Config**.
   The Services Config view is displayed.
3. In the **Data Sources** tab, click > **Archer**.
   The **Add Data Source** dialog is displayed.



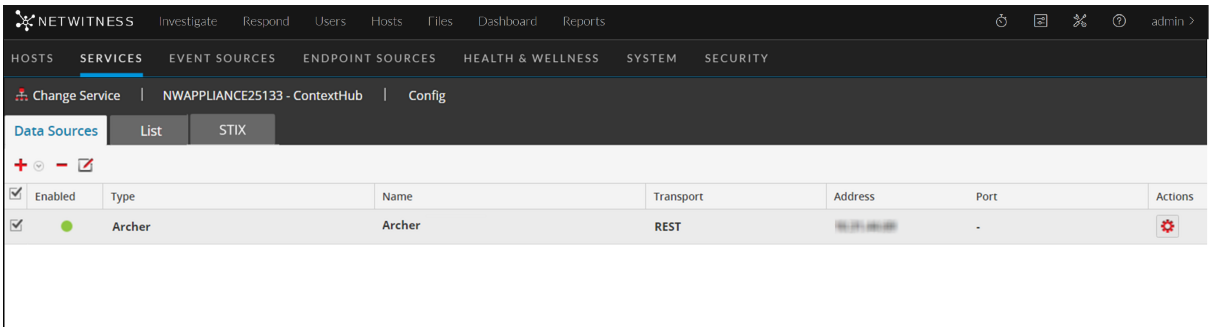
4. Provide the following information:
   - By default, the **Enable** checkbox is selected. If this option is unchecked, the save button is disabled, you cannot add the data source, and cannot view the contextual information.
   - **Context Highlighting**: This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled.

Note: You can disable the context highlighting globally in the Context Hub explorer view. After you disable this option, the entity values for all the data sources configured will not be highlighted if there are any contextual information.

- Enter the following fields:
  - **Name**: Enter a name for Archer data source.
  - **Host**: Enter the hostname or IP address where Archer server is installed.
  - **SSL**: By default this option is selected and enables SSL communication to Archer .
  - **Trust All Certificates**: Select this checkbox to add the data source without validating the certificate. If you uncheck this option, you need to upload a valid Archer server certificate for the connection to be successful.
  - **Port**: The default port is 443.
  - **Username**: Enter the Archer Server username.
  - **Password**: Enter the Archer Server password.
  - **Instance**: Enter the Instance name from which you want to extract data. An Archer instance is a single set up that includes unique content in a database, the connection to the database, the interface, and log-in. You might have individual instances for each office location or region or for development, test, and production environments. The Instance Database stores the Archer content for a specific instance.
  - **Context Base**: Enter the virtual directory name where the files are stored. For example, archer located at the Archer web address https://archer.company.com/**rsaarcher**/default.aspx. If the files are stored in the IIS default web address https://archer.company.com/default.aspx, then this field must be empty.
  - **Max. Concurrent Queries**: You can configure the maximum number of concurrent queries defined by the Context Hub service to be run against the configured data sources. The default value is 10.

5. Click **Test Connection** to test the connection between Context Hub and the Archer data source.
6. Click **Save**.

   Archer is added as a data source for Context Hub and is displayed in the **Data Sources** tab.



After adding the data source, you can configure data source settings. For instructions, see Configure Context Hub Data Source Settings . And View the contextual data in the Context Summary Panel of the Respond view or Investigate view. For instructions, see the *NetWitness Respond User Guide* and *Investigate User Guide*.

Configure Archer Data Source

After you have configured the required data sources you can customize the settings for the data sources based on your requirement.

To access and configure settings:

1. Go to **(Admin) > Services**.

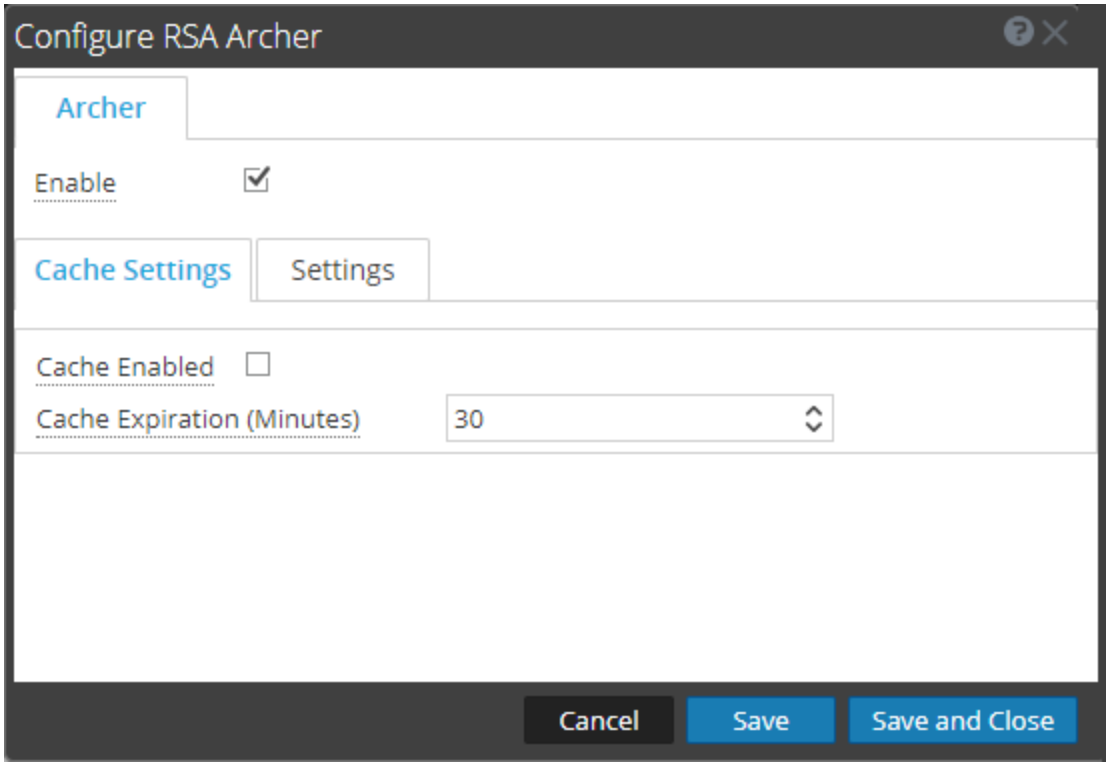
   The services view is displayed.

2. In the Services panel, select the Context Hub service and click > **View** > **Config**.

   The Services Config view of Context Hub is displayed.

3. Select the data source for which you want to configure the settings and click in the Actions column.

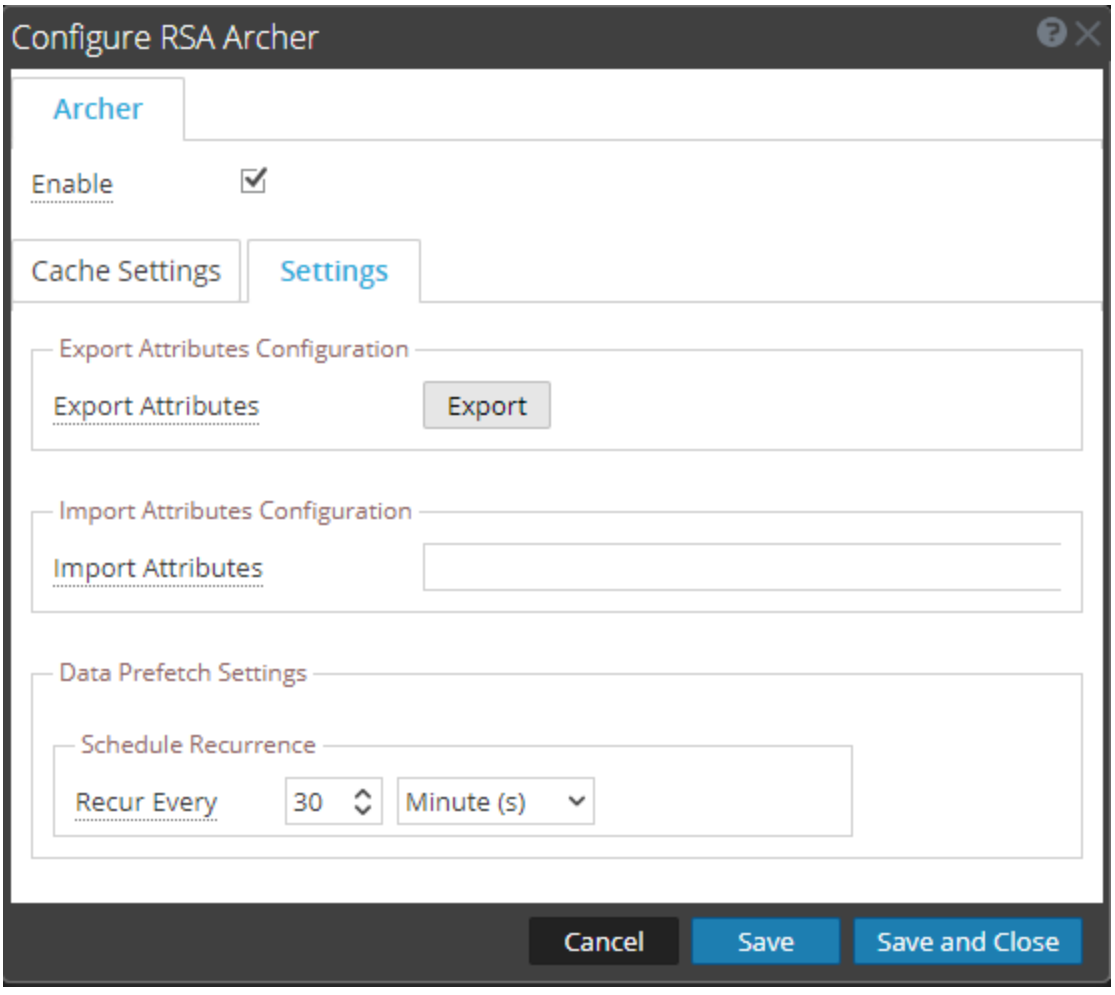
   The following screenshot is an example of the Configure Archer dialog:



4. In the **Settings** tab. Configure the following fields:

Field	Description
Enable	This option is enabled by default (checked) and can be used to enable or disable the response from the selected data source.
Cache Settings	Any lookup from Context Hub can be stored in the Context Hub cache for a configured time. Response to any subsequent matching request will be fetched from the Context Hub cache. Use this section to define the following cache settings for query lookup: • **Cache Enabled**: By default, this checkbox is selected and the query response is cached. • **Cache Expiration (Minutes)**: The maximum time the query lookup is retained in cache. The default time is 30 minutes and maximum is 7200 minutes that you can configure.

5. Click **Cache Settings**. Configure the following fields.



Field	Description
Export Attributes Configuration	In **Settings**, **Export Attributes Configuration**, click **Export** to export the Archer Attributes Configuration. These are the attributes visible in Context Lookup while viewing Archer details for a IP, Host, or Mac. A JSON configuration file gets downloaded and the order of the attributes in sync with the listing in the context panel is maintained in the JSON file.
Import Attributes Configuration	If you want to update or edit the configuration settings, in **Settings**, **Import Attributes Configuration**, click **Browse**. Select the JSON file containing the configuration attributes. The attributes appear in the Context Lookup panel when a user views the context, in the order which they were imported. **Note:** You can backup the previous attributes before importing any changes made to existing attributes.
Data Prefetch Settings	In **Settings**, **Data Prefetch Settings** helps prefetch the data. Configure the **Schedule Recurrence** to provide data faster when you hover over the intended entity in Respond.
Schedule Recurrence	In the **Recur Every** field, enter a value or use the drop-down to configure the recurrence for prefetch. The default time duration can be selected from the drop-down list for configuring the duration of recurrence. Available values are minutes, hours, days, or weeks.

6. Click any one of the following options:

- **Cancel** - select this option to cancel the changes.
- **Save** - select this option to save the changes.
- **Save and Close** - select this option to save and close the dialog.

Note: After you configure the data source settings, you can configure the Context Hub configuration parameters by navigating to  **(Admin) > Services**> **View** > **Explore** view. Make sure you restart the Context Hub service if you make any configuration changes in the Explore view.

Configure Active Directory as a Data Source

You can configure Active Directory (AD) as a data source for Context Hub using LDAP and use the Context Hub service to fetch contextual information from AD. Use the procedures in this topic to add AD as a data source for Context Hub service and configure the settings(if required) for AD.

Prerequisites

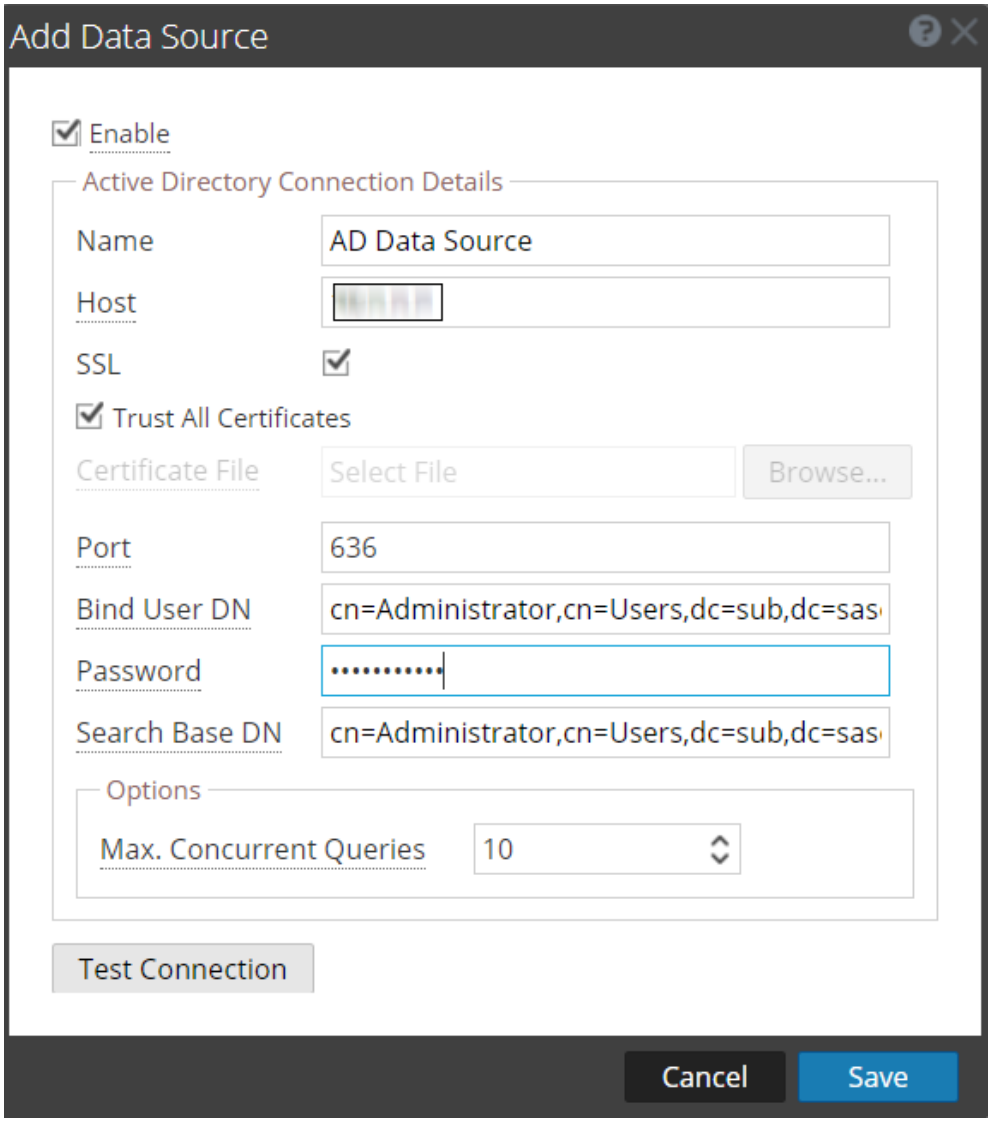
Before you configure Active Directory data source, ensure that:

- Context Hub service is available in **(Admin) > Services** view of NetWitness.
- AD is available and is running on Windows versions 2003, 2008, and 2012 are supported.

To add AD as a data source for Context Hub:

1. Go to **(Admin) > Services**.
   The services view is displayed.
2. Select the Context Hub service and click **> View> Config**.
   The Services Config View of Context Hub is displayed.

3. In the **Data Sources** tab, click + **> Active Directory**.
   The **Add Data Source** dialog is displayed.



You need to configure the Active Directory schema to replicate the following attributes to view the data in the Respond page:

- Employee ID
- Department
- Company
- Title
- Postal Code

All the other attributes replicate automatically.

4. Provide the following database connection details:

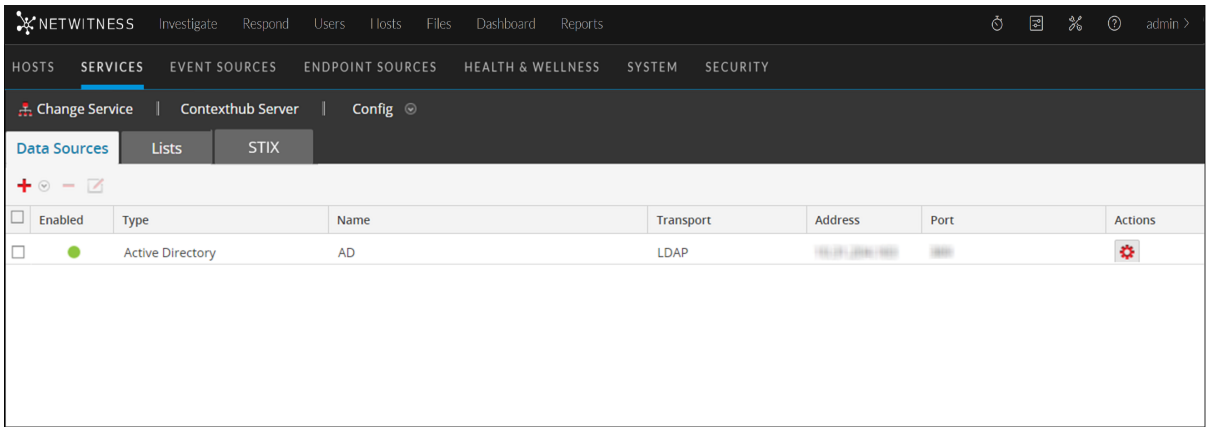
- By default, the **Enable** checkbox is selected. If this option is unchecked, the save button is disabled, you cannot add the data source, and cannot view the contextual information.
- Enter the following fields.
  - **Name:** Enter a name for the AD data source.
  - **Host:** Enter the host name or IP address of the AD.
  - **SSL:** By default this will be checked with 636 port number which will connect to the data source using Secure Sockets Layer (SSL) connection.
  - **Trust All Certificates**: Select this checkbox to add the data source without validating the certificate. If you uncheck this option, you need to upload a valid .cer or .crt format Active Directory server certificate for the connection to be successful. If you add multiple AD data sources with ssl, you should configure all the data sources with either a valid certificate or a Trust All Certificates.
  - **Port:** The default port is 636 with SSL and 389 without SSL.
    If you want to fetch data from multi-domains you can configure a single data source with the Global catalog port (3269 with SSL or 3268 without SSL).

    Alternately, for multi-domain, you can configure a single data source for each domain with the default port (389 with SSL or 636 without SSL).

    Multi-forest is a collection of multi-domains. If you want to fetch data from multi-forest you need to configure each forest with the Global catalog port (3269 with SSL or 3268 without SSL).
  - **Password:** Enter password of the user DN used to bind with AD.
  - **Bind User DN:** The distinguished name of the user that will authenticate to the search directory. For example, cn=Administrator,cn=Users,dc=sub,dc=saserver,dc=local.
  - **Search Base DN:** The base distinguished name, or base DN, identifies the entry in the directory from which searches are initiated; the base DN is often referred to as the search base. For example, dc=sub,dc=saserver,dc=local.

5. Click **Test Connection** to test the connection between Context Hub and the data source.
6. Click **Save**.
   AD is added as a data source for the configured Context Hub. The added AD data source is displayed

in the **Data Sources** tab.



After adding the data source, you can configure the data source settings. For instructions, see Configure Context Hub Data Source Settings .

Next steps

After completing the configuration, you can view the contextual data in the Context Summary Panel of the Respond view or Investigate view. For instructions, see the **Navigate to Context Summary Panel and View Additional Context** topic in the *Investigate User Guide*.

Configure NetWitness Endpoint as a Data Source

You can configure NetWitness Endpoint as a data source for Context Hub and use the Context Hub server to fetch contextual information from NetWitness Endpoint. Use the procedures in this topic to add NetWitness Endpoint as a data source for Context Hub service and configure the settings (if required) for NetWitness Endpoint.

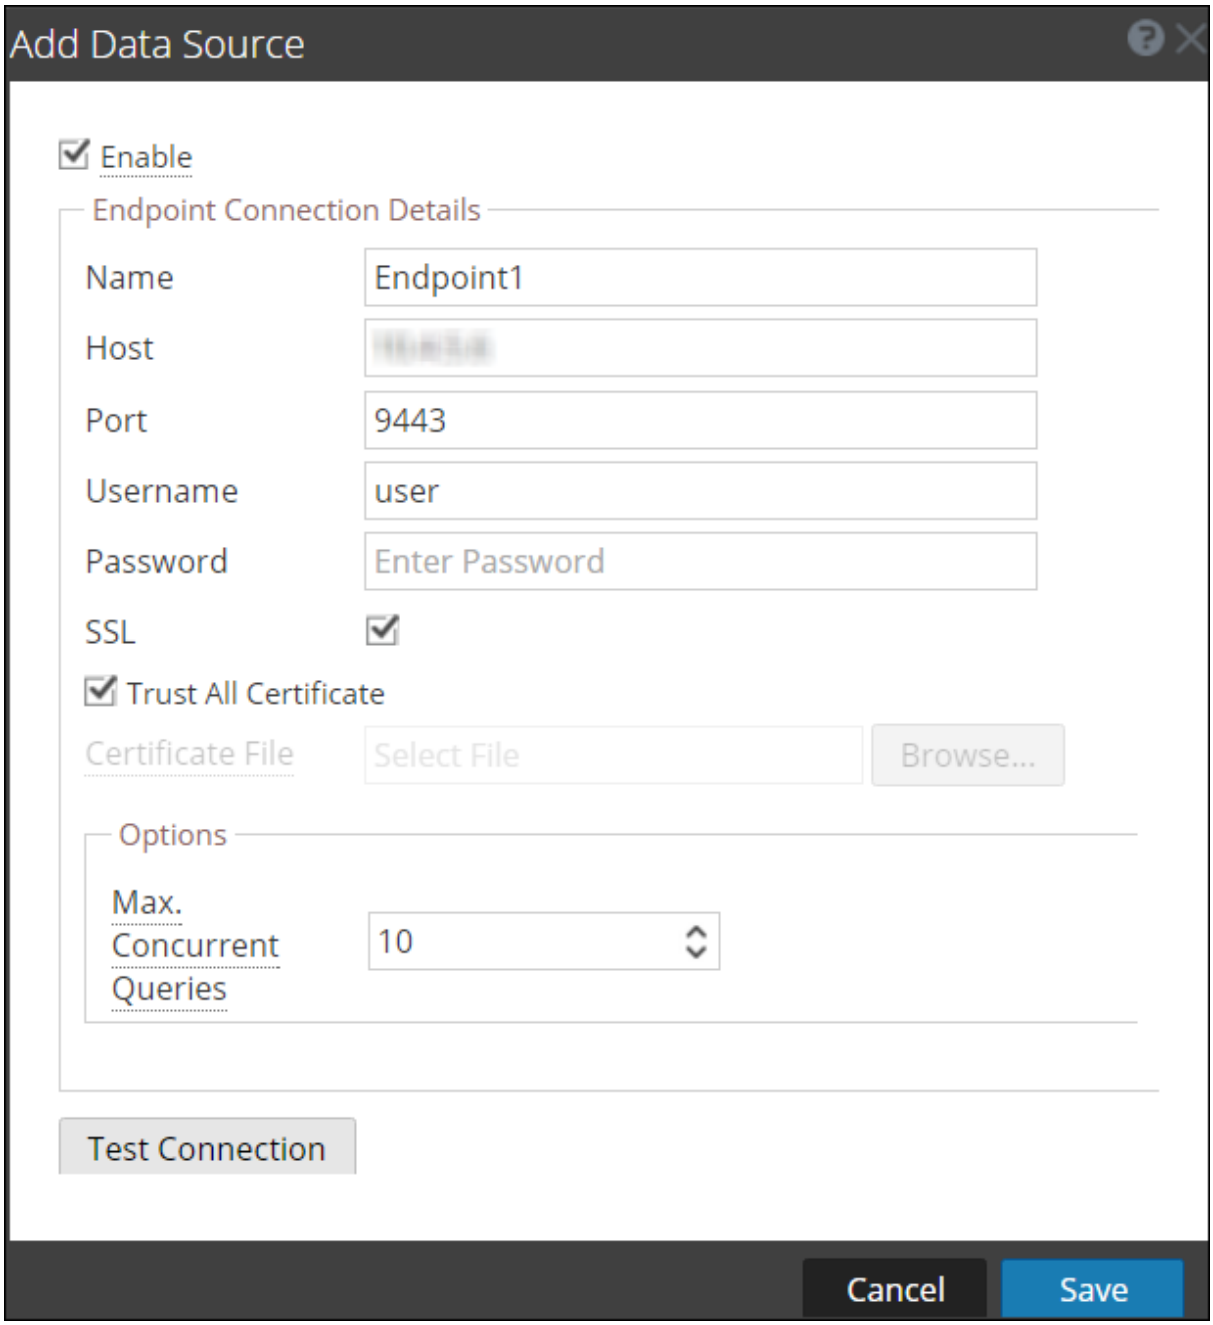
Prerequisites

Before you configure NetWitness Endpoint data source, ensure that:

- Context Hub service is available in  **(Admin) > Services**.view of NetWitness.
- NetWitness Endpoint (v4.1.1 to 4.3.0.5) is installed and configured.
  For more information on how to install, configure and for detailed information on NetWitness Endpoint, see the NetWitness Endpoint documents available at NetWitness Community.

To add NetWitness Endpoint as a data source for Context Hub:

1. Go to **(Admin) > Services**.
   The Services view is displayed.
2. Select the Context Hub service, and click > **View** > **Config**.
   The Services Config view is displayed.
3. In the **Data Sources** tab, click > **Endpoint**.
   The **Add Data Source** dialog is displayed.



4. Provide the following information:

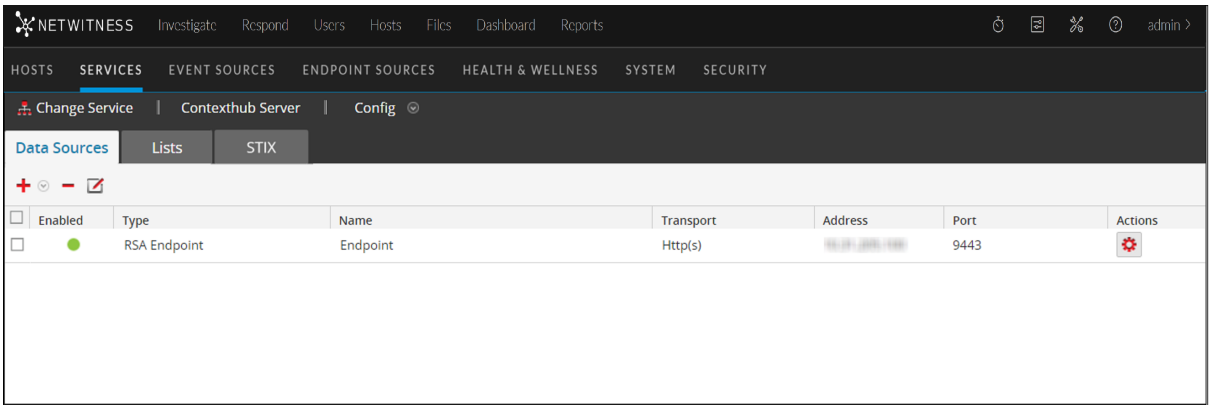
   - By default, the **Enable** checkbox is selected. If this option is unchecked, the save button is disabled, you cannot add the data source, and cannot view the contextual information.
   - **Context Highlighting** – This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled.

   Note: You can disable the context highlighting globally in the Context Hub explorer view. After you disable this option, the entity values for all the data sources configured will not be highlighted if there are any contextual information.

   - Enter the following fields:
     - **Name**: Enter a name for NetWitness Endpoint data source.
     - **Host**: Enter the hostname or IP address where NetWitness Endpoint API server is installed.
     - **Port**: The default port is 9443.
     - **SSL**: Select SSL if you want NetWitness to communicate with the host using SSL. This is enabled by default.
     - **Username**: Enter the NetWitness Endpoint API Server username.
     - **Password**: Enter the NetWitness Endpoint API Server password.
     - **Trust All Certificates**: Select this checkbox to add the data source without validating the certificate. If you uncheck this option, you need to upload a valid server generated or CA certificate to authenticate the connection with the supported formats of .cer or .crt of Base64 [PEM] encoded or DER encoded.
     - **Max. Concurrent Queries**: You can configure the maximum number of concurrent queries to be run against the configured data sources. The default value is 10.

5. Click **Test Connection** to test the connection between Context Hub and the NetWitness Endpoint.

6. Click **Save**.
   NetWitness Endpoint is added as a data source for Context Hub and is displayed in the **Data Sources** tab.



Next steps

After adding the data source, you can configure the settings. For more information, see Configure Context Hub Data Source Settings .

Also you can view the contextual data in the Context Summary Panel of the Respond view or Investigate view. For more information, see the *NetWitness Respond User Guide* and the *NetWitness Investigate User Guide*

Configure Respond as a Data Source

You can configure Respond as a data source for Context Hub and use the Context Hub service to fetch contextual information from Respond service. If Respond service is already configured, the configuration details are pre-populated while adding Respond as a data source. Use the procedures in this topic to add Respond as a data source for Context Hub service and configure the settings.

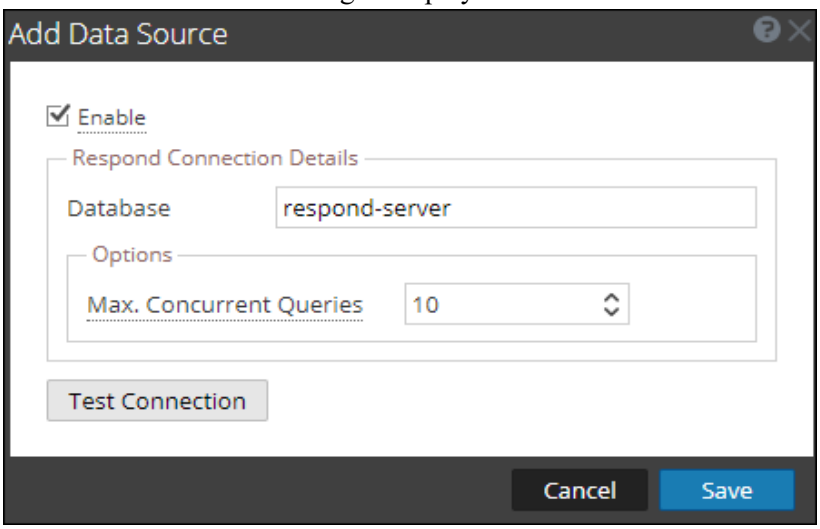
Prerequisites

Before you configure Respond data source, ensure that:

- Context Hub service is available in **(Admin) > Services** view of NetWitness.
- Respond service is available.

To add Respond as a data source for Context Hub:

1. Go to **(Admin) Services**.
   The services view is displayed.
2. Select the Context Hub service and click > **View** > **Config**.
   The Services Config View of Context Hub is displayed.
3. In the **Data Sources** tab, click > **Respond**.
   The **Add Data Source** dialog is displayed.

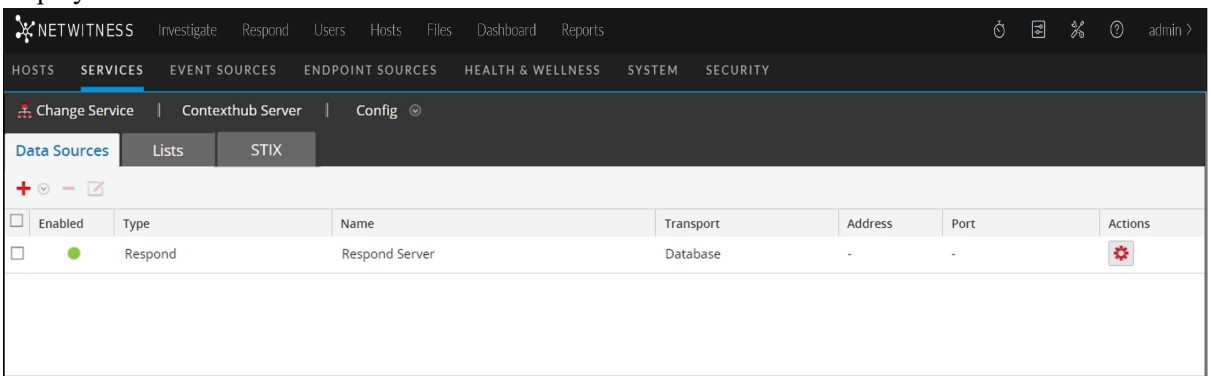


The required fields to configure the Respond data source are automatically updated.

You can enable or disable Context Highlighting. This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled.

Note: You can disable the context highlighting globally in the Context Hub explorer view. After you disable this option, the entity values for all the data sources configured will not be highlighted if there are any contextual information.

4. Click **Test Connection** to test the connection between Context Hub and the data source.
5. Click **Save**.
   Respond is added as a data source for the configured Context Hub. The added Respond data source is displayed in the **Data Sources** tab.



After adding the data source, you can configure the settings. For more information, see Configure Context Hub Data Source Settings .

Next steps

After completing the configuration, you can view the contextual data in the Context Summary Panel of the Respond view or Investigate view. For more information, see the *NetWitness Respond User Guide* and the *NetWitness Investigate User Guide*.

Configure File Reputation Server as a Data Source

File Reputation Server provides analysts the opportunity to view reputation status of files. By default, File Reputation is enabled in Additional Live Services section.

If Context Hub service is configured, File Reputation Server is automatically added as data source for Context Hub.

Prerequisites

Ensure that:

- Context Hub is enabled and the service is available in  **(Admin) > Services** view of NetWitness.
- Live Account is available.

Note: To create a Live Account, see the **Step 1. Create Live Account** topic in the *Live Services Management Guide*.

By default, **File Reputation** is enabled in **Additional Live Services** section. Before setting up File Reputation data source, make sure that you have signed in to your Live account with your Live Account Credentials and Context Hub is enabled. File Reputation is automatically added as a data source for context hub.

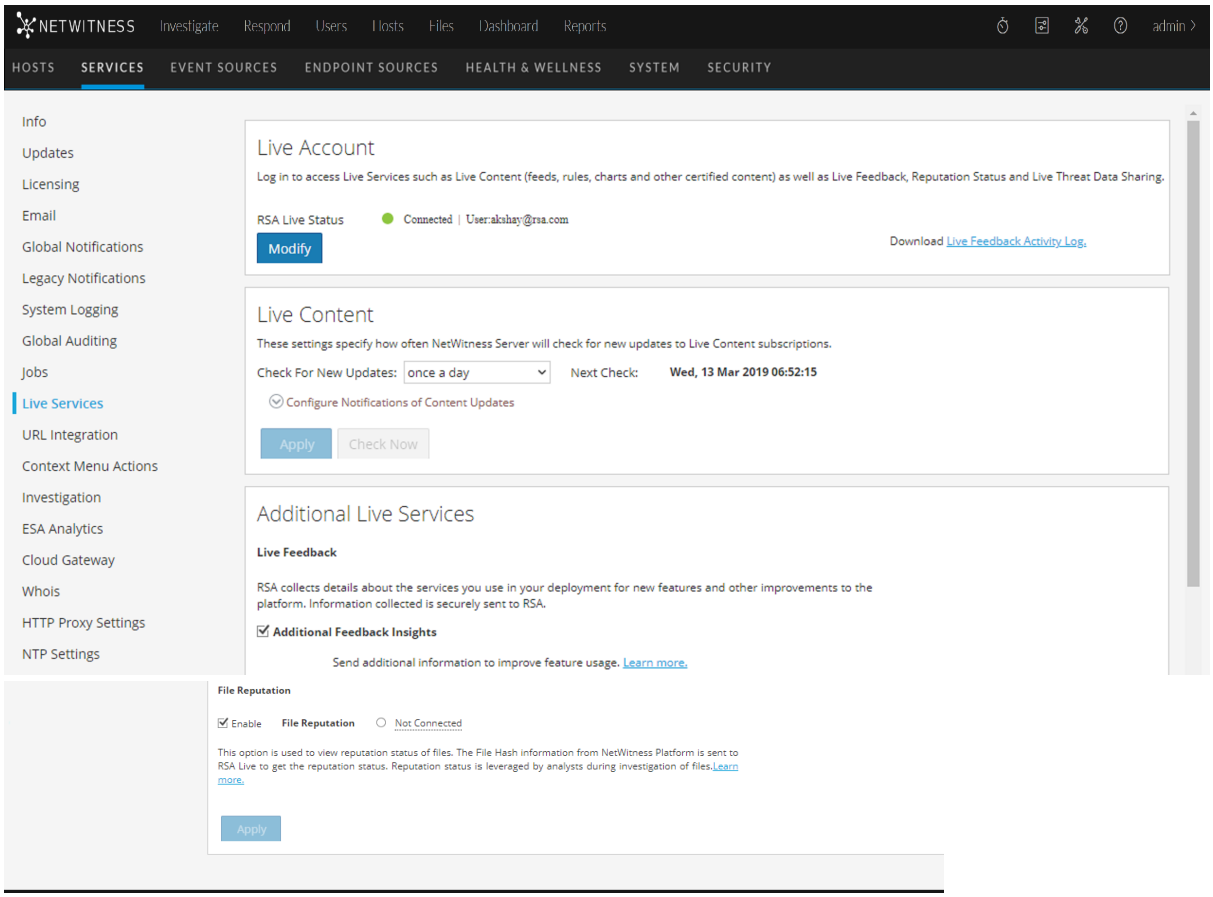
For information about configuring Live Account and Live Services, see the **Configure Live Services Settings** topic in the *System Configuration Guide*.

For information about configuring Context Hub service, see the **Step 1. Add the Context Hub Service** topic in the *Context Hub Configuration Guide*.

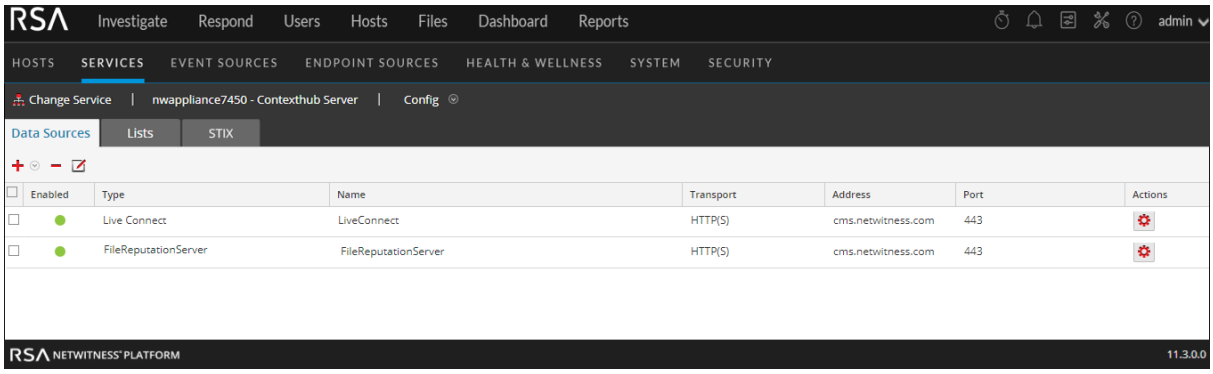
Enable or Disable File Reputation Data Source

To enable or disable File Reputation data source for Context Hub:

1. Go to  **(Admin) > System**.
2. In the left navigation pane, select **Live Services**.
3. In the **Additional Live Services** section, enable **File Reputation**.



4. Click **Apply**.
   File Reputation Server data source is enabled for Context Hub service.
5. To verify, go to the **Data Sources** tab and view the available sources.
   File Reputation source must be added to the list of available sources and the **Enabled** field must be a solid green circle ( ).



6. To disable File Reputation data source, disable **File Reputation** in Additional Live Services panel and click **Apply**.

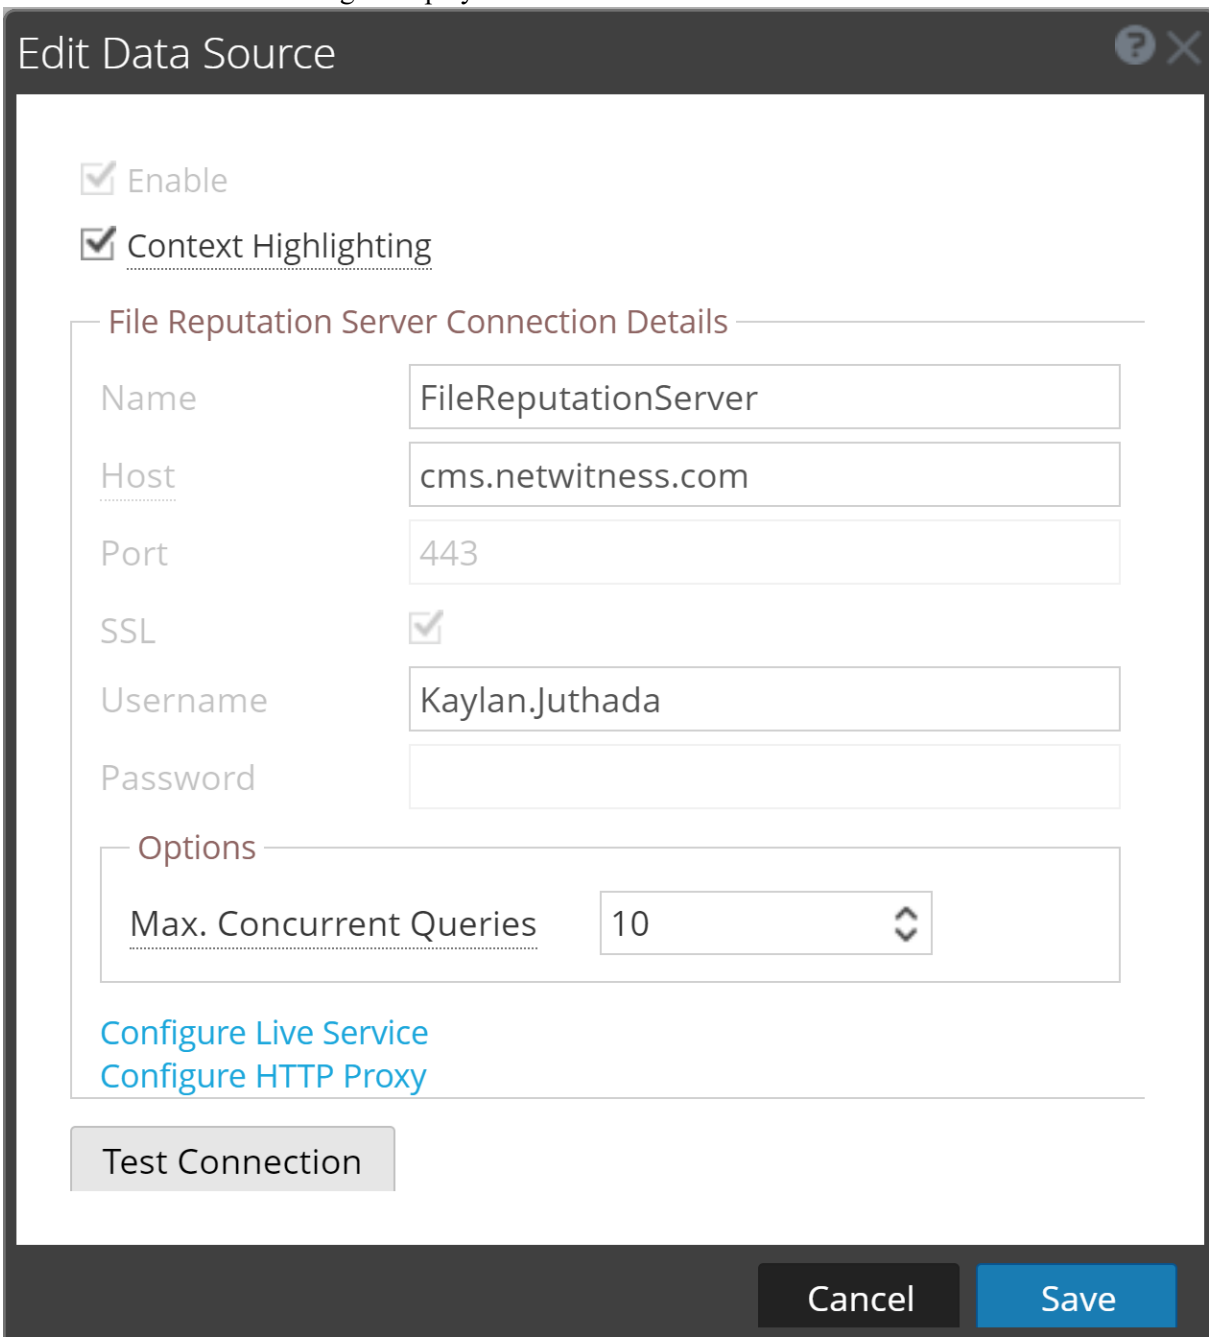
   File Reputation data source is disabled for Context Hub service.

Edit File Reputation Server Data Source Settings

To edit File Reputation Server data source for Context Hub:

1. Select  **(Admin) > Services**.
   The Services view is displayed.
2. In the **Services** panel, select the Context Hub service, and  **> View > Config**.
   The Services Config view is displayed.
3. In the **Data Sources** tab, select the File Reputation Server source and click .

The **Edit Data Source** dialog is displayed.



4. Edit the required fields:

Field	Description
Context Highlighting	This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled. **Note:** You can disable the context highlighting globally in the Context Hub explorer view. After you disable this option, the entity values for all the data sources configured will not be highlighted if there are any contextual information.
Max. Concurrent Queries	You can configure the maximum number of concurrent queries defined by the Context Hub service to be run against the configured data sources. The default value is 25.

5. To edit the Proxy settings, see **the HTTP Proxy Settings Panel** topic in the *System Configuration Guide*.
6. Click **Test Connection** to test the connection between Context Hub and the data source.
7. Click **Save** to save the settings.

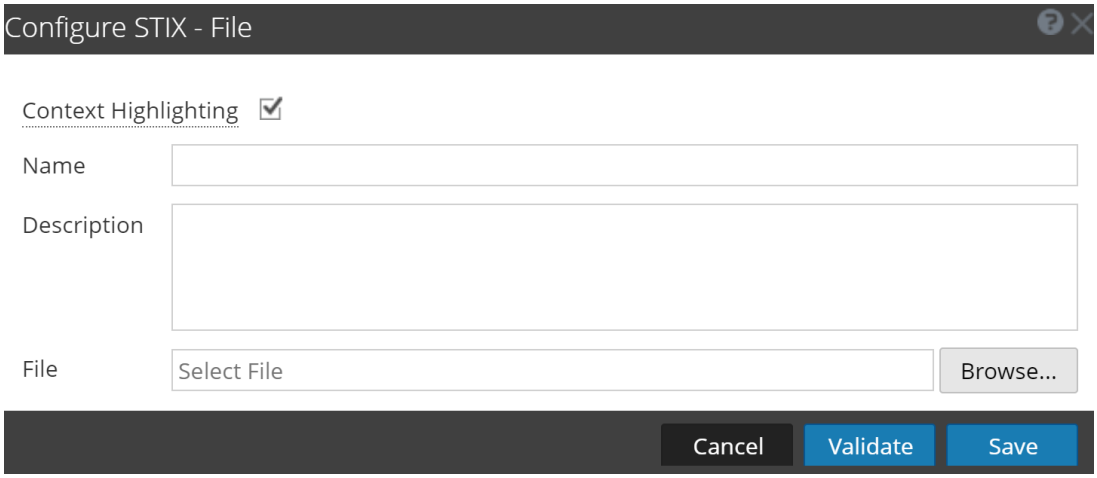
Configure STIX as a Data Source

You can configure Structured Threat Information eXpression (STIX) as a data source for Context Hub and use the Context Hub service to fetch contextual threat intelligence information from a STIX source.

Configure STIX File

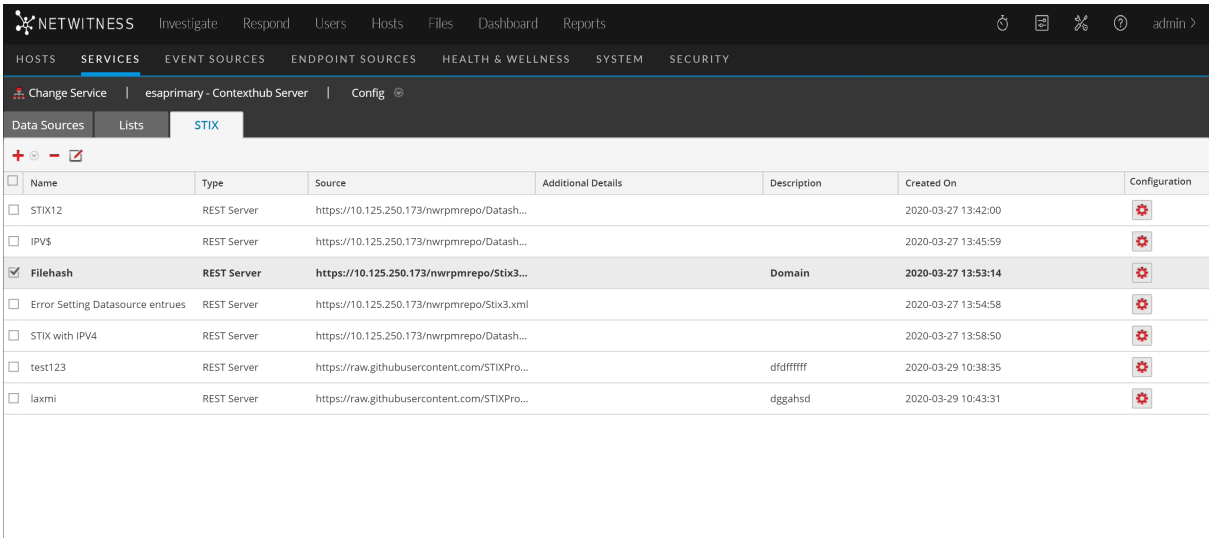
To add STIX as a data source for Context Hub:

1. Go to  **(Admin) > Services** .
   The Services view is displayed.
2. Select the Context Hub service and click  > **View** > **Config**.
   The Services config view of Context Hub is displayed.
3. Click the **STIX** tab, and click  .
4. Select **File** as data source.



4. Provide the following details:
   a. **Context Highlighting**: This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled.
   b. **Name**: Provide a name for the STIX file data source.
   c. **Description**: Provide description of the data source.
   d. **File**: Browse for the file you want to add as a data source.
5. Click **Validate** to verify the format of the file.

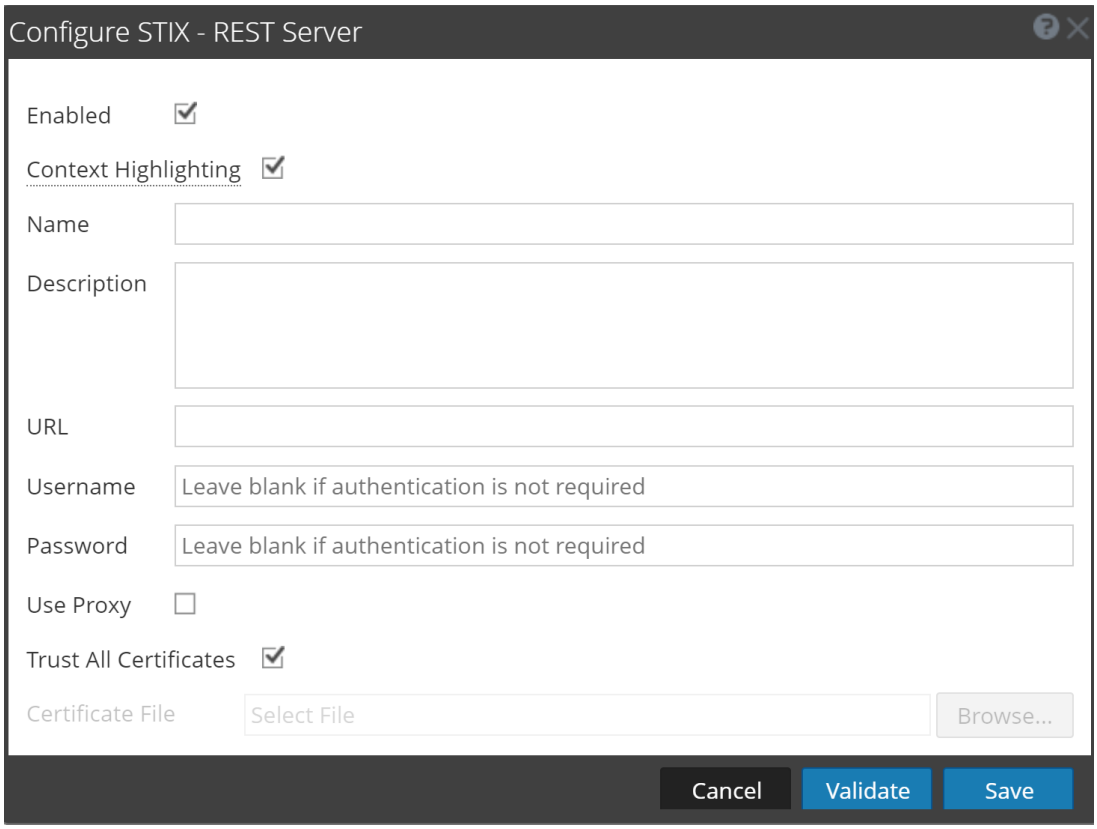
6. Click **Save** to configure the data source.
   The File is added as a data source for the configured Context Hub and is displayed in the **STIX** tab.



Configure REST Server

To add REST as a data source for Context Hub:

1. Go to **(Admin) > Services**.
   The Services view is displayed.
2. Select the Context Hub service and click > **View** > **Config**.
   The Services config view of Context Hub is displayed.
3. Click the **STIX** tab, and click .
4. Select **REST Server** as data source.



4. Provide the following details:
   a. **Enabled**: Select this checkbox to enable the connection.
   b. **Context Highlighting**: This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled.
   c. **Name**: Provide a name for the REST Server data source.
   d. **Description**: Provide a description for the data source.
   e. **URL**: Specify the URL to the STIX file to be hosted on the server.
   f. (Optional) **Username**: Enter the username for the REST server.
   g. (Optional) **Password**: Enter the password for the REST server.
   h. **Use Proxy**: Select this checkbox to use proxy.
   i. (Optional) **Trust All Certificates**: Select this checkbox if you want to trust all certificates and do not have a custom certificate.
   j. (Optional) **Certificate File**: Browse for the certificate file if you have not selected the Trust All certificates checkbox.
5. Click **Validate** to verify the connection parameters to the REST Server.

6. Click **Save** to configure the data source.
   The REST Server is added as a data source for the configured Context Hub and is displayed in the **STIX** tab.

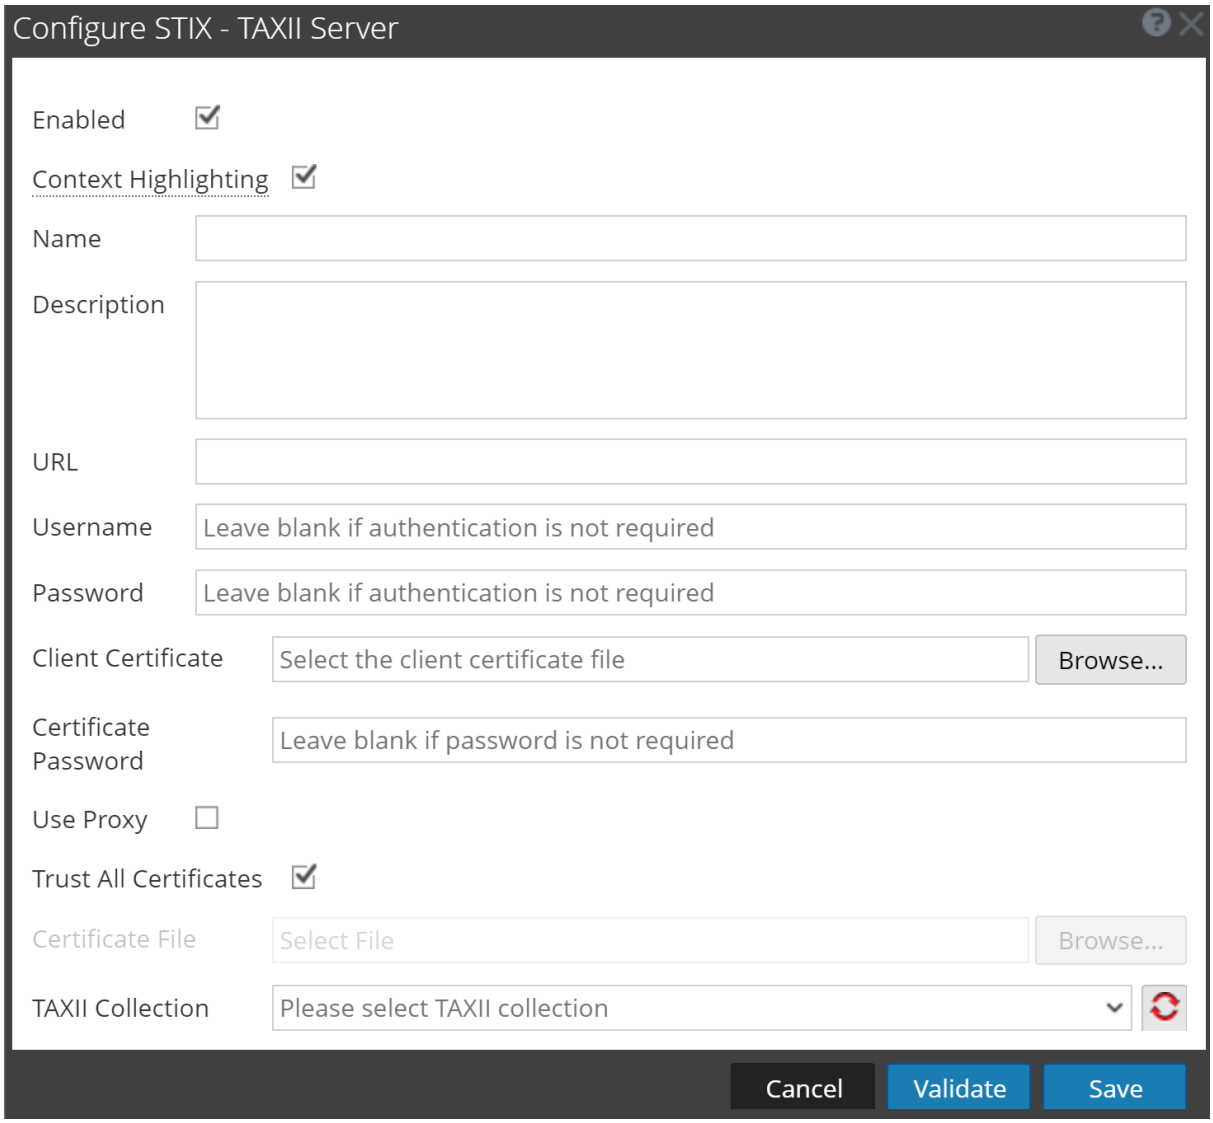
After adding the data source, you can configure additional settings. For more information, see Configure Context Hub Data Source Settings .

Configure TAXII Server

To add TAXII Server as a data source for Context Hub:

1. Go to  **(Admin) > Services**.
   The Services view is displayed.
2. Select the Context Hub service and click  > **View** > **Config**.
   The Services config view of Context Hub is displayed.
3. Click the **STIX** tab, and click .

4. Select **TAXII Server** as data source.



5. Provide the following details:
   a. **Enabled**: Select this checkbox to enable the connection.
   b. **Context Highlighting**: This highlights the meta values (in the **Investigate** > **Navigate**, **Events**, **Event details** and **Nodal graph**) for which the contextual information is available for this data source in the Context Hub. By default, this option is enabled.
   c. **Name**: Provide a name for the TAXII Server data source.
   d. **Description**: Provide a description for the data source.
   e. **URL**: Specify the discovery URL to the TAXII Server.
   f. (Optional) **Username**: Enter the username for the TAXII server.
   g. (Optional) **Password**: Enter the password for the TAXII server.

h. (Optional) **Client Certificate**: Browse to upload a pkcs12 format client certificate available on your local system.

i. (Optional) **Certificate Password**: Enter the password to the certificate, if it is password-protected.

j. (Optional) **User Proxy**: Select this checkbox to use proxy.

k. (Optional) **Trust All Certificates**: Select this checkbox if yoenabu want to trust all certificates and do not have a custom certificate.

l. (Optional) **Certificate File**: Browse for the certificate file if you have not selected the Trust All certificates checkbox.

m. **TAXII Collection**: Select the TAXII Collection name from the drop-down to automatically download the collection.

6. (Optional) Click  to manually retrieve the list of collections available in the TAXII server , if the collections are not downloaded automatically.
7. Click **Validate** to verify the connection parameters to the TAXII Server.
8. Click **Save** to configure the data source.
   The TAXII Server is added as a data source and is displayed in the **STIX** tab.

After adding the data source, you can configure additional settings. For more information, see Configure Context Hub Data Source Settings .

Note: You can disable the context highlighting globally in the Context Hub explorer view. After you disable this option, the entity values for all the data sources configured will not be highlighted if there are any contextual information.

Next steps

After completing the configuration, you can view the contextual data in the Context Summary Panel of the Respond view or Investigate view. For more information, see the *NetWitness Respond User Guide* and the *NetWitness Investigate User Guide*.

Configure REST API as a Data Source

You can configure REST API as a data source for Context Hub which allows analysts to use the Context Hub during an investigation to fetch contextual information from any accessible web service with an exposed REST API. Use the procedures in this topic to add REST API as a data source for Context Hub service and configure the settings (if required) for REST API. The maximum number of REST APIs supported by default is ten.

Prerequisites

Before you configure REST API data source, ensure that:

• Context Hub service is available in  (Admin) > **Services** view of NetWitness Platform.

• Third-party REST API web server is reachable.

To add REST API as a data source for Context Hub:

1. Go to (**Admin**) > **Services**.
   The Services view is displayed.

2. Select the Context Hub service and click  > **View** > **Config**.
   The Services Config view is displayed.

3. In the **Data Sources** tab, click  > **REST API**.
   The Add Data Source dialog is displayed.

4. Provide the following information:
   - By default, the **Enable** checkbox is selected. If this option is unchecked, the save button is disabled, you cannot add the data source, and cannot view the contextual information.
   - Enter the following fields.
     a. **Name** - Enter a name for the REST API data source.
     b. **URL** - Specify the REST API URL with the query parameters.
     The following table lists some examples of URL and associated parameters:

URL	Parameters
https://www.virustotal.com/vtapi/v2/file/report?apikey=<API Key>&resource=$metaValue	• **API key** – Specify your API key • **$metaValue** - You must specify the Test meta value in step 7 to test the REST API connection. For example, C5D0065B594A4775E26FA9875B21189F

URL	Parameters
https://www.virustotal.com/vtapi/v2/url/report?apikey= 1bfdf66f3b0f1875381830ec933f7151f5b48c1cbaa54ac625f2d5a74312e9a2&resource=$metaValue	• **API key** – Specify your API key • **$metaValue** - You must specify the Test meta value in step 7 to test the REST API connection. For example, textspeier.de

c. **Description** – Enter the description of the REST API data source.

d. **User name**: Enter the username of the REST API if it needs to be authenticated.

e. **Password**: Enter the password of the RSET API if it needs to be authenticated.

f. **Response Type**: Select the REST API response type. Possible values are:

- HTML
- JSON

g. **Test Meta value**: Enter the meta value to test the REST API connection.

Note:
- The meta value you specify will be replaced with `$metavalue` in the URL mentioned in step 2 for test connection.
- During context lookup, the `$metavalue` variable is replaced with meta value in the REST API call.

h. **SSL** - Enables SSL communication to REST API. By default, this option is selected.

i. (Optional) **Trust All Certificates**: Select this checkbox if you want to trust all certificates and do not have a custom certificate. By default, this option is enabled.

j. (Optional) **Certificate File**: Browse for the certificate file if you have not selected the Trust All certificates checkbox.

5. Click **Test Connection** to test the connection between Context Hub and the REST API data source.
6. Click **Next**.
7. In the **Response Preview** section, you can view the live response for the REST API configured, using the test meta value as a placeholder.
8. In the **Meta Mapping** section, select one or more meta type supported by the REST API to view context lookup (for meta values) in the Respond and Investigation views.
9. In the **Field Mapping** section, you can add friendly display name for the response field for which you want to perform context lookup. Do the following:

a. Click .

b. In the **Field** drop-down, enter the path for which you want to add a friendly name. As you type the path, the auto-suggest function looks for the matches in the JSON path returned in the response preview. You can also add the path that is not available in the response preview but available for other meta values.

c. **Value (from Preview)** field is filled automatically when you select the path.

d. In the **Display Name** field, enter the friendly name.

> **Note:** After configuration, only the fields that are mapped with friendly names are displayed during context lookup. If you do not map any fields, all fields in the JSON will be available during context lookup.

10. Click **Save** to configure the data source.

> **Note:** After you configure the data source settings, you can configure the Context Hub configuration parameters by navigating to  (**ADMIN**) > **Services**> **View** > **Explore** view.

Edit REST API Data Source

To edit REST API data source for Context Hub:

1. Go to  (**Admin**) > **Services**.
2. Select the Context Hub service and click  > **View** > **Config**.
3. In the Data Sources tab, click  > **REST API**.
4. Edit the required fields.
5. Click **Test Connection** to test the connection between Context Hub and the data source.
6. In the **Meta and Field Mapping** and **Cache settings**, edit the required fields.
7. Click **Save** to save the settings.

Delete REST API Data Source

To edit REST API data source for Context Hub:

1. Go to  (**Admin**) > **Services**.
2. Select the Context Hub service and click  > **View** > **Config**.
3. In the **Data Sources** tab, select the data source you want to delete.
4. Click .

Next steps

After completing the configuration, you can view the contextual data in the Context Summary Panel of the Respond view or Investigate view. For more information, see the *NetWitness Respond User Guide* and the *NetWitness Investigate User Guide*.

Configure Context Hub Data Source Settings

After you have configured the required data sources you can customize the settings for the data sources based on your requirement.

To access and configure settings:

1. Go to  **(Admin) > Services**.

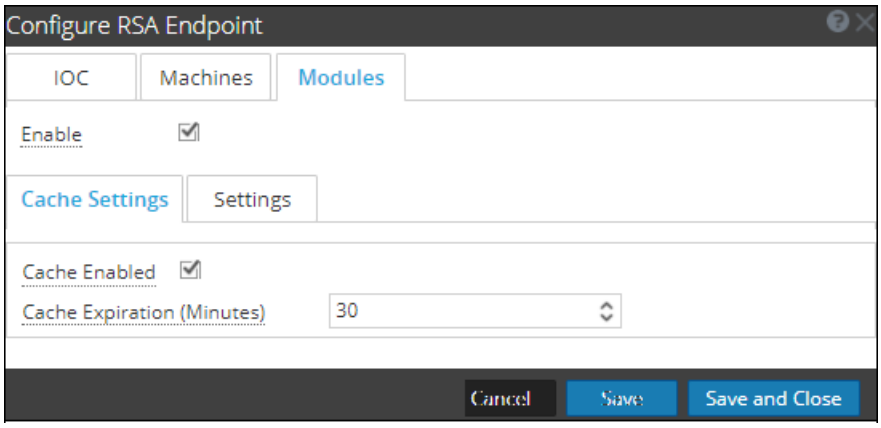
   The services view is displayed.

2. In the Services panel, select the Context Hub service and click > **View** > **Config**.

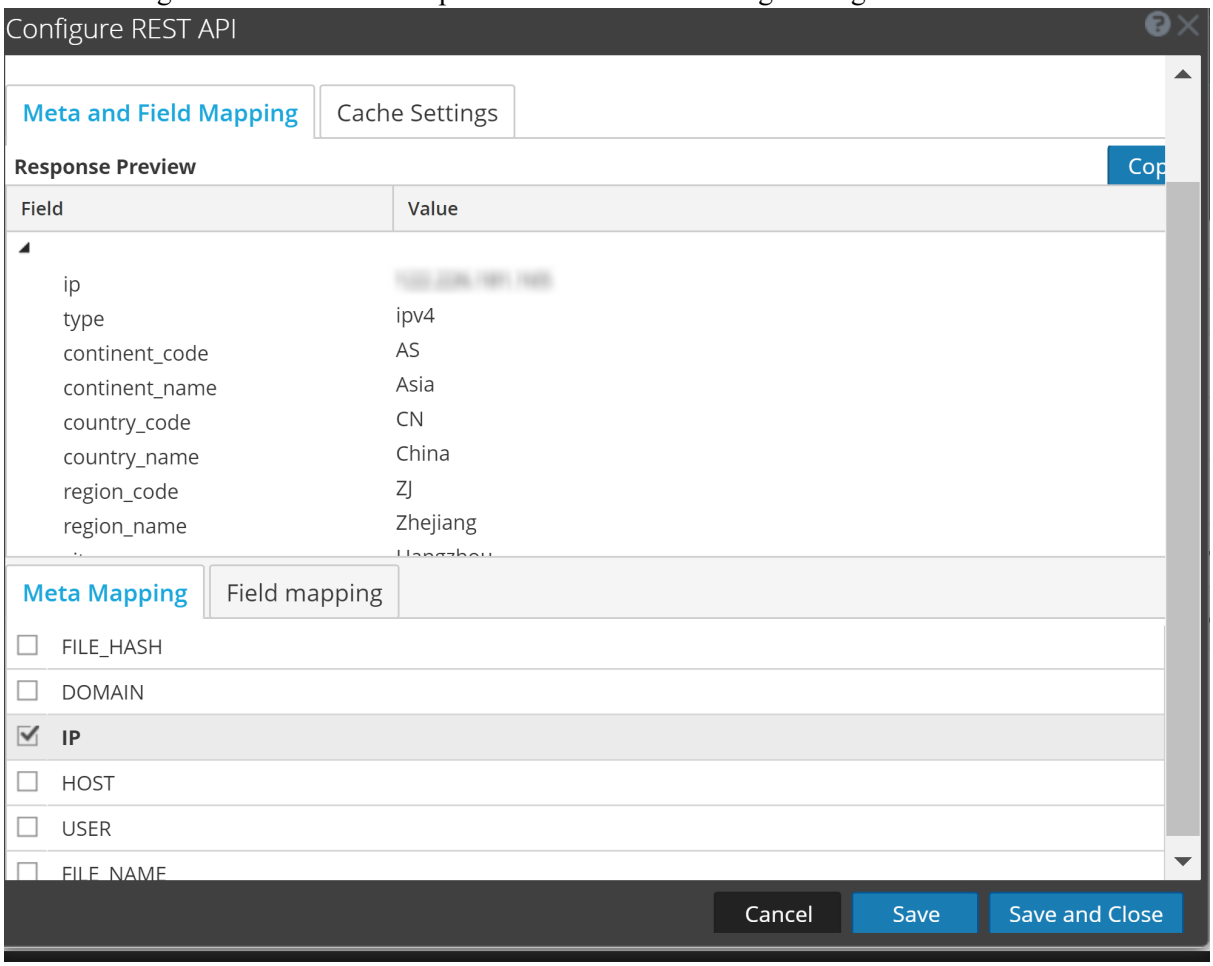
   The Services Config view of Context Hub is displayed.

3. Select the data source for which you want to configure the settings and click  in the Actions column.

   The following screenshot is an example of the NetWitness Endpoint settings dialog:



The following screenshot is an example of the REST API settings dialog:



4. Configure the following fields:

Field	Description
Enable	This option is enabled by default (checked) and can be used to enable or disable the response from the selected data source.

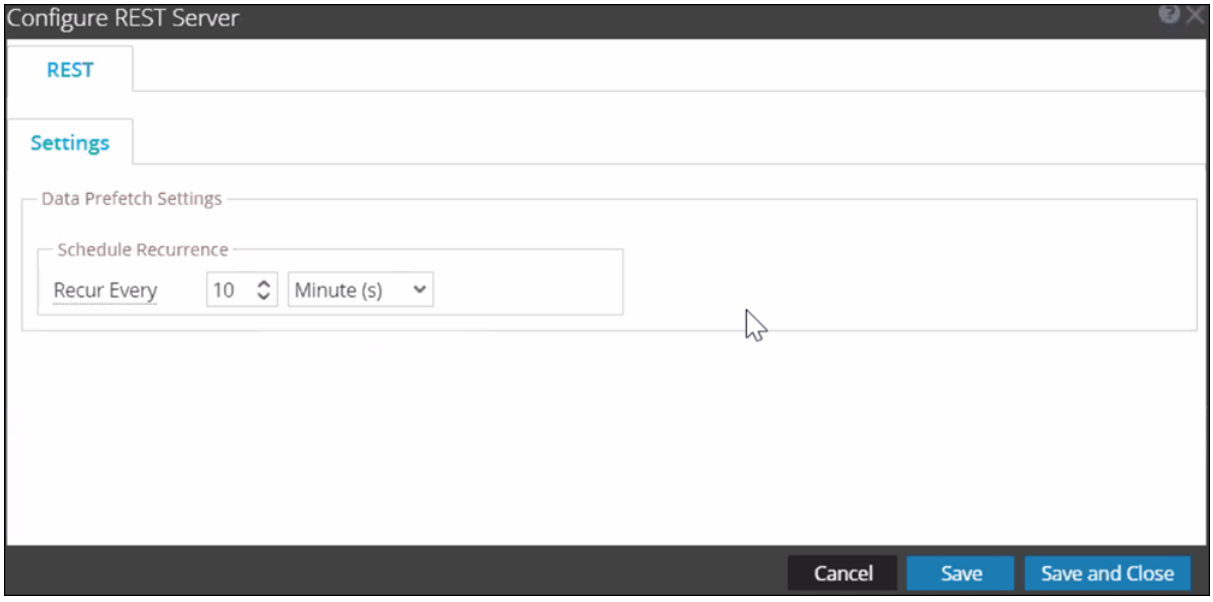
Field	Description
Cache Settings	Any lookup from Context Hub can be stored in the Context Hub cache for a configured time. Response to any subsequent matching request will be fetched from the Context Hub cache. Use this section to define the following cache settings for query lookup: • **Cache Enabled**: By default, this checkbox is selected and the query response is cached. • **Cache Expiration (Minutes)**: The maximum time the query lookup is retained in cache. The default time is 30 minutes and maximum is 7200 minutes that you can configure.
List value Expiration	**Enable**: Select Enable to define the number of days the list values must be available. By default, this option is disabled and the values are retained. **Time to Live (Days)**: Enter the number of days you want to the list values to be retained.
Meta Mapping	Any list stored in Context Hub should be made available for a lookup. The lookup in Context Hub is performed based on meta type or entities. Examples IP, HOST, MAC ADDRESS, DOMAIN, FILE_NAME, FILE_HASH, USER. **Meta Type**: Entities available in Context Hub. **Context Hub Fields**: Column headers from CSV file you have added when creating a list.
Meta and Field Mapping	**Response Preview**: You can view the live response for the REST API configured, using the test meta value as a placeholder. **Copy**: Copies the response preview. **Meta Mapping**: Any REST API stored in Context Hub should be made available for a lookup. The lookup in Context Hub is performed based on meta type or entities selected. Examples IP, HOST, MAC ADDRESS, DOMAIN, FILE_NAME, FILE_HASH, USER. **Meta Type**: Entities available in Context Hub. **Field Mapping**: Add friendly display name for the response field which is displayed during the context lookup.  - Add a field mapping. - Delete a field mapping. **Field**- Enter the path for which you want to add a friendly name. **Value (from Preview)**- The value is filled automatically based on the path. **Display Name**- Enter friendly name.
Minimum IIOC Score	The minimum IIOC score to be considered for fetching contextual information of NetWitness Endpoint modules.

Field	Description
Query Last (Days)	The duration (in days) for which the Context Data must be queried.
Limit	The maximum number of records to be displayed when Context Lookup is performed.
Recur Every	Configure recurring schedule to fetch and store contextual data for the required intervals.

5. Click any one of the following options:
   - **Cancel** - select this option to cancel the changes.
   - **Save** - select this option to save the changes.
   - **Save and Close** - select this option to save and close the dialog.

To access and configure settings for STIX data sources

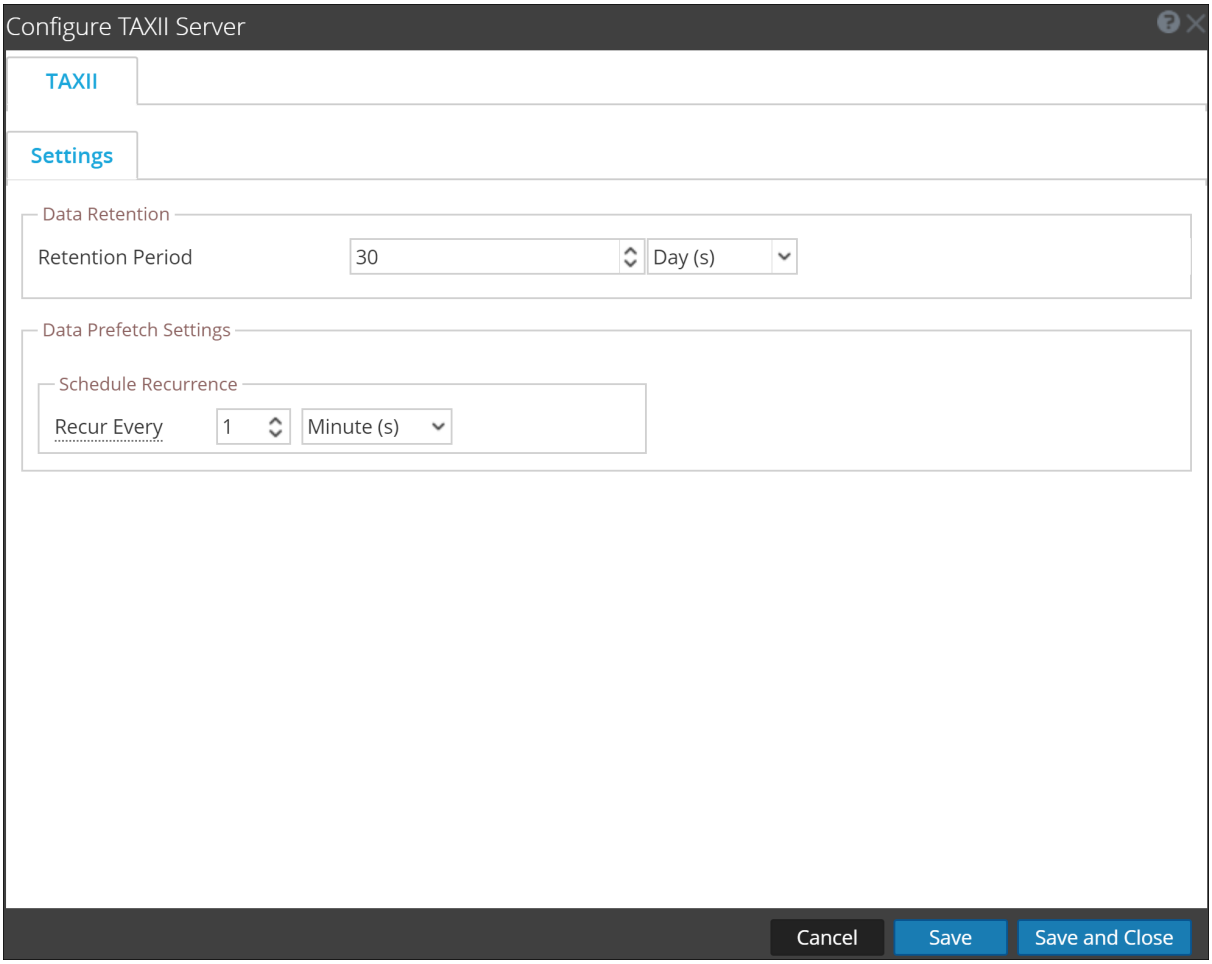
1. Go to **(Admin) > Services** and select Context Hub service.
2. Click > **View** > **Config** > **STIX** tab.
3. Select the data source for which you want to configure the settings and click in the Configurations column.
   a. For REST Server, the following fields are displayed.



Configure the following fields for REST Server.

Field	Description
Recur Every	Specify the time frequency in minutes, hours, days or weeks to retrieve the content from the data source.

b. For TAXII Server, the following fields are displayed.



Configure the following fields for TAXII Server.

Field	Description
Retention Period	Specify the number of days the content retrieved from the taxi data source must be retained before it is automatically deleted.
Recur Every	Specify the time frequency in minutes, hours, days or weeks to retrieve the content from the data source .

Based on the data source you select, the Response Groups differ. The following table describes the response groups for every data source.

Data Source (Connection)	Response Supported Groups	Field Settings
List	List	Meta Mapping Meta Type Context Hub Fields Settings Data Prefetch Settings Schedule Recurrence List Value Expiration Cache Settings Cache Enabled Cache Expiration (Minutes) [Min is 30 minutes Max is 7200 minutes]
Archer	Archer	Cache Settings Cache Enabled Cache Expiration (Minutes) Settings Export Attributes Configuration Export Attributes Data Prefetch Settings Schedule Recurrence
Active Directory	Users	Meta Mapping Meta Type Context Hub Fields Settings Data Prefetch Settings Schedule Recurrence List Value Expiration Cache Settings Cache Enabled Cache Expiration (Minutes)[Min is 30 minutes Max is 7200 minutes]

Data Source (Connection)	Response Supported Groups	Field Settings
Endpoint	IOC Machines Modules	Cache Settings Cache Enabled Cache Expiration (Minutes) Settings Context Panel Settings Cache Settings Cache Enabled Cache Expiration (Minutes) Settings Context Panel Settings Cache Settings Cache Enabled Cache Expiration (Minutes) Settings Minimum IIOC Score Context Panel Settings
Respond	Alerts Incidents	Context Panel Settings Data Prefetch Settings Query Last [Days Cache Settings Cache Enabled Cache Expiration (Minutes)
File Reputation Server	File Reputation Server	Cache Settings Cache Enabled Cache Expiration (Minutes)
TI	Displays information for STIX data sources.	IP address, email address, domain, filename, URL's, and file hash. **Note:** The context lookup for email address and URL will be displayed only if these metas are mapped. To map the metas, navigate to (Configure) > **System** > **Investigation** > **Context Lookup**.

Data Source (Connection)	Response Supported Groups	Field Settings
REST API	REST API	• Meta and Field Mapping • Meta Type • Field • Value (from Preview) • Display Name • Cache Settings • Cache Enabled • Cache Expiration (Minutes)

Note: After you configure the data source settings, you can configure the Context Hub configuration parameters by navigating to **(Admin) > Services**> **View** > **Explore** view. Make sure you restart the Context Hub service if you make any configuration changes in the Explore view.

Import or Export Lists for Context Hub

As an administrator you can import or export a list that is configured in the Context Hub service which can be used by an analyst. The file to be imported or exported is a CSV file and you can add multiple lists as Data Sources.

Prerequisites

Ensure that Context Hub is enabled and the service is available in  **(Admin) > Services** view of NetWitness.

Import a List

After you have imported a list, you can perform the following tasks:

- Import values to an existing list
- Add row to a list
- Edit a list name and description
- Edit a value from a list
- Delete a list
- Delete row from a list

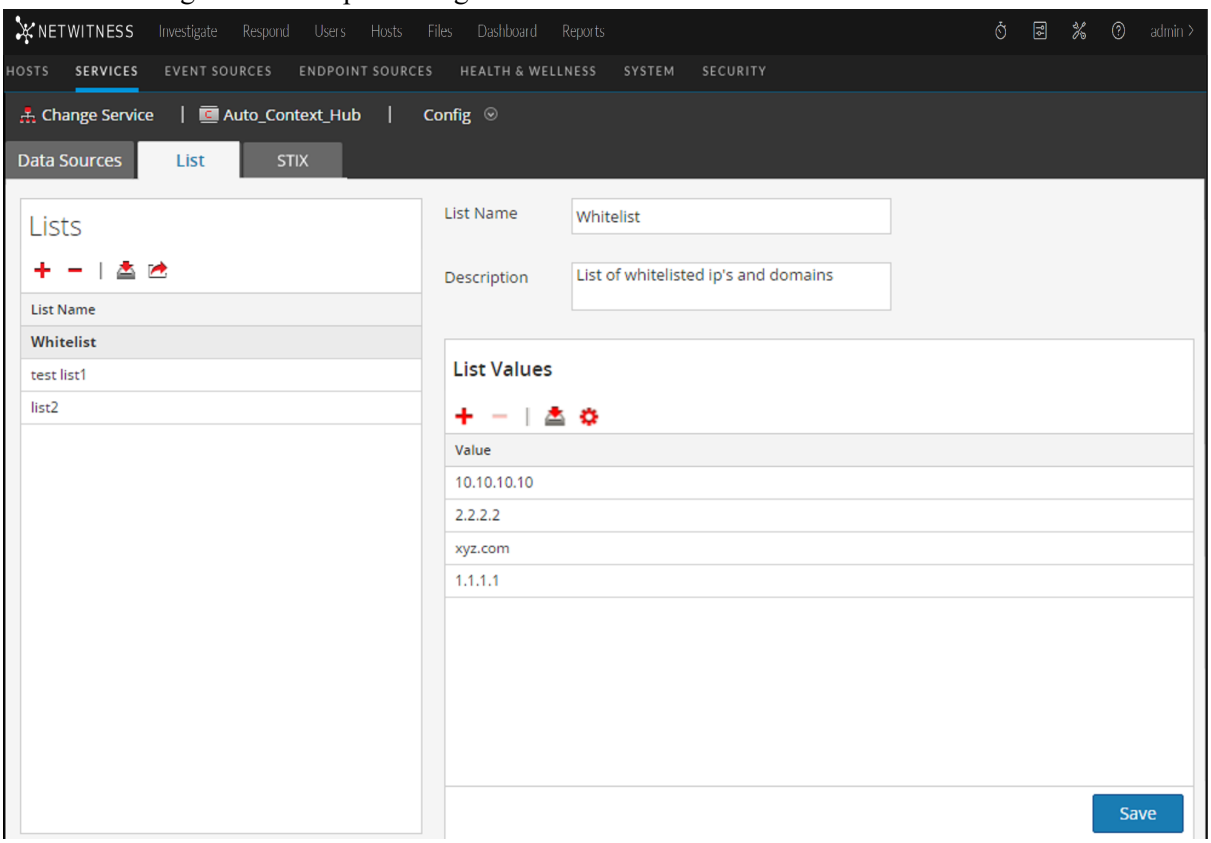
Note: You have to make the same changes to the relevant .CSV file, so that the changes get reflected the next time the schedule recurs. Otherwise, when you import values into an existing single-column or multi-column list, the data is overwritten from the source file when the schedule recurs. In case of a custom feed list, if the feed is edited or deleted, the corresponding Context Hub list also gets edited or deleted.

Import Single-Column List

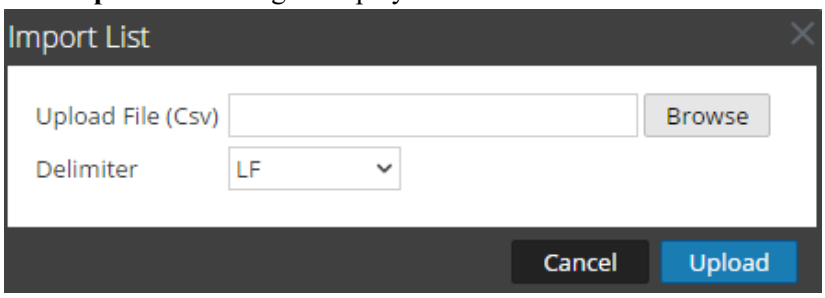
To import a list:

1. Select  **(Admin) > Services**.
   The services view is displayed.
2. In the **Services** panel, select the Context Hub service and click  > **View** > **Config**.
   The Services Config View of the Context Hub service is displayed.
3. Click the **Lists** tab.
   The Lists tab consists of the **Lists** panel and **List Values** panel.

The below image is an example of single-column list.



4. Click  on the **Lists** panel.
   The **Import List** dialog is displayed.



5. In the **Import List** dialog, complete the following steps:
   a. In the **Upload File (.CSV)** field, browse and select the `CSV` file.
   b. In the **Delimiter** field, select the delimiter to separate the values in a list from the options—**Comma**, **CR** (Carriage Return), and **LF** (Line Feed).
6. Click **Upload** to upload the CSV file to Context Hub.

These lists are considered as data sources for retrieving contextual information. But you can append to an existing multi-column list. The data will be appended only if the number of columns match.

Note: You cannot create a new multi column list by directly importing a CSV file. However, all the feeds that are converted into multi-colum lists will be displayed in the List tab. For information on how to import multi-column list, see Configure Lists as a Data Source

Import Values to an existing List

When you are importing into existing multi- column list the data is overwritten from the source file when the schedule recurs.

To import values to a list:

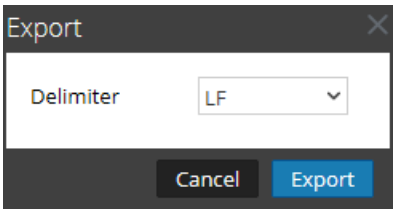
1. Go to  **(Admin) > Services**.
   The services view is displayed.
2. Service and click  > **View** > **Config**.
   The Services Config View of the Context Hub service is displayed.
3. Click the **Lists** tab.
   The Lists tab consists of the **Lists** panel and **List Values** panel.
4. In the Lists panel, select a list for which you want to import the values.
5. Click  on the **List Values** panel.
   The **Import List** dialog is displayed.
6. In the **Import List** dialog, complete the following steps:
   a. In the **Upload File (Csv)** field, browse and select the `CSV` file.
   b. In the **Delimiter** field, select the delimiter to separate the values in a list from the options—**Comma**, **CR**(Carriage Return), and **LF**(Line Feed).
7. Click **Upload** to upload the CSV file to NetWitness.

The list values are imported to the selected list. These lists are considered as data sources for retrieving contextual information. But you can append an existing multi column list. The data will be appended only if the number of column match.

Export List for Context Hub

To export a list:

1. On the **Lists** tab of the Services Config view of the Context Hub service, click  .
   The **Export** dialog is displayed.



2. In the **Delimiter** field, select the delimiter to separate the values in an exported list from the drop-down [**Comma**, **CR** (Carriage Return), and **LF** (Line Feed)].
3. Click **Export**.

In case of a single-column list, you can select the delimiter. And, in case of a multi-column list, the list is exported as CSV file to the local machine.

Note: When a custom feed is converted into a Context Hub list, you must map at least one meta key with one or more entity mapping for a column header with a meta. However, if you want to add or edit more entities you can do so by clicking ⚙.

Configure Meta Type Mapping for Context Hub

As an administrator you manage the mapping of Context Hub meta types with NetWitness meta keys.

The Context Hub service provides context lookup for meta values in the Respond and Investigation views. These meta values are grouped into meta types based on the category they belong to. For example, meta keys of NetWitness Respond and Investigation like `ip.src` and `ip.dst`are grouped into the meta type `IP` in Context Hub. The meta type `IP` is in turn mapped to metas like `alert.events.source.device.ip_address` and `alert.events.destination.device.ip_address` in the Respond database.

In the  (Admin) **> System > Investigation** view, the Context Lookup tab enables the administrator to configure the NetWitness meta keys and meta type mapping. The administrator can add or remove meta keys to the list of meta types supported by Context Hub.

The Context Hub service is pre-configured with default meta type and meta key mapping, which is expected to work with most deployments, unless there are some custom mappings created for your specific deployment.

Note: You cannot add a new Meta Type.

The default mapping is given below:

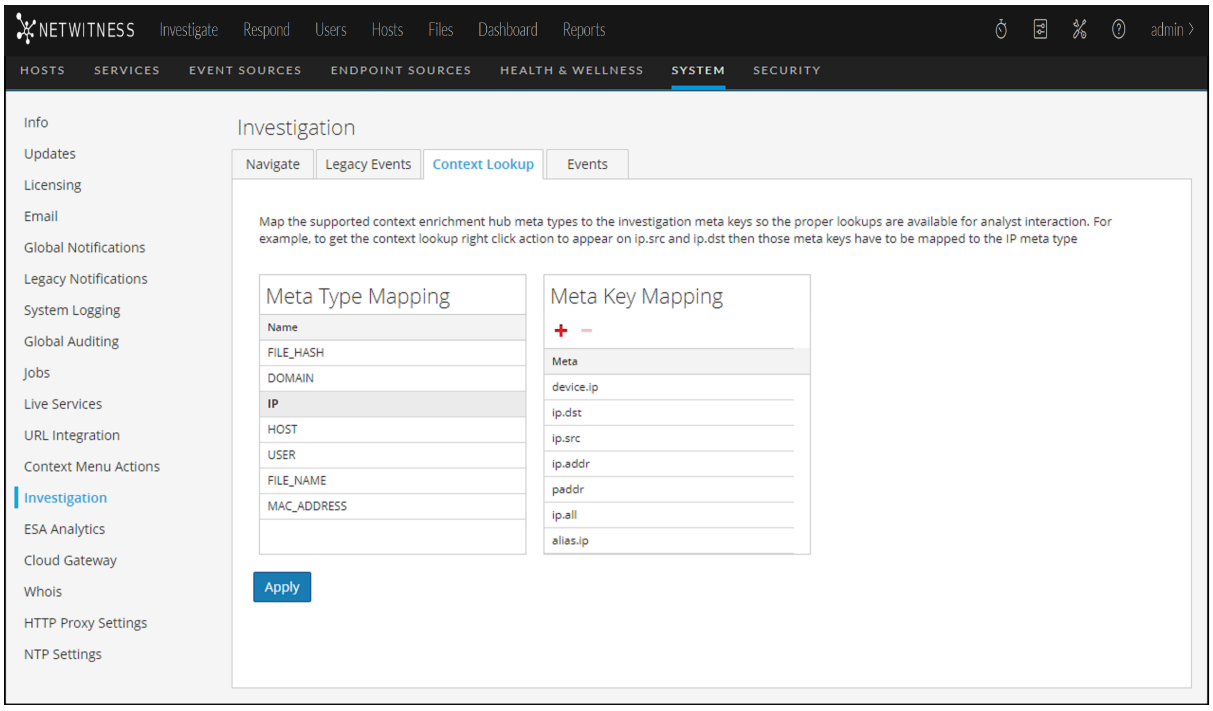
Meta Type Name	Meta Keys
IP	device.ip, ip.dst, ip.src, ip.addr, paddr, ip.all, alias.ip
USER	user.all, user.src, user.dst, username
DOMAIN	domain.src, domain.all, domain.dst
MAC_ADDRESS	eth.dst, eth.src, eth.all
FILE_NAME	filename, filename.all, filename.src, filename.dst, sourcefile
FILE_HASH	checksum.all, checksum, checksum.dst, checksum.src
HOST	device.host, alias.host, host.all

Procedure

To manage Investigation meta keys mapping:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Investigation**.

   The Investigation Configuration panel is displayed.
3. Select the **Context Lookup** tab.



4. Select a meta type to view the default meta keys that are mapped with this meta type.
5. To add a meta key, click + and enter the meta key.
6. To remove a meta key, select the meta key and click −.
7. To save the changes, click **Apply**.
8. In order to add a new meta, they need to be included in the Concentrator's custom index file. For example, if you want to add a meta "**fqdn**" then you need to add an new entry: **<key name="fqdn" description="Fully Qualified Domain Name="IndexValues" form-at="Text" valueMax="100" />** in the index file. For more information on how to include a new meta in the index file, see Index Customization topic in the *Core Database Tuning Guide*. After you add the new meta, you can view the contextual information on clicking the Pivot to investigate option in the Respond view.

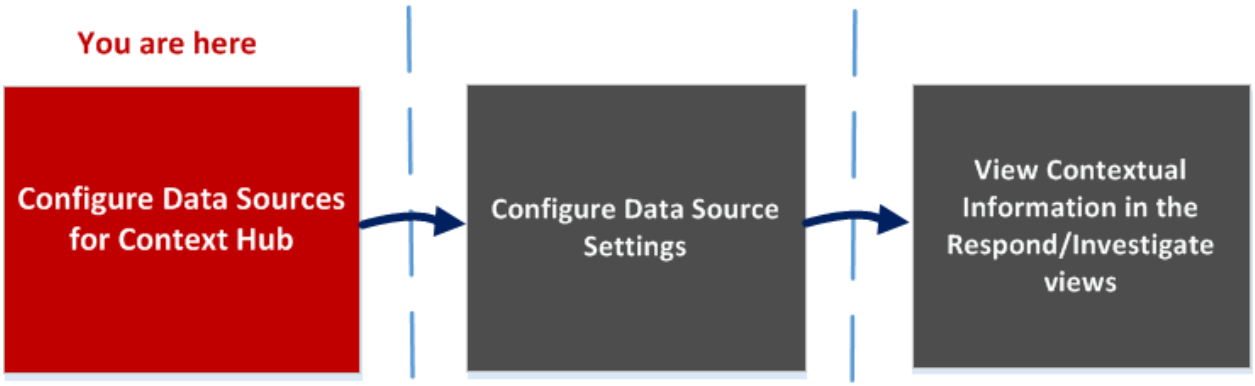
In case a new meta key is added, the Context Lookup menu option is enabled for the meta values under that meta key. For more information, see the "Investigation Configuration Panel" topic in the *System Configuration Guide*

Context Hub Data Sources Tab

In the **Data Sources** tab, you can configure one or more data sources for Context Hub service. Navigate to **(Admin) > Services** > Select Context Hub service > **View** > **Config** > **Data Sources** tab.

Workflow

This workflow shows the procedure to configure data sources for Context Hub service to view contextual information in the Respond / Investigate views.



- The first task is to add a data source
- The second task is to configure data sources settings to enhance your deployment. This task is optional as the settings for each data source is already configured with default values for optimal performance.
- And the third task is to view and analyze the contextual information in the Context Summary panel of the Respond or Investigate views.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure Data Sources for Context Hub*	Configure Lists as a Data Source Configure Archer as Data Source Configure Active Directory as a Data Source Configure NetWitness Endpoint as a Data Source Configure REST API as a Data Source Configure File Reputation Server as a Data Source Configure STIX as a Data Source
Administrator	Configure Hub Data Settings*	Configure Context Hub Data Source Settings

Role	I want to ...	Show me how
Analyst	View Contextual Information in Respond View	See the *NetWitness Respond User Guide*.
Analyst	Add, create and delete list from the Respond or Investigate View	See the *NetWitness Respond User Guide*. See the *Investigate User Guide*.
Analyst	Add or delete an entry from an existing list	See the *NetWitness Respond User Guide*.

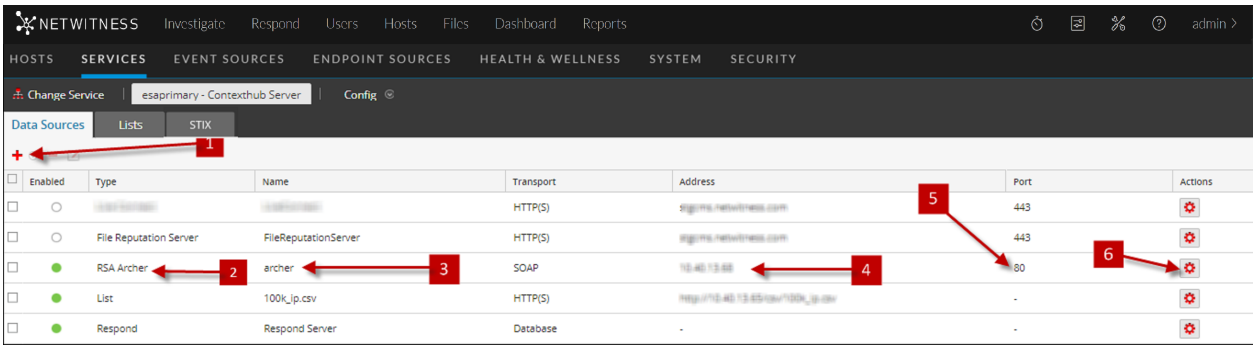
*You can complete this task here (that is in the Context Hub Data Sources Tab.)

Related Topics

- Configure Lists as a Data Source
- Configure Archer as Data Source
- Configure Active Directory as a Data Source
- Configure NetWitness Endpoint as a Data Source
- Configure Respond as a Data Source
- Configure REST API as a Data Source
- Configure File Reputation Server as a Data Source
- Configure STIX as a Data Source
- Configure REST API as a Data Source

Quick Look

The following example illustrates how to add a data source for Context Hub service.



1 Click + to display the **Add Data Source** dialog.
2 Displays the type of Data Source.
3 Name that identifies the Data Source.
4 The IP address or hostname of the data source.
5 The connection port for the data source.

6 Opens the **Configure Settings** dialog. You can view and edit the settings to be displayed on the Context Summary panel in the Respond or Investigate views.

7 Click **Test Connection** to verify that the host is connected to the Context Hub service.

Toolbar

The following table describes the toolbar actions.

Feature	Description
	Opens the Add Data Source dialog so that you can add a data source. You can add only one data source of each type. Except in case of Lists and Active Directory data sources which can be added in multiples. For detailed instructions to add a data source, see Configure Lists as a Data Source.
	Delete a data source. If you delete a data source, Context Hub does not consider the deleted service as a data source. All contextual information fetched previously will not be available.
	Opens the Edit Data Source dialog.
	Opens the Configure Settings dialog. You can view and edit the settings for the data sources. For description of each field in Configure Responses dialog, see Configure Context Hub Data Source Settings .

Data Source Configurations

The following table describes the listed configurations.

Feature	Description
Enabled	Indicates whether the data source is enabled or disabled. A solid colored green circle indicates that data source is enabled (). An blank white circle indicates that data source is disabled.
Type	The type of data source. For example, Lists, Archer, Active Directory, Endpoint, Respond, REST API or File Reputation server.
Name	The unique name to identify the data source. For example, Respond.
Address	The IP address or hostname of the data source.
Port	The connection port for the data source and vary based on the data source being added. For example, for Endpoint the port is 9443, for Lists the port is 80 and so on.

Context Hub Lists Tab

In the **Lists** tab, you can create and configure lists for Context Hub. Navigate to **(Admin) > Services** > Select Context Hub service > **View** > **Config** > **Lists** tab.

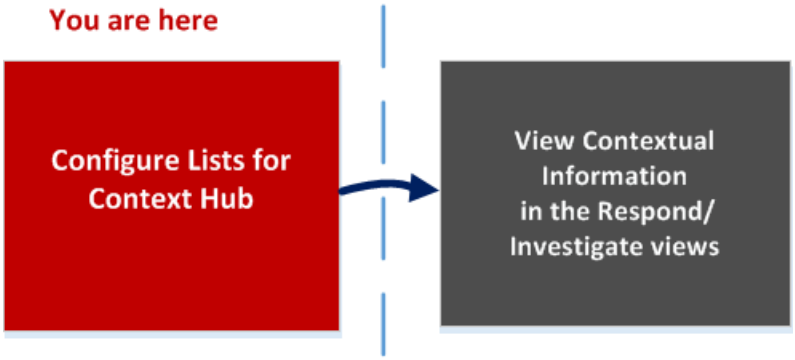
The Lists tab of the Context Hub service allows you to create one or more lists and add relevant list values to the list. These lists are automatically considered as data sources for the Context Hub service.

These lists may be populated with items either by importing external or custom feed CSV files or by adding meta values by using the option Add/Remove from List in Investigation and Respond views.

Note: You can also create lists and add list values from Respond and Investigation views. For more information, see the *NetWitness Respond User Guide* and the *NetWitness Investigate User Guide*.

Workflow

This workflow shows the procedure to configure lists for Context Hub service and to view contextual information in the Respond and Investigate views.



Creating one or more list is the first task in this workflow. The lists can contain supported metas such as an IP address, User, Host, Domain, MAC address, File Name or File Hash. The next task is to analyze or use the list data to view contextual data in Respond and Investigate views.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure List Data Source for Context Hub*	Configure Lists as a Data Source
Administrator/ Analyst	View Contextual Information in Respond View	See the *NetWitness Respond User Guide*.

Role	I want to ...	Show me how
Administrator/ Analyst	"Manage Lists and List Values in Investigation	See the *Investigate User Guide*.
Administrator/ Analyst	Create a list	See the *NetWitness Respond User Guide* and *Investigate User Guide*
Administrator/ Analyst	Update a list	See the *NetWitness Respond User Guide* and *Investigate User Guide*
Administrator/ Analyst	Delete list	See the *NetWitness Respond User Guide* and *Investigate User Guide*
Administrator/ Analyst	Import a list	Import or Export Lists for Context Hub
Administrator/ Analyst	Export list	Import or Export Lists for Context Hub

*You can complete this task here (that is in the Context Hub Lists Tab).

Related Topics

- Context Hub Data Sources Tab
- "Troubleshooting NetWitness Investigate" in the *NetWitness Investigate User Guide*

Quick Look

The following example illustrates how to add lists for Context Hub service.

The List tab consists of the **Lists** panel and **List Values** panel. The **Lists** panel has a toolbar with options to add, delete, import, and export lists. The entries under **List Name** are lists that are added or imported for the Context Hub service.

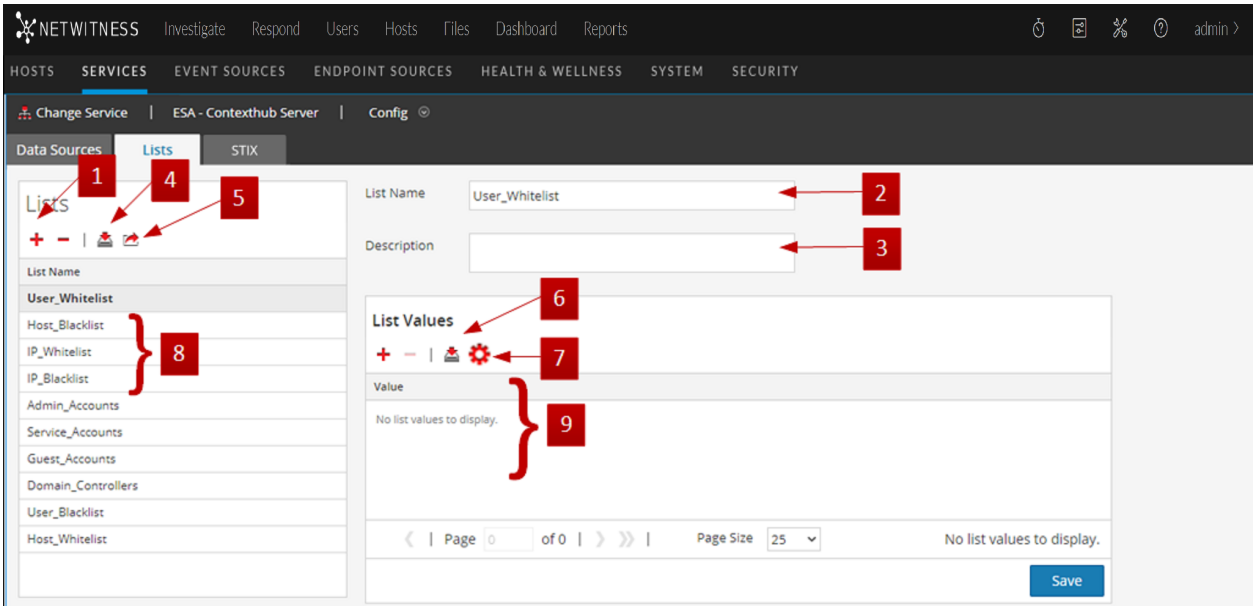
By default, 10 empty single-column lists are available in NetWitness11.1. These lists are empty and you need to add information to these lists. The out of the box 10 list names are used in ESA rules, for more information on ESA rules, see the *Alerting with ESA Correlation Rules User Guide*. For users upgrading from previous versions, they will be able to view these new lists in addition to their previously created lists. The lists available by default are:

- Admin_Accounts
- Guest_Accounts
- Service_Accounts
- User_Blacklist
- User_Whitelist
- Host_Whitelist

- Domain_Controllers
- IP_Blacklist
- IP_Whitelist
- Host_Blacklist

Note: If a list with the same name already exists prior to updating to or installing NetWitness12.0, then that list will be retained. Either rename that list before updating to 11.1 or update the contents in such a way that it can be used in ESA rules.

The lists are available in ESA rules tab in CONFIGURE > ESA Rules > Settings > Enrichment Sources. For more information on ESA rules, see the *Alerting with ESA Correlation Rules User Guide for Version 11.1*.
The **List Values** panel has a toolbar with options to add, delete, and import list values to the selected list. The entries under **Value** identify each list entry included in the list.



1	Click + to add a new list.
2	Name that identifies the list.
3	Description of the list.
4	Click to import list(s) to Context Hub.
5	Click to export a list to the local machine.
6	Click to import list values to selected list.
7	Click to add or edit entity mapping.

8	Displays the custom list(s) that are added to Context Hub.
9	Displays the list values that are added to the selected list.

Toolbar

The following table describes the toolbar actions.

Feature	Description
	Add a new list. For more information, see Configure Lists as a Data Source.
	Delete a list. If you delete a list from Context Hub, the list is no longer considered as a data source for retrieving contextual information.
	Import lists to Context Hub. For more information, see Import or Export Lists for Context Hub.
	Export a list to the local machine. For more information, see Import or Export Lists for Context Hub.

Note: You can select multiple lists at a time. Do one of the following:
1. Select a list, press and hold Ctrl key, and click the lists to be selected.
2. Select a list, press and hold Shift Key, and use arrow keys to select other lists.

List View Options

The following table describes the Lists configurations.

Feature	Description
List Name	Unique name to identify the list.
Description	Description of the list.
Save	Saves the changes made to the list.

Next steps

After completing the configuration, you can view the contextual data in the Context Summary Panel of the Respond view or Investigate view. For instructions, **Navigate to Context Summary Panel and View Additional Context** topic in the *Investigate User Guide*.

Context Hub STIX Tab

In the **STIX** tab, you can create and configure Structured Threat Information eXpression (STIX) data source for Context Hub. Navigate to  **(Admin) > Services**> Select Context Hub service > **View** > **Config** > **STIX** tab.

The STIX tab of the Context Hub service allows you to create one or more STIX, REST URLs, or TAXII data sources and edit them whenever required. When STIX is configured, Context Hub service automatically considers it as a data source.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure STIX Data Source for Context Hub*	Configure STIX as a Data Source
Administrator/ Analyst	View Contextual Information in Respond View	See the *NetWitness Respond User Guide*.

*You can complete this task here (that is in the Context Hub Lists Tab).

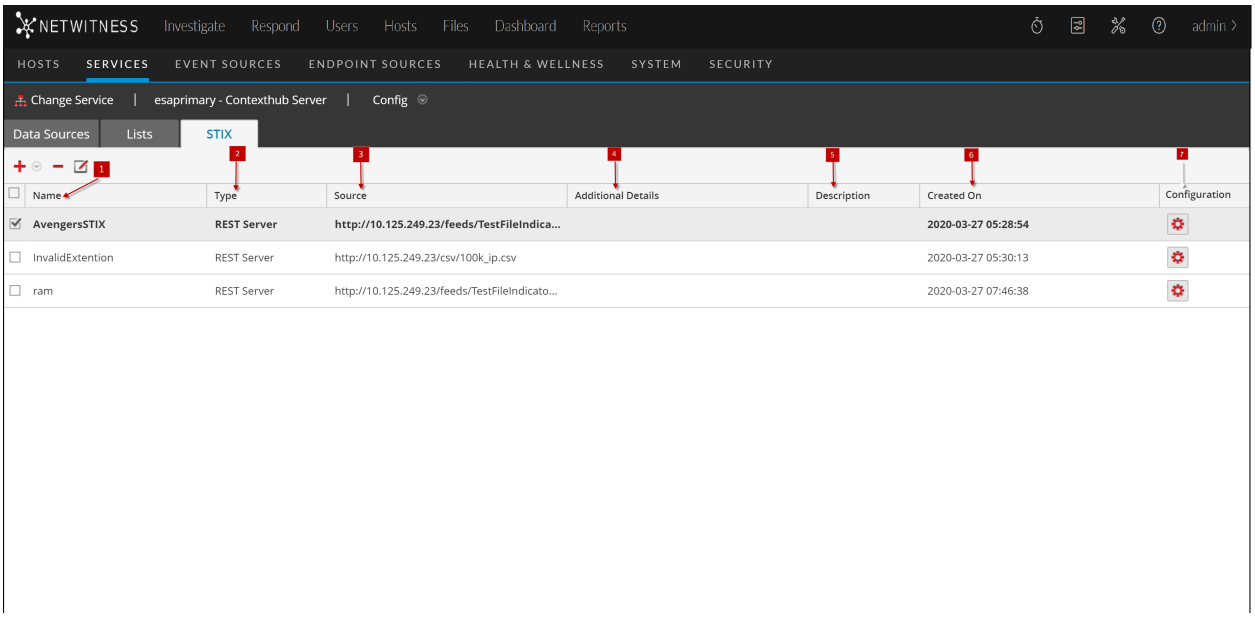
Related Topics

- Configure STIX as a Data Source

Quick Look

The following example illustrates how to add STIX to Context Hub service.

The STIX tab consists of add, delete and edit data sources options.



1	Name that identifies the added STIX source.
2	Type of data source - REST server, STIX or TAXII server.
3	The path of the source from which the STIX files are obtained.
4	Additional details related to the data source being added.
5	Description of the data source.
6	Date when the data source was created.
7	Click to edit the selected data source and configure advanced settings.

Toolbar

The following table describes the toolbar actions.

Feature	Description
+	Add a new data source such as File, REST Server, or TAXII Server. For more information, see Configure STIX as a Data Source .
−	Delete the selected data source.
✎	Edit the selected data source.

The following table describes the all the data source server configuration options.

Field	Description
Common Configuration Options	

Field	Description
Enabled	Select this checkbox to enable the configuration.
Name	Provide a name to the data source you want to add.
Description	Description of the data source.
Cancel	Click to revert the data source addition.
Validate	Click to verify the URL path to the Server.
Save	Click to save the configuration and add the required server as a data source.
REST Server Configuration Options	
URL	URL of the REST server.
Username (Optional)	Provide the username of the REST server if it needs to be authenticated.
Password (Optional)	Provide the password of REST server if it needs to be authenticated.
Trust All Certificates	Select this checkbox to trust all certificates.
Certificate File	Click browse to navigate to the location of the certificate file.
TAXII Server Configuration Options	
URL	URL of the TAXII server.
Username (Optional)	Provide the username of TAXII server if it needs to be authenticated.
Password (Optional)	Provide the password of REST server if it needs to be authenticated.
Client Certificate	Browse to upload a pkcs12 format client certificate available on your local system.
Certificate Password	Enter the password to the certificate, if it is password-protected.
UserProxy	Select this checkbox to enable proxy.
Trust All Certificates	Select this checkbox to trust all certificates.
Certificate File	Browse and select the certificate file.
TAXII Collections	Select the TAXII Collection name from the drop-down to automatically download the collection.
	Click to manually retrieve the list of available TAXII Servers, if the collections are not downloaded automatically.

Field	Description
STIX File Configuration Options	
File	Browse and select the STIX file.

Next steps

After completing the configuration, you can view the contextual data in the Context Summary Panel of the Respond view or Investigate view. For instructions, **Navigate to Context Summary Panel and View Additional Context** topic in the *Investigate User Guide*.

Troubleshooting

This topic provides information about possible issues that NetWitness users may encounter when setting up their ContextHub service in NetWitness.

Problem	Solutions
Prefetch for list fails if the list is created in append mode.	**Details:** The following error message is displayed in logs indicating that, entries in list exceeds the max allowed. `Error setting data source entries com.rsa.asoc.contexthub.exception.ContextHubException: total.entries.exceed.max` Also, Health & Wellness sets the stat - `Contexthub.Datasource.Health.Data-Sources-Health` to `Unhealthy` and displays the names of the lists for which prefetch has failed. For example, number of entries in the list are 50001 and number of records in the CSV file are 50001 (user did not change the csv since last prefetch.). Upper limit on number of entries in list is 100,000. Now on prefetch, Context Hub will try to append 50001 entries to the list but since 50001 + 50001 > 100,000, prefetch fails. **Solutions:** You should add only those entries in the .csv file which they wish to append to the existing .csv file. If, you do not want to append any entries to the list then perform one of the options, as applicable: • If you created the list with headers: remove all rows from the csv except the header. • If you created list without headers: you should have 0 rows in csv.
The Respond service is not able to send incidents to Archer with third party signed certificates.	As a workaround, you need to run a command to add a PEM certificate to the Respond trust store Run the following command on the Respond host: `security-cli-client --add-trusts -s respond-server -x <pem_certfilename> -u <username> -k <password>` Where: • `<pem_certfilename>` is the name of the certificate file. • `<username>` and `<password>` are your NetWitness administrator credentials.
SSL handshake with Archer certificate fails while adding it as a data source.	Use an archer generated certificate with the Trust All Certificates option configured.

Problem	Solutions
Pivot to Investigate option on the Respond page does not navigate to the correct link.	When you stop and restart the RabbitMQ server, the Pivot to Investigate option available on the respond screen is not visible. And the context panel for Pivot to Investigate reopens the same page. You need to restart the jetty service on the NetWitness Server, login to the NetWitness Server Host and enter the service jetty restart command.
When you import a list with missing quotes such as "172.16.0.0, the list is saved without any data to display.	This is because of the Apache bug (CSV-141), which does not parse CSV files with incorrect formats. To fix, import a list with correct quotes to avoid displaying an empty file. For example, “172.16.0.0”, “host.mycompany.com” and so on.
Increasing the limit settings for Alerts and Incidents leads to lookup error.	By default, the limit settings to view number of Alerts and Incidents is set to 50. If the limit is increased, the looked-up meta for alerts and incidents may lead to lookup error. This happens due to an internal database restriction. Make sure to keep the limit for viewing number of Alerts and Incidents to 50 or less.
Multiple incident and alert tabs are displayed while editing advanced configuration of Respond datasource.	You must delete duplicate entries in the Context Hub MongoDB and restart the MongoDB Context Hub Server. Perform the following: 1. Log in to control MongoDB on Admin Server. 2. Go to contexthub-server > ds_meta collection. **Note:** Make a note of '_id' value of duplicate entries for the incident and alerts documents that you want to delete. 3. In the ds_meta collection, delete 1 Incident and 1 Alert document. Once the duplicate entries are deleted only 1 Incident and 1 Alert type document will be available in the ds_meta collection. 4. Log in to application MongoDB on ESA Primary. 5. Go to contexthub-server, locate and delete the below collections: • ds_entries*_*<id-of-incident-doc-deleted-in-step 3> • ds_entries*_*<id-of-alert-doc-deleted-in-step 3> 6. Under bookmark_store collection, locate and delete documents with id same as _id of Incident and Alert from ds_meta collection. 7. Restart Context Hub Server.

NetWitness® Platform XDR

Version 12.0

Data Privacy Management Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

Data Privacy Overview **5**
- Data Obfuscation 5
- Data Retention Enforcement 6
- Audit Logging 6
- Components Covered by the Data Privacy Feature 7
- Data Privacy Feature Implementation by Component 7
- Component-Specific Configuration Guidelines 8

Recommended Configurations **10**
- Recommended Data Privacy Configuration 10
- Options for Data Retention Configurations 10
  - Data Storage With Data Retention Options in Effect 10
  - Option 1: No Original Data Saved to Disk, Only Hash Stored 12
  - Option 2: No Original or Obfuscated Values Stored: Not Recommended 13
- Optional Data Overwriting Options 13
  - Option 1: Limit Disk Space for Continuous Overwriting of Older Data 13
  - Option 2: Use Tiered Storage to Overwrite Data on a Scheduled Basis 14
  - Option 3: Purge Data Using String and Pattern Redaction Option 14
  - Limitations to Data Overwriting 15

Quick Start Procedures **16**
- Prepare to Configure Data Privacy 17
- Configure the Recommended Data Privacy Solution 19
  - Configure Metadata and Content Restrictions on Brokers, Concentrators, and Decoders 19
  - Add Data Privacy Officer and Analyst Accounts on the NetWitness Server 21
  - Configure Obfuscated Data on Decoders and Concentrators 22
  - Configure Data Retention on Concentrators and Decoders 23
  - Validate Data Privacy Protection 24

In-Depth Procedures **26**
- Configure Data Obfuscation 27
  - Configure the Decoder Hash Algorithm and Salt 27
  - Configure Language Keys 28
    - Example 1 29
    - Example 2 29
    - Example 3 30
  - Configure Metadata and Content Visibility Per User Role on Core Services 30
  - Configure Meta Keys Not Written to Disk Per Parser on a Decoder 34

Configure Data Retention 36
Data Retention 36
Deleting versus Retaining Log Data 36
Enable or Disable Cold Storage in a Log Storage Collection 37
Configure Log Retention and Storage on an Archiver 38
Schedule a Recurring Job to Check Data Retention Thresholds 38
Purge Data Using String and Pattern Redaction Option 40
Configure User Accounts for Use in Data Privacy 43
Customize the Default Administrators User Role at the Service Level 43
Add a User Account with the Aggregation User Role at the Service Level 43
Add Data Privacy Officer and Analyst Accounts on the NetWitness Server 44
Data Privacy References 46

Data Privacy Overview

This topic introduces the concept and implementation considerations for a data privacy officer or administrator who is managing exposure of privacy-sensitive data in NetWitness. In addition, information about recommended use cases is included.

Note: A data privacy plan touches on most components of NetWitness. The person who configures data privacy needs to understand NetWitness network components, configuration of NetWitness hosts and services as described in the *Host and Services Getting Started Guide*, and the types of information that need to be protected.

Regulatory mandates in some locations, for example the European Union (EU), require that information systems have a means of protecting privacy-sensitive data. Any data that could directly or indirectly identify "Who did what when?" may be considered privacy-sensitive data. A few examples are user names, email addresses, and host names. NetWitness provides a range of controls that customers can leverage to protect privacy-sensitive data. These controls can be used in a variety of combinations to protect privacy-sensitive data, without significantly reducing analytical capability.

A user role for a Data Privacy Officer (DPO) supports the management of privacy-sensitive data. The DPO can configure NetWitness to limit exposure of metadata and raw content (packets and logs) using a combination of techniques. The methods available to protect data in NetWitness include:

- Data Obfuscation
- Data Retention Enforcement
- Audit Logging

Data Obfuscation

NetWitness has configurable options for data obfuscation. Data privacy officers and administrators can specify which meta keys in their environment are privacy-sensitive and limit where the meta values and raw data for those keys are displayed in the NetWitness network. In place of the original values, NetWitness can provide obfuscated representations to enable investigation and analytics. In addition, DPOs and administrators can prevent persistence of privacy-sensitive meta values and raw logs or packets.

Three methods work together to implement data obfuscation:

- Obfuscation of meta values for privacy-sensitive meta keys with an optional salt. Meta keys configured as protected are represented by obfuscated values at the time of creation on a Decoder or Log Decoder; the obfuscated values are hashed and considered to be impossible to read. To implement, you must configure the Decoder and Log Decoder hash algorithm and salt, and configure privacy-sensitive language keys as protected on all Core services.
- Role-based access (RBAC) to the raw logs or packets and the privacy-sensitive meta values. The DPO can use roles with granular permission capabilities to restrict what an analyst versus a data privacy officer is able to view during configuration, analysis, and investigation. The *System Security and User Management Guide* provides in-depth coverage of the RBAC implementation in NetWitness. To implement, you must configure metadata and content visibility per role on individual Brokers, Concentrators, Decoders, Log Decoders, and Archivers.

- Preventing persistence of privacy-sensitive meta values and raw logs or packets. To implement, you must configure meta keys on parsers for individual Decoders and Log Decoders as transient.

Data Retention Enforcement

NetWitness can ensure that data is retained only as long as necessary or as specified. An administrator can configure data retention using age and time thresholds on a per-service basis. Schedulers running on each service automatically delete data meeting those thresholds. Once the data is deleted, it is no longer available through user interfaces, queries, or application programming interface (API) calls. Some of the NetWitness components also support purging of data through overwrites.

An administrator can manage data retention in several ways:

- Configure how long data persists in storage on the system.
- For Core services, strategically remove privacy-sensitive data that may have been written by configuring automatic removal of data of a specific age.
- Configure NetWitness so that original data is not sent or saved to the other components. If privacy-sensitive data makes its way into another database on the Reporting Engine, Malware Analysis, and NetWitness Servers, data retention can be managed there as well. This configuration for Event Stream Analysis is managed in the Services Explorer view.

Note: If a situation arises where the DPO decides that already collected data is privacy-sensitive after the system is functional, the administrator can manually overwrite the data from databases or files where the data is saved.

Audit Logging

Administrators can leverage audit logs that NetWitness creates using the Global Audit Logging feature. The audit logging feature generates audit log entries about many activities, and the following are examples of log entries that are relevant to data privacy:

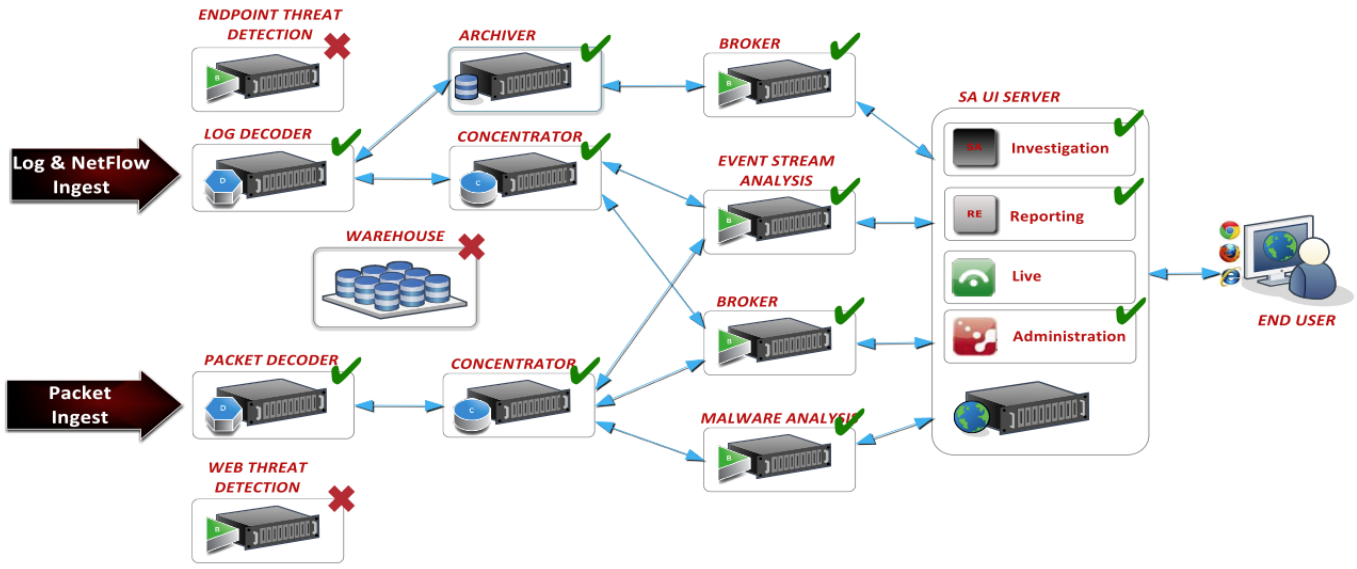
- Modifications to permissions and users assigned to roles.
- Failed and successful attempts to log on to NetWitness and log off.
- Data deletion.
- Data exports and downloads.
- Navigation by users to user interfaces and queries that users performed.
- Attempts (successful or not) to view or modify privacy-sensitive data, including an identification of the user who made the attempts.

All audit log entries are part of a standard audit trail for NetWitness. Administrators can configure NetWitness to forward audit logs to a designated destination, including third-party systems to provide additional filtering and reporting capabilities. For more information on Global Audit Logging, see "Configure Global Audit Logging" in the *System Configuration Guide*.

Components Covered by the Data Privacy Feature

The figure below identifies the NetWitness components covered by the Version 11.x data privacy feature with a green check mark. Components marked with an X are not supported by data privacy functions. The *NetWitness Getting Started Guide* provides a functional description of NetWitness components.

Note: NetWitness data privacy features are not supported for Warehouse and protected metadata can make it to Warehouse through Warehouse Connector, unless explicitly configured to be filtered out using Warehouse Connector Meta Filters. If protected metadata makes it to Warehouse, users having direct access to Warehouse can query such data. Data privacy officers need to prevent that through administrative, technical, and procedural controls outside of NetWitness.



Data Privacy Feature Implementation by Component

The following table identifies which data privacy features are supported for each NetWitness component. For each component, a check mark indicates if the component supports data obfuscation, data retention enforcement, data overwriting, and audit logging.

Component	Data Obfuscation	Data Retention Enforcement	Data Overwriting	Audit Logging
Ingestion				
Decoder	✓	✓	✓	✓
Log Decoder	✓	✓	✓	✓
Meta Aggregation				
Concentrator	✓	✓	✓	✓

Component	Data Obfuscation	Data Retention Enforcement	Data Overwriting	Audit Logging
Broker	n/a	✓ (stored in DPO cache only)[1]		✓
Real-Time Analysis				
Investigate	✓	✓ (stored in DPO cache only)[2]		✓
Event Stream Analysis	✓			✓
Malware Analysis	✓	✓		✓
Respond	✓	✓		✓
Reporting				
Reporting Engine	✓	✓		✓
Long-Term Analytics				
Archiver	✓	✓	✓(uncompressed)[3]	✓
Warehouse				

Notes:
1 - Brokers can cache data and this needs to be cleared by configuring an independent rollover and other removal of cache as required. The administrator can configure cache rollover for a Broker using the Scheduler in the Services Config view Files tab.
2 - Investigate and the NetWitness Server cache data, and this is cleared automatically every 24 hours.
3 - The overwriting procedure described in Configure Data Retention applies to uncompressed data.

Component-Specific Configuration Guidelines

NetWitness components and modules that obtain access to privacy-sensitive metadata and their obfuscated counterparts are Investigate, Event Stream Analysis (ESA), Malware Analysis, Respond, and Reports. They get access to data based on the permissions defined for the role to which the user belongs. The Administrator or DPO configures each Decoder or Log Decoder to identify meta keys that are flagged for obfuscation.

These components have additional guidelines to ensure that they function as expected with a data privacy scheme:

- **Event Stream Analysis**. When ESA receives privacy-sensitive data from NetWitness core, ESA passes on only the obfuscated version of the data. ESA does not store or show protected data. There are some additional guidelines for configuring advanced EPL rules and enrichment sources as well as information on how to remove sensitive data globally from all alerts (see "How ESA Handles Sensitive Data" in the *Alerting with ESA Correlation Rules User Guide*).

- **Malware Analysis**. Malware Analysis references certain meta keys during scoring, including `alias.host`, `client`, and others. To ensure no loss of analytical functionality Malware Analysis should be configured as a trusted client; that is, configured to connect to the NetWitness Core infrastructure with an account equivalent to a user in DPO role. Otherwise, if meta keys referenced by Malware Analysis do get tagged for obfuscation and are not accessible to Malware Analysis, some of the Indicators of Compromise (IOCs) may be rendered ineffective.
- **Respond Server service.** The Respond Server service uses a data privacy mapping file to display obfuscated data in alerts.(see "Obfuscate Private Data" in the *NetWitness Respond User Guide*) and has a configurable data retention period for alerts (see "Set a Retention Period for Alerts and Incidents" in the *NetWitness Respond User Guide*).
- **Reports.** In Reporting Engine, each Core service is added as two separate data sources, using the two separate service accounts; one data source has a service account representing the Data Privacy Officer role and the other data source has a service account representing a non-Data Privacy Officer role. "Configure Data Privacy for Reporting Engine" in the *Reporting Engine Configuration Guide* provides procedures to configure data privacy for Reporting Engine.

Recommended Configurations

This topic describes the recommended data privacy implementation for NetWitness and several additional use cases for managing exposure of privacy-sensitive data in NetWitness. Administrators can set up the NetWitness hosts and services to meet data privacy requirements for their environment. This section provides recommended configurations for data privacy and data retention.

Recommended Data Privacy Configuration

The recommended configuration to obtain the best analytical value with data obfuscation enabled is to define privacy-sensitive metadata and keep both original and obfuscated (hash) values of privacy-sensitive data on disk for Decoders, Log Decoders, Concentrators, and Brokers.

The assumption is that only a handful of metadata (approximately 10 meta keys) will be classified as protected and a FIPS 140-compliant algorithm for hashing will be used along with a salt to make reverse engineering the original value difficult. The recommended solution is SHA-256 with a salt length of at least 16 characters and up to 60 characters.

Note: By default, hash values are stored in binary format for faster response times and because it requires less storage space in the database when compared to saving them in string format. The recommended storage method is text/string.

Brokers and Investigate may have original and obfuscated data in cache due to data privacy officers using Investigate to confirm the original value to which the obfuscated value maps during investigations. Downstream services can also limit the use of the original sensitive values to in-memory processing so that data does not persist on disk in those downstream systems; this holds true for ESA and Malware Analysis.

The recommended solution to delete data when ready is the built-in and automatic data retention enforcement, which deletes data at a certain threshold. You can use this method for the following components: Decoder, Log Decoder, Log Collector, Archiver, Malware Analysis, NetWitness Respond, and Reporting Engine. You can manually configure Event Stream Analysis to support similar automatic data retention enforcement.

To manage cache storage, the NetWitness Server clears cache related to investigations of events every 24 hours. The Broker can also be configured to execute a periodic removal of locally stored cache.

Options for Data Retention Configurations

NetWitness provides alternative controls that the administrator can apply to enforce stronger restrictions on privacy-sensitive data storage when data obfuscation is enabled.

Data Storage With Data Retention Options in Effect

The following table summarizes where data is stored in the default configuration with no data privacy as well as for each data retention alternative. A checkmark indicates that privacy-sensitive data is saved on the component; a blank indicates that no privacy-sensitive data is stored on the component.

Component	Default Configuration	Data Storage Options		
	Original Data Stored	Original Data and Hash Stored (recommended)	Only Hash Stored	No Data Stored (all metadata is transient)
Ingestion				
Decoder	✓	✓		
Log Decoder	✓	✓		
Meta Aggregation				
Concentrator	✓	✓		
Broker	✓ (Cache only)	✓ (Cache only)		
Real-Time Analysis				
Investigate	✓	✓ (Cache only)		
Event Stream Analysis	✓			
Malware Analysis	✓			
Respond	✓			
Reporting				
Reporting Engine	✓	✓ (Optional)		
Long-Term Analytics				
Archiver	✓ (Optional)	✓ (Optional)		
Warehouse	✓ (Optional)	✓ (Optional)		
Content				
Live	n/a	n/a	n/a	n/a
Fraud Analysis				
RSA Fraud and Risk Intelligence Suite	n/a	n/a	n/a	n/a

Component	Default Configuration	Data Storage Options		
	Original Data Stored	Original Data and Hash Stored (recommended)	Only Hash Stored	No Data Stored (all metadata is transient)
End Point Protection				
NetWitness Endpoint	n/a	n/a	n/a	n/a

Notes:
Cache Only means that sensitive data is in the Broker or NetWitness Server cache. Configure Data Retention provides details about automated and manual clearing of cache.
Optional means that sensitive data storage does occur, but can be limited by optional configurations. For example, to limit where sensitive data is stored, do not enable DPO access for Reporting and do not aggregate original protected data into the Archiver.

Option 1: No Original Data Saved to Disk, Only Hash Stored

Administrators can eliminate the persistence of sensitive data to disk and store only an obfuscated value if the risk of exposure is too great. In this scenario, metadata generated during parsing on the Decoders and Log Decoders is used only in memory and not written to disk. Administrators can configure individual meta keys on a Decoder or Log Decoder as transient to ensure that sensitive metadata is not written to disk. Downstream services do not see original values and must use obfuscated values to conduct investigation and analytics.

To configure this data privacy scheme, data obfuscation must be enabled with hash values configured. You can configure individual meta keys on a Decoder or Log Decoder as transient to ensure that original values are not written to disk.

- Original values identified as sensitive are extracted from the raw data during parsing on the Decoder and Log Decoder and are accessible to the system during parsing (parsers, rules, feeds).
- The Decoder does not save the original values for meta keys identified as sensitive, storing only the hash of original values along with other non-sensitive metadata related to the event.

A side effect of these options is some loss in analytical capability, but you can configure these to suit the needs of your environment.

- By configuring all sensitive data as Transient, sensitive values are not persisted to disk, and the analytic capabilities using the original value are available at parse time only (parsers, rules, feeds).
- Event stream analysis (ESA) and Malware Analysis systems must rely only on the obfuscated meta values when doing their correlation and scoring respectively.
- Reporting Engine is limited to pulling reports using the non-sensitive and obfuscated values.

- The data privacy officer cannot view the original value, but can use the configured hash and salt to determine if an obfuscated value represents a specific known original value.

Option 2: No Original or Obfuscated Values Stored: Not Recommended

Administrators can eliminate the persistence of the original value to disk entirely if the risk of exposure is too great. As in Option 1, in this scenario, metadata generated during parsing on the Decoders and Log Decoders is used only in memory and not written to disk. Administrators can configure individual meta keys on a Decoder or Log Decoder as transient to ensure that sensitive metadata is not written to disk. Downstream services do not see original values and have no obfuscated values to conduct investigation and analytics.

To configure this data privacy scheme, configure individual meta keys on a Decoder or Log Decoder as transient to ensure that original values are not written to disk.

- Original values identified as sensitive are extracted from the raw data during parsing on the Decoder and Log Decoder and are accessible to the system during parsing (parsers, rules, feeds).
- The Decoder does not save not save the original values for meta keys identified as sensitive, storing only non-sensitive metadata related to the event.

A side effect of these options is significant loss in analytical capability, but you can configure these to suit the needs of your environment.

- By configuring all sensitive data as Transient, sensitive values are not persisted to disk, and the analytic capabilities using the original value are available at parse time only (parsers, rules, feeds). See Configure Data Retention.
- All downstream components have no visibility in the original values, obfuscated or otherwise.
- The data privacy officer has no visibility into the original value obfuscated or otherwise.

Optional Data Overwriting Options

Several options for overwriting data are available, and you should thoroughly understand each one before implementing data overwriting.

Option 1: Limit Disk Space for Continuous Overwriting of Older Data

If the desired data retention period to store the data, and therefore the amount of storage required for that data, is known, the size of the underlying hardware or the partition can be limited to that size. By reducing the hard drive storage or the partition size, the amount of free space available that has to be filled before new data overwrites it would also be limited. The newly ingested data continually overwrites the older data. Either solution must be done at deployment time to be effective.

Side effects of this option are:

- The removal of some disks will limit the number of resources available to distribute the I/O, causing some degradation in performance.

- The smaller partition size may cause some degradation in performance, but would alleviate some of the performance impact of removing disks.

Option 2: Use Tiered Storage to Overwrite Data on a Scheduled Basis

If overwriting of data is required on a scheduled automatic basis, you can configure the Decoders and Concentrators to use tiered storage. The tiered storage configuration provides a mechanism for invoking a script after a database file has been removed from the application but prior to its removal from the file system. If necessary, instead of moving the file to the second tier, or cold storage, (the intended function in a tiered storage use case), the script can use a utility like the CentOS `shred` utility to overwrite the file. This tool is less effective when the database is stored in a journaling file system like XFS, in which the Core database resides, and on a RAID logical drive like the ones with which the Core hosts connect.

Most other NetWitness components do not have this option; their data is stored in a database that does not support the tiered storage mechanism. The only other component that could use this overwrite method is the Reporting Engine since it saves reports and alerts as individual files. However, the Reporting Engine charts are stored in a database so they would be immune to this technique.

Option 3: Purge Data Using String and Pattern Redaction Option

Data purging provides a mechanism to strategically overwrite a specific subset of data from the system in case any sensitive data has been persisted either on purpose or by accident. The NetWitness `wipe` utility allows for unique patterns to be written over the data in the meta and packet databases for Core services, which may contain RAW packets or logs for existing sessions, based on a session identifier. All Core components have the capability to overwrite a subset of data that has been found by executing a query string, including regex patterns. The session identifiers resulting from the query are fed into the NetWitness `wipe` utility.

Note: This option is not available if the data in the Core database has been compressed (as typically done in Archiver deployments).

In most NetWitness components the database in use does not provide a built-in redaction or secure deletion mechanism. The Malware Analysis component can overwrite the data object in the database with the value `private` instead of deleting it during the data retention management process, but this is not meant to be a secure deletion mechanism.

Caution: Using this method on a large number of sessions has two drawbacks: it can be time-consuming and impact performance.

See Purge Data Using String and Pattern Redaction Option for procedure.

Limitations to Data Overwriting

There are limitations to the overwriting techniques described as Option 2 and 3. To perform the overwrite of the data in the disk sectors, the above options for overwriting and the overwrite command line tool provided as an alternative method (`shred`, a function of CentOS) make assumptions about the disk layout. NetWitness hosts use SSD drives and RAID configurations for performance and reliability reasons, and these inhibit the functionality of the overwrite techniques. If overwrite techniques alter SSD drives and RAID configurations in an attempt to increase security, there will inevitably be an associated performance cost reflected in ingest rates, query speeds, and potentially other areas. The command line tools available for overwrite are recommended only for special use cases when it is necessary to redact specific data. The tools are not for use in a real-time continuous method because of the potential performance cost that will be incurred.

Quick Start Procedures

This section provides end-to-end instructions for preparing to configure data privacy features, then completing the configuration of the recommended data privacy solution.

- Prepare to Configure Data Privacy
- Configure the Recommended Data Privacy Solution

Prepare to Configure Data Privacy

This topic provides general guidelines for planning and configuring data privacy policies in the NetWitness network. Before beginning configuration, you must understand the data that needs to be protected on your network and develop a plan. You will need to:

1. Identify the meta keys that hold privacy-sensitive data and need to be protected. This decision is based on requirements specific to your site.
2. Decide which users need access to privacy-sensitive metadata and raw content. The first decision is whether to separate the DPO and administrator roles for your site by configuring a custom administrators system role on Decoder and Log Decoders and removing the `dpo.manage` permission. By default, administrators have all permissions including the ability to configure the salted hash transform used to obfuscate data; some sites may want to reserve this access for data privacy officers. "Service User Roles and Permissions" in the *Hosts and Services Getting Started Guide* has more details on exactly what permissions each role has and the purpose of the permissions.
3. Plan the configuration changes you need to make in your NetWitness deployment to support adequate data privacy.
4. Assess how your configuration may impact out-of-the-box and custom content. For example, by default content available through Live for Reporting Engine is not geared toward obfuscated meta values.

In a single deployment, certain data-privacy configurations in the Core services must be the same. The following table lists these settings and uses a checkmark to identify the services for which the configuration must be the same. In the heading, the following abbreviations are used: D = Decoder, LD = Log Decoder, A = Archiver, C = Concentrator, and B = Broker.

	Configuration must be the same for these services:				
Setting	**D**	**LD**	**A**	**C**	**B**
Hash algorithm and salt for privacy-sensitive data	✓	✓			
Language key data privacy attributes in the custom index file (includes configuring keys as protected)	✓	✓	✓	✓	✓
Transient meta keys (not persisted on disk) per service and parser	✓	✓			
Meta data and raw content visibility per system user group. (The meta keys must exist in the custom index file.)	✓	✓	✓	✓	✓
User who has the `Aggregation` service user role assigned is added.*	✓	✓	✓		

* When trying to access data on an aggregate service, the Log Collector or Broker requests authentication. When prompted to enter user name and password, you must authenticate as a user who is assigned the `Aggregation` service role. "Services Security View - Aggregation Role" in the *Hosts and Services Getting Started Guide* provides detailed information about this role. Follow the instructions in "Add, Replicate or Delete a Service User" in the *Hosts and Services Getting Started Guide* to create a user and assign the new user the Aggregation service user role.

Configure the Recommended Data Privacy Solution

This topic tells administrators and data privacy officers how to configure the recommended data privacy solution in a NetWitness network. These are the basic steps to follow to configure the NetWitness system to identify sensitive data and determine who can see the sensitive data. The recommended configuration generates obfuscated values of certain original meta keys and then persists both the original and obfuscated data so that it is available to users assigned privileged role access.

This configuration has several parts:

1. Create two users with different levels of permissions. One user (the data privacy officer) can view all metadata and another user (an analyst) is restricted from seeing certain metadata and content with associated metadata.
2. Set up two transforms using a salt and hash to create an obfuscated version of original `username` and `ip.src` meta keys.
3. Configure data retention on the Decoder and Concentrator services.

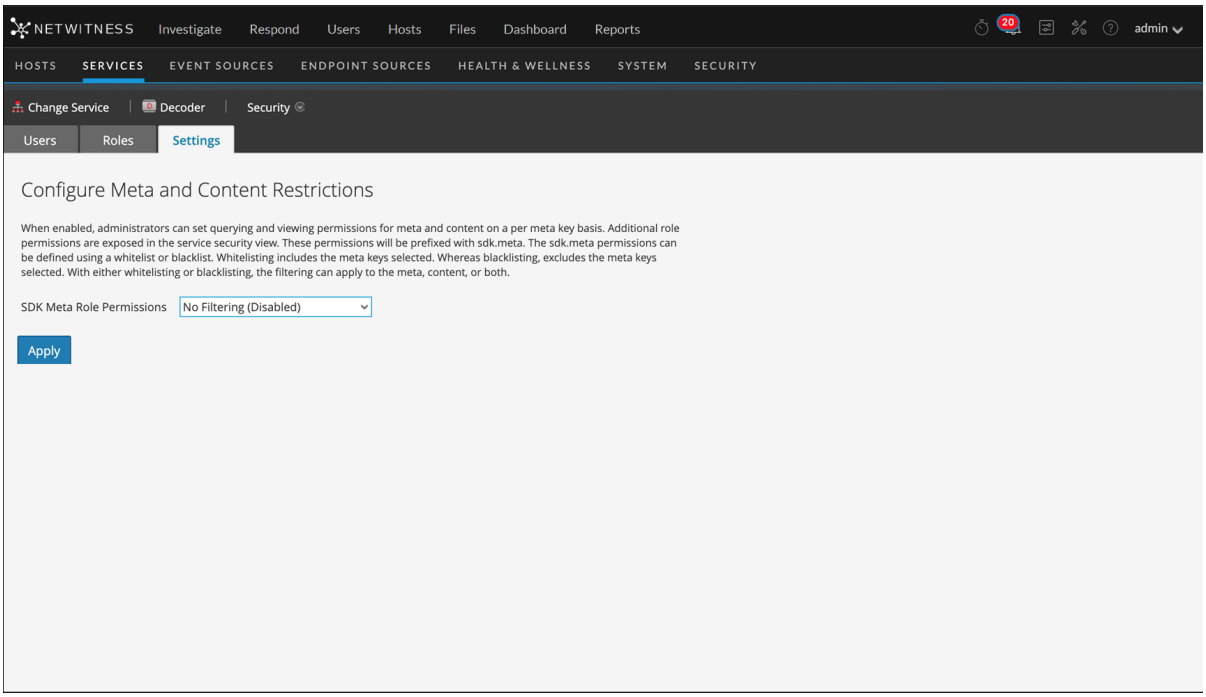
Note: The following conditions are required in order to complete this procedure:
-The Concentrator and Decoder must be added to the NetWitness Server using trusted connections.
-The NW Server version must be 11.x or later.
-The Core services must be 11.x or later.
-Aggregation must use Aggregators accounts on all Core services.

Configure Metadata and Content Restrictions on Brokers, Concentrators, and Decoders

Note: In version 11.6, if these 12 metas namely `'sessionid','nwe.callback_id', 'medium','session.split','ip.dst','ip.src','ipv6.src','ipv6.dst','tcp.dstport','tcp.srcport','udp.dstport', and 'udp.srcport'` are restricted then the group events option will be disabled.

To restrict the metadata and raw content that users can view, you must enable SDK system roles to allow more granular controls by configuring metadata and content restrictions on each service in the Services Security view.

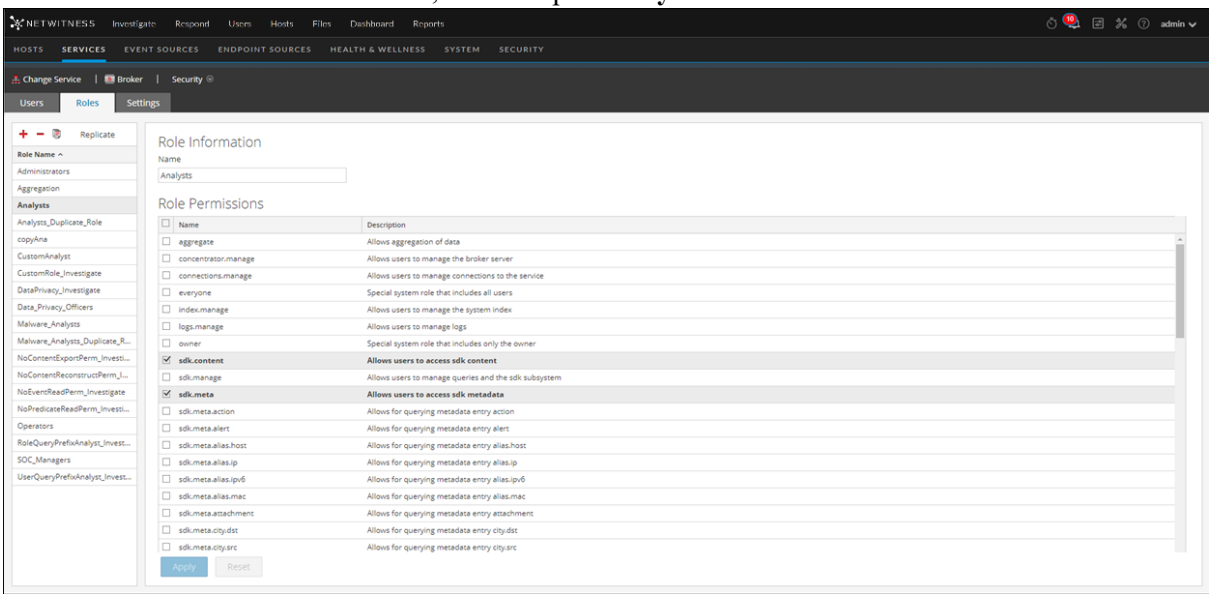
1. Go to  **(Admin)** > **Services**, select a service and then select  **> View > Security**.
2. Click the **Settings** tab.



3. In the **SDK Meta Role Permissions** field, select **Blacklist meta and content**. Click **Apply**.

   This allows the administrator to blacklist individual meta keys so that only the data privacy officer can see the meta keys and content. New roles per meta key are added to the Roles tab.

4. Click the **Roles** tab and select a role, for example **Analysts**.



5. In the **Roles** tab,

   a. Select the meta keys that you do not want analysts to see, for example, select `sdk.meta.username` and `sdk.meta.ip.src`.
   This restricts the analyst from seeing the privacy-sensitive meta keys `username` and `ip.src` as well as any content for any session that contains that metadata within it.

b. Ensure that `sdk.packets` is selected.
   If it is de-selected, analysts lose the ability to bulk export raw packets and logs. In NetWitness Platform 11.0 and later, RBAC just works for packets. Sessions that are restricted are just skipped during pcap generation in Investigate. Sessions that are allowed have packets returned. For more information on RBAC, see the *System Security and User Management Guide*.

c. Click **Apply**.

6. In the Roles tab, ensure that the `Data_Privacy_Officers` role has no `sdk.meta.values` selected. Click **Apply**.

   A DPO can view any metadata and any session.

   In the Roles tab, ensure that the `Aggregation` role has the following permissions: select `aggregate`, `sdk.content`, `sdk.meta`, and `sdk.packets`.

Add Data Privacy Officer and Analyst Accounts on the NetWitness Server

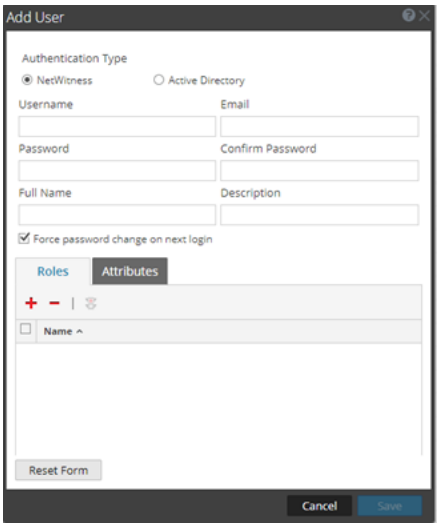
You must add two new user accounts in NetWitness at the system level to depict a privileged data privacy officer and a typical analyst. If the environment is configured using the default trusted connections, you do not need to create the new user accounts on the Core services (Brokers, Concentrators, and Decoders). When a user is created in the NetWitness Server, that user can log on to the services.

Note: The role name is required to exist on both the server and the services, and the role name for all must be identical. If you create a new custom role on the NetWitness Server, make sure to add it to all Core services as well.

1. Create a new user account for the data privacy officer:

   a. Go to **(Admin)** > **Security**, select the **Users** tab. In the **Users** tab toolbar, click .

   The Add User dialog is displayed.



   b. Create the new account with the following credentials.

Username = <new user name for logon, for example, DPOadmin>
Email = <new user's email, for example, DPOadmin@rsa.com>
Password = <new user's password for logging on, for example, RSAprivacy1!@>
Full Name = <new user's full name, for example, DPO Administrator>

c. Click the Roles tab, , and select the `Data_Privacy_Officers` role for the new user.

d. Select **Save**.

2. Create a new user account for the analyst with limited privileges:

a. In the **Services Security** view, select the **Users** tab. In the **Users** tab toolbar, click .

The Add User dialog is displayed.

b. Create the new account with the following credentials:

Username = <new user name for logon, for example, NonprivAnalyst>
Email = <new user's email, for example, NonprivAnalyst@rsa.com>
Password = <new user's password for logging on, for example, RSAprivacy2!@>
Full Name = <new user's full name, for example, Nonprivileged Analyst>

c. Click the Roles tab, , and select the `Analysts` role for the new user.

d. Select **Save**.

Configure Obfuscated Data on Decoders and Concentrators

This procedure creates the obfuscated values to provide to users who do not have access to the original values.

1. Configure a salt so that the obfuscated value becomes unique. Different companies may have analysts of the same first name and potentially the same login username, and using a salt limits the possibility of someone outside your organization determining your obfuscation mechanism. In this example, you use a simple salt and SHA-256, but the salt is configurable and the hash algorithm can be changed. For additional information, see Configure Data Obfuscation.

a. To define the salt and hash algorithm, go to  **(Admin) > Services**.

b. Select a Decoder in the **Admin Services** view and select  > **View > Config**.

c. Click the **Data Privacy** tab, and select hash algorithm (SHA-256). In the Salt field, type a hash, for example, **rsasecurity**, and click **Apply**.

2. Define the transforms, including the hash format, between the original meta key and obfuscated meta key on the Decoder. The default hash format is binary, but the recommended configuration calls for using the text/string format.

a. While still in the Services Config view, click the **Files** tab, and in the drop-down menu select **index-decoder-custom.xml.** (You can apply this same configuration to the Log Decoder in the `index-logdecoder-custom.xml` file.)

b. Enter the following lines in the available input area:

```
<?xml version="1.0" encoding="utf-8"?>
<language level="IndexNone" defaultAction="Auto">
<key name="username" description="Username" format="Text"
protected="true"><transform destination="username.hash"/></key>
<key name="username.hash" description="Username Hash"
format="Text"/>
<key name="ip.src" description="Source IP Address" format="IPv4"
protected="true"><transform destination="ip.src.hash"/></key>
<key name="ip.src.hash" description="Source IP Address Hash"
format="Text"/>
</language>
```

c. To restart the Decoder service, go to **(Admin)** > **Services**, find the Decoder service you want to restart, and select > **Restart**.
The service should automatically restart.

3. Define the meta keys on the Concentrator in the `index-concentrator-custom.xml` file:

a. Go to **(Admin)** > **Services**, select a Concentrator in the **Admin Services** view, and select > **View > Config**.

b. Click the **Files** tab, and in the drop-down menu select **index-concentrator-custom.xml**

c. Enter the following lines in the available input area:

```
<?xml version="1.0" encoding="utf-8"?>
<language level="IndexValues" defaultAction="Auto">
<key name="username" description="Username" format="Text"
level="IndexValues" protected="true"/>
<key name="username.hash" description="Username Hash"
format="Text" level="IndexValues" token="true"/>
<key name="ip.src" description="Source IP Address" format="IPv4"
level="IndexValues" protected="true"/>
<key name="ip.src.hash" description="Source IP Address Hash"
format="Text" level="IndexValues" token="true"/>
</language>
```

d. To restart the Concentrator service, go to **(Admin)** > **Services**, find the concentrator service you want to restart, and select > **Restart**.
The service should automatically restart.

Configure Data Retention on Concentrators and Decoders

Data retention configuration ensures that the data residing in the NetWitness Core components is deleted after a certain time. Configuring data retention on Concentrators and Decoders is not required for all environments, but it may be necessary to be in compliance with applicable laws and regulations. It is important to evaluate an appropriate retention period for your environment. The Data Retention Scheduler settings that you set apply to ALL data on a Concentrator or Decoder.

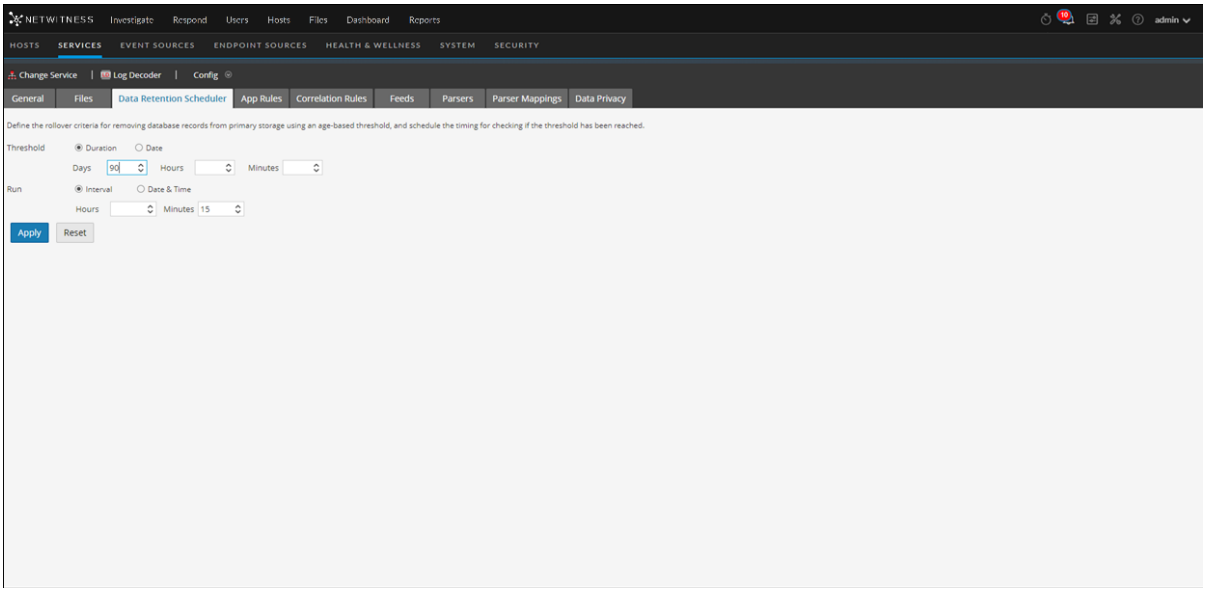
In the following example, NetWitness is configured to execute a check every 15 minutes to determine if the duration threshold has been met. If the threshold is met, NetWitness deletes data older than 90 days in the relevant databases.

Caution: The 90 day retention period is just an example. Adjust your rollover criteria depending on the location of the data and the applicable laws. In a strict data privacy environment, such as in Europe where laws require that Personally Identifiable Information (PII) not be saved or removed frequently, you may need to adjust the time.

This procedure is optional. If you do not set a time retention limit, the system automatically deletes the oldest data when the hard drive space is full.

(Optional) For each Concentrator and Decoder:

1. Navigate to the **Services Config** view **> Data Retention Scheduler** tab.



2. Define the data retention period. For example, set the **Threshold** to **Duration**, and type **90** in the **Days** field.
3. Define how often the scheduler checks to see if the threshold has been met. For example, set the runtime to **Interval** and select **15** in the **Minutes** field.
4. To save the configuration, click **Apply**.

Validate Data Privacy Protection

At this point, users have been added with roles that have permissions around specific types of metadata. The next step is to make sure the restricted user (the analyst) cannot view what the unrestricted user (the DPO) can. Also you need to ensure that the data retention configuration is limiting how long data is kept on the systems.

1. View role-based obfuscation in action:
   a. Log on as the unrestricted user (DPOadmin) and make sure this user can see all the data including the protected sensitive data `username` and `ip.src` along with any session that contains that metadata.
   b. Log off and the back on as the DPO user.

c. For each Decoder and Log Decoder, import a PCAP or logfile into the Services System view. Go to  **(Admin)** > Services, click a service, and use the **Upload Packet Capture File** option to upload a PCAP file that contains `username` and `ip.src` metadata.

d. When the import is complete, go to **Investigate > Navigate**, select the Concentrator connected to the Decoder to which the data was just imported.

e. Scroll down to make sure the `username` and `ip.src` meta keys and corresponding values are visible.

f. Click one of the green numbers next to a `username` or `ip.src` value and verify that the session loads in the Events view.

g. Make a note of the session ID to check when logging on as the restricted user.

h. Log off and log on as the restricted user (NonprivAnalyst).

i. Repeat steps c through f to verify that the user cannot see any `username` or `ip.src` metadata or sessions with that metadata including the one previously mentioned.

j. To jump to a specific session in the **Navigate** view, in the **Actions** menu, select **Go to Event** and enter the session ID.

2. Validate that the data retained in the database falls within the retention time configured in the Data Retention Scheduler.

a. Log off and log on as the unrestricted user (DPOadmin).

b. On the Concentrator, navigate to the **Services > Explore** view.

c. In the node tree, select the **database** node and then **stats**.

d. Observe the `meta.oldest.file.time` value and verify that this is not older then the threshold put on the data retention scheduler.

e. Change the service to the Decoder and repeat steps b through d, check for `stats meta.oldest.file.time` and `packet.oldest.file.time`.

In-Depth Procedures

This topic is a collection of procedures that a Data Privacy Officer uses to implement a data privacy plan for the NetWitness network. These procedures are part of an overall configuration, and are performed as needed to implement the data privacy plan and manage the flow of information in the network.

- Configure Data Obfuscation
- Configure Data Retention
- Configure User Accounts for Use in Data Privacy

Configure Data Obfuscation

This topic provides the procedures for configuring data obfuscation in NetWitness. In a single deployment, all Core service configurations for a data privacy solution must be the same; be sure to use the same hash and salt across all Decoders and Log Decoders.

Note: In order for data obfuscation to work, user accounts need to be configured as described in Configure User Accounts for Use in Data Privacy.

Configure the Decoder Hash Algorithm and Salt

Value hashing accomplished as part of the data privacy solution occurs at the time of meta key creation on the Decoder and Log Decoder. Both services have default settings for use with all meta keys whose values are transformed without a specified hash algorithm type or salt value. The initial NetWitness values for defaults are: hash algorithm (SHA-256) and salt (none).

Note: The SHA-1 algorithm and it is not available.

If you want to change the default settings, you can edit them in the Services Config view > Data Privacy tab or in the following nodes in the NetWitness Services Explorer view:

- `/decoder/parsers/transforms/default.type`

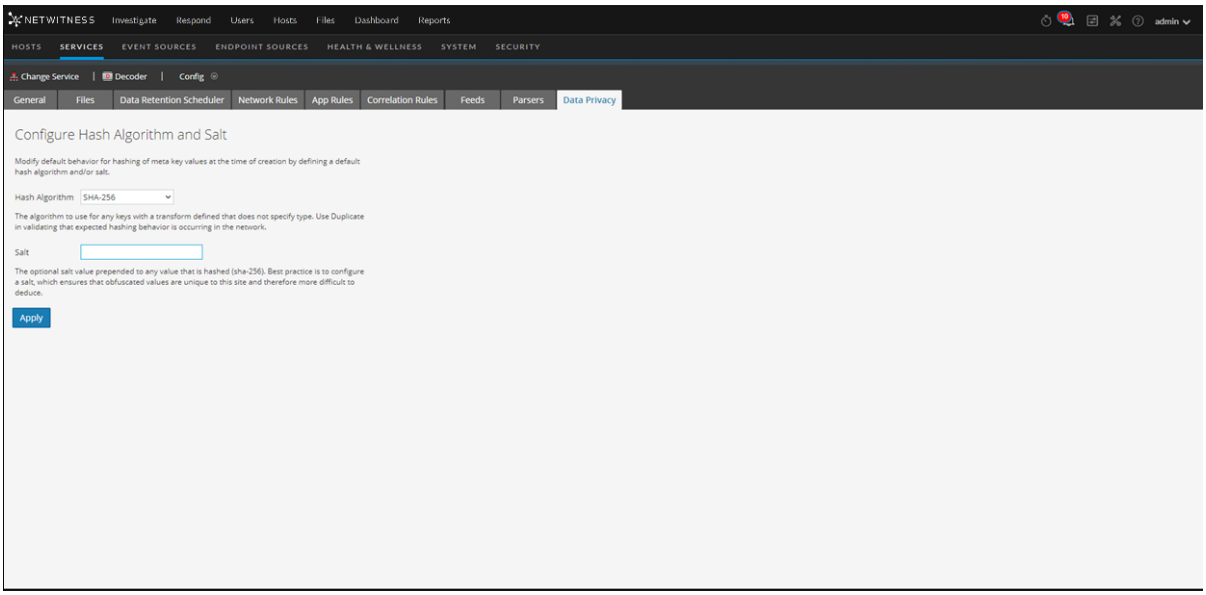
  The algorithm to use for any keys with a transform defined that does not specify `type`. The supported algorithms are: `duplicate` and `sha-256`.

- `/decoder/parsers/transforms/default.salt`

  The salt value prepended to any value that is hashed (`sha-256`). This value is optional, an empty salt is valid and produces an unsalted hash. The salt is not defined by default so that you can create a unique salt for your environment. In general, the longer and more complex the salt, the better the security. A salt of up to 60 characters can be used without any major impact. A salt of at least 16 characters is recommended.

To edit the default hash algorithm and salt:

1. Go to  **(Admin) > Services**.
2. In the **Services** view, select a Decoder or Log Decoder service, and select  **> View > Config > Data Privacy**.



3. In the **Configure Hash Algorithm and Salt** section, select a **Hash Algorithm** to use for any meta keys with a defined transform that does not specify type: `sha-256`. (A second algorithm, `duplicate`, is available for administrators to use in validating that expected hashing behavior is occurring in the network.)
4. (Optional) In the **Salt** field, enter a salt value to be prepended to any value that is hashed. This value is optional, an empty salt is valid and produces an unsalted hash. The salt is not defined by default so that you can create a unique salt for your environment. In general, the longer and more complex the salt, the better the security. Best practices for security purposes dictate a salt value that is no less than 100 bits or 16 characters in length. If a unique salt is required for each individual meta key, that needs to be configured in the index file as shown in example 3 below.
5. Click **Apply**.

   The new settings become effective immediately.

Configure Language Keys

The NetWitness Core Language has several language key attributes to facilitate data privacy. You can edit these attributes in the custom index file for each Decoder or Log Decoder. The custom index file (for example, `index-decoder-custom.xml`) is editable in the Services Configuration view > Files tab. After making changes in the index file, like the ones shown in the examples below, a service restart in a specific sequence is required.

Based on the data privacy requirements for your site, configure individual meta keys to be protected using the following `key` attributes:

- `protected`

  This attribute specifies that NetWitness should consider the values as protected and tightly control any release of the value. When propagating the protected attribute, NetWitness ensures that any downstream trusted system treats the values accordingly. Add this attribute to all services that create the protected values (that is, Decoder or Log Decoder) and any services that will provide trusted access (software development kit (SDK) query/values, aggregation) outside of Core services. The exception to this rule is that a Broker with no index file specified does not need to have the attribute added.

- `token`

  This attribute specifies that values for this key are stand-ins for another value and may not be visually interesting. The `token` attribute is informational, primarily for UI elements to display the value in a more useful or visually pleasing format.

- `transform`

  This child element of `key` indicates that any values for a given meta key should be transformed and the resulting value persisted in another meta key. The `transform` element is only required on Decoders and Log Decoders and is informational if specified on any other Core services. The `transform` element has the following attributes and children:

Name	Type	Description	Optional or Required
`destination`	attribute	Specifies the key name where the transformed value will be persisted.	required
`type`	attribute	The transform algorithm to apply. If not specified, the value of `/decoder/parsers/transforms/default.type` is used.	optional
`param`	child-element	A name/value pair, where each `param` element has a required attribute `name` and the child text is the value. The only supported `param` is used to specify a meta key-specific `salt` value. If not specified, the following value is used: `/decoder/parsers/transforms/default.salt`	optional

Example 1

On a Decoder or Log Decoder, mark `username` as protected and hash all values into `username.hash` with the default algorithm and salt:

```
<key name="username" description="Username" format="Text" protected="true">
<transform destination="username.hash"/></key>
```

Example 2

On a Concentrator, mark `username` as protected and `username.hash` as token:

```
<?xml version="1.0" encoding="utf-8"?> <language level="IndexNone"
defaultAction="Auto">
<key description="Username" format="Text" level="IndexValues" name="username"
protected="true"/>
<key description="Username Hash" format="Binary" level="IndexValues"
name="username.hash"
token="true"/> </language>
```

Example 3

On a Decoder or Log Decoder, mark `username` as protected and hash all values into `username.bin` with the specified algorithm and salt:

```
<key name="username" description="Username" format="Text" protected="true">
<transform destination="username.bin" type="sha-256">
param name="salt">0000</param> </transform></key>
```

Configure Metadata and Content Visibility Per User Role on Core Services

On individual Broker, Concentrator, Decoder, Log Decoder, and Archiver services being viewed in the Services Security view  **(Admin) > Services > Security**), administrators can configure the visibility of metadata and content based on the user group or role assigned to a user. This is called the SDK meta roles capability, and it is enabled by default. Administrators who want to configure metadata and content visibility per user must not disable the `sdk.content` permission (in the Roles tab). If the `sdk.content` permission has been disabled in the Roles tab, packets and raw logs are not visible to `system.roles` node. The `system.roles` node handles the filtering using the method configured in the Settings tab.

With `sdk.content` capability enabled, the next step is to select the method of filtering metadata and content in the Settings tab. Selecting a blacklist or whitelist option makes additional permissions for specific meta keys available in the Roles tab. The result is that administrators can choose a user role, such as analyst, in the Roles tab and select specific meta keys (and content) to be blacklisted or whitelisted for that user group. The permissions apply to any user in the user group.

For the Events view to display events correctly, the `sdk.meta.medium` meta key must be enabled to display values in the Type column (one of the default columns in the Events list). In addition, any meta keys used in a Query Prefix under User Settings in the Users tab or a PreQuery configured for Investigate must be enabled  **(Admin) > Security > Users** or  **(Admin) > Security > Roles**). If the meta key is not enabled for the user group, the column will not be displayed in the Events list.

Note: In the Version 11.4 and later Events panel, events are listed even when the user does not have permission to view the event reconstruction. In addition, the download options may still be available when a user does not have permission to view the event reconstruction. In these cases the user is allowed to download the events, but the output is an empty file because restricted permissions are enforced during the download.

The following table lists the options for filtering in the Settings tab and the numeric values used to disable (0) and the types of filtering (1 through 6). There is no need to know the numeric value unless configuring metadata and content visibility manually in the `system.roles` node.

system.roles Node Value	Settings Tab Option	Event Metadata	Original Event
0	**No Filtering.** System roles that define permissions on a per meta key basis are disabled.	Visible	Visible
1	**Whitelist meta and content.** By default no meta keys and no packets are visible. Selecting individual SDK meta roles per user group allows users to see metadata and packets for that SDK meta role.	Not Visible-Select to Show	Not Visible-Select to Show
2	**Whitelist only meta.** By default packets are shown, but no metadata is visible. Selecting individual SDK meta roles per user group allows users to see metadata for that SDK meta role.	Not Visible-Select to Show	Visible
3	**Whitelist only content.** By default metadata is visible, but no packets are visible. Selecting individual SDK meta roles per user group allows users to see packets for that SDK meta role.	Visible	Not Visible-Select to Show
4	**Blacklist meta and content.** By default all metadata and all packets are visible. Selecting individual SDK meta roles per user group prevents users from seeing metadata and packets for that SDK meta role.	Visible-Select to Hide	Visible-Select to Hide
5	**Blacklist only meta.** By default all metadata and all packets are visible. Selecting individual SDK meta roles per user group prevents users from seeing metadata for that SDK meta role.	Visible-Select to Hide	Visible
6	**Blacklist only content.** By default all metadata and all packets are visible. Selecting individual SDK meta roles per user group prevents users from seeing packets for that SDK meta role.	Visible	Visible-Select to Hide

Three factors determine what a user sees:

- The SDK meta role setting (blacklist or whitelist).
- The restricted meta keys configured for the group to which the user belongs.

- The meta keys in the session being analyzed. These include meta keys specified in a query prefix for a user at the service level, and for roles and users at the system level.

Caution: Be aware that with blacklisting, implicit trust is granted for all except the configured metadata. For a Decoder to have RBAC enabled and use implicit trust, it must only use a blacklist system setting; a whitelist setting will result in some issues with meta keys that are not explicitly enabled and therefore not visible. It is impossible to grant implicit trust under whitelist rules because the universe of meta keys cannot be known. If you want to use whitelisting, a workaround is to turn RBAC off for the Decoder and disable any user accounts from connecting directly to the Decoder if they should use RBAC.

Here is an example of how the SDK meta role configuration meshes with a Group that has restricted meta keys.

Configuration:

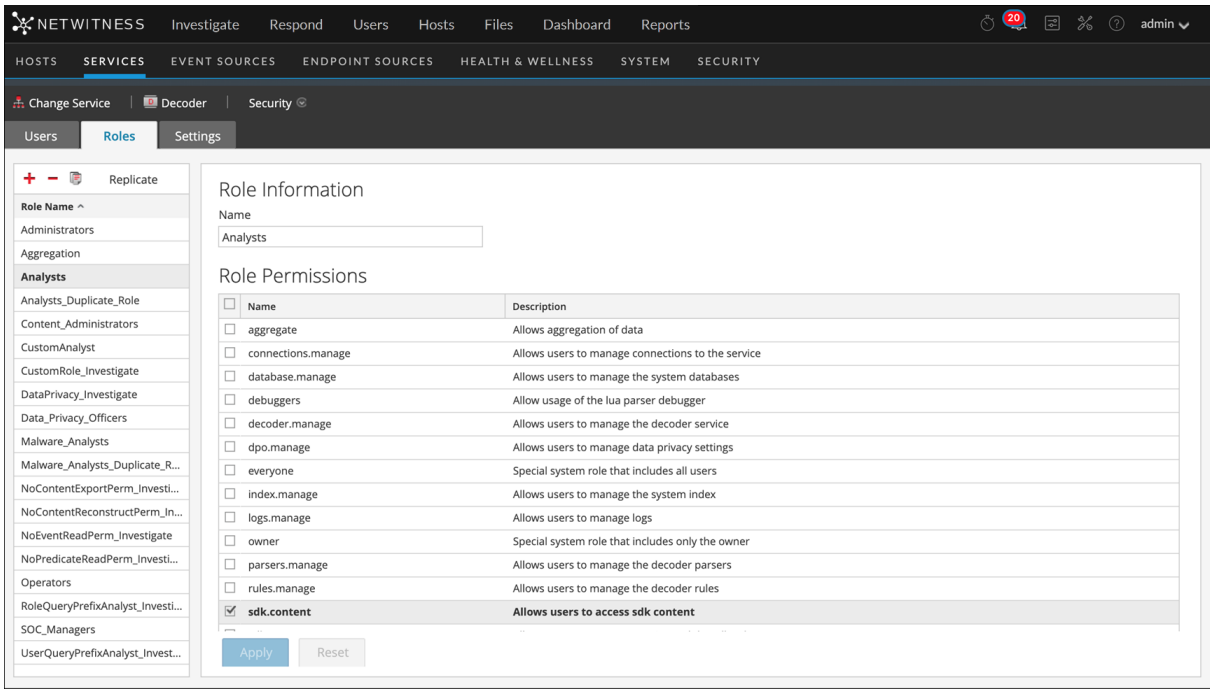
- The SDK meta role setting is **Blacklist meta and content**. With this option implemented, all metadata and all content (packets and logs) are visible by default.
- The administrator has restricted meta keys configured for the Analysts group to prevent viewing of sensitive data (for example, `username`).
- The packets and logs for any session that includes the `username` meta key are not visible to an analyst.

Result: Now a user who is a member of the Analyst Group investigates a session. Depending on the content of the session, the results are different:

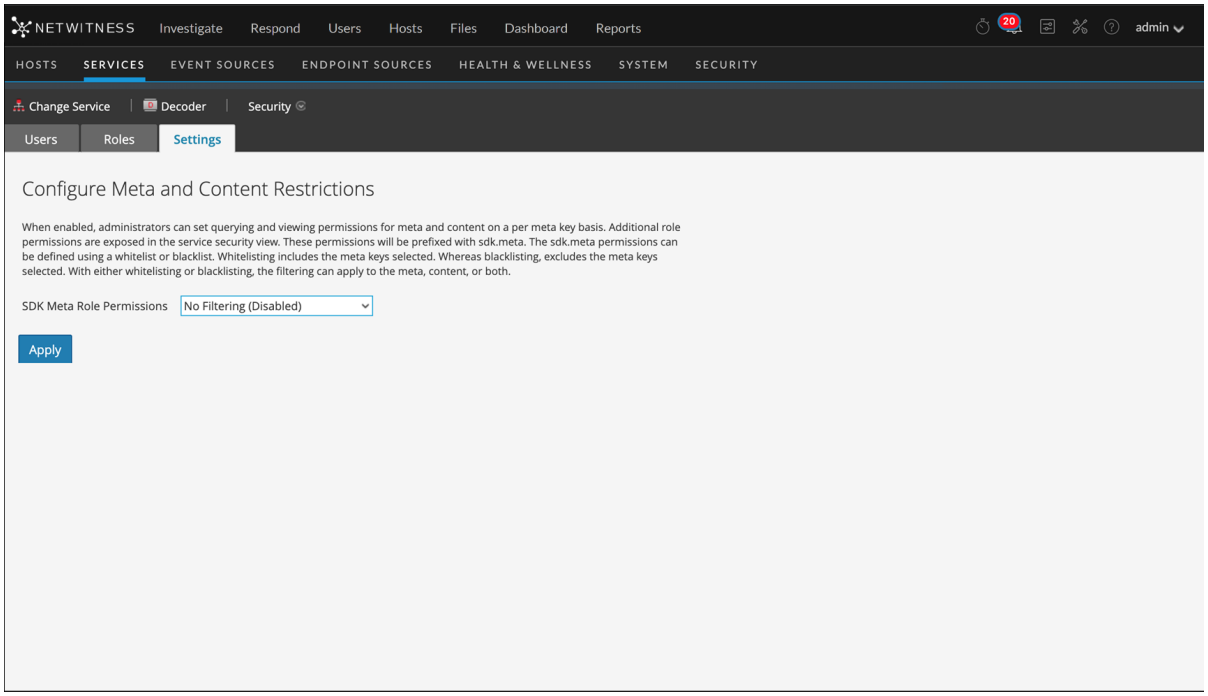
- Session 1 includes the following meta keys: `ip`, `eth`, `host`, and `file` The session does not include `username` so all packets and logs in the session are displayed.
- Session 2 includes the following meta keys, `ip`, `time`, `size`, `file`, and `username`, Because the session includes `username` , no packets or logs from the session are displayed for the analyst.

To configure metadata and content restrictions for a Decoder or Log Decoder:

1. Go to  **(Admin)** > **Services**.
2. In the **Services** view, select a Broker, Concentrator, Decoder, Log Decoder, or Archiver service and select  **> View > Security**. Click the **Roles** tab, select a role, and verify that the `sdk.content` permission is enabled.

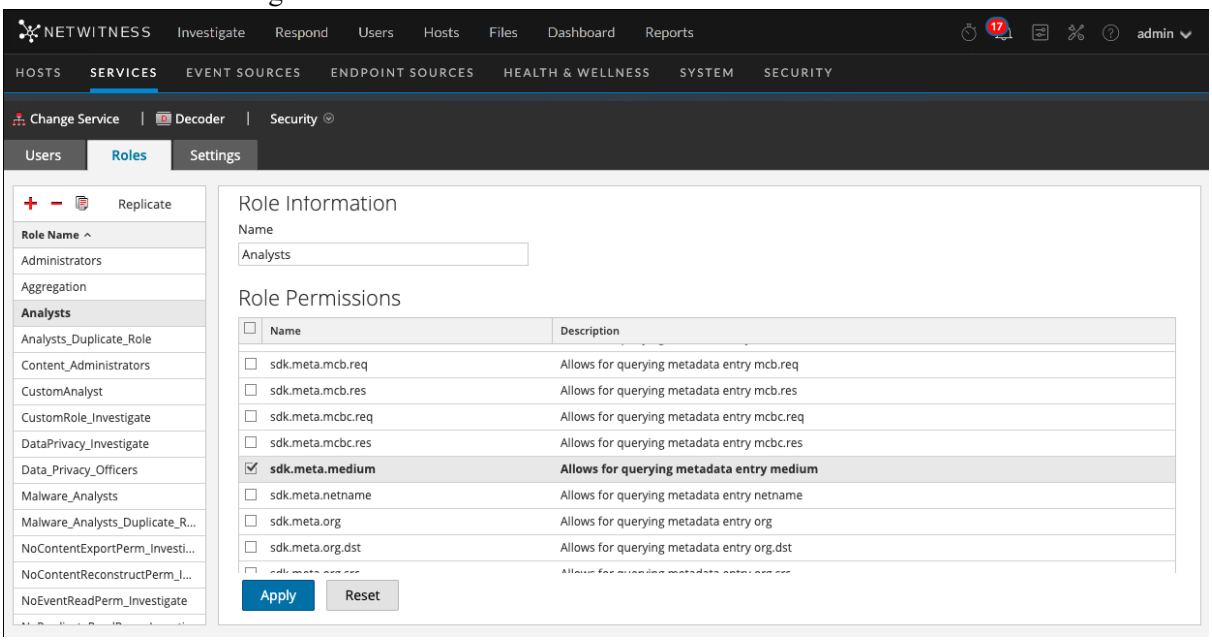


3. Click the **Settings** tab.



4. Select one of the filtering methods (blacklist or whitelist) and content types (meta and content, meta only, or content only), and click **Apply**.
5. Click the **Roles** tab and a role for which you want to allow content (whitelist) or block content (blacklist) as specified in the SDK Meta Role Permissions setting.

The Role Permissions for the selected role are displayed, and the SDK Meta Role Permissions are available for selection, for example, `sdk.meta.medium`. If you selected one of the whitelist options in the SDK Role Permissions setting, you must assign each SDK meta role to make the selected content visible to users assigned that SDK meta role. The `sdk.meta.medium` permission must be enabled in order for the Type column to be displayed in the Investigate > Events view. If you selected one of the blacklist options in the SDK Role Permissions setting, selected content will be hidden from users assigned that SDK meta role.



6. Select the SDK meta role permissions for users assigned this role. Click **Apply**.

The settings become effective immediately and apply to new packets and logs processed by the Decoder or Log Decoder.

Configure Meta Keys Not Written to Disk Per Parser on a Decoder

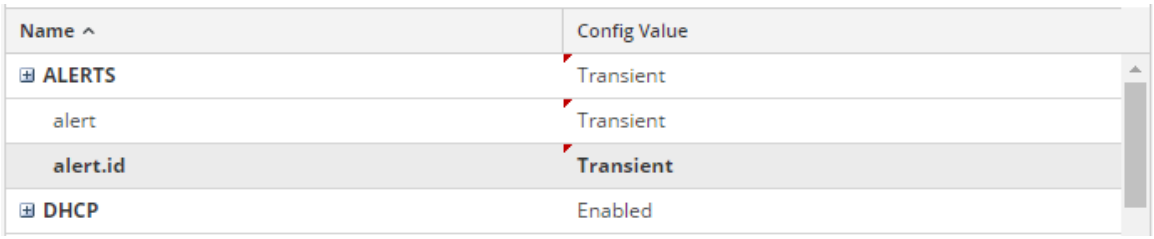
On a Decoder and Log Decoder, a Data Privacy Officer can configure individual meta keys that are not written to disk. To do so, the DPO specifies the meta keys as transient in the index and the parser configuration.

Note: The same capability was previously available on Log Decoders, and was configured when setting up parsers by modifying the table-map.xml file. Now it is integrated in the Services Config view.

To configure selected meta keys on individual parsers that will not written to disk:

1. Go to  **(Admin)** > **Services**.
2. In the **Services** view, select a Decoder or Log Decoder service and select  **> View > Config**.
3. In the **Parsers Configuration** section of the **General** tab, select a parser and then select **Transient** in the **Config Value** drop-down list. Access the list by clicking on the configuration value (Enabled, Disabled, or Transient).

The configuration change is marked by a red triangle.



4. Click **Apply**.

   The change is effective immediately. The parser configured as Transient will no longer store meta keys to disk.

Configure Data Retention

A NetWitness user with the role of Administrator can configure NetWitness to ensure that sensitive data has been removed after a specific retention period, regardless of system ingest rate. For instance, the policy might be to keep packets (both raw data and metadata) for no more than 24 hours, and to keep some logs (both raw data and metadata) for up to seven days. If sensitive data makes its way into another database on the Reporting Engine, Malware Analysis, Event Stream Analysis, and NetWitness Servers, data retention can be managed there as well. The administrator needs to set up each service individually across all NetWitness components (except Event Stream Analysis) based on policy and data privacy laws.

Sensitive data may also be in cache.

- Brokers can cache data and this needs to be cleared by configuring an independent rollover and other removal of cache as required. The administrator can configure cache rollover for a Broker by editing the `Scheduler` file in the Services Config view Files tab.
- Investigate and the NetWitness Server cache data, and this is cleared automatically every 24 hours.
- If the Data Privacy Officer (DPO) exports data, that is the same as saving data on the NetWitness Server in the jobs queue. To clear this data, the administrator or DPO should clean up the jobs queue on a regular basis.

Data Retention

You can schedule a recurring job for Decoder, Log Decoder, and Concentrator services in NetWitness to check if data is ready to be removed. The Data Retention Scheduler provides a means to configure basic scheduling (see below), and advanced Scheduler settings are also available by editing the `Scheduler` file in the Services Config view Files tab or the node in the Explorer view.

The Archiver has flexible data storage and retention options. You can place different types of log data into individual collections and manage them separately. These collections enable you to specify how much of the total storage space to use and how many days to store the logs in the collection. You can also determine whether to delete the log data or to move it to offline cold storage after it reaches the maximum specified storage space for the collection.

For example, you can put sensitive information in a collection and configure a limitation on how long to keep it, such as 30 days. To delete the data after 30 days, you would not enable warm or cold storage for that collection.

Deleting versus Retaining Log Data

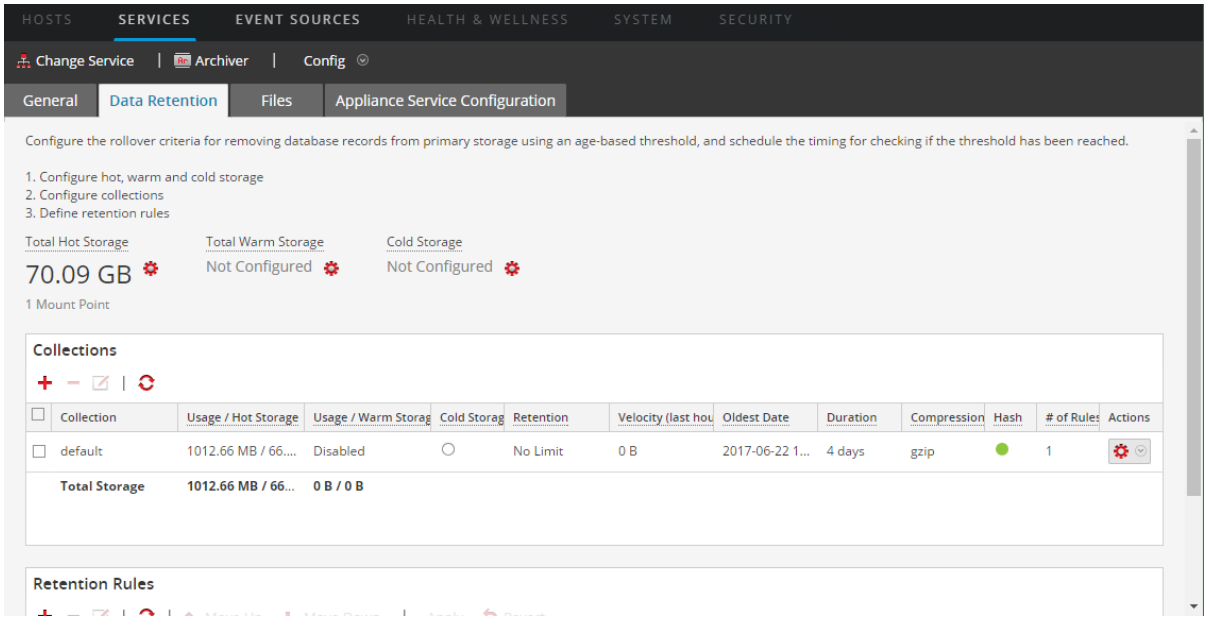
Administrators can configure hot, warm, and cold tiered storage on an Archiver. Cold storage contains the oldest log data that is either required for the operation of the business or mandated by regulatory requirements. When a collection reaches its retention limits for hot and warm storage, NetWitness deletes the log data from hot or warm storage. With cold storage configured, a copy goes into cold storage before the logs are deleted from hot or warm storage. You can choose whether to enable cold storage for each log storage collection. NetWitness does not manage cold storage.

Enable or Disable Cold Storage in a Log Storage Collection

When log data in a collection reaches its retention limits for hot and warm storage, you can delete it or move it to offline (cold) storage.

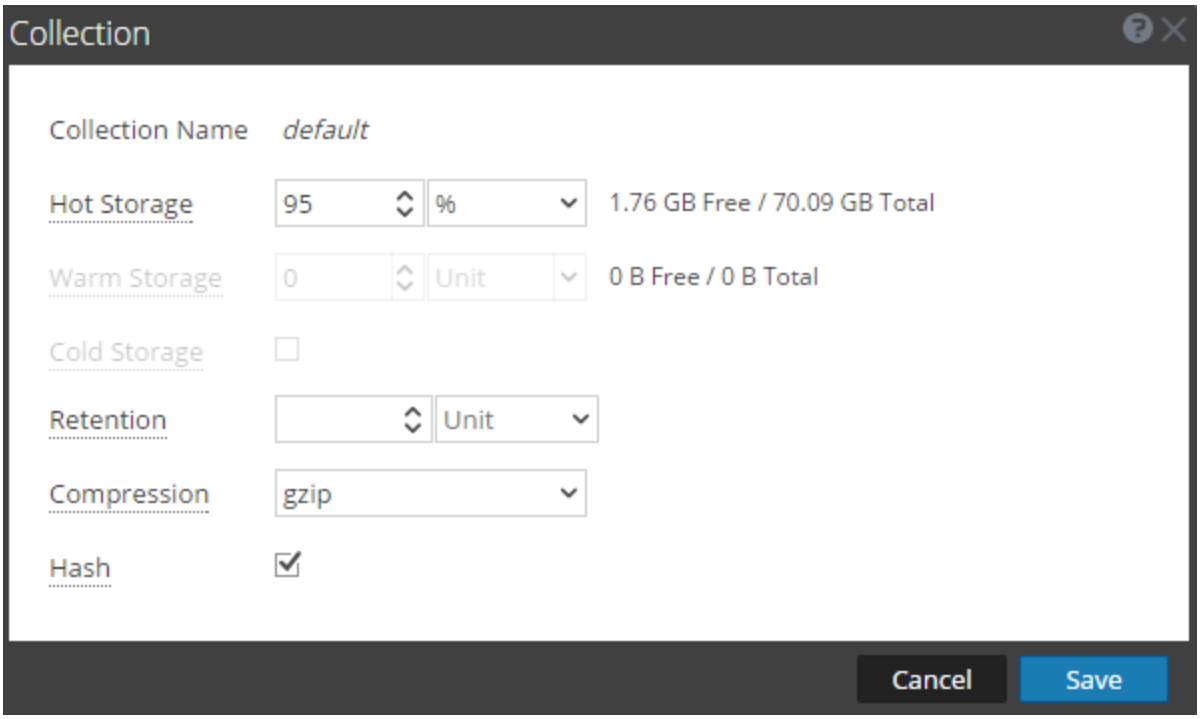
To enable or disable cold storage in a log retention storage collection on an Archiver:

1. Go to **(Admin)** > **Services**.
2. Select the Archiver service and select **> View > Config**.
3. Click the **Data Retention** tab.



4. In the **Collections** section of the Data Retention tab, select a collection and click .

   The Collection dialog is displayed.



Note: If the maximum storage size of the collection does not allow full data retention for the retention period specified, NetWitness deletes the data or it goes to warm or cold storage if specified in the collection.

5. Enable or disable cold storage:
   - To delete log data when the collection reaches its specified retention limits, clear the **Cold Storage** checkbox.
   - To move log data to offline storage when the collection reaches its specified retention limits, select the **Cold Storage** checkbox.
6. Click **Save**.

Configure Log Retention and Storage on an Archiver

To configure log retention and storage on an Archiver, see "Configure Archiver Storage and Log Retention" in the *Archiver Configuration Guide*.

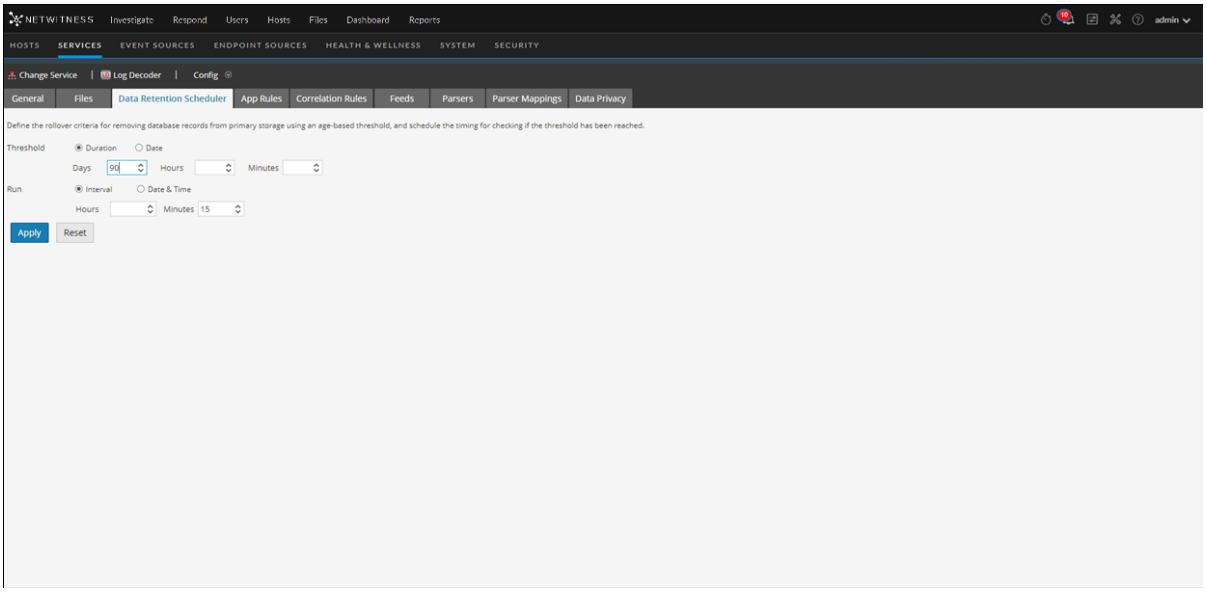
Schedule a Recurring Job to Check Data Retention Thresholds

The data retention scheduler configuration ensures that the data residing in the Decoder, Log Decoder, and Concentrator components is deleted after a certain time. For example, data retention on a Decoder might be configured to execute a check every 15 minutes to determine if the specified duration threshold has been met. If the threshold is met, NetWitness deletes data older then 4 hours in the relevant databases.

Caution: The schedule overwrites any previous schedule and becomes effective immediately. If the retention period is decreased, the data exceeding this retention period is removed.

For a Decoder, Log Decoder, or Concentrator:

1. Go to  **(Admin)** > **Services**.
2. In the **Services** view, select a Decoder, Log Decoder, or Concentrator service and select  **> View > Config**.
3. Click the **Data Retention Scheduler** tab.



4. Set the threshold based on the duration of time the data has been stored or the date on which the data was stored. Do one of the following:
   a. To define the duration of time that data can be stored before removal, select **Duration**, and then specify the number of days (365 maximum), hours (24 maximum), and minutes (60 maximum) that have elapsed since the time stamp on the data.
   b. To define the removal of data based on the date of the timestamp, select **Date**, and then specify the monthly date and time in the Calendar and Time fields.
5. Do one of the following to configure the **schedule for checking rollover criteria**:
   a. If you want to set a regular interval at which the scheduled database check occurs, select **Interval** and specify the **Hours** and **Minutes** between the scheduled checks.
   b. If you want to set a regular date and time at which the scheduled database check occurs, select **Date and Time** and specify the system clock time in `hh:mm:ss` format for the rollover.
      - To specify the day, select **Every Day**, **Weekdays,** or **Weekends**. The Scheduler defaults to **Every Day**.
      - To specify a different set of days of the week, select **Custom** and click on each day on which the database check occurs.

      Caution: The schedule overwrites any previous schedule and becomes effective immediately. If the retention period is decreased, the data exceeding this retention period

is removed.

6. Click **Apply** to complete the configuration.

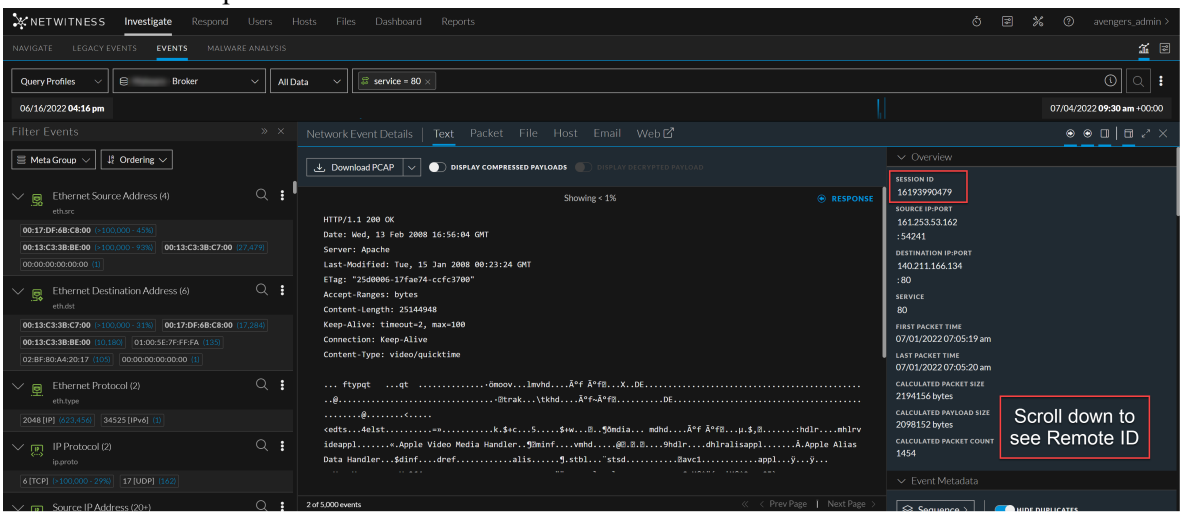
Purge Data Using String and Pattern Redaction Option

If you have data where you do not want it to be (for example, you have classified data in your unclassified Network Decoder), purging data using string and pattern redaction is the fastest way to wipe data. It is ideal for smaller amounts of data. This method requires that you identify the data to be purged by using a query in either the Investigate view, the REST API interface, or the NwConsole. Using the Investigate view requires additional steps to clear user interface cache on the admin server. This procedure includes the extra steps needed when you find the data to be purged using the Investigate view.

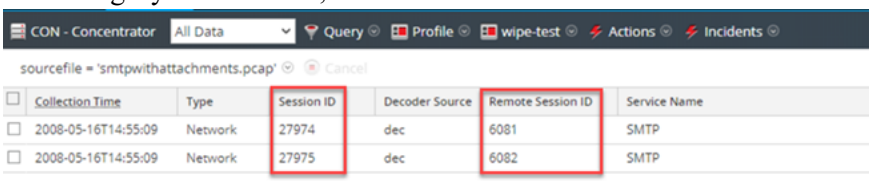
Note: If a large amount of data needs to be purged, it is best to start with the data on the storage component. Refer to Optional Data Overwriting Options for additional information.

To purge the data:

1. Go to **Investigate** > **Events** or **Investigate** > **Legacy Events** and find the events that you want to purge. In the Events view you can see the session ID in the Overview panel. You need to scroll down in the Event Meta panel to see the remote ID.

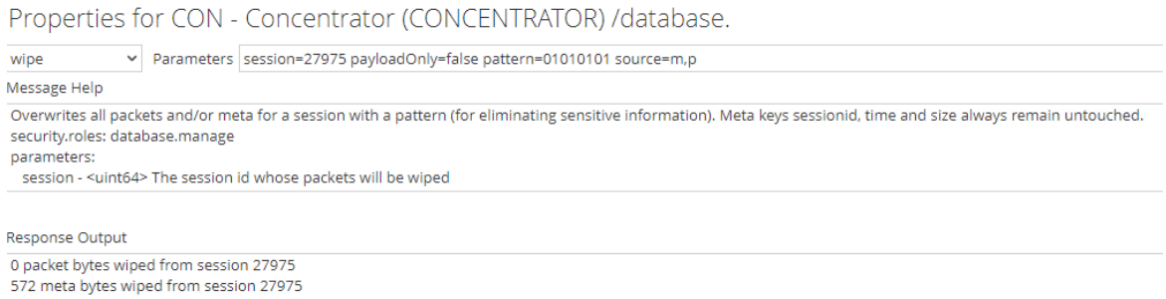


In the Legacy Events view, both IDs are visible in the Events list.

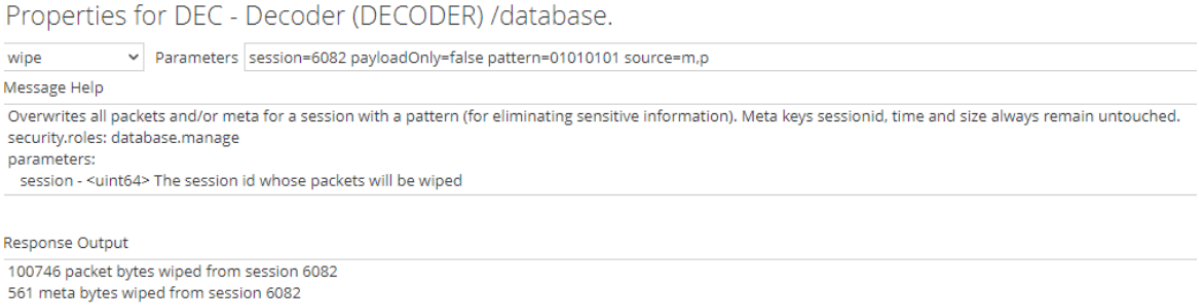


2. Make a note of the session ID and the remote ID of each event to be deleted. In the example above, the session IDs are 27974 and 27975. The remote IDs are 6081 and 6082.

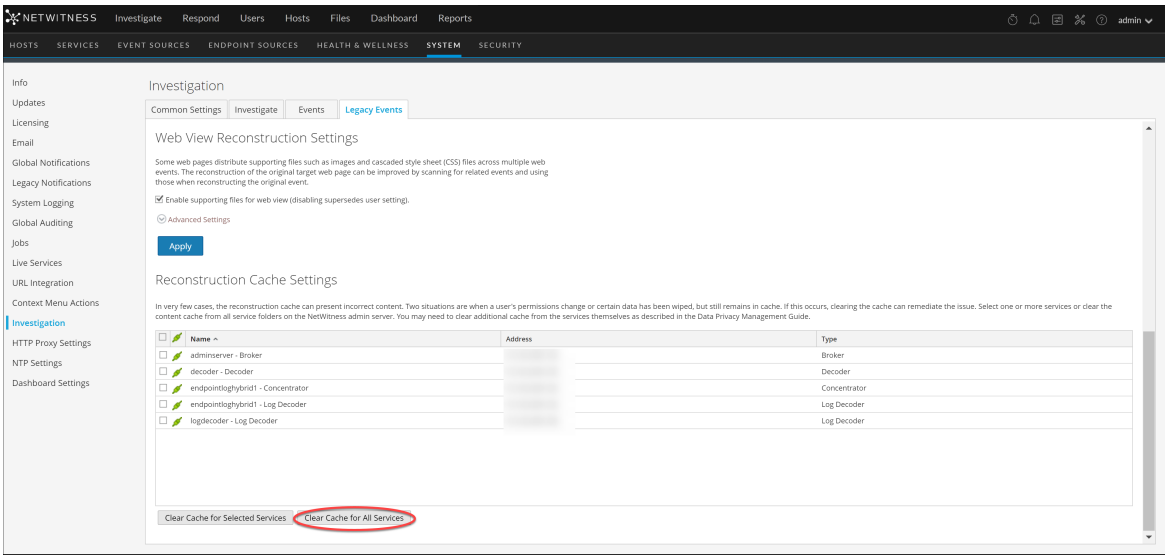
3. Starting with the Concentrator, view the service using the REST API. Refer to the *RESTful API User Guide* for instructions. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
   In the REST API, enter the `wipe` command with the following parameters:
   - **`session=<uint64>`**: The session ID whose packets will be wiped.
   - **`payloadOnly=<bool, optional>`**: If `payloadOnly` is `true`, only the packet payload will be overwritten. The default value is `true`.
   - **`pattern=<string, optional>`**: The pattern uses all zeros by default. If you use a string as your pattern, it will not overwrite any meta values that are not a string type. Therefore, it is best to keep the pattern as a numerical value.
   - **`metaList=<string, optional>`**: Comma-separated list of metadata to wipe. The default (empty) is all metadata.
   - **`source=<string, optional, {enum-any:m|p}>`**: The types of data to wipe: metadata, packets, or both. The default is packets only. The best option here is to wipe both metadata and packets `(m,p)`.



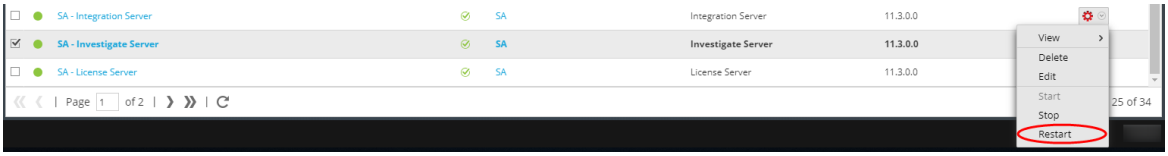
4. For each Network Decoder and Log Decoder in the path of the query, use the `wipe` command and the remote session IDs to overwrite the raw sessions on disk.



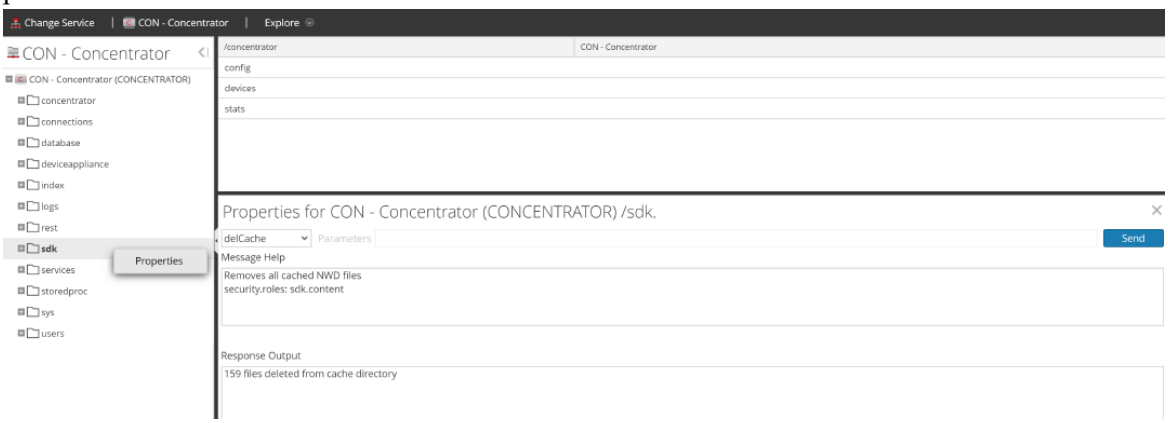
5. To ensure that the indexed meta values that were stored on the Concentrator index are removed, rebuild the index. This can take a long time but is necessary because the `wipe` command does not remove any data from the Concentrator index. Refer to the *Core Database Tuning Guide* for instructions.
6. (Optional) If the data was discovered using the Investigate view, additional steps are necessary:
   a. To clear the query path to prevent an analyst from bringing up an old copy of the session, go to: **(Admin)** > **System** > **Investigation** and select **Clear Cache for All Services**.



b. To clear the cache for a reconstruction in the Events view, restart the Investigate service.



c. To clear the cache for the Concentrator and Network Decoder, go to **(Admin)** > **Services** > **Concentrator** > right-click **sdk** > **Properties** and **(Admin)** > **Services** > **Decoder** > right-click **sdk** > **Properties** and execute the `delCache` command. The figure below shows this path for the Concentrator.



An analyst attempting to view the same session in Investigate will see that the raw data is unavailable for viewing and the metadata has been overwritten.

Configure User Accounts for Use in Data Privacy

This topic provides the procedures for configuring user accounts that work with data obfuscation in NetWitness. In order for data obfuscation to work, accounts and permissions for several types of users must be configured.

- Customize the default Administrators system role in NetWitness to remove permissions that should be available only to the Data Privacy Officer.
- Add two new user accounts at the system level to depict a data privacy officer and a typical analyst.
- Add a user account at the service level with the aggregation role so that Decoders and Log Decoders can aggregate data to a Concentrator or Broker.
- On the Reporting Engine, configure two separate service accounts. One service account for general purpose reporting that does not include any sensitive data and the other account for privileged users with access to all data including sensitive data. This procedure is described in "Configure Data Source Permissions" in the *Reporting Engine Configuration Guide*.

Customize the Default Administrators User Role at the Service Level

To separate the data privacy officer and administrator functions on each Decoder and Log Decoder, you need to remove the `dpo.manage` permission from a clone of the Administrators role.

1. Go to  **(Admin)** > **Services**, select a Decoder or Log Decoder, then select  **> View > Security**.
2. In the **Services Security** view, click the **Roles** tab, select **Administrators** and click .
   In the **Enter Role Name** dialog, enter a new role name such as Non_DPO_Administrators and click **Save**.
3. Select the new role.

   The Role Information is displayed for editing.
4. Clear the box next to **dpo.manage** so that it is no longer checked and click **Apply**.

   The permission to manage data privacy configuration is removed for the new role.
5. In the **Users** tab, select each user who has the **Administrators** role, and change their role to the cloned role.
6. Validate that the users with the modified Administrators role can login as with admin privileges.
7. Validate that the users with the modified Administrators role cannot configure metadata and content restrictions in the Settings tab.

Add a User Account with the Aggregation User Role at the Service Level

To ensure that Decoders and Log Decoders can aggregate data to a Concentrator or Broker:

1. **Go to** **(Admin) > Services, select a Decoder or Log Decoder, then select** **> View > Security.**
2. In the **Users** tab, add a user with the `Aggregation` role and click **Apply**.

Note: "Aggregation Role" in the *Hosts and Services Getting Started Guide* provides details about the application of this user role.

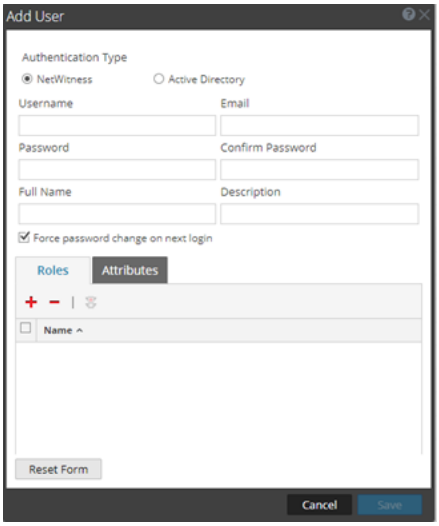
Add Data Privacy Officer and Analyst Accounts on the NetWitness Server

You need to add two new user accounts in NetWitness at the system level to depict a privileged data privacy officer and a typical analyst. If the environment is configured using the default trusted connections, you do not need to create the new user accounts on the Core services (Brokers, Concentrators, and Decoders). When a user is created in the NetWitness Server, that user can log on to the services.

Note: The role name is required to exist on both the server and the services, and the role name for all places must be identical. If you create a new custom role on the NetWitness Server, make sure to add it to all Core services as well.

1. Create a new user account for the data privacy officer:
   a. **Go to** **(Admin) > Security,** select the **Users tab** and click **.**

   The Add User dialog is displayed.



   b. Create the new account with the following credentials.

   Username = <new user name for logon, for example, DPOadmin>
   Email = <new user's email, for example, DPOadmin@rsa.com>
   Password = <new user's password for logging on, for example, RSAprivacy!@>
   Full Name = <new user's full name, for example, DPO Administrator>

c. Click the **Roles** tab, , and select the `Data_Privacy_Officers` role for the new user.

d. Select **Save**.

2. Create a new user account for the analyst with limited privileges:

a. In the  **(Admin) > Security** view, click the **Users** tab. In the **Users** tab toolbar, click .

The Add User dialog is displayed.

b. Create the new account with the following credentials:

Username = <new user name for logon, for example, NonprivAnalyst>
Email = <new user's email, for example, NonprivAnalyst@rsa.com>
Password = <new user's password for logging on, for example, RSAprivacy!@>
Full Name = <new user's full name, for example, Nonprivileged Analyst>

c. Click the **Roles** tab, , and select the `Analysts` role for the new user.

d. Select **Save**.

Data Privacy References

The following reference materials are available for management of data privacy and data retention. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

- See "Data Privacy Tab" in the *Decoder and Log Decoder Configuration Guide.*
- See "Data Retention Tab - Archiver" in the *Archiver Configuration Guide.*
- See "Data Retention Scheduler Tab"in the *Hosts and Services Getting Started Guide.*

NetWitness® Platform XDR

Version 12.0

Core Database Tuning Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

NetWitness Core Database Introduction **8**

NetWitness Platform Products Covered by this Guide 8

Frequently Used Terms 8

NetWitness Core Database History 9

Core Database Strengths and Weaknesses 9

Basic Database Configuration **11**

Find Help within the Core Service 11

Packet, Meta, and Session Storage 11

Index Storage 11

Tiered Database Storage 11

Archiver 12

Manifests 13

Search Historical Manifests 14

Advanced Database Configuration **16**

Database Configuration Nodes 16

packet.dir , meta.dir , session.dir 16

packet.dir.warm , meta.dir.warm , session.dir.warm 17

packet.dir.cold , meta.dir.cold , session.dir.cold 17

packet.file.size , meta.file.size , session.file.size 18

packet.files , meta.files , session.files 18

packet.free.space.min , meta.free.space.min , session.free.space.min 19

packet.index.fidelity , meta.index.fidelity 19

packet.integrity.flush , meta.integrity.flush , session.integrity.flush 19

packet.write.block.size , meta.write.block.size , session.write.block.size 19

packet.compression , meta.compression 20

packet.compression.level , meta.compression.level 20

hash.algorithm 20

hash.databases 20

hash.dir 20

Index Configuration Nodes 20

index.dir 21

index.dir.warm 21

index.dir.cold 21

index.slices.open 21

page.compression 21

save.session.count ... 21
reindex.enable ... 22
SDK Configuration Nodes ... 22
max.concurrent.queries ... 22
max.pending.queries ... 22
cache.window.minutes ... 22
max.where.clause.cache ... 23
max.unique.values ... 23
query.level.1.minutes , query.level.2.minutes , query.level.3.minutes ... 23
query.timeout ... 23
max.where.clause.sessions ... 23
max.query.groups ... 24
packet.read.throttle ... 24
cache.dir , cache.size ... 24
pin.cache.dir , pin.cache.size ... 24
parallel.values ... 24
parallel.query ... 25
Per-User Configuration Nodes ... 25
query.prefix ... 25
query.level ... 25
query.timeout ... 25
session.threshold ... 26
Scheduler ... 26
Example ... 26
Rollover ... 27
Synchronous Rollover ... 27
Asynchronous Rollover ... 27
Example ... 28
Snort Rules and Configuration ... 29
Configuration ... 29
Meta key usage ... 29
Rules ... 30
General options ... 30
Payload options ... 31
Non-payload options ... 32
Queries ... 33
query Syntax ... 33
where Clauses ... 35
Query Operators ... 36
Text Values ... 37

IP Addresses ... 37
MAC Addresses ... 37
Numeric Values ... 37
Bucketed Numeric Indexes ... 38
Numeric Value Aliases ... 38
Date and Time Expressions ... 38
Relative Time Points ... 39
Special Range Values ... 39
group by Clause (since 10.5) ... 40
order by Clause (since 10.5) ... 41
search parameter ... 42
values call ... 42
Parameters ... 43
values Flags ... 44
values Call Example ... 45
Values call and bucketing mode ... 46
Suggestion Mode ... 46
search parameter ... 46
msearch Call ... 46
msearch Flags ... 48
msearch Index Search Mode ... 48
Text Search Syntax ... 48
Search Syntax And Index Modes ... 48
msearch Tips ... 49
Stored Procedures ... 49
Use of Quotes in Query Syntax ... 49
hierarch Call ... 50
Index Customization ... 51
Index Configuration File Locations ... 51
Index configuration entries ... 51
Meta names ... 52
Data Types ... 52
Index Levels ... 53
Value Max ... 54
maxLength ... 54
Max Length without N-Grams ... 54
Max Length with N-Grams ... 54
minLength ... 55
ngrams ... 55
Numeric Bucketing ... 56

Key Value Aliases ... 56
Key Renaming ... 57
Entities ... 58
Entity Definition Rules ... 58
Entities in Brokers ... 59
Rebuilding the Index ... 60
Activating the Background Reindexer ... 60
Controlling the Background Reindexer ... 60
Background Reindexing Algorithm ... 60
Background indexer status ... 61
Effects on Aggregation ... 61
Forcing A Reindex ... 61
Optimization Techniques ... 62
Thresholds ... 62
Complex where Clauses ... 62
AND and OR ... 63
Use Case: Match a Large Subnet ... 63
Starting with NetWitness 11.4, all /8, /16, and /24 IPv4 subnets get their own index entries. So, no special action is needed to help optimize these queries. ... 63
Use Case: Substring Matching ... 64
Index Saves ... 64
Affects of Increasing the Save Interval ... 65
Affects of Decreasing the Save Interval ... 65
Working with valueMax ... 65
Parallelize Workloads ... 65
Index Rebuild ... 66
Scaling Retention ... 66
Increasing Packet and Meta Retention ... 66
Increasing Index Retention ... 66
Scaling Horizontally ... 67
Grouping Workloads ... 67
Cache Window ... 67
Time Limits ... 68
Query Priority ... 68
Rule Examples ... 70
Correcting invalid rules ... 70
Valid Syntax with the Modern Parser ... 71
Appendix A: Statistics ... 72
Statistics in /database/stats ... 72
Statistics in /index/stats ... 72

Statistics in /sdk/stats 73
Per-query Statistics 74
Appendix B: Index Inspect 75
Parameters 75
Response 75
Slice Summary 75
Per-Index Summary 75
Slice Summary Footer 76

NetWitness Core Database Introduction

This topic provides an overview of the NetWitness Platform Core database. The NetWitness Platform Core services contain a proprietary database developed specifically for use within the NetWitness Platform products. It bears little resemblance to traditional relational databases, and is not based on any off-the-shelf database technology. As such, many users find that there is a steep learning curve to understanding how the Core database works, and how to make best use of it. The purpose of this guide is to help NetWitness Platform users understand the database and use it to its fullest potential.

As a System Administrator, you can use this information to help plan your NetWitness Platform deployment, and to tune it for best performance. As an Analyst, you can use this guide to structure your analysis in ways that will return reports faster. As a Content Developer, you can use this guide to help write content that will be processed efficiently by the database system.

NetWitness Platform Products Covered by this Guide

This guide covers the capabilities of NetWitness Platform. The following NetWitness Platform components contain the Core database:

- Concentrator
- Archiver
- Decoder
- Log Decoder
- Workbench

Frequently Used Terms

Definitions for terms that are used throughout this document are presented here. The terms are listed in the order in which they enter the NetWitness Platform system:

- **Packet DB** : The packet database contains the raw captured data. On a Decoder, the packet database contains packets as captured from the network. Log Decoders use the packet database to store raw logs. The raw data stored in the packet database is accessible by a Packet ID, however, this ID is typically never visible to the end user.
- **Packet ID** : A number used to uniquely identify a packet or log in a packet database.
- **Meta DB** : The meta database contains items of information that are extracted by a Decoder or Log Decoder from the raw data stream. Parsers, rules, or feeds can generate meta items.
- **Meta ID** : A number used to uniquely identify a meta item in the meta database.
- **Meta Key** : A name used to classify the type of each meta item. Common meta keys include ip.src, time, or service.
- **Meta Value** : Each meta item contains a value. The value is what each parser, feed, or rule generates.

- **Session DB** : The session database contains information that ties the packet and meta items together into sessions.
- **Session** : On a packet Decoder, a session represents a single logical network stream. For example, a TCP/IP connection is one session. On a Log Decoder, each log event is one session. Each session contains the references to all the Packet IDs and Meta IDs that refer to the session.
- **Session ID** : A number used to uniquely identify sessions in the Session DB.
- **Index** : The index is a collection of files that provides a way to look up Session IDs using Meta Values.
- **Core Database** : This refers to the combination of the Packet, Meta, Session, and Index.

For syntax definitions, this document uses EBNF grammar definitions.

NetWitness Core Database History

NetWitness developed the Core database for use in packet capture systems. Early in the history of NetWitness, developers identified that existing database technologies would not be able to keep up with the high ingest rate inherent in full packet capture. Contemporary database technologies were not anywhere close to being able to keep up with capturing the number of sessions received every second, much less sorting every packet. Likewise, the volume of data meant that packet storage would need to be discarded and reused just as quickly as it was consumed. This was also a weakness of databases at the time. Thus, NetWitness created a database consisting of the packet, session, and meta databases.

In order to provide the analytical capabilities of NetWitness Investigator, a meta index was added to the NetWitness database. The index shared the same design goals as the original databases. It was designed to sustain a very high insert rate into a high number of very large indices.

The index has evolved considerably over the years. Early versions of the index were only capable of providing summary estimates about how many unique meta values were present in the meta database. Other versions have had great challenges in meeting acceptable query performance. For example, NetWitness 9.0 more frequently measured report times in minutes rather than seconds. The current version of the index is derived from the NetWitness 9.0 index, but has evolved considerably in order to meet performance expectations and to add new features.

Core Database Strengths and Weaknesses

Strengths:

- High sustained insert rates, without needing down time for bulk inserts.
- Decent query performance simultaneous with high insert rates.
- Automatic cleanup and rollover of old data with minimal fragmentation.
- Extremely high number of meta value indices: more than 100 enabled by default on a Concentrator.
- Ability to scale to Petabyte database sizes and Terabyte index sizes within a single node.

- Using meta key-value pairs, it is very flexible for storing arbitrary meta items within a session. Thus a session can be used to represent nearly any kind of data record.

Weaknesses:

- The query functionality is limited and low level.
- The packet, meta, and session DB schema is fixed, and all customization is done through custom meta keys and values.
- The database provides no transaction atomicity guarantees as you might expect to find in a SQL database.

Basic Database Configuration

This topic covers basic database configuration settings of NetWitness Core services. For information on how to configure the Core services by editing configuration files, see "Service Configuration Settings" in the *Host and Services Getting Started Guide* .

This document assumes that the reader has some familiarity with adjusting the configuration of a NetWitness Core service. To use this document, you should be familiar with one of the mechanisms for modifying the configuration tree of Core services. Examples of such mechanisms include the Explorer view of the Administration pages within the NetWitness Platform user interface, or the REST interface accessible on each service through a web browser.

Find Help within the Core Service

Each configuration item within a Core service has a built-in help description of what the item does. You can view this help information by hovering your mouse over the configuration item in the Explorer view. Each configuration item also indicates whether it can be changed without restarting or if a restart of the service is needed for the change to take effect.

Developers using the REST API can retrieve the help text for each configuration item by sending the `help` message to the configuration node path.

Packet, Meta, and Session Storage

Each of the packet, meta, and session databases are configured through the `/database/config` folder on each NetWitness Core service. Each database has a configurable parameter to specify where the Core service stores data. Packet, meta, and session databases follow a predictable pattern for all of their configuration entries. Configuration items for the packet database start with the prefix `packet` , meta database configuration starts with the prefix `meta` , and the session database configuration items start with the prefix `session` .

Index Storage

The index configuration is stored in the `/index/config` folder on each Core service.

Topics

- Tiered Database Storage
- Manifests

Tiered Database Storage

This topic describes tiered database storage and provides recommendations for Hot, Warm, and Cold tier storage.

Starting with version 10.4, the Archiver service has the capability to be configured to use tiered storage. The concept of tiered storage is to put the most recent data on a Hot tier, which is the fastest storage available on the Archiver.

All services use the Hot tier by default.

The next tier is known as Warm and is typically cheaper and slower storage, such as a network-attached storage (NAS). The Warm tier contains older data; how old depends upon how much storage is allocated on the Hot tier and the average ingest rate. When the Hot tier reaches max utilization, the natural progression is to move the oldest data from the Hot tier to the Warm tier. When configured correctly, this happens automatically and is invisible to the end user. Queries and data access happen automatically no matter what tier (Hot or Warm) the data resides on. However, there can be a performance impact when accessing data on the Warm tier as compared to the Hot tier, because access times on the Warm tier are typically slower.

In addition to Hot and Warm, there is also a Cold tier. The Cold tier is only used as a staging area for offline backup. NetWitness Core services do not access data on the Cold tier. NetWitness Core services move the oldest data to the Cold tier and consider it abandoned (the service no longer accesses the data). This data can then be backed up to long-term storage like tape for possible restoration months or even years later, depending on requirements. The backing up and subsequent removal of data on the Cold tier must be handled outside of NetWitness Core services via scripts or other processes.

If the Cold tier becomes full because external processes are not removing data in a timely manner, this causes the NetWitness Core service to eventually stop the ingestion of new data until the problem is corrected.

When moving data to the Cold tier, NetWitness recommends that the directory remain on the same mount point as where it is being moved from. Therefore, if the files are coming from the Warm tier, it is far better for performance reasons to set the Cold tier directory on the same file system. The reason for this is that the service attempts to simply move the file and directory to the Cold tier, which is a nearly instantaneous operation on the same file system. If the move fails, the fallback is to copy the data to the Cold tier, which takes more processing time and causes additional I/O contention on the tier from which it is being copied.

Archiver

The tiers of storage capabilities are used by the Archiver. You can configure Archiver to only use Hot storage (the default), Hot and Warm, or all three (Hot, Warm and Cold). All services must use Hot, you cannot configure a service to only use Warm. Data flows from Hot to Warm and finally to Cold. You can also skip Warm and go from Hot to Cold. If Cold (offline) storage is not configured, the oldest data is deleted on the last configured tier, which has been the standard operating procedure.

The typical Archiver deployment sets all the databases to unlimited size (packet.dir, meta.dir, session.dir, index.dir, and optionally the Warm tier variants), which means that the size specifier is left off or set to zero. This lets the databases and index grow unbounded. Instead of each database managing their own size and rolling out only when each individual database exceeds their configured size, Archiver rolls out everything together using the `/index sizeRoll` command. This enables the databases and index to roll out in unison. For more information on sizeRoll, see "Asynchronous Rollover" in Rollover .

Archiver is typically configured to place the index, session, meta, and packet (log) DB on the same volume, instead of multiple volumes like a Concentrator or Decoder. Although this can potentially cause more I/O contention when concurrent reads happen across multiple databases, it also maximizes overall retention. Because all databases are on the same volume, they are configured to roll out together, which minimizes orphaning of data. Decoder and Concentrator are configured for maximum I/O speed, but can suffer from estimates on the proper volume sizing.

For example, if the session DB is too large, it may have enough storage for six months of retention, whereas the meta DB and index only have retention for four months. Because the session, meta DB, and index are intricately tied together, the shortest retention period for all three define the overall retention period (in this case, four months). Retention of individual databases is mostly affected by factors beyond our control, such as traffic captured, meta generated (parsers, feeds, rules) and filtering. The databases are easily resized by a simple configuration change, but this usually also involves changes at the hardware and file system level to adjust partitions, which complicates dynamic resizing. Archiver avoids these problems by using a single volume for everything, with the trade-off of somewhat slower I/O speed.

Manifests

This topic describes manifest files and provides an example manifest for a meta DB file. It also describes manifest searching and provides an example manifest search.

Manifest files are created with every session, meta, and packet (log) DB file and index slice directory. A manifest file is a file that describes several key pieces of information about the data to which it refers. Manifest files are written as a JSON record. Manifest files travel with the data they represent from tier to tier. If the data they represent is deleted, the manifest file is also deleted, except in the following special case. If the service has `/database/config/manifest.dir` configured to a valid directory, at the point when the manifest data is deleted, a copy of the manifest file is placed into the directory pointed at by `manifest.dir` (the directory is created if it does not exist). This enables a NetWitness Platform feature called historical manifest searching.

The intention of this process is to keep historical manifest files for years, in one location for offline querying. As you might imagine from a service running for many years, this can potentially generate hundreds of thousands of files. This should not be a concern however, as the service automatically compresses files into a single archive in order to save space when they grow too numerous. Manifest files are very small and compress well.

Example manifest (`meta-000000023.nwmdb.manifest`) for a meta DB file:

```
{
    "filename" : "meta-000000023.nwmdb",
    "size" : 185153768,
    "fileTime" : 1403903940,
    "id1" : 150814110,
    "id2" : 159341086,
    "session1" : 4023382,
    "session2" : 4250442,
    "time1" : 1403903879,
    "time2" : 1404739851
}
```

```
filename  = The filename for the db file the manifest represents
size      = The size in bytes of the db file
fileTime  = The time the file was created
id1       = The starting id in the file (for this example, the starting meta
ID)
id2       = The last id in the file (for this example, the last meta ID)
session1  = The starting session ID of the first meta in the file
session2  = The last session ID of the last meta in the file
time1     = The POSIX time of the first "time" meta found in the file
time2     = The POSIX time of the last "time" meta found in the file
```

In this example manifest, the most important fields are `fileTime` , `time1` and `time2` . All three fields are written in POSIX time. `time1` and `time2` are the starting and stopping times of the meta recorded in the meta DB file `meta-000000023.nwmdb` . In particular, `fileTime` is always the time in which the file was created (not last modified). `time1` and `time2` are representative of the min and max range of the parsed data within the meta DB file. When doing historical searches by time, `time1` and `time2` are preferred over `fileTime` , when they are present. Manifest files for the other databases and index contain some different fields, but all have enough information to perform time based queries.

Search Historical Manifests

When manifests are collected in the directory pointed to by `manifest.dir` , it is assumed that the data they refer to was copied to the Cold tier and eventually backed up to offline storage. Because the historical manifests are still accessible by the service, this allows time-based queries to be performed on offline data, in order to determine what data needs to be restored for a given time range.

You can search manifests using the `/database manifest` command:

```
manifest: If a manifest directory is defined, it will allow operations on the
manifest files (such as a time based query) for database files in cold
storage.
security.roles: database.manage
parameters:
   op - <string, optional, {enum-one:query|compress}> The operation to perform
(defaults to query)
   time1 - <date-time, optional> The beginning time (UTC) for matching offline
database files
   time2 - <date-time, optional> The ending time (UTC) for matching offline
```

```
database files
   timeFormat - <string, optional, {enum-one:posix|simple}> Specify the time
format that is returned (posix, simple), default is posix
```

Example search:

```
/database manifest time1="2014-04-20 11:00:00" time2="2014-04-11
11:20:00" timeFormat=simple
```

The search returns all manifests that match the query:

```
[ filename=meta-000001691.nwmdb size=4843826176 fileTime="2014-Apr-20
11:06:34" id1=301555027452 id2=301733101896 session1=15352020201
session2=15361024200 time1="2014-Apr-20 11:05:34" time2="2014-Apr-20 11:16:34"
compression=gzip ]
[ filename=session-000001865.nwsdb size=268439552 fileTime="2014-Apr-20
11:06:35" id1=14674145801 id2=14682041000 metaId1=288217522208
metaId2=288370660984 packetId1=11733872441 packetId2=11741745303 ]
[ filename=session-000001866.nwsdb size=268439552 fileTime="2014-Apr-20
11:18:31" id1=14682041001 id2=14689936200 metaId1=288370660985
metaId2=288520616949 packetId1=11741745304 packetId2=11749618589 ]
```

The returned results can be used to correlate which files should be restored from backup for the given time range. For versions 10.4 and later, a service called Workbench can be used to take the restored files and provide a query interface over the restored data using one or more collections.

Setup of the Workbench service is beyond the scope of this document. For more information, see "Configure Data Backup and Restore" in the *Archiver Configuration Guide* .

Advanced Database Configuration

This topic explains the advanced configuration options of the NetWitness Core database.

The configuration options of the NetWitness Core database may change from one release to the next. However, many of the configuration items do not change frequently and are documented here. This is not an exhaustive list, since new features are added in every release, and they may require new configuration items. For the most up-to-date documentation, refer to the built-in help functionality of the NetWitness Core service.

Topics

- Database Configuration Nodes
- Index Configuration Nodes
- SDK Configuration Nodes
- Per-User Configuration Nodes
- Scheduler
- Rollover
- Snort Rules and Configuration

Database Configuration Nodes

This topic describes database configuration nodes. The following database configuration nodes are some of the advanced database configuration items of the NetWitness Core database that do not change frequently.

packet.dir,meta.dir,session.dir

This is the primary configuration entry for each database (also known as the Hot tier). It controls where in the file system the respective databases are stored. This configuration entry understands a complex syntax for specifying many directories as storage locations.

Configuration syntax:

```
config-value = directory, { ";" , directory } ;
directory    = path, [ ( "=" | "==" ) , size ] ;
path         = ? linux filesystem path ? ;
size         = number size_unit ;
size_unit    = "t" | "TB" | "g" | "GB" | "m" | "MB" ;
number       = ? decimal number ? ;
```

Example:

```
/var/lib/netwitness/decoder/packetdb=10
t;/var/lib/netwitness/decoder0/packetdb=20.5 t
```

The size values are optional. If set, they indicate the maximum total size of files stored there before databases roll over. If the size is not present, the database does not automatically roll over, but its size can be managed using other mechanisms.

The use of = or == is significant. The default behavior of the databases is to automatically create directories specified when the Core service starts. However, this behavior can be overridden by using the == syntax. If == is used, the service does not create any directories. If the directories do not exist when the service starts, the service does not successfully start processing. This gives the service resilience against file systems that are missing or unmounted when the host boots.

If you modify the size of a directory in use, the size takes effect immediately, as long as it is larger. If the size is smaller, it is ignored if it is more than 10 percent smaller than the existing size. This prevents an accidental mistype that causes a enormous loss of data. For example, if the packet database was configured for `12 TB` and someone mistyped it as `12 GB`, the database would end up deleting over 11 TBs of data in order to shrink it down to just 12 GB. Instead, the database ignores the `12 GB` setting and logs a warning, so that the error can be caught quickly. Of course, if the size specified is actually correct and more than a 10 percent difference from the existing size, the only recourse for it to take effect is to restart the service. When it starts back up, it assumes the size is correct and adjusts the database to the new size by rolling out the oldest data until the new size is reached. If you actually do want to adjust the size downward and by more than 10 percent without restarting the service, you need to modify the size multiple times, each time adjusting it by less than 10 percent. Watch the service logs to know when the database has adjusted to the new size, as it only adjusts the total database size when the latest file being written has been closed.

If new directories get added or deleted (semicolon separated), they do not take effect until the service restarts.

packet.dir.warm, meta.dir.warm, session.dir.warm

These settings are optional and are used for Warm tier storage on an Archiver. By default, they are blank and unused. If configured, they follow the same format and behavior as `packet.dir`, `meta.dir`, and `session.dir` (see _ `packet.dir`, `meta.dir`, and `session.dir` _ above). When configured, the oldest file on the Hot tier moves to the Warm tier when no available space remains in the Hot tier.

packet.dir.cold, meta.dir.cold, session.dir.cold

These settings are optional and are used to move files from either a Hot or Warm tier storage system to the Cold tier directory specified. Specifically, this setting is nothing more than a directory, there are no size specifiers. However, the defined path name has a few special format specifiers that you can use to name the directory with the date of the data in it.

```
%y   = The year of the data being moved to the cold tier
%m   = The month of the data being moved to the cold tier
```

```
%d   = The day of the data being moved to the cold tier
%h   = The hour of the data being moved to the cold tier
%##r = A block of time within a day. So %12r would create two blocks, 00 and
01\. 00 for all data in the AM, 01 for all PM data
```

Example setting:

```
packet.dir.cold = /var/lib/netwitness/archiver/database1/alldata/cold-
storage-%y-%m-%d-%8r
```

For the setting above, if a log database file was about to be moved to cold storage and it was created on `2014-03-02 15:00:00` , it would be moved to the following directory on the Cold tier:

```
/var/lib/netwitness/archiver/database1/alldata/cold-storage-2014-03-02-01
```

The last number `01` needs some explanation. The `%8r` specifier breaks the hours of the day into 24 / 8 = 3 parts. The first eight hours of the day would be block `00` , so 12 a.m. to 8 a.m. The next eight hours are from 8 a.m. to 4 p.m. and are assigned block `01` . Since the data being moved to cold storage was created at 3 p.m., it falls into block `01` . The `%r` format specifier is useful for backing up files with a granularity somewhere between a day `%d` and a single hour `%h` . The Cold storage directory is created on demand and is defined by the data being moved when the format specifiers are used.

The ability to add a date to the path of the data is just a convenience added for backup and restore. It is a way of tagging the data with a date in the path.

packet.file.size , meta.file.size , session.file.size

This controls the size of the files created with each database. It is normally not necessary to change these values as the default values typically work well. This setting takes effect immediately for subsequent files.

packet.files , meta.files , session.files

This setting controls the number of files held open by the database. You can increase this value to improve performance: however, the operating system has an overall limit on the number of files that service can keep open. If this limit is exceeded, an error is reported and the service does not function. This setting takes effect immediately.

In versions 10.6 and later, the default value for packet.files, meta.files, and session.files is `auto` and the service manages the number of open files based on this criteria:

1. Number of collections
2. Amount of system memory

When set to `auto` , the number is dynamic and you can view it in the logs when it changes. NetWitness recommends that you leave this value as `auto` and do not change it to a specific number.

packet.free.space.min , meta.free.space.min , session.free.space.min

This setting provides a safety limit on the minimum free space that exists on the paths specified by the packet.dir, meta.dir, and session.dir directories, respectively. This setting is used to prevent the service from running out of space in the event that other programs have filled up the space that should be dedicated to each of the databases. This setting takes effect immediately.

packet.index.fidelity , meta.index.fidelity

This setting controls how frequently packet ID locations and meta ID locations are indexed. This setting can be increased to reduce the amount of space needed by each packet or meta nwindex file, but increasing the setting reduces the speed at which individual packets or meta items can be located. This setting takes effect immediately.

The session database does not have a fidelity setting because it does not generate index files.

packet.integrity.flush , meta.integrity.flush , session.integrity.flush

This setting controls whether the database forces a sync operation on the file system when it is finished writing a file. The default value is `sync` , which means when a file is closed there will be a significant delay while the data writes to non-volatile storage. It may be necessary to set this to `normal` in order to achieve higher sustained write rates, especially on a Decoder. This setting takes effect on the next file created. Therefore, it is expected that at least one more sync will happen if the value was just changed to `normal` .

If packet drops are occurring and packet.integrity.flush is set to `sync` , set it to `normal` and monitor. Keep the session and meta flush settings on `sync` . If packet drops are still problematic, then set all three to `normal` and monitor.

packet.write.block.size , meta.write.block.size , session.write.block.size

The block size represents how much data is allocated at a time within each database file. Larger block sizes can provide higher throughput and compression ratios, and can improve the rate at which items can be retrieved from the database sequentially. However, larger block sizes have a detrimental impact on random read speed for compressed packet and meta items. This setting takes effect immediately.

packet.compression , meta.compression

These parameters control whether the databases compress data. Compression reduces the amount of storage needed by each database, but it can have a major detrimental impact on the speed at which items are written to the database, and the speed at which items are retrieved from the database. Changes take effect immediately on the next file creation. Make a note that you can not compress the `pcapng` packet db format.

As of version 10.4, the valid values for this parameter are `gzip` , `bzip2` , `lzma` , or `none` . `gzip` is the preferred algorithm when compression is used, because it provides a good balance between performance and space savings. Both `bzip2` and `lzma` can achieve better space savings, but the tradeoff in speed is substantial and likely should only be considered for low ingest speeds and when storage space is at a premium.

packet.compression.level , meta.compression.level

You can use these settings to further refine how the compression algorithms behave. They have no effect when compression is disabled. The valid values are between 0–9\. The default value of zero means let the software pick the best setting for speed and compression. The values between 1 and 9 are used as a sliding scale between performance (1) and compression (9). The value of 9 typically gives you the best compression for a given algorithm, but the worst performance. Somewhere in the middle is usually the best setting, which is what zero picks.

hash.algorithm

This setting controls how the database files are hashed. The default value is `none` , so no hashing is performed. The valid values are `none` , `sha256` , `sha1` , or `md5` . Database files can be hashed to provide evidence that they have not been tampered with since they were closed. Hashing is time intensive and affects ingest performance when enabled. This change takes effect immediately.

hash.databases

This setting controls which databases are hashed. Valid values are `session` , `meta` , and `packet` and are comma separated when hashing multiple databases. This change takes effect immediately.

hash.dir

This setting is normally empty, which means the hash file is created in the same directory as the database file that was hashed. If this setting is defined, the hash file is written to the directory specified instead. This could be some form of write-once storage for resilience against hash tampering.

Hash files are small XML files containing the hex encoded hash along with metadata about the database file that was hashed.

Index Configuration Nodes

This topic describes index configuration nodes. The following index configuration nodes are some of the advanced database configuration items of the NetWitness Core database that do not change frequently.

index.dir

The index.dir setting controls where the files used by the index are stored. This setting supports the same syntax as the packet.dir , meta.dir , and session.dir settings.

index.dir.warm

The Warm tier storage for index slices. This setting supports the same syntax as packet.dir.warm , meta.dir.warm , and ession.dir.warm .

index.dir.cold

The Cold tier storage for index slices. This setting supports the same syntax as packet.dir.cold , meta.dir.cold , and session.dir.cold .

index.slices.open

This setting controls the number of index slices held open by the index. Index slices are opened automatically as needed by queries. When queries complete, the index engine may hold the slices open so that subsequent queries execute faster. The most recently created slices are the slices that will be held open, since they are mostly likely to be used by queries.

If queries against the index require the index to open slices, then they will execute slower than if the slices were already open. Therefore, this parameter should be tuned such that most queries executed against the index will work on open slices. However, each open index slice consumes some resources, such as file handles and memory. If there are too many index slices open, the overall performance of the service can suffer.

You should set this parameter so that the open index slices will cover most of the time ranges that most queries will need. For example, if most queries are over the past two weeks, and there are index slices created every 8 hours, then there are 14 days x 3 slices per day, or 42 slices created over the past two weeks. Thus, you could set index.slices.open to 42 so that only slices that are likely to be used are held open.

If this parameter is set to 0, then all slices are held open until the next index save. In this scenario, the only thing limiting the number of slices open in the process is the number of slices in the index.

page.compression

Deprecated. Versions of the NetWitness Core index between 9.8 and 10.2 supported two different index compression algorithms, and you can choose between them using this setting. As of 10.3, the only recommended value is the default of huffhybrid .

save.session.count

This setting controls how often the index is automatically saved when new sessions are inserted. If the value of save.session.count is greater than 0, any time more than save.session.count sessions are added to the index, the index automatically saves itself. If the save.session.count is set to 0, this feature is disabled and the index will not automatically save itself when new sessions are added to the index.

`save.session.count` can be used to implement an automatic save pattern that is based on the volume of data that enters the index. This is useful because it allows a lightly loaded system to generate save points less often.

For more information on the topic of index saves, see the section in this guide on Optimization Techniques .

reindex.enable

This setting controls the operation of the background reindexer .

SDK Configuration Nodes

This topic describes the SDK configuration nodes that affect the database. There are some additional configuration items in each Core service that affect the database, but do not actually affect how the database stores or retrieves data. These settings exist in the `/sdk/config` folder.

max.concurrent.queries

This setting controls how many query operations are allowed on the database simultaneously. Allowing more simultaneous query operations can improve overall responsiveness for more users, but if the query load of the Core service is very I/O bound, having a high max.concurrent.queries value can have a detrimental effect. The recommended value is near the number of cores on the system, including hyper threading. Thus, for an appliance with 16 cores, the value should be somewhere close to 32\. Subtract a few for aggregation threads and general system response threads. Subtract a few more if this is a hybrid system (for example, both a Decoder and Concentrator running on the same appliance). There is no magic number, but somewhere between 16 and 32 should work well.

max.pending.queries

This setting controls the backlog size for the query engine of the database. Larger values allow the database to queue more operations for execution. A queued query does not make progress on its execution, so it may be more useful to make the system produce errors when the queue is full, rather than allowing the queue to grow very large. However, on a system that is primarily performing batch operations such as reports, there may be no detrimental effect to having a large queue.

cache.window.minutes

This setting controls a feature of the query engine that is intended to improve query responsiveness when there are a large number of simultaneous users. For more information on cache window, see Optimization Techniques .

max.where.clause.cache

The where clause cache controls how much memory can be consumed by query operations that need to produce a large temporary data set to evaluate sorting or counting. If the where clause cache size is overflowed, the query still works, but it is much slower. If the where clause cache is too large, it is possible for queries to allocate so much memory that the service would be forced into swap or run out of memory. Thus, this value multiplied by the max.concurrent.queries should always be much less than the size of physical RAM. This setting understands sizes in the form of a number followed by a unit, for example `1.5 GB`.

max.unique.values

The maximum unique values limits how much memory can be consumed by the SDK Values function. SDK Values produces a sorted list of unique values. In order to produce accurate results, it may need to merge together large numbers of unique values from many slices. This merged set of values must be held in memory, so this parameter exists to put a limit on how much memory the merged value set can consume. The default value will limit memory usage to approximately 1/10th of total RAM.

query.level.1.minutes, query.level.2.minutes, query.level.3.minutes

These settings are available in 10.4 and earlier versions.

In versions 10.4 and earlier, the Core database supports three query priority levels. Each user is assigned to one of the priority levels. Therefore, there are up to three groups of users that can be defined for the purposes of performance tuning. These settings control how long each user level is allowed to execute the queries. For example, lower privileged users may have a lower value so that they are not able to use all the resources of the Core service with long-running queries.

query.timeout

This setting is available in 10.5 and later versions.

Query levels have been replaced in versions 10.5 and later with per user account query timeouts. For trusted connections, these timeouts are configured on the NetWitness Platform server. For accounts on Core services, there is a new config node under each account called `query.timeout`, which is the maximum amount of time in minutes that each query can run. Setting this value to zero means no query timeout will be enforced by the Core service.

max.where.clause.sessions

This setting will be deprecated in a future release. Use `max.query.memory` to limit overall query memory usage.

This setting imposes a limit on how many sessions can be scanned by a single query. For example, if a user selects all meta from the database, the database stops processing results once the number of sessions read for the query reaches this configuration value. The value of 0 disables this limit.

The number of sessions needed to fully process a query is equal to the number of sessions that match the WHERE clause of the query, assuming that all terms in the where clause have a suitable index. If there are terms in the where clause that are not indexed, the database has to read more sessions and meta, and reaches this limit sooner.

max.query.groups

This setting will be deprecated in a future release. Use `max.query.memory` to limit overall query memory usage.

This setting imposes a limit on the number of unique groups collected in a single query. For example, if a query has a group by clause with multiple metas that have high unique value counts, the amount of memory needed for that query could easily outpace the amount of RAM available on the server. Thus, this limit exists to prevent out-of-memory conditions from happening.

Setting a value of 0 disables this limit.

packet.read.throttle

This is a decoder-only setting that affects the access to the packets database. When packet.read.throttle is set to a value greater than 0, the decoder attempts to throttle packet reads when it detects packet contention on the packet database. Higher numbers provide more throttling. Changes takes effect immediately.

cache.dir , cache.size

All NetWitness Platform Core services maintain a small file cache of raw content extracted from the device. These parameters control the location (cache.dir) and size (cache.size) of this cache.

pin.cache.dir , pin.cache.size

All NetWitness Platform Core services provide a separate file cache of raw content that is marked (or pinned) for long-term retention. These parameters control the location (pin.cache.dir) and size (pin.cache.size) of this cache. You can configure the maximum size of the pinned cache at `/sdk/config/pin.cache.size`. The default size is zero that prevents pinning of new sessions once storage decreases below 100 MB. If you configure it with a non-zero value, then it will reduce the cached files when the total size exceeds the configured size by deleting the oldest pinned sessions.

parallel.values

This setting is available in 10.5 and later versions.

This setting allows SDK-values operations to be executed in parallel. If this is set to 0, it will disable parallel execution. If it is set to a value greater than 0, it represents the number of threads created when each SDK-values operation is executed. The maximum value is the number of logical CPUs available when the process started.

Setting a higher value for parallel.values is useful when there are small numbers of simultaneous users, since it will allow for more complex Investigations to be executed more quickly. If there are many simultaneous users, it is better to use a low value here, since there will be many independent SDK-values operations executed simultaneously.

parallel.query

This setting is available in 10.5 and later versions.

This configuration is similar to the parallel.values setting in that the maximum value is the number of logical CPUs. Setting parallel.query to a specific value should take into account the number of simultaneous users to maximize CPU utilization without consistently exceeding available resources.

Setting a higher value for parallel.query is useful when there are small numbers of simultaneous users and queries, since it will allow more complex queries to be executed more quickly. If there are many simultaneous users and queries, it is better to use a low value, since there will be many independent SDK-query operations executed simultaneously.

Query operations are limited by the meta database read rate, so setting parallel.query to a value higher than 4 is unlikely to produce dramatically better results than the default value of 0\. The best number to use for parallel.query will depend on the type of storage attached. Experiment with different values of parallel.query to determine the best results for your storage system.

Per-User Configuration Nodes

This topic describes the per-user configuration nodes. There are settings that influence the actions users are allowed to perform on the database. These settings are stored in the configuration tree at `/users/accounts/<username>/config` , where `<username>` is the name of the user to which the settings apply.

query.prefix

A query prefix applies a filter to every query operation that the user performs. This is implemented by taking the `query.prefix` values and appending it the the where clause of each query using the logical && (and) operator. For more information on Where Clauses, see Queries .

query.level

This setting is available in 10.4 and earlier versions.

The query.level setting assigns the query level that the users have for every query they perform. These influence whether their queries are limited by the query.level.1.minutes, query.level.2.minutes, or query.level.3.minutes.

query.timeout

This setting is available in 10.5 and later versions.

The query.timeout setting assigns the maximum amount of time in minutes that a user can run each query. For trusted connections, these timeouts are configured on the NetWitness Platform server. For accounts on Core services, this setting is stored in the configuration tree at `/users/accounts/<username>/config` , where `<username>` is the name of the user to which the setting applies. When this value is set to zero, the Core service does not enforce the query timeout.

session.threshold

The session.threshold setting assigns a maximum session threshold for the user. If set, this threshold value is assigned to all values calls that the user performs. A detailed discussion of both the values call and thresholds is covered in this guide.

Scheduler

This topic provides a brief introduction to the scheduler and explains how to schedule commands. All NetWitness Core services come with a built-in scheduler found under `/sys/config/scheduler`. To use the scheduler, you add the command you want to run periodically using one of two messages:

`/sys/config/scheduler addInter` - Add a command to run at the specified interval (every N hours, minutes or seconds)

or

`/sys/config/scheduler addMil` - Add a command to run at the specified time of day or even specific days of the week

Example

For example, suppose that you have a use case to delete all packet data that is greater than seven days old. Since you cannot configure the packet.dir setting to rollout data based on a time interval, you need to schedule the `/database timeRoll` command to run every so often. For this example, create a timeRoll to run every 20 minutes:

```
addIter minutes=20 pathname=/database msg=timeRoll params="type=packet days=7"
```

This command adds a scheduled task (it is persisted between restarts of the service) to run every 20 minutes, on the `/database` node, and ages out all packet data older than seven days. The `params` parameter is used to pass all the parameters to the command specified (in this case `timeRoll`). Notice how it quotes all the embedded parameters (`type` and `days`) so they are not interpreted as parameters to be passed to the outer `addIter` command. If the parameters inside `params` need to use quotes, you must escape the inner quotes with a backslash. You can rewrite it with embedded quotes, which does not alter the command in any way:

```
addIter minutes="20" pathname="/database" msg="timeRoll"
params="type=\"packet\" days=\"7\""
```

This command works identically to the original, but demonstrates how to escape complicated parameter passing. Additional useful scheduler commands are:

`/sys/config/scheduler print` - Print all scheduled commands (you can also see them by doing an `ls` on the scheduler node).

`/sys/config/scheduler delSched` - Delete a scheduled command by passing in the identifier shown in the `print` (or `ls`) command.

This is a brief introduction to the scheduler. For more information on command parameters, send the `help` message to the scheduler node and pass in the command name via the `msg` parameter. For more information, see the "Services Explore View" topic in the *Host and Services Getting Started Guide* .

Rollover

This topic describes the two rollover mechanisms. The database operates as a first-in, first-out (FIFO) queue. New data is always appended to the database, and the oldest data is automatically removed as needed. Data that is in the middle of the database is immutable, meaning it cannot be modified.

There are two mechanisms to for rollover: synchronous and asynchronous.

Synchronous Rollover

Synchronous rollover refers to rollover settings that are applied in response to a write operation on the database. That means data is removed from the database in direct response to the need to write new data. Synchronous rollover is configured by setting size values on the configuration for `packet.dir` , `meta.dir` , `session.dir` , and `index.dir` .

Synchronous rollover on the packet, meta, and session databases can occur within any write operation. Synchronous rollover on the index occurs when the index is saved.

Asynchronous Rollover

Asynchronous rollover refers to database file removal that occurs when an explicit rollover command is issued to the database. Most commonly this type of rollover is scheduled to run periodically using the built-in scheduler of the Core service. The user can also explicitly request it.

The asynchronous rollover command is the `sizeRoll` message present on the `/index` and `/database` nodes of the configuration tree. The message on the `/database` node does size rollover on packet, meta, and session databases only, while the message on the `/index` node can do simultaneous rollover on both the index and the packet, meta, and session databases.

The `sizeRoll` command has the following parameter syntax:

```
size-roll-params    = {type-param, space}, (max-size-param | min-free-param |
max-percent-param), {max-size-warm-param, space}
type-param          = "type=", {type-flag} , { ",", type-flag } ;
type-flag           = "packet" | "meta" | "session" ;
max-size-param      = "maxSize=", number, {space}, unit ;
max-percent-param   = "maxPercent=", number, {space}, unit ;
min-free-param      = "minFree=", number, {space}, unit ;
max-size-warm-param = "maxSizeWarm=", number, {space}, unit ;
unit                = "t" | "TB" | "g" | "GB" | "m" | "MB" ;
```

```
number              = ? decimal number ? ;
percentage          = ? number between 0 and 100 ? ;
```

The `type` parameter controls the databases to consider for removing the oldest data based on total size or space remaining. If type is not specified on the `/index` sizeRoll, only the index is considered for rollover operations.

The `maxSize` parameter sets a current maximum size of the database or index. If the database is larger than this size, oldest data is deleted first (or moved to the Warm or Cold tier, depending on the configuration) until total size is less than `maxSize` . The `sizeRoll` operation determines which data is oldest out of all the databases and the index based on session IDs. Sessions or index entries with lowest session IDs are deleted first, possibly including removing meta and packet databases that are orphaned by removing entries from the session database. The index data is rolled out if the sessions that it refers to are removed.

The `maxSizeWarm` parameter sets a current maximum size on the *Warm* tier, but otherwise behaves identically to the `maxSize` parameter. When data is rolled out on the Warm tier, it is moved to the Cold tier (if configured) or deleted.

The `maxPercent` parameter sets a maximum percentage of all the volumes of all databases passed in `type` parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes.

The `minFree` parameter sets a minimum allowed free space on the volumes before oldest data is deleted.

Each call to the sizeRoll operation provides a single pass through the database to delete files. When the operation completes, the current size utilization of the database will have met the criteria specified by the `maxSize` , `maxPercent` , or `minFree` parameters and the optional `maxSizeWarm` . Therefore, this operation can be scheduled periodically to ensure that the database can continue to operate uninterrupted.

Example

The following example shows a typical sizeRoll scheduler entry for an Archiver:

```
pathname=/index minutes=5 msg=sizeRoll params="type=meta,session,packet
maxSize=25TB maxSizeWarm=150TB"
```

This scheduler entry specifies that every five minutes the database ensures that the max size of the meta, session, packet, and index does not exceed 25 terabytes on the Hot tier and does not exceed 150 terabytes on the Warm tier.

Snort Rules and Configuration

Snort® rules and configuration are added to the parsers/snort directory for Investigator and Decoder. Decoder supports the payload detection capabilities of Snort rules. The rules files must have the extension `.rules` and the configuration files must have the extension `.conf` . The Decoder implementation of Snort rules is centered on using the content strings defined in a Snort rule as a token. Once a token is matched, the rule header and additional rule options can be evaluated. Currently, rules that do not define any content (via `content` or `uricontent` rule options) are not supported.

Configuration

The configuration files are loaded prior to loading rules.

Variable Definitions	Description
`ipvar`	The full language for defining IP Address variables is supported, including lists, CIDR, and negation.
`portvar`	The full language for defining IP Address variables is supported, including lists, ranges, and negation.
`var`	Not supported; use `ipvar` or `portvar` .

Action Definitions	Description
`ruletype`	The definition of additional `ruletype` s is supported. However, only rules that have a base rule type of **alert** are supported.

General Configuration	Description
`nopcre`	This configuration option disables all rules with **pcres**.

Meta key usage

Starting with the 11.3 release, an attempt has been made to align better the Snort parser's meta key usage with that of other parsers. As of the 11.3 release, the default mode operation continues to write to the legacy key set (consistent with previous releases). To use the aligned key set, set the `udm` option to `true` for the Snort parser in the `parser.options` configuration node. Refer to the **General options** section, below, for a description of how the two modes differ.

Rules

- Any rule that does not properly parse is ignored.
- Any valid Snort rule should successfully parse; however,there are rule options that are not supported by Decoder that are not fully parsed.

Snort rules are parsed and loaded when PCS is loaded (any import/capture in Investigator, initial capture start and parser reload in Decoder).

Section	Description
Header	The header conditions are evaluated when a rule receives the first token callback for a stream. The header is evaluated once per stream, and prevents any further consideration of a rule against a specific stream if the conditions are not met.
Actions	The specified action or a rule must be defined (either one of the native Snort actions or defined in the configuration using the `ruletype` statement) for the rule to be considered valid. Decoder only utilizes rules with alert actions.
Protocols	Decoder supports the current Snort protocol keywords (`tcp` , `udp` , `icmp` , `ip`).
IP Addresses	The full language for defining IP addresses is supported, including lists, CIDR, and negation.
Port Numbers	The full language for defining port numbers is supported, including lists, ranges and negation.
Direction Operator	The directional operator supports the from-to (`->`) and bidirectional (`<>`) values. The to-from (`<-`) value is invalid and will cause the rule to fail to load.

General options

Snort rule general options can result in different meta keys being written depending on whether the Snort parser is in legacy mode or not.

Aligned key mode:

Option	Description
`msg`	If the rule matches, the `msg` value is added as `sig.name` meta.
`sid`	If the rule matches, the `sig.id` value is added as meta.
`classtype`	If the rule matches, the `classtype` name is added as `threat.cat` meta.

Option	Description
`priority`	If the rule matches and it has a `priority` option, it is used for the value of the `risk.num` meta.

Legacy key mode:

Option	Description
`msg`	If the rule matches, the `msg` value is added as `risk.info` , `risk.warning` , or `risk.suspicious` meta, depending on rule priority.
`sid`	If the rule matches, the `sid` value is added as meta.
`classtype`	If the rule matches, the `classtype` name is added as `threat.cat` meta.
`priority`	If the rule matches and it has a `priority` option, it is used to determine the type of risk meta associated with the `msg` value.

Payload options

Decoder supports the following payload rule options:

Option	Description
`content`	The `content` option creates a token for Decoder to match. Only tokens of three or more bytes are accepted. It is also important to note that Decoder differs from Snort in that rules are evaluated across the payload of the reconstructed stream and not just a single packet. This can result in differences in rules matches between Snort and Decoder, especially when considering positional options.
`nocase`	Currently not supported. This option is ignored and case-sensitive matching is used.
`depth`	This option is applied to the distance of the token from the beginning of the stream. If the token position is greater than this value, it is not a match.
`offset`	This option is applied to the distance of the token from the beginning of the stream. If the token position is less than this value, it is not a match.
`distance`	This option is applied to the distance of the token from the end of the previous token match. If the relative token position is less than this value, it is not a match.
`within`	This option is applied to the distance of the token from the end of the previous token match. If the relative token position is greater than this value, it is not a match.

Option	Description
`http_uri`	Any token that hits is verified to fall within an `http_uri` as indicated by the HTTP parser. No URI normalization is applied.
`uricontent`	There is no URI normalization applied. Otherwise, this is equivalent to the content option with the `http_uri` modifier.
`pcre`	Currently, PCREs are only applied to URIs and must specify the U option.

Non-payload options

Option	Description
`flow`	Verifies that the rule is only applied to the client or server stream.
`to_client`	Limits the rule to only matching on a stream that Decoder has defined as Server.
`from_server`	Synonym for `to_client` .
`from_client`	Limits the rule to only matching on a stream that Decoder has defined as Client.
`flowbits`	Maintains state per session and are reset at the end of each session.
`set`	When the rule matches, the specified flowbit is set.
`unset`	When the rule matches, the specified flowbit is cleared.
`toggle`	When the rule matches, the specified flowbit is flipped.
`isset`	When the rule is evaluated, the specified flowbit state must be set for the rule to match.
`isnotset`	When the rule is evaluated, the specified flowbit state must not be set for the rule to match.
`noalert`	Prevents the rule from generating meta data if it matches.

Queries

This topic covers the database query syntax. There are three main mechanisms for performing queries in the database, the `query` , `values` , and `msearch` calls on the `/sdk` folder on each Core service.

The `query` call returns meta items from the meta database, possibly using the index for fast retrieval.

The `values` call returns groups of unique meta values sorted by some criteria. It is optimized to return a subset of the unique values sorted by an aggregate function such as count.

The `msearch` call takes text search terms as it's input, and returns matching sessions that match the search terms. It can search within indexes, meta, raw packets, or raw logs.

`query` Syntax

The `query` message has the following syntax:

```
query-params     = size-param, space, query-param, {space, start-meta-param},
{space, end-meta-param}, {space, search-param} ;
size-param       = "size=", ? integer between 0 and 1,677,721 ? ;
query-param      = "query=", query-string ;
start-meta-param = "id1=", metaid ;
end-meta-param   = "id2=", metaid ;
search-param     = "search=", search-string ;
metaid           = ? any meta ID from the meta database ? ;
```

The `id1` , `id2` , and `size` parameters form a paging mechanism for returning a large number of results from the database. Their usage mostly benefits developers who are writing applications directly against the NetWitness Core database. Normally, results are returned in the order of oldest to newest data (higher meta IDs are always more recent). In order to return results from most recent to oldest, reverse the IDs such that `id1` is larger than `id2` . This has a slight performance penalty, because the where clause must be completely evaluated before processing in reverse order can begin.

When size is left off or set to zero, the system streams back all results without paging. For the RESTful interface, this results in the full response to be returned with chunked-encoding. The native protocol returns the results over multiple messages.

The `query` parameter is a `query` command string with its own NetWitness-specific syntax:

```
query-string      = select-clause {, where-clause} {, group-by-clause {,
order-by-clause } } ;
select-clause     = "select ",  ( "*" | meta-or-aggregate {, meta-or-
aggregate} ) ;
```

```
where-clause      = " where ",  { where-criteria } ;
meta-or-entity    = (meta_key | entity) ;
meta-or-aggregate = meta-or-entity | aggregate_func, "(", meta-or-entity, ")"
;
aggregate_func    = "sum" | "count" | "min" | "max" | "avg" | "distinct" |
"first" | "last" | "len" | "avglen" | "countdistinct" ;
group-by-clause   = " group by ", meta-key-list
meta-key-list     = meta-or-entity {, meta-key-list}
order-by-clause   = " order by ", order-by-column
order-by-column   = meta-or-aggregate { "asc" | "desc" } {, order-by-column}
```

The `select` clause allows you to specify either `*` to return all the meta in all the sessions that match the where clause, or a set of meta field names and aggregate functions to select a subset of the meta with each session.

The `select` clause may contain entity names in the place of meta key names. If an entity name is in the select clause, meta items returned by the query will have their key name set to the entity name, rather than their actual meta key name stored in the session. Thus, the names of the meta items returned in the query will match the names of the metas in the select clause. For example, if there is an entity `ip` that consists of `ip.dst` and `ip.src` , then a query containing `select ip` will only return `ip` fields, with nothing to distiguish `ip.dst` meta items from `ip.src` meta items in the result set.

The `select` clause may contain renamed meta key names. Any fields appearing in the result set as a result of a renamed key in the `select` clause will be returned with the meta key name matching the name used in the `select` clause. For example, if the key `port_src` is used to rename `tcp.srcport` , then a query containing `select port_src` will only return `port_src` fields, even if the underlying meta had type `tcp.srcport` .

Note: Usage of renamed meta key pairs in the `select` clause cannot be combined with fixed-size result paging for a query. Doing so causes discrepancies in the results returned to Brokers. The reason for the discrepancies is that Concentrators cannot return only one of the key values of a renamed meta key pair and still preserve the correctness of the result set for the requested size. Hence, the Concentrator omits renamed meta key pair results to preserve the correctness of the result set, which causes the Broker to pull the result from the next Concentrator and advance the IDs that are returned.

Example: `select ip.proto,ipv6.proto` cannot be combined with `size=10` (a paging query) `size=10 flags=0 threshold=0 query="select time,ip.src,ip.dst, ip.proto,ipv6.proto,eth.type,size,payload,lifetime,client,did`

The aggregate functions have the following effect on the query result set.

Function	Result
`sum`	Add all meta values together; only works on numbers
`count`	The total number of meta fields that would have been returned

Function	Result
min	The minimum value seen
max	The maximum value seen
avg	The average value for the number
distinct	Returns a list of all unique values seen
countdistinct	Returns the number of unique values seen. countdistinct is equivalent to the number of metas that would have been returned by the distinct function.
first	Returns the first value seen
last	Returns the last value seen
len	Converts all field values to a UInt32 length instead of returning the actual value. This length is the number of bytes to store the actual value, not the length of the structure stored in the meta database. For example, the word "NetWitness" returns a length of 10. All IPv4 fields, like ip.src , return 4 bytes.
avglen	Returns a single value which is the average value returned from the len function. The result is always a float64 value.

where Clauses

The where clause is a filter specification that allows you to select sessions out of the collection by using the index.

Syntax:

```
where-criteria    = criteria-or-group, { space, logical-op, space, criteria-
or-group } ;
criteria-or-group = criteria | group ;
criteria          = (meta-key | entity), ( unary-op | binary-op meta-value-
ranges ) ;
group             = ["~"], "(" where-clause ")" ;
logical-op        = "&&" | "||" ;
unary-op          = "exists" | "!exists" ;
binary-op         = "=" | "!=" | "<" | ">" | ">=" | "<=" | "begins" |
"contains" | "ends" | "regex" ;
meta-value-ranges = meta-value-range, { ",", meta-value-range } ;
meta-value-range  = (meta-value | "l" ), [ "-", ( meta-value | "u" ) ] ;
```

```
meta-value          = number | quoted-value | ip-address | mac-address |
relative-time ;
number              = ? any numeric value ? | ( '"' text '"' )
quoted-value        = ( '"' text '"' ) | ( '"' date-time '"' ) ;
relative-time       = "rtp(" , time-boundary , "," ,  positive-integer , time-
unit, ")" ;
time-boundary       = "earliest" | "latest" | "now" ;
positive-integer    = ? any non-negative integral number ?
time-unit           = "s" | "m" | "h" ;
```

When specifying rule criteria, the `meta-value` part of the clause is expected to match the type of the meta specified by the `meta-key` . For example, if the key is `ip.src` the `meta-value` should be an IPv4 address. Entity names are allowed in any location where a meta-key name is required.

Queries using a `meta-key` name will match meta items corresponding both to the `meta-key` name as well as to the names of any "renames" specified for the key. See "Key Renaming" under the Index Customization topic for details on key renaming.

Query Operators

The following table describes the function of each operator.

Operator	Function
=	Match sessions containing the meta value exactly. If a range of values is specified, any of the values is considered a match.
!=	Matches all sessions that would not match the same clause as if it were written with the = operator.
<	For numeric values, matches sessions containing meta with the numeric value less than the right side. If the right side is a range, the first value in the range is considered. If multiple ranges are specified, the behavior is undefined. For text metas, a lexicographical comparison is performed.
<=	Same behavior as < , but sessions containing meta that equals the value exactly are also considered matches.
>	Similar to the < operator, but matches sessions where the numeric value is greater than the right side. If the right side is a range, the last value in the range is considered for the comparison.
>=	Same behavior as > , but sessions containing meta that equals the value exactly are also considered matches.
begins	Matches sessions that contain text meta value that starts with the same characters as the right side.

Operator	Function
`ends`	Matches sessions that contain text meta that ends with the same characters as the right side.
`contains`	Matches sessions that contain text meta that contains the substring given on the right side.
`regex`	Matches sessions that contain text meta that matches the regex given on the right side. The regex parsing is handled by boost::regex.
`exists`	Matches sessions that contain any meta value with the given meta key.
`!exists`	Matches sessions that do not contain any meta value with the given meta key.
`length`	Matches sessions that contain text meta values of a certain length. The expression on the right side must be a non-negative number.

Text Values

The system expects quoted text values. Unless it can be parsed as a time (see below), a quoted value is interpreted as text.

> It is also important to quote any text value that may contain – so that it is not interpreted as a range.

For text values, the backslash character `\` is used as an escape value. This character is used when you need to search for a value containing quote characters. If you need to search for a backslash character, then the backslash itself must be escaped, as `\\` . Note that if you are wrapping the query parameters within another language, such as the parameter fields of the REST interface, you may need to add additional escape levels as required by whatever API or interface you are using to interact with the core service.

IP Addresses

IP addresses can be expressed using standard text representations for IPv4 and IPv6 addresses. In addition, the query can use CIDR notation to express a range of addresses. If CIDR notation is used, it is expanded to the equivalent value range.

MAC Addresses

A MAC address can be specified using standard MAC address notation: `aa:bb:cc:dd:ee:ff`

Numeric Values

In a where clause, you can specify numeric search values. Numbers should not be surrounded by quotes.

Bucketed Numeric Indexes

Meta keys indexed with bucketing can be used like any other numeric search value. Under most situations such searches will return sessions that have a meta value that exactly matches the requested search criteria.

Special behavior is invoked for queries that select only `sessionid` , for example a query of the form `select sessionid where size = 2048` . Selecting `sessionid` explicitly bypasses all meta database read operations, and only returns index information. If selecting `sessionid` only, and if the numeric value specified is exactly equal to one of the bucket values, then the system will return all sessions that match somewhere in the bucket, rather than an exact match. For example, the search term `size = 2048` will match all sessions in the 2 KB bucket, which is the range from 2048 to 3171 bytes. However, if the search values does not match a bucket values, then the system will return only matches for the exact byte value. For example, the search term `size = 2049` will only match sessions with a size meta value exactly 2049. In this mode of operation, specifying a non-bucket value in a where clause is slower than searching within a bucket value. The 'where' clause parameter to the `values` API also invokes this optimization.

Using bucketed values in other forms of `query` does not invoke special behavior. The same is true for the `msearch` API. For those APIs, the use of a bucketed index in the where clause is evaluated accurately, without special meaning applied to bucket values. To search within an entire bucket using these APIs, specify the bucket range explicitly. For example `size=2048-3171` .

More information on how to tell if an index is bucketing is in the topic Index Customization .

Numeric Value Aliases

For numeric values, aliases specified in the index can be used in a query as a quoted string in place of where a literal numeric value would be used; e.g.,

```
select * where service = "NFS"
```

Numeric value aliases can be used anywhere a numeric literal might be used: as a single value, as the beginning or end of a range, or in a comma-delimited list of values (and/or ranges).

Refer to the topic Index Customization for details of how value aliases can be specified in the index.

Date and Time Expressions

In NetWitness Platform, dates are represented using Unix epoch time, which is the number of seconds since Jan 1, 1970 UTC. In queries, you can express the time as this number of seconds, or you can use the string representation. The string representation for the date and time is `"YYYY-mmm-DD HH:MM:SS"` . A three-letter abbreviation represents the month. You can also express the Month as a two-digit number, 01-12.

Time values must be quoted.

All times specified in queries are expected to be in UTC.

Relative Time Points

Relative time points allow a where clause to reference a value at some fixed offset relative to the earliest or latest time metas seen in the collection. It can also be used to reference a point in time relative the the current time.

A relative time point expression has the syntax `rtp(boundary, duration)` .

The boundary is either `earliest` , `latest` , or `now` .

The duration is an expression of hours, minutes, or seconds. For example, `24h` , `60m` , or `60s` . When the boundary is `earliest` , the duration represents the amount of time **after** the earliest time present in the collection. If the boundary is `latest` , the duration represents the amount of time **before** the earliest time present in the collection. If the boundary is `now` , the duration represents the amount **before** the current time.

When the boundary is `now` , the system clock of the Core service host is used to determine what time it is.

Boundary can be represented as 0 seconds if you wish to specify the relative time point with no duration offset. This is most useful in the case of the `now` boundary, since it is possible that the highest, latest, time observed in the collection may be much earlier than the current time.

Relative time points can only be used in SDK operations, where there is a collection from which to get the boundaries for earliest and latest time metas.

Relative time points only work on indexed meta types. The default indexed meta types are `time` and `event.time` .

Examples:

```
Last 90m of collection time:
time = rtp(latest, 90m) - u

First 2 days of event time:
event.time = l - rtp(earliest, 48h)

Events added in the last hour:
time = rtp(now, 60m) - rtp(now,0s)
```

Special Range Values

Ranges are normally expressed with the syntax * `smallest` * - * `largest` *, but there are some special placeholder values you can use in range expressions. You can use the letter `l` to represent the lower-bound of the all meta values as the start of the range, and `u` to represent the upper bound. The bounds are determined by looking at the smallest or largest meta value found in the index out of all the meta values that have already entered the index.

If you use the `l` or `u` tag, it should be unquoted.

For example, the expression `time = "2014-may-20 11:57:00" - u` would match all time from that 2014-may-20 11:57:00 to the most recent time found in the collection.

Notice that it is easy to confuse a range expression with a text string. Make sure that text values that contain − are quoted, and that hyphens within range expressions are not within quoted text.

`group by` Clause (since 10.5)

The query API has the ability to generate aggregate groups from the results of a query call. This is done using a `group by` clause on the query. When `group by` is specified, the result set for the query is subdivided into groups. Each group of results is uniquely identified by the meta values indicated in the group by clause.

For example, consider the query `select count(ip.dst)` . This query returns a count of all ip.dst metas in the database. However, if you add a `group by` clause, like this: `select count(ip.dst) group by ip.src` , the query returns a count of the ip.dst metas found for each unique ip.src.

As of version 10.5, you can utilize up to 6 meta fields in a `group by` clause.

The `group by` clause shares some of the same functionality as the `values` call, but it offers significantly more advanced groups at the expense of longer query times. Producing the results of a grouped query involves reading the meta from the meta database for all sessions that match the `where` clause, while a values call can produce its aggregates by reading the index only.

The contents of each group returned by the query are defined by the `select` clause. The `select` clause can contain any of the aggregate functions or meta fields selected. If multiple aggregates are selected, the result of the aggregate function is defined for each group. If nonaggregate fields are selected, the meta fields are returned in batches for each group.

The result set of a `group by` query is encoded with the following rules:

1. All meta items associated with a group are delivered with the same group number.
2. The first meta items returned to the group identify the group key. For example, if the `group by` clause specifies `group by ip.src` , then the first meta item of each group will be an `ip.src` .
3. The normal, nonaggregate meta items are returned after the `group key` , but they all will have the same group number as the group key metas.
4. The aggregate result meta fields for each group are returned next.
5. All fields within a group are returned together. Different group results will not be interleaved.

If one of the `group by` meta items is missing from one of the sessions matched by the `where` clause, that meta field is treated as a NULL for the purposes of that group. When the results for that group are returned, the NULL-valued parts of the group key will be omitted from the group's results, since the database has no concept of NULL.

The semantics of a `group by` query differ from a SQL-like database in terms of what meta fields are returned. SQL databases require you to select the `group by` columns explicitly in the `select` clause if you want them to be returned in the result set. The NetWitness Core database always implicitly returns the group columns first.

A query with a `group by` clause honors the result set `size` parameter if one is provided. However, due to the nature of the grouping, it puts an additional burden on the caller to page and reform groups if a fixed-size result set is requested. For this reason, you should not specify an explicit result size when making a `group by` call. By not specifying an explicit size, the entire result set will be delivered as partial results.

`group by` clauses allow results to be grouped by an entity definition.

The following table describes the database honors configuration parameters that limit I/O or memory impact of a group by query.

Parameter	Function
/sdk/config/max.query.groups	This is the limit on how many groups can be held in memory to calculate aggregates. This parameter allows you to limit the overall memory usage of the query.
/sdk/config/max.where.clause.sessions	This is the limit on how many sessions from the where clause can be processed in a query. This parameter allows you to set a limit on the number of sessions that have to be read from the meta and session databases to resolve a query.

order by Clause (since 10.5)

An order by clause can be added to a query that contains a group by clause. The order by clause causes the set of grouped results to be returned in sorted order.

An order by consists of a set of items to sort by in ascending or descending order. Sorting can be performed on any data field that will be returned in the result set. This includes meta specified by the select clause, aggregate function results specified by the select clause, or group by meta fields.

The order by clause can sort over many columns. There is no limit on the number of order by columns allowed in the query; but a practical limit exists in that each of the order by columns must refer to something returned by the select clause or group by clause. The multiple column sort is imposed lexicographically, meaning that if two groups have equal values for the first column, then they are sorted by the second columns. If they are equal in the second column, they are sorted by the third column, and so on for however many order by columns are provided. Groups that do not contain any of the metas referenced by the order by clause are sorted first in the result set in the case of an ascending sort, and last in the case of a descending sort.

The NetWitness Core database is unique in that the groups of results returned by a query may each have many values for a selection. For example, it is possible to select all meta items that match a meta type and organize them into groups, and it is possible to use the distinct() function to return groups of distinct meta values. If an order by clause references one of the fields in the group that has multiple values, the sorting order is applied as follows:

1. Within each group, the fields with multiple matching values are ordered by the ordering clause
2. All the groups are sorted by comparing the first occurrence of the ordered field found within each group

The order by clause is only available in queries that have a group by clause, since groups are required to organize the meta fields into distinct records. If you wish to sort an arbitrary query as if there were no grouping applied, use group by sessionid . This ensures that results are returned in groups of distinct sessions or events.

group by clauses are naturally returned in ascending group key order; but, an order by clause can be used to return groups in a different order.

If an order by column does not specify asc or desc , the default ordering is ascending.

Examples:

```
select countdistinct(ip.dst) GROUP BY ip.src ORDER BY countdistinct(ip.dst)
select countdistinct(ip.dst) GROUP BY ip.src ORDER BY countdistinct(ip.dst)
desc
select countdistinct(ip.dst),sum(size) GROUP BY ip.src ORDER BY sum(size)
desc, countdistinct(ip.dst)
select sum(size) GROUP BY ip.src, ip.dst ORDER BY ip.dst desc
select user.dst,time GROUP BY sessionid ORDER BY user.dst
select * GROUP BY sessionid ORDER BY time
```

`search` parameter

The `query` API supports a `search` parameter to perform free text searching. The syntax of the search parameter is identical to the search parameter utilized by the `msearch` API call. Refer to the `msearch` documentation for a description of the search field syntax.

The `search` parameter acts as an extension of the `where` clause in the `query` parameter. This means that the `query` and `search` parameters work together. Use the `query` parameter to specify the `select` clause, the `group by` clause, or the `order by` clause. Any `where` clause criteria specified in the `query` parameter are combined with the search filter as if they were joined with an `AND` operation.

Searches through the `query` API are always done against indexed meta, in case-insensitive mode. It has the same behavior as specifying flags `si,sm,ci` to the `msearch` API.

`values` call

The index provides a low-level `values` function to access the unique meta values that have been stored in the index. This function allows developers to perform more advanced operations on groups of unique meta values.

The `values` call parameter syntax:

```
values-params          = field-name-param, space, where-param, space, size-
param, {space, flags-param} {space, start-meta-param}, {space, end-meta-
param}, {space, threshold-param}, {space, aggregate-func-param}, {space,
aggregate-field-param}, {space, min-param}, {space, max-param}, {space,
search-param} ;
field-name-param       = "fieldName=", (meta-key | entity) ;
where-param            = "where=", where-clause ;
```

```
size-param            = "size=", ? integer between 1 and 1,677,721 ? ;
start-meta-param      = ? same as query message ?
end-meta-param        = ? same as query message ?
flags-param           = "flags=", {values-flag, {"," values-flag} } ;
values-flag           = "sessions" | "size" | "packets" | "sort-total" |
"sort-value" | "order-ascending" | "order-descending" ;
threshold-flag        = "threshold=", ? non-negative integer ? ;
aggregate-func-param  = "aggregateFunction=", { aggregate-func-flag } ;
aggregate-func-flag   = "count" | "sum" ;
aggregate-field-param = "aggregateFieldName=", ( meta-key | entity ) ;
min-param             = "min=", meta-value ;
max-param             = "max=", meta-value ;
search-param          = "search=", search-string ;
```

The `values` call provides the function of returning a set of unique meta values for a given meta key. For each unique value, the `values` call can provide an aggregate total count. The function used to generate the total is controlled by the flags parameter.

Parameters

The following table describes the function of each parameter.

Parameter	Function
`fieldName`	This is the meta key name for which you retrieve unique values. For example, if `fieldName` is `ip.src` , this function returns the unique source IP values in the collection. Entities can be used for the field name, in which case the result is defined as the combined set of field values for all the referenced meta keys. If the `fieldName` refers to a key with rename references, the result is defined as the combined set of field values for the given meta key name plus all of the references' meta keys.
`where`	This is a `where` clause which filters the set of sessions for which the unique values are returned. For example, if the `fieldName` is `ip.src` , and the `where` clause is `ip.src = 192.168.0.0/16` , only values in the range of `192.168.0.0` to `192.168.255.255` are returned. For information on the `where` clause syntax, see *Where Clauses* .
`size`	The size of the set of unique values to return. This function is optimized to return a small subset of the possible unique values in the database.

Parameter	Function
id1 , id2	These optional parameters limit the scope of the search for unique values to a specific region of the meta database and the index. Setting the id1 and id2 parameters to a limited range of the meta database is very important to running searches quickly on large collections.
flags	Flags control how the values are sorted and totaled. Flags are described in the following Values Flags section.
threshold	Setting the threshold parameter allows the values call to short-cut collection of the total associated with each value once the threshold is reached. By providing a threshold, the caller can reduce the amount of index and meta items that must be retrieved from the database. If the threshold parameter is omitted or set to 0, this optimization is not used.
aggregateFunction	Optional parameter used to change the default behavior from counting sessions, packets, or size to counting or summing the numeric field defined by aggregateFieldName . Both parameters must be specified when either is defined. Pass either sum or count to specify which behavior to perform.
aggregateFieldName	The meta field on which to perform the aggregateFunction . Both aggregateFunction and aggregateFieldName parameters must be specified when the aggregate flag is set. Performing a values call using one of the aggregate functions can be significantly slower than a values call that collects totals of sessions, packets, or size. The reason for this is that each session that matches the where clause must be retrieved from the meta database. This scan causes a large portion of the query to be I/O bound on the meta DB volumes. The time taken to run an aggregate values call is linearly proportional to the number of sessions that match the where clause.
min , max	The minimum and maximum value that should be returned from the call. These parameters are used to iterate (or page) over an extremely large number of values, typically more values than could be returned from a single call. Primarily used in conjuction with the flags sort-value, sort-ascending such that the highest value returned would be used in a subsequent call as the min parameter value. The values are exclusive. If min="rsa" was specified and rsa was a valid value, rsa would not be returned; instead, the next highest value would be returned.
search	Text search pattern to be used to further refine the where parameter

values Flags

The flags parameter controls how the values call operates. There are three groups of flags that correspond to the different modes of operation as shown in the following table.

Flag	Description
sessions , size , packets	The values call allows you to specify one of these flags to determine how the total for each value is calculated. If the flag is sessions , the values call returns a count of sessions that contain each value. If the flag is size , the values call totals the size of all sessions that contain each unique value, and reports the total size for each unique value. If the flag is packets , the values call totals the number of packets in all sessions that contain each unique value, and then reports that total for each unique value.
sort-total , sort-value	These flags control how results are sorted. If the flag is sort-total , the result set is sorted in order of the totals collected. If the flag is sort-value , the results are returned in order of the sorting order of the values.
order-ascending , order-descending	These flags control the sort order of the result set. For example, if sorting by total in descending order, the values with the greatest total are returned first.
suggest	Enables suggestion mode for the values API. All other flags are ignored if this flag is set
database-scan	Make the values call bypass the index and instead collect unique values as if where traversing the meta database. This mode is slow on most cases, but it can be fast if the where clause matches very few sessions.
ignore-cache , clear-cache	These flags control result set caching on the set of values returned by this call. Normally these should not be used.

values Call Example

The values call is used extensively by the Navigation view in NetWitness Platform. The default view generates calls that look like this:

```
/sdk/values id1=198564099173 id2=1542925695937 size=20 flags=sessions,sort-total,order-descending threshold=100000 fieldName=ip.src where="time=\"2014-May-20 13:12:00\"-\"2014-May-21 13:11:59\""
```

In this example, the Navigation view requests unique values for ip.src . It requests unique values of ip.src in the time range given. It asks for the count of sessions that match each ip.src , and the results are the top 20 ip.src values when sorted by the number total count of sessions in descending order. In addition, the Navigation view has a meta ID range in order to provide an optimization hint to the query engine.

Values call and bucketing mode

When a values call is executed with a `fieldName` parameter that specifies a bucketed indexed meta, the system will only return the bucket values present within the rest of the criteria. This has the side effect of producing counts and totals that represent all sessions within each returned bucket. This is useful because it summarizes size meta into groups that represent human-readable ranges like 1 MB, 2 MB, and so one.

When the number of sessions scanned by the values call drops below 1000 sessions, the values call operates in meta-scanning mode, and at that point it returns exact values for numeric value indexes, regardless of the bucketing setting on the index.

Suggestion Mode

The `values` call has an additional execution mode that is used to provide suggested search values. In this mode of operation, the `values` call only identifies unique values stored with the given meta key name. It provides these results within milliseconds. To achieve this it does not provide any session counts, it will not utilize any other sort flags. The suggestion mode does utilize the `where` parameter to refine suggestions, but it only utilizes the time range clause if provided. Other portions of the `where` clause are not utilized to refine suggestion.

Suggestion mode is enabled by setting the `suggest` flag in the `flags` parameter.

Suggestion mode gives special meaning to the `min` parameter. The `min` parameter can contain the starting point for the suggested values. The return values of `suggest` mode will only include values that start with the text provided in the `min` parameter.

`search` parameter

The `values` API supports a `search` parameter to perform free text searching. The syntax of the search parameter is identical to the search parameter utilized by the `msearch` API call. Refer to the `msearch` documentation for a description of the search field syntax.

The `search` parameter acts as an extension of the `where` parameter. This means that the `where` and `search` parameters work together. Any `where` parameter specified is combined with the search filter as if they were joined with an `AND` operation.

Searches through the `values` API are always done against indexed meta, in case-insensitive mode.

The `values` API only operates on index entries for most requests, in order to provide fast totals over a large numbers of events. When the search parameter operates over in-exact indexes, such as N-Gram indexes, it may include sessions with near matches instead of narrowing the search to exact matches.

`msearch` Call

The index provides a low-level `msearch` function to perform text searches against all meta types. This type of search does not require users to define their queries in terms of known meta types. Instead, it searches all parts of the database for matches. `msearch` is used by the Events view text search. See the "Filter and Search Results in the Events View" topic in the *Investigation and Malware Analysis Guide* for detail on the accepted search forms and examples.

`msearch` parameters:

```
msearch-params  = search-param, {space, where-param}, {space, limit-param},
{space, size-param}, {space, flags-param};
search-param    = "search=", ? free-form search string ? ;
where-param     = "where=", ? optional where clause ? ;
limit-param     = "limit=", ? optional session scan limit ? ;
size-param     = "size=", ? optional result count limit ? ;
flags           = "flags=", {msearch-flag, {"," msearch-flag} };
msearch-flag    = "sp" | "sm" | "si" | "ci" | "regex" ;
```

The `msearch` algorithm works as follows:

1. A set of sessions is identified from the index by finding the intersection of three sets:
   - (Set 1) All sessions in the database
   - (Set 2) Sessions that match the `where` clause parameter
   - (Set 3) If the `si` flag is specified, sessions that indexed values that match the search string parameter.
2. If the search specifies the `sm` parameter, all meta items from the set of sessions identified in step 1 are read and scanned to see if they match the search string parameter. The meta items will be read from the service nearest to the point where the search was executed. For example, if the search is performed on a Broker, the meta items may be read from the Concentrator nearest to the broker, but if the search is performed on an Archiver the meta items will be read from the Archiver itself.
3. If the search specifies the `sp` parameter, all raw packet or log entries from the set of sessions identified in step 1 are read and scanned to see if they match the search string parameter. The packets will be read from the service nearest to the point where the search was executed. For example, if the search is performed on a Concentrator, the packet data will be read from the Decoder, but if the search is performed on an Archiver, the packet data will be read from the Archiver itself.
4. Matches from step 2 and step 3 are returned as they are found, up to the point where the `limit` parameter is reached or the `size` count is reached, whichever occurs first. The `limit` parameter specifies the maximum number of sessions for which meta and packet data will be scanned. If `limit` is not specified, the entire set of sessions determined in step 1 is scanned. The `size` parameter specifies the maximum number of results that will be returned. In practice, the `size` parameter acts more as a suggestion. It is possible that slightly more results than specified will be returned, but fewer results will never be returned. If the `size` parameter is not specified, all results matching the search will be returned.

`msearch` Flags

Flag	Description
`sp`	Scans raw packet data
`sm`	Scans all meta data
`si`	Does index lookups for all search parameters before scanning meta
`ci`	Performs a case insensitive search. Returned results are case-preserving.
`regex`	Treats the search parameter as a regular expression. Only a single regular expression can be specified, but the regular expression may be arbitrarily complex.

`msearch` Index Search Mode

Using the index search mode, specified by using the `si` flag, causes results to be returned significantly faster than any other mode. The main limitation of this mode is that it only returns matches on text terms that match value-indexed meta values.

- The `si` parameter must be combined with the `sm` flag. The `si` parameter implies the search only matches indexed meta.
- The `si` parameter can be used with regex searches, however only text indexed values will match. IP addresses and numbers will not match the regex.

Text Search Syntax

The search parameter given to `msearch` is composed of 1 or more words, separated by whitespace. For example, searching for `foo` returns sessions that contain the word foo.

If multiple terms are provided for the search, they are implicitly considered to be an AND operation. For example, searching for `foo bar` returns sessions that contain both foo AND bar. Sessions that contain only foo or only bar are filtered out. If you want to search for sessions containing any of two or more terms, you must explicitly separate the terms with the word OR. For example, searching for `foo OR bar` returns sessions that contain either foo or bar.

Search Syntax And Index Modes

The searches given to the `msearch` command are interpreted according to the index level on all the indexes. `msearch` works on the value-indexed keys in the index. Search terms provided to `msearch` will find values that are an exact match to values that were indexed.

As of version 11.1, there are new index modes available that allow `msearch` to locate text that is not an exact match to the search input. `msearch` supports wildcard searches on the word meta index, if the word meta index has the ngram option enabled. For details on the ngram option, see the topic Index Customization.

The wildcard search allows the use of the `*` and `?` characters as wildcards in search terms. The `*` can stand for 0 or more characters, while the `?` may stand for any single character. To search for those characters in an N-gram enabled index, you may escape them with a backslash character.

If the word index has the 'edge' N-gram option enabled, then it can be used to locate searches that end in a wildcard. This means it is only useful for finding text that begins with a known prefix.

If the word index has the 'all' N-gram option enabled, then wildcards may appear anywhere in the search term.

This table summarizes the relationship between word index level, and the types of searches that `msearch` will locate.

Search input	Non-indexed	IndexValues	IndexValues with Edge N-grams	IndexValues with All N-grams
"foo"	no match	"foo"	words starting with "foo"	words containing "foo"
"foo*"	no match	literal "foo*"	words starting with "foo"	words starting with "foo"
"*foo"	no match	literal "*foo"	no match	words ending with "foo"
"*foo*"	no match	literal "*foo*"	words staring with foo	words containing "foo"
"foo*"	no match	literal "foo*"	literal "foo*"	literal "foo*"

`msearch` Tips

- Always use the `where` clause to specify a time range for the search.
- To search for IP address ranges, specify them in the `where` clause.
- Use the `limit` parameter when not using the index search mode. Without it, there will be an extremely large amount of data read by the meta and packet databases.

Stored Procedures

The `query` and `values` calls provide more low-level search functionality. For more advanced use cases, server-side stored procedures exist.

Use of Quotes in Query Syntax

The query parser does not care whether you use single or double quotes within a query statement. A single- or double-quoted value is treated as text meta.

The query parser attempts to make sense of whatever you put in the statement. It is not very strict about what it will accept.

For example:

`reference.id=4752`

This clause identifies sessions that have a `reference.id` meta value that has a *numeric* value of 4752.

`reference.id='4752'` or `reference.id="4752"`

This clause identifies sessions that have a `reference.id` meta value that has a *string* value of `4752` .

However, the query engine implicitly compares numbers and strings that look like numbers as equal when the values are semantically the same. So it works with either syntax.

For most efficient performance, however, it is always a good idea to construct the queries such that the query syntax matches the data types generated by the parser.

For example, if the parser is creating `reference.id` as a numeric data type (such as `uint32` or `uint64`), then use the numeric syntax.

If the parser is creating `reference.id` as a text data type, then use the string syntax.

`hierarch` Call

The `hierarch` call returns information about the hierarchy of devices attached to the collection represented in this database.

A hierarchy consists of this device, plus any devices that this device is connected to. For each device, the contents of the `/sys/stats` folder is returned. This information includes the device name, it's UUID, and it's version information.

The hierarchy command returns it's information as a MessagePack object, which may be translated into different representations depending on what API you are using to access the Core service. For example, using the REST API it is translated to a JSON object.

For devices that connect to upstream devices, such as a Broker or Concentrator, the hierarchy message will contain a `devices` member. The `devices` member is an array that holds the contents of the `hierarch` message as executed on each upstream device. In this way, the `hierarch` message forms a hierarchical directory of all services that the device connects to, both directly and indirectly.

Index Customization

This topic describes how to use the custom index file to customize the index. Each NetWitness NextGen service is installed with a default index configuration that is intended to cover the index needs for most users of the product. However, it is possible to index new meta keys in order to use the index with custom content that generated custom meta.

Index Configuration File Locations

The index customization is accomplished by making changes to the custom index file. The location of this file is **/etc/netwitness/ng/index-<servicename>-custom.xml** , where **<servicename>** corresponds to the name of the product that you are customizing. For example, the Concentrator custom index file is /etc/netwitness/ng/index-concentrator-custom.xml.

Concentrator products also include a file that describes the default index configuration: **/etc/netwitness/ng/index-concentrator.xml** . This file is useful as a template to show how the custom index file is formatted.

If you make customizations to the index in the custom index file, those customizations override any conflict with the default index configuration.

You can make changes to the custom index file while the service is running. When the service receives an index save command, the changes to the custom index file are read and applied to the index.

Changes to the index can only be applied to new incoming data. Data cannot be retroactively reindexed with a new custom index configuration, except by Rebuilding the Index .

Index configuration entries

The custom index file is an XML document. The root element of this document is the `language` element, and inside there are elements for each meta key to describe each custom index. Each element of the custom index configuration looks like this:

```
<key name="did" description="Decoder Source" level="IndexValues" format="Text"
valueMax="100"  />
```

Definitions for each attribute in this element:

Attribute	Description
name	The name of the meta key that will be indexed
description	A human-readable description for the meta type
level	The type of index that will be created for this meta key
valueMax	The maximum unique values that will be stored for this key per slice
format	The format of the data held by all meta items with this meta key name

Attribute	Description
bucket	Enable size bucketing
ngrams	Enable ngram generation
threshold	Threshold for approximate value merging on ngram indexes
defaultAction	Default Navigate view action for this key: Open, Closed, Auto, Hidden

The next few sections examine these parameters in greater detail.

Meta names

The meta name used by the index refers to the meta key name present within every meta item in the meta database. These meta names are generated by the Decoders when parsing. Parsers can choose to generate meta with any meta key name. Therefore, the custom index allows you to choose which of the meta items generated by the Decoder are indexed.

Meta key names can be 16 characters long, and contain only letters or the '.' character.

Data Types

When the Decoder generates meta items, it assigns a data type. Each parser can choose the data type of the meta it generates. However, there are recommended and standard data types for each of the default meta keys. For example, ip.src and ip.dst are stored as the IPv4 meta type, and alias.host is stored as the Text meta type. Each parser must agree on the data format for each meta key generated by the Decoder.

When adding a custom index to the Concentrator, the data type of the custom index must match the format of the data generated by the Decoder. If the types do not match, the Concentrator attempts to convert the meta generated into the type specified for the custom index. However, these conversions sometimes fail, and the resulting index can produce undefined results.

Likewise, when many Decoders and Concentrators work together as part of a NetWitness installation, they must all agree on the types for each meta key. Conflicts of meta types between NetWitness NextGen services can lead to undefined behavior.

The following table shows the metadata types supported by the NetWitness NextGen services.

Type	Size in bytes	Description
Int8	1	Signed 8-bit integer
UInt8	1	Unsigned 8-bit integer
Int16	2	Signed 16-bit integer
UInt16	2	Unsigned 16-bit integer
Int32	4	Signed 32-bit integer
UInt32	4	Unsigned 32-bit integer
Int64	8	Signed 64-bit integer

Type	Size in bytes	Description
UInt64	8	Unsigned 64-bit integer
UInt128	16	Unsigned 128-bit integer
Float32	4	32-bit floating point number, single precision
Float64	8	64-bit floating point number, double precision
TimeT	8	Unix epoc timestamp
Binary	1-255	Arbitrary binary data
Text	1-255	UTF-8 Encoded text data
IPv4	4	IPv4 address bytes
IPv6	16	IPv6 address bytes
MAC	6	MAC Address bytes

When defining a custom index, it is important to use the best data type for the meta. For example, never store IP addresses as Text, since the Text representation takes more bytes than the IPv4 representation.

Index Levels

There are three levels, or types, of indexing: IndexNone, IndexKeys, and IndexValues.

IndexNone

This type of custom index is not really an index at all. Custom index entries with the IndexNone level exist only to define and document the meta key. IndexNone entries can be used in custom Decoder indices to enforce a specific data type for a meta key across all the parsers on a Decoder.

IndexKeys

This type of custom index indicates that the index only keeps track of sessions that contain meta items with this meta key name. However, it does not index any unique values in the meta database for the meta key.

Key-level indices take much less storage space, memory, and CPU time to manage, but they require a lot more work from the query engine when you perform query or values operations using them.

If used in a where clause, a meta key indexed at the key level can only be used to resolve operations such as exists or !exists.

IndexValues

This type of custom index keeps sessions that contain each individual unique value for the meta key. This type of index is also known as a "full index".

This type of index is needed for efficient processing of most where clauses, and for use of this meta key as the fieldName parameter of a values call.

Value Max

Value max is a parameter that can have a very significant impact on the accuracy and performance of a Value-level index.

As a Decoder parses packets or logs, it is allowed to create meta of any type with any value. Usually, these meta items are created from data copied directly out of the packet or log. Therefore, anyone can create unique meta values in response to nearly any event.

The performance of the index is directly dependent on the number of unique values it has found for each meta key. As the number of unique values increases, the rate at which new meta is indexed can decrease, and the speed with which queries are completed decreases. Since any person can influence the creation of unique meta values, it is possible for any person to affect the performance of the index.

The value max parameter limits the number of unique values that can enter the index. Therefore, a malicious user cannot flood the system with a large number of unique values in an attempt to make the NetWitness system not work.

It is important to set a value max on any meta key that may have its value influenced directly by incoming packets or logs.

The value max applies only to values added since the last index save operation.

The limit for how high value max can be set varies from version to version and on the amount of RAM available to the NetWitness NextGen service. As of 10.3, the recommended ceiling for value max is 5,000,000 for any meta key. If there are a lot of custom indices, then the value max may have to be lower.

maxLength

The max length parameter is used exclusively on the `word` meta type. The meaning of the `maxLength` parameter depends on whether the index is storing N-grams, as indicated by the `ngrams` parameter. The default and recommended value for `maxLength` is 5.

Max Length without N-Grams

If N-Gram support is turned off, then the `maxLength` parameter indicates that search terms need to be truncated so that they will match truncated values in the index and meta database. If this is the case, the `maxLength`**must** be less than or equal to the corresponding setting for **/decoder/parsers/config/token.max.length** on the Log Decoder service that is generating `word` token metas. The index will use the `maxLength` to properly interpret search terms fed into the `msearch` SDK function.

Max Length with N-Grams

If N-Gram support is turned on, by setting `ngrams="Edge"` or `ngrams="All"` , then the maxLength parameter controls the maximum length of N-Grams extracted from the meta item. In this scenario, the maxLength does not have to match the length of `word` meta items generated on the Log Decoder.

minLength

The minimum length parameter is used exclusively on the `word` meta type. It only has an effect when N-grams are generated. It indicates the smallest length N-gram that will be extracted from the `word` meta items. The default and recommended minimum N-Gram length is 3, which means that searches against the `word` index must have at least 3 characters.

ngrams

The `ngrams` parameter is used exclusively on the `word` meta type. N-gram indexes extract information that allow for fast lookup of searches that only match part of the word. For example they allow for finding 'ball' inside the word 'basketball'. If set to the value of `all` , then the index will create entries for all N-grams within the word meta values. The minimum value of N is specified by `minLength` , and the maximum value of N is specified by `maxLength` .

The `ngrams` parameter also supports the value `edge` , which indicates the index will only store N-grams that appear at the beginning of a word. Edge N-grams are useful for type-ahead search matching, and take less space than storing all N-grams. However they are not useful to locate matches inside the word or at the end of the word.

The `ngrams` parameter supports the `` `allvalue` `` value for the text format meta keys. It means that the index for a meta key will store `` `all` `` N-grams within the meta values and also `` `IndexValues` `` limited by ValueMax.

This index type enhances the search capability on overflowed index values due to Value Max limits. The N-grams index provides the ability to search any meta value and the Values index provides the ability to retrieve top N available values.

The `` `minLength` `` parameter specifies the minimum value of N, the `` `maxLength` `` parameter specifies the maximum value of N, and the `` `ValueMax` `` parameter specifies maximum unique values. The following are some guidelines to follow while using these parameters:

- It is recommended to set `minLength=3` and `maxLength=3` for compact index storage of N-grams and also use ValueMax to limit value index storage. When compared to Text format keys indexed by IndexValues and ValueMax=0 (unlimited) this N-gram index configuration provides better search functionality with compact index storage and memory usage.
- The `contains` operation in queries runs faster for meta keys indexed with this Ngram index type when compared to IndexValues.
- As the index type uses both N-grams and IndexValues for the same meta key, it increases the index memory and the index storage usage for the meta key and eventually reduces index retention. Hence it is recommended to choose this index type `only` for desired meta keys to consider storage and index retention.
- When you switch to this N-gram type and if the new behavior is required on the whole index, you must perform a re-index.

N-gram indexing has a major impact on the functionality of the text indexes. Using the N-gram settings of 'all' or 'allvalue' N-grams with maxLength 3, a meta key index will consume approximately 2 times more space than if N-grams were not enabled.

Note: Note: When you switch from ngrams 'allvalue' or 'all' to IndexValues, then you may need to consider re-index as index slices created before the configuration change would be ngram indexes and the values call would return ngrams.

N-gram indexing has a major impact on the functionality of the text indexes. Using the N-gram settings of `all` N-grams with maxLength 3, a meta key index will consume approximately 2 times more space than if N-grams were not enabled.

In the default index configuration, only the `word` meta key has N-gram indexing enabled. This meta key is used to index text tokens extracted from unparsed logs on the Log Decoder.

The N-gram index mode supports a 'threshold' tunable parameter that controls the precision of the index. The threshold is used to merge similar index values together depending on how closely the set of indexed sessions matches. Values greater than 0 and less than or equal to 1.0 are accepted. A value of 1.0 means that the index will only merge values if they were found in the same set of sessions. Higher values mean that the index will merge fewer values, at the expense of requiring more time to create the index during aggregation. Lower values mean the index will merge more values together, at the expense of longer search execution time due to more database access. The threshold parameter does **not** affect search accuracy.

Numeric Bucketing

Indexes on meta formats that are unsigned integers, specifically UInt32 and UInt64, can make use of size bucketing to improve performance.

Size bucketing rounds down the size values in the index to their nearest traditional byte unit of information. Enabling this option on a numeric index reduces the number of unique values to track in the index, which improves aggregation and query performance.

The bucketing option is enabled by the boolean parameter `bucket` on the key element. `bucket` may have the value `0` , for off or `1` for on. The default is `0` .

Examples of bucket number values:

Raw Value	Value Stored in Index	Explanation
0 - 1,023	0 - 1,023	Values 0-1023 are stored unmodified
1,024 - 1,048,575	1 KB, 2 KB, 3 KB ... 1,023 KB	Values under 1 MB are stored in 1 KB buckets
1,048,576 - 1,073,741,823	1 MB, 2 MB, 3 MB ... 1,023 MB	Values under 1 GB are stored in 1 MB buckets
1,073,741,824 - 1,099,511,627,775	1 GB, 2 GB, 3 GB ... 1,023 GB	Values under 1 TB are stored in 1 GB buckets

Key Value Aliases

Value aliases can be specified for keys. Value aliases are text representations that correspond to specific values for a key. These text representations may be easier to remember and more convenient to display. Aliases can be used in the rule/query language (see Queries) and are accessible via the SDK.

Value aliases are specified using the `aliases` and `alias` elements:

```
<key description="Service Type" format="UInt32" level="IndexValues"
```

```
name="service" valueMax="75" defaultAction="Open">
        <aliases>
                <alias format="$alias" value="0">OTHER</alias>
                <alias format="$alias" value="20">FTPD</alias>
                <alias format="$alias" value="21">FTP</alias>
                <alias format="$alias" value="22">SSH</alias>
                <alias format="$alias" value="23">TELNET</alias>
                <alias format="$alias" value="25">SMTP</alias>
                ⋮
        </aliases>
</key>
```

Key Renaming

The index language supports the concept of key renaming. This feature is used to provide backwards compatibility for new key names to deprecate and replace old key names. A renaming is achieved by adding `rename` elements to the key. This has the effect of indicating the parent key renames the key in the rename element. For example, the key definition below defines a new key named `port_src` that renames the key `tcp.srcport` .

```
<key name="port_src" description="Source Port" format="UInt16"
level="IndexValues">
        <rename name="tcp.srcport"/>
</key>
```

The rename element indicates to the database that uses of the parent key, in this case `port_src` , will include both meta items with type port_src and meta items with type tcp.srcport. Thus, new meta items can be added to the database and queried using `port_src` , and such queries will return information that was previously stored in tcp.srcport as well.

The rename element accepts a single attribute, `name` , that refers to a previously defined key.

Keys referred by rename elements must have the same type as the parent key.

Keys referred by rename elements must have the same index level as the parent key.

If a key is redefined in a custom index file, and the redefined key contains rename elements, then those rename elements replace any previously defined rename elements.

Note: Usage of renamed meta key pairs in the `select` clause cannot be combined with fixed-size result paging for a query. For more information, see the Queries topic.

Entities

The index configuration is used to define entities. Entities provide a convenient way to work with several meta keys at the same time. An entity definition is an alias that groups together the results from other meta keys. You can use an entity definition anywhere you would use a normal meta name. The primary use for entities is to organize similar meta types into a single, easier to use, meta type. For example, the default NextGen database language includes distinct meta types for IP source and IP destination. You could define an entity that represents the combined set of all IP sources and destinations using an `entity` element:

```
<entity name="ip.all" description="any ip entity">
       <keyref name="ip.src"/>
       <keyref name="ip.dst"/>
</entity>
```

The `entity` element accepts the following attributes:

Name	Description
name	(Required) The name of the entity
description	(Optional) A description of the entity
defaultAction	(Optional) Navigate view action for this entity: Open, Closed, Auto, Hidden

Entity definitions create new entries in the NextGen service language. Since they are returned in the SDK `language` call, they can be used by older client applications that are not directly aware of the concept of entities.

Each entity definition must contain one or more `keyref` elements. The keyref element only allows a single `name` attribute that must refer to a real meta `key` element defined somewhere else in the device's language. The `keyref` is also allowed to refer to meta types defined in the default language.

Meta entities can be utilized in application rules, but are not supported in network rules as meta available is too limited.

Entity Definition Rules

- All the keys referenced by an entity must have the same data type
- All the keys referenced by an entity must have the same index level
- An entity name cannot conflict with any existing meta type
- Keyrefs must refer to meta key names that are defined earlier in the index configuration

Entities in Brokers

Brokers will inherit entity definitions from up-stream devices, in the same way that meta key definitions are inherited. If the upstream devices attached to the broker do not all have the same set of entities defined, the Broker will log a warning. All upstream devices should have the same entity configuration. A broker operating with mismatched entity definitions may produce undefined behavior.

Rebuilding the Index

Under normal operation, changes made to the index configuration for a service are only applied to new data that enters the collection. Rebuilding the index over all the data in the collection is a time-consuming process because it requires all of the meta database storage to be read from disk.

In version 11.0 and later, it is possible to rebuild the index while the service is online. Version 11.0 services rebuild indexes in the background whenever the service detects that part of the session and meta databases is unindexed.

Activating the Background Reindexer

The background reindexer is activated whenever the service starts. During startup, the indexer checks for gaps between sessions that are indexed and sessions that are present in the session and meta database. If any gaps are found, the background reindexer begins reindexing the session and meta database on the service.

Examples of events that may activate the background reindexer:

1. A power failure or crash occurred, rendering the last slice of the index corrupt. The corrupt data is deleted at startup, leaving a gap in the index.
2. Index data is forcibly deleted, either by doing an index reset or deleting files from the filesystem.

Controlling the Background Reindexer

The operation of the background indexer is controlled by the configuration node `/index/config/reindex.enable`. If `reindex.enable` is set to true, the next time the service starts the reindexer will operate. If `reindex.enable` is set to false, the reindexer will not start the next time the service starts, but will continue to operate until the service is restarted.

Background Reindexing Algorithm

The operation of the background indexer is as follows:

1. The index examines the ranges of sessions that are present in the index and compares them to the ranges of sessions that have valid meta data. Any discrepancies between the two are considered gaps.
2. The gaps in the index are subdivided into slices based on the current value of `/index/config/save.session.count`.
3. For each slice that is missing, a temporary index is created in one of the directories specified by `/index/config/index.dir`. The slices are reindexed in reverse numerical order. Thus, the most recently collected sessions are indexed first.
4. Once the slice is completely reindexed, it is moved into its valid location in the online index. If the reindexed slice belongs on the warm tier, it is moved to the warm tier.
5. The newly indexed data appears as part of the collection.

Background indexer status

The stat node `/index/stats/updater.state` indicates the current state of the background reindexer. This node will say `running`, `not running`, or `failed`. If the status is `failed`, check the service log for more diagnostic information.

Effects on Aggregation

Services that perform aggregation utilize the index to keep track of sessions that have already been aggregated. If the index does not have enough information to begin aggregation, aggregation will be offline until enough slices have been reindexed. During this time the aggregation status for the upstream device will indicate that it is waiting on aggregation.

Forcing A Reindex

To force the index on a service to be rebuilt:

1. Ensure that `/index/config/reindex.enable` is true.
2. Reset the index by using the `reset` message on the service. For example: `/concentrator/reset index=1` will restart the service and delete all the index files.
3. Wait for the service to restart. Background reindexing will start.
4. The most recently collected data will be available for queries as soon as the index slice representing those sessions has been reindexed.

Optimization Techniques

This topic describes optimization techniques for the NetWitness Core database. The NetWitness Core database is set up to work with a wide variety of work loads by default. However, like any database technology, its performance can be very sensitive to both the nature of the data being ingested, and the nature of the searches that the user performs against the database.

Thresholds

Thresholds are a useful optimization that can have a dramatic effect on how fast results are returned to the NetWitness Platform Navigate view. Thresholds are applied to the `values` call. For more information about the `values` call, see Queries .

The threshold defines a limit to how much of the database is retrieved from disk in order to produce a count. For most queries, the number of sessions that match the `where` clause is very large. For example, selecting all the log events for just one hour running at 30,000 events per second matches 108,000,000 sessions.

NetWitness introduced the threshold feature based on the observation that most cases where a count of sessions is required do not have to have results that are accurate down to the very last session. For example, when looking at the top 20 IP addresses present over the past hour, it is not very important if the report indicates that an IP value matched 10,000,000 or 10,000,001 sessions exactly. The estimate is good enough. In these scenarios, we can make an estimate for the value of the count returned when our count exceeds the threshold parameter. When the threshold is reached, the remaining count is estimated, and the results are sorted based on the estimated counts, if necessary.

Complex `where` Clauses

The amount of time it takes for the NetWitness Core database to produce a result is dependent on the complexity of the query. Queries that align directly with the indexes present on the meta can be resolved quickly, but it is very easy to write queries that cannot be resolved quickly. Sometimes, queries that cannot be returned quickly can be processed by the Core database and the index differently to produce much more satisfying results for the customer.

It is useful to know the relative *cost* of each part of the `where` clause. A clause with a high cost takes longer to execute. In the following table, the query operations are ordered in terms of their relative cost, from lowest to highest.

Operation	Cost
exists, !exists	Constant
= , !=	Logarithmic in terms of the number of unique values for the meta key, linear in terms of the number of unique elements that match a range expression
< , > , <= , >=	Logarithmic in terms of unique value lookup, but more likely to be linear since the expression matches a large range of values
`begins` , `ends` , `contains`	Linear in terms of the number of unique values for the meta key

Operation	Cost
regex	Linear in terms of the number of unique values for the meta key with a high per-value cost

AND and OR

When constructing a where clause, keep in mind that constructing many terms using the AND operator can have a beneficial affect on the performance of a query. Any time that multiple criteria can be used to filter down the set of sessions matching the where clause, there is less work for the query to do. Likewise, each OR clause creates a larger set of sessions to process for each query.

As a general rule of thumb, the more AND clauses in the query, the faster it completes, but the more OR clauses in the query, the slower it completes.

Use Case: Match a Large Subnet

It is common for users to construct queries that attempt to include or exclude a large subnet. This type of query is common because the users are trying to include or exclude some large portion of their internal network from their investigation.

Starting with NetWitness 11.4, all /8, /16, and /24 IPv4 subnets get their own index entries. So, no special action is needed to help optimize these queries.

For versions of NetWitness 11.3 and earlier, it is slower for the query engine to resolve this query using the ip.src or ip.dst indices alone. The issue arises from the fact that a where clause such as this:

```
ip.src = 10.0.0.0/23
```

Actually must be interpreted as:

```
ip.src = 10.0.0.0 || ip.src = 10.0.0.1 || ip.src = 10.0.0.2 || ... || ip.src =
10.0.1.255
```

Thus, the index could have to create a where clause with 512 terms.

The solution to this problem is to use the Decoder to tag common networks of interest using application rules. For example, you could create meta items with an application rule that looks like this:

```
name=internal rule="ip.src = 10.0.0.0/23" order=3 alert=network
```

This rule creates meta items in the meta key network with the value internal for any IP address in the 10.0.0.0/23 network.

The where clause could be expressed as:

```
network = "internal"
```

Assuming there is a value-level index on the network meta data, the index does not have to expand this query into anything more complex, and the sessions matching the desired subnet are matched very quickly.

Use Case: Substring Matching

Using the operators `begins` , `ends` , `contains` , and `regex` in a `where` clause can be very slow if there are a large number of unique values for the meta key. Each of these operators is evaluated independently against each unique value. For example, if the operator is `regex` , the `regex` must be run independently against each unique value.

To work around this, the most effective strategy is to reorganize the meta items such that the user does not have to use a substring match.

For example, consider if the users are attempting to find the host name within a URL somewhere in the session. The users might write a `where` clause such as:

```
url contains 'www.rsa.com'
```

In this scenario, it is likely that the `url` meta key contains one unique value for every session that was captured by the Decoder, and therefore has a huge number of unique values. In this case, the contains operation is slow.

The best approach is to identify the part of meta data they are attempting to match, and move the matching into the content parser.

For example, if there is meta data being generated for each URL, a parser could also break down the URL into its constituent components. For example, if the Decoder generates URL meta data with the value `http://www.rsa.com/netwitness` , it could also generate `alias.host` meta data with the value `www.rsa.com` . Queries could be performed using:

```
alias.host = 'www.rsa.com'
```

Since the substring operator is no longer needed, the query is much faster.

Index Saves

The Core index is subdivided by save points, also known as slices. When the index is saved, all the data in the index is flushed to disk, and that portion of the index is marked as read-only. Saves serve two functions:

- Each save point represents a place where the index could be recovered in the case of a power failure.
- Periodically saving can ensure that the portion of the index that is actively being updated does not grow larger than RAM.

Save points have the effect of partitioning the index into independent, non-overlapping segments. When a query must cross over multiple save points, it must re-execute parts of the query and merge the results together. This ultimately makes the query take longer to complete.

By default, for installations of versions 10.5 and later, a save is performed on the Core index every time 600,000,000 sessions are added to the database. This interval is set by the index configuration parameter `save.session.count` . For more information, see Index Configuration Nodes .

Older versions of NetWitness Platform, or systems that have been upgraded from versions prior to 10.5, use a time-based save schedule that saves the index every 8 hours. You can see the current save interval by using the scheduler editor in the NetWitness Admin UI for the service. The default entry looks like this:

```
hours=8 pathname=/index msg=save
```

By adjusting the interval, you can control how often saves are created.

Affects of Increasing the Save Interval

By increasing the save interval, save points are created less frequently, and therefore fewer save points exist. This has a positive effect on query performance, because it becomes less likely that queries traverse slices, and when slices do have to be traversed, there are not as many to traverse.

There are downsides to increasing the save interval though. First, the Concentrator is more likely to hit the `valueMax` limit set on any of the indices. Second, the recovery time in the event of a forced shutdown or power failure is increased. And third, the aggregation rate may suffer if the index slice grows too large to fit in memory.

Affects of Decreasing the Save Interval

By decreasing the save interval, it is possible to avoid hitting the `valueMax` limits while maintaining a full value index for meta data that contains a large number of unique values. Decreasing the save interval does have a detrimental impact on query performance, since more slices are created.

Working with `valueMax`

The `valueMax` limitation can be frustrating to customers who want to index all possible unique meta. Unfortunately that is not possible in the general case. Meta keys exist that can have arbitrary random data from anywhere on the Internet, and all unique values cannot be indexed.

However, it is possible to work around some of the limitations of `valueMax` by using key level indices instead of value indices. Key level indices are not influenced by `valueMax` .

It is possible to use the Navigate view on a meta key indexed at the key level. The database uses value level indices in the `where` clause where possible, but meta database scanning is used to resolve unique values for the `values` call. This approach works well when the `where` clause provides an effective filter to limit search scope to a small number of sessions, perhaps less than 10,000 sessions.

In cases where the `valueMax` is reached, the users can perform a database scan on their queries to ensure no relevant values were dropped. This feature is accessible in the Investigator 9.8 client via the right-click menu on the Navigation view. Although the meta database scan takes a long time, it reassures the customer that they are not missing anything in their reports.

Parallelize Workloads

When the customer is using a lot of reports, ensure that they are making full use of the parallel executing options within Reporting Engine. Likewise, ensure that the number of `max.concurrent.queries` is appropriate for the hardware.

The NetWitness Platform Navigate view has the ability to run the components of its output in parallel, which can have a significant impact on the perceived performance of the NetWitness Core service.

Index Rebuild

In rare cases, a Core service might benefit from an index rebuild. Examples:

- The NetWitness Core service has index slices created by a very old version of the product and has not rolled out any data in more than six months.
- The index was configured incorrectly, and the customer wants to re-index all meta with a new index configuration.
- The traffic load into the Core service was very light, and the save interval was large, causing more slices than needed to be generated.

In these cases, an index rebuild may provide performance improvements. To do so, you must send the message `reset` with the parameter `index=1` to the `/decoder` folder on a Decoder, the `/concentrator` folder on a Concentrator, or the `/archiver` folder on an Archiver.

Be aware that a full reindex takes days to complete on a fully loaded Concentrator, and possibly weeks on a full Archiver.

Scaling Retention

There are several ways to improve the retention of the NetWitness Core database. Retention refers to the period of time that is covered by data stored in the database.

The first step in analyzing retention is to determine which part of the database is the limiting factor in terms of retention. The packet, meta, and session databases provide the `packet.oldest.file.time` , `meta.oldest.file.time` , and `session.oldest.file.time` stats in the `/database/stats` folder to show the age of the oldest file in the database. The index provides the `/index/stats/time.begin` stat to show the oldest session time stored in the index.

Increasing Packet and Meta Retention

The primary mechanism for increasing retention on these databases is adding more storage. If adding more storage to the NetWitness Core service is not possible, then it may be necessary to use the compression options on the packet and meta database to reduce the amount of data each database writes.

If meta retention is a concern, you may want to remove unneeded content from the Decoder generating meta. Many parsers generate meta that the customer does not need to store long term.

Increasing Index Retention

Usually the index has longer retention than the databases, but with a complex custom index the index retention may be shorter. Usually the easiest course of action is to remove unneeded value-level indices from the custom config, or perhaps override some of the default value-level indices with key-level indices.

It is also possible to scale the index by adding additional index storage. However, the index storage should be extended using solid-state drives only.

Scaling Horizontally

Starting in version 10.3, Concentrators and Archivers have the ability to be clustered using group aggregation. Group aggregation allows a single Decoder to feed sessions to multiple Concentrators or Archivers in a load-balanced manner. Group aggregation enables the query and aggregation workload to be split among an arbitrarily large pool of hardware.

For more information, see the "Group Aggregation" topic in the *Deployment Guide* .

Grouping Workloads

The NetWitness Core database works much better when all the users of the system are working within the same region of the database. Since the database is fed data from the Decoder with a first-in-first-out scheme, data in the database typically is clustered together according to the time it was captured and stored. Therefore, the database works better when all users are working on the same time period of data.

It is not always possible for all users to be working on the same time period simultaneously. The NetWitness Core database can handle that use case, but it is slow to do so because it must alternate between having different periods of time in RAM. It is not possible to have all of the time periods in RAM at the same time. Typically less than 1 percent of the database and less than 10 percent of the index fits in RAM.

To make NetWitness Platform work for the customer, it is important to get the customer to organize their users into groups that tend to work on the same time ranges. For example, users who do daily monitoring over the most recent data may be one user group. Perhaps there is another user group that does queries further back in time as part of an investigation. And perhaps another set of users do reports over large periods of time. Attempting to serve all the groups from a single database can lead to frustration and long wait times for results to be produced. However, if the different use cases can be spread to different Concentrator hardware, the perceived performance can be much better. In this case, it may be beneficial to utilize more Concentrator services with less RAM and CPU power rather than a single large and expensive Concentrator intended to meet all needs.

Cache Window

Consider this sequence of events:

1. At 9:00 a.m., user "kevin" logs in to a Concentrator and requests a report on the last one hour of collection time.
2. The Concentrator retrieves reports for the time range 8:00 a.m. to 9:00 a.m.
3. At 9:02 a.m., user "scott" logs in to the same Concentrator and also requests a report on the last one hour of collection time.
4. The Concentrator retrieves reports for the time range 8:02 a.m. to 9:02 a.m.

Notice that even though both users were looking at time ranges that were close together, the work done by the Concentrator to produce reports for Kevin could not be re-sent to Scott, since the time ranges are slightly different. The Concentrator had to re-calculate most of the reports for Scott.

The setting `cache.window.minutes` on the `/sdk` node allows you to optimize this situation. When a user logs in, the point in time representing the most recent data for the collection only moves forward in increments of the the number of minutes in this setting.

For example, assume the `/sdk/config/cache.window.minutes` is 10\. Re-evaluating the above action changes the sequence of events.

1. At 9:00 a.m., user "kevin" logs in to a Concentrator and requests a report on the last one hour of collection time.
2. The Concentrator retrieves reports for the time range 8:00 a.m. to 9:00 a.m.
3. At 9:02 a.m., user "scott" logs in to the same Concentrator and also requests a report on the last one hour of collection time.
4. The Concentrator retrieves reports for the time range 8:00 a.m. to 9:00 a.m.
5. At 9:10 a.m., user "scott" re-loads the reports for the last one hour of collection time.
6. The Concentrator retrieves reports for the time range 8:10 a.m. to 9:10 a.m.

The report returned in step 3 falls in the cache window, so it is returned instantaneously. This gives Scott the impression that the Concentrator is very fast.

Thus, larger `cache.window` settings improve perceived performance, at the cost of introducing small delays until the latest data is available to search.

Time Limits

When a query is running on the NetWitness Core database for a very long time, the Core service dedicates more and more CPU time and RAM to that query in order to get it to complete faster. This can have a detrimental impact on other queries and aggregation. In order to prevent lower privileged users from using more than their share of the Core service resources, it is a good idea to put time limits on the queries run by normal users.

Query Priority

Some calls in `/sdk` support a `queryPriority` parameter. This parameter is used when the query execution pool is full and other queries are waiting to be executed. Those waiting queries are executed as resources become available and in the order of their priority. The `queryPriority` parameter uses the same scale as the Linux `nice` command. When the `queryPriority` parameter is passed to a query on a downstream device (ie. Broker, Concentrator), it is passed along with other parameters to the upstream devices (ie. Decoder, Log Decoder).

If this parameter is not passed, the calls will default to the current user's `query.priority` setting. See Per-User Configuration Nodes .

This parameter can be used on the following calls: `language` , `msearch` , `packets` , `query` , `timeline` and `values` .

All queries and rule conditions in NetWitness Core services must follow these guidelines:

All string literals, value aliases, and time stamps must be quoted. Do not quote numbers, MAC, or IP addresses.

- `extension = 'torrent'`
- `time='2015-jan-01 00:00:00'`
- `service=80`
- `ip.src = 192.168.0.1`

The space on the right and the left of an operator is optional. For example, you can use `service=80` or `service = 80`.

Rule Examples

The following table shows examples of rule conditions. You can use rule conditions for log retention collections in an Archiver and for application, network, and correlation rules on a Decoder, Log Decoder, or Concentrator. Rule conditions are also used in all `WHERE` clauses in all Core database queries.

For detailed information on rule syntax in NetWitness Platform, see **WHERE Clauses** in the Queries .

Rule Name	Condition
ComplianceDevices	`device.group='PCI Devices' || device.group='HIPPA Devices'`
HighValueWindows	`device.group='Windows Compliance'`
MediumValueWindows	`device.type='winevent_nic' && msg.id='security_4624_security'`
LowValueWinLogs	`device.type='winevent_nic' && msg.id='security_4648_security'`
LowValueProxyLogs	`device.class='proxy' && msg.id='antivirus_license_expired'`
GeneralWindows	`device.type='winevent_nic'`

Correcting invalid rules

Since version 11.0, NetWitness Platform has been using a parser for rules and queries that strictly defines valid syntax. When a Core service encounters invalid syntax, it writes a warning in the NetWitness Platform logs indicating the error.

NetWitness Platform 11.0 and later do not support parsing of legacy syntax rules (as Security Analytics 10.6 did).

After you update to NetWitness Platform 11.0 (or later), rules with invalid syntax are highlighted in the user interface. The Rule Editor provides additional tooltips. After you fix the rules, the highlights disappear. See "Fix Rules with Invalid Syntax" in the *Decoder and Log Decoder Configuration Guide* .

The `/decoder/config/rules/rule.errors` and `/concentrator/config/rules/rule.errors` stats, introduced in 10.6, contain the count of rules with errors. If `rule.errors` is nonzero, NetWitness Platform generates a Health and Wellness alert to indicate that you need to fix the rules.

Valid Syntax with the Modern Parser

- All text types must quote literal values. Example: `username = 'user1'`
- Quotes can use single or double quotes; but they must match. (You cannot start with a single quote and finish with a double quote.)
- If the literal value has a quote, you can escape it (using a backslash) or use a different starting quote character. Both of the following examples are valid: `username = "User's"`, `username = 'User\'s'`

The following are valid syntax rules using the modern parser:

- To use a backslash in a literal string, escape it using an extra backslash: `\\`
- All time types should use quotes for dates in this form: `time = 'YYYY-MM-DD HH:MM:SS'`
- All time types that are the number of seconds since EPOCH (Jan 1, 1970), should not be quoted. Example: `time = 1448034064`
- **Everything** else is unquoted: IP addresses, MAC addresses, numerics, and so on. Example: `service = 80 && ip.src = 192.168.1.1/16`

Appendix A: Statistics

This topic describes statistics used to monitor system operation. The Core services provide a very large number of statistics for monitoring the operation of the system. Some of them are useful for monitoring performance, while some of them exist for monitoring the operation of the system or for debugging purposes.

Statistics in /database/stats

The following table shows the meaning of the statistics in **/database/stats** .

Statistic	Meaning
`meta.bytes` , `packet.bytes` , `session.bytes`	The total size of data (in bytes) stored in each database
`meta.first.id` , `packet.first.id` , `session.first.id`	The first meta ID, packet ID, and session ID, respectively, stored in the database
`meta.last.id` , `packet.last.id` , `session.last.id`	The last meta ID, packet ID, and session ID, respectively, stored in the database
`meta.oldest.file.time` , `packet.oldest.file.time` , `session.oldest.file.time`	The creation date of the oldest file in each database
`meta.rate` , `packet.rate` , `session.rate`	The count of the number of meta, packet, and session objects added to each database over the last second
`meta.total` , `packet.total` , `session.total`	The total number of meta, packet, and session objects within each database
`meta.volume.bytes` , `packet.volume.bytes` , `session.volume.bytes`	The approximate total volume size (in bytes) for all directories used by each database
`meta.free.space` , `packet.free.space` , `session.free.space`	The approximate total unused space (in bytes) across all directories used by each database

Statistics in /index/stats

The following table shows the meaning of the statistics in **/index/stats** .

Statistic	Meaning
`checkpoint.page` , `checkpoint.summary`	The last objects stored the last time an index save was created (debugging)

Statistic	Meaning
index.bytes	An approximate measure of how much disk space is required by index files
index.last.load.time	The timestamp when the current index configuration was loaded from the index configuration files
memory.used	An approximate measure of how much memory is occupied by the index
page.first.id, summary.first.id	The first page and summary object stored in the index (debugging)
page.last.id, summary.last.id	The last page and summary object stored in the index (debugging)
page.total, summary.total	Number of pages and summaries in the index (debugging)
session.first.id	The ID of the first session indexed
session.last.id	The ID of the last session indexed
sessions.since.save	The number of sessions currently held by the current index slice
values.added	The number of unique values added to the current index slice
slices.total	The number of slices in the index
time.begin	The oldest time meta indexed
time.end	The most recent time meta indexed
updater.state	The status of background reindexer

Statistics in /sdk/stats

The following table shows the meaning of the statistics in **/sdk/stats**

Statistic	Meaning
cache.window.time.begin	The beginning of the current time enforced by cache.window.minutes
cache.window.time.end	The end of the current time enforced by cache.window.minutes
queries.active	The number of queries currently executing in the index
queries.queued	The number of queries waiting for execution
values.calls	The number of calls made to the "values" function since the process was started

Statistic	Meaning
values.calls.cached	The number of calls made to the "values" function that were resolved by the values call result cache

Per-query Statistics

SDK operations, such as query and values, provide information about their execution status in /sdk/config/stats/queries/<handleid> , where <handleid> is a unique identifier for the query operation.

The following table shows the meaning of per-query statistics.

Statistic	Meaning
channel.path	This stat provides a link to the connection channel over which the operation is communicating. This channel is used to communicate results back to the client.
query.type	The type of operation being performed, such as queries or values
query	The complete set of parameters given to the query
query.progress	The percentage of the query execution that has completed
query.status	A message describing what stage of the query execution is currently occurring
running.since	The time at which the query began execution
user	The user name that executed the query

Appendix B: Index Inspect

The NetWitness Core database index has a built-in debugging feature called `inspect` that provides detailed information about the composition of its indexes. The index inspect feature is located at `/index/inspect` in every Core service configuration tree. Services that do not actually have an index, like Broker, do not have the `/index/inspect` feature.

Parameters

Options

Type: String This parameter may be set to the value `all-slices` to collect inspect information about every slice in the index. If it is not set, information on the current, most recently created slice is returned.

Collecting information on all slices may take a very long time to complete if there are many index slices.

Response

Inspect returns many rows of key value pairs that represent the state of the index.

Slice Summary

The first row returned for every slice is a summary with the following values.

`session1`	The first session ID indexed in the slice
`session2`	The last session ID indexed in the slice
`meta1`	The first meta ID in the first session indexed in the slice
`meta2`	The last meta ID in the last session indexed in the slice

Per-Index Summary

There will be per-index summary rows returned for each index. Only value-level indexes are reported.

`key`	The meta key name for the index
`pathname`	The path on disk to this index
`values`	The number of unique values stored in this index
`summaries`	The number of summary entries occupied by this index in the summary.db file
`pages`	The number of page entries occupied by this index in the page.db file
`sessions`	The number of sessions that had a value that was inserted into this index
`size`	The cumulative "size" meta values for all sessions that inserted a value into this index

`packets`	The cumulative count of packets for all sessions that inserted a value into this index
`summary1`	The first summary ID used by this index
`summary2`	The last summary ID used by this index
`session1`	The first session ID referenced by this index
`session2`	The last session ID referenced by this index

Slice Summary Footer

The last row in each inspect report contains cumulative statistics for all the indexes in the slice.

`totalKeys`	The number of indexed meta types
`totalValues`	The number of unique values tracked by all indices in this slice
`totalMemory`	An approximate total of the memory needed to open this index slice

NetWitness® Platform XDR

Version 12.0.0.0

Decoder Configuration Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Decoder and Log Decoder Quick Setup **14**

Configure Common Settings on a Decoder **17**

Configure Capture Settings 19

Select a Network Adapter 19

Configure a Decoder to Begin Capturing Data Automatically 21

Configure Optional Capture Settings 22

(Optional) Configure System-Level (BPF) Packet Filtering 24

Examples 26

Testing 26

Conversions 27

(Optional) Configure a Decoder to Capture Data Across All Types of Network Interfaces 27

(Optional) Configure Meta-Only Decoders 30

(Optional) Configure Selective Network Data Collection 31

Use Predefined Policies 32

Create New Policies from Predefined Ones 36

Create Custom Policies 40

Verify Published Policies 44

Troubleshooting Policy Deployments That Fail 52

Unpublish Policies 53

Delete Policies 55

Supported Protocols for Selective Network Data Collection 56

(Optional) Configure a Decoder to Write Standard pcap-formatted Files 58

(Optional) Multiple Adapter Packet Capture 59

Configure Multiple Adapter Packet Capture 59

Per Interface Configuration 60

Stat Nodes 60

Per-Interface Stats 60

Aggregate Stats 61

Pool Stats 61

capture.port Meta Key 61

Special Devices 62

Packet arrived out of order Message 62

Multi-Assembler Modes 62

(Optional) Internet Content Adaptation Protocol Capture 64

ICAP Capture Options 64

(Optional) Data Plane Development Kit Packet Capture 64

Advantages ... 64
How it Works ... 65
Migrate PF_RING Devices to DPDK ... 65
Manually Move a Capture Interface to DPDK ... 66
Check the device status and stats of DPDK interfaces ... 68
Check link status and stats from NwConsole ... 69
Utilizing Receive Side Scaling with DPDK ... 70
Advantages ... 70
Configuring Receive Side Scaling ... 71
How it Works ... 71
Configure the Decoder Capture Failover in Load Balance Deployments ... 71
Configuration Failover on DPDK Capture Interfaces ... 71
Disable Wake On LAN (WOL) in BIOS settings ... 72
(Optional) Remove DPDK on Capture Interfaces ... 72
Validate Link Status ... 72
DPDK Interface Configuration and Statistics ... 73
(Optional) Preserve VLAN Tags When Using the Packet MMAP Capture Interface ... 73
(Optional) Process Raw Syslog Data without Priority Field ... 78
(Optional) Configure Decoder to Support OpenAppID ... 79
Enable and Disable Parsers and Log Parsers ... 82
Start and Stop Data Capture ... 85
Configure Decoder Rules ... 86
Rule Processing ... 86
Rule and Query Guidelines ... 87
Rule Examples ... 87
Invalid Rules ... 88
General Syntax Guidelines ... 88
Capture Rule Syntax ... 89
Configure Capture Rules ... 91
Import Rules from a File and Export Rules ... 93
Push Rules to Other Services ... 95
Change Execution Order of Rules ... 97
Restore a Rule Snapshot from History ... 97
Configure Application Rules ... 99
Configure Correlation Rules ... 103
Configure Network Rules ... 107
Supported Meta Keys in Network Rule Conditions ... 107
Fix Rules with Invalid Syntax ... 112
Decoder Commands for Managing Rules ... 114
add Command ... 114
merge Command ... 115

Methods of Sending a List of Rules to a Service ... 115
Send a NetWitness Rule File ... 116
Send Numbered Parameters ... 116
Ordering Rules When Pushing ... 117
replace Command ... 118
clear Command ... 118
delete Command ... 118
validate Command ... 118
Configure Parsers and Feeds ... 119
Configure Parsers ... 119
Upload and Delete Custom Parsers ... 120
Upload Parsers to a Decoder or Log Decoder ... 120
Manage Upload Jobs ... 122
Delete Deployed Parsers ... 122
Enable and Configure the Entropy Parser ... 123
Entropy Parser Configuration in the Concentrator Custom Index File ... 126
Flex Parsers ... 128
Arithmetic Functions ... 129
Language Definition ... 129
Common Parser Operations ... 131
Match Port and Identify Immediately ... 131
Match Port and Delay Identification ... 131
Match Token and Identify Immediately ... 132
Match Multiple Tokens ... 132
Match Token and Create Metadata ... 133
General Functions ... 134
General Functions Language Definition ... 134
Logging Functions ... 136
Language Definition ... 136
Nodes ... 137
Nodes Language Definition ... 137
Payload Functions ... 141
Language Definition ... 141
Regex ... 143
Language Definition ... 143
String Functions ... 144
String Functions Language Definition ... 144
GeoIP2 Parsers ... 147
Lua Parsers ... 149
HTTP Parsers ... 150
Visibility into HTTP/2 Sessions ... 151

Decoder Snort Detection 152
Using Snort Rules in Decoders 152
Initial Decoder Snort Configuration 152
Snort Parser Output 154
Rule Statistics 155
Configuration Directives 155
Variable Definitions 155
ruletype 156
Other Configuration Entries 156
config detection: debug 156
config nopcre 156
config classification 156
Snort Parser Capabilities 156
Rule Sections 156
Content 157
Flow 158
Fast Pattern 158
PCRE 158
HTTP Parsing 158
File Data 159
Binary (byte) Directives 159
Default option / fallback 160
Packet Scanning/Scoping 160
Known Limitations for Snort Parser 160
Decoder Token Match Limits 160
Flow Bits 160
DCE Preprocessor 160
Thresholds and Filters 160
Raw Byte Scanning 160
IPv6 161
Performance Considerations 161
Content Tokens 161
Negation 161
Search Parser 162
Search Methods 162
Syntax 162
Parameters 162
Example 163
Wireless LAN Configuration 164
Troubleshooting Parsers 164
Lua Parser Errors 164
Results of Automatically Disabling a Parser 165
Resetting the Parser 165

Configure Feeds 165
Custom Feed Definition File Structure 166
Sample Feed Definition File 166
Define Multiple Values in a Single Field 167
Feed Definition Equivalents for Custom Feed Wizard Parameters 167
Sample Files for a MetaCallback Feed Using CIDR Index Range for IPv4 and IPv6 169
Feed Definitions File 171
Create a Custom Feed 172
Create a STIX Custom Feed 183
Create an Identity Feed 189
Import the SSL Certificate 196
Cannot Verify Identity Feed URL 196
Edit, Upload, or Remove a Feed 198
Create Custom Meta Keys Using a Custom Feed 202
Add a Custom Meta Key in the Log Decoder 202
Deploy a Log Decoder Feed in Live 203
Add the Custom Meta Key Entry in the Concentrator Custom Index file 206
Investigate a Custom Meta Key 207
Additional Procedures 207
Verify the Custom Meta Keys on ESA 207
Update the Schema in Archiver 208
Update the Schema in Warehouse Connector 209
Update the Schema in Reporting Engine 209
Decoder and Log Decoder Additional Procedures 211
Configure High Speed Packet Capture Capability (Version 11.6 and Later) 211
Hardware Prerequisites 213
Software Prerequisites 213
Decoder Installation 213
Install the Decoder Service on Host that Will Perform the Capture 213
Select the Capture Interfaces and Assign Interfaces to DPDK 213
Verify the Decoder Packages Installation 213
Select the Decoder Operating Mode 214
Configure the Decoder for Normal Mode (Default Mode) 214
Configure the Decoder for 10G Mode 214
Performance Tuning Parameters 216
Storage Considerations 217
Packet Retention Requires Extremely High Sustained Throughput 217
Using SAN and Other Storage Configurations 217
Optimize Read/Write Operations When Adding New Storage 217
Parsing and Content Considerations 219
Best Practices 219

Use Case 1: 10G Mode - Egress, Deep Packet Inspection 220
Tested Live Content 220
Not Tested 220
Other 220
Aggregation Adjustments Based on Tested Live Content 221
Configure the Decoder for NDR Mode 221
Performance Tuning Parameters 223
Storage Considerations 223
NDR Mode Assumes No Packet Retention 223
Parsing and Content Considerations 224
Parsing at Speeds Greater than 10G 224
Best Practices 224
Use Case 1: NDR Mode - Egress, General Purpose 225
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers 225
Tested Live Content 225
Not Tested 225
Other 225
Use Case 2: NDR Mode - Egress, Data Exfiltration 226
Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers 226
Tested Live Content 226
Not Tested 226
Other 226
Use Case 3: NDR Mode - Lateral Movement 226
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers 226
Tested Live Content 226
Not Tested 227
Other 227
Configure High Speed Packet Capture Capability (Version 11.6 and Later) 228
Hardware Prerequisites 229
Software Prerequisites 229
Decoder Installation 229
Install the Decoder Service on Host that Will Perform the Capture 230
Select the Capture Interfaces and Assign Interfaces to DPDK 230
Verify the Decoder Packages Installation 230
Select the Decoder Operating Mode 230
Configure the Decoder for Normal Mode (Default Mode) 231
Configure the Decoder for 10G Mode 231
Performance Tuning Parameters 232
Storage Considerations 233
Packet Retention Requires Extremely High Sustained Throughput 233
Using SAN and Other Storage Configurations 233
Optimize Read/Write Operations When Adding New Storage 233

Parsing and Content Considerations 235
Best Practices 235
Use Case 1: 10G Mode - Egress, Deep Packet Inspection 236
Tested Live Content 236
Not Tested 237
Other 237
Aggregation Adjustments Based on Tested Live Content 237
Configure the Decoder for NDR Mode 237
Performance Tuning Parameters 239
Storage Considerations 240
NDR Mode Assumes No Packet Retention 240
Parsing and Content Considerations 240
Parsing at Speeds Greater than 10G 240
Best Practices 240
Use Case 1: NDR Mode - Egress, General Purpose 241
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers 241
Tested Live Content 241
Not Tested 241
Other 242
Use Case 2: NDR Mode - Egress, Data Exfiltration 242
Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers 242
Tested Live Content 242
Not Tested 242
Other 242
Use Case 3: NDR Mode - Lateral Movement 242
Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers 242
Tested Live Content 243
Not Tested 243
Other 243
Configure a Log Decoder to Accept Protobuf 244
Configure Session Split Timeouts 246
Configure Syslog Forwarding to Destination 249
Configure Transaction Handling on a Decoder 251
Transactions Off 251
Transactions Represented as Meta Items 252
Transactions Split Sessions 252
Configure Data Export 253
Decrypt Incoming Packets (TLS 1.2) 254
Step 1: Validate That The Network Decoder Captures Encrypted Traffic 256
Step 2: Obtain Private Keys from Managed Servers 256
Extracting Certificates for the Decoder SSL Decryption 256
Extract private key from Apache 256

Extract Private Key and Certificate from Windows Servers 257
Step 2 (a): Extract PKCS12 file from Internet Information Services (IIS) 257
Using IIS Manager App 257
Using Microsoft Management Console (MMC) 257
Step 2 (b): Extract PKCS12 file from Exchange Servers 257
Step 2 (c): Extract Private Key From a pfx File Using OpenSSL 258
Recommendations 258
Step 3: Validate That The Private Key Cipher Suite is Supported 258
Supported Cipher Suites 258
Unsupported Cipher Suites 259
Step 4: Confirm HTTPS Parser is Enabled on Decoders 263
Step 5: Upload the Supported Private Keys to Decoders 263
Upload Multiple Premaster and Private Keys 265
Parameters for Managing Keys 268
Return Values 269
Encryption Keys 269
Premaster Key 269
Private Keys or PEM Files 269
Step 6: Reload HTTPS Parser on Decoders 270
Step 7: View Decrypted Sessions 271
Viewing Unencrypted Traffic 271
Troubleshooting 272
Decryption of Secure SMTP 272
Performance Considerations 273
TLS Certificate Hashing 275
JA3 and JA3S TLS Fingerprints 276
Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder 276
Community ID 277
Decrypt Incoming Packets TLS 1.3 278
Non-FIPS Builds 278
Install or update Non-FIPS builds 279
(Option 1) Instructions to install nonfips rpms during an upgrade 279
(Option 2) Instructions to install nonfips rpms after upgrade 280
TLS 1.3 Keys and Cipher suites 280
Cipher Suites 280
Steps to configure Network Decoder to decrypt TLS 1.3 traffic 281
Step 1: Validate that the Network Decoder Captures Encrypted TLS 1.3 Traffic 281
Step 2: Confirm HTTPS Parser is Enabled on Decoders 281
Step 3: Integrate Network Decoder to receive keys from Thirdparty TLS 1.3 key forwarder 282
Option 1: Configure F5 BIG-IP Local Traffic Manager (LTM) 282

Option 2: Using sslKeys API ... 282
Parameters for Managing Keys ... 283
Usage: ... 284
Return Values ... 284
Step 4: (Optional) Upload TLS 1.3 keys to validate Decryption ... 284
TLS 1.3 Keys ... 285
Step 5: Validate TLS 1.3 keys are received on Decoder ... 285
Step 6: View Decrypted Sessions ... 286
Viewing Unencrypted Traffic ... 286
Option 1 : Using Event Analysis UI ... 286
Option 2 : Using RESTful API on Decoder service ... 286
Troubleshooting ... 287
Additional Features ... 287
Edit Decoder System Configuration ... 288
Enable CPU Usage Statistics for Installed Content ... 290
Enable Parser Mappings ... 291
Enable IP Address to Event Source Mapping ... 291
Update IP to Event Source Mapping ... 292
Read IP to Event Source Type Mappings ... 294
Edit IP to Event Source Type Mappings ... 294
Delete IP to Event Source Type Mappings ... 295
Sort the Hostname or Event Source Type ... 295
Import IP to Event Source Mapping Entries ... 296
Export IP to Event Source Mapping Entries ... 296
Search IP to Event Source Mapping Entries ... 297
Enable or Disable Lua and Flex Parsing Systems ... 298
Map IP Address to Service Type for Log Parsing ... 300
Map an IP Address to a Service Type ... 300
IPdevice Command ... 300
Map an IP Address to a Time Zone ... 302
tzinfo Command ... 302
iptmzone Command ... 303
Examples ... 303
Change the Date format ... 303
Missing Year Support ... 305
Latent log during transition to new year ... 306
Logs from forward time zones (or skewed clocks) just before new year transition ... 306
Latent logs received where a leap day cannot be successfully assigned to an appropriate leap year ... 306
Limitations ... 306

Examples 306
Obtain Log Files from a Pre-11.0 Log Decoder 308
Upload a Log File to a Log Decoder 310
Upload a Packet Capture File 311
Simultaneous Import and Capture 312
F5 BIG IP - NetWitness Perfect Forward Secrecy Inspection Visibility 313
Prerequisites 313
Deployment 314
BIG-IP LTM Configurations 314
HTTPS Monitor 315
Pool & Node 317
HTTP Profile 318
Server SSL Profile 320
Internal Virtual Server 322
iRule Session keys copying and Sideband communication 324
Configuration 324
iRule NetWitness-sideband-irule 324
Troubleshooting Packet Drops (11.x and above) 329
Quick configuration Checks to avoid packet drops 329
To check and tune the configuration: 329
Information required to troubleshoot packet drops 330
What if the REST port is inaccessible? 330
How do I troubleshoot packet drops? 331
Navigate to the Packet drops tool using the REST interface 331
Introduction to Drops tool Charts 331
Symptoms Checklist 332
Symptom 1: Higher traffic ingestion rates for the content deployed would cause packet drops 332
Resolution: 332
Symptom 2: Packet Database Write backup would cause packet drops 333
Resolution: 334
Symptom: Session Database Write backup would cause packet drops 337
Resolution: 338
Symptom: Session Pool Parsing delays causes packet drops (Parsing issues) 339
Resolution: 341
Symptom: Session Parsing stuck can cause packet drops 343
Resolution: 343
Symptom: Lua Parser Warnings can cause packet drop issues 343
Decoder and Log Decoder References 345
Services Config View - Capture Policies Tab 346
Services Config View - Edit Capture Policies Wizard 349

Services Config View - Data Privacy Tab 353
Services Config View - Data Retention Scheduler 355
Services Config View - Feeds Tab 357
Upload Feeds Dialog 360
Services Config View - Files Tab 362
Services Config View - General Tab 364
Services Config View - Parsers Tab 373
Services Config View - Parser Mappings Tab 375
Data Export Tab 378
Services Config View - Rules Tabs 380
App Rules Tab 384
Correlation Rules Tab 388
Network Rules Tab 391
Services System View - Decoders 395

Decoder and Log Decoder Quick Setup

A basic NetWitness network includes at minimum Brokers, Concentrators, and Decoders. Brokers aggregate data from Concentrators, and Concentrators consume data from at least one Network Decoder or Log Decoder. The basic network may include both types of Decoders. Network Decoders are usually referred to as Decoders, and they capture network data in packet form. Log Decoders capture log data as events.

Adding a Decoder makes it visible and available for use with NetWitness Administration, Live Services, and Investigate. To add a service in NetWitness, you select the service type, provide service connection information, and validate that the service can be reached. The *Hosts and Services Getting Started Guide* provides the information you need to understand and install all NetWitness services.

After the services are added, you need to configure each service. This is the preferred order for configuring your system:

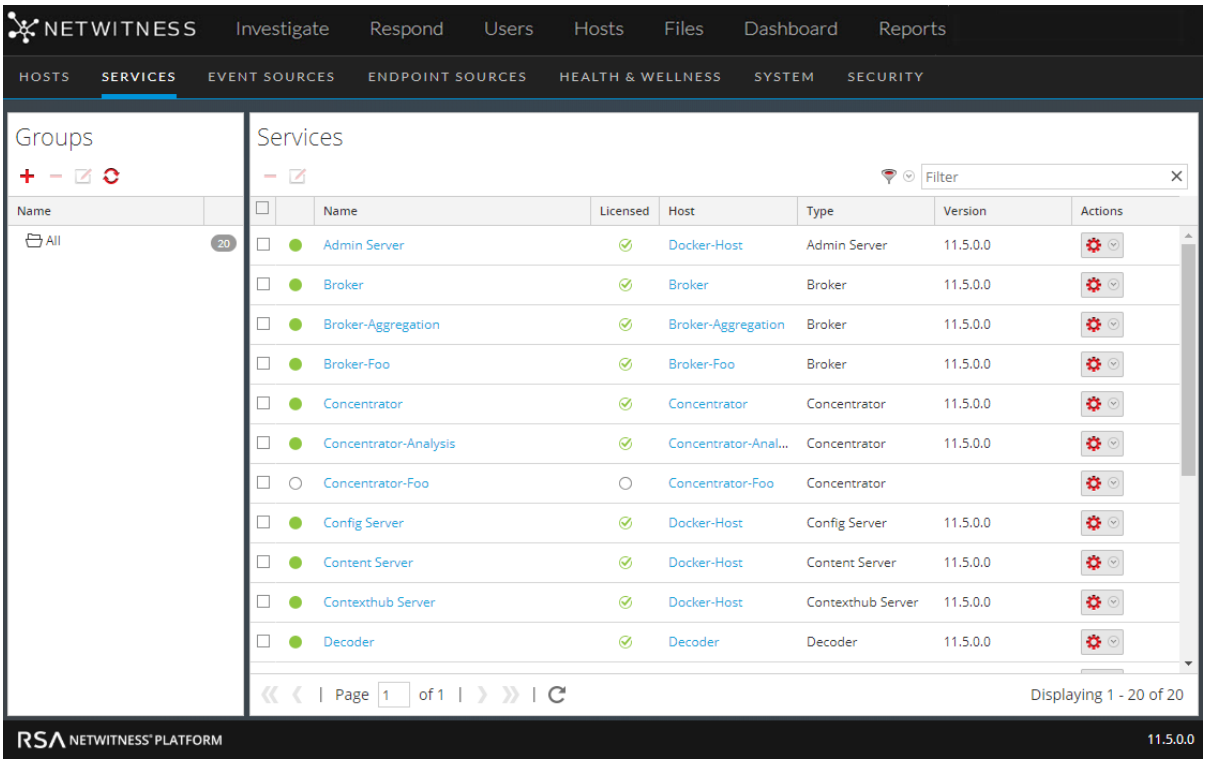
1. Decoders
2. Log Decoders
3. Concentrators (refer to the *Broker and Concentrator Configuration Guide*)
4. Brokers (refer to the *Broker and Concentrator Configuration Guide*)

Note: A Log Decoder is a special type of Decoder, which is configured and managed in a similar way to a Decoder. Most of the information in this guide refers to both types of Decoders. "Decoder" refers to both types of Decoders. Information that applies exclusively to Network Decoders or Log Decoders is clearly identified.

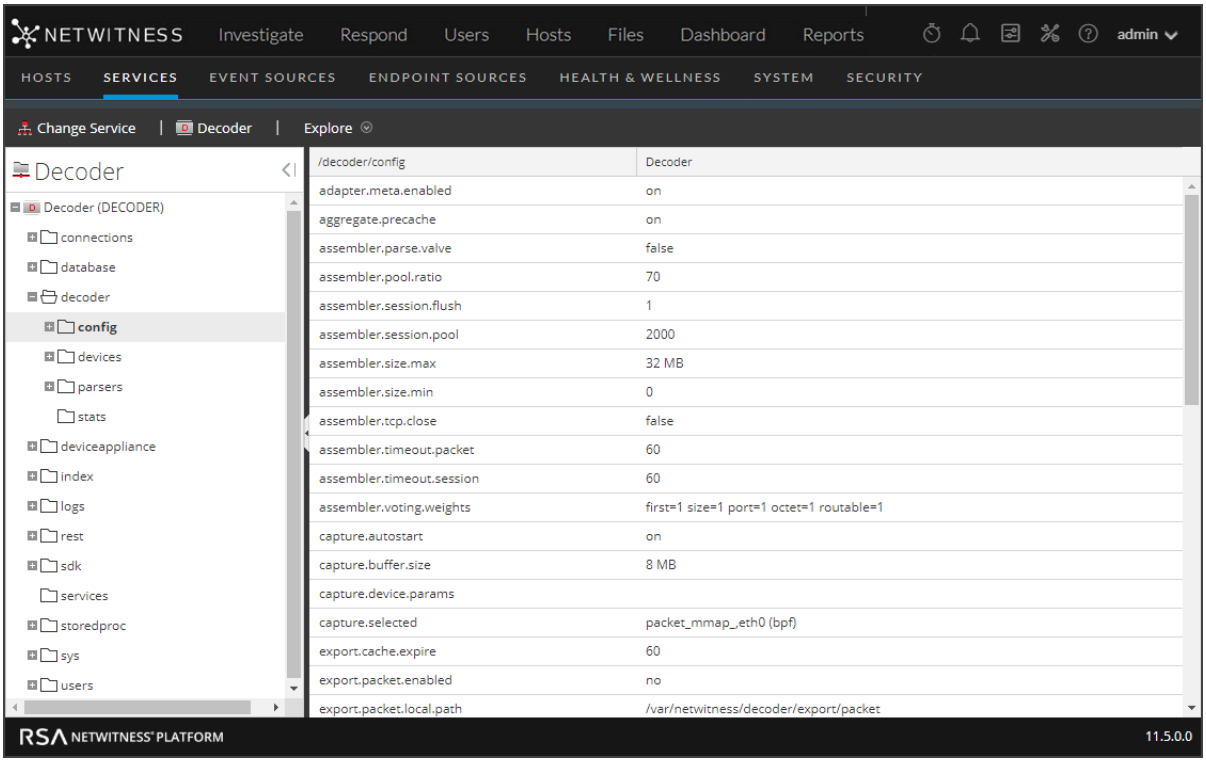
Basic configuration of the Decoder involves selecting a network adapter interface and starting data capture.

In addition, you can configure each Decoder to control the type of traffic captured using rules, feeds, and parsers. Advanced configuration tasks enable additional features that are relevant to specific applications. For example, configure a 10G Decoder, create custom meta keys, or decrypt incoming packets.

The easiest way to configure all of the required Decoder and Log Decoder settings is to use the options in the NetWitness user interface. For the most part, configuration is performed in the Administration Services view ( (Admin) > Services).



Administrators who feel comfortable working outside of the user interface can configure the basic parameters as well as advanced settings by editing database nodes in the Decoder node tree using the Services Explore view.

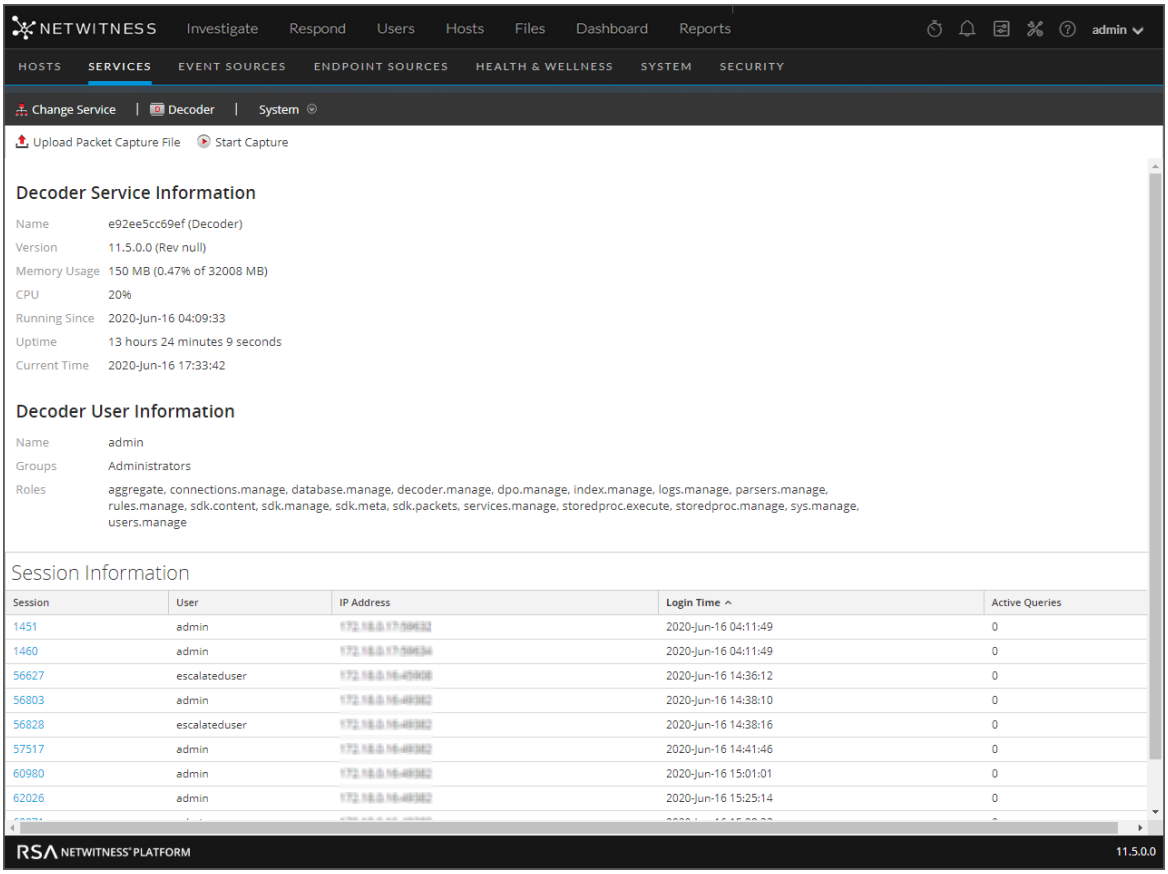


Perform Initial Quick Setup

This procedure accomplishes the initial, basic configuration of a Decoder, and starts data capture. When the basic setup is complete, the Decoder begins capturing data for the Concentrator to consume.

To configure a Decoder and start capturing data:

1. Assign a network interface for capturing data. For details, see "Select a Network Adapter" in Configure Capture Settings.
2. Do one of the following:
   a. To start capture, select the Decoder and  **> View > System**. In the toolbar click .



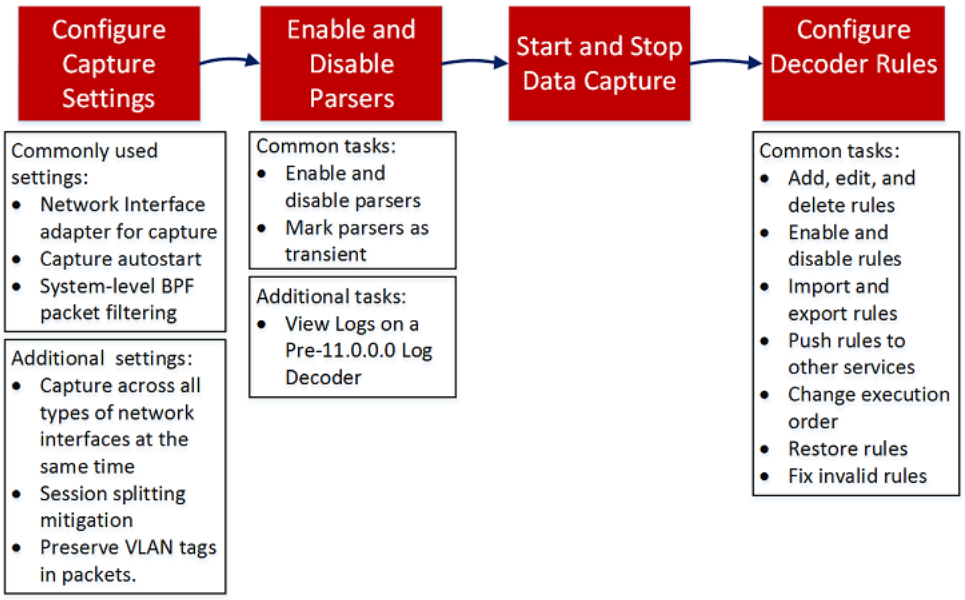
   b. To enable Capture Autostart, see "Configure a Decoder to Begin Capturing Data Automatically" in Configure Capture Settings.
   The Decoder begins capturing data for consumption by a Concentrator. For additional configuration options, refer to Configure Common Settings on a Decoder Configure Common Settings on a Decoder and Decoder and Log Decoder Additional Procedures

Configure Common Settings on a Decoder

This section introduces commonly used configuration settings on a Decoder with procedures and background information. After you have completed Decoder and Log Decoder Quick Setup, you can refine your configuration by using parsers, feeds, and rules to limit the captured data.

Note: A Log Decoder is a special type of Decoder, which is configured and managed in a similar way to a Decoder. Most of the information in this guide refers to both types of Decoders. "Decoder" refers to both types of Decoders. Information that applies exclusively to Network Decoders or Log Decoders is clearly identified.

The following workflow illustrates commonly used settings and breaks the configuration process into four steps.

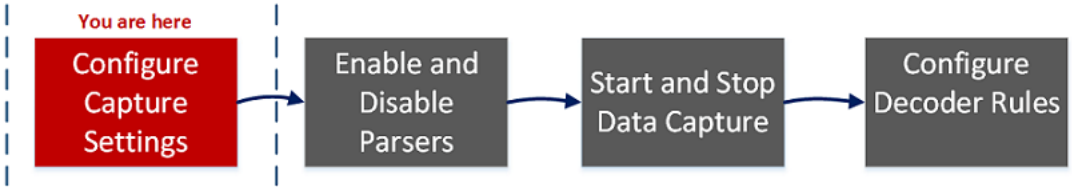


Configuration Step	Description
Configure Capture Settings	When initially setting up the Decoder, configuring the network adapter interface is required. Additional optional capture settings are available; one that is frequently used is Capture Autostart.
Enable and Disable Parsers and Log Parsers	View the parsers that have been downloaded and deployed from Live, and manage which ones are enabled or disabled.
Start and Stop Data Capture	When a Decoder starts up, it automatically begins aggregating data if Capture Autostart is enabled. When autostart is not enabled, you can start and stop data capture manually.

Configuration Step	Description
Configure Decoder Rules	Capture rules can add alerts or contextual information to sessions or logs. They can also define which data a Decoder or Log Decoder filters out. By default, no capture rules are defined when you first configure NetWitness. Unless rules are specified and the rules are valid, the packets are not filtered. You can deploy the latest rules from Live as described in the *Live Services Management Guide*. You can define capture rules at any time, and you can fix rules that use invalid syntax (Fix Rules with Invalid Syntax).

Configure Capture Settings

When initially setting up the Decoder, configuring the network adapter interface is required. Additional optional capture settings are available; two that are frequently used are the Berkeley Packet Filter, and Capture Autostart.



Besides the basic network adapter interface setup, you may decide to use one of the special-purpose configurations described in (Optional) Preserve VLAN Tags When Using the Packet MMAP Capture Interface or (Optional) Configure a Decoder to Capture Data Across All Types of Network Interfaces

The rest of the capture settings have default values chosen to be effective in most cases (see a detailed list in Services Config View - General Tab). You can adjust these in some circumstances, for example, if Customer Support advises a change. You can edit the capture settings at any time.

Select a Network Adapter

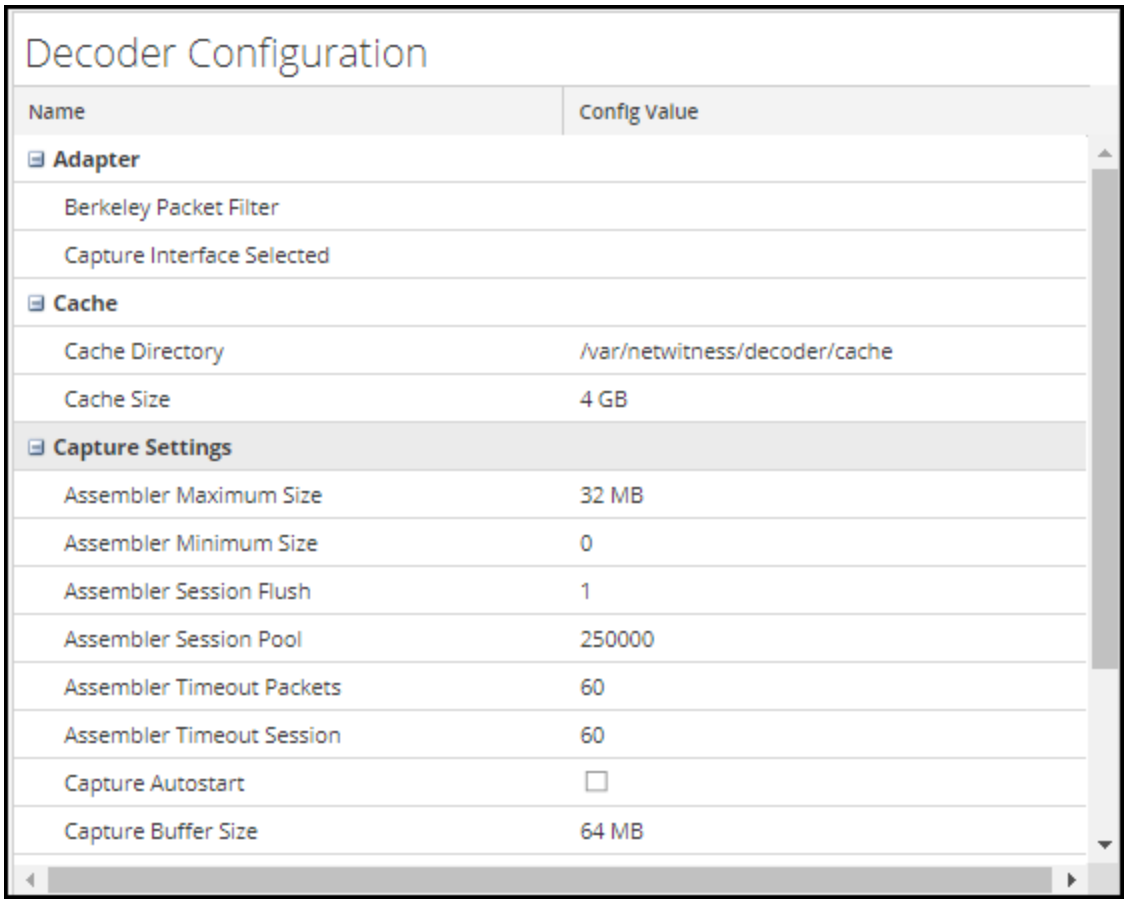
The table below describes the Network Adapter settings for a Decoder. The system administrator sets the default network adapters when the Decoder is installed. Consult your System Administrator for more information.

Adapter Parameter	Description
Berkeley Packet Filter	Berkeley Packet Filters (BPF) are applied to the packet stream before the packets are copied to the Decoder adapter for analysis. This allows unwanted traffic to be efficiently discarded. However, any packets discarded are not accounted for in any Decoder statistics (capture rate, packets dropped, and packets filtered and total packets).

Adapter Parameter	Description
Capture Interface Selected	Select an adapter through which the Decoder captures packets. For the lower speed internal capture interface, use the `packet_mmap_,7,eth1` adapter, which corresponds to the monitor port located on the motherboard. There are six additional capture ports: • `packet_mmap_,1,lo (bpf)` • `packet_mmap_,2,eth2 (bpf)` • `packet_mmap_,3,eth3 (bpf)` • `packet_mmap_,4,eth4 (bpf)` • `packet_mmap_,5,eth5 (bpf)` • `packet_mmap_,8,ALL (bpf)` There are three wireless capture services available: • `packet_netmon_ (Microsoft Netmon)` • `packet_mac80211_ (Linux mac80211)` • `packet_airport_ (Mac OS X AirPort)`
Capture Interface Selected for Log Decoder	The following capture service is available: • `log_events,Log Events`

To configure the network adapter on a Decoder:

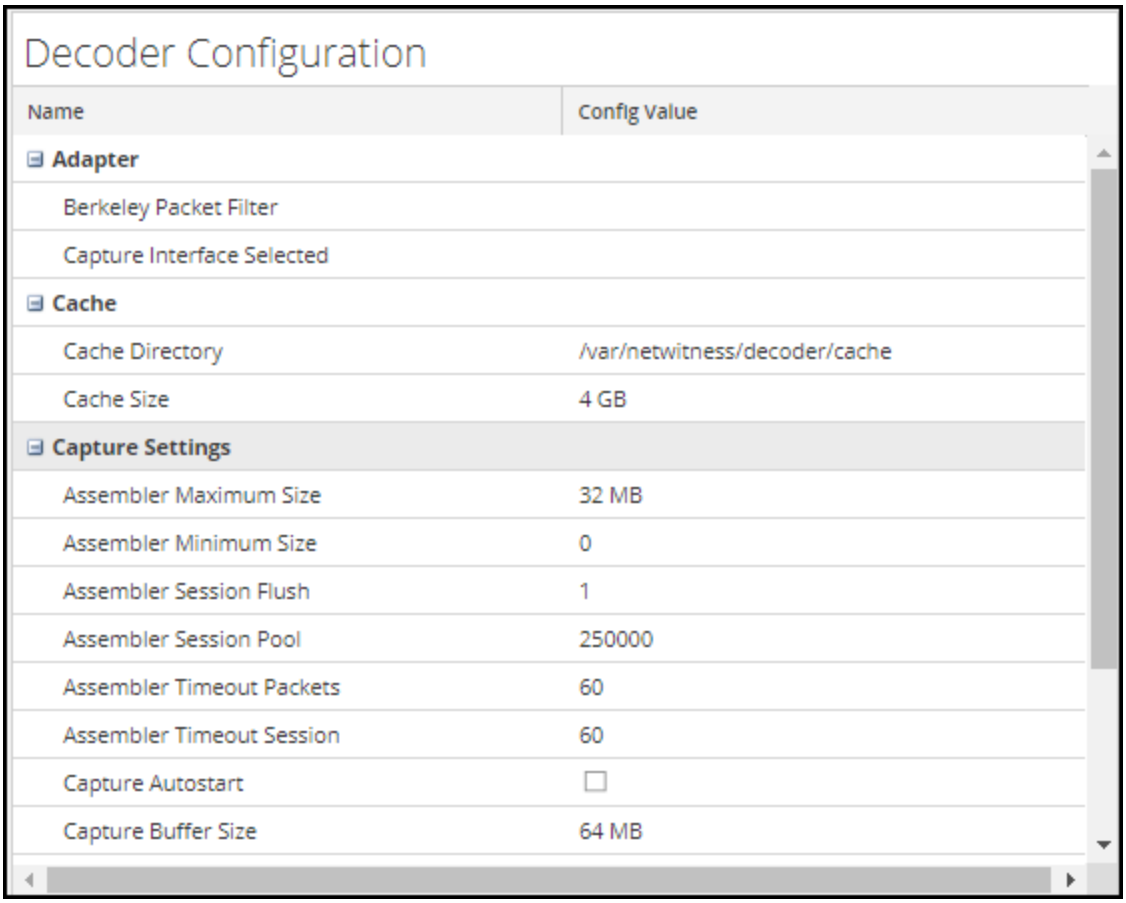
1. Go to  (Admin) **> Services**.
2. In the **Administration Services view**, select the Decoder and  **> View > Config.**
   The Services Config view is displayed with the General tab open.



3. In the **Capture Interface Selected** field, select the network adapter that best suits the Decoder.
4. To save the changes, click **Apply.**
5. If necessary to put the changes into effect, navigate back up to the **Administration Services view**, select the Decoder, and select > **Restart**.

Configure a Decoder to Begin Capturing Data Automatically

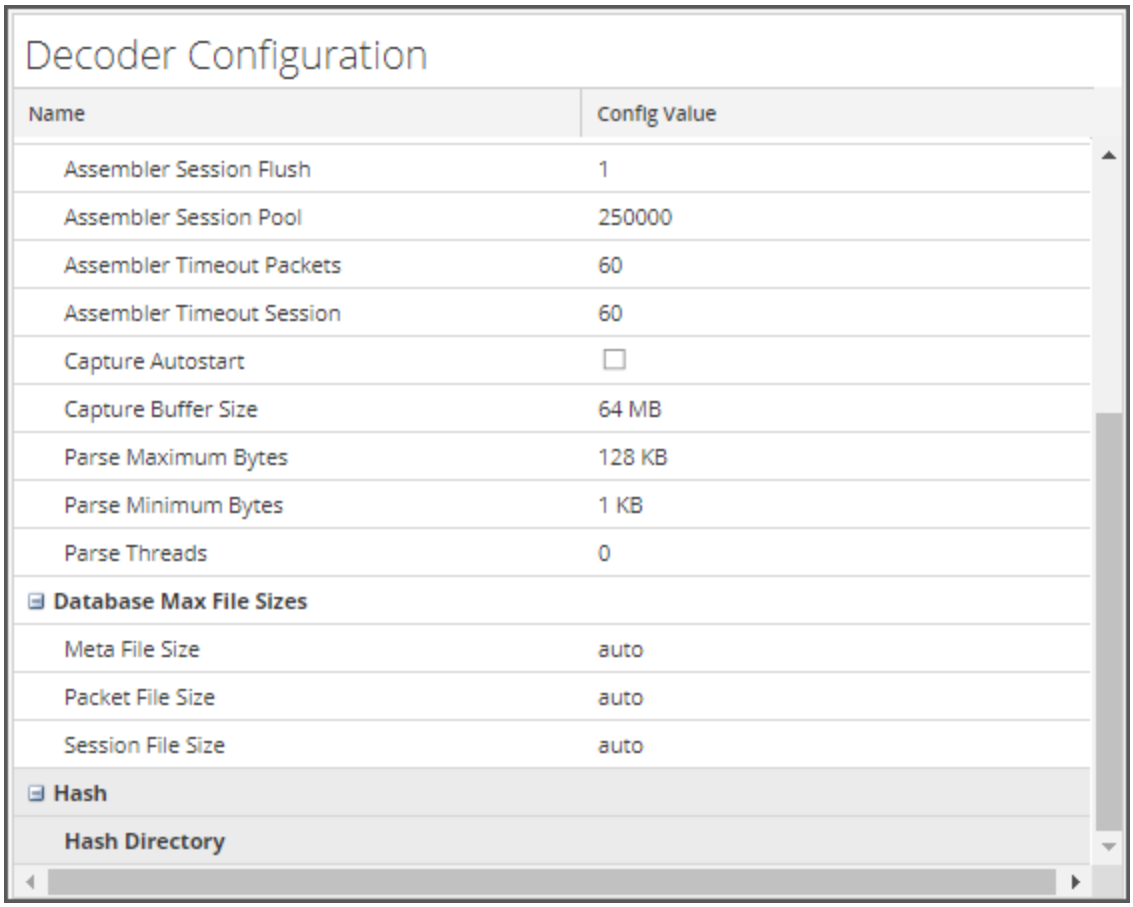
1. Go to (Admin) **> Services**.
2. In the **Administration Services view**, select the Decoder and **> View > Config.**
   The Services Config view is displayed with the General tab open.



3. Under **Capture Settings**, select the **Capture Autostart** checkbox.
4. To save the changes, click **Apply.**
5. If necessary to put the changes into effect, navigate back up to the **Administration Services view**, select the Decoder, and select > **Restart**.

Configure Optional Capture Settings

1. Go to (Admin) > **Services**.
2. In the **Administration Services view**, select the Decoder and > **View** > **Config**.
   The Services Config view is displayed with the General tab open.



3. If you want to apply a system-level filter to the packet stream before the packets are copied to the Decoder adapter for analysis, configure the Berkeley Packet Filter as described in (Optional) Configure System-Level (BPF) Packet Filtering.
4. In the **Capture Settings** sections, review the default values. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change. See Services Config View - General Tab for an explanation of these settings.
5. In the **Database Max File Sizes** section, review the default values. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change. See Services Config View - General Tab for an explanation of these settings.
6. In the **Hash** section, define a directory for hash files if you are using this feature. See Services Config View - General Tab for an explanation of these settings.

(Optional) Configure System-Level (BPF) Packet Filtering

You can use Berkeley Packet Filters to control which packets and logs are processed by a Decoder.

Berkeley Packet Filters (BPF) are applied to the packet stream before the packets are copied to the Decoder adapter for analysis. This allows unwanted traffic to be efficiently discarded. These discarded packets are not accounted for in any Decoder statistics (capture rate, packets dropped, and packets filtered and total packets).

The Decoder also supports system-level packet filtering defined using `tcpdump/libpcap` syntax. Specifying a `Libpcap` filter can efficiently reduce packet volume based on Layer 2 - Layer 4 attributes. A `Libpcap` filter is appropriate for use when a Decoder is receiving a traffic volume that is placing a load against the physical resources of the platform. In this scenario, the Decoder may consistently drop packets and have a large number of capture pages available(`/decoder/stats/capture.pagefree` is high).
The following is an example of a `libpcap` filter to keep only packets that do not have both source and destination addresses in the 10.21.0.0/16 subnet.

```
not (src net 10.21.0.0/16 and dst net 10.21.0.0/16)
```

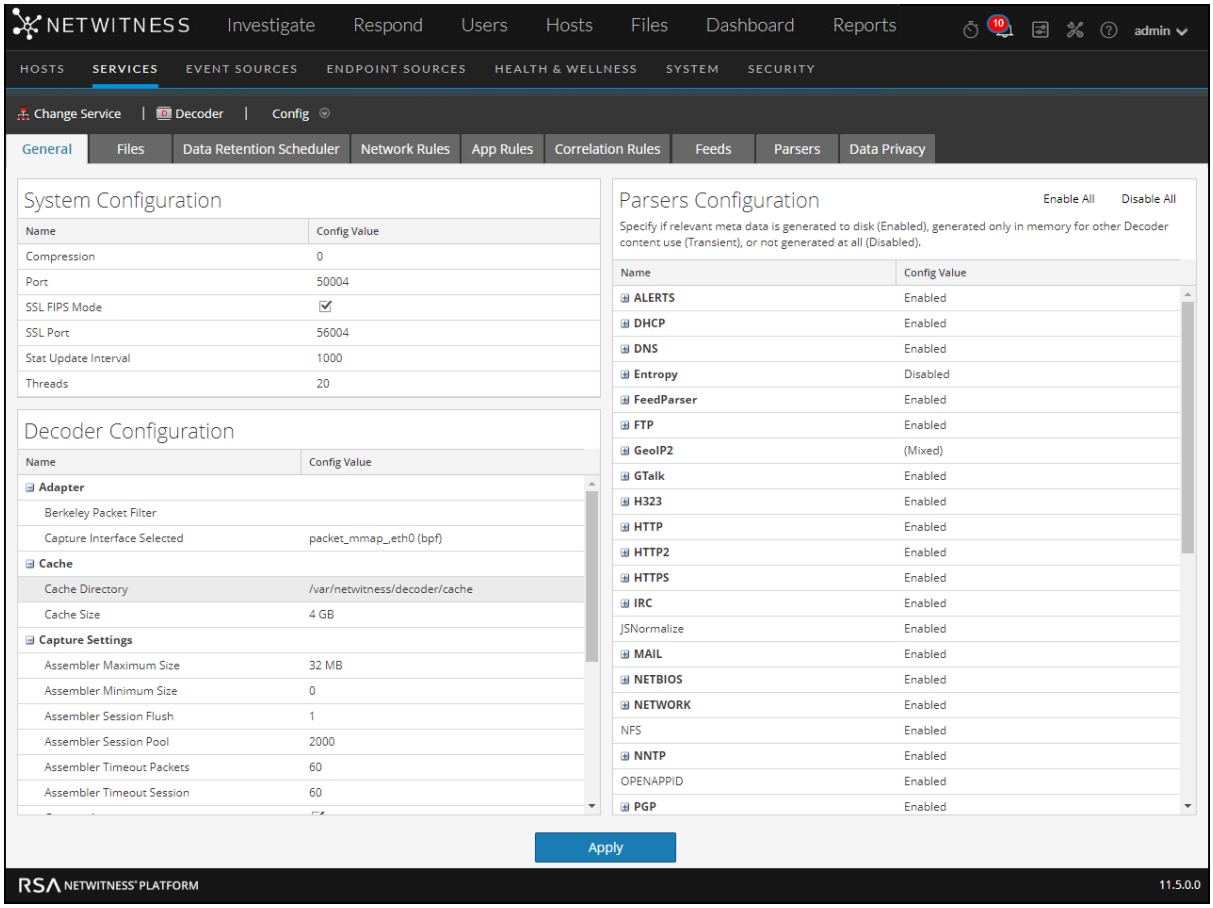
For a full reference of the `Libpcap` filter syntax, see the main pages for:

- tcpdump (http://www.tcpdump.org/tcpdump_man.html).
- pcap-filter (http://www.unix.com/man-page/FreeBSD/7/pcap-filter/).

To add a system-level Berkeley Packet Filter:

1. Go to  (Admin) **> Services.**

2. In the Administration Services view, select a Decoder service and  **> View > Config.**
   The Services Config view is displayed with the General tab open.



3. In the **Decoder Configuration Section**, under **Adapter**, click in the field next to **Berkeley Packet Filter**.
4. Type only one filter in the field. If you want to filter multiple items, join multiple expressions using `and`. Several examples are provided below.
   The user interface validates input at the time you enter your filter string.
5. To save the filter, click **Apply**.
   If the syntax is correct, a confirmation message is displayed.
   If the syntax is incorrect, a **Packet filter is not valid** message is displayed and a corresponding log message will follow in the log messages on the Decoder:
   ```
   164474800 2020-May-01 19:03:08 warning Decoder Failed to parse filter 'example_badrule': syntax error
   ```
6. To activate the filter, you must stop and start capture on the Decoder:
   a. Change the **Config** view to the **System** view.
   b. Click **Stop Capture**.
   c. Click **Start Capture**.
      The active filter will be displayed in the Decoder logs.

Examples

These are several filter examples:

- Drop packets to or from any address in the 10.21.0.0/16 subnet:
  `not (net 10.21.0.0/16)`
- Drop packets that have both source and destination addresses in the 10.21.0.0/16 subnet:
  `not (src net 10.21.0.0/16 and dst net 10.21.0.0/16)`
- Drop packets that are from 10.21.1.2 or are headed to 10.21.1.3.
  `not (src host 10.21.1.2 or dst host 10.21.1.3)`
- Combine both IP and HOST:
  `not (host 192.168.1.10) and not (host api.wxbug.net)`
- Drop all port 53 traffic, both TCP & UDP:
  `not (port 53)`
- Drop only UDP port 53 traffic:
  `not (udp port 53)`
- Drop all IP protocol 50 (IPSEC) traffic:
  `not (ip proto 50)`
- Drop all traffic on TCP ports 133 through 135.
  `not (tcp portrange 133-135)`

The following filters combine some of the above to demonstrate how to put multiple directives into one filter:

- Drop any port 53(DNS) traffic sourced from 10.21.1.2 or destined to 10.21.1.3.
  `not (port 53) and not (src host 10.21.1.2 or dst host 10.21.1.3)`
- Drop any traffic using IP proto 50 or port 53 or any traffic from net 10.21.0.0/16 destined to net 10.21.0.0/16
  `not (ip proto 50 or port 53) or not (src net 10.21.0.0/16 and dst net 10.21.0.0/16)`

Caution: The use of parentheses can have a large and potentially disruptive effect on the use of Packet Filters. As a best practice, keep "not" operations outside of parentheses and always test your rules before deploying them. Failure to properly format your rules (despite input validation) can cause a packet filter to drop ALL traffic or behave in other unexpected ways. This is due to the way packet `Libpcap` filters work and is not the result of any logic within NetWitness software.

Testing

BPF filters can and should be tested using either `tcpdump` or `windump` to ensure that they will provide the expected behavior before implementing them. This example shows a test of a filter using windump:

```
windump -nni 2 not (port 53 or port 443) or not (ip proto 50)
```

Conversions

If for the sake of performance, you have decided that an existing network rule filter would be better running as a System-Level Packet Filter, you can convert it. There are a few things to remember when doing conversions.

- `&&` becomes `and`
- `alias.ip` becomes `host` if a single host or `net` if a network.
- `ip.src` becomes `src host` if a single host or `src net` if a network.
- `ip.dst` becomes `dst host` if a single host or `dst net` if a network.
- Use CIDR notation when listing a network (that is, 10.10.10.0/24).
- `||` becomes `or`
- `!` becomes `not`
- Multiple rules must be joined with `and`.

The manual for TCPDump also gives examples of filters and strings that can be used: http://www.tcpdump.org/tcpdump_man.html

Additionally, the following site provides an excellent reference for BPF-style packet filters: http://biot.com/capstats/bpf.html

Caution: If you are capturing `vlan` tagged packets, above standard bpf filter may not work. For example, if you use `not (udp port 123)` to filter `vlan` tagged NTP traffic on udp port 123, it will not work. This is because the bpf filter machinery is simple and does not account for protocols not referenced in the rule. So the OS executing the bpf filter will look for the `udp port` values at the byte offset they would occur in a standard Ethernet/udp packet; but the optional `vlan` tag fields in the Ethernet header pushes those values by 4 bytes, thus the bpf filter rule will fail. To fix it, you need to change the bpf filter to: `not (vlan and udp port 123)`.

(Optional) Configure a Decoder to Capture Data Across All Types of Network Interfaces

The `packet_mmap_,ALL` adapter is capable of capturing across all types of network interfaces at the same time. For example, this can include things like physical network interfaces over different media types and tunnel interfaces.

The default behavior of the `ALL` adapter is to capture from all interfaces from the system, except for the hard-coded defaults of `lo`, `eth0`, and `em1`.

You can select any subset of the capture interfaces by editing the Decoder configuration node `/decoder/config/capture.device.params` to include an `interfaces=` parameter. The `interfaces` parameter contains a comma-separate list of interfaces that are used for capture. Instead of using all interfaces for capture, only the specified interfaces are used.

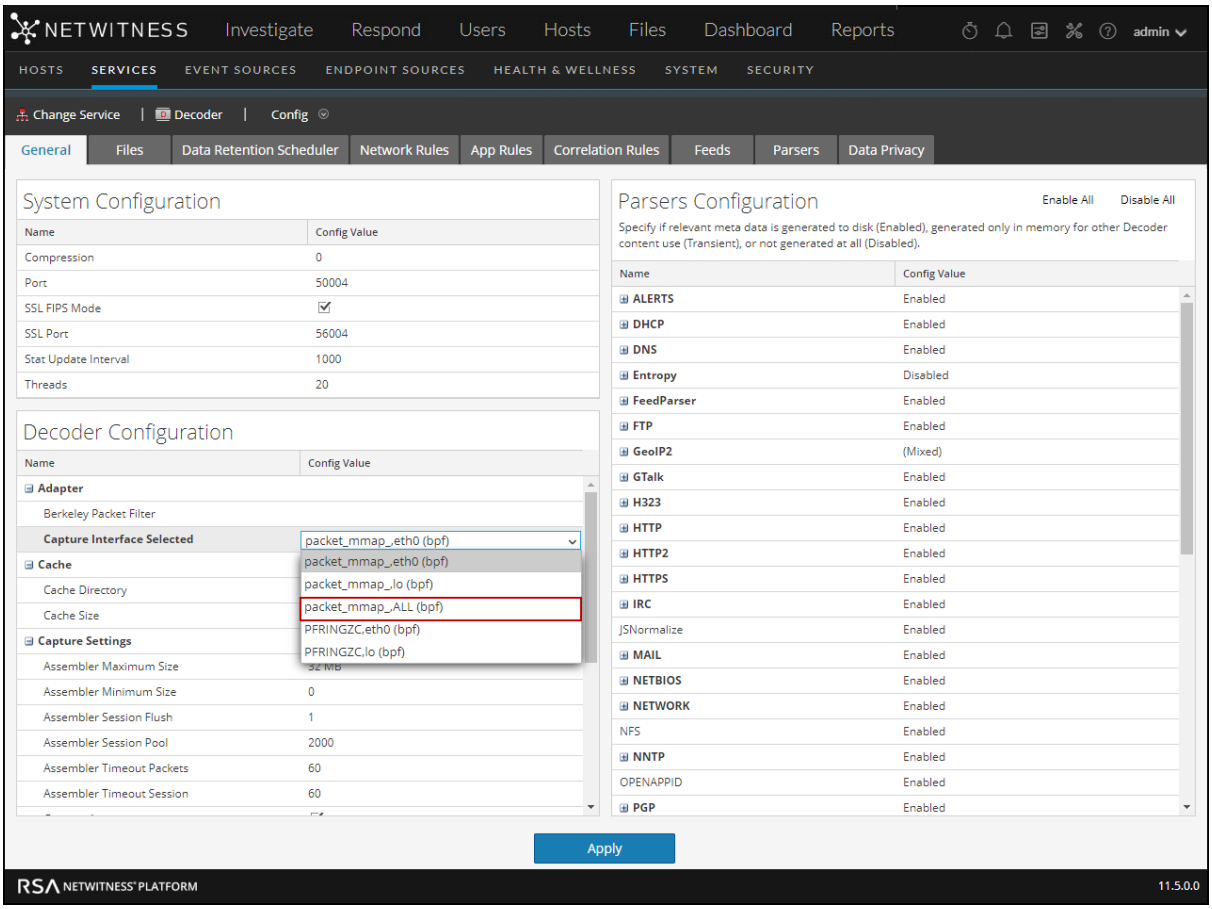
For example, if you want to force capture on interfaces `em1`, `em2`, and `em4`, and ignore `em3`, you can select the `packet_mmap_,ALL` adapter, and then add this line to `capture.device.params`: `interfaces=em1,em2,em4`

Note: Decoder automatically performs decapsulation of Virtual Extensible LAN (VXLAN) protocol from network traffic on UDP-4789 so that parsing can take place on the decapsulated Ethernet frames. During the decapsulation, the VXLAN ID is extracted and stored in the VXLAN meta key. By default, the VXLAN ID extraction is enabled. To disable, navigate to the `VlanGre` parser in the Parser Configuration settings. Click the drop-down list for `vlan` in the **Config Value** column and select **Disabled**. Each parser is configurable in the Services Config View - General Tab. The Parser Configuration panel provides a way to enable or disable parsers to use on Decoders in addition to limiting the metadata that the parser creates.

Note: Using the `interfaces` parameter to select `eth0`, `lo`, or `em1` overrides the default behavior, which is to drop traffic from those ports.

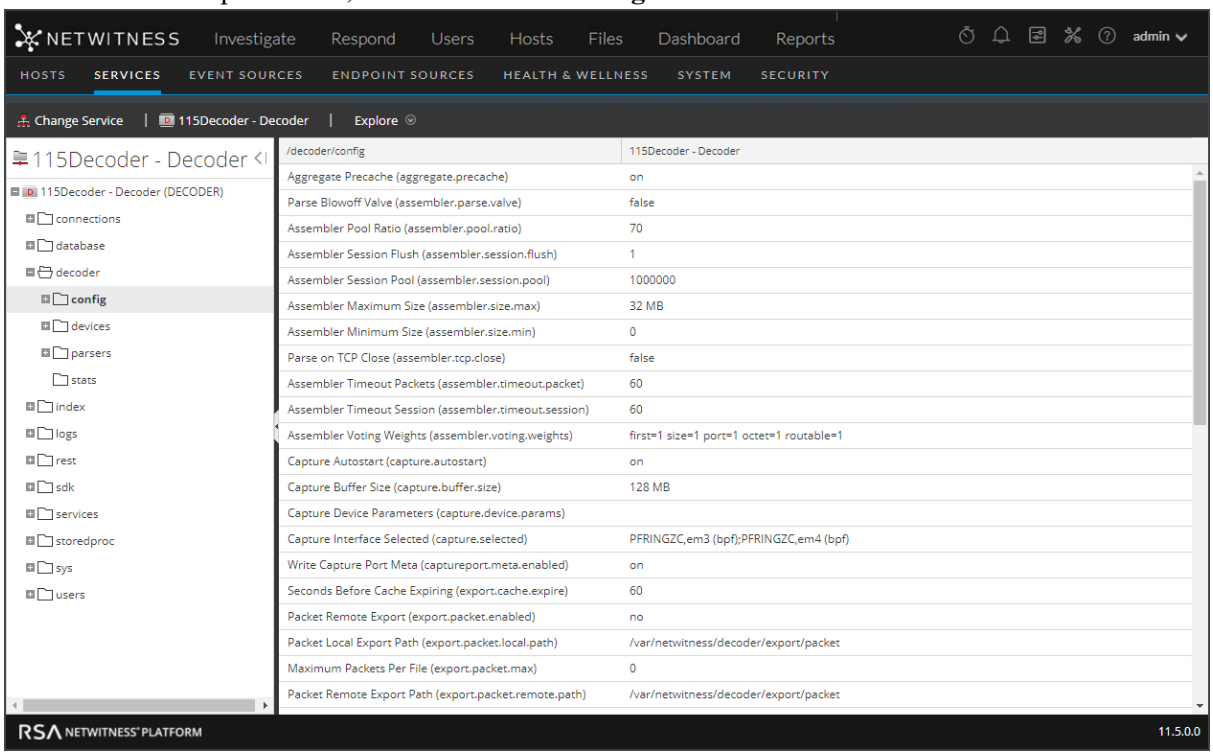
To configure the `packet_mmap_,ALL` adapter to capture from specific interfaces instead of all interfaces:

1. Go to (Admin) > **Services**, select the Decoder service and > **View > Config.**
2. In the **Services Config view**, set **Capture Interface Selected** to `packet_mmap_,ALL` adapter.

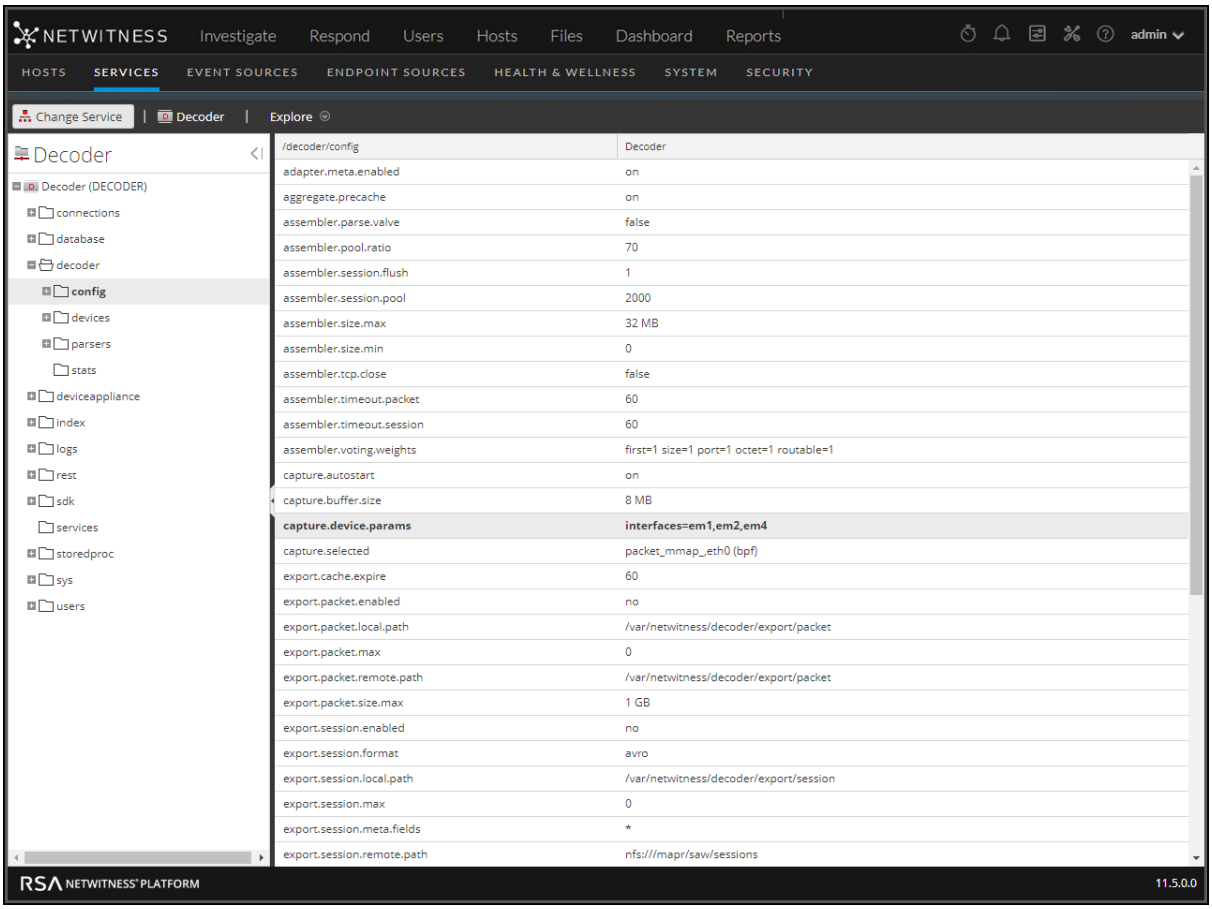


3. Click **Apply**, and then restart the Decoder service.

4. Go to (Admin) > **Services**, select the Decoder service, and click > **Config** > **Explore**.
5. In the Services Explore view, select **decoder** > **config**.



6. Click in the values column next to `capture.device.params`, type **`interfaces=em1,em2,em4`**, and press **Enter**.



The change goes into effect immediately; only traffic on `em1`, `em2`, and `em4` interfaces is captured.

(Optional) Configure Meta-Only Decoders

You can configure Decoders so that packets and logs can be processed, and then dropped before they are written to disk. This is called a Meta-Only Decoder, which uses the Meta-Only license, and can save storage space (however, analysts cannot reconstruct events in Investigate if you use this option). The configuration option `/decoder/config/packet.write.disabled` controls this feature. If this option is set to `true`, all packets are dropped after parsing, so they are never written to the database. This applies to both Log and Network Decoders. The ingested logs and packets flow through the system normally so that parsing and other operations are not impacted. The default setting is `false`, which preserves the normal behavior of writing packets to disk.

Note: You must purchase the Meta-Only license before you can realize the full benefit of using this option. For information about purchasing licenses, see the *Licensing Management Guide for RSA NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

To configure a Meta-Only Decoder where packets and logs are parsed and not written to disk:

1. Go to  (Admin) > **Services** and select a Decoder.
2. Click  > **View** > **Explore**, and in the left panel, expand **decoder** and click **config**.

3. In the right pane, go to `packet.write.disabled`, and change the value from `false` to `true`.

Note: You can make this configuration update while capture is running, but the update does not take affect until capture is restarted.

For an example of how to use the Meta-Only license to apply centrally-managed capture policies across your Network Decoders, see (Optional) Configure Selective Network Data Collection

(Optional) Configure Selective Network Data Collection

Selective network data collection gives administrators the ability to apply centrally managed capture policies across their Network Decoders. This results in better use of service resources, including hard drive space, which leads to more predictable costs and lessens the burden of managing multiple services. Administrators can determine which traffic is stored and how it is stored by using policies. Each policy contains a list of supported base protocols and definitions for handling any other protocols that are detected.

Note: Administrators can use this feature to create policies for metadata-only while using the cost-effective Meta-Only license, without having to pay for additional storage or throughput licensing. To use this license, work with your RSA Customer Support Representative to purchase the license, and then enable the functionality as described in (Optional) Configure Meta-Only Decoders.

The administrator can choose to deploy predefined policies that capture:

- All base and other protocols (Full Capture - All Protocols)
- Only metadata on all base and other protocols (Capture Meta Only - All Protocols)
- Only metadata on all base protocols and drop all other protocols (Capture Meta on Base Protocols, Drop all other protocols)
- All base protocols and only metadata on all other protocols (Full Capture on Base Protocols, Meta only on all other protocols)

The predefined policies are not configurable. The only way you can edit predefined policies is by assigning services (such as Decoders) to deploy them.

Administrators can create custom policies to give further control over the deployment. A base set of protocols are available for alterations by the administrator, allowing you to choose what level of capture you prefer on a per-protocol basis. If you are only making slight changes, a good start for customization is to clone one of the predefined polices and alter it. These centrally-managed policies are then applied to services (Network Decoders) to allow handling multiple use cases across your environment.

Here are some example use cases and how an administrator would deploy these.

Use Case 1: Predefined Policy for Internal Assets

To gain further visibility into internal lateral traffic going East-West through the environment, a Network Decoder is deployed in the interior of the network to capture this traffic. Applying full packet capture for this segment may generate too much data to manage at a reasonable cost point. Instead, the predefined "Capture Meta Only – All Protocols" policy can be applied to achieve the same parsing and detection while only saving the metadata, limiting the amount of storage required of the solution. In this situation, the only steps you must complete is to choose the predefined policy, edit it, and then enable it by assigning it to the appropriate internal services (Decoders). For an example of how to set up this policy, see Use Predefined Policies.

Use Case 2: Custom Policy for DMZ Assets

To protect the public-facing web servers, a Network Decoder is deployed into the demilitarized zone (DMZ). The network communications between these web servers and other network devices are a subset of the base protocols. To make sure you can harvest forensic evidence for any communications related to these high-value assets, you would want to do full packet capture on HTTP, SMTP, FTP, and DNS protocols. You still want visibility into any other traffic these servers are engaged in, so you would prefer metadata generated for all other protocols. The suggested steps to create this custom policy would have you choose the predefined policy "Full Capture on Base Protocols, Meta only on all other protocols" and then clone it to adjust it to the protocols aligned with the DMZ scenario. During the clone process, the name and description can be updated to identify the new purpose of the policy. The define policy step allows all the base protocol rules inside the policy except HTTP, SMTP, FTP, and DNS to be altered from "Collect All" to "Collect Meta Only" to achieve the desired policy. The last steps are assigning the policy to the appropriate DMZ services (Decoders). For an example of how to set up this policy, see Create New Policies from Predefined Ones.

Note: When a Decoder is being managed by a selective collection policy, it is possible that the Decoder specifications in the policy could be overwritten if the Decoder configuration is edited outside of the policy. To try and avoid this situation, the following warning message is displayed at the top of the Parsers section of the Decoder Configuration page: `This Decoder is being managed by the following capture policy: <policy-name>. Some parsers are required by the capture policy, and changes made here may overwrite that policy. Review of the capture policy is recommended.`
Another warning message is displayed at the top of the Application Rule section: `This Decoder is being managed by the following capture policy: <policy-name>. Some application rules are required by the capture policy, and changes made here may overwrite that policy. Review of the capture policy is recommended.`

The following sections describe configuring and using selective network data collection.

- Use Predefined Policies
- Create New Policies from Predefined Ones
- Create Custom Policies
- Unpublish Policies
- Delete Policies
- Supported Protocols for Selective Network Data Collection

For detailed information about the Selective Collection user interface, see:

- Services Config View - Capture Policies Tab
- Services Config View - Edit Capture Policies Wizard

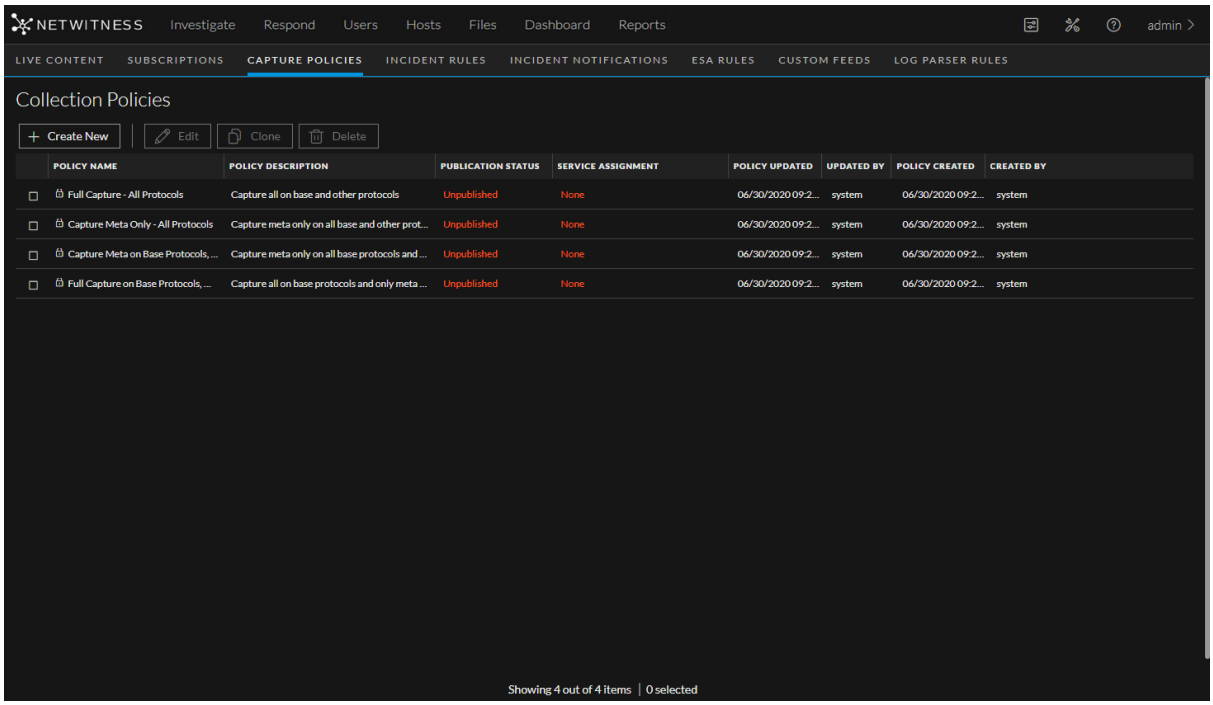
Use Predefined Policies

This procedure describes how to use the predefined policies. You can use this procedure to configure a policy for internal assets. For example, as described in Use Case 1: Predefined Policy for Internal Assets , to gain further visibility into internal lateral traffic going East-West through the environment, you could deploy a Network Decoder in the interior of your network to capture the traffic. Applying full packet capture for this segment may generate too much data to manage at a reasonable cost point. Instead, use the predefined "Capture Meta Only – All Protocols" policy to achieve the same parsing and detection while only saving the metadata, limiting the amount of storage required for this solution.

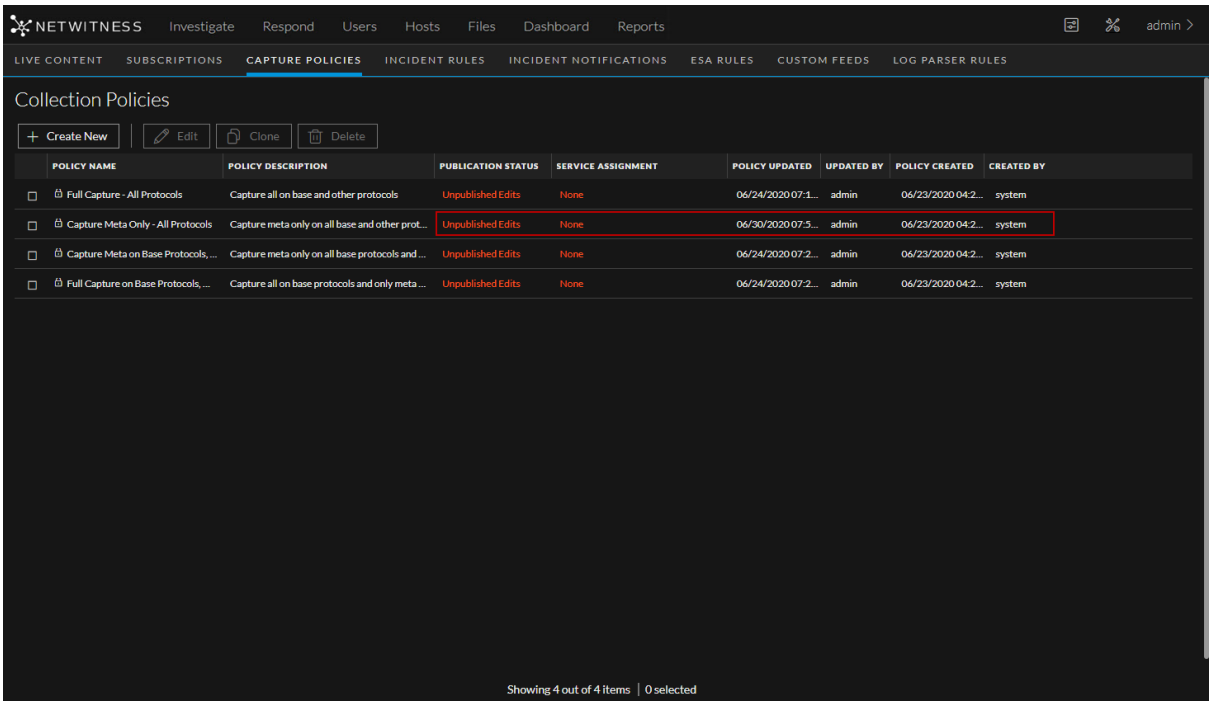
In this situation, the only steps you must complete is to select the "Capture Meta Only – All Protocols" predefined policy to edit it, enable it by assigning it to the appropriate Decoders, and then publish the policy to start collecting the data.

To work with a predefined policy:

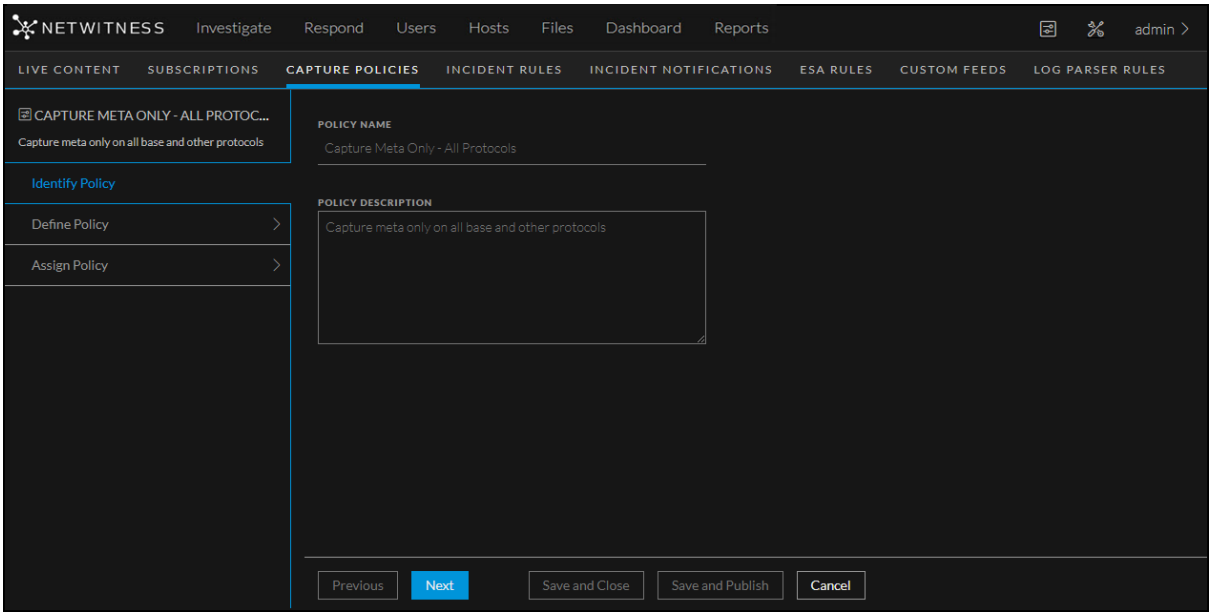
1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.



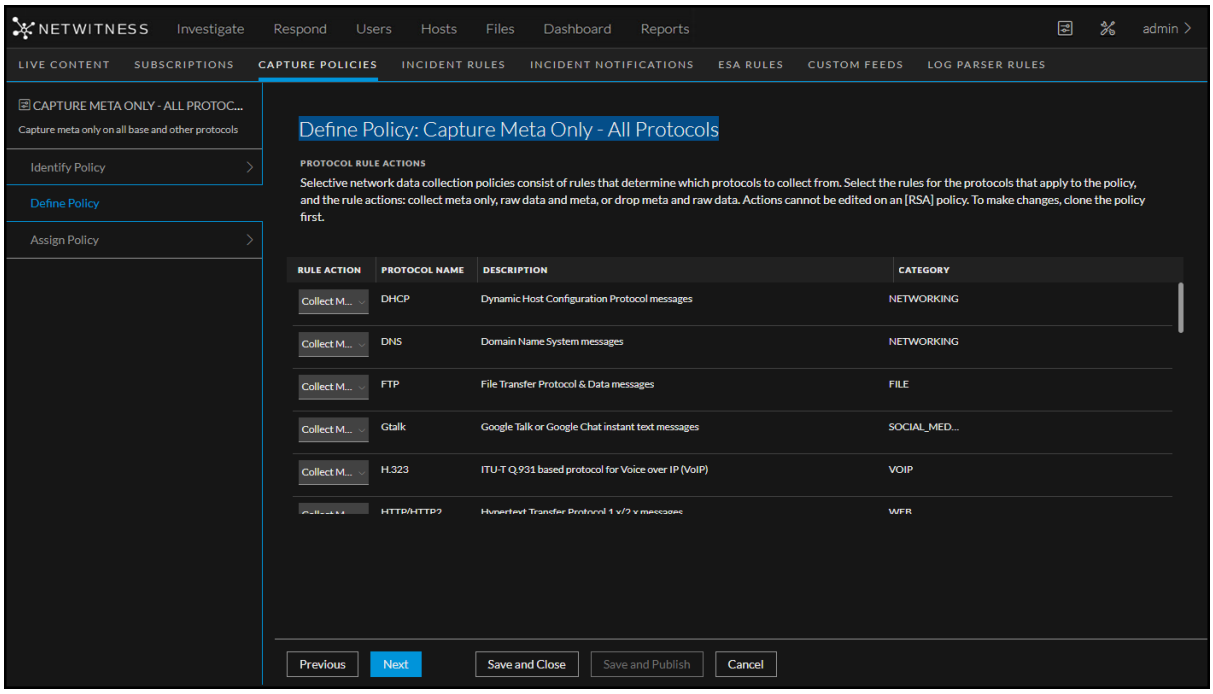
Notice that the PUBLICATION STATUS column lists each of the predefined policies with the status of **Unpublished**. The only way to edit a predefined policy is to assign Decoders to apply the policy. If changes are made to the Decoder assignments but are not published, the status changes to **Unpublished Edits**, the date the change was made is displayed in the POLICY UPDATED column, and the user name of the person who made the change is displayed in the UPDATED BY column. In predefined policies, the CREATED BY field is always **system**.
The following figure shows the statuses of the policies.



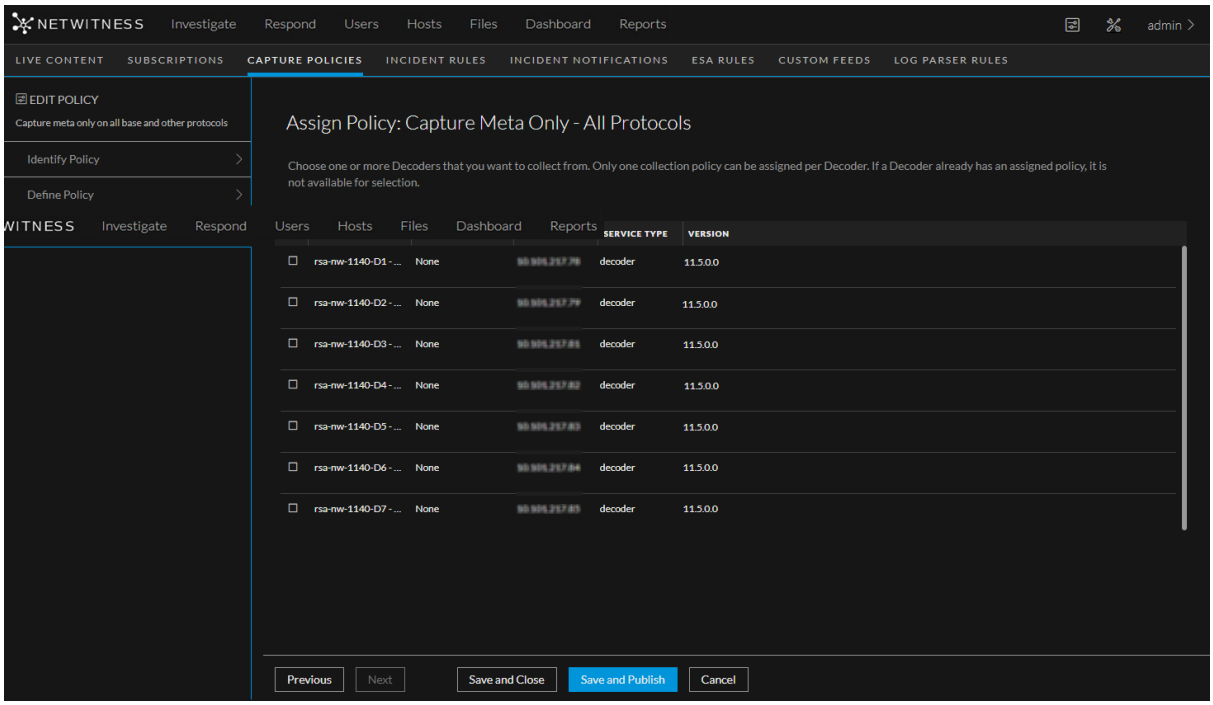
2. Select **Capture Meta Only - All Protocols**, and then click **Edit**.
The editing view is displayed with the **Identify Policy** option active.



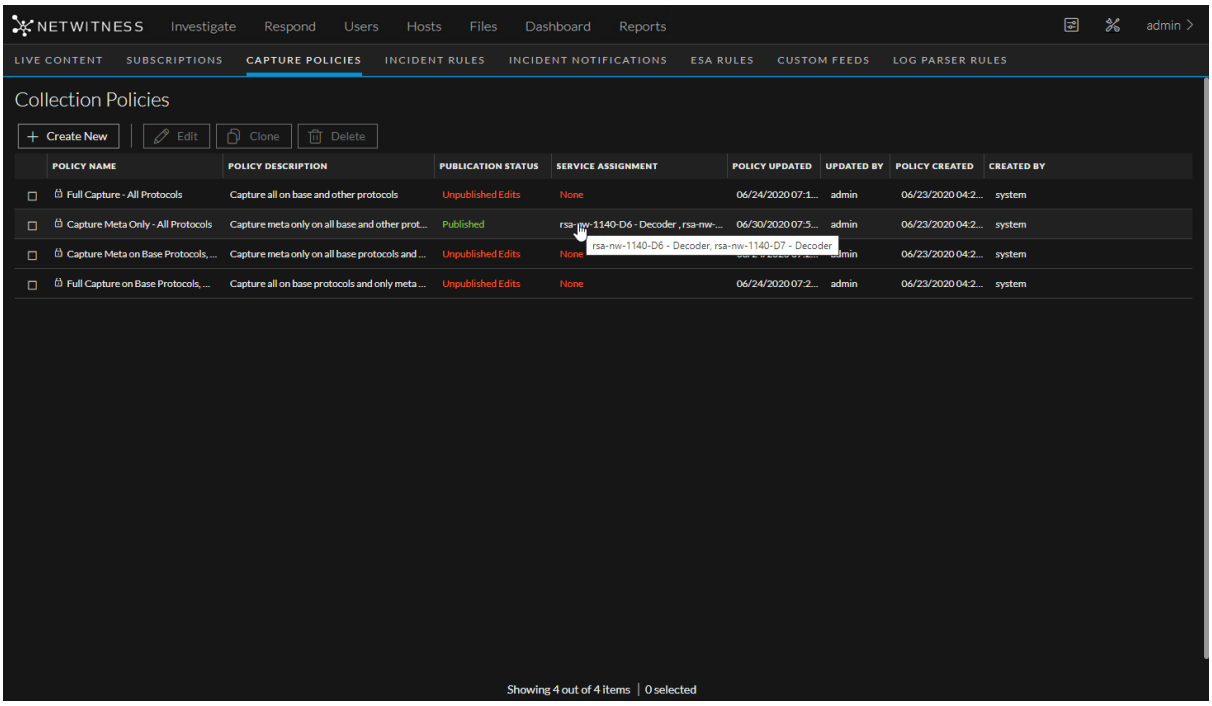
3. Because this is a predefined policy, you cannot change the policy name or description. Click **Next**.
The Define Policy: Capture Meta Only - All Protocols page is displayed.



4. Because this is a predefined policy, you cannot edit this page. Review the rules on this page to understand the actions that will be enabled when you publish this policy. Click **Next**.
   The Assign Policy page is displayed.



5. Select the Decoders that you want to collect from, and then click **Save and Publish**.
   The publication status moves to **Updating**. It changes to **Published** when the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as **Published**, so even if not all of the Decoders are available (and the publication status is not **Published**), data will begin to be collected on the Decoders that are available. When the status moves to **Published**, all the services are displayed in the Service Assignment column. You can hover over the names of the Decoders to see their full names, as shown in the following figure.



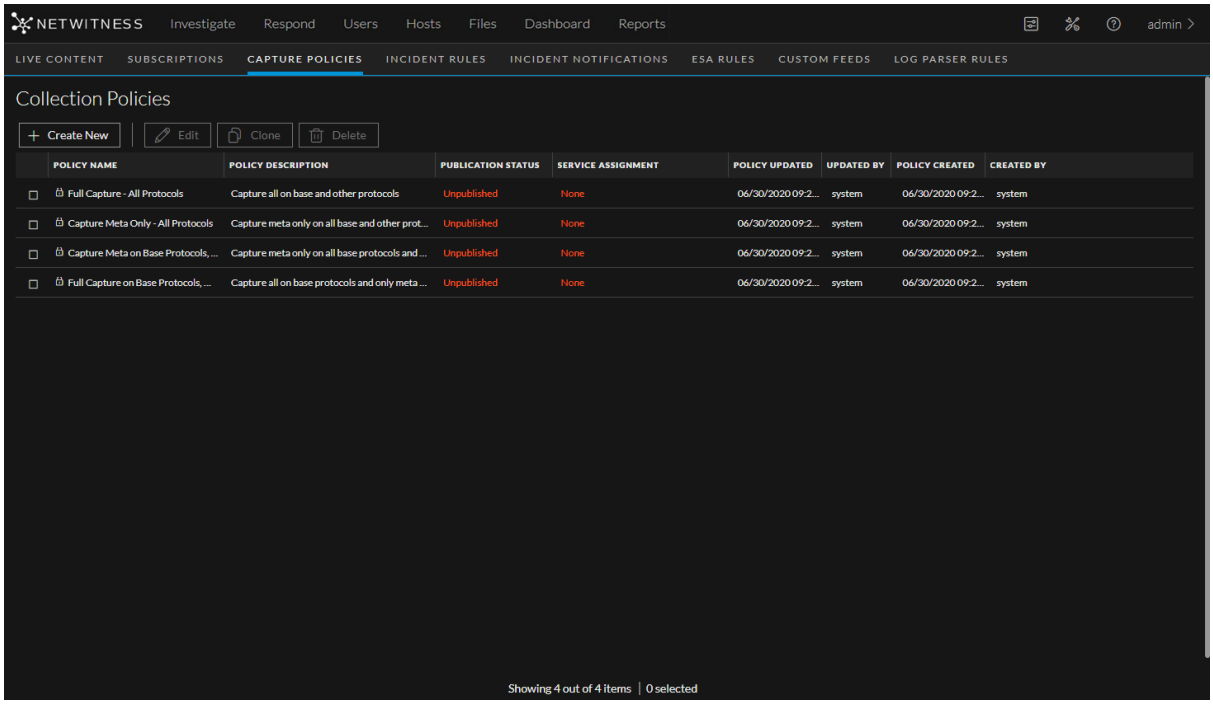
Create New Policies from Predefined Ones

This procedure describes the steps for creating a new policy from a predefined one, and provides an example of how you might use this policy.

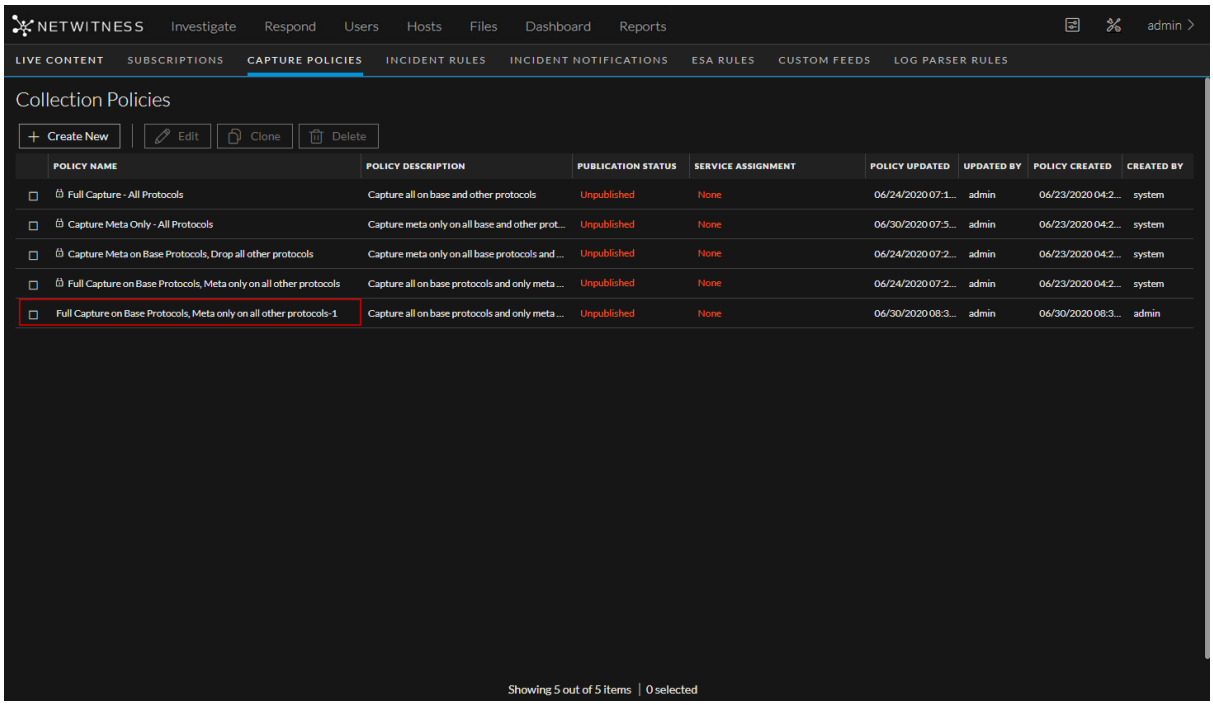
For example, as described in Use Case 2: Custom Policy for DMZ Assets , suppose a Network Decoder has been deployed into the demilitarized zone (DMZ) to protect public-facing web servers. The network communications between these web servers and other network devices are a subset of the base protocols. To make sure they can harvest forensic evidence for any communications related to these high value assets, you would want to get full packet capture on HTTP, SMTP, FTP, and DNS protocols. To provide visibility into any other traffic these servers are engaged in, you would want to generate metadata for all other protocols.

To configure a policy that would achieve the results described here, you would select the predefined policy "Full Capture on Base Protocols, Meta only on all other protocols," clone it, and then edit it to align the protocols with the DMZ scenario. During the clone process, you can update the name and description to identify the new purpose of the policy. In the step where you define the policy, you can change all the base protocol rules inside the policy (except HTTP, SMTP, FTP, and DNS) from "Collect All" to "Collect Meta Only" to achieve the desired results. The last steps assign the policy to the appropriate DMZ Decoders.

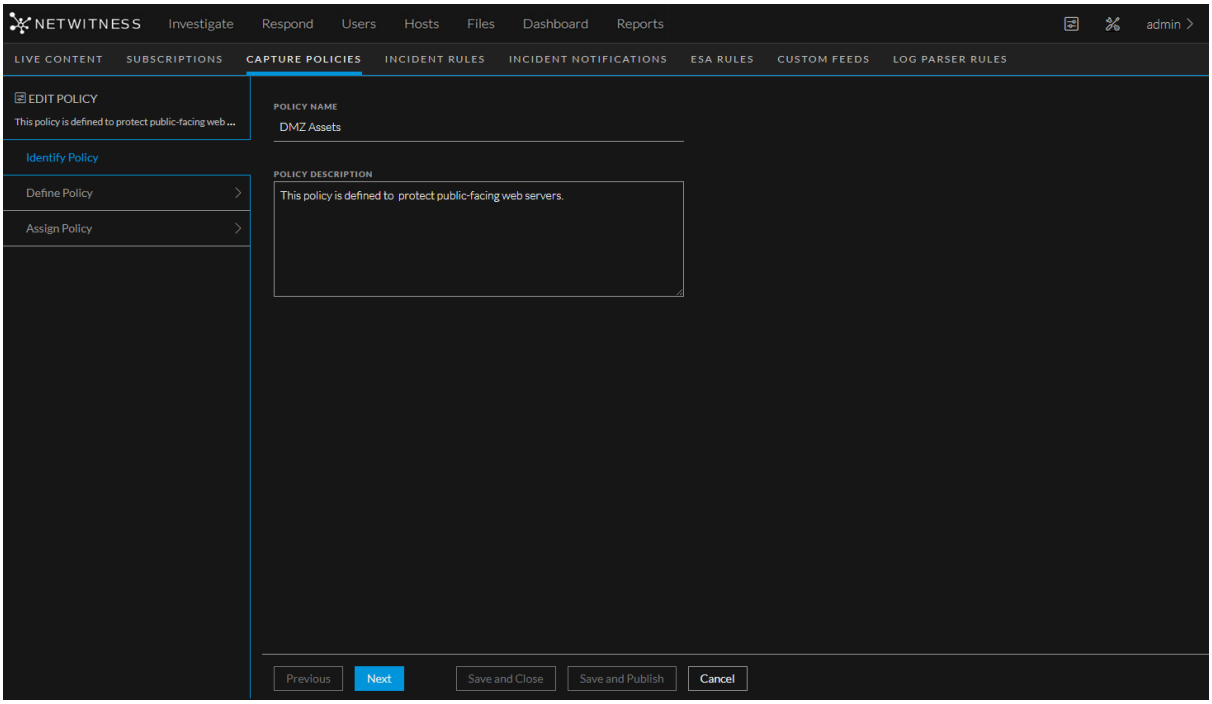
1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.



2. Select **Full Capture on Base Protocols, Meta only on all other protocols** and then click **Clone**. The Collection Policies view is displayed with the cloned policy at the bottom of the list, identified by appending -1 to the policy name.

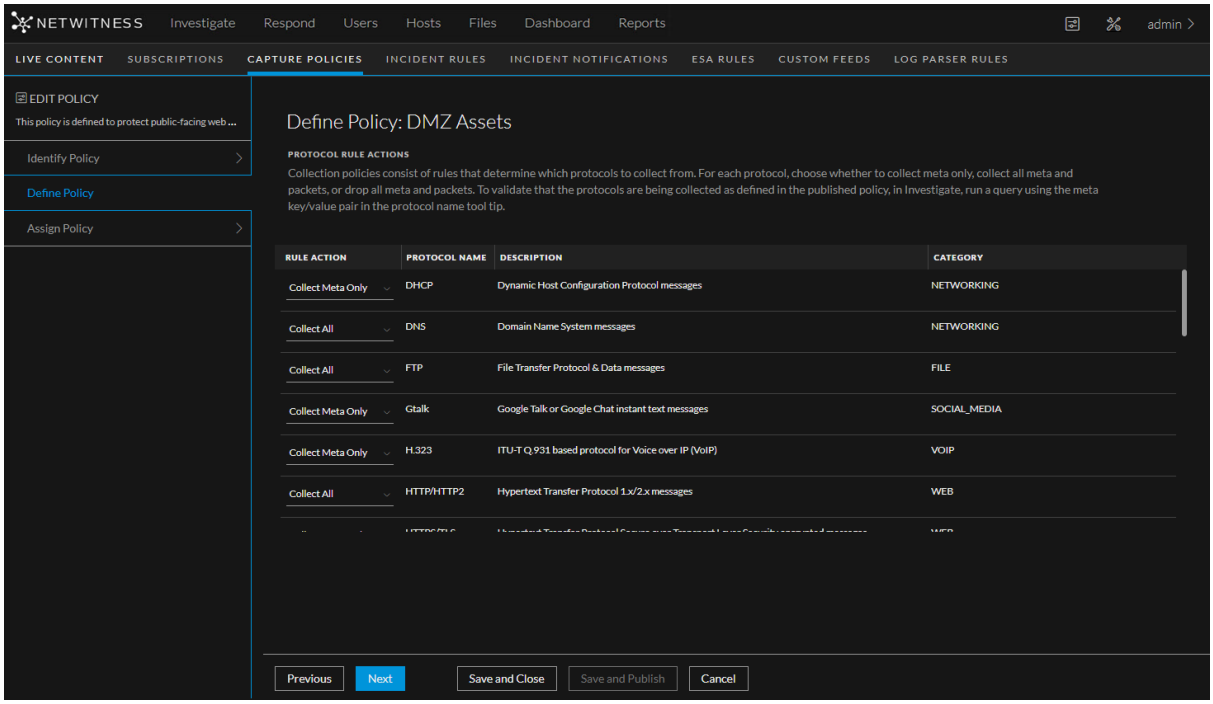


3. Select the cloned policy and click **Edit**.
   The editing view is displayed with the **Identify Policy** option active.
   In this example, we are naming this policy "DMZ Assets" and providing the description "This policy is defined to protect public-facing web servers."

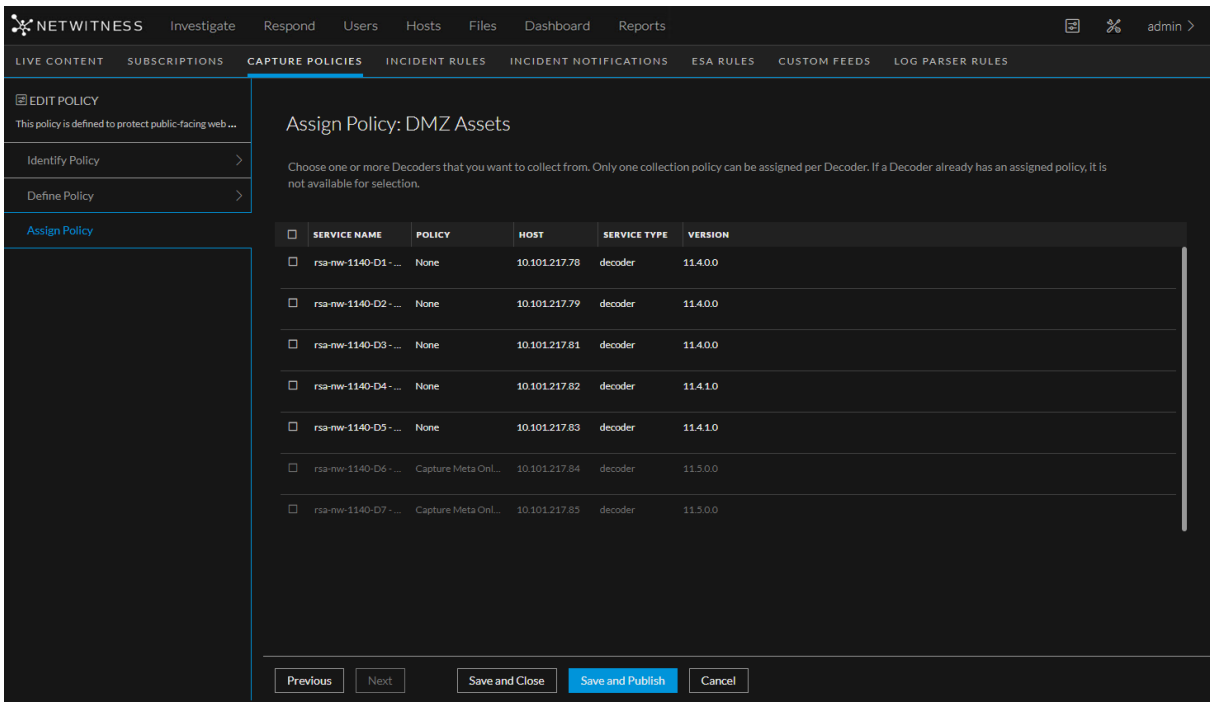


4. Click **Next**. The Define Policy view is displayed, with a list of rules you can use to define the policy. You can select the following options for rule actions:
   - **Collect Meta Only**: Collect metadata
   - **Drop All**: Drop metadata and network packets
   - **Collect All**: Collect metadata and network packets

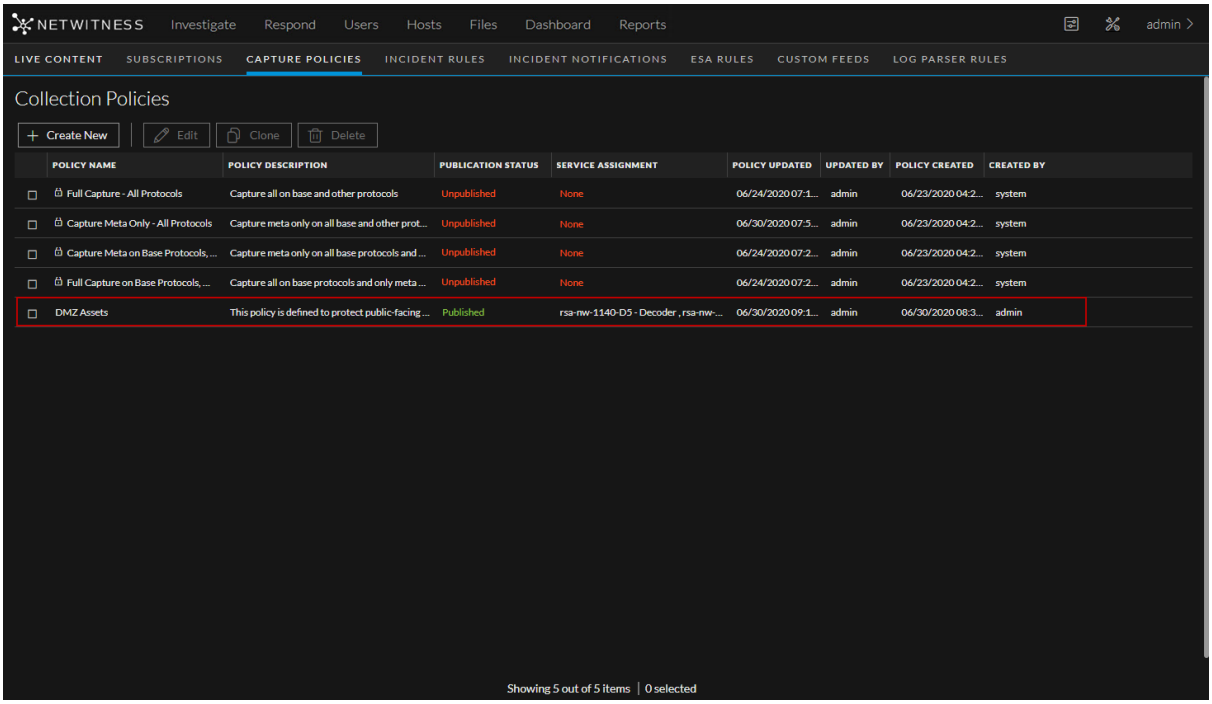
To collect the targeted data for the DMZ Assets policy, we are going to change all the base protocol rules inside the policy (except HTTP, SMTP, FTP, and DNS) from **Collect All** to **Collect Meta Only**, as shown below.



5. Click **Next**.
   The Assign Policy page is displayed.



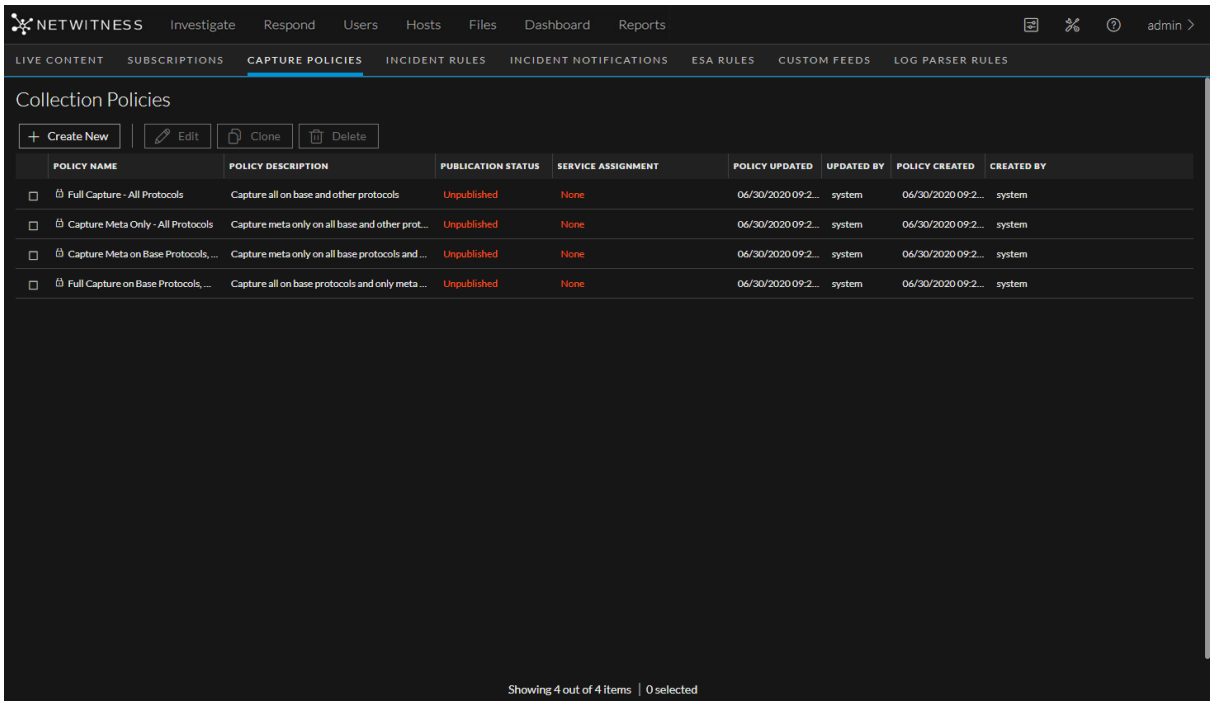
6. A Decoder can only be assigned to one policy. Decoders that are not available are greyed-out, because they are already assigned to other policies. Select the Decoders to collect the targeted data, and then click **Save and Publish**. The publication status moves to **Updating**. It changes to **Published** when the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as **Published**, so even if not all of the Decoders are available (and the publication status is not **Published**), data will begin to be collected on the Decoders that are available.
   Your new policy is displayed in the Collection Polices view, with the services displayed in the Service Assignment column.



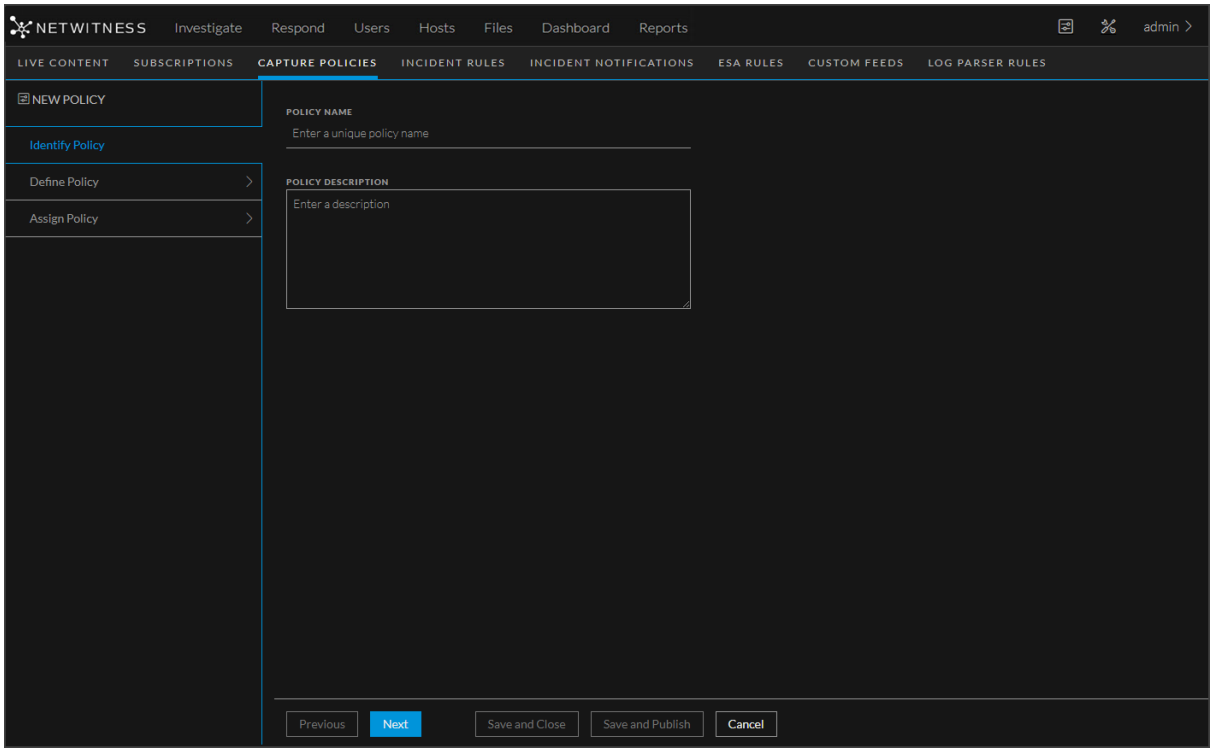
Create Custom Policies

You can create a custom Selective Network Decoder policy, where you can select specific protocols on which to collect data, as well as the rules that apply to them.

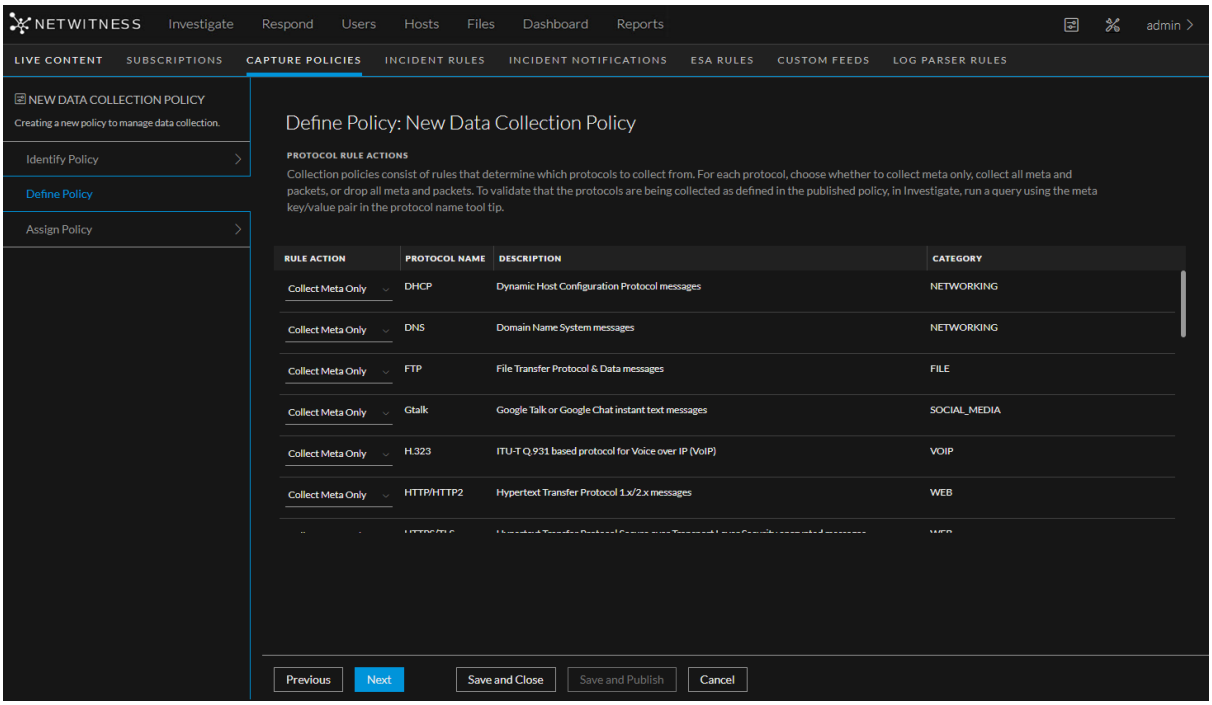
1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.



2. Click **Create New**.
   The editing view is displayed with the **Identify Policy** option active.

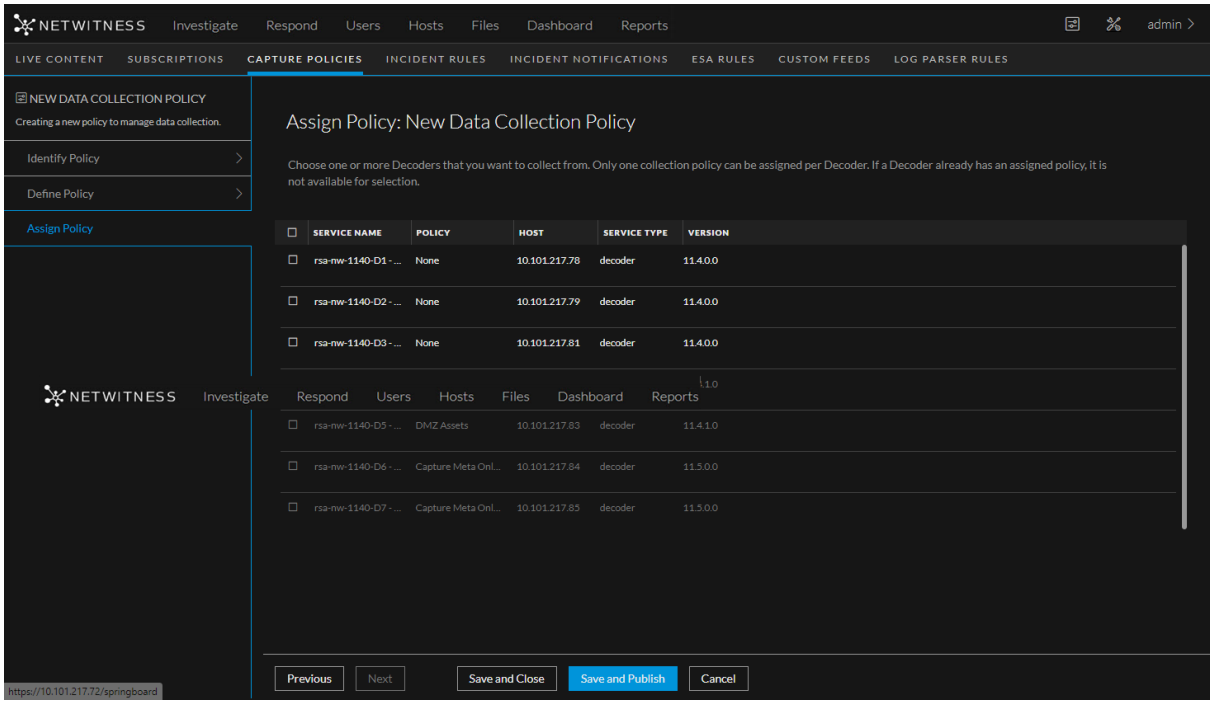


3. Enter a name and description for your policy, and then click **Next**.
   The Define Policy view is displayed.



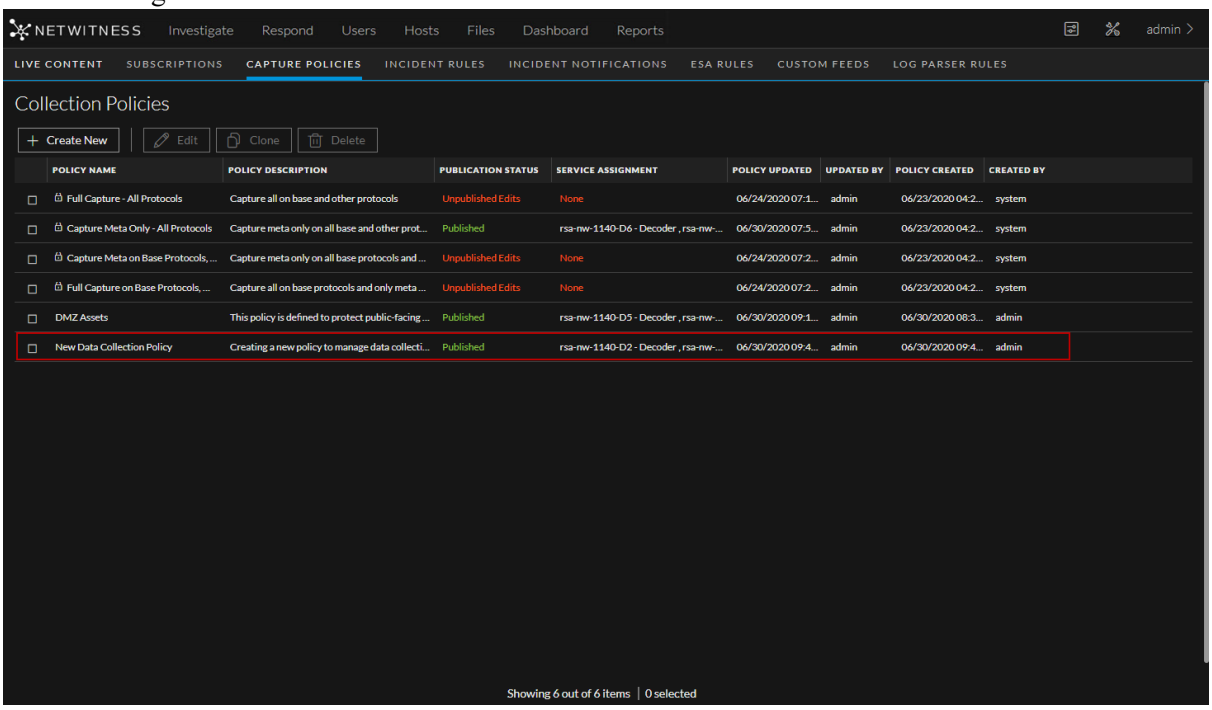
4. Define the actions for each of the protocols. The options for **RULE ACTION** are:
   - **Collect Meta Only**: Collect metadata
   - **Drop All**: Drop metadata and network packets
   - **Collect All**: Collect metadata and network packets

5. Click **Next**.
   The Assign Policy view is displayed.



6. A Decoder can only be assigned to one policy. Decoders that are not available are greyed-out, because they are already assigned to other policies. Select the Decoders to collect the targeted data, and then click **Save and Publish**. The publication status moves to **Updating**. It changes to **Published** when the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as **Published**, so even if not all of the Decoders are available (and the publication status is not **Published**), data will begin to be collected on the Decoders that are available.

Your new policy is displayed in the Collection Polices view, with the services displayed in the Service Assignment column.

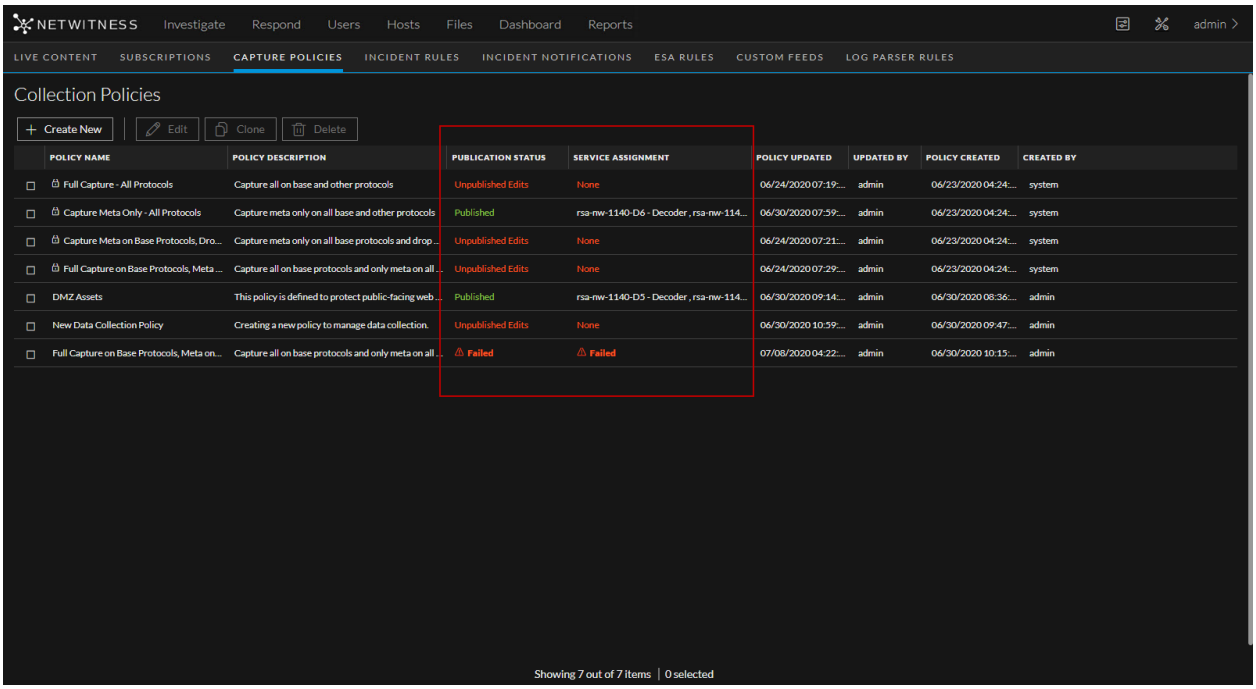


Verify Published Policies

When a policy has been published, the policy is applied to the Decoders that were designated in the policy assignment step. The Decoders that are assigned to the policy start collecting the data defined in the policy. There are a few ways to validate the policy is functioning as expected.

View Publication Status

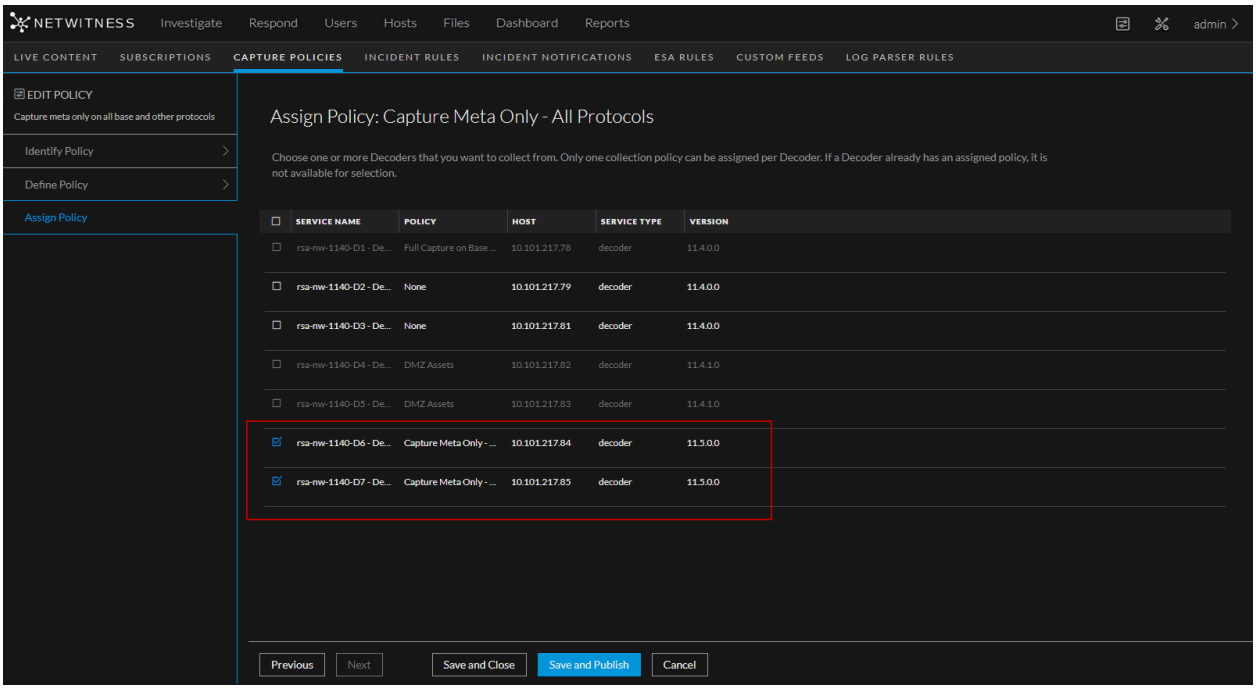
Publication status is indicated in the Capture Polices view in the PUBLICATION STATUS column. You can also see which Decoders are assigned to the policy in the SERVICE ASSIGNMENT COLUMN.



The publication statuses are:

- **Unpublished** - the rule actions of the policy have been defined and Decoders may have been assigned to the policy, but the policy has not been deployed to any Decoders.
- **Unpublished Edits** - policies that have been updated but are not deployed to Decoders. Clicking **Save and Close** after making edits keeps them from automatically being published, allowing the administrator to wait until a specific outage window to publish them. This would display as **Unpublished** (for policies not previously published to Decoders) or **Unpublished Edits** (for policies that have been previously published to Decoders, but the updates have not been deployed to Decoders).
- **Published** - the policy has been successfully deployed to all the Decoders that you selected. However, the Decoders that are available start gathering data immediately, before the policy is displayed as Published, so even if not all of the Decoders are available (and the publication status is not Published), data will begin to be collected on the Decoders that are available.
- **Failed** - policies that failed to execute, for example, if a Decoder is offline or a system is down.

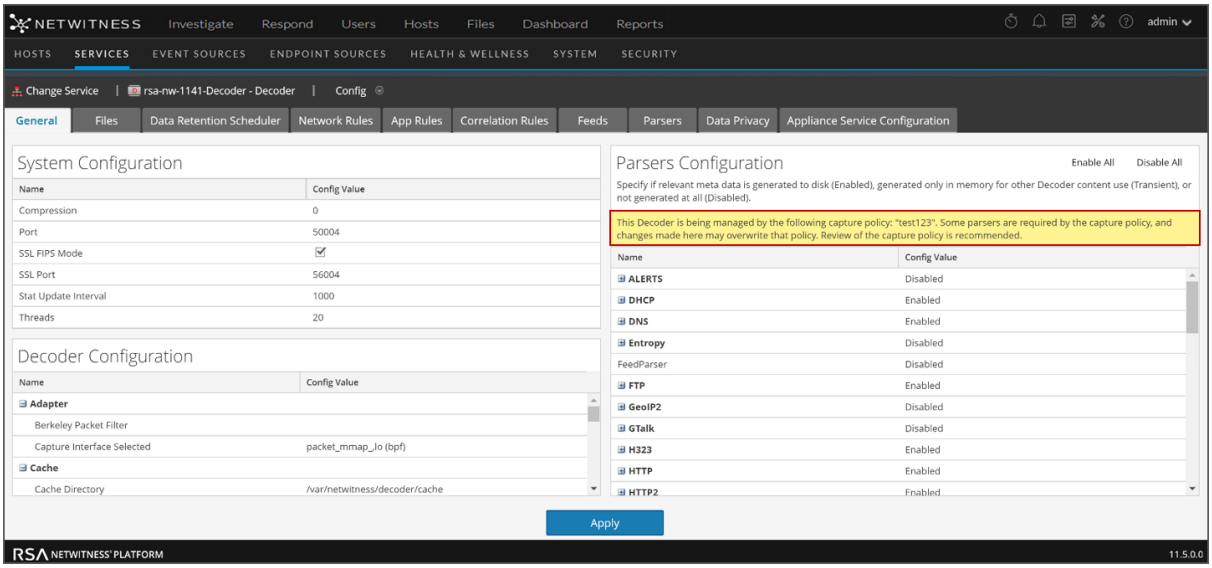
You can also select a published policy, and then click **Edit** to verify the Decoders that are assigned to the policy.



Verify Policies in Admin > Services

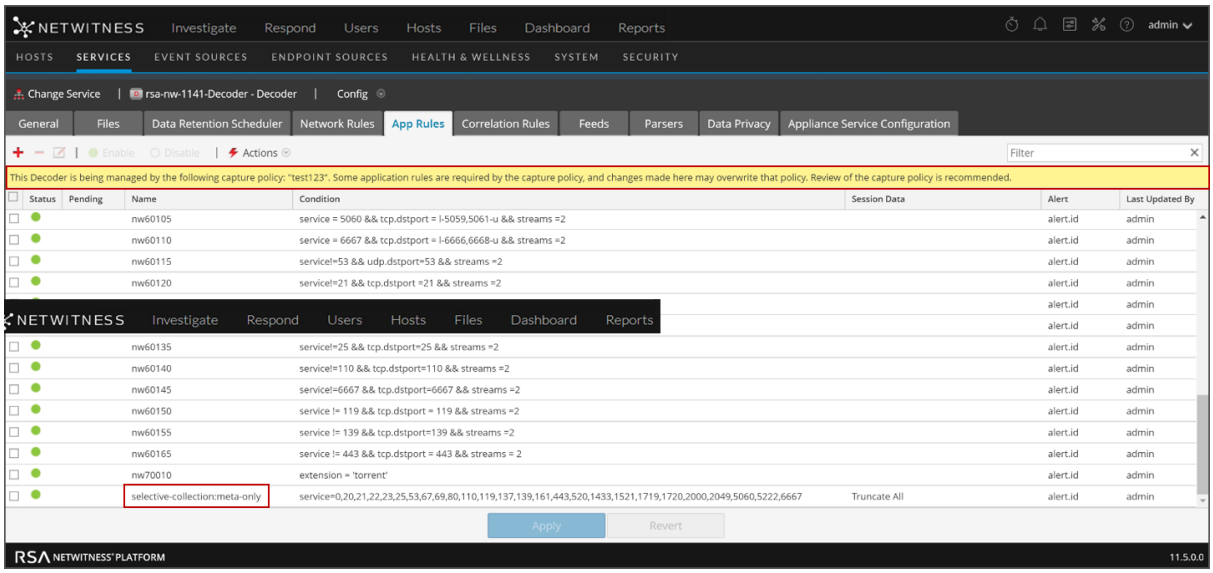
1. As an administrator, log in to the NetWitness user interface to view the Decoders that have been assigned to the policy.

2. Go to  (Admin) > **Services** > **<decoder name>** >  **(Actions)** > **View** > **Config**. On the General tab, in the right panel, there is a list of parsers that are installed on the Decoder by default. If a policy is administering this Decoder and requires any of those parsers to function, a message is displayed as shown below.



3. On the **App Rules** tab, a similar message indicates that some of the rules are required for the policy to function. You can search for "selective collection" to find the specific rules required, as the rule names contain that information.

Note: In Investigate, you must have the **Alert** session option selected for the selective collection policy to be displayed.



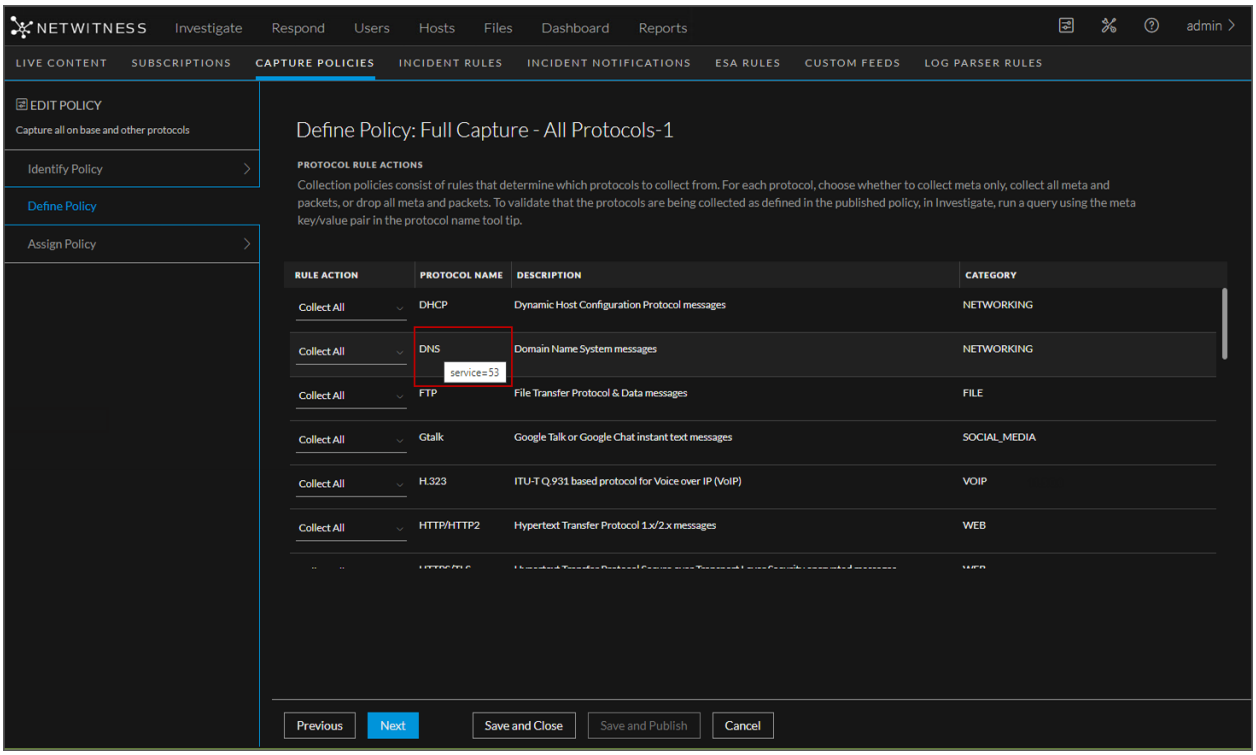
By default, the selective collection rules are added to the bottom of the rule set when the policy is applied. However, an administrator can directly affect the policy by altering the selective policy rules or by moving conflicting rules (rules with the same services) above the selective collection rules. In both cases, those changes take effect when they are saved locally on the individual Decoders, until the next time the policy is published from the central configuration page, when that policy will override changes made individually on the Decoders.

Note: There currently is no feedback in the centralized configuration page denoting that the configuration related to the policy has been altered locally on the Decoders. To avoid this situation, manage all the policy rules from the centralized configuration page. If that is not possible, as additional application rules outside of collection policies may be required, follow standard security practices, including separation of duties, to limit people who need access to alter application rules and general parsers.

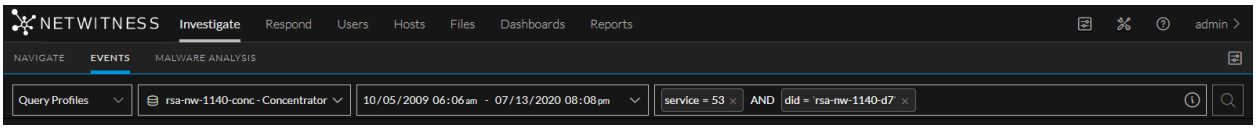
Verify Policies in the Investigate View

You can further verify that the policy is working as expected (the protocols are being collected as published) in the Investigate view. In the centralized configuration page for selective collection, the administrator can get a translation of what the query parameter would be for each individual protocol when viewing the rules inside the policy.

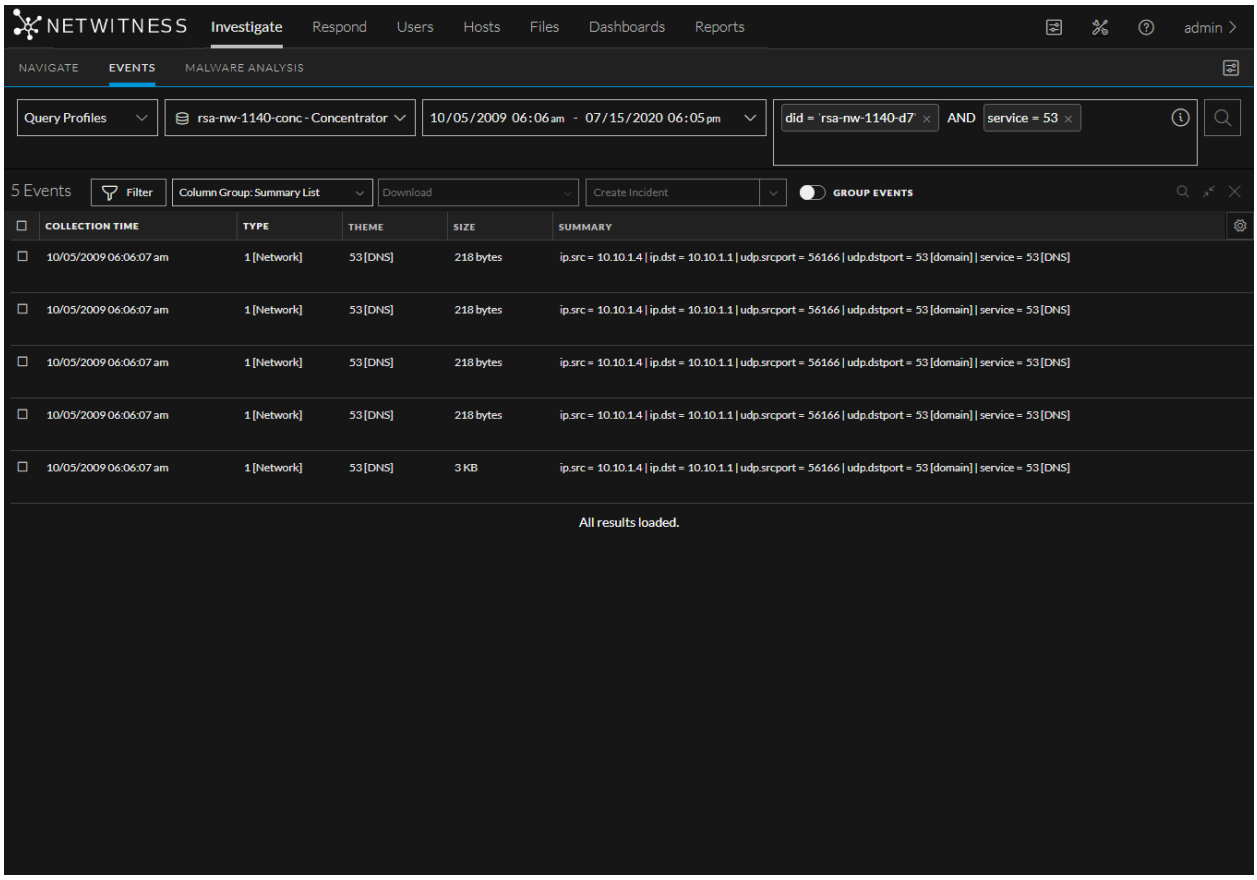
For example, in (Configure) > **Capture Policies** > **<select a policy>** > **Edit** > **Define Policy**, if you hover over the DNS protocol name, `service=53` is displayed.



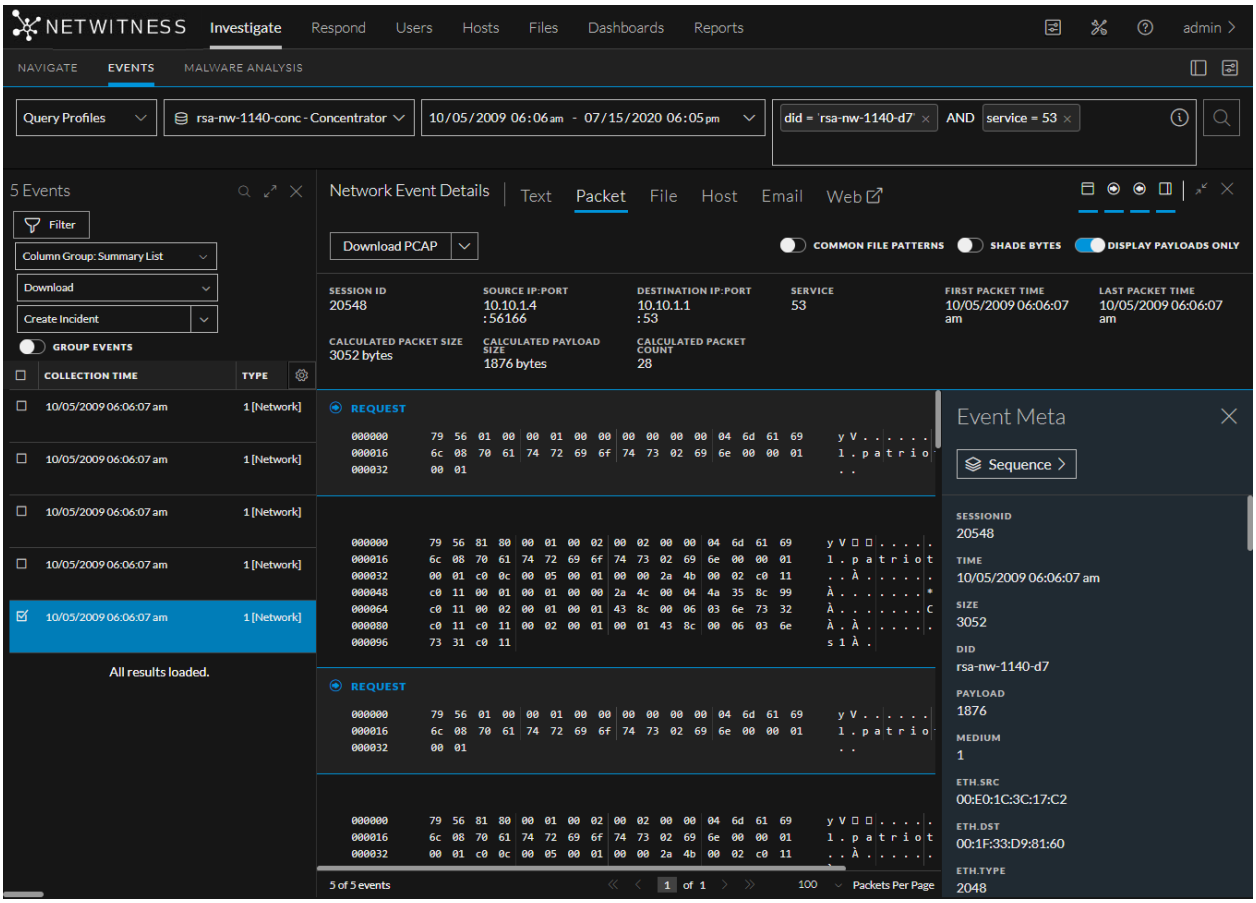
To determine if the appropriate action has been taken for DNS traffic, in **Investigate** > **Events** you can add the query parameter (for example, `service=53`) as a filter in the query bar and run the query against the Broker or Concentrator the Network Decoder is associated with for the time the policy has been active on the Network Decoder. To make sure you are only seeing the traffic from the Decoders that you want, you can add the Decoder Source, or `did` meta key as another filter in the query bar.



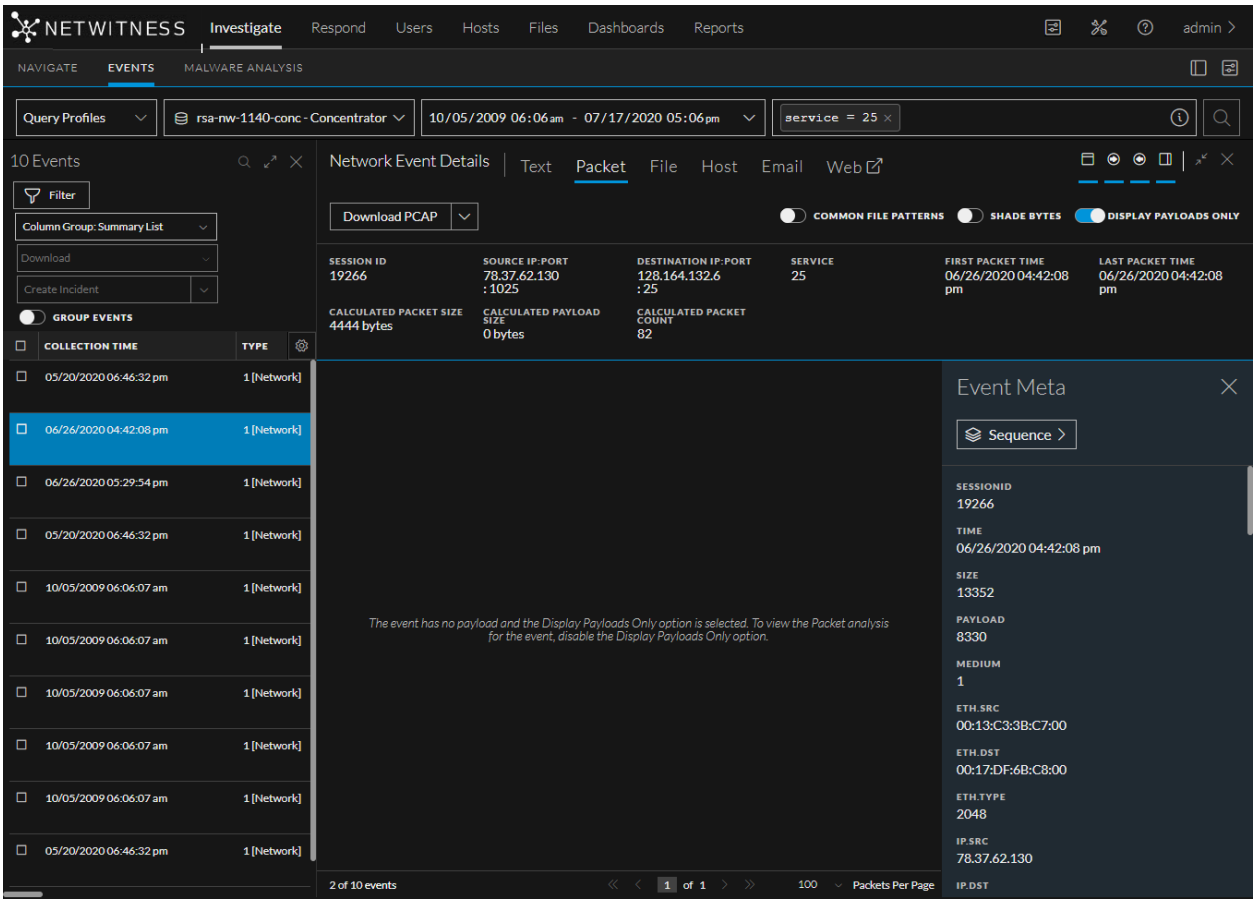
After running the query, if there are any related results, those events are displayed in the table (unless you applied an action to drop them).



If you chose the action to `collect all` you can click on one of the events and review the metadata and the payload information for that packet session.

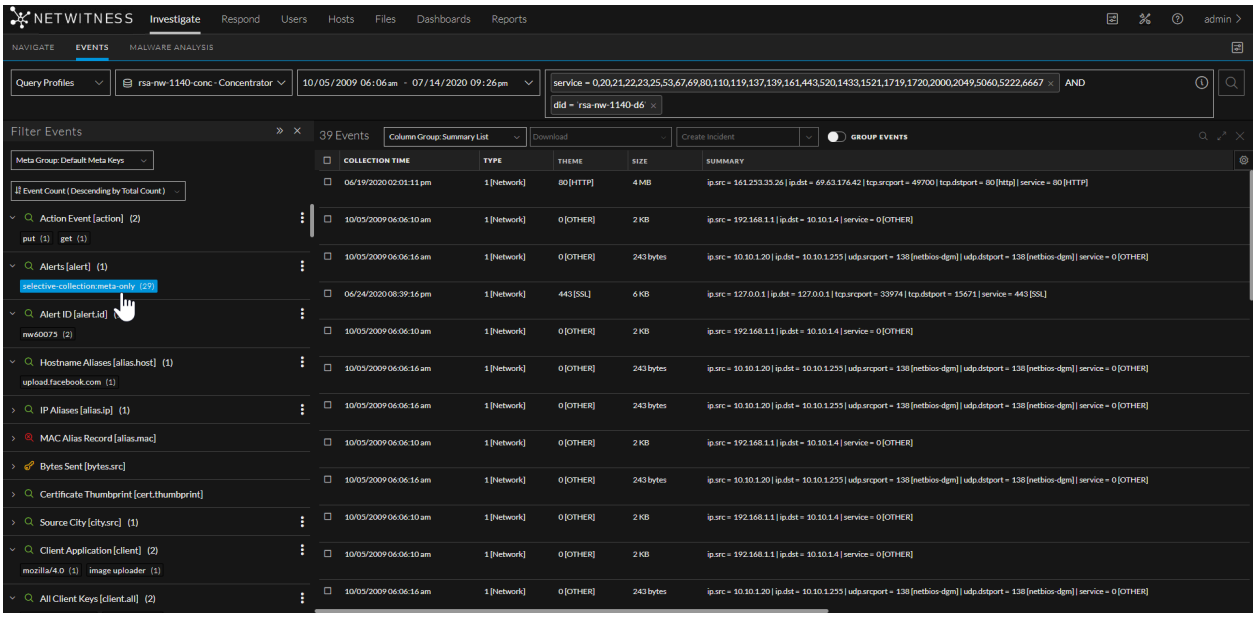


If you chose the action to `collect meta only`, you can click on one of the events and review the metadata, however, a message would be displayed instead of the packet payloads, as shown in the following figure.



To expedite this process, you can also take the application rules from one of the Network Decoders the policy is applied to and use it as the filter in the **Investigate** > **Events** query bar.

The following figure shows an entire `collect meta only` configuration by using the application rule shown here, created on the Decoder (in Verify Policies in Admin > Services):
`(service=0,20,21,22,23,25,53,67,69,80,110,119,137,139,161,443,520,1433,1521,1719,1720,2000,2049,5060,5222,6667) truncate`:



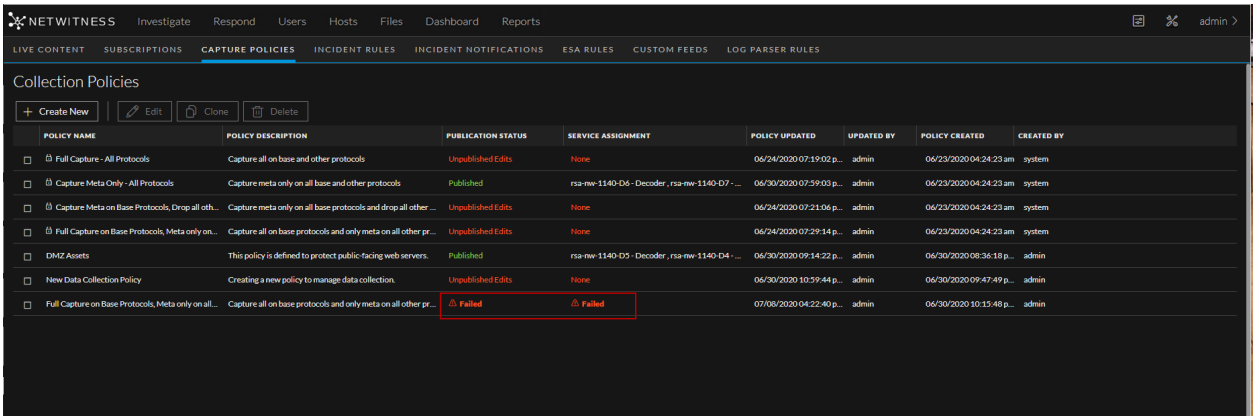
Note: If the findings do not align with what is expected from the policy, make sure that there are not other rules on the individual Decoders that might be conflicting with the selective collection policy rules. In addition, make sure when checking in Investigate that you are looking at only the traffic associated with the Decoders related to the policy.

Troubleshooting Policy Deployments That Fail

Policy deployment can fail for the following reasons:

- A Decoder service assigned to the policy is not running
- A Decoder host system is down when you try to deploy the policy

When a policy deployment fails, the Failed icon is displayed in the PUBLICATION STATUS and SERVICE ASSIGNMENT columns as shown in the following figure.



Use the following options to research publication failures:

Check the status of the Decoders

Check to see if the Decoder system is down, or if the Decoder service has stopped capturing data.

1. Select the policy, click **Edit**, and go to the Assign Policy page to determine which Decoders are assigned to the policy,
2. Go to  (Admin) > **Services**, select the Decoder, click  > **View** > **System** to investigate the status of the Decoder.

Review the content-server.log file

You can review the content in this file to troubleshoot errors. This file is located in the following directory on the NW Server host:

```
/var/log/netwitness/content-server/content-server.log
```

The following example shows a content-server log snippet that is displayed when publishing a policy fails:

```
2020-07-08 16:22:45,287 [ nioEventLoopGroup-2-2]
INFO     DataAccess|Finished publishing policy 'Full Capture on Base
Protocols, Meta only on all other protocols-1' outcome:FAILED
```

This is an example of a content server log snippet for successfully publishing a policy:

```
2020-06-23 17:40:50,373 [ nioEventLoopGroup-2-7]
INFO     DataAccess|Finished publishing policy 'Policy 1' outcome:SUCCESS
```

Review `/var/log/messages` on Decoder Systems

You can view publication status in `/var/log/messages` on a Decoder system.

The following is an example of successful publication in `/var/log/messages`:

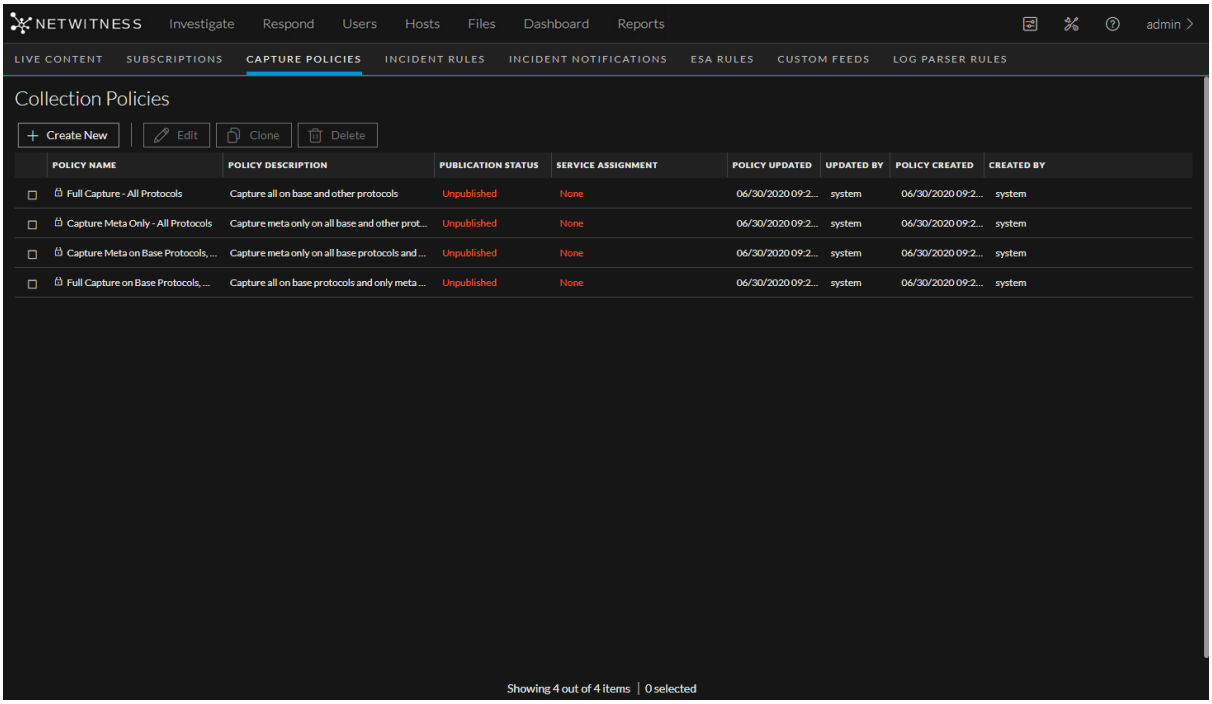
```
Jun 15 18:35:17 rsa-nw-1141-Decoder-1 NwDecoder[8260]: [Rules] [audit] User
admin (session 26182, 10.237.169.43:55380) has added application rule
'selective-collection:meta-only' at position 38
```

Unpublish Policies

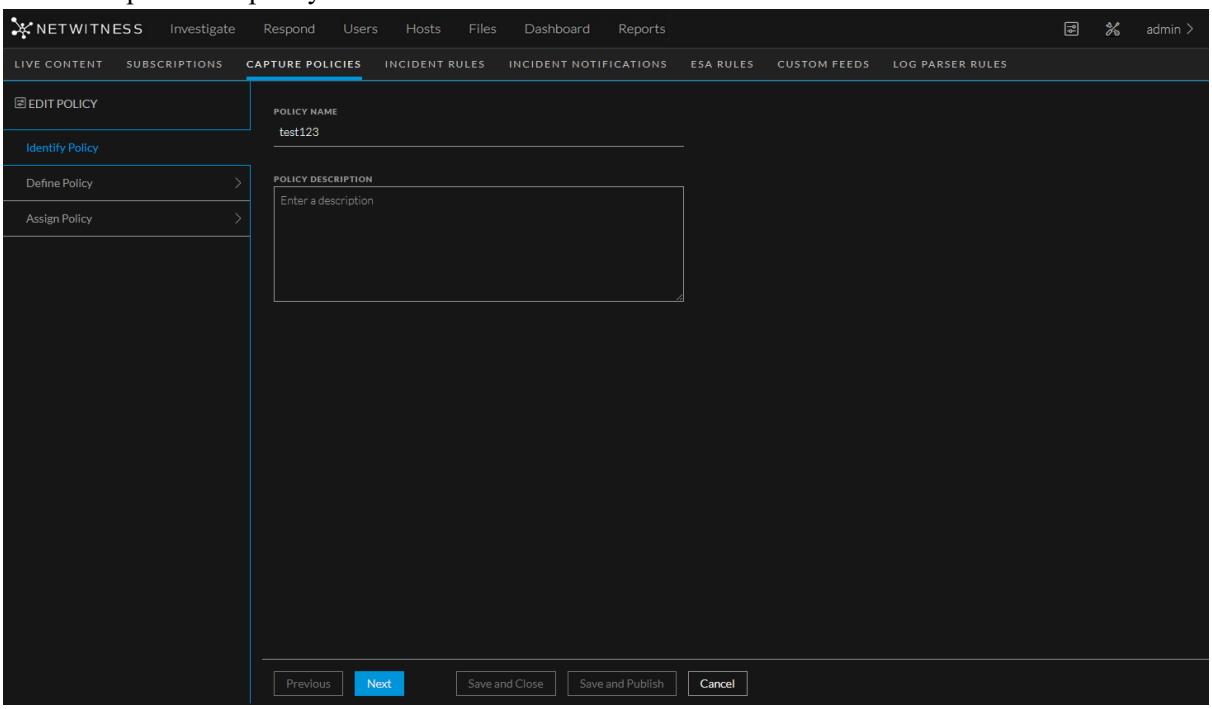
You can unpublish a policy by unassigning the Network Decoders from the policy, which preserves the policy but makes it inactive so that the data it defines is not being collected. You can then change or remove a portion of policy, or remove the entire policy from the Decoders to which they are assigned.

To unpublish a policy:

1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.

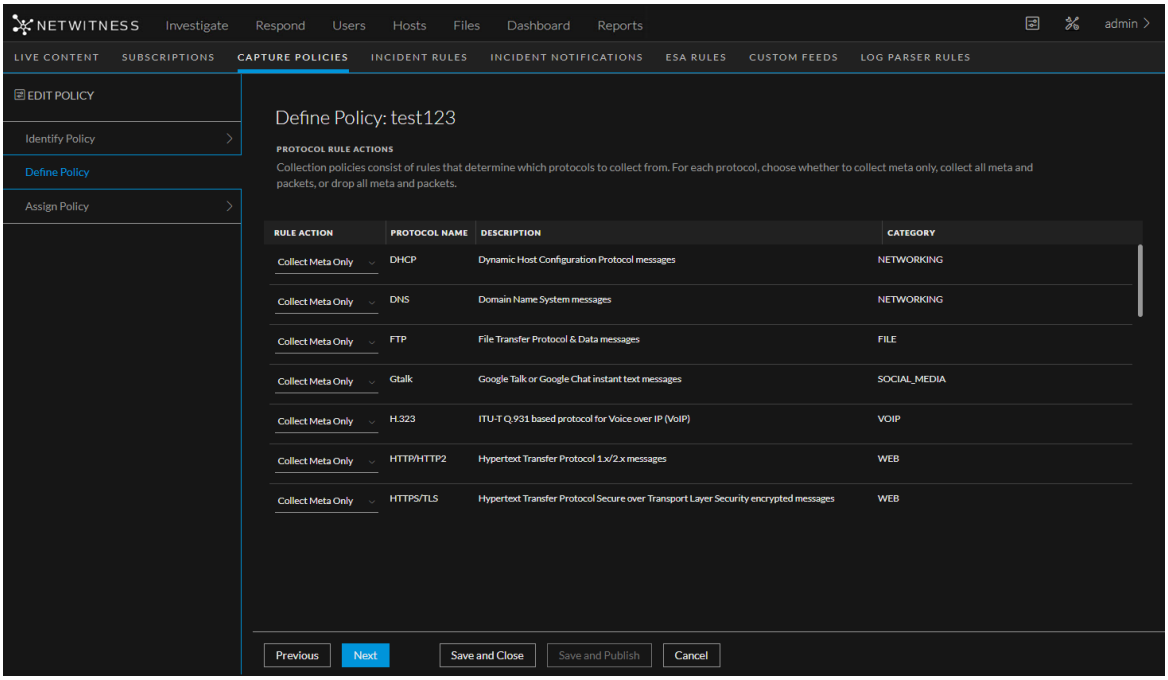


2. Select the published policy and click **Edit**.



   a. To change the name or description of the policy, update the policy name or description in the Identify Policy view.

b. To change the protocols that are collected, click **Next** and then edit the rule actions in the Define Policy view.



c. To unassign or reassign Decoders to this policy, click **Next**.

- To fully unpublish this policy, deselect all the Decoders and then click **Save and Publish**.
- To change the Decoders that are assigned to this policy, select or deselect the Decoders as needed and then click **Save and Publish.**

When the changes to the policy status have been completed and the Decoders updated, the policy status is displayed in green text as **Published** on the Collection Policies tab. If you unassigned all the Decoders without assigning any new ones, the publication status is **Unpublished Edits**.

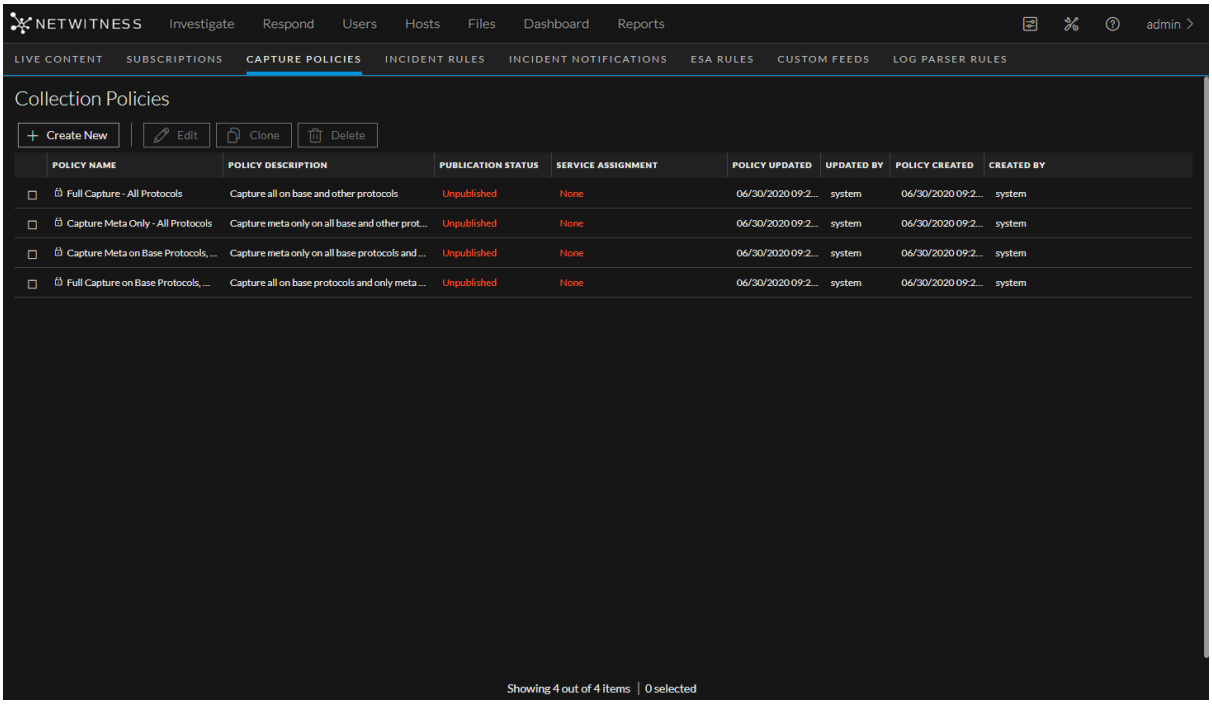
Delete Policies

You can only delete customized policies. Predefined policies can be unpublished, but cannot be deleted. You can only delete one policy at a time. When you delete a policy, any Decoders that were assigned to the policy immediately stop applying the rules in the policy.

For information about how to unpublish a policy, see Unpublish Policies.

To delete a policy:

1. In the NetWitness user interface, go to (Configure) and select the **CAPTURE POLICIES** tab.



2. Select a customized policy and click **Delete**.
   The policy is removed from the system.

Caution: If you delete a policy that is published to Decoders, a warning is displayed. If you click **Accept** in the warning, that policy is removed from the Decoders, the collection status is changed, and the policy is deleted.

Supported Protocols for Selective Network Data Collection

The protocols in this list are supported for selective network data collection. These protocols have descriptive tool tips in the user interface. You can choose how to handle each of these protocols individually by dropping, capturing metadata only, or capturing all data related to them. There is also an additional rule for each policy that handles all other unidentified protocols not included in this list. The policies apply the collect or drop actions based on protocol characteristics, regardless of the port that they are captured on.

Category	Protocol	Description
Database	TNS	Transparent Network Substrate Oracle database networking protocol messages
	TDS	Tabular Data Stream for SQL server communications
Email	SMTP	Simple Mail Transport Protocol messages

Category	Protocol	Description
	POP3	Post Office Protocol v3 messages (Email clients use to retrieve mail from server)
File	SMB	Server Message Block (or CIFS) network communication protocol file sharing communications
	NFS	Network File System protocol messages
	FTP	File Transfer Protocol messages
	TFTP	Trivial File Transfer Protocol messages
Networking	SNMP	Simple Network Management Protocol messages
	DNS	Domain Name System messages
	DHCP	Dynamic Host Configuration Protocol messages
	NETBIOS	Network Basic Input/Output System networking protocol messages
Remote Access	SSH	Secure Shell secure remote connections
	TELNET	Teletype network remote interface protocol messages
	SCCP	Cisco Skinny Client Control Protocol terminal control protocol messages
Routing	RIP	Routing Information Protocol distance-vector routing protocol messages
Social Media	GTalk	Google Talk or Google Chat instant text messages
	IRC	Internet Relay Chat protocol text messages
	NNTP	Network News Transfer Protocol for Usenet news articles
VoIP	RTP	Real-time Transport Protocol audio and video streaming protocol messages
	SIP	Session Initiation Protocol text-based messaging application control protocol messages
	H323	ITU-T Q.931 based protocol for Voice over IP (VoIP)
Web	HTTP/HTTP2	HyperText Transfer Protocol 1.x/2.x messages
	HTTPS/TLS	HyperText Transfer Protocol Secure over Transport Layer Security or Secure Sockets Layer encrypted messages

Note: When selective network data collection is set up to collect metadata only for the HTTPS/TLS protocol, it truncates the traffic after the SSL/TLS handshake to ensure that all the certificates and other details are retained for analysis purposes while the encrypted packet payloads are dropped.

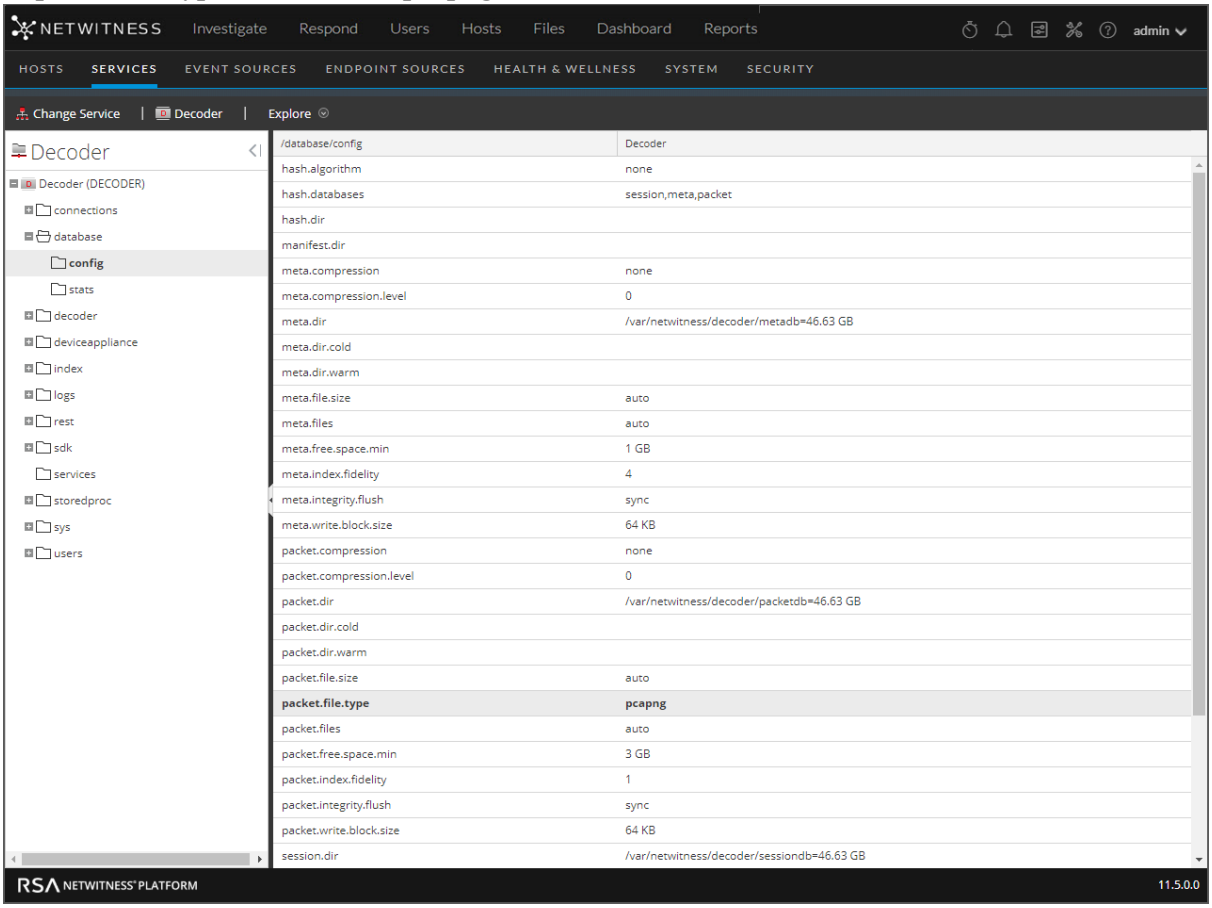
(Optional) Configure a Decoder to Write Standard pcap-formatted Files

To provide a more open database format, the Network Decoder can write standard pcap-formatted files. You can enable pcapng-formatted database files with the configuration node:
`/database/config/packet.file.type = 'pcapng' or 'netwitness'`

Note: This capability is enabled by default if you install 11.2 or later. If you upgrade from a previous version to 11.2, you must enable pcapng-formatted database files manually, which can result in an approximate 4% decrease in disk space (as the pcapng files require more space than the NetWitness nwdb files).

To enable writing standard pcap-formatted files:

1. Go to (Admin) > **Services**, select a Network Decoder service, and then select > **View > Explore**.
2. Go to **database** > **config**.
3. In **packet.file.type**, the default is **pcapng**.



4. To change the packet file type to NetWitness formatting, type **netwitness** and press **Enter**. This change will take effect immediately on the next packet file that is created.

Note: In the pcapng database file format, the data is in clear text, and is not obfuscated by our proprietary format, which can improve security.

Caution: Please do not touch any files in the packet database directories! You must not read or edit any `pcapng` file in the packet database directories, as they are always in use while Decoder is running. Decoder always expects full and exclusive access to those files, and other processes reading those files prevent normal Decoder operation. The proper way to access the `pcapng` files is to set up a cold storage directory. This allows Decoder to copy `pcapng` files to the cold storage directory before deletion. At that point, you are responsible for managing the `pcapng` files, including making sure that the cold storage volume never fills up. Keep in mind that copying the `pcapng` files to cold storage requires a non-trivial amount of I/O and could interfere with packet capture. Cold storage for `pcapng` is not supported at 10G speeds.

(Optional) Multiple Adapter Packet Capture

Beginning with 11.5, the Network Decoder can capture from multiple interfaces simultaneously. This functionality allows Network Decoders to capture from multiple physical Network Interface Cards (NICs), or multiple ports on multiple interfaces, while leveraging the same network rules, application rules, and parsers for each NIC. The benefit of capturing from multiple physical NICs is that this method multi-threads the capture process for each adapter interface, which also multi-threads network rule evaluation.

For example, prior to 11.5, defining `capture.selected=PFRINGZC, em3` along with `capture.device.params=device=zc:em3,zc:em4` allocated a single thread for the Network Decoder to capture any traffic that came into either `em3` or `em4` interfaces. This also meant that all network rules were evaluated in a single thread.

However, in 11.5, configuring `capture.selected=PFRINGZC,em3; PFRINGZC,em4` allocates two threads for the Network Decoder to capture the traffic collected on either `em3` or `em4`, which provides more resources to the capture pipeline.

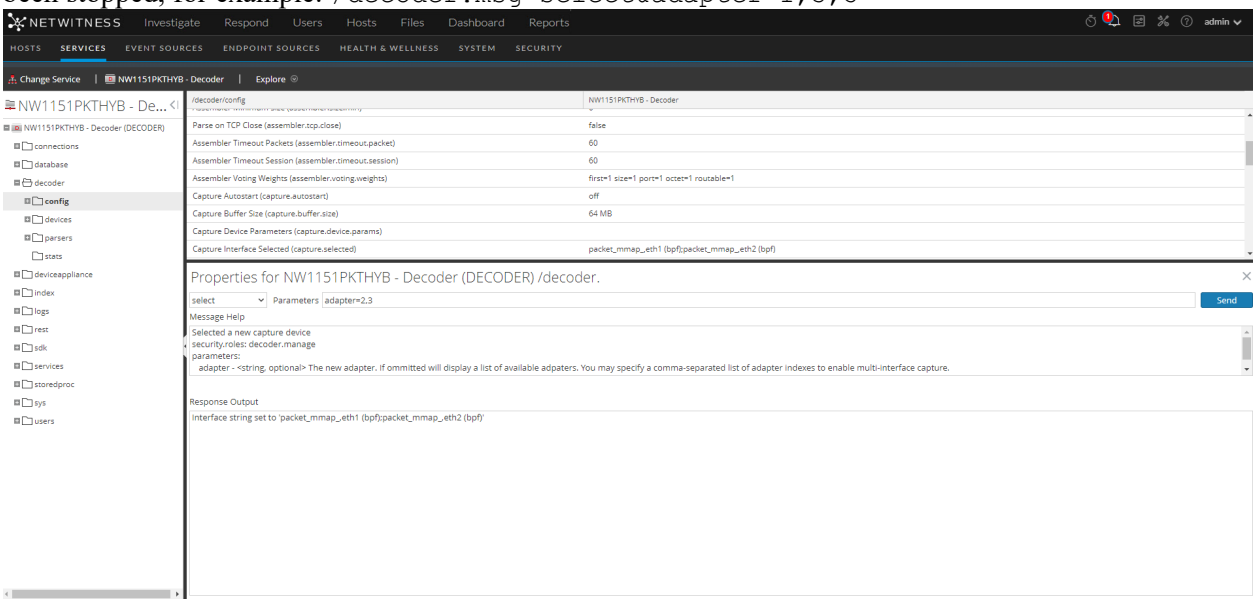
In this topic, the term "adapter" refers to a capture device within the Decoder. An adapter's name consists of the device (for example, the software API that the Decoder uses to interact with an adapter), and an interface (the physical device). A device can have multiple interfaces, and an interface can be compatible with multiple devices.

Configure Multiple Adapter Packet Capture

Note: Changes to capture device configuration do not take effect until service restart.

You can define multiple adapter packet capture for Decoders in the config setting `/decoder/config/capture.selected`. It accepts a semi-colon separated list of adapters, for example: `bpf,en5 (bpf);null_device,Null Capture`. This setting can be edited while capture is running, however, the changes do not take effect until the service has been restarted.

You can use the `/decoder?msg=select` command to select one or more adapters after capture has been stopped, for example: `/decoder?msg=select&adapter=1,3,5`



Note: As of 11.5, this configuration is not yet available in the `decoder` > `config` page. Use the `decoder` > `explore` page or the REST API. You can the Explore view by selecting a Decoder in the user interface and going to  (actions)> **View** > **Explore**. You can access the REST API by opening a browser and specifying the IP address of the host, for example, `https://<decoder-ip-address>:50106`.

Per Interface Configuration

The `/device/config/device.capture.params` configuration is a global configuration setting, and is applied to every interface when capture is started.

There is also a configuration setting for individual interfaces, which is `/decoder/devices/<devicename>/<interfacename>/config`. Currently, the only option available in this configuration setting is `capture.params`. Any value in this setting will override the global configuration setting.

Stat Nodes

As with configuration, there are per-interface stats and aggregate stats (similar to the global configuration). When there is a single device configured, these stat nodes function as in previous versions. When there is more than one device configured, the stat settings `/decoder/stats/capture.device` and `/decoder/stats/capture.selected` read `multi`.

Per-Interface Stats

The stats for each interface can be found in:
`/decoder/devices/<devicename>/<interfacename>/stats`

The captured stats are:

- capture.avg.size
- capture.dropped
- capture.filtered
- capture.header.bytes
- capture.kept
- capture.payload.bytes
- capture.received
- capture.total.bytes
- pool.packet.captured

Aggregate Stats

The capture stats under /decoder/stats represent an aggregate of the individual interface stats.

In the following table, for example, capture.avg.size is the maximum average size among all the interfaces, where capture.dropped is the total captures that were dropped for all the interfaces.

capture.rate	SUM
capture.avg.size	MAX
capture.dropped	SUM
capture.filtered	SUM
capture.header.bytes	SUM
capture.kept	SUM
capture.payload.bytes	SUM
capture.received	SUM
capture.total.bytes	SUM
pool.packet.captured	MIN

Pool Stats

All the network packet and session pool pages are tracked in /decoder/stats/pool/.

The stats under /decoder/stats/pool represents the total number of capture pages configured in Decoder. These include packet pool and session pool pages.

capture.port Meta Key

Each session is given a capture.port meta key which identifies the adapter from which the packets in the session were captured. This meta key can optionally be turned off with the /decoder/config/captureport.meta.enabled option. This meta key is on by default for Network Decoders, and is off by default for Log Decoders.

IMPORTANT: The `capture.port` meta key does not work with the `packet_mmap_,ALL` adapter. You must specify individual ports while using `capture.port`. The `packet_mmap_,ALL` adapter is capable of capturing across all types of network interfaces at the same time. For more information, see (Optional) Configure a Decoder to Capture Data Across All Types of Network Interfaces.

Note: If you turn the `capture.port` meta key off, you must perform a capture restart for the change to take affect since this is a configuration change.

Special Devices

The nwimport device is for handling PCAP imports while simultaneously capturing network data. For more information on PCAP imports, see Upload a Packet Capture File.

`nwimport`: This device is not selectable but does have stat and config nodes. The config node for `nwimport` is ignored.

`Packet arrived out of order` Message

This message means that the Decoder is seeing old packets. This can be a side-effect of running import while capture is running. This message is usually suppressed at the beginning of import and then unsuppressed at the end. However, packets can still be in the state of being processed even after the upload of the imported file has completed. This means that this warning can be unsuppressed before the processing of the imported packets is completed, which results in displaying this message.

For information about updating Decoder configurations, see the *RESTful API User Guide for RSA NetWitness Platform*. For information about stopping and restarting capture, see Configure Capture Settings. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Multi-Assembler Modes

The process that reassembles captured packets into streams is known as the Assembler. You can customize the Assembler operation to suit the packet capture configuration. It works in two modes (`on` and `off`) that are configured using the `/decoder/config/assembler.threading.enabled` variable. The value of `assembler.threading.enabled` only takes effect if multi-adapter packet capture is enabled.

Note: If you change the value of `assembler.threading.enabled`, you must perform a capture restart.

OFF Mode: If assembler threading is `off`, then one Assembler is used to assemble all packets captured from all interfaces. Packets from a stream may appear on any capture interface and the single assembler will put them back together.

- Assembler threading must be `off` to allow assembler to reconstruct streams that are capture across multiple capture interfaces.
- When assembler threading is `off`, the single Assembler instance may limit the maximum packet rate that the Decoder can process.

ON Mode: If assembler threading is `on`, there will be a separate Assembler for each capture interface. Each Assembler will only reassemble the packets it receives from it's capture interface.

- This mode of operation enables higher throughput because each Assembler instance operates on a dedicated processor.
- Each logical stream of packets must be ingested by one interface only. For example, the streams on each capture interface could come from different networks.

For compatibility with previous releases, the Multi Assembler mode is `off` by default.

When importing packets, the imported packets are fed to one of the existing assemblers, or the single Assembler instance when Assembler threading is `off`. Importing packets does not create a separate Assembler process.

(Optional) Internet Content Adaptation Protocol Capture

Internet Content Adaptation Protocol (ICAP) is a service protocol that encapsulates HTTP messages into ICAP Messages and forwards them to an ICAP server for processing.

NwDecoder supports capturing ICAP Messages and converting the HTTP requests and responses into packets. The capture device is named `icap`, `ICAP Server`, and can be used simultaneously with other capture devices (i.e., packet capture).

ICAP Capture Options

The capture options are specified on Decoder in the configuration node `/decoder/config/capture.params`

You must restart the capture before any changes take effect.

- `reqmod=<bool>` - By default, the Decoder will process both `REQMOD` and `RESPMOD` messages. Since the `RESPMOD` contains both the request and response, it may be advantageous for Decoder to ignore `REQMOD` messages and only process `RESPMOD`. If you set the `reqmod=false`, the Decoder will generate a request packet and response packets from a single `RESPMOD` message. The default behavior is equivalent to `reqmod=true`

- `client_ip=<string>` - You can configure some ICAP Clients to include the originating IP address in the ICAP headers. The name of the header with this information can vary, so you can use this setting to specify the name of the client IP header. For example, `client_ip=X-Client-IP` would configure the Decoder to extract the IP from the `X-Client-IP` header. Ignoring this setting will disable searching for the header.

(Optional) Data Plane Development Kit Packet Capture

The NetWitness Decoder supports capturing packets through the Data Plane Development Kit, also known as DPDK. DPDK is a set of libraries and drivers that support applications that work on directly on packets. DPDK's web site, www.dpdk.org, provides more information.

Advantages

The Decoder can use the DPDK to ingest packets from Ethernet interfaces. It is an alternative to existing packet capture options built into the Decoder. DPDK offers these advantages:

- It is entirely open source, with no binary-only driver components, unlike PF_RING.
- It is directly supported by multiple vendors, including our main vendor, Intel.
- Since the kernel-level interface is in the mainline Linux kernel, there are no separate driver packages to compile and distribute.
- DPDK devices are completely separate from the Linux network stack, which avoids unintentional interaction between Linux networking and packet capture.
- DPDK is capable of capturing at greater than 10G, unlike the `packet_mmap`.

How it Works

DPDK completely detaches the Ethernet interface from the Linux kernel. It does this by doing the following:

- It uses a Linux User I/O driver mechanism to expose the raw PCI-Express device to normal Linux programs. This means that it can take the raw memory space mapped to the PCI device and expose it, so that a normal Linux program can directly manipulate the hardware. There are a couple of drivers built into the Linux mainline kernel that can accomplish this. Decoder is designed to use `vfio_pci_generic`.

  Note: Only one driver can be attached to an Ethernet interface at a time. When `vfio_pci_generic` is attached to an interface, the Linux network driver for that interface is detached.

- The DPDK provides its own user-space implementation of the hardware drivers for a wide variety of Ethernet cards. When a DPDK application like Decoder is loaded, it can attach directly to the PCI device and manipulate it using DPDK-specific library functions. When a device is attached to DPDK, the Decoder can see it and automatically sets the device up with the DPDK Capture Device option.

Migrate PF_RING Devices to DPDK

For configurations using multiple adapter packet capture (available in 11.5 and later), you must use the manual procedure outlined in Manually Move a Capture Interface to DPDK . For example, the migration (`dpdk migrate`) does not support a Decoder when capture is configured with: `capture.interface=PFRINGZC,em3; PFRINGZC,em4;packet_mmap,em2`.

For configurations prior to 11.5, the procedure that uses the devices parameter to capture from multiple `PF_RING`-supported interfaces (for example, `ixgbe` and `i40e`) will be converted to using DPDK, along with the multiple adapter packet capture, so that one capture thread is used per interface.

For example, the migration will work if configured as follows:

- For multi-interface capture: `capture.interface=PFRINGZC,em3` along with `capture.device.params=device=zc:em3,zc:em4`
- For single interface capture: `capture.interface=PFRINGZC,em3`

Note: `dpdk migrate` only functions with `PFRINGZC`, not `mmap` interfaces.

If you are using `PF_RING` to capture packets on your Decoder, you can use the Decoder's `dpdk migrate` command to migrate your existing PF_RING capture devices to DPDK. This command uses your currently-selected capture device configuration to determine which network interfaces are moved to DPDK control.

1. Go to  (Admin) > Services and select a Decoder.
2. Click  > **View** > **Explore**, right-click **decoder** and select **Properties**.
3. From the command drop-down list, select the `dpdk` command, type the parameter `migrate`, and click **Send**.

4. In the command output window, the changes that will be made on the Decoder to perform the migration are displayed. If everything looks correct for the migration, add the parameter `commit=1` to the command after `migrate` to actually make the changes on the Decoder.



1. At the reboot prompt, reboot the host to make the driver loading changes.
2. (Optional) Run the `/decoder/devices/<devicename>/msg=prune` command to remove configuration options for devices and interfaces that no longer exist after the migration.

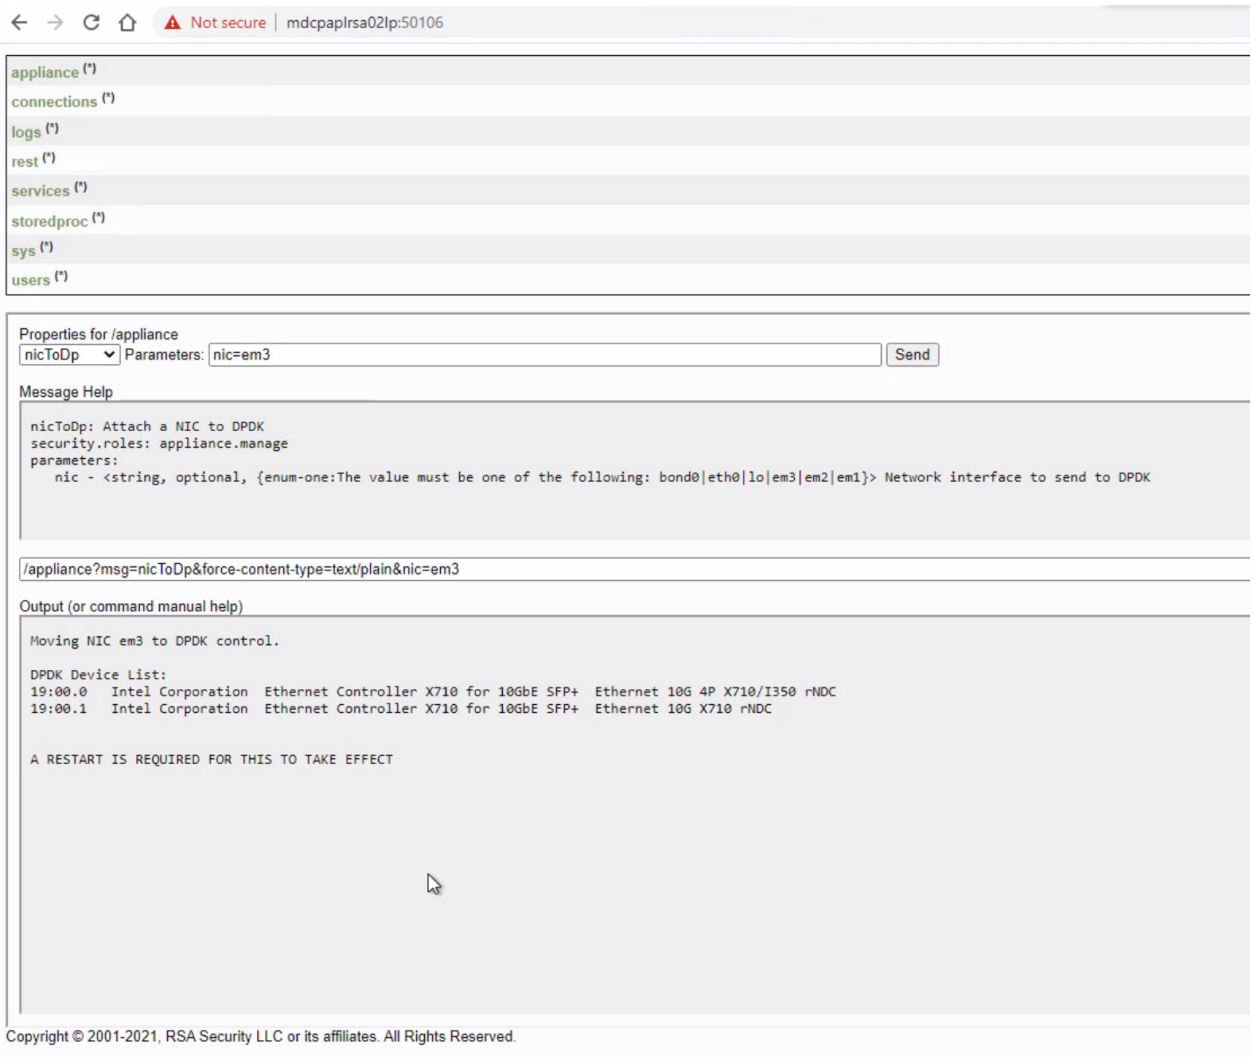
> **Note:** The `prune` command removes only the undetected interface folders. Also, it ignores the `interfaces`, `stats`, `nwimport`, and `flow_events` folders available in `/decoder/devices/`.

Manually Move a Capture Interface to DPDK

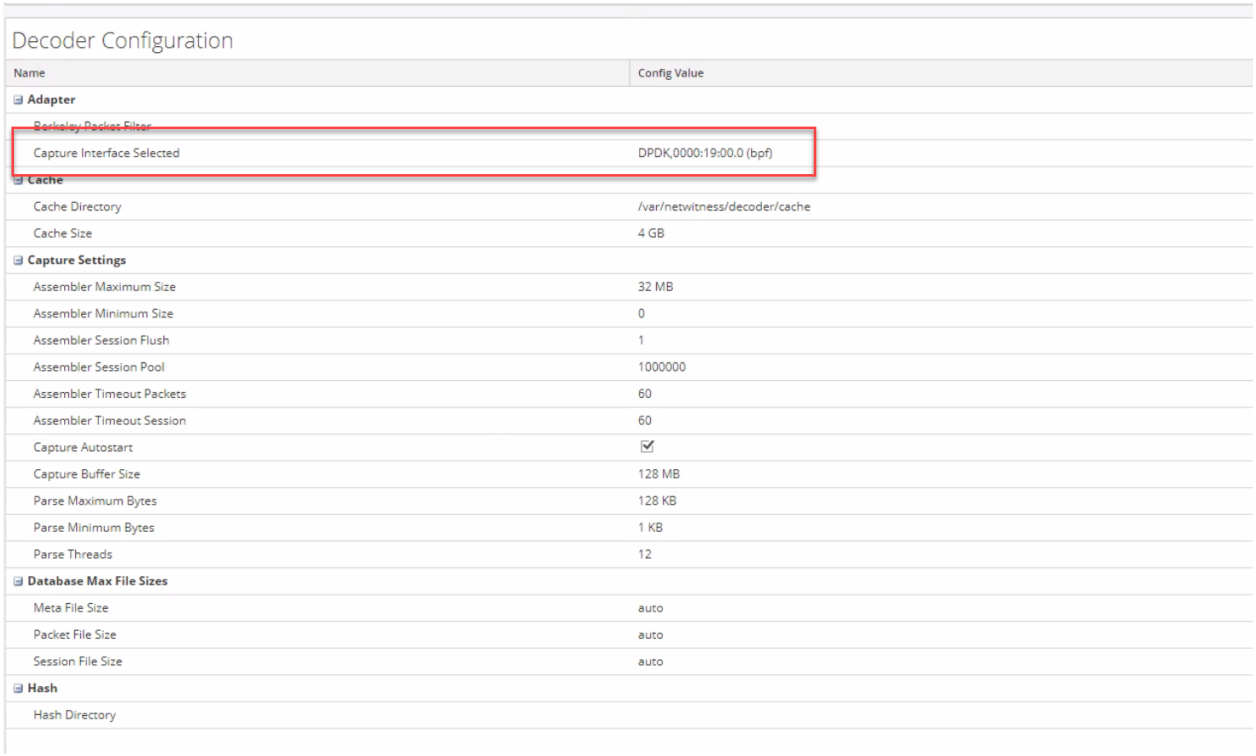
1. Run the new host service command `/appliance/nicToDp`, that converts a named network interface and sends it to DPDK.
   For example, if you provide `nicToDp` an interface name like `p1p1` it detaches it and sets it up for DPDK to use. It does the following:
   a. Identifies the PCI address of the named interface.

   > **Note:** From this point on, the Ethernet interface is identified by its PCI address alone, which is a physical characteristic based on where the Ethernet interface is installed in the computer.

   b. Creates configuration files that describe which PCI addresses DPDK is allowed to use.
   c. Sets up a systemd unit that binds the configured PCI addresses to DPDK instead of to the default Linux drivers on the next system reboot. The following figure shows a successful command.



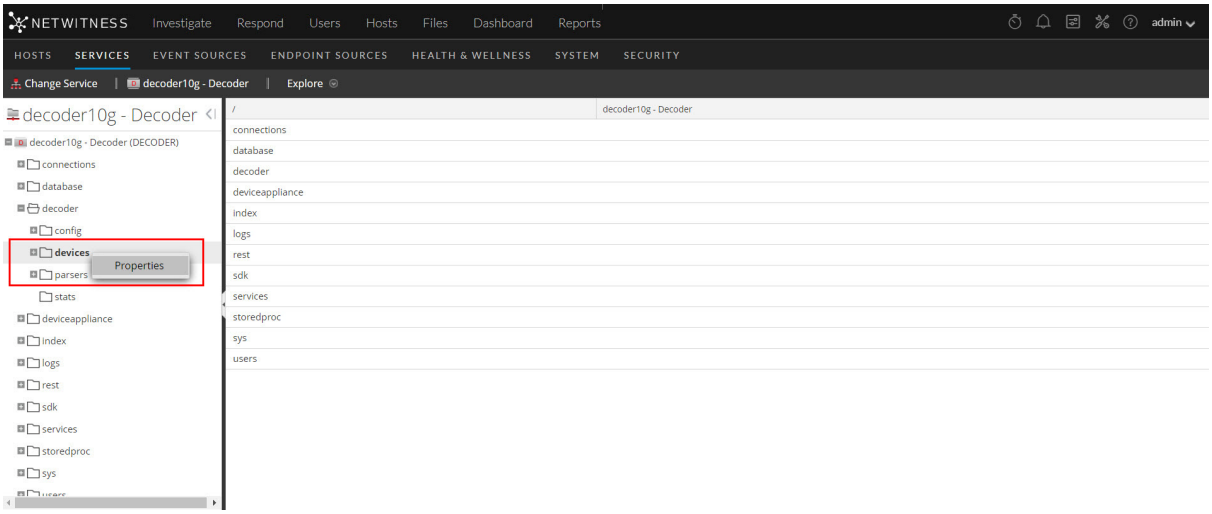
2. Reboot the Decoder host. The reboot accomplishes the following:
   a. Ensures that live capture on the interfaces stops before the network driver is unloaded.
   b. Binds the `vfio_pci_generic` driver to the configured interfaces.
   c. Pre-allocates some Huge Page memory for each interface, as required by the DPDK libraries.
3. After the reboot, the Decoder does not start capturing, since the network interface it had been using to capture has been moved to DPDK. Change the Decoder's capture configuration to select the DPDK device that is now available. For information about how to change the capture configuration, see Configure Capture Settings.
4. Restart the service. The following figure shows the Decoder configuration after a restart.



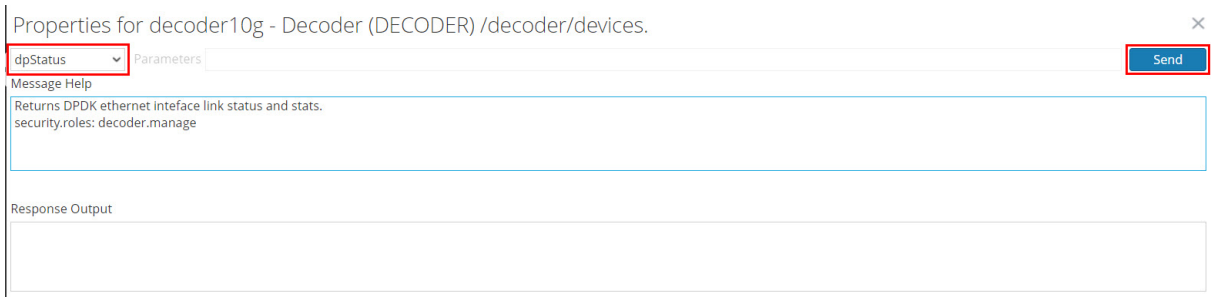
Check the device status and stats of DPDK interfaces

When you capture using DPDK interfaces, the link status and stats of the DPDK can be obtained using the command `/decoder/devices dpStatus`. The traditional commands like `ethtool` or `ifconfig`, which are used for other capture interfaces, will not fetch the link status and stats of DPDK devices.

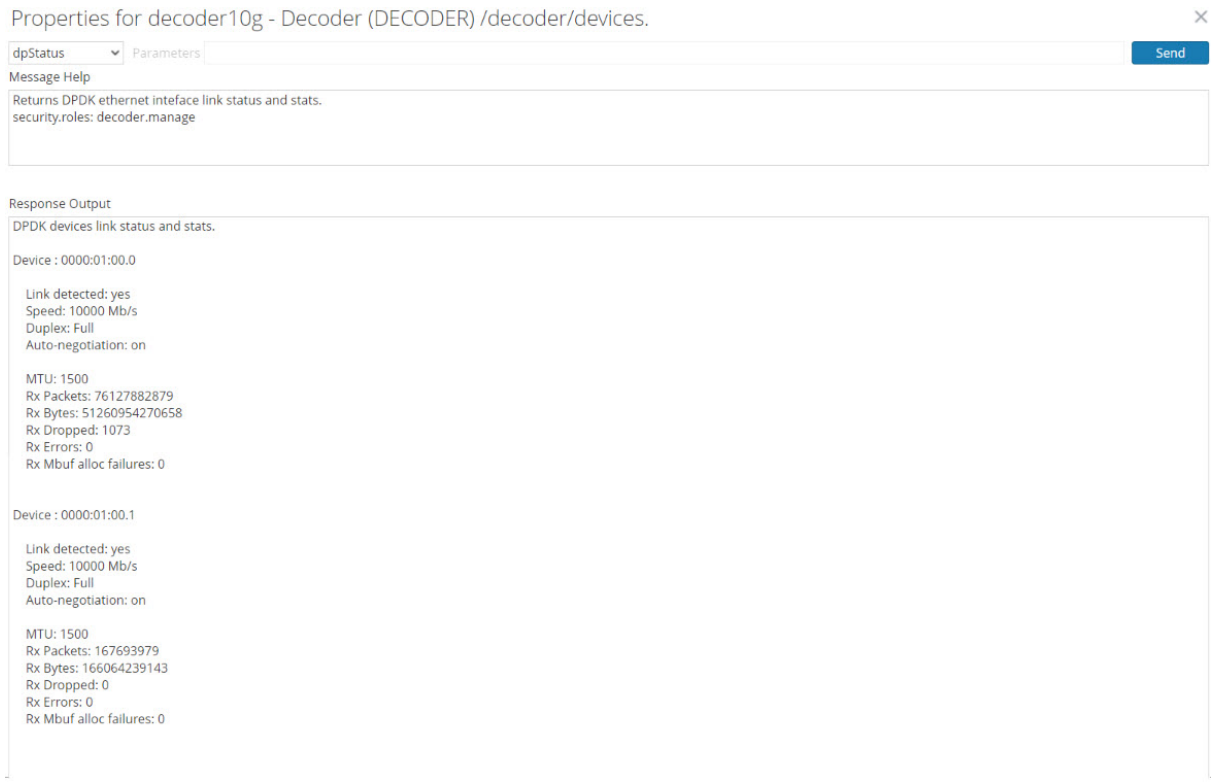
1. Go to (Admin) > **Services**.
2. In the Administration Services view, select the Decoder and **View > Explore**.
3. Navigate to **decoder > devices**, right click on **devices** and select **Properties**.



4. Select **dpStatus** from the drop-down and click **Send**.



The link status and stats of DPDK devices are displayed in the **Response Output** section.



Check link status and stats from NwConsole

Run `/decoder/devices dpStatus` on NwConsole. Sample output when capture is in **Start** state:

```
[localhost:56004] /> send /decoder/devices dpStatus
DPDK devices link status and stats.

Device : 0000:01:00.0

    Link detected: yes
    Speed: 10000 Mb/s
    Duplex: Full
    Auto-negotiation: on

    MTU: 1500
    Rx Packets: 416843663
    Rx Bytes: 281110914127
    Rx Dropped: 314
    Rx Errors: 0
    Rx Mbuf alloc failures: 0


Device : 0000:01:00.1

    Link detected: yes
    Speed: 10000 Mb/s
    Duplex: Full
    Auto-negotiation: on

    MTU: 1500
    Rx Packets: 597764
    Rx Bytes: 373053232
    Rx Dropped: 0
    Rx Errors: 0
    Rx Mbuf alloc failures: 0
```

Sample ouput when the capture is in **Stop** state:

```
Device : 0000:01:00.0
Link detected: no
Device : 0000:01:00.1
Link detected: no
```

Utilizing Receive Side Scaling with DPDK

Receive Side Scaling (RSS) efficiently distributes the network receive processing across multiple CPUs in a multi-processor system. RSS ensures that the processing associated with a given connection stays on the assigned CPU.

IMPORTANT: Receive Side Scaling is supported on DPDK devices using the ixgbe or i40e device drivers only.

Advantages

When RSS is enabled, the DPDK capture devices split the incoming packet stream into multiple queues providing following benefits:

- Each queue is processed with a dedicated capture worker thread. As there are more threads processing the capture work, it allows more CPU time available for Berkeley Packet Filters (BPF) rule and network rule evaluation.

- Each capture worker thread can feed multiple assemble worker threads simultaneously which allows more CPU time available for assembly.
- Each CPU thread processes only a fraction of the packet stream. However, the total number of packets processed from all queues working in parallel increases the overall output (packets processed per second) of Decoder.

Configuring Receive Side Scaling

To enable RSS in a DPDK controlled NIC, set the `/decoder/devices/DPDK/(PCI_ID)/config/rss` value to a value greater than `1`. The typical values of RSS must be a small power of 2, for example, `4` or `8`. This value determines the number of CPU threads to be allotted for packet processing during the next capture. Similarly, if assembler threading is enabled, a matching number of assembler threads will be created to process the packets on each incoming queue.

For example, the following sample configuration will utilize a total of eight threads (four capture threads + four assembler threads) for multi-adapter packet capture with DPDK and multi-assembler threading turned on. This is the maximum recommended value for an eight core Decoder CPU.

`assembler.threading.enabled` set to `on` or `true`

```
/decoder/devices/DPDK/0000:01:00.1/config/rss=2
/decoder/devices/DPDK/0000:01:00.0/config/rss=2
```

To minimize packet drops, it is recommended that you keep the total number of capture and assembler threads less than the number of cores per CPU.

How it Works

When RSS is enabled, Decoder utilizes the RSS hardware acceleration on the capture NIC to apply a symmetric RSS hash function. The hash function sends packets from both directions of a given stream to the same queue. It enables the packets received on that queue to be processed by a dedicated assembler thread. This eliminates the need for Decoder to merge streams from different RSS queues and reassemble the streams. Hence, it is recommended that you enable multi-threaded assembly when using RSS. To enable, set `/decoder/config/assembler.threading.enabled` to `true` or `on`.

Configure the Decoder Capture Failover in Load Balance Deployments

The DPDK capture in the decoder service provides the ability for load balancers to failover and avoid data loss when decoder capture is off or when the service is shut down due to maintenance activity. When the decoder capture is stopped, the DPDK capture interface link status is down. This provides the ability for the traffic ingestion physical load balancers (for example, Keysight Ixia, Gigamon) which monitor interface links to switch to a different decoder and avoid data loss.

Configuration Failover on DPDK Capture Interfaces

Refer to the (Optional) Data Plane Development Kit Packet Capture *topic for instructions to configure DPDK capture on the decoder*.

- You can migrate existing PF_RING devices to DPDK

  (Or)

- Manually move capture interfaces to DPDK

Disable Wake On LAN (WOL) in BIOS settings

Some load balancers monitor interface link status through NIC link lights. You must disable WOL as the interface will not be down if WOL is enabled. To disable WOL in the BIOS, follow the BIOS instructions.

Example: f2 -> Device Settings -> "Integrated NIC 1" (capture interfaces) -> "NIC Configuration" -> "Wake on LAN" = Disabled

Note: When you restart the device and if capture does not auto-start, the lights on the capture interface are on. But if you turn on capture and turn it off again the lights will be off.

(Optional) Remove DPDK on Capture Interfaces

In case you need to revert DPDK on capture interfaces you can use the Appliance service `dpToNic` command.

- If you run this command without arguments, it lists the ethernet devices controlled by DPDK.
  - **Example:** /appliance dpToNic

    DPDK Device List:

    82:00.1 Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection Ethernet Server Adapter X520-2
- If you specify an ethernet interface using the 'eth' parameter, it removes the interface from DPDK so that when the system reboots it will be treated as a normal Linux network interface.
  - Example: /appliance dpToNic eth=82:00.1

    Detaching PCI Device 82:00.1 from DPDK.

    DPDK Device List: A RESTART IS REQUIRED FOR THIS TO TAKE EFFECT

Validate Link Status

- On Capture Start
  - During the decoder capture start, the service logs DPDK interface link status
    - [DPDK] eth dev 0000:05:00.0 link is up
  - On the upstream traffic ingest Load Balancer, the link status would be detected.
    - If it is a Linux-based device, you can use the ethtool <interface name> command to check link status
- On Capture Stop
  - During the decoder capture stop, the service logs that capture is stopped on DPDK interfaces
    - [Thread] Stopped thread: Capture Work (DPDK, 0000:05:00.0) id: 31214
  - On the upstream traffic ingest Load Balancer, the link status would not be detected.
    - If it is a Linux-based device, you can use the ethtool <interface name> command to check link status.

DPDK Interface Configuration and Statistics

- The `dpdk-devbind --status` command displays the status of DPDK and other known network interfaces. It displays the PCI domain, bus, slot, and function, along with a text description of the device.
  - Example: `dpdk-devbind --status`
    - 0000:01:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' drv=vfio-pci unused=ixgbe_zc
    - 0000:01:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' drv=vfio-pci unused=ixgbe_zc
- The default MTU is 2048. If you need to adjust the MTU for capture, add the `snaplen` parameter to `/decoder/config/capture.device.params` .
- VLAN offload is disabled.
- The DPDK interface capture stats are updated under /decoder/devices/dpdk/
  - Example: /decoder/devices/DPDK/0000:01:00.1/stats/
    - Received packets (Rx packets) - capture.received
    - Received bytes (Rx bytes) - capture.total.bytes
    - Received dropped (Rx dropped) - capture.dropped

For information on checking device status and stats of DPDK interfaces, refer to Check the device status and stats of DPDK interfaces

(Optional) Preserve VLAN Tags When Using the Packet MMAP Capture Interface

When capturing traffic containing VLAN tags, you may need to configure the Packet MMAP capture interface to preserve the VLAN tags in the packets (VLAN fixup). By default, the network capture hardware removes the tags. Performing this procedure preserves the tags in the packets, and the tag values are parsed into VLAN meta data for further analysis.

There are two mechanisms for enabling the VLAN fixup.

- Option 1: Set `vlan-fix=true` within `capture.device.params`. This option performs the VLAN fixup on all traffic entering the Decoder. This option is appropriate in most cases, since it is assumed that all the traffic will be VLAN tagged.This mechanism works on either single-interface mode, or on all-interfaces mode. This option overrides the VLAN fixup settings on individual interfaces; even interfaces that are not configured to do VLAN fixup will have the feature enabled.
- Option 2: Use the `interfaces` parameter within `capture.device.params` on a per-device basis. The `interfaces` parameter accepts a comma-separated list of interface names on which to capture packets. By adding `:vlan` to an interface name, you can enable the VLAN fixup on individual

interfaces. If the interface does not have the `:vlan` suffix added, then it will not perform the VLAN fixup.

After editing this parameter, you must restart capture on the Decoder in order for changes to `capture.device.params` to take effect.

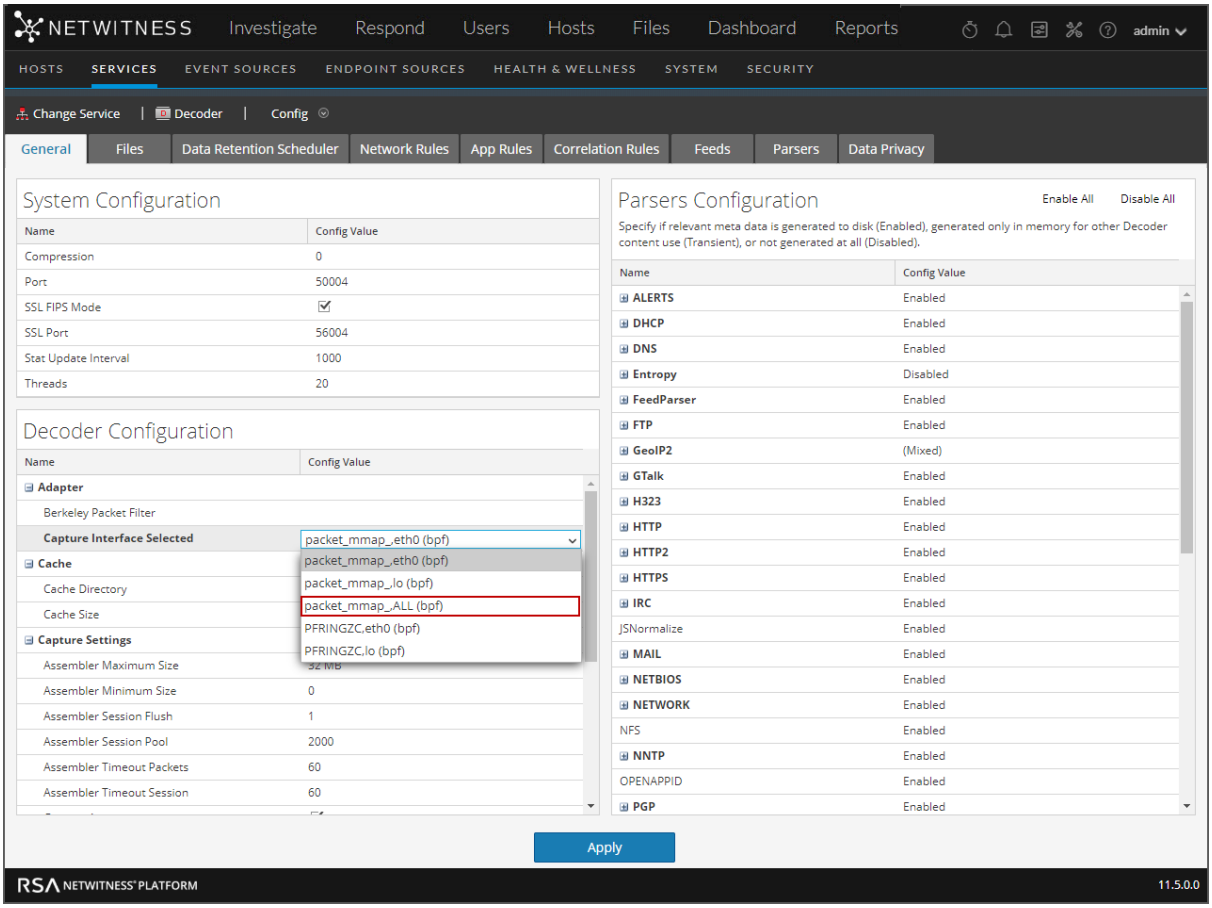
These are `vlan` examples of both options. If you need to pass multiple settings for `capture.device.params`, use the following synax. Notice that quotes are needed to delineate whitespace. For more information about how to use quotes for whitespace, see the "Connecting to a Service" topic in the *NwConsole User Guide for RSA NetWitness Platform*.
`name1="value1" name2="value2"`.

Parameter	Value	Effect
`capture.device.params`	`vlan-fix=true`	VLAN fixup always performed on all interfaces. The default value is `vlan-fix=false`.
`capture.device.params`	`interfaces=eth0:vlan,eth1`	VLAN fixup performed on traffic capture on `eth0` interface only
`capture.device.params`	`interfaces=eth0:vlan,eth1 vlan-fix=true`	VLAN fixup always performed because the `vlan-fix` setting overrides the interfaces setting.

To configure the `packet_mmap_` adapter to preserve the VLAN tags in packets:

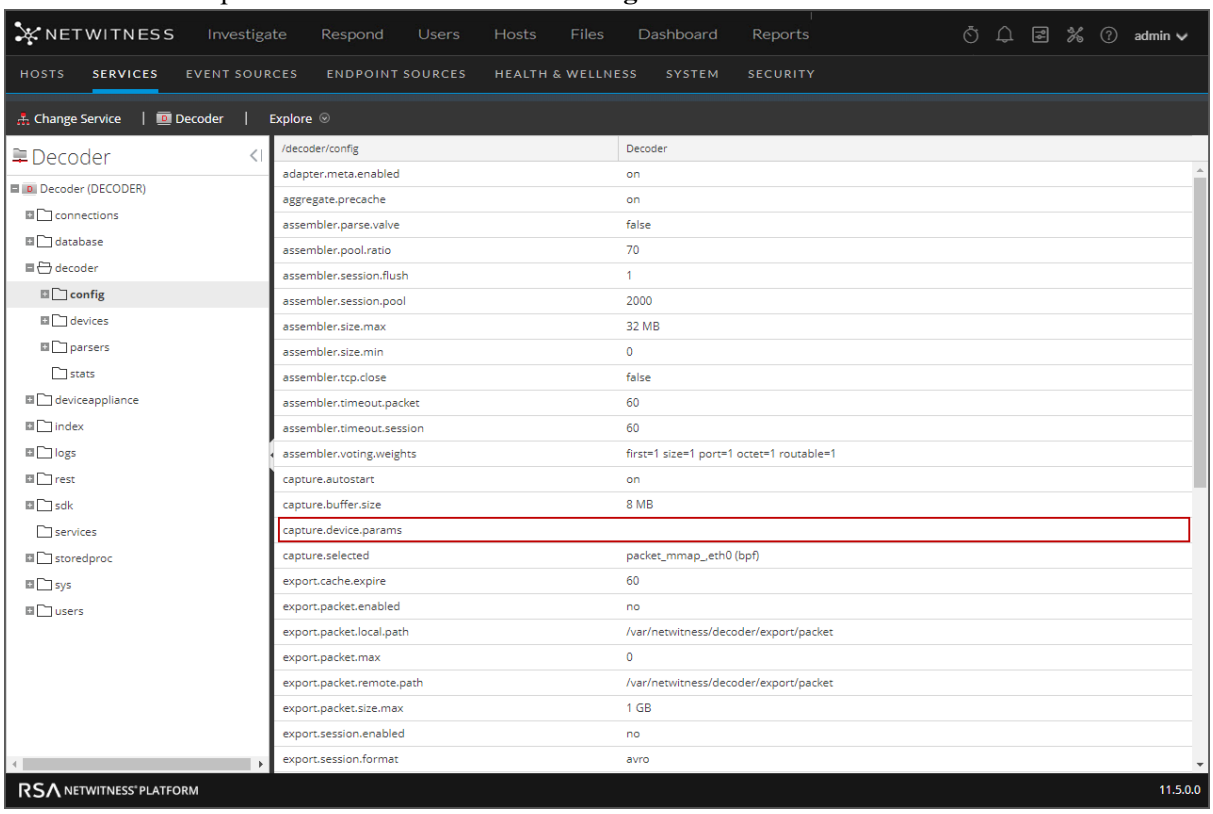
1. Go to  (Admin) , select the Decoder service and  **> View > Config.**

2. In the **Decoder Configuration** panel, set **Capture Interface Selected** to the `packet_mmap_,ALL` adapter.



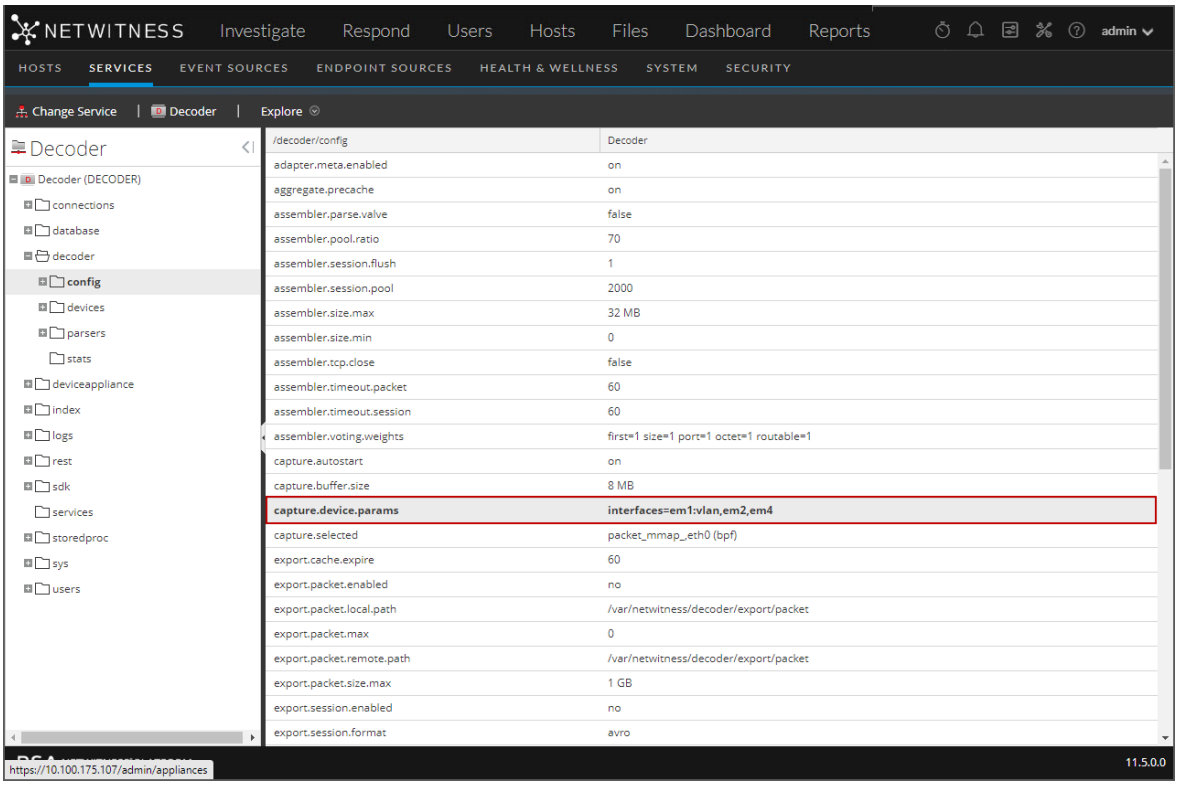
3. To go to the **Explore** view, click **Config** in the toolbar and select **Explore** in the drop-down list.

4. In the Services Explore view select **decoder** > **config**.



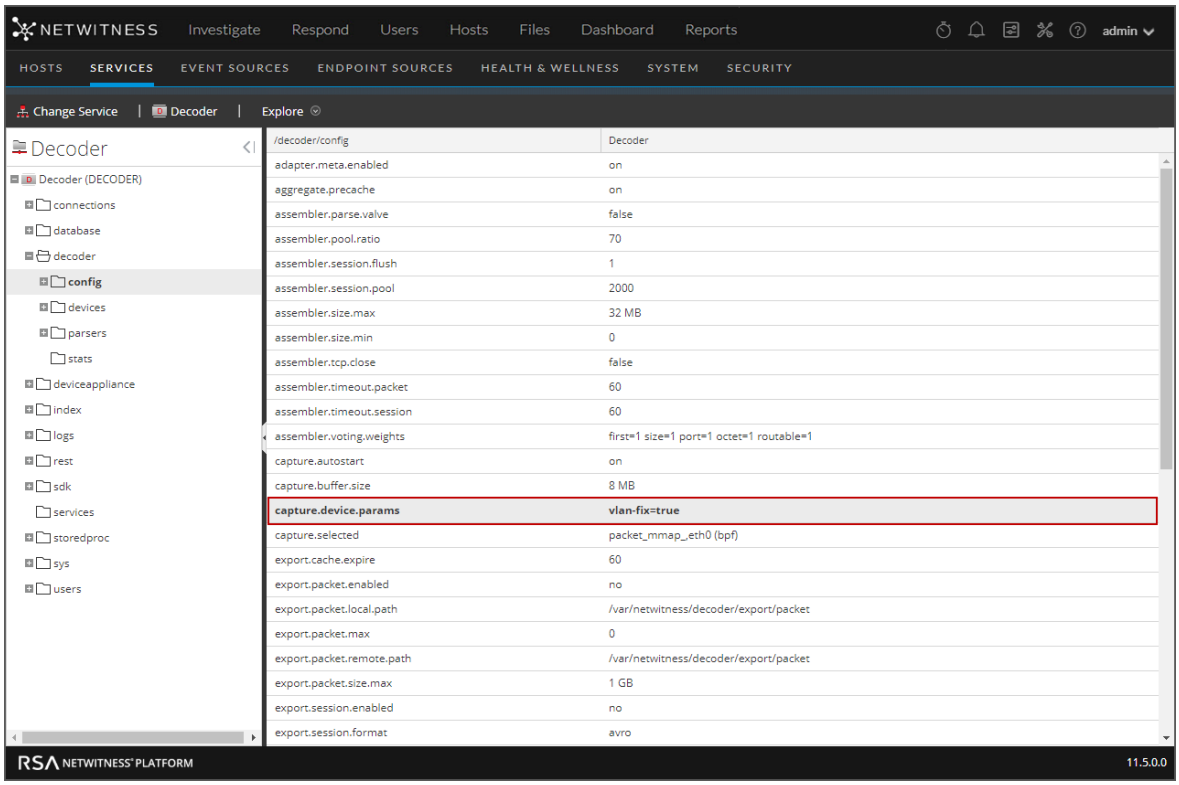
5. Click in the values column next to `capture.device.params`, and do one of the following:

   - To preserve VLAN tags on an interface in the interfaces list, add **`:vlan`** after the interface name and press **Enter**.
     For example, this specifies that VLAN tags are preserved on `em1`, but not on `em2` and `em4`:
     `interfaces=em1:vlan,em2,em4`



The change goes into effect after restarting capture; only traffic on `em1` has the VLAN tags preserved.

- To preserve VLAN tags on all interfaces, enter the following and press **Enter**:
  `vlan-fix=true`



VLAN tags are preserved on all capture interfaces.

6. Select Start Capture (Start Capture). The change takes affect after capture is restarted.

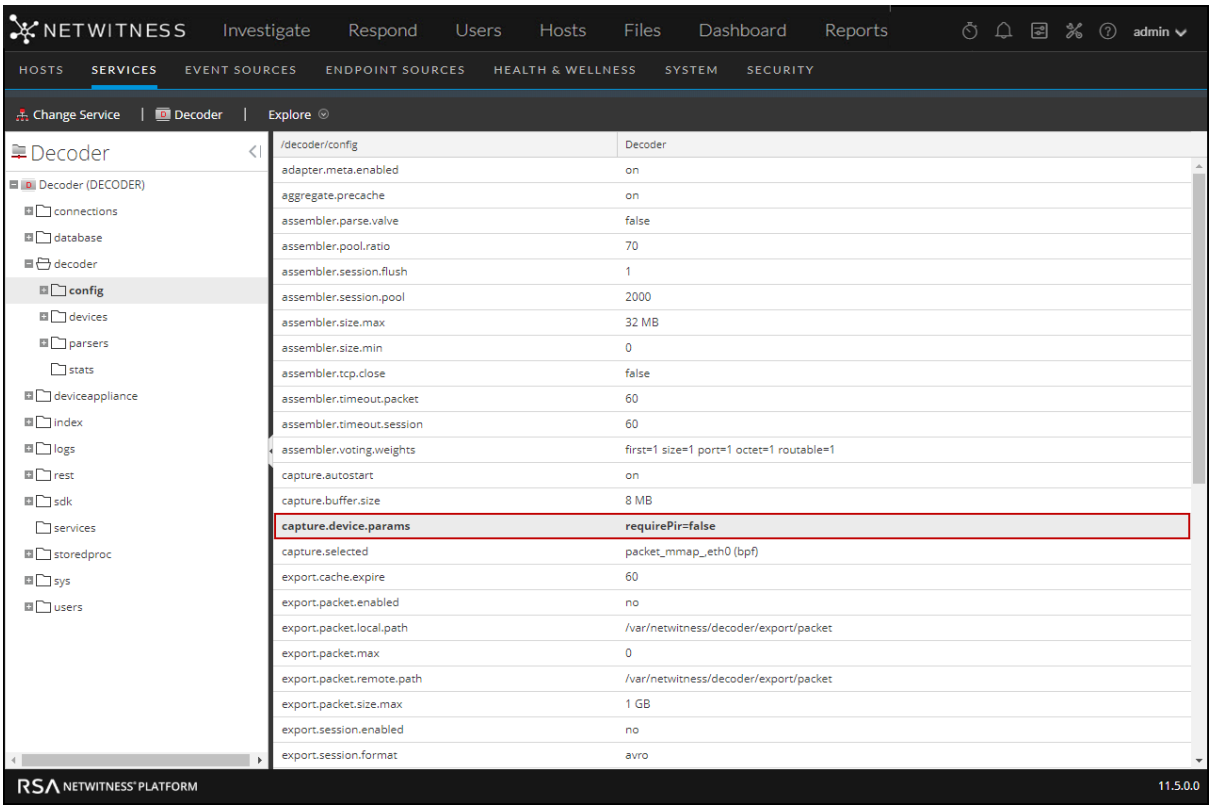
(Optional) Process Raw Syslog Data without Priority Field

You have the option to process raw syslog data that does not contain a valid priority (`PRI`) field.

To configure a Log Decoder to process syslog without a Priority field:

1. Go to (Admin) > **Services**, select the Log Decoder service and > **View > System**.
2. Select Stop Capture (Stop Capture).
3. From the drop-down menu, where System is shown, select **Explore**.
4. Select **decoder** > **config**.
5. In the `capture.device.params` field, add the following text, and then click **Enter** to save the changes:

   `requirePri=false`



6. From the drop-down menu, where Explore is shown, select **System**.
7. Select Start Capture (Start Capture). The change takes affect after capture is restarted.

(Optional) Configure Decoder to Support OpenAppID

The NetWitness Decoder can identify applications using the OpenAppID detectors. OpenAppID from Cisco is an application-layer network security plug-in for Snort (an open source network intrusion detection system). It is a set of open source Lua libraries (detectors) that identifies applications in the network traffic.

When Network Decoder detects an application with an OpenAppID detector, the `appid` meta is created. The value of the `appid` meta helps in identifying the application as it is the application's entry in the `appMapping.data` file. The `appMapping.data` file is part of the OpenAppID distribution and contains metadata of the application detectors.

Note: By default, all detectors are disabled to avoid any negative impact on the performance.

To configure the OpenAppID detectors:

IMPORTANT: Ensure that you enable only the required OpenAppID detectors. Enabling all the OpenAppID detectors can reduce the performance. Before you enable any detector, understand your application detection requirements and enable only the specific detectors. For example, enabling the `service_*` detector or any other detector that attempts to match short generic patterns can cause severe impact on the performance.

1. Download the OpenAppID source package (`snort-openappid.tar.gz`) from the Snort Downloads (https://www.snort.org/downloads) page.
2. Upload the OpenAppID source package to the network decoder (`http://<networkdecoder>:50104/decoder/parsers/upload`), which can be accessed using the REST interface.
   Once uploaded, decoder will extract the package to an appropriate folder. The default folder is `/etc/netwitness/ng/odp`.
3. To enable detectors, modify the `/decoder/parsers/config/openappid.enabled` file with a combination of detector filenames or wildcards, separated by commas. For example, `client_*, payload_slashdot.lua, payload_reddit.lua`.
   The following table describes the detector group filenames and their functions. You can extract the complete set of detectors from the OpenAppID source package (`snort-openappid.tar.gz`) or it can be viewed on the Decoder file system at `/etc/netwitness/ng/odp/lua`.

Detector Group Filename	Description
`client_*`	Detectors for client applications (for example, Facebook, MacAppStore).
`content_group*`	Detectors for a specific type of network attribute such as protocol, port, or service (for example, port_services which contains many services identified by their common ports).
`payload_*`	Detectors for the host based traffic (for example, Hulu, Reddit).
`service_*`	Detectors for service applications (for example, Minecraft, VNC).
`ssl_host_group_*`	Detectors for SSL hosts logically grouped together (for example, 334 which contains many ecommerce websites).

4. Ensure that you specify the detector filenames correctly. If you specify an incorrect filename, the detector will not load. The following table lists some good and bad file naming conventions and their status on subsequent loading.

Detector Group Filename	Status
`client_9P.lua`	loaded
`client_9p.lua`	not loaded (Linux only)
`client_9P`	not loaded
`client_9P*`	loaded
`client_`	loaded

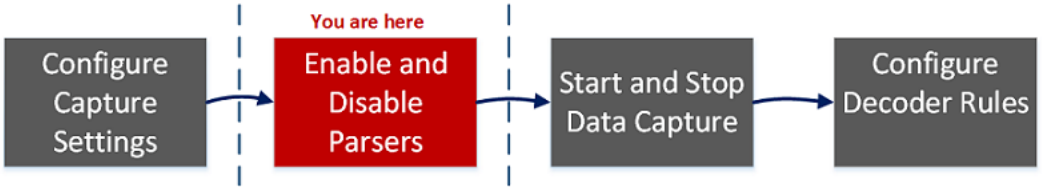
Detector Group Filename	Status
9P.*	
lient_ 9P.lua	not loaded
lient_ 9P.*	not loaded
*lient_ 9P.lua	loaded

5. Once the required detectors are enabled, reload the parser to load the detectors. The following example message will confirm that the detectors and enabled successfully.
```
2020-May-14 17:03:37 [OpenAppID] 59 detectors loaded out of 59
```

Enable and Disable Parsers and Log Parsers

Administrators can see which parsers have been downloaded from Live and deployed on a Decoder or Log Decoder, see which of these have been enabled, and enable or disable parsers and log parsers.

The following figure illustrates commonly used settings on a Decoder. For a quick basic setup with only the required steps, see Decoder and Log Decoder Quick Setup.



You should only download and deploy the parsers you need for the following reasons:

- There is an impact on performance as you increase the number of deployed parsers.
- The more parsers you deploy, the more meta data created, which impacts data retention.
- Not having extra (unnecessary) log parsers deployed reduces the potential for mis-identification of messages.

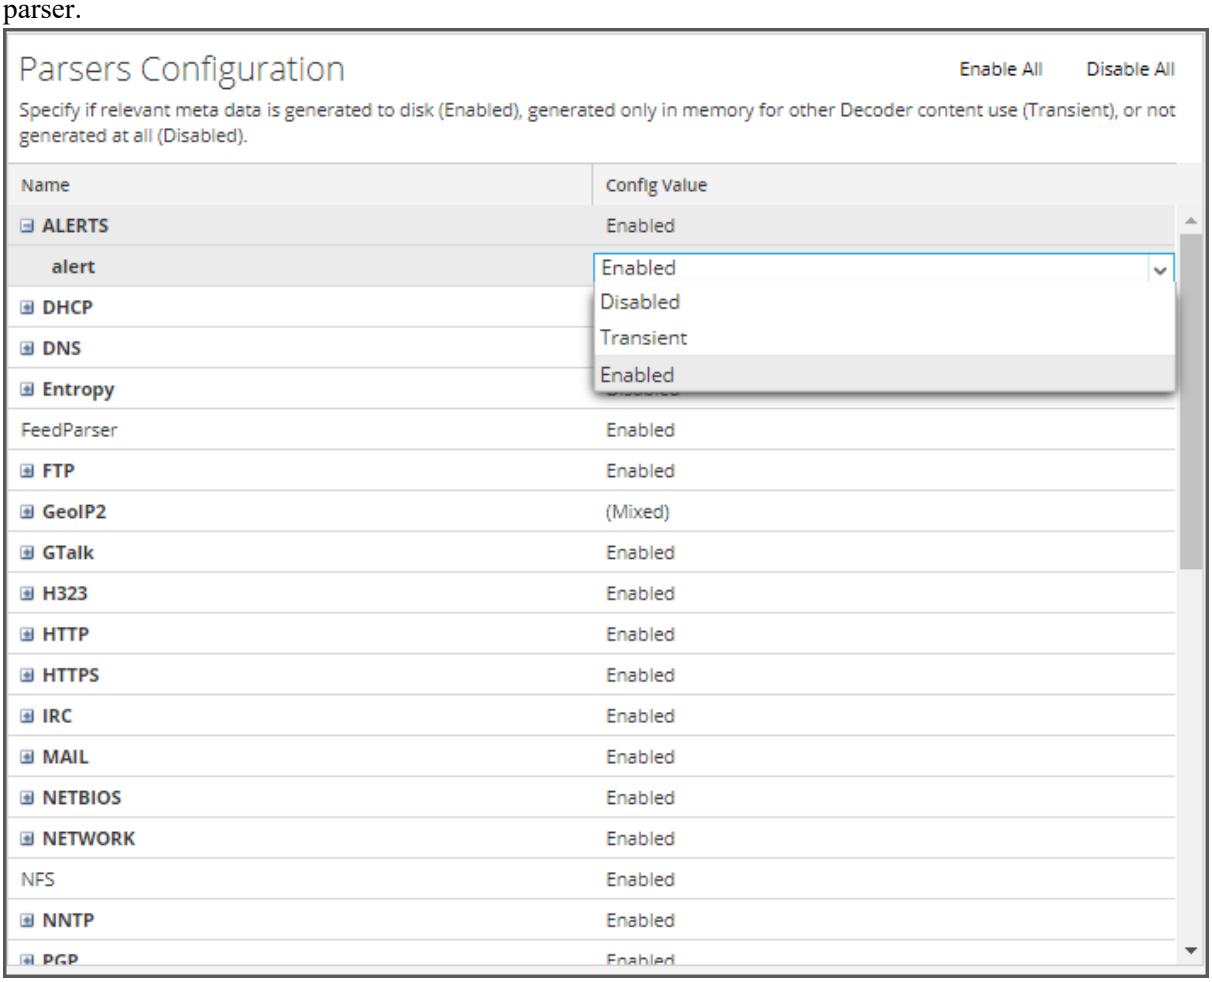
The Parsers Configuration panel provides a way to select parsers to use on the Decoder. Within some parsers, you can also configure the metadata that the parser creates. These are the options in the Parsers Configuration panel.

Option	Description
Enable All **Disable All**	These options provide a way to quickly select either all parsers or no parsers.
Name	The names of parsers available to the Decoder. A plus sign indicates that the metadata generated by the parser is configurable. Clicking the plus sign displays the metadata that the parser can create.
Config Value	A drop-down list changes the setting for the parser or metadata to **Enabled**, **Disabled**, or **Transient**. • When **Enabled**, the Decoder is using the parser to filter traffic. • When **Transient**, the Decoder is using the parser to filter traffic, and the generated metadata is not stored on disk. The transient metadata is available in memory to additional content (that is, parsers, feeds, and application rules) on that Decoder. This helps administrators to protect certain data and is usually done as part of a data privacy plan (see the *Data Privacy Management Guide*). • When **Disabled**, the Decoder is not using the parser. If the generated metadata for the parser is configurable, clicking the plus sign to expand the parser displays configurable meta keys and the same drop-down list selects the meta key the parser will create.

Note: For a Log Decoder, you must have previously deployed log parsers from Live. See the **Find and Deploy Live Resources** topic in the *Live Services Management* Guide for details. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

To enable or disable a parser, or to view the status for each parser:

1. Go to (Admin) > **Services**.
2. In the **Administration Services view**, select a Log Decoder or a Decoder, and **>View** > **Config**.
3. In the **Parsers Configuration** panel, look for the Decoder parser or the Log Decoder event source parser.



4. In the **Config Value** column, note the current status for your parser.

   You can update the status of any individual parser by selecting its **Config Value** and selecting **Disabled**, **Transient**, or **Enabled** from the drop-down menu. Alternatively, you can select **Enable All** or **Disable All** to update the status for all of your log parsers at once.

5. Click **Apply**.

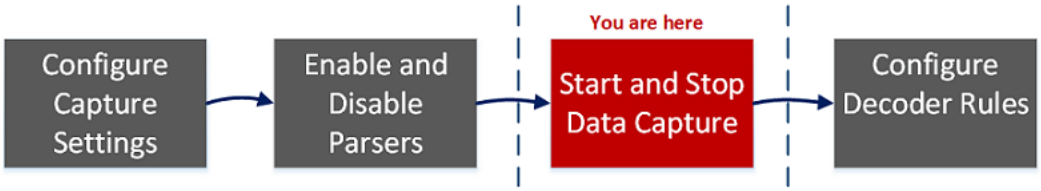
When you click **Apply**, note that all parsers are reloaded into NetWitness. The status for each parser is updated, based on your selections.

Start and Stop Data Capture

When a Decoder starts up, it automatically begins aggregating data if **Capture Autostart** is enabled. When autostart is not enabled, you can start and stop data capture manually.

Note: The Capture Configuration Settings in the Service Config view for a Decoder determine whether Capture Autostart is enabled.

The following figure illustrates commonly used settings on a Decoder. For a quick basic setup with only the required steps, see Decoder and Log Decoder Quick Setup. You may want to stop and start capture at other times, for example, before you shut down the service.



To start and stop capture:

1. Go to (Admin) > **Services**.
2. Select a Decoder or Log Decoder service, and select **> View > System**.
3. In the toolbar, click **Start Capture**.

   If the service is a Decoder, it begins capturing packets. If the service is a Log Decoder, it begins capturing logs.

   When packet or log capture is in progress, the option in the toolbar changes to **Stop Capture**, and the option to upload a file is unavailable.
4. Whenever you want to discontinue traffic capture on a Decoder, click **Stop Capture**.

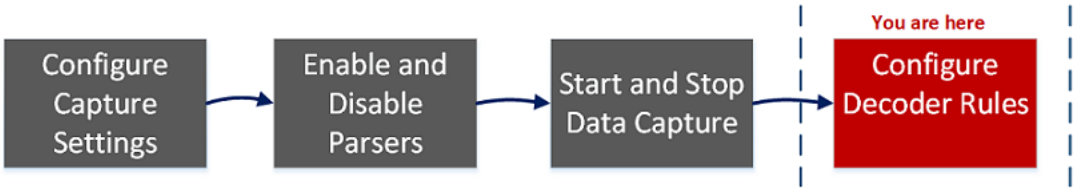
   Packet or log capture ceases, and the option to upload a file to the service is again available.

Note: When you stop the Log Decoder service while capture is running, all events currently in Log Decoder memory will be processed and persisted. Should an issue arise where it is necessary to quickly shutdown the service, use the Services Explore view to stop capture (`/decoder stop`), passing the parameters `flush=false` before stopping the Log Decoder service. For further information, see the "Services Explore View" in the *Host and Services Getting Started Guide*.

Configure Decoder Rules

This topic provides procedures for creating and managing rules for Decoder or Log Decoder traffic capture in the **Services Config view** > **Rules** tabs . Services Config View - Rules Tabs provides details about the Rules tab options.

The following figure illustrates commonly used settings on a Decoder. For a quick basic setup with only the required steps, see Decoder and Log Decoder Quick Setup.



Capture rules can add alerts or contextual information to sessions or logs. They can also define which data is filtered out by a Decoder or Log Decoder. Rules are created for specific metadata patterns, which result in predefined actions when matches are found. For example, to keep all traffic that fits certain criteria, but discard all other traffic, you can create a rule to perform the necessary actions. When applied, rules affect both packet capture file importing, as well as live network capture.

Rule and Query Guidelines provides guidelines that all queries and rule conditions in NetWitness Core Services must follow.

By default, no rules are defined when you first install NetWitness. Until rules are specified, the packets are not filtered. You can deploy the latest rules from Live. You can define three types of rules: Network Rules, Application Rules, and Correlation Rules.

- Network rules are applied at the packet level and are made up of rule sets from Layer 2, Layer 3, and Layer 4. Multiple rules can be applied to the Decoder. Rules can be applied to multiple layers (for example, when a network rule filters out specific ports for a specific IP address). Network rules are only available on Network Decoders.
- Application rules are applied at the session level. If the first rule listed is not a match, the Decoder then attempts to match the next rule listed, until a match is found.
- Correlation rules are applied over a configurable sliding time window. When a match is found, the service creates a new super session that identifies other sessions that match the rule, then creates a session list for analysis.

The two most common uses of rules are:

- To alert, and thereby create a custom alert meta value, when certain conditions are found.
- To filter out certain types of traffic that do not add value to the analysis of the data.

Groups of capture rules form rule sets, which you can import and export. This feature enables use of multiple rule sets for various scenarios. You can import the exported rule set, in the form of an .nwr file, to other NetWitness services, simplifying the deployment and configuration of multiple services.

Rule Processing

These are the principles governing capture rule processing:

- Multiple rules can be applied to the Decoder.
- Capture rules are executed one after the other, in sequence.
- Rule processing stops when all rules are processed or after a rule configured to stop rule processing is matched.
- A default rule can be used to either include or exclude all traffic not otherwise selected by a rule. A default rule, if used, must always be placed at the bottom of the rule list. Otherwise, rule processing stops as soon as the default rule is evaluated since, by definition, all traffic is selected by the default rule.
- When rule processing stops, the session is saved using the configured session options and debug options.

Rule and Query Guidelines

All queries and rule conditions in RSA NetWitness Core services must follow these guidelines:

All string literals and time stamps must be quoted. Do not quote numbers, MAC, or IP addresses.

- `extension = 'torrent'`
- `time='2015-jan-01 00:00:00'`
- `service=80`
- `ip.src = 192.168.0.1`

Note: The space on the right and the left of an operator is optional. For example, you can type a rule as `service=80` or `service = 80` .

Rule Examples

The following table shows examples of rule conditions. You can use rule conditions for log retention collections in an Archiver and for application, network, and correlation rules on a Decoder, Log Decoder, or Concentrator. Rule conditions are also used in all `WHERE` clauses in all Core database queries.

For detailed information on rule syntax in NetWitness, see "`WHERE` Clauses" in the "Queries" section of the *Core Database Tuning Guide* .

Rule Name	Condition
ComplianceDevices	`device.group='PCI Devices' || device.group='HIPPA Devices'`
HighValueWindows	`device.group='Windows Compliance'`
MediumValueWindows	`device.type='winevent_nic' && msg.id='security_4624_security'`

Rule Name	Condition
LowValueWinLogs	`device.type='winevent_nic' && msg.id='security_4648_security'`
LowValueProxyLogs	`device.class='proxy' && msg.id='antivirus_license_expired'`
GeneralWindows	`device.type='winevent_nic'`

Invalid Rules

NetWitness uses a rule parser that strictly defines valid syntax for rules and queries. When a Core service encounters invalid syntax, it writes a warning in the NetWitness logs indicating the error.

Note: NetWitness 11.x does not support parsing of legacy syntax rules (as version 10.6 did). After you update to NetWitness 11.x, rules with invalid syntax are highlighted in the user interface, and no rules will be applied until the invalid rules are corrected. The Rule Editor provides additional tooltips. After you fix the rules, the highlights disappear. See Fix Rules with Invalid Syntax.

The `/decoder/config/rules/rule.errors` and `/concentrator/config/rules/rule.errors` stats, contain the count of rules with errors. If `rule.errors` is nonzero, NetWitness generates a Health and Wellness alert to indicate that you need to fix the rules.

General Syntax Guidelines

- All text values must quote literal values. Example: `user = 'user1'`
- Quotes can use single or double quotes; but they must match. (You cannot start with a single quote and finish with a double quote.)
- If the literal value has a quote, you can escape it using double quotes. Example: `user = "User's"`

The following are valid syntax rules:

- All time values should use quotes for dates in this form: `time = 'YYYY-MM-DD HH:MM:SS'`
- All time values that are the number of seconds since EPOCH (Jan 1, 1970), should not be quoted. Example: `time = 1448034064`
- Everything else is unquoted: IP addresses, MAC addresses, numerics, and so on. Example: `service = 80 && ip.src = 192.168.1.1/16`

Capture Rule Syntax

Capture rules compare fields to values or to other fields. This is an example of a simple expression with a meta key on the left side of the operator and a value on the right side.

```
ip.dst=192.168.1.1
```

The syntax allows a meta key on the right side of the operator in Decoders and Log Decoders for application and network rules. Meta key comparison does not apply in the `where` clause in queries. This is an example of a simple expression with a meta key on the left side of the operator and a meta key on the right side.

```
ip.src=ip.dst
```

Rules that include a meta key comparison support renamed meta keys; if a rule queries a meta key that has been renamed, the rule is parsed for the renamed meta key. For example, if the meta key `ip_dst` is used in a rule, it is transparently mapped to the renamed meta key: `ip.dst`. Existing rules that include original keys will trigger alerts that include data for the renamed meta key.

This is an example of a rule that finds packets having the same `ip.src` address and `ip.dst` address on a Decoder, and generates an alert on the Concentrator.

```
alert=ioc name=testRule8 rule="ip.src=ip.dst" order=38
```

This rule would generate an error because `eth.src` and `ip.src` are incompatible formats.

```
rule="eth.src=ip.src" name="testRule99" alert=ioc
```

Values can be expressed as discrete values, a range of values, an upper or lower bound, or a combination of these three. You can create a greater than or less than comparison, and test equality or inequality against a range of values or an upper/lower bound.

`key 0-5` (a range of values)

`key = 0-u` is the same as `key >= 0` (upper bound, greater than or equal to)

The following table summarizes the operators on meta keys.

Left Operand Format	Operator	Right Operand Format	Description
any	=	compatible with left operand	Equality operator. You can use values or meta keys on the right side of the equality operator.
any	!=	compatible with left operand	Inequality operator. You can use values or meta keys on the right side of the inequality operator.
any	<	compatible with left operand	Less than operator. You can use values or meta keys on the right side of this operator.

Left Operand Format	Operator	Right Operand Format	Description
any	`<=`	compatible with left operand	Less than or equal to operator. You can use values or meta keys on the right side of this operator.
any	`>`	compatible with left operand	Greater than operator. You can use values or meta keys on the right side of this operator.
any	`>=`	compatible with left operand	Greater than or equal to operator. You can use values or meta keys on the right side of this operator.
text	`contains`	text	Find values that contain the right operand. You can use meta keys or values on the right side of this operator.
text	`begins`	text	Find values that begin with the right operand. You can use meta keys or values on the right side of this operator.
text	`ends`	text	Find values that end with the right operand. You can use meta keys or values on the right side of this operator.
text	`length`	integer	Find strings of a certain length. You can use meta keys or values on the right side of this operator.
any	`count`	integer	Find values with a specific number of occurrences within the session. You can use meta keys or values on the right side of this operator.
any	`ucount` and `unique`	integer	Finds a number of uniquely occurring values. You can use meta keys or values on the right side of this operator. For example, if the results include instances of a meta key with five unique values and three of the same value, the `ucount` is six.
N/A	`exists`	any	Finds any values for the meta key. You can use meta keys or values on the right side of this operator.
N/A	`!exists`	any	Finds any sessions in which the meta key does not occur. You can use meta keys or values on the right side of this operator.
text	`regex`	text	Finds values matching a regular expression. You can use values on the right side of this operator.

The following table summarizes other syntax elements used in rules.

Syntax element	Description
*	Default rule. By using an asterisk (*) as the sole character in a rule, that rule will select all traffic.
u	Upper bound of a range of times, IP addresses, or numeric formats. For example, to select all TCP ports above 40000, the syntax would be: `tcp.port = 40000-u`
l	Lower bound of a range of times, IP addresses, or numeric values. For example, to select all TCP ports below 40000, the syntax would be: `tcp.port = l-40000`
- (dash)	Denotes a range. This is only applicable to time values, IP or MAC addresses, or numeric values. Separate the lower and upper bounds of the range with a dash (-) character. For example, to select TCP ports between 25 and 443, the syntax would be: `tcp.port = 25-443`
, (comma)	Denotes a list of ranges or values or meta keys. Single values may be used as well as any combination of ranges and upper or lower bounds. Single meta keys may be used in a list. Meta keys and literal values cannot both appear on the right-hand side of an operator. For example, the following is valid syntax: `tcp.port = 1-10,25,110,143-225,40000-u`
( )	Grouping operator. An expression can be enclosed in parentheses to create a new logical expression. For example, the following would select traffic on port 80 to/from 192.168.1.1 OR traffic on port 443 to/from 10.10.10.1: `(alias.ip=192.168.1.1 && tcp.port=80) || (alias.ip=10.10.10.1 && tcp.port=443)`
~	Logical NOT operator, a negation of an expression.
&&	Logical AND operator, a conjunction of two expressions.
||	Logical OR operator, a disjunction of two expressions.

Configure Capture Rules

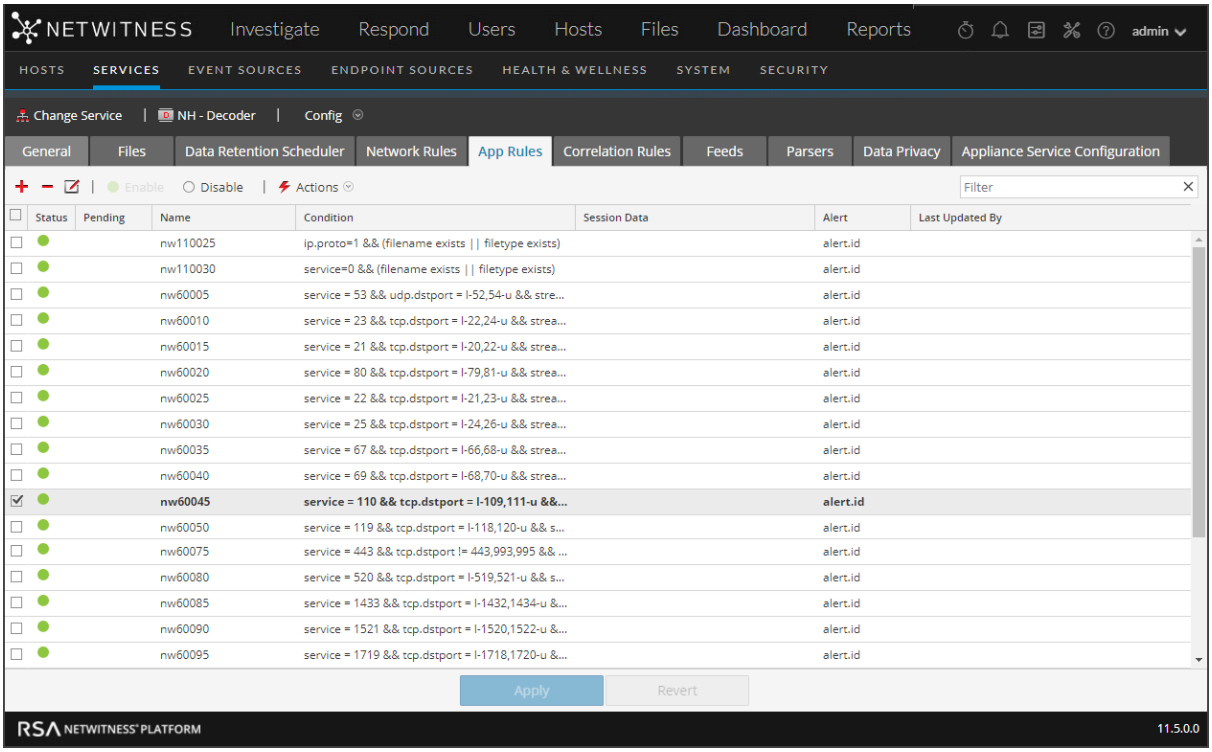
The Decoder and Log Decoder rules are editable in the Services Config view. While each type of rule (network, application, and correlation) has its own tab; the functions are similar for all types of rules. You can:

- Add, edit, and delete rules
- Enable and disable rules
- Change the execution sequence of rules
- Import rules from a file
- Export rules to a file

- Push rules to another service
- Revert or apply rule changes
- Restore one of the last ten rule configurations from a snapshot

To configure rules in the Rules tabs

1. Go to (Admin) > **Services**.
2. In the **Services** view, select a Decoder service and > **View** > **Config**.
3. In the **Services Config** view, select one of the Rules tabs: Network Rules, App Rules, or Correlation Rules.
   The rules list for the selected rule type is displayed.



Each type of rule has a list with slightly different columns and different parameters. Several basic guidelines apply to all rule management activities:

- The rules are executed in the sequence they are displayed in the list. To change the execution sequence of rules, drag and drop rules to the appropriate location in the list or use the context menu options to arrange the rules in the list.
- To select a single row, click the row.
- To select a group of adjacent rows, click the first, then shift-click the row at the end of the group.
- To select multiple non-adjacent rows, click the first, then control-click the others.
- When editing rules in the Rules tab, you must apply the configuration changes in order to activate.

- Until changes are applied, you can discard edits to the list and revert to the unedited rules.
- Once rules are applied, you can recover the last ten rules configurations using the **History** option in the **Actions** menu.

To add a rule in any Rules tab, do one of the following:

- Click  .
- Right-click a rule, and select **Insert Above** or **Insert Below** from the context menu.
  The Rule Editor dialog for that type of rule is displayed.

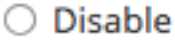
To remove a rule:

1. From any Rules tab, select the rules to remove from the rules list.
2. Click .
   The selected rules are removed from the list, but still exist on the service.

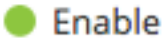
To edit a rule

1. From any Rules tab, select the rule to edit.
2. Click  or double-click the rule row.
   The Rule Editor dialog for that type of rule is displayed.

To disable a rule:

1. From any Rules tab, select the rules to disable.
2. Click  .
   The status changes to disabled in the rules list, but the rule is still enabled on the service.

To enable a rule:

1. From any Rules tab, select the rules to enable.
2. Click  .
   The status changes to enabled in the rules list, but the rule is still disabled on the service.

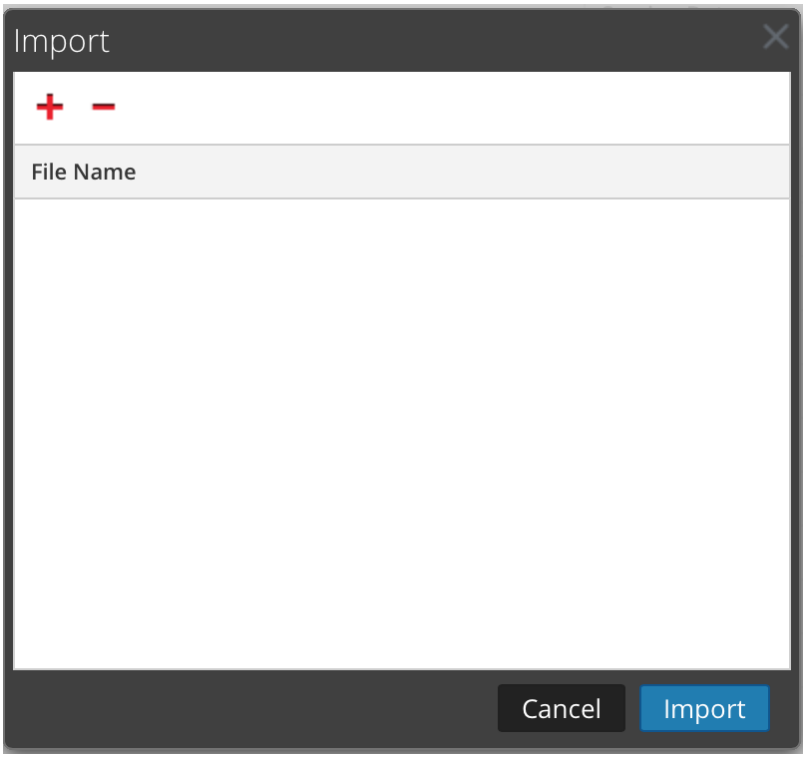
Import Rules from a File and Export Rules

You can import network, application, and correlation rules to a Decoder from a file that contains rules of the same type. After the rules are imported, you can edit and manage them as you would any other rules.

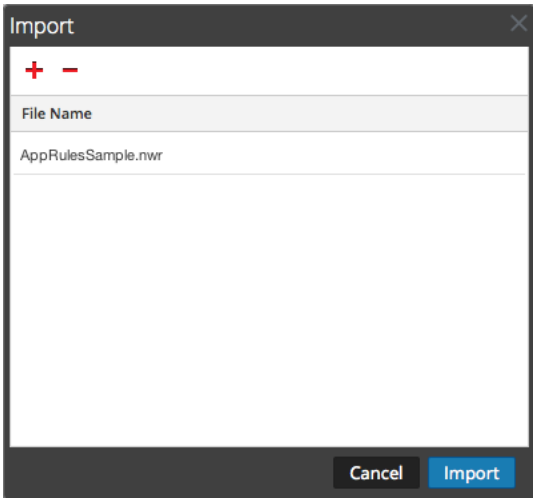
When you attempt to import a group of rules, NetWitness Administration checks the type of rules imported. If you are successful, a message displays the number of rules imported. If the rule type differs from the active tab type, the rules are not imported. You must re-import the rules under the correct tab or select another file to import.

To import rules to a service:

1. From any Rules tab, select Actions > Import .
   The Import dialog is displayed.

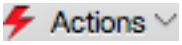


2. Click + .
   A view of the directory structure is displayed.
3. Choose one or more NetWitness rules (.nwr) files to import, and click **Open**.
   The file is added to the list in the Import dialog.



4. Click **Import**.
   The rules are imported into the user interface. Imported rules have a red corner in each edited column.
5. Edit or reorder the rules if needed.
6. To save the rules to the service, click **Apply**.
   The rules for the service are updated with the changes.

To export a rule to a file:

1. To export a subset of the rules, select the rules to be exported.
2. Do one of the following:
   - In the toolbar, select 

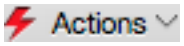
 > Export > Selection. (**Export > All** exports all rules in the rules list even if you have a subset selected for export.)
   - Right-click the selected rules and select **Export Selection**.

     A prompt for the filename is displayed.
3. Enter the filename and click **Export**.
   The .**nwr** file is downloaded.

Push Rules to Other Services

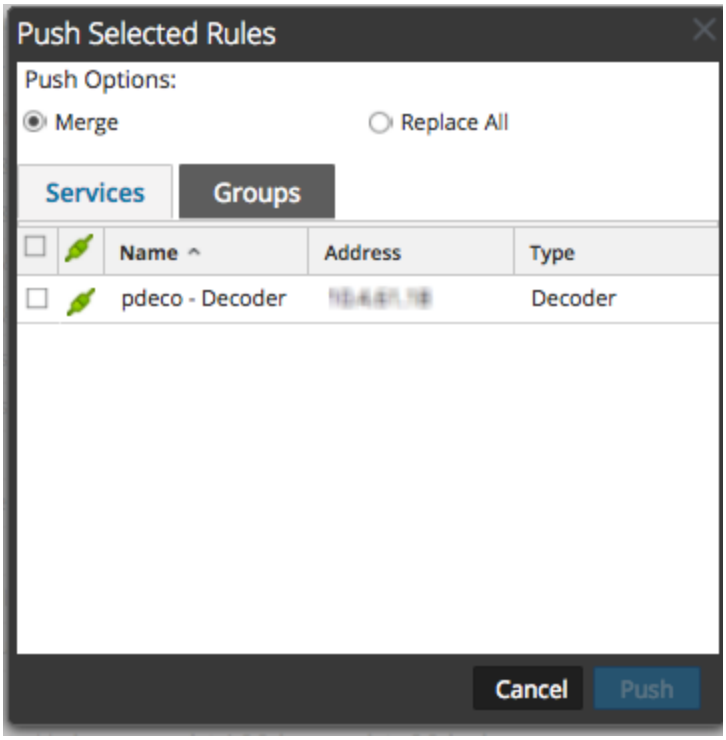
You can apply (push) rules or selected rules to other services (Decoders or Log Decoders) or service groups. When you push all rules to other services, all rules on the target services are removed and replaced with all of the rules on the source service.

To push selected rules from this Decoder to other Decoders:

1. From any Rules tab, select the rules that you want to push to another Decoder.
2. Do one of the following:
   - Select 

 > Push > Selection.

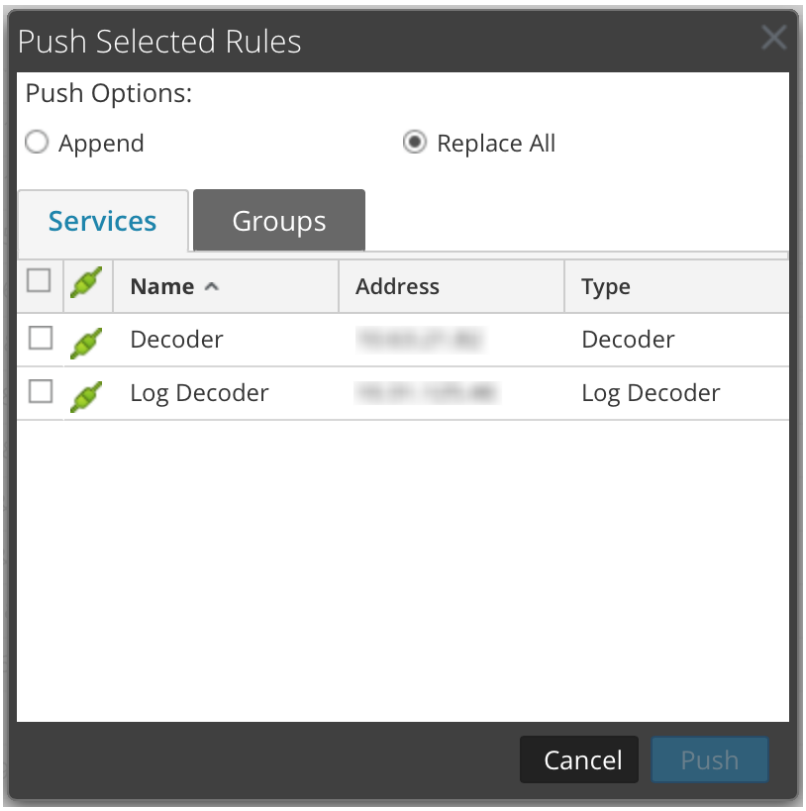
- Right-click the selected rules and select **Push Selected Rules**.
  The Push Selected Rules dialog is displayed.



3. Select a Push Option:
   - Select **Replace All** to delete all rules on the target services and replace them with the selected rules. This is the default selection.
   - Select **Merge** to merge the selected rules with the existing rules on the target services.
4. On the **Services** tab, select the target services to receive the pushed rules, or select the groups of services from the **Groups** tab.
5. Click **Push**.
   The rules are pushed to the selected services and become effective immediately.

To push all rules from this Decoder to other Decoders:

1. From any Rules tab, select Actions **> Push > All**.
   (**Push > All** pushes all rules in the rules list even if you have a subset selected to push.) The Push Selected Rules dialog is displayed.



2. On the **Services** tab, select the target services to receive the pushed rules, or select the groups of services from the **Groups** tab.
3. Click **Push**.
   All rules from the target services are deleted and replaced with all of the rules from source service. The rules become effective immediately.

Change Execution Order of Rules

Capture rules are applied in the order they are displayed in the rules list. To reorder rules, use either of these methods:

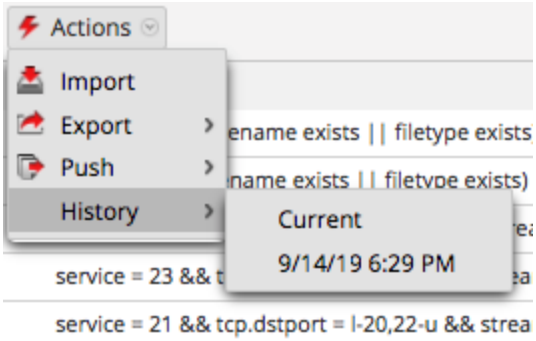
- Drag and drop the rules in the appropriate location in the rules list.
- Right-click a rule to display the context menu, and use the **Cut** and **Paste** options.

Restore a Rule Snapshot from History

NetWitness keeps the last ten snapshots of rules applied to a service.

To restore a rules snapshot from history:

1. Select Actions > **History**.
   A submenu of snapshots is displayed.



2. Select the snapshot time from the submenu.
   The rules from the snapshot are loaded into the rules list, replacing the current set. But the current set is still in use on the service.
3. To apply the rules to the service, click **Apply**.
   The rules are applied to the service.

Configure Application Rules

Application layer rules are applied at the session level. The following are sample application rules.

To truncate packets carried via Server Message Block protocol (SMB), create a rule as follows:

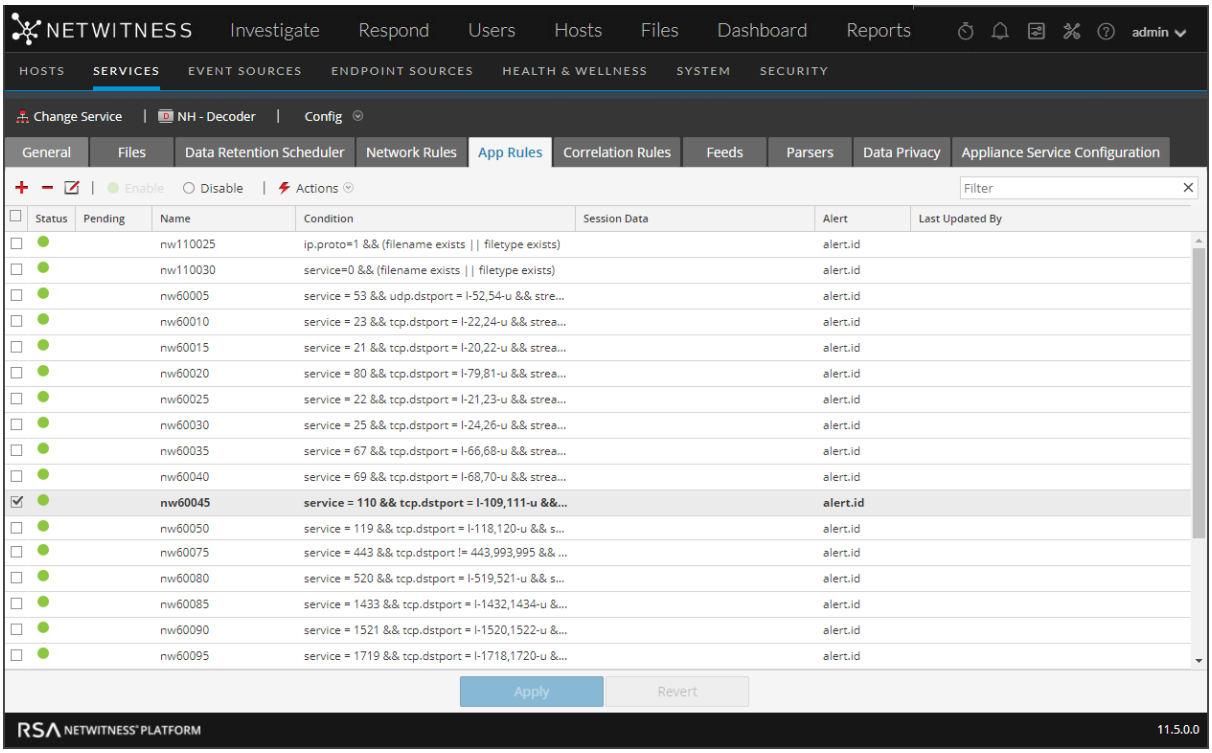
- Rule Name: Truncate SMB
- Condition: `service=139`
- Rule Action: Truncate All

To retain email to and from a specific e-mail address, create a rule as follows:

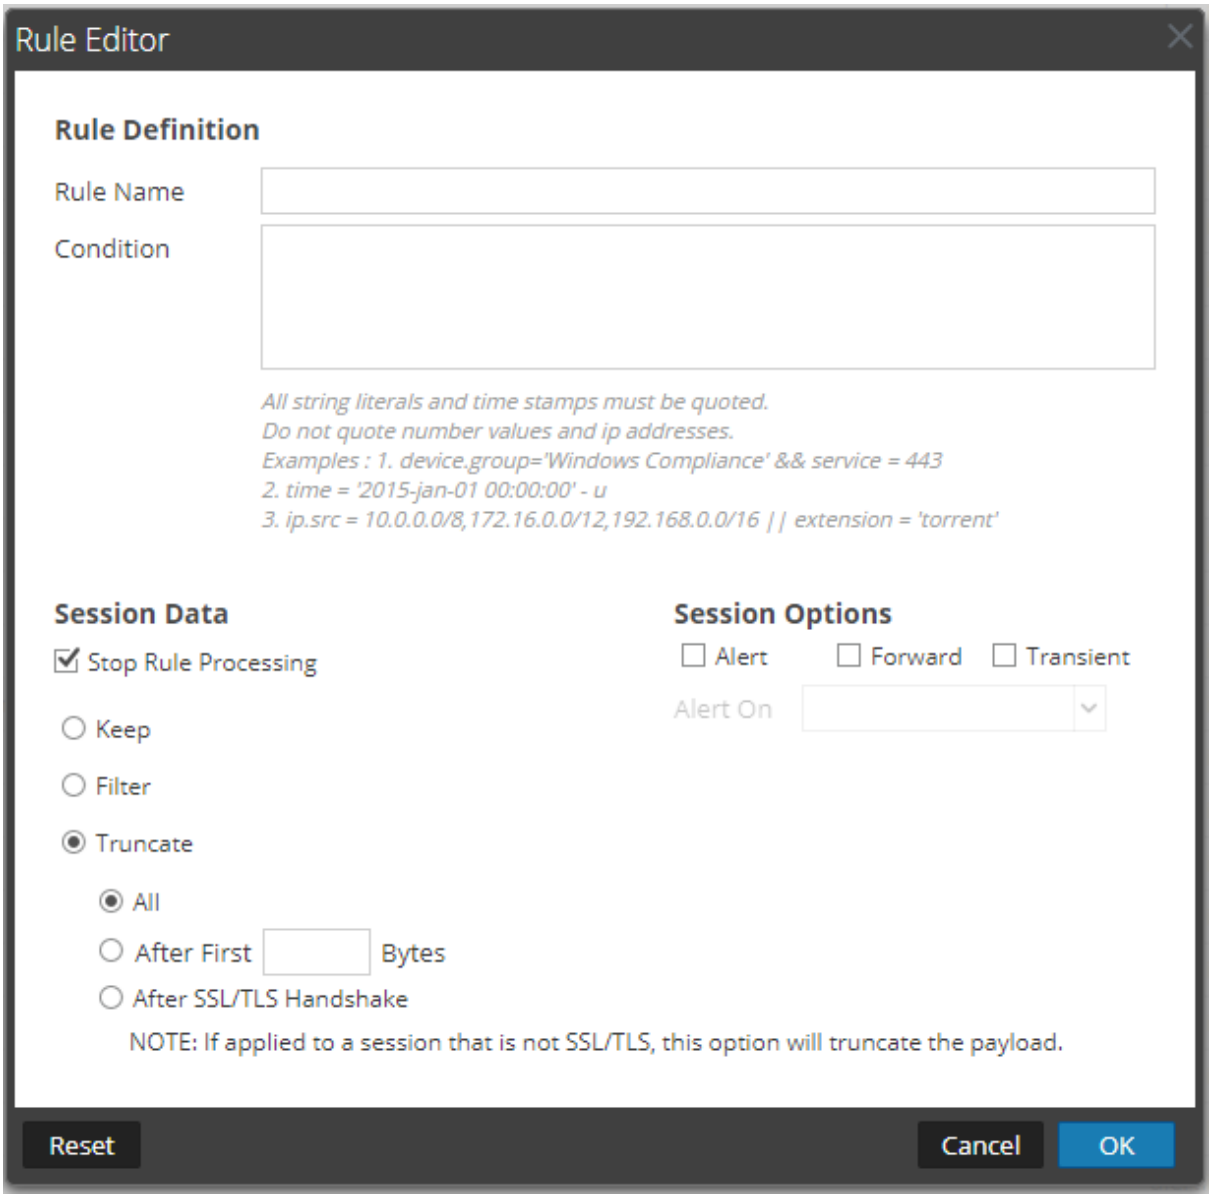
- Rule Name: Retain Email Tom Jones
- Condition: `email='Tom.Jones@TheShop.com'`
- Rule Action: Keep

To add or edit an application rule:

1. Go to  (Admin) > **Services**.
2. Select a Decoder or Log Decoder service and  > **View** > **Config**.
   The Systems Config view for the selected service is displayed.
3. Select the **App Rules** tab.



4. Do one of the following:
   - If adding a new rule, click  .
   - If editing a rule, select the rule from the rules list and click .
5. The Rule Editor Dialog is displayed with application rule parameters.



   a. In the **Rule Name** field, type a name for the rule. For example, for a rule that truncates all SMB, type **Truncate SMB**.
   b. In the **Condition** field, build the rule condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the window actions. As you build the rule definition, NetWitness displays syntax errors and warnings. For example, to truncate all SMB, type **service=139**.

All string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details.

c. If you want rule evaluation to end with this rule, check the **Stop Rule Processing** checkbox.

d. In the **Session Data** section, choose one of the following actions to apply when a matching packet is found:

- **Keep**: The packet payload and associated meta are saved when they match the rule.
- **Filter**: The packet is not saved when it matches the rule.
- **Truncate**: Select a truncate option to execute when a packet matches the rule. The example uses the **All** option.
- **Truncate All** to save the packet headers and associated metadata, and do not save the packet payload.
- **Truncate After First <n> Bytes** to save the packet headers and associated metadata, and do not save the packet payload after the specified first <n> bytes, where <n> is a number of bytes.
- **Truncate SSL/TLS Handshake** to truncates the payload for all sessions except in the case of an SSL/TLS session, where the SSL exchange is preserved, but the rest of the payload is not saved. This option is for use with SSL parsers.

e. In the **Session Options** section, do any of the following:

- **To generate a custom alert** when a session metadata matches the rule, enable the Alert flag and select the name of the alert meta from the **Alert On** drop-down list.
- **To perform syslog forwarding** when the log matches the rule, enable the **Forward** flag. Make sure that:
  - You have enabled both the Alert and Forward flags to carry out syslog forwarding.
  - The name of the rule mentioned in the Rule Editor dialog matches the syslog forwarding destination name specified in the Log Decoder > View > Explore > `/decoder/config/logs.forwarding.destination` parameter
- **To prevent the alert metadata that is created from being written to the disk**, enable the **Transient** flag.

6. To save the rule and add it to the grid, click **OK**.

   The rule is added at the end of the grid or inserted where you specified in the context menu. The plus sign is displayed in the **Pending** column.

7. Check that the rule is in the correct execution sequence with other rules in the grid. If necessary, move the rule.

8. To apply the updated rule set to the Decoder or Log Decoder, click **Apply**.

NetWitness saves a snapshot of the currently applied rules, then applies the updated set to the Decoder and removes the pending indicator from the rules that were pending.

Monitor Application Rules

The Decoder and Log Decoder keep track of how many times each application rule matches a session. These stats can be viewed by connecting to the Decoder or Log Decoder Explore view and viewing the properties on the `/decoder/config/rules/application` folder. Then, send the command "`statdump`" to that folder. The output of this message is a listing of the number of times each application rule is hit. The listing is ordered in the same order as the contents of the rule definitions in the `/decoder/config/rules/application` folder. For example, on a system with three application rules:

```
0001: hits=6543 loaded=true
0002: hits=9294 loaded=true
0003: hits=43 loaded=true
```

The hit counters for the application rules are reset whenever the parsers are reloaded.

Configure Correlation Rules

Basic Correlation Rules are applied at the session level and alert the user to specific activities that may be occurring in their environment. NetWitness applies correlation rules over a configurable sliding time window. When the conditions are met, alert metadata is created for this activity and there is a visible indicator of the suspicious activity.

The following are sample correlation rules illustrating two use cases and the syntax.

Objective: In sessions where `port.dst` exists, if there is any combination of ip.src and ip.dst where the count of unique instances of `port.dst > 5` within one minute, then alert. To achieve this objective, create a rule as follows:

- Rule Name: IPv6 Vertical TCP Port Scan 5
- Rule: `port.dst exists`
- Instance Key: `ip.src,ip.dst`
- Threshold: `u_count(port.dst)>5`
- Time Window: 1 min

Objective: In sessions where `action==login` and `error==fail`, if there is any combination of `ip.src` and `ip.dst` that appears in more than 10 sessions within five minutes, then alert. To achieve this objective, create a rule as follows:

- Rule Name: IPv4 Potential Brute Force 10
- Rule: action=`'login' && error='fail'`
- Instance Key: `ip.src,ip.dst`
- Threshold: `count()>10`
- Time window: 5 mins

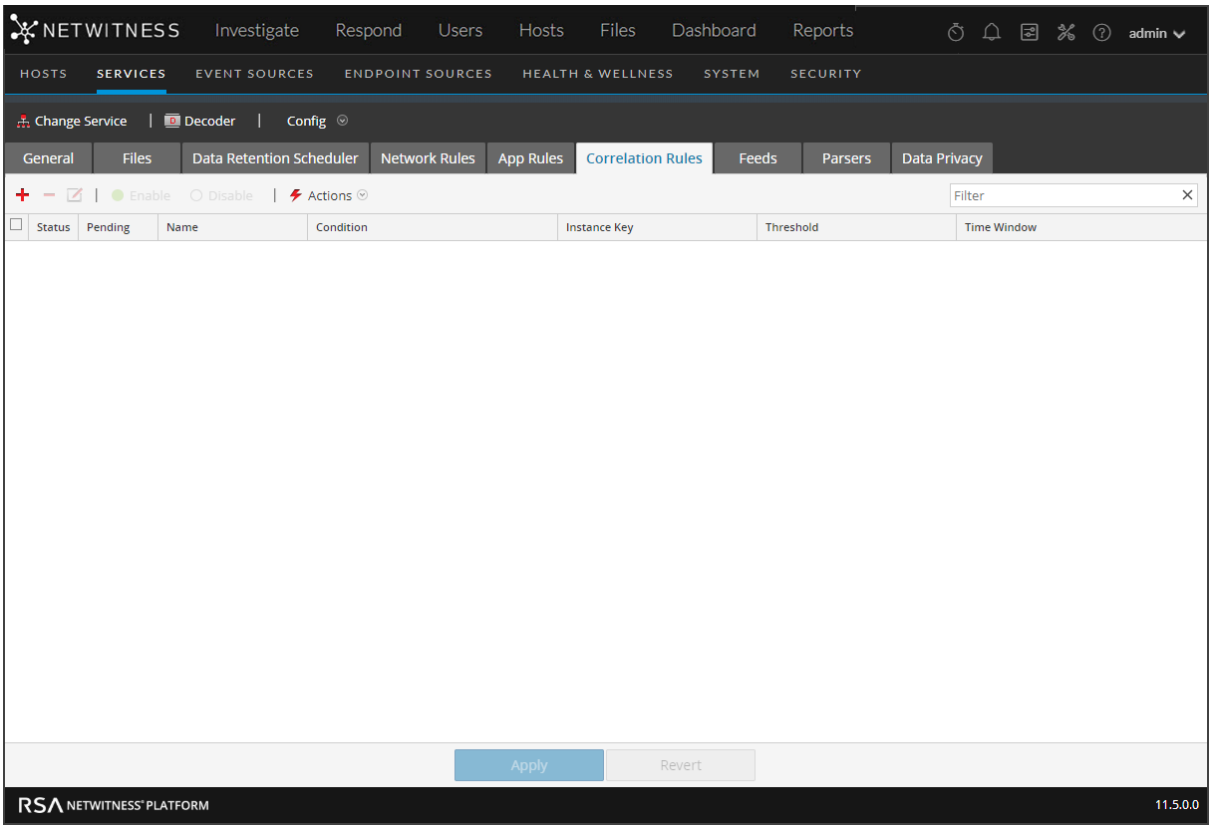
Both sample rules have the same instance key: `ip.src` and `ip.dst`. Because we are looking for unique combinations of `ip.src` and `ip.dst` that match the correlation condition, **`ip.src`** and **`ip.dst`** are **primary keys**.
Threshold can include an **associated key** that identifies the meta type that we are counting to determine if the condition is satisfied. In the first example, the associated key specified in Threshold is `port.dst`. We are counting unique instances of `port.dst` for every `ip.src/ip.dst` pair. In the second example, the associated key is not specified in the Threshold because it is merely a count of sessions. It is helpful to think of this scenario as counting unique session IDs and the associated meta is implicitly session.id. We are counting unique `session.id` for every `ip.src/ip.dst` pair.

Invalid use case: In sessions where (rule), if there is any combination of `ip.src` and `ip.dst` that have a unique count of `ipv6.dst > 5` within (time window), then alert. This case does not work because the associated key `ipv6.dst` is an IPv6 meta type. IPv4 and IPv6 meta types are not permitted to be used as associated keys.

To add or edit a correlation rule

1. Go to  (Admin) > **Services**, select a service, and click  > **View** > **Config**.
   The Service Config view for the selected service is displayed.

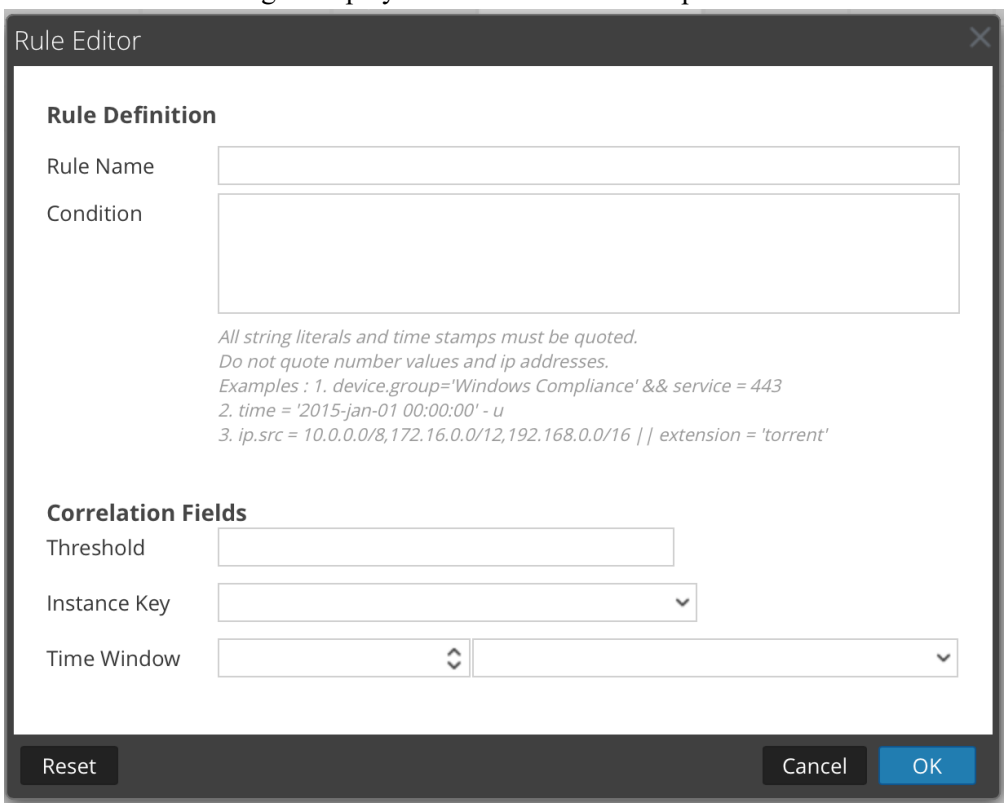
2. Select the **Correlation Rules** tab.



3. In the **Correlation Rules** tab, do one of the following:

   - If adding a new rule, click .

- If editing a rule, select the rule from the rules grid and click .
  The Rule Editor dialog is displayed with correlation rule parameters.



4. In the **Rule Name** field, type a name for the rule. For example, to create the sample rule, **IPv6 Vertical TCP Port Scan 5**.
5. In the **Condition** field, build the rule condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the window actions. As you build the rule definition, syntax errors and warnings are displayed by NetWitness. For example, to create the sample rule, type **tcp.dstport exists**. When this condition is matched, the session data action is performed.
   All string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details.
6. In the **Threshold** field, use one of the threshold parameters to specify the minimum number of occurrences required to create a correlation session and an associated key if required. The associated key cannot be an IPv4 or IPv6 meta type.
   - `u_count(associated_key)` = the count of unique values of the specified key
   - `sum(associated_key)` = the values of the specified key
   - `count` = number of sessions (no associated key is specified)
7. In the **Instance Key** field, select the target indicator to base the event upon. This can be a single key or a compound key (two primary keys, separated by a comma).

8. In the **Time Window**, set the duration during which the threshold must be reached to create a correlation session.
9. To save the rule and add it to the grid, click **OK**.
   The rule is added at the end of the grid or inserted where you specified in the context menu. The plus sign is displayed in the **Pending** column.
10. Check that the rule is in the correct execution sequence with other rules in the grid. If necessary, move the rule.
11. To apply the updated rule set to the service, click **Apply**.
    NetWitness saves a snapshot of the currently applied rules, then applies the updated set to the Decoder or Log Decoder.

Configure Network Rules

Network rules are applied at the packet level on a Decoder and are made up of rule sets from Layer 2, Layer 3, and Layer 4. Multiple rules can be applied at the packet level to a Decoder. Network rules can apply to multiple network layers (for example, when a network rule filters out specific ports for a specific IP address). Network rules do not apply to Log Decoders, they apply only to Network Decoders.

You can create and manage network rules in the Services Config view > Network Rules tab.

Note: Because network rules are applied on the packet level, you must specify both the source and destination meta keys in the rule condition. This is because the packet flow can occur on both sides at the packet level, while the directions are still not determined. For example, if you want to filter traffic on port 553 and port 55553, the condition should be written as follows:
`port.src=553 || port.dst=553 || port.src=55553 || port.dst=55553`
You must also specify both the source and destination meta keys in the conditions for `ip.src` and `ip.dst`. Specifying only source or destination in the condition will not work as expected.

Supported Meta Keys in Network Rule Conditions

The following table describes the meta keys that NetWitness supports for use in network rule conditions.

Meta Key	Description
`eth.addr`	Ethernet source or destination address. Commonly known as the MAC address.
`eth.dst`	Destination Ethernet address. This is the same as the Ethernet address field except that it selects only packets where the destination address matches the selected value (s).
`eth.src`	Same as Ethernet destination except that it focuses on the source address.
`eth.type`	Ethernet frame type.
`hdlc.type`	Frame type of the HDLC frame.
`alias.ip`	IPv4 source or destination address in standard form. IP addresses can be entered in CIDR notation for subnets.
`ip.dst`	Destination IPv4 address in standard form. IP addresses can be entered in CIDR notation for subnets.
`ip.proto`	IPv4 protocol field.
`ip.src`	Source IPv4 address in standard form. IP addresses can be entered in CIDR notation for subnets.
`alias.ipv6`	IPv6 source or destination address in hex format. Generally IPv6 addresses are written as eight groups of four hex digits, thus expressing the entire 128 bit address length. Supports notation to represent multiple blocks of 0000 in an address. Does not support CIDR notation.

Meta Key	Description
ipv6.dst	Destination IPv6 address in hex format.
ip.proto	IPv6 protocol field. This maps to the Next Header field in the IPv6 header and uses the same values as the IPv4 protocol field.
ipv6.src	Source IPv6 address in hex format.
port.dst	Destination TCP port.
tcp.port	TCP source or destination port.
port.src	Source TCP port.
port.dst	Destination UDP port.
udp.port	UDP source or destination port.
port.src	Source UDP port.

The following are sample network rules.

To truncate all SSL from the source port, create a rule as follows:

- Rule Name: Truncate SSL
- Condition: `port.src=443`
- Rule Action: Truncate

To filter subnet traffic, create a rule as follows:

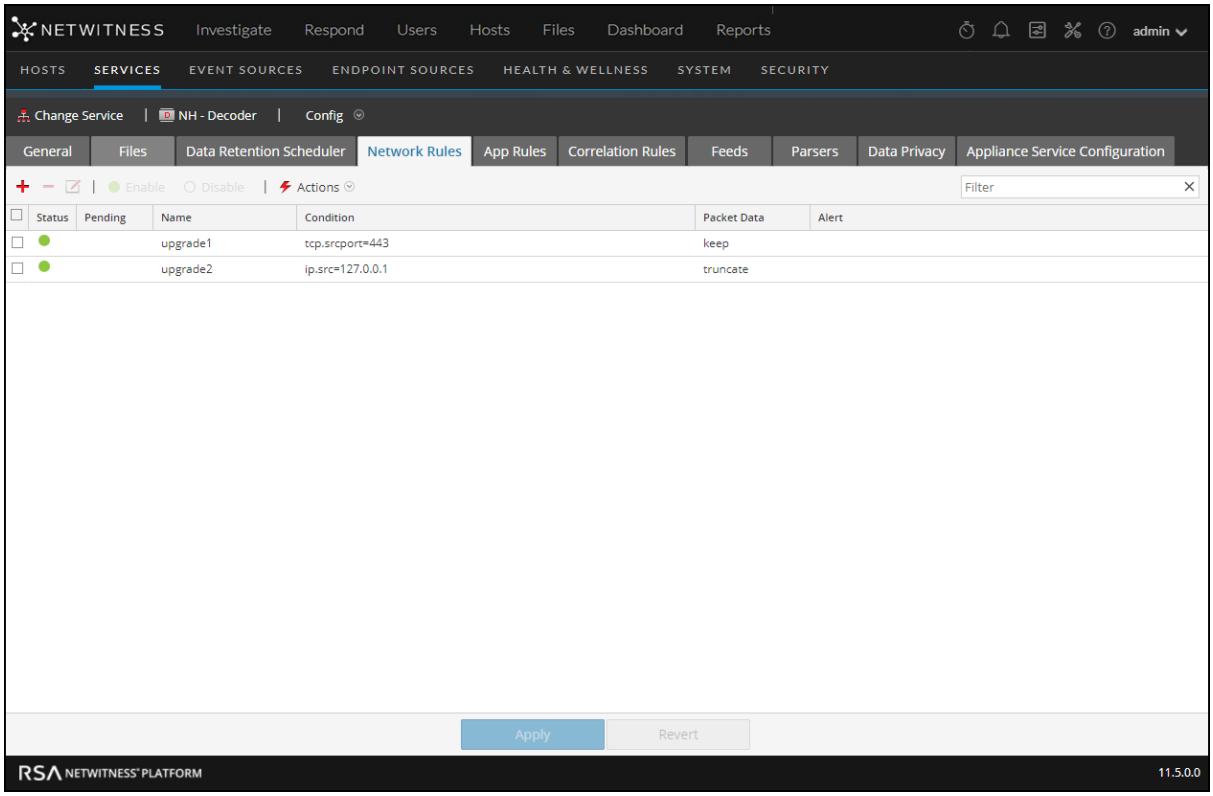
- Rule Name: Subnet Filter
- Condition: `alias.ip=192.168.2.0/24`
- Rule Action: Filter

Meta entities, which provide a way to work with several meta keys at the same time, can be used in application rules, but are not supported in network rules as the metadata available are too limited. For more information on meta entities, see the *Core Database Tuning Guide*.

To add or edit a network rule:

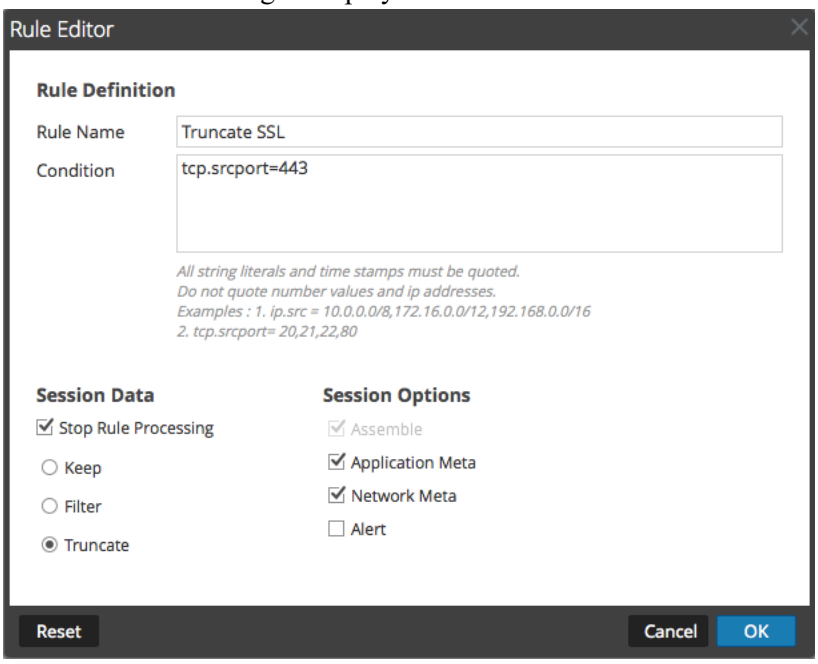
1. Go to  (Admin) > **Services**, select a Decoder service, and  **> View > Config**.
   The Services Config view for the selected service is displayed.

2. Select the **Network Rules tab**.
   The Network Rules tab is displayed.



3. In the **Network Rules** tab, do one of the following:
   - If adding a new rule, click .

- If editing a rule, select the rule from the rules list and click .
  The Rule Editor dialog is displayed.



4. In the **Rule Name** field, provide a name for the rule. For example, for a rule that truncates all SSL from the source port, type **SSL Truncate**.
5. In the **Condition** field, build the rule condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta values from the window actions. As you build the rule definition, NetWitness displays syntax errors and warnings. For example, to truncate all SSL from the source port, **`tcp.srcport=443`**.
   All string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details. Supported Meta Keys in Network Rule Conditions describes the meta keys that NetWitness supports for use in network rule conditions.
6. If you want rule evaluation to end with this rule, select the **Stop Rule Processing** checkbox.
7. In the **Session Data** section, choose one of the following actions to apply when a matching packet is found:
   - **Keep**: The packet payload and associated metadata are saved when they match the rule.
   - **Filter**: The packet is not saved when it matches the rule.
   - **Truncate**: The packet payload is not saved when it matches the rule, but packet headers and associated metadata are retained.

8. In the **Session Options** section, select all options that apply of these four.
   - **Assemble**: The assembler assembles the packet chain when it matches the rule.
   - **Network Meta:** The packet generates network metadata when it matches the rule.
   - **Application Meta**: The packet generates application metadata when it matches the rule.
   - **Alert**: The packet generates a custom alert when metadata matches the rule.

9. To save the rule and add it to the rules list, click **OK**.
   The rule is added at the end of the list or inserted where you specified in the context menu.

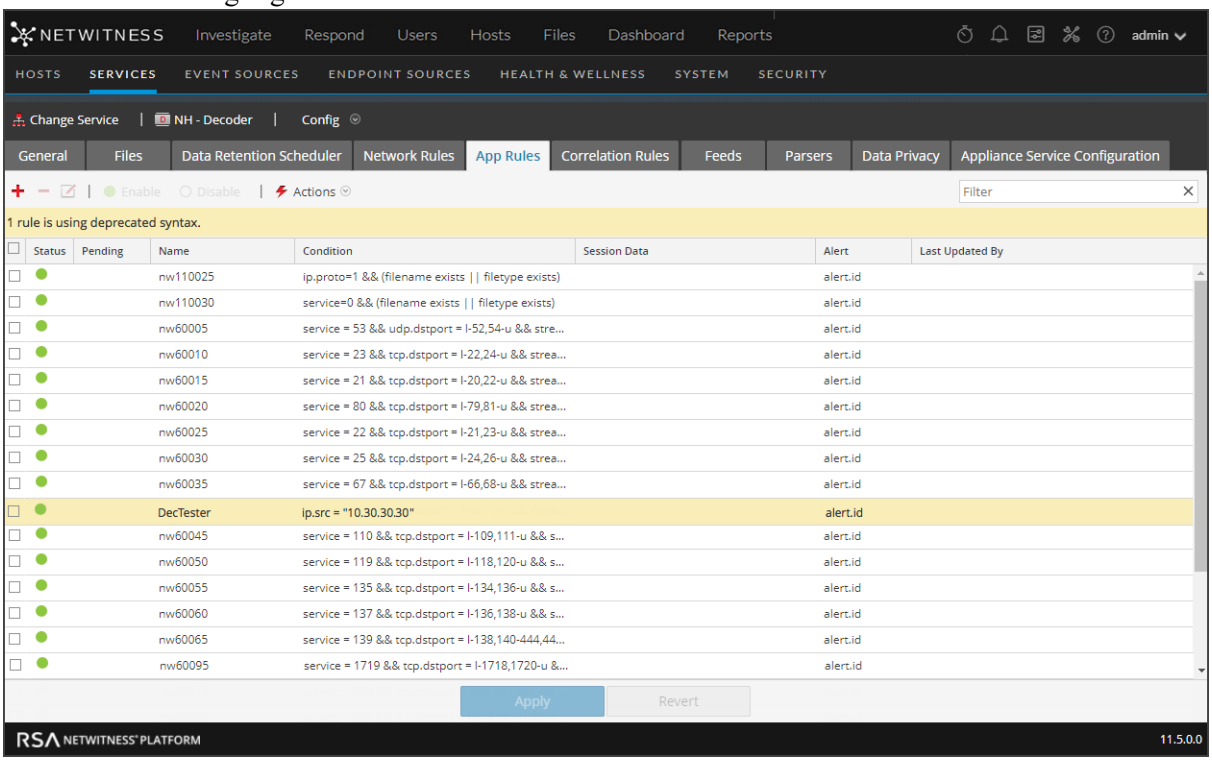
10. Check that the rule is in the correct execution sequence with other rules in the list. If necessary, move the rule.
11. To apply the updated rule set to the Decoder, click **Apply**.
    NetWitness saves a snapshot of the currently applied rules, then applies the updated set to the Decoder and removes the pending indicator from the rules that were pending.

Fix Rules with Invalid Syntax

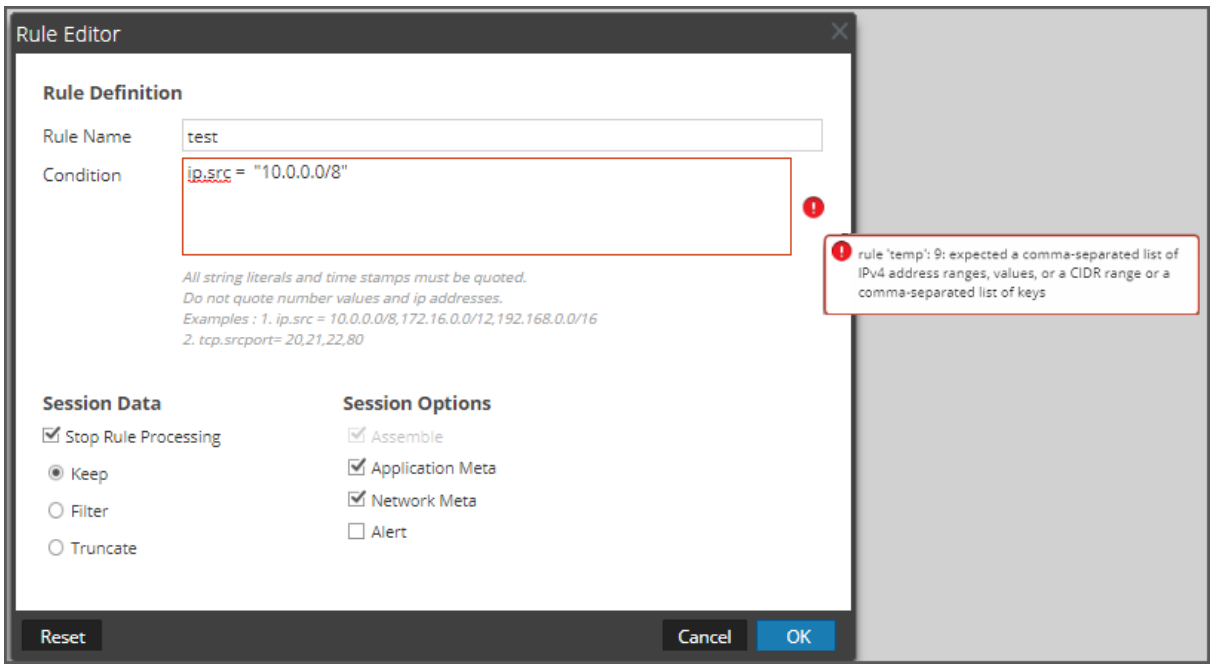
After an update to NetWitness 11.x, the user interface highlights any rules with invalid syntax. The Rule Editor provides additional tooltips. After you fix the rules, the highlights disappear. Configure Decoder Rules provides guidelines that all queries and rule conditions in NetWitness must follow.

To correct rules with invalid syntax:

1. Go to (Admin) > **Services**.
2. In the **Services** view, select a Decoder and **> View > Config**.
3. In the **Services Config** view, select one of the Rules tabs: Network Rules, App Rules, or Correlation Rules.
   The Rules tab for the selected rule type shows the number of rules using invalid syntax and the invalid rules are highlighted.



4. Select an invalid rule and click .
   The Rules Editor shows additional information for the invalid rule.



5. In the **Condition** field, correct the rule syntax.
   All string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details.
   For example, if the invalid rule condition is `ip.src="10.30.30.30"`, correct the syntax by removing the quotes: `ip.src=10.30.30.30`
6. Do one of the following:
   - To correct the rule individually, click **Save**.
     The corrected rule is applied independently to the Decoder. The corrected rule appears on the Rules tab without highlights.
   - To correct the rule and apply the rule to the Decoder later with other rules, click **OK**.
     The corrected rule appears on the Rules tab without highlights. The rule is not applied to the Decoder.

Decoder Commands for Managing Rules

In the NetWitness Core database, the Rules tree holds the main functionality related to managing rules for all Core services that have rules: Concentrators, Decoders, Log Decoders, and Archivers. Although you can manage rules in the NetWitness user interface, advanced users may prefer to manage rules using a command line to add, merge, replace, delete, and validate rules on a service. This section provides a brief overview of the commands and their usage. These are the available commands:

- `add` - Adds a single rule at the specified position.
- `clear` - Deletes all existing rules in the current node on the service. For example, using the command in `/decoder/config/rules/application` node deletes all existing application rules on the Decoder.
- `delete` - Deletes one or more rules at a specified position and count.
- `merge` - Merges a pushed rule set with an existing rule set. Existing rules that match the incoming rules (by name or rule) are replaced; otherwise, rules are inserted by the position indicated as described in merge Command.
- `replace` - Deletes all existing rules and replaces them with the incoming rule set.
- `validate` - Validates the syntax of a rule, but does not validate the meta keys.

add Command

The `add` command adds the rule to the existing rule set. Formatting is important because the API uses double quotes in the rule language and also uses double quotes as parameters to all NetWitness APIs. Therefore, you must escape any double quotes in the rule itself by preceding it with a backslash (\) character. This is the syntax of the command:

```
add rule=<string> name=<string> alert=<string, optional> atPos=<uint32, optional>
```

- `rule` is the rule to add. Be sure to place double quotes around any rules with white space and to escape any double quotes that are part of the rule with a backslash.
- `name` is the name of the rule.
- `alert` is the alert for the rule (if any).
- `atPos` is the position at which the rule should be added (1 based). Zero is the top of the list and any number larger than the current size of the list is appended to the list.

This is an example of command to add a rule using NwConsole

```
send /decoder/config/rules/application add rule="ip.src exists" order=1 alert=ioc name=testrule
```

For example, take the following rule:

```
alias.host = "myPC" && country.src="china","russian federation"
```

To add this as a rule, you would need to send the parameters as follows:

```
rule="alias.host = \"myPC\" && country.src=\"china\",\"russian federation\"" name=myRule filter
```

Notice how all the double quotes had to be escaped inside the rule parameter. A simple trick to make this more readable is to use single quotes inside the rule. Single and double quotes are interchangeable in the rule and query language, but not in parameters for the API (only double quotes are supported there). Therefore, this is more readable:

```
rule="alias.host = 'myPC' && country.src='china','russian federation'"
name=myRule filter
```

merge Command

The `merge` command is used to merge an incoming list of rules with the existing rules on the service. This is how it works:

- It finds existing rules that match via the name OR via a matching rule, updates the existing rule name, and keeps the same position.
- It inserts new rules into the rule list based on the NUMBER position. If the number is zero, it goes to the top of the list.
- It processes the rules in the order received so if you have two rules numbered zero, the second rule is processed after the first and claims the top spot. All existing rules are pushed down two places. Any numbers higher than existing rule positions are appended after the last existing rule and numbered in sequence.
- Any non-numbered rule is appended after the last existing rule and numbered in sequence.

This is the syntax of the merge command:

```
merge --file-data=<string> --file-format<string>
```

- `file-data` is the full path and name of the rules file to merge.
- `file-format` is the format of the rules file. Valid values are `params-list`, `string, params`, `binary`, and `params-binary`.

Methods of Sending a List of Rules to a Service

There are two ways to send a list of rules. You can send them as a `.nwr` (NetWitness Rule) file or as a numbered set of parameters, each number indicates the position to insert the rule at as well as the encoded rule. If you want to see the current list of rules on a service, you need to run the `ls` command on the rule category (for instance, application rules on a Decoder are found in `/decoder/config/rules/application`).

This is an example of commands to list the existing rules using NwConsole:

```
login <hostname>:50004 <username> <password>
cd /decoder/config/rules/application
ls
```

This is another example to list existing rules in NwConsole:

```
send /decoder/config/rules/application ls
```

This is an example of the command to point to network rules in the RESTful port, which supports a basic `admin HTML` app.

```
http[s]://<decoder>:50104/decoder/config/rules/network
```

Send a NetWitness Rule File

Let's start with an example `nwr` file, each rule must be on a separate line:

```
rule="ip.src=192.168.0.1" name=first keep
rule="ip.src=192.168.1.1" name=second alert=ioc
rule="ip.src=192.168.2.1" name=third filter
```

To push and merge rules using NwConsole, use the following commands:

```
login <hostname>:50004 <username> <password>
send /decoder/config/rules/application merge --file-data=/root/App_
Rules.nwr --file-format=params-list
```

To replace the existing rules with the rules in the file, instead of using the `merge` command, use the `replace` command.

```
send /decoder/config/rules/application replace --file-
data=<pathname> --file-format=params-list
```

To merge the rules in an `nwr` file using the RESTful port, you can use a `curl` command that pushes the rules:

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-
stream" --data-binary @<pathname> -X POST
"http://<hostname>:50104/decoder/config/rules/application?msg=merge"
```

The examples are pushing application rules. To push network rules, send the rules to `/decoder/config/rules/network`. For correlation rules, send the rules to `/decoder/config/rules/correlation`.

Send Numbered Parameters

The other way to send a list of rules is to send them as numbered parameters. The difficulty with this method is remembering to escape the quotes within each numbered rule. Though it is only a problem if you are trying to do it by hand. For instance, to send the same rules above as parameters via NwConsole, use the following command:

```
send /decoder/config/rules/application merge
1="rule=\"ip.src=192.168.0.1\" name=first keep"
2="rule=\"ip.src=192.168.1.1\" name=second alert=ioc"
3="rule=\"ip.src=192.168.2.1\" name=third filter"
```

This command is hard to read because you have to escape the inner quotes with a backslash (\). Otherwise, these two commands accomplish the same thing. Merging or adding three rules in positions 1, 2 and 3. If you think the above was hard to read, this is what the equivalent `curl` command looks like:

```
curl -u "<username>:<password>"
"http://<hostname>:50104/decoder/config/rules/application?msg=merge&1=rule%3D
%22ip.src%3D192.168.0.1%22%20name%3Dfirst%20keep&2=rule%3D%22ip.src%3D192.168
.1.1%22%20name%3Dsecond%20alert%3Dioc&3=rule%3D%22ip.src%3D192.168.2.1%22%20n
ame%3Dthird%20filter"
```

For more details on how to escape double quotes inside parameters, see add Command.

Ordering Rules When Pushing

Pushed rules are ordered in one of two ways. When passing as parameters, the number of each parameter determines the insertion order. If it is not actually a number, `merge` checks for an `order` parameter within the rule itself and uses that value if found.

Note: Using `order` is the only way to set the order with a `.nwr` file. If neither a number nor an `order` parameter is found, there are no guarantees of the insertion order.

Example

A Decoder has the following application rules installed; notice the numbering is ALWAYS consecutive and starts at 1:

```
0001 : rule="ip.src = 192.168.0.1 || ip.dst = 192.168.0.1 || alias.host = 'My-PC'" name=first keep
0002 : rule="ip.src=192.168.1.1" name=second alert=ioc
0003 : rule="ip.src=192.168.2.1" name=third filter
```

And you want to merge the following four rules:

```
rule="ip.src=192.168.3.1" name=third keep
rule="ip.dst=192.168.4.1" name=NewRule filter order=0
rule="alias.host = 'pc1','pc2'" name=filterTheseNames filter order=append
rule="service=80,443" name=web filter order=3
```

Use any method to push your rules and this is what you end up with:

```
0001 : rule="ip.dst=192.168.4.1" name=NewRule filter order=1
0002 : rule="ip.src = 192.168.0.1 || ip.dst = 192.168.0.1 || alias.host = 'My-PC'" name=first keep order=2
0003 : rule="service=80,443" name=web filter order=3
0004 : rule="ip.src=192.168.1.1" name=second alert=ioc order=4
0005 : rule="ip.src=192.168.3.1" name=third keep order=5
0006 : rule="alias.host = 'pc1','pc2'" name=filterTheseNames filter order=6
```

Are there any surprises here? This is how each rule was processed.

`1. rule="ip.src=192.168.3.1" name=third keep`

This rule had the same name as an existing rule on the Decoder (third). So the rule updated the existing rule, `changing _filter_ to _keep_`.

`2. rule="ip.dst=192.168.4.1" name=NewRule filter order=0`

This rule is new and had `order=0` in it, which means insert at the very top.

`3. rule="alias.host = 'pc1','pc2'" name=filterTheseNames filter order=append`

This rule had a non-number `append` for `order`, therefore, it went to the end of the list. You can accomplish the same thing by giving a very large number, like `999999`.

`4. rule="service=80,443" name=web filter order=3`

This rule is last but has `order=3`, therefore, if it does not match an existing rule by name or the text of the rule itself, it should be placed in position 3. And there it is, the third rule in the list. Any rules that follow were pushed further down.

replace Command

The `replace` command removes all existing rules and replaces them with the incoming rule list. Refer to merge Command for details on how to format the incoming rule list and how ordering works.

This is an example of the `replace` command using a NetWitness Rule File :

```
send /decoder/config/rules/application replace --file-data=/root/Decoder-AppRules.nwr --file-format=string
```

This is an example of the `replace` command using Numbered Parameters :

```
send /decoder/config/rules/application replace 1="rule=\"ip.src exists\" name=\"test rule\" order=1 alert=ioc"
```

clear Command

The `clear` command removes all existing rules on the service. This is an example of the command:

```
send /decoder/config/rules/application clear
```

delete Command

The `delete` command deletes one or more rules on the service.

```
delete atPos <uint32> count <uint32, optional>
```

- `atPos` deletes the rule at the given position. Rules are numbered starting with 1 and go in sequential order.
- `count` deletes one or more rules starting `atPos`. This is an optional parameter defining the number of rules to delete starting `atPos`. The default value is 1.

This example of the command deletes four rules beginning at position 0003:

```
send /decoder/config/rules/application delete atPos=0003 count=4
```

validate Command

The `validate` command takes the provided rule and verifies that it parses correctly. Keep in mind that this command cannot verify whether language keys and entities are valid.

```
validate rule <string>
```

`rule` - is the name of the rule to validate. Make sure to place double quotes around any rules with white space.

Configure Parsers and Feeds

Parsers and feeds are responsible for analyzing the packets and logs when captured or imported in Decoders. Most commonly, they are used for static metadata extraction and service identification. The flexible definition allows custom extension of the core defined services to provide extra service type identification and metadata extraction. This is important due to the volume of custom applications that are used on networks.

See the following sections for details about configuring parsers and feeds.

- Configure Parsers
- Configure Feeds

Configure Parsers

NetWitness has a set of native parsers that are defined by the system, and also provides the option to add additional parsers. Each parser is configurable in the Services Config View - General Tab. The Parser Configuration panel provides a way to enable or disable parsers to use on Decoders in addition to limiting the metadata that the parser creates.

There are also several types of custom configurable parsers:

- GeoIP2 – This parser associates IP addresses with geographical locations. For new installations and upgrades, the GeoIP2 parser is enabled by default. For more information on these parsers, see GeoIP2 Parsers.
- Search – This parser is user-configured to generate metadata by scanning for pre-defined keywords and regular expressions.
- FLEXPARSE (deprecated) – This is a generic parser definition language for extending the existing application protocol support of the Decoder. By default this parser is disabled (see Enable or Disable Lua and Flex Parsing Systems).
- Lua – This parser is defined using the Lua scripting language for extending the existing application protocol support of the Decoder.
- Log – This application parser supports the Log Decoder and is configured to generate metadata by scanning log files.
- Snort® – This parser supports the payload detection capabilities of Snort IDS rules. Snort rules and configuration are added to the `parsers/snort` directory for Investigation and Decoder (see Decoder Snort Detection).

In the Services Config view > Parsers tab, you can view deployed parsers on a Decoder, upload parsers, and delete deployed parsers. The user interface includes an Indicator if the parser originated from Live Services, installed through NetWitness, or uploaded manually. Parsers can be added and removed while a Decoder is running without affecting capture.

Note: To pass options to parsers, you must first give the name of the parser and then the options to be passed in this format:
`<ParserName>="<ParserOptions>"<Whitespace><ParserName2>="<Parser2Options>"`
Each `ParserName=Value` option must be separated by whitespace. Normally, the `Value` must have double quotes around it. The `Value` itself can sometimes list multiple `Option=Value` pairs, each separated by whitespace, and if those values have whitespace, they must be in escaped double quotes. To escape a quote, place a backslash before it: `\"`.
This is an example of defining options for Parser1, Parser2, and Parser3:

```
Parser1="Option1=\"Option1 Value With Space\" Option2=Option2ValueNoSpace"
Parser2="Option1=Value" Parser3="op1=val1 op2=val2 op3=\"another value\""
```

In addition, you can download parsers using NetWitness Live Services.

Upload and Delete Custom Parsers

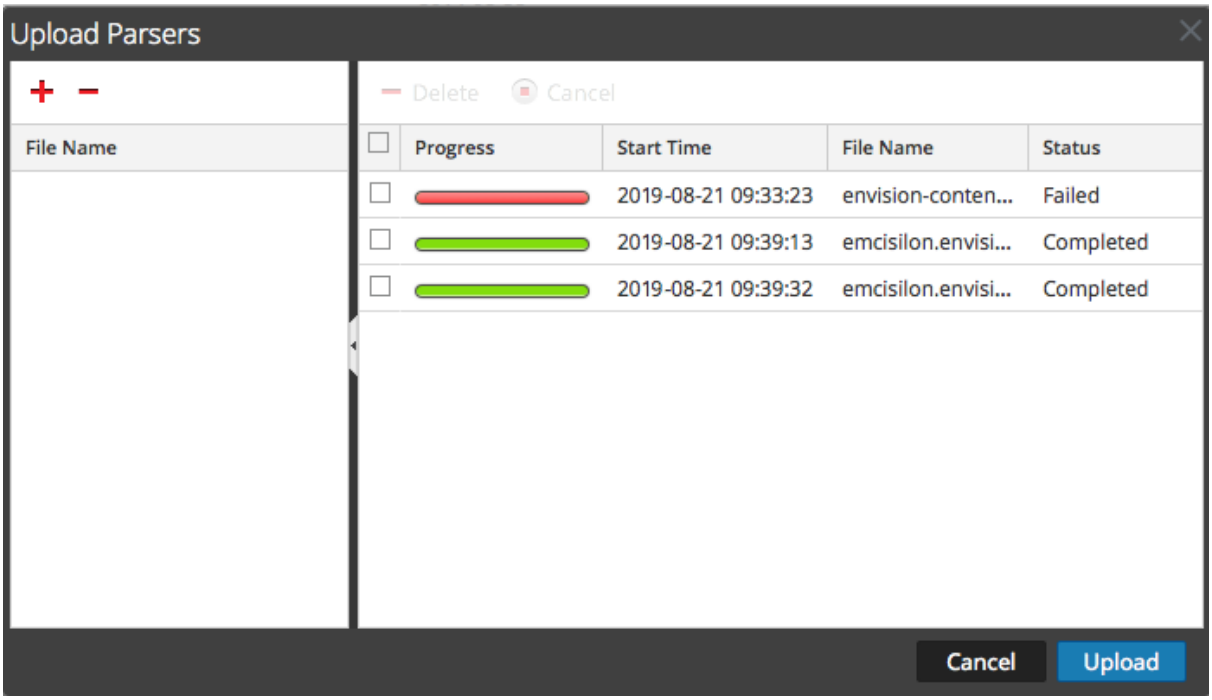
RSA NetWitness has the ability to upload parsers from your local system and delete these parsers.

Upload Parsers to a Decoder or Log Decoder

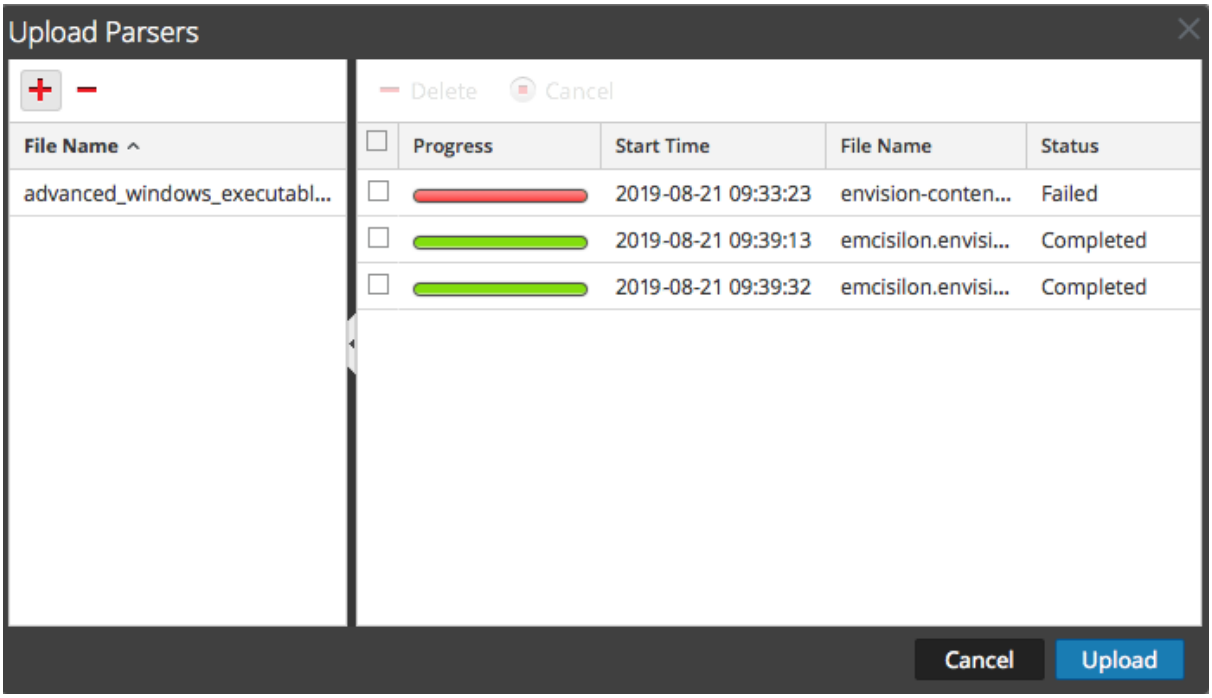
The Upload option in the Service Config view > Parsers tab displays the Upload Parsers dialog, in which you can manage the uploading of parsers to a Decoder or Log Decoder. In the File list, you prepare a list of parsers for uploading. You can add files from a directory structure, and delete files from the list if you decide that you don't want to upload a particular file. When the list is ready, clicking Upload starts the upload process.

1. Go to  (Admin) > **Services**, select a service, and click  **> View > Config.**
   The Config view for the selected service is displayed.
2. Click the **Parsers** tab.

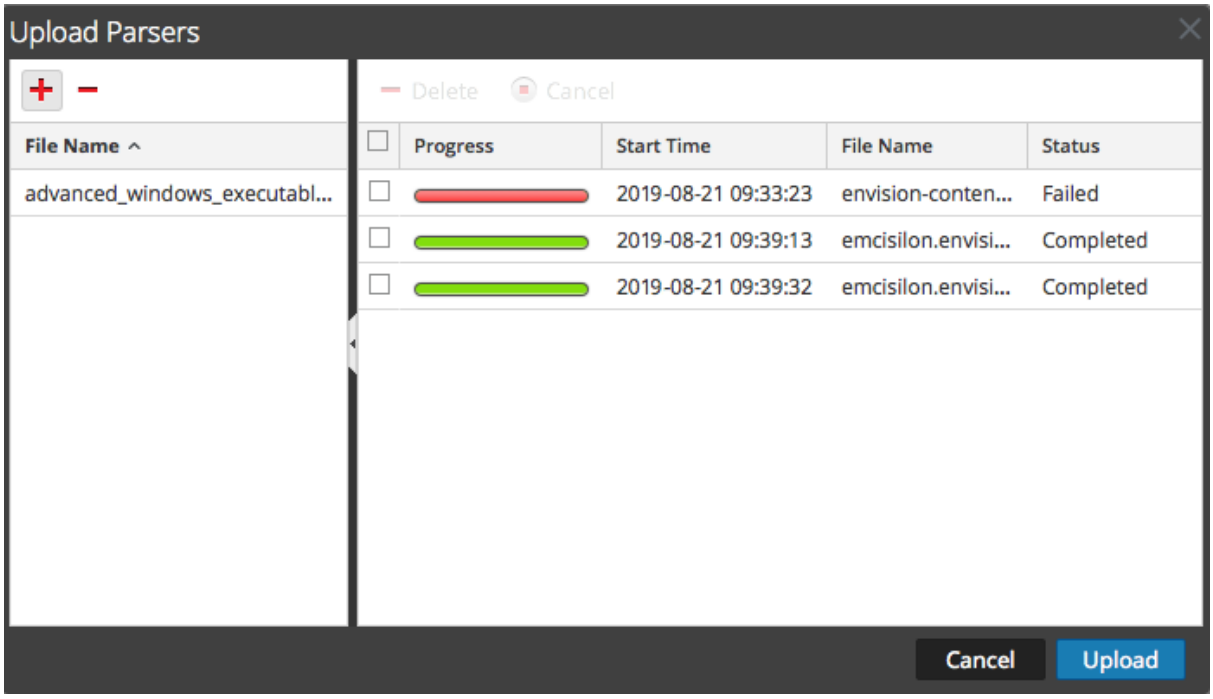
3. Click Upload.
   The Upload Parsers dialog is displayed.



4. Click .
   A file selection dialog is displayed.
5. Select the **.flex**, **.parser**, and **.lua** files to be updated, and click **Open**.
   The dialog closes, and the selected files are displayed in the File list.



6. Click **Upload**.
   The Upload Job grid shows the progress of the upload jobs with each job representing a file being uploaded.



7. Use any of the Upload grid tools to manage the upload of selected jobs: pause and resume, cancel, and delete.
   Once a job is complete, it is deployed on the Decoder and listed with the deployed parsers in Parsers tab.

Manage Upload Jobs

You can use any of the Upload grid tools to manage the upload of selected jobs: pause, resume, cancel, and delete.

- To cancel uploading a set of parsers while the upload is in queue or progress, click Cancel.
- To pause uploading a set of parsers, if the upload is not yet complete, click Pause.
- To resume uploading a set of parsers after a pause, click Resume.
- To delete an upload job, click .

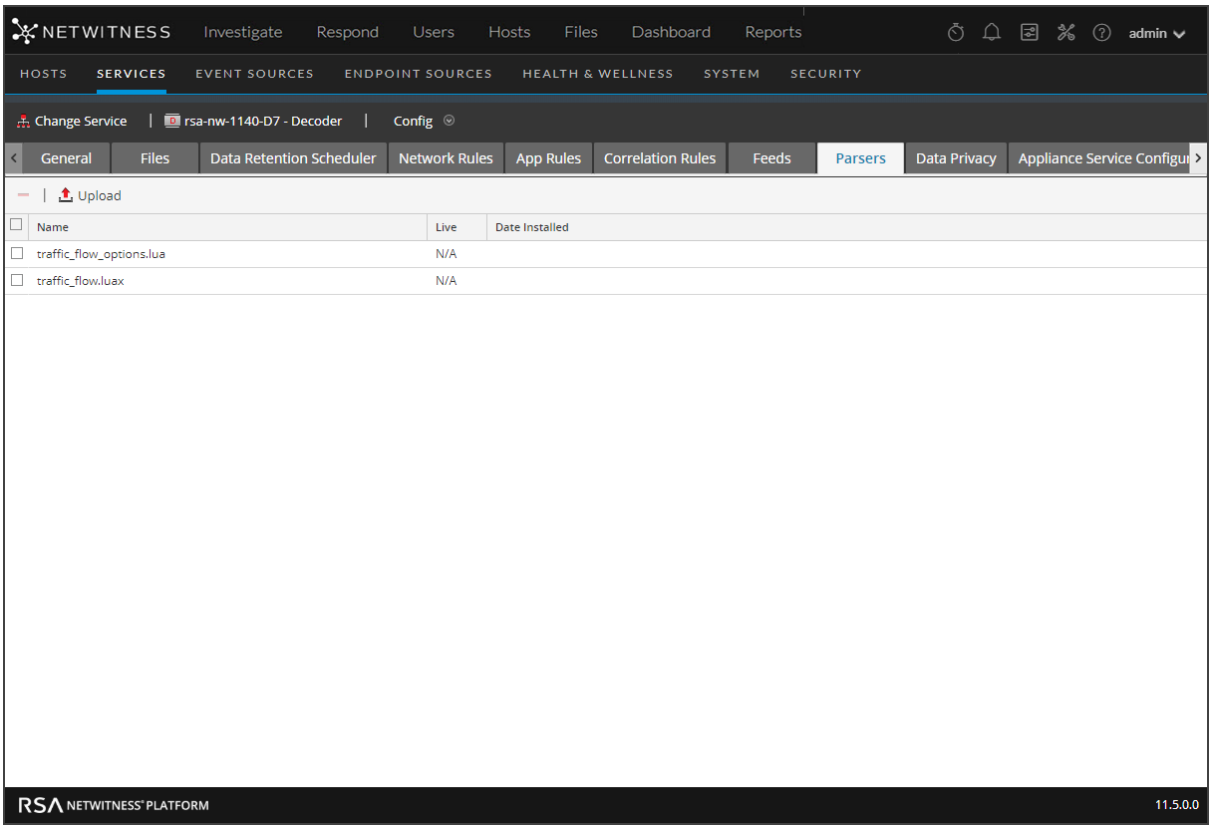
Delete Deployed Parsers

The Delete option in the Service Config view > Parsers tab provides a way to delete deployed parsers from a Decoder or Log Decoder. Parsers can be added and removed while a Decoder is running without affecting capture.

Note: Unless otherwise stated, any reference to Decoders applies to Log Decoders as well.

To delete a parser from a Decoder:

1. Go to (Admin) > **Services**, select a Decoder, and > **View** > **Config.**
   The Services Config view for the selected service is displayed.
2. Click the **Parsers** tab.



3. In the **Parsers** tab, select one or more parsers to delete.
4. Click .
   A dialog requests confirmation that you want to delete the parsers.
5. If you want to delete the parsers, click **Yes**.
   The parsers are removed from the Decoder immediately.

Enable and Configure the Entropy Parser

Beginning with NetWitness 11.0, the administrator can configure a Decoder to use a NetWitness native parser, known as the Entropy parser. When the Entropy parser is enabled, analysts have visibility into channels that are trying to blend in with other traffic, but do not follow normal protocol behavior. This helps to identify channels that do not conform to the normal environment traffic baseline, and may be worthy of investigation.

The parser creates meta keys, based on statistics collected by the native NetWitness parser, that help to identify behavior of any channel that is getting lots of network traffic. When the parser is first enabled, the analyst needs to become familiar with overall behavior for the different channels seen in a captured session to understand the frequency of bytes and the normal client and server payload. Once the normal behavior is known, analysts can use the meta keys to find behavior that does not match the expected.

By default, the Entropy parser generates 10 additional meta keys that do not add significantly to the load on a Decoder, and are useful for this specialized case. The parser is disabled by default.

Enable indexing if you have interest in exploring interesting sessions based on payload byte analysis of the packets. By default, to make indexing easier, the normal `Float32` value for `entropy.req` and `entropy.res` is multiplied by 10k and stored in a `UInt16` (thus giving four digits of precision, 0 to 10,000).

However, if you define the `entropy.*` fields in the Decoder language to be `Float32`, the Decoder will store it as a float with a range of 0.0 to 1.0. Take care to change the language everywhere if you decide to keep it as a `Float32`.

RSA does not recommend indexing as a `Float32` because of the high unique counts due to minute changes in precision.

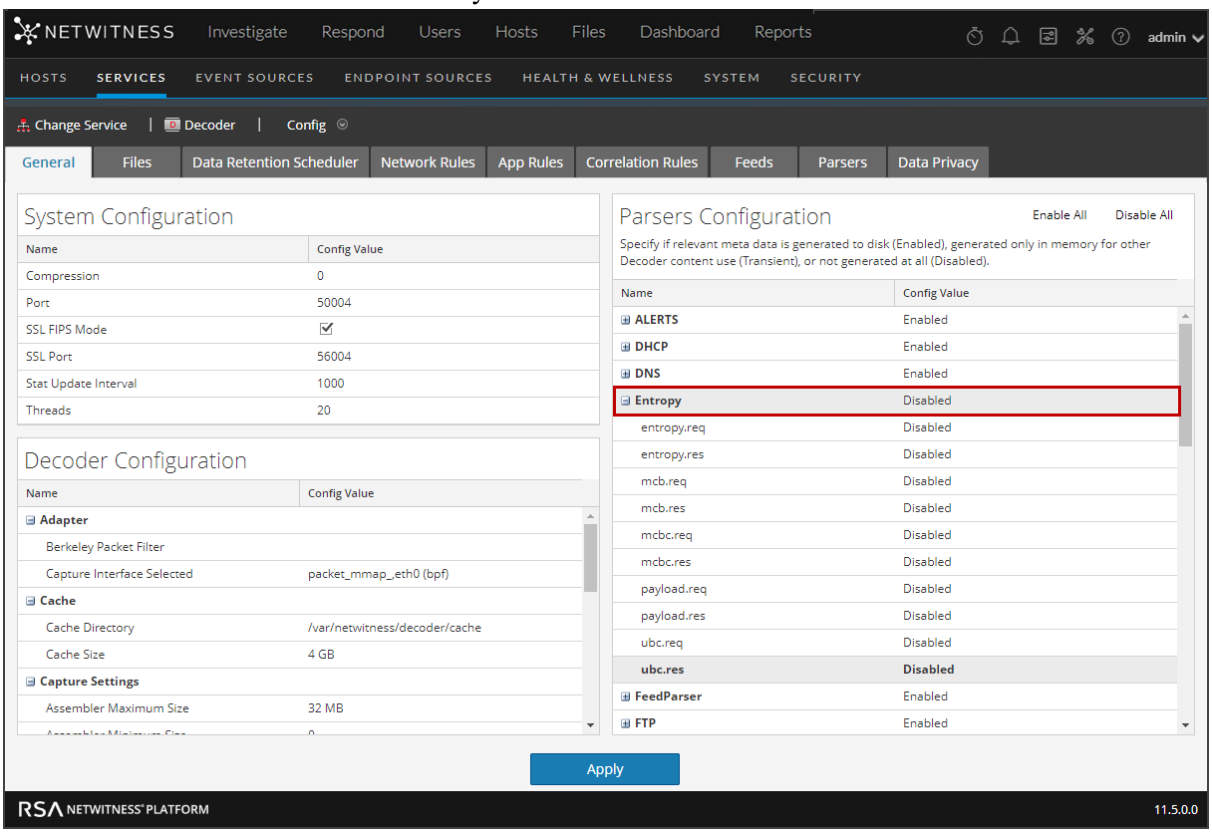
These are the new meta keys generated by the Entropy parser by default:

- `entropy.req` and `entropy.res`: These meta keys capture entropy using the Shannon entropy equation, which has a floating point value as a result. The floating point value of 0 to 1.000 is multiplied by 10000 and written in NetWitness as `UInt 16`, an unsigned integer of 0 through 10000. .
- `mcb.req` and `mcb.res`: The most common byte is simply which byte for each side (0 thru 255) was seen the most.
- `mcbc.req` and `mcbc.res`: The most common byte count is the number of times the most common byte (above) was seen in the session streams.
- `ubc.req` and `ubc.res`: - Unique byte count is the number of unique bytes seen in each stream. 256 would mean all byte values of 0 thru 255 were seen at least once.

To enable and configure the Entropy parser on a Decoder:

1. Go to  (Admin) > **Services**, select a Decoder service and click  **View > Config**.
   The Services Config view for the selected Decoder is displayed.
2. The Entropy parser is disabled by default. Click the drop-down list for **Entropy** in the **Config Value** column, and select **Enabled**. If you want to disable some of the meta keys, click the drop-down list

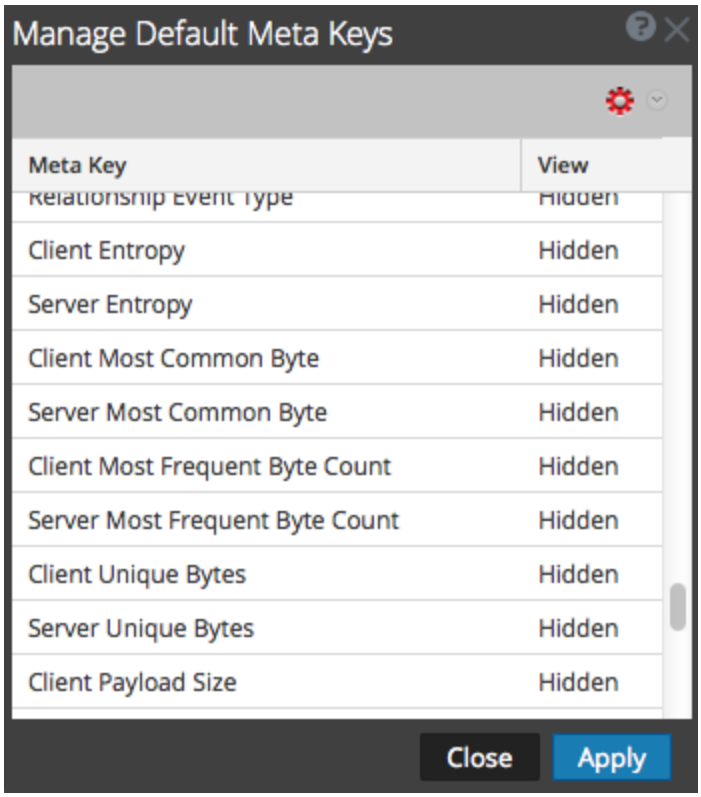
and select **Disabled** next to the meta key.



3. Click **Apply**.
   The Entropy parser is enabled and begins creating the new meta keys as configured in the Concentrator custom index file.

4. In the Service Config view select the Concentrator that is aggregating traffic from this Decoder. Select **View > Files** and open the Custom Index file for the Concentrator. Look for the Entropy parser meta keys to see if they are included and uncommented.
   By default the keys are commented out and therefore not enabled. To enable that part of the language the administrator needs to copy that part of index file into the `index-concentrator-custom.xml` and uncomment the `key description` line for each meta key. An example of the custom index file with the Entropy parser keys and instructions is shown below in Entropy Parser Configuration in the Concentrator Custom Index File.

5. With the Entropy meta keys enabled, they are available to analysts in Investigate, but hidden by default. To make the meta keys visible in the Investigate Values view, edit the default meta keys in the Default Meta Keys dialog so that they are open instead of hidden. You can manage these meta key the same way you manage other meta keys.



Entropy Parser Configuration in the Concentrator Custom Index File

The following is an excerpt of the Concentrator Index file lines that the administrator must copy to the custom index file. The comments provide guidance on configuring the parser.

```
<!-- This section is commented out because it's only used by the Entropy
parser which is disabled by default. To enable this part of the language, copy
to index-concentrator-custom.xml and uncomment the keys. HOWEVER, take note
that depending on how the Entropy parser is configured, the entropy.req and
entropy.res format might be a Float32 instead of a UInt16. So make sure to
change to the correct type if necessary.-->

<!-- Entropy parser meta - enable indexing if you have interest in exploring
this for interesting sessions based on payload byte analysis of the packets.
By default, to make indexing easier, the normal Float32 value for entropy.req
and entropy.res is multiplied by 10k and stored in a UInt16 (thus giving 4
digits of precision, 0 to 10,000). However, if you define the entropy.* fields
in the Decoder language to be Float32, it will store it as a float with a
range of 0.0 to 1.0. Take care to change the language everywhere if you decide
to keep it as a Float32. We do not recommend indexing as a Float32 because of
the high unique counts due to minute changes in precision. -->

<!--

<key description="Entropy Request (Client)" format="UInt16" level="IndexNone"
name="entropy.req" valueMax="10001"/>

<key description="Entropy Response (Server)" format="UInt16" level="IndexNone"
name="entropy.res" valueMax="10001"/>

-->
```

```
<!-- The most common byte is simply which byte for each side (0 thru 255) was
seen the most -->

<!--

<key description="Most Common Byte Request" format="UInt8" level="IndexNone"
name="mcb.req"/>

<key description="Most Common Byte Response" format="UInt8" level="IndexNone"
name="mcb.res"/>

-->

<!-- The most common byte count is the number of times the most common byte
(above) was seen in the session streams -->

<!--

<key description="Most Common Byte Count Request" format="UInt32"
level="IndexNone" name="mcbc.req" valueMax="500000"/>

<key description="Most Common Byte Count Response" format="UInt32"
level="IndexNone" name="mcbc.res" valueMax="500000"/>

-->

<!-- Unique byte count is the number of unique bytes seen in each stream. 256
would mean all byte values of 0 thru 255 were seen at least once -->

<!--

<key description="Unique Byte Count Request" format="UInt16" level="IndexNone"
name="ubc.req"/>

<key description="Unique Byte Count Response" format="UInt16"
level="IndexNone" name="ubc.res"/>

-->

<!-- The payload size metrics are the payload sizes of each session side at
the time of parsing. However, in order to keep indexing from having high
unique counts (bad for performance), the two payload size metas below are
indexed in buckets. -->

<!--

<key description="Payload Size Request" format="UInt32" level="IndexNone"
bucket="true" name="payload.req" valueMax="500000"/>

<key description="Payload Size Response" format="UInt32" level="IndexNone"
bucket="true" name="payload.res" valueMax="500000"/>

-->
```

Flex Parsers

There are two kinds of Flex parsers:

- **Service identification based solely on port.** These are parsers that use only the source or destination ports to identify the session application type (service). These are the most basic and easiest to define.
- **Service identification based on a found token(s).** These parsers use tokens to identify the service type. This is also an easy way to expand which service types are identified. These are important when identifying non-internet standard applications. These parsers require that the protocol has a definable token that can uniquely identify the service type.

Five common parser operations are:

- Match Port and Identify Immediately
- Match Port and Delay Identification
- Match Token and Identify Immediately
- Match Multiple Tokens
- Match Token and Create Metadata

Detailed language information and samples are provided in this topic. This topic describes the XML schema used to define a FlexParse file. The SML node, attribute, and values referenced in descriptive text are **bold**. The root node of every file must be the **parsers** node. Under that node there can be any number of parser nodes. Each parser node defines a single parser. A parser node can have an optional **declaration** node and any number of **match** nodes.

Topics

- Arithmetic Functions
- Common Parser Operations
- General Functions
- Logging Functions
- Nodes
- Payload Functions
- Regex
- String Functions

Arithmetic Functions

This topic defines language for the flex parser arithmetic functions.

This topic defines language for the flex parser arithmetic functions. All numbers are 64-bit unsigned values and subject to both underflow and overflow, depending on the operation.

Language Definition

The following table provides language definitions.

Node Name	Attribute Name	Description
`and`		Performs bitwise AND between two numbers.
	`name`	Variable to AND result into.
	`value`	Number to AND into result.
`or`		Performs bitwise OR between two numbers.
	`name`	Variable to OR result into.
	`value`	Number to OR into result.
`increment`		Performs ADDITION of two numbers.
	`name`	Variable containing the initial value AND to receive ADDITION results.
	`value`	Number to ADD to initial value.
`decrement`		Performs SUBTRACTION of two numbers.
	`name`	Variable containing initial value AND to receive SUBTRACTION results.
	`value`	Number to SUBTRACT from initial value.
`divide`		Performs DIVISION of two numbers.
	`name`	Variable containing the initial value AND to receive DIVISION results.
	`value`	Number by which to divide the initial value. Division by zero generates an error and stops any further processing of the current session by this parser.
`modulo`		Performs MODULO of two numbers.
	`name`	Variable containing the initial value AND to receive MODULO results.

Node Name	Attribute Name	Description
	value	Number by which to divide the initial value. Division by zero generates an error and stops any further processing of the current session by this parser.
multiply		Performs MULTIPLICATION of two numbers.
	name	Variable containing the initial value AND to receive MULTIPLICATION results.
	value	Number by which to MULTIPLY the initial value.
shiftleft		Performs a binary shift left.
	name	Variable containing the initial value AND to receive shift results.
	value	Number of bits to shift by.
shiftright		Performs a binary shift right.
	name	Variable containing the inital value AND to receive shift results.
	value	Number of bits to shift by.

Common Parser Operations

This topic provides some examples of common parser operations.

This topic includes five common parser operations.

Match Port and Identify Immediately

```
<?xml version="1.0" encoding="utf-8"?>
<parsers
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
     xsi:noNamespaceSchemaLocation="parsers.xsd">
          <parser name="CustApp" desc="Acme Custom App" service="45324">
               <declaration>
                    <port name="port" value="45324" />
               <declaration>
               </match name="port">
                    <identify />
               </match>
          </parser>
</parsers>
```

Match Port and Delay Identification

```
<?xml version="1.0" encoding="utf-8"?>
<parsers
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="parsers.xsd">
     <parser name="MSRPC" desc="Microsoft RPC protocol" service=135">
          <declaration>
               <port name="port" value="135" />
               <number name="state" scope="session" />
               <session name="end" value="end" />
          </declaration>
          <match name="port">
               <assign name="state" value="1" />
          </match>
          <match name="end">
               <if name="state" equal="1" />
                    <identify />
               </if>
          </match>
     </parser>
```

```
</parsers>
```

Match Token and Identify Immediately

```
<?xml version="1.0" encoding="utf-8?>
<parsers
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="parsers.xsd">
     <parser name="RDP" desc="Remote Desktop Protocol" service="3389">
          <declaration>
               <token name="signature" value="Cookie: mstshash=" />
          </declaration>
          <match name="signature">
               <identify />
          </match>
     </parser>
</parsers>
```

Match Multiple Tokens

```
<?xml version="1.0" encoding="utf-8"?>
<parsers
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
     xsi:noNamespaceSchemaLocation="parsers.xsd">
     <parser name="MyServiceMultiToken" desc="Multiple Tokens" service="333">
          <declaration>
               <number name="state" scope="stream" />
               <token name="user" value="USER " />
               <token name="pass" value="PASS " />
               <session name="session" value="end" />
          </declaration>
          <match name="user">
               <or name="state" value="1" />
          </match>
          <match name="pass">
               <or name="state" value="2" />
          </match>
          <match name="session">
               <if name="state" equal="3">
                    <identify />
               </if>
          </match>
```

```
     </parser>
</parsers>
```

Match Token and Create Metadata

```
<?xml version="1.0" encoding="utf-8"?>
<parsers xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="parsers.xsd">
     <parser name="SHELL" desc="Command Shell Identification">
          <declaration>
               <token name="cmd.exe" value=" (C) Copyright 1985-2001 Microsoft
               Corp" options="linestart" />
               <meta name="client" key="client" format="Text" />
          </declaration>
          <match name="cmd.exe"
               <register  name="client" value="MS Command Shell" />
          </match>
     </parser>
</parsers>
```

General Functions

This topic defines language for the flex parser general functions.

General Functions Language Definition

Node Name	Attribute Name	Description
apptype		Gets the currently defined service type for the current session.
	name	A number variable to receive the current service type.
identify		Marks the session with the parser's service type if the service type has not already been identified.
assign		Assigns a value to a variable.
	name	The unique identifier assigned to the item in the declaration section.
	value	Optional. If specified, the action defined in the match is only applied when the declaration matches the given value.
getmeta		Retrieves the value of meta that generated a callback. This function will return empty results (0, zero length string) if called when there was no meta callback.
	name	The variable to receive the value of the meta key that generated the callback.
gettoken		Returns the current matched token.
	name	A string variable to receive the current matched token. If there is no current token, the variable is assigned an empty string.
end		This terminates the execution of the current **match** section.
if		Compares two values. If the comparison is true, executes any sub-actions. Comparisons can be **number** or **string** types, as long as both values are the same type.
	name	The unique variable identifier assigned to the item in the **declaration** section.
	equal notequal less lessequal greater greaterequal and or	The operation value to compare. If true, any sub-actions are executed.
register		Adds metadata to the session.

Node Name	Attribute Name	Description
	name	The unique identifier of a meta variable to be created, as defined in the **declaration** section.
	value	The value of the metadata to be created.
while		Compares two values and executes any sub-actions if the comparison is true. Comparisons can be **number** or **string** types, as long as both values are the same type.
	name	The unique variable identifier assigned to the item in the declaration section.
	equal notequal less lessequal greater greaterequal and or	Specifies the operation value to compare. If true, any sub-action is executed. The **and** and **or** attributes signify bitwise operations and can only be applied to **number** variables.
call		Executes the specified **match** element. This can be any match element defined in the same flex parser regardless of how it was declared.
	value	The name of the match element, or a string variable containing the name of a match element. • If the match element name is specified, the parser will not load if the named matched element doesn't exist. • If a string variable is specified, the **call** element will execute any child elements that it may have if the string value resolves to a match element after executing the named match element. • If no **match** element can be found matching the string value, no action is taken.

Logging Functions

This topic defines language for the flex parser logging functions.

Logging functions provide a means for a flex parser to write to the system log. Logging functions can be extremely useful when creating a new flex parser, but should be kept to an absolute minimum when a flex parser is deployed to a production system.

Language Definition

Node Name	Attribute Name	Description
`failure`		Logs a message to the system log with the log level **Failure.**
	`value`	A string to include as the log message.
`warning`		Logs a message to the system log with the log level **Warning.**
	`value`	A string to include as the log message.
`info`		Logs a message to the system log with the log level **Info**.
	`value`	A string to include as the log message.
`debug`		Logs a message to the system log with the log level **Debug.**
	`value`	A string to include as the log message.

Nodes

This topic defines language for the flex parser nodes.

Nodes Language Definition

Node Name	Attribute Name	Description
`parsers`		The root node in each definition file.
	`xmins:xsi`	Defines the namespace to use for the schema inclusion. This attribute is not required; however, language definition is not possible without it. This node must have the following value: http://www.w3.org/2001/XMLSchema-instance
	`xsi:noNamespaceSchemaLocation`	Defines the XSD schema validation file used to validate the language definition. This attribute is not required; however, language definition is not possible without it. This node must have the following value: parsers.xsd
`parser`		The node that defines a single parser definition. This node must be directly under the `parsers` node. There can be more than one per file.
	`name`	The name that uniquely identifies the parser. This name should be short and succinct. This is used by the system to allow enabling and disabling. It should contain only the letters [a-z] and [A-Z].
	`desc`	Provides a friendly description of what the parser does.
	`service`	The unique number assigned to the session when identified.
`declaration`		Delineates the definition. Each of these definitions can have an associated `match` entry.

Node Name	Attribute Name	Description
token		Specifies a definition for identifying a token somewhere in the session protocol. This defines a `match` callback when the specified tokens are encountered in a session payload. The `read` position is set to the byte immediately following the matched token.
	name	This is a unique identifier for the declaration.
	value	This is the exact token value to be identified.
	options	Options specify that the token should start on a new line or at an end of a line (`linestart` or `linestop`).
meta-callback		Registers a callback for the flex parser whenever meta of a specific format is created. This can be further qualified to generate callbacks only for sessions that have been identified as a specific apptype (for example, 80 for HTTP).
	name	Name of the match element to be executed when a callback occurs. (String)
	key	Name of the meta key that generates callbacks. (String)
	format	The data type of the meta key that will generate the meta.
	apptype	The meta callback is only generated if the session being parsed has been identified with the specified apptype. (Unsigned Integer, Optional)
number		Defines a numeric variable that can be referenced elsewhere within the parser definition. All numeric values are 64-bit unsigned values.
	name	This is a unique identifier for the declaration.

Node Name	Attribute Name	Description
	`scope` (optional)	Specifies when to reset the variable. This can either be for each side of a two-sided session or only after a new session is detected. The possible values are **`global`**, **`constant`**, **`stream`**, and `session` (default).
`string`		Defines a numeric variable that can be referenced elsewhere within the parser definition.
	`name`	This is a unique identifier for the declaration.
	`scope` (optional)	Specifies when to reset the variable. This can either be for each side of a two-sided session or only after a new session is detected. The possible values are **`global`**, **`constant`**, **`stream`**, and `session` (default).
port		Defines a **match** callback when a session is encountered using the specified port. The **read** position is set to the first byte of the first stream (client) in the session.
	`name`	This is a unique identifier for the declaration.
	`value`	This is the port number to identify.
`session`		Defines a `match` callback for session begin/end events. These events only occur if a token for the parser is encountered in the session.
	`name`	This is a unique identifier for the declaration.
	`value`	Specifies that processing takes place at the beginning of a new session or at the end of a session (`begin` or `end`).
`stream`		Defines a `match` callback for stream begin/end events. These events only occur if a token for the parser is encountered in the stream.
	`name`	This is a unique identifier for the declaration

Node Name	Attribute Name	Description
	`value`	Specifies that processing takes place at the beginning or at the end of a stream (`begin` or `end`).
`function`		Defines a `match` section that can be used as a generic function. No callbacks are associated with this declaration.
	`name`	This is a unique identifier for the declaration.
`meta`		Defines the type of data that the parser will create.
	`key`	Specifies the key name. The key needs to be 1-16 bytes in size.
	`format`	Specifies the variant type (for example, **`Text`**, **`IPv4`**, **`UInt32`**). Refer to the SDK documentation for a full list.
`pattern`		Defines a regular expression variable for use by the `regex` function
	`name`	This is a unique identifier for the declaration.
	`scope` (optional)	Specifies when to reset the variable. This can be for each side of a two-sided session or only after a new session is detected. Possible values are **`global`**, **`constant`**, **`stream`**, and `session` (default).
	`value` (optional)	Specifies a regular expression to assign to the pattern variable. This attribute is only valid when the **`scope attribute`** is set to `constant`.
`match`		The possible entries for taking an action once a match criterion has been found for a declaration. These nodes can be nested to provide deeper logic. There are several categories of execution elements (functions) that can appear as children of a match element: • General • Arithmetic • String • Payload

Payload Functions

This topic defines language for the flex parser payload functions.

These functions operate on a `read` position, set at the beginning of a `match` element.

Language Definition

Node Name	Attribute Name	Description
`find`		Searches the stream payload starting at the read position for a provided string value. If the value is found, the offset from the read position is returned. Any child elements will then execute. If not found, any child elements will not execute.
	`name`	A `number` variable to receive the offset from the `read` position where the match begins.
	`value`	A string to find.
	`length` (optional)	A limit to the length of the payload to be searched. If a limit is not provided, the remainder of the payload is searched. It is recommended to always use the smallest value possible here in order to reduce the effect on performance.
`install-decoder`		To enable tokens to match on payload data that may be fragmented or otherwise encoded. A scan decoder can be installed to preprocess a section of the payload before it is scanned for tokens. An example would be an HTTP response that uses the chunked transfer encoding with gzip content encoding. By parsing the HTTP header, the necessary type, offset, and length parameters can all be set, after which the HTTP response payload would appear to the token scanning as if neither encoding had been applied. However, this incurs significant overhead.
	`type`	The type of decoder to install. Valid options are: gzip, deflate, chunked, chunked-gzip, chunked-deflate.
	`offset`	Offset from the current read position to begin decoding.
	`length`	The maximum payload length to decode.
`isdecoding`		Tests whether an installed decoder is currently active. If so, any children of this function will execute. This function has no parameters.

Node Name	Attribute Name	Description
`move`		Moves the `read` position forward in the current stream by a specified number of bytes. If there is sufficient data in the stream, the `read` position is updated and any child elements will then execute. If not found, the `read` position remains unchanged and any child elements will not execute.
	`value`	The number of bytes to move the `read` position.
	`direction` (optional)	The direction to move the current read position. Can be `forward` (default) or **`reverse.`**
`packetid`		Returns the id of the packet for the current read position. It is possible for the result to be 0, which indicates that the packet id could not be determined.
	`name`	A number variable to receive the current packet id.
`payload-position`		Returns the current read position. This is a zero based index into the stream payload.
	`name`	A number variable to receive the current read position.
`read`		Reads a specified number of bytes starting at the `read` position into a variable. If there is sufficient data in the stream, the `read` position is updated, the data read assigned, and any child elements will then execute. If not found, the `read` position remains unchanged and any child elements will not execute.
	`name`	The name of a `string` or `number` variable to receive stream data. If a `number` variable is provided, the bytes read are interpreted as a single unsigned numeric value.
	`length`	The number of bytes to read from a stream.
	`endianess` (optional)	The byte ordering to use when reading into a number variable. Can be `big` (default) or `little`. The attribute is invalid when reading into a `string` variable.

Regex

This topic defines language for the flex parser regex node.

Regex searches the stream payload starting at the `read` position for matches to a provided regular expression. If matches are found, the offset from the `read` position and, optionally the matched string, is returned. Any child elements execute. If no matches are found, child elements do not execute.

Language Definition

Attribute Name	Description
`name`	A `number` variable to receive the offset from the `read` position where the match begins.
`value`	A regular expression to find.
`length` (optional)	A limit to the length of the payload to be searched. If a limit is not provided, the remainder of the payload is searched. It is recommended to always use the smallest value possible here in order to reduce the effect on performance.
`found` (optional)	The name of a `string` variable to receive a matched string.

String Functions

This topic provides language definitions for the flex parser string functions.

String Functions Language Definition

Node Name	Attribute Name	Description
`append`		Attaches a number or string to the end of a `string` variable.
	`name`	The unique identifier of a string variable to which the specified value is to be attached.
	`value`	A number or string to attach.
`find`		Searches a string for a provided string value. If it is found, the position is returned and any child elements will execute. Otherwise, child elements will not execute.
	`name`	A `number` variable to receive the zero-based position, where the provided value string was found in the **in** string.
	`value`	A string to find.
	`in`	A string to search.
	`length` (optional)	A limit to the length of the **in** string to be searched. If a limit is not provided, all of **in** will be searched.
`length`		Assigns the length of a string to a `number` variable.
	`name`	A `number` variable to receive the length of the specified string.
	`value`	A string value whose length is to be determined.
`regex`		Searches a string for matches to the provided regular expression. If a match is found, the position and, optionally, the matching string is returned. Any child elements will then execute. If not found, any child elements will not execute. Regular expression operations can adversely affect system performance.
	`name`	A number variable to receive the zero-based position, where the provided regular expression matched in the in string.
	`value`	A regular expression to be searched for.

Node Name	Attribute Name	Description
	in	A string to search.
	length (optional)	A limit to the length of the in string to be searched. If a limit is not provided, all of **in** will be searched.
	found (optional)	The name of a string variable to receive the matched string.
substring		At least one of the optional attributes from and length must be specified.
	name	The unique identifier of a string variable to receive the extracted value.
	value	A string value from which to extract a substring.
	from (optional)	The zero-based position from which to begin the substring. If not specified, it defaults to zero.
	length (optional)	The number of characters to extract. If not specified, it defaults to the remaining length of the string.
tolower		Converts a string to all lowercase letters.
	name	The name of a string variable to process.
toupper		Converts a string to all uppercase letters.
	name	The name of a string variable to process.
urldecode		Decodes a string containing url-encoded characters.
	name	A string variable to receive the decoded string.
	value	A url-encoded string to decode.
base64decode		Decodes a base-64 encoded string.
	name	A string variable to receive the decoded string.
	value	A url-encoded string to decode.
uudecode		Decode a uuencoded string.
	name	A string variable to receive the decoded string.
	value	A uuencoded string. The header and trailing lines should not be included.
quotedprintabledecode		Decode a Quoted-printable encoded string.
	name	A string variable to receive the decoded string.
	value	A quoted-printable encoded string.

Node Name	Attribute Name	Description
`convert-ebcdic`		Convert an EBCDIC string to its ASCII equivalent.
	`name`	A string variable to receive the decoded string.
	`value`	A url-encoded string to decode.

GeoIP2 Parsers

This topic describes the GeoIP2 parser for Decoders. This parser converts IP addresses into geographic locations, such as the country name and city where the IP address is typically found.

Note: In version 11.3 and later, the native GeoIP2 parser replaces the GeoIP parser (which has been permanently removed). The GeoIP2 parser provides the same basic functionality as the GeoIP parser as well as many enhancements. For example, it converts IP addresses into geographic locations, provides the latest Maxmind GeoIP package, and supports IPv6 addresses as well as IPv4.

Available in NetWitness version 11.2 or later, the GeoIP2 Parser is enabled by default for upgrades and new installations. The GeoIP2 parser provides the latest Maxmind GeoIP package and supports IPv6 addresses as well as IPv4.

To edit the GeoIP2 parser configuration:

1. Go to  (Admin) > **Services**.
2. In the **Administration services** view, select a Log Decoder or a Decoder.
3. Click the settings icon () and select **View** > **Config**. In the Parsers Configuration panel, select **GeoIP2** to view and update configuration options.
4. Define the IP addresses to lookup. The GeoIP2 parser enables the following IP addresses by default: `ip.src`, `ip.dst`, `ipv6.src`, and `ipv6.dst`. You can update options by using `parsers.options` to remove or add new IP addresses. For example, you can edit `parsers.options` and pass a comma-separated list of IP addresses to use as follows:
   `GeoIP2="ipaddr=ip.src,ip.dst,ipv6.src,ipv6.dst,alias.ip"`
   This adds a new IP address to lookup called `ip.addr`. However, since `alias.ip` does not end in `.src` or `.dst`, the parser will elect to place the GeoIP2 metadata generated in meta keys without a `.src` or `.dst` suffix. So, you would see country, city, and so on, after the `alias.ip` metadata.
5. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Note: The list you pass for `alias.ip` replaces the default list. So, if you pass `alias.ip=ip.src`, it generates only GeoIP2 metadata for `ip.src`, and generates no metadata for other IP addresses.

Note: `parsers.options` is used for passing options to multiple parsers. So if you add GeoIP2 to it, you should not delete any other options being passed to other parsers (like Entropy).

The following table provides the full list of metadata that the GeoIP2 parser can potentially generate and indicates which metadata is or is not enabled by default:

Enabled by Default	Not Enabled
country, country.src, country.dst	latdec, latdec.src, latdec.dst
	longdec, longdec.src, longdec.dst
domain, domain.src, domain.dst	isp, isp.src, isp.dst
org, org.src, org.dst	city, city.src, city.dst

You can enable the other metadata using the standard parser configurations.

Note: By disabling some metadata by default, the GeoIP2 parser does not work the same as the GeoIP parser (which did not, by default, disable any metadata it generated). If you need any of the disabled metadata, you need to enable them (once only) for each Decoder, after upgrading to 11.2 or later. Keep in mind that the `isp` and `org` meta keys usually produce an equivalent value to `domain`.

Lua Parsers

There are a number of Lua parsers available from Live. See RSA Content for:

- A complete list of these parsers
- Their interdependencies
- The Flex parsers that are subsumed by each Lua parser

Five common parser operations are:

- Match Port and Identify Immediately
- Match Port and Delay Identification
- Match Token and Identify Immediately
- Match Multiple Tokens
- Match Token and Create Metadata

HTTP Parsers

The HTTP parser is a native parser that is used for Decoders to parse both requests and responses in HTTP messages. In version 11.4 and later, the HTTP parser provides a decompression option.

The decompression option mimics the decompression option in the Lua HTTP parser and is controlled by the `'decompression'` option for the HTTP parser. Parser options are set in the `/decoder/parsers/config/parsers.options` configuration node. To set an option on the HTTP parser, you append an `HTTP=""` clause to the `parsers.options` field so that HTTP can understand the `'decompression'` option.

For example, you could add `HTTP="decompression=true"` to the parser option list to enable decompression of all HTTP compressed bodies.

You can use the following values in the `'decompression'` field.

Value	Description	Example Entry in parsers.options
`true`	Decompress all bodies	`HTTP="decompress=true"`
`false`	Do not decompress. This is the default.	`HTTP="decompress=false"`
`1`	Decompress `application/*` content	`HTTP="decompress=1"`
`2`	Decompress `audio/*` content	`HTTP="decompress=2"`
`4`	Decompress `font/*` content	`HTTP="decompress=4"`
`8`	Decompress `image/*` content	`HTTP="decompress=8"`
`16`	Decompress `message/*` content	`HTTP="decompress=16"`
`32`	Decompress `model/*` content	`HTTP="decompress=32"`
`64`	Decompress `text/*` content	`HTTP="decompress=64"`
`128`	Decompress `video/*` content	`HTTP="decompress=128"`

The numeric values can be combined by addition to search for multiple types of content. For example, if you want to decompress application and text content, use `1 + 64 = 65`, which becomes `HTTP="decompress=65"`.

To set the decompression option:

1. Go to  (Admin) > **Services** and select a Decoder, and in the actions menu (), select **View** > **Explore**.
2. Expland **decoder** > **parser**s and select **config**.
3. In **parsers.options**, append `HTTP="decompress=<option from table>"`.

4. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Visibility into HTTP/2 Sessions

You can search for metadata items derived from headers in the HTTP/2 stream to gain visibility into HTTP/2 sessions.

To turn on header parsing for HTTP/2 sessions:

1. Go to **ADMIN** > **Services** and select a Decoder, and in the actions menu (), select **View** > **Explore**.
2. Expland **decoder** > **parser**s and select **config**.
3. In **parsers.options**, append `HTTP2="headers=true"`.
4. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Decoder Snort Detection

Snort is a free and open-source intrusion prevention system that uses a rule-based language to detect malicious network traffic. The NetWitness Platform Decoder offers compatibility with Snort detection rules, sometimes referred to as Snort signatures. NetWitness supports importing existing Snort rules into the Decoder, extending the known threats that can be detected. There are many sources for Snort rules, as it is an open source solution. In most cases, the behavior of Snort rules in the Decoder is the same as the rules in other network analysis tools. However, the Decoder does not have coverage for all Snort rule options. The topics in this section highlight NetWitness Platform Decoder-specific details of how to set up Snort parsers, how Snort rules are handled, and explain cases where Decoder might differ in behavior from Snort.

Using Snort Rules in Decoders

The Snort Detection engine is implemented as a built-in native parser in Decoders. The parser's identifier is `Snort`, and can be found in the Decoder Parsers Configuration page along with the other native parsers enabled by default. The Snort parser reads files from the directory `/etc/netwitness/ng/parsers/snort`. All files in this directory are read during the Decoder service restart. Files that end in the extension `.conf` are treated as configuration files. The format for configuration files matches the configuration file format for Snort, although not all configuration directives are supported. Further details are in Configuration Directives . Files that end in the extension `.rules` are treated as Snort rules files. A rules file can contain one or more rules, and the Snort directory can contain more than one rules file. The rule file format is the same as for Snort itself. The Snort rule files can be loaded into the Decoder by any of the following methods:

1. The Decoder service starts.
2. The `reload` message is sent to the `/decoder/parsers` folder. This allows for rules to be reloaded while capture is running.

However, note the following:

- Any rule that does not properly parse is ignored.
- Any valid Snort rule should successfully parse; however, there are rule options that are not supported by Decoders which are not fully parsed. For more information, see Snort Parser Capabilities .

Note: RSA recommends that the maximum number of Snort rules in use per Decoder is 1,000.

Initial Decoder Snort Configuration

Follow these steps to set up an initial configuration of Snort on a Decoder.

1. Edit the `snort.conf` file in the `/etc/netwitness/ng/parsers/snort` directory, and configure it according to the instructions in Configuration Directives , which describes the details of the variables inside the configuration file. If the variables `HOME_NET, EXTERNAL_NET` and `HTTP_PORTS` are not defined, they default to the value of `any`.
   A sample `snort.conf` file is shown below, but depending on where the Decoder is placed in the network, and which threats need to be captured, you might need to alter the configuration.
   For example, you can be more specific about the internal subnets by defining `ipvar HOME_NET [192.168.0.0/16, 10.0.0.0/8, 172.16.0.0/12]`. However, if the Decoder is capturing

traffic on a DMZ, those subnets might not be accurate.
Also, if `HOME_NET` has been defined, and the `EXTERNAL_NET` has also been narrowed down to `ipvar EXTERNAL_NET !$HOME_NET`, then you may miss lateral movement traffic (for example, `192.168/16` to `192.168/16`) within your environment. In general we suggest you align this configuration file with how you configure the `traffic_flow` configuration for the corresponding Traffic Flow Lua parser.

An example `snort.conf` file that is a good starting point:

```
#Configure the network addresses you are protecting
ipvar HOME_NET any
#Configure the network addresses external to your environment
ipvar EXTERNAL_NET any
#Configure any specific network ports
portvar HTTP_PORTS any
```

2. Generate or upload a Snort rules file into the same directory (`/etc/netwitness/ng/parsers/snort`). Read the Snort documentation at https://www.snort.org/#documents for details on how to construct rules. There are many free community-generated Snort rules available as well, that can be uploaded from https://www.snort.org/downloads/#rule-downloads. Several other sites publish Snort signatures as responses to new threats that are found. The basic rule structure and supported sections are mentioned in Rule Sections.

   Note: Only the Snort V2.x rules are supported. There are some rule options that are not supported, which are described in Snort Parser Capabilities .

3. Make sure the Snort configuration and rules files do not cause any errors when initialized by the Decoder by running the service restart command on the Decoder.
   The reload will check in the Snort rules without restarting, but does not log any messages about whether there are any issues with the rules or configuration.

   a. In the NetWitness user interface, go to  (Admin) > **Services** > **Decoder** >  > **View** > **System**.

   b. Select **Shutdown Service** and confirm. (The Decoder will automatically restart.)

   c. Go to > **Services** > **Decoder** >  > **View** > **Logs**.

   d. Search for Snort to validate that the rules files loaded.
      The log will indicate if the rules files are loaded completely (`full`), partially (`partial`), or if any rules railed to load (`failed`).
      The following are example log entries, using `mysnort.rules`, with two fully supported rules in the file. If any of the rules are marked partial or failure, examine the rules to determine why the load was not successful.

```
Loaded mysnort.rules, full 2, partial 0, failures 0
Loaded 2 snort rules, 2 small tokens, 0 with pcres, 0 partial, 0
unsupported
```

      When you search for Snort, you can find details on configuration variables and warning messages

if they are not defined. The following is an example log entry for a `snort.conf` file that is not fully populated with the minimal IP and port variables.

```
Undefined snort ip variables (will default to any): EXTERNAL_NET, HOME_
NET
Undefined snort port variables (will default to any): HTTPS_PORTS
```

e. Alternatively, you can perform the service restart and review of the logs from the Decoder console:

```
systemctl restart nwdecoder
tail -f /var/log/messages | grep -i snort
```

4. Import or capture network traffic that is known to trigger the Snort rules to validate they are functioning.
5. Validate that the Snort rules are getting triggered by the network traffic or by sample uploaded PCAPs. The Snort statistics can be used to track which rules have been evaluated, compared to generated metadata. For more information, see Rule Statistics . The default metadata generated by Snort rules are: `sig.id`, `sig.name`, `threat.category`, `threat.source`, and `risk.num`. The mapping of Snort rule information to NetWitness Platform metadata, and how to change it, is described in Snort Parser Output.

Snort Parser Output

When a Snort rule matches a session, it produces meta items. The meta items generated may change depending on the configuration of the Snort parser. Snort parser meta key usage has been updated with a new option for the Snort parser. The option, `Snort="udm=true"` (set to true by default), uses the Aligned Unified Data Model (UDM) key set. For information about UDM, see https://community.rsa.com/community/products/netwitness/rsa-content/udm. By default, the Aligned UDM key set is used.

Note: To pass options to parsers, you must first give the name of the parser and then the options to be passed in this format:
`<ParserName>="<ParserOptions>"<Whitespace><ParserName2>="<Parser2Options>"`
Each `ParserName=Value` option must be separated by whitespace. Normally, the `Value` must have double quotes around it. The `Value` itself can sometimes list multiple `Option=Value` pairs, each separated by whitespace, and if those values have whitespace, they must be in escaped double quotes. To escape a quote, place a backslash before it: `\"`.
This is an example of defining options for Parser1, Parser2, and Parser3:

```
Parser1="Option1=\"Option1 Value With Space\" Option2=Option2ValueNoSpace"
Parser2="Option1=Value" Parser3="op1=val1 op2=val2 op3=\"another value\""
```

The following keys are generated in the Aligned Unified Data Mode:

- `sig.id` with the value of the rule's `sid` field
- `threat.category` with the value of the rule's `classtype` field
- `sig.name` with the value of the rule's `msg` field
- `risk.num` with the value of the rule's priority

The following keys are generated in Legacy Key Mode:

- `alert.id` with the value of rule's `sid` field
- `threat.category` with the value of the rule's `classtype` field
- `risk.info`, `risk.suspicious`, or `risk.warning` depending on the priority of the rule. The value of the risk meta is the rule's `msg` field.

If you want to use the legacy key set which contains keys that are consistent with previous releases for Snort parser meta keys, follow these steps.

1. In the NetWitness User Interface, go to **ADMIN** > **Services**.
2. Select a Decoder and then click ⚙ > **View** > **Explore**.
3. In the left panel, select **decoder** > **parsers** > **config**.
4. In the right panel, in `parser.options`, add `Snort="udm=false"`.

Rule Statistics

The Snort parser maintains counters to monitor the activity of each rule. These counters can be retrieved using the running `snrtStat` on the Decoder's `/decoder/parsers` node, which can be accessed through the NwConsole utility or the REST interface. The statistics reported for each rule are:

- `evaluations`: This counts the number of times the rule was evaluated. A rule can be evaluated once for each stream in a session. Typically a rule will only match one side of a session, and will be evaluated, at most, once. The `content` strings determine if a rule will be evaluated in a session. A `content` pattern must match somewhere in the session in order for it to be evaluated.
- `hits`: This counts the number of times a rule matched and the corresponding alert meta item was created.

To review the Snort statistics when using the Decoder console:

1. Log in to the console locally or remotely
2. Type the command `NwConsole`
3. Log in to `localhost:50004 <admin> <netwitness>` (change login credentials accordingly)
4. Enter `/decoder/parsers snrtStat`

The output is a JSON structure listing the number of evaluations and hits performed, along with the options evaluated, for each Snort ID.

Configuration Directives

Decoder supports a limited number of configuration directives.

Variable Definitions

Decoder supports definitions of variables that hold the value of IP ranges using `ipvar` or Port ranges using `portvar`. For example, Snort rules usually make use of `HOME_NET, EXTERNAL_NET` and `HTTP_PORTS` variables. These have to be defined in your configuration file. If they are undefined, they default to the special value `any`.

ruletype

The definition of additional rule types is supported. However, only rules that have a base rule type of `alert` are supported. Rule type definitions follow the same definition format as Snort configuration.

Other Configuration Entries

config detection: debug

Adding this configuration line turns on debug messages that are generated as the rule is being processed. These debug messages are generally emitted only emitted when a rule matches a session by one of the content patterns, but then fails to match other options within the rule. This is useful for debugging complex filters in Snort rules, to determine why they may not be matching as expected.

config nopcre

Adding this configuration line turns off the PCRE (regular expression) capability.

config classification

Classification configuration influences how meta is generated when a rule matches a session. Meta output from Snort rule matches is described in Snort Parser Output.

Snort Parser Capabilities

The general format of a Snort rule is that it contains a rule header and rule options. The rule header consists of the rule action, protocol type, source-destination criteria (IP addresses, ports), and the direction of traffic. The rule options consist of the message (`msg`) which describe what the rule has detected, the references related to the threat or where the rule was generated, the classification of the rule, the unique signature (`rule`) identifier, and a long list of attributes (for example, `flow`, `content`, `pcre`) to define how the traffic is identified. The specific details on which header and rule options are supported are outlined in the remainder of this section.

Rule Sections

Section	Description
Header	The header conditions are evaluated when a rule receives the first token callback for a stream. The header is evaluated once per stream, and prevents any further consideration of a rule against a specific stream if the conditions are not met.
Actions	The specified action or a rule must be defined (either one of the native Snort actions, or defined in the configuration using the `ruletype` statement) for the rule to be considered valid. The Decoder only uses rules with alert actions.
Protocols	The Decoder supports the current Snort protocol keywords (`tcp`, `udp`, `icmp`, `ip`).
IP Addresses	The full language for defining IP addresses is supported, including lists, CIDR, and negation.
Port Numbers	The full language for defining port numbers is supported, including lists, ranges and negation.

Section	Description
Direction Operator	The directional operator supports the from-to (‘->’) and bidirectional (‘<>’) values. The to-from (‘<-’) value is invalid and causes the rule to fail to load.

Content

The Snort `content` statement makes up the bulk of the Snort pattern detection capability. For best performance, Decoder strives to map `content` patterns to the Token Parser used by most NetWitness parsers. Therefore, the `content` strings are the main mechanism to activate the rule matching engine on a session for a particular rule. If the rule does not have any content patterns, it is effectively unsupported. The Snort parser allows for very short strings patterns in content, but be aware that a rule that contains only very short content may have to be evaluated for every session, which will make the Snort Engine run much more slowly.

The Snort parser supports most forms of `content` statements, including the following modifiers. All modifiers effect the last `content` pattern declared in the rule.

Option	Description
`nocase`	Case-insensitive pattern matching is used.
`offset`	This option is applied to the distance of the token from the beginning of the packet. If the number of bytes between the offset point and the start of the token is **less** than this value, it is not a match.
`depth`	This option is applied to the distance of the token from the current offset position, or the beginning of the packet if no offset is set. If the number of bytes between the 'offset' point and the beginning of the token is **greater** than this value, it is not a match.
`distance`	This option is applied to the distance of the token from the current Detection Offset End (DOE) point. If the number of bytes between the current token and the DOE is **less** than this value, the token is not a match.
`within`	This option is applied to the distance of the token from the current Detection Offset End (DOE) point. If the number of bytes between the current token and the DOE is **greater** than this value, the token is not a match.

These content modifiers are supported if the Snort parser is used in conjuction with AppHTTP native parser:

Option	Description
`http_client_body`	The token must appear in the body of an HTTP request.
`http_cookie`	The token must appear in an HTTP cookie header.
`http_header`	The token must appear any where in in an HTTP request or response header.

Option	Description
`http_method`	The token must appear in the HTTP method field. The detection engine does not currently restrict the method tokens, such as `GET` or `POST`, to the HTTP method field, so these matches will work even if `AppHTTP` is turned off.
`http_uri`	The token must appear in the **normalized** HTTP header. Non-normalized header matches are not supported.

Further detail on the HTTP related modifiers is described below in HTTP Parsing.

Flow

The flow directive verifies that the rule is only applied to the client or server stream. It supports the following options.

Option	Description
`to_client`	Limits the rule to only matching on a stream that a Decoder has defined as Server.
`from_server`	Synonym for `to_client`.
`from_client`	Limits the rule to only matching on a stream that a Decoder has defined as Client.

Fast Pattern

The Snort parser uses `content` tagged with `fast_pattern` to indicate tokens that should always be used as tokens in the Decoder's token scanner. Decoder supports the `fast_pattern` and `fast_pattern:only` directives, but it does not support the `fast_pattern:offset,length` directive.

PCRE

The Snort parser implements `pcre` regex matching using the same libpcre library as other tools. It supports most of the flags documented for the `pcre` tag, with a few minor exceptions as described later in this document.

HTTP Parsing

The Snort parser relies on the native HTTP parser to notify it when HTTP is being parsed. The HTTP parser feeds information to the Snort engine to let it know which regions of the session should be used to evaluate the `http content` flags. The `pcre` implementation has option flags analogous to the `http content` options, and these are supported as well. If the native HTTP parser is disabled, then these features are also disabled in the Snort parser.

File Data

The Snort parser supports the `file_data` tag for HTTP response bodies only. This functionality relies on the native HTTP parser being enabled. If the Native HTTP parser has the decompression feature turned on, the Snort parser also uses decompressed responses to implement rule matching inside of `file_data` regions. You can enable the HTTP decompression feature by adding the token `HTTP="decompress=true"` to the configuration field `/decoder/config/parsers.options`.

The Decoder has a Javascript normalization parser that functions similarly to the `normalize_javascript` feature in Snort's HTTP preprocessor. It also affects data presented to Snort rules using `file_data`. To enable the Javascript normalization parser, add the token `JSNormalize="unescape=true"` to the configuration field `/decoder/config/parsers.options`.

The Snort parser emulates the `file_data` behavior in the same way as Snort. Directives that appear after `file_data` are implicitly evaluated whenever a content match is found inside of a `file_data` region. Because of this, `file_data` has side effects on other statements that appear after it in a rule.

Binary (byte) Directives

The Snort parser supports the `byte_test`, `byte_extract` and `byte_jump` directives.

The detection engine supports most options on these directives:

- The options for extracting the value, such as the offset, the little endian / big endian options, and relative options are supported. The relative flag causes the extraction of the byte value to occur relative to the DOE (Detection Offset End) pointer.
- `byte_jump` and `byte_extract` support the `align` and `multiplier` options to manipulate the extracted value.

The '`byte_extract` directive stores values in variable names that can be used to substitute a number into a few specific locations elsewhere in the rule. These locations are supported:

- The `value` field in a subsequent `byte_test`
- The value of a `within` statement on a subsequent content pattern
- The value of a `distance` statement on a subsequent content declaration.

The `byte_jump` option explicitly moves the DOE (Detection Offset End) pointer for use by subsequent directives, like `byte_test`, content marked with distance or within, or PCRE patterns marked as relative.

There are some less common byte operations that are not supported:

- The byte options do not support the `string` option for converting text into a binary value
- The byte options do not support the `bitmask` option.
- The byte options do not support the `dce` option.

Default option / fallback

When the Snort parser encounters rule options it does not yet support, it skips those options and assumes they did not prevent the rule from matching. In such cases it relies on the implementation of options it does support, such as the content or pcre statements, to determine if the rule matches. This behavior means that the Snort parser is more likely to fail in the `false-positive` direction rather than the `false-negative` direction when using features that Decoder does not yet support.

Packet Scanning/Scoping

Depending on how the rule is formed, it may be intended to match content entirely within single packets, or it may be intended to match data that spans multiple packets. The Snort parser attempts to handle both cases. Rule statements that are implicitly packet-scoped may be evaluated for every packet available in a session at parse time, and the Snort parser correctly scopes such rules so they only match if the elements of the rule match in a single packet. Conversely, if the rule uses a feature that implies reassembly and reconstruction, such as `file_data`, the Snort parser will allow the rule to be evaluated for the reconstructed portion of the stream that the rule references. This allows the Snort parser to emulate the various buffers that the Snort product provides.

Known Limitations for Snort Parser

Decoder Token Match Limits

The Snort Parser is activated by token scanning in Decoder, and therefore is subject to the same limitations on matches as all Decoder parsers. Configuration items such as `parse.bytes.min`, `parse.bytes.max`, timeouts, and size limits all apply to the Snort parser. The Decoder configuration guide describes these parameters in great detail. The Snort parser is implemented on top of reassembled sessions only so that features that require an assembled session can be supported.

Flow Bits

Flow bit directives, which start with the keyword `flowbits`, are currently mostly ignored by the Snort detection engine. Statements in rules that test flow bits currently always pass. However, the detection engine will honor flowbit directives that specify `noalert`. Rules marked as `noalert` never generate meta, even if they match.

DCE Preprocessor

The Snort parser does not currently implement any DCE preprocessors, so any features that rely on the DCE preprocessor are not implemented.

Thresholds and Filters

The Snort parser implements the `threshold` filters option only. It does not yet implement the `detection_filter` filter.

Raw Byte Scanning

The Snort parser does not currently allow rules to individually bypass the normalization or decoding using tags like `rawbytes`.

IPv6

The Snort rule parser does not currently allow rules to utilize IPv6 addresses as part of their header.

Note: Only the Snort V2.x rules are supported.

Performance Considerations

The performance of the Snort parser is dependent entirely on the rules that are being evaluated. Therefore it will vary greatly from one installation to another.

Content Tokens

Content tokens form the main entry point into the Snort engine. Hence they are the main factor affecting rules performance. The more times that a content string is found in the packet stream, the more work the Snort parser performs, and the more CPU time it uses. Therefore, it's important to use content tokens that are long, unique, and unlikely to appear in random data. Short tokens will spuriously activate the Snort parser, wasting CPU time.

Negation

Statements like `content` and `pcre` support negation, meaning they only allow the rule to match if a content string is not found or if a regex doesn't match. These are usually detrimental to performance because they still tend to contain strings that will activate rule evaluation, yet will often fail since they are intended to filter out patterns. For example, a negated `pcre` search still has to be evaluated if it is negated, and if it is used to filter out common text. it will produce lots of CPU activity on traffic that does not actually match the Snort rule.

Search Parser

The Search Parser is a custom parser used to generate metadata by scanning for predefined keywords and regular expressions. The parser searches the payload of a reconstructed session for string matches and can execute a regular expression search. You can configure the parser by editing the **search.ini** file.

Caution: The search parser can have a significant impact on system performance. It is important that both the search mechanism and the data to which it is applied to be well understood before creating new search definitions and enabling the search parser.

The search definition is used across all protocols. There are three basic search methods:

- Keyword: Search a stream for a specific set of words
- Pattern: Search a stream for a regular expression match
- Keyword + Pattern: Search a stream for a regular expression if it contains any of a given set of keywords.

Search Methods

The Search parser uses three basic search methods:

- Keyword: Search a stream for a specific set of words.
- Pattern: Search a stream for a regular expression match.
- Keyword+Pattern: Search a stream for a regular expression if it contains any of a given set of key words.

Syntax

```
Maxrecon=<max_size>Maxsearch=<max_ssearch_length>MatchLimit=<max_matches_per_
stream
Search Name
Services=<service_id_list>Keywords=<keyword_list>|Pattern=<expression>Case=0|1
Proximity=<number_of_bytes>Recon=0|1
Raw=0|1
```

Parameters

Parameters used in this command:

Parameter	Description
`autocheck`	Automatically fixes all problems without prompting
`header Only`	Check/display the header of each file

Parameter	Description
chatty	Displays a hex dump of every object in the file (huge amount of data)
dump#-#	Indicates a zero-based object or range of objects in the file to output in hex to the console

Example

Following is an example of the command:

To check all NetWitness database files located in the Collection named Default. If any problems are found, the command will describe the problem and ask if you would like to fix it.

```
dbcheck C:\Documents and Settings\User\My Documents\NetWitness\
Investigations\Default\*.nw*
```

Wireless LAN Configuration

One of the files available for editing in the Services Config view > Files tab is **wlan-config.xml**, the wireless LAN configuration file.

It controls the 802.11 parsers. Its chief purpose is to control decryption of raw 802.11 frames captured by the Decoder. This file is optional. If decryption of 802.11 traffic is not desired, there is no need to create the file.

There are five link-level parsers related to wireless LAN packet capture:

- IEEE 802.11 parser (data frames and beacons only)
- Radiotap w/ 802.11 header
- Absolute Value Systems (AVS) w/ 802.11 header
- Prism II w/ 802.11 header
- CACE's "Per Packet Information" (PPI) w/ 802.11 header

The 802.11 wireless parsers introduced in 9.8 all share a single configuration file. This wlan-config.xml file is used to define any wireless access points the user may have in the network, and its primary purpose is to control decryption. The BSSID of the access point and the SSID that it's authoritative for is added to this file as well as all of the active default keys used by the access point.

Troubleshooting Parsers

This topic provides guidance for troubleshooting issues related to parsers.

Lua Parser Errors

Lua parsers occasionally generate errors. If a parser enters a state where it generates multiple errors, these errors can hinder performance. Beginning with 11.4, a new option is available for Lua parsers that instructs the Decoder to automatically disable the parser after a configurable number of errors.

The value is set on the `/decoder/parsers/config/parsers.options` configuration node as shown in the following example:

```
Entropy="log2=true" GeoIP2="ipaddr=ip.src,ip.dst,ipv6.src,ipv6.dst" error_
parser="errorMax=5" addy_parser="errorMax=10"
```

This configuration node enables you to set options for different parsers. In this example, the Lua parser `error_parser` is configured with a maximum error count of 5, and the `addy_parser` to 10. The `errorMax` setting has a valid range of values from 0 (meaning the feature is disabled) to 65,535, and takes effect when parsers are loaded or reloaded.

To disable a Lua parser after a defined number of errors:

1. Go to **ADMIN** > **Services**, select a Decoder and then select  > **View** > **Explore**.
2. In the left panel, expand **decoder** > **parsers**, and select **config**.
3. In **parsers.options**, add the following command, where `<any_parser>` is the Lua parser for which you want to limit errors, and `<n>` is the number of errors to which the parser is limited:
   `<any_parser>="errorMax=<n>"`

4. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**.

Results of Automatically Disabling a Parser

When a parser is automatically disabled, a log message is generated (per parser thread) and states that the parser has been disabled, as shown in the following example:

```
(W) 2019-Apr-25 16:25:33 [Parse] Lua parser 'error_parser' has been disabled
because it exceeded the configured error threshold (5)
```

If detailed statistics for parsers are enabled, the text of the last error is populated in the parser's detailed statistics under `/decoder/parsers/definitions/<parser-name>`. Also, a new attribute has been added to the XML returned by the `/decoder/parsers?msg=schema` call. When the parser is in an error state, the attribute `error` is set to `1` and the `enabled` attribute is set to `0`. When the parser is reset, the values are reset to `0` and `1` respectively.

Resetting the Parser

If a parser is disabled because it has exceeded the error threshold, it can be reset by reloading the parser. This causes the statistics to reset so that if the faulty parser is still in place, it will function again until the error threshold is met.

Configure Feeds

NetWitness uses feeds to create metadata based on externally defined metadata values. A feed is a list of data that is compared to sessions as they are captured or processed. For each match, additional metadata is created. This data could identify and classify malicious IPs or incorporate additional information such as department and location based on internal network assignments. Some examples of feeds include threat feeds to identify BOTNets, DHCP mappings, or even Active Directory (AD) information such as physical location or logical department.

You can use the Live module in NetWitness to obtain feeds from outside sources. "Live Content in NetWitness" in the *Live Services Management Guide* provides an overview of the Live content management tool.

Within the NetWitness user interface, you can view the list of currently deployed feeds, along with an indicator if a feed that originated from Live was installed through NetWitness or manually. Feeds can be added, removed, and updated while a Decoder is running without affecting capture.

Custom Feed Definition File Structure

The NetWitness Custom Feed wizard allows creation and deployment of custom Decoder feeds based on deterministic logic that offers the meta keys specific to the selected Decoders and Log Decoders. Although the wizard guides users through the process to create both on-demand and recurring feeds, it is helpful to understand the form and content of a feed file when you create a feed.

Before you create or modify custom feeds, keys, or data types, ensure that you have a good understanding of the relationship between keys and datatypes in databases. It is important that any custom keys that you create match the corresponding datatypes. If your feed utilizes a custom meta key, you must define that key's data type in `index-decoder-custom.xml`. This is to ensure that the key data is generated with correct data format.

Also, to enable investigation on the custom meta key, you must define its index on the concentrator, using the `index-concentrator-custom.xml` file. For more information on how to write the `index-decoder-custom.xml` or `index-concentrator-custom.xml` file, see the topic "Index Customization" in the *Core Database Tuning Guide for RSA NetWitness Platform*.

Feed filenames in RSA NetWitness are in the form `<filename>.feed`. To create a feed, NetWitness requires a feed data file in `.csv` or `.xml` format and a feed definition file in `.xml` format, which describes the structure of a feed data file. The Custom Feed wizard can create the feed definition file based on a feed data file, or based on a feed data file and the corresponding feed definition file.

The files that you use to create an on-demand feed must be stored on your local file system. The files used to create a recurring feed must be stored at an accessible URL, whence NetWitness can fetch the most current version of the file for each recurrence. After a NetWitness feed is created, you can download the feed to your local file system, edit the feed files, and then edit the NetWitness feed to use the updated feed files.

Sample Feed Definition File

This is an example of a feed definition file named `dynamic_dns.xml`, which NetWitness creates based on your entries in the Custom Feed wizard. It defines the structure of the feed data file named `dynamic_dns.csv`.

Note: The feed file path should be `.csv` regardless of the Feed Type (Default or STIX).

```
<?xml version="1.0" encoding="utf-8"?>
 <FDF xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="feed-definitions.xsd">

   <FlatFileFeed name="Dynamic DNS Domain Feed"
   path="dynamic_dns.csv"
   separator=","
   comment="#"
   version="1">

   <MetaCallback
   name="alias.host"
   valuetype="Text"
   apptype="0"
   truncdomain="true"/>

   <LanguageKeys>
            <LanguageKey name="threat.source" valuetype="Text" />
            <LanguageKey name="threat.category" valuetype="Text" />
            <LanguageKey name="threat.desc" valuetype="Text" />
```

```
    </LanguageKeys>

    <Fields>
    <Field index="1" type="index" key="alias.host" />
    <Field index="4" type="value" key="threat.desc" />
    <Field index="2" type="value" key="threat.source" />
    <Field index="3" type="value" key="threat.category" />
    </Fields>
    </FlatFileFeed>

  </FDF>
```

Define Multiple Values in a Single Field

The information in this section applies to NetWitness version 11.3.1 and later.

Custom feeds support multiple values in a single field. This allows feeds to generate multiple values of the same meta.

For example:

```
<Fields>

     <Field index="1" type="index" key="lc.cid" />

     <Field index="2" type="value" key="user" separator=";" openchar="("
     closechar=")"/>

</Fields>
```

The second field is defined to accept multiple values in the corresponding CSV file. The following is an example of the defined values:

```
clc1,(bob;tom;sam) clc12,susan clc123,doris
```

This definition means that when we see the `lc.id` with a value of `cic1`, three `username` meta values are generated with the values of `bob`, `tom` and `sam`.

The separator in the field definition for a multi-value field **must** be different than the separator for the feed file itself. Escaping the separator character in field values is **not** supported.

Feed Definition Equivalents for Custom Feed Wizard Parameters

The NetWitness Custom Feed wizard provides options to define the structure of the data feed file. These correspond directly to attributes in the feed definition (`.xml`) file.

NetWitness Parameter	Feed Definition File Equivalent
(Define Feed Tab) Feed Type	Select: **Default** - to define a feed based on a `.csv` formatted feed data file. **STIX** - to define a feed based on STIX formatted `.xml` file.
(Define Feed Tab) Feed Task Type	Select: **Adhoc** - to create an on-demand feed. **Recurring** - to update the `.csv` or `.xml` file persistently and store it in a location accessible by NetWitness, so NetWitness downloads a file at regular intervals and pushes it to the downstream devices.

NetWitness Parameter	Feed Definition File Equivalent
(Define Feed tab) **Name**	The custom feed name in the feed data file. It corresponds to the `flatfeedfile name` attribute in the feed definition file. For example, Dynamic DNS Test Feed. **Note:** You can use special characters to define the name of the custom feed.
(Define Feed tab) **File**/ **Browse**	This is the name of the feed data file. It corresponds to the `flatfeedfile path` attribute in the feed definition file. For example, `dynamic_dns.csv`.
(Advanced Options tab) **XML Feed File**	The name of the feed definition file. For example, `dynamic_dns.xml`.
(Advanced Options tab) **Separator**	The separator character used to separate attributes in the feed data file. It corresponds to the `flatfeedfile separator` in the feed definition file. For example, a comma.
(Advanced Options tab) **Comment**	The character used to identify a comment in the feed data file. It corresponds to the `flatfeedfile comment` attribute in the feed definition file. For example, `#`.
(Define Columns tab, Define Index) **Type**	The type of lookup value in the index position of the feed data file. **IP** means that each row in the feed data file contains an IP address in the lookup value position. The IP value is in dotted-decimal format (for example, 10.5.187.42). **IP Range** means that each row in the feed data file contains a range of IP addresses in the lookup value position. The IP range is in CIDR format (for example, 192.168.2.0/24). **Non IP** means that the each row in the feed data file contains a metadata value other than IP address in the lookup value position. The Service Type and Truncate Domain, and Callback Keys fields become active for a Non IP index.
(Define Columns tab, Define Index) **CIDR**	Specifies that the IP value in the lookup position is in CIDR format. The **CIDR** attribute sets the IP address format in the field to Classless Inter-Domain Routing (CIDR) notation.
(Define Columns tab, Define Index) **Service Type**	For a Non IP index, the integer service type to filter meta lookups. It corresponds to the `MetaCallback apptype` attribute in the feed definition file. A value of 0 indicates no filtering by service type.
(Define Columns tab, Define Index) **Truncate Domain**	For a Non IP index, for meta values that contain domain names (for example, hostnames), the system can strip off the host specific element in the data. Truncate Domain corresponds to the `MetaCallback truncdomain` attribute. If the value is `www.example.com`, it is truncated to `example.com`. A value of **False** selects no truncation, and **True** selects truncation.

NetWitness Parameter	Feed Definition File Equivalent
(Define Columns tab, Define Index) **Callback Keys**	For a Non IP index, the available meta keys to match on instead of ip.src/ip.dst (the defaults for IP index type) are selectable from the drop-down list. The Callback Key corresponds to the `MetaCallback name` attribute, and the index column of the csv file must contain data that can match the chosen meta key. For example, if the `user` meta key is chosen, the index column of the `.csv` file needs to be populated with users to be matched.
(Define Columns tab, Define Index) **Index Column**	Identifies the column in the feed data file that provides the lookup value for the row. Each position in each row of the feed data file is identified by a **Field index** attribute in the feed definition file. A field with an index of **1** is the first entry in a row, the second field has an index of **2**, the third field has an index of **3**, and so on.
(DEFINE VALUES) **Key**	The name of the `LanguageKey`, as defined in the feed definition file, for which meta is created from this row of the feed data file. It corresponds to the `Field key` attribute in the feed definition file. A key applies only to a field whose type is set to `value`. In the feed definition file, there is a list of `LanguageKeys` from `index.xml`, or a summary name if Source Name and Destination Name are used. For example, `reputation` is a summary name for `reputation.src` and `reputation.dst`. This value is referenced by the Field key attribute.

Sample Files for a MetaCallback Feed Using CIDR Index Range for IPv4 and IPv6

These sample files demonstrate how to use CIDR index ranges for IPv4 and IPv6 in custom MetaCallback feeds. As with other custom feeds, you must create feed data file in .csv format, and a feed definition file in .xml format.

Note: Using MetaCallback feeds with CIDR index ranges is supported only through the Advanced Configuration wizard or the REST interface.

The following example shows the contents of both a `.csv` file and an `.xml` file for a MetaCallback feed using CIDR index ranges for IPv4 or IPv6.

`.csv file:`

```
192.168.0.0/24, Sydney
192.168.1.0/24, Melbourne
```

`.xml file:`

```
<?xml version="1.0" encoding="UTF-8"?>
<FDF xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="feed-definitions.xsd">
<FlatFileFeed name="ip_test" path="ip_test.csv" separator="," comment="#">
      <MetaCallback name="DstIP" valuetype="IPv4" apptype="0"
truncdomain="false">
            <Meta name="ip.dst"/>
      </MetaCallback>
      <LanguageKeys>
            <LanguageKey name="alert" valuetype="Text" />
      </LanguageKeys>
```

```
      <Fields>
            <Field index="1" type="index" range="cidr"/>
            <Field index="2" type="value" key="alert" />
      </Fields>
</FlatFileFeed>
</FDF>
```

Note: To configure a CIDR index range for feeds with single or multiple MetaCallbacks of value type IPv4 or IPv6, the field of type index MUST contain a range attribute with `range="cidr"`. Also, configuring "cidr" index ranges for feeds with MetaCallbacks of multiple different value types is not supported.

Feed Definitions File

This topic introduces the feed definitions file, which is available for editing in the Services Config view > Files tab. One of the files available for editing in the Services Config view > Files tab is **`feed-definitions.xml`**, the feed definitions file.

You can define feeds in the `feed-definitions.xml` file. The Decoder uses an XML schema to define feed messages when it creates a binary .feed file from the feeds defined here.

For details on the feed definition language, refer to Custom Feed Definition File Structure

Create a Custom Feed

You can create a custom feed using the Custom Feed wizard. To complete this procedure, you need a feed data file in `.csv` or `.xml` format. If you also have an associated feed definition file in `.xml` format, which describes the structure of the feed data file, you can use the feed definition file to create a feed. The Custom Feed wizard can create the feed based on a feed data file, or based on a feed data file and corresponding feed definition file.

Before you create or modify custom feeds, keys, or data types, ensure that you have a good understanding of the relationship between keys and datatypes in databases. It is important that any custom keys that you create match the corresponding datatypes. If your feed utilizes a custom meta key, you must define that key's data type in `index-decoder-custom.xml`. This is to ensure that the key data is generated with correct data format.

Also, to enable investigation on the custom meta key, you must define its index on the concentrator, using the `index-concentrator-custom.xml` file. For more information on how to write the `index-decoder-custom.xml` or `index-concentrator-custom.xml` file, see the topic "Index Customization" in the *Core Database Tuning Guide for RSA NetWitness Platform*.

Note: For information about STIX and creating a STIX custom feed, see "Create a STIX Custom Feed" in the *Decoder and Log Decoder Configuration Guide*.

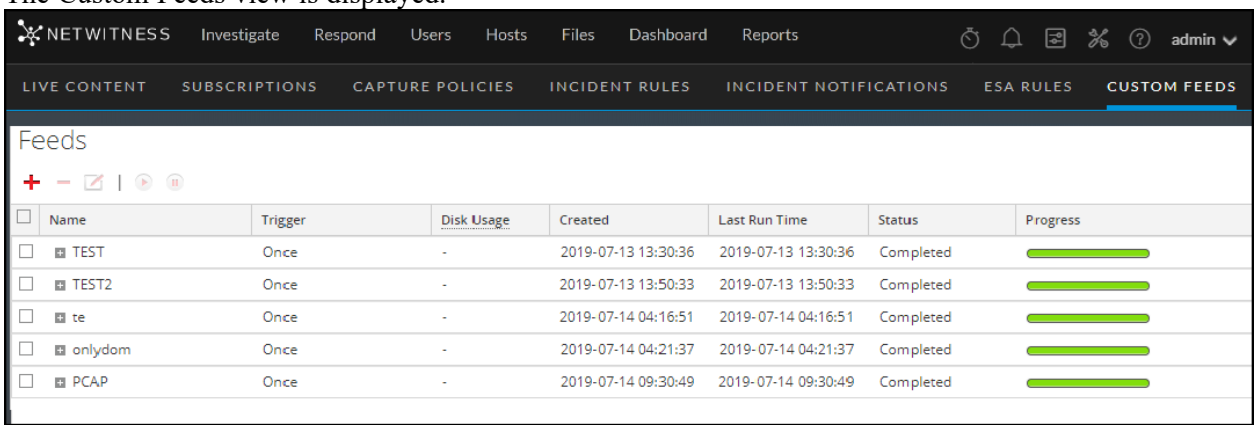
The feed data file and optionally the feed definition file (`.xml`) must be available on the local file system for an on-demand custom feed. For a recurring custom feed, the files must be available at a URL that is accessible to the NetWitness server.

Note: When you create a source and destination-based feed on a Log Decoder, it only populates the source meta key. You cannot use a range-based or CIDR feed. You must list every single IP address. To resolve this issue, create two different feeds using IP addresses and you can use CIDR in these feeds.

To create a custom feed:

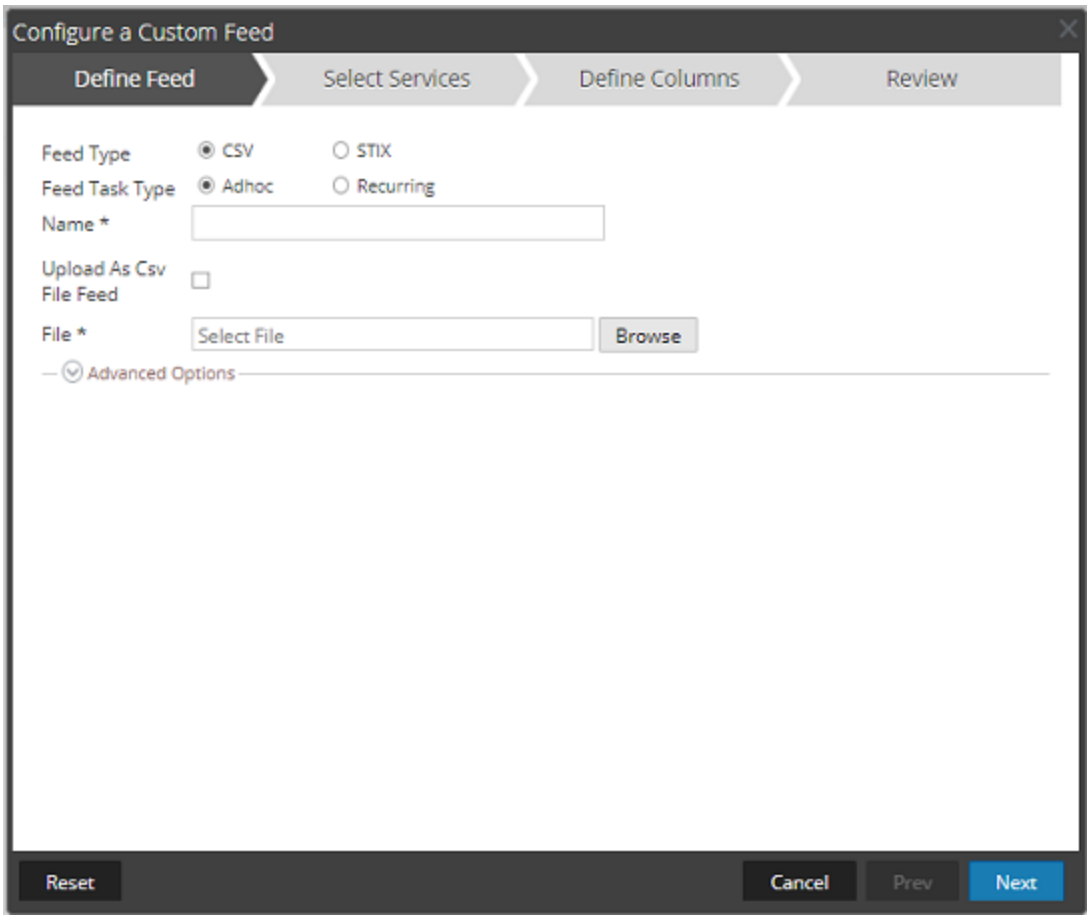
1. Go to **Configure > Custom Feeds.**

   The Custom Feeds view is displayed.



2. In the **Feeds** panel, click + > **Custom Feed** > **Next**.

   The Configure a Custom Feed wizard is displayed, with the Define Feed form open.



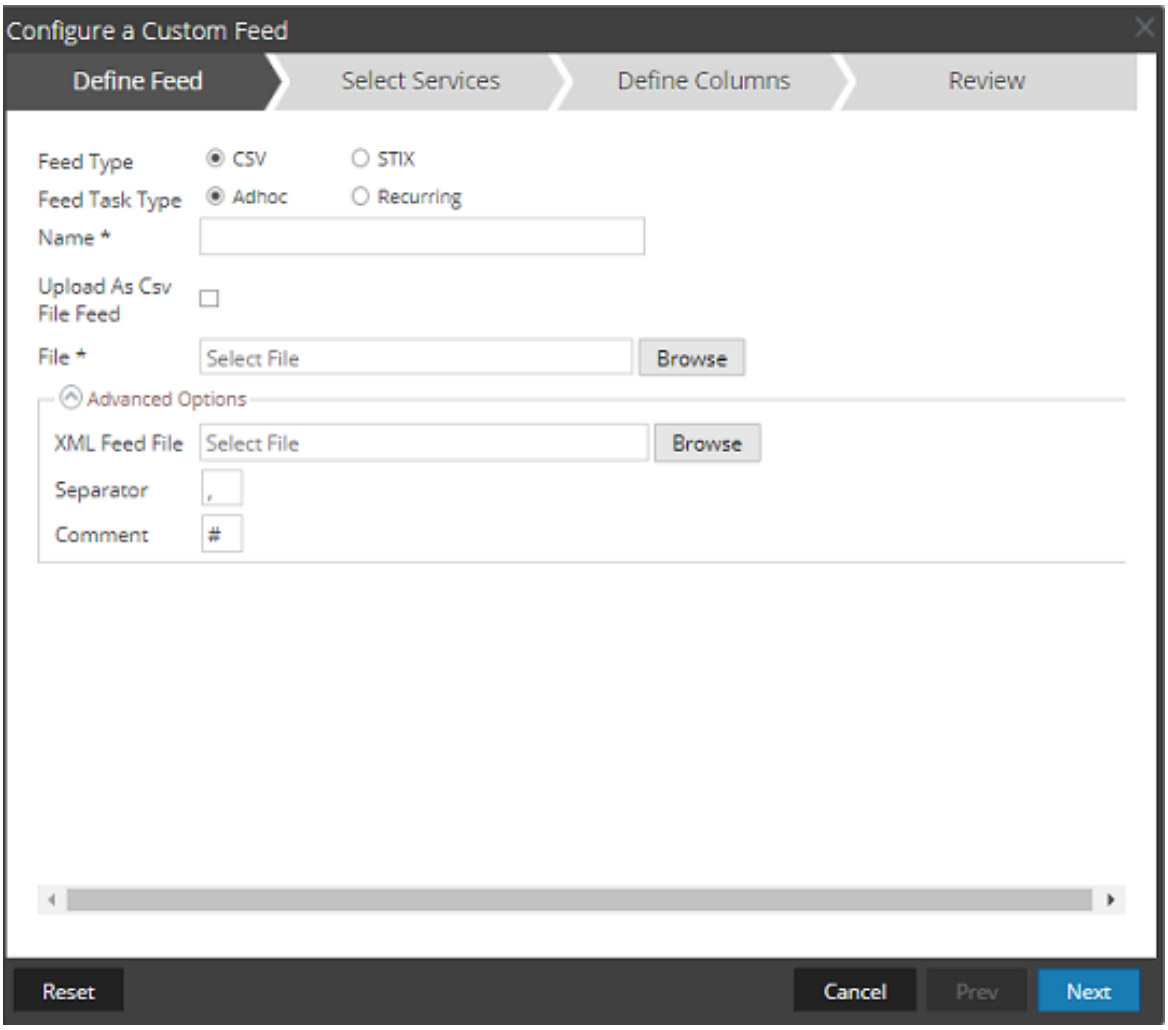
3. Select the Feed Type: **CSV** or **STIX**.
4. To define a feed based on a `.csv` formatted feed data file, select **CSV** (which is the default) in the **Feed Type** field.
5. To define an on-demand feed task that executes once, select **Adhoc** in the **Feed Task Type** field and do one of the following:
   a. (Conditional) To define a feed based on a CsvFileFeed file, select the **Upload as Csv File Feed** checkbox, type the feed **Name**, select a .csv content file from the local file system, and click **Next**. If you do not select the checkbox, the .csv file will be a FlatFileFeed file.

      Note: When you select the Upload as Csv File Feed checkbox, the XML feed options under Advanced are unavailable.

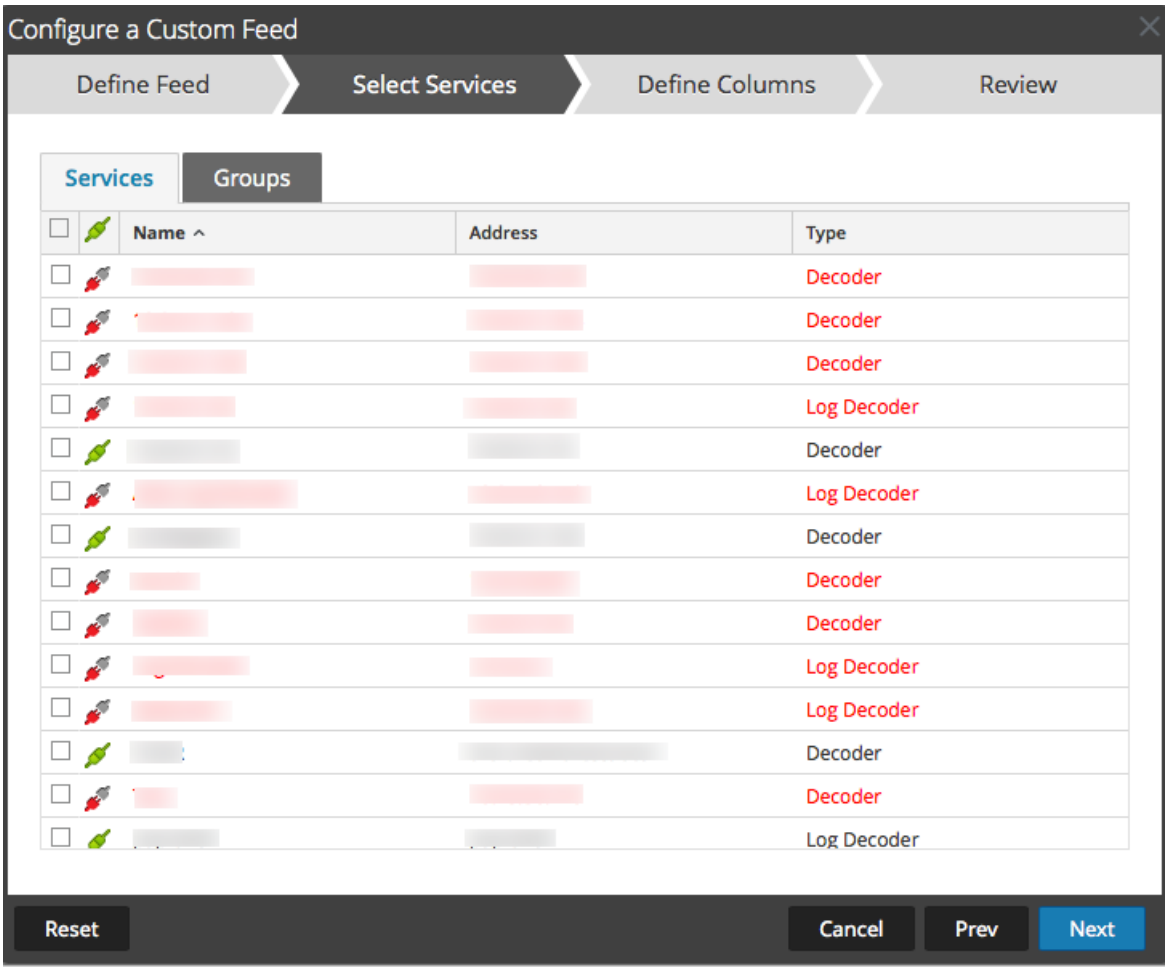
   b. (Conditional) To define a feed based on an XML feed file, select **Advanced Options**.

      Note: Ensure that the Upload as Csv File Feed checkbox is deselected.

      The Advanced Options are displayed:



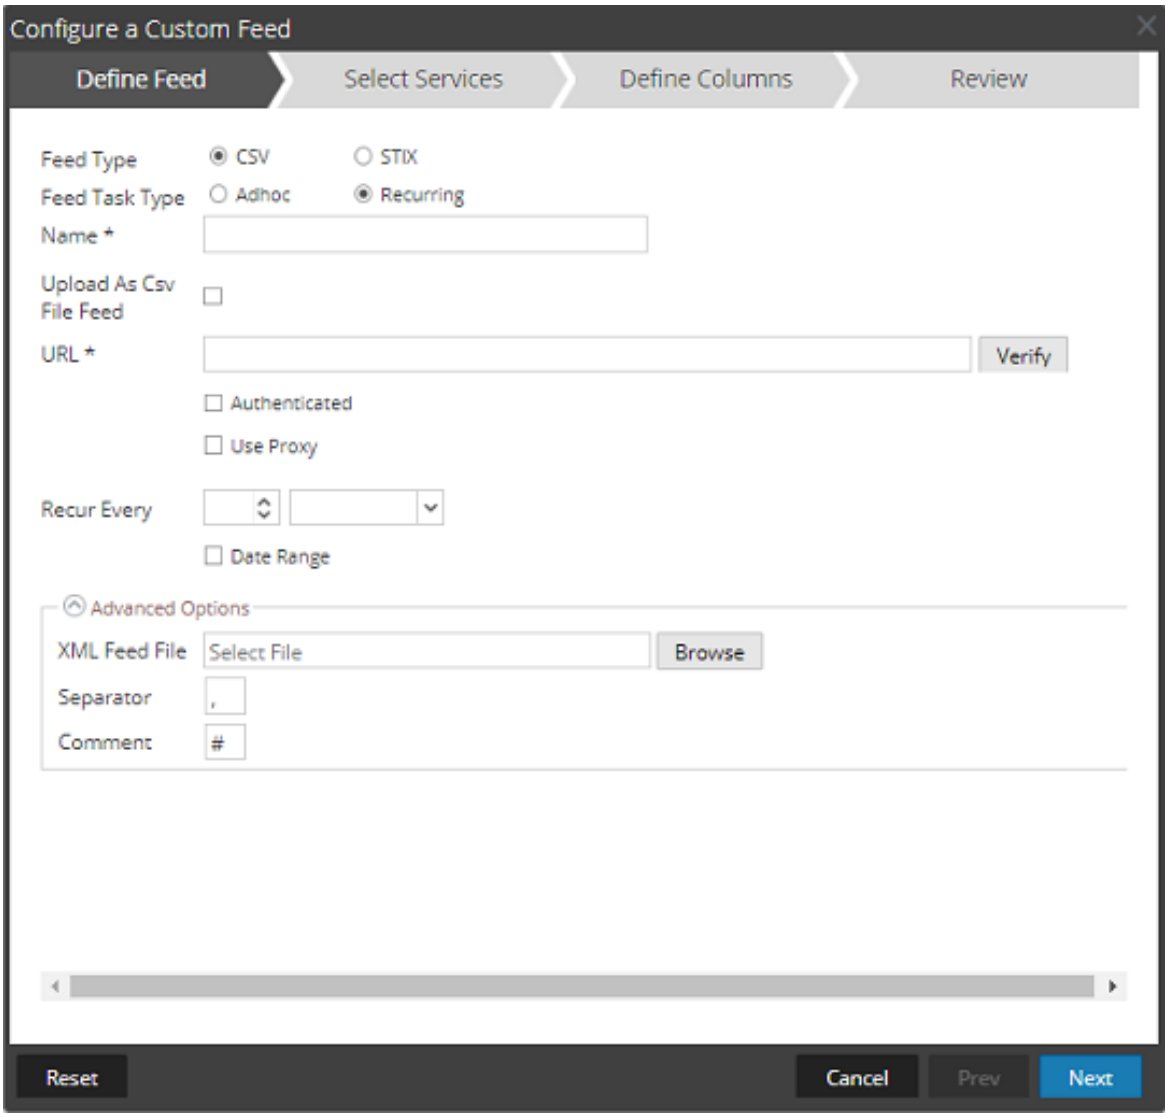
c. Select an XML feed file from the local file system, choose the separator (default is comma), specify the comment characters used in the feed data file (default is #), and click **Next**.
The Select Services form is displayed. This is an example of the form for a feed based on a feed data file with no feed definition file. If you are defining a feed based on a feed definition file, the Define Columns tab is not needed.



6. To define a recurring feed task that executes repeatedly at specified intervals, during a specified date range:

   a. In the Define Feed form, select **Recurring** in the **Feed Task Type** field.

      The Define Feed form includes the fields for a recurring feed.



b. In the **URL** field, enter the URL where the feed data file is located, for example, `http://<hostname>/<feeddatafile>.csv`, and click **Verify**. NetWitness verifies the location where the file is stored in order to enable checking for the latest file automatically before each recurrence.

c. (Optional) If the URL has restricted access and requires authentication using your username and password, select **Authenticated**.
NetWitness provides your user name and password for authentication to the URL.

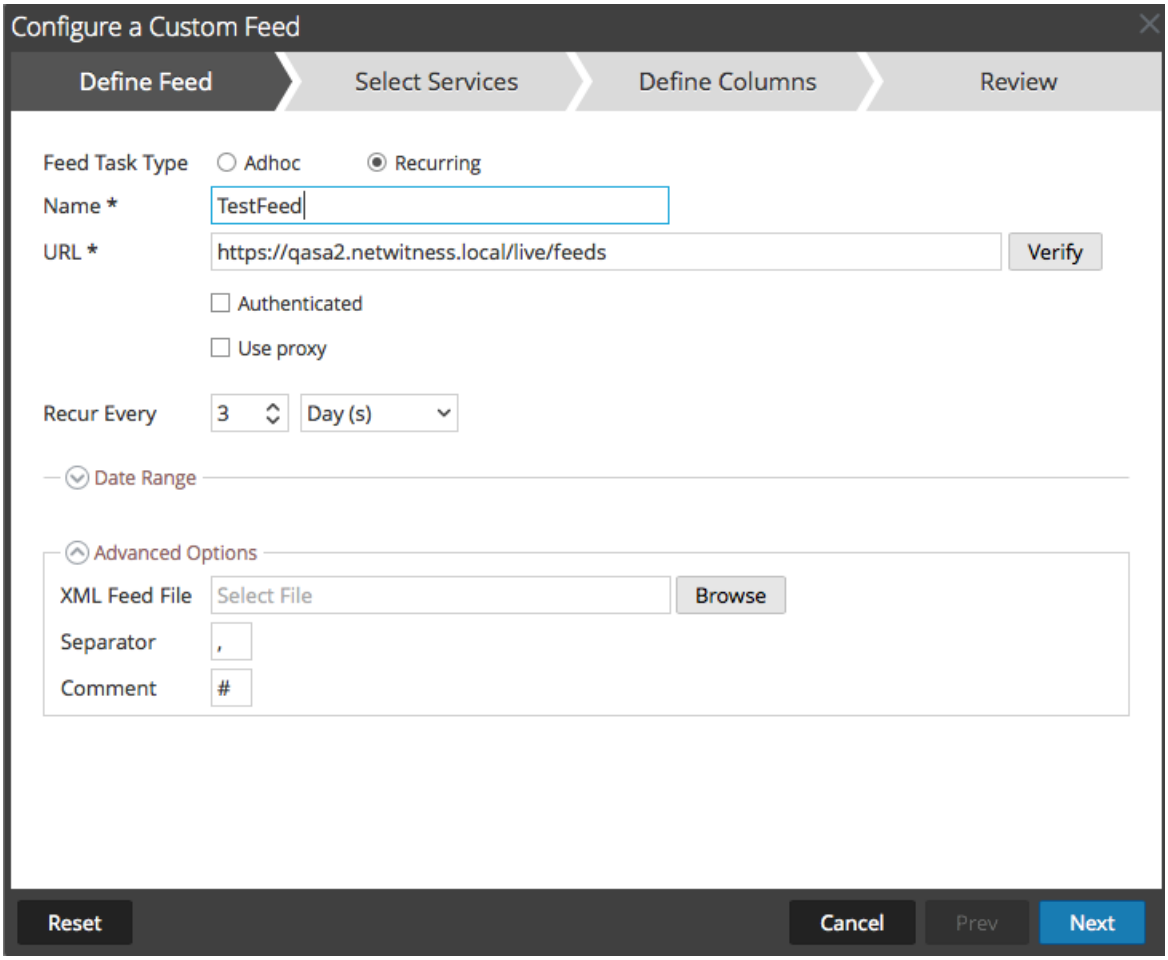
d. If you want the NetWitness server to access the Feed URL through a proxy, select **Use Proxy**. For more information on configuring a proxy, see "Configure Proxy for NetWitness" in the *System Configuration Guide*. By default, the **Use Proxy** checkbox is not set.

e. To define the interval for recurrence, do one of the following:

- Specify the number of minutes, hours, or days between recurrences of the feed.

- Specify recurrence every week, and select the days of the week.

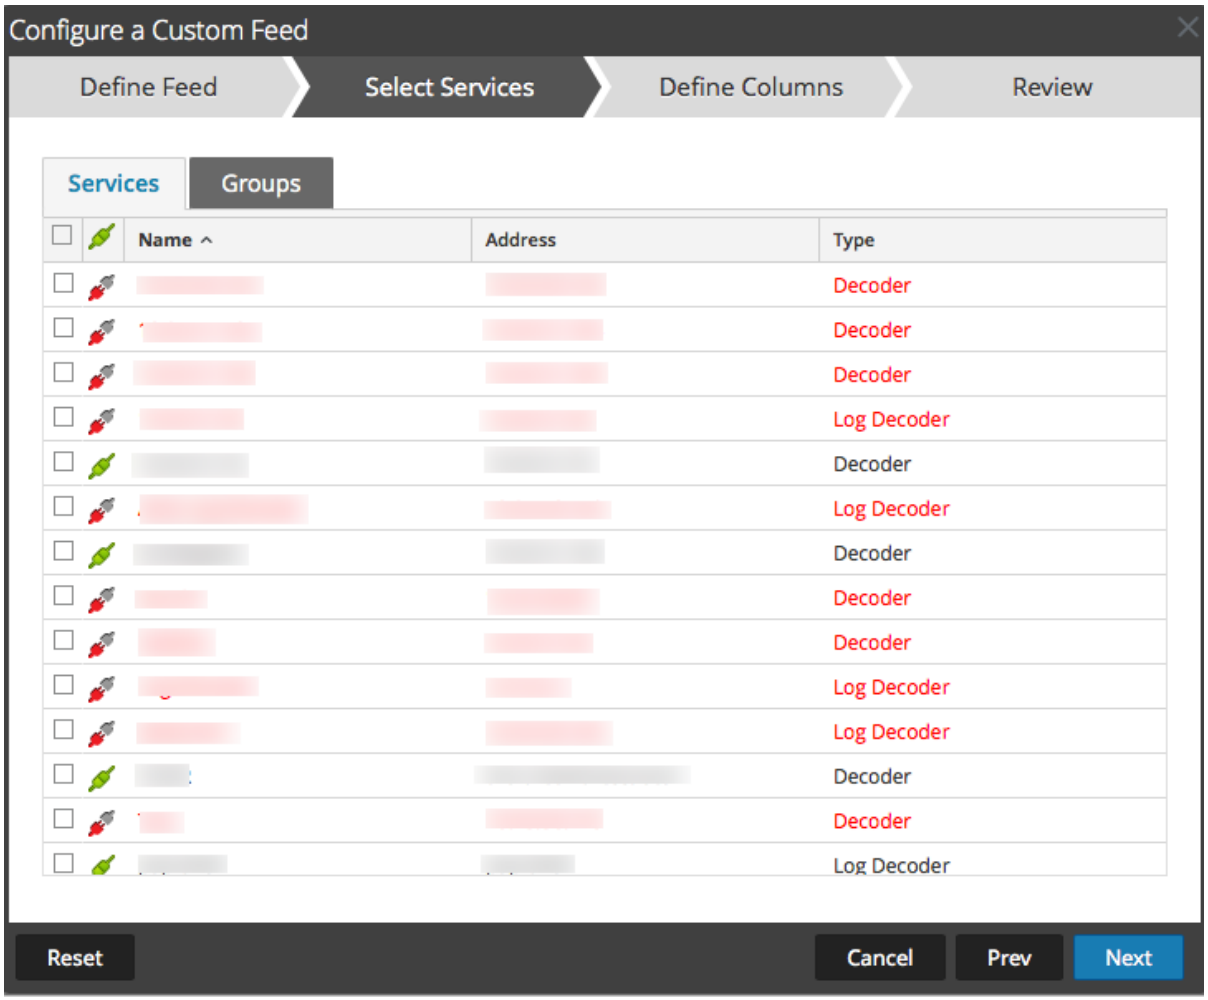
f. To define the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.



7. (Conditional) If you want to define a feed based on an XML feed file:

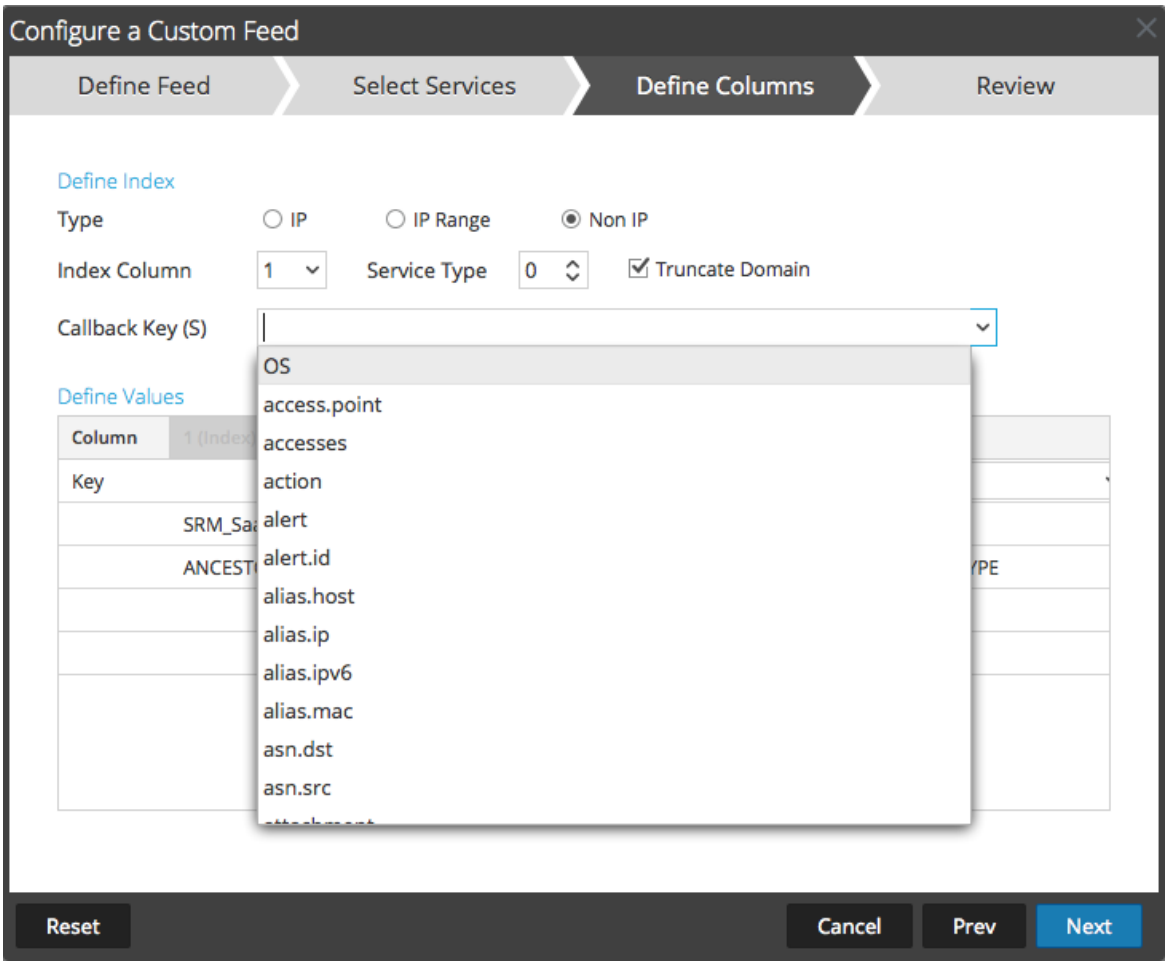
   a. Type the feed **Name**, select **Advanced Options**. The Advanced Options fields are displayed.

   b. Select an XML feed file from the local file system, choose the **Separator** (default is comma), specify the **Comment** characters used in the feed data file (default is #) and click **Next**. The Select Services form is displayed.

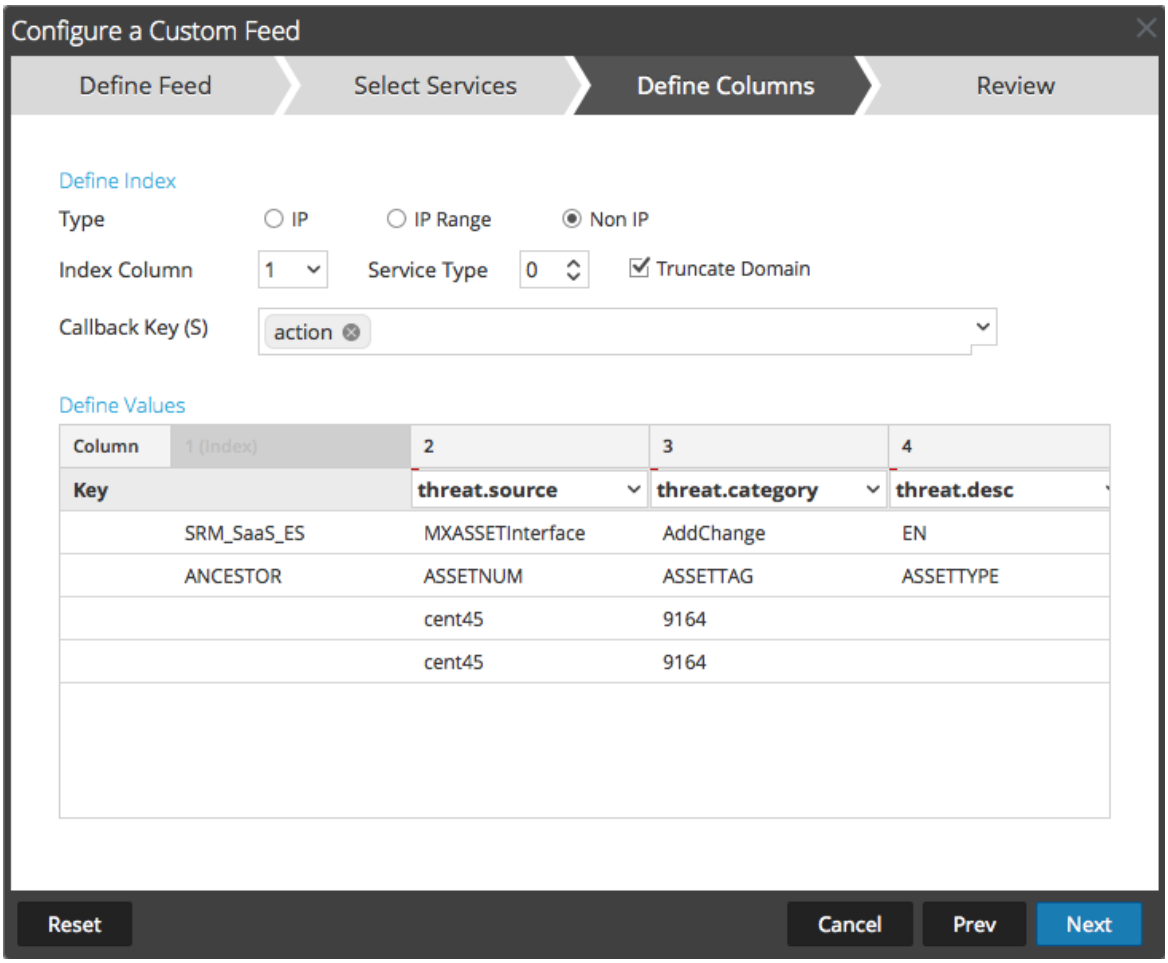


8. To identify services on which to deploy the feed, do one of the following:
   a. Select one or more Decoders and Log Decoders, and click **Next**
   b. Click the **Groups** tab and select a group. Click **Next**.

      The Define Columns form is displayed.
9. To map columns in the Define Columns form:
   a. Define the Index type: **IP**, **IP Range**, or **Non IP**, and select the index column.
   b. (Conditional) If the index type is **IP** or **IP Range** and the IP address is in CIDR notation, select **CIDR**.
   c. Conditional) If the index type is **Non IP**, additional settings are displayed. Select the service type and **Callback Keys,** and optionally select the **Truncate Domain** option.

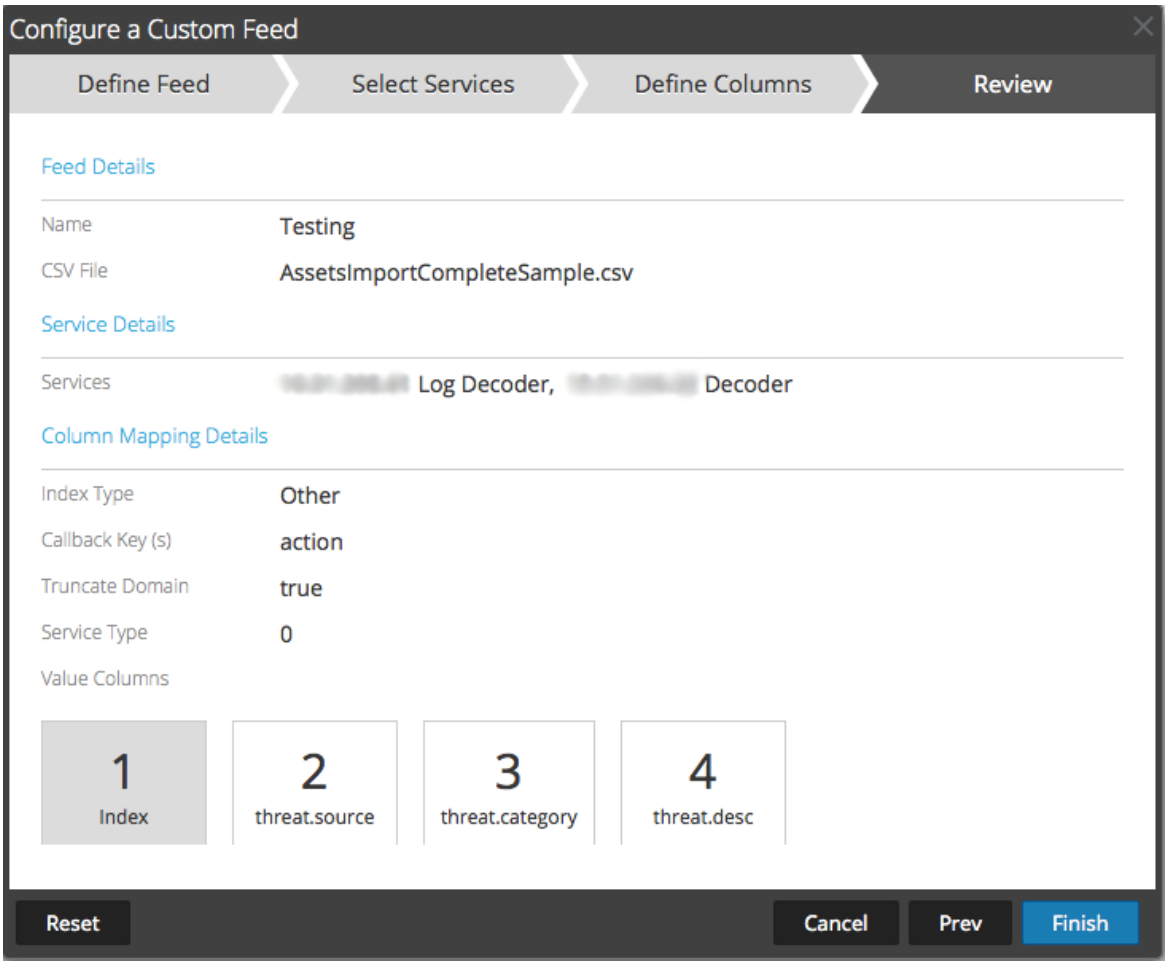


d. Select the language key to apply to the data in each column from the drop-down list. The meta keys displayed in the drop-down list is based on the meta keys available for the service define values. You can also add other meta keys based on advanced expertise.

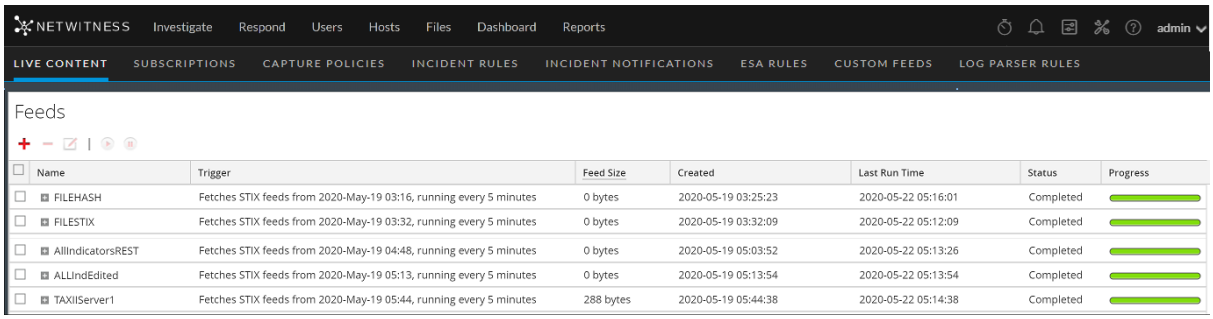


e. Click **Next**.

The Review form is displayed.



10. Anytime before you click **Finish**, you can:
    - Click **Cancel** to close the wizard without saving your feed definition.
    - Click **Reset** to clear the data in the wizard.
    - Click **Next** to display the next form (if not viewing the last form).
    - Click **Prev** to display the previous form (if not viewing the first form)
11. Review the feed information, and if correct, click **Finish**.
12. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file are listed in the Feed grid and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
LIVE CONTENT
SUBSCRIPTIONS
CAPTURE POLICIES
INCIDENT RULES
INCIDENT NOTIFICATIONS
ESA RULES
CUSTOM FEEDS
LOG PARSER RULES
Feeds
Name
Trigger
Feed Size
Created
Last Run Time
Status
Progress
FILEHASH
Fetches STIX feeds from 2020-May-19 03:16, running every 5 minutes
0 bytes
2020-05-19 03:25:23
2020-05-22 05:16:01
Completed
FILESTIX
Fetches STIX feeds from 2020-May-19 03:32, running every 5 minutes
0 bytes
2020-05-19 03:32:09
2020-05-22 05:12:09
Completed
AllIndicatorsREST
Fetches STIX feeds from 2020-May-19 04:48, running every 5 minutes
0 bytes
2020-05-19 05:03:52
2020-05-22 05:13:26
Completed
ALLIndEdited
Fetches STIX feeds from 2020-May-19 05:13, running every 5 minutes
0 bytes
2020-05-19 05:13:54
2020-05-22 05:13:54
Completed
TAXIIServer1
Fetches STIX feeds from 2020-May-19 05:44, running every 5 minutes
288 bytes
2020-05-19 05:44:38
2020-05-22 05:14:38
Completed

Create a STIX Custom Feed

Structured Threat Information Expression (STIX™) is a structured language for describing cyber threat information so it can be shared, stored, and analyzed in a consistent manner. For more information about STIX, see https://stixproject.github.io/.

You can create a custom feed using a STIX-formatted feed data file (`.xml`) in RSA NetWitness. NetWitness supports Structured Threat Information Expression (STIX) 1.0, 1.1 and 1.2 versions only.

Caution: If a STIX recurring feed is configured and you update Security Analytics from 10.6.x to NetWitness 11.x, you must re-configure the STIX recurring feed.

In NetWitness, STIX feeds of type Indicator or Observable that contain properties such as the IP addresses, File hashes, Domain names, URIs and Email addresses are supported. The property values in the Equals operator are supported. Attributes such as Type and Title are also read from the STIX. A STIX file with a single STIX_Package is supported.

TAXII (Trusted Automated eXchange of Indicator Information) is the main transport mechanism for cyber threat information represented in STIX. Using the TAXII services, organizations can share cyber threat information in a secure and automated manner.

The STIX and TAXII communities work closely together to ensure that they continue to provide a full stack for sharing threat intelligence.

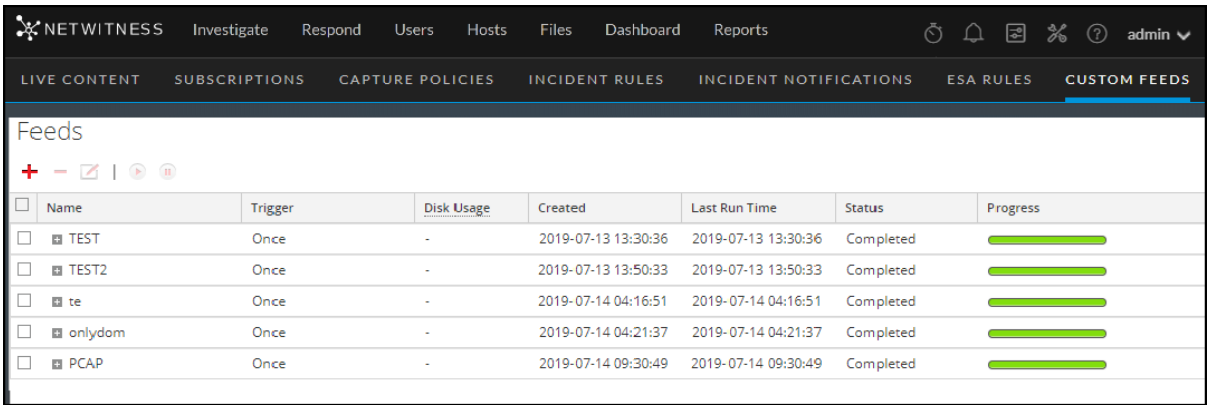
Apart from the TAXII server, STIX data can also reside on a REST server and you can fetch the STIX file from the REST server by providing the URL of the REST server. For example, `http://stixrestserver.internal.com`.

The STIX feed data file and optionally the feed definition file, both in `.xml` format must be available on the local file system for an on-demand custom feed. For a recurring custom feed, the files must be available at a URL that is accessible to the NetWitness server.

To create a STIX custom feed:

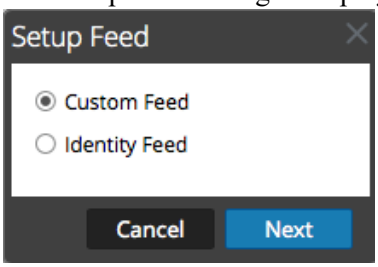
1. Go to **(Configure) > CUSTOM FEEDS**.

   The Custom Feeds view is displayed.



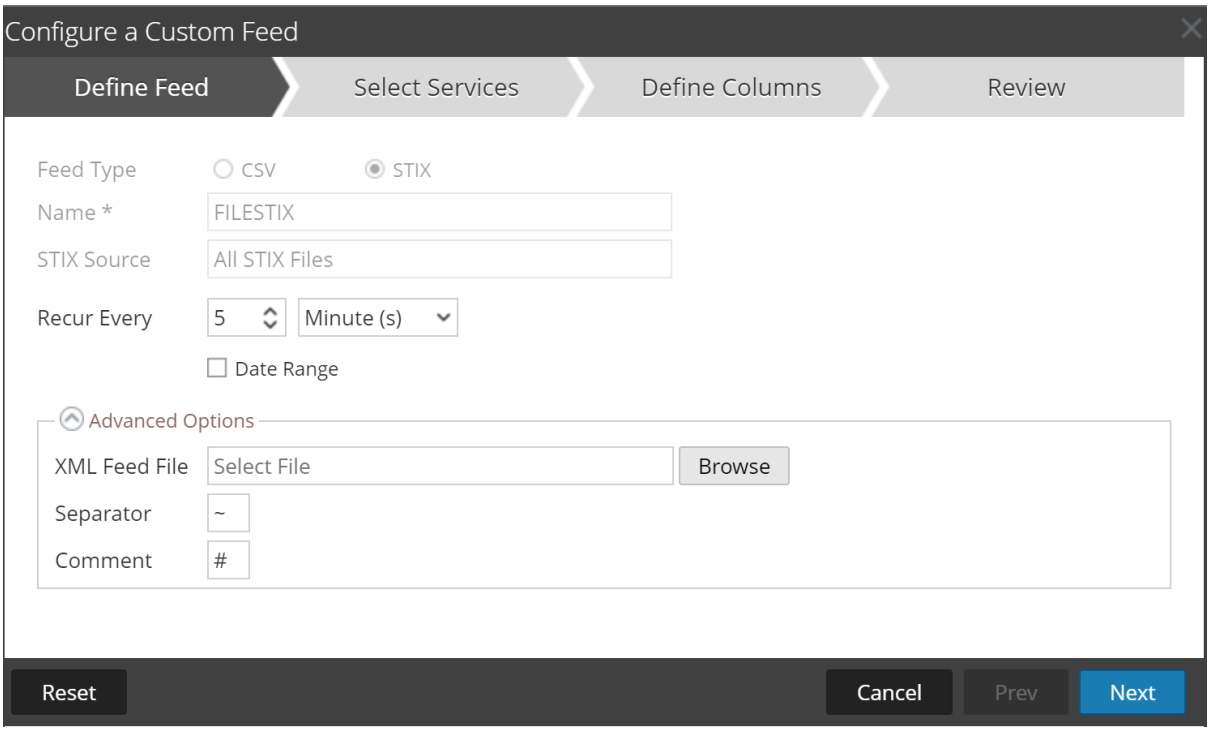
2. In the toolbar, click .

The Setup Feed dialog is displayed.



3. To select the feed type, click **Custom Feed** and **Next**.

   The Configure a Custom Feed wizard is displayed, with the Define Feed dialog open.



4. Enter the following details:

   a. **Feed Type**: Select **STIX**, to define a feed based on a STIX formatted `.xml` file.

   b. **Name**: type the feed name, to define a feed based on STIX formatted `.xml` file.

   c. **STIX Source**:Select a STIX data source from the drop-down which is added in Context Hub.

   d. **Recur Every**: Specify a recurring feed task that executes repeatedly at specified intervals.

   Note: NetWitness verifies the connection to the server, so that NetWitness can check for the latest file automatically before each recurrence.

   e. **Date Range**: Select the checkbox and specify the date range for the feed task to recur.

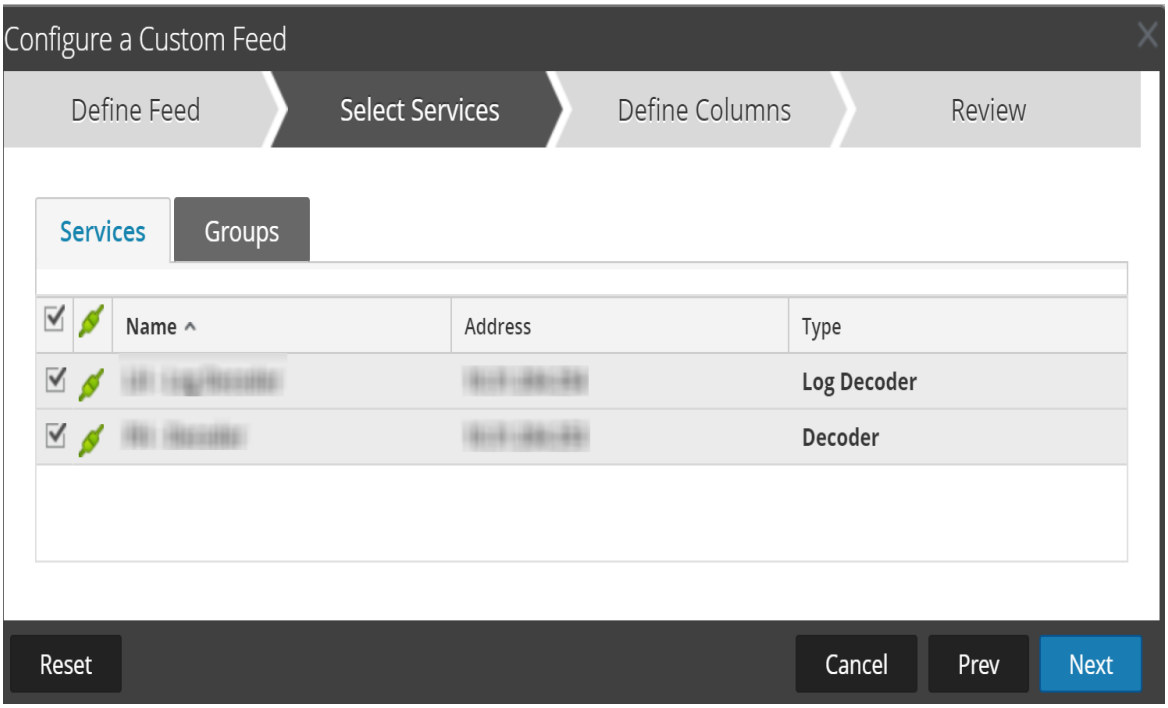
5. (Optional) Select **Advanced Options**,to define a feed based on an XML feed file.

   a. **XML Feed file**: Browse and select an XML feed file from the local file system.

   b. **Separator**: Choose a separator (default is comma).

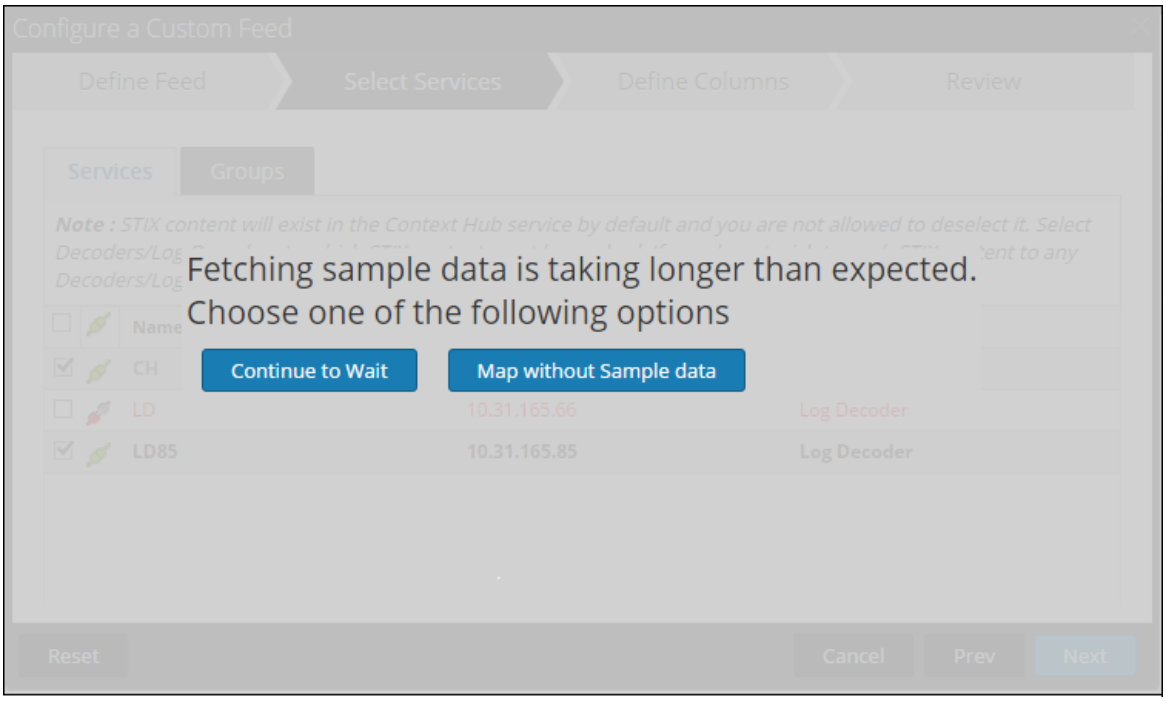
c. **Comment**: Specify the comment characters used in the feed data file (default is #).

6. Click **Next**.
7. The Select Services dialog is displayed. This is an example of the form for a feed based on a feed data file with no feed definition file. If you are defining a feed based on a feed definition file, the Define Columns tab is not needed.



8. To identify services on which to deploy the feed, do one of the following:
   a. Select one or more Decoders and Log Decoders, and click **Next**.
   b. In case of STIX feed, Context Hub will be selected by default and you are not allowed to deselect it. In addition, you can select one or more Decoders and Log Decoders and click **Next** or Click the **Groups** tab and select a group. Click **Next**.

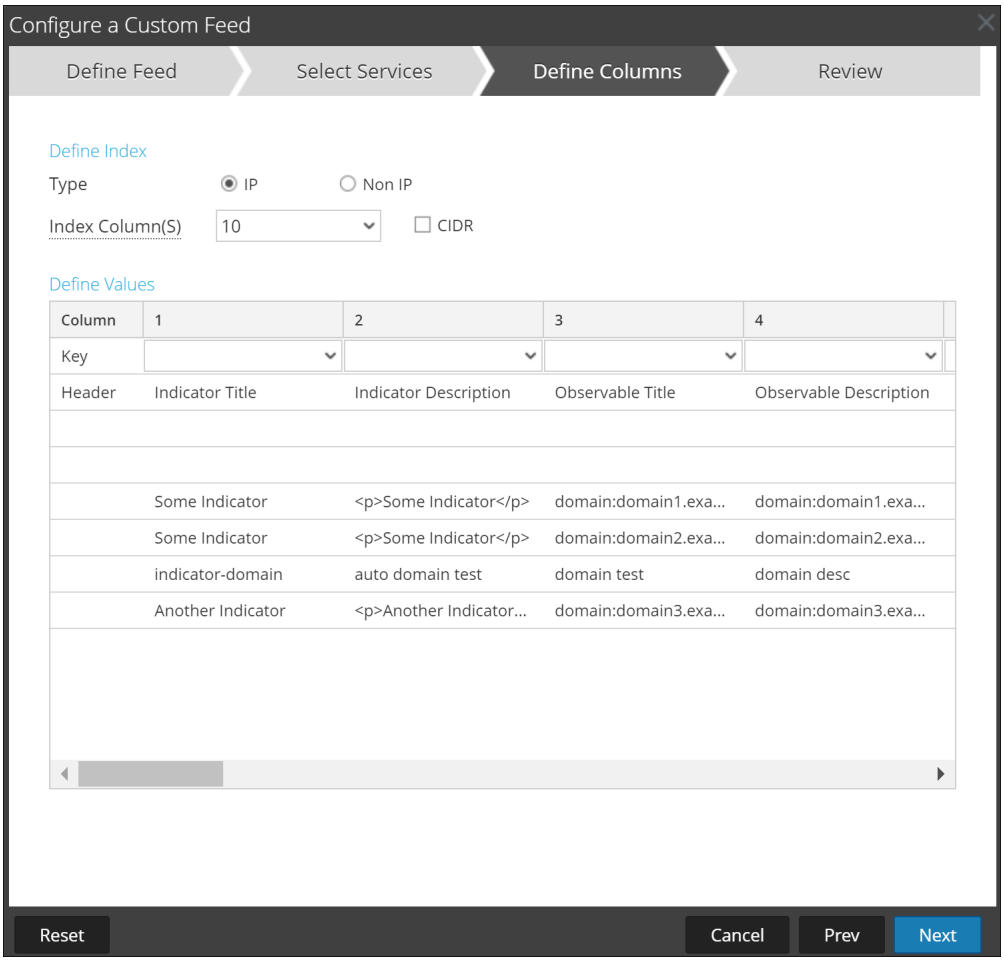
      If the data from the STIX server is large, the following message is displayed:



- If you click **Continue to Wait**, it continues to wait till the sample data is fetched or timeout (10 minutes) whichever is sooner. In case of timeout no sample data is retrieved even after 10 minutes.
- If you click **Map Without Sample data**, the mapping column is displayed without any sample data.

The Define Columns dialog is displayed.

9. To map columns in the Define Columns form:
   a. Define the Index type: **IP**, **IP Range**, or **Non IP**, and select the index column.
   b. (Optional) If the index type is **IP** or **IP Range** and the IP address is in CIDR notation, select **CIDR**.
   c. (Optional) If the index type is **Non IP**, additional settings are displayed. Select the service type and **Callback Keys,** and optionally select the **Truncate Domain** option.

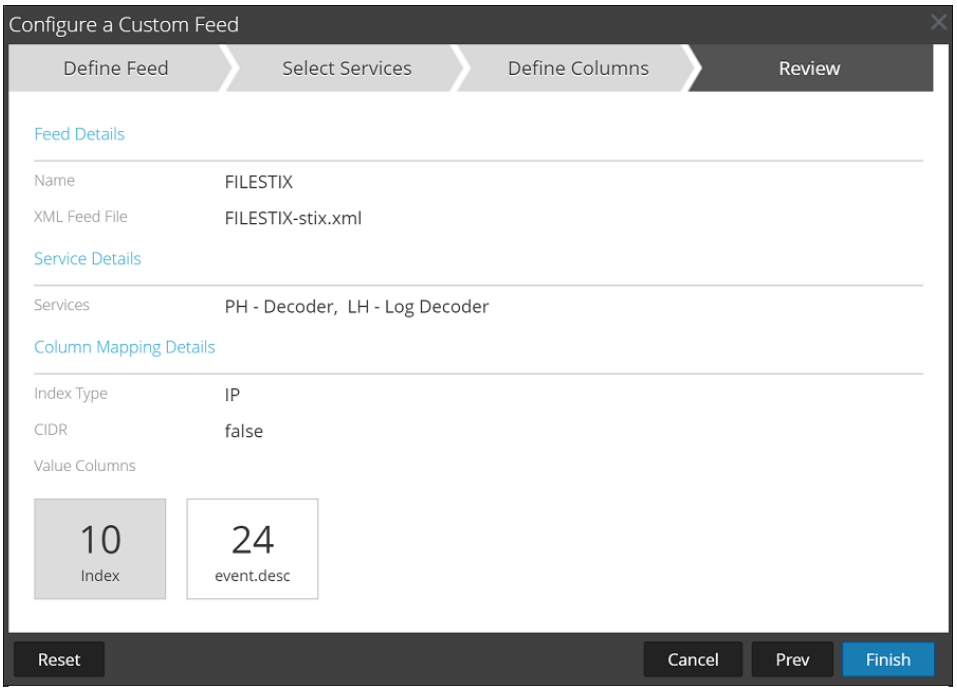


Note:
- If the **Index Type** is Non IP, you can select multiple index columns in the **Index Column(S)**. The values from all the selected columns are merged in the first index column that you selected and the merged values are pushed to the Log Decoder for parsing. For example, in the **Index Column(S)** if you select 2,4,7 as index columns the values from the 2,4 and 7 columns are merged in the column 2 and the values are pushed to Log Decoder for parsing.
- Indexing cannot be done for the columns such as Indicator Title, Indicator Description, Observable Title, Observable Description, as the look up cannot be performed for those columns.

d. Select the language key to apply to the data in each column from the drop-down list. The meta displayed in the drop-down list is based on the meta available for the service define values. You can also add other meta based on advanced expertise.

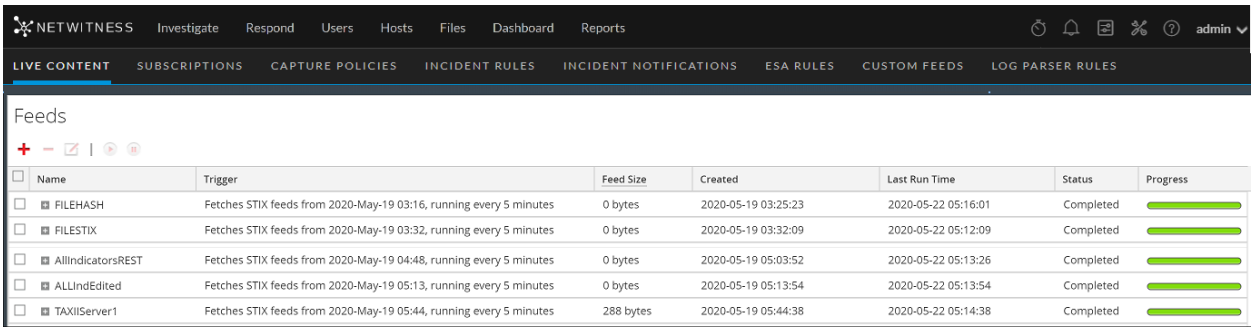
e. Click **Next**.

   The Review dialog is displayed.



10. Anytime before you click **Finish**, you can:
    - Click **Cancel** to close the wizard without saving your feed definition.
    - Click **Reset** to clear the data in the wizard.
    - Click **Next** to display the next dialog (if not viewing the last form).
    - Click **Prev** to display the previous dialog (if not viewing the first form)
11. Review the feed information, and if correct, click **Finish**.

Upon successful creation of the feed definition file, the Create Feed wizard closes, the feed and corresponding token file are listed in the Feed grid, and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.



Note: Health and Wellness raises alerts when the available heap memory of Context Hub server is critically low. If the status of Context Hub server is Unhealthy due to low memory. For more information on how to troubleshoot `OutOfMemoryError` on `Contexthub` Server, refer to "Troubleshooting" in the *Live Services Management Guide*.

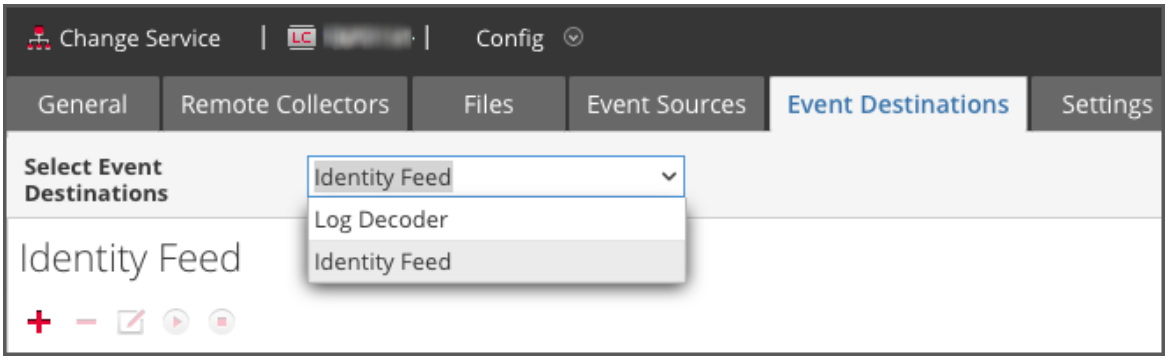
Create an Identity Feed

You can create an Identity feed and populate it to selected Decoders and Log Decoders. In order to create an identity feed, you need to have:

- A Log Collector service with an Identity Feed Event Processor
- A Log Collector service with Windows Collection configured and enabled

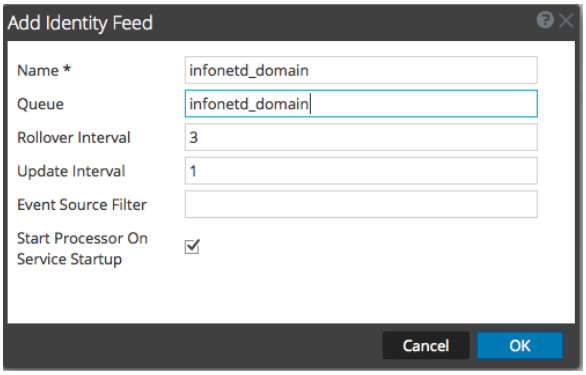
To create an identity feed:

1. Add a destination for the feed.
   a. Go to **(Admin) > Services** and in the **Services**.
   b. In the list of services, select a **Log Collector** service, and select **View > Config**.
   c. Select the **Event Destinations** tab.
   d. In the Select **Event Destinations** field, select **Identity Feed**.



   e. Click and enter a unique name for the feed.

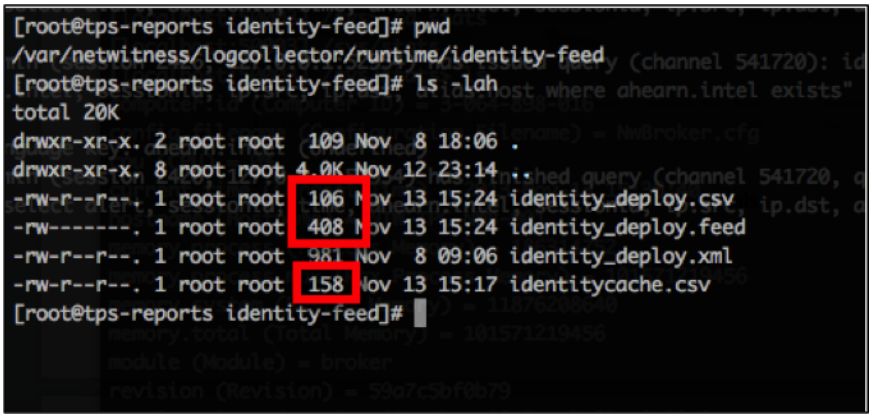
   The Queue name identifies the feed within the Log Collector. Use the name of the feed for the Queue.



   f. Click **OK**.
2. Test generation of messages.

a. Have users log into Windows boxes on the domain to generate the appropriate log messages on the domain controllers for testing.

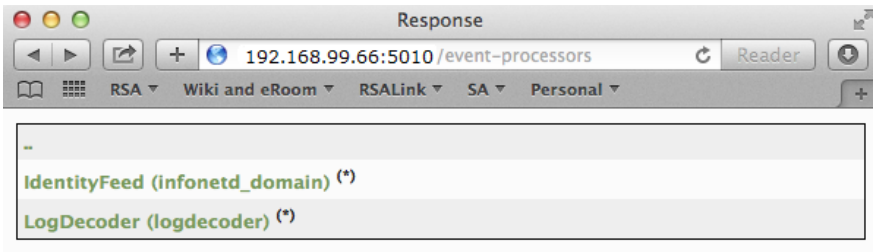
b. Verify that data is written to the feed files. SSH to the Log Decoder/Collector or Virtual Log Collector being configured. Navigate to `/var/netwitness/logcollector/runtime/identity-feed` and verify that the `Identity_deploy` files are getting populated with data.



c. Open up a web browser (Non-Internet Explorer browsers preferred) and log in to the REST interface of the Log Collector. Use administrative credentials when logging in. For example, if the IP address of your Log Collector is 192.168.99.66, the URL would be:

- SSL not enabled: **http://192.168.99.66:50101/event-processors**
- SSL enabled: **https://192.168.99.66:50101/event-processors**

The browser screen should look like this:

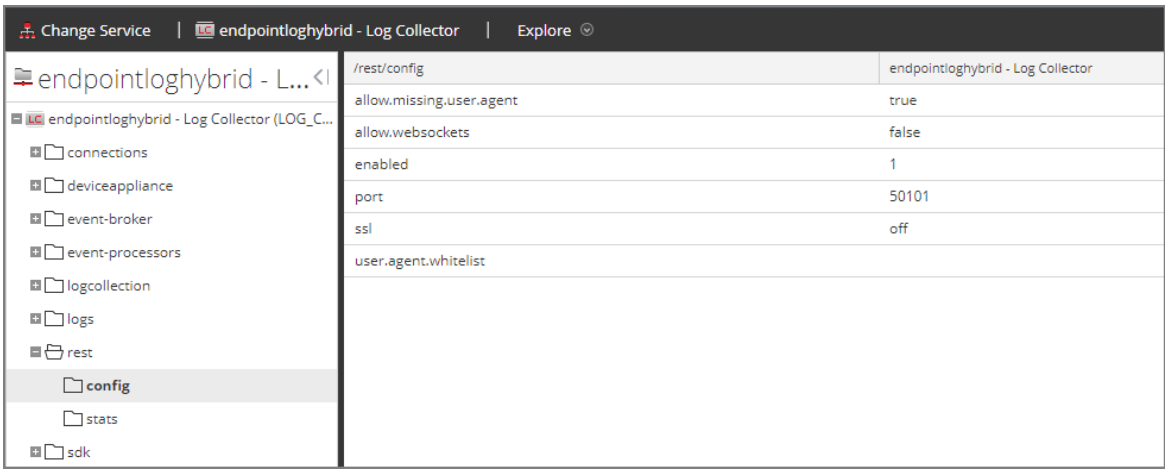


The screen contains the name of the identity feed you created earlier (`infonetd_domain`, in this example).

For the identity feed to function correctly, port 50101 must be active on the Log Collector, and you must determine whether SSL encryption is active.

d. Go to  **(Admin) > Services > <*Log Collector being setup*>**  **> View > Explore**.

e. In the left pane, expand **rest > config**.



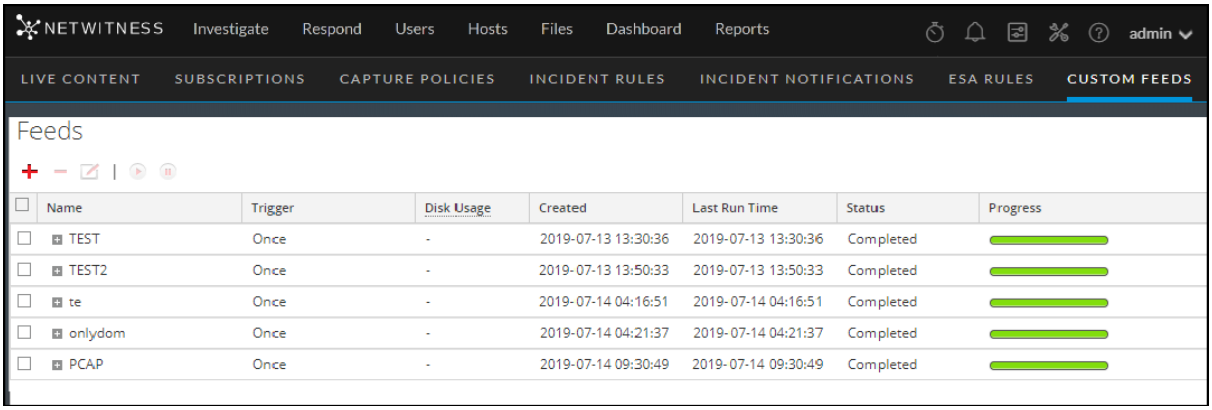
For REST to be active, **enabled** must be set to **1**.

f. Note the value for **ssl**. If SSL should be enabled for your environment, this must be set to **on**.

Note: If you changed the setting for either the **enabled** or **ssl** option you must restart the Log Collector service before moving forward.

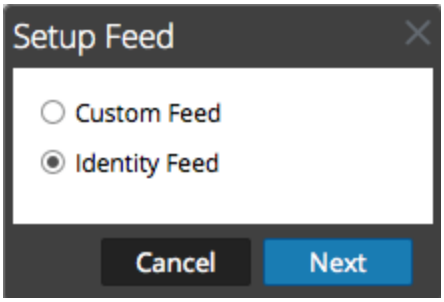
3. Go to **(Configure) > Custom Feeds**.

The Feeds dialog is displayed.



4. In the toolbar, click .

The Setup Feed dialog is displayed.



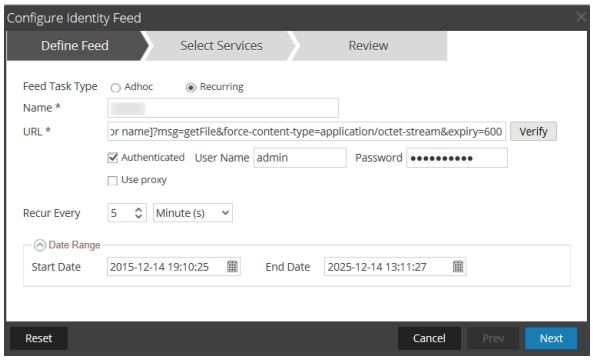
5. Make sure **Identity Feed** is selected and click **Next**.

The Configure Identity Feed panel opens with the **Define Feed** tab displayed.

6. (Conditional) You can create an on-demand or recurring feed.

   - To define an on-demand Identity feed task that executes once, select **Adhoc** in the **Feed Task Type** field, type the feed **Name**, and browse for and open the feed.
   - To define a recurring Identity Feed task that executes on a recurring basis, select **Recurring** in the **Feed Task Type** field.

     The **Define Feed** dialog includes the fields for a recurring feed.



Note: NetWitness verifies the location where the file is stored, so that NetWitness can check for the latest file automatically before each recurrence.

7. Enter a value and verify the URL field.

   a. In the **URL** field, enter the URL where the feed data file is located. This is the REST API interface that was setup earlier. Make sure you have the following information to construct the URL:

      - The IP address of the Log Collector being used to construct the Identity Feed file.
      - The identity queue name, as set in step 2c.
      - Whether or not SSL is enabled on the Log Collector REST port, as set in step 2f.

      You can construct this value as follows:

      - SSL enabled: `https://<LogCollector>:50101/event-processors/<ID Event processor name>?msg=getFile&force-content-type=application/octet-stream&expiry=600`
      - SSL not enabled: `http://<LogCollector>:50101/event-processors/<ID Event processor name>?msg=getFile&force-content-type=application/octet-stream&expiry=600`

      So, using the example from earlier, the complete value that you would enter into this field is as follows:

      ```
      http://192.168.99.66:50101/event-processors/infonetd_
      domain?msg=getFile&force-content-type=application/octet-
      stream&expiry=600?msg=getFile&force-content-type=application/octet-
      stream&expiry=600
      ```

b. For the URL verification to work correctly, it is important that the NetWitness UI server can access the Log Collector's REST API port (50101). This can be tested by going to the NetWitness UI server via SSH. Once there, run the following command:

- SSL enabled: `curl -vk https://<ip of log collector>:50101`
- SSL not enabled: `curl -v http://<ip of log collector>:50101`

If the `curl` command does not connect then there may be a network firewall or routing issue between the NetWitness UI server and the Log Collector.

Example of a bad connection:

```
* About to connect() to 192.168.99.66 port 50105 (#0)
* Trying 192.168.99.66... No route to host
* couldn't connect to host
* Closing connection #0
curl: (7) couldn't connect to host
```

Example of a good connection:

```
* About to connect() to 192.168.99.66 port 50105 (#0)
* Trying 192.168.99.66... connected
* Connected to 192.168.99.66 (192.168.99.66) port 50105 (#0)
> GET / HTTP/1.1
> User-Agent: curl/7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7
NSS/3.19.1 Basic ECC zlib/1.2.3 libidn/1.18 libssh2/1.4.2
> Host: 192.168.99.66:50105
> Accept: */*
>
< HTTP/1.1 401 Unauthorized
< Content-Length: 71
< Connection: Keep-Alive
< Pragma: no-cache
< Expires: -1
< Cache-Control: no-cache, no-store, must-revalidate
< WWW-Authenticate: Basic realm="NetWitness"
< Content-Type: text/xml; charset=utf-8
<
<?xml version="1.0" encoding="utf-8"?>
<error>401 Unauthorized</error>
* Connection #0 to host 192.168.99.66 left intact
* Closing connection #0
```

8. The REST API requires a username and password when attempting to pull the `identity_deploy.csv` file from the Log Collector. This can be any username and password that is available on the service itself. For more information, see the "Services Security View" topic in the *Hosts and Services Guide*.

   To see which accounts are available, go to **(Admin) > Services > <log collector being setup> > Actions > View > Security**.

   Under the Users table, you see all the users that can be used in this step. It is suggested that a separate user account is created specifically for this setup, and is used nowhere else in the environment, for added security. For details, see "Add a User and Assign a Role" in the *System Security and User Management Guide*. (Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

9. To define the recurrence interval, do one of the following:

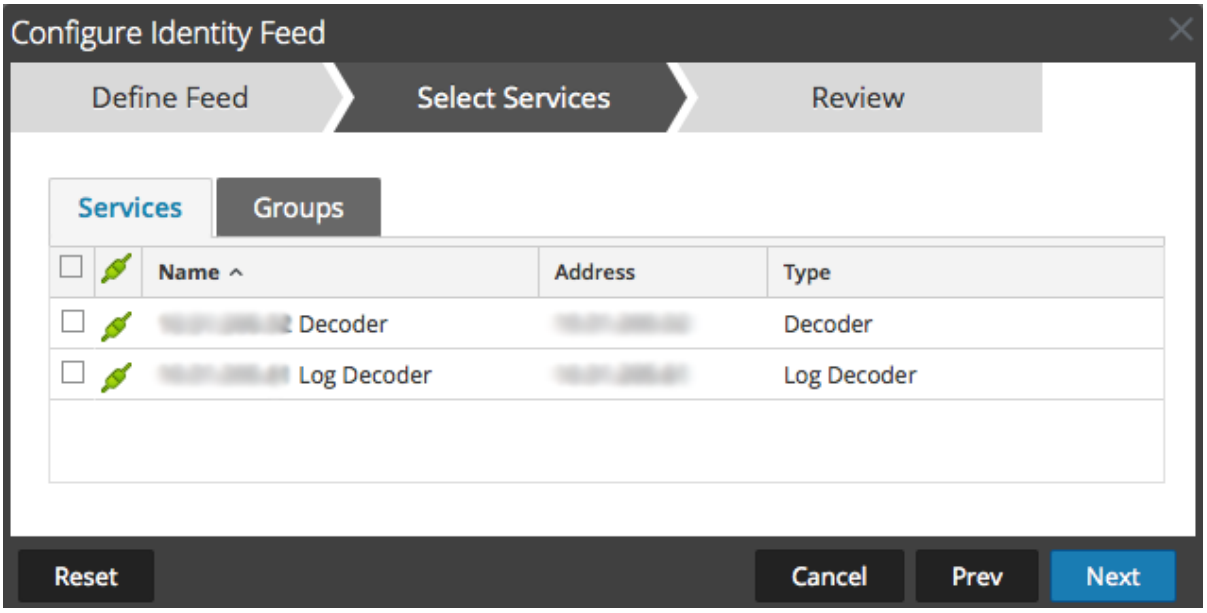
- Specify the number of minutes, hours, or days between recurrences of the feed.
- Enter the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.

10. If using SSL encryption, you need to install the REST API SSL certificate for the Log Collector into the NetWitness UI server. For more information, see Import the SSL Certificate.

    If, after importing the SSL certificate, the verification of the URL still fails, see Cannot Verify Identity Feed URL.

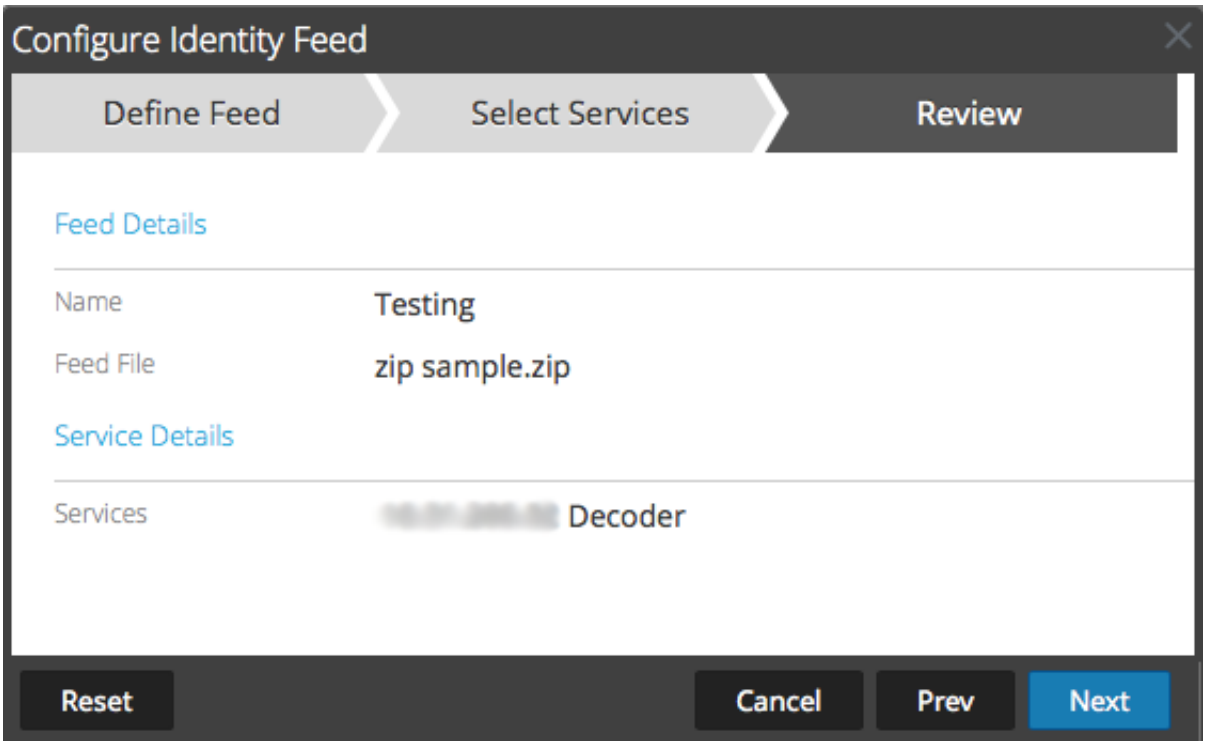
11. Click **Verify** to verify your identity feed configuration before you proceed to the Select Services dialog.
12. Click **Next**.

    The Select Services dialog is displayed.



13. To identify services on which to deploy the feed, select one or more Decoders and Log Decoders and click **Next**.
14. Click the **Groups** tab, select a group, and click **Next**.

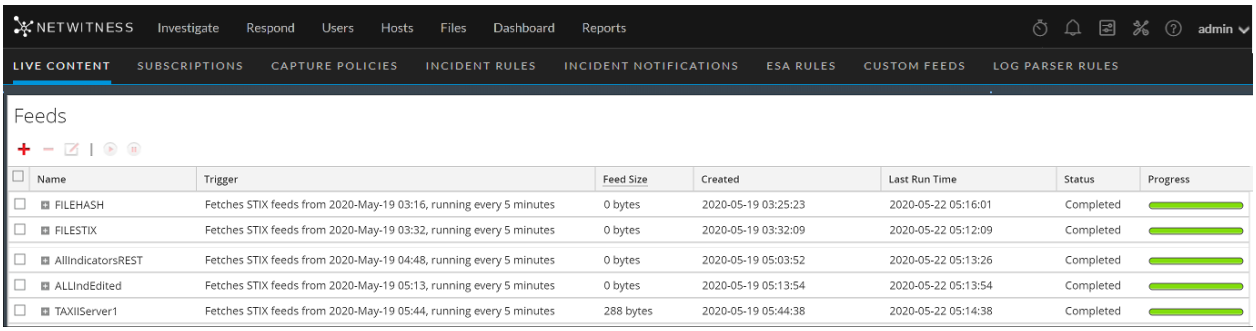
    The Review dialog is displayed.



Note: If a group of devices with Decoders and Log Decoders is used to create recurring or custom feeds and this group is deleted, you can edit the feed and add a new group to the feed.

15. Anytime before you click **Finish**, you can:
    - Click **Cancel** to close the wizard without saving your feed definition.
    - Click **Reset** to clear the data in the wizard.
    - Click **Next** to display the next form (if not viewing the last form).
    - Click **Prev** to display the previous form (if not viewing the first form).
16. Review the feed information, and if correct, click **Finish**.

Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file are listed in the Feed grid and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.



Import the SSL Certificate

If SSL is configured on the Identity feed's Log Collector, follow these steps to import the Log Collector's SSL certificate into the NetWitness UI server key store. If this certificate is not imported, the NetWitness UI server will be unable to pull the Identify feed file from the Log Collector.

1. To pull the SSL certificate off the Log Collector, SSH into the NetWitness UI server and run the following command:

   ```
   echo -n | openssl s_client -connect <HOST>:<PORT> | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/<SERVERNAME>.cert
   ```

   This command saves the SSL certificate to `/tmp/<SERVERNAME>.cert`. For example:

   ```
   echo -n | openssl s_client -connect 192.168.99.66:50101 | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/logcollector.cert
   ```

2. To import the SSL certificate into the NetWitness UI server, SSH into the UI server and run the following command:

   ```
   keytool -importcert -alias <name an alias for the cert> -file <the cert file pathname> -keystore /etc/pki/java/cacerts
   ```

   For example:

   ```
   keytool -importcert -alias logcollector01 -file /tmp/logcollector.cert -keystore /etc/pki/java/cacerts
   ```

3. The system requests a password. Enter the password for the keystore on the NetWitness UI server, not for the jetty keystore. The default password is **changeit**.

4. Restart **jettysrv** to allow jetty to read the new certificate in the store.

Cannot Verify Identity Feed URL

If the Identity feed URL cannot be verified, and you are using SSL, make sure you followed the steps in Import the SSL Certificate.

If there are issues, it is possible that the internal name of the certificate does not match the hostname of the Log Collector. The following procedure checks this.

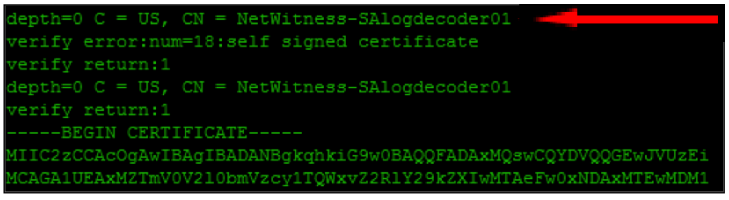
1. SSH to the NetWitness UI server.
2. Run the following command to output the CN name of the SSL cert:

   ```
   echo -n | openssl s_client -connect <log decoder>:50101 | sed -ne '/BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
   ```

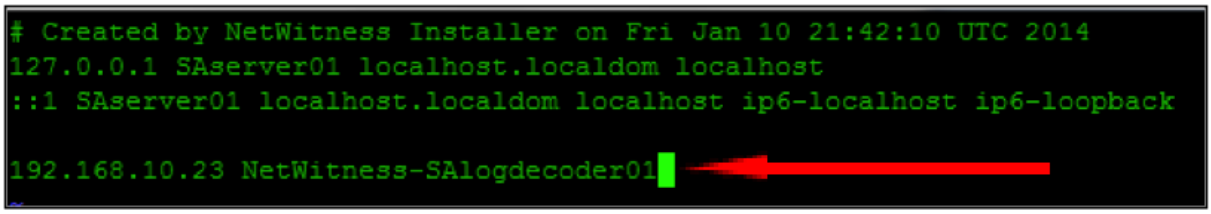
   For example:

   ```
   echo -n | openssl s_client -connect salogdecoder01:50101 | sed -ne '/BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
   ```

3. Retrieve the CN name of the SSL certificate.



4. Edit the `/etc/hosts` file and add the IP address and CN name to the file.



5. Restart the network service on the appliance.
6. Confirm that the name placed in the **/etc/hosts** file is used instead of the FQDN or IP address in the Identity feed URL.
7. Re-verify the Identity feed URL.

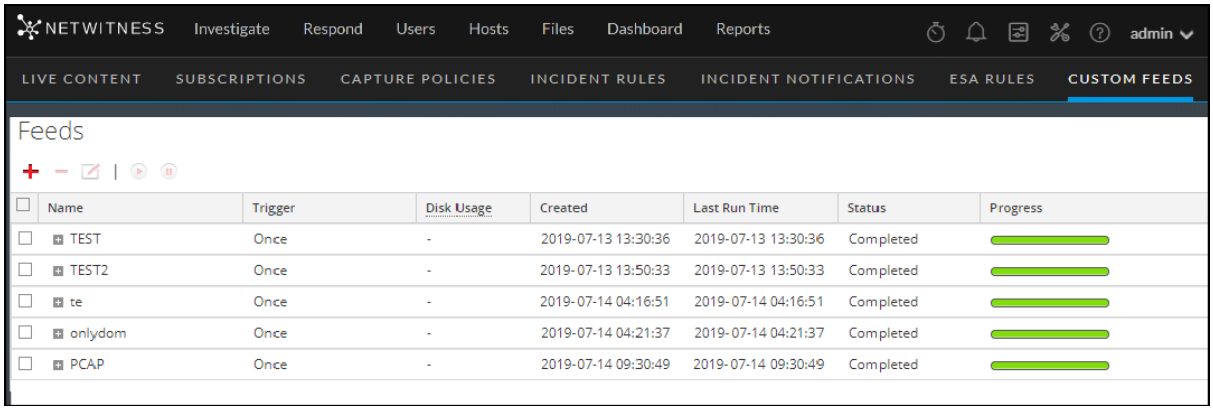
Edit, Upload, or Remove a Feed

You can upload a feed, edit an existing feed, or remove a feed.

To edit an existing feed:

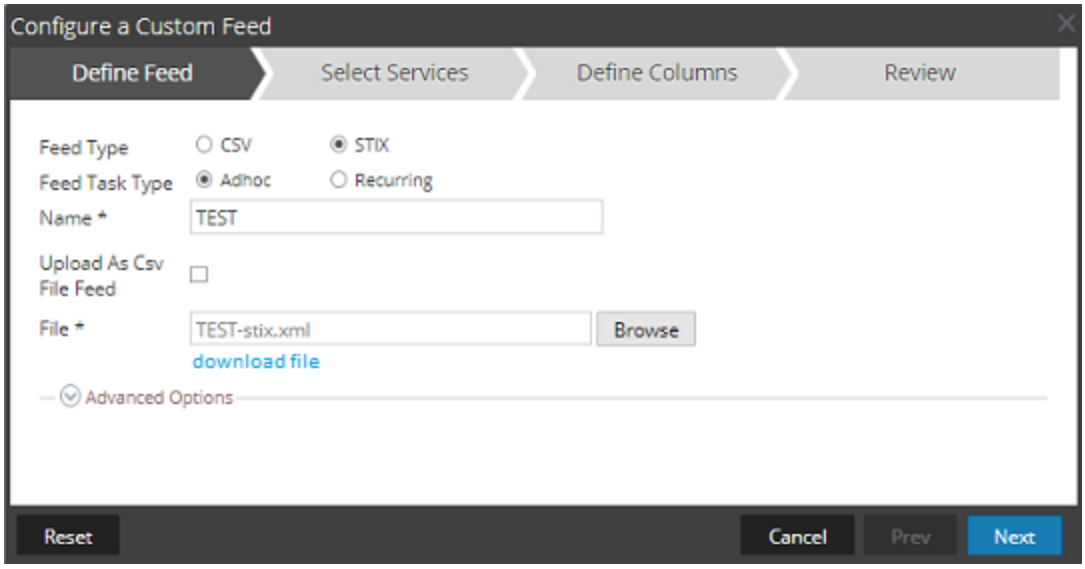
1. Go to (Configure) > **Custom Feeds**.

   The Feeds view is displayed.



2. In the toolbar, select a feed and click .

   The Configure Custom Feed or Configure Identity Feed panel opens in the Custom Feed wizard.



3. If you want to edit the feed file:

   a. Click **download file**.

      For an Identity feed, the .zip file is downloaded. For a custom feed, the .csv or .xml file is downloaded to your local file system. For a STIX feed, the .xml file is downloaded to your local file system.

b. Edit and save the file.

c. In the **Define Feed** tab, browse for and open the edited file.

4. Edit any other parameters in the **Define Feed** tab, **Select Services** tab, and **Define Columns** tab that apply to the type of feed.
5. Anytime before you click **Finish**, you can:
   - Click **Cancel** to close the wizard without saving your changes.
   - Click **Reset** to clear the data in the wizard.
   - Click **Next** to display the next form (if not viewing the last form).
   - Click **Prev** to display the previous form (if not viewing the first form).
6. In the **Review** tab, review the feed information, and if correct, click **Finish**.

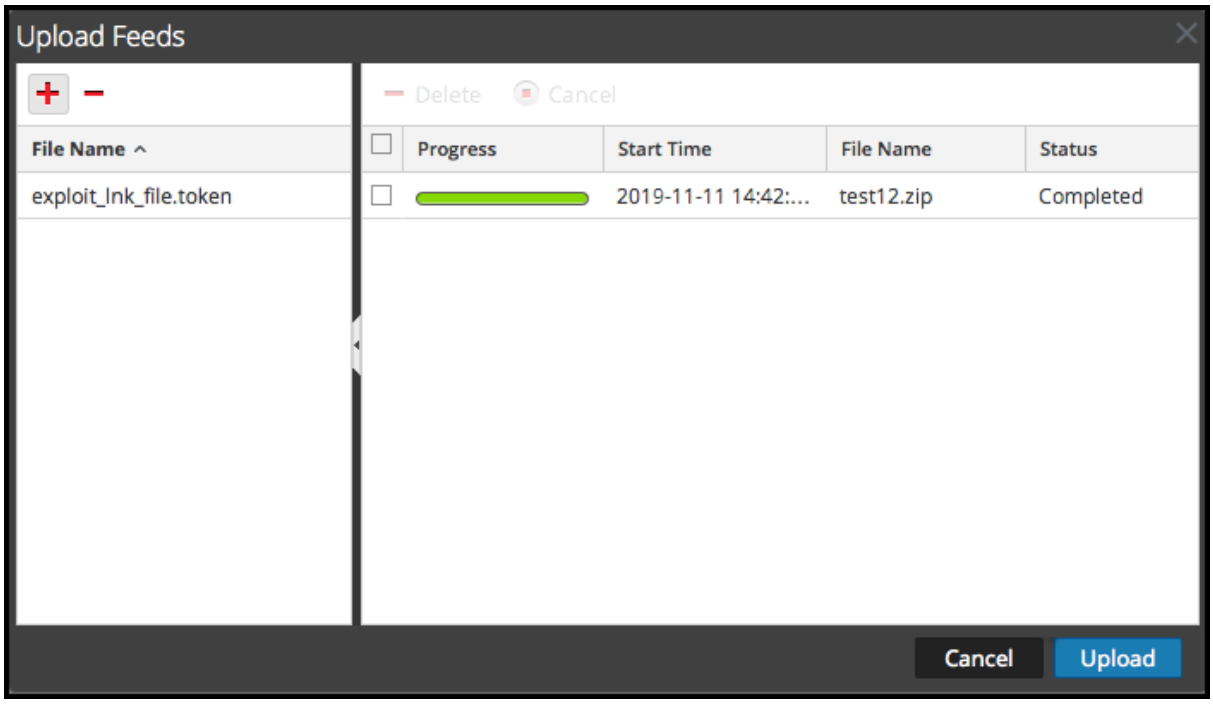
The feed is recreated with the updated file and new feed specifications. The feed is added to the Feeds list and progress bar tracks completion. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file is listed in the Feeds list. You can expand or collapse the entry to see how many services are included, and which services are successful.

To upload a feed to a Decoder or Log Decoder:

1. Go to  (Admin) > **Services**.
2. Select a Decoder or Log Decoder service and click  **> View > Config**.

   The Services Config view is displayed with the General tab open.
3. Select the **Feeds** tab.
4. In the Feeds tab toolbar, click .

   The Upload Feeds dialog is displayed.



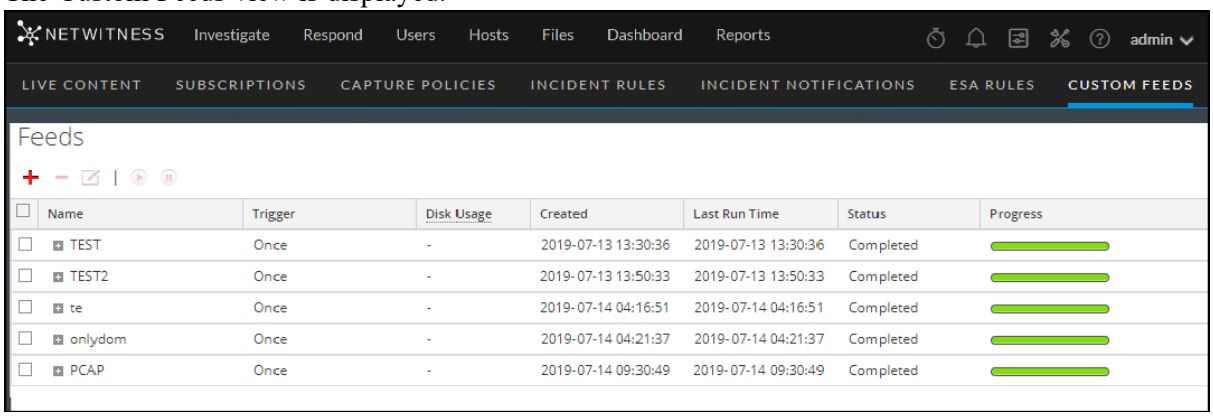
5. In the **File** grid, click + and select a feed file. Supported files are *.feed, *.token, and *.filter.
6. Select the feed file from the **File** list and click **Upload**.

   The Upload Job list is updated to show the progress and status of the uploaded feed.

To remove a feed:

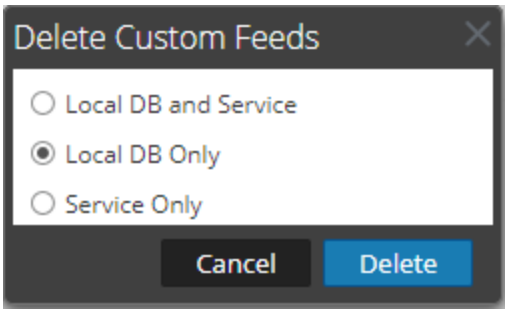
1. Go to (Configure) > **Custom Feeds**.

   The Custom Feeds view is displayed.



2. In the toolbar, select a feed and click −.

   The Delete Custom Feeds dialog is displayed.



You can select one of the following options to delete the feed:

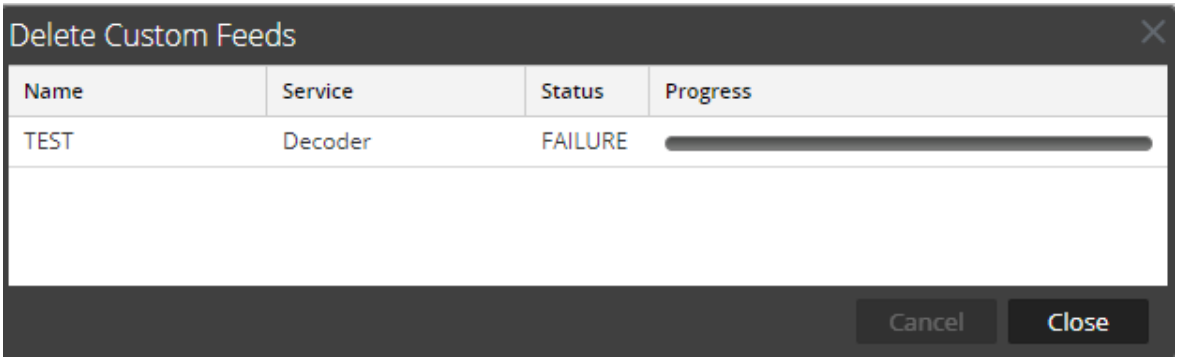
- If you choose to delete the feed from **Local DB and Service**, the feed is deleted from both the service and the local NetWitness box. The deleted feed will no longer be seen on the NetWitness user interface.
- If you choose to delete the feed from **Local DB Only**, the feed is deleted from the local NetWitness box. The deleted feed will not be seen on the NetWitness user interface; however, the last deployed version of the feeds will be present on the service. The undeployed feeds will be deleted forever.
- If you choose to delete the feed from **Service Only**, the feed is deleted from the service. The deleted feed will appear on the NetWitness user interface and can be deployed again.

3. Select where you want to delete the feed and click **Delete**.

   A warning dialog is displayed.

4. Click **yes** to confirm that you want to delete the feed from the select areas.

   - If you chose to delete the feed from the **Local DB Only**, the feed is deleted.
   - If you chose to delete the feed from the **Local DB and Service** or **Service Only**, the Delete Custom Feeds view is displayed showing the progress of the deletion on the service.



Create Custom Meta Keys Using a Custom Feed

This topic provides information on how to add custom meta keys using a custom feed in the Log Decoder, and highlights the configuration changes to reflect the custom meta keys in the Concentrator, ESA, Archiver, Warehouse Connector, and Reporting Engine schema. You can create custom meta keys to retrieve data, to investigate and analyze the logs and packets. Custom meta keys enable you to add an enrichment context for the log and packet data.

Here is an example of creating a custom meta key in the Log Decoder. In this scenario, an organization wants to track the location of an asset such as a printer. So, a custom meta key **source location** is introduced, which indicates the location of the asset, for example Printer1, which is located in the 'Fifth Floor A wing'.

Note: Custom meta keys can be created in the Decoder as well. Select the `index-decoder-custom.xml` file when you create a custom meta key in the Decoder.

Add a Custom Meta Key in the Log Decoder

To add custom meta keys using custom feed:

1. Go to  (Admin) > **Services** .
2. Select a Log Decoder service and click  > **View** > **Config** > **Files** tab > **index-logdecoder-custom.xml.**

```
<Language>
  <?xml version="1.0" encoding="utf-8"?>
  <Language level="IndexNone" defaultAction="Auto">
  <!-- Reserved Meta key for Feed -->
  <Key description="Source Location" level="IndexNone"
name="location.src" format="Text"/>
 </Language>
 name = Name of the key (max is 16 chars)
```

3. Restart the Log Decoder service. In the Services view, click  > **Restart**.

Note: Following are the NetWitness reserved keywords. Do not use these keywords as custom meta keys as it can cause display error in the NetWitness user interface due to CSS conflicts referencing to the same keywords.
risky
safe
unsafe
ecat
black
ecat-text
bold
machine-name
critical
item
row
list
ioc-info
status-info
details

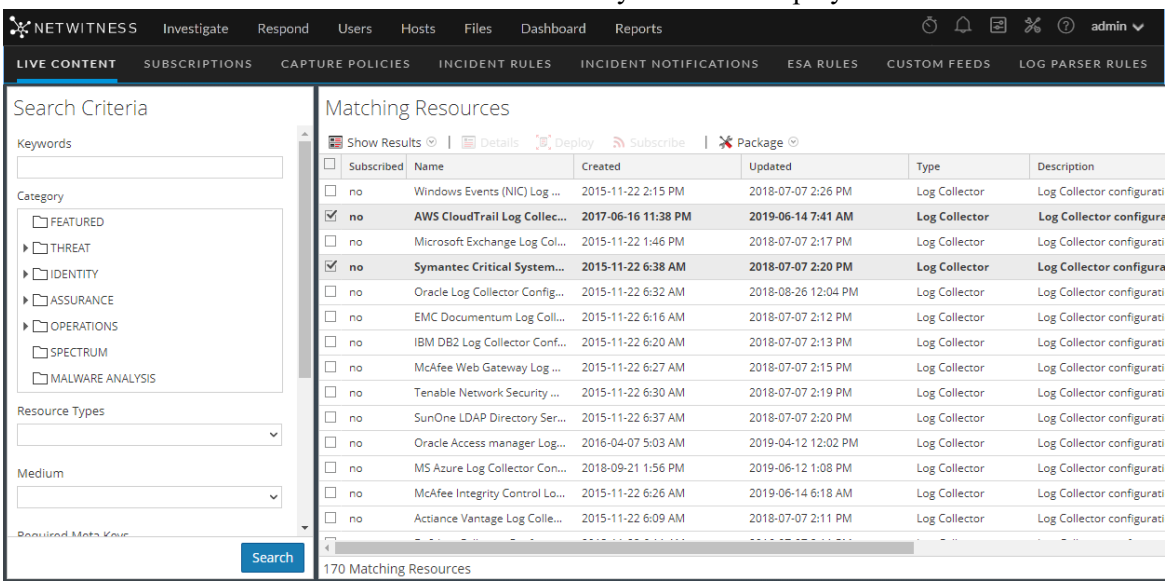
```
score
cards
intro-subject
context-intro
im-risk-score-type
im-priority-type
im-score-type
incident-details
cs-name
im-score-critical
im-score-high
im-score-medium
im-score-low
im-score-lowest
im-score-neutral
alert-score-high
ecat-score
loading-msg
content-loading
live-title
live-text
```

Deploy a Log Decoder Feed in Live

To deploy the feed in the live environment:

1. Go to  (Configure) > **Live Content**
2. Select a group of resources, or a previously created resource package. To select a resource or group of resources:
   a. In the **Live Search View,** browse Live resources (for example, search for the **Log Collector** resource Type).
   b. In the **Matching Resources** panel, select **Show Results > Grid**.

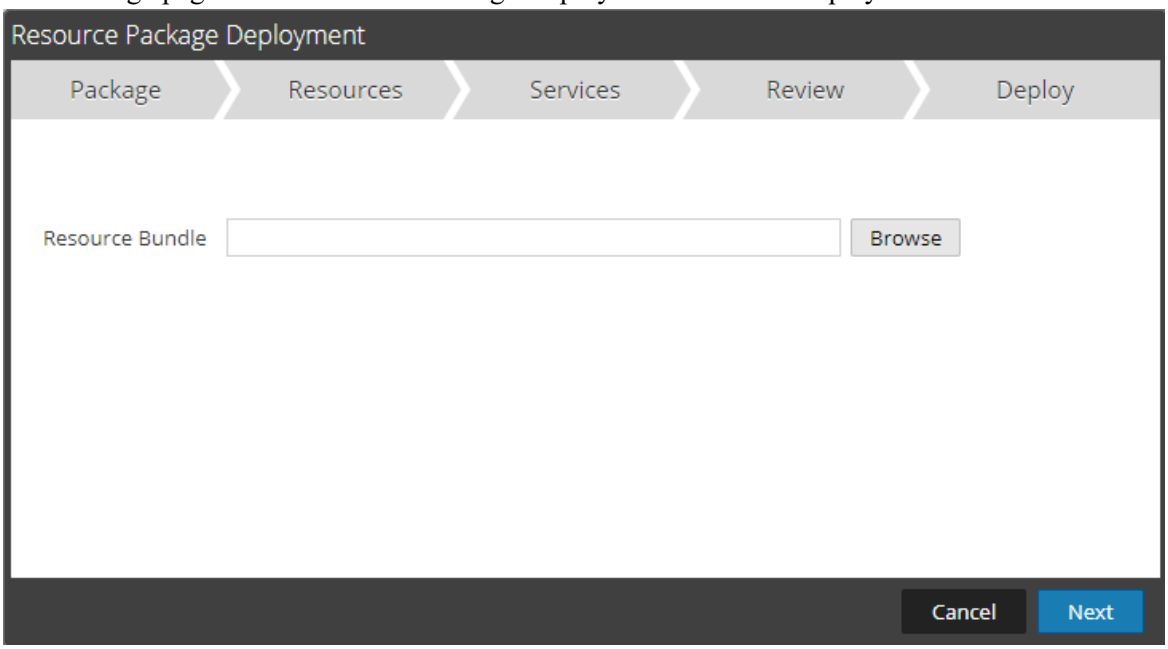
c. Select the checkbox to the left of the resources that you want to deploy.



d. In the Matching Resources toolbar, click Deploy.

3. To select a resource package to deploy:

a. In the **Live Search** view **- Matching Resources** toolbar, select **Package** > **Deploy** :
The Package page of the Resource Package Deployment wizard is displayed.



b. Click **Browse** and select a package from your network.

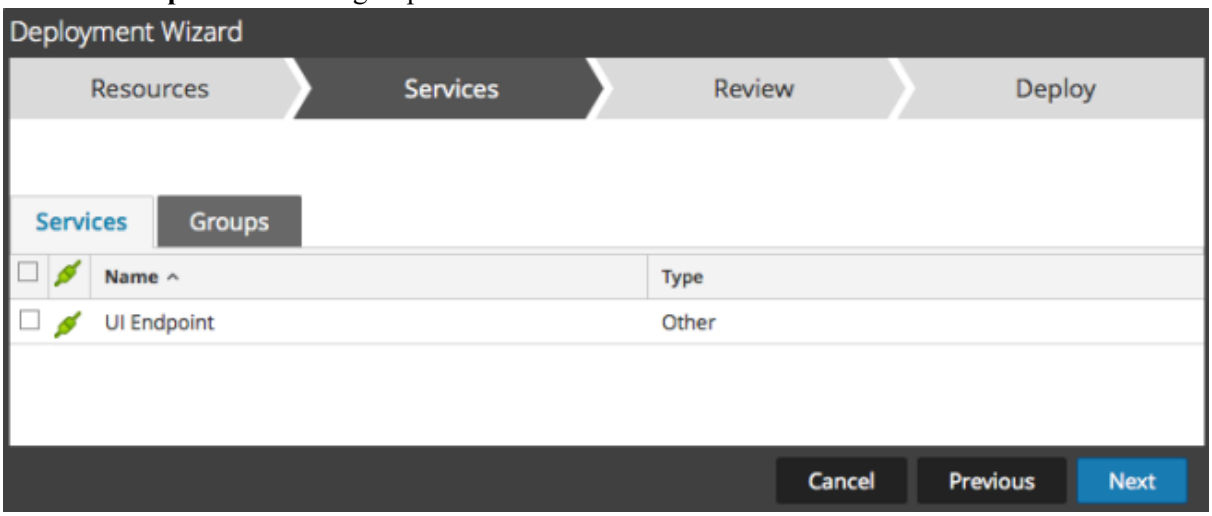
c. Click **Open.**
At this point, whether you are deploying a package or a group of resources, the Deployment Wizard opens, and the Resources page is displayed.

3. Click **Next**.
The **Services** page is displayed. It has two tabs, **Services** and **Groups**, which provide a list of

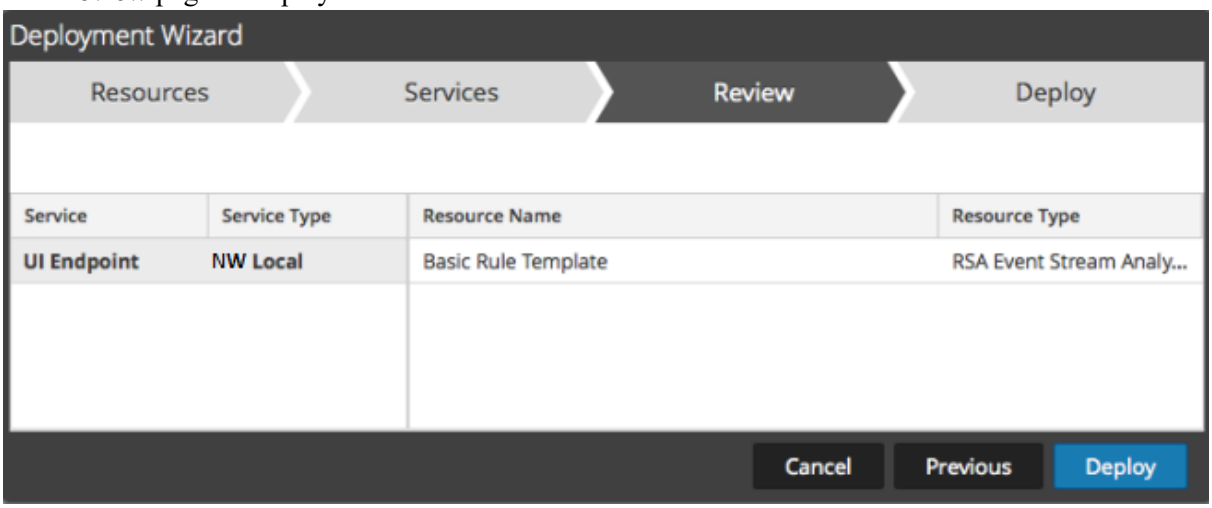
services and service groups that are configured in the (Admin) > Services view. The columns are a subset of the columns available in the Services view.

Note: The Live server is "smart" about deploying resources to Services. For example, it does not deploy resources that have a Medium of packets to any Log Decoders. This means that only applicable content resources are deployed to each Service.

4. Select the services to which you want to deploy the content. You can select any combination of services and service groups.
   Use the **Services** tab to select individual services, list of services and service groups that are configured in the Admin Services view.
   Use the **Groups** tab to select groups of services.



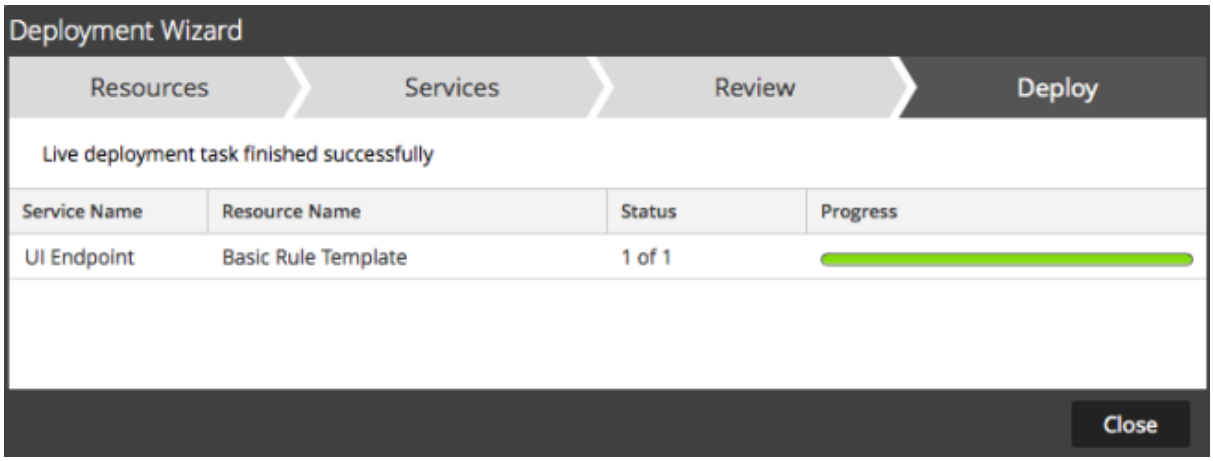
5. Click **Next**.
   The **Review** page is displayed.



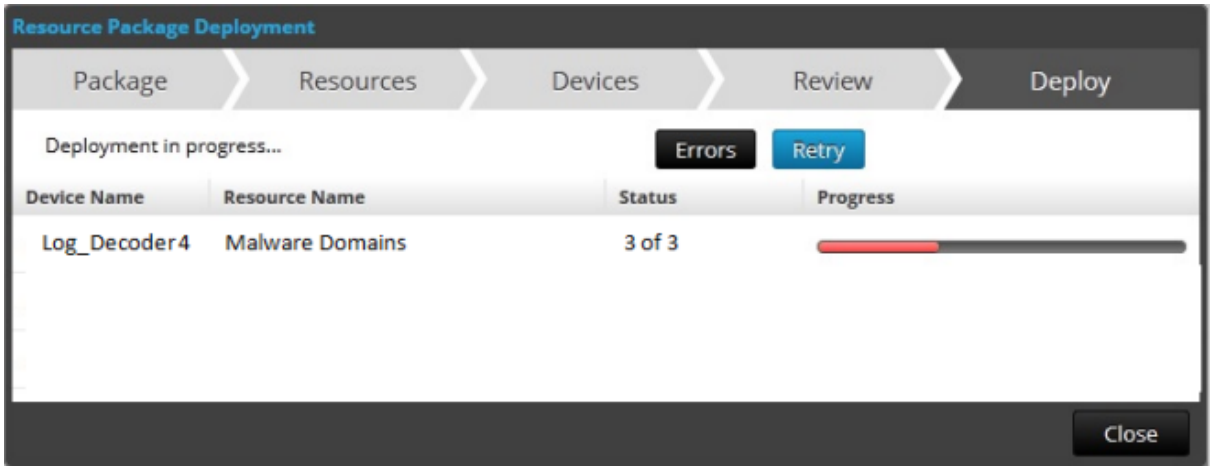
   Make sure that you have selected correct resources and the services to which you want to deploy them.

6. Click **Deploy**.
   The **Deploy** page is displayed. The Progress bar turns green when you have successfully deployed

the resources to the selected services.



If you try to deploy resources and services that are not compatible, NetWitness displays the Errors and Retry buttons, which you click to review the errors and retry the deployment.



7. Click **Close**.

Note: The Source IP should be indexed by selecting the type as 'IP' as the ip.src. and ip.dst are in IPv4 format.

In this scenario, a custom meta key location.src (location source) is added by indexing the hostname (alias.host). In this example, the printer hostname are populated in meta key 'alias.host'. Select **alias.host** as callback key, and index type as 'Non IP' in the Feed Wizard as shown below. In the Define Values section, select the custom meta key from the drop down menu.

Add the Custom Meta Key Entry in the Concentrator Custom Index file

To add the custom meta key entry in the Concentrator custom index file:

1. Go to (Admin) > **Services** > **Concentrator**.
2. Click > **View** > **Config** > **Files** tab > **index-concentrator-custom.xml.**

3. Add the custom meta key entry in the Concentrator index file.

```
<Language>
 <?xml version="1.0" encoding="utf-8"?>
 <Language level="IndexNone" defaultAction="Auto">
  <!-- Reserved Meta key for Feed -->
  <Key description="Source Location"  level="IndexValues"
name="location.src" format="Text"  valueMax="10000"
defaultAction="Open"/>
 </Language>
```

4. To restart the Concentrator service, in the Services view, click > **Restart**.

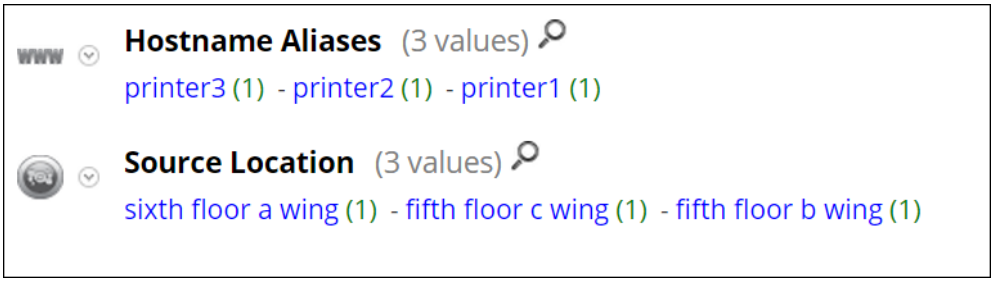
Note: In case of the Broker, the Broker derives its index from the Concentrator from which it aggregates. So you do not need to create custom meta in the Broker. If you have not indexed the meta key in the Concentrator, the Broker does not display the meta key in Investigate.

Investigate a Custom Meta Key

Note: You have to log out and log back in to the NetWitness User Interface in order to view the custom meta key in Investigate.

To investigate a custom meta key:

1. Go to **Investigate**.
   The Investigate dialog, which provides services, is displayed.
2. Select a Concentrator service, and click **Navigate**.



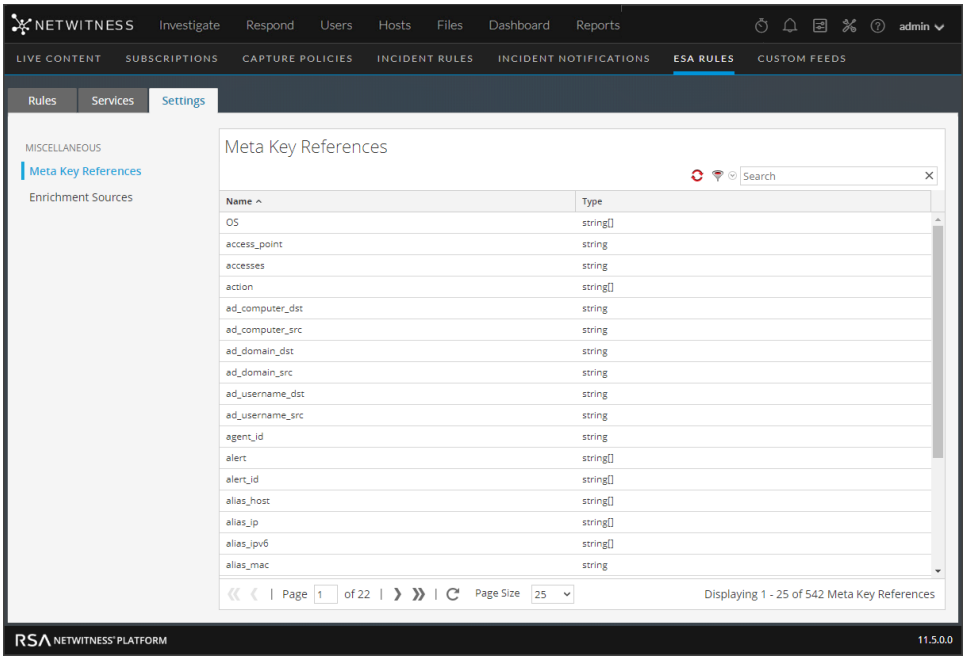
Additional Procedures

The following procedures must be executed if you have Warehouse Connector, Archiver, Reporting Engine, and ESA configured.

Verify the Custom Meta Keys on ESA

After you add custom meta keys on the Concentrator, you can verify that your meta keys are updated on ESA.

1. Go to (Configure) > **ESA RULES** > **Settings** tab.



2. In the Meta Key References, click the Meta Re-Sync (Refresh) icon ( ).

3. Verify that the custom meta keys appear on ESA. If you do not see the meta keys, you may need to restart the Concentrator.

Update the Schema in Archiver

If you want to configure the Archiver, using the new custom meta keys, you need to update the Archiver schema in the Reporting Engine. To update the Archiver schema in Reporting Engine:

1. Go to **(Admin)** > **Services** > **Archiver**.
2. Select > **View** > **Config** > **Files** > **index-archiver-custom.xml**.
3. Add the custom meta key entry in the Archiver index file.

```
<Language>
  <?xml version="1.0" encoding="utf-8"?>
  <Language level="IndexNone" defaultAction="Auto">
  <!-- Reserved Meta key for Feed -->
  <Key description="Source Location" level="IndexValues"
name="location.src" format="Text"
  valueMax="10000" defaultAction="Open"/>
 </Language>
```

4. To restart the Archiver service, click > **Restart.**
   The Archiver schema is updated with the custom meta key.

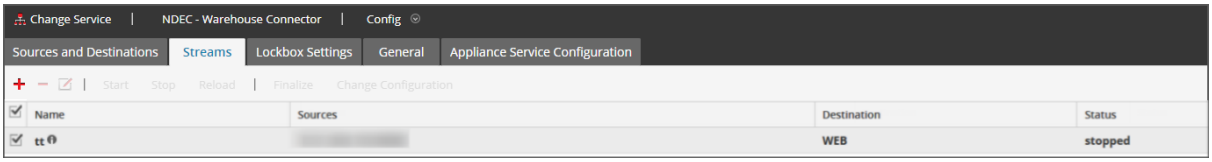
Update the Schema in Warehouse Connector

If you want to configure the Warehouse Connector with custom metadata and use it in a Warehouse Connector report then you need to update the Warehouse Connector schema in the Reporting Engine.

If the Log Decoder or Decoder, where the custom meta key is added, is one of the sources in the Warehouse Connector stream, you need to update the schema in the Warehouse Connector.

To update the Warehouse Connector schema in the Reporting Engine:

1. Go to (Admin) > **Services** > **Warehouse Connector**.
2. Click > **View** > **Config**.
   The Services Config view of Warehouse Connector is displayed.
3. Click the **Streams** tab.
4. Select the stream and then click **Reload**.
   The Warehouse Connector pulls the schema from the downstream devices (Log Decoder/Decoder).



For more information on streams, see "Configure Streams" in the *Warehouse Connector Configuration Guide*.

Update the Schema in Reporting Engine

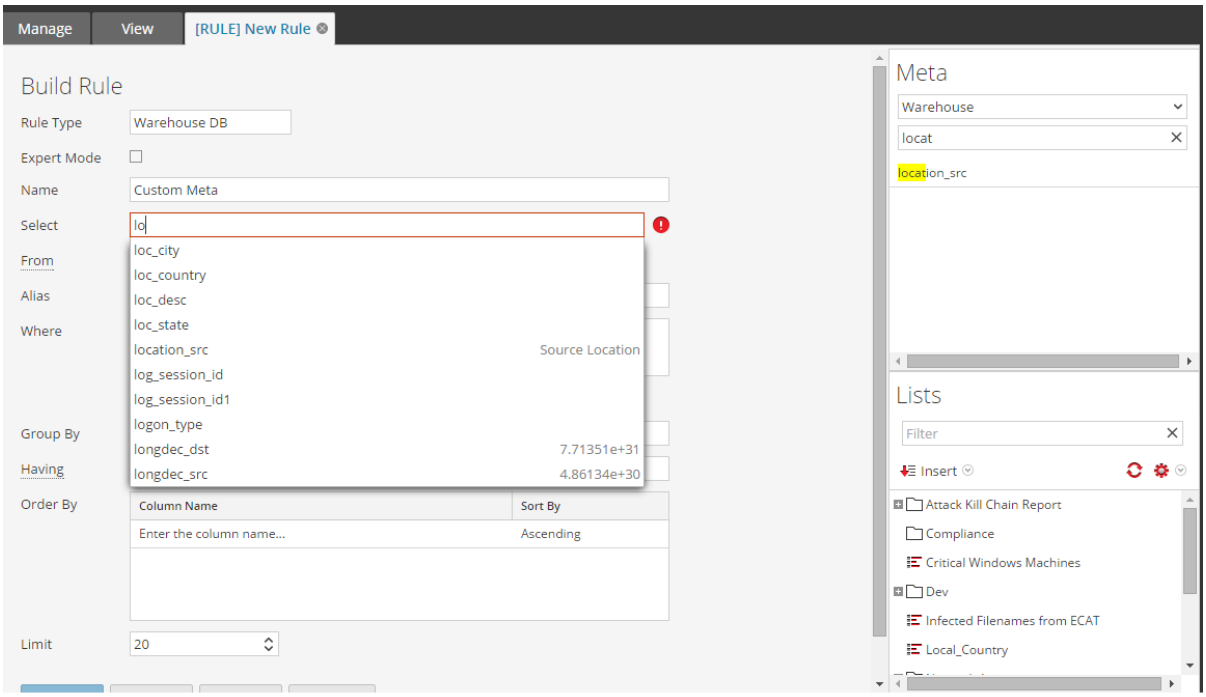
To update the schema in Reporting Engine:

1. Go to **(Admin)** > **Services** > **Reporting Engine.**
2. Click > **Restart.**

Note: Restart the Reporting Engine or wait for thirty minutes for the schema to be updated.

To view the custom meta key:

1. Navigate to **Monitor** > **Reports** > **Rules**.
2. In the toolbar, click .
3. Select **Warehouse DB**.
4. In the **Build Rule** tab, search for the custom meta from the right panel.
   The custom meta key is displayed.


Manage
View
[RULE] New Rule
Build Rule
Rule Type
Warehouse DB
Expert Mode
Name
Custom Meta
Select
lo
From
Alias
Where
Group By
Having
Order By
Limit
20
loc_city
loc_country
loc_desc
loc_state
location_src
Source Location
log_session_id
log_session_id1
logon_type
longdec_dst
7.71351e+31
longdec_src
4.86134e+30
Column Name
Sort By
Enter the column name...
Ascending
Meta
Warehouse
locat
location_src
Lists
Filter
Insert
Attack Kill Chain Report
Compliance
Critical Windows Machines
Dev
Infected Filenames from ECAT
Local_Country

Decoder and Log Decoder Additional Procedures

This topic explains the additional procedures an administrator could choose to follow which are not essential for the configuration of the Decoder or Log Decoder.

Topics

- Configure High Speed Packet Capture Capability (Version 11.6 and Later)
- Configure a Log Decoder to Accept Protobuf
- Configure Session Split Timeouts
- Configure Syslog Forwarding to Destination
- Configure Transaction Handling on a Decoder
- Configure Data Export
- Decrypt Incoming Packets
- Edit Decoder System Configuration
- Enable CPU Usage Statistics for Installed Content
- Enable Parser Mappings
- Enable or Disable Lua and Flex Parsing Systems
- Map IP Address to Service Type for Log Parsing
- Map an IP Address to a Time Zone
- Obtain Log Files from a Pre-11.0 Log Decoder
- Upload a Log File to a Log Decoder
- Upload a Packet Capture File

Configure High Speed Packet Capture Capability (Version 11.6 and Later)

This topic guides administrators on how to tune a Network Decoder specifically for high speed packet capture using NetWitness 11.6. This applies when capturing packets on 10G, 40G, or higher speed interface cards.

IMPORTANT: Packet capture at high speeds requires careful configuration and pushes the Decoder hardware to its limits, so you must read this entire topic while implementing a high speed capture solution.

RSA NetWitness provides support for high-speed collection on the Decoder. You can capture network packet data from higher speed networks and optimize your Network Decoder to capture network traffic up to 40 Gbit/s. These network capture speeds can be easily achieved on a Decoder but make a note that not all NetWitness Platform features operate at these rates. The actual throughput of Decoder depends on the following factors:

- **Amount of packets filtered vs. the amount of packets retained**
  It is important to filter the traffic that enters the Decoder to minimize packet drops. You can use Network Rules or a BPF Rule to filter the traffic before it enters advanced features such as stream reassembly and deep packet inspection. For more information, see Configure Network Rules and (Optional) Configure System-Level (BPF) Packet Filtering.
- **Average packet size**
  Decoder capturing small packets, for example 64 byte packets, can decrease the overall throughput of the Decoder. If average packet sizes are small, the line rate of Decoder ingest will fall significantly below the network link speed. It occurs due to the amount of per-packet overhead between each Ethernet frame.
- **Shared system resources**
  If the Decoder host is also hosting other applications or add-on features that manipulate data and perform other analysis, it can lower the overall throughput of the Decoder. Even if these add-on applications are not utilizing the CPU, these will utilize the shared resources like memory bandwidth and CPU interconnect bandwidth.

Enhancements that facilitate capture in these high speed environments include the following:

- Utilization of the DPDK capture drivers to leverage the commodity 10G and 40G Intel NIC cards for high-speed capture environments. For more information on how to configure DPDK, see (Optional) Data Plane Development Kit Packet Capture.
- Introduction of `assembler.parse.valve` configuration, which automatically disables application parsers when certain thresholds are exceeded, to limit risk of packet loss. When the application parsers are disabled, network layer parsers are still active. When stats fall below exceeded thresholds, application parsers are automatically re-enabled.
- Utilization of Berkeley Packet Filters (BPF) in 10G environments. For more information, see (Optional) Configure System-Level (BPF) Packet Filtering in Configure Capture Settings.
- Decoder will utilize the symmetric RSS to distribute capture loads among multiple CPU cores. For more information see, "*Utilizing Receive Side Scaling with DPDK*" in (Optional) Data Plane Development Kit Packet Capture
- Decoder will distribute the session reassembly work among multiple CPU cores as well, provided that the packets fed into each assembly are naturally segregated using the RSS features or as a result of being captured on different physical interfaces.

Note: Beginning with 11.5, the Network Decoder can capture from multiple interfaces simultaneously. This functionality allows Network Decoders to capture from multiple physical Network Interface Cards (NICs) while leveraging the same network rules, application rules, and parsers for each NIC. You can use this feature in all capture environments. For more information, see (Optional) Multiple Adapter Packet Capture.

Note: Enabling intra-session for HTTP pipelining is not supported when capturing at 10G rates on a single Decoder as the HTTP_Lua parser required to function can cause dropped packets at that rate of ingest. In this case, the load should be spread across multiple Concentrator-Decoder pairs.

Hardware Prerequisites

- A Series 5 or Series 6 Decoder
- An Intel 82599-based ethernet card, such as the Intel x520 or an Intel i40e-based card like the Intel XL710. All RSA-provided 10G cards meet this requirement such as the following:
  - All SMC-10GE cards provided by RSA.
  - A Dell Network Daughter Card using an Intel controller to provide 10G network interfaces. This is included in all Series 5 hardware.
- Sufficiently large and fast storage to meet the capture requirement of packets in case you set up in a mode other than meta-only. Storage considerations are covered later in this topic.
- If the Decoder will reassemble and handle high numbers of sessions, for example 30000 sessions per second or more, then at least two physical storage volumes are required. This will allow one volume to be used for writing session data while the other to read session data and deliver it for down-stream analytics.
- Each Network Decoder configured with a minimum of 2 DACs or SAN connectivity when set up in mode other than meta-only.

Software Prerequisites

RSA NetWitness Platform 11.6 or later.

Decoder Installation

Perform the following steps to install the Decoder.

Install the Decoder Service on Host that Will Perform the Capture

The Network Decoder service package contains the necessary software components to utilize the features described in the document. The Network Decoder service package and all other required packages are installed during the Decoder installation.

Select the Capture Interfaces and Assign Interfaces to DPDK

You must choose the physical network connections that will be used for packet ingest. The physical interface type will depend on the type of network traffic sent to the Decoder. The interface you select must match the physical characteristics of the Decoder feed (or network link) used to receive packets. For example, if you want to capture at 40G speeds on a cable, then you need a 40G interface capable of attaching to that cable.

For information on how to assign interfaces to DPDK, see (Optional) Data Plane Development Kit Packet Capture.

Verify the Decoder Packages Installation

Perform the following steps to verify the Decoder package installation:

- After you configure the interfaces to use with DPDK, ensure that you reboot the Decoder host.
- After the reboot, ensure that the DPDK interfaces appear in the Decoder's capture adapter list.

Select the Decoder Operating Mode

You can configure Decoder to handle any type of capture scenario that ranges from a very deep application-level inspection to very fast and simple network connection tracking. The Decoder includes the following three built-in templates that serve as a starting point for configuring your capture needs:

1. **Normal**: It is the default Decoder mode with no defined resource allocation and all tunable parameters set to default. It can be a good option if your capture rate is less than 5 Gbit/s or if you want to run a large amount of deep inspection content such as application parsers, feeds, OpenAppID detectors, complex Snort rules, and so on. For more information, see Configure the Decoder for Normal Mode (Default Mode).
2. **10G**: This mode is available in NetWitness 10.4 and later. You can use this mode if your capture rate is under 10 Gbit/s. This mode assumes that most of the packets are retained and saved in the packet database. Make a note that not all parsers can run at 10 Gbit/s speeds as some parsers are too complex for high speeds. For more information, see Configure the Decoder for 10G Mode.
3. **NDR**: This mode is to capture beyond the 10 Gbit/s speeds, but below 40Gb/sec. Make a note of the following points while using this mode:
   a. By default, this mode assumes that packets are not retained as the rate at which this mode consumes storage is significantly high.
   b. By default, the ingested packets are dropped. You must insert the network rules before the drop statement to pick portions of the network stream and keep the incoming packet feed for further analysis.
   c. Make a note that at higher speeds only a limited amount of deep packet inspection content can be used.
   d. Ensure that you enable the multi-thread assembly while using this mode.

   For more information, see Configure the Decoder for NDR Mode.

Configure the Decoder for Normal Mode (Default Mode)

By default, the Decoder runs in the normal mode and requires no additional configuration. The default normal mode captures up to 5Gb/sec with large amounts of deep packet inspection while storing network sessions.

Configure the Decoder for 10G Mode

The instructions in this section are to configure the Decoder for capture speeds up to 10Gb/sec with medium amounts of deep packet inspection while storing network sessions.

To configure the Decoder:

1. From the **Decoder Explore** view, right-click **Decoder** and select **Properties**. In the properties drop-down menu, select **reconfig** and enter the following parameters:
   `op=10g`
   Ensure that the correct parameter is displayed in the output.

2. Add the `update=true` parameter and run **reconfig** again to save the configuration changes. For example, your final parameters can be `op=10g update=true`.
3. From the **Decoder Explore** view, right-click **database** and select **Properties**.
4. In the **Properties** drop-down menu, select **reconfig** and enter the following parameters:
   `update=1 op=10g`
   These parameters adjust the packet database to use very large file sizes and Direct I/O.
5. Perform the following steps to select capture interfaces:
   a. In the **Decoder Explore** view, right-click on **Decoder** and select **Properties**.
   b. In the **Properties** drop-down menu, click **select** and run the command to view the available adapters. For example, this might produce output like the list below. The actual list of adapters will vary depending on the hardware available on the Decoder host.

      ```
      1: packet_mmap_,eth0
      2: DPDK,0000:41:00.0
      3: DPDK,0000:41:00.1
      ```

   c. To select a single interface for capture, enter the parameter `adapter=N` where `N` is the capture interface. For example, using the list above you might choose to capture on the first DPDK interface using the `adapter=2` parameter.
   d. To select more than one interface for capture, enter the parameter `adapter=N,M,...` where `N,M` and so on are the capture interfaces. For example, using the list above you might choose to capture on both DPDK interfaces using the `adapter=2,3` parameter.
6. (Optional) Application parsing is extremely CPU intensive and can cause the Decoder to drop packets. To mitigate application parsing-induced drops, you can set `/decoder/config/assembler.parse.valve` to `true`. These are the results:
   - When session parsing becomes a bottleneck, application parsers (HTTP, SMTP, FTP, and others) are temporarily disabled.
   - Sessions are not dropped when the application parsers are disabled, just the fidelity of the parsing performed on those sessions.
   - Sessions parsed when the application parsers are disabled still have associated network meta (from the network parser).
   - The statistic `/decoder/parsers/stats/blowoff.count` displays the count of all sessions that bypassed application parsers (network parsing is still performed).
   - When session parsing is no longer a potential bottleneck, the application parsers are automatically re-enabled.
   - The assembler session pool should be large enough that it is not forcing sessions.
   - You can determine if sessions are being forced by the statistic `/decoder/stats/assembler.sessions.forced` (it will be increasing). Also `/decoder/stats/assembler.sessions` will be within several hundred of `/decoder/config/assembler.session.pool`.

7. (Optional) If you need to adjust the MTU for capture, add the `snaplen` parameter to `capture.device.params`. Unlike previous releases, the `snaplen` does not need to be rounded up to any specific boundary. The Decoder automatically adjusts the MTU set on the capture interfaces.
8. The following configuration parameters are deprecated and no longer necessary.
   - The `core= parameter` in `capture.device.params`
   - Any configuration files under `/etc/pf_ring` directory
   - Separate `device= parameters` in `capture.device.params`. All multi-interface selection is performed with the **select** command described in Step 5 (b).

   Note: An Ethernet device installed post imaging must be added to DPDK if you want to use it as a capture interface. Similarly, it also require configuration if used as a network interface, or for system tools to access it without manual configuration.

Performance Tuning Parameters

By default, the following tunable parameters are disabled. It is recommended that you enable these to achieve high capture rates and consistent performance.

1. You can use the BPF filter to perform fast filtering of the packets. The BPF filter is the fastest way to remove packets from ingest. By dropping the unwanted traffic as early as possible, you can reduce the workload on the Decoder and ensure that essential packets are not dropped. For more information on BPF/PCAP filter, see Configure Capture Settings.
2. You can turn on the Receive Side Scaling (RSS) feature for your network capture interface. The RSS feature splits the traffic coming into the interface in separate queues. It allows each queue to be handled by a different thread, and therefore run on a different CPU core. This provides more CPU time to execute per-packet operations like evaluating BPF and Network rules. RSS makes handling of higher packet rates easier.
   The correct value for RSS will depends on the number of CPUs available on the host. A good starting point for the number of RSS threads depends on how many Cores per CPU socket are present in your Decoder host. Take the number of Cores present on a single CPU, and divide it by 2. For example, on a 12 core processor, you might use up to 6 RSS queues. You can distribute the RSS unequally between interfaces. For example you can assign 4 RSS queues to a busy interface and 2 to a less busy interface.

   IMPORTANT: Ensure that a Network Interface on a Decoder host is typically attached to only one CPU socket. Therefore, you must count only the CPU cores of the CPU socket that is attached to the network interfaces.
   Setting higher numbers of RSS queues is possible, but there are diminishing returns if the total number of RSS queues spawns more Capture and Assembly threads than there are physical cores on the Decoder host.

   For more details on how to configure RSS, see "*Utilizing Receive Side Scaling with DPDK*" in (Optional) Data Plane Development Kit Packet Capture.

Storage Considerations

Packet Retention Requires Extremely High Sustained Throughput

When capturing at 10G line rates, the storage system holding the packet and meta databases must be capable of sustained write throughput of 1400 MBytes/s. Make a note that many SANs are not capable of achieving these speeds.

Using SAN and Other Storage Configurations

The Decoder allows any storage configuration that can meet the sustained throughput requirement. The standard 8 Gbit FC link to a SAN is not sufficient to store packet data at 10G; in order to use a SAN it may be required to perform aggregation across multiple targets using a software-RAID Scheme. Thus environments using SAN are required to configure connectivity to the SAN using multiple FCs.

Optimize Read/Write Operations When Adding New Storage

A 10G Decoder is optimized to stagger read and write operations across multiple volumes so that the current file being written is on a different volume from the next file that will be written. This allows maximum throughput on the raid volume when reading data from the last file being written while writing the current file on a different volume. However, if volumes are added after a Decoder has been in use, the ability to stagger is limited because one or more volumes are already full so the new volume is the only place new files can be written.

To remedy this situation, an administrator can run a `stagger` command on an existing NetWitness database (packet, log, meta, or session), that has at least two volumes, to stagger the files across all volumes in the most optimal read/write pattern. The major use case is when new storage is added to an existing Decoder and you want to stagger the volumes before restarting capture.

The configuration nodes for this command are the session, meta, and packet databases. Each of these lives under `/database/config`, which is usually a root node. The config nodes for a Decoder are:

- `/database/config/packet.dir`
- `/database/config/meta.dir`
- `/database/config/session.dir`

The *NetWitness Platform Core Database Tuning Guide* has information on how those configurations are formatted.

The `stagger` command is typically only useful for a 10G Decoder and usually just for the packet database. Maximum performance is achieved for storing and retrieving packets when multiple volumes are present. In this scenario, the Decoder always fills the volume with the most free space. When the volumes are roughly the same size, this results in a staggered write pattern, which allows maximum throughput for reading and writing across all volumes. However, this only naturally occurs when multiple packet storage volumes are present at the time the Decoder is first deployed.

A typical use case is adding more storage to an existing Decoder to increase retention. However, when adding storage to an deployment that has already filled the existing volumes with stored packets, the Decoder will naturally fill the new storage with packets before rolling out any packets on the existing storage. This results in a suboptimal read/write pattern because most reads will occur on the same volume that is currently being written to. In a 10G deployment, reads are blocked from the volume when writes are occurring. This does not stop ALL reads on that volume, because the file is buffered in memory before being written, but it does result in suboptimal read performance.

With the `stagger` command, you can add more storage and then have the service naturally stagger the files across ALL volumes (existing and new) so that read performance is optimized.

Caution: This command should only be performed after the storage is mounted and the Decoder configured to use it (for example, after adding the mount point(s) to `packet.dir`).

The downside to this command is it can take some time to stagger and the Decoder should not be capturing during the stagger operation.

Recommended workflow:

1. Add all storage and configure mount points.
2. Add new storage mount points to `packet.dir` (or `session.dir`/`meta.dir`) and restart service (very important!).
3. Ensure capture is stopped.
4. Run the stagger operation. You must not terminate the connection that initiated the stagger operation until the operation is complete. If you run stagger from NwConsole, run the `timeout 0` command before sending the `stagger` command. This will prevent the normal 30 second command timeout.
5. Start capture after the `stagger` command finishes.

The following are the parameters for the stagger command:

- `type` - The database that will be staggered (session, meta, or packet). Typically only the packet database is useful for staggering, but it is possible to do the session or meta database when multiple volumes are present for those databases. Since the session and meta databases write far less data than the packet database, typically staggering those databases results in less noticeable performance gains.
- `dryRun` - If `true` (the default), will only return a description of the operations that would be performed. If `false`, then the files will actually be moved to an optimal read/write pattern. You MUST pass false to actually stagger the files.

Example usage from NwConsole:

```
login <decoder>:50004 <username> <password>
timeout 0
send /database stagger type=packet dryRun=false
```

Note: If you run this command using the RESTful API, pass the additional parameter `expiry=0` to prevent a timeout from the service. You will also need to ensure the HTTP client does not disconnect before the operation completes.

Parsing and Content Considerations

Parsing raw packets at high speeds presents unique challenges. Given the high session and packet rates, parsing efficiency is paramount. A single parser that is inefficient (spends too long examining packets) can slow the whole system down to the point where packets are dropped at the card.

For initial 10G testing, start with only native parsers (except SMB/WebMail). Use the native parsers to establish baseline performance and with little to no packet drops. Do not download any Live content until this has been done and the system is proven to capture without issue at high speeds.

After the system has been operational and running smoothly, Live content should be added very slowly, especially the parsers.

Best Practices

Whether you are updating a currently deployed system or deploying a new system, it is recommended you use the following best practices to minimize risk for packet loss. One caveat is that if you are updating a current 10G deployment but not adding any additional traffic. For example, a current Decoder capturing off a 10G card at 2G sustained should see no difference in performance, unless part of the update also entails adding additional traffic for capture.

- Incorporate baseline parsers (except SMB/Webmail, both of which generally have high CPU utilization) and monitor to ensure little to no packet loss.
- When adding additional parsers, add only one or two parsers at a time.
- Measure performance impact of newly added content, especially during peak traffic periods.
- If drops start occurring when they did not happen before, disable all newly-added parsers and enable just one at a time and measure the impact. This helps pinpoint individual parsers causing detrimental effects on performance. It may be possible to re-factor it to perform better or reduce its feature set to just what is necessary for the customer use case.
- Although lesser performance impacts, feeds should also be reviewed and added in a phased approach to help measure performance impacts.
- Application Rules also tend to have little observable impact, though again, it is best not to add a large number of rules at once without measuring the performance impact.
- If you regularly get a timeout message in the Investigate > Events view, such as `The query on channel 192577 was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time. If there are no messages indicating any specific warnings then try increasing the Core Query Timeout from the default 5 minutes to 10 minutes as described in "View Query and Session Attributes per Role" section of the *System Security and User Management Guide*.

Also, making the recommended configuration changes outlined in the Configuration section will help minimize potential issues.

Use Case 1: 10G Mode - Egress, Deep Packet Inspection

In this setup the goal is to ingest at 10G sustained line rates, perform Deep Packet Inspection (DPI), store metadata, and store raw packets for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- MA content (7 Lua parsers, 1 feed, 1 application rule)
- 4 feeds (alert ids info, nwmalwaredomains, warning, and suspicious)
- 41 application rules
- DNS_verbose_lua (disable DNS)
- fingerprint_javascript_lua
- fingerprint_pdf_lua
- fingerprint_rar_lua
- fingerprint_rtf_lua
- MAIL_lua (disable MAIL)
- SNMP_lua (disable SNMP)
- spectrum_lua
- SSH_lua (disable SSH)
- TLS_lua
- windows_command_shell
- windows_executable

Not Tested

- SMB_lua, native SMB disabled by default
- html_threat

Other

- HTTP_lua reduces the capture rate from >9G to <7G. At just under 5G this parser can be used in place of the native without dropping (in addition to the list above).
- xor_executable pushes parse CPU to 100% and the system can drop significantly due to parse backup.

Aggregation Adjustments Based on Tested Live Content

A 10G Decoder can serve aggregation to a single Concentrator while running at 10G speeds. Deployments using Malware Analysis, Event Stream Analysis, Warehouse Connector, and Reporting Engine are expected to impact performance and can lead to packet loss.

For the tested scenario, the Concentrator aggregates between 45 and 70k sessions per second. The 10G Decoder captures between 40 and 50k sessions per second. With the content identified above, this is about 1.5 to 2 million meta per second. Due to the high volume of session rates, the following configuration changes are recommended:

- Nice aggregation on the Concentrator limits the performance impact on the 10G Decoder. The following command turns on nice aggregation.
  `/concentrator/config/aggregate.nice = true`
- Due to the high volume of sessions on the Concentrator, you may consider activating parallel values mode on the Concentrator by setting `/sdk/config/parallel.values` to `16`. This improves Investigation performance when the number of sessions per second is greater than 30,000.
- If multiple aggregation streams are necessary, aggregating from the Concentrator instead has less impact on the Decoder.
- Further review for content and parsing is required for deployments where you want to use other NetWitness components (Warehouse Connector, Malware Analysis, ESA, and Reporting Engine).

Configure the Decoder for NDR Mode

The instructions in this section are to configure the Decoder for capture speeds more than 10Gb/sec but less than 40Gb/sec with small amounts of DPI while storing only metadata.

To configure the Decoder:

1. From the **Decoder Explore** view, right-click **Decoder** and select **Properties**. In the properties drop-down menu, select **reconfig** and enter the following parameters:
   `op=ndr`
   Ensure that the correct parameter is displayed in the output.
2. Add the `update=true` parameter and run **reconfig** again to save the configuration changes. For example, your final parameters will be `op=ndrupdate=true`.
3. From the **Decoder Explore** view, right-click **database** and select **Properties**.
4. In the **Properties** drop-down menu, select **reconfig** and enter the following parameters:
   `update=1 op=ndr`
   These parameters adjust the packet database to use very large file sizes and Direct I/O.
5. Perform the following steps to select capture interfaces:
   a. In the **Decoder Explore** view, right-click on **Decoder** and select **Properties**.
   b. In the **Properties** drop-down menu, click **select** and run the command to view the available adapters. For example, this might produce output like the list below. The actual list of adapters will vary depending on the hardware available on the Decoder host.

```
1: packet_mmap_,eth0
2: DPDK,0000:41:00.0
3: DPDK,0000:41:00.1
```

c. To select a single interface for capture, enter the parameter `adapter=N` where `N` is the capture interface. For example, using the list above you might choose to capture on the first DPDK interface using the `adapter=2` parameter.

d. To select more than one interface for capture, enter the parameter `adapter=N,M,...` where `N,M` and so on are the capture interfaces. For example, using the list above you might choose to capture on both DPDK interfaces using the `adapter=2,3` parameter.

6. (Optional) Application parsing is extremely CPU intensive and can cause the Decoder to drop packets. To mitigate application parsing-induced drops, you can set `/decoder/config/assembler.parse.valve` to `true`. These are the results:

   - When session parsing becomes a bottleneck, application parsers (HTTP, SMTP, FTP, and others) are temporarily disabled.
   - Sessions are not dropped when the application parsers are disabled, just the fidelity of the parsing performed on those sessions.
   - Sessions parsed when the application parsers are disabled still have associated network meta (from the network parser).
   - The statistic `/decoder/parsers/stats/blowoff.count` displays the count of all sessions that bypassed application parsers (network parsing is still performed).
   - When session parsing is no longer a potential bottleneck, the application parsers are automatically re-enabled.
   - The assembler session pool should be large enough that it is not forcing sessions.
   - You can determine if sessions are being forced by the statistic `/decoder/stats/assembler.sessions.forced` (it will be increasing). Also `/decoder/stats/assembler.sessions` will be within several hundred of `/decoder/config/assembler.session.pool`.

7. (Optional) If you need to adjust the MTU for capture, add the `snaplen` parameter to `capture.device.params`. Unlike previous releases, the `snaplen` does not need to be rounded up to any specific boundary. The Decoder automatically adjusts the MTU set on the capture interfaces.

8. The following configuration parameters are deprecated and no longer necessary.

   - The `core= parameter` in `capture.device.params`
   - Any configuration files under `/etc/pf_ring` directory
   - Separate `device= parameters` in `capture.device.params`. All multi-interface selection is performed with the **select** command described in Step 5 (b).

Note: An Ethernet device installed post imaging must be added to DPDK if you want to use it as a capture interface. Similarly, it also require configuration if used as a network interface, or for system tools to access it without manual configuration.

Performance Tuning Parameters

By default, the following tunable parameters are disabled. It is recommended that you enable these to achieve high capture rates and consistent performance.

1. You can use the BPF filter to perform fast filtering of the packets. The BPF filter is the fastest way to remove packets from ingest. By dropping traffic that you don't need to retain as early as possible, you can reduce the workload on the Decoder and ensure that essential packets are not dropped. For more information on BPF/PCAP filter, see Configure Capture Settings.
2. You can turn on the Receive Side Scaling (RSS) feature for your network capture interface. The RSS feature splits the traffic coming into the interface in separate queues. It allows each queue to be handled by a different thread, and therefore run on a different CPU core. This provides more CPU time to execute per-packet operations like evaluating BPF and Network rules. RSS makes handling of higher packet rates easier.
   The correct value for RSS will depends on the number of CPUs available on the host. A good starting point for the number of RSS threads depends on how many Cores per CPU socket are present in your Decoder host. Take the number of Cores present on a single CPU, and divide it by 2. For example, on a 12 core processor, you might use up to 6 RSS queues. You can distribute the RSS unequally between interfaces. For example you can assign 4 RSS queues to a busy interface and 2 to a less busy interface.

   IMPORTANT: Ensure that a Network Interface on a Decoder host is typically attached to only one CPU socket. Therefore, you must count only the CPU cores of the CPU socket that is attached to the network interfaces.
   Setting higher numbers of RSS queues is possible, but there are diminishing returns if the total number of RSS queues spawns more Capture and Assembly threads than there are physical cores on the Decoder host.

   For more details on how to configure RSS, see "*Utilizing Receive Side Scaling with DPDK*" in (Optional) Data Plane Development Kit Packet Capture.

Storage Considerations

NDR Mode Assumes No Packet Retention

The base or default configuration for NDR mode starts with all packet writes disabled. So, in this scenario the storage throughput requirements are relatively less. If you choose to turn on packet retention, be aware that it is relatively easy to overwhelm the I/O throughput of most storage solutions with a 40G network feed.

Parsing and Content Considerations

Parsing at Speeds Greater than 10G

As network ingest speeds increase, less CPU time is available to examine each packet or each session. It means that only limited amount of parsing can be performed. Fortunately, you only need to account for the amount of traffic that enters Decoder's parsing sub-system. For example, if you are ingesting 40 Gbit/s but filter out 30 Gbit/s of traffic, you only need to allow 10 Gbit/s of traffic to move through the parsers. At raw 40 Gbit/s speeds, only the "well-behaved" content works. The "well-behaved" content can defined as:

- Content that only activate when a large, uncommon search token is registered. It includes the Snort rules with a large, uncommon "fast-pattern" content field. For more information on fast-pattern and Snort rules, see Decoder Snort Detection.
- Content that generates a finite number of metas per session.
- Parsers that do not activate on session begin/end events.

The overall maximum throughput of the reassembly and parsing systems on Decoder is about 100,000 sessions (or streams) per second. However, in practical terms, very few downstream analytical or database services can handle such activity. If you find that your session rate is too high, or that your down-stream services cannot handle how many sessions the Decoder is generating, consider filtering out lower value traffic as it is ingested.

Best Practices

The NDR mode has a network rule configured to drop all incoming traffic by default. In general, when capturing above 10Gb/sec, you must limit the amount of traffic analyzed using DPI. These use cases in this document are example guidelines that do not guarantee the actual throughput of the Decoder. As mentioned previously in this document, the actual throughput of the Decoder depends on the following factors:

- **Amount of packets filtered vs. the amount of packets retained**
  It is important to filter the traffic that enters the Decoder to minimize packet drops. You can use Network Rules or a BPF Rule to filter the traffic before it enters advanced features such as stream reassembly and DPI. For more information, see Configure Network Rules and (Optional) Configure System-Level (BPF) Packet Filtering.
- **Average packet size**
  Decoder capturing small packets, for example 64 byte packets, can decrease the overall throughput of the Decoder. If average packet sizes are small, the line rate of Decoder ingest will fall significantly below the network link speed. It occurs due to the amount of per-packet overhead between each Ethernet frame.
- **Shared system resources**
  If the Decoder host is also hosting other applications or add-on features that manipulate data and perform other analysis, it can lower the overall throughput of the Decoder. Even if these add-on applications are not utilizing the CPU, these will utilize the shared resources like memory bandwidth and CPU interconnect bandwidth.

Whether you are updating a currently deployed system or deploying a new system, it is recommended you use the following best practices to minimize risk for packet loss.

- Incorporate baseline parsers (except SMB/Webmail, both of which generally have high CPU utilization) and monitor to ensure little to no packet loss.
- When adding additional parsers, add only one or two parsers at a time.
- Measure performance impact of newly added content, especially during peak traffic periods.
- If drops start occurring when they did not happen before, disable all newly-added parsers and enable just one at a time and measure the impact. This helps pinpoint individual parsers causing detrimental effects on performance. It may be possible to re-factor it to perform better or reduce its feature set to just what is necessary for the customer use case.
- If you regularly get a timeout message in the Investigate > Events view, such as `The query on channel 192577 was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time. If there are no messages indicating any specific warnings then try increasing the Core Query Timeout from the default 5 minutes to 10 minutes as described in "View Query and Session Attributes per Role" section of the *System Security and User Management Guide*.

Use Case 1: NDR Mode - Egress, General Purpose

Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- The **`NETWORK`** native parsers
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.

- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see Configure Capture Settings.

Use Case 2: NDR Mode - Egress, Data Exfiltration

Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- Native parsers - HTTP, HTTPS(SSL), SMTP, DNS, FTP, SFTP/SSH, and VNC
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see Configure Capture Settings.

Use Case 3: NDR Mode - Lateral Movement

Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- Native parsers - Kerberos, SMB, VNC, and SFTP/SSH
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Disable the Snort rule related to SMB traffic (for example, `M.HackTool.SMB.Impacket-Obfuscation.[Service Names]`). It is recommended because large regex parameters can cause Snort parser to utilize more than 50% CPU.
- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see Configure Capture Settings.

Configure High Speed Packet Capture Capability (Version 11.6 and Later)

This topic guides administrators on how to tune a Network Decoder specifically for high speed packet capture using NetWitness 11.6. This applies when capturing packets on 10G, 40G, or higher speed interface cards.

IMPORTANT: Packet capture at high speeds requires careful configuration and pushes the Decoder hardware to its limits, so you must read this entire topic while implementing a high speed capture solution.

RSA NetWitness provides support for high-speed collection on the Decoder. You can capture network packet data from higher speed networks and optimize your Network Decoder to capture network traffic up to 40 Gbit/s. These network capture speeds can be easily achieved on a Decoder but make a note that not all NetWitness Platform features operate at these rates. The actual throughput of Decoder depends on the following factors:

- **Amount of packets filtered vs. the amount of packets retained**
  It is important to filter the traffic that enters the Decoder to minimize packet drops. You can use Network Rules or a BPF Rule to filter the traffic before it enters advanced features such as stream reassembly and deep packet inspection. For more information, see Configure Network Rules and (Optional) Configure System-Level (BPF) Packet Filtering.
- **Average packet size**
  Decoder capturing small packets, for example 64 byte packets, can decrease the overall throughput of the Decoder. If average packet sizes are small, the line rate of Decoder ingest will fall significantly below the network link speed. It occurs due to the amount of per-packet overhead between each Ethernet frame.
- **Shared system resources**
  If the Decoder host is also hosting other applications or add-on features that manipulate data and perform other analysis, it can lower the overall throughput of the Decoder. Even if these add-on applications are not utilizing the CPU, these will utilize the shared resources like memory bandwidth and CPU interconnect bandwidth.

Enhancements that facilitate capture in these high speed environments include the following:

- Utilization of the DPDK capture drivers to leverage the commodity 10G and 40G Intel NIC cards for high-speed capture environments. For more information on how to configure DPDK, see (Optional) Data Plane Development Kit Packet Capture.
- Introduction of `assembler.parse.valve` configuration, which automatically disables application parsers when certain thresholds are exceeded, to limit risk of packet loss. When the application parsers are disabled, network layer parsers are still active. When stats fall below exceeded thresholds, application parsers are automatically re-enabled.

- Utilization of Berkeley Packet Filters (BPF) in 10G environments. For more information, see (Optional) Configure System-Level (BPF) Packet Filtering in Configure Capture Settings.
- Decoder will utilize the symmetric RSS to distribute capture loads among multiple CPU cores. For more information see, "*Utilizing Receive Side Scaling with DPDK*" in (Optional) Data Plane Development Kit Packet Capture
- Decoder will distribute the session reassembly work among multiple CPU cores as well, provided that the packets fed into each assembly are naturally segregated using the RSS features or as a result of being captured on different physical interfaces.

Note: Beginning with 11.5, the Network Decoder can capture from multiple interfaces simultaneously. This functionality allows Network Decoders to capture from multiple physical Network Interface Cards (NICs) while leveraging the same network rules, application rules, and parsers for each NIC. You can use this feature in all capture environments. For more information, see (Optional) Multiple Adapter Packet Capture.

Note: Enabling intra-session for HTTP pipelining is not supported when capturing at 10G rates on a single Decoder as the HTTP_Lua parser required to function can cause dropped packets at that rate of ingest. In this case, the load should be spread across multiple Concentrator-Decoder pairs.

Hardware Prerequisites

- A Series 5 or Series 6 Decoder
- An Intel 82599-based ethernet card, such as the Intel x520 or an Intel i40e-based card like the Intel XL710. All RSA-provided 10G cards meet this requirement such as the following:
  - All SMC-10GE cards provided by RSA.
  - A Dell Network Daughter Card using an Intel controller to provide 10G network interfaces. This is included in all Series 5 hardware.
- Sufficiently large and fast storage to meet the capture requirement of packets in case you set up in a mode other than meta-only. Storage considerations are covered later in this topic.
- If the Decoder will reassemble and handle high numbers of sessions, for example 30000 sessions per second or more, then at least two physical storage volumes are required. This will allow one volume to be used for writing session data while the other to read session data and deliver it for down-stream analytics.
- Each Network Decoder configured with a minimum of 2 DACs or SAN connectivity when set up in mode other than meta-only.

Software Prerequisites

RSA NetWitness Platform 11.6 or later.

Decoder Installation

Perform the following steps to install the Decoder.

Install the Decoder Service on Host that Will Perform the Capture

The Network Decoder service package contains the necessary software components to utilize the features described in the document. The Network Decoder service package and all other required packages are installed during the Decoder installation.

Select the Capture Interfaces and Assign Interfaces to DPDK

You must choose the physical network connections that will be used for packet ingest. The physical interface type will depend on the type of network traffic sent to the Decoder. The interface you select must match the physical characteristics of the Decoder feed (or network link) used to receive packets. For example, if you want to capture at 40G speeds on a cable, then you need a 40G interface capable of attaching to that cable.

For information on how to assign interfaces to DPDK, see (Optional) Data Plane Development Kit Packet Capture.

Verify the Decoder Packages Installation

Perform the following steps to verify the Decoder package installation:

- After you configure the interfaces to use with DPDK, ensure that you reboot the Decoder host.
- After the reboot, ensure that the DPDK interfaces appear in the Decoder's capture adapter list.

Select the Decoder Operating Mode

You can configure Decoder to handle any type of capture scenario that ranges from a very deep application-level inspection to very fast and simple network connection tracking. The Decoder includes the following three built-in templates that serve as a starting point for configuring your capture needs:

1. **Normal**: It is the default Decoder mode with no defined resource allocation and all tunable parameters set to default. It can be a good option if your capture rate is less than 5 Gbit/s or if you want to run a large amount of deep inspection content such as application parsers, feeds, OpenAppID detectors, complex Snort rules, and so on. For more information, see Configure the Decoder for Normal Mode (Default Mode).
2. **10G**: This mode is available in NetWitness 10.4 and later. You can use this mode if your capture rate is under 10 Gbit/s. This mode assumes that most of the packets are retained and saved in the packet database. Make a note that not all parsers can run at 10 Gbit/s speeds as some parsers are too complex for high speeds. For more information, see Configure the Decoder for 10G Mode.
3. **NDR**: This mode is to capture beyond the 10 Gbit/s speeds, but below 40Gb/sec. Make a note of the following points while using this mode:
   a. By default, this mode assumes that packets are not retained as the rate at which this mode consumes storage is significantly high.
   b. By default, the ingested packets are dropped. You must insert the network rules before the drop statement to pick portions of the network stream and keep the incoming packet feed for further analysis.

c. Make a note that at higher speeds only a limited amount of deep packet inspection content can be used.

d. Ensure that you enable the multi-thread assembly while using this mode.

For more information, see Configure the Decoder for NDR Mode.

Configure the Decoder for Normal Mode (Default Mode)

By default, the Decoder runs in the normal mode and requires no additional configuration. The default normal mode captures up to 5Gb/sec with large amounts of deep packet inspection while storing network sessions.

Configure the Decoder for 10G Mode

The instructions in this section are to configure the Decoder for capture speeds up to 10Gb/sec with medium amounts of deep packet inspection while storing network sessions.

To configure the Decoder:

1. From the **Decoder Explore** view, right-click **Decoder** and select **Properties**. In the properties drop-down menu, select **reconfig** and enter the following parameters:
   `op=10g`
   Ensure that the correct parameter is displayed in the output.
2. Add the `update=true` parameter and run **reconfig** again to save the configuration changes. For example, your final parameters can be `op=10g update=true`.
3. From the **Decoder Explore** view, right-click **database** and select **Properties**.
4. In the **Properties** drop-down menu, select **reconfig** and enter the following parameters:
   `update=1 op=10g`
   These parameters adjust the packet database to use very large file sizes and Direct I/O.
5. Perform the following steps to select capture interfaces:
   a. In the **Decoder Explore** view, right-click on **Decoder** and select **Properties**.
   b. In the **Properties** drop-down menu, click **select** and run the command to view the available adapters. For example, this might produce output like the list below. The actual list of adapters will vary depending on the hardware available on the Decoder host.

      ```
      1: packet_mmap_,eth0
      2: DPDK,0000:41:00.0
      3: DPDK,0000:41:00.1
      ```

   c. To select a single interface for capture, enter the parameter `adapter=N` where `N` is the capture interface. For example, using the list above you might choose to capture on the first DPDK interface using the `adapter=2` parameter.
   d. To select more than one interface for capture, enter the parameter `adapter=N,M,...` where `N,M` and so on are the capture interfaces. For example, using the list above you might choose to capture on both DPDK interfaces using the `adapter=2,3` parameter.

6. (Optional) Application parsing is extremely CPU intensive and can cause the Decoder to drop packets. To mitigate application parsing-induced drops, you can set `/decoder/config/assembler.parse.valve` to `true`. These are the results:
   - When session parsing becomes a bottleneck, application parsers (HTTP, SMTP, FTP, and others) are temporarily disabled.
   - Sessions are not dropped when the application parsers are disabled, just the fidelity of the parsing performed on those sessions.
   - Sessions parsed when the application parsers are disabled still have associated network meta (from the network parser).
   - The statistic `/decoder/parsers/stats/blowoff.count` displays the count of all sessions that bypassed application parsers (network parsing is still performed).
   - When session parsing is no longer a potential bottleneck, the application parsers are automatically re-enabled.
   - The assembler session pool should be large enough that it is not forcing sessions.
   - You can determine if sessions are being forced by the statistic `/decoder/stats/assembler.sessions.forced` (it will be increasing). Also `/decoder/stats/assembler.sessions` will be within several hundred of `/decoder/config/assembler.session.pool`.
7. (Optional) If you need to adjust the MTU for capture, add the `snaplen` parameter to `capture.device.params`. Unlike previous releases, the `snaplen` does not need to be rounded up to any specific boundary. The Decoder automatically adjusts the MTU set on the capture interfaces.
8. The following configuration parameters are deprecated and no longer necessary.
   - The `core= parameter` in `capture.device.params`
   - Any configuration files under `/etc/pf_ring` directory
   - Separate `device= parameters` in `capture.device.params`. All multi-interface selection is performed with the **select** command described in Step 5 (b).

   Note: An Ethernet device installed post imaging must be added to DPDK if you want to use it as a capture interface. Similarly, it also require configuration if used as a network interface, or for system tools to access it without manual configuration.

Performance Tuning Parameters

By default, the following tunable parameters are disabled. It is recommended that you enable these to achieve high capture rates and consistent performance.

1. You can use the BPF filter to perform fast filtering of the packets. The BPF filter is the fastest way to remove packets from ingest. By dropping the unwanted traffic as early as possible, you can reduce the workload on the Decoder and ensure that essential packets are not dropped. For more information on BPF/PCAP filter, see Configure Capture Settings.
2. You can turn on the Receive Side Scaling (RSS) feature for your network capture interface. The RSS feature splits the traffic coming into the interface in separate queues. It allows each queue to be

handled by a different thread, and therefore run on a different CPU core. This provides more CPU time to execute per-packet operations like evaluating BPF and Network rules. RSS makes handling of higher packet rates easier.
The correct value for RSS will depends on the number of CPUs available on the host. A good starting point for the number of RSS threads depends on how many Cores per CPU socket are present in your Decoder host. Take the number of Cores present on a single CPU, and divide it by 2. For example, on a 12 core processor, you might use up to 6 RSS queues. You can distribute the RSS unequally between interfaces. For example you can assign 4 RSS queues to a busy interface and 2 to a less busy interface.

IMPORTANT: Ensure that a Network Interface on a Decoder host is typically attached to only one CPU socket. Therefore, you must count only the CPU cores of the CPU socket that is attached to the network interfaces.
Setting higher numbers of RSS queues is possible, but there are diminishing returns if the total number of RSS queues spawns more Capture and Assembly threads than there are physical cores on the Decoder host.

For more details on how to configure RSS, see "*Utilizing Receive Side Scaling with DPDK*" in (Optional) Data Plane Development Kit Packet Capture.

Storage Considerations

Packet Retention Requires Extremely High Sustained Throughput

When capturing at 10G line rates, the storage system holding the packet and meta databases must be capable of sustained write throughput of 1400 MBytes/s. Make a note that many SANs are not capable of achieving these speeds.

Using SAN and Other Storage Configurations

The Decoder allows any storage configuration that can meet the sustained throughput requirement. The standard 8 Gbit FC link to a SAN is not sufficient to store packet data at 10G; in order to use a SAN it may be required to perform aggregation across multiple targets using a software-RAID Scheme. Thus environments using SAN are required to configure connectivity to the SAN using multiple FCs.

Optimize Read/Write Operations When Adding New Storage

A 10G Decoder is optimized to stagger read and write operations across multiple volumes so that the current file being written is on a different volume from the next file that will be written. This allows maximum throughput on the raid volume when reading data from the last file being written while writing the current file on a different volume. However, if volumes are added after a Decoder has been in use, the ability to stagger is limited because one or more volumes are already full so the new volume is the only place new files can be written.

To remedy this situation, an administrator can run a `stagger` command on an existing NetWitness database (packet, log, meta, or session), that has at least two volumes, to stagger the files across all volumes in the most optimal read/write pattern. The major use case is when new storage is added to an existing Decoder and you want to stagger the volumes before restarting capture.

The configuration nodes for this command are the session, meta, and packet databases. Each of these lives under `/database/config`, which is usually a root node. The config nodes for a Decoder are:

- `/database/config/packet.dir`
- `/database/config/meta.dir`
- `/database/config/session.dir`

The *NetWitness Platform Core Database Tuning Guide* has information on how those configurations are formatted.

The `stagger` command is typically only useful for a 10G Decoder and usually just for the packet database. Maximum performance is achieved for storing and retrieving packets when multiple volumes are present. In this scenario, the Decoder always fills the volume with the most free space. When the volumes are roughly the same size, this results in a staggered write pattern, which allows maximum throughput for reading and writing across all volumes. However, this only naturally occurs when multiple packet storage volumes are present at the time the Decoder is first deployed.

A typical use case is adding more storage to an existing Decoder to increase retention. However, when adding storage to an deployment that has already filled the existing volumes with stored packets, the Decoder will naturally fill the new storage with packets before rolling out any packets on the existing storage. This results in a suboptimal read/write pattern because most reads will occur on the same volume that is currently being written to. In a 10G deployment, reads are blocked from the volume when writes are occurring. This does not stop ALL reads on that volume, because the file is buffered in memory before being written, but it does result in suboptimal read performance.

With the `stagger` command, you can add more storage and then have the service naturally stagger the files across ALL volumes (existing and new) so that read performance is optimized.

Caution: This command should only be performed after the storage is mounted and the Decoder configured to use it (for example, after adding the mount point(s) to `packet.dir`).

The downside to this command is it can take some time to stagger and the Decoder should not be capturing during the stagger operation.

Recommended workflow:

1. Add all storage and configure mount points.
2. Add new storage mount points to `packet.dir` (or `session.dir`/`meta.dir`) and restart service (very important!).
3. Ensure capture is stopped.
4. Run the stagger operation. You must not terminate the connection that initiated the stagger operation until the operation is complete. If you run stagger from NwConsole, run the `timeout 0` command before sending the `stagger` command. This will prevent the normal 30 second command timeout.
5. Start capture after the `stagger` command finishes.

The following are the parameters for the stagger command:

- `type` - The database that will be staggered (session, meta, or packet). Typically only the packet database is useful for staggering, but it is possible to do the session or meta database when multiple volumes are present for those databases. Since the session and meta databases write far less data than the packet database, typically staggering those databases results in less noticeable performance gains.

- `dryRun` - If `true` (the default), will only return a description of the operations that would be performed. If `false`, then the files will actually be moved to an optimal read/write pattern. You MUST pass false to actually stagger the files.

Example usage from NwConsole:

```
login <decoder>:50004 <username> <password>
timeout 0
send /database stagger type=packet dryRun=false
```

Note: If you run this command using the RESTful API, pass the additional parameter `expiry=0` to prevent a timeout from the service. You will also need to ensure the HTTP client does not disconnect before the operation completes.

Parsing and Content Considerations

Parsing raw packets at high speeds presents unique challenges. Given the high session and packet rates, parsing efficiency is paramount. A single parser that is inefficient (spends too long examining packets) can slow the whole system down to the point where packets are dropped at the card.

For initial 10G testing, start with only native parsers (except SMB/WebMail). Use the native parsers to establish baseline performance and with little to no packet drops. Do not download any Live content until this has been done and the system is proven to capture without issue at high speeds.

After the system has been operational and running smoothly, Live content should be added very slowly, especially the parsers.

Best Practices

Whether you are updating a currently deployed system or deploying a new system, it is recommended you use the following best practices to minimize risk for packet loss. One caveat is that if you are updating a current 10G deployment but not adding any additional traffic. For example, a current Decoder capturing off a 10G card at 2G sustained should see no difference in performance, unless part of the update also entails adding additional traffic for capture.

- Incorporate baseline parsers (except SMB/Webmail, both of which generally have high CPU utilization) and monitor to ensure little to no packet loss.
- When adding additional parsers, add only one or two parsers at a time.
- Measure performance impact of newly added content, especially during peak traffic periods.
- If drops start occurring when they did not happen before, disable all newly-added parsers and enable just one at a time and measure the impact. This helps pinpoint individual parsers causing detrimental effects on performance. It may be possible to re-factor it to perform better or reduce its feature set to just what is necessary for the customer use case.
- Although lesser performance impacts, feeds should also be reviewed and added in a phased approach to help measure performance impacts.

- Application Rules also tend to have little observable impact, though again, it is best not to add a large number of rules at once without measuring the performance impact.
- If you regularly get a timeout message in the Investigate > Events view, such as `The query on channel 192577 was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time. If there are no messages indicating any specific warnings then try increasing the Core Query Timeout from the default 5 minutes to 10 minutes as described in "View Query and Session Attributes per Role" section of the *System Security and User Management Guide*.

Also, making the recommended configuration changes outlined in the Configuration section will help minimize potential issues.

Use Case 1: 10G Mode - Egress, Deep Packet Inspection

In this setup the goal is to ingest at 10G sustained line rates, perform Deep Packet Inspection (DPI), store metadata, and store raw packets for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- MA content (7 Lua parsers, 1 feed, 1 application rule)
- 4 feeds (alert ids info, nwmalwaredomains, warning, and suspicious)
- 41 application rules
- DNS_verbose_lua (disable DNS)
- fingerprint_javascript_lua
- fingerprint_pdf_lua
- fingerprint_rar_lua
- fingerprint_rtf_lua
- MAIL_lua (disable MAIL)
- SNMP_lua (disable SNMP)
- spectrum_lua
- SSH_lua (disable SSH)
- TLS_lua
- windows_command_shell
- windows_executable

Not Tested

- SMB_lua, native SMB disabled by default
- html_threat

Other

- HTTP_lua reduces the capture rate from >9G to <7G. At just under 5G this parser can be used in place of the native without dropping (in addition to the list above).
- xor_executable pushes parse CPU to 100% and the system can drop significantly due to parse backup.

Aggregation Adjustments Based on Tested Live Content

A 10G Decoder can serve aggregation to a single Concentrator while running at 10G speeds. Deployments using Malware Analysis, Event Stream Analysis, Warehouse Connector, and Reporting Engine are expected to impact performance and can lead to packet loss.

For the tested scenario, the Concentrator aggregates between 45 and 70k sessions per second. The 10G Decoder captures between 40 and 50k sessions per second. With the content identified above, this is about 1.5 to 2 million meta per second. Due to the high volume of session rates, the following configuration changes are recommended:

- Nice aggregation on the Concentrator limits the performance impact on the 10G Decoder. The following command turns on nice aggregation.
  `/concentrator/config/aggregate.nice = true`
- Due to the high volume of sessions on the Concentrator, you may consider activating parallel values mode on the Concentrator by setting `/sdk/config/parallel.values` to `16`. This improves Investigation performance when the number of sessions per second is greater than 30,000.
- If multiple aggregation streams are necessary, aggregating from the Concentrator instead has less impact on the Decoder.
- Further review for content and parsing is required for deployments where you want to use other NetWitness components (Warehouse Connector, Malware Analysis, ESA, and Reporting Engine).

Configure the Decoder for NDR Mode

The instructions in this section are to configure the Decoder for capture speeds more than 10Gb/sec but less than 40Gb/sec with small amounts of DPI while storing only metadata.

To configure the Decoder:

1. From the **Decoder Explore** view, right-click **Decoder** and select **Properties**. In the properties drop-down menu, select **reconfig** and enter the following parameters:
   `op=ndr`
   Ensure that the correct parameter is displayed in the output.

2. Add the `update=true` parameter and run **reconfig** again to save the configuration changes. For example, your final parameters will be `op=ndrupdate=true`.
3. From the **Decoder Explore** view, right-click **database** and select **Properties**.
4. In the **Properties** drop-down menu, select **reconfig** and enter the following parameters:
   `update=1 op=ndr`
   These parameters adjust the packet database to use very large file sizes and Direct I/O.
5. Perform the following steps to select capture interfaces:
   a. In the **Decoder Explore** view, right-click on **Decoder** and select **Properties**.
   b. In the **Properties** drop-down menu, click **select** and run the command to view the available adapters. For example, this might produce output like the list below. The actual list of adapters will vary depending on the hardware available on the Decoder host.

      ```
      1: packet_mmap_,eth0
      2: DPDK,0000:41:00.0
      3: DPDK,0000:41:00.1
      ```

   c. To select a single interface for capture, enter the parameter `adapter=N` where `N` is the capture interface. For example, using the list above you might choose to capture on the first DPDK interface using the `adapter=2` parameter.
   d. To select more than one interface for capture, enter the parameter `adapter=N,M,...` where `N,M` and so on are the capture interfaces. For example, using the list above you might choose to capture on both DPDK interfaces using the `adapter=2,3` parameter.
6. (Optional) Application parsing is extremely CPU intensive and can cause the Decoder to drop packets. To mitigate application parsing-induced drops, you can set `/decoder/config/assembler.parse.valve` to `true`. These are the results:
   - When session parsing becomes a bottleneck, application parsers (HTTP, SMTP, FTP, and others) are temporarily disabled.
   - Sessions are not dropped when the application parsers are disabled, just the fidelity of the parsing performed on those sessions.
   - Sessions parsed when the application parsers are disabled still have associated network meta (from the network parser).
   - The statistic `/decoder/parsers/stats/blowoff.count` displays the count of all sessions that bypassed application parsers (network parsing is still performed).
   - When session parsing is no longer a potential bottleneck, the application parsers are automatically re-enabled.
   - The assembler session pool should be large enough that it is not forcing sessions.
   - You can determine if sessions are being forced by the statistic `/decoder/stats/assembler.sessions.forced` (it will be increasing). Also `/decoder/stats/assembler.sessions` will be within several hundred of `/decoder/config/assembler.session.pool`.

7. (Optional) If you need to adjust the MTU for capture, add the `snaplen` parameter to `capture.device.params`. Unlike previous releases, the `snaplen` does not need to be rounded up to any specific boundary. The Decoder automatically adjusts the MTU set on the capture interfaces.
8. The following configuration parameters are deprecated and no longer necessary.
   - The `core= parameter` in `capture.device.params`
   - Any configuration files under `/etc/pf_ring` directory
   - Separate `device= parameters` in `capture.device.params`. All multi-interface selection is performed with the **select** command described in Step 5 (b).

   Note: An Ethernet device installed post imaging must be added to DPDK if you want to use it as a capture interface. Similarly, it also require configuration if used as a network interface, or for system tools to access it without manual configuration.

Performance Tuning Parameters

By default, the following tunable parameters are disabled. It is recommended that you enable these to achieve high capture rates and consistent performance.

1. You can use the BPF filter to perform fast filtering of the packets. The BPF filter is the fastest way to remove packets from ingest. By dropping traffic that you don't need to retain as early as possible, you can reduce the workload on the Decoder and ensure that essential packets are not dropped. For more information on BPF/PCAP filter, see Configure Capture Settings.
2. You can turn on the Receive Side Scaling (RSS) feature for your network capture interface. The RSS feature splits the traffic coming into the interface in separate queues. It allows each queue to be handled by a different thread, and therefore run on a different CPU core. This provides more CPU time to execute per-packet operations like evaluating BPF and Network rules. RSS makes handling of higher packet rates easier.
   The correct value for RSS will depends on the number of CPUs available on the host. A good starting point for the number of RSS threads depends on how many Cores per CPU socket are present in your Decoder host. Take the number of Cores present on a single CPU, and divide it by 2. For example, on a 12 core processor, you might use up to 6 RSS queues. You can distribute the RSS unequally between interfaces. For example you can assign 4 RSS queues to a busy interface and 2 to a less busy interface.

   IMPORTANT: Ensure that a Network Interface on a Decoder host is typically attached to only one CPU socket. Therefore, you must count only the CPU cores of the CPU socket that is attached to the network interfaces.
   Setting higher numbers of RSS queues is possible, but there are diminishing returns if the total number of RSS queues spawns more Capture and Assembly threads than there are physical cores on the Decoder host.

   For more details on how to configure RSS, see "*Utilizing Receive Side Scaling with DPDK*" in (Optional) Data Plane Development Kit Packet Capture.

Storage Considerations

NDR Mode Assumes No Packet Retention

The base or default configuration for NDR mode starts with all packet writes disabled. So, in this scenario the storage throughput requirements are relatively less. If you choose to turn on packet retention, be aware that it is relatively easy to overwhelm the I/O throughput of most storage solutions with a 40G network feed.

Parsing and Content Considerations

Parsing at Speeds Greater than 10G

As network ingest speeds increase, less CPU time is available to examine each packet or each session. It means that only limited amount of parsing can be performed. Fortunately, you only need to account for the amount of traffic that enters Decoder's parsing sub-system. For example, if you are ingesting 40 Gbit/s but filter out 30 Gbit/s of traffic, you only need to allow 10 Gbit/s of traffic to move through the parsers. At raw 40 Gbit/s speeds, only the "well-behaved" content works. The "well-behaved" content can defined as:

- Content that only activate when a large, uncommon search token is registered. It includes the Snort rules with a large, uncommon "fast-pattern" content field. For more information on fast-pattern and Snort rules, see Decoder Snort Detection.
- Content that generates a finite number of metas per session.
- Parsers that do not activate on session begin/end events.

The overall maximum throughput of the reassembly and parsing systems on Decoder is about 100,000 sessions (or streams) per second. However, in practical terms, very few downstream analytical or database services can handle such activity. If you find that your session rate is too high, or that your down-stream services cannot handle how many sessions the Decoder is generating, consider filtering out lower value traffic as it is ingested.

Best Practices

The NDR mode has a network rule configured to drop all incoming traffic by default. In general, when capturing above 10Gb/sec, you must limit the amount of traffic analyzed using DPI. These use cases in this document are example guidelines that do not guarantee the actual throughput of the Decoder. As mentioned previously in this document, the actual throughput of the Decoder depends on the following factors:

- **Amount of packets filtered vs. the amount of packets retained**
  It is important to filter the traffic that enters the Decoder to minimize packet drops. You can use Network Rules or a BPF Rule to filter the traffic before it enters advanced features such as stream reassembly and DPI. For more information, see Configure Network Rules and (Optional) Configure System-Level (BPF) Packet Filtering.
- **Average packet size**
  Decoder capturing small packets, for example 64 byte packets, can decrease the overall throughput of

the Decoder. If average packet sizes are small, the line rate of Decoder ingest will fall significantly below the network link speed. It occurs due to the amount of per-packet overhead between each Ethernet frame.

- **Shared system resources**
  If the Decoder host is also hosting other applications or add-on features that manipulate data and perform other analysis, it can lower the overall throughput of the Decoder. Even if these add-on applications are not utilizing the CPU, these will utilize the shared resources like memory bandwidth and CPU interconnect bandwidth.

Whether you are updating a currently deployed system or deploying a new system, it is recommended you use the following best practices to minimize risk for packet loss.

- Incorporate baseline parsers (except SMB/Webmail, both of which generally have high CPU utilization) and monitor to ensure little to no packet loss.
- When adding additional parsers, add only one or two parsers at a time.
- Measure performance impact of newly added content, especially during peak traffic periods.
- If drops start occurring when they did not happen before, disable all newly-added parsers and enable just one at a time and measure the impact. This helps pinpoint individual parsers causing detrimental effects on performance. It may be possible to re-factor it to perform better or reduce its feature set to just what is necessary for the customer use case.
- If you regularly get a timeout message in the Investigate > Events view, such as `The query on channel 192577 was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time. If there are no messages indicating any specific warnings then try increasing the Core Query Timeout from the default 5 minutes to 10 minutes as described in "View Query and Session Attributes per Role" section of the *System Security and User Management Guide*.

Use Case 1: NDR Mode - Egress, General Purpose

Generate NetFlow Style Meta Only + Small Subset of Snort Rules + All Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- The **`NETWORK`** native parsers
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see Configure Capture Settings.

Use Case 2: NDR Mode - Egress, Data Exfiltration

Generate NetFlow Style Meta Only + Exfiltration Focused Specific Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- Native parsers - HTTP, HTTPS(SSL), SMTP, DNS, FTP, SFTP/SSH, and VNC
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see Configure Capture Settings.

Use Case 3: NDR Mode - Lateral Movement

Generate NetFlow Style Meta Only + Small Subset of Snort Rules + Specific Native Parsers

In this setup the goal is to ingest at rates higher than 10G sustained line rates, perform DPI with only native network parsers (non-Lua), and store only metadata for some time.

Tested Live Content

All (not each) of the following parsers can run at 10G on the test data set used:

- Native parsers - Kerberos, SMB, VNC, and SFTP/SSH
- 35 Snort rules for FireEye red team tool detection

Not Tested

All Lua parsers.

Other

The following considerations are recommended while running the Decoder in the NDR mode:

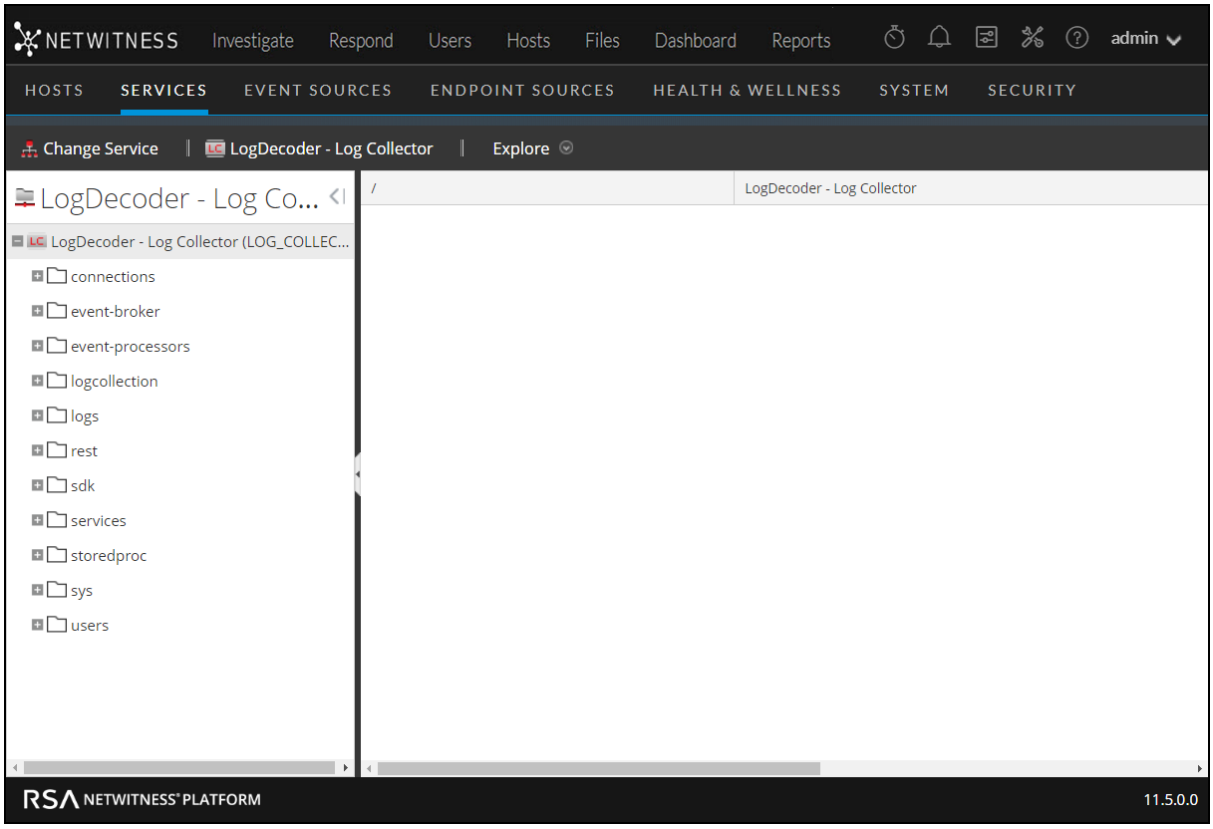
- Disable the Snort rule related to SMB traffic (for example, `M.HackTool.SMB.Impacket-Obfuscation.[Service Names]`). It is recommended because large regex parameters can cause Snort parser to utilize more than 50% CPU.
- Some native parsers should be disabled in some environments. For example, the DNS parser might generate too many small sessions or too much meta per session.
- It is recommended to keep the session rate less than 35,000 sessions per second. This is to limit the packet drops to less than 1% at such high throughput rates.
- Configure rules to filter the traffic that you do not require. For more information on BPF/PCAP filter, see Configure Capture Settings.

Configure a Log Decoder to Accept Protobuf

There are occasions when you want to analyze log files that are in protobuf (Protocol Buffer) format. You can configure a Log Decoder with a Log Collector service to accept logs in protobuf (Protocol Buffer) format.

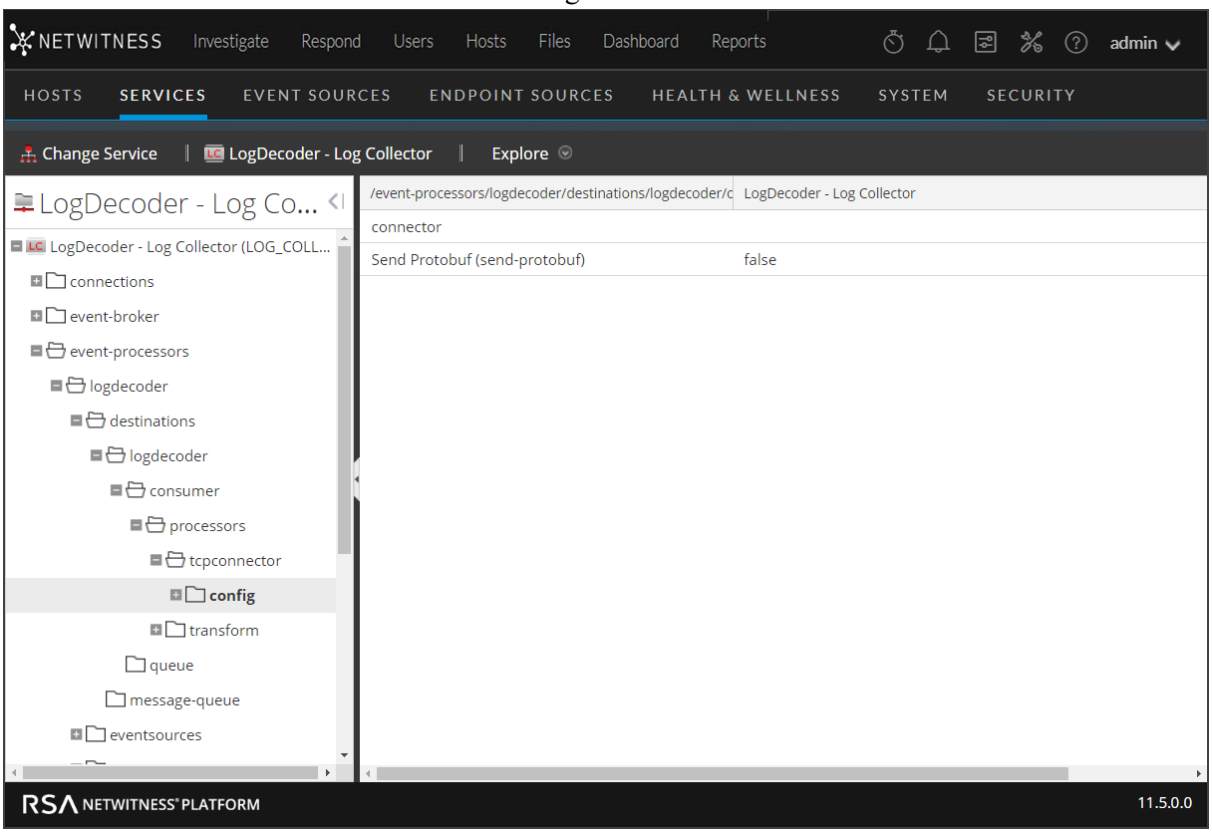
To import a log file to a Log Decoder:

1. Go to **(Admin) > Services**.
2. Select a Log Decoder with a Log Collector service in the **Service** list, and select **> View > Explore**.
   The Explore view for the Log Decoder -Log Collector is displayed.



3. Navigate to **`event-processors/logdecoder/destinations/logdecoder/consumer/processors/tcpconnector/config`**

Your screen should look similar to the following.



4. For the **send-protobuf** field, select **false**, and change the value to **true.**
5. Navigate to `event-processors/logdecoder/destinations/logdecoder/consumer/processors/tcpconnector/config/connector/channel/tcp` and change the **port** value to **50202**.
6. Navigate to `event-processors/logdecoder/destinations/logdecoder/consumer/processors/tcpconnector/config/connector/event` and change the following parameters:
   - Clear the **delimiter** field
   - Change **format** to **%text%**

Configure Session Split Timeouts

The default behavior of the Decoder is to automatically end sessions that exceed a configured size or have been inactive for a period of time. When the session is ended due to timeout, any subsequent packets received in that session appear to be stored in a new session. You can mitigate the effect of session splitting due to long periods of inactivity between packets using this procedure.

When a Decoder session exceeds a configured size (32MB by default, the `/decoder/config/assembler.size.max`) or has been inactive for a period of time, the session is split. NetWitness has the previous packet and the next packet and can propagate session state from the initial session fragment to the subsequent session fragment.

Each session fragment is annotated (`session.split` meta) such that it can be identified and associated with other fragments from the actual network session. Directionality as determined by the initial session reduces the occurrence of fragments having reversed directionality.

If there is a gap in time between packets large enough that there are no longer any packets for the session in memory, the session is removed from the Decoder. If a subsequent packet shows up after this occurs, a new session is created with no context to the preceding session. The issue is the inability to continue a session when we encounter a gap between packets of a session that is larger than the packets we are buffering (based upon available memory and timeout configurations). Once the last packet of a session is removed from memory, the session is also removed, and with it the necessary context for ensuring consistent directionality.

There are two timeout settings in a Network Decoder, `/decoder/config/assember.timeout.session` and `assembler.timeout.packet`. Both default to 60 seconds. The setting `assembler.timeout.session` controls how long a session lives in Assembler without receiving another packet. The setting `assembler.timeout.packet` controls how long a session waits before getting parsed. If the session is kicked out of Assembler before this timeout, then it automatically goes to parsing.

The session timeout is the number of seconds since the last packet was added to that session. Therefore, this timeout resets on every packet added to that session. The packet timeout is the number of seconds since the very first packet for that session was added (in other words, the packet that created the session). This is never reset and once the timeout expires, the session is parsed.

The important point is a session can be parsed but still remain in Assembler. A session in Assembler can still have packets added to it, even if it has already been parsed. Packets added after the session is parsed will never be seen by parsers, but they will be attached to the session and can be viewed by a subsequent `/sdk content` or `/sdk packets` call.

After a session is parsed, the session AND its metadata are written to disk. At this point, they can be aggregated and "seen" by `sdk` commands. Packets are written in order of capture and are not reordered by what session they belong to. Nor are they necessarily written when the session and meta data are written.

You can disable both timeout nodes, `/decoder/config/assembler.timeout.session` and `assembler.timeout.packet`, by setting them to zero in the Services Explore view.

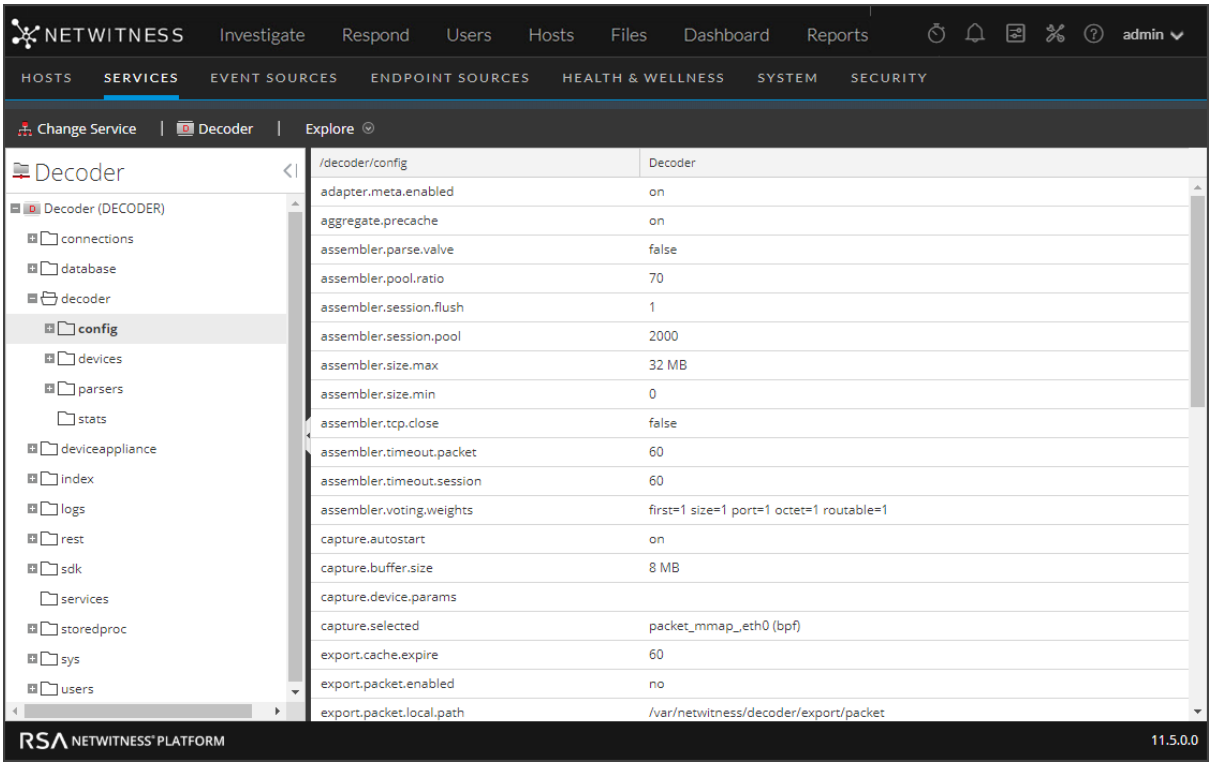
If both timeouts are disabled, the sessions are still split due to time or size expiration. However, the Decoder keeps track of the network stream for as long as it has sufficient memory. Thus, when more packets arrive on the same network stream, the Decoder adds `split` meta items to the subsequent sessions. Using a combination of the `split` metadata and the stream key, it is possible to reconstruct the network stream from the multiple sessions.

The length of time for which sessions are tracked is limited by the number of session pool entries available on the Decoder, and therefore the actual time window varies according to the rate at which new sessions are added. If new sessions are added at a high rate, the size of the time window decreases. The size of the pool is set using the configuration entry `/decoder/config/assembler.session.pool`, which sets the maximum number of sessions that will be tracked at a time.

The `/decoder/stats/assembler.timespan` statistic allows you to see when the Decoder is no longer tracking session splits because the ingest rate is too high and the Decoder does not have enough memory to track. This statistic shows the number of seconds tracked within the session table, which is the effective time window in which the Decoder can link together sessions. Under normal operation this statistic matches the value of `/decoder/config/assembler.timeout.session`, but when running in Time Split mode, the `/decoder/stats/assembler.timespan` statistic grows or shrinks depending on the ingest rate.

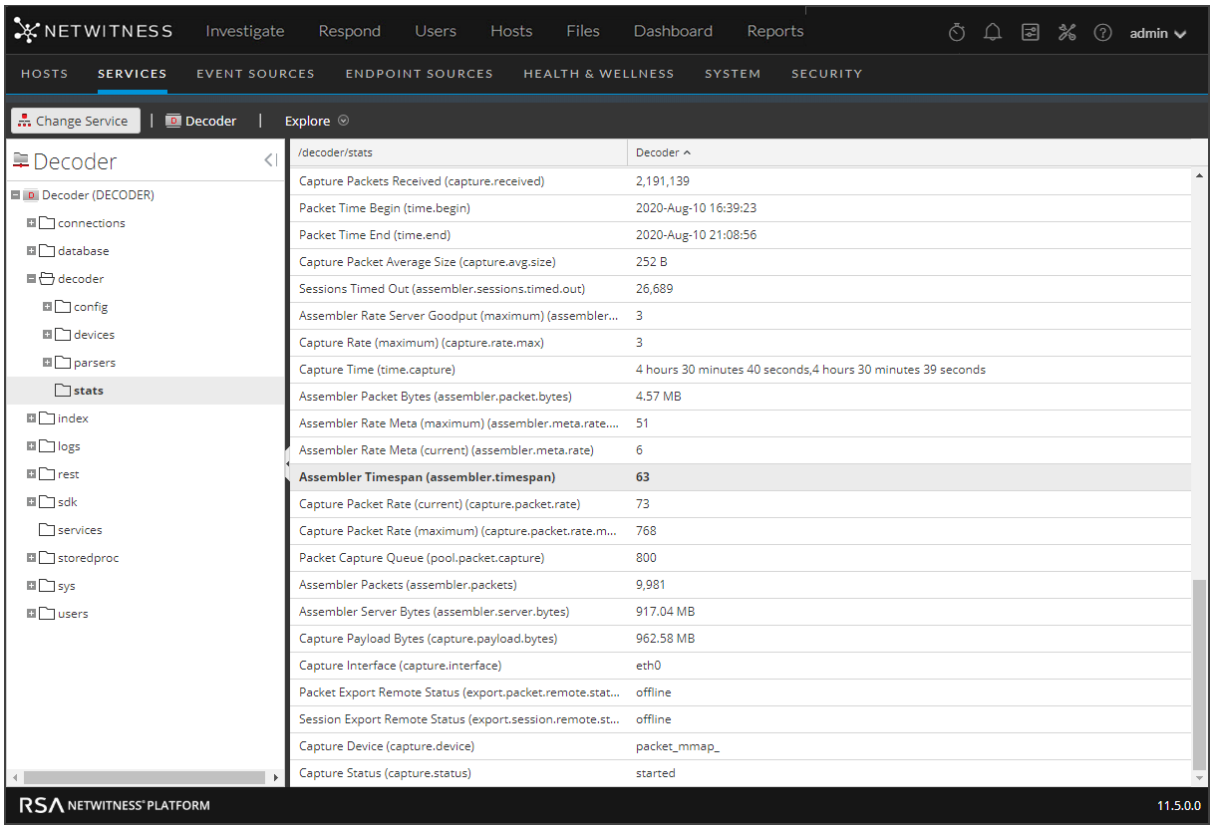
To configure Time Split mode, set the following configuration parameters and restart the Decoder:

1. In the (Admin) > **Services** view, select the Decoder service and **> View > Explore.**
2. In the Services Explore view select **decoder** > **config**.



3. Click in the **Value** column next to the parameter and set these two parameters :
   `/decoder/config/assembler.session.flush = 0`
   `/decoder/config/assembler.timeout.session = 0`
4. To see when the Decoder is no longer tracking session splits because the ingest rate is too high and the Decoder does not have enough memory to track, view the `/decoder/stats/assembler.timespan` statistic, in the Services Explore view select **decoder** >

stats.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Change Service
Decoder
Explore
Decoder
Decoder (DECODER)
connections
database
decoder
config
devices
parsers
stats
index
logs
rest
sdk
services
storedproc
sys
users
/decoder/stats
Decoder
Capture Packets Received (capture.received) 2,191,139
Packet Time Begin (time.begin) 2020-Aug-10 16:39:23
Packet Time End (time.end) 2020-Aug-10 21:08:56
Capture Packet Average Size (capture.avg.size) 252 B
Sessions Timed Out (assembler.sessions.timed.out) 26,689
Assembler Rate Server Goodput (maximum) (assembler... 3
Capture Rate (maximum) (capture.rate.max) 3
Capture Time (time.capture) 4 hours 30 minutes 40 seconds,4 hours 30 minutes 39 seconds
Assembler Packet Bytes (assembler.packet.bytes) 4.57 MB
Assembler Rate Meta (maximum) (assembler.meta.rate.... 51
Assembler Rate Meta (current) (assembler.meta.rate) 6
Assembler Timespan (assembler.timespan) 63
Capture Packet Rate (current) (capture.packet.rate) 73
Capture Packet Rate (maximum) (capture.packet.rate.m... 768
Packet Capture Queue (pool.packet.capture) 800
Assembler Packets (assembler.packets) 9,981
Assembler Server Bytes (assembler.server.bytes) 917.04 MB
Capture Payload Bytes (capture.payload.bytes) 962.58 MB
Capture Interface (capture.interface) eth0
Packet Export Remote Status (export.packet.remote.stat... offline
Session Export Remote Status (export.session.remote.st... offline
Capture Device (capture.device) packet_mmap_
Capture Status (capture.status) started
RSA NETWITNESS® PLATFORM
11.5.0.0

Configure Syslog Forwarding to Destination

In addition to collecting Syslog messages, you can configure the Log Decoder to forward Syslog messages to another Syslog receiver.

Note: No configuration is necessary to collect Syslog messages on the Log Decoder.

NetWitness forwards Syslog messages after it has parsed the messages and before it writes the messages to the Log Decoder.

Note: You must configure Syslog Forwarding using the steps defined in this topic using the **Explore** view.

The Log Decoder must be in the **Started** state before you can configure Syslog Forwarding. To configure Syslog Forwarding:

1. Configure Log Decoder application layer rules (Application rules) to tag Syslog messages with metadata that instructs NetWitness to forward the messages:

   a. In the **Services** view, select a Log Decoder, and in the Actions column, select > **View** > **Explore**.

   b. Go to the **/decoder/config/rules/application** node, right-click **application,** and click **Properties**.

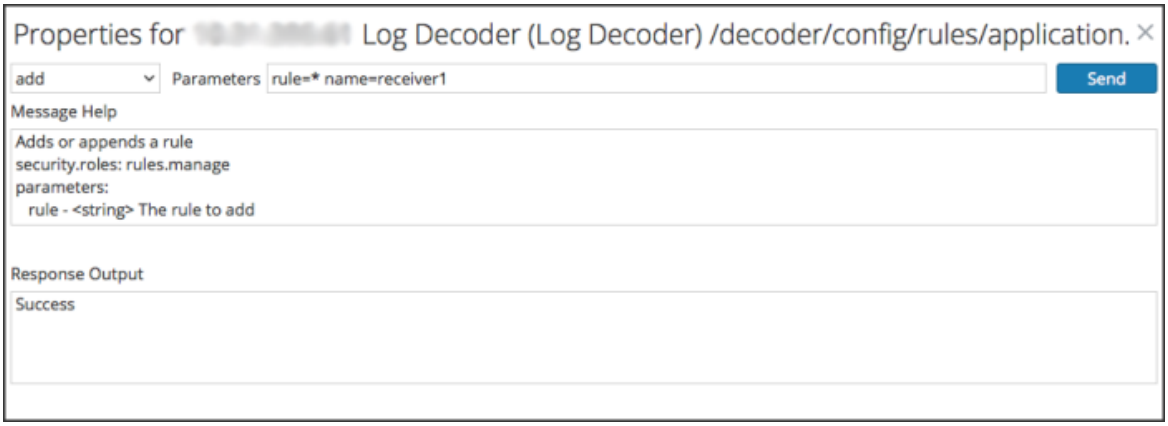
   c. In the **Properties** view, specify the **add** command with the following parameters:

   `rule=<query> name=<name>`

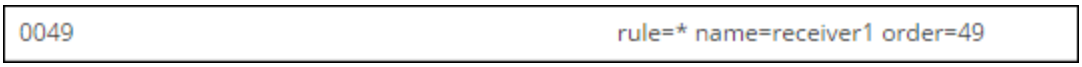
   Example 1: `rule=* name=receiver1`

   Example 2: `rule="device.type='winevent_nic'" name=receiver`)

   d. Click **Send**.



NetWitnesscreates the `name=receiver1 rule=* order=<n>` rule. NetWitness inserts the order number (for example, `order=49`) based on when you set up the rule.



e. Go to the `/decoder/config/rules/application` node and click the `name=receiver1 rule=* order=49` rule.

f. Add **alert forward** parameters to the rule parameters.

```
rule=* name=receiver1 order=49 alert forward
```

or

```
rule=* name=receiver1 forward
```

All other rule parameters have the same meaning as they do in other application rules.

The following Application rule example selects all logs with the * rule. It creates a transient alert meta with the value "receiver1" and tags the entire log for forwarding it to the syslog forwarding destination. You can define as many different forwarding rules as you need with the same name or unique names.

2. Define Syslog forwarding destinations and enable forwarding.

a. In the **Services** view, select a Log Decoder, and  **> View > Explore**.

b. Syslog forwarding destinations are defined in the configuration node.

```
/decoder/config/logs.forwarding.destination
```

This configuration node contains one or more name/value pairs. The name corresponds to the name parameter in the application rule that you used to tag logs with forwarding meta. The value is a colon-separated triple of transport, host, and port followed by an optional formatting parameter.

```
name=(udp|tcp|tls):host:port[:(retainsource|rfc3164)]
```

The first parameter indicates the transport protocol and must be one of udp, tcp, or tls. Specifying udp will forward logs via RFC 3164 / RFC 5426 UDP syslog protocol. Specifying tcp will forward logs via a TCP connection with RFC 6587 framing. Specifying tls will forward logs in accordance with RFC 5425.

The host is an IPv4 address, IPv6 address, or host name.

The port is the port to which the logs are sent. This is typically port 514 for UDP syslog, and 6514 for TLS connections. There is no standard port assignment for syslog over TCP.

Optionally, `retainsource` or `rfc3164` can be specified at the and of the destination string to indicate that additional formatting and information should be included with each log forwarded. Specifying `retainsource` will include z-connector headers at the beginning of the log based and will be populated by the `time`, `device.(ip|ipv6|host)`, and `lc.cid` meta and is best used for forwarding to other log decoders. The `rfc3164` option will prepend a valid RFC3164 header to all events forwarded constructed of the `syslog.pri`, `time`, and `device.(ip|ipv6|host)` meta. In both cases, the original log text is unmodified.

Example forwarding destination:

```
gears=tls:gears.netwitness.local:6514
```

Example forwarding over tcp to blackout on port 514 with z-connector headers:

```
fwdrule=tcp:blackout.netwitness.local:514:retainsource
```

In the `/decoder/config/logs.forwarding.destination` parameter, specify the destination. For example:

TLS Connections: `receiver1=tls:receiver1.netwitness.local:6514`

UDP Connections: `receiver1=udp:receiver1.netwitness.local:514`

TCP Connections: `receiver1=tcp:receiver1.netwitness.local:514`



Note:
You can configure:
- Multiple rules to forward logs to the same destination.
- Multiple rules to forward logs to multiple destination.

For TLS connections, the certificate of the forwarding destination must be validated. The certificate authority that signed the destination's certificate must be present in the Log Decoder's CA trust store and the certificate must reside on the destination or Syslog receiver. Refer to "Configure Certificates" in the *Log Collection Configuration Guide* for information about manipulating the Log Decoder's CA trust store. (Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

c. In the `/decoder/config/logs.forwarding.enabled` parameter, specify **true**.



Configure Transaction Handling on a Decoder

Administrators can configure a Decoder to subdivide incoming sessions into smaller transaction sessions when using Lua parsers designed to create transactions. The feature allows analysts to perform analytics on the split sessions in downstream services such as Investigate.

Caution: Use caution when enabling this feature when you are above the standard ingest rate supported by the Network Decoder as it may start to cause backups with aggregation and dropped packets.

Transaction Handling

The Decoder service configuration node has a new parameter for configuration of transaction handling: `/decoder/parsers/config/parser.transaction.mode`. This node controls the behavior of the Decoder when a parser defines a transaction within a network session.

The values for `parser.transaction.mode` correspond to the operating modes:

- `off` (transactions off)
- `meta` (transactions represented as meta items)
- `split` (transactions split sessions)

Transactions Off

When transactions mode is off, any application-level transactions created by parsers are ignored, and nothing is stored in the collection to represent the transaction.

Transactions Represented as Meta Items

In this mode of operation, when a parser generates an application-level transaction, a new meta item of type `{{trans}}` is added to the session in which the transaction occurred. The `{{trans}}` meta item contains a list of other meta items that constitute the transaction.

Transactions Split Sessions

In this mode of operation, when a parser generates an application-level transaction, the session is split. The session splitting is accomplished by:

1. A new session item is created.
2. Network meta items are copied from the parsed session into the new session.
3. Meta items marked in the transaction are moved from the original session to the new session.

The following meta items are duplicated into the split session from the session that was parsed:

- `time`
- `medium`
- `eth.src`
- `eth.dst`
- `eth.type`
- `ip.proto`
- `ip.src`
- `ip.dst`
- `ipv6.src`
- `ipv6.dst`
- `ip.proto`
- `port.src`
- `port.dst`
- `tcp.flags`
- `udp.srcport`
- `udp.dstport`
- `service`
- `udp.srcport`
- `udp.srcport`
- `tls.premaster`

Configure Data Export

This exports meta data and raw logs from the Decoder or Log Decoder and converts it to Logstash JSON object. To enable data export you must configure the export connector through Logstash.

To configure the data export:

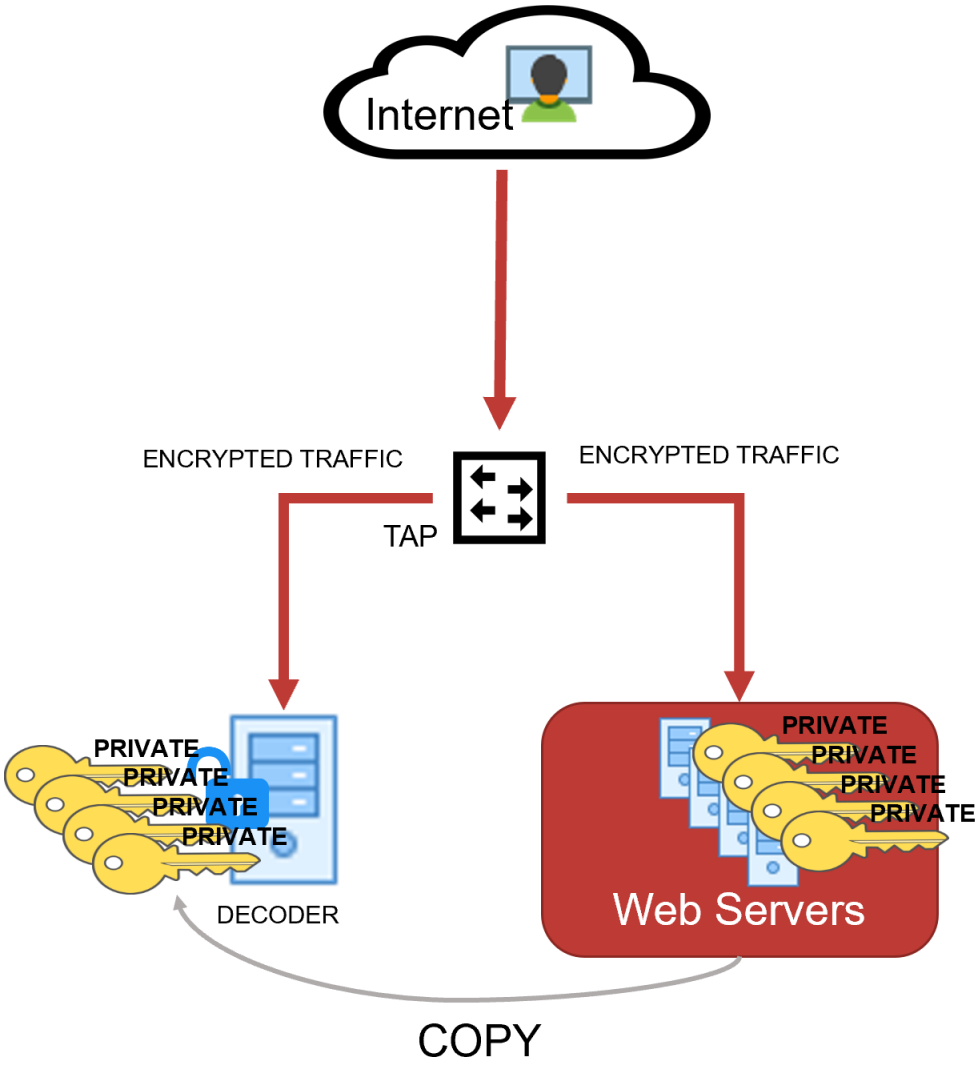
1. Go to  **(Admin) > Services** and select a Decoder or Log Decoder.
2. Select  **> View > Config**.
3. In the Configuration view, select the **Data Export** tab.
4. In the Data Export list, click the name of the host.

The Event Source tab is displayed with export connector typespec selected.

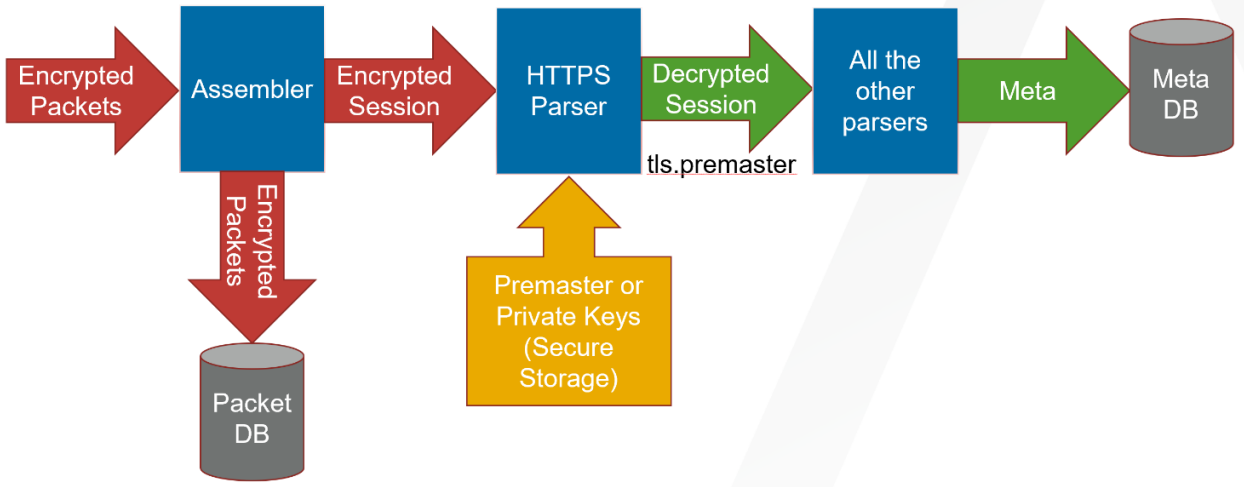
For more information on configuring export connector, see Configure Logstash Event Sources in NetWitness.

Decrypt Incoming Packets (TLS 1.2)

Administrators can configure a Network Decoder to decrypt incoming network packets. This capability is currently supported with encrypted communications to managed servers using TLS 1.2 or earlier and certain cipher suites, which are described in Supported Cipher Suites . The following figure illustrates the supported configuration of a Network Decoder deployed as a passive network device where it captures network traffic for managed servers. The Network Decoder is able to decrypt the traffic by using private keys uploaded to it from the managed servers .



As the network traffic is collected and passed through the system, the encrypted packets are stored as captured even if the system has been able to decrypt them. As the Decoder parses the encrypted traffic, it uses any private keys that have been uploaded, and determines if the network traffic can be decrypted. If the Decoder is able to decrypt the traffic, it generates an additional meta key, defined as `tls.premaster`, which contains the decryption key for that particular session. Whether the traffic is decrypted or not, the metadata `service` is tagged as 443 or SSL/TLS traffic, as the Decoder is port agnostic, even though it is marked with the default port number. The following figure shows the basic progression through the NetWitness platform as the packets are being decrypted and the session metadata is generated.



Follow the steps below to configure the Network Decoder to natively decrypt this network traffic.

- Step 1: Validate That The Network Decoder Captures Encrypted Traffic
- Step 2: Obtain Private Keys from Managed Servers
- Step 3: Validate That The Private Key Cipher Suite is Supported
- Step 4: Confirm HTTPS Parser is Enabled on Decoders
- Step 5: Upload the Supported Private Keys to Decoders
- Step 6: Reload HTTPS Parser on Decoders
- Step 7: View Decrypted Sessions

Note: Currently, you cannot review the decypted payloads directly in the in the user interface. However, you can review them using the RESTful API interface and the NwConsole. However, the Decoders being able to decrypt traffic enables them to detect and generate potentially malicious communications that would otherwise not be detected.

See Troubleshooting for resolving any issues you might encounter while performing these steps.

Step 1: Validate That The Network Decoder Captures Encrypted Traffic

First, make sure the Decoder is actually capturing the SSL/TLS traffic from the managed servers that you want to use to gain visibility into the encrypted sessions. You can do this by searching in the Investigate views (Navigate or Events) using the IP addresses of the servers as the destination IP addresses, and the service as SSL/TLS.

For example, if you are searching for HTTPS traffic to 192.168.1.2 and 192.168.1.3, you can run the following query:

```
service = 443 && ip.dst = 192.168.1.2, 192.168.1.3
```

Conduct this search on the Concentrator that aggregates metadata from the Decoder that is expected to capture this network segment, or on a Broker that may have access to all collected metadata.

(Version 11.6 or later) In case of simultaneous ingestion of the encrypted and decrypted traffic streams in the Decoder, ensure that multi-thread assembly is enabled. This is to prevent mixing of packets from different sessions by single thread assembly. For more information, see the following topics:

- Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder
- (Optional) Multiple Adapter Packet Capture

Step 2: Obtain Private Keys from Managed Servers

Private keys are the asymmetric keys generated by services that accept TLS traffic. They are normally stored in PEM files. These keys are used during the TLS handshake to encrypt the premaster symmetric key that is used for the rest of the payload encryption.

Note: The method to export private keys from web servers (from which to decrypt traffic) varies depending on the product. Refer to the product documentation for instructions for exporting private keys.

Extracting Certificates for the Decoder SSL Decryption

Extracting the proper key materials from popular servers is required to leverage the Decoder built-in decryption capability. The following are a few examples to extract private keys and certificates from popular servers like Apache, Microsoft IIS, and Exchange.

Note: These procedures can also be used as references to extract keys from other servers as well.

Extract private key from Apache

The private key file location would be referenced in the main Apache configuration file, which is `httpd.conf` or `apache2.conf`. The config directive `SSLCertificateKeyFile` will specify the path on your server where private keys are stored.

If Apache is using OpenSSL default certificates and keys, then locate private keys (which are usually saved to the folder /usr/local/ssl by default). You can run the command openssl version –a to find OPENSSLDIR where your server is saving the keys.

If the private key is password protected, then it can be extracted using the below command with the password.

```
openssl rsa -in yourcertificate.pem -out serverkey.pem
```

Validate the located private key using the following command.

```
openssl rsa -in serverkey.pem -check -noout
```

Extract Private Key and Certificate from Windows Servers

Windows servers provide the interface to extract certificates and private keys in PKCS12 format files (.pfx format). It is recommended to use password protection while exporting private keys.

After exporting keys to the pfx file, you can use the OpenSSL tool to extract private keys as PEM files and upload them to Decoder.

Step 2 (a): Extract PKCS12 file from Internet Information Services (IIS)

There are multiple ways to extract certificates from IIS and the following are few examples.

Using IIS Manager App

- Open Internet Information Services Manager App.
- Under IIS > Navigate to Select Certificates > Select Features view > Select the required web or application server certificate.
- Right-click the certificate and select the Export option. In the export wizard, select the file path to save the pfx file (PKCS12 format)
- Example: C:\ServerCerts\servercrt.pfx.
- Enter the password and confirm the password in the export.
- This would create password protected pfx file which contains a private key and associated certificate.

Using Microsoft Management Console (MMC)

- Open Microsoft Management Console (MMC).
- In the Console Root expand Certificates (Local Computer).
- If the Certificates folder is not available then navigate to File > Add Snap In > Add Certificates.
- Locate your server certificate, it is located in the Web Server or Personal sub-folders.
- Right-click the certificate, identified by the Common Name, select Export, and follow the wizard. Use password to protect the pfx file.
- This would create password protected pfx file which contains a private key and associated certificate.

Step 2 (b): Extract PKCS12 file from Exchange Servers

Refer the Export a Certificate from an Exchange Server document for information on how to extract the pfx file from the Exchange server.

Step 2 (c): Extract Private Key From a pfx File Using OpenSSL

OpenSSL tool provides API to export keys in PEM format.

1. Extract private key in PEM format, at the openssl prompt use the password to decrypt the pfx file.
   ```
   openssl pkcs12 -in servercrt.pfx -nocerts -nodes -out serverkey.pem
   ```
2. Validate the private key generated
   ```
   openssl rsa -in serverkey.pem -check -noout
   ```
3. Extract cert in PEM format to validate compatibility with private key
   ```
   openssl pkcs12 -in servercrt.pfx -nokeys -out serverkey.crt
   ```
4. Validate private key and certificate modulus. The modulus should be the same for both the private key and its certificate.
   ```
   openssl rsa -noout -modulus -in serverkey.pem
   openssl x509 -noout -modulus -in serverkey.crt
   ```

Recommendations

Once private keys are extracted from servers and imported to Decoder, it is recommended to secure or delete the extracted PEM file of the private key (serverkey.pem) generated by the above OpenSSL commands.

Step 3: Validate That The Private Key Cipher Suite is Supported

As a passive collection device, the Network Decoder can only decrypt ciphers that use the RSA key exchange. The list of supported keys, and whether the keys are FIPS compliant, are listed in Supported Cipher Suites and Unsupported Cipher Suites . If you are expecting to decrypt less-secure channels of communication, you might need to disable FIPS. However, this is a system change to the entire Decoder, not just the parsing of captured traffic, so take caution before disabling FIPS. If you require more-secure cipher suites for the managed server communications, search for "SSL decrypt" in our third party partner integrations section at RSA Ready Community (https://community.rsa.com/community/products/rsa-ready).

Note: If FIPS is enabled, the list of ciphers for decryption is restricted to only those that are FIPS-compliant.

Supported Cipher Suites

The following table lists which cipher-suites are supported using private keys.

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_RSA_WITH_AES_256_GCM_SHA384	AES256-GCM-SHA384	[0x9d]	TLSv1.2	Kx=RSA	Compliant
TLS_RSA_WITH_AES_256_CBC_SHA256	AES256-SHA256	[0x3d]	TLSv1.2	Kx=RSA	Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_RSA_WITH_AES_256_CBC_SHA	AES256-SHA	[0x35]	SSLv3	Kx=RSA	Compliant
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	CAMELLIA256-SHA	[0x84]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_3DES_EDE_CBC_SHA	DES-CBC3-SHA	[0x0a]	SSLv3	Kx=RSA	Compliant
TLS_RSA_WITH_AES_128_GCM_SHA256	AES128-GCM-SHA256	[0x9c]	TLSv1.2	Kx=RSA	Compliant
TLS_RSA_WITH_AES_128_CBC_SHA256	AES128-SHA256	[0x3c]	TLSv1.2	Kx=RSA	Compliant
TLS_RSA_WITH_AES_128_CBC_SHA	AES128-SHA	[0x2f]	SSLv3	Kx=RSA	Compliant
TLS_RSA_WITH_SEED_CBC_SHA	SEED-SHA	[0x96]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	CAMELLIA128-SHA	[0x41]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_IDEA_CBC_SHA	IDEA-CBC-SHA	[0x07]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_RC4_128_SHA	RC4-SHA	[0x05]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_WITH_DES_CBC_SHA	DES-CBC-SHA	[0x09]	SSLv3	Kx=RSA	Non-Compliant
TLS_RSA_EXPORT_WITH_DES40_CBC_SHA	EXP-DES-CBC-SHA	[0x08]	SSLv3	Kx=DES	Non-Compliant

Unsupported Cipher Suites

The following table lists which cipher-suites are not supported using private keys.

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	ECDHE-RSA-AES256-GCM-SHA384	[0xc030]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	ECDHE-ECDSA-AES256-GCM-SHA384	[0xc02c]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	ECDHE-RSA-AES256-SHA384	[0xc028]	TLSv1.2	Kx=ECDH	Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	ECDHE-ECDSA-AES256-SHA384	[0xc024]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDHE-RSA-AES256-SHA	[0xc014]	SSLv3	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	ECDHE-ECDSA-AES256-SHA	[0xc00a]	SSLv3	Kx=ECDH	Compliant
TLS_DHE_DSS_WITH_AES_256_GCM_SHA384	DHE-DSS-AES256-GCM-SHA384	[0xa3]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	DHE-RSA-AES256-GCM-SHA384	[0x9f]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	DHE-RSA-AES256-SHA256	[0x6b]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256	DHE-DSS-AES256-SHA256	[0x6a]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_256_CBC_SHA	DHE-RSA-AES256-SHA	[0x39]	SSLv3	Kx=DH	Compliant
TLS_DHE_DSS_WITH_AES_256_CBC_SHA	DHE-DSS-AES256-SHA	[0x38]	SSLv3	Kx=DH	Compliant
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	DHE-RSA-CAMELLIA256-SHA	[0x88]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA	DHE-DSS-CAMELLIA256-SHA	[0x87]	SSLv3	Kx=DH	Non-Compliant
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	ECDH-RSA-AES256-GCM-SHA384	[0xc032]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	ECDH-ECDSA-AES256-GCM-SHA384	[0xc02e]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	ECDH-RSA-AES256-SHA384	[0xc02a]	TLSv1.2	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	ECDH-ECDSA-AES256-SHA384	[0xc026]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	ECDH-RSA-AES256-SHA	[0xc00f]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	ECDH-ECDSA-AES256-SHA	[0xc005]	SSLv3	Kx=ECDH/ECDSA	Compliant

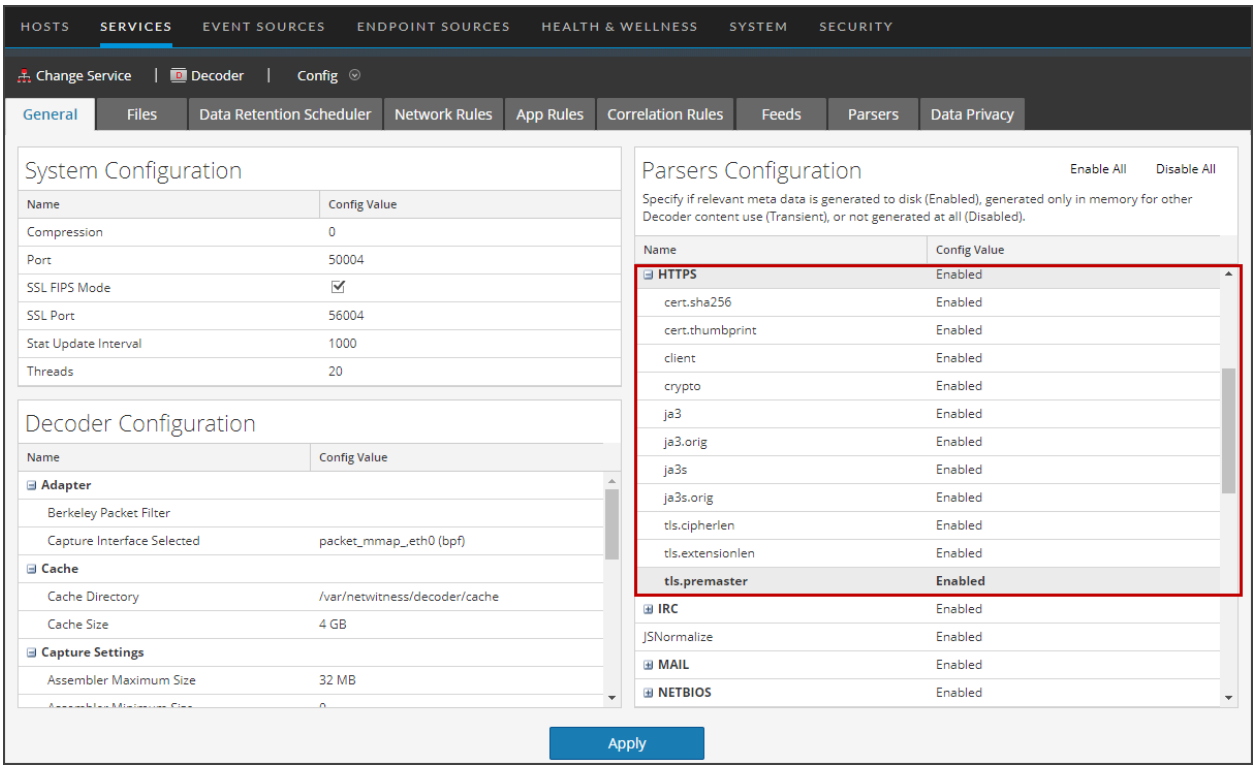
Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	ECDHE-RSA-DES-CBC3-SHA	[0xc012]	SSLv3	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	ECDHE-ECDSA-DES-CBC3-SHA	[0xc008]	SSLv3	Kx=ECDH	Compliant
TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	EDH-RSA-DES-CBC3-SHA	[0x16]	SSLv3	Kx=DH	Compliant
TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	EDH-DSS-DES-CBC3-SHA	[0x13]	SSLv3	Kx=DH	Compliant
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	ECDH-RSA-DES-CBC3-SHA	[0xc00d]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	ECDH-ECDSA-DES-CBC3-SHA	[0xc003]	SSLv3	Kx=ECDH/ECDSA	Compliant
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	ECDHE-RSA-AES128-GCM-SHA256	[0xc02f]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	ECDHE-ECDSA-AES128-GCM-SHA256	[0xc02b]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	ECDHE-RSA-AES128-SHA256	[0xc027]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	ECDHE-ECDSA-AES128-SHA256	[0xc023]	TLSv1.2	Kx=ECDH	Compliant
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDHE-RSA-AES128-SHA	[0xc013]	SSLv3	Kx=ECDH	Compliant
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	ECDHE-ECDSA-AES128-SHA	[0xc009]	SSLv3	Kx=ECDH	Compliant
TLS_DHE_DSS_WITH_AES_128_GCM_SHA256	DHE-DSS-AES128-GCM-SHA256	[0xa2]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	DHE-RSA-AES128-GCM-SHA256	[0x9e]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	DHE-RSA-AES128-SHA256	[0x67]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_DSS_WITH_AES_128_CBC_SHA256	DHE-DSS-AES128-SHA256	[0x40]	TLSv1.2	Kx=DH	Compliant
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	DHE-RSA-AES128-SHA	[0x33]	SSLv3	Kx=DH	Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_DHE_DSS_WITH_AES_128_CBC_SHA	DHE-DSS-AES128-SHA	[0x32]	SSLv3	Kx=DH	Compliant
TLS_DHE_RSA_WITH_SEED_CBC_SHA	DHE-RSA-SEED-SHA	[0x9a]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_DSS_WITH_SEED_CBC_SHA	DHE-DSS-SEED-SHA	[0x99]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA	DHE-RSA-CAMELLIA128-SHA	[0x45]	SSLv3	Kx=DH	Non-Compliant
TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA	DHE-DSS-CAMELLIA128-SHA	[0x44]	SSLv3	Kx=DH	Non-Compliant
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	ECDH-RSA-AES128-GCM-SHA256	[0xc031]	TLSv1.2	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	ECDH-ECDSA-AES128-GCM-SHA256	[0xc02d]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	ECDH-RSA-AES128-SHA256	[0xc029]	TLSv1.2	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	ECDH-ECDSA-AES128-SHA256	[0xc025]	TLSv1.2	Kx=ECDH/ECDSA	Compliant
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	ECDH-RSA-AES128-SHA	[0xc00e]	SSLv3	Kx=ECDH/RSA	Compliant
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	ECDH-ECDSA-AES128-SHA	[0xc004]	SSLv3	Kx=ECDH/ECDSA	Compliant
TLS_ECDHE_RSA_WITH_RC4_128_SHA	ECDHE-RSA-RC4-SHA	[0xc011]	SSLv3	Kx=ECDH	Non-Compliant
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	ECDHE-ECDSA-RC4-SHA	[0xc007]	SSLv3	Kx=ECDH	Non-Compliant
TLS_ECDH_RSA_WITH_RC4_128_SHA	ECDH-RSA-RC4-SHA	[0xc00c]	SSLv3	Kx=ECDH/RSA	Non-Compliant
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	ECDH-ECDSA-RC4-SHA	[0xc002]	SSLv3	Kx=ECDH/ECDSA	Non-Compliant
TLS_DHE_RSA_WITH_DES_CBC_SHA	EDH-RSA-DES-CBC-SHA	[0x15]	SSLv3	Kx=RSA	Non-Compliant
TLS_DHE_DSS_WITH_DES_CBC_SHA	EDH-DSS-DES-CBC-SHA	[0x12]	SSLv3	Kx=DSS	Non-Compliant
TLS_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA	EXP-EDH-RSA-DES-CBC-SHA	[0x14]	SSLv3	Kx=DSS	Non-Compliant

Cipher Suite Name (RFC)	Name (OpenSSL)	Cipher Suite	TLS Version	KeyExch.	FIPS
TLS_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA	EXP-EDH-DSS-DES-CBC-SHA	[0x11]	SSLv3	Kx=DSS	Non-Compliant

Step 4: Confirm HTTPS Parser is Enabled on Decoders

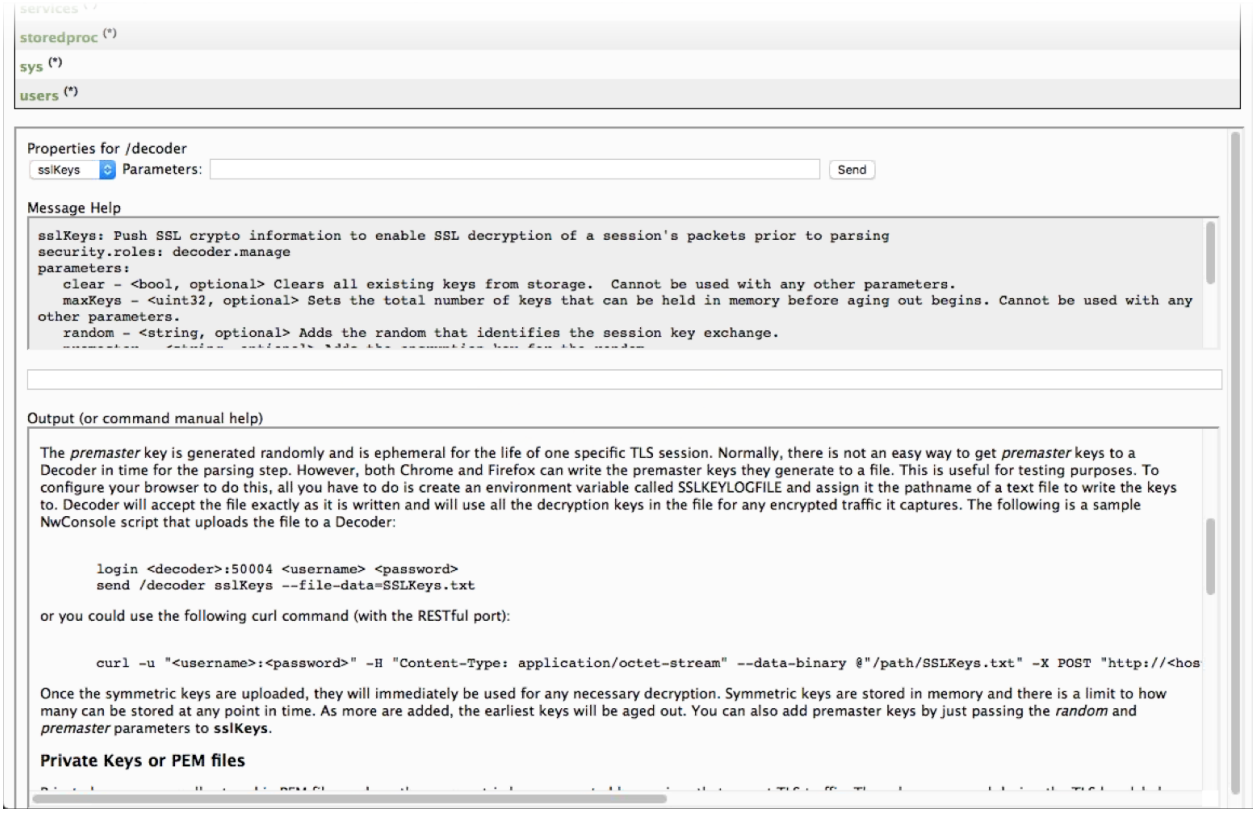
The decryption process requires the native HTTPS parser to be enabled to examine the TLS handshake and to compare the uploaded encryption keys to determine if they decrypt the traffic. This can be validated in (Admin) > **Services** > **Decoder** > **Config** > Parsers Configuration panel as shown in the following figure.



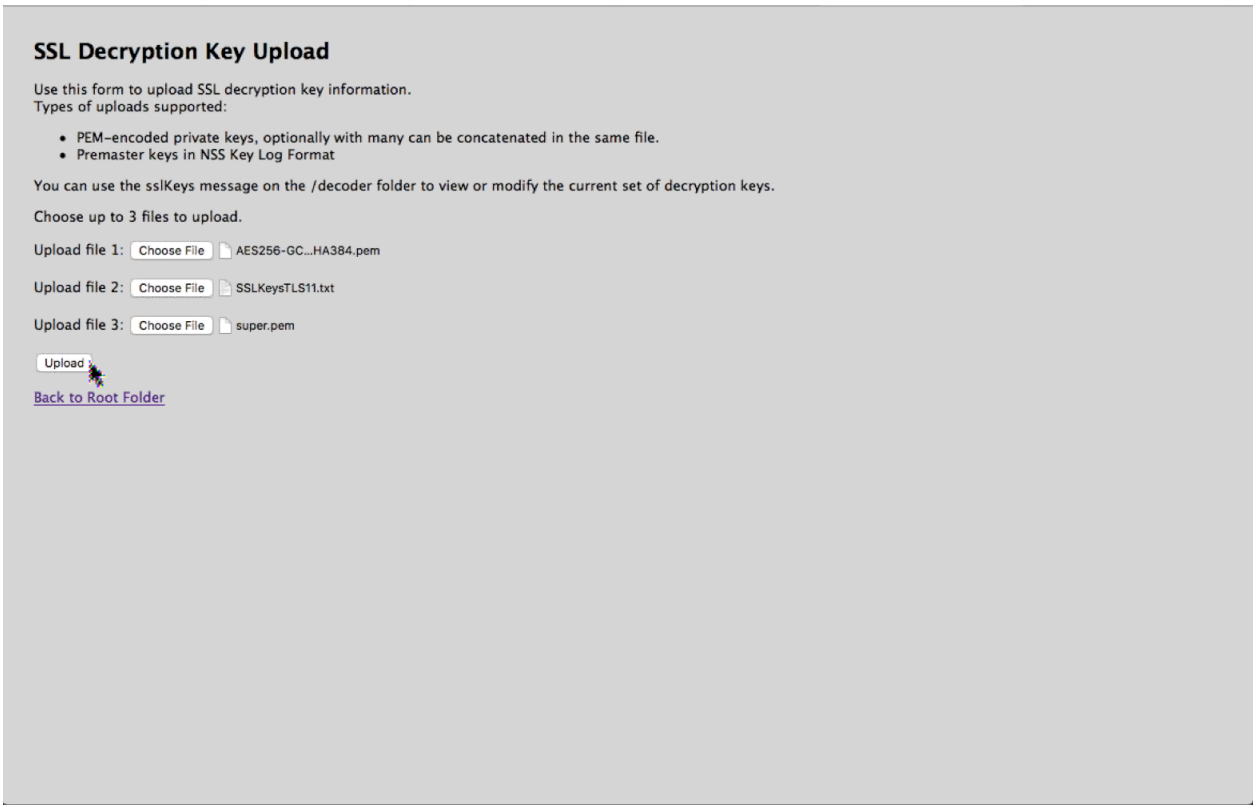
Step 5: Upload the Supported Private Keys to Decoders

The `sslKeys` command provides a way to upload premaster or private keys to the Decoder, so that captured encrypted packets that match the keys can be decrypted before parsing. Administrators configure the Decoder by entering the `sslKeys` command using the NwConsole command line interface or the Decoder RESTful interface. After upload, you must issue a parser reload command so that the newly installed key becomes visible to the HTTPS parser. Now, all TLS handshakes that use that private key will be able to be decrypted by the Decoder.

For information about how to use the NwConsole command line interface, see the *NwConsole User Guide for RSA NetWitness Platform*. For information about how to use the RESTful interface, see the *RESTful API User Guide for RSA NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.



The RESTful interface form at the path `/decoder/sslkeys` allows uploading a single PEM-encoded private key, a single file containing multiple private keys concatenated together, or a single file of multiple premaster keys.

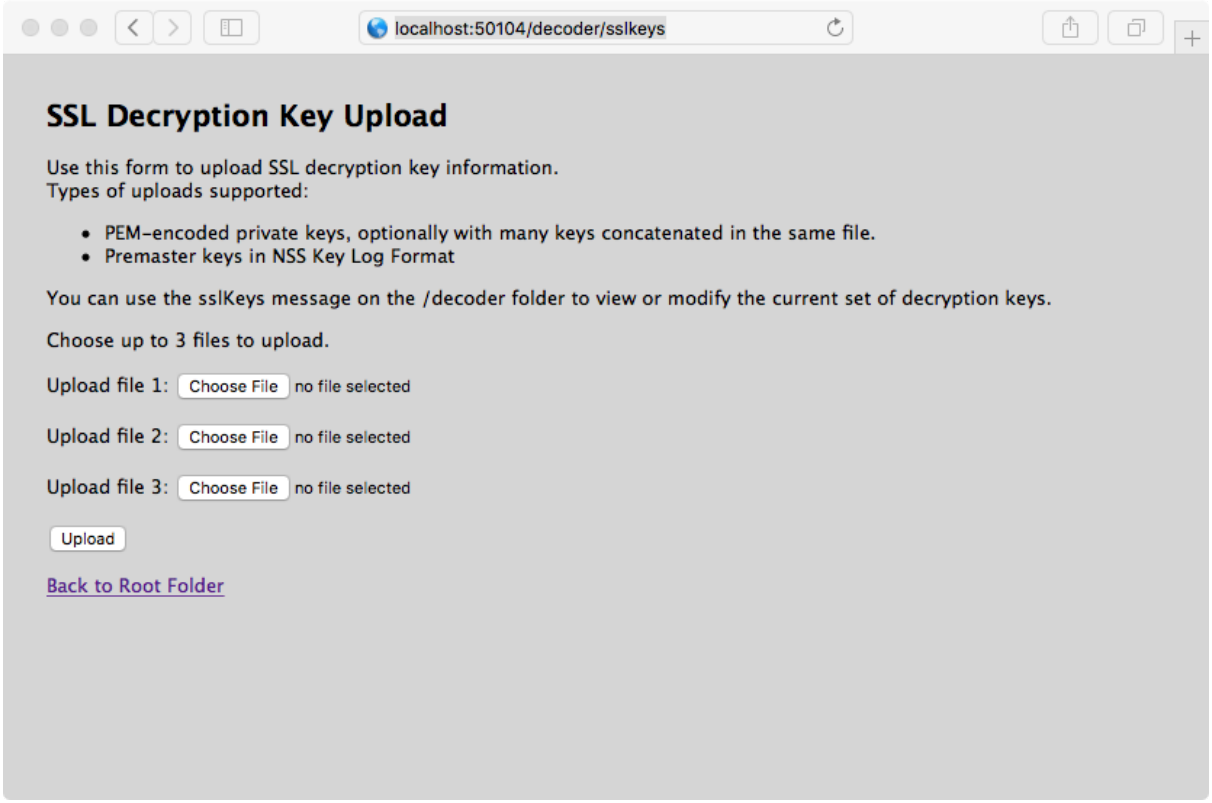


Upload Multiple Premaster and Private Keys

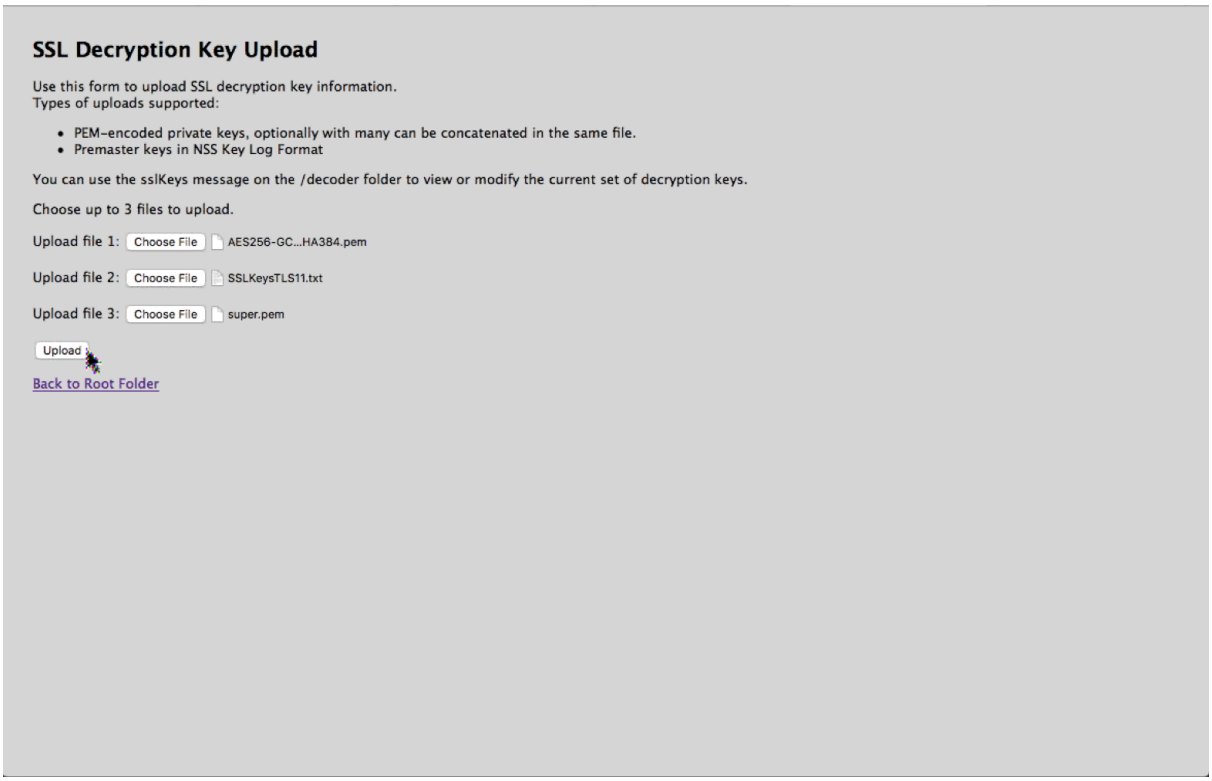
Note: As of 11.5, uploading premaster keys is not officially supported yet for real-time capture, but can be used to validate that decryption is functioning properly on the Decoder.

You can use the RESTful interface form to facilitate uploading of multiple keys, both premaster and private at the same time.

1. Open the RESTful API in your browser and navigate to this path on the Decoder that you want to configure: `/decoder/sslkeys`.

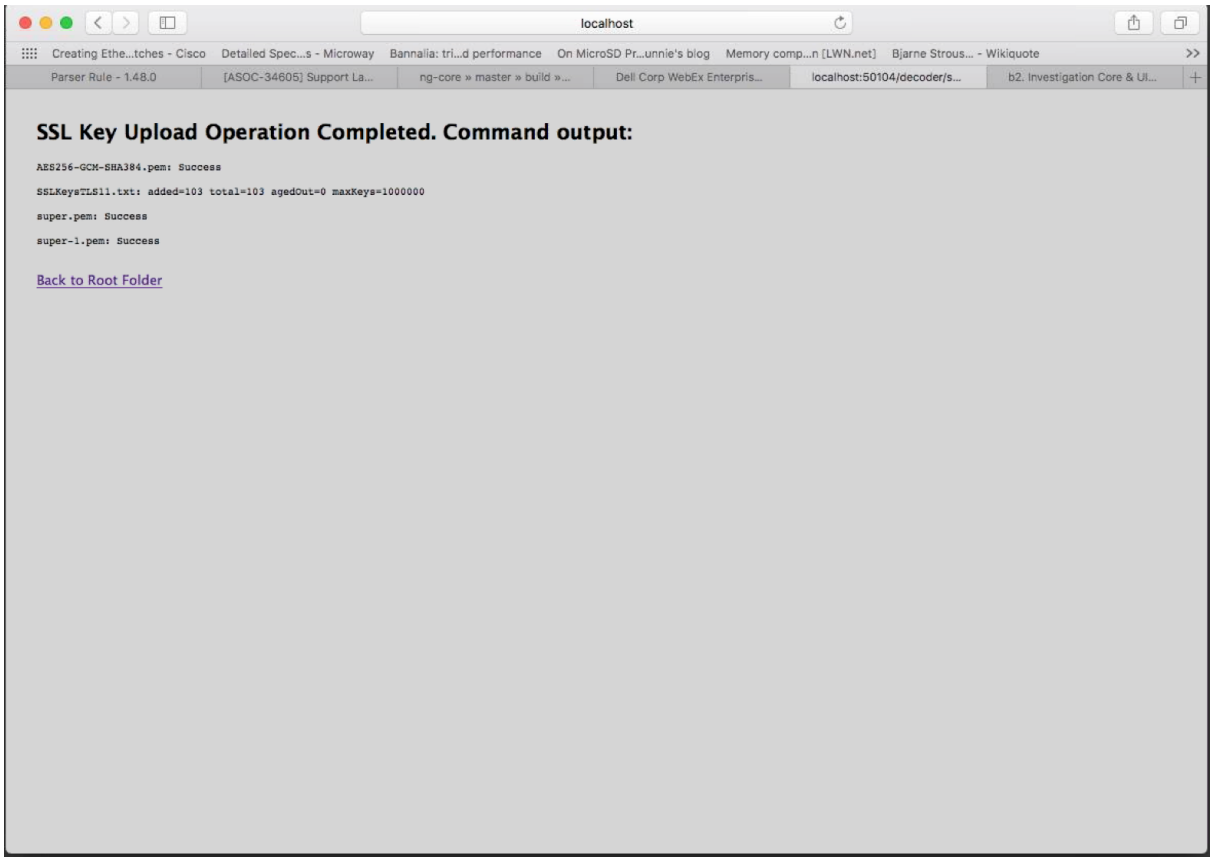


2. Next to **Upload File 1**, click **Choose File** and locate the premaster key file or PEM file that you want to upload on your local file system.
3. (Optional) Repeat for **Upload File 2** and **Upload File 3**.



4. Click **Upload**.

   The files are uploaded to the Decoder and results are displayed in the form.



Parameters for Managing Keys

The `sslKeys` command has several parameters for managing premaster and private keys. This is the full list of parameters:

Parameter	Description
`clear`	Removes all premaster keys from memory. Does not delete any PEM files installed on the system.
`maxKeys`	Changes the maximum number of premaster keys that are stored in memory.
`listPems`	Returns a list of all installed private key PEM files.
`deletePem`	Deletes the named PEM file from the file system. You can pass this parameter more than once to remove multiple files.
`random`	The random hash used to identify the premaster key.
`premaster`	The premaster key that will be installed for the previous `random` parameter. They must show up in pairs and `random` must be first.

Return Values

Most `sslKeys` commands return name/value pairs of statistics about the premaster keys in memory. The statistics are listed in the following table.

Name	Description
`added`	The number of premaster keys just added during this command.
`total`	The total number of premaster keys loaded in memory.
`agedOut`	The total number of premaster keys that were removed during this command; this is not a lifetime stat.
`maxKeys`	The current maximum allowed premaster keys

Encryption Keys

The `sslKeys` command accepts two types of encryption keys:

- Premaster key: the symmetric key used in the TLS payload stream for encryption and decryption.
- Private key: the asymmetric private key used during the TLS handshake that encrypts the premaster.

Premaster Key

The premaster key is generated randomly and is ephemeral for the life of one specific TLS session. Normally, there is not a good way to get premaster keys to a Decoder in time for the parsing step. However, both Chrome and Firefox can write the premaster keys they generate to a file. This is useful for testing purposes. To configure your browser to do this, create an environment variable called `SSLKEYLOGFILE` and assign it the pathname of a file to which the keys will be written. The Decoder will accept the file exactly as written and will use all the decryption keys in the file for any encrypted traffic it captures.

This is a sample NwConsole script that uploads the file to a Decoder:

```
login <decoder>:50004 <username> <password>
send /decoder sslKeys --file-data=SSLKeys.txt
```

This is an example using a curl command (with the RESTful port) to upload the file to a Decoder:

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-stream" --
data-binary @"/path/SSLKeys.txt" -X POST
"http://<hostname>:50104/decoder?msg=sslKeys"
```

After the symmetric keys are uploaded, they will immediately be used for any necessary decryption. Symmetric keys are stored in memory and there is a limit to how many can be stored at any point in time. As more keys are added, the earliest keys will be aged out. You can also add premaster keys by just passing the `random` and `premaster` parameters to `sslKeys`.

Private Keys or PEM Files

Private keys are normally stored in PEM files and are the asymmetric keys generated by services that accept TLS traffic. These keys are used during the TLS handshake to encrypt the premaster symmetric key that will be used for the rest of the payload encryption.

For example, if you have a web server whose traffic you want visibility into, you need to upload the private key it uses to encrypt traffic. You only need to do this once, as it is stored permanently (or until removed by a delete command). Private keys are automatically encrypted before storing to protect them. After upload, you must issue a parser reload command so that the newly installed key becomes visible to the HTTPS parser. Now, all TLS handshakes that use that private key will be able to be decrypted by the Decoder.

Note: Not all ciphers suites use a "known" private key (for example, Ephemeral Diffie Hellman). Encrypted traffic with those ciphers cannot be decrypted unless the premaster key is uploaded to the Decoder before the session is parsed.

These are some sample commands that upload a PEM file to be used for decryption.

Using NwConsole:

```
send /decoder sslKeys pemFilename=MyKey.pem --file-data=/path/MyKey.pem
```

Using the RESTful interface (you must provide the `pemFilename` parameter in the URL):

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-stream" --
data-binary @"/path/MyKey.pem" -X POST
"http://<hostname>:50104/decoder?msg=sslKeys&pemFilename=MyKey.pem"
```

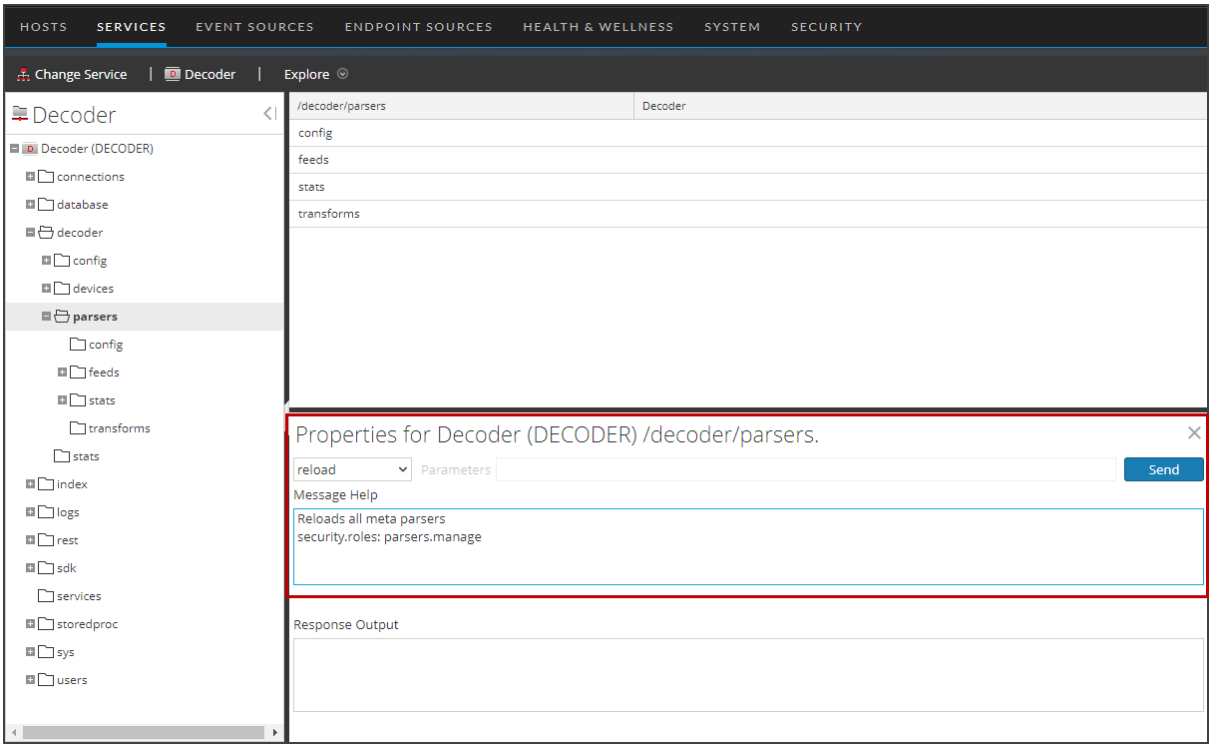
Step 6: Reload HTTPS Parser on Decoders

After uploading the private keys, you must issue a parser reload command so that the newly-installed key becomes visible to the HTTPS parser. All TLS handshakes that use that private key are now able to be decrypted by the Decoder.

To issue a parser reload command:

1. Go to  (Admin) > **Services** and select the Decoder.
2. Click  > **View** > **Explore** > **decoder**, right-click **parsers** and select **Properties**.

3. In the Properties for Decoder (DECODER) /decoder/parsers pane, select reload from the drop-down menu and click **Send**.



Step 7: View Decrypted Sessions

All storage of the packets is still encrypted, so when you search for any decrypted traffic, your query should still focus on `service=443` or HTTPS (HTTP over SSL/TLS). If any packets were able to be decrypted, `tls.premaster` metadata is associated with the session that has the premaster key required for decrypting that session. The session will also have crypto metadata denoting the cipher suite, with crypto formated as `tls.premaster`.

Viewing Unencrypted Traffic

If packets are decrypted during the parse stage, encrypted packets are written to disk, and the matching premaster key used for decrypting is written to the `tls.premaster` meta key, analysts can view the unencrypted packets using the `tls.premaster` meta key.

One Decoder API that you can use to see the unencrypted packets is the `/sdk/content` RESTful service. You need to know the Session ID of the encrypted packets and the `flags` parameter masked to the value 128 (or 0x80 in hex). Point your browser to the Decoder RESTful interface and type in the following command, substituting the actual Session ID for `<id>`:

`http://<decoder>:50104/sdk/content?session=<id>&flags=128&render=text`

The Decoder returns a simple web page showing the packets after they are decrypted.

If you want to see what the packets look like encrypted, type in one of the following commands, substituting the Session ID for <id>:

`http://<decoder>:50104/sdk/content&session=<id>&render=text`

`http://<decoder>:50104/sdk/content&session=<id>&flags&render=text`

For more information on the `/sdk/content` service, see the manual page for `/sdk content`.

Alternatively, you can use the NetWitness Investigator thick client to decrypt and display the network sessions, as it can communicate to a Concentrator or have PCAPs and private keys loaded locally.

If you are not able to view the encrypted sessions, check Troubleshooting. However, if you are able to see the `tls.premaster` metadata, that confirms that the Decoder is able to see inside the encrypted session, which enables parsers to see the unencrypted packet payload and create metadata accordingly. For these decrypted sessions, look for additional metadata around usernames, passwords, or HTTP body content that would be obfuscated when encrypted, which would result in metadata not being generated.

Troubleshooting

If you are still not able to decrypt or view the decrypted payloads after attempting the steps above, try these steps.

1. Validate that your private key is using an acceptable cipher-suite.
2. Make sure your Network Decoder is capturing the traffic matching the server for which you have the private key. To check, you can upload the private key and PCAP into Wireshark and see if it is able to decrypt it.

Decryption of Secure SMTP

In version 11.3, NetWitness supports opportunistic SSL/TLS decryption, which addresses RFC 3207 (https://tools.ietf.org/html/rfc3207). You can add an HTTPS parser option that provides a .csv list of destination ports of the session where the STARTTLS command will be searched, with at least one encryption key that has been uploaded. This enables the STARTTLS functionality.

Note: The .csv list must include all the destination ports of a session. If a destination port is not in the list, the search for STARTTLS will not be started.

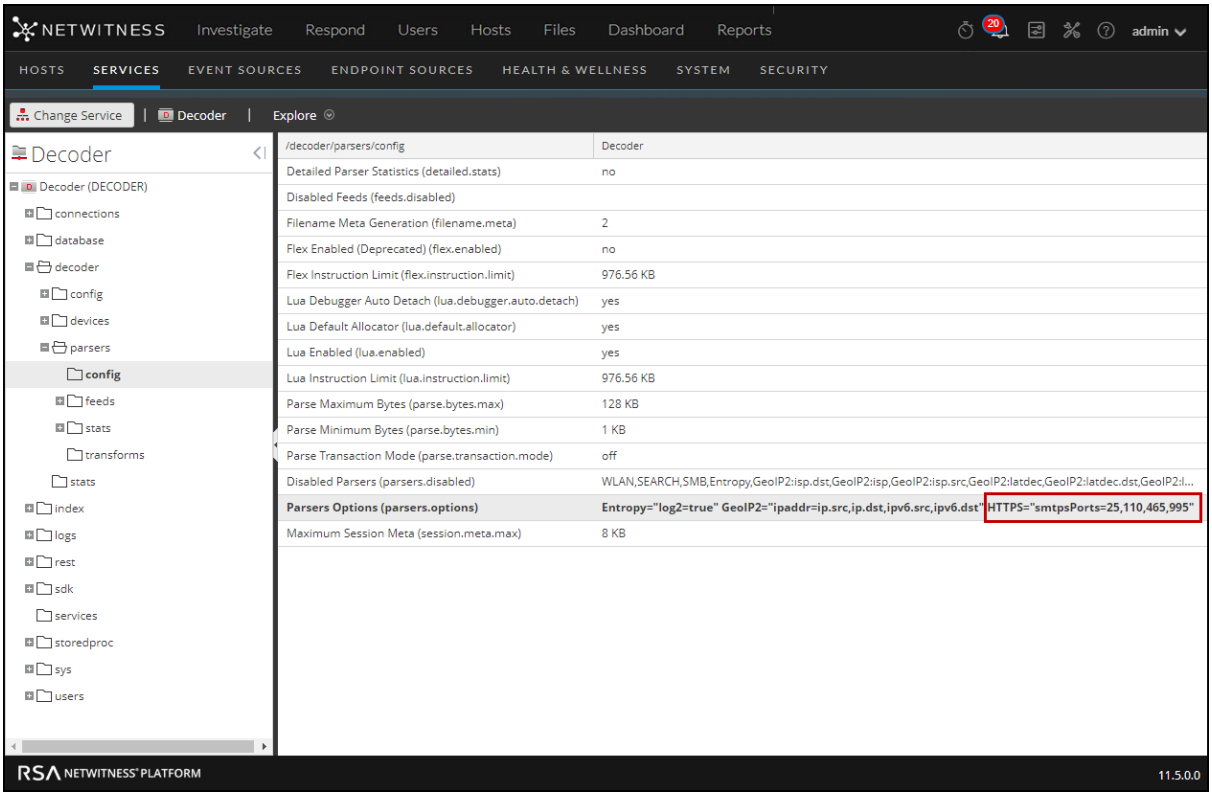
The STARTTLS search is limited to the first 1024 bytes of the session's payload, to prevent performance degradation. If STARTTLS is found, the parser uses the next client packet with payload as the start of a TLS negotiation.

To add an HTTPS parser option for secure SMTP decryption:

1. In the NetWitness User Interface, go to **(Admin)** > **Services**.
2. Select a Log Decoder service and click > **View** > **Explore**.
3. In the left panel, select **decoder** > **parsers** > **config**.
4. In **parsers.options**, enter an HTTPS parser option in the following format that provides a .csv list of the destination ports of the session where the STARTTLS command is searched:

   `HTTPS="smtpsPorts=<port#1,port#2,port#n>"`

The following image shows an example of an HTTPS parser option for secure SMTP decryption.



Performance Considerations

Decrypting packets in real time requires extra work in the parsing stage. Before implementing this feature, plan carefully to ensure the incoming traffic bandwidth does not overwhelm the available compute power. You may need more Decoders to decrypt traffic than you would need if not decrypting.

Packets captured on a Decoder normally have a timeout of ~60 seconds in the assembly stage before they are sent to the parsing step. If the Decoder is under memory pressure due to very high bandwidth, the lifetime of the packets in Assembler may be shortened. To alleviate this situation, you can configure a longer timeout value and increase the amount of memory available to hold packets in Assembly. Also, in order to perform decryption of the packets, the Decoder must receive the decryption key before the parsing stage.

Note: Currently, only TLS 1.2 and earlier protocols can be decrypted

With no feeds loaded, the following parsers enabled, and 50% of the sessions being decrypted, a Decoder can process traffic at 3 Gbps .

Parser Name	Description
SYSTEM	Session Details
NETWORK	Network Layer
ALERTS	Alerts
GeoIP	Geographic data based on ip.src and ip.dst
GeoIP2	Geographic data by default based on IPv4 (`ip.src`, `ip.dst`) and IPv6 (`ipv6.src`, `ipv6.dst`) meta keys
HTTP	Hyper Text Transport Protocol (HTTP)
HTTP_Lua	Hyper Text Transport Protocol (HTTP) Lua
FTP	File Transfer Protocol (FTP)
TELNET	TELNET Protocol
SMTP	Simple Mail Transport Protocol (SMTP)
POP3	Post Office Protocol (POP3)
NNTP	Network News Transport Protocol (NNTP)
DNS	Domain Name Service (DNS)
HTTPS	Secure Socket Layer (SSL) Protocol
MAIL	Standard E-Mail Format (RFC822)
VCARD	Extracts VCARD fullname and email information
PGP	Identifies PGP blocks within network traffic
SMIME	Identifies SMIME blocks within network traffic
SSH	Secure Shell (SSH)
TFTP	Trivial File Transfer Protocol (TFTP)
DHCP	Dynamic Host Configuration Protocol (DHCP and BOOTP)
NETBIOS	Extracts NETBIOS computer name information.
SNMP	Simple Network Management Protocol (SNMP)
NFS	Network File System (NFS) protocol
RIP	Routing Information Protocol (RIP).
TDS	MSSQL and Sybase database protocol (TDS)

Parser Name	Description
TNS	Oracle database protocol (TNS)
IRC	Internet Relay Chat (IRC) protocol
RTP	Real Time Protocol (RTP) for audio/video
SIP	Session Initiation Protocol (SIP)
H323	H.323 Teleconferencing protocol
SCCP	Cisco Skinny Client Control Protocol
GTalk	Google Talk (GTalk)
VlanGre	Vlan ID and GRE/EtherIP tunnel addresses
BITTORRENT	BitTorrent File Sharing Protocol
FIX	Financial Information eXchange Protocol
GNUTELLA	Gnutella file sharing protocol
IMAP	Internet Message Access Protocol
MSRPC	Microsoft Remote Procedure Call protocol
RDP	Remote Desktop Protocol
SHELL	Command Shell Identification
TLSv1	TLSv1
SearchEngines	A parser that extracts search terms
FeedParser	External Feed Parser

TLS Certificate Hashing

The Network Decoder can produce hashes of certificates that are seen in the packet stream. These hashes are the SHA-1 value of any DER-encoded certificate encountered during a TLS handshake. The hashes produced can be used to compare network traffic with hashes from public SSL blacklists, such as the one from sslbl.abuse.ch.

The TLS Certificate hashing feature is disabled by default. To enable TLS Certificate hashing:

1. Go to  **(Admin)** > **Services**, select a Network Decoder service and  > **View > Explore.**
2. Select **decoder** > **parsers** > **config**.
3. In the values column next to `parsers.options`, type `HTTPS="cert.sha1=true"`.
   When this option is enabled, the SHA-1 is stored as a text value in the `cert.thumbprint` meta key.

4. The change takes effect on parser reload. In the left panel, right-click **parsers** and click **Properties**. In the drop-down menu, select **reload** and then click **Send**

The SHA-256 hash of the certificate can be enabled by adding the parser option:

```
HTTPS="cert.sha256=true"
```

to a Network Decoder's `/decoder/parsers/config/parsers.options` configuration.

When this option is enabled, SHA-256 is stored as a text value in the meta keys:

```
cert.sha256
cert.thumbprint
```

JA3 and JA3S TLS Fingerprints

The Network Decoder can produce the JA3 value of TLS clients and the JA3S value of TLS servers that are observed in a Network session. The values that are produced conform to the values generated by the open source JA3 tools (https://github.com/salesforce/ja3). The JA3 features are disabled by default.

To enable JA3 features:

1. In the NetWitness Platform User Interface, go to  **(Admin)** > **Services**.
2. Select a Decoder service and click  > **View** > **Explore**.
3. In the left panel, select **decoder** > **parsers** > **config**.
4. In **parsers.options**, add `HTTPS="ja3=true ja3s=true"`.
5. In the left panel, right click on **parsers** and select **properties**.
6. In the right panel, in the lower pane, set the value in the first drop down to **Reload**.
7. Click **Send**.

JA3 and JA3S can be enabled independently using these options. When enabled, the JA3 is stored as a text value in the meta key `ja3`. The JA3S is stored as a text value in the meta key `ja3s`.

Additional meta keys are activated using a mechanism similar to the JA3S and JA3S meta items. For example, `ja3.orig` is enabled by adding `ja3.orig=true` to the HTTPS parser options. This is also true for `ja3s.orig`, `tls.extensionlen`, and `tls.cipherlen`.

`ja3.orig`, `tls.extensionlen`, and `tls.cipherlen` are only created if JA3 is enabled.

`ja3s.orig` is only created if JA3S is enabled.

Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder

In previous versions, ingesting encrypted and decrypted traffic streams of the same session to a Decoder causes packets from both the sessions to mix on the Decoder as reassembly is performed by the same thread. It cause session parsing and content extraction inaccuracies.

In NetWitness 11.6, using multi-adapter capture and multi-thread assembler features, you can configure the Decoder to receive encrypted and decrypted streams on separate adapters. The multi-thread assembler feature allows decoder to assemble packets from its corresponding capture work thread. It keeps the packets from encrypted and decrypted sessions separate during assembly and avoid inaccuracies in session parsing and content extraction.

To configure the Decoder for the simultaneous ingestion of the encrypted and decrypted traffic streams:

1. Configure Decoder to capture from multiple adapters simultaneously, see (Optional) Multiple Adapter Packet Capture.
2. Set `/decoder/config/assembler.threading.enabled=on` to enable multi-thread assembler.
3. Restart the Decoder service.

Community ID

The Network Decoder generates Community ID flow hash values that are compatible with the Community ID specification defined by `https://github.com/corelight/community-id-spec`.

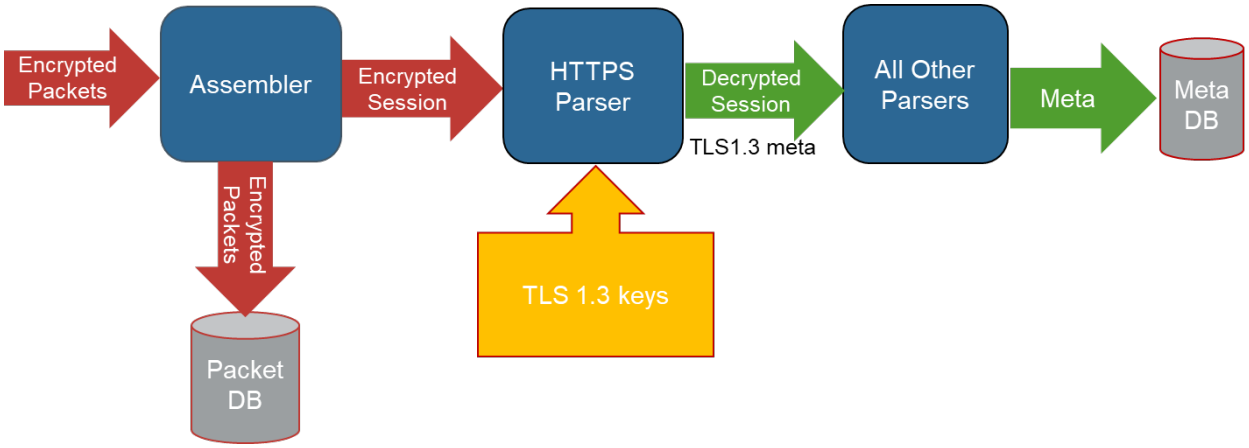
The Community ID is written to a text-formatted meta key named `community`.

The Community ID is generated on sessions containing TCP over IPv4, TCP over IPv6, UDP over IPV4, and UDP over IPV6.

Decrypt Incoming Packets TLS 1.3

Administrators can configure a `nonfips` Network Decoder to decrypt incoming TLS 1.3 network packets. This capability is currently supported with encrypted communications to managed servers using TLS 1.3 or earlier.

For TLS 1.2 and earlier protocols refer Decrypt Incoming Packets (TLS 1.2)



As the network traffic is collected and passed through the system, the encrypted packets are stored as captured even if the system can decrypt them. As the Decoder parses the encrypted traffic, it uses TLS 1.3 keys that have been pushed or uploaded, and determines if the network traffic can be decrypted. If the Decoder is able to decrypt the traffic, it generates an additional meta keys, defined as `tls.client.hdshk, tls.server.hdshk, tls.client.app, tls.server.app, tls.client.early` which contains the decryption keys for that particular session. Whether the traffic is decrypted or not, the metadata service is tagged as 443 or SSL/TLS traffic, as the Decoder is port agnostic, even though it is marked with the default port number.

Non-FIPS Builds

Network Decoder TLS 1.3 decryption is supported ONLY through special `nonfips` rpm builds which run in NON FIPS mode. The `nonfips` builds use crypto library which has pending FIPS certification in NIST, hence they are labeled as `nonfips`. If you are restricted to run only FIPS in your environment then TLS 1.3 decryption feature is not supported.

The following nonfips builds for components Decoder, Appliance , Concentrator and Broker are required to be installed, as they provide dependencies and support to TLS 1.3 decryption in session parsing and during content reconstruction.

Decoder (rsa-nw-decoder-nonfips rpm)	Decrypt on Session Parsing and Content Reconstruction
Concentrator (rsa-nw-concentrator-nonfips rpm)	Decrypt on Content Reconstruction

Broker (rsa-nw-broker-nonfips rpm)	Decrypt on Content Reconstruction
Appliance (rsa-nw-appliance-nonfips rpm)	Installs dependencies and scripts for TLS 1.3 Decryption

Install or update Non-FIPS builds

The `rsa-nw-appliance` rpm now includes a script `nw-nonfips-override-generator.sh` to generate `nonfips` override descriptors. It will be installed at the path: `/usr/local/sbin/nw-nonfips-override-generator.sh`

The script automates the descriptor override generation steps highlighted in the wiki **Hotfixes with Descriptor Overrides**.

PR: https://github.rsa.lab.emc.com/asoc/ng-core/pull/2933

(Option 1) Instructions to install nonfips rpms during an upgrade

Step 1: Upgrade Node-Zero

- Follow upgrade instructions for node-zero from UI or command line, this would install the latest repo under `/var/netwitness/common/repo/`

Step 2: Before upgrading other Node-X hosts with core services, run the script on Node-Zero.

- `/usr/local/sbin/nw-nonfips-override-generator.sh`
  - This would invoke script in default mode which uses latest installed version ex: 12.0.0.0 and repo /var/netwitness/common/repo/12.0.0.0/RSA .
  - The script generates and validates nonfips descriptor override files.

Step 3: Refresh Node-Zero to apply applicable nonfips rpms (broker, appliance) using the command:

- `nw-manage --refresh-host --host-key <node-zero-ip-address>`
- Reboot Node-Zero

Step 4: Upgrade each Node-X

- (Option 1) : Upgrade Node-X via Command line
  - Start upgrade host Node-X with the following command from Node-Zero.
    - Example: upgrade-cli-client –-upgrade --host-key <Host UUID or IP address> --version 12.0.0.0
  - Reboot the Node-X
- (Option 2): Upgrade Node-X from UI:
  - Continue upgrading from UI and follow instructions to upgrade the remaining hosts.
- If applicable, all the hosts would have the nonfips rpms (decoder, concentrator, broker, appliance) installed.

(Option 2) Instructions to install nonfips rpms after upgrade

After upgrading Node-Zero and all other Node-X hosts, and if you wish to install nonfips rpms at later then following are instructions:

- Step 1: Run the script on Node-Zero
  - `/usr/local/sbin/nw-nonfips-override-generator.sh`
    - This would invoke script in default mode which uses latest installed version ex: 12.0.0.0 and repo /var/netwitness/common/repo/12.0.0.0/RSA .
    - The script generates and validates nonfips descriptor override files.
- Step 2: Command line update:
  - Refresh all the hosts (including Node-Zero) to apply applicable nonfips rpms (decoder, concentrator, broker, appliance).
    - nw-manage --refresh-host --host-all
    - Reboot all the hosts.

Note: The `/usr/local/sbin/nw-nonfips-override-generator.sh` script would generate logs while running the commands, and also update `/var/log/messages` using the logger command.

TLS 1.3 Keys and Cipher suites

According to RFC 8446 the TLS 1.3 protocol doesn't support Private key based decryption and it only support a set of ephemeral symmetric keys also called as key secrets for the life of one specific TLS 1.3 session. The TLS 1.3 session also encrypts handshake, all the TLS records after Server Hello are encrypted , and at each stage a corresponding key is used for the encryption ex: client handshake, server handshake, client application and server application etc.. . So in order for decrypting a TLS 1.3 session all the required session keys `tls.client.hdshk, tls.server.hdshk, tls.client.app, tls.server.app, tls.client.early`(optional) are needed to be available.

Cipher Suites

TLS 1.3 protocol recommends to support only strong and advanced Cipher suites. The Network Decoder supports all the advertised cipher suites from RFC-8446, however on integration with Thirdparty key forwarder the supported cipher suites might be limited to list supported by Thirdparty server. Example only GCM ciphers are supported etc..

The following is list of cipher suited supported in Network Decoder

- TLS_AES_128_GCM_SHA256
- TLS_AES_256_GCM_SHA384
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_128_CCM_SHA256
- TLS_AES_128_CCM_8_SHA256

Steps to configure Network Decoder to decrypt TLS 1.3 traffic

- Step 1: Validate that the Network Decoder Captures Encrypted TLS 1.3 Traffic
- Step 2: Confirm HTTPS Parser is Enabled on Decoders
- Step 3: Integrate Network Decoder to receive keys from Thirdparty TLS 1.3 key forwarder
- Step 4: (Optional) Upload TLS 1.3 keys to validate Decryption
- Step 5: Validate TLS 1.3 keys are received on Decoder
- Step 6: View Decrypted Sessions

Step 1: Validate that the Network Decoder Captures Encrypted TLS 1.3 Traffic

First, make sure the Decoder is actually capturing the SSL/TLS traffic from the managed servers that you want to use to gain visibility into the encrypted sessions. You can do this by searching in the Investigate views (Navigate or Events) using the IP addresses of the servers as the destination IP addresses, and the service as SSL/TLS.

For example, if you are searching for HTTPS traffic to 192.168.1.2 and 192.168.1.3, you can run the following query:

```
service = 443 && ip.dst = 192.168.1.2, 192.168.1.3
```

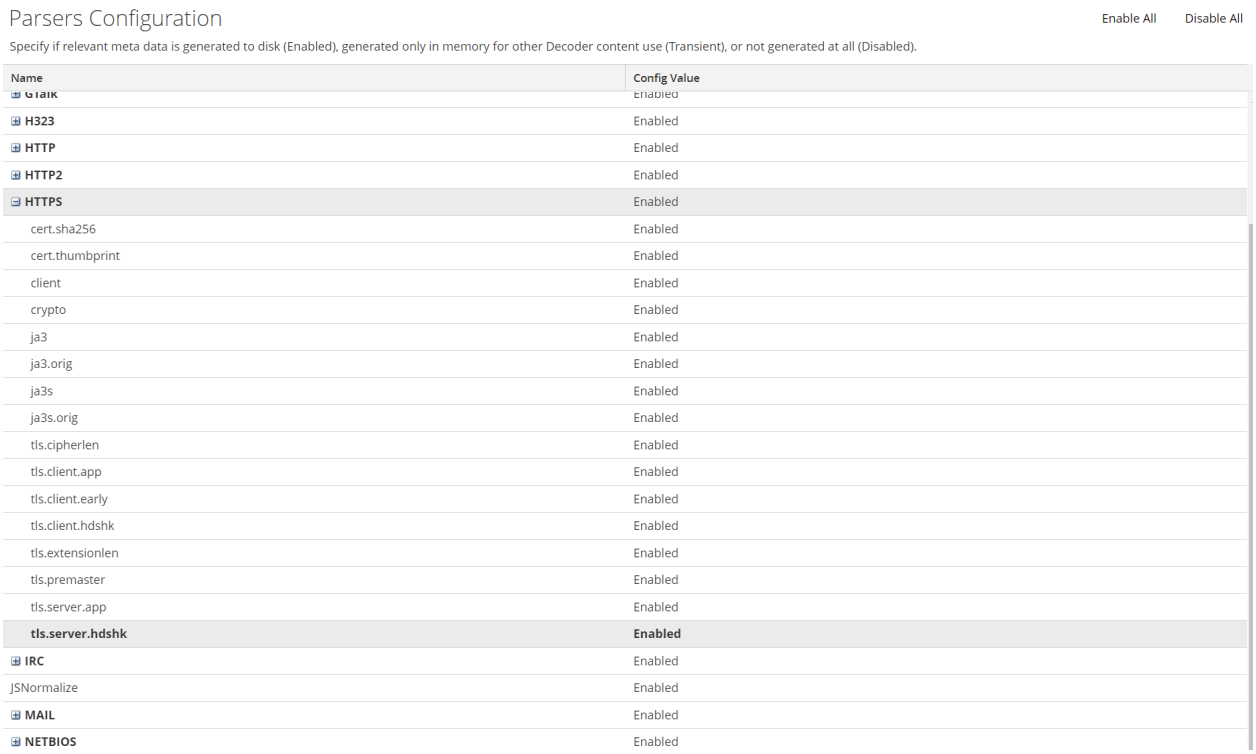
Conduct this search on the Concentrator that aggregates metadata from the Decoder that is expected to capture this network segment, or on a Broker that may have access to all collected metadata.

(Version 11.6 or later) In case of simultaneous ingestion of the encrypted and decrypted traffic streams in the Decoder, ensure that multi-thread assembly is enabled. This is to prevent mixing of packets from different sessions by single thread assembly. For more information, see the following topics:

- Perform Simultaneous Ingestion of the Encrypted and Decrypted Traffic Streams to Decoder
- (Optional) Multiple Adapter Packet Capture

Step 2: Confirm HTTPS Parser is Enabled on Decoders

The decryption process requires the native HTTPS parser to be enabled to examine the TLS handshake and to compare the uploaded encryption keys to determine if they decrypt the traffic. This can be validated in **Admin > Services > Decoder > Config** > Parsers Configuration panel as shown in the following figure.



If TLS_Lua parser is enabled then meta version: TLS 1.3 is registered for TLS 1.3 encrypted traffic.

Step 3: Integrate Network Decoder to receive keys from Thirdparty TLS 1.3 key forwarder

TLS 1.3 protocol only support ephemeral keys which are unique per session and Network Decoder need to receive those keys in time for parsing step.

Network Decoder provides `sslKeys` API which can receive the keys forwarded or injected to Decoder on RESTful interface.

A Thirdparty TLS 1.3 key forwarder which acts as Man-In-The-Middle (MITM) gateway for TLS 1.3 may intercept SSL handshake and forward the TLS 1.3 keys to Network Decoder.

Example: F5 BIG-IP Local Traffic Manager.

Option 1: Configure F5 BIG-IP Local Traffic Manager (LTM)

If you are using BIG-IP LTM in your environment for application traffic management, load balancing etc.. and if you would like to decrypt and analyze TLS 1.3 traffic handled by BIG-IP LTM, then refer the following section to Configure F5 BIG-IP LTM and forward keys to Network Decoder.

Refer F5 BIG IP - Netwitness Perfect Forward Secrecy Inspection Visibility for more information.

Option 2: Using sslKeys API

If you are using other Man-In-The-Middle (MITM) gateway which can intercept TLS 1.3 handshake and forward the keys, then you can refer the following information and parameters to send keys to `sslKeys` API on Network Decoder.

The `sslKeys` API is used to push SSL crypto information to enable SSL decryption of a session's packets prior to parsing.

Parameters for Managing Keys

The `sslKeys` command has several parameters for managing premaster and private keys. This is the full list of parameters:

Parameter	Description
clear	Removes all premaster and TLS 1.3 keys from memory. Will not delete any PEM files installed on the system.
maxKeys	Changes the maximum number of premaster and TLS 1.3 keys that will be stored in memory.
counts	Returns key counts for sslKeys. Cannot be used with any other parameters.
listPems	Returns a list of all installed private key PEM files.
deletePem	Deletes the named PEM file from the file system. You can pass this parameter more than once to remove multiple files.
random	The random hash used to identify the premaster key or TLS 1.3 keys.
premaster	The premaster key that will be installed for the previous random parameter. They must show up in pairs and random must be first.
CETS	The TLS 1.3 client early traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
CHTS	The TLS 1.3 client handshake traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
SHTS	The TLS 1.3 server handshake traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
CTS0	The TLS 1.3 client traffic secret 0 key (application) that will be installed for the previous random parameter. It must show up with random parameter.
STS0	The TLS 1.3 server traffic secret 0 key (application) that will be installed for the previous random parameter. It must show up with random parameter.

Usage:

The Client random and TLS 1.3 keys can be passed as parameters to `sslKeys`, ex:

- Using Native port 56004:
  - `/decoder sslKeys random=`<value> `CETS=`<value> `CHTS=`<value> `SHTS=`<value> `CTS0=`<value> `STS0=`<value>
- Using RESTful port 50104:
  - `/decoder?msg=sslKeys&random=`<value>`&CETS=`<value>`&CHTS=`<value>`&SHTS=`<value>`&CTS0=`<value>`&STS0=`<value>
- Replace <value> with base64 encoded TLS 1.3 key value in the above cases.

Note: For most of the sessions the CETS (CLIENT_EARLY_TRAFFIC_SECRET) key may not be generated or used by TLS clients and it is normal. However, if the other keys (handshake and application keys) are not generated and not available then TLS 1.3 decryption will NOT be successful.

Return Values

Most commands return name/value pairs of statistics about the premaster and TLS 1.3 keys in memory. The statistics are listed in the following table.

Parameter	Description
added	The number of premaster and TLS 1.3 keys just added during this command.
total	The total number of premaster and TLS 1.3 keys loaded in memory.
agedOut	The total number of premaster and TLS 1.3 keys that were removed during this command (this is not a lifetime stat).
maxKeys	The current maximum allowed premaster and TLS 1.3 keys.
premaster	The total number of premaster keys loaded in memory.
tls13	The total number of TLS 1.3 keys loaded in memory.

Step 4: (Optional) Upload TLS 1.3 keys to validate Decryption

Note: Uploading TLS 1.3 keys is not supported for real-time capture, but can be used to validate that decryption is functioning properly on the Decoder.

The `sslKeys` command provides a way to upload TLS 1.3 keys to the Decoder, so that captured encrypted packets that match the keys can be decrypted before parsing. Administrators configure the Decoder by entering the `sslKeys` command using the NwConsole command line interface or the Decoder RESTful interface. Now, all TLS handshakes that use the TLS 1.3 keys and will be able to be decrypted by the Decoder.

For information about how to use the NwConsole command line interface, see the *NwConsole User Guide for RSA NetWitness Platform*. For information about how to use the RESTful interface, see the *RESTful API User Guide for RSA NetWitness Platform*. Go to the *Master Table of Contents* to find all RSA NetWitness Platform 11.x documents.

You can use the RESTful interface form to facilitate uploading of multiple keys, both premaster and private at the same time.

1. Open the RESTful API in your browser and navigate to this path on the Decoder that you want to configure:

   `/decoder/sslkeys`.

2. Click **Choose File** Next to **Upload File 1**, and locate the TLS 1.3 keys file that you want to upload on your local file system.
3. Click **Upload**. The files are uploaded to the Decoder and results are displayed in the form.

TLS 1.3 Keys

Similar to `premaster` key, capturing TLS 1.3 keys and uploading to Decoder in time for parsing is not an easy way. However, well known browsers like Chrome , Firefox and IE can write the TLS 1.3 keys they generate to a file and this is useful for testing purposes. To configure your browser to do this, create an environment variable called SSLKEYLOGFILE and set with a pathname to text file to write the keys. Decoder will accept the file exactly as it is written and will use all the decryption keys in the file for any encrypted traffic it captures.

The following is a sample NwConsole script that uploads the file to a Decoder:

```
login <decoder>:50004 <username> <password>
send /decoder sslKeys --file-data=SSLKeys.txt
```

Optionally the parameters can be passed in as list in the upload file, ex:

```
send /decoder sslKeys --file-data=SSLKeys.txt --file-format=params-list
```

or you could use the following curl command (with the RESTful port):

```
curl -u "<username>:<password>" -H "Content-Type: application/octet-stream" --
data-binary @"/path/SSLKeys.txt" -X POST
"http://<hostname>:50104/decoder?msg=sslKeys"
```

Once the symmetric keys are uploaded, they will immediately be used for any necessary decryption. Symmetric keys are stored in memory and there is a limit to how many can be stored at any point in time. As more are added, the earliest keys will be aged out. You can also add available TLS 1.3 keys by just passing the `random`, `CETS` , `CHTS` , `SHTS` , `CTS0` and `STS0` parameters to `sslKeys`.

Note: For most of the sessions the CLIENT_EARLY_TRAFFIC_SECRET key may not be generated or used by TLS clients and it is normal. However, if handshake and application keys are not generated and not available then TLS 1.3 decryption wouldn't be successful.

Step 5: Validate TLS 1.3 keys are received on Decoder

It is advised to validate if Decoder is receiving keys to Decrypt TLS 1.3 sessions. As Decryption wouldn’t be successful if the keys are not available.

The following are few ways:

1. Decoder generates the following TLS 1.3 meta per session

`tls.client.hdshk`, `tls.server.hdshk`, `tls.client.app`, `tls.server.app`, and `tls.client.early`

2. The `sslKeys` API provides an option to display current key stats using counts=true .

Example: `/decoder sslKeys counts=true`

```
premasterKeys: 478
tls13Keys: 341
total: 819
maxKeys: 1000000
```

Step 6: View Decrypted Sessions

All storage of the packets is still encrypted, so when you search for any decrypted traffic, your query should still focus on `service=443` or HTTPS (HTTP over SSL/TLS). If any packets were able to be decrypted, the `tls.client.hdshk`, `tls.server.hdshk`, `tls.client.app`, `tls.server.app`, and `tls.client.early` metadata is associated with the session that has the TLS 1.3 keys required for decrypting that session. The session will also have crypto metadata denoting the cipher suite.

Viewing Unencrypted Traffic

If packets are decrypted during the parse stage, encrypted packets are written to disk, and the matching TLS 1.3 keys used for decrypting is written to the TLS 1.3 meta, analysts can view the unencrypted packets using the TLS 1.3 meta keys.

Option 1 : Using Event Analysis UI

Starting 12.0 the Event Analysis View support Displaying Decrypted and Encrypted Payloads of TLS sessions when TLS keys meta are registered.

Option 2 : Using RESTful API on Decoder service

The Decoder API that you can use to see the unencrypted packets is the `/sdk/content` RESTful service. You need to know the Session ID of the encrypted packets and the `flags` parameter masked to the value 128 (or 0x80 in hex). Point your browser to the Decoder RESTful interface and type in the following command, substituting the actual Session ID for `<id>`:

`http://<decoder>:50104/sdk/content?session=<id>&flags=128&render=text`

The Decoder returns a simple web page showing the packets after they are decrypted.

If you want to see what the packets look like encrypted, type in one of the following commands, substituting the Session ID for <id>:

`http://<decoder>:50104/sdk/content&session=<id>&render=text`

`http://<decoder>:50104/sdk/content&session=<id>&flags&render=text`

For more information on the `/sdk/content` service, see the manual page for `/sdk content`.

If you are not able to view the encrypted sessions, check Troubleshooting. However, if you are able to see the TLS 1.3 metadata, that confirms that the Decoder is able to see inside the encrypted session, which enables parsers to see the unencrypted packet payload and create metadata accordingly. For these decrypted sessions, look for additional metadata around usernames, passwords, or HTTP body content that would be obfuscated when encrypted, which would result in metadata not being generated.

Troubleshooting

If you are still not able to decrypt or view the decrypted payloads after attempting the steps above, try these steps.

1. Validate if Decoder is receiving TLS 1.3 keys. Refer #Step 5 listed above.
2. Validate that your TLS 1.3 session key is using an acceptable cipher-suite.
3. Make sure your Network Decoder is capturing the traffic matching the server for which you have the TLS 1.3 keys. To check, you can upload the TLS 1.3 keys and PCAP into Wireshark and see if it is able to decrypt it.

Additional Features

The following features which are supported in TLS 1.2 are also supported in TLS 1.3 decryption.

1. TLS Certificate Hashing
2. JA3 and JA3S TLS Fingerprints
3. Performance Considerations

For more details refer Decrypt Incoming Packets (TLS 1.2)

Edit Decoder System Configuration

When a service is first added to NetWitness, default values for the system configuration parameters are in effect. In most cases, the default values for compression, statistics update interval, and number of threads in the thread pool are set at a good point for optimal system performance. You do not need to edit these setting unless an RSA Customer Support technician advises you to change them.

System Configuration

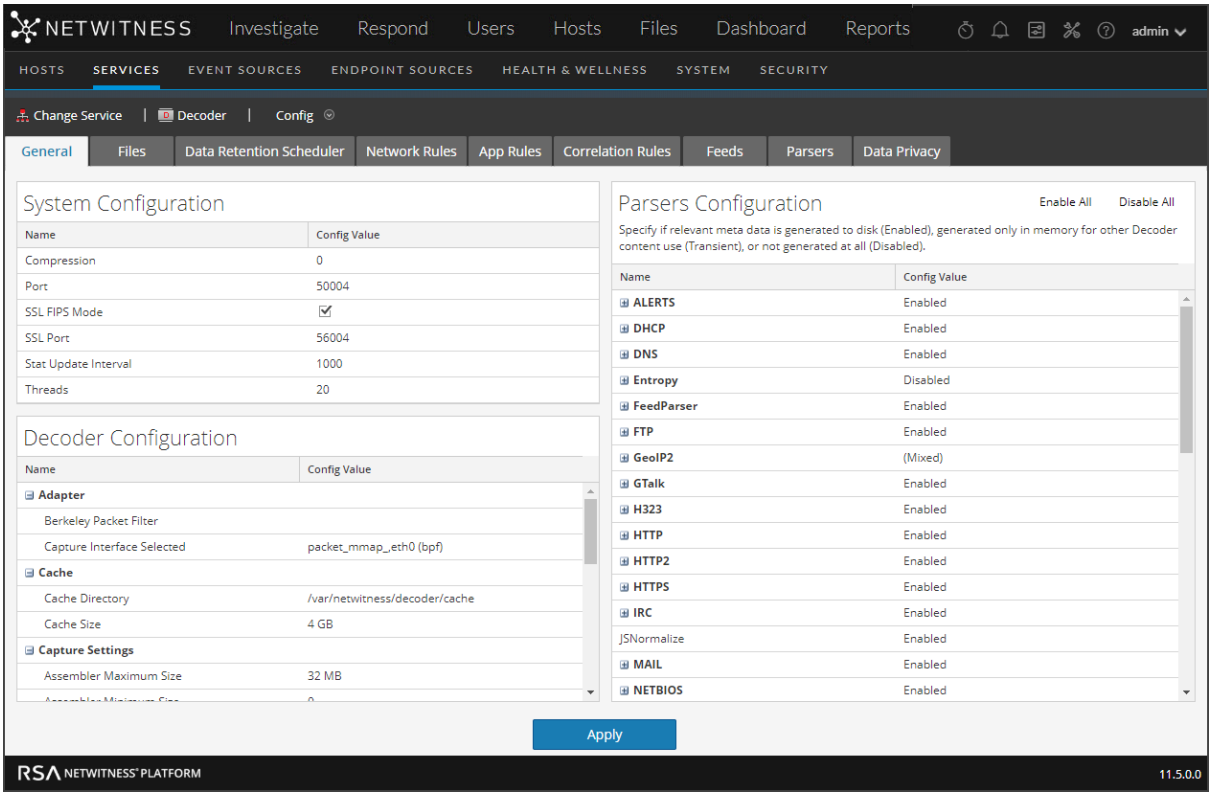
Name	Config Value
Compression	0
Port	50004
SSL FIPS Mode	☐
SSL Port	56004
Stat Update Interval	1000
Threads	20

One parameter that you may want to change for your environment is the SSL setting, which by default is not enabled. When enabled, the security of data transmission is managed by encrypting information and providing authentication with SSL certificates.

To edit system configuration parameters for a Decoder or Log Decoder:

1. Go to  **(Admin) > Services**.

2. Select a Decoder or Log Decoder service, and select  **> View > Config**.
   The Services Config view for the service is displayed with the General tab open.



3. Under **System Configuration**, click in a field that you want to edit **(Compression**, **Port**, **SSL FIPS Mode**, **SSL Port**, **Stat Update Intervals**, or **Threads**). Type a new value.
4. When finished editing, click **Apply.**
   The settings become effective immediately.

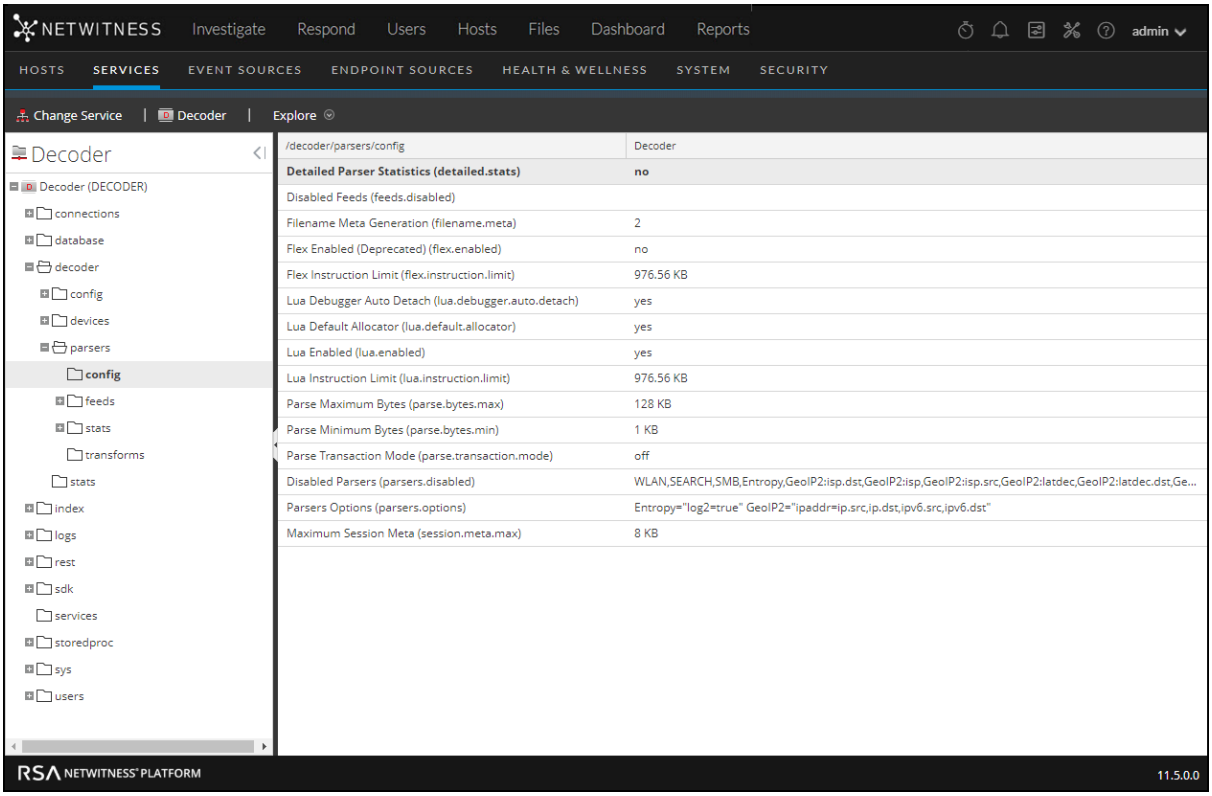
Enable CPU Usage Statistics for Installed Content

The Decoder provides CPU utilization statistics for all installed content, which you can use to reveal how much CPU time is used by parsers, feeds, application rules, and lexical scanning. The statistics are visible as Stat nodes in the service tree from the Explore view when `/decoder/parsers/config/detailed.stats` is enabled and the Decoder is capturing the stats.

Each piece of content is accounted as a single percentage value (0-100) regardless of the number of parse threads running. The percentage represents an average of the CPU utilization for the content across all threads.

To enable usage statistics monitoring:

1. In the NetWitness User Interface, go to **(Admin)** > **Services**.
2. Select a Decoder service and click > **View** > **Explore**.
   The Explore view for the Decoder is displayed.
3. Open `/decoder/parsers/config/` and in the right pane, select the `detailed.stats` parameter.



4. Change the value to **yes**. If the Decoder is not capturing data, start capture.
   When you open the Decoder Stats node in the Explore view, the new statistic is visible.

Enable Parser Mappings

This topic tells administrators how to enable event source mapping on a Log Decoder.

The Log Collector discovers the event source type on a per-message basis. If the correct parser is not identified for the event source, a small percentage of logs may be misidentified. The misclassified messages do not populate event source rules and alerts, and the reports do not have the correct data. If there are multiple event source types associated with an IP address, it makes it difficult for the parsers to identify the exact event source from which the logs are generated.

If you map an IP address to its event source type, the Log Decoder can identify the event source from which the log is generated. When messages are delivered to the Log Decoder from a mapped event source, only the assigned parsers are queried to find event matches.

You can assign event source types to IPV4, IPV6, or the hostname value of the event source. You can also assign multiple event source types to a single IP address. You can also use the Log Collector ID when different event source types with the same IP address are sent to different Log Collectors.

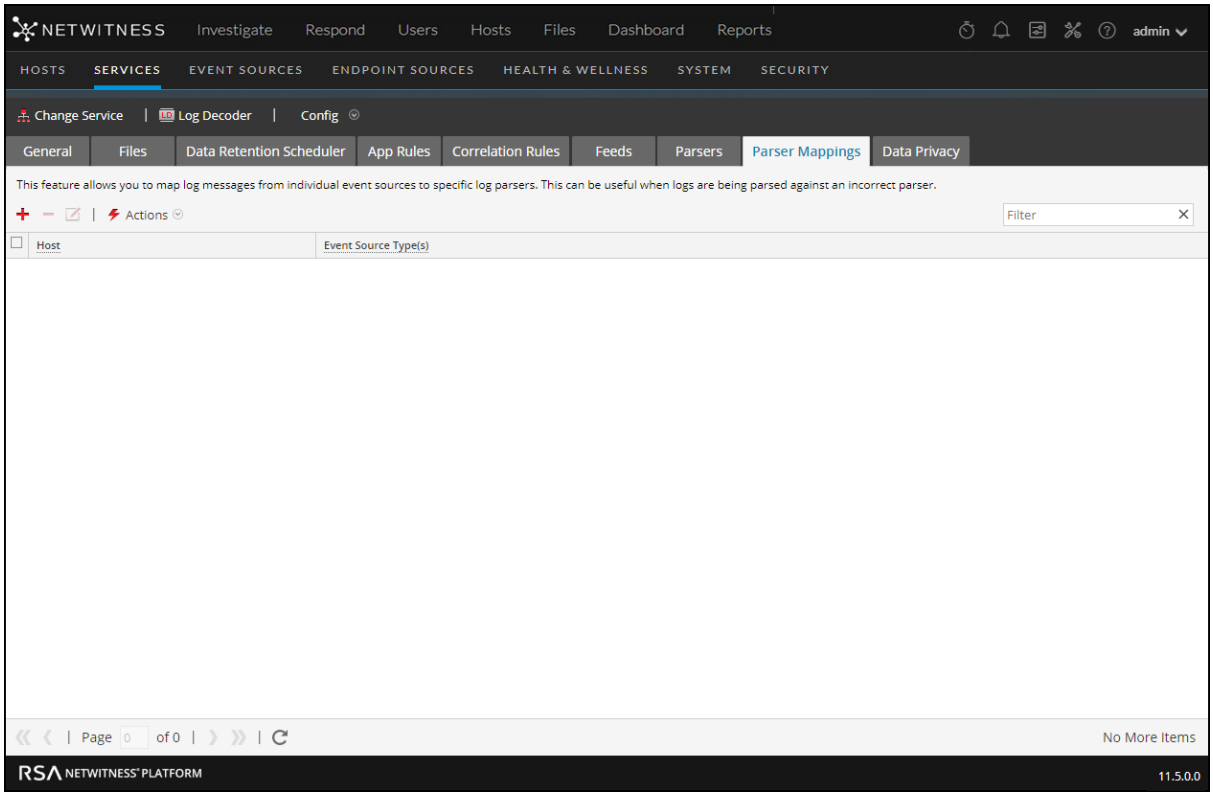
Note: You can also enable parser mapping functions by navigating to  (Admin) **> Event Sources > Discovery.**

Enable IP Address to Event Source Mapping

To enable an IP address to event source mapping:

1. Go to  **(Admin) > Services** and select a Log Decoder.
2. Select  **> View > Config**.
3. In the Configuration page, select the **Parser Mappings** tab.

The Parser Mappings tab is displayed in the Services Config view.



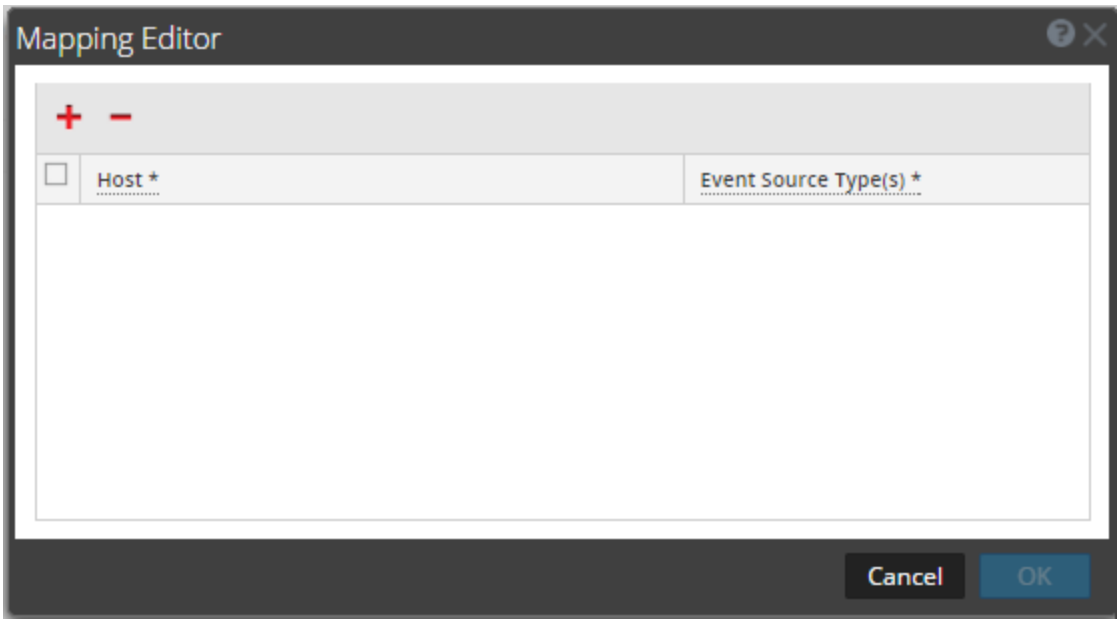
Update IP to Event Source Mapping

To update an IP to event source mapping:

1. Go to (Admin) > **Services**.
2. Select a **Log Decoder**, and in the **Actions** column, select > **View** > **Config**.

   The Services Config view is displayed.
3. Select the **Parsers Mapping** tab.
4. Click .

   The Mapping Editor is displayed.



5. Any of the following mappings can be defined:
   - **One Host and One Event Source Type**

     In the **Host** field, enter the hostname.

     For example: `10.0.0.1`
   - In the **Event Sources(**s) field, enter the event source type.

     For example: `apache`
   - **One Host and One or More Event Source Types**

     In the **Host** field, enter the hostname.

     For example: `10.0.0.1`
   - In the **Event Source(s)** field, enter the event source type.

     For example: `apache,sap,aix`
   - **One Host, One Log Collector, and One Event Source Type**

     In the **Host** field, enter the hostname and Log Collector ID.

     For example: `10.0.0.1,LC-1`
   - In the **Event Source(s)** field, enter the event source type.

     For example: `apache`
   - **One Host, One Log Collector ID, and One or More Event Source Types**

     In the **Host** field, enter the hostname and Log Collector ID.

     For example: `10.0.0.1,LC-1`
   - In the **Event Source(s)** field, enter the event source type.

     For example: `apache,sap,aix`

Note: The event source types are processed in the order you enter the parsers and if one or more parsers matches a log, the first parser in the list is queried. The Host/IP can be IPv4, IPv6, or Hostname.

6. Click **OK**.

   The Parser Mapping is added.

7. To cancel the parser mappings selection, click **Cancel**.

Read IP to Event Source Type Mappings

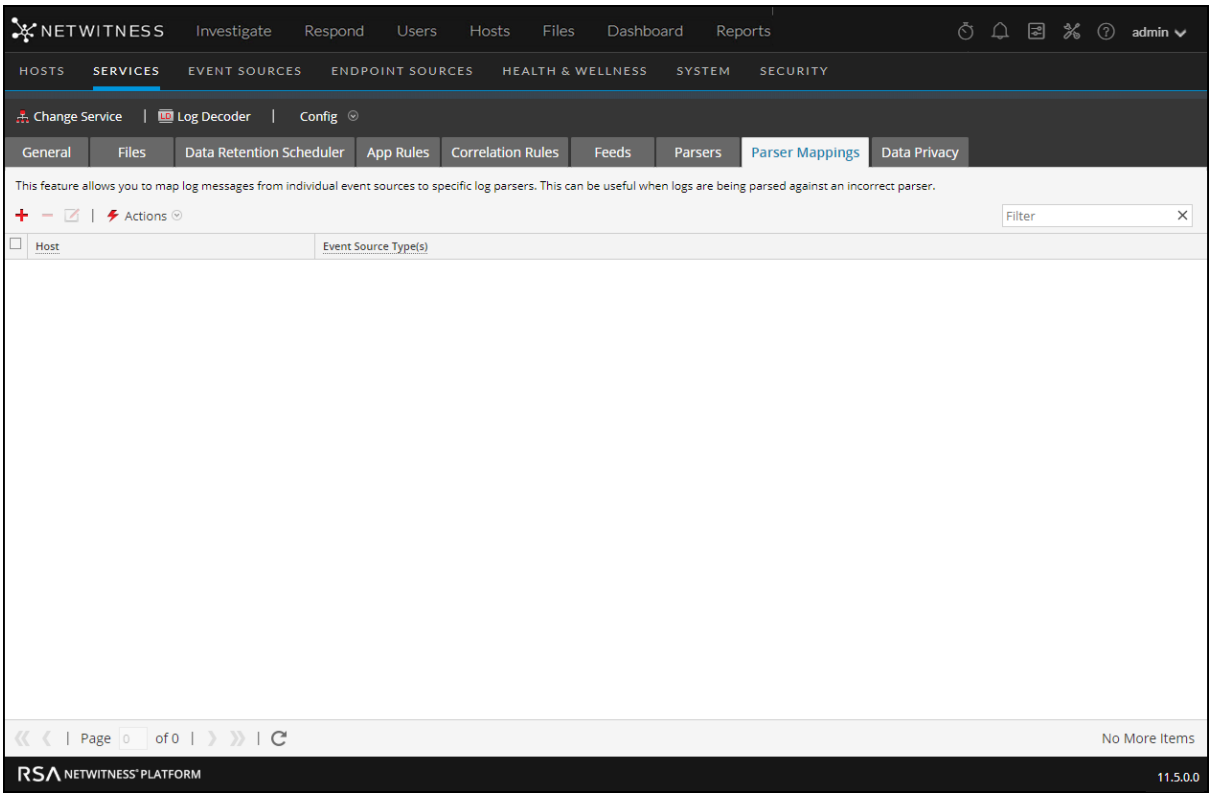
To read an IP to event source type mappings:

1. Go to **(Admin) > Services**, and select a Log Decoder service.
2. In the Actions column, select **> View > Config**.

   The Services Config view is displayed.

3. Select the **Parsers Mapping** tab.

   The mappings are displayed.



Edit IP to Event Source Type Mappings

To edit IP to event source type mappings:

1. Go to **(Admin) > Services**, and select a Log Decoder service.
2. In the Actions column, select **> View > Config**.

   The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. Select the mapping you want to edit.

   Note: You can only edit one mapping at a time.
5. Click .
6. In the **Event Source(s)** field, modify the event source(s).

   Note: The host is not editable and the field is disabled.
7. Click **OK** to accept the edited Event Source.
8. To cancel the changes, click **Cancel**.

Delete IP to Event Source Type Mappings

To delete IP to event source type mappings:

1. Go to **(Admin) > Services**, and select a Log Decoder service.
2. In the Actions column, select **> View > Config**.

   The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. Select the mapping you want to delete.
5. Click .

   The mapping is deleted and the grid is refreshed.
6. To cancel the changes, click **Cancel**.

Sort the Hostname or Event Source Type

To sort the hostname or event source type:

1. Go to **(Admin) > Services**, and select a Log Decoder service.
2. In the Actions column, select **> View > Config**.

   The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. To sort a column, click in the column header.

Event Source Types are applied for your selected IP address. Logs are parsed against the parsers in the order they are listed.

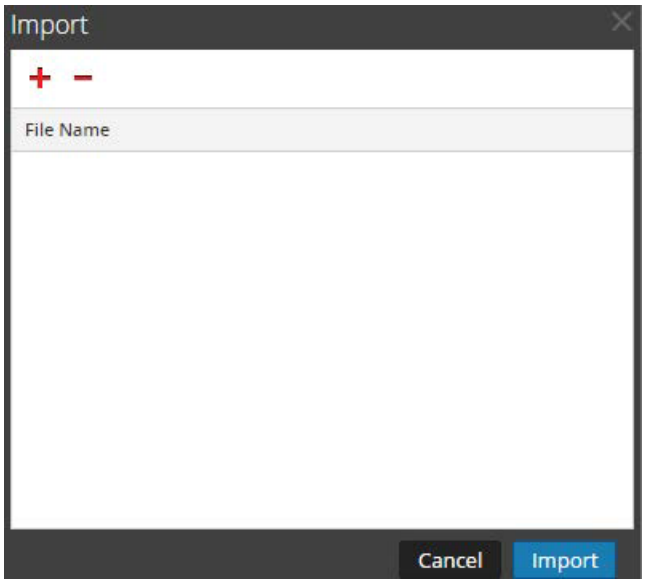
Import IP to Event Source Mapping Entries

To import IP to event source mapping entries:

1. Go to **(Admin) > Services**, and select a Log Decoder service.
2. In the Actions column, select **> View > Config**.

   The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. Select **Actions > Import**.

   The Import dialog is displayed.



5. Click .
6. Select the file you want to import and click **OK**.
7. To load the parser, click **Import**.

Note: You can only import one .csv file at a time.

Export IP to Event Source Mapping Entries

To export IP to event source mapping entries:

1. Go to **(Admin) > Services**, and select a Log Decoder service.
2. In the Actions column, select **> View > Config**.

   The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. Select the mappings you want to export.

5. Select **Actions > Export > Selection**.

   The Export Selection dialog is displayed.



6. Enter the file name and click **Export**.

Search IP to Event Source Mapping Entries

To search IP to event source mapping entries:

1. Go to **(Admin) > Services**, and select a Log Decoder service.
2. In the Actions column, select **> View > Config**.

   The Service Config view is displayed.
3. Select the **Parser Mappings** tab.
4. In the Parsers Mappings toolbar, enter the Host or Event Source in the **Filter** field.
5. Click **Enter**.

The Hosts or Event Sources that match the names entered in the **Filter** field are displayed.

Enable or Disable Lua and Flex Parsing Systems

This topic tells administrators how to enable or disable Lua and Flex parsing systems on a Decoder or Log Decoder. Flex parsers are deprecated and disabled by default.

The settings to enable or disable Lua and Flex parsing systems are configured correctly by default and you do not typically have to change them. However, you may need to adjust these settings at the request of RSA Customer Care or for troubleshooting purposes.

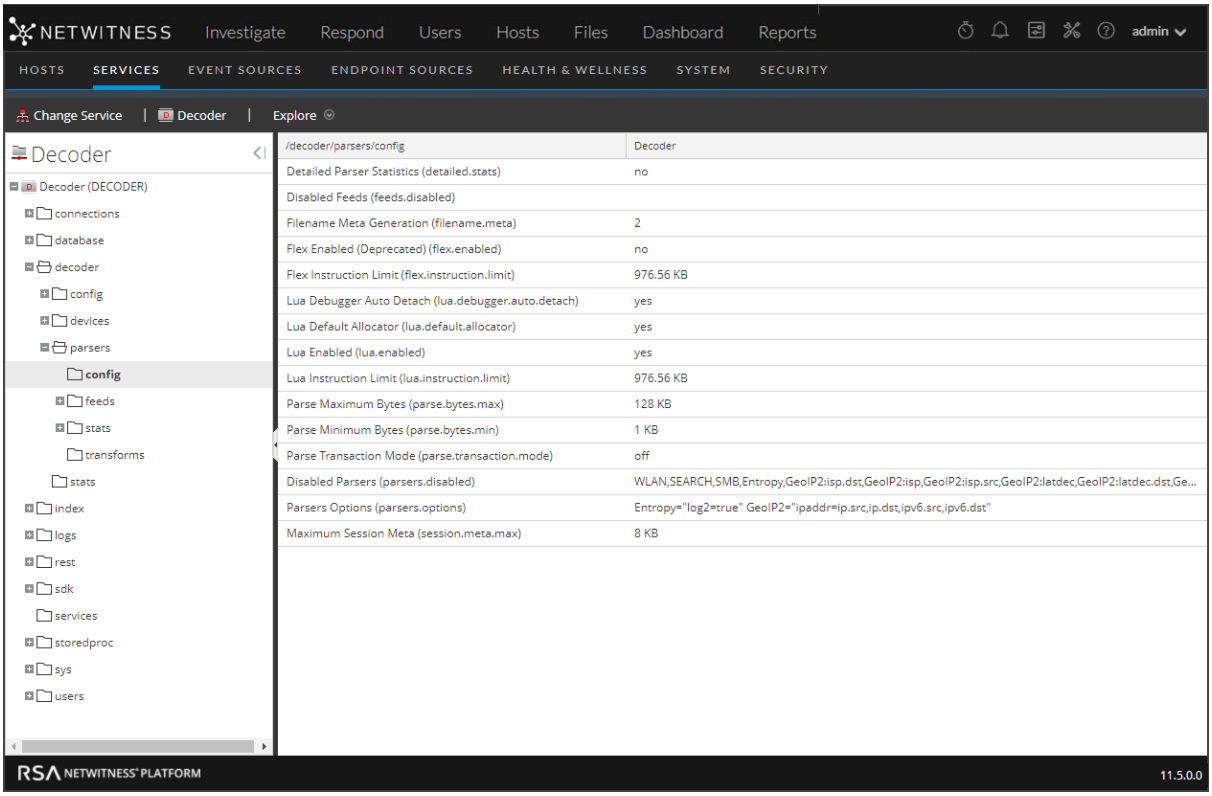
In addition to configuring individual parsers, you can enable and disable all Lua parsing as well as all Flex parsing in the Services Explore view. You enable and disable the Lua parsing and Flex parsing systems settings separately, but they work in the same way.

- If you **disable** the Lua or Flex parsing system, the corresponding parsing system is disabled and no parsers are loaded.
- If you **enable** the Lua or Flex parsing system, the corresponding parsing system is enabled and individual parsers are enabled and disabled following the current individual configurations.

To enable or disable Lua and Flex parsing systems on a Decoder or Log Decoder:

1. Go to  **(Admin) > Services.**
2. Select a Decoder or Log Decoder and  **> View > Explore.**
   The Services Explore view for the selected service is displayed.

3. In the Node list, navigate to and select **`/decoder/parsers/config`**.



4. In the values panel:
   - To enable the Lua parsing system, in the value field for `lua.enabled`, type **yes**.
   - To disable the Lua parsing system, in the value field for `lua.enabled`, type **no**.
   - To enable the Flex parsing system, in the value field for `flex.enabled`, type **yes**.
   - To disable the Flex parsing system, in the value field for `flex.enabled`, type **no**.

Map IP Address to Service Type for Log Parsing

This topic describes the procedure to map an IP address to a service type for log parsing.

The Log Collector discovers event source type on a per-message basis. If the correct parser is not used for the specific event source, the messages that are common between event source types are misclassified. The misidentified messages will not populate service rules and alerts, and the reports will not have proper information. Also, if there are multiple services associated with an IP address, it can be difficult for the parsers to identify the exact service from which the log is generated.

If you map an IP address to its services, the Log Decoder can identify the service from which the log is generated. When messages come into the log decoder from a mapped service, the assigned parsers are loaded to find event matches.

You can assign service types to IPV4, IPV6 or hostname value of the event source. You can also assign multiple service types to a single IP address. You can also use the CollectorID when different service types with the same IP address are sent to different collectors.

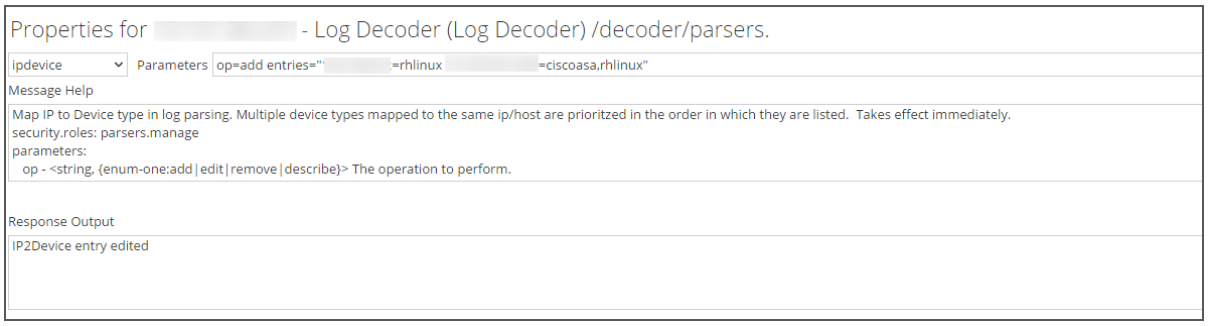
Map an IP Address to a Service Type

To map an IP address to a service type, do the following:

1. Go to  **(Admin) > Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select  **> View > Explore**.
3. Go to **/decoder/parsers** node, right-click **parsers**, and select **Properties**.
4. In the **Properties** view, specify the **ipdevice** command with the following parameters:

   `op=add/remove entries="`***`ipaddress`***`=`***`service`***`"`

   for example, `op=add entries="10.100.201.300=ciscoasa"`
5. Click **Send**.



IPdevice Command

In the `ipdevice` command, three operations are available:

- `add`: This operation adds or updates entries in the ipdevice map. Multiple space delimited address/type pairs may be specified.

  `op=add entries="`***`address`***`=`***`service type`***`"`

- `remove`: This operation removes entries from the ipdevice map. Multiple space delimited address/type pairs may be specified.

  `op=remove entries="`***`address`***`"`

- `describe`: This operation returns the values currently in the ipdevice map.

Map an IP Address to a Time Zone

This topic discusses NetWitness support for Event Time.

Often times logs do not fully specify timestamps and may be missing time zone information. To properly normalize such timestamps to UTC, the Log Decoder provides the ability to associate devices from a specific address (IPv4 or IPv6) or hostname to a time zone or a fixed offset.

Timestamp consideration for log files:

- Some logs are in UTC format.
- Some logs that are not in UTC include a timezone offset value to adjust the time accordingly.
- For logs that use local time, with no offset provided, you can create a source mapping to manually adjust times for logs from that event source.

Three time zone formats are currently accepted and are shown in the following examples:

- Olson format: `America/Anguilla`
- POSIX format: `AST2:45ADT0:45,M4.1.6/1:45,M10.5.6/2:45`
- Offset by Minutes format: `= -500`

NetWitness maps the device address (IPv4 or IPv6) or hostname to a specific time zone or offset. Event time meta that is parsed from a log that is from a mapped address and does not include an offset or time zone as part of the timestamp is adjusted to UTC according to the mapping.

To map an IP address to a time zone, do the following:

1. Go to  **(Admin) > Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select **> View > Explore**.
3. Go to **/decoder/parsers** node, right click **Parsers**, and select **Properties**.
4. In the **Properties** view, specify the `iptmzone` command with the following parameters:

   `op=add entries="`***`ipaddress`***`=`***`timezone`***`"`

   for example: `op=add entries="10.10.10.10=Africa/Addis Ababa"`
5. Click **Send**.

tzinfo Command

You can view the strings that can be used to specify the time zone by using the `tzinfo` command.

To view the time zone strings, do the following:

1. Go to  **(Admin) > Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select **> View > Explore**.
3. Go to **/decoder/parsers** node, right click **Parsers**, and select **Properties**.

4. In the **Properties** view, specify the `tzinfo` command with the following parameters:

   `op=tznames`

5. Click **Send**.

The list of time zone strings is returned in the **Response Output** text box.

iptmzone Command

In the `iptmzone` command, three operations are available:

- `add`: This operation adds or updates entries in the iptmzone map. Multiple space delimited address/type pairs may be specified.

  `op=add entries="address=time zone"`

- `remove`: This operation removes entries in the iptmzone map. Multiple space delimited address/type pairs may be specified.

  `op=remove entries="address"`

- `describe`: This operation returns the values currently in the iptmzone map.

Examples

The following examples provide instances for mapping IP addresses to time zones:

- If you want to map two different entries with different IPV4 values and time zone, enter the following parameter in the **iptmzone** command and click **Send**

  `op=add entries="10.10.10.10=America/Anguilla 10.10.10.11=Pacific/Rarotonga"`

- If you want to remove an entry for a single IPV4 value and time zone, enter the following parameter in the **iptmzone** command and click **Send**.

  `op=remove entries="10.5.245.9"`

- If you want to create a single entry for an IPV6 value and time zone, enter the following parameter in the **iptmzone** command and click **Send**.

  `op=add entries="2001:DB8:85A3::8A2E:370:7334=America/Anguilla"`

- If you want to create a single entry to map an IPV4, IPV6, or hostname with the Minute Offset, Olson, or POSIX format, enter the following string in the **iptmzone** command and click **Send**.

  ```
  op=add entries="10.168.0.2=America/Anguilla
  2001:DB8:85A3::8A2E:370:7334=0500 nwappliance21=EST5EDT,M3.2.0/2,M11.1.0"
  ```

Change the Date format

The Log Decoder parsing engine uses the `event.time` meta key to keep track of the time an event occurs for incoming log messages. You can use the `mdformat` for the **iptmzone** command to change the date format. The Log Decoder currently has the ability to change the dates in the logs to have the following format:

```
mdy or dmy (month/day/year or day/month/year)
```

Notes:

- Event time metadata in a parser does not account for `dmy` and `mdy`.
- Currently, there is no functionality to detect this scenario and adjust parsing automatically.

To change the date format, do the following:

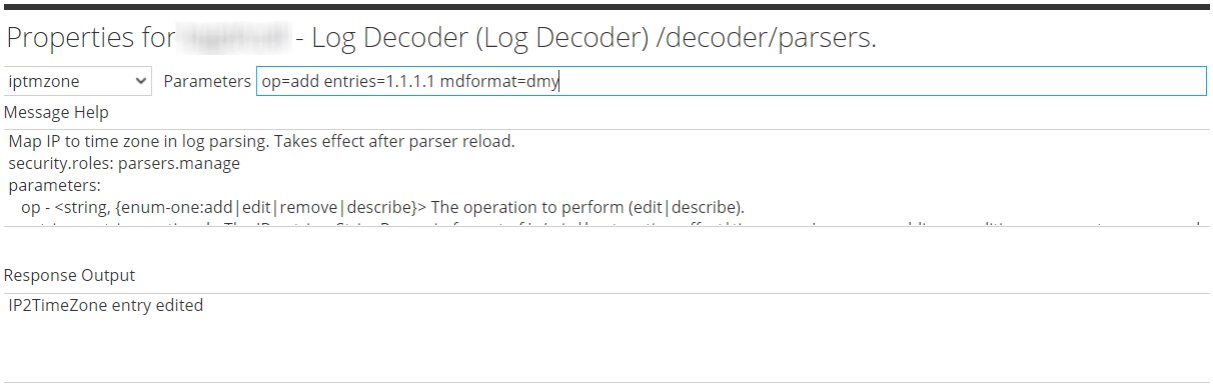
1. Go to **(Admin) > Services**.
2. In the **Services** view, select a Log Decoder, and in the **Actions** column, select; **> View > Explore**.
3. Go to **/decoder/parsers** node, right click **Parsers**, and select **Properties**.
4. In the **Properties** view, select **iptmzone** from the drop-down list, and specify the command with the following parameters:

   `op=add entries="<ipaddress>" mdformat=[dmy | mdy | none]`

   for example: `op=add entries="1.1.1.1" mdformat=dmy`

   Where:

   - `ipaddress` is the Device IP address, and
   - `mdformat` can be **dmy**, **mdy**, or **none**.

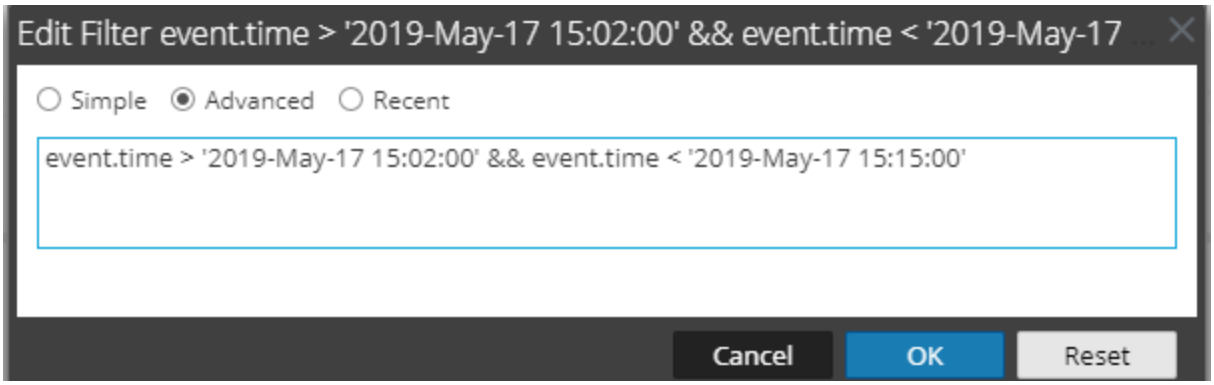


5. Click **Send**.

NetWitness maps the IP address along with the date format in the Log Decoder. Event time meta items are updated according to their respective mappings.

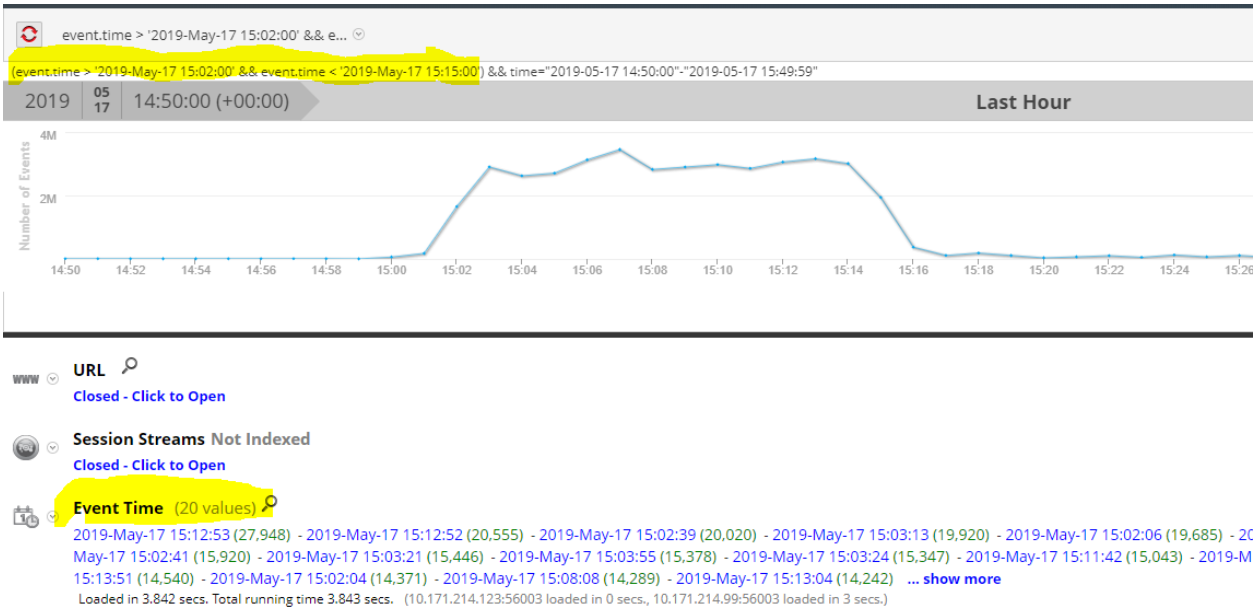
Event Time Filter Example

If you want to see logs that have an event time between 3:02 PM and 3:15 PM on May 17, 2019, create the following filter:



Note: The query builder may tell you "invalid expression," but that is a validation error in the User Interface. The query works because even though this is a text (string) search, the system actually compares the ASCII values. So, for example the expression **A < B** would be true.

Below is a sample of the event meta values returned from the query:



Missing Year Support

If there is no year associated with a format string, the parser attempts to populate the year heuristically. In most cases, the value is populated based on the current year. The current year is assigned as the message year (UTC), as per the clock on the Log Decoder.

The following are the various other scenarios:

1. Latent log during transition to new year.
2. Logs from forward time zones (or skewed clocks) just before new year transition.
3. Latent logs received where a leap day cannot be successfully assigned to an appropriate leap year.

Latent log during transition to new year

If the day is more than 31 days into the future, the year is decremented.

For example,

The message date is Dec-31. The current date is 2018-Jan-1. The temporary assigned date, 2018-Dec-31, will be more than 31 days into the future.

This will cause the year component to be decremented to 2017, resulting in a reasonable message time.

Logs from forward time zones (or skewed clocks) just before new year transition

If the day is more than 334 days into the past, the year is incremented.

For example,

The message date is Jan-1. The current date is 2017-Dec-31. The temporary assigned date, 2017-Jan-1, will be more than 334 days into the past.

This will cause the year component to be incremented to 2018, resulting in a reasonable message time.

Latent logs received where a leap day cannot be successfully assigned to an appropriate leap year

If the day (Feb 29) is invalid (say, the assigned year is NOT a leap year), the relative position to the current time cannot be calculated. To do this, the year is decremented in an attempt get a valid time stamp.

Because the position of the leap day relative to the transition to the new year, along with the expectation that no logs would be received this far into the future, we do not ever make an attempt to increment the year to produce a valid time stamp.

After the year is decremented, the same logic is then followed (refer to statement 1 and 2). If the result is still an invalid day. The parser throws and a valid message time cannot be parsed.

Limitations

Note: Currently, there is no mechanism to assign a year to the import of logs.

As this pertains to leap days, if you find the correct leap year, an inconsistent data would result. For example, if a set of messages is two years old and during a leap year, only a single day would remain which is accurate. Because of the heuristic described above, the remaining set of messages would have time stamps one year too.

For data consistency, usually you cannot find valid leap years beyond 1 month and less than 11 months.

Examples

The following examples provide instances for logs with year and logs without year:

- If the log is with year, event time is generated as below:

```
%emcavamar: 1704^^2017-04-07^^07:50:02^^1^^<event-source NodeID="avamar"
ProgramName="com.avamar"/>^^1270626^^SYSTEM^^ERROR^^OK^^MCS:DPN_
Proxy^^Internal server error^
<MESSAGE
     id1="1:01"
     id2="1"
     eventcategory="1605020000"
     functions="&lt;@msg:*PARMVAL($MSG)&gt;&lt;@event_time:*EVNTTIME
     ($MSG,'%W-%G-%F %H:%U:%O',fld2,fld3)&gt;"
     content="&lt;fld1&gt;^^&lt;fld2&gt;^^&lt;fld3&gt;^^&lt;fld4&gt;^^
     &lt;&lt;event-source NodeID=&quot;&lt;hostname&gt;&quot;
     ProgramName=&quot;&lt;fld8&gt;&quot;/&gt;^^&lt;fld9&gt;^^&lt;cate
     gory&gt;^^&lt;severity&gt;^^&lt;event_
     type&gt;^^&lt;agent&gt;^^&lt;event_description&gt;" />
```

- If the log is without year, event time is still generated, The current year is assigned as the message year (UTC), as per the clock on the Log Decoder.

```
%emcavamar: 1704^^04-07^^07:50:02^^1^^<event-source NodeID="avamar"
ProgramName="com.avamar"/>^^1270626^^SYSTEM^^ERROR^^OK^^MCS:DPN_
Proxy^^Internal server error^
<MESSAGE
     id1="1:01"
     id2="1"
     eventcategory="1605020000"
     functions="&lt;@msg:*PARMVAL($MSG)&gt;&lt;@event_time:*EVNTTIME
     ($MSG,'%G-%F %H:%U:%O',fld2,fld3)&gt;"
     content="&lt;fld1&gt;^^&lt;fld2&gt;^^&lt;fld3&gt;^^&lt;fld4&gt;^^
     &lt;&lt;event-source NodeID=&quot;&lt;hostname&gt;&quot;
     ProgramName=&quot;&lt;fld8&gt;&quot;/&gt;^^&lt;fld9&gt;^^&lt;cate
     gory&gt;^^&lt;severity&gt;^^&lt;event_
     type&gt;^^&lt;agent&gt;^^&lt;event_description&gt;" />
```

Obtain Log Files from a Pre-11.0 Log Decoder

NetWitness 11.0. added the capability to view a small sampling of recent logs for specific devices through detail tabs of the Discovery View. By default, Log Decoders prior to 11.0 do not have the necessary configuration to enable this feature, but a few minor changes can make it available.

To enable logs preview for a pre-11.0 Log Decoder, follow these steps on the Log Decoder:

1. Go to (Admin) > **Services** > select a **Log Decoder**, then select > **View** > **Config**.
2. Click the **Files** tab and select **index-logdecoder-custom.xml** from the drop-down menu.
3. Add the following three lines at the end of the file (before the closing language tag):

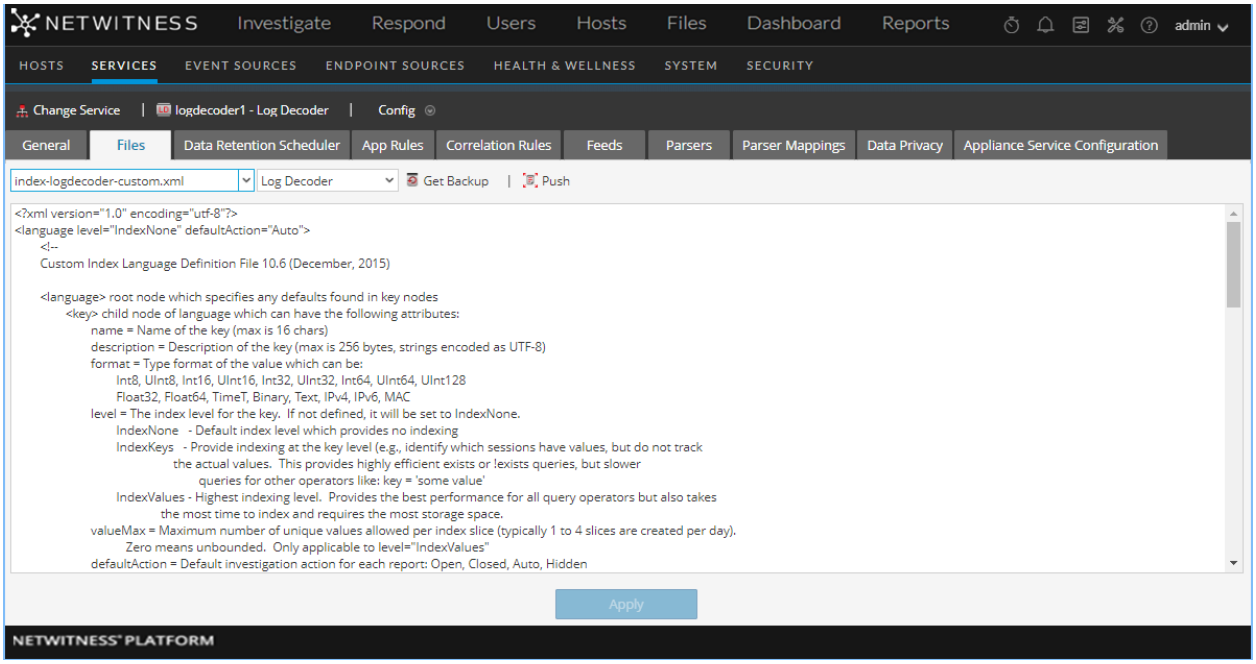
```
<key description="Device IP" level="IndexValues" name="device.ip" format="IPv4"
valueMax="100000" defaultAction="Open"/>
<key description="Device IPv6" level="IndexValues" name="device.ipv6" format="IPv6"
valueMax="100000" defaultAction="Open"/>
<key description="Device Host" level="IndexValues" name="device.host" format="Text"
valueMax="100000" defaultAction="Open"/>
```

4. Click **Apply**.
5. Restart the Log Decoder service as follows.

   Select Log Decoder service > **Explore** > **decoder** > **Properties** > **reset**. You select **reset** from a drop down menu.

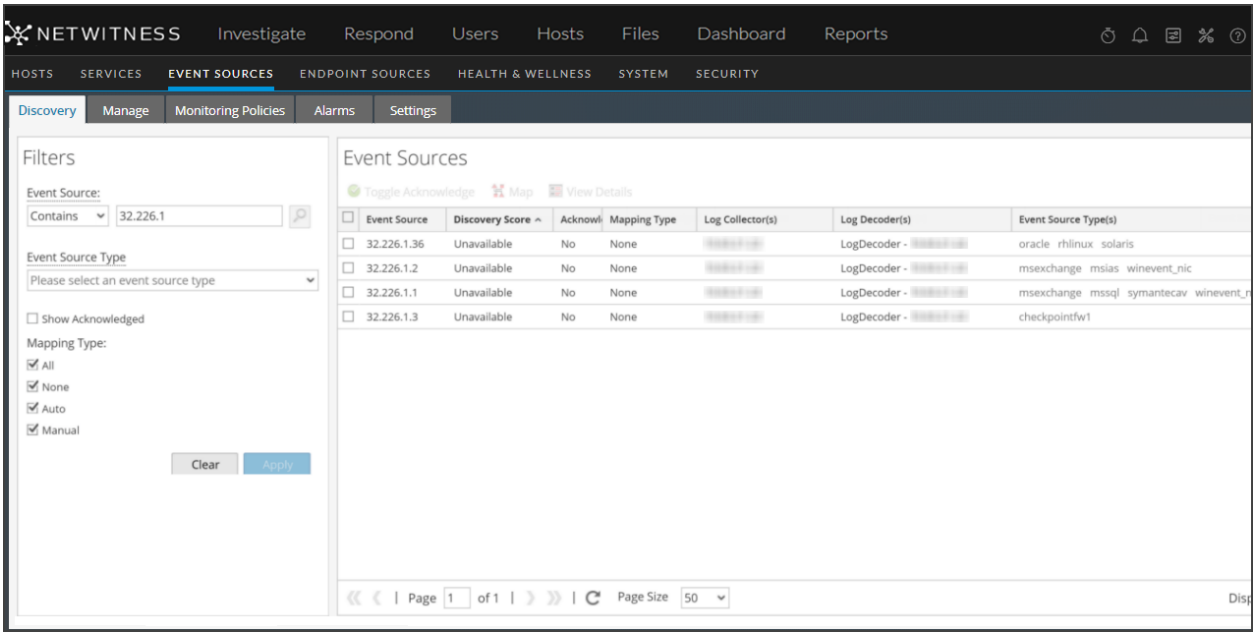
Click **Send** after you select reset.

This is an example of the **index-logdecoder-custom.xml** file.



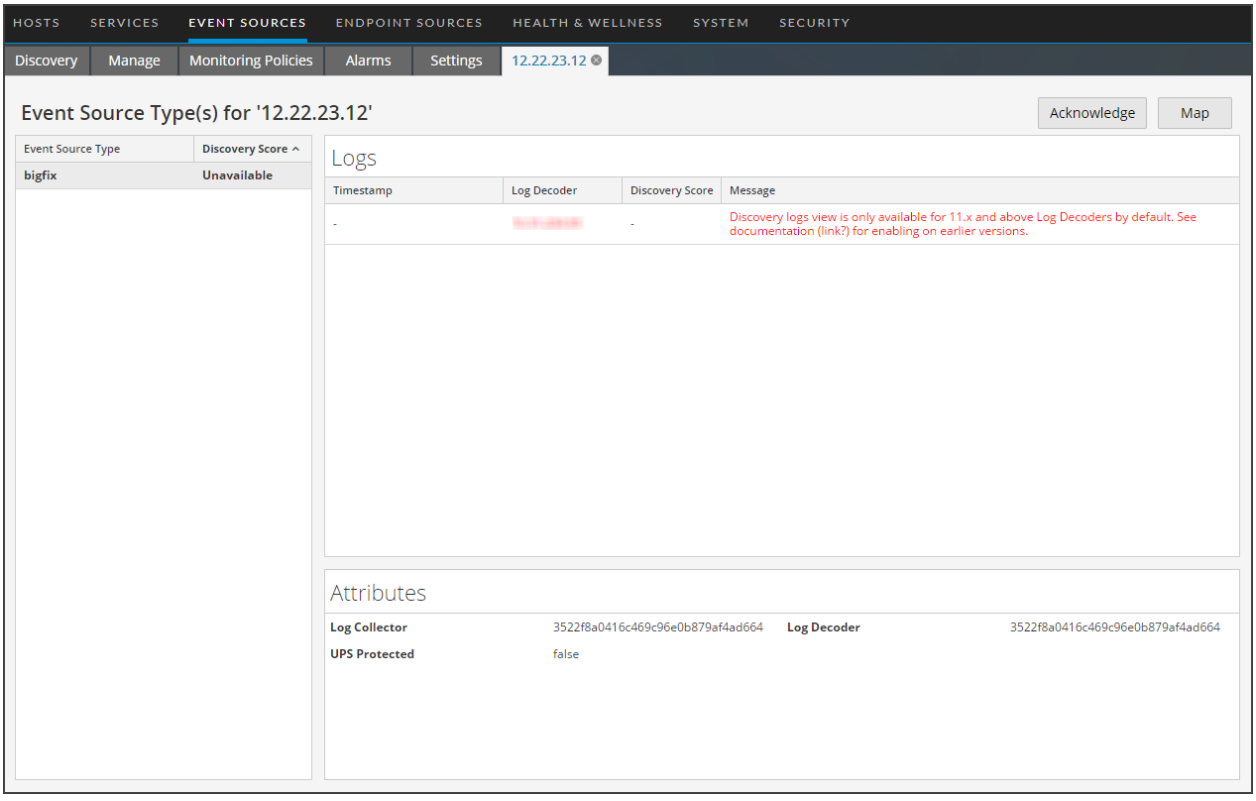
Note: Discovery Scores are only available for 11.x and above Log Decoders. Discovery Scores for pre-11.x Log Decoders are displayed as Unavailable.

The following example shows the Discovery Score as **Unavailable** in the **Details** view for a pre-11.0 Log Decoder.



Note: Device logs are only available for 11.x and above Log Decoders.

The following example shows the message that is displayed in the Logs panel for a pre-11.0 Log Decoder.



Upload a Log File to a Log Decoder

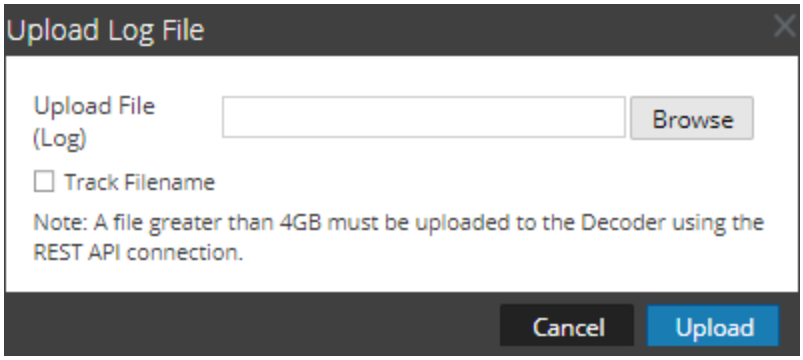
There are occasions when you want to analyze a log file that is not available on the service you are using. You can upload a log file captured on another service to NetWitness. Log filenames are of the type **.log**.

When a log file is uploaded to a Log Decoder, the Log Decoder analyzes and generates meta for each log it contains. These logs are added to the already decoded logs on the Log Decoder and are available for analysis. NetWitness includes a filename tracking option that makes searching for a particular set of logs easier. When the log file is uploaded with file tracking, the Log Decoder adds meta to each log based on the uploaded filename. You can then filter sessions for analysis using that meta.

The option to upload a log file is dimmed when other Log Decoder operations prevent an upload from occurring, for example, when the Log Decoder is capturing logs.

To import a log file to a Log Decoder:

1. Go to **(Admin)** **Services**.
2. Select a Log Decoder in the **Service** grid, and select **> View > System**.
   The Services System view for the Log Decoder is displayed.
3. In the toolbar, click **Upload Log File**.



4. To choose a log file, click **Browse**.
   A directory view is displayed.
5. Select the log file that you want to upload.
   The filename is displayed in the **Upload File** field.
6. If you want the Log Decoder to add meta to the logs based on the filename, click the checkbox next to **Track Filename**.
7. To upload the file, click **Upload**.
   The selected file is uploaded and a status message indicates that the file is uploaded. The log file is available for analysis.

Upload a Packet Capture File

There are occasions when you want to analyze a packet capture file that is not available on the service you are using. You can upload a file captured on another service to NetWitness. Supported packet capture file types are `pcap` and `pcap.gz`.

When a packet capture file is uploaded to a Decoder, the Decoder creates sessions from the packet capture file packets. These sessions are added to the already decoded sessions on the Decoder and are available for analysis. NetWitness includes a filename tracking option that makes searching for a particular set of sessions easier. When the packet capture file is uploaded with file tracking, the Decoder adds meta to the sessions based on the uploaded filename. You can then filter sessions for analysis using that meta.

The option to upload a packet capture file is dimmed when other Decoder operations prevent an upload from occurring; for example, when the Decoder is capturing packets.

To select and upload a packet capture file:

1. Go to **(Admin)** > **Services**.

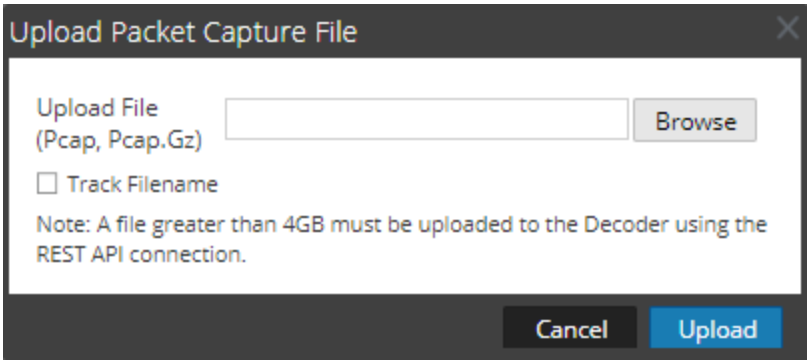
   The Administration Services view is displayed.

2. Select the Decoder name, and **> View > System**.

   The Services System view for the Decoder is displayed.

3. In the toolbar, click **Upload Packet Capture File**.

   The **Upload Packet Capture File dialog** is displayed.



4. To choose a capture file, click **Select**.

   A directory view is displayed.

5. Browse the directory and select the packet capture file that you want to upload.

   The filename is displayed in the **Upload File (pcap, pcap.gz)** field.

6. If you want the Decoder to add meta to the sessions based on the filename, click the checkbox next to **Track Filename**.

7. To upload the file, click **Upload**.

   A progress bar shows upload progress.

Upload time varies depending on the size of the file. When the file upload is complete, a status message is displayed. The file is now available for investigation.

Simultaneous Import and Capture

You can import or upload PCAPs while the capture is running using RESTful API. Make a note of the following points before you start importing during capture:

- The source file metadata is not created for imported PCAPs when the capture is running. If capture is stopped, the source file metadata is created as in older versions.
- The NWD files can be imported, but not while the capture is running. If you attempt to import an NWD file while the capture is running, the Decoder returns an error. An active import will block starting/stopping of import and capture.
- The file tracking meta is not created during the simultaneous import and capture.

F5 BIG IP - NetWitness Perfect Forward Secrecy Inspection Visibility

This topic discusses Perfect Forward Secrecy (PFS) Inspection Visibility with F5 BIG-IP Local Traffic Manager (LTM) and NetWitness Decoder integration.

NetWitness Decoder allows administrators to receive session keys from PFS sessions and decrypt each session traffic. It can receive session keys through its RESTful port, the PFS encrypted traffic through its capture interface, and supports the traffic decryption for inspection.

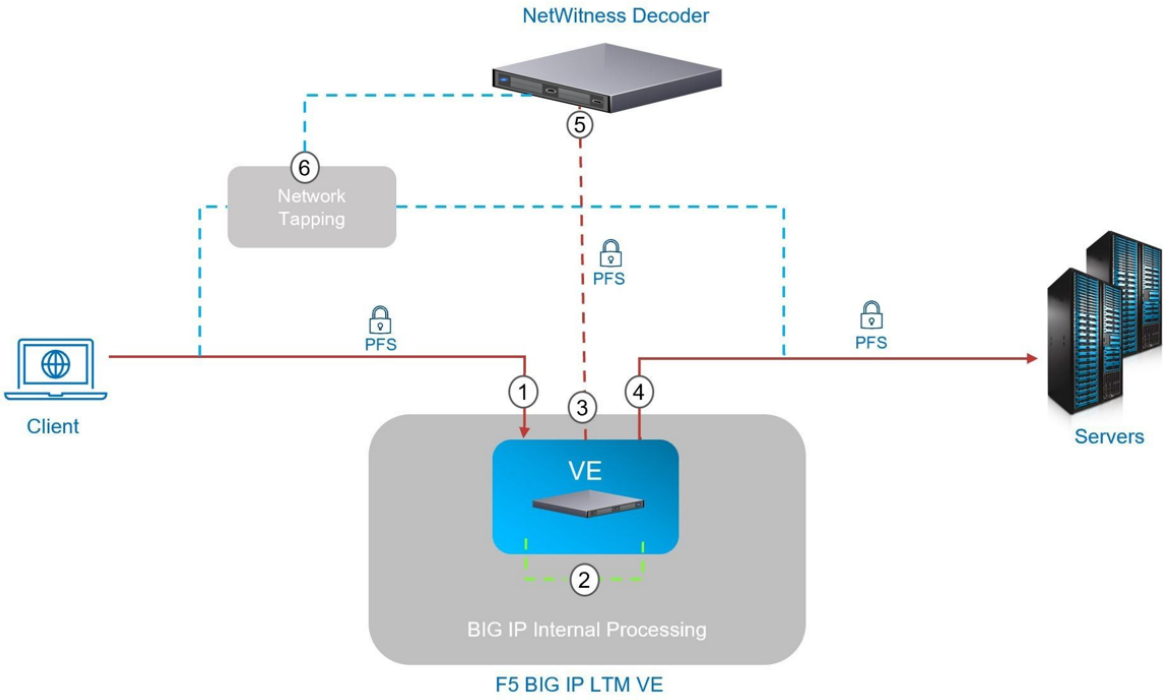
This solution utilizes functionality in iRule to extract session keys and deliver them to NetWitness Decoder via a sideband connection.

Prerequisites

A few prerequisites need to be considered in your environment for the integration of BIG-IP LTM and NetWitness Decoder.

- Deploy and Configure BIG IP LTM External Virtual Server (version 15.1.5 and above)
  - Configure Client SSL Profile with TLS 1.3 enabled.
  - Configure Server SSL Profile with TLS 1.3 enabled.
- Deploy and Configure NetWitness Decoder NON-FIPS (version 12.0 and above)
  - Install Decoder `nonfips` builds which support TLS 1.3 decryption.
  - Allow NetWitness Decoder RESTful port (default 50104) in your Network Firewall to receive keys from BIG-IP LTM.
    - NetWitness Decoder receives keys via HTTPS on RESTful port.
  - Create an Administrator account on the Decoder Device with role permissions `security.roles: decoder.manage` to forward TLS keys to Decoder.
    - Example name: bigip-service
  - Consider deploying NetWitness Decoder close to BIG-IP LTM so that it's feasible to mirror the encrypted traffic via TAP or Switch and avoid Network Latency in receiving TLS keys.

Deployment



1. A client connects to BIG-IP LTM using TLS with PFS cipher. The External Virtual Server has `clientssl` profile to decrypt traffic for inspection.
2. While inspecting the traffic an `iRule` on External Virtual Server collects the session keys and makes a `sideband` connection to Internal Virtual Server with the NetWitness Decoder as a `pool member`. The Internal Virtual Server has a `serverssl` profile to ensure traffic between BIG-IP LTM and NetWitness Decoder is protected since session keys are forwarded.
3. The Internal Virtual Server has a `HTTP profile` applied to manage the key transfer via HTTPS between BIG-IP LTM and NetWitness Decoder. `Oneconnect` profile is utilized and session keys will be sent on existing connections if available to reduce the connection setup overhead with NetWitness Decoder.
4. The External Virtual Server also has a `serverssl` profile to re-encrypt the traffic back to server pool.
5. The NetWitness Decoder receives the TLS keys from BIG-IP LTM on its RESTful port and decrypts the traffic ingested.
6. The Payload is not sent from BIG-IP LTM to NetWitness Decoder, this is received either from port mirroring configurations or a Network TAP infrastructure as shown above.

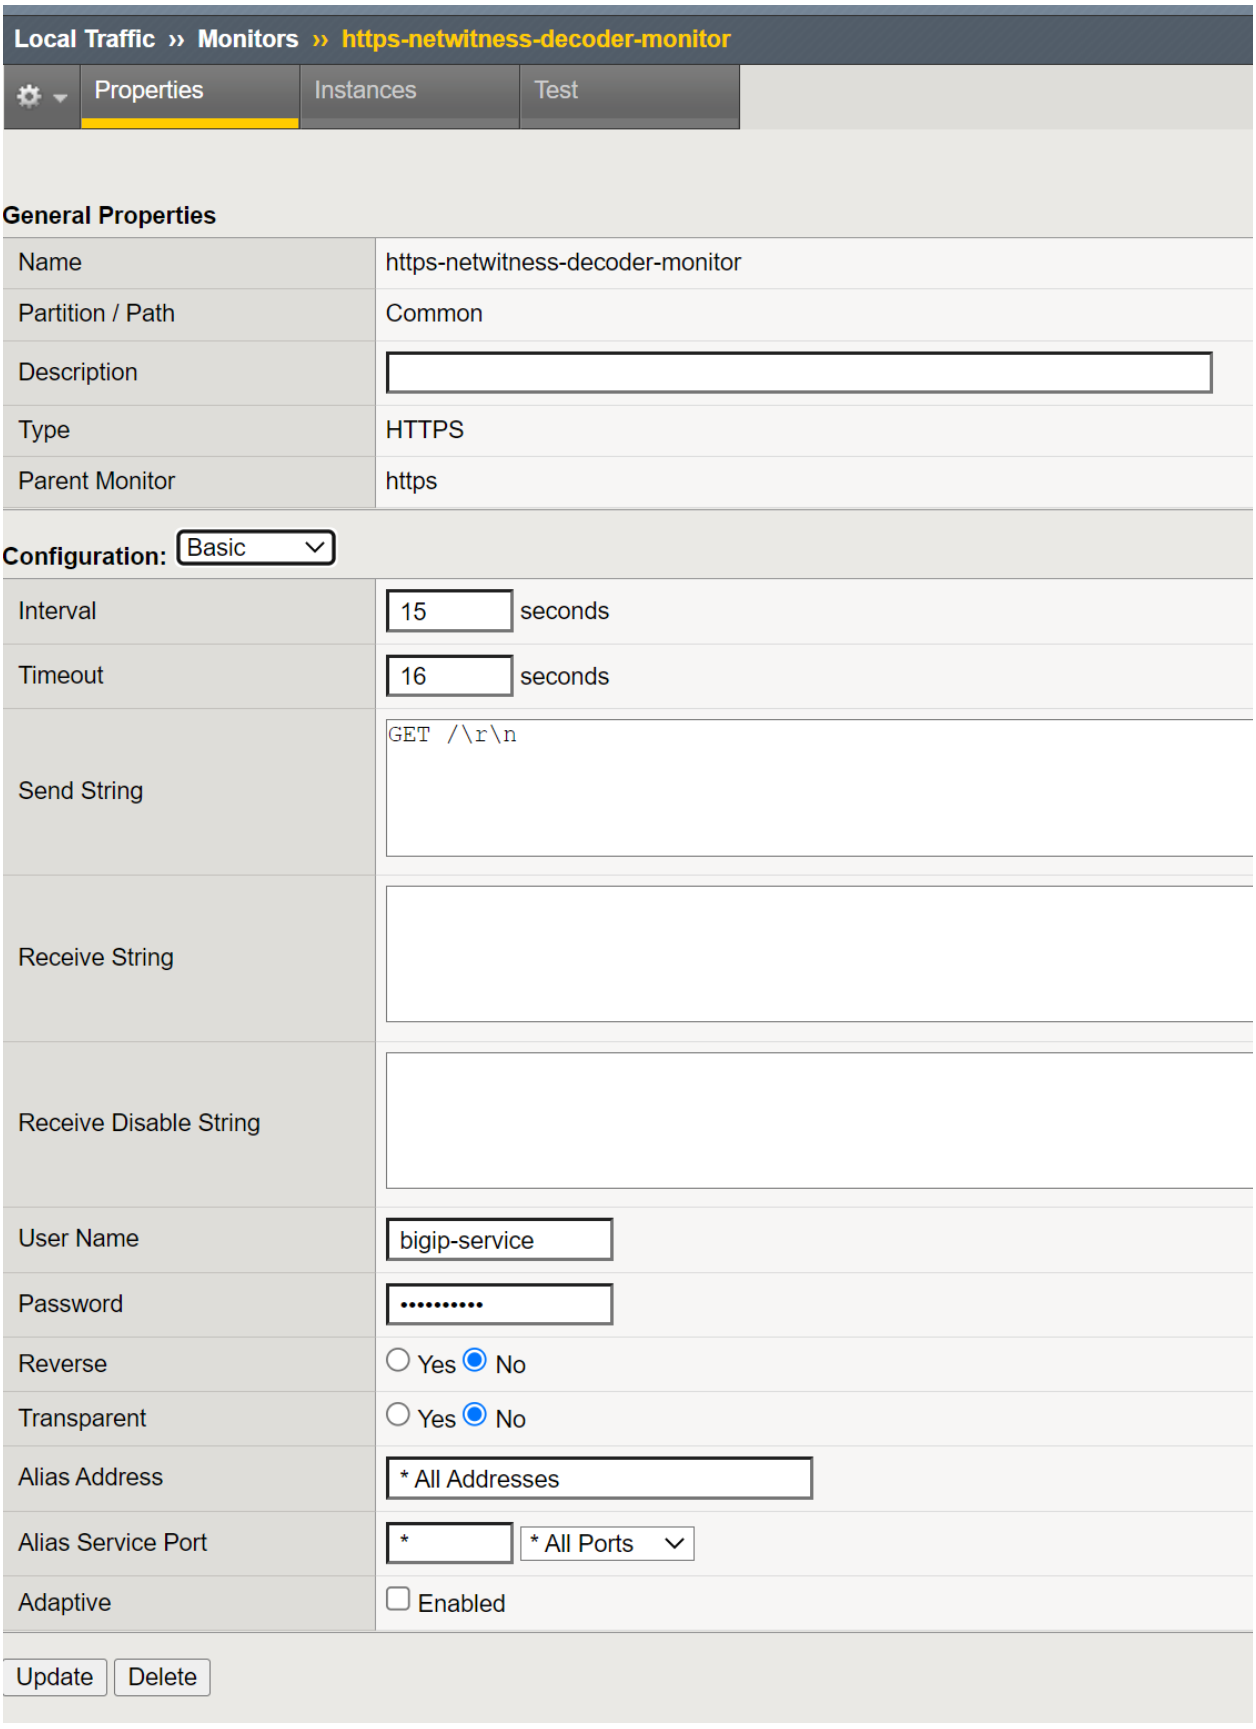
BIG-IP LTM Configurations

The BIG-IP configuration objects required include the following:

- HTTPS Monitor for monitoring NetWitness Decoder Node and Pool status (e.g online, offline etc…)
- Node and Pool objects for NetWitness Decoder.
- HTTP profile for managing key transfer to NetWitness Decoder via HTTPS.
- Server SSL profile for HTTPS sideband connection to NetWitness Decoder.
- Internal Virtual Server for receiving session keys and sideband communication with NetWitness Decoder.
- iRule: for session keys copying and sideband communication with Internal Virtual Server.

HTTPS Monitor

- Configure HTTPS monitor for NetWitness Decoder.
- Set the User Name and Password of the Administrator account created on Decoder device for BIG-IP transactions. Example Configuration using bigip-service account.

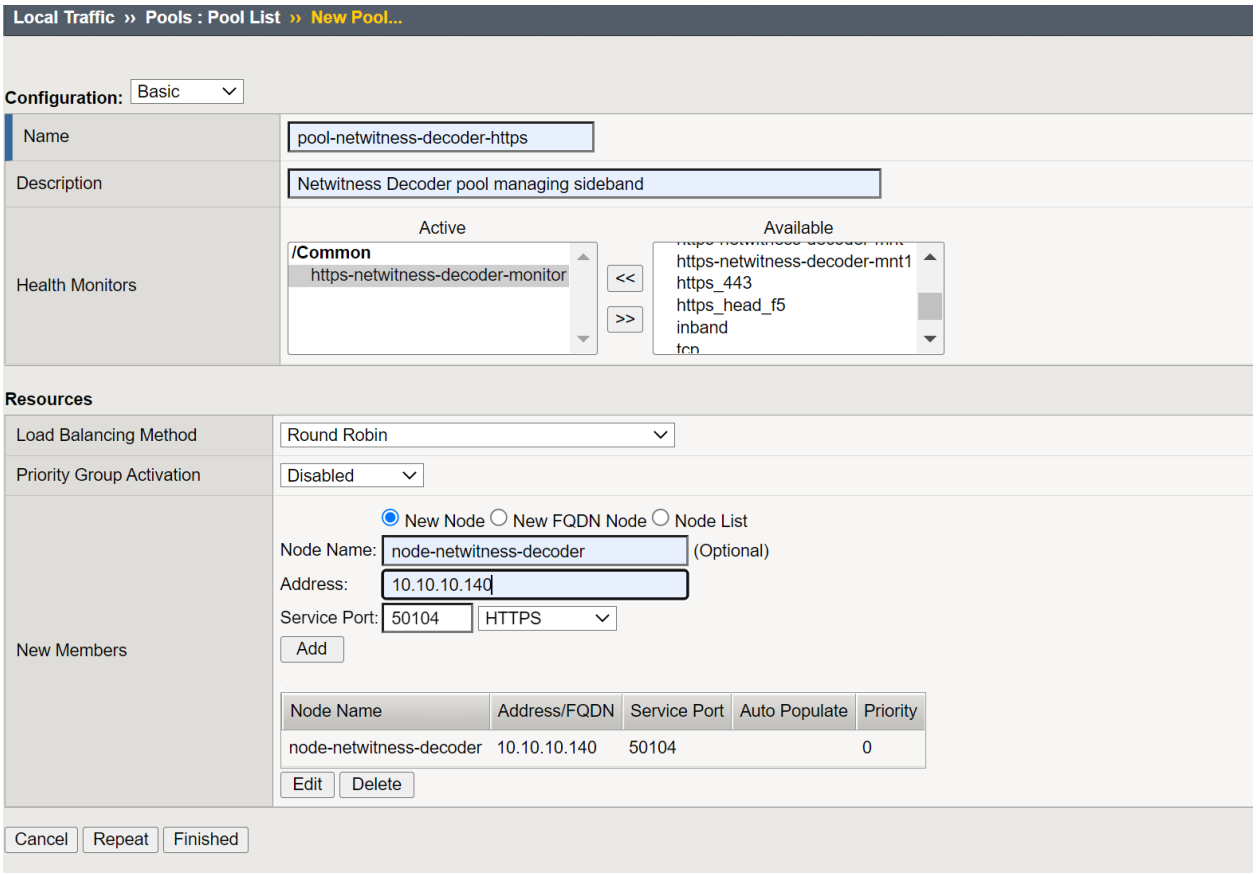

Local Traffic » Monitors » https-netwitness-decoder-monitor
Properties
Instances
Test
General Properties
Name
https-netwitness-decoder-monitor
Partition / Path
Common
Description
Type
HTTPS
Parent Monitor
https
Configuration: Basic
Interval
15
seconds
Timeout
16
seconds
Send String
GET /\r\n
Receive String
Receive Disable String
User Name
bigip-service
Password
Reverse
Yes
No
Transparent
Yes
No
Alias Address
* All Addresses
Alias Service Port
*
* All Ports
Adaptive
Enabled
Update
Delete

```
ltm monitor https /Common/https-netwitness-decoder-monitor
{
adaptive disabled
defaults-from /Common/https
destination *.50104
interval 15
ip-dscp 0
password <encrypted-bigip-service-account-password>
recv none
recv-disable none
send "GET /\r\n"
ssl-profile /Common/serverssl-secure
time-until-up 0
timeout 16
username bigip-service
}
```

Pool & Node

Configure Pool and Node objects.

- Select **Health Monitors** `https-NetWitness-decoder-monitor`
- Add **New Node** with the following:
  - Set **Node Name** example: `node-NetWitness-decoder`
  - **Address**: <Decoder IP Address >, **Service Port**: 50104 HTTPS and then click **Add**.
- Click **Finished** to create Pool with a Node.




```
ltm pool /Common/pool-netwitness-decoder-https {
members {
/Common/node-netwitness-decoder:50104 {
address 10.10.10.140
}
}
monitor /Common/https-netwitness-decoder-monitor
}
```

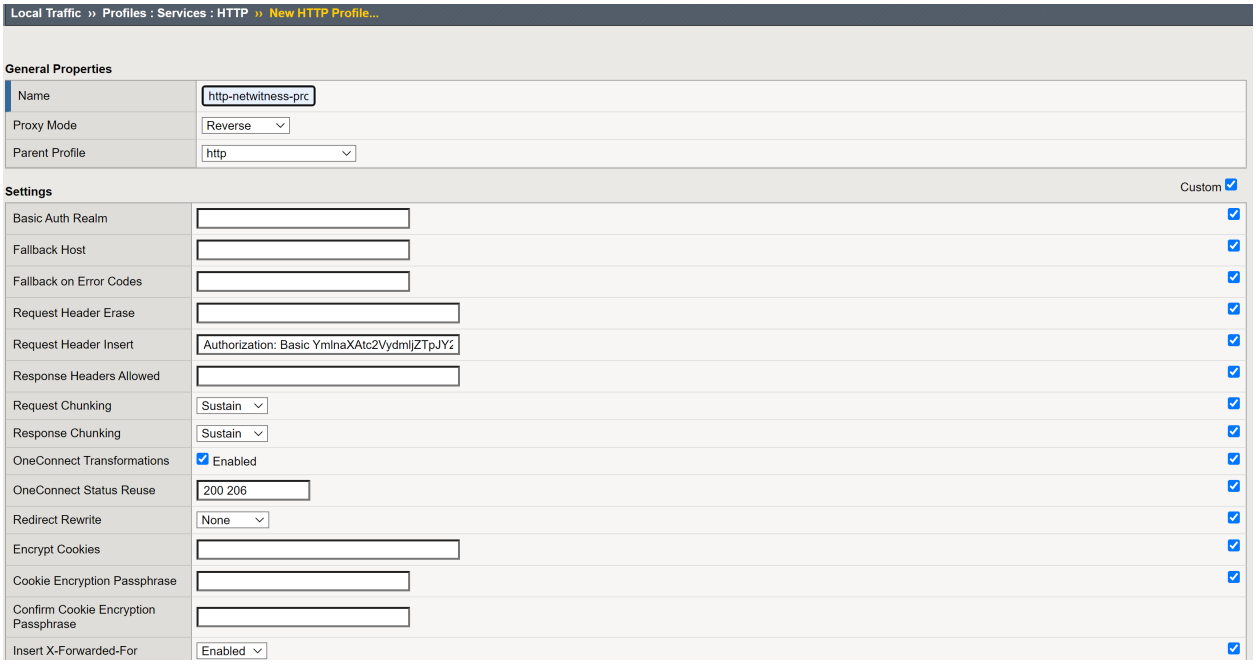
HTTP Profile

- Configure HTTP Profile **Name** `http-NetWitness-profile` to authenticate NetWitness Decoder device and send TLS keys.
- Set **Parent Profile** as **http**
- Check the **Custom** check box to edit the Settings.
- Set **Request Header Insert** field with basic Authorization header string
  - `"Authorization: Basic <Base64 encoded stirng (username:password)>"`
  - Example: `"Authorization: Basic YmlnaXAtc2VydmljZTpJY2FuU2VleW91QDEyMyE="`

- Where
  - username - Administrator account on Decoder device (example: bigip-service)
  - password - The account password

- Set **Insert X-Forwarded-For** as **Enabled** , this would add X-Forwarded-For header with the original Client IP of the session.

Rest of the other configurations can remain as default values.

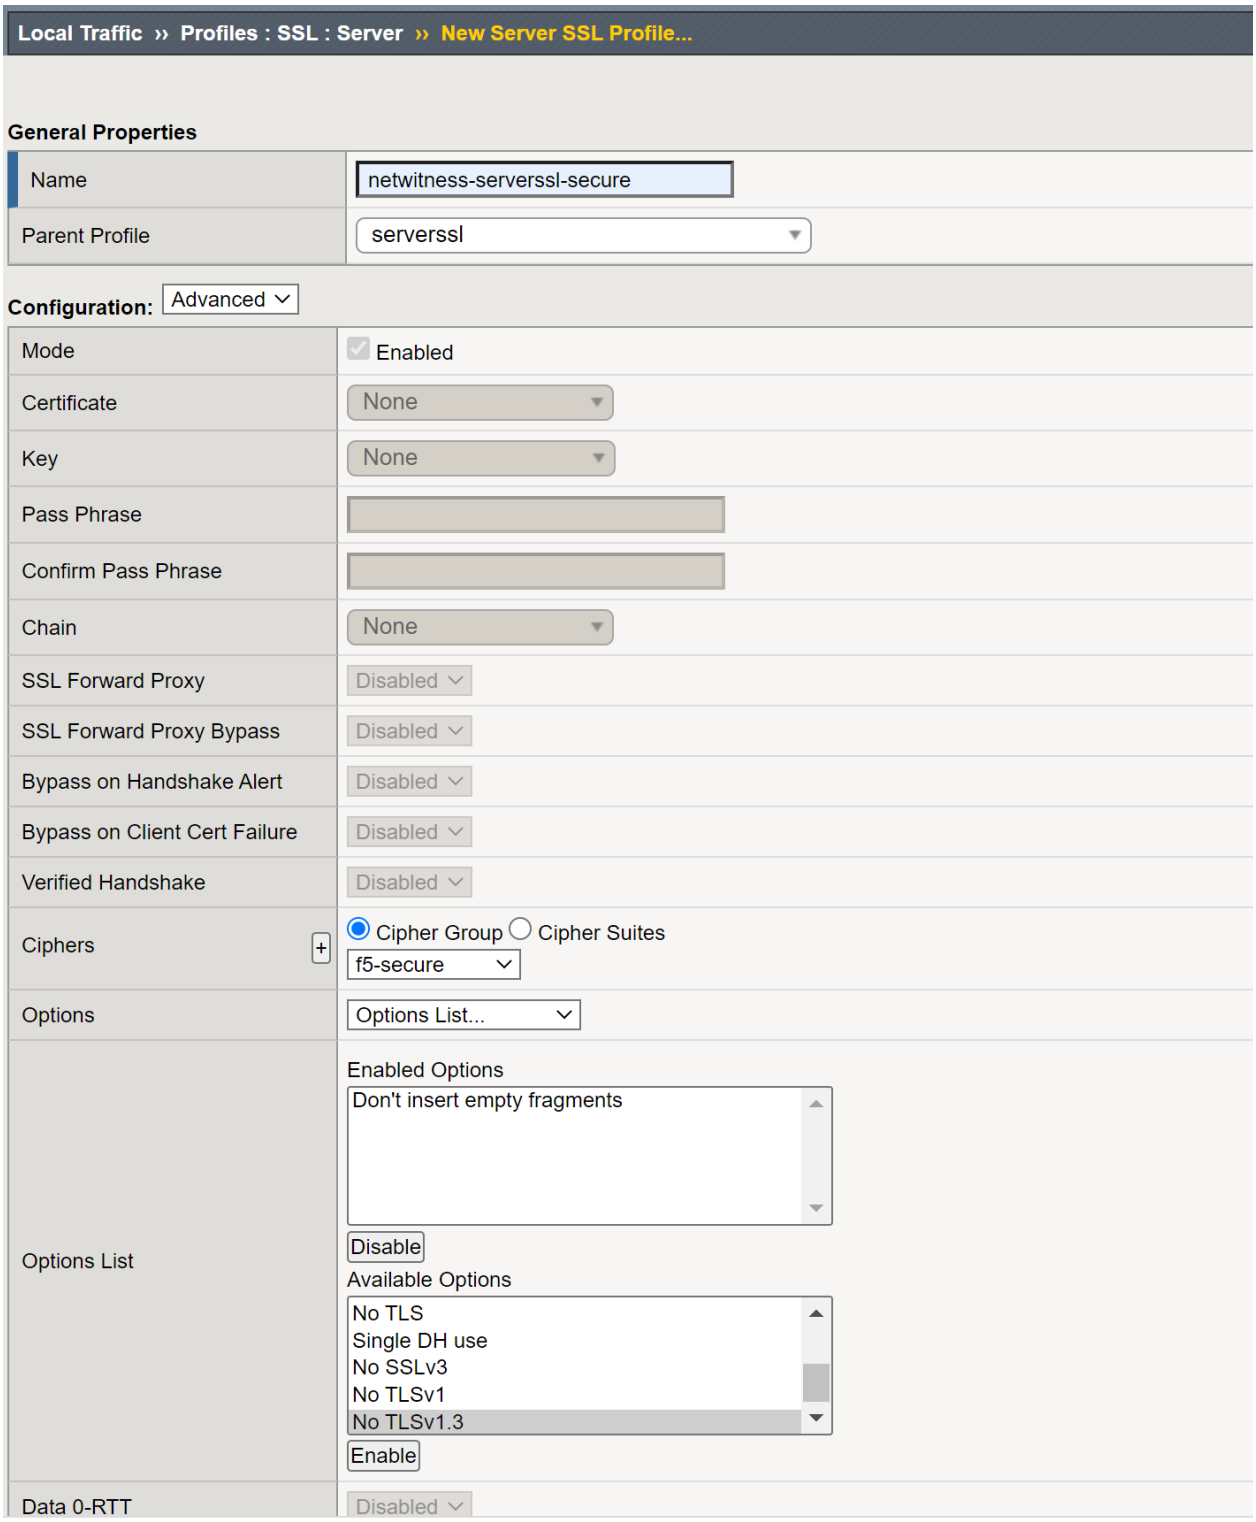



```
ltm profile http /Common/http-NetWitness-profile {
accept-xff disabled
app-service none
basic-auth-realm none
defaults-from /Common/http
encrypt-cookies none
enforcement {
known-methods { CONNECT DELETE GET HEAD LOCK OPTIONS POST PROPFIND PUT
TRACE UNLOCK }
max-header-count 64
max-header-size 32768
max-requests 0
pipeline allow
rfc-compliance disabled
truncated-redirects disabled
unknown-method allow
}
fallback-host none
```

```
fallback-status-codes none
header-erase none
header-insert "Authorization: Basic
YmlnaXAtc2VydmljZTpJY2FuU2VleW91QDEyMyE="
hsts {
include-subdomains enabled
maximum-age 16070400
mode disabled
preload disabled
}
insert-xforwarded-for enabled
lws-separator none
lws-width 80
oneconnect-status-reuse "200 206"
oneconnect-transformations enabled
proxy-type reverse
redirect-rewrite none
request-chunking sustain
response-chunking sustain
response-headers-permitted none
server-agent-name BigIP
sflow {
poll-interval 0
sampling-rate 0
}
via-request preserve
via-response preserve
xff-alternative-names none
}
```

Server SSL Profile

1. Set server-ssl profile **Name** `NetWitness-serverssl-secure`
2. Select **Advanced** from **Configuration**.
3. Set **f5-secure** from **Ciphers > Cipher Group**.
4. (Optional) **Disable** `no-TLS1.3` in options which enables TLS 1.3 between BIG-IP and NetWitness Decoder.
5. Click **Update** to create `server-ssl` profile.

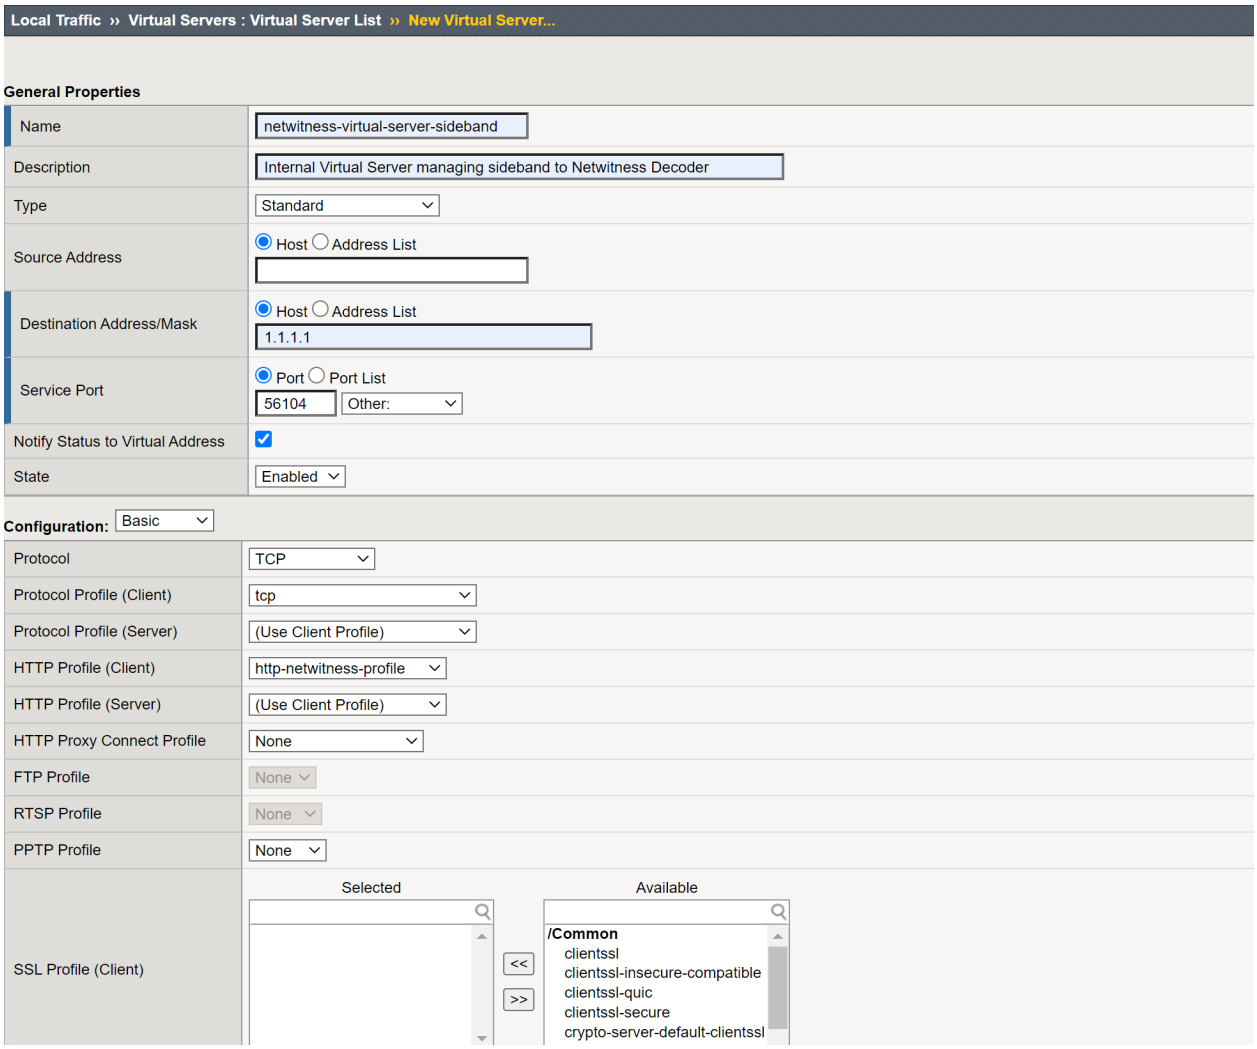



```
ltm profile server-ssl /Common/NetWitness-serverssl-secure {
app-service none
cipher-group /Common/f5-secure
ciphers none
```

```
defaults-from /Common/serverssl
options { no-ssl }
renegotiation disabled
}
```

Internal Virtual Server

1. Create an Internal Virtual server to manage sideband connection to NetWitness Decoder and forward keys.
2. Set **Name** as `NetWitness-virtual-server-sideband` and add **Description** for virtual server.
3. Let the Source Address be blank, and it would default to 0.0.0.0.
4. Set **Destination Address/Mask** to 1.1.1.1; this is a non-routable IP address and is used internally.
5. Set **Service Port** to 56104 , a virtual non-routable port used internally.
6. Select **HTTP Profile (Client)** `http-NetWitness-profile` that was created earlier. By default, Server inherit Client profile.
7. Select **SSL Profile (Server)** `NetWitness-serverssl-secure` which was created earlier. This would be used by the Pool `pool-NetWitness-decoder-https` when connecting to NetWitness Decoder.
8. Select **Source Address Translation** Auto-Map.
9. Select **OneConnect Profile** oneconnect.
10. Select **Default Pool** `pool-NetWitness-decoder-https` which was created earlier.
11. Click **Finished** to create Internal Virtual Server.




```
ltm virtual /Common/NetWitness-virtual-server-sideband {
description "Internal Virtual Server managing sideband to NetWitness
Decoder"
destination /Common/1.1.1.1:56104
ip-protocol tcp
mask 255.255.255.255
pool /Common/pool-NetWitness-decoder-https
profiles {
/Common/http-NetWitness-profile { }
/Common/NetWitness-serverssl-secure {
context serverside
}
/Common/oneconnect { }
/Common/tcp { }
}
```

```
serverssl-use-sni disabled
source 0.0.0.0/0
source-address-translation {
type automap
}
translate-address enabled
translate-port enabled
}
```

iRule Session keys copying and Sideband communication

Configuration

1. Configure and add the following `netwitness-sideband-irule` **iRule** in your **External Virtual Server**, as an Administrator
2. Set **variable** `static::virtual_server` in `RULE_INIT` with name of **Internal Virtual Server** of your environment.
   - Example `set static::virtual_server "NetWitness-virtual-server-sideband"`
3. (Optional) If you would like to disable extracting keys for Client side communication, then comment the routine CLIENTSSL_HANDSHAKE.
4. (Optional) If you would like to disable extracting keys for Server side communication, then comment the routine SERVERSSL_HANDSHAKE.
5. (Optional) For any additional troubleshooting uncomment the log messages.

iRule `NetWitness-sideband-irule`

```
# NetWitness sideband iRule.


# The script extracts TLS keys from PFS (Perfect Forward Secrecy) communication between client and server, then it forwards
# the keys to internal virtual server via sideband connection. The internal virtual server then send the keys to
# NetWitness Decoder for TLS Decryption.
#
# Configuration Steps:
#
#   1. (Required) As an administrator you would need to set variable static::virtual_server with name of internal virtual server of your environment.
#   2. (Optional) If you would like to disable extracting keys for Client side communication, then comment the routine CLIENTSSL_HANDSHAKE.
```

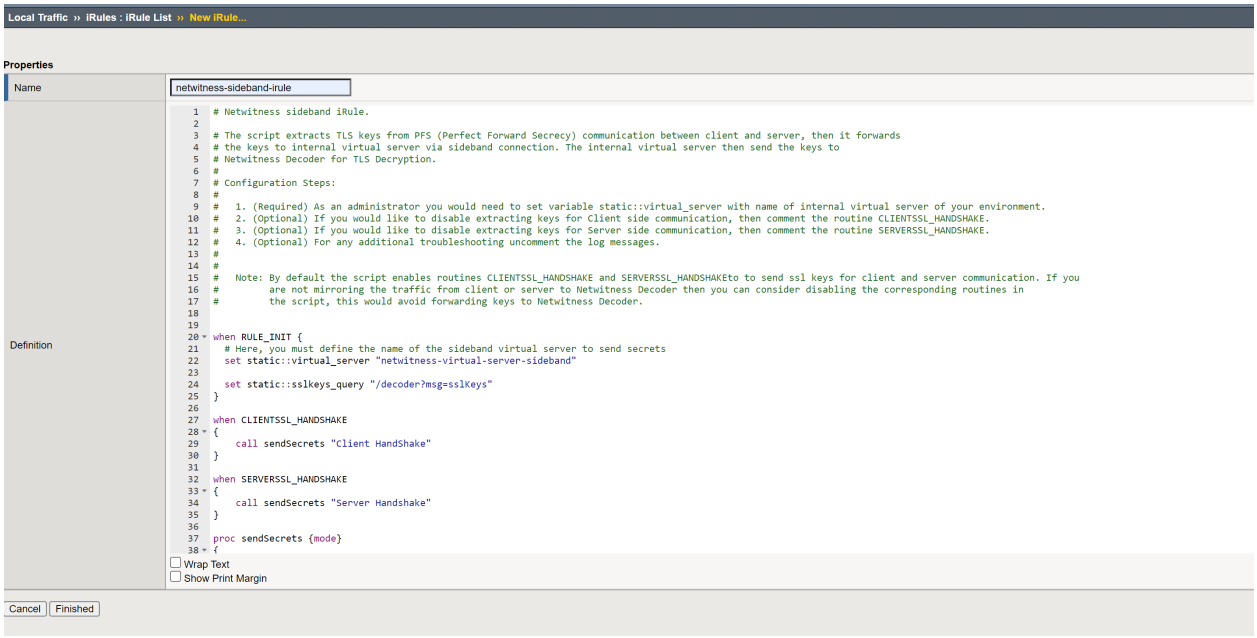
```
# 3. (Optional) If you would like to disable extracting keys for Server side communication, then comment the routine SERVERSSL_HANDSHAKE.
# 4. (Optional) For any additional troubleshooting uncomment the log messages.
#
#
# Note: By default the script enables routines CLIENTSSL_HANDSHAKE and SERVERSSL_ HANDSHAKEto to send ssl keys for client and server communication. If you
#       are not mirroring the traffic from client or server to NetWitness Decoder then you can consider disabling the corresponding routines in
#       the script, this would avoid forwarding keys to NetWitness Decoder.
when RULE_INIT {
# Here, you must define the name of the sideband virtual server to send secrets
set static::virtual_server "NetWitness-virtual-server-sideband"
set static::sslkeys_query "/decoder?msg=sslKeys"
}
when CLIENTSSL_HANDSHAKE
{
call sendSecrets "Client HandShake"
}
when SERVERSSL_HANDSHAKE
{
call sendSecrets "Server Handshake"
}
proc sendSecrets {mode}
{
if {[SSL::cipher version] == "TLSv1.3" } {
#log local0. "$mode: CLIENT_HANDSHAKE_TRAFFIC_SECRET [SSL::clientrandom] [SSL::tls13_ secret client hs]"
set secrets "random=[SSL::clientrandom]"
if { [SSL::tls13_secret client early] ne "" } {
set secrets "$secrets&CETS=[SSL::tls13_secret client early]"
}
set secrets "$secrets&CHTS=[SSL::tls13_secret client hs]&SHTS=[SSL::tls13_secret server hs]&CTS0=[SSL::tls13_secret client app]&STS0=[SSL::tls13_secret server app]"
if { [catch {call sidebandSend $secrets} ] } {
#log local0. "$mode $static::virtual_server sideband is not reachable, ssl version [SSL::cipher version]"
return
}
```

```
}
else
{
set secrets "random=[SSL::clientrandom]&premaster=[SSL::sessionsecret]"
if { [catch {call sidebandSend $secrets} ] } {
#log local0. "$mode $static::virtual_server sideband is not reachable, ssl version [SSL::cipher version]"
return
}
}
}
proc sidebandSend {secrets} {
#log local0. "$static::virtual_server launching sidebandSend ..."
set cmp_unit [TMM::cmp_unit]
# Key for session table
set key "${cmp_unit}_conn_${static::virtual_server}"
# Session table data for $key for this dest_addr
set conn [session lookup dest_addr $key]
if { $conn eq "" } {
set conn [connect -timeout 1000 -idle 300 -status conn_status $static::virtual_server]
if { $conn_status ne "connected" }{
#log local0. "Failed to connect to virtual_server $static::virtual_server"
return
}
session add dest_addr $key "$conn" 300
}
else {
# Attempt sideband connection re-use
set conn_info [connect info -status $conn]
set conn_state [lindex [lindex $conn_info 0] 0]
if { $conn_state ne "connected" }{
set conn [connect -timeout 1000 -idle 300 -status conn_status $static::virtual_server]
if { $conn_status ne "connected" }{
#log local0. "Failed to connect on connection re-use to virtual_server $static::virtual_server"
return
}
session add dest_addr $key "$conn" 300
```

```
}
}
#log local0. "Sending keys to virtual_server $static::virtual_server ..."
set send_bytes [send -timeout 1000 -status send_status $conn "GET $static::sslkeys_query&$secrets HTTP/1.0\r\n\r\n"]
#log local0. "Sent keys to virtual_server $static::virtual_server ..."
recv -timeout 1 $conn
}
```

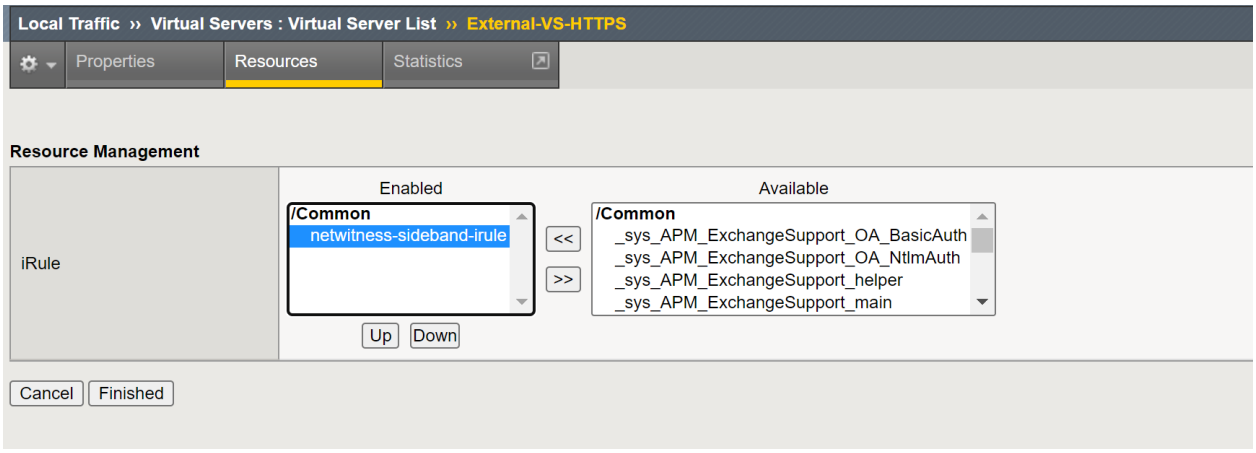
Copy

1. **Create iRule** in iRule List create wizard
2. Copy the above iRule content to **Definition**
3. Click **Finished** to save iRule



Example: Add iRule to External Virtual Server

Select **iRule** `netwitness-sideband-irule` and click **Finished** to add iRule.


Local Traffic ›› Virtual Servers : Virtual Server List ›› External-VS-HTTPS
Properties
Resources
Statistics
Resource Management
iRule
Enabled
/Common
netwitness-sideband-irule
Up
Down
<<
>>
Available
/Common
_sys_APM_ExchangeSupport_OA_BasicAuth
_sys_APM_ExchangeSupport_OA_NtlmAuth
_sys_APM_ExchangeSupport_helper
_sys_APM_ExchangeSupport_main
Cancel
Finished

Troubleshooting Packet Drops (11.x and above)

Packet drops can occur if there is backup happening at Packet Pool or Session Pool on the Decoder, and eventually, Decoder runs out of free capture pages resulting in drops.

Quick configuration Checks to avoid packet drops

The Decoder would log warning messages when it encounters packet drops. These logs contain possible reasons for drops, and you can solve some of the drop symptoms through simple configuration checks.

To check and tune the configuration:

In most cases, the decoder configuration parameters would be different from the default configuration or with the hardware deployed. So, make sure to check the following and fix configuration issues.

1. When `/var/log/messages` logs show drops with reason "check capture configuration, packet sizes".
   - NwDecoder[15913]: [Packet] [warning] **Packet drops encountered, packet capture** (737626/737628): **check capture configuration, packet sizes and rates**
   - For 10G decoder run

     ```
     /decoder reconfig op=10g update=1
     ```

     Or

     Explore view | /decoder | properties | select reconfig | op=10g update=1 | send
   - For normal decoder run

     ```
     /decoder reconfig update=1
     ```

     Or

     Explore view | /decoder | properties | select reconfig | update=1 | send
   - The Decoder service needs to be restarted for changes to be effective.
   - Monitor decoder for drops
2. When `/var/log/messages` logs show drops with reason "check packet database configuration, iostats, packet and content calls"
   - NwDecoder[74030]: [Packet] [warning] **Packet drops encountered, packet write** (717957/723314): **check packet database configuration, iostats, packet and content calls**
   - Packet drops encountered (884632/884642): **check session & meta database configuration, iostats and sdk activity**.
   - For 10G decoder run
     - `/database reconfig op=10g update=1`

       Or

- Explore view | /database | properties | select reconfig | op=10g update=1 | send

- For normal decoder run
  - `/database reconfig update=1`

    Or
  - Explore view | /database | properties | select reconfig | update=1 | send
- The Decoder service needs to be restarted for changes to be effective.
- If we are still seeing drops, then a few more changes are required. There could be i/o bound waits for database writes on the Decoder.
  - Set `/database/config/packet.integrity.flush=normal`
  - Set `/database/config/session.integrity.flush=normal`
  - Set `/database/config/meta.integrity.flush=normal`
- Monitor decoder for drops.

3. As of NetWitness Platform version 11.3.1.1, `sosreport` retrieves service and database reconfig information as well as what settings were active when the `sosreport` was retrieved. This information can be used for cross-checking purposes and can be found in the `..../sos_commands/rsa_nw_rest` directory (service-reconfig, database-reconfig, ls<service>).

Information required to troubleshoot packet drops

- Monitor Packet Drops Tool Output (Highly recommended)
- This can be accessed through REST port `http://<decoder>:50104/sdk/app/packetdrops`
- By default, the tool looks for drops in last 24 hrs and also provides options to search drops based on time ranges.
- Enable detailed stats on decoder `REST /decoder/parsers/config/detailed.stats=yes`

What if the REST port is inaccessible?

- Enable detailed stats on decoder

  `/decoder/parsers/config/detailed.stats=yes`
- Wait for new drops and then collect `sosreport`. The `sosreport` on Decoder would collect few stats db files.
- These files can be copied to your local decoder `/var/netwitness/decoder/statdb` and restart Decoder.
- Access packet drops tool on your local Decoder.

How do I troubleshoot packet drops?

The packet drops tool would help narrow down the possible cause for drops. There could be various cases involved that need verification.

The following sections describe how to use the Packet drops tool and analyze its results.

Navigate to the Packet drops tool using the REST interface

- The tool would search drop instances and list them with links.
- Start investigating the latest drop instances where the drop count is high.
- Look at other drop instances and find the pattern for drops.

Introduction to Drops tool Charts

- Incremental Packet Drops - Displays packet drops count at that instance of time
- Capture - Displays traffic ingestion rates Capture.rate in Mbps and Calculated capture rate (instantaneous rate) in Mbps.
- Packet Pool - Displays graphs for the following packet pool stats:
  - capture - /decoder/stats/pool.packet.capture → Number of free packet pages available for capture
  - assembler - /decoder/stats/pool.packet.assembler → Number of packet pages waiting to be assembled.
  - write - /decoder/stats/pool.packet.write → Number of packet pages waiting to be updated to Packet Database.
  - pool - /decoder/stats/assembler.packet.pages → Number of packet pages held in the assembler
  - export - /decoder/stats/pool.packet.export → Number of packet pages waiting to be exported
- Session Pool
  - This displays graphs for following session pool stats
  - parse - /decoder/parsers/stats/pool.session.parse → Number sessions waiting to be parsed.
  - write - /decoder/stats/pool.session.write → Number of session pages waiting to be written.
  - correlate - /decoder/stats/pool.session.correlate → Number of session pages waiting to be correlated.
  - queue.sessions.total - /decoder/parsers/stats/queue.sessions.total → The total number sessions in parse threads and queues.
  - export - /decoder/stats/pool.session.export Number of session pages waiting to be exported.

When packet drops occur, the chart shows that instance of time in yellow.

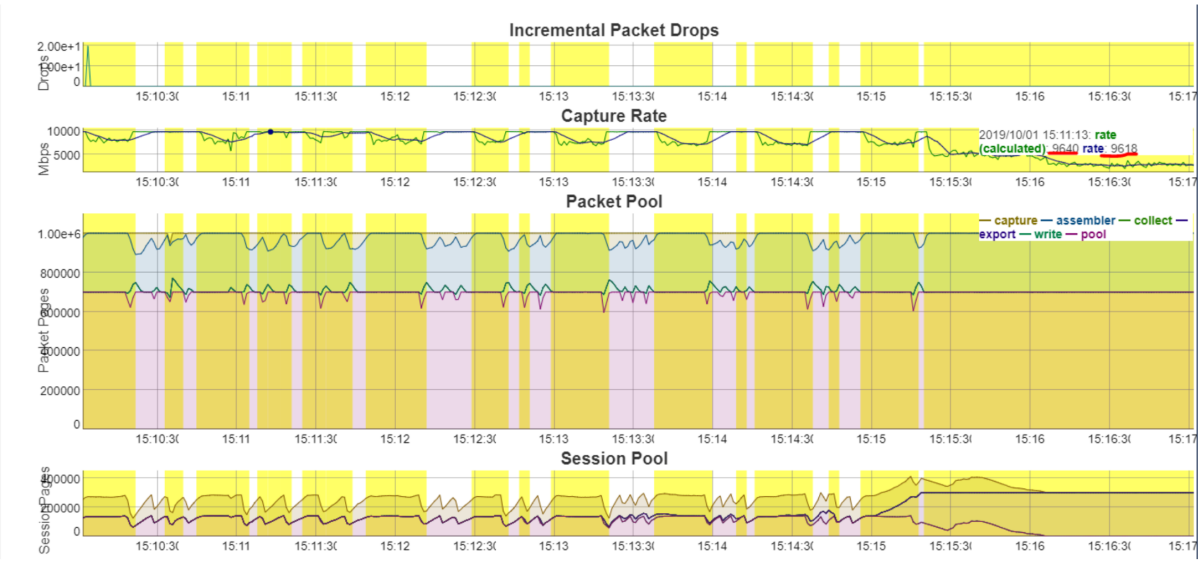
Symptoms Checklist

We recommend following the symptom checks in the order listed below, and this helps to troubleshoot all the symptoms associated with current packet drops.

Symptom 1: Higher traffic ingestion rates for the content deployed would cause packet drops

If the traffic ingestion rate -(capture) > 8 Gbps for the content (higher number of parsers) deployed on Decoder, this would cause packet drops.

ex: Below screenshot shows capture rate > 8 Gbps (9.6 Gbps)



Resolution:

Consider the following options:

- Split the traffic ingestion into multiple decoders so that the ingestion rate would be ~4-5 Gbps

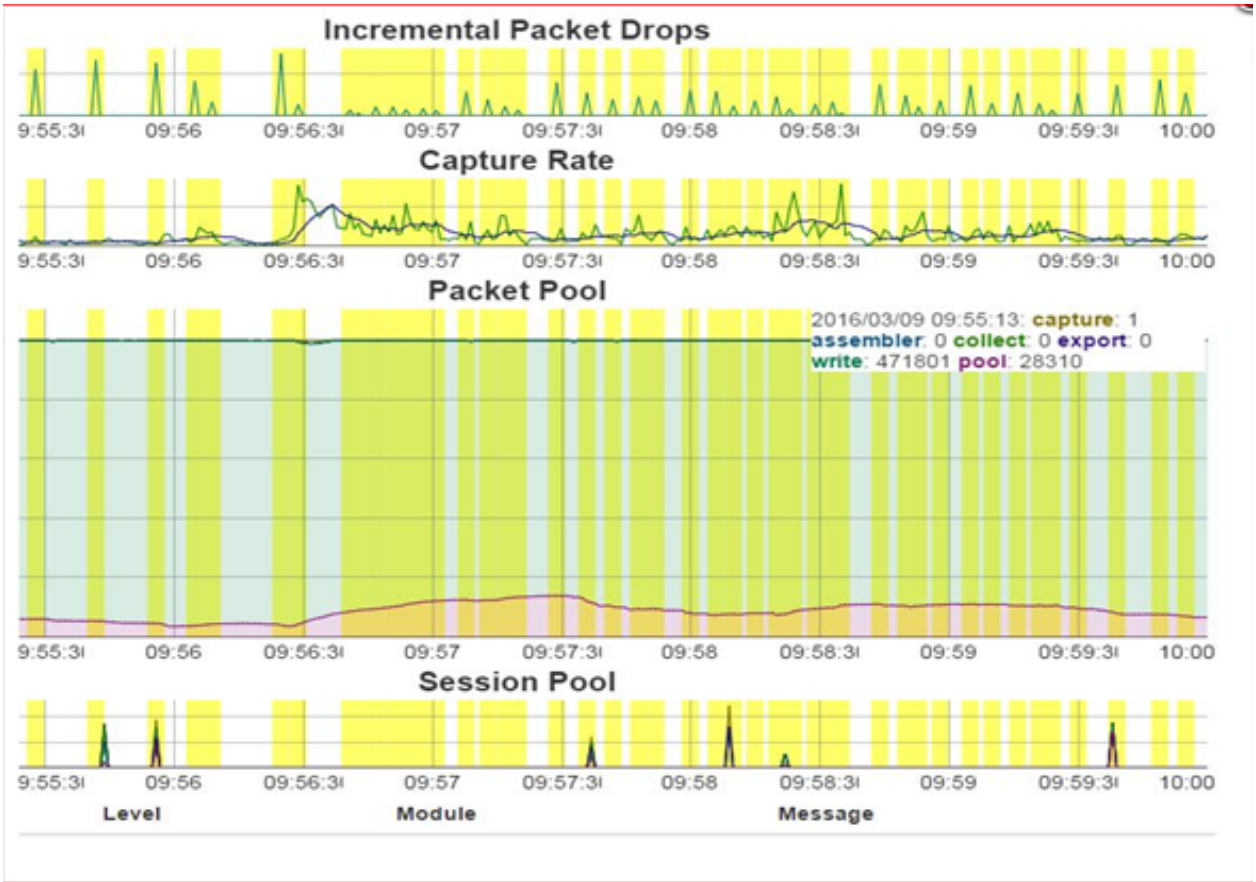
  Or

- Start with baseline supported 10G content and then add parsers one by one until no drops are observed.
  - Refer Parsing and Content Considerations section on Configure High Speed Packet Capture Capability (Version 11.6 and Later)

Symptom 2: Packet Database Write backup would cause packet drops

Packet Pool Write - write (pool.packet.write), if write backup increases, then packet write delays would be causing the drops

ex: Below screenshot displays, there are many packet pages (471K) waiting on the write queue.



Decoder logs would throw warnings like the below:

- NwDecoder[74030]: [Packet] [warning] **Packet drops encountered, packet write** (717957/723314): **check packet database configuration, iostats, packet and content calls**

Top -Hp <Decoder PID> would display decoder packet write thread waiting on Disk.

```
%Cpu(s):  0.9 us,  0.7 sy,  0.0 ni, 90.3 id,  8.1 wa,  0.0 hi,  0.0 si,  0.0 st
KiB Mem : 13148560+total, 27355252 free, 72468064 used, 31662292 buff/cache
KiB Swap:  4194300 total,  4190452 free,     3848 used. 58113828 avail Mem

  PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
131928 root      20   0   45.8g   4.6g 545220 S   9.0  3.6 452:10.36 PFRING Balancer
131851 root      20   0   45.8g   4.6g 545220 S   1.7  3.6 723:05.92 Parse Work 2
131859 root      20   0   45.8g   4.6g 545220 S   1.7  3.6 723:46.83 Parse Work 10
131850 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 722:14.60 Parse Work 1
131853 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 724:39.44 Parse Work 4
131855 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 722:58.13 Parse Work 6
131857 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 724:08.50 Parse Work 8
131858 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 724:45.26 Parse Work 9
131860 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 723:19.92 Parse Work 11
131861 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 722:48.85 Parse Work 12
131862 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 723:51.35 Parse Work 13
131863 root      20   0   45.8g   4.6g 545220 S   1.3  3.6 723:17.89 Parse Work 14
131852 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 723:02.83 Parse Work 3
131854 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 722:38.98 Parse Work 5
131856 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 722:40.62 Parse Work 7
131864 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 721:45.72 Parse Work 15
131865 root      20   0   45.8g   4.6g 545220 S   1.0  3.6 725:29.23 Parse Work 16
131838 root      20   0   45.8g   4.6g 545220 S   0.7  3.6 917:20.65 CaptureWork
131839 root      20   0   45.8g   4.6g 545220 S   0.7  3.6 526:43.29 Assem Work
131840 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  28:04.24 Packet Export W
131841 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  42:24.57 Collection ID W
131843 root      20   0   45.8g   4.6g 545220 D   0.7  3.6 303:38.71 Session Write W
131844 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  41:22.46 Session Id Work
131845 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  41:23.96 Session Correla
131849 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  41:19.76 Session Export
132932 root      20   0   45.8g   4.6g 545220 S   0.7  3.6  49:56.33 Request Handler
131771 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  31:57.08 stat updates
131842 root      20   0   45.8g   4.6g 545220 D   0.3  3.6 846:47.73 Packet Write Wo
131847 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  61:44.11 Session Index W
131848 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  38:53.76 Session Stream
131925 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  23:09.92 Timestamper
138159 root      20   0   45.8g   4.6g 545220 S   0.3  3.6  48:02.02 Request Handler
131770 root      20   0   45.8g   4.6g 545220 S   0.0  3.6  20:36.10 NwDecoder
131773 root      20   0   45.8g   4.6g 545220 S   0.0  3.6  20:24.78 service listen
131790 root      20   0   45.8g   4.6g 545220 S   0.0  3.6  51:13.34 Request Handler
131798 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   0:01.70 REST executor
131837 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   0:03.07 Request Handler
131846 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   0:00.10 FeedLoader
131868 root      20   0   45.8g   4.6g 545220 S   0.0  3.6   2:18.88 amqp thread
```

iotop tool can show Disk i/o activity: rpm is available here and can be installed: `http://mirror.centos.org/centos/7/os/x86_64/Packages/iotop-0.6-4.el7.noarch.rpm`

Example iotop results where decoder packet write thread is blocked on IO 99% of decoder IO time and its throughput is just ~400KB/s

```
iotop -o -d 2
```

```
Total DISK READ :       0.00 B/s | Total DISK WRITE :     463.05 K/s
Actual DISK READ:       0.00 B/s | Actual DISK WRITE:       0.00 B/s
  TID  PRIO  USER     DISK READ  DISK WRITE  SWAPIN     IO>    COMMAND
 3373 be/4 root        0.00 B/s  431.79 K/s  0.00 % 99.99 % NwDecoder [Packet Write Wo]
 3374 be/4 root        0.00 B/s   31.26 K/s  0.00 % 99.99 % NwDecoder [Session Write W]
```

Resolution:

- Make sure database configuration tuning is applied as suggested in the section To check and tune the configuration:

- Check I/O stats on the Decoder using the command "`iostat -mNx 1`". Refer to How do you get statistics on I/O performance?
  - If % `iowait` is > 10% then decoder packet db writes have higher i/o waits
  - If % util for packetdb goes greater than 95 and wMB/s < 1000 , then Disk write throughput is low and the Disk where packetdb exists needs to be replaced.
  - If `iotop` is installed the disk io activity can be monitored through 'iotop –o –d 2'
  - For a 10G decoder we recommend packet db disk write throughput to be 1300 MB/s (~10Gbps) for better write performance.
- Lot of Content calls to extract packets or content can cause packet write issue.
  - Check drops tool logs or `/var/log/messages` or sosreport logs for SDK-Content Calls.
  - You can also use NwConsole **topQuery** command on messages logs to identify Content calls.
  - Set `/decoder/sdk/config/packet.read.throttle=100` (a higher value) so that packet write would get preference.
  - Check service invoking SDK-Content calls and reduce the content calls.
- Kernel and Driver compatibility issues
  - Check if the firmware is updated according to Kernel version. If not update firmware.

How do you get statistics on I/O performance?

1. The command you want to run for near real time statistics on I/O usage is "`iostat -N -x -m 1`". For detailed information on the output of iostat, type "`man iostat`". If the columns do not line up, you can leave off the `-N` option, but you should probably run it once so you can see what disk groups correspond to which databases.

	%user	%nice	%system	%iowait	%steal	%idle
avg-cpu:	2.41	0.00	0.22	5.98	0.00	91.39

Device	rrqm/s	wrqm/s	r/s	w/s	rMB/s	wMB/s	avgrq-sz	avgqu-sz	await	svctm	%util
sda	0.00	0.00	0.00	1.00	0.00	0.00	8.00	0.00	3.00	3.00	0.30
sdb	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_root	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

VolGroup-lv_home	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_swap	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_nwhome	0.00	0.00	0.00	1.00	0.00	0.00	8.00	0.00	3.00	3.00	0.30
VolGroup-lv_tmp	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_vartmp	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
VolGroup-lv_varlog	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
sdc	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
sdd	0.00	0.00	242.00	0.00	30.25	0.00	255.97	1.98	8.14	4.13	100.00
sde	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
index-index	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
concentrator-root	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
concentrator-sessiondb	0.00	0.00	141.00	0.00	17.62	0.00	255.94	0.81	5.84	5.77	81.30
concentrator-metadb	0.00	0.00	102.00	0.00	12.75	0.00	256.00	1.17	11.39	9.75	99.40

2. In the above sample output, you can see the session and meta DBs (last two lines) are reading in about 30 MB/s combined. You can also see that `sdd` corresponds to those two databases, since the numbers combined roughly add up. The important takeaway is the last column, %util. This means: Percentage of CPU time during which I/O requests were issued to the device (bandwidth utilization for the device). Device saturation occurs when this value is close to 100%. Therefore, the disks are currently saturated with read requests and the databases are currently I/O bound. When measuring performance of the SA Core services, it's important to know if the software is I/O bound. This would be the limiting factor for increasing performance as the current hardware is running at capacity.
3. To determine which physical volume the index is on, run "`vgdisplay -v index`". In the above example, it's /dev/sde. To see a list of volumes, run "`vgdisplay -s`".
4. Run "`ls -l /dev/mapper`" to see the drive mappings.

Symptom: Session Database Write backup would cause packet drops

Session Pool Write - write (pool.session.write), if session write backup increases, then session write delays would be causing the drops

ex: Below screenshot displays there are many session pages (198K) waiting on session write queue



Decoder logs would also throw warnings like below:

- Packet drops encountered (884632/884642): **check session & meta database configuration, iostats and sdk activity**.

Top -Hp <Decoder PID> would display decoder packet write thread waiting on Disk.

```
Cpu(s):  0.9 us,  0.7 sy,  0.0 ni, 90.3 id,  8.1 wa,  0.0 hi,  0.0 si,  0.0 st
iB Mem : 13148560+total, 27355252 free, 72468064 used, 31662292 buff/cache
iB Swap:  4194300 total,  4190452 free,     3848 used. 58113828 avail Mem

  PID USER      PR  NI    VIRT    RES    SHR S %CPU %MEM     TIME+ COMMAND
31928 root      20   0   45.8g   4.6g 545220 S  9.0  3.6 452:10.36 PFRING Balancer
31851 root      20   0   45.8g   4.6g 545220 S  1.7  3.6 723:05.92 Parse Work 2
31859 root      20   0   45.8g   4.6g 545220 S  1.7  3.6 723:46.83 Parse Work 10
31850 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 722:14.60 Parse Work 1
31853 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 724:39.44 Parse Work 4
31855 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 722:58.13 Parse Work 6
31857 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 724:08.50 Parse Work 8
31858 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 724:45.26 Parse Work 9
31860 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 723:19.92 Parse Work 11
31861 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 722:48.85 Parse Work 12
31862 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 723:51.35 Parse Work 13
31863 root      20   0   45.8g   4.6g 545220 S  1.3  3.6 723:17.89 Parse Work 14
31852 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 723:02.83 Parse Work 3
31854 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 722:38.98 Parse Work 5
31856 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 722:40.62 Parse Work 7
31864 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 721:45.72 Parse Work 15
31865 root      20   0   45.8g   4.6g 545220 S  1.0  3.6 725:29.23 Parse Work 16
31838 root      20   0   45.8g   4.6g 545220 S  0.7  3.6 917:20.65 CaptureWork
31839 root      20   0   45.8g   4.6g 545220 S  0.7  3.6 526:43.29 Assem Work
31840 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  28:04.24 Packet Export W
31841 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  42:24.57 Collection ID W
31843 root      20   0   45.8g   4.6g 545220 D  0.7  3.6 303:38.71 Session Write W
31844 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  41:22.46 Session Id Work
31845 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  41:23.96 Session Correla
31849 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  41:19.76 Session Export
32932 root      20   0   45.8g   4.6g 545220 S  0.7  3.6  49:56.33 Request Handler
31771 root      20   0   45.8g   4.6g 545220 S  0.3  3.6  31:57.08 stat updates
```

iotop tool can show Disk i/o activity: rpm is available here and can be installed: `http://mirror.centos.org/centos/7/os/x86_64/Packages/iotop-0.6-4.el7.noarch.rpm`

Example iotop results where decoder session write thread is blocked on IO 99% of decoder IO time and its throughput is just ~30KB/s

iotop –o –d 2

```
Total DISK READ :       0.00 B/s | Total DISK WRITE :     463.05 K/s
Actual DISK READ:       0.00 B/s | Actual DISK WRITE:       0.00 B/s
  TID  PRIO  USER     DISK READ  DISK WRITE  SWAPIN     IO>    COMMAND
 3373 be/4 root        0.00 B/s  431.79 K/s  0.00 % 99.99 % NwDecoder [Packet Write Wo]
 3374 be/4 root        0.00 B/s   31.26 K/s  0.00 % 99.99 % NwDecoder [Session Write W]
```

Resolution:

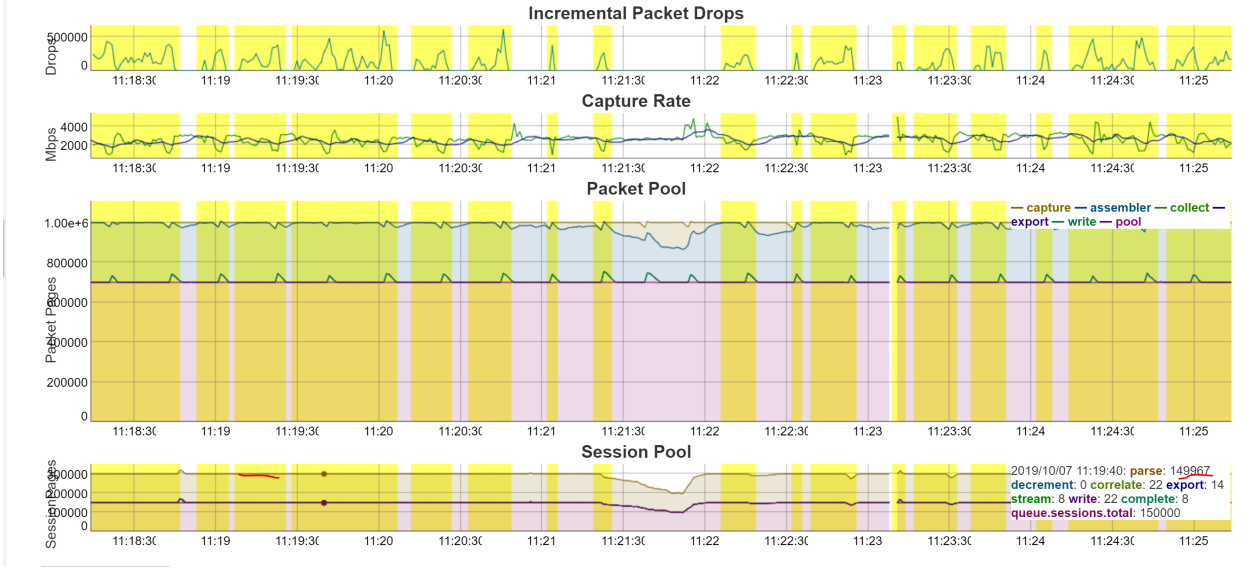
- Make sure database configuration tuning is applied as suggested in section To check and tune the configuration:
- Aggressive aggregators from Decoder
  - Check aggregators using `/decoder whoAgg` or Explore `/decoder | properties | select whoAgg | send`
  - Many aggregators with nice disabled (false) in their configuration can cause session pool backup.
    - For Archiver and Warehouse Connectors nice should be enabled (true).
    - For Concentrator nice should be disabled (false).
  - If there is one device aggregating from Decoder then it would mostly be Concentrator, don't enable nice on Concentrator.

  - If there is an Archiver or Warehouse Connector then set `aggregate.nice=true` On Warehouse Connectors and Archivers aggregating from Decoder.
- Check I/O stats on the Decoder using command "`iostat -mNx 1`". Refer to How do you get statistics on I/O performance?
  - If % `iowait` is > 10% then decoder session db and metadb writes have higher i/o waits
  - If % util for `sessiondb` or `metadb` goes greater than 95 and wMB/s < 1000 , then Disk write throughput is low and the Disk where `sessiondb` and `metadb` exists needs to be replaced.
  - If `iotop` is installed the disk io activity can be monitored through '`iotop -o -d 2`'
  - For a 10G decoder we recommend `sessiondb` and `metadb` write throughput to be 1300 MB/s (~10Gbps) for better session write performance.
- Lot of Content calls to extract meta or content can cause session write issue.
  - Check drops tool logs or `/var/log/messages` or `sosreport logs` for SDK-Content Calls
  - You can also use NwConsole **topQuery** command on messages logs to identify Content calls.
  - Set `/decoder/sdk/config/packet.read.throttle=100` (a higher value) so that packet write would get preference.
  - Check service invoking SDK-Content calls and reduce the content calls.
- Kernel and Driver compatibility issues
  - Check if the firmware is updated according to the Kernel version. If not, update firmware.

Symptom: Session Pool Parsing delays causes packet drops (Parsing issues)

Session Pool Parse - parse (pool.session.parse): If session parse backup increases, then sessions would be queued on assembler and would be causing the drops.

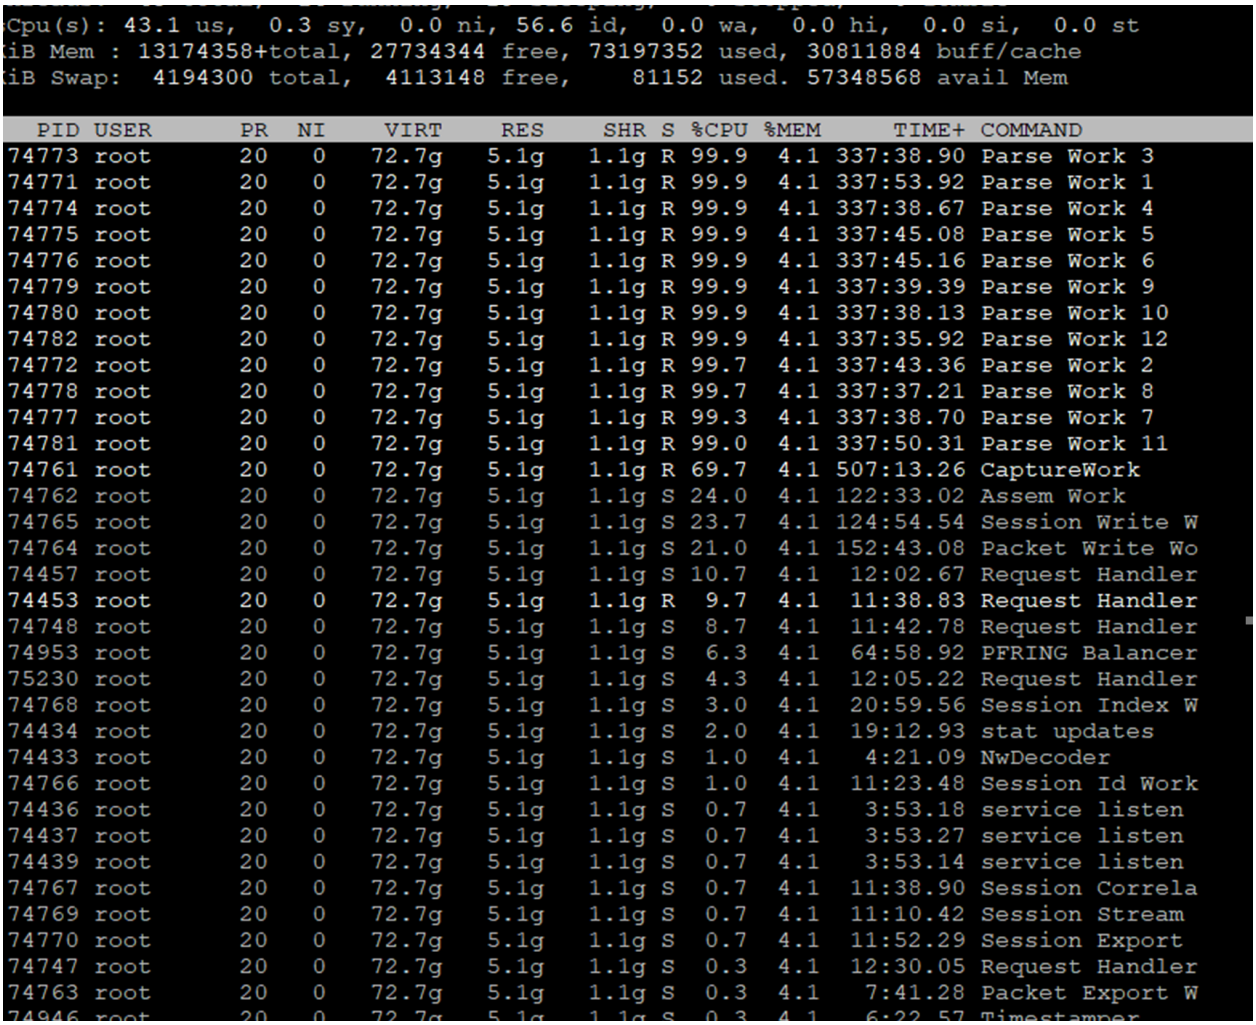
ex: Below screenshot displays there are many sessions waiting to be parsed - Session Pool | parse stat.



Decoder logs would throw warnings like the below:

NwDecoder[42946]: [Packet] [warning] Packet drops encountered, packet assemble (910178/910179): **check session pool (following log), line and session rates**

top -Hp <Decoder pid> would display decoder Parse Work (Parser Threads) are so busy with high CPU usage.



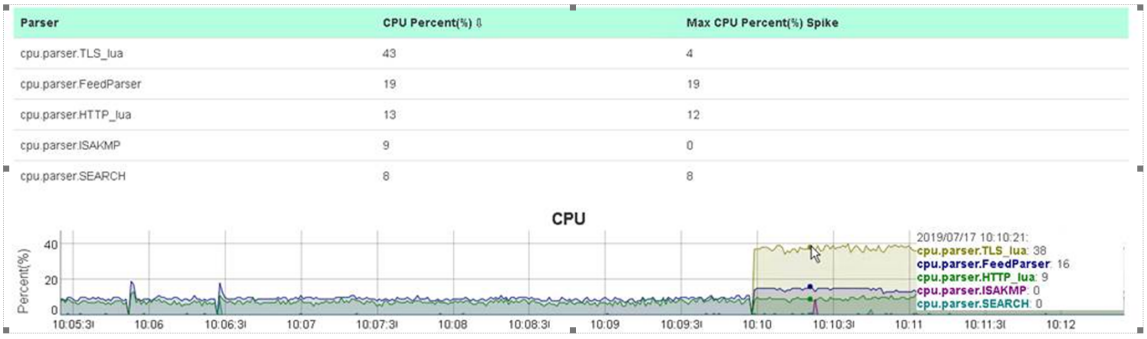
```
Cpu(s): 43.1 us,  0.3 sy,  0.0 ni, 56.6 id,  0.0 wa,  0.0 hi,  0.0 si,  0.0 st
iB Mem : 13174358+total, 27734344 free, 73197352 used, 30811884 buff/cache
iB Swap:  4194300 total,  4113148 free,    81152 used. 57348568 avail Mem

  PID USER      PR  NI    VIRT    RES    SHR S %CPU %MEM     TIME+ COMMAND
74773 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:38.90 Parse Work 3
74771 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:53.92 Parse Work 1
74774 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:38.67 Parse Work 4
74775 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:45.08 Parse Work 5
74776 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:45.16 Parse Work 6
74779 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:39.39 Parse Work 9
74780 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:38.13 Parse Work 10
74782 root      20   0   72.7g   5.1g   1.1g R 99.9  4.1 337:35.92 Parse Work 12
74772 root      20   0   72.7g   5.1g   1.1g R 99.7  4.1 337:43.36 Parse Work 2
74778 root      20   0   72.7g   5.1g   1.1g R 99.7  4.1 337:37.21 Parse Work 8
74777 root      20   0   72.7g   5.1g   1.1g R 99.3  4.1 337:38.70 Parse Work 7
74781 root      20   0   72.7g   5.1g   1.1g R 99.0  4.1 337:50.31 Parse Work 11
74761 root      20   0   72.7g   5.1g   1.1g R 69.7  4.1 507:13.26 CaptureWork
74762 root      20   0   72.7g   5.1g   1.1g S 24.0  4.1 122:33.02 Assem Work
74765 root      20   0   72.7g   5.1g   1.1g S 23.7  4.1 124:54.54 Session Write W
74764 root      20   0   72.7g   5.1g   1.1g S 21.0  4.1 152:43.08 Packet Write Wo
74457 root      20   0   72.7g   5.1g   1.1g S 10.7  4.1  12:02.67 Request Handler
74453 root      20   0   72.7g   5.1g   1.1g R  9.7  4.1  11:38.83 Request Handler
74748 root      20   0   72.7g   5.1g   1.1g S  8.7  4.1  11:42.78 Request Handler
74953 root      20   0   72.7g   5.1g   1.1g S  6.3  4.1  64:58.92 PFRING Balancer
75230 root      20   0   72.7g   5.1g   1.1g S  4.3  4.1  12:05.22 Request Handler
74768 root      20   0   72.7g   5.1g   1.1g S  3.0  4.1  20:59.56 Session Index W
74434 root      20   0   72.7g   5.1g   1.1g S  2.0  4.1  19:12.93 stat updates
74433 root      20   0   72.7g   5.1g   1.1g S  1.0  4.1   4:21.09 NwDecoder
74766 root      20   0   72.7g   5.1g   1.1g S  1.0  4.1  11:23.48 Session Id Work
74436 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1   3:53.18 service listen
74437 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1   3:53.27 service listen
74439 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1   3:53.14 service listen
74767 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1  11:38.90 Session Correla
74769 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1  11:10.42 Session Stream
74770 root      20   0   72.7g   5.1g   1.1g S  0.7  4.1  11:52.29 Session Export
74747 root      20   0   72.7g   5.1g   1.1g S  0.3  4.1  12:30.05 Request Handler
74763 root      20   0   72.7g   5.1g   1.1g S  0.3  4.1   7:41.28 Packet Export W
74946 root      20   0   72.7g   5.1g   1.1g S  0.3  4.1   6:22.57 Timestamper
```

Resolution:

- Enable parser detailed stats on decoder `/decoder/parsers/config/detailed.stats=yes`, wait for new drop instance.
- Click on 

to analyse parsing activity.
- Click on Top Parsers CPU usage, this displays top 5 parsers using higher cpu usages.
  - Try to disable the parsers which are using high CPU usage and wouldn't be parsing the larger percentage of traffic and monitor decoder for drops due to parsing.
    - For ex: parser like TLS_lua , TLD_Lua has higher CPU usage than HTTP_lua (which parsers larger percent of traffic).



Parser	CPU Percent(%) ⇩	Max CPU Percent(%) Spike
cpu.parser.TLS_lua	43	4
cpu.parser.FeedParser	19	19
cpu.parser.HTTP_lua	13	12
cpu.parser.ISAKMP	9	0
cpu.parser.SEARCH	8	8

- Repeat the above step till packet drops are reduced or resolved
- Notify parser Content owner (Customer - if it is their custom content or RSA Content team if it is RSA parser) so that a fix may be provided for parser issue.

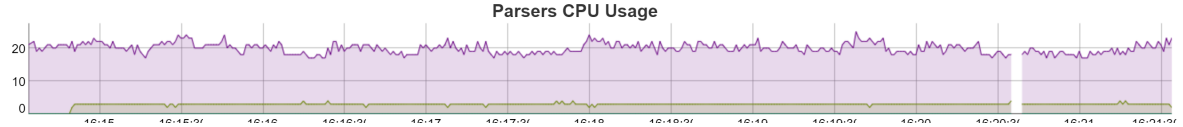
- Click Top Parsers Token Callback counts, this displays top 5 parsers having higher token callback counts during the drop window.
- Try to disable the parsers which are using high Token Callback count and wouldn't be parsing the larger percentage of traffic. Monitor decoder for drops due to parsing.

Top Parsers TokenCallback count

Parser	TokenCallbacks Count ⇩	Max TokenCallbacks Count Spike
TLS_lua	66656555	314517
HTTP_lua	14169879	151699
china_chopper	10648606	116502
rekaf	10604218	115998
Form_Data_lua	8871783	97752

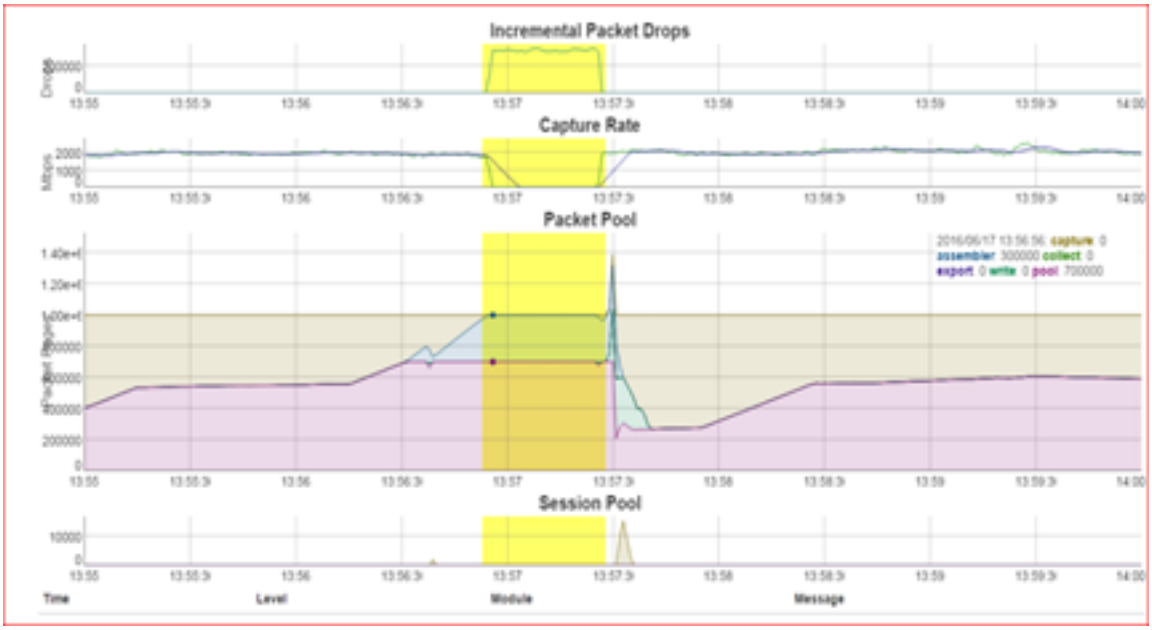
- For ex: parser like TLS_lua has higher Token Callback count than HTTP_lua (which parsers larger percent of traffic)

- Click Top Parsers Meta Callback counts, this displays top 5 parsers having higher meta callback counts during drop window.
- The parsers which displays higher Token Callback count and does not generate much meta and wouldn't parse larger percent of traffic, needs to be disabled.
  - Disable such parsers and monitor Decoder.
- Click Top Parsers Memory Usage, high Lua memory usage (>1 Gb per parser) or memory spike can cause packet drops.
  - Disable Lua Parsers which are using higher memory usage > 1 Gb and monitor Decoder for drops.
- Complex App rules including "regex" or "contains" are CPU intensive
  - Check Parsers CPU Usage graph, if there is high CPU usage for cpu.app.rules stat, then check the app rules configured and tune them.



Symptom: Session Parsing stuck can cause packet drops

Session parsing got stuck at a parser would show symptom of graph like below where queue.sessions.total =1 on Session pool.



Decoder would log warning message like:

NwDecoder[42946]: [Packet] [warning] Packet drops encountered, packet assemble (910178/910179): **check session pool (following log), line and session rates**.

Resolution:

- Enable parser detailed stats on decoder `/decoder/parsers/config/detailed.stats=yes`, wait for new drop instance.
- Click on Analyse Parser Stats to analyse parsing activity.
- Click on Top Parsers CPU usage, this displays top 5 parsers using higher cpu usages.
  - The parser which shows CPU usage when Session Pool `queue.sessions.total =1` would be the one got stuck or looping.
  - Disable the parser and monitor for drops . Notify parser Content owner to get the fix.

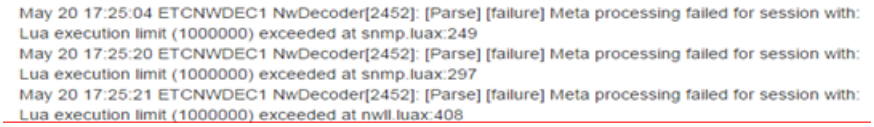
Symptom: Lua Parser Warnings can cause packet drop issues

A lua parser which throws warnings and errors can affect the parse thread and its cpu cycles, hence if there is any lua parser which is throwing warnings then it needs to be fixed.

Resolution:

- Identify and disable the lua parser throwing errors or warnings and monitor decoder.
- Create a content ticket so that content owners can provide a fix or workaround.

ex: Lua Parser throwing parse failure for execution limits exceeding.



Decoder and Log Decoder References

This is a collection of references, which provide information about the user interface for Decoders and Log Decoders in NetWitness, with references to the procedures that describe the work you can do in that part of the user interface.

Topics

- Services Config View - Capture Policies Tab
- Services Config View - Edit Capture Policies Wizard
- Services Config View - Data Privacy Tab
- Services Config View - Data Retention Scheduler
- Services Config View - Feeds Tab
- Services Config View - Files Tab
- Services Config View - General Tab
- Services Config View - Parsers Tab
- Services Config View - Parser Mappings Tab
- Services Config View - Rules Tabs
- Services System View - Decoders

Services Config View - Capture Policies Tab

In the Capture Policies tab ((Configure) > **Capture Policies**, administrators can configure selective network data collection policies.

Selective network data collection gives administrators the ability to apply centrally managed capture policies across their Network Decoders. This results in better use of Decoder resources, including hard drive space, which leads to more predictable costs and lessens the burden of managing multiple services. You can determine which traffic is stored and how it is stored by using policies. Each policy contains a list of supported base protocols and definitions for handling any other protocols that are detected. The administrator can choose to deploy predefined policies that capture:

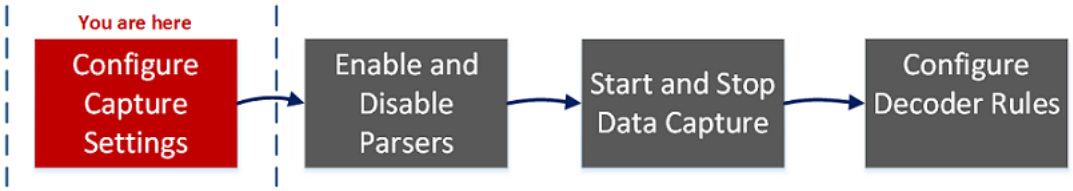
- All base and other protocols (Full Capture - All Protocols)
- Only metadata on all base and other protocols (Capture Meta Only - All Protocols)
- Only metadata on all base protocols and drop all other protocols (Capture Meta on Base Protocols, Drop all other protocols)
- All base protocols and only metadata on all other protocols (Full Capture on Base Protocols, Meta only on all other protocols)

The predefined policies are not configurable. The only way you can edit predefined policies is by assigning services (such as Decoders) to deploy them.

Administrators can create custom policies to give further control over the deployment. A base set of protocols is available for alterations by the administrator, allowing you to choose what level of capture you prefer on a per-protocol basis. If you are only making slight changes, a good start for customization is to clone one of the predefined polices and alter it. These centrally managed policies are then applied to services (Network Decoders) to allow handling multiple use cases across your environment.

Workflow

The following figure depicts the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	use predefined selective collection policies*	(Optional) Configure Selective Network Data Collection
Administrator	create new policies from predefined ones*	(Optional) Configure Selective Network Data Collection

User Role	I want to...	Documentation
Administrator	create custom policies*	(Optional) Configure Selective Network Data Collection
Administrator	verify policies	(Optional) Configure Selective Network Data Collection
Administrator	unpublish policies*	(Optional) Configure Selective Network Data Collection
Administrator	delete policies*	(Optional) Configure Selective Network Data Collection

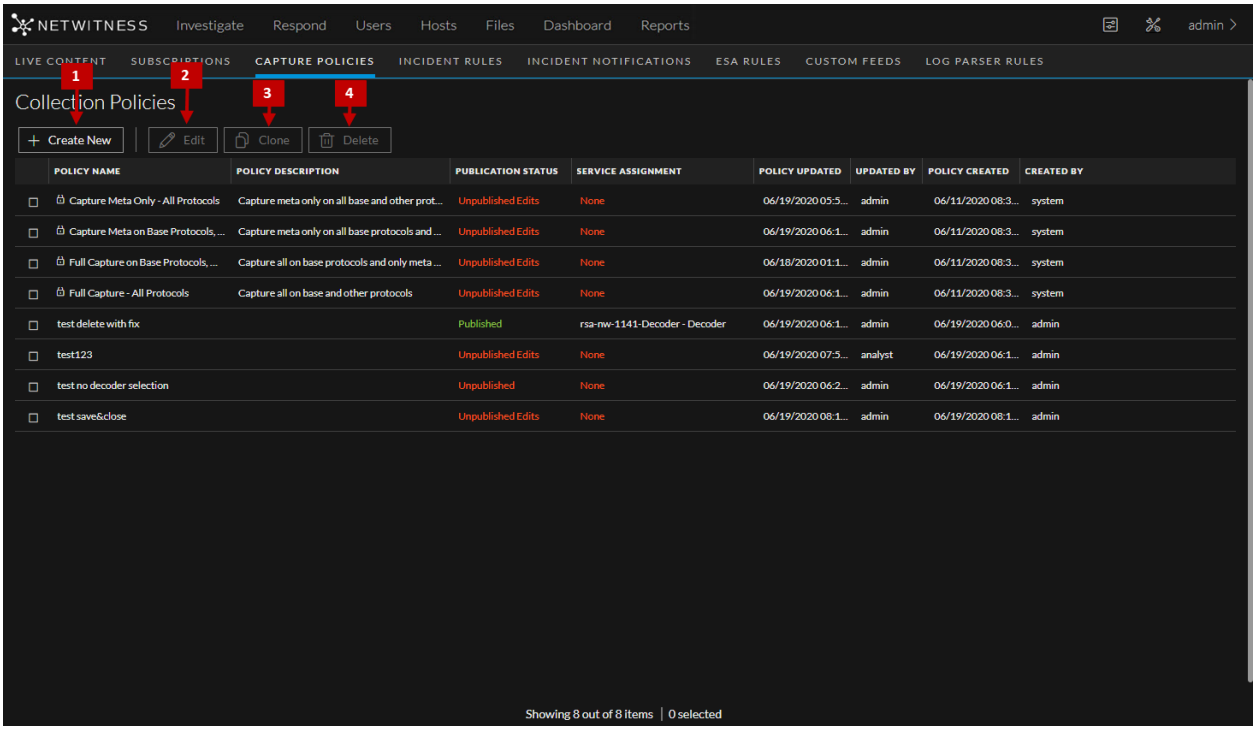
*You can complete these tasks here.

Related Topics

- Services Config View - Edit Capture Policies Wizard

Quick Look

This is an example of the Capture Policies tab.



1 Create New - Provides a wizard to create a new policy.
2 Edit - Provides a wizard to edit an existing custom policy.
3 Clone - Selects an existing policy (predefined or custom) and copies it to create a new policy.
4 Delete - Removes policies.

The following table describes the contents of the Capture Policies tab.

Column	Description
Policy Name	The name of the policy (not editable for predefined policies).
Policy Description	The description of the policy (not editable for predefined policies).
Publication Status	• **Unpublished**: policies that have not been deployed to Decoders yet. • **Unpublished Edits**: policies that have been updated but are not deployed to Decoders. Clicking **Save and Close** after making edits keeps them from automatically being published, allowing the administrator to wait until a specific outage window to publish them. This would display as **Unpublished** (for policies not previously published to Decoders) or **Unpublished Edits** (for policies that have been previously published to Decoders, but the updates have not been deployed to Decoders). • **Published**: policies which have been deployed to Decoders. • **Failed**: policies that failed to execute, for example, if a Decoder is offline or a system is down. For information about how to troubleshoot failed policies, see "Troubleshooting Policy Deployments That Fail" in (Optional) Configure Selective Network Data Collection.
Service Assignment	Identity of the Decoder or service on which the policy is run.
Policy Updated	Date the policy was edited.
Updated By	Role of the person who updated the policy.
Policy Created	Date on which the policy was created.
Created By	Role of the person who created the policy.

Services Config View - Edit Capture Policies Wizard

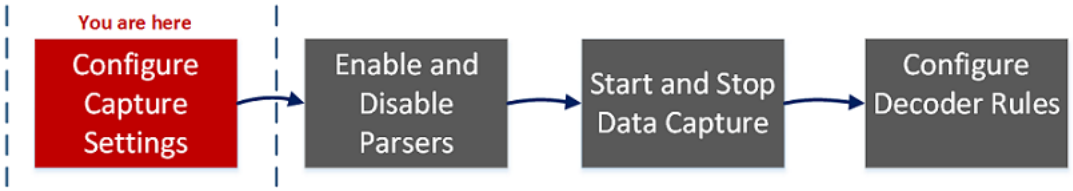
In the Edit Policies view ((Configure) > **Capture Policies** > **Edit**, administrators can configure selective network data collection policies. The panels for creating a new policy and editing an existing policy are the same except that in the editing wizard, the definitions of the existing policy are displayed.

You can edit existing policies to customize them for your environment. A base set of protocols are available for alterations, allowing you to choose what level of capture you prefer on a per-protocol basis. If you are only making slight changes, a good start for customization is to clone one of the predefined polices and alter it.

Note: A set of predefined policies are available immediately. The only way you can edit a predefined policy is by assigning Decoders to the policy to collect the targeted data.

Workflow

The following figure shows the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	create new policies from predefined ones*	(Optional) Configure Selective Network Data Collection
Administrator	create custom policies*	(Optional) Configure Selective Network Data Collection
Administrator	unpublish policies*	(Optional) Configure Selective Network Data Collection
Administrator	delete policies*	(Optional) Configure Selective Network Data Collection

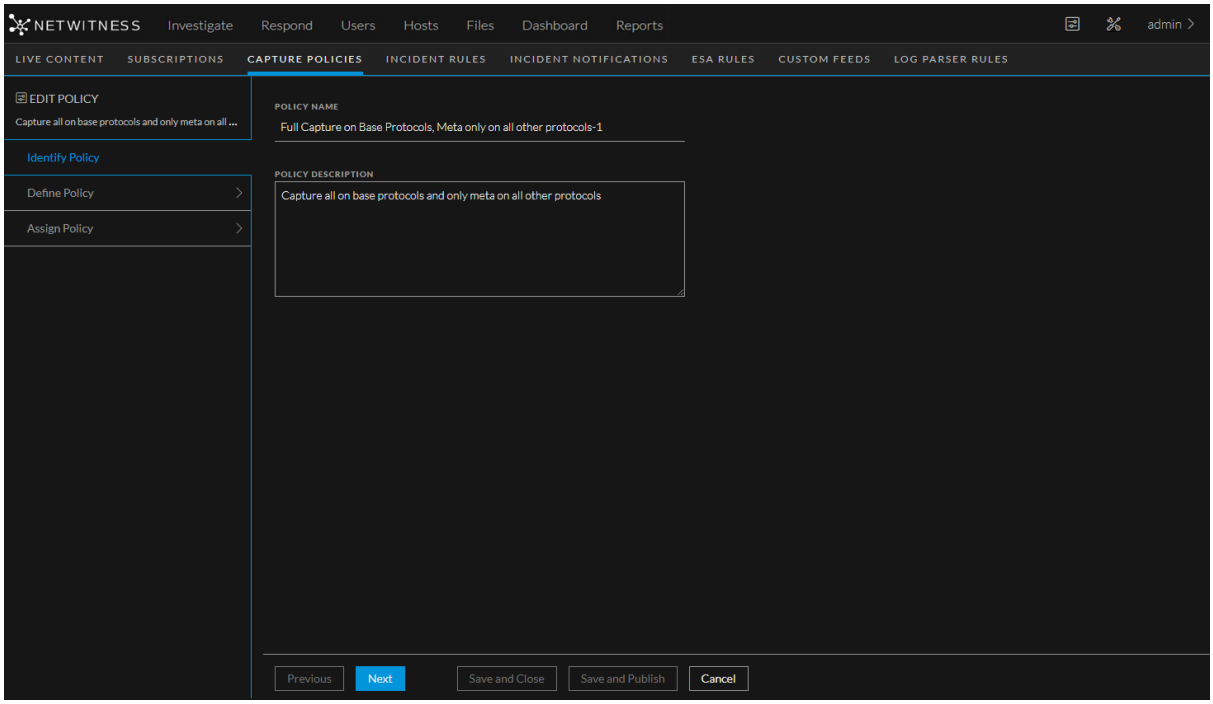
*You can complete these tasks here.

Related Topics

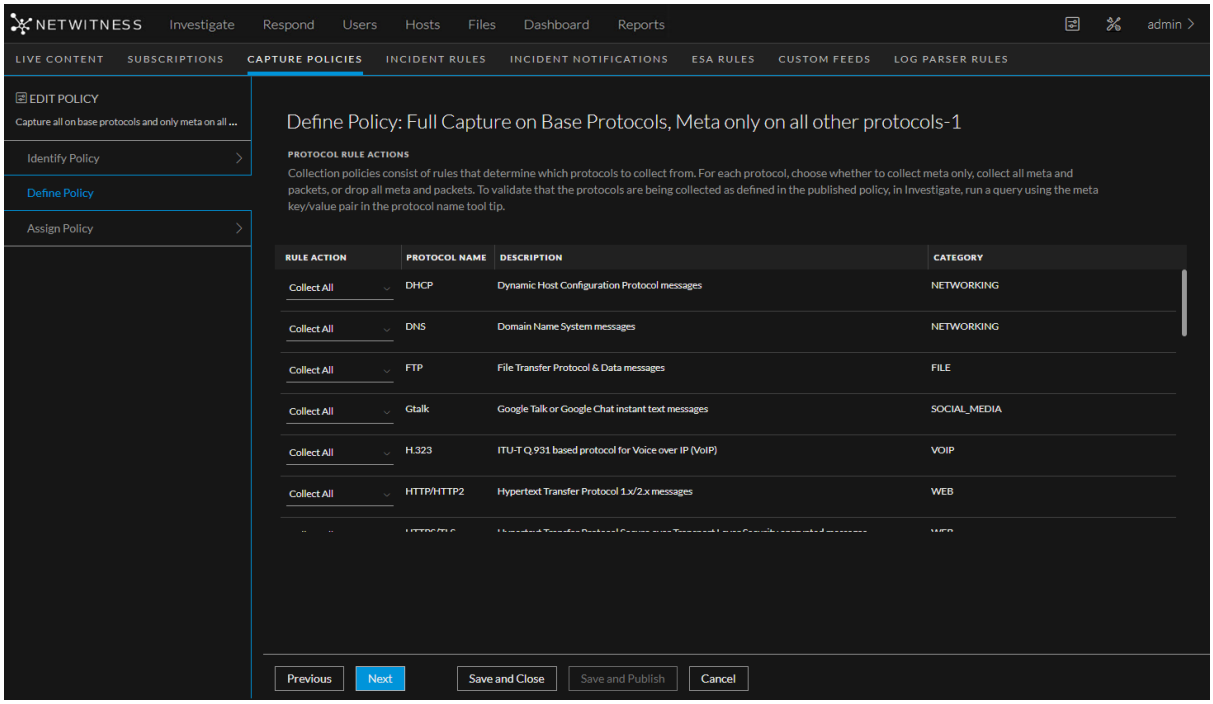
- Services Config View - Capture Policies Tab

Quick Look

1. To access the Edit Policy wizard, go to (Configure) > **Capture Policies**, select a policy, and click **Edit**.
   The Identify Policy page is displayed, where you can change the name and description of the policy.



2. To view the Define Policy page, click **Next**.



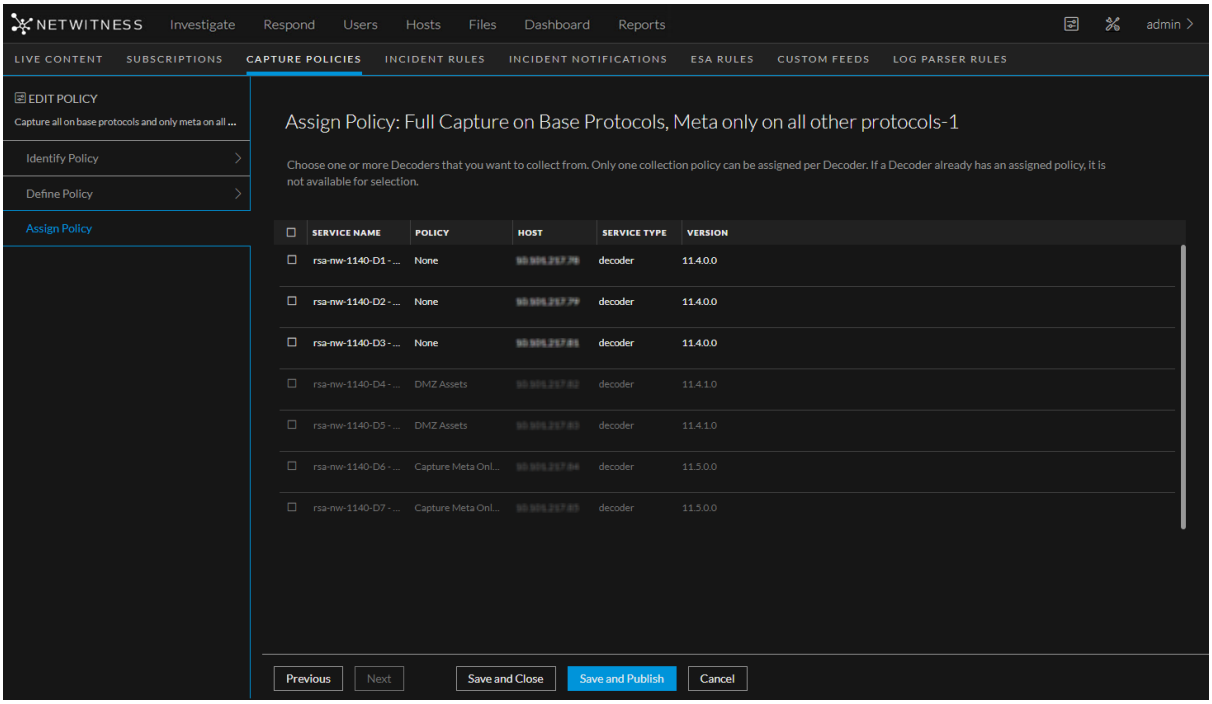
In the Define Policy page, you can define rule actions for the protocols for which you are interested in collecting data. The rule actions are:

- **Collect Meta Only**: Collect metadata
- **Drop All**: Drop metadata and network packets
- **Collect All**: Collect metadata and network packets

You can click **Save and Close** to save your policy without deploying it on Decoders. The following table describes the actions of the buttons on this page.

Button	Description
Previous	Go to the previous page of the wizard
Next	Keep updates from the current page and go to the next page of the wizard.
Save and Close	Save the policy definitions without publishing the policy.
Save and Publish	Save the policy definitions and publish the policy to Decoders to begin data collection.
Cancel	Cancel the definitions and return to the Capture Policies tab.

3. Click Next to go to the Assign Policy page.



This is where you assign the policy to Decoders to capture the data. The following table describes the columns in this page.

Column	Description
Service Name	The name of a service. For 11.5, Decoder is the only available service.
Policy	The policy that a Decoder is assigned to.
Host	The IP address or name of the host system for the Decoder.
Service Type	The type of service the policy is published to. For 11.5, the service type is **decoder**.
Version	The version of NetWitness Platform that is installed on the service host.

Services Config View - Data Privacy Tab

In the Data Privacy tab (**(Admin) > Services > Select a Decoder or Log Decoder >** > **Config > Data Privacy tab)**, administrators can configure data privacy parameters for certain Core services. For the Decoder and Log Decoder, you can set the default hash algorithm and salt.

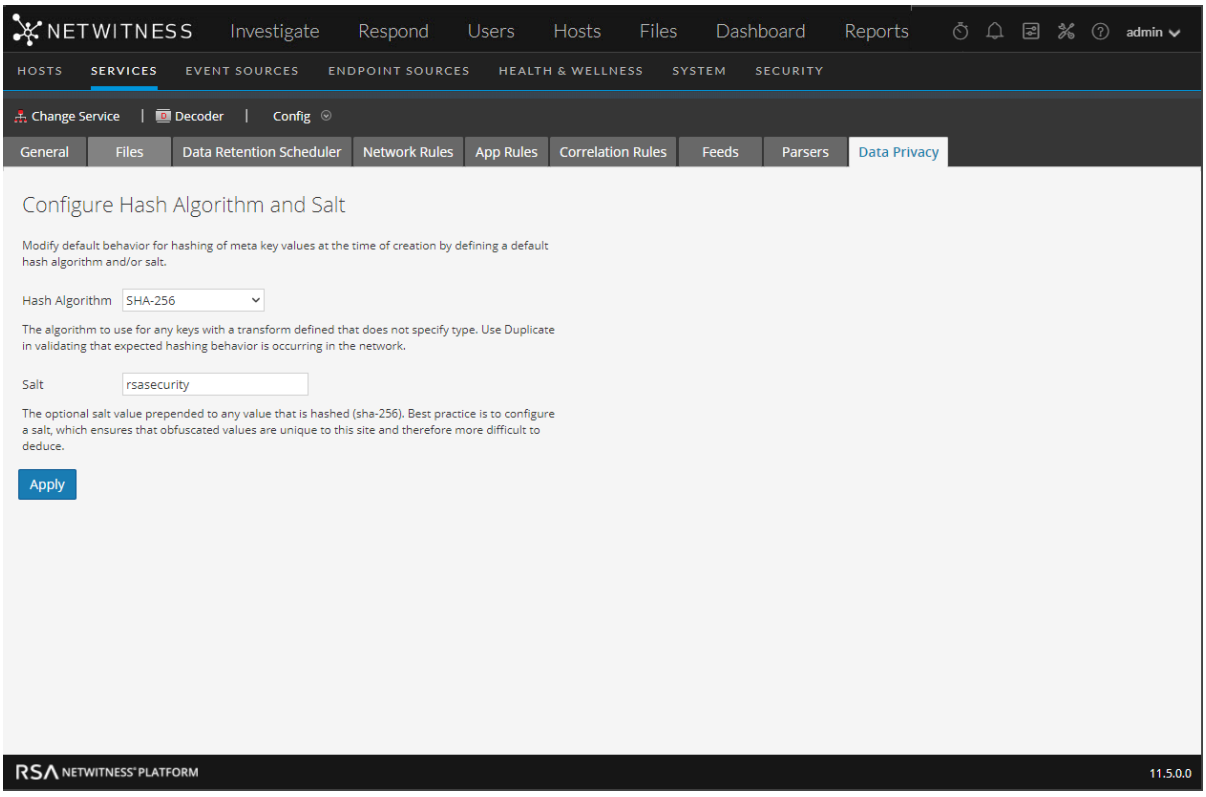
What do you want to do?

User Role	I want to ...	Documentation
Administrator	configure hash algorithm and salt	"Configure the Hash Algorithm and Salt" in the *Data Privacy Management Guide.* (Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder

Quick Look



The Data Privacy tab has the Configure Hash Algorithm and Salt configuration settings. The following table describes the parameters in this tab.

Parameter	Description
Hash Algorithm	Displays a drop-down list of hash algorithms to use for any keys with a transform that does not specify algorithm type. Possible values are SHA-256 and Duplicate. Duplicate is a special algorithm available for administrators to use when validating that expected hashing behavior is occurring in the network.
Salt	Indicates the optional salt value prepended to any value that is hashed. Best practices for security purposes dictate a salt value that is no less than 100 bits or 16 characters in length. Configuring a value ensures that obfuscated values are unique to this site and therefore more difficult to deduce. For more information on this field, see "Configure Data Obfuscation" in the *Data Privacy Management* guide.
Apply	Applies any changes.

Services Config View - Data Retention Scheduler

In the Data Retention Scheduler tab, you can set the rollover criteria for removing database records from primary storage using an age-based threshold. You can also schedule the timing to check whether the threshold is reached.

To access the Data Retention Scheduler tab, go to (Admin) > **Services** > select a **Decoder** or **Log Decoder** service and click > **View** > **Config** > **Data Retention** tab.

What do you want to do?

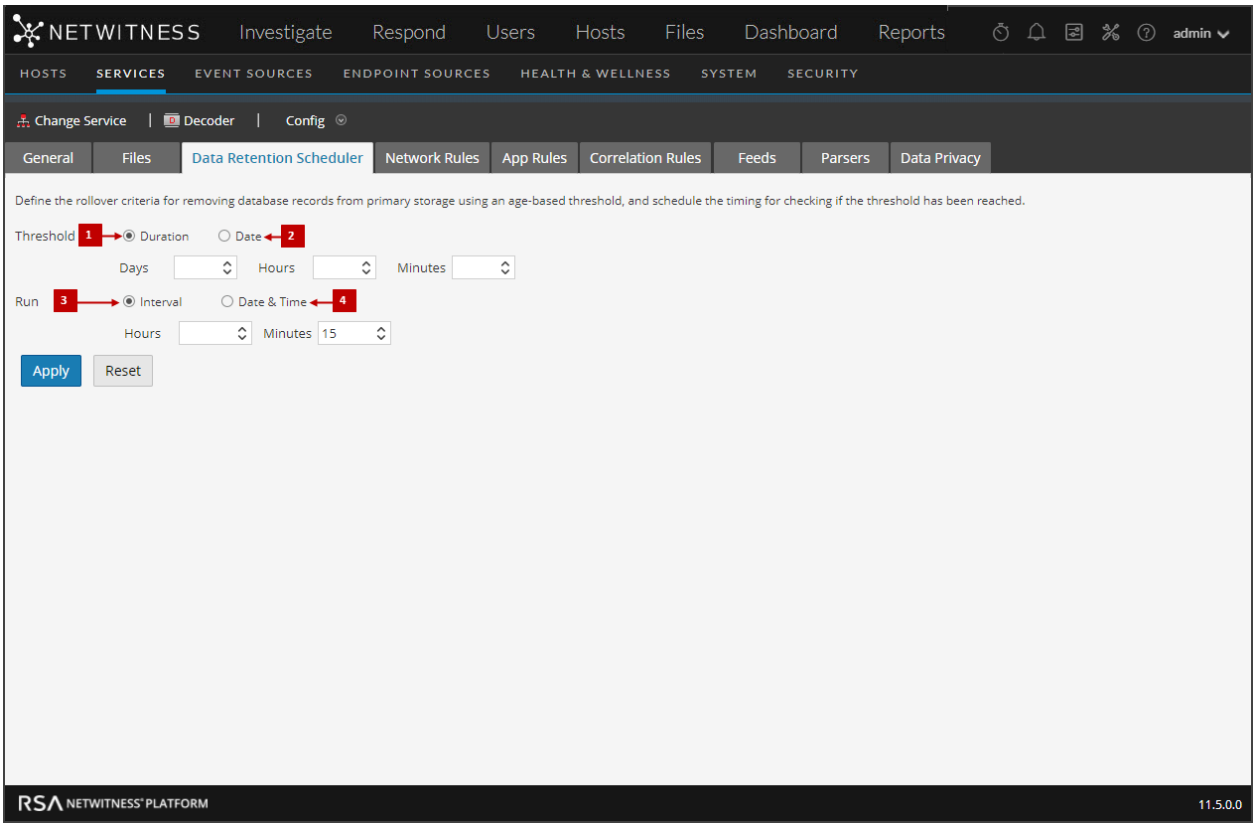
User Role	I want to...	Documentation
Administrator	Schedule the timing to see if the threshold is reached.	Configure Transaction Handling on a Decoder

Related Topics

- Configure Common Settings on a Decoder
- Decoder and Log Decoder Quick Setup

Quick Look

This is an example of the Data Retention Scheduler tab.



1 **Threshold Duration**: Removes database files older than the selected number of days, minutes, or hours.

2 **Threshold Date**: Removes database files older than the selected UTC date (YYYY-MM-DD-HH:MM:SS) that are not compatible with minutes, hours, or days parameters.

3 **Run Interval**: Indicates the number of hours between executions.

4 **Run Date and Time**: Defines which days of the week to execute the scheduler, as well as time of execution in HH:MM:SS format for the local time of the service.

Services Config View - Feeds Tab

Feeds and parsers are Lua programs loaded and compiled when either processing capture files in NetWitness Investigate or capturing data with Decoders. Most commonly, they are used for static meta extraction and service identification.

Note: Pre-11.0 versions of NetWitness used FLEXPARSE programs in addition to Lua programs; Flexparsers are deprecated in NetWitness 11.0. Unless otherwise stated, any reference to Decoders applies to Log Decoders as well.

NetWitness uses feeds to create metadata based on externally defined meta values. A feed is a list of data that is compared to sessions as they are captured or processed. For each match, additional metadata is created. This data can identify and classify malicious IPs or incorporate additional information such as department and location based on internal network assignments. Some examples of feeds include threat feeds to identify BOTNets, DHCP mappings, or even active directory information such as physical location or logical department.

Feeds can be added, removed, and updated while a Decoder is running without affecting capture. The Feeds tab (**(Admin) > Services >** select a Decoder or Log Decoder service and click **> View > Config > Feeds tab**) provides a user interface for managing feeds on Decoders.

What do you want to do?

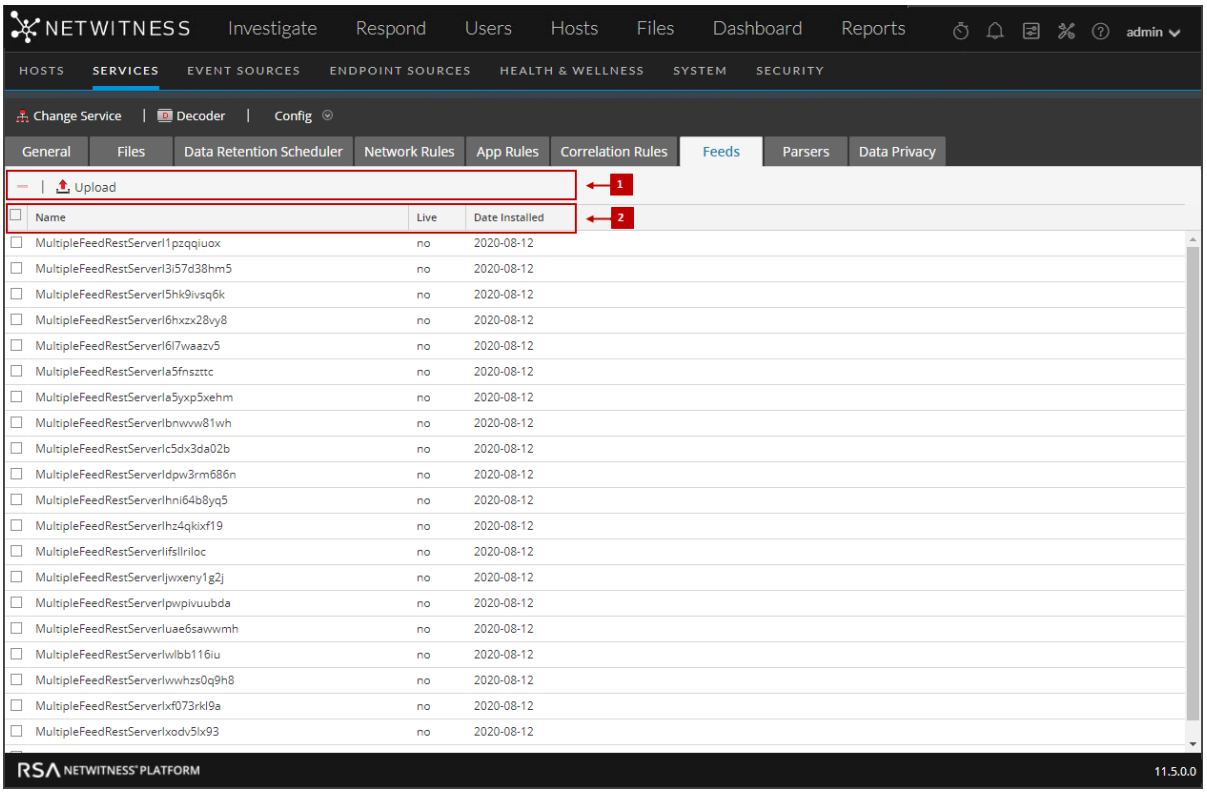
User Role	I want to...	Documentation
Administrator	configure feeds	Configure Parsers and Feeds
Administrator	enable and disable parsers	Enable and Disable Parsers and Log Parsers

Related Topics

- Configure Common Settings on a Decoder
- Decoder and Log Decoder Quick Setup
- Upload Feeds Dialog

Quick Look

This is an example of the Feeds tab.



1 Feeds Tab Toolbar - Provides options to work with feeds in the grid
2 Feed List - Lists all feeds that are currently deployed on the Decoder

Feeds Tab Toolbar

Feature	Description
Upload	Displays the Upload Feeds dialog.
−	Deletes the selected feeds.

Feeds List

The Feeds list provides a listing of all currently deployed feeds for the Decoder.

Column	Description
Name	The name of the feed or the feed file.

Column	Description
Live	Indicates if the feed originated from Live. Possible values are **Yes**, **No**, or **N/A**. • **Yes** = Installed through Live • **No** = Installed through NetWitness • **N/A** = The feed has no attributes file created by NetWitness to track the installation date. The feed may have been installed manually, not through NetWitness or Live Services. Manually installed feeds still function properly.
Date Installed	The date the feed was pushed to the service.

Upload Feeds Dialog

This topic describes the features of the Upload Feeds dialog in the Services Config view > Feeds tab.

The **Upload** option in the Services Config view > Feeds tab displays the Upload Feeds Dialog, in which you can manage the uploading of feeds to a Decoder or Log Decoder.

What do you want to do?

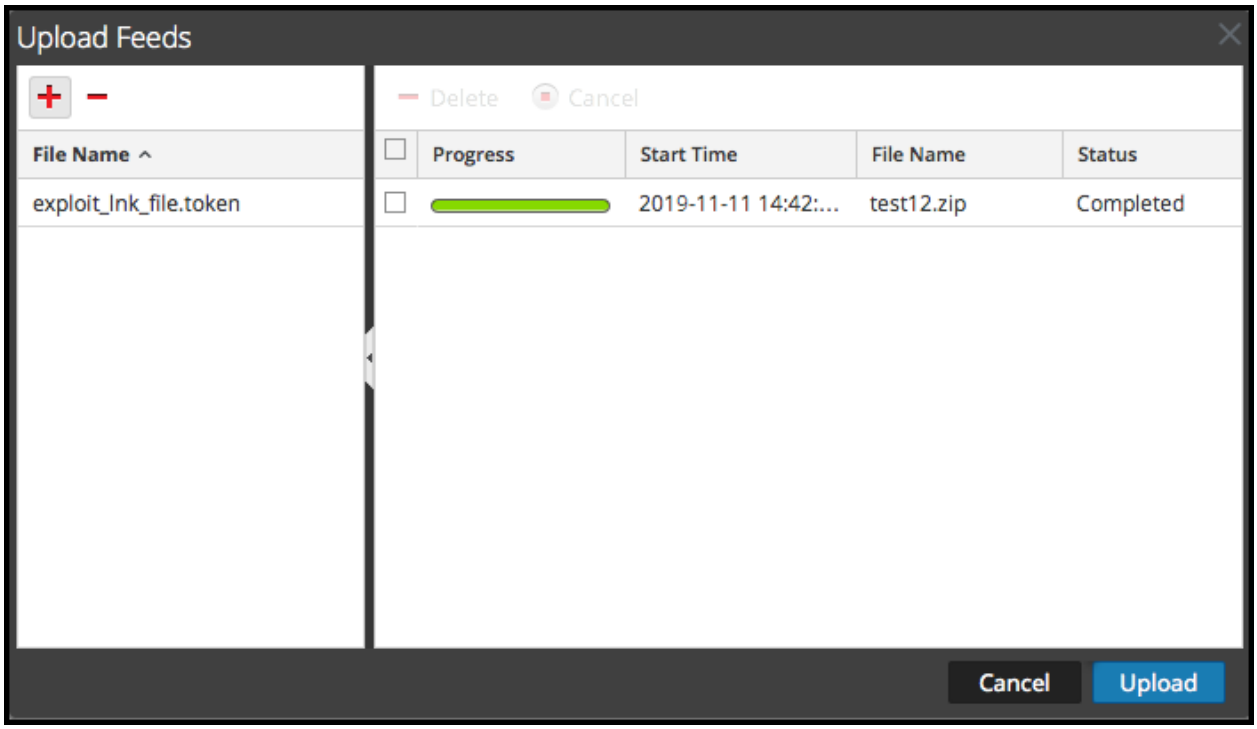
User Role	I want to...	Documentation
Administrator	prepare a list of feeds for upload	Edit, Upload, or Remove a Feed
Administrator	view and delete upload jobs	Edit, Upload, or Remove a Feed

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder

Quick Look

This is an example of the Upload Feeds dialog.



File List

The File list is the place to prepare a list of feeds for uploading. You can add files from a directory structure, and delete files from the list if you decide that you don't want to upload a particular file. When the list is ready, clicking **Upload** starts the upload process.

Feature	Description
	Opens a view of the directory structure where you can select files to add to the File list.
	Deletes the selected files from the File list.
File Name	Lists the feed files you have added from a file system in preparation for uploading to a Decoder. When you click **Upload**, the files listed here are uploaded.

Upload Job List

The Upload Job list provides a view of upload jobs started by clicking **Upload**.

Feature/Column	Description
Delete	Deletes an upload job.
Progress	Displays progress of an upload job.
Start Time	Displays the start time of an upload job.
File Name	Lists filename of the feed being uploaded.
Status	Displays the status of upload job.

Upload Feeds Dialog Buttons

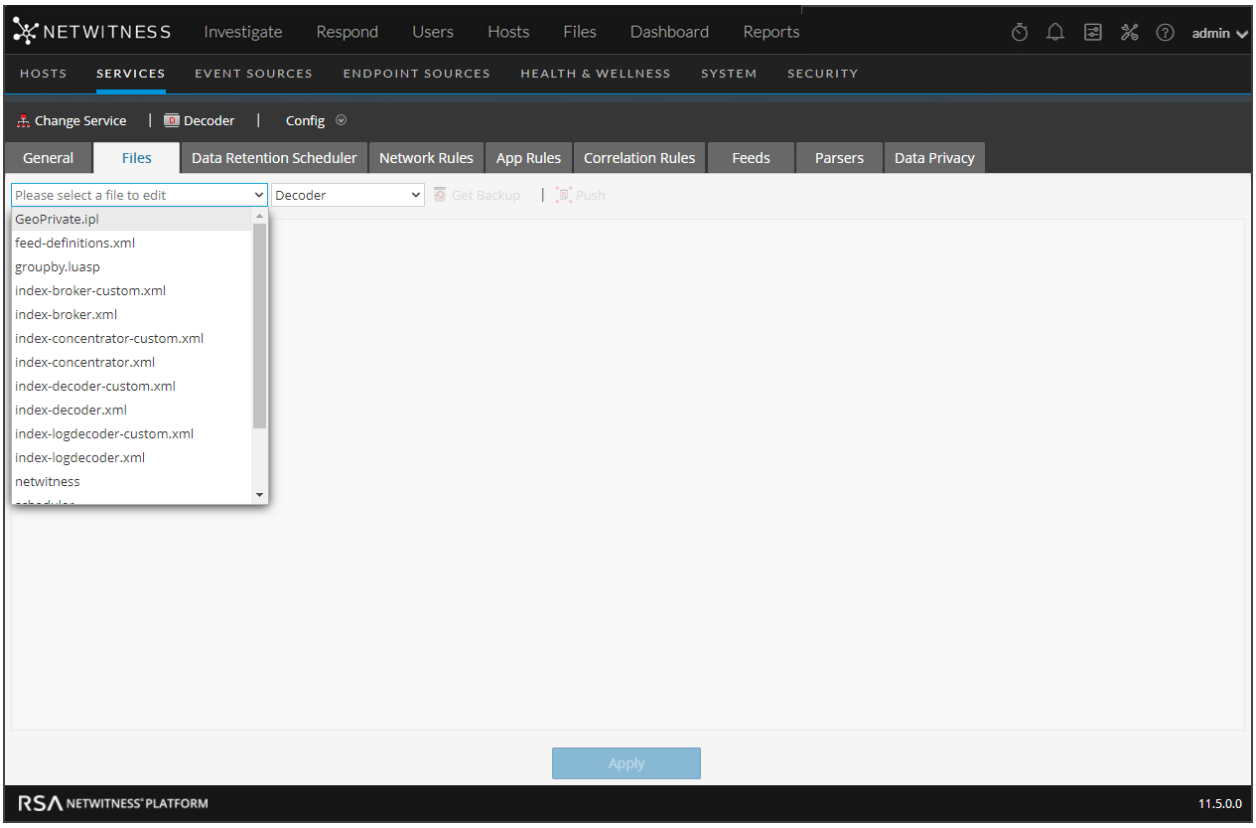
Feature	Description
Cancel	Closes the Upload Feed dialog.
Upload	Starts uploading the feed files listed in the File list. Each feed is listed in a separate row in the Upload job list.

Services Config View - Files Tab

The Decoder and Log Decoder configuration files are visible and editable in the (Admin) > Services > Config view > Files tab. "Edit Core Services Configuration Files" in the *Hosts and Services Getting Started Guide* provides general instructions for editing files. (Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

Like other Core services, both the Decoder and Log Decoder have an index file, and may also have a crashreporter, netwitness, and scheduler. The Decoder and Log Decoder index files are named **`index-decoder-custom.xml`** and **`index-logdecoder-custom.xml`.**

Note: `Table-map.xml` and `table-map-custom.xml` are available only for Log Decoders with log content installed.



What do you want to do?

User Role	I want to...	Documentation
Administrator	obtain log files from pre-11.0 Log Decoder	Obtain Log Files from a Pre-11.0 Log Decoder

User Role	I want to...	Documentation
Administrator	edit files and parsers	• Feed Definitions File • Flex Parsers • GeoIP2 Parsers • Lua Parsers • Decoder Snort Detection • Search Parser • Wireless LAN Configuration

Related Topics

- Configure Common Settings on a Decoder
- Decoder and Log Decoder Quick Setup
- Create Custom Meta Keys Using a Custom Feed

Quick Look

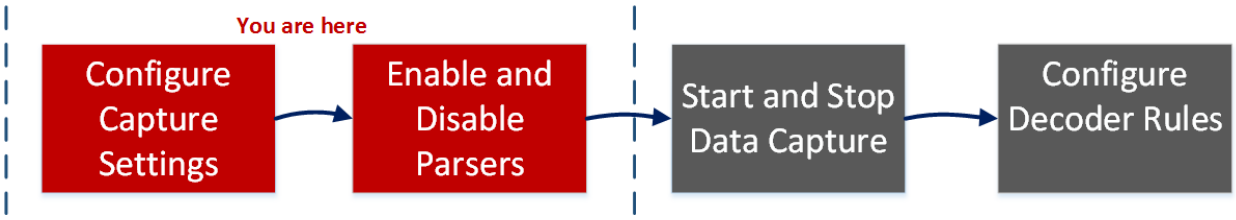
Filename	Description
`GeoPrivate.ipl`	This fixed parser takes the IP addresses and converts them to geographical locations. The locations are displayed through the Google Earth display.
`feed-definitions.xml`	Used to create custom feeds, this is the XML schema used by the Decoder to define a **feed** message when it creates a **.feed** file.
`traffic_flow_options.lua`	Used to provide directionality information. Update this file with environment-specific internal and external subnets for the Lua parser to create proper directionality in metadata. The parser is described in RSA Content for RSA NetWitness Platform.
`search.ini`	This is the Search Parser configuration file. The Search Parser is a custom parser, used to generate metadata by scanning for pre-defined keywords and regular expressions.
`wlan-config.xml`	This is the wireless LAN configuration file (9/9/2009). This file controls the 802.11 parsers. Its chief purpose is to control decryption of raw 802.11 frames captured by the Decoder.

Services Config View - General Tab

The General tab for a Decoder in the Services Config view provides a way to manage basic service configuration, configure data capture, and select the parsers that are applied to the captured data. To access the General tab, go to **Admin > Services >** select a Decoder or Log Decoder and click **> View > Config > General tab**.

Workflow

The following figure depicts common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	configure capture settings*	Configure Capture Settings
Administrator	manage parsers and log parsers*	Enable and Disable Parsers and Log Parsers
Administrator	start and stop data capture	Start and Stop Data Capture
Administrator	configure rules	Configure Decoder Rules

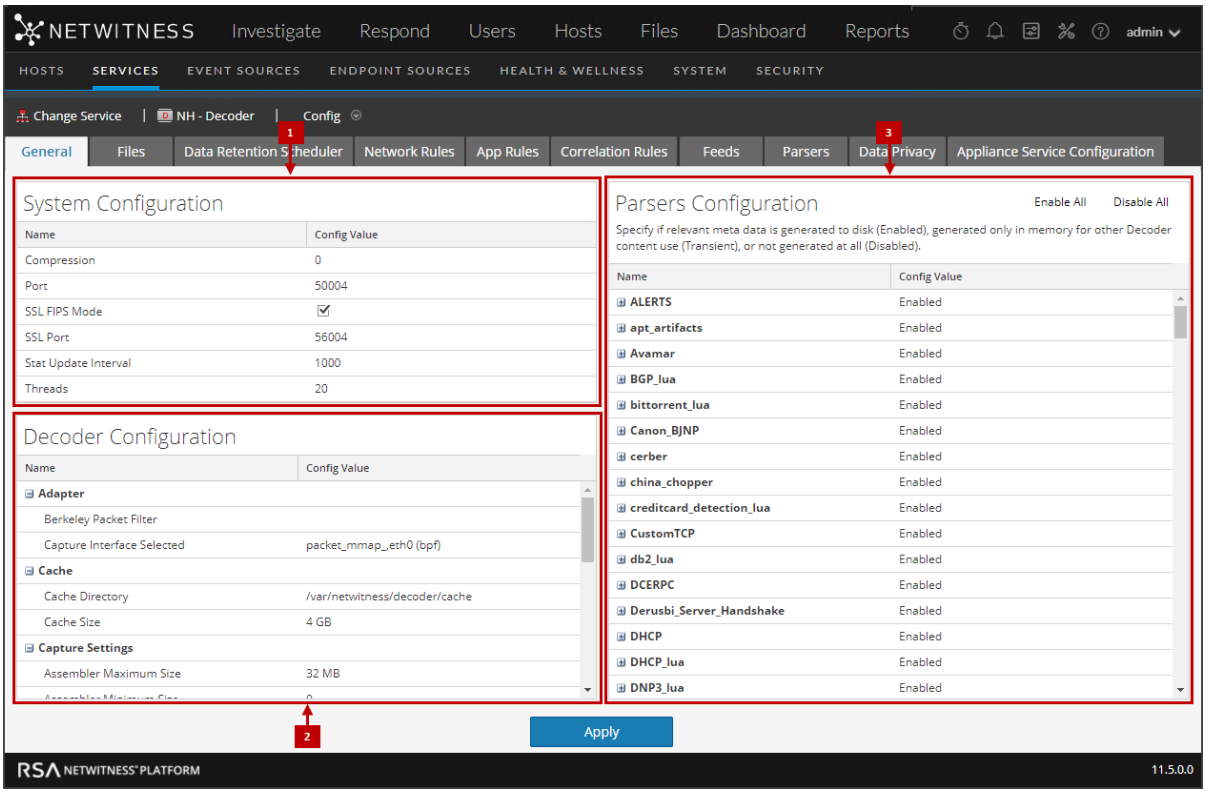
*You can complete these tasks here.

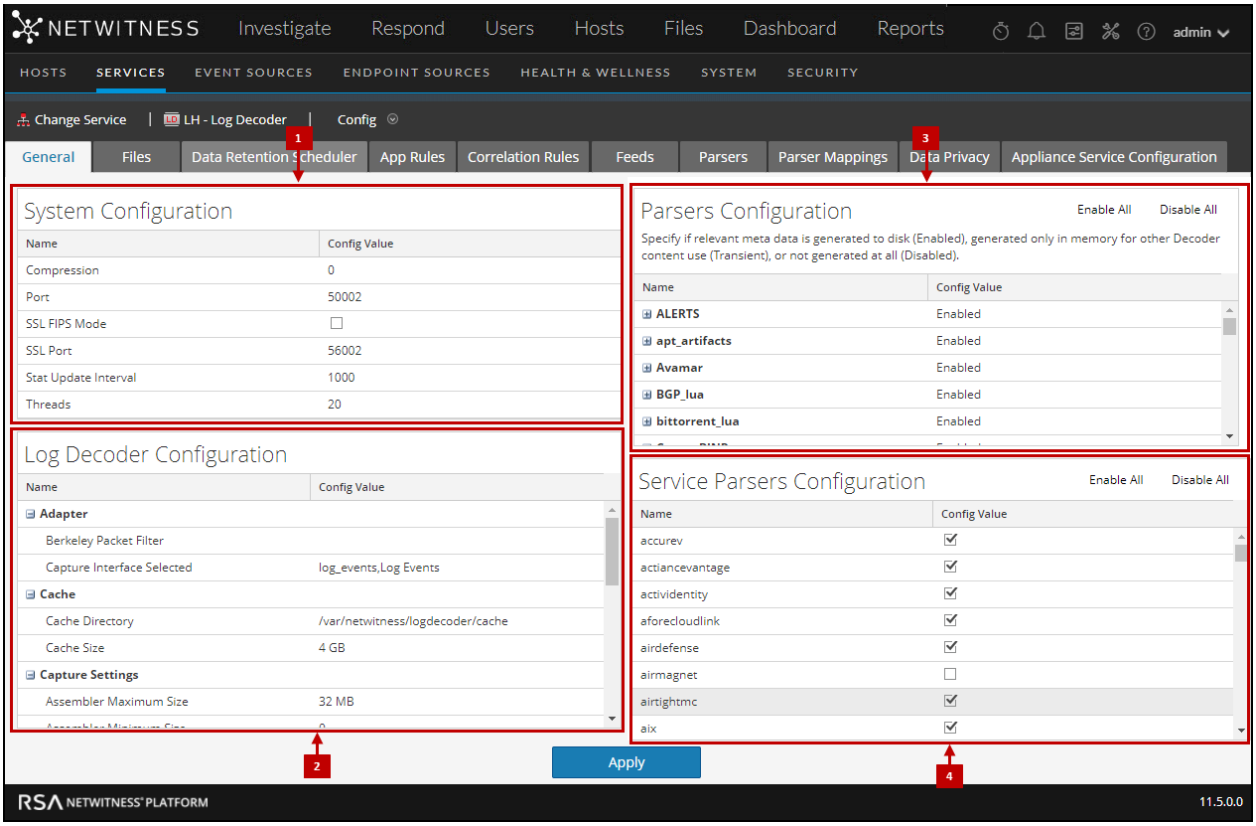
Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- Configure Parsers and Feeds

Quick Look

The first figure is an example of the General tab for a Decoder. The second is the General tab for a Log Decoder.



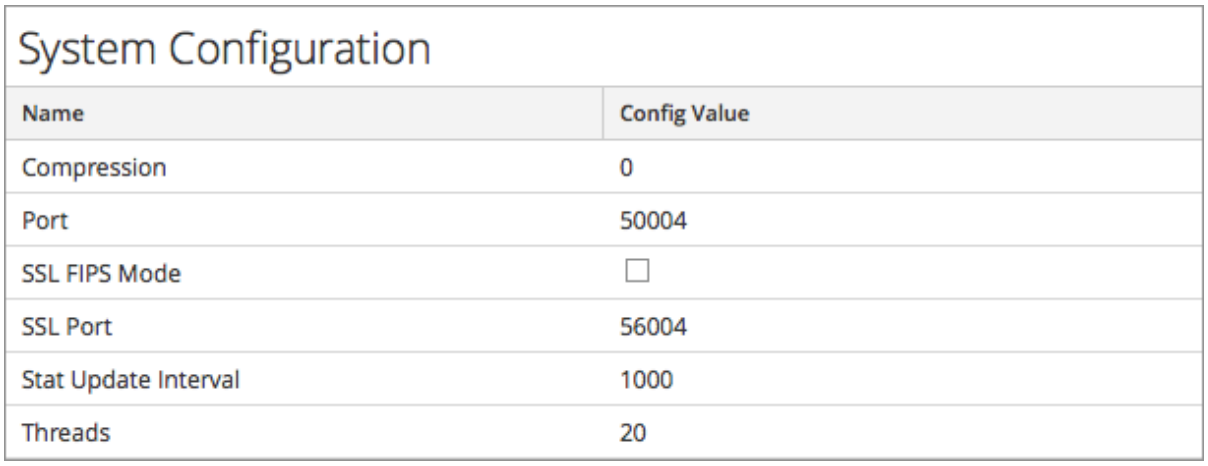


1 System Configuration - Manages service configuration for a Decoder or a Log Decoder.

2 Decoder Configuration or Log Decoder Configuration - Lets you view and edit service configuration parameters for a Decoder or Log Decoder.
3 Parsers Configuration - Lets you select parsers to use on the Decoder or Log Decoder.
4 Service Parsers Configuration (Log Decoders only) - Lets you select service parsers to use on the Log Decoder.

System Configuration Section

The System Configuration section manages service configuration for a Decoder. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

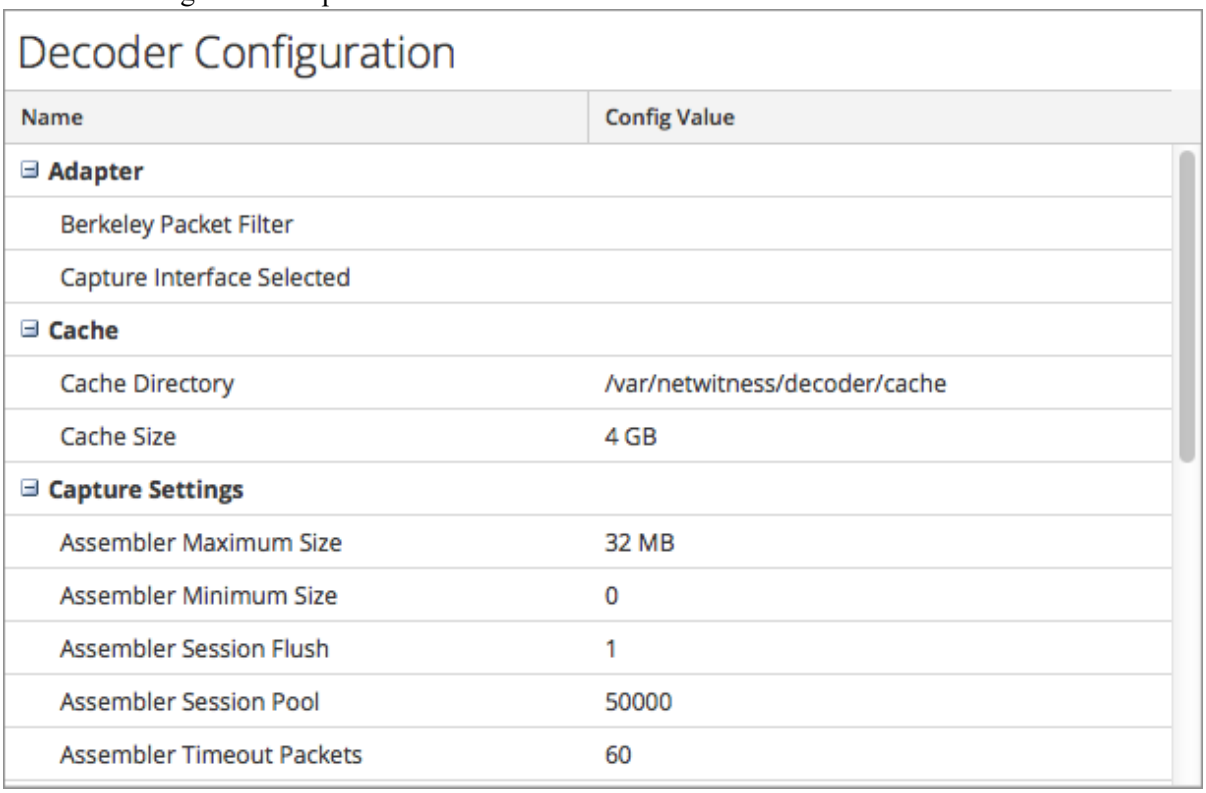


The System Configuration section has these parameters.

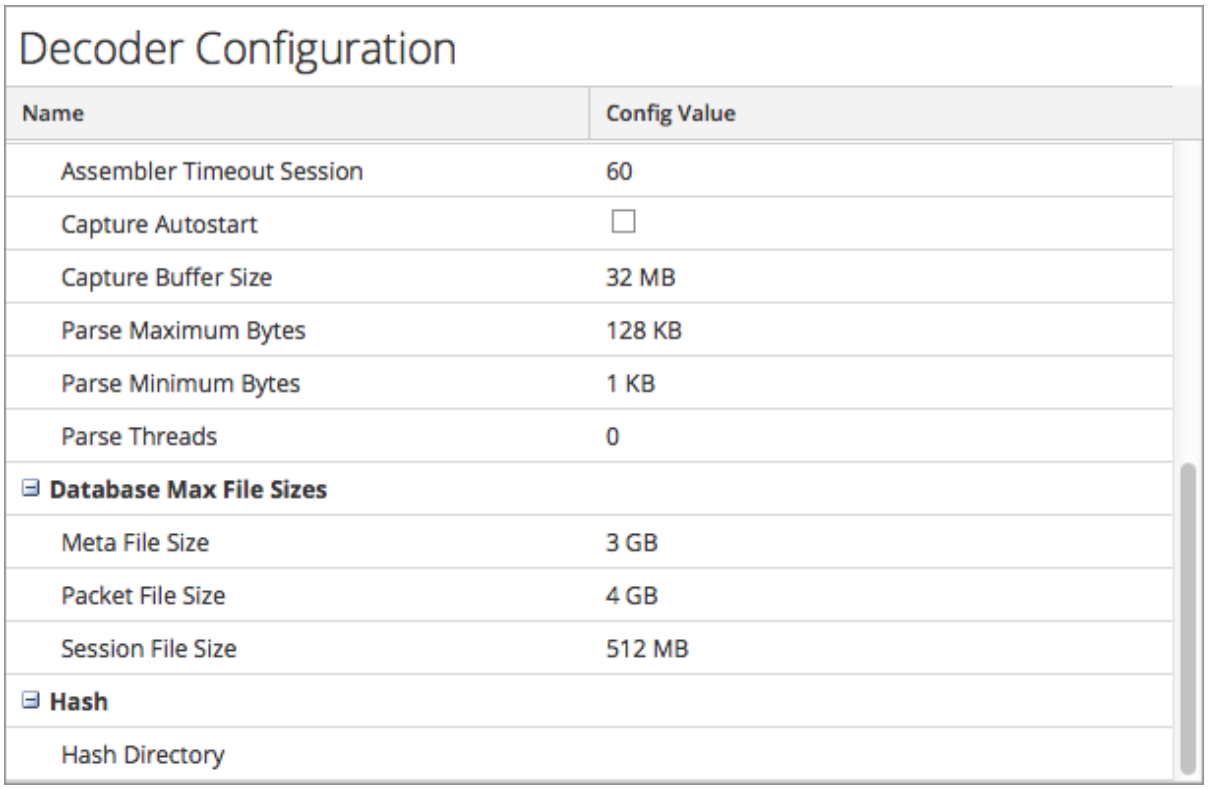
Parameter	Description
Compression	The minimum number of bytes that must be transmitted per response before compression. A setting of 0 disables compression. The default value is **0**. A change in value is effective immediately for all subsequent connections.
Port	Determines the port used by the service. **Note:** If you change the port number, ensure that you restart the service.
SSL FIPS mode	If enabled, all the data transferred in the network will be encrypted using SSL.
SSL Port	Indicates the port used for encrypting using SSL.
Stat Update Interval	The number of milliseconds between statistic updates on the system. Lower numbers cause more frequent updates and can slow down other processes. The default value is **1000**. A change in value is effective immediately.
Threads	The number of threads in the thread pool to handle incoming requests. A setting of **0** lets the system decide. A change takes effect on service restart.

Decoder Configuration Section

The Decoder Configuration section provides a way to view and edit service configuration parameters for a Decoder or Log Decoder. When a service is first added, default values are in effect. You can edit these values to manage traffic capture.



Scrolling to the bottom of the section reveals these additional Decoder Configuration parameters.



Adapter Section

Adapter parameters configure the network interface for capture as described in Configure Capture Settings.

Cache Section

Cache parameters configure the cache directory and size for session cache files. The following table describes the cache settings. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

Cache Parameter	Description
Cache Directory	The directory where session cache files are stored. The default value is `/var/netwitness/decoder/cache`. Change takes effect immediately.
Cache Size	The maximum size, in Megabytes (MB), that all files in the cache directory can attain before the oldest files are deleted. Once the threshold is reached, the cache size is reduced by 10%. The default value is **4 GB**. Change takes effect immediately.

Capture Settings Section

The Capture Settings section provides a way to configure operational capture settings. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

Capture Settings Parameter	Description
Assembler Maximum Size	Specifies the maximum size in bytes that a session's packet data size can attain. The default value is **32 MB**. Change takes effect immediately.
Assembler Minimum Size	Specifies the minimum size in bytes that a session must have in order to generate metadata. A value of **0** means every session has metadata generated. The default value is **0.** Change takes effect immediately.
Assembler Session Flush	Specifies whether a session is removed from the assembler when the session's last chain is removed from the assembler. The default value is **1**. • **2** = if the first packet of a session times out of assembler, the session is removed from assembler after parsing is complete. Any subsequent packets for this session create a new session in assembler. • **1** = If the last chain of a session times out of assembler, the session is removed from assembler. Any subsequent packets for this session create a new session in assembler. • **0** = If the last chain of a session times out of assembler, the session is left in assembler until it times out. Any subsequent packets for this session are filtered Change takes effect on service restart.
Assembles Session Pool	Specifies the number of entries in the session pool. The default value is **350000**. Change takes effect on service restart.
Assembler Timeout Packets	Specifies the number of seconds before a packet or chain is timed out. T default value is **60**. Change takes effect immediately.
Assembler Timeout Session	Specifies the number of seconds before a session is timed out. Default value is **60**. Change takes effect immediately.
Capture Autostart	Specifies whether capture begins automatically each time Decoder is started. When checked, the value = yes. When unchecked, the value = no. The default value is **no**. Change takes effect immediately.
Capture Buffer Size	The capture memory buffer allocation in Megabytes. Default value is **64 MB**. Change takes effect on service restart.
Parse Maximum Bytes	The maximum number of bytes to scan a stream for additional tokens. When the first token is found, the stream is scanned up to the set number of bytes, but no further. A setting of **0** removes the early termination and the full stream is scanned regardless of size. The default value is **128 KB**. Change takes effect immediately.
Parse Minimum Bytes	The minimum number of bytes to scan a stream for the first token. If no token is found within the set number of bytes, scanning is terminated. A setting of **0** removes the early termination and the full stream is scanned regardless of size. The default value is **1 KB**. Change takes effect immediately.

Capture Settings Parameter	Description
Parse Threads	The number of parse threads to use for session parsing. A value of **0** means let the server decide. The default value is **0**. Change takes effect on service restart.

Database Max File Sizes Section

The Database Max File Sizes section controls the maximum file size for various databases. When a service is first added, default values are in effect and should be changed only in special circumstances, for example, if Customer Support advises a change.

File Size Parameter	Description
Meta File Size	The maximum size of meta database files in Megabytes. The default value is **10 MB**. Change takes effect on service restart.
Packet File Size	The maximum size of packet database files in Megabytes. The default value is **10 MB**. Change takes effect on service restart.
Session File Size	The maximum size of session database files in Megabytes. The default value is **100 MB**. Change takes effect on service restart.

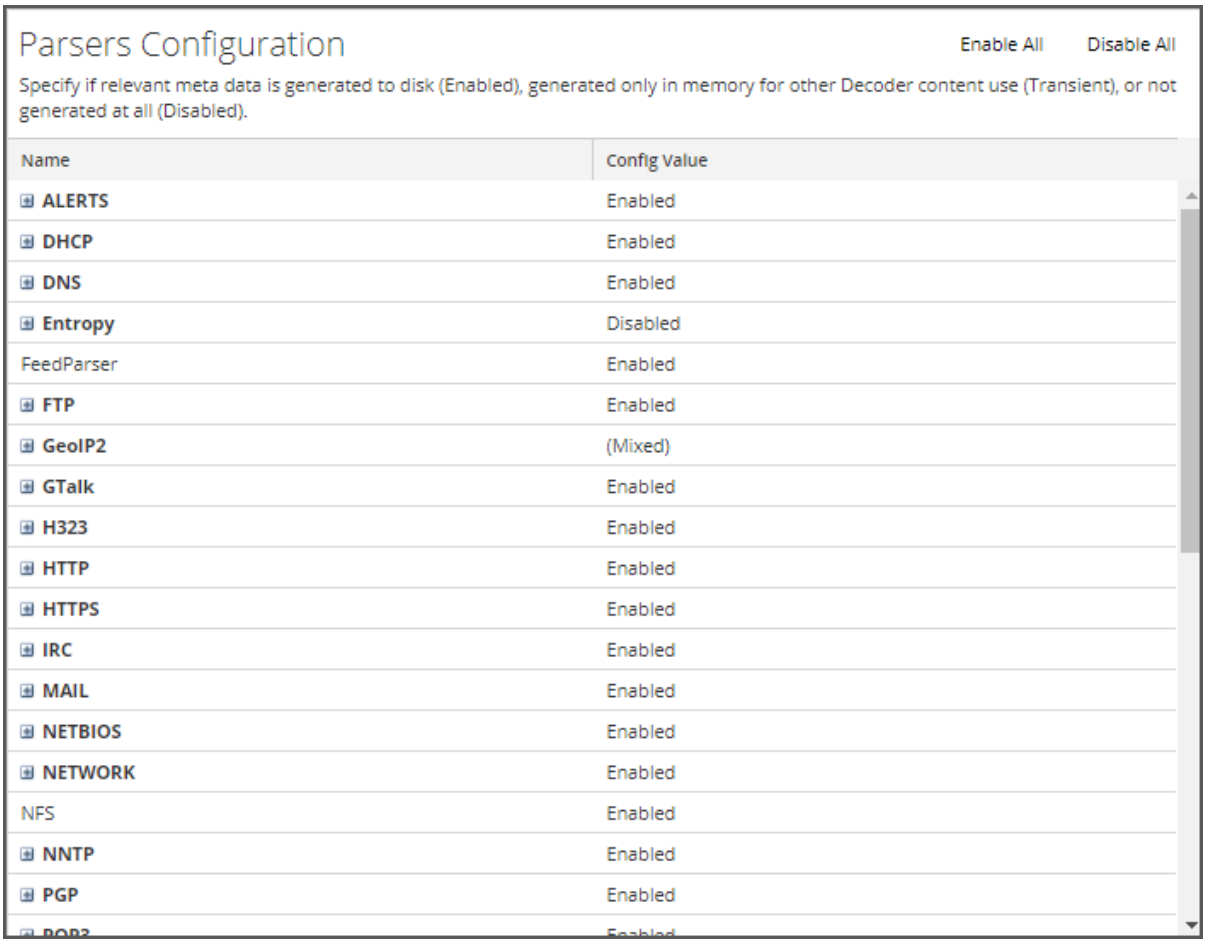
Hash Section

The Hash section settings control data base file hashing options. There is a small performance penalty when hashing.

Hash Parameter	Description
Hash Directory	The server directory where all hash files are written. If empty, each hash file is written to the same directory as the file being hashed. The default value is blank. Change takes effect on service restart.

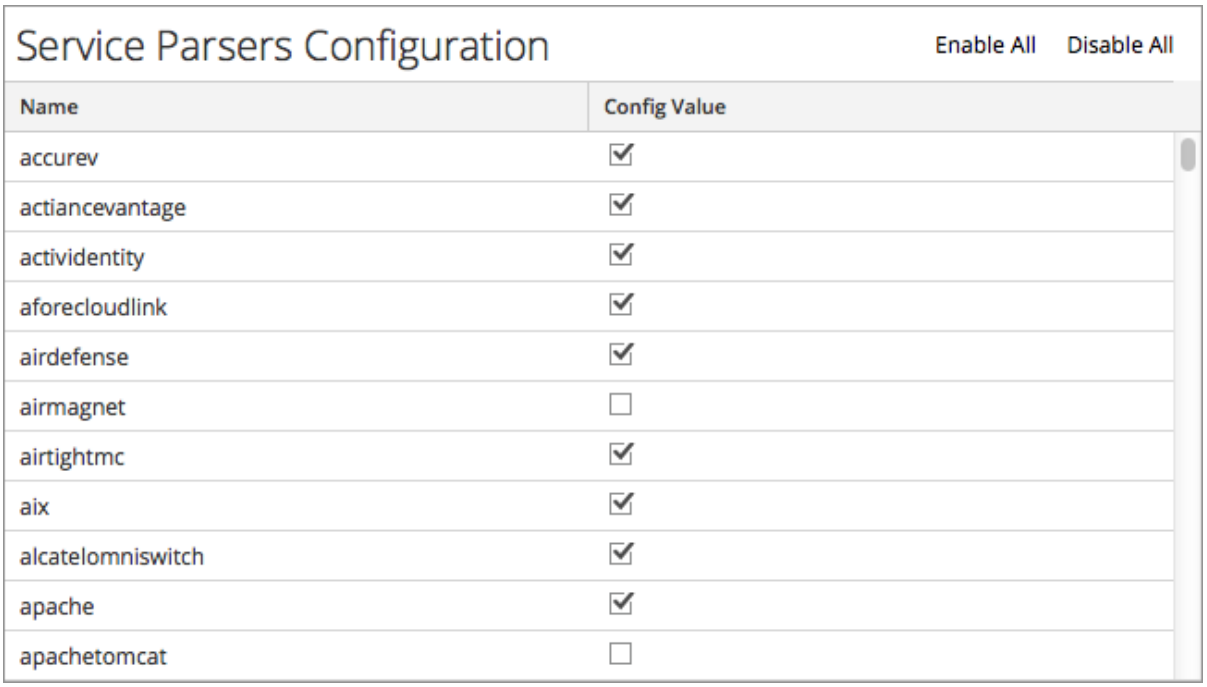
Parsers Configuration Panel

The Parsers Configuration panel provides a way to select parsers to use on the Decoder. Within some parsers, you can also configure the metadata that the parser creates. See Enable and Disable Parsers and Log Parsers for detailed information and procedures.



Service Parsers Configuration Section for Log Decoder

The Service Parsers Configuration section provides a way to select Service parsers to use on the Log Decoder.


Service Parsers Configuration
Enable All
Disable All
Name
Config Value
accurev
actiancevantage
actividentity
aforecloudlink
airdefense
airmagnet
airtightmc
aix
alcatelomniswitch
apache
apachetomcat

Services Config View - Parsers Tab

In the Services Config View > Parsers tab, you can view deployed parsers on a Decoder or Log Decoder, upload parsers, and delete deployed parsers. Parsers can be added and removed while a Decoder or Log Decoder is running without affecting capture.

To access the Parsers tab, go to  (Admin) > **Services** > select a **Decoder** or **Log Decoder** service and click  > **View** > **Config** > **Parsers** tab.

What do you want to do?

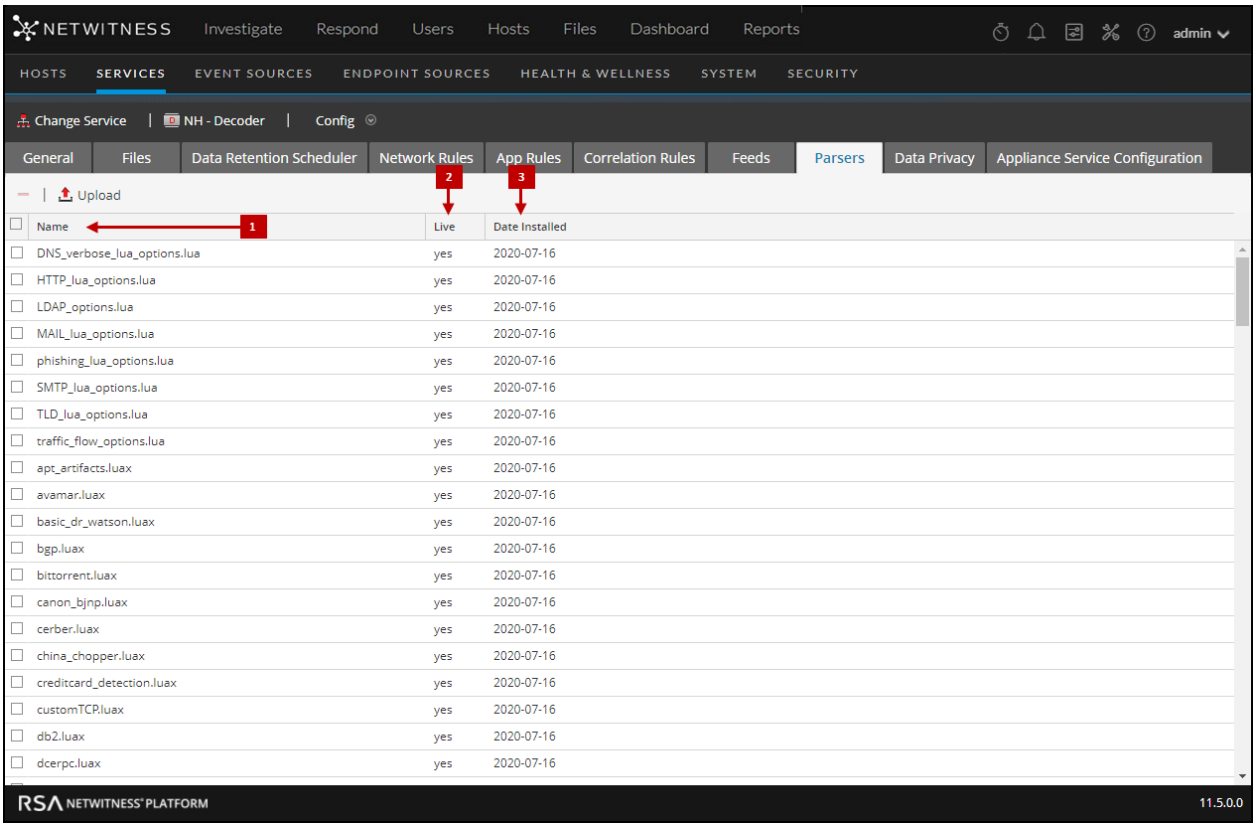
User Role	I want to...	Documentation
Administrator	View deployed parsers.	Enable and Disable Parsers and Log Parsers
Administrator	Upload parsers to a Decoder or Log Decoder.	Enable and Disable Parsers and Log Parsers

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- Upload and Delete Custom Parsers

Quick Look

This is an example of the Parsers tab. The Parsers grid lists all parsers that are currently deployed on the Decoder.



1 **Name**: The name of the parser or the parser file.

2 **Live**: Indicates if the parser originated from Live. Possible values are **Yes**, **No**, or **N/A**.

- **Yes** = Installed through Live Services.
- **No** = Installed through NetWitness.
- **N/A** = The parser has no attributes file created by NetWitness to track the installation date. The parser may have been installed manually, not through NetWitness or Live Services.

3 **Date Installed**: The date the parser was pushed to the service.

Parsers Tab Toolbar

The Parsers Tab toolbar has options to work with parsers in the grid.

Feature	Description
Upload	Enables you to upload parsers to a Decoder or Log Decoder.
−	Requests confirmation that you want to delete the selected parsers. You can select **No** to cancel the deletion, or select **Yes** to delete the selected parsers.

Services Config View - Parser Mappings Tab

This topic provides a description of the configurable options for a Log Decoder in the Parser Mappings tab.

In the Parser Mappings Administrators can configure log parser mappings for Log Decoder services. To access the Parser Mappings tab, go to (Admin) > **Services** > select a Log Decoder service and click > **View** > **Config** > **Parser Mappings** tab.

Note: You can also configure log parser mappings for Log Decoder services by navigating to **Admin > Services > Event Sources > Discovery**.

This feature is intended to track a subset of event sources that is parsing against the wrong parser.

What do you want to do?

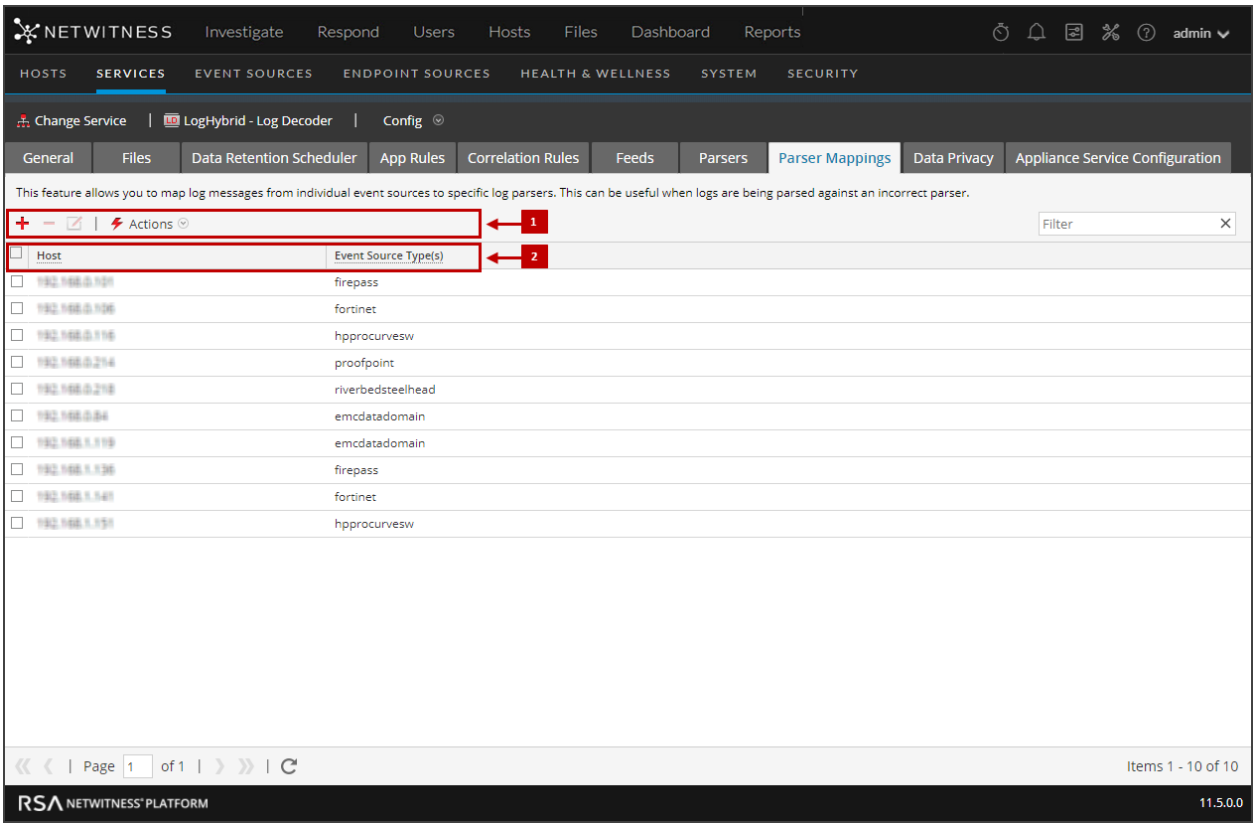
User Role	I want to...	Documentation
Administrator	Manage IPs for Event Source Mapping.	Enable Parser Mappings

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder

Quick Look

This is an example of the tab.



1 Parser Mappings Toolbar - Provides options to work with parser mappings in the list
2 Parser Mappings List - Lists all parsers that are currently mapped on the Log Decoder

Parser Mappings Toolbar

The Parser Mappings toolbar has options to work with parser mappings in the grid.

Feature	Description
	Add a parser mapping.
	Delete the selected parser mapping.
	Edit a parser mapping.
	Refresh the list of parser mappings.
	Display the Actions menu. • **Import** - Import a parser mapping to a file. • **Export** - Save a parser mapping to a file.

Parser Mappings List

The Parser Mappings list displays all parsers that are currently mapped on the Log Decoder.

Parameter	Description
Host	Displays the IP address of the host.
Event Source	Displays the event sources that are parsing incorrectly.

Parser Mappings Editor Dialog

The Parser Mappings Editor dialog allows you to update an IP to event source mapping.

To access the Parser Mappings Editor dialog, in the Services Config view for a Log Decoder, select the Parser Mappings tab.

Data Export Tab

The Data Export tab (**Admin** > **Services** > select a Decoder or Log Decoder and click > **View** > **Config** > **Data Export** tab) exports meta data and raw logs from the Decoder or Log Decoder and converts it to Logstash JSON object by configuring the export connector through Logstash.

Prerequisites

Make sure you have log collector service to configure Logstash in the UI.

What do you want to do?

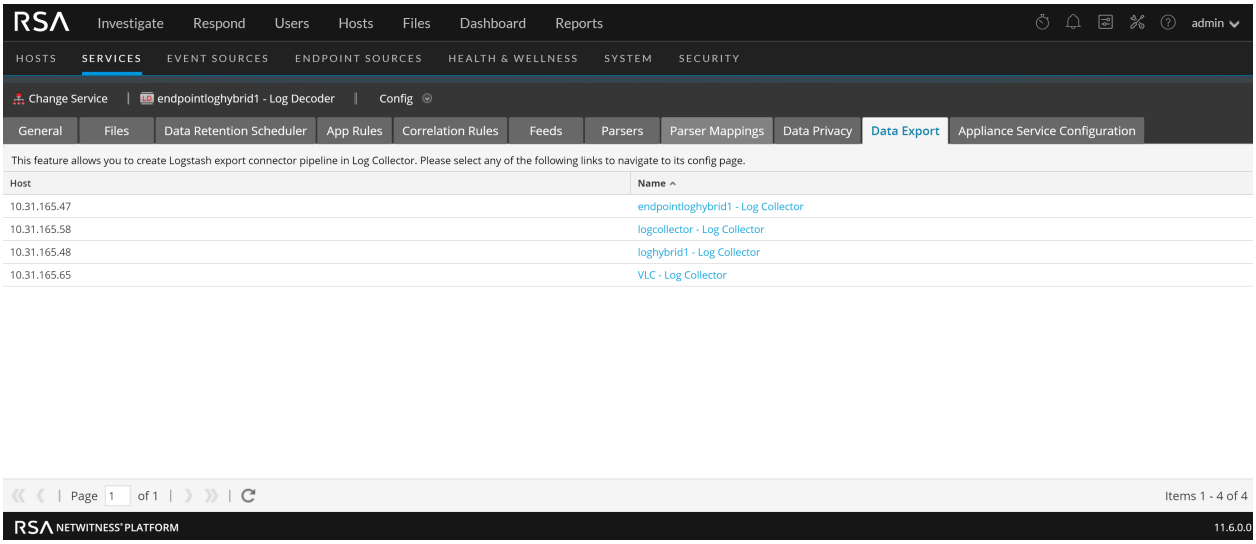
User Role	I want to...	Documentation
Administrator	Configure Export Connector	Configure Logstash Event Sources in NetWitness

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- Configure Decoder Rules
- Services Config View - Rules Tabs

Quick Look

The following figure shows an Data Export tab.



The data export list displays all the Log Collectors in your deployment.

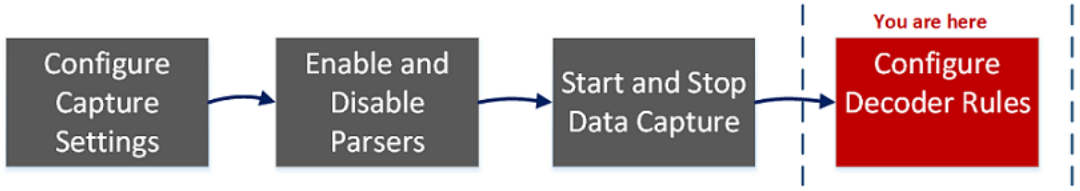
Parameter	Description
Host	Displays the IP address of the host.
Name	Displays the name of the Log Collector.

Services Config View - Rules Tabs

The Rules tabs in the Services Config view (**(Admin) > Services >** select a service and click **> View > Config**) enable you to define and manage capture rules. Each type of rule has a list with slightly different columns and different parameters in the Rule Editor dialog. Application and correlation rules apply to both Decoders and Log Decoders. Network rules apply only to Network Decoders.

Workflow

The following figure depicts the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	configure capture settings	Configure Capture Settings
Administrator	manage parsers and log parsers	Enable and Disable Parsers and Log Parsers
Administrator	start and stop data capture	Start and Stop Data Capture
Administrator	configure rules*	Configure Decoder Rules
Administrator	import, export, or push a rule*	Configure Decoder Rules
Administrator	enable or disable a rule*	Configure Decoder Rules
Administrator	add, edit, or delete a rule*	Configure Decoder Rules

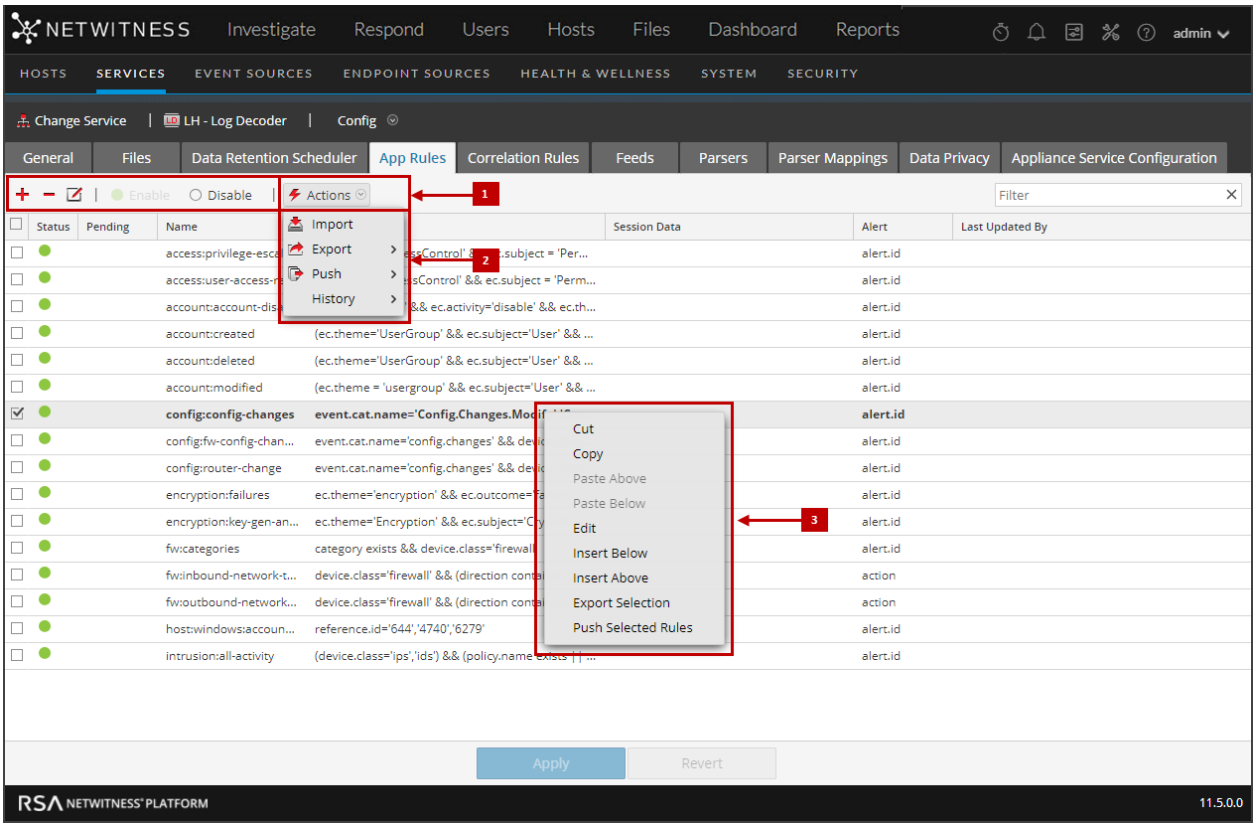
*You can complete these tasks here.

Related Topics

- Configure Common Settings on a Decoder
- Decoder and Log Decoder Quick Setup
- App Rules Tab
- Correlation Rules Tab
- Network Rules Tab

Quick Look

This is an example of the App Rules tab.



1 Rules Tab Toolbar - Provides options to work with rules in the list
2 Rules Actions Menu - Provide options to manage sets of rules
3 Rules List Context Actions - Displays the Rules List Context Menu

Rules Tab Toolbar

The toolbar is the same for all Config view > Rules tabs.

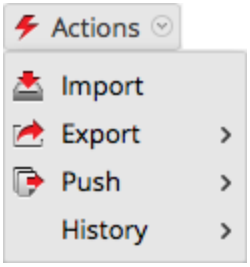


Feature	Description
Actions	Displays the **Actions** menu.
+	Adds a new rule to a service.
−	Deletes a rule from a service.
	Allows rule modification.
Disable	Disables a rule (without deleting the rule).
Enable	Enables (reactivates) a rule.

Feature	Description
Filter	The input field for a search string. NetWitness filters the rules dynamically as you type a search string. Clicking **x** clears the input field, restoring the unfiltered view.
Apply	Saves the changes made to rules and applies the configured rules to a service. Until you apply changes, it is possible to reload the rules as they were before current modifications.
Revert	Discards unsaved changes to the list and reverts to the unedited rules.

Rules Actions Menu

The Actions menu has options that help to manage sets of rules.



Option	Description
Import	Imports a set of rules into the user interface so that it can be applied to a service. You can edit the rules before applying.
Export	Saves selected rules or all rules to an .nwr file on the client machine.

Push	Allows rules to be applied to other services (Decoders or Log Decoders) or Decoders belonging to a service group. When pushing, the rules can either be merged (update existing rules and append new ones) or replaced.
	• **Push > All.** Pushes all rules to other services. All rules on the target services are removed and replaced with all of the rules on the source service.
	• **Push > Selection.** Pushes selected rules to other services. You have two options:
	  • **Replace.** Deletes all rules on the target services and replaces them with the selected rules from the source service.
	  • **Merge.** Merges the selected rules with the existing rules on the target services

History	Displays the last ten snapshots of rules applied through NetWitness. You can select and apply (restore) a snapshot to the Decoder at anytime.

Rules List Context Actions

Within a rules list, right-clicking a row displays the Rules list context menu.

Option	Description
Cut	Deletes the current rule.

Option	Description
Copy	Copies the current rule.
Paste Above	Pastes the copied rule above the current rule.
Paste Below	Pastes the copied rule below the current rule.
Edit	Edits the current rule.
Insert Below	Inserts imported rules below the current rule.
Insert Above	Inserts imported rules above the current rule.
Export Selection	Exports the selected rules.
Push Selected Rules	Pushes the selected rules to other services.

App Rules Tab

The App Rules tab (**Admin > Services >** select a Decoder or Log Decoder and click **> View > Config > App Rules tab**) enables you to manage application rules. NetWitness applies application rules at the session level.

What do you want to do?

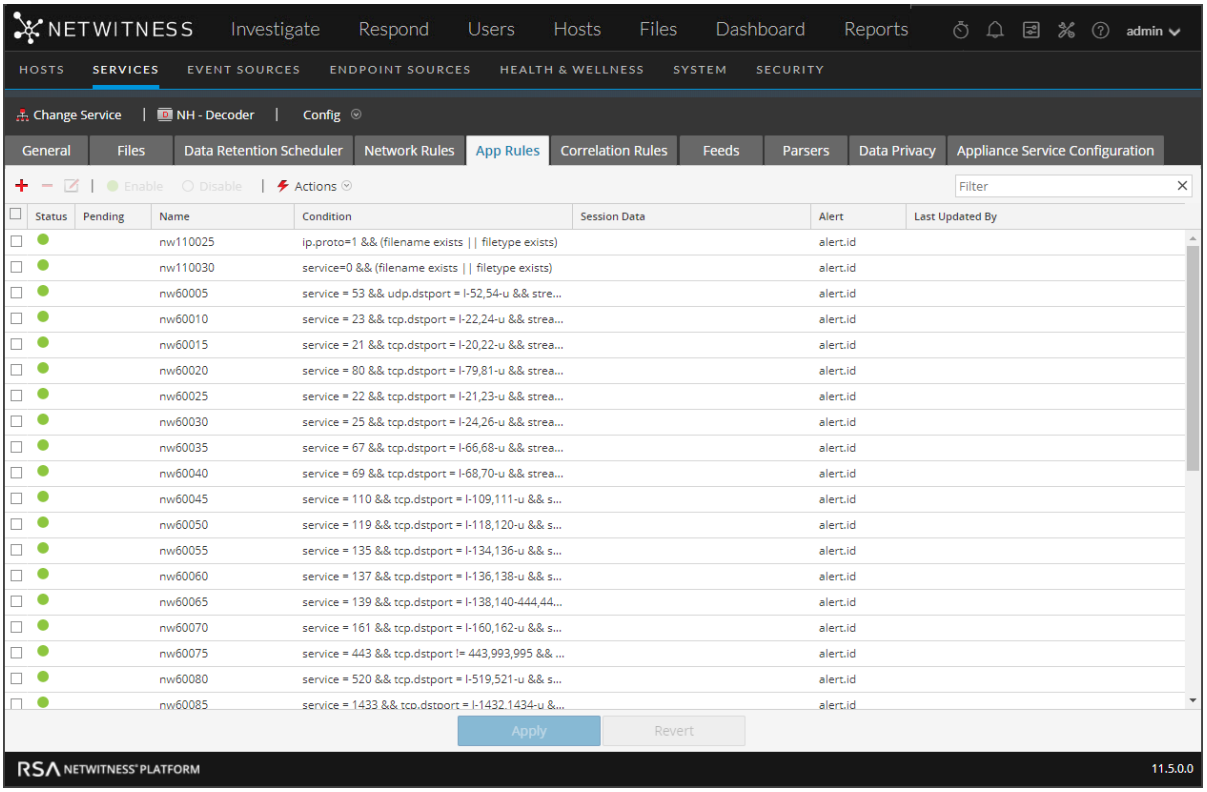
User Role	I want to...	Documentation
Administrator	add or edit application rules	Configure Application Rules

Related Topics

- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- Configure Decoder Rules
- Services Config View - Rules Tabs

Quick Look

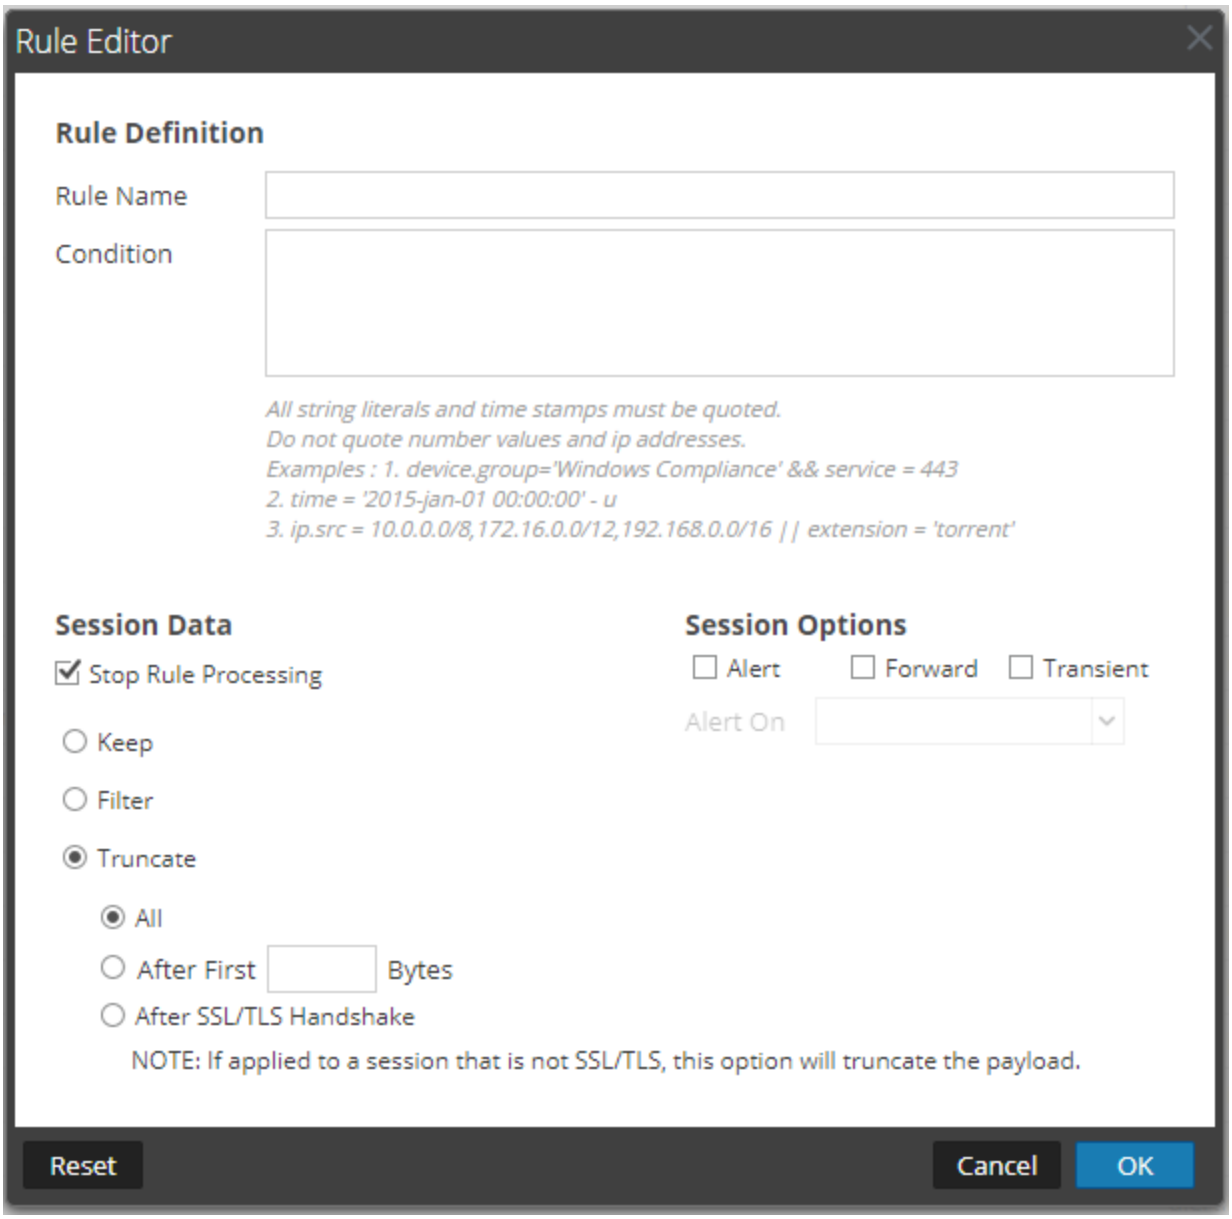
The following figure shows an App Rules tab and the table describes the columns.



Column	Description
Pending	This column indicates whether a rule has pending changes. Rules that are currently active on the Decoder have no indicator. If the rule is new or has been modified, the column contains . Once the rules are applied, the pending indicator is removed.
Name	This is the rule name, a descriptive identifier for the rule.
Condition	This is the definition of the condition that triggers an action when matched.
Session Data	This column displays the Session Data action taken when a packet matches the rule. Possible values are **Filter**, **Keep**, or **Truncate**.
Alert	This column displays the name of the custom alert that the Decoder generates when metadata matches the rule.
Status	This column indicates whether the rule is enabled or disabled with a circle icon. If the circle is filled green, the rule is enabled. If the circle is empty, the rule is disabled.

Rule Editor Dialog

The following figure shows the Rule Editor dialog for an application rule.



The Rule Editor dialog provides the fields and options needed to define an application rule.

Field	Description
Rule Name	The descriptive name that identifies the rule.
Condition	The definition of the condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the Intellisense window actions. As you build the rule definition, Intellisense displays syntax errors and warnings. All string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details.

The following table describes the Session Data actions and options.

Action	Description
Stop Rule Processing	If checked, further rule evaluation ends if the rule is matched, and the session is saved in accordance with the session action. If not checked, rule evaluation continues until all rules are evaluated.
Keep	The packet payload and associated metadata are saved when they match the rule.
Filter	The packet is not saved when it matches the rule.
Truncate	**Truncate All** – truncates all session payload bytes. The packet payload is not saved when it matches the rule, but packet headers and associated metadata are retained. This is the default truncation option. **Truncate After First <n> Bytes** – truncates the session payload bytes after the specified first <n> bytes, where <n> is an integer. The packet payload is not saved after <n> bytes when it matches the rule, but packet headers and associated metadata are retained. **Truncate SSL/TLS After Handshake** – truncates the payload for all sessions except in the case of an SSL/TLS session, where the SSL exchange is preserved, but the rest of the payload is not saved. This option is for use with SSL parsers.
Alert and **Alert On**	If **Alert** is checked, the packet generates a custom alert when metadata matches the rule. You can select the name of the alert in the **Alert On** field.
Forward	Enables the performance of syslog forwarding when the log matches the rule.
Transient	Prevents the alert metadata that is created from being written to the disk.

The following table describes Rule Editor dialog actions.

Action	Description
Reset	Resets the contents of the dialog to their values before editing; changes are discarded.
Cancel	Cancels any edits and closes the Rule Editor dialog.
OK	Saves the new rule or edited rule, and adds it to the rules grid. The Rule Editor dialog closes.
Save	(Rules with deprecated syntax only) Applies a corrected rule individually to the Decoder service. See Fix Rules with Invalid Syntax.

Correlation Rules Tab

The Correlation Rules tab (**(Admin)** > select a service and click **> View > Config > Correlation Rules tab**) enables you to manage correlation rules. Basic correlation rules are applied at the session level and alert the user to specific activities that may be occurring in their environment. NetWitness applies correlation rules over a configurable sliding time window.

What do you want to do?

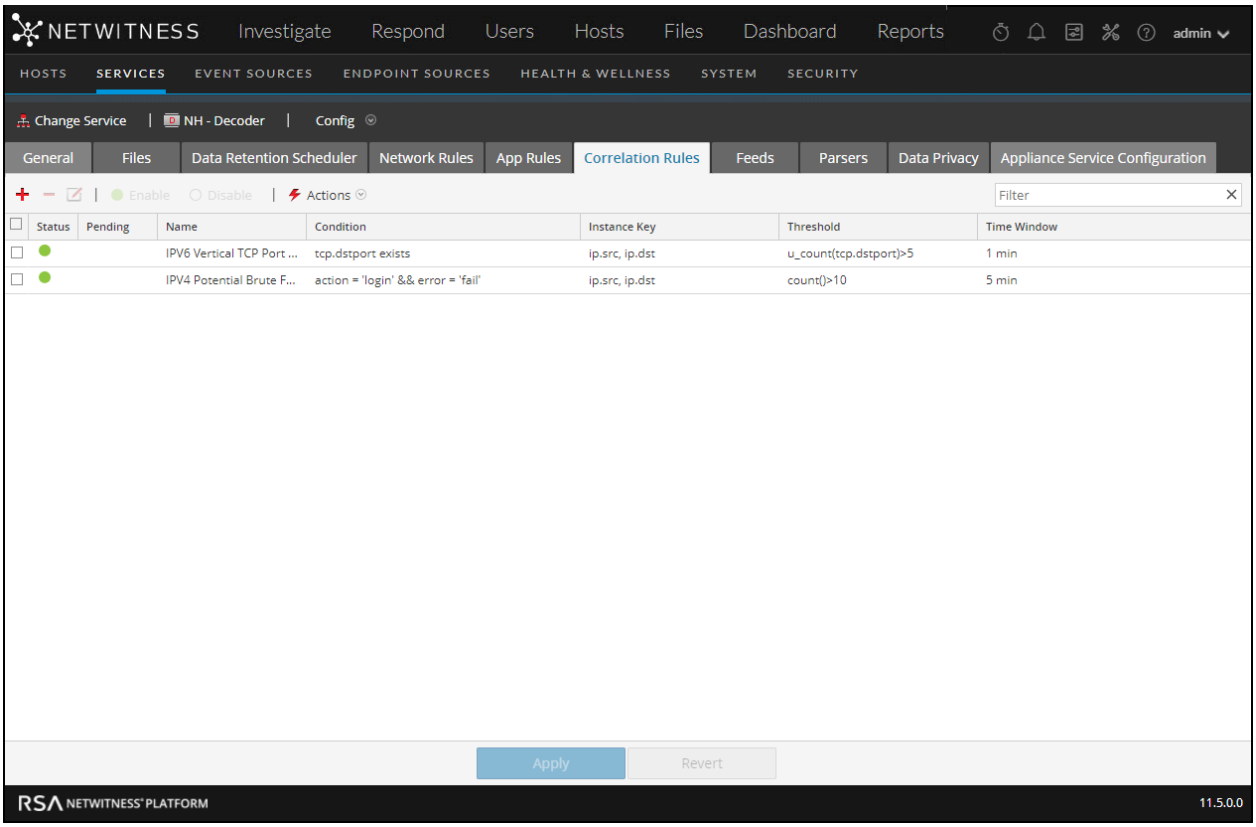
User Role	I want to...	Documentation
Administrator	add or edit a correlation rule	Configure Correlation Rules

Related Topics

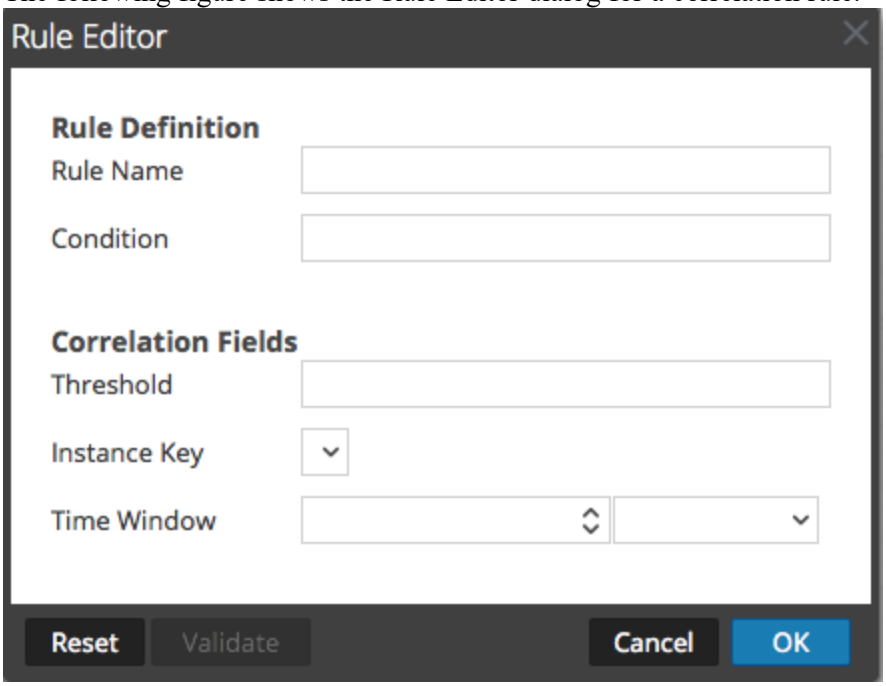
- Configure Common Settings on a Decoder
- Decoder and Log Decoder Quick Setup
- Configure Decoder Rules
- Services Config View - Rules Tabs

Quick Look

The following figure shows the Correlation Rules tab.



The following figure shows the Rule Editor dialog for a correlation rule.



The following table describes the Correlation Rules tab columns.

Column	Description
Pending	This column indicates whether a rule has pending changes. Rules that are currently active on the Decoder have no indicator. If the rule is new or has been modified, the column contains . Once the rules are applied, the pending indicator is removed.
Name	This is the descriptive name for the rule.
Condition	This is the definition of the condition that triggers an action when matched. In conditions, all string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details.
Instance Key	This is the target indicator to base the event upon. It can be a single primary key, such as ip.src or a compound primary key such as `ip.src,ip.dst`.
Threshold	This is the minimum number of occurrences required to trigger a correlation session and can include a associated key that identifies the meta type that were are counting to determine if the condition is satisfied. The correlation engine cannot use IPv4 or IPv6 as an associated meta type. Use one of these three arguments: • `u_count(associated_key)` = the count of unique values of the specified key. A key is required. • `sum(associated_key)` = the values of the specified key. a key is required. • `count()` = number of sessions, no associated key used. If included, it is ignored.
Time Window	This is the duration in hours, minutes, or seconds within which the threshold must be reached to trigger a correlation session.
Status	This column indicates whether the rule is enabled or disabled with a circle icon. If the circle is filled green, the rule is enabled. If the circle is empty, the rule is disabled.

The **Rule Editor** dialog provides the fields and options needed to define a network rule. The fields correspond exactly to the grid columns.

Action	Description
Reset	Resets the contents of the dialog to their values before editing; changes are discarded.
Cancel	Cancels any edits and closes the Rule Editor Dialog.
OK	Saves the new rule or edited rule, and adds it to the rules grid. The Rule Editor Dialog closes.
Save	(Rules with deprecated syntax only) Applies a corrected rule individually to the Decoder service. See Fix Rules with Invalid Syntax.

Network Rules Tab

The Network Rules tab (**(Admin) > Services >** select a Decoder and click **> View > Config > Network Rules tab**) enables you to manage network rules. NetWitness applies network rules at the packet level. Network rules consist of rule sets from Layer 2, Layer 3, and Layer 4. Multiple rules can be applied to the Decoder. Rules can be applied to multiple layers (for example, when a network rule filters out specific ports for a specific IP address). Network rules apply only to Network Decoders.

What do you want to do?

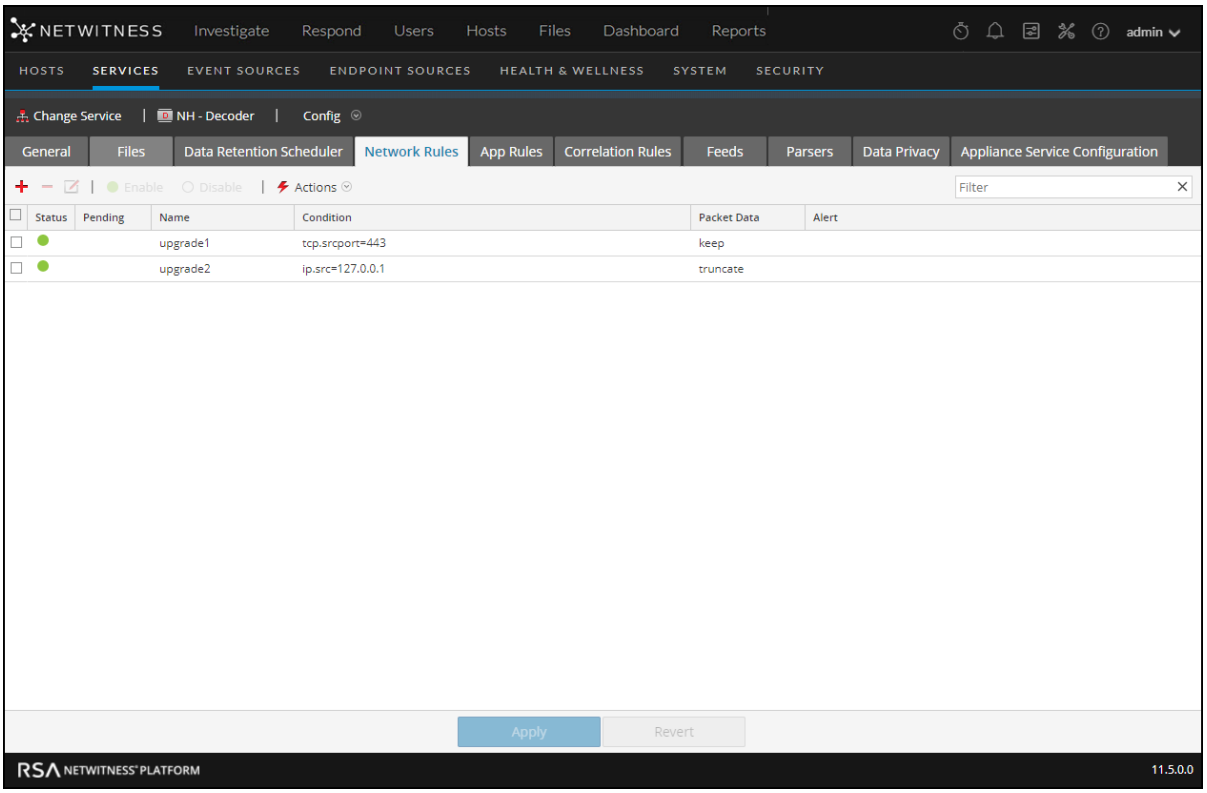
User Role	I want to...	Documentation
Administrator	add, edit, or fix network rules	Configure Network Rules

Related Topics

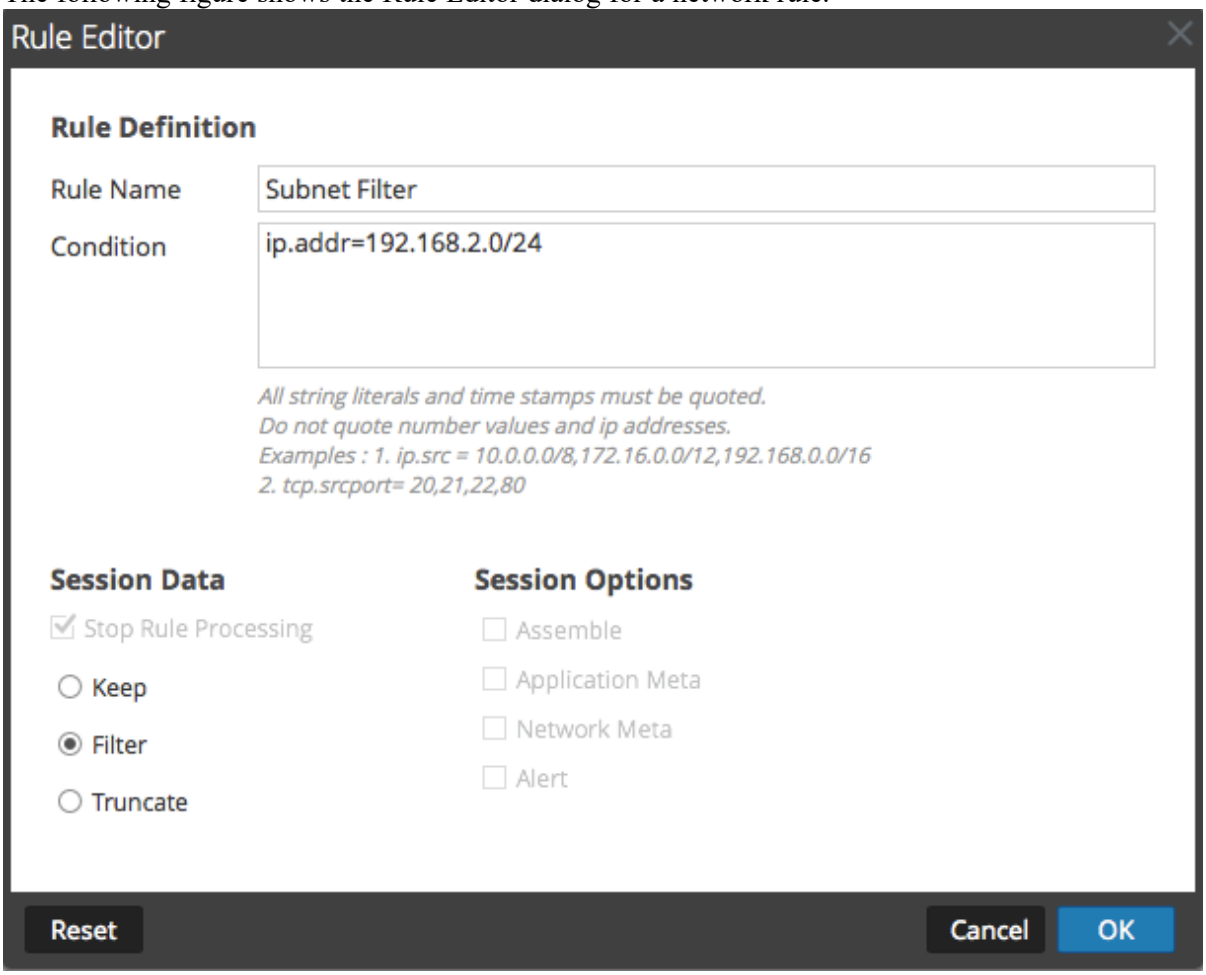
- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- Configure Decoder Rules
- Services Config View - Rules Tabs

Quick Look

The following figure shows the Network Rules tab.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Change Service
NH - Decoder
Config
General
Files
Data Retention Scheduler
Network Rules
App Rules
Correlation Rules
Feeds
Parsers
Data Privacy
Appliance Service Configuration
Enable
Disable
Actions
Filter
Status
Pending
Name
Condition
Packet Data
Alert
upgrade1
tcp.srcport=443
keep
upgrade2
ip.src=127.0.0.1
truncate
Apply
Revert
RSA NETWITNESS PLATFORM
11.5.0.0

The following figure shows the Rule Editor dialog for a network rule.



The following table describes the columns in the Network Rules grid.

Column	Description
Pending	This column indicates whether a rule has pending changes. Rules that are currently active on the Decoder have no indicator. If the rule is new or has been modified, the column contains . Once the rules are applied, the pending indicator is removed.
Name	This is the rule name, a descriptive identifier for the rule.
Condition	This is the definition of the condition that triggers an action when matched.
Packet Data	This column displays the Session Data action taken when a packet matches the rule. Possible values are **Filter**, **Keep**, or **Truncate**.
Alert	This column indicates whether the Decoder generates a custom alert when metadata matches the rule. Possible values are **Enabled** or **Disabled**.
Status	This column indicates whether the rule is enabled or disabled with a circle icon. If the circle is filled green, the rule is enabled. If the circle is empty, the rule is disabled.

The **Rule Editor** dialog provides the fields and options needed to define a network rule.

The following table describes the Rule Definition fields.

Field	Description
Rule Name	The descriptive name that identifies the rule.
Condition	The definition of the condition that triggers an action when matched. You can type directly in the field or build the condition in this field using meta from the Intellisense window actions. As you build the rule definition, Intellisense displays syntax errors and warnings. In conditions, all string literals and time stamps must be quoted. Do not quote number values and IP addresses. Configure Decoder Rules provides additional details. This section also describes the meta keys that NetWitness supports for use in network rule conditions.

The following table describes the Session Data actions.

Action	Description
Stop Rule Processing	If checked, further rule evaluation ends if the rule is matched, and the session is saved as indicated. If not checked, rule evaluation continues until all rules are evaluated.
Keep	The packet payload and associated meta are saved when they match the rule.
Filter	The packet is not saved when it matches the rule.
Truncate	The packet payload is not saved when it matches the rule, but packet headers and associated meta are retained.

The following table describes the session options.

Action	Description
Assemble	If checked, the assembler assembles the packet chain when it matches the rule.
Network Meta	The packet generates network metadata when it matches the rule.
Application Meta	The packet generates application metadata when it matches the rule.
Alert	The packet generates a custom alert when metadata matches the rule.

The following table describes Rule Editor dialog actions.

Action	Description
Reset	Resets the contents of the dialog to their values before editing; changes are discarded.
Cancel	Cancels any edits and closes the Rule Editor dialog.
OK	Saves the new rule or edited rule, and adds it to the rules grid. The Rule Editor dialog closes.
Save	(Rules with deprecated syntax only) Applies a corrected rule individually to the Decoder service. See Fix Rules with Invalid Syntax.

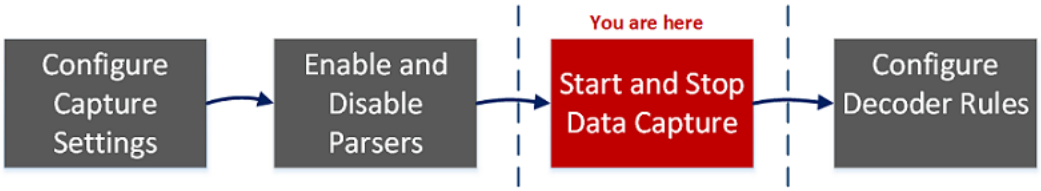
Services System View - Decoders

A Log Decoder is a special type of Decoder, and is configured and managed in a similar way to a Decoder. Therefore, most of the information in this section refers to both types of Decoders. Differences for Log Decoders are noted.

To reach the Services System view, go to **(Admin) > Services >** select a Decoder or Log Decoder > > **View** > **System**.

Workflow

The following figure depicts the workflow for common Decoder configuration tasks with the steps you can complete in this view highlighted.



What do you want to do?

User Role	I want to...	Documentation
Administrator	configure capture settings	Configure Capture Settings
Administrator	manage parsers and log parsers	Enable and Disable Parsers and Log Parsers
Administrator	start and stop data capture*	Start and Stop Data Capture
Administrator	upload packet capture and log files*	Upload a Log File to a Log Decoder Upload a Packet Capture File
Administrator	reset log stats, perform host tasks, shutdown the service, shutdown the appliance service, and reboot the host*	*Hosts and Services Getting Started Guide*
Administrator	configure rules	Configure Decoder Rules

*You can complete these tasks here.

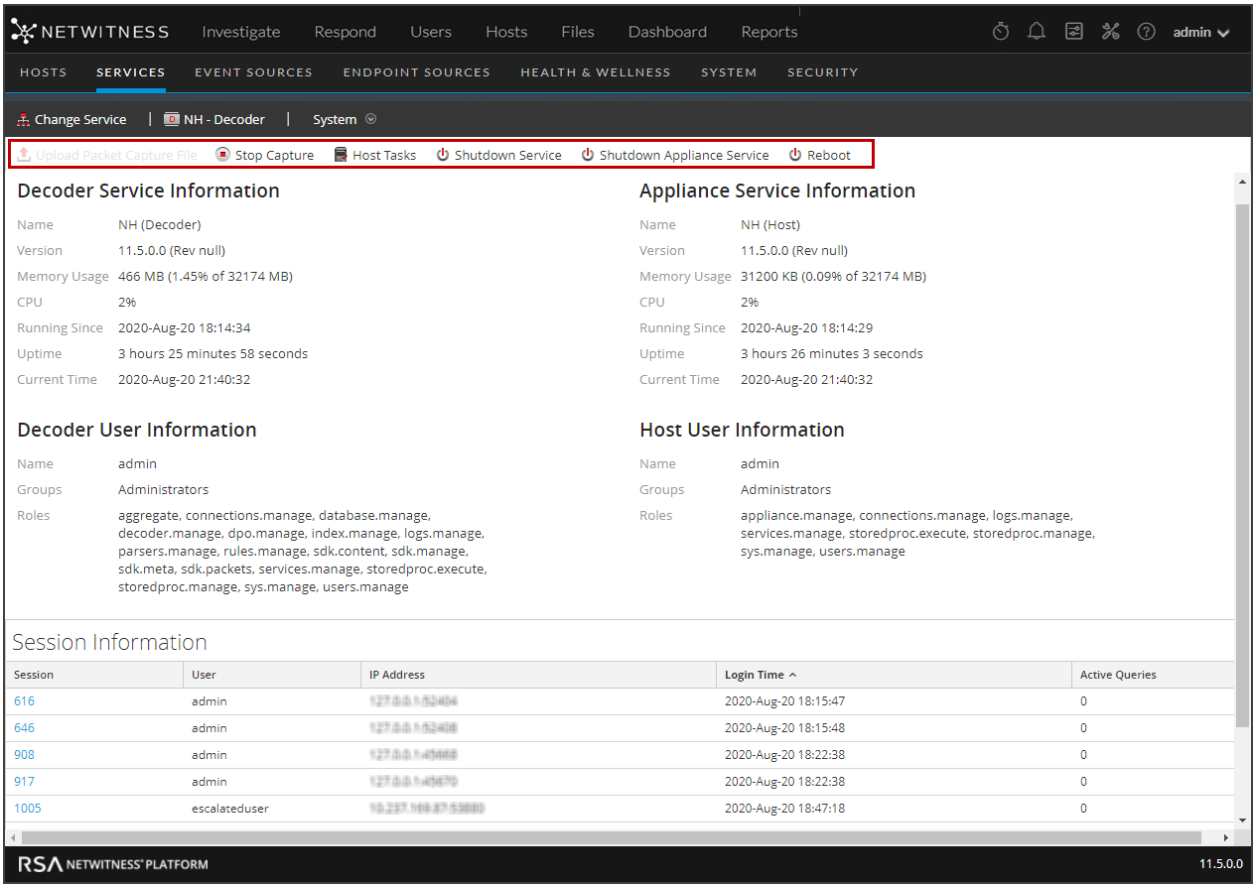
Related Topics

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

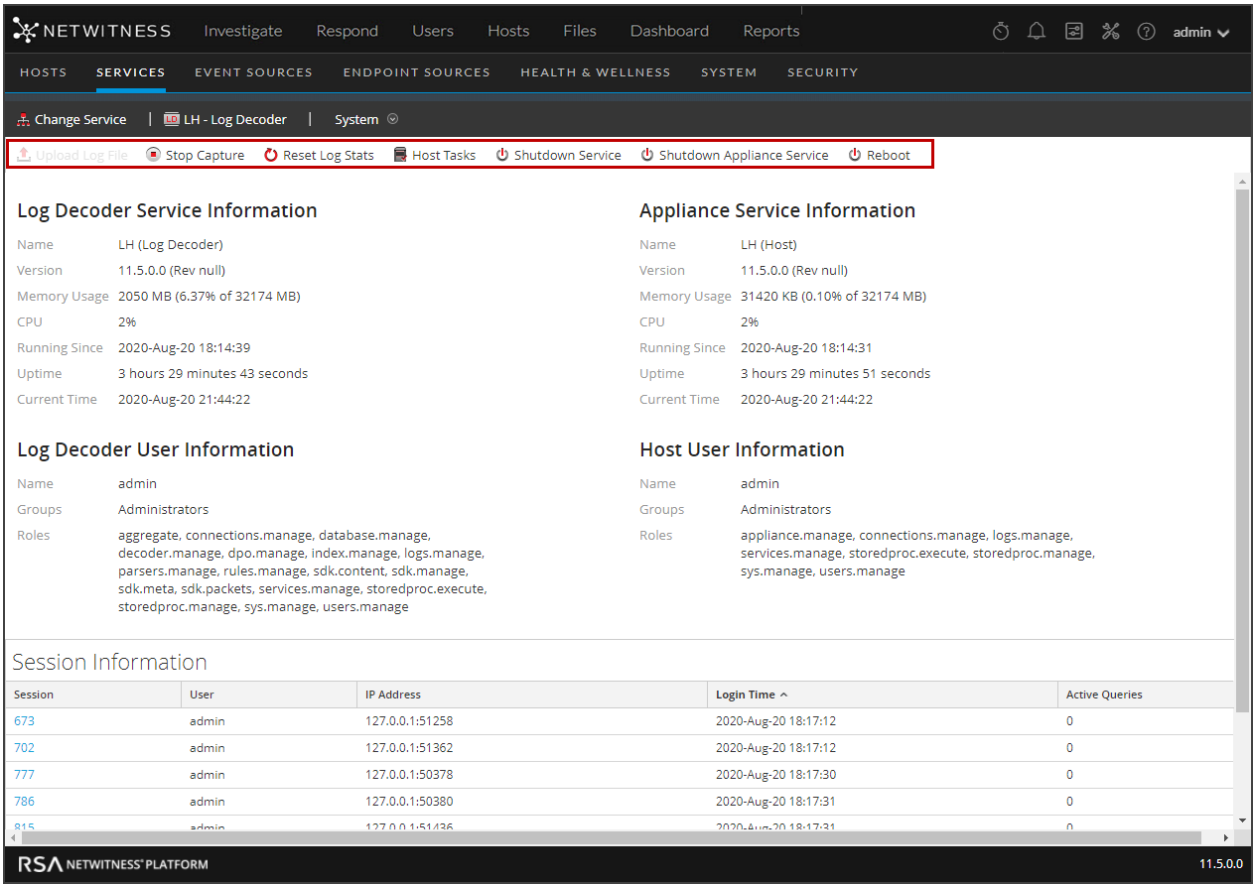
- Decoder and Log Decoder Quick Setup
- Configure Common Settings on a Decoder
- "Services System View" in the *Hosts and Services Getting Started Guide*

Quick Look

This is an example of the Services System view for a Decoder.

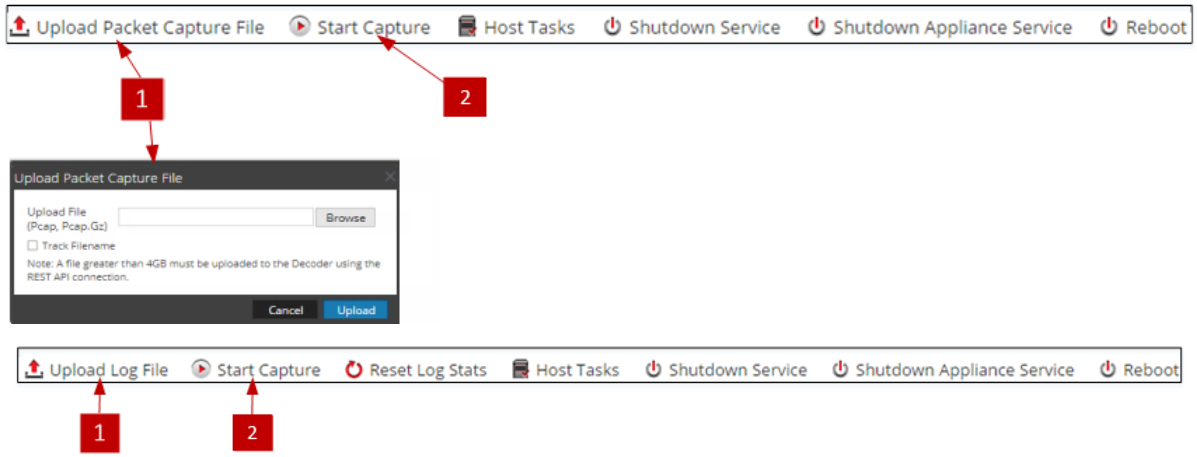


This is an example of the Services System view for a Log Decoder.



Service Info Toolbar

These two toolbars illustrate the options specific to Decoders and Log Decoders.



In addition to the common options in the Services System view toolbar, you can start and stop capture of packets or logs. The upload file options are different for the standard Decoder (packet capture file) and the Log Decoder (log file).

Action	Description
Upload Packet Capture File	Displays a dialog that provides a way to select a packet capture (.pcap) file for upload to the selected Decoder. For more information, see Upload a Packet Capture File. **Note:** This option does not apply to Log Decoders.
Upload Log File	Displays a dialog that provides a way to select a log (.log) file for upload to the selected Log Decoder. For more information, see Upload a Log File to a Log Decoder.
Start/Stop Capture	Starts packet capture on the selected Decoder. When packet capture is in progress, the option in the toolbar changes to Stop Capture, and the option to upload a file is unavailable.

NetWitness® Platform XDR

Version 12.0

Deployment Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

November, 2022

Contents

The Basics **6**

Basic Deployment 7

Process 7

NetWitness High-Level Deployment Diagrams 8

Deployment Optional Setup Procedures **9**

Analyst User Interface 9

Features and Limitations 9

Use Case 10

Deployment 10

Group Aggregation 12

RSA Group Aggregation Deployment Recommendations 12

Advantages of Using Group Aggregation 12

Configure Group Aggregation 14

Prerequisites 14

Set up Group Aggregation 15

New Health and Wellness 18

Hybrid Categories on Series 6 (R640) Hardware 21

NW Server Deployment on ESA Hardware 21

Second Endpoint Server 22

Warm Standby NW Server Host 23

Procedures 23

Planned Fail-Over Scenario 24

Required Fail-Over Scenario without Hardware Replacement 24

Required Fail-Over Scenario with Hardware Replacement 24

Set Up Secondary NW Server in Standby Role 25

Fail Over Primary NW Server to Secondary NW Server with Same IP Address 38

Fail Over Primary NW Server to Secondary NW Server with Different IP Address 39

SSO 41

Reporting Engine 42

UCF 43

PAM 43

ECAT 43

RSA NetWitness Orchestrator (By Demisto) 46

Audit Logging 47

Health and Wellness 47

Malware Analysis 47

Windows Legacy Collection 48
Fail Back Secondary NW Server to Primary NW Server 49
Network Architecture and Ports 50
NetWitness Network Architecture Diagram 50
NetWitness Network (Packets) Architecture Diagram with Ports 51
NetWitness Logs Architecture Diagram with Ports 52
Event Stream Analysis Network (Packets) Architecture Diagram with Ports 53
Event Stream Analysis (Logs) Architecture Diagram with Ports 54
NetWitness Firewall Requirements Summary 55
Comprehensive List of NetWitness Host, Service, and iDRAC Ports 59
NW Server Host (Primary and Warm Standby NW Server Host) 60
Analyst UI Host 61
Archiver Host 62
Broker Host 63
Concentrator Host 64
Endpoint Log Hybrid 65
Endpoint Relay Server 67
Event Stream Analysis (ESA) Host 68
New Health and Wellness 69
New Health and Wellness on Different Subnet 69
iDRAC Ports 70
Log Collector Host 71
Log Decoder Host 72
Log Hybrid Host 73
Log Hybrid - Retention Host 75
Malware Host 76
Network Decoder Host 77
Network Hybrid Host 78
UEBA Host 79
NetWitness Endpoint Architecture 80
NetWitness Endpoint 4.4 Integration with NetWitness Platform 80
NetWitness Endpoint Architecture with Ports 81
How to Change UDP Port for Endpoint Log Hybrid 81
Task 1 - Tell All Agents to Use a New UDP Port 81
Task 2 - Update the Port on All Endpoint Log Hybrid Hosts in Your Environment 82
Site Requirements and Safety 84
Intended Application Uses 84
Service 84
Safety Information 84
Site Selection 84

Equipment Handling Practices 84
Power and Electrical Warnings 85
Rack Mount Warnings 85
Cooling and Air Flow 85

The Basics

This guide describes the basic requirements of a NetWitness deployment and outlines optional scenarios to address the needs of your enterprise. Even in small networks, planning can ensure that all goes smoothly when you are ready to bring the hosts online.

Note: This document refers to several additional documents available on NetWitness Community Link. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

There are many factors you must consider before you deploy NetWitness. The following items are just some of these factors. You need to estimate growth and storage requirements when you consider these factors:

- The size of your enterprise (that is, the number of locations and people that will use NetWitness)
- The volume of network data and logs you need to process
- The performance each NetWitness user role needs to do their jobs effectively.
- The prevention of downtime (that is, how to avoid a single point of failure).
- The environment in which you plan to run NetWitness
  - RSA Physical Hosts (software running on hardware supplied by RSA)
  - Software Only provided by RSA:
    - On-Premises (On-Prem) Virtual Hosts
      See the *NetWitness Virtual Host Installation Guide* for detailed instructions on how to deploy on-prem virtual hosts.
    - VCloud:
      - Amazon Web Services (AWS)
        See the *NetWitness AWS Installation Guide* for detailed instructions on how to deploy virtual hosts in AWS.
      - Azure
        See the *NetWitness Azure Installation Guide* for detailed instructions on how to deploy virtual hosts in Azure.
      - Google Cloud Platform (GCP)
        See the *NetWitnessGoogle Cloud Platform Installation Guide* for detailed instructions on how to deploy virtual hosts in Google Cloud.

Basic Deployment

Before you can deploy NetWitness you need to:

- Consider the requirements of your enterprise and understand the deployment process.
- Have a high-level picture of the complexity and scope of a NetWitness deployment.

Process

The components and topology of a NetWitness network can vary greatly between installations, and should be carefully planned before the process begins. Initial planning includes:

- Consideration of site requirements and safety requirements.
- Review of the network architecture and port usage.
- Support of group aggregation on Archivers and Concentrators, and virtual hosts.

When ready to begin deployment, the general sequence is:

- For RSA Physical Hosts:
  1. Install physical hosts and connect to the network as described in the NetWitness Hardware Setup Guides and the *NetWitness Physical Host Installation Guide*.
  2. Set up licensing for NetWitness as described in the *NetWitness Licensing Guide*.
  3. Configure individual physical hosts and services as described in *NetWitness Host and Services Getting Started Guide*. This guide also describes the procedures for applying updates and preparing for version upgrades.
- For On-Prem virtual hosts, follow the instructions in the *NetWitness Virtual Host Setup Guide*.
- For AWS, follow the instructions in the *NetWitness AWS Installation Guide*.
- For Azure, follow the instructions in the *NetWitness Azure Installation Guide*.
- For Google Cloud, follow the instructions in the *NetWitnessGoogle Cloud Platform Installation Guide*.

When updating hosts and services, follow recommended guidelines under the "Running in Mixed Mode" topic in the *NetWitness Host and Services Getting Started Guide*.

You should also become familiar with Hosts, Host Types, and Services as they are used in the context of NetWitness also described in the *NetWitness Host and Services Getting Started Guide*.

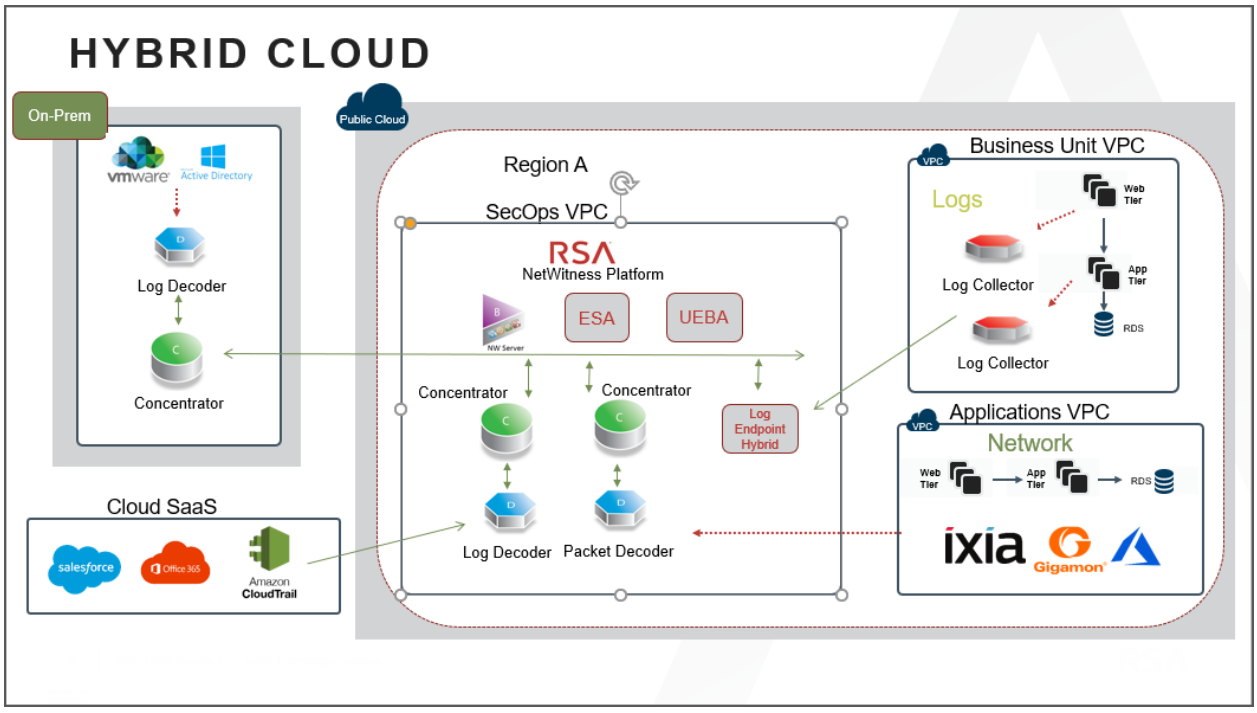
NetWitness High-Level Deployment Diagrams

NetWitness is inherently modular. Whether organizations are looking to deploy on-premise or in the cloud, the NetWitness components are decoupled in a way which allows flexible deployment architectures to satisfy a variety of use cases.

The following figure is an example of a hybrid cloud deployment, where the base of the components are residing within the SecOps VPC. Centralizing these components make management easier while keeping network latency to a minimum.

Network, log, and endpoint traffic could then be aggregated up to the SecOps VPC. The on-premise location would function just like a normal physical deployment and would be accessible for investigations and analytics.

Cloud SaaS visibility could be captured from a Log Decoder residing in either the cloud or on-premise locations.



Deployment Optional Setup Procedures

You can deploy NetWitness with the following options.

Analyst User Interface
Group Aggregation
New Health and Wellness Search
Hybrid Categories on Series 6 (R640) Hardware
NW Server Deployment on ESA Hardware
Second Endpoint Server
Warm Standby NW Server

Analyst User Interface

The Analyst User Interface (UI) gives you access to a subset of features in the NetWitness Platform UI that you can set up in individual locations when you deploy NetWitness Platform in multiple locations. It is designed to reduce latency and improve the performance that can occur when accessing all functionality from the Primary User Interface on the NW Server Host (Primary UI).

You can have can have multiple Analyst UI instances provisioned in the same manner as the other NW component hosts.

Features and Limitations

Each Analyst UI host:

- Can be deployed to specific organizational groups. For example: the Americas, EMEA, APAC, Tier 1 Analysts, Tier 3 Analysts.
- If Analyst UI hosts are deployed regionally, you have the capability of querying those regional brokers directly (less latency), instead of than having to route through the Primary UI.
- Helps distribute load off the Primary UI.
- Has its own Reporting Engine (RE).
- If it becomes unavailable for any planned or unplanned reason, it will not affect the Primary UI or any other Analyst UI instances.
- Provides the same pre-query filter verification, Data Privacy protection, and RBAC functionality as the Primary UI.
- Points back to the primary NW Server for authentication and configuration.
- Does not have access to any administrative functions. All administration functions take place on the Primary UI.
- Does not allow you to create or manage Content (that is, ESA rules, app rules, feeds). All Content creation and management takes place on the Primary UI.

Use Case

Large environments that include Geo distribution with a single data center and multiple NW Servers require Analyst UI instances in all their NetWitness locations or managed entities.

For example, if an Analyst UI is deployed for the EMEA SOC team, analysts can query their EMEA NetWitness Platform hosts directly. If the EMEA team has Broker hosts and Concentrator hosts within the region, the Analyst UI can connect and query them instead of connecting back to Primary user Interface (Primary UI).

Deployment

You must install the **Analyst UI** service category on a dedicated host and you install it in the same manner as any component service category on a host.

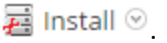
See the "Task 2 - Install 11.7 on Other Component Hosts" in the NetWitness Platform Installation Guides for instructions on how to install any component service. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

After you provision the Analyst UI host (that is after you run the nwsetup-tui for the component host designated for the Analyst UI), complete the following steps to install the Analyst UI service category on the provisioned host.

1. Log in to **NetWitness** and go to  (Admin) > **Hosts**.
   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background.

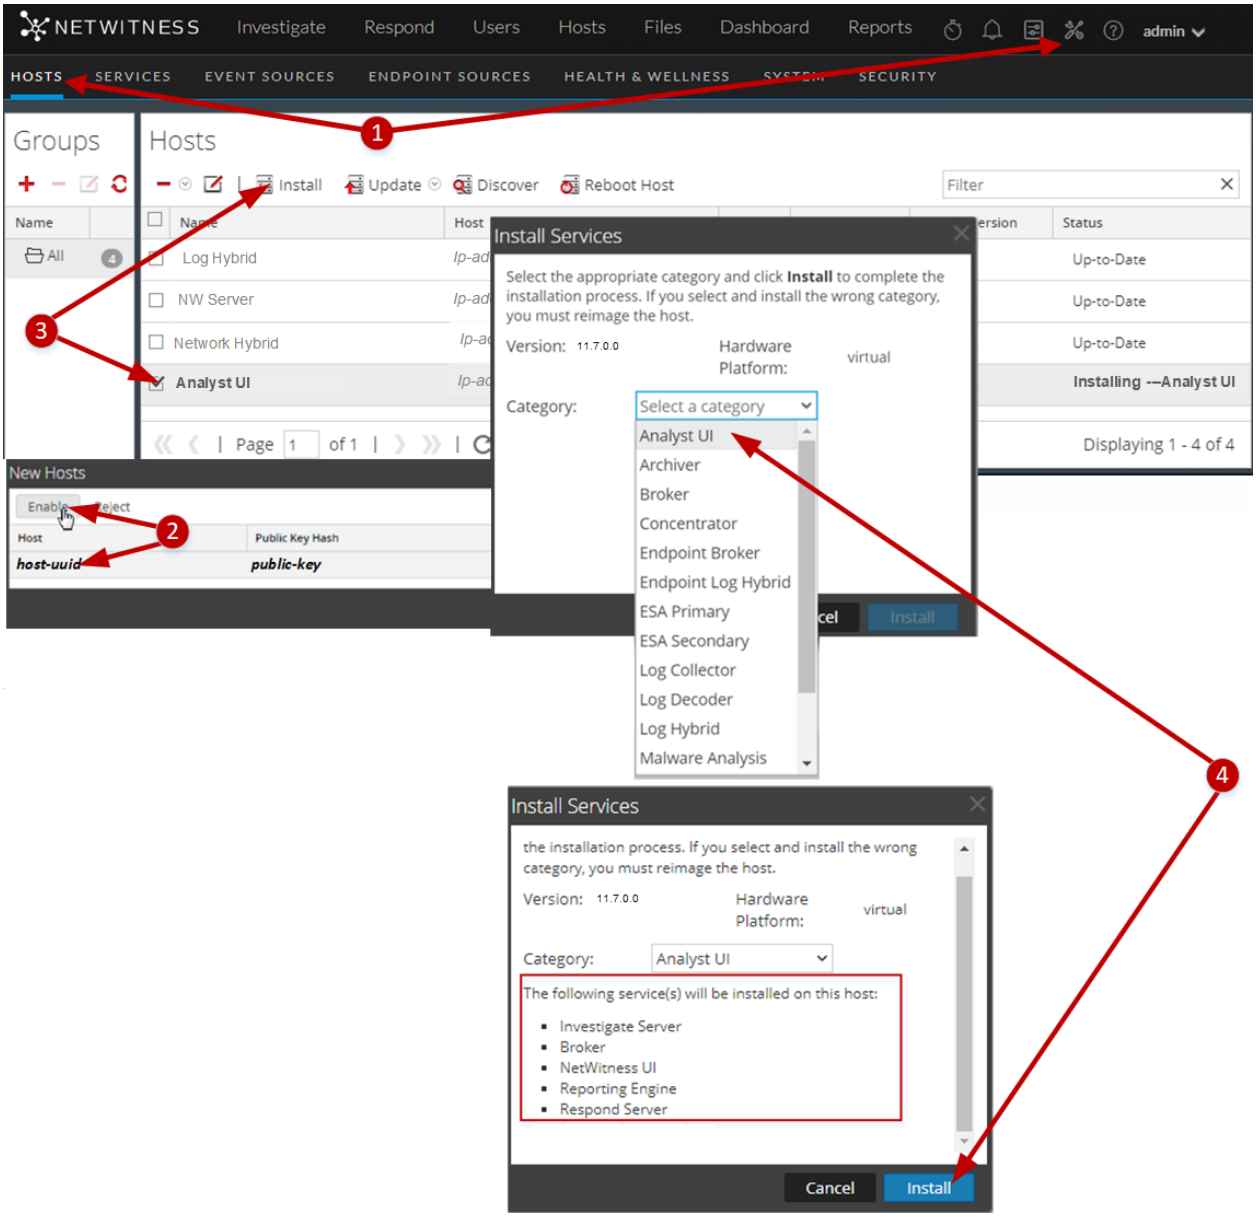
   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

2. Select the host in the **New Hosts** dialog and click **Enable**.
   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

3. Select that host in the **Hosts** view (for example, **Analyst UI**) and click 

   .
   The **Install Services** dialog is displayed.

4. Select **Analyst UI** in **Category** and click **Install**.



5. Configure NetWitness Platform for each Analyst UI instance.Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

   a. Make sure that each Analyst UI instance is connected to the correct local Reporting Engine and has the appropriate Investigation parameters set. The *Getting Started Guide for NetWitness Platform* describes the default Analyst UI Dashboard and how you manage dashboards.

   Note: You must add data sources to each Reporting Engine instance to execute Reports and Charts on an Analyst UI. See "Configure the Data Sources" in the*Reporting Engine Configuration Guide for NetWitness Platform* for instructions.

b. Configure whether to normalize alerts for any Respond Server (NW Server or Analyst UI) by enabling or disabling alert normalization. "Configure Analyst UI for Respond Server Alert Normalization" in the *NetWitness Respond Configuration Guide for NetWitness Platform* tells you how to configure Respond Server alert normalization for the Analyst UI.

Group Aggregation

You use Group Aggregation to configure multiple Archiver or Concentrator services as a group and share the aggregation tasks between them. You can configure multiple Archiver services or Concentrator services to efficiently aggregate from multiple Log Decoder services to improve query performance on the data:

- Stored in the Archiver.
- Processed through the Concentrator.

RSA Group Aggregation Deployment Recommendations

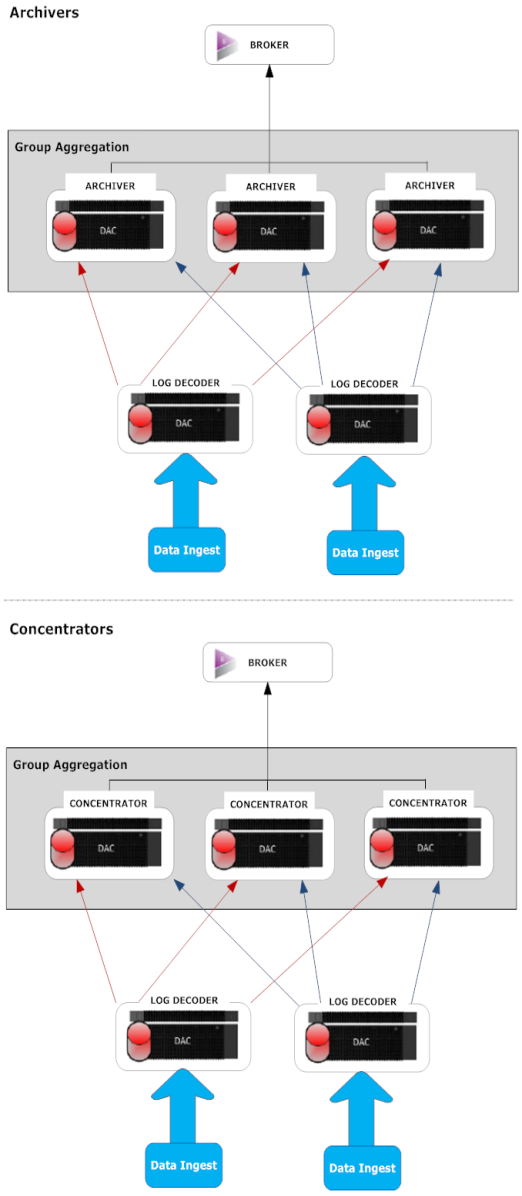
RSA recommends the following deployment for Group Aggregation:

- 1 - 2 Log Decoders
- 3 - 5 Archivers or Concentrators

Advantages of Using Group Aggregation

- Increases the speed of NetWitness queries.
- Improves the performance of aggregate queries (Count and Sum) on the environment.
- Enhances investigation service performance.
- Gives you the option of storing data for a longer duration for investigation purposes.

The following diagram illustrates Group Aggregation.



You can have any number of Archivers or Concentrators grouped together and form an aggregation group. The Archiver or Concentrator services in the group divide all the aggregated sessions between them based on the number of sessions defined in the Aggregate Max Sessions parameter.

For example, in an aggregation group containing two Archiver services or two Concentrator services with the Aggregate Max Sessions parameter set to 10,000, the services would divide the session between themselves as illustrated in the following table.

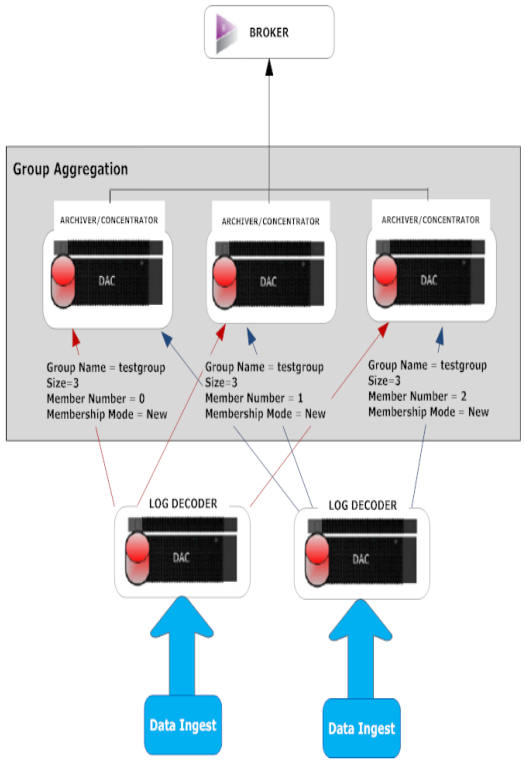
Archiver 0 or Concentrator 0	Archiver 1 or Concentrator 1
1 - 9,999	10,000 - 19,999
20,000 - 29,999	30,000 - 39,999
40,000 - 49,999	50,000 - 59,999

Configure Group Aggregation

Complete this procedure to configure multiple Archiver or Concentrator services as a group and share the aggregation tasks between them.

Prerequisites

Plan the network design for group aggregation. The following figure is an example of a group aggregation setup.



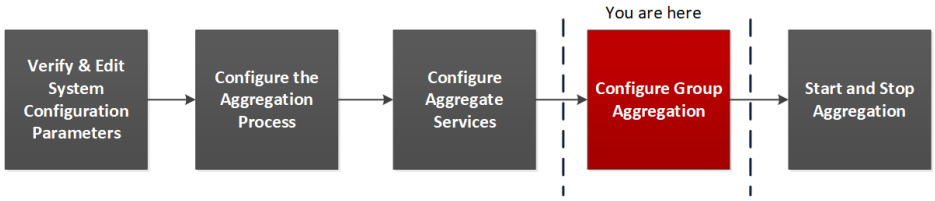
Ensure that you understand the Group aggregation parameters in the following table, and create a group aggregation plan.

Parameter	Description
Group Name	It determines the group to which the Archiver or Concentrator belongs. You can add any number of groups aggregating data from a Log Decoder. The Group Name parameter is used by the Log Decoder to identify which Archiver or Concentrator services are working together. All Archiver or Concentrator services in the group should have the same group name.
Size	It determines the number of Archiver or Concentrator services in the aggregation group.
Member Number	It determines the position of the Archiver or Concentrator in the aggregation group. For a group of size N, member number from 0 to N-1 must be set on each of the Archiver or Concentrators services in the aggregation group. For example: If the size of the aggregation group is 2, the member number of one of the Archiver or Concentrator service should be set to 0 and the member number of the other Archiver or Concentrator should be set to 1.
Membership Mode	There are two membership modes: • **New**: Adding a new Archiver or Concentrator service as a member to the existing aggregation group or creating an aggregation group. The Archiver or Concentrator service does not aggregate any existing sessions from the service as other members of the group would have already aggregated all the sessions on the service. This Archiver or Concentrator service will only aggregate new sessions as they appear on the service. • **Replace**: Replacing an existing aggregation group member. The Archiver or Concentrator will begin aggregation from the oldest session available on the service it is aggregating from.

Note: The Membership Mode parameter has an effect only when no sessions have been aggregated from the service. After some sessions are aggregated, this parameter has no effect.

Set up Group Aggregation

This workflow shows the procedures you complete to configure group aggregation.



Complete the following steps to set up group aggregation.

1. Configure multiple Archiver or Concentrator services in your environment. Make sure that you add the same Log Decoder as data source to all the services.

2. Perform the following on all the Archiver or Concentrator services that you want to be part of aggregation group:

   a. Go to **(Admin) > Services**.

   b. Select the Archiver or Concentrator service, and select > **View** > **Config**.

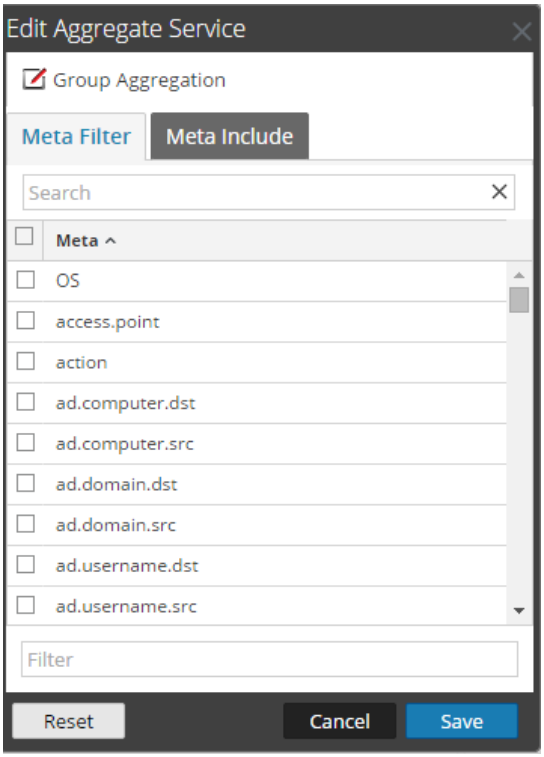
   The Service Config view of the Archiver or Concentrator is displayed.

   c. In the **Aggregate Services** section, select **Log Decoder**.

   d. Click Toggle Service to change the status of the Log Decoder to offline if it is online.

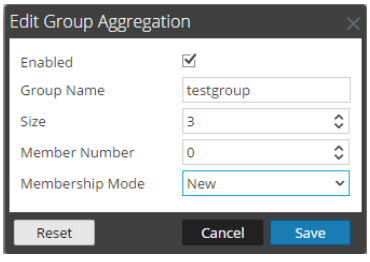
   e. Click .

   The **Edit Aggregate Service** dialog is displayed.



   f. Click Group Aggregation.

   The **Edit Group Aggregation** dialog is displayed.



   g. Select the **Enabled** checkbox and set the following parameters:

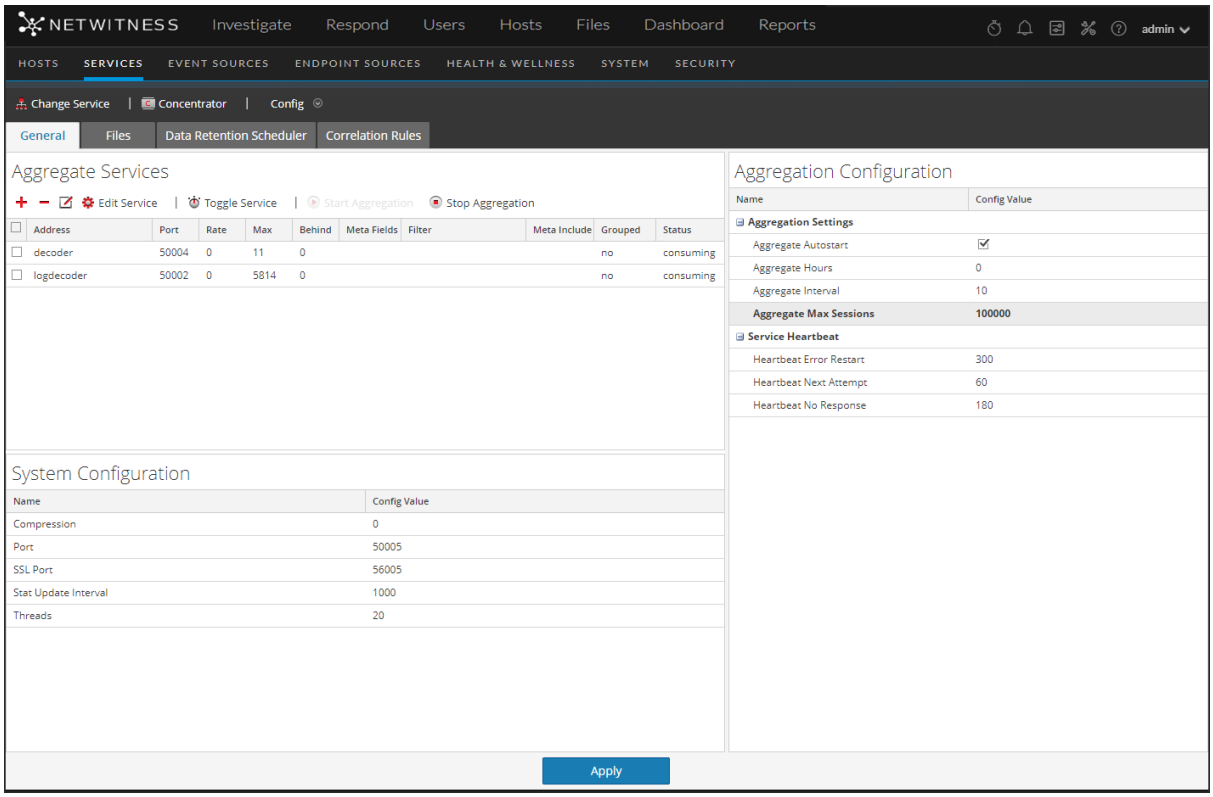
- In the **Group Name** field, type the group name.
- In the **Size** field, select the number of Archiver or Concentrator services in the aggregation group.
- In the **Member Number** field, select the position of the Archiver or Concentrator in the aggregation group.
- In the **Membership Mode** drop-down menu, select the mode.

h. Click **Save**.

i. In the Service Config view, click **Apply**.

j. Perform **Step b** to **Step i** on all other Archiver or Concentrator services that need to be part of group aggregation.

3. In the **Aggregation Configuration** section, set the **Aggregate Max Sessions** parameter set to **10000**.



New Health and Wellness

New Health and Wellness is an advanced monitoring and alerting system that provides insights on the operational state of the host and services in your deployment, and helps identify potential issues.

System Requirements

The following tables list the memory, disk, and CPU recommended for the New Health and Wellness based on the size of the deployment.

Note: The recommended values might differ when you install and try the new features and enhancements.

Caution: If the New Health and Wellness node is on a different subnet, you must open the respective NetWitness Platform hosts port. For more information, see "New Health and Wellness on Different Subnet" section in the Network Architecture and Ports.

Standalone Virtual Host

Minimum memory for a standalone virtual host is 16 GB.

Each NetWitness platform host writes 150 MB of New Health and Wellness metrics data into Elasticsearch data per day. For example, if you have 45 NetWitness Platform hosts then 6.6 GB of metrics data is written to Elasticsearch per day.

CPU	Memory
4 cores	16 GB

Physical Host

Deployment Size	Memory	CPU	DISK per day
Small (~5-10 hosts / 20-40 services)	16 GB	15%	1.5 GB
Medium (~150-200 hosts / 300- 400 services)	18 GB	15%	29 GB
Large (~250-300 hosts / 500-600 services)	22 GB	15%	44 GB

Based on the resources available, you can deploy the 11.7 New Health and Wellness feature on any one of the following, listed in the order of preferred deployment method with most preferred first:

- Standalone virtual host (Most preferred recommendation to ensure no performance impact on any other functionality of deployed nodes)

- Physical host:
  - Broker
  - Admin Server
  - ESA

Installing New Health and Wellness enables all hosts in your deployment to start sending metrics to monitor New Health and Wellness. After you deploy New Health and Wellness , see the "Monitor New Health and Wellness " topic in the *System Maintenance Guide* for instructions on how to configure and use this feature.

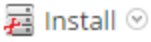
Please direct any New Health and Wellness feedback to nw.health.wellness.feedback@rsa.com.

After you provision the New Health and Wellness host, complete the following steps to install the **New Health and Wellness** service category on the provisioned host.

1. Log in to NetWitness and go to  (Admin) > **Hosts**.
   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background.

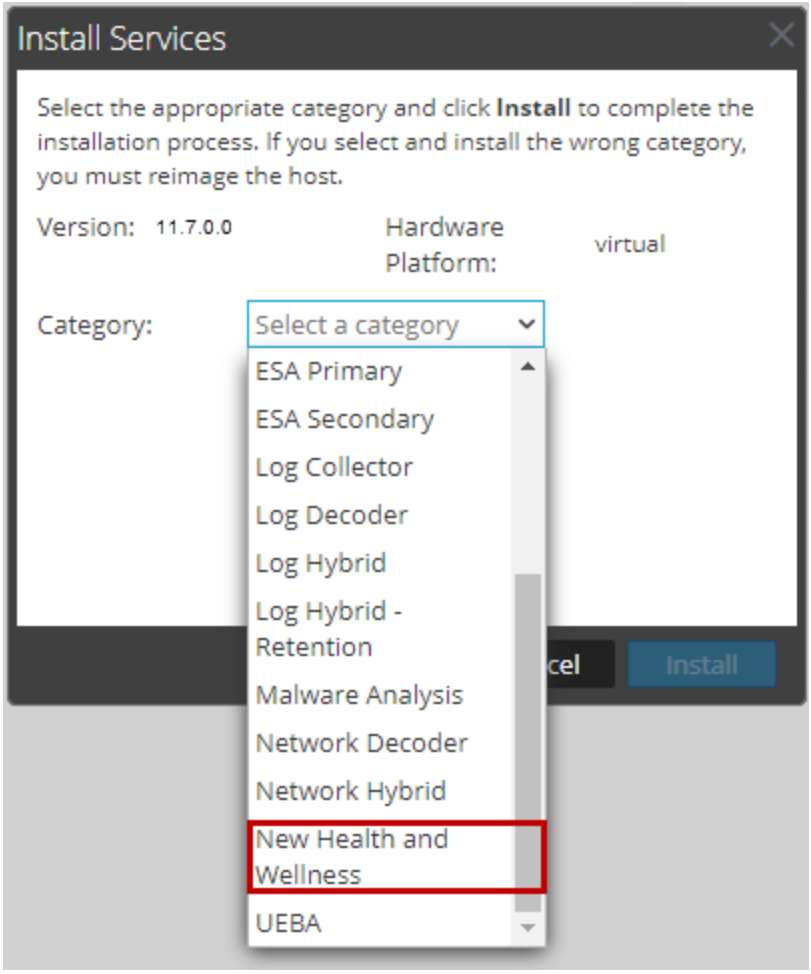
   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   Note: If you are not installing New Health and Wellness on a standalone virtual host ignore step 2.

2. Select the host in the **New Hosts** dialog and click **Enable**.
   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.
3. Select that host on which New Health and Wellness should be installed in the **Hosts** view (for example, **New Health and Wellness**) and click 

.

The **Install Services** dialog is displayed.



4. Select **New Health and Wellness** in **Category** and click **Install**.
5. Refresh all hosts to write to elastic search.
   a. SSH to the NW Server host.
   b. Run the following commands.
      `nw-manage --refresh-host --host-all`
      This may take few minutes for the changes to take effect based on the number of hosts in your deployment.

Note: (For Standalone Virtual host only) After you review your initial datastore configuration, you may determine that you need to add a new volume. For information on adding a new volume see “Add New Volume and Extend Existing File Systems” topic in the *Virtual Host Installation Guide*.

Note: After you have installed New Health and Wellness, for some reason, if you want to uninstall New Health and Wellness, you must refer to "Uninstall New Health and Wellness" in the *Upgrade guide*.

Hybrid Categories on Series 6 (R640) Hardware

You can install Hybrid Categories such as Log Hybrid and Network (Packet) Hybrid service categories on a Series 6 (R640) Physical host. This gives you the ability to attach multiple PowerVault external storage devices to the Series 6 (R640) Physical host.

NW Server Deployment on ESA Hardware

You now have the option to deploy the NW Server host on Series 6 Analytics hardware. The Series 6 Analytics Hardware has more memory and storage capacity than the standard Core appliance on which NW Server has typically been deployed. This results in better overall responsiveness and larger retention capacity for Report Engine.

Note: You can install the NW Server on ESA hardware, but you cannot co-locate any ESA services (categories) with the NW Server on this hardware.

Second Endpoint Server

Complete the following procedure to deploy a second Endpoint Server.

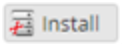
1. Set up a new host in NetWitness Platform.
   - For a physical host, complete steps 1 to 16 in "Install NetWitness Platform" under "Installation Tasks" in the *Physical Host Installation Guide for NetWitness Platform 11.7*.
   - For a virtual host, complete steps 1 to 6 in "Step 4. Install NetWitness Platform" in the *Virtual Host Installation Guide for NetWitness Platform 11.7*.

     Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
2. SSH to the host that you set up in step 1.
3. Submit the following command string.
   ```
   mkdir -p /etc/pki/nw/nwe-ca
   ```

   Note: You do not need to modify permissions.
4. Copy the following two files from the previously deployed endpoint server to the new/second endpoint server:
   ```
   /etc/pki/nw/nwe-ca/nwerootca-cert.pem
   /etc/pki/nw/nwe-ca/nwerootca-key.pem
   ```
5. nstall Endpoint on the host.
   a. Log in to NetWitness Platform and go to  (Admin) > **Hosts**.
   The **New Hosts** dialog is displayed with the Hosts view grayed out in the background.

      Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the new host in the **New Hosts** dialog and click **Enable**.
   The New Hosts dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the Hosts view (for example, Endpoint Server II) and click 

.
   The **Install Services** dialog is displayed.

   d. Select **Endpoint** in **Host Type** and click **Install**.

Warm Standby NW Server Host

The Warm Standby NW Server duplicates the critical components and configurations of your active NW Server host to increase reliability.

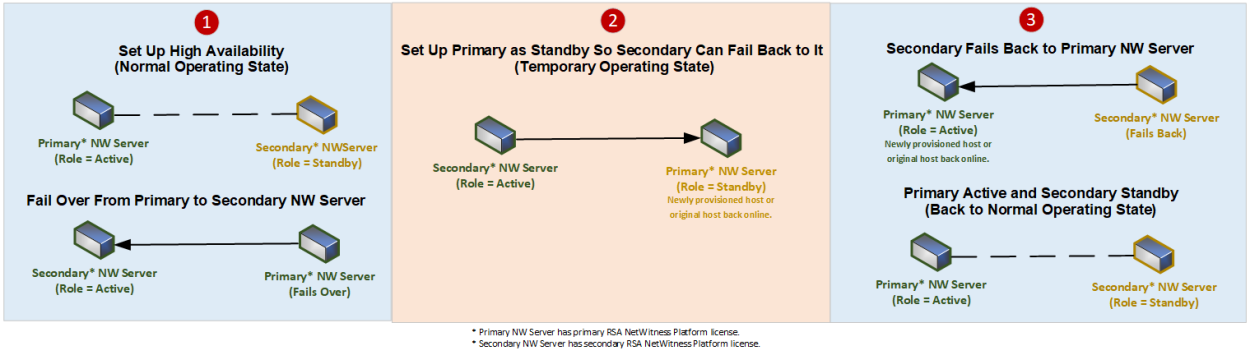
A secondary NW Server remains in the standby role and, when configured, receives backups of the primary NW Server in the active role at regular intervals. If the primary NW Server fails (goes offline), the fail-over procedure must be executed allowing the secondary NW Server to assume the active role.

In version 11.5 and later, we support a secondary NW Server that has a different IP address from the primary NW Server. Having the same IP address for both the primary NW Server and secondary NW Server is no longer necessary.

When you set up a secondary NW Server as a Warm Standby, a failure or scheduled switch from the primary NW Server to the secondary NW Server is referred to as a fail-over. You fail back to return to the normal operating state (that is, primary NW Server in the active role and the secondary NW Server in the standby role).

The following diagram illustrates the fail-over and fail-back process.

1. Set up secondary NW Server as standby (initial setup). This is the normal operating state.
2. The primary NW Server fails over to the secondary NW Server. After the fail-over, get the primary NW Server back online and set it up in the standby role. This is a temporary operating state.
3. Fail the secondary NW Server back to the primary. The primary NW Server is back to the active role and secondary is back to the standby role. This is the normal operating state.



Note: When you set up the secondary NW Server, follow the same administrative procedures, for example, for upgrade and maintenance, as the procedures for the primary active NW Server.

Procedures

Complete the following task to set up a secondary NW Server in the standby role for fail-over:

- Set up a secondary NW Server in the standby role.

Complete the following tasks when required to maintain high availability:

- Fail over the primary NW Server to secondary NW Server.
- Fail back the secondary NW Server to primary NW Server.

Planned Fail-Over Scenario

This scenario occurs when you schedule a fail over (see **Planned Fail-Over** under step 3 in the Fail Over primary NW Server to Secondary NW Server procedure). You should not need do anything after the fail-over completes.

Required Fail-Over Scenario without Hardware Replacement

This scenario occurs when the primary NW Server fails (see *Required Fail-Over* under step 3 in the Fail Over Primary NW Server to Secondary NW Server topic), but you are able to recover it easily without re-imaging (for example, the active NW Server has corrupt or insufficient RAM). You do not need to run the `nwsetup-tui` and you do not need to contact NetWitness Customer Support to reestablish correct licensing when:

1. The active (primary NW Server) fails over to the Standby (secondary NW Server) and that secondary host temporarily assumes the role of the active NW Server.
2. You fix the problem with the primary NW Server (for example, install new RAM) and fail back to it from the secondary host.

Required Fail-Over Scenario with Hardware Replacement

This scenario occurs when the active NW Server completely fails and the hardware requires replacement, for example you receive a Return Merchandise Authorization (RMA). You need to run reconfigure the host with the `nwsetup-tui` and contact NetWitness Customer Support to reestablish licensing. If you choose to rebuild the replacement host as a temporary standby (for example, until your scheduled fail-back occurs), you must answer "**Yes**" to the **Standby Host Recovery Mode** `nw-setup-tui` prompt when configuring this temporary standby for failing back (see step 4 in the Set Up Secondary NW Server in Standby Role procedure for the context of this prompt).

Set Up Secondary NW Server in Standby Role

1. Before you install a secondary NW Server host for the standby role, make sure that:
   a. The primary NW Server is running 11.7.
   b. All component hosts are running 11.7
      If you are:
      - Installing NetWitness Platform 11.7, follow the instructions in the *NetWitness Platform Physical Host Installation Guide for Version 11.7*.
      - Updating from 11.x to 11.7, follow the instructions in *NetWitness Platform Update Guide for Version 11.x to 11.7*.
        Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
2. Create a base image on the secondary NW Server:
   a. Attach media (ISO) to the host.
      See the *NetWitness Platform Build Stick Instructions* for more information.
      - Physical media - use the ISO to create bootable flash drive media the **Etcher®** or another suitable imaging tool etch an Linux file system on the USB drive. See the *NetWitnessBuild Stick Instructions* for information on how to create a build stick from the ISO. Etcher is available at: https://etcher.io.
      - iDRAC installations - the virtual media type is:
        - **Virtual Floppy** for mapped flash drives.
        - **Virtual CD** for mapped optical media devices or ISO file.
   b. Log in to the host and reboot it.

```
login: root
Password:
Last login: Tue Sep 19 13:27:15 on tty1
[root@saserver ~]# reboot
```

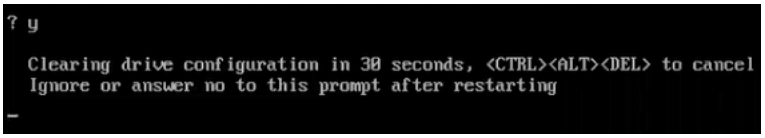
c. Select **F11** (boot menu) during reboot to select a boot device and boot to the connected media. After some system checks during booting, the following **Welcome to NetWitness Platform 11.7** installation menu is displayed. The menu graphics will render differently if you use a physical USB flash media.

d. Select **Install NetWitness Platform 11.7** (default selection) and press **Enter**.
The Installation program runs and stops at the **Enter (y/Y) to clear drives** prompt that asks you to format the drives.



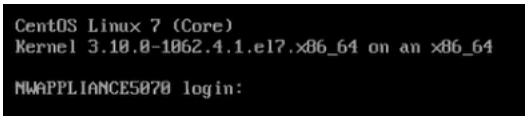
Caution: You must respond y or Y to this prompt even if the host does not have an internal RAID configuration or the installation will fail.

e. Type **Y** to continue.
The default action is No, so if you ignore the prompt and it will select No in 30 seconds and will not clear the drives. The **Press enter to reboot** prompt is displayed.



The system displays the all installation tasks it is performing. This can take a minute or so.After it completes the tasks, the installation program reboots the host.

Caution: Do not reboot the attached media (media that contains the ISO file, for example a build stick).

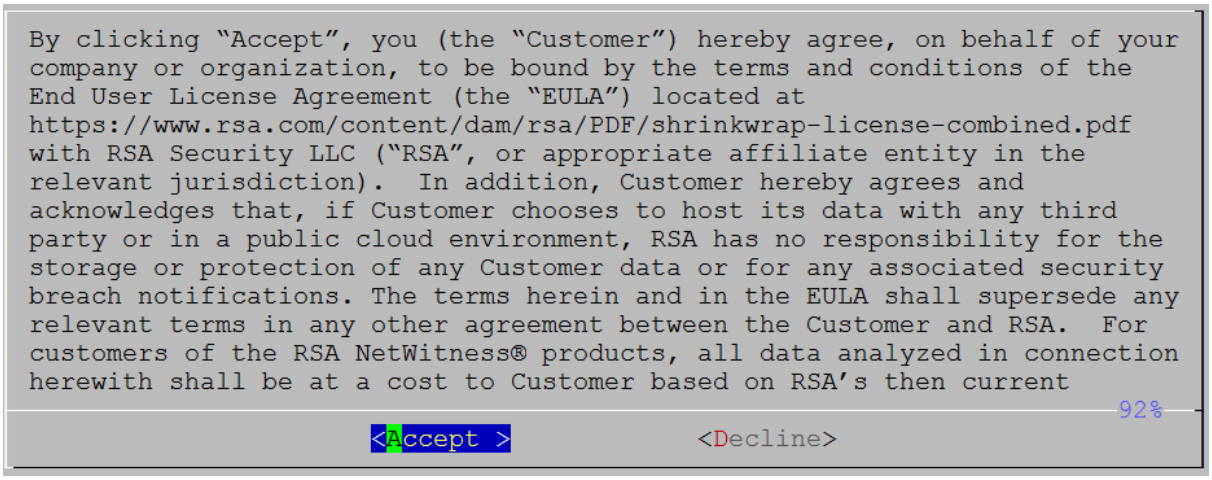


f. Log in to the host with the `root` credentials.

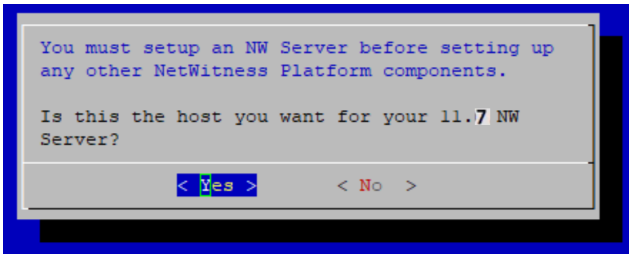
2. Run the `nwsetup-tui` command.

Note: 1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, use the Tab key to move to and from commands (such as **<Yes>**, **<No>**, **<OK>**, and **<Cancel>**. Press **Enter** to register your command response and move to the next prompt.
2.) The Setup program adopts the color scheme of the desktop or console you use access the host.
3.) During the Setup program, when you are prompted for the network configuration of the host, be sure to specify the same network configuration that was used for the original installation of 11.x on this host (it must be exactly the same).

This initiates the `nwsetup-tui` (Setup program) and the EULA is displayed.

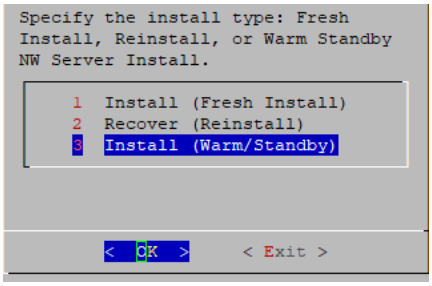


3. Tab to **Accept** and press **Enter**.
   The **Is this the host you want for your 11.7 NW Server** prompt is displayed.

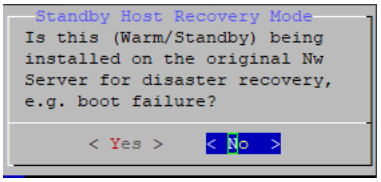


   Your response to this prompt identifies a host as either the primary or secondary during a fresh install (and the selected response stays constant regardless of the current or future role, that is active or standby of the host).

4. Tab to **Yes** and press **Enter**.
   The **Install or Recover** prompt is displayed.



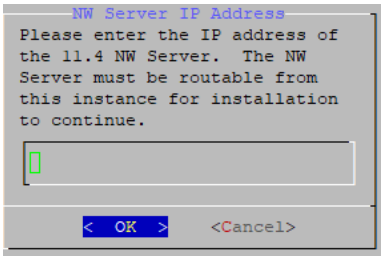
5. Tab to **3 Install (Warm Standby)** and press **Enter**.
   The Standby Host Recovery Mode prompt is displayed.



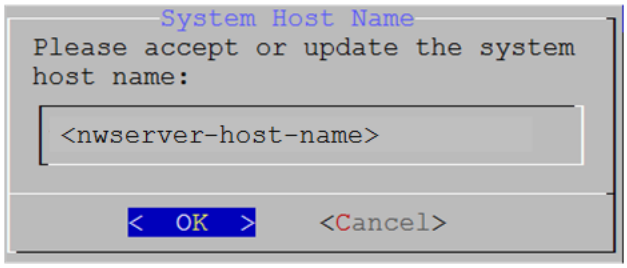
6. Tab to:

   - **No** and press **Enter** to set up a secondary NW Server with the standby role (most common scenario).
   - **Yes** and press **Enter** to set up a host that was previously used as a primary NW Server with the standby role so you can execute a fail-over and fail-back (less common scenario).

     The NW Active Server IP Address prompt is displayed.



7. Type the IP Address of the NW Server in the active role, tab to **OK**, and press **Enter**.
   The **Host Name** prompt is displayed



Caution: If you include "." in a host name, the host name must also include a valid domain name.

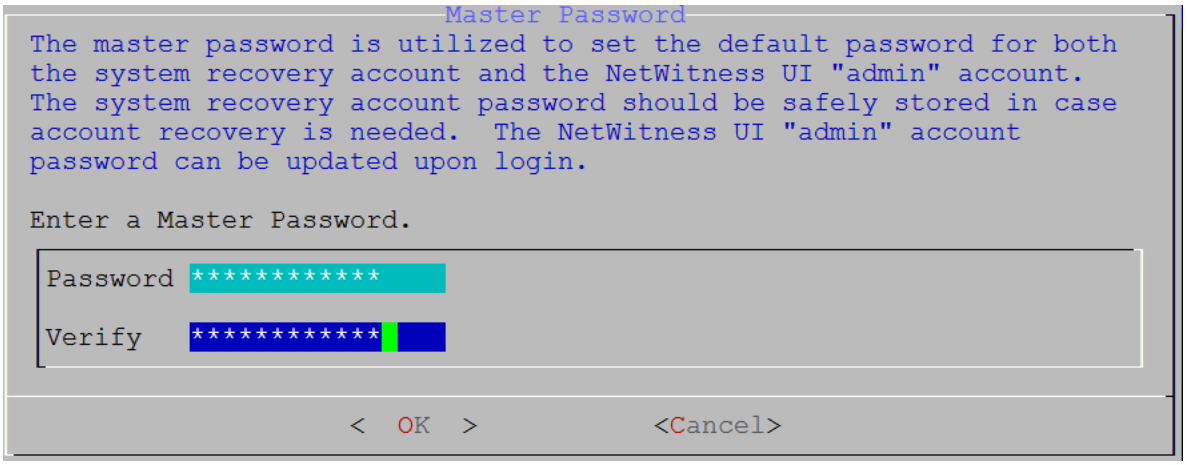
8. Press **Enter** if want to keep this name. If not edit the host name, tab to **OK**, and press **Enter** to change it.
   The **Master Password** prompt is displayed.

   Note: You must use the same Master and Deploy Admin credentials fot the Warm Standby NW Server Host that you used for the Active NW Server Host.

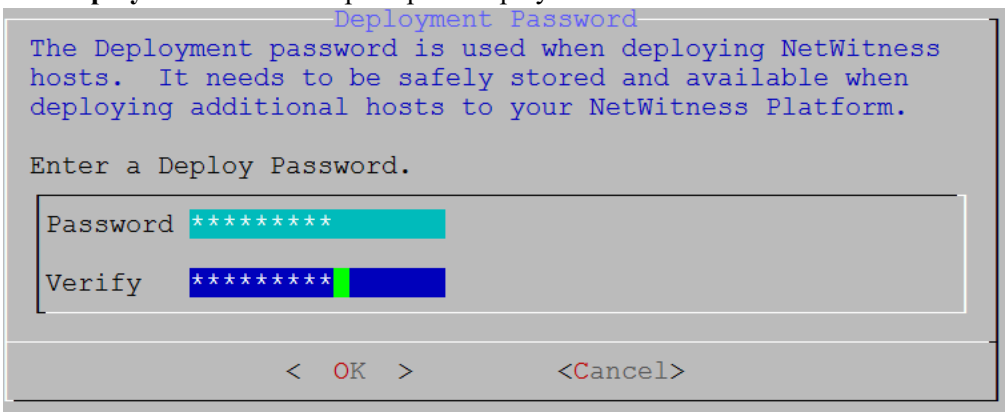
   The following list of characters are supported for Master Password and Deployment Password:

   - Symbols: **! @ # % ^ +**
   - Lowercase Characters: **a-z**
   - Uppercase Characters: **A-Z**

   No ambiguous characters are supported for Master Password and Deployment Password. For example: space { } [ ] ( ) / \ ' " ` ~ ; : .< > -

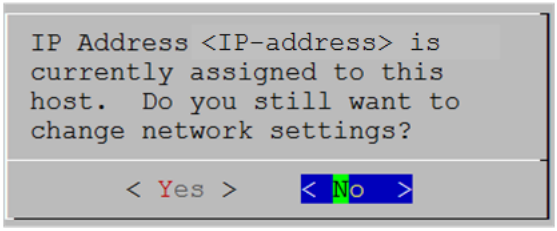


9. Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.
   The **Deployment Password** prompt is displayed.



10. Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**. One of the following conditional prompts is displayed.

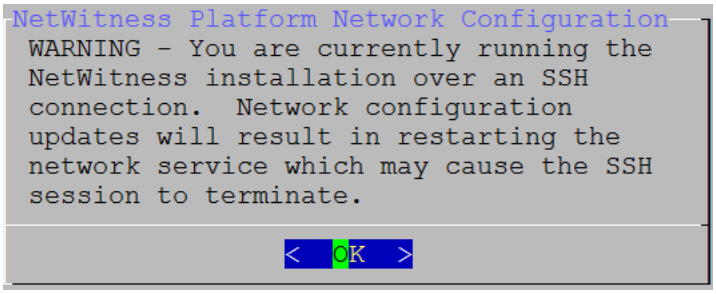
    - If the Setup program finds a valid IP Address for this host, the following prompt is displayed.



      Press **Enter** if you want to use this IP and avoid changing your network settings. Tab to **Yes** and press **Enter** if you want to change the IP configuration found on the host.

    - If you are using an SSH connection, the following warning is displayed.

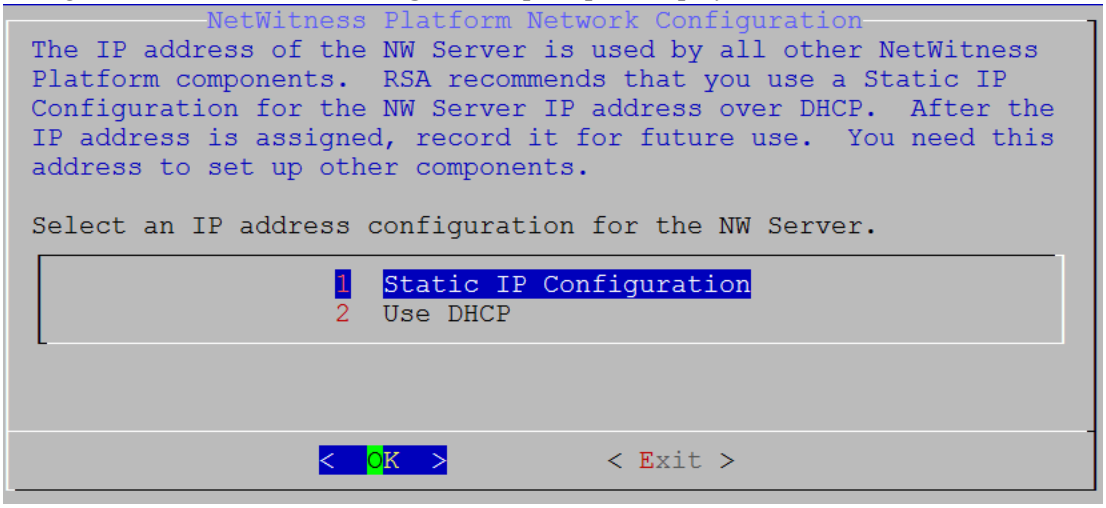
> **Note:** If you connect directly from the host console, the following warning will not be displayed.



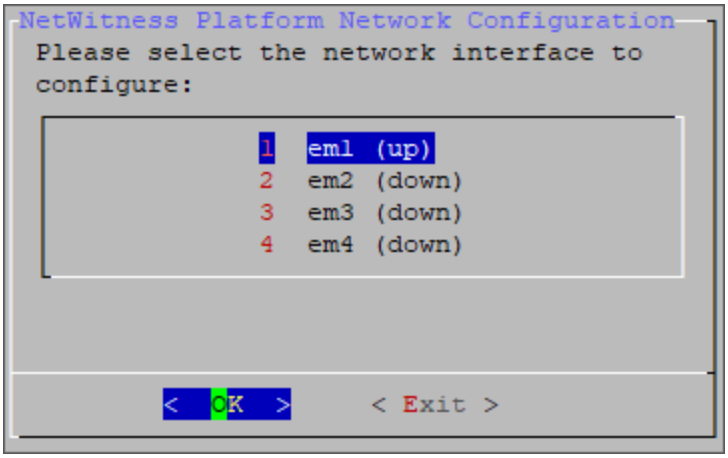
      Press **Enter** to close warning prompt.

    - If the Setup Program found an IP configuration and you chose to use it, the **Update Repository** prompt is displayed. Go to step 12 to and complete the installation.

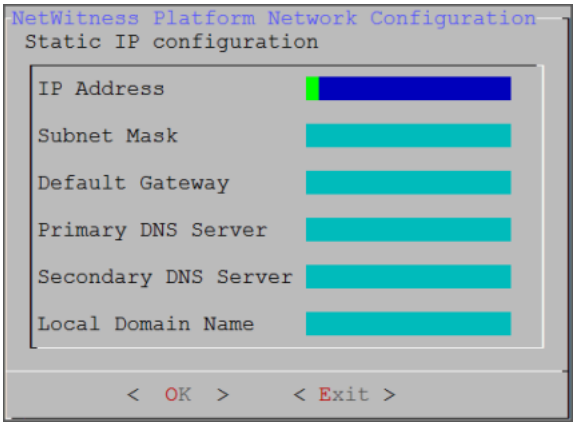
    - If the Setup Program did not find an IP configuration or if you chose to change the existing IP configuration, the **Network Configuration** prompt is displayed.



11. Tab to **OK** and press **Enter** to use **Static IP**.
    If you want to use DHCP, down arrow to **2 Use DHCP** and press **Enter**.
    The **Network Configuration** prompt is displayed.



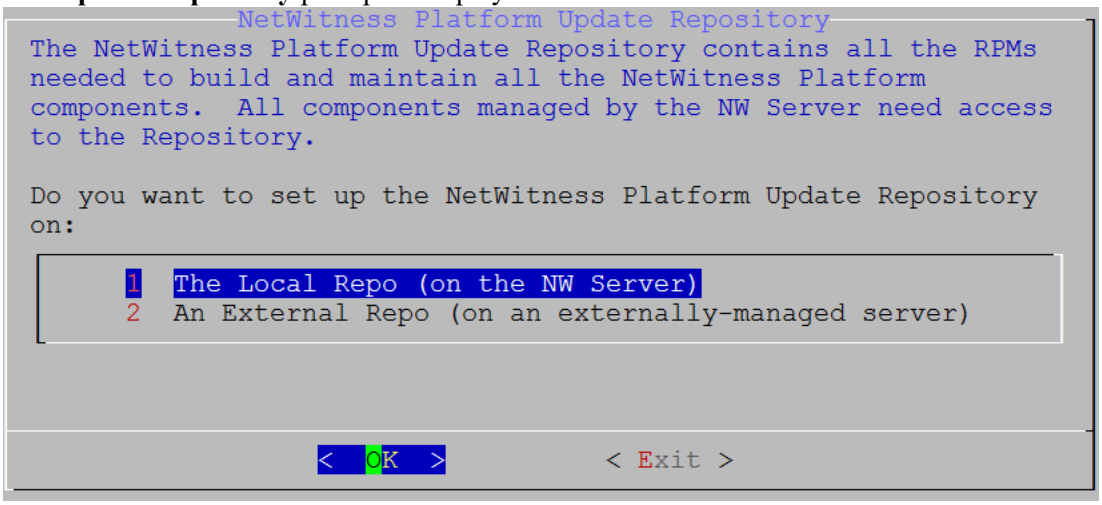
12. Down arrow to the network interface you want, tab to **OK**, and press **Enter**. If you do not want to continue, tab to **Exit.**
    The following **Static IP Configuration** prompt is displayed.



13. Type the configuration values (using the down arrow to move from field to field), tab to **OK**, and press **Enter**. If you do not complete all the required fields, an `All fields are required` error message is displayed (**secondary DNS Server** and **Local Domain Name** fields are not required). If you use the wrong syntax or character length for any of the fields, an `Invalid <field-name>` error message is displayed.

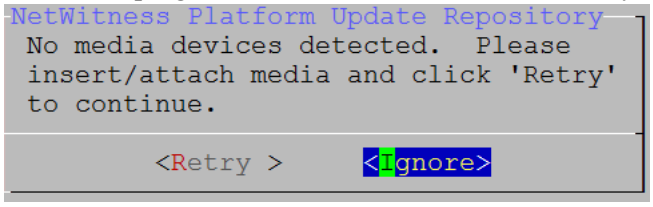
    Caution: If you select **DNS Server**, make sure that the DNS Server is correct and the host can access it before proceeding with the installation.

    The **Update Repository** prompt is displayed.

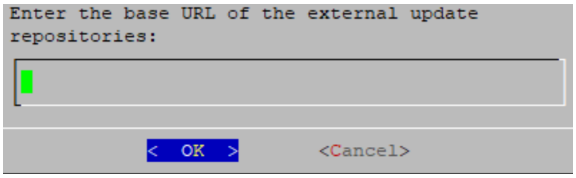


14. Press **Enter** to choose the **Local Repo** on the NW Server.
    If you want to use an external repo, down arrow to **External Repo**, tab to **OK**, and press **Enter**.

    - If you select **1 The Local Repo (on the NW Server)** in the Setup program, make sure that you have the appropriate media attached to the host (media that contains the ISO file, for example a build stick) from which it can install NetWitness 11.7
    - .0.0. If the program cannot find the attached media, you receive the following prompt.



- If you select **2 An External Repo (on an externally-managed server)**, the UI prompts you for a URL. The repositories give you access to NetWitness updates and CentOS updates. Refer to "Appendix B. Create an External Repo" in the *Physical Host Installation Guide* for instructions on how to create this repo and its external repo URL so you can enter it in the following prompt.

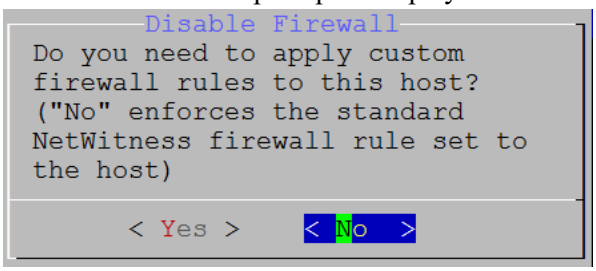


Enter the base URL of the NetWitness external repo and click **OK.** The **Start Install** prompt is displayed.

See "Set Up an External Repository with RSA and OS Updates" under "Hosts and Services Procedures" in the *NetWitness Platform Hosts and Services Getting Started Guide* for instructions.

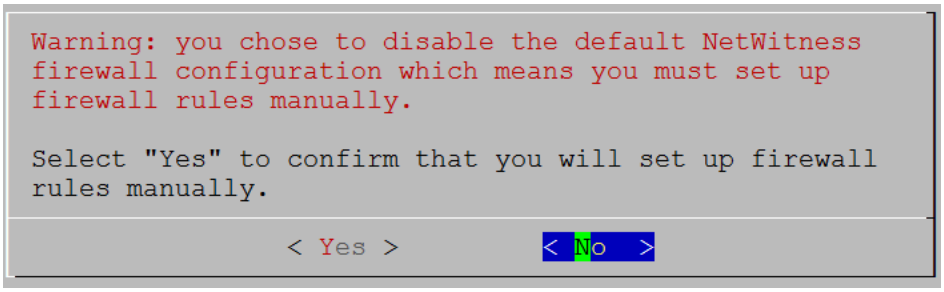
Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

The Disable firewall prompt is displayed.

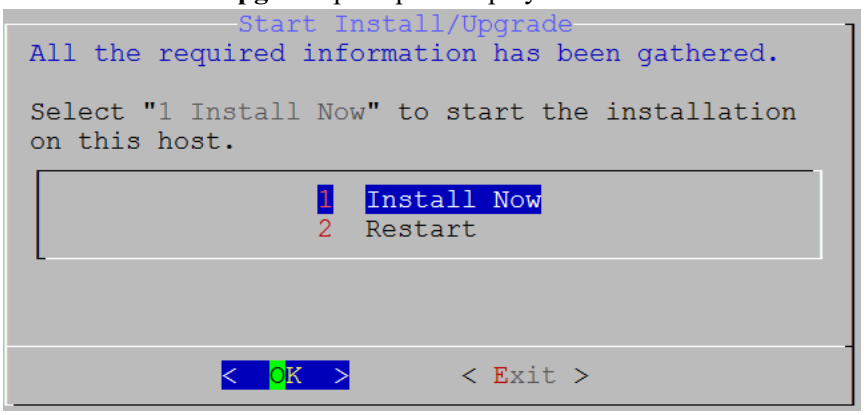


15. Tab to **No** (default), and press **Enter** to use the standard firewall configuration. Tab to **Yes**, and press **Enter** to disable the standard firewall configuration.

If you select **Yes**, confirm your selection(select **Yes** again) or select **No** to use the standard firewall configuration.

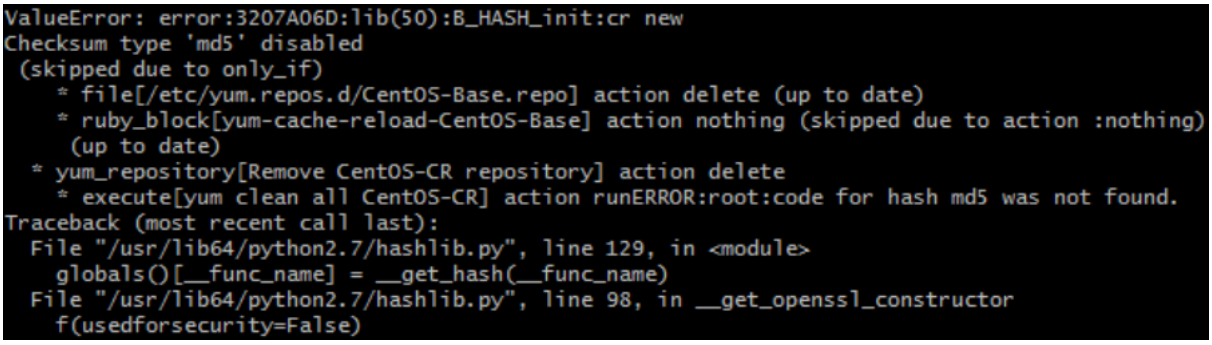


The **Start Install/Upgrade** prompt is displayed.



16. Press **Enter** to install 11.7 on the NW Server.
    When **Installation complete** is displayed, you have installed the 11.7 NW Server on this host.

> **Note:** Ignore the hash code errors similar to the errors shown in the following figure that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.



17. License the secondary NW Server.

    a. Log in to the secondary NW Server User Interface, click  (Admin) > **System** > **Info**, and note the **License Server ID** under **Version Information**.

    b. SSH to the primary NW Server.

    c. Edit the `/opt/netwitness/flexnetls/local-configuration.yaml` file and add the `back up hostid` (that is, the **License Server ID**).
    This is an example of the section of the `local-configuration.yaml` file before you add the **License Server ID**.

```
# Hostid of the backup server, if in fail over configuration.
#backup-hostid:
```

    This is an example of the section of the `local-configuration.yaml` file after you add the MAC address (for example, `000c2918c80d`) of the Warm Standby NW Server Host.

```
# Hostid of the backup server, if in fail over configuration.
backup-hostid: "000c2918c80d"
```

d. Restart the fneserver service.
```
systemctl restart flexnetls-RSALM
```

e. (Conditional) If your NetWitness Platform deployment is prohibited from accessing the Internet (Air Gap), you must:

i. Download the capability request from NetWitness Platform User Interface.

ii. Upload the request to FNO.

iii. Upload the response from FNO to the NetWitness Platform User Interface.

18. Schedule the backup of the primary NW Server and the copying of this backed-up data to the secondary NW Server.

a. SSH to the primary NW Server.

b. Submit the following commands.
```
/opt/rsa/saTools/bin/schedule-standby-admin-data-sync -d -i <warm-standby-admin-server-ip>
```

Note: If the Warm Standby server becomes active, and if it will be accessed by any NW host with a NAT IP address, the NAT IP address must be added to the command:
```
/opt/rsa/saTools/bin/schedule-standby-admin-data-sync -d -i <warm-standby-admin-server-ip>-p<warm-standby-NAT-admin-server-ip>
```

This backs up the primary NW Server data and copies the backup archive file to the secondary NW Server daily for future fail-over use. It also schedules the backup and copy to execute on a daily basis. You can display help for the `schedule-standby-admin-data-sync` script with the following command string.
```
/opt/rsa/saTools/bin/schedule-standby-admin-data-sync --help
```
This returns the following help to which you can refer to customize the host data backup (such as backup frequency).
```
Schedule Data Synch between AdminServer and Standby AdminServer
Script also executes a synchronization each time.
Usage:
schedule-standby-admin-data-sync command [options]
```

```
Commands:

-h, --help                  Display Help
-d, --daily                 Schedule daily data
                            synchronization
-w, --weekly                Schedule weekly data
                            synchronization
-c, --custom <crontab       Schedule custom data
formatted>                  synchronization
                            i.e. to schedule for
                            midnight on 1st and
                            10th of the month:
                            '0 0 1,10 * *'
-i, --standby-ip <IP        (required) IP
address>                    address of standby
                            NW Server
-p, --standy-ip-public <NAT (optional) NAT-based
IP address>                 IP address of
                            standby NW Server
-v, --verbose               Enable verbose
                            output
```

Fail Over Primary NW Server to Secondary NW Server with Same IP Address

Initially, the primary NW Server fails over to the secondary NW Server. When the primary NW Server is back up, the secondary NW Server fails over to the primary NW Server, and that is referred to as a fail-back. When it is possible for the secondary NW Server to have the same IP address as the primary NW Server after failover, complete the following procedure to fail over from the primary NW Server to the secondary NW Server.

1. SSH to the secondary NW Server.

IMPORTANT: If you have installed New Health & Wellness on Admin Server, perform the following:
- Install New Health & Wellness on Standby NW server.
- Back up the Search category (New Health and Wellness) from Primary NW server to standby NW server using the NRT:

```
nw-recovery-tool --export --dump-directory <Backup_directory_path> --category Search
```

You must manually copy files from Primary NW server to Standby NW server after running the above command.

2. Run the `nw-failover` script with the `--make-active` argument:
   ```
   nw-failover --make-active
   ```
3. Complete the following steps for planned and required fail-over.
   a. SSH to the primary NW Server.
   b. Run the fail-over script with the `--make-standby` argument to assign the standby role to the primary NW Server:
      ```
      nw-failover --make-standby
      ```
   c. Shut down the primary NW Server.

   Note: If you have a catastrophic failure, you may need to provision a new host or re-image the primary NW Server and complete the Set Up secondary NW Server in Standby Role procedure for this host to create a new primary NW Server, so that you can fail back to it.
4. Validate that the component hosts have connected to the new active NW Server by running the following command on the new active NW Server:
   ```
   nw-manage --check-hosts-status
   ```

   Note: Do not continue to the next step until all component hosts have successfully connected to the new active NW Server. Rerun `nw-manage --check-hosts-status` as needed to continue verification.
5. Follow the IP address change procedures for the now-active NW Server in "Change Host Network Configuration" in the *System Maintenance Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Note: If the primary NW Server is going to remain online, you must update its IP address so that it no longer conflicts with the Warm Standby NW Server IP address when it is changed. Follow the IP address change procedures for the now-active NW Server in "Change Host Network Configuration" in the *System Maintenance Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

6. Shut down the primary NW Server, or reboot the primary NW Server and set it up as the new secondary NW Server by completing the steps in Set Up Secondary NW Server in Standby Role for this host to create a new primary NW Server, so that you can fail back to this host.

IMPORTANT: After successful fail over to Standby NW server, restore the Search category (New Health & Wellness) using NRT:
`nw-recovery-tool --import --dump-directory <Backup_directory_path> --category Search`

Fail Over Primary NW Server to Secondary NW Server with Different IP Address

If your secondary NW Server will maintain a different IP address from your primary NW Server after a fail-over event, for example, if the secondary NW Server is located in a different data center from the primary NW Server, follow these steps to fail over from the primary NW Server to the secondary NW Server.

1. SSH to the secondary NW Server.

IMPORTANT: If you have installed New Health & Wellness on Admin Server, perform the following:
- Install New Health & Wellness on Standby NW server.
- Back up the Search category (New Health and Wellness) from Primary NW server to standby NW server using the NRT:
`nw-recovery-tool --export --dump-directory <Backup_directory_path> --category Search`
You must manually copy files from Primary NW server to Standby NW server after running the above command.

2. Run the failover script:
   `nw-failover --make-active`
3. Complete the following steps for planned and required fail-over.
   a. SSH to the primary NW Server.
   b. Run the fail-over script with the `--make-standby` argument to assign the standby role to the primary NW Server:
      `nw-failover --make-standby`
   c. Shut down the primary NW Server.

      Note: If you have a catastrophic failure, you may need to provision a new host or re-image the primary NW Server and complete the Set Up secondary NW Server in Standby Role procedure for this host to create a new primary NW Server, so that you can fail back to it.

4. If you are running in a mixed-version environment, log in to each component host that is still on a version prior to 11.7 and execute the following:
   - `netconfig --update-dns --dns <active-nw-server-ip-address>`
   - `sed -Ei 's/^master:.*/master: <active-nw-server-ip-address>/g' /etc/salt/minion`
   - `systemctl restart salt-minion`

   Note: If you have a component host that uses a NAT IP address to reach the NW Server, substitute the `<NAT-IP-address>` for the `<active-nw-server-ip-address>` in both bullet items above.

5. Shut down the primary NW Server, or reboot the primary NW Server and set it up as the new secondary NW Server by completing the steps in Set Up Secondary NW Server in Standby Role for this host to create a new primary NW Server, so that you can fail back to this host.

IMPORTANT: After successful fail over to Standby NW server, restore the Search category (New Health & Wellness) using NRT:
`nw-recovery-tool --import --dump-directory <Backup_directory_path> --category Search`

Note: If the secondary NW Server had any custom certificates, you must re-apply them after failover. For information about how to re-apply custom certificates, see "(Optional) Use a Custom Server Certificate" in the *System Security and User Management Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Note: After upgrading the NW Server host or a component host to 11.7 or later version, review the contents of the `/etc/hosts.user` file for any obsolete host entries. The `/etc/hosts.user` file contains system and user-generated entries that are not managed by NetWitness Platform. However, entries from `/etc/hosts.user` are merged with NetWitness Platform-generated host mappings to create and update `/etc/hosts`. To avoid conflicts with NetWitness Platform-generated mappings, and to avoid generating connectivity errors resulting from an IP address change, RSA recommends that you remove any entries in `/etc/hosts.user` that include a non-loopback IP address of a NetWitness Platform host. After updating `/etc/hosts.user`, you must refresh the system by running the following command:
`nw-manage --refresh-host --host-key <ID, IP, hostname or display name of host>`

Note: While changing a host's IP address or during failover, component hosts can become disconnected from NW Server hosts. Follow these steps to reconnect a host system to its NW Server system.
1. Log in to the component host using SSH or the console.
2. Run the command `nw-manage --override-nws-ip --ipv4 <current IP address of the NW Server.`
When this command completes, the component host is reconnected to the NW Server at the specified IP address.

Follow the steps in the sections that apply to your environment.

- SSO
- Reporting Engine
- UCF
- PAM
- ECAT
- RSA NetWitness Orchestrator (By Demisto)
- Audit Logging
- Health and Wellness
- Malware Analysis
- Windows Legacy Collection

SSO

Update Configuration for Single Sign-On

Note: You must disable SSO configurations ONLY when NW Server IP is changed.

When the host network is configured with a new IP address, the SSO configurations also must be updated.

To do this:

1. Disable the SSO configuration using `nw-shell` after failover from new IP.
   To resolve this issue you must disable SSO manually, using the following commands:
   a. SSH to `admin server node`.
   b. Connect to `nw-shell`.
   c. Connect to admin server service using the `connect --service admin-server` command.
   d. Log in to admin server using the `login` command.
   e. Enter the admin username and password.
   f. Execute the following commands:
      - `cd /rsa/security/authentication/web/saml/sso-enabled`
      - `set false`
      - `logout`
      - `exit`

- `systemctl restart rsa-nw-admin-server`



2. Change the host IP address to the new IP.
3. Generate the new metadata and reupload it in ADFS. For more information, see the "Configure SAML 2.0 provider settings for portals" topic in the Microsoft documentation.

For more information, see the "Troubleshooting" topic in the *System Security and User Management Guide.*

Reporting Engine

Update Configuration for Reporting Engine

Note: You must update the Reporting Engine configurations ONLY when NW Server IP is changed.

When the host network is configured with a new IP address, you must update and verify the Reporting Engin configurations. The hostname for NetWitness configurations under the Output Actions must be updated with the new IP.

To manually configure the new IP, perform the following steps:

1. Log in to NetWitness Platform.
2. Navigate to  (Admin) > **Services > Reporting Engine > View > Config**.
3. Click the **Output Actions** tab.

4. Add the new IP address in the **Hostname** field.
5. Click **Apply**.

UCF

To enable UCF to communicate with NetWitness Platform:

1. On the UCF server, execute the `runConnectionManager.bat` file (the same file that is used for adding connection details).
2. Select **Option #2, Edit endpoints**.
3. Select the NW Server connection from the options that are displayed.
4. When you are prompted for Host Address (the old IP address is shown in parentheses) enter the new IP address.

   Note: Do not change any other setting.

PAM

If you have PAM configured, after the failover, you must configure the system again using the instructions in the "Configure PAM Login Capability" topic in the *System Security and User Management Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

ECAT

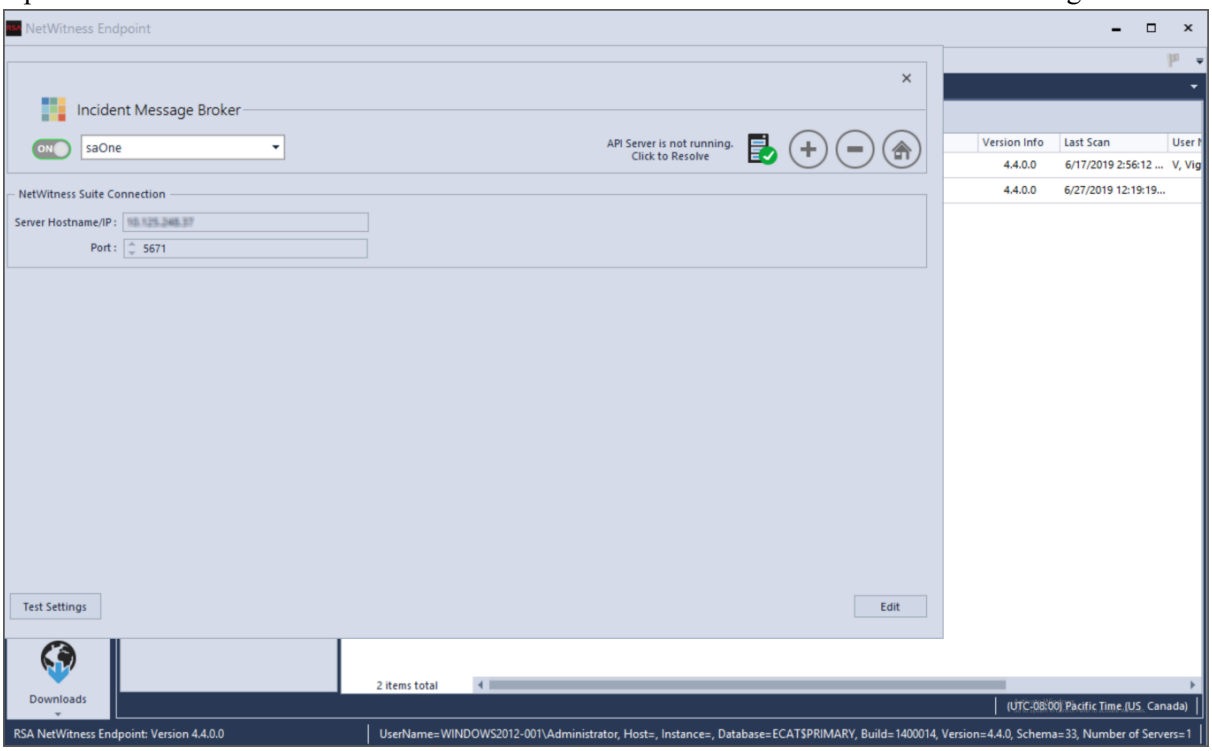
Update the following services:

- Incident Message Broker
- NetWitness Suite
- Syslog Server Settings

Incident Message Broker

1. Log in to the NetWitness Endpoint user interface and go to **Configure** > **Monitoring and External Components Configuration** > **Incident Message Broker**.

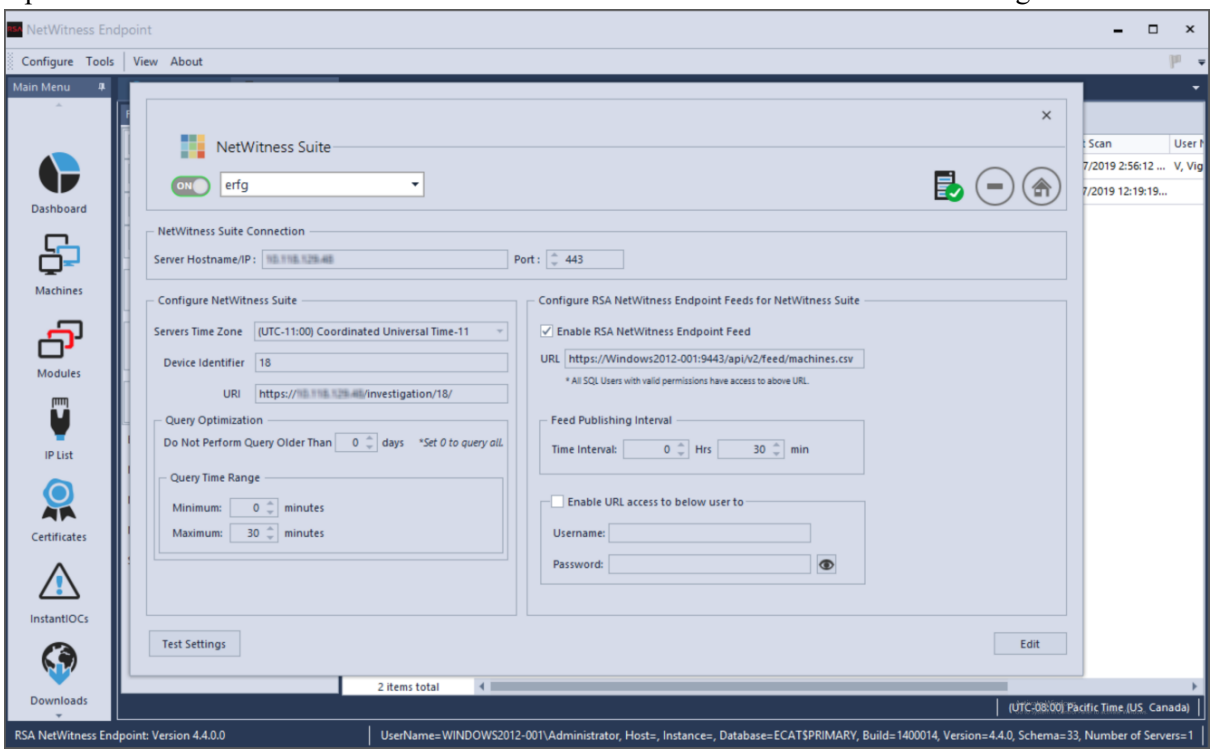
2. Update the server Hostname and IP Address to the current active server and test the settings.



NetWitness Suite

1. Log in to the NetWitness Endpoint user interface and go to **Configure** > **Monitoring and External Components Configuration** > **Netwitness Suite**.

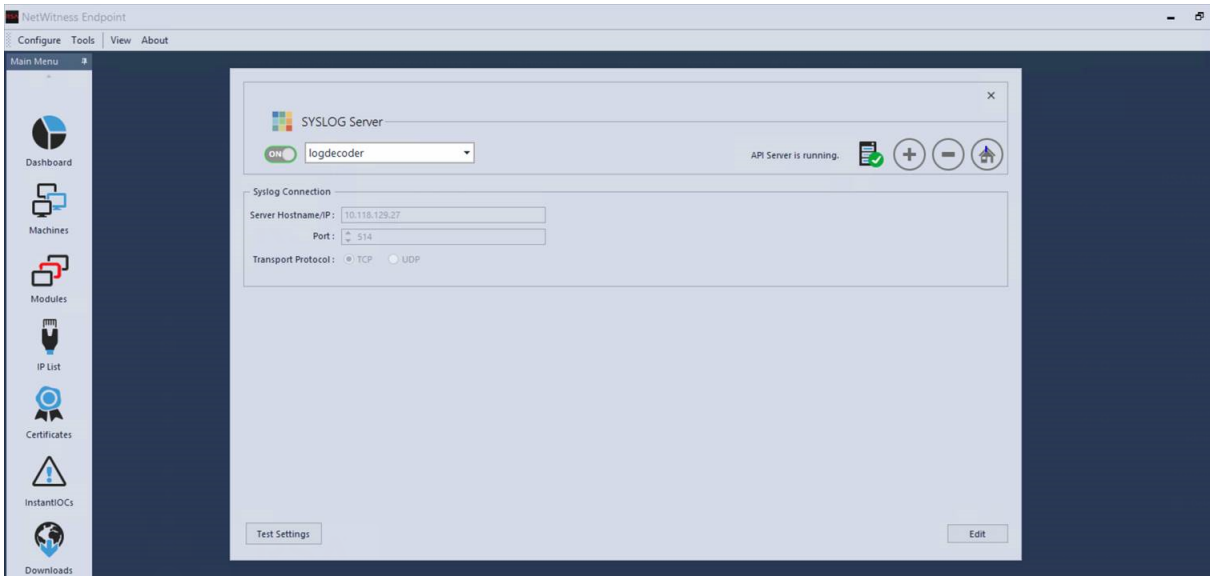
2. Update the server Hostname and IP address to the current active server and test settings.



Syslog Server Settings

If you are forwarding syslog messages to a NetWitness Platform Log Decoder, update the syslog server settings to point to the new IP address of the Log Decoder host.

1. Log in to the NetWitness Endpoint user interface and go to **Configure** > **Syslog Server**.

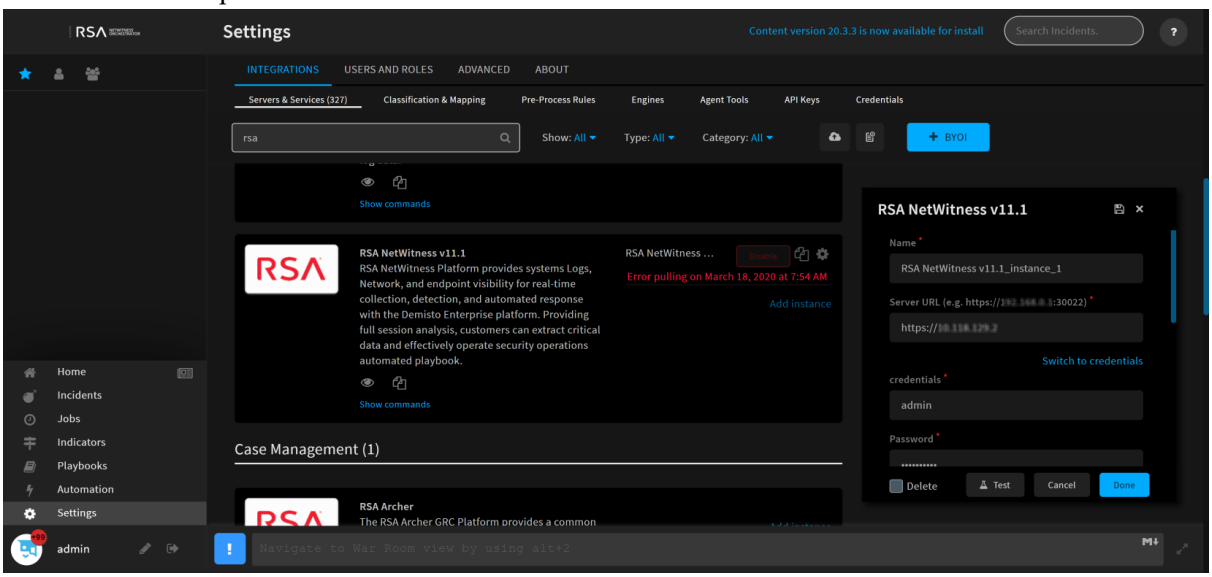


2. Select **logdecoder**, and in **Server Hostname/IP**, enter the new IP address of the Log Decoder host.

RSA NetWitness Orchestrator (By Demisto)

Update the Current Active NW Server to Fetch Respond Incidents and Alerts

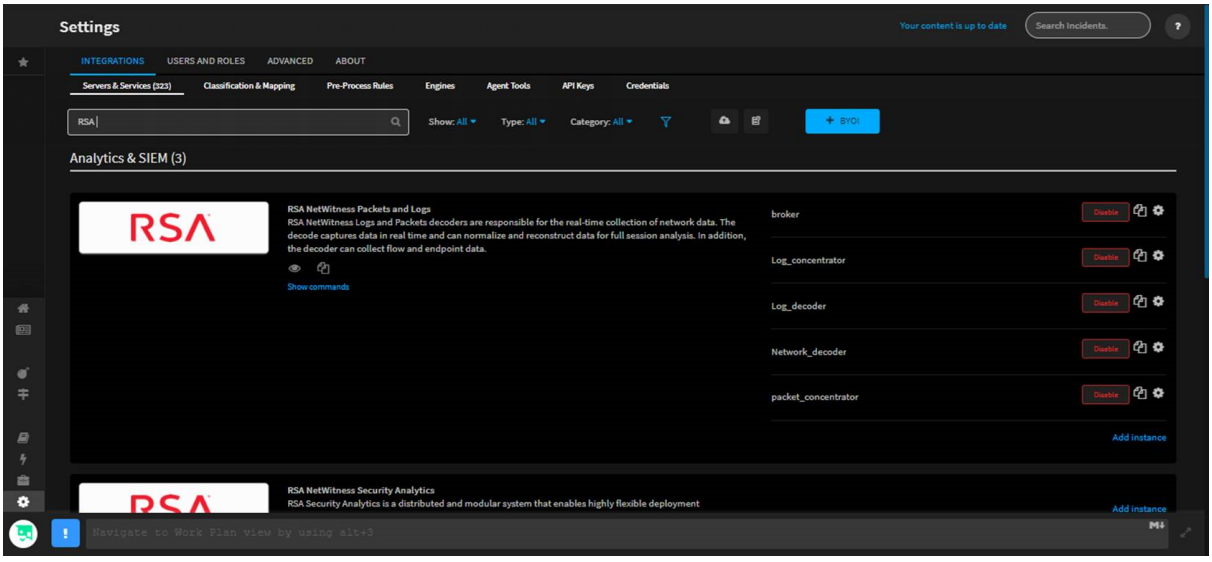
1. Log in to Orchestrator and go to **Settings** > **server&services**.
2. Edit the RSA NetWitness V11.1 instance by updating the server URL to the current active NW Server to fetch respond incidents and alerts.



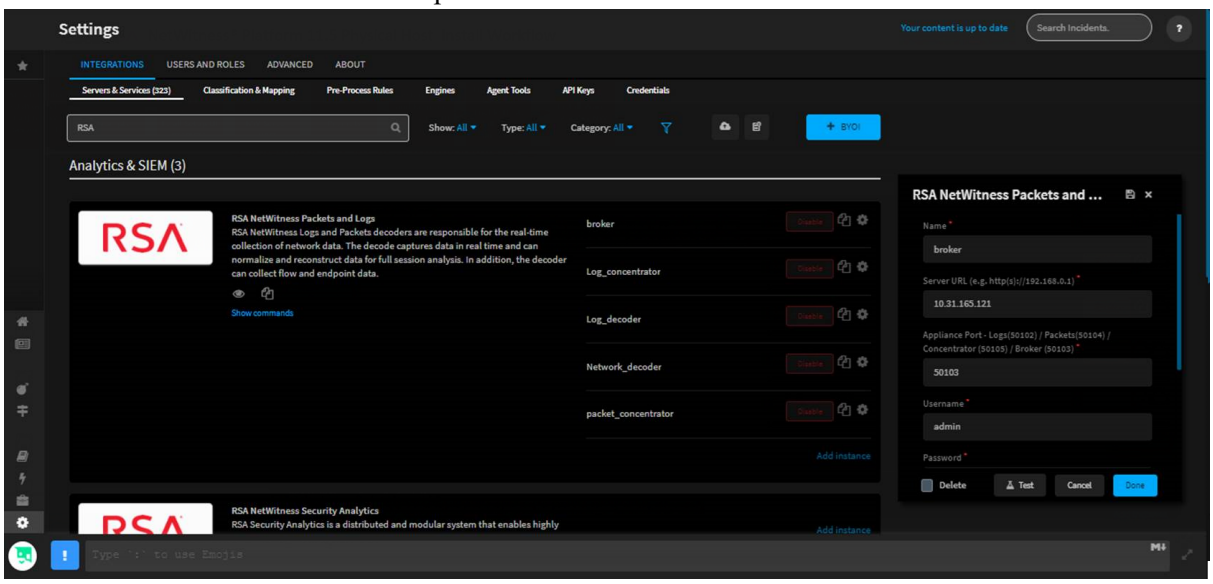
Update Component Hosts Acting as Data Sources

If you change the IP address of a component host, for example, a Concentrator, Network or Log Decoder, or Broker, that is acting as data source to the Orchestrator, update the following settings to point to the new IP address of the host.

1. Log in to Orchestrator and go to **Settings** > **server&services** and select the component host.



2. Enter the new IP address of the component host in **Server URL** and click **Done**.



Audit Logging

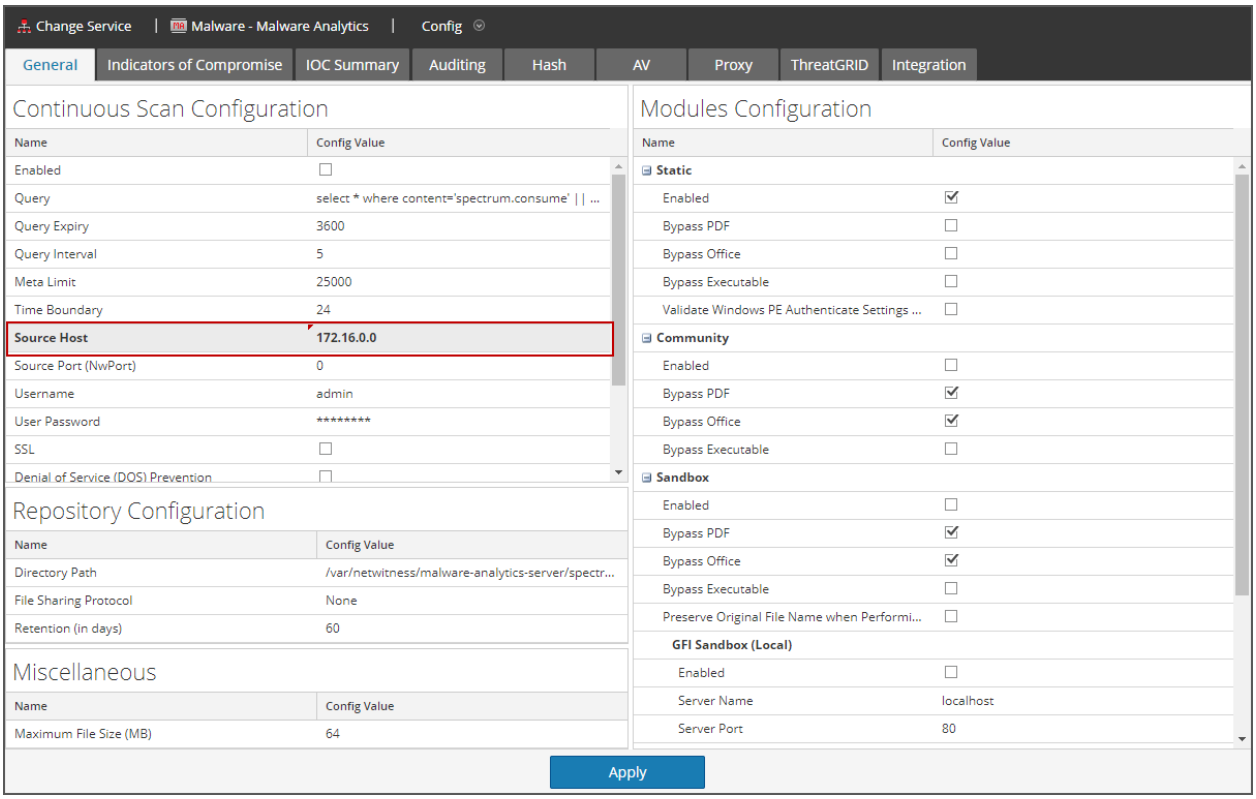
If you have changed the IP address of the NW Server, you must reconfigure audit logging. For instructions, see "Configure Global Audit Logging" in the *System Configuration Guide*.

Health and Wellness

If you have any Health and Wellness rules that contain IP addresses that have been changed, you must update those rules with the new IP addresses. For information about managing Health and Wellness rules, see "Monitor Health and Wellness using NetWitness Platform UI" in the *System Maintenance Guide*.

Malware Analysis

Source host IP address changes are not updated in the NetWitness user interface for Malware Analysis continuous scan configurations. You must manually update this configuration in the Malware Analysis Config view > **General** > **Continuous Scan Configuration** and update the source host IP address to the new host IP address.



Windows Legacy Collection

On occasion, you may need to change the IP address of your Windows Legacy Collector. You may also need to edit any Destination Groups that you have configured.

Change WLC IP Address

The following procedure describes how to change the IP address for your system.

1. Log onto the Windows Legacy Collector system and manually change the IP address on the system.
2. In the UI, confirm that the Log Collector service corresponding to the WLC system shows up in error (Red). It might take some time for it to reflect the changed status.
3. On the NetWitness Server, use the **nw-manage** utility to view the host information for the WLC using the following command:

   ```
   nw-manage --list-hosts
   ```

   Sample output from running the command is shown here:

   ```
   {
   "id" : "fdb8150c-e040-459e-8cc5-3c60ec2c65ae",
   "displayName" : "WLC-HOST-104",
   "hostname" : "10.101.216.102",
   "ipv4" : "10.101.216.102",
   "ipv4Public" : null
   } ]
   ```

   You use the value of **"id"** from your output in the following step.

4. Use the **nw-manage** utility to change the IP address of the WLC. For the **host-id** argument, use the value for the **"id"** that you noted from step 3. For the **ipv4** value, use the new IP Address to which you are changing.

   ```
   nw-manage --update-host --host-id "fdb8150c-e040-459e-8cc5-3c60ec2c65ae" --ipv4 10.101.216.105
   ```

5. After you see the message that the previous command ran successfully, go to the NetWitness Server UI and verify that the WLC service is running without any errors.

Edit Destination Groups For Log Collectors and VLCs

The Windows Legacy Collector is often configured with Destination Groups to forward events to Log Collectors or Virtual Log Collectors. If the IP address of any such Destination LC or VLC is changed, the Windows Legacy Collector can no longer forward events. To remediate this, you must edit the Destination groups for the WLC, making sure to select the new LC or VLC IP Address.

Note: If you added any content to the `/etc/hosts` file on the primary server, the contents of that file are available under `/var/netwitness/standby-data/unmanaged/etc` on the failover server. You can manually copy those files to the `/etc/hosts` file on the failover server after the failover is complete.

Fail Back Secondary NW Server to Primary NW Server

After a fail-over from the primary NW Server to the secondary NW Server, you need to fail back to your original setup of the primary NW Server in the active role and the secondary NW Server in the standby role.

Essentially, you follow the same steps described under Fail Over Primary NW Server to Secondary NW Server to fail back to your original setup (that is primary NW Server-active and secondary NW Server-standby). The difference is that you now need to fail over from the secondary NW Server to the primary NW Server.

Note: The secondary SA becomes active in case of a fail over. You should re-run the `schedule-standby-admin-data-sync` from the secondary SA before failing back to the primary SA.

Network Architecture and Ports

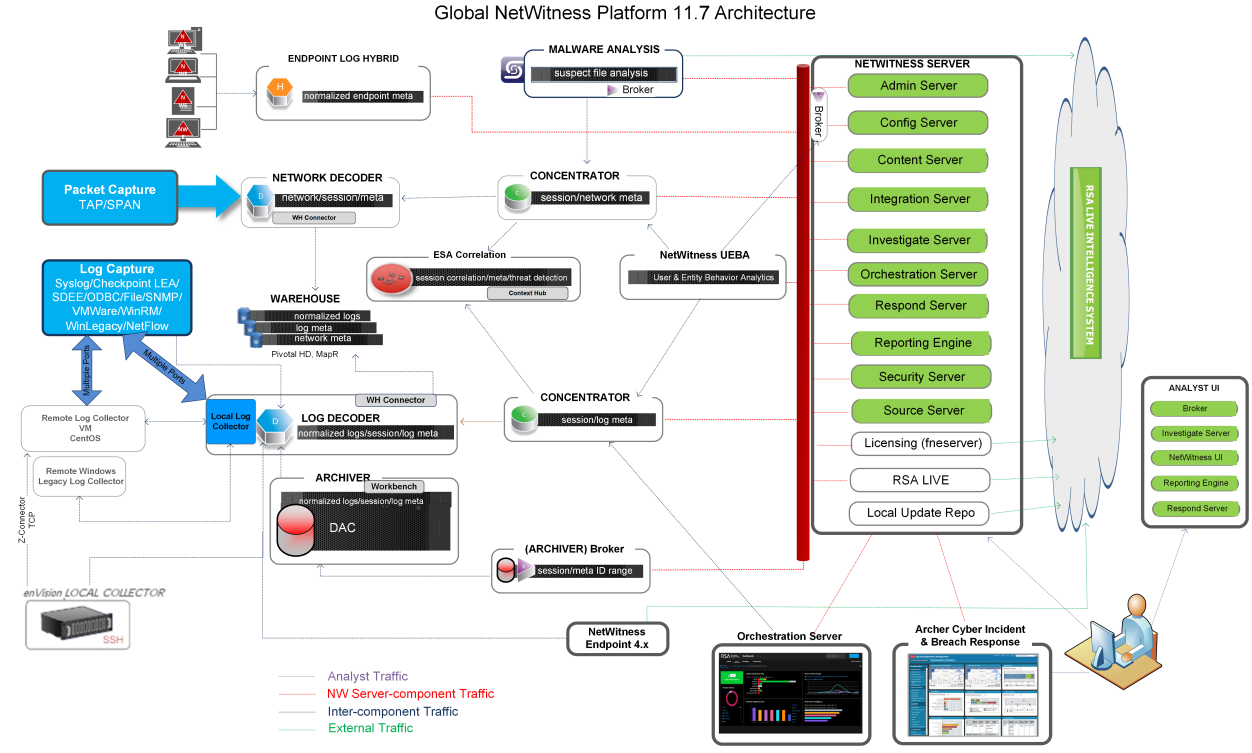
Refer to the following diagrams and port tables to ensure that all the relevant ports are opened for components in your NetWitness deployment to communicate with each other.

See NetWitness Endpoint Architecture at the end of this topic for individual Endpoint Architectural diagrams.

NetWitness Network Architecture Diagram

The following diagram illustrates the NetWitness network architecture including all of its component products.

Note: NetWitness core hosts must be able to communicate with the NetWitness Server (Primary Server in a multiple server deployment) through UDP port 123 for Network Time Protocol (NTP) time synchronization.

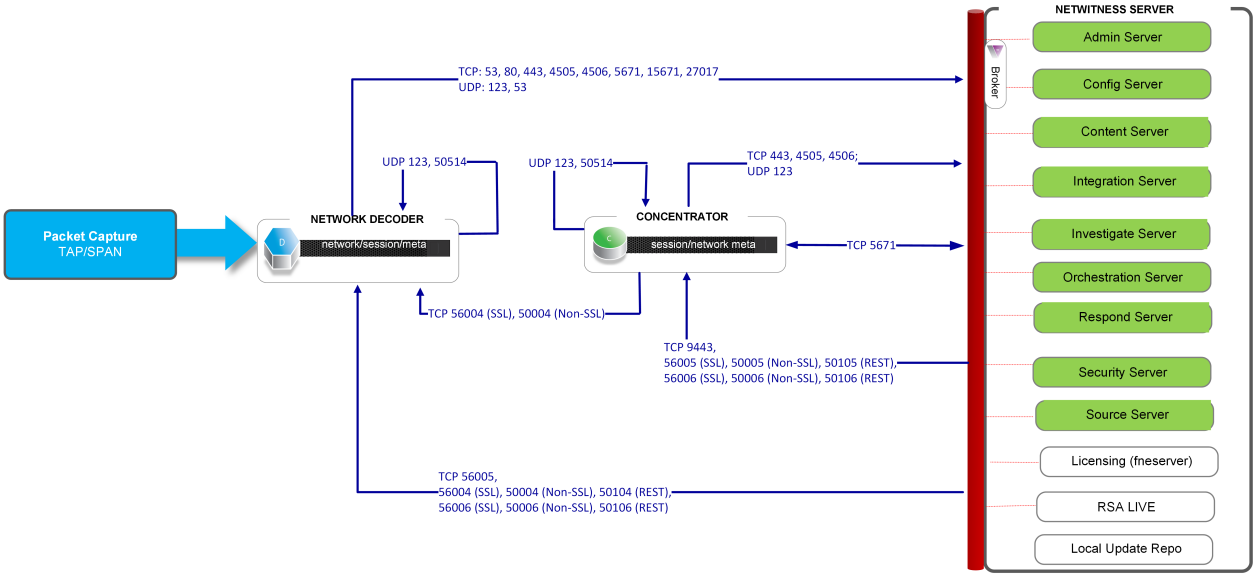


RSA NETWITNESS PLATFORM

Note:
Admin, Config, Content, Integration, Investigate, Orchestration, Respond, Security, and Source services come online automatically when you deploy the NW Server.
The core service hosts use the Local Update Repository on the NetWitness Server to get the rpm packages (that is, version updates).
NW Endpoint needs to access https://cms.netwitness.com to download Live Feeds.
RSA recommends that you use the Broker at the top of your deployment hierarchy for UEBA data source.
See *RSA NetWitness Platform Cloud Behavioral Analytics Gateway Configuration Guide* for information on the Cloud Gateway service.

NetWitness Network (Packets) Architecture Diagram with Ports

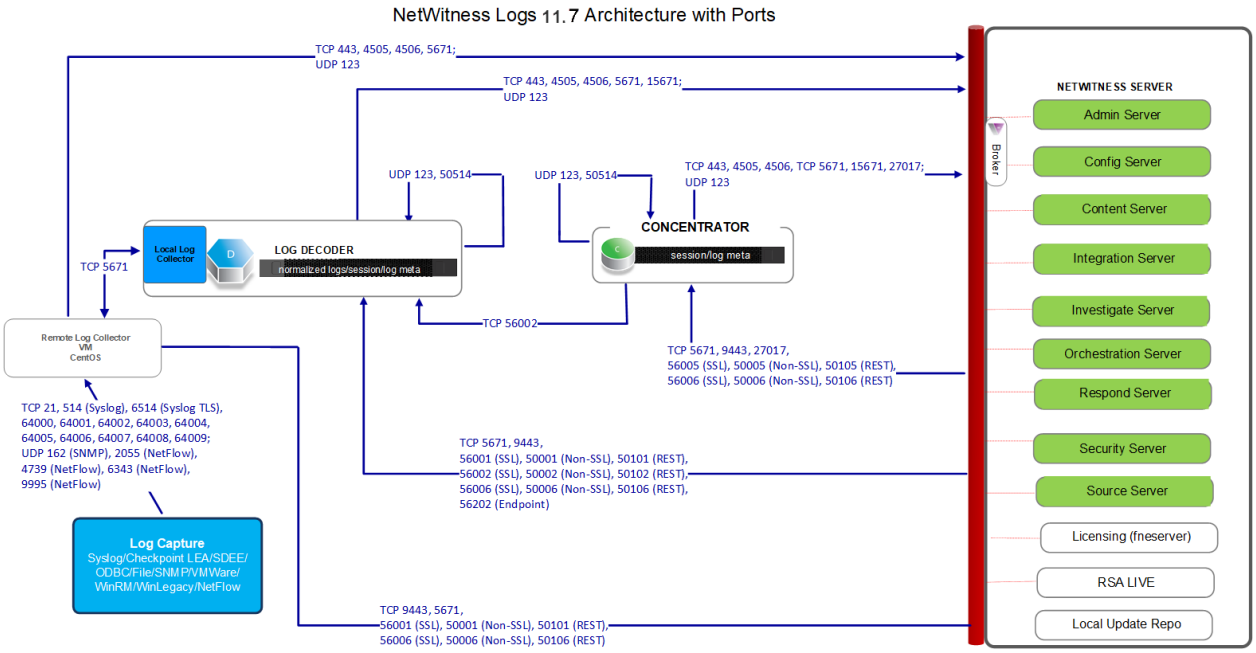
NetWitness Network 11.7 Architecture with Ports



RSA NETWITNESS NETWORK

Notes:
Admin, Config, Content, Integration, Investigate, Orchestration, Respond, Security, and Source services come online automatically when you deploy the NW Server.
The core service hosts use the Local Update Repository on the NetWitness Server to get the rpm packages (that is, version updates).

NetWitness Logs Architecture Diagram with Ports



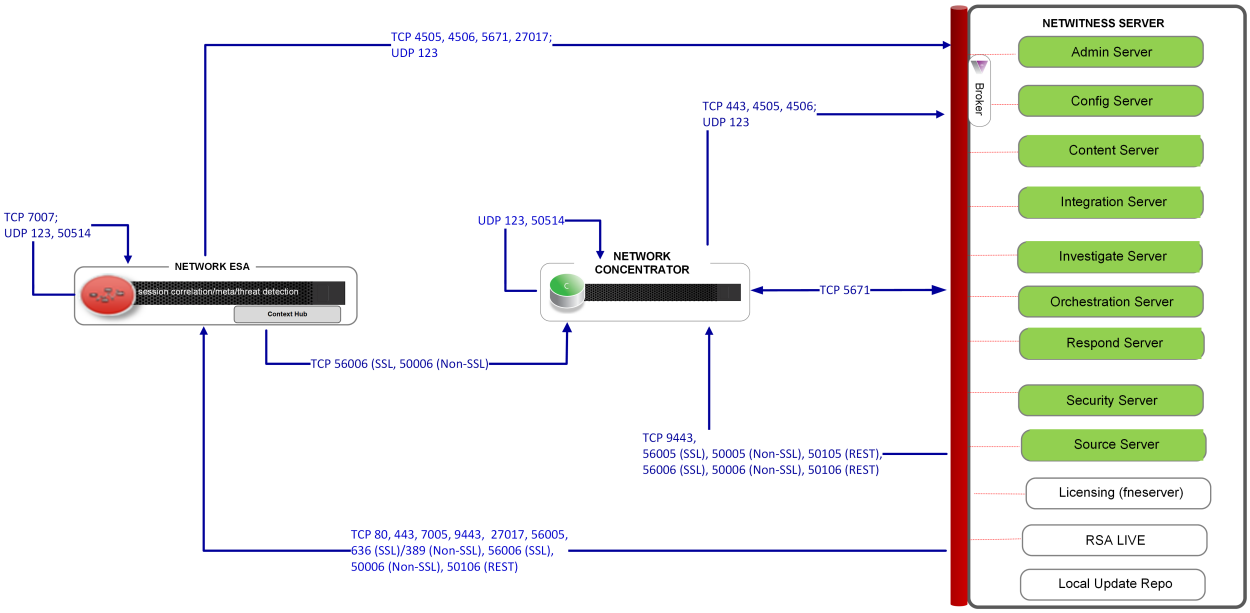
RSA NETWITNESS LOGS

Note:
Admin, Config, Content, Integration, Investigate, Orchestration, Respond, Security, and Source services come online automatically when you deploy the NW Server.
The core service hosts use the Local Update Repository on the NetWitness Server to get the rpm packages (that is, version updates).

Event Stream Analysis Network (Packets) Architecture Diagram with Ports

The following diagram illustrates the Event Stream Analysis network architecture with packet capture.

Event Source Analysis (ESA) Network 11.7 Architecture with Ports



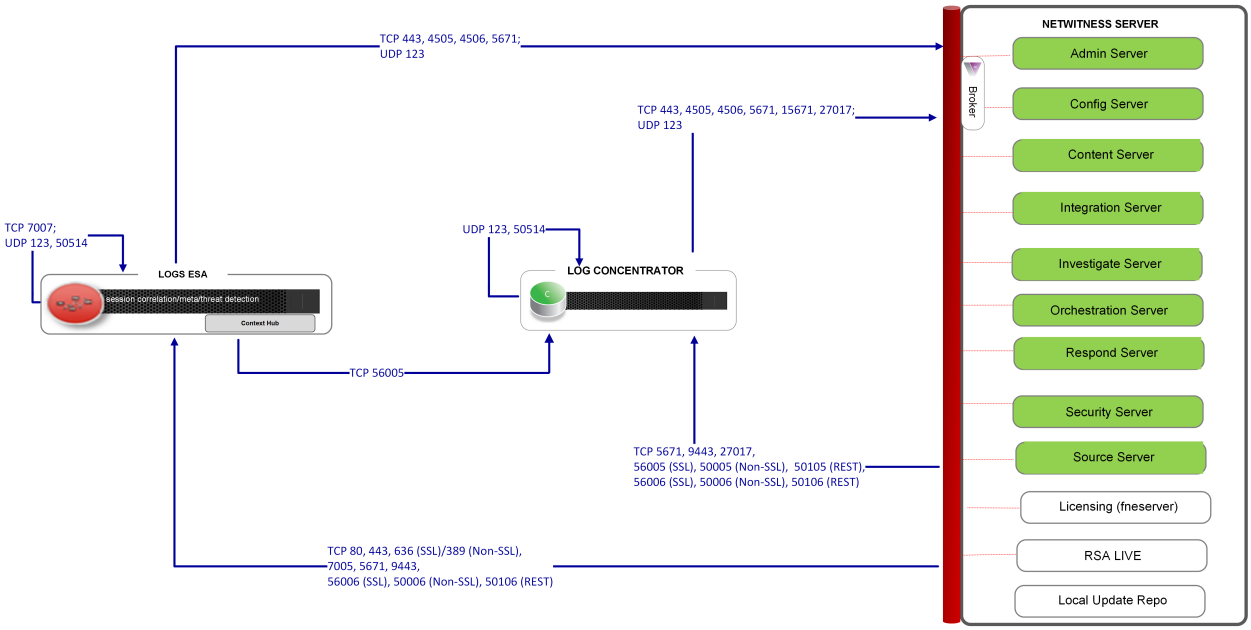
RSA NETWITNESS NETWORK

Notes:
Admin, Config, Content, Integration, Investigate, Orchestration, Respond, Security, and Source services come online automatically when you deploy the NW Server. The core service hosts use the Local Update Repository on the NetWitness Server to get the rpm packages (that is, version updates). Port 7005 is used by Context Hub as its HTTP port for REST calls, and is only installed on ESA Primary hosts.

Event Stream Analysis (Logs) Architecture Diagram with Ports

The following diagram illustrates the Event Stream Analysis network architecture with log collection.

Event Stream Analysis (ESA) Logs 11.7 Architecture with Ports



RSA NETWITNESS LOGS

Note:
Admin, Config, Content, Integration, Investigate, Orchestration, Respond, Security, and Source services come online automatically when you deploy the NW Server. The core service hosts use the Local Update Repository on the NetWitness Server to get the rpm packages (that is, version updates). Port 7005 is used by Context Hub as its HTTP port for REST calls, and is only installed on ESA Primary hosts.

NetWitness Firewall Requirements Summary

The following table lists all the ports that need to be open in your firewall by host.

Note: The "NW Server" host ports apply to both the Primary and Warm Standby NW Server. Synchronization between the Primary and Warm Stanby is done through TCP Port 22.

Source Host	Destination Host	Ports
NW Server	ESA Primary	**TCP:** 22, 5671, 7005 **UDP:**123
NW Server	ESA	**TCP:** 22, 5671 **UDP**: 123
NW Server	Network Decoder	**TCP:** 22, 5671, 50004 (Non-SSL), 50006 (Non-SSL), 50104 (REST), 50106 (REST), 56004 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Broker	**TCP:** 5671, 50003 (Non-SSL), 50006 (Non-SSL), 50103 (REST), 50106 (REST) 56003 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Concentrator (Network & Logs)	**TCP:** 22, 5671, 50005 (Non-SSL), 50006 (Non-SSL), 50105 (REST), 50106 (REST), 56005 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Network Hybrid	**TCP:** 22, 5671, 50004 (Non-SSL), 50005 (Non-SSL), 50006 (Non-SSL), 50104 (REST), 50105 (REST), 50106 (REST), 56004 (SSL), 56005 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Log Decoder	**TCP:** 22, 5671, 50001 (Non-SSL), 50002 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Log Hybrid	**TCP:** 22, 5671, 50001 (Non-SSL), 50002 (Non-SSL), 50005 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50105 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56005 (SSL), 56006 (SSL) **UDP:** 123

Source Host	Destination Host	Ports
NW Server	Log Hybrid - Retention	**TCP:** 22, 5671, 50001 (Non-SSL), 50002 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50105 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Endpoint Log Hybrid	**TCP:** 22, 5671, 7050, 7054, 50001 (Non-SSL), 50002 (Non-SSL), 50005 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50105 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56005 (SSL), 56006 (SSL), 56202 (Endpoint) **UDP:** 123
NW Server	VLC	**TCP:** 22, 5671, 50001 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50106 (REST), 56001 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Archiver	**TCP:** 22, 514, 5671, 6514, 50006(Non-SSL), 50007 (Non-SSL), 50008 (Non-SSL), 50106 (REST), 50107 (REST), 50108 (REST), 56006 (SSL), 56007 (SSL), 56008 (SSL) **UDP:** 123, 514
NW Server	Malware	**TCP:** 22, 5671, 5432, 50003 (Non-SSL), 50006 (Non-SSL), 50103 (REST), 50106 (REST), 56003 (SSL), 56006 (SSL), 60007 **UDP:** 123
NW Server	UEBA	**TCP:** 22, 15671, 5671, 443 **UDP:** 123
ESA	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017, 444 **UDP:** 123, 53
ESA	Active Directory	**TCP:** 389 (Non-SSL), 636 (SSL)
ESA	Archer	**TCP:** 80 (Non-SSL), 443 (SSL),
ESA Secondary	ESA Primary	**TCP:** 27017
ESA Primary or Secondary	Concentrator	**TCP:** 50005 (Non-SSL), 56005 (SSL)
Network Decoder	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017, 444 **UDP:** 53, 123

Source Host	Destination Host	Ports
Concentrator (Network & Logs)	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017, 444 **UDP:** 53, 123
Network Hybrid	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017, 444 **UDP:** 53, 123
Log Decoder	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017, 444 **UDP:** 53, 123
Log Hybrid	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671,15671, 27017, 444 **UDP:** 53, 123
Log Hybrid - Retention	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671,15671, 27017, 444 **UDP:** 53, 123
VLC	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671,15671, 27017, 444 **UDP:** 53, 123
VLC	Log Collector	**TCP:** 5671
Log Collector	VLC	**TCP:** 5671
Endpoint Log Hybrid	NW Server	**TCP:** 53, 80, 443, 5671, 4505, 4506, 15671, 27017, 444 **UDP:** 53, 123
Endpoint Log Hybrid	Log Decoder	**TCP:** 50202 (Non-SSL), 50102 (REST), 56202 (SSL) **UDP:** 514
Endpoint Agent	Log Decoder	**TCP:** 514, 6514 **UDP:** 514
Endpoint Agent	Endpoint Log Hybrid	**TCP:** 443 **UDP:** 444
UEBA	NW Server	**TCP:** 53, 80, 443, 444, 4505, 4506, 5671, 15671, 27017, 50003 (Broker-Non-SSL), 50103 (Broker/REST), 56003 (Broker/SSL) **UDP:** 53, 123
UEBA	Concentrator	**TCP:** 50005 (Non-SSL), 50105 (REST), 56005 (SSL)

Source Host	Destination Host	Ports
www connections		
NW Server	cms.netwitness.com download.rsasecurity.com rsasecurity.subscribenet.com update.netwitness.com	**TCP:** 443
ESA (Primary & Secondary)	cms.netwitness.com	**TCP:** 443
Malware	panacea.threatgrid.com cloud.netwitness.com	**TCP:** 443

Comprehensive List of NetWitness Host, Service, and iDRAC Ports

Note: For ports used in event collection through the NetWitness Logs, see the *Log Collection Configuration Guide for RSA NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

This section contains the port specifications for the following hosts:

NW Server Host	iDRAC Ports
Analyst UI Host	Log Collector Host
Archiver Host	Log Decoder Host
Broker Host	Log Hybrid Host
Concentrator Host	Log Hybrid - Retention
Endpoint Log Hybrid Host	Malware Host
Endpoint Relay Server	Network Decoder Host
Event Stream Analysis Host	Network Hybrid Host
New Health & Wellness	UEBA Host

NW Server Host (Primary and Warm Standby NW Server Host)

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	NW Server	TCP 443, 80	nginx - NetWitness UI
Admin Workstation	NW Server	TCP 15671	RabbitMQ Management UI
Admin Workstation	NW Server	TCP 22	SSH Primary to Standby NW Server synchronization port.
NW Hosts	NW Server	TCP 444	Node-infra-server check
NW Hosts	NW Server	TCP 53 UDP 53	DNS
NW Hosts	NW Server	TCP 15671	RabbitMQ Management UI
NW Hosts	NW Server	TCP 4505, 4506	Salt Master Ports
NW Hosts	NW Server	TCP 443	RSA Update Repository
NW Hosts	NW Server	TCP 5671	RabbitMQ-amqp
NW Hosts	NW Server	UDP 123	NTP
NW Hosts	NW Server	TCP 27017	MongoDB
NW Server	cms.netwitness.com	TCP 443	Live Content Management System (CMS) Library
NW Server	download.rsasecurity.com	TCP 443	RSA Licensing
NW Server	rsasecurity.subscribenet.com	TCP 443	RSA Licensing
NW Server	update.netwitness.com	TCP 443	Netwitness Software Updates
NW Server	NFS Server	TCP 111, 2049, UDP 111, 2049	iDRAC Installations
NW Server	NW Hosts	UDP 123	NTP
NW Server	NW Endpoint	TCP 443, 9443	For NW Endpoint 4.x integrations
NW Hosts	NW Server	TCP 444	nw-node-infra service failover check

Analyst UI Host

Source Host	Destination Host	Destination Ports	Comments
Analyst UI	NW Server	TCP 7006	The Content Server is listening on this port.
Analyst UI	NW Server	TCP 7009	The Admin Server is listening on this port.
Analyst UI	NW Server	TCP 7012	The Integration Server is listening on this port.
Analyst UI	NW Server	TCP 7015	The Source Server is listening on this port.
Analyst UI	NW Server	TCP 7016	The License Server is listening on this port.
NW Hosts	Analyst UI	TCP 5671	RabbitMQ-amqp
Analyst UI	NW Server	UDP 123	NTP
Analyst UI	NW Server	TCP 444	nw-node-infra service failover check

Archiver Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Archiver	TCP 15671	RabbitMQ Management UI
Archiver	NW Server	TCP 15671	RabbitMQ Management UI
Archiver	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Archiver	TCP 22	SSH
NW Server	Archiver	TCP 50008 (Non-SSL), 56008 (SSL), 50108 (REST)	Archiver Application Ports
NW Server	Archiver	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Archiver	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
NW Server	Archiver	TCP 50007 (Non-SSL), 56007 (SSL), 50107 (REST)	Workbench Application Ports
Archiver	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Archiver	NW Server	TCP 444	nw-node-infra service failover check

Broker Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Broker	TCP 15671	RabbitMQ Management UI
Broker	Concentrator	TCP 50005 (Non-SSL), 56005	Concentrator Application Port
Broker	NW Server	TCP 15671	RabbitMQ Management UI
Broker	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Broker	TCP 22	SSH
NW Server	Broker	TCP 50003 (Non-SSL), 56003 (SSL), 50103 (REST)	Broker Application Ports
NW Server	Broker	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Broker	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Broker	NW Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Endpoint Broker	NW Server	TCP 443	RSA Update Repository
Broker	NW Server	TCP 444	nw-node-infra service failover check

Concentrator Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Concentrator	TCP 15671	RabbitMQ Management UI
Concentrator	Log Decoder	TCP 50002 (Non-SSL), 56002 (SSL)	Log Decoder Application Port
Concentrator	Network Decoder	TCP 56004, 50004 (Non-SSL)	Network Application Port
Concentrator	NW Server	TCP 15671	RabbitMQ Management UI
Concentrator	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Concentrator	TCP 22	SSH
NW Server	Concentrator	TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports
Malware	Concentrator	TCP TCP 50005 (Non-SSL), 56005 (SSL)	Malware
NW Server	Concentrator	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Concentrator	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Concentrator	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Concentrator	NW Server	TCP 444	nw-node-infra service failover check

Endpoint Log Hybrid

Source Host	Destination Host	Destination Ports	Comments
Endpoint Agent	Endpoint Log Hybrid	TCP 443 UDP 444	NGINX HTTPS NGINX UDP. If UDP port 444 is not acceptable in your environment, see How to Change UDP Port for Endpoint Log Hybrid.
Endpoint Agent	Log Decoder or Virtual Log Collector	TCP 514 (Syslog) UDP 514 (Syslog) TLS 6514	Windows Log Collection
Endpoint Log Hybrid	Log Decoder (External)	TCP 50102 (REST) 56202 (Protobuf SSL) 50202 (Protobuf)	To forward meta to an external Log Decoder
Endpoint Log Hybrid	NW Server	TCP 443	RSA Update Repository
NW Server	Endpoint Log Hybrid	TCP 7050	UI web traffic
Endpoint Log Hybrid	NW Server	TCP 5671	Message Bus
Endpoint Log Hybrid	NW Server	TCP 27017	MongoDB
NW Server	Endpoint Log Hybrid	TCP 7054	UI web traffic
NW Server	NFS Server	TCP 111, 2049 UDP 111, 2049	iDRAC Installations
NW Server	Endpoint Log Hybrid	TCP 50001 (Non-SSL), 56001 (SSL), 50101(REST)	Log Collector application ports
NW Server	Endpoint Log Hybrid	TCP 50002 (Non-SSL), 56002 (SSL), 56202 (Endpoint), 50102 (REST)	Log Decoder application ports
Admin Workstation	Endpoint Log Hybrid	TCP 15671	RabbitMQ Management UI

Source Host	Destination Host	Destination Ports	Comments
Endpoint Log Hybrid	NW Server	TCP 15671	RabbitMQ Management UI
Endpoint Log Hybrid	NW Server	TCP 444	nw-node-infra service failover check

Endpoint Relay Server

Source Host	Destination Host	Destination Ports	Comments
Endpoint Agent	Relay Server	TCP 443	To forward host data to the Relay Server
Endpoint Log Hybrid	Relay Server	TCP 443	Pull host data from the Relay Server

Event Stream Analysis (ESA) Host

Note: The ports in this table are for the ESA Primary an ESA Secondary hosts. The Content Hub, Correlation and ESA Analytics services are co-located on the ESA Primary host. The Correlation and ESA Analytics services are co-located on the ESA Secondary host.

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	ESA	TCP 15671	RabbitMQ Management UI
ESA Primary and Secondary	NW Server	TCP 15671	RabbitMQ Management UI
ESA Primary and Secondary	NW Server	TCP 443	RSA Update Repository
Admin Workstation	ESA	TCP 22	SSH
NW Server, ESA Secondary	ESA Primary	TCP 27017	MongoDB
NW Server	ESA Primary	TCP 7005	Context Hub Launch Port - (ESA Primary)
NW Server	ESA	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
ESA Primary and Secondary	cms.netwitness.com	TCP 443	Live
ESA Primary and Secondary	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
ESA Primary and Secondary	Active Directory	636 (SSL)/389 (Non-SSL)	
NW Server	ESA	80 (HTTP)/ 443 (HTTPS)(REST)	
ESA Primary	Archer	443 (SSL)/80 (Non-SSL)	
ESA Primary	ESA Primary	TCP 7007	Launch Port
ESA Primary and Secondary	NW Server	TCP 444	nw-node-infra service failover check

New Health and Wellness

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	New Health and Wellness	TCP 22	SSH
Admin Workstation	New Health and Wellness	TCP 5601	Kibana UI
NW Hosts	New Health and Wellness	TCP 9200	Elasticsearch REST API Port
NW Server	New Health and Wellness	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
NW Server	New Health and Wellness	TCP 15671	RabbitMQ Management UI
NW Server	New Health and Wellness	TCP 7018	Metrics Server Launch Port
NW Server	New Health and Wellness	TCP 7020	Node Infra Server Launch Port

New Health and Wellness on Different Subnet

If the New Health and Wellness is on a different subnet, you must open the respective NetWitness Platform hosts port.

Example:
New Health and Wellness is on subnet A: 10.10.1.0/24 and LogHybrid host is on subnet B: 10.10.2.0/24. In this case, you must open ports for Log Decoder, Log Collector, Concentrator on Metrics Server (New Health and Wellness host) to allow ports in the firewall for communication.

Source Host	Destination Host	Destination Ports	Comments
New Health and Wellness	Log Decoder	50002(Non-SSL),56002 (SSL)	Log Decoder Application Ports
New Health and Wellness	Log Collector	50001(Non-SSL),56001 (SSL)	Log Collector Application Ports
New Health and Wellness	Concentrator	50005(Non-SSL)/56005 (SSL)	Concentrator Application Ports

iDRAC Ports

Port	Function	Comments
22*	SSH	Default, configurable port through which iDRAC listens for connections
443*	HTTP	Default, configurable port through which iDRAC listens for connections
5900*	Virtual Console keyboard and mouse redirection, Virtual Media, Virtual Folders, and Remote File Share.	Default, configurable port through which iDRAC listens for connections
111, 2049	TCP	NetWitness Platform hosts to NFS Server
111, 2049	UDP	NetWitness Platform hosts to NFS Server

Log Collector Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Collector	TCP 15671	RabbitMQ Management UI
Log Collector	NW Server	TCP 15671	RabbitMQ Management UI
Log Collector	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Collector	TCP 22	SSH
Log Collector	Log Event Sources	See *Log Collection Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.	
Log Event Sources	Log Collector	TCP 514 (Syslog) UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)"	Log Collection Ports
Log Event Sources	Log Collector	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008,64009	Log Collection FTP/S Ports
NW Server	Log Collector	TCP 50001 (Non-SSL), 56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Collector	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Collector	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Collector	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC installations
Log Collector	Virtual Log Collector	TCP 5671	In Pull Mode
Virtual Log Collector	Log Collector	TCP 5671	In Push Mode
Log Collector	NW Server	TCP 444	nw-node-infra service failover check

Log Decoder Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Decoder	TCP 15671	RabbitMQ Management UI
Log Decoder	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Decoder	TCP 22	SSH
Log Decoder	Log Event Sources	See *Log Collection Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.	
Log Event Sources	Log Decoder	TCP 514 (Syslog), UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)	Log Collection Ports
Log Event Sources	Log Decoder	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008, 64009	Log Collection FTP/S Ports
NW Server	Log Decoder	TCP 50001 (Non-SSL),56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Decoder	TCP 50002 (Non-SSL), 56002 (SSL),56202 (Endpoint), 50102 (REST)	Log Decoder Application Ports
NW Server	Log Decoder	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Decoder	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Decoder	Log Collector	TCP 6514	
Log Decoder	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Log Decoder	NW Server	TCP 15671	RabbitMQ Management UI
Log Decoder	NW Server	TCP 444	nw-node-infra service failover check

Log Hybrid Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Hybrid	TCP 15671	RabbitMQ Management UI
Log Hybrid	NW Server	TCP 15671	RabbitMQ Management UI
Log Hybrid	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Hybrid	TCP 22	SSH
Log Collector	Log Event Sources	See *Log Collection Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.	
Log Event Sources	Log Hybrid	TCP 514 (Syslog), UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)	Log Collection Ports
Log Event Sources	Log Hybrid	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008, 64009	Log Collection FTP/S Ports
NW Server	Log Hybrid	TCP 50001 (Non-SSL), 56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Hybrid	TCP 50002 (Non-SSL), 56002 (SSL), 56202 (Endpoint), 50102 (REST)	Log Decoder Application Ports
NW Server	Log Hybrid	TCP TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports

Source Host	Destination Host	Destination Ports	Comments
NW Server	Log Hybrid	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Hybrid	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Hybrid	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Log Hybrid	NW Server	TCP 444	nw-node-infra service failover check

Log Hybrid - Retention Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Hybrid - Retention	TCP 15671	RabbitMQ Management UI
Log Hybrid - Retention	NW Server	TCP 15671	RabbitMQ Management UI
Log Hybrid - Retention	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Hybrid - Retention	TCP 22	SSH
Log Collector	Log Event Sources	See *Log Collection Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.	
Log Event Sources	Log Hybrid - Retention	TCP 514 (Syslog), UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)	Log Collection Ports
Log Event Sources	Log Hybrid - Retention	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008, 64009	Log Collection FTP/S Ports
NW Server	Log Hybrid - Retention	TCP 50001 (Non-SSL), 56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Hybrid - Retention	TCP 50002 (Non-SSL), 56002 (SSL), 56202 (Endpoint), 50102 (REST)	Log Decoder Application Ports
NW Server	Log Hybrid - Retention	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Hybrid - Retention	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Hybrid - Retention	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Log Hybrid - Retention	NW Server	TCP 444	nw-node-infra service failover check

Malware Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Malware	TCP 15671	RabbitMQ Management UI
Malware	NW Server	TCP 15671	RabbitMQ Management UI
Malware	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Malware	TCP 22	SSH
NW Server	Malware	TCP 60007	Malware Application Ports
NW Server	Malware	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Malware	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
NW Server	Malware	TCP 5432	Postgresql
NW Server	Malware	TCP 56003 (SSL), 50103 (REST)	Broker Application Ports
Malware	panacea.threatgrid.com	TCP 443	Threatgrid
Malware	cloud.netwitness.com	TCP 443	Community evaluation / Opswat
Malware	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Malware	NW Server	TCP 444	nw-node-infra service failover check

Network Decoder Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Network Decoder	TCP 15671	RabbitMQ Management UI
Network Decoder	NW Server	TCP 15671	RabbitMQ Management UI
Network Decoder	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Network Decoder	TCP 22	SSH
NW Server	Network Decoder	TCP 56004 (SSL), 50104 (REST), 50004 (Non-SSL)	Network Decoder Application Ports
NW Server	Network Decoder	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Network Decoder	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Network Decoder	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Network Decoder	NW Server	TCP 444	nw-node-infra service failover check

Network Hybrid Host

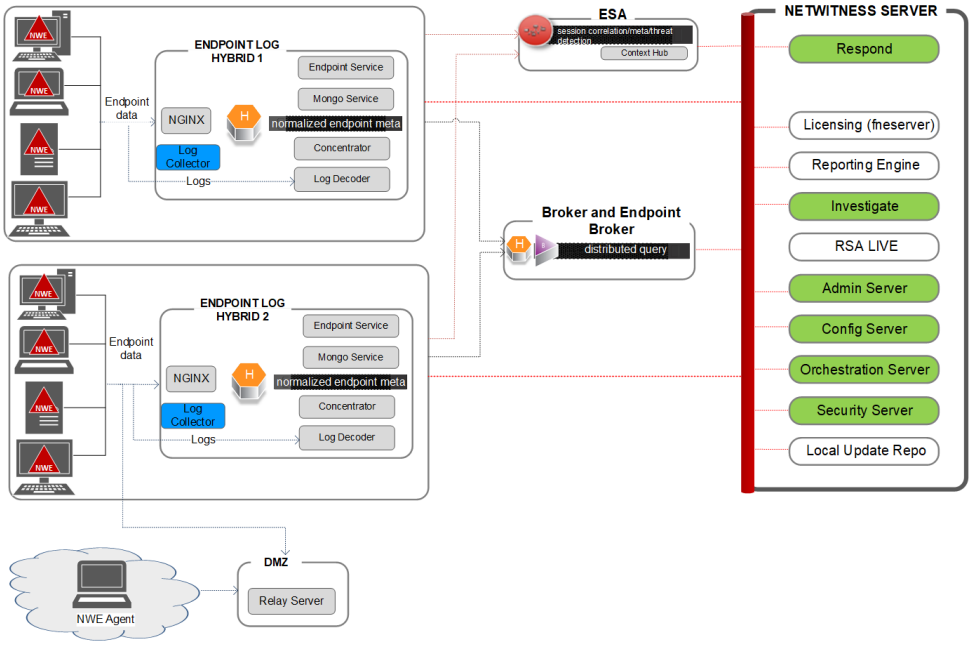
Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Network Hybrid	TCP 15671	RabbitMQ Management UI
Network Hybrid	NW Server	TCP 15671	RabbitMQ Management UI
Network Hybrid	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Network Hybrid	TCP 22	SSH
NW Server	Network Hybrid	TCP 56004 (SSL), 50104 (REST), 50004 (Non-SSL)	Network Decoder Application Ports
NW Server	Network Hybrid	TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports
NW Server	Network Hybrid	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Network Hybrid	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Network Hybrid	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Network Hybrid	NW Server	TCP 444	nw-node-infra service failover check

UEBA Host

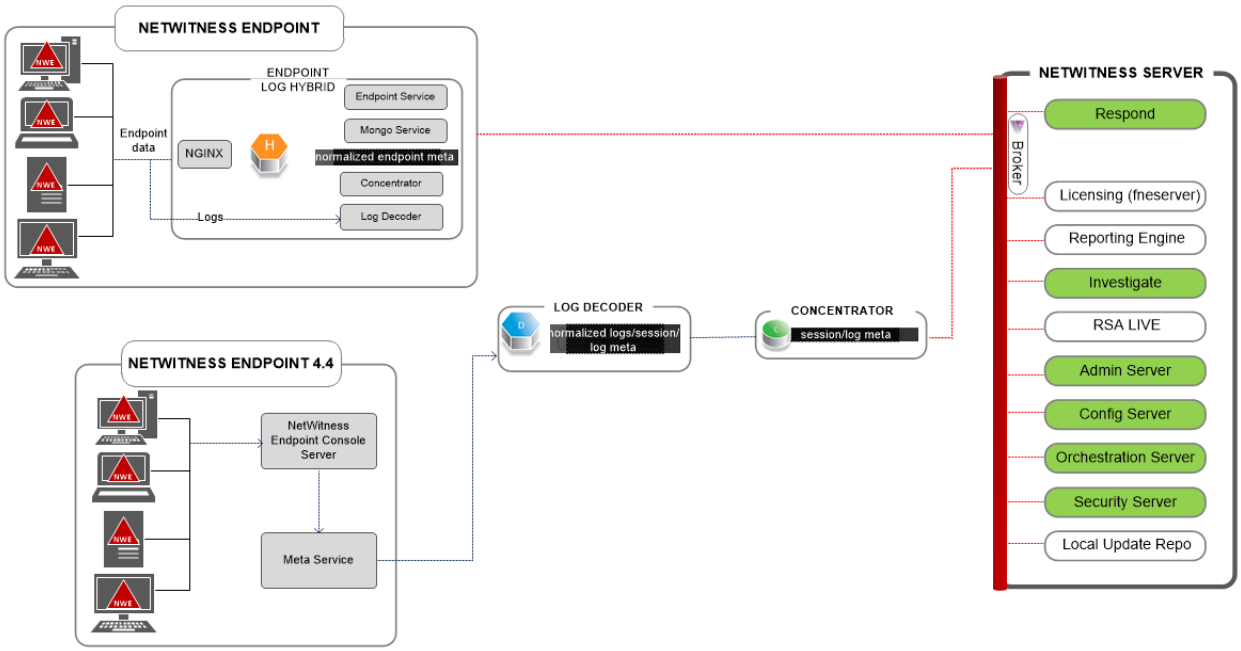
Source Host	Destination Host	Destination Ports	Comments
UEBA Server	NW Server	TCP 443	RSA Update Repository
UEBA Server	Broker	TCP 56003 (SSL), 50103 (REST)	Broker Application Ports
UEBA Server	Concentrator	TCP TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports
Admin Workstation	UEBA Server	443	UEBA Monitoring
Admin Workstation	UEBA Server	22	SSH
Admin Workstation	UEBA Server	TCP 15671	RabbitMQ Management UI
UEBA Server	NW Server	TCP 15671	UEBA Alerts forwarding to Respond
NW Server	NFS Server	TCP 111, 2049 UDP 111, 2049	iDRAC Installations
NW Server	UEBA Server	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts
UEBA Server	NW Server	TCP 444	nw-node-infra service failover check

NetWitness Endpoint Architecture

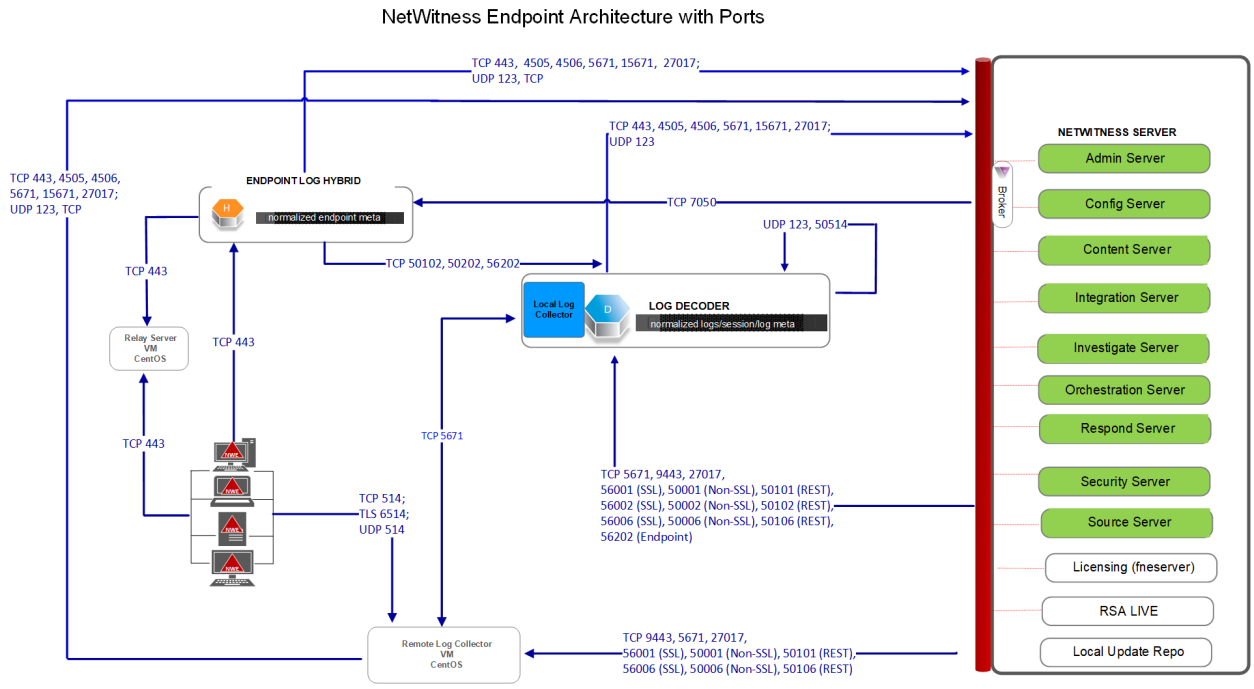
NetWitness Endpoint Architecture



NetWitness Endpoint 4.4 Integration with NetWitness Platform



NetWitness Endpoint Architecture with Ports



For more information on the services running on Endpoint Log Hybrid, see *RSA NetWitness Endpoint Configuration Guide*.

How to Change UDP Port for Endpoint Log Hybrid

The following steps tell you how to change the Endpoint Log Hybrid default UDP port 444 if it is not acceptable in your environment. 555 is the example this procedures uses as a replacement for 444 UDP port.
There are two tasks you need to o change the Endpoint Log Hybrid default UDP port 444:

Task 1 - Tell All Agents to Use a New UDP Port

Task 2 - Update the Port on All Endpoint Log Hybrid Hosts in Your Environment

Note: If you did not select the custom firewall rules option when you ran the nwsetup-tui, NetWitness platform overwrites the firewall rules are overwritten after a period of time. Please refer to the following Knowledge Base Article 00036446 (https://community.netwitness.com/t5/netwitness-knowledge-base/how-to-add-custom-firewall-rules-after-nwsetup-tui-has-completed/ta-p/5900) if this is the case.

Task 1 - Tell All Agents to Use a New UDP Port

Complete the following steps to update the UDP port in the default Enterprise Data Replication (EDR) policy, and all other policies you have, to tell all agents to use a new UDP port.

1. In the **NetWitness** menu, select  (Admin) > **Endpoint Sources** > **Policies**.
   The **Policies** view is displayed.
2. Select the **Default EDR Policy** and click **Edit** from the toolbar.
3. roll down to find the **UDP PORT** and change the value (for example, change from **444** to **555**).
4. Click **Publish Policy** at the bottom of the view.

Task 2 - Update the Port on All Endpoint Log Hybrid Hosts in Your Environment

SSH to each Endpoint Log Hybrid host in your environment with `admin` credentials and make the following updates.

1. Update the `iptables` rules to allow `555` in place of `444`.
   a. Replace `444` with `555` in the following file.
      ```
      vi /etc/sysconfig/iptables
      ```
   b. Restart `iptables` with the following command string.
      ```
      systemctl restart iptables
      ```
   c. Verify the change with the following command string.
      ```
      iptables -L -n
      ```
      The following is an example of what is displayed for a correct change.
      ```
      ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp multiport dports 555 /*
      EndpointNginxPort */ ctstate NEW
      ```
2. Update the SELinux policy. 555 is a privileged port, so you must update SELinux policy to allow this port.
   a. Run the following command string.
      ```
      semanage port -a -t http_port_t -p udp 555
      ```
      If you received any python errors or warnings, ignored them.
   b. Verify the change with the following command string.
      ```
      semanage port -l | grep http_port_t
      ```
      The following is an example of what is displayed for a correct change.
      ```
      http_port_t udp 555, 444
      ```
   c. (Optional) Remove `444`.
3. Update nginx config.
   a. Edit the following file.
      ```
      vi /etc/nginx/nginx.conf
      ```
   b. Search for the following string.
      ```
      listen 444 udp;
      ```
   c. Replace `444` with `555`.

d. Restart nginx with the following command string.
```
systemctl restart nginx
```

4. Verify that agents are communicating over the new port.

a. Run the following command string.
```
tcpdump -i eth0 port 555
```

b. Wait for 30 seconds because the port sends out a beacon every 30 seconds. If t everything is working correctly, information similar to the following will be displayed.
```
09:20:12.571316 IP 10.40.15.103.60807 > EPS1.rsa.lab.emc.com.dsf: UDP,
length 20
09:20:12.572433 IP EPS1.rsa.lab.emc.com.dsf > 10.40.15.103.60807: UDP,
length 1
```
Both lines must be returned. One is the size request (`20` bytes) and the other is the response size (`1` byte).

Site Requirements and Safety

Make sure that you read this topic thoroughly and observe all warnings and precautions prior to installing or maintaining your NetWitness devices.

Intended Application Uses

This product was evaluated as Information Technology Equipment (ITE) that may be installed in offices, schools, computer rooms, and similar indoor commercial type locations. This device is not intended for any connection to an outdoor type cable.

Service

There are no user-serviceable components inside of this device. Please contact Customer Care in the event of a malfunction. In a fault condition, high temperatures may arise inside the system causing an alarm signal. In the event of the alarm signal, immediately disconnect the device from the power source and contact Customer Care. Further operation of the device will be unsafe and may cause personal injury or property damage.

Safety Information

Site Selection

The system is designed to operate in a typical office environment. Choose a site that is:

- Clean, dry, and free of airborne particles (other than normal room dust).
- Well-ventilated and away from sources of heat, including direct sunlight and radiators.
- Away from sources of vibration or physical shock.
- Isolated from strong electromagnetic fields produced by electrical devices.
- In regions that are susceptible to electrical storms, we recommend you plug your system into a surge suppressor.
- Provided with a properly grounded wall outlet.
- Provided with sufficient space to access the power supply cords, because they serve as the product′s main power disconnect.

Equipment Handling Practices

Reduce the risk of personal injury or equipment damage by:

- Conforming to local occupational health and safety requirements when moving and lifting equipment.
- Using mechanical assistance or other suitable assistance when moving and lifting equipment.

- Reducing the weight for easier handling by removing any easily detachable components.

Power and Electrical Warnings

Caution: The power button, indicated by the standby power marking, DOES NOT completely turn off the system AC power; 5V standby power is active whenever the system is plugged in. To remove power from system, you must unplug the AC power cord(s) from the wall outlet.

- Do not attempt to modify or use an AC power cord if it is not the exact type required. A separate AC cord is required for each system power supply.
- This product contains no user-serviceable parts. Do not open the system.
- When replacing a hot-plug power supply, unplug the power cord to the power supply being replaced before removing it from the server.

Rack Mount Warnings

- The equipment rack must be anchored to an unmovable support to prevent it from tipping when a server or piece of equipment is extended from it. The equipment rack must be installed according to the rack manufacturer's instructions.
- Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.
- Extend only one piece of equipment from the rack at a time.
- To avoid risk of potential electric shock, a proper safety ground must be implemented for the rack and each piece of equipment installed in it.

Cooling and Air Flow

Installation of the equipment should be such that the amount of air flow required for safe operation of the equipment is not compromised.

NetWitness® Platform XDR

Version 12.0

Endpoint Agent Installation Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

December, 2022

Contents

Introduction to Endpoint Agent Installation **4**
- Supported Operating Systems 4
- Hardware Requirements 5
- Installation Flowchart 5

Generate an Endpoint Agent Packager **6**

Generate Endpoint Agent Installers **9**

Deploy and Verify Endpoint Agents **10**
- Deploying Agents (Windows) 10
  - Verifying Windows Agents 10
- Deploying Agent (Linux) 10
  - Verifying Linux Agents 11
- Deploying Agent (Mac) 11
  - Verifying Mac Agents 11

Uninstall Agents **13**
- Uninstalling Agent using UI (Supported from agent version 11.5.1 and later) 13
- Uninstalling Agent Manually 17
  - Uninstalling Windows Agent 17
  - Uninstalling Linux Agent 17
  - Uninstalling Mac Agent 17

Upgrade Agents **19**
- Upgrading 11.5.1 or Later Agents Versions Using UI 19
- Upgrading 11.3.x or Later Agents Versions 22

Recommendations for Installing Agents in Virtual Desktop Infrastructure (VDI) Environment **23**

Troubleshooting **24**
- Packager Issue 24
- Agent Upgrade via UI Issues 24
- Agent Uninstall via UI Issues 25

Introduction to Endpoint Agent Installation

Note: The information in this guide applies to Version 11.1 and later.

Hosts can be laptops, workstations, servers, physical or virtual, where a supported operating system is installed. An Endpoint Agent can be deployed on a host with either a Windows, Mac, or Linux operating system. The installation process involves:

1. (Optional) Configuring the Relay Server

Note: You must set up the default relay server before generating the Agent packager. Whenever the Relay server configuration is modified, agent policy is updated automatically. For more information on configuring the relay server, see *Endpoint Configuration Guide*.

2. Generating an agent packager
3. Generating the agent installer

You can run the agent installer specific to your operating system to deploy agents on the hosts. The agents collect endpoint data and tracking events from these hosts. It monitors key behaviors related to process, file, registry, console, and network, and forwards them as events to the Endpoint Server over HTTPs.

Note: The Endpoint agent can operate either in Insights or Advanced mode depending on the policy configuration. For more information, see the *NetWitness Endpoint Configuration Guide*.

Supported Operating Systems

Note: From version 12.0 and higher, NetWitness Endpoint agents run on ARM devices running on Windows 10 and 11.

Windows	**Linux** *(The agent software runs only on x86_64 architecture)*	**macOS**
Windows 11	CentOS 7.x and 8.x	macOS Monterey (12)
Windows 10 (32 and 64-bit) (up to version 21H2)	Red Hat Enterprise Linux 7.x and 8.x	macOS Big Sur (11)
Windows 8.1 (32 and 64-bit)	SUSE Linux Enterprise Server 12 SP3, 12 SP4, 12 SP5 and 15 SP1	macOS Catalina (10.15)
Windows 8 (32 and 64-bit)	Ubuntu 16.04 LTS, 18.04 LTS, and 20.04 LTS	macOS Mojave (10.14)
Windows 7 (32 and 64-bit)		macOS High Sierra (10.13)
Windows Server 2022 Windows Server 2022 Core		macOS Sierra (10.12)

Windows Server 2019 Windows Server 2019 Core		OS X El Capitan (10.11)
Windows Server 2016		OS X Yosemite (10.10)
Windows Server 2012 R2		OS X Mavericks (10.9)
Windows Server 2012		
Windows Server 2008 R2 (32 and 64-bit)		
Windows 10 Kiosk Mode (64-bit)		

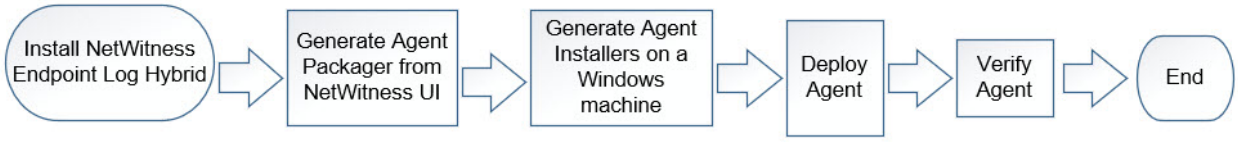
Hardware Requirements

The following are the minimum hardware requirements to run an agent in a host (laptops, workstations, servers, physical or virtual):

- 256 MB RAM
- 300 MB disk space
- Single-core CPU

Installation Flowchart

The following flowchart illustrates the Endpoint agent installation process:



Generate an Endpoint Agent Packager

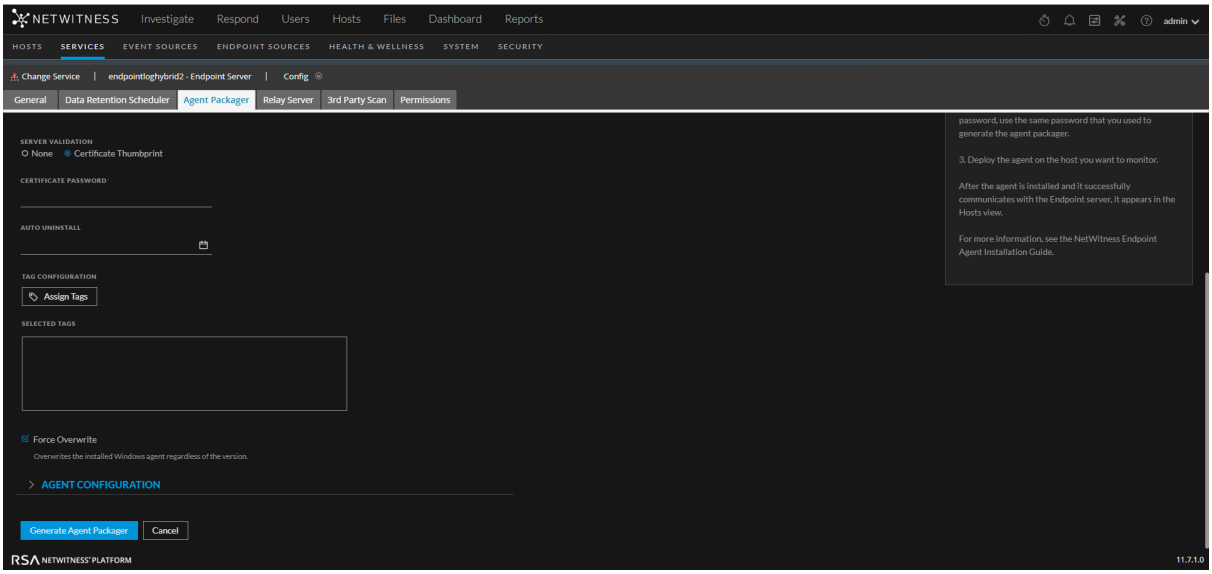
To generate an agent packager to collect endpoint data from hosts:

1. Log in to NetWitness.

   Type `https://<NW-Server-IP-Address>/login` in your browser to get to the NetWitness Login screen.

2. Click  **(Admin) > Services**.

3. Select the **Endpoint Server** service and click  **> View > Config > Agent Packager** tab.
   The Agent Packager tab is displayed.



4. Enter the values in the following fields:

Field	Description
Endpoint Server	Displays all the available Endpoint servers in the deployed.
Endpoint Server Forwarder (Optional)	The optional Endpoint Server Forwarder allows you to enter an alternative Fully Qualified Domain Name (FQDN) or IP address on which the sever can be reached in the case that agents need to go through a NAT or similar in order to reach the Endpoint Server. If specified forwarder is not available, agent will eventually fall back to the packaged address.
HTTPS Port	Port number. For example, 443.
Server Validation	Determines how the agent validates the Endpoint Server certificate: • None – The agent will not validate the server certificate. • Certificate Thumbprint – default selection. The agent identifies the server by validating the thumbprint of the Root CA of the server certificate.
Certificate Password	Password used to download the packager. The same password is used while generating the agent installer. **Note:** The password must be minimum seven characters long and a combination of uppercase and lowercase letters, numbers, and special characters. For example, **Admin@123**.
Auto Uninstall	Date and time the agent automatically uninstalls. You can leave it blank if not required.
Tag Configuration	When you click **Assign Tags** under Tag Configuration, you can do any of the following: • Create new tags and assign them to the hosts. • Select already existing tags and assign them to the hosts. For more information, see Investigate Hosts.
Force Overwrite	Overwrites the installed Windows agent regardless of the version. If this option is not selected, the same installer can be run multiple times on a system, but installs the agent only once. If you enable this option, make sure that you provide the same service name and driver service name as the previously installed agent, while creating a new agent. **Note:** If you want to force overwrite with MSI, run the following command: `msiexec /fvam <msifilename.msi>` After you move an agent from one deployment to another, using Force Overwrite to change the agent incurs an 8-hour delay in communication between the agent and its Endpoint Server on the new deployment. To eliminate the delay, uninstall the agent from the old deployment, and reinstall the agent on the new deployment.

Field	Description
Agent Configuration	
Note: The following Service and Driver fields are applicable only for Windows.	
Service	
Service Name	Name of the agent service. For example, NWEAgent.
Display Name	Display name of the agent service. For example, NWE Agent.
Description	Description of the agent service. For example, NetWitness Endpoint.
Driver	
Driver Service Name	Name of the driver service. For example, NWEDriver.
Driver Display Name	Display name of the driver service. For example, NWE Driver.
Driver Description	Description of the driver service. For example, NetWitness Endpoint Driver.
Generate Agent	Generates an agent packager.

5. Click **Generate Agent**.

   This downloads an agent packager (**AgentPackager.zip**) on the host where you are accessing the NetWitness user interface.

Generate Endpoint Agent Installers

To generate endpoint agent installers to deploy on hosts:

Note: Use a Windows machine to execute the agent packager file.

1. Unzip the **AgentPackager.zip** file. It includes the following:
   - **agents** folder – Contains executables for Linux, Mac, and Windows.
   - **config** folder – Contains configuration file and the certificates required to communicate between the Endpoint Server and the agent.
   - **AgentPackager.exe** file.
2. Run the **AgentPackager.exe** file as administrator by right-clicking the file and selecting **Run as administrator**.
3. Enter the same password used while generating the agent packager and press **Enter**.
   This creates the following installers in the root folder:
   - nwe-agent-package.exe (for Windows)
   - NWE000032.msi (for Windows 32-bit)
   - NWE000064.msi (for Windows 64-bit)
   - nwe-agent.pkg (for Mac)
   - nwe-agent.x86_64.rpm (for RPM based Linux 64-bit)
   - nwe-agent.x86_64.deb (for Debian based Linux 64-bit)

   Note: The MSI files should not be renamed.

Deploy and Verify Endpoint Agents

This section provides instruction on how to deploy and verify agents.

Note: By default, the agent is installed in the Insights mode. Depending on the policy assigned, the agent can operate in Insights or Advanced mode. Make sure you review the policy before deploying the agent. For more information, see *NetWitness Endpoint Configuration Guide*.

Deploying Agents (Windows)

To deploy the agent, run the **nwe-agent-package.exe** file on the hosts you want to monitor.

Verifying Windows Agents

After deploying the Windows agents, you can verify if a Windows agent is running by using any of the following methods:

- Using the NetWitness UI

  The Hosts view contains the list of all hosts with an agent. You can look for the host name on which the agent is installed.

  Note: Click **Hosts** or press F5 to refresh the list for latest data.

- Using Task Manager

  Open Task Manager and look for service name that you configured while generating the agent packager on the host machine.

- Using Services.msc

  Open `Services.msc` in run and look for the service name that you configured while generating the agent packager on the host machine.

Deploying Agent (Linux)

To deploy the agent on the hosts you want to monitor:

RPM based Linux

Run the **nwe-agent.x86_64.rpm** (for 64-bit) file. To run the command, open Terminal on the Linux machine and run the following command as `root`:

```
rpm -iv <installer file name>.rpm
```

For example, using the default installer file name, you can enter the following command:

`rpm -iv nwe-agent.x86_64.rpm` (for x86_64 architecture)

Note: To upgrade RPM based Linux agents, run **rpm -Uvh nwe-agent.x86_64.rpm**.

Debian based Linux

Run the **nwe-agent.x86_64.deb** (for 64-bit) file. To run the command, open Terminal on the Linux machine and run the following command as `root`:

```
dpkg -i <installer file name>.deb
```

For example, using the default installer file name, you can enter the following command:

`dpkg -i nwe-agent.x86_64.deb` (for x86_64 architecture)

(Enter the administrator password when prompted.)

Note: To upgrade Debian based Linux agents, run **dpkg -i nwe-agent.x86_64.deb**.

Verifying Linux Agents

After deploying the Linux agents, you can verify if a Linux agent is running by using any of the following methods:

- Using the NetWitness UI

  The Hosts view contains the list of all hosts with an agent.

  Note: Click **Hosts** or press F5 to refresh the list for latest data.

- Using Command Line

  Run the following command to get the PID:

  ```
  pgrep nwe-agent
  ```

- To check the NetWitness Endpoint version, run the following command:

  ```
  cat /opt/rsa/nwe-agent/config/nwe-agent.config | grep version
  ```

Deploying Agent (Mac)

To deploy the agent, run the `nwe-agent.pkg` file on the hosts you want to monitor. On macOS version Catalina (10.15) and higher, you need to move the `nwe-agent.pkg` file to a folder with sufficient access privileges (e.g., `/tmp`) and install the agent from there.

Verifying Mac Agents

After deploying the Mac agents, you can verify if a Mac agent is running by using any of the following methods:

- Using the NetWitness UI

  The Hosts view contains the list of all hosts with an agent.

  Note: Click **Hosts** or press F5 to refresh the list for the latest data.

- Using Activity Monitor

  Open Activity Monitor (/Applications/Utilities/Activity Monitor.app) and look for NWEAgent.
- Using Command Line

  Run the following command to get the PID

  ```
  pgrep NWEAgent
  ```
- To check the NetWitness Endpoint version, run the command:

  ```
  grep a /var/log/NWEAgent.log | grep NWEAgent | grep Version
  ```

Uninstall Agents

You can use one of the following methods to uninstall Endpoint agents. Select a method based on your current Endpoint agent version.

- Uninstalling Agent using UI (Supported from agent version 11.5.1 and later)
- Uninstalling Agent Manually

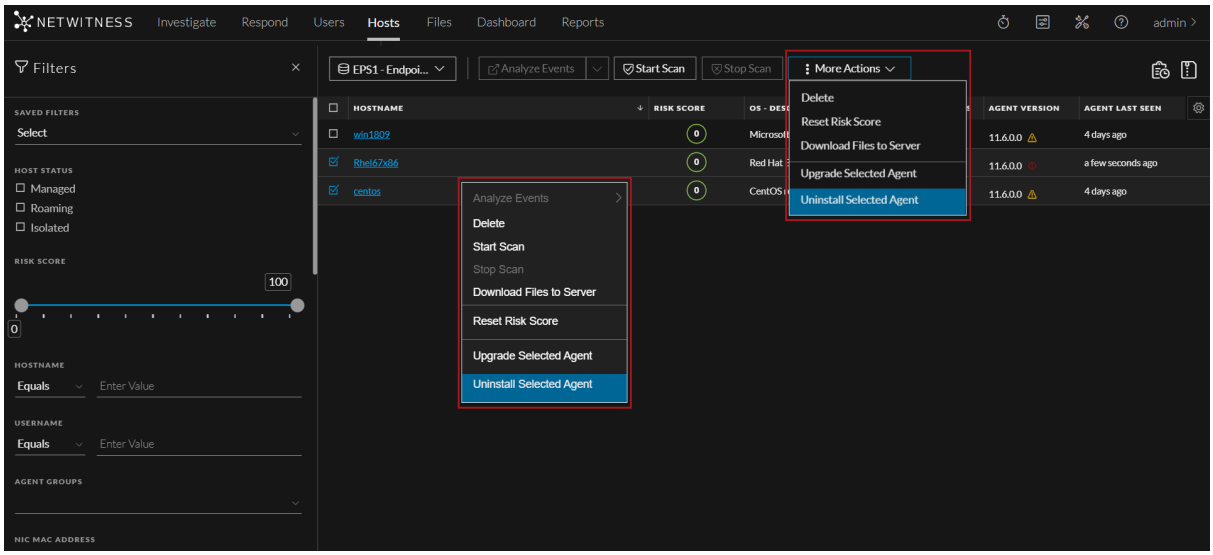
Uninstalling Agent using UI (Supported from agent version 11.5.1 and later)

Note: To uninstall an agent using UI, the agent must be a managed agent with version 11.5.1 or later.

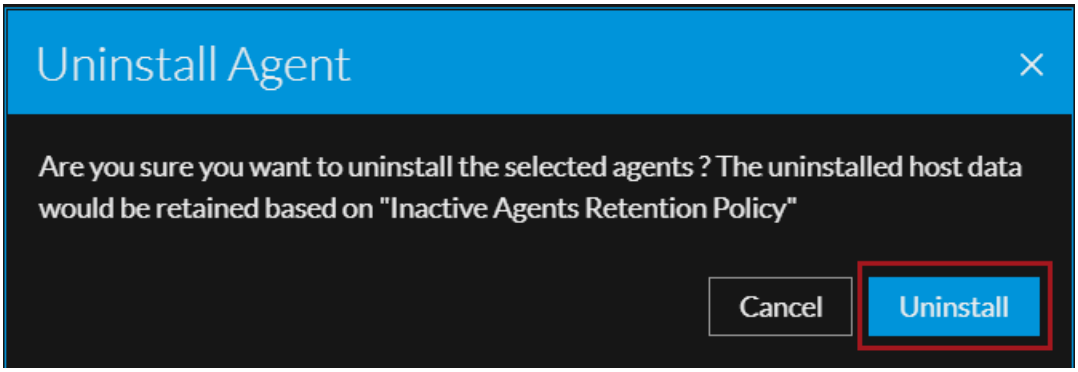
You can uninstall agent using UI by going to **Hosts** and performing one of the following options.

Uninstall one or more agents from Hosts view

1. Select one or more hosts from which you want to uninstall the agents.
2. Selected **More Actions > Uninstall Selected Agent** from the toolbar.



3. In the **Uninstall Agent** dialog, click **Uninstall.**

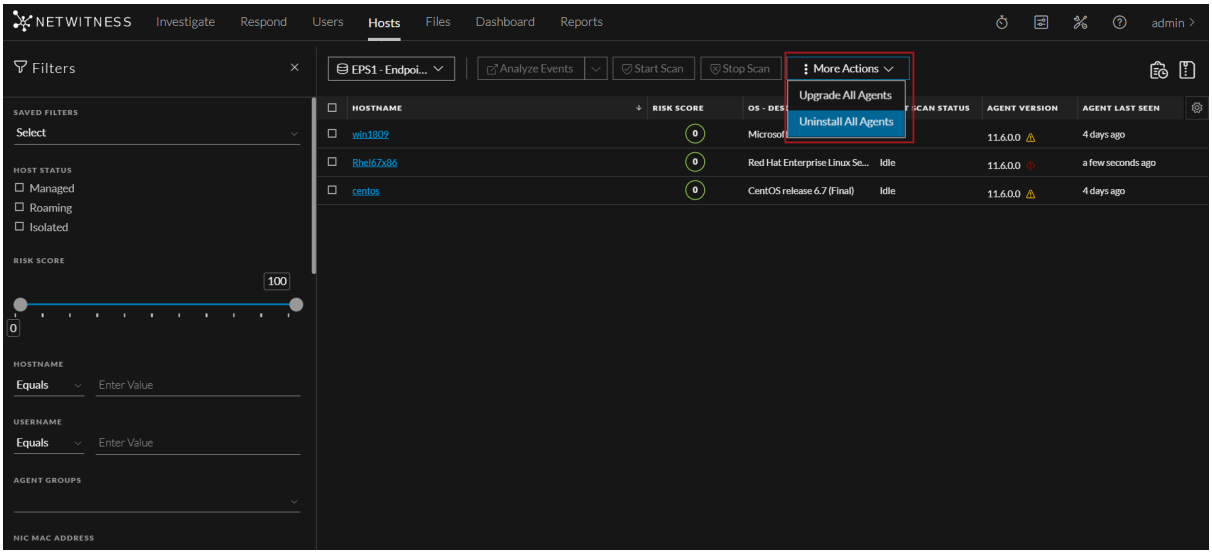


Uninstall all the agents from Hosts view

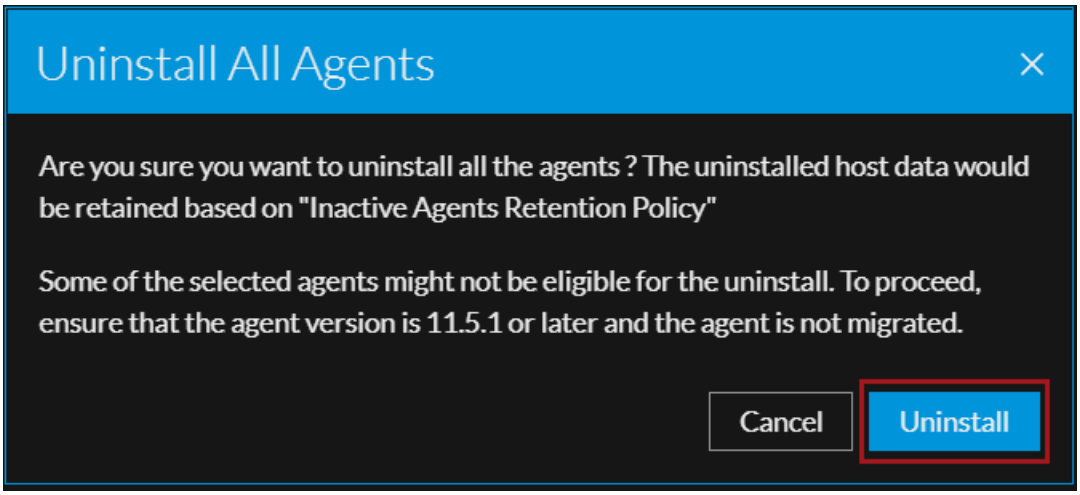
Note: Uninstall All Agents via UI is also supported for Broker service.

1. Select **More Actions > Uninstall All Agents** from the toolbar.

Note: For the uninstall all agents option, you do not need to select the hosts. **Upgrade All Agents** / **Uninstall All Agents** are the default options on the More Actions drop-down. If you select one or more hosts, the More Actions drop-down shows Upgrade Selected Agent / Uninstall Selected Agent as the available options.



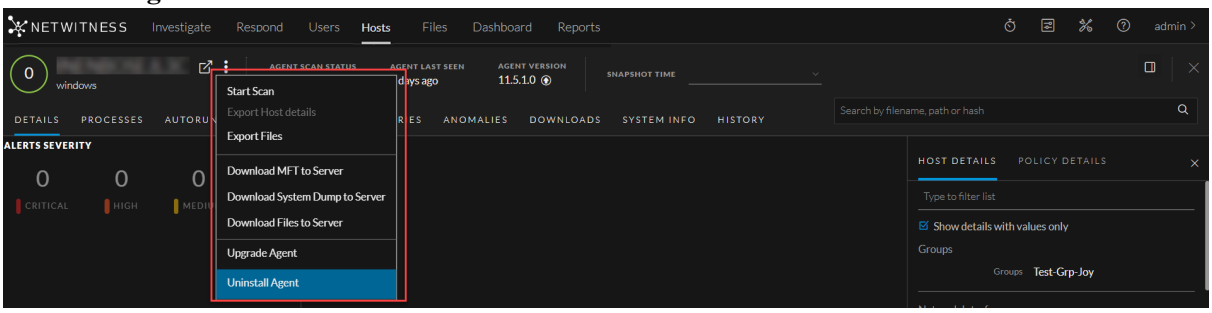
2. In the **Uninstall All Agents** dialog, click **Uninstall.**



Note: Uninstall All Agents process will require more time to complete depending upon the number of agents selected based on the filters applied.

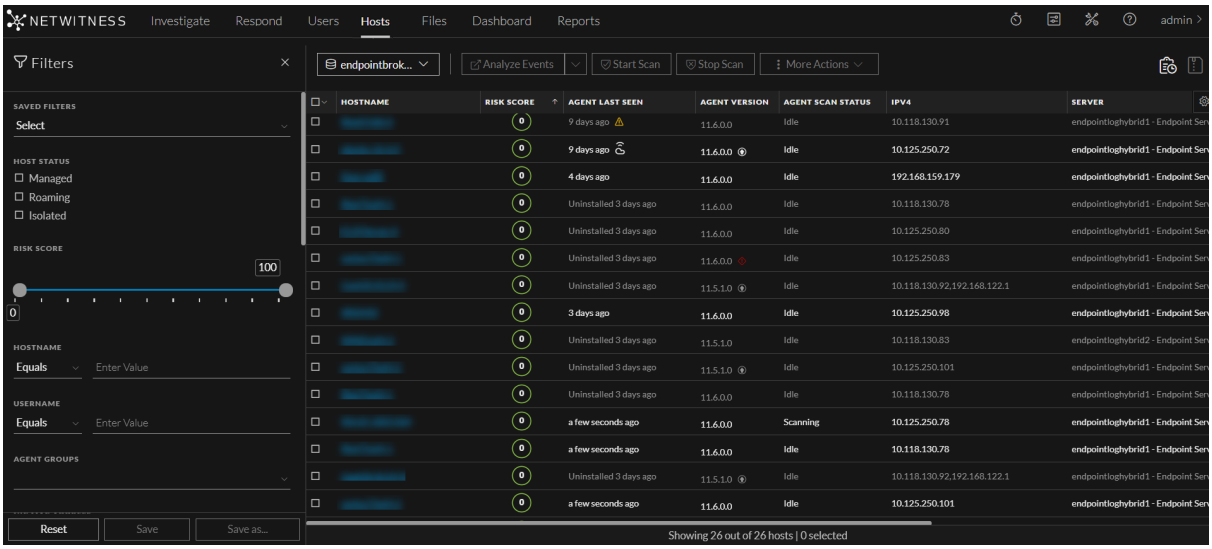
Uninstall an agent from Host details view

1. Select the hostname to open the host details, click ⋮ (**More Actions**) beside the hostname, and select **Uninstall Agent**.



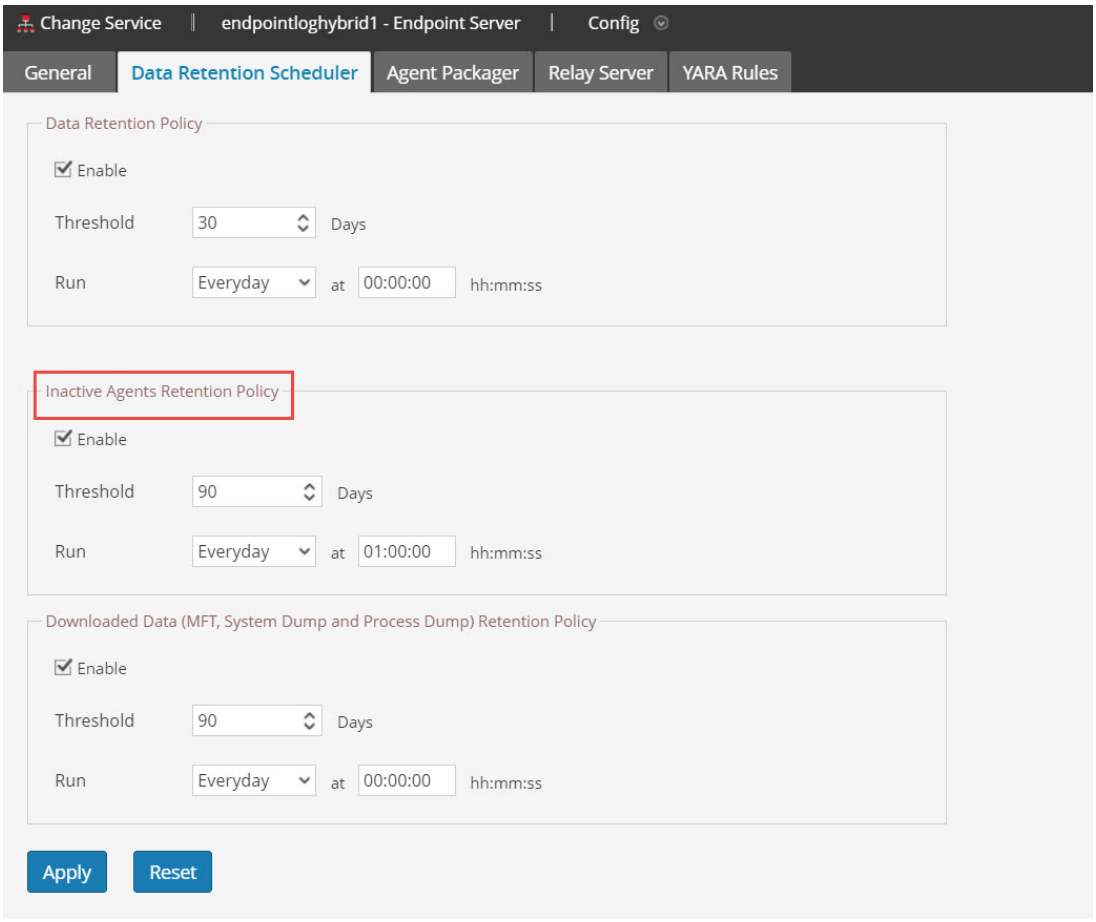
2. In the **Uninstall Agent** dialog, click **Uninstall.**

3. After the agent is uninstalled, **Agent Last Seen** is updated and the host is grayed out in the **Hosts** view.



Note: Once an agent is uninstalled, only **Delete** and **Analyze Events** actions can be performed on it.

IMPORTANT: If you are uninstalling agent using UI, the uninstalled host data would be retained based on **Inactive Agents Retention Policy.**



Uninstalling Agent Manually

Uninstalling Windows Agent

Run the following command as administrator:

```
msiexec /x{63AC4523-5F19-42F0-BC43-97C8B5373589}
```

Uninstalling Linux Agent

For RPM based Linux, run the following command as root:

```
rpm -ev nwe-agent
```

For Debian based Linux, run the following command as root:

```
dpkg -r nwe-agent
```

Uninstalling Mac Agent

Run the following commands:

1. `sudo launchctl unload /Library/LaunchDaemons/com.rsa.nwe.agent.daemon.plist`
2. `sudo rm -Rf /usr/local/nwe`
3. `sudo rm -Rf '/Library/Application Support/NWE'`
4. `sudo rm –Rf /Library/LaunchDaemons/com.rsa.nwe.agent.daemon.plist`
5. `sudo pkgutil --forget com.rsa.pkg.nwe`

Upgrade Agents

You can use one of the following methods to upgrade Endpoint agents. Select a method based on your current Endpoint agent version. If you have upgraded the endpoint server recently, you must restart it to see **Upgrade Selected Agent** and **Upgrade All Agents** options on the UI.

- Upgrading 11.5.1 or Later Agents Versions Using UI
- Upgrading 11.3.x or Later Agents Versions

Upgrading 11.5.1 or Later Agents Versions Using UI

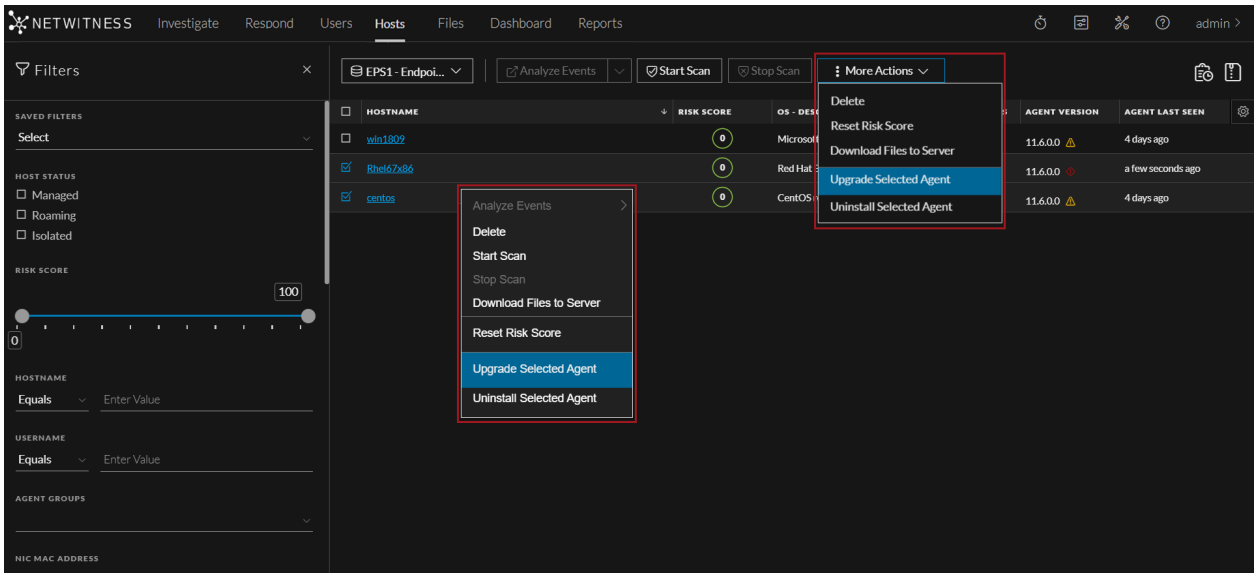
Note: To upgrade agent from UI, ensure the service user account has both `endpoint-server.agentupdate.manage` and `endpoint-server.ca.manage` permissions. For more information on how to assign roles and permissions, see "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.

From 11.6, you can upgrade the Endpoint agents using the UI. Agent version must be 11.5.1 or later to upgrade using the UI.

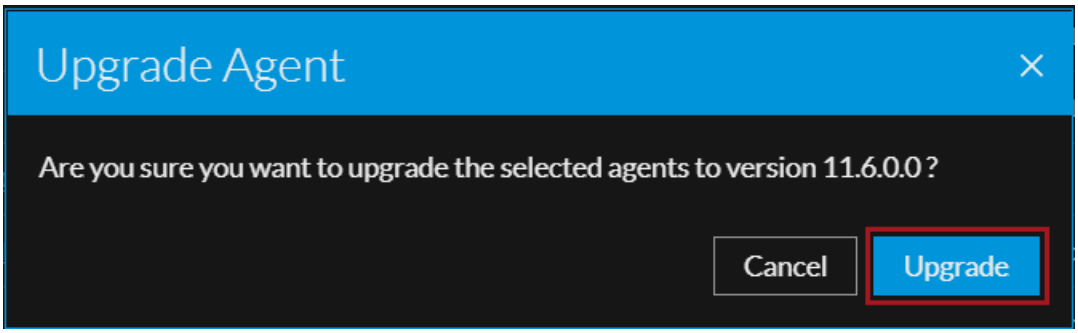
You can upgrade agents using UI by going to **Hosts** and performing one of the following options.

Upgrade one or more agents from Hosts view

1. Select one or more hosts and select **More Actions > Upgrade Selected Agent**



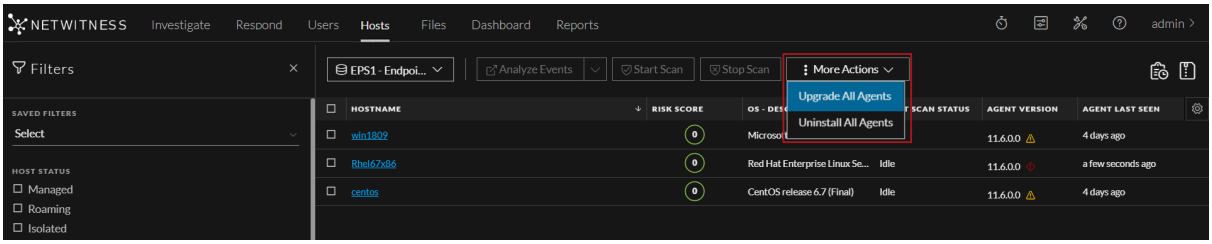
2. In the **Upgrade Agent** dialog, click **Upgrade**.



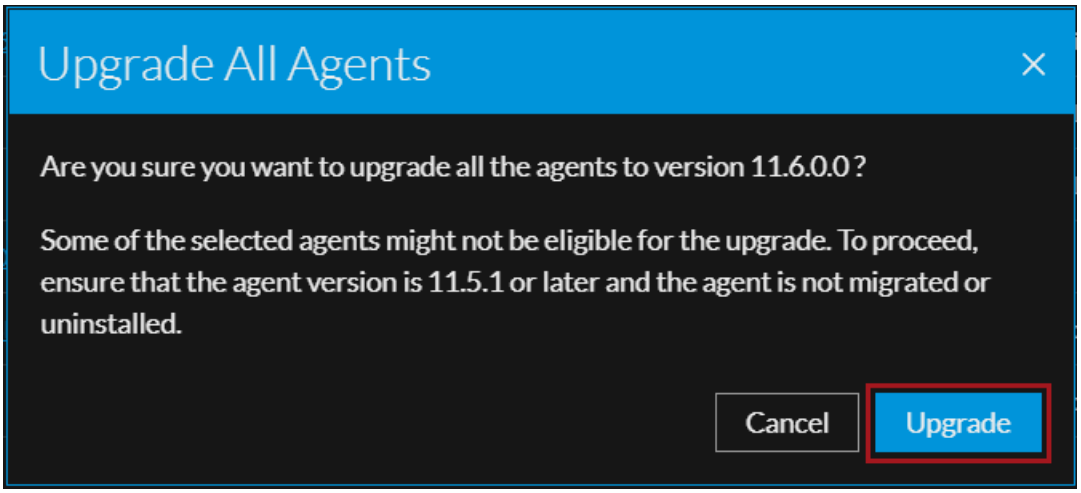
Upgrade all agents from Hosts view

1. Select **More Actions > Upgrade All Agents** from the toolbar, to perform bulk agent upgrade.

Note: For the upgrade all agents option, you do not need to select the hosts. **Upgrade All Agents** / **Uninstall All Agents** are the default options on the More Actions drop-down. When you select one or more hosts, the More Actions drop-down shows Upgrade Selected Agent / Uninstall Selected Agent as the available options.



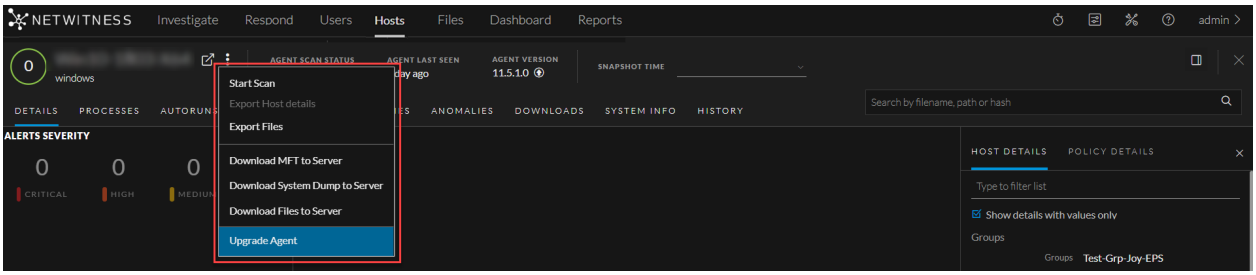
2. In the **Upgrade All Agents** dialog, click **Upgrade**.



Note: Upgrade All Agents process will require more time to complete, depending upon the number of agents selected based on the filters applied.

Upgrade an agent from Hosts details page

Select the hostname to open the host details, click ⋮ (**More Actions**) beside the hostname, and select **Upgrade Agent**.



1. In the **Upgrade Agent** dialog, click **Upgrade**.

Note: To upgrade an agent using UI, the agent must be a managed agent with version 11.5.1 or later.

The following upgrade status icons are displayed in the **Agent Version** column.

Icons	Upgrade Status	Description
	Upgrade Available	An agent upgrade is available.
	Agent Upgrade Pending	The agent has not yet received the upgrade command.
	In Progress	The agent is being upgraded.
	Failed	Agent upgrade failed. View details in the agent history page.
No Icon	No Upgrade Available	No upgrade available for the agent. Refer to agent eligibility.

You can view the list of commands issued to the agents (by the server or actions performed by any analyst) in the Host view and Host details. By default, commands are sorted based on the command time.

To view the commands:

1. Go to **Hosts**.
2. Do any one of the following,

   - To view all commands, click . You can also filter commands.
     The Agent History view is displayed.
   - To view commands specific to a particular host:
     - Click the host for which you want to view the commands.
     - In the Host details view, click **History** tab. You can also filter commands.
       The History view is displayed.

For more information on Agent History, see the "Investigating Hosts" topic in the *NetWitness Endpoint User Guide*.

Upgrading 11.3.x or Later Agents Versions

You can upgrade the 11.3.x and later versions of Endpoint agent to 11.6 or later.

Note: In a multi-server Endpoint deployment, during an agent upgrade, make sure that the correct Endpoint server is mentioned in the respective agent policy. In case the agent uses the default policy, ensure to use the agent packager downloaded from the respective Endpoint server to which it is communicating. Using Agent packager from different Endpoint server for agent upgrade will result in migrating the agents to another Endpoint server.

Note: For a subsequent installation or upgrade, use the same service and driver service name.

To upgrade from 11.3.x and later, download the 11.6 agent packager, deploy and verify agents. For more information, see Generate an Endpoint Agent Packager and Deploy and Verify Endpoint Agents.

Upgrade from 4.4.0.9 and 4.4.1.x is supported only for version 11.3. For more information, see *NetWitness Endpoint 4.4.1.x to NetWitness Platform 11.3 Migration Guide*.

To upgrade an agent from 4.4.0.x to 11.6, uninstall the agent and perform a fresh installation.

Recommendations for Installing Agents in Virtual Desktop Infrastructure (VDI) Environment

Agent ID is generated based on various parameters, such as security identifier (SID) and SMBIOS Universal Unique Identifier (UUID). A SMBIOS UUID is a 128-bit number used to uniquely identify a host.

Note: While cloning the VDI image where an agent is already installed, the agent ID automatically changes for Windows and Mac agents if `uuid.action = keep` is not set in the .vmx file. For more information, see Configure a Virtual Machine to change the UUID.
For Linux agents, the agent ID does not change automatically on VDI clone.

When you clone a VDI image:

- If you do not change the agent ID for each VDI clone, make sure that the SMBIOS UUID remains the same.
- If you change the agent ID for each VDI clone, make sure that the SMBIOS UUID is also changed.

To avoid duplication of agent IDs, make sure that the SMBIOS UUID changes on the following VDIs:

- Citrix XenServer
- VMWare Workstation
- VMware vCloud Director

  For more information, see VMware Knowledge Base.
- vCenter hosted ESXi Server

To get the SMBIOS UUID on a Windows virtual host, execute the following command:

```
wmic csproduct get UUID
```

Troubleshooting

This section provides information about possible issues when using the NetWitness Endpoint.

Packager Issue

Issue	Failed to generate the agent installers.
Explanation	Some encryption software may create additional files that fails to generate the agent installers.
Resolution	Copy the packager to a machine that does not have antivirus or encryption software and then generate the agent installers.

Issue	Failed to generate agent installers for MAC.
Explanation	Agent packager `AgentPackager.exe` fails to generate MAC agent installer (`nwe-agent.pkg`) with the error message "Failed to generate table of content for package" or "Failed to create config file C:\AgentPackager (4)\agents\mac\Plugins\NWEInstallerPlugin.bundle\Contents\Resources\config.cfg".
Resolution	Run the `AgentPackager.exe` as administrator by right-clicking the file and selecting **Run as Administrator**.

Issue	Agent packager generates temporary agent installers for MAC.
Explanation	Agent packager `AgentPackager.exe` generates MAC agent installer as `nwe-agent_tmp.pkg` instead of `nwe-agent.pkg`.
Resolution	Run the `AgentPackager.exe` as administrator by right-clicking the file and selecting **Run as Administrator**. The MAC agent package `nwe-agent.pkg` will be generated as expected

Agent Upgrade via UI Issues

Issue	Agent upgrade not available.
Explanation	1. The agent version might not be supported for upgrade from UI. Agent version has to be 11.5.1 or later for upgrade via UI. 2. Logged in user may not have appropriate permission for upgrade from UI. 3. Agent version is up-to date.
Resolution	1. Use the manual upgrade method for upgrading the version.

	2. Use admin user which have following permissions, endpoint-server.agentupdate.manage and endpoint-server.ca.manage. 3. No upgrade required. Agent version is already up-to date and no further upgrade is available.

Issue	Agent upgrade is in Pending state.
Explanation	The hosts for which the command is shown in pending could be in offline or inactive state.
Resolution	Ensure Hosts/agents are communicating with Endpoint server directly or via Relay server for it to receive the upgrade command from server. Verify "Agent Last seen" time in Host listing page.

Issue	Agent upgrade failed.
Explanation	1. Agent upgrade fail with any of the following reasons: • Service Name or Driver name mismatch • Checksum mismatch • Installer size mismatch 2. Agent installer could not be created.
Resolution	1. Retry upgrading the agent. If it continues to fail, use the manual upgrade method. 2. Check the permissions for the user initiated the upgrade command. User can be viewed in **Host** > **Agent History** page. User should have following permissions, `endpoint-server.agentupdate.manage` and `endpoint-server.ca.manage`.

Issue	Linux agent upgrade fails with an error, sudo not found.
Resolution	Make sure that the sudo package is installed on the agent machine.

Agent Uninstall via UI Issues

Issue	Agent uninstalled failed.
Explanation	Agent uninstall fails due to several unknown reasons.
Resolution	1. Retry uninstalling the agent. 2. Use the manual uninstall method. For more information, see Uninstalling Agent Manually.

NetWitness® Platform XDR

NetWitness Endpoint 4.4.0.x to NetWitness Platform 11.4 Migration Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Introduction **4**

Migration Flowchart 4

How to Find the List of Documents 5

Migrating NetWitness Endpoint 4.4.0.x to NetWitness Platform 11.3 and Later **6**

Task 1 - Plan your NetWitness Deployment 6

Task 2 - Set up NetWitness Platform 11.3 and Later 6

Task 3 - Configure NetWitness Platform for Endpoints 7

Task 4 - Import NetWitness Endpoint 4.4.0.x Configurations 7

Task 5 - Set up Other NetWitness Endpoint 4.4.0.x Configurations 7

Task 6 - Deploy Agents 8

Task 7 - Verify the Agent Migration 8

Importing NetWitness Endpoint 4.4.0.x Configurations to NetWitness Platform **9**

Introduction

This guide provides instructions on how to migrate an existing NetWitness Endpoint 4.4.0.x to NetWitness Platform 11.3 and later.

The migration broadly involves:

- Planning and setting up NetWitness Platform. Contact your account manager for sizing the deployment. For more information on sizing guidelines, see the *Virtual Host Installation Guide*.
- Importing file status, certificate status, and blocked hashes from NetWitness Endpoint 4.4.0.x to NetWitness 11.3 and later.

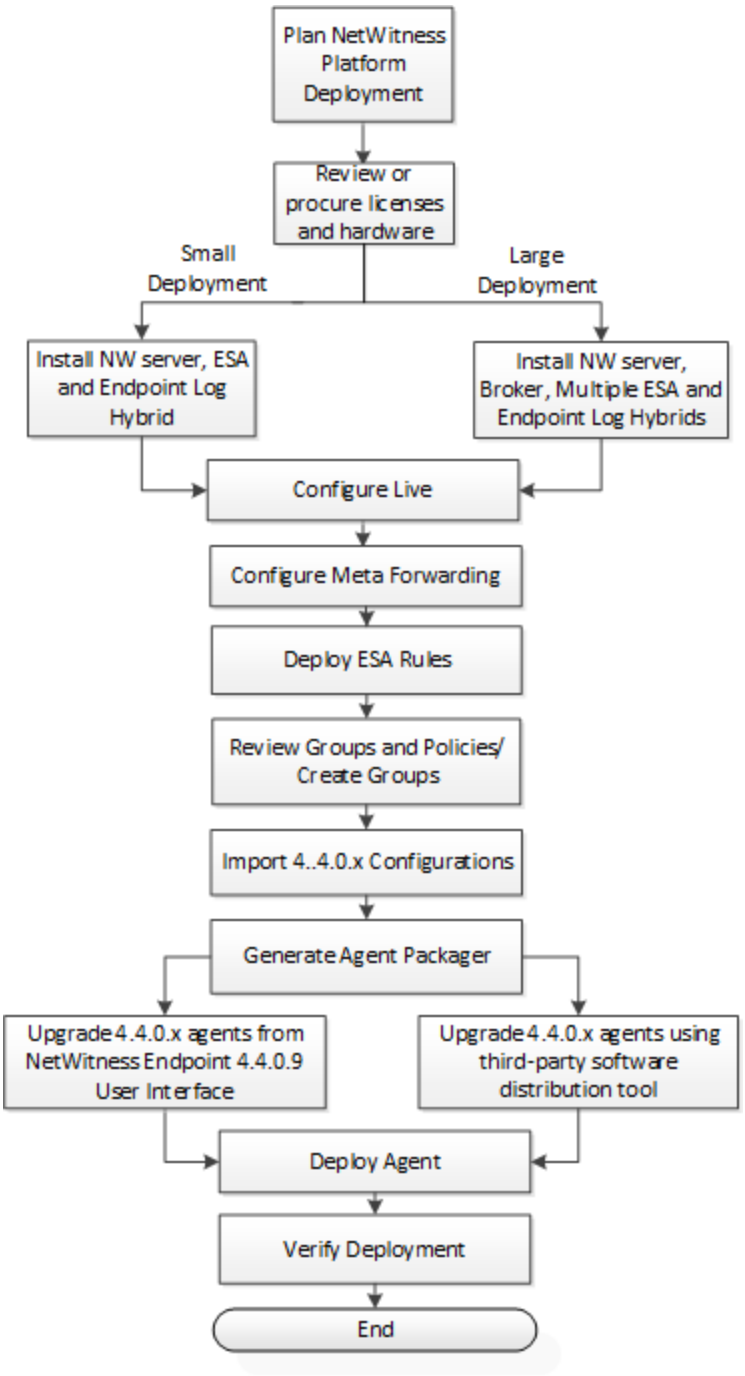
  Note: Endpoint scan and tracking data in the SQL database, users, IIOCs, custom queries, bad certificates, bad domains, bad IPs, and bad file hashes in NetWitness Endpoint 4.4.0.x cannot be migrated.

- Upgrading 4.4.0.x agents to 11.3 and later.

For more information on the new features of NetWitness Endpoint, see the *Release Notes*.

Migration Flowchart

The following flowchart illustrates the migration process:



Note: The instructions for each step is provided in the sections below.

How to Find the List of Documents

If you want to view the NetWitness Platform product documentation, use the following link:
https://community.rsa.com/t5/rsa-netwitness-platform/ct-p/netwitness-documentation/version/11.3

Migrating NetWitness Endpoint 4.4.0.x to NetWitness Platform 11.3 and Later

This topic describes tasks required to migrate NetWitness Endpoint 4.4.0.x to NetWitness 11.3 and later.

Task 1 - Plan your NetWitness Deployment

1. Review the Endpoint architecture and choose one of the following deployments based on number, distribution, and location of endpoints, and data collected from the agents. For more information, see the "NetWitness Endpoint Architecture" topic in the *Deployment Guide*.
   - Small deployment includes NetWitness Server, an Endpoint Log Hybrid, and Event Stream Analysis (ESA).
   - Large deployment includes NetWitness Server, one or more Endpoint Log Hybrids, one or more ESA.

     If you have multiple Endpoint Log Hybrids, install an Endpoint Broker. Once installed, it automatically queries all Endpoint servers in your deployment and provides a consolidated view of all Endpoint servers.

     Note: RSA recommends that you install the Endpoint Broker on the NetWitness Broker host. It is also supported on a separate host, NetWitness Server, and on an Endpoint Log Hybrid.
2. Review existing licenses and hardware, and procure them as required.

   Make sure you have the required hardware for your deployment. For more information, see the "Supported Hardware" topic in the *Physical Host Installation Guide*.

   If you have existing NetWitness Endpoint 4.4.0.x licenses, you do not need to procure a new license. Your existing NetWitness Endpoint 4.4.0.x licenses will be available on myRSA as a NetWitness 11.3 and later license, with a different serial/license number.

   For more information, see the *Licensing Management Guide*.

Task 2 - Set up NetWitness Platform 11.3 and Later

1. Configure the ports on your firewall. For more information, see the "Network Architecture and Ports" topic in the *Deployment Guide*.
2. Install the following NetWitness components:
   - NetWitness Server
   - Endpoint Log Hybrid
   - ESA
   - Endpoint Broker (for more than one Endpoint Log Hybrids)

   For more information, see the *Physical Host Installation Guide*.

3. If you have NetWitness Endpoint 4.4.0.x licenses:
   - With NetWitness - Map your entitlements to the existing license server.
   - Without NetWitness:

     a. Install NetWitness (see Task 2 - Set up NetWitness Platform 11.3 and Later)

     b. Obtain License Server ID

     c. Map entitlements
4. Verify if licenses are reflected on **ADMIN > System > Licenses**.

Task 3 - Configure NetWitness Platform for Endpoints

- Configure your RSA Live account and make sure the File Reputation service is enabled. For more information, see the *Live Services Management Guide*.
- Create users and assign appropriate roles. For more information, see the *System Security and User Management Guide*.
- Configure Endpoint Meta forwarding on Endpoint Log Hybrid. For more information, see the *NetWitness Endpoint Configuration Guide*.
- Deploy ESA rules.

  The existing NetWitness Endpoint 4.4.0.x IIOCs are available as out-of-the-box Application rules and automatically available on installation.

  You must configure the ESA Correlation server with an Endpoint Concentrator and deploy the ESA content for risk score calculation. For more information, see the *ESA Configuration Guide*.
- Review the default Agent Endpoint (EDR) policy and create groups as required. If you want to enable agents for log collection, review and apply the Windows Log policy.

  For more information, see the *NetWitness Endpoint Configuration Guide*.

Task 4 - Import NetWitness Endpoint 4.4.0.x Configurations

Import file status, certificate status, and blocked hashes from NetWitness Endpoint 4.4.0.x to NetWitness Platform. For more information, see Importing NetWitness Endpoint 4.4.0.x Configurations to NetWitness Platform.

Task 5 - Set up Other NetWitness Endpoint 4.4.0.x Configurations

You have to manually set up the following configurations:

- Deploy Blacklisted IP addresses and other feeds that are relevant for your deployment from RSA Live through the NetWitness user interface.

- (Optional) For any other external threat feeds, such as blacklisted IP address, domain, and checksum, that you may want to use to tag endpoint metadata, see the "Create a Custom Feed" topic in the *Live Services Management Guide*.

  For example, to notify an analyst about any communication from a host or file to a certain blacklisted IP address, domain, or hash, create a feed on the Log Decoder, and tag appropriate sessions for investigation and alerting .

- (Optional) Review custom IIOCs and write Endpoint rule. For more information, see the "Custom Endpoint Rule for Risk Scoring" topic in the *NetWitness Endpoint Configuration Guide*.

Task 6 - Deploy Agents

1. Generate an agent packager from NetWitness 11.3 and later.
2. Copy the agent packager (`AgentPackager.zip`) to a Windows machine and generate the 11.3 and later agent installers.
3. Do one of the following to upgrade 4.4.0.x agents to 11.3 and later:
   - If you have NetWitness Endpoint 4.4.0.9 Console Server, copy the agent installers to the NetWitness Endpoint Console Server, and upgrade from the NetWitness Endpoint user interface.
   - If you have 4.4.0.0 or 4.4.0.8, copy the agent installers, and use the third-party software distribution tool.
4. Deploy the agents.

For more information, see the *NetWitness Endpoint Agent Installation Guide*.

Task 7 - Verify the Agent Migration

After the agent migration, verify the following:

- Agents are able to communicate with the Endpoint Server and are listed in the Investigate > Hosts view.
- Perform a scan and make sure that the snapshot details are displayed in the Investigate > Host Details view.
- Hosts metadata is available in Investigate > Navigate and Events view.

  If the Windows Log collection is enabled, make sure that the Windows logs are available in the Navigate and Events view.

- File reputation, file status, risk scores are available in Hosts and Files view.

For detailed information, see the *NetWitness Endpoint Configuration Guide*.

Importing NetWitness Endpoint 4.4.0.x Configurations to NetWitness Platform

You can import file status, certificate status, and blocked hashes from NetWitness Endpoint 4.4.0.x to NetWitness using the MigrationHelper python script.

Note: The MigrationHelper python script must be run only on a Windows host.

You can download the script from RSA Link:

RSA NetWitness > **Downloads** > **RSA NetWitness Platform** > **Version 11.3** > **Tools**.

Prerequisites

To run the python script:

- Install Python 3.6.x or later on a Windows host that can connect to the NetWitness Endpoint 4.4.0.x primary database.
- Install pyodbc by downloading the wheel file from https://pypi.org/project/pyodbc/#files, and run the following command:

  ```
  pip install wheel-file.whl
  ```

- If json and os.path libraries are not available on Python installation, install these libraries by downloading the corresponding wheel file from https://pypi.org/, and run the following command:

  ```
  pip install wheel-file.whl
  ```

Import File and Certificate Status

Note: If the certificate status is graylisted in NetWitness Endpoint, this status is not exported as graylist is not supported for certificates in NetWitness Platform 11.3 and later.

1. Run the `MigrationHelper` python script.

   Note: Run this script from any host that has access to NetWitness Endpoint primary database.

2. Enter the following:
   a. Database server host name or IP address (for example, 10.40.40.10)
   b. Database name (for example, ECATPrimary)
   c. Database credentials
3. Enter the path to store the exported files and press **Enter**. Make sure that the path exist. The file and certificate status are exported to JSON files.
4. Log in to the Context Hub server and copy the exported file to the `/var/netwitness/contexthub-server/data/` directory.
5. On the NW server, run the `nw-shell` command from the command line.

Note: Make sure all Endpoint servers on NetWitness 11.3 and later are online while importing data.

6. Run the `login` command and enter the credentials.
7. Connect to the Context Hub server using the following command:

   ```
   connect --service contexthub-server
   ```

8. Run the following commands to import the file status:

   ```
   cd contexthub/file/status/import
   show
   invoke <file path>/FileStatus.json
   ```

   Note: `<file path>` is the path in the Context Hub server where the file is saved. The Context Hub server is located in the ESA primary host.

9. Run the following commands to import the certificate status:

   ```
   cd contexthub/certificate/status/import
   show
   invoke <file path>/CertificateStatus.json
   ```

   Note: `<file path>` is the path in the Context Hub server where the file is saved. The Context Hub server is located in the ESA primary host.

10. Check the progress of the import in the `/var/log/netwitness/contexthub-server/contexthub-server.log` file.

    Once the import is complete, a message `Imported File status successfully` or `Imported Certificate status successfully` is displayed in the log file.

If you want to unblock the imported 4.4.0.x blocked files:

1. On the NW server, run the `nw-shell` command from the command line.
2. Run the `login` command and enter the credentials.
3. Connect to the Context Hub server using the following command:

   ```
   connect --service contexthub-server
   ```

4. Run the following commands to unblock the file status:

   ```
   cd contexthub/file/status/unblock
   invoke <checksum of blocked file>
   ```

NetWitness® Platform XDR

Version 12.0

NetWitness Endpoint Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

NetWitness Endpoint **7**
About NetWitness Endpoint 7
Endpoint Agent Data Flow 8
Agent Modes **11**
Endpoint Log Hybrid Configuration **13**
Deploying Endpoint Application Rules and ESA Correlation Rules **17**
Custom Endpoint Rule for Risk Scoring 17
Add a custom Application Rule 18
Add a custom ESA Rule 19
Add the rule to RiskConfig 20
Configuring Metadata Forwarding **23**
Start Metadata Forwarding to the Log Decoder 24
Stop Metadata Forwarding to the Log Decoder 25
Remove Metadata Forwarding 25
Endpoint Metadata Mappings 25
JSON Schema for Metadata Mappings 25
View the Metadata Mappings 26
Add or Modify Metadata Mappings 28
View the Custom Metadata Mappings 28
Endpoint Sources **29**
Groups 29
Policies 29
Group Ranking 30
Example 1 30
Example 2 30
Example 3 32
Default Agent Endpoint (EDR) Policy 33
Default Windows Log Policy 34
Default File Log Policy 34
Creating Groups and Policies **36**
Create a Group 36
Construct a Policy 39
Create an EDR Policy 39
Create a Windows Log Policy 43
Create a File Log Policy 45
Replace Windows SFTP Agents 49

Managing Groups **50**
View Group Details 50
Filter Endpoint Groups 50
Edit a Group 51
Delete a Group 51
Managing Policies **53**
View Policy Details 53
Filter Policies 53
Edit a Policy 54
Delete a Policy 55
Conflict Resolution 56
Change Policy Ordering for Groups **57**
Edit Ranking 57
Simulation Examples 59
Agent Endpoint Policies Examples 60
File Log Policies Simulation Example 63
The SIMULATE Slider 68
69
Manage Role Permissions at Endpoint Server Level **71**
To configure permissions at the endpoint server level: 71
Managing Inactive Agents **73**
Configure Retention Policy for Memory Dumps and MFT **74**
(Optional) Installing and Configuring Relay Server **75**
Installing the Relay Server 77
Installation Media 77
Relay Server Host System Requirements 77
Configuring the Relay Server 79
Configure YARA **81**
Configure OPSWAT **83**
Integrating NW Endpoint 4.4.0.2 or Later with NW Platform **85**
Endpoint References **88**
General Tab 89
Data Retention Scheduler Tab 91
Agent Packager Tab 94
Quick Look 94
Relay Server Tab 96
Workflow 96
What do you want to do? 96

Quick Look 96
Features 97
Endpoint Sources - Groups 99
Create Group 101
Define Group 102
Apply Policies 103
Ranking Groups 104
Endpoint Sources - Policies 106
Create Policy 108
Panels for Log File Policy 109
Define Connection Settings 109
Define File Policy Settings 111
Define Policy Panel for Agent Endpoint Policy 113
Define Policy Panel for Windows Logs Policy 117
General Tab 120
General Tab 122
Troubleshooting 124
Agent Communication Issues 124
Packager Issues 125
Health and Wellness Issues 125
Endpoint Issues 125
Disk Usage and Mongo Issues 126
Log Decoder Issues 126
File Log Policy Issues 127
Invalid Policy or Bad Connection Issues 127
Reset File Collection Bookmarks 127
Missing Log Collectors and Event Sources in the User Interface 128
Relay Server Issues 128
Test Connection Issues 128
129
Installation Issues 129
YARA Issues 130
OPSWAT Issues 130
Files can not be scheduled for scan 130
OPSWAT not configured on all endpoint servers 130
OPSWAT scan results not getting updated 131
Appendices 132
Reset File Collection Bookmarks 132
Construct a JSON File to Identify Agents and Event Source Types for Reset 132
Reset Bookmarks 133

How to Find Agent IDs and Source Types 134
How to Find Endpoint Service IDs 135
Currently Supported File Log Event Source Types 136
Specify UNC (Universal Naming Convention) Paths 137
Secure the UNC Path Location 137
Share a folder between machines in a domain 137
Share a folder between machines in a Workgroup 141

NetWitness Endpoint

Note: The information in this topic applies to NetWitness Version 11.1 and later.

About NetWitness Endpoint

NetWitness Platform provides an endpoint detection and response solution that continuously monitors the behavior of all endpoints in and outside the network to provide deep visibility and analysis of executables and processes. It helps to detect new, unknown, and targeted attacks, highlights suspicious activity for investigation, exposes anomalous behaviors, and determines the scope of compromise to help analysts respond to advanced threats faster. During investigation, the analyst can use the visual indication of threat level to assess the risk of endpoints.

As part of this solution, NetWitness introduces **Endpoint Log Hybrid** that:

- Collects and manages endpoint (host) data from Windows, Mac, and Linux hosts.
- Collect log files and Windows logs from Windows hosts.
- Generates metadata to correlate endpoint data with sessions from other events sources, such as logs and network.

Analysts can:

- Perform instant scans for detailed insights of the host behavior at any point in time.
- Analyze the scope of the attack across hosts and network through integrated metadata.
- Quickly triage and focus their investigation by managing suspect and legitimate files.
- Perform multiple checks of file legitimacy to determine if a file is malicious, including checking file certificates and hashes.
- Blacklist malicious files and then block them across all hosts in the network to prevent future execution of this file on any host.
- Download Master File Table (MFT), system dump, and process dump for forensic investigation.
- Isolate host from the network to safely investigate possible threats within the host.

Endpoint Log Hybrid receives data from the Endpoint Agents. The following services run on the Endpoint Log Hybrid:

- **Endpoint Server**: Manages data received and stores it in a database. It parses the events, generates metadata, and forwards it to the Log Decoder through protobuf. You can deploy up to 6 Endpoint Log Hybrid hosts. For a consolidated view of all endpoint data from multiple Endpoint Log Hybrid hosts, install the Endpoint Broker. You can add only one broker in a NetWitness platform deployment which serves up to 6 Endpoint Log Hybrid hosts. Multiple Endpoint Log Hybrids are required to share certificates by copying them from the primary Endpoint Log Hybrid to the secondary Endpoint Log Hybrid using the `SCP` command.

  Note: You may need to install your Endpoint Server on separate hardware from your Log Decoder. If you are only using NW Platform for collecting and analyzing logs, you can co-locate your Endpoint Server on the same physical hardware as your Log Decoder. For more information, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness Platform*.

  If you exceed these guidelines, the amount of disk space usage and CPU might become so high as to create alarms for your Endpoint Server in Health and Wellness. If you notice this, and are running both log collection and EDR scans, you can use Throttling to control the amount of data coming into the Log Decoder.

  If that doesn't help, NetWitness recommends that you move your Endpoint Server onto separate hardware from that used by your Log Decoder.

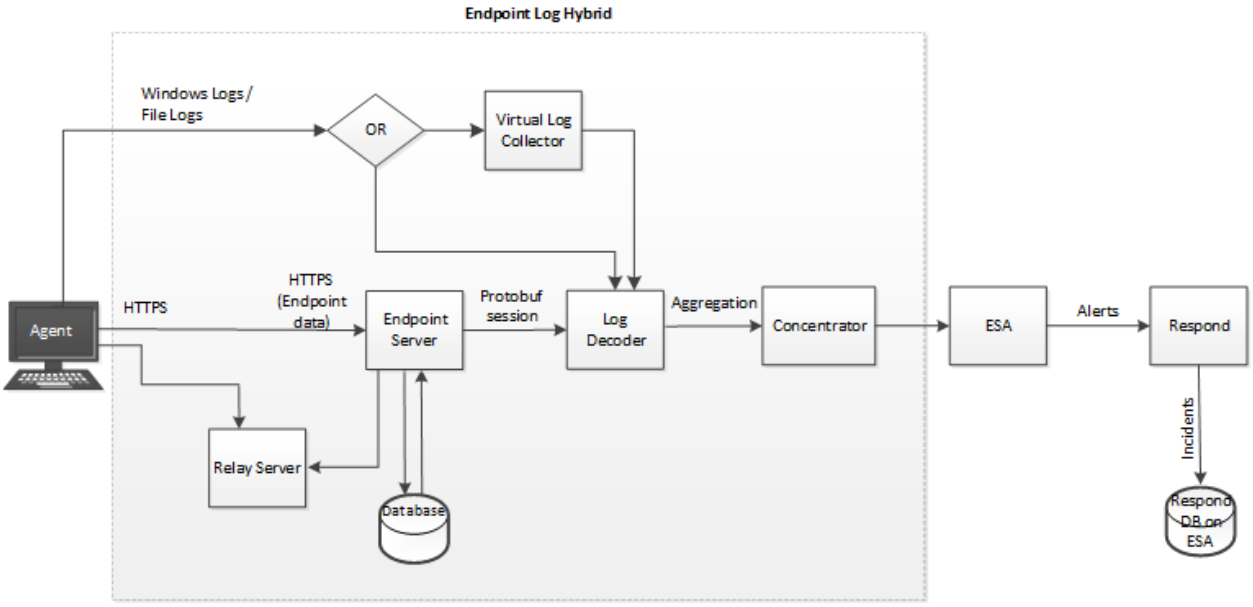
- **Log Decoder**: Captures data from the Endpoint Server and processes the metadata.
- **Concentrator**: Aggregates metadata from the Log Decoder and makes it available for all upstream components like Investigate, Reporting Engine, Respond, and Event Stream Analysis similar to NetWitness Decoder and Concentrator.
- **Log Collector**: Collects logs from all event sources that are supported for the log collection in the NetWitness Platform.

In addition to the above services, the Endpoint Log Hybrid leverages the following services:

- **Event Stream Analysis (ESA)**: Creates alerts from ESA rules for Endpoint data.
- **Endpoint Broker**: Provides a consolidated view of all Endpoint servers in a multiple Endpoint Log Hybrid deployment.

Endpoint Agent Data Flow

The following figure shows the endpoint data flow from the agent to the NetWitness:

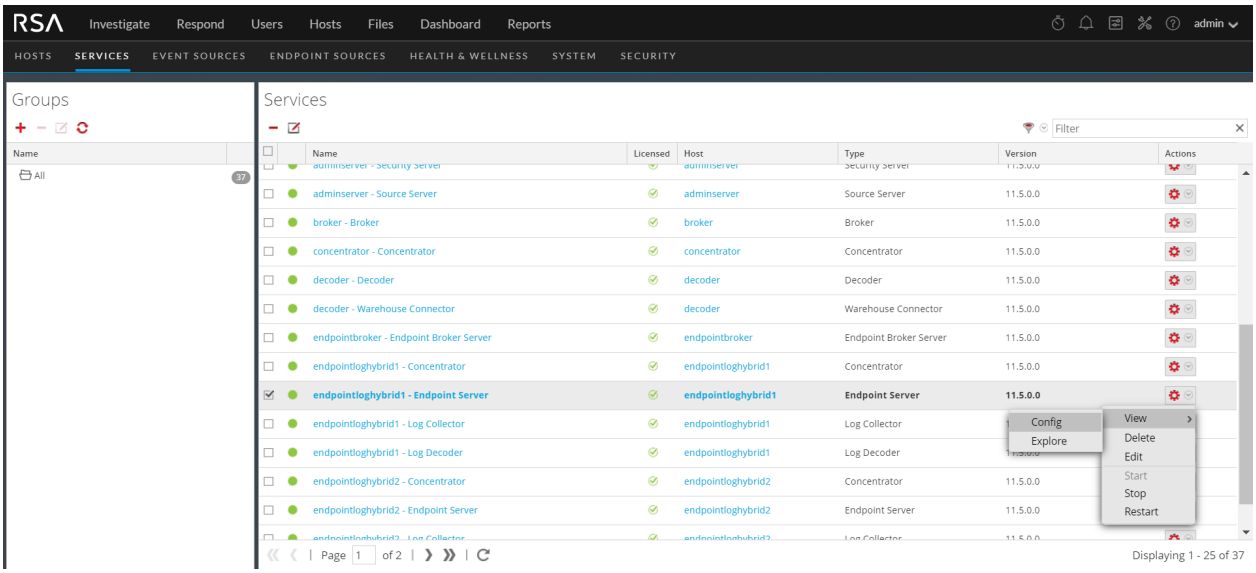


The *Hosts and Services Getting Started Guide* provides the information you need to understand and install all the NetWitness services.

Basic configuration involves:

- Installing agents on hosts
- Deploying the ESA rules from the Endpoint Rule Bundle
- Creating groups and policies
- Configuring Endpoint metadata forwarding and retention policies
- Defining health and wellness policies to monitor Endpoint Server
- Installing and configuring Relay Server

You can configure the required settings in the NetWitness user interface under Administration Services Config view (**(Admin) > Services > Endpoint Server > Config**).


RSA
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Groups
Name
All
37
Services
Filter
Name
Licensed
Host
Type
Version
Actions
adminserver - Source Server
adminserver
Source Server
11.5.0.0
broker - Broker
broker
Broker
11.5.0.0
concentrator - Concentrator
concentrator
Concentrator
11.5.0.0
decoder - Decoder
decoder
Decoder
11.5.0.0
decoder - Warehouse Connector
decoder
Warehouse Connector
11.5.0.0
endpointbroker - Endpoint Broker Server
endpointbroker
Endpoint Broker Server
11.5.0.0
endpointloghybrid1 - Concentrator
endpointloghybrid1
Concentrator
11.5.0.0
endpointloghybrid1 - Endpoint Server
endpointloghybrid1
Endpoint Server
11.5.0.0
endpointloghybrid1 - Log Collector
endpointloghybrid1
Log Collector
endpointloghybrid1 - Log Decoder
endpointloghybrid1
Log Decoder
endpointloghybrid2 - Concentrator
endpointloghybrid2
Concentrator
11.5.0.0
endpointloghybrid2 - Endpoint Server
endpointloghybrid2
Endpoint Server
11.5.0.0
Config
Explore
View
Delete
Edit
Start
Stop
Restart
Page
1
of 2
Displaying 1 - 25 of 37

Agent Modes

Note: The information in this topic applies to NetWitness Version 11.3 and later.

In NetWitness 11.3 and later, the Endpoint agent can operate either in Insights or Advanced mode depending on the policy configuration. For more information on policy configuration, see the *NetWitness Endpoint Configuration Guide*. You can have both Insights and Advanced agents in a single deployment.

There is no license required for the Insights agent. However, you must procure a license for an Advanced agent. For more information on licensing, see the *Licensing Management Guide*.

The following table list the features supported for Insights and Advanced agents:

Feature	Insights Agent	Advanced Agent
Full System Scan	No	Yes (agent version 11.6.1 or higher)
Scan data - Processes, Autoruns, Files, Drivers, Libraries, and System Information	Yes - Windows, Mac, and Linux	Yes - Windows, Mac, and Linux
Tracking data - Process, File, Registry, Network, and Console	No	Yes - Windows and Mac Registry and Console events are applicable only for Windows.
Expanded network visibility	Yes - Windows	Yes - Windows
Anomaly detection - Image Hooks, Kernel Hooks, Registry Discrepancies, and Suspicious Threads	No	Yes - Windows
Windows log collection	Yes	Yes
File log collection	Yes - Windows	Yes - Windows
Threat detection content - ESA, Application Rules	Yes	Yes
Analysis of downloaded file	Yes	Yes
File status - Whitelist, Blacklist, Graylist, and Neutral	Yes	Yes
File remediate (Block)	No	Yes - Windows
Process visualization	No	Yes
Live connect	Yes	Yes

Feature	Insights Agent	Advanced Agent
File reputation service (Third-party lookup)	Yes	Yes
Risk score for hosts	No	Yes
MFT, process dump, and system dump download	No	Yes - Windows
Automatic file download	Yes	Yes
Automatic Memory DLL Download	No	Yes
Manual file download	No	Yes
Network isolation	No	Yes - Windows
Relay server	Yes	Yes

Endpoint Log Hybrid Configuration

This topic provides the high-level tasks required to configure the Endpoint Log Hybrid.



Tasks	Description
Install the Endpoint Log Hybrid	For a physical host: See *Install Endpoint Log Hybrid* topic under Post Installation Tasks on *Physical Host Installation Guide* For a virtual host: See *Install Endpoint Log Hybrid* topic under Step 5. Post Installation Tasks on *Virtual Host Setup Guide*.
Configure one or more Endpoint Log Hybrids	For a physical host: See *Configuring Multiple Endpoint Log Hybrids* topic under Post Installation Tasks on *Physical Host Installation Guide* For a virtual host: See *Configuring Multiple Endpoint Log Hybrids* topic under Step 5. Post Installation Tasks on *Virtual Host Setup Guide*.
Add hosts to the Endpoint Log Hybrid	For a physical host: See *Add Hosts to the Endpoint Log Hybrid* topic under Post Installation Tasks on *Physical Host Installation Guide* For a virtual host: See *Add Hosts to the Endpoint Log Hybrid* topic under Step 5. Post Installation Tasks on *Virtual Host Setup Guide*.
Deploy Application and ESA Rules	See Deploying Endpoint Application Rules and ESA Correlation Rules.

Tasks	Description
Configuring Metadata Forwarding	Similar to logs and packets, you can view Endpoint metadata in the Navigate and Events view. You can also generate reports and alerts for the Endpoint data. By default, the Endpoint Meta option is disabled. The agent must be installed with the Endpoint Meta option enabled to forward metadata.

Tasks	Description
Install Agents on Hosts	The Endpoint agent installer is generated using the Packager tab under  **(Admin) > Services > Config > Endpoint Server** from the NetWitness user interface. The Packager is a zip file that contains executables and configuration files for generating agent installer for Linux, Mac, and Windows operating systems. You can install only one version of the agent on a host. If you have a previous version of an agent installed (for example, 4.4), uninstall this agent to install the 12.0 agent. After the agent is installed, it appears on the **Hosts** view. By default, the Endpoint data is posted for the first time. To collect subsequent Endpoint data, you have to either schedule a scan or perform ad hoc scan. It retrieves data, such as drivers, processes, DLLs, files (executables), services, autoruns, security information, anomalies, system configurations, and scripts found on the host. **Note:** If a Windows host has proxy setup enabled and configured, the Endpoint agent communicates to the Endpoint server through proxy. If a host has both Automatic (Auto detect and Setup Script) and Manual proxy settings enabled and configured, the order of precedence is as follows: 1. Manual proxy setup 2. Automatic proxy setup - Use setup script 3. Automatic proxy setup - Automatically detect settings
Install and Configure the Relay Server	See (Optional) Installing and Configuring Relay Server.

Tasks	Description
3rd Party Scan	Configure 3rd party scans like YARA and OPSWAT. See Configure YARA and Configure OPSWAT
Permissions	Grant or Revoke role permissions for the selected Endpoint Server. See Manage Role Permissions at Endpoint Server Level
Manage Groups and Policies	To efficiently manage and update endpoint agent configurations, you can group the agents, and manage their behavior using policies.
Enable Reputation Status	Reputation Status is enabled by default in an NetWitness Platform 11.3 and later deployment and displays information about the file. For troubleshooting, see the *Live Services Guide*.
Risk Score	Risk Score is calculated and obtained from NetWitness Respond for hosts and files. For more information, see the *NetWitness Respond Configuration Guide*.
	Define data retention policies to optimally store and manage the Endpoint data based on the age of the Endpoint data or the storage size. By default, 30 days of agent data is retained.
Managing Inactive Agents	By default, agents (including all the collected Endpoint data) that have not communicated with the Endpoint Server for 90 days will be automatically deleted.
Configure Retention Policy for Memory Dumps and MFT	Define retention policy to optimally store and manage downloaded system dump, process dump, and MFT. By default, 90 days of data is retained.
Investigate Endpoint data	You can investigate the Endpoint data in the **Hosts** and **Files** views. For more information, see the *NetWitness Endpoint User Guide*.

Deploying Endpoint Application Rules and ESA Correlation Rules

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The existing IIOCs from NetWitness Endpoint 4.4.0.x are now available as OOTB Endpoint Application rules tagged as Indicators of Compromise, Behaviors of Compromise, Enablers of Compromise, and Analysis.File. Application rules for Endpoint are automatically available on installation of NetWitness 11.3 and later.

For Endpoint risk score, every Application rule must have an ESA rule that generates alerts used for the risk score calculation. A set of OOTB ESA rules are available as Endpoint Rule Bundle. You must specify the Endpoint data sources (Concentrators) and deploy the ESA Rules from the Endpoint Rule Bundle. For more information, see "Deploy Endpoint Risk Scoring Rules on ESA" section in the *ESA Configuration Guide*.

If the Application rule key value matches with ESA rule then an alert is triggered which is used to compute the risk score and an incident is raised when risk score exceeds the defined threshold limit.

Note: If you are upgrading from an existing Endpoint Log Hybrid to 11.3 or later, you must deploy the Application rules from RSA Live. During deployment, you must specify Endpoint Log Hybrid Log Decoder service. In case of multiple Endpoint servers, select all the Endpoint Log Hybrid Log Decoder services. For more information, see the *Live Services Management Guide*.

You can view the application rules that are deployed in  **(Admin) >Services > Endpoint Log Hybrid - Log Decoder** > **Config** > **App Rules** and application rules that were triggered in **Investigate** > **Navigate**> **Endpoint Log Hybrid - Concentrator** > **App rules**.

The Endpoint ESA rules generate alerts with the severity; Critical, High, and Medium. You can view the alerts on:

- Risk Details tab - You can view Critical, High and Medium alerts for a host or file on **Hosts** > **Risk Details** or **Files** > **Risk Details**.
- Respond view : You can view only critical and high severity alerts on **NetWitness Respond** > **Alerts**.

Custom Endpoint Rule for Risk Scoring

If you have custom IIOCs in NetWitness Endpoint 4.4.0.x, you need to create these custom Endpoint rules. Once you have created your custom Application rule, you must create the custom ESA Rule for risk score calculation and update the RiskConfig file in MongoDB.

To create a custom Endpoint rule, perform the following tasks:

1. Add a custom Application Rule
2. Add a custom ESA Rule
3. Add the rule to RiskConfig

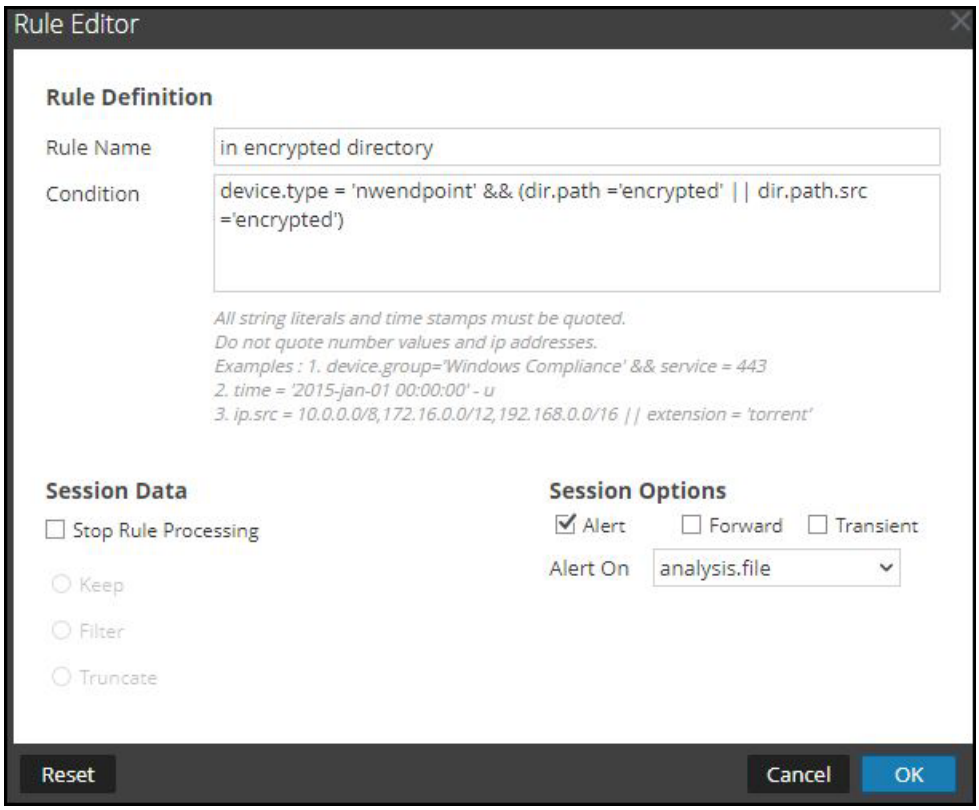
Add a custom Application Rule

To add a custom application rule:

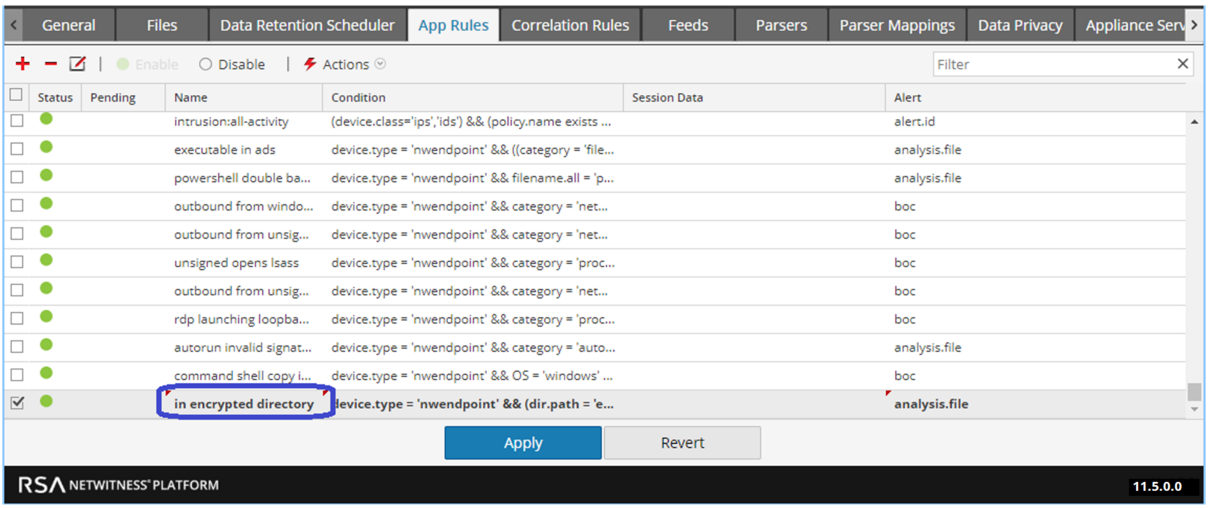
1. Complete steps 1-11 in "Configure Application Rules" topic of *Decoder and Log Decoder Configuration Guide*.

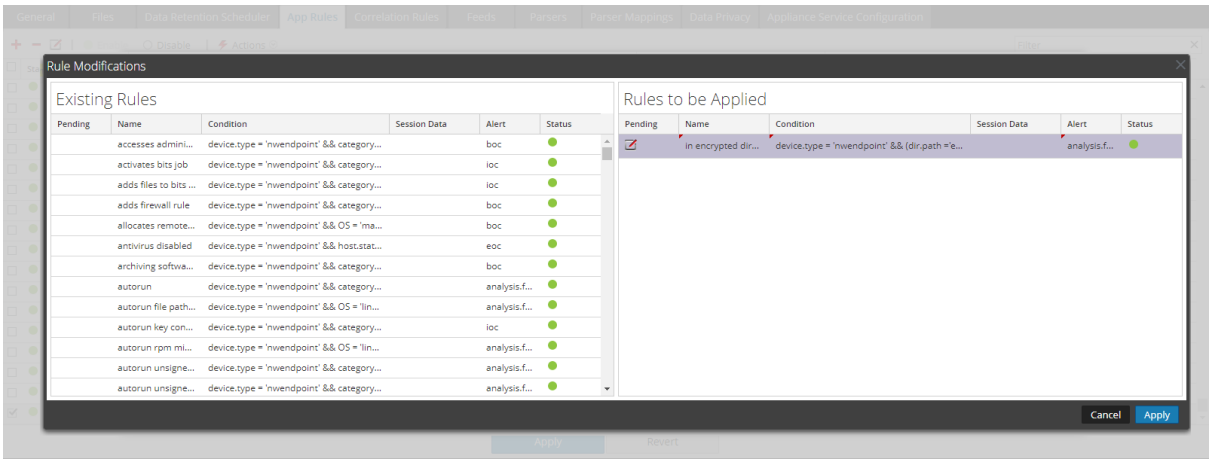
   Note: You must be familiar with the metakeys tagged as (Indicators of Compromise, Behaviors of Compromise, Enablers of Compromise, and Analysis.File) on which an alert will be generated. In the example below, the alert is generated on Analysis.File metakey for In Encrypted Directory rule.

   Following is an example of an Application Rule created for In Encrypted Directory alert.



2. After a custom application rule is added successfully, select the newly created rule (For example, In Encrypted Directory alert) and click **Apply**.





In case of multiple Endpoint servers, you must create this custom Application rule on every Endpoint Hybrid Log Decoder service.

Add a custom ESA Rule

To add a custom ESA rule, perform the following.

1. SSH to the Admin Server.
2. Create a new JSON file (for example,`in encrypted directory.json`) with the custom ESA rule definition in the below format.

```
{
"id": "In Encrypted Directory","key": "analysis.file",
"value": "in encrypted directory",
"title": "In Encrypted Directory",
"type": "ENDPOINT",
"enabled": true,
```

```
"description": "End Point rule for In Encrypted Directory",
"severity": "MEDIUM"
}
```

The following table describes the fields that define a rule.

Fields	Description
id	The name of the ESA Rule. For example, In Encrypted Directory.
key	The metakey on which an alert would be generated. For example, alert is generated on analysis.file metakey for In Encrypted Directory rule.
value	Specify the value. The value must exactly match with the App rule name. For example, in encrypted directory.
title	The name of the alert. For example, In Encrypted Directory.
type	Specify the type of rule. For custom endpoint rule, the type must be ENDPOINT.
enabled	The status of the rule. Specify true, if the rule should be considered for risk scoring.
description	The description of the rule.
severity	The severity of the rule; critical, high or medium.

3. To enter shell mode of nw-shell, execute the following command:

   ```
   nw-shell
   ```

4. Connect to ESA correlation-server using the following command:

   ```
   connect --service correlation-server
   ```

5. Login to the ESA correlation-server using the following command:

   ```
   login
   ```

   Note: You must provide Administrator credentials.

6. Navigate to the API xpath using the following command:

   ```
   cd correlation/keyvalue/settings/set
   ```

7. Run the API using the following command:

   ```
   invoke --file <absolute-path-to-rule-definition-file>
   ```

   Note: You must specify the absolute path to the rule definition file. For example, `invoke --file /root/rule.json`

Add the rule to RiskConfig

After you create the custom Application rule and the ESA rule, you must update the RiskConfig in mongoDB.

To update the riskconfig file, perform the following:

1. SSH to Admin Server.
2. Create a JavaScript file (For example, `in-encrypteddirectory- rule.js)` with the custom ESA rule definition in the below format.

```
db.risk_rule.insertMany(
[ {
"name" : "In Encrypted Directory",
"enabled" : true,
"handler" : "Default",
"entities" : {

},
"metas" : {
"File" : [
{
"meta" : "checksum_src",
"name" : "filename_src",
"weight" : NumberInt(100)
}
],
"Host" : [
{
"meta" : "agent_id",
"name" : "alias_host",
"weight" : NumberInt(100)
}
]
},
"_class" : "com.rsa.asoc.respond.pipeline.risk.rules.AlertScoringRule"
} ]
)
```

The following table describes the fields that define a rule.

Field	Description
name	The name of the ESA rule.

Field	Description
enabled	The flag to enable or disable risk scoring. Specify true to enable risk scoring.
handler	The value of this should be Default.
entities	The value of this should be empty.
metas > Files > meta	The metakey for a file for which score should be calculated.
metas > Files > name	The name of the metakey of the file identity.
metas > Files > weight	By default the weight value is 100.
metas > Host > meta	The metakey for a host for which score should be calculated.
metas > Host > name >	The name of the metakey of the host identity.
metas > Host > weight	By default the weight value is 100.
_class	This is used for internal purpose, do not change.

3. Insert the new rule into the riskconfig file on mongoDB using following command:

```
mongo respond-server --authenticationDatabase admin -u deploy_admin -p
<deploy_admin-user-password> in-encrypted-directory-rule.js
```

4. Confirm if ESA rule is updated successfully in the riskconfig, using following command

```
mongo respond-server --authenticationDatabase admin -u deploy_admin -p
<deploy_admin-user-password> --eval "db.risk_rule.find({ "name": /.*In
Encrypted Directory.*/i })"
```

5. Restart the Respond server for the changes to take effect.

```
service rsa-nw-respond-server restart
```

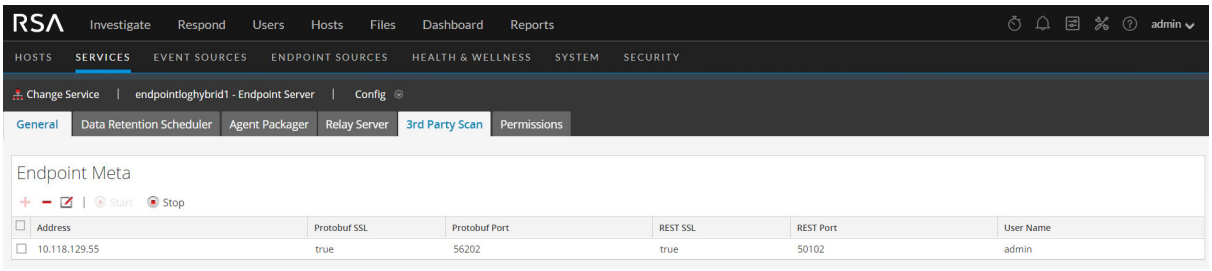
After you create a custom Endpoint rule and update the risk configuration file, whenever an event is generated for the new rule (For example, In Encrypted Directory) an alert will be generated and the risk score is calculated for the host and file.

Configuring Metadata Forwarding

To view the metadata, you must enable the metadata forwarding while installing the Endpoint Log Hybrid. The Endpoint metadata is displayed in the NetWitness Investigate (**Navigate** and **Events** views) similar to Logs and Packets. For information on metadata mappings, see Endpoint Metadata Mappings.

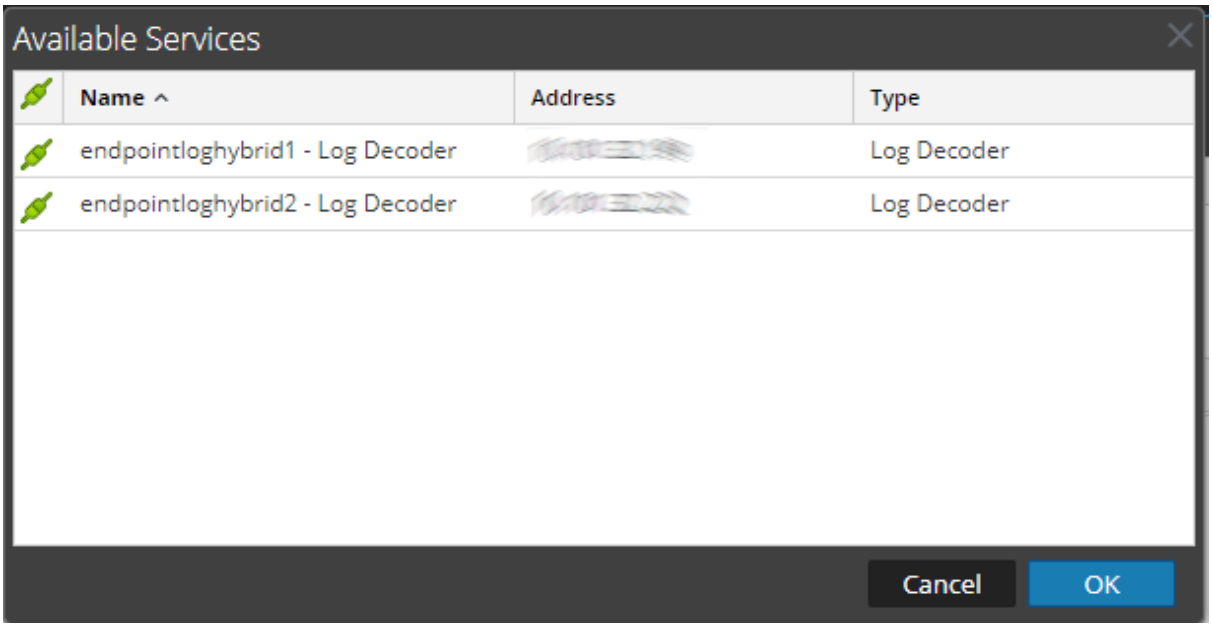
To configure metadata forwarding:

1. Go to **(Admin) > Services**.
2. In the Services view, select the **Endpoint Server** service.
3. Click and select > **View** > **Config**.
4. Click the **General** tab.



5. Click in the toolbar.

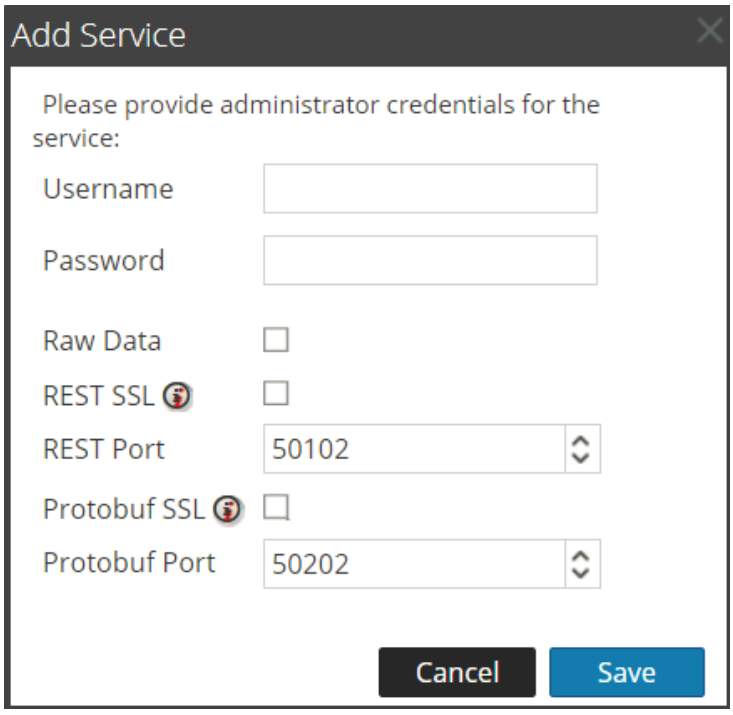
   The Available Services dialog is displayed.



6. Select a Log Decoder service and click **OK**.

   The Add Service dialog is displayed.

Note: You can add only one Log Decoder service.



7. Enter the administrator credentials for authentication.
8. (Optional) If you enable Raw Data, a brief summary of the session is forwarded along with the metadata.
9. (Optional) If you have enabled SSL on the REST port in the Log Decoder, select the **REST SSL** option. By default, the REST port for non-SSL is 50102.
10. Select the **Protobuf SSL** option to enable SSL on Protobuf. By default, the Protobuf port is 50202.
11. Click **Save**.

After configuring the metadata forwarding, make sure to:

- Start the capture on the Log Decoder
- Start the aggregation on the Concentrator
- Add the Log Decoder as a service in the **Concentrator**

Start Metadata Forwarding to the Log Decoder

1. In the Endpoint Meta config > General view, select the service.
2. Click Start
   The Endpoint Server starts forwarding the metadata to the Log Decoder.

Stop Metadata Forwarding to the Log Decoder

1. In the Endpoint Meta config > General view, select the service.
2. Click Stop.
   The Endpoint Server stops forwarding the metadata to the Log Decoder.

Remove Metadata Forwarding

Note: Make sure you stop the service, before removing the metadata forwarding.

1. In the Endpoint Meta config view, select the service.
2. Click −.
3. Click **Apply**.

Endpoint Metadata Mappings

You can view the default metadata mappings or modify the metadata mappings for endpoints.

JSON Schema for Metadata Mappings

All metadata mappings is configured using the JSON schema. The following is a sample JSON schema:

```
{
"metaKeyPairs" : [
     {
     "metaKeyPairsCategory" : "",
     "keyPairs" : [
          {
          "endpointJpath" : "",
          "metaName" : "",
          "type" : "",
          "enabled" : true
          },
          {
          "endpointJpath" : "",
          "metaName" : "",
          "type" : "",
          "enabled" : true
```

```
            }
        ]
        }
]
}
```

The following APIs are used to view or modify the metadata mappings:

- `get-default` - Returns the default configurations for the endpoint metadata mappings.
- `get-custom` - Returns the custom configurations for the endpoint metadata mappings.
- `set-custom` – Helps customize the endpoint metadata mappings.

View the Metadata Mappings

To view the endpoint metadata mappings:

1. On the NW server, run the `nw-shell` command from the command line.
2. Run the `login` command and enter the credentials.
3. Connect to the Endpoint Server using the following command:
   `connect --host <IP address> --port <number>`

Note: The default port is 7050.

4. Run the following commands:

```
cd endpoint/meta
cd get-default
invoke
```

The following screen shows the default metadata mappings:

```
    {
      "endpointJpath" : "users/sessionType",
      "metaName" : "logon_type",
      "type" : "text",
      "enabled" : true
    },
    {
      "endpointJpath" : "hostFileEntries/hosts",
      "metaName" : "dhost",
      "type" : "text",
      "enabled" : true
    },
    {
      "endpointJpath" : "securityConfigurations",
      "metaName" : "event_state",
      "type" : "text",
      "enabled" : true
    }
  ]
},
{
  "metaKeyPairsCategory" : "MACHINE_IDENTITY",
  "keyPairs" : [
    {
      "endpointJpath" : "_id",
      "metaName" : "agent.id",
      "type" : "text",
      "enabled" : true
    },
```

To disable a default metadata mapping:

Enter the same endpointJpath value and set the enabled parameter to `false`.

For example, if the endpointJpath is `Category` and enabled parameter is `true`, enter the same endpointJpath and set the enable parameter to `false`.

```
{
  "metaKeyPairsCategory" : "COMMON",
  "keyPairs" : [
    {
      "endpointJpath" : "Category",
      "metaName" : "category",
      "type" : "text",
      "enabled" : true
    },
```

Note: Do not modify the metaKeyPairsCategory in the schema; “COMMON”, “COMMON_MACHINE”, “COMMON_MACHINE_FOR_EVENTS”.

To change the metadata name or metadata type:

Enter the same endpointJpath value and specify values for the metaName and type.

Note: The metaName must exist in the table-map.xml of the Log Decoder, index-concentrator.xml or index-concentrator-custom.xml file of the Concentrator, for the metaName to appear on the Investigate view.

Add or Modify Metadata Mappings

To add or modify the metadata mappings, run the `set-custom` API. The metaKeyPairs configuration provided in the JSON file should match the JSON schema of the default configuration received through the `get-default` API.

1. On the NW server, run the `nw-shell` command from the command line.
2. Run the `login` command and enter the credentials.
3. Connect to the Endpoint Server using the following commands:

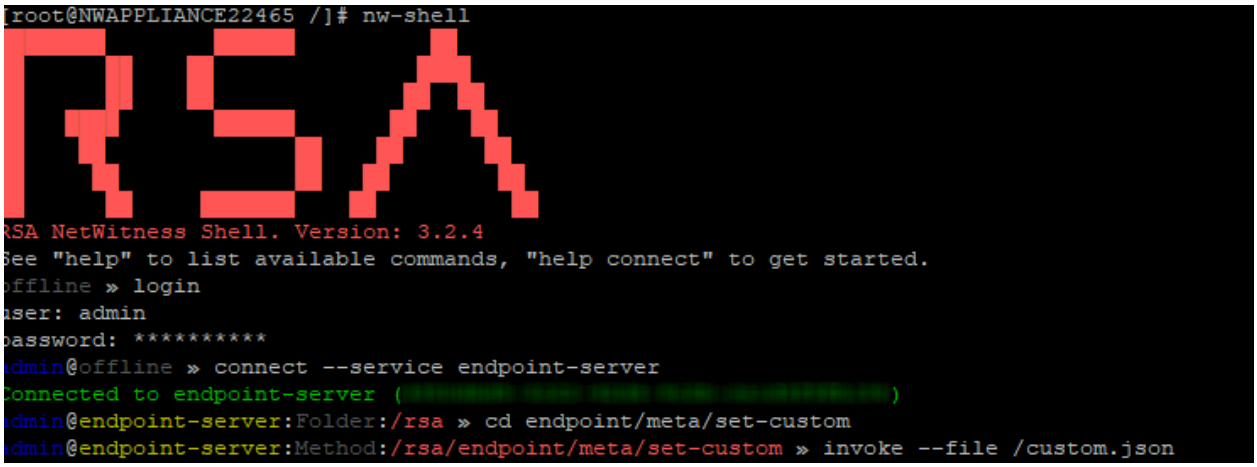
   ```
   connect --service endpoint-server
   ```

   Note: The default port number is 7050.

4. Run the following commands:

   ```
   cd endpoint/meta
   cd set-custom
   invoke --file <json file>
   ```

You can add new metaKeys by adding entries to the file that will be uploaded using the `set-custom` API. The following example shows how to add a new metadata mapping:



View the Custom Metadata Mappings

To view the custom metadata mappings, run the `get-custom` API, and then invoke commands.

Note: The `get-custom` API will return values only if the metadata mappings are modified using the `set-custom` API.

Endpoint Sources

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The Endpoint agents deployed in your environment may be large in number and geographically distributed. To efficiently manage and update configurations automatically, agents can be organized into smaller subsets called **Groups**.

Groups

Groups can be created based on Machine OU, IP address (IPv4 and IPv6), host names, operating system type, and operating system description. You can create groups based on your requirements. For example, you can group all agents running on Windows 2016 Server and IP ranging from 10.40.10.1 to 10.40.10.200. For more information on creating groups, see Creating Groups and Policies.

Note: All agents that are not part of any group use the default policy settings.

Policies

To manage the behavior of agents in a group, you can apply a set of rules called **Policies**. The RSA NetWitness supports three types of policies for endpoints: **Agent Endpoint**, **Agent File Logs**, and **Agent Windows Logs** policies. The following default policies are available on installation.

Note: RSA recommends that you review these default policies before deploying agents.

- Default Agent Endpoint (EDR) Policy
- Default Windows Log Policy
- Default File Log Policy

You can either assign the default policies to a group, modify the default policy, or create custom policies based on your organizational requirements.

Note: You cannot edit the default policy for Windows Logs nor for File Logs.

You can do the following through a policy:

- Define the agent mode - Insights or Advanced
- Configure scan schedule and settings
- Configure automatic file download
- Configure endpoint settings, such as which Endpoint server the agents should communicate, port details, and beacon intervals
- Configure response actions such as blocking
- Configure Windows and File Log collection

For example, you can create a policy to schedule scan and enable blocking. For more information on creating policies, see Create an EDR Policy, Create a File Log Policy, or Create a Windows Log Policy.

Group Ranking

When a group is created, a rank is associated with every group based on the creation order. If an agent belongs to multiple groups, to handle conflicting configurations, you can reorder the groups to change the ranking, and the policy associated with the highest ranked group takes precedence.

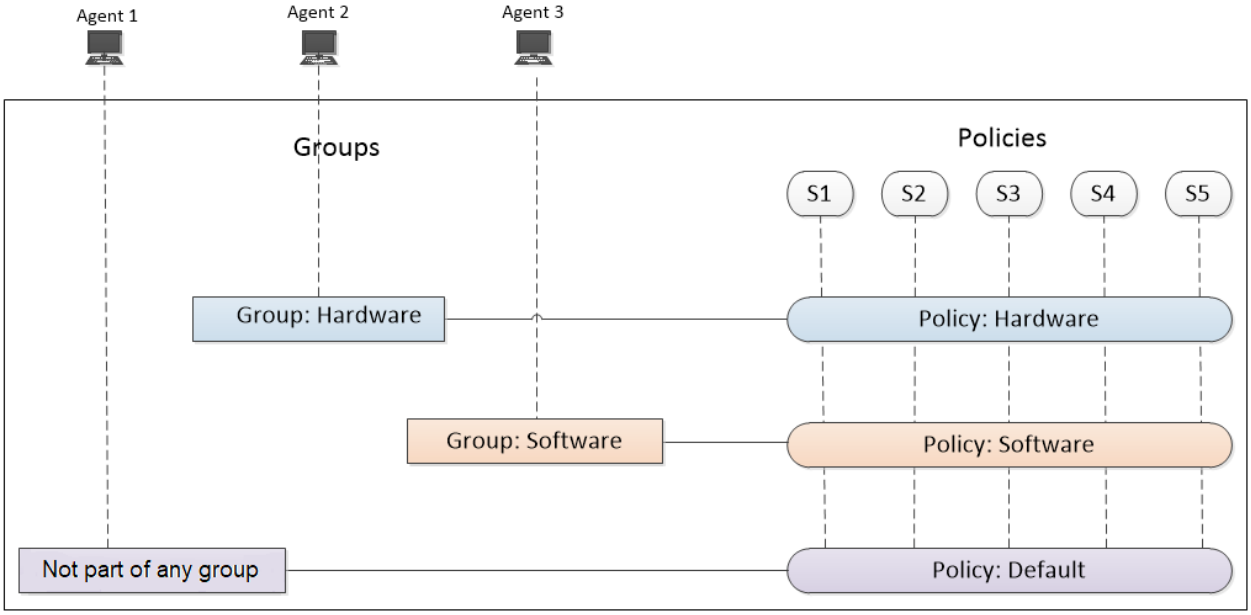
Example 1

A **Server** group contains 100 hosts with a default Agent Endpoint policy. Amongst these, if 20 hosts require further investigations, analyst can:

1. Create a temporary group with a static list of these 20 hosts.
2. Create or apply any policy to this group that will not impact any other hosts.
3. Edit the ranking for the new group, moving to the top of the Ranking list (making sure it is above the existing Server group).
4. After investigation is done, delete this group. The hosts are revalidated and assigned to the appropriate group based on the ranking.

Example 2

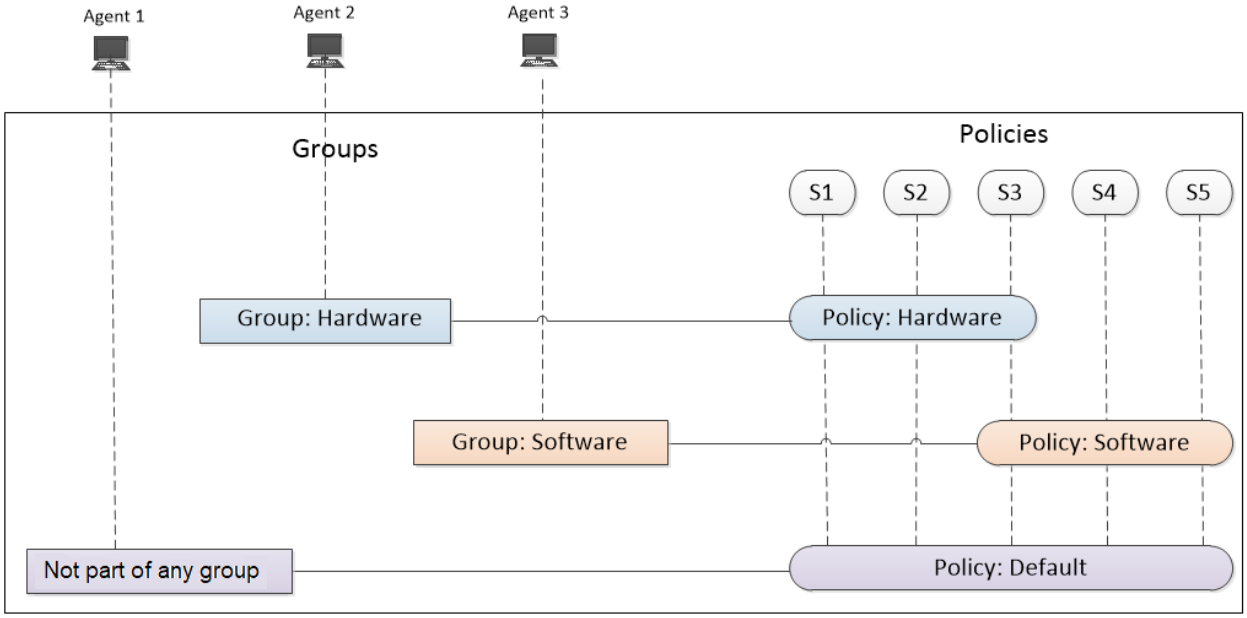
Case 1:



* S1, S2, S3, S4, and S5 represents policy settings, such as run scheduled scan, agent mode, scan settings, response actions, and so on.

Each agent is a part of a unique group that is associated with a policy, where each policy has all settings S1, S2, S3, S4, and S5 defined. For example, Agent 2 is a part of the group Hardware, where all settings in the policy Hardware are applicable.

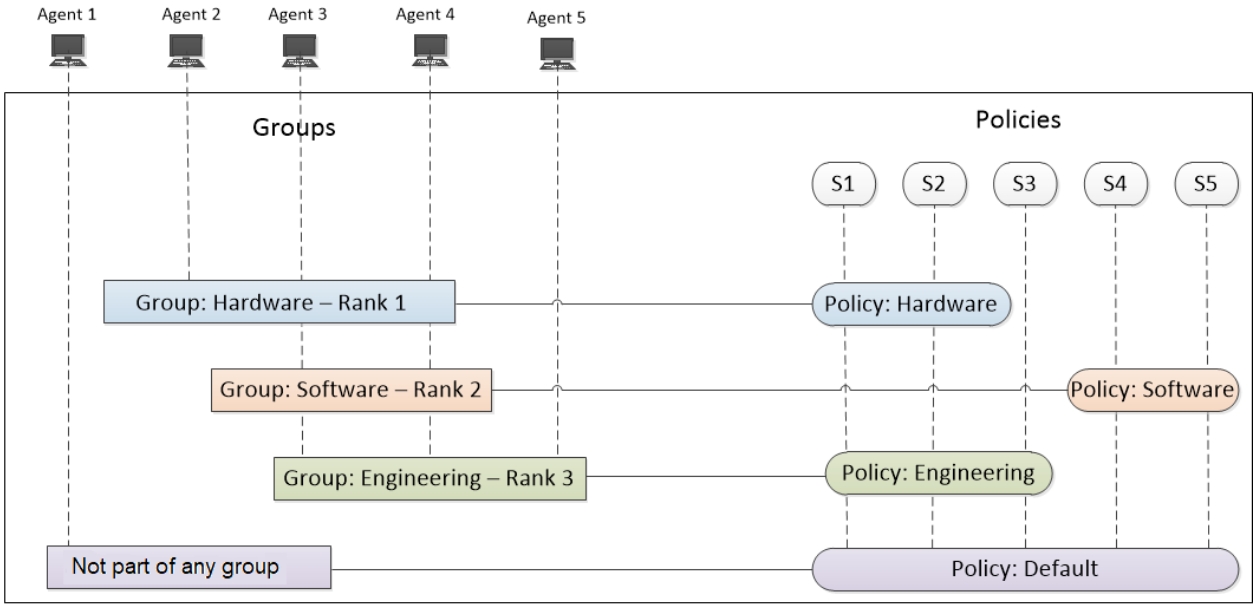
Case 2:



* S1, S2, S3, S4, and S5 represents policy settings, such as run scheduled scan, agent mode, scan settings, response actions, and so on.

In the policy Hardware, S4 and S5 settings are not defined, and hence the agent 2 inherits settings S4 and S5 from the default policy.

Case 3:



* S1, S2, S3, S4, and S5 represents policy settings, such as run scheduled scan, agent mode, scan settings, response actions, and so on.

- Agent 2 is a part of the highest ranked group Hardware, and with the policy Hardware. The agent 2 inherits settings S3, S4, and S5 from the default policy as they are not defined in the policy Hardware.

- Agent 3 is a part of Hardware, Software, and Engineering groups, and with the policy Software. The agent considers the settings S4 and S5 from the policy Software, and the remaining undefined settings are inherited as follows:
  - S1 and S2 from the policy Hardware, which is associated with the highest ranked group.
  - S3 from the policy Engineering, which is the next ranked group.
- Agent 4 is a part of Hardware, Software, and Engineering groups, and with the policy Engineering.
  - Though settings S1 and S2 are defined in policy Engineering, the agent 4 considers the settings S1 and S2 from the policy Hardware as it is associated with the highest ranked group.
  - S4 and S5 from the policy Software, which is the associated with the next highest ranked group.
  - S3 from the policy Engineering.

The following are some of the key points:

- If an agent is not assigned to any group, default policies are applied.
- A policy can be assigned to multiple groups. However, a group can only have one policy of each type (Agent Endpoint and Agent Windows Logs).
- An agent can belong to multiple groups. The policy is derived based on the ranking of the group as shown in the above example (case 3).
- If all settings are defined in a single policy, and it is the highest ranked policy for an agent, no policy settings from other ranked groups are inherited (case 1).
- If there are any undefined settings in the policy, the settings from the default policy is considered as shown in the example above (case 2 and 3).
- If an agent falls into more than one group, its complete set of policy attributes is determined as follows:
  - It takes all settings from the highest ranked policy that applies.
  - Any settings that are not set in the highest ranked policy are taken from the next highest ranked policy that applies.
  - If there are still unset attributes, they are taken from the default policy.
  - If there are any conflicts, the higher ranked policy wins.

Example 3

Assume the following:

- Agent A belongs to below two groups, **Production Servers** and **All Windows Hosts**.
- The Production Servers group has the **Schedule scan set and no blocking** policy assigned, and it has the following settings:
  - Schedule Scan : Enabled
  - Effective Date: 2019-03-08

  - Start Time: 09:00
  - Scan Frequency: Every 1 week
  - CPU Maximum: 45 %
  - Virtual Machine Maximum: 20 %
  - Blocking: Disabled
- The All Windows Hosts group has the **EDR for All Windows** policy applied, which has the following settings:
  - Scan Master Boot Record: Disabled
  - Blocking: Enabled
- The **Production Servers** group is ranked higher than the **All Windows Hosts** group for EDR policies. Keep in mind that ranking only applies to policies of the same source type: that is, all EDR policies are ranked, and all Windows Logs policies are ranked separately.

Agent A gets its final policy configuration as per the ranking of the groups (and associated policies) to which it belongs:

- The agent uses the schedule set in the **Schedule scan set and no blocking** policy.
- Scan Master Boot Record is disabled, because that is set in the **EDR for All Windows** policy.
- Blocking is disabled: since there is a conflict, the value in the higher ranked policy is used.
- All other attributes are set based on values in the Default EDR policy.
- Note that if you wanted Blocking to be enabled, you could change the group ranking so that All Windows Hosts is higher than Production Servers: in this case, Production Servers would win the conflict, and Blocking would be enabled for Agent A.

Default Agent Endpoint (EDR) Policy

When an agent is installed, it operates in an Insights mode until a policy is assigned. The following are the default EDR policy settings:

Settings	Fields	Default Value
Scan Schedule	Run Scheduled Scan	Disabled
	Effective Date	Current date
	Scan Frequency	Every week
	Start Time	09:00 (this is 9 AM)
	CPU Maximum	25%
	Virtual Machine Maximum	10%

Settings	Fields	Default Value
Agent Mode	Monitoring mode	Advanced
Scan Settings	Scan Master Boot Record	Disabled
	Auto Scan New Systems When Added	Disabled
File Download Settings	Automatic File Download	Enabled
	Automatic Memory DLL Download	Enabled
	Signature	Exclude All Signed
	File Size Limit	1 MB
Response Action Settings	Blocking	Disabled
	Network Isolation	Disabled
Endpoint Server Settings	Endpoint Server	The agent considers the default Endpoint Server that is configured during packager generation.
	Server Alias (Optional)	
	HTTPS Port	443
	HTTPS Beacon Interval	15 Minutes
	UDP Port	444
	UDP Beacon Interval	30 Seconds

Default Windows Log Policy

The following are the default Windows Log policy settings:

Settings	Fields	Default Value
Windows Log Settings	Status	Disabled
	Protocol	TLS
	Send Test Log	Disabled

Default File Log Policy

The following are the default File Log policy settings:

Settings	Fields	Default Value
File Log Settings	Status	Disabled
	Protocol	TLS
	Send Test Log	Disabled

Creating Groups and Policies

Note: The information in this topic applies to NetWitness Version 11.3 and later.

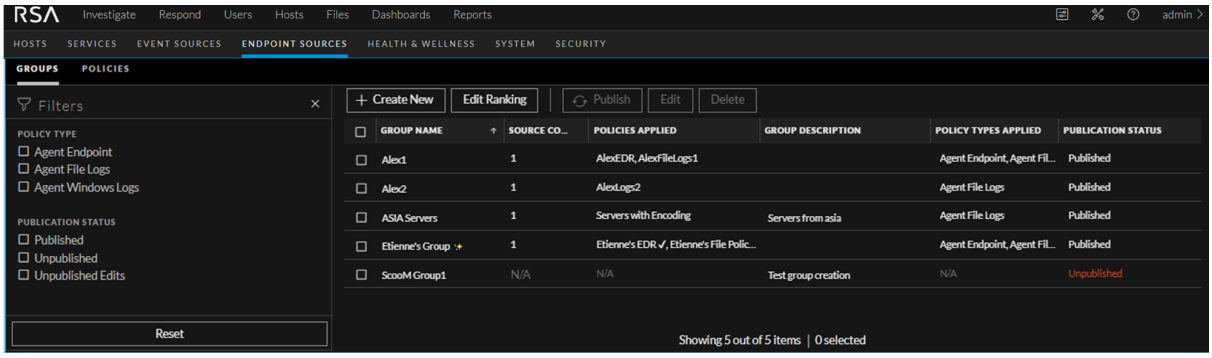
The following sections provide instructions on how to create groups and policies.

- Groups: Create a Group
- Policies:
  - Construct a Policy
  - Create an EDR Policy
  - Create a Windows Log Policy
  - Create a File Log Policy
  - Replace Windows SFTP Agents

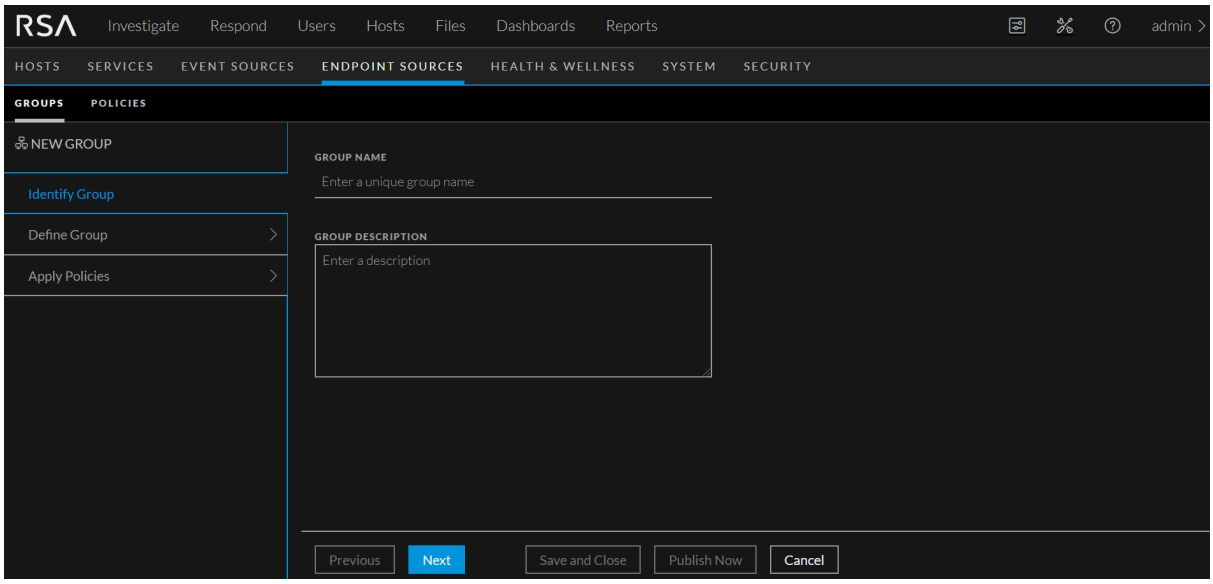
Create a Group

To create a group:

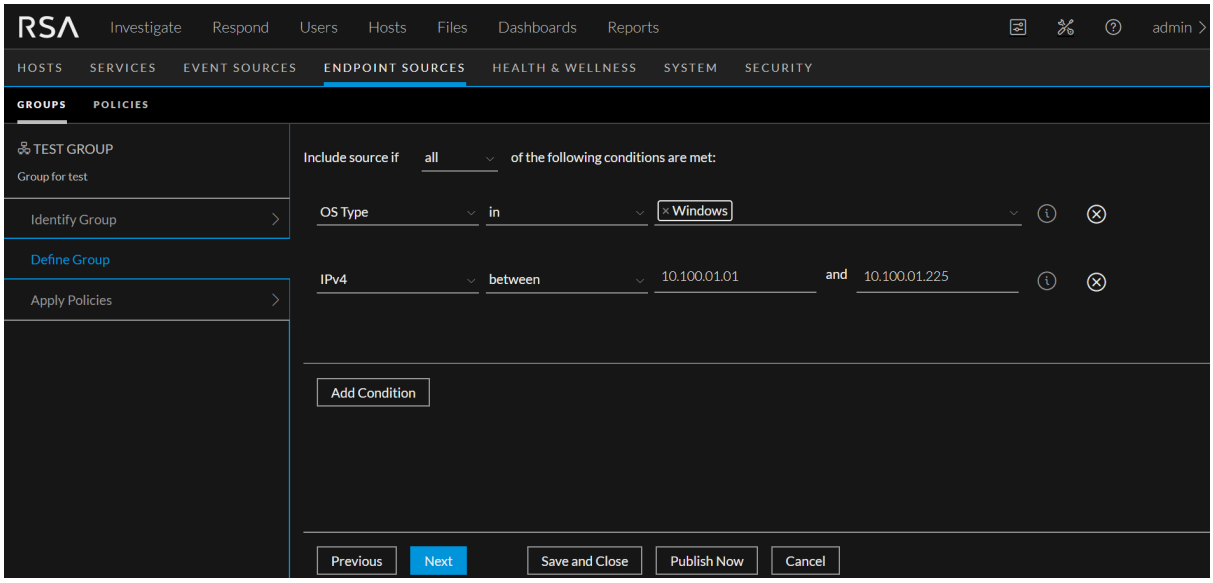
1. Go to **(Admin) > Endpoint Sources** view.
2. In the left panel, select the **Groups** tab.



3. In the toolbar, click **Create New**.
4. In the **New Group** panel, enter a group name and group description, and click **Next**.



5. Specify the logical statements that define the condition for an agent to be included in the group. Each logical statement consists of: parameter, operator, and values to match.



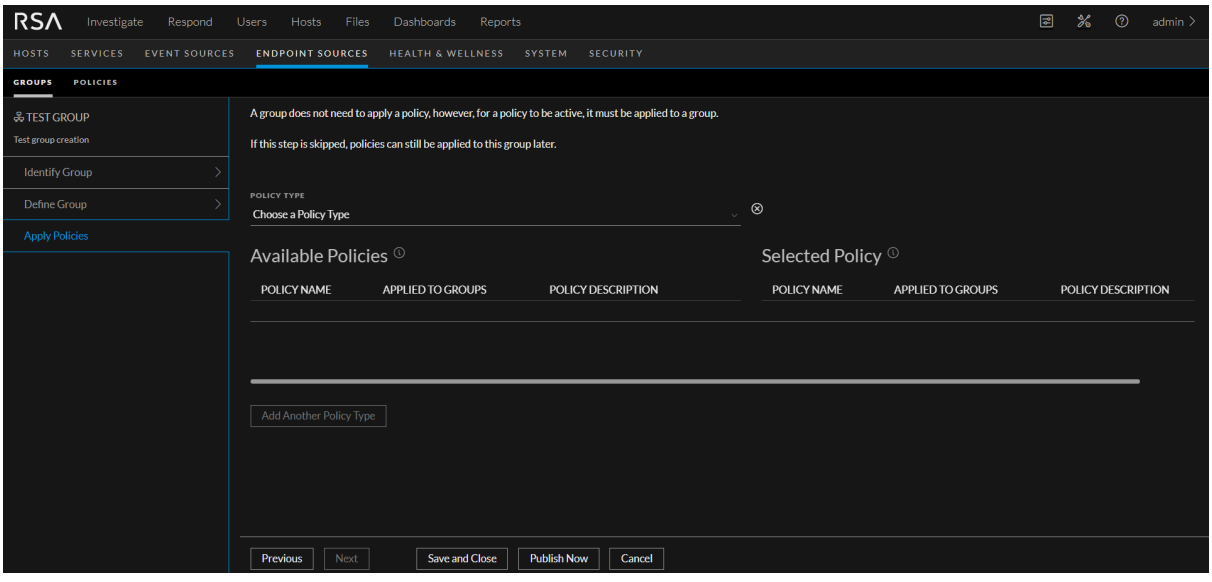
- In the **Include source if ___ of the conditions are met** field, select either **all** or **any**.
- For each logical statement, select the required options:

Item	Description
Parameter	The parameter can be OS Type, OS Description, Host Name, IPv4, IPv6, Machine OU or Tag. • OS Type, OS Description, Host Name: The value you enter should reference hardware or virtual machines that are running endpoint agents. • IPv4 or IPv6: Enter valid IP addresses as either ranges or as a set of IP addresses to include or exclude. • Machine OU: Enter a valid Machine OU(Organizational Unit) name. • Tag: Enter a valid tag that already exists. **Note:** If you do not want to include certain IP addresses, use the **Not in** operator, and enter the IP addresses separated by a space or a comma.
Operator	The choice of values is dependent upon the parameter you chose. For example, if your parameter is OS Type, the only operator available is **in**.
Value or values to match	The value or values to match. For the OS Type parameter, you can choose one or more values from the drop-down list. For all other parameters, you can enter free-form text. **Note:** Although you can enter any text for values, the system validates your entries when you attempt to proceed to another screen, and will not allow you to proceed until values are valid.

6. Continue adding conditions until you have completely specified the new group. After you have added all conditions, click **Next** to proceed.
7. (Optional) Click **Apply Policies** and select the source type from the drop-down list.

   Policies with the selected source type are displayed below **Available Policies**.

   Select a policy by clicking . Skip this step if you want to apply a policy to the group at a later time.



Note: You can attach only one policy per source type to a group. That is, you cannot attach more than one Agent Endpoint policy to a single group, nor more than one Agent Windows Logs policy.

For more information on creating policies, see Create an EDR Policy, Create a Windows Log Policy, or Create a File Log Policy.

8. Do one of the following:
   - Click **Save and Close** to save the settings and return to the Groups view. The publication status is displayed as **Unpublished** in the Groups view.

     Note: You can select an unpublished group and click **Publish** to publish a group.
   - Click **Publish Now** to publish the group.

Construct a Policy

When you create a policy, you should keep in mind the way groups and Agents can inherit values from other policies. The simplest way to construct a policy is to set values for all of the available settings for that policy type. If you do this, you can assign the policy to one or more groups, and all of the agents in those groups can receive the settings defined in the policy itself.

However, you do not need to set all possible available settings within a single policy. In the Group Ranking section, there are examples that show where values can come from, based on the ranking order.

Note: Remember that a group can have no more than one of each policy type assigned to it: it can be assigned 0 or 1 Agent Endpoint, Agent File Logs, and Agent Windows Logs policies.

If an agent is only a member of one group, that agent's settings are as follows:

- If a value is set in the policy assigned to the group, that value is used.
- If a value is not set in the policy assigned to the group, the value set in the default policy is used.

For agents that are members of multiple groups, it is a bit more complicated. For these agents, the settings are evaluated from highest priority member group (as set on the Edit Rankings page) to the lowest priority member group. If a parameter is set in a higher group, it is not overwritten, event if the same parameter is also set in a lower group. For example, assume an Agent is part of three groups:

- If a value is set in the highest ranked policy, the agent uses that value.
- If a value is not set in the highest ranked policy, but is set in the second-highest-ranked policy, the agent uses the value from the second-highest-ranked policy.
- If a value is not set in either the first- or second-highest ranked policy, but is set in the third-highest-ranked policy, the agent uses the value from the third-highest-ranked policy.
- If a value is set in the highest-ranked policy and either or both of the other policies, the value is taken from the highest-ranked policy.
- If a value is not set in any of the three policies assigned to the agent, the agent uses the value set in the default policy.

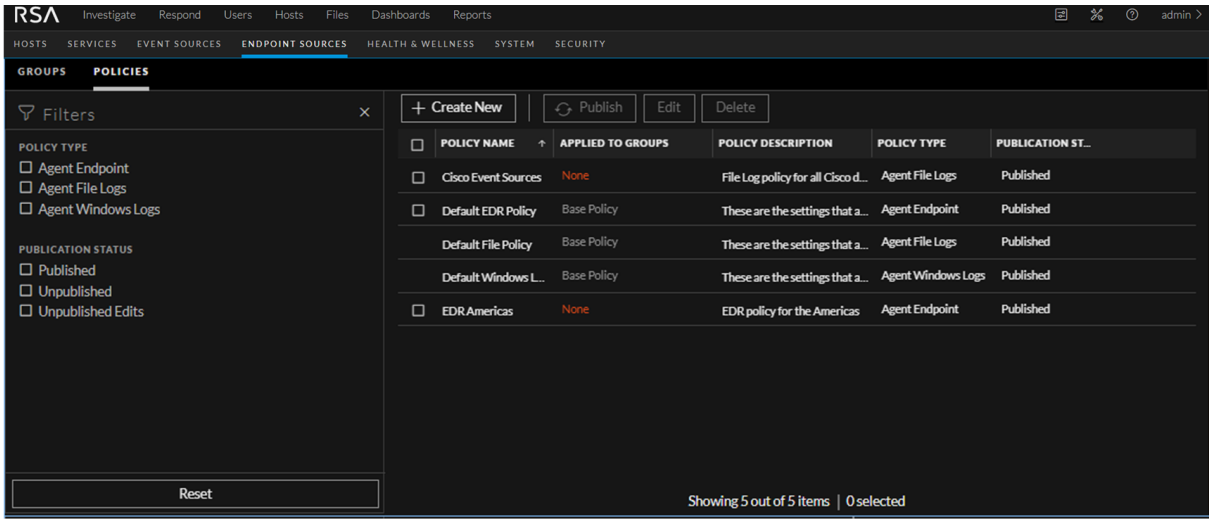
Create an EDR Policy

While creating a policy, note the following:

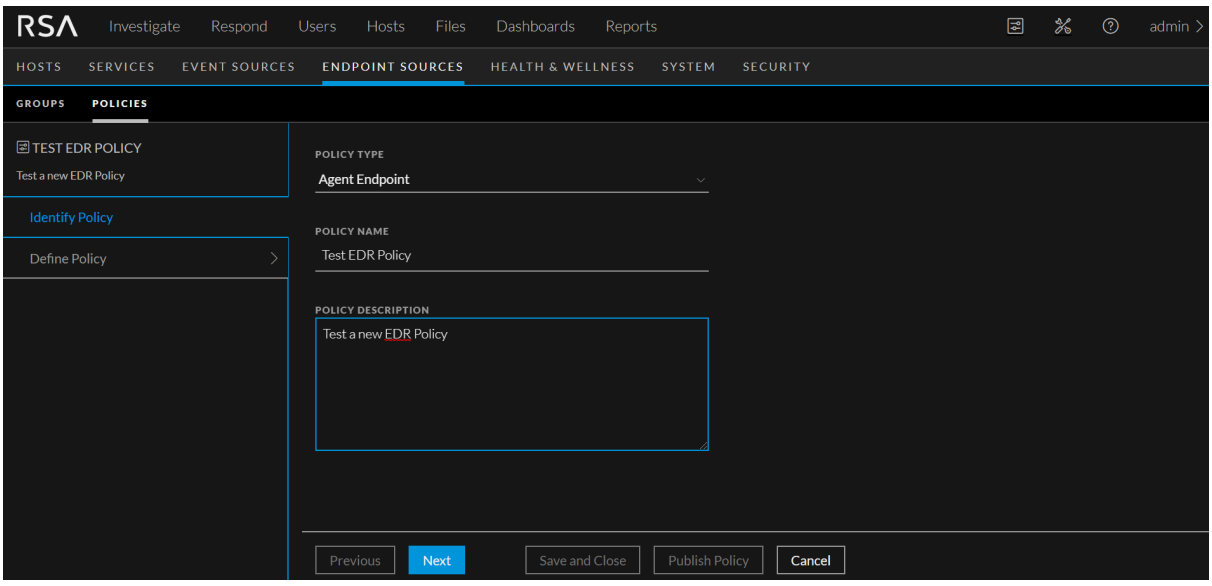
- Whenever you choose a setting, it is added to the **Selected Settings** panel.
- To clear any of your selected settings, click  to remove that setting.
- At any point in the wizard, you can choose **Save and Close**, so that you can return to complete the policy at a later time.

To create an EDR policy:

1. Go to  **(Admin) > Endpoint Sources**.
2. Click **Policies**. The available policies are displayed.



3. Click **Create New** to add a new policy.
4. In the New Policy panel, do the following:

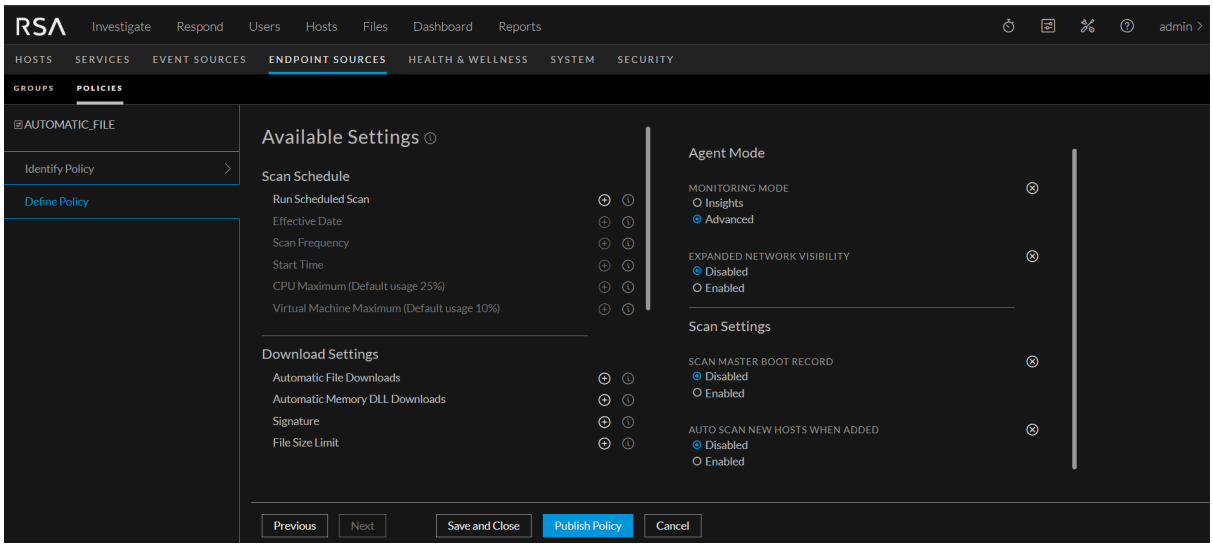


- Select **Agent Endpoint** as the source type from the drop-down list.
- Enter the policy name.
- Enter a description for the policy.

5. Click **Next**.

6. Click  to select a setting from list of **Available Settings**. After you click, the specific setting is moved under the **Selected Settings** panel. You need to enter the required values for the selected settings. For details, see Define Policy Panel for Agent Endpoint Policy.

Note: You do not need to set all possible available settings within a single policy. The complete list of settings for an agent is derived from one or more groups to which that agent belongs. This is described in more detail in Construct a Policy.



- In the Scan Schedule category, enable **Run Scheduled Scan** to configure the scan, and set any of the available parameters based on your needs. For more details, see Define Policy Panel for Agent Endpoint Policy.
- In the **Agent Mode** category, you can set the following:
  - Select the monitoring mode of the agent - Insights or Advanced.
  - Select **Expanded Network Visibility**:

Prerequisite

Note: For Expanded Network Visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Services Security View - Aggregation Role" in the *Hosts and Services Getting Started Guide for RSA NetWitness Platform*.

  - To enable network tracking and monitoring on Windows hosts in Insights mode.

Note: Reports UDP connections for Windows 10 version 1803 in Insights mode.

- To optimize the frequency of agents sending network events for both Insights and Advanced modes.
- The endpoint agents should have the **Expanded Network Visibility** option enabled for enriched network events from NetWitness Platform packet deployment with endpoint data. In the **Events** View > **Hosts** tab, the endpoint data related to the network event selected will be automatically displayed if it exists. For more information, see "Host Information" in the *Investigate User Guide*.

Note: For network tracking in Insights mode, verify that the Windows Management Instrumentation (WMI) service is enabled.

- In the Scan Settings category, you can enable either or both of the following actions:
  - Enable **Scan Master Boot Record** to include Master Boot Record (MBR) details in scheduled scans.
  - Enable **Auto Scan New Systems When Added** to automatically queue a scan for any host that does not have any snapshot data
- In the Download Settings category, you can set the following:
  - Enable **Automatic File Downloads** to automatically download files based on the Signature and file size. By default this option is enabled.
  - Enable **Automatic Memory DLL Downloads** to automatically download all the memory DLLs that are detected during a scan. By default this option is enabled.
  - Select the **Signature** type to limit the download of files based on the signature (not available for Linux systems).
  - Specify the **File Size Limit** to limit the download of files based on the file size. The file size limit should be between 1 KB - 10 MB.
- In the Response Action Settings category you can enable or disable the following actions:
  - Enable **Blocking** to prevent the execution of a malicious file on any host.
  - Enable **Network Isolation** to provide an option to isolate a compromised host during investigation.
- In the Endpoint Server Settings, configure your server:
  - Add the Endpoint server that the agent will communicate from the drop-down list.

    Note: If you do not select an Endpoint Server, the agent uses the default Endpoint Server that is configured during packager generation.
  - (Optional) Enter an alternative hostname or IP address.
  - Enter the HTTPS port used for communication.
  - Specify the HTTPS beacon interval.
  - Enter the UDP port used for communication.
  - Specify the UDP beacon interval.

- **Advanced Setting** - For NetWitness Support staff only.

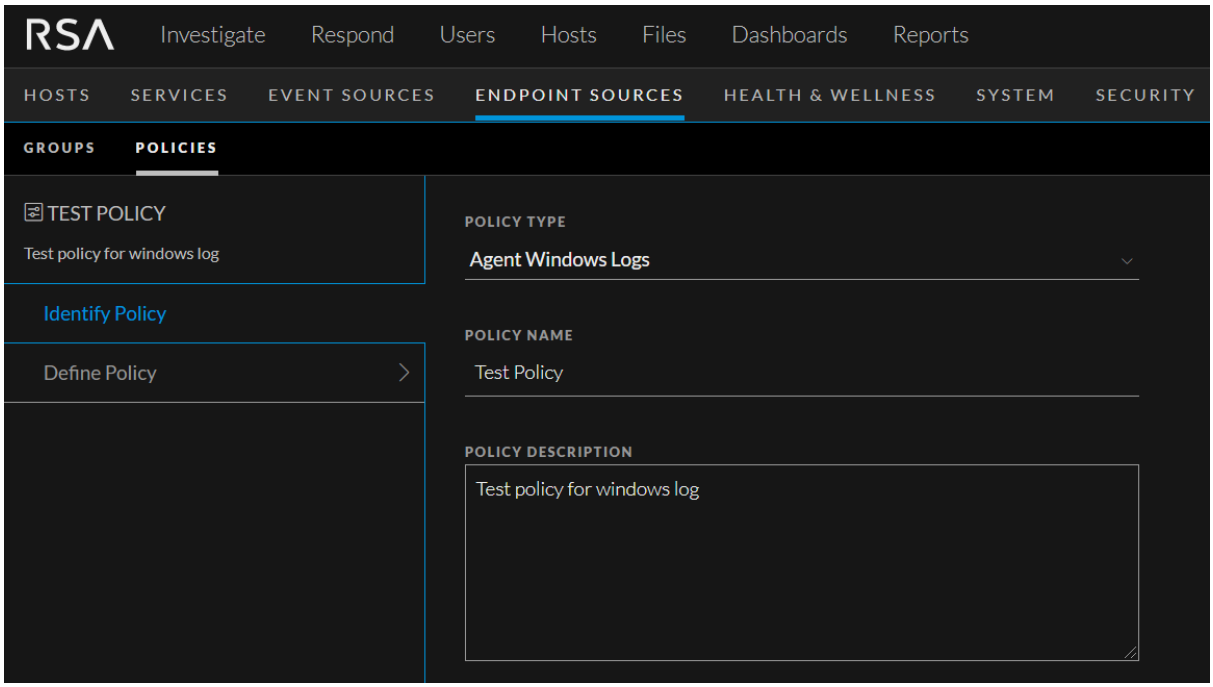
  IMPORTANT: It is strongly recommended not to use the Advanced Configuration unless advised to do so by NetWitness.

7. Do one of the following:
   - Click **Save and Close** to save the settings and return to the Policies view. The policy will be listed under the **Unpublished** category.
   - Click **Publish Policy** to publish the policy.

Create a Windows Log Policy

To create a Windows Log policy:

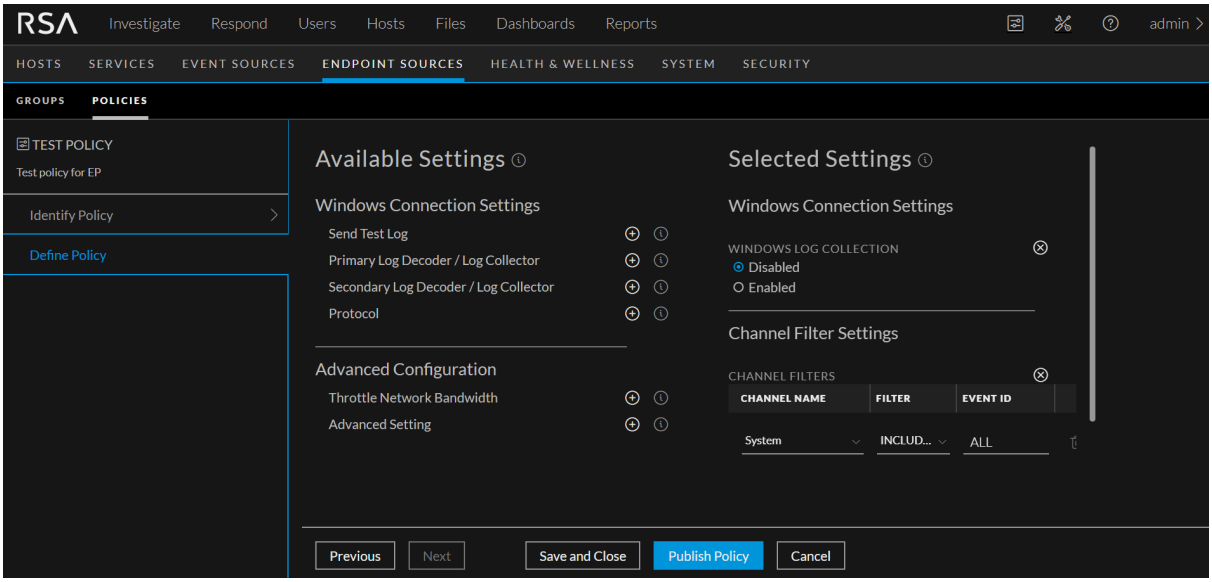
1. Go to **(Admin) > Endpoint Sources**.
2. Click **Policies**. The available policies are displayed.
3. Click **Create New** to add a new policy.
4. In the New Policy panel, do the following:



   - Select **Agent Windows Logs** as the source type from the drop-down list.
   - Enter the policy name.
   - Enter a description for the policy.
5. Click **Next**.

6. Click  to select a setting from list of **Available Settings**. After you click a setting it is moved under the **Selected Settings** panel. You need to enter the required values for the selected settings.

Note: You do not need to set all possible available settings within a single policy. The complete list of settings for an agent is derived from one or more groups to which that agent belongs. This is described in more detail in Construct a Policy.



- Select **Windows Log Collection** to enable Windows Log collection. By default, this option is disabled.
- Enable **Send Test Log** to send a test log. By default, this option is disabled.
- Select **Primary Log Decoder / Log collector** to forward logs from the drop-down list.
- (Optional) Select **Secondary Log Decoder / Log collector** to forward logs from the drop-down list.

  Note: When the Endpoint Agent is configured to use the UDP protocol and the Primary Log Decoder/ Remote Log Collector is not reachable, the secondary Log Decoder or Log Collector is not functional. The logs are not forwarded to the secondary Log Decoder or Log Collector when the primary is down, thus resulting in the event loss.

- Select **Protocol** from the drop-down list. The available options are UDP, TCP, and TLS. By default, the protocol is TLS.
- Add **Channel Filters** and select the channels from which the logs are collected from the drop-down list. You can add or remove a channel filter and specify individual Event IDs.

- **Advanced Configuration**
  - **Throttle Network Bandwidth:** use this setting to limit network bandwidth that the Agent uses to connect to NetWitness. This setting is disabled by default: click **Enabled** to turn it on, and then enter a value in kilobits per second.
    - If not set, Agent does not do any network throttling.
    - If set to a positive value ***x***, agent limits network bandwidth to ***x*** kbps.
  - **Advanced Setting** - For NetWitness Support staff only.

    IMPORTANT: It is strongly recommended not to use the Advanced Configuration unless advised to do so by NetWitness.

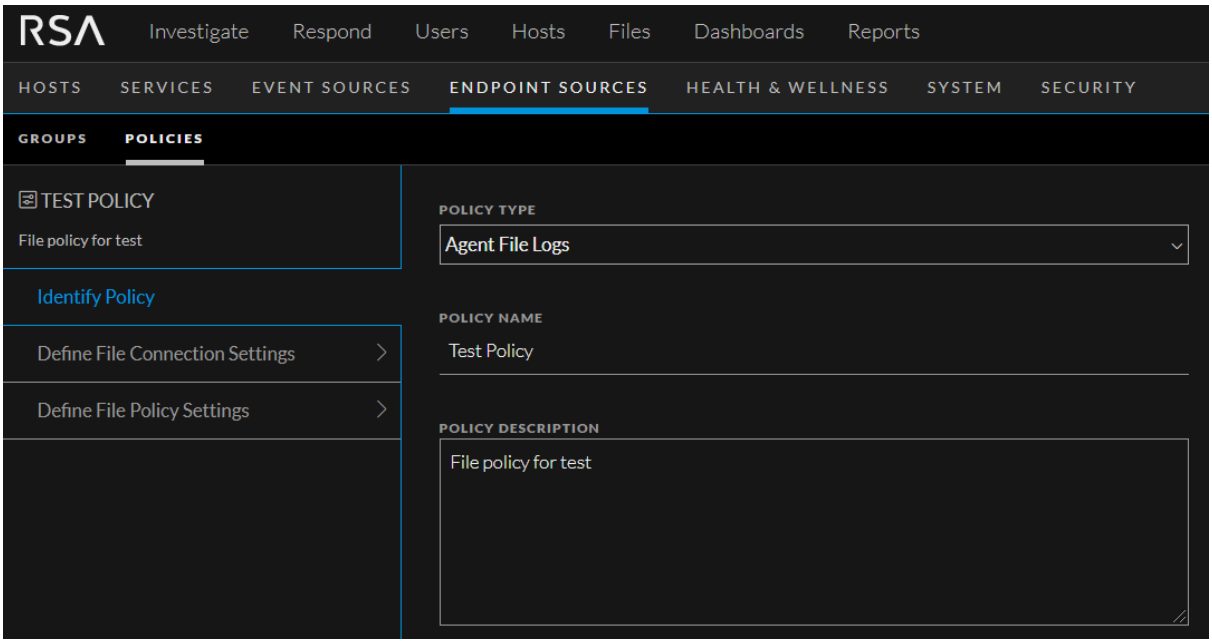
7. Do one of the following:
   - Click **Save and Close** to save the settings and return to the Policies view. The policy will be listed under the **Unpublished** category.
   - Click **Publish Policy** to publish the policy.

Create a File Log Policy

Note: You cannot create File Log policies while the system is in mixed mode. Until all Endpoint servers are updated to 11.4, the Agent File Logs options on the policy create, assign policy, and edit ranking pages are disabled.

To create a File Log policy:

1. Go to  **(Admin) > Endpoint Sources**.
2. Click **Policies**. The available policies are displayed.
3. Click **Create New** to add a new policy.
4. In the New Policy panel, do the following:

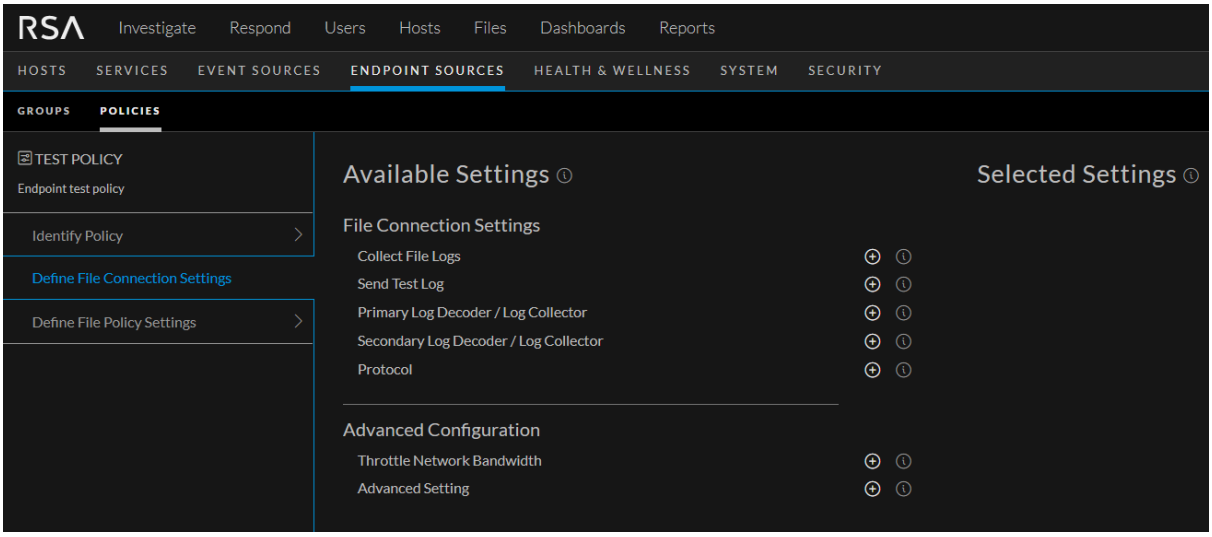


- Select **Agent File Logs** as the source type from the drop-down list.
- Enter the policy name.
- Enter a description for the policy.

5. Click **Next**.

6. Click  to select a setting from list of **Available Settings**. After you click a setting, the specific setting is moved under the **Selected Settings** panel. You need to enter the required values for the selected settings.

Note: You do not need to set all possible available settings within a single policy. The complete list of settings for an agent is derived from one or more groups to which that agent belongs. This is described in more detail in Construct a Policy.



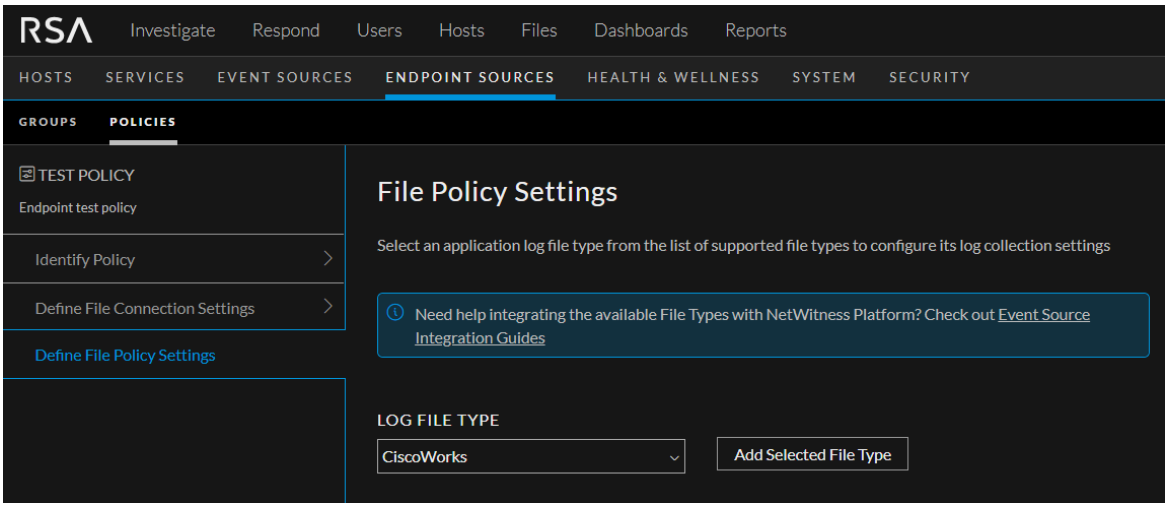
- Enable **Collect File Logs** to collect file logs on endpoints assigned to this policy. By default, this option is disabled.
- Enable **Send Test Log** to send a test log. By default, this option is disabled.
- Select **Primary Log Decoder / Log collector** to forward file logs from the drop-down list.
- (Optional) Select **Secondary Log Decoder / Log collector** to forward file logs from the drop-down list.

  Note: When the Endpoint Agent is configured to use the UDP protocol and the Primary Log Decoder/Remote Log Collector is not reachable, the secondary Log Decoder or Log Collector is not functional. The logs are not forwarded to the secondary Log Decoder or Log Collector when the primary is down, thus resulting in the event loss.

- Select **Protocol** from the drop-down list. The available options are UDP, TCP, and TLS. By default, the protocol is TCP.
- **Advanced Configuration**
  - **Throttle Network Bandwidth:** use this setting to limit network bandwidth that the Agent uses to connect to NetWitness. This setting is disabled by default: click **Enabled** to turn it on, and then enter a value in kilobits per second.
    - If not set, Agent does not do any network throttling.
    - If set to a positive value ***x***, agent limits network bandwidth to ***x*** kbps.
  - **Advanced Setting** - For NetWitness Support staff only.

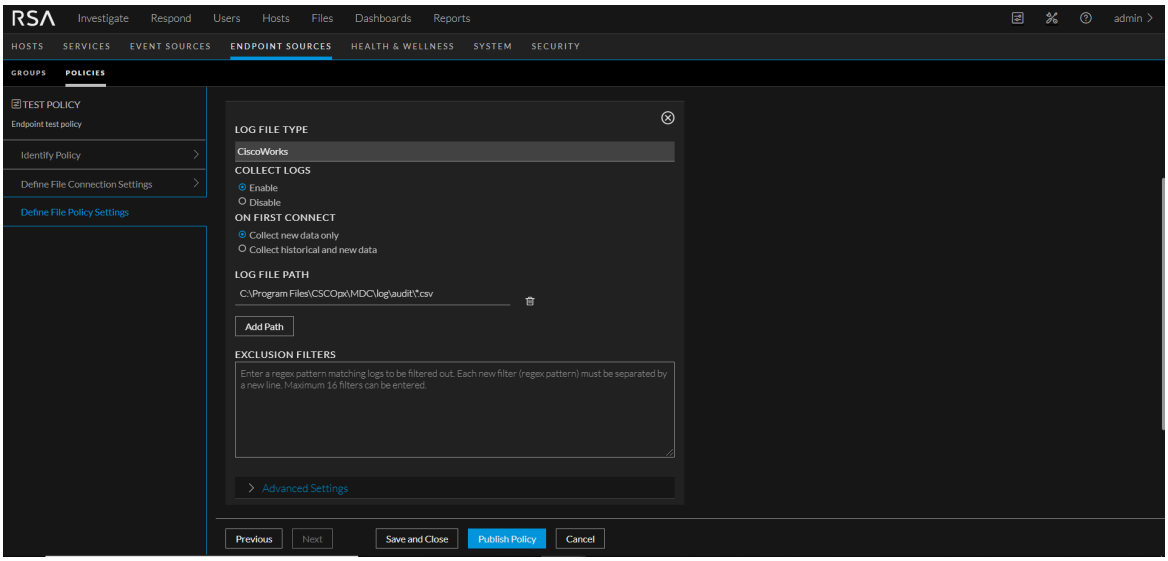
    IMPORTANT: It is strongly recommended not to use the Advanced Configuration unless advised to do so by NetWitness.

7. Click **Next**, and add the file type or types for the policy.
   a. Select a log file type from the list, then click **Add Selected File Type** to configure its log collection settings.



   For a list of currently supported event source types, see Endpoint Sources - Policies.

   b. Choose values for the available parameters.



- **Collect Logs**: Select **Enable** or **Disable**.
- **On First Connect**: Choose whether or not to collect older, historical data, or just new data. The default setting is to collect new data only.

  Note: New data is data that is collected starting from when you configure log collection for the specified event source.

- (Optional) **Log File Path**: Add a path for where the log files are stored. This is only necessary if the log files are not stored in the standard directory for the selected event source type. To add a path, click **Add Path** and enter a pathname.

  Note: This can be a Universal Naming Convention (UNC) pathname. (`\\host-name\share-name\file-path`).

  Click **Add Path** again to add another path. You can add as many paths as you like.

- **Exclusion Filters**: You can enter a newline-separated list of exclusion filters, which specify log files from which RSA should not be collecting data.

  Note: The filter needs to be entered as a valid regex string, or the system will not allow you to save it.

- Advanced Settings: note that most users do not need to set these parameters.
  - **Source Alias**: For most installations, this value is not needed. Use this to specify a unique event source name, in the case you have multiple event sources of the same type, for example two IBM WebSphere MQ event sources in the same NetWitness installation.
  - **File Encoding**: Select a type of file encoding. You can choose from a wide variety of encodings. The default value is **UTF-8/ASCII**.

For more details on the available parameters, see Panels for Log File Policy.

c. You can add more file types to the policy. After you have added all your file types, proceed to the next step.

8. Do one of the following:

- Click **Save and Close** to save the settings and return to the Policies view. The policy will be listed under the **Unpublished** category.
- Click **Publish Policy** to publish the policy.

Replace Windows SFTP Agents

Note that you might want to replace the SFTP Agent for Windows with the NetWitness Endpoint Agent for collection from file event sources. If so, perform the following procedure:

1. Using the File Collection Policy wizard (as described below), configure file collection for your event sources.
2. Verify that collection is working.
3. On each of your event sources, stop the SFTP service.

Managing Groups

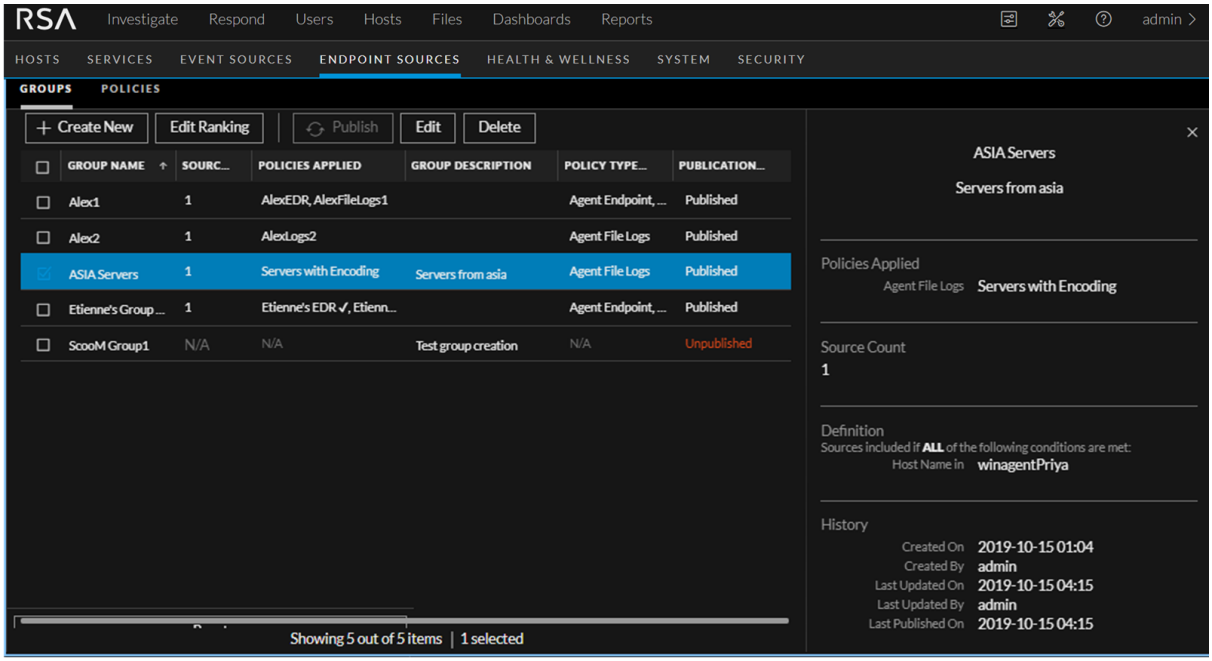
Note: The information in this topic applies to NetWitness Version 11.3 and later.

You can view group details, edit group details, filter endpoint groups, delete groups, and edit group ranking. For details on how to create groups, see Create a Group.

View Group Details

To view properties of the selected group:

1. Go to **(Admin) > Endpoint Sources**.
2. In the left panel, select the **Groups** tab. The details, such as group name, source count, policies applied, group descriptions, source type applied, and publication status are displayed. For more details on these columns, see Endpoint Sources - Groups.
3. Click the row to view the properties in the right-panel.



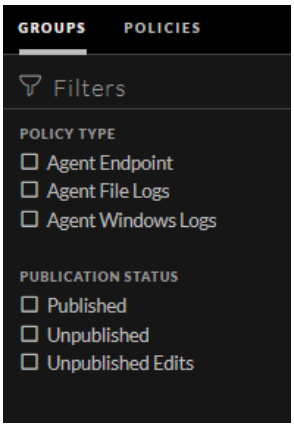
Filter Endpoint Groups

The Filters Panel allows you to filter the list of displayed groups, based on the one of the following source type:

- Agent Endpoint
- Agent Windows Logs

Additionally, you can sort based on publication status:

- Published - Groups that are published to use.
- Unpublished - Groups that are saved but not published.
- Unpublished Edits - Groups that are previously published and edited later and saved, but not published.



The Filters panel can be hidden or displayed:

- To hide, click the icon at the top-right of the panel.
- To display if hidden, click the icon in the toolbar.

Click **Reset Filters** to remove the currently applied filter criteria.

Edit a Group

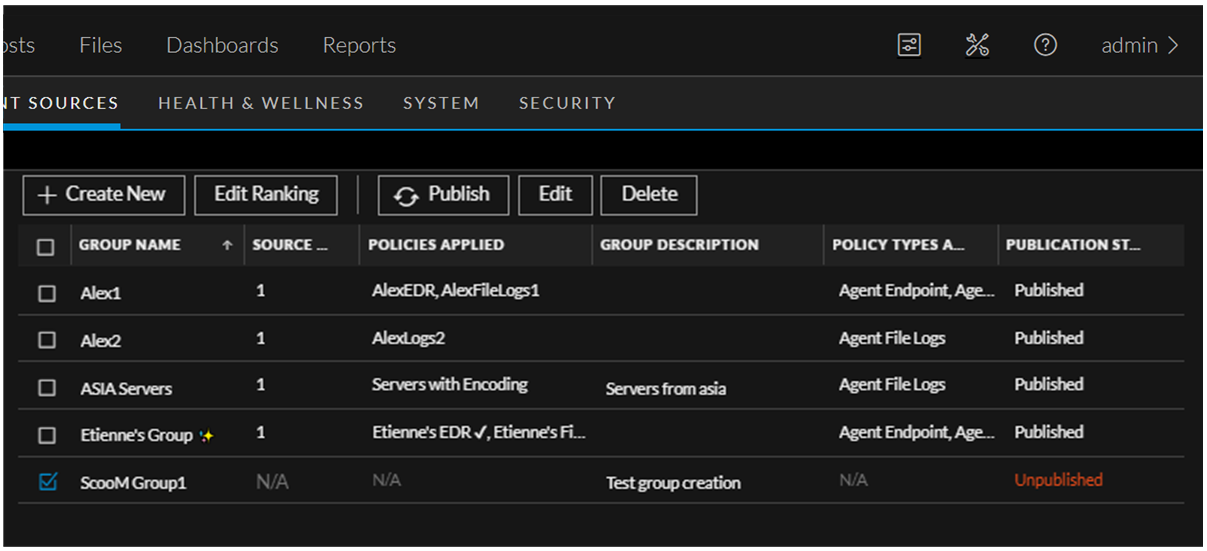
You can edit the properties of the group at any point in time. To edit properties of a group:

1. Go to **(Admin) > Endpoint Sources**.
2. Select a group and click **Edit**.
3. Edit the group details as required.
4. Do one of the following:
   - Click **Save and Close** to save the changes and return to the Groups view. The group will be listed under the **Unpublished Edits** category.
   - Click **Publish Now** to publish the changes.

Delete a Group

To delete a group:

1. Go to **(Admin) > Endpoint Sources**.
2. The **Groups** tab and available groups are displayed.



3. Select one or more groups and click **Delete**.
4. Click **Delete**. The confirmation message is displayed.
5. In the Delete Groups dialog, click **Delete Group(s)** to permanently delete the selected groups.

Managing Policies

Note: The information in this topic applies to NetWitness version 11.3 and later. Policies are not applicable on standalone agents.

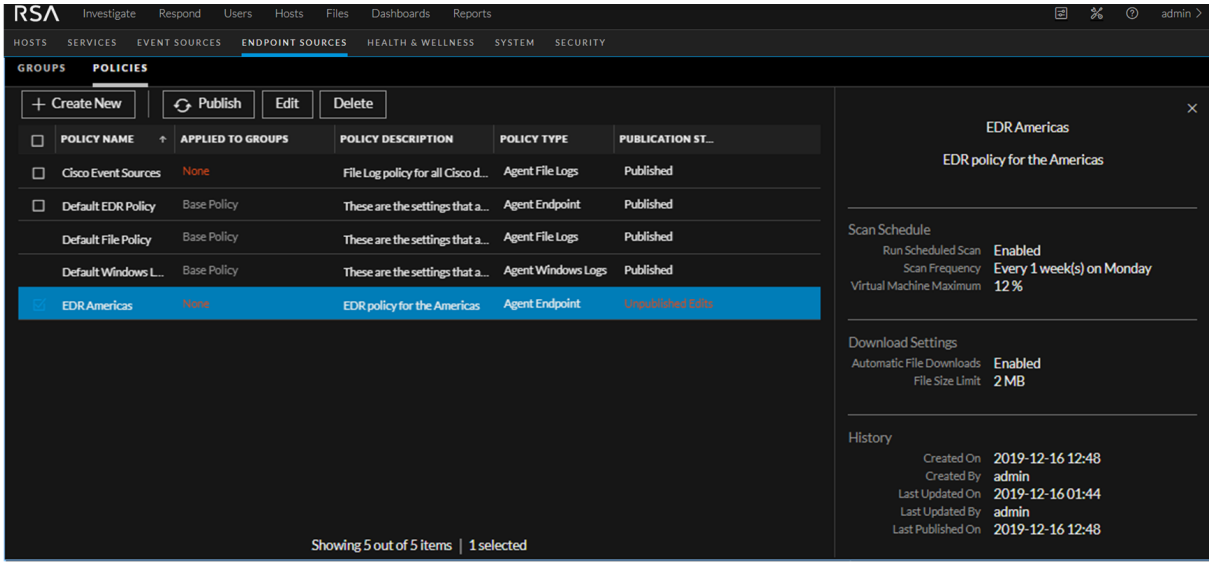
You can view, edit, filter, and delete policies, as detailed in the following sections:

- View Policy Details
- Filter Policies
- Edit a Policy
- Delete a Policy
- Conflict Resolution

View Policy Details

To view properties of the selected policy:

1. Go to **(Admin) > Endpoint Sources**.
2. In the left panel, select the **Policies** tab. The details, such as policy name, applied to groups, policy description, source type, and publication status are displayed. For more details on these columns, see Endpoint Sources - Policies.
3. Click the row to view details about selected policy in right pane.



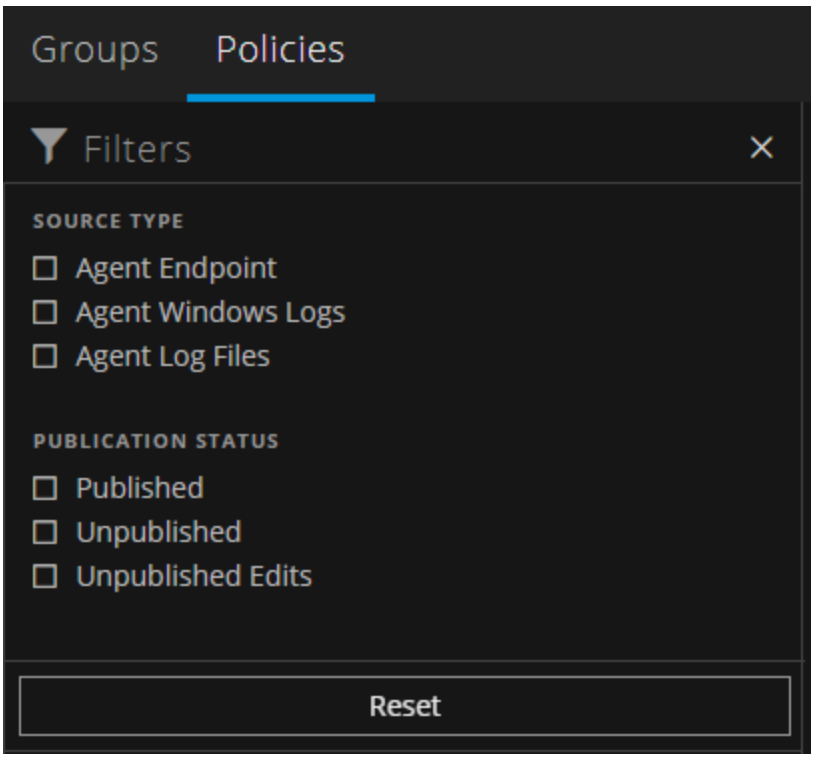
Filter Policies

The Filters Panel allows you to filter the list of displayed policies, based on the source type. You can filter on any combination of the following:

- Agent Endpoint
- Agent File Logs
- Agent Windows Logs

Additionally, you can filter based on publication status:

- Published: Policies that are published to use.
- Unpublished: Policies that are saved but not published.
- Unpublished Edits: Policies that are previously published and edited later and saved, but not published.



The Filters panel can be hidden or displayed:

- To hide, click the icon at the top-right of the panel.
- To display if hidden, click the icon in the toolbar.

Click **Reset Filters** to remove the currently applied filtering criteria.

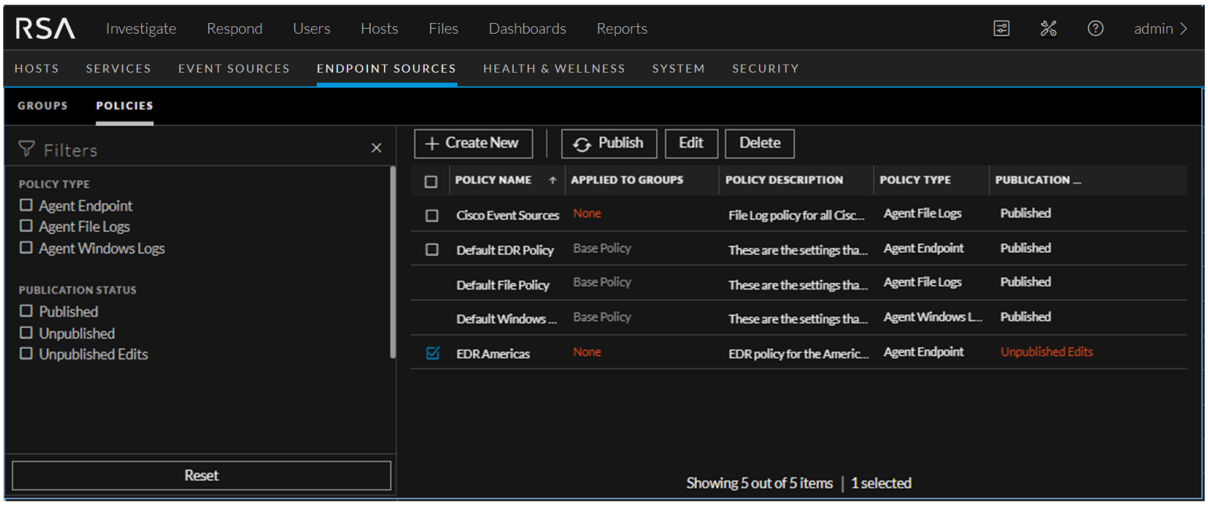
Edit a Policy

You can edit the settings of the default Agent Endpoint and custom policies. The default Agent Windows Log policy cannot be edited.

Note: For the default EDR policy, you cannot edit the source type, policy name, and policy description. However, you can edit the details in the Define Policy panel.

To edit a policy:

1. Go to **(Admin) > Endpoint Sources**, and select the **Policies** tab.
2. Select a policy and click **Edit**.

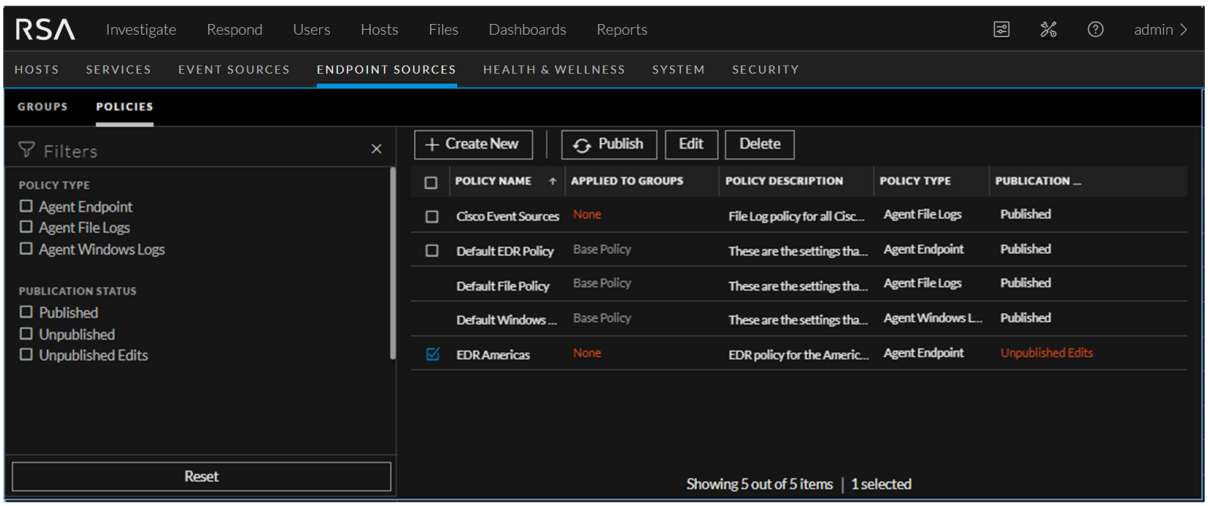


3. Edit the policy details as required.
4. Do one of the following:
   - Click **Save and Close** to save the changes and return to the Policies view. The policy will be listed under the **Unpublished Edits** category.
   - Click **Publish Policy** to publish the changes.

Delete a Policy

To delete a policy:

1. Go to **(Admin) > Endpoint Sources**.
2. Click the **Policy** tab. The available policies are displayed.



3. Select one or more policies and click **Delete**.

   The confirmation message is displayed.

4. In the Delete Policies dialog, click **Delete Policy(ies)** to permanently delete the selected policies.

Conflict Resolution

An endpoint can be in more than one group, and can thus have more than one Agent Endpoint, Agent File Logs, or Windows Logs policy applied to it. In this case, there may be conflicting settings that could be applied to the endpoint.

For example, an endpoint that is in two Groups could have two, different File Log policies applied to it. In this case, some of the settings could have conflicting values. The value that is actually applied to the endpoint is determined by the highest-ranked policy that contains a value for that setting.

For example, assume there is an endpoint that has 2 Agent File Log policies applied to it:

- LF Policy One: Log File Type is webgateway, and **File Encoding** is set to UTF-8
- LF Policy Two: Log File Type is webgateway, and **File Encoding** is set to Local Encoding

How NetWitness assumes the webgateway logs are encoded is dependent upon which policy is ranked higher:

- If Policy One is ranked higher than Policy Two, NetWitness treats the logs as having UTF-8 encoding.
- If Policy Two is ranked higher than Policy One, NetWitness treats the logs as having Local Encoding.

For an example using EDR policies, see Simulation Examples, which shows how you can preview the settings that would be applied before actually changing any policy rankings.

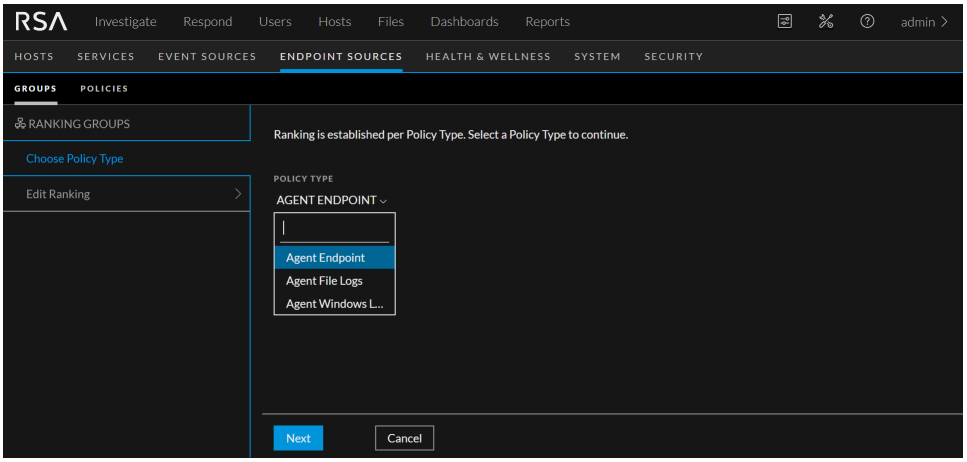
Change Policy Ordering for Groups

An endpoint agent can be included in multiple groups. And these groups can have different policies applied to them. In this case, you can edit the ordering or ranking of policies, to specify a hierarchy for your policies.

Edit Ranking

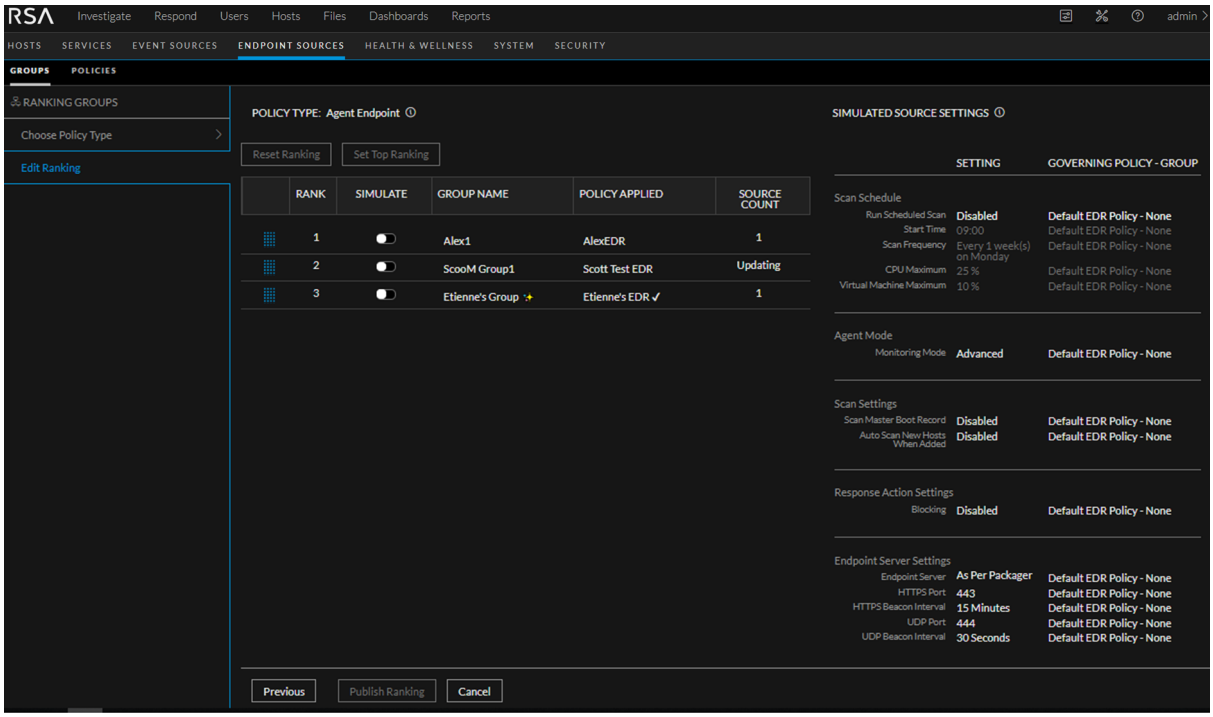
To edit the ordering or ranking of a group:

1. Go to  **(Admin) > Endpoint Sources**.
2. Select the Groups tab and click **Edit Ranking**.
3. Select one of the following source type for the drop-down list:
   - **Agent Endpoint** to rank the groups associated with Agent Endpoint type policies.
   - **Agent Windows Logs** to rank the groups associated with Agent Windows Log type policies
   - **Agent Log Files** to rank the groups associated with agents that are using File Collection.



4. Click **Next**.

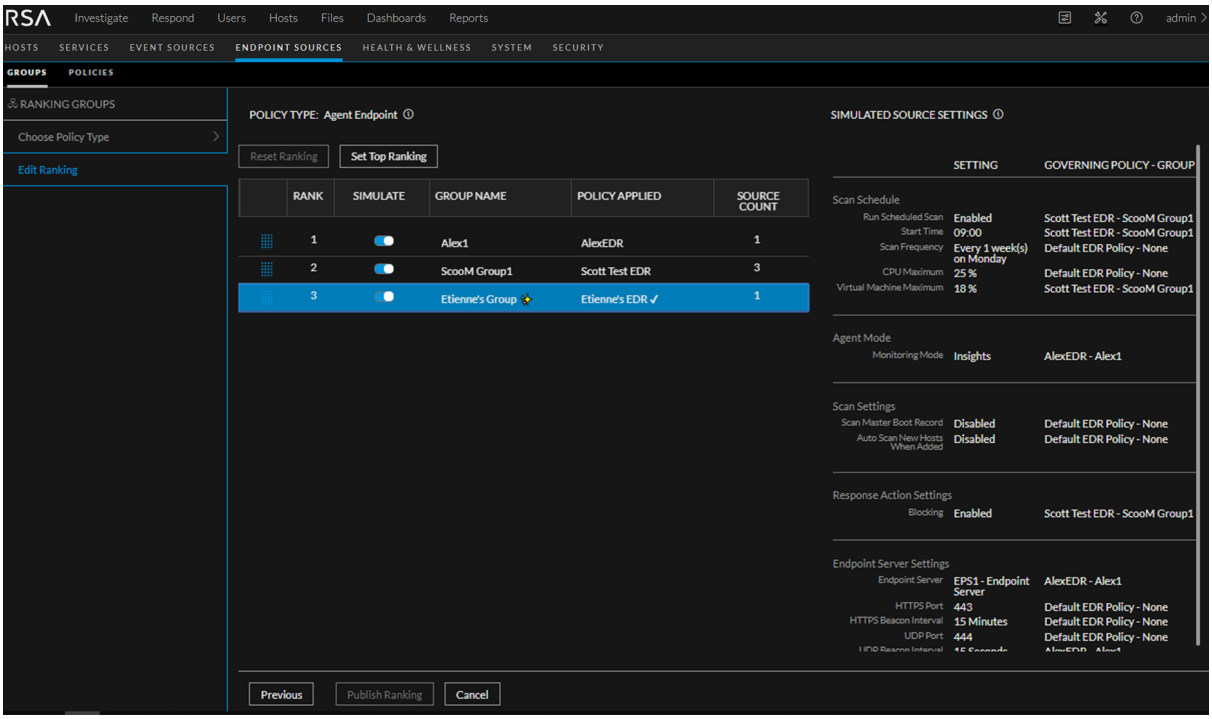
   The Ranking view is displayed:



You can simulate your policy settings and how they affect the endpoints within their groups. This gives you the ability to preview how your changes to the ranking will affect the policy settings applied to each hypothetical agent.

5. You can manipulate the sliders to simulate different options. You can select the slider in the **Simulate** column for each group into which a hypothetical agent would fall, and drag it to the right to simulate turning on that policy.

   This image shows all policies being simulated:



6. Reorder your groups as necessary.
   a. Select anywhere within a group's row.
   b. Drag the group up or down to change the priority. Priority decreases from top to bottom.
   c. Repeat moving groups until they are ordered as you prefer.

   Note: To move any group to the top, select the group and click **Set Top Ranking**.

7. As you change the rankings for your groups, you can preview how the policy settings would change based on your new rankings. For example, assume each of the following scenarios:
   - If you have a hypothetical agent that belongs in group 1, simulate only group 1 to see which policies would affect that agent.
   - If you have a hypothetical agent that belongs in group 1 and group 2, simulate both group 1 and group 2 to see how policies would be applied for that agent.
8. After you have specified the optimal ranking order, you can publish the new ranking.

Simulation Examples

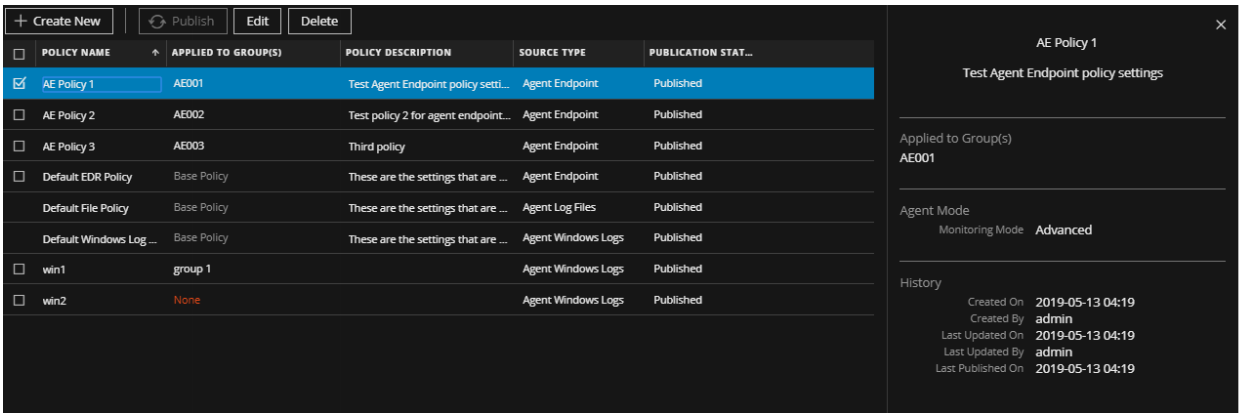
This section contains simulation examples for Agent Endpoint policies and Agent Log policies. Note that Windows policies have the same behavior as Agent Endpoint policies.

Agent Endpoint Policies Examples

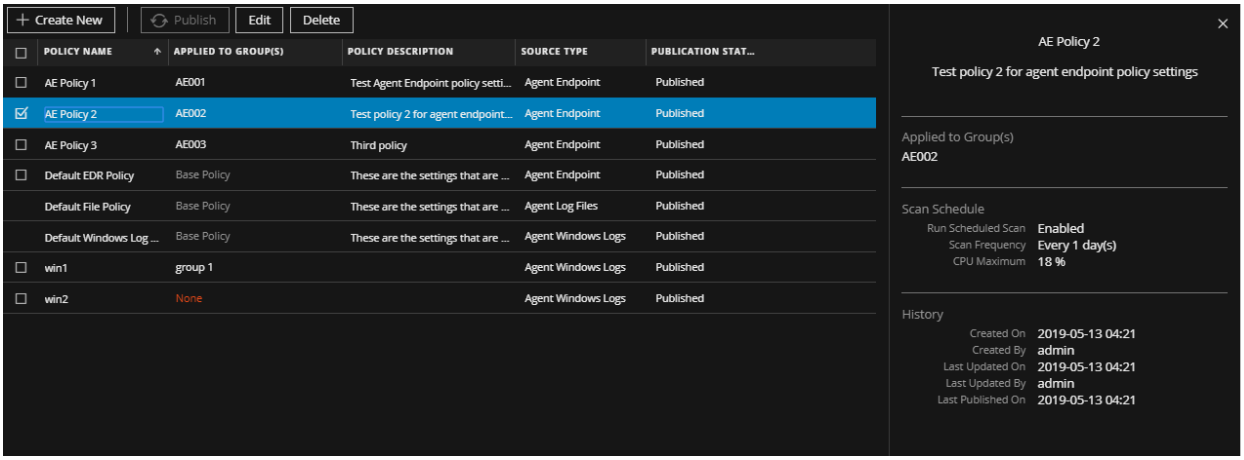
The **Simulated Source Settings** panel shows each individual policy setting, along with the governing policy for each setting.

Examine **AE Policy 1**, **AE Policy 2** and **AE Policy 3** to see which values are set in those policies.

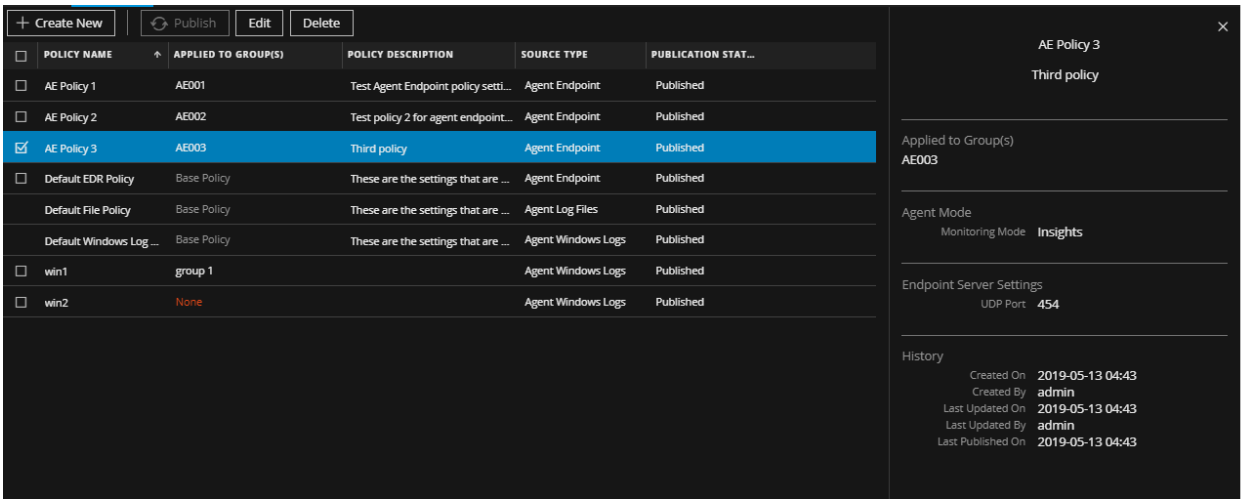
AE Policy 1 has Agent Mode set:



AE Policy 2 has a scan schedule set, as well as limiting the CPU maximum to 18%:

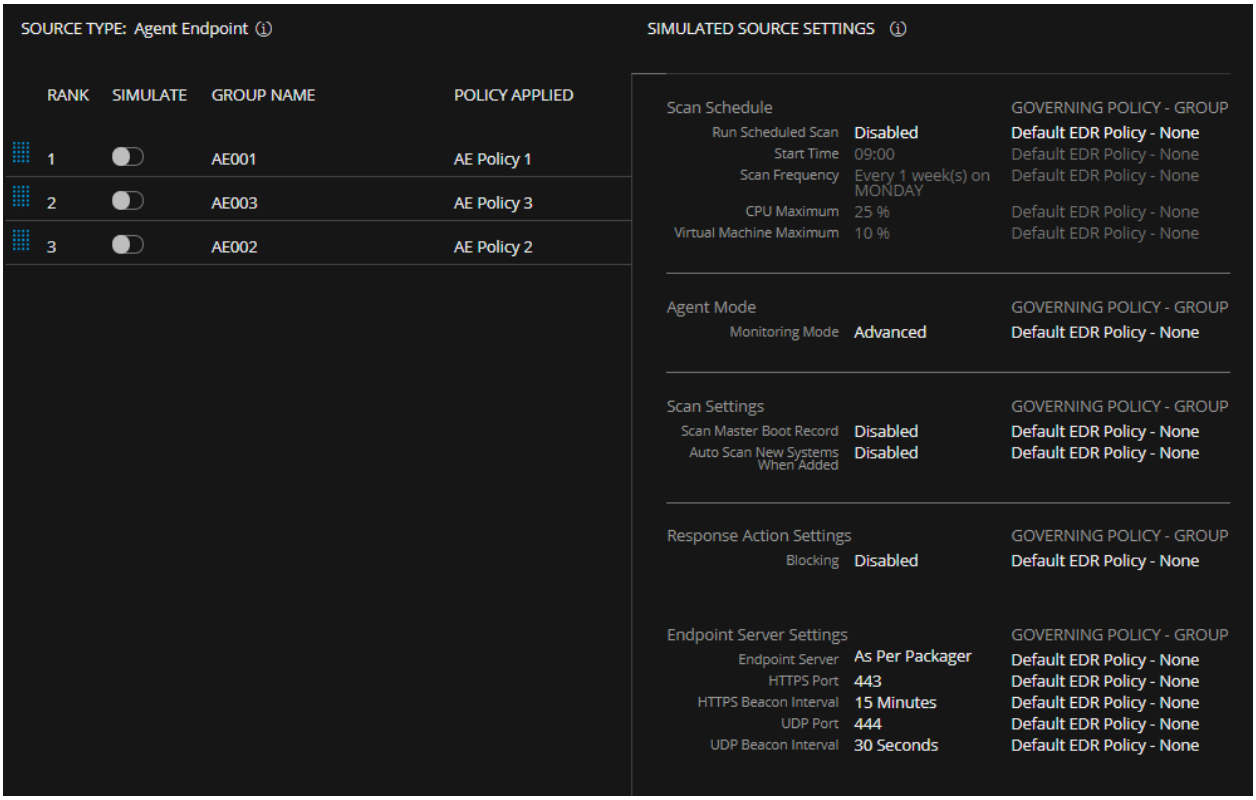


AE Policy 3 has Agent Mode set (to **Insights**), and sets the UDP Port to **454**:



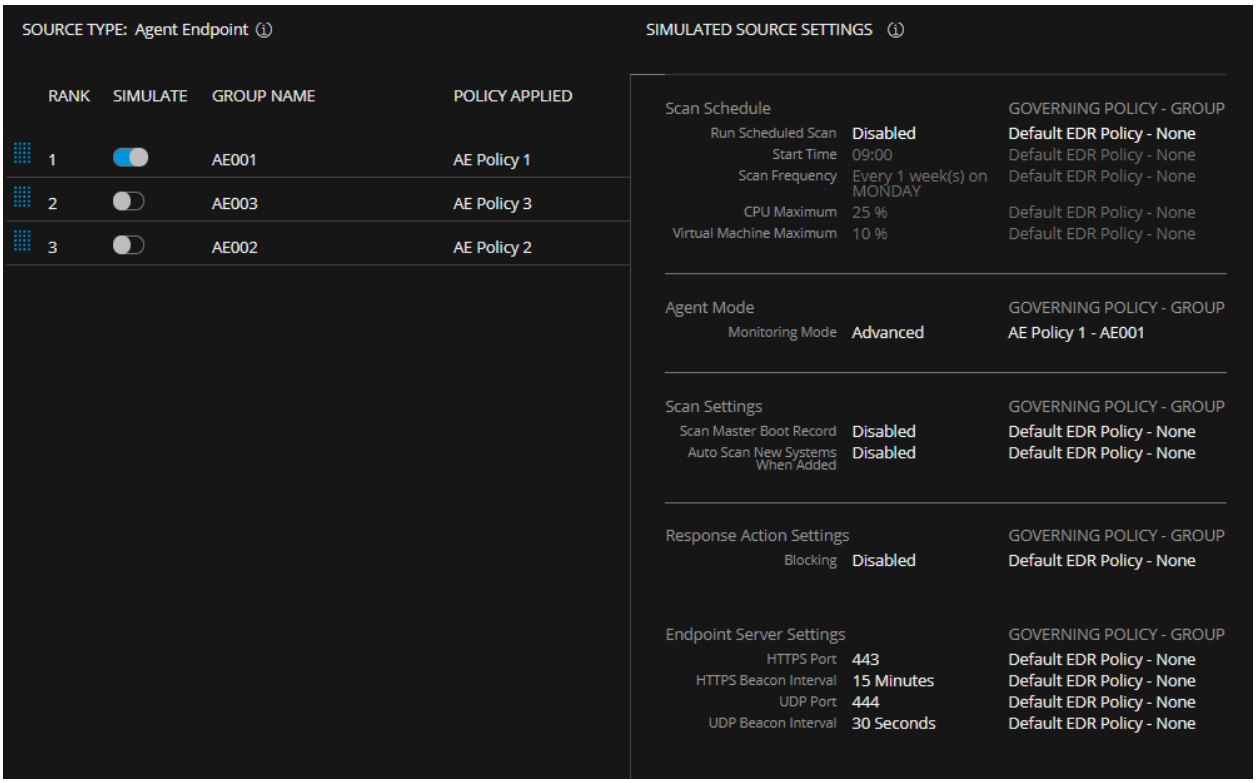
No Settings Applied

When none of the policies are simulated, you can see that the Default EDR Policy governs all behavior.



Simulate a Single Policy

When you select the AE001 Group, you can see that the values set in AE Policy 1 govern the behavior Agent Mode, as well as the Default EDR Policy being used for all unset parameters.



Simulate Multiple Policies

When you select multiple groups and policies, you can see the effects of each policy, based on the current ranking.



You can see that Agent Mode is set to **Insights**: this is because **AE Policy 3** is ranked above **AE Policy 1**, so the higher ranked policy's setting is used. None of the other parameters are set in more than a single policy, so for those, each policy's setting is used. Since **AE Policy 1** only has Agent Mode set—and **AE Policy 3** ranks higher and sets a different value—**AE Policy 1** does not govern any of the EDR settings. And finally, for parameters not set in any of the simulated policies, the Default EDR Policy settings are used.

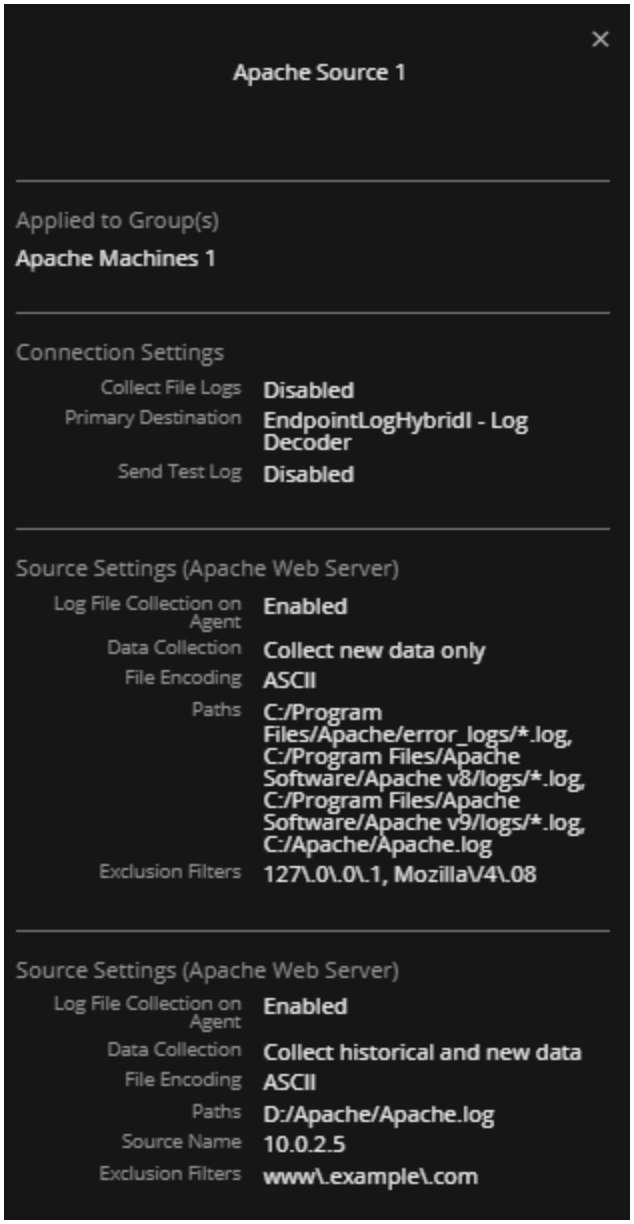
File Log Policies Simulation Example

The collection of parameter values into a complete policy works a bit differently for File Log policies than for EDR or Windows policies. To determine which values are applied from which policies, note the following:

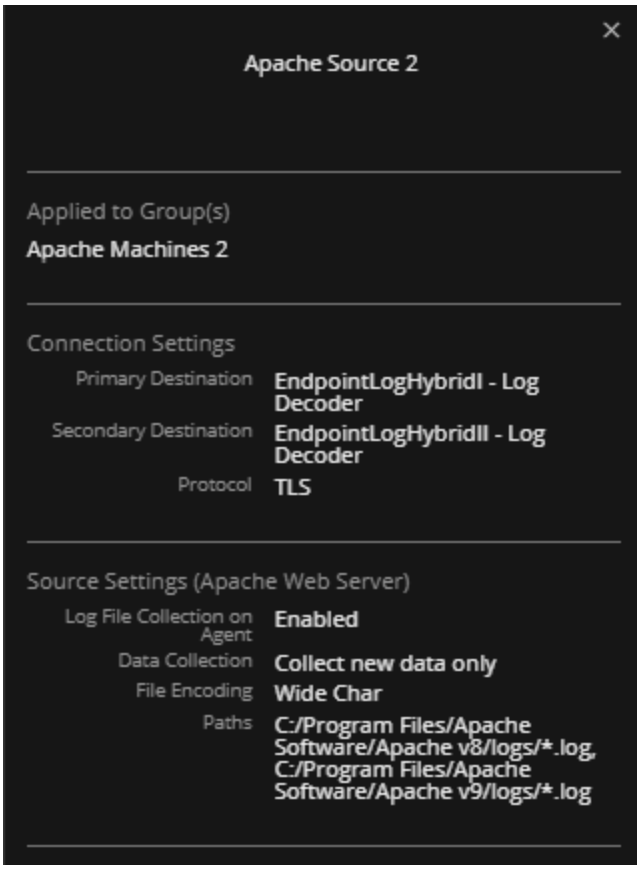
- Each event source type acts as a separate setting. That is, if a policy has values for both Apache and RSA Access Manager, for example, each of those event source types is treated as a separate set of values.
- When an endpoint inherits values, it might not only get values from the highest ranked policy that has a value set. It might, for example, inherit Apache values from one policy, and RSA Access Manager settings from another.
- If you consider a set of File Log policies that all include settings for the same event sources, then they behave the same as EDR and Windows policies.

Consider the following example, where there are three log file policies—two with Apache source types and one with MS SQL event source type.

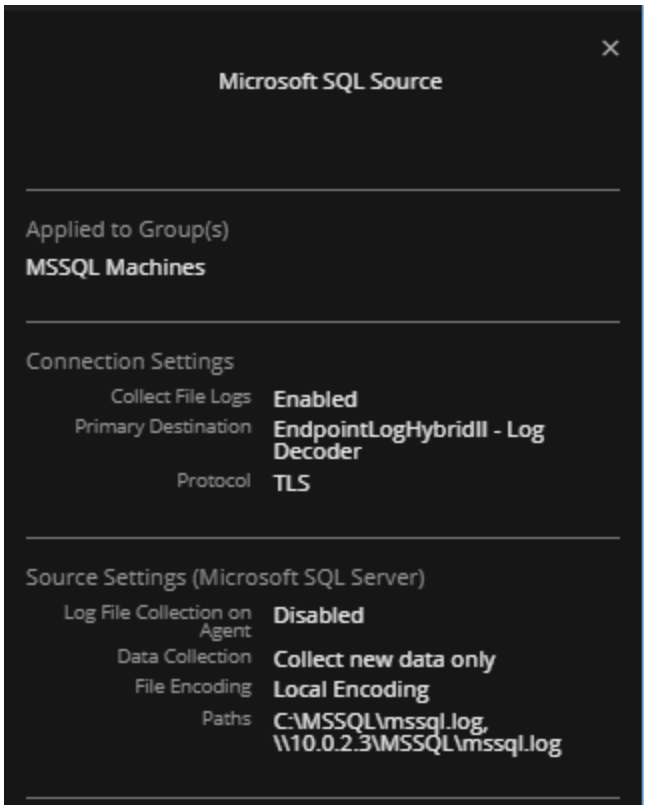
- File policy **Apache Source 1**:



- File policy **Apache Source 2**:

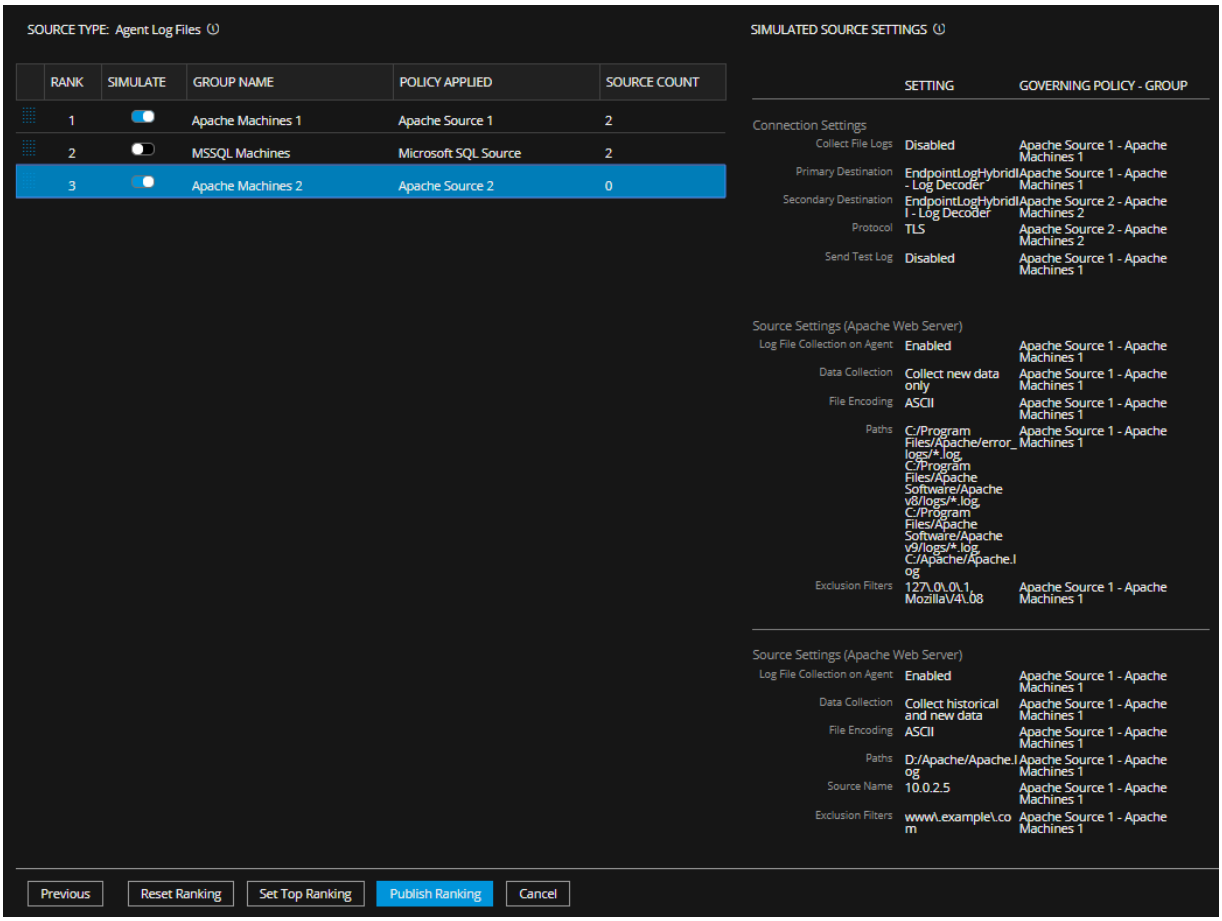


- File policy **Microsoft SQL Source**:



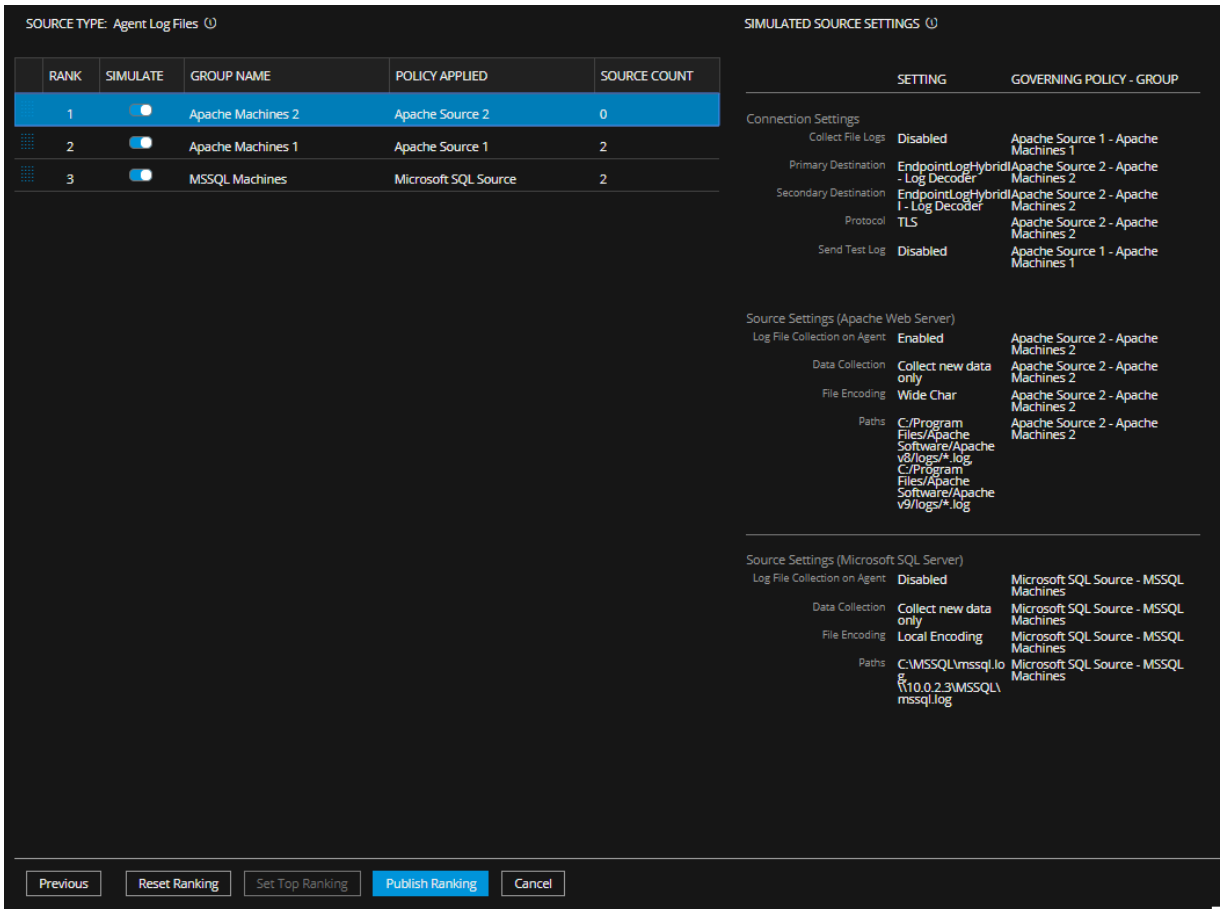
Examine two of the possible ranking orders:

- Simulate policies 1: **Apache Source 1**, 2: **Apache Source 2**. Note that in this case, the MS SQL policy is not simulated.



In this case, all Apache settings for the group are inherited from the **Apache Source 1** policy.

- Move the **Apache Source 2** policy higher than **Apache Source 1** policy, and add the **Microsoft SQL Source** policy. Simulate policies 1: **Apache Source 2**, 2: **Microsoft SQL Source**, 3: **Apache Source 1**.



In this case, all Apache settings for the group are inherited from the **Apache Source 2** policy, and groups also get the MS SQL settings.

So, the Apache settings are inherited from the highest ranked Apache policy **only**, but the source settings as a whole are **combined** to include settings from each event source type.

The SIMULATE Slider

The slider has two positions:

- On: 
- Off: 

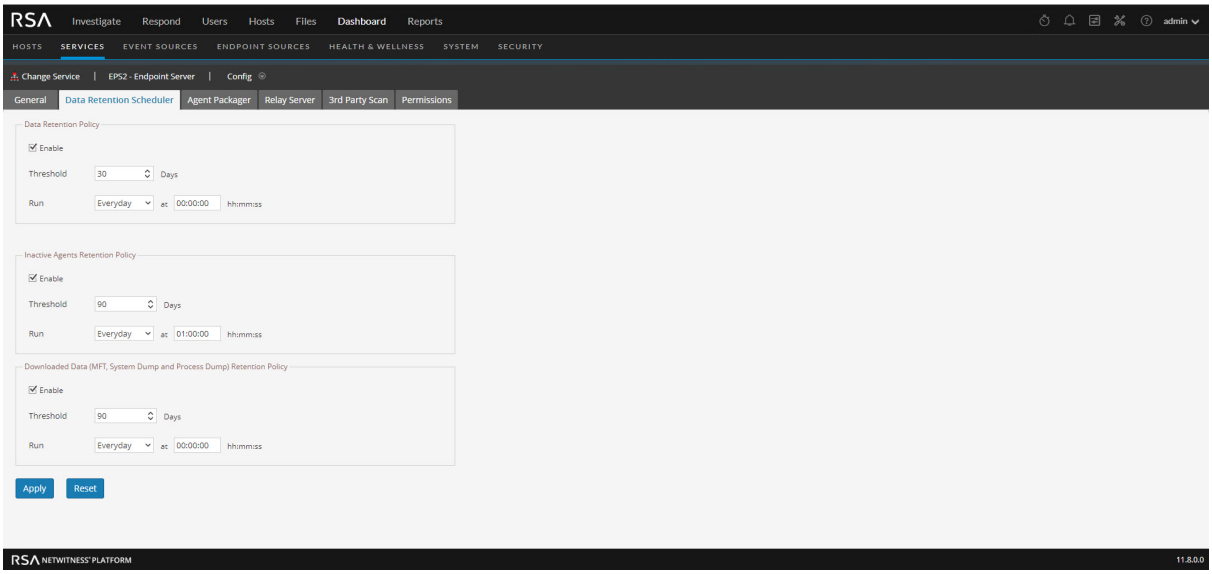
If the simulate slider is ON for a policy, that policy's values are factored into the complete set of governing settings. If the slider is OFF for a policy, the setting for that policy have no effect on the list of the governing settings.

Configuring Data Retention Policy

An administrator can configure the retention policies to retain the Endpoint data based on the age or the storage size. By default, days and size-based retention policies are enabled.

To change the configuration for age-based retention:

1. Go to **(Admin) > Services**.
2. In the Services view, select the **Endpoint Server** service.
3. Click and select > **View** > **Config**.
4. Click the **Data Retention Scheduler** tab.

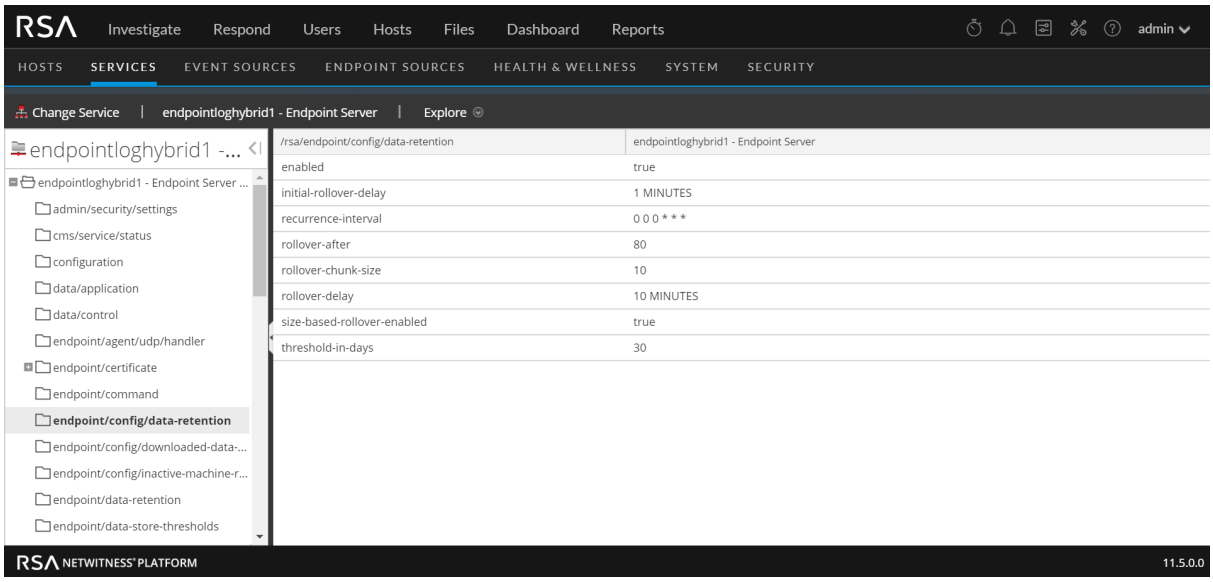


5. In the **Data Retention Policy** panel, by default, the **Threshold** is set to 30 days, and **Run** to Everyday. This means only 30 days of Endpoint data is retained and the older data is deleted from the database.
6. Click **Apply**.

To change the configuration for size-based retention:

By default, for the size-based retention, the `rollover-after` value is set to 80 and `rollover-chunk-size` is set to 10. This means that when the storage size exceeds 80 percent of the space allocated for the disk partition, 10 percent of the older Endpoint data is deleted from the database. However, you can change these values as follows:

1. Go to **(Admin) > Services**.
2. In the Services view, select the **Endpoint Server** service.
3. Click and select > **View > Explore**. The Explore view is displayed:

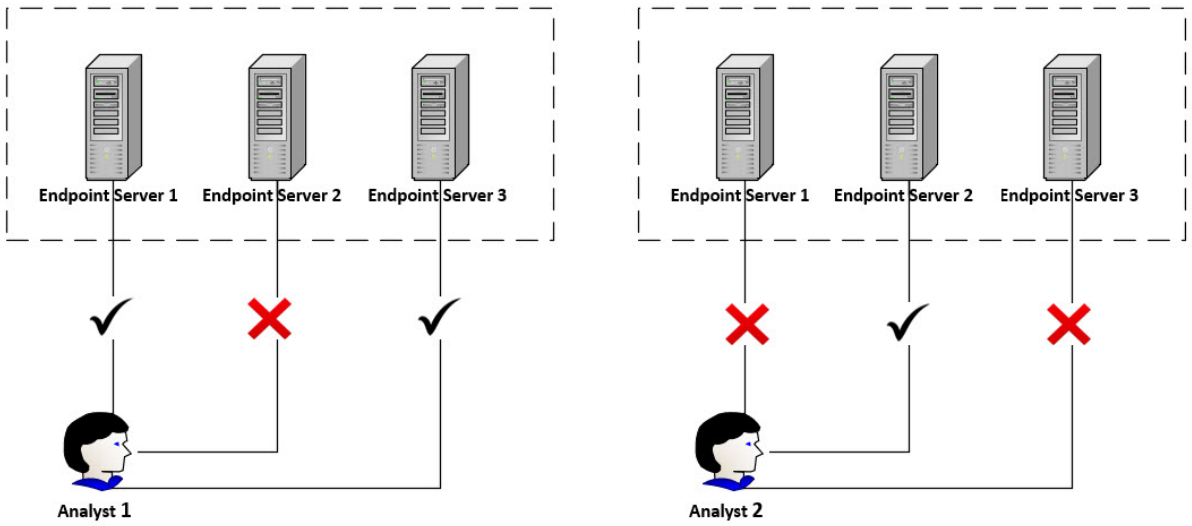


4. In the left panel, select **endpoint/config/data-retention**.
5. Edit the configurations based on your requirements.

Manage Role Permissions at Endpoint Server Level

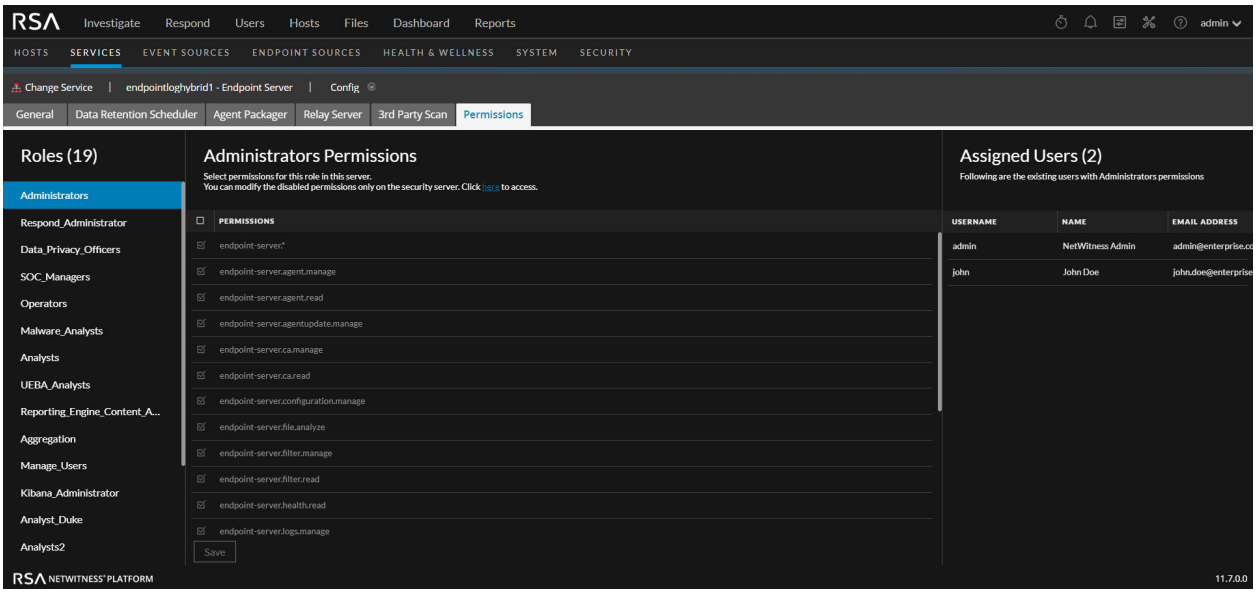
NetWitness Platform version 11.7 and later provides more flexibility for managing RBAC (Role-Based Access Control). This enhancement enables you to grant/revoke access for any role at individual endpoint server levels in an environment rather than all. On the **admin > Config** view for any endpoint server, you can navigate to the **Permissions** tab and manage a significant number of permissions for the selected endpoint server. A new permission called **endpoint-server.file.analyze** lets the user perform file analysis activities(Analyze Files, Save a Local Copy and Initiate OPSWAT Scans), allowing more flexible options to manage user permissions. Refer to Manage Users with Roles and Permissions on the *System Security and User Management Guide* for more information.

The following illustrations show how you can set permissions at the endpoint server level for an analyst.



To configure permissions at the endpoint server level:

1. Go to **(Admin) > Services**.
2. In the Services view, select the Endpoint Server service.
3. Click  and select **View > Config**.
4. Navigate to **Permissions** tab.
5. Select a role from the **Roles** pane on the left. (For more information on managing roles, refer (Optional) Add a Role and Assign Permissions in *System Security and User Management Guide*.)
6. Select the permissions from the **PERMISSIONS** section (Center pane).
7. Click **Save**. The **Assigned Users** pane on the right shows the list of users assigned with the selected role and permissions.
8. Repeat steps **5 – 7** if you want to modify permissions to any role. Make sure you save the changes before navigating out of the current role.



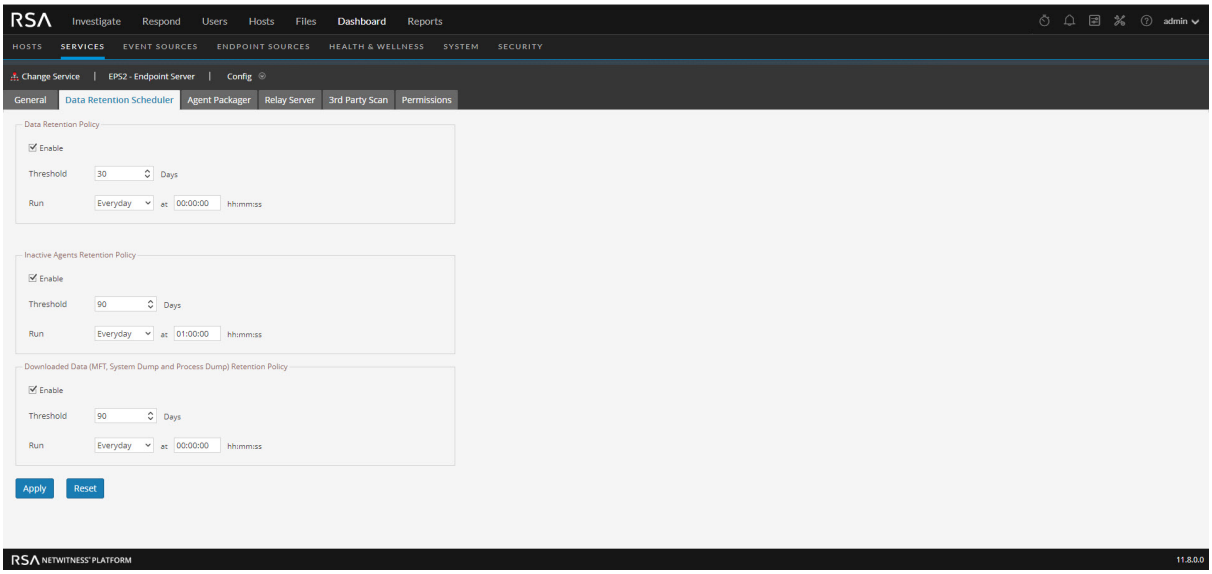
Note: Some of the permissions on the Permissions tab are disabled, and you can modify them only on the Security server (Admin > Security). All other permissions listed here are editable on both the Security server and the Permissions tab. For more information on managing permissions and on the NetWitness Platform, see *System Security and User Management Guide*.

Managing Inactive Agents

An administrator can configure the inactive agent retention policy to delete data of agents that are inactive, from the Endpoint Server. On deletion, the Endpoint Server stops collecting data from these agents. By default, this option is enabled.

To configure the inactive agent retention policy:

1. Go to  **(Admin) > Services**.
2. In the Services view, select **Endpoint Server**.
3. Click  and select > **View** > **Config**.
4. Click the **Data Retention Scheduler** tab.



5. In the **Inactive Agents Retention Policy** panel, by default, the **Threshold** is set to 90 days and **Run** to Everyday. This means that the data of agents that have not communicated with the Endpoint server for 90 days is deleted from the database.
6. Click **Apply**.

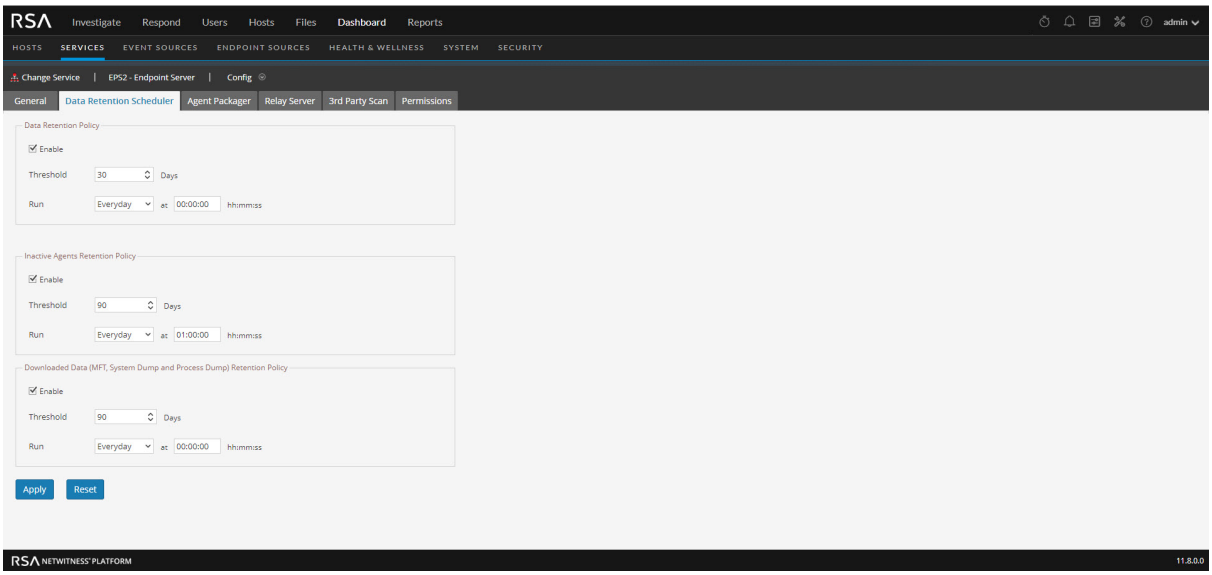
Note: The Inactive Agents Retention Policy is not applicable for NetWitness Endpoint 4.4.0.2 or later agents.

Configure Retention Policy for Memory Dumps and MFT

An administrator can configure the retention policy to delete the downloaded system dump, process dump, and Master File Table (MFT) from the Endpoint server based on the number of days. By default, this option is enabled.

To configure the retention policy:

1. Go to  **(Admin) > Services**.
2. In the Services view, select **Endpoint Server**.
3. Click  and select > **View** > **Config**.
4. Click the **Data Retention Scheduler** tab.



5. In the **Memory Dumps and MFT Retention Policy** panel, by default, the **Threshold** is set to 90 days and **Run** to every day. This means only 90 days of data is retained and the older data is deleted from the Endpoint server.
6. Click **Apply**.

(Optional) Installing and Configuring Relay Server

Note: The information in this topic applies to NetWitness Version 11.3.1 and later.

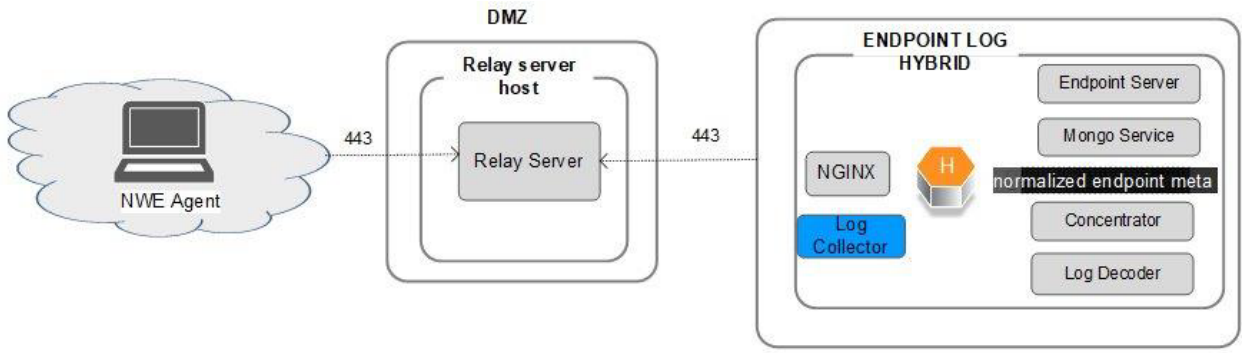
Relay Server (referred to as RAR in RSA Endpoints) extends NetWitness Platform's visibility into endpoints when they are outside the corporate network. The Relay Server deployed in a cloud or DMZ relays the endpoint data between the hosts and the Endpoint Server. The hosts that are outside the corporate network send the endpoint data to the configured Relay Server and the corresponding Endpoint server pulls the data.

Note: If you have Windows hosts that are outside the corporate network, the log data is not sent to the Relay Server.

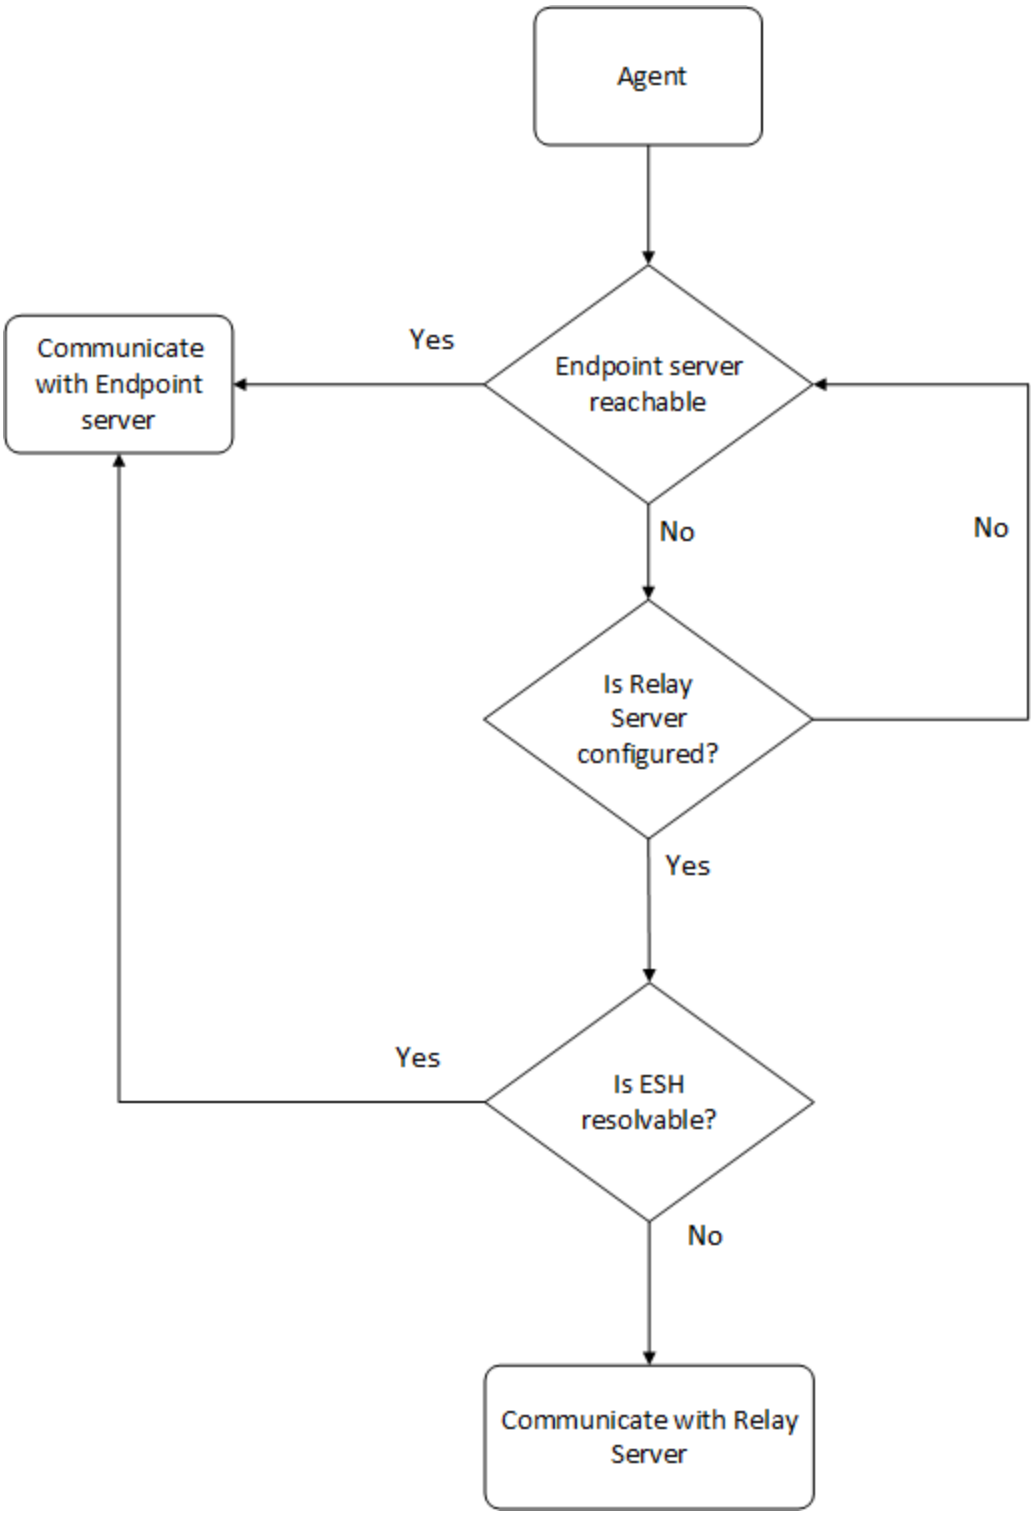
You can configure a Relay Server on the Endpoint Server Config view. Once the Relay Server is configured, the policy for the host is automatically updated and you can view the Relay Server settings on the **Host Details** view > **Policy Details** panel.

You can configure a single Relay Server with one or more Endpoint servers. In this case, the Relay Server ensures that the endpoint data reaches the Endpoint Server configured in the policy.

The following describes the architecture of the Relay Server.



The following flowchart explains how the host switches to the Relay Server.


Agent
Endpoint server reachable
Yes
Communicate with Endpoint server
No
No
Is Relay Server configured?
Yes
Yes
Is ESH resolvable?
No
Communicate with Relay Server

Installing the Relay Server

The Relay Server installer contains binaries, certificates, configuration files, and the installation script required to install the Relay Server.

IMPORTANT:
- The Relay Server version must match with the corresponding NetWitness Endpoint Server version. If you plan to upgrade a Relay Server to a newer version, first upgrade the Endpoint Server, then download the Relay Server installer, and run the installer script.
- Operating System updates and general system hardening on the Relay Server must be managed by the customer according to standard best practices. The Relay server package does not contain OS updates and the operating system will not be updated as part of the standard NetWitness update process.
- Do NOT run the nwsetup-tui script to install the Relay Server. Follow the instructions in this document only as Relay Server is an independent server and not part of NetWitness UI (Admin > Hosts).

Installation Media

The Relay Server can be installed only on a CentOS 7 or NetWitness Platform 11.4.0.0 base image which is available for download from Download Central (https://download.rsasecurity.com). Also, make sure that the Relay server host is connected to internet to download the required dependencies.
For more information on deploying Relay Server host on a:

- DMZ - see "Step 1a. Deploy the Virtual Host to create VM" in the *Virtual Host Installation Guide*.
- Cloud
  - see "Step 1. Deploy NW Server Host" in the *Azure Installation Guide*.
  - see "AWS Deployment" in the *AWS Installation Guide*.

Relay Server Host System Requirements

Agents	RAM	CPU Cores	Disk	Ideal Beacon Interval
20000	32 GB	4 cores	200GB	5 min

To install the Relay Server:

1. Log in to NetWitness Platform.
2. Click  **(Admin) > Services**.
3. Select the **Endpoint Server** service and click  > **View** > **Config** > **Relay Server** tab.
4. In the **Download Installer** section, enter the installer password and click **Download** to download the Relay Server installer file (**RelayInstaller.zip**).
5. Copy the Relay Server installer file (**RelayInstaller.zip**) to the Relay Server host.

6. Unzip the **RelayInstaller.zip** file on the Relay Server host. For example:

   ```
   /home/RelayInstaller.zip
   unzip <installer path>
   ```

7. Set up the execution permission using the following command:

   ```
   chmod +x install.sh
   ```

8. Run the installer script using the following command:

   ```
   ./install.sh
   ```

   The **All necessary RPMs will be installed without further** prompts is displayed.

9. Enter **Y** to continue the installation.

   The password prompt is displayed.

10. Enter the password.

    Note: Make sure you enter the same password you set while downloading the Relay Server installer.

    Note: In case if you are re-installing the Relay Server host. **Do you wish to update the list** prompt is displayed.
    - Enter **Y** to update the Endpoint server IPs.

    Enter the Endpoint Server IPs prompt is displayed.

11. Enter all the Endpoint server IPs you plan to configure with the Relay server with comma separated.

If the Relay Server installation is successful, you can check the status of the services:

- Check if the Relay Server is up and running:

  ```
  systemctl status rsa-nw-relay-server
  ```

- Check if Ngnix is running:

  ```
  systemctl status ngnix
  ```

You can also update Endpoint Server IPs without installing the Relay Server.

To update Endpoint Server IPs without installing the Relay Server:

1. Run the following command:

   ```
   bash /var/netwitness/relay-configure-allowed-hosts.sh
   ```

   The list of all the configured Endpoint server IPs is displayed and **Do you wish to update the list** prompt is displayed.

2. Enter **Y** to update the list of Endpoint server IPs.

   Enter the Endpoint Server IPs prompt is displayed.

3. Enter a comma-separated list of all the Endpoint Server IPs to update.

   The list of updated IPs is displayed.

Configuring the Relay Server

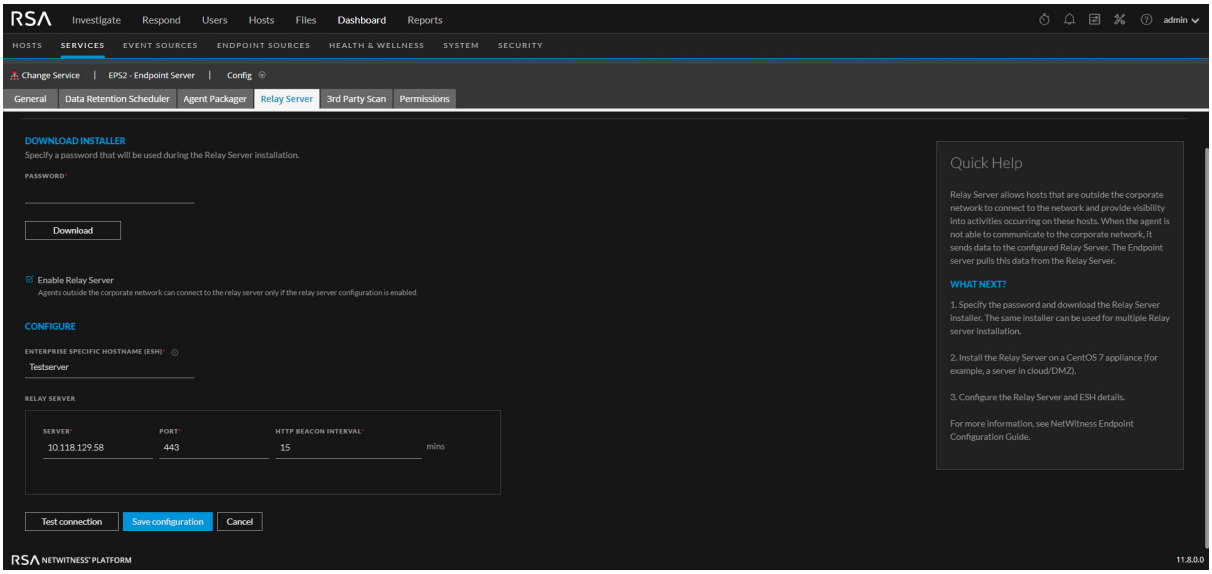
Make sure you have installed the Relay Server.

Note: During Relay Server host installation, firewalld is configured to allow incoming connections only on TCP ports 443 and 22.

To configure the Relay Server:

1. Log in to NetWitness Platform.
2. Click **(Admin) > Services**.
3. Select the **Endpoint Server** service and click > **View** > **Config** > **Relay Server** tab.

   The Relay Server tab is displayed.



4. Select the **Enable Relay Server** check box to enable the Relay Server configuration.

   Note: To disable the Relay Server, clear the **Enable Relay Server** check box.

   Caution: Before you disable the Relay Server configuration, if the hosts will be always roaming make sure to migrate these hosts to an alternate Endpoint server configured with a different Relay server. Else these hosts will not be able to connect back to the corporate network. When you disable the configuration, the Relay Server settings are removed from the EDR policy.

5. In the **Configure** section:

a. Enter the ESH.

IMPORTANT: This should a hostname that can only be resolved on the corporate internal DNS and not on public internet DNS (ex: 1.1.1.1 or 8.8.8.8). It is used to determine if the host running the Endpoint agent client is connected to the internal network (physically on premise or via VPN).

b. Specify the Relay **Server**, **Port** and **HTTP Beacon Interval**.

IMPORTANT: The Relay Server needs to be a hostname that can be resolved correctly on both internal DNS (with the internal IP) and public DNS (with the public IP)

6. Click **Test Connection** to check if the Relay Server is reachable.
7. Click **Save Configuration** to save the configuration.

Note: Before you modify the Relay Server configuration, perform any one of the following:
- Make sure that the hosts are inside the corporate network so that the policy with the Relay Server configuration is applied.
- If hosts will always be roaming, then migrate these hosts to an alternate Endpoint server configured with a different Relay Server.

IMPORTANT: You must change the root password after you deploy the Relay Server host.

Configure YARA

YARA helps analysts in identifying and classifying malware in a simple and effective manner. By default YARA is disabled, administrators can configure and enable YARA. You are required to provide YARA rule files and specify the path where these rule files are stored. Ensure that the Endpoint server and YARA rules directory are in the same appliance.

Note: The YARA rules directory should have the permission as 'netwitness'.

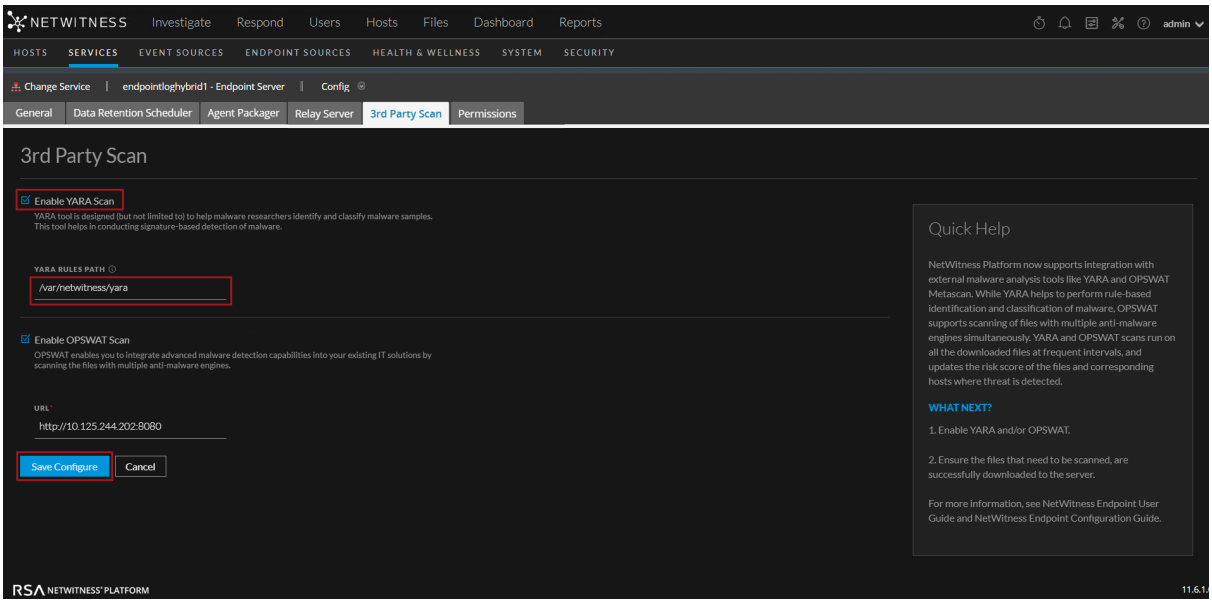
Note: For new installations, YARA will be included as part of the orchestration. If you are upgrading to NetWitness Platform 11.6 or later, YARA is included automatically. In case of an upgrade, deploy the following Yara App rules from Live-

- Yara Rule Matched
- Process with Matched Yara Rule

IMPORTANT: Ensure that YARA is configured and YARA RULES PATH is same across all the Endpoint servers.

To change the configuration for YARA:

1. Go to  **(Admin) > Services**.
2. In the Services view, select the **Endpoint Server** service.
3. Click  and select **View** > **Config**.
4. Click the **3rd Party Scan** tab.



5. Select **Enable YARA Scan.**

6. In **YARA RULES PATH**, specify the directory path where the YARA rule files are stored.

   Note: You can add any number of YARA rule files in this directory. Each rule file can have more than one rule. All downloaded files are scanned by all the YARA rule files.

7. Click **Save Configure**.

Configure OPSWAT

OPSWAT (MetaDefender Core) provides advanced malware detection capabilities by scanning files with multiple anti-malware engines simultaneously. OPSWAT is disabled by default. As an administrator, you can enable and configure OPSWAT on the endpoint servers.

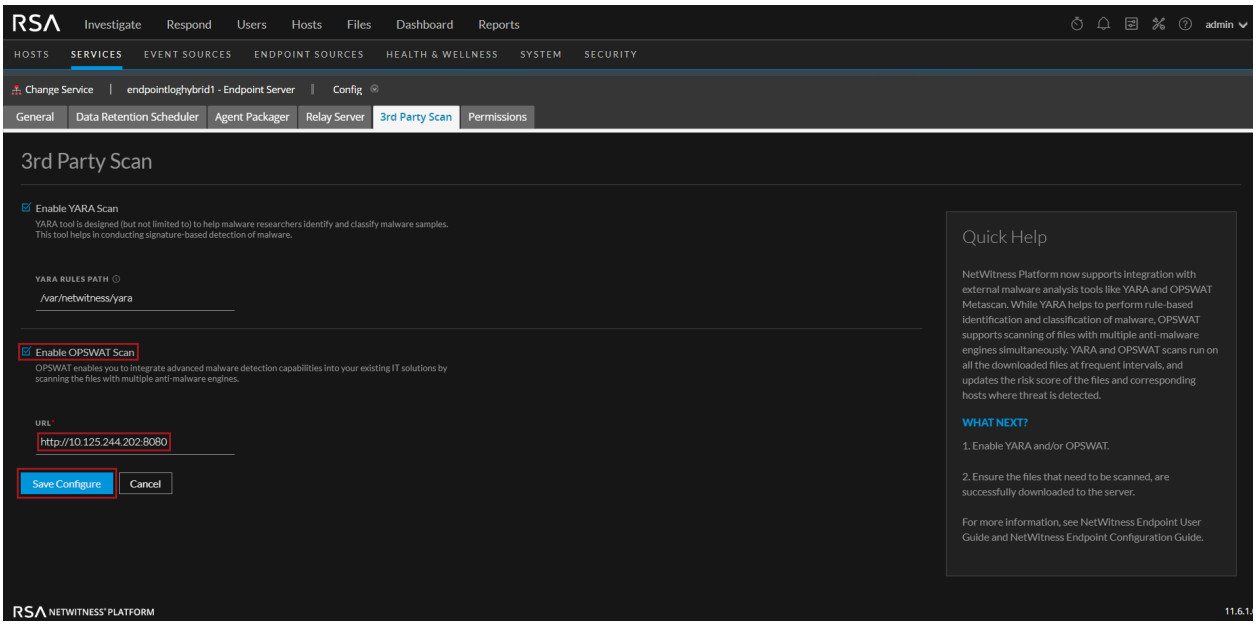
Note: You need to deploy the following OPSWAT app rules from Live services to log decoder for successful configuration of OPSWAT on endpoint servers:

- opswat reported infected
- opswat reported suspicious
- process with opswat reported suspicious
- process with opswat reported infected

IMPORTANT: Ensure that OPSWAT is configured in all endpoint servers and the configuration parameters are consistent across all endpoint servers.

To enable and configure OPSWAT:

1. Go to **(Admin) > Services**.
2. In the Services view, select the Endpoint Server service.
3. Click  and select **View > Config**.
4. Click the **3rd Party Scan** tab.
5. Select **Enable OPSWAT Scan**.
6. Enter the URL.
7. Click **Save Configure.**


RSA
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Change Service
endpointloghybrid1 - Endpoint Server
Config
General
Data Retention Scheduler
Agent Packager
Relay Server
3rd Party Scan
Permissions
3rd Party Scan
Enable YARA Scan
YARA tool is designed (but not limited to) to help malware researchers identify and classify malware samples.
This tool helps in conducting signature-based detection of malware.
YARA RULES PATH
/var/netwitness/yara
Enable OPSWAT Scan
OPSWAT enables you to integrate advanced malware detection capabilities into your existing IT solutions by scanning the files with multiple anti-malware engines.
URL*
http://10.125.244.202:8080
Save Configure
Cancel
Quick Help
NetWitness Platform now supports integration with external malware analysis tools like YARA and OPSWAT Metascan. While YARA helps to perform rule-based identification and classification of malware, OPSWAT supports scanning of files with multiple anti-malware engines simultaneously. YARA and OPSWAT scans run on all the downloaded files at frequent intervals, and updates the risk score of the files and corresponding hosts where threat is detected.
WHAT NEXT?
1. Enable YARA and/or OPSWAT.
2. Ensure the files that need to be scanned, are successfully downloaded to the server.
For more information, see NetWitness Endpoint User Guide and NetWitness Endpoint Configuration Guide.
RSA NETWITNESS PLATFORM

Integrating NW Endpoint 4.4.0.2 or Later with NW Platform

You can configure the Endpoint Metadata for the NetWitness Endpoint 4.4.0.2 by integrating the Meta Integrator service in the NetWitness Endpoint 4.4.0.2 directly to a Log Decoder. You can view the Endpoint metadata in the **Investigate > Navigate** and **Events** views. This integration includes the following steps:

- Enable Metadata Forwarding
- Enable Machines to Forward Metadata

Enable Metadata Forwarding

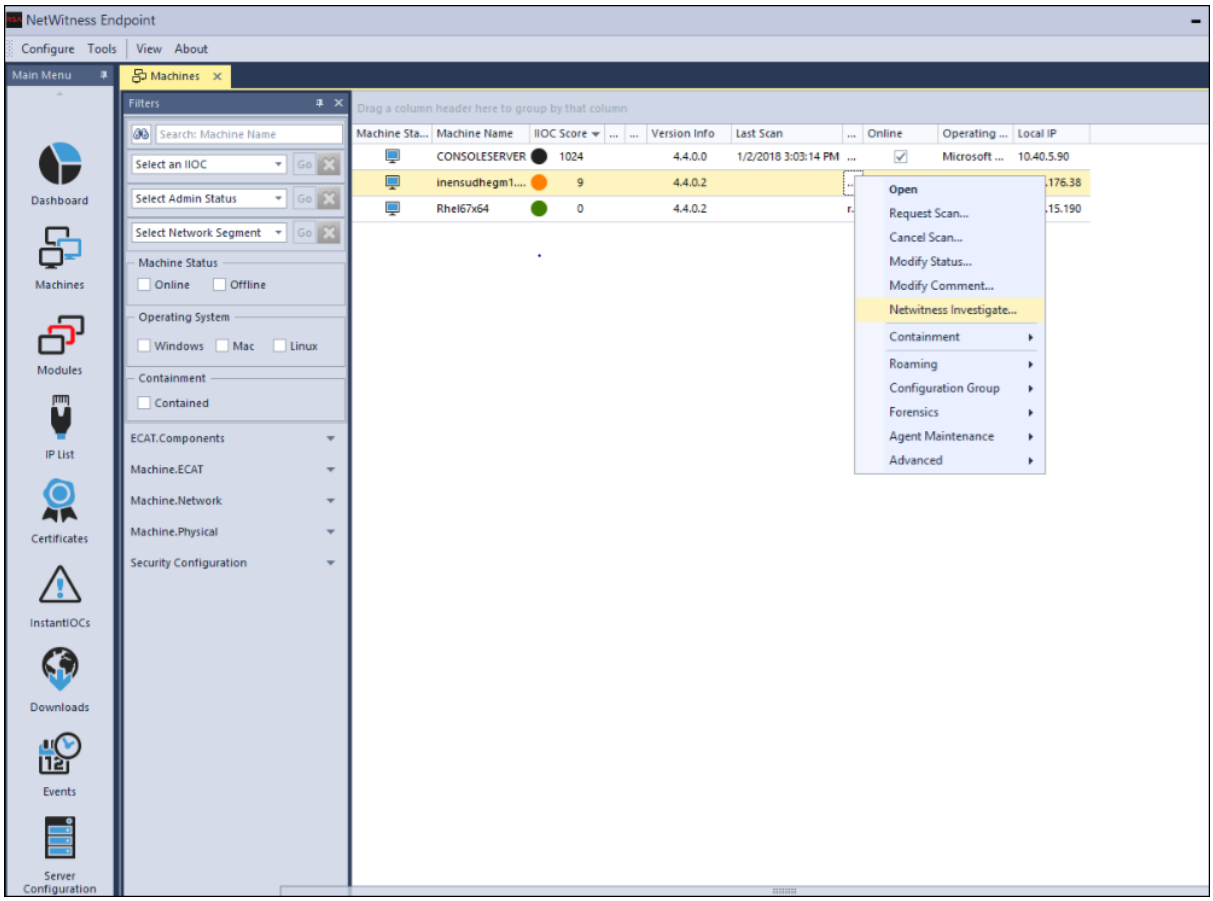
To enable the Metadata Integrator service for the selected NetWitness Endpoint 4.4.0.2 agents, run the following command:

```
ConsoleServer.exe /nw-investigate enable
```

Enable Machines to Forward Metadata

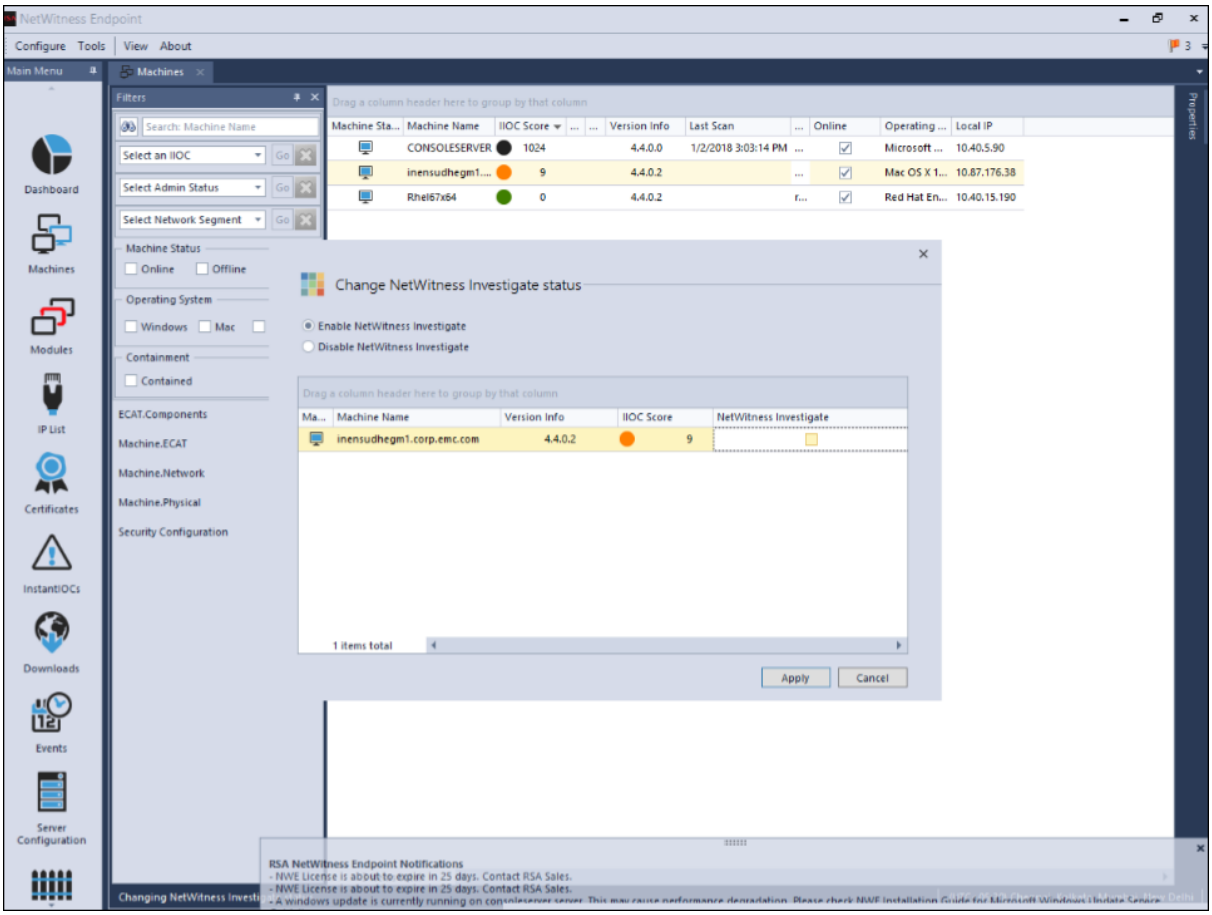
After you enable the Metadata Forwarding using any one of the above options, perform the following to enable the machines to forward metadata to the Log Decoder.

1. Open the NetWitness Endpoint 4.4.0.2 user interface.
2. Click **Machines** from the left panel. The list of available machines are displayed.



3. Select machines for which you want to forward metadata to the NetWitness Endpoint Server.
4. Right-click and select the **NetWitness Investigate** option.

   The Change NetWitness Investigate Status dialog is displayed.



5. Select the **Enable NetWitness Investigate** option.
6. Click **Apply**.
7. To verify if the **Enable NetWitness Investigate** option is enabled, repeat step 4.

Endpoint References

This section is intended to help you understand the purpose of the Services Config View for the Endpoint Server. For each configuration, there is a brief introduction and a What Do You Want To Do table with links to related procedures. In addition, it includes workflow and Quick Look sections to highlight important features in the user interface.

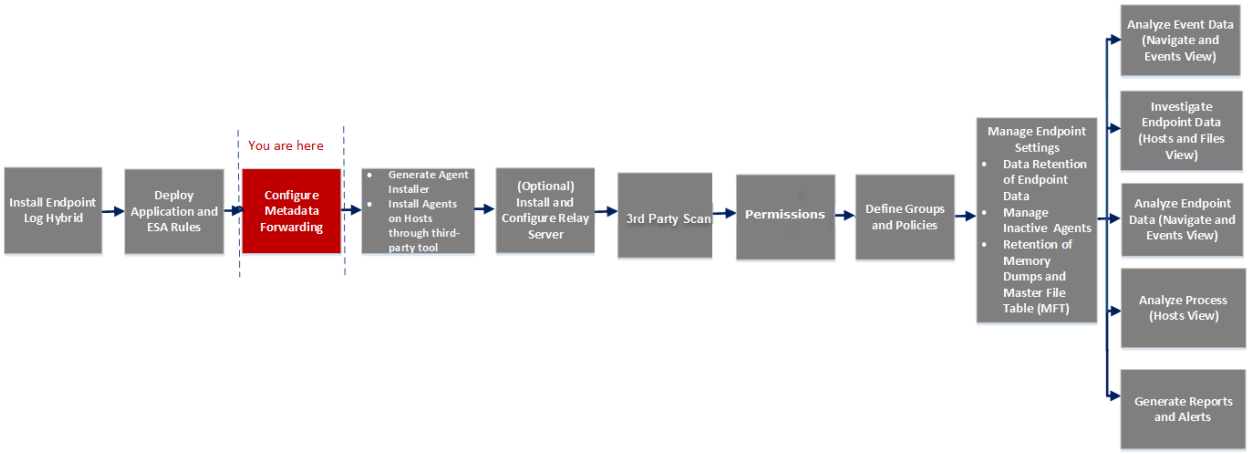
You can view the complete service nodes in tree form in the Services Explore view. For more information, see the "Services Explore View" topic in the *Hosts and Services Getting Started Guide*.

General Tab

In the **General** tab, you can configure the Endpoint metadata forwarding for multiple endpoint servers. To access this view:

1. Go to **(Admin) > Services**.
2. In the Services view, select **Endpoint Server** for which you want to configure the metadata forwarding.
3. Click and select > **View** > **Config**.
4. Click the **General** tab.

Workflow



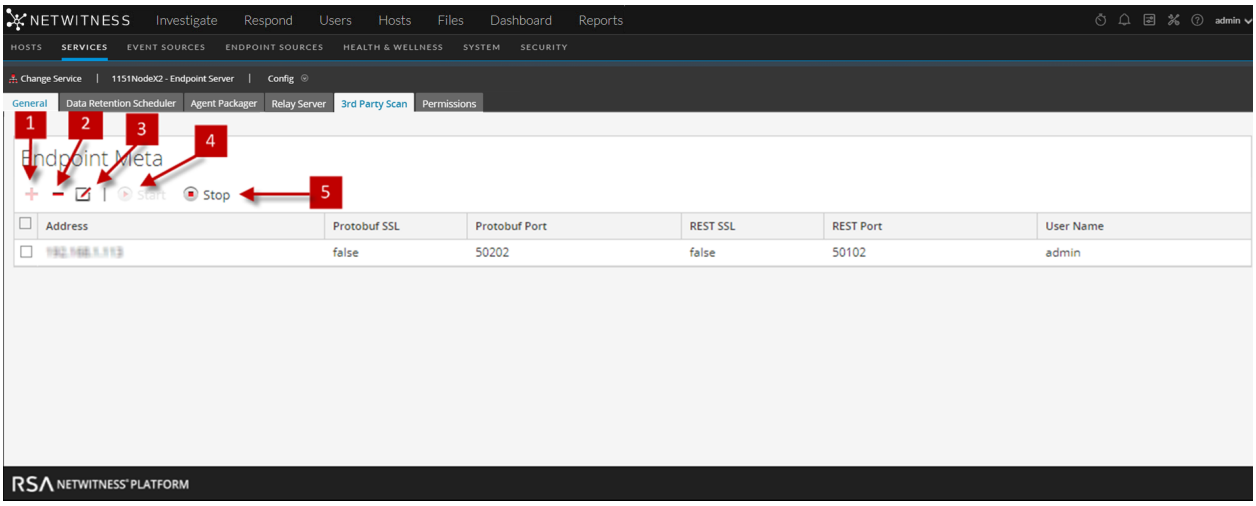
What do you want to do?

User Role	I want to ...	Show me how
Administrator	Configure Endpoint Metadata Forwarding for the NetWitness Endpoint Agents*	Configuring Metadata Forwarding
Administrator	Configure Endpoint Metadata Forwarding for the NetWitness Endpoint 4.4.0.2 or later Agents*	Integrating NW Endpoint 4.4.0.2 or Later with NW Platform

*You can perform this task in the current view

Quick Look

The following figure is an example of the General tab.



1 Click + to view the Available Services dialog.

2 Click − to delete the added service.

3 Click to edit the information for the added service.

4 Click Start to start the Endpoint metadata forwarding.

5 Click Stop to stop the Endpoint metadata forwarding.

The following table describes the fields in the General tab.

Field	Description
Address	Displays the IP address of the Log Decoder.
Protobuf SSL	Indicates if SSL is enabled on Protobuf. By default, this option is disabled.
Protobuf Port	Displays the port used for Protobuf. By default, the port is 50202.
REST SSL	Indicates if SSL is enabled on the REST port in the Log Decoder. By default, this option is disabled.
REST Port	Displays the port used for REST communication. The default value is 50102 (for non-SSL) and value 56102 (for SSL).
User Name	Displays the user name.
Raw Data	Sends a brief summary of the session along with the metadata if enabled. By default, this option is disabled.

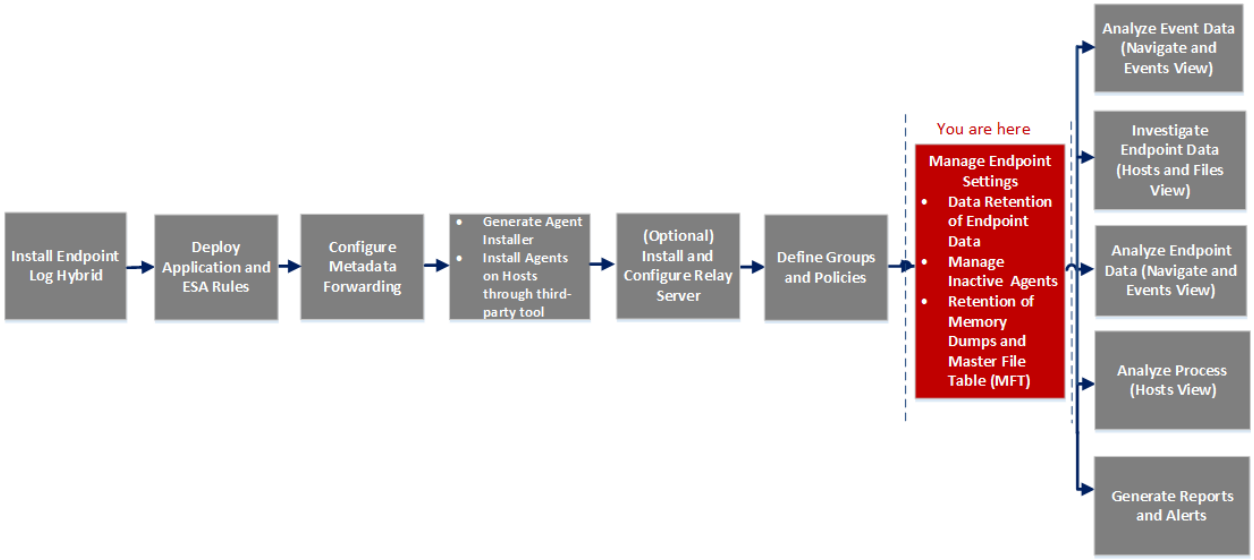
Data Retention Scheduler Tab

In the **Data Retention Scheduler** tab, you can configure data retention, inactive agents, and memory dump and MFT retention policies for multiple endpoint servers. To access this view:

1. Go to  **(Admin) > Services**.
2. In the Services view, select **Endpoint Server**.
3. Click  and select > **View** > **Config**.
4. Click the **Data Retention Scheduler** tab.

Repeat the above steps to configure data retention settings for multiple endpoint servers.

Workflow



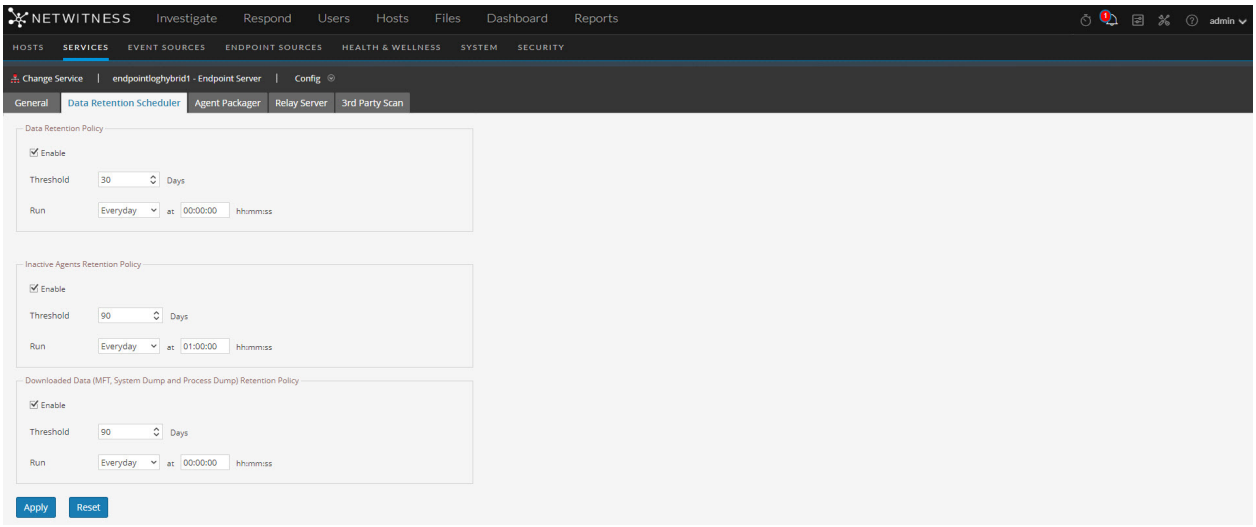
What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure Data Retention Policy*	
Administrator	Configure Inactive Agents Policy*	Managing Inactive Agents
Administrator	Configure Memory Dumps and MFT Policy*	Configure Retention Policy for Memory Dumps and MFT

*You can perform this task in the current view

Quick Look

The following figure is an example of the Data Retention Scheduler tab.



Features

The following table lists the fields for data retention policy.

Field	Description
Enable	Enables the configuration for the data retention policy. By default, this option is enabled.
Threshold	Displays the number of days the Endpoint data is retained in the database. By default, the Threshold is set to 30 days. The data older than 30 days is deleted from the database.
Run	Displays the schedule for running the data retention job. By default, the database check occurs everyday at 00:00:00 AM. You can select the frequency from the drop-down list (Everyday, Weekdays, Weekends, or Custom, where Custom allows you to select one or more specific days of the week) and time to run the job.
Apply	Overwrites any previous schedule for the data retention policy and applies the new schedule immediately.
Reset	Resets the schedule to the default settings.

The following table lists the fields for inactive agents retention policy.

Fields	Description
Enable	Enables the configuration for the inactive agents policy. By default, this option is enabled.
Threshold	Displays the number of days the inactive agents are retained in the Endpoint Server. By default, the threshold value is 90 days.

Fields	Description
Run	Displays the schedule for running the inactive agents retention job. By default, the database check occurs everyday at 00:00:00 AM. You can select the frequency from the drop-down list (Everyday, Weekdays, Weekends, or Custom, where Custom allows you to select one or more specific days of the week) and time to run the job.
Apply	Overwrites any previous schedule for the inactive agents retention policy and applies the new settings immediately.
Reset	Resets the schedule to the default settings.

The following table lists the fields for memory dumps and MFT retention policy.

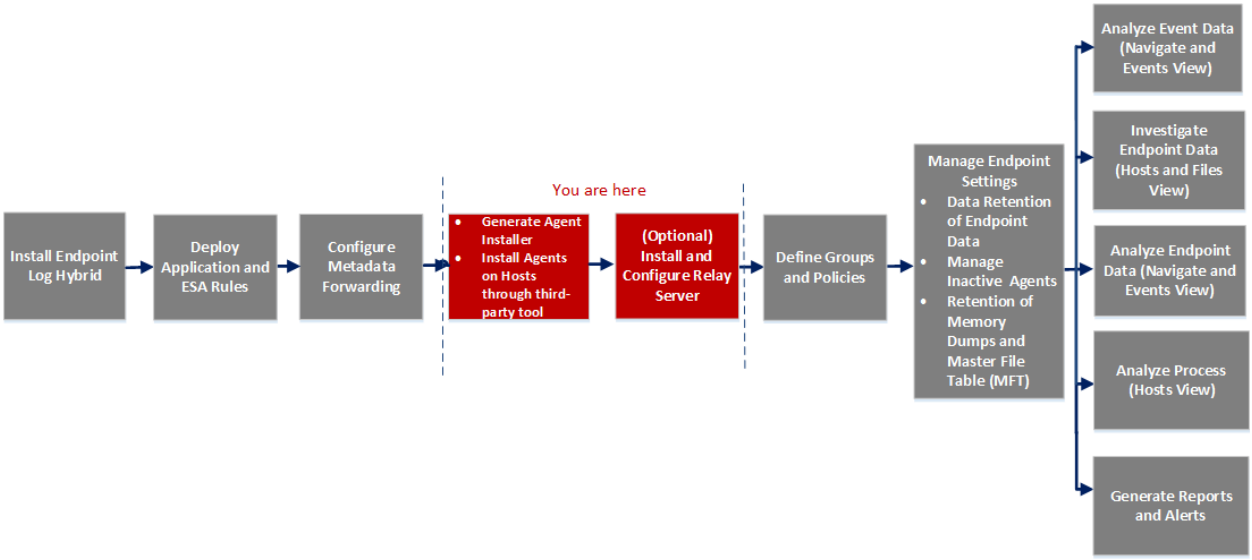
Fields	Description
Enable	Enables the configuration for the memory dumps and MFT policy. By default, this option is enabled.
Threshold	Displays the number of days the memory dumps and MFT are retained in the Endpoint Server. By default, the Threshold is set to 90 days. The data older than 90 days is deleted from the Endpoint Server.
Run	Displays the schedule for running the retention job. By default, the database check occurs everyday at 00:00:00 AM. You can select the frequency from the drop-down list (Everyday, Weekdays, Weekends, or Custom, where Custom allows you to select one or more specific days of the week) and time to run the job.
Apply	Overwrites any previous schedule for the retention policy and applies the new schedule immediately.
Reset	Resets the schedule to the default settings.

Agent Packager Tab

In the **Agent Packager** tab, you can generate an agent packager and agent installer. To access this view:

1. Go to **(Admin) > Services**.
2. In the Services view, select **Endpoint Server**.
3. Click and select > **View** > **Config**.
4. Click the **Packager** tab.

Workflow



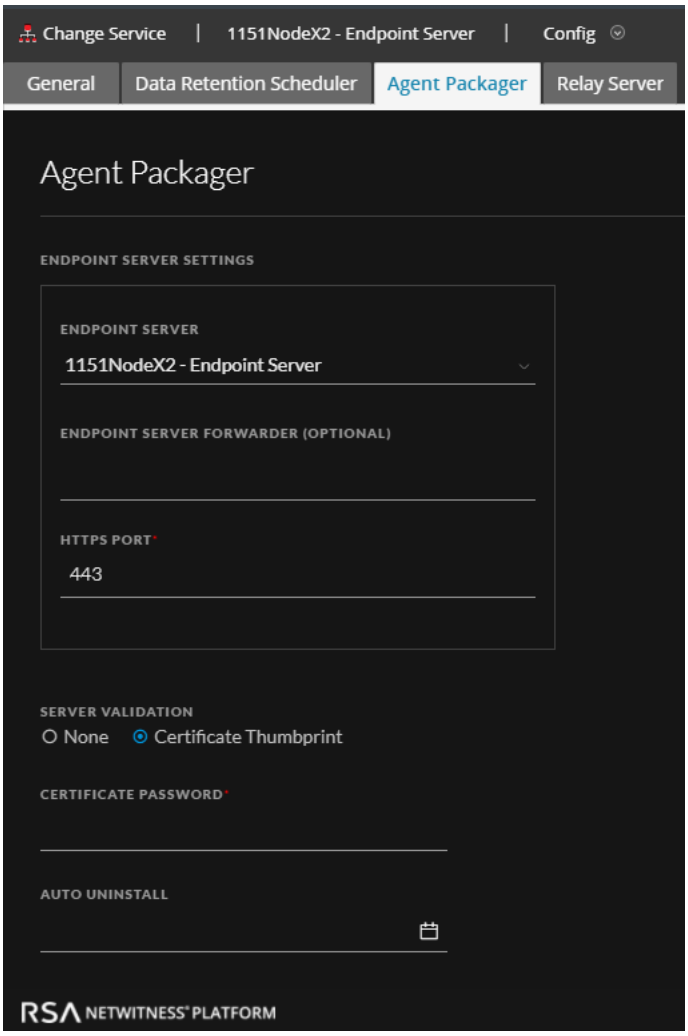
What do you want to do?

User Role	I want to ...	Show me how
Administrator	Generate an Agent Packager for Endpoint Data Collection*	*NetWitness Endpoint Agent Installation Guide*
Administrator	Generate an Agent Installer*	
Administrator	Install and Configure Relay Server*	*NetWitness Endpoint Configuration Guide*

*You can perform this task in the current view

Quick Look

The following figure is an example of the Agent Packager tab.


Change Service | 1151NodeX2 - Endpoint Server | Config
General
Data Retention Scheduler
Agent Packager
Relay Server
Agent Packager
ENDPOINT SERVER SETTINGS
ENDPOINT SERVER
1151NodeX2 - Endpoint Server
ENDPOINT SERVER FORWARDER (OPTIONAL)
HTTPS PORT*
443
SERVER VALIDATION
None
Certificate Thumbprint
CERTIFICATE PASSWORD*
AUTO UNINSTALL
RSA NETWITNESS PLATFORM

Relay Server Tab

In the **Relay Server** tab, you can download and configure Relay Server . To access this view:

1. Go to **ADMIN** > **Services**.
2. In the Services view, select **Endpoint Server**.
3. Click  and select > **View** > **Config**.
4. Click the **Relay Server** tab.

Workflow



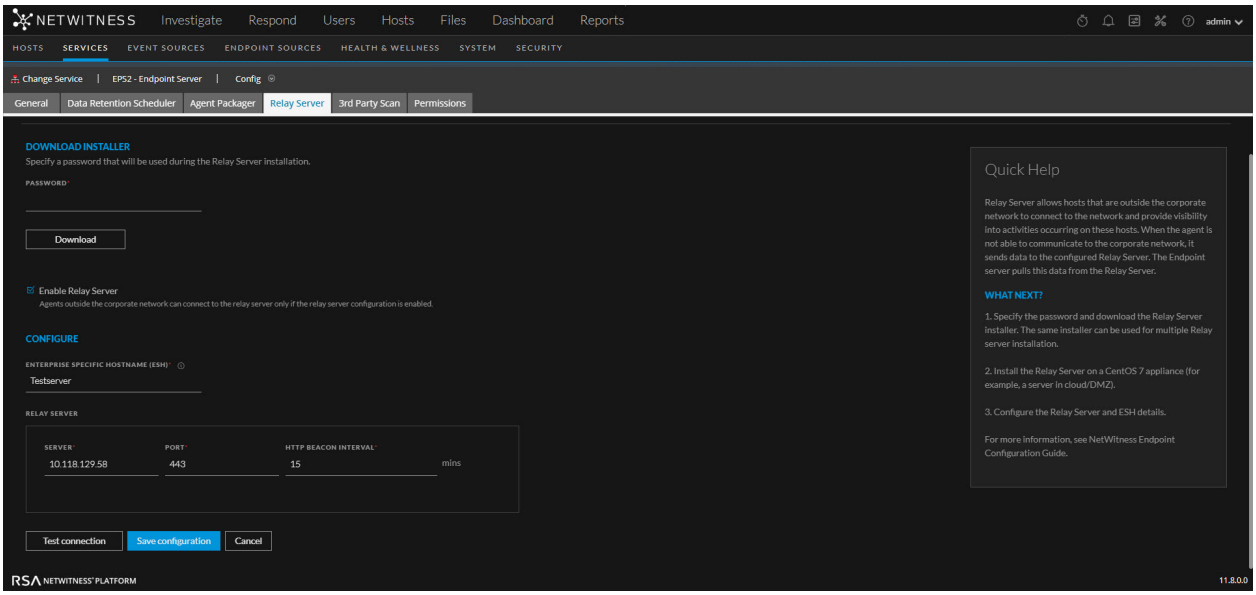
What do you want to do?

User Role	I want to ...	Show me how
Administrator	Install Relay Server*	(Optional) Installing and Configuring Relay Server
Administrator	Configure Relay Server*	(Optional) Installing and Configuring Relay Server

* You can perform this task in the current view.

Quick Look

The following figure is an example of the Relay Server tab.



Features

The following table lists the fields for Relay Server tab.

Field	Description
Download Installer	
Password	Enter the relay server installation password. For example, netwitness. Password must be minimum of 3 characters and can contain alphanumeric and special characters. **Note:** You must provide the same password when prompted during Relay Server installation.
Download	Click to download the Relay Server installer.
Enterprise Specific Hostname (ESH)	Enter the hostname which can be resolved only within the corporate network.
Enable Relay Server	Check Enable Relay Server for Agents outside the corporate network to connect to the configured relay server. By default, this option is disabled.
Configure	
Server	Hostname or IP of the Relay Server.
Port	Port number. For example, 443.
HTTP Beacon Interval	Enter the Interval value ranges from 60 – 1440 minutes.
Test Connection	Click test connection to check if the relay server is reachable by the agents.

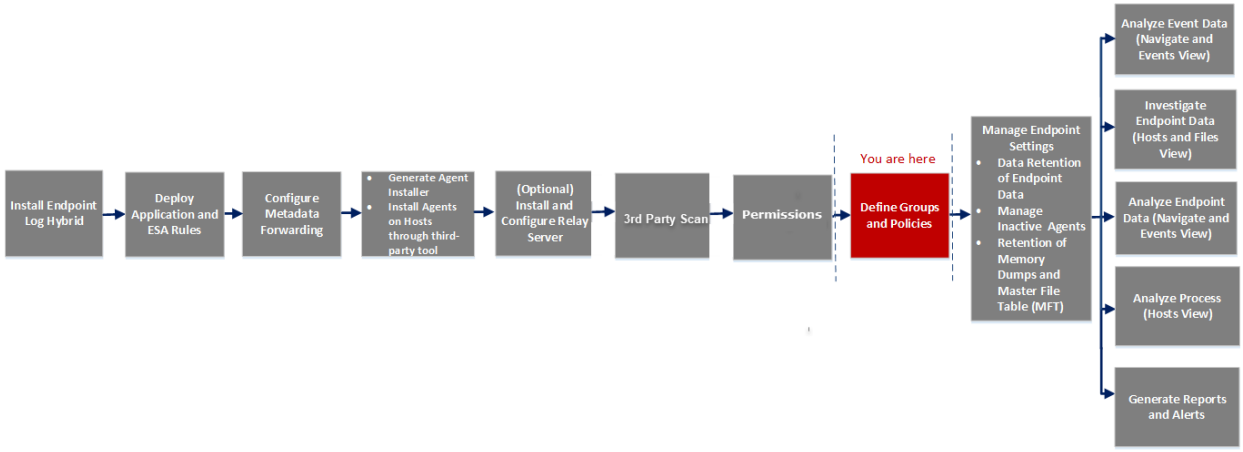
Field	Description
Save Configuration	Saves the relay server configuration.
Cancel	Cancel the unsaved changes.

Endpoint Sources - Groups

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The **(Admin) > Endpoint Sources** view contains two tabs: **Groups** and **Policies**.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Administrator	create new groups*	Create a Group
Administrator	edit groups*	Edit a Group
Administrator	edit ranking*	Managing Groups
Administrator	delete groups*	Delete a Group
Administrator	view default policies	Default Agent Endpoint (EDR) Policy
Administrator	create an EDR policy	Create an EDR Policy
Administrator	create a Windows Log policy	Create a Windows Log Policy
Administrator	edit policies	Edit a Policy
Administrator	delete policies	Delete a Policy

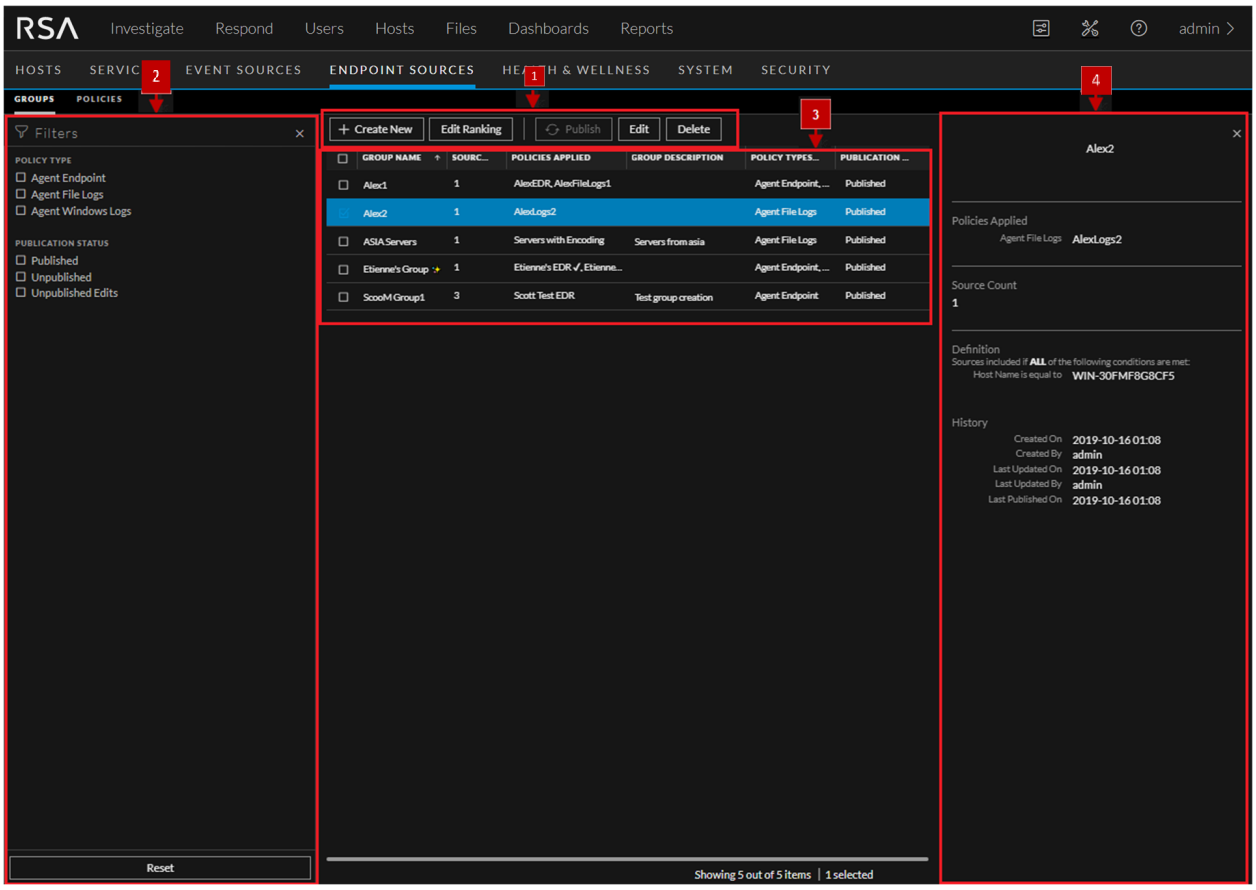
*You can perform this task in the current view.

Related Topics

- Endpoint Sources
- Managing Policies

Quick Look

Below is an example of the Groups tab:



1 **Toolbar**

- **Create New**: Lets you create a new group. For more information, see Create a Group
- **Edit Ranking**: Lets you edit the ranking of groups. For more information, see Managing Groups
- **Publish**: Publishes the selected group or groups.
- **Edit**: Lets you edit the details of an existing group. For more information, see Edit a Group.
- **Delete**: Deletes the selected group or groups permanently. For more information, see Delete a Group.

2 **Filter Pane**

- **Filters**: You can filter groups based on Policy Type and Publication Status.

  To hide, click the  icon at the top-right of the panel. To display if hidden, click the  icon in the toolbar.

- **Reset**: Removes the currently applied filter criteria.

For more information, see Filter Endpoint Groups.

3 **Groups List Pane**

- **Group name**: Name of the group.
- **Source Count**: Number of hosts that are currently members of the group.
- **Policies applied**: Lists the policies applied to this group.
- **Group description**: Description of the group.
- **Policy Types Applied:** Type of policies applied to the group: Agent Endpoint, Agent File Logs, Agent Windows Logs, or any combination of these.
- **Publication Status**: Status of the group - Published or Unpublished.

You can also sort on any column. If you mouse over a column header, a sort icon is displayed: . Click the icon to sort by the selected column.

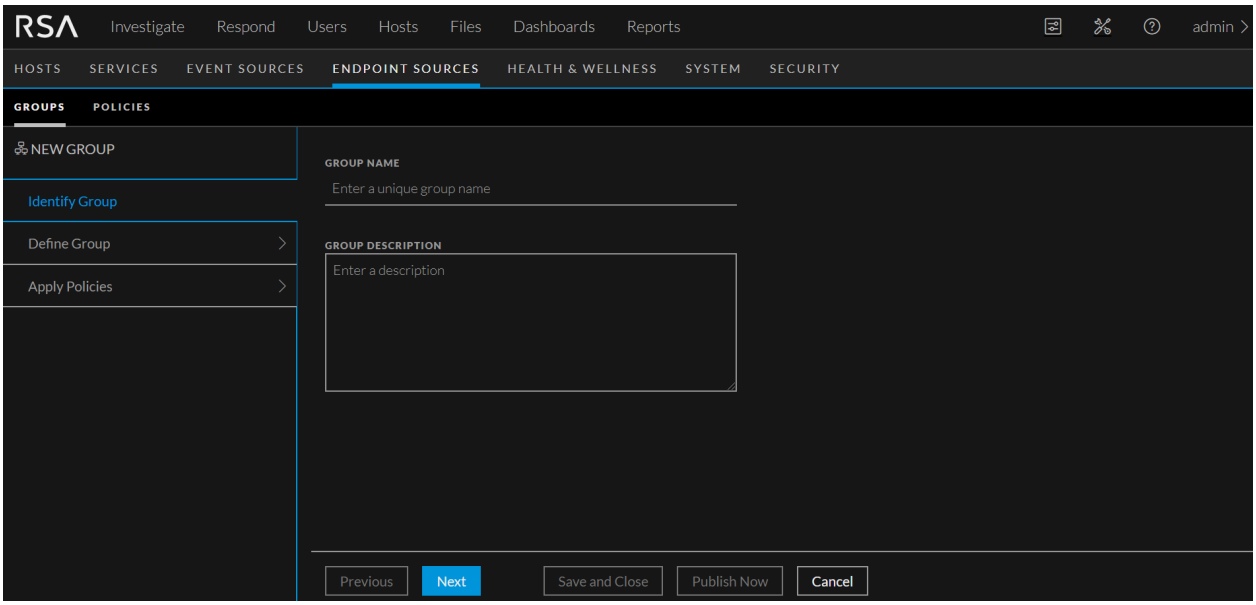
4 **Group Details Pane**

Displays the properties of the selected group.

Note: Click the row to view the Properties panel for a group.

Create Group

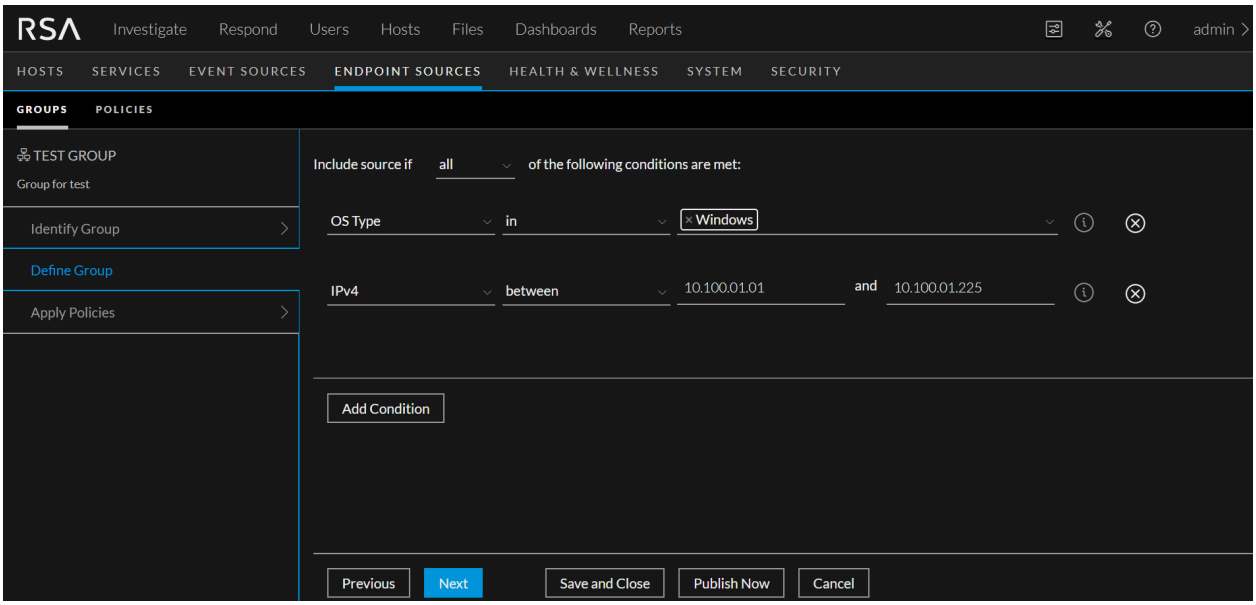
Below is an example of the Create Group dialog. The table describes the information and options in the Create Group dialog.



Field	Description
Group Name	Name of the group. The name should be unique.
Group Description	Description of the group and should not exceed 8000 characters.

Define Group

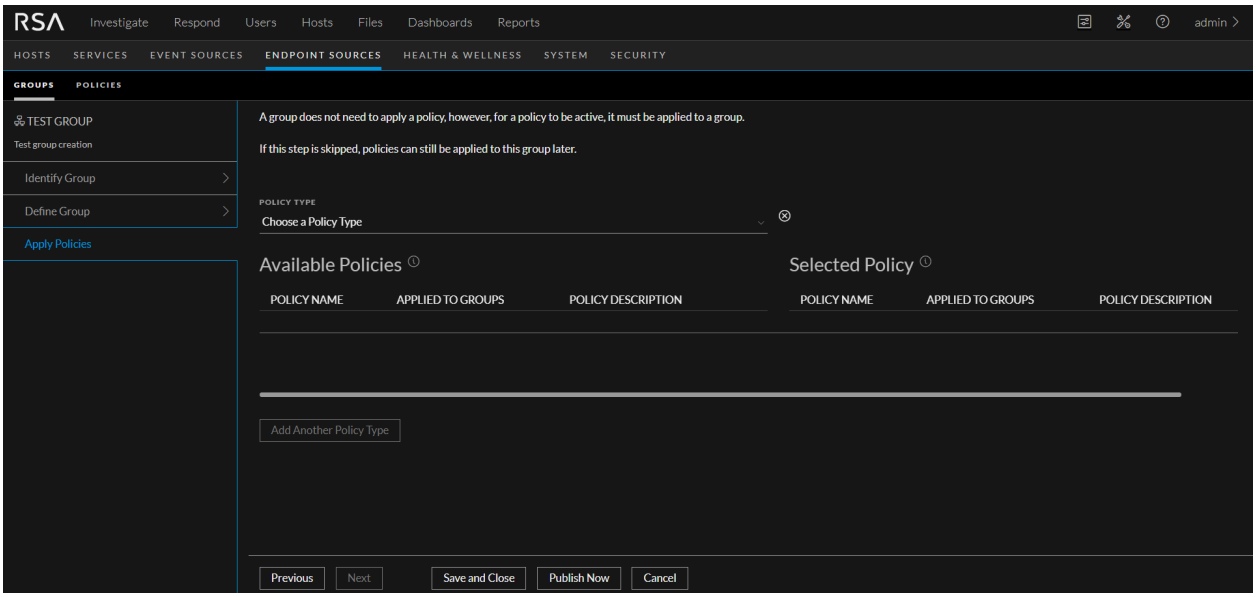
Below is an example of Define Group panel. The table describes the information and options in the Define Group panel:



Field	Description
Include source if ...of the conditions are met	Defines the conditions for an agent to be included in the group. Available options are all or any.
Parameter	The parameter can be OS Type, OS Description, Host Name, IPv4, IPv6 or Machine OU. • OS Type - Type of operating system. Available options are: Windows, Linux, and MacOS. • OS Description - Description of the operating system. The description should not exceed 256 characters. Available operators are: is equal to, contains, start with, and ends with. For example, Microsoft Windows 10 Enterprise. • Host name - Name of the host. The host name can contain only alphanumeric characters. Available operators are: is equal to, contains, start with, ends with, and in. For example, DESKTOP-QQPDNG3. • IPv4 and IPv6 - IP address. Available operators are: between, in, not in, and between. For example, 10.40.15.220. • Machine OU - Name of the Machine OU. Available operators are: is equal to, contains. For example, OU=Win10, OU=Contractors, DC=corp, DC=com **Note:** If you do not want to include certain IP addresses, use the **Not in** operator, and enter the IP address separated by a space or a comma.
Operator	The choice of values is dependent upon the parameter you chose. For example, if your parameter is OS Type, the only operator available is **in**.
Value or values to match	The value or values to match. For the OS Type parameter, you can choose one or more values from the drop-down list. For all other parameters, you can enter free-form text. **Note:** Although you can enter any text for values, the system validates your entries when you attempt to proceed to another screen, and will not allow you to proceed until values are valid.
Add condition	Lets you add another condition.

Apply Policies

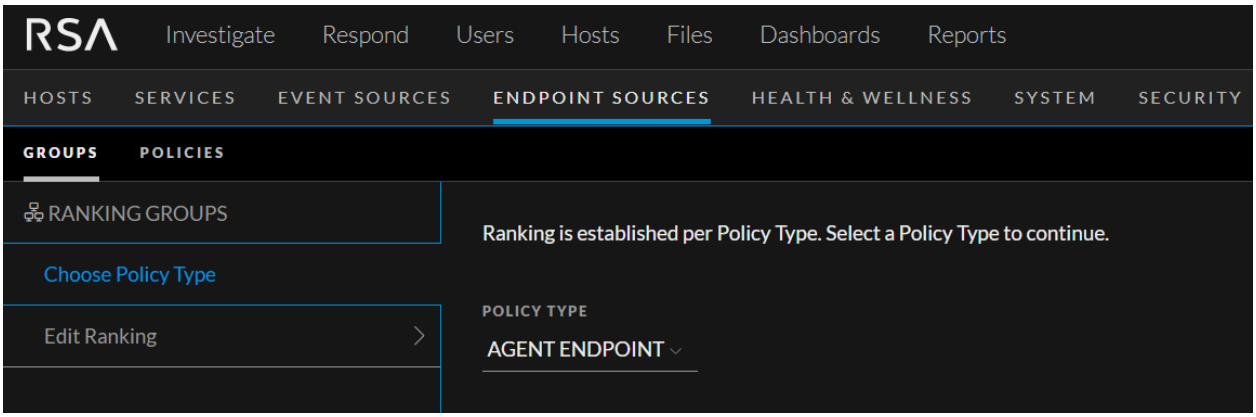
Below is an example of Apply Policies panel. The table describes the information and options in the Apply Policies panel:



Field	Description
Source Type	Defines the source type for the group. Available options are Agent Endpoint and Agent Windows Logs.
Available Policies	List the available policies associated with the source type.
Selected Policies	List the policies selected.
Add Another Source Type	Lets you add another source type.
Save and Close	Saves the settings and closes the Create Group dialog.
Publish Now	Publishes the created group.

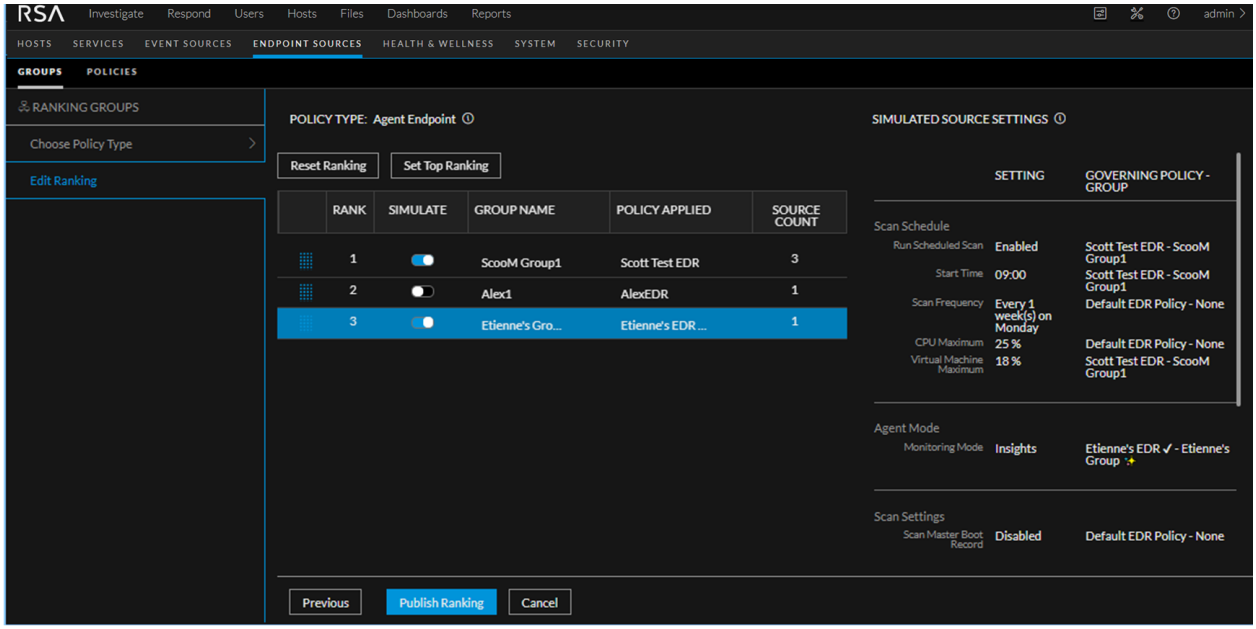
Ranking Groups

Below is an example of the Ranking Groups dialog. The table describes the information and options in the Ranking Groups dialog.



Field	Description
Source Type	Establishes ranking for the source type. Available options are Agent Endpoint and Agent Windows Logs.

Below is an example of the Edit Ranking panel.



From this panel, you can do the following:

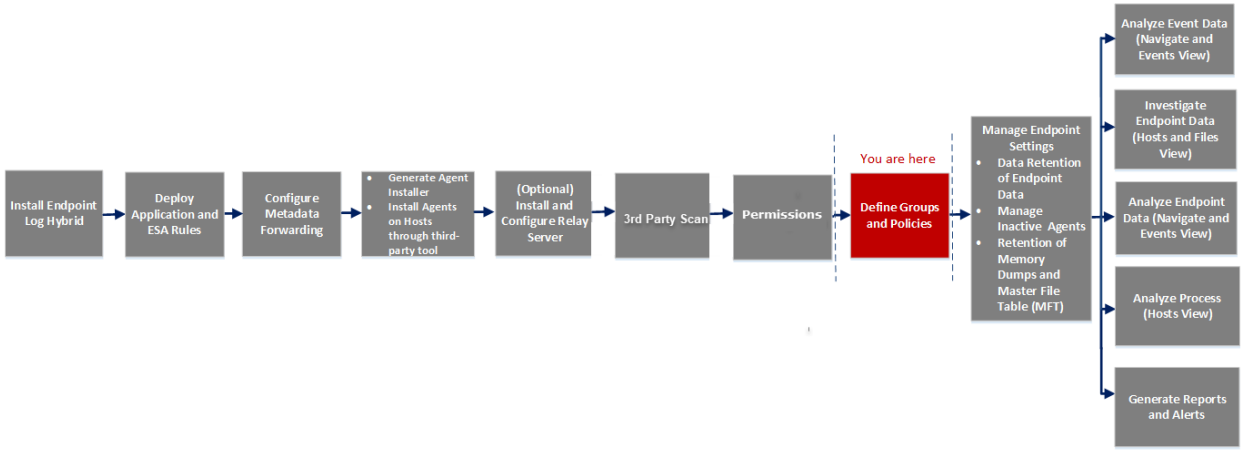
- Drag the group up or down to change the priority. Priority decreases from top to bottom.
- Turn the Simulate slider on or off, to simulate your policy settings and how they affect the endpoints within their groups. For more details, see Simulation Examples.
- Use the available buttons to perform actions:
  - **Reset Ranking**: Resets the ranking to the original order.
  - **Set Top Ranking**: Moves the selected group to the top.
  - **Previous**: Navigates to the Choose Source Type panel.
  - **Publish Ranking**: Lets you edit the details of an existing group. For more information, see Edit a Group.
  - **Cancel**: Discards the changes and returns to the Groups tab.

Endpoint Sources - Policies

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The **(Admin) > Endpoint Sources** view contains two tabs: **Groups** and **Policies**.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Administrator	create new groups	Create a Group
Administrator	edit groups	Edit a Group
Administrator	edit ranking	Managing Groups
Administrator	delete groups	Delete a Group
Administrator	view default policies*	Default Agent Endpoint (EDR) Policy
Administrator	create an EDR policy*	Create an EDR Policy
Administrator	create a Windows Log policy*	Create a Windows Log Policy
Administrator	create a File Log policy*	Create a File Log Policy
Administrator	edit policies*	Edit a Policy
Administrator	delete policies*	Delete a Policy

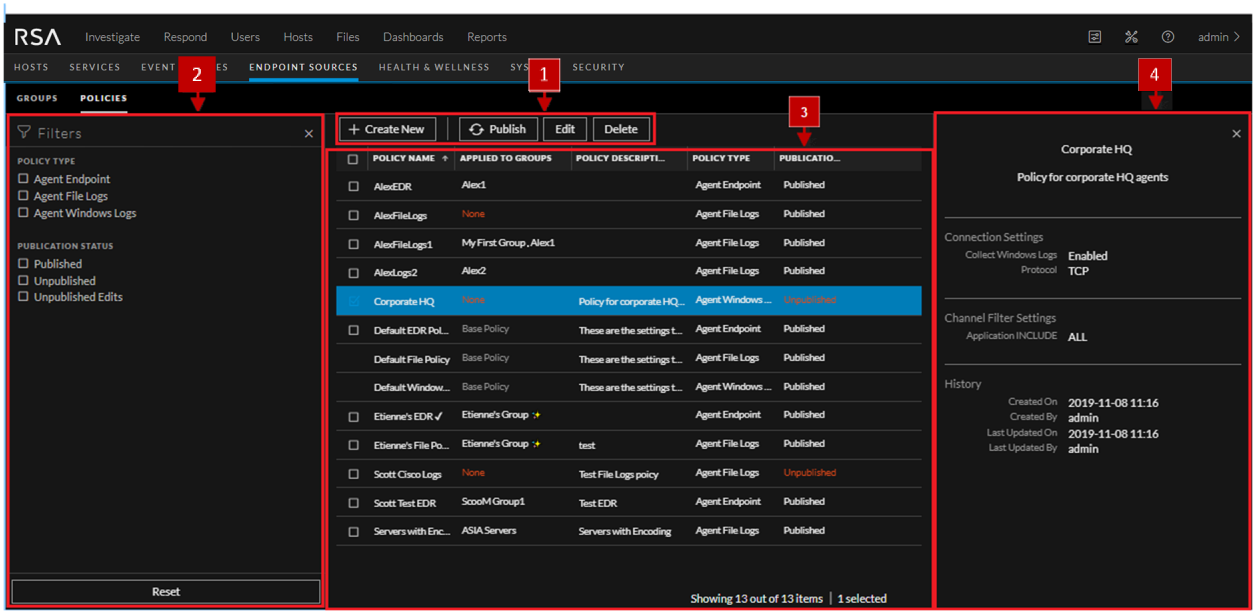
*You can perform this task in the current view

Related Topics

- Endpoint Sources
- Managing Groups

Quick Look

Below is an example of the Policies tab:



1 **Toolbar**

- **Create New**: Lets you create a new policy. For more information, see Managing Policies.
- **Publish**: Publishes the selected policy.
- **Edit**: Lets you edit the details of an existing policy. For more information, see Edit a Policy.
- **Delete**: Deletes the selected policies permanently. For more information, see Delete a Policy.

2 **Filter Panel**

- **Filters**: You can filter policies based on Policy Type and Publication Status.

  To hide, click the icon at the top-right of the panel. To display if hidden, click the icon in the toolbar.

- **Reset**: Removes the currently applied filter criteria.

For more information, see Filter Policies.

3 **Policies List Panel**

Policy View. Displays the policy details:

- **Policy name**: Name of the policy.
- **Applied to groups**: Lists the group to which this policy is applied.
- **Policy description**: Displays the first portion of the description.
- **Policy type**: Displays the policy type: Agent Endpoint, Agent File Logs, or Agent Windows Logs.
- **Publication Status**: Status of the policy: Published or Unpublished.

You can also sort on any column. If you mouse over a column header, a sort icon is displayed: . Click the icon to sort by the selected column.

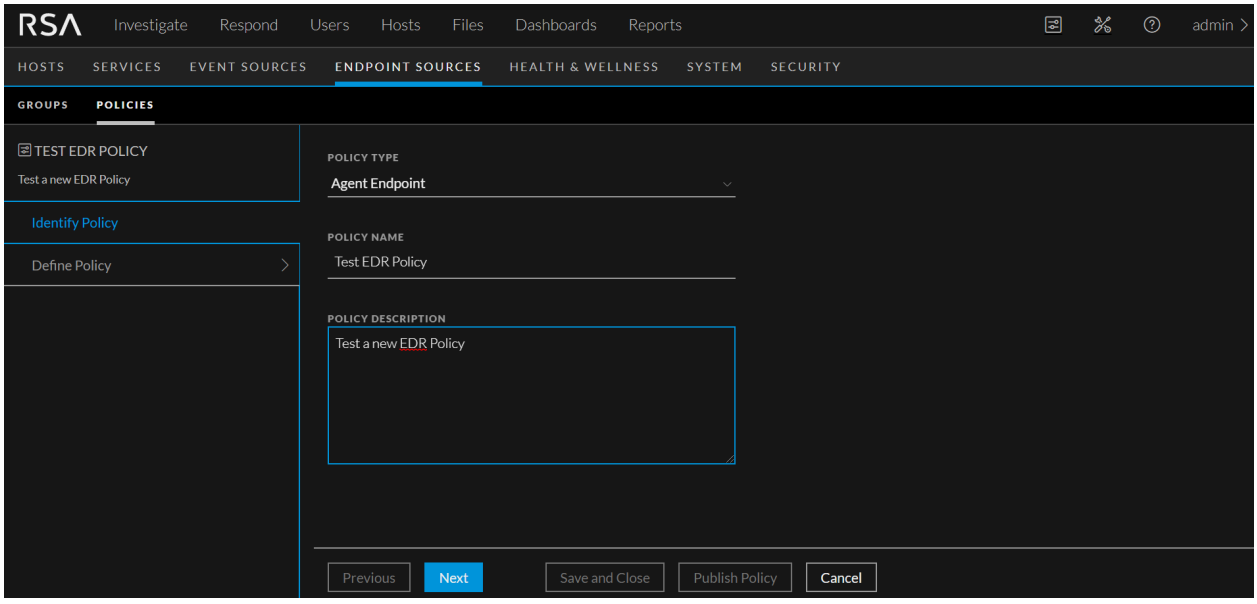
4 **Policy Details Panel**

Displays the properties of the selected policy.

Note: To view the Properties panel for a policy, click the Policy Name.

Create Policy

Below is an example of the Create Policy dialog. The table describes the information and options in the Create Policy dialog.



Field	Description
Policy Type	Displays the type for the policy. Available options are Agent Endpoint, Agent File Logs, and Agent Windows Logs.

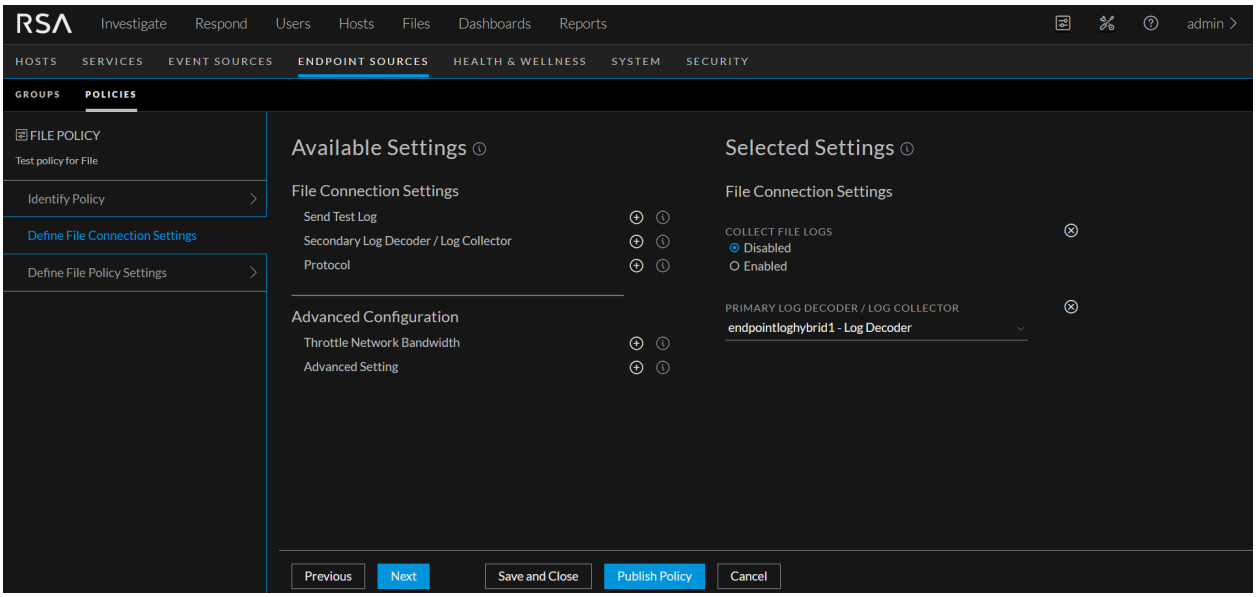
Field	Description
Policy Name	Name of the policy. The name should be unique.
Policy Description	Description of the policy. Description should not exceed 8000 characters.

Panels for Log File Policy

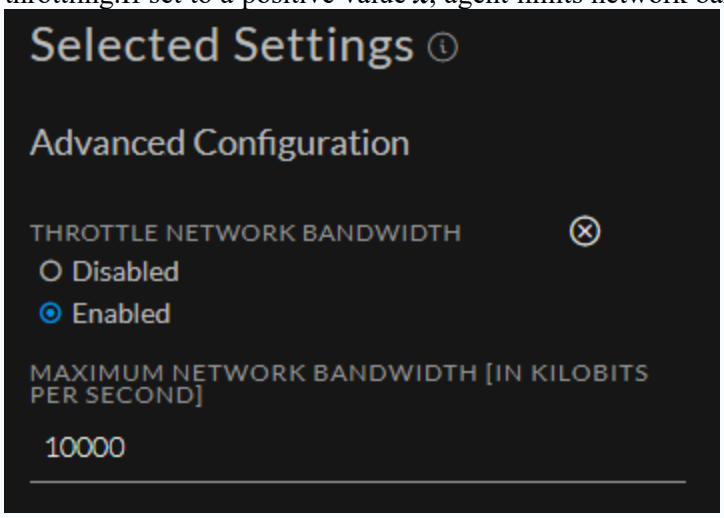
There are two panels for defining the parameters for an Agent Log File Policy: **Define Connection Settings** and **Define File Policy Settings**.

Define Connection Settings

Below is an example of Define Connection Settings panel. The table describes the information and available options.

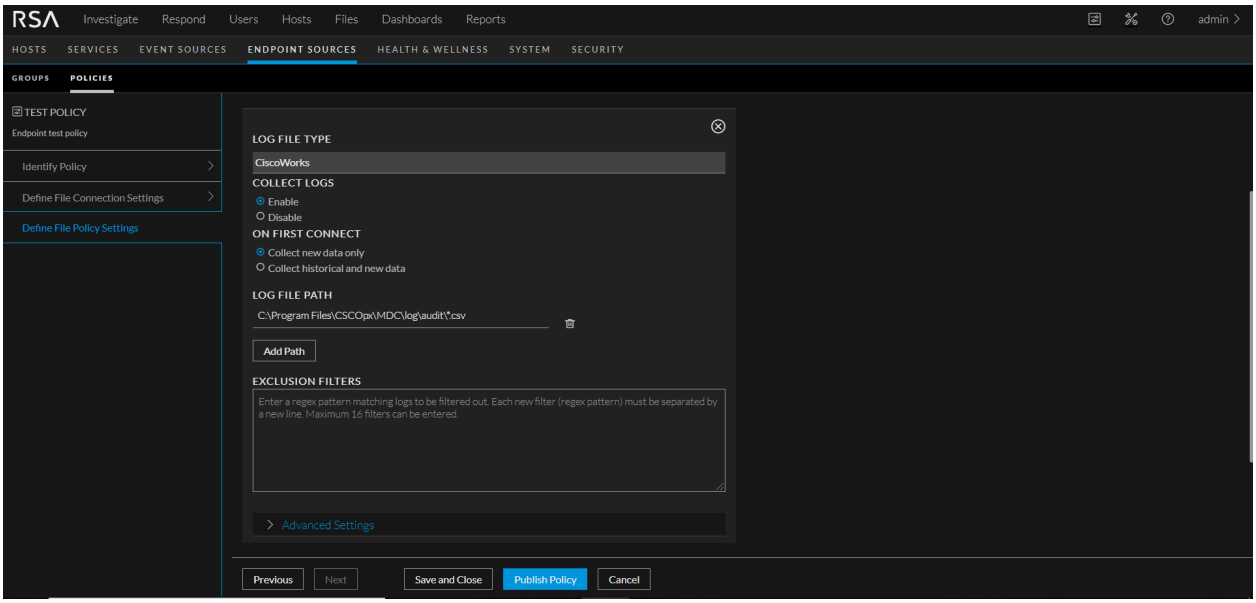


Field	Description
Collect File Logs	If enabled, the log file collection capability of the agent is activated. Logs are collected and forwarded to the RSA NetWitness as they are generated. If disabled, no defined event source logs are collected. **Note:** This option must be enabled for any file event sources to be collected.
Send Test Log	If enabled, a sample log is sent to the configured server when the policy is loaded to test connectivity. This allows to test the configuration before standard logs are available. By default, this option is disabled.

Field	Description
Primary Log Decoder / Log Collector	The primary Log Decoder or Log Collector to which the collected file logs will be forwarded.
Secondary Log Decoder / Log Collector	If the primary Log Decoder or Log Collector is not reachable, collected file logs are forwarded to the secondary Log Decoder or Log Collector. **Note:** The NetWitness cannot detect failures when UDP protocol is used.
Protocol	Select the transport protocol that is used to forward the collected file logs to the NetWitness servers. The following options are available: • SSL: Recommended, but also the most resource-intensive option. • TCP: Sends the logs in clear text over a reliable TCP connection. May be acceptable within a corporate network. • UDP: Sends the log in clear text over a non-guaranteed UDP connection. This is the least resource intensive option. **Note:** Resource intensity is dependent upon the Log Decoder, since there is only a single connection per agent.
Advanced Configuration	
Throttle Network Bandwidth	Use this setting to limit network bandwidth that the Agent uses to connect to NetWitness. This setting is disabled by default: click **Enabled** to turn it on, and then enter a value in kilobits per second.If not set, Agent does not do any network throttling.If set to a positive value ***x***, agent limits network bandwidth to ***x*** kbps. 

Field	Description
Advanced Setting	**Caution:** It is strongly recommended not to use this setting unless advised to do so by RSA.

Define File Policy Settings



Field	Description
Log File Type	From the drop-down menu, select the type of event source to be monitored. The list of available event source is based on all the event source types defined on your RSA NetWitness. You can add event source types using the Live Services module. For details, see "Find and Deploy Live Resources" in the *Live Services Management Guide*.
Collect Logs	If enabled, log files for this file type instance are collected and forwarded to the NetWitness. File collection must be enabled on each source applying this policy for these specific logs to be collected.
On First Connect	Determines whether the NetWitness Agent collects all logs or only newly created logs located in the specified paths upon initial collection. In both cases, new logs are collected. **Note:** Historical logs cannot be collected after an agent has begun collecting logs.

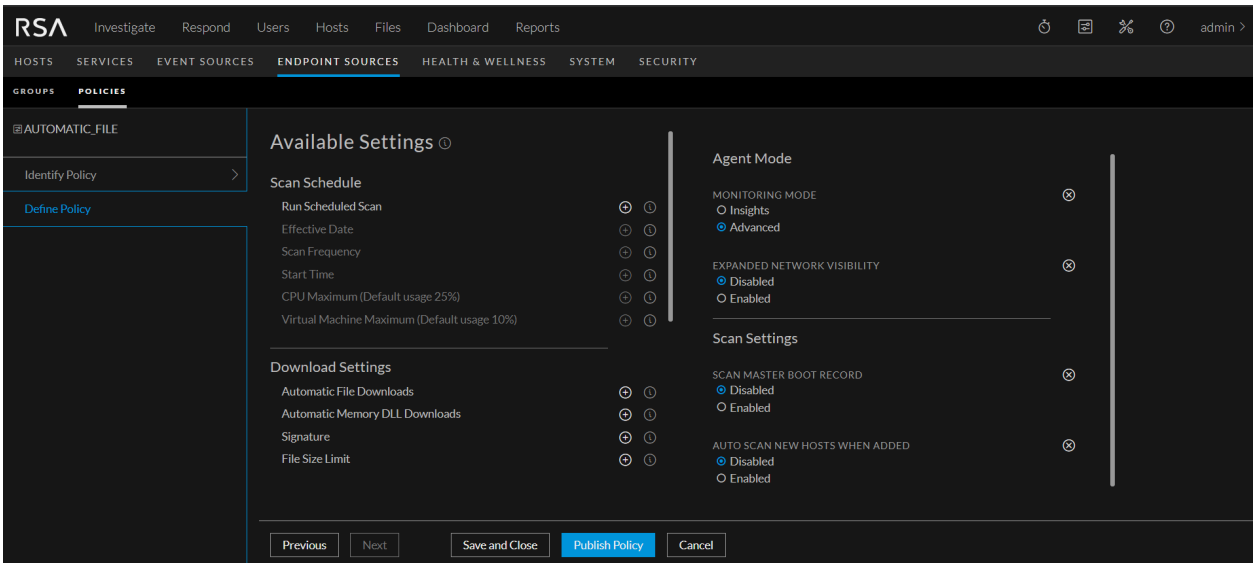
Field	Description
Log File Path	One or more paths to be used by the agent to locate the log files. Represents the location of the log files to be read. **Note:** The Path value cannot end at a directory—the final portion of the path must represent a file name or set of files (using wildcard characters). You can use wildcards for both files and directories. Each source is limited to entry of 16 paths. This setting must include a path and a file spec. For example: `C:\Program Files\apache-tomcat-*\logs\*.log`. In this case, the file spec is all files with a ".log" extension in the specified path. If you cannot use wildcards to specify multiple files, you can add additional paths to accommodate the differences in path locations on a specific endpoint agent. This might be due to installation locations or version information. Only the paths with valid locations and files on the specific endpoint agent are used, and the others are ignored. For many event source types, there is a default path. If so, you only need to enter a path if the log files are not stored in the standard directory for that event source type. **Note:** This can be a standard Windows pathname (such as `C:\Program Files\Apache\error_logs\logfile.log`) or a UNC (Universal Naming Convention) pathname (`\\host-name\share-name\file-path`). For more details about UNC paths, see Endpoint Sources - Policies below.
Exclusion Filters	An optional list of regex patterns which can be used to filter out any logs that match the patterns. Each separate filter should be entered on a new line. Each source is limited to 16 exclusion filters. **Note:** Each filter needs to be entered as a valid regex string, or the system does not allow you to save it.
Advanced Settings	

Field	Description
Source Alias	Optionally, enter a hostname, IPv4 or IPv6 address to identify individual sources. This is recommended when there are two or more sources of the same type on the same server: For example, a server that runs two instances of Apache web server. **Note:** This value only rarely needs to be entered. One example is if you have more than one Web Server, and they are running different Apache servers. Note the following: • If you enter a value for this parameter, the event source is applicable to a single Endpoint server. • This optional address or hostname is included in the meta for any logs originating from this source. This can be used by analysts to assist in identifying the source. • Set a value for this parameter if two sources of the same event source type are configured in the same policy. • This setting is not commonly needed: it is only useful if the policy is only applied to a single endpoint.
File Encoding	Specifies the type of character encoding of the log files. If **Local Encoding** is selected, the NetWitness Agent uses the default encoding of the Windows machine upon which it is running. This setting must match the encoding of the log files, or they will not be processed correctly. **Note:** UTF-8/ASCII is recommended (and the default). UTF-8 is a super-set of ASCII Note that all logs are re-encoded to UTF-8 before being sent to the NetWitness.

For a list of the currently supported types, see Currently Supported File Log Event Source Types.

Define Policy Panel for Agent Endpoint Policy

Below is an example of Define policy panel. The table describes the information and options for Agent Endpoint policy:



Settings	Description
	Scan Schedule
Run Scheduled Scan	Run a scheduled scan if you want to receive regular snapshots from a host. Scan snapshots provide detailed information about processes and files loaded on the memory. By default, this option is disabled. You can also run a manual scan from the **Hosts** view. **Note:** The following scan schedule options are available only when the scan schedule is enabled. The values entered are specific to the agent time zone.
Effective Date	Date when the policy takes effect. If you do not want this policy to take effect as soon as it is applied to a group and published, set an effective date that is in the future. By default, this is set to the current date.
Scan Frequency	Determines how often the scheduled scan runs on a host. By default, this is set to every week. Every network is different and the frequency should balance the needs of the analysts for current data, availability to review the data, and how systems deal with the load of the generated data. Select Days or Weeks: • Days: Select the number of days of the scan frequency. You can set a schedule to scan every n days, where n is 1, 2, 3, 4, 5, 6, 10, 15, or 20. For example, to scan every third day, select 3. • Weeks: Select after how many weeks the policy scan should be initiated and on which day of the week the policy scan should initiate. For example, to scan every other Wednesday, choose 2 and W.
Start Time	Time when the scheduled scan starts to run on a host. By default, this is set to 9:00. This is the local host time, meaning that scans across a global network will not run all at once. Note that the time is in 24 hour format. To set a time of 7:30 PM, select 19:30.

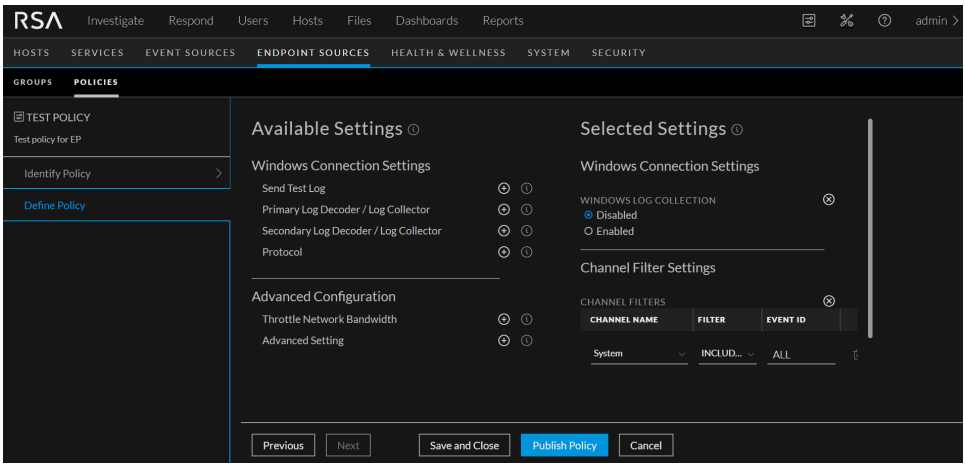
Settings	Description
CPU Maximum	Amount of CPU the agent can use to run scheduled scans on physical hosts. By default, the value is set at 25%. Increasing the CPU maximum increases the speed of scan snapshot retrieval. Drag the slider to specify the maximum CPU usage by the created policy. Minimum value is 5%. Use the slider to select the maximum CPU processing power to use for the scan. Note that the higher the percentage, the less CPU is available for other tasks on the host.
Virtual Machine Maximum	Amount of CPU the agent can use to run scheduled scans on virtual machines. By default, the value is set at 10%. Increasing the virtual machine maximum value increases the speed of scan snapshot retrieval. Drag the slider to specify the maximum Virtual Machine usage by the policy. Minimum value is 5%. Use the slider to select the maximum CPU processing power to use for the scan. Keep in mind that the higher the percentage, the less CPU is available for other tasks running on the virtual machine.
Agent Mode	
Monitoring mode	Allows you to specify whether an agent should operate in Insights (free) or Advanced mode (license). By default, it is set to Advanced.
(Version 11.5 and Later) Expanded Network Visibility	Enables network tracking and monitoring on Windows hosts in Insights mode. **Note:** For network tracking in Insights mode, verify that the Windows Management Instrumentation (WMI) service should be enabled. It also optimizes the frequency of agents sending network events for network packet correlation for both Insights and Advanced modes. By default, this option is disabled.
Scan Settings	
Scan Master Boot Record	Includes Master Boot Record (MBR) details in scheduled scans. By default, this option is disabled. This can help to identify when an operating system boot sequence is compromised. However, not all modifications to the MBR are malicious, as they could be made to provide encryption or enforce licensing of certain legitimate software.
Auto Scan New Systems When Added	Automatically scans when a new host is added. By default, this option is disabled. If this option is disabled, no snapshot data is displayed in the **Hosts** view until a manual or scheduled scan is run on these hosts. Existing hosts will not be affected. **Note:** Enabling this option on a new deployment when this policy is applied to a large number of hosts may result in a large number of simultaneous scans that cause performance degradation.
Download Settings	
Automatic File Download	Automatically download the files to the NetWitness Endpoint Server based on the file size and signature. If a file is present on multiple hosts or multiple Endpoint Servers, only one instance of the file is downloaded. By default this option is enabled.

Settings	Description
Automatic Memory DLL Downloads	From version 11.6.1 and higher, all memory DLLs that are detected during a scan, will be automatically downloaded regardless of the file size. This option is enabled by default.
Signature	Limits the download of files based on the signature. The options are : • Exclude All Signed - Downloads all the unsigned files to the NetWitness Endpoint Server and exclude all the signed files. • Exclude only Microsoft and Apple signed - Downloads all the unsigned files and exclude the files signed by Microsoft and Apple. • Include All- Downloads all the signed and unsigned files. **Note:** . In case of Linux, **Exclude all signed** and Exclude **Microsoft and Apple** signed options will download the files that are not part of any installed RPMs or files which are part of RPM but the hashes does not match with RPM.
File Size Limit	Limits the download of files based on the file size. The File size should be between 1 KB and 10 MB. By default, file size lesser than or equal to 1 MB are downloaded automatically.
Response Action Settings	
Blocking	Allows an analyst to prevent the execution of a malicious file on any host running an Advanced mode agent. By default, this option is disabled. File blocking will not be enforced if it is disabled by policy, which might be desirable to ensure that there are no performance side effects on systems where CPU or IO performance is critical. **Note:** Blocking is only supported on Windows agent (in Advanced mode) with NetWitness version 11.3 and later.
Network Isolation	Allows an analyst to block hosts that are compromised from connecting to the network. This controls the spread of an attack and help analyze the malware behavior after the network isolation. All attempted network connections are monitored and reported to the Endpoint Server. By default, this option is disabled. **Note:** Network isolation is only supported on Windows agent (in Advanced mode) with NetWitness version 11.4 and later.
Endpoint Server Setting	
Endpoint Server	Displays all available Endpoint servers in the deployed. **Note:** If you do not select an Endpoint Server, the agent uses the default Endpoint Server that is configured during packager generation.
Server Alias (Optional)	The optional server alias allows you to enter an alternative hostname or IP address on which the server can be reached in the case that agents need to go through a NAT or similar in order to reach the Endpoint Server.

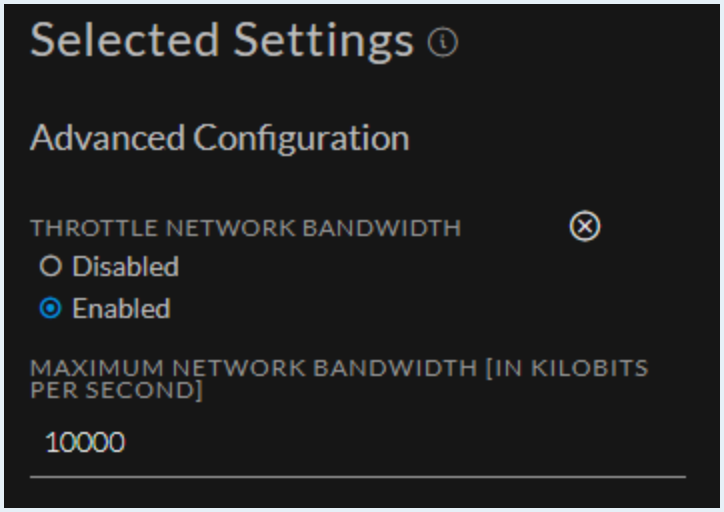
Settings	Description
HTTPS Port	Port number used for HTTPS communication. By default, the port is set to 443. If you want to change this port, make sure that it matches the server configuration. If you enter the wrong port, the agents can no longer communicate with the Endpoint server and the system will be non-functional.
HTTPS Beacon Interval	Determines how often an agent can communicate with the Endpoint server over HTTPS. By default, the value is set to 15 minutes. The default method of beaconing is UDP. Beaconing is used as a method of keep-alive to know if a host is online and to allow hosts to respond faster than the fallback HTTPS beacon time.
UDP Port	Port number used for UDP communication. By default, the port is set to 444. If you want to change this port, make sure that it matches the server configuration. Entering the wrong port results in loss of functionality and effects performance.
UDP Beacon Interval	Determines how often an agent can communicate with the Endpoint server over UDP. By default, the value is set to 30 seconds.

Define Policy Panel for Windows Logs Policy

The table describes the information and options for Agent Windows Logs policy:



Settings	Description
Windows Log Collection	If enabled, logs from the Windows hosts are collected and forwarded to the NetWitness Platform. By default, this option is disabled.
Send Test Log	If enabled, a sample log is sent to the configured server when the policy is loaded to test connectivity. This allows to test the configuration before standard logs are available. By default, this option is disabled.

Settings	Description
Primary Log Decoder / Log collector	Primary NetWitness Log Decoder or Log Collector to which the collected Windows logs are forwarded.
(Optional) Secondary Log Decoder / Log collector	If the primary Log Decoder or Log Collector is not reachable, the collected Windows logs are forwarded to the secondary Log Decoder or Log Collector. **Note:** NetWitness cannot detect failures when UDP protocol is used.
Protocol	Select whether TLS, TCP, or UDP transport protocol is used to forward the collected Windows logs to the NetWitness Platform servers. By default, the protocol is TCP.
Channel Filters	Configure which Windows Log events to collect by selecting a channel, filter condition, and the relevant event IDs. You can either select common channels, such as Security or System from the drop-down list, or create custom channels by entering the channel name. By default, all events are collected from a selected channel. To collect a subset of events from that channel replace 'ALL' with the relevant Event IDs. Select **INCLUDE** if only events with the listed Event IDs should be collected or select **EXCLUDE** to collect all events except for these events.
Advanced Configuration	
Throttle Network Bandwidth	Use this setting to limit network bandwidth that the Agent uses to connect to NetWitness. This setting is disabled by default: click **Enabled** to turn it on, and then enter a value in kilobits per second. • If not set, Agent does not do any network throttling. • If set to a positive value ***x***, agent limits network bandwidth to ***x*** kbps. 

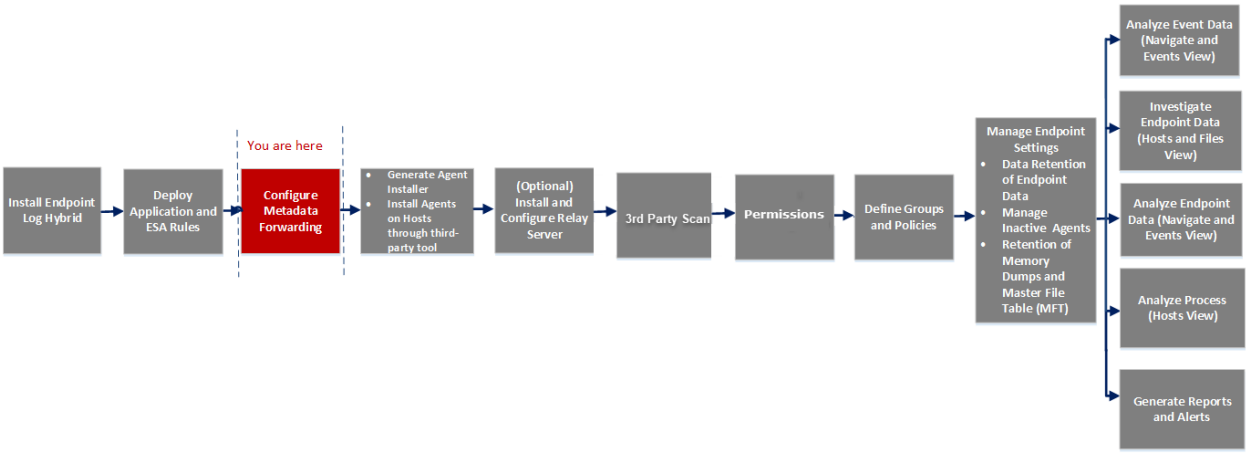
Settings	Description
Advanced Setting	**Caution:** It is strongly recommended not to use this setting unless advised to do so by RSA.

General Tab

In the **General** tab, you can configure the Endpoint metadata forwarding for multiple endpoint servers. To access this view:

1. Go to **(Admin) > Services**.
2. In the Services view, select **Endpoint Server** for which you want to configure the metadata forwarding.
3. Click and select > **View** > **Config**.
4. Click the **General** tab.

Workflow



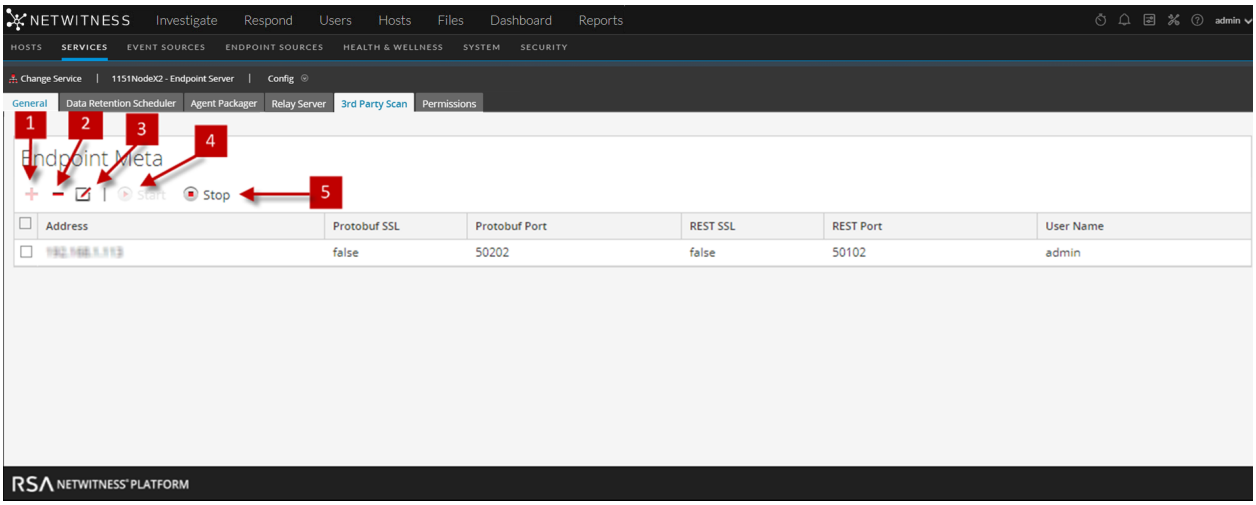
What do you want to do?

User Role	I want to ...	Show me how
Administrator	Configure Endpoint Metadata Forwarding for the NetWitness Endpoint Agents*	Configuring Metadata Forwarding
Administrator	Configure Endpoint Metadata Forwarding for the NetWitness Endpoint 4.4.0.2 or later Agents*	Integrating NW Endpoint 4.4.0.2 or Later with NW Platform

*You can perform this task in the current view

Quick Look

The following figure is an example of the General tab.



1. Click + to view the Available Services dialog.
2. Click − to delete the added service.
3. Click to edit the information for the added service.
4. Click Start to start the Endpoint metadata forwarding.
5. Click Stop to stop the Endpoint metadata forwarding.

The following table describes the fields in the General tab.

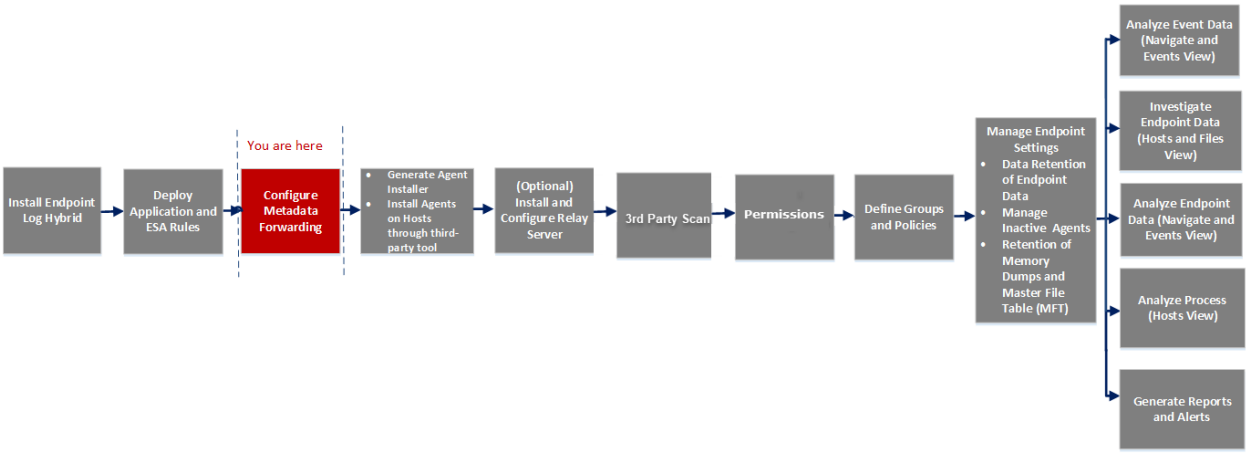
Field	Description
Address	Displays the IP address of the Log Decoder.
Protobuf SSL	Indicates if SSL is enabled on Protobuf. By default, this option is disabled.
Protobuf Port	Displays the port used for Protobuf. By default, the port is 50202.
REST SSL	Indicates if SSL is enabled on the REST port in the Log Decoder. By default, this option is disabled.
REST Port	Displays the port used for REST communication. The default value is 50102 (for non-SSL) and value 56102 (for SSL).
User Name	Displays the user name.
Raw Data	Sends a brief summary of the session along with the metadata if enabled. By default, this option is disabled.

General Tab

In the **General** tab, you can configure the Endpoint metadata forwarding for multiple endpoint servers. To access this view:

1. Go to  **(Admin) > Services**.
2. In the Services view, select **Endpoint Server** for which you want to configure the metadata forwarding.
3. Click  and select > **View** > **Config**.
4. Click the **General** tab.

Workflow



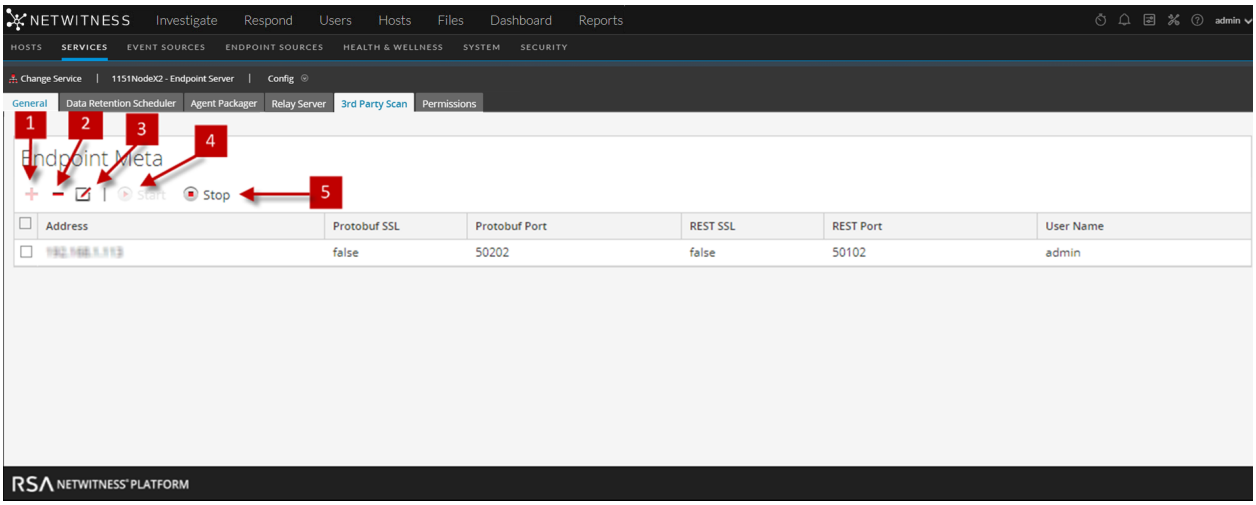
What do you want to do?

User Role	I want to ...	Show me how
Administrator	Configure Endpoint Metadata Forwarding for the NetWitness Endpoint Agents*	Configuring Metadata Forwarding
Administrator	Configure Endpoint Metadata Forwarding for the NetWitness Endpoint 4.4.0.2 or later Agents*	Integrating NW Endpoint 4.4.0.2 or Later with NW Platform

*You can perform this task in the current view

Quick Look

The following figure is an example of the General tab.



1. Click + to view the Available Services dialog.
2. Click − to delete the added service.
3. Click to edit the information for the added service.
4. Click Start to start the Endpoint metadata forwarding.
5. Click Stop to stop the Endpoint metadata forwarding.

The following table describes the fields in the General tab.

Field	Description
Address	Displays the IP address of the Log Decoder.
Protobuf SSL	Indicates if SSL is enabled on Protobuf. By default, this option is disabled.
Protobuf Port	Displays the port used for Protobuf. By default, the port is 50202.
REST SSL	Indicates if SSL is enabled on the REST port in the Log Decoder. By default, this option is disabled.
REST Port	Displays the port used for REST communication. The default value is 50102 (for non-SSL) and value 56102 (for SSL).
User Name	Displays the user name.
Raw Data	Sends a brief summary of the session along with the metadata if enabled. By default, this option is disabled.

Troubleshooting

This section provides information about possible issues when using RSA NetWitness Endpoint.

Agent Communication Issues

Issue	Agent **Last Seen Time** column is not updated in the UI.
Explanation	The issue could be due to any one of the following: • Agent is inactive • Agent data is not processed if the **Endpoint.Health.Overall-Health** statistic shows **Unhealthy** due to which all the agent data including agent last seen time is not updated.
Resolution	See the resolution for these statistics in the Health and Wellness Issues section.

Issue	Agent is unable to communicate with the Endpoint Server.
Explanation	This could be due to one of the following reasons: • Agent is inactive. • Endpoint Server settings is incorrect in the agent packager or policy configuration, or not available for communication. • Endpoint Server or Nginx Server is not running . • Firewall or IP table rules are blocking the connection between the host and Endpoint Server.
Resolution	• Check if the Endpoint Server and Nginx Server are reachable. • If the Endpoint Server settings are incorrect, uninstall the agent, download the agent packager, and reinstall the agent. • Update firewall or IP table rules, if required.

Issue	Agent takes a long time to scan.
Explanation	Sometimes, the NetWitness Endpoint scan takes a long time to complete. This is because of the CPU usage by other antivirus programs (such as Windows Defender, McAfee, Norton, and so on) that may be installed on the agent machines.
Resolution	It is recommended to whitelist the `<service.exe>` (name provided in the packager, by default, the service name is `NWEAgent.exe`) file in the antivirus suite.

Issue	You want to change the responsiveness of the Agent.

Explanation	Depending on your installation, you can adjust Beaconing intervals to change how responsive your agents are.
Resolution	If resources are not a concern, you can lower the HTTPS Beacon Interval and UDP Beacon Intervals. If resources are a concern and responsiveness of the agent is not, you can increase these intervals.

Issue	Agent is unable to generate network tracking events in Insights mode.
Explanation	Verify that Windows Management Instrumentation (WMI) service is running.
Resolution	• Run `Services.msc` and look for Windows Management Instrumentation (WMI) service. • Go to properties and change the **Startup type** to Automatic.

Packager Issues

Message	`Failed to load the client certificate.`
Issue	Incorrect certificate password.
Explanation	While generating the agent installer, the certificate password does not match with the one provided while downloading the agent packager from the UI.
Resolution	Specify the correct certificate password.

Health and Wellness Issues

Endpoint Issues

Behavior	The health check of the **Endpoint.Health.Overall-Health** statistic shows **Unhealthy**.
Issue	Endpoint Server service or required resources are not available or not in a usable state. This could be due to one of the following reasons: • Unable to forward Endpoint meta data to the Log Decoder. • Endpoint Log Hybrid disk usage reaches the specified limit. • Mongo DB is down or excessive read and write errors during processing.
Resolution	• For Log Decoder issue, see Log Decoder Issues section. • For disk usage and Mongo DB issues, see Disk Usage and Mongo Issues section.

Disk Usage and Mongo Issues

Behavior	The health check of the **Data.Application.Connection-Health Application**, **Data Store Disk Usage** or **Data Persistence** for Endpoint Server shows Unhealthy.
Issue	• **Data.Application.Connection-Health Application** or **Data Persistence** shows Unhealthy, if Mongo service is down or fails due to authentication. • **Data Store Disk Usage** shows **Unhealthy**, if Endpoint Server Mongo storage size has exceeded the threshold. By default, the server automatically delete the old data when it reaches 80% of the disk space.
Resolution	• For **Data.Application.Connection-Health Application** or **Data Persistence** issue, you must check the Endpoint server logs (`/var/log/netwitness/endpoint-server/endpoint-server.log`) and Mongo logs (`/var/log/mongodb/mongod.log`), and: ◦ If the issue is due to authentication, you must reissue the certificate. For more information, see "Service Certificate Reissue" section in the *System Maintenance Guide*. ◦ If the issue is due to Mongo service is down, you must restart the Mongo. • For **Data Store Disk Usage** issue, you must increase the storage or configure data retention settings to clear the old data. For more information, see .

Log Decoder Issues

Behavior	Endpoint metadata is not available in the **Investigate > Navigate** or **Events** view.
Issue	The health check of the **Log Decoder Buffer** and **Meta Forward** shows **Unhealthy** in the Health and Wellness.
Explanation	The issue could be due to any of the following reasons: • Log Decoder capture is not started. • Concentrator aggregation is not started. • Log Decoder connection issue. • Log Decoder buffer usage is beyond the specified limit.
Resolution	Make sure that: • Capture is enabled on the Log Decoder. • Aggregation is enabled on the Concentrator. • Meta forwarding is configured properly. **Note:** Make sure **Capture Autostart** is enabled in the Service Config view for Log

	Decoder and **Aggregate Autostart** is enabled in the Service Config view for Concentrator.

File Log Policy Issues

Invalid Policy or Bad Connection Issues

Issue	Policies can be invalid for a variety of reasons. Some examples: • No sources found if the policy is enabled. • Invalid or missing typespec file • No destination is reachable for a file log policy event source type Additionally, if capture is stopped on the destination Log Decoder, Endpoint Agents will send an error to the Endpoint Server saying that they failed to connect. Also, if there is a lot of data to be processed for Agents collecting File data (when File Policy is enabled) , there is a possibility that Log Decoder buffer becomes full. If this happens, the Log Decoder cannot process any requests from the Agents communicating via EPS.
Explanation	The system is dynamic in nature, which means its state can change: event sources can lose their connection, typespec files can be altered or deleted, and other changes can occur that can invalidate a previously valid policy.
Resolution	To help identify the specific issue, check the log file on the Endpoint Server that reports the error: `/var/log/netwitness/endpoint-server/endpoint-server.audit.log` Relevant errors will be listed as **FileLogError** in the log file. If you experience this issue, you can do the following: 1. Try to identify and target higher-value data, thus limiting the total amount of data being processed. 2. Enable throttling in the File policy to smooth out the peaks in usage. 3. If you really do need to process more data on a regular basis, consider server-side hardware upgrades.

Reset File Collection Bookmarks

Issue	If the system is not configured correctly, NetWitness might collect logs and not be able to parse them. Or, files might get sent, but for some reason, not make it to the Log Decoder (for example if communication is via UDP and there is a network connectivity issue). In these and other cases, you can reprocess these "missing" log files.

Explanation	For whatever reason, you may need to reprocess logs from the beginning of the file.
Resolution	Reset bookmarks for an event source type using the procedure described here: Reset File Collection Bookmarks.

Missing Log Collectors and Event Sources in the User Interface

Issue	Some log collectors or event sources seem to be missing from the list of available items.
Explanation	The Filter drop-down menus (types, log collectors, and log decoders) only show values that are in the event sources database, rather than all possible values. For example, if you have a log collector that has not yet collected any logs, then it is missing from the list.
Resolution	Collect logs from a specific log collector and event source, and then they should appear as items in the appropriate menu.

Relay Server Issues

Test Connection Issues

Issue	Relay Server test connection failed.

Resolution	1. Check if the hostname or IP and port of the Relay Server are correct.
	2. Make sure that the hostname or IP of the Relay Server is resolvable from the Endpoint Server. Perform the following:
	a. In the Endpoint Log Hybrid console, verify if the Relay Server is reachable using the following command:
	`nc -zvw3 <relayhost> <relayport>`
	If the Relay Server is not reachable contact your Administrator.
	b. If the Relay Server is reachable, verify if the correct Relay Server installer is used by getting the Endpoint Server revision ID from the Relay Server host (`/var/log/relay-install.log`) and check the Endpoint Server RPM on Endpoint Log Hybrid using the following command:
	`rpm -qa | grep <Endpoint Server Revision ID>`
	c. Make sure if the Relay Server is installed and running.
	• Verify the Relay Server installation logs using the following command:
	`/var/log/relay-install.log`
	• Verify the status of Relay Server using the following command:
	`systemctl status rsa-nw-relay-server`

Issue	Test fails when installing relay server in cloud, using CentOS 7 configuration

Resolution	1. Check if you have entered the suggested port numbers.
	2. If you have entered any other port number than the suggested one
	3. Make sure that the hostname or IP of the Relay Server is resolvable from the Endpoint Server. Perform the following:
	a. In the Endpoint Log Hybrid console, verify if the Relay Server is reachable using the following command:
	`nc -zvw3 <relayhost> <relayport>`
	If the Relay Server is not reachable contact your Administrator.
	b. If the Relay Server is reachable, verify if the correct Relay Server installer is used by getting the Endpoint Server revision ID from the Relay Server host (`/var/log/relay-install.log`) and check the Endpoint Server RPM on Endpoint Log Hybrid using the following command:
	`rpm -qa | grep <Endpoint Server Revision ID>`
	c. Make sure if the Relay Server is installed and running.
	• Verify the Relay Server installation logs using the following command: `/var/log/relay-install.log`
	• Verify the status of Relay Server using the following command: `systemctl status rsa-nw-relay-server`

Issue	Relay Server installer generation fails with an error message ‘Unable to download the installer. Retry after sometime’.
Explanation	Dependencies of the Relay Server are not resolved or downloaded completely.
Resolution	You must retry the download after 5-10 minutes. If the download still fails even after all dependencies are downloaded in the Endpoint Server, contact the RSA Customer Support. **Note:** You can check ‘Finished downloading all Relay Server dependencies’ message in the Endpoint Server logs at `/var/log/netwitness/endpoint-server/endpoint-server.log`, to see if the dependencies are downloaded. If the download fails due to yum related issues, then you must clean yum repo using the command `yum clean all` and restart the Endpoint Server.

Installation Issues

Issue	Relay Server installation fails due to missing or corrupted dependencies.
Resolution	Re-download the installer dependencies, perform the following:

	1. Go to **(Admin) > Endpoint Server** service > select > **View** > **Explore**. 2. In the Endpoint server configuration, make sure **endpoint.relay.installer.download-on-restart** boolean is set to true (by default it is true). 3. Restart the Endpoint server using the following command: `systemctl restart rsa-nw-endpoint-server` Fresh dependencies will be downloaded to the local directory in the Endpoint Server. This may take few minutes. 4. Download the Relay Installer. 5. Run the Relay Server Installation Script. For more information, see (Optional) Installing and Configuring Relay Server.

YARA Issues

Issue	Failure in saving YARA configuration.
Explanation	The rule-folder path is under some Linux user home-directory such as /root or /home/user1.
Resolution	Choose other paths such as /var or /tmp etc. or change the owner of the directory to 'netwitness' user.

OPSWAT Issues

Files can not be scheduled for scan

Issue	Some files cannot be scheduled for the scan. See Endpoint Configuration Guide > Troubleshooting
Explanation	The issue could be due to any one of the following: 1. One of the endpoint servers is either down or not responsive. 2. OPSWAT is not configured in one or more of the endpoint servers.
Resolution	1. Ensure all the endpoint servers are up and running. 2. Ensure OPSWAT is configured in all the endpoint servers.

OPSWAT not configured on all endpoint servers

Issue	OPSWAT is not configured on appropriate Endpoint servers.
Explanation	The issue could be due to any one of the following:

Resolution	1. OPSWAT is not configured in any of the endpoint servers. 2. One or more endpoint servers are down.
	1. Ensure OPSWAT is configured in all the endpoint servers. 2. Ensure all the endpoint servers are functioning.

OPSWAT scan results not getting updated

Issue	OPSWAT scan is scheduled successfully, but no results are getting updated.
Explanation	The issue could be due to any one of the following: 1. Either you have selected unsupported file extensions, or some files are larger than the configured file size limit. 2. OPSWAT server is either down or not responsive.
Resolution	1. Ensure only the files with supported extensions are selected, and file size has not exceeded the configured maximum file size. 2. Contact OPSWAT support for additional information.

Appendices

Reset File Collection Bookmarks

In cases where issues have caused logs to be lost, or not correctly sent to the Log Decoder, you can resend messages in log files by resetting the bookmarks for those log files.

Note: For security reasons, RSA does not allow resetting bookmarks from the agents. Rather, you must do so from an Endpoint Server.

The following procedure describes how to reset bookmarks for file collection logs.

Note: Currently, you can reset bookmarks for all sources or just one specific source, by providing a list in a JSON file.

Construct a JSON File to Identify Agents and Event Source Types for Reset

First, you need to construct a JSON file using the following structure:

```
{

     "agentIds": [],
     "sourceType" : ""

}
```

where:

- `agentIds`: a list of the IDs for one or more Endpoint Agents: these are the individual agents on which the source log files reside.
- `sourceType`: this is a list of the file event source type or types for which you want the log file bookmarks to be reset.

For details on finding agent IDs and source types, see How to Find Agent IDs and Source Types below.

For example, the following source code snippet could be used to delete bookmarks for **all** sources on 3 agents:

```
{

     "agentIds": ["43F27B6E-A02D-955A-9607-2DFC5D17B6E7",
        88AD4B2C-192B-B50E-A125-C05B801301AA"
        "3899038D-8F42-BC93-5BA7-ECBFC309D6A3"],
     "sourceType": "ALL"

}
```

Similarly, the following source code snippet could be used to delete bookmarks for **apache** sources on 3 agents:

```
{
```

```
    "agentIds": ["43F27B6E-A02D-955A-9607-2DFC5D17B6E7",
       88AD4B2C-192B-B50E-A125-C05B801301AA"
       "3899038D-8F42-BC93-5BA7-ECBFC309D6A3"],
    "sourceType": "apache"

}
```

Reset Bookmarks

Perform the following steps to reset the bookmarks that you specified in a JSON-formatted file:

1. SSH to the NetWitness Admin Server.
2. Run `nw-shell` command. for details about using the NetWitness shell, see the *Shell User Guide*, available in RSA Link.
3. After nw-shell starts, connect to an Endpoint Server service, using the following command:

   `connect --service endpoint-server.`*`serviceID`*

   where *serviceID* is identifier for the Endpoint Server that hosts the agents you are changing. See How to Find Endpoint Service IDs for details on how to retrieve the service ID.
4. Change to the directory where the reset command resides:

   `cd endpoint/command/reset-bookmark`
5. Login with an administrator account.
   a. Type the login command:

      `login`

   b. Enter the user name for your admin account.

   c. Enter the password for your admin account.
6. Run the reset command: you need to provide the JSON path and filename that you created earlier.

   `invoke --file <path and filename for JSON>`

   For example:

   `invoke --file /tmp/test.json`

The bookmarks for each log file identified in your JSON file are reset. The following image shows an example NetWitness Shell session:

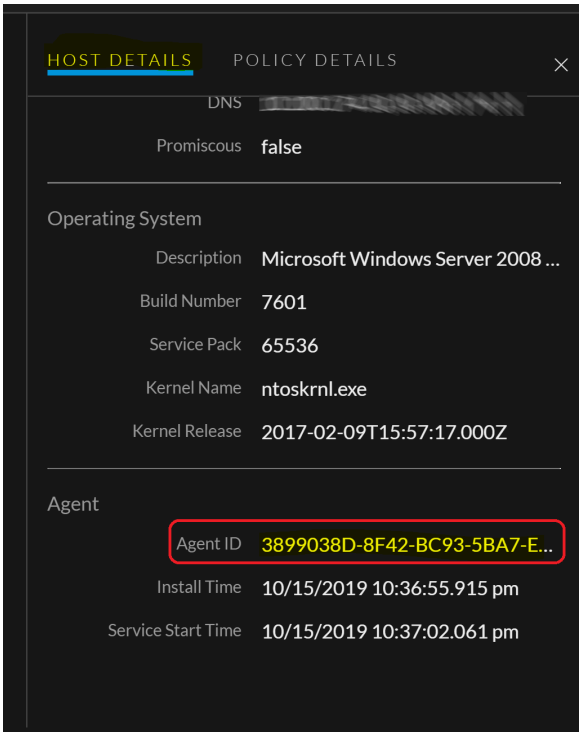
```
[root@SA ~]# nw-shell

RSA Netwitness Shell. Version: 5.15.0-SNAPSHOT

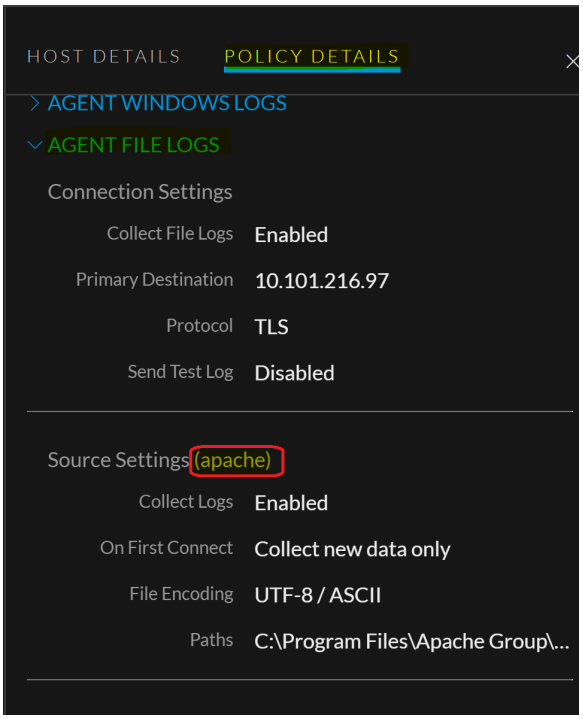

offline » connect --service endpoint-server.38909c2f-7a9b-415a-b567-f49a19cf250e
INFO: Connected to endpoint-server (38909c2f-7a9b-415a-b567-f49a19cf250e)
endpoint-server:Folder:/rsa » cd endpoint/command/reset-bookmark
endpoint-server:Method:/rsa/endpoint/command/reset-bookmark » login
user: admin
password: **********
admin@endpoint-server:Method:/rsa/endpoint/command/reset-bookmark » invoke --file /tmp/test.json
```

How to Find Agent IDs and Source Types

To find the Agent IDs for agents, go to **Hosts** > ***<select an Agent>***, then click the **Host Details** panel, and scroll down to the **Agent** section, where the Agent ID is shown:



To find the source types, go to **Hosts** > ***<select an Agent>***, then click the **Policy Details** panel, expand Agent File Logs, view the Source Settings for the source type name to use:



How to Find Endpoint Service IDs

You can retrieve the service ID for an Endpoint Server by using SSH to connect to it.

To retrieve the service ID for an Endpoint Server:

1. SSH to the NetWitness Endpoint Server for which you need to retrieve the ID. The IP address is available under  **(Admin) > Hosts**. The IP address for each host is listed in the **Host** column of the table.
2. View the file that contains the ID by running the following command:

   ```
   cat /etc/netwitness/endpoint-server/service-id
   ```

   It returns the Endpoint Server ID, for example:

   ```
   38909c2f-7a9b-415a-b567-f49a19cf250e
   ```

Currently Supported File Log Event Source Types

The following event source types are currently supported:

Apache Tomcat	Apache Web Server	CA Siteminder
CiscoWorks	CiscoWorks	Citrix XenMobile Device Manager
Courion Password Courier	EMC NetWorker	EMC Symmetrix
GlobalSCAPE EFT Server	IBM TAM WebSEAL	IBM WebSphere
IBM WebSphere MQ	JBoss Application Server	Kaspersky Anti-Virus
McAfee Endpoint	Microsoft DHCP	Microsoft Exchange
Microsoft Exchange 2007	Microsoft Exchange 2010	Microsoft Exchange 2013
Microsoft Exchange 2016	Microsoft Exchange SMTP	Microsoft Forefront Threat Management Gateway
Microsoft IAS (TVM)	Microsoft ISA	Microsoft ISA 2006
Microsoft ISA PF	Microsoft SQL Server	MSIAS
Oracle Access Management	Oracle iPlanet Web Server	Oracle WebLogic
Oracle WebLogic Audit Recorder	Perforce	Perforce AL
Perforce P4D	Rapid7 NeXpose	RIM Blackberry Enterprise Server
RSA Access Manager	RSA ACE Server	RSA ACE Server
RSA ACE Server AM	RSA ACE Server AMX	RSA Adaptive Auth (Hosted)
RSA Certificate Manager	RSA Federated Identity Manager	SAP ERP Central Component
Steel-Belted Radius Accounting	Steel-Belted Radius Authentication	SunOne LDAP Directory Server
Trend Micro IMSS	Trend Micro IWSS	Trend Micro IWSS Audit
VMware View	Windows DNS Debug Logs	

Specify UNC (Universal Naming Convention) Paths

During configuration of a Log File Policy, you can specify the log file path. You can set one or more paths to be used by the agent to locate the log files. The path can be a standard Windows pathname (such as `C:\Program Files\Apache\error_logs\logfile.log`) or a UNC (Universal Naming Convention) pathname (`\\host-name\share-name\file-path`). This topic describes how to specify a UNC path.

Secure the UNC Path Location

When you use a UNC path to collect log data on a remote system, make sure that you secure the UNC path location. One solution that works with minimal risk in a Windows domain environment is to do the following:

1. Create a share on the directory on the computer where the log data exists in isolation.
2. Name the share to something like **LOGDATA$** for example. (Shares can be hidden from curious browsers by adding a "$" to the end of the share name).
3. Remove all the default share permissions except local admin so it can be changed.
4. Add a share permission for the agent computer system. This allows any user on the agent system to access the shared location. On the agent system collecting the remote log data, nothing else should be required to properly collect the log data from the UNC path.

Note: You may not be able to view the UNC directory contents from file explorer. Seek advice from your IT or security group for additional guidence to setting up and securing a UNC directory share.

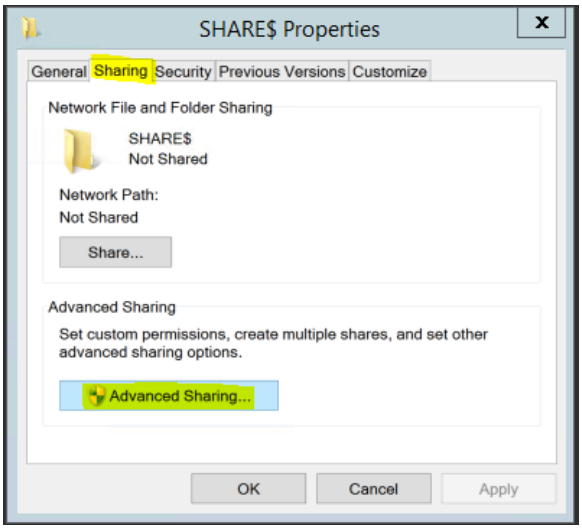
The following procedures describe how to:

- Share a folder between machines in a domain, and
- Share a folder between machines in a Workgroup

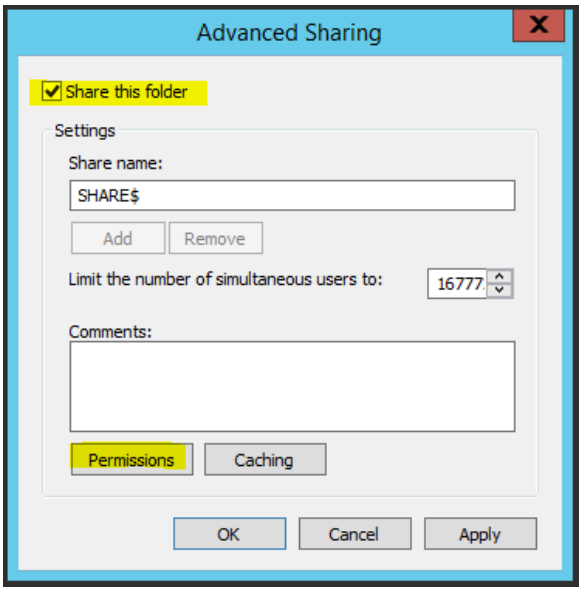
Share a folder between machines in a domain

This procedure describes how to share a folder between Windows machines that are both in the same domain.

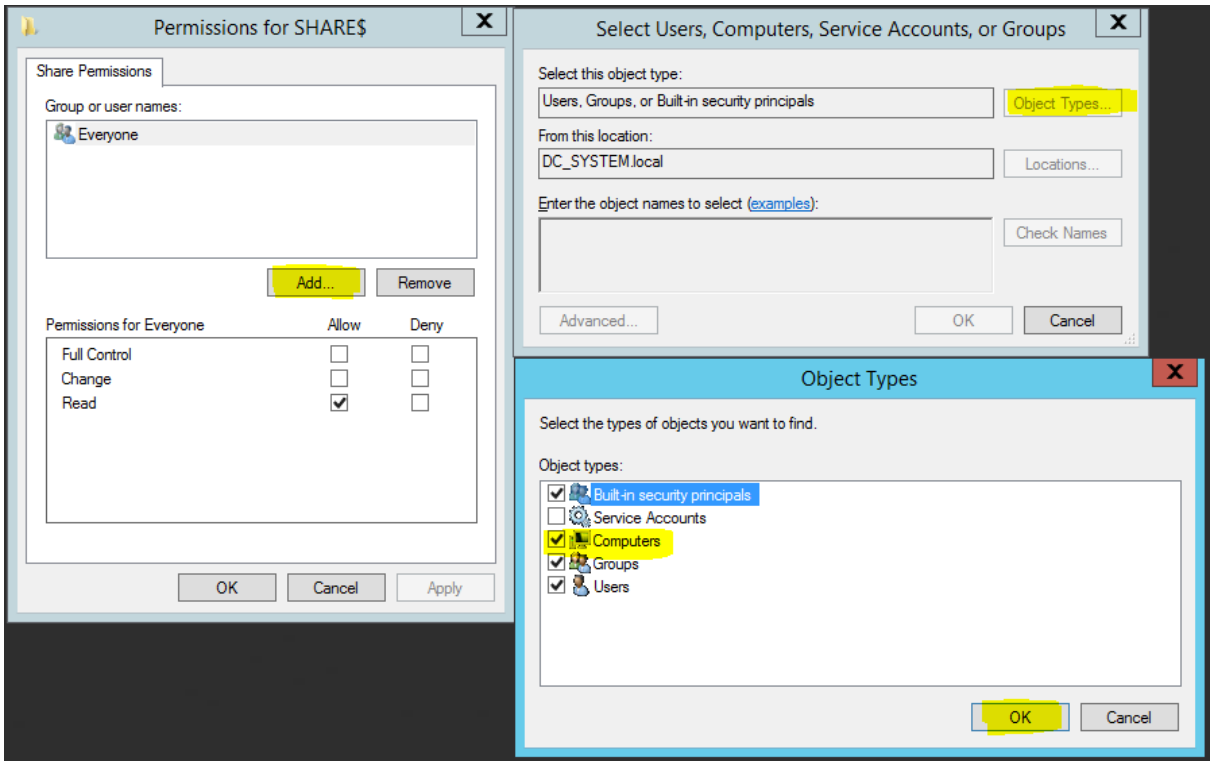
1. Log on to a Domain member machine that contains the logs folder you want to collect using an agent.
2. Right click on the folder you want to share with the agent to collect logs from, and click **Properties**: **SHARE$** in this example.



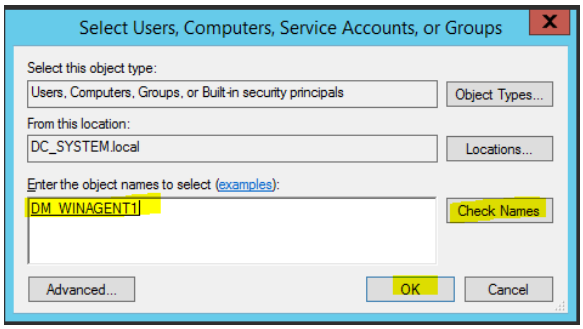
3. Click Advanced Sharing, select Share this folder, and then click Permissions.



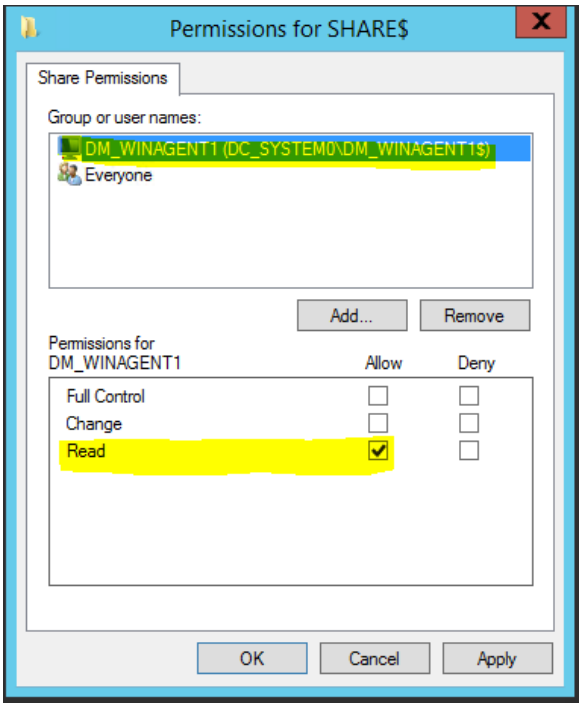
4. Click **Add**, and then on the next window click **Object Types**, check **Computers**, hit **OK**.



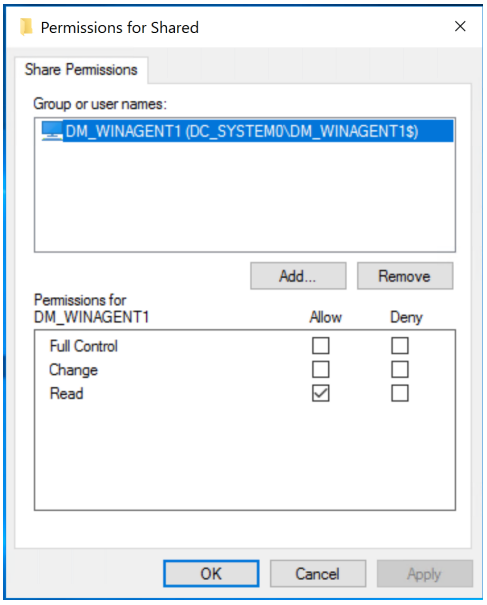
5. Search for the Agent computer name as shown below and click **OK**.



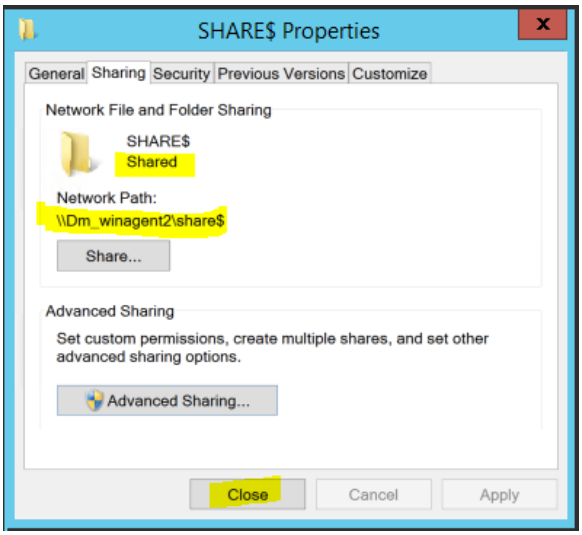
6. In the Permissions screen, provide **Read** permission to the agent, click **Apply**, and then click **OK**.



7. Remove **Everyone** from the Share Permissions list.



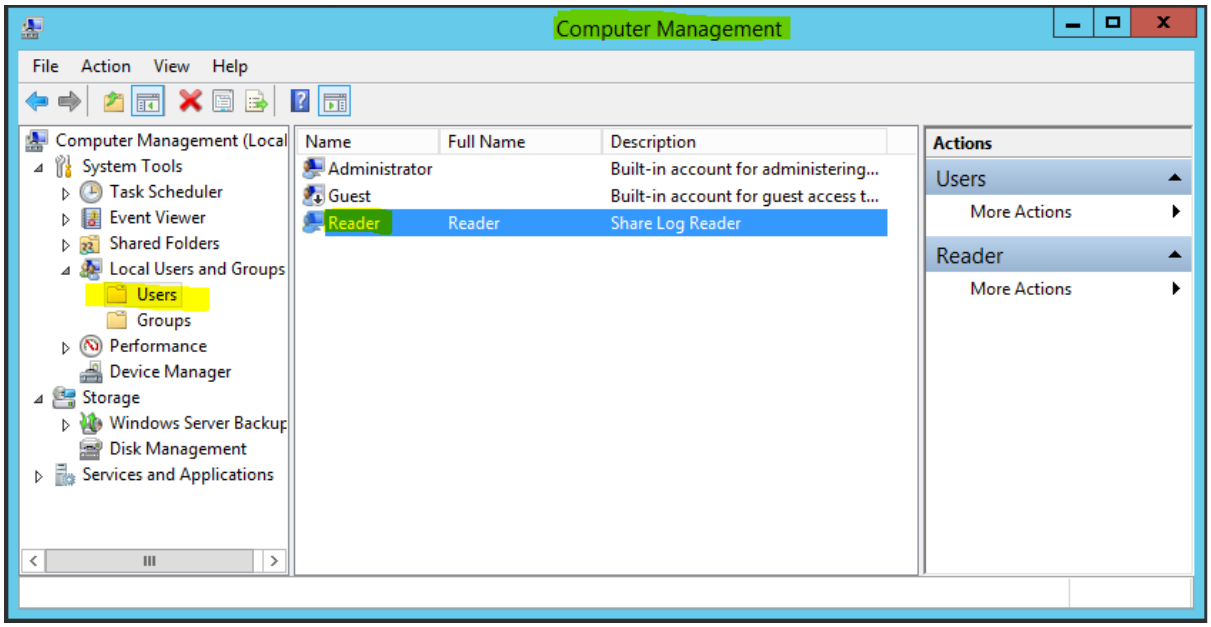
8. Verify that now the **SHARE$** folder status is correct, and note the network path so that you can enter it later, when you configure the policy that will use this shared folder.



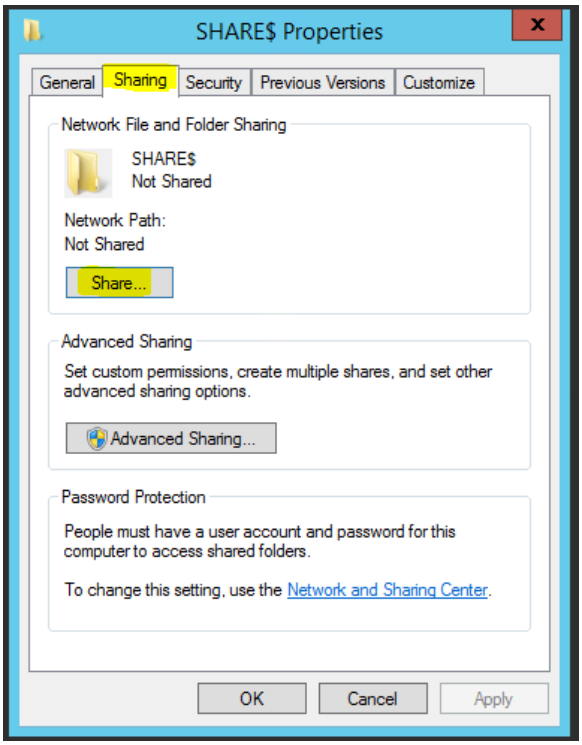
Share a folder between machines in a Workgroup

This procedure describes how to share a folder between Windows machines that are both in the same Workgroup.

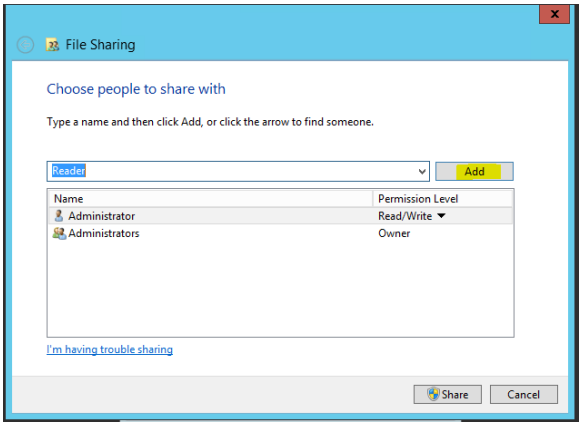
1. Log on to a workgroup machine that contains the logs folder you want to collect using an agent.
2. Create a non-admin user for log collection: **Reader** in this example.



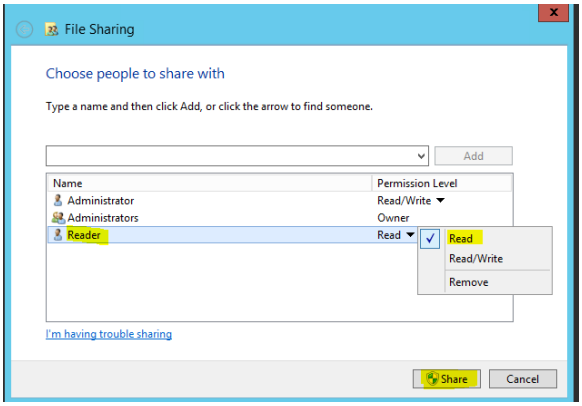
3. Right click on the folder you want to share with the agent to collect logs from, and click **Properties**: **SHARE$** in this example.
4. Click the **Sharing** tab, then click **Share**.



5. Choose the newly-created user and click **Add**.



6. Select the **Read** permission and click **Share**.



7. Log onto the Agent to add credentials, so that the system can read logs from the shared folder.
   a. Download the `psexec` tool from the Microsoft web site.
   b. Run the following command:

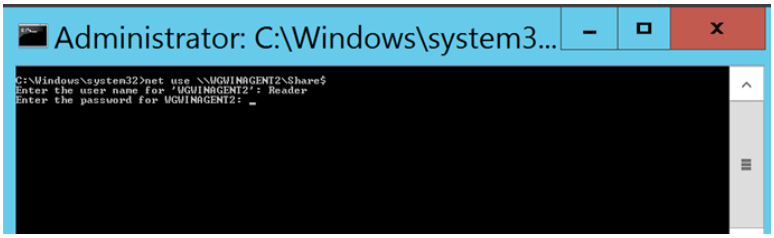
      ```
      psexec -i -s cmd.exe
      ```

      A new command window opens, running as **system**.
   c. In the command window, run the following command to cache credentials for the newly-created user to access logs on shared folder from a workgroup machine:

      ```
      net use \\hostname of machine with logs\Share$
      ```

      For example:

      ```
      net use \\WGWINAGENT2\Share$
      ```
   d. Provide the username and password for the non-admin user created earlier (in step 2).



This command adds credentials to read logs from the shared folder on the workgroup machine.

NetWitness® Platform XDR

Version 12.0

RSA Endpoint Integration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

NetWitness Endpoint Integration **4**

Integration Options 4

Integration Methods 4

NetWitness Endpoint Metadata Integration 5

Built-in NetWitness Endpoint Lookup 6

NetWitness Endpoint Alerts and Indicators of Compromise 6

Configure NetWitness Endpoint Alerts to Respond **7**

Configure NetWitness Endpoint to Forward NetWitness Endpoint Alerts 8

Configure Contextual Data from NetWitness Endpoint Through Recurring Feed **11**

Enable the NetWitness Endpoint Feed for NetWitness 11

Export the NetWitness Endpoint SSL Certificate 14

Configure the NetWitness Concentrator Service 15

Configure the Recurring Custom Feed Task in NetWitness 16

Configure Endpoint Alerts Through Syslog into a Log Decoder **20**

Configure NetWitness Endpoint to Send Syslog Output to NetWitness 21

Edit the Table Mapping in table-map-custom.xml 22

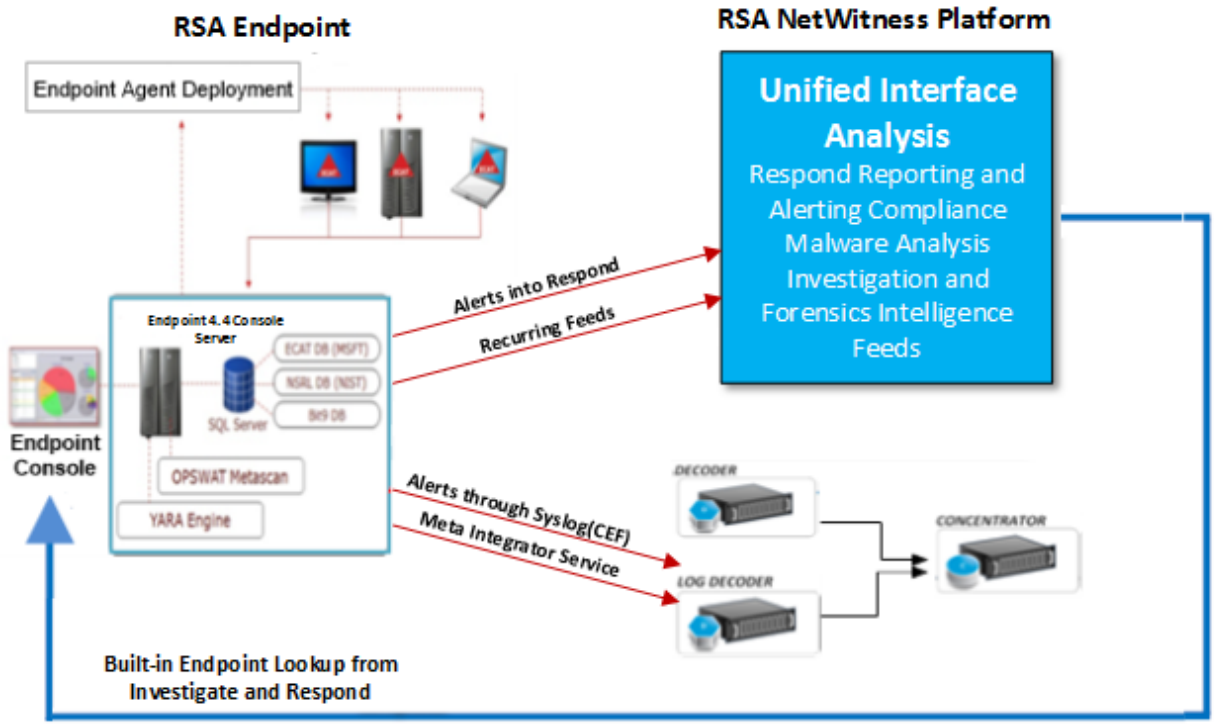
Configure the NetWitness Suite Concentrator Service 24

NetWitness Endpoint Integration

NetWitness customers who are using NetWitness Endpoint 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later can integrate into NetWitness 11.x in several different ways.

Note: In Version 11.2 and later, the following components are rebranded:
- NetWitness Suite to NetWitness Platform
- Packet Decoder to Network Decoder

Integration Options



Integration Methods

The following are the NetWitness Endpoint integration methods:

- Configure Endpoint Alerts through Respond
- Configure Contextual Data from Endpoint through Recurring Feed
- Configure Endpoint Alerts through Syslog into a Log Decoder
- Configuring Meta Integrator service in the NetWitness Endpoint 4.4.0.2 or later directly to a Log Decoder

Endpoint alerts into NetWitness Respond. This integration provides the capability for forwarding Endpoint alerts to Respond.

Contextual data from Endpoint through a NetWitness Live recurring feed. This integration can enrich the session displayed in NetWitness Investigation with contextual information; some examples include the host operating system, MAC address, IIOC score, and other data that may not be present in the log or packet data.

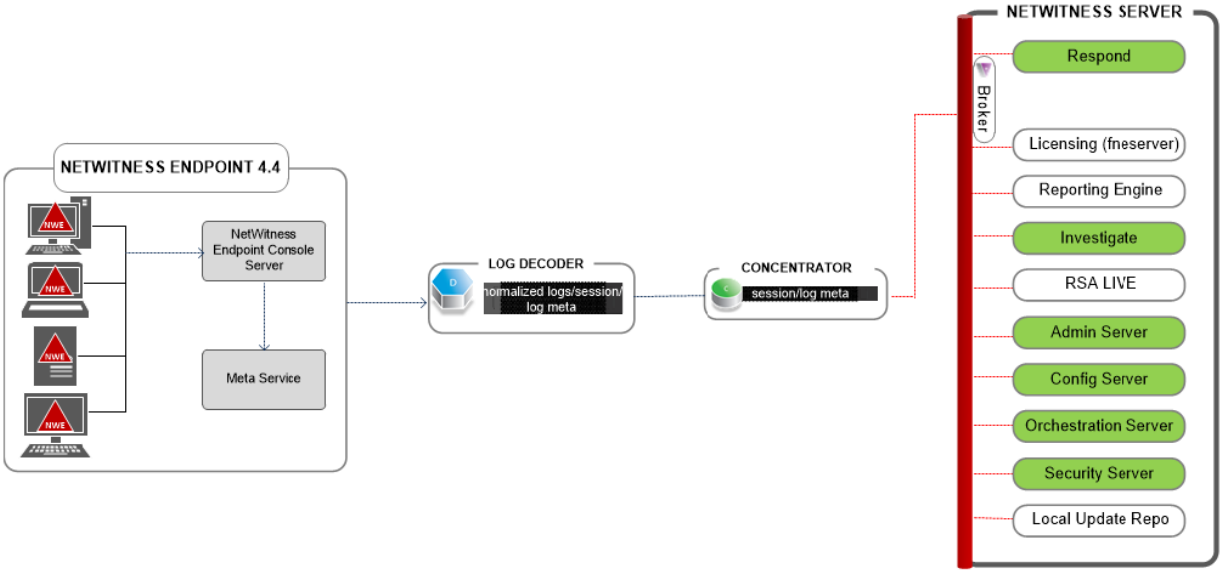
NetWitness Endpoint alerts through Syslog (CEF) into NetWitness Log Decoders. This integration provides the capability to forward Endpoint events through Syslog and to correlate the events with other log or packet metadata in the NetWitness ecosystem.

(For Version 11.1 and later) **NetWitness Endpoint directly to a Log Decoder.** This integration lets you view the Endpoint metadata in the **Investigate > Navigate** and **Event Analysis** view similar to Logs and Packets.

Note: For information on NetWitness Endpoint 4.4.0.x integration with NetWitness, see *NetWitness Endpoint Configuration Guide*.

NetWitness Endpoint Metadata Integration

The NetWitness Platform provides seamless integration allowing Endpoint metadata to be included into the NetWitness work flow. This lets analyst to investigate an incident and respond using packet, log, and endpoint metadata. The endpoint metadata provides further indicators and context related to a host, user, process, or file. It also provides tracking data that provide data of what has transpired with a host, user, process, or file.



Built-in NetWitness Endpoint Lookup

With the NetWitness Endpoint user interface (UI) installed on the same machine where the analyst is using a browser to access NetWitness, the built-in NetWitness Endpoint Lookup from NetWitness Investigation and NetWitness Respond provides right-click access to the NetWitness Endpoint console server for the following meta keys: IP address (ip-src, ip-dst, ipv6-src, ipv6-dst, orig_ip), host (alias-host, domain.dst), client, and file-hash. These are described in the "Launch an External Lookup of a Meta Key" topic in *Investigation and Malware Analysis User Guide* and the "View Alerts" topic in *NetWitness Respond User Guide*.

NetWitness configuration is not required for endpoint lookup when you are using one of the built-in parsers, NetWitness Endpoint or CEF, and you have not customized the default meta keys used when loading metadata in Investigation. For more information, see "Manage and Apply Default Meta Keys in an Investigation" topic in the *Investigation and Malware Analysis User Guide*.

Note: The exception occurs if you customize NetWitness by editing the display setting for the default meta keys in Investigation, add meta keys to the table-map-custom.xml file, or customize NetWitness Endpoint feeds. Some configuration is required to add the custom meta keys to the context menu NetWitness Endpoint Lookup in the **(Admin)** > System view as described in the "Add Custom Context Menu Actions" topic in the *System Configuration Guide*.

NetWitness Endpoint Alerts and Indicators of Compromise

NetWitness Endpoint IIOC (Instant Indicator of Compromise) is a database query that NetWitness Endpoint runs on collected NetWitness Endpoint scan data to determine the presence of potential malware on scanned hosts. NetWitness Endpoint 4.1.2 or later ships with IOCs that users can enable and mark as alertable. NetWitness Endpoint runs IOC queries regularly on new scan data, which is collected and stored in the database. If the IOC query is satisfied, this indicates a potential indicator of compromise, and the event can be reported to a user or sent to an external system as an alert.

Possible types of alerts are:

- Machine alert: This alert indicates that the machine in question is suspicious.
- Module alert: This alert indicates that a module, such as a file, a DLL, or an executable, is suspicious. It contains details about the module in question.
- Event alert: This alert represents any other suspicious activity detected by NetWitness Endpoint that does not fall into the above categories.

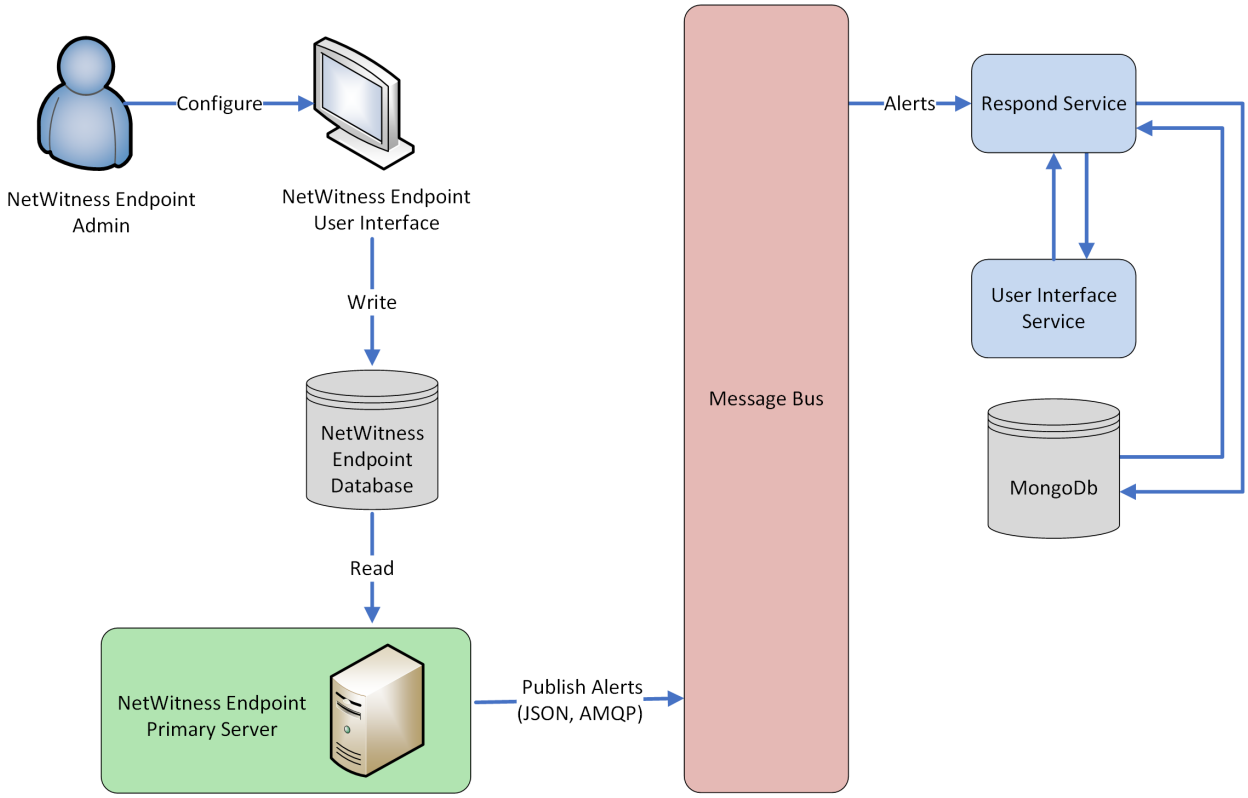
Each of these alert types can be sent to NetWitness.

Configure NetWitness Endpoint Alerts to Respond

This procedure is required to integrate NetWitness Endpoint with NetWitness so that the NetWitness Endpoint alerts are picked up by the Respond component of NetWitness and displayed in the **Respond > Alerts** view.

Note: NetWitness supports NetWitness Endpoint versions 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later for NetWitness Respond integration. For more information, see the "NetWitness Suite Integration" topic in the *NetWitness Endpoint User Guide*.

The diagram below represents the flow of NetWitness Endpoint alerts to the Respond Incident List view of NetWitness and its display in the **Respond > Alerts** view.



Prerequisites

Ensure that you have the following:

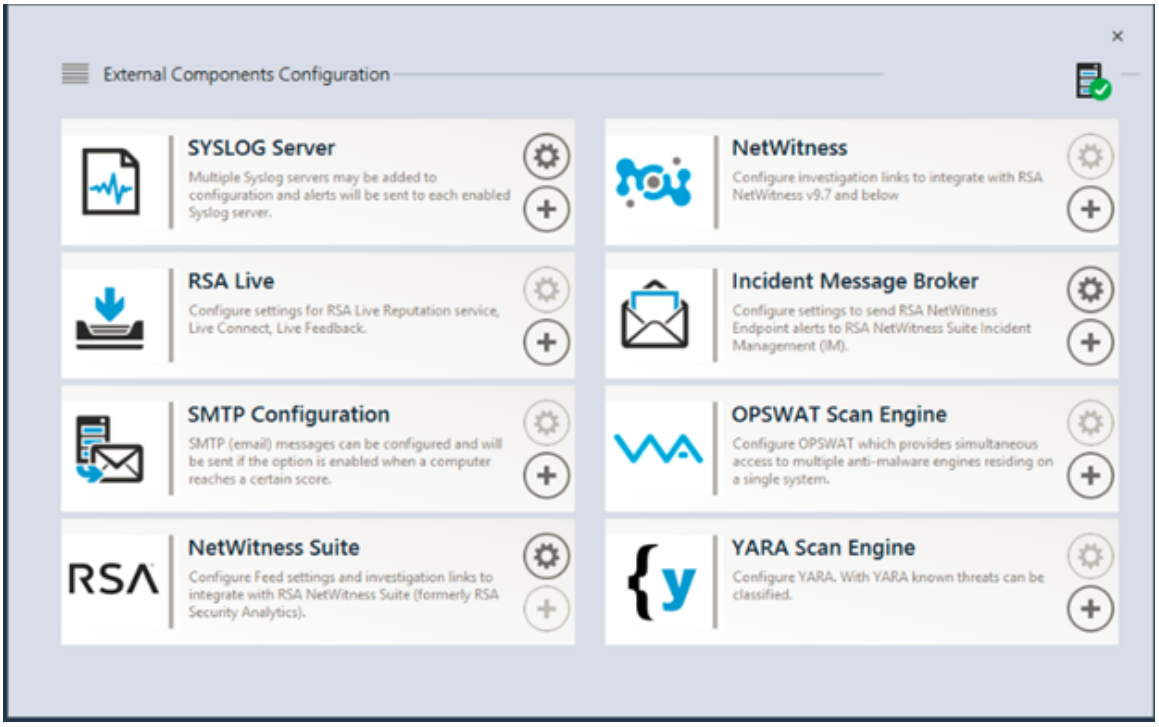
- The Respond service is installed and running on NetWitness.
- NetWitness Endpoint 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later is installed and running.

Configure NetWitness Endpoint to Forward NetWitness Endpoint Alerts

To configure NetWitness Endpoint to send alerts to Respond to the NetWitness user interface:

1. In the NetWitness Endpoint user interface, click **(Configure) > Monitoring and External Components**.

   The External Components Configuration dialog is displayed.



2. From the components listed, select **Incident Message Broker** and click + to add a new IM broker.
3. Enter the following fields:
   a. **Instance Name**: Enter a unique name to identify the IM broker.
   b. **Server Hostname/IP address**: Enter the Host DNS or IP address of the IM broker (NetWitness Server).
   c. **Port number**: The default port is 5671.
4. Click **Save**.
5. Navigate to the **ConsoleServer.exe.config** file in **C:\Program Files\RSA\ECAT\Server**.
6. Modify the virtual host configurations in the file as follows:
   `<add key="IMVirtualHost" value="/rsa/system" />`

Note: In NetWitness 11.0 and 11.1, the virtual host is "/rsa/system". For version 10.6.x and below, the virtual host is "/rsa/sa".

7. Restart the API Server and Console Server.
8. To set up SSL for Respond Alerts, perform the following steps on the NetWitness Endpoint primary console server to set the SSL communications:
   a. Export the NetWitness Endpoint CA certificate to .CER format (Base-64 encoded X.509) from the personal certificate store of the local computer (without selecting the private key).
   b. Generate a client certificate for NetWitness Endpoint using the NetWitness Endpoint CA certificate. (You MUST set the CN name to ecat. Run cmd.exe console with Administrator rights.)

      ```
      makecert -pe -n "CN=ecat" -len 2048 -ss my -sr LocalMachine -a sha1 -sky
      exchange -eku 1.3.6.1.5.5.7.3.2 -in "NweCA" -is MY -ir LocalMachine -sp
      "Microsoft RSA SChannel Cryptographic Provider" -cy end -sy 12
      c:\client.cer
      ```

      Note: In the previous code sample, if you upgraded to version 4.3 (or later) from a previous version and did not generate new certificates, you should substitute `"EcatCA"` for `"NweCA"`. Or, if your current operating system has PowerShell version 5.1 or later, you can use the following code sample:

      ```
      PS C:\> New-SelfSignedCertificate -KeyExportPolicy Exportable -Subject
      "CN=ecat" -KeyAlgorithm RSA -KeyLength 2048 -CertStoreLocation
      "cert:\LocalMachine\My" -HashAlgorithm SHA256 -KeySpec KeyExchange -
      TextExtension @("2.5.29.37={text}1.3.6.1.5.5.7.3.2,1.3.6.1.5.5.7.3.1") -
      Provider "Microsoft RSA SChannel Cryptographic Provider" -KeyUsage
      DigitalSignature, KeyEncipherment, KeyAgreement -Signer (Get-ChildItem -
      Path Cert:\LocalMachine\My\ -DnsName NweCA)-NotAfter (Get-Date).AddYears
      (5); Export-Certificate –Cert (Get-ChildItem –Path
      Cert:\LocalMachine\My\ -DnsName ecat) –FilePath C:\Client.cer
      ```

   c. Make a note of the thumbprint of the client certificate generated in step b. Enter the thumbprint value of the client certificate in the IMBrokerClientCertificateThumbprint section of the ConsoleServer.Exe.Config file as shown.

      ```
      <add key="IMBrokerClientCertificateThumbprint"
      value="896df0efacf0c976d955d5300ba0073383c83abc"/>
      ```

9. On the NetWitness Server, copy the NetWitness Endpoint CA certificate file in .CER format into the import folder:

   `/etc/pki/nw/trust/import`

10. Issue the following command to initiate the necessary Chef run:
    `orchestration-cli-client --update-admin-node`
    This appends all of those certificates into the truststore.
11. Restart the RabbitMQ server:
    `systemctl restart rabbitmq-server`
    The NetWitness Endpoint account should automatically be available on RabbitMQ.
12. Import the `/etc/pki/nw/ca/nwca-cert.pem` and `/etc/pki/nw/ca/ssca-cert.pem` files from the NetWitness Server and add them to the Trusted Root Certification stores in the Endpoint Server.

Troubleshooting

This section suggests how to resolve problems you may encounter when you configure NetWitness Endpoint alerts to Respond.

Known Issues	Solutions
Orchestration fails on admin node.	You must copy and paste the content of NweCA or EcatCA certificate in `/etc/rabbitmq/ssl/truststore.pem` and restart the Rabbitmq service.

Configure Contextual Data from NetWitness Endpoint Through Recurring Feed

You can configure NetWitness Endpoint data in NetWitness to provide contextual data from NetWitness Endpoint to Decoder and Log Decoder sessions. This configuration adds contextual meta values in addition to the instant IOC alerts that can be used to build correlations to other metadata in the NetWitness ecosystem.

Administrators can configure NetWitness to consume system scan contextual data from NetWitness Endpoint through a NetWitness Live recurring feed. This integration can enrich the session from a Decoder or Log Decoder with contextual information displayed in NetWitness Investigation; some examples include the host operating system, MAC address, IIOC score, and other data that may not be present in the log or packet data into sessions from a Decoder or Log Decoder.

Note: Although this feature is targeted for customers with a Network Decoder, a recurring feed can also be implemented in Log Decoders.

Caution: In an environment with many NetWitness Endpoint hosts, using recurring feed may result in decreased performance on the NetWitness ingest devices (Decoder and Log Decoder).

Prerequisites

- Version 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later NetWitness Endpoint Console server and NetWitness Server Version 10.4 and above installed.
- Version 11.0 or 11.1 NetWitness Decoder and Concentrator connected to the NetWitness Server in the network.

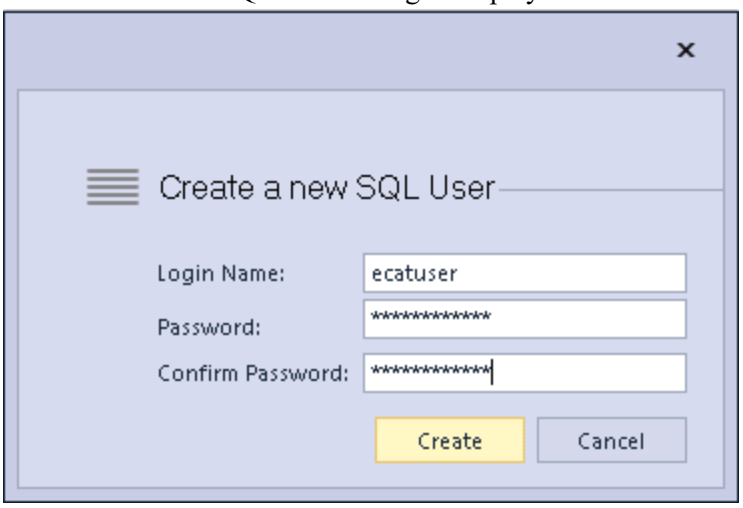
To Configure Contextual Data from NetWitness Endpoint Through Recurring Feed, perform the following:

1. Enable the NetWitness Endpoint Feed for NetWitness in the NetWitness Endpoint User Interface.
2. Export the NetWitness Endpoint CA Certificate from the NetWitness Endpoint Console server and Import it into NetWitness trust store.
3. Configure the NetWitness Concentrator service to define which meta keys are indexed.
4. Create a recurring feed in NetWitness Live.

Enable the NetWitness Endpoint Feed for NetWitness

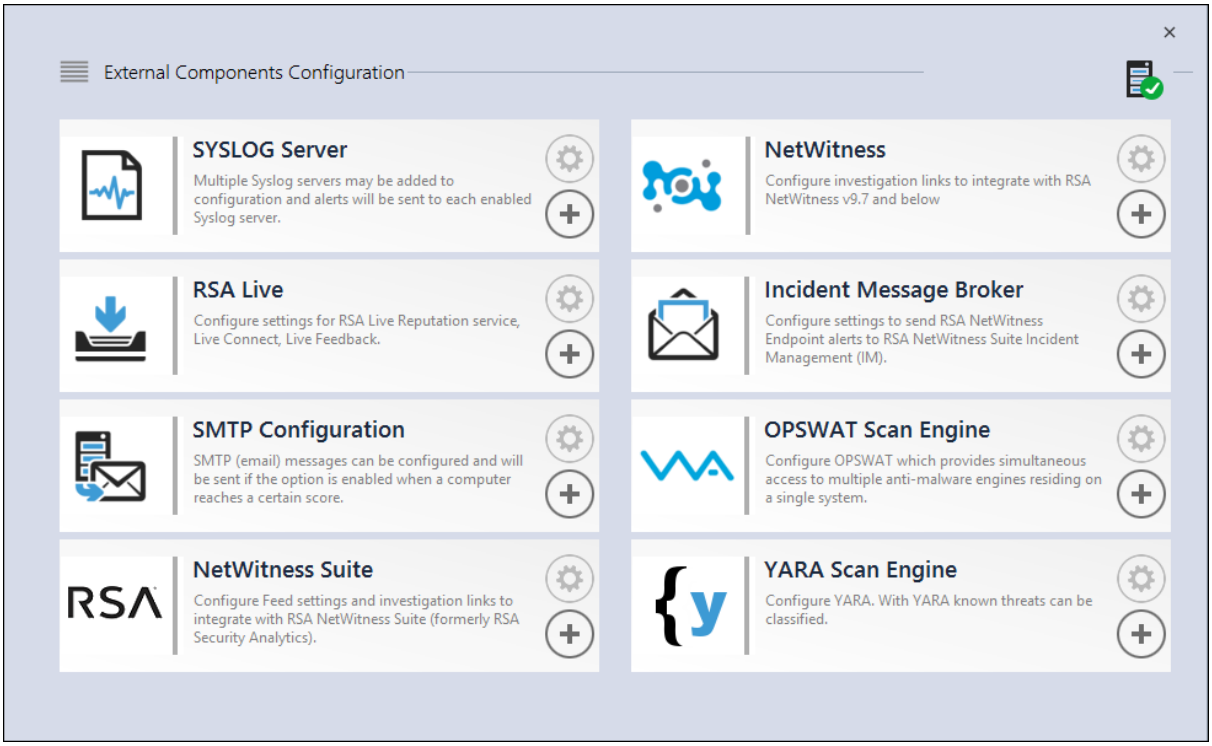
1. In the NetWitness Endpoint user interface, create SQL user in NetWitness Endpoint:
   a. Open the NetWitness Endpoint user interface and log on using the proper credentials.
   b. From the menu bar, select  **(Configure) > Manage Users and Roles**, right-click in the pane, and select **create sql user**.

The Create a new SQL User dialog is displayed.

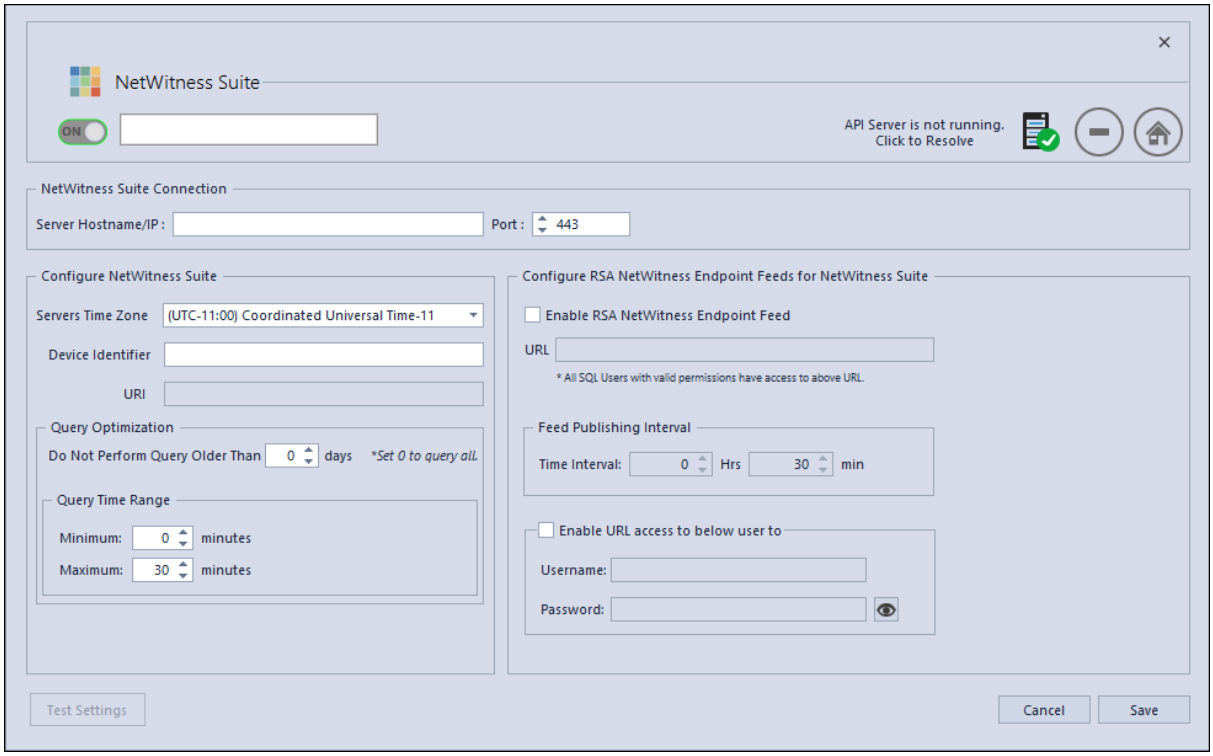


c. Enter the **Login Name** and **Password** and click **Create**.

2. From the menu bar, select **(Configure) > Monitoring and External Components**.
   The External Components Configuration dialog is displayed.



3. In NetWitness, click +.
   The NetWitness dialog is displayed.



4. In the **NetWitness** panel, in **On**, enter the name to identify the NetWitness component.
5. In the **NetWitness Connection** panel, perform the following.
   a. In the **Server Hostname/IP** field, enter the host name or IP address of the NetWitness Server.
   b. In the **Port** field, enter the port number. By default port number is 443.
6. In the **Configure NetWitness** panel, perform the following:
   a. In the **Servers Time Zone** field, select the time zone for the component from the drop-down list.
   b. In the **Device Identifier** field, enter the NetWitness concentrator device ID.

   Note: You can find the Device Identifier in NetWitness when you look up a Concentrator or Broker in **Investigation > Navigate > <*Concentrator or Broker Name*>**. The Device Identifier is the number in the URL after "investigation." For example, in the URL `https://<IP address>investigation/319/navigate/values`, the Device Identifier is **319**.

   The **URI** field is populated when you click **Save**.
7. In the **Query Optimization** panel, in the **Do Not Perform Query Older Than** field, enter the number of days to limit the query period. Enter **0** if you want to discard this feature.

8. In the **Query Time Range** panel, perform the following:
   a. In the **Minimum** field, enter the number of minutes for the minimum query time range. This value is used to automatically increase the time range submitted to NetWitness. This ensures that a query returns a positive response if the NetWitness Endpoint Agent's reported time is slightly different than NetWitness Endpoint's time.
   b. In the **Maximum** field, enter the number of minutes to limit the time range. This value is used to automatically limit the time range submitted to NetWitness, so that a query does not overload the NetWitness Server.
9. In the **Configure NetWitness Endpoint Feeds for NetWitness** panel, perform the following:
   a. Select **Enable NetWitness Endpoint Feed**.
   b. In the **URL** field, enter the SQL **Username** and **Password** (configured in step 1) to access the location of the feed.
   The **URL** field is populated when you click **Save**.
   c. Enter the time interval for the frequency at which feeds are published.
10. In the **Feed Publishing Interval** panel, in the **Time Interval** field, select the time interval in **hrs** and **mins** for the frequency at which feeds are published.
11. In the **Enable URL access to below user to** panel, enter the **Username** and **Password** of the NetWitness Endpoint user.
12. Click **Save**.
    A feed is created.

Export the NetWitness Endpoint SSL Certificate

Note: This procedure works only for NetWitness 10.5 and above because Java 8 support was added for 10.5. If you are using an earlier version of NetWitness, refer to the applicable version of this guide.

To export the NetWitness Endpoint CA certificate from the NetWitness Endpoint Console server and copy it to the NetWitness host:

1. Log on to the NetWitness Endpoint Console.
2. Open **MMC**.
3. Add a certificate snap-in for **Computer account**.
4. Export the certificate named **`NweCA`** (in NetWitness Endpoint 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later fresh install) or **`EcatCA`** (in NetWitness Endpoint upgraded from previous version to 4.3.0.4 or 4.3.0.5).
   a. Export without a private key.
   b. Export in DER encoded binary X.509 (`.CER`) format.
   c. Name it **`NweCA.cer`** (in NetWitness Endpoint 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later fresh install) or **`EcatCA.cer`** (in NetWitness Endpoint upgraded from previous version to 4.3.0.4 or 4.3.0.5).

5. Copy the NetWitness Endpoint CA certificate to the NetWitness host:
   - For NetWitness Endpoint 4.3.0.4, 4.3.0.5 or 4.4 fresh installation:
     `scp NweCA.cer root@<sa-machine>:.`
   - For NetWitness Endpoint upgraded from previous version to 4.3.0.4 or 4.3.0.5:
     `scp EcatCA.cer root@<sa-machine>:.`
6. To import the NetWitness Endpoint Endpoint CA certificate into the NetWitness Trusted store, navigate to java directory. Enter the following commands:
   - For NetWitness Endpoint fresh installation:
     `keytool -import -v -trustcacerts -alias nweca -file ~/NweCA.cer -keystore /etc/pki/java/cacerts -storepass changeit`
   - For NetWitness Endpoint upgraded from previous version:
     `keytool -import -v -trustcacerts -alias ecatca -file ~/EcatCA.cer -keystore /etc/pki/java/cacerts -storepass changeit`

   When prompted for certificate update confirmation, enter **Yes**.
7. Create a file `jetty.user` in the `/etc/default` directory and add the following text.
   `JAVA_OPTIONS="${JAVA_OPTIONS} -Djdk.security.allowNonCaAnchor=true"`
8. On the NetWitness host, do one of the following:
   - For NetWitness Endpoint 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later fresh installation, edit**`/etc/hosts.user`** to map the IP address of the NetWitness Endpoint Console server to the name **NweServerCertificate** by adding the following line to the file:

     `<ip-address-ecat-cs> NweServerCertificate`
   - For NetWitness Endpoint upgraded from previous version to 4.3.0.4 or 4.3.0.5, edit **`/etc/hosts.user`** to map the IP address of the upgraded NetWitness Endpoint Console server to the name **`ecatserverexported`** by adding the following line to the file:

     `<ip-address-ecat-cs> ecatserverexported`
9. SSH to ADMIN Server and run the following command.
   `nw-manage --refresh-host --host-all`
10. To restart NetWitness, enter the following command:
    service jetty restart

Configure the NetWitness Concentrator Service

1. Log on to NetWitness and go to  **(Admin) > Services**.
2. Select a **Concentrator** from the list and select **View > Config.**
3. Select the **Files** tab, and from the **Files to Edit** drop-down menu, select **index-concentrator-custom.xml**.
4. Add the following NetWitness Endpoint meta keys to the file and click **Apply**. Make sure that this file contains the XML sections already; if the lines are not included, add them. The following lines are examples; make sure the values match your configuration and the column names you included in

the feed definition, where:
description is the name of the meta key you want to display in NetWitness Investigation.
level is "IndexValues"
name matches the column name of the CSV file that NetWitness uses while defining the recurring feed (see the table in *Configure the Recuring Custom Feed Task in NetWitness* below).

```
<key description="Gateway" format="Text" level="IndexValues" name="gateway"
valueMax="250000" defaultAction="Open"/>
<key description="Risk Number" format="Float64" level="IndexValues"
name="risk.num" valueMax="250000" defaultAction="Open"/>
<key description="Strans Addr" format="Text" level="IndexValues"
name="stransaddr" valueMax="250000" defaultAction="Open"/>
<key description="Domain" format="Text" level="IndexValues" name="domain"
valueMax="250000" defaultAction="Open"/>
<key description="User Account" format="Text" level="IndexValues"
name="username" valueMax="250000" defaultAction="Open"/>
<key description="Ecat Connectiontime" format="Text" level="IndexValues"
name="ecat.ctime" valueMax="250000" defaultAction="Open"/>
<key description="Ecat Scantime" format="Text" level="IndexValues"
name="ecat.stime" valueMax="250000" defaultAction="Open"/>
```

5. Restart the Concentrator to activate the custom key updates.

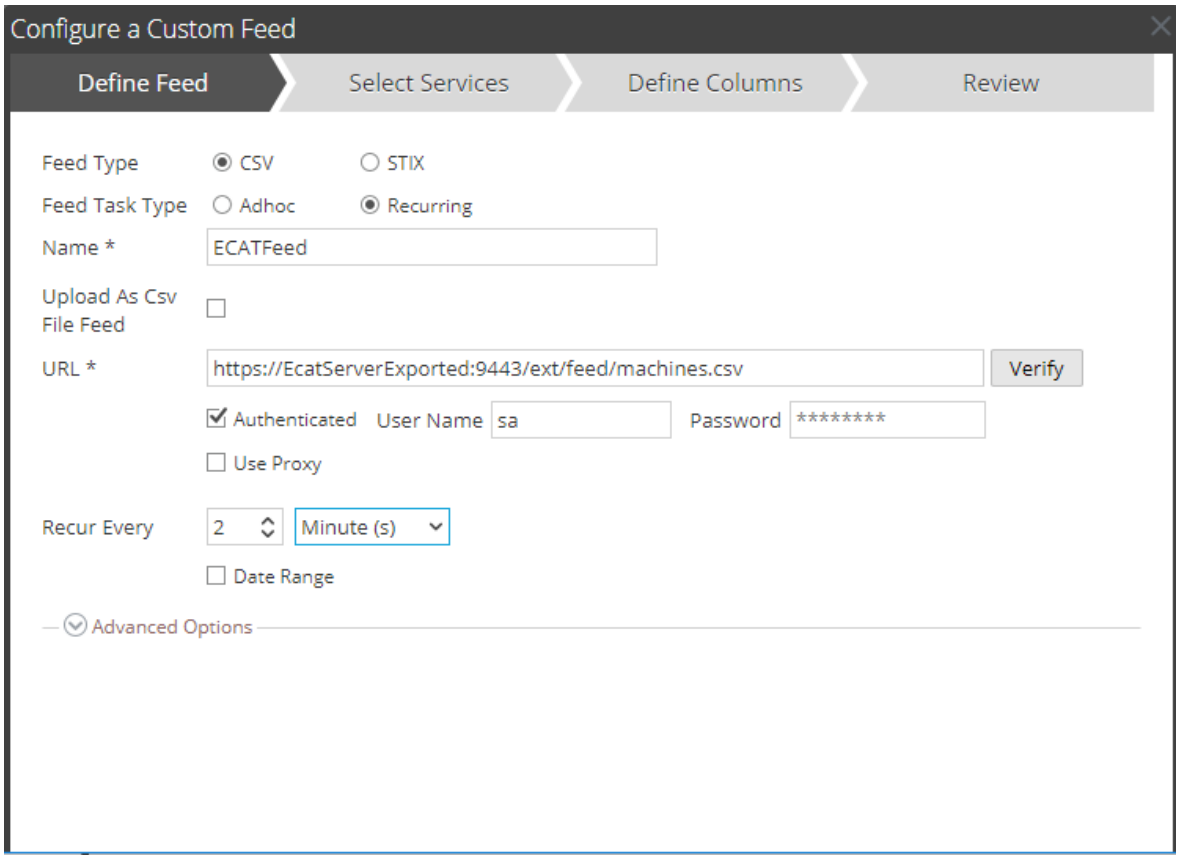
Configure the Recurring Custom Feed Task in NetWitness

1. Log on to NetWitness and go to  **(Configure) > Custom Feeds**.
   The Feeds view is displayed.
2. In the toolbar, click .
   The Setup Feed dialog is displayed.
3. In the Setup Feed dialog, select **Custom Feed** and click **Next**.
   The Configure a Custom Feed wizard is displayed, with the Define Feed form open.
4. In the **Define Feed**, perform the following:
   a. In the **Feed Type** field, select **CSV**.
   b. In the **Feed Task Type** field, select **Recurring**.
   c. In the **Name** field, enter the name of the feed. For example, EndpointFeed.
   d. Enable the **Upload As Csv File Feed** checkbox to upload the feed as a CSV file.
   e. In the **URL** field, enter the URL with the hostname of the Windows server on which NetWitness Endpoint is installed:
      - For NetWitness Endpoint 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later fresh installation, use the URL **`https://NweServerCertificate:9443/api/v2/feed/machines.csv`**.
      - For NetWitness Endpoint upgraded from previous version to 4.3.0.4 or 4.3.0.5, use the URL **`https://ecatserverexported:9443/api/v2/feed/machines.csv`**.

f. Enable the **Authenticated**checkbox and enter the username and password as noted in *Enable the ECAT Feed* above.

g. Click **Verify** to check if NetWitness can reach the web resource.

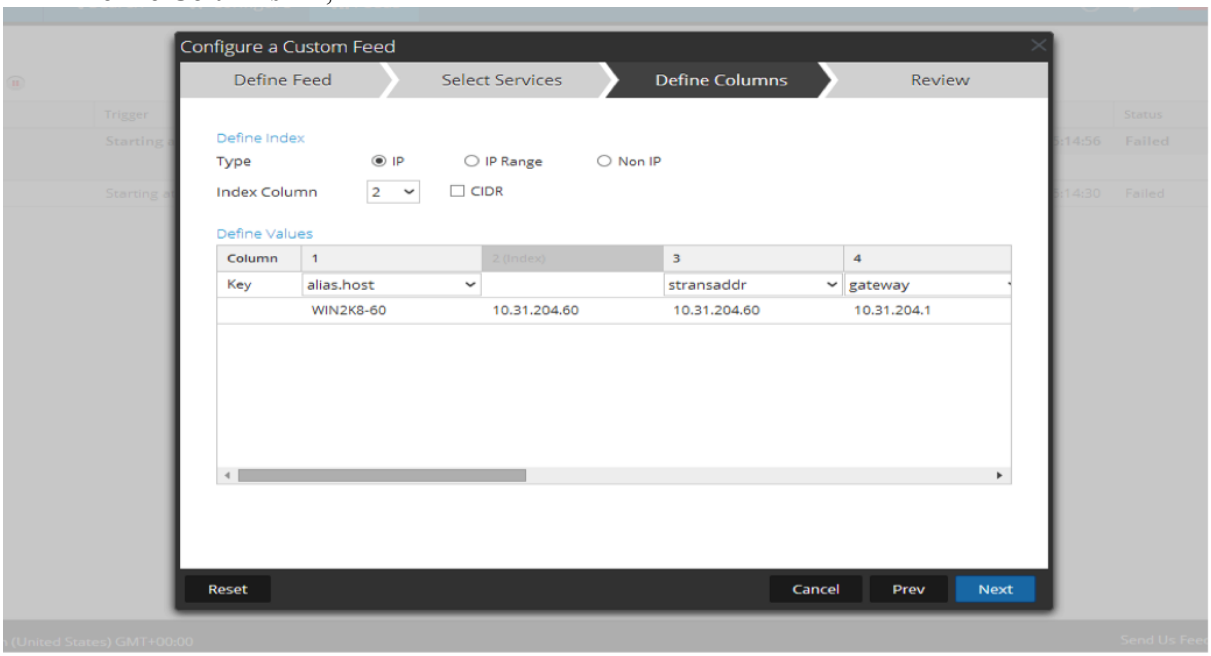
.



h. Define the schedule and click **Next**

5. In the **Select Services** tab, select the Decoder or groups to consume the feed. Click **Next**.

6. In the **Define Columns** tab, enter the column names as shown in the table below and save the feed.

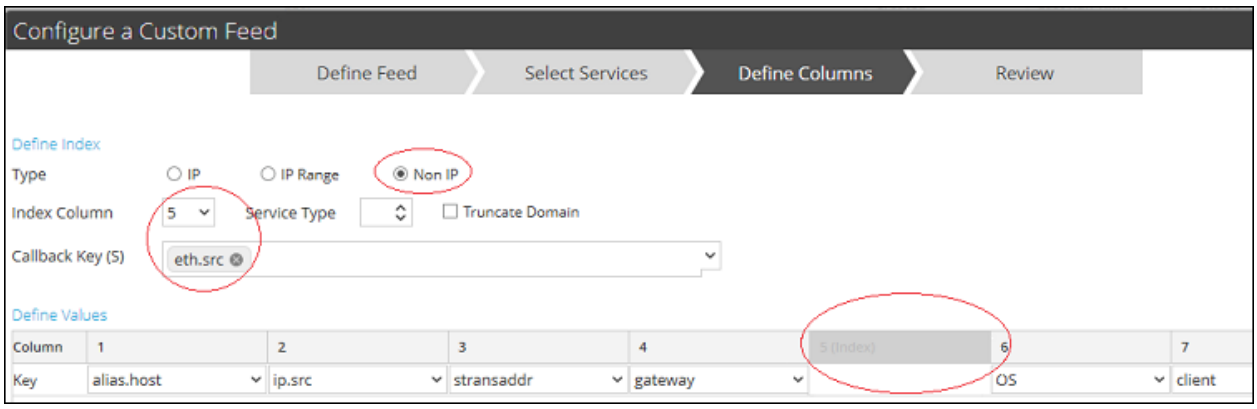


The following table shows the columns in the CSV file for the NetWitness Endpoint feed.

Column	Name	Description	Column Name in NetWitness (Meta Key Name)
1	MachineName	Host name of the Windows agent	alias.host
2	LocalIp	IPv4 address	IP type (indexed column)
3	RemoteIp	Far end IP as seen by the router	stransaddr
4	GatewayIp	IP of the gateway	gateway
5	MacAddress	MAC address	eth.src
6	OperatingSystem	Operating system used by the Windows Agent	OS
7	AgentID	Agent ID of the host (unique ID assigned to the agent)	client
8	ConnectionUTCTime	Last time when the agent connected to NetWitness Endpoint server	ecat.ctime
9	Source Domain	Domain	domain.src
10	ScanUTC time	Last time when the agent was scanned	ecat.stime

Column	Name	Description	Column Name in NetWitness (Meta Key Name)
11	UserName	Username of the client machine	username
12	Machine Score	Score of the agent indicating the suspicious level	risk.num

Note: In the table, the recommended index setting is LocalIp. However, if the LocalIp for NetWitness Endpoint Agent PC is allocated by a DHCP Server and the DHCP lease has expired, and if the IP is then re-allocated to another PC, the metadata created by the feed will be incorrect. To avoid this risk, use the machine name or the Mac address instead of the localIP address as the Feed's index. For example, to use a Mac address, you could enter the values as shown in the following figure.



Result

When viewing feed data in NetWitness, upon a match of the indexed value (ip.src), meta data is populated in Investigation, Reporting, and Alerting Interfaces.

Configure Endpoint Alerts Through Syslog into a Log Decoder

You can configure the use of NetWitness Endpoint data in NetWitness to provide NetWitness Endpoint alerts through Syslog into Log Decoder sessions. This generates metadata that is used by NetWitness Investigation, Alerts, and Reporting Engine.

For NetWitness networks that are consuming logs, this integration of NetWitness Endpoint with NetWitness pushes NetWitness Endpoint events to the Log Decoder through common event format (CEF) syslog messages and generates metadata that is used by NetWitness Investigation, Alerts, and Reporting Engine. The use case for this integration is SIEM Integration to allow centralized event management, correlation of NetWitness Endpoint events with other Log Decoder data, NetWitness reporting on NetWitness Endpoint events, and NetWitness alerting of NetWitness Endpoint events.

Prerequisites

The following are required for this integration:

- Version 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later NetWitness Endpoint UI.
- NetWitness Server Version 11.1 is installed.
- Version 10.4 or later NetWitness Log Decoder and Concentrator connected to the NetWitness Server in the network.
- Port UDP- 514 or TCP - 1514 open from NetWitness Endpoint server to Log Decoder in the firewall.

Procedure

1. Deploy the required parser (CEF or rsaecat) to the Log Decoder as described in the "Manage Live Resources" topic in *Live Services Management*. After you deploy the parser, make sure the parser is enabled. For more information, see "Services Config View - General Tab" in the *Malware Analysis Configuration Guide*.

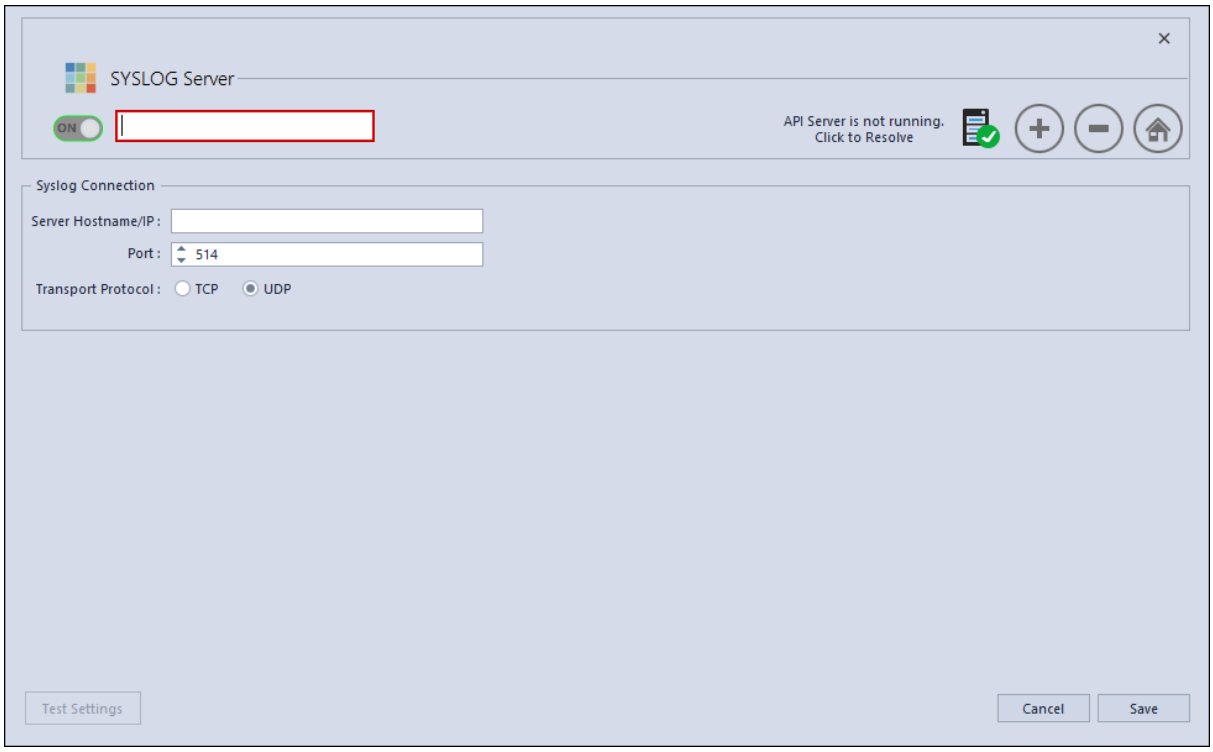
Note: Use only one of these parsers. When the CEF parser is deployed, it supersedes the NetWitness Endpoint parser, and all CEF messages into NetWitness are processed by the CEF parser. Enabling both parsers is an unnecessary burden on performance.

2. Configure NetWitness Endpoint to send syslog output to NetWitness and generate NetWitness Endpoint alerts to the Log Decoder.
3. (Optional) Edit the table mapping in `table-map-custom.xml` and the `index-concentrator-custom.xml` to add fields based on user preferences for metadata to be mapped to NetWitness.

Configure NetWitness Endpoint to Send Syslog Output to NetWitness

To add the Log Decoder as a Syslog external component and generate NetWitness Endpoint alerts to the Log Decoder:

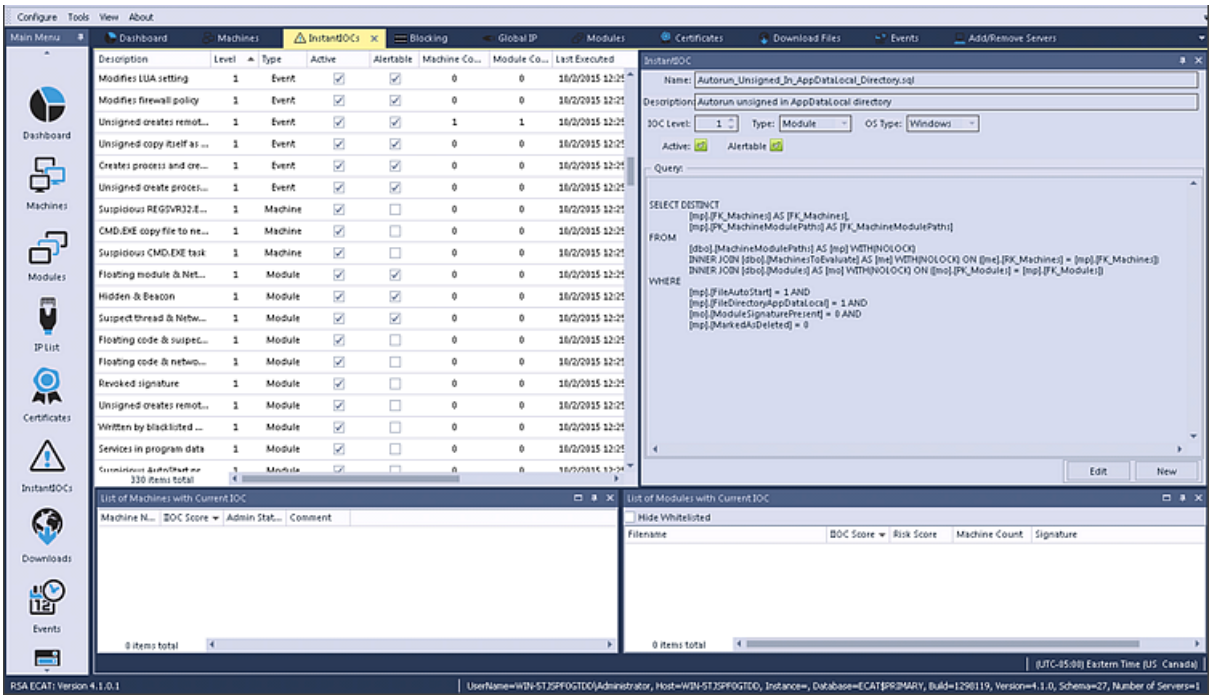
1. Open the NetWitness Endpoint user interface and log on using the proper credentials.
2. From the menu bar, select **(Configure) > Monitoring and External Components**.
   The External Components Configuration dialog is displayed.
3. In **SYSLOG Server**, click .
   The SYSLOG Server dialog is displayed.



4. In the **NetWitness** panel, in **On**, enter the descriptive name for the Log Decoder.
5. In the **Syslog Connection** panel, perform the following to enable Syslog messaging:

   Server Hostname/IP = The hostname DNS or IP address of the NetWitness Log Decoder
   Port = 514
   Transport Protocol = Select **UDP** or **TCP** as appropriate for your Syslog server for the transport protocol.

6. Click **Save**.
7. Open the **InstantIOCs** window in the NetWitness Endpoint UI and, in the **Alertable** column, click to enable each IIOC for which you want alerts sent to the Log Decoder.



When the instant IOCs are triggered, Syslog alerts from the NetWitness Endpoint server are sent to the Log Decoder. Log Decoder alerts are then aggregated to the Concentrator. These events are injected into the Concentrator as metadata.

Edit the Table Mapping in table-map-custom.xml

In the default NetWitness table-map.xml provided by NetWitness, the meta keys in the `table-map.xml` file are set to `Transient`. In order to view the meta keys in Investigation, the keys must be set to `None`. To make changes to the mapping, you must add the entries to the `table-map-custom.xml` on the Log Decoder.

This is the list of meta keys in `table-map.xml`.

NetWitness Endpoint Fields	NetWitness Mapping	Transient in NetWitness
agentid	client	No
CEF Header Hostname Field	alias.host	No
CEF Header Product Version	version	No
CEF Header Product Name	Product	Yes
CEF Header Severity	severity	Yes

NetWitness Endpoint Fields	NetWitness Mapping	Transient in NetWitness
CEF Header Signature ID	event.type	No
CEF Header Signature Name	event.desc	No
destinationDnsDomain	ddomain	Yes
deviceDnsDomain	domain	Yes
dhost	host.dst	No
dst	ip.dst	No
end	endtime	Yes
fileHash	checksum	No
fname	filename	No
fsize	filename.size	No
gatewayip	gateway	Yes
instantIOCLevel	threat.desc	No
instantIOCName	threat.category	No
machineOU	dn	No
machineScore	risk.num	No
md5sum	checksum	No
os	OS	No
port	ip.dstport	No
protocol	protocol	Yes
Raw Message	msg	Yes
remoteip	stransaddr	Yes
rt	alias.host	No
sha256sum	checksum	No
shost	host.src	No
smac	eth.src	No
src	ip.src	No
start	starttime	No
suser	user.dst	No
timezone	timezone	No

NetWitness Endpoint Fields	NetWitness Mapping	Transient in NetWitness
totalreceived	rbytes	Yes
totalsent	bytes.src	No
useragent	user.agent	No
userOU	org	Yes

The following seven keys are not in `table-map.xml`; to use these keys in NetWitness you need to add them to `table-map-custom.xml`, and set the flags to `None`.

NetWitness Endpoint Fields	NetWitness Mapping	Transient in NetWitness
moduleScore	cs.modulescore	Yes
moduleSignature	cs.modulesign	Yes
Target module	cs.targetmodule	Yes
YARA result	cs.yararesult	Yes
Source module	cs.sourcemodule	Yes
OPSWATResult	cs.opswatresult	Yes
ReputationResult	cs.represult	Yes

Here are the entries to be added to the `table-map-custom.xml` if required.

```
<mapping envisionName="cs_represult" nwName="cs.represult" flags="None"
envisionDisplayName="ReputationResult"/>
 <mapping envisionName="cs_modulescore" nwName="cs.modulescore" format="Int32"
flags="None" envisionDisplayName="ModuleScore"/>
 <mapping envisionName="cs_modulesign" nwName="cs.modulesign" flags="None"
envisionDisplayName="ModuleSignature"/>
 <mapping envisionName="cs_opswatresult" nwName="cs.opswatresult" flags="None"
envisionDisplayName="OpswatResult"/>
 <mapping envisionName="cs_sourcemodule" nwName="cs.sourcemodule" flags="None"
envisionDisplayName="SourceModule"/>
 <mapping envisionName="cs_targetmodule" nwName="cs.targetmodule" flags="None"
envisionDisplayName="TargetModule"/>
 <mapping envisionName="cs_yararesult" nwName="cs.yararesult" flags="None"
envisionDisplayName="YaraResult"/>
```

Note: Restart the Log Decoder or reload the log parsers for the changes to take effect.

Configure the NetWitness Suite Concentrator Service

1. Log on to NetWitness and go to  **(Admin) > Services**.
    1. Select a **Concentrator** from the list and select **View > Config.**
2. Select the **Files** tab, and from the **Files to Edit** drop-down list, select **index-concentrator-custom.xml**.

3. Add the NetWitness Endpoint meta keys to the file and click **Apply**. Make sure that this file contains the XML sections already; if the lines are not included, add them.
4. Restart the Concentrator.
5. To add the Concentrator as a data source in the Reporting Engine, in the  **(Admin) > Services** view, select the **Reporting Engine** and Select **View> Config > Sources.**
   NetWitness Endpoint meta is populated in Reporting Engine, and you can run reports by selecting the appropriate meta keys.

Example

Note: The following lines are examples; make sure the values match your configuration and the column names you included in the feed definition, where:
description is the name of the meta key you want to display in NetWitness Investigation.
level is "IndexValues"
name is the NetWitness Endpoint meta key name from the table below

```
<language>
<key description="Product" format="Text" level="IndexValues" name="product"
valueMax="250000" defaultAction="Open"/>
 <key description="Severity" format="Text" level="IndexValues" name="severity"
valueMax="250000" defaultAction="Open"/>
 <key description="Destination Dns Domain" format="Text" level="IndexValues"
name="ddomain" valueMax="250000" defaultAction="Open"/>
 <key description="Domain" format="Text" level="IndexValues" name="domain"
valueMax="250000" defaultAction="Open"/>
 <key description="Destination Host" format="Text" level="IndexValues" name="host.dst"
valueMax="250000" defaultAction="Open"/>
 <key description="End Time" format="TimeT" level="IndexValues" name="endtime"
valueMax="250000" defaultAction="Open"/>
 <key description="Checksum" format="Text" level="IndexValues" name="checksum"
valueMax="250000" defaultAction="Open"/>
 <key description="Filename Size" format="Int32" level="IndexValues"
name="filename.size" valueMax="250000" defaultAction="Open"/>
 <key description="Gateway" format="Text" level="IndexValues" name="gateway"
valueMax="250000" defaultAction="Open"/>
 <key description="Domain OU" format="Text" level="IndexValues" name="dn"
valueMax="250000" defaultAction="Open"/>
 <key description="Risk Number" format="Float64" level="IndexValues" name="risk.num"
valueMax="250000" defaultAction="Open"/>
 <key description="ReputationResult" format="Text" level="IndexValues"
name="cs.represult" valueMax="250000" defaultAction="Open"/>
 <key description="Module Score" format="Text" level="IndexValues"
name="cs.modulescore" valueMax="250000" defaultAction="Open"/>
 <key description="Module Sign" format="Text" level="IndexValues" name="cs.modulesign"
valueMax="250000" defaultAction="Open"/>
 <key description="opswat result" format="Text" level="IndexValues"
name="cs.opswatresult" valueMax="250000" defaultAction="Open"/>
 <key description="source module" format="Text" level="IndexValues"
name="cs.sourcemodule" valueMax="250000" defaultAction="Open"/>
 <key description="Target Module" format="Text" level="IndexValues"
name="cs.targetmodule" valueMax="250000" defaultAction="Open"/>
 <key description="yara result" format="Text" level="IndexValues" name="cs.yararesult"
valueMax="250000" defaultAction="Open"/>
 <key description="Protocol" format="Text" level="IndexValues" name="protocol"
valueMax="250000" defaultAction="Open"/>
 <key description="Event Time" format="TimeT" level="IndexValues" name="event.time"
valueMax="250000" defaultAction="Open"/>
 <key description="Source Host" format="Text" level="IndexValues" name="host.src"
```

```
valueMax="250000" defaultAction="Open"/>
 <key description="Start Time" format="TimeT" level="IndexValues" name="starttime"
valueMax="250000" defaultAction="Open"/>
 <key description="Timezone" format="Text" level="IndexValues" name="timezone"
valueMax="250000" defaultAction="Open"/>
 <key description="Received Bytes" format="UInt64" level="IndexValues" name="rbytes"
valueMax="250000" defaultAction="Open"/>
 <key description="Agent User" format="Text" level="IndexValues" name="user.agent"
valueMax="250000" defaultAction="Open"/>
 <key description="Source Bytes" format="UInt64" level="IndexValues" name="bytes.src"
valueMax="250000" defaultAction="Open"/>
 <key description="Strans Address" format="Text" level="IndexValues" name="stransaddr"
valueMax="250000" defaultAction="Open"/>
 </language>
```

Result

Analysts can:

- Create NetWitness alerts based on NetWitness Endpoint events by configuring NetWitness Endpoint events as an enrichment source.
- Create ESA rules using NetWitness Endpoint meta as described in the "Add Rules to the Rules Library" topic in the *Alerting Using ESA Guide*.
- Report on NetWitness Endpoint events using NetWitness Endpoint meta as described in the "Configure a Rule" topic in the *Reporting Guide*.
- View NetWitness Endpoint alerts in NetWitness Respond as described in the "View Alerts" topic in *NetWitness Respond User Guide*.
- View NetWitness Endpoint meta keys in Investigation along with standard NetWitness core meta keys as described in the "Conduct an Investigation" topic in *Investigation and Malware Analysis User Guide*.

NetWitness® Platform XDR

Version 12.0

NetWitness Endpoint Quick Start Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

What is NetWitness Endpoint?

NetWitness Endpoint is an endpoint detection and response tool that continuously monitors the behavior of all endpoints in the network to provide deep visibility and analysis of executables and processes. It helps to detect new, unknown, and targeted attacks, highlights suspicious activity for investigation, exposes anomalous behaviors, and determines the scope of compromise to help analysts respond to advanced threats faster.

About this Guide

This guide provides end-to-end instructions to configure NetWitness Endpoint and to use Endpoint features.

NetWitness Platform Documentation in NetWitness Community Portal

NetWitness Platform product documentation is organized along functional lines. If you are looking for a specific guide or version, go to the a.

Use these links to view the NetWitness Platform documentation. Both links provide the same documentation, in these two formats:

- HTML Guides include the latest information for currently supported versions: NetWitness Platform Documentation.
- PDF Guides provide the information for a specific version: NetWitness Platform PDFs

Use these links to access documentation that is not related to a particular version of the software:

- Hardware setup guides: https://community.netwitness.com/t5/netwitness-platform-hardware/tkb-p/netwitness-hardware-documentation
- Documentation for Content such as feeds, parsers, application rules, and reports: https://community.netwitness.com/t5/netwitness-platform-threat/tkb-p/netwitness-threat-intelligence.

Getting Started

The following tasks can be performed in any sequence.

Description	References
	

Description	References
View information about product updates, improvements, and known issues.	*Release Notes*
Understand NetWitness Endpoint.	"Getting Started with NetWitness Platform" and "Investigate" in the *NetWitness Platform Getting Started Guide*

Setup and Installation

Fresh Installation

The following tasks must be performed in the sequence listed.

Description	References
	
Obtain a license for Endpoint Log Hybrid.	Licensing Management Guide
Review the supported hardware.	"Supported Hardware" in the *Physical Host Installation Guide*
Review the Endpoint architecture; plan your deployment based on the number of endpoints, distribution, and location of these endpoints; and choose one of the following deployments: • Single Endpoint server • Multiple Endpoint server	"NetWitness Endpoint Architecture" in the *Deployment Guide*
Configure the ports on your firewall.	"Network Architecture and Ports" in the *Deployment Guide*
Install NetWitness Server and other components. For a single Endpoint server deployment, you need to install - NetWitness Server, Endpoint Log Hybrid, and ESA. For a multiple Endpoint server, in addition to the above components, you need to install - Additional Endpoint Log Hybrid, NetWitness Broker along with Endpoint Broker installed on it.	- *Physical Host Installation Guide* for instructions on how to set up physical hosts - *Virtual Host Installation Guide* for instructions on how to set up virtual hosts

Description	References
Install Endpoint Log Hybrid.	"NetWitness Endpoint" in the *Physical Host Installation Guide*
Review the services installed.	*Hosts and Services Getting Started Guide*
Note: Review the default policies and modify them accordingly. Install Endpoint agent on hosts.	"Endpoint Sources" topic in the *Endpoint Configuration Guide* *NetWitness Endpoint Agent Installation Guide*

Upgrade

The following tasks must be performed in the sequence listed.

Description	References
	
Install the 11.5 Relay Server	"Install the 11.5 Relay Server" in the Upgrade Guide
Upgrade Endpoint Agents	"Upgrade Agents" in the *Endpoint Agent Installation Guide*

Configuration

The following tasks can be performed in any sequence.

Description	References
	
Understand NetWitness Endpoint and high-level tasks required for configuration.	"NetWitness Endpoint Overview and Endpoint Server Configuration" in the *Endpoint Configuration Guide*
Review Groups and Policies for agents.	"Endpoint Sources" topic in the *Endpoint Configuration Guide*

Description	References
Set up the RSA Live account and verify if the ESA content and Application Rules for Endpoint are available. **Note:** The file reputation service is automatically enabled on RSA Live.	*Live Services Management Guide*
Create role-based access control (RBAC).	"Role Permissions" in the System Security and User Management Guide
Configure data retention policy.	"Configure Data Retention" in the *Endpoint Configuration Guide*
Manage inactive agents.	"Manage Inactive Agents" in the *Endpoint Configuration Guide*
Configure Memory Dumps and MFT retention policy.	"Configure Retention Policy for Memory Dumps and MFT" in the *Endpoint Configuration Guide*

Investigation

The following tasks can be performed in any sequence.

Description	References
	
Understand how investigation works.	"How NetWitness Investigate Works" in the *NetWitness Investigate User Guide*
Configure investigate views.	"Configuring NetWitness Investigate Views and Preferences" in the *NetWitness Investigate User Guide*
Begin an investigation in different Investigate views.	"Beginning an Investigation" in the *NetWitness Investigate User Guide*
Review best practices for files and hosts and set up your Investigate view for investigation.	"Best Practices" under Investigating Files and Investigating Hosts in the*NetWitness Endpoint User Guide*
Investigate files.	"Investigating Files" in the *NetWitness Endpoint User Guide*
Investigate hosts.	"Investigating Hosts" in the *NetWitness Endpoint User Guide*

Description	References
Investigate process.	"Investigating Hosts" in the *NetWitness Endpoint User Guide*
Analyze downloaded files.	"Analyzing Downloaded Files" in the *NetWitness Endpoint User Guide*
Change file status and remediate.	"Changing File Status or Remediate" in the *NetWitness Endpoint User Guide*
Analyze events.	"Analyzing Events" in the *NetWitness Endpoint User Guide* "Analyzing Raw Data and Metadata in the Events View", "Investigating Metadata in the Navigate View", and "Examining Raw Events in the Events View" in the *NetWitness Investigate User Guide*
Perform host forensics.	"Performing Host Forensics" in the *NetWitness Endpoint User Guide*
Isolate hosts from network.	"Isolating Hosts from Network" in the *NetWitness Endpoint User Guide*
Investigate using REST APIs	*API User Guide*

Respond and Reporting

The following tasks can be performed in any sequence.

Description	References
	
Respond to Endpoint incidents.	*NetWitness Respond User Guide*
View reports related to Endpoint data.	*Reporting User Guide*

Maintenance

The following tasks can be performed in any sequence.

Description	References
	
Monitor health and wellness.	*System Maintenance Guide*

Integration (for Legacy NetWitness Endpoint)

The following tasks can be performed in any sequence.

Description	References
	
Configure NetWitness Endpoint 4.4.x metadata with NetWitness.	"Integrating NetWitness Endpoint 4.4.0.2 or Later with NetWitness Platform" topic in the *Endpoint Configuration Guide*
Configure integrated operation of NetWitness Endpoint 4.4.x with NetWitness.	NetWitness Endpoint Integration Guide

NetWitness® Platform XDR

Version 12.0

NetWitness Endpoint User Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

November, 2022

Contents

Introduction to Endpoint Investigation **7**
- Endpoint Metadata 7
- Risk Score 9
  - Severity of Alerts 11
  - Global and Local Risk Score 12
  - Automated Incident Creation Based on Risk Score 13
- File Reputation 13
- File Status 13
- Remediation 14
- Network Isolation 14

Focusing on Endpoint Analysis **15**

Investigating Files **17**
- Best Practices 17
- View Files 18
- Filter Files 19
- Add and Sort Columns in the Table 20
- Analyze Files Using the Risk Score 21
- Analyze Hosts with File Activity 23
- Analyze Files Using YARA 23
- Analyze Files Using OPSWAT 25
  - Manual Scan - Scan selected files 26
  - Manual Scan - Scan all files 27
- Launch an External Lookup for a File 29
- Set Files Preference 30
- Export Global Files 30
- Analyze Certificates 31
  - Change the Certificate Status 32
  - Filter Certificates 33
- Resetting Risk Score of Files 34

Investigating Hosts **36**
- Best Practices 36
- View Hosts 37
- Manage Hosts Using Tags 38
  - Manage Tags 38
    - To Create Tags: 38

To Delete Tags: 40
Assign tags 41
Create and Assign Tags When Generating the Agent Packager 43
Unassign tags 46
What happens next after unassigning or deleting tags from hosts? 47
View Agent History 47
Filter Hosts 48
Adding and Sorting Columns in the Table 54
Scan Hosts 55
Standalone Scan on Air-gapped Windows Hosts 57
Standalone scan workflow: 58
Generate the scan configuration file 58
Install Endpoint Agent and Register for Standalone Scan 59
Start a Standalone scan 59
Upload the Standalone scan result file 60
Analyze Hosts Using the Risk Score 60
Analyze Host Details 62
Filter Host Details 63
Search Files on Host 63
Analyze Processes 65
Analyze Autoruns 67
Analyze Files 67
Analyze Libraries 68
Analyze Drivers 68
Analyze Anomalies 68
Analyze System Information 69
Analyze History 69
Export Host Details or Files to JSON File 70
Launch an External Lookup for a File 70
Delete a Host 71
Deleting Hosts with Older Agent Versions 72
Set Hosts Preference 72
Export Host Attributes 73
Migrate Hosts 74
Analyzing Risky Users 74
Resetting Risk Score of Hosts 75
Investigating a Process 77
Best Practices 77
Analyze a Process 78
Analyze Events for a Process 81

Changing File Status or Remediate **83**
Import File Hashes using the Block Hash tool 85
JSON File Format 85
Analyzing Downloaded Files **88**
Download Files to Server 88
Save Downloaded Files 90
Analyze Downloaded Files 90
Performing Host Forensics **93**
Download Master File Table 93
Analyze Downloaded MFT 95
System and Process Memory Dump 97
Download Files Using Full Path or Wildcard 99
Filter Downloaded Files 102
Save Downloaded File 102
Delete Downloaded Files 103
Analyzing Events **104**
Analyze Events from Files View 104
Analyze Events from Hosts View 104
Text Analysis for an Endpoint Event 105
Isolating Hosts from Network **107**
Edit Exclusion List 108
Release Isolated Hosts 109
NetWitness Endpoint with Third-Party Antivirus Products **111**
Troubleshooting NetWitness Endpoint **112**
General Issues 112
Multi-server Issue 113
Hosts View Issues 113
Files View Issues 114
Policy Issue 114
Driver Issue 115
Download Issue 115
File Reputation Service Issue 115
Risk Scoring for Hosts or Files Issue 115
Endpoint Broker/Server Issue 116
NetWitness Endpoint Reference Materials **118**
Files View 119
File Details View 121
Hosts View 123
Hosts View - Details Tab 127

Hosts View - Process Tab 131
Process Details 135
Hosts View - Autoruns Tab 136
Hosts View - Files Tab 140
Hosts View - Drivers Tab 144
Hosts View - Libraries Tab 148
Hosts View - Anomalies Tab 152
Image Hooks 156
Kernel Hooks 156
Suspicious Threads 157
Registry Discrepancies 157
Hosts View - Downloads Tab 158
MFT Viewer 162
Hosts View - System Information Tab 165
System Information Panel 167
Hosts View - Agent History Tab 169

Introduction to Endpoint Investigation

NetWitness Investigate provides data analysis capabilities in NetWitness, so that analysts can analyze packet, log, endpoint, and UEBA data, and identify possible internal or external threats to security and the IP infrastructure. This guide helps analysts perform investigations of endpoint data using NetWitness Investigate.

Note: In Version 11.1 and later, the Hosts and Files views provide a view into endpoint data. Earlier versions offer access to endpoint data using a standalone NetWitness Endpoint server.

For more information, see the *NetWitness Endpoint Quick Start Guide*, the *NetWitness Investigate Quick Start Guide*, and the *NetWitness Investigate User Guide*.

Endpoint Metadata

Endpoint metadata is generated when hosts are scanned and when there are real-time activities on the hosts. You can view the following categories of sessions when metadata forwarding is enabled:

Operating System	Scan Categories & Real-time events	Tracking Categories

Windows	file, service, dll, process, task, autorun, machine, kernel hook, image hook, registry discrepancies, suspicious threads and removable device(USB) detection	• Process event - Reports any process related activities, such as `openprocess`, `openosprocess`, `createprocess`, `createremotethread`, `openbrowserprocess`.
		• File event - Reports any file related activities by an executable, such as `readdocument`, `writetoexecutable`, `renameexecutable`, `selfdeleteexecutable`, `openphysicaldrive`.
		• Registry event - Reports activities that result in registry creation or modification, such as `modifyservicesimagepath`, `modifyfirewallpolicy`, `createservicesimagepath`, `createsecuritycenterconfiguration`, `modifybadcertificatewarningsetting`, `Modifies Startup Folder Location`, `Modifies Winlogon Registry Settings`, `Registers Time Provider Dll`, `Registers Port Monitor Dll`, `Registers Netsh helper Dll`, `Registers AppInit Dll`, `Registers AppCert Dll`,
		• System event - Reports connection of removable devices(USB devices), IP change and boot events such as, `removableDeviceConnected`, `removableDeviceDisconnected`
		• Network event - TCP/UDP and incoming/outgoing.
		  • Reports outbound and inbound network connections on all supported Windows platforms.
		  • Reports IPv4 and IPv6 connections.
		• Console event (for Windows 8 and later) - User input that is entered into a console application, such as `cmd.exe`, `powershell.exe`, is captured and reported with the context `console.local`.
		Commands executed by `cmd.exe`, `powershell.exe` as a result of inter-process communication through anonymous pipes are captured and reported with the context `console.remote`.

Operating System	Scan Categories & Real-time events	Tracking Categories
		For example, `Get-Item -Path Registry::HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion`.
Linux	file, autrorun, loaded library, systemd, process, cron, initd, and machine	-

Mac	file, daemon, process, task, dylib, autorun, and machine	• Process event - Reports any process related activities, such as `openprocess`, `createprocess`, `openosprocess`, `openbrowserprocess`, `allocateremotememory`, `createremotethread`.
		• File event - Reports any file related activities by an executable, such as `writetoexecutable`, `renameexecutable`, `createautorun`, `deleteexecutable`, `selfdeleteexecutable`, `writetoplist`, `writetosudoers`, `createbrowserextension`.
		• Network event - TCP/UDP and incoming/outgoing.
		  • Reports outbound and inbound network connections on all supported Mac operating system.
		  • Reports IPv4 and IPv6 connections.

For more information on metadata, meta keys, meta values, and meta entities, see the *NetWitness Investigate User Guide*.

Risk Score

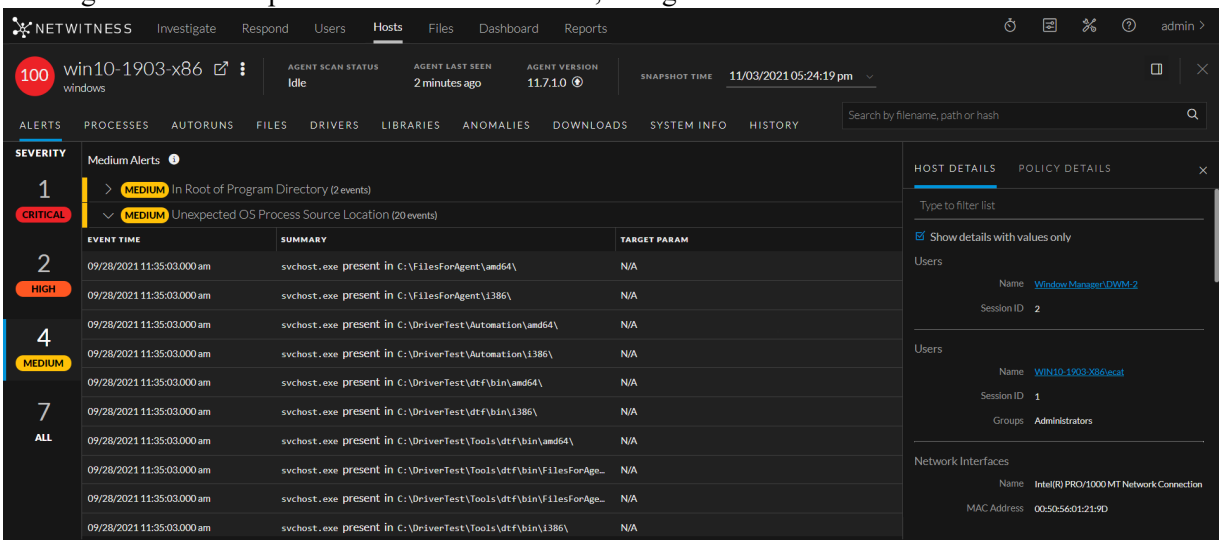
Analysts can use the risk score to begin an investigation on hosts and files. NetWitness uses a proprietary algorithm to calculate the risk scores ranging from 0 to 100. A subset of alerts associated with hosts and files contribute to the risk score calculation. Analysts can review critical and high alerts associated with a risk score to identify strong evidence of malicious activity and take required action.

Note: If you have an Insights agent, you can view the risk score for files but not for hosts. To view the risk score for hosts, upgrade to the Advanced agent. For more information, see the *NetWitness Endpoint Configuration Guide*.

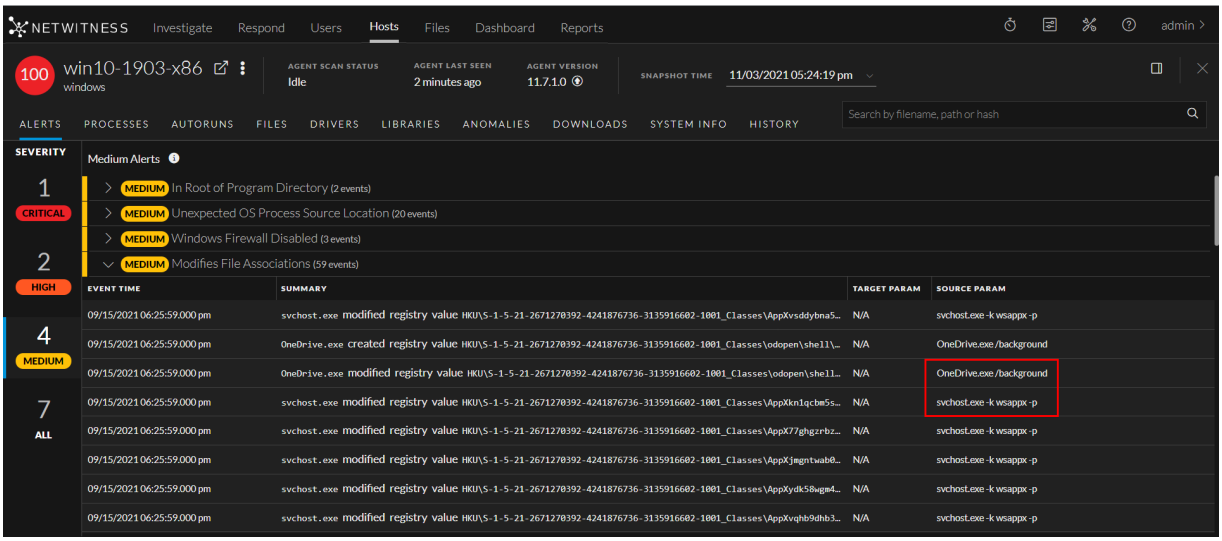
The following factors contribute to the risk score:

- **Distinct Alerts.** Any host or file activities that are suspicious or malicious generate alerts. Only the distinct alerts are used for risk score calculation.
- **Severity of Alerts.** Severity of alerts, such as critical, high, and medium.

  This figure is an example of a host with 1 Critical, 2 High and 4 Medium distinct alerts.

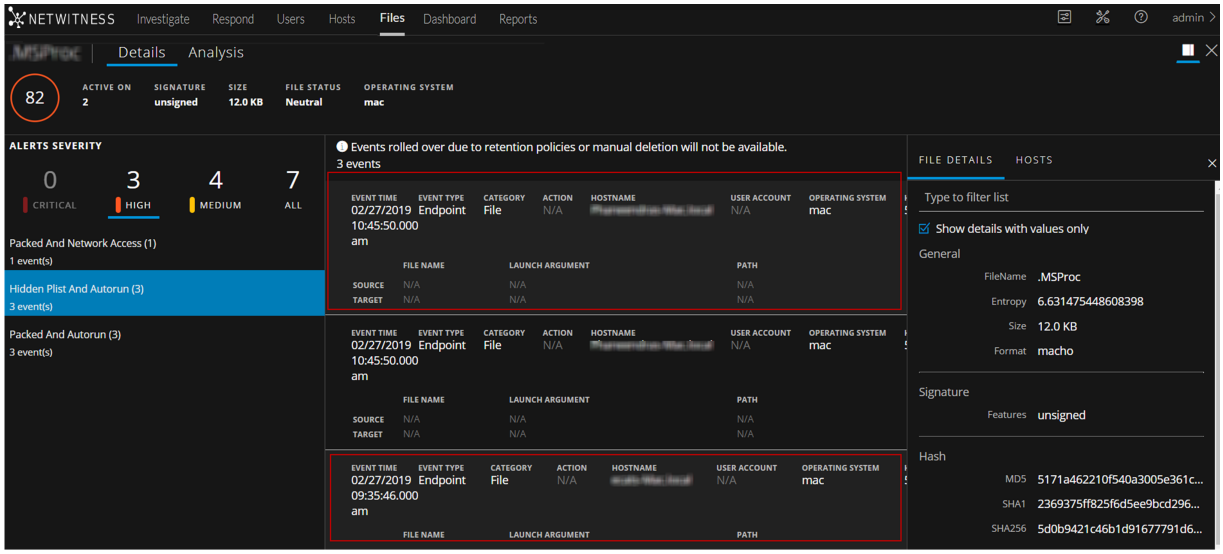


All the distinct alert shown in the above example can be for the same file or different files. For example, Modifies File Associations alert is triggered for files, such as `svchost.exe` and `OneDrive.exe`.



This figure is an example of files with distinct alerts. Each file can have a multiple distinct alerts. The

files can have a same alert name being triggered by two different hosts as shown below.



The risk score is reset when you perform any of the following actions:

- Whitelist or blacklist a file after investigation. The risk score of a file is set to 0 on whitelisting and set to 100 on blacklisting.
- If the alerts or events triggered by the host or files on the host are false positive, you make changes to the Endpoint Application rules or ESA rules and reset the risk score.

Besides the above factors, the risk score is reset when a file no longer matches Yara rules in the subsequent scans

Note: When you whitelist a file or reset the risk score, the alerts that contributed to the risk score are not shown in the Host Details tab.

The host risk score depends on the risk score of all the files on the host. When you change the file status or reset the file risk score, the host risk score is recalculated. For example, the score for all the hosts on which a blacklisted file is present is recalculated and becomes 100. If the host is not found to be infected, you can reset the host risk score. This deletes the alerts contributed to the risk score and does not impact the global file score. For more information on changing the file status, see Changing File Status or Remediate.

Note: For the risk score calculation, the ESA Correlation server must be configured with an Endpoint Concentrator. The application rules are automatically deployed on installation. For an upgrade, you must deploy the application rules from RSA Live. For more information, see the *NetWitness Endpoint Configuration Guide*.

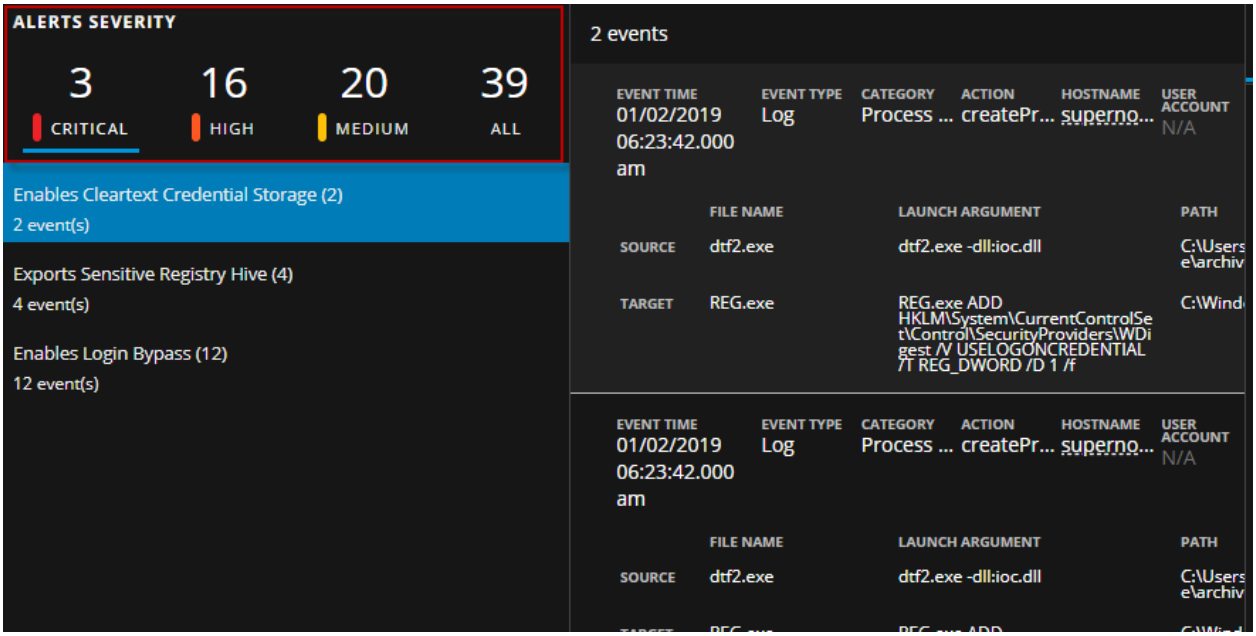
Note: For the accurate risk score calculation, the default multi-valued meta keys are required on the ESA Correlation service. For more information, see "Configure Meta Keys as Arrays in ESA Correlation Rule Values" section in the *ESA Configuration Guide*.

Severity of Alerts

The following table depicts the risk score range based on the associated alert severity:

Severity	Color	Risk Score Range
Critical	Red	100
High	Orange	70-99
Medium	Yellow	31-69
Low	Green	0-30

The following is an example of alerts contributing to the risk score:



In the above example, there are three distinct critical alerts. For each alert type, associated events are displayed. You can see that the "Enables Cleartext Credential Storage" alert was triggered twice. The details of the two events are displayed with the metadata information. For more information on severity alerts and metadata information, see Analyze Hosts Using the Risk Score and Analyze Files Using the Risk Score.

Global and Local Risk Score

Analysts can get better context on file activities on hosts using the global risk score and the local risk score of a file.

Global Risk Score - The global risk score is an aggregate of all suspicious and malicious activities performed by the file across all hosts. This score indicates the potential threat posed by the file across the NetWitness Platform.

Local Risk Score - The local risk score is calculated on suspicious or malicious activities performed by the file on a specific host. The local risk score is used for the host risk score calculation.

For more information on the global and local risk score, see Investigating Files and Investigating Hosts.

Automated Incident Creation Based on Risk Score

By default, a threshold is set for the risk score to control the generation of incidents and alerts in NetWitness Respond. For more information on configuring the threshold limit, see the *NetWitness Respond Configuration Guide*.

File Reputation

The File Reputation service available on RSA Live checks the reputation of every file hash against an extensive database of known file hashes updated in real-time. The file reputation is displayed on the Investigate and Respond views.

The reputations for a file hash are:

Reputation	Description
Malicious	File hash is labeled as malicious.
Suspicious	File hash is suspected to be malicious.
Unknown	File hash is not known.
Known	File hash information is known to the file reputation service and does not have any previous bad record.
Known Good	File hash information is known good, such as files signed by Microsoft or RSA.
Invalid	File hash format is invalid.

The suspicious or malicious files are available for further analysis in the **Investigate > Navigate** view and **Investigate > Events** view. For more information on the file reputation service, see the *Live Services Management Guide*.

Note: The File Reputation service supports maximum of 10 million files for a reputation of file hash.

File Status

To help analysts triage and focus on their investigation, NetWitness provides capabilities to manage suspect and legitimate files. For example, you can whitelist files that are legitimate (such as security products), or blacklist files based on known threats and investigation.

A file can be classified as follows:

- Blacklist: File that is marked suspicious, such as when ransomware is found by scan.
- Graylist: File that is marked for a later review.
- Whitelist: File that is legitimate and is not to be considered for risk scoring.
- Neutral: Default status.

For more information, see Changing File Status or Remediate.

Remediation

If a file is malicious or infected, you can block the file to prevent future execution on any host. Remediation helps to:

- Stop or reduce the spread of identified malware, such as viruses, trojans, rootkits, worms, spyware, and adware.
- Identify attempted breach points to aid in deeper analysis; all events are time-stamped allowing analysts to trace backward to identify the entry point.
- Remove unwanted software, such as adware, which can potentially mask real malware.
- Stop all actions possible by the loader.

You can block files with the following file extension: EXE, COM, SYS, DLL, SCR, OCX, BAT, PS1, VBS, VBE, and VB. For more information, see Changing File Status or Remediate.

Network Isolation

If you suspect that a host is potentially compromised with the threat still being active, you can isolate the host from the network and safely investigate possible threats within the host. By isolating the host, you can control the spread of an attack and analyze the malware behavior. When a host is isolated, only connection to the following IP addresses are allowed:

- Endpoint Server, Relay Server, DNS, DHCP, Gateways, 0.0.0.0, and 255.255.255.255.
- Other IP addresses that you include in the exclusion list.

In the isolated state, all events are reported to the Endpoint Server retaining full visibility into activities on the host. You can continue investigation by requesting scans, downloading MFT, files, and so on. The following metadata is added to the network events:

- network.isolated - indicates that the host is isolated.
- network.connectallowed - indicates that the network connection is allowed as the IP address is included in the exclusion list.
- network.connectblocked - indicates that the network connection is blocked.

Note: If the agent is enabled for log or file collection, make sure that you add the Log Decoder IP addresses in the exclusion list while you isolate the host.

For more information, see Isolating Hosts from Network.

Focusing on Endpoint Analysis

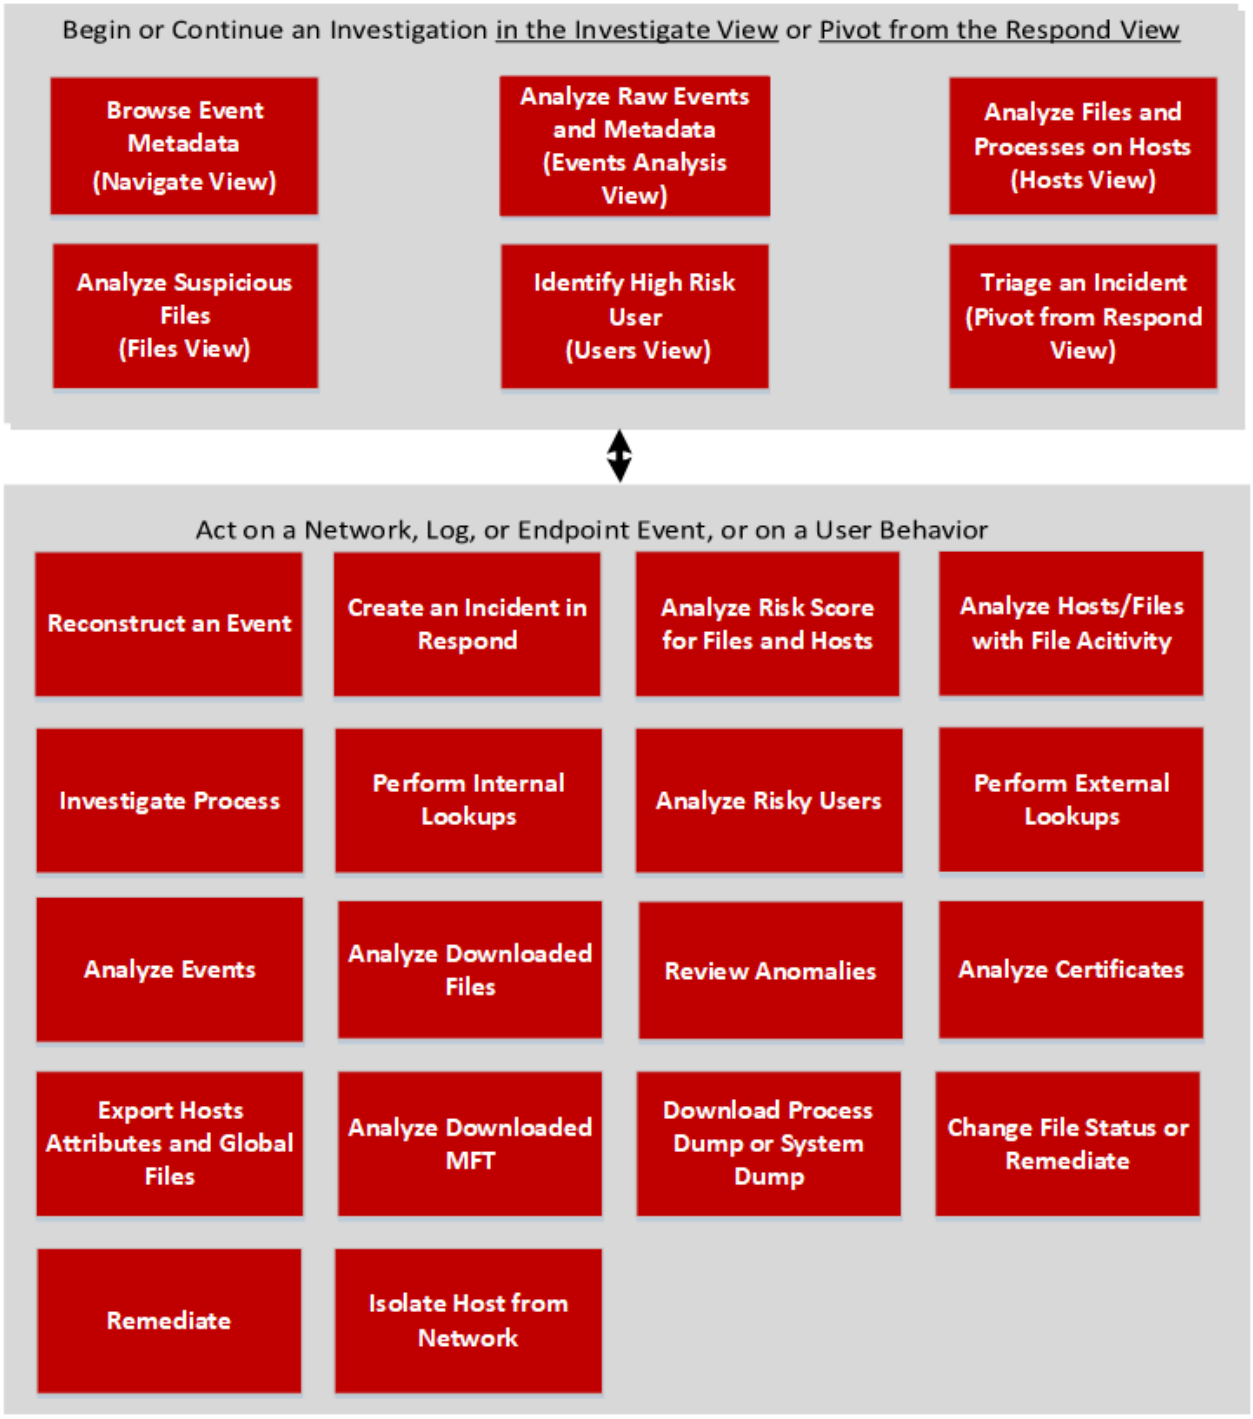
This guide provides the information needed to conduct an investigation that is focused on endpoint data from configured hosts. Analysts who conduct analysis using Investigate need to have the appropriate system roles and permissions set up for their user accounts. An administrator must configure roles and permissions as described in Roles and Permissions for Endpoint Analysts. For more information on roles and permissions, see the *System Security and User Management Guide*.

To hunt for information on hosts that have the agent running, begin the investigation in the Hosts view (**Hosts**). For every host, you can see processes, drivers, DLLs, files (executables), services, anomalies, and autoruns that are running, and information related to logged-in users. (See Investigating Hosts.)

You can begin the investigation on files in your deployment in the Files view (**Files**). (See Investigating Files.)

Note: To access the Hosts and Files views, you must have the `endpoint-server.filter.manage` permission.

Analysts use the Hosts and Files views to investigate or perform analysis on hosts or files using attributes such as IP address, host name, Mac address, risk score, and so on. This figure shows the high-level capabilities of an endpoint investigation. The top box are all the possible starting points, and the lower box shows the tasks that you can accomplish from different starting points.


Begin or Continue an Investigation in the Investigate View or Pivot from the Respond View
Browse Event Metadata (Navigate View)
Analyze Raw Events and Metadata (Events Analysis View)
Analyze Files and Processes on Hosts (Hosts View)
Analyze Suspicious Files (Files View)
Identify High Risk User (Users View)
Triage an Incident (Pivot from Respond View)
Act on a Network, Log, or Endpoint Event, or on a User Behavior
Reconstruct an Event
Create an Incident in Respond
Analyze Risk Score for Files and Hosts
Analyze Hosts/Files with File Acitivity
Investigate Process
Perform Internal Lookups
Analyze Risky Users
Perform External Lookups
Analyze Events
Analyze Downloaded Files
Review Anomalies
Analyze Certificates
Export Hosts Attributes and Global Files
Analyze Downloaded MFT
Download Process Dump or System Dump
Change File Status or Remediate
Remediate
Isolate Host from Network

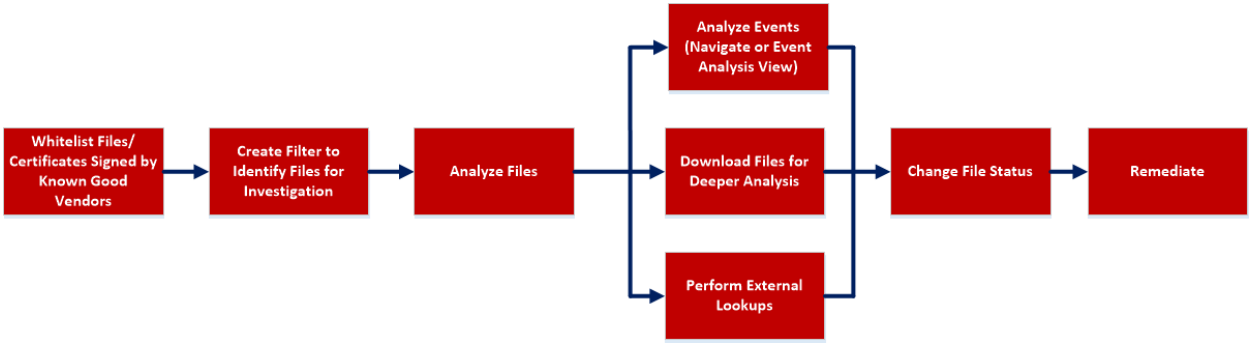
Investigating Files

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The Files view provides a holistic view of all files in your NetWitness platform. You can apply various filters, sort, and categorize files into different status to reduce the number of files for analysis and identify suspicious or malicious files.

Best Practices

The following are some best practices and tips that may help you investigate efficiently to identify and isolate threats or attacks:



- Whitelist all files signed by RSA, Microsoft, and any other known good vendors. Use the filters to list the files and change the status of all these files to whitelist. For more information, see Filter Files and Changing File Status or Remediate.

  Note: Some Microsoft signed files are restricted from whitelisting as there is a potential risk of them being used for malicious purposes. To view the list, see Files Restricted from Whitelisting.

- Change the status of certificate and the associated files automatically. For more information, see Analyze Certificates.
- Filter to exclude whitelisted, files with valid signature, known good files based on reputation status. For more information, see Filter Files.
- Lookup Google or VirusTotal with the filename or hash to get more information about a suspected file. For more information, see Launch an External Lookup for a File.
- Analyze the files using one or more of these indicators:

  a. Risk score - Displays the risk score for a file. Analysts can view the associated alerts and events for further investigation. For more information, see Analyze Files Using the Risk Score.

  b. On Hosts - Indicates the number of hosts on which a file exist. If a file is present on fewer hosts with a high risk score, it may require further investigation. You can also sort or filter using On Hosts column to narrow down the search during investigation. For more information, see Analyze Hosts with File Activity.

c. File status - To manage suspected and legitimate files, analysts can use the file status to manage. For more information on the various file status, see Changing File Status or Remediate.

d. Reputation status - Indicates the reputation of a file hash for analyst to narrow-down the files to investigate. For more information, see File Reputation.

e. Signature - A valid signature on a file signed by a trusted vendor, such as Microsoft and Apple indicates that the file is not a risk. If a file is unsigned, it may be malicious, and needs investigation.

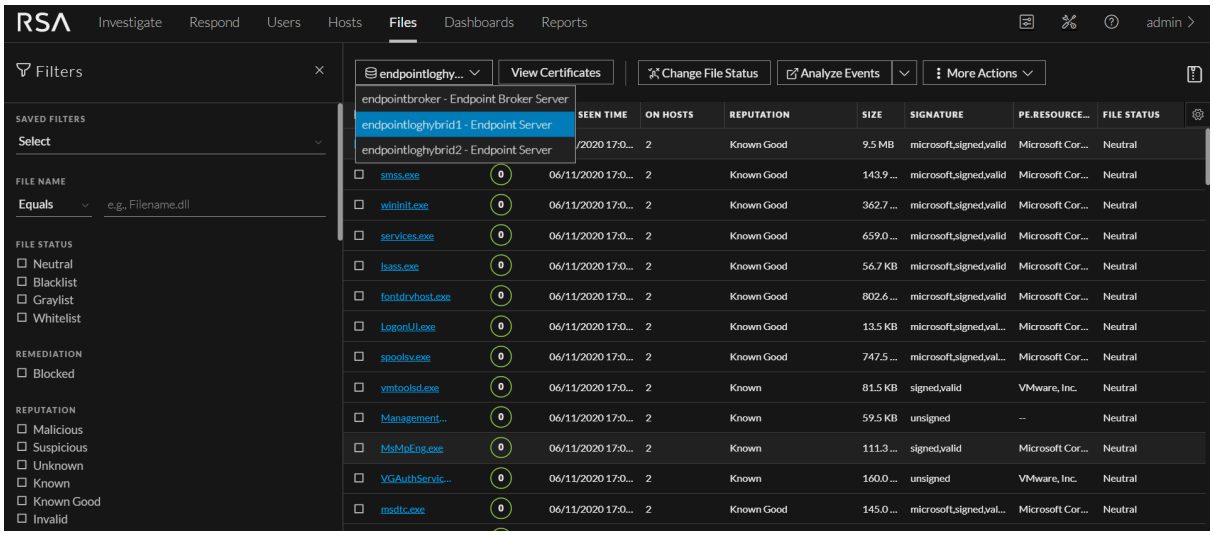
f. File name - Many trojans write random file names when dropping their payloads to prevent an easy search across the hosts in the network based on the filename. For example, if a file is named `svch0st.exe`, `scvhost.exe`, or `svchosts.exe`, it indicates that the legitimate Windows file named `svchost.exe` is being mimicked.

- Investigate a particular file name or hash by pivoting to Navigate or Events view to view context, file activity on different hosts, and any file transfers across the network through packet data. For more information, see Analyzing Events.
- Investigate files using a rule-based detection technique. YARA helps to identify threats effectively using easy-to-create malware descriptions called YARA rules. For more information, see Analyze Files Using YARA.
- Download suspicious files to the server for deeper analysis. For more information, see Analyzing Downloaded Files.
- Change the status of the file (blacklist or graylist), and block an infected or malicious file. For more information, see Changing File Status or Remediate.

View Files

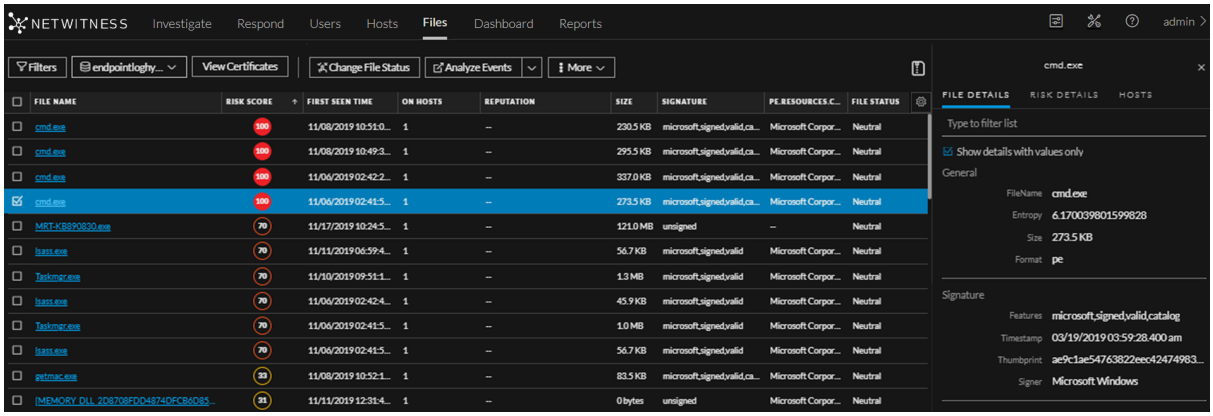
You can view all files present on a specific Endpoint server or consolidated list of all files on multiple Endpoint servers using the Endpoint Broker for analysis. To view files:

1. Go to **Files**.
2. Select one of the following:



- **Endpoint Broker Server** to view all files across all Endpoint servers.
- **Endpoint Server** to view files on a specific Endpoint server.

3. Select the file that you want to analyze.
4. Click a row to view the following details:

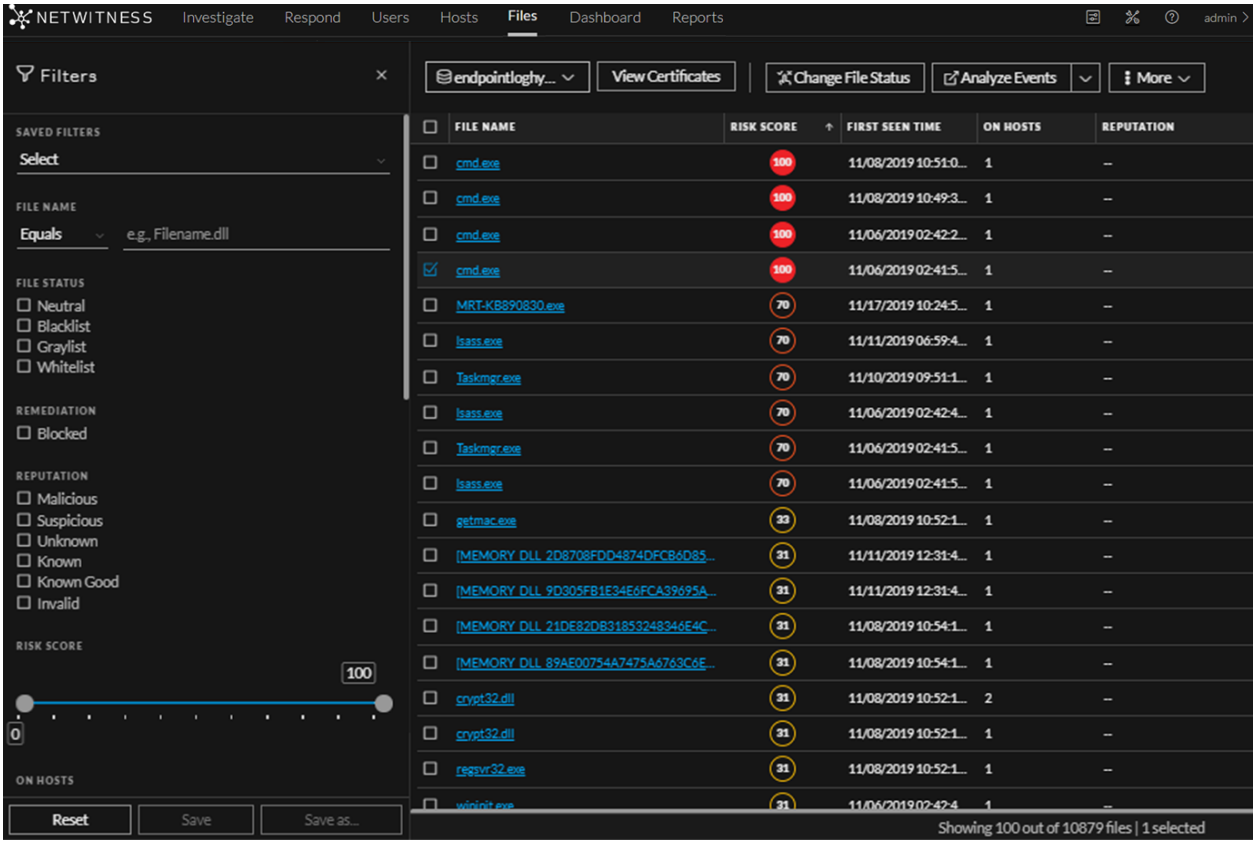


- **File Details** displays the file information. For more information, see Launch an External Lookup for a File.
- **Risk Details** displays the distinct alerts associated with the risk score. For more information, see Analyze Files Using the Risk Score.
- **Hosts** displays the number of hosts on which file activities are present. For more information, see Analyze Hosts with File Activity.

Filter Files

You can narrow down the investigation by filtering files using file name, on hosts, file status, risk score, remediation, reputation status, operating system, size, entropy, format, signature, company name, checksum (MD5 and SHA256), downloaded status, and YARA rules.

Note: While filtering on a large data set, use at least one indexed field with the `Equals` operator for better performance. The following fields are indexed in the database - Filename, MD5, SHA256, Operating System, First Seen Time, Format, File Status, On Host, and Reputation Status.



Select the parameters in the Filters tab. Click **Save** to save the search and provide a name (up to 250 alphanumeric characters). The filter is added to the Saved Filters list. To delete a filter, hover over the name and click .

Note: Special characters are not allowed in the filter name except underscore (_) and hyphen (-) while saving the filter.

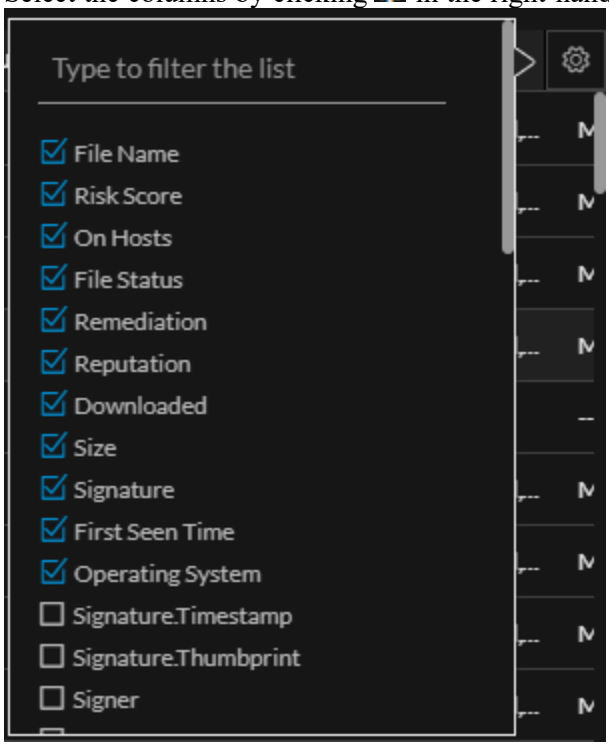
For example, to filter the files based on file reputation, select the reputation status in the Filter panel.

Note: For the file size, 1 KB is calculated as 1024 bytes. For example, if the actual size of the file is 8421 bytes, the UI will display it as 8.2 KB instead of 8.22 KB. It is recommended to search using the bytes format when using the `Equals` operator.

Add and Sort Columns in the Table

By default, the Files view displays a few columns, and files are sorted based on the risk score. To add or remove columns:

1. Go to **Files**.
2. Select the columns by clicking in the right-hand corner.



3. Scroll down or enter the keyword to search and select the required columns.
4. To sort the column in ascending or descending order, click the arrow on the column header.

Analyze Files Using the Risk Score

To analyze files using the risk score:

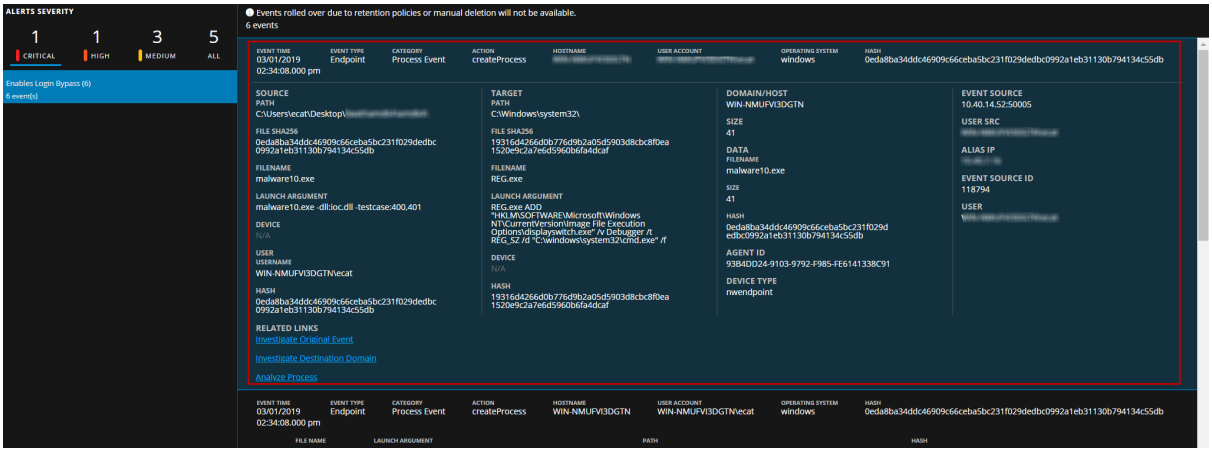
1. Go to **Files**.

   The Files view is displayed.
2. In the Server drop-down list, select the Endpoint server or Endpoint Broker server to view the files.
3. Select the file and do any of the following.
   - Click a row to view the risk associated with the file in the **Risk Details** panel.
   - Click the hostname to investigate the host.

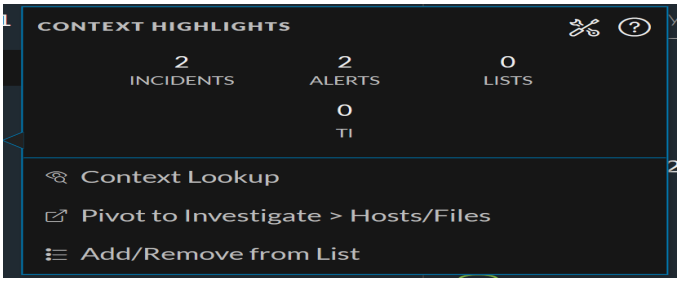
     The Details tab is displayed.
4. In the **Alert Severity** panel, click the alert severity, such as **Critical**, **High**, **Medium**, or **All**.
   The list of distinct alerts is displayed along with the total number of events associated with the alert.
5. Click an alert to view the associated events.

Note: Only the latest 1000 events are displayed.

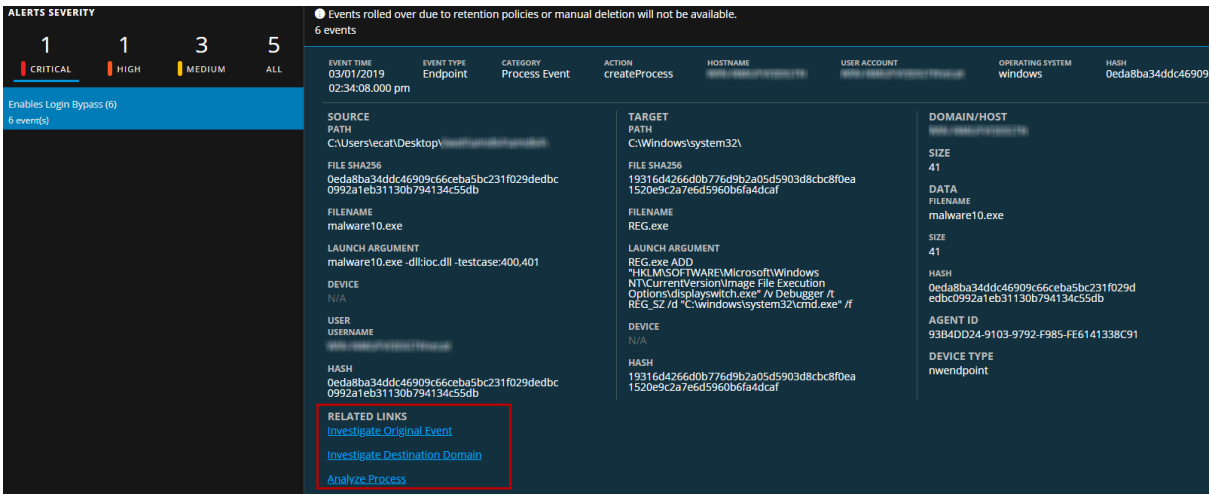
6. To view all the metadata associated with a specific event, click the event header. The information such as source path, target path, filename, and others is displayed.



7. Click on the down arrow or right-click one of the meta values for IP, Hostname, Mac, File name, File hash, User, and Domain to view additional information about the specific metadata. A dialog box displays a list of the data sources that have context data available for meta value. These are the possible data sources: NetWitness Endpoint, Incidents, Alerts, Hosts, Files, and Feeds.



8. To investigate the original event and destination domain of the event, do any of the following:



- To reconstruct an event in a readable form that matches the original, click the **Investigate Original Event** link highlighted in blue. For more information on event reconstruction, see the *NetWitness Investigate User Guide*.
- For details about the elements associated with an event, click the **Investigate Destination Domain** link highlighted in blue. For more information on Contextual Information for an Event, see the *NetWitness Investigate User Guide*.

  Note: Investigate Destination Domain link is not displayed if there is no domain.dst event.
- To view a list of processes captured on the hosts and investigate a particular process, click the **Analyze Process** link highlighted in blue. For more information on process analysis, see Investigating a Process.

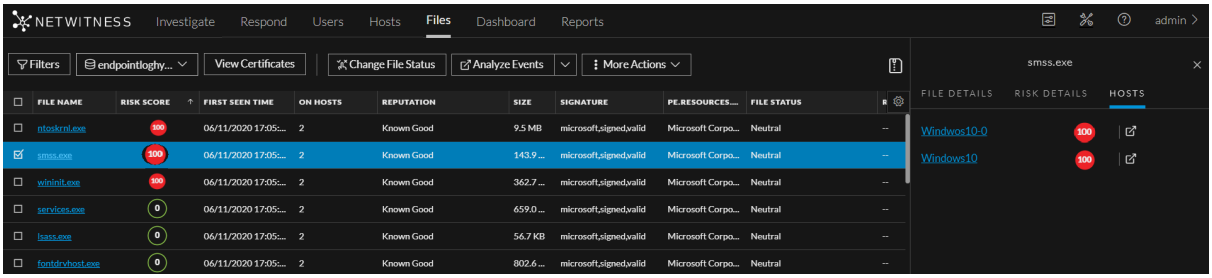
  Note: Analyze Process link is not displayed if there is no createprocess event.

Analyze Hosts with File Activity

To view the list of hosts on which a file exist, do the following:

Note: By default, the system detects the best data source for the On Hosts aggregation. To change the data source, in the Explore view, modify the investigate service ID under `endpoint/investigate`.

1. In the **Files** tab, click the row for the file you want to analyze.
2. In the right panel, click the **Hosts** tab. The list of hosts along with the risk score are displayed.



3. Click the host name to open the host details.
4. Click  to analyze events on the host in the Events view. For more information, see Analyzing Events.

Analyze Files Using YARA

YARA helps analysts with rule-based detection capabilities in identifying and classifying malware. You can easily create malware descriptions using YARA rules. Administrators must enable and configure YARA on the Endpoint server. To learn more about enabling and configuring YARA, refer to NetWitness Endpoint Configuration Guide.

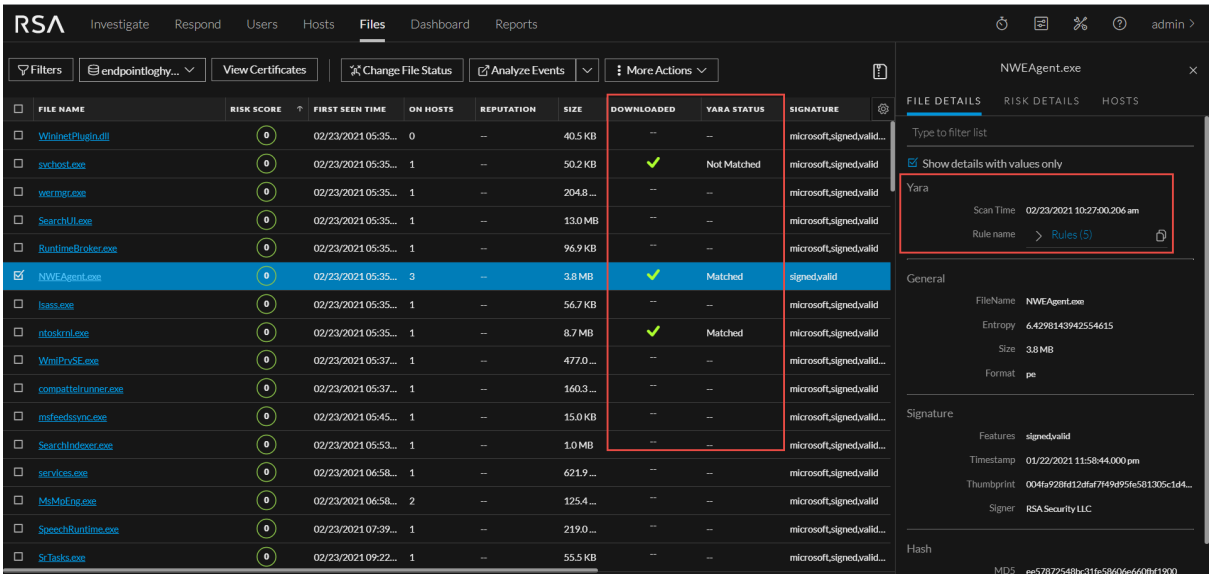
Files must be downloaded to the Endpoint server and YARA scans the downloaded files automatically. The scan results are displayed under **YARA STATUS** on the **Files** tab.

To analyze the scanned files,

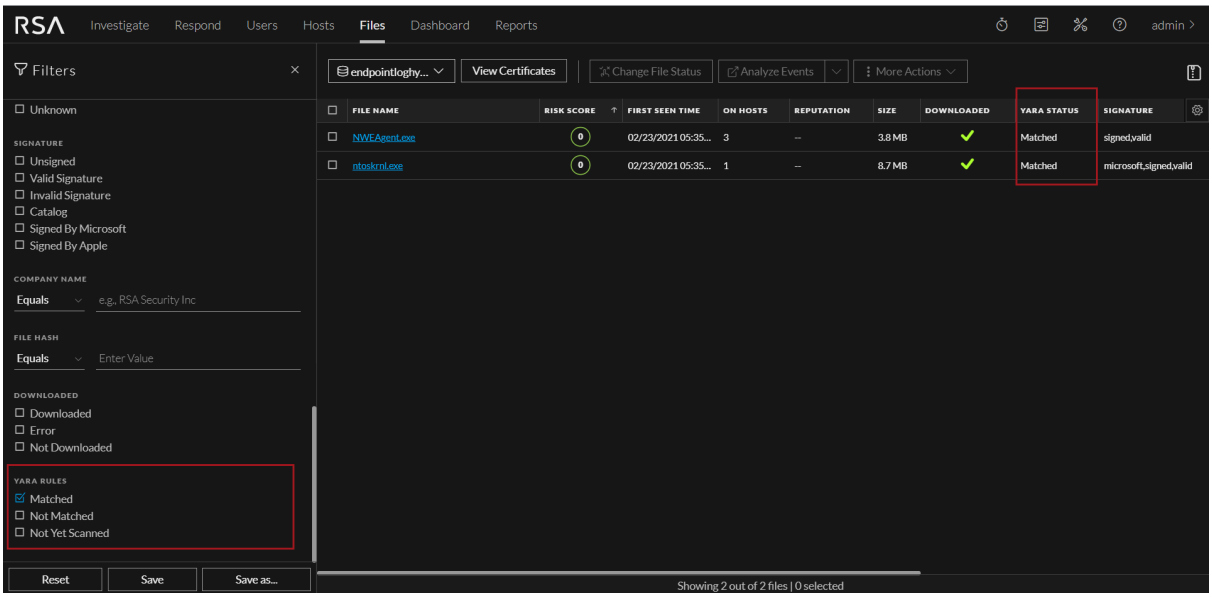
1. Go to **Files**.
2. Select the Endpoint server from the server drop-down , to view files.
3. Select a file that is downloaded to the Endpoint server and do any of the following:
   - Click a row to view the YARA scan details associated with the file in the **File Details** panel.
   - If any file matches one or more YARA rules, the **File Details** panel displays the matching rule names besides scan time.
   - **YARA STATUS** field displays the status of the YARA scan. Following are the available statuses and their definitions
     - **Matched** - The file matches with one or more YARA rules
     - **Not Matched** - The file does not match with any of the YARA rules
     - **Not Yet Scanned** - There is no scan performed for this file. Files will be scanned automatically once they have been downloaded. After the scan YARA status will be updated to either Matched or Not Matched

   Note: When a file matches with a YARA rule, high severity alerts are triggered and, the file's risk score is updated. In the subsequent scans, If the same file doesn't match with a YARA rule, the risk score will be reset.

   Note: If a downloaded file has an error, it will not be scanned by YARA and, the **Downloaded** column will display the file download status as **Error**.



4. On the **Filters** pane, scroll to the **YARA RULES** section. This section provides options to filter the files based on YARA scan status:
   - Select **Matched** to view the files that match YARA rules.
   - You can also view the files that do not match YARA rules or not yet scanned against YARA rules, by selecting **Not Matched** or **Not Yet Scanned** from the **YARA RULES** section.

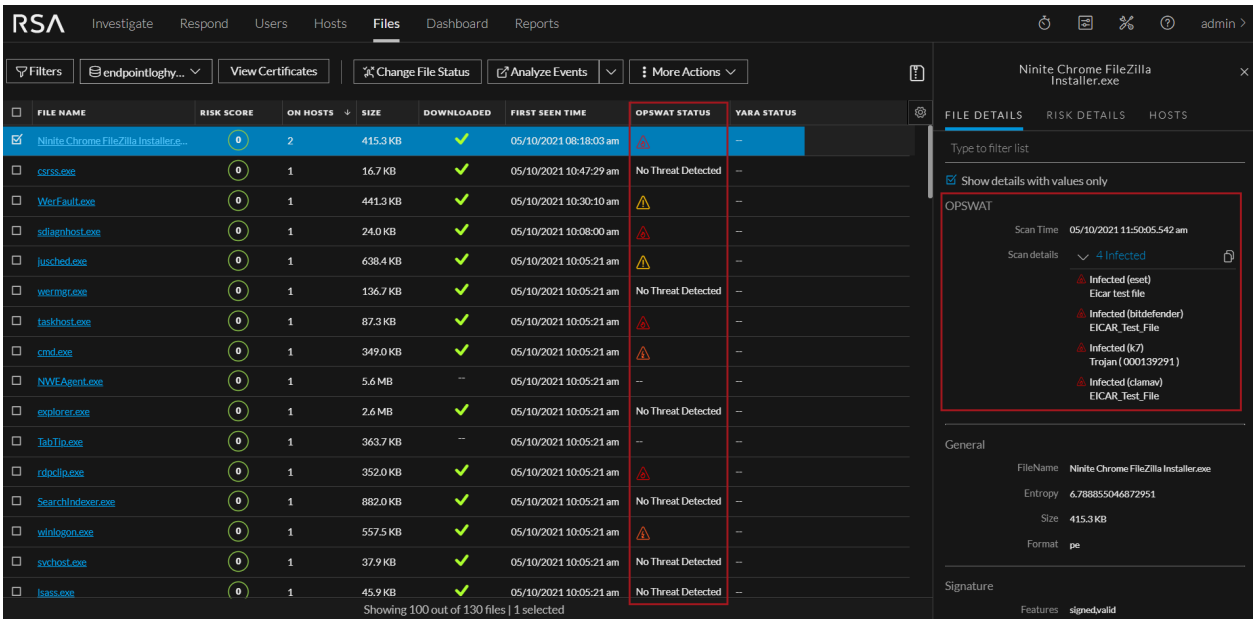


For more information on investigating with YARA, see *NetWitness Investigate User Guide.*

Analyze Files Using OPSWAT

OPSWAT (MetaDefender Core) provides advanced malware detection capabilities by scanning files with multiple anti-malware engines simultaneously. Administrators need to enable and configure OPSWAT on the Endpoint server. To learn more about enabling and configuring OPSWAT, see *NetWitness Endpoint Configuration Guide*.

All downloaded files (executable) will automatically be sent to OPSWAT server for scanning once OPSWAT is enabled and configured on the endpoint servers.



Scan files with OPSWAT

Downloaded files are automatically sent to the OPSWAT server for scanning. However, you can also initiate the scan manually using options under the**More Actions** menu. Executable files with the following file extensions, ***pe***, ***script***, ***macro***, and ***elf*** are supported by this feature. The maximum file size limit is set to 10MB by default. You can increase or decrease it if required.

Automatic scan:

The endpoint server will automatically send all (executable) downloaded files to the OPSWAT server. The scan results will be displayed under the**OPSWAT STATUS** column on the**Files** tab.

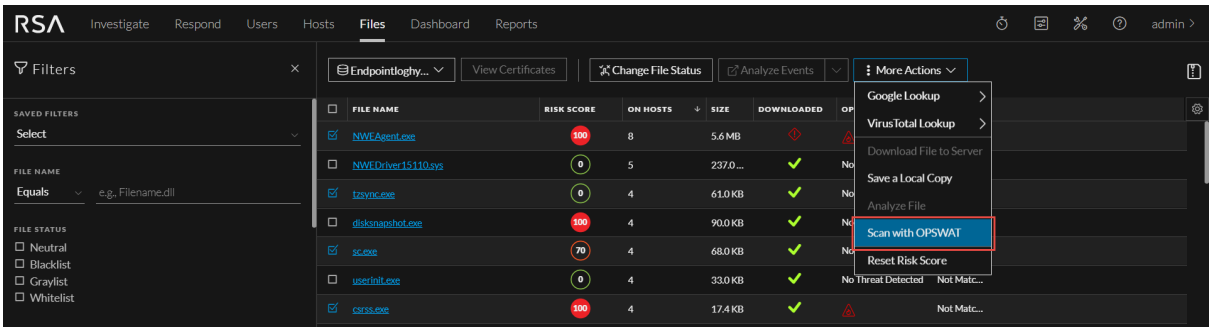
Manual Scan:

You can also manually initiate an OPSWAT scan using the options under the**More Actions** menu on the **Files** tab.

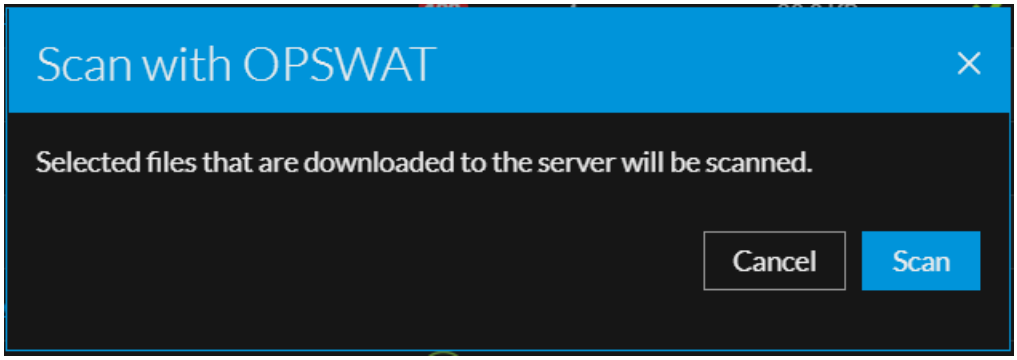
The **More Actions** menu provides the following options:

Manual Scan - Scan selected files

1. Select the files that need to be sent to the OPSWAT server for scanning.
2. Select **More Actions > Scan with OPSWAT**.

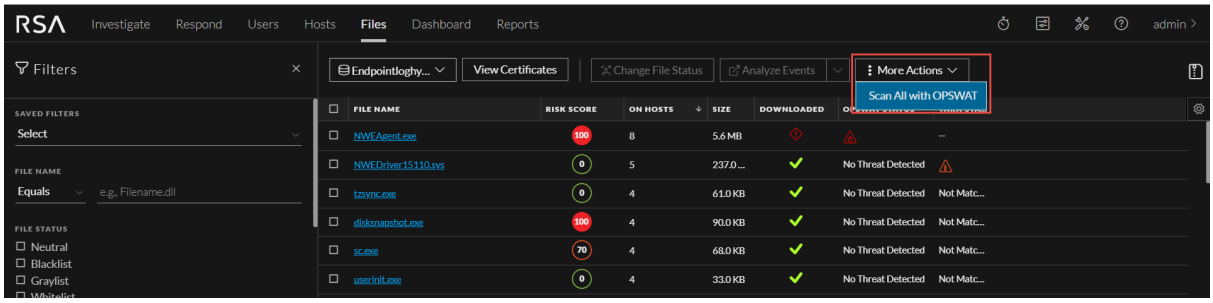


3. Click **Scan** on the confirmation pop-up.

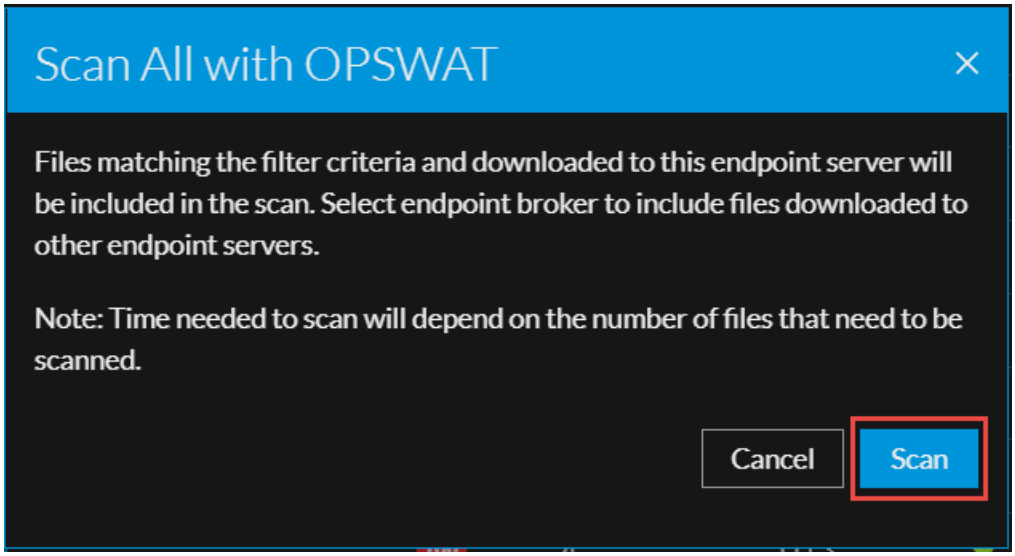


Manual Scan - Scan all files

1. Select **More Actions > Scan All with OPSWAT**.



2. Click **Scan** on the confirmation pop-up.



View OPSWAT Scan Results

The scan results will be displayed under the**OPSWAT STATUS** column on the**Files** tab as follows:

 - File is infected

 - Suspicious file

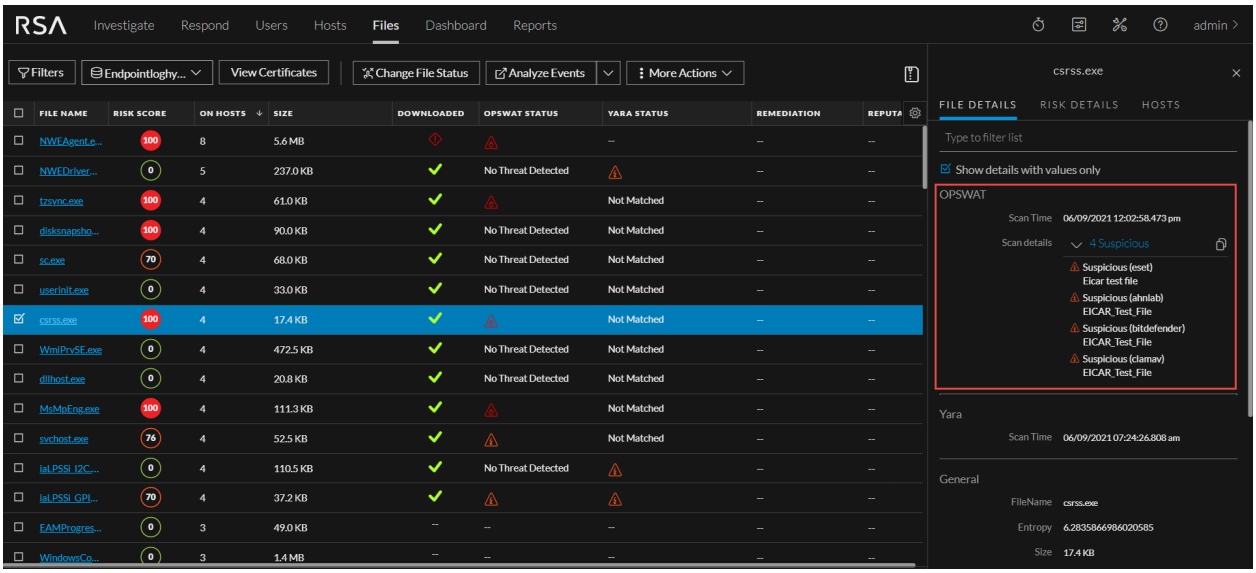
 - Scan failed; see troubleshooting section for more information.

 - The file is clean as of the last scanned time.

 - Not yet scanned.

Click on a file to view scan results under the**FILE DETAILS** panel.

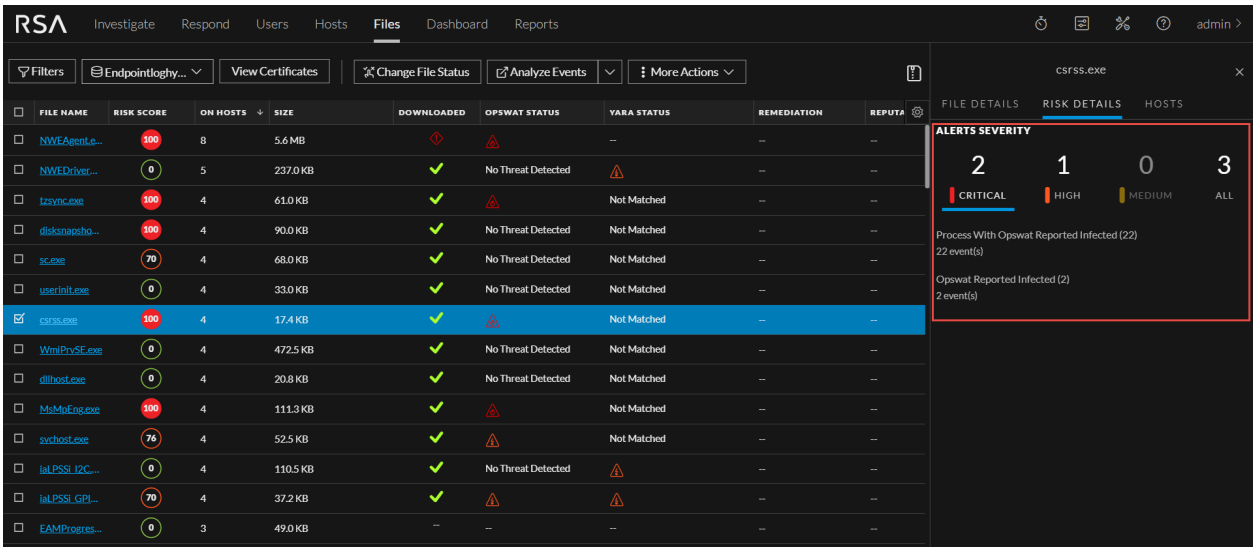
- **Scan Time**: States the last scanned time.
- **Scan details**: States whether a file is infected or suspicious or no threat detected, names of anti-malware engines that identified the threat.



Alerts and Impact on Risk scores

Alerts: Critical alerts will be triggered when OPSWAT finds a file as infected, and High severity alerts when a file is found suspicious.

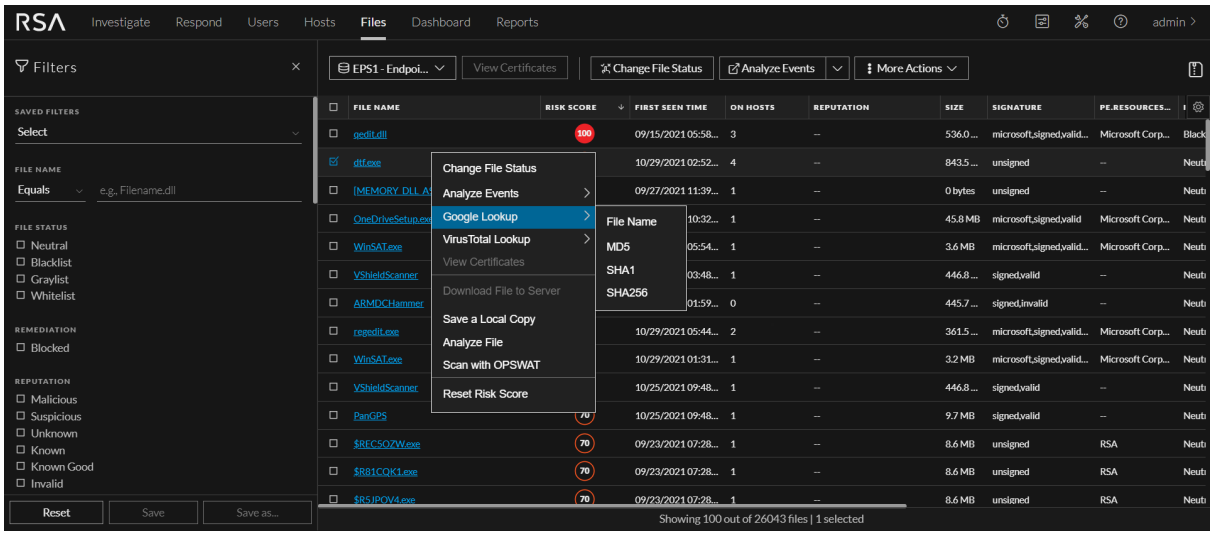
Risk score: If a file is found infected or suspicious by OPSWAT, the risk score of that file and the corresponding host will be increased.



Launch an External Lookup for a File

While analyzing a file, you can search Google or VirusTotal with the filename or hash to get more information about the file. To launch the search:

1. Go to **Files**.
2. View the details of the file name and hash from the table MD5, SHA1, and SHA256 columns, or view the details in the File Details tab on the right panel.
3. Select one or more files, and right-click or in the **More Actions** drop-down list in the toolbar, do the following:



   a. Select **Google Lookup** and perform a search on the filename, MD5, SHA1, or SHA256.

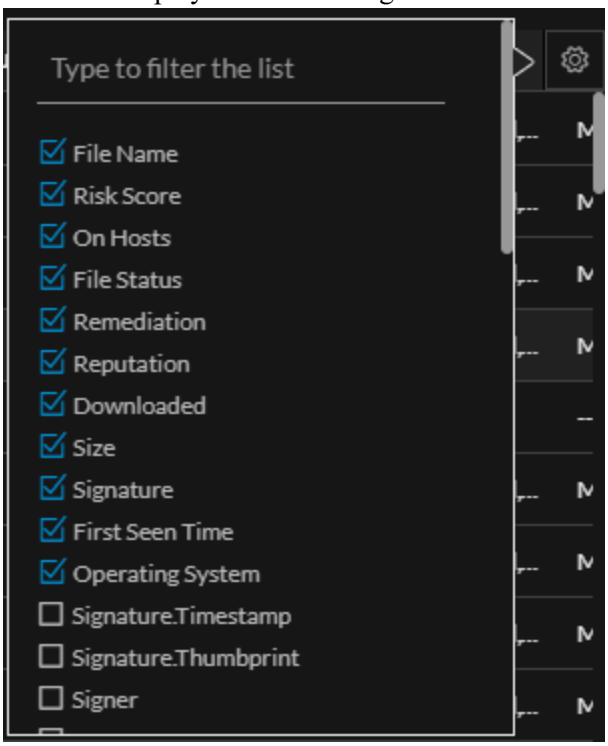
b. Select **VirusTotal Lookup** and perform a search on MD5, SHA1, or SHA256.

Note: To open files in multiple tabs, make sure you enable the pops-up in the browser.

Set Files Preference

By default, the Files view displays a few columns and the files are sorted based on the risk score. If you want to view specific columns and sort data on a specific field:

1. Go to **Files**.
2. Select the columns by clicking in the right-hand corner. The following example shows the drop-down list displayed while adding columns:



3. Sort the data on the required column.

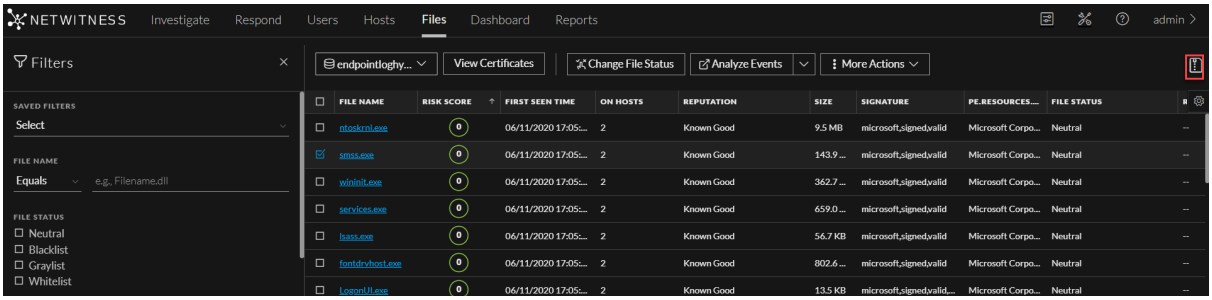
Note: The selections you make here become your default view every time you log in to the Files view.

Export Global Files

To extract the list of global files to a comma-separated values (csv) file:

Note: While filtering on a large data set, use at least one indexed field with the `Equals` operator for better performance. You can export up to 100k files at a time.

1. Go to **Files**.
2. Filter the files by selecting the required filter option.
3. Add columns by clicking  in the right-hand corner.
4. Click  to export the files to a csv file.



You can either save or open the CSV file.

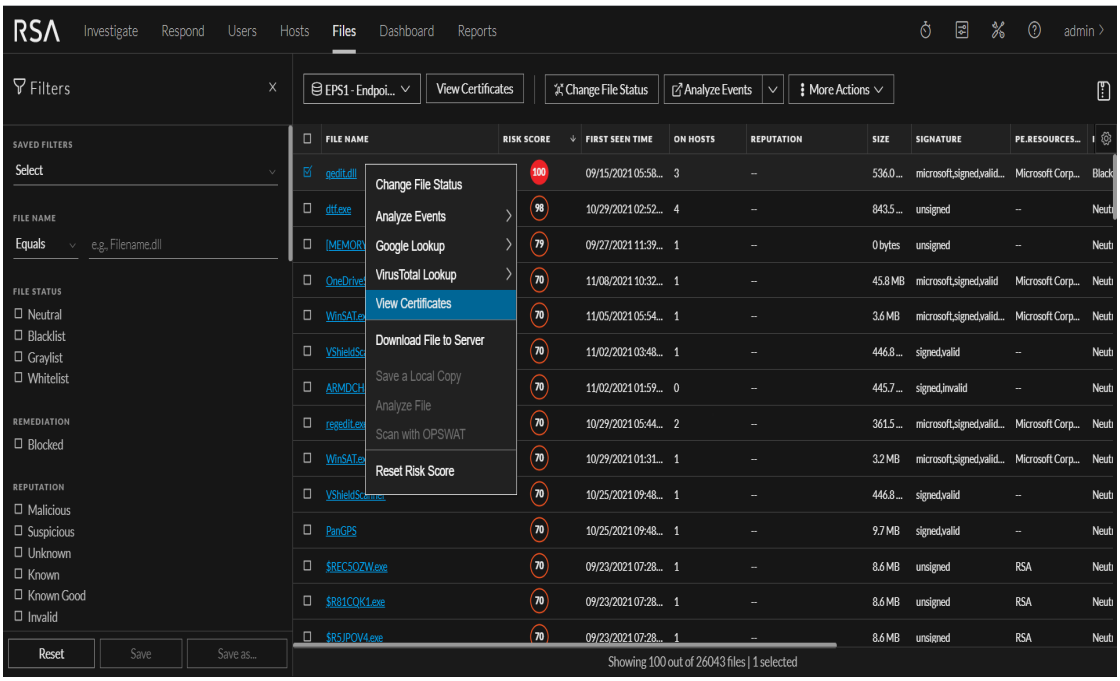
Analyze Certificates

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The Certificates view provides a list of code-signing certificates reported by hosts found in your deployment and their associated properties. You can select the certificates under a specific Endpoint server.

To view the certificates in an Endpoint server:

1. Go to **Files**.
2. From the drop-down menu, select the Endpoint server to view certificates present on that server. To view a consolidated list of certificates, select the Endpoint Broker server.
3. Select a file and do one of the following:



- Right-click and select **View Certificates** from the context menu.
- Click **View Certificates** in the toolbar.

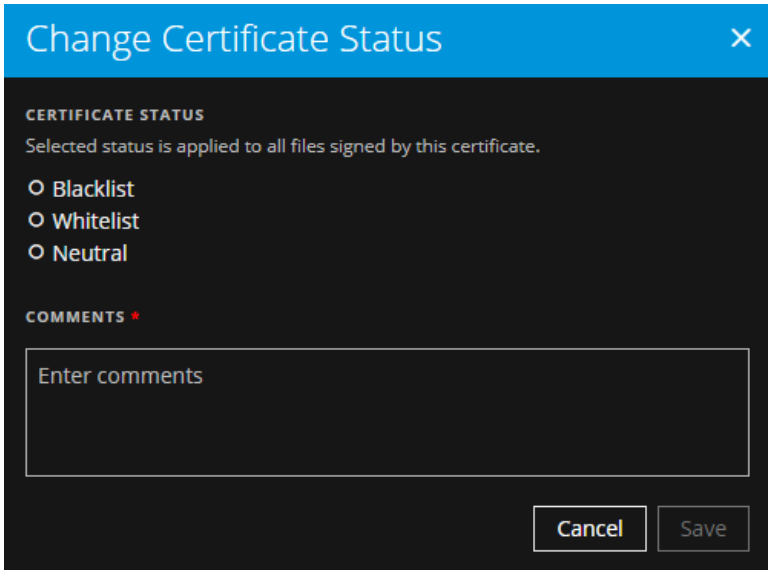
Change the Certificate Status

You can assign a Whitelist status to the certificate signed by certain trusted vendors and this status can be automatically applied to all files that is signed by this certificate. If you consider abc a trusted vendor, you can set the status for the certificates signed by abc as Whitelist.

Similarly, you can also set the certificate status as Blacklist or Neutral. If a company's certificate is stolen or compromised, you can blacklist this certificate and remediate.

To change the certificate status:

1. Select a certificate, and click **Change Certificate Status**.



2. In the Change Certificate Status dialog, select a status - Blacklist, Whitelist, or Neutral.

   Note: If you have manually updated a file status in the Files or Hosts view, changing the status in the Certificate view does not impact the file status as the manual update takes precedence. For example, if you have whitelisted the file vmci.sys that is signed by VMware, Inc. in the Files or Hosts view, and you have blacklisted VMware, Inc. in the Certificate view, the file vmci.sys remains Whitelisted though the certificate is blacklisted.

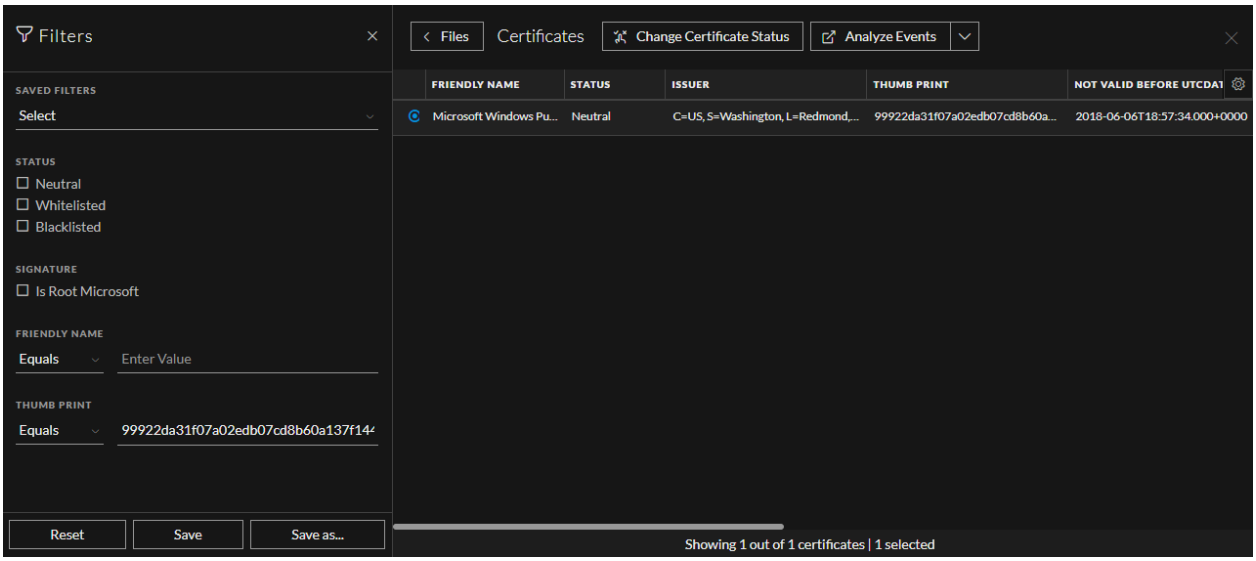
3. Add a comment and click **Save**.

4. Click **< Files** to go to the Files view.

Note: In a multi-server environment, changing the status of a certificate in one endpoint server updates the respective files in other endpoint servers. For example, if a certificate status is set to Blacklist on one endpoint server, all files signed by this certificate are set to Blacklisted on all endpoint servers.

Filter Certificates

You can filter certificates on status, signature, friendly name, and thumb print.



Click **Save** to save the filter and provide a name (up to 250 alphanumeric characters). The filter is added to the Saved Filters list. To delete a filter, hover over the name, and click .

Note: Special characters are not allowed except underscore (_) and hyphen (-) while saving the filter.

Resetting Risk Score of Files

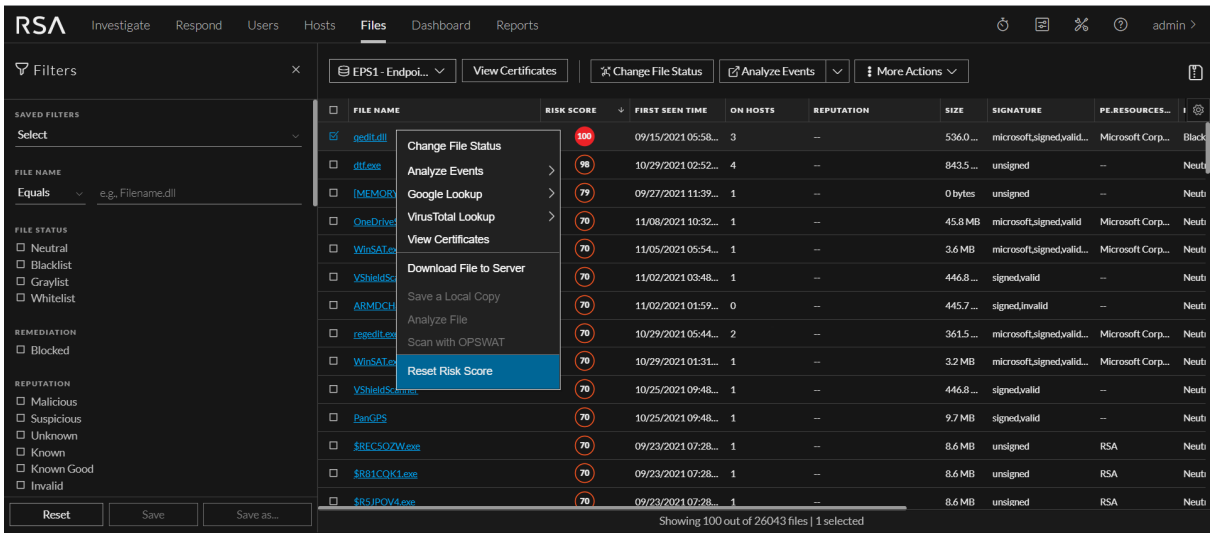
You can reset the risk score for a file in these situations:

- If the alerts or events triggered by the host or a file are considered to be false positive, you can make required changes to the Endpoint Application rules or ESA rules.
- After you take required action on a malicious file.

When you reset the risk score, the risk calculation for the file is deleted and score is set to 0. The risk score on all the hosts on which this file exists is recalculated. You can reset the risk score for a single file or multiple files.

To reset the risk score of a file:

1. Go to **Files**.
2. Select the Endpoint Server or Endpoint Broker.
3. Select one or more files and do one of the following:



- Right-click and select **Reset Risk Score** from the context menu.
- Click **More Actions** > **Reset Risk Score** in the toolbar.

All the alerts associated with the score are deleted.

Note: You can select a maximum of 100 files to reset the score.

4. Refresh the page to view and confirm if the file's score is reset. This may take sometime for changes to take effect.

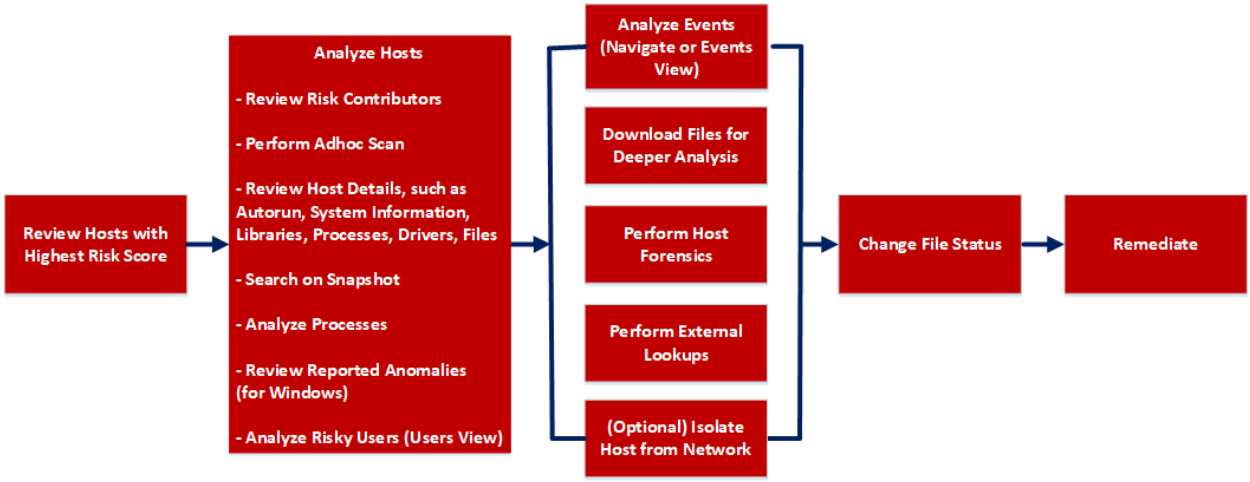
Investigating Hosts

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The Hosts view allows you to investigate on a host, which includes scan details, tracking events related to alerts, anomalies, and process details.

Best Practices

The following are some best practices and tips that may help you investigate efficiently to identify and isolate threats or attacks:



- Review hosts with highest risk score and analyze the alerts contributing to the risk. Review the entities, such as file name, processes involved in the alerts. For more information, see Analyze Hosts Using the Risk Score.
- Review files or processes that created this suspected file, and check if any other files are accessed or created in the Events view. For more information, see Analyzing Events.
- Review hosts for rare files in the **On Hosts** column. If a file is present on 100 hosts, it can be legitimate. If a file is present on fewer hosts with a high risk score, it may be malicious and needs further investigation.
- Filter to exclude hosts on host status, risk score, hostname, and so on. For more information, see Filter Hosts.
- Search Google or VirusTotal with the file hash and review any reported activities. For more information, see Launch an External Lookup for a File.

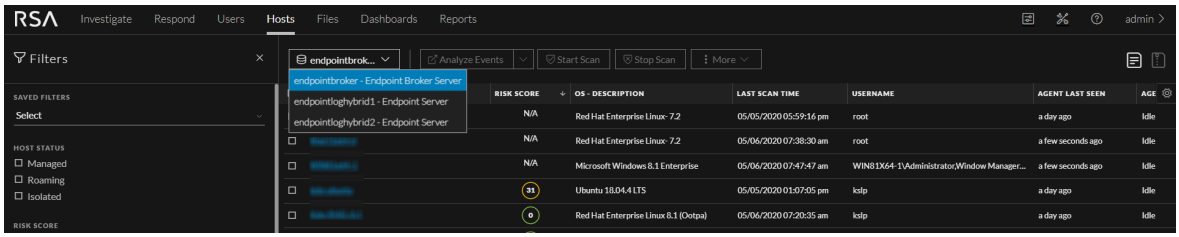
- Review the processes, autoruns, files, libraries, drivers, and system information. For example,
  - Search for files in known malware locations. For example,
    - C:\Windows\.
    - C:\Users\<name>\AppData\<uncommon folder>.
    - C:\Users\<name>\AppData\Local\Temp.
    - C:\Windows\Temp\.
  - Search for a particular file name or hash and review the snapshot to check when the file was first seen.
  - Review any network connections established by the process, such as:
    - Domain or IP address.
    - Ports used (common (80 and 443) versus uncommon ports (8080 , 8888, and 3465)) and check if the ports are listening actively.
  - Check the file compile time. If the date is recent, it could be malicious.
  - Check the file creation time on the host.
- Review reported anomalies, such as suspicious threads, kernel hooks, image hooks, and registry discrepancies. For more information, see Analyze Anomalies.
- Launch Process Analysis to view the sequence of activities performed on the host by the file or process. For more information, see Analyze Processes.
- Download suspicious files to the server for deeper analysis. For more information, see Analyzing Downloaded Files.
- Download MFT, process, or system dump to the server for forensic investigation. For more information, see Performing Host Forensics.
- After investigation if a file is found to be malicious you change the status of the file (blacklist or graylist) and block infected or malicious file. For more information, see Changing File Status or Remediate.
- (Optional) If you suspect that a host is potentially compromised with the threat still being active, you can isolate the host from the network and safely investigate possible threats within the host. For more information, see Isolating Hosts from Network.

View Hosts

You can view all hosts present on a specific Endpoint server or consolidated list of all hosts on multiple Endpoint servers using the Endpoint Broker for analysis. By default, hosts are sorted based on the risk score. To view the hosts:

1. Go to **Hosts**.
2. Select from the following:

- **Endpoint Broker Server** to view all hosts across all Endpoint servers. When querying, the Endpoint Broker ignores Endpoint servers that are offline. If the Endpoint server is online but is not responding, the Endpoint Broker waits for 10 seconds, and ignores if it does not respond.
- **Endpoint Server** to view hosts on a specific Endpoint server.



3. Select a host that you want to analyze.
4. Click a row to view the following details:
   - **Host Details** displays the host information such as Network Interfaces, operating system, hardware and others.
   - **Risk Details** displays the distinct alerts associated to the risk score and the alerts severity. Click **Critical**, **High**, **Medium**, or **All** to display all the alerts. For more information, see Analyzing Risky Users.
5. Click **Show next 100 hosts** to view other hosts.
6. Click the host name to investigate the scan results. For more information, see Analyze Host Details.

Manage Hosts Using Tags

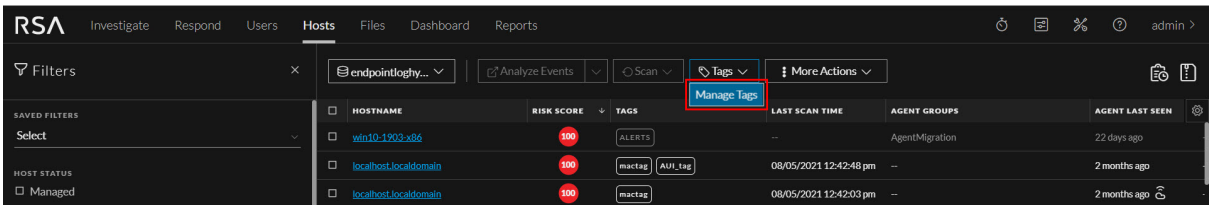
From version 11.7 and later, NetWitness Platform allows you to create Tags to manage the hosts effectively. Tags are custom texts that you can create and assign to hosts for identifying them. A tag can contain alphabets, numbers and special characters(Except \ **' , [ ] "** and **space**). You can use these tags to create host groups. You can also filter hosts by tags using the filters pane on the Hosts screen.

Manage Tags

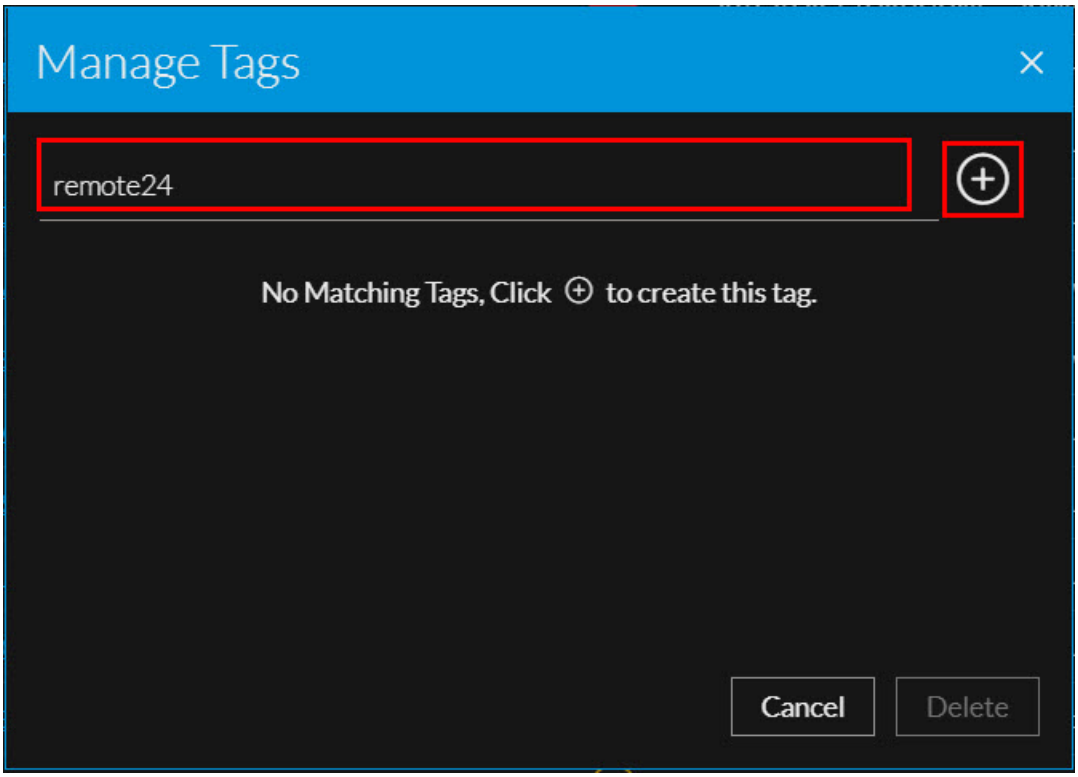
The Manage Tags option allows you to create and delete tags without selecting any hosts.

To Create Tags:

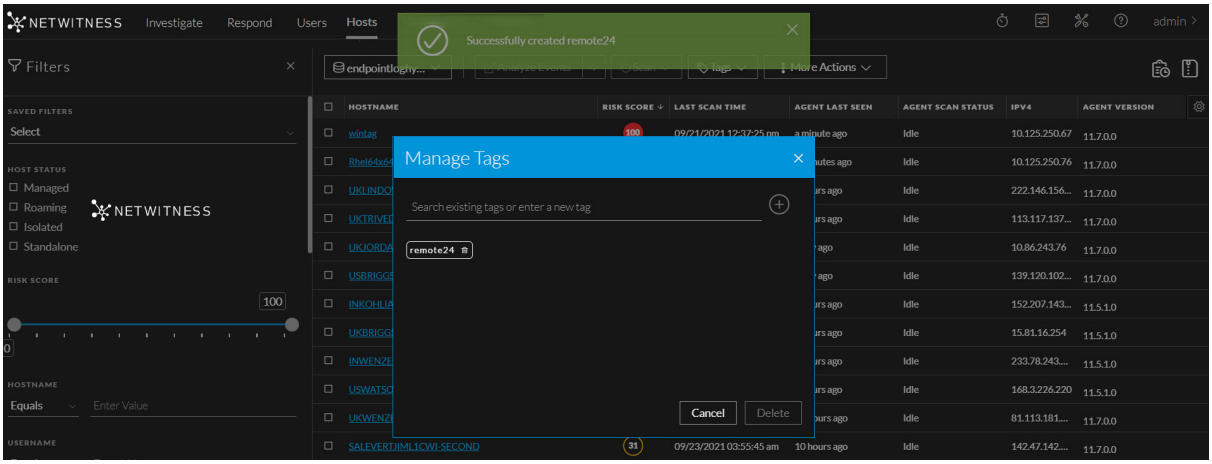
1. Click **Tags > Manage Tags** on the Hosts tab.



2. Enter a valid tag in the text field on the **Manage Tags** pop-up and click .



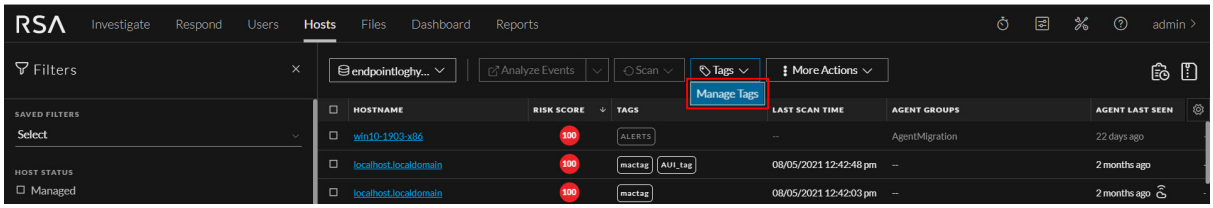
- The tag is created, and a success message will appear.



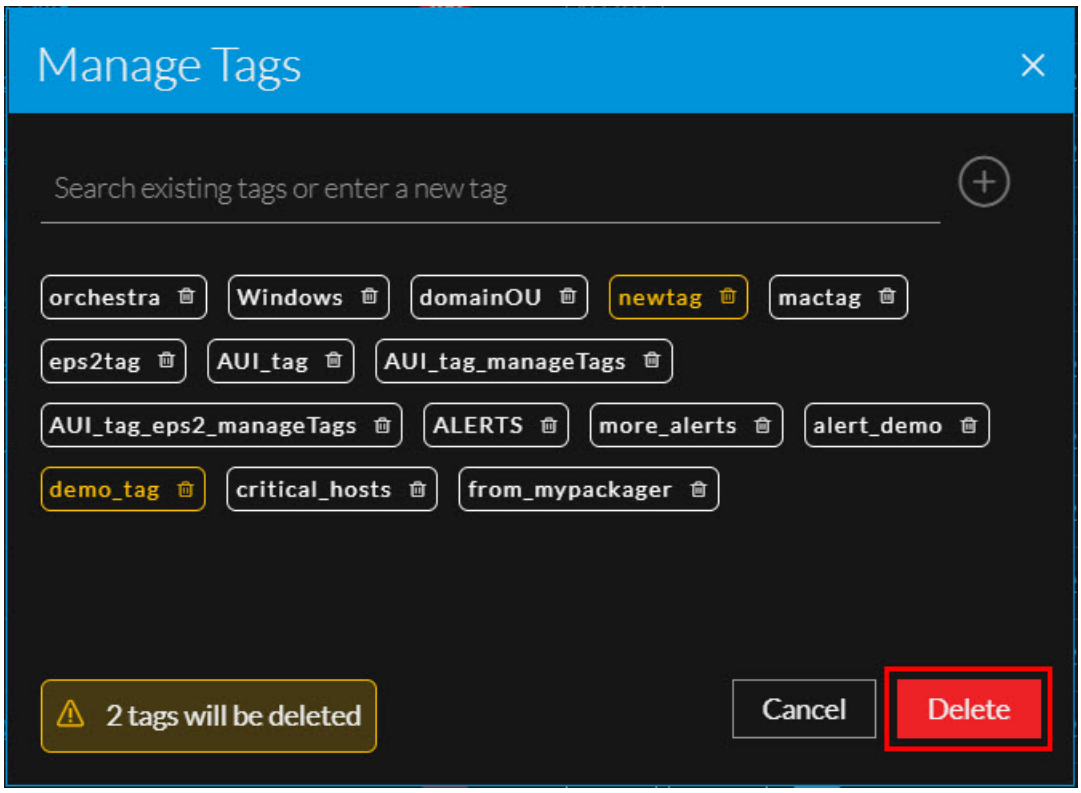
3. Repeat step 2 to create more tags.

To Delete Tags:

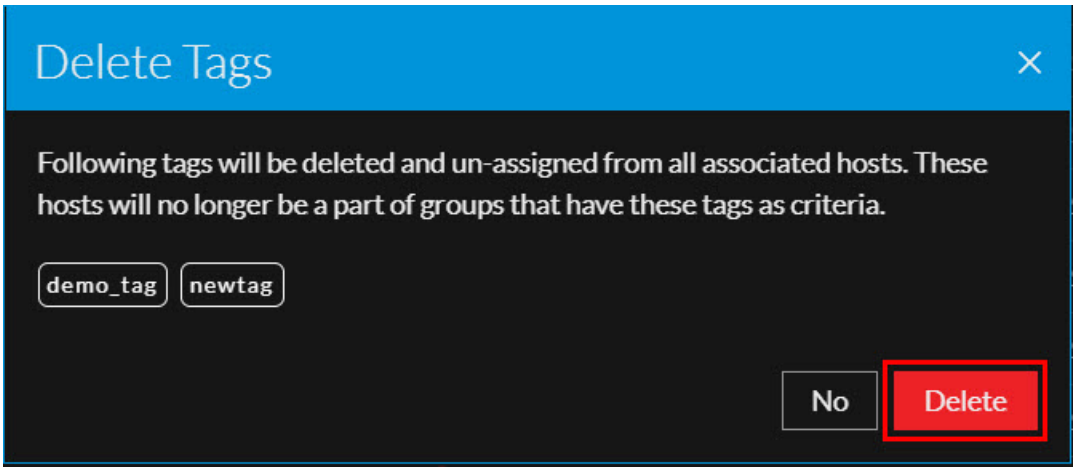
1. Click **Tags > Manage Tags** on the Hosts tab.



2. Select the tags that you want to delete and click **Delete** on the **Manage Tags** pop-up.



3. Click **Delete** on the **Delete Tags** confirmation pop-up.



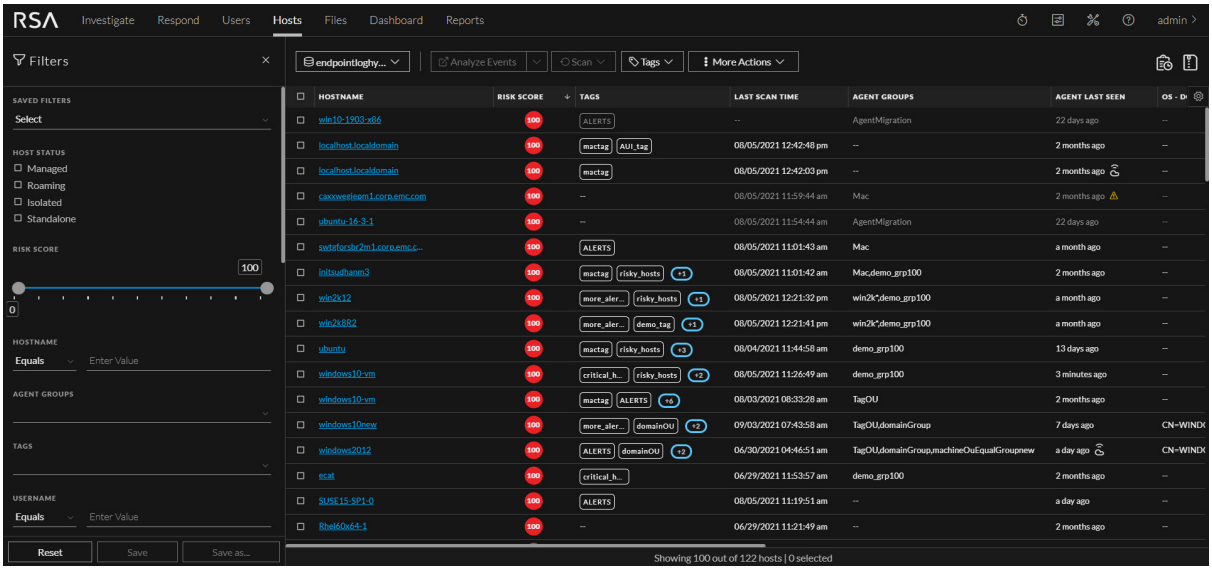
Note: The tags you deleted will also be unassigned from all associated hosts. If the associated hosts are a part of groups created using these tags as one of the criteria, then the hosts will no longer be a part of those groups. Refer What happens next after unassigning or deleting tags from hosts? for more information.

Assign tags

1. Select one or more hosts on the Hosts tab.
2. Select **Tags > Assign Tags** from the menu.

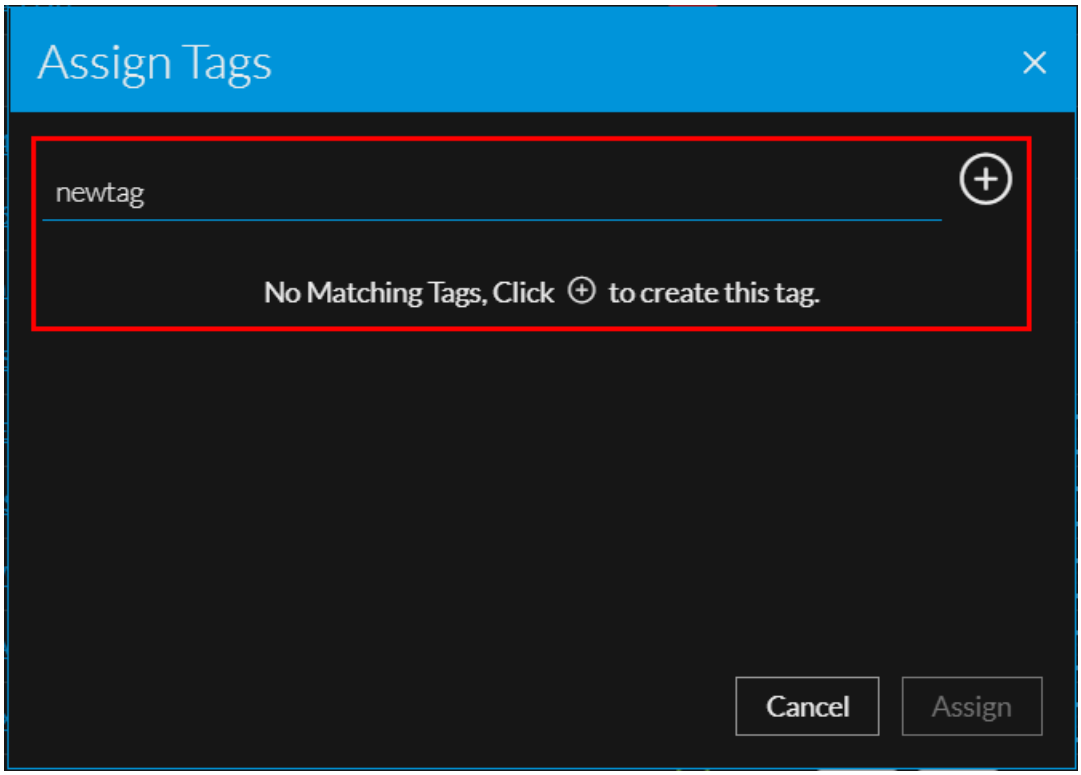
   (or)

   - Select a host, right click and select **Assign Tags**.



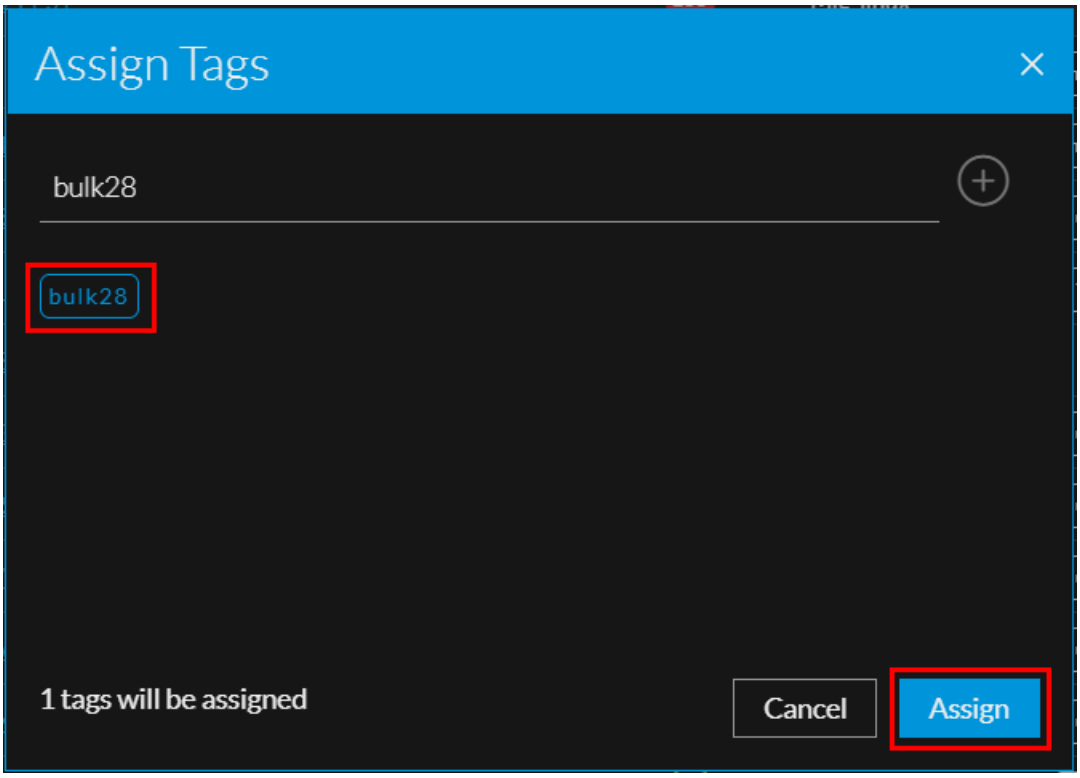
3. The Assign Tags pop-up will appear and shows all the existing tags. You can either

- Enter a valid tag and click  to create it. The newly created tag is selected by default.

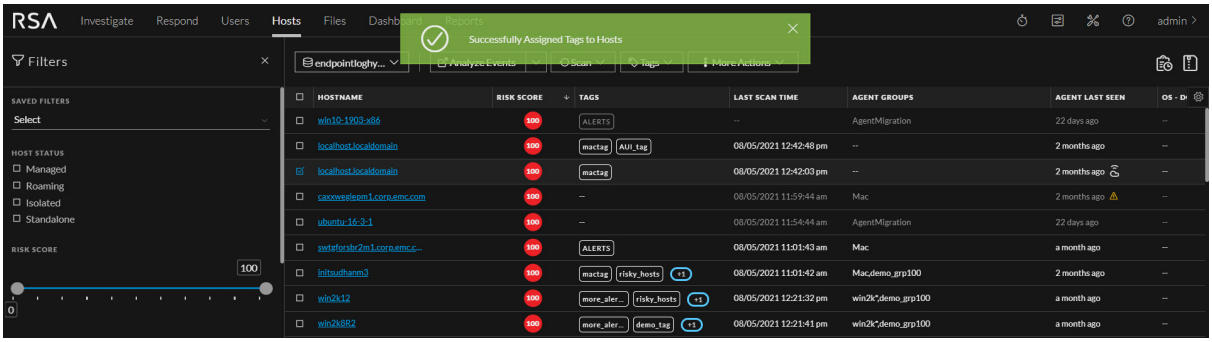


(Or)

- Search for an existing tag using the text field and select it if required.



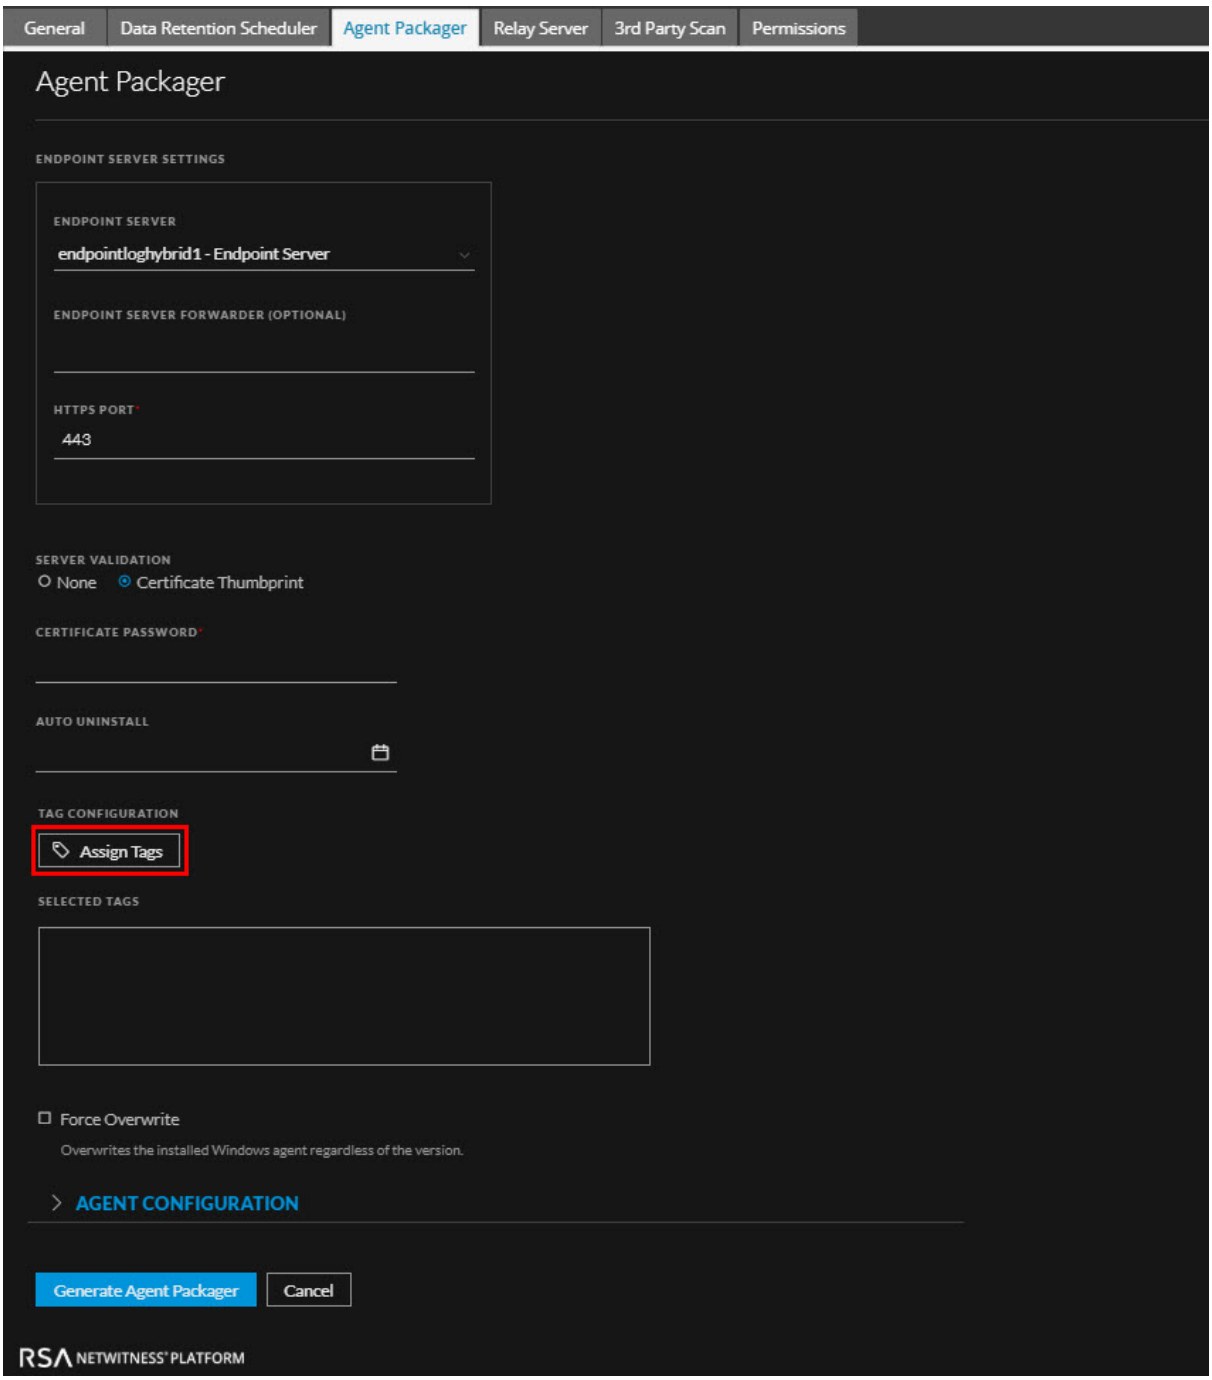
4. Click **Assign**. These tags will be assigned to the selected hosts.



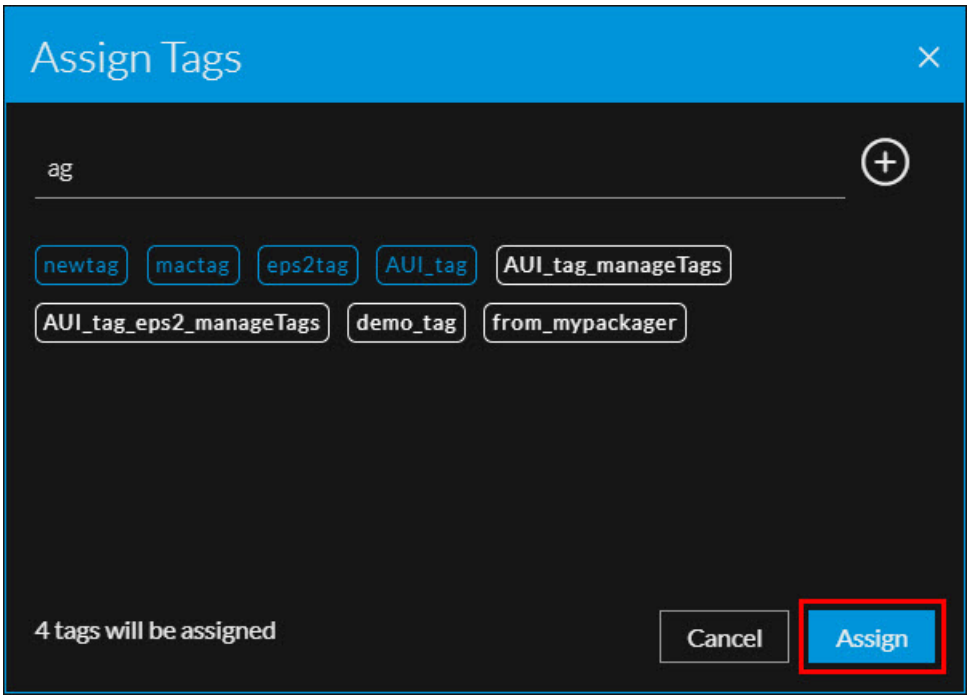
Create and Assign Tags When Generating the Agent Packager

You can add tags to the hosts while installing the Endpoint agents. When generating the agent packager on the Agent Packager tab, you can either create new tags or select already existing tags. These tags will automatically be assigned to the host in which the agent will be installed.

1. Click **Assign Tags** under **TAG CONFIGURATION** on the Agent Packager tab.



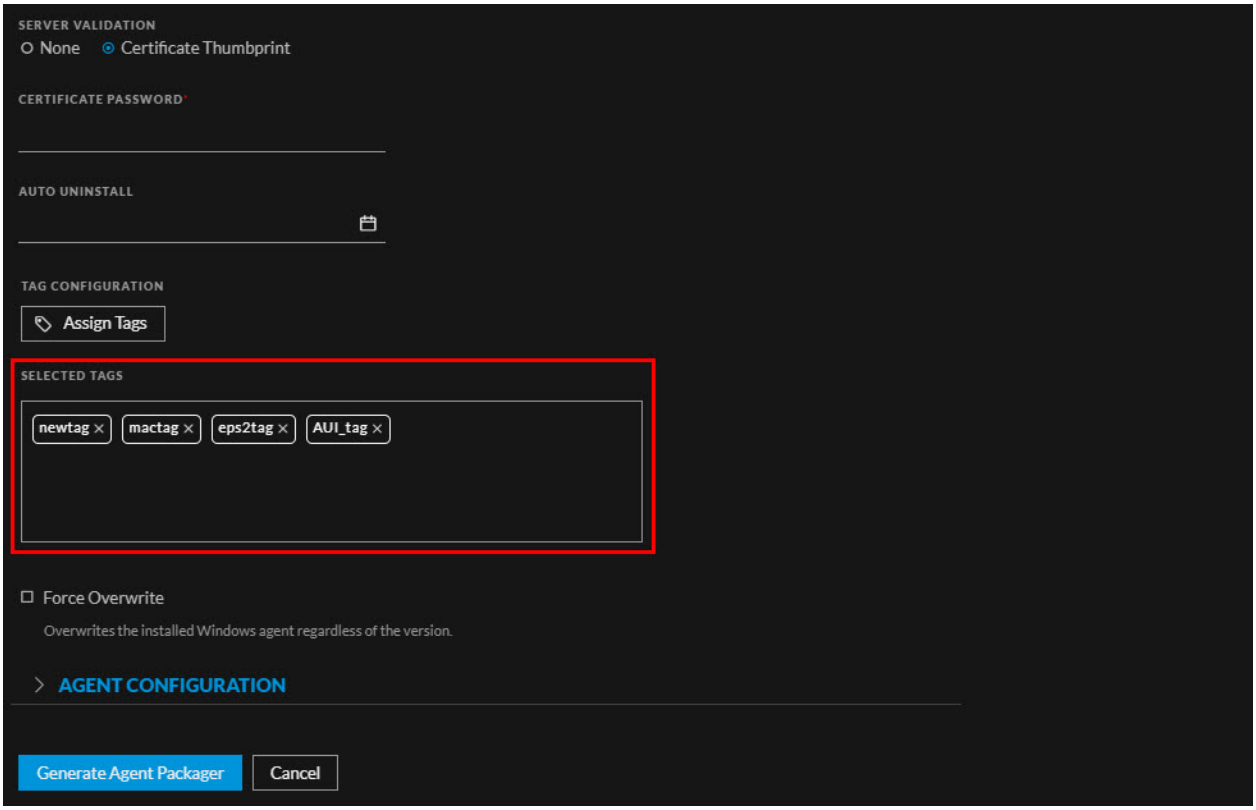
2. On the **Assign Tags** pop-up,
   - Search for an existing tag using the text field and click **Assign**.



(or)

- Enter a new tag in the text field and click  to add it to the selection. Click **Assign**.

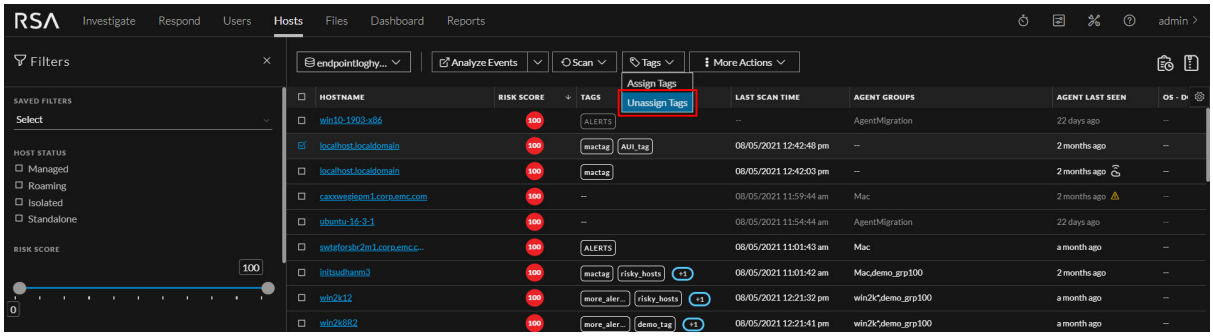
3. Assigned tags will appear under **SELECTED TAGS**.



Note: The tags added while generating the Agent Packager are applicable only to the newly installed agents and not to the manually upgraded agents.

Unassign tags

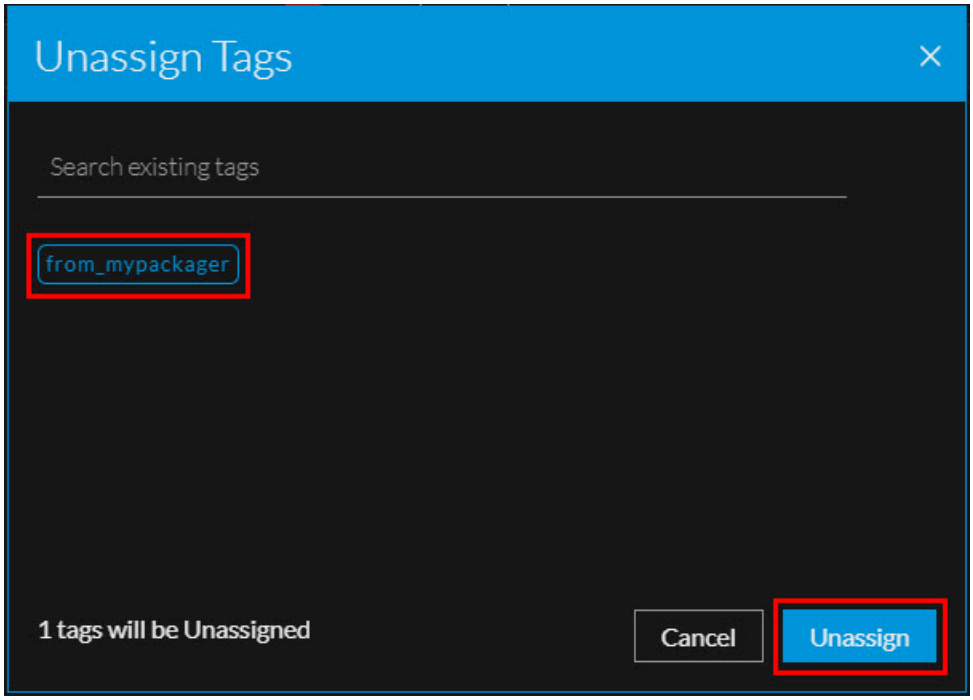
1. Select one or more hosts on the Hosts tab.
2. Select **Tags > Unassign Tags** from the menu.



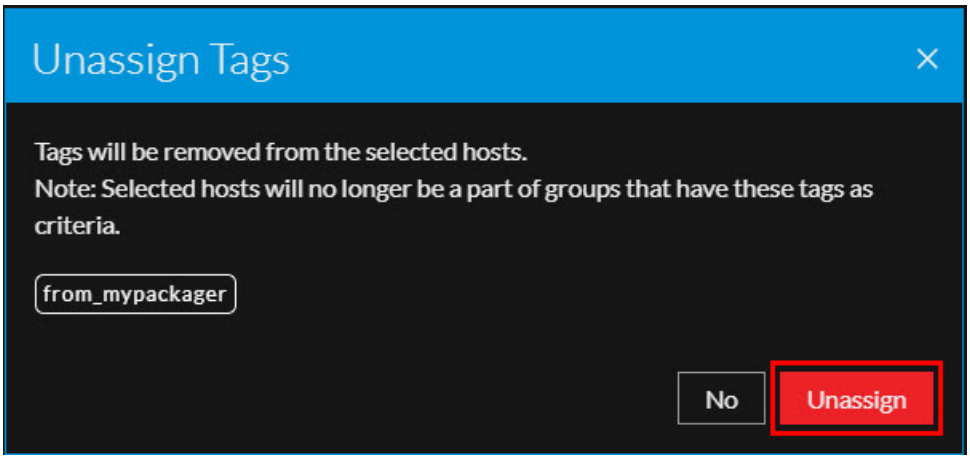
(or)

- Right click on a host and select **Unassign** Tags. The **Unassign Tags** pop-up will appear and shows all the tag assigned to the selected hosts.

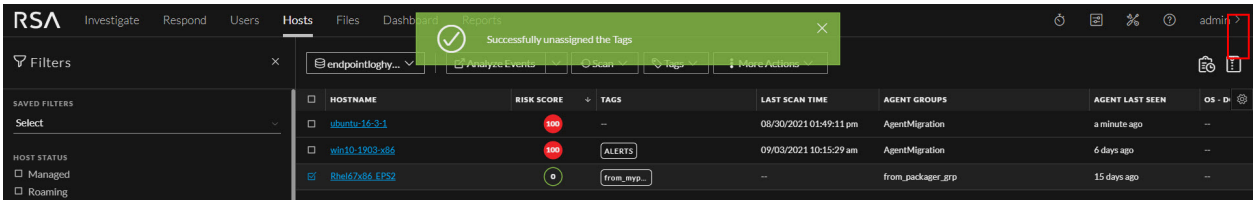
3. Search for the tags you want to unassign, select them and click **Unassign**.



4. Click **Unassign** on the confirmation pop-up.



5. These tags will be Unassigned from the selected hosts.



What happens next after unassigning or deleting tags from hosts?

Unassigning/deleting tags from hosts will immediately initiate group & policy evaluation. For example, if a host is a part of the group created using only the unassigned/deleted tag as criteria, the host will no longer be a part of that group. Refer to the following scenarios to understand more.

Example Scenario 1: Assume you create a group with a couple of tags as grouping criteria, all the hosts with these tags assigned will be a part of the group. And, if you delete these two tags(or unassigning from the hosts), the hosts may no longer be a part of this group.

Example Scenario 2: Assume you create a group with a couple of tags and IP addresses as grouping criteria; all the hosts assigned with these tags and IP addresses will be a part of the group. And, if you delete(or unassigning from the hosts) the tags alone, the hosts may still be a part of this group as they are still grouped using the IP addresses.

View Agent History

You can view the list of commands issued to the agents (by the server or actions performed by any analyst) in the Host view and Host details. By default, commands are sorted based on the command time.

To view the commands:

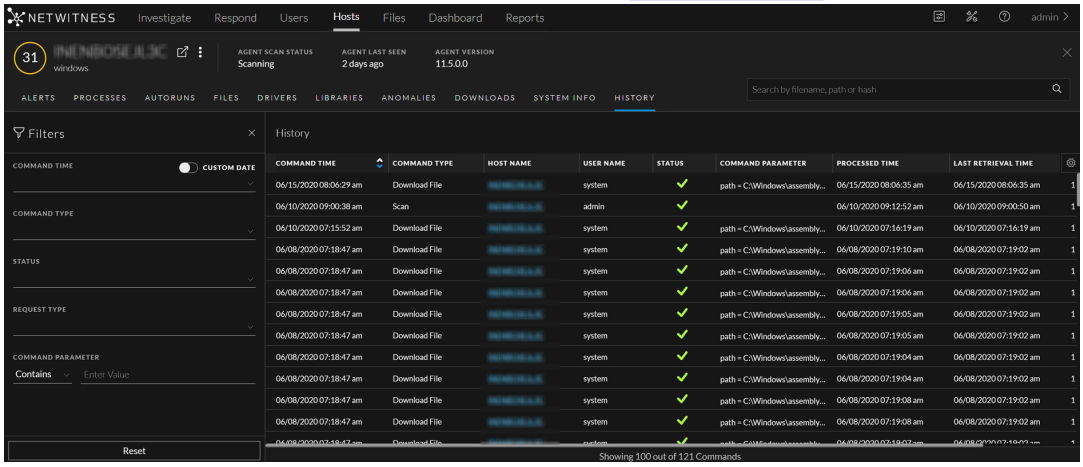
1. Go to **Hosts**.
2. Do any one of the following,

   - To view all commands, click . You can also filter commands, for more information see Filter Hosts.

     The Agent History view is displayed. For more details, see Analyze History.

   - To view commands specific to a particular host:
     - Click the host for which you want to view the commands.
     - In the Host details view, click **History** tab. You can also filter commands, for more information see Filter Host Details.
       The History view is displayed. For more details, see Analyze History.

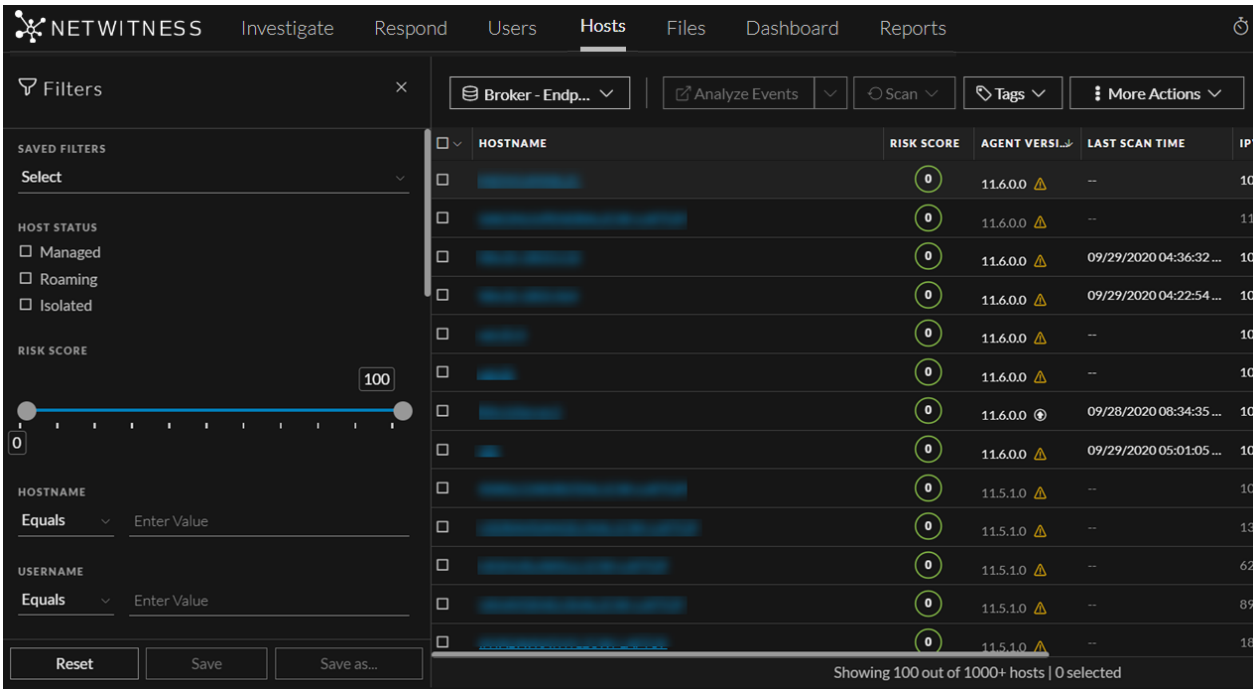


Filter Hosts

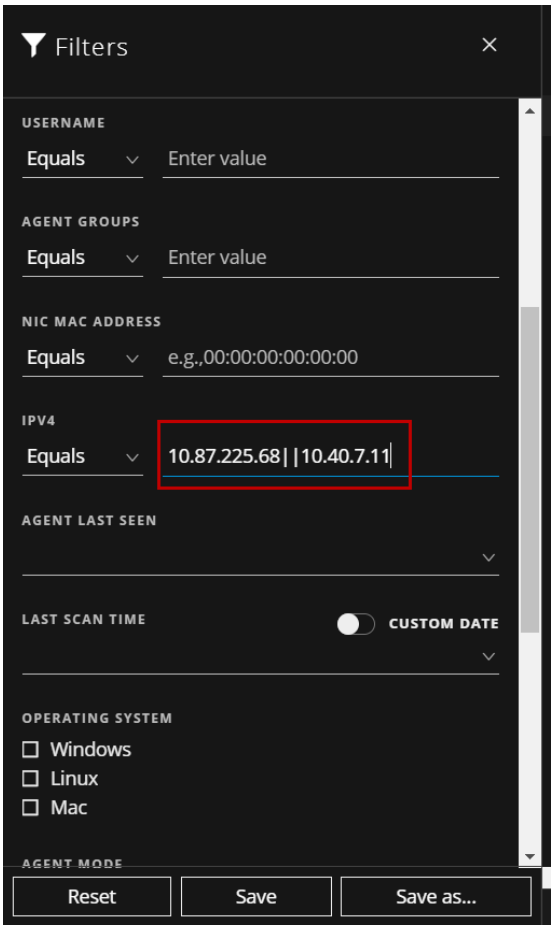
You can filter hosts on agent version, agent ID, agent mode, agent upgrade, agent last seen, last scan time, operating system, hostname, username, Mac address, risk score, IPV4, driver error code, security configurations, agent groups, and host status - managed, roaming, and isolated.

In the **Host** view > click , to filter the commands on command type, status, host name, request type, command parameter and command time. In the Command Time field, you can filter by custom date range.

Note: While filtering on a large amount of data, use at least one indexed field with the `Equals` operator for better performance. The following fields are indexed in the database - Hostname, IPv4, Operating System, Last Scan Time, and Risk Score.

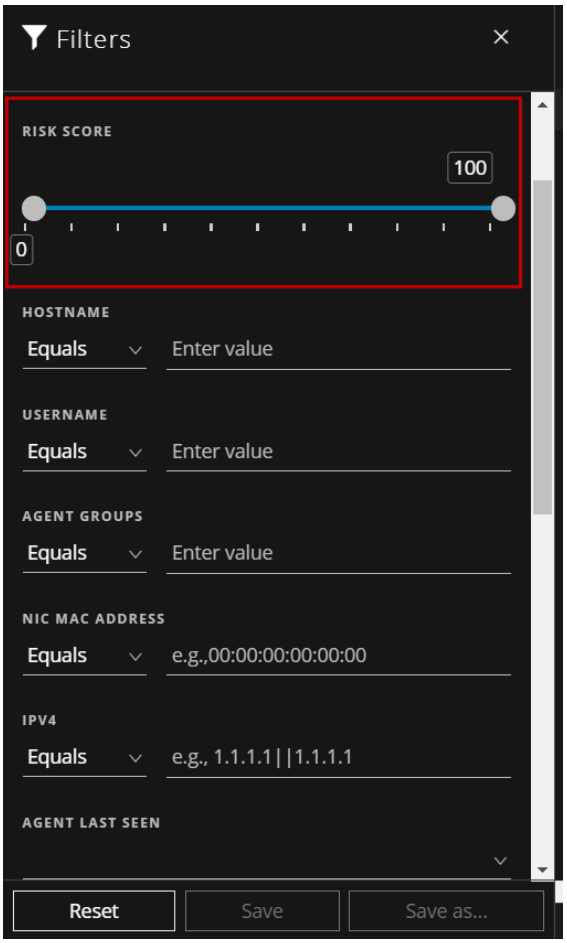


To search multiple values within a field, set the filter option to `Equals`, and use `||` as a separator. For example, using `Equals` operator for multiple IPV4 values with a separator `||`.



To filter on the agent last seen or last scan time, select the option from the drop-down list. If you select 3 Hours ago for the Last Scan Time, the result displays hosts that were last scanned 3 hours ago or earlier.

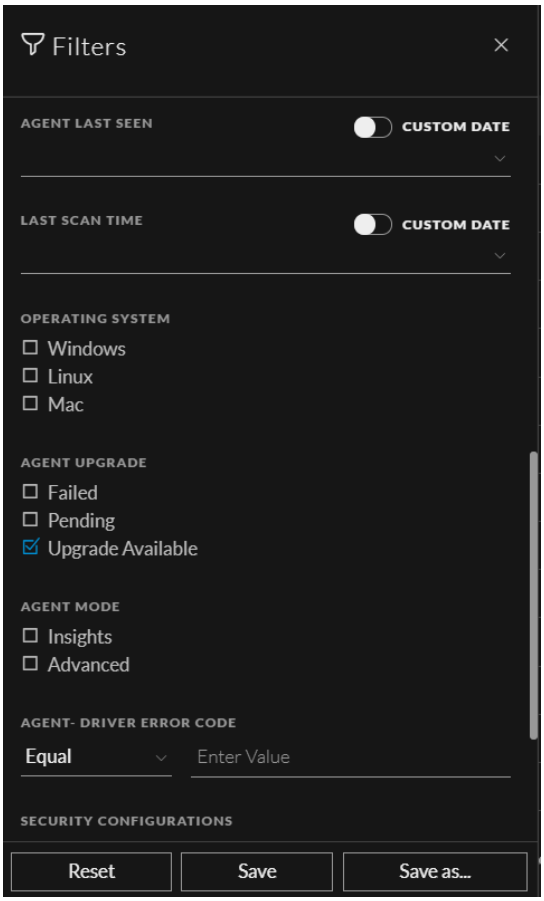
To filter on the risk score, use the slider to increase or decrease the values between 0 to 100.



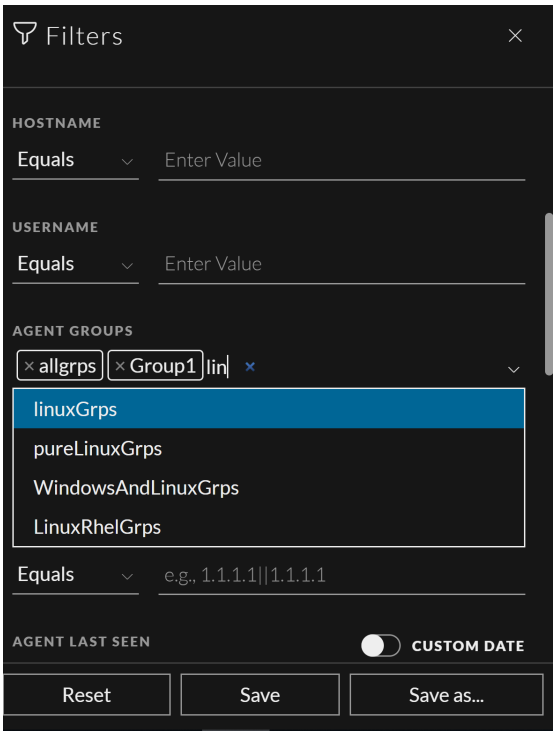
Click **Save** to save the filter and provide a name (up to 250 alphanumeric characters). The filter is added to the Saved Filters panel on the left. To delete a filter, hover over the filter name and click 🗑.

Note: Special characters are not allowed except underscore (_) and hyphen (-) while saving the filter.

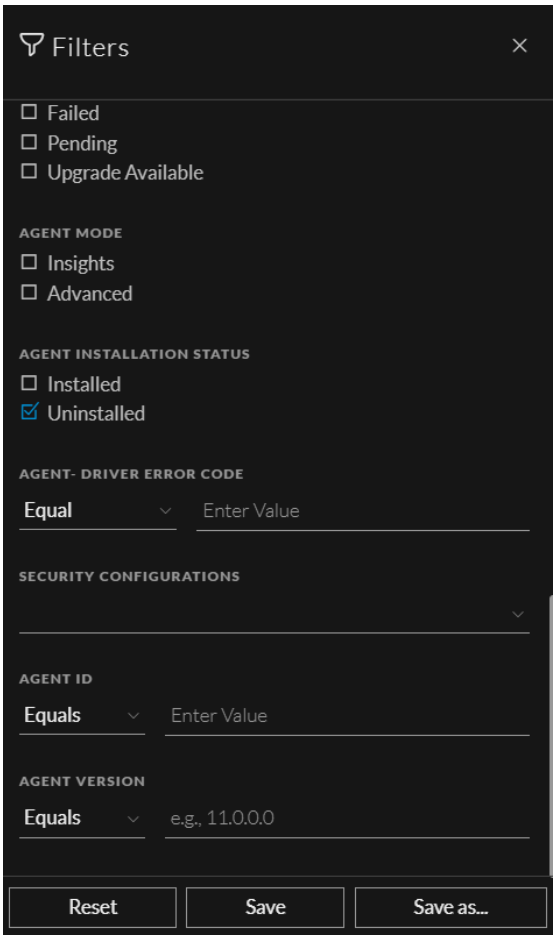
To filter the agents based on the upgrade status, select one of the upgrade statuses. For example, select the **Upgrade Available** checkbox to get the list of agents available for an upgrade.



To filter the agents based on single or multiple agent groups, select a group from the drop-down list. You can also search the name of the groups from this list.



To filter the agents based on the installation status, select one of the installation statuses. For example, select the **Uninstalled** checkbox to get the list of agents for which the uninstall is initiated or successfully completed.

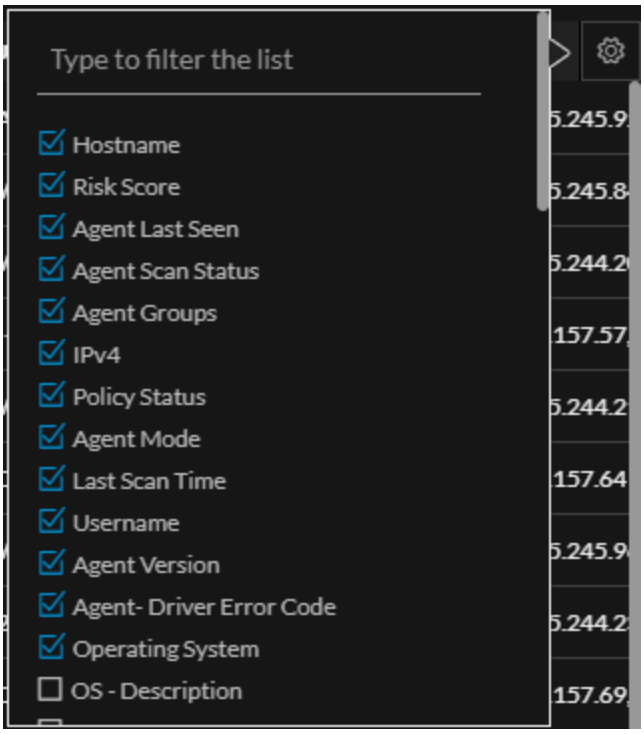


You can also filter the commands on command type, status, host name, request type, command parameter and command time (In which you can filter by custom date range), by clicking .

Adding and Sorting Columns in the Table

By default, the Hosts view displays a few columns and the hosts are sorted based on the risk score. To add or remove columns:

1. Go to **Hosts**.
2. Select the columns by clicking  in the right-hand corner.



3. Scroll down or enter the keyword to search for the column.
4. Click the arrow on the column header to sort the column in ascending or descending order.

Scan Hosts

You may want to perform an on-demand scan if you want to get the latest snapshot of the host.

You can either choose to perform a quick scan or a full system scan.

Quick scan - Scans all executable files that are loaded in memory. Both Insights and Advanced agents support quick scan.

Full System Scan - Scans all fixed drives or the system drive You can perform a full system scan only on advanced agents that are in version 11.6 or later. Native executables are included in the full system scan, by default

When hosts are scanned, the Endpoint Agent retrieves the following data that can be used for investigation:

- Drivers, processes, DLLs, files (executables), services, autoruns, anomalies, host file entries, and scheduled tasks running on the host.
- System information such as network share, installed Windows patches, Windows tasks, logged-in users, bash history, and security products installed.

To perform a Quick Scan:

1. Go to **Hosts**.
2. Select one or more hosts (up to 100) at a time, and do one of the following:

- Click **Scan > Start Scan** from the menu (Or)
- Right-click and select **Start Scan** from the context menu

3. Click **Start Scan** on the pop-up. Quick scan is initiated for executable modules loaded in memory.

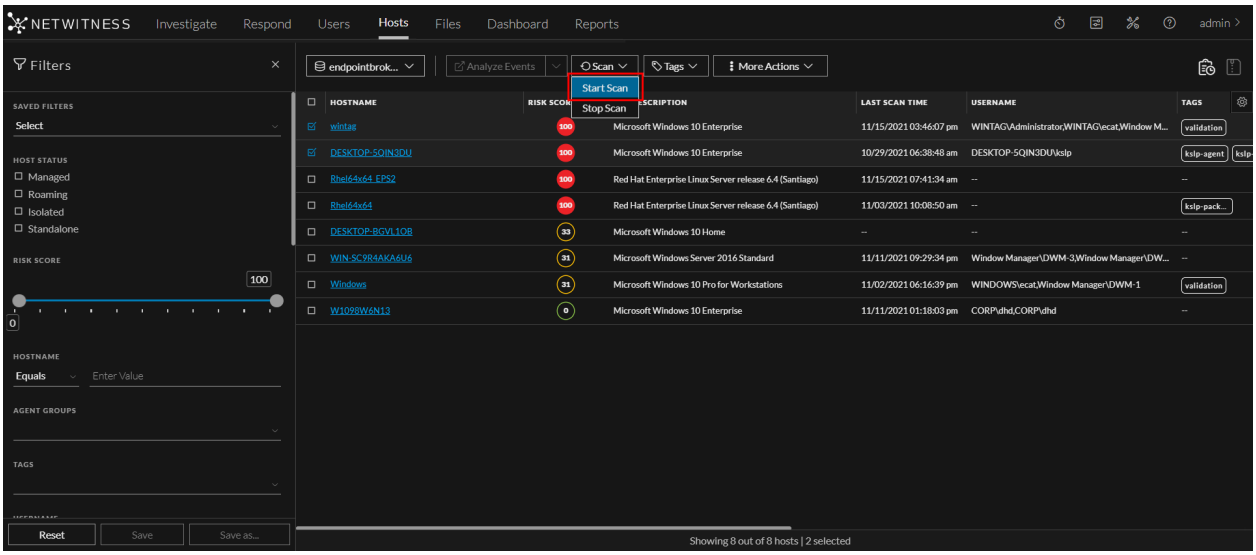
The following are the scan statuses:

Status	Description
Idle	No scan is in progress.
Scanning	Scan is in progress.
Pending	Scan request is sent to the server, and the agent will receive the request the next time it communicates with the server.
Cancel	Stop request is sent to the server, and the agent will receive the request the next time it communicates with the server.

Note: By default, the scan utilizes 25% of the CPU. You can click CPU Maximum and select a value from 5% to 100%. Increasing the CPU Maximum limit reduces the scan time but could lead to more CPU usage.

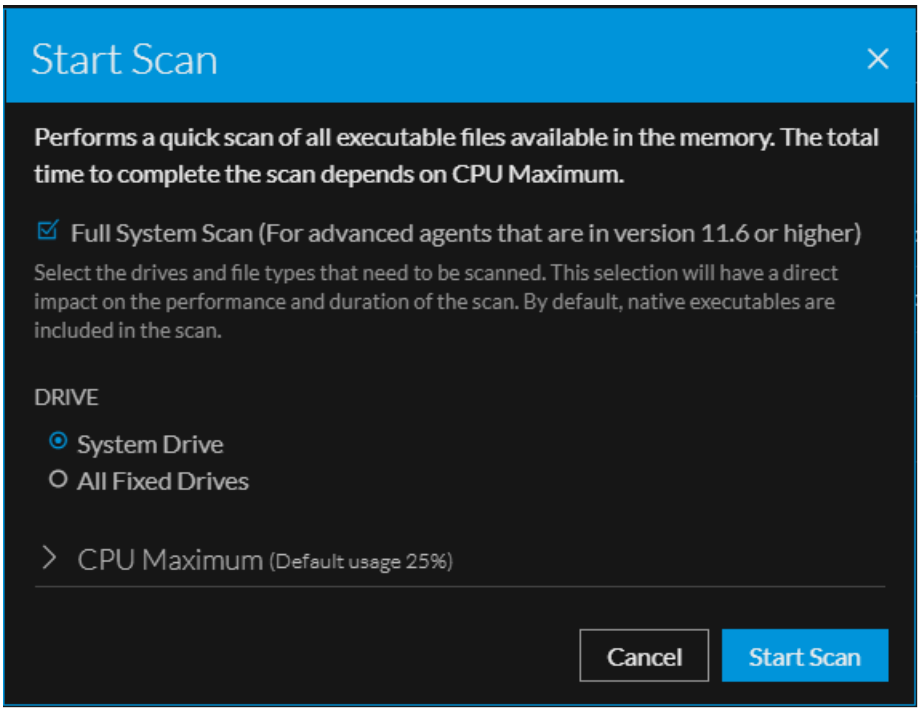
To perform a Full System Scan:

1. Go to **Hosts**.
2. Select one or more hosts (up to 100) at a time for an on-demand scan, and do one of the following:
   - Click **Scan > Start Scan** from the menu bar (Or)
   - Right-click and select **Start Scan** from the context menu



1. From the **Start Scan** pop-up, select **Full System Scan (Only on advanced agents that are 11.6 or higher.)**

2. Select **System Drive**(Default selection) or **All Fixed Drives**
3. Click **Start Scan** on the pop-up.



Note: An Endpoint server supports up to 10k Full System Scans in a rollover period.

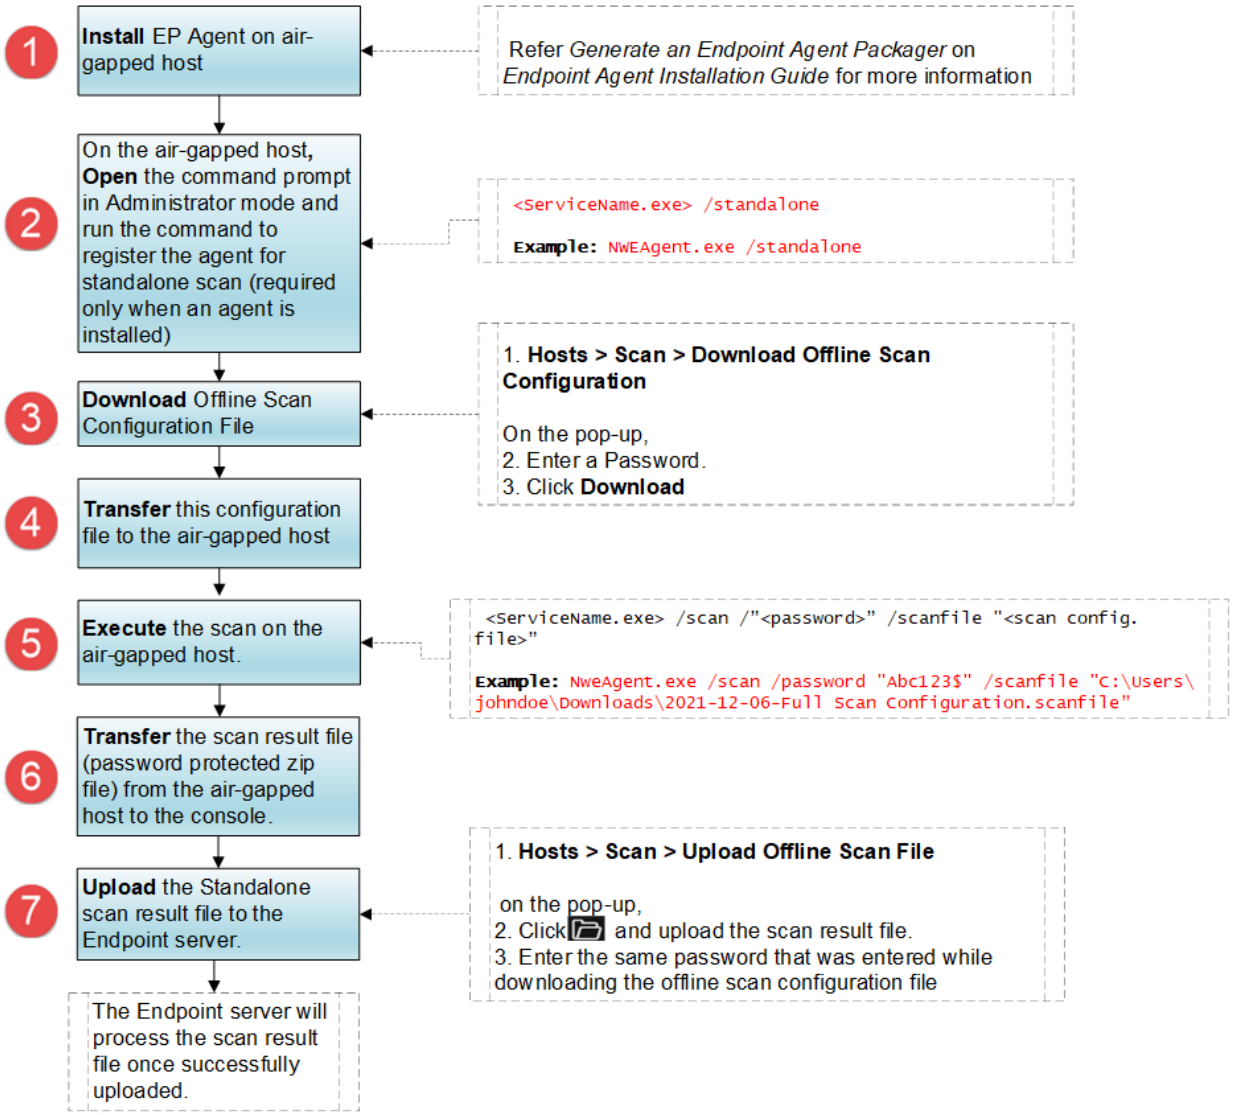
Standalone Scan on Air-gapped Windows Hosts

The Standalone scan feature allows administrators to run scans on the air-gapped Windows hosts that are disconnected from the network. Administrators can download the scan command once from UI and execute it on multiple hosts. For the best utilization of resources, we recommend running standalone scan every two weeks. Policies do not apply on air-gapped hosts, and features such as downloading MFT to the server, upgrading agents through UI, downloading a file to the server are not available for standalone agents.

The scan process involves two files:

- **Offline Scan Configuration** - Contains the configuration information needed to run the scan.
- **Scan Results File** – This contains the results of the Scan, which you can upload using the **Scan > Upload Offline Scan File** option on the **Hosts** view. This file will be imported and processed by NetWitness.

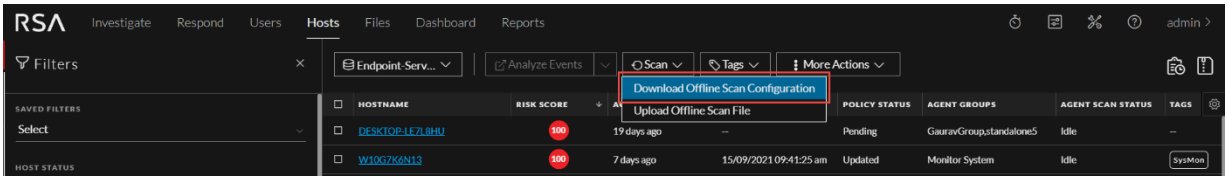
Standalone scan workflow:



Note: Both Download Offline Scan Configuration and Upload Offline Scan File options are available only on the Endpoint server view. These options can't be accessed from Broker view.

Generate the scan configuration file

1. Click **Scan > Download Offline Scan Configuration** on the Hosts screen.

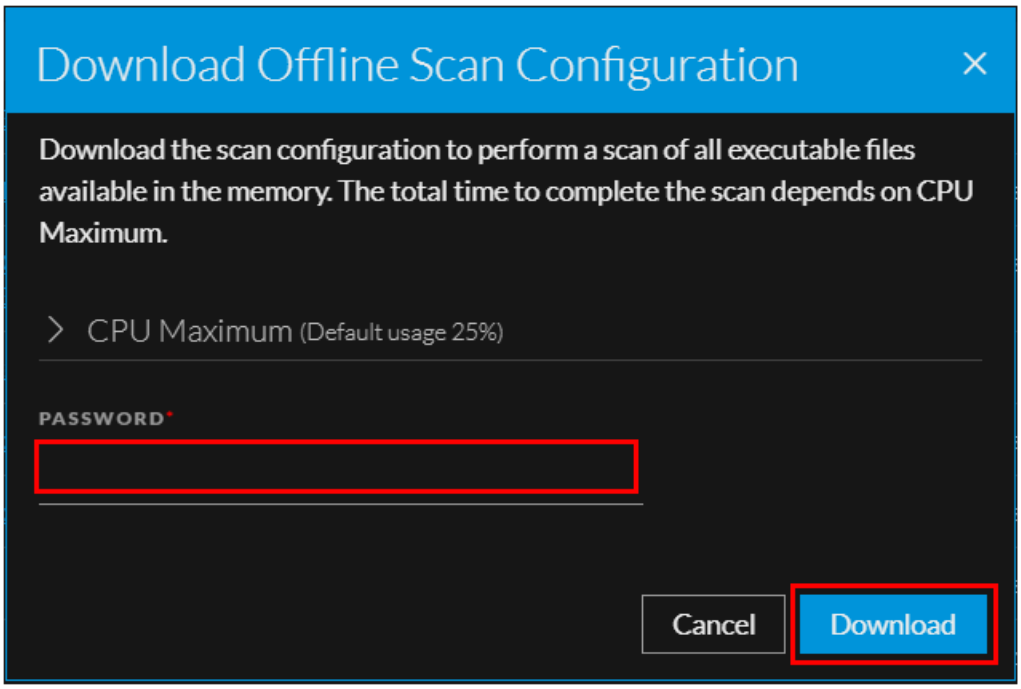


On the Download Offline Scan Configuration pop-up,

2. (Optional) **Select CPU Maximum**.

3. Enter a Password. (not more than 31 characters long)

4. Click **Download**.



5. Transfer the Offline Scan Configuration file to the air-gapped host.

Install Endpoint Agent and Register for Standalone Scan

1. Install the Endpoint Agent on the air-gapped host. Refer to *Endpoint Agent Installation Guide* for more information.
2. Register the agent for standalone scan (required only when an agent is installed)
   - Open command prompt in administrator mode and execute the following command:

     `ServiceName.exe /standalone`

     Example: `NweAgent.exe /standalone`

Start a Standalone scan

1. Open the command prompt in Administrator mode, on the air-gapped host.

2. Execute the scan using the following command(syntax):

```
ServiceName.exe /scan /password ”<password>” /scanfile "<scan config. file>”
```

Example: `NweAgent.exe /scan /password "Abc123$" /scanfile "C:\Users\johndoe\Downloads\2021-12-06-Full Scan Configuration.scanfile"`

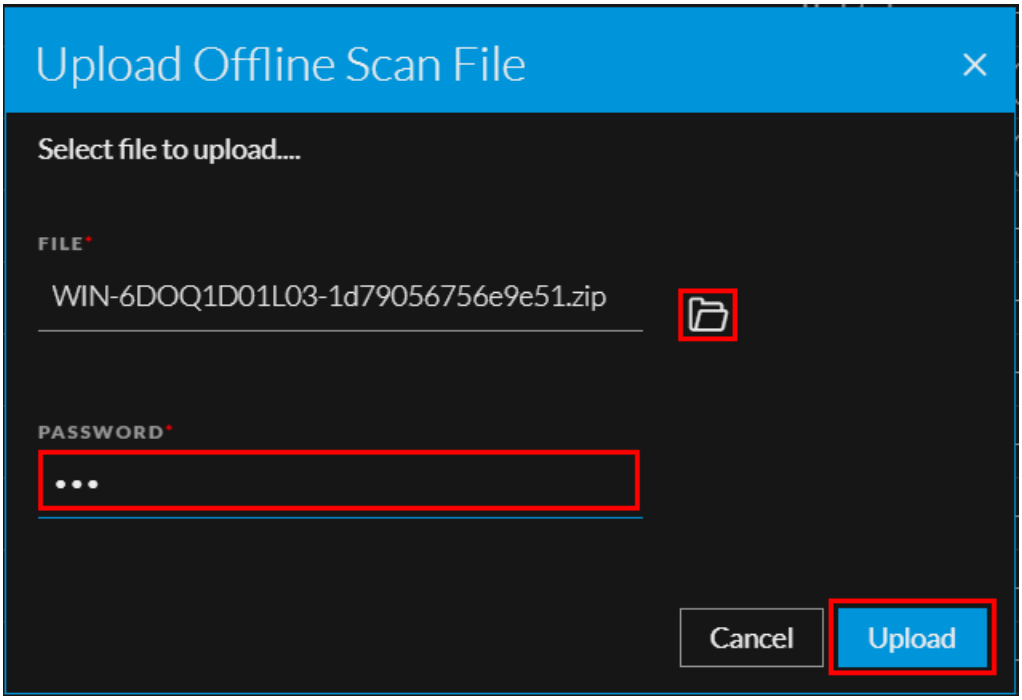
- `<password>` is the password entered while generating the Offline Scan Configuration File.

- `<scan config file>` is name of the scan configuration file with full path.

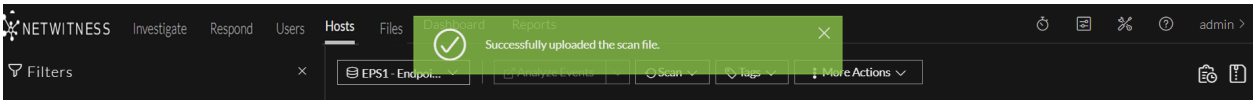
3. Wait until the scan is completed.
4. Transfer the scan result file(password protected .zip file) back to the machine to upload to the UI.

Upload the Standalone scan result file

1. Click **Scan > Upload Offline Scan File** on the Hosts screen.



2. Click  and upload the scan result file.
3. Enter the same password that was entered while downloading the offline scan configuration file.
4. The Endpoint server will process the scan result file once successfully uploaded.



Note: Standalone agents can only be upgraded manually using the Endpoint agent packager. Refer to *Generate an Endpoint Agent Packager* on *Endpoint Agent Installation Guide* for more information.

Analyze Hosts Using the Risk Score

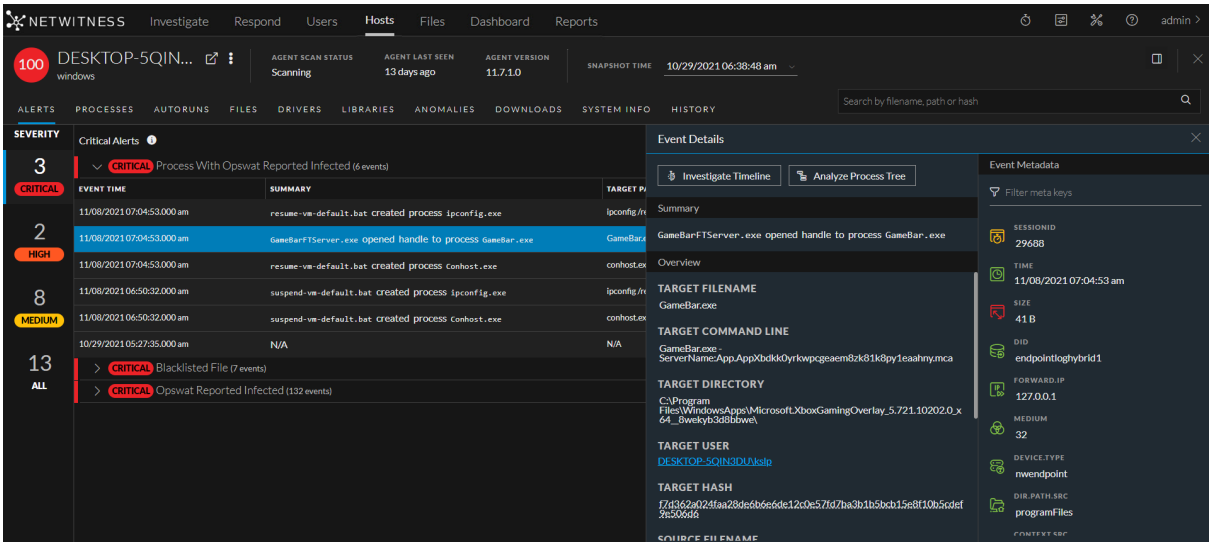
You can investigate a host by analyzing the risk contributors such as alerts and events to look for suspicious or malicious activity.

To analyze the hosts using the risk score:

1. Go to **Hosts**.

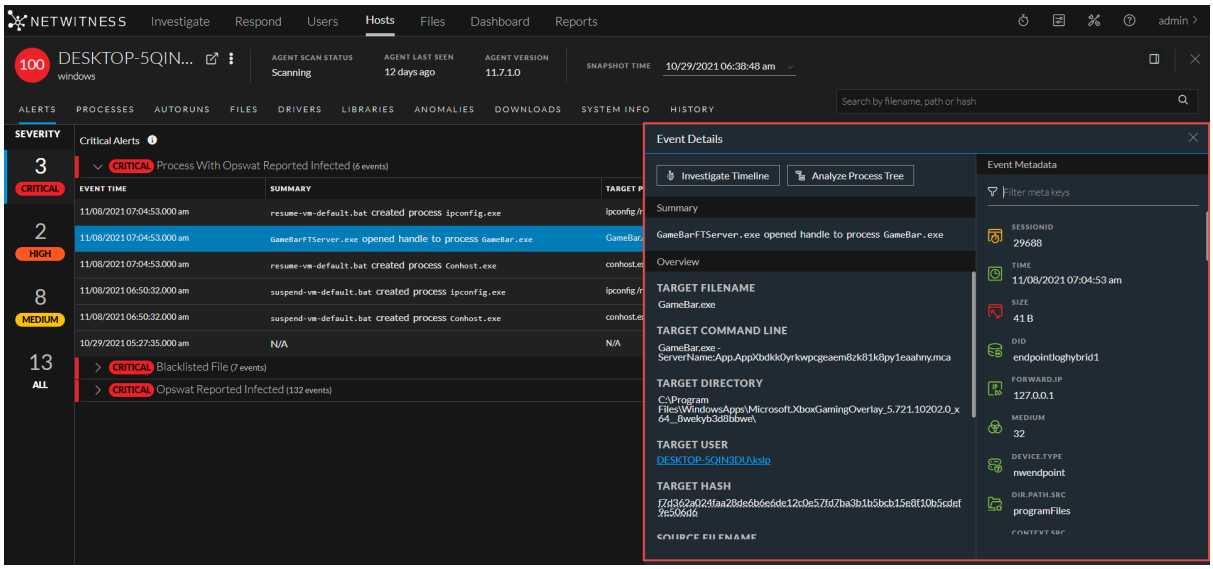
   The Hosts view is displayed.
2. In the Server drop-down list, select the Endpoint server or Endpoint Broker server to view the hosts.
3. Select the host and do any of the following.
   - Click a row to view the risk associated with the host in the **Risk Details** panel.
   - Click the hostname to investigate the host.
4. In the **Alerts** panel, click the alert severity such as **Critical**, **High**, **Medium**, or **All**.

   The list of distinct alerts is displayed along with the total number of events associated with the alert.
5. Click an alert to view the associated events.

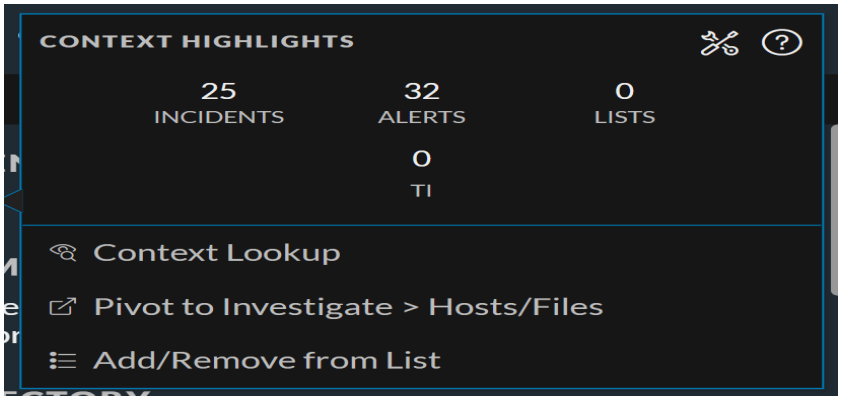


Note: Only the latest 1000 events are displayed.

6. To view all the metadata associated with a specific event, click on an event. The information such as summary, target path, filename, and others is displayed on the **Event Details** panel.



7. Event Metadata such as IP, Hostname, Mac, Filename, File hash, User, and Domain is displayed on the Event Details panel. Scroll down to view all related information. These are the possible data sources: NetWitness Endpoint, Incidents, Alerts, Hosts, Files, and Feeds.



8. To investigate the original event and destination domain of the event, do any of the following:

   - To investigate the events in a specific time frame, click the **Investigate Timeline** on the **Event Details** panel. For more information, see the *NetWitness Investigate User Guide*.
   - To view a list of processes captured on the hosts and investigate a particular process, click the **Analyze Process Tree** on the **Event Details** panel. For more information on process analysis, see Investigating a Process.

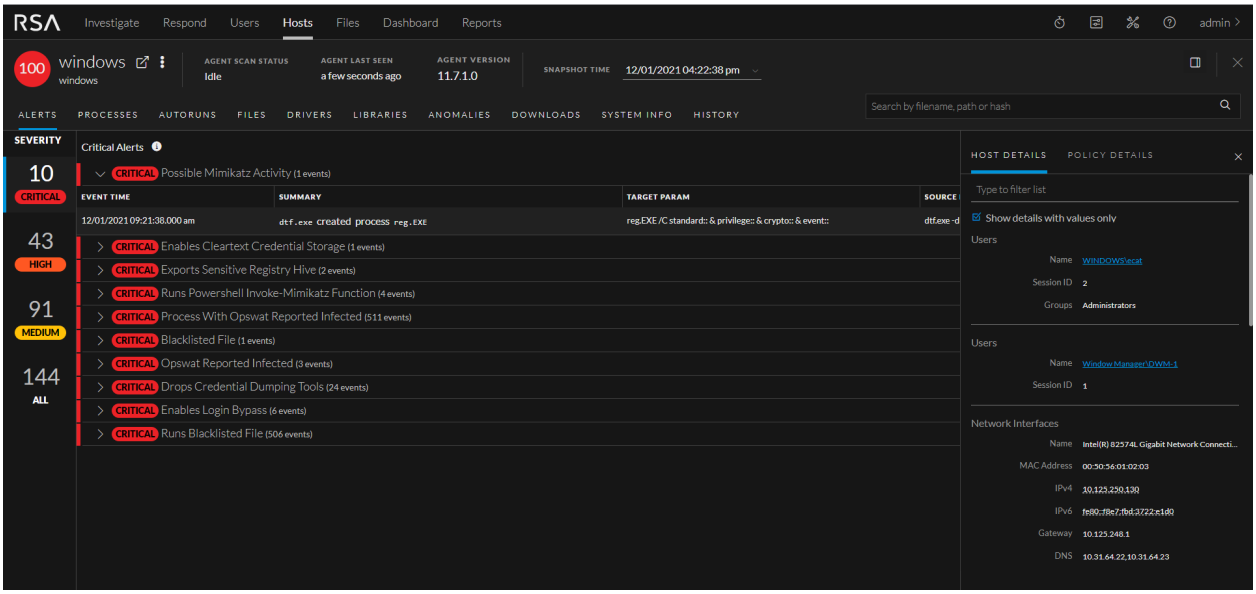
Analyze Host Details

To look for suspicious files on a host, click the host name and view the details of the host, or start an on-demand scan to get the most recent information. On the right-hand panel, you can view the following:

- **Host Details** displays the host information, such as Network interface, operating system, hardware and others.

- **Policy Details** displays the complete resolved policy settings.

For more details, see Hosts View - Details Tab.



Filter Host Details

In the Processes, Autoruns, Files, Drivers, Libraries, and Anomalies tabs, you can filter the processes or files on file status, reputation, file or process name, signature, and risk score. Click **Save** to save the filter and provide a name (up to 250 alphanumeric characters). The filter is added to the Saved Filters panel on the left. To delete a filter, hover over the filter name and click 🗑.

In the **Host** view > **Files** tab, you can filter the files available on host, and files deleted from host. The result of files deleted from host depends on the data retention policy configured in the **Endpoint Config** view > **Data Retention Scheduler** tab. By default, data retention policy is configured for 30 days, this means only 30 days of deleted files are stored in the Endpoint server. These filter options are disabled if **All Files Available on Host** toggle is disabled.

In the **Host** view > **History** tab, you can filter the commands on command type, status, host name, request type, command parameter and command time. In the Command Time field, you can filter by custom date range.

Click **Save** to save the filter and provide a name (up to 250 alphanumeric characters).The filter is added to the **Saved Filters** panel on the left. To delete a filter, hover over the filter name and click 🗑.

Note: Special characters are not allowed except underscore (_) and hyphen (-) while saving the filter.

Search Files on Host

To investigate a host or to check if it is infected with a known malware, you can search for occurrences of the file name, file path, or SHA-256 checksum.

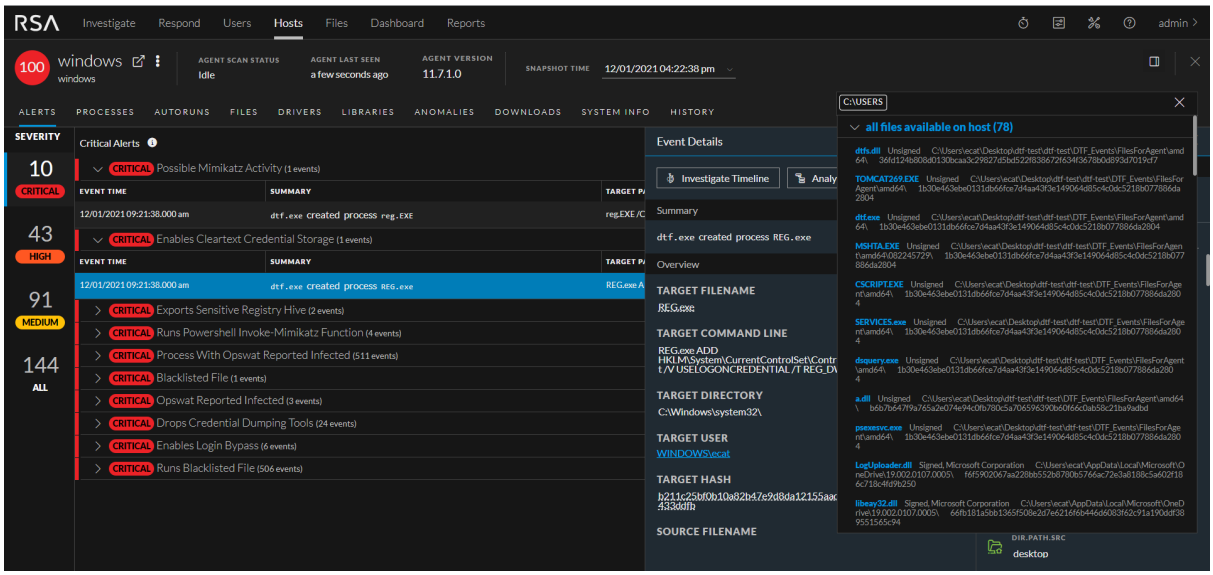
Note: To search for a SHA-256 checksum, provide the entire hash string in the search box.

The result displays the matching files present on the host in **All Files Available on Host** category and in the respective **snapshot** category with the details, such as file name, signature information,and checksum. In addition, the **snapshot** category displays the system interaction, for example, ran as process, library, autorun, service, task, or driver. To view more details, click the **filename** or **system interaction** link.

Example, a user has clicked and executed a malicious attachment through a phishing email, and downloaded it to `C:\Users`. To investigate this file:

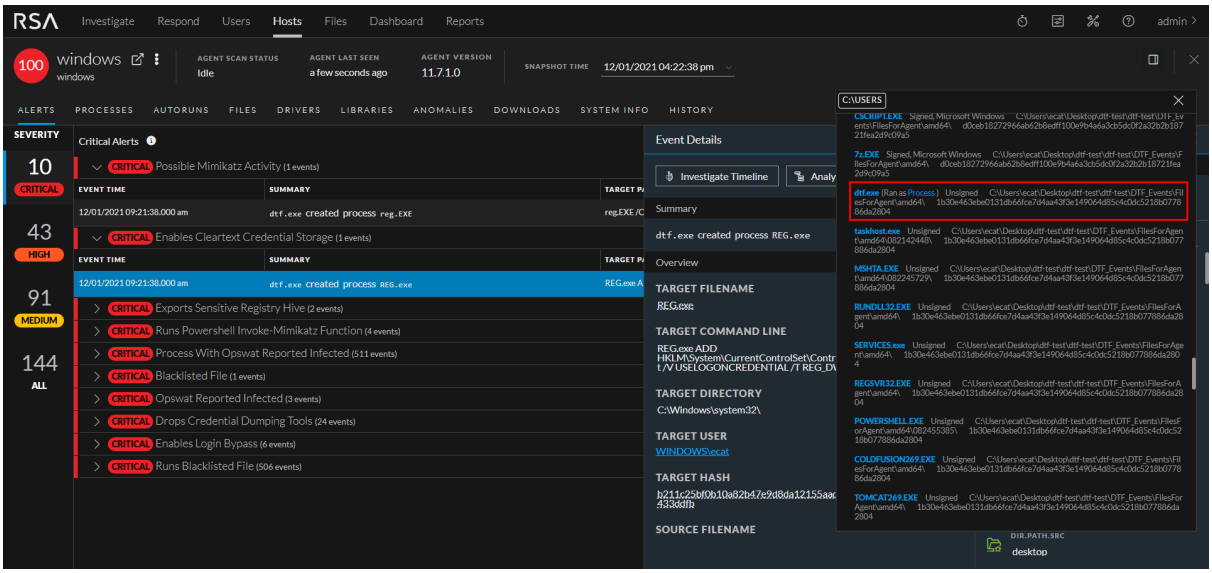
1. Go to **Hosts**.
2. Select the host that you want to investigate or select the Endpoint Broker server to investigate all the hosts.
3. In the **Details** tab, enter the file path `C:\Users` in the search box.

   The search displays a maximum of 100 results of the executables in this folder. In this example, there are some unsigned file that might be malicious. If the search is executed on an Endpoint Broker server, it queries all the Endpoint servers.



   This file is run as a Process.

4. To view details of this file, click **Process** in the result.

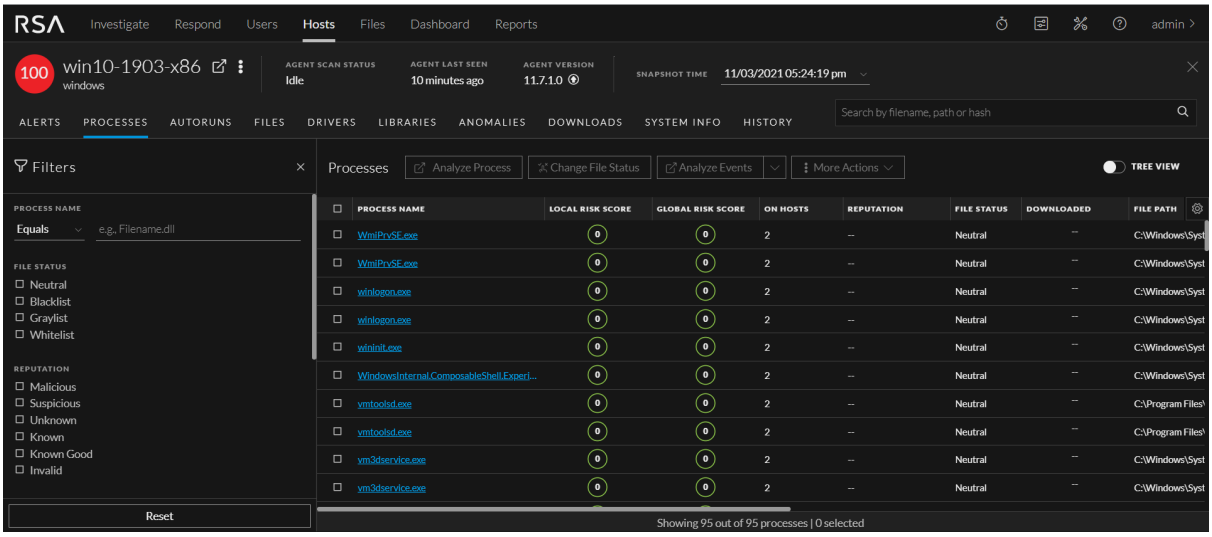


This opens the Process tab where you can view the process details.

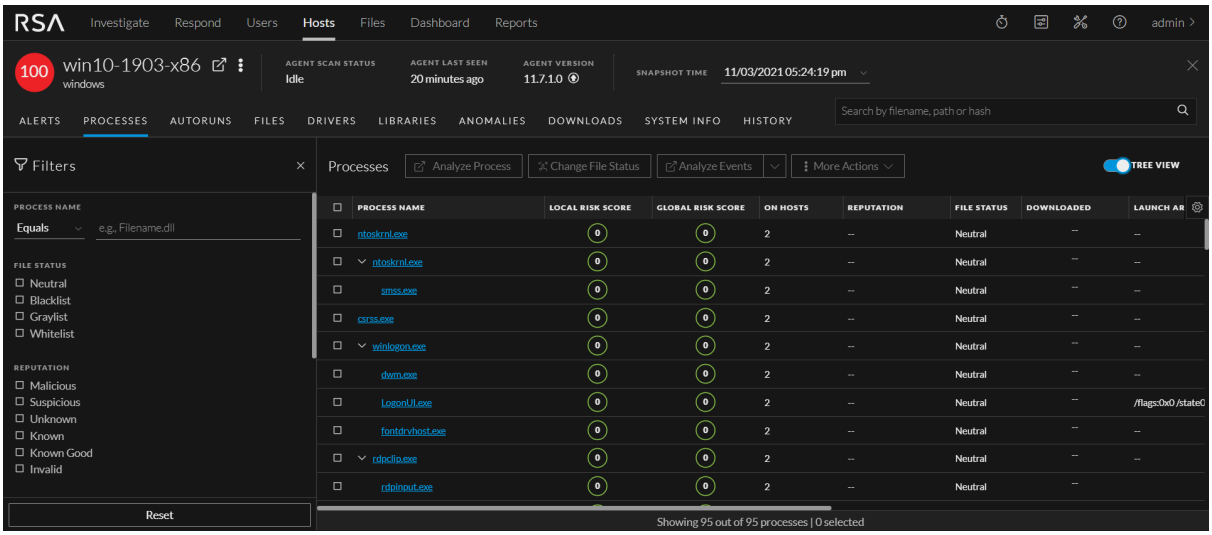
Analyze Processes

To analyze the process:

1. In the **Hosts** details, select the **Processes** tab.

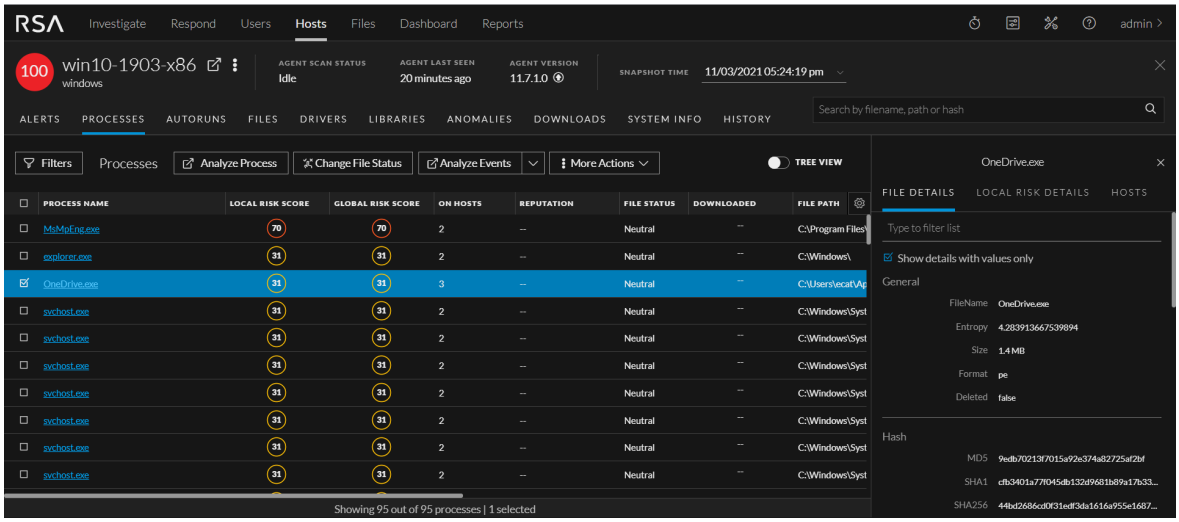


To view the process tree, click the toggle switch. The following is an example of the tree view:

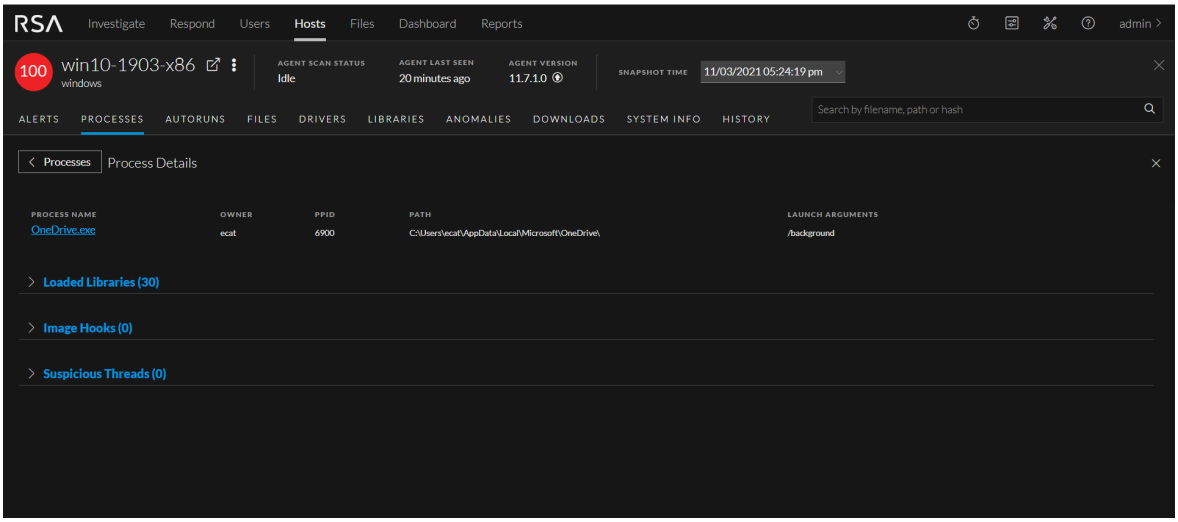


2. In the **Processes** Tab, do one of the following:

   - Click a row to view the properties of a process in the right panel.



- Click the process name to view the process details of a specific process.



When reviewing processes, it is important to see the launch arguments. Even legitimate files can be used for malicious purposes, so it is important to view all of them to determine if there is any malicious activity.

For example,

- `rundll32.exe` is a legitimate Windows executable that is categorized as a good file. However, an adversary may use this executable to load a malicious DLL. Therefore, when viewing processes, you must view the arguments of the `rundll32.exe` file.
- `LSASS.EXE` is a child to `WININIT.EXE`. It should not have child processes. Often malware use this executable to dump passwords or mimic to hide on a system (`lass.exe`, `lssass.exe`, `lsasss.exe`, and so on).
- Most legitimate user applications like Adobe and Web browsers do not spawn child processes like `cmd.exe`. If you encounter this, investigate the processes.

You can view the sequence of activities performed on the host by the file or process using the process analysis. For more information, see Investigating a Process.

Analyze Autoruns

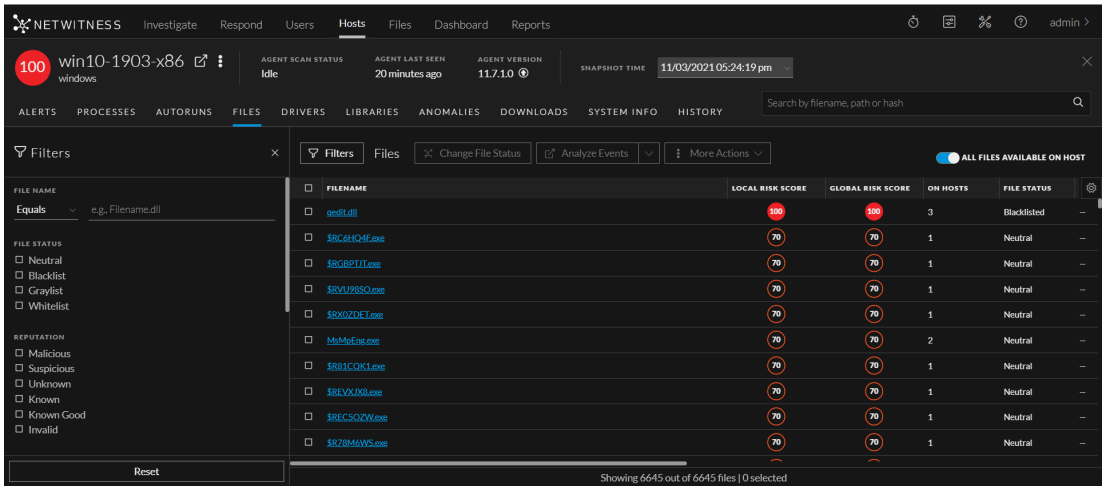
In the Hosts details, select the **Autoruns** tab. You can view the autoruns, services, tasks, and cron jobs that are running for the selected host.

For example, in the Services tab, you can look for the file creation time. The compile time is found within each portable executable (PE) file in the PE header. The time stamp is rarely tampered with, even though an adversary can easily change it before deploying to a victim's endpoint. This time stamp can indicate if a new file is introduced. You can compare the time stamp of the file against the created time on the system to find the difference. If a file was compiled a few days ago, but the time stamp of this file on the system shows that it was created a few years ago, it indicates that the file is tampered.

Analyze Files

To analyze the files, you can do either one of the following based on your requirement.

- In the **Hosts** view, select the **Files** tab.
  You can view the list of all files (reported as part of scan and tracking) on the host including the deleted files.



- To view the files reported as part of scan snapshot, you must disable **All Files Available On Host** toggle and select the scan time form the **Snapshot** drop-down list.

Example for analyze files, many trojans write random filenames when dropping their payloads to prevent an easy search across the endpoints in the network based on the filename. If a file is named `svch0st.exe`, `scvhost.exe`, or `svchosts.exe`, it indicates that the legitimate Windows file named `svchost.exe` is being mimicked.

Analyze Libraries

In the Hosts details, select the **Libraries** tab. You can view the list of libraries loaded at the time of scan.

For example, a file with high entropy gets flagged as packed. A packed file means that it is compressed to reduce its size (or to obfuscate malicious strings and configuration information).

Analyze Drivers

In the Hosts details, select the **Drivers** tab. You can view the list of drivers running on the host at the time of scan.

For example, using this panel, you can check if the file is signed or unsigned. A file that is signed by a trusted vendor such as Microsoft and Apple, with the term `valid`, indicates that it is a good file.

Analyze Anomalies

Note: This tab is available only for advanced agent.

In the Hosts details, select the **Anomalies** tab. You can view the following details for the selected host:

- Image hooks - Hooks found in executable images (user-mode or kernel-mode) - IAT, EAT, Inline, exceptionHandler.

- Kernel hooks - Hooks found on kernel objects (such as Driver Object [Pointers, IRP_MJ, SSDT, IDT, and so on]). This also includes filter devices.
- Suspicious threads - Threads whose starting address points to memory DLLs or floating code. The threads could be running with either user-mode or kernel-mode privileges. These threads could run malicious code inside a trusted application to execute their own code.
- Registry discrepancies - The Windows registry is a hierarchical database that stores configuration settings and options on Microsoft Windows operating systems. It contains settings for low-level operating system components and for applications running on the platform: the kernel, device drivers, services, SAM, user interface, and third party applications all use the registry. The discrepancies between low-level parsing with Win32 registry API are reported.

Note: Anomalies is applicable only for Windows hosts.

For example, hooking is used to intercept calls in a running application and to capture information related to the API invocations. Malicious programs can implant hooks in various system applications for different purposes, such as hiding files, directories, registry entries, intercepting users keystrokes to establish a stealthy communication channel with the attacker.

Analyze System Information

In the Hosts details, select the **System Information** tab. This panel lists the agent system information. For Windows operating system, the panel displays the host file entries and network shares of that host.

For example, malware might use host file entries to block antivirus updates.

Analyze History

In the Host details, select the **History** tab. This tab lists the commands along with the respective status and additional details.

When you review the history, look for the command status and retrieval count to check if the agent retrieved the commands.

Below are some examples:

- A file download command is issued, but the file is deleted on the host. In this case status of the command is failed as the file is not downloaded.
- The retrieval count increases, but the command is not processed. This happens when an analyst requests a large number of files (For example, MFT, system dump, or process dump), and the connection breaks when the agent uploads these files.
- If the agent command is not retrieved, the agent is either offline or busy processing other commands (For example, uploading a system dump). In this case, the status of the command shows pending.

To view more details, click the **Hostname** link highlighted in blue. The Hosts details view is displayed. In the case of MFT, download file, system dump, and process dump command types, **Downloads** tab is displayed with details such as file name, type, status, size, downloaded time and SHA256 of the file, when you click on the **Hostname** link.

Export Host Details or Files to JSON File

Note: Export Host details option is disabled if there is no snapshot time.

To export host details or files to JSON file:

1. Go to **Hosts**.
2. Select the hostname to open the host details.
3. Click (**More**) beside the hostname and do any of the following:
   - To export the scan data categories for the host, select **Export Host Details**. This exports files such as:
     - allfiles.json - This file consists of the file name, file path, signature, file checksum, and so on that is reported as part of scan and tracking.
     - fileContext.json - This file consists of the file name, file path, signature, file checksum, and so on that is reported during the host scan.
     - machinedetails.json - This file consists of the machine details, including hardware, operating system, interfaces, and so on, along with the agent details like version, policy details.

   Note: If Endpoint Broker is selected and a host is communicated with multiple Endpoint servers, during the host details export, all files and details of the host are exported from the Endpoint server where the selected snapshot is stored.

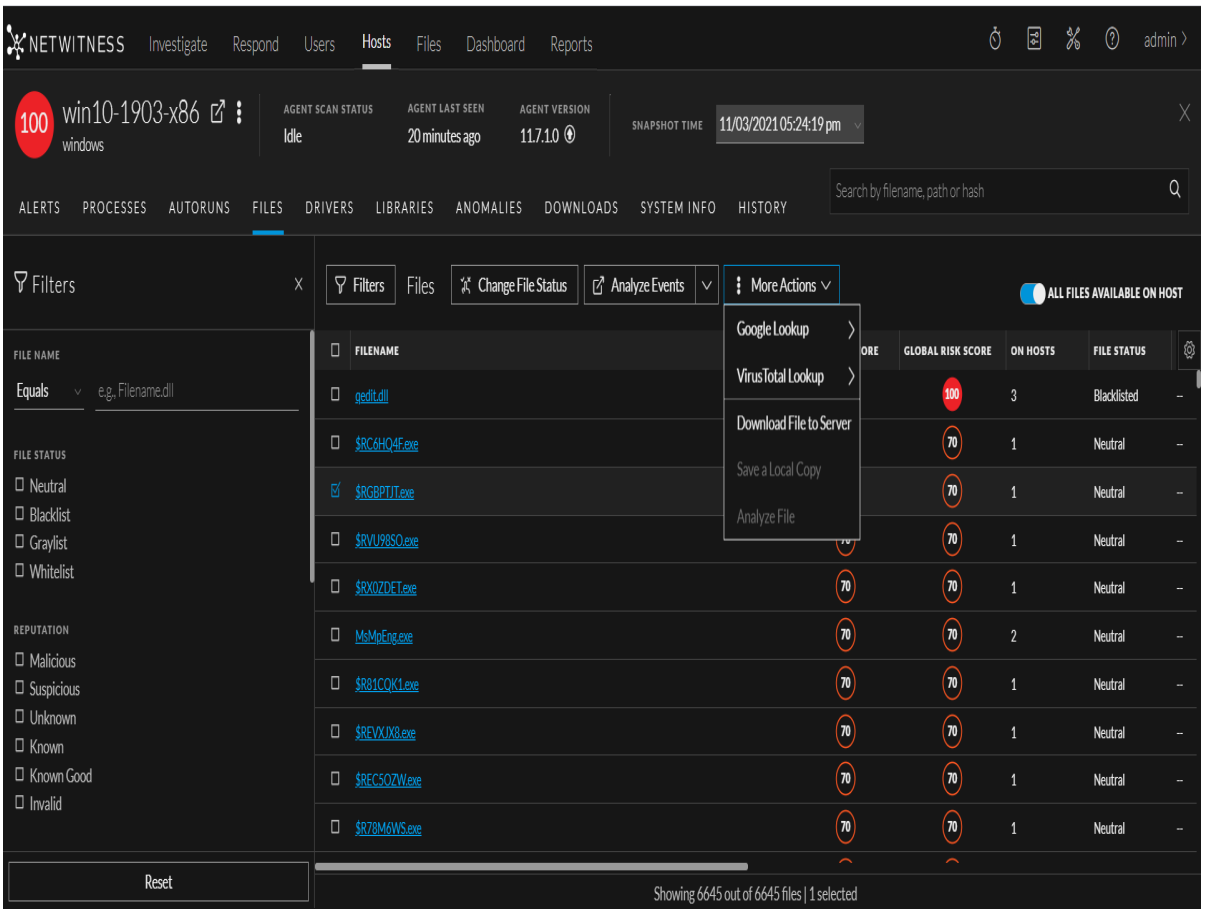
   Note: allfiles.json file is exported irrespective of the selected snapshot.

   - To export all the files available on the host, select **Export Files**. This exports:
     - allfiles.json - This file consists of the file name, file path, signature, file checksum, and so on that is reported as part of scan and tracking.

Launch an External Lookup for a File

While analyzing a file, you can search Google or VirusTotal with the filename or hash to get more information about the file. To launch the search:

1. Go to **Hosts** > **Host Details** (Autorun, Files, Drivers, Libraries, or Anomalies tab).
2. Right-click one or more files, or in the **More Actions** drop-down list in the toolbar, do the following:



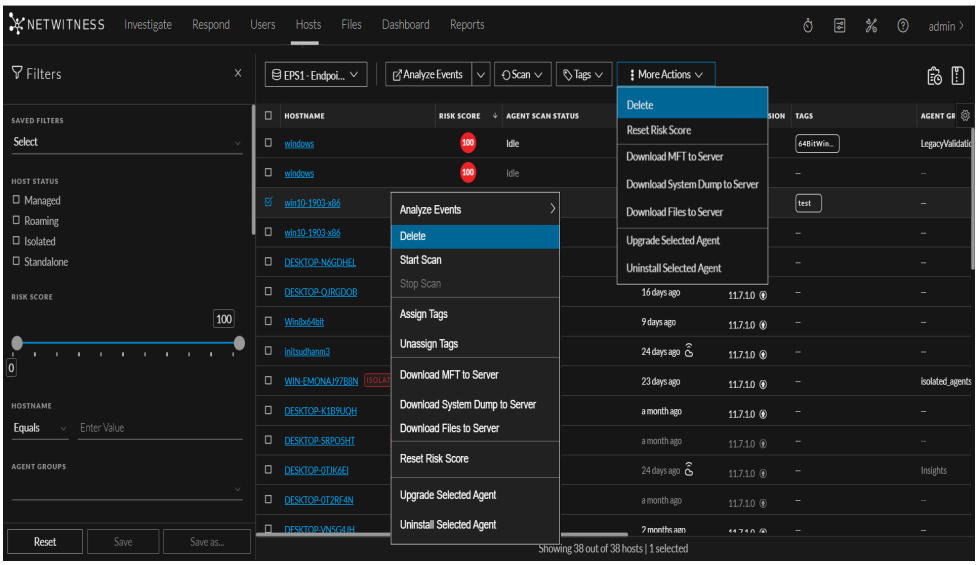
- Select **Google Lookup** to perform a search on the filename, MD5, SHA1, or SHA256.
- Select **VirusTotal Lookup** to perform a search on MD5, SHA1, or SHA256.

Note: To open files in multiple tabs, make sure you enable the pops-up in the browser.

Delete a Host

If the agent is uninstalled on a host or if you no longer require the host scan data, you can manually delete this host from the Hosts view. Deleting a host deletes all scan data associated with the host. To delete hosts:

1. Go to **Hosts**.
2. Select the hosts that you want to delete from the Hosts view and do one of the following:



- Right-click and select **Delete** from the context menu.
- Click **More** drop-down list in the toolbar and select **Delete**.

Note: If you accidentally delete a host from the Hosts view, the Endpoint Server forbids all requests from this agent. The agent must be uninstalled manually from the host and reinstalled for it to appear on the Hosts view.

Deleting Hosts with Older Agent Versions

After upgrading the 11.1.x and 11.2.x agents to 11.3 or later, if you want to delete the hosts with older versions:

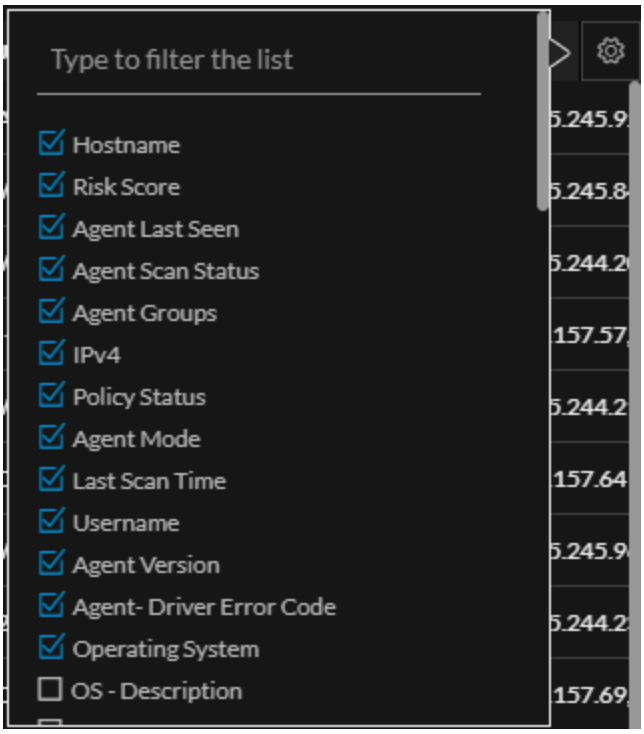
1. Go to **Hosts** view.
2. Filter the hosts based on the Agent version, and delete these hosts.

   If you do not delete, the hosts are deleted based on the Data Retention Policy settings.

Set Hosts Preference

By default, the Hosts view displays a few columns and the hosts are sorted based on the risk score. If you want to view specific columns and sort data on a specific field:

1. Go to **Hosts**.
2. Select the columns by clicking in the right-hand corner. The following example shows the drop-down list displayed while adding columns:



3. Scroll down or enter the keyword to search for the column in the displayed list.
4. Sort the data on the required column.

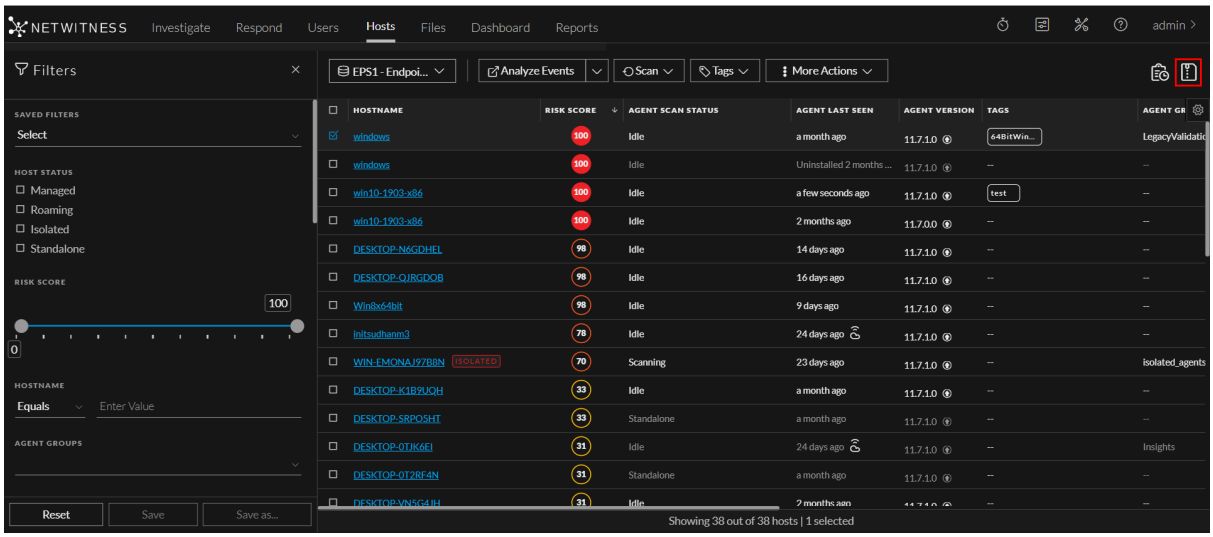
Note: The selections you make here become your default view every time you log in to the Hosts view.

Export Host Attributes

You can export up to 100,000 host attributes at a time. To extract the host attributes to a csv file:

1. Go to **Hosts**.
2. Filter the hosts by selecting the required filter options.
3. Add columns by clicking in the right-hand corner.

4. Click  to export the host attributes to a csv file.



You can either save or open the csv file.

Migrate Hosts

Hosts can be migrated from one Endpoint server to another using groups and policy associated with the host. If a host is migrated, the Server column shows as **Migrated**. On all the tabs within the Hosts view, the message Host is migrated to <Server-name> is displayed. You view the host details by clicking the <Server-name>. The risk score of a migrated host is displayed on all Endpoint servers where it is present.

Note: Some of the actions are disabled for the migrated host on the selected server, such as start scan, start stop, analyze events, and others. If you want to perform the required action, select the Endpoint server to which the host is migrated.

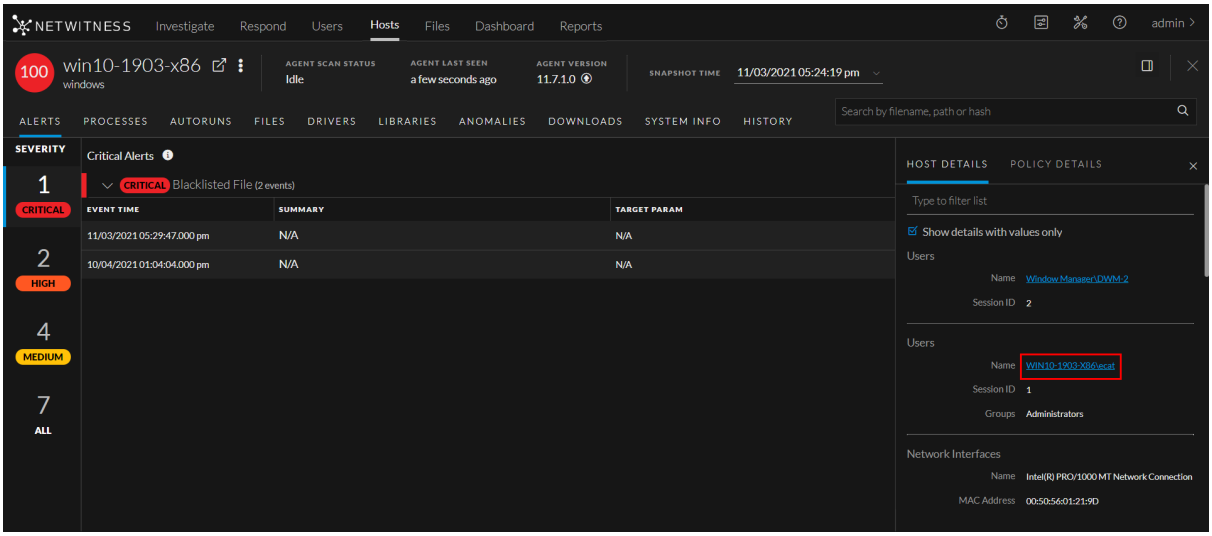
Note: To view only managed hosts, select the **Show Only Managed Agents** option in the Filters panel.

Analyzing Risky Users

If you have NetWitness UEBA installed, you can view the alerts associated with users logged in on the host. To analyze risky users:

1. Go to **Hosts**.
2. Click the host name you want to analyze.
3. In the **Host Details** panel, under the Users category, click the name.

This opens the Entities tab for investigation in a new tab.



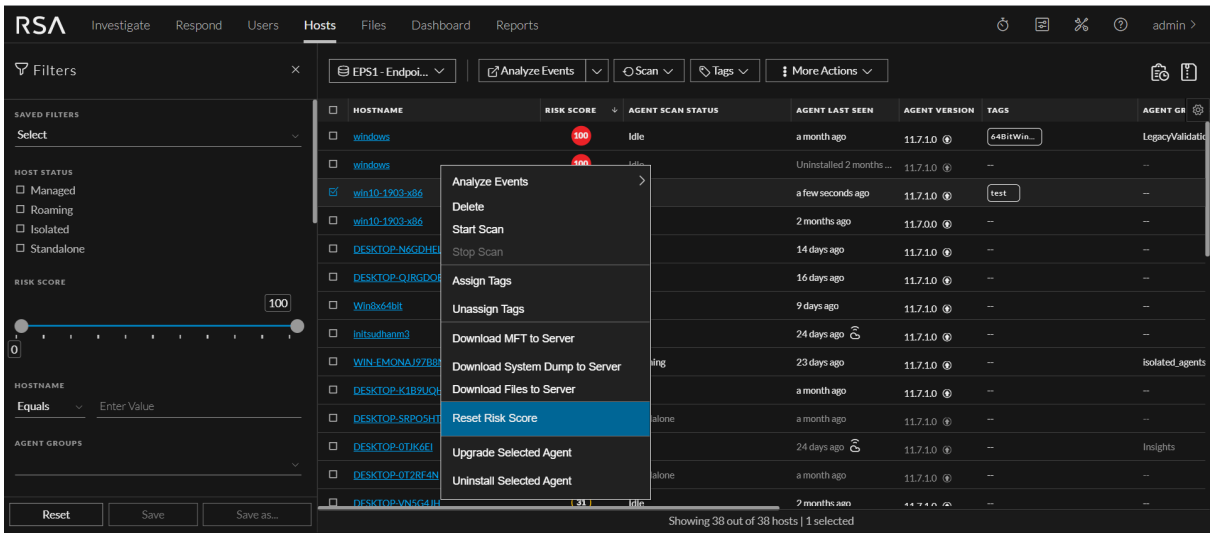
Resetting Risk Score of Hosts

You can reset the risk score for a host in these situations:

- If the alerts or events triggered by the host or files on the host are false positive, you can make changes to the Endpoint Application rules or ESA rules.
- After you take required action on the host for malicious file activities contributing to the risk score. When you reset the risk score, all the risk calculation for the host is deleted. When you reset the host's risk score, it does not change the file's risk score. You can reset the score for a single host or multiple hosts.

To reset the risk score of a host:

1. Go to **Hosts**.
2. Select the Endpoint Server or Endpoint Broker.
3. Select one or more hosts and do one of the following:



- Right-click and select **Reset Risk Score** from the context menu.
- Click **More Actions** > **Reset Risk Score** in the toolbar.

All the alerts associated with the score are deleted.

Note: You can select a maximum of 100 hosts to reset the score.

4. Refresh the page to view and confirm if the host's score is reset. This may take sometime for changes to take effect.

Investigating a Process

Note: The information in this topic applies to NetWitness Version 11.3 and later.

Analysts can perform process analysis to investigate a particular process behavior to:

- Understand the entire process event chain, process parent-child relationships, and all associated events in a timeline view.
- Analyze important process attributes, such as username, launch arguments, reputation, file status, signer, signature, risk score, and file path.

The Analyze Process view provides a list of processes captured on hosts in a parent-child hierarchical format over a time range. The process tree is created from the tracking event type "Process event" where the action meta key is `createProcess`. The agent reports new events for the same createProcess if the following parameters change:

- Parent process filename
- Child process filename
- Launch arguments
- User name

If the above parameters do not change, the event is reported only once every eight hours.

Best Practices

When reviewing a host for malicious activity, there are a few key things to review while looking for malicious processes.

- Process Name - When reviewing running processes on a host, check for the name of the program that looks suspicious. Sometimes malware uses random names, such as `wzuduje.exe`. In some cases, the names might be misleading such as `adob3.exe`, `scvhost.exe`, or `Microsoft.exe`. Being familiar with Windows processes and any type of internal tool that might be used throughout the environment, also helps you to identify potentially malicious or suspicious files.
- File Path - Similar to knowing normal and key Windows processes, knowing what path the processes originate from is a key to detect certain processes that imitate the legitimate process. For instance, if you see `svchost.exe` running on a system from `C:\Users\<username>\AppData\Roaming\adobe\` (which is a valid file path), and knowing that the legitimate Windows process originates from `C:\Windows\System32\`, you can determine that the `svchost.exe` file starting from the `C:\Users\<username>\AppData\Roaming\adobe\` directory is the suspicious one. To help determine further identification of a suspicious process, review the Autoruns tab to see if this process is running as an autorun, service, or task.
- File Signature - When a software package is created, it has a valid digital signature. The following are a few exceptions:
  - If a process that is running is not digitally signed, it does not automatically confirm that the file is malicious.

  - While files may have a valid signature, it does not mean that they are legitimate. There are instances of software identified as a Potentially Unwanted Program (PUP) or Adware, which can have a valid signing certificate.
- On Hosts - Indicates the number of hosts on which a file exist. If a file is present on fewer hosts with a high risk score, it may be malicious and needs further investigation.
- Reputation - Leveraging the reputation service is a way to find malicious processes.
- Analyze events - For further insight to a process, you can analyze console events, network events, file events, process events, and registry events.
  - Network events - Look for any suspicious domains to which the process is connecting. Sometimes malware creates legitimate connections to a known site, such as google.com, bing.com to hide its activity on the network. Look for connections to Dynamic DNS domains where a lot of known malicious activity resides. During analysis, consider uncommon processes making direct connections to an IP address or to a uncommon port number.
  - File and process events - Review process interactions that have occurred on the system with the suspected file. You can look for key events such as `writeToExecutable`, `renameExecutable`, and `createRemoteThread`, which indicate suspicious behavior.
- Leverage other methods
  - Look up with Google - You can search the file name or hash value against Google to determine if the file is malicious.
  - Look up with VirusTotal – You can search the hash value against the VirusTotal to determine if the file is malicious between multiple AV vendors.
  - Download file – Download and analyze a file to find indicators such as compile time, imported DLLs, section names, and performing string searches. Look for TLD values (.com, .net, .biz) or debug information of a compiled binary (.pdb), which can be easily changed or forged.
  - Time stamp values – Review modified, accessed, and created dates associated with the binary. Review how long a file has been residing on a host. While this value is correct most of the time, attackers can change the time stamp values of a file.

Analyze a Process

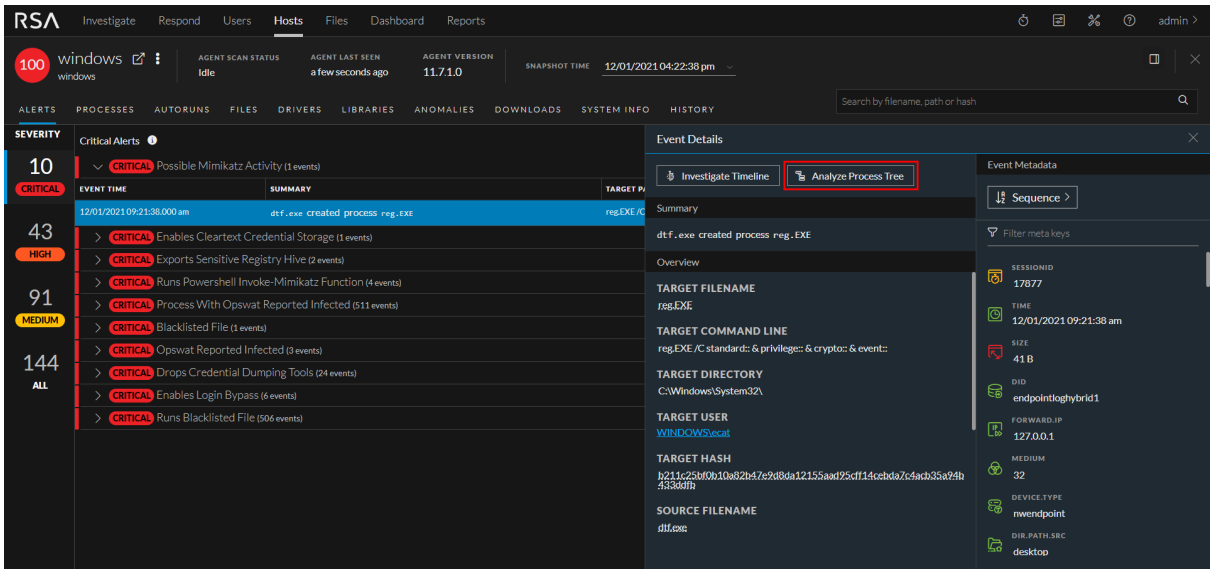
Note: Linux is not supported for analyzing a process.

To analyze the process:

1. Go to **Hosts** and click on a host.
2. Click on an event associated with an alert on the Host Details view. The **Event Details** panel appears.
3. To analyze process activities of a file, do one of the following:
   - In the **Event Details** panel, click **Analyze Process Tree**.
   - Select the **Processes** tab and do one of the following:

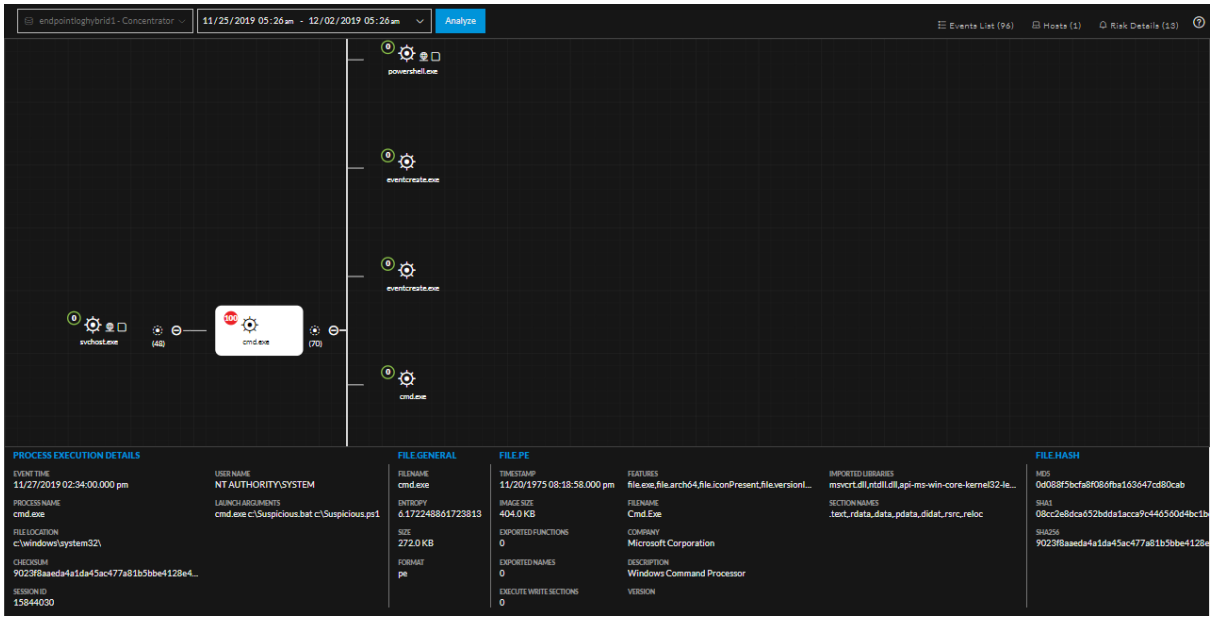
- Right-click a process and select **Analyze Process** from the context menu.
- Click **Analyze Process** in the toolbar.

In the following example, the file `dtf.exe` has created proces `reg.exe`.



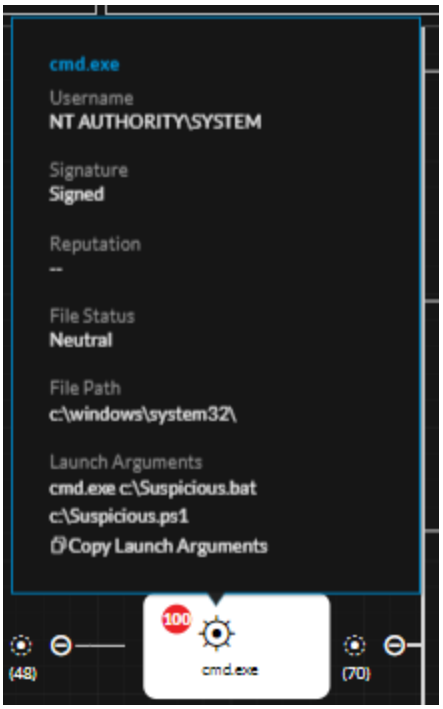
Clicking **Analyze Process** displays the process visualization. For each node, the process name, risk score, and type of activity the selected process has performed (network, file, or registry) are displayed. Optionally, you can change the time range to view data.

You can view the properties, such as process execution details, file properties of the selected process in the bottom of the view.



Note: No result is displayed in the process visualization view if there is no data for last seven days or if there is no `createprocess` event.

4. On the right side of the process visualization view:
   - Click **Events List** to view the associated events. You can also filter events based on the events category. For more information on filtering, see Analyze Events for a Process.
   - Click **Hosts** to view the hosts on which this file is present and the associated risk score. For more information, Analyze Hosts with File Activity.
   - Click **Risk Details** to view the list of distinct alerts, such as Critical, High, Medium and All. For more information, see Analyze Hosts Using the Risk Score.
5. Hover over the process name to analyze important process attributes, such as username, launch arguments, reputation, file status, signer, signature, and file path.

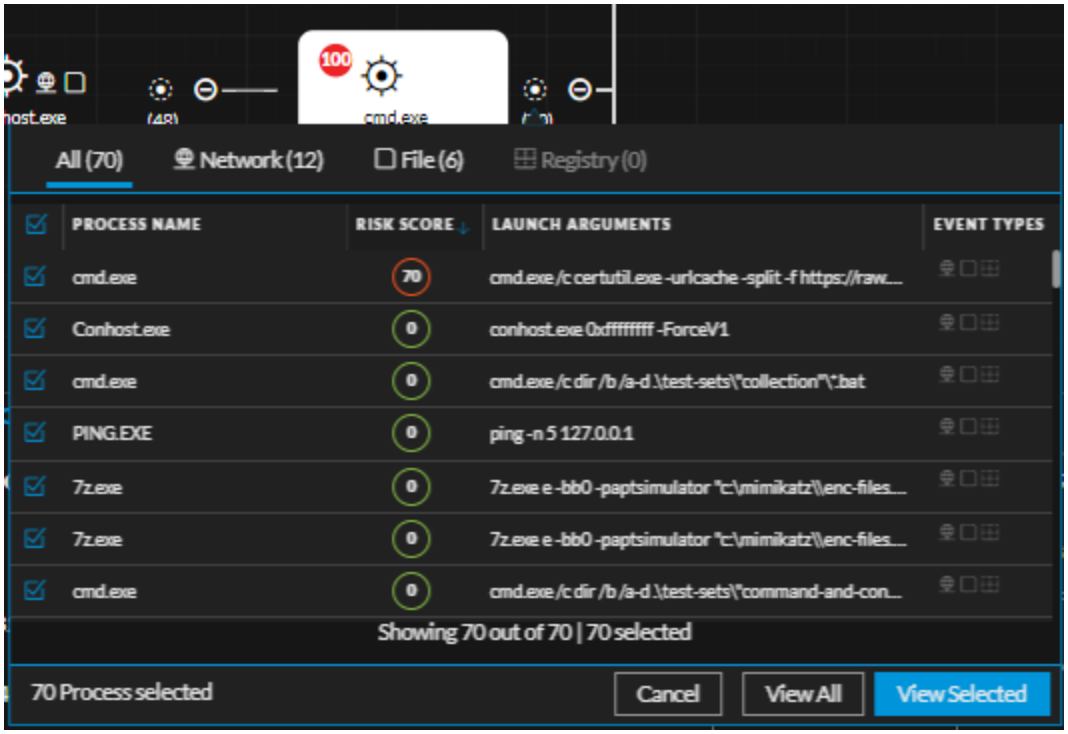


6. Click  to view the child processes. The Process selection dialog is displayed with the child processes associated with the process based on the risk score. You can filter the result on the event type by clicking icons on the top panel. When no matching event types are available, these filter options are disabled.

   Depending on the type of event, the icons are highlights in the Event Types column.

   a. Click **View All** to view all child processes or select the required processes and click **View selected**. The associated events and properties are displayed in the right panel.

b. Click  to change the process selection and click  to collapse the view.

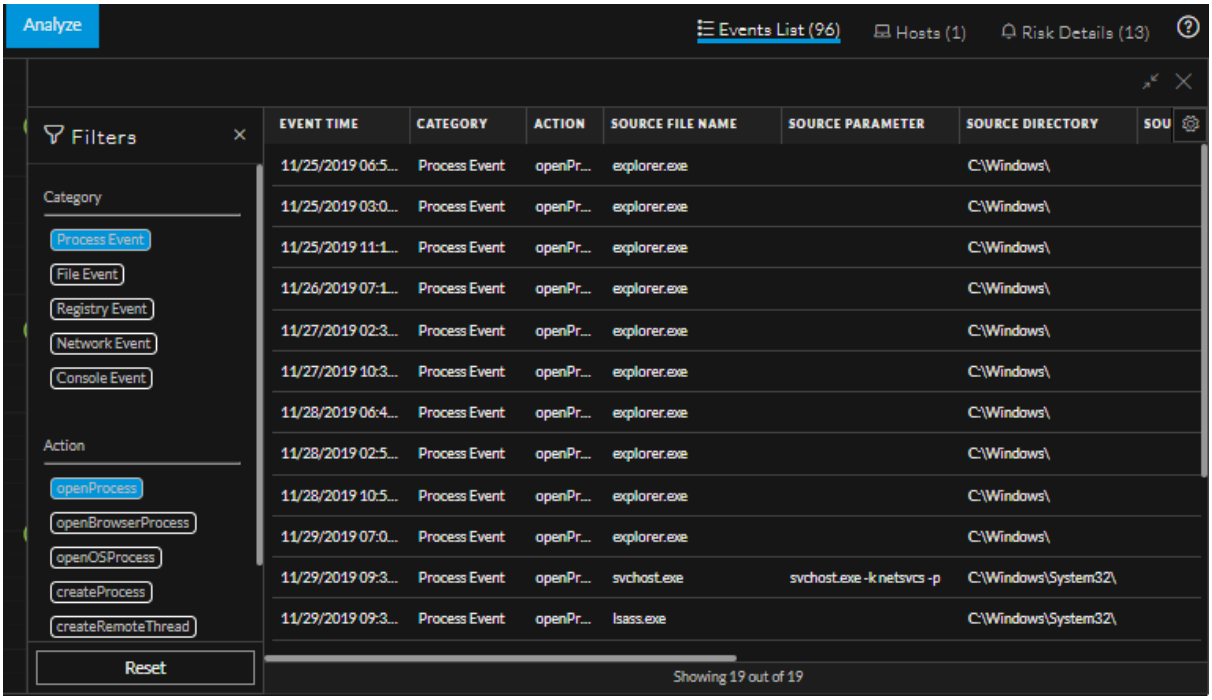


Analyze Events for a Process

To analyze events for the selected process:

1. Perform steps 1 to 3 in Analyze a Process.
2. In the process visualization, click the **Events** tab.
3. To narrow down the search to find any suspicious indicators, behaviors, or specific type of event, filter on a set of matched events based on a category - Process, File, Registry, Network Event, or Console Event (for Windows).

For example, to view only process events, select the **Process Event** category, and filter on action.



The result displays the sequence of activities involving this process for the selected filters.

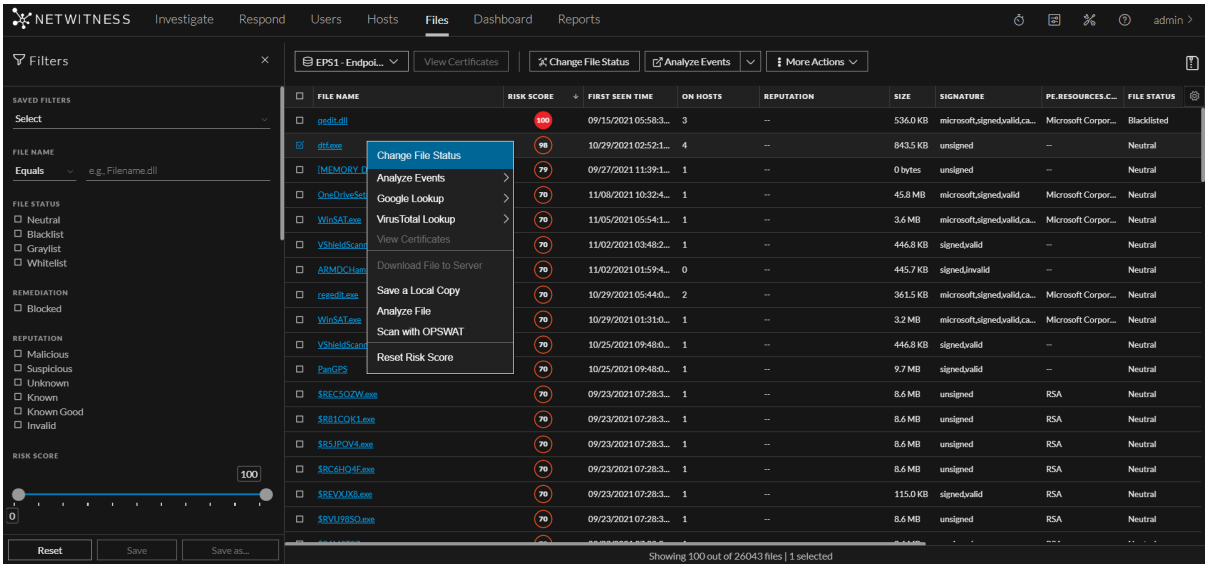
Note: For the console events, the context for local and remote are available only if the data is sent from 11.4 or later agents.

Changing File Status or Remediate

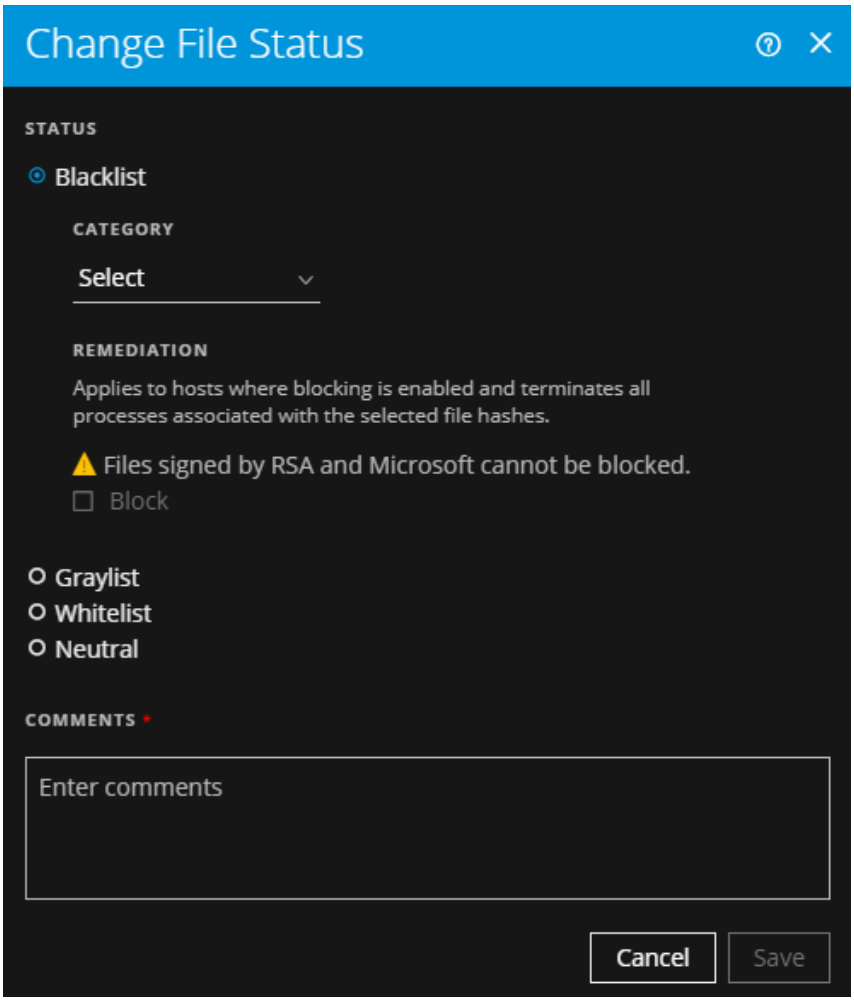
Note: By default, the blocking option is disabled in the policy. To enable blocking, in the policy configuration, change the **Blocking** option to **Enabled** under Response Action Settings. For more information, see the *NetWitness Endpoint Configuration Guide*.

To change the status of a file:

1. Do one of the following:
   - Go to **Hosts** (Processes, Autoruns, Files, Drivers, Libraries, or Anomalies tab).
   - Go to **Files**.
2. Select one or more files and do one of the following:



   - Right-click and select **Change File Status** from the context menu.
   - Click **Change File Status** in the toolbar.
3. In the Change File Status dialog, select a status - **Blacklist**, **Whitelist**, **Graylist**, or **Neutral**.



Note: You cannot whitelist certain Microsoft files, such as `cscript.exe`, `wscript.exe`, `cmd.exe`, `bash.exe`, as there is a potential risk of them being used for malicious purposes. For more information, see Files Restricted from Whitelisting.

If you select Blacklist or Graylist, the following options are displayed:

a. **Category:** Select the appropriate category type: **Generic Malware**, **APT: Advanced Persistent Threats**, **Attacker Tool**, **Unidentified**, **Ransomware**.

Caution: Before blocking, make sure that you review the file because this may cause the system or software to be unusable.

a. **Remediate:** Select **Block** to block the file.

Note:
You cannot block the following:.
- Memory DLL and floating code
- Files that are signed by Microsoft or RSA.

To delete a blocked file, users can log in to the host and execute the delete command using the elevated command prompt.

4. Add a comment and click **Save**.

You can change the status of only 100 files at a time. When the status is changed, it impacts the file status on all hosts on which the file is present. The status is sent as a session under the **File** category, and available for investigation. If the file is seen in subsequent scan or tracking, the corresponding sessions contain a meta value with the file status (except Neutral).

Import File Hashes using the Block Hash tool

The Block Hash tool allows you to import a set of file hashes which can be set to block state and change the file bias status (whitelist, blacklist, and graylist). The tool allows you to block the imported file hashes (suspicious, invalid, and malicious) and prevent them from opening or executing on the hosts. You can block up to a maximum of **50,000** file hashes using this tool.

Note: For more information on changing the file status or blocking the file, see Changing File Status or Remediate.

IMPORTANT:
- Enter only valid SHA256 hashes; otherwise, the blocking functionality might break.
- Do not block any file hashes signed by RSA, Microsoft, and Apple. It might make your Endpoints unresponsive.
- Make sure the number of hashes entered in the JSON file is less than the number of the available limit of hashes that can be blocked (the maximum limit is **50,000**).
For example: If 100 file hashes are already blocked as part of NetWitness deployment, you can only block 49,900 more file hashes using this tool.

JSON File Format

The example below describes the JSON file format for blocking and blacklisting the file hashes.

Sample demoblock.JSON:

```
[{
"checksums":
["1b30e463ebe0131db66fce7d4aa43f3e149064d85c4c0dc5218b077886da2804","67fa30e4
63ebe0131db66fce7d4aa43f3e149064d85c4c0dc5218b077dsbhb561","78vbba909e463ebe0
131dsdsdb66fce7d4aa43f3e1dsdsd49064d85dsdsman61n"],
"fileStatus": "Blacklist",
"comment”: "File blocking set through new tool",
"remediationAction": "Block"
}]
```

The example below describes the JSON file format only for blacklisting the file hashes.

Sample demoblock.JSON:

```
[{
"checksums":
["2b30e463ebe0131db66fce7d4aa43f3e149064d85c4c0dc5218b077886da2800","97fa30e4
63ebe0131db66fce7d4aa43f3e149064d85c4c0dc5218b077dsbhb500","38vbba909e463ebe0
131dsdsdb66fce7d4aa43f3e1dsdsd49064d85dsdsman68c"],
"fileStatus": "Blacklist",
"comment": "File status change set through new tool",
```

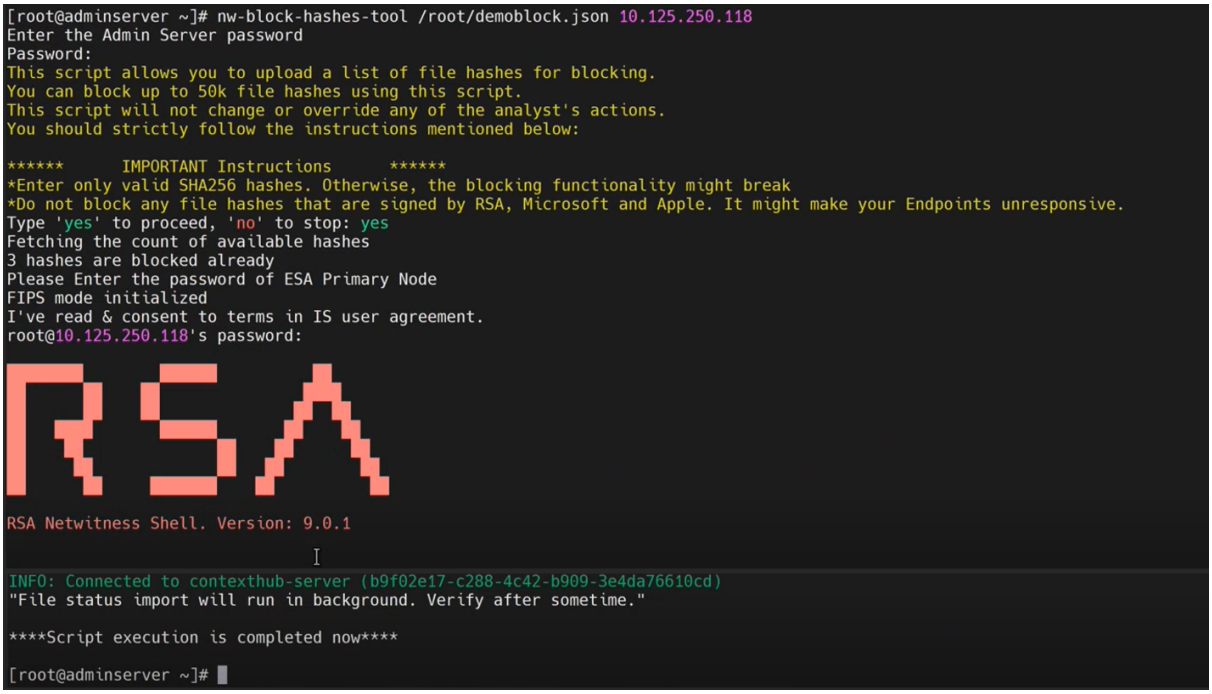
```
}]
```

To block the file hashes using the Block Hash tool:

1. SSH to node 0 and copy the JSON file (containing the file hashes to be blocked) stored in it.
2. Run the tool.

   ```
   nw-block-hashes-tool <absolute path of json residing on node 0> <ESA node IP/hostname>
   ```

   Note: Enter the JSON file path in `<absolute path of json residing on node 0>` and enter the ESA node IP in `<ESA node IP/hostname>`.
   For Example: `nw-block-hashes-tool /root/demoblock.json 10.125.250.118`.



3. Enter the Admin Server password.
4. Follow the instructions displayed on the screen. Enter any one of the following options when prompted.
   - **yes**: Enter **yes** to proceed with the execution of the script.
   - **no**: Enter **no** to stop the execution of the script.

   Note: At this stage, the tool validates the number of incoming file hashes through the JSON file and the number of existing file hashes already blocked in the deployment. If the total of the file hashes (incoming through JSON file and existing as blocked in the deployment) exceeds the maximum limit of **50,000**, the tool stops the execution, and it will not proceed further.

Files Restricted from Whitelisting

To view or update the files that are restricted from whitelisting, do the following:

1. On the NW server, run the `nw-shell` command from the command line.
2. Run the `login` command and enter your credentials.
3. Connect to the Endpoint Server using the following command:

   ```
   connect endpoint-server
   ```
4. Run the following commands to view the list of files:
   - `cd endpoint/file/status/restricted/get`
   - `invoke Whitelist`
5. Run the following commands to add files to the list:
   - `cd endpoint/file/status/restricted/get`
   - `invoke '{"id":"<filename>","restrictedStatus":["Whitelist"], "enable":true}`
6. Run the following commands to delete files from the list:
   - `cd endpoint/file/status/restricted/update`
   - `invoke '{"id":"<filename>","restrictedStatus":["Whitelist"], "enable":false}`

Analyzing Downloaded Files

To perform a deep analysis of suspicious files, you can manually or automatically download the file to the server.

Note: Saving or analyzing downloaded file works the same way irrespective of whether the file is downloaded manually or automatically.

Note: Downloaded files are stored in the Endpoint Server which may fill up the disk space. To utilize the storage efficiently without impacting the health of Endpoint Server, NetWitness recommends you to configure an external storage mount, so all the Endpoint Server can use the configured location to store the downloaded data.
By default, all files are downloaded to `/var/netwitness/endpoint-server/<files>/`. If you want to change the location, make sure that you have **endpoint-server.configuration.manage** permissions and do the following:
1. In the Explore view, go to **endpoint/download**,
2. In the base-path, provide the location of the directory.

Caution: By default, the status **File Download Disk Usage** stats in the **Health and Wellness** view shows unhealthy if the disk usage reaches 60% and the file download stops automatically when the disk usage is 70%. You can customize the warning or fatal thresholds in the  **> Services > Endpoint Server > view > Explore > rsa.endpoint.file-download-disk-thresholds.warning-percent** and **rsa.endpoint.file-download-disk-thresholds.fatal-percent** parameters respectively.

For the downloaded file, you can:

- Search for strings in the executable
- View text content for scripts
- View imported libraries and functions
- Save a local copy for further analysis

Download Files to Server

Downloading file to server is not supported for memory DLL and floating code.

Note: Downloading files may take significant time. Additional requests to the agent during download are queued and processed when the download is complete.

Automatic File Download

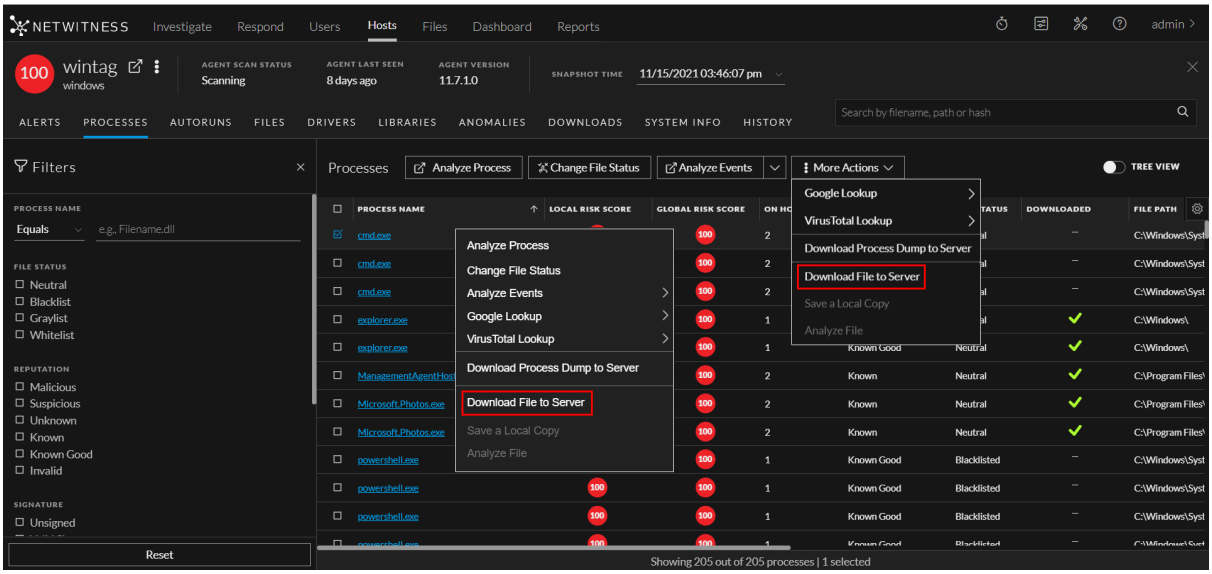
By default, the files that are unsigned and size lesser than or equal to 1 MB are downloaded automatically to the NetWitness Endpoint server. And, only single copy of each file is downloaded automatically. You can limit the volume of files to be downloaded in the  **> Endpoint Sources > Policies** tab, so the files matching certain criteria are only downloaded automatically. For more information on automatic file download settings, see "Create an EDR Policy" section in the *NetWitness Endpoint Configuration Guide*.

The status of the download is displayed in the **Files** tab > **Downloaded** column.

Manual File Download

To manually download files to the server from the Hosts view:

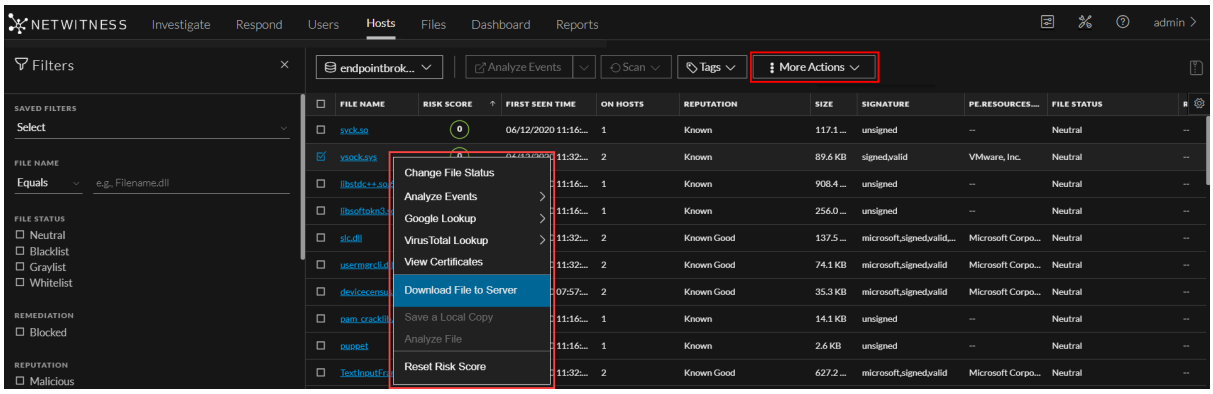
1. Go to **Hosts**.
2. Select the hostname to open the Host Details view.
3. In any of the Processes, Autoruns, Files, Drivers, Libraries, or Anomalies tabs, select the file, and do one of the following:



- Right-click and select **Download File to Server** from the context menu.
- Select **Download File to Server** from the **More Actions** drop-down list in the toolbar.

To download files to the server from the Files view:

1. Go to **Files**.
2. Select the file and do one of the following:



- Right-click and select **Download File to Server** from the context menu.
- Select **Download File to Server** from the **More Actions** drop-down list in the toolbar.

The status of the download is displayed in the Downloaded column. The download statuses are Downloaded, Not downloaded, and Error.

Save Downloaded Files

You can retrieve a downloaded file and save it to your local file system for further analysis. Downloaded files are stored in the server in the configured location. This option is enabled only if the file is downloaded to the server.

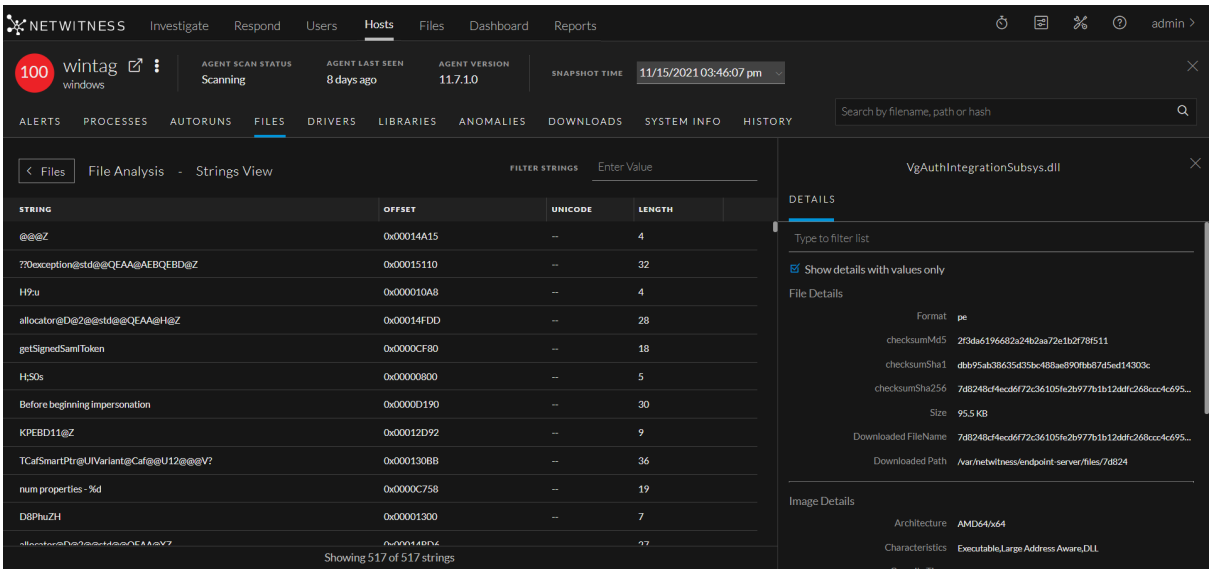
To save a file:

1. Go to **Hosts** or **Files** .
2. Right-click the file you want to save and select **Save a Local Copy**.
3. Browse the location and click **Save**.

Analyze Downloaded Files

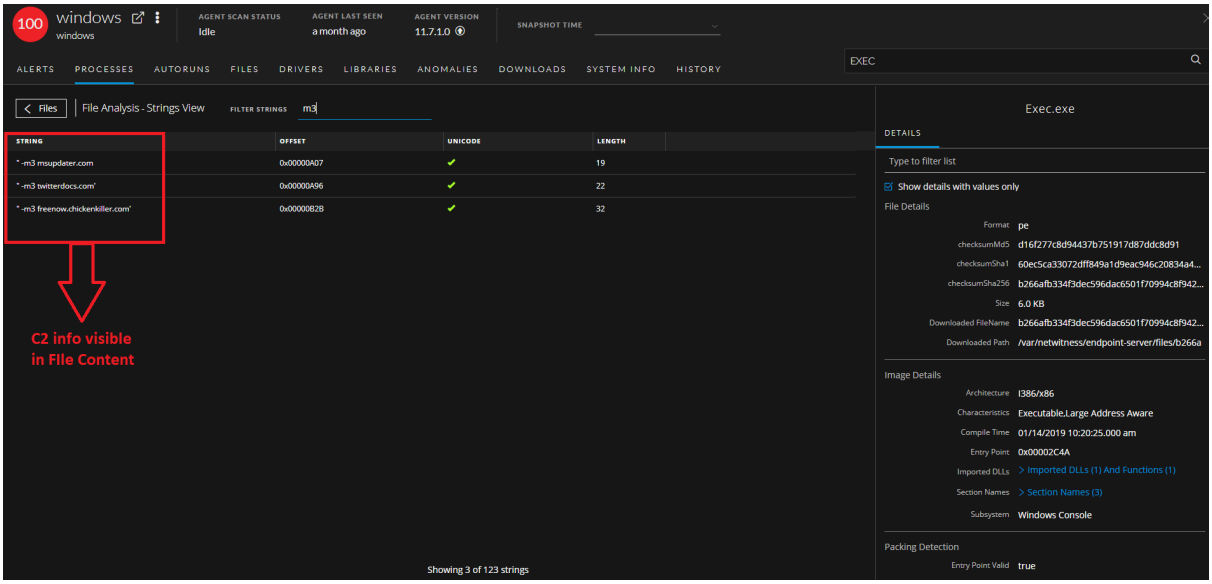
You can use the **Analyze File** option to view detailed information about a downloaded file. This option is enabled only if the file is downloaded to the server. To analyze a file:

1. Go to **Hosts** or **Files**.
2. Right-click the downloaded file and select **Analyze File**. The File Analysis view opens and properties of the are is displayed in the right panel.

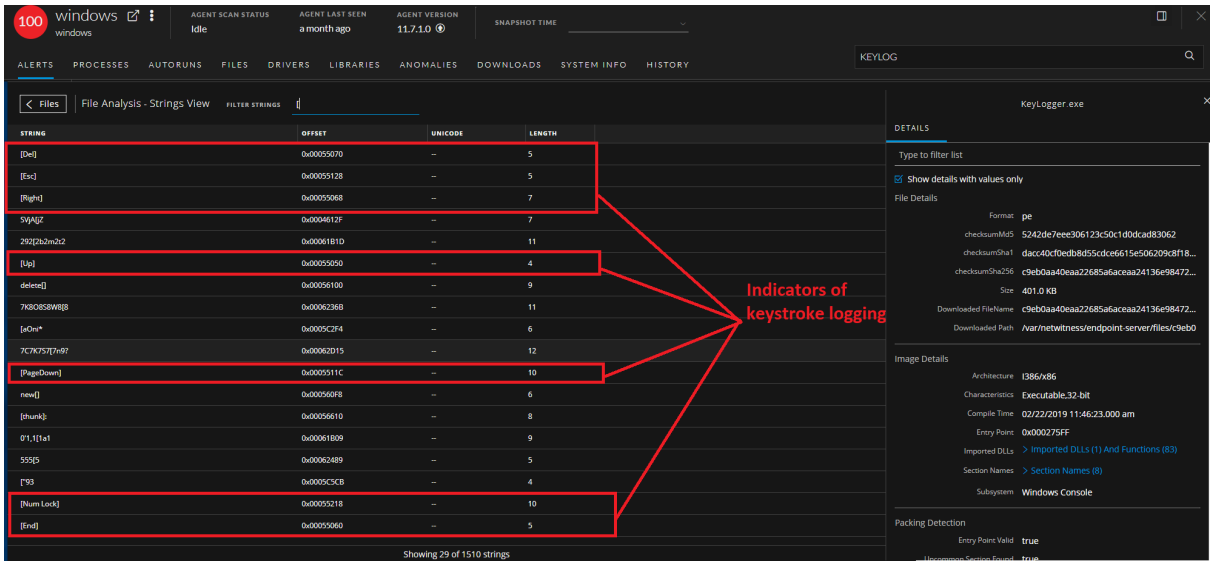


3. View strings in the file in the Strings view while analyzing an executable (such as `macho`, `pe`, `elf`). This view contains the string, offset in the binary, unicode, and the length of the string. You can search for or filter on a specific string value in the **Filter String** field.
4. View the text content of the file and look for any suspicious behavior in the script file.

   For example, if the file contains C2 information in the form of domain names or IP addresses, it is highly suspicious.



If you see unprintable keyboard keys listed within the file, such as: [F1], [F2]…[Page Up], [Enter], [ESC], and so on, that may be indicative of a keystroke logger.


100
windows
windows
AGENT SCAN STATUS
Idle
AGENT LAST SEEN
a month ago
AGENT VERSION
11.7.1.0
SNAPSHOT TIME
ALERTS
PROCESSES
AUTORUNS
FILES
DRIVERS
LIBRARIES
ANOMALIES
DOWNLOADS
SYSTEM INFO
HISTORY
KEYLOG
Files
File Analysis - Strings View
FILTER STRINGS
STRING
OFFSET
UNICODE
LENGTH
[Del] 0x00055070 -- 5
[Esc] 0x00055128 -- 5
[Right] 0x00055068 -- 7
5VjAJjZ 0x0004612F -- 7
292[2b2m2r2 0x00061B1D -- 11
[Up] 0x00055050 -- 4
delete[] 0x00056100 -- 9
7K8O8S8W8J8 0x0006236B -- 11
[aOn* 0x0005C2F4 -- 6
7C7K7S7J7n9? 0x00062D15 -- 12
[PageDown] 0x000551 1C -- 10
new[] 0x000560F8 -- 6
[thunk]: 0x00056610 -- 8
01,1[1a1 0x00061B09 -- 9
555[5 0x00062489 -- 5
[*93 0x0005C5CB -- 4
[Num Lock] 0x00055218 -- 10
[End] 0x00055060 -- 5
Showing 29 of 1510 strings
Indicators of keystroke logging
KeyLogger.exe
DETAILS
Type to filter list
Show details with values only
File Details
Format pe
checksumMd5 5242de7eee306123c50c1d0dcad83062
checksumSha1 dacc40cf0edb8d55cdce6615e506209c8f18...
checksumSha256 c9eb0aa40eaa22685a6aceaa24136e98472...
Size 401.0 KB
Downloaded FileName c9eb0aa40eaa22685a6aceaa24136e98472...
Downloaded Path /var/netwitness/endpoint-server/files/c9eb0
Image Details
Architecture I386/x86
Characteristics Executable,32-bit
Compile Time 02/22/2019 11:46:23.000 am
Entry Point 0x000275FF
Imported DLLs > Imported DLLs (1) And Functions (83)
Section Names > Section Names (8)
Subsystem Windows Console
Packing Detection
Entry Point Valid true

Performing Host Forensics

Note: The information in this topic applies to NetWitness Version 11.4 and later.

You can perform the following forensic investigation on a host:

- Master File Table (MFT)
- System Dump
- Process Dump

Note: This is applicable only for Windows agent (in Advanced mode) with NetWitness Platform version 11.4. Downloading system dump files may take significant time. Additional requests to the agent during system dump download are queued and processed when the download is complete. MFT, system dump, and process dump downloads are not supported for agents communicating through Relay server.

Note: MFT, system dump, and process dump are stored in the Endpoint Server which may fill up the disk space. For large deployments, to utilize the storage efficiently without impacting the health of Endpoint Server, NetWitness recommends you to configure an external storage mount, so all the Endpoint Server can use the configured location to store the downloaded data.
By default, all files are downloaded to `/var/netwitness/endpoint-server/<file type>/`, where `<file type>` is MFT, system dump, or process dump. If you want to change the location, make sure that you have endpoint-server.configuration.manage permissions and do the following:
1. In the Explore view, go to **endpoint/download**.
2. In the base-path, provide the location of the directory.

Download Master File Table

Master File Table contains metadata of every file on the host. It keeps track of information, such as filename, size, timestamps, permissions, and location of the file on the host. It consists of two sets of timestamps - Standard Information ($SI) and File Name ($FN). Each set has the following timestamps - creation, access, update, and modification.

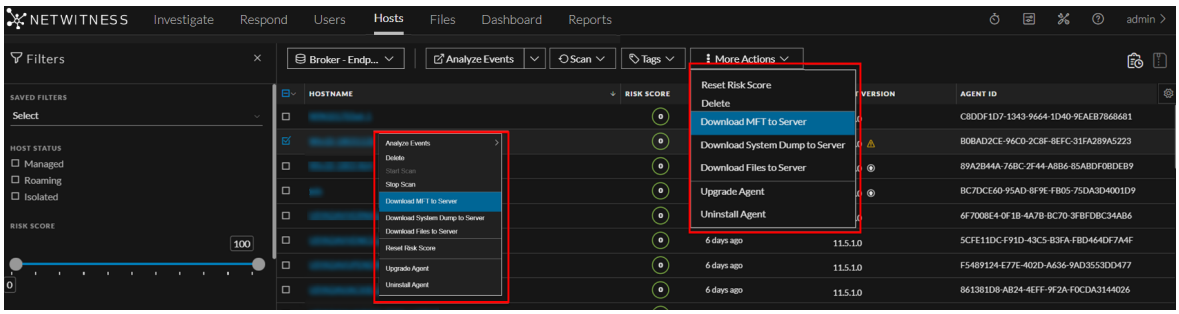
Time stomping is a technique that modifies the timestamps for a file (creation, access, update, and modification time) to mimic files that are in the same folder, making it difficult to identify suspicious files on a host. To perform forensic investigation of a suspicious file, you can download and analyze the MFT, and focus on files that are time stomped. For more information, see Analyze Downloaded MFT.

During MFT analysis, you can also search for suspicious filenames, and also files that were created before or after a known malicious event. You can also download files from the MFT viewer for further analysis.

Download MFT to Server

To download MFT to the server from the Hosts view:

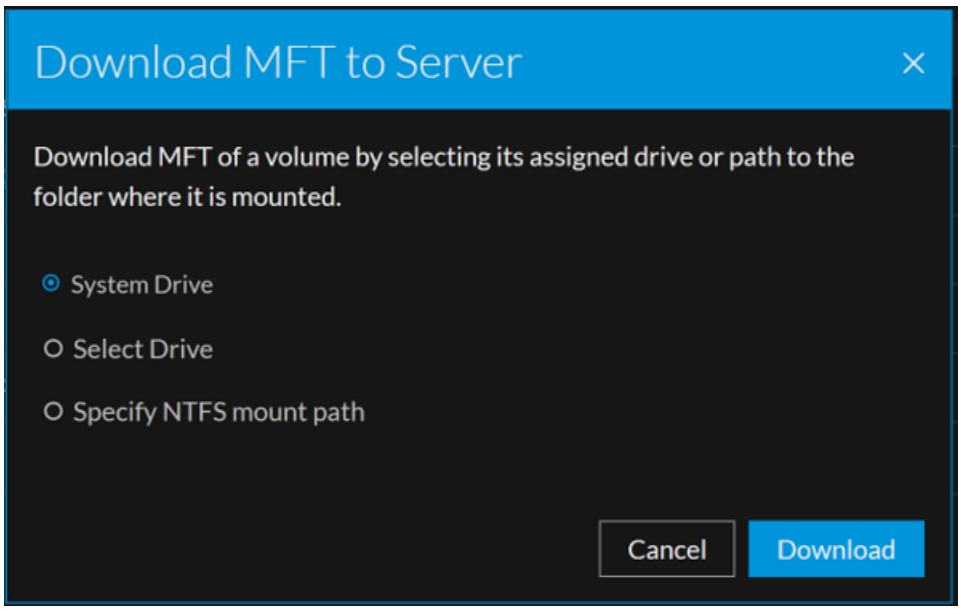
1. Go to **Hosts** and do one the following:
   - Select one or more hosts and select **Download MFT to Server** from the right-click context menu, or from the **More Actions** drop-down list in the toolbar.



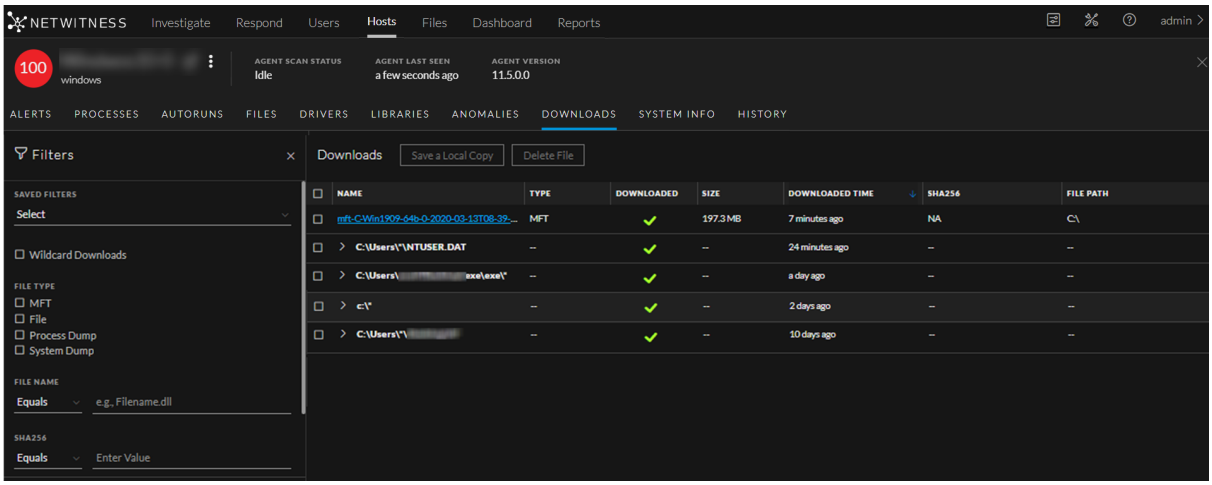
   - Select the hostname to open the host details, click ⋮ (**More**) beside the hostname, and select **Download MFT to Server**.

2. In the **Download MFT to Server** dialog, select one of the following:
   - **System Drive** - to download MFT to the system drive.
   - **Select Drive** - to download MFT on assigned drive.
     You can select any drive from the **Drive** drop-down list. By default the selected drive is C.
   - **Specify NTFS mount path** - to download MFT on the path to the folder where it is mounted

   Click **Download**.



3. View details of the downloaded MFT in the **Downloads** tab within the host details. For more information, see Hosts View - Downloads Tab.



Analyze Downloaded MFT

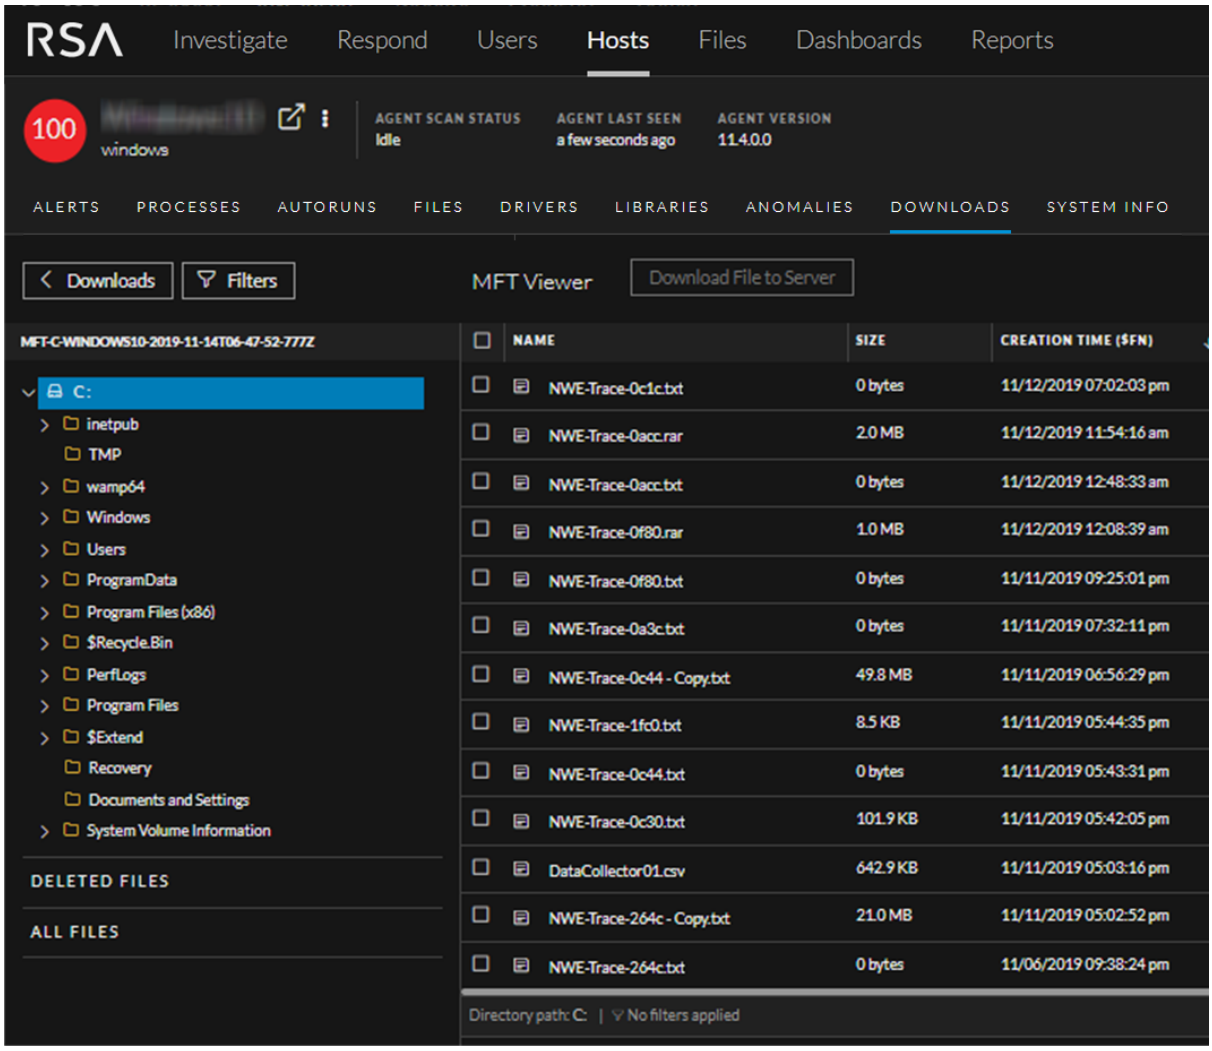
You can use the MFT viewer to begin analysis where you can search for files based on file name, time stamps, and identify files that are timestomped.

View MFT

To view the content of the downloaded MFT:

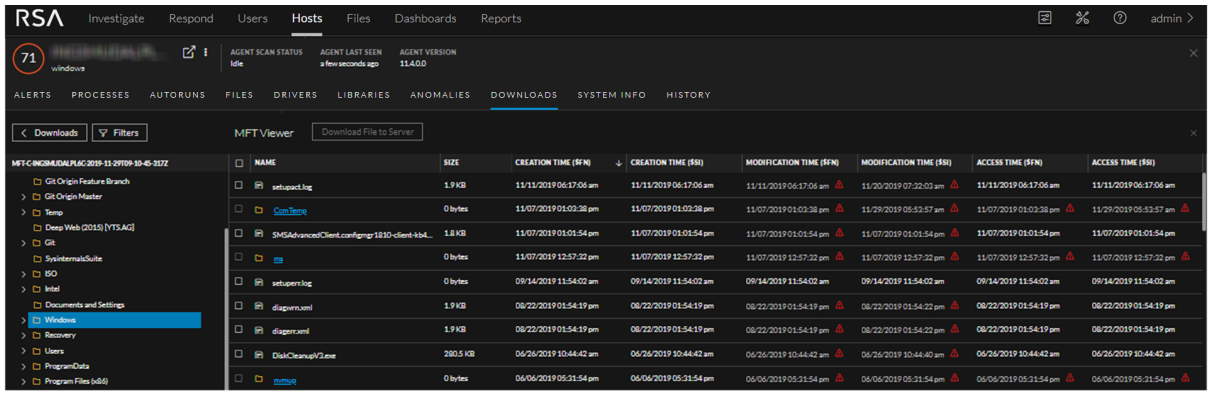
1. Go to **Hosts**.
2. Select the hostname to open the host details and select the **Downloads** tab.
3. Click the file name. The MFT viewer is displayed.

   All available files are displayed in a tree view similar to the Windows Explorer in the All Files folder. The Deleted Files folder contains a sequential list of all deleted files.



4. Click ▶ to view the folder structure. Click the row to view the folder content.

   The details of the MFT is displayed in the table. By default, the table is sorted on the creation time ($FN). If the $SI and $FN timestamps are different, the columns are highlighted in red (⚠) indicating that it is time stomped.

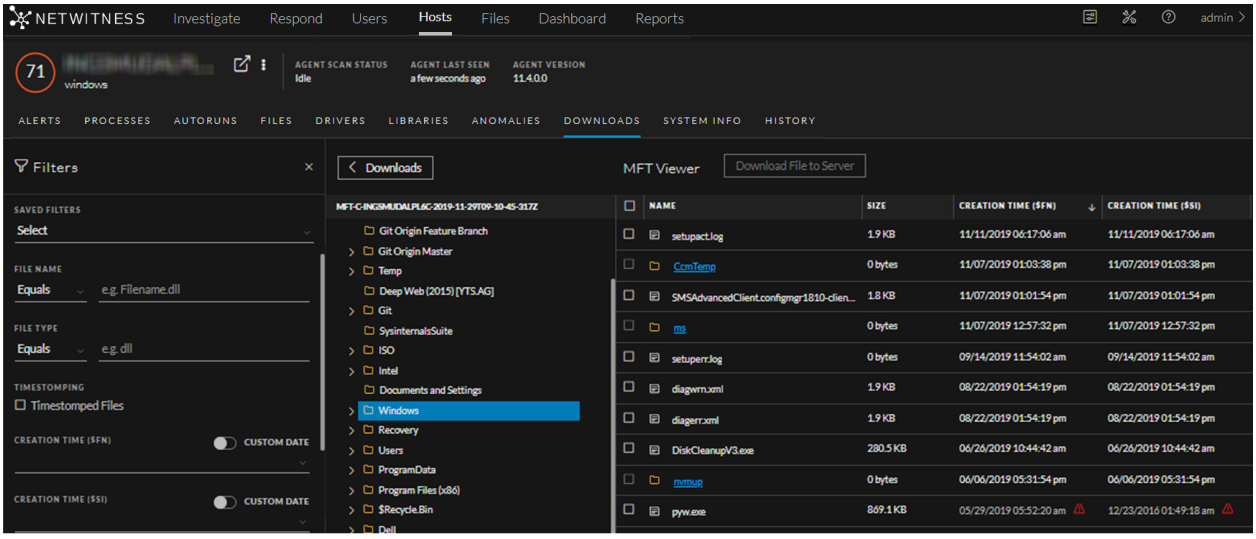


5. Select one or more files and click **Download File to Server** on the toolbar to download files to the server.

> **Note:** Downloading a folder is not supported and hence the option is grayed out for folders.

Filter MFT

You can filter files on file name, creation time ($FN), creation time ($SI), access time ($FN), access time ($SI), update time ($FN), update time ($SI), modified time ($FN), and modified time ($SI).



Click **Save** to save the filter and provide a name (up to 250 alphanumeric characters). The filter is added to the Saved Filters panel on the left. To delete a filter, hover over the filter name and click .

> **Note:** Special characters are not allowed except underscore (_) and hyphen (-) while saving the filter.

To filter, save, and delete MFT, see Filter Downloaded Files, Save Downloaded File, and Delete Downloaded Files.

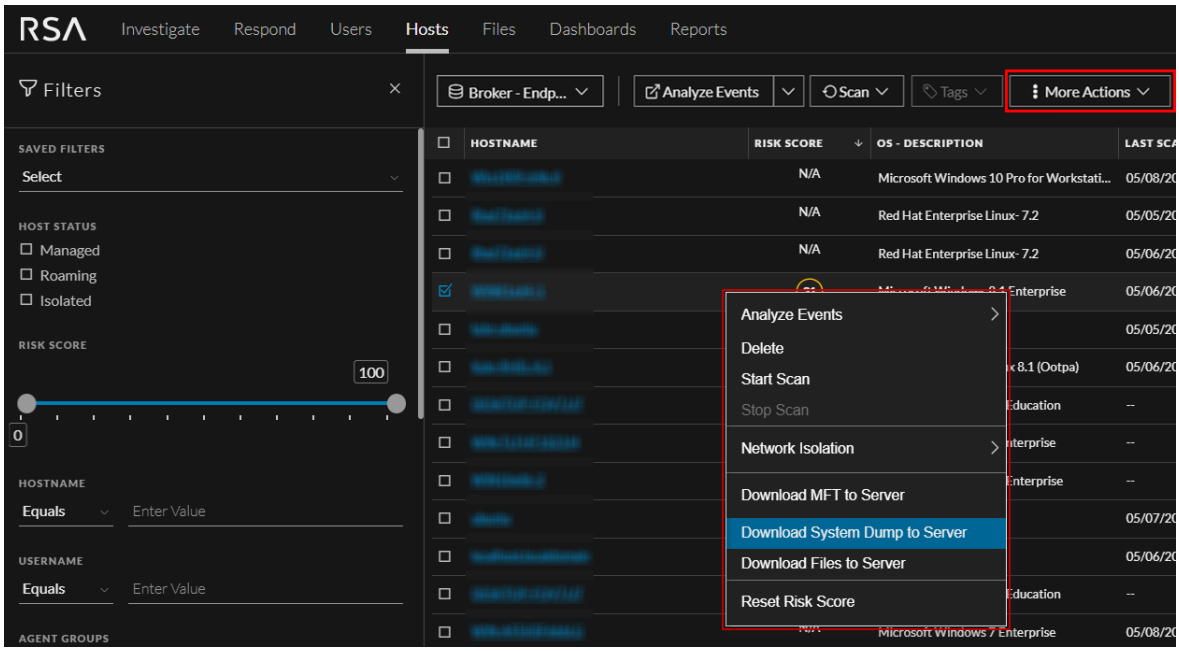
System and Process Memory Dump

To perform forensic investigation during an incident response, you can request a memory dump of a host or a process running on the host. You can analyze these dumps using third-party tools, such as Volatility, Rekall.

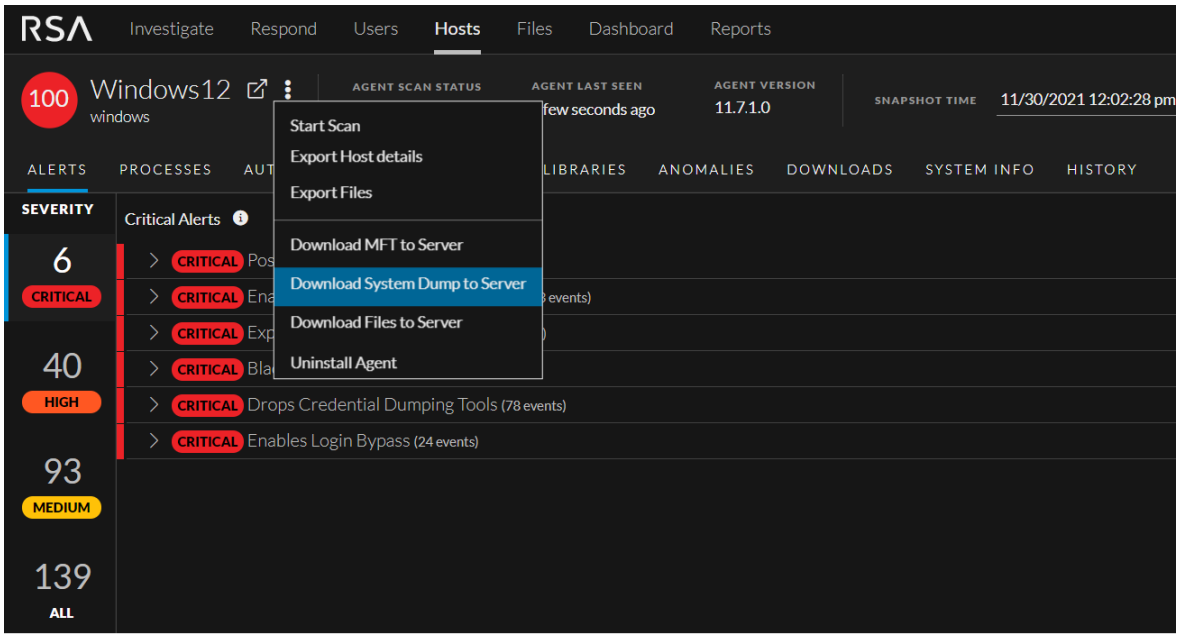
Download System Dump to Server

To download system dump to the server from the Hosts view:

1. Go to **Hosts** and do one the following:
   - Select a host and select **Download System Dump to Server** from the right-click context menu, or from the **More** drop-down list in the toolbar.



- Select the hostname to open the host details and select **Download System Dump to Server** from the **More** option besides the hostname.

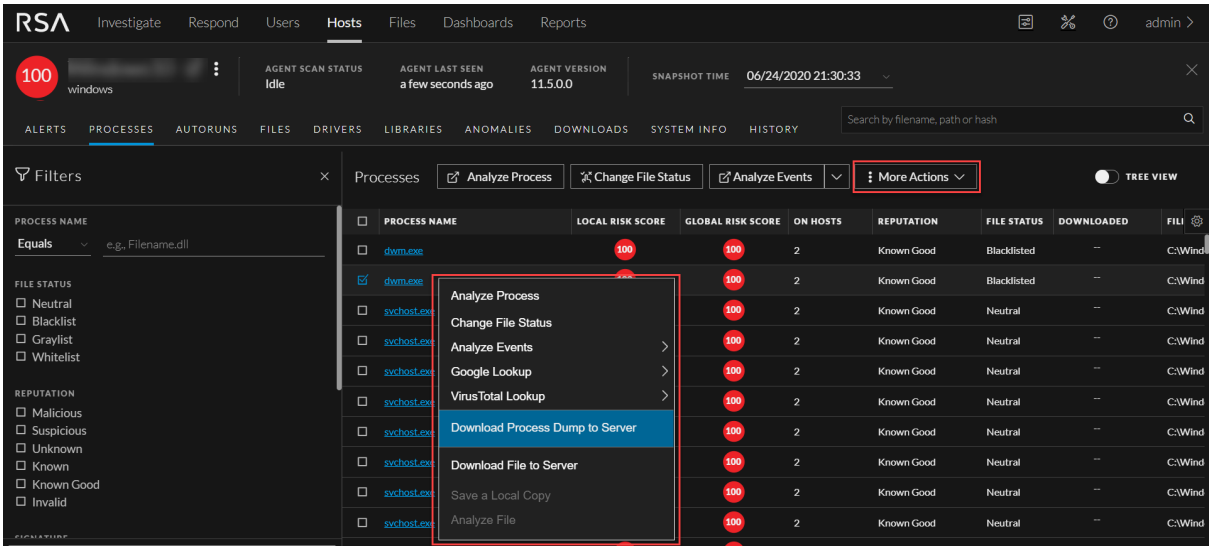


2. View the details of the downloaded system dump in the **Downloads** tab within the host details. For more information, see Hosts View - Downloads Tab.

Download Process Dump to Server

To download process dump to the server:

1. Go to **Hosts**.
2. Select the hostname to open the host details.
3. In the Processes, Libraries, or Anomalies tab, select **Download Process Dump to Server** from the right-click context menu, or from the **More Actions** drop-down list in the toolbar.



4. View the details of the download process dump in the **Downloads** tab within the host details. For more information, see Hosts View - Downloads Tab.

To filter, save, and delete system dump or process dump, see Filter Downloaded Files, Save Downloaded File, and Delete Downloaded Files.

The following are some errors you might encounter during system and process dump download:

Issue	Explanation
Parameter is incorrect.	The process for which the dump is requested might be running with a different process ID.
Element not found	The process for which the dump is requested is no longer active.
java.io.IOException:Unable to unwrap data, invalid status [CLOSED]	Connection to the agent is interrupted.
java.net.SocketTimeoutException	The network is slow or the system is down.
One or more arguments are not correct	Agent might be in the Insight mode or driver is not running.

Download Files Using Full Path or Wildcard

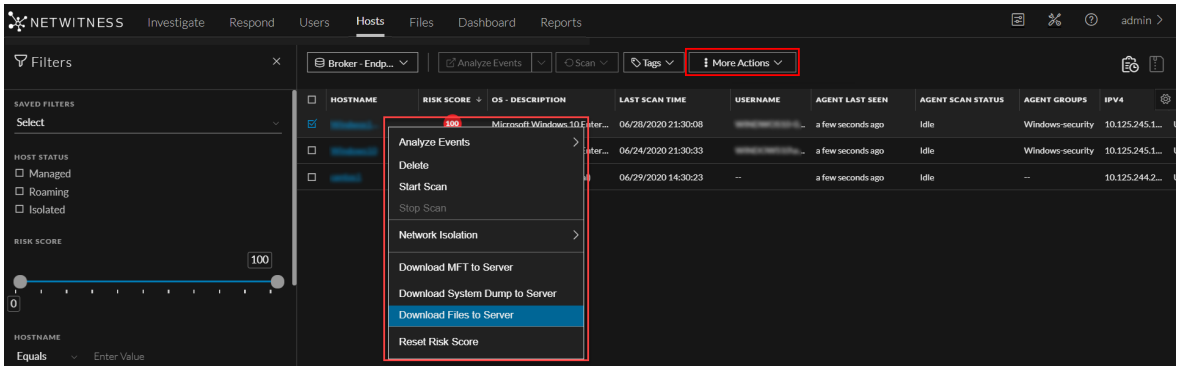
You can manually download files that help in investigations by either providing full path of the file or using wildcard.

Note: This is applicable only for agents in Advanced mode with NetWitness Platform version 11.5 and later.

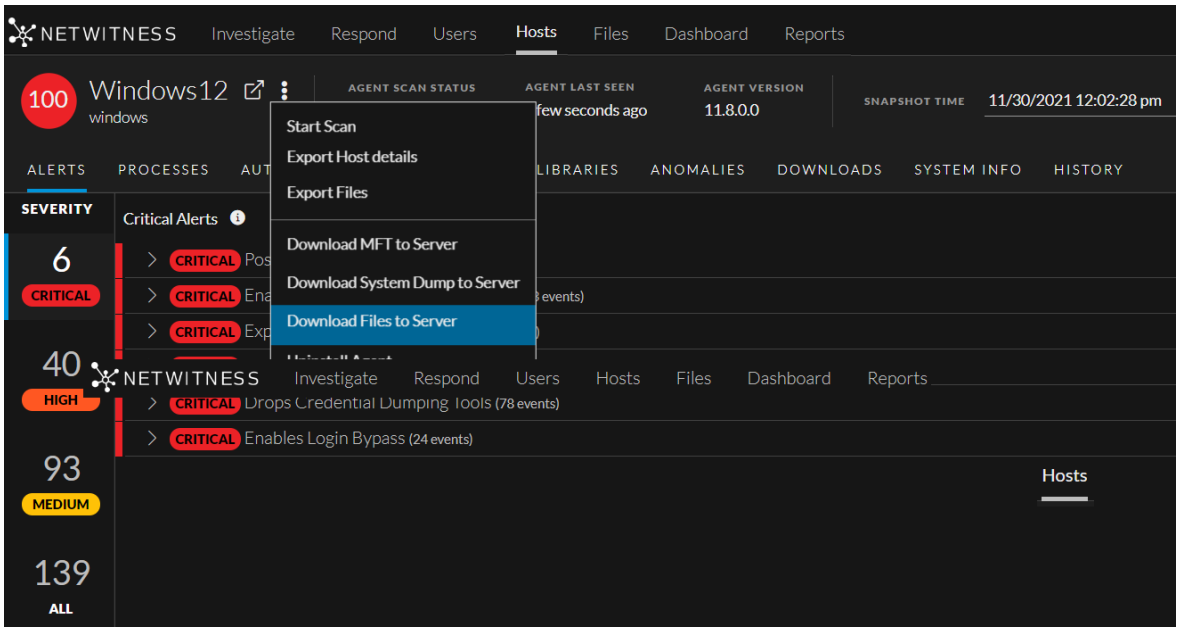
To download files to the server:

1. Go to **Hosts** and do one of the following:

   - Select one or more hosts from the same operating system, and select **Download Files to Server** from the right-click context menu, or from the **More Actions** drop-down list in the toolbar. You can download files from only top 100 selected hosts at a time.



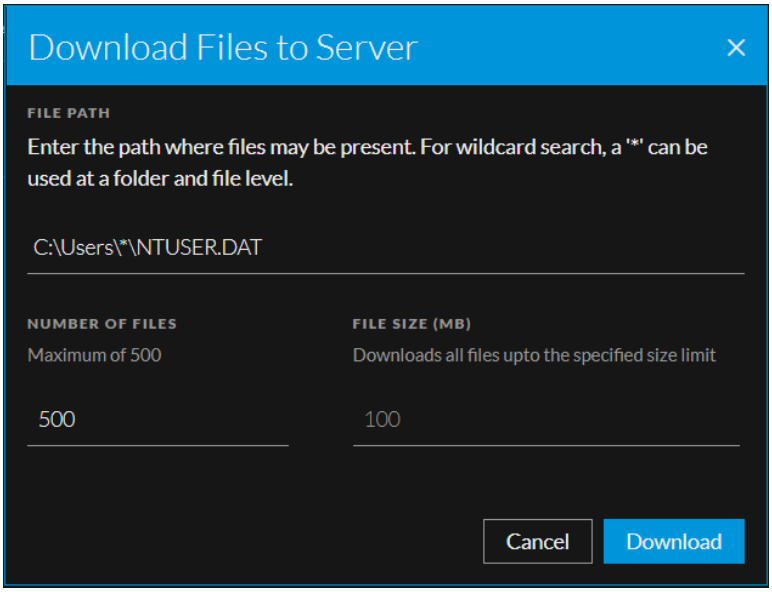
   - Select the hostname to open the host details, click ⋮ (**More**) beside the hostname, and select **Download Files to Server**.



2. In the **Download Files to Server** dialog, enter the full path where the files may be present or search using wildcard. For wild card search, you can use a maximum of two `*`, one at a folder level and the other at a file level.

   For example, to retrieve the registry hive, you can enter the full path, `C:\Windows\System32\config\SYSTEM`.

If you want to retrieve user settings and configuration preferences for all users, download all files using the wildcard `C:\Users\*\NTUSER.DAT`.

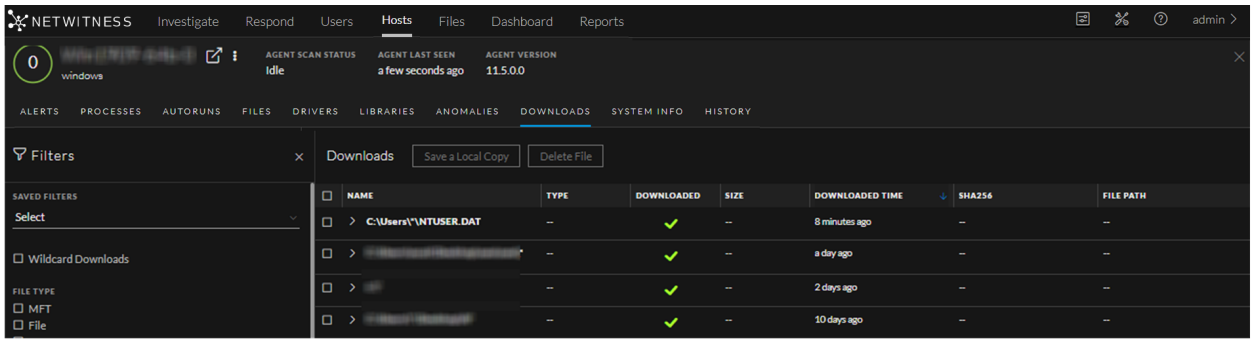


3. For wildcard search, enter the number of files to download and size of the file. By default, the number of files is set to 10 and file size is set to 100 MB. For example, if the maximum number of files is set to 10 and file size is set to 10 MB, first 10 files within 10 MB are downloaded.

> **Note:** You can set a limit to the Maximum Number of Files field on the explore page of the Endpoint server (**> [Endpoint server] > explore > endpoint/command > max-file-count**). By default, the limit is set to 100, and you can change it to any value between 100 - 1000 in each Endpoint server. In broker view, if the Endpoint servers have different max-file-count, the lesser value will be taken as the limit.

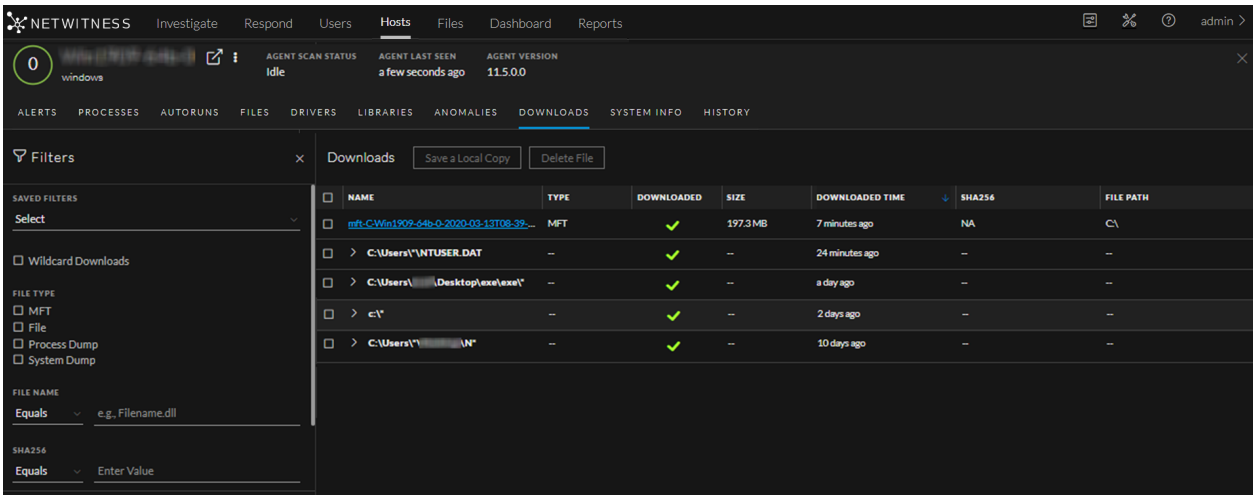
4. Click **Download**.

All files downloaded as a part of wildcard search are grouped together based on the search criteria. For example, all files downloaded using `C:\Users\*\NTUSER.DAT` are grouped, and you can click to expand and view all files under this group. You can sort the groups on the downloaded time and view the status of the download in the **Downloaded** column.



Filter Downloaded Files

You can filter the downloaded files on wildcard downloads, file type, file name, SHA256 (for files), and downloaded time. In the Downloaded Time field, you can also filter by custom date.



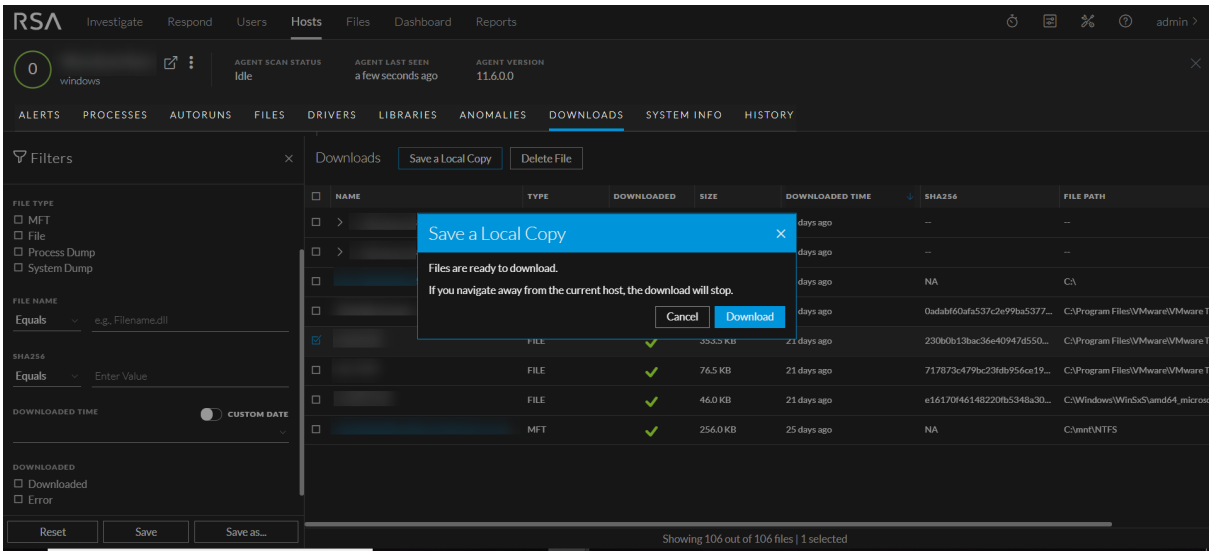
Click **Save** to save the filter and provide a name (up to 250 alphanumeric characters). The filter is added to the Saved Filters panel on the left. To delete a filter, hover over the filter name and click .

Note: Special characters are not allowed except underscore (_) and hyphen (-) while saving the filter.

Save Downloaded File

You can retrieve the downloaded file and save it to your local file system for further analysis. To save the file:

1. Go to **Hosts**.
2. Select the hostname to open the host details and select the **Downloads** tab.
3. Right-click the file you want to save and select **Save a Local Copy** from the context menu or from the toolbar.



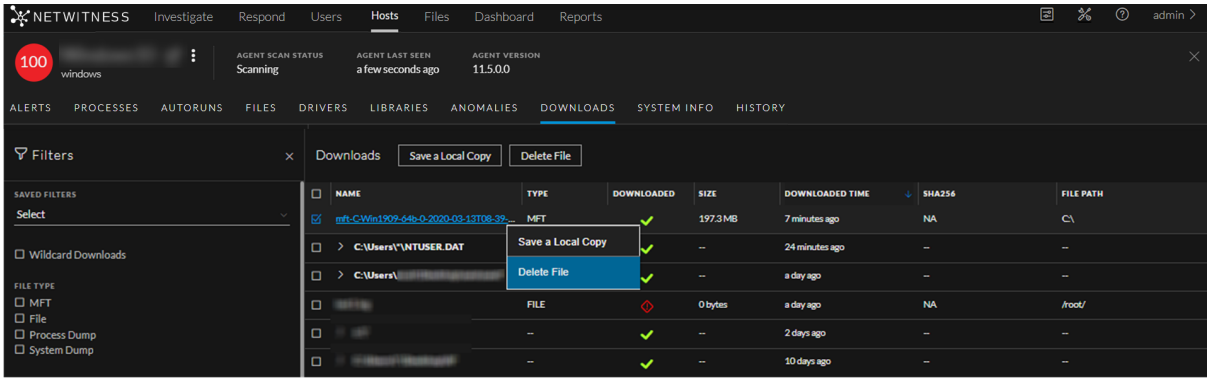
4. In the **Save a Local Copy** dialog, click **Download**.

Note: For wildcard downloads, select a file from the group that are downloaded successfully to save a local copy. You cannot save multiple files in the group at a time or save files with errors.

Delete Downloaded Files

If you want to delete the downloaded file from the server:

1. Go to **Hosts**.
2. Select the hostname to open the host details and select the **Downloads** tab.
3. Right-click one or more files you want to delete, and select **Delete File** from the context menu or from the toolbar.



Note: For wildcard downloads, you can select the group to delete all files that are downloaded.

Analyzing Events

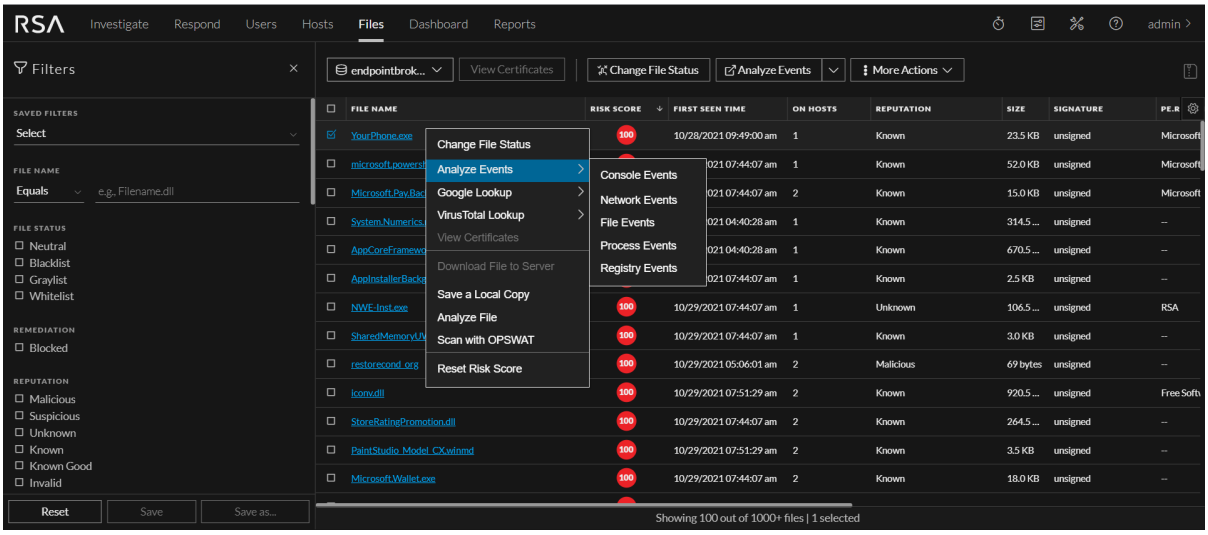
If you need to investigate a particular host, IP address, username, filename, or hash to look for related activity across a time range, you can pivot to Navigate view to get the entire context of the activity. By default, the time range is set to 7 days. You can change the time range.

Note: By default, the system detects the best data source to pivot to Navigate view. To change the data source, modify the investigate service ID under endpoint or investigate in the Explore view.

Analyze Events from Files View

To investigate a particular filename or hash (SHA256 and MD5):

1. Go to **Files**.
2. Select the file you want to analyze and do one of the following:
   - Right-click and select **Analyze Events** from the context menu.
   - Click **Analyze Events** in the toolbar.



This opens the Navigate view with data related to the file. For more information on analyzing events in the Navigate and Events views, see the *NetWitness Investigate User Guide*.

Note: If the values are not indexed, the results take time to load. For more information, see Troubleshooting NetWitness Endpoint.

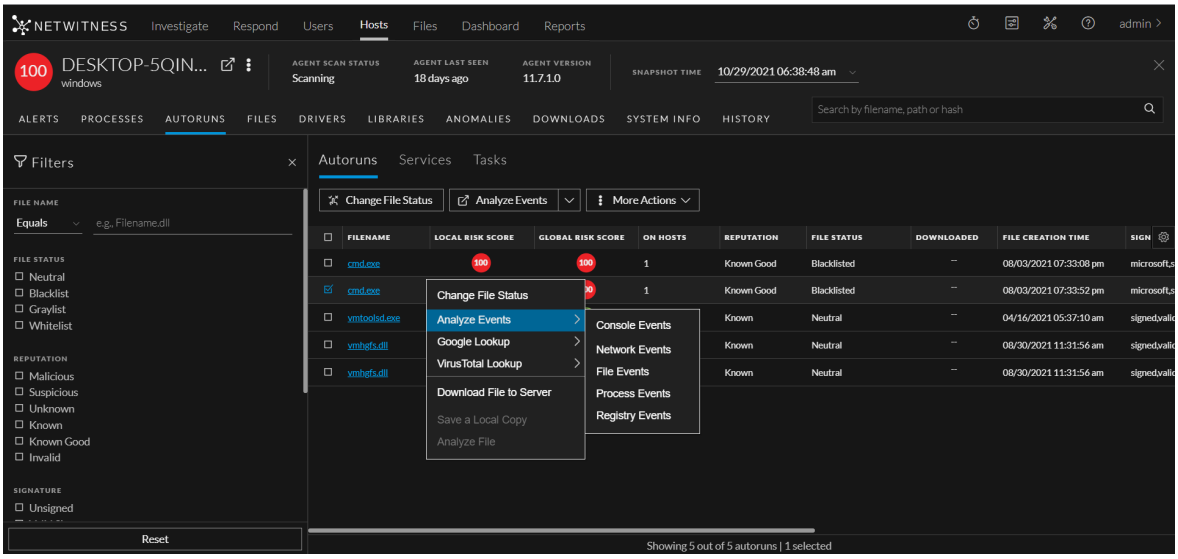
Analyze Events from Hosts View

To investigate a particular host, IP address (IPV4), or username:

1. Go to **Hosts**.
2. Do one of the following:

- Right-click a host, select **Analyze Events** from the right-click context menu or in the toolbar, and select a specific event type (such as network events, file events) that you want to view.
- Select the hostname to open the host details. Right-click a file or in the toolbar, select **Analyze Events**, and select a specific event type that you want to view.

  The following figure is an example of the Autoruns tab.



  This opens the Navigate view with data related to the file.

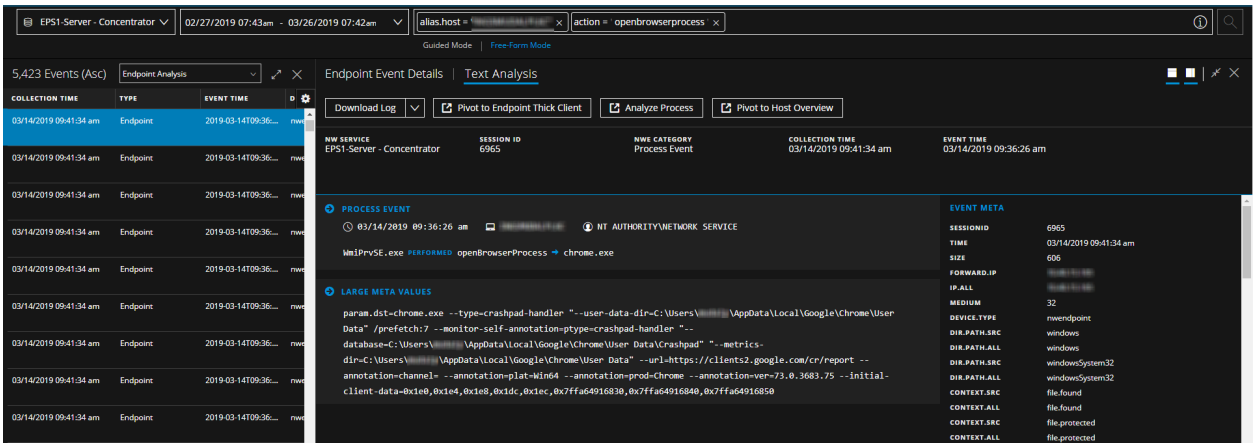
For more information on analyzing events in the Navigate and Events views, see the *NetWitness Investigate User Guide*.

Text Analysis for an Endpoint Event

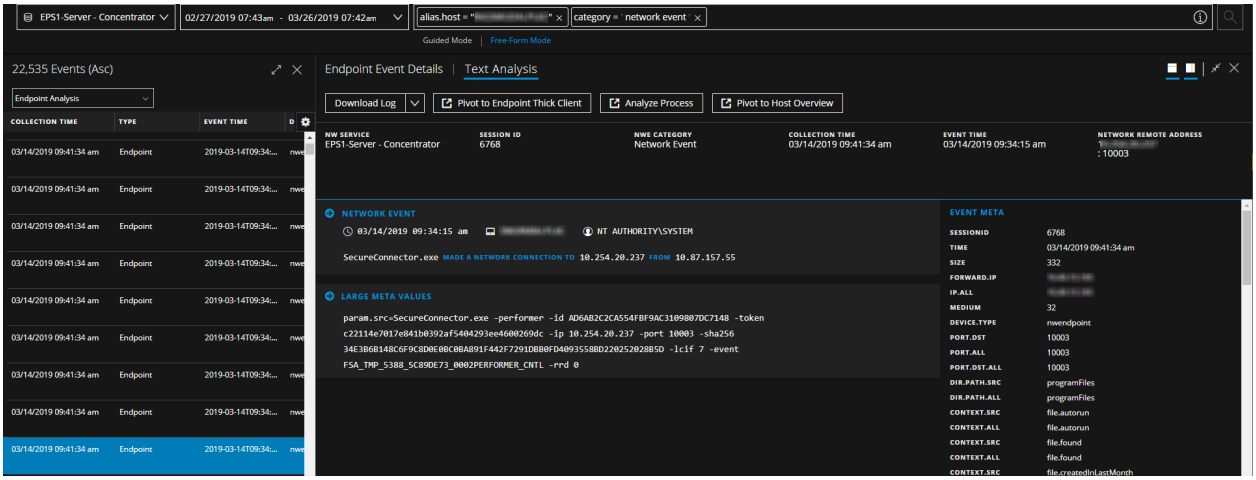
You can view all Endpoint events in their original text format in the Events view Event List panel. When you click an event in the Event list panel, the adjacent panel shows the Text Analysis. Pagination controls add flexibility when paging through the reconstructed text of an event. The Text Analysis displays the following:

- Event Header, which provides summary information about the event. (Version 11.5 and Later) The event header includes host name, process, and user name details in addition to other event details if the selected Endpoint event contains these metadata.
- Options for exporting - log, csv, xml, and json formats.
- Option to pivot to the Endpoint Thick Client to analyze the meta value.
- Option to analyze process details associated with the event.
- Option to view the host details for further analysis.

Below is an example of the Process event for Endpoint. The text in the Text Analysis panel explains that a source process `WmiPrvSE.exe` opened a browser process named `chrome.exe`. In the events, if there is a meta value that exceeds 255 characters, the value is displayed in the Large Meta Values panel.



Below is an example of the Network event:



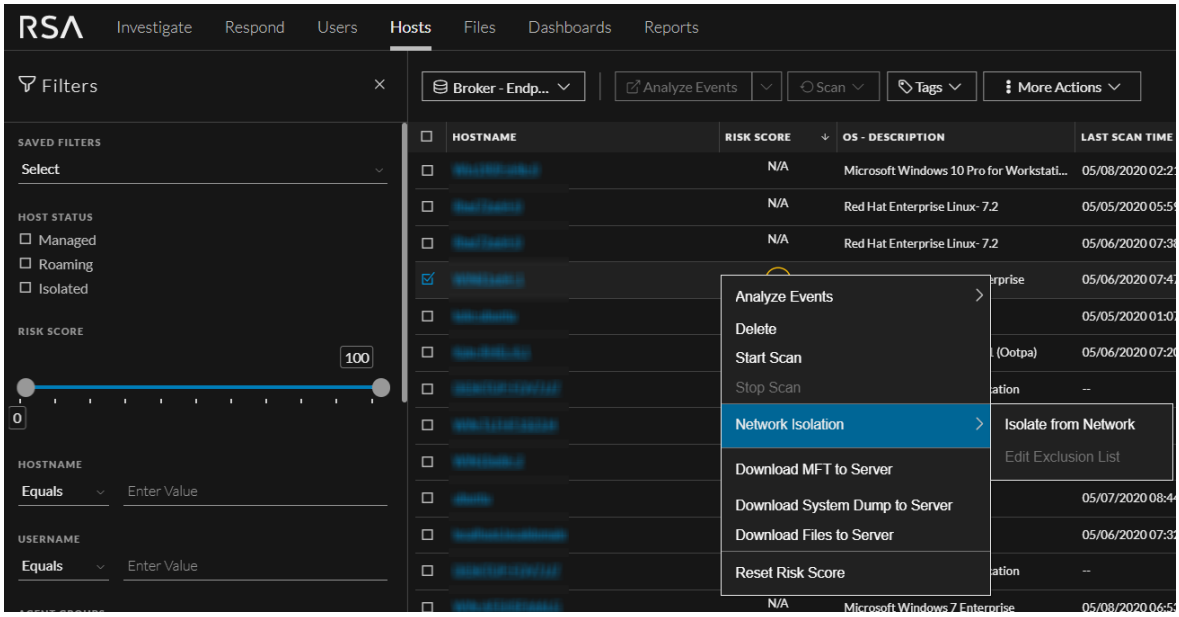
For more information on Events view, see the *NetWitness Investigate User Guide*.

Isolating Hosts from Network

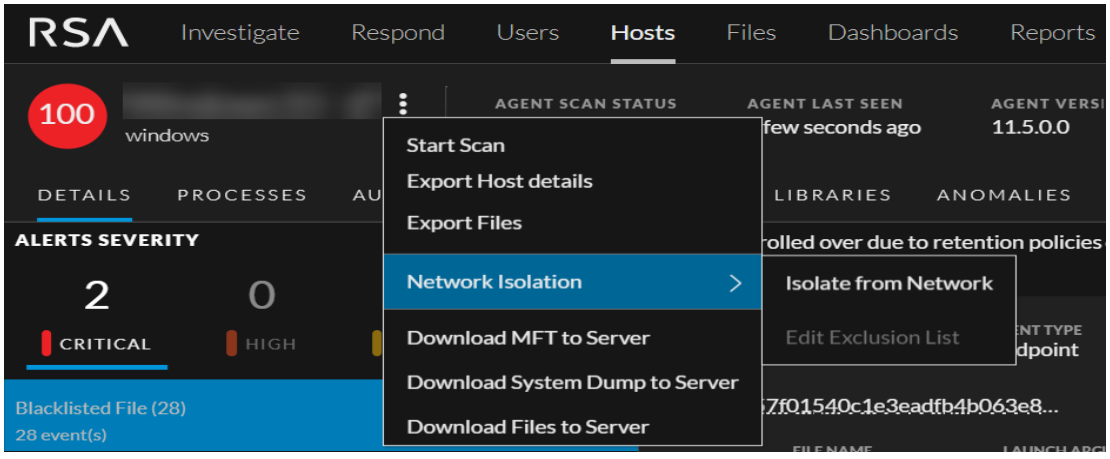
Note: By default, the network isolation option is disabled in the policy, and you cannot view options mentioned in this section. To enable network isolation, in the policy configuration, select **Enabled** in the **Network Isolation** option under Response Action Settings. For more information, see the *NetWitness Endpoint Configuration Guide*.

To isolate a host from the network:

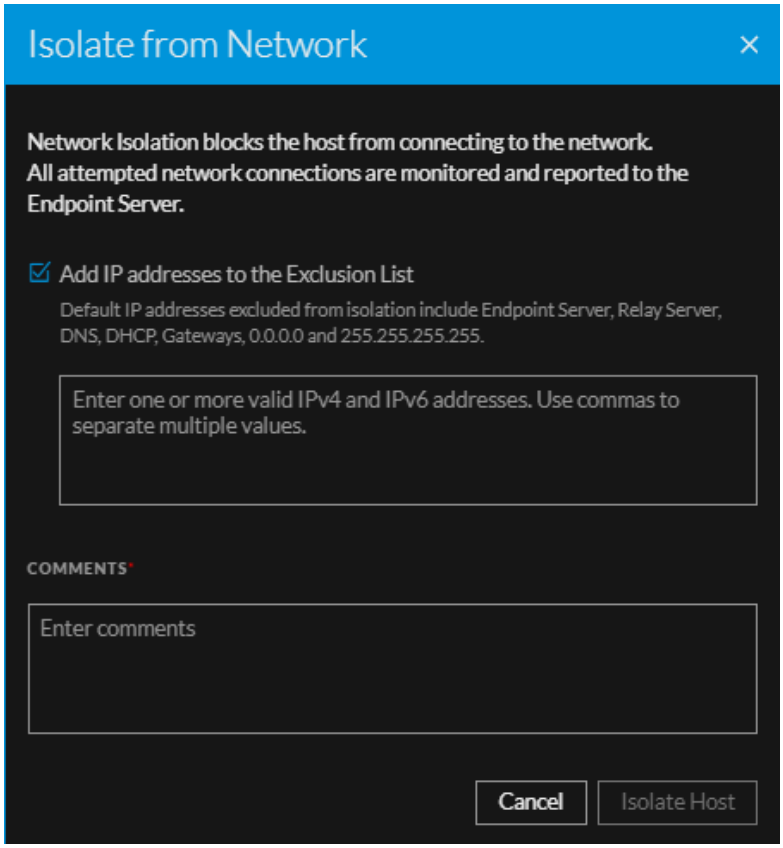
1. Go to **Hosts** and do one of the following:
   - Select a host and select **Network Isolation > Isolate from Network** from the right-click context menu, or from the **More** drop-down list in the toolbar.



   - Select the hostname to open the host details, click ⋮ (**More**) beside the hostname, and select **Network Isolation > Isolate from Network**.



2. In the Isolate from Network dialog, by default, a set of IP addresses are excluded from isolation. For more information, see Network Isolation. To add IP addresses to the list, select the **Add your IPs to Exclusion List** checkbox. You can enter up to 100 IP addresses separated by comma.

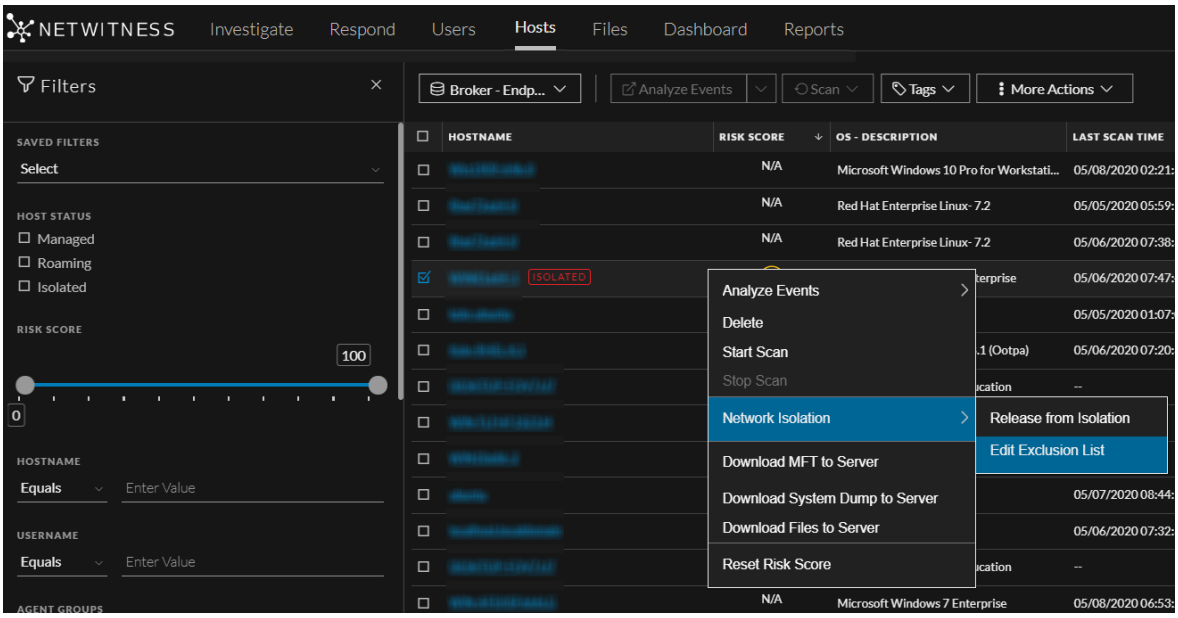


3. Enter comments.
4. Click **Isolate Host**.

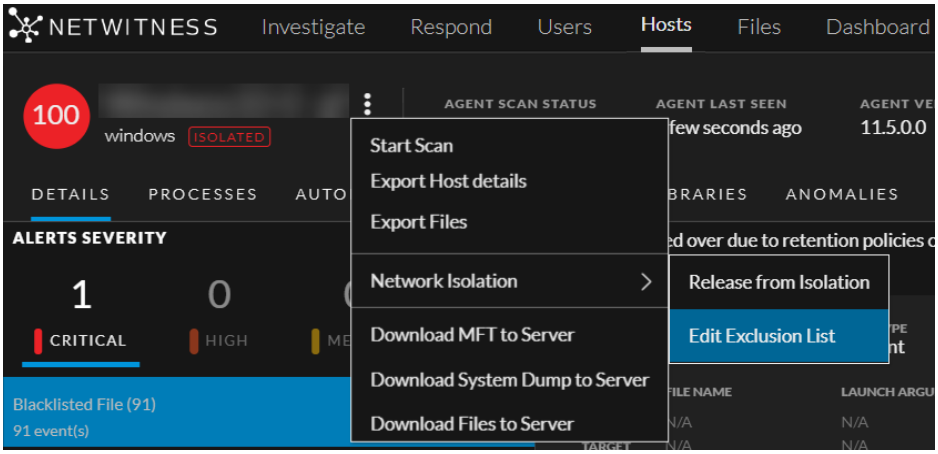
Edit Exclusion List

To edit the exclusion list:

1. Go to **Hosts** and do one of the following:
   - Select a host and select **Network Isolation > Edit Exclusion List** from the right-click context menu, or from the **More** drop-down list in the toolbar.



- Select the hostname to open the host details, click ⋮ (**More**) beside the hostname, and select **Network Isolation > Edit Exclusion List**.

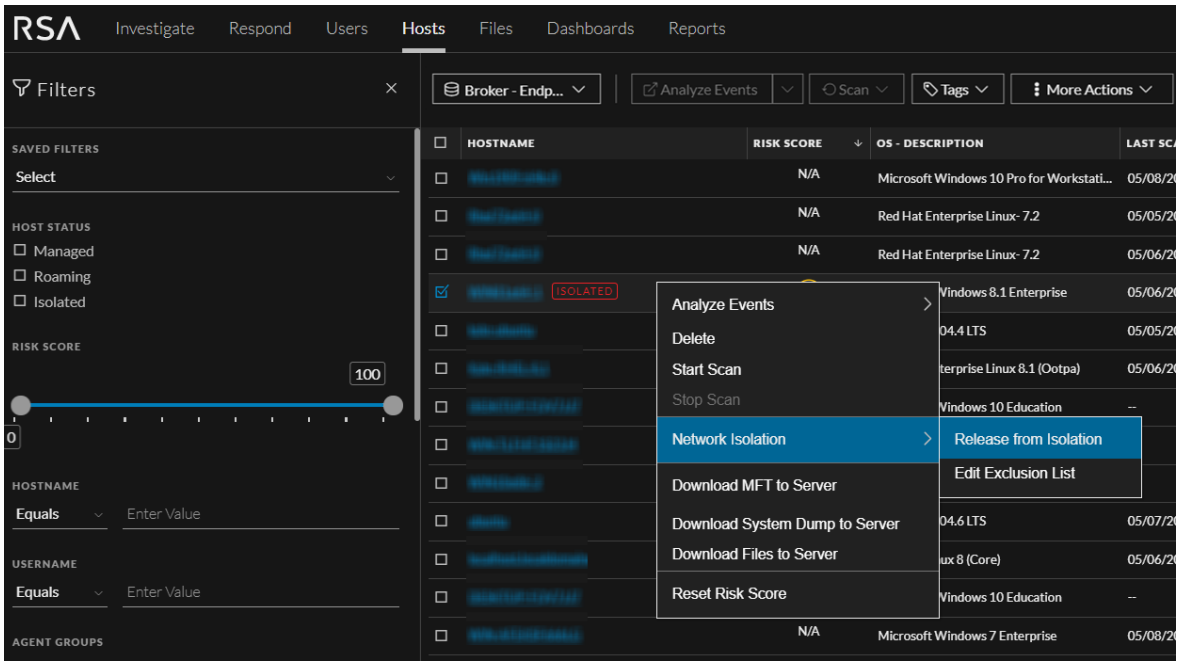


2. Add or modify the IP addresses in the list.
3. Enter comments and click **Save**.

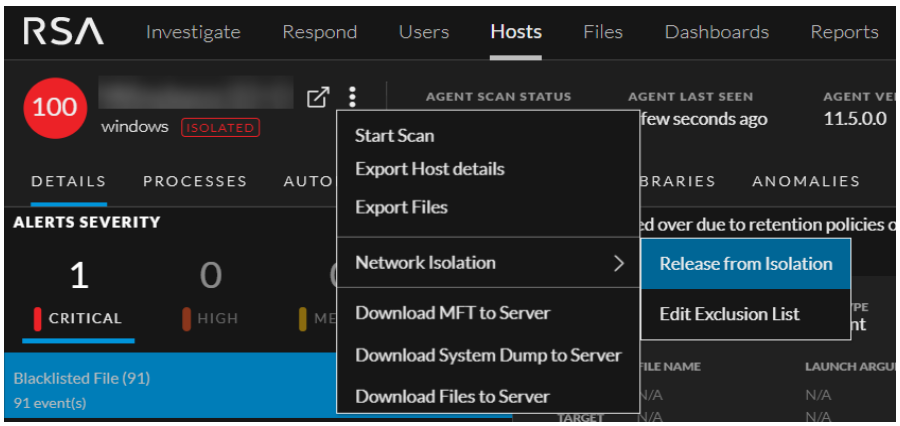
Release Isolated Hosts

Releasing the isolated host restores the network connection and removes IP addresses added to the Exclusion list. To release the host from isolation:

1. Go to **Hosts** and do one of the following:
   - Select a host and select **Network Isolation > Release from Isolation** from the right-click context menu, or from the **More** drop-down list in the toolbar.



- Select the hostname to open the host details, click ⋮ (**More**) beside the hostname, and select **Network Isolation > Release from Isolation**.



2. Enter comments and click **Release Host**.

NetWitness Endpoint with Third-Party Antivirus Products

If you want the NetWitness Endpoint agent to coexist with any of the security products, make sure to whitelist the agent. For more information, see the respective third-party product documentation.

Troubleshooting NetWitness Endpoint

This section provides information about possible issues when using NetWitness Endpoint.

General Issues

Issue	Some of the hosts or files data are not displayed when Endpoint Broker is selected for querying.
Solution	The Endpoint Broker aggregates data from all Endpoint Servers, which responds within 10 seconds. You must increase the query timeout value to see the result of Endpoint server that is online. Perform the following: 1. Go to  **(Admin) > Endpoint Broker** service. 2. Click  > **View** > **Explore**. 3. Click **endpoint/broker** node. 4. In the **query-timeout** field increase the value, for example, 30 seconds.

Issue	The Endpoint Agent is unable to communicate with the Endpoint Server. The connection may not be established due to any of the following issues: • UDP • HTTPS • Firewall

Solution	• To verify the UDP or HTTPS connection, you must verify the connection between Windows Endpoint Agent and Endpoint Server:
	1. Go to System32 folder using the following command:
	`cd C:\Windows\System32`
	2. Execute the following command:
	`<Agent Service name>.exe /testnet`
	For example, `NWEAgent.exe /testnet`
	• If the issue is with the firewall, check the incoming and outgoing firewall rules.

Issue	The Endpoint Agent is unable to communicate with the Log Decoder. The connection may not be established due to any of the following issues: • UDP • TCP • TLS

Solution	• Firewall

	• To verify the UDP, TCP, and TLS connection, you must verify the connection between Windows Endpoint Agent and the Log Decoder:
	1. Go to System32 folder using the following command:
	`cd C:\Windows\System32`
	2. Execute the following command:
	`<Agent Service name>.exe /testlognet`
	For example, `NWEAgent.exe /testlognet`
	• If the issue is with the firewall, check the incoming and outgoing firewall rules.

Multi-server Issue

Issue	Agent is not communicating with the Endpoint Server after migration.
Solution	Check the Nginx logs of the Endpoint Server to which the agent has migrated, and if the agent is communicating with error code 403, that means the certificate of the first Endpoint Server and second Endpoint Server are different. This is because during the installation of second Endpoint Server, the certificate of first Endpoint Server is not copied to the second Endpoint Server. Reinstall the second Endpoint Server by copying the certificate of first Endpoint Server, and reinstall the agent. For more information, see the *Physical Host Installation Guide*.

Hosts View Issues

Message	`An error has occurred. The Endpoint Server may be offline or inaccessible.`
Issue	When attempting to access the Hosts or Files view, the view opens with the message.
Explanation	Endpoint Server or Nginx Server is not running. Check the status of the Endpoint Server under  **(Admin) > Services** or check if the Endpoint Server host IP address is registered with the Admin Server. For more information, see the *Physical Host Installation Guide* or *Virtual Host Installation Guide*. If the service is not running, start the Endpoint Server.

Issue	Hosts view shows 'No Results Found.' error in the following scenario: • Host A belongs to Endpoint server A and does not exist in Endpoint server B. • Endpoint server B is selected in Hosts view. • In a new window/tab, open the Host details page (of Host A) • Navigate to Hosts view using **[*Host Name*] > Actions > Pivot to Investigate**

	> Hosts/Files `Endpoint server B is selected by default, and the page shows 'No results found.' error.`
Explanation	This is expected behavior. The Endpoint server that is selected first is considered to be current /active throughout the session.
Workaround	Ensure the correct Endpoint server is selected, and `Pivot to Investigate > Hosts/Files` will behave normally.

Issue	For MFT download: 1. The request fails with file not found error for drives/mount paths created on the machine after agent is installed. 2. Incorrect MFT is downloaded for the provided NTFS mount path.
Solution	Make sure that the agent version is 11.6 or later.

Files View Issues

Issue	Unable to analyze events from **Hosts** and **Files** view.
Explanation	Other than Broker or Concentrator, if any aggregation service, such as Archiver, is aggregating data from the Log Decoder that is configured for metadata forwarding from any Endpoint server, clicking **Analyze Events** from **Hosts** and **Files** view for this Endpoint server may not work. To resolve this issue: **Note:** To get the investigate-service-id: 1) Go to **(Admin) > Services > Concentrator** service. 2) Click > **View** > **Explore** tab. 3) Expand the **sys/stats** node list. 4) In the UUID filed, copy the value. 1. Go to **(Admin) > Services > Endpoint Server** service. 2. Click > **View** > **Explore** tab. 3. In the **endpoint/investigate** field, specify the **investigate-service-id**.

Policy Issue

Issue	Policy status in the Policy Details panel is not updated or shows Policy Unavailable/Permission Required.
Explanation	**Policy Unavailable** - Hosts belong to previous versions, such as NetWitness Platform 11.1 or 11.2, where a policy is not applied.

	Permission Required - If you do not have permissions, see the "Role Permissions" topic in the *System Security and User Management Guide*.

Issue	Policy Status shows error.
Explanation	Policy may have wrong configurations. Check the error description, logs in Endpoint server, and audit logs for details. Contact your system administrator with the error details.

Driver Issue

Issue	While loading the driver on the host, an error is encountered.
Explanation	Check the driver error code in the **Agent-Driver Error Code** column under **Hosts** view. Contact your system administrator with the error code.

Download Issue

Issue	Downloads (Files, MFT, System/Process dumps, etc.) fail at times.
Explanation	Downloads fail when there is not sufficient disk space on the Endpoint Server.
Workaround	Clean up some disk space and try downloading again. We recommend you keep sufficient disk space before initiating any download.

File Reputation Service Issue

Issue	When you configure RSA Live for the first time and the File Reputation service is not connected.
Solution	You must manually enable the File Reputation service. To enable the File Reputation service: 1. Go to  **(Admin) > System > Live Services**. 2. In the **Additional Live Services** section, select the enable File Reputation check box. 3. Click **Apply**.

Risk Scoring for Hosts or Files Issue

Issue	NetWitness Endpoint takes a long time to process risk scoring for Hosts or Files.

Solution	Check the backlog of alerts for risk scoring. 1. SSH to the ESA Primary appliance. 2. Execute the following command: `mongo respond-server --authenticationDatabase admin -u deploy_admin -p <deploy_admin_password> --eval 'db.staging.find({"$or":[{state:"STAGED"},{state :"WORKING"}]}).count()' --quiet` The backlog count is displayed. If the backlog count is 1 million or greater, you must disable the risk scoring and Endpoint ESA alerts. 3. To disable risk scoring: a. Go to **(Admin) > Services > Respond** service. b. Click > **View** > **Explore**. c. Expand the **respond/scheduled/jobs** node list. d. In the **risk-scoring-enabled** field, set the value to **false**. 4. To disable Endpoint ESA alerts: a. To disable NetWitness Endpoint ESA alerts generation for severity; Critical, High and Medium. i. Go to **(Configure) > ESA Rules**. The Configure view is displayed with the Rules tab open. ii. In the **Options** panel, under Deployments, select the Endpoint deployment to delete. A confirmation dialog is displayed. iii. Click **Yes**. b. To disable only Medium severity NetWitness Endpoint ESA alerts: i. Go to **(Admin) > ESA Correlation** service (on which Endpoint deployment is added). ii. Click > **View** > **Explore**. iii. Expand the **correction/alert** node list. iv. In the **transient-enabled** field, set the value to **false**.

Endpoint Broker/Server Issue

Issue	User have access to one Endpoint server and unable to access the other.
Explanation	Reach out to the administrator and check if you have access to that endpoint servers. Request for access if required.

Issue	In the Endpoint Server Broker view, the user can scan hosts that belong to a particular Endpoint server but cannot scan hosts that belong to another Endpoint server.
Explanation	User may not have access to that endpoint server. Check with the Administrator to see if the user has access.

NetWitness Endpoint Reference Materials

This section is intended to help you understand the purpose and application of NetWitness **Investigate** > **Hosts** view and **Files** view. For each view, there is a brief introduction and a What Do You Want To Do table with links to related procedures. In addition some of the reference materials include workflows and Quick Looks to highlight important features in the user interface.

- Files View
- Hosts View
- Hosts View - Details Tab
- Hosts View - Process Tab
- Hosts View - Autoruns Tab
- Hosts View - Files Tab
- Hosts View - Drivers Tab
- Hosts View - Libraries Tab
- Hosts View - Anomalies Tab
- Hosts View - Downloads Tab
- Hosts View - System Information Tab
- Hosts View - Agent History Tab

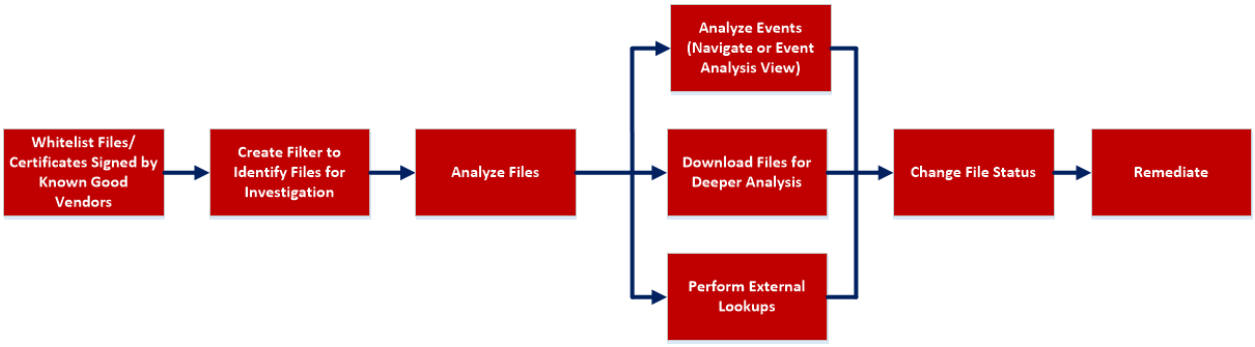
Files View

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Files view provides a holistic view of all files in your deployment. To access this view, go to **Files**. By default, the Files view displays 100 files. To display more files, click **Load More** at the bottom of the page.

You can either view files specific to an Endpoint server or view all files from multiple Endpoint servers by selecting the Endpoint Broker.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	whitelist files and certificates signed by known good vendors*	Analyze Certificates
Threat Hunter	create filter to identify files for investigation*	Filter Files
Threat Hunter	analyze files*	Investigating Files
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate

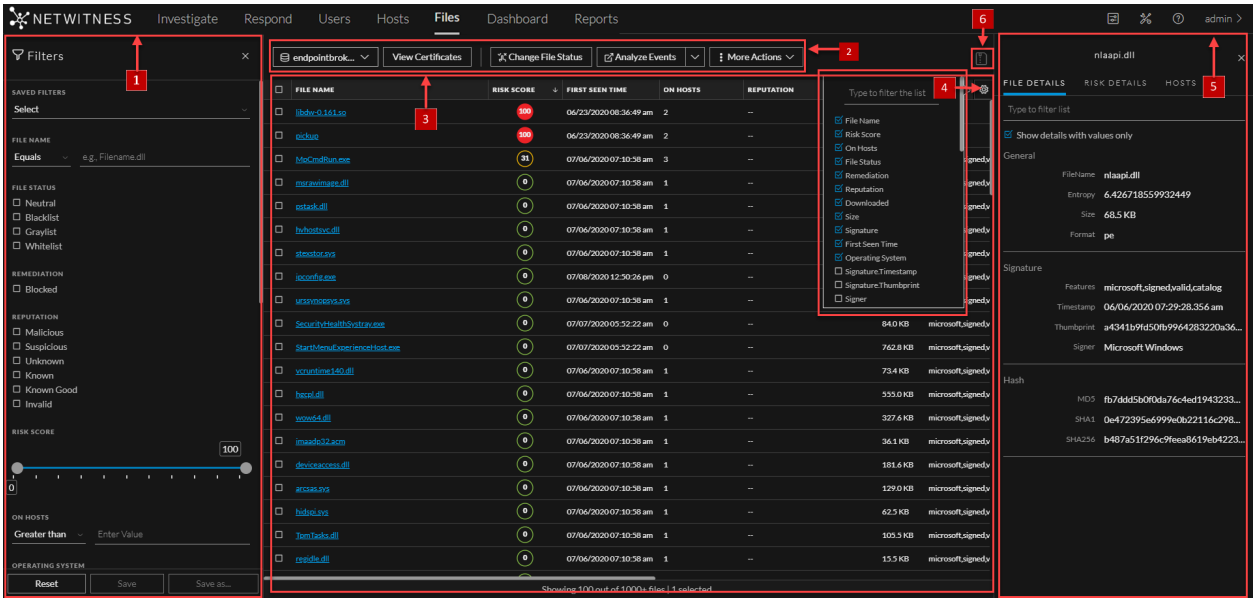
*You can perform this task in the current view

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Analyzing Events
- Analyze Certificates
- Changing File Status or Remediate

Quick Look

Below is an example of the Files view:

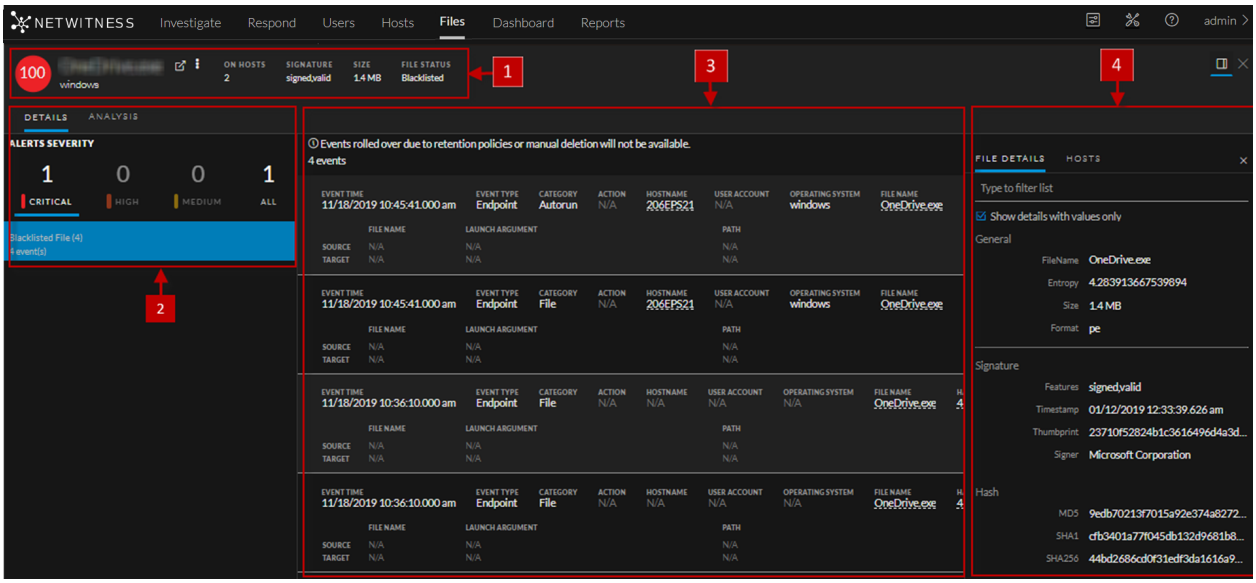


1 **Filter Files**. You can filter the files by selecting the options in the Filters panel and create filters. For more information, see Filter Files.

2	**Actions in the toolbar:** **Server drop-down list** - You can select the Endpoint server or Endpoint Broker server to view the hosts. **View Certificates** - Provides a list of code-signing certificates reported by hosts found in your deployment and their associated properties. For more information, see Analyze Certificates. **Change File Status** - Provides capabilities to manage suspect and legitimate files and block malicious or infected files to prevent future execution of the file on any host. For more information, see Changing File Status or Remediate. **Analyze Events** - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events. **More Actions** - Provides options to: • Perform external lookups. • Download files to server, save a local copy, and analyze files for deeper analysis. • Reset risk score. **Note:** You can perform the above actions from the right-click context menu.
3	**Sort Columns.** Lets you sort on column titles.
4	**Settings Menu.** You can set Files view preferences by selecting columns from the Settings menu. For more information, see Set Files Preference.
5	**Show/Hide File Properties Panel.** Click a row to show or hide the File Properties panel. It displays the following tabs: **File Details** - Displays the file information. **Risk details** - Displays the distinct alerts associated with the risk score. **Hosts** - Displays the top 100 hosts based on the risk score on which the file is present. For more information, see Analyze Hosts with File Activity.
6	**Export to CSV** - Extracts global files to a CSV file. For more information, see Export Global Files.

File Details View

To access this view, go to **Files**, and select a file. Below is an example of the File Details view:



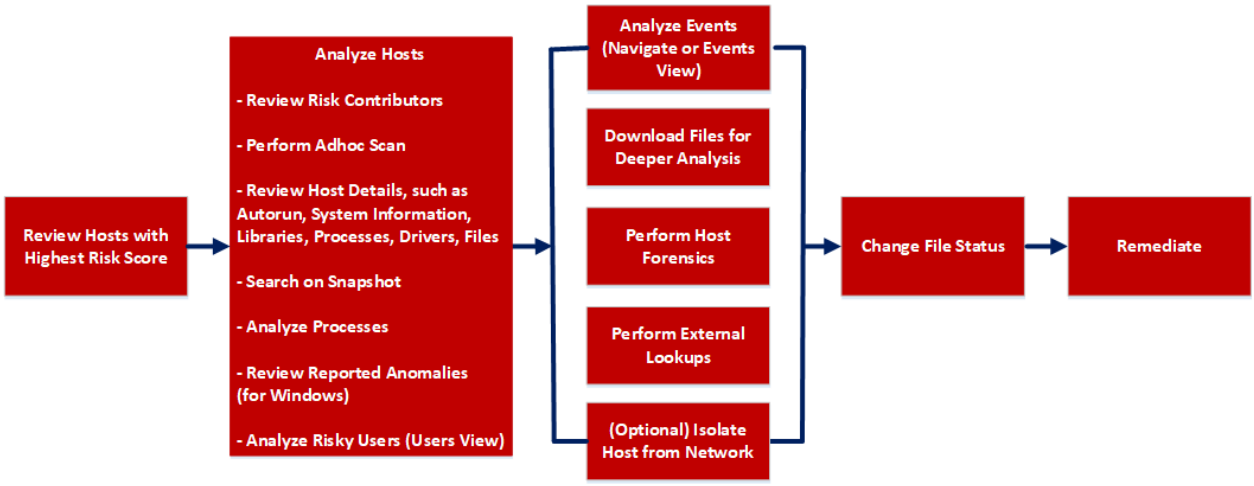
1	**Agent and Scan Details**. You can view the following agent and scan details of the selected host: **Host name** - Name of the host. For example, `WIN-ABC`. **Risk score** - Risk score of the host. **Operating System** - Operating system on which the agent is running (Linux, Windows, or Mac). **Analyze Events** - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events. **More** - Provides options to perform external lookups. **On Hosts** - Indicates the number of hosts on which a file exist. **Signature** - Provides signatory information. **Size** - Size of the file. **File Status** - Status of the file. For example, `Neutral`.
2	**Details** tab - Displays list of distinct alerts, such as Critical, High, Medium and All, along with the total number of events associated with the alert. **Analysis** tab - Provides detailed information about a downloaded file. For more information, see Analyzing Downloaded Files.
3	Displays events for an alert and metadata associated with a specific event.
4	**Show/Hide File Properties Panel.** Click a row to show or hide the File Properties panel. It displays the following tabs: **File Details** - Displays the file information. **Hosts** - Displays the hosts on which file activities are present. For more information, see Analyze Hosts with File Activity.

Hosts View

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Hosts view provides a list of all hosts with an Endpoint agent installed. To access this view, go to **Hosts**. By default, hosts are sorted based on the risk score.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details*	Analyze Host Details
Threat Hunter	search on snapshot*	Search Files on Host
Threat Hunter	analyze processes*	Investigating a Process
Threat Hunter	review reported anomalies*	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users*	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate
Threat Hunter	filter files*	Filter Host Details
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT*, system dump*, or process dump	Performing Host Forensics

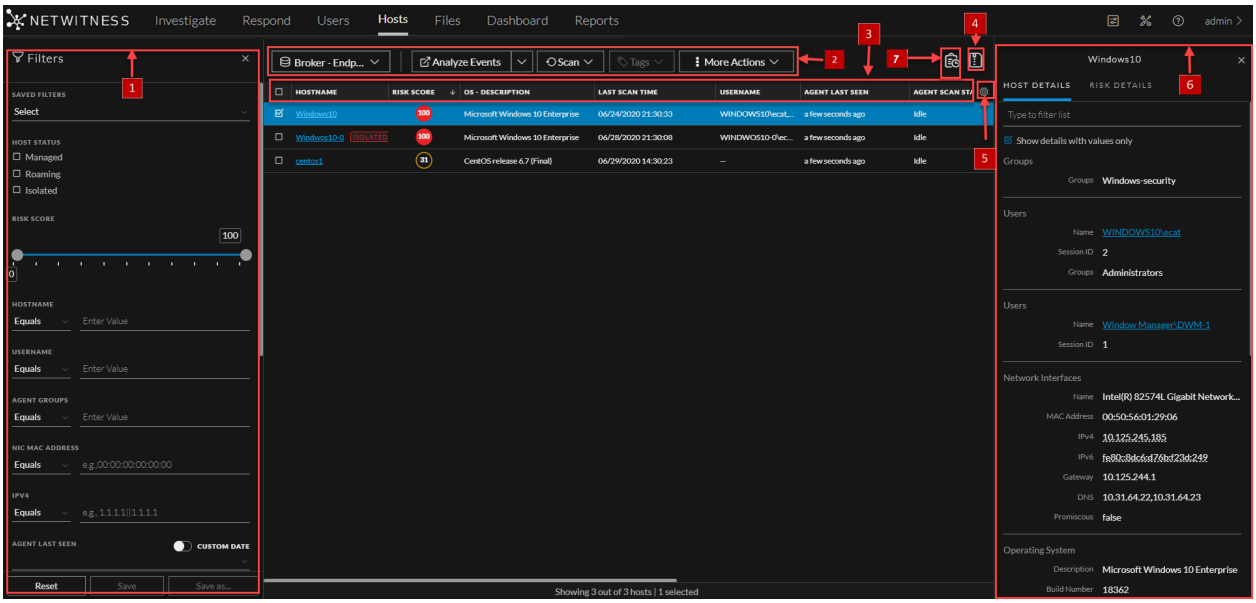
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Changing File Status or Remediate
- Investigating a Process
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Hosts view:



1	**Filter Hosts**.You can filter the hosts by selecting the options in the Filters panel and create filters. For more information, see Filter Hosts.
2	**Actions in the toolbar:** **Server drop-down list** - You can select the Endpoint server or Endpoint Broker server to view the hosts. **Analyze Events** - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events. **Start Scan** - Starts a scan for the selected hosts. **Stop Scan** - Stops a scan for the selected hosts. **More Actions** - Provides options to: • Reset risk score. • Delete - Lets you delete hosts manually from the user interface. After deletion, the Endpoint server does not process any request from this host. For more information, see Delete a Host. • Download MFT to the server. For more information, see Performing Host Forensics. • Download System Dump to the server. For more information, see System and Process Memory Dump. • Isolation host from the network. For more information, see Isolating Hosts from Network. **Note:** You can perform the above actions from the right-click context menu.
3	**Sort Columns**. Lets you sort on column titles.
4	**Export to CSV** - Extracts host attributes to a CSV file. For more information, see Export Host Attributes.

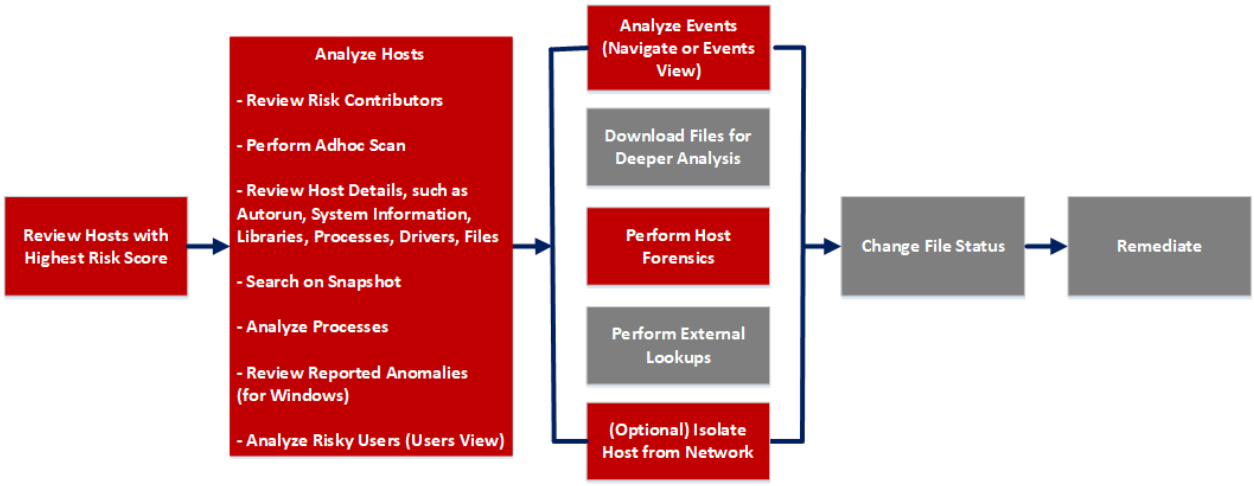
5	**Settings Menu.** You can set Hosts view preferences by selecting columns from the Settings menu. For more information, see Set Hosts Preference.
6	**Show/Hide Host Properties Panel.** Click a row to show or hide the Host Properties panel. It displays the following tabs: **Host details** - Displays the host information such as Network Interfaces, operating system, hardware and others. **Risk details** - Displays the distinct alerts associated with the risk score.
7	**View Agent History** - Displays the list of commands issued to the agent. For more information, see View Agent History.

Hosts View - Details Tab

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Details tab provides details of the selected host. To access this view, go to **Hosts** view, and select a host.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details*	Analyze Host Details
Threat Hunter	search files on host*	Search Files on Host
Threat Hunter	analyze processes*	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users*	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis	Analyzing Downloaded Files
Threat Hunter	perform external lookups	Launch an External Lookup for a File
Threat Hunter	change file status or remediate	Changing File Status or Remediate
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT*, system dump, or process dump*	Performing Host Forensics

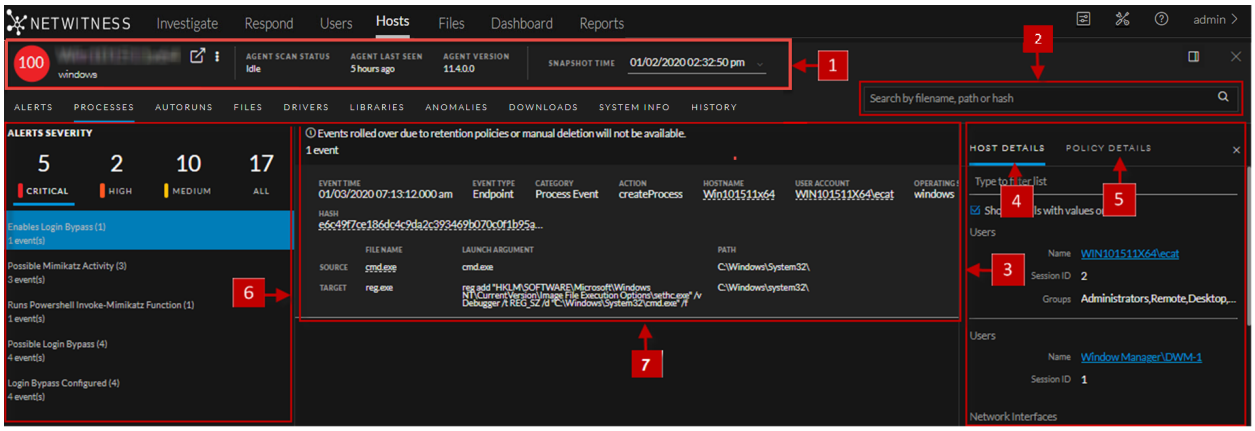
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Details tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

Analyze Events - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events.

More - Provides options to:

- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

Snapshot Time - Lists scanned time stamps. To view the scan history, you can select the snapshot time from the drop-down menu.

2 Search files on host. Lets you search the files on the host (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.

3 **Show/Hide Right Panel** - Displays host and policy details panel.

4 **Host Details Panel** - Displays all properties of the selected host. It is grouped as follows:

Groups - Groups on which the host is added on.

User - Information related to the user.

Network Interfaces - Network adapter information, such as Mac Address, Gateway.

Operating System - Operating system version and build information.

Agent - Agent-related information, such as agent ID, driver error code, install time, and agent mode.

Hardware - Information related to the architecture.

Locale - Time zone and language that is local to the host.

5 **Policy Details Panel** - Displays the following:

- Policy Status -
  - Updated - Host has the latest policy.
  - Pending - Policy is resolved but the latest policy is not updated on the host. When the host communicates with the Endpoint server next time, the latest policy is applied if there are no errors.
  - Unavailable - Hosts that belong to previous versions, such as NetWitness Platform 11.1 or 11.2, or the source server is not installed.
  - Error - Problem applying the latest policy along with the error description.
- Blocked Hashes Status -
  - Updated - Host has the latest blocked hashes.
  - Pending - Hashes are blocked but the latest hashes are not updated on the host. When the host communicates with the Endpoint server next time, the latest hashes are applied if there are no errors.
- Evaluated Time - Time when the Endpoint server evaluated the policy.
- Relay Server. Displays the Relay Server details.
  - Server - Host name or IP address of the Relay Server.
  - Port - Port number.
  - HTTP Beacon Interval - HTTP beacon interval value in minutes.
- Complete resolved policy settings. For more information, see "Managing Policies" in the *NetWitness Endpoint Configuration Guide*.

Note: The values that are not set in the policy are not displayed.

6 **Alerts Severity** - Displays list of distinct alerts, such as Critical, High, Medium and All, along with the total number of events associated with the alert.

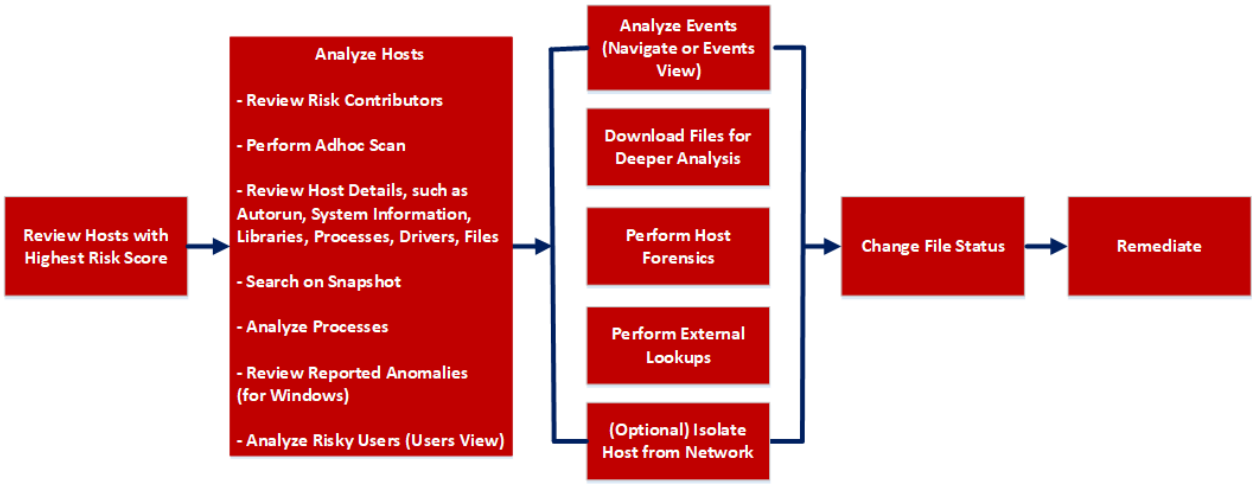
7 Displays events for an alert and metadata associated with a specific event. For more information, see Analyze Hosts Using the Risk Score.

Hosts View - Process Tab

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Process panel provides a list of processes running on the host. To access this tab, select a host from the **Hosts** view and click the **Process** tab.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search on snapshot*	Search Files on Host
Threat Hunter	analyze processes*	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate
Threat Hunter	filter files*	Filter Host Details
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT, system dump, or process dump*	Performing Host Forensics

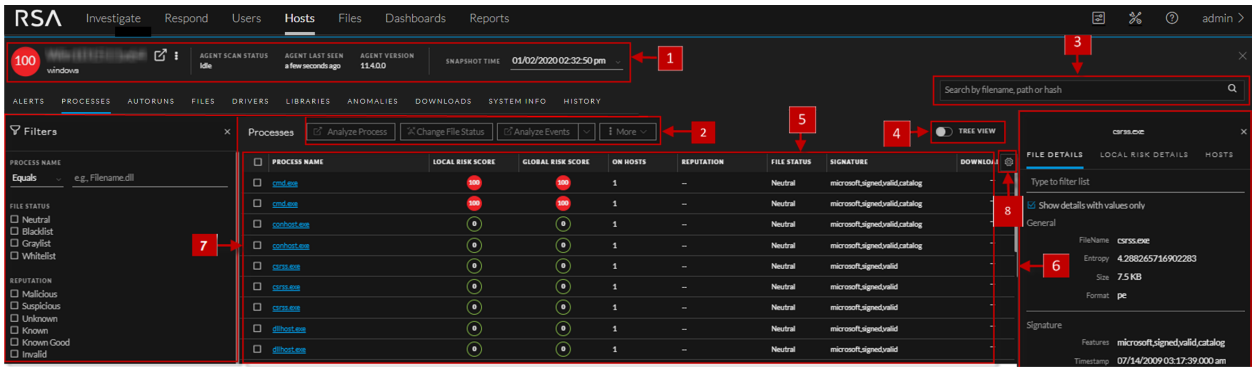
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Changing File Status or Remediate
- Investigating a Process
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Process tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

More - Provides options to:

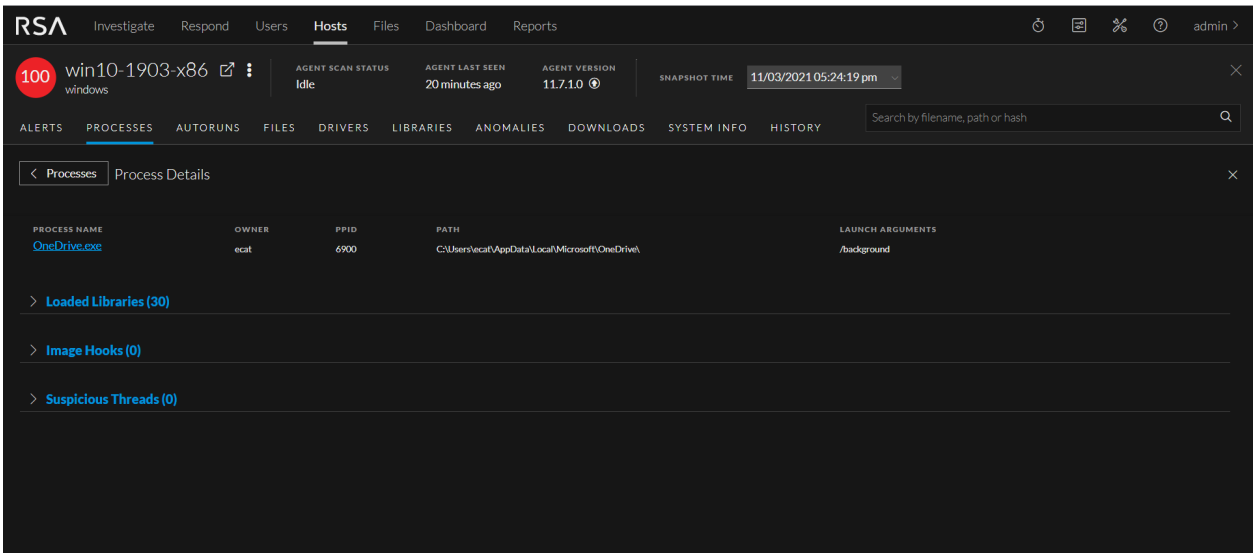
- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

Snapshot Time - Lists scanned time stamps. To view the scan history, you select the snapshot time from the drop-down menu.

2	**Actions in the toolbar:** **Analyze Process** - Lets you perform process analysis to investigate a particular process behavior, and understand the entire process event chain, process parent-child relationships, and all associated events. For more information, see Investigating a Process. **Change File Status** - Provides capabilities to manage suspect and legitimate files and block malicious or infected file to prevent future execution of the file on any host. For more information, see Changing File Status or Remediate. **Analyze Events** - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events. **More** - Provides options to: • Perform external lookups. • Download files to server, save a local copy, and analyze files for deeper analysis. • Download process dump to server. **Note:** You can perform some of the above actions from the right-click context menu.
3	**Search on Snapshots**. Lets you search on all snapshots (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.
4	**Toggle.** Lets you toggle between List view and Tree view.
5	**Process panel** - Displays process information, such as process name, local risk score, global risk score, On Hosts, reputation status, file status, and others.
6	**Show/Hide Right Panel** - Displays the following properties of a process in the right panel: • **File Details** - Displays all properties of the selected process. It is grouped as follows: **General** - General information about the file, such as file name, entropy, size, and format. **Signature** - Provides signatory information. **Hash** - Hash type of the file (MD5, SHA1, and SHA256). **Time** - Time when the file was created, modified, or accessed. **Location** - Location of the file. **Process** - Details of the process, such as image size and PID. **Image** - Image details loaded by the process. • **Local Risk Details** - Displays the alerts associated with the local risk score, such as Critical, High, Medium and All. • **Hosts** - Displays the top 100 hosts based on the risk score on which the file is present.
7	**Filter Files**. You can filter processes by selecting the options in the Filters panel and create filters. For more information, see Filter Host Details.
8	**Settings Menu.** You can set Hosts view preferences by selecting columns from the Settings menu. For more information, see Set Hosts Preference.

Process Details

Clicking the process name displays the process details of a specific process as shown in the following figure:



Field	Description
Process Name	Name of the process. For example, `server.exe`.
PID	ID of the process. For example, `492`.
Path	Path of the file associated with the process on the disk. For example, `C:\Windows\System32`.
Launch Arguments	Command line arguments passed to the process when it is launched. For example, `-k LocalServiceNoNetwork`.

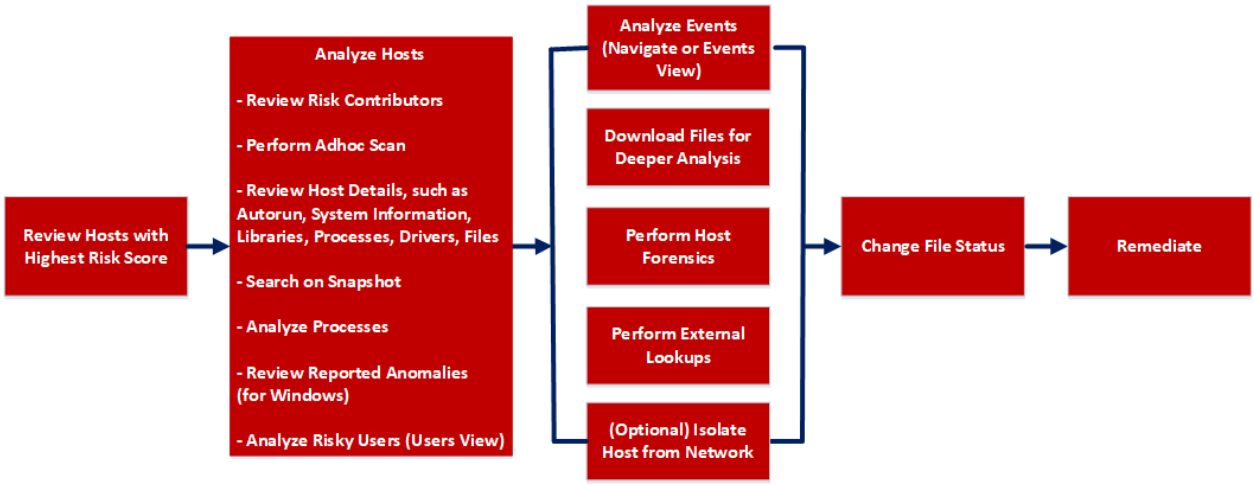
- List of loaded libraries for the selected process, such as DLLs (for Windows), Dylibs (for Mac), or .SO (for Linux).
- List of autoruns (if configured).
- List of image hooks and suspicious threads (for Windows).

Hosts View - Autoruns Tab

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Autoruns panel provides a list of autoruns, services, tasks, and cron jobs running on the host. To access this tab, select a host from the **Hosts** view and click the **Autoruns** tab.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search on snapshot*	Search Files on Host
Threat Hunter	analyze processes	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate
Threat Hunter	filter files	Filter Host Details
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT, system dump, or process dump*	Performing Host Forensics

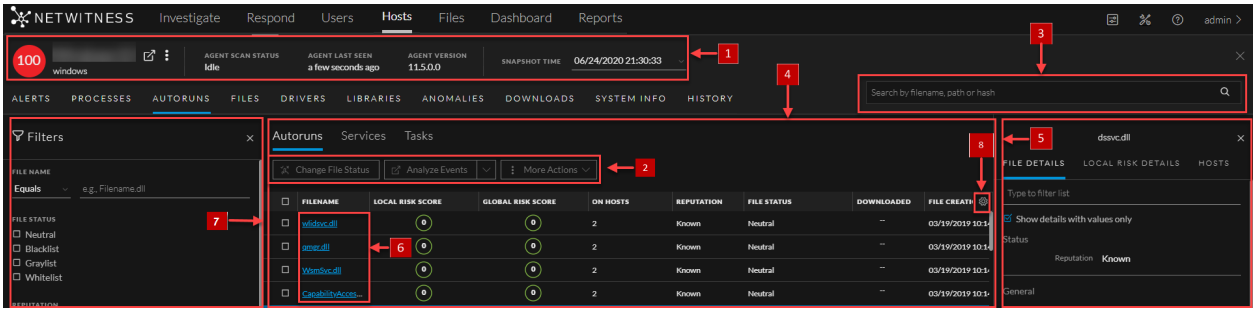
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Changing File Status or Remediate
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Autoruns tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

More - Provides options to:

- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

Snapshot Time - Lists scanned time stamps. To view the scan history, you can select the snapshot time from the drop-down menu.

2 **Actions in the toolbar:**

Change File Status - Provides capabilities to manage suspect and legitimate files and block malicious or infected file to prevent future execution of the file on any host. For more information, see Changing File Status or Remediate.

Analyze Events - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events.

More Actions - Provides options to:

- Perform external lookups.
- Download process dump to server.
- Download files to server, save a local copy, and analyze files for deeper analysis.

Note: You can perform some of the above actions from the right-click context menu.

3 **Search on Snapshots**. Lets you search on all snapshots (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.

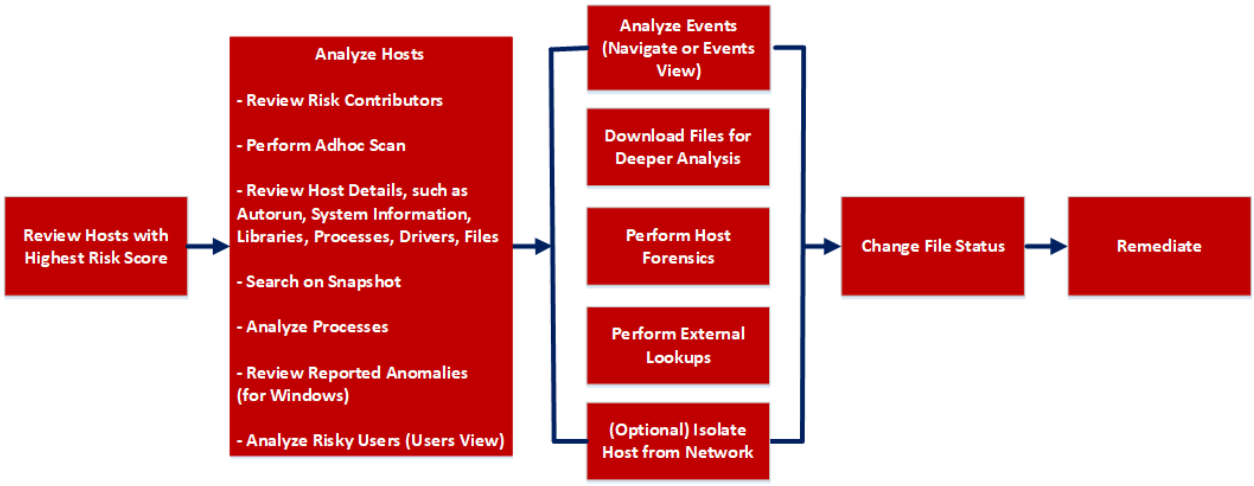
4	**Details Panel** - Displays the following tabs: • **Autoruns** - Files that are executed at start-up. • **Services** - Files that are running as a service for the selected host. • **Tasks/Cron jobs** - Files that are configured to run as scheduled tasks along with the trigger.
5	**Show/Hide Right Panel** - Displays the following properties in the right panel: • **File Details** - Displays all properties of the selected process. It is grouped as follows: **General** - General information about the file, such as file name, entropy, size, and format. **Signature** - Provides signatory information. **Hash** - Hash type of the file (MD5, SHA1, and SHA256). **Time** - Time when the file was created, modified, or accessed. **Location** - Location of the file. **Autoruns/Services/Tasks** - Details related to autrouns, services, or tasks. • **Local Risk Details** - Displays the alerts associated with the local risk score, such as Critical, High, Medium and All. • **Hosts** -Displays the top 100 hosts based on the risk score on which the file is present.
6	Clicking a filename lets you navigate to the Files view for further analysis.
7	**Filter Files**. You can filter files by selecting the options in the Filters panel and create filters. For more information, see Filter Host Details.
8	**Settings Menu.** You can set Hosts view preferences by selecting columns from the Settings menu. For more information, see Set Hosts Preference.

Hosts View - Files Tab

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Files tab displays all files on the host including the files deleted within last 30 days. To access this tab, select a host from the **Hosts** view and click the **Files** tab. By default, it displays 100 files. To display more files, click **Load More** at the bottom of the page.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search files on host*	Search Files on Host
Threat Hunter	analyze processes	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate
Threat Hunter	filter files*	Filter Host Details
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT*, system dump, or process dump*	Performing Host Forensics

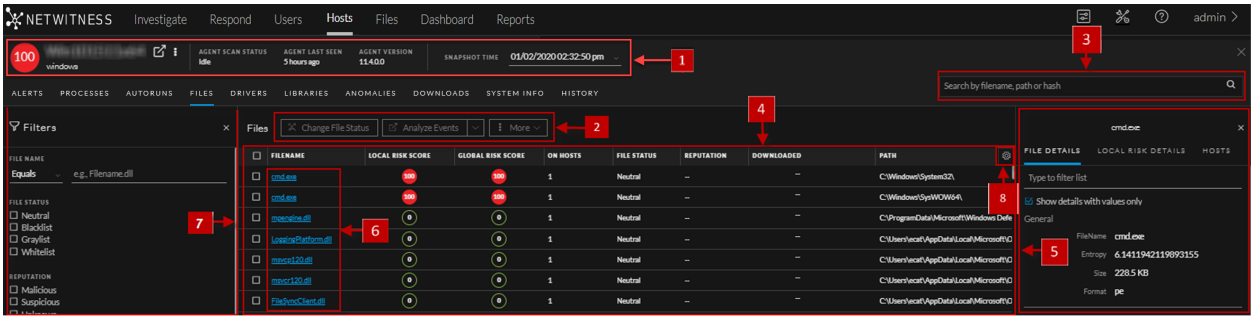
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Changing File Status or Remediate
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Files tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

More - Provides options to:

- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

Snapshot Time - Lists scanned time stamps. To view the scan history, you can select the snapshot time from the drop-down menu.

2 **Actions in the toolbar:**

Change File Status - Provides capabilities to manage suspect and legitimate files and block malicious or infected file to prevent future execution of the file on any host. For more information, see Changing File Status or Remediate.

Analyze Events - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events.

More Actions - Provides options to:

- Perform external lookups.
- Download files to server, save a local copy, and analyze files for deeper analysis.

Note: You can perform some of the above actions from the right-click context menu.

3 Search files on host. Lets you search the files on the host (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.

4 **All Files Available on Host** - Lists all files (reported as part of scan and tracking) on the host. By default, **All Files Available on Host** toggle is enabled for Windows and Mac.

5 **Details Panel** - Displays information, such as filename, local risk score, global risk score, on hosts, reputation status, file status, package details and others.

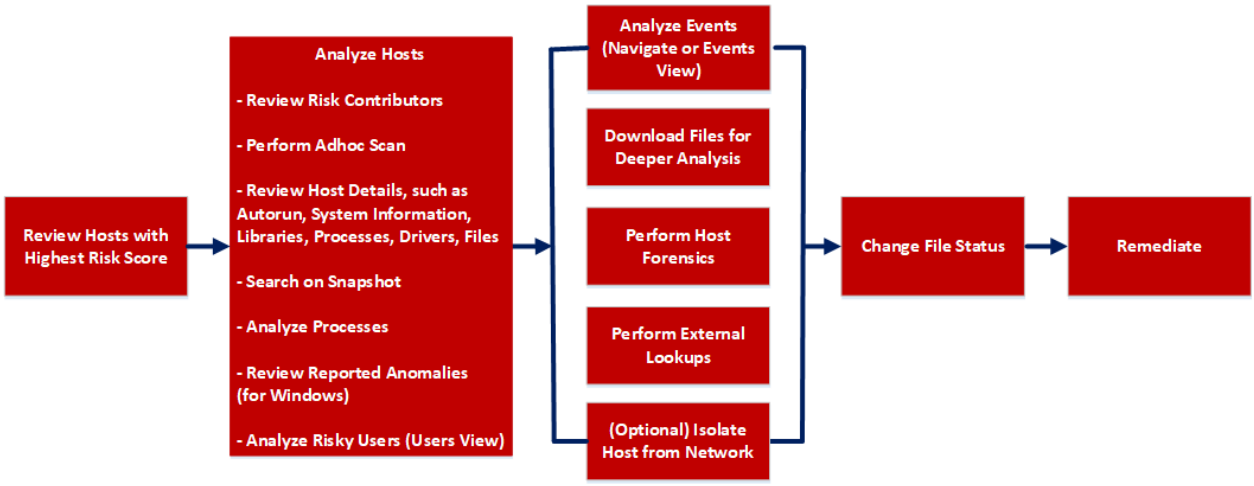
6	**Show/Hide Right Panel** - Displays the following properties in the right panel: • File Details - Displays all properties of the selected process. It is grouped as follows: **General** - General information about the file, such as file name, entropy, size, and format. **Signature** - Provides signatory information. **Hash** - Hash type of the file (MD5, SHA1, and SHA256). **Time** - Time when the file was created, modified, or accessed. **Location** - Location of the file. • Local Risk Details - Displays the alerts associated with the local risk score, such as Critical, High, Medium and All. • Hosts - Displays the top 100 hosts based on the risk score on which the file is present.
7	Clicking a filename lets you navigate to the Files view for further analysis.
8	**Filter Files**. You can filter files by selecting the options in the Filters panel and create filters. **Note:** In the **Deleted** column, a trash icon appears next to the deleted file. The **Deleted** column is not displayed if you disable **All Files Available On Host**. For more information, see Filter Host Details.
9	**Settings Menu.** You can set Hosts view preferences by selecting columns from the Settings menu. For more information, see Set Hosts Preference.

Hosts View - Drivers Tab

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Drivers tab lists the drivers running on the hosts at the time of scan. To access this tab, select a host from the **Hosts** view and click the **Drivers** tab.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search on snapshot*	Search Files on Host
Threat Hunter	analyze processes	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate
Threat Hunter	filter files*	Filter Host Details
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT*, system dump, or process dump*	Performing Host Forensics

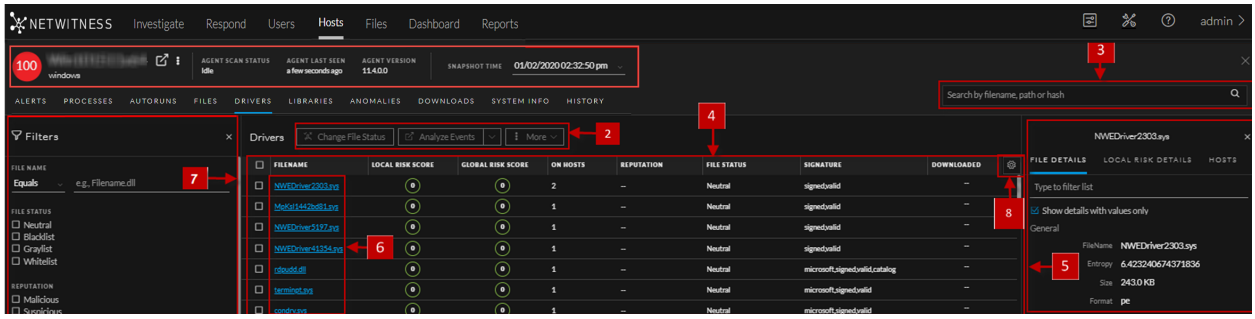
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Changing File Status or Remediate
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Drivers tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

More - Provides options to:

- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

Snapshot Time - Lists scanned time stamps. To view the scan history, you can select the snapshot time from the drop-down menu.

2 **Actions in the toolbar:**

Change File Status - Provides capabilities to manage suspect and legitimate files and block malicious or infected file to prevent future execution of the file on any host. For more information, see Changing File Status or Remediate.

Analyze Events - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events.

More - Provides options to:

- Perform external lookups.
- Download files to server, save a local copy, and analyze files for deeper analysis.

Note: You can perform some of the above actions from the right-click context menu.

3 **Search on Snapshots**. Lets you search on all snapshots (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.

4 **Details Panel** - Displays information, such as filename, local risk score, global risk score, on hosts, reputation status, file status, and others.

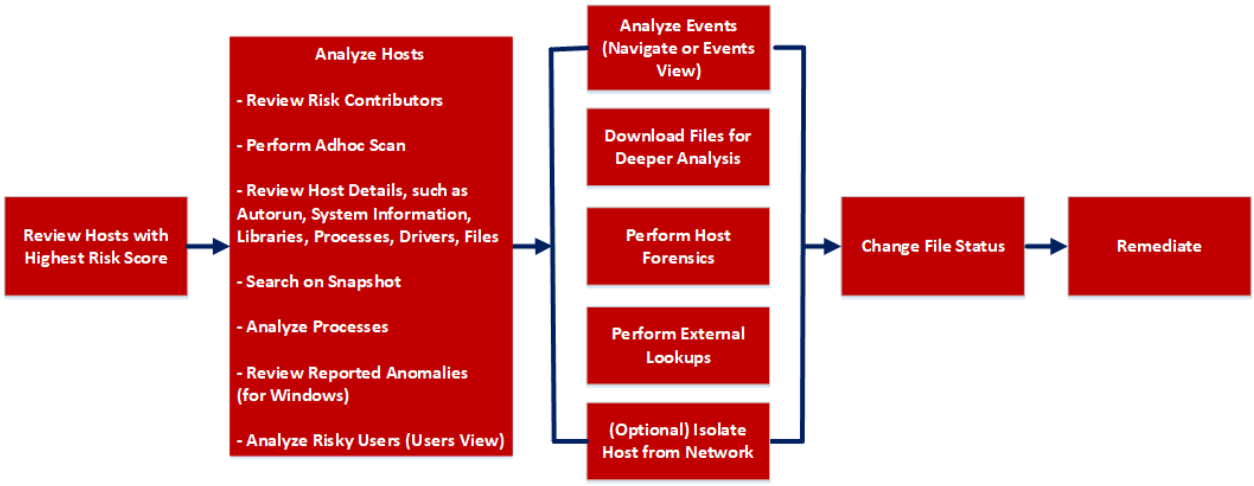
5	**Show/Hide Right Panel** - Displays the following properties in the right panel: • **File Details** - Displays all properties of the selected process. It is grouped as follows: **General** - General information about the file, such as file name, entropy, size, and format. **Signature** - Provides signatory information. **Hash** - Hash type of the file (MD5, SHA1, and SHA256). **Time** - Time when the file was created, modified, or accessed. **Location** - Location of the file. **Image** - Loaded image. • **Local Risk Details** - Displays the alerts associated with the local risk score, such as Critical, High, Medium and All. • **Hosts** - Displays the top 100 hosts based on the risk score on which the file is present.
6	Clicking a filename lets you navigate to the Files view for further analysis.
7	**Filter Files**. You can filter files by selecting the options in the Filters panel and create filters. For more information, see Filter Host Details.
8	**Settings Menu.** You can set Hosts view preferences by selecting columns from the Settings menu. For more information, see Set Hosts Preference.

Hosts View - Libraries Tab

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The Libraries tab lists the libraries loaded at the time of scan. To access this tab, select a host from the **Hosts** view and click the **Libraries** tab.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search on snapshot*	Search Files on Host
Threat Hunter	analyze processes	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate
Threat Hunter	filter files*	Filter Host Details
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT, system dump, or process dump*	Performing Host Forensics

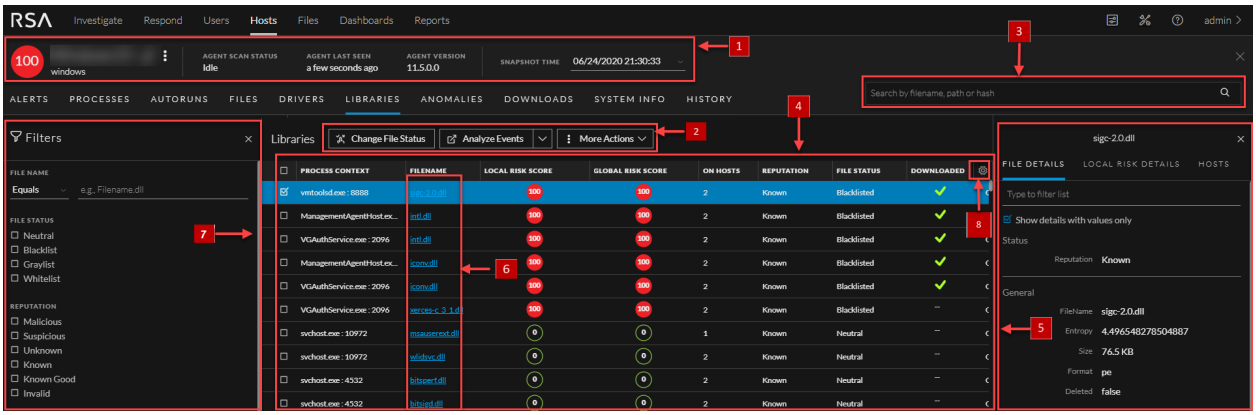
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Changing File Status or Remediate
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Libraries tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

More - Provides options to:

- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

Snapshot Time - Lists scanned time stamps. To view the scan history, you can select the snapshot time from the drop-down menu.

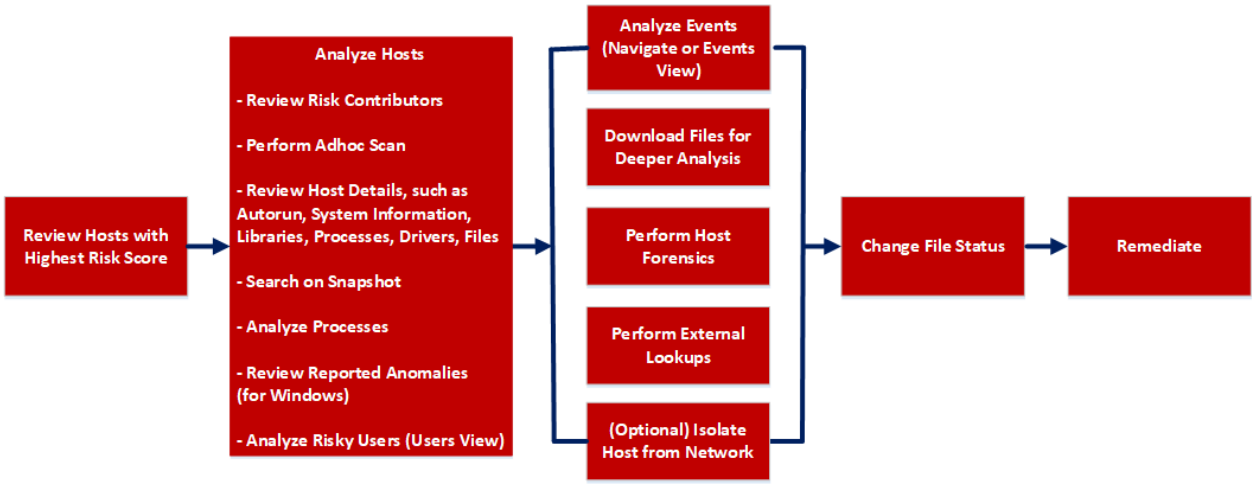
2	**Actions in the toolbar:** **Change File Status** - Provides capabilities to manage suspect and legitimate files and block malicious or infected file to prevent future execution of the file on any host. For more information, see Changing File Status or Remediate. **Analyze Events** - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events. **More Actions** - Provides options to: • Perform external lookups. • Download process dump to server. • Download files to server, save a local copy, and analyze files for deeper analysis. **Note:** You can perform some of the above actions from the right-click context menu.
3	**Search on Snapshots**. Lets you search on all snapshots (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.
4	**Details Panel** - Displays information, such as process context, filename, local risk score, global risk score, on hosts, reputation status, file status, and others.
5	**Show/Hide Right Panel** - Displays the following properties in the right panel: • **File Details** - Displays all properties of the selected process. It is grouped as follows: **General** - General information about the file, such as file name, entropy, size, and format. **Signature** - Provides signatory information. **Hash** - Hash type of the file (MD5, SHA1, and SHA256). **Time** - Time when the file was created, modified, or accessed. **Location** - Location of the file. **Process** - Details of the process, such as image size and PID. • **Local Risk Details** - Displays the alerts associated with the local risk score, such as Critical, High, Medium and All. • **Hosts** - Displays the top 100 hosts based on the risk score on which the file is present.
6	Clicking a filename lets you navigate to the Files view for further analysis.
7	**Filter Files**. You can filter files by selecting the options in the Filters panel and create filters. For more information, see Filter Host Details.
8	**Settings Menu.** You can set Hosts view preferences by selecting columns from the Settings menu. For more information, see Set Hosts Preference.

Hosts View - Anomalies Tab

Note: The information in this topic applies to NetWitness Version 11.3 and later.

The Anomalies panel provides a list of image hooks, suspicious threads, kernel hooks, and registry discrepancies running on the host. To access this tab, select a host from the **Hosts** view and click the **Anomalies** tab.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score*	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search on snapshot*	Search Files on Host
Threat Hunter	analyze processes	Investigating a Process
Threat Hunter	review reported anomalies*	Analyze Anomalies

User Role	I want to ...	Show me how
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events*	Analyzing Events
Threat Hunter	download files for deeper analysis*	Analyzing Downloaded Files
Threat Hunter	perform external lookups*	Launch an External Lookup for a File
Threat Hunter	change file status or remediate*	Changing File Status or Remediate
Threat Hunter	filter files	Filter Host Details
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT, system dump, or process dump*	Performing Host Forensics

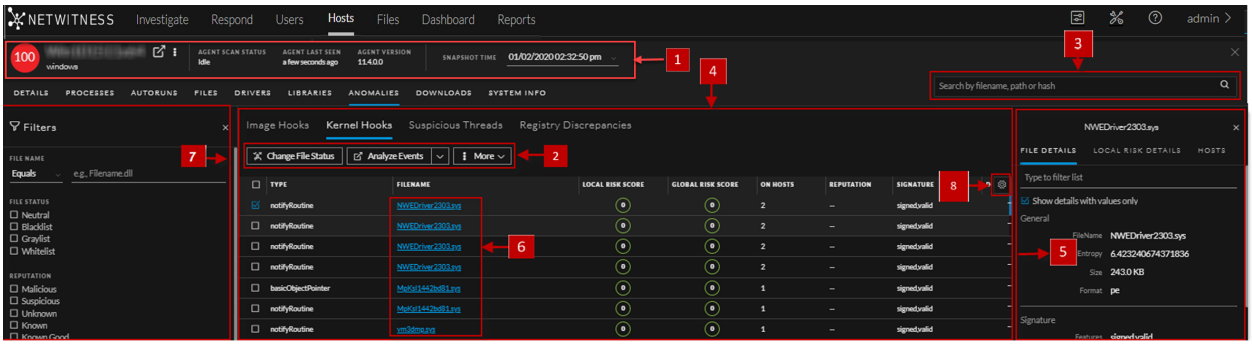
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts
- Analyzing Downloaded Files
- Changing File Status or Remediate
- Analyzing Events
- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Anomalies tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

More - Provides options to:

- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

Snapshot Time - Lists scanned time stamps. To view the scan history, you can select the snapshot time from the drop-down menu.

2	**Actions in the toolbar:** **Change File Status** - Provides capabilities to manage suspect and legitimate files and block malicious or infected file to prevent future execution of the file on any host. For more information, see Changing File Status or Remediate. **Analyze Events** - Lets you investigate a particular host, IP address, username, filename, or hash to get the entire context of the activity. For more information, see Analyzing Events. **More** - Provides options to: • Perform external lookups. • Download process dump to server. • Download files to server, save a local copy, and analyze files for deeper analysis. **Note:** You can perform some of the above actions from the right-click context menu.
3	**Search on Snapshots**. Lets you search on all snapshots (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.
4	**Details Panel** - Displays the following tabs: • Image Hooks • Kernel Hooks • Suspicious Threads • Registry Discrepancies
5	**Show/Hide Right Panel** - Displays the following properties in the right panel: • **File Details** - Displays all properties of the selected process. It is grouped as follows: **General** - General information about the file, such as file name, entropy, size, and format. **Signature** - Provides signatory information. **Hash** - Hash type of the file (MD5, SHA1, and SHA256). **Time** - Time when the file was created, modified, or accessed. **Location** - Location of the file. **Image Hooks/Kernel Hooks/Suspicious Threads/Registry Discrepancies** - Details related to image hooks, kernel hooks, suspicious threads, or registry discrepancies. • **Local Risk Details** - Displays the alerts associated with the local risk score, such as Critical, High, Medium and All. • **Hosts** - Displays the top 100 hosts based on the risk score on which the file is present.
6	Clicking a filename lets you navigate to the Files view for further analysis.
7	**Filter Files**. You can filter files by selecting the options in the Filters panel and create filters. For more information, see Filter Host Details.

8 **Settings Menu.** You can set Hosts view preferences by selecting columns from the Settings menu. For more information, see Set Hosts Preference.

Image Hooks

Image hooks found in executable image are displayed in the following columns.

Columns	Description
Type	Type of the hook . Possible values are - inline, iat, eat, or exception Handler.
Local Risk Score	Risk score of suspicious or malicious activities performed by the file on a specific host.
Global Risk Score	Aggregated score of all suspicious and malicious activities performed by the file across all hosts.
Reputation	Reputation of a file hash. The statuses are - Malicious, Suspicious, Unknown, Known, Known Good, and Invalid.
Signature	Provides signatory information.
Downloaded	Indicates the status of the downloaded file - Downloaded, Not Downloaded, and Error.
Hooked Process	Process in which hooks are placed.
Hooked Filename	Name of the file that was modified by the hook.
Hooked Symbol	Symbol in which the hook is performed.

Kernel Hooks

Hooks found on kernel objects are displayed in the following columns.

Category	Description
Type	Type of kernel object which was modified. Possible values are: objectInitializer,basicObjectPointer, majorFunction, invalidObject, fastIO, notifyRoutine, attachedDevice, device, miniPort, sdt, sysEnter, or type.idt.
Driver name	Name of the driver which placed the hooks.
Local Risk Score	Risk score of suspicious or malicious activities performed by the file on a specific host.
Global Risk Score	Aggregated score of all suspicious and malicious activities performed by the file across all hosts.

Category	Description
Reputation	Reputation of a file hash. The statuses are - Malicious, Suspicious, Unknown, Known, Known Good, and Invalid.
Signature	Provides signatory information.
Downloaded	Indicates the status of the downloaded file - Downloaded, Not Downloaded, and Error.
Object Function	Name of the object function hooked into.
Hooked File Name	Name of the file that was modified by the hook.

Suspicious Threads

Threads whose service table was hooked are displayed in the following columns.

Category	Description
Start Address	Start Address - Start address of the thread.
DLL Name	Name of the DLL.
Local Risk Score	Risk score of suspicious or malicious activities performed by the file on a specific host.
Global Risk Score	Aggregated score of all suspicious and malicious activities performed by the file across all hosts.
Reputation	Reputation of a file hash. The statuses are - Malicious, Suspicious, Unknown, Known, Known Good, and Invalid.
Process	File name and PID of the process in which thread is running.
Downloaded	Indicates the status of the downloaded file - Downloaded, Not Downloaded, and Error.
Signature	Provides signatory information.
Thread ID	ID of the running thread.
Thread Environment Block	Address of the thread environment block.

Registry Discrepancies

Configuration settings and options on Microsoft Windows operating systems that are stored are displayed in the following columns.

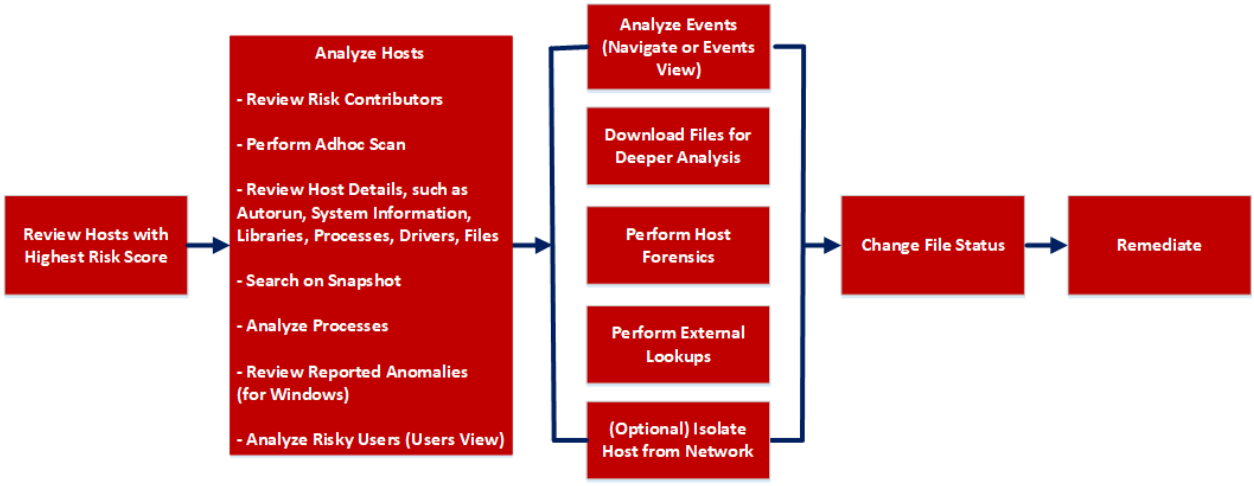
Category	Description
Hive	Name of the registry hive when possible, otherwise it displays the hive ID. Possible values are: `hkeyClassesRoot`, `hkeyCurrentUser`, `hkeyLocalMachine`, `hkeyUsers`, or `hkeyPerformanceData`.
Reason	Type of registry discrepancy. Possible values are: notFound, embeddedNull, accessDenied, parentIsHidden, or dataMismatch.
Registry Path	Registry path that is affected. The value is separated by a @ character.
Raw Type	Value type found in the low-level parsing.
Raw Data	Value data extracted from the low-level parsing.
API Type	Value type from the Win32 registry API.
API Data	Value data from the Win32 registry API.

Hosts View - Downloads Tab

Note: The information in this topic applies to NetWitness Version 11.4 and later.

The Downloads tab provides information about all downloads (MFT, files, system dump, and process dump) performed on the host. To access this tab, select a host from the **Hosts** view and click the **Downloads** tab.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts	Investigating Hosts
Threat Hunter	perform adhoc scan	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search on snapshot	Search Files on Host
Threat Hunter	analyze processes	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events	Analyzing Events
Threat Hunter	download files for deeper analysis	Analyzing Downloaded Files
Threat Hunter	perform external lookups	Launch an External Lookup for a File
Threat Hunter	change file status or remediate	Changing File Status or Remediate
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT, download files, system dump, or process dump*	Performing Host Forensics

*You can perform this task in the current view.

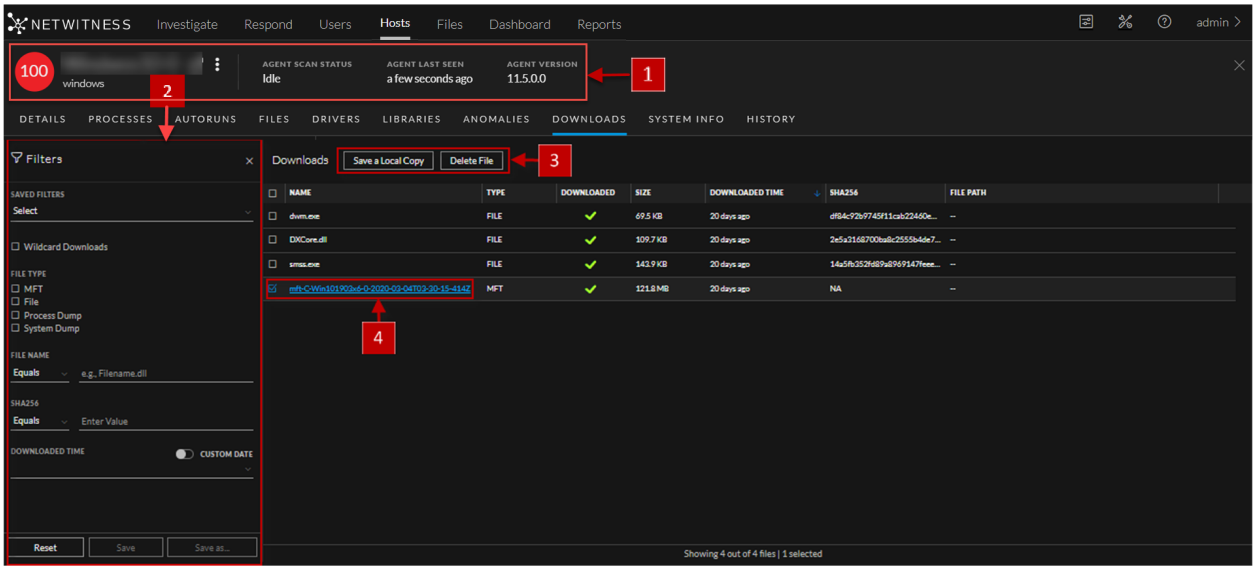
Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts

- Performing Host Forensics
- Isolating Hosts from Network

Quick Look

Below is an example of the Downloads tab:



1	**Agent and Scan Details**. You can view the following agent and scan details of the selected host: **Host name** - Name of the host. For example, `WIN-ABC`. **Risk score** - Risk score of the host. **Operating System** - Operating system on which the agent is running (Linux, Windows, or Mac). **Agent Scan Status** - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts. **Agent Last Seen** - Time when the agent last communicated with the Endpoint server. **Agent Version** - Version of the agent. For example, `11.3.0.0`. **More** - Provides options to: • Start a scan for the selected hosts. For more information, see Scan Hosts. • Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes. • Isolation host from the network. For more information, see Isolating Hosts from Network. • Download MFT to the server. For more information, see Performing Host Forensics. • Download files to the server. For more information, see Download Files Using Full Path or Wildcard. • Download System Dump to the server. For more information, see System and Process Memory Dump.
2	**Filter Files**. You can filter downloaded files by selecting the options in the Filters panel and create filters. For more information, see Performing Host Forensics.
3	**Actions in the toolbar:** **Save a Local Copy** - Lets you retrieve the downloaded MFT and save it to your local file system for further analysis. **Delete File** - Deletes the downloaded MFT from the server. For more information, see Performing Host Forensics.
4	**View MFT Details**. Click the filename to view the MFT details. For more information, see MFT Viewer.

The table displays the following information:

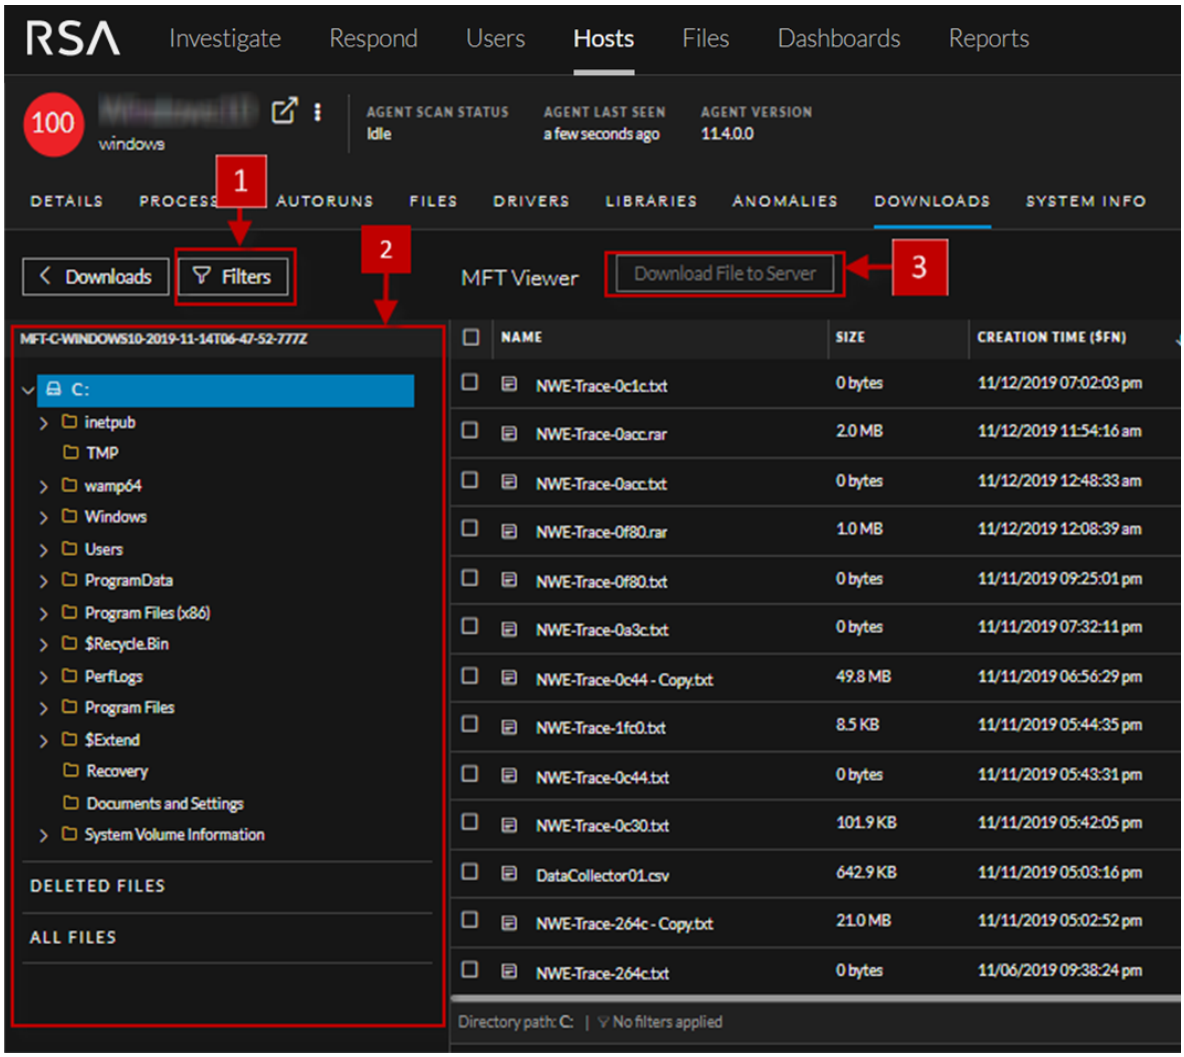
Column	Description
File Name	Name of the file that is downloaded. For example, `VGAuthService.exe`.
Type	Type of file downloaded - MFT, file, memory dump.

Column	Description
Downloaded	Status of the download:  - Download successful  - Processing the downloaded file  - Errors including download failed  - Errors downloading one or more files in the group.
Size	Size of the downloaded file.
Downloaded Time	Time when the MFT was downloaded.
SHA256	SHA256 of the file. **Note:** This is applicable only for files.

MFT Viewer

You can analyze the downloaded MFT using the MFT Viewer. For more information, see Analyze Downloaded MFT.

Below is an example of the MFT Viewer:



1 **Filter Files**. You can filter files by selecting the options in the Filters panel and create filters. For more information, see Filter MFT.

2 **Folder Details**. Lets you view the content of the MFT.

3 **Download File to Server**. Downloads files to the server.

The table displays the following information:

Column	Description
Name	Name of the file. For example, `dtf.exe`.
Size	Size of the file.
Creation Time ($FN)	File Name ($FN) creation time.
Creation Time ($SI)	Standard Information ($SI) creation time.

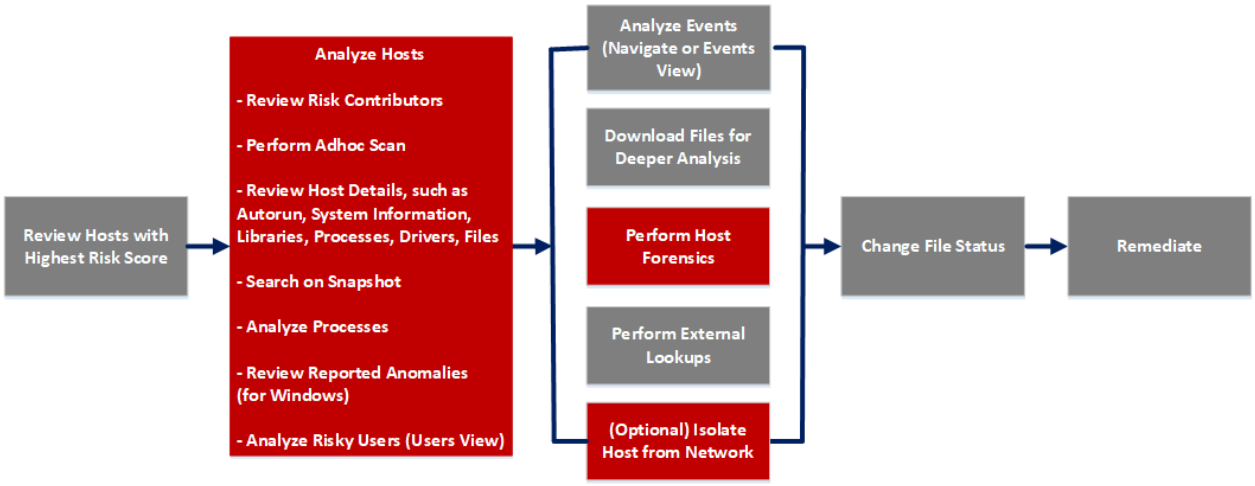
Column	Description
Modification time ($FN)	$FN modified time.
Modification time ($SI)	$SI modified time.
Access time ($FN)	$FN access time.
Access time ($SI)	$SI access time.
Update time ($FN)	$FN updated time.
Update time ($SI)	$SI updated time.
Full Path	Path of the file.
Allocated Size	File size on the disk.
Archive	Indicates if a file is archived.
Compressed	Indicates if a file is compressed.
Encrypted	Indicates if a file is encrypted.
Hidden	Indicates if a file is hidden.
Directory	Indicates if it is a directory.
Extension	Type of the file. For example, exe, pdf, txt.

Hosts View - System Information Tab

Note: The information in this topic applies to NetWitness Version 11.1 and later.

The System Information tab lists the agent system information. To access this tab, select a host from the **Hosts** view and click the **System Information** tab.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	review hosts with highest risk score	Analyze Hosts Using the Risk Score
Threat Hunter	analyze hosts*	Investigating Hosts
Threat Hunter	perform adhoc scan*	Scan Hosts
Threat Hunter	review host details	Analyze Host Details
Threat Hunter	search on snapshot*	Search Files on Host
Threat Hunter	analyze processes	Investigating a Process
Threat Hunter	review reported anomalies	Analyze Anomalies
Threat Hunter	analyze risky users	Analyzing Risky Users
Threat Hunter	analyze events	Analyzing Events
Threat Hunter	download files for deeper analysis	Analyzing Downloaded Files

User Role	I want to ...	Show me how
Threat Hunter	perform external lookups	Launch an External Lookup for a File
Threat Hunter	change file status or remediate	Changing File Status or Remediate
Threat Hunter	isolate host from network*	Isolating Hosts from Network
Threat Hunter	download MFT, system dump, or process dump*	Performing Host Forensics

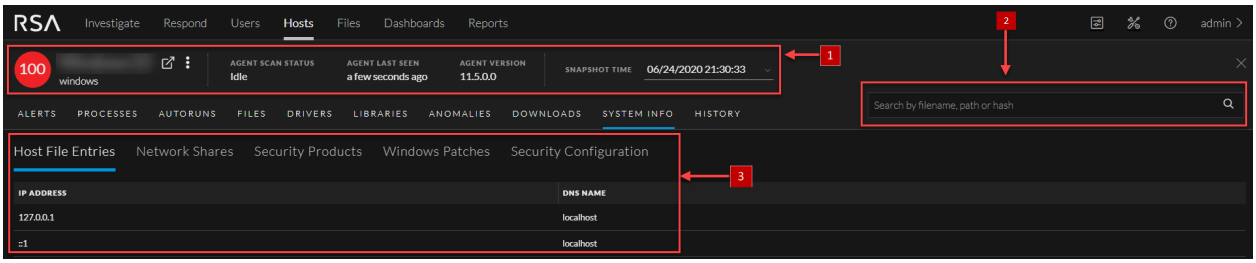
*You can perform this task in the current view.

Related Topics

- Focusing on Endpoint Analysis
- Investigating Hosts

Quick Look

Below is an example of the System Information tab:



1	**Agent and Scan Details**. You can view the following agent and scan details of the selected host: **Host name** - Name of the host. For example, `WIN-ABC`. **Risk score** - Risk score of the host. **Operating System** - Operating system on which the agent is running (Linux, Windows, or Mac). **Agent Scan Status** - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts. **Agent Last Seen** - Time when the agent last communicated with the Endpoint server. **Agent Version** - Version of the agent. For example, `11.3.0.0`. **More** - Provides options to: • Start a scan for the selected hosts. For more information, see Scan Hosts. • Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes. • Isolation host from the network. For more information, see Isolating Hosts from Network. • Download MFT to the server. For more information, see Performing Host Forensics. • Download System Dump to the server. For more information, see System and Process Memory Dump. **Snapshot Time** - Lists scanned time stamps. To view the scan history, you can select the snapshot time from the drop-down menu.
2	**Search on Snapshots**. Lets you search on all snapshots (file name, file path, and SHA-256 checksum). For more information, see Search Files on Host.
3	**System Information Panel** - See System Information Panel.

System Information Panel

The System Information panel displays the following tabs:

Tabs	Description
Host File Entries	All network redirections written in the host file. For example, IP Address - `10.10.10.3` and DNS Name - `localhost,localhost.localdomain,localhost4,localhost4.localdomain4`
Network Shares	Network name of the shared resource (for Windows only). For example, Name - `Admin$`, Description - `Remote Admin`, Path - `C:\`, Permissions - `None`, Type - `disk, special`, Max Users - `4294967295`, Current Users - `0`.
Security Products	Installed security products (for Windows only). For example, Display Name - `Windows Defender`, Instance - `D68DDC3A-831F-4FAE-9E44-DA132C1ACF46`, Features - `Enabled`, Type - `antiVirus`.

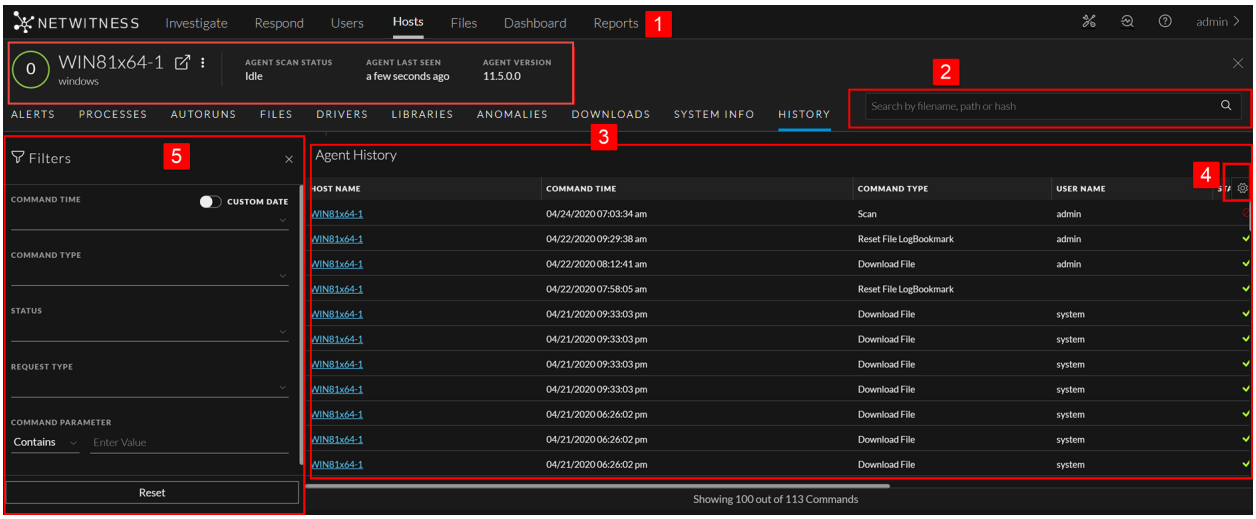
Tabs	Description
Windows Patches	List of patches applied by Windows update (for Windows only). For example, `KB2959936`.
Security Configuration	Security configuration details on the host. For example, firewall disabled or enabled, smart screen filter disabled or enabled. This field is only applicable for Windows and Mac.

Hosts View - Agent History Tab

The Agent History tab lists the commands along with the respective status and additional details.

Quick Look

Below is an example of the Agent History tab:



1 **Agent and Scan Details**. You can view the following agent and scan details of the selected host:

Host name - Name of the host. For example, `WIN-ABC`.

Risk score - Risk score of the host.

Operating System - Operating system on which the agent is running (Linux, Windows, or Mac).

Agent Scan Status - Current status of the scan - Idle, Scanning, Starting Scan, or Stopping Scan. For more information, see Scan Hosts.

Agent Last Seen - Time when the agent last communicated with the Endpoint server.

Agent Version - Version of the agent. For example, `11.3.0.0`.

More - Provides options to:

- Start a scan for the selected hosts. For more information, see Scan Hosts.
- Extracts host attributes and endpoint data to a JSON file of the selected snapshot. For more information, see Export Host Attributes.
- Isolation host from the network. For more information, see Isolating Hosts from Network.
- Download MFT to the server. For more information, see Performing Host Forensics.
- Download System Dump to the server. For more information, see System and Process Memory Dump.

2 **Search files on host**. Lets you search the files on the host (file name, file path, and SHA-256 checksum).

3 **Details Panel**- Displays information, such as:

- **Command Time** - Command issued time.
- **Command Type** - Type of the command (Identity, scan, stop scan, download file, MFT, process dump, system dump, start isolation, update isolation exclusion list, stop isolation, reset file logbookmark, and download multiple files, agent upgrade, and uninstall agent) issued.
- **User Name** - User who issued the command. For example, Analyst, System.
- **Status** - Status (success, pending, expired, failed, or cancelled) of the command issued.

Note: If the command's status is expired, it means that the agent is unable to process the command even after five retries.

- **Command Parameter** - Parameters associated with the command. For example, Command parameter for command type Download File is `path = C:\Windows\System32\ | filename = cmd.exe | hash = 6f88fb88ffb0f1d5465c2826e5b4f523598b1b8378377c8378ffebc171bad18b`

Note: Command types such as identity, scan, stop scan, stop isolation, system dump do not contain any associated command parameters.

- **Processed Time** - Time at which the command is completed, pending, expired, failed, or cancelled.
- **Last Retrieval time** - Last time when the command is issued to the agent.
- **Total Retrieval** - The number of times the command is issued to the agent.

Note: After you upgrade to NetWitness version 11.5, the commands executed in the previous versions are displayed automatically. The fields such as last retrieval time, total retrieval, and user do not contain any values. For system generated commands, the user field value shows as system.

4 **Filter Files**. You can filter commands by selecting the options in the Filters panel. For more information, see Filter Host Details.

5 **Settings Menu.** You can set History view preferences by selecting columns from the Settings menu.

NetWitness® Platform XDR

Version 12.0

Alerting with ESA Correlation Rules User Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Getting Started with ESA **11**

How ESA Generates Alerts 12

Data Source Configuration Changes 12

An Endpoint Risk Scoring Rules Bundle is available in NetWitness Platform 12

Best Practices 13

Understand Event Stream Analysis Rule Types 13

Best Practices for Writing Rules 14

Best Practices for Working with RSA Live Rules 15

Best Practices for Deploying Rules 15

Best Practices for System Health 16

Troubleshoot ESA 17

Troubleshoot ESA Correlation Services 17

Troubleshoot RSA Live Rules for ESA 19

Troubleshoot ESA Rules 20

SMTP Notification Error Example 27

Integration-Server SMTP Notification Error Example 27

Example ESA Correlation Server Warning Message for Missing Meta Keys 27

Multi-Valued Warning Message Example 28

Single Value Warning Message Example 28

Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event_source_id 28

Steps to Troubleshoot Memory Issues with an ESA Service Offline 30

Step 1: Verify that your Host Is Running 30

Step 2: View Detailed Statistics in Health & Wellness 30

Step 3: Bring up your ESA Services 35

Step 4: Check the Alerts and Events Volume 35

View Alert Summaries 35

View Events Matched 36

Step 5: Disable and Repair the Rule that Caused Issues 37

Disable Rules 37

Edit Rules 37

Deploy Rules 37

Verify that the Rules are Enabled 37

(Optional) Check the ESA Correlation Log Files for More Information 38

ESA Rule Troubleshooting with Nw-Shell 38

Find Your Engine Name for Nw-Shell 38

Connect to an ESA Correlation Server 39
View the Contents of a Named Window 39
See the Method Input and Output 40
Obtain Correlation Server Metrics for ESA Rule Deployment Troubleshooting Using Nw-Shell 41
View Memory Metrics for Rules 44
Prerequisites 44
View Health Statistics and Trends for ESA Correlation in New Health & Wellness 44
View Memory Metrics for an ESA Correlation Service in Health & Wellness 46
View Memory Metrics for an ESA Correlation Service and its ESA Rules 47
How ESA Handles Sensitive Data 50
How ESA Treats Sensitive Data from Core Services 50
Advanced EPL Rule 50
Enrichment Source 51
How to Remove Sensitive Meta Keys Globally from All Alerts 51
ESA Rule Types 53
Sample Rules 53
Endpoint Risk Scoring Rules Bundle 53
Trial Rules Mode 54
ESA Permissions 54
Practice with Sample Rules 55
Rule Library 56
Practice with Sample Rules 57
Working with Trial Rules 58
Deploy Rules as Trial Rules 58
Add Rules to the Rule Library 60
Download Configurable RSA Live ESA Rules 61
Prerequisites 61
Download RSA Live ESA Rules 61
Customize an RSA Live ESA Rule 64
Prerequisites 65
Configure Parameters for an RSA Live ESA Rule 65
Add a Rule Builder Rule 66
Step 1. Name and Describe the Rule 66
Prerequisites 66
Name and Describe a Rule 66
Step 2. Build a Rule Statement 68
Example 68
Prerequisites 68
Build a Rule Statement 69

To Add a Whitelist 71
To Add a Blacklist 71
Example: Blacklist 72
Example: Strict Pattern Matching and Using the Is Not Null Operator 73
Example Results 76
Example: Grouping the Rule Results 77
Example: Working with Numeric Operators 78
Step 3. Add Conditions to a Rule Statement 80
Example 80
Add Conditions to a Rule Statement 80
Example 81
Validate an ESA Rule 82
Working with Rules 86
Edit, Duplicate or Delete a Rule 86
Edit a Rule 86
Duplicate a Rule 86
Delete a Rule 86
Filter or Search for Rules 87
Prerequisites 87
Filter Rules 87
Search for Rules 88
Import or Export Rules 88
Import ESA Rules 88
Export ESA Rules 89
Choose How to be Notified of Alerts 90
Notification Methods 90
Email Notifications 91
Syslog 91
Script Alerter 91
Add Notification Method to a Rule 92
Prerequisites 92
Add a Notification Method to a Rule 92
Add a Data Enrichment Source 94
Example Rule with Enrichments 94
Enrichment Sources 96
Configure a Context Hub List as an Enrichment Source 97
Prerequisites 97
Configure a Context Hub List as an Enrichment Source 97
Configure an In-Memory Table as an Enrichment Source 99
Prerequisites 100

Configure an Ad hoc In-Memory Table 100
Add a Recurring In-Memory Table 103
Add an Enrichment to a Rule 103
Deploy Rules to Run on ESA 105
How an ESA Rule Deployment Works 105
ESA Rule Deployment Steps 106
Step 1. Add an ESA Rule Deployment 106
Prerequisites 106
Step 2. Add an ESA Service 108
Step 3. Add Data Sources 110
Step 4. Add and Deploy Rules 114
(Optional) Add a Data Source Filter 116
Deploy the Endpoint Risk Scoring Rules Bundle 119
Additional ESA Rule Deployment Procedures 120
Replace an ESA Service in an ESA Rule Deployment 120
Remove an ESA Service from an ESA Rule Deployment 120
Add an ESA Service to an ESA Rule Deployment 121
Edit a Data Source in an ESA Rule Deployment 121
Add or Remove a Data Source 123
Remove a Data Source from an ESA Rule Deployment 123
Add a Data Source to an ESA Rule Deployment 123
(Optional) Add or Remove a Data Source Filter 124
Remove a Data Source Filter from an ESA Rule Deployment 124
Add a Data Source Filter to an ESA Rule Deployment 124
Adjust Maximum Sessions for the Data Source Filter 125
Edit or Delete a Rule in a Deployment 125
Edit a Rule 125
Delete a Rule 126
Stop Deployment 126
Edit the ESA Rule Deployment Name or Delete a Deployment 126
Edit the ESA Rule Deployment Name 127
Delete an ESA Rule Deployment 127
Show Updates to an ESA Rule Deployment 128
Fast Deployment and Full Deployment of ESA Rules 129
View ESA Stats and Alerts 132
View Stats for an ESA Service 132
View ESA Stats 132
Enable or Disable Rules 133
Refresh the Statistics 133
View a Summary of Alerts 134

Add an Advanced EPL Rule **137**
Prerequisites 137
Add an Advanced EPL Rule 137
Validate an Advanced EPL Rule 140
Event Processing Language (EPL) 144
ESA Annotations 146
@RSAContext Annotation (11.5 and later) 146
Prerequisites 146
Single-Column Context Hub Lists 146
Multi-Column Context Hub Lists 147
Automatic Context Hub List Updates 149
Single Column Context Hub List Update Example 149
Multi-Column Context Hub List Update Example 149
@RSAAlert Annotation 149
@RSAPersist Annotation 151
@UsesEnrichment (10.6.1.1 and later) 151
@Name 151
@Audit Annotation 152
Example Advanced EPL Rules 153
Example #1: 153
EPL #1: 153
EPL #2: 154
Example #2: 154
EPL #3: 154
EPL #4: Using NamedWindows and match recognize 155
EPL #5: Using Every @RSAAlert(identifiers={"user_src"}) 155
Example #3: 156
EPL #6: @RSAAlert(identifiers={"ip_src"}) 156
EPL #7: @RSAAlert(identifiers={"ip_src"}) 157
Example #4: 157
EPL #8: using time_batch 157
Example #5: 158
EPL #9: using timer:interval 158
Example #6: 158
EPL #10: using timer and Lockout 158
Example #7: 159
EPL #11: @RSAAlert(oneInSeconds=0) 159
Configure an In-Memory Table Using an EPL Query **161**
Workflow 161
Prerequisites 162

Procedure 162
Example 163
Step 1: Create the Enrichment 163
Step 2: Create Your Rule 165
Rule Statement 165
Rule Logic with Enrichment Added 166
ESA Alert References 167
Rules Tab 168
What do you want to do? 168
Related Topics 168
Quick Look 169
Rules Tab Options Panel 170
What do you want to do? 170
Related Topics 170
Quick Look 170
Rules Section 170
Deployments Section 170
Rule Library Panel 172
What do you want to do? 172
Related Topics 172
Quick Look 172
Rule Library Toolbar 173
Rule Library List 173
Rule Builder Tab 175
What do you want to do? 175
Related Topics 175
Quick Look 175
Conditions Section 178
Notifications Section 180
Enrichments Section 180
Debug Option 181
Test Rule Section 181
Syntax 183
Build a Statement Dialog 184
What do you want to do? 184
Related Topics 184
Quick Look 184
Advanced EPL Rule Tab 189
What do you want to do? 189
Related Topics 189

Quick Look 189
Notifications Section 192
Enrichments Section 193
Test Rule Section 193
Syntax 195
Rule Syntax Dialog 196
Quick Look 196
Deployment Panel 198
What do you want to do? 198
Related Topics 198
Quick Look 199
ESA Services Section 199
Data Sources Section 200
Position Tracking Information 201
Editing Position Tracking Information 201
Use Case Scenario 203
Importing Position Tracking Information 204
Deployment Options 205
ESA Rules Section 206
Data Source Filter (Optional) Section 207
Create Data Source Filter Dialog - Simple 208
Create Data Source Filter - Advanced 209
Deploy ESA Services Dialog 211
What do you want to do? 211
Related Topics 211
Quick Look 211
Deploy ESA Rules Dialog 213
What do you want to do? 213
Related Topics 213
Quick Look 213
Updates to the Deployment Dialog 215
What do you want to do? 215
Related Topics 215
Quick Look 215
Services Tab 217
What do you want to do? 217
Related Topics 217
Quick Look 217
ESA Services Panel 218
General Stats Panel 218

Deployed Rule Stats Panel 219
Settings Tab 221
What do you want to do? 221
Related Topics 221
Quick Look 221
Meta Key References 222
Viewing the List of Meta Entities 222
Enabling Meta Entity in the ESA Correlation Server 222
Meta Entity Usage Example 223
Building Rules with Custom Meta Entities 224
Prerequisites 224
Enrichment Sources 224

Getting Started with ESA

This topic covers quick start topics for NetWitness Event Stream Analysis (ESA) to help you get started in using ESA. The following topics are designed to assist you in working with ESA Correlation Rules.

- Best Practices helps you to understand how to best set up, deploy, and create rules.
- Troubleshoot ESA helps you to troubleshoot different aspects of ESA, including rule writing and deployment.
- View Memory Metrics for Rules helps you to work with memory metrics to understand memory usage for ESA services.

In NetWitness version 11.5 and later, There are only two services that can run on an ESA host:

- **ESA Correlation (ESA Correlation rules):** Creates alerts from ESA rules.
- **Contexthub Server (Context Hub):** Runs only on an ESA primary host. Contexthub Server provides enrichment lookup capability in the Respond and Investigate views. For information, see the *Context Hub Configuration Guide*.

Note: The Event Stream Analytics Server (ESA Analytics) service is not supported in NetWitness Platform version 11.5 and later.

The first service is the ESA Correlation service that creates alerts from ESA rules, also known as ESA Correlation Rules, which you create manually or download from Live.

In NetWitness 11.3 and later, the ESA Correlation service replaces the Event Stream Analysis service and is also known as ESA Correlation Server. The ESA Correlation service provides the same services as the Event Stream Analysis service with the added benefit of enabling you to specify different data sources for your ESA correlation rules. Like the Event Stream Analysis service, the ESA Correlation service installs on the ESA Primary and ESA Secondary host types.

The second service is the Contexthub Server service, which provides enrichment lookup capabilities in the Respond and Investigate views. It runs only on an ESA Primary host. For information, see the *Context Hub Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

IMPORTANT: The NetWitness server (Admin server), ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

How ESA Generates Alerts

The ESA Correlation service runs rules that specify criteria for problem behavior or threatening events in your network. When ESA detects a threat that matches rule criteria, it generates an alert.

To generate alerts, ESA performs the following functions:

1. Gathers data
2. Runs ESA rules against the data
3. Captures events that meet rule criteria
4. Generates alerts for those captured events

Data Source Configuration Changes

In NetWitness version 11.3 and later, the ESA Correlation service enables you to specify different data sources for different sets of rules. Instead of adding data sources, such as Concentrators, to the entire ESA Correlation service, you can specify different data sources for each ESA rule deployment. An ESA rule deployment includes an ESA Correlation service with its associated data sources and a set of ESA rules. For example, you may want to use Concentrators with HTTP packet data in one deployment and Concentrators with HTTP log data in another deployment. For more detailed information, see Deploy Rules to Run on ESA.

An Endpoint Risk Scoring Rules Bundle is available in NetWitness Platform

An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. Endpoint risk scoring rules only apply to NetWitness Endpoint. You can add the Endpoint Risk Scoring Rules Bundle to an ESA rule deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) in the ESA Rule Deployment.

The ESA Correlation service can process endpoint risk scoring rules, which generate alerts that are used in risk scoring calculations to identify suspicious files and hosts. To turn on risk scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. For instructions, see "Deploy Endpoint Risk Scoring Rules on ESA" in the *ESA Configuration Guide*. For complete information on configuring NetWitness Endpoint, see the *NetWitness Endpoint Configuration Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Best Practices

Best practices provide guidelines to help you write and manage rules, deploy rules, and maintain system health for your ESA services.

Understand Event Stream Analysis Rule Types

The ESA Correlation service provides advanced stream analytics such as correlation and complex event processing at high throughputs and low latency. It is capable of processing large volumes of disparate event data from Concentrators. However, when working with Event Stream Analysis, you should be aware of the factors that affect resource usage in order to create effective rules.

Each event that is received by ESA is evaluated to determine if it may trigger a rule. There are three types of rules that can be deployed in order to determine what the ESA engine should do with the incoming event. Each of these rule types have different impacts on system resource utilization. All three rule types may be created via the Rule Builder, Advanced Event Processing Language (EPL) rules, or downloaded via RSA Live. The table below lists the rule type and the impact this rule may have on system resources.

Rule Type	Description
Simple Filter Rule	This rule has no correlation to other events. At ingestion time, this rule is evaluated against a set of conditions, and if those conditions are met an alert is generated. If no conditions match, the event is quickly released by the engine to free up memory usage. These rules do not take up memory since the events are not retained beyond the initial evaluation. The memory resource usage does not increase as more simple filter rules are deployed. However, if the filter condition is too generic, it is possible that this rule can generate too many alerts, which will strain the system resources for the storage and retrieval of these alerts. For example, you might write a rule to generate an alert when HTTP network activity arrives over a non-standard HTTP port.
Event Window Rule	This rule evaluates a set of events over a time period for specific conditions. At ingestion time, the rule is evaluated against a set of conditions. If those conditions are met, the event is retained in memory for a specific amount of time. After the specified time passes, the events are removed from the time window if the number of events collected does not meet the threshold to trigger an alert. The memory consumption of such rules is highly dependent on the incoming event rate (traffic), the amount of data per event, and the time length specified in the event window. Each matching event is retained in memory until the time window has passed, so the longer the time window, the greater the potential volume. For example, you might write a rule that generates an alert if a user has five failed login attempts within a ten minute time frame.

Rule Type	Description
Followed By Rule	This rule evaluates a chain of incoming events to determine if the sequence of events matches a particular condition. At ingestion time, the rule is evaluated against a set of conditions. If the conditions are met, one of two actions occurs: • If this is the first event of the sequence, a new event thread is started, and the event is retained as the head of the sequence. • If the event belongs to an existing event thread, it is added to that sequence. In both cases, the event is retained in memory. The amount of resource usage is particularly sensitive to the customer environment for this type of rule. If the filter condition generates many event threads, resources are consumed for each new thread (in addition to the event). Additionally, if the end of the event thread is never met (that is, an alert is never generated), then the entire event is saved in memory indefinitely. For example, you might write a rule to generate an alert when a user fails to log in to a server, then performs a successful login, and then creates a new account.

Note: ESA sends alerts to NetWitness Respond for processing and the alerts are eventually stored in a database. If your rule creates too many alerts, it can slow down another part of the system.

When writing and deploying rules, you should be aware that rule memory usage and alert generation consume system resources. The sections below are designed to help you keep your usage at a healthy level and monitor for problems if systems are becoming overloaded.

Best Practices for Writing Rules

These are general guidelines for writing rules.

- **Create alerts for actionable events**. The purpose of an alert should be to notify you of an event that requires immediate and specific action. For events that do not require action, or only require you to have awareness of the event, you can create a report.
- **Validate your ESA rules within the Rule Builder or Advanced EPL Rule Builder before you deploy them**. To prevent errors in your rules and confirm that they generate the expected alerts, you can test the rule logic with JSON data within the rule builders. This capability is available in NetWitness Platform version 11.5 and later. For more information, see Validate an ESA Rule and Validate an Advanced EPL Rule .
- **Configure new rules as trial rules so you can observe how they react in your environment**. If you deploy new rules as trial rules, they will be disabled if the configured memory threshold is exceeded.

- **Configure Alert notifications only after your rule testing and tuning is complete**. This can help ensure you do not get flooded with notifications if a rule behaves differently than you expect.
- **Rules need to be specific so that you limit resource usage**. Use the following guidelines to limit usage:
  - Make the filters on the rule exclude all but the necessary events for the rule to fire accurately.
  - Make the size of your windows (window time for correlation) as small as possible.
  - Limit the events that you include in the window: For example, if you only want to see IDS events, ensure that you only include those events in your time window.
- **Add Memory Thresholds to ESA rules that use memory**. For example, if a rule contains windows or pattern matching, configure a memory threshold for that rule. If the rule goes over the allotted memory threshold, it gets disabled individually and an error is displayed for that rule on the  **(Configure) > ESA Rules > Services** tab. This capability is available in NetWitness Platform version 11.5 and later.
- **Rules need to be tuned to an alert level that is manageable**. If you are flooded with alerts, then the purpose and of an alert is lost. For example, maybe you want to know about encrypted traffic to other countries. But, you could limit the list to countries that are known risks. This limits the volume of alerts to a level you can manage.

For more best practice information for writing ESA rules, see ESA Rule Writing Best Practices.

Best Practices for Working with RSA Live Rules

These are guidelines for RSA Live Rules.

- **Deploy RSA Live rules in small batches**. Not every rule is suited to every environment. The best way to ensure your RSA Live rules are successful is to deploy them in small batches so you can test them in your environment. If you deploy small batches, it's much easier to tell if a particular rule has an issue.
- **Read the rule descriptions provided with RSA Live rules**. ESA rules are not “one size fits all.” Not all rules will work in your environment. The rule descriptions tell you which parameters you will need to modify to successfully deploy a rule in your environment.
- **Set your parameters**. RSA Live rules have parameters that need to be modified. If you do not modify your parameters, the rule may not work or it may exhaust your memory.
- **Deploy new rules as trial rules so you can observe how they react in your environment.** If you deploy new rules as trial rules, they will be disabled if the configured memory threshold is exceeded. For more details, see Working with Trial Rules.

Best Practices for Deploying Rules

These are general guidelines for deploying rules.

- **Deploy rules in small batches so you can observe how they react in your environment.** Not all environments are the same, and a rule will need to be tuned for memory usage, alert volume, and effective detection of events.
- **Test rules before you configure alert notifications**. Configure Alert notifications only after your rule testing and tuning is complete. This can help ensure you do not get flooded with alerts if a rule behaves differently than you expect.
- **Monitor system health as a part of your deployment process**. When you deploy rules, monitor your system's health as a part of your deployment process. You can view total memory usage for your ESA in the Health and Wellness tab. For more information, see "View Detailed Statistics in Health and Wellness" in Troubleshoot ESA.

Best Practices for System Health

These are general guidelines for system health.

- **Set up new rules as trial rules**. A common issue is that new rules may cause memory issues. To prevent this, you can set up new rules as trial rules. If the configured memory threshold is met, all trial rules are disabled to prevent the system from running out of memory. For more information about trial rules, see Working with Trial Rules.
- **Set up thresholds in Health & Wellness to alert you if memory usage is too high**. There are metrics in NetWitness Health & Wellness that track memory usage. You can set up alerts and notifications to send you an email if those thresholds are crossed. For more information about the memory statistics you can view, see View Memory Metrics for Rules.
- **Monitor memory metrics for each rule in Health & Wellness**. For each rule, you can view the estimated memory usage in Health & Wellness. You can use this information to ensure that rules do not use too much memory. For more information about the memory statistics you can view, see View Memory Metrics for Rules.

Troubleshoot ESA

This section describes common issues that may occur while using ESA, and it suggests common solutions to these problems.

Troubleshoot ESA Correlation Services

Problem	Possible Causes	Solutions
On the NetWitness Platform Dashboard, the ESA service appears in red to indicate it is offline. In the **(Configure) > ESA Rules** view, the following message appears: "The Service is either offline or not reachable."	Several	When an ESA Correlation service is offline, there are many possible causes. However, a common issue is that you have created a rule that uses excessive memory and causes the ESA service to fail. To troubleshoot this problem, see Steps to Troubleshoot Memory Issues with an ESA Service Offline. Other common causes might be that your firewall is blocking the connection between the ESA and NetWitness Platform, or the ESA Correlation service machine may be down. To bring up ESA Services: Go to **(Admin) > Services**, select your ESA service, and then select **> Start**. If your ESA service is stopping and restarting in a loop, you may need to call Customer Support to get the services to start.

Problem	Possible Causes	Solutions
On the NetWitness Platform Dashboard, the ESA service appears in red to indicate it is offline. In the **(Configure) > ESA Rules** view, the following message appears: "The Service is either offline or not reachable.	Configuration issues	If your system has been recently upgraded, you may have made a configuration error. Go to **(Admin) > Services**, select your ESA service and then select **> Edit**. In the Edit Service dialog, click **Test Connection**. If the connection fails, you likely have a configuration error. Attempt to fix your configuration error and try again.
Suddenly, an ESA Correlation service is completely unresponsive and appears to have crashed. The ESA Correlation Log file shows a Too Many Open Files error.	Connectivity issues	The most likely cause for this error is a connectivity issue between the ESA Correlation service and a data source used in an ESA rule deployment, such as a Concentrator or Decoder. 1. Restart the ESA Correlation service. 2. Go to **(Admin) > Services**, select your ESA service and then select **> Restart**. 3. Check the connectivity from ESA to the data source. Look for connection errors in the ESA Correlation logs. You can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`. For example: `Error: com.rsa.netwitness.streams. RecordStreamException: admin@<ip address>:56005:: com.rsa.netwitness.streams.RecordStreamException: connect::com.rsa.asoc.transport.nw.session. NextgenException: Failed to connect to <ip address>:56005` If there are network connectivity issues, fix the issues and then restart the ESA Correlation service again to see if it fixes the problem. If you have a data source with intermittent connectivity, you should remove it from the ESA rule deployment.

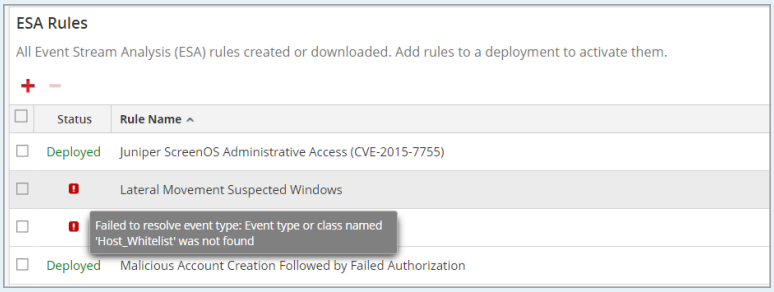
Troubleshoot RSA Live Rules for ESA

Problem	Possible Causes	Solutions
I imported a group of rules from RSA Live, and now my ESA service is crashing. Why?	You may not have configured the parameters for the RSA Live rule to tune it for your environment.	Each rule in RSA Live has a description that includes the parameters you must configure and prerequisites for your environment. Review this description to see if the rule is appropriate for your environment. To ensure that you deploy rules safely in your environment, configure new rules as trial rules to test them in your environment. Trial rules add a safeguard for testing new rules. For details on this, see Deploy Rules as Trial Rules.

I imported a group of rules from RSA Live, and while the rules deployed without errors, they were later disabled.	Not all RSA Live rules are meant for every environment. You may not have the correct meta in your ESA for the rule to run.	You can verify that a rule was disabled by going to  **(Configure) > ESA Rules > Services > Deployed Rule Stats**. If the rule is disabled, the green icon does not display next to the rule.
		If a rule deployed correctly but was disabled, check the logs for exceptions related to the rule. Specifically, check to see if the rules were disabled due to missing meta. To do this, go to the ESA Correlation logs. You can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`.
		Then, search for a message similar to the following:
		`"Property named '<meta_name>' is not valid in any stream"`
		For example, you might see:
		`Failed to validate filter expression '(medium=1 and streams=2 or medium=3...(238 chars)': Property named 'tcp_flags_seen' is not valid in any stream`
		If a similar message displays, you may need to add a custom meta key to the Log Decoder or Concentrator. To do this, follow these instructions: "Create Custom Meta Keys Using Custom Feed" in the *Decoder and Log Decoder Configuration Guide*.

Troubleshoot ESA Rules

Problem	Possible Causes	Solutions
I have an ESA rule that is not getting deployed and is not creating alerts.	A meta key that the rule uses is a string array type, but it shows as a string type on ESA.	Check to see if any string array meta keys that the rule uses are configured as string array types on ESA. Go to  **(Configure) > ESA Rules > Settings** tab (Meta Key References). • If it shows `string[]`, it is configured as a string array type on ESA. This is fine. • If it shows `string` without the brackets, it is configured as a string type and you need to fix it on ESA. In the ESA Correlation service Explore view, go to `correlation/stream`. Add string array meta keys to the **multi-valued** list to allow them to be used as an array in ESA rules. Go back to the Meta Key References and click the refresh icon (). Verify that the meta keys with a string array type show a value of `.string[]`. For additional details, see "Configure Meta Keys as Arrays in ESA Correlation Rules" in the *ESA Configuration Guide*.

Problem	Possible Causes	Solutions
I created a rule, and I checked the syntax. The rule looked fine. When I went to deploy the rule, I got an error. Why?	You may not have the correct meta to deploy the rule.	Check the Meta Key References. You may not have the correct meta to deploy the rule. Check the ESA Correlation service log files to see which meta keys are missing: `/var/log/netwitness/correlation-server/correlation-server.log` In NetWitness version 11.3.0.2 and later, you can check the ESA rule status in the ESA rule deployment (Go to  **(Configure) > ESA Rules > Rules** tab. In the options panel on the left, select a deployment and look in the **ESA Rules** section). If a disabled rule has an error message, it shows in the Status field. Hover over the rule to view the error message tooltip.  In the above example, the error message shows: "Failed to resolve event type: Event type or class named "Host_Whitelist" was not found." In this case, a Context Hub list called "Host_Whitelist" that is used by the rule is not available. For more information on context hub lists, see the *Context Hub Configuration Guide*. For more information, see the ESA Rules Section section of the Deployment Panel reference.
I set up notifications for a rule, but we are not receiving them. The `correlation-server.log` file does not show any errors. Why?	Correlation-server successfully sent the notification messages to integration-server, but when integration-server tried to send the notifications to their destination, it failed.	When troubleshooting notifications, check both the ESA Correlation service log files (`/var/log/netwitness/correlation-server/correlation-server.log`) AND the Integration-Server log files on the NetWitness Server (`/var/log/netwitness/integration-server/integration-server.log`). For an example, see Integration-Server SMTP Notification Error Example. **Note:** For any notification-related troubleshooting, check the `integration-server log` file in addition to the log file of the service creating the notification.

Problem	Possible Causes	Solutions
I created a rule with an enrichment, added an SMTP notification, and deployed my rule. We are not receiving SMTP notifications. Why?	You do not have a template that met the criteria to parse the events.	Check the ESA Correlation service log files to see if the SMTP notification failed: `/var/log/netwitness/correlation-server/correlation-server.log`. For more details on the notification error, check the Integration-Server log file on the NetWitness Server (also known as Node 0, Admin server, or NWServer): `/var/log/netwitness/integration-server/integration-server.log`. If you use an ESA rule that has an enrichment, such as a Context Hub list, you must create a custom template. You can duplicate a default template and adjust it for your enrichment. See SMTP Notification Error Example below for a notification error example. For information on creating a custom ESA template, see "Define a Template for ESA Alert Notifications" in the *System Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
I created a custom rule (via the Rule Builder or Advanced EPL), and my rule is not firing. Why?	You may have connectivity issues.	Check the Offered Rate statistic on the **(Configure) > ESA Rules > Services** tab. Select the ESA service and then look at the statistics on the tab for the Deployment. If the Offered Rate is zero, then the ESA service is not receiving data from Concentrators. Check the ESA Correlation log files for connectivity issues: `/var/log/netwitness/correlation-server/correlation-server.log`. If the offered rate is not zero, the meta key name and type used in the rule likely doesn't match the meta key present in events. Check to see if the meta key name and type used in the rule is valid by searching for the meta key name in **(Configure) > ESA Rules > Settings** tab (Meta Key References).

Problem	Possible Causes	Solutions
I created a custom rule (via the Rule Builder or Advanced EPL), and my rule is not firing. Why?	There may be a problem with the rule.	**If a specific rule is not firing**, go to **(Configure) > ESA Rules > Services** to see if the rule was disabled. In the **Deployed Rule Stats** section, a rule that is disabled displays a clear enabled button (instead of the green enabled button). You can also check Events Matched field. Go to **(Configure) > ESA Rules > Services**. From there, you can see the number of events that were matched in the **Events Matched** column. If no events matched, check the logic of your rule for errors. For example, check the syntax for uppercase and lowercase errors, and check the time window. If the rule still doesn't fire, consider simplifying the logic of the rule to see if it fires when there is less complexity.
After a recent upgrade, I am not seeing alerts and I am seeing disabled rules.	There may be a problem with the ESA rule deployment.	Deploy the ESA rule deployments again. ESA Rule Deployment Steps provides more information on deploying rules using the ESA Correlation service. If this does not resolve the issue, check the ESA Correlation log files for more information: `/var/log/netwitness/correlation-server/correlation-server.log`.
After an update or upgrade to 11.3.0.2 or later, if I try to make an adjustment to some rules, I get an error when trying to save them.	The Ignore Case option may be selected for a meta key that does not contain alphabetic values, such as IP address.	In NetWitness Platform 11.3.0.2 and later, the Ignore Case option has been removed from the ESA Rule Builder - Build a Statement dialog for meta keys that do not contain text values. Adding Ignore Case on meta keys which do not contain alphabetic values causes additional processing to occur for no added benefit. In the ESA Rule Builder - Build a Statement dialog, check to see if you have any meta keys that do not contain alphabetic characters, for example, ip_src and ip_dst. If you do, clear the **Ignore Case** checkbox for those meta keys and try to save the rule again.

Problem	Possible Causes	Solutions
After an upgrade to 11.3.0.2 or later, I see a warning message in the ESA Correlation service log file showing a difference between the multi-valued and default-multi-valued parameter meta key values. Why?	You do not have the required meta keys on ESA Correlation that the Endpoint, UEBA, and Live content rules need to work.	If you want to use the latest Endpoint, UEBA, and Live content rules, add the necessary meta keys to the `multi-valued` and `single-valued` parameter fields. For detailed information and instructions, see "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*. For example warning messages, see Example ESA Correlation Server Warning Message for Missing Meta Keys.
Meta keys marked as sensitive for data privacy are still included in notifications and alerts for some rules.	In ESA rules that do not select every piece of meta from the session (that is, using **`select *`**), you may see that data privacy (if enabled) and the **Pivot to Investigate > Navigate** link accessed from a context tooltip in the Respond Incident Details view does not work.	The following steps apply to all released versions of NetWitness 11.3 and later. In 11.4 and later, you do not need to follow these steps for data privacy. However, you need to follow these steps if you want to enable the **Pivot to Investigate > Navigate** link accessed from a context tooltip in the Respond Incident Details view. 1. Add the ESA generated `event_source_id` meta key to the `index-concentrator-custom.xml` file. 2. Add the `event_source_id` meta key to the SELECT statement within any ESA rule that does not select every piece of metadata from the session. 3. After the Concentrator changes take effect, redeploy the ESA rule deployment that contains the ESA rule. To do this, see Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event_source_id. For NetWitness 11.4 and later, to resolve the data privacy issue, see How to Remove Sensitive Meta Keys Globally from All Alerts.

Problem	Possible Causes	Solutions
The **Pivot to Investigate > Navigate** link does not work in a context tooltip accessed from Respond.	In ESA rules that do not select every piece of meta from the session (that is, using `select *`), you may see that data privacy (if enabled) and the **Pivot to Investigate > Navigate** link accessed from a context tooltip in the Respond Incident Details view does not work.	For NetWitness 11.3 and later (including 11.4), see Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event_source_id.
I added a data source filter to the data sources in my ESA rule deployment. It was working fine and then all of a sudden, I stopped receiving alerts.	If there are any adjustments to an application rule on the Decoders that are mapped to the data sources used in your data source filter, the filter does not work for that rule since the application rule used in the filter no longer exists.	The optional data source filter is available in NetWitness 11.5 and later. If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.

Problem	Possible Causes	Solutions
I added a data source filter to the data sources in my ESA rule deployment and I see an "Invalid header size" error while communicating with Core services in the ESA Correlation log file.	You are filtering out a large portion of the traffic and less sessions match the aggregation criteria.	In the Explore view node list for an ESA Correlation service, select **correlation > stream**. Decrease the `max-sessions` parameter from 10,000 to a lower session count and restart the ESA Correlation service. For step-by-step instructions, see Adjust Maximum Sessions for the Data Source Filter.
I created a rule using a Context Hub list and it was working properly, but suddenly, the rule stopped firing and ESA is not processing any rules.	If a Context Hub list that is used by ESA Correlation is renamed or deleted, ESA will not be able to access the list and may halt processing for all rules.	Do not rename or delete a Context Hub list that is used in a deployed ESA rule. ESA Correlation will also not be able to access a Context Hub list if you delete it and then add it again with the same name while the rule is deployed. If you rename a Context Hub list or recreate the Context Hub list with the same name, update the ESA rules that use that Context Hub list, and then redeploy the ESA rule deployments that contain those rules.

SMTP Notification Error Example

The following SMTP notification error example is an excerpt from a `correlation-server.log` file, which shows an error message for sending notifications with unsupported templates. In this example, there is a rule that is configured with the GeoIP enrichment, which has a hash table as one of its fields (the GeoIPLookup meta). Because the default SMTP template is only designed to deal with metas that are either singular values or arrays that contain only singular values, such as `"ip.src":"1.1.1.1"` and `"action":["fw:inbound-network-traffic"]`, sending the email notification fails due to the array containing a hash table.

```
FTL stack trace ("~" means nesting-related):

- Failed at: ${value!""} [in template "smtp.ftl" in macro "value_of" at line
1, column 152]

- Reached through: @value_of metadata[key] [in template "smtp.ftl" at line 85,
column 141]

----

...

For "${...}" content: Expected a string or something automatically convertible
to string (number, date or boolean), or "template output" , but this has
evaluated to an extended_hash (LinkedHashMap wrapped into
f.t.DefaultMapAdapter):

==> value!"" [in template "smtp.ftl" at line 1, column 154]
```

Integration-Server SMTP Notification Error Example

The following SMTP notification error example is an excerpt from an `integration-server.log` file, which shows a failure when the Integration-server attempts to send an email notification to the email notification server. In this case, you should check the email notification server configuration in the Global Notifications settings (**(Admin) > System > Global Notifications > Servers** tab).

```
2019-10-09 18:53:42,015 [-SMTP-5c45c867e4b03b89a49b78ba] WARN
Notification|SMTP dispatch failed (Reason: Sending the email to the following
server failed : email.server.com:25)

2019-10-09 18:53:42,100 [-SMTP-5c45c867e4b03b89a49b78ba] WARN
SystemOperation|Failed to forward ResolvedNotification
{server=5c45c867e4b03b89a49b78ba, destination=5c45c854e4b03b89a49b78b9,
content-length=30681}

java.lang.IllegalArgumentException: org.apache.commons.mail.EmailException:
Sending the email to the following server failed : email.server.com:25
```

Example ESA Correlation Server Warning Message for Missing Meta Keys

If you see a warning message in the ESA Correlation server error logs for missing multi-valued meta keys, there is a difference between the `default-multi-valued` parameter and `multi-valued parameter` meta key values, and the new Endpoint, UEBA, and Live content rules will not work. The same is true for missing single-valued meta keys. Completing the "Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules" procedure in the *ESA Configuration Guide* should fix the issue.

Multi-Valued Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[alert, alert_id,
browserprint, cert_thumbprint, checksum, checksum_all, checksum_dst,
checksum_src, client_all, content, context, context_all, context_dst,
context_src, dir_path, dir_path_dst, dir_path_src, directory,
directory_all, directory_dst, directory_src, email_dst, email_src,
feed_category, feed_desc, feed_name, file_cat, file_cat_dst, file_
cat_src, filename_dst, filename_src, filter, function, host_all,
host_dst, host_orig, host_src, host_state, ip_orig, ipv6_orig, OS,
param, param_dst, param_src, registry_key, registry_value, risk,
risk_info, risk_suspicious, risk_warning, threat_category, threat_
desc, threat_source, user_agent] are still MISSING from multi-valued
```

Single Value Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[accesses,
context_target, file_attributes, logon_type_desc, packets] are still
MISSING from single-valued
```

Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event_source_id

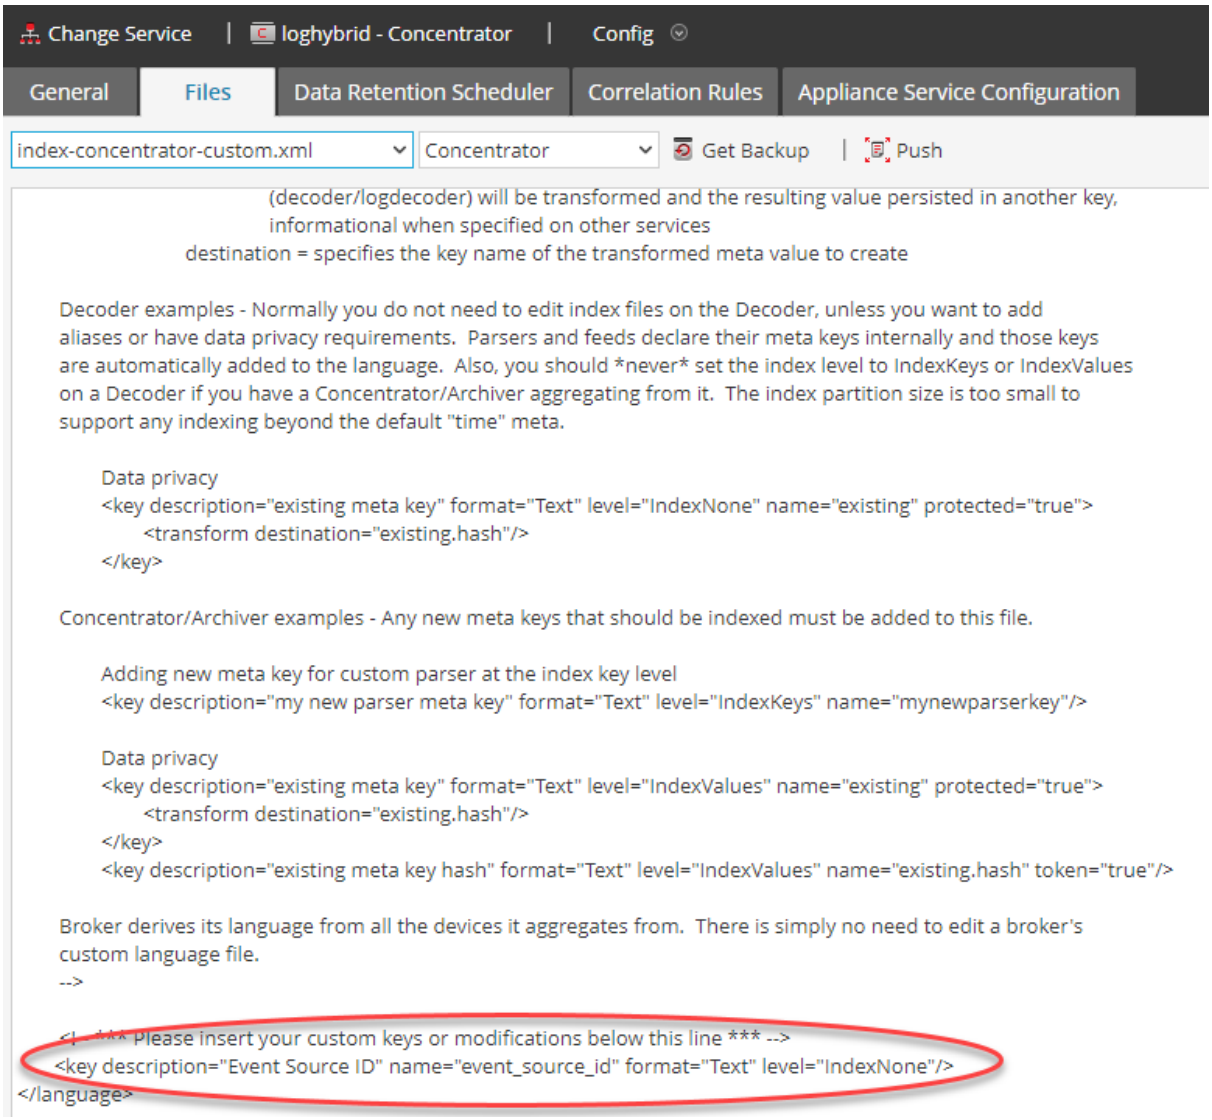
In ESA rules that do not select every piece of meta from the session (that is, using **`select *`**), you may see that data privacy (if enabled) and the **Pivot to Investigate > Navigate** link accessed from a context tooltip in the Respond Incident Details view does not work.

The following steps apply to all released versions of NetWitness 11.3 and later. In 11.4 and later, you do not need to follow these steps for data privacy, instead, see How to Remove Sensitive Meta Keys Globally from All Alerts. However, you need to follow these steps if you want to enable the **Pivot to Investigate > Navigate** link accessed from a context tooltip in the Respond Incident Details view.

Note: Do not use any Esper keyword as custom meta keys since this causes an error while creating an ESA Rule. For Esper keywords, see Reserved keywords.

1. Update the **`index-concentrator-custom.xml`** file to include the ESA generated **`event_source_id`** meta key. If you do not add the meta key, ESA cannot recognize it and the rule will fail to deploy.
   The following figure shows the file configured for the custom meta key "**`Event Source ID`**" with index settings of "**`IndexNone`**" with a format of "**`Text`**".

   <key description="Event Source ID" name="event_source_id" format="Text" level="IndexNone"/>



To save and deploy the new setting on the NetWitness host, select the **Apply** button. To force the change, restart the Concentrator service or you can wait until the next polling interval for the change to be recognized.

The XML file can also be deployed to other NetWitness hosts by clicking the **Push** button and selecting the destination NetWitness host. Only deploy the XML file to a NetWitness host that runs that service (that is, other Concentrators).

2. Update any ESA rule that selects only certain meta from the session to include the ESA generated **`event_source_id`** meta key. Add the **`event_source_id`** meta key to the **`SELECT`** statement. See the following example rule.Example Rule with event_source_id

@RSAAlert SELECT user_dst, reference_id, hostname, event_source_id FROM Event(device_class='Windows Hosts', reference_id IN ('4624' , '4625'), user_dst IS NOT NULL, user_dst NOT LIKE '%$', user_dst NOT IN ('ANONYMOUS LOGON','SYSTEM')) .win:time_batch(10 Minutes) GROUP BY user_dst HAVING COUNT(distinct hostname) >= 15;

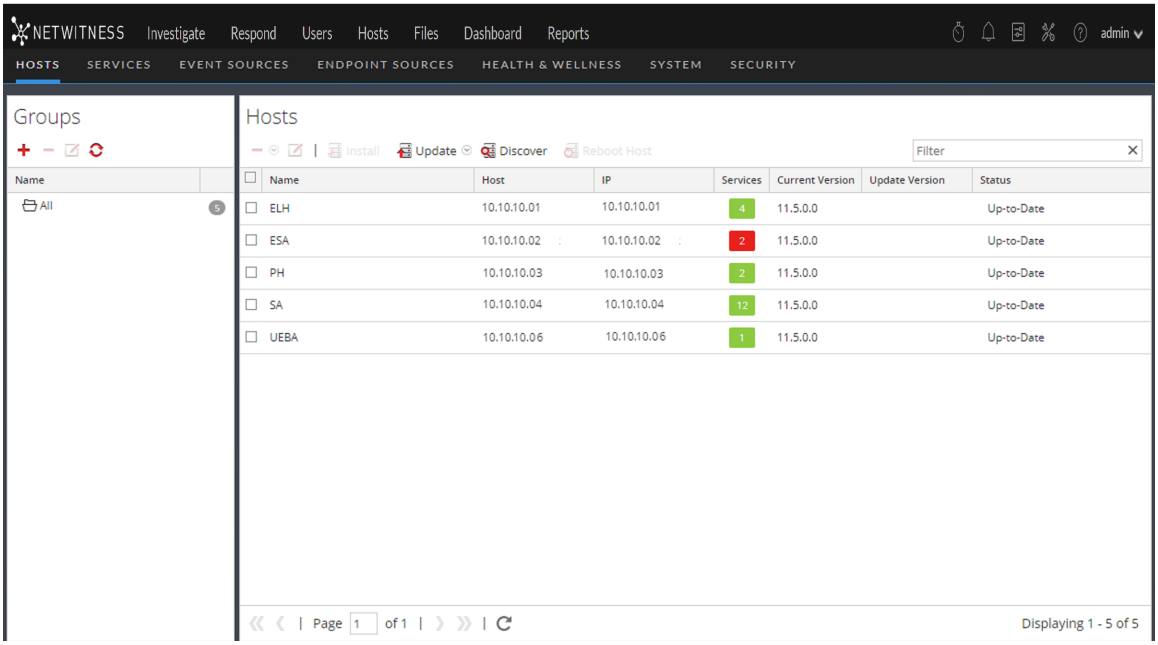
3. After the Concentrator changes takes effect, redeploy the ESA rule deployment that contains the ESA rule.

For additional information, see How ESA Handles Sensitive Data. For information on the stragegy and benefits of obfuscating data, see the *Data Privacy Management Guide*.

Steps to Troubleshoot Memory Issues with an ESA Service Offline

Step 1: Verify that your Host Is Running

The first step to troubleshooting is to ensure that your host is running. To do this, go to **(Admin) > Hosts**. If the host is down, the system parameters will not display (updating host information can sometimes be delayed), the **Services** display in red, and you may see an error message.

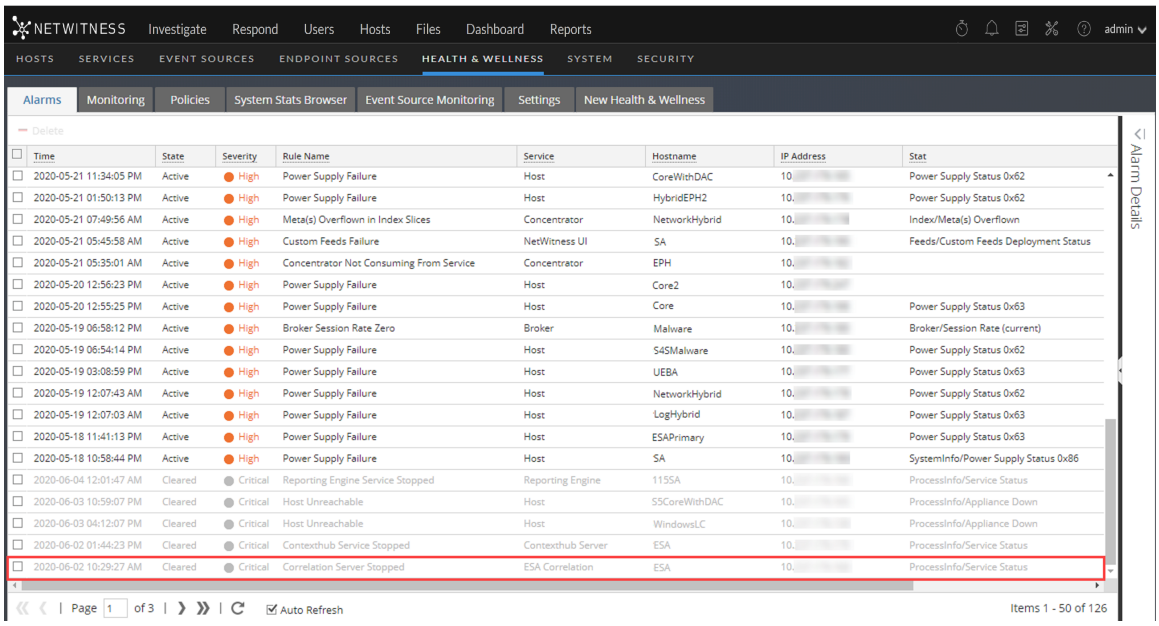


If your host is down, contact your NetWitness Administrator to restart it. Otherwise, go to Step 2.

Step 2: View Detailed Statistics in Health & Wellness

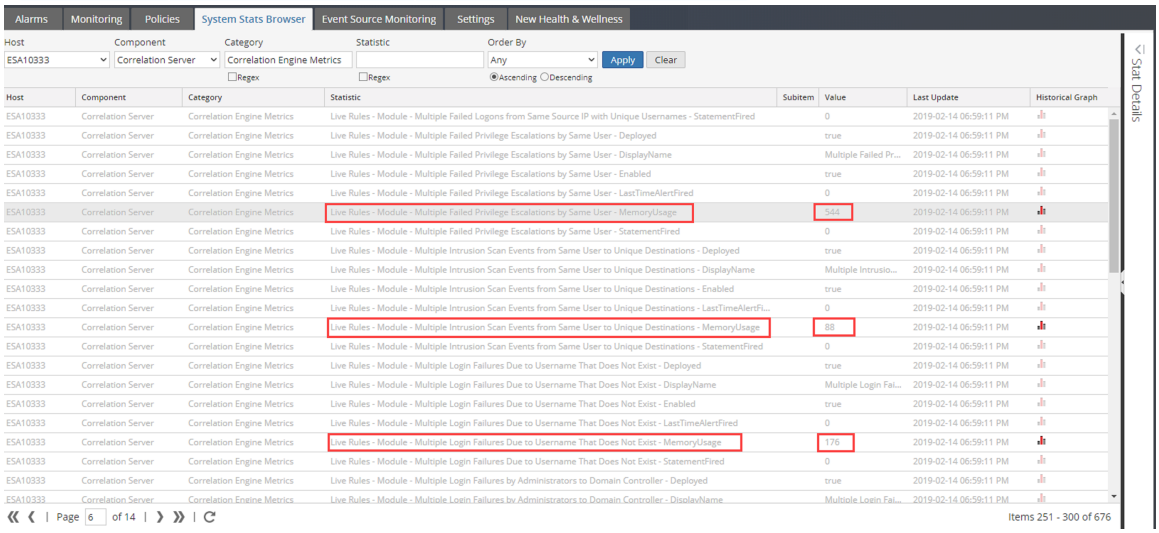
If your ESA service is down, you can go to Health & Wellness and view the **last known metrics** to see where potential issues are occurring. The most common problem is that your ESA service is exceeding memory thresholds, which causes it to stop or fail. In NetWitness 11.5 and later, see also View Health Statistics and Trends for ESA Correlation in New Health & Wellness.

1. Go to **(Admin) > Health & Wellness > Alarms** to see if the ESA triggered any alarms. Look for the following alarms for ESA Correlation:
   - Correlation Server in Critical State
   - Correlation Server in Unhealthy State
   - Correlation Server Stopped



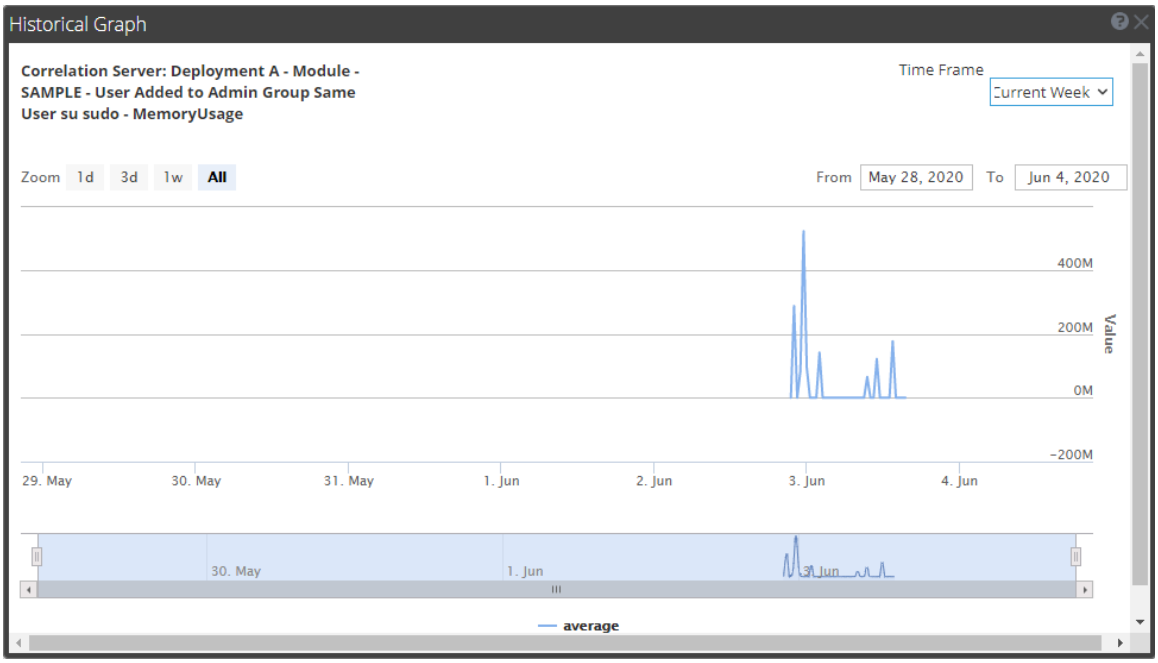
2. Go to **(Admin) > Health & Wellness > System Stats Browser** to see the memory metrics for each rule's performance. To view the metrics, enter the following and click **Apply**:

Host	Component	Category
<your host>	Correlation Server	Correlation Engine Metrics

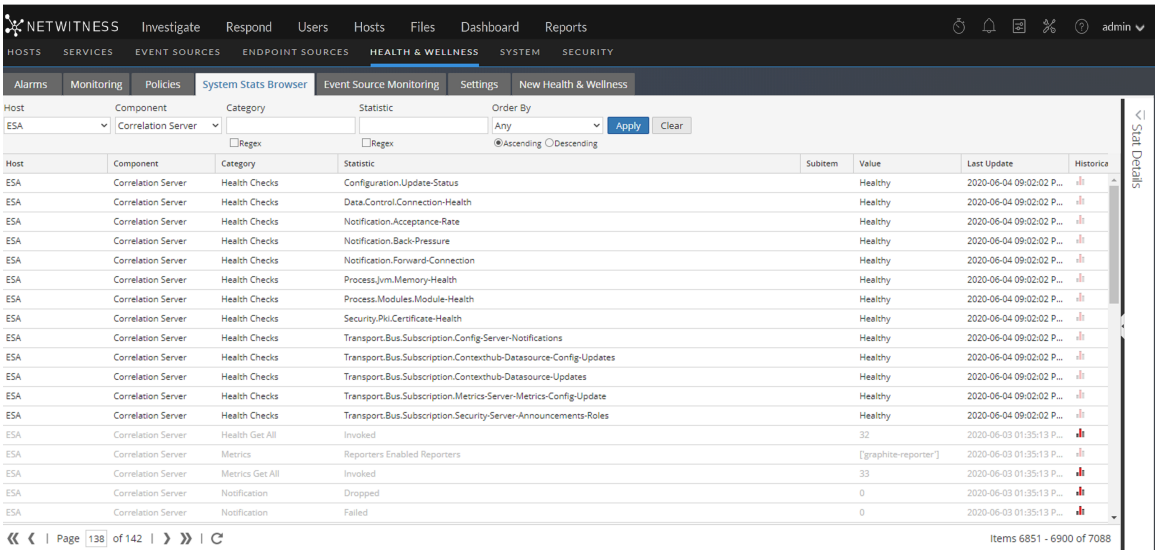


The name of the rule is in the **Statistic** column and the memory usage in bytes is in the **Value** column.

3. Click to view a historical view of memory usage for the rule in the **Historical Graph** column.



4. In the **System Stats Browser**, you can also see details of your ESA Correlation service performance.

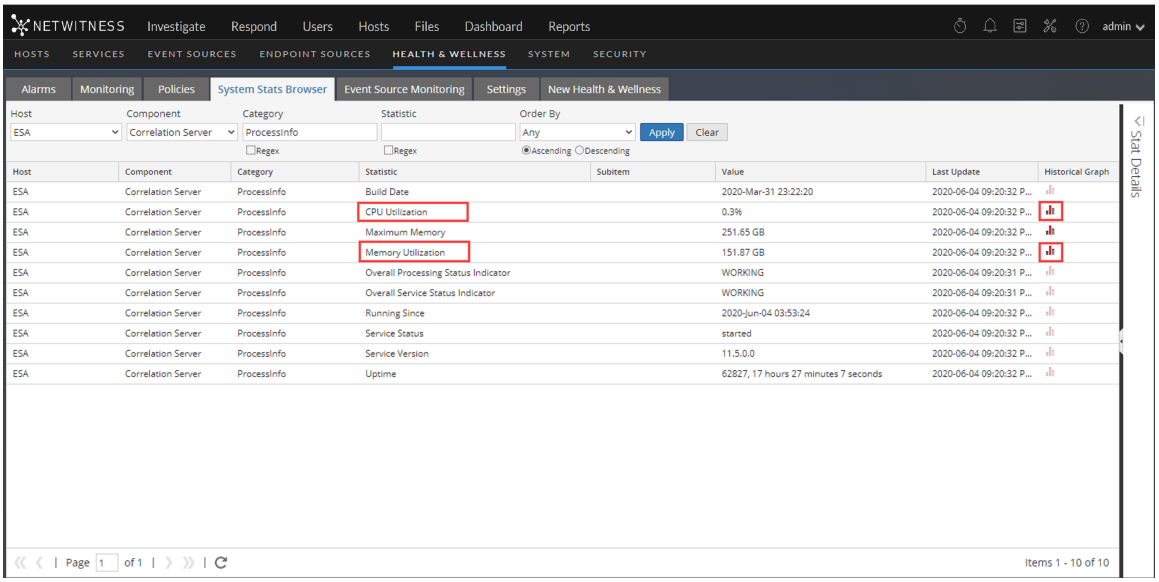


Select your host, and use the following filters to view the following statistics:

Host	Component	Category	Statistic	Example
<your host>	Host	SystemInfo	CPU Utilization	49.74%
<your host>	Host	SystemInfo	Memory Utilization	72.87%

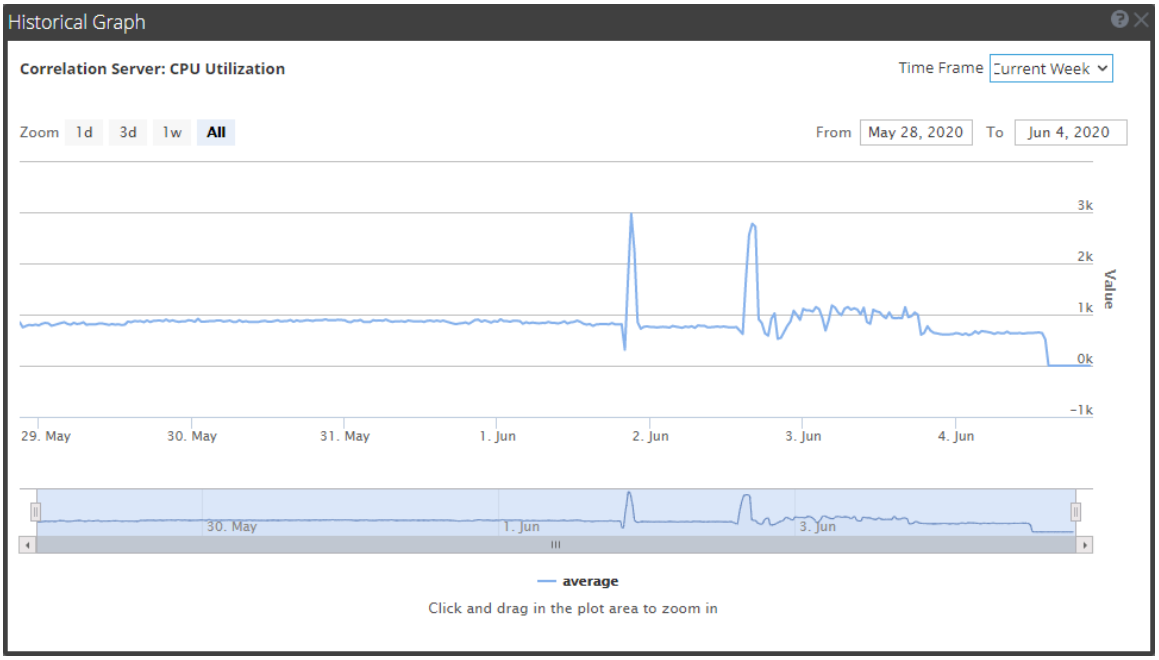
Host	Component	Category	Statistic	Example
<your host>	Host	SystemInfo	Used Memory	183.14 GB
<your host>	Host	SystemInfo	Total Memory	251.65 GB
<your host>	Host	SystemInfo	Uptime	289546, 3 days 8 hours 25 minutes 46 seconds
<your host>	Correlation Server	Process jvm	Memory Total Max	163 GB
<your host>	Correlation Server	Process jvm	Memory Total Used	13.50 GB
<your host>	Correlation Server	ProcessInfo	CPU Utilization	0.3%
<your host>	Correlation Server	ProcessInfo	Maximum Memory	251.65 GB
<your host>	Correlation Server	ProcessInfo	Memory Utilization	151.87 GB

The following figure shows the location of the ESA Correlation service CPU and Memory Utilization statistics.

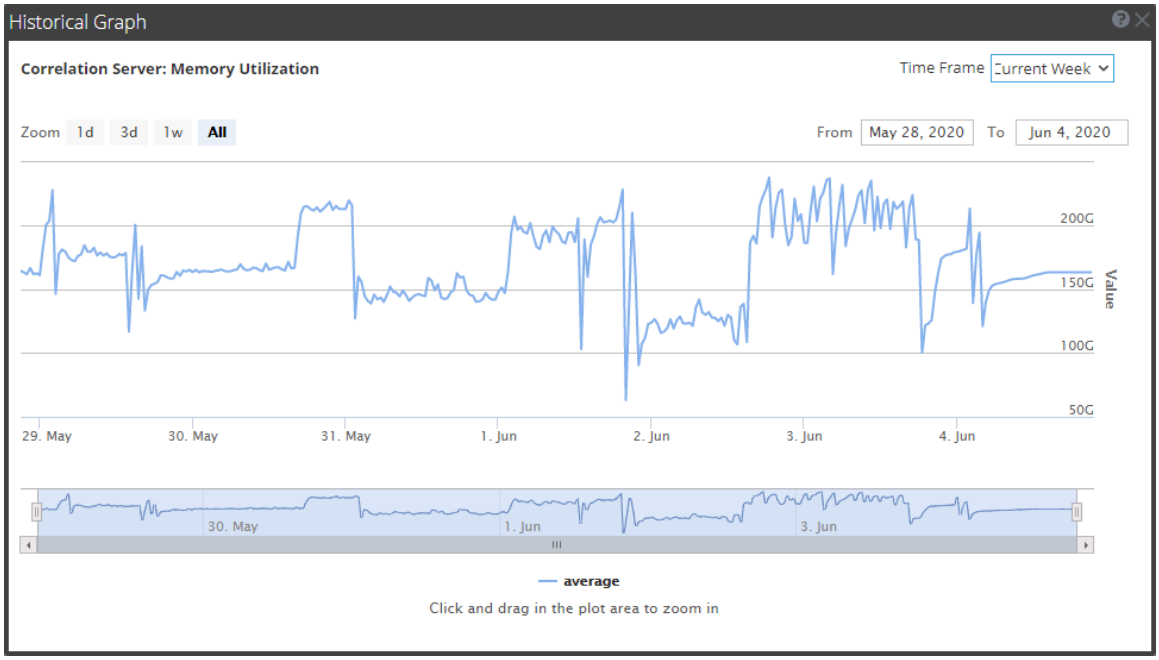


5. Click  to view a historical view of CPU and memory utilization.

   The following figure shows the historical graph of **CPU utilization**.



The following figure shows the historical graph of **Memory Utilization**.



If you are having a problem with memory or CPU utilization, continue to step 3.

Step 3: Bring up your ESA Services

1. Go to **(Admin) > Services**, select your ESA service, and then select **> Start**.
2. Return to the ESA Service to troubleshoot which rules have created memory issues.

If your ESA service is stopping and restarting in a loop, you may need to call Customer Support to get the services to start.

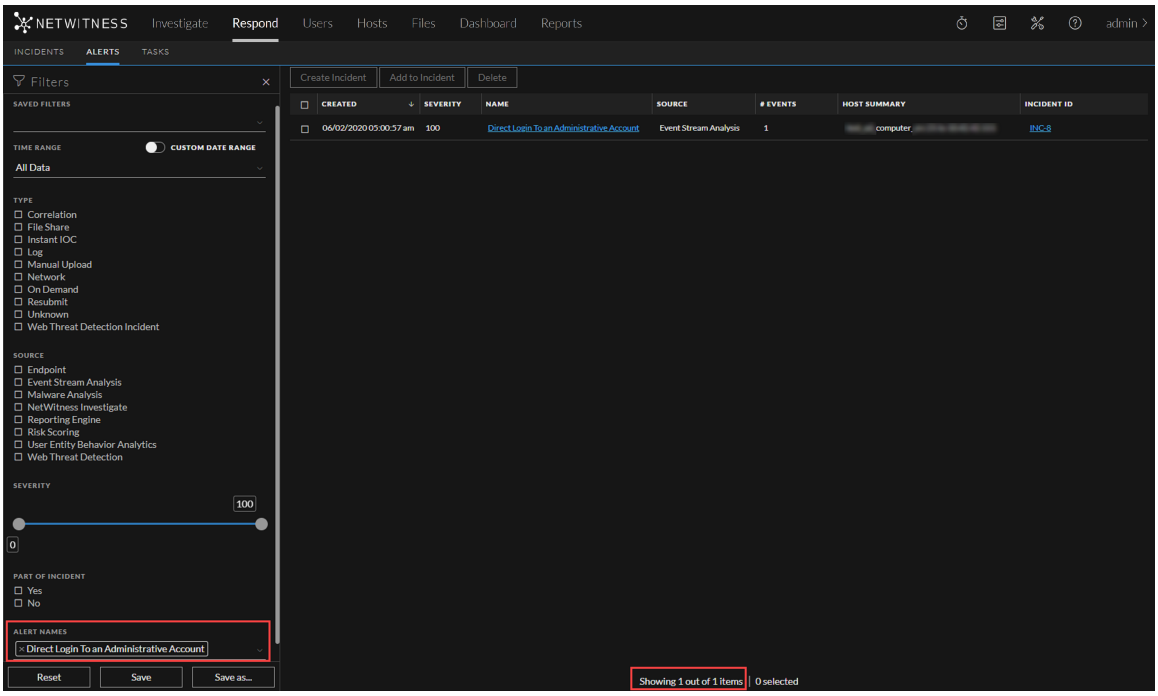
If you are able to start your ESA service without a shutdown, continue to step 4.

Step 4: Check the Alerts and Events Volume

After you are able to restart your ESA service without an immediate shutdown, you can review the stats for your rules to see which rules are consuming too many resources. Sometimes, ESA services fail because a rule is generating too many alerts or a rule is matching too many events. Check for both of these issues if you have determined that memory usage is causing your ESA service to shut down.

View Alert Summaries

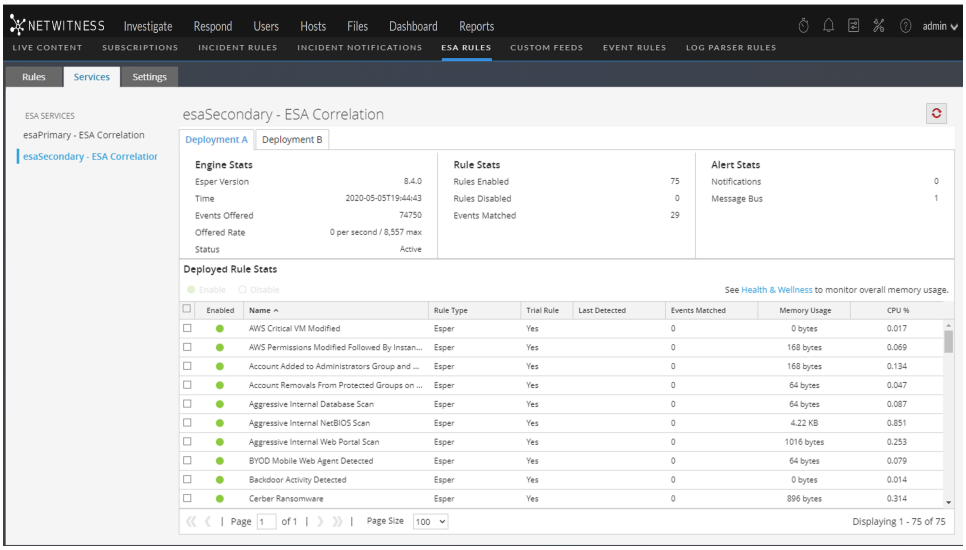
Rules that generate a high volume of alerts can overwhelm the system and cause it to fail or restart. To view the alert summaries, go to **Respond > Alerts**. In the **Filters** panel on the left, in the **Alert Names** section, select the alert name for the rule. The number of alerts with that name appears at the bottom of the Alerts list results. If the number is significantly high for a particular rule, you need to disable the rule and rewrite it to be more efficient.



To clear your filter, click **Reset**.

View Events Matched

Sometimes a rule matches too many events, which can use up excessive memory. This typically occurs if you create a large event window where a great number of events accumulate without triggering an alert. This is a problem because each event is stored in memory while the rule waits for the alert to trigger. To check for this issue, go to **(Configure) > ESA Rules > Services**. From there, you can see the number of events that were matched in the **Events Matched** column for the deployment. If a high number of events were matched for a given rule, you can investigate the rule further to see if you can make it more efficient.



Step 5: Disable and Repair the Rule that Caused Issues

Once you have determined the rules that need to be rewritten, disable them and rewrite rules so that they don't generate such a high volume of alerts or events. For pointers on how to write more efficient rules, see Best Practices.

Disable Rules

1. To disable rules, go to **(Configure) > ESA Rules > Services**, and select the rules you want to disable in the **Deployed Rules Stats** field.
2. Select **Disable** to disable the rules.

Edit Rules

1. To repair the rules, go to **(Configure) > ESA Rules > Rules tab > Rule Library**.
2. For each rule that you repair, do the following:
   a. Select the rule to edit and then select **> Edit**.
   b. Edit the rule to be more efficient. For instructions on creating rules, see Add Rules to the Rule Library
   c. When you are satisfied with your rule, you can save the rule as a trial rule to ensure that any memory issues do not affect ESA services performance. To do this, follow the steps listed in Working with Trial Rules.

Deploy Rules

1. Go to **(Configure) > ESA Rules > Rules** tab.
2. In the options panel on the left, select the deployment that contains the rule.
3. In the Deployment view, the rule that you changed shows a status of Updated. Click **Deploy Now**. The rule status changes to Deployed.

Verify that the Rules are Enabled

After you deploy the ESA rules, they should automatically show as enabled. If not, you can enable the rules.

1. Go to **(Configure) > ESA Rules > Services** tab, and select the ESA service in the options panel.
2. On the deployment tab for the deployment that contains the rules, in the Deployed Rule Stats section, look at the status of the rules in the Enable column. Enabled rules show a green circle. If the rules show a white circle, you can enable the rules.
3. To enable rules, select the rules you want to enable and select **Enable** above the table.

(Optional) Check the ESA Correlation Log Files for More Information

Once you verify that your services are down and some potential causes for the system going down, check to see if the service is stopping and restarting in a loop. To do this, go to the ESA Correlation logs. You can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`.

ESA Rule Troubleshooting with Nw-Shell

Note: This procedure applies to NetWitness Platform 11.3 and later versions.

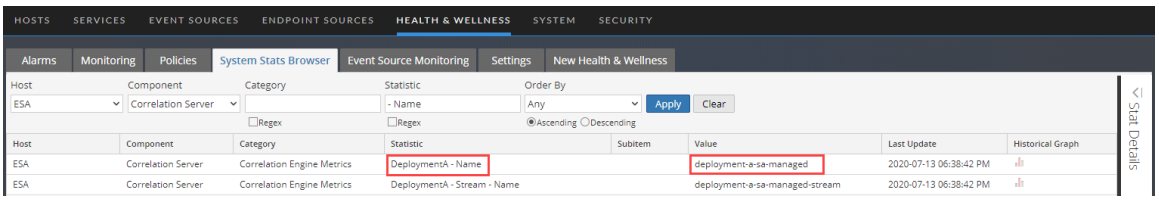
The ESA Correlation service replaces the Event Stream Analysis service in NetWitness version 11.3 and later. As a result of this change, some settings are no longer available in the user interface. In addition to the standard troubleshooting methods available, you can use the **nw-shell** utility to perform advanced troubleshooting of the ESA Correlation service and rules. For detailed information on the nw-shell utility, see the *NetWitness Shell User Guide*.

- Find Your Engine Name for Nw-Shell
- Connect to an ESA Correlation Server
- View the Contents of a Named Window
- See the Method Input and Output

Find Your Engine Name for Nw-Shell

Follow these steps to find your engine name using your ESA rule deployment name. Your engine name is required for ESA Nw-Shell troubleshooting. Locate the names of each deployment that you plan to troubleshoot.

1. In the NetWitness UI, go to **(Admin) > Health & Wellness > System Stats Browser**.
2. In the System Stats Browser, use the following filters and then click **Apply**.
   a. **Host**: Select your ESA host.
   b. **Component**: Select **Correlation Server**.
   c. **Statistic**: Type **- Name** (put a space between the dash and name).
3. In the **Statistic** column, locate your deployment followed by **- Name**. The name in the **Value** field is `YOUR ENGINE NAME`.



In the above example, the deployment name is DeploymentA and the engine name is `deployment-a-sa-managed`.

Connect to an ESA Correlation Server

1. Log in to nw-shell:
   a. Connect via SSH to the NW server (head unit).
   b. After logging in and getting the command prompt, type `nw-shell`.
2. Connect to ESA:
   a. Go to your ESA physical host and type the command:
   `cat /etc/netwitness/correlation-server/service-id`
   b. Go back to the NW server and nw-shell and connect to the correlation server:
   `connect --service correlation-server.ID`
3. Log in to ESA with the admin user credentials.
   a. Type `login`.
   b. Enter the username and password of the admin user.

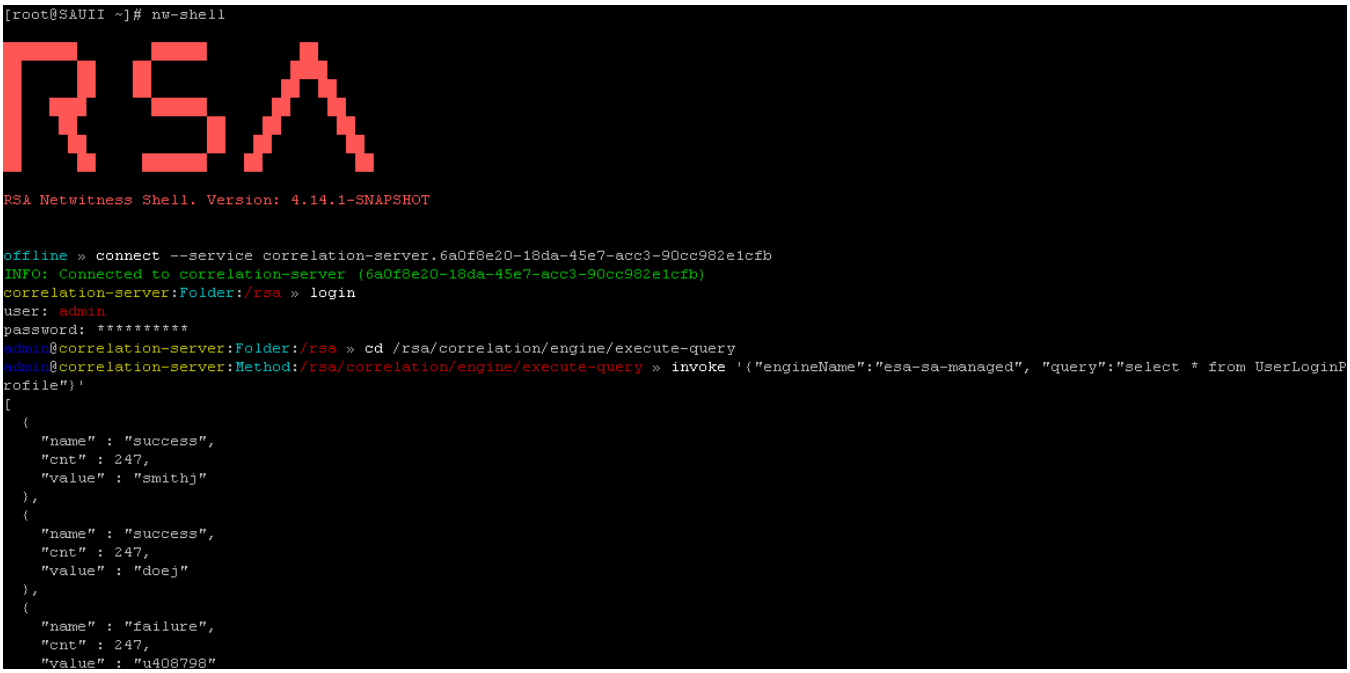


View the Contents of a Named Window

Use the execute-query method to view the contents of a named window.

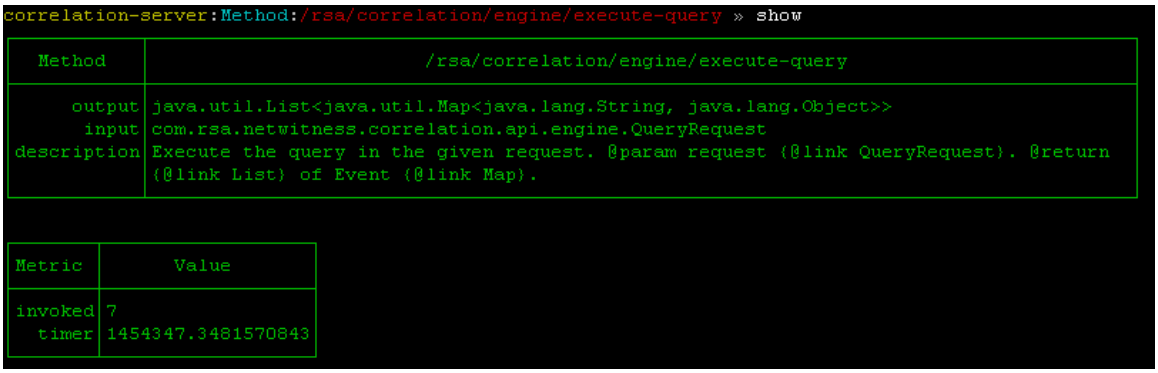
1. After you are connected to the ESA correlation service and authenticated, type:
   `cd /rsa/correlation/engine/execute-query`
2. Type `invoke '{"engineName":"<YOUR ENGINE NAME>", "query":"<YOUR QUERY>"}`
   - Where `<YOUR ENGINE NAME>` = the engine name that you located in Find Your Engine Name for Nw-Shell.
   - Where `<YOUR QUERY>` = the select statement into the named window.

Example: `invoke '{"engineName":"esa-sa-managed", "query":"select * from UserLoginProfile"}'`



See the Method Input and Output

Type `show` at the command line to see the input and output expected. All method invocation must be prefaced by `invoke`.



Obtain Correlation Server Metrics for ESA Rule Deployment Troubleshooting Using Nw-Shell

Note: This procedure is available in NetWitness Platform version 11.4.1 and later.

You can use Nw-Shell to view ESA Correlation Server metrics for each of your ESA rule deployments. These metrics show the number of sessions behind for the deployment data sources as well as the memory usage for the rules in the deployment.

1. Find the engine name to use for Nw-Shell. See Find Your Engine Name for Nw-Shell.
2. Connect to an ESA Correlation Server. See Connect to an ESA Correlation Server.
3. After you are connected to the ESA correlation service and authenticated, type:
   `cd /rsa/correlation/service/stats/get-condensed-metrics`
4. Type `invoke '<YOUR ENGINE NAME>'`
   Where `<YOUR ENGINE NAME>` is the engine name that you located in Find Your Engine Name for Nw-Shell.

Here is an example of the metrics output that you can obtain for your ESA rule deployment:

```
{
  "engineName" : "esa",
  "eventsOffered" : 1650,
  "maxEventsRate" : 2.337019271237945,
  "eventsRate" : {
    "count" : 1650,
    "oneSecRate" : 0.02129597415982235,
    "meanRate" : 0.05795554387594279,
    "oneMinuteRate" : 0.15485195471608212,
    "fiveMinuteRate" : 0.12419048320215775,
    "fifteenMinuteRate" : 0.11923922260543295
  },
  "streamMetrics" : {
    "pollingRate" : {
      "count" : 30100,
      "meanRate" : 1.0572440914676082,
      "oneMinuteRate" : 1.1497867142612552,
      "fiveMinuteRate" : 1.1213207743063618,
      "fifteenMinuteRate" : 1.1178158863433476
    },
    "positionTracking" : 475,
    "polling" : 30100,
    "bufferedRecords" : 0,
    "incomingRecords" : {
      "count" : 1650,
      "meanRate" : 0.05796815004652209,
      "oneMinuteRate" : 0.16275810590703457,
      "fiveMinuteRate" : 0.1253818211054517,
      "fifteenMinuteRate" : 0.11956773913684943
    },
    "outgoingRecords" : {
      "count" : 1650,
      "meanRate" : 0.05796815072883537,
      "oneMinuteRate" : 0.16275810590703457,
      "fiveMinuteRate" : 0.1253818211054517,
      "fifteenMinuteRate" : 0.11956773913684943
    },
    "sourceMetrics" : {
      "nw://admin@10.10.10.01:50005?compression=0&compressionLevel=6" : {
        "bufferedRecords" : 0,
        "incomingRecords" : {
          "count" : 1650,
          "meanRate" : 0.05796819356968545,
          "oneMinuteRate" : 0.16275810590703457,
          "fiveMinuteRate" : 0.12538181849661215,
          "fifteenMinuteRate" : 0.11956716772097528
```

```
        },
        "outgoingRecords" : {
          "count" : 1650,
          "meanRate" : 0.05796819300115741,
          "oneMinuteRate" : 0.16275810590703457,
          "fiveMinuteRate" : 0.12538181849661215,
          "fifteenMinuteRate" : 0.11956716772097528
        },
        "sessionsBehind" : 2,
        "sessionLastId" : "1645",
        "sessionRate" : "0",
        "lastReceivedSessionId" : 1645
      }
    }
  },
  "ruleMetrics" : [
    {
      "ruleName" : "create_persist",
      "memoryUsage" : 104,
      "cpuLockedTimePercentage" : 50.286,
      "cpuLockedTimeNanos" : 138253,
      "statementFired" : 0,
      "alertsFired" : 0
    },
    {
      "ruleName" : "test_persist",
      "memoryUsage" : 72,
      "cpuLockedTimePercentage" : 39.44,
      "cpuLockedTimeNanos" : 108433,
      "statementFired" : 0,
      "alertsFired" : 0
    },
    {
      "ruleName" : "test-per",
      "memoryUsage" : 0,
      "cpuLockedTimePercentage" : 10.275,
      "cpuLockedTimeNanos" : 28249,
      "statementFired" : 0,
      "alertsFired" : 0
    }
  ]
}
```

View Memory Metrics for Rules

This topic tells ESA rule writers how to view memory metrics for an ESA Correlation service and its associated ESA rules. You can see estimated memory usage for each rule running on a server, and you can use this information to modify your rule statements and conditions if they use too much memory.

Rules can sometimes consume more memory than you expect, causing ESA to slow down or stop. To see approximately how much memory a rule is using, you can view estimated memory usage for each rule in the Health & Wellness System Stats browser (you need permissions to access this module). You can use this information to modify your rules to be more efficient.

At a high level, you need to complete the following steps to use memory metrics to troubleshoot memory usage for rules:

1. Ensure you have the correct permissions to view the Health & Wellness module. For information on roles and permissions, see ESA Permissions.
2. View the memory statistics in Health & Wellness.
3. (Recommended) Configure Health & Wellness ESA policies to send an email if memory thresholds are exceeded. See "Manage Policies" in the *System Maintenance Guide* for instructions on sending email notifications.
4. Use the memory metrics data to modify rules to be more efficient, if necessary.

Note: You can also view memory metrics for ESA rules in the  **(Configure) > ESA Rules > Services** tab. See View Stats for an ESA Service.

Prerequisites

The following are requirements for using memory metrics:

- You must have the appropriate permissions to view Health & Wellness statistics.
- (Recommended) Configure the ESA Health & Wellness policy to send an email when memory thresholds are exceeded.

Note: Memory Metrics is always on for the ESA Correlation service; you do not have to enable it.

View Health Statistics and Trends for ESA Correlation in New Health & Wellness

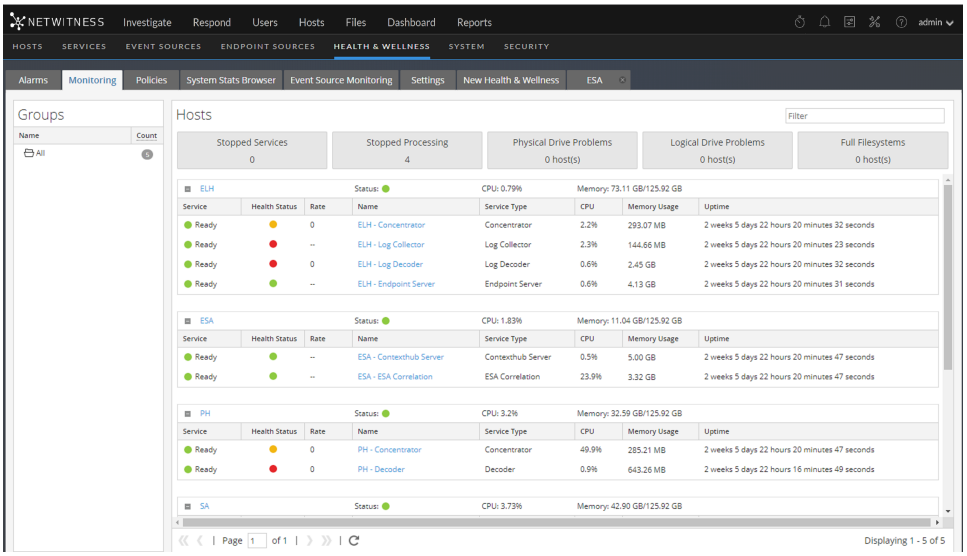
In NetWitness version 11.5 and later, New Health & Wellness provides improved and intuitive dashboards, monitors, and visualizations. The ESA Correlation Overview dashboard provides health statistics and trends on ESA rule deployments.



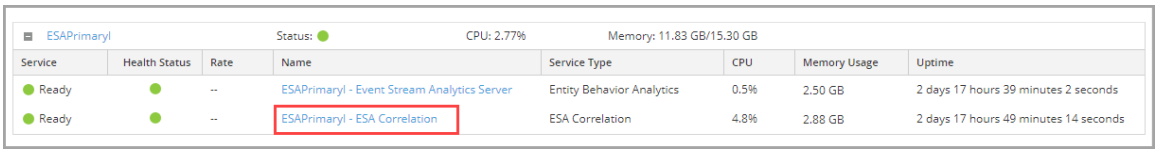
For more information, see "Monitor New Health and Wellness" and "Appendix A: New Health and Wellness Dashboards / ESA Correlation Overview Dashboard" in the *System Maintenance Guide*.

View Memory Metrics for an ESA Correlation Service in Health & Wellness

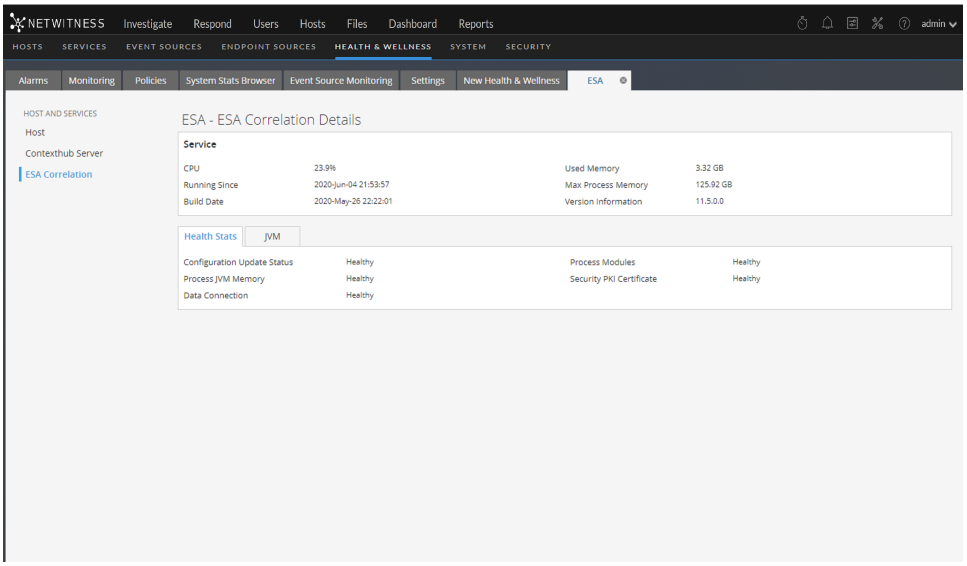
1. Go to  **(Admin) > Health & Wellness > Monitoring** tab.



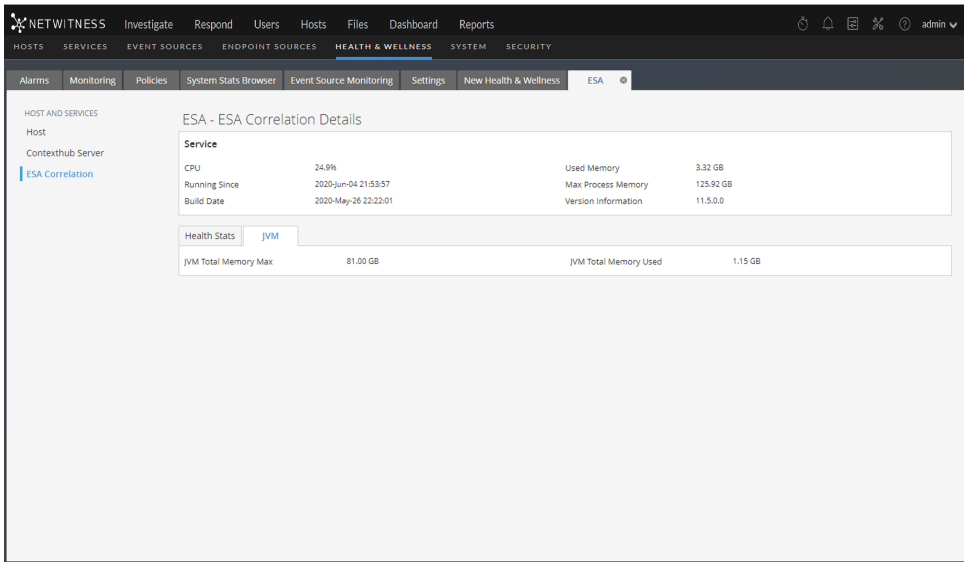
2. Locate your host and click the link in the **Name** field for your ESA Correlation service, for example, ESA- ESA Correlation.



3. On the tab for your ESA host, click the **Health Stats** tab.
You can view the health status of the ESA Correlation service.



4. Click the **JVM** tab.
   You can view the JVM total memory used by the selected ESA Correlation service.

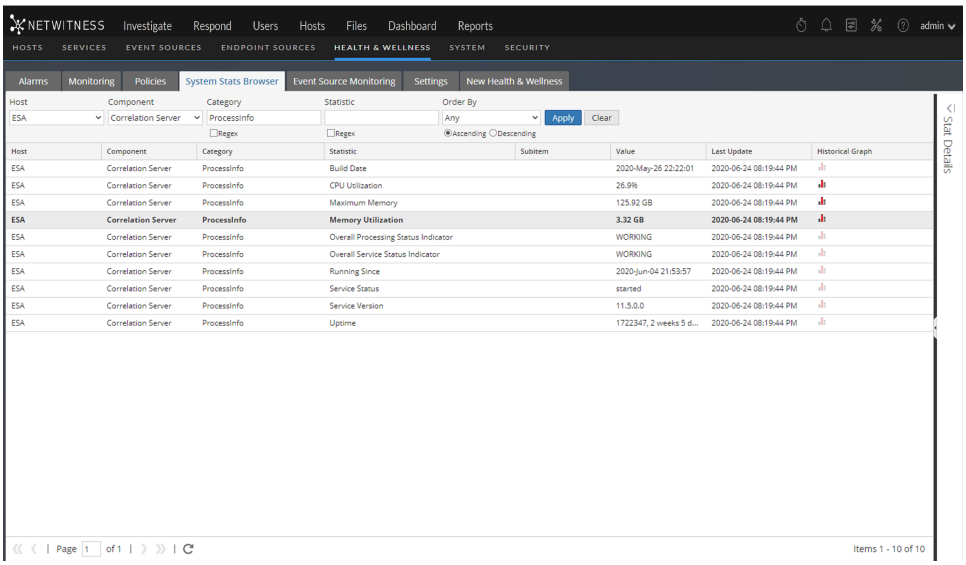


Note: You can also view memory metrics for the ESA Correlation service in the **(Configure) > ESA Rules > Services** tab. See View Stats for an ESA Service.

View Memory Metrics for an ESA Correlation Service and its ESA Rules

1. Go to **(Admin) > Health & Wellness > System Stats Browser**.
2. To view memory metrics for an ESA Correlation service, in the **Host** field, select your ESA host. Select **Correlation Server** for **Component**, enter **ProcessInfo** for **Category**, and then click **Apply**.

Host	Component	Category
<your host>	Correlation Server	ProcessInfo



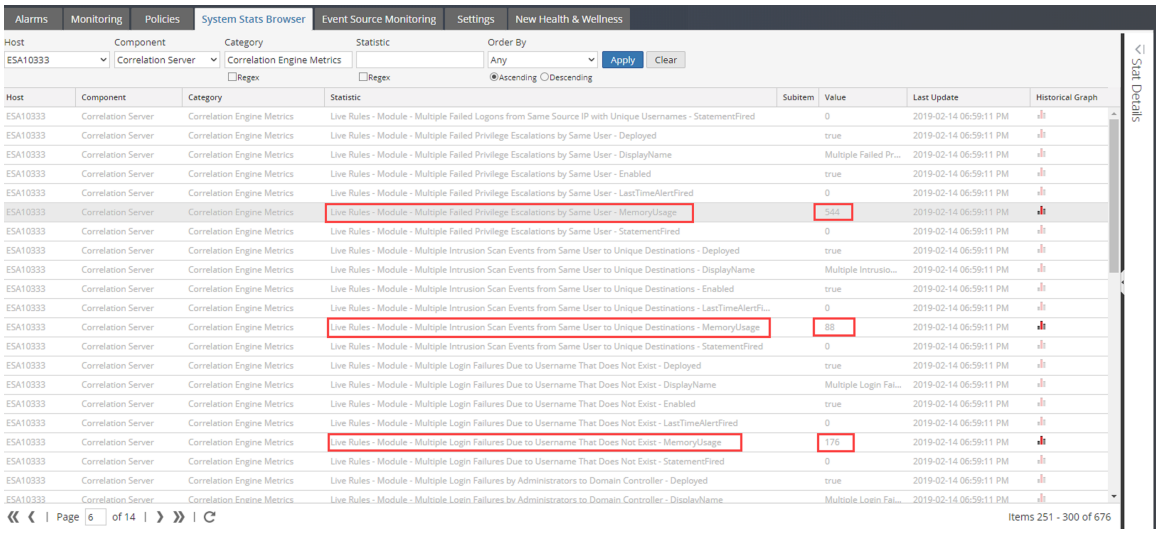
The **Memory Utilization** statistic shows the total memory in use by the ESA Correlation service.

3. To view the historical memory usage for the ESA Correlation service, click the **Historical Graph** icon.



4. To view the memory metrics for individual rules, in the **Category** field, enter **Correlation Engine Metrics** and click **Apply**.

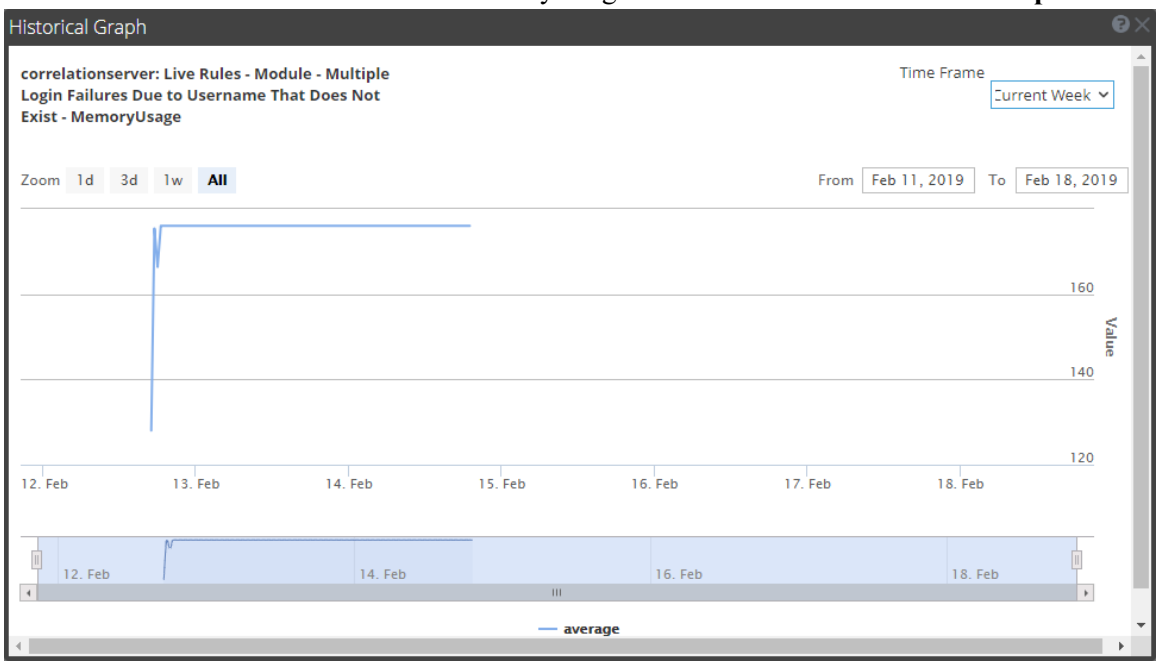
Host	Component	Category
<your host>	Correlation Server	Correlation Engine Metrics



The name of the rule is in the **Statistic** column appended with **MemoryUsage** and the memory usage in bytes is in the **Value** column.

Note: The **Last Update** field reflects when Health & Wellness polls ESA. However, the Memory Metrics is not synchronized with the Health & Wellness polling. For example, if the memory threshold is exceeded on 2/10/19 at 12 p.m., but Health & Wellness polls at 2/10/19 at 12:10 p.m., the **Last Update** field will display a timestamp of 2/10/19 12:10 p.m.

5. Click  to view a historical view of memory usage for the rule in the **Historical Graph** column.



Note: You can also view memory metrics for ESA rules in the  **(Configure) > ESA Rules > Services** tab. See View Stats for an ESA Service.

How ESA Handles Sensitive Data

This topic explains how ESA treats sensitive data, such as usernames or IP address, that it receives from Core services. The Data Privacy Officer (DPO) role can identify meta keys that contain sensitive data and should display obfuscated data. ESA does not display or store sensitive meta. Consequently, ESA will not pass sensitive data to NetWitness Respond.

Optionally, ESA can add an obfuscated version of the sensitive data to an event. For example, the DPO identifies `user_dst` as sensitive. ESA can add an obfuscated version, such as `user_dst_hash`, to an event. The obfuscated meta is not sensitive, so ESA will display and store it the same way as any other non-sensitive meta.

For more information on the strategy and benefits of obfuscating data, see the *Data Privacy Management Guide*.

This topic explains the following:

- How ESA treats sensitive data it receives from Core services
- How to prevent sensitive data leaks in an Advanced EPL rule
- How to remove sensitive meta keys from global alerts

How ESA Treats Sensitive Data from Core Services

When ESA receives sensitive data from Core services, ESA passes on only the obfuscated version of the data. ESA does not store or show sensitive data.

The following features are impacted:

- Outputs – ESA does not forward sensitive data to outputs, which include alerts, notifications, and MongoDB storage.
- Advanced EPL rules – If an EPL statement creates an alias for a sensitive meta key, sensitive data will leak. This topic illustrates how this happens so you can avoid it.
- Enrichments – If a sensitive meta key is used in the join condition, sensitive data will leak. This topic illustrates how this happens so you can avoid it.

Advanced EPL Rule

If an EPL query statement renames a sensitive meta key, the data will not be protected.

ESA identifies a sensitive meta key by the name:

- `ip_src` is the sensitive meta key.
- `ip_src_hash` is the non-sensitive, obfuscated version.

To support data privacy, the sensitive meta key must not be renamed in an EPL query. If a sensitive meta key is renamed, the data will no longer be protected.

For example, in a rule such as select `ip_src as ip_alias...`, `ip_alias` contains the sensitive data but it is not protected because ESA only knows about `ip_src`, not `ip_alias`. In this case, IP addresses would not be obfuscated. Real values would be displayed.

Enrichment Source

When a sensitive meta key is used in a join condition, sensitive data can be displayed.

The enrichment database, which is the other part of the join condition, has one column that matches the sensitive meta key. This cross reference is to actual values not obscured values. Consequently, actual values are displayed.

In the following example, both parts of the join condition are highlighted.



- `ip_src` contains sensitive data.
- `ipv4` will be added to the alert and exposed as non-sensitive data.

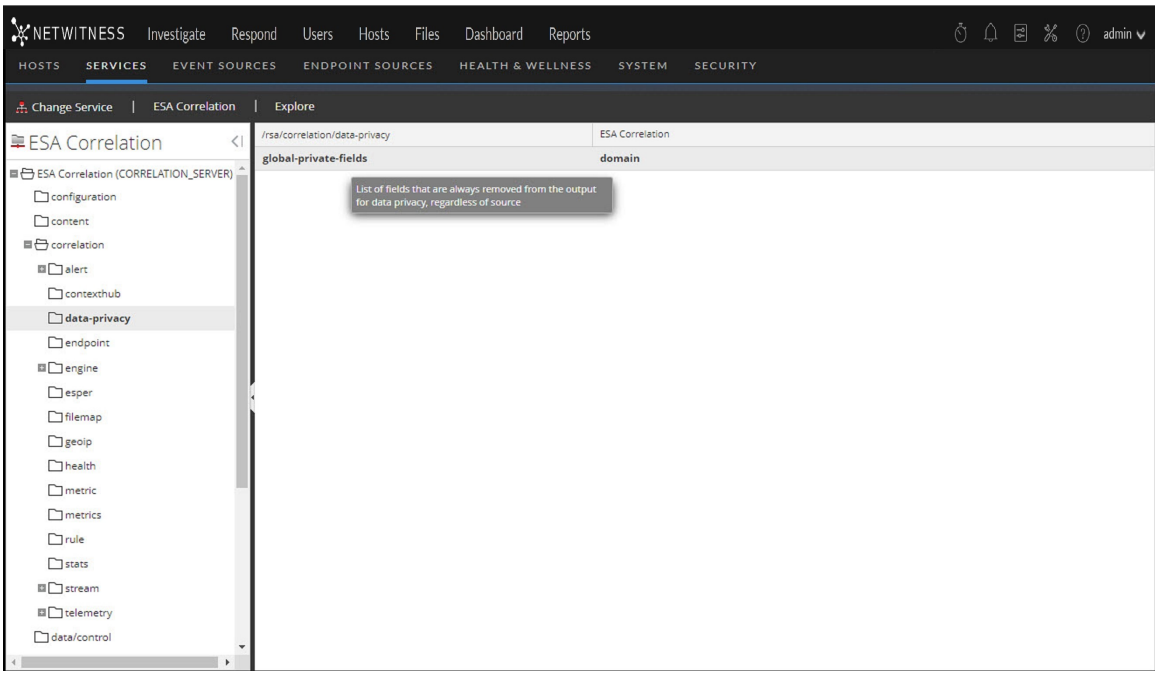
Because the `ipv4` value is the same as the `ip_src` value, `ipv4` contains and displays sensitive data.

How to Remove Sensitive Meta Keys Globally from All Alerts

Note: This procedure applies only to ESA Correlation Rules in NetWitness Platform 11.4 and later versions.

For data privacy reasons, it may be necessary to remove some sensitive meta keys from the alert output globally, regardless of the data source. In the ESA Correlation service, you can set the **`global-private-fields`** parameter to remove the meta keys from all alert output.

1. Go to **(Admin) > Services**, and in the Services view, select an ESA Correlation service and then select **> View > Explore**.
2. In the Explore view node list for the ESA Correlation service, select **correlation > data-privacy**.
3. In the **`global-private-fields`** parameter, add the sensitive meta keys that you want removed from all alerts.



The changes are effective immediately.

For more information on the strategy and benefits of obfuscating data, see the *Data Privacy Management Guide*.

ESA Rule Types

This topic describes each type of ESA rule, when to use them and the permissions each role has with them. The following table lists each type, describes it, and explains when to use it.

Rule Type	Description	When to Use
RSA Live ESA	RSA Live has a catalog of ESA rules that you can download and modify to run in your network.	Download RSA Live ESA rules to leverage rules that are already built. Modify the configurable parameters to customize to meet your requirements.
Rule Builder	In the rule builder, you define rule criteria in an easy-to-use interface.	Use the rule builder to create your first rules. You choose many of the rule conditions from lists.
Advanced EPL	With the Event Processing Language (EPL), you define rule criteria by writing a query.	Use advanced EPL rules to define rule criteria in the EPL syntax.
Endpoint Rule Bundle	An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. The rules in this bundle only apply to NetWitness Endpoint.	If you have NetWitness Endpoint, you can configure risk scoring to identify suspicious files and hosts. To turn on risk scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. For instructions, see "Deploy Endpoint Risk Scoring Rules on ESA" in the *ESA Configuration Guide*. For complete information on configuring NetWitness Endpoint, see the *NetWitness Endpoint Configuration Guide*.

Sample Rules

Sample Rule Builder rules come with NetWitness and appear in the Rule Library. Use sample rules to get comfortable working with rules before creating your own. You can safely edit and deploy these sample rules.

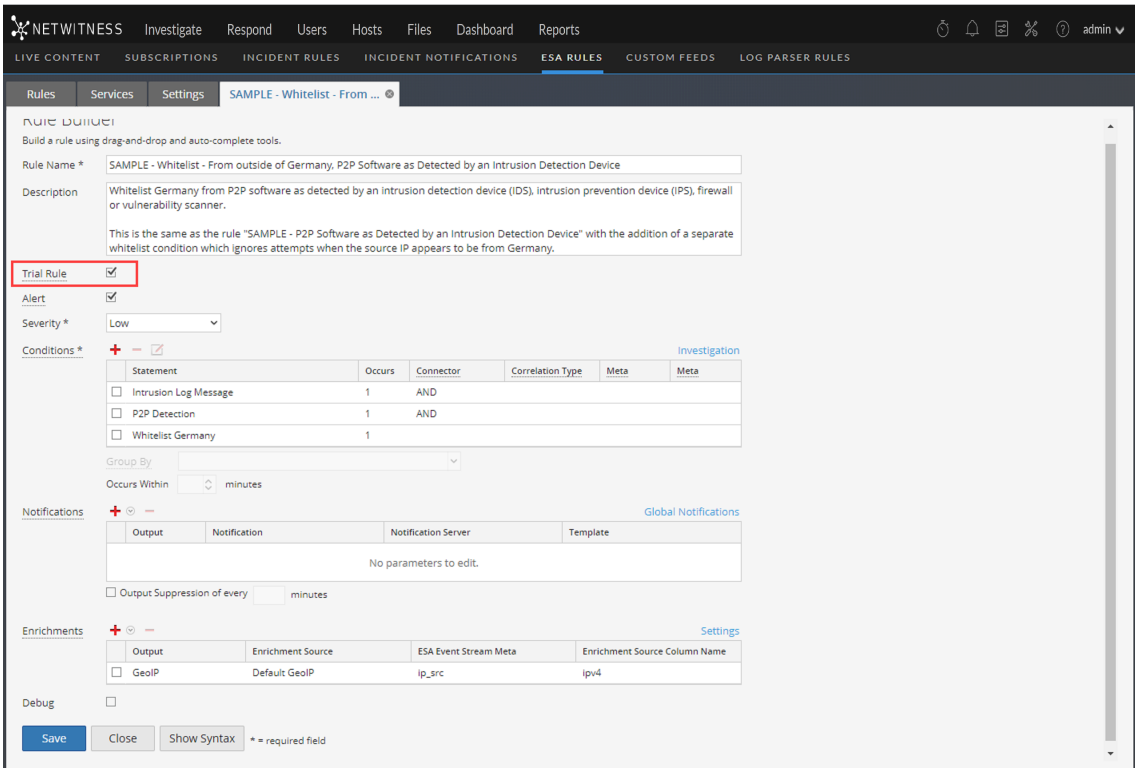
Endpoint Risk Scoring Rules Bundle

An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. These rules appear in the Rule Library with the sample rules. Endpoint risk scoring rules only apply to NetWitness Endpoint. You can add the Endpoint Risk Scoring Rules Bundle to an ESA rule deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) in the ESA Rule Deployment.

Trial Rules Mode

For any type of rule, you can select the Trial Rule setting as an additional safeguard. Trial rules get disabled if they exceed a memory threshold set by the administrator. Run a rule in trial mode to monitor memory usage and to disable the rule automatically if it uses more memory than the threshold allows.

The following figure shows the Trial Rule setting in the Rule Builder.



ESA Permissions

This topic lists all ESA permissions and shows which permissions are assigned to each pre-configured NetWitness role. User access is restricted based on roles and permissions assigned to roles.

- Administrators
- Operators
- Analyst
- Security Operations Center (SOC) Managers
- Malware Analysts (MA)
- Data Privacy Officer

There are four permissions for ESA:

- **Access Alerting Module**: Is required for any permission
- **View Rules**: Allows view-only permission for rules in the Rule Library
- **View Alerts**: Allows view-only permission for alerts ESA generates
- **Manage Rules**:Allows you to view, create, edit, and delete rules

The following table lists permissions for ESA and the roles to which they are assigned. Use this table to see how each role can work with rules and alerts.

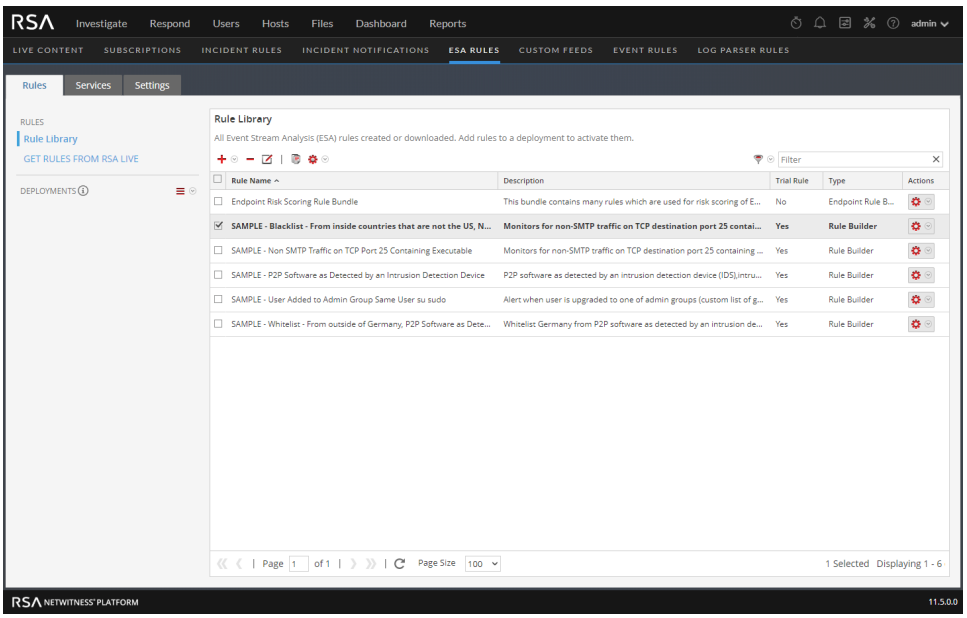
Permission	Administrators	Operators	Analysts	SOC Mgrs	MAs	DPOs
Access Alerting Module	Yes	Yes	Yes	Yes		Yes
View Rules	Yes	Yes		Yes		Yes
View Alerts	Yes		Yes	Yes		Yes
Manage Rules	Yes	Yes		Yes		Yes

For more information on roles and permissions, see the *System Security and User Management Guide*.

Practice with Sample Rules

NetWitness comes with sample rules so analysts can become familiar with how rules look before they create their own rules. Use the sample rules to become familiar with the Rule Builder and to practice editing and deploying a rule.

Sample rules are installed in the Rule Library, which contains every rule you download or create. The following figure shows sample rules in the Rule Library.



These are the available sample rules:

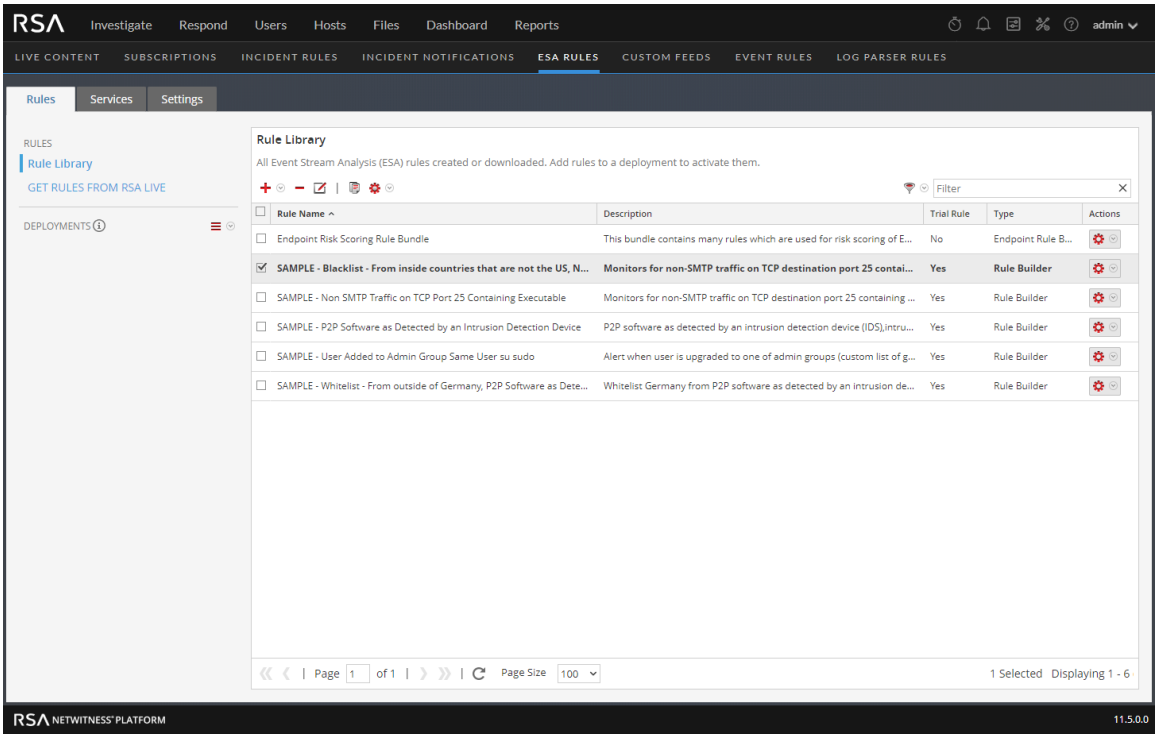
- SAMPLE - Blacklist - From inside countries that are not the US, Non SMTP Traffic on TCP Port 25 Containing Executable
- SAMPLE - Non SMTP Traffic on TCP Port 25 Containing Executable
- SAMPLE - P2P Software as Detected by an Intrusion Detection Device
- SAMPLE - User Added to Admin Group Same User su Sudo
- SAMPLE - Whitelist - From outside of Germany, P2P Software as Detected by an Intrusion Detection Device.

Each name begins with SAMPLE to distinguish the rules that are installed with NetWitness from the rules you download and create.

Rule Library

The Rule Library shows the following information for a rule:

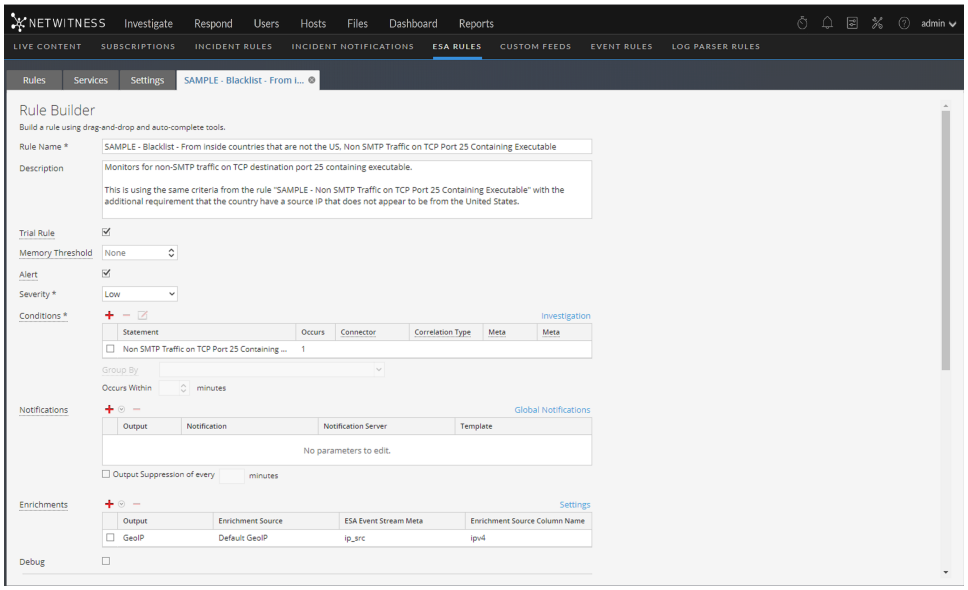
- **Name** summarizes the data or events the rule collects.
- **Description** explains the rule in more detail, although only the beginning shows in the Rule Library.
- **Trial Rule** indicates if trial mode is enabled or disabled for the rule.
- **Type** shows the origin of the rule, built in Rule Builder or Advanced EPL, downloaded from RSA Live, or Endpoint Rule Bundle.



Practice with Sample Rules

1. Go to **(Configure) > ESA Rules**.
   The ESA Rules view is displayed with the Rules tab open.

2. In the **Rule Library**, double-click a sample rule or select a sample rule and click .
   The rule is opened in Rule Builder.



3. To practice with a sample rule, refer to the following topics for detailed descriptions and procedures:
   - To familiarize yourself with the Rule Builder user interface, see Rule Builder Tab for a description of each field.
   - To learn how to edit a rule, see Add a Rule Builder Rule for a step-by-step procedure.
   - To deploy sample rule, see Deploy Rules to Run on ESA to learn how to associate the rule with an ESA service.

After you practice with sample rules, you will be able to download, create, and deploy your own rules.

Working with Trial Rules

The ESA Correlation service is capable of processing large volumes of disparate event data from Concentrators. However, when working with ESA Correlation rules, it is possible to create rules that use excessive memory. This can slow your ESA service or even cause it to shut down unexpectedly. To ensure that rules do not use excessive memory, you can enable them as trial rules. You should disable the trial rule setting only after testing the new rule in your environment during times of both normal and peak network traffic.

You can set a global threshold of the percentage of memory that trial rules may use. If that configured memory threshold is exceeded, all trial rules are disabled automatically. To configure the memory threshold, see "Change Memory Threshold for All Trial Rules" in the *ESA Configuration Guide*.

For suggestions on creating more efficient rules, see "Best Practices for Writing Rules" in Best Practices.

By default, new rules and RSA Live rules that you import are configured as trial rules. As a best practice, when you edit an existing rule, select the Trial Rule option, which allows you to deploy the rule with an added safeguard.

Note: Run a rule as a trial rule long enough to assess the performance during normal and peak network traffic.

Deploy Rules as Trial Rules

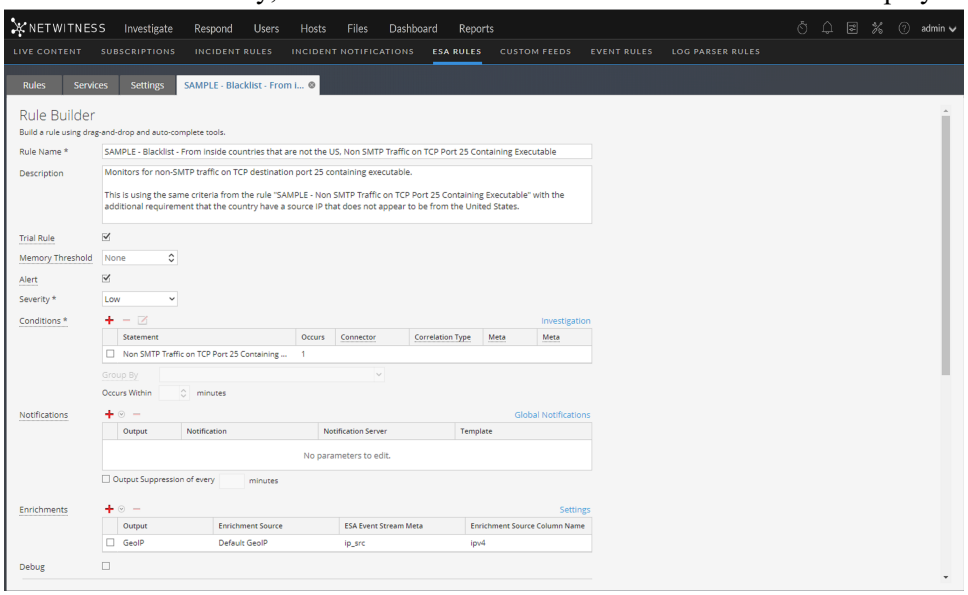
This topic explains to administrators how to enable trial rules when creating new rules or editing rules. Trial rules are automatically disabled if a specified total JVM memory utilization threshold is exceeded.

In NetWitness 11.4 and later, ESA trial rules no longer change status after an upgrade or deployment.

For example, if you change the status of a trial rule to disabled ((Configure) > ESA Rules > Services tab) and redeploy the ESA rule deployment ((Configure) > ESA Rules > Rules tab), the trial rule remains disabled.

1. Go to **(Configure) > ESA Rules**.
   The Configure ESA Rules view is displayed with the Rules tab open.

2. From the Rule Library, choose to add or edit a rule. The rule builder is displayed in a new tab.



3. To make a new or existing rule a trial rule, select the **Trial Rule** checkbox.
4. Add the rule conditions or modify the rule as needed. For instructions on editing rules, see Add Rules to the Rule Library.
5. Click **Save**.
6. Ensure that trial rules are enabled for your ESA and that you are satisfied with the thresholds configured for trial rules.
   The memory threshold is set in the configuration file. To configure it, see "Change Memory Threshold for All Trial Rules" in the *ESA Configuration Guide*.
   - The threshold is configured per ESA and is a percentage of Java Virtual Memory.
   - The configuration parameter, `fatal-percentage`, has a default value of 90.
7. Optionally, you can set up the policies in Health and Wellness to send you an email notification if the total JVM memory utilization threshold is exceeded.

The next time you deploy the rule, it runs in trial rule mode.

Note: If a trial rule is disabled, you will need to go to the **(Configure) > ESA Rules > Services** tab to re-enable the trial rules. For more instructions on re-enabling trial rules on a service, see View ESA Stats and Alerts.

Add Rules to the Rule Library

This topic explains how to add each type of rule to the rule library. You must add a rule to the Rule Library before you can deploy it. Permission to manage rules is required for all tasks in this section. To add rules, you can download them from ESA Live, create a rule via the Rule Builder, or write advanced EPL rules.

For more details on each of these procedures, see:

- Download Configurable RSA Live ESA Rules
- Add a Rule Builder Rule
- Add an Advanced EPL Rule

In addition to deploying a rule, you can edit, duplicate, import, export, and remove a rule in the Rule Library. For details on these procedures, see Working with RulesWorking with Rules

Download Configurable RSA Live ESA Rules

This topic explains how to download configurable rules from the NetWitness Live Content Management System so you can customize them to meet your needs.

RSA Live contains a catalog of rules. Each rule has configurable parameters so you can customize the rule for your environment. If RSA Live has a rule to detect events that you want to detect in your network, download the rule to save time. You can edit the configurable parameters and save the rule in your Rule Library. For detailed information about each rule, including whether the rule is for logs, packets, or both, see "RSA ESA Rules" at the following link:
https://community.netwitness.com/t5/netwitness-platform-threat/rules/ta-p/677884

This is an example of how each RSA Live ESA rule is described on RSA Live:

Rule Name	Description
Logins across Multiple Servers	Detects logins from the same user across 3 or more separate servers within 5 minutes. The time window and number of unique destinations are configurable.

As the name shows, the rule looks for logins across multiple servers. The description explains the rule criteria in more detail and specifies which parameters you modify.

Note: When a rule description includes a configurable parameter, the default setting for the parameter is used. In the example rule, the description states 5 minutes. However, the time window is configurable so 5 is the default number of minutes.

Prerequisites

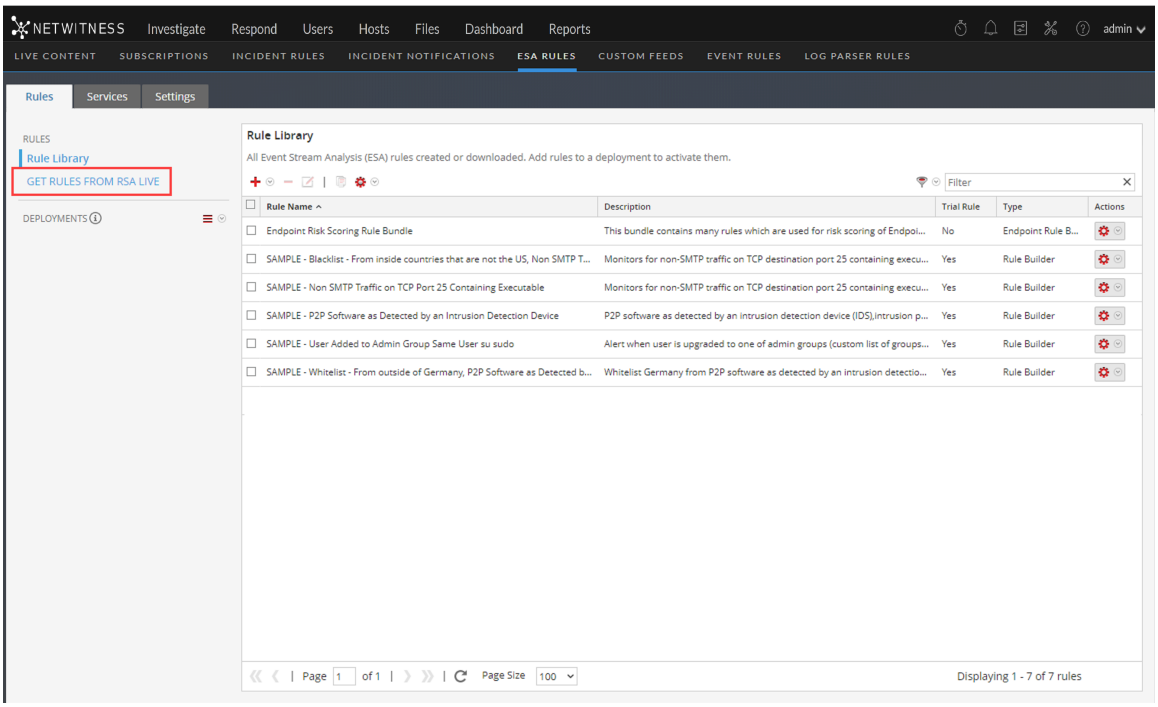
These are the prerequisites for downloading configurable RSA Live ESA rules;

- Have permission to manage rules
- Create a Live Account. See the *Live Services Management Guide* for details.
- Set up Live on NetWitness. See the *Live Services Management Guide* for details.
- Update your meta keys. See "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*.

Download RSA Live ESA Rules

Caution: Before you deploy the latest Live ESA rules, update your meta keys. See "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*.

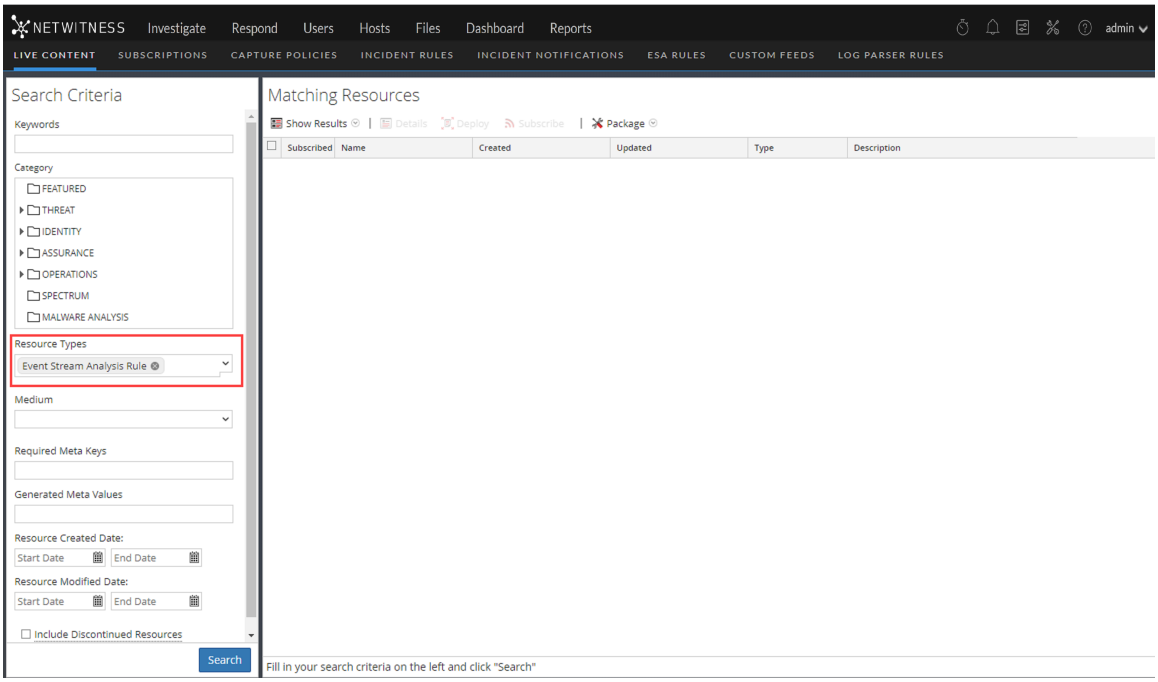
1. Go to  **(Configure) > ESA Rules**.
   The Rules tab is displayed.



2. In the options panel, click **Get Rules from RSA Live**.

   The Live Content Search view is displayed. (Alternatively, you can go to **(Configure) > Live Content**.)

3. In **Search Criteria**, for **Resource Types** select **Event Stream Analysis Rule**.

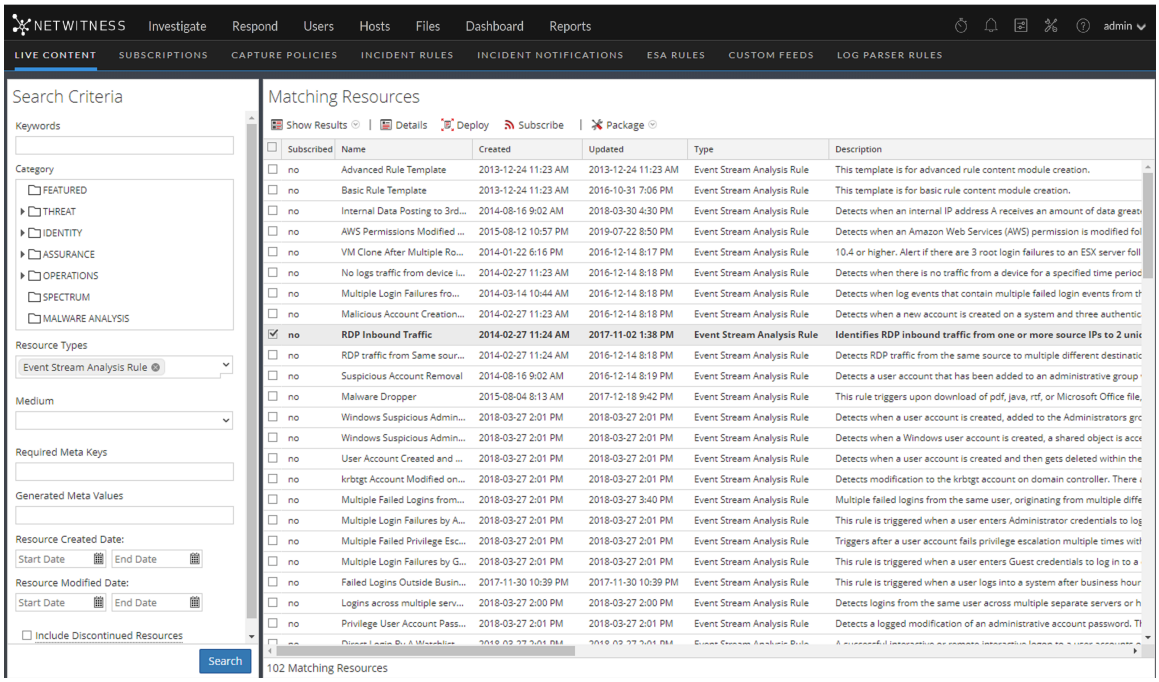


4. Specify any of the following criteria to find a rule to configure for your environment. For detailed information about each rule, including whether the rule is for logs, packets, or both, see "RSA ESA

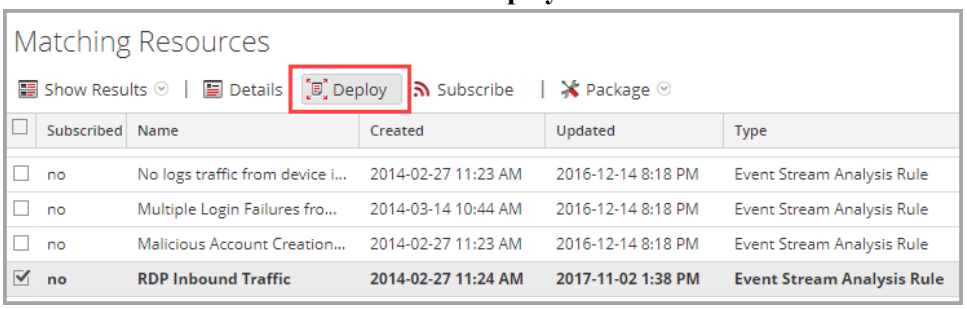
Rules" at the following link: https://community.netwitness.com/t5/netwitness-platform-threat/rules/ta-p/677884
For a detailed description of the search criteria, see "The Live Search View" in the *Live Services Management Guide*.

- Keywords
- Category
- Resource Types (Event Stream Analysis Rule)
- Medium (Log, Log and Packet, or Packet)
- Required Meta Keys
- Generated Meta Values
- Resource Created Date
- Resource Modified Date
- Include Discontinued Resources

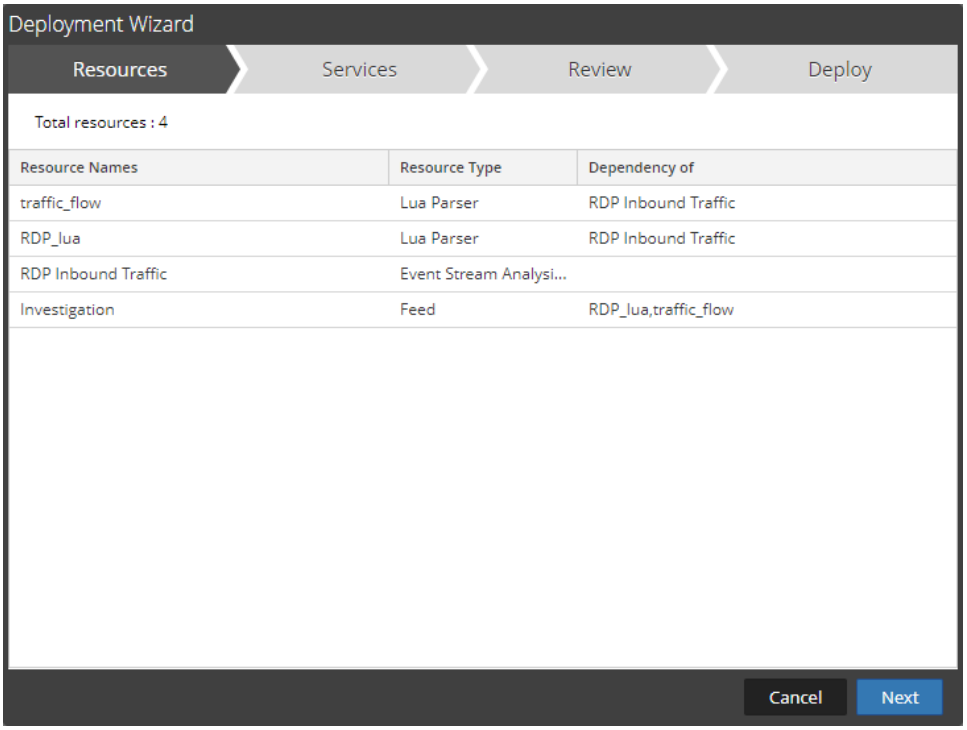
5. Click **Search**. Rules that match the search criteria are displayed in Matching Resources.



6. Select each rule to download and click **Deploy**.



The Deployment Wizard is displayed.



7. Follow the steps in the wizard. If you need more information, see "Deploy Resources in Live" in the *Live Services Management Guide*.

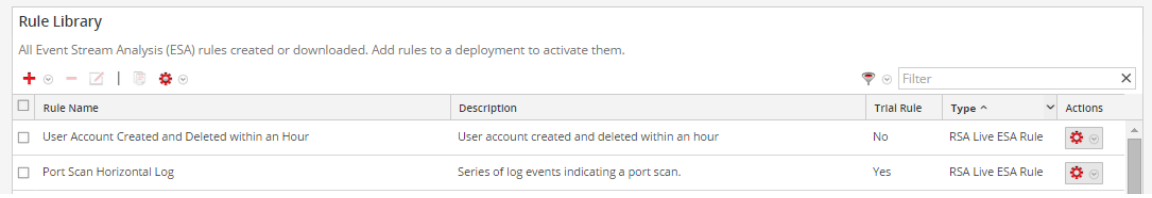
When you finish the steps in the wizard, the selected rules are displayed in the Rule Library.

Customize an RSA Live ESA Rule

This topic explains how to configure parameters in an RSA Live ESA rule. When you download an RSA Live ESA rule, the rule appears in the Rule Library which includes the following columns:

- Rule Name
- Description
- Trial Rule

- Type
- Actions



The type is RSA Live ESA Rule.

Prerequisites

- Administrator, Operator, SOC Manager, or DPO role permissions are required.
- Rules must be downloaded to the Rule Library.

Configure Parameters for an RSA Live ESA Rule

1. Go to **(Configure) > ESA Rules > Rules** tab.
2. In the **Rule Library**, double-click an RSA Live ESA Rule or select the rule and click . The RSA Live ESA Rule tab is displayed.
3. (Optional) Change the following fields:
   - Rule Name
   - Description
   - Trial Rule (Enabled by default. NetWitness recommends you run a rule as a trial rule long enough to assess the performance during normal and peak network traffic.)
   - Alert (This option applies to 11.3 and later.) Select Alert to send an alert to Respond. Clear the checkbox if you do not want to send an alert to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.
   - Severity
   - Notifications
   - Enrichments
4. To configure the rule for your environment, in the **Parameters** section replace the default in the **Value** Column.



5. Click **Save**.

Add a Rule Builder Rule

Each ESA rule is designed to detect something in your network and to generate an alert for it:

- User activity that is not allowed, such as attempting to download software that is not sanctioned
- Suspicious behavior, such as mass audit clearing
- Known malicious threats, such as worm propagation or a password-cracking tool

There are two methods to design a rule in ESA:

- **Rule Builder** is an easy-to-use interface. You provide a meta key and value, then select choices from lists to complete the criteria.
- **Advanced EPL** allows you to write queries in the Event Processing Language. You must know EPL syntax.

If you know EPL, you can use either method. If you do not know EPL, you should use Rule Builder. These topics explain the Rule Builder.

Step 1. Name and Describe the Rule

This topic provides instructions to identify a rule, indicate if it is a trial rule and assign a severity level. When you add a new rule, the first information to provide is a unique name and description of what the rule detects. After you save the rule, this information is displayed in the Rule Library.

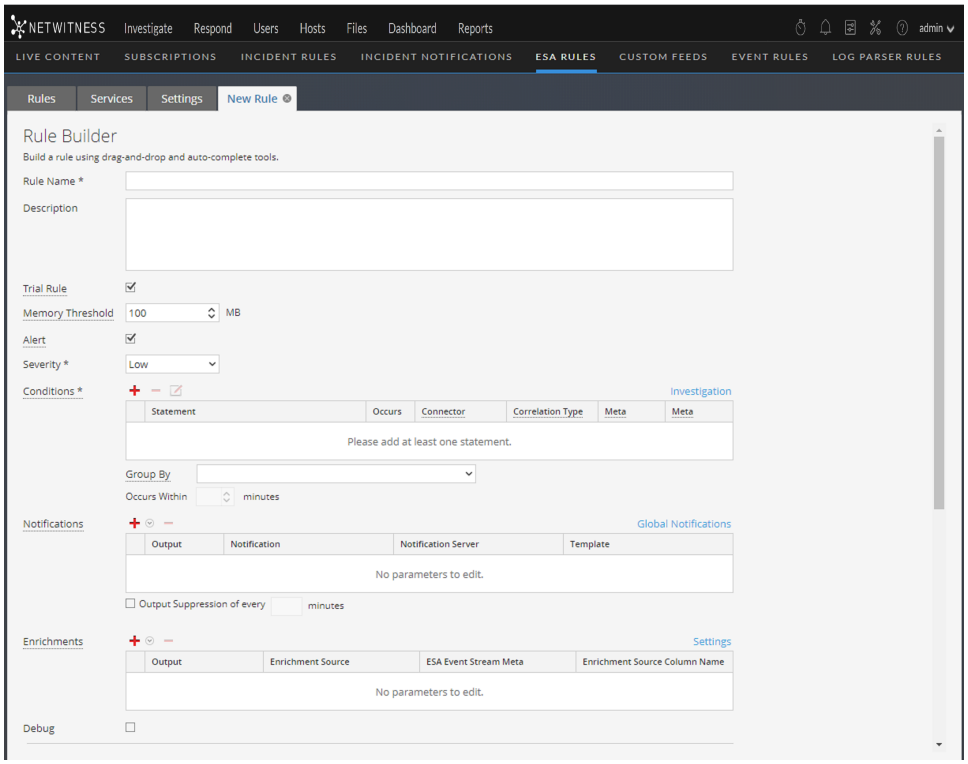
Prerequisites

You must have permission to manage rules. See ESA Permissions.

Name and Describe a Rule

1. Go to  **(Configure) > ESA Rules > Rules** tab.

2. In the **Rule Library**, select  **> Rule Builder**.
   The New Rule tab is displayed.



3. Type a unique, descriptive name in the **Rule Name** field.
   This name will appear in the Rule Library so be specific enough to distinguish the rule from others.
4. In the **Description** field, explain which events the rule detects.
   The beginning of this description will appear in the Rule Library.
5. By default, new rules are configured as a Trial Rule. A trial rule automatically disables the rule if all trial rules collectively exceed the memory threshold. If you are editing an existing rule, you can select **Trial Rule** to safely test the rule edits.
   Use trial rule mode as a safeguard to see if a rule runs efficiently and to prevent downtime caused by running out of memory. For more information, see Working with Trial Rules.
6. (This option applies to 11.5 and later.) Enter a **Memory Threshold** for a rule that uses memory, such as a rule that contains windows or pattern matching. If the configured memory threshold is exceeded, the rule gets disabled individually and an error is displayed for that rule on the  (Configure) > ESA Rules > Services tab. The Memory Threshold option works for trial rules and non-trial rules. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.
7. (This option applies to 11.3 and later.) Select **Alert** to send an alert to Respond. Clear the checkbox if you do not want to send an alert to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.
8. For **Severity**, classify the rule as Low, Medium, High or Critical.

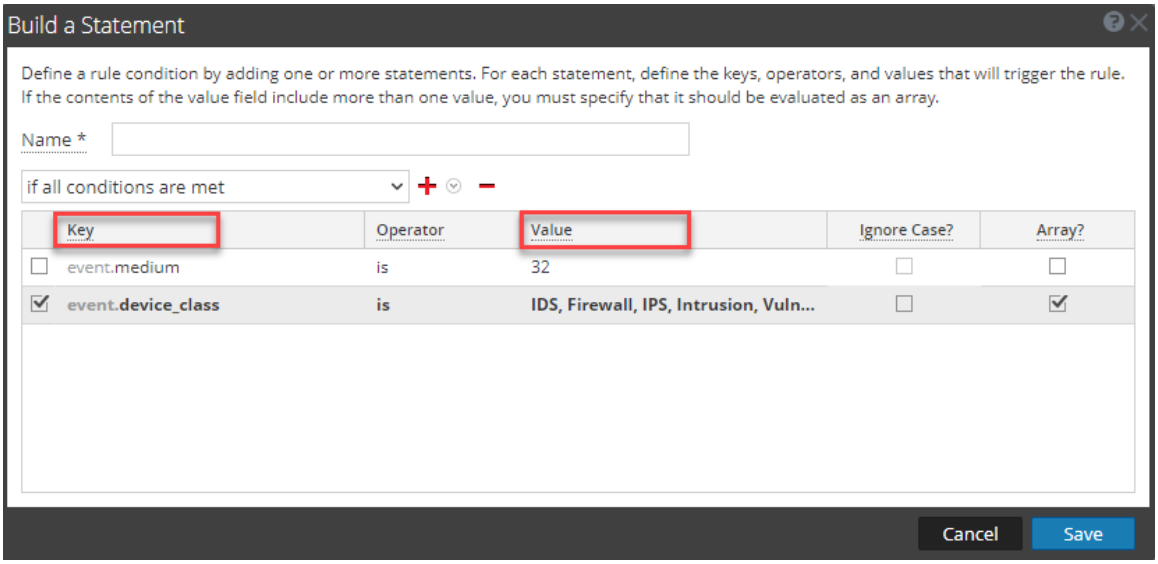
Step 2. Build a Rule Statement

This topic provides instructions to define rule criteria in Rule Builder by adding statements. A statement is a logical grouping of rule criteria in the Rule Builder. You add statements to define what a rule detects.

Example

The following graphic shows an example of a Rule Builder statement.

Every statement contains a key and value. Then, you build logic around the pair by selecting an option in each other field.



Prerequisites

To build a rule statement, you must know the meta key and the meta value.

For a complete list of meta keys, go to **(Configure) > ESA Rules > Settings > Meta Key References**.

Meta entities are not currently supported, such as:

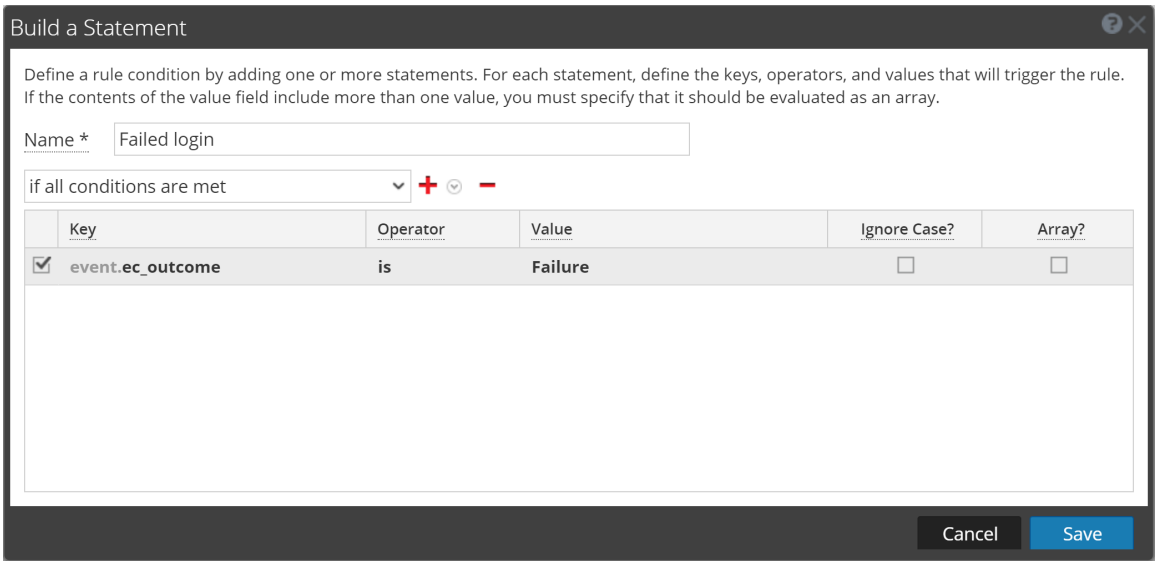
```
fullname.all
eth.all
ip.all
ipv6.all
port.src.all
port.dst.all
dir.path.all
org.all
geoip.all
port.all
```

```
domain.all
email.all
filename.all
directory.all
checksum.all
param.all
context.all
attack.all
analysis.all
compromise.all
inv.all
outcome.all
ec.all
user.all
host.all
client.all
```

Caution: If you add meta entities to your rule, they cannot get data from the data sources, so they do not trigger alerts.

Build a Rule Statement

1. Go to  **(Configure) > ESA Rules**.
   The Rules tab is displayed by default.
2. In the **Rule Library**, click  **> Rule Builder** or edit an existing Rule Builder rule.
   The Rule Builder view is displayed.
3. In the **Conditions** section, click .
   The Build Statement dialog is displayed.



4. **Name** the statement. Be clear and specific. The statement name will appear in the Rule Builder.
5. From the drop-down list, select which circumstances the rule requires:
   - if **all conditions** are met
   - if **one of these conditions** are met
6. Specify the criteria for the statement:
   a. For **Key,** type the name of the **Meta Key**.
   b. For **Operator** specify the relationship between the meta key and the value you will provide for it. The operator that you use depends on the metadata type.
   The choices are: is, is not, is not null, is greater than (>), is greater than or equal to (>=), is less than (<), is less than or equal to (<=), is one of (For array type meta), is not one of (For array type meta), contains, not contains, begins with, ends with
   c. Type the **Value** for the meta key.
   Do not add quotes around a value. Separate multiple values with a comma.
   d. The **Ignore Case?** field is designed for use with string and string array values. By choosing the **Ignore Case** field, the query will treat all string text as a lowercase value. This ensures that a rule that searches for the user named Johnson would trigger if the event contains "johnson," "JOHNSON," or "JoHnSoN." (For best rule performance, only use the **Ignore Case** option when necessary.)
   e. The **Array?** field indicates if the contents of the Value field represent one or more than one value.
   Select the Array checkbox if you entered multiple, comma-separated values in the **Value** field. For example, "ec_activity is Logon, Logoff" requires you to select the Array checkbox.
7. To use another meta key in the statement, click , select **Add Meta Condition** and repeat step 6.
8. To add a whitelist, click  and select **Add Whitelist Condition**.

9. To add a blacklist, click  and select **Add a Blacklist Condition**.
10. To save the statement, click **Save**.

To Add a Whitelist

You use a whitelist to ensure that specified entities are excluded from triggering the rule. Whitelists can be based on geographic location, in-memory enrichment, or Context Hub list sources. For example, if you want to create a rule that only triggers for IP addresses outside of the US, you can create a whitelist of US IP addresses.

1. After you add a meta condition, click  and select **Add Whitelist Condition**.
2. In the **EnterWhitelist Name** field, select an enrichment source. Any in-memory enrichment, Context Hub list, or a named window in Esper can be used as the source for a whitelist.
3. For the subcondition:
   a. If you used a GeoIP source for the whitelist, ipv4 is automatically entered for the subcondition. Enter the meta value for the corresponding value field. For example, enter *ipv4 is ip_src* to ensure the GeoIP records are selected based upon the ip_src being found in the GeoIP lookup database. In addition, if you used a GeoIP source for the whitelist, you might want to add a subcondition to specify the geographic region to exclude from the rule results. For example, to specify that the country code must be USA, enter "*CountryCode is US*".
   b. If you used a Context Hub list for the whitelist, select a column name from the list, then select an operator and enter the meta value for the corresponding value field.

Note: An *Entity* is a specified piece of meta, such as IP address, MAC address, user, host, domain, file name, or file hash.

To Add a Blacklist

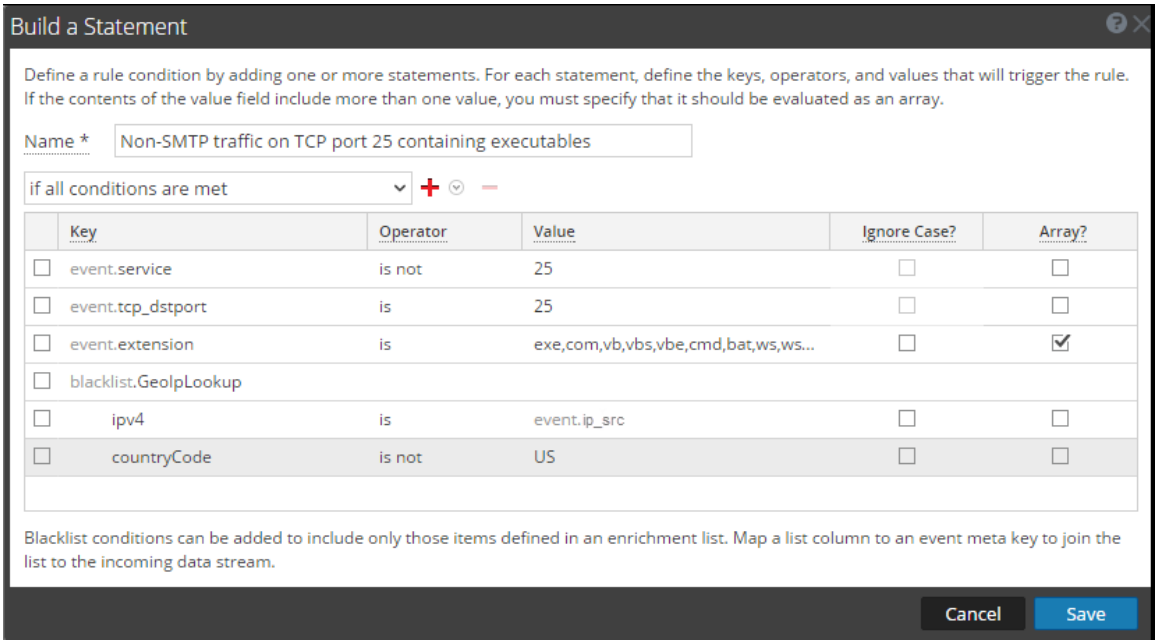
You use a blacklist to ensure that specified entities trigger the rule. Blacklists can be based on geographic location, in-memory enrichment, or Context Hub list sources. For example, you can specify that the rule only includes results from Germany.

1. After you add a meta condition, click  and select **Add Blacklist Condition**.
2. In the **Enter Blacklist Name** field, select an enrichment source. Any in-memory enrichment, Context Hub list, or a named window in Esper can be used as the source for a blacklist.
3. For the subcondition:
   a. If you used a GeoIP source for the blacklist, ipv4 is automatically entered for the subcondition. Enter the meta value for the corresponding value field. For example, enter ipv4 is ip_src to ensure the GeoIP records are selected based upon the ip_src being found in the GeoIP lookup database. In addition, if you used a GeoIP source for the blacklist, you might want to add a subcondition to specify the geographic region to include in the rule results. For example, to specify that the rule only includes results for Germany, enter "*CountryCode is DE*".

b. If you used a Context Hub list for the blacklist, select a column name from the list, then select an operator and enter the meta value for the corresponding value field.

Example: Blacklist

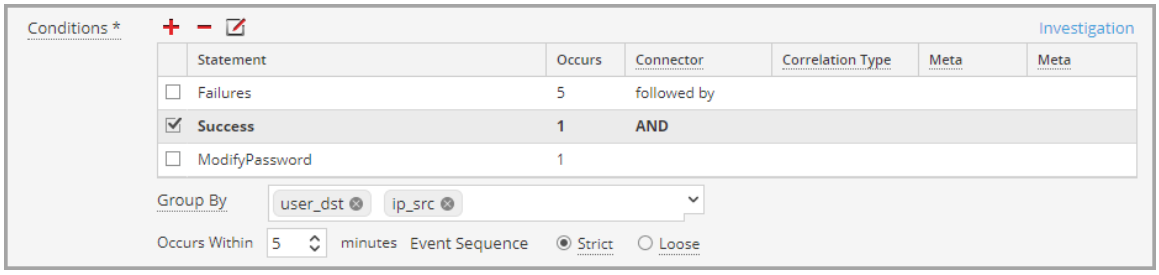
The following statement shows a blacklist statement for a rule that monitors for non-SMTP traffic on TCP destination port 25 containing an executable from countries that are outside of the United States.



Statement	Description
service is not 25	The traffic is not SMTP traffic.
tcp_dstport is 25	The traffic is running on TCP port 25.
extension is exe, com,vb,vbs,vbe,cmd,bat,ws,wsf,src,sh	The file extension is an executable.
GeoIpLookup	The blacklist is based on a GeoIPLookup source.
ipv4 is ip_src	The GeoIP records are selected based upon the ip_src being found in the GeoIP lookup database.
countryCode is not US	When looking up the IP address Event.ip_src in the GeoIP database, the record it returns does not contain "US" in the countryCode field.

Example: Strict Pattern Matching and Using the *Is Not Null* Operator

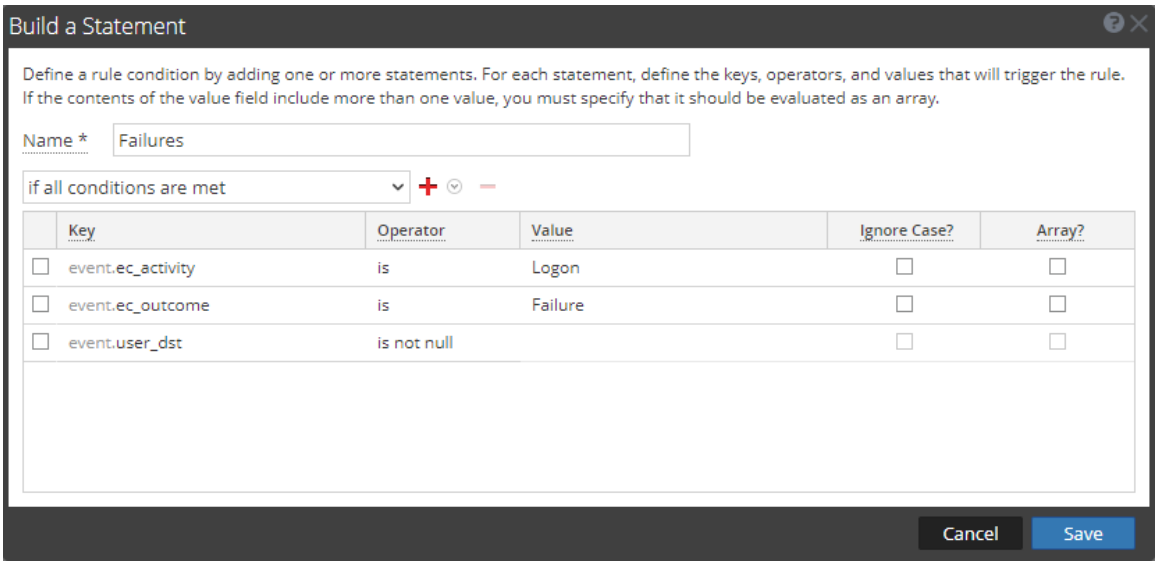
The following example uses the ability to exclude null values and create a strict pattern match to ensure that it returns the expected rule results. The following conditions make up the rule:



Rule Condition	Description
Failures	This condition searches for five failed logins with a "followed by" connector, meaning that the condition (Failures) must be followed by the next condition (Success).
Success	This condition searches for one successful login.
ModifyPassword	This condition searches for an instance where the password is modified.
GroupBy: user_dst, ip_src	The GroupBy field ensures that all the previous conditions are grouped by the user_dst meta (the user destination account) and ip_src. This is important to the construction of the rule because the rule attempts to find a case where a user has attempted to log into the same destination account multiple times, finally logged in successfully, and then changed the password. Grouping by ip_src ensures that the user logged in from the same machine attempted to log into an account multiple times. The rule may give unexpected results if you do not group the results.
Occurs within 5 minutes	The time window for the events to occur is five minutes. If the events occur outside of this time window, the rule does not trigger.
Event Sequence: Strict	The event sequence is configured for a strict pattern match. This means that the pattern must match exactly as it is specified with no intervening events. Strict pattern matching allows you to ensure that the Esper engine only generates alerts for rules that exactly match the pattern you want to find. For example, a common rule might be to search for five failed logins followed by a successful login. If you select a loose pattern match, this rule will trigger if there are any number of successful logins between the failed logins. Since the point of the rule is to find frequent *and* sequential login attempts, a strict match is required to ensure that you get the results you expect.

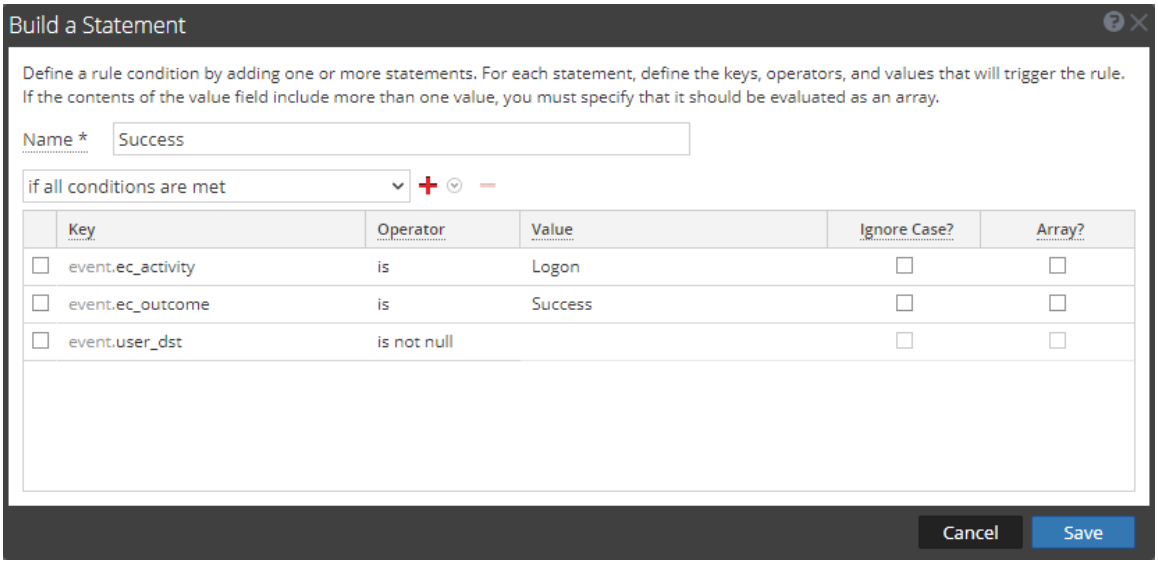
Note: Each of these conditions is explained in further detail in the sections below.

For each condition, a statement is built in the Rule Builder. The following statement makes up the Failures condition:



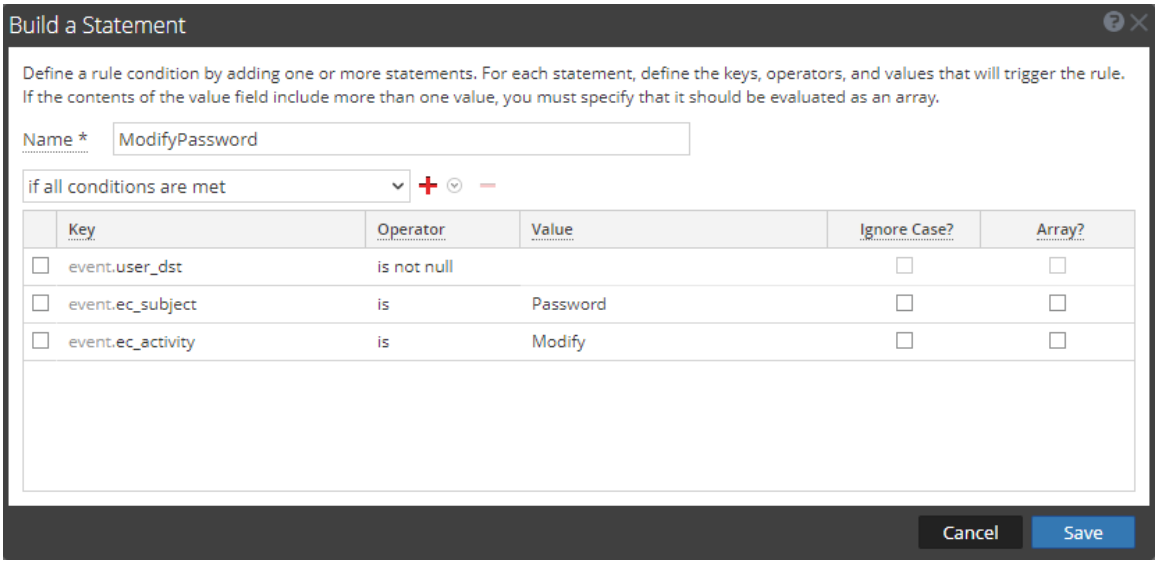
Rule Statement	Description
ec-activity is Logon	Identifies activity that attempts to log on to a system. The **Ignore Case** field is designed for use with string and string array values. By choosing the **Ignore Case** field, the query will treat all string text as a lowercase value. You may want to use this field if you are unsure what case may be used when logging a particular event. For best rule performance, only use the **Ignore Case** option when necessary.
ec_outcome is Failure	Identifies activity outcome logged as "failure."
user_dst is not null	Ensures that the condition is only true if user_dst is populated. The **is not null** operator allows you to ensure that a field returns a value. You may want to use this field when a rule depends on a particular field returning a value. For example, you want to create a rule that identifies the same user attempting to log into the same destination account multiple times (potentially a password-guessing attack). If the field that represents the user destination account is empty, you don't want the rule to trigger. To ensure the field contains a value, you use the **is not null** operator.

The following statement makes up the Success condition:



Rule Statement	Description
ec_activity is Logon	Identifies logon activity.
ec_outcome is Success	Identifies a logon that is successful.
user_dst is not null	Ensures that user destination account field must be populated for the condition to be true.

The following statement makes up the ModifyPassword condition:

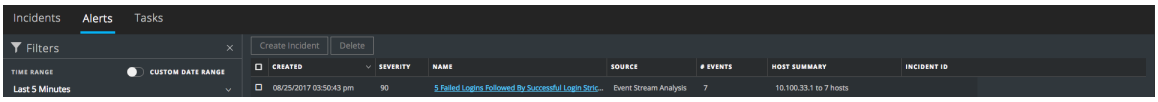


Rule Statement	Description
user_dst is not null	Ensures the user destination account field must be populated for the condition to be true.
ec_subject is Password	Identifies a subject of Password.
ec_activity is Modify	Identifies activity where the password was modified.

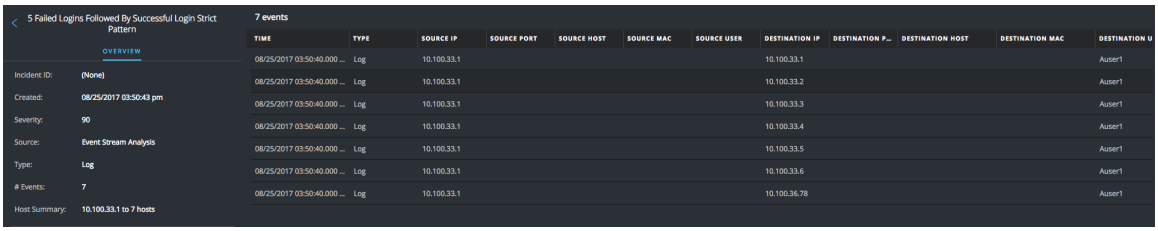
Example Results

When the alert fires for the above example rule, you can see that the rule triggered for seven events, and that each event contains a user. You can also see that the events follow a strict pattern: five failed login events, followed by a successful login event, followed by a modification to the account.

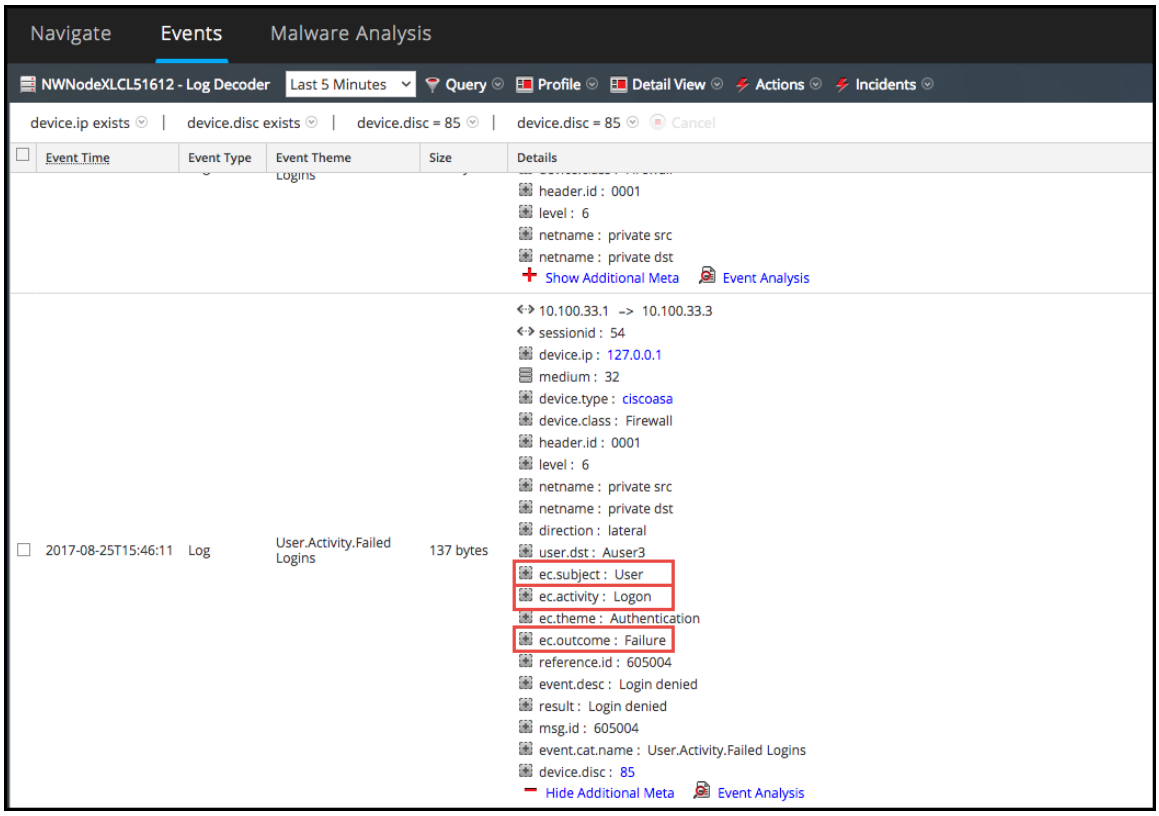
The following figure shows the alert in the Respond Alerts List view.



The next figure shows the events in the alert in the Respond Alert Details view.



Drilling down into the Investigation view by clicking on the source for one of the events, you can see the case for each of the string values.

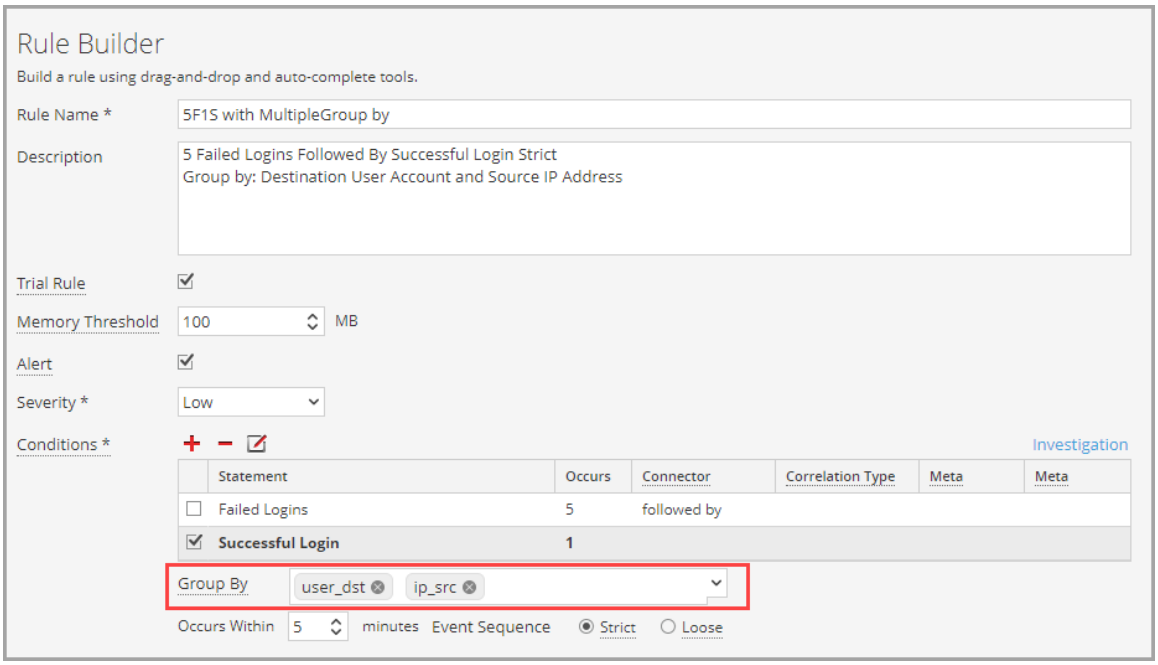


Example: Grouping the Rule Results

The **Group By** field allows you to group and filter rule results. For example, suppose that there are three user accounts; Joe, Jane, and John and you use the **Group By** meta, user_dst. The result will show events grouped under the accounts for Joe, Jane, and John.

You can also group by multiple keys, which can further filter rule results. For example, you might want to group by user destination account and machine to see if a user logged into the same destination account from the same machine attempts to log into an account multiple times. To do this, you might group by user_dst and ip_src.

The following example shows a rule grouped by user_dst and ip_src.

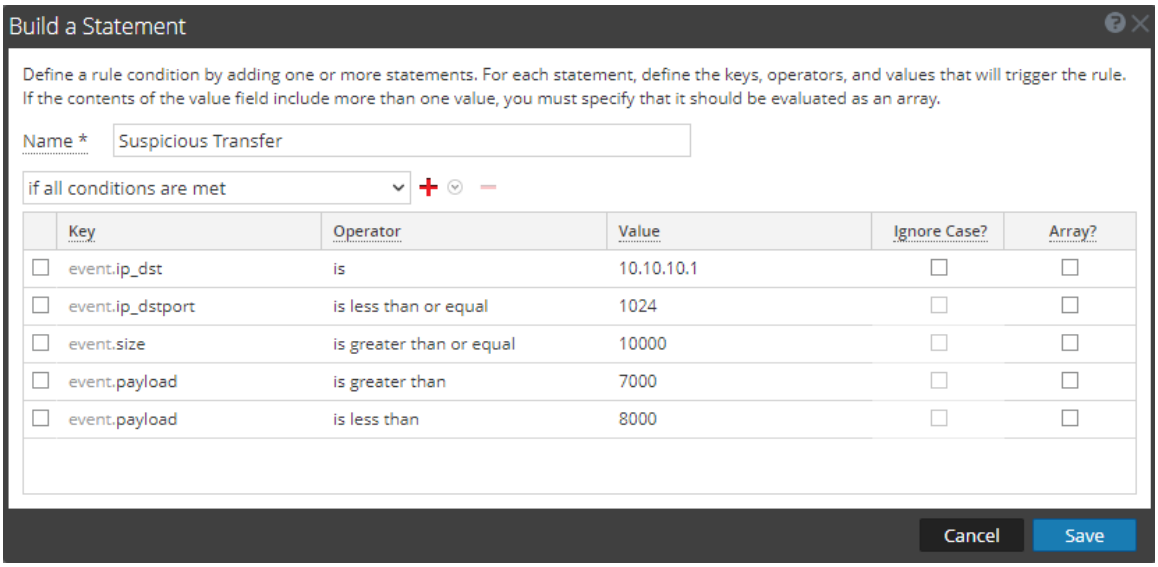


Rule Condition	Description
Failed Logins	Identifies five failed login attempts (must be followed by the next condition; that is, the five failed logins must be followed by a successful login).
Successful Login	Identifies one successful login.
Group By: user_dst and ip_src	Groups the rule results by user_dst (user destination account) and ip_src (IP address of the machine that the user is logging in from). This allows the rule to look for a user logged in from the same machine to the same destination account, resulting in a much more targeted rule result.
Occurs within 5 minutes with a strict pattern match	The events must occur within five minutes, and the pattern matching is strict, meaning it must follow the pattern exactly for the rule to trigger.

Example: Working with Numeric Operators

Numeric operators allow you to write rules against numeric values, such as specifying that a value is greater than, less than, or equal to a specific value. This is useful particularly for cases where you might want to specify a numeric threshold, that is, *payload is greater than 7000*.

The following example attempts to identify a data transfer to a particular destination through the common ports where the transfer size is high and the payload is in a suspicious range.



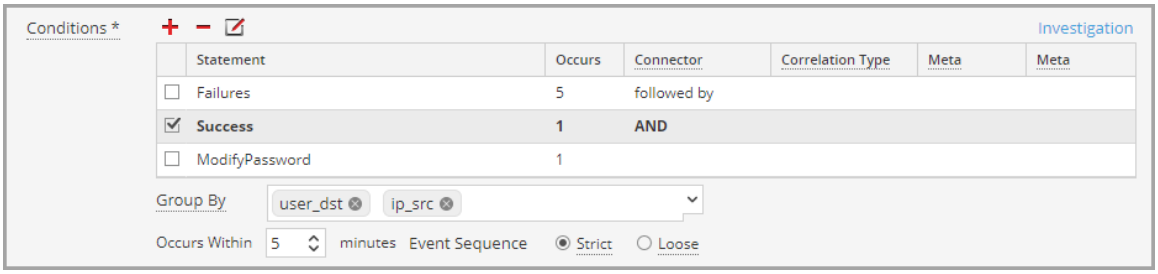
Rule Statement	Description
ip_dst is 10.10.10.1	The destination port is 10.10.10.1.
ip_dstport is greater than or equal to 1024	The destination port is in a commonly used port range, 1024 or greater.
size is greater than or equal to 10000	The size of the transfer is 10000 or greater, which is a suspiciously large transfer.
payload is greater than 7000	The payload is between 7000 and 8000, which is a suspiciously large payload.
payload is less than 8000	The payload is between 7000 and 8000, which is a suspiciously large payload.

Step 3. Add Conditions to a Rule Statement

This topic provides instructions to add conditions, such as specifying a certain time frame, to a rule statement. When you build a statement, you specify what a rule detects. You add conditions to make further stipulations, such as how many times or when the criteria must occur.

Example

The following graphic shows an example of the conditions for Rule Builder statements. Combined, the statements and conditions comprise the rule criteria.



This rule detects 5 failed logon attempts followed by one successful logon, which could be the sign that someone has hacked into user account. This is the criteria for the rule:

A. 5 failed logons are required.

B. 1 successful logon must follow the failures

C. A password was changed.

D. All events must occur within 5 minutes.

E. Group alerts by user (user_dst), because steps A and B must be performed on the same user destination account. Also, group by machine (ip_src) to ensure that the user logged in from the same machine attempts to log into an account multiple times.

F. The match is a strict pattern, meaning that the pattern must match exactly with no intervening events.

Add Conditions to a Rule Statement

1. In the **Conditions** section, select a statement and click .
2. For **Occurs**, enter a value to specify how many occurrences are required to meet the rule criteria.
3. If you have multiple statements, in the **Connector** field select a logical operator to join one statement to another:
   - followed by
   - not followed by

- AND
- OR

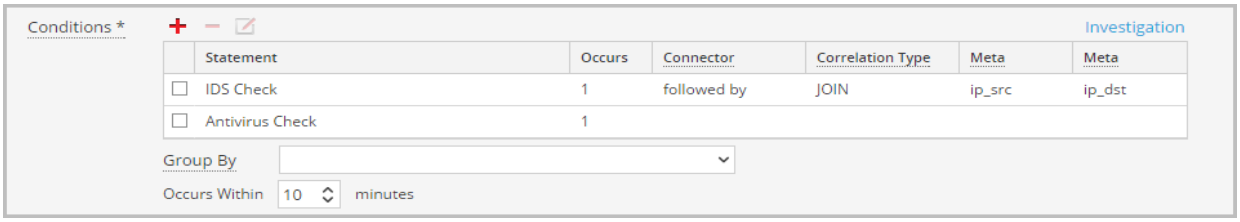
4. **Correlation Type** applies only to **followed by** and **not followed by**. If you choose a correlation type of SAME, select one meta to correlate on, and if you choose a correlation type of JOIN, select two meta to correlate on. You may want to use JOIN if you are trying to correlate on meta from two different data sources. For example, say you want to correlate an AV alert with an IDS alert. See the examples below for a use case where two meta from different sources are joined.
5. If events must happen within a specific timeframe, enter a number of minutes in the **Occurs Within** field.
6. Choose whether the pattern must follow a **Strict** match or a **Loose** match. If you specify a strict match, this means that the pattern must occur in the exact sequence you specified with no additional events occurring in between.

   For example, if the sequence specifies five failed logins (F) followed by a successful login (S), this pattern will only match if the user executes the following sequence: F,F,F,F,F,S. If you specify a loose match, this means that other events may occur within the sequence, but the rule will still trigger if all of the specified events also occur. For example, five failed login attempts (F), followed by any number of intervening successful login attempts (S), followed by a successful login attempt might create the following pattern: F,S,F,S,F,S,F,S,F,S which would trigger the rule despite the intervening successful logins.
7. Choose the fields to group by from the dropdown list. The **Group by** field allows you to group and evaluate the incoming events.
   For example, in the rule that detects 5 failed logon attempts followed by 1 successful attempt, the user must be the same, so user_dst is the **Group By** meta key. You can also group by multiple keys. Using the previous example, you might want to group by user and machine to ensure that the same user logged in from the same machine attempts to log in to an account multiple times. To do this, you might group by user_dst and ip_src.

Example

The following graphic shows an example of the conditions for a rule that allow you to evaluate the same entities across multiple devices so you can accomplish complex use cases. For example, you can create a rule that triggers if an IDS (Intrusion Detection System) alert is followed by an AV(Anti-virus) alert for the same workstation. The work station key is not the same between the two (IDS & AV) sources, so you can perform a JOIN in order to evaluate the different entities.

In the IDS alert, the workstation is identified by the source IP address from the IDS alert, and would be compared to the destination IP address from the AV alert.



This is the criteria for the rule:

A. An IDS alert occurs.

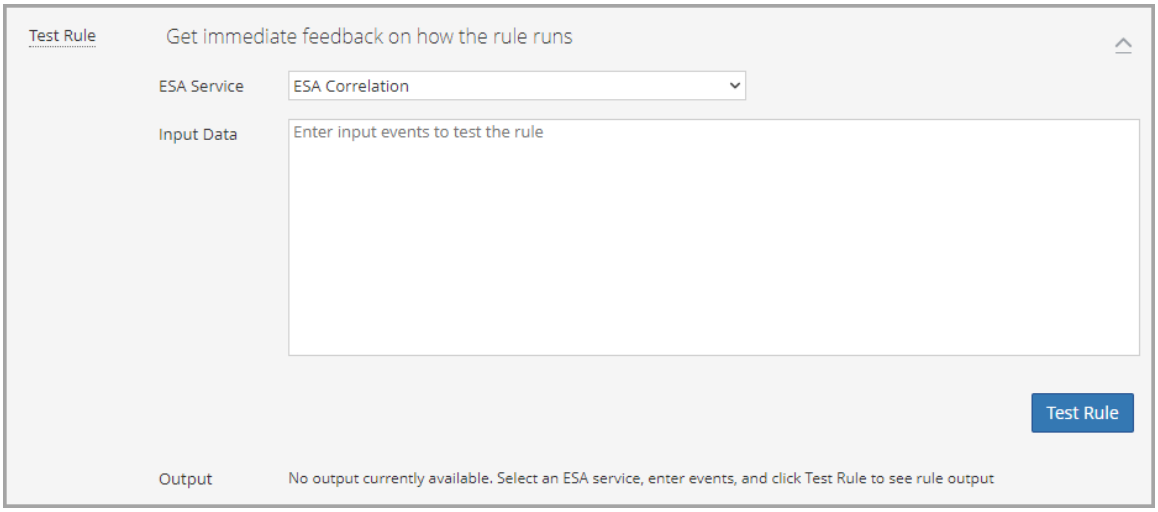
B. The destination IP address from the AV alert and source IP address for the workstation from the IDS alert are joined to allow you to view the same entities across different sources.

C. An Antivirus alert follows the IDS alert.

Validate an ESA Rule

You can confirm that an ESA rule generates the expected alerts by testing the rule logic using JSON input data. You can view the alerts in the output, but this test does not send any alert notifications.

1. If you are not already in the rule, go to **(Configure) > ESA Rules > Rules** tab and in the **Rule Library**, open the ESA rule that you want to test.
2. Scroll down to the **Test Rule** section.



3. In the **ESA Service** field, select the ESA Correlation service to process the rule. Use the same ESA Correlation service that you plan to use in the ESA rule deployment that contains the rule.
4. In the **Input Data** field, enter the input events to test the rule. Download the events from the Investigate view in JSON format, copy the events, and paste them in this field. You can do this from the Investigate > Navigate view or the Investigate > Events view.

   To download the events from the Investigate > Navigate view:

   a. In the main menu, go to **Investigate > Navigate** in a new tab, select a data source, and click **Navigate**.

   b. In the Navigate view, click **Load Values** and click a meta value to filter the events.

c. Save the events as meta in the JSON file format [Save Events > Meta > (name the file) > Export Meta Format: choose JSON].

d. In the toolbar click the **(Jobs)** icon and then click **View Your Jobs**.

e. In the Jobs panel, download your extracted meta, for example: **`investigation-2020-May-19-08-30-20.json`**.

f. Go to back to the **(Configure) > ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.

To download the events from the Investigate > Events view:

a. In the main menu, go to **Investigate > Events** in a new tab.

b. In the Events view, enter a query for the ESA rule test.

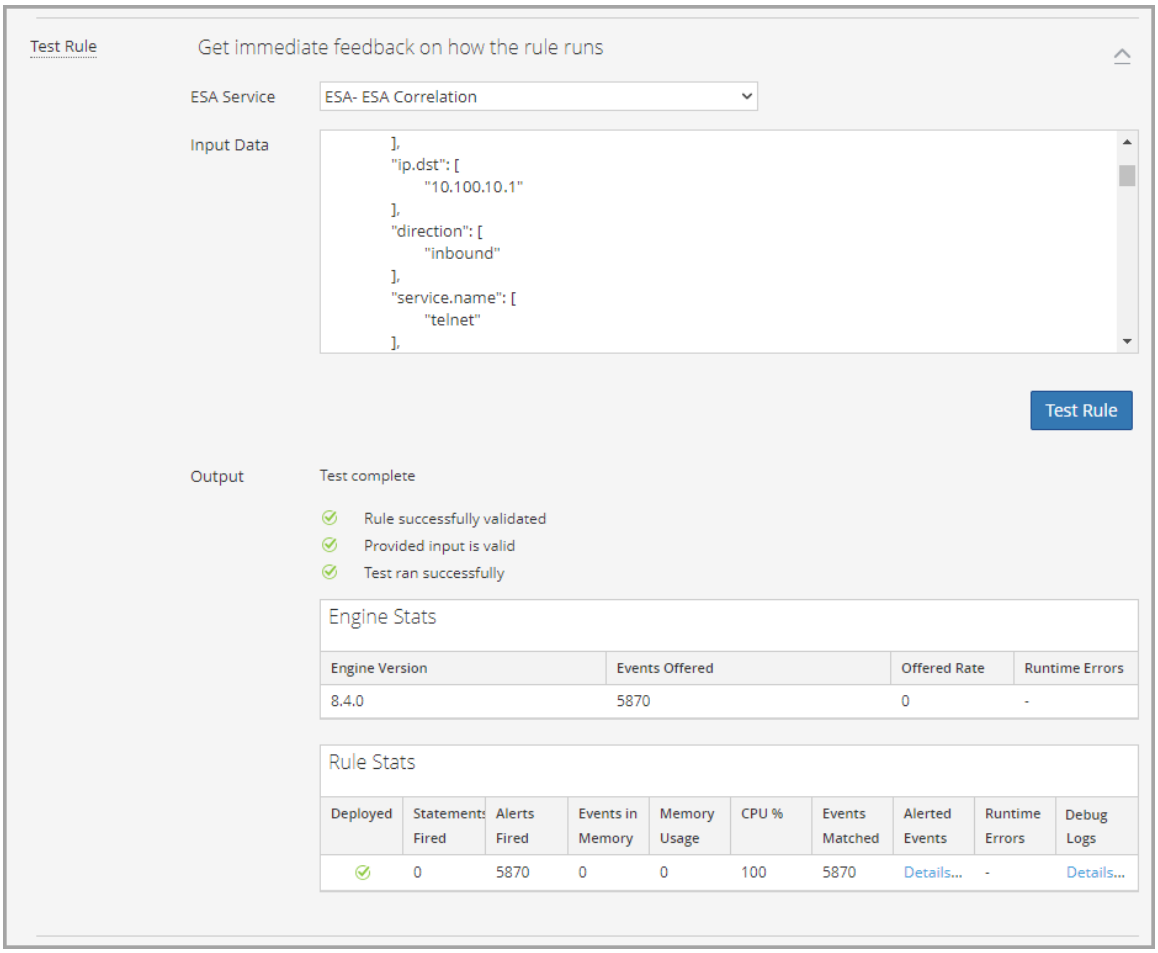
c. Select the events to use and in the **Download** or **Download All** menu, select **Visible Meta as JSON** or **All Meta as JSON**, depending on the size of your selection.

d. In the main menu, go to **Dashboards** and in the toolbar click the **(Jobs)** icon and then click **View Your Jobs**.

e. In the Jobs panel, download your extracted meta, for example: **Concentrator_ALL_EVENTS_ALL_META.json**.

f. Go to back to the **(Configure) > ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.

5. Click **Test Rule**. The **Output** field shows the output of your rule and you can determine if the results meet your requirements.



The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.

Field	Description
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Working with Rules

This topic discusses additional procedures you can perform on rules. You may want to perform any of the following procedures:

- Edit, Duplicate or Delete a Rule
- Filter or Search for Rules
- Import or Export Rules

Edit, Duplicate or Delete a Rule

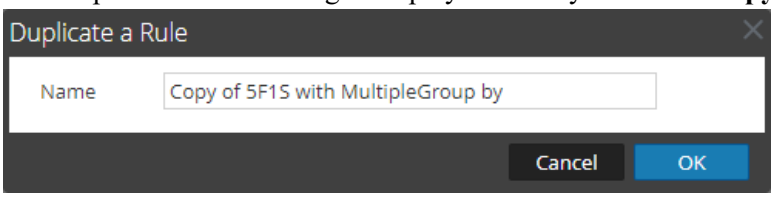
This topic provides instructions to edit, duplicate, or delete an Event Stream Analysis (ESA) rule. When you edit a rule, ESA applies the updated criteria going forward. No changes are made to previously generated alerts.

Edit a Rule

1. Go to **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the **Rule Library**, select the rule you want to edit and click .
   Depending on the rule type, the respective rule tab is displayed.
3. Modify the required parameters.
4. Click **Save**.

Duplicate a Rule

1. In the **Rule Library**, select the rule you want to duplicate and click .
2. The Duplicate a Rule dialog is displayed. The system adds **Copy of** in front of the rule name.

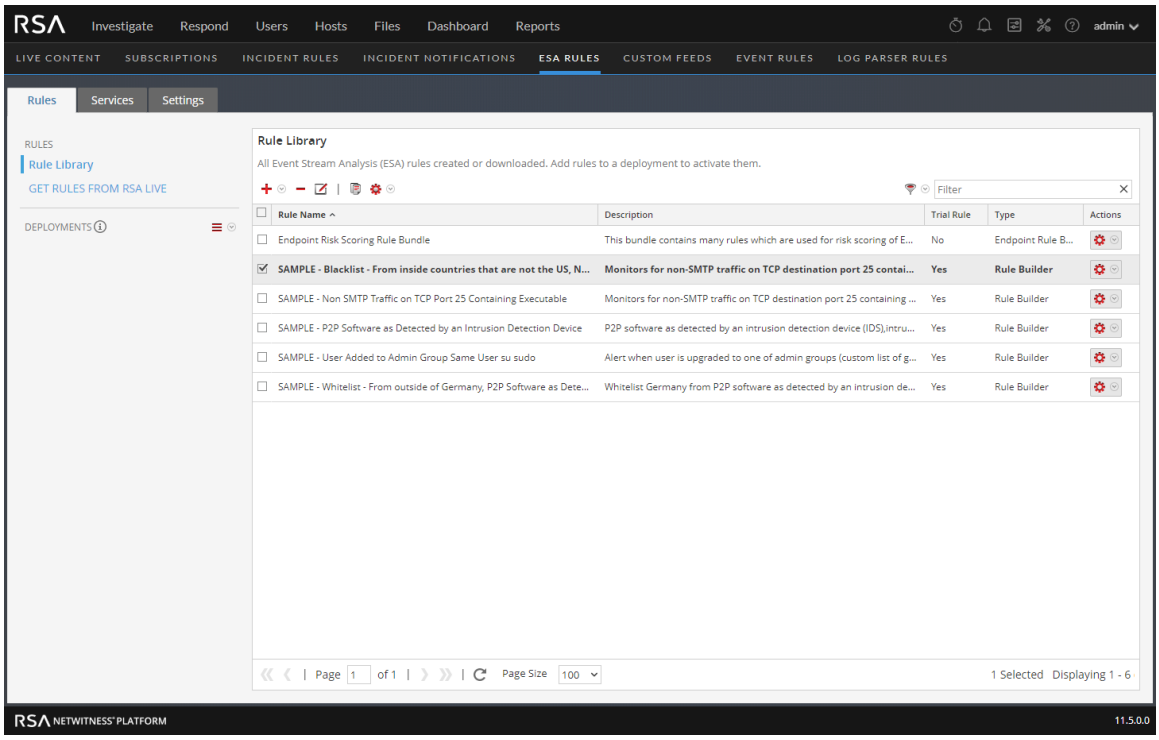


3. In the **Name** field, type a unique name for the duplicate rule and click **OK**.

A duplicate rule with the new name is added to the Rule Library.

Delete a Rule

1. Go to **(Configure) > ESA Rules > Rules**.
   The Rules tab is displayed.



2. In the Rule Library, select one or more rules and click .
   A warning dialog is displayed.
3. Click **Yes**.
   A confirmation message that the rule is deleted successfully is displayed and the selected rule is deleted from the Rule Library.

Filter or Search for Rules

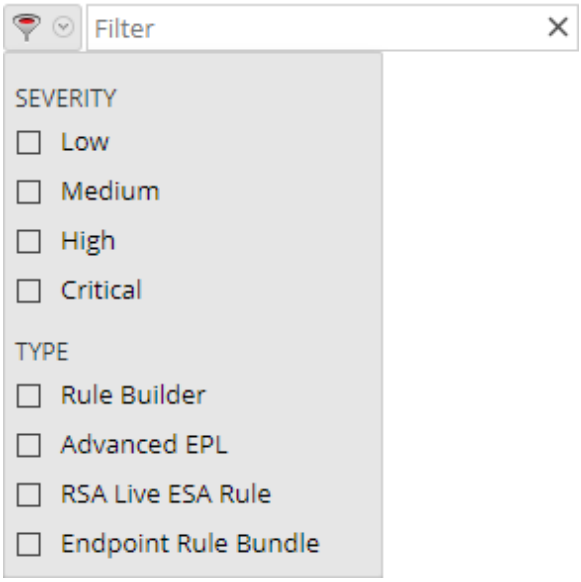
This topic shows analysts how to specify the type of rules that display in the Rule Library.

Prerequisites

Make sure that you understand the Rule Library view components. For more information, see Rule Library Panel.

Filter Rules

1. Go to **(Configure) > ESA Rules**.
   The Rules tab is displayed by default.
2. In the **Rule Library** panel toolbar, click and select the severity and type of rules that you would like to appear in the Rule Library list. The following figure shows the Filter drop-down list.



The selected rule types appear in the list.

Search for Rules

1. Go to **(Configure) > ESA Rules**.
   The Rules tab is displayed by default.
2. In the **Rule Library** panel toolbar, type a rule name in the Filter field.
   The Rule Library panel lists the rules that match the names entered in the Filter field.

Import or Export Rules

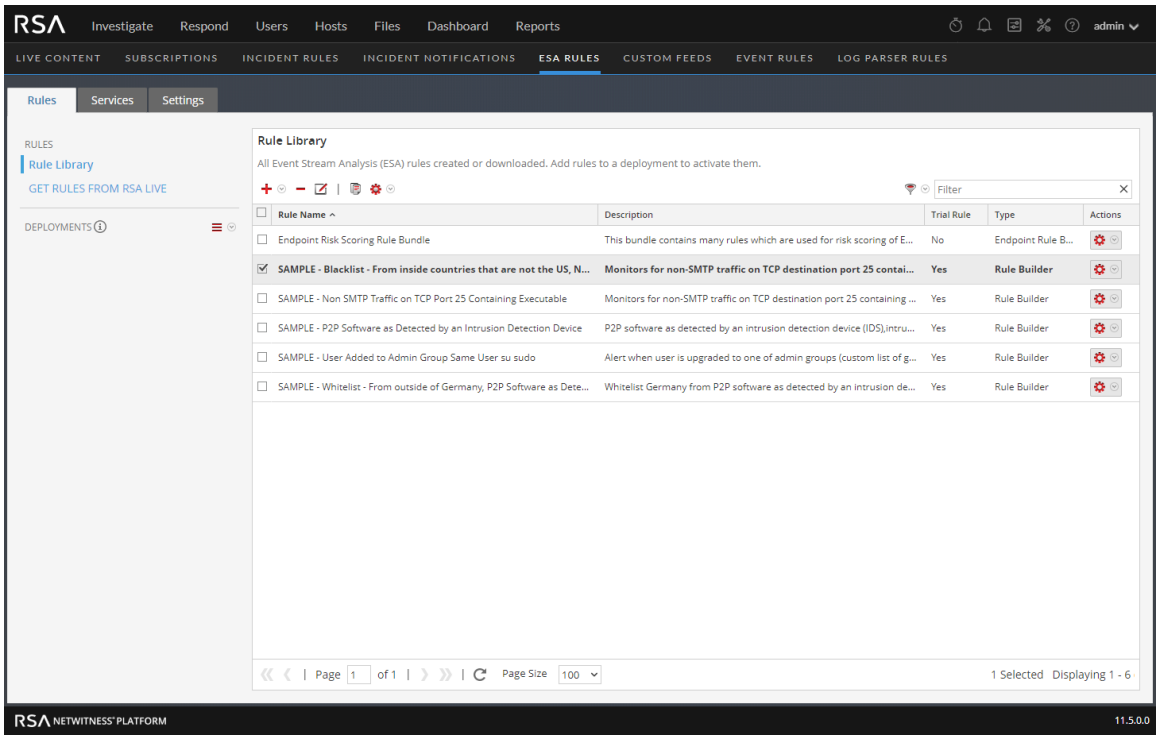
The topic provides instructions to import ESA rules from a NetWitness instance and to export ESA rules to your hard drive so you can keep a local copy.

If you exported a rule in an earlier version of NetWitness, the following conditions apply when you import the rule in version 10.5 or later:

- Exported in version 10.3 – You cannot import rules to version 10.5 or later.
- Exported in version 10.4 – You can import rules to version 10.5 or later.

Import ESA Rules

1. Go to **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.



2. In the **Rules Library** toolbar, select > **Import**.
   The Import ESA Rules dialog is displayed.



3. Click **Browse** to browse and select the file containing the ESA rules.
4. Click **Import**.

Export ESA Rules

1. Select an ESA rule or multiple rules and select > **Export** in the Rule Library toolbar.
   A warning dialog is displayed.
2. Click **Yes**.
   The Export Rules dialog is displayed.
3. In the **Enter File Name** field, type a filename for the file with the ESA rules and click **Export**.
   The file is exported as a binary file to your machine.

Note: The binary file cannot be edited.

Choose How to be Notified of Alerts

This topic explains the different notification methods and how to add a notification method to a rule. Administrator, SOC Manager or DPO role permissions are required for all tasks in this section.

When a rule triggers an alert, ESA can send a notification in the following ways:

- Email
- Syslog
- Script

To configure a notification, you configure these components:

- **Notification Server**: The notification server is the source of the notifications. After you configure a notification server, you can add it to a rule. When the rule triggers an alert, the rule will use that server to send alert notifications.
- **Notifications**: These are the outputs (destinations) of the notifications, which can be email, script, and Syslog. When you design a rule, you can specify the notification for an alert.
- **Templates**: The message format of an alert notification is defined in a template.

If you use an ESA rule that has an enrichment, such as a Context Hub list, you must create a custom template. You can duplicate a default template and adjust it for your enrichment. For more information, see Troubleshoot ESA Rules. For information on creating a custom template, see see "Configure Meta Keys as Arrays in ESA Correlation Rules" in the *System Configuration Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

Alert suppression and alert rate regulation are two features that Event Stream Analysis provides. Alert suppression ensures that multiple emails are not sent out for the same alert. For example, consider a rule to detect failed user logins. If you set the alert suppression to three minutes, you will see only the alerts generated in that time frame. This is fewer than the number of alerts you would see without alert suppression. Some alerts can be duplicates. With alert suppression, emails are not sent for duplicate alerts. This ensures the inbox is not flooded with redundant alert notifications.

Alert rate regulation is a preventive measure to ensure that alerts from misconstrued rules do not flood the system. This ensures that ESA does not send more than the configured limit of emails within one minute.

Notification servers, notifications, and templates are configured in the Administration System view. For more information, see "Configure Notification Servers", "Configure Notification Outputs", and "Configure Templates for Notifications" in the *System Configuration Guide.*

Notification Methods

When a rule triggers an alert, ESA can send a notification in the following ways:

- Email
- Syslog
- Script

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

Email Notifications

ESA Correlation can send notifications to users through email about various system events.

To configure these email notifications, you need to:

- Configure the SMTP email server as an output provider. For instructions, see "Configure the Email Settings as Notification Server" in the *System Configuration Guide*.
- Set up an email account to receive notifications. For instructions, see "Configure Email as a Notification" in the *System Configuration Guide*.
- Configure a template for email notification. For instructions, see "Configure Global Notifications Templates" in the *System Configuration Guide*.

Syslog

Event Stream Analysis can send events and consolidate logs in Syslog format to a Syslog server.

To configure these Syslog notifications, you need to:

- Configure Syslog server settings as an output provider. For instructions, see "Configure a Syslog Notification Server" in the *System Configuration Guide*.
- Configure Syslog message format as an output action. For instructions, see "Configure Syslog as a Notification" in the *System Configuration Guide*.
- Configure a template for Syslog. For instructions, see "Configure Global Notifications Templates" in the *System Configuration Guide*.

Script Alerter

Apart from the alert notifications ESA allows users to run scripts in response to ESA alerts.

Scripts enable you to do custom integration with applications that exist in your environment. For example, if you want to open an incident ticket from an application when a specific alert is triggered, Script Alerter lets you write a script that calls the application API and has ESA invoke it when the specific ESA rule is triggered. You can configure a FreeMarker template to define what details you want to extract from the output of the ESA rule and pass it as command line arguments to the script.

To use the Script Alert, you need to:

- Configure the user identity and other details that are required to execute the script. For instructions, see "Configure Script as a Notification Server" in the *System Configuration Guide*.

- Define the Script. For instructions, see "Configure Script as a Notification" in the *System Configuration Guide*.
- Configure a template for the script. For instructions, see "Configure Global Notifications Templates" in the *System Configuration Guide*.

Add Notification Method to a Rule

This topic tells administrators how to add a notification, such as email, to a rule. ESA uses the notification method when it generates an alert for an event that meets rule criteria.

You add a notification to a rule so ESA can let you know when a rule triggers an alert. Although the notification fields are not required, it is a best practice to add a notification to a rule.

When you add a notification method to a rule, you select the following information:

- Output
- Notification
- Notification Server
- Template

Prerequisites

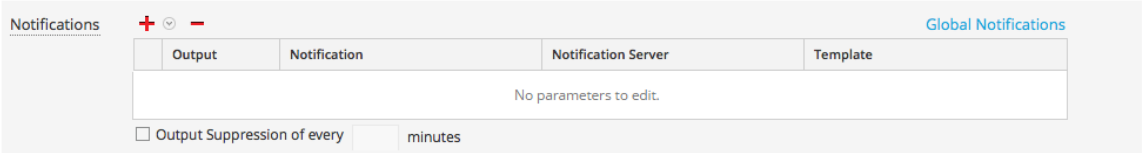
- Your role must have permission to manage rules.
- The rule must exist.
- The notification method must be configured with a supported server and template:

  Go to **(Admin) > System > Global Notifications**.

  For detailed procedures, see the *System Configuration Guide*.

Add a Notification Method to a Rule

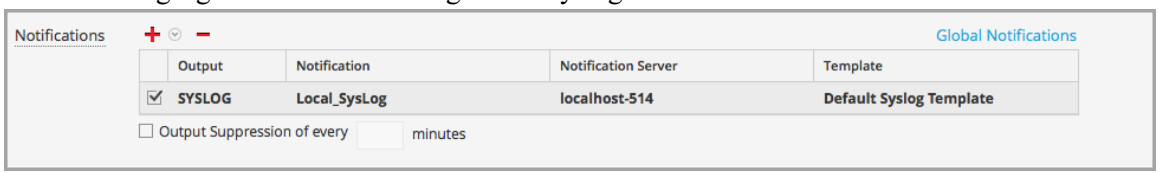
1. Go to **(Configure) > ESA Rules > Rules** tab.
2. In the **Rule Library**, click to add a new rule or select an existing rule and click . Depending on the rule type, the Rule Builder or Advanced EPL tab is displayed. The Notifications section is the same for both tabs.



3. Click and select the **Output** for the alert:

- Email
- SNMP (This option is not supported in NetWitness 11.3 and later.)
- Syslog
- Script

4. Double-click the **Notification** field and select the name of a previously configured output.
   For example, Level 1 Analyst could be the name of an email notification that goes to the L1-Analysts email distribution group.
5. Double-click the **Notification Server** field and select the server that sends the notification.
6. Double-click the **Template** field and select a format for the alert.
   The following figure shows the settings for a Syslog notification.

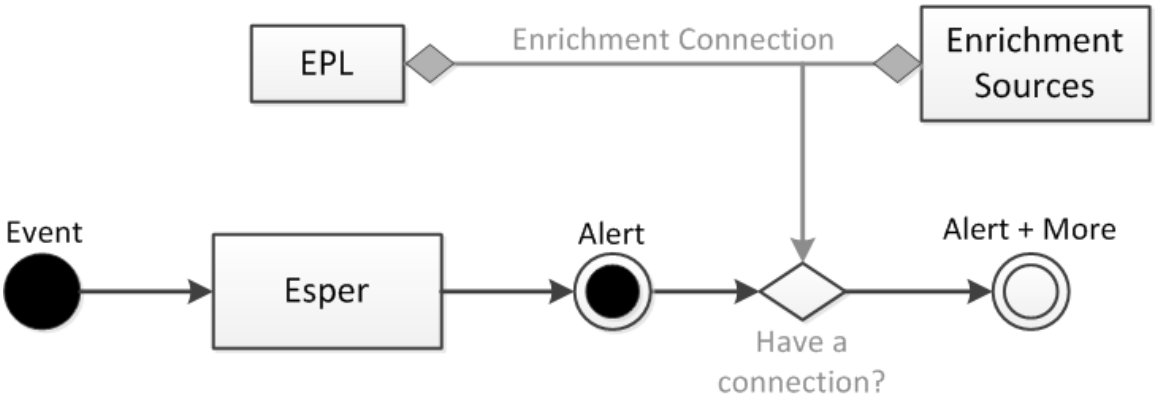


7. If you want to specify frequency, select **Output Suppression**, then enter the number of **minutes**.
8. If you want to add another notification, repeat steps 3-7.
9. Click **Save**.
   When ESA generates an alert for an event that matches the rule criteria, you will be notified of the alert via each notification method added to the rule.

Add a Data Enrichment Source

This topic tells how to add a previously configured enrichment source to a rule. When ESA creates an alert, information from the source gets included in it.

Enrichments provide the ability to include contextual information into correlation logic and alert output. Without enrichments, all information included in an ESA alert is from a Core service. With enrichments, you can request for look ups into a variety of sources and include the results into the outgoing alerts. The following figure illustrates the enrichment feature.



Enrichment configuration is made up of two logical units:

- Enrichment Sources – These are data stores of contextual information.
- Enrichment Connections – These act as connectors between alert meta and source columns.

ESA allows you to make connections between Event Processing Language (EPL) statements and enrichment sources. Once the connections are established, the system joins the selected fields from the alert output with the information in the sources and uses the matching data to enrich the alert that is sent out. ESA can connect with the following sources:

- Esper Named Windows
- MaxMindGeoIP Database

Note: The geoIP enrichment source can neither be created nor deleted. It is provided out of the box to the user.

Example Rule with Enrichments

The following example rule illustrates how ESA enrichments can enhance alerts.

```
@RSAAlert @Name("simple") SELECT * FROM Event(ec_theme='Login Failure')
```

This rule generates an alert for every logon failure and thus if the following (simplified) event stream is received at ESA:

sessionid	ec_theme	username	ip_src	ip_dst	host_dst
1	Login Success	dshrute	23.xx.23x.16		
2	Login Failure	jhalpert	23.xx.23x.16	31.1x.x9.1x8	www.facebook.com

An alert without an enrichment with the following constituent `events` might be generated in response to the second session:

```
{
     "events": [
          {
               "username": "jhalpert",
               "host_dst": "www.facebook.com",
               "ip_dst": "31.1x.x9.1x8",
               "sessionid": 2,
               "ec_theme": "Login Failure",
               "esa_time": 1406148964130,
               "ip_src": "23.xx.23x.16"
          }
     ]
 }
```

The JSON output shows all the information available for inclusion into an ESA notification using an appropriate FreeMarker template. For instance, the template expression `${events[0].username}` would evalute to `jhalpert`.

With enrichments, the same deployment, with the same event stream, can generate the alert shown below.

```
{"events": [
     {
          "username": "jhalpert",
          "host_dst": "www.facebook.com",
          "GeoIpLookup": [
          {
               "city": "Cambridge",
               "longitude": -71,
               "countryCode": "US",
               "areaCode": 617,
               "metroCode": 506,
               "region": "MA",
               "dmaCode": 506,
               "ipv4Obj": "/23.xx.23x.16",
               "countryName": "United States",
               "postalCode": "02142",
               "ipv4": "23.xx.23x.16",
               "latitude": 42,
               "organization": "Verizon Business"
          }
     ],
     "orgchart": [
          {
               "supervisor": "mscott",
               "name": "James Halpert",
               "extension": 3692,
               "location": "Scranton",
               "department": "Sales",
               "id": "jhalpert"
          }
     ],
```

```
    "ip_dst": "31.1x.x9.1x8",
    "sessionid": 2,
    "LoginRegister": [
        {
            "username": "dshrute",
            "ip_src": "23.xx.23x.16"
        }
    ],
    "ec_theme": "Login Failure",
    "esa_time": 1406155218912,
    "ip_src": "23.xx.23x.16"
    }
]}
```

The system pulls contextual data to make the alert more meaningful.

To include the name of the supervisor and the name of the user with the last successful login in the ESA notification, this example includes the following template expressions: `${events[0]["orgchart"][0].supervisor}` gives the name of the supervisor of the employee in the alert and `${events[0]["LoginRegister"][0].username}` gives the name of the user with the last successful logon from the same `ip_src` (using a stream based Named Window).

Enrichment Sources

This topic explains options for adding an external data source to provide additional information in alerts. Enrichment sources provide additional information in alerts. For example, an in-memory table can provide a full name, title, office location, and employee number if a user matches rule criteria. The following types of enrichment sources are available:

- Context Hub List (Preferred)
- In-Memory Table (Ad hoc only)
- GeoIP

Note: Database, Database Connection, Warehouse Analytics, and Recurring In-Memory Tables as enrichment sources are not supported for the ESA Correlation service in NetWitness 11.3 and later.

It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources. You can share Context Hub List enrichment sources across the NetWitness. You can only use the In-Memory Table with ESA. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources.

Note: The geoIP enrichment source can neither be created nor deleted. It is provided out of the box to the user.

Configure a Context Hub List as an Enrichment Source

Note: The information in this topic applies to NetWitness Version 11.1 and later.

This topic provides instructions on how to configure a Context Hub list as an enrichment source for ESA. Once a Context Hub list is added as an enrichment source, analysts can use the configured list as a statement condition when creating an ESA rule. Any changes made to the list from within Context Hub are automatically reflected in the enrichment source in real-time. For example, you could create a list of IP addresses in Context Hub and then use that list as either a blacklist or whitelist as part of a correlation rule condition. Any subsequent changes made to the IP list in Context Hub will be reflected in the enrichment source in real-time, to ensure the correlation rule operates with a constantly updating set of information.

Prerequisites

Before configuring a Context Hub list as an enrichment source, the list must first be created as a data source in Context Hub. Any list created in Context Hub is supported and the lists may contain string or numeric values, including IP addresses. For information on creating a list as a data source in Context Hub, see the *NetWitness Context Hub Configuration Guide*.

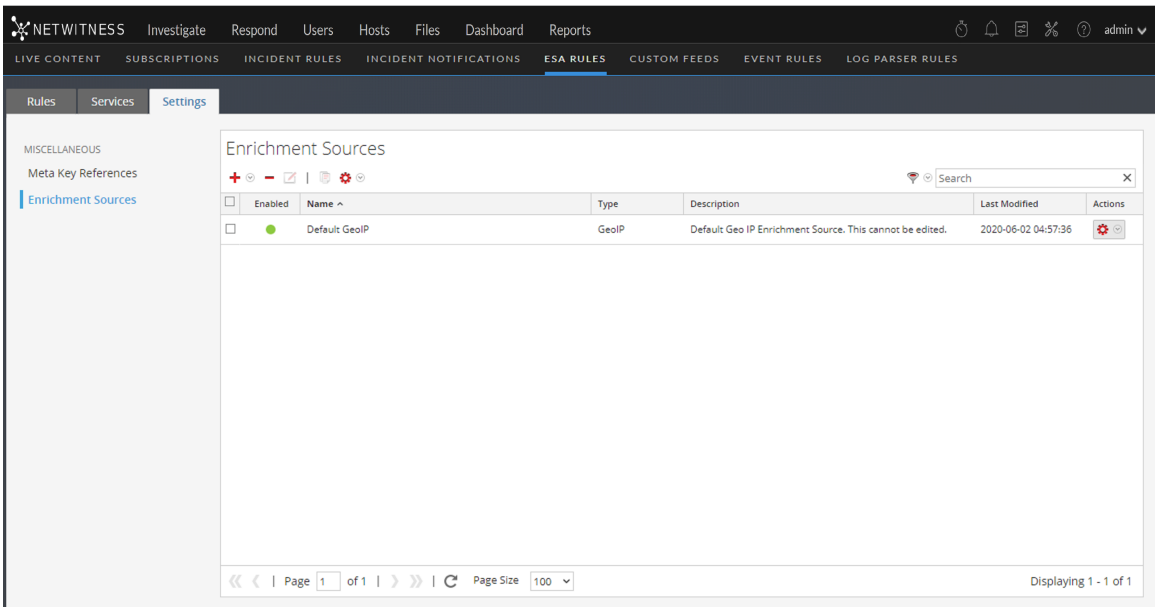
Caution: When creating a Context Hub list for use as an enrichment source, the list name and its field names cannot include any spaces or special characters, or start with a number. If you do not follow this naming convention, when you attempt to add the list as an enrichment source in ESA, an error message will be displayed and you will not be allowed to add the list.

IMPORTANT: If you rename a Context Hub list or recreate the Context Hub list with the same name, update the ESA rules that use that Context Hub list, and then redeploy the ESA rule deployments that contain those rules.

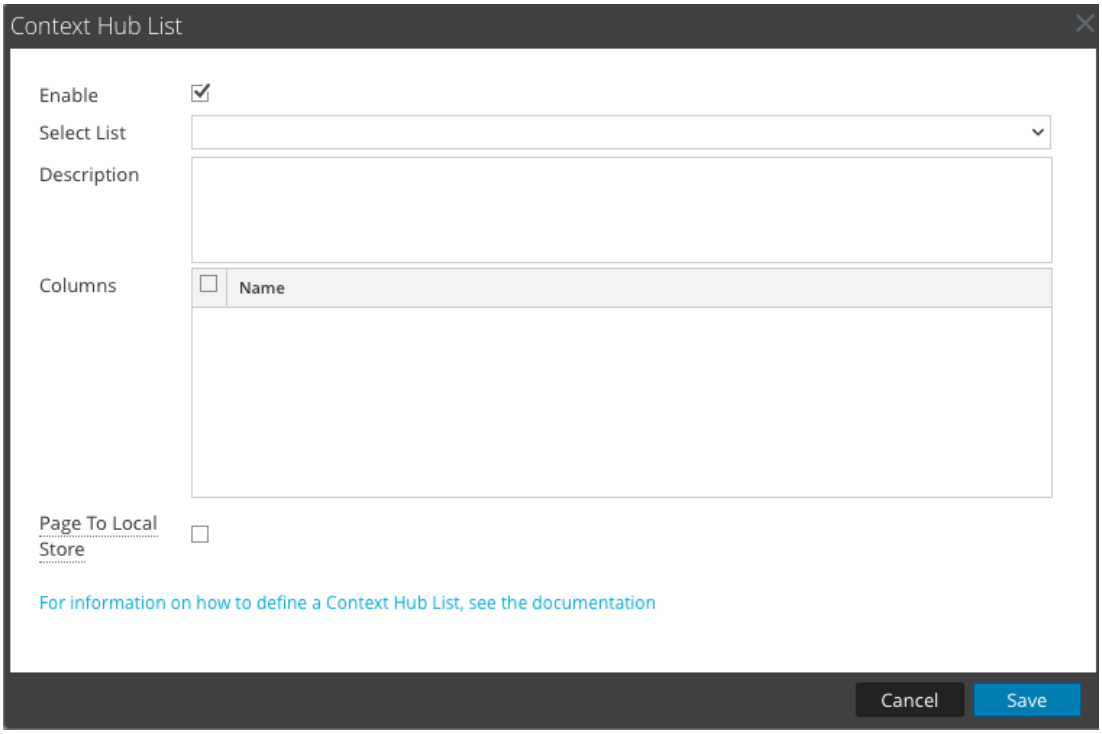
Configure a Context Hub List as an Enrichment Source

1. Go to  **(Configure) > ESA Rules > Settings** tab.
2. In the options panel, select **Enrichment Sources**.

   The Enrichment Sources panel is displayed.



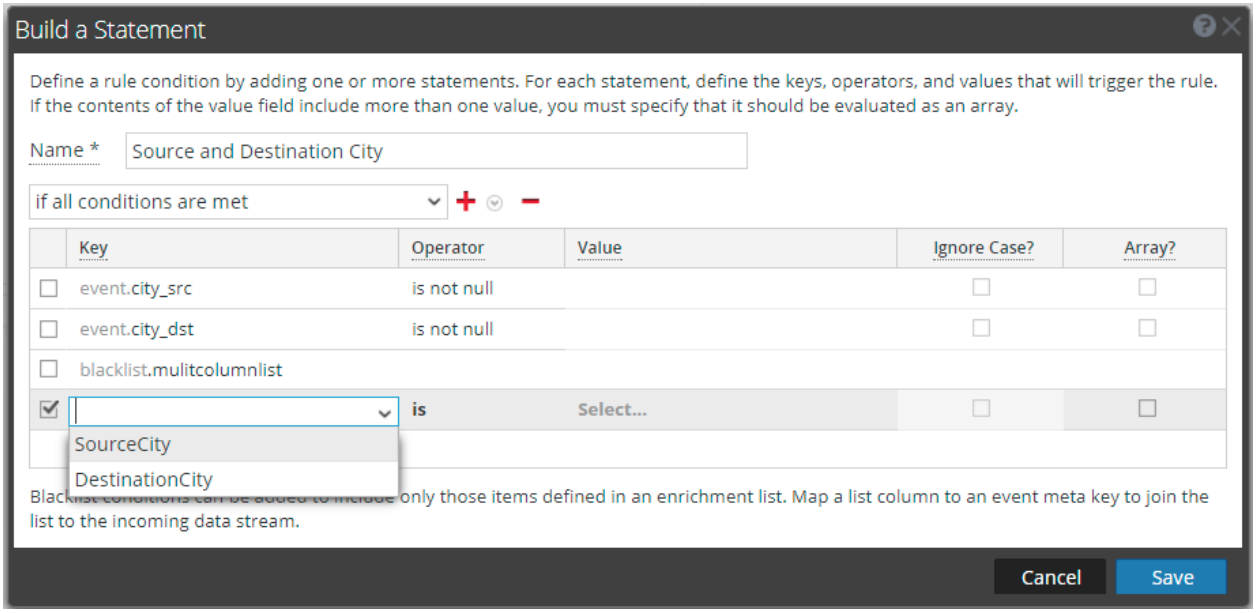
3. From the  drop-down menu, select **Context Hub**.



4. Select **Enable** to enrich alerts with a Context Hub list. This is selected by default. If disabled, the alerts will not be enriched with the configured Context Hub list.
5. Select the desired Context Hub list from the **Select List** drop-down menu of pre-configured lists.
6. (Optional) In the **Description** field, type a brief description about the selected Context Hub list. The text entered here is displayed on the Enrichment Sources panel.

7. In the **Columns** field, all columns included in the selected Context Hub list are listed. Click to enable or disable the columns in the list that you wish to include when using this list as an enrichment source in an alert.
8. (Optional) Click to enable the **Page To Local Store** option. This option is useful if you have a very large list and performance is affected. If this is the case, enabling this option will write a copy of the Context Hub list to the local disk to improve performance.
9. Click **Save**.
   The Context Hub list is configured. You can now add it to an ESA rule as part of a condition statement as either a blacklist or a whitelist condition.

The following figure illustrates adding a Context Hub list as part of a condition statement. In this example, a context Hub list named "multicolumnlist" was added as a blacklist condition. The list contains two columns, SourceCity and DestinationCity. The next step would be to select one of the column names as the subcondition and then specify the operator and enter the meta value for the corresponding value field.



For complete details for adding a whitelist or blacklist to a condition statement, see Step 2. Build a Rule Statement.

To add a Context Hub list as a condition to an existing rule, select to edit the desired rule in the Rule Library, then add a condition in the Conditions section and select to add a whitelist or blacklist condition to the new condition statement.

Configure an In-Memory Table as an Enrichment Source

This topic provides instructions on how to configure an in-memory table. When you configure an in-memory table, you upload a .CSV file as an input to the table. You can associate this table with a rule as an enrichment source. When the associated rule generates an alert, ESA will enrich the alert with relevant information from the in-memory table.

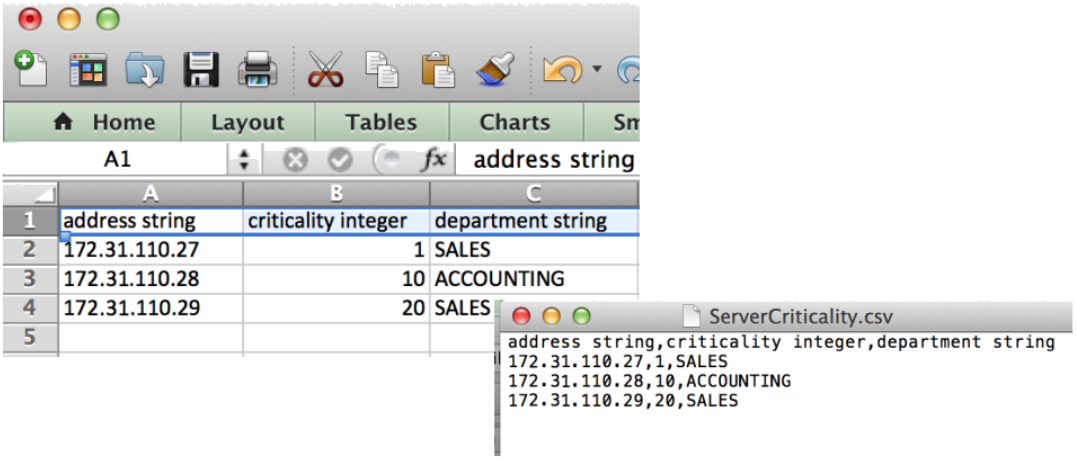
For example, a rule could be configured to detect when a user tries to download freeware and to identify the person by user ID in the alert. The alert could be enriched with additional information from an in-memory table that contains details such as full name, title, office location and employee number.

Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see Configure a Context Hub List as an Enrichment Source.

Prerequisites

- The column name in the .CSV file cannot have whitespace characters.
  For example *Last_Name* is correct, and *Last Name* is incorrect.
- The .CSV file must begin with a header line that defines fields and types.
  For example, *address string* would define the header field as *address*, and the type as *string*.

The following shows a valid .CSV file represented as a .CSV and as a table.

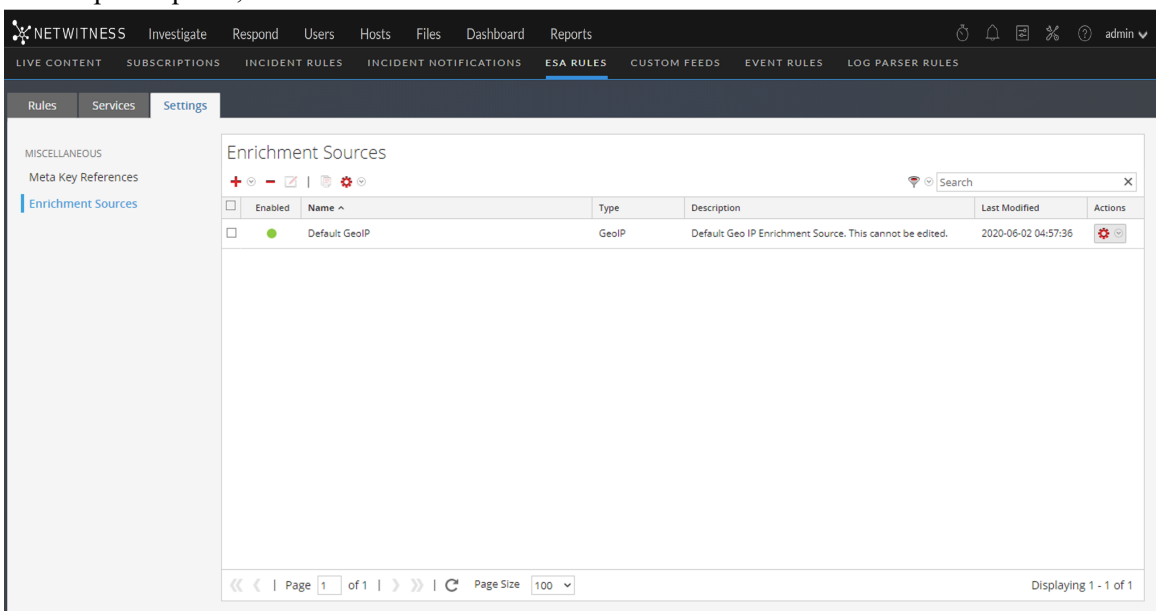


Configure an Ad hoc In-Memory Table

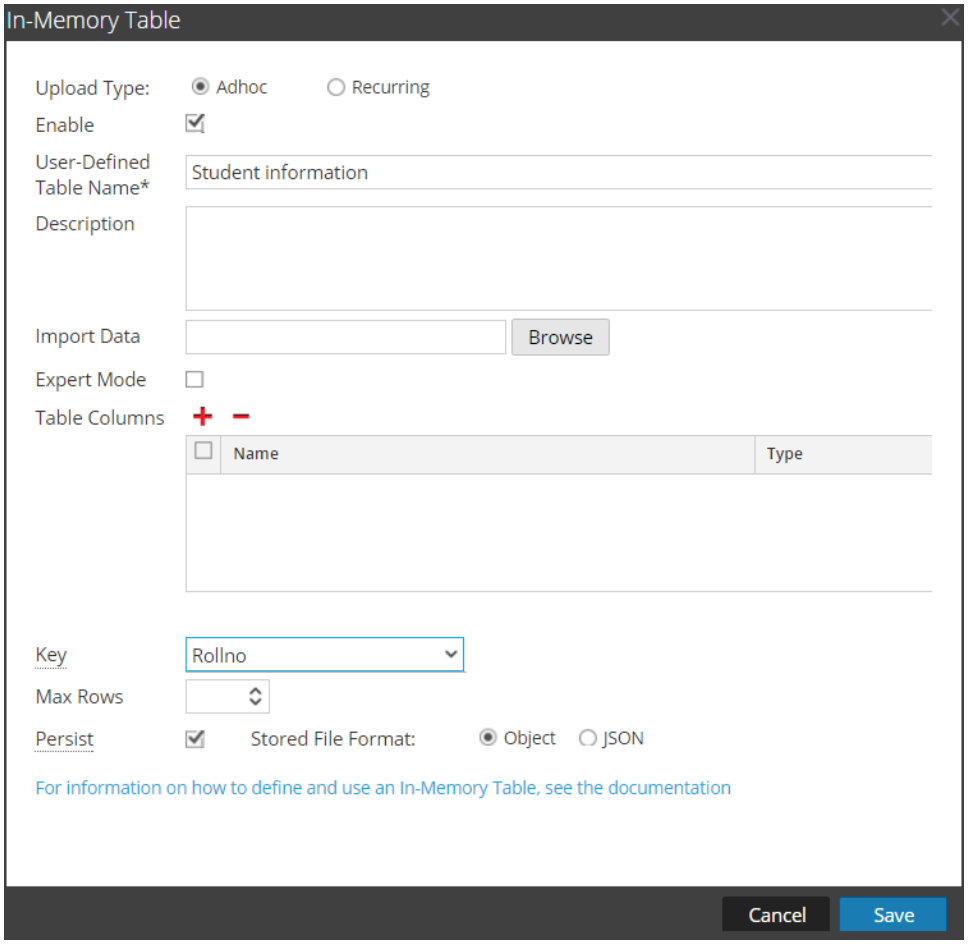
Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see Configure a Context Hub List as an Enrichment Source.

1. Go to **(Configure) > ESA Rules**.
   The Configure view is displayed with the ESA Rules tab open.
2. Click the **Settings** tab.

3. In the options panel, select **Enrichment Sources**.



4. In the **Enrichment Sources** section, click  > **In-Memory Table**.



5. Describe the in-memory table:
   a. Select **Ad hoc**.
   b. By default, **Enable** is selected. When you add the in-memory table to a rule, alerts will be enriched with data from it.
   If you add an in-memory table to a rule but do not want alerts to be enriched, deselect the checkbox.
   c. In the **User-Defined Table Name** field, type a name, such as Student Information, for the in-memory table configuration.

   Note: Do not use any Esper keyword as **User-Defined Table Name** since this causes an error while using this enrichment in the ESA Rule. For Esper keywords, see Reserved keywords.

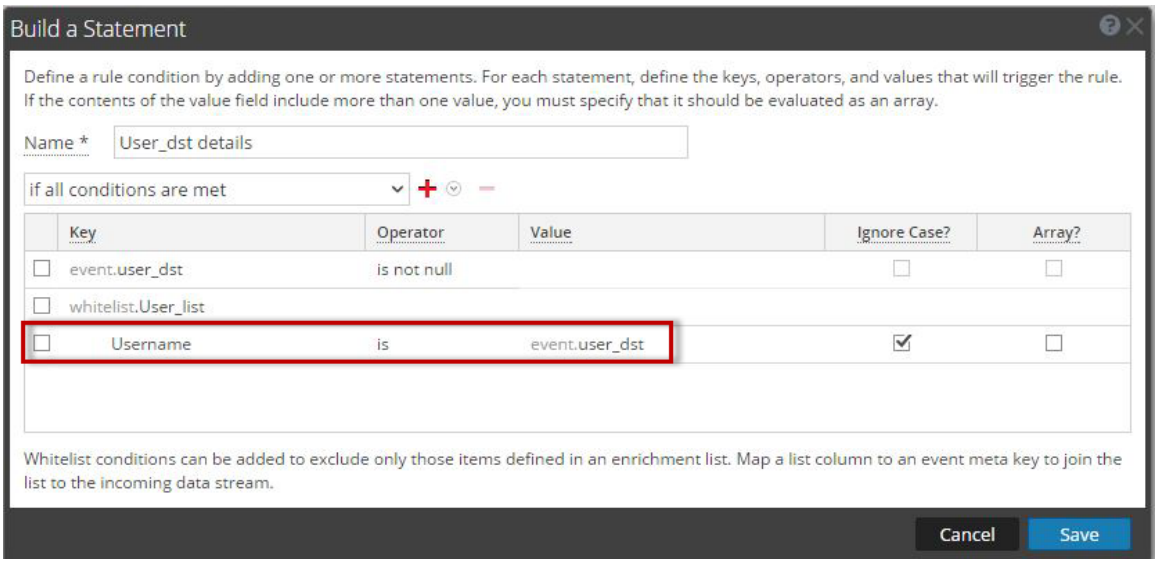
   d. If you want to explain what the enrichment adds to an alert, type a **Description** such as:
   When an alert is grouped by Rollno, this enrichment adds student information, such as name and marks.
6. In the **Import Data** field, select the .CSV file that will feed data to the in-memory table.
7. If you want to write an EPL query to define an advanced in-memory table configuration, select **Expert Mode**.
   The Table Columns are replaced by a **Query** field.
8. In the **Table Columns** section, click + to add columns to the in-memory table.
9. If a valid file is selected in the Import Data field, the columns populate automatically.

Note: If you selected Expert mode, a Query field is displayed instead of Table Columns.

10. In the **Key** drop-down menu, select the field to use as the default key to join incoming events with the in-memory table when using a CSV-based in-memory table as an enrichment. By default, the first column is selected. You can also later modify the key when you open the in-memory table in enrichment sources.
11. In **Max Rows** drop-down menu, select the number of maximum number of rows that can reside in the in-memory table at a particular instance.
12. Select **Persist** to preserve the in-memory table on disk when the ESA service stops and to re-populate the table when the service restarts.
13. In **Stored File Format** field, do one of the following:
    - Select **Object**, if you want to store the file in a binary format.
    - Select **JSON**, if you want to store the file in a text format.
    By default, **Object** is selected.
14. Click **Save**.
    The adhoc in-memory table is configured. You can add it to a rule as an enrichment or part of the rule condition. See Add an Enrichment to a Rule.

When you add an in-memory table, you can add it to a rule as an enrichment or as a part of the rule condition. For example, the following rule uses an in-memory table as a part of the rule condition to create a whitelist, and it also uses an in-memory table of details in the user_dst file to enrich the alert that is displayed.

The rule shows the in-memory table as a whitelist rule condition:



Next, the alert is enriched with the User_list in-memory table:



Therefore, the user_dst in-memory table is used to create a whitelist, and it is also used to enrich the data in the alert if the alert is triggered.

Add a Recurring In-Memory Table

Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see Configure a Context Hub List as an Enrichment Source.

It is preferable to use Context Hub List enrichment sources for ESA rules instead of In-Memory Table enrichment sources. You can share Context Hub List enrichment sources across the NetWitness Platform. You can only use the In-Memory Table with ESA.

Note: Database, Database Connection, Warehouse Analytics, and Recurring In-Memory Tables as enrichment sources are not supported for the ESA Correlation service in NetWitness 11.3 and later.

Add an Enrichment to a Rule

This topic tells how to add a previously configured enrichment source to a rule. When ESA creates an alert, information from the source gets included in it.

Adding an enrichment to a rule allows you to request for look ups into a variety of sources and include the results in the outgoing alerts, giving you a more detailed alert. This procedure requires role permissions for Administrator, DPO, and SOC Manager.

Note: This procedure does not apply to adding a Context Hub list as an enrichment to a condition statement in an existing rule. For information see Configure a Context Hub List as an Enrichment Source.

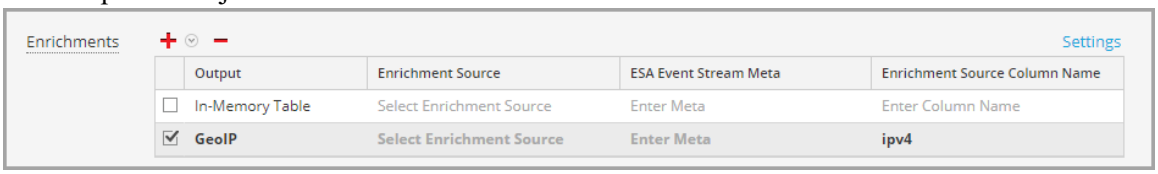
To add an enrichment to a rule:

1. Go to **(Configure) > ESA Rules**.
2. In the **Rule Library** view, do one of the following:
   - Double-click a rule.
   - Select a rule and click in the **Rule Library** toolbar.

   The Rule Builder panel is displayed in a new NetWitness tab.
3. In the **Enrichments** section, click and select any of the following enrichment types:
   - In-Memory Table
   - GeoIP

   Note: If you use a GeoIP source, ipv4 is automatically populated, and is not editable.

   The enrichment types that you have selected are displayed in the table.
4. For the added enrichment type, perform the following:
   - In the **Output** column, select the type that you have configured.
   - In the **Enrichment Source** drop-down list, select the enrichment source defined.
   - In the **ESA Event Stream Meta** field, type the event stream meta key whose value will be used as one operand of join condition.



   - In the **Enrichment Source Column Name** field, type the enrichment source column name whose value will be used as another operand of the join condition.
5. Select **Debug**. This adds an @Audit('stream') annotation to the rule. This is useful when debugging the Esper rules.
6. Click **Show Syntax** to test if the defined ESA rule is valid.
7. Click **Save**.

For details on parameters and their descriptions, see Rule Builder Tab.

Deploy Rules to Run on ESA

This section explains how an ESA Rule Deployment works and how to set up a deployment to run a group of ESA rules. Administrator, SOC Manager, or Data Privacy Officer role permissions are required for all procedures in this section.

To create an ESA rule deployment, you need to perform the steps described in ESA Rule Deployment Steps.

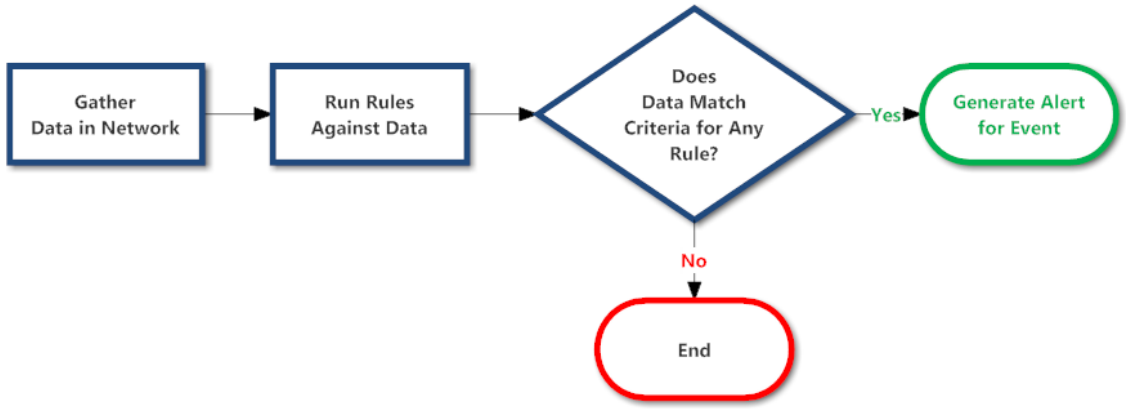
How an ESA Rule Deployment Works

An ESA rule deployment consists of an ESA service, one or more data sources, and a set of ESA rules. When you deploy rules, the ESA service runs them to detect suspicious or undesirable activity in your network. Each ESA rule detects a different event, such as when a user account is created and deleted within one hour.

The ESA service performs the following functions:

1. Gathers **data** in your network
2. Runs ESA **rules** against the data
3. Applies rule **criteria** to data
4. Generates an **alert** for the captured event

The following graphic shows this workflow:



In addition, you may want to perform other steps on your deployment, such as replacing an ESA service, changing a data source, editing or deleting a rule from the deployment, renaming or deleting the deployment, or showing updates to the deployment. For descriptions of these procedures, Additional ESA Rule Deployment Procedures.

ESA Rule Deployment Steps

This topic explains how to add an ESA rule deployment, which includes an ESA service with its associated data sources and a set of ESA rules. You can add an ESA rule deployment to organize and manage ESA services and rules. Think of the deployment as a container for these components:

- An ESA service
- One or more data sources (This is available in version 11.3 and later.)
- A set of ESA rules

For example, if you add a Spam Activity deployment it could include an ESA London service, Concentrators with the appropriate data, and a set of ESA rules to detect suspicious email activity.

Note: An ESA rule deployment can have only one ESA service. You can, however, use the same ESA service in multiple deployments.
In NetWitness version 11.2 and earlier, the ESA service is the Event Stream Analysis service. In version 11.3 and later, it is the ESA Correlation service.

To add an ESA rule deployment, you need to complete the following procedures:

- Step 1. Add an ESA Rule Deployment
- Step 2. Add an ESA Service
- Step 3. Add Data Sources
  - (Optional) Add a Data Source Filter
- Step 4. Add and Deploy Rules

Step 1. Add an ESA Rule Deployment

Prerequisites

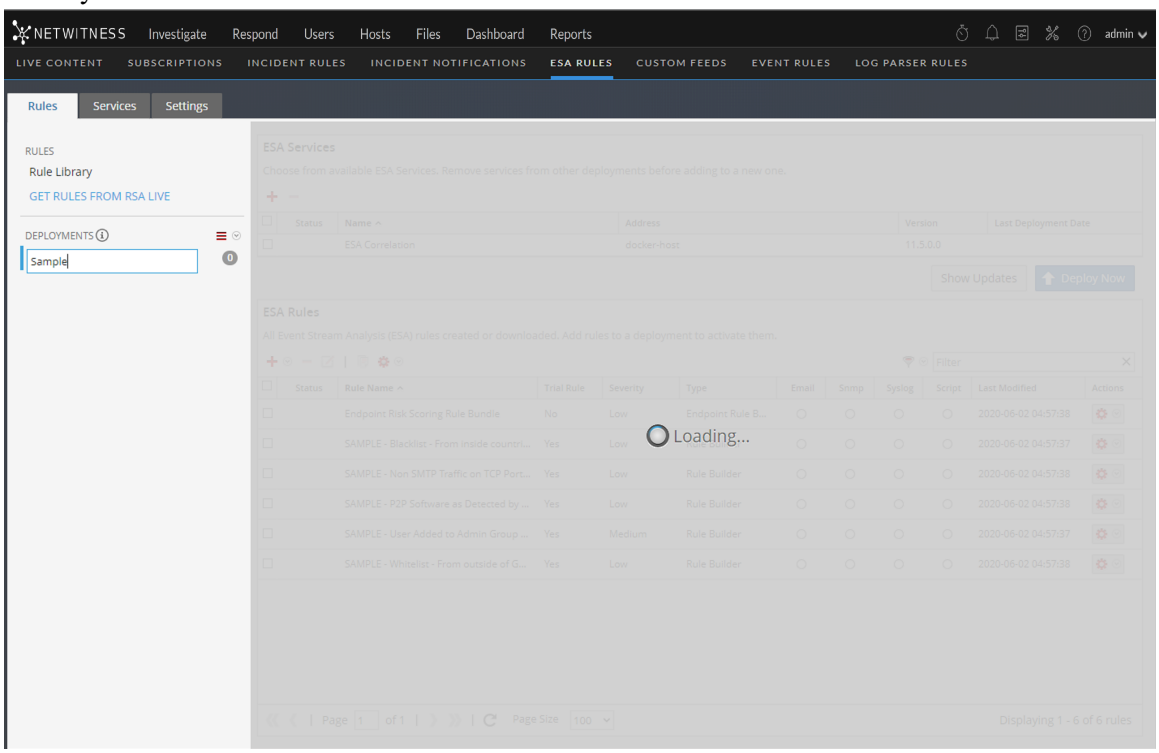
The following are required to add an ESA rule deployment:

- The ESA service must be configured on the host.
- Rules must be in the Rule Library. See Add Rules to the Rule Library.

To add an ESA rule deployment:

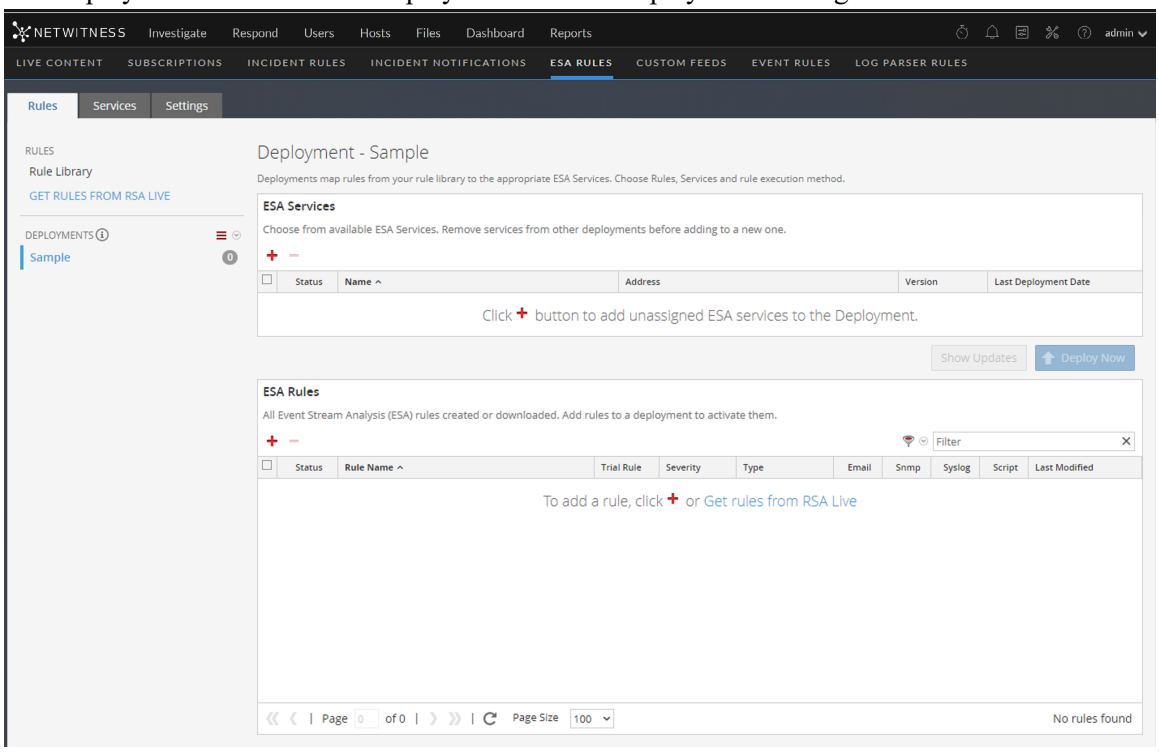
1. Go to  **(Configure) > ESA Rules**.
   The Rules tab is displayed.
2. In the options panel on the left, next to Deployments, select  **> Add** and type a **name** for the deployment. The naming convention is up to you. For example, it could indicate the purpose or

identify an owner.



In NetWitness 11.3 and later, the deployment names that you choose appear on the deployment tabs in the **(Configure) > ESA Rules > Services** tab.

3. Press **Enter**.
   The deployment is added. The Deployment view is displayed on the right.



Step 2. Add an ESA Service

The ESA service in an ESA rule deployment gathers data in your network and runs ESA rules against the data. The goal is to capture events that match rule criteria, then generate an alert for the captured event.

An ESA rule deployment can have only one ESA service. You can, however, use the same ESA service in multiple deployments. For example, ESA London could be in these deployments simultaneously:

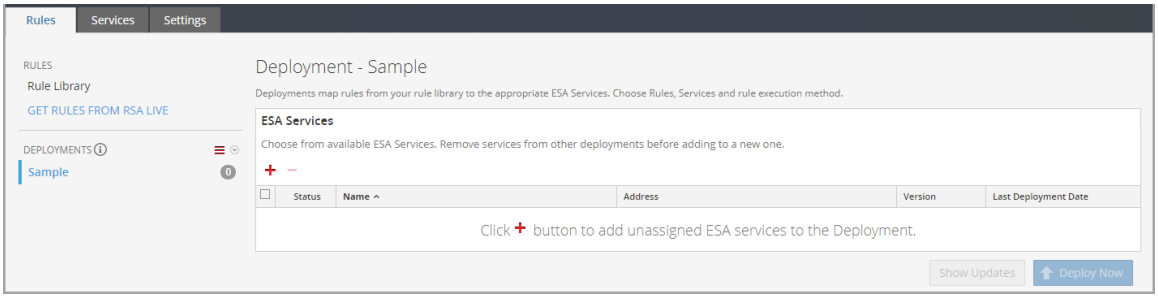
- Deployment EUR, which includes one set of rules
- Deployment CORP, which includes another set of rules.

Changes made to an ESA rule deployment do not take effect until you click Deploy Now. For example, Deployment EUR could include the ESA London service and a set of 25 rules. If you replace the ESA London service with the ESA Paris service, the next time you deploy Deployment EUR, the 25 rules will be removed from ESA London and added to ESA Paris.

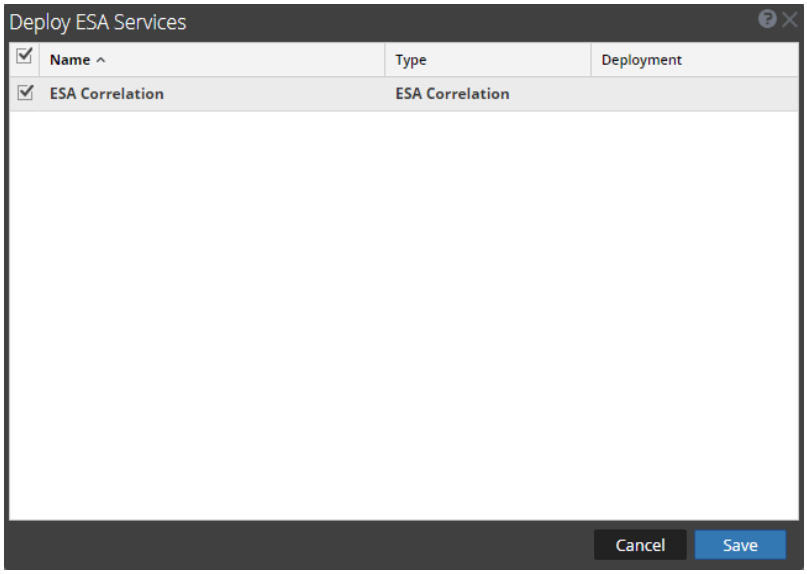
Deleting an ESA rule deployment immediately removes the rules from the ESA service. If an ESA service is not part of any deployment, the ESA service does not have any rules.

To add an ESA service:

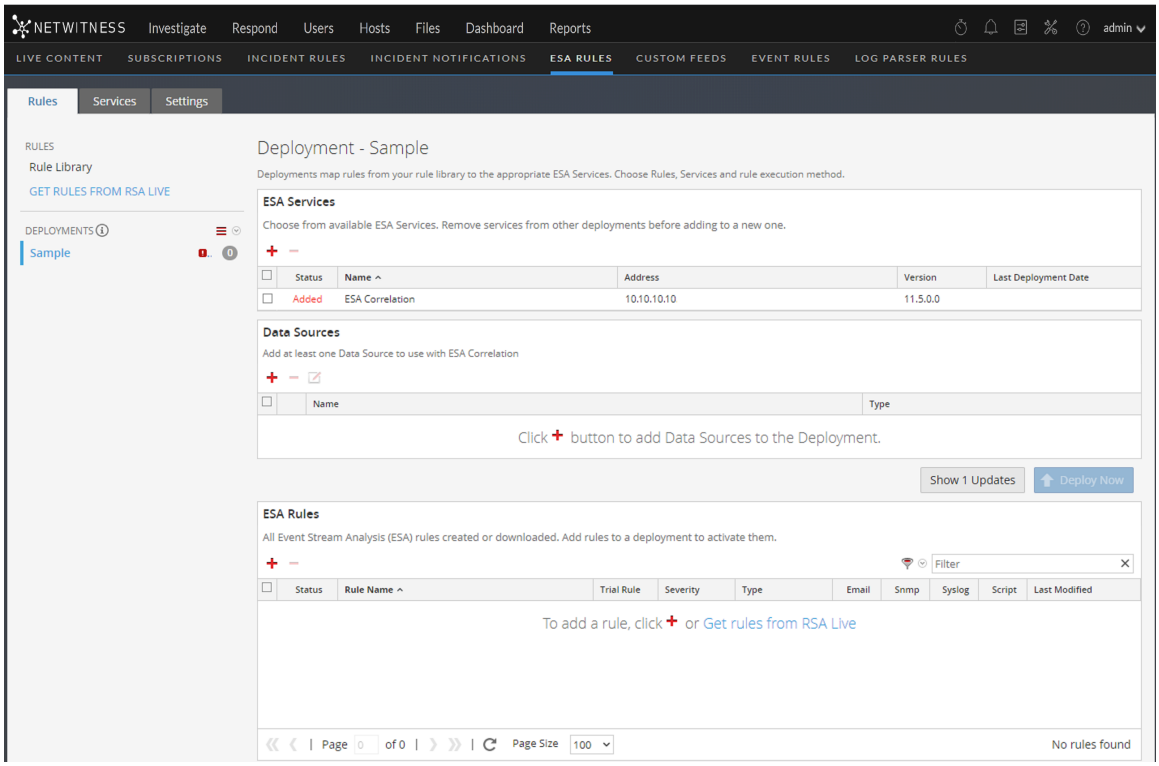
1. Go to **(Configure) > ESA Rules > Rules** tab.
2. In the options panel, select a deployment:



3. In the **Deployment** view, click + in **ESA Services**.
   The Deploy ESA Services dialog lists each configured ESA.



4. Select an ESA service and click **Save**.
   The Deployment view is displayed. The ESA service is listed in the **ESA Services** section, with the status **Added**.

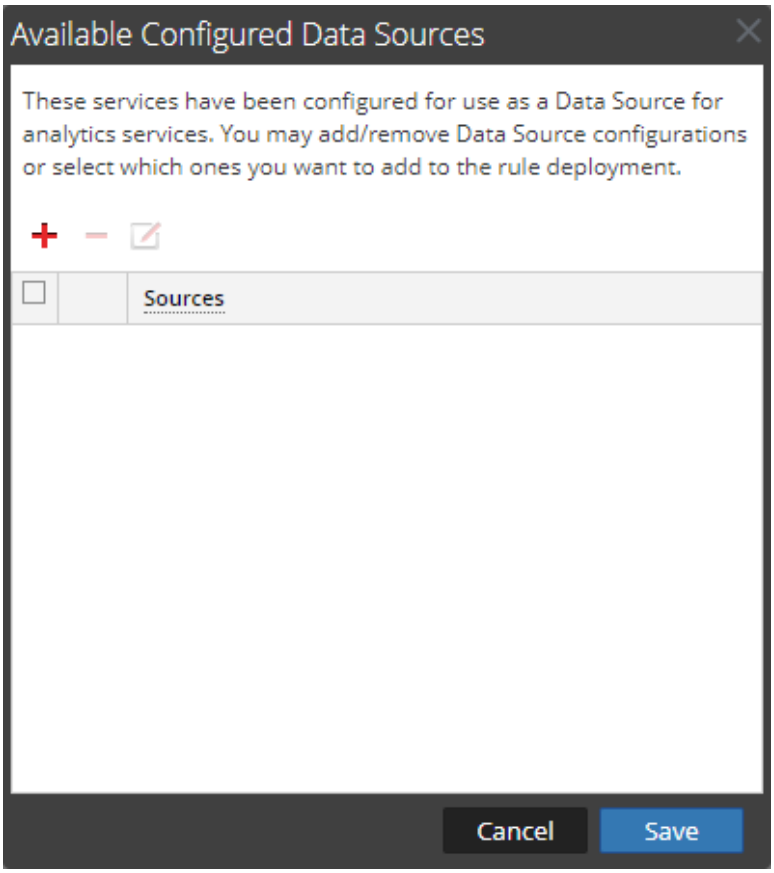


Step 3. Add Data Sources

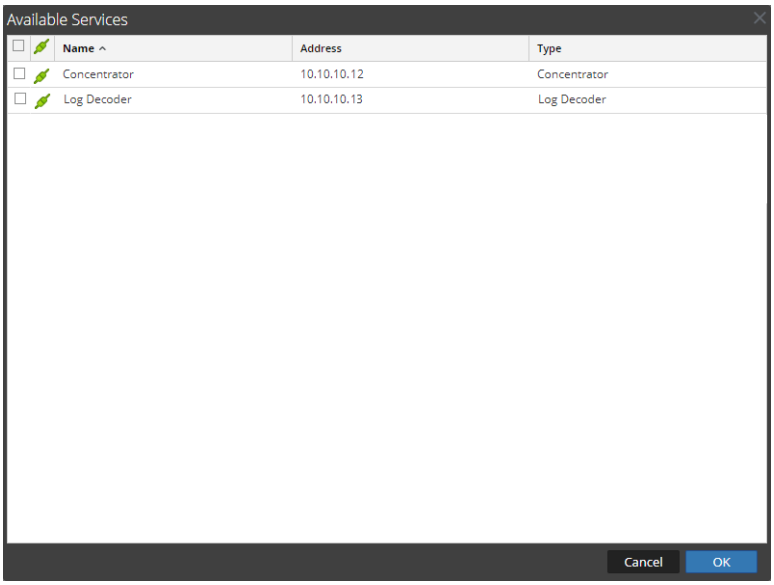
Note: This option is available in version 11.3 and later.

You can select one or more data sources, such as Concentrators, to use for your selected ESA Service. This enables you to specify different data sources for each deployment. For example, you may want to use Concentrators with HTTP packet data in one deployment and Concentrators with HTTP log data in another deployment.

1. Go to **(Configure) > ESA Rules > Rules** tab.
2. In the options panel, select a deployment.
3. Configure one or more data sources for your deployment. Do the following for each data source:
   a. In the **Deployment** view **Data Sources** section, click .
   The **Available Configured Data Sources** dialog lists the services that have been configured for use as a data source.



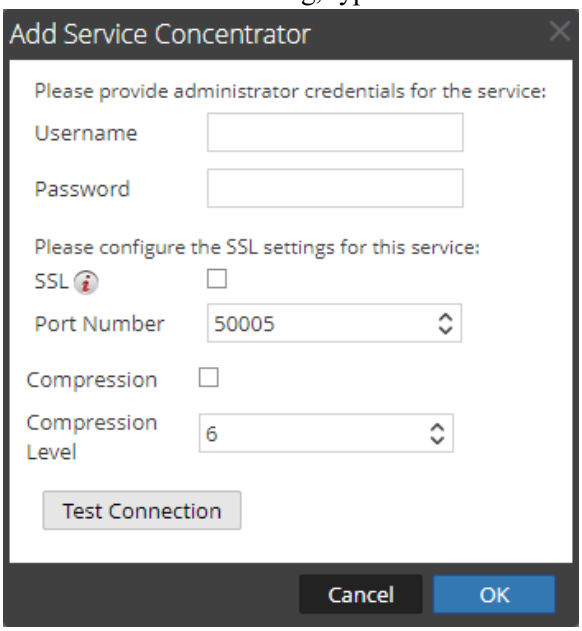
   b. To add a data source configuration, click .
   The **Available Services** dialog lists the available data sources from the (Admin) > Services view, such as Concentrators.



Note: You can add a Log Decoder as a data source for ESA, but it is better to add a Concentrator to take advantage of undivided aggregation as the Decoder may have other processes aggregating from it.

c. In the **Available Services** dialog, select a data source, such as a Concentrator, and click **OK**.

d. In the **Add Service** dialog, type the Administrator username and password for the data source.



e. To enable the SSL or Compression options, select the corresponding checkboxes.

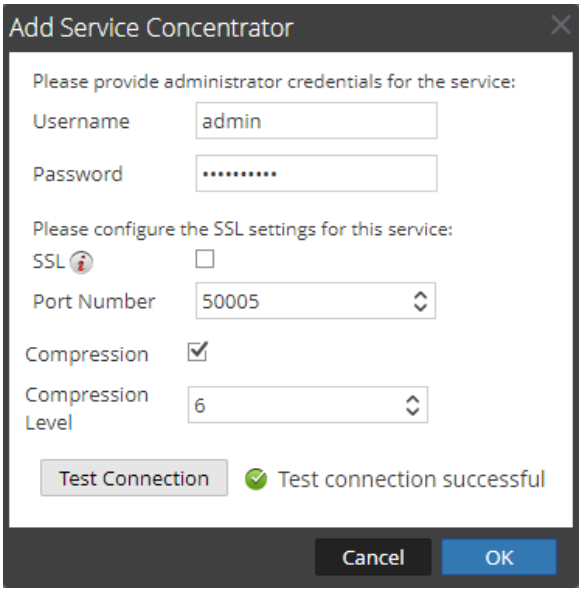
f. (Optional) You have the option to adjust the Compression Level for Concentrators on ESA in NetWitness 11.3 and later. To enable compression, select the **Compression** checkbox. You can set the **Compression Level** for a Concentrator from 0-9:

- Compression Level = **0** (If compression is enabled, it allows Core Services to control the amount of compression.)
- Compression Level = **1** (It uses the lowest amount of compression and has the highest performance.)
- Compression Level = **9** (It uses the highest amount of compression and has the worst performance.)

Somewhere in the middle between 1 and 9 is usually the best setting, which is what you get when you select a compression level of 0. For more detailed information, see the *Core Database Tuning Guide*.

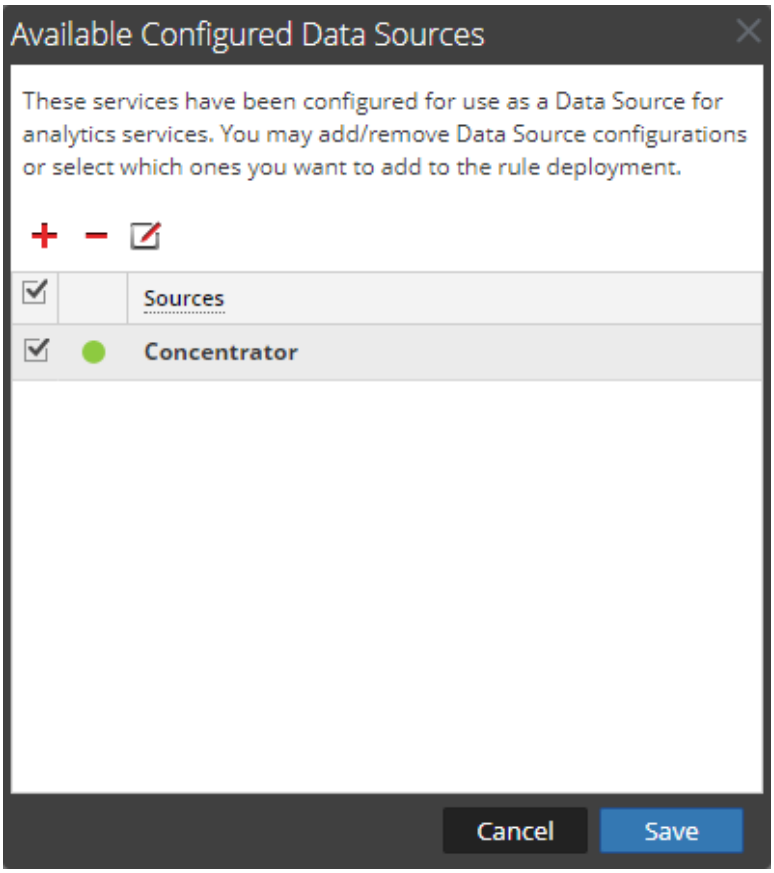
Note: If you make any ESA service, data source, or ESA rule changes to an ESA rule deployment, you need to redeploy the deployment. For example, if you change the configuration of a data source in an ESA rule deployment, you must redeploy all the ESA rule deployments that contain that data source.
When you set the compression level for a Concentrator on ESA, it sets the same compression level for that Concentrator for ESA Correlation Rules.

g. Click **Test Connection** to make sure that it can communicate with the ESA service.



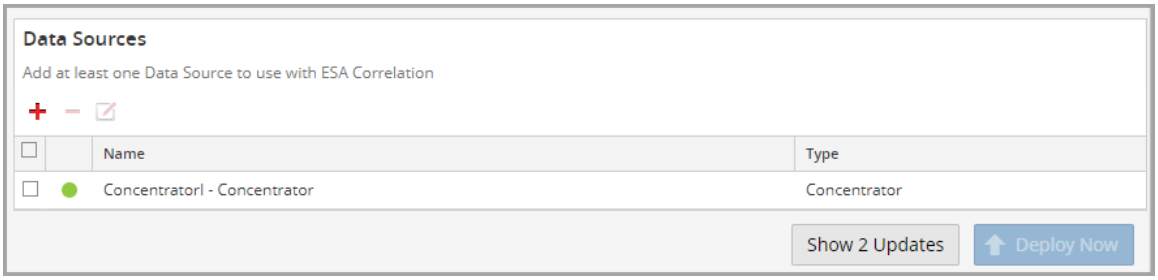
h. Click **OK**.
After you configure your data sources and they appear in the **Available Configured Data Sources** dialog, you can use them for your deployment.

4. In the **Available Configured Data Sources** dialog, select at least one data source to use for the deployment.



A solid colored green circle indicates a running service and a white circle indicates a stopped service.

5. Click **Save**.
   In the Deployment view **Data Sources** section, the selected data sources are added to the deployment. The **Deploy Now** button activates after an ESA service, a data source, and rules are added to an ESA rule deployment.



6. (Optional) If you have a medium to large NetWitness Platform deployment and you have high throughputs, you can create a filter to forward only the data relevant to this deployment to ESA. Before deploying the ESA rule deployment, add the data source filter. See (Optional) Add a Data Source Filter.

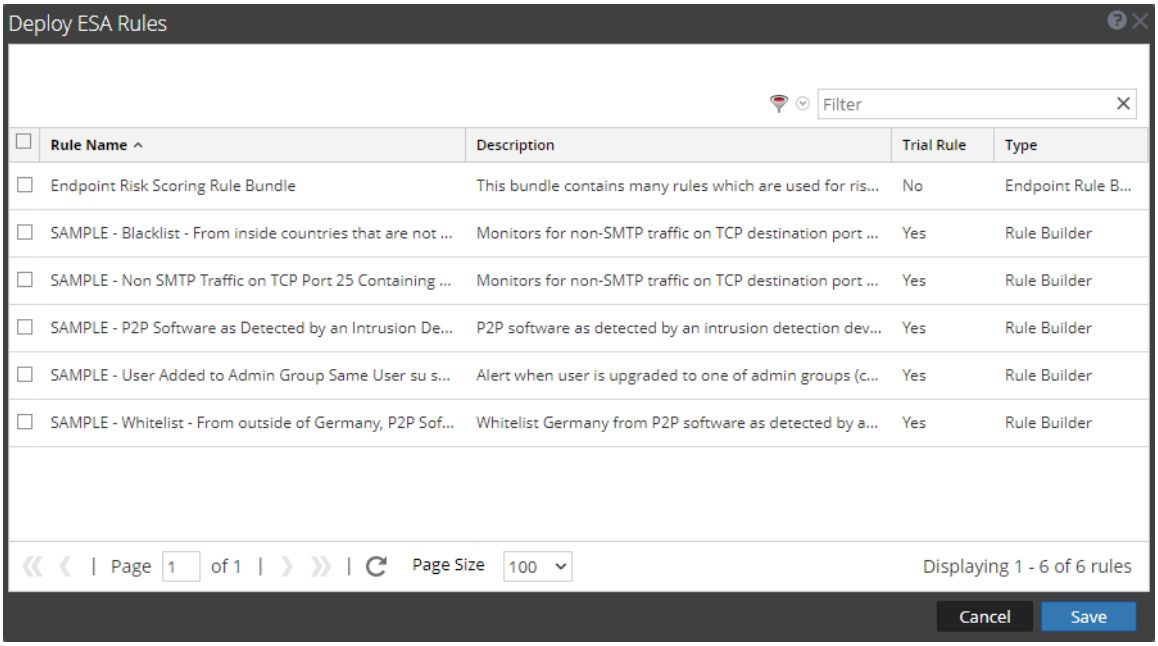
Step 4. Add and Deploy Rules

This topic explains how to add ESA rules to an ESA rule deployment and then deploy the rules on ESA. Each ESA rule has unique criteria. The ESA rules in an ESA rule deployment determine which events ESA captures, which in turn determine the alerts you receive.

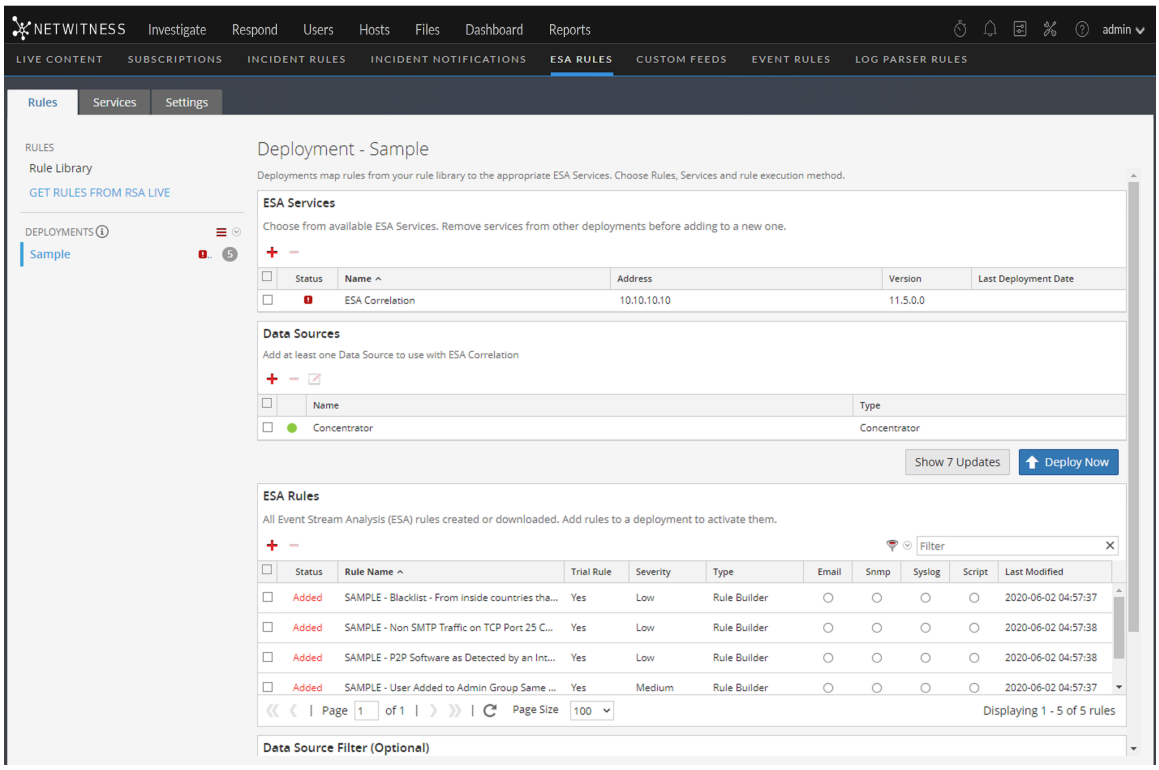
For example, Deployment A includes ESA Paris and, among others, a rule to detect file transfer using a non-standard port. When ESA Paris detects a file transfer that matches the rule criteria, it captures the event and generates an alert for it. If you remove this rule from Deployment A, ESA will no longer generate an alert for such an occurrence.

To add and deploy rules:

1. Go to **(Configure) > ESA Rules > Rules** tab.
2. In the options panel, select a deployment.
3. In the **Deployment** view, click in **ESA Rules**.
   The Deploy ESA Rules dialog is displayed and shows each rule in your Rule Library:



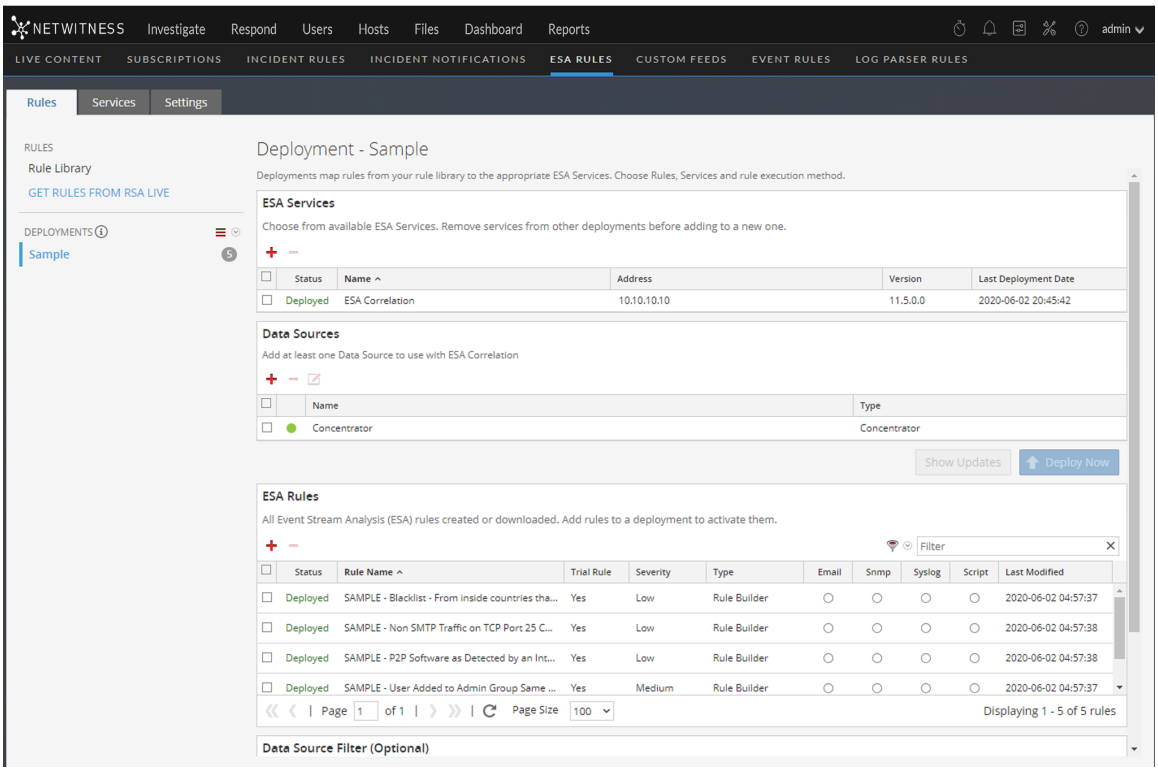
4. Select rules and click **Save**.
   The Deployment view is displayed and the **Deploy Now** button is enabled.



5. The rules are listed in the ESA Rules section.
   - In the Status column, **Added** is next to each new rule.
   - In the Deployments section,  indicates there are updates to the deployment.
   - The total number of rules in the deployment is on the right.



6. Click **Deploy Now**.
   The ESA Correlation service runs the rule set. After ESA correlation completes the processing of each rule in the deployment, the status changes to **Deployed**.



(Optional) Add a Data Source Filter

Note: This option is available in NetWitness Platform version 11.5 and later.

To improve performance, you can add an optional data source filter to your ESA rule deployment so that only the data relevant to the deployment is forwarded to ESA. The filter is comprised of application rules, which are applied to the Decoders mapped to your selected data sources. There are two options that you can use to create the filter:

- The **Simple** data source filter option enables you to select the application rules to be included in the filter query. The application rules that you select must be enabled on the Decoders that feed the data sources in your ESA rule deployment. ESA Correlation uses the filtered event data to process the ESA rules. This is the procedure described here.
- The **Advanced** data source filter option enables you to add a data source query directly. The individual application rule queries must be separated by an “or” condition. For more information on creating and writing Decoder rules, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.

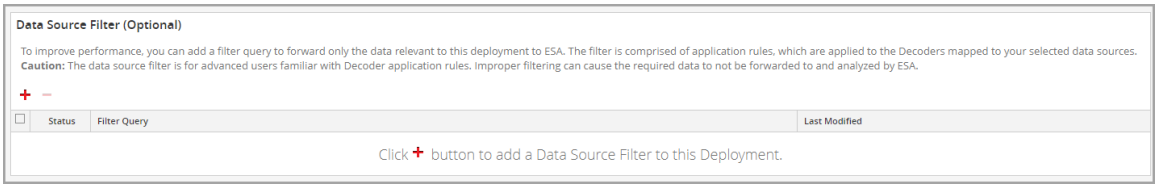
Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.

IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.

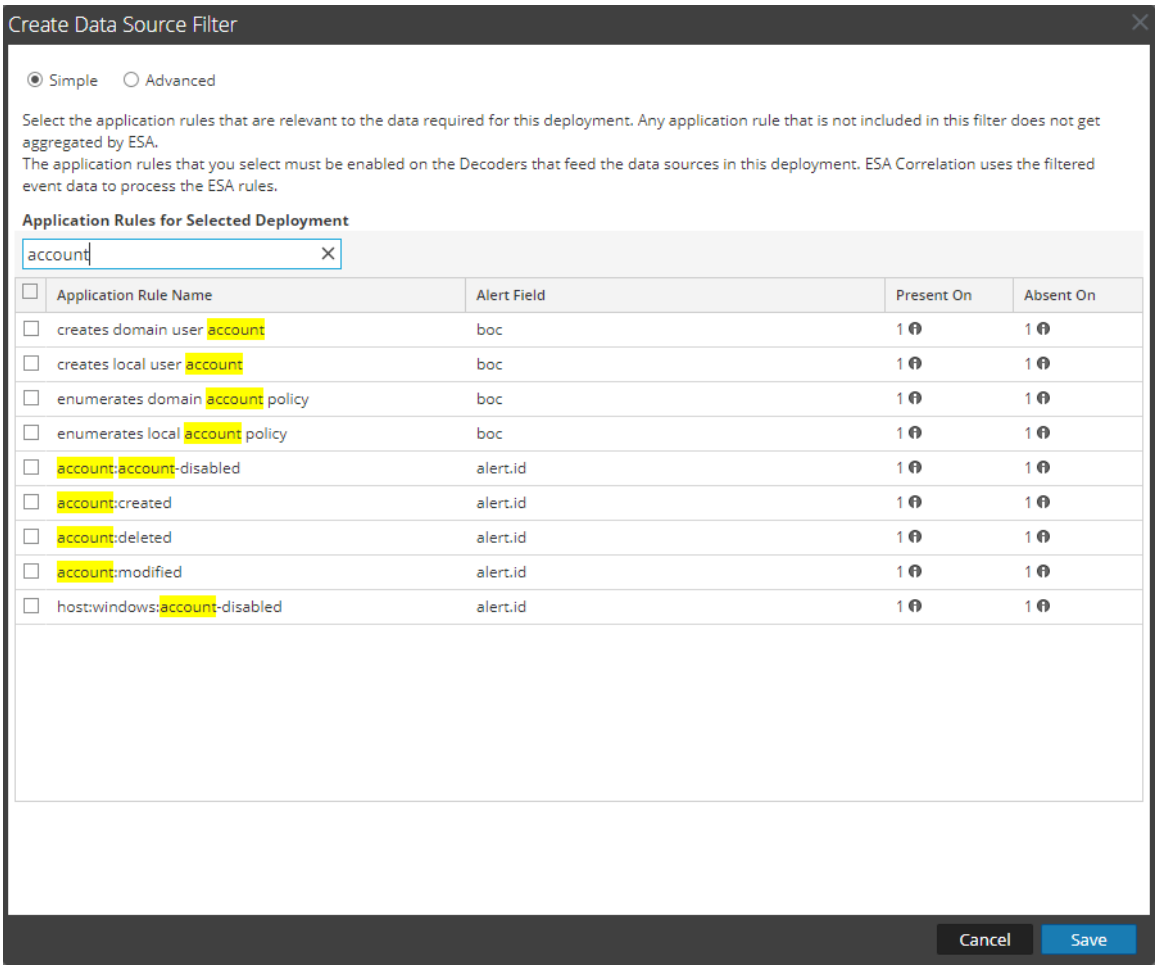
To add a Simple data source filter:

1. Go to **(Configure) > ESA Rules > Rules** tab.
2. In the options panel, select a deployment.
3. In the **Deployment** view **Data Source Filter (Optional)** section, click .

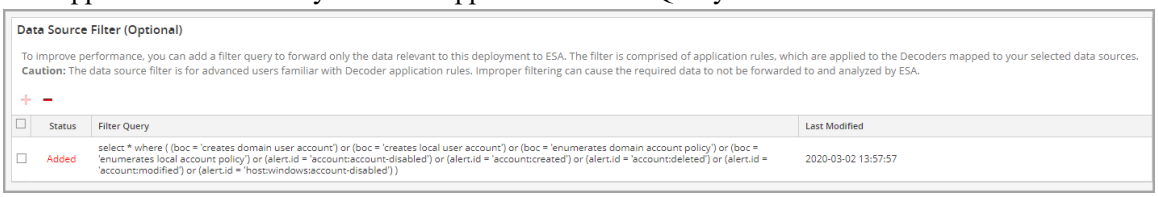


The **Create Data Source Filter** dialog lists the application rules that are available to filter the events in the data sources in your ESA rule deployment.

4. Select the application rules that you want to use to filter the data sources in your deployment and click **Save**. The **Alert Field** shows the meta key used in the alert. **Present On** shows the number of Decoders mapped to the data sources that have the rule. **Absent On** shows the number of mapped Decoders that do not have the rule. If present, hover over the icon to view the names of the Decoders. You can use the filter to help locate the rules. For example, you can type "account" to search for application rules that contain that word.



The application rules that you select appear in a Filter Query with the status of Added.



5. When you are ready to deploy your ESA rule deployment, click **Deploy Now**.
   The ESA Correlation service runs the rule set. After ESA Correlation completes the processing of each rule in the deployment, the status changes to **Deployed**. The data source filter status changes to **Deployed** when the filter is actively streaming only the relevant data as defined in the filter query.

   Troubleshooting Information: When filtering out a large portion of the traffic, you may see an "Invalid header size" error while communicating with Core services in the ESA Correlation log file. Decrease the `max-sessions` parameter from 10,000 to a lower session count. See Adjust Maximum Sessions for the Data Source Filter.

Deploy the Endpoint Risk Scoring Rules Bundle

An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. Endpoint risk scoring rules only apply to NetWitness Endpoint. You can add the Endpoint Risk Scoring Rules Bundle to an ESA rule deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) in the ESA Rule Deployment.

The ESA Correlation service can process endpoint risk scoring rules, which generate alerts that are used in risk scoring calculations to identify suspicious files and hosts. To turn on risk scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. For instructions, see "Deploy Endpoint Risk Scoring Rules on ESA" in the *ESA Configuration Guide*. For complete information on configuring NetWitness Endpoint, see the *NetWitness Endpoint Configuration Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Additional ESA Rule Deployment Procedures

In addition to deploying an ESA service and rules, you may want to perform other steps on your ESA rule deployment, such as replacing an ESA service, changing a data source, editing or deleting a rule from the deployment, renaming or deleting the deployment, or showing updates to an ESA rule deployment.

Note: You cannot edit or duplicate an Endpoint Risk Scoring Rules Bundle.

In NetWitness version 11.3 and later, you can add or remove a data source from a deployment. In NetWitness 11.3.0.2 and later, you can edit a data source in an ESA rule deployment. This enables you to change the data source password, SSL, port, and compression settings.

- Replace an ESA Service in an ESA Rule Deployment
- Edit a Data Source in an ESA Rule Deployment (This option is available in NetWitness version 11.3.0.2 and later.)
- Add or Remove a Data Source (This option is available in version 11.3 and later.)
- (Optional) Add or Remove a Data Source Filter (This option is available in version 11.5 and later)
- Edit or Delete a Rule in a Deployment
- Stop Deployment
- Edit the ESA Rule Deployment Name or Delete a Deployment
- Show Updates to an ESA Rule Deployment
- Fast Deployment and Full Deployment of ESA Rules

Each of the following procedures starts in the Rules tab [**(Configure) > ESA Rules > Rules** tab].

Anytime you make changes to an ESA rule deployment, you must redeploy it for the changes to take effect. To redeploy the deployment, click the **Deploy Now** button for that deployment.

Replace an ESA Service in an ESA Rule Deployment

An ESA rule deployment can have only one ESA service, but you can replace it at any time with another ESA service. You can use the same ESA service in multiple deployments.

Remove an ESA Service from an ESA Rule Deployment

1. Go to **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the options panel, under **Deployments**, select a deployment.
3. In the **ESA Services** section, select a service and click in the toolbar.
   A confirmation dialog is displayed.

4. Click **Yes**.
   The service is removed from the deployment.
5. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

Add an ESA Service to an ESA Rule Deployment

1. To add an ESA Service to an ESA rule deployment, see Step 2. Add an ESA Service. For the ESA Correlation service in NetWitness Respond 11.3 and later, you must add at least one data source to the service. See Step 3. Add Data Sources.
2. After you finish making changes to the ESA rule deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the ESA rule deployment is redeployed.

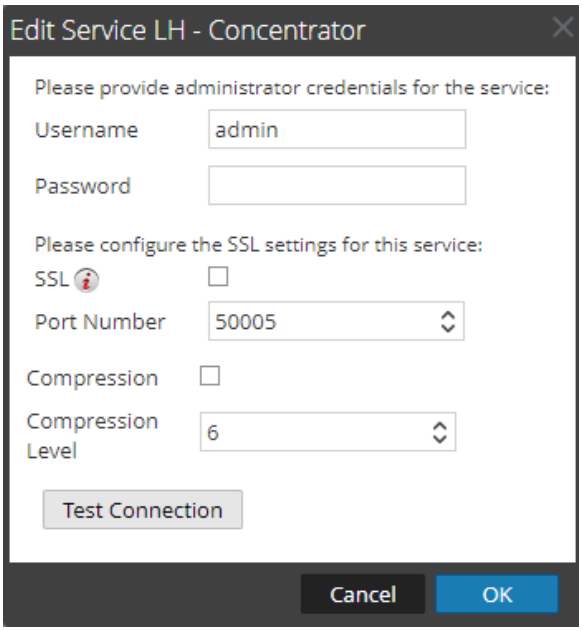
Edit a Data Source in an ESA Rule Deployment

Note: This procedure applies to NetWitness Platform 11.3.0.2 and later versions.

You can change the configuration of a data source in an ESA rule deployment. You can change the data source password, SSL, port, and compression settings. When a data source password changes, it is important to change the password on the data source so that ESA can continue to communicate with the data source.

Note: If you make any ESA service, data source, or ESA rule changes to an ESA rule deployment, you need to redeploy the deployment. For example, if you change the configuration of a data source in an ESA rule deployment, you must redeploy all the ESA rule deployments that contain that data source.

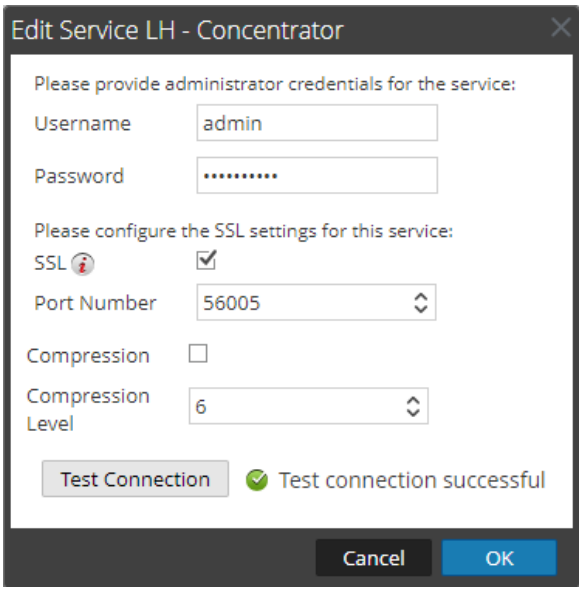
1. Go to  **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the Rules tab options panel, under **Deployments**, select a deployment.
3. In the **Data Sources** section, select a data source and click in the toolbar.
4. In the **Edit Service** dialog, type the Administrator username and password for the data source. If the password changed on the data source, enter the new password here.



5. To enable the SSL or Compression options, select the corresponding checkboxes.
6. (Optional) You have the option to adjust the Compression Level for Concentrators on ESA in NetWitness 11.3 and later. To enable compression, select the **Compression** checkbox. You can set the **Compression Level** for a Concentrator from 0-9:
   - Compression Level = **0** (If compression is enabled, it allows Core Services to control the amount of compression.)
   - Compression Level = **1** (It uses the lowest amount of compression and has the highest performance.)
   - Compression Level = **9** (It uses the highest amount of compression and has the worst performance.)

   Somewhere in the middle between 1 and 9 is usually the best setting, which is what you get when you select a compression level of 0. For more detailed information, see the *Core Database Tuning Guide*.

   Note: When you set the compression level for a Concentrator on ESA, it sets the same compression level for that Concentrator for ESA Correlation Rules.
7. Click **Test Connection** to make sure that it can communicate with the ESA service.



8. Click **OK**.
9. After you finish making changes to the deployment, click **Deploy Now** to redeploy the ESA rule deployment. The changes take effect on ESA after the deployment is redeployed. You can view the update information in the Updates to the Deployments dialog. See Show Updates to an ESA Rule Deployment.

Add or Remove a Data Source

Note: This option is available in NetWitness version 11.3 and later.

Remove a Data Source from an ESA Rule Deployment

1. Go to **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the Rules tab options panel, under **Deployments**, select a deployment.
3. In the **Data Sources** section, select the data source and click in the toolbar.
   The data source is removed from the deployment.
4. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

Add a Data Source to an ESA Rule Deployment

1. To add a data source, see Step 3. Add Data Sources.
2. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

(Optional) Add or Remove a Data Source Filter

To improve performance, you can add an optional data source filter to your ESA rule deployment so that only the data relevant to the deployment is forwarded to ESA. The filter is comprised of application rules, which are applied to the Decoders mapped to your selected data sources.

You cannot edit a data source filter. To modify a data source filter, you must remove the filter, add a new filter, and then redeploy the ESA rule deployment.

Note: This option is available in NetWitness version 11.5 and later.

Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.

IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.

Remove a Data Source Filter from an ESA Rule Deployment

If you plan to replace your data source filter with an adjusted filter, you may want to copy the filter query in the data source filter before you remove it and compare it with the new query.

1. Go to  **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the Rules tab options panel, under **Deployments**, select a deployment.
3. In the **Data Source Filter (Optional)** section, select the filter and click in the toolbar.
   The filter is removed from the deployment.
4. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

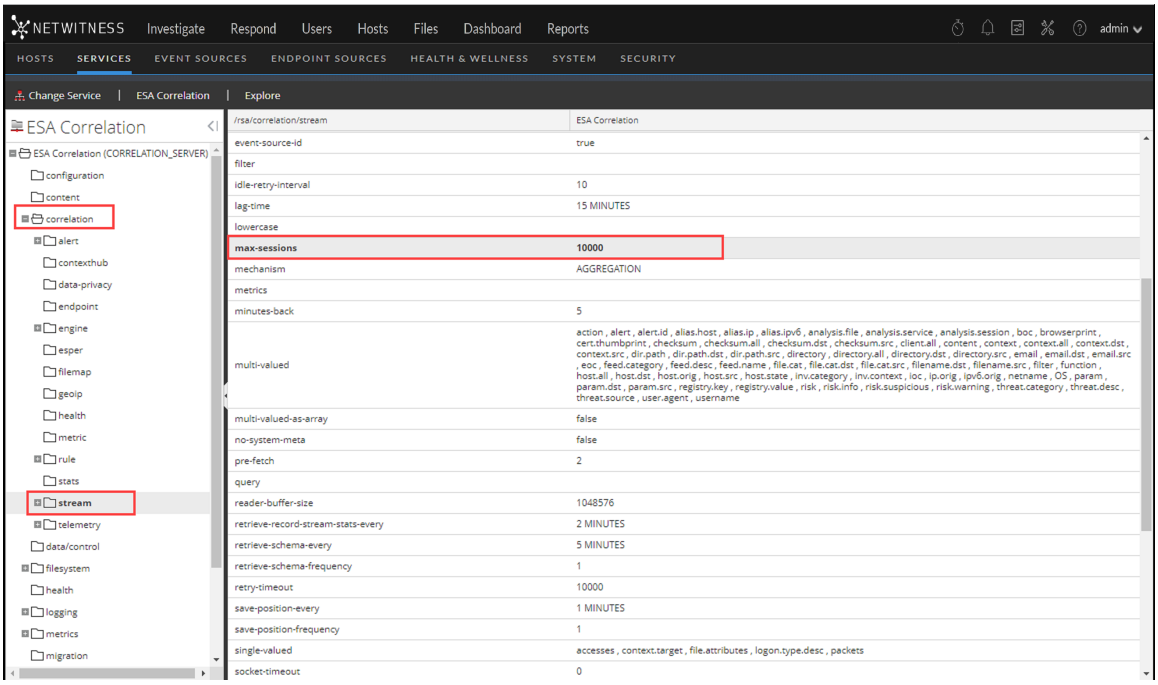
Add a Data Source Filter to an ESA Rule Deployment

1. To add a simple data source filter, see (Optional) Add a Data Source Filter. If necessary, you can use the advanced filter instead of the simple filter to add a data source query directly. The individual application rule queries must be separated by an "or" condition. For more information on creating and writing Decoder rules, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.
2. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

Adjust Maximum Sessions for the Data Source Filter

When filtering out a large portion of the traffic, you may see an "Invalid header size" error while communicating with Core services in the ESA Correlation log file. (You can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`). Lower the `max-sessions` parameter until you no longer see the error in the log. The more you filter out the traffic, the lower you should set the `max-sessions` parameter.

1. In the Explore view node list for an ESA Correlation service, select **correlation > stream**.



2. In `max-sessions`, lower the value until you no longer see the error in the ESA Correlation log file. The default value is **10000**.
3. Restart the ESA Correlation service. Go to **(Admin) > Services**, select the ESA Correlation service, and then select **> Restart**.

Edit or Delete a Rule in a Deployment

In an ESA rule deployment, you can edit and delete rules to customize the deployment.

Edit a Rule

1. Go to **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the Rules tab options panel, under **Deployments**, select a deployment.
3. In the **ESA Rules** panel, double-click a rule to open it in a new tab.

4. Modify the rule, then click **Save**.
   The rule is saved.
5. Click **Deploy Now** to redeploy the deployment.
   The changes take effect on ESA after the deployment is redeployed.

Delete a Rule

1. Go to  **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the options panel, under **Deployments**, select a deployment.
3. In the **ESA Rules** panel, select a rule and click in the toolbar.
   A confirmation dialog is displayed.
4. Click **Yes**.
   The rule is deleted.
5. Click **Deploy Now** to redeploy the deployment.
   The changes take effect on ESA after the deployment is redeployed.

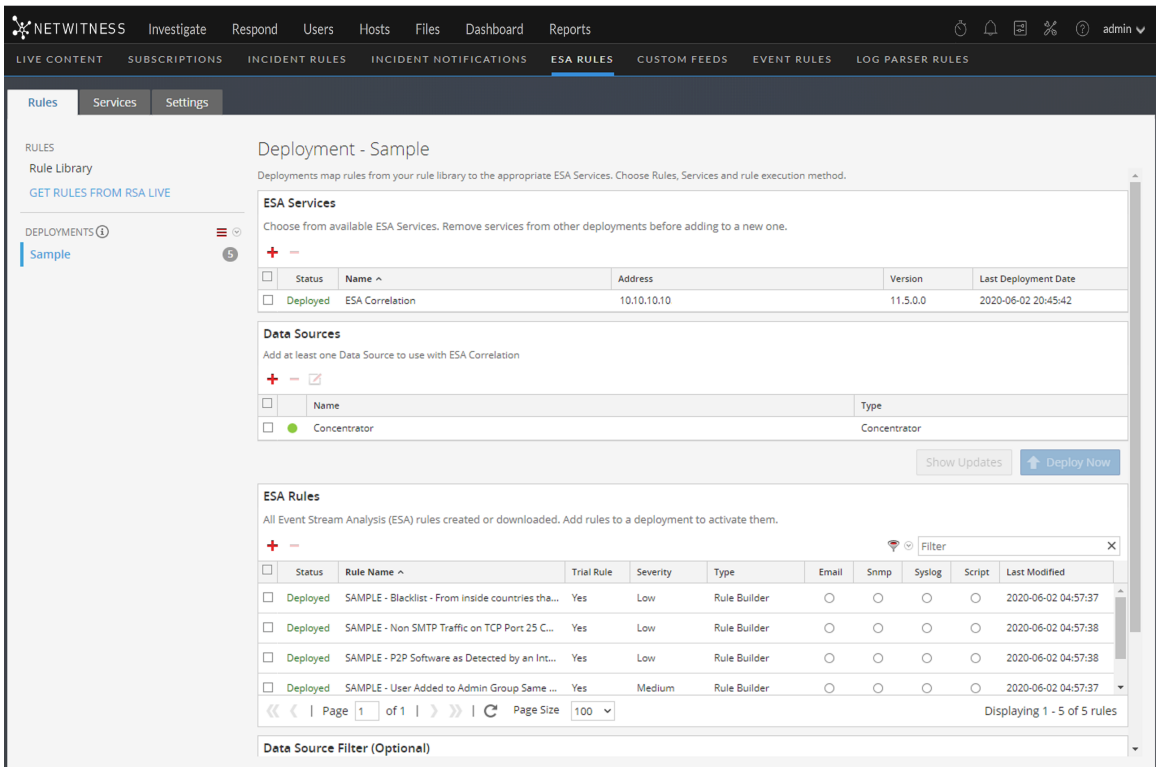
Stop Deployment

1. Go to  **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.
2. In the options panel, under **Deployments**, select a deployment.
3. In the **ESA Rules** panel, select a rule and click in the toolbar.
   A confirmation dialog is displayed.
4. Click **Yes**.
   The rule is to stopped.
5. Click **Stop Deployment** to stop the deployment.

Edit the ESA Rule Deployment Name or Delete a Deployment

To access the deployments:

1. Go to  **(Configure) > ESA Rules**.
   The Configure view is displayed with the Rules tab open.
2. In the options panel, under **Deployments**, select a deployment.
   The Deployment view is displayed.



Edit the ESA Rule Deployment Name

1. In the options panel, under **Deployments**, select a deployment.

   The Deployment view is displayed.

2. Select  > **Edit**.

   The deployment name is made available for editing.

3. Enter the new deployment name.

4. Click **Deploy Now** to redeploy the deployment.
   The changes take effect on ESA after the ESA rule deployment is redeployed. In NetWitness 11.3 and later, the deployment names that you choose appear on the deployment tabs in the  **(Configure) > ESA Rules > Services** tab.

Delete an ESA Rule Deployment

1. In the options panel, under **Deployments**, select a deployment.

   The Deployment view is displayed.

2. Select  > **Delete**.

   A confirmation dialog is displayed.

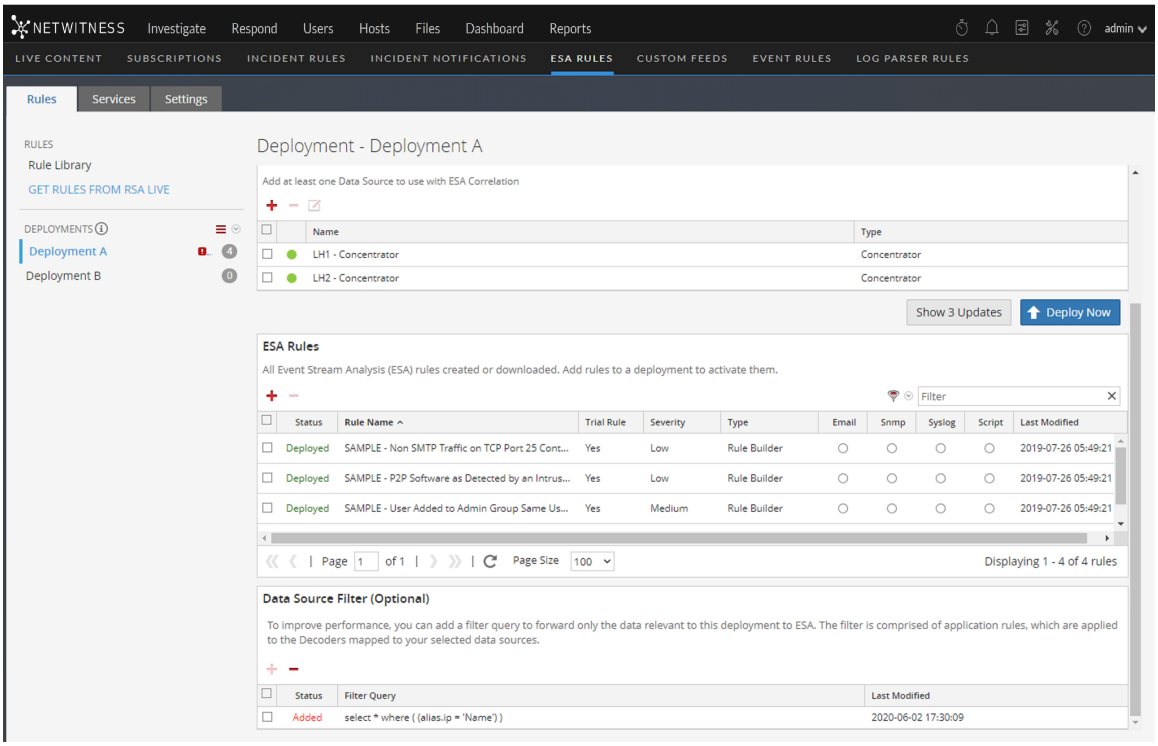
3. Click **Yes**.

   The deployment is deleted.

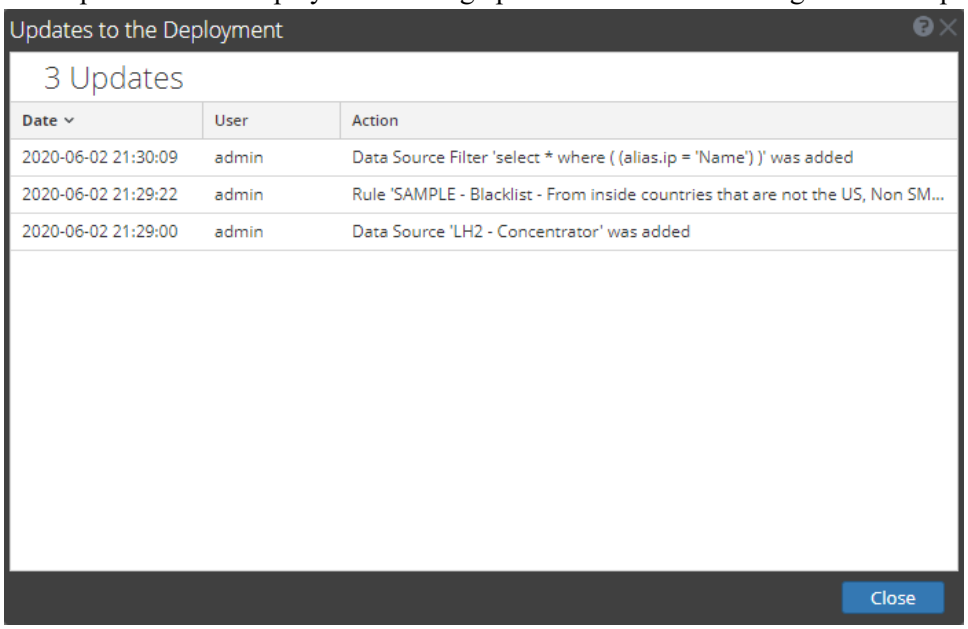
Show Updates to an ESA Rule Deployment

You can view changes to an ESA rule deployment, such as adding or removing rules. When there is a change to a deployment, the update icon ( ) appears next to the name of the deployment in the Rules tab options panel.

1. Go to **(Configure) > ESA Rules**.
   The Rules tab is displayed.
2. In the options panel, under **Deployments** click **Show Updates** on the far right.



The Updates to the Deployments dialog opens and shows the changes to the deployment.



3. Click **Close**.

Fast Deployment and Full Deployment of ESA Rules

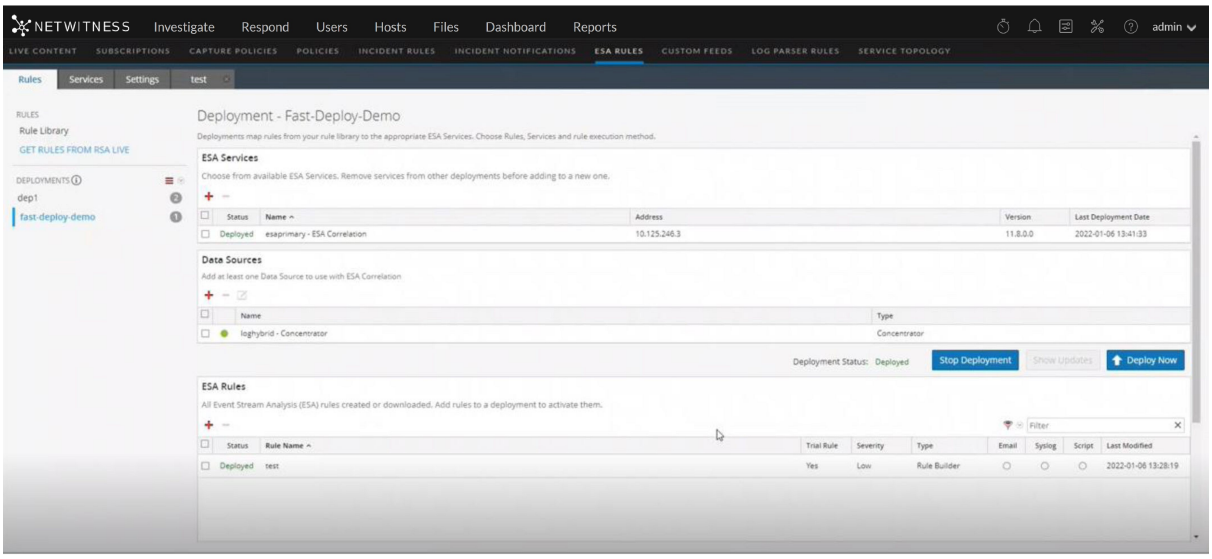
Only changes to an ESA rule is eligible for a fast deployment. This means if an ESA rule has been added, modified or deleted only then you can perform a fast deployment. In case, along with rules, if anything else has been changed such as a data source has been added, then the Fast Deployment option will not be displayed as it will not be eligible. The Fast Deployment option can be used to quickly load the latest rule that runs entirely on the Esper Engine hence reducing the need to redeploy the engine again in a shorter time where all the related settings of the selected deployment are deleted and the deployment is recreated from scratch.

The Full Re-Deploy can be used if:

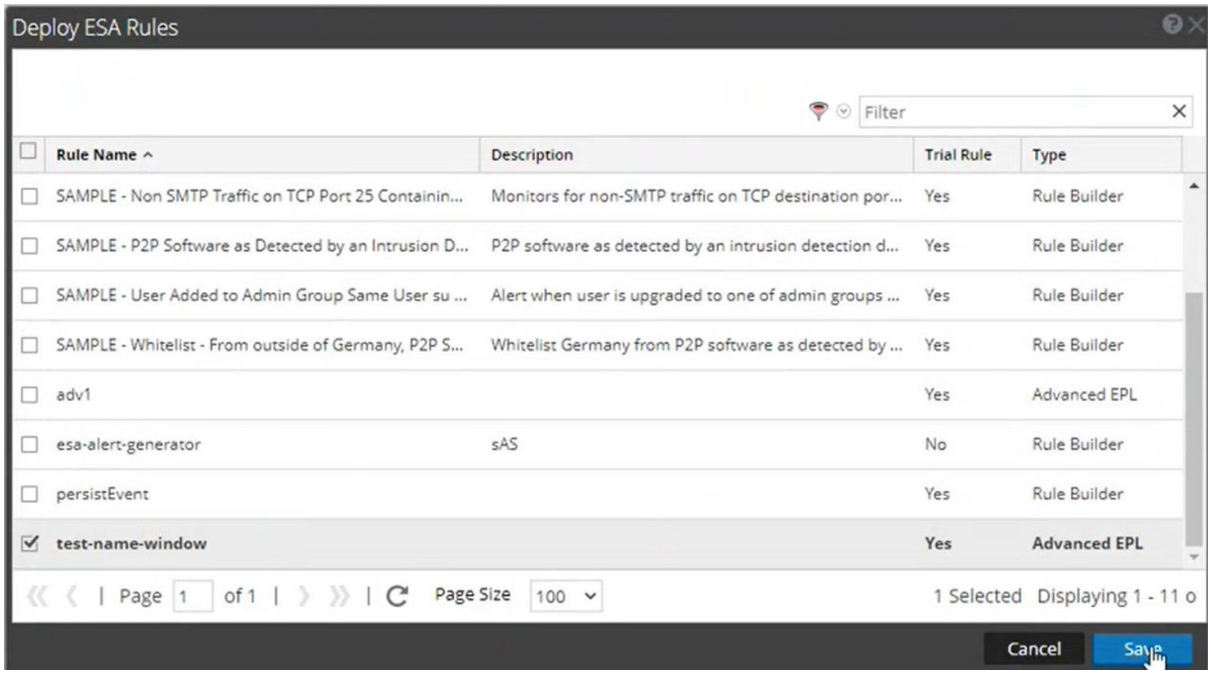
- The previously deployed rules failed, and the analyst wants to redeploy the same rules again.
- There are other changes made to the ESA Services, the Data Sources along with the ESA Rules.
- If the analyst wants to tear down the entire engine and recreate it.

To perform Fast Deployment:

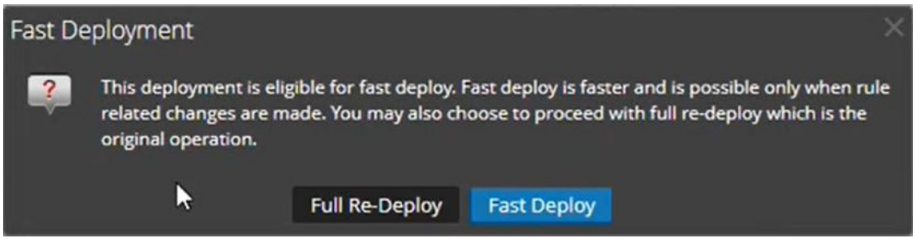
1. Go to  (**Configure**) > **ESA Rules**.
   The Rules tab is displayed.



2. Under ESA Rules, click  to deploy a new rule.
3. Select a rule you want to deploy from the list and click **Save**.



4. Click **Deploy Now**.
   A confirmation message is displayed.



5. Click **Fast Deploy** to begin the deployment of the new ESA Rule.

Note: The Fast Deploy button will be displayed only if the ESA rule has been added, deleted or updated, not if you have made updated to the ESA Services or Data Sources.

View ESA Stats and Alerts

When ESA generates alerts, you can view details about how the rules performed, such as statistics on the engine, rule, and alert, and you can also view information on which rules are enabled or disabled. For instructions on viewing ESA stats, see View Stats for an ESA Service

When your ESA generates alerts, you can view the results in the Respond Alerts List view. This enables you to see trends and understand both the volume and frequency of alerts. For instructions on viewing alerts, see View a Summary of Alerts

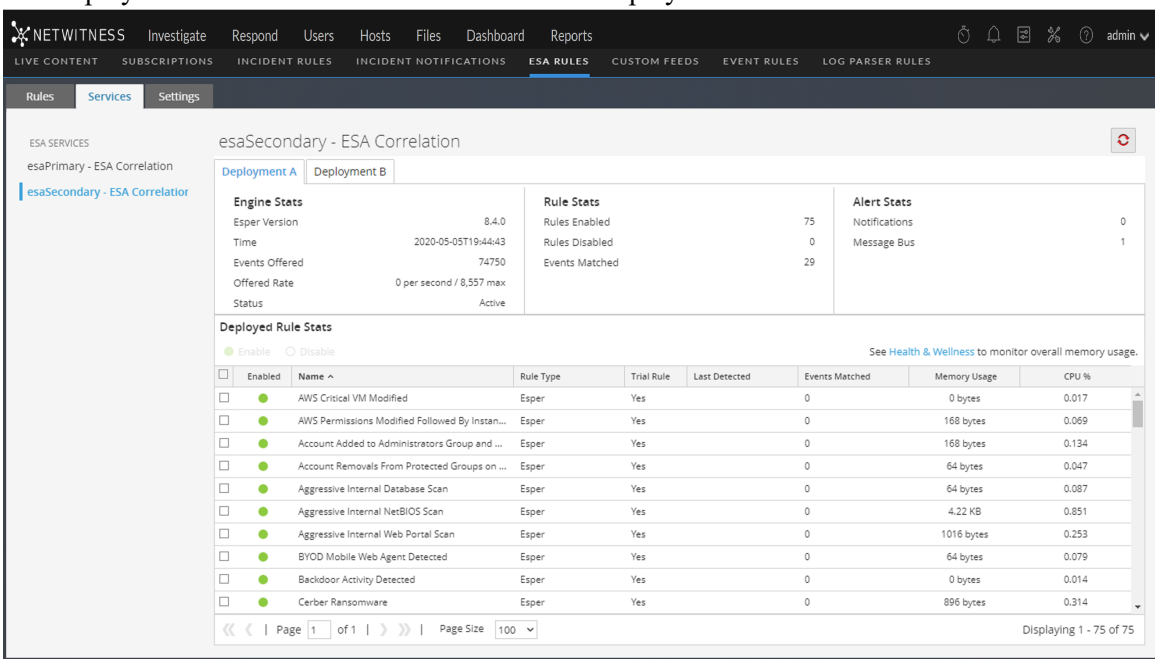
View Stats for an ESA Service

This topic describes how to view the deployment statistics (stats) for an ESA Correlation service. This procedure is useful when you are attempting to determine the effectiveness of a rule or troubleshoot an ESA rule deployment.

Caution: When you modify and re-deploy an ESA rule deployment, all of the stats are removed from that deployment. The generated alerts are not removed from NetWitness Respond.

View ESA Stats

1. Go to **(Configure) > ESA Rules > Services** tab.
2. From the **ESA Services** list on the left, select a service.
   The deployment stats for the selected service are displayed.



3. (This option applies to NetWitness version 11.3 and later.) In the Deployment view under the ESA Correlation service name, select the tab of the deployment you would like to view. For example,

select the Deployment A tab to view the stats for deployment A. Select the Deployment B tab to view the status for deployment B.

4. Review the following sections of ESA stats.
   For a complete description of each statistic in each section, see Services Tab.
   - **Engine Stats**
   - **Rule Stats**
   - **Alert Stats**
5. In the **Deployed Rule Stats**, review details about the rules deployed on the ESA.
   For a complete description of each column in each section, see Services Tab.
   - If the rule is enabled or disabled
   - What the rule name is
   - The type of rule
   - If the rule is running in Trial Rule mode
   - Last detected
   - Events matched
   - The amount of memory used by the rule
   - The percentage of the deployment CPU used by the rule (available in NetWitness version 11.5 and later)
6. To monitor overall memory usage and health of your ESA Correlation service, click **Health & Wellness**.

Enable or Disable Rules

1. In the **Deployed Rule Stats** panel, select a rule from the grid.
2. Click Enable to enable the rule, or click Disable to disable the rule.
   The Services tab is refreshed to show the changes, which take effect immediately.

Refresh the Statistics

The Services tab does not update statistics automatically unless you enable or disable a rule. To ensure you view current statistics:

1. Click in the upper right corner to refresh the information.
2. View the updated information.

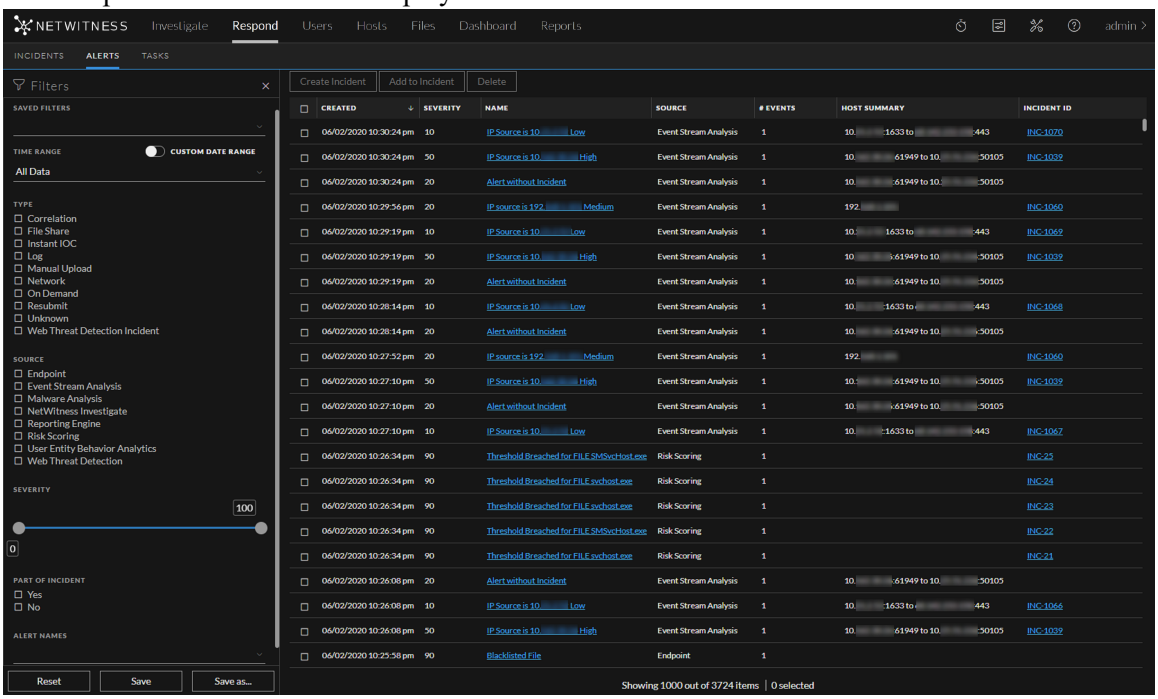
View a Summary of Alerts

In the Repond view, you can browse through various alerts from multiple sources. You can filter the alerts list to show only alerts of interest, such as by Alert Name, alert source, and a specific time range from the following sources:

1. Detect AI
2. Endpoint
3. Event Stream Analysis
4. Malware Analysis
5. NetWitness Investigate
6. Reporting Engine
7. Risk Scoring
8. User Entity Behavior Analysis
9. Web Threat Detection

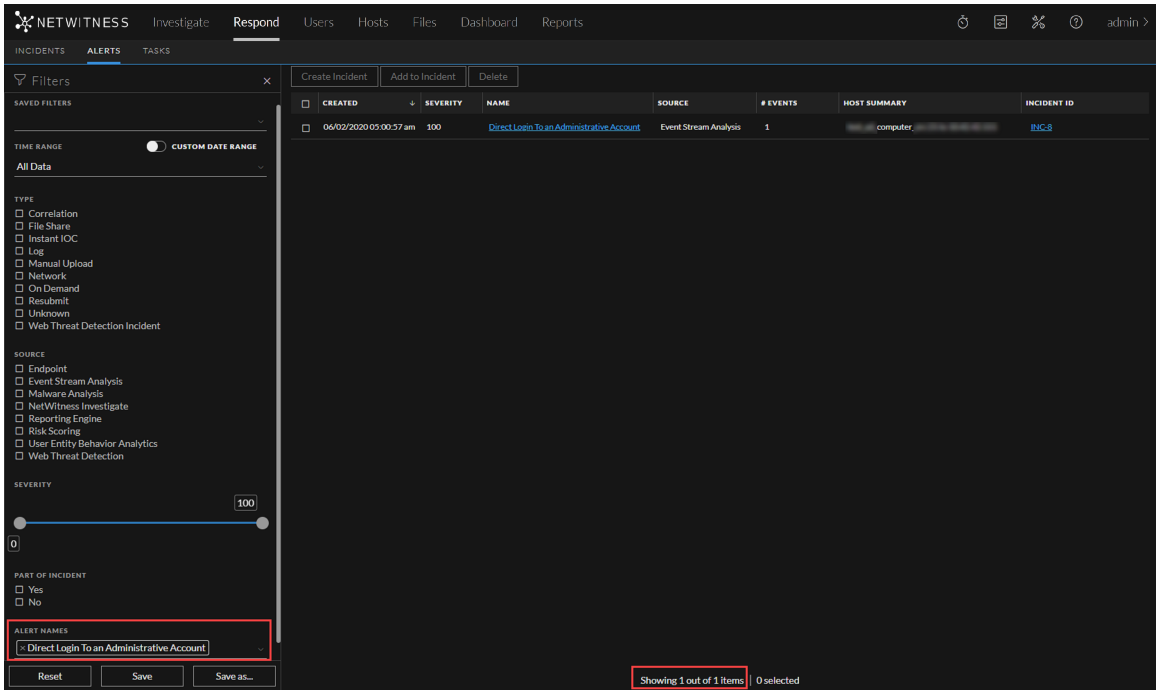
Perform the following steps to use the functionalities provided in the Respond view.

1. Go to **Respond > Alerts**.
   The Respond Alerts List view displays a list of all NetWitness alerts.



2. In the **Filters** panel on the left, you can filter the alerts list to view specific alerts for a specific time frame. For example, in the Alert Names section, you can select an alert for an ESA rule, such as

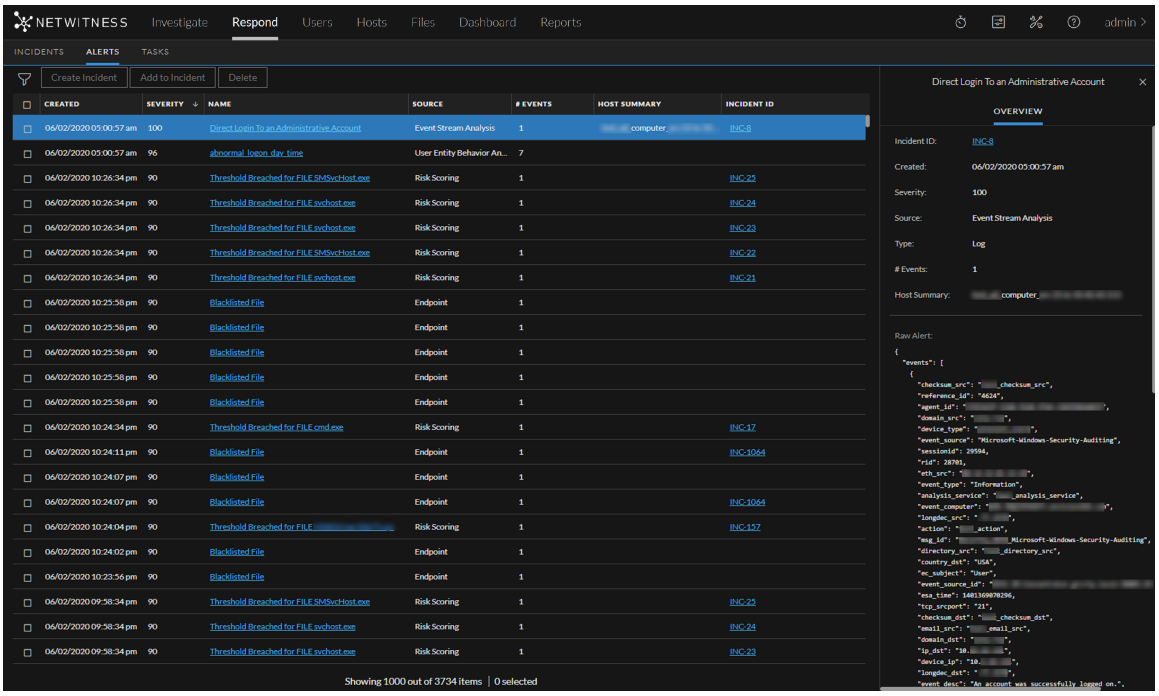
Direct Login to an Administrative Account, and leave the Time Frame set to Last Hour.
The alerts list to the right shows a list of alerts that match your filter selection along with a count of the alerts at the bottom of the alerts list.



The alerts list shows information about each of the alerts.

- **Created**: Displays the date and time when the alert was created in the source system.
- **Severity**: Displays the level of severity of the alert. The values are from 1 to 100.
- **Name**: Displays a basic description of the alert.
- **Source**: Displays the original source of the alert.
- **# of Events**: Indicates the number of events contained within an alert.
- **Host Summary**: Displays details of the host, like the host name from where the alert was triggered.
- **Incident ID**: Shows the incident ID of the alert. If there is no incident ID, the alert does not belong to an incident.

3. You can click an alert in the list to open an **Overview** panel on the right where you can view raw alert metadata.



For more information about filtering alerts and viewing alert details, see the *NetWitness Respond User Guide*.

Add an Advanced EPL Rule

This topic provides instructions to define rule criteria by writing an EPL query. EPL is a declarative language for handling high-frequency time-based event data. It is used to express filtering, aggregation, and joins over possibly sliding windows of multiple event streams. EPL also includes pattern semantics to express complex temporal causality among events.

Write an advanced EPL rule when rule criteria is more complex than what you can specify in Rule Builder.

It is outside the scope of this guide to explain EPL syntax.

- For EPL Documentation, see http://www.espertech.com/esper/esper-documentation/
- For the EPL Online Tool, see http://esper-epl-tryout.appspot.com/epltryout/mainform.html

For best practices on writing advanced EPL rules, see ESA Rule Writing Best Practices.

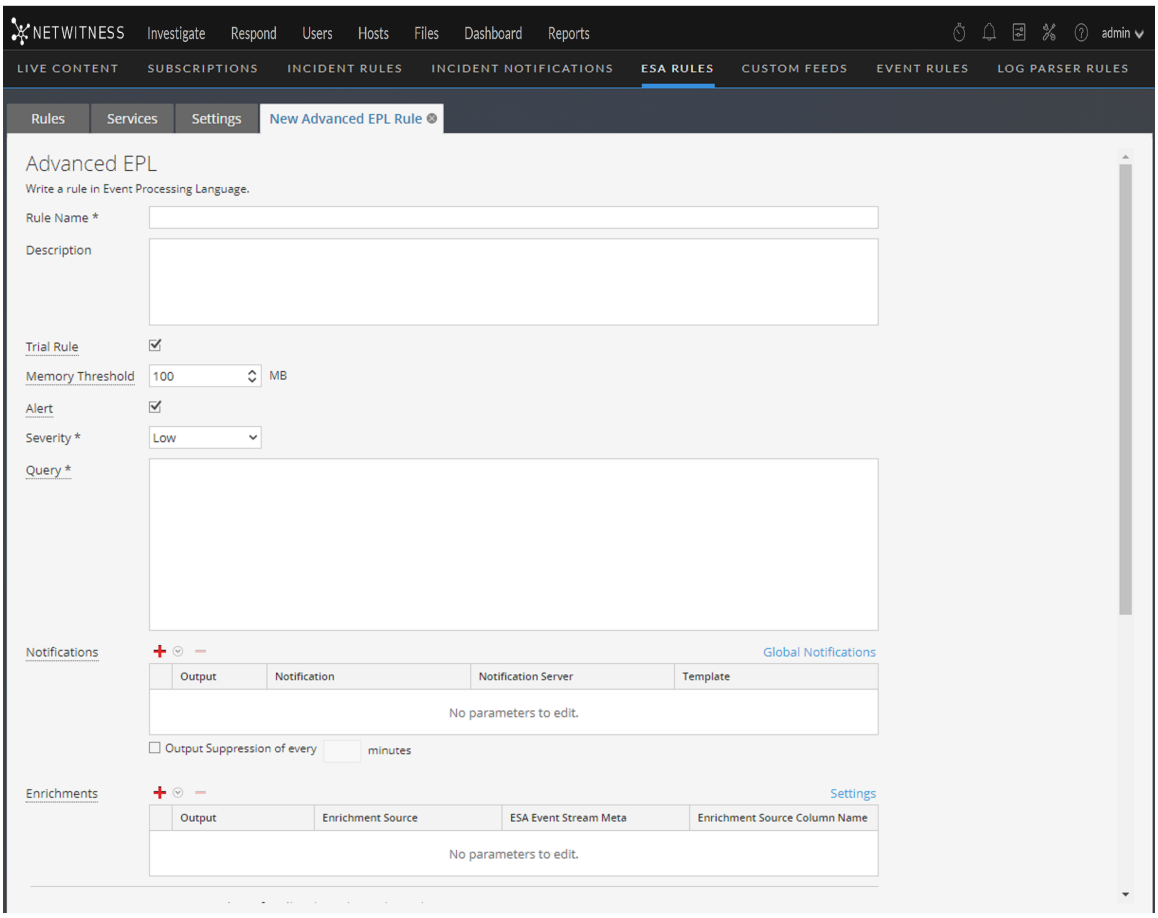
Prerequisites

The following are prerequisites for adding an advanced rule:

- You must know Event Processing Language (EPL).
- You must understand ESA Annotations to mark which EPL statements are linked to generating alerts.

Add an Advanced EPL Rule

1. Go to  **(Configure) > ESA Rules**.
2. In the **Rule Library**, select  **> Advanced EPL**.



3. Type a unique, descriptive name in the **Rule Name** field.

   This name will appear in the Rule Library so be specific enough to distinguish the rule from others.

4. In the **Description** field, explain which events the rule detects.

   The beginning of this description will appear in the Rule Library

5. Select **Trial Rule** to automatically disable the rule if all trial rules collectively exceed the memory threshold.

   Use trial rule mode as a safeguard to see if a rule runs efficiently and to prevent downtime caused by running out of memory. For more information, see Working with Trial Rules.

6. (This option applies to 11.5 and later.) Enter a **Memory Threshold** for a rule that uses memory, such as a rule that contains windows or pattern matching. If the configured memory threshold is exceeded, the rule gets disabled individually and an error is displayed for that rule on the (Configure) > ESA Rules > Services tab. The Memory Threshold option works for trial rules and non-trial rules. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.

7. (This option applies to 11.3 and later.) Select **Alert** to send an alert to Respond. Clear the checkbox if you do not want to send an alert to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.

8. For **Severity**, classify the rule as Low, Medium, High or Critical.
9. To define rule criteria, write a **Query** in EPL.

Note: For all meta key names, use an underscore not a period. For example, `ec_outcome` is correct but `ec.outcome` is not.

Supported meta entities:

S. No	Supported Meta Entities	Description
1	fullname_all	
2	eth_all	
3	ip_all	Combines all the IPv4 meta keys.
4	ipv6_all	Combines all the IPv6 meta keys.
5	port_src_all	
6	port_dst_all	
7	dir_path_all	
8	org_all	
9	geoip_all	
10	port_all	
11	domain_all	
12	email_all	
13	filename_all	
14	directory_all	
15	checksum_all	
16	param_all	
17	context_all	
18	attack_all	
19	analysis_all	
20	compromise_all	
21	inv_all	
22	outcome_all	
23	ec_all	
24	user_all	

25	host_all	
26	client_all	

10. For dynamic statement name generation in ESA, you must enclose the meta keys in curly brackets and include this annotation in the syntax:

    ```
    @Name("RIG {ip_src} {alias_host} {ec_activity}")
    ```

    where,

    - RIG is the static part of the statement name
    - {ip_src}, {alias_host}, {ec_activity} is the dynamic part of the statement name

    Note: If any of the metas in the dynamic part of the statement name has a null value, it is displayed as a static text.

    If a rule should generate an alert, include this ESA annotation in the syntax:

    ```
    @RSAAlert
    ```

    For more information on ESA Annotations, see ESA Annotations.

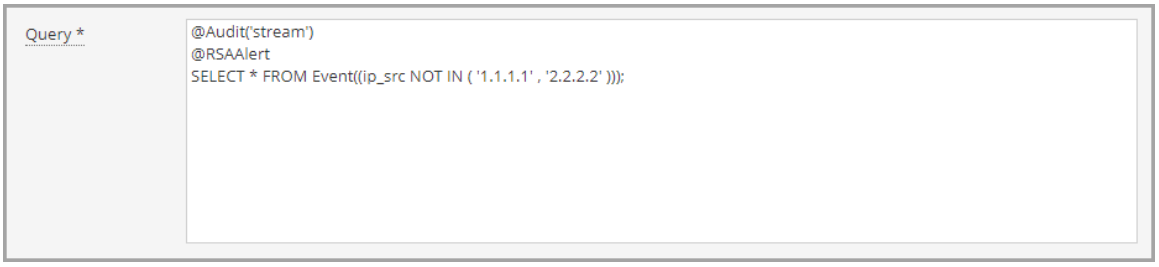
Validate an Advanced EPL Rule

You can confirm that an ESA rule generates the expected alerts by testing the rule logic using JSON input data. You can view the alerts in the output, but this test does not send any alert notifications.

If you want to view all of the debug information, include an `@Audit('stream')` annotation to your rule query and view the Debug Log in the test output. To enable auditing you require to add `@Audit` to the rule and set the logging level for the Esper audit package to INFO. This can be done by creating `correlation-server.yml` file under `/etc/netwitness/correlation-server` with this content and restarting the correlation-server service with systemctl restart rsa-nw-correlation-server:

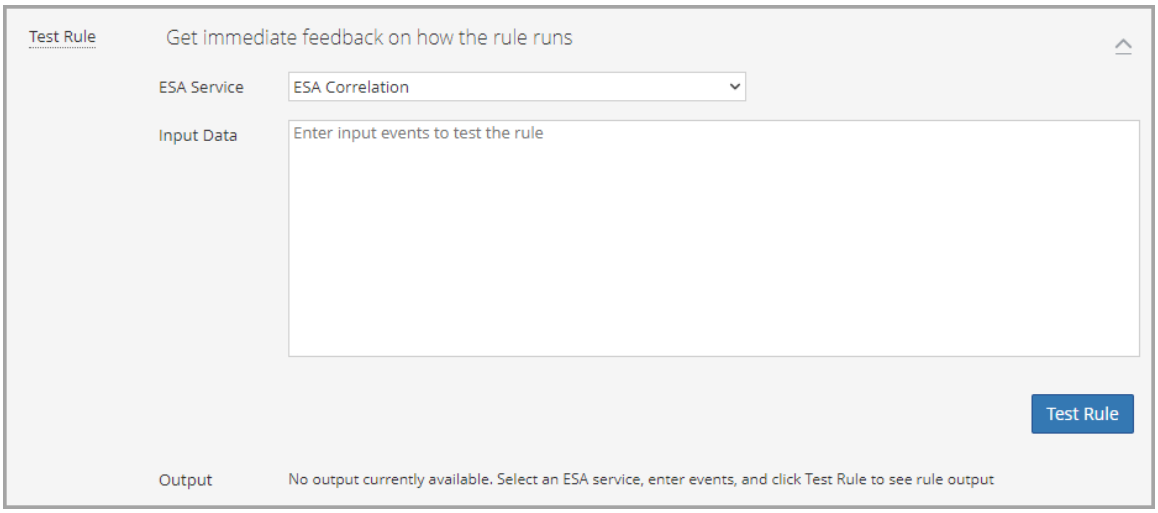
```
logging:
   level:
     com.espertech.esper.audit: INFO
```

The following basic example query contains the `@Audit('stream')` annotation and queries for events that do not have a source IP of 1.1.1.1 or 2.2.2.2.



1. If you are not already in the rule, go to **(Configure) > ESA Rules > Rules** tab and in the **Rule Library**, open the ESA rule that you want to test.

2. Scroll down to the **Test Rule** section.



3. In the **ESA Service** field, select the ESA Correlation service to process the rule. Use the same ESA Correlation service that you plan to use in the ESA rule deployment that contains the rule.
4. In the **Input Data** field, enter the input events to test the rule. Download the events from the Investigate view in JSON format, copy the events, and paste them in this field. You can do this from the Investigate > Navigate view or the Investigate > Events view.

   To download the events from the Investigate > Navigate view:

   a. In the main menu, go to **Investigate > Navigate** in a new tab, select a data source, and click **Navigate**.
   b. In the Navigate view, click **Load Values** and click a meta value to filter the events.
   c. Save the events as meta in the JSON file format [Save Events > Meta > (name the file) > Export Meta Format: choose JSON].
   d. In the toolbar click the **(Jobs)** icon and then click **View Your Jobs**.
   e. In the Jobs panel, download your extracted meta, for example: **`investigation-2020-May-19-08-30-20.json`**.
   f. Go to back to the **(Configure) > ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.

   To download the events from the Investigate > Events view:

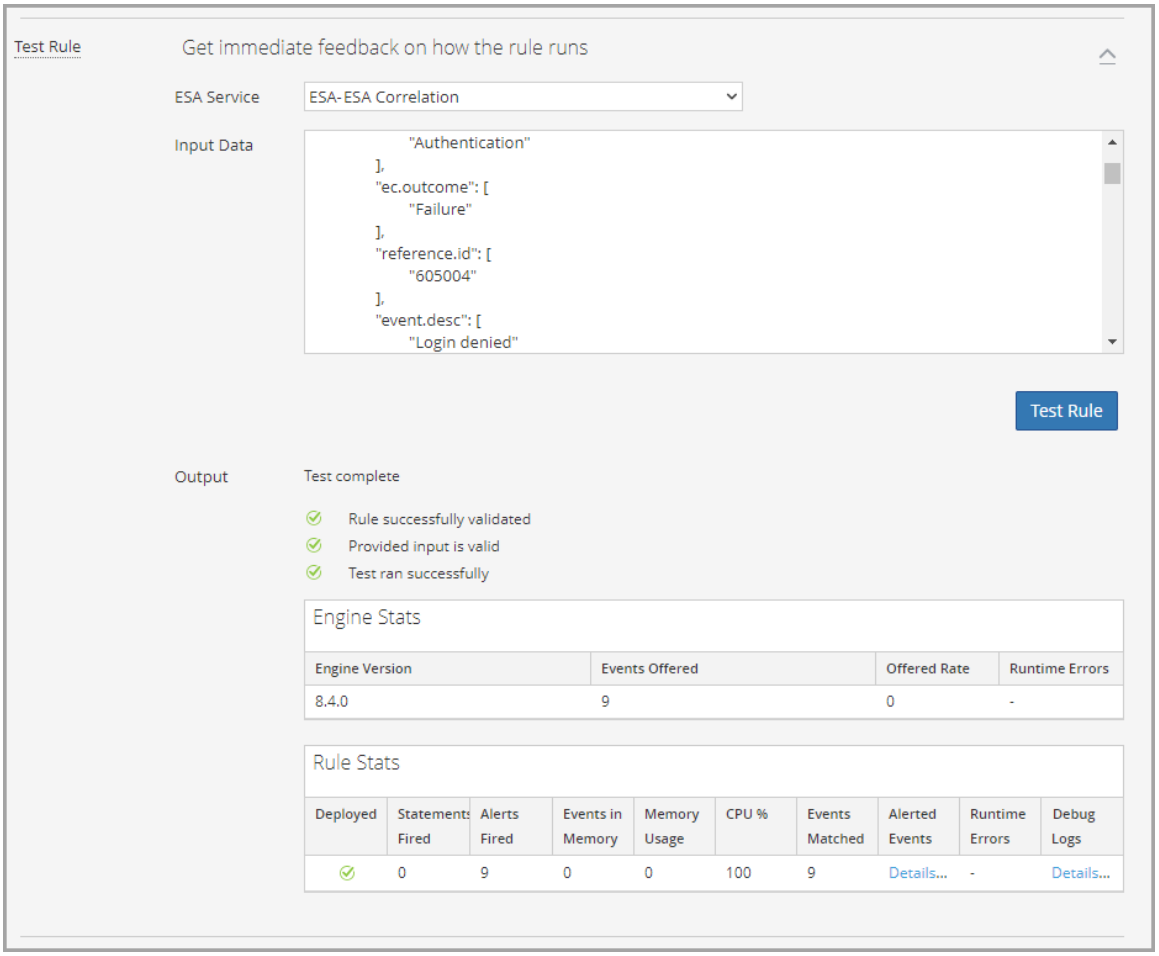
   a. In the main menu, go to **Investigate > Events** in a new tab.
   b. In the Events view, enter a query for the ESA rule test.
   c. Select the events to use and in the **Download** or **Download All** menu, select **Visible Meta as JSON** or **All Meta as JSON**, depending on the size of your selection.

d. In the main menu, go to **Dashboards** and in the toolbar click the **(Jobs)** icon and then click **View Your Jobs**.

e. In the Jobs panel, download your extracted meta, for example: **Concentrator_ALL_EVENTS_ALL_META.json**.

f. Go to back to the **(Configure) > ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.

5. Click **Test Rule**. The **Output** field shows the output of your rule and you can determine if the results meet your requirements.



Note: If you are testing any Rule that has meta key defined as type 'short', the Test Rule will not generate alert for the event.

The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service

Field	Description
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Event Processing Language (EPL)

This topic describes Event Processing Language (EPL), a declarative language for dealing with high frequency time-based event data. ESA uses Event Processing Language (EPL), a declarative language for dealing with high frequency time-based event data. It is used for express filtering, aggregation, and joins over possibly sliding windows of multiple event streams. EPL also includes pattern semantics to express complex temporal causality among events. It can perform, but is not limited to, the following functions:

- Filter Event
- Alert Suppression
- Compute percentages or rations
- Average, count, min and max for a given time window
- Correlate events arriving in multiple stream
- Correlate events that arrive out of order
- On-Off Windows
- Followed-by and Not Followed-by support
- Regex filter support

Databases require explicit querying to return meaningful data and are not suited to push data as it changes. The developer must implement the temporal and aggregation logic himself. By contrast, the EPL engine provides a higher abstraction and intelligence and can be thought of as a database turned upside-down. Instead of storing the data and running queries against stored data, EPL allows applications to store queries and continuously run the data through. Response from the EPL engine is real-time when conditions occur that match user defined queries.

For best practices on writing advanced EPL rules, see ESA Rule Writing Best Practices.

Advanced ESA rules require correct character case, but in the Investigate Navigate view all characters are converted to lowercase. However, the meta may not be lowercase despite appearances in the Investigate Navigate view. To ensure you are using the correct case, you can make a strict pattern match for better performance.
Strict Pattern Match Example

```
@RSAAlert(oneInSeconds=0) SELECT * FROM Event( (medium IN ( 1 ) AND filetype IN ( 'pdf' , 'windows_executable' , 'x86 pe' , 'windows executable' )) ).win:time(5 Minutes) MATCH_RECOGNIZE ( MEASURES E1 as e1_data , E2 as e2_data PATTERN (E1 E2) DEFINE E1 as (E1.filetype IN ('pdf')), E2 as (E2.filetype IN ( 'pdf' , 'windows_ executable' , 'x86 pe' , 'windows executable' )) );
```

Caution: Care should be taken to only add the case-insensitive *toLowerCase()* function on meta keys as needed. The *toLowerCase()* function can cause significant performance decreases. Consider checking the Investigate Events view or the Event Analysis view to see the real character case for meta fields and avoid unnecessary usage of the function.

For the purposes of online help, basic statements are used to illustrate how to set up ESA; however, for more information about writing EPL statements, the http://www.espertech.com site provides tutorials and examples.

Note: In NetWitness version 11.5, ESA Correlation supports Esper version 8.4.0. In version 11.4, ESA Correlation supports Esper version 8.2.0 and in version 11.3, ESA Correlation supports Esper version 7.1.0.

ESA Annotations

This topic describes annotations that NetWitness provides to use in advanced EPL rules.

For best practices on writing advanced EPL rules, see ESA Rule Writing Best Practices.

@RSAContext Annotation (11.5 and later)

The @RSAContext annotation can be used in advanced rules to dynamically add or remove data from a Context Hub list after the rule fires. For example, you can create a rule that automatically adds an IP address to a blacklist and removes it from a whitelist.

You can update a single-column or a multi-column Context Hub list. The @RSAContext annotation also performs error handling when the Context Hub list cannot be reached.

Prerequisites

Before deploying a rule using the @RSAContext annotation, the list must exist in Context Hub. For information on creating a Context Hub list, see the *Context Hub Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

IMPORTANT: If you rename a Context Hub list or recreate the Context Hub list with the same name, update the ESA rules that use that Context Hub list, and then redeploy the ESA rule deployments that contain those rules.

Single-Column Context Hub Lists

The @RSAContext annotation uses the following format for a single column Context Hub list:

```
@RSAContext (list="<single_column_list>", action=<DELETE_ENTRY or ADD_ENTRY>,
onError = <STOP_ALL_RULE_PROCESSING_AND_WAIT or IGNORE_ERROR_AND_CONTINUE>,
fields={"LIST=<meta_key>"})
```

The @RSA Context parameters for a **single-column** Context Hub List are described in the following table.

Parameter	Description
`list`	Where `<single_column_list>` is the name of the single-column Context Hub list (whitelist or blacklist).
`action`	You can add an entry to or remove an entry from the list: • `DELETE_ENTRY` - Removes an entry from the list. • `ADD_ENTRY` - Adds an entry to the list. You can only have one action per `@RSAContext` entry.

Parameter	Description
onError	Identifies how to handle errors, for example, when the Context Hub list is not available or is full. You can choose one of the following options: • STOP_ALL_RULE_PROCESSING_AND_WAIT: If any error occurs with the selected actions on the Context Hub list, stop all rule processing including data aggregation and retry indefinitely until the actions are successfully executed. Halting processing for one Context Hub list results in halting processing for all Context Hub lists. Select this option if it is important to update a Context Hub list before processing any more rules. For example, you may have a blacklist that must be updated before continuing rule processing. • IGNORE_ERROR_AND_CONTINUE: If any error occurs with the selected actions on the Context Hub list, ignore the error and the @RSAContext entry, and continue processing. If it is not a Context Hub error, assume that the error occurs before reaching Context Hub, and retry until the actions are successfully executed. Select this option if continuing rule processing is more important than updating the Context Hub list.
fields	Maps a meta key value to the column in the Context Hub list. Single-column lists always show LIST= in capital letters.

The first example deletes a source IP address from an IP whitelist. However, if the list is not available, the source IP address is not removed from the list and processing is continued.

```
@RSAContext (list="ip_whitelist", action=DELETE_ENTRY, onError = IGNORE_ERROR_
AND_CONTINUE, fields={"LIST=ip_src"})
```

The second example adds a source IP address to an IP blacklist. However, if the list is not available, all rule processing stops including data aggregation, and the ESA Correlation service retries indefinitely until the entry is added to the specified Context Hub list.

```
@RSAContext (list="ip_blacklist", action=ADD_ENTRY, onError = STOP_ALL_RULE_
PROCESSING_AND_WAIT, fields={"LIST=ip_src"})
```

Multi-Column Context Hub Lists

The @RSAContext annotation uses the following format for a multi-column Context Hub list:

```
@RSAContext (list="<multi_column_list>", action=<DELETE_ENTRY or ADD_ENTRY>,
onError = <STOP_ALL_RULE_PROCESSING_AND_WAIT or IGNORE_ERROR_AND_CONTINUE>,
fields={"<column1>=<meta_key1 or constant value1>",<column2>=<meta_key2 or
constant value2>",<column n>=<meta_key n or constant value n>",})
```

The @RSA Context parameters for a **multi-column** Context Hub List are described in the following table.

Parameter	Description
list	Where <multi_column_list> is the name of the multi-column Context Hub list.
action	You can add an entry to or remove an entry from the list: • DELETE_ENTRY - Removes an entry from the list. • ADD_ENTRY - Adds an entry to the list. You can only have one action per @RSAContext entry.
onError	Identifies how to handle errors, for example, when the Context Hub list is not available or is full. You can choose one of the following options: • STOP_ALL_RULE_PROCESSING_AND_WAIT: If any error occurs with the selected actions on the Context Hub list, stop all rule processing including data aggregation and retry indefinitely until the actions are successfully executed. Halting processing for one Context Hub list results in halting processing for all Context Hub lists. Select this option if it is important to update a Context Hub list before processing any more rules. For example, you may have a blacklist that must be updated before continuing rule processing. • IGNORE_ERROR_AND_CONTINUE: If any error occurs with the selected actions on the Context Hub list, ignore the error and the @RSAContext entry, and continue processing. If it is not a Context Hub error, assume that the error occurs before reaching Context Hub, and retry until the actions are successfully executed. Select this option if continuing rule processing is more important than updating the Context Hub list.
fields	Maps a meta key values or constant values to the columns in the Context Hub list. Multi-column lists show the names of each column in the fields parameter.*

The first example deletes a user and the associated source and destination IP addresses from a multi-column IP whitelist. However, if the list is not available, the specified fields are not removed from the list and processing is continued.

```
@RSAContext (list="MultColumn_whitelist", action=DELETE_ENTRY, onError =
IGNORE_ERROR_AND_CONTINUE, fields={"source=ip_src","destination=ip_
dst","user='smith'"})
```

The second example adds a user and the associated source and destination IP addresses to a multi-column IP blacklist. However, if the list is not available, all rule processing stops including data aggregation, and the ESA Correlation service retries indefinitely until the entry is added to the specified Context Hub list.

```
@RSAContext (list="MultColumn_blacklist", action=ADD_ENTRY, onError = STOP_
ALL_RULE_PROCESSING_AND_WAIT, fields={"source=ip_src","destination=ip_
dst","user='smith'"})
```

Automatic Context Hub List Updates

When using the @RSAContext annotation in your rules, string array meta keys (shown as `string[]` in the Meta Key References on the ESA Rules > Settings tab) are inserted and removed as separate lines in the Context Hub list. This enables lookups in the Investigate and Respond views.

Single Column Context Hub List Update Example

A single column Context Hub list has one column with the string array meta key `alias_host`. If the rule using that list fires and `alias_host` has three values (Google, Yahoo, and Dell), ESA Correlation adds three rows to the Context Hub list:

Google
Yahoo
Dell

Multi-Column Context Hub List Update Example

A multi-column Context Hub list has three columns. If the rule using that list fires and the meta keys have the following values:

- `username` (string array) with values Aimee and Chris
- `ip_dst` (string) with value 10.10.10.10
- `alias_host` (string array) with values Dell, Google, and Yahoo

ESA Correlation adds six rows to the Context Hub list:

Aimee, 10.10.10.10, Dell
Aimee, 10.10.10.10, Google
Aimee, 10.10.10.10, Yahoo
Chris, 10.10.10.10, Dell
Chris, 10.10.10.10, Google
Chris, 10.10.10.10, Yahoo

@RSAAlert Annotation

The @RSAAlert annotation is used to mark which EPL statements are linked to generating alert notifications. It is designed to work with the alert notification suppression feature in the Rule Builder user interface.

The @RSAAlert annotation can be useful when working with alert notifications, especially if you want to filter notifications, such as sending one notification for each user that triggers an alert.

For example, suppose you want to generate alert notifications for login failures. You could add the following statement:

*@RSAAlert select * from event(msg_id="login_fail")*

Event number	Message ID	username	src_ IP	Time
1	login_fail	alice	1.2.3.4	10:00
2	login_fail	alice	1.2.3.4	10:01
3	login_fail	alice	6.7.8.9	10:01
4	login_fail	bob	1.2.3.4	10:01
5	login_fail	alice	1.2.3.4	10:03

For the above statement, five alert notifications are generated.

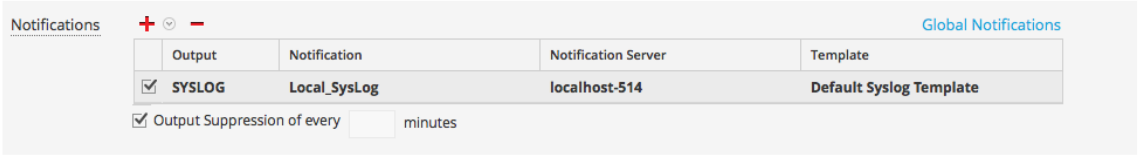
However, suppose you wanted to modify the statement to generate one alert for each separate username. You can use the *identifier* attribute. For example, the statement *@RSAAlert(identifier={"username"}) SELECT* FROM Event(msg_id="login_fail")* generates one notification for the first alert for "bob" and one for the first alert for "alice." Subsequent alerts for "bob" and "alice" are ignored.

You can further distinguish the users by adding details via the identifier variable. For example, you can distinguish by user and IP address using the following statement: *@RSAAlert(identifier={"username", "src_ip"}) SELECT* FROM Event(msg_id="login_fail")*. Then, you would see notifications generated by user name and IP address (one alert for "alice" at 1.2.3.4, another alert for "alice" at 6.7.8.9, and an alert for "bob" at 1.2.3.4).

To use identifiers with Alert Notification Suppression:

The @RSAAlert annotation is designed to work with the alert notification suppression feature in the Rule Builder user interface. To do this:

1. Create a rule in the Rule Builder user interface, and select the alert suppression feature when configuring notifications.



2. Copy the code from the Rule Builder rule into a new advanced rule.
3. Configure the advanced rule to include identifiers (as described above) and save the advanced rule.
4. Delete the original rule builder rule.

@RSAPersist Annotation

The @RSAPersist annotation is used to mark a named window as an ESA managed window for persistence. By marking the named window as an ESA managed window, ESA periodically writes the contents of the window to disk and restores them back if the window is undeployed and redeployed. The systems take a snapshot just before the ESA rule deployment is undeployed and the window is removed. Conversely, it restores the window contents from the snapshot just after the deployment is redeployed. This ensures that the contents of the window are not lost if the deployment state is altered or if the ESA service goes down.

For example, consider a named window, `DHCPTracker` that holds a mapping from IP addresses to each assigned hostname.You can annotate the statement with the @RSAPersist annotation as:

```
@RSAPersist
 create window DHCPTracker.std:unique(ip_src) as (ip_src string, alias_
host string);
 insert into DHCPTracker select IP as ip_src, HostName as alias_host from
DHCPAssignment(ID=32);
```

Note: All windows definitions are not suitable for persistence. `@RSAPersist` annotation must be used with care. If the window has timed-records or if it depends on time based constraints it is very likely that the reverted snapshots will not restore it to the correct state. Also, any changes to the window definition will invalidate the snapshots and reset the window to a blank state. The system does not do any semantic analysis to determine if the changes to the window definition are conflicting or not. Note that other parts of a deployment (that is, other than the particular CREATE WINDOW call that defines the window) may change, without invalidating the snapshots.

Caution: (This caution applies only to NetWitness Platform versions 11.3.x, 11.4.x, and 11.5.0.0.) To avoid data being overwritten, if you have a rule with a named window, do not disable and re-enable it. Instead, undeploy and redeploy the ESA rule deployment that contains the rule.

@UsesEnrichment (10.6.1.1 and later)

The @UsesEnrichment can be used in advanced EPL rules to reference enrichments. In order to synchronize enrichments with ESA, all enrichment dependencies in EPL rules must be referenced with the @UsesEnrichment annotation.

The @UsesEnrichment annotation uses the following format:

```
@UsesEnrichment(name= '<enrichment_name>')
```

For example, the following EPL references a whitelist enrichment:

```
@UsesEnrichment(name = 'Whitelist')
@RSAAlert
SELECT * FROM Event(ip_src NOT IN (SELECT ip_address FROM Whitelist))
```

@Name

The @Name is the statement name defined in ESA advanced rules. It is used to dynamically generate statement names in ESA alerts. The statement name of only an alert triggering statement is displayed. This annotation has meta keys enclosed in curly brackets.

The @Name annotation uses the following format:

```
@Name("<static_part_of_statement_name> {meta_key1} {meta_key2}…")
```

For example, the following EPL references meta keys *ip_src* and *user_name* whose values will be dynamically generated.

@Name(“Login Event to {ip_src} by {user_name}”)

Note: You can specify any number of meta keys in the statement for dynamic statement name generation.
The length of individual meta key is limited to 64, after which the value is truncated and appended with “…”.
The length of the dynamic generation of statement name is limited to 128, after which the value is truncated to 128 and appended with “…”. All the remaining values post truncation will be treated as static values.

@Audit Annotation

Add the `@Audit('stream')` annotation to your ESA rules to print alerts to the ESA logs for troubleshooting. This is useful when debugging the Esper rules. The @Audit('stream') annotation provides debug information for the next statement in the rule. For example:

```
@Audit('stream')
@RSAAlert
SELECT * FROM Event((ip_src NOT IN ( '1.1.1.1' , '2.2.2.2' )));
```

In the above example, `@RSAAlert` is only necessary if the statement needs to send an alert to the Respond view.

In NetWitness Platform 11.5 and later, you can test rules in the rule builders. If you add the `@Audit ('stream')` annotation to an advanced EPL rule, you can view the Debug Log in the test output. For more information, see Validate an Advanced EPL Rule .

Example Advanced EPL Rules

Following are the examples of Advanced ESA rules. Each example has multiple ways of implementing the same use-case.

For best practices on writing advanced EPL rules, see ESA Rule Writing Best Practices.

Example #1:

Create a user account and delete the same user account in 300s. User information is stored in user_src meta.

EPL #1:

Rule Name	CreateAndDelete Useraccount1
Rule Description	Create a user account followed by an action to delete the same user account in 300 seconds.

Rule Code	`@RSAAlert(oneInSeconds=0)`
	`SELECT * FROM Event(ec_subject='User'`
	`AND ec_outcome='Success'`
	`AND user_src is NOT NULL`
	`AND ec_activity IN ('Create', 'Delete')`
	`).win:time(300 seconds)`
	`match_recognize (partition by user_src`
	`measures C as c, D as d`
	`pattern (C D)`
	`define`
	`C as C.ec_activity='Create' ,`
	`D as D.ec_activity='Delete');`

Note	• Filter events needed for pattern in given time frame. Filter conditions should be such that only required events are passed to match recognize function. In this case, they are create and delete user account Events. That is, Event(ec_subject='User' AND ec_outcome='Success' AND user_src is NOT NULL AND ec_activity IN ('Create', 'Delete') • Partition by creates buckets. In this case, Esper creates buckets per value of user_src. And hence value of user_src is common between both events. • Define the pattern you want. Right now it is set to Create Followed by Delete. You can do multiple creates followed by delete (C+ D). Pattern is very similar to regular expression. • Most efficient use case. • The 'loose' pattern match of (C+ D) will result in decreased performance. Unless you need to include all C events within the generated alert, keep the strict pattern match of (C D). See the Esper documentation for more details.

EPL #2:

Rule Name	CreateAndDeleteUseraccount2
Rule Description	Create a user account followed by an action to delete the same user account in 300 seconds.

Rule Code	`@RSAAlert(oneInSeconds=0)`
	`SELECT * from pattern[every (a= Event(ec_subject='User' AND ec_outcome='Success' AND user_dst is NOT NULL AND ec_activity IN ('Create'))`
	`->`
	`( Event(ec_subject='User' AND ec_outcome='Success' AND user_dst is NOT NULL AND ec_activity IN ('Create') AND user_src = a.user_src)`
	`) )where timer:within(300 Sec) ];`

Note	• Lets say same user is created twice and deleted once in that order. Then the above pattern will fire 2 alerts. • A thread is created for every User creation. • There is no way to control threads. It is important to have time bounds and preferably small intervals. • If you do not need every first event to start a new thread and match with the subsequent second event, then add suppression syntax of @SuppressOverlappingMatches after the pattern keyword. See the Esper documentation for more details.

Example #2:

Detect pattern where user created followed by login by same user and user is deleted in end. In case of windows logs user info is stored in either user_dst or user_src depending on event.

user_src(create) = user_dst(Login) = user_src(Delete)

EPL #3:

Rule Name	CreateUserLoginandDeleteUser
Rule Description	Detect a pattern where a user creates a User account followed by login by the same user followed by deletion of the User account.

Rule Code	`@RSAAlert(oneInSeconds=0)`
	`SELECT * FROM Event(ec_subject='User'`
	`and ec_activity in ('Create','Logon','Delete')`
	`and ec_theme in ('UserGroup', 'Authentication')`
	`and ec_outcome='Success'`
	`).win:time(300 seconds)`
	`match_recognize (measures C as c, L as l, D as d`
	`pattern (C L D)`
	`define`
	`C as C.ec_activity = 'Create',`

	L as L.ec_activity = 'Logon' AND L.user_dst = C.user_src, D as D.ec_activity = 'Delete' AND D.user_src = C.user_src);
Note	• Since user_src/user_dst is not common across all events we can't use partition. It will be 1 single bucket running 1 pattern at a time. For example, for user 1 and 2 if the stream of events are C1C2L1D1, C1L1C2D1, there will be no alert because C1 thread got reset by C2. Alert will be fired only if C1L1D1 are in order and no other event either from same user or other user falls in between. • Another solution would be to use Named Window and merge user_dst and user_src into single column and then run match recognize. (EPL #3). • Pattern can also be used. You might get more alerts than expected. (EPL #4).

EPL #4: Using NamedWindows and match recognize

Rule Name	CreateUserLoginandDeleteUser
Rule Description	Detect a pattern where a user creates a User account followed by login by the same user followed by deletion of the User account.
Rule Code	@Name('NormalizedWindow') create window FilteredEvents.win:time (300 sec) (user String, ecactivity string, sessionid Long); @Name('UsersrcEvents') Insert into FilteredEvents select user_src as user, ec_activity as ecactivity, sessionid from Event (ec_subject='User' and ec_activity in ('Create','Delete') and ec_theme in ('UserGroup', 'Authentication') and ec_outcome='Success' and user_src is not null); @Name('UsrdstEvents') Insert into FilteredEvents select user_dst as user, ec_activity as ecactivity, sessionid from Event(ec_subject='User' and ec_activity in (Logon') and ec_theme in ('UserGroup', 'Authentication') and ec_outcome='Success' and user_dst is not null); @Name('Pattern') @RSAAlert(oneInSeconds=0, identifiers={"user"}) select * from FilteredEvents match_recognize (partition by user measures C as c, L as l, D as d pattern (C L+D) define C as C.ecactivity= 'Create', L as L.ecactivity= 'Logon', D as D.ecactivity='Delete');

EPL #5: Using Every @RSAAlert(identifiers={"user_src"})

Rule Name	CreateUserLoginandDeleteUser

Rule Description	Detect a pattern where a user creates a User account followed by login by the same user followed by deletion of the User account.
Rule Code	`SELECT a.time as time,a.ip_src as ip_src,a.user_dst as user_dst,a.ip_dst as ip_dst,a.alias_host as alias_host from pattern [every (a=Event (ec_subject='User' and ec_activity='Create' and ec_theme='UserGroup' and ec_outcome='Success') -> (Event(ec_subject='User' and ec_activity='Logon' and ec_theme='Authentication' and user_src=a.user_dst) -> b=Event(ec_subject='User' and ec_activity='Delete' and ec_theme='UserGroup' and user_dst=a.user_dst))) where timer:within(300 sec)];`

Example #3:

Excessive login failures from same sourceIP.

EPL #6: @RSAAlert(identifiers={"ip_src"})

Rule Name	ExcessLoginFailure
Rule Description	The same user tried logging in from the same Source IP and faced login failures.
Rule Code	`@RSAAlert(oneInSeconds=0) SELECT * FROM Event ( ip_src IS NOT NULL AND ec_activity ='Logon' AND ec_outcome = 'Failure').win:time_batch(300 seconds) GROUP BY ip_src HAVING COUNT(*) = 10;`
Note	• Uses time _batch: Looks at events in batches(tumbling window). Every event matching the filter criteria will be kept for the specified time window. • "GROUP BY" clause aggregates events within the data window by ip_src and HAVING clause instructs a count of 10 events with the same ip_src must occur within the time window. • One of the issues with tumbling windows is that events occurring towards the end of the batch might not lead to an alert. In the below sequence of events at t=301 even though 10 login failures occurred for the same login in the last 300 secs, there will be no alert because the batch of events was dropped at t=300.

Time t	Login Failures for Specific Users	Alert	Time Batch
0	0	0	1
295	6	0	1
299	3	0	1
301	1	0	2
420	6	0	2
550	3	0	2
600	0	0	3
720	6	0	3

	850 3 0 3 900 1 1 3 ends and 4 begins • Above problem can be resolved using win:time windows (EPL#7)instead of win:time_length_batch windows. • Outer group by is to control events when time elapses. Say you have 9 events at end of 60 secs, Esper engine will push those 9 events to listener. Group by and count will restrict it since count is not equal to 10. • Time and count can be modified as needed.

EPL #7: @RSAAlert(identifiers={"ip_src"})

Rule Name	ExcessLoginFailure
Rule Description	The same user tried logging in from the same Source IP and faced login failures.

Rule Code	`@RSAAlert(oneInSeconds=0)`
	`SELECT * FROM Event ( ip_src IS NOT NULL AND ec_activity ='Logon'`
	`AND ec_outcome = 'Failure').win:time(300 seconds) GROUP BY ip_src HAVING COUNT(*) = 10;`

Note	• This is a sliding window and hence after an alert is fired for a set of events they can be used for another alert as well until time has passed. • If 10 events were involved in causing the alert only the last event will appear. • Events are not removed from the time window. You could use output rate limiting. See the Esper documentation for more details.

Example #4:

Multiple failed logins from multiple different users from same source to same destination, a single user from multiple different sources to same destination.

EPL #8: using time_batch

Rule Name	MultiplefailedLogins
Rule Description	There are multiple failed logins for the following cases: - From multiple users from same source to same destination. - Single user from multiple sources to the same destination.

Rule Code	`@RSAAlert(oneInSeconds=0)`
	`SELECT * FROM Event (`
	`ec_activity='Logon'`
	`AND ec_outcome='Failure'`
	`AND ip_src IS NOT NULL`

	AND ip_dst IS NOT NULL AND user_dst IS NOT NULL) .win:time_batch(300 seconds) group by ip_src,ip_dst having count(distinct user_dst) >= 5;
Note	• ip.dst and ip.src are common across all events. • user_dst is unique for all events. • Alert is fired when there are at least 5 different users try to login from same ip.src and ip.dst combination.

Example #5:

No Log traffic from a device in a given timeframe.

EPL #9: using timer:interval

Rule Name	NoLogTraffic
Rule Description	There is no log traffic observed from a device in a given time frame.
Rule Code	SELECT * FROM pattern [every a = Event(device_ip IN ('10.0.0.0','10.0.0.1') AND medium = 32) -> (timer:interval (3600 seconds) AND NOT Event(device_ip = a.device_ip AND device_type = a.device_type AND medium = 32))];
Note	• Rule only detects sudden loss of traffic. It won't alert if there is no traffic to begin with. You need at least 1 event for rule to alert. • List of device ip address or device hostnames as input. Only these systems will be tracked. • Time input is required. Alert is fired when time interval between events exceeds input time.

Example #6:

Multiple Failed Logins NOT followed by a Lockout event by the same user.

EPL #10: using timer and Lockout

Rule Name	FailedloginswoLockout
Rule Description	There are multiple failed logins that are not followed by Lockout event by the same user.

Rule Code	SELECT * FROM pattern [every-distinct(a.user_dst, a.device_ip, 1 msec) (a= Event(ec_activity='Logon' and ec_outcome='Failure' and user_dst IS NOT NULL) -> [2](Event (device_ip =a.device_ip and ec_activity='Logon' and ec_outcome='Failure' and user_dst=a.user_dst) AND NOT Event((ec_activity='Logon' and ec_outcome='Success' and device_ip = a.device_ip and user_dst=a.user_dst) or (ec_activity='Lockout' and device_ip = a.device_ip and user_dst=a.user_dst)))) where timer:within(60 seconds) -> (timer:interval(30 seconds) and not Event(device_ip=a.device_ip and user_dst=a.user_dst and ec_activity='Lockout'))];
Note	• Above query detects the absence of a Lockout Event after the occurrence of 2 failed logins from same user. • The occurrence of the multiple failed logins are timed and are assumed to occur within a certain period of time. Also, in-practice the Lockout event is assumed to occur within a short time after the occurrence of the last failed login event because the threshold value of Failed logins per user is set in a given domain. • In current query, every distinct will suppress new thread for combination of user and device for 1 millisec. • Time allowed for 3 failed logins is 60 secs since 1st failed attempt. Wait period for lockout event to occur is 30 secs

Example #7:

Custom functions to perform LIKE and REGEX operations for ARRAY elements.

EPL #11: @RSAAlert(oneInSeconds=0)

Rule Name	MatchLikeRegex
Rule Description	There are custom functions to perform LIKE and REGEX comparisons of array meta keys.
Rule Code	SELECT * FROM pattern[e1=Event(matchLike(alias_host, "10.0.0.%")) AND e2=Event(matchRegex(alias_host, "10\.0\.0\.1[0-9][0-9]")) where timer:within(5 Minutes)];

Note:
1. "**.**" in meta keys should be replaced with ("_").
2. All patterns should be time bound.
3. Use appropriate tags in front of statements, for example:

```
@RSAPersist:
@RSAAlert:
```

For additional details you can refer to:

- EPL Documentation: http://www.espertech.com/esper/esper-documentation/
- EPL Online Tool: http://esper-epl-tryout.appspot.com/epltryout/mainform.html

Configure an In-Memory Table Using an EPL Query

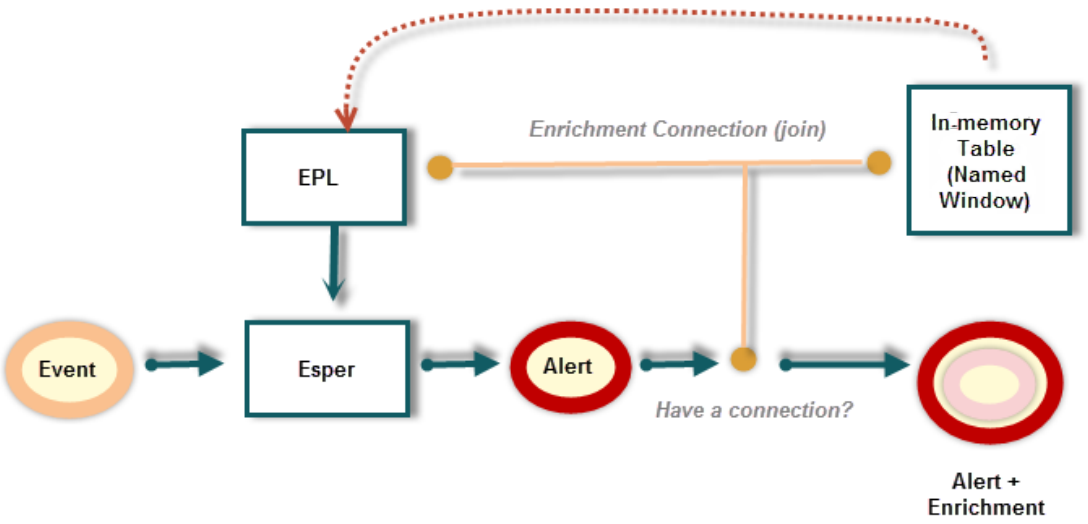
Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see Configure a Context Hub List as an Enrichment Source.

When you use an In-Memory Table configuration in expert mode, you can create an enrichment source or named window based on an Esper query. This allows you to have more control over the content and create more dynamic content. When you do this, an EPL query constructs the named window to capture interesting states from the event stream.

Workflow

The following shows the workflow for creating a query using a named window:

1. The event is sent to the Esper Engine.
2. An EPL query is generated.
3. An alert is triggered.
4. The query checks to see if there is a connection between the event and the Named Window.
5. If there is a connection, the query that populates the Named Window is run and populated.
6. The content from the Named Window is added to the alert content and sent or displayed (depending on your settings).



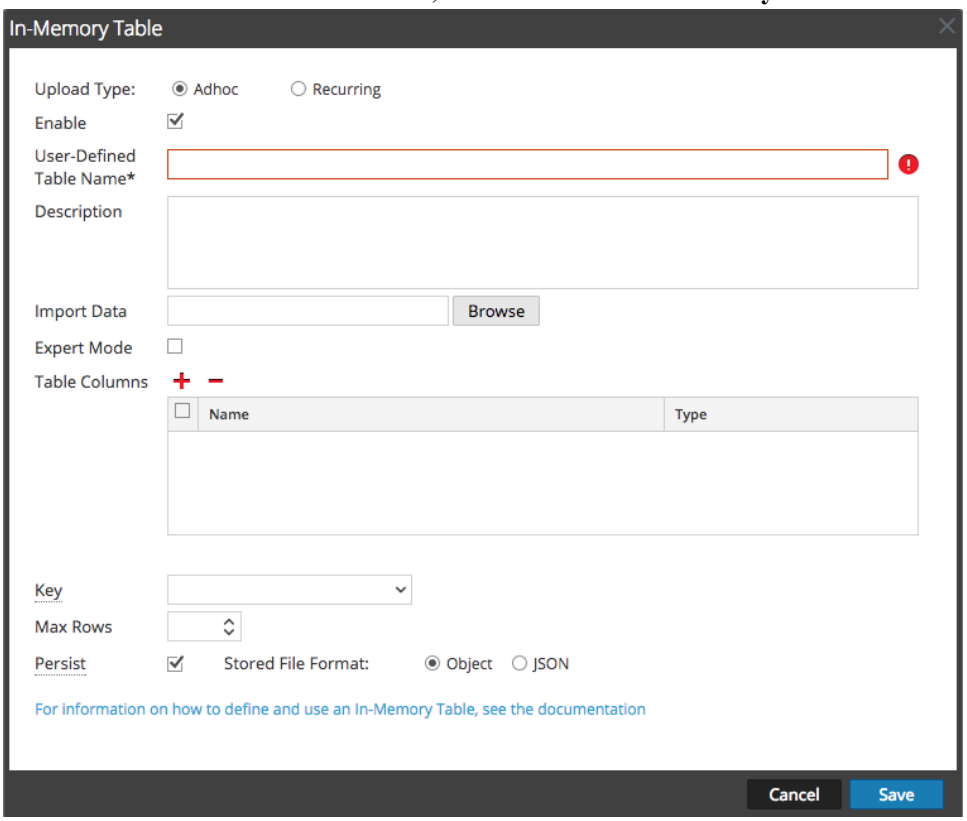
Prerequisites

- The meta used in the EPL statement must exist in the data.
- You must create well-formed EPL statements.

Procedure

Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules.

1. Go to **(Configure) > ESA Rules**.
   The Configure view is displayed with the Rules tab open.
2. Click the **Settings** tab.
3. In the options panel, select **Enrichment Sources**.
4. In the **Enrichment Sources** section, click > **In-Memory Table**.



5. Select **Adhoc**.
   By default, Enable is selected. When you add the in-memory table to a rule, alerts will be enriched with data from it.
6. In the **User-Defined Table Name** field, type a descriptive name to describe the in-memory table.
7. If you want to explain what the enrichment adds to an alert, enter information in the **Description** field.
   This description displays when you view the list of enrichments from the Enrichment Sources view, so it's a good idea to enter a thorough description as a best practice. Doing this allows other users to understand the content of the enrichment without opening it to examine its contents.
8. Select **Expert Mode** to define an advanced in-memory table configuration by writing an EPL query. The Table Columns are replaced by a **Query** field.
9. Select **Persist** to preserve the in-memory table on disk when the ESA service stops and to re-populate the table when the service restarts.
10. Enter the EPL query in the **Query** field. The query should be well-formed, and it's a good idea to test it before entering it in the field.
11. Click **Save**.

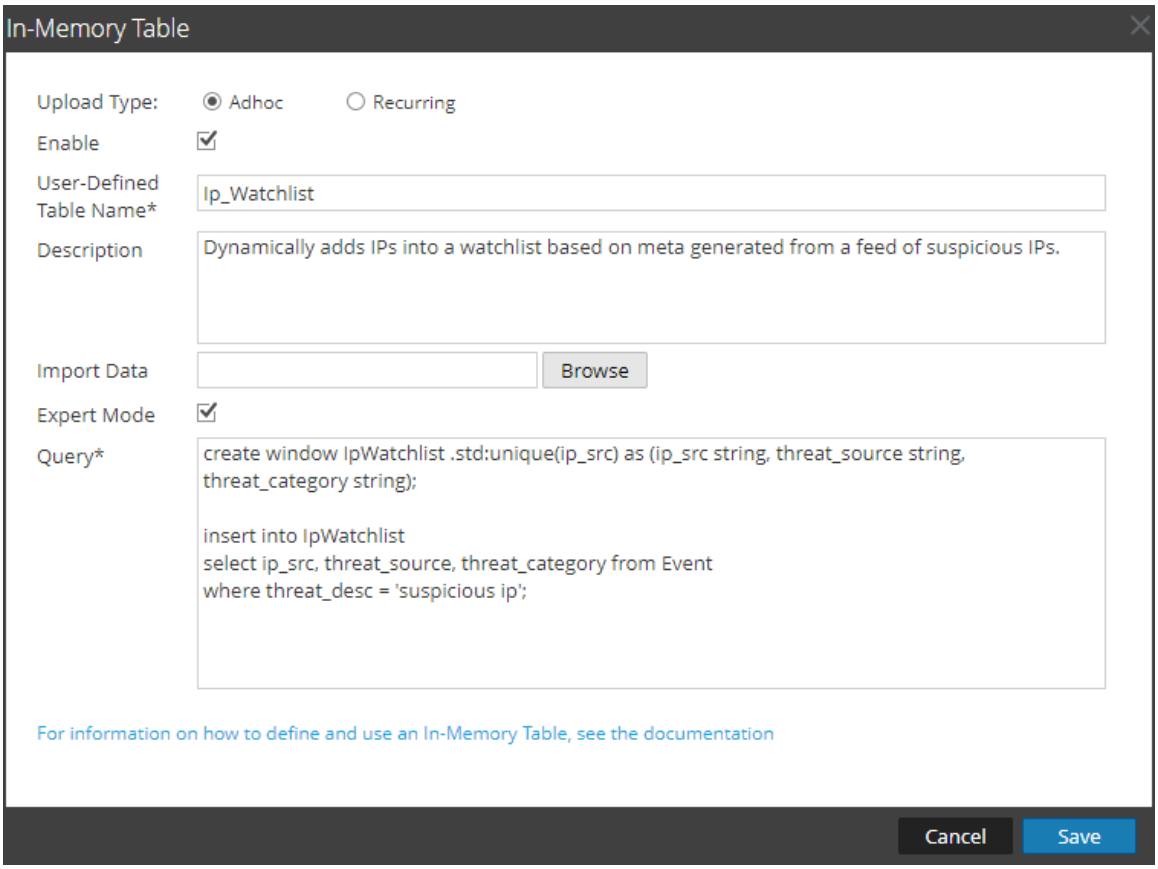
Example

For example, you want to know when an IPS or IDS is giving five or more inbound events with an event identified with malicious code. Additionally, you would like to know when the source IP of those events has been identified as suspicious by other sources. This information helps to more quickly triage the event and determine whether the alert is a true positive.

Step 1: Create the Enrichment

In this example, this enrichment is a watchlist of IPs that have been identified as suspicious by third party sources or by internal staff. The meta of `threat_desc equal to 'suspicious ip'` is generated when a match to a feed occurs. This meta can be matched and output based on a log, packet, or endpoint event.

The enrichment should look like the following:



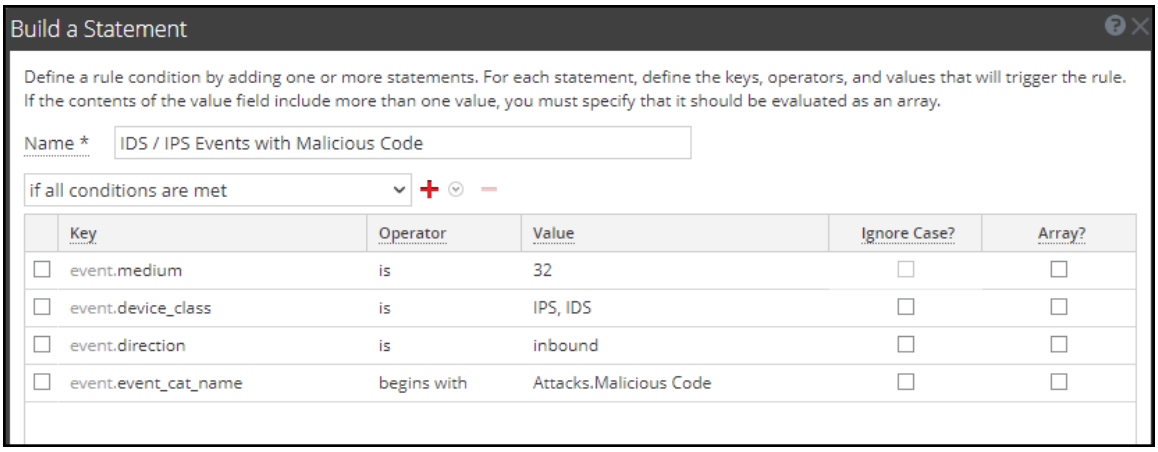
Parameters	Description
Upload Type	Adhoc
IP_Watchlist	IP_Watchlist
Description	Dynamically populated whitelist based on a feed of IPs that are considered suspicious.
Expert Mode	Selected

Query	`create window IpWatchlist .std:unique(ip_src) as (ip_src string, threat_source string, threat_category string);`
	`insert into IpWatchlist`
	`select ip_src, threat_source, threat_category from Event`
	`where threat_desc = 'suspicious ip';`

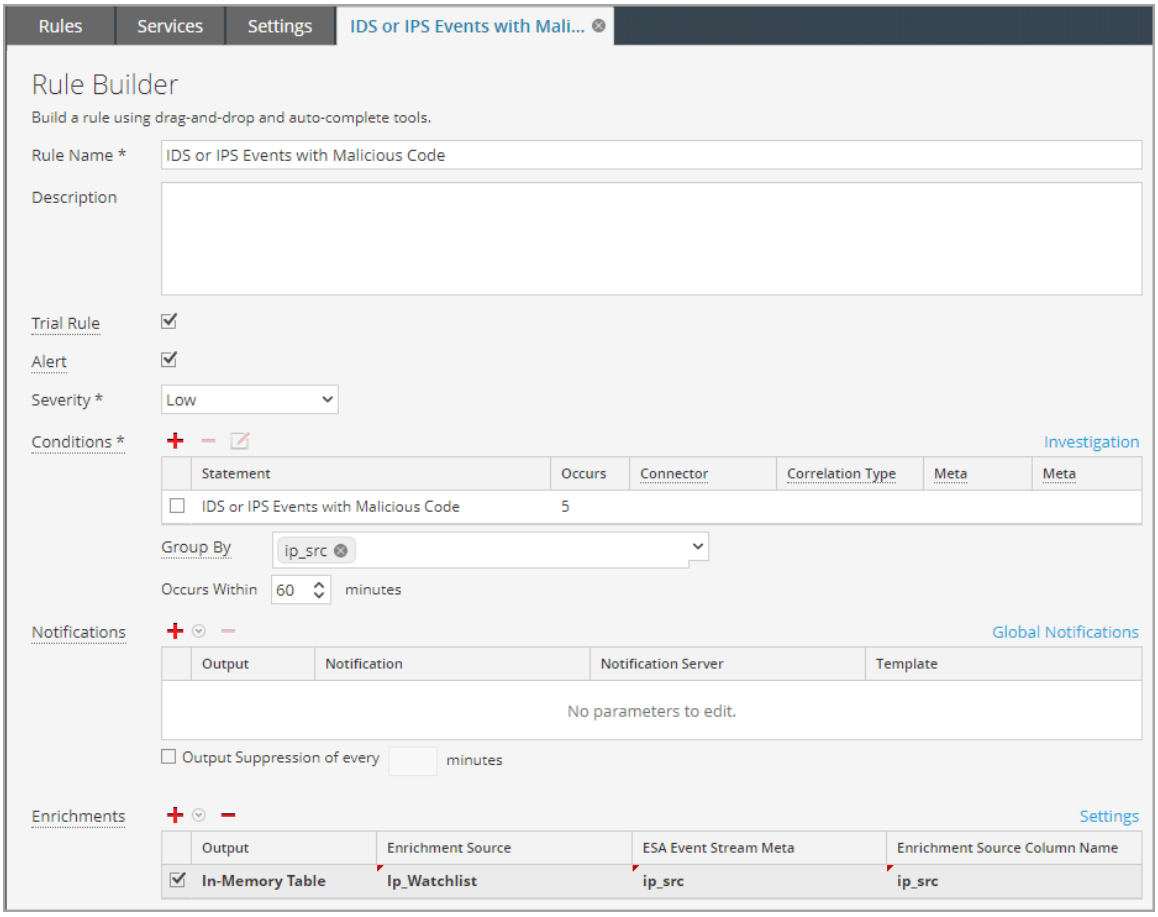
Step 2: Create Your Rule

First, you need to create your ESA Correlation rule. This example rule looks for inbound IPS or IDS log events with the `event_cat_name` beginning with `Attacks.Malicious Code`. If five or more events for the same `ip_src` occur within 60 minutes, then an alert will be triggered. If an `ip_src` from the Enrichment equals the `ip_src` from the alert, then that alert will be enriched with additional meta. In this case, the analyst would see the values for `threat_source` and `threat_category` in the raw alert. `Threat_category` would indicate the type of malware and `threat_source` would indicate the entity that has reported the ip as suspicious. The analyst could use this information to do additional research or escalate to the next tier for creation of a possible incident.

Rule Statement



Rule Logic with Enrichment Added



ESA Alert References

In Event Stream Analysis (ESA), you configure and deploy ESA rules to get alerted about potential network threats.

These topics explain the user interface for ESA Correlation rules.

- Rules Tab
- Rules Tab Options Panel
- Rule Library Panel
- Rule Builder Tab
- Build a Statement Dialog
- Advanced EPL Rule Tab
- Rule Syntax Dialog
- Deployment Panel
- Deploy ESA Services Dialog
- Deploy ESA Rules Dialog
- Updates to the Deployment Dialog
- Services Tab
- Settings Tab

Rules Tab

The Rules tab enables you to configure ESA rules and deployments.

What do you want to do?

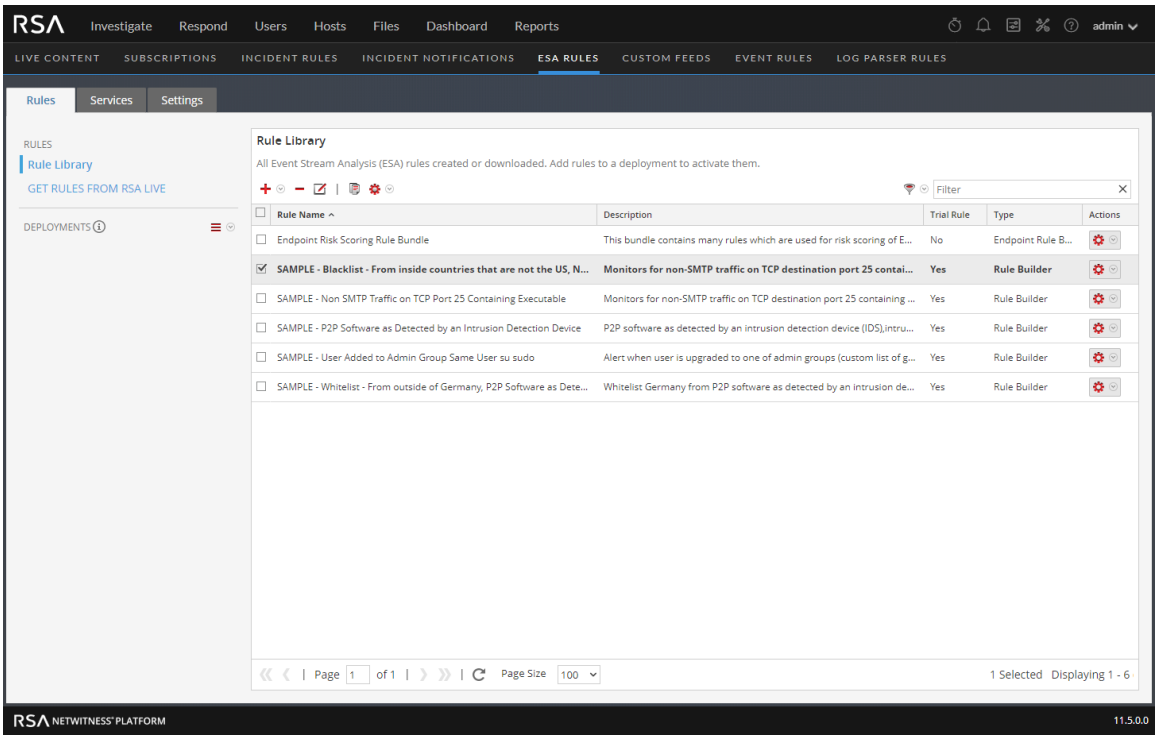
Role	I want to ...	Show me how
Content Expert	View types of rules.	ESA Rule Types
Content Expert	Deploy Trial Rules.	Working with Trial Rules
Content Expert	Create a rule.	Add Rules to the Rule Library
Content Expert	Deploy a rule.	Deploy Rules to Run on ESA

Related Topics

- Getting Started with ESA

Quick Look

The Rules tab is displayed when you go to **(Configure) > ESA Rules**.
The following figure shows the Rules tab.



The Rules tab is divided into three sections:

- Rules Tab Options Panel
- Rule Library Panel
- Deployment Panel

Rules Tab Options Panel

In the **Rules** tab options panel to the left, you can view ESA rules in the Rule Library and create ESA rule deployments.

What do you want to do?

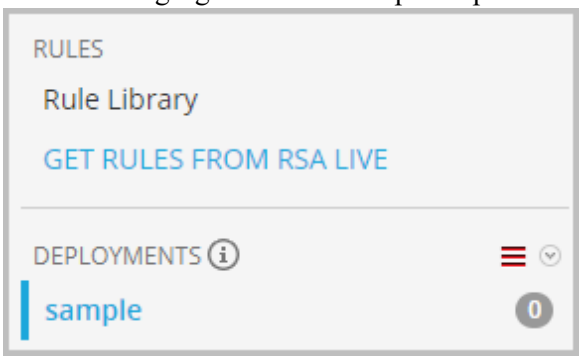
Role	I want to ...	Show me how
Content Expert	View an ESA rule.	Add Rules to the Rule Library
Content Expert	Create an ESA rule deployment.	ESA Rule Deployment Steps

Related Topics

- Working with Rules

Quick Look

The following figure shows the options panel in the **Rules** tab.



There are two sections in the options panel: Rules and Deployments.

Rules Section

The Rules section contains two options. **Rule Library** is selected by default, and when it's selected, the Rule Library view is displayed within the tab. **Get Rules From RSA Live** navigates to the Live Search view, where you can search for rules.

Deployments Section

The Deployments section lists ESA rule deployments and indicates whether there are updates to the deployments. From this section, deployments can be added, deleted, edited, and refreshed. Selecting a deployment from the list displays the Deployment panel within the tab. The following table describes the features of this section.

Feature	Description
	Displays a drop-down menu from which you can choose to add, edit, or delete an ESA rule deployment. You can also refresh the list of deployments to see if there are any new updates to the list.
	Indicates whether there are any updates to the deployment.
	Indicates the number of rules in the deployment.

Rule Library Panel

The Rule Library panel allows you to manage rules.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Add an ESA rule.	Add a Rule Builder Rule
Content Expert	Edit, duplicate, or delete an ESA rule.	Edit, Duplicate or Delete a Rule
Content Expert	Import or export ESA rules.	Import or Export Rules
Content Expert	Filter the ESA rules list.	Filter or Search for Rules

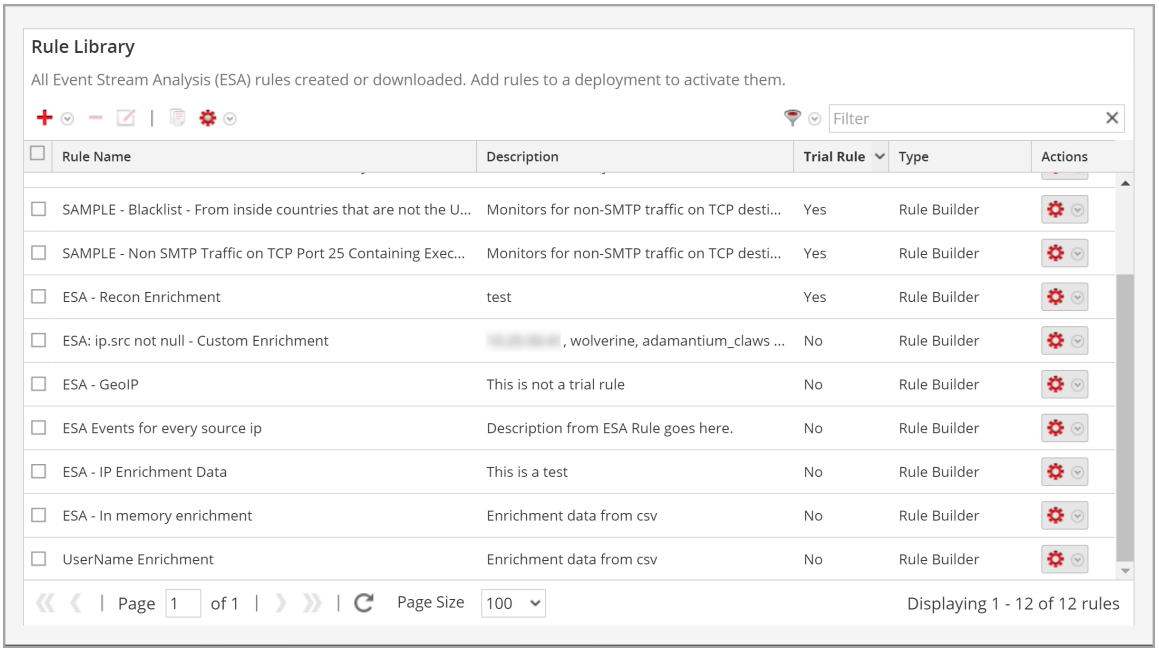
Related Topics

- Add an Advanced EPL Rule

Quick Look

To access this view, go to **(Configure) > ESA Rules**. The Rules tab is displayed and the Rule Library panel is on the right.

The following figure shows the Rule Library panel.



The Rule Library panel includes the following components:

- Rule Library toolbar
- Rule Library list

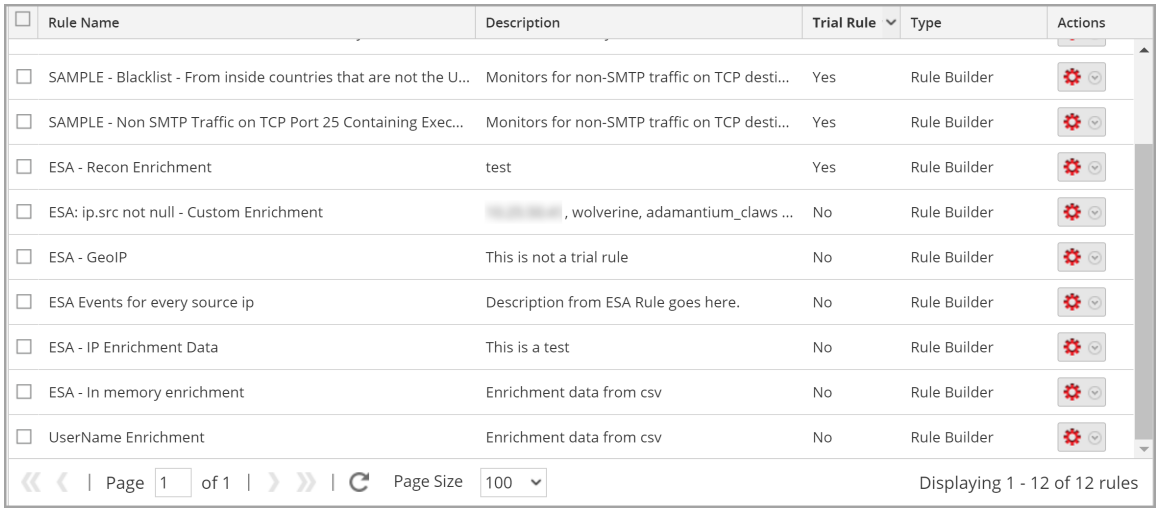
Rule Library Toolbar

The Rule Library toolbar allows you to add, delete, edit, duplicate, filter, export, and import ESA rules. The following figure shows the icons for these actions.



Rule Library List

The following figure shows the Rule Library list.

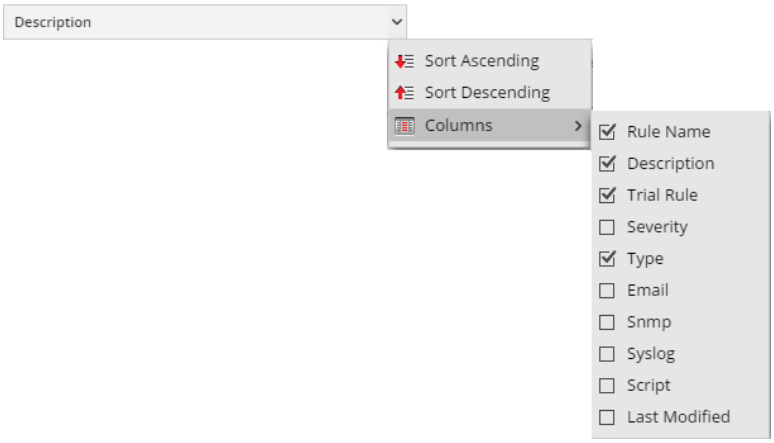


The Rule Library list shows all of the ESA rules. The following table lists the columns in the Rule Library list and their description.

Column	Description
Rule Name	Purpose of the ESA rule.
Description	Summary of what the ESA rule detects.
Trial Rule	Deployment mode to see if the rule runs efficiently.
Type	The type of rule. For more information, see ESA Rule Types.
Actions ( )	Menu to delete, edit, duplicate, or export the selected rule.
Severity	Threat level of alert triggered by the rule.
Email	Indicates whether an alert notification for the rule is sent by email. This column is not visible by default.

Column	Description
SNMP	Indicates whether an alert notification for the rule is sent using SNMP. This column is not visible by default. (ESA SNMP notifications are not supported in NetWitness version 11.3 and later.)
Syslog	Indicates whether an alert notification for the rule is sent using Syslog. This column is not visible by default.
Script	Indicates whether an alert notification for the rule executes a script. This column is not visible by default.
Last Modified	The date and time when the ESA rule was last modified. This column is not visible by default.

To display columns which aren't visible by default, hover over the title of a column and click the **v** on the right. This opens a drop-down menu in which you can sort the contents of the column or choose which columns you want to see in the Rule Library list.



Rule Builder Tab

The Rule Builder tab enables you to define a Rule Builder rule.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Define a Rule Builder rule.	Add a Rule Builder RuleStep 1. Name and Describe the Rule
Content Expert	Define rule criteria.	Step 2. Build a Rule Statement
Content Expert	Add conditions to the rule.	Step 3. Add Conditions to a Rule Statement
Content Expert	Test the ESA rule logic.	Validate an ESA Rule

Related Topics

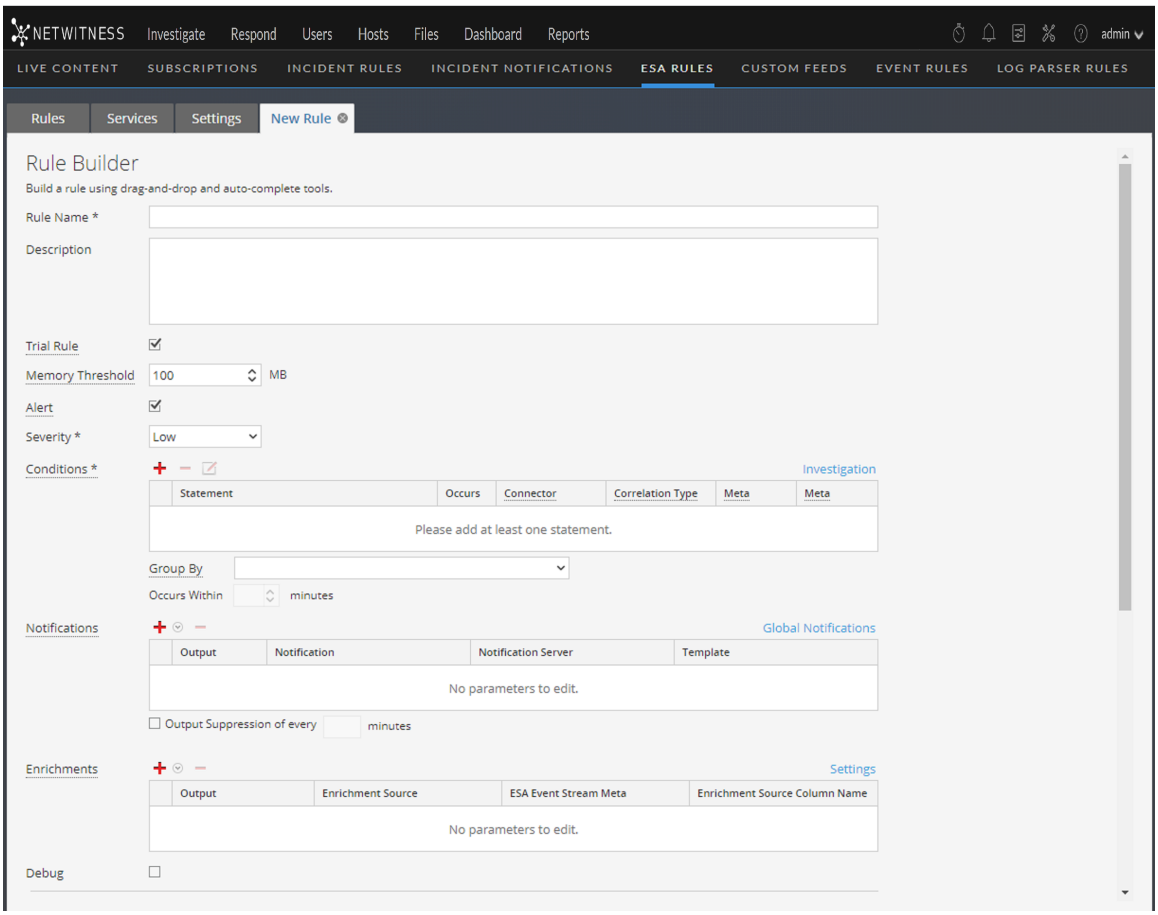
- Add an Advanced EPL Rule

Quick Look

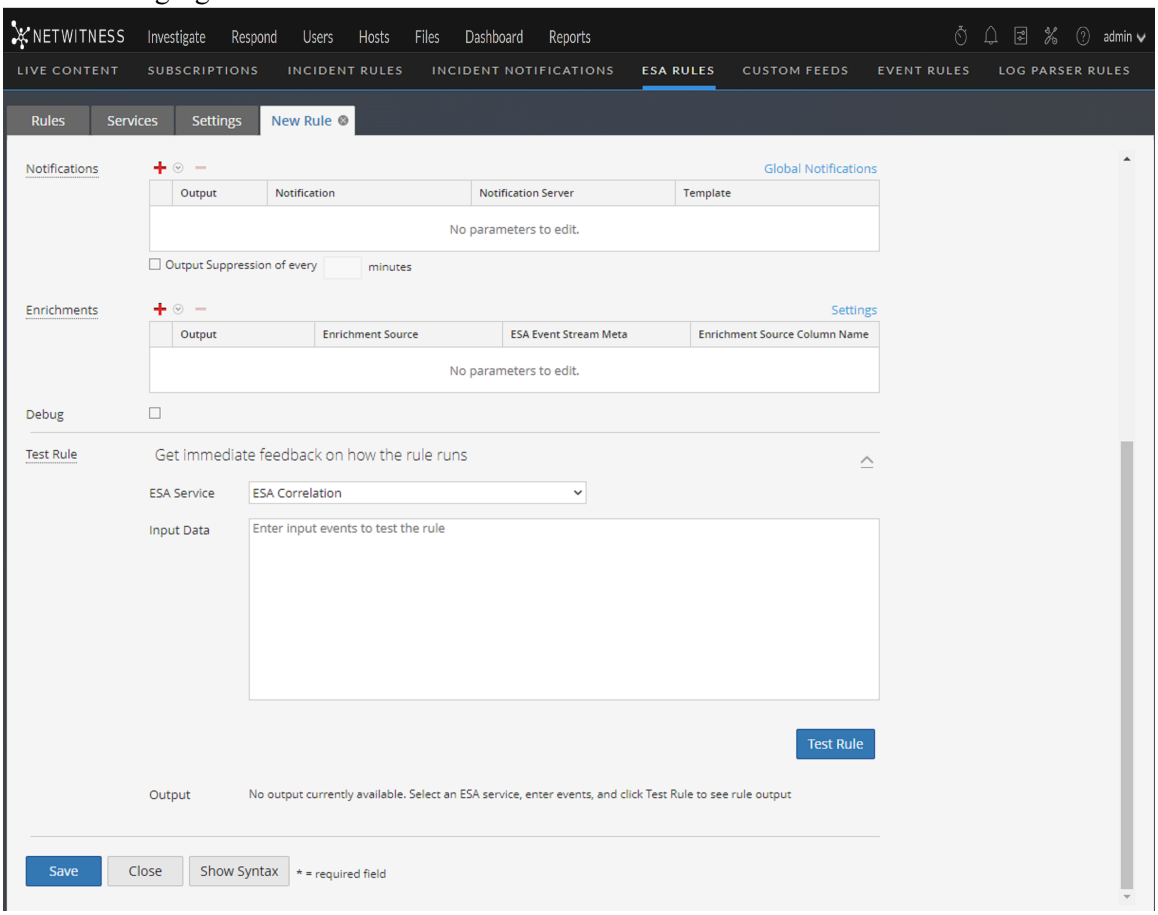
To access the Rule Builder tab:

1. Go to  **(Configure) > ESA Rules**.
   The Rules tab opens by default.
2. In the **Rule Library** toolbar, select  **> Rule Builder**.
   The Rule Builder tab is displayed.

The following figure shows the Rule Builder tab.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
LIVE CONTENT
SUBSCRIPTIONS
INCIDENT RULES
INCIDENT NOTIFICATIONS
ESA RULES
CUSTOM FEEDS
EVENT RULES
LOG PARSER RULES
Rules
Services
Settings
New Rule
Rule Builder
Build a rule using drag-and-drop and auto-complete tools.
Rule Name *
Description
Trial Rule
Memory Threshold
100
MB
Alert
Severity *
Low
Conditions *
Investigation
Statement
Occurs
Connector
Correlation Type
Meta
Meta
Please add at least one statement.
Group By
Occurs Within
minutes
Notifications
Global Notifications
Output
Notification
Notification Server
Template
No parameters to edit.
Output Suppression of every
minutes
Enrichments
Settings
Output
Enrichment Source
ESA Event Stream Meta
Enrichment Source Column Name
No parameters to edit.
Debug

The following figure shows the Rule Builder tab scrolled down with the Test Rule section in view.



The following table lists the parameters in the Rule Builder tab.

Field	Description
Rule Name	Purpose of the ESA rule.
Description	Summary of what the ESA rule detects.
Trial Rule	Deployment mode to see if the rule runs efficiently.
Memory Threshold	(This option applies to version 11.5 and later.) The maximum memory usage allowed for this rule in MB. Add Memory Thresholds to ESA rules that use memory. For example, if a rule contains windows or pattern matching, configure a memory threshold for that rule. If the configured memory threshold is exceeded, it gets disabled individually and an error is displayed for that rule on the (Configure) > ESA Rules > Services tab. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.
Alert	(This option applies to version 11.3 and later.) When selected, the alert is sent to Respond. If the checkbox is cleared, an alert will not be sent to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.

Field	Description
Severity	Threat level of alert triggered by the rule.

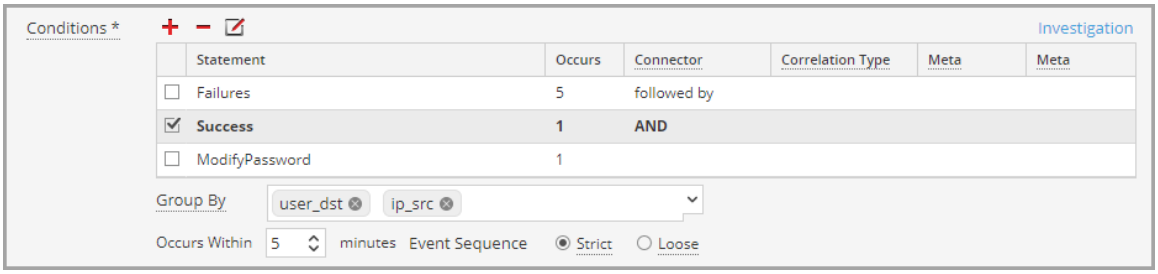
The Rule Builder includes the following components:

- Conditions Section
- Notifications Section
- Enrichments Section
- Debug Option
- Test Rule Section

Conditions Section

In the Conditions section of the Rule Builder tab, you define what the rule detects.

The following figure shows the Conditions section.



The following table lists the parameters of the Conditions section.

Parameter	Description
+	Add a statement.
−	Remove selected statement.
✎	Edit selected statement.
Statement	Logical group of conditions for one operation.
Occurs	Alert frequency if the condition is met. This specifies that there must be at least that many events that satisfy the criteria in order to trigger an alert. The time window in minutes binds the Occurs count.

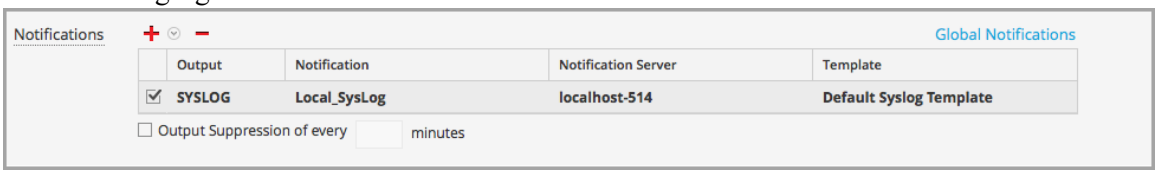
Parameter	Description
Connector	Options to specify relationship among the statements: • followed by • not followed by • AND • OR The Connector joins two statements with AND, OR, followed by, or not followed by. When followed by is used, it specifies that there is a sequencing of those events. AND and OR build one large criteria. The followed by creates distinct criteria that occurs in sequence.
Correlation Type	**Correlation Type** applies only to **followed by** and **not followed by**. If you choose a correlation type of SAME, select one meta to correlate on, and if you choose a correlation type of JOIN, select two meta to correlate on. You may want to use JOIN if you are trying to correlate on meta from two different data sources. For example, say you want to correlate an AV alert with an IDS alert.
Meta	Enter the meta condition if choosing a correlation type of SAME or JOIN (as described above).
Meta	Enter the second meta condition if choosing a correlation type of JOIN (as described above). For example, The destination IP address from the AV alert and source IP address for the workstation from the IDS alert are joined to allow you to view the same entities across different sources.
occurs within minutes	Time window within which the conditions must occur.
Event Sequence	Choose whether the pattern must follow a *strict* match or a *loose* match. If you specify a strict match, this means that the pattern must occur in the *exact* sequence you specified with no additional events occurring in between. For example, if the sequence specifies five failed logins (F) followed by a successful login (S), this pattern will only match if the user executes the following sequence: F,F,F,F,F,S. If you specify a loose match, this means that other events may occur within the sequence, but the rule will still trigger if all of the specified events also occur. For example, five failed login attempts (F), followed by any number of intervening successful login attempts (S), followed by a successful login attempt might create the following pattern: F,S,F,S,F,S,F,S,F,S which would trigger the rule despite the intervening successful logins.
Group By	Select the meta key by which to group results from the dropdown list. For example, suppose that there are three users; Joe, Jane, and John and you use the Group By meta, **user_dst** (user_dst is the meta field for the user destination account). The result will show events grouped under the user destination accounts, Joe, Jane, and John. You can also group by multiple keys. For example, you might want to group by user and machine to see if a user logged in from the same machine attempts to log into an account multiple times. To do this, you might group by user_dst and ip_src.

Notifications Section

In the Notifications section, you can choose how to be notified when ESA generates an alert for the rule.

For more information on the alert notifications, see Add Notification Method to a Rule.

The following figure shows the Notifications section.



Parameter	Description
	To add an alert notification type.
	To delete the selected alert notification.
Output	Alert notification type. Options are: • Email • SNMP (This option is not supported in NetWitness version 11.3 and later.) • Syslog • Script
Notification	Name of previously configured output, such as an email distribution list.
Notification Server	Name of server that sends the output.
Template	Name of template for the alert notification.
Output Suppression of every	Option to specify alert frequency.
Minutes	Alert frequency in minutes.

Enrichments Section

In the Enrichments section, you can add a data enrichment source to a rule.

For more information on the enrichments, see Add an Enrichment to a Rule.

The following figure shows the Enrichments section.



Parameter	Description
	To add an enrichment.
	To delete the selected enrichment.
Output	Enrichment source type. Options are: • In-Memory Table (Ad hoc only - Recurring In-Memory Tables are no longer supported in version 11.3 and later.) • GeoIP
Enrichment Source	Name of previously configured enrichment source, such as a .CSV filename for an In-Memory Table.
ESA Event Stream Meta	ESA meta key whose value will be used as one operand of join condition.
Enrichment Source Column Name	Enrichment source column name whose value will be used as the other operand of the join condition. For an in-memory table, If you configured a key when creating a .CSV-based enrichment, this column automatically populates with the selected key. However, you can change it if you like. For a GeoIP enrichment source, ipv4 is automatically selected.

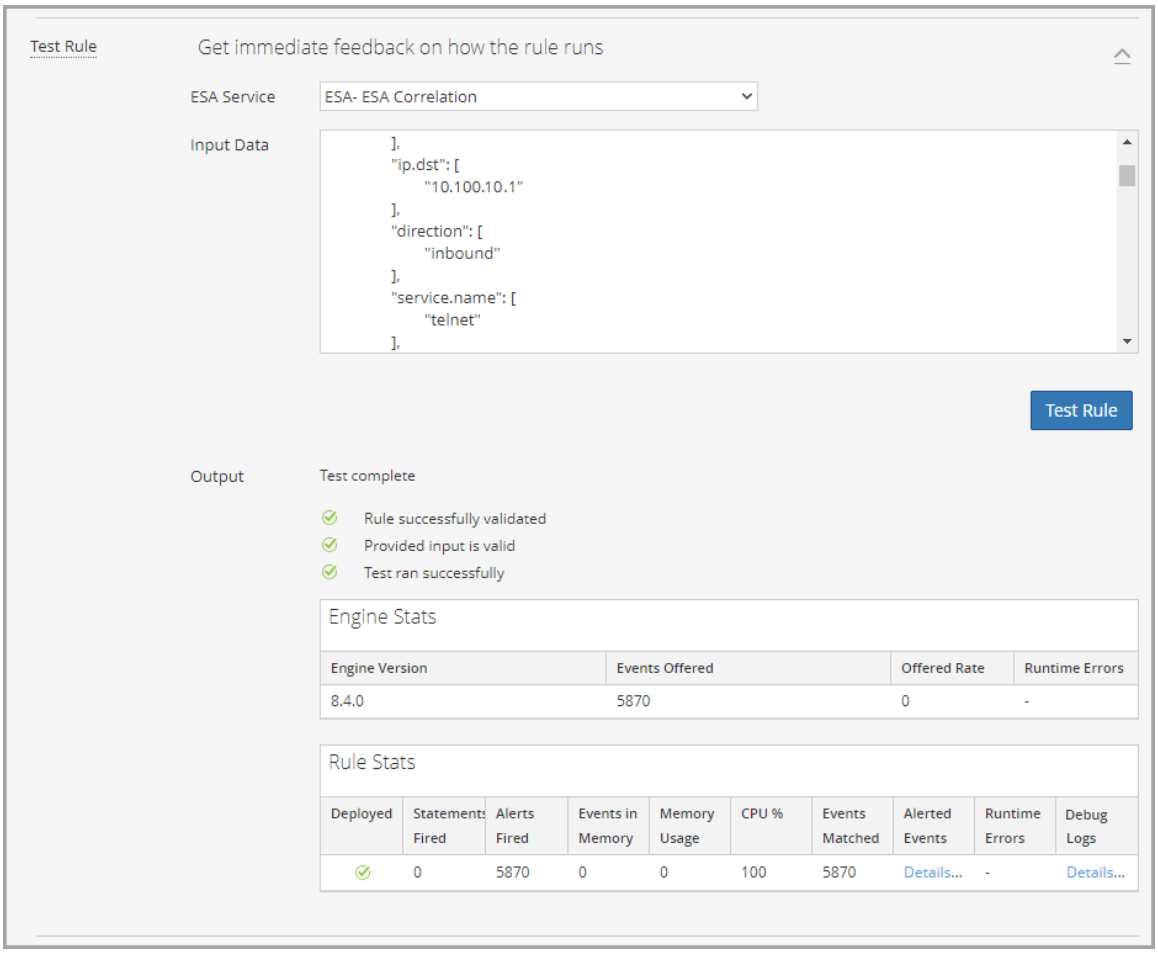
Debug Option

Select the Debug option to print alerts to the ESA logs for troubleshooting. This adds an @Audit ('stream') annotation to the rule. This is useful when debugging the Esper rules.

Test Rule Section

Note: The Test Rule section is available in NetWitness Platform 11.5 and later.

In the Test Rule section, you can validate your ESA rule to determine if the rule logic is working as expected before deploying the rule.



Field	Description
ESA Service	Select the ESA Correlation service to process the rule.
Input Data	Enter the input events to test the rule. Download the events from the Investigate view in JSON format, copy the events, and paste them in this field.
Output Data	After you select an ESA Correlation service, input data, and click the Test Rule button, you can view the output of the rule here and verify that the rule is working according to your requirements. You can view the alerts in the output, but this test does not send any alert notifications.

The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service

Field	Description
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Syntax

Click **Show Syntax** to view the EPL syntax of conditions, statements, and debugging parameters. It also provides a warning when the syntax is invalid. For more information, see Rule Syntax Dialog.

Build a Statement Dialog

The Build a Statement dialog allows you to construct a condition statement when creating a new Rule Builder rule.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Configure a rule statement.	Step 2. Build a Rule Statement
Content Expert	Add conditions to the rule.	Step 3. Add Conditions to a Rule Statement

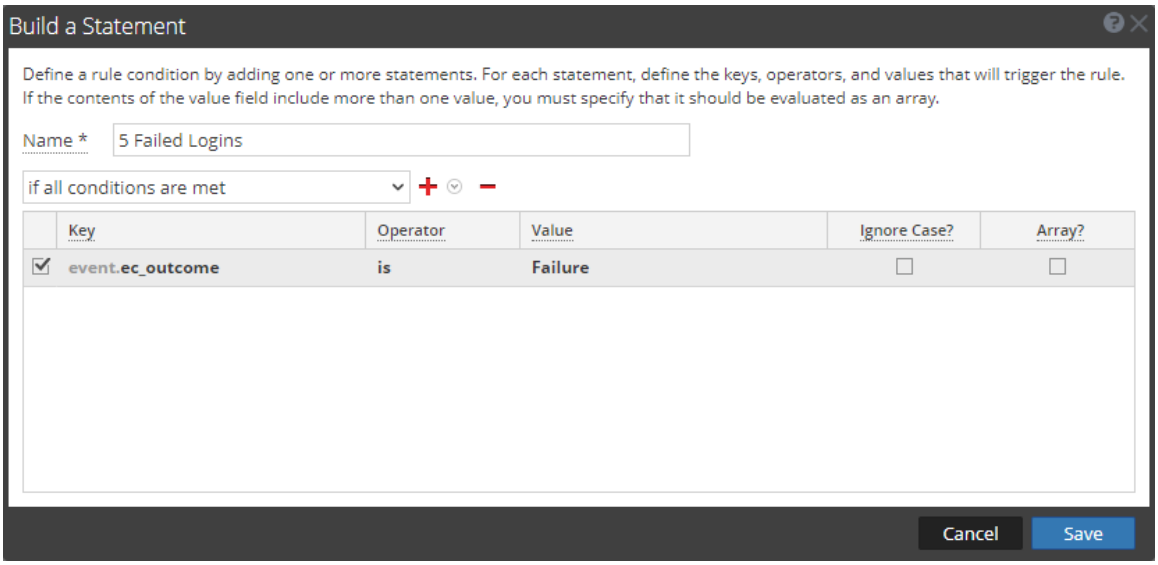
Related Topics

- Add a Rule Builder Rule

Quick Look

To access the Build a Statement dialog:

1. Go to  **(Configure) > ESA Rules**.
   The Configure ESA Rules view is displayed with the Rules tab open.
2. In the **Rule Library** toolbar, select  **> Rule Builder**.
   A New Rule tab is displayed..
3. In the **Conditions** section, click .
   The Build a Statement dialog is displayed.



The following table describes the parameters in the Build a Statement dialog.

Parameter	Description
Name	Purpose of the statement.
Select	Conditions the rule requires. There are two options: • If all conditions are met • If any of these conditions are met
Key	Key for ESA to check in the rule statement.

Parameter	Description
Operator	Relationship between the meta key and value for the key: • is • is not • is not null • is greater than (>) • is greater than or equal to (>=) • is less than (<) • is less than or equal to (<=) • is one of (For array type meta) • is not one of (For array type meta) • contains • not contains • begins with • ends with
Value	Value for ESA to look for in the key.
Ignore Case?	This field is designed for use with string and array of string values. By choosing the **Ignore Case** field, the query will treat all string text as a lowercase value. This ensures that a rule that searches for the user named Johnson would trigger if the event contains "johnson," "JOHNSON," or "JoHnSoN."
Array?	Choice to indicate if contents of Value field represent one value or multiple values: • Select the box to indicate multiple values. • Clear the box to indicate one value.
	Add a statement. You can add a meta condition, whitelist condition, or blacklist condition.
	Delete selected statement.
Save	Add statement to the Conditions section of the Rule Builder tab.

The following table shows the operators you can use in the Rule Builder:

Operator	Required Value	Usage	Example	Meaning
is	Singular string value	The meta key is equal to the *value* field.	***user_dst*** is John Doe.	***user_dst*** is equal to the string *"John Doe"*.

Operator	Required Value	Usage	Example	Meaning
is	Array string value	The meta key is equal to one of the elements of the *value* field.	***user_dst*** is John, Doe, Smith.	***user_dst*** is equal either to the string *"John"* or to the string *"Doe"* or to the string *"Smith"* (Note, the spaces are stripped.).
is not	Singular string value	The meta key is not equal to the *value* field.	***size*** is not 200.	***size*** is not equal to the number 200 (size is a numeric value).
is not	Array string value	The meta key is not equal to any of the elements of the *value* field.	***size*** is not 200, 300, 400.	***size*** is equal neither to *200* nor to *300* nor to *400*.
is not null	N/A (looks for any value)	The meta key value is not null.	***user_dst*** is not null.	***user_dst*** is a meta that contains a value.
is greater than (>)	Number	The numeric value of the meta key is greater than the number in the *value* field.	***payload*** is greater than 7000.	***payload*** is a numeric value that is greater than 7000.
is greater than or equal to (>=)	Number	The numeric value of the meta key is greater than or equal to the number in the *value* field.	***payload*** is greater than or equal to 7000.	***payload*** is a numeric value that is greater than or equal to 7000.
is less than (<)	Number	The numeric value of the meta key is less than the number in the *value* field.	***ip_dstport*** is less than 1024.	***ip_dstport*** is a numeric value that is less than the numeric value 1024.
is less than or equal to (<=)	Number	The numeric value of the meta key is less than or equal to the number in the *value* field.	***ip_dstport*** is less than or equal to 1024.	***ip_dstport*** is a numeric value that is less than or equal to numeric value 1024.
is one of	Array string value	The meta key is one of the array string values in the *value* field.	***alias_host*** is one of Facebook, UTube, Instagram.	***alias_host*** is one of the array string values *Facebook, UTube, Instagram*.
is not one of	Array string value	The meta key is not one of the array string values in the *value* field.	***alias_host*** is not one of Facebook, UTube, Instagram.	***alias_host*** is not one of the array string values *Facebook, UTube, Instagram*.

Operator	Required Value	Usage	Example	Meaning
contains	String	The *value* field is a substring of the meta key. (This operator is only available for a string-valued meta key).	***ec_outcome*** contains failure.	***ec_outcome*** is a string that contains the substring *"failure"*.
not contains	String	The *value* field is not a substring of the meta key (This operator is only available for a string-valued meta key).	***ec_outcome*** not contains failure.	***ec_outcome*** is a string that does not contain the substring*"failure"*.
begins with	String	The *value* field is the beginning of the meta key (This operator is only available for a string-valued meta key).	***ip_dst*** begins with 127.0.	***ip_dst*** is a string that starts with *"127.0"*.
ends with	String	The *value* field is the end of the meta key (This operator is only available for a string-valued meta key).	***user_dst*** ends with son.	***user_dst*** is a string that ends in*"son"*.

Note: Terms in ***bold italics*** are Meta that may not exist in all customer environments.

Advanced EPL Rule Tab

The Advanced EPL Rule tab enables you to define rule criteria with an Event Processing Language (EPL) query.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Define an Advanced EPL rule.	Add an Advanced EPL Rule
Content Expert	Test the Advanced EPL rule logic.	Validate an Advanced EPL Rule
Content Expert	See examples of an Advanced EPL Rule.	Example Advanced EPL Rules
Content Expert	See best practices for writing Advanced EPL Rules.	ESA Rule Writing Best Practices

Related Topics

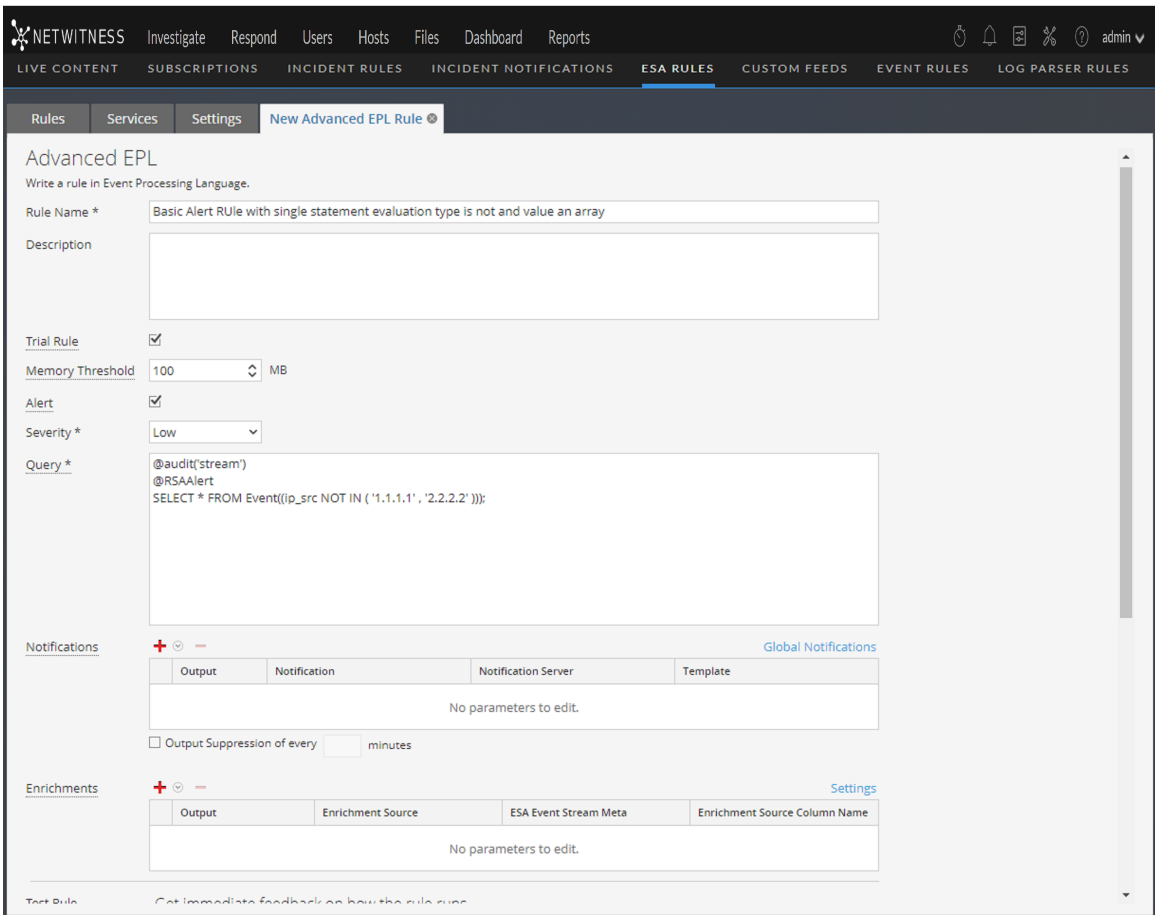
- Add a Rule Builder Rule
- Enrichment Sources

Quick Look

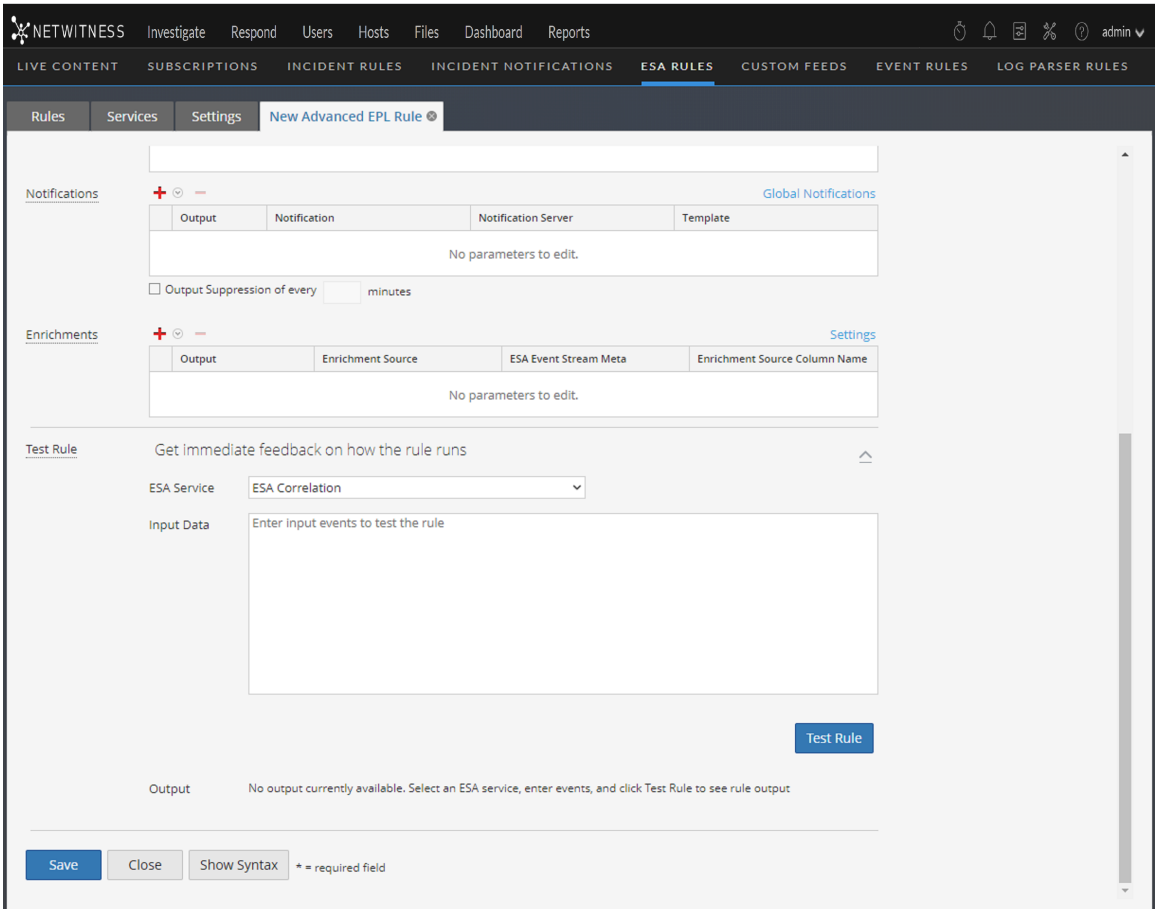
To access the Advanced EPL Rule tab:

1. Go to  **(Configure) > ESA Rules**.
   The Configure view is displayed with the Rules tab open by default.
2. In the **Rule Library** toolbar, select  **> Advanced EPL**.
   The Advanced EPL Rule tab is displayed.

The following figure shows the Advanced EPL Rule tab.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
LIVE CONTENT
SUBSCRIPTIONS
INCIDENT RULES
INCIDENT NOTIFICATIONS
ESA RULES
CUSTOM FEEDS
EVENT RULES
LOG PARSER RULES
Rules
Services
Settings
New Advanced EPL Rule
Advanced EPL
Write a rule in Event Processing Language.
Rule Name *
Basic Alert RUle with single statement evaluation type is not and value an array
Description
Trial Rule
Memory Threshold
100
MB
Alert
Severity *
Low
Query *
@audit('stream')
@RSAAlert
SELECT * FROM Event((ip_src NOT IN ('1.1.1.1' , '2.2.2.2')));
Notifications
Global Notifications
Output
Notification
Notification Server
Template
No parameters to edit.
Output Suppression of every
minutes
Enrichments
Settings
Output
Enrichment Source
ESA Event Stream Meta
Enrichment Source Column Name
No parameters to edit.

The following figure shows the Advanced EPL Rule tab scrolled down with the Test Rule section in view.



The following table lists the parameters in the Advanced EPL Rule tab.

Parameters	Description
Rule Name	Purpose of the ESA rule.
Description	Summary of what the ESA rule detects.
Trial Rule	Deployment mode to see if the rule runs efficiently.
Memory Threshold	(This option applies to version 11.5 and later.) The maximum memory usage allowed for this rule in MB. Add Memory Thresholds to ESA rules that use memory. For example, if a rule contains windows or pattern matching, configure a memory threshold for that rule. If the configured memory threshold is exceeded, it gets disabled individually and an error is displayed for that rule on the (Configure) > ESA Rules > Services tab. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.

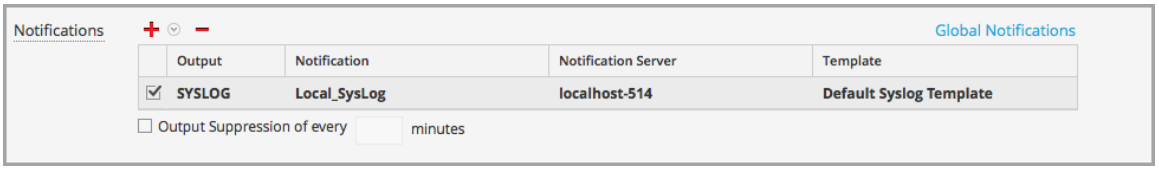
Parameters	Description
Alert	(This option applies to version 11.3 and Later.) When selected, the alert is sent to Respond. If the checkbox is cleared, an alert will not be sent to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.
Severity	Threat level of alert triggered by the rule.
Query	EPL query that defines rule criteria.

Notifications Section

In the Notifications section, you can choose how to be notified when ESA generates an alert for the rule.

For more information on the alert notifications, see Add Notification Method to a Rule.

The following figure shows the Notifications section.



Parameter	Description
	To add an alert notification type.
	To delete the selected alert notification type.
Output	Alert notification type. Options are: • Email • SNMP (This option is not supported in NetWitness version 11.3 and later.) • Syslog • Script
Notification	Name of previously configured output, such as an email distribution list.
Notification Server	Name of server that sends the output.
Template	Name of template for the alert notification.
Output Suppression of every	Option to specify alert frequency.
Minutes	Alert frequency in minutes.

Enrichments Section

In the Enrichments section, you can add a data enrichment source to a rule.

For more information on the enrichments, see Add an Enrichment to a Rule.
The following figure shows the Enrichments section.

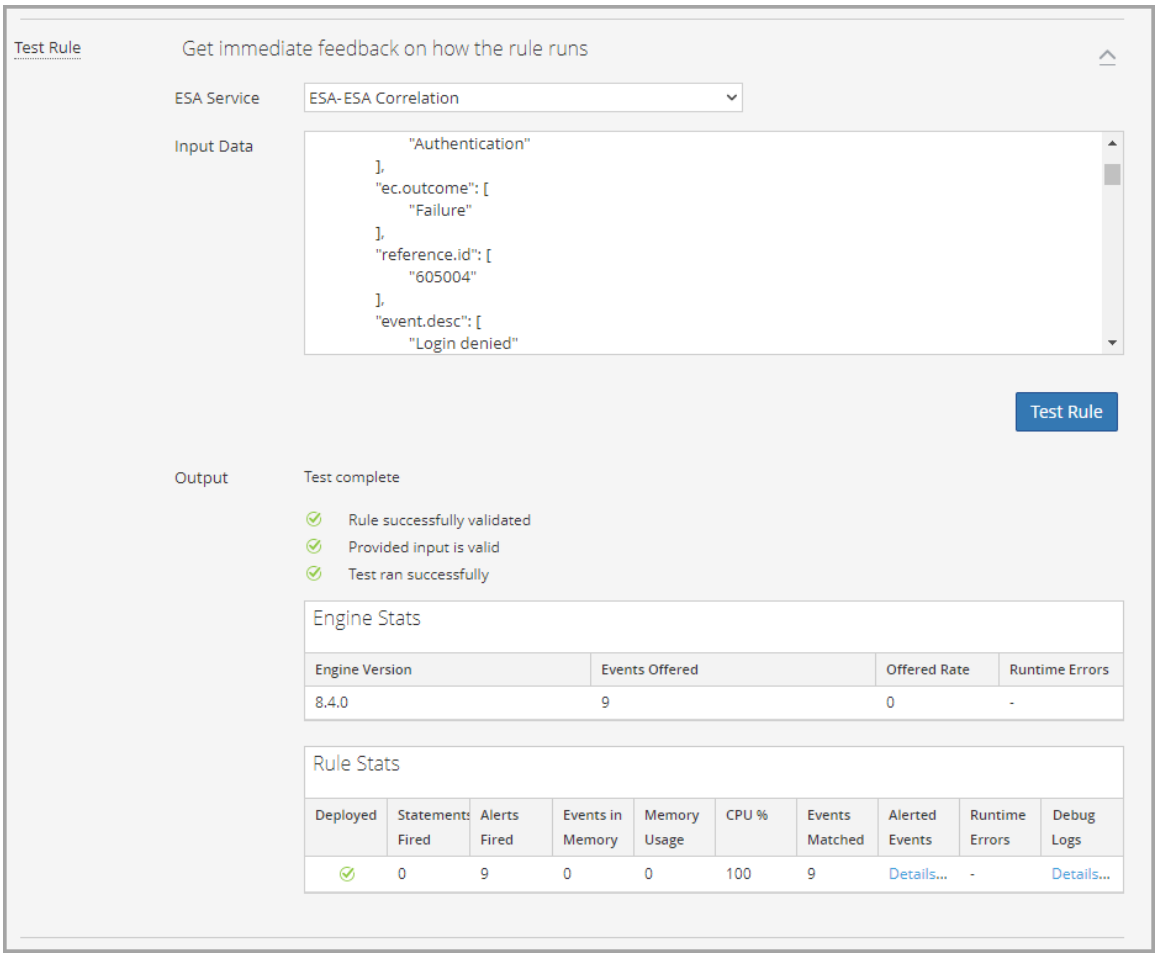


Parameter	Description
	To add an enrichment.
	To delete the selected enrichment.
Output	Enrichment source type. Options are: • In-Memory Table (Ad hoc only - Recurring In-Memory Tables are no longer supported in version 11.3 and later.) • GeoIP
Enrichment Source	Name of previously configured enrichment source, such as a .CSV filename for an In-Memory Table.
ESA Event Stream Meta	ESA meta key whose value will be used as one operand of join condition.
Enrichment Source Column Name	Enrichment source column name whose value will be used as the other operand of the join condition.

Test Rule Section

Note: The Test Rule section is available in NetWitness Platform 11.5 and later.

In the Test Rule section, you can validate your ESA rule to determine if the rule logic is working as expected before deploying the rule.



Field	Description
ESA Service	Select the ESA Correlation service to process the rule.
Input Data	Enter the input events to test the rule. You can download the events from the Investigate view in JSON format, copy the events, and paste them in this field.
Output Data	After you select an ESA Correlation service, input data, and click the Test Rule button, you can view the output of the rule here and verify that the rule is working according to your requirements. You can view the alerts in the output, but this test does not send any alert notifications. If you want to view all of the debug information for the test, include an `@Audit('stream')` annotation to your rule query.

The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service

Field	Description
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Syntax

Click **Show Syntax** to view the EPL syntax of conditions, statements, and debugging parameters. It also provides a warning when the syntax is invalid. For more information, see Rule Syntax Dialog.

Rule Syntax Dialog

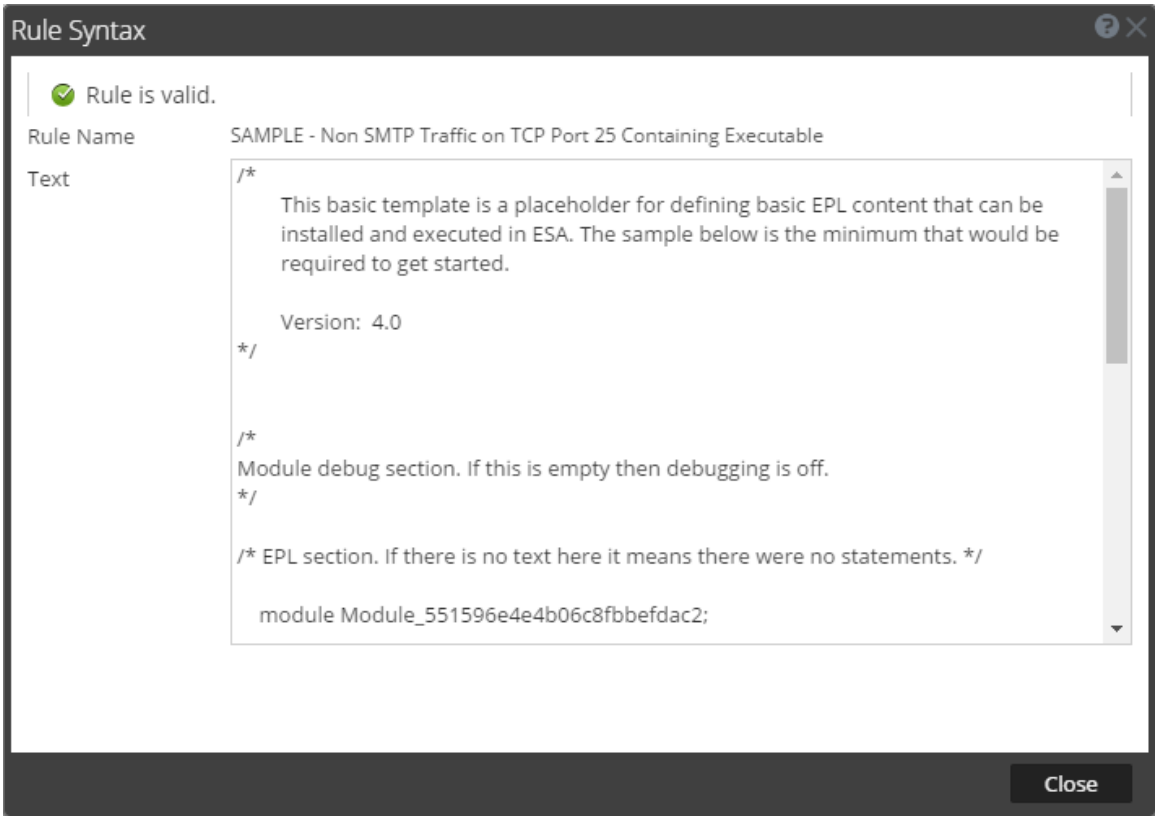
This topic describes the features of the Rule Syntax dialog. The Rule Syntax dialog displays the EPL syntax of conditions, statements, and debugging parameters, and provides a warning when the syntax is invalid.

Quick Look

To access this dialog:

1. Go to **(Configure) > ESA Rules**.
2. In the **Rule Library** view, do one of the following:
   a. Click and select **Advanced EPL** or **Rule Builder**.
   b. Double-click an existing rule.
   c. Select an existing rule and click in the **Rule Library** toolbar.
   d. In the row of an existing rule, select **> Edit**.
   The new or existing rule is displayed in a new tab, available to edit.
3. Click **Show Syntax** at the bottom of the tab.

The following figure shows an example of the Rule Syntax dialog showing a valid rule.



The following table describes the Rule Syntax dialog parameters.

Parameters	Description
Rule is valid or **Validation error in rule**	Indicates whether the rule syntax is valid or needs to be changed.
Rule Name	Displays the name of the rule.
Text	Displays the EPL syntax of conditions, statements, and debugging parameters if the rule is valid.

Deployment Panel

ESA rule deployments map rules from your rule library to the appropriate ESA Services and data sources. The Deployment panel ( (Configure) > ESA Rules > Rules tab) enables you to create and configure ESA rule deployments that specify:

- ESA Services
- Data Sources (This is available in NetWitness version 11.3 and later.)
- ESA Rules

When you are ready to start aggregating data and generating alerts from an ESA rule deployment, you deploy the ESA rule deployment to activate it.

Note: An ESA rule deployment can have only one ESA service. You can, however, use the same ESA service in multiple deployments.
In NetWitness Platform version 11.2 and earlier, the ESA service is the Event Stream Analysis service. In version 11.3 and later, it is the ESA Correlation service.

What do you want to do?

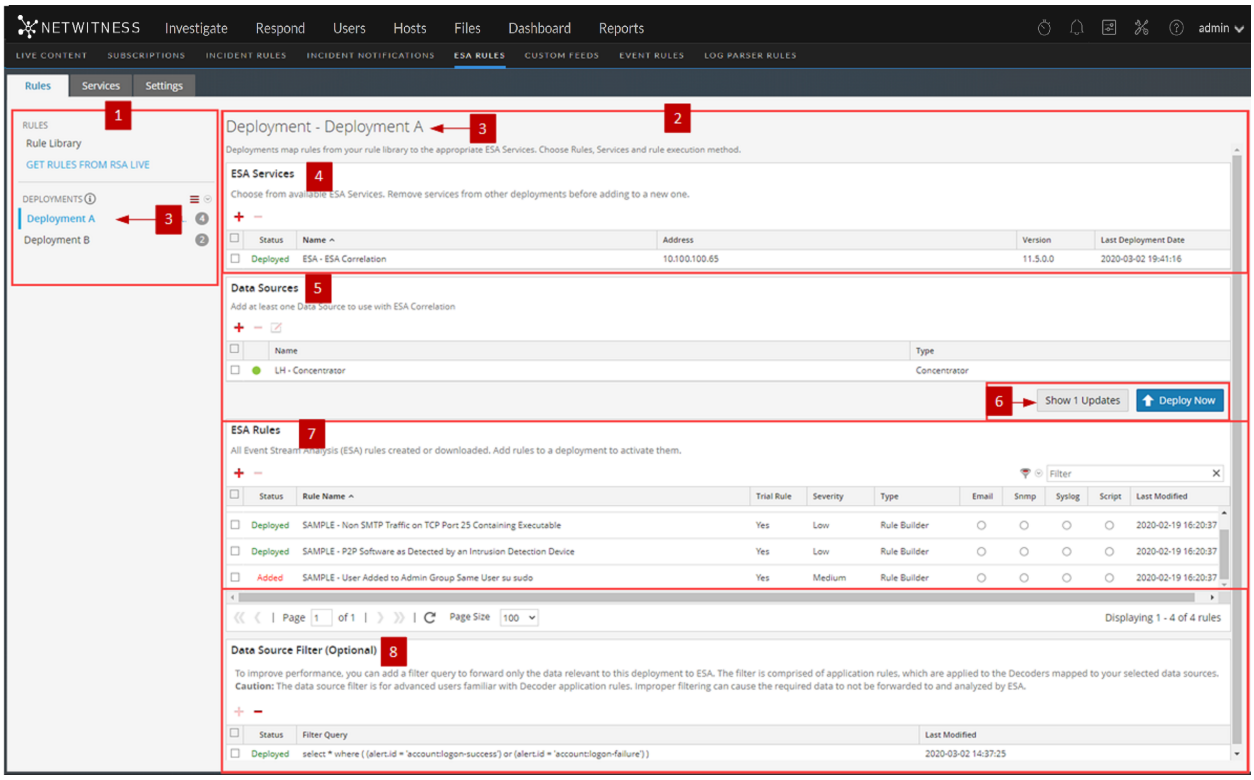
Role	I want to ...	Show me how
Content Expert	Add an ESA rule deployment.	ESA Rule Deployment Steps
Content Expert	Manage deployments.	Additional ESA Rule Deployment Procedures

Related Topics

- View Stats for an ESA Service

Quick Look

The following figure shows the Deployment panel.



1	Shows the options panel. Click the deployment in the options panel to view the deployment in the Deployment panel on the right.
2	Shows the Deployment panel.
3	Shows the ESA rule deployment name.
4	The ESA Services section shows the ESA Correlation service that processes the ESA rules and creates alerts.
5	The Data Sources section shows the data sources, such as Concentrators, that provide the data for the deployment.
6	The deployment options enable you to view the changes to the deployment (Show Updates) and deploy the ESA rule deployment on ESA (Deploy Now). The (Configure) > ESA Rules > Services tab shows the status of the deployment after it is deployed.
7	The ESA Rules section shows the rules that are used to trigger alerts in the deployment.
8	The Data Source Filter (Optional) section enables you to add a filter query to forward only the data relevant to this deployment to ESA. The data source filter is for advanced users familiar with Decoder application rules.

ESA Services Section

In the ESA Services section, you can manage each ESA service in the deployment.

The following table describes the actions you can perform in the ESA Services section.

Task	Description
+	Adds an ESA service to the deployment.
−	Removes the selected ESA service from the deployment.

The following table describes the columns in the ESA Services section.

Title	Description
Status	Indicates if the deployment status is **Added**, **Deployed**, **Updated**, or **Failed**.
Name	Name of the ESA service.
Address	IP address of the host where the ESA service is installed.
Version	Version of the ESA service.
Last Deployment Date	The date and time when the ESA service was last deployed.

Data Sources Section

Note: This option is available in NetWitness Platform version 11.3 and later.

In the Data Sources section, you can select one or more data sources, such as Concentrators, to use for your selected ESA Service.

The following table describes the actions you can perform in the Data Sources section.

Task	Description
+	Adds a data source for the selected ESA service to the deployment.
−	Removes a data source for the selected ESA service from the deployment.
	(This option is available in NetWitness Platform version 11.3.0.2 and later.) Enables you to change the configuration of a data source in an ESA rule deployment. You can change the data source password, SSL, port, and compression settings. When a data source password changes, it is important to change the password on the data source so that ESA can continue to communicate with the data source. **Note:** If you make any ESA service, data source, or ESA rule changes to an ESA rule deployment, you need to redeploy the deployment. For example, if you change the configuration of a data source in an ESA rule deployment, you must redeploy all the ESA rule deployments that contain that data source. When you set the compression level for a Concentrator on ESA, it sets the same compression level for that Concentrator for ESA Correlation Rules.

The following table describes the columns in the Data Sources section.

Title	Description
(Status)	Shows the status of the data source. A solid colored green circle indicates a running service and a white circle indicates a stopped service.
Name	Shows the name of the data sources used by the selected ESA service. You can specify the data sources separately for each ESA rule deployment.
Type	Shows the type of the data sources. Data sources can be Concentrators or Decoders. It is important that you choose data sources that have the appropriate data for the rules in the deployment. For example, if you have NetWitness Endpoint and you want to deploy the Endpoint Risk Scoring Rules Bundle, you must choose endpoint data sources.

Note: You can add a Log Decoder as a data source for ESA, but it is better to add a Concentrator to take advantage of undivided aggregation as the Decoder may have other processes aggregating from it.

Position Tracking Information

The ESA Correlation service continuously streams data from the data sources like decoders (log and network), and concentrators. ESA retrieves events from the data sources, and applies rules to generate alerts to detect malicious activities. When you deploy a data source, ESA starts processing information from the latest available session, by default. Position Tracking Information enables you to visualize the progress of the sessions that ESA has processed, and provides information on the session IDs and the time/date when the events were processed.

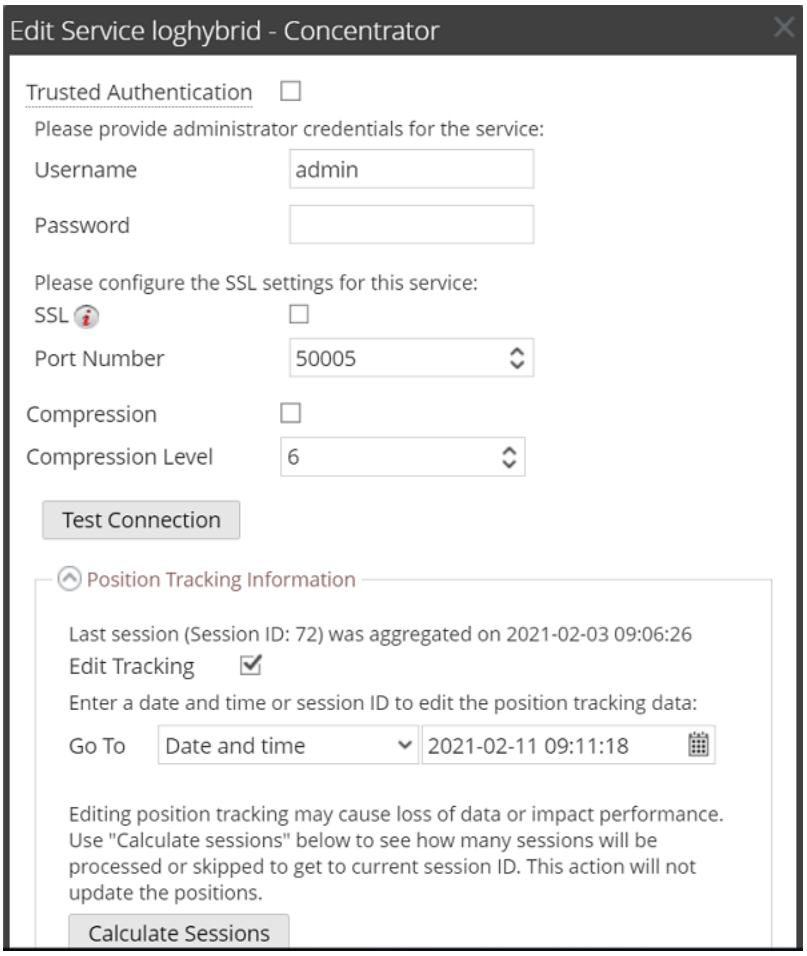
Edit Position Tracking Information enables you to:

- Visualize the number of sessions that a particular ESA data source has already analyzed, review the number of sessions ESA would process after you edit the position tracking, and plan your work.
- Edit the tracking position information based on:
  - Date and Time (Collection Time)
  - Session ID
- Edit position tracking for multiple data sources before you deploy them.
- Calculate the number of sessions that the ESA Correlation Service is scheduled to process for a particular data source to either process, reprocess, or skip sessions with respect to the current position of the data source.

Note: The Edit Position tracking feature with the `Date and Time` option works based on the profile time settings in the Netwitness UI. This time-zone based time from the UI is converted to UTC, and is sent to the core, to retrieve the corresponding session ID for that time stamp.
Example: If the UI follows IST, the UI converts it to UTC and sends it to the core. The session ID is fetched for the specific UTC time stamp, and set to position tracking at deployment.

Editing Position Tracking Information

The following figure shows the Edit Position Tracking window.



To edit position tracking information:

1. Select the specific data source from the **Data Source** menu.
2. Click  .

The **Edit Service** window is displayed.

3. Select the **Trusted Authentication** check-box, or enter your administrative credentials (username and password).
4. In the **Position Tracking Information** menu, click the **Edit Tracking** check-box to select it.
a. If you want to edit the position tracking information based on date and time stamp:
In the **Go To** text field, select **Date and Time** and enter the date and time.
The ESA Correlation service starts processing the events from the date and time that you entered.

b. If you want to edit the position tracking information, based on the session ID:
In the **Go To** text field, select **Session ID** and enter the session ID.
The ESA Correlation service starts processing the events from the session ID that you entered.

5. Click **Calculate Sessions** to calculate the number of sessions that will be processed with respect to the existing position of the data source, if any.
6. Click **OK**.

7. Click **Deploy Now**.
The tracking position information will be deployed to the ESA Correlation service, only when the deployment is successfully completed.

Note: After you deploy, in the **Data Sources** menu, click on , to view the edit tracking information. The default time-out associated with this information is 1 minute.

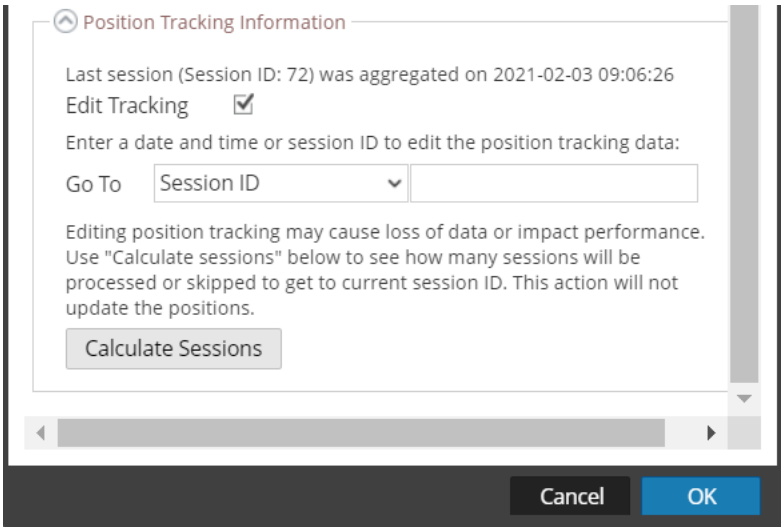
Use Case Scenario

This section provides information about how you can use position tracking information in a real world scenario.

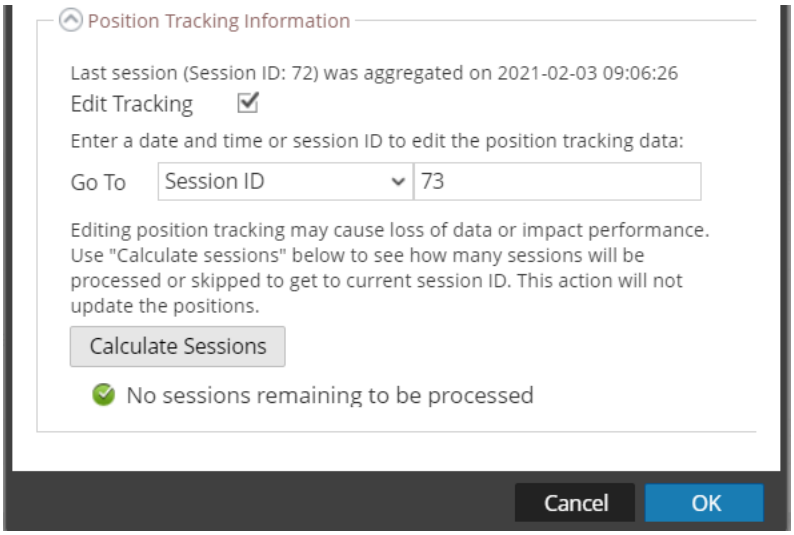
If you have deployed a data source with a total of 72 sessions that ESA has already processed, and if you want to start processing the events from the beginning, or go back with respect to the time or sessions ID.

1. Click **Edit Tracking**.
2. Enter 1 in the **Session ID** text box.
3. Click **Calculate Sessions**.
All the 72 sessions will be reprocessed.

The following image shows the use case scenario.



4. Enter a future Session ID.
The **No sessions remaining to be processed** message is displayed.



Note: Editing the tracking information is optional. If you add a new data source to an existing ESA deployment, and you do not edit the tracking information, ESA follows the default behavior to process events.

Importing Position Tracking Information

You can migrate the settings of position tracking for one or more data sources at the same time from an existing deployment, using the import function.

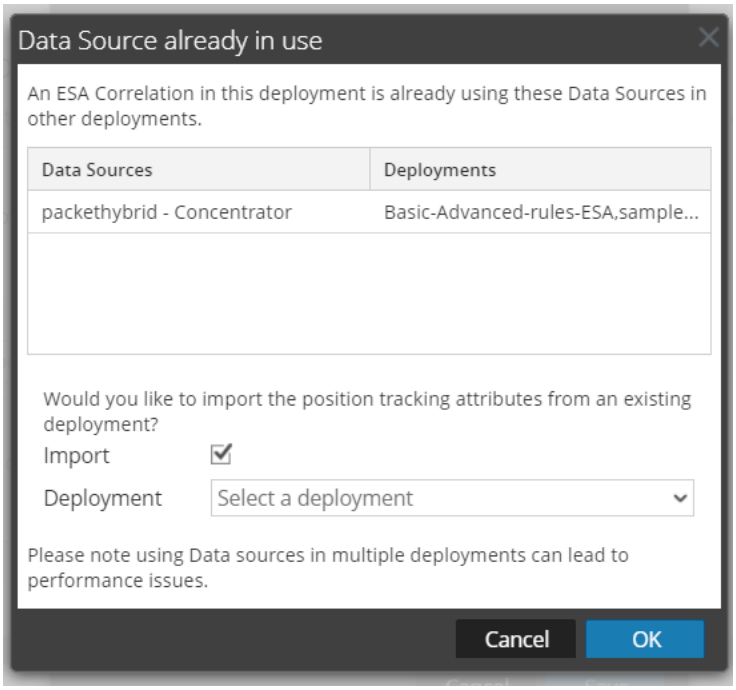
To import position tracking information from an existing deployment:

1. Go to **Configure** > **ESA Rules** > **Deployment** tab.
2. Create a new deployment. See ESA Rule Deployment Steps
3. Add ESA Services. See Step 2. Add an ESA Service

The Data Sources panel will be enabled.

4. Add Data Sources. See Step 3. Add Data Sources
5. In the Data Sources tab, click + to add a configured data source.
The Available Configured Data Sources window appears.
6. Select the Data Sources and click **Save**.The Data Sources window displays the data sources that are already in use by the other deployments.

The following figure shows the Import functionality.



7. Select the **Import** check-box.

8. In the **Deployment** dropdown box, select the deployment from which you want to import the position tracking attributes from an existing deployment. The dropdown menu lists all of the deployments with which data sources are shared. Review the list of data sources in the **Data Sources** menu, and click **OK**.

9. Add an ESA Rule. See Step 4. Add and Deploy Rules

10. In the Data Sources tab, click **Deploy Now**.
The deployment retrieves the position tracking information from the selected pre-existing deployment, and applies it to the new deployment.

Note: Ensure that you do not delete the deployment from where you imported the position tracking information, before you perform **Deploy Now** on the newly created or edited deployment. The import position tracking information function is disabled by default. Sharing data sources between multiple deployments can lead to performance issues.

Deployment Options

There are two deployment options below the Data Sources section. These options apply to the entire ESA rule deployment.

The following table describes these deployment options.

Task	Description
Show Updates	Enables you to view a history of updates to the deployment.
Deploy Now	Activates the ESA rule deployment. The selected ESA service starts aggregating data and generating alerts using the specified ESA rules in the deployment. You need to add ESA Rules to the deployment before deploying the ESA rule deployment.

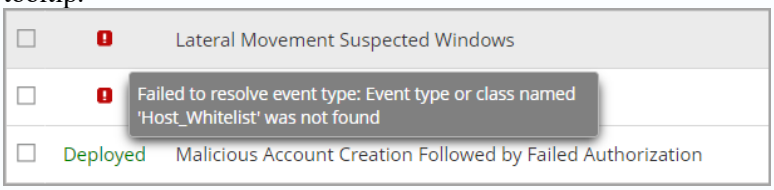
ESA Rules Section

In the ESA Rules section, you manage rules in the deployment. This section lists all rules that are currently in the deployment.

The following table describes the actions you can perform in the ESA Rules section.

Task	Description
+	Opens the Deploy ESA Rules dialog, where you can select a rule.
−	Removes the selected ESA rules from the deployment.
	Filters the list of rules.
Filter	Enables you to search for a rule.

The following table describes the columns in the ESA Rules section.

Title	Description
Status	Indicates the rule status: • **Deployed** - the rule is deployed. • **Updated** - the rule has been updated since the last deployment. • **Added** - the rule has been added since the last deployment. • **Disabled** - the rule is disabled due to an error in the rule or an error during the deployment of the rule. In NetWitness Platform version 11.3.0.2 and later, if a disabled rule has an error message, it shows ! in the Status field. Hover over the rule to view the error message tooltip. 
Rule Name	Describes the purpose of the ESA rule.
Trial Rule	Indicates whether the rule is Deployment mode to see if the rule runs efficiently.
Severity	Shows the threat level of alert triggered by the rule.
Type	Shows the type of the ESA rule. For more information, see ESA Rule Types.

Title	Description
Email, SNMP, Syslog, Script	Indicates which notification types are used for alerts generated by the rules. (ESA SNMP notifications are not supported in NetWitness version 11.3 and later.)
Last Modified	Shows the date and time when the ESA rule was last modified.

Data Source Filter (Optional) Section

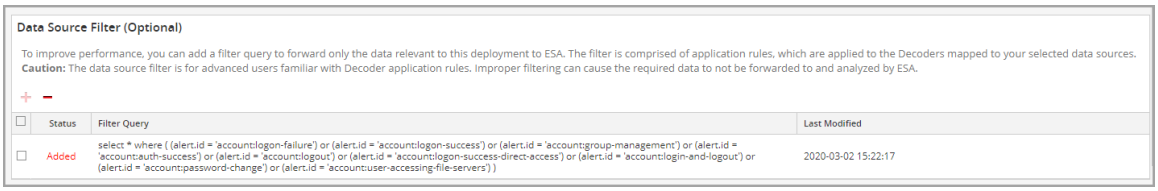
Note: This option is available in NetWitness version 11.5 and later.

The data source filter is optional. If you have a medium to large NetWitness Platform deployment and you have high throughputs, you can add a filter query to forward only the data relevant to this deployment to ESA. This helps to filter out certain types of traffic that does not add value to the analysis of the data in the ESA rule deployment.

Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.

IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.



The following table describes the columns in the Data Source Filter (Optional) section.

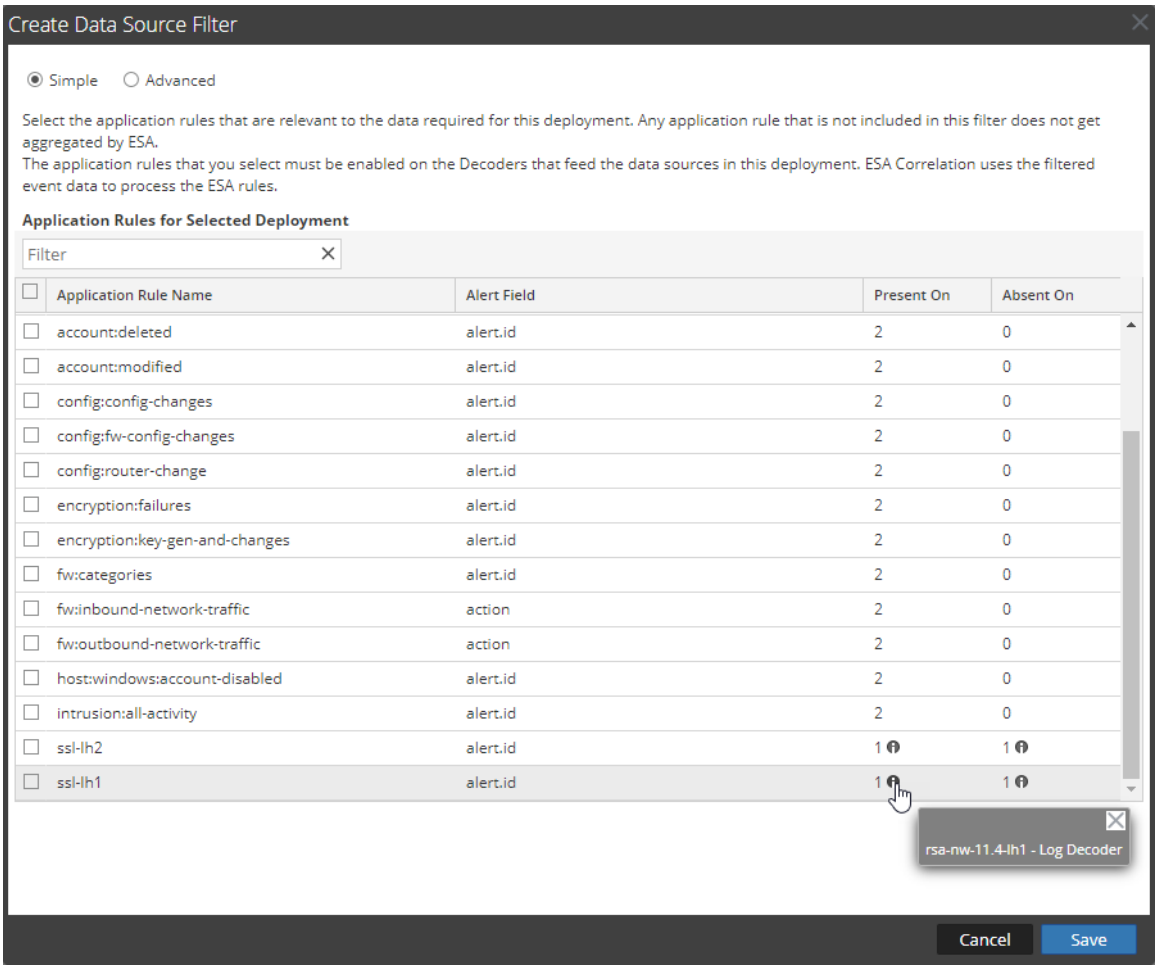
Title	Description
☐	Enables you to select and remove a data source filter.
Status	Indicates the data source filter status: • **Added** - the filter is added to the deployment, but it is not yet deployed. • **Deployed** - the filter is actively streaming only the relevant data as defined in the filter query.

Title	Description
Filter Query	Shows the query that contains the application rules. Each application rule query is separated by an "or" condition. You can use either the simple or advanced option to create the data source filter. You can only have one data source filter per deployment for your selected data sources.
Last Modified	Shows the date that the filter was added to the ESA rule deployment.

Create Data Source Filter Dialog - Simple

When you create the data source filter, you select application rules to be included in the filter query. The application rules that you select must be enabled on the Decoders that feed the data sources in this deployment. ESA Correlation uses the filtered event data to process the ESA rules.

Caution: The data source filter is for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.



The following table describes the columns in the Create Data Source Filter dialog.

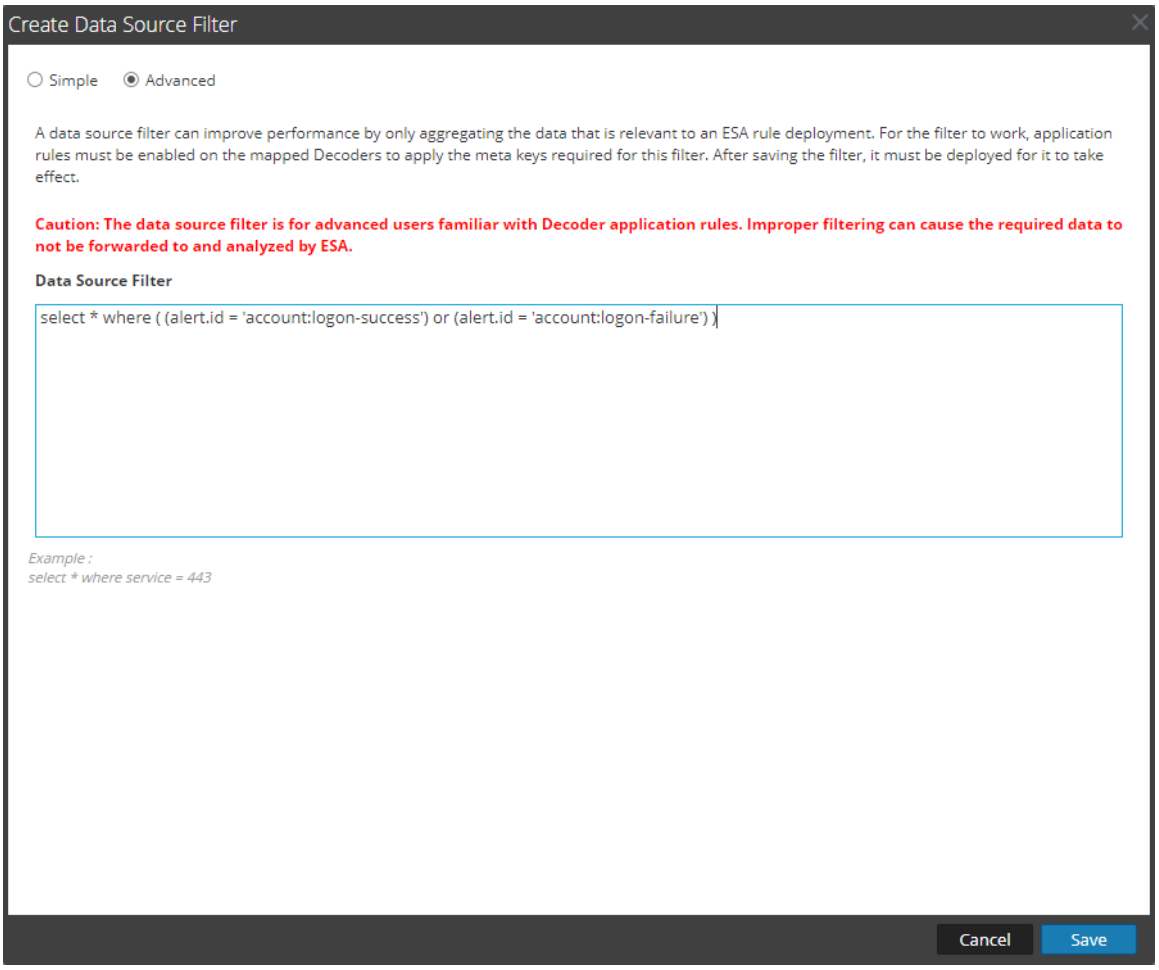
Title	Description
Filter	Enables you to search for a rule. For example, you can type "account" to search for application rules that contain that word.
Application Rule Name	Shows the names of the application rules on the Decoders mapped to the data sources in this deployment.
Alert Field	Shows the meta key used in the alert.
Present On	Shows the number of Decoders mapped to the data sources in this deployment that *have* the listed application rule. If the rule is not present on all of the Decoders, you can hover over the icon to view the names of the Decoders that contain the rule.
Absent On	Shows the number of Decoders mapped to the data sources in this deployment that *do not have* the listed application rule. If a rule is not present on all of the Decoders, you can hover over the icon to view the names of the Decoders that do not contain the rule.

For more information, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.

Create Data Source Filter - Advanced

If necessary, you can use the advanced filter instead of the simple filter to add your data source filter query directly. The individual application rule queries must be separated by an "or" condition. For more information on creating and writing Decoder rules, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.

Caution: The data source filter is for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.


Create Data Source Filter
Simple
Advanced
A data source filter can improve performance by only aggregating the data that is relevant to an ESA rule deployment. For the filter to work, application rules must be enabled on the mapped Decoders to apply the meta keys required for this filter. After saving the filter, it must be deployed for it to take effect.
Caution: The data source filter is for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.
Data Source Filter
select * where ((alert.id = 'account:logon-success') or (alert.id = 'account:logon-failure'))
Example :
select * where service = 443
Cancel
Save

Deploy ESA Services Dialog

The Deploy ESA Services dialog displays all ESA services available to be added to an ESA rule deployment.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Configure an ESA rule deployment.	ESA Rule Deployment Steps
Content Expert	Add a service.	ESA Rule Deployment Steps
Content Expert	Add data sources. (This is available in 11.3 and later.)	ESA Rule Deployment Steps
Content Expert	Add and deploy rules.	ESA Rule Deployment Steps

Related Topics

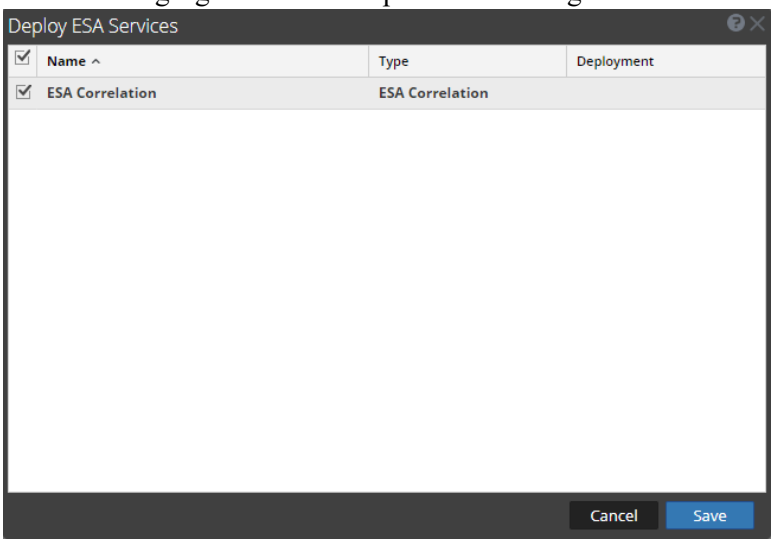
- Additional ESA Rule Deployment Procedures
- View Stats for an ESA Service

Quick Look

To access this dialog:

1. Go to  **(Configure) > ESA Rules**.
   The Rules tab opens by default.
2. In the options panel, under the **Deployment** section, select or add a deployment.
3. In the **ESA Services** panel, click .
   The Deploy ESA Services dialog is displayed.

The following figure is an example of this dialog.



The following table describes the parameters of the Deploy ESA Services dialog.

Parameters	Description
Name	Displays the name of configured ESA services.
Deployment	Displays the ESA rule deployments to which the service has already been added.

Deploy ESA Rules Dialog

The Deploy ESA Rules dialog enables you to filter and select rules to deploy to an ESA service.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Configure an ESA rule deployment.	Step 1. Add an ESA Rule Deployment
Content Expert	Deploy a rule	Step 4. Add and Deploy Rules

Related Topics

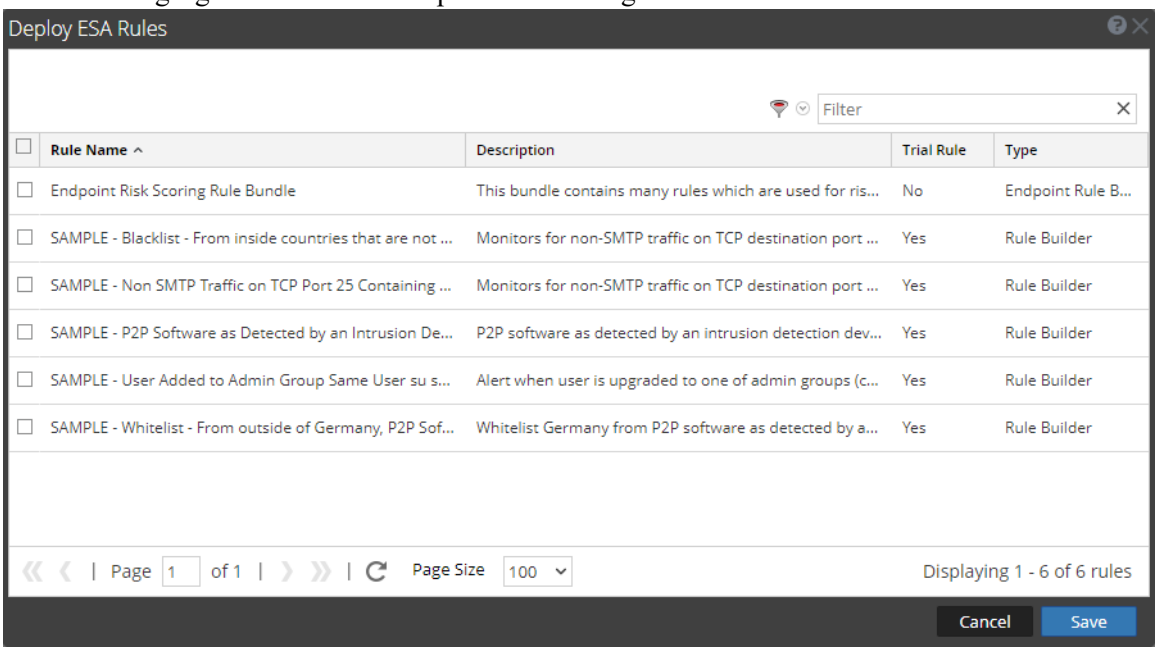
- Additional ESA Rule Deployment Procedures

Quick Look

To access this dialog:

1. Go to  **(Configure) > ESA Rules**.
   The Rules tab opens by default.
2. In the options panel, under the **Deployment** section, select or add a new deployment by clicking  **> Add**.
3. If you add a new deployment, type the name of the deployment in the box in the options panel.
4. In the **ESA Rules** panel, click .
   The Deploy ESA Rules dialog is displayed.

The following figure shows an example of this dialog.



The following table describes the parameters of the Deploy ESA Rules dialog.

Parameters	Description
	Filters the list of rules based on severity and type. The text box beside this icon filters based on rule name.
Rule Name	Displays the name of the rule.
Description	Describes the rule.
Trial Rule	Indicates whether or not the rule is a trial rule.
Type	Indicates the type of rule: RSA Live ESA, Advanced EPL, or Rule Builder.

Updates to the Deployment Dialog

The Updates to the Deployment dialog displays changes made to the deployment, such as adding a rule or service. Deployment updates are indicated by the update icon ( ) next to the name of the deployment in the Rules tab options panel.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Deploy rules to run on ESA.	ESA Rule Deployment Steps
Content Expert	Edit or delete an ESA rule deployment.	Edit the ESA Rule Deployment Name or Delete a Deployment
Content Expert	View deployment updates.	Show Updates to an ESA Rule Deployment

Related Topics

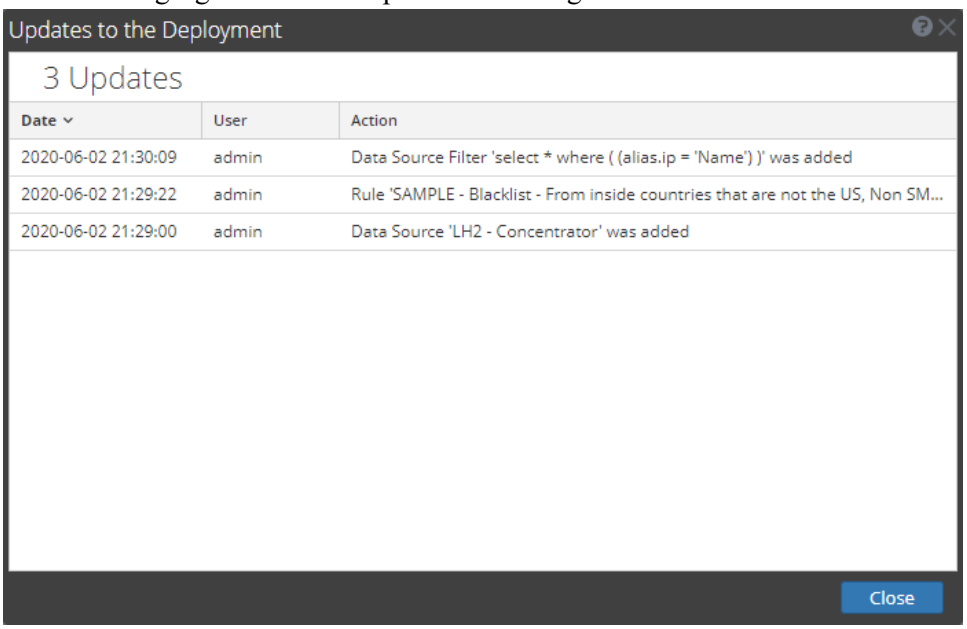
- Replace an ESA Service in an ESA Rule Deployment
- Add or Remove a Data Source
- Edit or Delete a Rule in a Deployment

Quick Look

To access this dialog:

1. Go to **(Configure) > ESA Rules**.
   The Rules tab opens by default.
2. In the options panel, under the **Deployments** section, select or add a deployment.
3. In the **Deployment** panel, click **Show Updates**.
   The Updates to the Deployment dialog is displayed.

The following figure is an example of this dialog.



The Updates to the Deployment dialog displays the number of updates at the top of the dialog. The following table describes the parameters of this dialog.

Parameters	Description
Date	Displays the day and time of the update.
User	Displays the user who made the update.
Action	Describes the update.

Services Tab

This topic provides an overview of the **(Configure) > ESA Rules > Services** tab. The Services tab shows the status of the deployments on each ESA service.

What do you want to do?

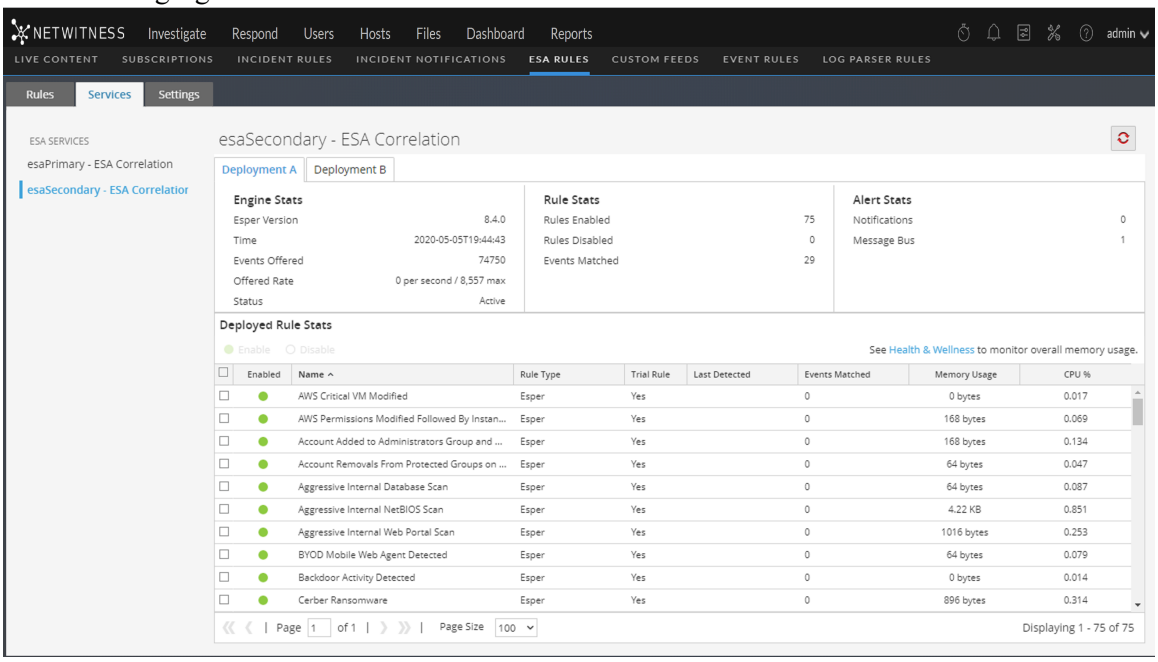
Role	I want to ...	Show me how
Content Expert	Troubleshoot Services Tab.	Troubleshoot ESA
Content Expert	View deployment Stats for an ESA Service.	View Stats for an ESA Service

Related Topics

- View a Summary of Alerts

Quick Look

The following figure shows the Services tab:



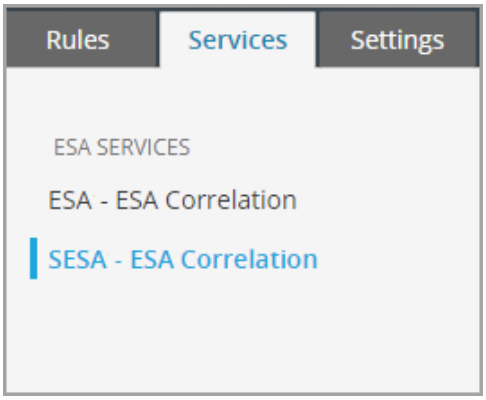
(This option is available in NetWitness version 11.3 and later.) If an ESA Correlation service has multiple deployments, under the service name, you will see a tab for each deployment. In the above example, there are two deployment tabs, Deployment A and Deployment B. Each tab displays information specific to that deployment.

The Services tab has the following sections:

- ESA Services panel (on the left)
- General Stats panel (top right)
- Deployed Rule Stats panel (bottom right)

ESA Services Panel

The ESA Services panel lists the name of each ESA service added to NetWitness.



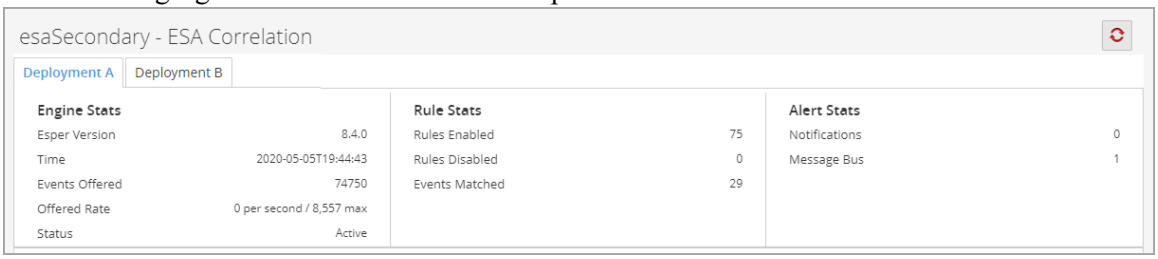
General Stats Panel

The General Stats panel provides information on the Esper engine, rules, and alerts.

The General Stats panel contains the following sections:

- Engine Stats
- Rule Stats
- Alert Stats

The following figure shows the General Stats panel.



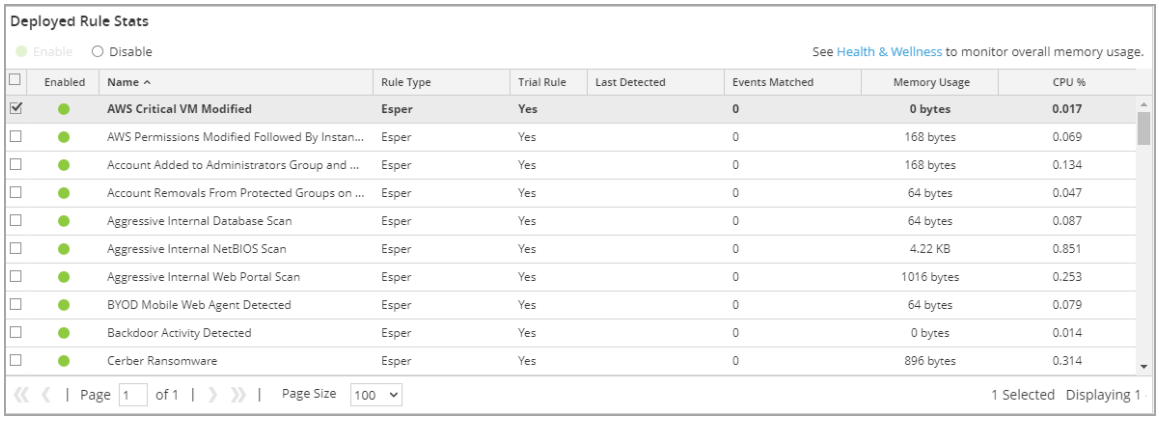
The following table lists and describes the parameters in each section.

Sections	Parameter	Description
Engine Stats	Esper Version	Esper version running on the ESA service
	Time	Time when the last event was sent to Esper Engine
	Events Offered	Number of events processed by the ESA service since the last service start
	Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events.
	Status	Shows the status of the deployment. A status of **Active** means that the deployment is active. A status of **Inactive** means that there was probably an error starting the deployment. Check the error log file for more information: `/var/log/netwitness/correlation-server/correlation-server.log`.
Rule Stats	Rules Enabled	Number of rules enabled
	Rules Disabled	Number of rules disabled
	Events Matched	Total number of events matched to all rules on the ESA service
Alert Stats	Notifications	The total number of notifications sent by email, SNMP, syslog, or script for the deployment. (ESA SNMP notifications are not supported in NetWitness Platform version 11.3 and later.)
	Message Bus	The total number of alerts sent to Respond for the deployment

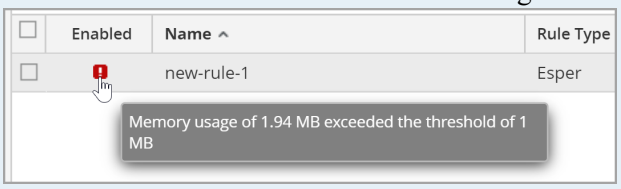
Deployed Rule Stats Panel

The Deployed Rule Stats panel provides details on the rules that are deployed on the ESA service.

The following figure shows the Deployed Rule Stats panel.



The table lists the various parameters in the view and their description.

Parameters	Description
Enable	Enables a rule that was disabled.
Disable	Disables a rule that was enabled.
Health & Wellness link	Enables you to monitor overall memory usage and health of your ESA Correlation service.
Enabled	Indicates whether the rule is enabled or disabled. A green circle icon indicates that the rule is enabled. A white circle icon indicates that the rule is disabled. If a disabled rule has an error message, it shows in the Enabled field. Hover over the icon to view the error message tooltip. The following example shows that the rule was disabled because it exceeded the configured memory threshold for that rule. 
Name	Name of the ESA rule.
Rule Type	(This field applies to version 11.3 and later.) **Endpoint** indicates a rule from the Endpoint Risk Scoring Bundle and **Esper** indicates Esper-specific rules, such as Rule Builder and Advanced EPL rules.
Trial Rule	Indicates if the rule is running in trial rule mode.
Last Detected	The last time alert was triggered for the rule.
Events Matched	The total number of events that matched the rule.
Memory Usage	The total amount of memory used by the rule. **Note:** The Endpoint Risk Scoring Rules Bundle rules do not show memory usage.
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each. (This field is available in version 11.5 and later.) **Note:** The Endpoint Risk Scoring Rules Bundle rules do not show CPU usage.

Settings Tab

This topic describes the components of the **(Configure) > ESA Rules > Settings** tab. In the Settings tab, you can perform the following tasks:

- MetaKey Referances
- Viewing the List of Meta Entities
- Enabling Meta Entity in the ESA Correlation Server
- Building Rules with Custom Meta Entities
- Configure a data enrichment source

What do you want to do?

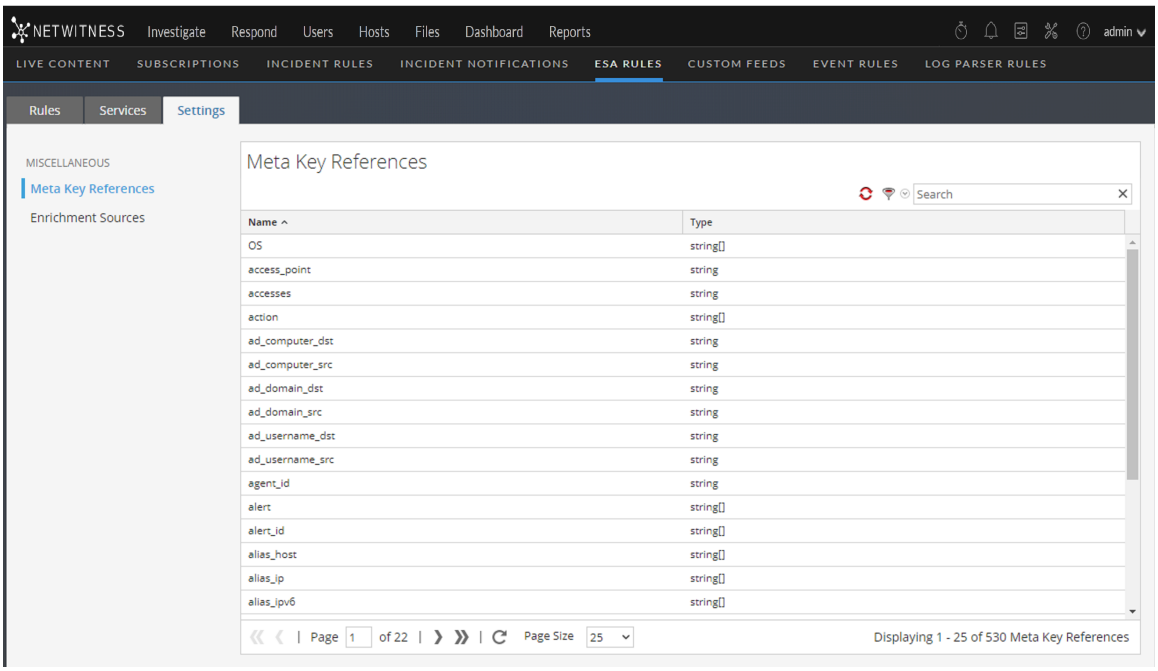
Role	I want to ...	Show me how
Content Expert	Configure an in-memory table as an enrichment source. (Recurring In-Memory Tables are no longer supported in version 11.3 and later.)	Configure an In-Memory Table as an Enrichment Source
Content Expert	Configure a Context Hub list as an enrichment source.	Configure a Context Hub List as an Enrichment Source

Related Topics

- Add a Data Enrichment Source
- Rule Builder Tab.

Quick Look

The following figure shows the Meta Key References section in the Settings tab.



Meta Key References

The Meta Key References section lists each Meta key used by the ESA Correlation server and the type of value the key requires.

Meta key entities are configured to be a part of Event Schema. After you upgrade from NetWitness® Platform 11.5, or previous releases to NetWitness® Platform 11.6, you can enable the string [] meta keys entities. You can create rules and configure alerts based on the meta key entities that you select. You can also add meta entities to create rules. The meta entities retrieves data from the data sources to trigger alerts.

Note: Only the meta keys comprising of string [] are supported.

Viewing the List of Meta Entities

Perform the following steps to view the list of meta entities that are available in the system:

1. Click to **Configure** > **ESA RULES** > **Settings** > **Meta Key References.**

The Meta Key References page displays all the meta keys that are enabled.

2. Enter the name of the meta key in the Search text box to view specific meta entities.

Enabling Meta Entity in the ESA Correlation Server

After an upgrade to NetWitness 11.6 or later, you need to enable the meta entities in the ESA Correlation server to create rules and configure alerts based on the meta key entities.

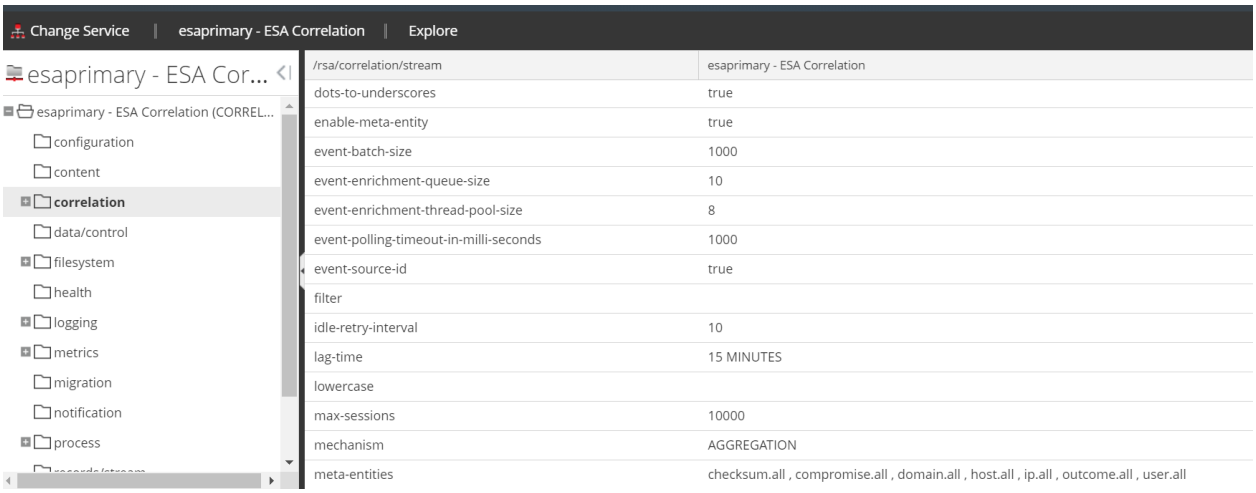
To enable meta entities in the ESA Correlation server:

1. Click (Admin) **Admin** > **Services**.

2. In the Services view, select the ESA Correlation service.

3. Click the settings icon and go to **View** > **Explore**.

The ESA Correlation Configuration page lists all the services.



4. Click **enable-meta-entity** and set it to **true**.

Note: By default, **enable-meta-entity** is set to **false**. Ensure that you enable only the specific meta entitiy using which you want to build a rule. Enabling multiple meta entities affects the system performance.

5. Add custom meta entities by appending or editing the meta-entities field values under explore view.

In the Explore view node list for the ESA Correlation service, select **Correlation** > **Stream**.All the stream settings are listed. You can add custom meta entities to the list as per the requirement. Once you enable custom meta entities, you can view them under Meta Key References. Once you enable the custom meta key entities, the rules will start evaluating the meta key entities and start to evaluate and trigger alerts as per the configured conditions. For more information, see Viewing the List of Meta Entities

The following meta entities are listed in the ESA Correlation Configuration page:

- checksum.all
- compromise.all
- domain.all
- host.all
- ip.all
- outcome.all
- user.all

Meta Entity Usage Example

ip_all, and *host_all* are some of the examples of meta entities.

For example, the *ip_all* meta entity can be used to check for any malicious IPs. It combines all the IPv4 meta keys (*ip_src*,*alias_ip*, *ip_dst*, and *forward_ip*) and creates a string array to retrieve information and trigger alerts.

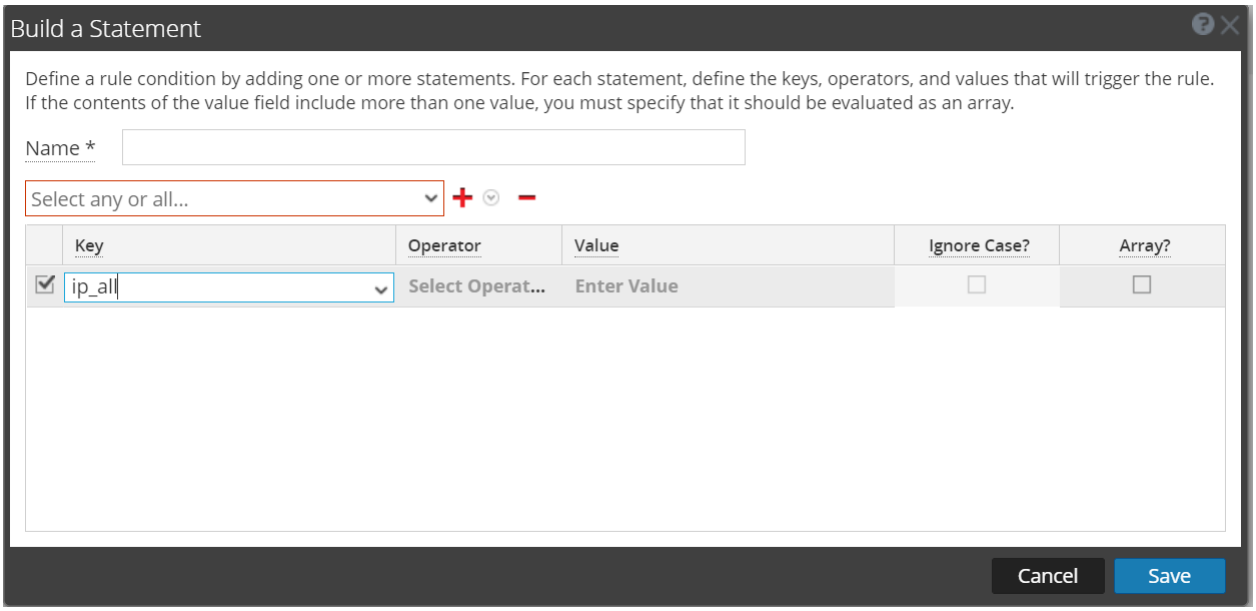
Note: You can add and enable custom meta entities (string arrays) to the list.

Building Rules with Custom Meta Entities

This topic provides instructions to define rule criteria in Rule Builder by adding statements. A statement is a logical grouping of rule criteria in the Rule Builder. You add statements to define what a rule detects.

The following graphic shows an example of a Rule Builder statement using meta entities.

Every statement contains a key and value. Then, you build logic around the pair by selecting an option in each other field.



Prerequisites

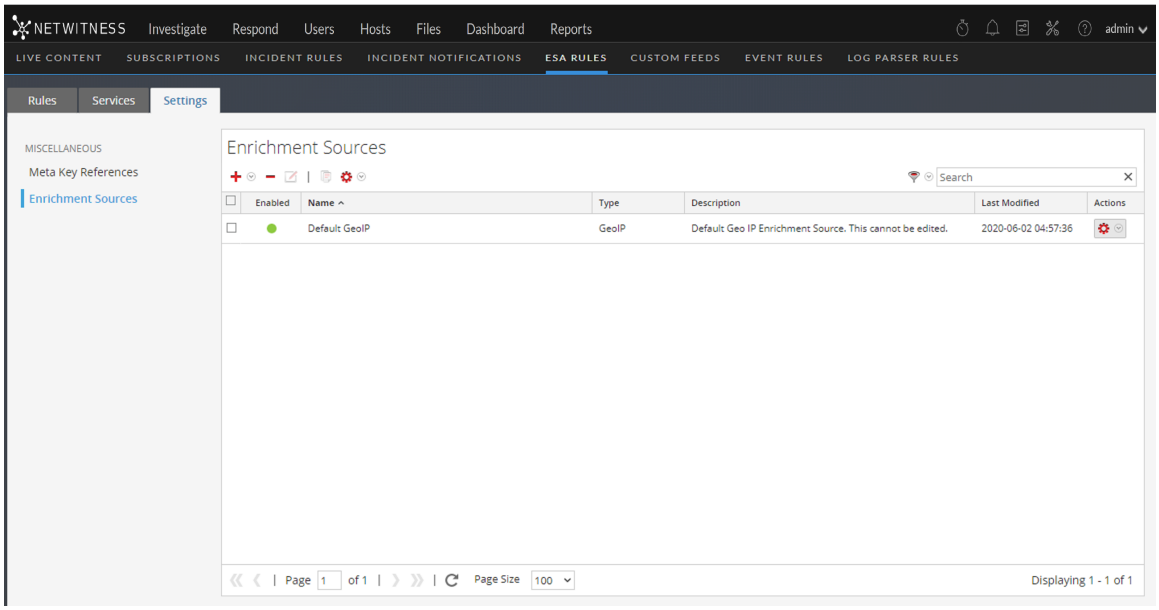
To build a rule statement, you must know the meta key and the meta value. For a complete list of meta keys, go to **Configure** > **ESA Rules** > **Settings** > **Meta Key References**.For more information about building rules with meta entities, see Rule Builder Tab.

Enrichment Sources

In the Enrichment Sources section, you can use the following external data sources:

- GeoIP
- In-Memory Table (Add hoc only - Recurring In-Memory Tables are no longer supported in version 11.3 and later.)
- Context Hub

The following figure shows the Enrichment Sources section in the Settings tab.



NetWitness® Platform XDR

Version 12.0

ESA Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Event Stream Analysis Overview **5**

Upgrade Considerations for ESA Analytics 6

Configure ESA Correlation Rules **7**

Data Source Configuration Changes 7

Endpoint Risk Scoring Rules Bundle 7

ESA Correlation Rules Configuration Workflow 9

Prerequisites 9

Procedure 10

ESA Correlation Health and Wellness Monitoring 11

Upgrade Considerations for ESA Hosts 11

Trial Rule Status Changes 11

Upgrade Considerations for ESA Rule Deployments (from 11.2.x and Earlier only) 12

Upgrade Considerations for ESA Rule Deployments (from 11.3.x) 12

ESA Rule Deployments Upgrade Example (from 11.2.x and Earlier only) 13

Additional ESA Correlation Rules Procedures **17**

Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys 17

Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules 19

Adjust Custom ESA Rule Builder and ESA Advanced Rules 20

Example ESA Correlation Server Warning Message for Missing Meta Keys 21

Multi-Valued Warning Message Example 21

Single Value Warning Message Example 21

Configure Meta Keys on New ESA Hosts to Match Upgraded ESA Hosts in the Same NetWitness Platform Network 21

Configure Advanced Settings for an ESA Correlation Service 23

Access Advanced Settings for an ESA Correlation Service 23

Enable or Disable Sending ESA Rule Alerts to the Respond View 24

Enable ESA Correlation Service Debugging for All Rules 26

Configure Maximum Events per Alert for All Rules 27

Adjust Maximum Sessions for the ESA Data Source Filter 27

Configure Meta Keys as Arrays in ESA Correlation Rule Values 29

Determine if a Meta Key is a String Array Type on ESA 29

Add the String Array Type Meta Key to ESA 30

Verify that the String Array Type Meta Key is Configured Correctly on ESA 30

Required String Array Meta Keys on the ESA Correlation Service 31

Remove Sensitive Meta Keys Globally from All Alerts for Data Privacy 32

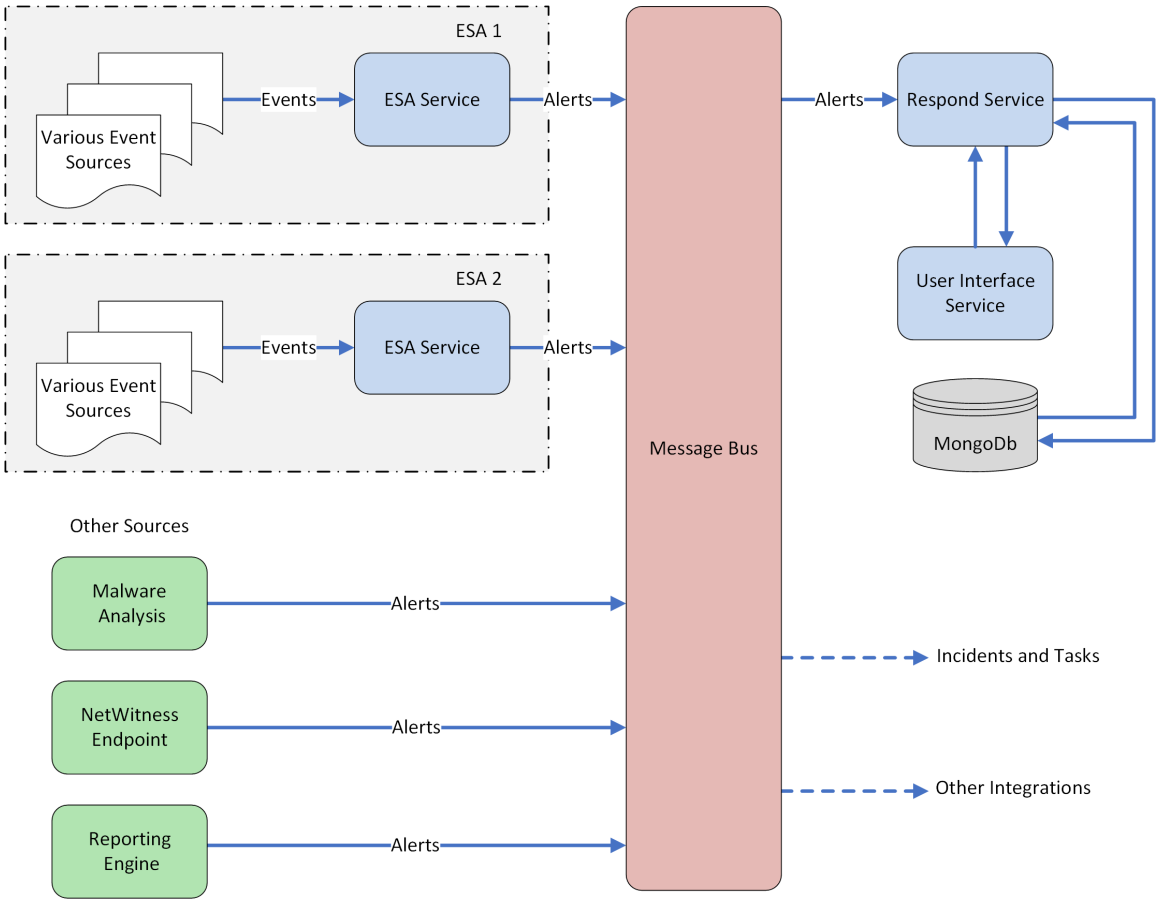
Configure Character Case for Advanced ESA Rules 33
Deploy Endpoint Risk Scoring Rules on ESA 35
Important Considerations when Deploying the Endpoint Risk Scoring Rules Bundle 35
Deploy the Endpoint Risk Scoring Rules Bundle on ESA 36
Change the Endpoint Risk Scoring Rule Bundle in a Deployment 43
View the Status of the Endpoint Risk Scoring Rules Deployment 44
Disable or Enable Individual Endpoint Risk Scoring Rules 45
Change Memory Threshold for ESA Rules 46
Change Memory Threshold for All Trial Rules 46
Change Memory Threshold for Individual Trial Rules and Non-Trial Rules 47
Start, Stop, or Restart ESA Service 49
Start the ESA Service 49
Stop the ESA Service 49
Restart the ESA Service 49
Start the ESA Service from the Command Line 49
Stop the ESA Service from the Command Line 49
Restart the ESA Service from the Command Line 50
View Audit Logs and Verify ESA Component Versions 51
View Audit Logs for Rules 51
Create Action 51
Update Action 52
Remove Rule Action 52
Delete Deployment Action 52
ESA Audit Logs on NW Server (11.5 and Later) 53
Verify ESA Correlation Version 54
References - Previous ESA Versions 55
Services Config View Data Sources Tab (11.2 and Earlier) 56
Services Config View Advanced Tab (11.2 and Earlier) 60
Alert Engine Settings 62
Event Stream Engine Settings 63
Whois Lookup Service Configuration (11.1.x to 11.4.x) 64
ESA Analytics Mappings (11.1.x to 11.4.x) 68
Toolbar 71
ESA Analytics Mappings 71
Module Settings (11.1.x to 11.4.x) 74
Configurations 75
Warm-Up State 78

Event Stream Analysis Overview

NetWitness Event Stream Analysis (ESA) provides advanced stream analytics such as correlation and complex event processing at high throughputs and low latency. It is capable of processing large volumes of disparate event data from Concentrators.

ESA's advanced Event Processing Language allows you to express filtering, aggregation, joins, pattern recognition, and correlation across multiple disparate event streams. Event Stream Analysis helps perform powerful incident detection and alerting.

The following diagram shows the high-level data workflow:



In NetWitness version 11.5 and later, There are only two services that can run on an ESA host:

- **ESA Correlation (ESA Correlation rules):** Creates alerts from ESA rules.
- **Contexthub Server (Context Hub):** Runs only on an ESA primary host. Contexthub Server provides enrichment lookup capability in the Respond and Investigate views. For information, see the *Context Hub Configuration Guide*.

Note: The Event Stream Analytics Server (ESA Analytics) service is not supported in NetWitness Platform version 11.5 and later.

The first service is the ESA Correlation service that creates alerts from ESA rules, also known as ESA Correlation Rules, which you create manually or download from Live.

In NetWitness 11.3 and later, the ESA Correlation service replaces the Event Stream Analysis service and is also known as ESA Correlation Server. The ESA Correlation service provides the same services as the Event Stream Analysis service with the added benefit of enabling you to specify different data sources for your ESA correlation rules. Like the Event Stream Analysis service, the ESA Correlation service installs on the ESA Primary and ESA Secondary host types.

The second service is the Contexthub Server service, which provides enrichment lookup capabilities in the Respond and Investigate views. It runs only on an ESA Primary host. For information, see the *Context Hub Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

IMPORTANT: Mixed mode is not supported for ESA hosts in NetWitness Platform version 11.5 and later. The NetWitness server, ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

Upgrade Considerations for ESA Analytics

The Event Stream Analytics Server (ESA Analytics) service is not supported or available in NetWitness version 11.5 and later. The Whois Lookup Configuration and ESA Analytics Mapping panels are no longer in the user interface [ (Admin) > System].

Note: Event Stream Analysis (ESA) is not end of life. ESA Correlation rules and the ESA Correlation service are supported. ESA Analytics, which is used for Automated Threat Detection, is different from ESA Correlation Rules and is EOL. In its place, you can use ESA Correlation as is offers more functional capabilities and better performance.

Configure ESA Correlation Rules

This topic provides high-level tasks to configure NetWitness Event Stream Analysis (ESA) Correlation Rules using the ESA Correlation service.

IMPORTANT: Mixed mode is not supported for ESA hosts in NetWitness Platform version 11.5 and later. The NetWitness server, ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

Data Source Configuration Changes

In NetWitness version 11.3 and later, the ESA Correlation service enables you to specify different data sources for different sets of rules. Instead of adding data sources, such as Concentrators, to the entire ESA Correlation service, you can specify different data sources for each ESA rule deployment. An ESA rule deployment includes an ESA Correlation service with its associated data sources and a set of ESA rules. For example, you may want to use Concentrators with HTTP packet data in one deployment and Concentrators with HTTP log data in another deployment. For more detailed information, see Deploy Rules to Run on ESA.

In NetWitness Platform 11.5 and later, you can add an optional data source filter to the data sources in your ESA rule deployments to improve performance. This allows your data sources to be filtered further so that only the data relevant to the deployment is forwarded to ESA. The filter is comprised of application rules, which are applied to the Decoders mapped to your selected data sources.

Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.

IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.

For more information, see “(Optional) Add a Data Source Filter” in ESA Rule Deployment Steps.

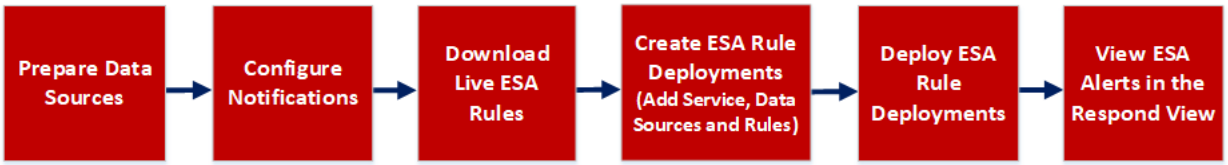
Endpoint Risk Scoring Rules Bundle

An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. Endpoint risk scoring rules only apply to NetWitness Endpoint. You can add the Endpoint Risk Scoring Rules Bundle to an ESA rule deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) in the ESA Rule Deployment.

The ESA Correlation service can process endpoint risk scoring rules, which generate alerts that are used in risk scoring calculations to identify suspicious files and hosts. To turn on risk scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. For instructions, see Deploy Endpoint Risk Scoring Rules on ESA. To configure NetWitness Endpoint, see the *NetWitness Endpoint Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

ESA Correlation Rules Configuration Workflow

The following diagram shows the high-level workflow for configuring ESA Correlation Rules with the ESA Correlation service.



ESA Rule Deployments are groups of ESA Rules processed by an ESA service to create alerts. In NetWitness 11.3 and later, the ESA Correlation service processes the ESA rules and creates alerts.

Before you can configure ESA Correlation Rules, install and configure the data sources (Concentrators) to use for the ESA rules. For example, you may have a Concentrator with HTTP packet data and another with Windows Log data. Next, configure the global notification methods that content experts can use for the ESA rules. For example, they may want to send an email notification when a rule creates an alert.

The NetWitness Live Content Management System (known as *Live*) is a valuable source of the latest internet security resources for NetWitness customers. RSA Live contains an extensive library of ESA rules to detect threats that you can use to save time. Download the rules for the events that you want to detect in your network to the ESA Rule Library and adjust them as needed for your network environment.

After you prepare your data sources and download Live ESA rules, you can create one or more ESA rule deployments. An ESA rule deployment contains an ESA service, one or more data sources, and a set of ESA rules. For example, you can create an ESA rule deployment that contains an ESA Correlation service, a Concentrator with HTTP packet data, and a set of ESA rules for HTTP packet data. When you are ready to have the ESA service run the rule set, you deploy the ESA rule deployment, which places the rules on ESA.

After you deploy an ESA rule deployment, verify that you can view the ESA alerts in the Respond view (Respond > Alerts).

Prerequisites

Make sure that you:

- Install the ESA Correlation service in your network environment.
- Install and configure one or more Concentrators in your network environment.
- Download or ensure that you have access to the Alerting with Correlation Rules User Guide for version 11.3 or later. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Procedure

The following table shows the high level tasks required to configure ESA Correlation Rules.

Tasks	Reference
1. Prepare data sources, such as Concentrators, to use for your ESA Correlation Rules.	Refer to Broker and Concentrator Configuration Guide.
2. Configure notifications for the ESA Correlation service.	Refer to Notification Methods.
3. Download Event Stream Analysis rules using Live. Configure the Live ESA Rule parameters for your environment.	Refer to Download Configurable RSA Live Rules.
4. Create ESA rule deployments*: Choose ESA Rules and the appropriate ESA service to use in the ESA rule deployment. For NetWitness 11.3 and later, you must also choose the data sources to use for these rules.	Refer to ESA Rule Deployment Steps.
5. Deploy ESA rule deployments.*	Refer to ESA Rule Deployment Steps.
6. View ESA alerts in the Respond view.	Refer to the NetWitness Respond User Guide.

*ESA rule deployments are groups of ESA Rules that are processed by an ESA service, such as the ESA Correlation service in NetWitness version 11.3 and later.

For additional optional advanced ESA Correlation Rules configuration procedures, see Additional ESA Correlation Rules Procedures.

For more information on alerting with ESA Correlation rules best practices, creating rules, working with trial rules, adding data enrichment sources, viewing statistics for an ESA service, and troubleshooting, see the Alerting with ESA Correlation Rules User Guide.

ESA Correlation Health and Wellness Monitoring

In NetWitness version 11.5 and later, New Health and Wellness provides improved and intuitive dashboards, monitors, and visualizations. The ESA Correlation Overview dashboard provides health statistics and trends on ESA rule deployments.



For more information, see "Monitor New Health and Wellness" and "Appendix A: New Health and Wellness Dashboards / ESA Correlation Overview Dashboard" in the System Maintenance Guide.

Upgrade Considerations for ESA Hosts

Mixed mode is not supported for ESA hosts in NetWitness Platform version 11.5 and later.

IMPORTANT: The NetWitness server (Admin server), ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

Trial Rule Status Changes

In NetWitness Platform 11.4 and later, ESA trial rules no longer change status after an upgrade or deployment. For example, if you change the status of a trial rule to disabled [(Configure) > ESA Rules > Services tab] and redeploy the ESA rule deployment [(Configure) > ESA Rules > Rules tab], the trial rule remains disabled. Previously, ESA trial rules could change status after an upgrade or when they were redeployed.

Upgrade Considerations for ESA Rule Deployments (from 11.2.x and Earlier only)

Caution: After upgrading to NetWitness version 11.3 or later, due to the ESA Correlation service data source changes, there are necessary data changes to migrated ESA rule deployments.

After the upgrade to version 11.3 or later, migrated ESA rule deployments change as follows:

1. If an ESA rule deployment contains two services before 11.3, the deployment splits into two deployments. You can only have one ESA Correlation service in an ESA rule deployment in version 11.3 or later.
2. If an ESA service has multiple ESA rule deployments before 11.3, they combine into one deployment in version 11.3 or later. You can still access your old deployments. For a detailed example, see ESA Rule Deployments Upgrade Example (from 11.2.x and Earlier only).

To support Endpoint and UEBA content as well as changes to ESA rules from Live, a data change from single-value (string) to multi-value (string array) is required for several meta keys within the ESA Correlation service. Some single-value meta keys are also required. See Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys.

Upgrade Considerations for ESA Rule Deployments (from 11.3.x)

Caution: Before upgrading from NetWitness 11.3.x to 11.4 or later, it is important to delete ESA rule deployments that do not contain an ESA Correlation service. The remaining ESA rule deployments should have been deployed at least once with the ESA Correlation service.

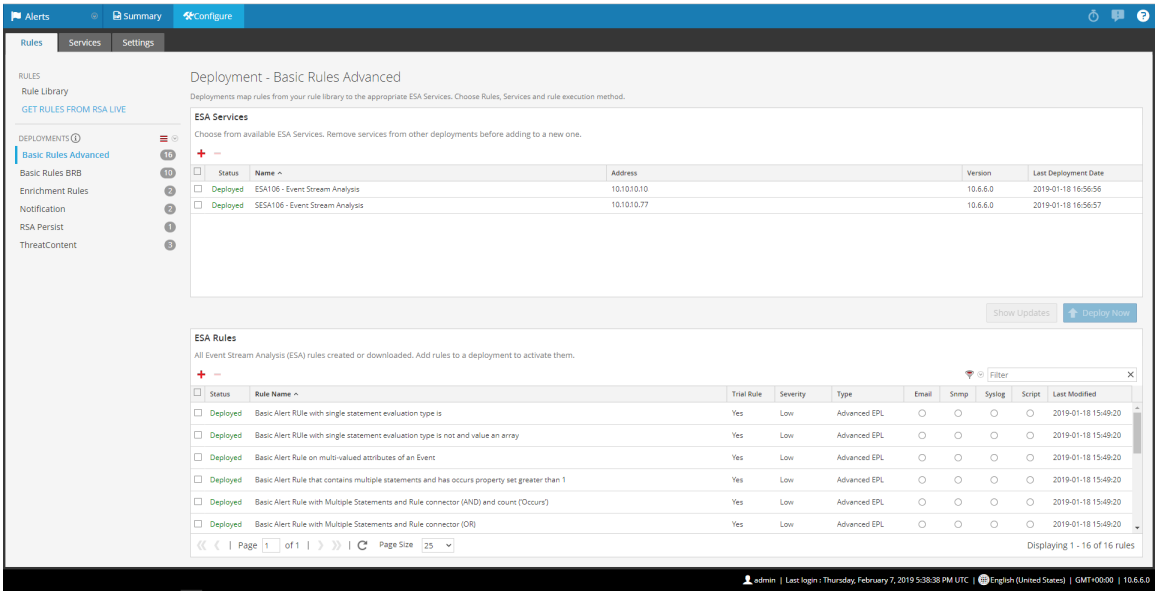
ESA Rule Deployments Upgrade Example (from 11.2.x and Earlier only)

Note: This example applies only to upgrades from NetWitness version 11.2.x and earlier to 11.3.x or 11.4.x.

to NetWitness 10.6.6 ESA rule deployments before and after an upgrade to version 11.3 or later. In this example, before the upgrade there are six 10.6.6 ESA rule deployments. Four deployments have both ESA primary and ESA secondary Event Stream Analysis services. One deployment has only the ESA primary service and another one has only the ESA secondary service. Each 10.6.6 ESA rule deployment before the upgrade has a set of rules as shown in the following table:

10.6.6 ESA Rule Deployment (Before Upgrade)	# of Rules	ESA Primary	ESA Secondary
Basic Rules Advanced	16	x	x
Basic Rules BRB	10	x	x
Enrichment Rules	2	x	x
Notification	2		x
RSA Persist	1	x	
Threat Content	3	x	x
Total # of Rules		**32**	**33**

The following figure shows the 10.6.6 ESA rule deployments before the upgrade.



In this example, after the upgrade to 11.3, there are only two functional ESA rule deployments; one has the ESA primary (ESAP) ESA Correlation service and the other has the ESA secondary (ESAS) ESA Correlation service. The ESAP - ESA Correlation service deployment contains 32 rules, which is the total of all of the rules in the 10.6.6 deployments that contained the ESA primary service (16 + 10 + 2 + 1 + 3). The ESA secondary ESA Correlation service contains 33 rules, which is the total of all the rules in the 10.6.6 deployments that contained the ESAS Service (16 + 10 + 2 + 2 + 3). Since the RSA Persist deployment only contained one ESA primary service in 10.6.6, that rule was added to the ESAP 11.3 deployment. Since the Notification 10.6.6 deployment contained one ESA secondary service, the rule was added to the 11.3 ESAS deployment.

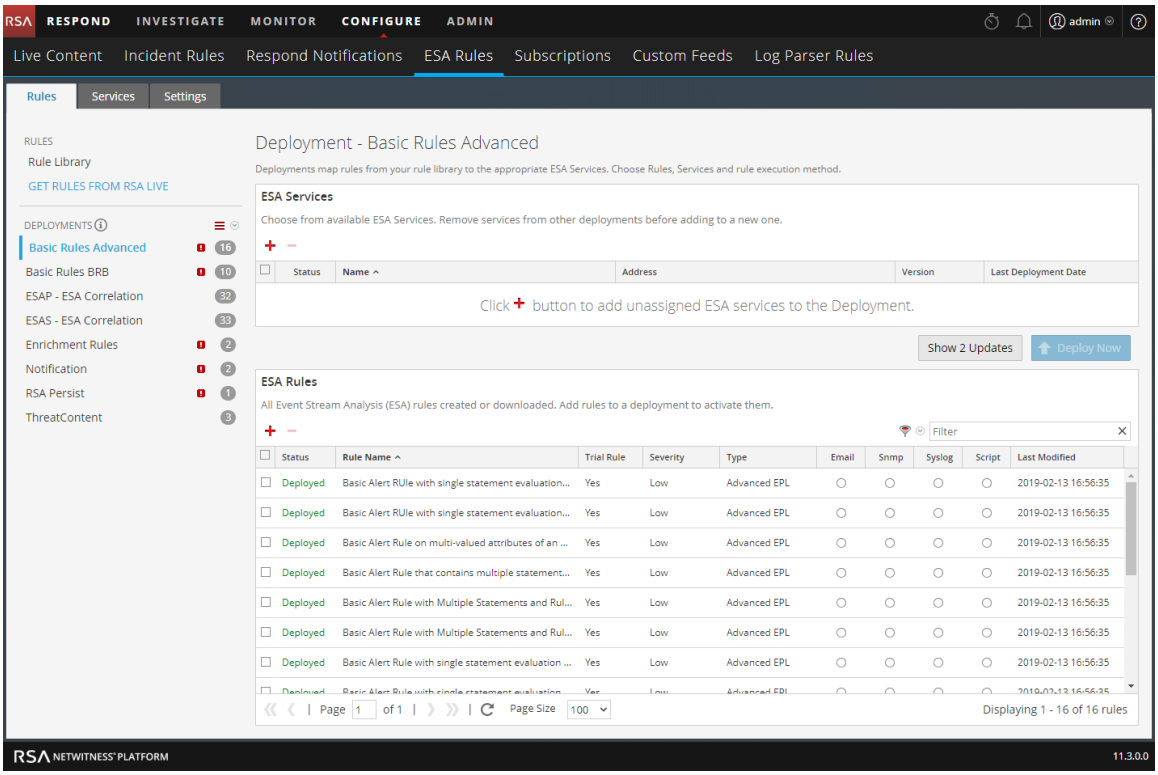
The following table shows the 11.3 ESA rule deployments after upgrade, the number of rules in each deployment, and which deployments have the ESA primary and ESA secondary ESA Correlation services.

11.3 ESA Rule Deployment (After Upgrade)	# of Rules	ESA Primary	ESA Secondary
Basic Rules Advanced	16		
Basic Rules BRB	10		
ESAP - ESA Correlation	**32**	**x**	
ESAS - ESA Correlation	**33**		**x**
Enrichment Rules	2		
Notification	2		
RSA Persist	1		
Threat Content	3	-	-
Total # of Rules		**32**	**33**

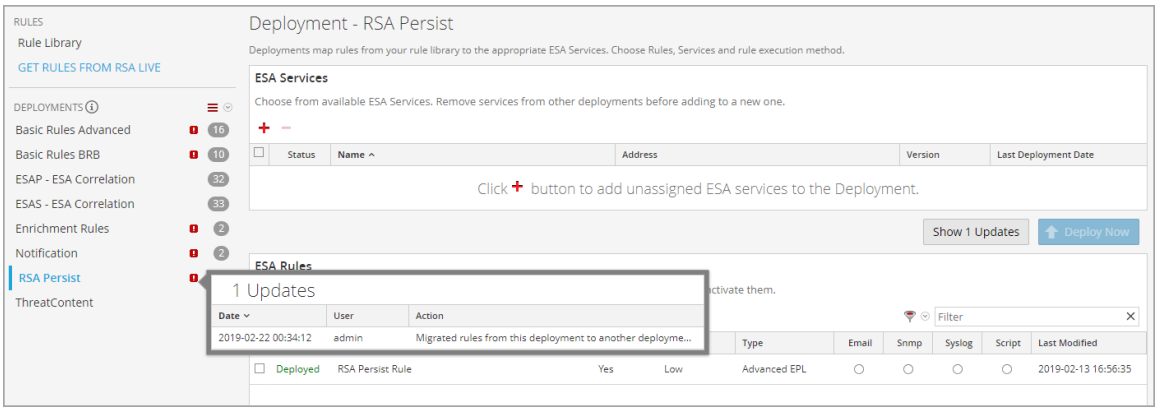
In this example, all of the 10.6.6 ESA rule deployments were moved to 11.3, but they have no services after the upgrade. The only functional 11.3 ESA rule deployments are ESAP and ESAS. The 10.6.6 ESA rule deployments are preserved in case you want to use them.

Caution: Before upgrading from NetWitness 11.3.x to 11.4, it is important to delete ESA rule deployments that do not contain an ESA Correlation service. The remaining ESA rule deployments should have been deployed at least once with the ESA Correlation service.

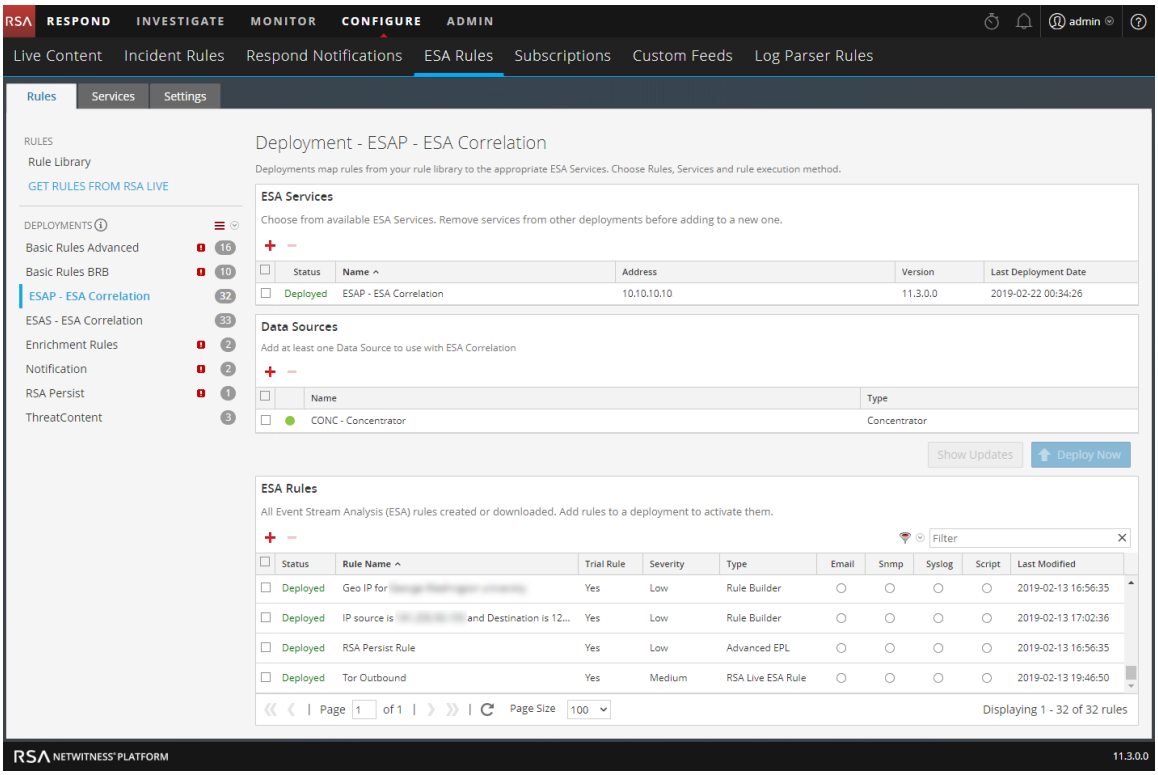
The following figure shows the 11.3 ESA rule deployments after the upgrade. Notice that the Basic Rules Advanced deployment from 10.6.6 does not have any ESA services, but it still contains the original 16 rules.



The 10.6.6 ESA rule deployments show update messages detailing the changes as shown in the following figure.



The following figure shows the ESA primary rule deployment after the upgrade (ESAP - ESA Correlation). Notice that the ESAP - ESA Correlation deployment has only the ESA primary ESA Correlation service.



In NetWitness 11.3 and later, you can only have one ESA Correlation service per deployment, but an ESA Correlation service can be in more than one ESA rule deployment.

Additional ESA Correlation Rules Procedures

This topic is a collection of individual procedures, which an Administrator may perform at any time and they are not required to complete the initial setup of ESA Correlation Rules.

Use this section when you are looking for instructions to perform a specific task after the initial setup of ESA.

- Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys
- Configure Advanced Settings for an ESA Correlation Service
  - Enable or Disable Sending ESA Rule Alerts to the Respond View
  - Enable ESA Correlation Service Debugging for All Rules
  - Configure Maximum Events per Alert for All Rules
  - Adjust Maximum Sessions for the ESA Data Source Filter
  - Configure Meta Keys as Arrays in ESA Correlation Rule Values
  - Remove Sensitive Meta Keys Globally from All Alerts for Data Privacy
- Configure Character Case for Advanced ESA Rules
- Deploy Endpoint Risk Scoring Rules on ESA
- Change Memory Threshold for ESA Rules
- Start, Stop, or Restart ESA Service
- View Audit Logs and Verify ESA Component Versions

Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys

Note: This procedure applies only to ESA Correlation Rules in NetWitness Platform 11.3.0.2 and later versions.

To support Endpoint, UEBA, and RSA Live content, a data change from single-value (string) to multi-value (string array) is required for several meta keys within the ESA Correlation service for 11.3 and later. Additional string meta keys are required within the ESA Correlation service for 11.3.0.2 and later.

If the meta keys used for your ESA rules are different from the required default multi-value meta keys, your ESA rules continue to work, but you should update your ESA rules to use the required meta keys as soon as possible to ensure that your rules continue to deploy properly.

Note: On a new installation of ESA on 11.3.0.2 and later, no ESA rule adjustments are necessary.

The ESA Correlation service has the following multi-valued (string array) and single-valued (string) parameters:

- **multi-valued** - Shows the string array meta keys currently used for your ESA rules.
  - For an upgrade to NetWitness Platform 11.3.0.2 and later, it shows the existing string array meta keys before the upgrade. (This parameter is equivalent to the Event Stream Analysis service ArrayFieldNames parameter in NetWitness Platform versions 11.2 and earlier.)
  - For a new installation of NetWitness Platform 11.3.0.2 and later, it contains all the required string array meta keys for the latest version.
- **single-valued** - Shows the string meta keys currently used for your ESA rules.
  - For an upgrade to NetWitness Platform 11.3.0.2 and later from versions prior to 11.3, this parameter value is empty.
  - For a new installation of NetWitness Platform 11.3.0.2 and later, it contains all the required string meta keys for the latest version.
- **default-multi-valued** - Shows the required string array meta keys for the latest version.
  - For a new installation of NetWitness Platform 11.3.0.2 and later, this parameter value is empty.
- **default-single-valued** - Shows the required string meta keys for the latest version.
  - For a new installation of NetWitness Platform 11.3.0.2 and later, this parameter value is empty.

Note: If you have the same value in the `single-valued` and `multi-valued` parameter fields, the `single-valued` meta key value takes precedence over the `multi-valued` meta key value.

To use the latest Endpoint, UEBA, and Live content rules, you must update the **multi-valued** parameter on the ESA Correlation service to include all of the meta keys in the **default-multi-valued** field. You must also update the **single-valued** parameter field to include all of the meta keys in the **default-single-valued** field. To do this, follow the Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules procedure.

Caution: Any changes that you make to the `multi-valued` parameter may cause an error when you deploy your existing rules. You can update the `multi-valued` parameter, resync your meta keys, and update the ESA rules at your convenience. You may want to add a couple meta keys at a time to reduce the number of reported errors.

Note: If you are using multiple ESA Correlation services, the `multi-valued` and `single-valued` parameters should be the same on each ESA Correlation service.

In NetWitness Platform 11.3.0.2 and later, ESA automatically adjusts the operator in the rule statement when there is a change from string to string array, but you still may need to make manual adjustments to adjust for the string array changes.

To change the string type meta keys to string array type meta keys manually, see Configure Meta Keys as Arrays in ESA Correlation Rule Values.

To use the latest Endpoint, UEBA, and Live content rules, the following default **multi-valued** meta keys are required on the ESA Correlation service in NetWitness Platform version 11.3 and later:

```
action , alert , alert.id , alias.host , alias.ip , alias.ipv6 , analysis.file
, analysis.service , analysis.session , boc , browserprint , cert.thumbprint ,
checksum , checksum.all , checksum.dst , checksum.src , client.all , content ,
context , context.all , context.dst , context.src , dir.path , dir.path.dst ,
dir.path.src , directory , directory.all , directory.dst , directory.src ,
email , email.dst , email.src , eoc , feed.category , feed.desc , feed.name ,
file.cat , file.cat.dst , file.cat.src , filename.dst , filename.src , filter
, function , host.all , host.dst , host.orig , host.src , host.state ,
inv.category , inv.context , ioc , ip.orig , ipv6.orig , netname , OS , param
, param.dst , param.src , registry.key , registry.value , risk , risk.info ,
risk.suspicious , risk.warning , threat.category , threat.desc , threat.source
, user.agent , username
```

The following default **single-valued** meta keys are also required on the ESA Correlation service in NetWitness Platform 11.3.0.2 and later:

```
accesses , context.target , file.attributes , logon.type.desc , packets
```

If you used any meta keys in the ESA rule notification templates from the Required String Array or String Meta Keys list, update the templates with the meta key changes. See "Configure Global Notification Templates" in the *System Configuration Guide*.

Note: Advanced EPL rules may get disabled and are not automatically updated so they must be fixed manually.

For additional troubleshooting information, see “Troubleshoot ESA” in the *Alerting with ESA Correlation Rules User Guide*.

Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules

To use the latest Endpoint, UEBA, and Live content rules, you must update the **multi-valued** parameter field on the ESA Correlation service to include all of the meta keys in the **default-multi-valued** field. You must also update the **single-valued** parameter field to include all of the meta keys in the **default-single-valued** field.

Caution: Any changes that you make to the `multi-valued` parameter may cause an error when you deploy your existing rules. You can update the `multi-valued` parameter, resync your meta keys, and update the ESA rules at your convenience. You may want to add a couple meta keys at a time to reduce the number of reported errors.

Note: If you see a warning message in the ESA Correlation server error logs for missing multi-valued meta keys, there is a difference between the `default-multi-valued` parameter and `multi-valued parameter` meta key values, and the new Endpoint, UEBA, and Live content rules will not work. The same is true for missing single-valued meta keys. Completing this procedure should fix the issue. For example warning messages, see Example ESA Correlation Server Warning Message for Missing Meta Keys.

1. After an upgrade to 11.3.0.2 or later, go to  **(Admin) > Services**, and in the Services view, select an ESA Correlation service and then select  **> View > Explore**.
2. In the Explore view node list for the ESA Correlation service, select **correlation > stream**.

3. Compare the **multi-valued** parameter meta keys with the required **default-multi-valued** meta keys. Copy and paste the missing string array meta keys from the **default-multi-valued** parameter to the **multi-valued** parameter. (You may want to copy only a couple meta keys at one time to reduce the number of reported errors).
4. Copy and paste the string meta keys from the **default-single-valued** parameter to the **single-valued** parameter.
5. If you have multiple ESA Correlation services, make the same meta key changes on each ESA Correlation service.
6. Apply the changes on the ESA Correlation services:
   a. Go to  **(Configure) > ESA Rules** and click the **Settings** tab.
   b. In the Meta Key References, click the Meta Re-Sync (Refresh) icon ().
7. If you are using any of the **default-multi-valued** or **default-single-valued** meta keys in your ESA Advanced rules, update the rule syntax. See also Adjust Custom ESA Rule Builder and ESA Advanced Rules.
8. If you used any meta keys in the ESA rule notification templates from the **default-multi-valued** parameter list, update the templates with the meta key changes. See "Configure Global Notification Templates" in the *System Configuration Guide*.
9. Deploy your ESA rule deployments.
10. Check your rules for error messages in the ESA Rules section of the ESA rule Deployment or check the ESA Correlation error logs for errors.
    - (This option is available in NetWitness verson 11.3.0.2 and later.) To access the error messages in the ESA rule deployment, go to  **(Configure) > ESA Rules > Rules** tab, select a deployment in the options panel on the left, and go to the **ESA Rules** section. If the ESA rule status shows "Disabled" or shows the  icon in the Status column, you need to determine the issue to fix the rule. If a disabled rule has an error message, it shows  in the Status field. You can hover over the rule to view the error message tooltip without going to the error log.
    - To access the ESA Correlation service logs, you can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`.

Adjust Custom ESA Rule Builder and ESA Advanced Rules

Update your ESA Rule Builder and ESA Advanced rules to work with the string and string array meta keys listed in the `default-multi-valued` and `default-single valued` parameter fields for the ESA Correlation service. You can add additional meta keys to the `multi-valued` and `single-valued` parameters.

For example, if you use ec.outcome as a single-valued meta key in your ESA rule as shown below:

```
@RSAAlert
SELECT * FROM Event((ec_outcome IN ( 'Success' )))
.win:time_length_batch(2 Minutes, 2)
```

```
HAVING COUNT(*) >= 2;
```

If you add ec.outcome to the **multi-valued** parameter field, you need to update your rule as shown below:

```
@RSAAlert
SELECT * FROM Event(( 'Success' = ANY( ec_outcome ) ))
.win:time_length_batch(2 Minutes, 2)
HAVING COUNT(*) >= 2;
```

For more information, see Configure Meta Keys as Arrays in ESA Correlation Rule Values.

Example ESA Correlation Server Warning Message for Missing Meta Keys

If you see a warning message in the ESA Correlation server error logs for missing multi-valued meta keys, there is a difference between the `default-multi-valued` parameter and `multi-valued parameter` meta key values, and the new Endpoint, UEBA, and Live content rules will not work. The same is true for missing single-valued meta keys. Completing the Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules procedure should fix the issue.

Multi-Valued Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[alert, alert_id,
browserprint, cert_thumbprint, checksum, checksum_all, checksum_dst,
checksum_src, client_all, content, context, context_all, context_dst,
context_src, dir_path, dir_path_dst, dir_path_src, directory,
directory_all, directory_dst, directory_src, email_dst, email_src,
feed_category, feed_desc, feed_name, file_cat, file_cat_dst, file_
cat_src, filename_dst, filename_src, filter, function, host_all,
host_dst, host_orig, host_src, host_state, ip_orig, ipv6_orig, OS,
param, param_dst, param_src, registry_key, registry_value, risk,
risk_info, risk_suspicious, risk_warning, threat_category, threat_
desc, threat_source, user_agent] are still MISSING from multi-valued
```

Single Value Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[accesses,
context_target, file_attributes, logon_type_desc, packets] are still
MISSING from single-valued
```

Configure Meta Keys on New ESA Hosts to Match Upgraded ESA Hosts in the Same NetWitness Platform Network

Note: If you are upgrading from version below 11.4.x.x, make sure that meta keys on new ESA hosts are updated to match upgraded ESA Hosts in the Same NetWitness Platform Network.

If you have one or more ESA hosts in a NetWitness Platform network, which were upgraded from a version before 11.3.0.2 to 11.6, and you add a new ESA host, you must configure the meta keys on the new ESA host to match the other ESA hosts. All ESA Correlation services on the same NetWitness Platform network must have the same Meta Key configurations.

1. For each ESA Correlation service on an upgraded ESA host and for the ESA Correlation service on the newly installed ESA host:
   a. Open a new tab, go to  **(Admin) > Services**, and in the Services view, select the ESA Correlation service and then select  **> View > Explore**.
   b. In the Explore view node list for the ESA Correlation service, select **correlation > stream**.
2. Ensure that the **`multi-valued`** and **`single-valued`** meta key values are the same on each of the upgraded ESA Correlation services.
3. Ensure that the **`multi-valued`** and **`single-valued`** meta key values on the newly installed ESA host are the same as those on the upgraded services.
4. To apply any changes on the ESA Correlation services, go to  **(Configure) > ESA Rules** and click the **Settings** tab. In the Meta Key References, click the **Meta Re-Sync (Refresh)** icon ( ).
5. If you updated the ESA Correlation services, redeploy the ESA rule deployments.

For more information, see Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys.

Configure Advanced Settings for an ESA Correlation Service

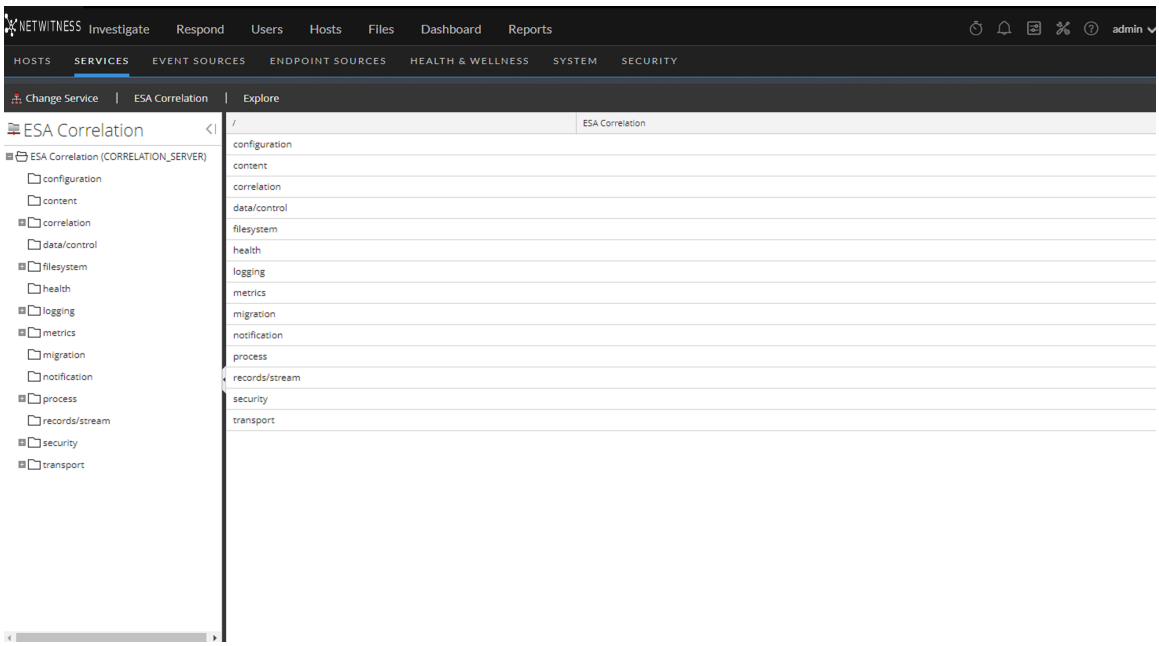
These procedures are optional and they apply only to ESA Correlation Rules.

In the Explore view for an ESA Correlation service, you can manage sending ESA rule alerts to the Respond view, turn on debugging for all rules, configure the events to preserve for rules with multiple events, and configure meta keys as string array values on ESA.

- Enable or Disable Sending ESA Rule Alerts to the Respond View
- Enable ESA Correlation Service Debugging for All Rules
- Configure Maximum Events per Alert for All Rules
- Adjust Maximum Sessions for the ESA Data Source Filter
- Configure Meta Keys as Arrays in ESA Correlation Rule Values
- Remove Sensitive Meta Keys Globally from All Alerts for Data Privacy

Access Advanced Settings for an ESA Correlation Service

1. Go to  **(Admin) > Services**.
   The Services view is displayed.
2. In Services view, select an ESA Correlation service and then select  **> View > Explore**.
   The Explore view is displayed.



Enable or Disable Sending ESA Rule Alerts to the Respond View

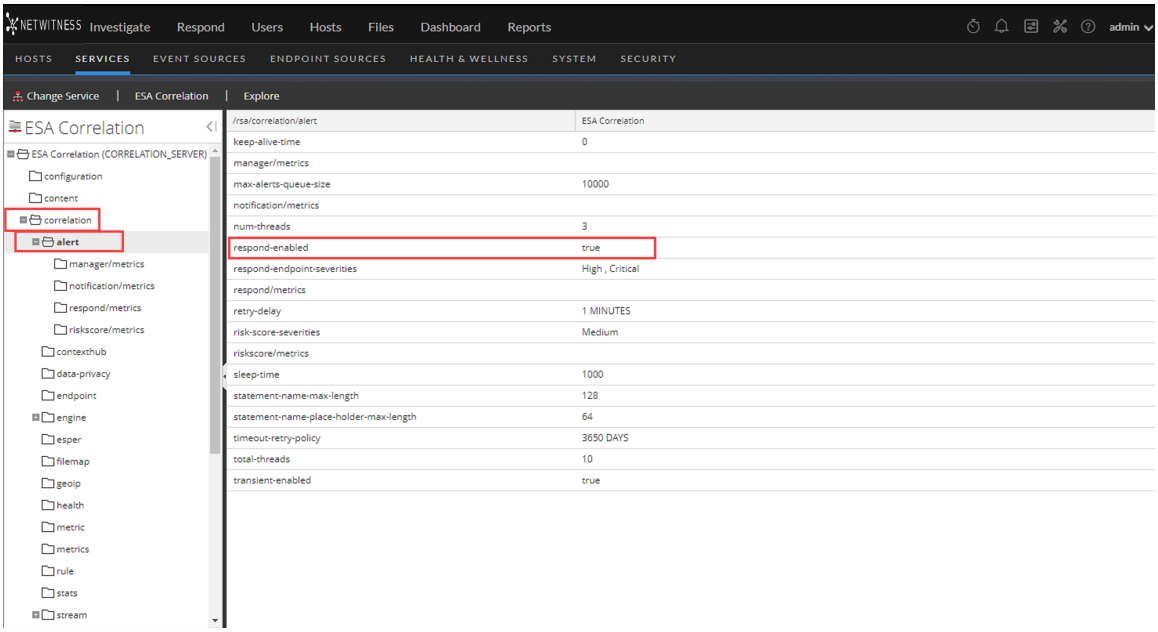
ESA gathers data, runs ESA Correlation rules against the data, captures events that meet rule criteria, and creates alerts for those captured events. You can view those alerts in the Respond view.

Before an ESA Correlation rule alert can go to the Respond view, both of the following settings must be enabled:

1. For all rules, the ESA Correlation service must have the `respond-enabled` parameter set to **true**. (The default is true.)
2. For an individual rule, the ESA Correlation rule must have the **Alert** option selected in the rule builder for that rule.

To enable or disable alert forwarding to the Respond view for ALL ESA Correlation rules:

1. In the Explore view node list for an ESA Correlation service, select **correlation > alert**.



2. To allow all ESA Correlation Rule alerts to go to the Respond view, set `respond-enabled` to **true**. Alerts for ESA rules that have the **Alert** option selected are visible in the Respond view.
3. To stop all ESA Correlation Rule alerts from going to the Respond view, set `respond-enabled` to **false**.
   ESA Correlation Rules do not go to the Respond view, even if you select the **Alert** option in the rule.
   The changes take effect immediately.

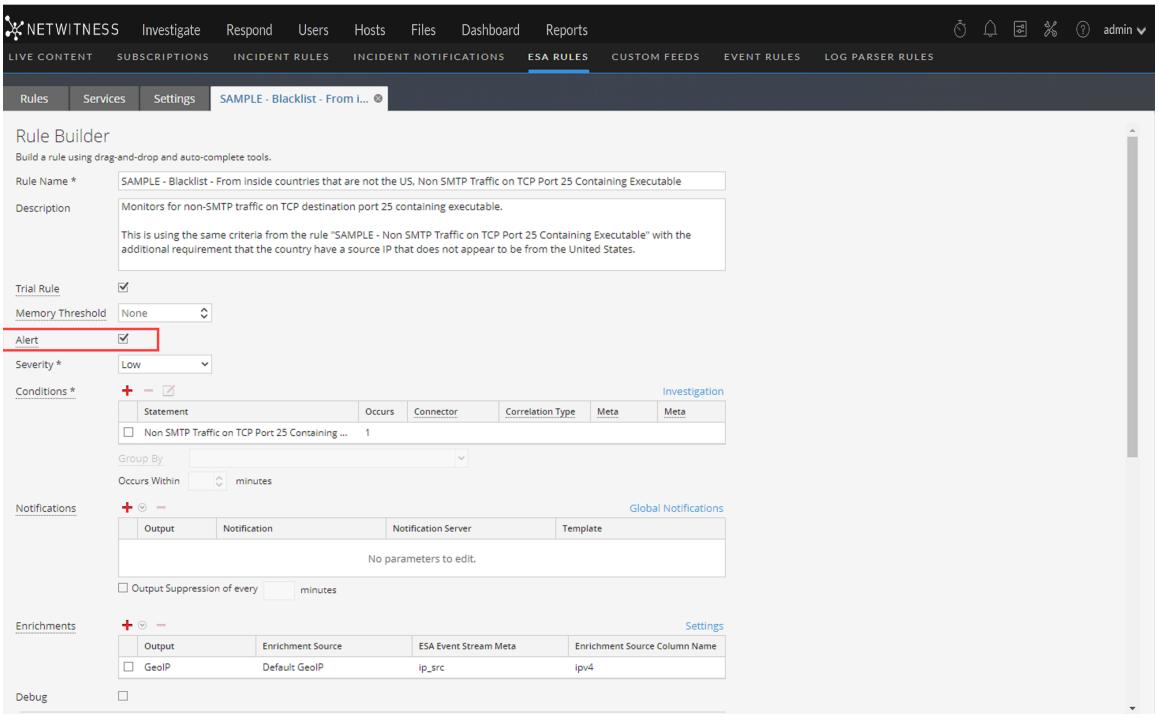
Note: The `respond-enabled` parameter is equivalent to the **Forward Alerts On Message Bus** option in the Event Stream Analysis service in NetWitness Platform version 11.2 and earlier.

To send or not send alerts to the Respond view for a single ESA Correlation rule:

Content experts managing the ESA Correlation rules can decide whether to send alerts to the Respond view for each rule.

1. Go to **(Configure) > ESA Rules > Rules** tab.
   The Rules tab is displayed.

2. In the **Rule Library**, select the rule you want to edit and click .
   Depending on the rule type, the respective rule tab is displayed.



   - To turn on Respond alerts for a rule, select the **Alert** checkbox.
   - To turn off Respond alerts for a rule, clear the **Alert** checkbox.

3. Click **Save**.

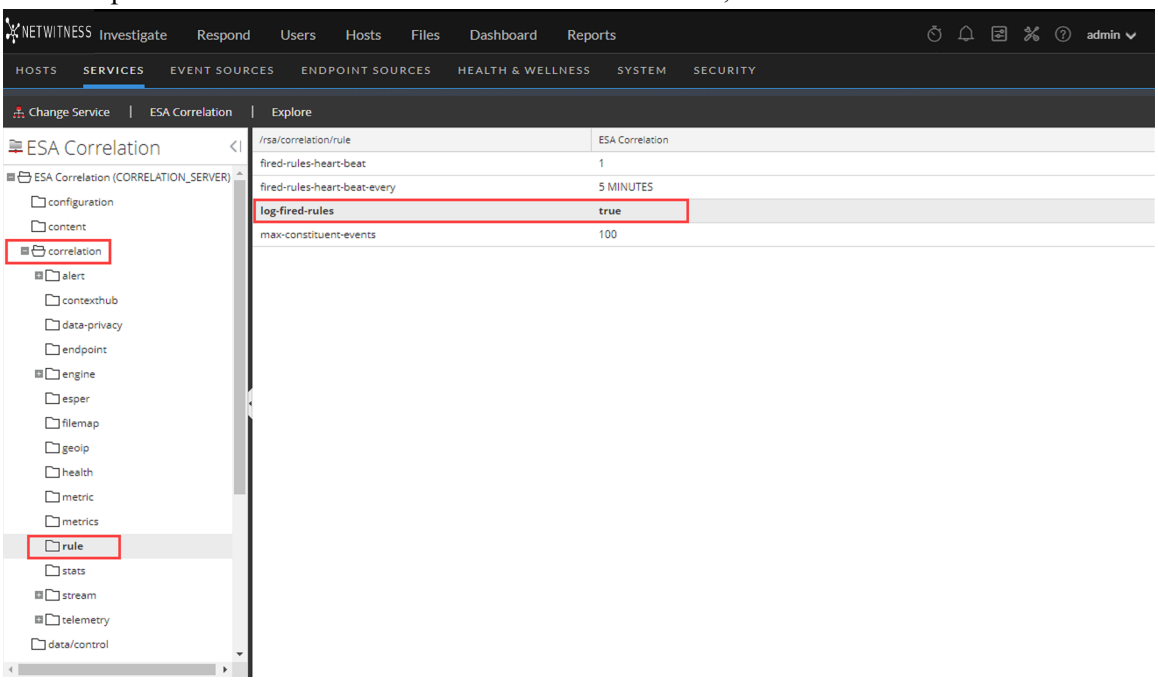
   For more information, see the *Alerting with ESA Correlation Rules User Guide*.

Enable ESA Correlation Service Debugging for All Rules

You can turn on debugging for all ESA rules to see if rules are creating (firing) alerts and data is being processed properly by the ESA Correlation service. This can also be helpful when writing or fixing global notification templates, such as syslog or email. You can see the actual content of an alert before sending the notification.

When you disable ESA Correlation service debugging for all rules, you can still turn on debugging for an individual rule at any time.

1. In the Explore view node list for an ESA Correlation service, select **correlation > rule**.

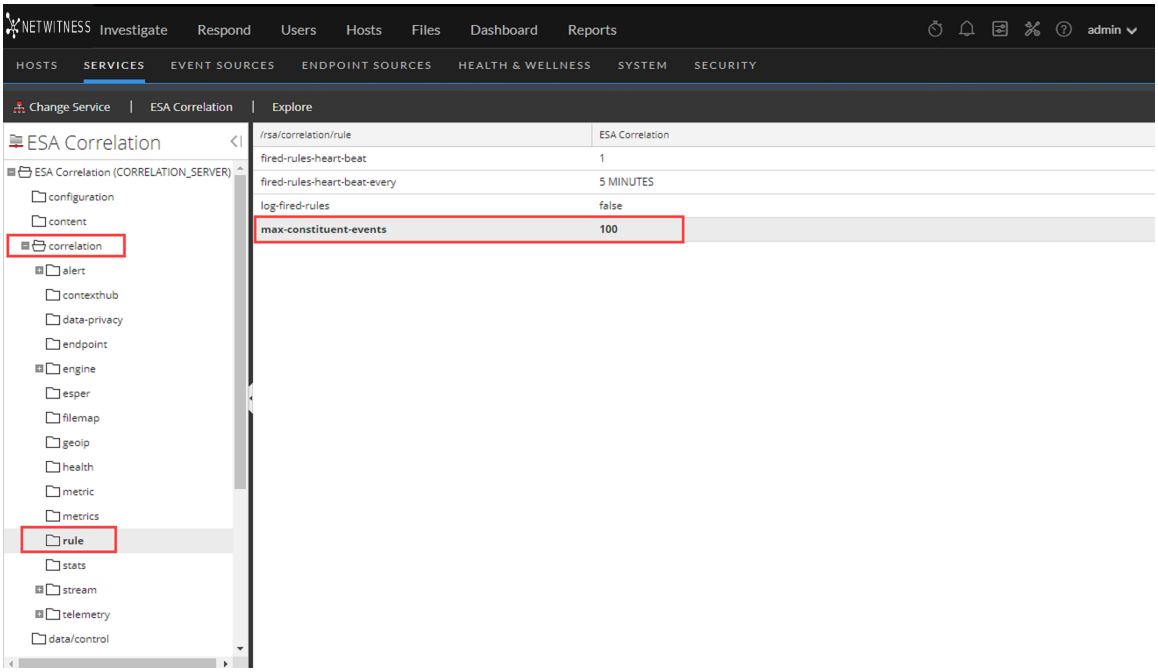


2. Set `log-fired-rules` to **true** to print alerts to the /var/log/netwitness/correlation-server/correlation-server.log for troubleshooting. This is the same as the Debug option in the rule builders for individual ESA rules except that this option enables debugging for all rules.
3. When you are ready to turn off debugging for all ESA rules, set `log-fired-rules` to **false**. The changes take effect immediately.

Note: The `log-fired-rules` parameter is equivalent to the **Debug Rules?** option in the Event Stream Analysis service in NetWitness Platform version 11.2 and earlier.

Configure Maximum Events per Alert for All Rules

1. In the Explore view node list for an ESA Correlation service, select **correlation > rule**.



2. For rules that contain multiple events, in `max-constituent-events`, enter how many of the associated events to preserve. For example, if a rule fires an alert with 200 associated events and this parameter is set to 100, only the first 100 are preserved by ESA, the rest are dropped. The default value is **100**.
   The changes take effect immediately.

Note: The `max-constituent-events` parameter is equivalent to the **Max Constituent Events** option in the Event Stream Analysis service in NetWitness Platform version 11.2 and earlier.

Adjust Maximum Sessions for the ESA Data Source Filter

Note: This procedure applies only to NetWitness Platform version 11.5 and later.

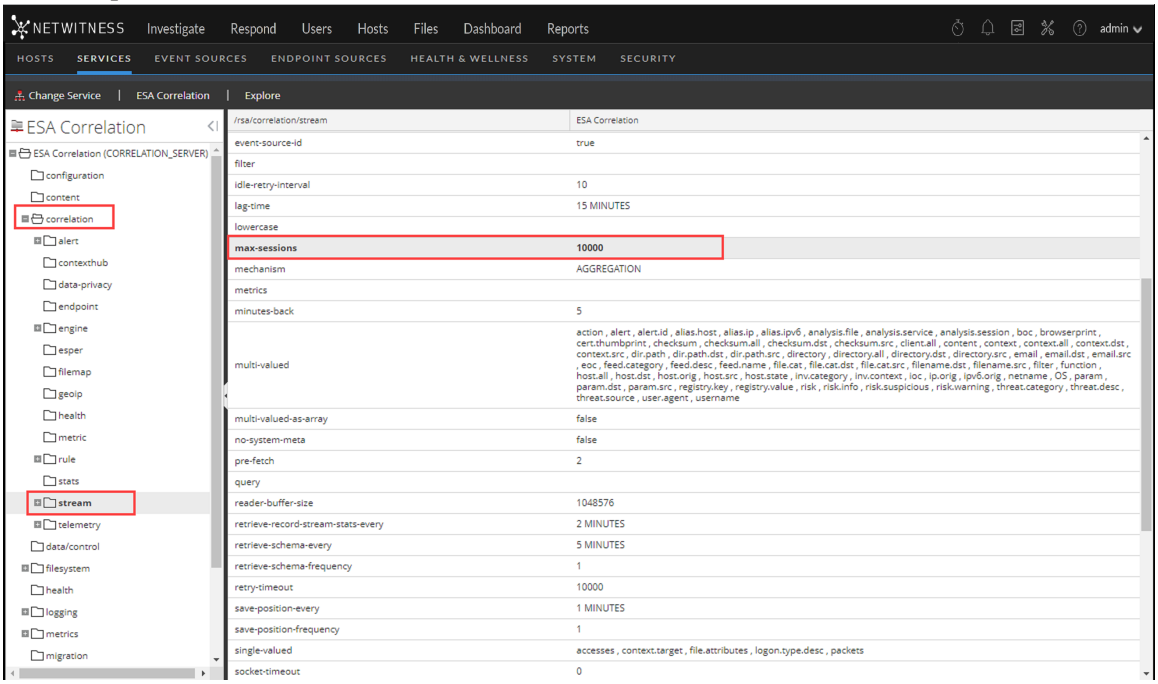
In NetWitness Platform 11.5 and later, you can add an optional data source filter to the data sources in your ESA rule deployments to improve performance. This allows your data sources to be filtered further so that only the data relevant to the deployment is forwarded to ESA. The filter is comprised of application rules, which are applied to the Decoders mapped to your selected data sources.

Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.

When filtering out a large portion of the traffic, you may see an "Invalid header size" error while communicating with Core services in the ESA Correlation log file. (You can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`). Lower the `max-sessions` parameter until you no longer see the error in the log. The more you filter out the traffic, the lower you should set the `max-sessions` parameter.

1. In the Explore view node list for an ESA Correlation service, select **correlation > stream**.



2. In `max-sessions`, lower the value until you no longer see the error in the ESA Correlation log file. The default value is **10000**.

3. Restart the ESA Correlation service. Go to  **(Admin) > Services**, select the ESA Correlation service, and then select  **> Restart**.

IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.

For more information, see "(Optional) Add a Data Source Filter" in the *Alerting with ESA Correlation Rules User Guide*.

Configure Meta Keys as Arrays in ESA Correlation Rule Values

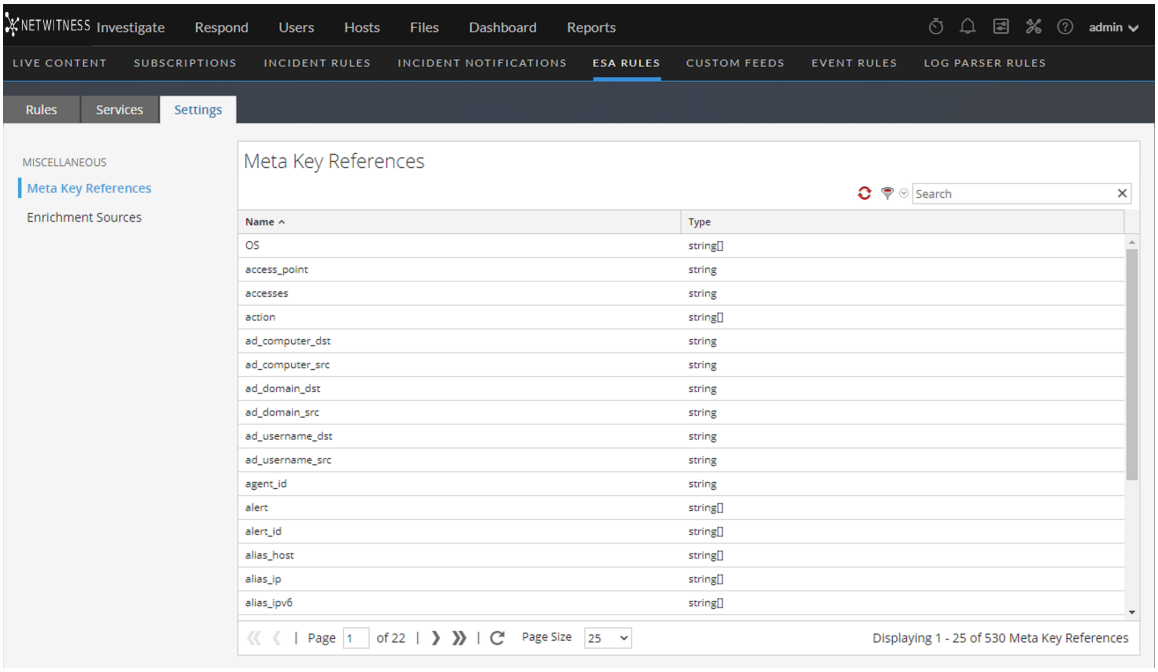
A common reason for an ESA rule to generate an error during deployment is because a meta key in the rule is a string array type, but it shows as a string type on ESA. To prevent or fix this issue, do the following:

- Determine if a Meta Key is a String Array Type on ESA
- Add the String Array Type Meta Key to ESA
- Verify that the String Array Type Meta Key is Configured Correctly on ESA

Caution: Changing string to string array type is not necessary for all fields. To support Endpoint, UEBA, and RSA Live content, specific string array (multi-value) and string (single-value) meta keys are required. See Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys.

Determine if a Meta Key is a String Array Type on ESA

1. Go to **(Configure) > ESA Rules** and click the **Settings** tab.



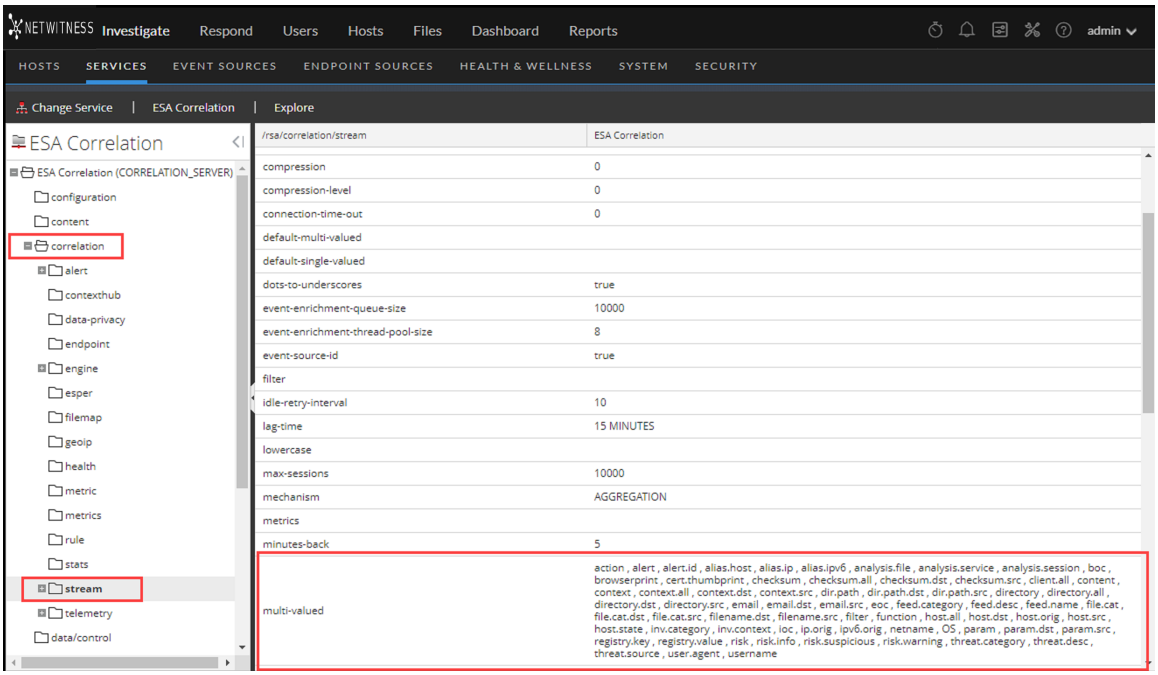
2. In the Meta Key References, for each meta key that is a string array type, locate the meta key in the Name field and then check the value.
   - If it shows `string[]`, it is configured as a string array type on ESA. This is fine.
   - If it shows `string` without the brackets, it is configured as a string type and you need to fix it on ESA. Go to Add the String Array Type Meta Key to ESA.

Caution: Changing string to string array type is not necessary for all fields. To support Endpoint, UEBA, and RSA Live content, specific string array (multi-value) and string (single-value) meta keys are required. See Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys.

Add the String Array Type Meta Key to ESA

Caution: Any changes that you make to the `multi-valued` parameter may cause an error when you deploy your existing rules. If you add a meta key to the `multi-valued` parameter field that you use in other ESA rules, ensure that those rules are using the string array syntax.

1. In the Explore view node list for an ESA Correlation service, select **correlation > stream**.
2. Add string array meta keys to the **`multi-valued`** list to allow them to be used as an array in ESA rules.



3. Verify the configuration on ESA. Go to Verify that the String Array Type Meta Key is Configured Correctly on ESA.

Note: The `multi-valued` parameter is equivalent to the `arrayFieldNames` parameter in the Event Stream Analysis service in NetWitness version 11.2 and earlier.

Verify that the String Array Type Meta Key is Configured Correctly on ESA

1. Go back to **(Configure) > ESA Rules** and click the **Settings** tab.
2. In the Meta Key References, click the Meta Re-Sync (Refresh) icon ( ).
3. Verify that the meta keys with a string array type show a value of `string[]`.

Required String Array Meta Keys on the ESA Correlation Service

To use the latest Endpoint, UEBA, and Live content rules, the following default **multi-valued** meta keys are required on the ESA Correlation service in NetWitness Platform version 11.3 and later:

```
action , alert , alert.id , alias.host , alias.ip , alias.ipv6 , analysis.file
, analysis.service , analysis.session , boc , browserprint , cert.thumbprint ,
checksum , checksum.all , checksum.dst , checksum.src , client.all , content ,
context , context.all , context.dst , context.src , dir.path , dir.path.dst ,
dir.path.src , directory , directory.all , directory.dst , directory.src ,
email , email.dst , email.src , eoc , feed.category , feed.desc , feed.name ,
file.cat , file.cat.dst , file.cat.src , filename.dst , filename.src , filter
, function , host.all , host.dst , host.orig , host.src , host.state ,
inv.category , inv.context , ioc , ip.orig , ipv6.orig , netname , OS , param
, param.dst , param.src , registry.key , registry.value , risk , risk.info ,
risk.suspicious , risk.warning , threat.category , threat.desc , threat.source
, user.agent , username
```

The following default **single-valued** meta keys are also required on the ESA Correlation service in NetWitness Platform 11.3.0.2 and later:

```
accesses , context.target , file.attributes , logon.type.desc , packets
```

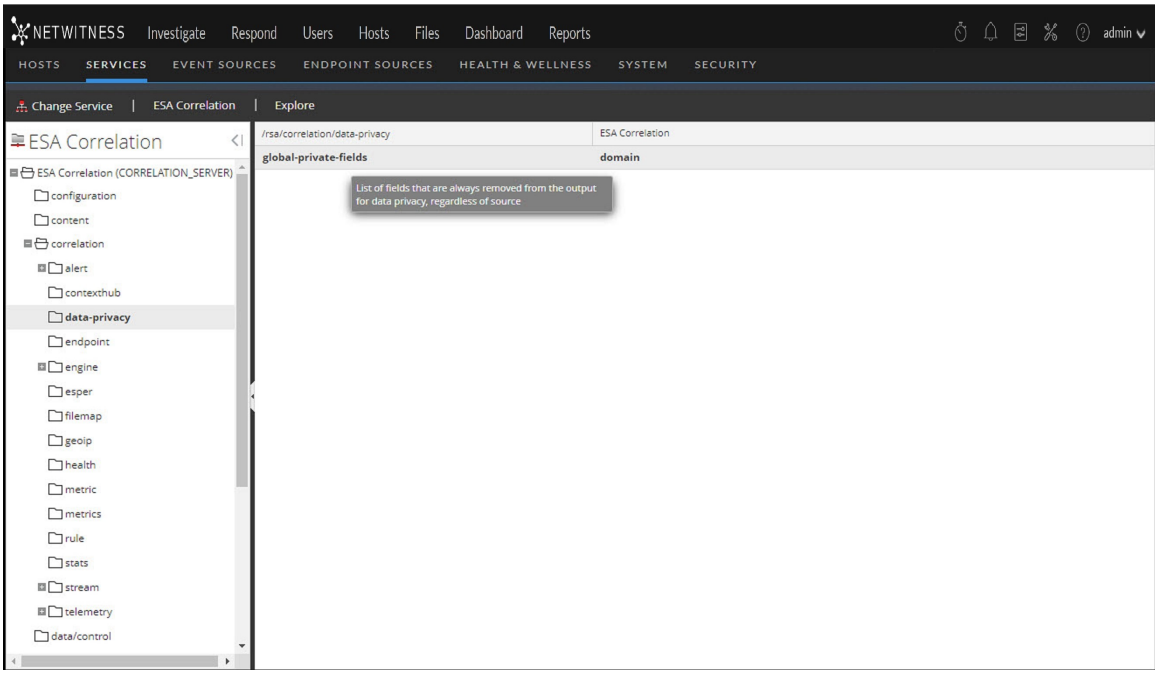
Note: Check the `default-multi-valued` and `default-single-valued` parameters on your ESA Correlation service for the latest required fields. For more information, see Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys.

Remove Sensitive Meta Keys Globally from All Alerts for Data Privacy

Note: This procedure applies only to ESA Correlation Rules in NetWitness Platform 11.4 and later versions.

For data privacy reasons, it may be necessary to remove some sensitive meta keys from the alert output globally, regardless of the data source. In the ESA Correlation service, you can set the **`global-private-fields`** parameter to remove the meta keys from all alert output.

1. Go to  **(Admin) > Services**, and in the Services view, select an ESA Correlation service and then select  **> View > Explore**.
2. In the Explore view node list for the ESA Correlation service, select **correlation > data-privacy**.
3. In the **`global-private-fields`** parameter, add the sensitive meta keys that you want removed from all alerts.



The changes are effective immediately.

For more information, see "How ESA Handles Sensitive Data" in the *Alerting with ESA Correlation Rules Configuration Guide*. For more information on the strategy and benefits of obfuscating data, see the *Data Privacy Management Guide*.

Configure Character Case for Advanced ESA Rules

Note: This procedure applies only to ESA Correlation Rules in NetWitness Platform 11.3.0.2 and later versions, however, it is not supported in version 11.3.1.0.

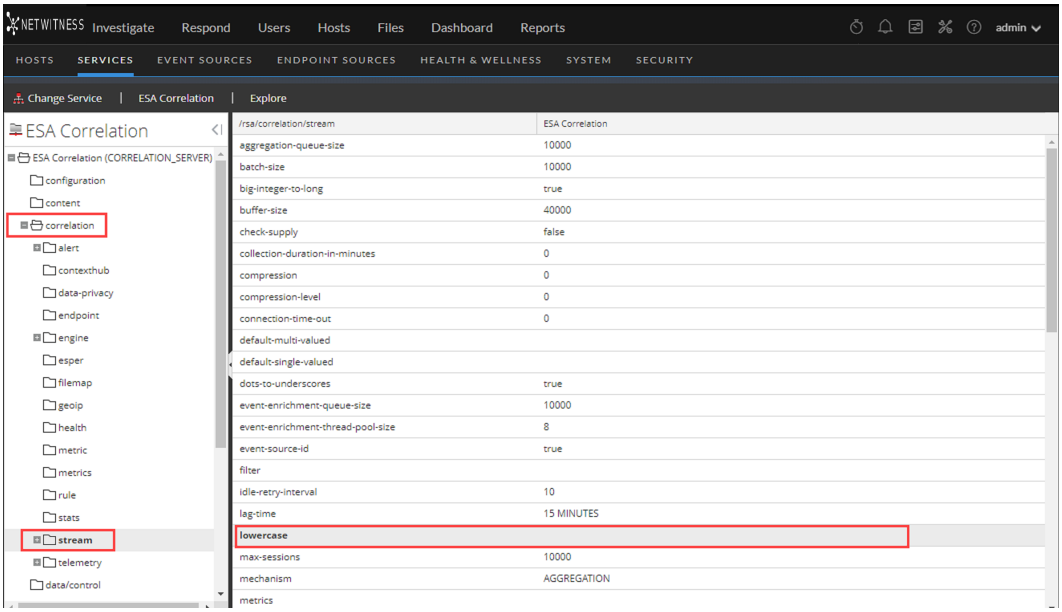
Advanced Event Processing Language (EPL) rules require correct character case, but in the Investigate Navigate view all characters are converted to lowercase. However, the meta keys may not be lowercase despite appearances in the Investigate Navigate view. To ensure you are using the correct case, you can use the `toLowerCase`() function. However, care should be taken to only add the case-insensitive `toLowerCase()` function on string and string array meta keys as needed. The `toLowerCase()` function can cause significant performance decreases. Consider checking the Investigate Events view or the Event Analysis view to see the real character case for meta fields and avoid unnecessary usage of the function. For more information, see "Event Process Language (EPL)" in the *Alerting with ESA Correlation Rules User Guide*.

You can optimize your rule performance by identifying the meta keys used most often in your environment. Instead of using the `toLowerCase()` function with the original meta key, replace the meta key throughout the rule with `<meta.key>_lower`. You can also use the special case-insensitive meta keys in your Rule Builder rules.
For example, you can configure ESA Correlation to use `filename_lower` (which is case insensitive) instead of using the original `filename` meta key. In your rule, replace `filename` with `filename_lower`.

To configure special case-insensitive meta keys to use in your ESA rules:

1. Go to  **(Configure) > ESA Rules > Rules** tab. In your ESA rule deployments, identify any ESA rules using the `toLowerCase()` function more than ten times for a particular string or string array meta key. Keep track of these ESA rules and meta keys.
2. Go to  **(Admin) > Services**, select an ESA Correlation service, and then select  **> View > Explore**.

3. In the Explore view node list, select **correlation > stream**.
   Notice that there is a **`lowercase`** parameter with empty values.



4. Update the **`lowercase`** parameter with the string or string array meta keys identified in step 1 using a comma separated list, for example: `protocol,alias.host,action,alert`

   Note: String and string array are the only data types supported for the ESA Correlation service **`lowercase`** parameter.

   Use NWDB format (decimal), NOT Esper format (underscore). Do not press **Enter** to commit or it will put in a return. Instead, click another parameter.

5. After you add all of the meta keys, validate the meta keys on ESA.

   a. Go to **(Configure) > ESA Rules > Settings tab > Meta Key References** and click the Meta Re-Sync (Refresh) icon ( ).

   b. Search for **`_lower`** or **`<meta key>_lower`**, for example: `protocol_lower`.

   - The meta keys with a string array type should show a value of `string[]`.
   - The meta keys with a string type should show a value of `string` (without the brackets).

6. Update all of your ESA rules that use `.toLowerCase` meta keys and replace them with `<meta key>_lower` (Example: filename_lower IN ('svchost.exe'))
7. Deploy the ESA rule deployment again.

Note: If you remove a meta key from the `lowercase` parameter list and re-sync the meta key references, you also need to update the rules that use the corresponding lowercase meta key (`<meta.key>_lower`).

Deploy Endpoint Risk Scoring Rules on ESA

Note: The Information in this topic applies to NetWitness 11.3 and later.

Endpoint Risk Scoring Rules only apply to NetWitness Endpoint.

The ESA Correlation service processes and deploys endpoint risk scoring rules. These rules generate alerts that are used in risk scoring calculations to identify suspicious files and hosts. To turn on Risk Scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. An *Endpoint Risk Scoring Rules Bundle* comes with NetWitness along with the sample ESA rules. The Endpoint Risk Scoring Bundle contains approximately 400 rules. You add this rule bundle to an ESA Rule Deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) during ESA Rule Deployment.

For complete information on configuring NetWitness Endpoint, see the *NetWitness Endpoint Configuration Guide*. For more information about ESA rule deployments, see "Deploy Rules to Run on ESA" in the *Alerting with ESA Correlation Rules User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Important Considerations when Deploying the Endpoint Risk Scoring Rules Bundle

- If you add the Endpoint Risk Scoring Bundle to an ESA rule deployment, the deployment should have data sources with endpoint data.
- An ESA rule deployment can have only one ESA Correlation service. You can, however, use the same ESA Correlation service in multiple deployments.
- If you have two ESA Correlation services with the same endpoint data sources, deploy the Endpoint Risk Scoring Rules Bundle on only one of them.

Deploy the Endpoint Risk Scoring Rules Bundle on ESA

Caution: Before you deploy the Endpoint risk scoring rules, update your meta keys. See Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys.

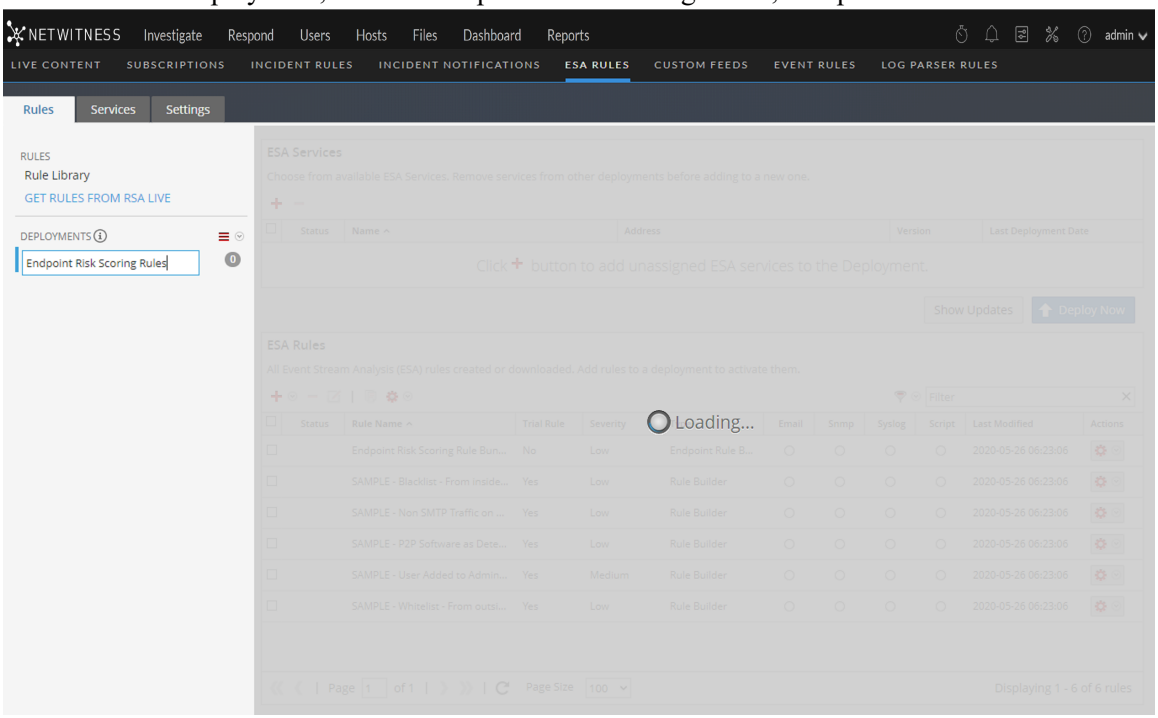
When you deploy the Endpoint Risk Scoring Rules Bundle in an ESA rule deployment, the ESA Correlation service gathers endpoint data in your network and runs endpoint risk scoring rules against the data. The goal is to capture events that match rule criteria, then generate alerts for the captured events.

The following procedure shows how to create an ESA rule deployment with the Endpoint Risk Scoring Rules Bundle and deploy it. If you already have an ESA rule deployment with endpoint data sources, you can add the Endpoint Risk Scoring Rules Bundle to the existing deployment.

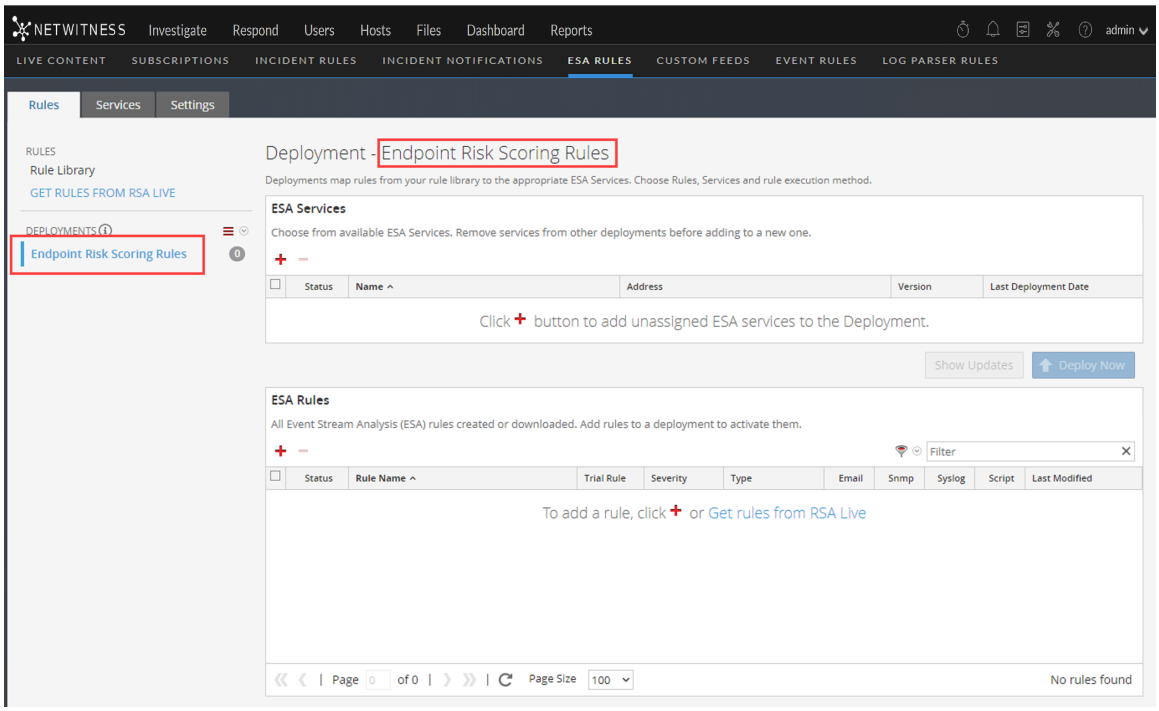
To create and deploy an ESA rule deployment with the Endpoint Risk Scoring Bundle:

1. Go to **(Configure) > ESA Rules > Rules tab**.

2. In the options panel on the left, next to Deployments, select **> Add** to add a deployment, type a name for the deployment, such as Endpoint Risk Scoring Rules, and press **Enter**.



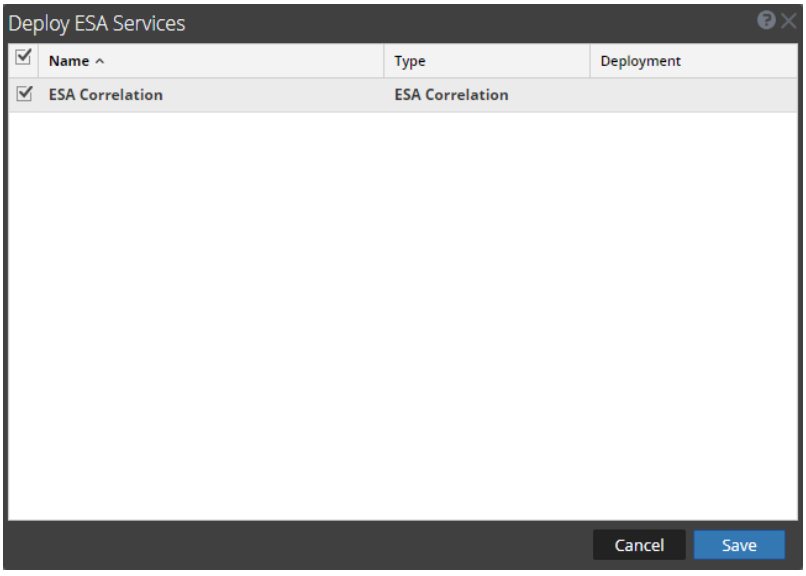
The deployment is added and the Deployment view is displayed on the right.



The deployment name that you choose also appears on a deployment tab in the **(Configure) > ESA Rules > Services** tab, where you can view deployment details and statistics.

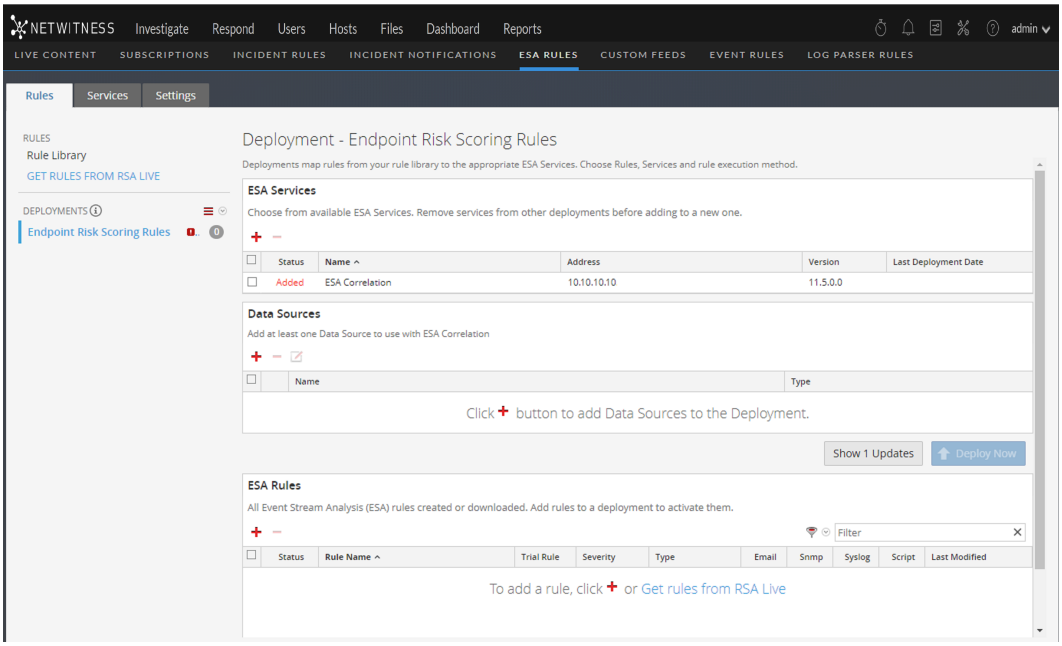
3. Add an ESA Correlation service:

   a. In the **Deployment** view **ESA Services** section, click + .
   The Deploy ESA Services dialog lists each configured ESA Correlation service.



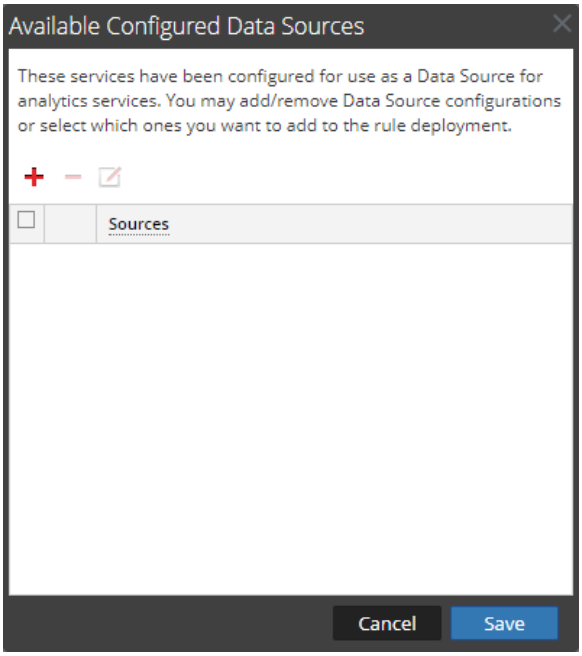
   b. Select an ESA Correlation service and click **Save**.
   The Deployment view is displayed. The ESA Correlation service is listed in the **ESA Services**

section with an **Added** status.



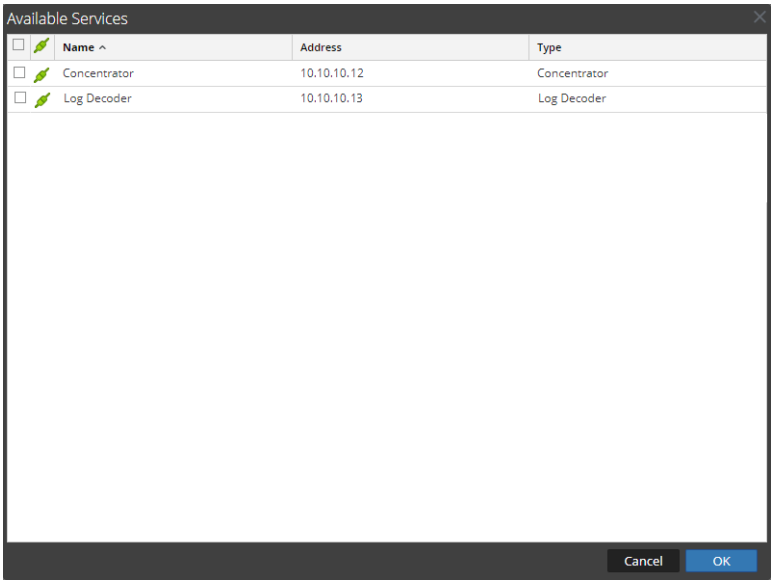
4. Add one or more data sources with endpoint data:

   a. In the **Deployment** view **Data Sources** section, click + .
   The **Available Configured Data Sources** dialog lists the services that have been configured for use as a data source.



   b. To add a data source configuration, click + .
   The **Available Services** dialog lists the available data sources from the (Admin) > Services

view.

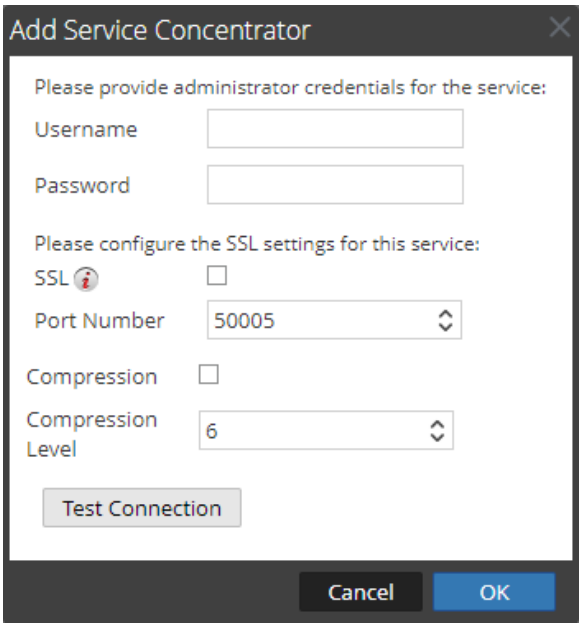


If an endpoint data source (service) that you are looking for is not in the list, see the *Hosts and Services Getting Started Guide* for instructions on how to install a service on a host.

Note: You can add a Log Decoder as a data source for ESA, but it is better to add a Concentrator to take advantage of undivided aggregation as the Decoder may have other processes aggregating from it.

c. In the **Available Services** dialog, select an endpoint data source, such as a Concentrator, and click **OK**.

d. In the **Add Service** dialog, type the Administrator username and password for the endpoint data source.



e. To enable the SSL or Compression options, select the corresponding checkboxes.

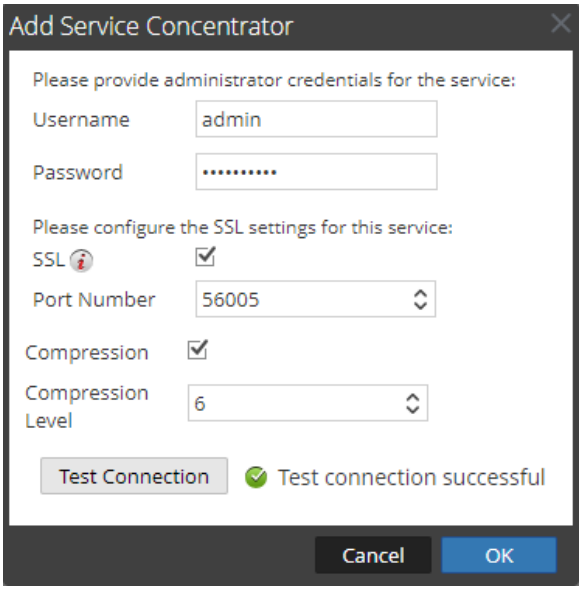
f. (Optional) You have the option to adjust the Compression Level for Concentrators on ESA in NetWitness 11.3 and later. To enable compression, select the **Compression** checkbox. You can set the **Compression Level** for a Concentrator from 0-9:

- Compression Level = **0** (If compression is enabled, it allows Core Services to control the amount of compression.)
- Compression Level = **1** (It uses the lowest amount of compression and has the highest performance.)
- Compression Level = **9** (It uses the highest amount of compression and has the worst performance.)

Somewhere in the middle between 1 and 9 is usually the best setting, which is what you get when you select a compression level of 0. For more detailed information, see the *Core Database Tuning Guide*.

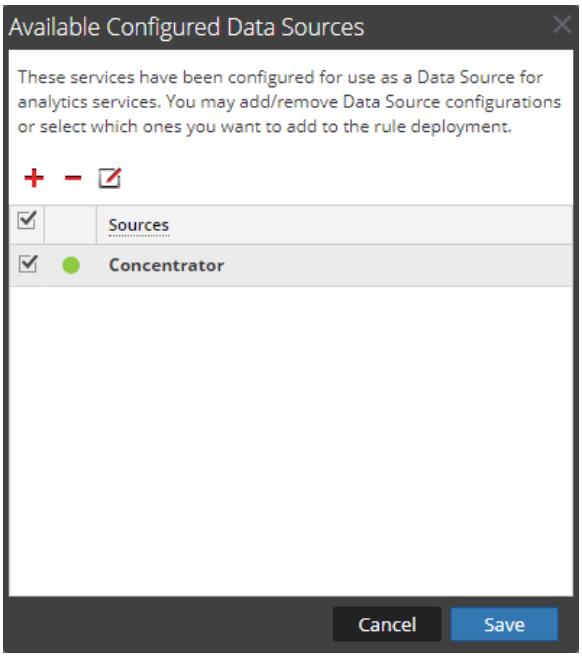
Note: If you make any ESA service, data source, or ESA rule changes to an ESA rule deployment, you need to redeploy the deployment. For example, if you change the configuration of a data source in an ESA rule deployment, you must redeploy all the ESA rule deployments that contain that data source.

g. Click **Test Connection** to make sure that it can communicate with the ESA Correlation service.



h. Click **OK**.
After you configure your endpoint data sources and they appear in the **Available Configured Data Sources** dialog, you can use them for your deployment.

i. In the **Available Configured Data Sources** dialog, select at least one endpoint data source to use for the deployment.



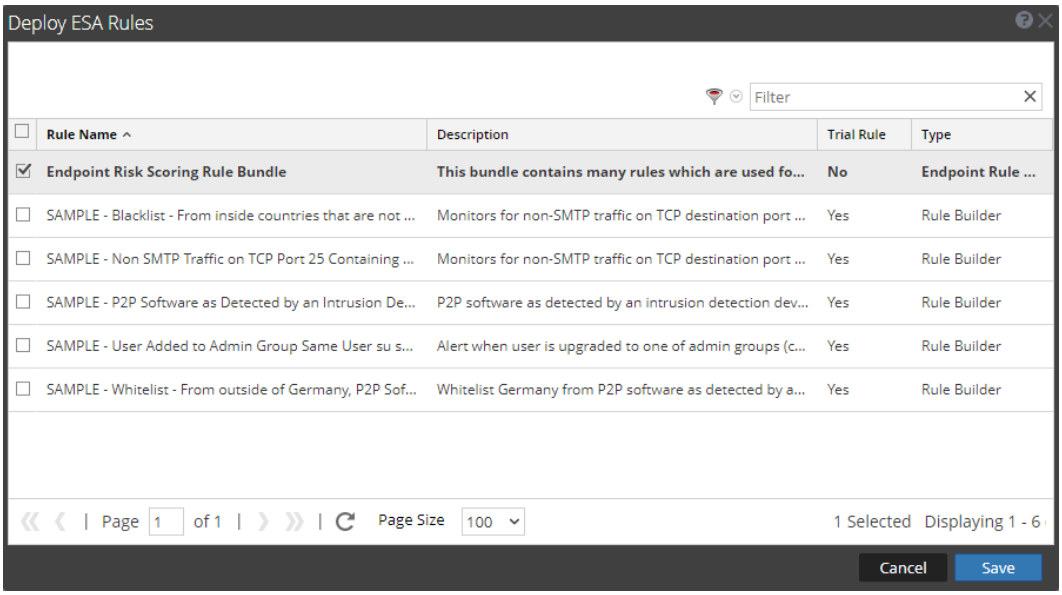
A solid colored green circle indicates a running service and a white circle indicates a stopped service.

j. Click **Save**.
In the Deployment view **Data Sources** section, the selected data sources are added to the deployment. The **Deploy Now** button activates after a service, data source, and rules are added to an ESA rule deployment.

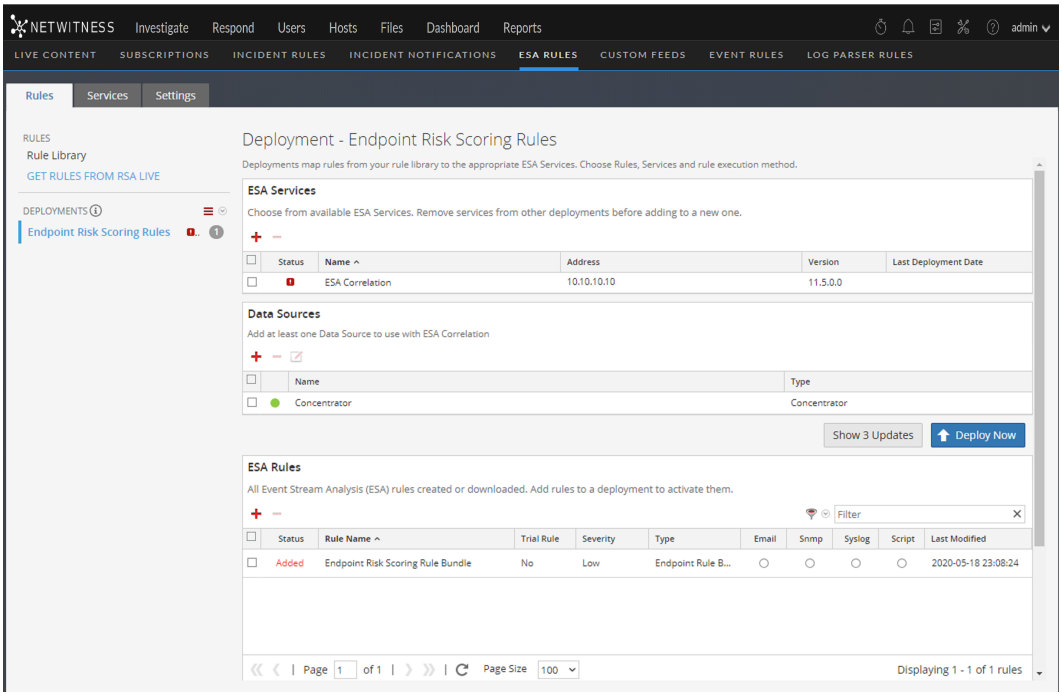


5. Add the Endpoint Risk Scoring Rules Bundle:

a. In the **Deployment** view ESA Rules section, click .
The Deploy ESA Rules dialog is displayed and shows each rule in your Rule Library.

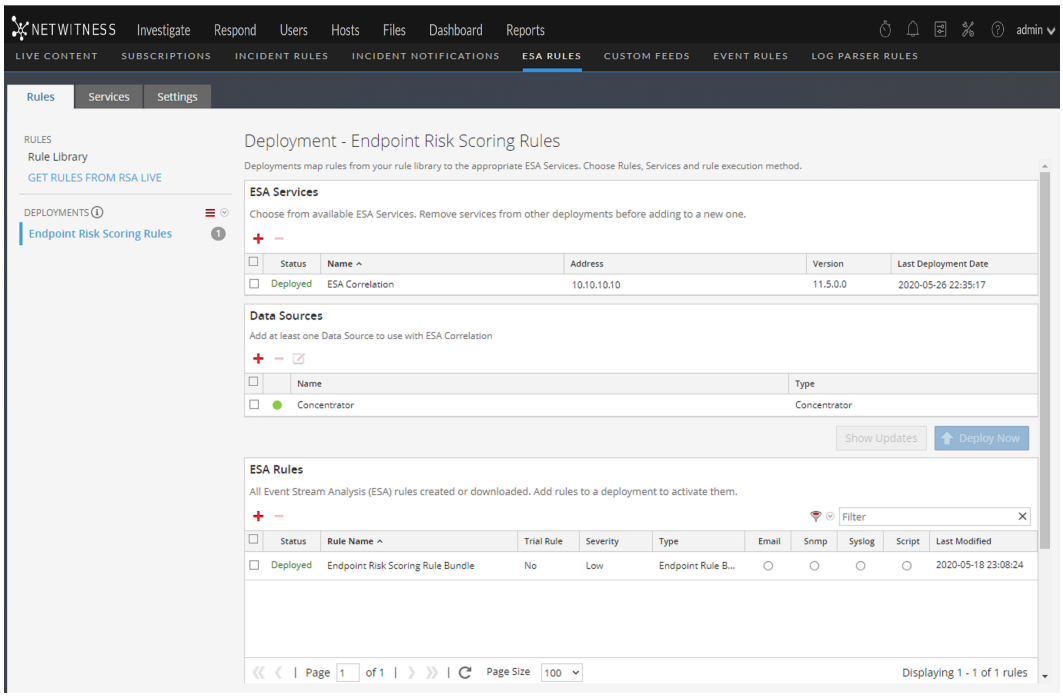


b. Select the Endpoint Risk Scoring Rule Bundle and click **Save**.
The Deployment view is displayed and the **Deploy Now** button is enabled.



The Endpoint Risk Scoring Rules Bundle is listed in the **ESA Services** section with an **Added** status.

6. Click **Deploy Now**.
The ESA Correlation service runs the rules in the Endpoint Risk Scoring Rules Bundle. The status of the bundle changes to **Deployed**.



You can now view information and statistics on the **(Configure) > ESA Rules > Services** tab. See View the Status of the Endpoint Risk Scoring Rules Deployment.

Change the Endpoint Risk Scoring Rule Bundle in a Deployment

You cannot edit or duplicate the Endpoint Risk Scoring Rules Bundle. After the bundle is deployed, you can enable and disable individual rules within the bundle. See Disable or Enable Individual Endpoint Risk Scoring Rules.

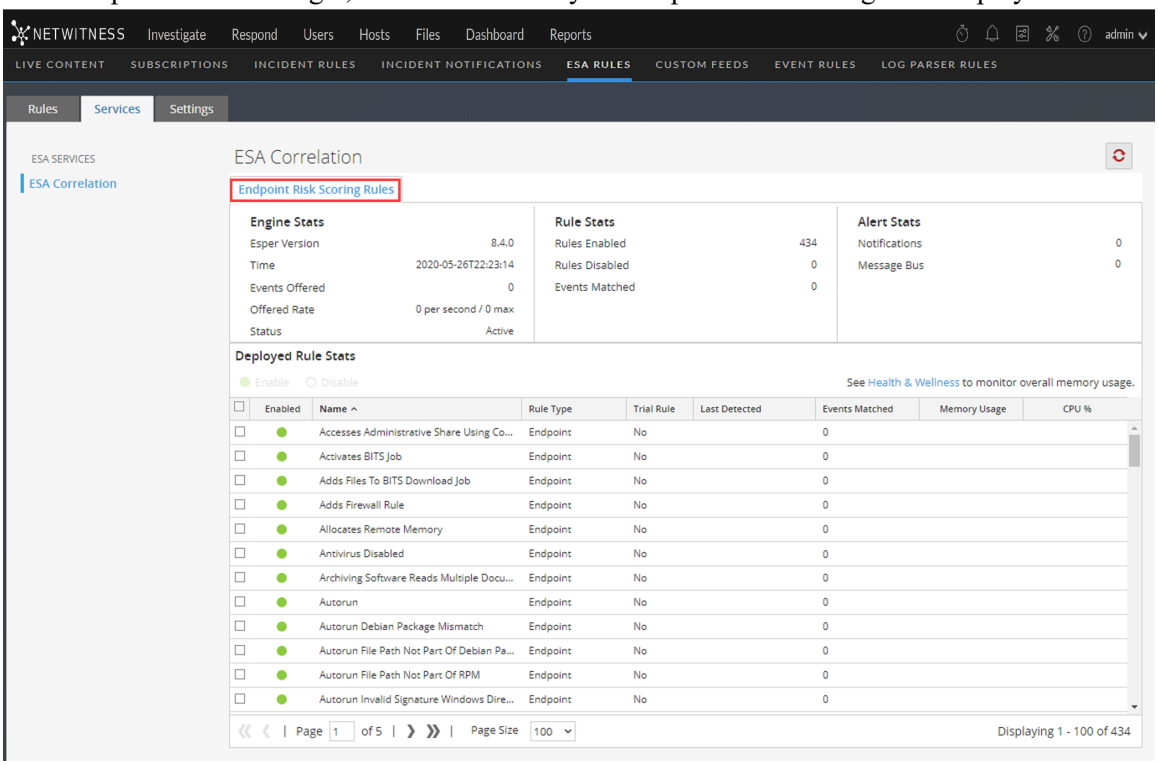
When you make changes to the ESA Rule Deployment containing the Endpoint Risk Scoring Rules Bundle, such as changing the endpoint data sources or changing compression levels, you must redeploy it for the changes to take effect. To redeploy, click the **Deploy Now** button for that deployment.

Caution: Deleting an ESA Rule Deployment with an Endpoint Risk Scoring Rule Bundle stops the Risk Scoring alerts that are used in risk scoring calculations to identify suspicious files and hosts.

For more information about changing ESA rule deployments, see "Additional ESA Rule Deployment Procedures" in the *Alerting with ESA Correlation Rules User Guide*.

View the Status of the Endpoint Risk Scoring Rules Deployment

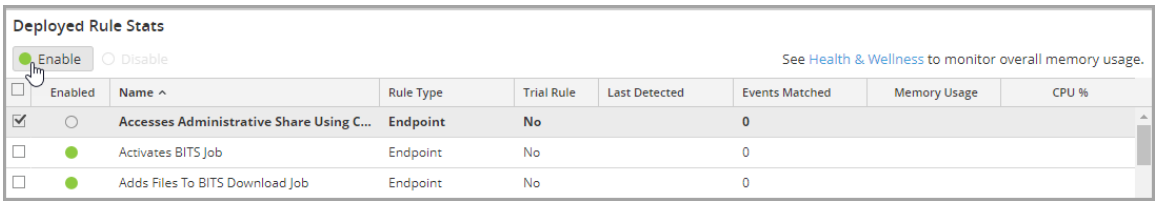
1. Go to the ESA Rules Services tab (**(Configure) > ESA Rules > Services**).
2. In the options panel on the left, select your ESA Correlation service.
   Your deployment name shows on a tab to the right, for example, Endpoint Risk Scoring Rules. If you see multiple tabs on the right, select the tab for your endpoint risk scoring rules deployment.



3. In the **Engine Stats**, **Rules Stats** and **Alert Status** sections, look at the statistics related to the deployment, such as Rules Enabled, Rules Disabled, and Events Matched, which show the total numbers for the deployment.
4. In the **Deployed Rules Stats** section, look at the following details for each Endpoint Risk Scoring Rule:
   - **Enable**: Indicates the enabled status. A green circle icon indicates that the rule is enabled. A white circle icon indicates that the rule is disabled.
   - **Name**: Shows the name of the rule.
   - **Rule Type**: Endpoint indicates a rule from the Endpoint Risk Scoring Bundle and Esper indicates Esper-specific rules, such as Rule Builder and Advanced EPL rules.
   - **Last Detected**: Shows the last time an alert was triggered for the rule.
   - **Events Matched**: Shows the total number of events that matched the rule.

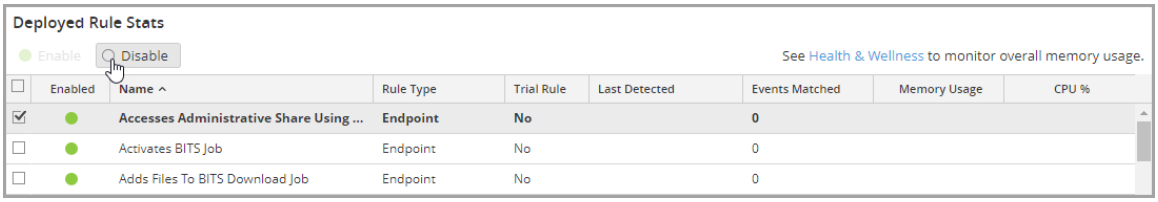
Disable or Enable Individual Endpoint Risk Scoring Rules

1. Go to the ESA Rules Services tab (**(Configure) > ESA Rules > Services**).
2. In the options panel on the left, select your ESA Correlation service.
   Your deployment name shows on a tab to the right, for example, Endpoint Risk Scoring Rules. If you see multiple tabs on the right, select the tab for your endpoint risk scoring rules deployment.
3. In the **Deployed Rules Stats** section, do one of the following:
   - To enable rules, select the rules that you want to enable in the rules list and click the **Enable** button above the list.



   The selected rules are enabled and a message shows that the rules enabled successfully.

   - To disable rules, select the rules that you want to disable in the rules list and click the **Disable** button above the list.



   The selected rules are disabled and a message shows that the rules disabled successfully.

Change Memory Threshold for ESA Rules

The following procedures pertain to setting memory thresholds for ESA rules to prevent them from using excessive memory.

Change Memory Threshold for All Trial Rules

This procedure is optional and applies only to ESA Correlation Rules.

Administrators can increase or decrease the memory threshold for trial rules. Threshold refers to the ESA memory usage, which includes ESA base memory, trial rules, and non-trial rules. When the threshold is exceeded, all deployed trial rules on an ESA service are disabled.

You use trial rules to see if a rule runs efficiently and does not use excessive memory, which can impact performance or force the service to shut down.

By default, the memory threshold is 90, which is the percentage of Java Virtual Memory (JVM).

- The memory threshold is per ESA, not per rule.
- When the memory threshold is exceeded, all trial rules running on the ESA are automatically disabled.
- The ESA configuration has the following parameters for trial rules:
  - `fatal-percentage`: If memory rises above this percentage, ESA disables trial rules. For example, if fatal-percentage is set to 90, when memory rises above 90 percent, ESA disables trial rules.
  - `check-every`: This parameter determines how often ESA checks the fatal-percentage to disable trial rules.

For more information, see "Work with Trial Rules" in the *Alerting with ESA Correlation Rules User Guide*.

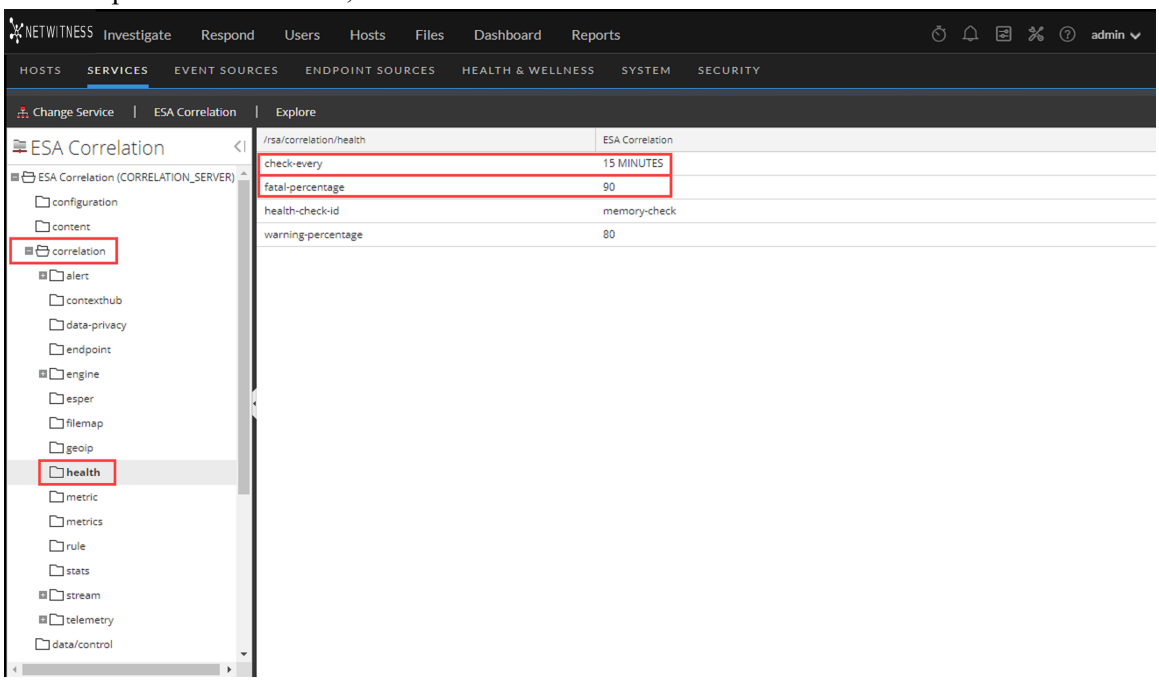
Prerequisites

A role with administrative privileges must be assigned to you.

To change memory threshold for trial rules:

1. Log on to NetWitness as admin.
2. Go to  **(Admin) > Services**.
3. Select the ESA Correlation service and then select  **> View > Explore**.

4. In the Explore view node list, select **correlation > health**.



5. In the right panel, in `fatal-percentage`, type a percentage of JVM that trial rules on the ESA cannot exceed.
   The new memory threshold takes effect immediately.
6. If necessary, you can also adjust the `check-every` parameter, which determines how often ESA checks the fatal-percentage to disable trial rules. By default, ESA checks the fatal-percentage every 15 minutes.

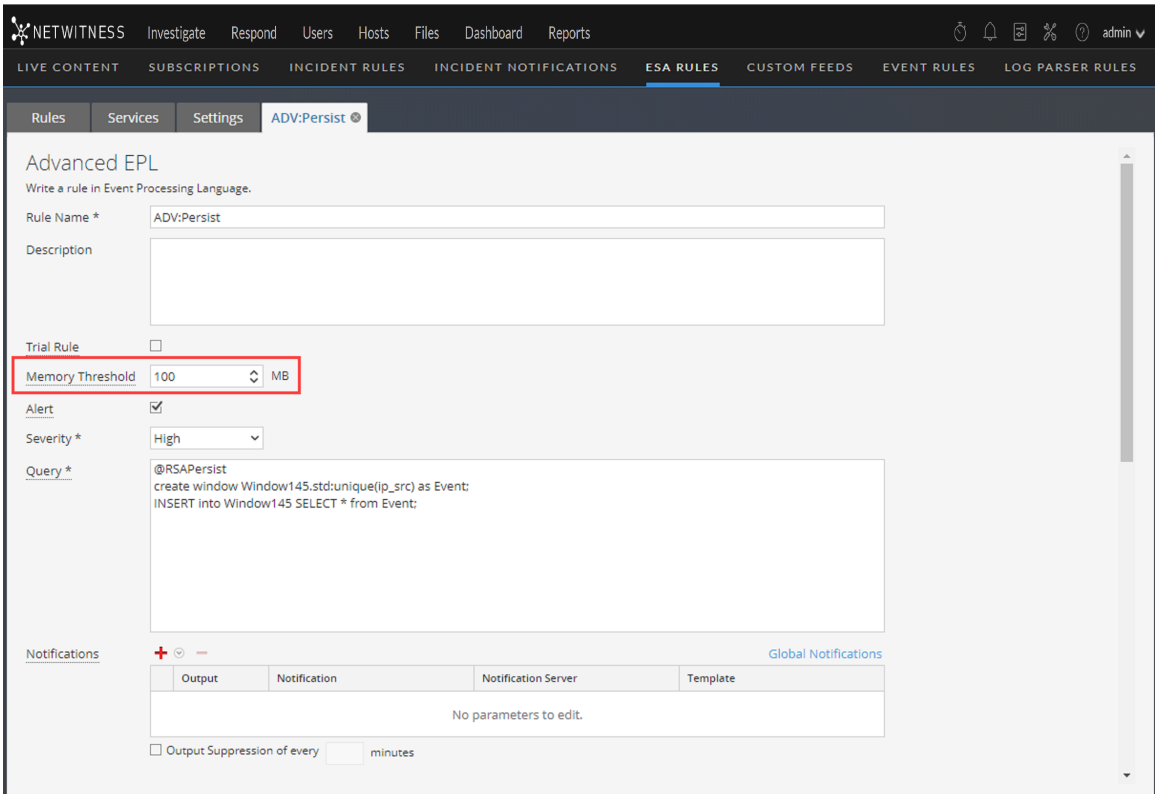
Change Memory Threshold for Individual Trial Rules and Non-Trial Rules

Note: This option is available in NetWitness Platform version 11.5 and later.

In addition to setting a memory threshold for all trial rules, you can set a memory threshold individually for both trial rules and non-trial rules. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set. You should configure a memory threshold for rules that use memory, such as a rule that contains windows or pattern matching. If the configured memory threshold is exceeded, the rule gets disabled individually and an error is displayed for that rule on the (Configure) > ESA Rules > Services tab.

1. Go to **(Configure) > ESA Rules >Rules** tab.
2. In the **Rule Library**, select the rule you want to configure and click .
   The rule details are displayed.

3. In the **Memory Threshold** field, add the maximum memory usage allowed for this rule in MB. **100 MB** is the default for new rules.



4. Click **Save**.
5. When you are finished changing the memory thresholds for individual rules in an ESA rule deployment, redeploy the deployment.
   For more information, see the *Alerting with ESA Correlation Rules User Guide*.

Start, Stop, or Restart ESA Service

This topic provides instructions to start, stop, or restart the ESA Correlation service from the NetWitness user interface and from the command line. These procedures apply to ESA Correlation Rules.

Start the ESA Service

1. Log on to NetWitness as admin.
2. Go to  **(Admin) > Services**.
3. Select the ESA Correlation service and then select  **> Start**.

Stop the ESA Service

1. Log on to NetWitness as admin.
2. Go to  **(Admin) > Services**.
3. Select the ESA Correlation service and then select  **> Stop**.

Restart the ESA Service

1. Log on to NetWitness as admin.
2. Go to  **(Admin) > Services**.
3. Select the ESA Correlation service and then select  **> Restart**.

Start the ESA Service from the Command Line

1. Use ssh to connect to the ESA Correlation service and log in as the root user.
2. Type the following command and press **ENTER**:
   ```
   systemctl start rsa-nw-correlation-server
   ```

Stop the ESA Service from the Command Line

1. Use ssh to connect to the ESA service and log in as the root user.
2. Type the following command and press **ENTER**:
   ```
   systemctl stop rsa-nw-correlation-server
   ```

Restart the ESA Service from the Command Line

1. Use ssh to connect to the ESA service and log in as the root user.
2. Type the following command and press **ENTER**:
   ```
   systemctl restart rsa-nw-correlation-server
   ```

View Audit Logs and Verify ESA Component Versions

This topic provides details about audit logging and instructions to verify the versions of the ESA components installed. These procedures apply to ESA Correlation Rules.

View Audit Logs for Rules

Audit logging allows you to view details about rules that are created and changed in NetWitness. There are local audit logs in each of the services in NetWitness. When Global Audit Logging is configured, NetWitness audit logs collect in a centralized system that converts them into the required format and forwards them to an external syslog system.

For details on how to access your local audit logs, see "Local Audit Log Locations" in the *System Configuration Guide*. To set up Global Audit Logging, see "Configure Global Audit Logging" in the *System Configuration Guide*.

The following Syslog global audit log examples show create, update, remove rule, and delete deployment actions for the ESA Correlation service (correlation-server).

Create Action

```
09-17-2018 08:59:50 System3.Info 10.0.0.0 Sep 17 15:59:54 esaprimary
{deviceVendor=RSA, deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} DataAccess{action=create, success=true,
identity=admin, parameters={EngineSettings=}}09-17-2018 08:59:50 System3.Info
10.0.0.0 Sep 17 15:59:54 esaprimary {deviceVendor=RSA, deviceVersion=11.3.0.0,
deviceService=correlation-server, deviceServiceId=1abc9465-d0d4-48a9-9205-
414066fabc2f, deviceProduct=NetWitness} DataAccess{action=create,
success=true, identity=admin, parameters={EngineSettings=}}09-17-2018 08:59:50
System3.Info 10.0.0.0 Sep 17 15:59:54 esaprimary {deviceVendor=RSA,
deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} API{action=/rsa/correlation/module/settings/set,
success=true, identity=admin, parameters={Arguments=[ModuleSettings(id=null,
name=a-d-v:multiple_failed-login_successful-login-rule-module,
displayName=ADV: Multiple_FailedLogin_SuccessfulLogin, enabled=true,
eplStatements=[module GHmoduleId15;@Name('GHmoduleName15') @Description
('GHmoduleDesc15') @RSAAlert(oneInSeconds=0, identifiers=
{"user_dst"}
) SELECT * FROM Event(ec_outcome in ('Success', 'Failure') AND ec_
activity='Logon').win:time(5 min) match_recognize (measures F as f_array, S as
s pattern (F F F F F+ S+) define F as F.ec_outcome= 'Failure', S as S.ec_
outcome= 'Success');], queries=[], maxConstituentEvents=null,
logFiredRules=null, trial=false, alert=ModuleSettings.Alert
(respondEnabled=true, severity=9, notificationReasons=[], uniqueIdentifiers=
[], rateLimit=RateL...09-17-2018 08:59:50 System3.Info 10.0.0.0 Sep 17
15:59:54 esaprimary {deviceVendor=RSA, deviceVersion=11.3.0.0,
deviceService=correlation-server, deviceServiceId=1abc9465-d0d4-48a9-9205-
414066fabc2f, deviceProduct=NetWitness} DataAccess{action=create,
success=true, identity=admin, parameters={ModuleSettings=}}
```

Update Action

```
09-17-2018 08:54:21 System3.Info 10.0.0.0 Sep 17 15:54:25 esaprimary
{deviceVendor=RSA, deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} DataAccess{action=update, success=true,
identity=admin, parameters={EngineSettings=5b9fce315068213b17760553}}09-17-
2018 08:54:21 System3.Info 10.0.0.0 Sep 17 15:54:25 esaprimary
{deviceVendor=RSA, deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} DataAccess{action=update, success=true,
identity=admin, parameters={EngineSettings=5b9fce315068213b17760553}}09-17-
2018 08:54:21 System3.Info 10.0.0.0 Sep 17 15:54:25 esaprimary
{deviceVendor=RSA, deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} API{action=/rsa/correlation/engine/settings/set,
success=true, identity=admin, parameters={Arguments=[EngineSettings(id=null,
name=endpoint-sa-managed, displayName=endpoint, description=endpoint,
enabled=true, eventType=Event, instanceId=1abc9465-d0d4-48a9-9205-
414066fabc2f, streamId=5b9fce314a5b1f5951babc29, moduleIds=
[5b9fce314a5b1f5951babc2a, 5b9fce314a5b1f5951babc2b],
enableStatementMetric=null)]}}
```

Remove Rule Action

```
09-17-2018 09:01:11 System3.Info 10.0.0.0 Sep 17 16:01:15 esaprimary
{deviceVendor=RSA, deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} API{action=/rsa/correlation/stream/settings/remove,
success=true, identity=admin, parameters={Arguments=
[5b9fcf7a4a5b1f5951babc2c]}}09-17-2018 09:01:11 System3.Info 10.0.0.0 Sep 17
16:01:15 esaprimary {deviceVendor=RSA, deviceVersion=11.3.0.0,
deviceService=correlation-server, deviceServiceId=1abc9465-d0d4-48a9-9205-
414066fabc2f, deviceProduct=NetWitness} API
{action=/rsa/correlation/stream/settings/remove, success=true, identity=admin,
parameters={Arguments=[5b9fcf7a4a5b1f5951babc2c]}}09-17-2018 09:01:11
System3.Info 10.0.0.0 Sep 17 16:01:15 esaprimary {deviceVendor=RSA,
deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} DataAccess{action=remove, success=true,
identity=admin, parameters={StreamSettings=5b9fcf7a4a5b1f5951babc2c}}
```

Delete Deployment Action

```
09-17-2018 09:02:45 System3.Info 10.0.0.0 Sep 17 16:02:50 esaprimary
{deviceVendor=RSA, deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} API{action=/rsa/correlation/engine/settings/remove,
success=true, identity=admin, parameters={Arguments=
[5b9fcfcb4a5b1f5951babc2f]}}09-17-2018 09:02:45 System3.Info 10.0.0.0 Sep 17
16:02:50 esaprimary {deviceVendor=RSA, deviceVersion=11.3.0.0,
deviceService=correlation-server, deviceServiceId=1abc9465-d0d4-48a9-9205-
414066fabc2f, deviceProduct=NetWitness} API
{action=/rsa/correlation/engine/settings/remove, success=true, identity=admin,
parameters={Arguments=[5b9fcfcb4a5b1f5951babc2f]}}09-17-2018 09:02:45
```

```
System3.Info 10.0.0.0 Sep 17 16:02:50 esaprimary {deviceVendor=RSA,
deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} DataAccess{action=remove, success=true,
identity=admin, parameters={EngineSettings=5b9fcfcb4a5b1f5951babc2f}}09-17-
2018 09:02:45 System3.Info 10.0.0.0 Sep 17 16:02:50 esaprimary
{deviceVendor=RSA, deviceVersion=11.3.0.0, deviceService=correlation-server,
deviceServiceId=1abc9465-d0d4-48a9-9205-414066fabc2f,
deviceProduct=NetWitness} API{action=/rsa/correlation/engine/stop,
success=true, identity=admin, parameters={Arguments=[madhavi-sa-managed]}}
```

Each log contains the following parameters:

- **Time stamp**: Time the rule was modified. Example: `09-17-2018 08:54:21`
- **System Info**: Information about the system where the action was performed, such as IP address. Example: `10.0.0.0`
- **deviceVersion**: Version of your ESA service. Example: `11.3.0.0`
- **deviceService**: Example: `correlation-server`
- **action**: Examples: `create`, `update`, `remove`
- **Parameters**: Placeholder for the following keys:
  - **Epl Module Identifier (moduleIds)**: unique identifier for the rules. Example: `5b9fce314a5b1f5951babc2a, 5b9fce314a5b1f5951babc2b`
  - **enabled**: Shows if the rule is enabled or not. Example: `enabled=true`
  - **respondEnabled**: Shows if alerts from this rule can go to the Respond view. Example: `respondEnabled=true`
  - **trial**: Displays if the rule is configured as a trial rule or not. Example: `trial=false`
  - **EplStatements**: Displays the rule syntax. Example:

    ```
    eplStatements=[module GHmoduleId15;@Name('GHmoduleName15') @Description
    ('GHmoduleDesc15') @RSAAlert(oneInSeconds=0, identifiers=

    {"user_dst"}

    ) SELECT * FROM Event(ec_outcome in ('Success', 'Failure') AND ec_
    activity='Logon').win:time(5 min) match_recognize (measures F as f_array,
    S as s pattern (F F F F F+ S+) define F as F.ec_outcome= 'Failure', S as
    S.ec_outcome= 'Success');]
    ```

  - **identity**: Example: `admin`

ESA Audit Logs on NW Server (11.5 and Later)

In NetWitness Platform 11.5 and later, in addition to the audit logs available on ESA Correlation-server, new audit logs on the NW Server (SA_SERVER) show when users add, modify, filter, delete, export, and import ESA rules in the Rule Library. The NW Server audit logs also show when users add, modify, and deploy ESA rule deployments. Modifications to an ESA rule deployment include adding, deleting, or updating a rule in a deployment as well as adding a data source or an ESA Correlation service to a deployment.

Verify ESA Correlation Version

1. Use ssh to connect to the ESA Correlation service and log in as the root user.
2. Type the following command and press ENTER:
   ```
   rpm -qa | grep rsa-nw-correlation-server
   ```
   The ESA Correlation server version is displayed.

References - Previous ESA Versions

This section is a collection of references, which describe the user interface for ESA Configuration in previous versions of NetWitness.

Note: ESA Analytics is not supported in NetWitness 11.5 and later versions.

See the following topics for details:

- Services Config View Data Sources Tab (11.2 and Earlier)
- Services Config View Advanced Tab (11.2 and Earlier)
- ESA Analytics Mappings (11.1.x to 11.4.x)
- Module Settings (11.1.x to 11.4.x)
- Whois Lookup Service Configuration (11.1.x to 11.4.x)

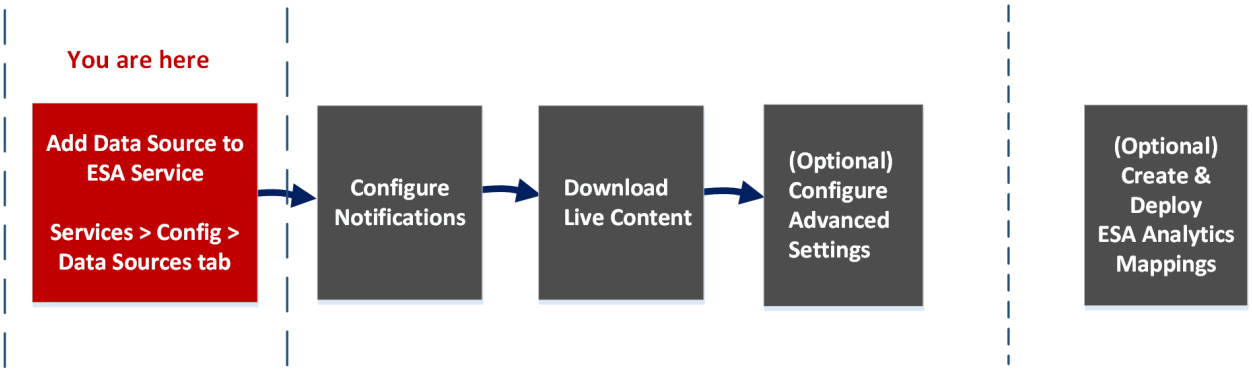
Services Config View Data Sources Tab (11.2 and Earlier)

Note: The information in this topic applies ONLY to NetWitness version 11.2 and earlier.
For version 11.3 and later, see Data Source Configuration Changes.
ESA Analytics is not supported in NetWitness Platform 11.5 and later versions.

The **Services Config view > Data Sources** tab of an ESA service enables you to configure the sources that ESA uses to analyze data. An ESA service ingests data from Concentrators to detect incidents and alert analysts to potential threats.

Workflow

This workflow shows the overall process for configuring ESA. It also shows where configuring data sources is located in the process.



In NetWitness 11.2 and earlier, ESA has two services, the Event Stream Analysis service (ESA Correlation Rules) and the Event Stream Analytics Server service (ESA Analytics). The first four procedures shown pertain to configuring the Event Stream Analysis service:

- Add Data Source to ESA Service*
- Configure Notifications
- Download Live Content
- (Optional) Configure Advanced Settings

The last procedure is separate from the others and pertains to creating mappings for the ESA Analytics services to start automatically detecting advanced threats:

- (Optional) Create and Deploy ESA Analytics Mappings

What do you want to do?

Role	I want to ...	Show me how
Administrator	Add a Concentrator as a data source to the Event Stream Analysis Service *	See "Configure ESA Correlation Rules" and "Step 1. Add a Data Source to an ESA Service" in the *ESA Configuration Guide for version 11.2.*
Administrator	Configure Notifications	See "Notification Methods" in the *Alerting with ESA Correlation Rules User Guide for version 11.2.*
Administrator	Download Live Content	See "Download Configurable RSA Live Rules" in the *Alerting with ESA Correlation Rules User Guide for version 11.2.*
Administrator	Configure Advanced Settings	See "Step 2. Configure Advanced Settings for an ESA Service" in the *ESA Configuration Guide for version 11.2.*

*You can complete these tasks here (that is in the Services Config view Data Sources tab).

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

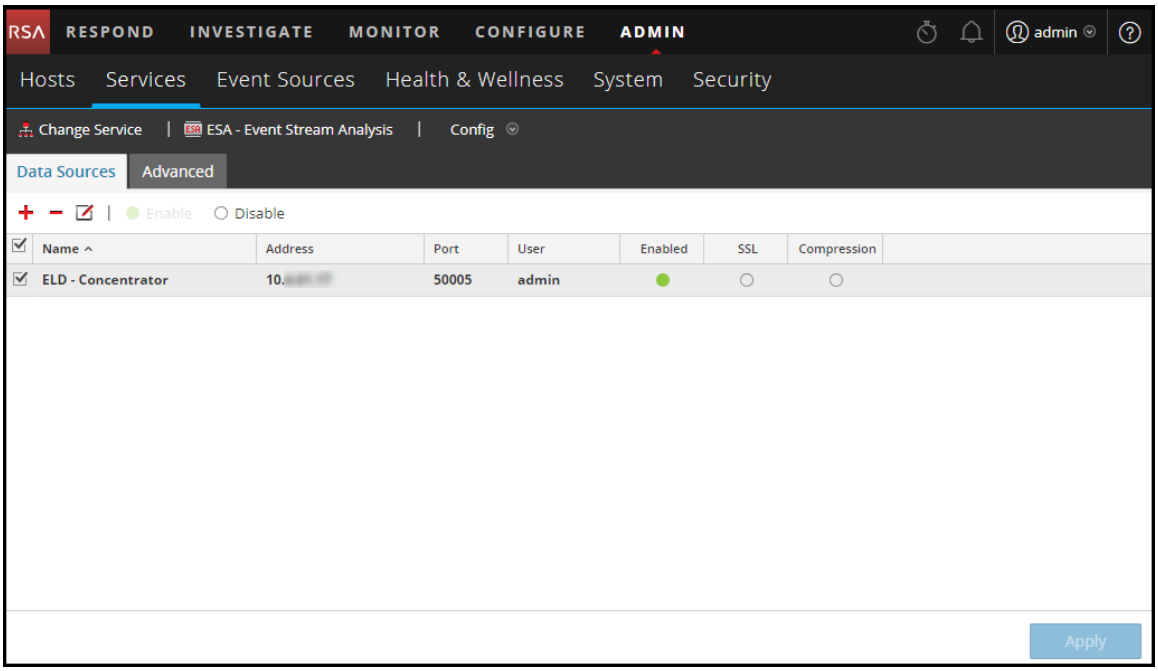
Related Topics

- See "Add or Update a Host" in the *Host and Services Getting Started Guide*.

Quick Look

To access the Data Sources tab, go to **Admin** > **Services** > (Select an ESA service) >  > **View** > **Config**.

The following figure shows the Services Config view Data Sources tab for an ESA service.



Toolbar

The following table describes the options in the toolbar.

Option	Description
+	Adds a new data source to the ESA service.
−	Deletes a data source from the ESA service.
	Edits a data source. You must have the username and password credentials for the service in order to make changes.
Enable	Enables the selected data source.
Disable	Disables the selected data source.

Data Sources

The Data Sources list shows all of the data sources added to the ESA service. The following table describes the columns the Data Sources list.

Column	Description
Name	The name of the data source service.
Address	The address of the data source service.
Port	The port used by the data source.

Column	Description
User	The user connected with the data source.
Enabled	Indicates if the data source is enabled.
SSL	Indicates if SSL communication is enabled.
Compression	Indicates if compression is enabled.

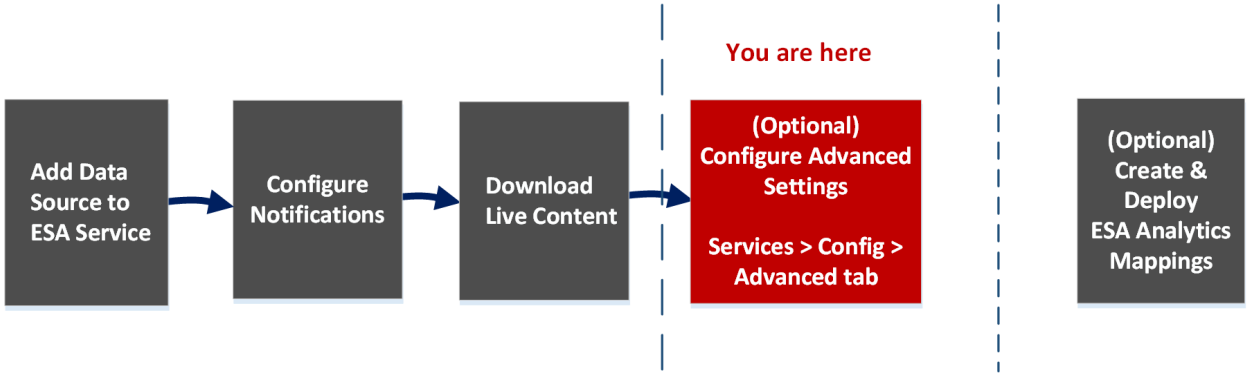
Services Config View Advanced Tab (11.2 and Earlier)

Note: The information in this topic applies ONLY to NetWitness version 11.2 and earlier.
For version 11.3 and later, see Configure Advanced Settings for an ESA Correlation Service.
ESA Analytics is not supported in NetWitness Platform 11.5 and later versions.

The **Services Config view > Advanced** tab of an ESA service enables you to configure advanced settings. In the Advanced view, you can configure advanced settings to improve performance, to preserve events for rules with multiple events, to buffer events in memory, and to set the number of events to be stored on the ESA.

Workflow

This workflow shows the overall process for configuring ESA. It also shows where configuring advanced settings is located in the process.



In NetWitness 11.2 and earlier, ESA has two services, the Event Stream Analysis service (ESA Correlation Rules) and the Event Stream Analytics Server service (ESA Analytics). The first four procedures shown pertain to configuring the Event Stream Analysis service:

- Add Data Source to ESA Service
- Configure Notifications
- Download Live Content
- (Optional) Configure Advanced Settings*

The last procedure is separate from the others and pertains to creating mappings for the ESA Analytics services to start automatically detecting advanced threats:

- (Optional) Create and Deploy ESA Analytics Mappings

What do you want to do?

Role	I want to ...	Show me how
Administrator	Add a Concentrator as a data source to the Event Stream Analysis Service	See "Configure ESA Correlation Rules" and "Step 1. Add a Data Source to an ESA Service" in the *ESA Configuration Guide for version 11.2*.
Administrator	Configure Notifications	See "Notification Methods" in the *Alerting with ESA Correlation Rules User Guide for version 11.2*.
Administrator	Download Live Content	See "Download Configurable RSA Live Rules" in the *Alerting with ESA Correlation Rules User Guide for version 11.2*.
Administrator	Configure Advanced Settings *	See "Step 2. Configure Advanced Settings for an ESA Service" in the *ESA Configuration Guide for version 11.2*.

*You can complete these tasks here (that is in the Services Config view Advanced tab).

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

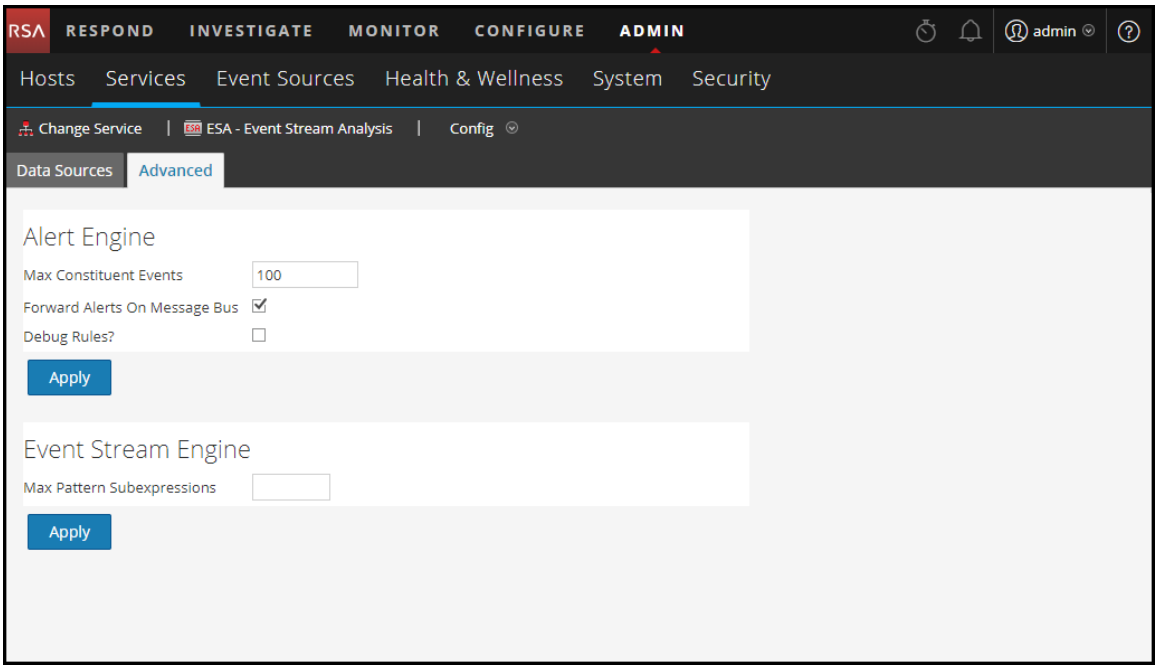
Related Topics

- See "Add or Update a Host" in the *Host and Services Getting Started Guide*

Quick Look

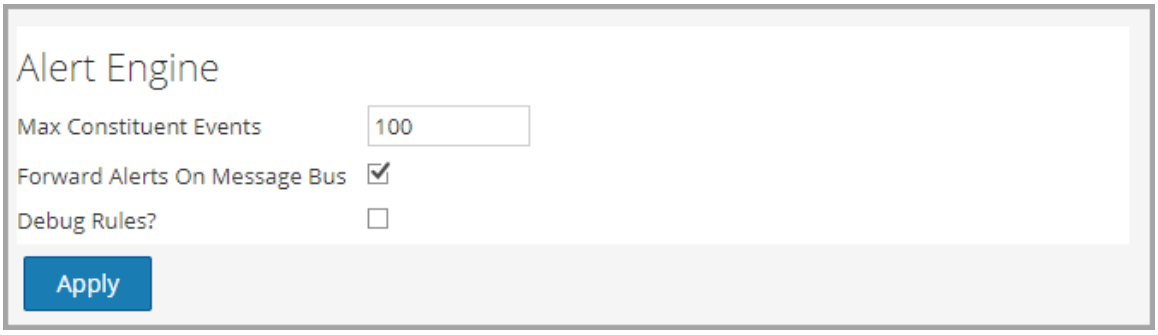
To access the Advanced tab, go to **Admin > Services >** (Select an ESA service) >  **> View > Config**.

The following figure shows the Services Config view Advanced tab for an ESA service.



Alert Engine Settings

In the Alert Engine section, you specify values to preserve events for rules that choose multiple events. The following figure shows the Alert Engine section.



The following table lists the parameters in the Alert Engine section and their descriptions.

Parameter	Description
Max Constituent Events	For rules that contain multiple events, this configuration value determines how many of the associated events are preserved. For example, if a rule fires an alert with 200 associated events and this parameter is set to 100, only the first 100 are preserved by ESA, the rest are dropped. The default value is 100.
Forward Alerts On Message Bus	To forward ESA alerts for NetWitness Respond, you must select this option. The ESA alerts generated will be sent to the Message Bus and subsequently to Respond. This option is selected by default. You may want to ensure that the Respond Server service is running.
Debug Rules?	Selecting enables debugging rules.

Event Stream Engine Settings

In the Event Stream Engine section, you specify details to improve performance. The following figure shows the Event Stream Engine section.



The following table lists the parameter in the Event Stream Engine section and its description.

Parameter	Description
Max Pattern Subexpressions	Certain rules require ESPER to maintain subexpressions in memory before deciding to fire them or not. These subexpressions consume memory and if left unchecked may cause the service to go down with memory exhaustion. This parameter is a safety measure that keeps such memory hogging rules under check. If a rule exceeds the specified number of subexpressions, its processing is delayed. The default value is 0, which means this setting is disabled. You must set a value if there are service stability issues.

Whois Lookup Service Configuration (11.1.x to 11.4.x)

Note: The information in this topic applies ONLY to NetWitness® Platform versions 11.1.x to 11.4.x. The Whois Lookup Service and ESA Analytics are not supported in NetWitness Platform 11.5 and later versions.

In the Whois Lookup Configuration panel (Admin > System > Whois), you configure a connection to the Whois Lookup service for your preconfigured ESA Analytics modules used in RSA Automated Threat Detection. The Whois Service enables you to get accurate data about domains that you connect to. In order to ensure effective scoring, it is important that you configure the Whois service settings.

You must have an RSA Live account to use this service.

If you configured a Live account in the Live Services panel (Admin > System > Live Services), the Whois Lookup Service is automatically configured for you. You just need to check the connection of the Whois Lookup service.

Note: If you do not have an RSA Live account, you can create one at the RSA Live Registration Portal: https://cms.netwitness.com/registration/ The *Live Services Management Guide* provides additional information.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure the Whois Lookup service.	See "Configure the Whois Lookup Service" in the *ESA Configuration Guide for NetWitness Platform 11.4*.
Administrator	Check the connection of the Whois Lookup service.	See "Configure the Whois Lookup Service" in the *ESA Configuration Guide for NetWitness Platform 11.4*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

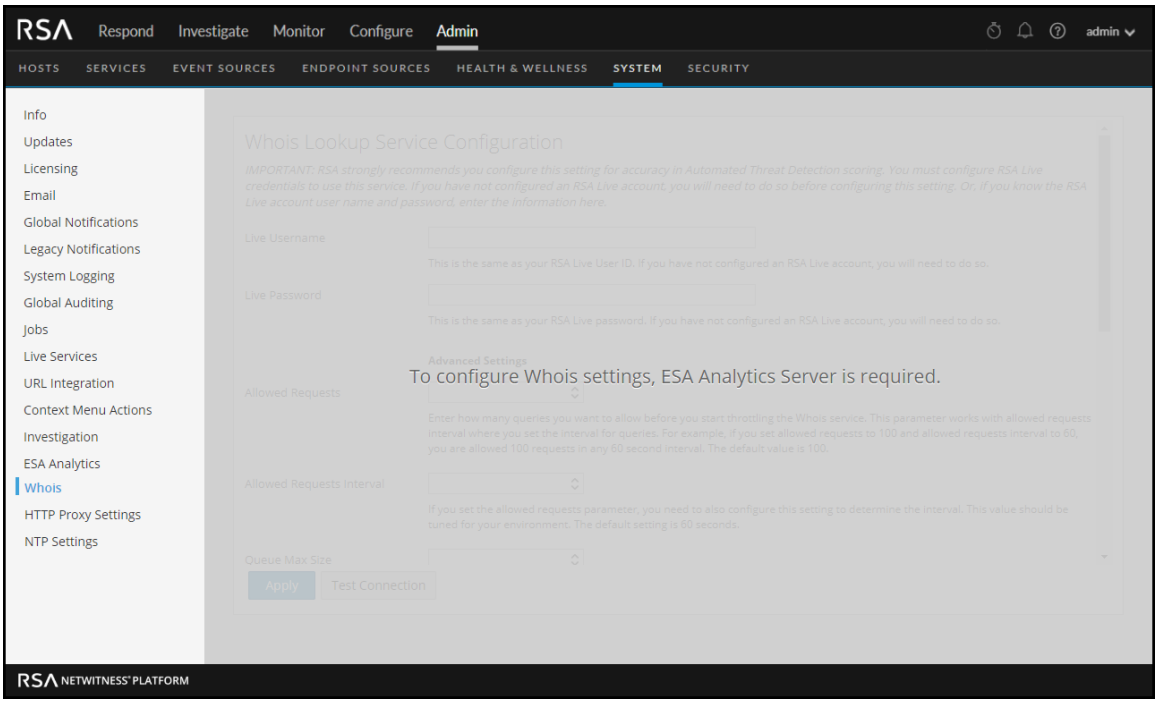
Related Topics

- See "ESA Analytics Mappings" in the *ESA Configuration Guide for NetWitness Platform 11.4*.

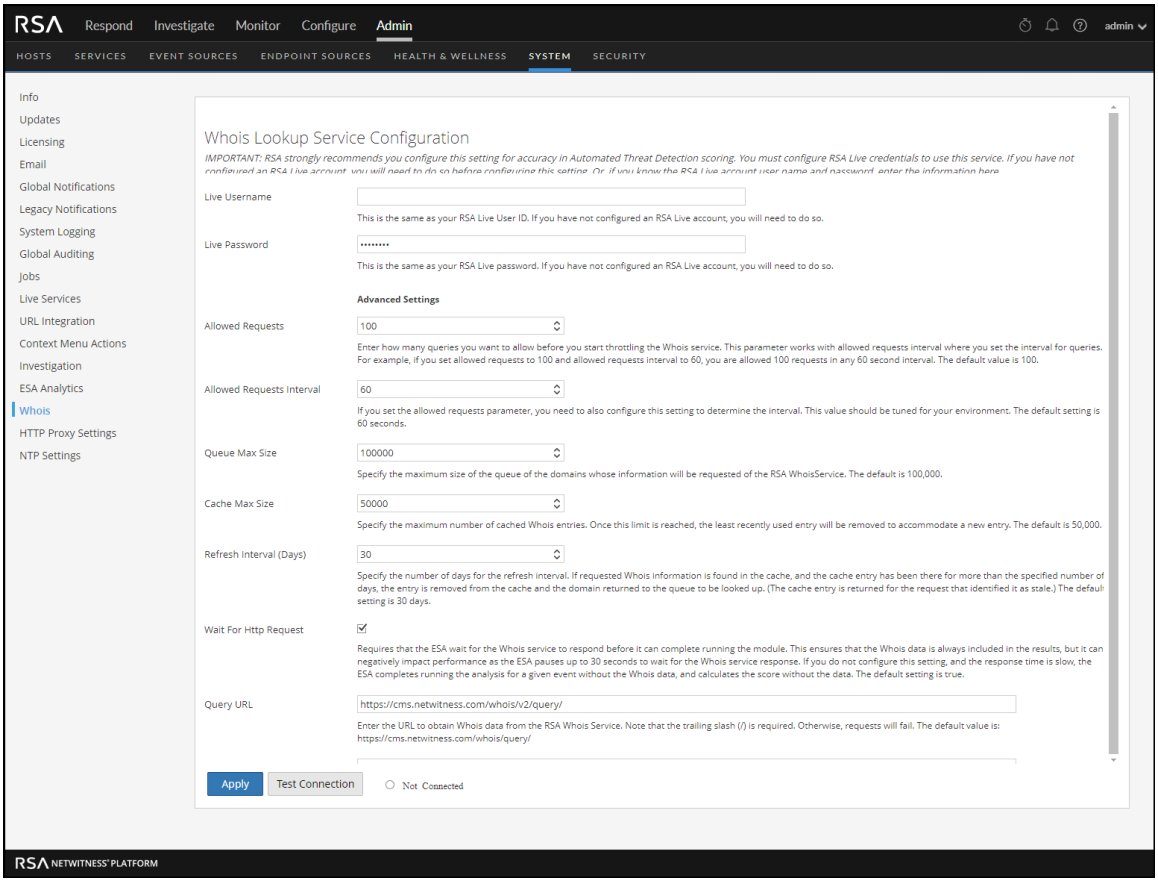
Quick Look

To access the Whois Lookup Service Configuration, go to Admin > System and in the options panel, select Whois.

The ESA Analytics Server service must be available (shows a green circle) in the Admin > Services view. If you do not have an ESA Analytics Server service available, you will see the following panel.



If you have an ESA Analytics Server service available, you will see the following panel.



The following table describes the listed Whois Lookup Service configuration settings.

Parameter	Description
Live Username	**Required only if you did not already configure the Whois Lookup service.** Enter the authentication credential for the Whois Server. This is the same as your RSA Live User ID. If you have not configured an RSA Live account, you will need to do so. The default value is "whois."
Live Password	**Required only if you did already configure the Whois Lookup service.** Enter the authentication credential for the Whois Server. This is the same as your RSA Live password. If you have not configured an RSA Live account, you will need to do so. The default value is null.
Allowed Requests	(Optional) Enter how many queries you want to allow before you start throttling the Whois service. This parameter works with **Allowed Requests Interval** (in seconds), where you set the interval for queries. For example, if you set **Allowed Requests** to 100 and **Allowed Requests Interval** to 60, you are allowed 100 requests in any 60 second interval. The default value is 100.
Allowed Requests Interval	(Optional) If you set the **Allowed Requests** parameter, you need to also configure this setting to determine the interval. This value should be tuned for your environment. The default setting is 60 seconds.
Queue Max Size	(Optional) Specify the maximum size of the queue of the domains whose information will be requested of the Whois Service. The default is 100,000.
Cache Max Size	(Optional) Specify the maximum number of cached Whois entries. Once this limit is reached, the least recently used entry will be removed to accommodate a new entry. The default is 50,000.
Refresh Interval (Days)	(Optional) Specify the number of days for the refresh interval. If requested Whois information is found in the cache, and the cache entry has been there for more than the specified number of days, the entry is removed from the cache and the domain returned to the queue to be looked up. (The cache entry is returned for the request that identified it as stale.) The default setting is 30 days.

Parameter	Description
Wait For HTTP Request	(Optional) Requires that the ESA wait for the Whois service to respond before it can complete running the module. This ensures that the Whois data is always included in the results, but it can negatively impact performance as the ESA pauses up to 30 seconds to wait for the Whois service response. If you do not configure this setting, and the response time is slow, the ESA completes running the analysis for a given event without the Whois data, and calculates the score without the data. The default setting is true.
Query URL	(Optional) Enter the URL to obtain Whois data from the Whois service. The trailing slash ('/') is required. Otherwise, requests will fail. The default value is: `https://cms.netwitness.com/whois/v2/query/`
Authentication URL	(Optional) Enter the URL to obtain authentication tokens from the Whois service. The default value is: `https://cms.netwitness.com/authlive/authenticate/WHOIS`

ESA Analytics Mappings (11.1.x to 11.4.x)

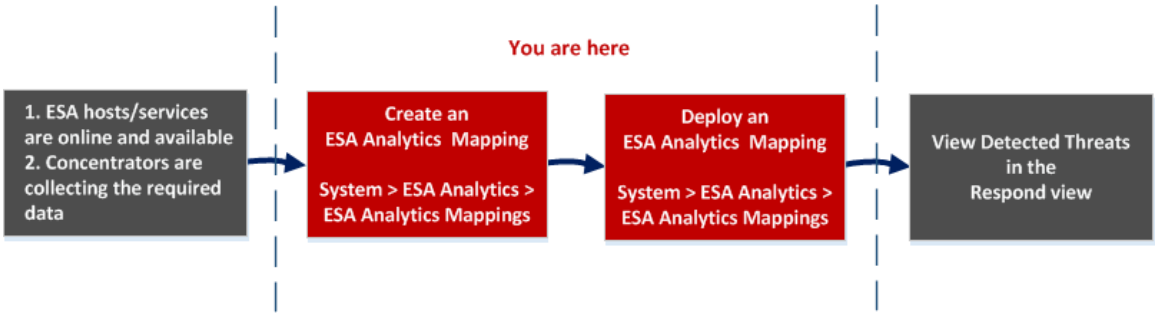
Note: The information in this topic applies ONLY to NetWitness versions 11.1.x to 11.4.x. ESA Analytics is not supported in NetWitness 11.5 and later versions.

In the ESA Analytics Mappings panel (Admin > System > ESA Analytics), you define how the RSA Automated Threat Detection functionality should automatically detect advanced threats. You can analyze the data that resides on one or more Concentrators by selecting a preconfigured ESA Analytics module.

To better utilize your network resources and reduce unnecessary data flow, you can map multiple data sources, such as Concentrators, to available ESA Analytics services in order to process data more efficiently and take advantage of additional capacity.

Workflow

This workflow shows the process for creating and enabling an ESA Analytics mapping to start automatically detecting advanced threats.



Before you create an ESA Analytics mapping, ensure that the ESA hosts and services that you want to use for your mappings are online and available. All of the services need to be in sync with a consistent time source. Also ensure that the Concentrators are collecting the required data. When you create an ESA Analytics mapping, you select an ESA Analytics module to map, such as Suspicious Domains. Then you select the data sources, such as Concentrators, to use for that module along with an ESA Analytics service to process the data. When you are ready to start aggregating data, you deploy the mapping. Analysts can view detected threats for that module in the Respond view.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Verify that the ESA hosts and services are online and available.	Admin > Hosts and Admin > Services See *Hosts and Services Getting Started Guide*.
Administrator	Ensure that the Concentrators are collecting the required data.	See *Broker and Concentrator Configuration Guide*

Role	I want to ...	Show me how
Administrator	Create ESA Analytics mappings.*	See "Mapping ESA Data Sources to Analytics Modules" in the *ESA Configuation Guide for NetWitness Platform 11.4*.
Administrator	Deploy ESA Analytics mappings.*	See "Mapping ESA Data Sources to Analytics Modules" in the *ESA Configuation Guide for NetWitness Platform 11.4*.
Administrator, Analyst	View detected threats.	See *NetWitness Respond User Guide*.

*You can complete these tasks here (that is in the ESA Analytics Mappings panel).

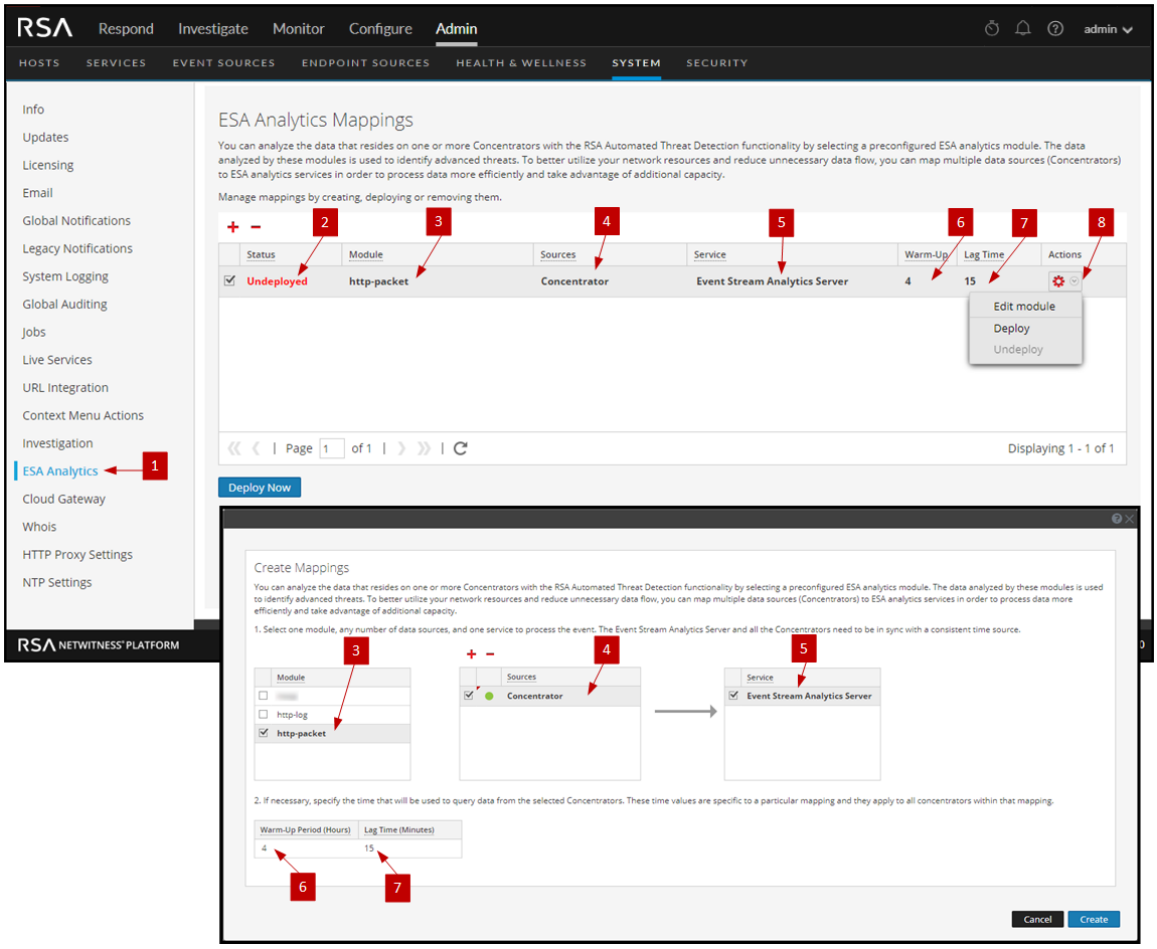
Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Related Topics

- "Configure ESA Analytics" in the *ESA Configuation Guide for NetWitness Platform 11.4*
- "Update a Mapping" in the *ESA Configuation Guide for NetWitness Platform 11.4*
- "Undeploy a Mapping" in the *ESA Configuation Guide for NetWitness Platform 11.4*
- "Delete a Mapping" in the *ESA Configuation Guide for NetWitness Platform 11.4*
- "Change the Warm-up Period and Lag Time" in the *ESA Configuation Guide for NetWitness Platform 11.4*
- Module Settings (11.1.x to 11.4.x)

Quick Look

The following example illustrates an ESA Analytics mapping. The configuration defines the data sources for the selected module and the ESA Analytics service that will process the events from those data sources.



1	Displays the ESA Analytics Mappings panel.
2	Shows the status of the ESA Analytics mapping.
3	The name of the module that is mapped.
4	Data sources, such as Concentrators, assigned to the mapping.
5	ESA Analytics service that processes the data for the mapping.
6	Warm-up period configuration (in hours) on the data sources for the mapping.
7	Lag configuration (in minutes) on the data sources for the mapping.
8	Actions for changing module settings, deploying module mappings, and undeploying module mappings.

Toolbar

The following table describes the toolbar actions.

Icon / Button	Description
	Opens the Create Mappings dialog where you can create an ESA Analytics mapping. Create a separate mapping for each module. After creating and reviewing the mappings, you deploy them.
	Deletes an ESA Analytics Mapping. • You can delete a mapping with a status of Undeployed at any time. Since a mapping in the Undeployed state is not deployed and is not running, it does not affect data aggregation. • Deleting a deployed mapping clears the configuration on the ESA server, reverts the deployment for that mapping, and stops pulling data from the data source for that module. You should undeploy a mapping with a status of Deployed before deleting it.
Deploy Now	After you create your mappings, you need to deploy them in order to start aggregating data for the modules. You can select one or more mappings with a status of Undeployed to deploy.

Note: If you want to make changes to a deployed mapping, such as adding or removing Concentrators or changing the service, you must undeploy and delete the existing mapping and then create and deploy a new mapping for that module.

ESA Analytics Mappings

The following table describes the listed ESA Analytics mappings.

Title	Description
	To select an individual mapping, select the checkbox next to the mapping.
Status	Shows the status of the mapping. There are two statuses: **Undeployed** - An undeployed mapping maps an ESA Analytics module to sources and an ESA Analytics service. It does not start aggregating data for the module until you deploy the mapping. **Deployed** - A deployed mapping is deployed and running. In a deployed mapping, the selected ESA Analytics service uses query-based aggregation to collect the appropriate filtered events for the selected module from the Concentrators.

Title	Description
Module	Indicates the selected ESA Analytics module. An ESA Analytics module is a pipeline composed of activity objects that enrich an event with additional information through mathematical computations. The module resides within the ESA Analytics service.
Sources	Sources are the data sources, such as Concentrators, from which ESA will aggregate the data for the specified module.
Service	Indicates the ESA Analytics service that will process the data for the specified module. The selected service needs to be in sync with a consistent time source.
Warm-Up Period (Hours)	Specifies a warm-up duration (in hours). A warm-up period is required to allow Automated Threat Detection to "learn" your traffic. The warm-up period should run when typical traffic is running. During this time, alerting for your module mapping is suppressed. The Warm-up Period primes the module with historical data and guarantees that the specified number of hours of data collection completes before sending alerts. NetWitness provides preconfigured ESA Analytics modules. Each module type has a default warm-up period defined, which you can adjust to your environment, if necessary. After this warm-up period, alerts can be viewed. For more information about Warm-up Period and Lag time, see Module Settings (11.1.x to 11.4.x).

Title	Description
Lag Time (Minutes)	Specifies a constant time delay in minutes, which is added to avoid losing events being processed by the data sources during periods of heavy activity. For example, Concentrator performance varies depending on factors such as incoming load, ongoing queries, and indexing. Due to these factors, a Concentrator may not aggregate events in real-time, which leads to the delay. The Lag parameter gives the Concentrator a chance to finish aggregating all of the data. After the warm-up period completes, data aggregation continues at **Current (System) Time - Lag Time**. This is useful when a Concentrator is slow in aggregating data. The Lag time guarantees that the module does not process data that arrives to the Concentrator within the Lag time window so there is adequate delay to ensure all events that get generated in the enterprise can be processed by the module. For example, if Lag time is 30 minutes, and the current time is 2:00 PM, the Concentrator starts pulling records at 1:30 PM. The Lag time window, 30 minutes in this example, remains constant as time advances. When the current time advances to 2:01 PM, the Concentrator pulls the next minute of data at 1:31 PM, and so on. **Important**: The Lag time defines the buffer between the current time and the time when the module ingests the data. **Caution:** NetWitness recommends that Administrators adjust the Lag parameter dynamically based on the performance of each of the individual Concentrators to avoid missing any events during aggregation. For more information about Warm-up Period and Lag time, see Module Settings (11.1.x to 11.4.x).
	Enables you to select additional actions for the selected module mapping: • **Edit Module** - Enables you to configure the warm-up period and lag time for the selected module mapping. • **Deploy** - Deploys the selected module mapping. The specified ESA Analytics service starts pulling data from the data sources for that module. • **Undeploy** - Undeploys the selected module mapping. The specified ESA Analytics service stops pulling data from the data sources for that module. **Caution:** Undeploying a mapping with a status of Deployed will affect data aggregation for that module.

Module Settings (11.1.x to 11.4.x)

Note: The information in this topic applies ONLY to NetWitness versions 11.1.x to 11.4.x.
ESA Analytics is not supported in NetWitness 11.5 and later versions.

After you create or deploy a module mapping in the ESA Analytics Mappings panel (Admin > System > ESA Analytics), you have the option to change some module configurations for that mapping.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Change the warm-up period for an undeployed module mapping.	"Change the Warm-up Period and Lag Time" in the *ESA Configuation Guide for NetWitness Platform 11.4*
Administrator	Change the warm-up period for a module mapping during the warm-up period.	"Change the Warm-up Period and Lag Time" in the *ESA Configuation Guide for NetWitness Platform 11.4*
Administrator	Change the warm-up period for a module mapping after the warm-up period is complete.	"Change the Warm-up Period and Lag Time" in the *ESA Configuation Guide for NetWitness Platform 11.4*

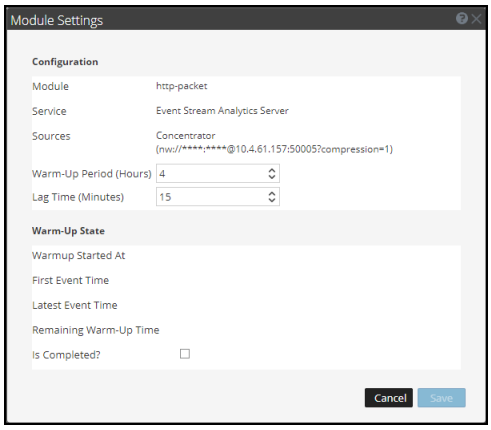
Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Related Topics

- "Mapping ESA Data Sources to Analytics Modules" in the *ESA Configuation Guide for NetWitness Platform 11.4*
- ESA Analytics Mappings (11.1.x to 11.4.x)

Quick Look

To access the module settings, in the ESA Analytics Mappings panel, select the mapping that you want to change and in the **Actions** column, select  **> Edit Module**. The Module Settings dialog has a Configurations section and a Warm-Up State section.



Configurations

The Configurations section enables you to change the Warm-Up Period and Lag Time configurations. The following table describes the settings available for an ESA Analytics module mapping.

Field	Description
Module	Shows the name of the mapped module.
Service	Shows the ESA Analytics service that processes the data for the mapping.
Sources	Shows the mapped data sources and the URLs used to communicate with ESA.

Field	Description
Warm-Up Period (Hours)	Specifies a warm-up duration in hours. A warm-up period is required to allow Automated Threat Detection to "learn" your traffic. The warm-up period should run when typical traffic is running. During this time, alerting for your module mapping is suppressed. The Warm-up Period primes the module with historical data and guarantees that the specified number of hours of data collection completes before sending alerts. NetWitness provides preconfigured ESA Analytics modules. Each module type has a default warm-up period defined, which you can adjust to your environment, if necessary. After this warm-up period, alerts can be viewed. You can update the Warm-Up Period of a deployed module mapping depending on whether or not the warm-up period is complete: • **During the warm up period** - You can add hours to the warm-up period or subtract any remaining warm-up time. • **The warm-up period is complete** - You can add hours to the warm-up period by adding the difference between the current time and the First Event Time to the hours that you want to add. For example, a warm-up period of 10 hours is complete and the First Event Time shows 12:00:00. The current (system) time is 16:00:00 (4 hours later) and you want to add 5 more hours to the warm-up time. To do this, you need to add 9 hours (4+5=9) to the warm-up period of 10, so you would set the new warm-up period to 19 hours. You cannot decrease the warm-up period if it is complete, unless you delete the mapping and create a new one. The Warm-up Period value is specific to a particular mapping and it applies to all Concentrators within that mapping after you deploy it. If a Concentrator is shared between two modules with different warm-up times, the Concentrator uses separate Warm-up Period values for each module mapping.

Field	Description
Lag Time (Minutes)	Specifies a constant time delay in minutes, which is added to avoid losing events being processed by the data sources during periods of heavy activity. For example, Concentrator performance varies depending on factors such as incoming load, ongoing queries, and indexing. Due to these factors, a Concentrator may not aggregate events in real-time, which leads to the delay. The Lag parameter gives the Concentrator a chance to finish aggregating all of the data. When you specify a Lag time, the first time the module deploys, data aggregation starts at **Current (System) Time - Lag Time - Warm-Up Time**. For example, if the current time is 2:00 PM, Lag time is 30 minutes, and Warm-up time is 4 hours, when the module deploys for the first time, data collection starts at 9:30 AM (2:00 PM - .5 hour - 4 hours). After the warm-up period completes, data aggregation continues at **Current (System) Time - Lag Time**. This is useful when a Concentrator is slow in aggregating data. The Lag time guarantees that the module does not process data that arrives to the Concentrator within the Lag time window so there is adequate delay to ensure all events that get generated in the enterprise can be processed by the module. For example, if Lag time is 30 minutes, and the current time is 2:00 PM, the Concentrator starts pulling records at 1:30 PM. The Lag time window, 30 minutes in this example, remains constant as time advances. When the current time advances to 2:01 PM, the Concentrator pulls the next minute of data at 1:31 PM, and so on. **Important**: The Lag time defines the buffer between the current time and the time when the module ingests the data. The Lag time value is specific to a particular mapping and it applies to all Concentrators within that mapping after you deploy it. If a Concentrator is shared between two modules with different Lag times, the Concentrator uses separate Lag values for each module mapping. **Caution:** NetWitness recommends that Administrators adjust the Lag parameter dynamically based on the performance of each of the individual Concentrators to avoid missing any events during aggregation. To determine the correct Lag Time, add together the following to get an environmental lag time: 1. **Log or Packet Latency** - This is the time it takes for the Log Decoder to receive the logs or the (Packet) Decoder to receive packets. For example, the Log Decoder may get logs every 20 minutes. In this case, you would want to set Lag time to at least 20 minutes, preferably 25 minutes, so that you do not miss events. 2. **Aggregation Latency** - This is the time it takes to get the data from the Log Decoder to the Concentrator. 3. **Other Buffer** - Add in any additional time delay specific to your

Field	Description
	environment.

Warm-Up State

The Warm-Up State section provides information about the warm-up state, which you can use to determine the appropriate adjustments to the warm-up period.

Field	Description
Warmup Started At	The time when the first event was processed by the ESA Analytics module from the data source.
First Event Time	The time that the first event occurred. The warm-up time is based on this time.
Latest Event Time	The time that the latest event occurred.
Remaining Warm-Up Time	The number of hours remaining in the warm-up period.
Is Completed?	Indicates whether the warm-up period is complete. If it is true, the warm-up period is complete. If it is false, the module is still warming up and you can view the number of hours remaining in the Remaining Warm Up Time field.

NetWitness® Platform XDR

Version 12.0

Event Sources Management

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

NetWitness Event Sources **6**

Event Source Management **7**

Workflow 7

Prerequisites 8

Automatic Mapping 8

Navigate to Event Source Management 8

How Alarms and Notifications Work 9

Large Email Notifications 10

High and Low Thresholds Both Triggered 11

Automatic Alerting 11

Common Scenarios for Monitoring Policies 12

Managing Event Source Groups **14**

Definitions 14

Manage Tab Details 14

Default Groups 14

Creating Event Source Groups 15

Procedure 15

Examples 16

Simple Example 16

Complex Example 17

Creating Event Source Group Form 18

Parameters 18

Rule Criteria 19

Acknowledging and Mapping Event Sources 20

Acknowledge Event Source Types 21

Manually Map Event Source Types 21

Viewing Logs from Pre-11.0 Log Decoder 22

Editing or Deleting Event Source Groups 22

Edit an Event Source Group 22

Delete an Event Source Group 22

Removing Idle Event Sources 23

Creating an Event Source and Editing Attributes 25

Mandatory Attributes 26

Create an Event Source 26

Update Attributes for an Event Source 27

Bulk Editing Event Source Attributes 27

Importing Event Sources 28

Import Event Source Attributes ... 29
Troubleshooting the Import File ... 31
Exporting Event Sources ... 31
Sorting Event Sources ... 32
Monitoring Policies ... 34
Configuring Event Source Group Alerts ... 34
Create an Alert Policy for an Event Source Group ... 34
Set and View the Thresholds for an Alert Policy ... 35
Setting Up Notifications ... 36
Prerequisites ... 36
Add Notifications for an event source group ... 36
Disabling Notifications ... 38
Prerequisites ... 38
Disable Notifications ... 38
Configuring Automatic Alerting ... 40
Prerequisites ... 40
Configure Automatic Alerting ... 40
Viewing Event Source Alarms ... 42
Sort the Alarms Information ... 42
Filter Alarms by Type ... 42
Event Source Management References ... 44
Discovery Tab ... 44
Manage Tab ... 49
Filter Panel ... 50
Groups Panel ... 52
Event Sources Panel ... 52
Sorting ... 54
Historical Graph for System Stats ... 54
Parameters ... 55
Examples of Zooming In ... 55
Manage Event Source Tab ... 56
Event Sources View ... 63
Create/Edit Group Form ... 66
Details View ... 67
Manage Parser Mappings ... 69
Best Practices ... 71
Advanced Configuration ... 71
Alarms Tab ... 72
Monitoring Policies Tab ... 74
Event Groups Panel ... 76

Thresholds Panel 76
Notifications Panel 77
Settings Tab 80
About Automatic Alerting 81
Features 83
ESM Troubleshooting & Appendix 85
Alarms and Notifications Issues 85
Alarms 85
Automatic Alarms 85
Manual Alarms 85
Notifications 85
Automatic Notifications 86
Manual Notifications 86
Duplicate Log Messages 86
Details 86
Clean Up Duplicate Messages 86
Troubleshooting Feeds 87
Details 87
How it Works 87
Feed File 88
Troubleshooting Feeds 88
10.5 Log Decoders 88
Feed File Existence 88
Group Meta Populated on LD 88
Device Group Meta on Concentrator 89
SMS Log File 89
Verify Logstats data is getting Read & Published by ESMReader & ESMAggregator 90
Configure JMX Feed Generator Job Interval 92
Import File Issues 93
Negative Policy Numbering 93
Details 93
Clean Up Duplicate Messages 94
Viewing Logs from Pre-11.0 Log Decoder 94

NetWitness Event Sources

Event sources are network sources that send information about events to the NetWitness. They can be physical devices, such as laptops, network switches or firewall, and virtual or cloud-based applications, such as VMware. For example:

- An Apache HTTP Server
- Amazon Web Services CloudTrail
- A Barracuda Web Application Firewall
- A connection to Dropbox
- An Oracle Database
- A VMware vCenter Server

You first configure all of your event sources so that they can communicate with the NetWitness. NetWitness provides configuration guides for many common event sources, using a variety of collection methods (such as Syslog or file collection). After you have your event sources configured, use the information in this guide to manage them going forward.

Event Source Management

The Event Sources module in NetWitness provides an easy way to manage event sources and configure alerting policies for your event sources.

Workflow

This workflow shows the overall process for managing event sources, and configure monitoring for them. It also shows where configuring alarms and alerts settings are located in the process.



Prerequisites

There are two permissions that affect Event Source Management:

- **View Event Sources** is needed for users to view event sources, their attributes, and their thresholds and policies.
- **Modify Event Sources** allows users to add, edit, and otherwise update event sources.

For details, see the following topics:

- The *Roles Tab* topic available in the **System Security and User Management** guide > **References** > **Administration Security View** > **Roles Tab**.
- The *Role Permissions* topic describes the built-in NetWitness system roles, which control access to the user interface. Available in the **System Security and User Management** guide > **How Role-Based Access Control Works**.
- The *Manage Users with Roles and Permissions* topic describes how to manage users in NetWitness, using roles and permissions. Available in the **System Security and User Management** guide > **Manage Users with Roles and Permissions**.

Automatic Mapping

Introduced in NetWitness version 11.1, the system maps incoming events to a type based on previous logs received from that address, reducing the mis-parsing of messages and reducing the number of items that need attention in the Discovery workflow. The User Interface (UI) indicates that an address has been auto-mapped in the Discovery workflow.

Navigate to Event Source Management

You can view the details about your existing event source groups by doing the following:

1. Go to  **(Admin) > Event Sources**.
2. Click any of the following:
   - The **Discovery** tab. Use this tab to review the event source types that NetWitness has discovered for each address and the system's confidence of how likely it is that they were identified completely accurately.
   - The **Manage** tab. Use this tab to add, edit, and delete event source groups and view details for your existing event source groups.
   - The **Monitoring Policies** tab. Use this tab to view or edit your event source alerting configuration.

- The **Alarms** tab. Use this tab to see the details of the alarms that have been generated. Alarms are generated when event sources exceed or fall below their set thresholds.
- The **Settings** tab. Use this tab to view or change the behavior for automatic alerts.

Note: When the system receives logs from an event source that does not currently exist in the Event Source List, NetWitness adds the event source to the list. Additionally, if it matches the criteria for any existing group, it becomes part of that group.

How Alarms and Notifications Work

The Event Source module in NetWitness displays alarms and sends notifications based on alarms that are triggered.

For alarms, consider the following:

Alarms are of two types: **automatic** (triggered when baselines are exceeded or not met) and **manual** (configured using thresholds).

- **Automatic**: If you turn on automatic alerts, the system reports alarms for **all** event sources that go above or below their normal baselines by the required amount. You can specify the over / under percentage on the Settings Tab.
- **Manual**: The system sends an alert whenever an event source exceeds the thresholds in the policy for the associated groups.
- Alarms appear on the UI, in the Alarms Tab.

For notifications, consider the following:

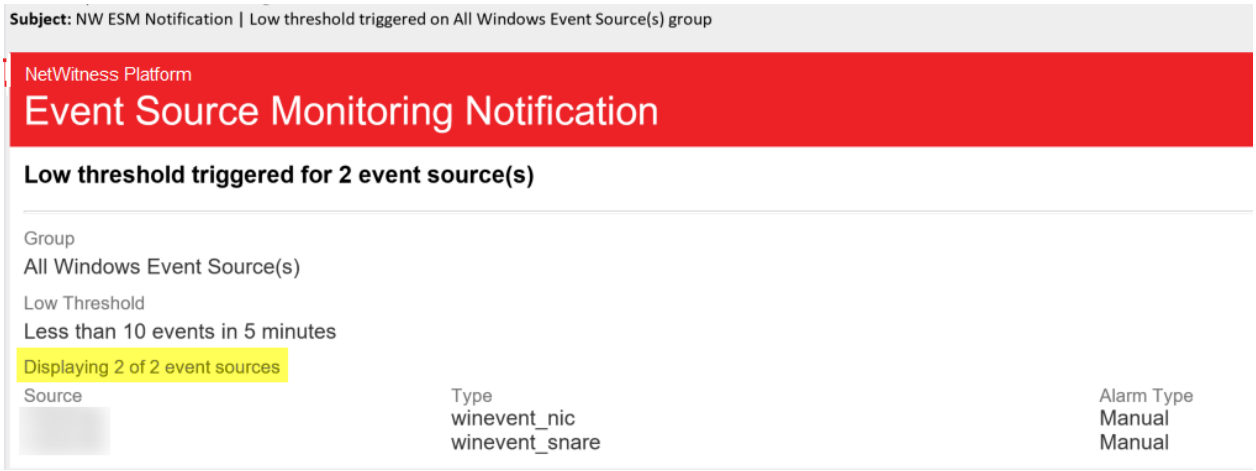
- To receive manual notifications (through email, SNMP or Syslog):
  - Specify a policy for an event source group.
  - Set a high or low (or both) threshold.
  - Enable the policy.
- To receive automatic (baseline) notifications:
  - Baseline alerting must be on. This is turned on by default.
  - You must enable notifications from automatic monitoring. See Configuring Automatic Alerting for details.
  - The event source that triggers the alarm must be in a group that has a policy enabled.
- If you have automatic alerting turned on, and you have configured a policy and threshold for a group:
  - If the event source goes outside its baseline, you see an automatic alert and receive a notification.
  - If the event source goes outside its thresholds, you see a manual alert and receive a notification.
  - If both occur (threshold and baseline exceeded or not met), you receive two alarms (visible on the Alarms tab) and a notification that indicates both alarms. That notification will list the event source that double alarmed twice; one listing indicating it was an automatic alarm.

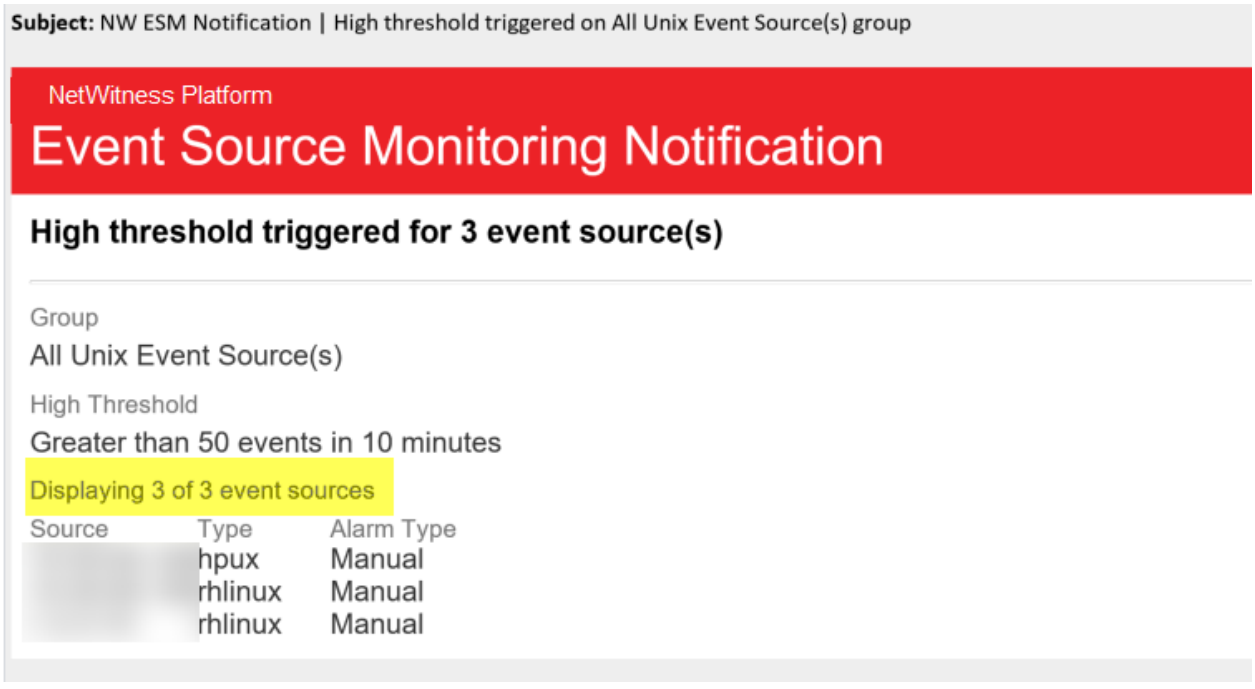
Large Email Notifications

If you have set up email notifications, keep in mind that the email can grow very large, depending on the number of event sources in the notification.

If the number of event sources in the alarmed state exceeds 10,000, then the email notification contains the details for only the first 10,000 and a total count. This is to ensure that the email is successfully delivered.

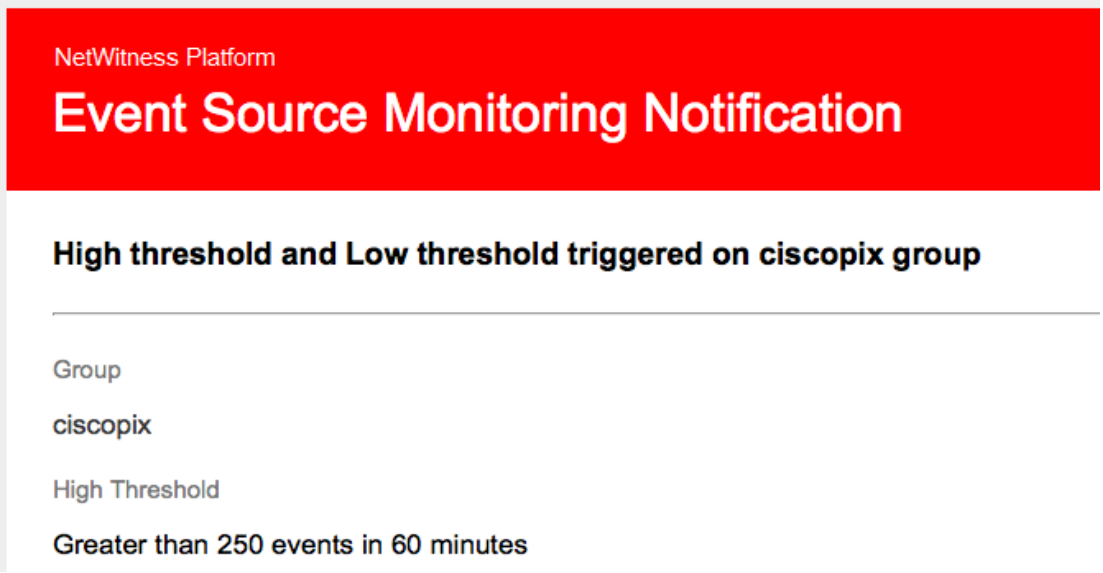
The following examples show a low threshold triggered for two event source groups and a high threshold triggered for three event source groups.





High and Low Thresholds Both Triggered

Occasionally, both the high and low alarms are triggered for a particular event source group. The easiest way to see when this happens is to read the email header, which clearly states when both thresholds are triggered, as shown in this image:



In this example, the header states, "High threshold and Low threshold triggered on ciscopix group." To see the details for the low threshold event sources, you may need to scroll down past hundreds, or even thousands, of the high threshold event sources.

Automatic Alerting

This topic describes automatic alerts, which are based on baseline settings.

Note: Automatic alerting, and all of the parameters that determine its behavior, are currently in Beta testing.

You can set up policies and thresholds for your event source groups. You do this so that you receive notifications when the thresholds are not met. NetWitness also provides an automatic way to receive alarms, if you do not want to set up thresholds to generate alarms.

To trigger automatic alerts, you can use baseline values. This way, you do not need to set up numerous group thresholds and policies in order to receive alerts. Any anomalous amount of messages trigger alerts, without needing to do any configuration (except for turning on automatic alerting).

Note the following:

- After you begin collecting messages from an event source, it takes the system approximately a week to store a baseline value for that event source. After this initial period, the system alerts you when the number of messages for a period are above or below the baseline by a set amount. By default, this amount is 2 standard deviations above or below the baseline.

- Base your high and low deviation settings on how "regular" your event sources behave. That is, if you expect little or no variance in the number of messages that arrive for a given time (for example, 8 to 9 am on a weekday), then you can set a low value for the Deviation. Conversely, if you often see peaks and valleys, set the Deviation value higher.
- If you enable a policy, but do not have any thresholds set, then you can still receive automatic (baseline) notifications, as long as you have turned on automatic alerting.

Common Scenarios for Monitoring Policies

Typically, organizations monitor their event sources in "buckets" based on how critical the event sources are. One typical example is as follows:

- There is a group of PCI devices, and it is critical to know if any of these devices stop sending messages (or send too few messages) within a half hour.
- There is a group of Windows devices, and it is useful to know if any of these devices stop sending messages after four hours.
- There is a group of quiet devices that do not typically send a lot of messages, but you would like to know if they do not send anything for 24 hours.

Many organizations may have a network that resembles this example. You may have more or different categories, but this example is used to discuss this feature.

You may have dozens or even hundreds of event source groups, and still only have a few groups for which you need to set thresholds and alerts.

Note: If an Event Source is a member of multiple groups that have alerting configured, it will only alert on the first matching group in the ordered list. (The Monitor Policies tab presents an ordered list of your groups.)

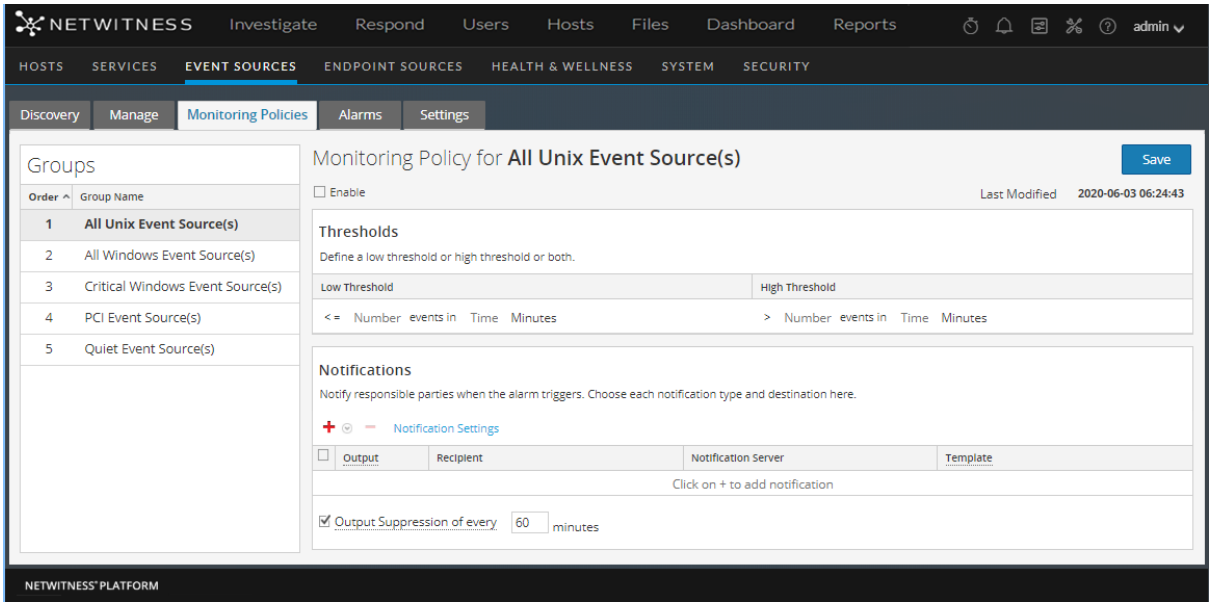
Ordering the Groups

Note: To change the order of the groups, drag and drop a group to its new location. The higher a group is listed, the higher the precedence for that group's thresholds: NetWitness checks the thresholds in the order provided in this panel. Thus, your highest priority groups should be at the top of this list.

The first thing to keep in mind is how to order your groups on the Monitoring Policies page. Assuming that you have the three groups mentioned above, you should order them as follows:

1. Quiet event sources. Having this group first ensures that you will not get numerous false alerts.
2. High priority PCI event sources. The highest priority devices should be after the quiet devices
3. Windows event sources. The time range is longer (four hours versus a half hour) for these devices than for the PCI devices. Therefore, they should come after the PCI devices.
4. All event sources. Optionally, you could set thresholds for all devices as a catch-all. This ensures that your entire network is operating as expected. For the catch-all group, you do not need to specify any

thresholds—you can use automatic alerting to generate alarms for the event sources in this group.



In the figure above, note the following:

- The groups are ordered as discussed in the previous section.
- The threshold for PCI devices is to alert if the number of messages coming in to NetWitness is fewer than 10 messages in 30 minutes.
- A low threshold is defined, but not a high threshold. This is typical for many use cases.

After you have set up and ordered your groups and begun to receive alerts, you may need to adjust the order. Use these guidelines to help you adjust the ordering:

- If you receive more notifications than you need, you can move the group down in the order. Similarly, if you are getting too few notifications, move the group up towards the top.
- If you notice that one event source is creating more alerts than it should, you can move it to another group, or create a new group for that event source.

Managing Event Source Groups

Definitions

When dealing with event source groups in NetWitness, note the following:

- An **event source** is essentially the combination of values for all of its attributes.
- An **event source group** is the set of event sources that match a set of criteria that are defined for that group.

For example, you might have the following groups:

- A group named **Windows Devices**, consisting of all the event source types associated with Microsoft Windows event sources (`winevent_nic`, `winevent_er`, and `winevent_snare`).
- A group named **Low Priority Services**, consisting of all services where the Priority attribute has been set lower than 5.
- A group named **U.S. Sales Servers**, where you gather event sources located in the U.S.A. and having an Organization attribute of Sales, Finance, or Marketing.

Manage Tab Details

The Manage tab in the Event Source module provides an easy way to manage event sources. In this tab, you can:

- Set up event source groups in a consistent way.
- Work with event source attributes in a consistent, straightforward manner.
- Easily search through your entire set of event sources.
- Bulk edit and update your event sources and event source groups.

You can view the details about your event source groups by doing the following:

1. Go to  **(Admin) > Event Sources**.
2. Select the **Manage** panel to see the details for your existing event source groups.

Note: When the system receives logs from an event source that does not currently exist in the Event Source List, NetWitness automatically adds the event source to the list. Additionally, if it matches the criteria for any existing groups, it becomes part of that group.

Default Groups

NetWitness has several default groups. You can customize these as required and use them as templates for creating new groups.

The default groups are as follows:

- All Event Sources
- All Unix Event Sources
- All Windows Event Sources
- Critical Windows Event Sources
- PCI Event Sources
- Quiet Event Sources

You can edit any of these groups to investigate the rules that define the groups.

Note: You cannot edit or delete the **All** event source group.

Creating Event Source Groups

Administrators must receive notifications when event sources are no longer being collected by NetWitness. They need to be able to configure how long the event sources can be quiet (that is, not collect any log messages) before sending a notification based on different factors.

NetWitness provides event source groups so that you can group similarly important devices together. You can create groups based on attributes that you imported from your CMDB (configuration management database), or by manually choosing event sources to add to the group.

For example, these are some of the types of event source groups that you can create:

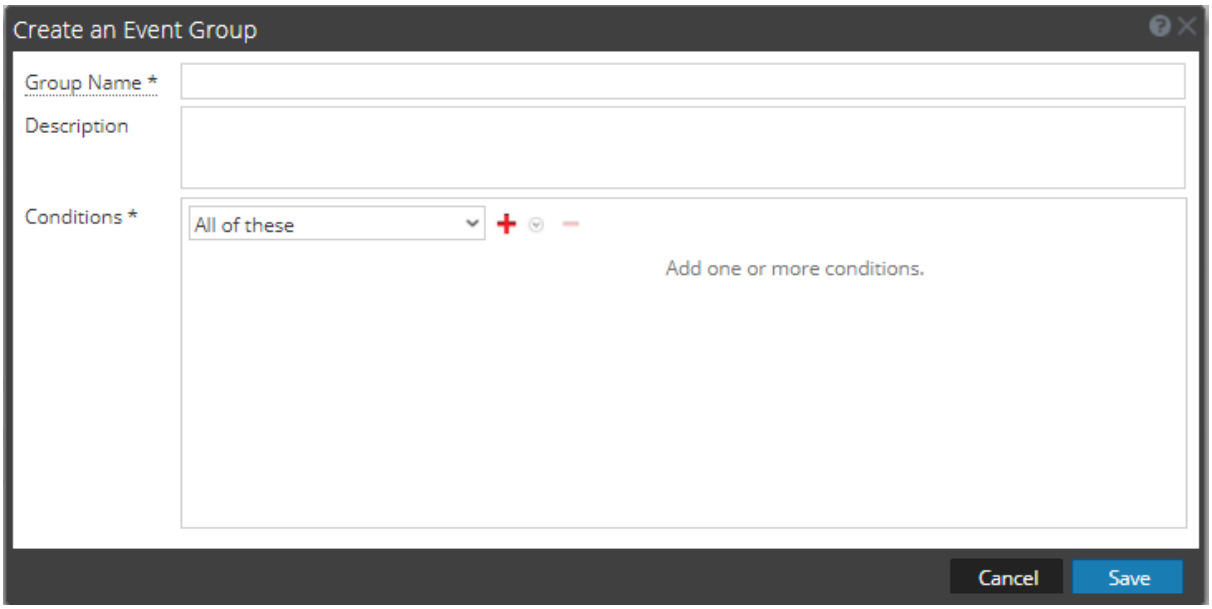
- PCI sources
- Windows Domain Controllers
- Quiet sources
- Finance Servers
- High Priority devices
- All Windows sources

Procedure

To create an Event Source group:

1. Go to  **(Admin) > Event Sources**.
2. In the **Manage** panel, click  .

   The Create an Event Group dialog is displayed.



3. Enter a Group Name.
4. Enter a Description.
5. Click + to add a condition. Continue adding conditions as necessary. For details on constructing conditions, see Create/Edit Group Form.
6. Click **Save**.

   The new group is listed in the **Manage** panel.

Examples

This section describes a simple example, and then discusses how to set up a more complex set of rules.

Simple Example

If you want to create an event source group that contains all of your high priority event sources, this example describes the necessary steps.

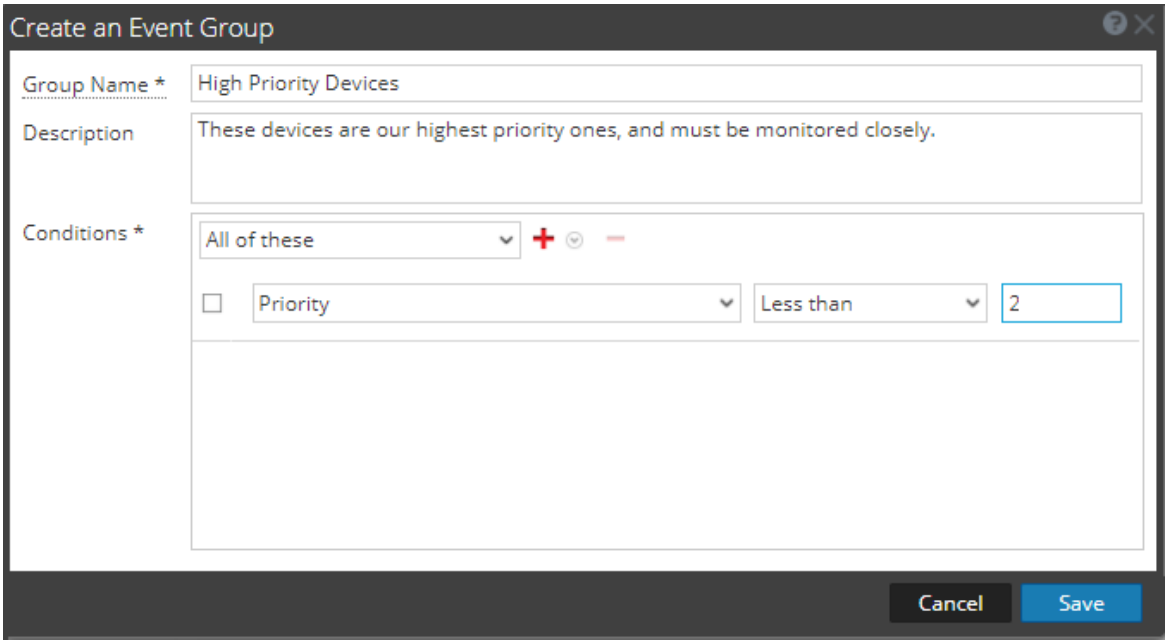
1. Go to **(Admin) > Event Sources**.
2. In the **Manage > Groups** panel, click + .
3. Enter **High Priority Devices** for the Group Name.
4. Enter a description, such as, "These devices are our highest priority ones, and must be monitored closely."
5. Leave **All of these** selected and click + to add a condition.
6. Select **Add condition** from the drop-down menu.

a. Select an Attribute: **Priority**.

b. Select an Operator: **Less than**.

c. Enter a value: **2**.

The following figure displays the updated Edit Event Group dialog.



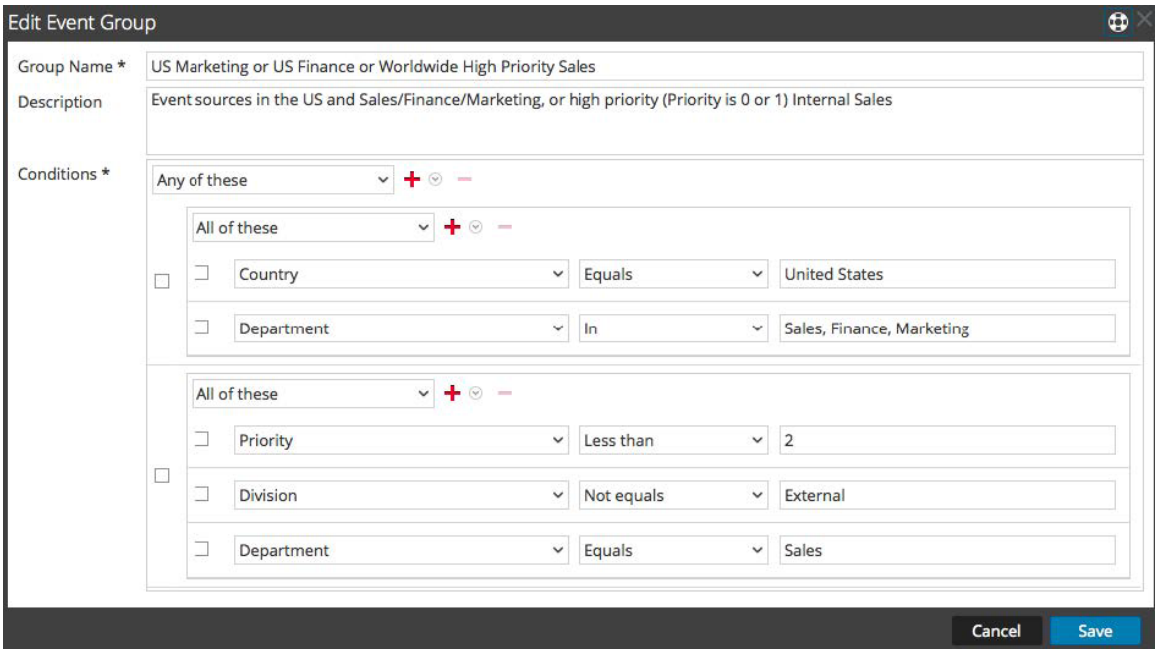
7. Click **Save**.

Complex Example

In this example, you want to create a fairly complex rule: match event sources that are in the United States, and in either the Sales, Finance, or Marketing departments. Also, match worldwide internal, high priority Sales event sources. High Priority is assumed to be where the priority is 1 or 0. Logically, the definition is as follows:

```
(Country=United States AND (Dept.=Sales OR Dept.=Finance OR
Dept.=Marketing))
 OR
 (Priority < 2 AND Division != External AND Dept.=Sales)
```

The following figure is an example of the criteria for creating such an Event Source Group.



Creating Event Source Group Form

The Create Event Source Group form is displayed when you are creating or editing an Event Source Group.

Parameters

The following table describes the fields on the Create/Edit an Event Group form.

Field	Description
Group Name	This field is required, and appears throughout the NetWitness UI as the identifier for the group.
Description	An optional description to help describe the purpose or details for the group.
Tools 	The following items are available on the toolbar: • **Add (+)**: clicking the **Add** displays a menu where you can choose to add a condition or a group. • **Remove (−)**: removes the selected rule or group of rules from the list. When you add a new group, that has the effect of creating nested levels of conditions.
Conditions	Described below, in the **Rule Criteria** table.
Cancel / Save	**Cancel** and **Save** options are available in the form.

Rule Criteria

The rules that you specify determine the event sources that will become part of this event source group. A rule consists of the following:

- Grouping: how the rule interacts with other rules
- Attribute: which attribute the rule is matching against
- Operator: how the rule matches the attribute
- Value: the attribute value used for the rule

The following table provides details on these rule constructors.

Rule Constructor	Details
Grouping	You can group conditions, in order to create complex rules for an event source group. The following choices are available when grouping your rules: • **All of these**: logically equivalent to AND • **Any of these**: logically equivalent to OR • **None of these**: logically equivalent to NOT If you are creating a simple group, and specifying a single condition, you can leave the default value (**All of these**) selected.
Attribute	This contains a drop-down list, consisting of all event source attributes. The attributes are displayed by the section to which they belong. For example, all of the **Identification** attributes are displayed first, followed by the **Properties**, **Importance**, and so on.

Rule Constructor	Details

Operator	Choose from the following options:
	• **Equals**: matches the provided value
	• **Not equals**: returns event sources whose specified attribute not equal to the provided value
	• **In**: provide a list of values in comma separated format, and event sources that match any of the provided values are included. For example:
	`Where IP in 10.25.50.146, 10.25.50.248`
	This condition returns event sources that have either `10.25.50.146 or 10.25.50.248` as their IP attribute.
	• **Not in**: similar to **In**, except that it matches items whose attribute is not equal to any of the listed values.
	• **Like**: matches items that begin with the provided string. For example:
	`Where Event Source Type Like Apache`
	This condition returns event sources whose Event Source Type begins with `Apache`.
	• **Not like**: similar to **Like**, except that it matches items whose attribute does not begin with the provided string.
	• **Greater than**: matches items whose attribute is greater than the provided value. For example, if you specify Priority Greater than 5, the condition would match any item with a priority of 6 or higher.
	• **Less than**: similar to **Greater than**. Matches items whose attribute is less than the provided value.

Value	Enter a value or group of values. The value type depends on the attribute for the condition. For example, for IPv6, you need to specify a value in IPv6 format.

Acknowledging and Mapping Event Sources

In NetWitness version 11.1, NetWitness introduced Automatic Mapping. The system automatically maps incoming events to a type based on previous logs received from that address, reducing the number of items that need attention in the Discovery workflow. The UI indicates that an address has been auto-mapped in the Discovery workflow.

Acknowledge Event Source Types

The Discovery tab lets you review the event source types that NetWitness has discovered for each address and the system's confidence of how likely it is that they were identified accurately. If the discovered event source types are correct, you can acknowledge to filter out that event source from the view by default. If incorrect, you can set the allowed event source types for a particular address so that future logs will parse against the correct parsers.

To acknowledge event sources:

1. Go to  **(Admin) > Event Sources**.

   The Discovery tab is displayed.
2. Select one or more event sources.
3. Click **Toggle Acknowledge**.

Note the following:

- After Event Sources are Acknowledged, they are no longer displayed in the Event Source Type(s) column.
- The **Toggle Acknowledge** button behaves as follows:
  - If the Acknowledged state for all of the selected event sources is the same, all values are toggled. That is, if you select only event sources with **Yes** in Acknowledged column, the value changes to **No** for all of them. Similarly if they all have **No** in the Acknowledged column, the value changes to **Yes** for all selected event sources.
  - If you select a multiple event sources, and the value for some is **Yes** and for other it is **No**, when you click **Toggle Acknowledge**, all of the values are set to **Yes** for the selected event sources.

Note: Acknowledged Event Sources are not displayed by default.

Manually Map Event Source Types

When discovered event source types are not completely accurate, you can manually map the parsers to obtain additional information.

To map one or more event sources:

1. Go to  **(Admin) > Event Sources**.

   The Discovery tab is displayed.
2. Select one or more event sources.
3. Click 

.

   The Manage Parser Mappings dialog box is displayed.

Note: For event sources that were created manually, the Manage Parser Mappings window has an empty Display Name in the Log Parsers column. To view the missing display names, close the Manage Parser Mappings dialog box, and then reopen it.

4. Add or remove parser mappings, and change the priority order, based on the needs of your organization. For more details, see Manage Parser Mappings .

Note: Discovery scores for the mapped Event Sources are listed in the Event Source Type(s) column from the lowest to highest discovery scores. Discovery scores range from 0 (least confident) to 100 (most confident).

Viewing Logs from Pre-11.0 Log Decoder

NetWitness 11.0 added the capability to view a small sampling of recent logs for specific devices through detail tabs of the Discovery View. By default, Log Decoders prior to 11.0 do not have the necessary configuration to enable this feature, but a few minor changes can make it available. For more details, see Viewing Logs from Pre-11.0 Log Decoder.

Editing or Deleting Event Source Groups

You may occasionally need to remove an event source group. For example, if you close an office, and you had a group consisting of all the event sources in that office, you can remove the group, since none of those event sources will send information to NetWitness.

Similarly, you may need to change some of the conditions that are used to populate the group.

Note: You cannot edit the event source group name. After you create a group, that name exists as long as the group itself exists.

Edit an Event Source Group

1. Go to  **(Admin) > Event Sources**.
2. In the **Manage** panel, select an existing Event Source Group.
3. Click .

   The Edit Event Group dialog is displayed.
4. Modify any of the details, or add, edit or remove conditions as necessary.
5. Click **Save**.

Delete an Event Source Group

Note the following:

- You can delete any group except for the **All** group, which lists all configured event sources in the system.
- If you delete a group, the associated policy for that group also gets deleted automatically.
- If there are any event sources that belong **only** to the deleted group, they would no longer have a policy alarm associated with them. Remember that event sources can belong to multiple groups.
- Deleting a group has no effect on baseline alarms.

To delete an event source group:

1. Go to  **(Admin) > Event Sources**.
2. In the **Manage** panel, select an existing Event Source Group.
3. Click .

   A confirmation dialog is displayed.
4. Click **Yes** to delete the group.

Removing Idle Event Sources

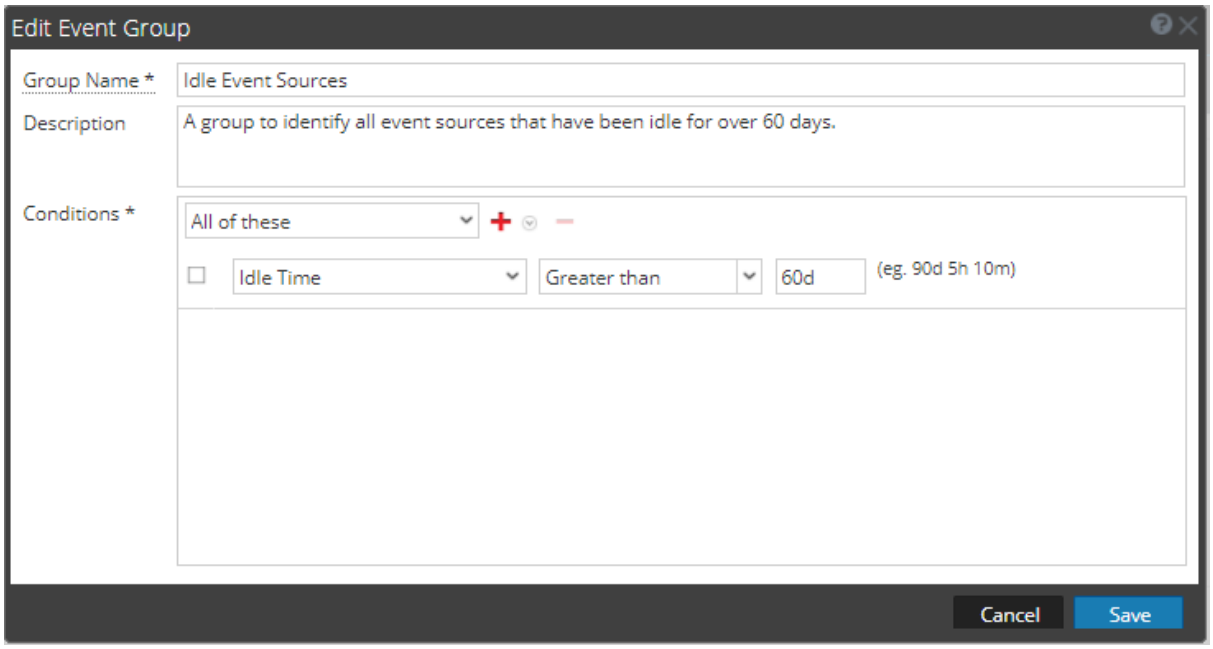
Periodically, you may want to update your set of event sources, and remove ones that are no longer being used. You can use the **Idle Time** parameter to do this.

Note: The information in this topic applies to NetWitness Version 11.2 and later.

To remove idle event sources:

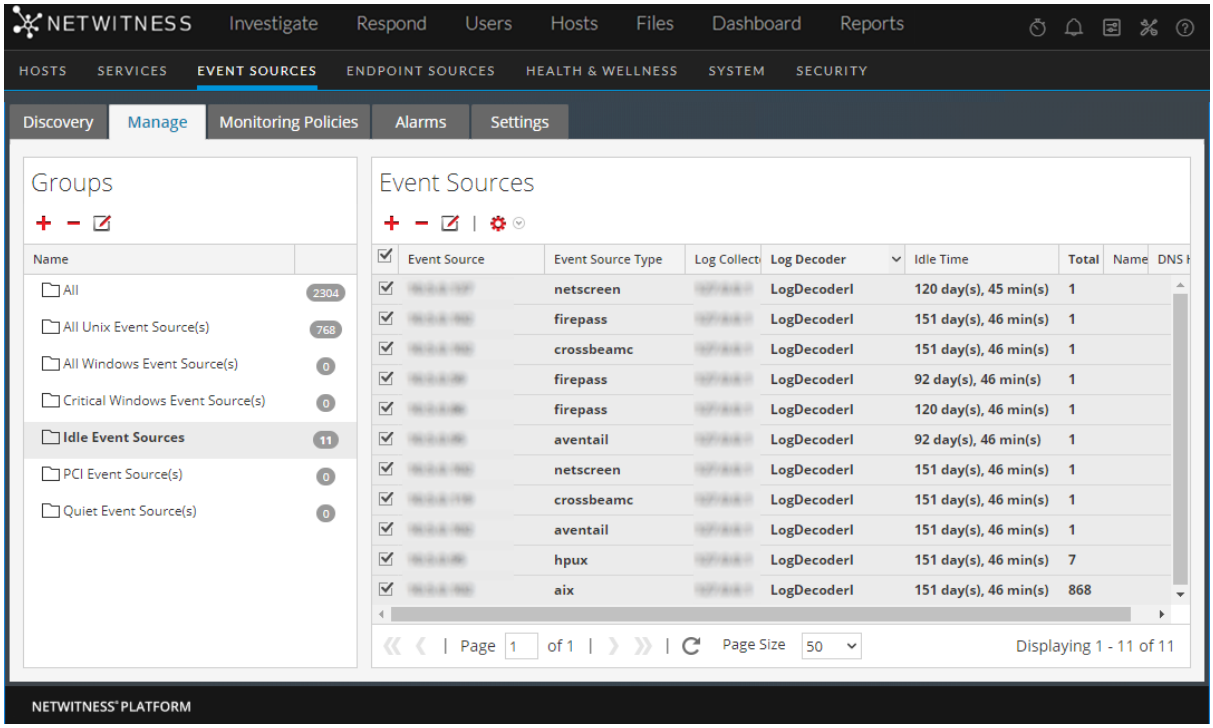
1. Go to  **(Admin) > Event Sources**.
2. In the **Manage** panel, click  .

   The Create an Event Group dialog is displayed.
3. Fill in the name and description as you like, and add a condition that uses the **Idle Time** parameter, as shown here:



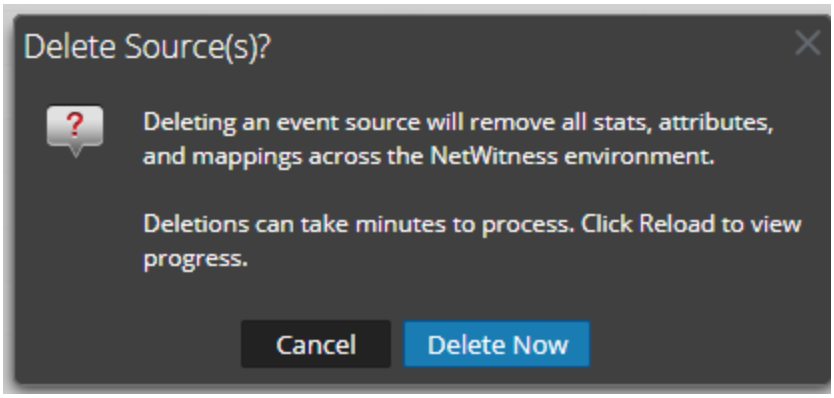
In this example, we have set the condition to identify event sources that have been idle for at least 60 days.

4. Save the new group, then select it in the Groups panel.
5. Select some or all event sources in the group. The following screen shows all event sources from this group selected.



6. In the Event Sources panel, click − to delete the selected, idle event sources.

A confirmation message appears:



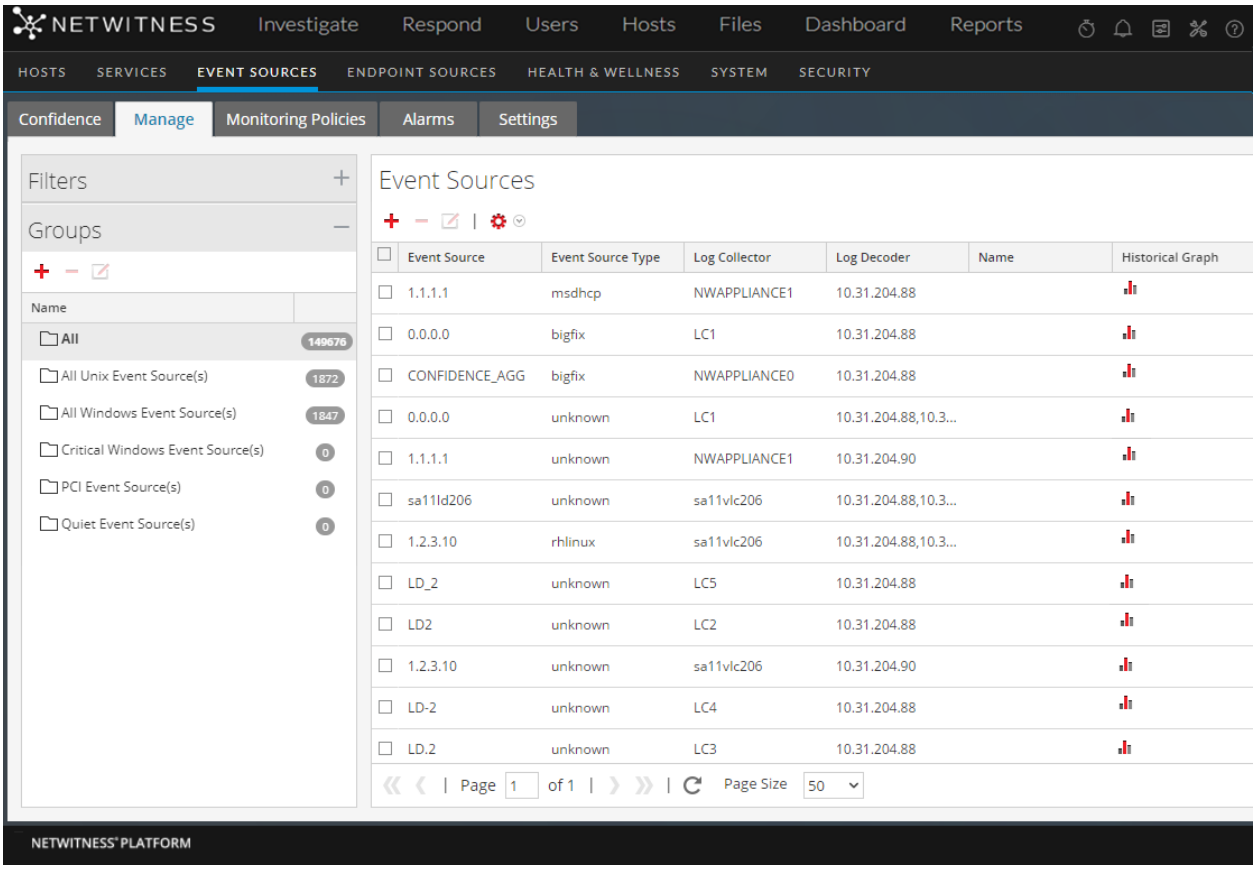
7. Click Delete Now to confirm your intention to delete the selected event sources.

If, in the future, an event source that has been removed sends logs, a new event source will be created.

Creating an Event Source and Editing Attributes

You can organize your event sources into groups. You do this by entering values for various attributes for each event source. For example, for all of your high priority event sources, you could set the **Priority** to 1. You can see details about the available attributes on the Manage Event Source Tab.

The following figure shows an example of the Event Sources panel:



Event source attributes are a combination of auto-populated and user-entered information. When an event source sends log information to NetWitness, it is added to the list of event sources, and some basic information is auto-populated. At any time after that, users can add or edit details for other event source attributes.

Mandatory Attributes

The following identification attributes are handled specially: **IP**, **IPv6**, **Hostname**, **Event Source Type**, **Log Collector**, and **Log Decoder**. If you create an event source manually, you can enter these values. After you save the event source, these values can no longer be changed.

Event sources can also be auto-discovered; any event source that sends messages to the Log Decoder will be added to the list of event sources. If you edit the attributes for an auto-discovered event source, you cannot edit any of these fields.

Note that not all of these fields are mandatory. To uniquely identify an event source, the following information is required:

- IP or IPv6 or Hostname, and
- Event Source Type

Additionally, NetWitness uses a hierarchy for IP, IPv6, and Hostname. The order is as follows:

1. IP
2. IPv6
3. Hostname

If you enter event sources manually, then you need to keep this order in mind, otherwise, you may end up with duplicates when messages are received from the event sources that you manually added.

All other attributes (such as Priority, Country, Company, Vendor, and so on) are optional.

Create an Event Source

1. Go to  **(Admin) > Event Sources**.
2. Select the **Manage** tab.
3. In the **Event Sources** panel, click  to open the details screen, which contains all of the event source attributes.

   The Manage Event Source Tab is displayed.
4. Enter or change the values for any attributes.
5. Click **Save**.

Note: The Discovery Score is listed as **Unavailable** for manually-added event sources. The score remains as **Unavailable** until the event source begins sending information to the NetWitness

Update Attributes for an Event Source

1. Go to  **(Admin) > Event Sources**.
2. Select the **Manage** tab.
3. In the **Event Sources** panel, select an event source from the list.
4. In the **Event Sources** panel, click  to open the details screen, which contains all of the event source attributes.

   The Manage Event Source Tab is displayed.
5. Enter or change the values for any attributes, except for certain attributes that cannot be altered after you have entered them.
6. Click **Save**

Bulk Editing Event Source Attributes

You can select multiple event sources, or an entire group, or even all event sources for bulk editing. For example, you might want to change the Priority or the Manager for a large number of your event sources.

Note: You cannot select individual event sources across displayed pages. For example, if you have a group with 225 event sources, and your Page Size is 50, you can only select event sources from the currently displayed 50 items.

If you want to edit items that span multiple pages, you can do the following:

- In the browser, increase the page size (the maximum is 500 entries on a single page). If your page size is small, you might be able to get all of your items on a single page.
- Create a new event source group that contains only the items that you want to bulk edit. Then, you can select all items for that group, rather than selecting individual items.
- Bulk edit incrementally. On the first page, select the items that you want to edit. Make your edits, then go to the next page and repeat the process, until you have made all of your changes.

Note: Mandatory fields cannot be edited; IP, IPv6, Hostname, Event Source Type, Log Collector, and Log Decoder.

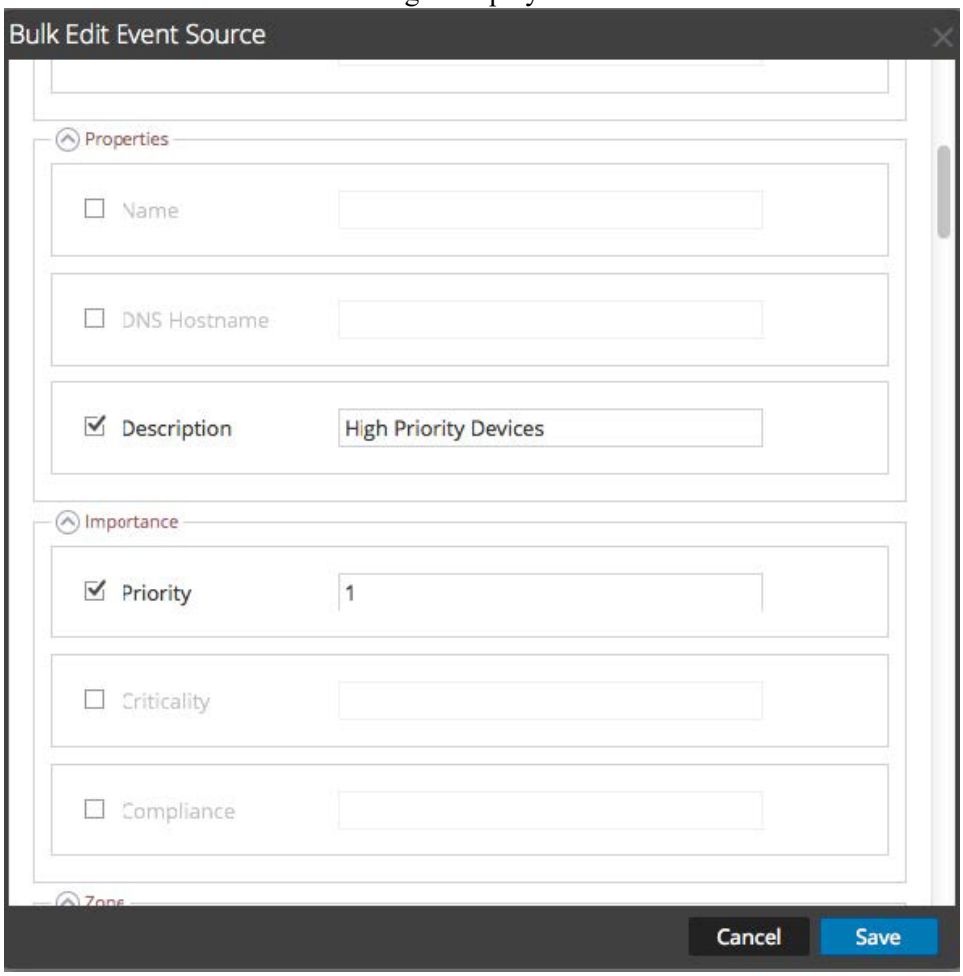
To bulk edit attributes for Event Sources:

1. Go to  **(Admin) > Event Sources**.
2. Select the **Manage** tab.
3. Optionally, select an event source group.
4. In the **Event Sources** panel, select one or more event sources to edit.

Note: To select all event sources, select the box next to the **Actions** column in the last (far-right) column of the list table.

5. Select the **Edit** icon from the menu bar.

   The Bulk Edit Event Source dialog is displayed.



6. Enter values for any of the available attributes. In the screen shot above, the Name and Priority attributes have been updated.
7. When you have updated as many attributes as required, click **Save**.

Importing Event Sources

You can import event source attributes from a CSV-formatted file. To import information from a configuration management database (CMDB), a spreadsheet, or other type of file, first convert or save the information to a CSV file.

Note: The following identification attributes are handled specially: **IP**, **IPv6**, **Hostname**, **Event Source Type**, **Log Collector**, and **Log Decoder**. If you import an event source that includes a different value for any of these fields (when compared with the value in NetWitness), the original value in NetWitness will **not** be overwritten.

The imported attributes are associated with the matched Event Source and are available for use in rules to create Event Source Groups.

NetWitness treats the import file as the correct, complete record. This assumption leads to the following behaviors related to importing event source attributes:

- By default, when you import attributes, the system updates attributes for existing event sources only.
- If the event source exists in the import file, but not in NetWitness, the attributes for that event source are ignored. That is, NetWitness does **not** create a new event source for these attributes.
- If the event source exists in both the import file and NetWitness, values for that event source are overwritten.
- If an attribute is blank in the import file, it clears the corresponding attribute in NetWitness.
- If an attribute is not specified in the import file, then the corresponding attribute is ignored in NetWitness (that is, it is **not** cleared).

Note: There is a difference between a blank attribute vs. one that is not specified at all. If an attribute is specified but blank, the assumption is that it is meant to be blank, and NetWitness clears that attribute for the corresponding event source. However, if an attribute is not specified at all, it is assumed that no change is expected.

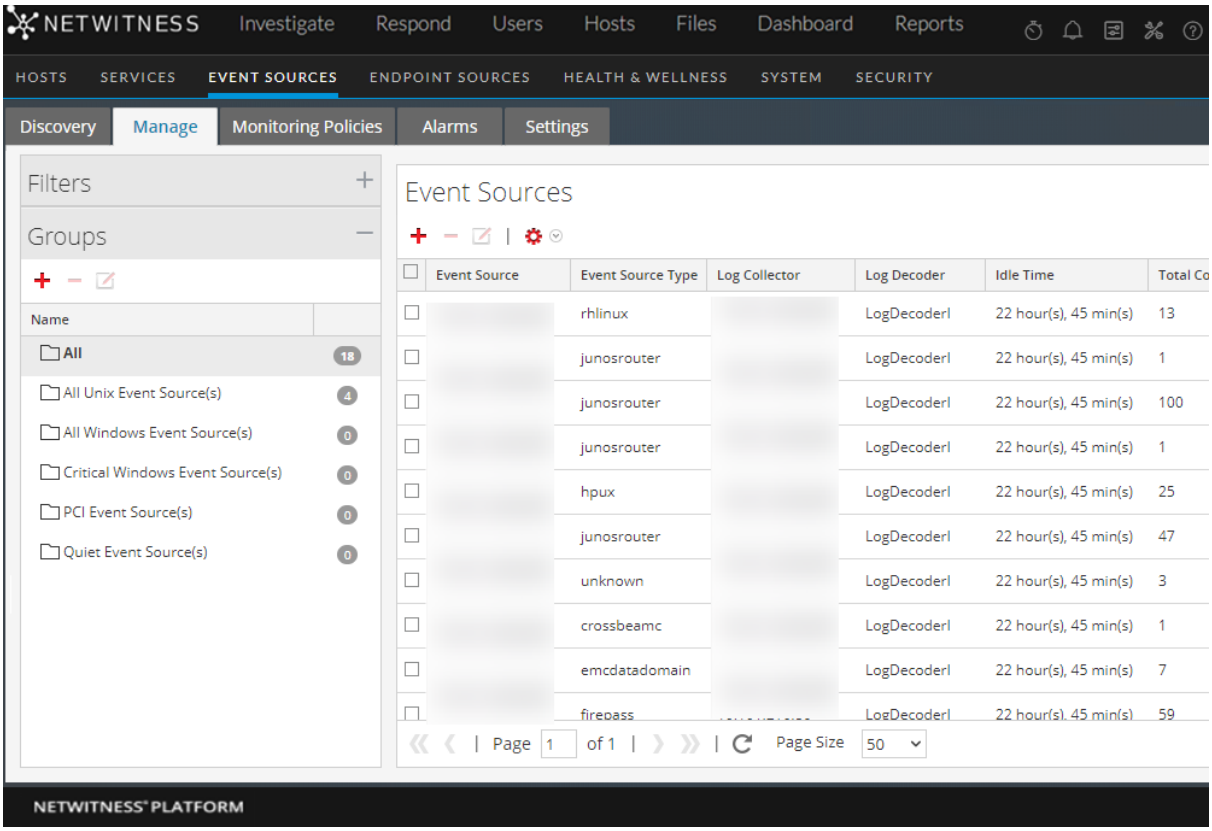
The above behaviors are the defaults—you can change the behavior as specified in the following procedure.

Import Event Source Attributes

To import Event Source attributes from a file:

1. Go to  **(Admin) > Event Sources**.
2. Select the **Manage** tab.

   The Event Sources Manage tab is displayed.



3. From the Import/Export menu in the toolbar ( ), select **Import** (Import).
   The Import Event Sources dialog is displayed.



4. Navigate to the import file, and select the appropriate boxes:
   - **Default**: The default behavior is described above.
   - **Add only**: Imports an attribute only if the corresponding field in NetWitness is blank. Thus, no existing values will be overwritten.

- **Do not clear values**: Does not clear attribute values in NetWitness for items in the import file that are blank.
- **Add Unknown Sources**: Adds new event sources based on items in the import file.

Note: You can select multiple options.

5. Click **Import**.
6. Click **Yes** in the confirmation dialog to perform the import.

Troubleshooting the Import File

If your import file is not formatted correctly, or is missing required information, an error is displayed, and the file is not imported.

Check the following:

- If you are adding unknown sources, each line in the file must contain a combination of the required attributes:
  - IP or IPv6 or Hostname, and
  - Event Source Type
- The first line of the file must contain header names, and the names must match the names in NetWitness. To get a list of correct column names, you can export a single event source. Examine the exported CSV file: the first row of the file contains the correct set of attribute/column names.

If your import file is not formatted correctly, or is missing required information, an error is displayed, and the file is not imported.

Exporting Event Sources

You can export all or some of your event sources, along with their corresponding attributes, to a CSV file.

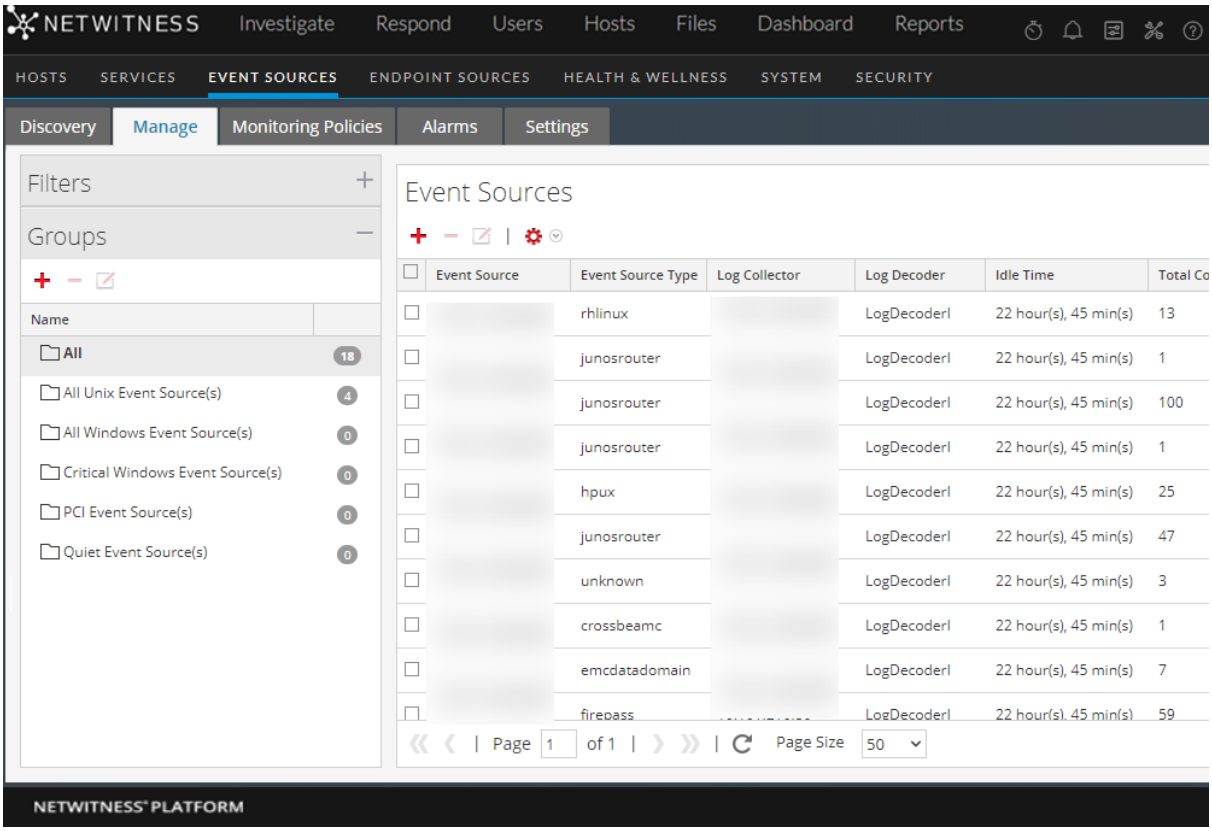
Note the following:

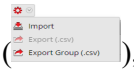
- The exported CSV includes all attribute columns.
- The exported CSV includes a header line at the top, listing each column name.
- You can export all entries in a group.
- You can export all entries (select the **All** group).
- You can select entries and export only those entries.

To export your Event Sources:

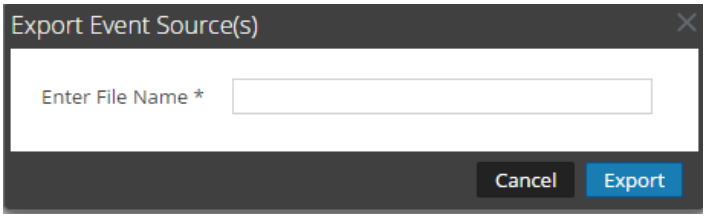
1. Go to  **(Admin) > Event Sources**.
2. Select the **Manage** tab.

   The Event Sources Manage tab is displayed.



3. Select the group that contains the event sources to export.
4. Select as many event sources as you need. Alternatively, you can export the entire group: to export the entire group, you do not need to select any individual event sources.
5. From the Import/Export menu in the toolbar (), select **Export (.csv)** or **Export Group (.csv)**.

   The Export Event Sources dialog is displayed.



6. Enter a file name and click **Export**.

The event source attributes are saved to the file name you specified, in a CSV format.

Sorting Event Sources

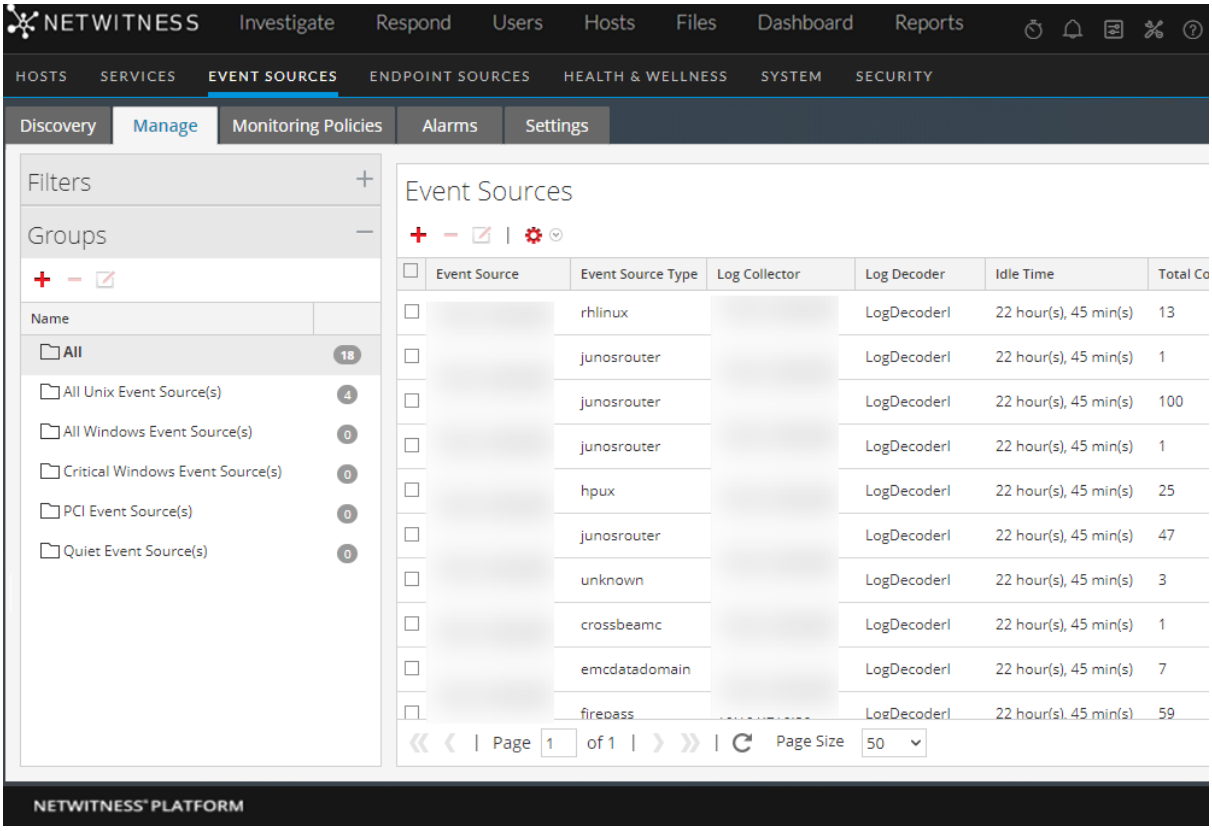
The event sources panel displays attributes for the currently selected event source group. You can configure the list of attributes that are displayed, as well as sort the list on any of the displayed attributes.

Note: The entire list is sorted, not just the items displayed on the current page. (The navigation bar at the bottom of the page shows how many pages exist for this list of event sources.)

To sort your event sources:

1. Go to **(Admin) > Event Sources**.
2. Select the **Manage** tab.

   The Event Sources Manage tab is displayed.



3. To sort a column, click  in the column header.

   The Sort Options drop-down menu is displayed.
4. Select the sort order that you want.

Monitoring Policies

Use the Monitoring Policies view to manage alert configuration for your event source groups.

You can create policies that alert on event source groups, by setting thresholds and notifications:

- Thresholds set ranges for frequency of log messages. You can specify a low threshold, a high threshold, or both.
- Notifications describe how and where to send alerts when thresholds are not met.
- You combine thresholds and notifications to create alerts based on the frequency you specify.
- If automatic alerting is enabled (it is by default), you can create and enable a policy ***without*** setting any thresholds. If you then turn on automatic notifications, notifications will be sent whenever an event source in the group is above or below its baseline by the specified amount.

For example, let's say that you have created an event source group that consists of all your Windows event sources based in the United Kingdom. You could specify a policy that alerts you whenever fewer than 1000 events per 30 minutes arrive.

Note: In addition to, or instead of setting up monitoring policies for your event source groups, you can Configuring Automatic Alerting to view alarms when the number of messages for an event source are outside of the normal bounds.

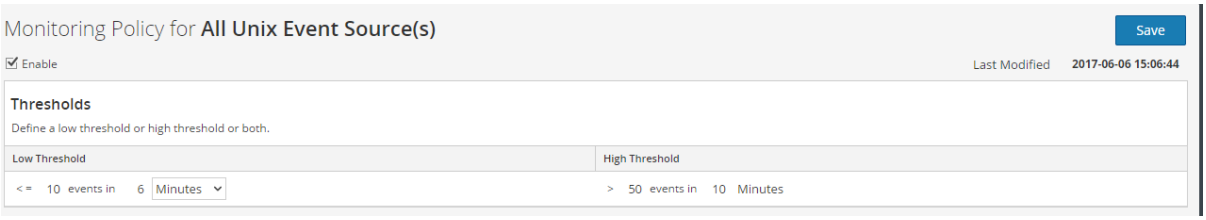
Configuring Event Source Group Alerts

Each event source group can have its own alerting policy. This includes setting the thresholds for when to alert, and setting the notification type when an alert is triggered. This topic describes the steps involved in creating an alert policy for an event source group.

Create an Alert Policy for an Event Source Group

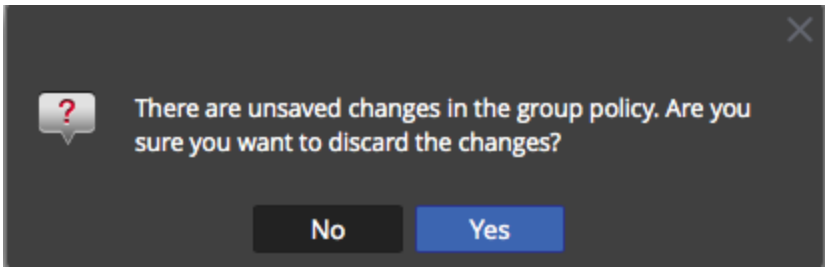
1. Go to  **(Admin) > Event Sources**.
2. Select the **Monitoring Policies** tab.
3. In the **Event Groups** panel, select a group.
4. Enter values for the Low Threshold and High Threshold fields.

   This is an example of alert thresholds.



5. Select **Enable** and click **Save** to enable the alert policy that you have configured.

Note: If you make changes to a policy, and attempt to exit the page before you save your changes, an Unsaved Changes warning message is displayed:

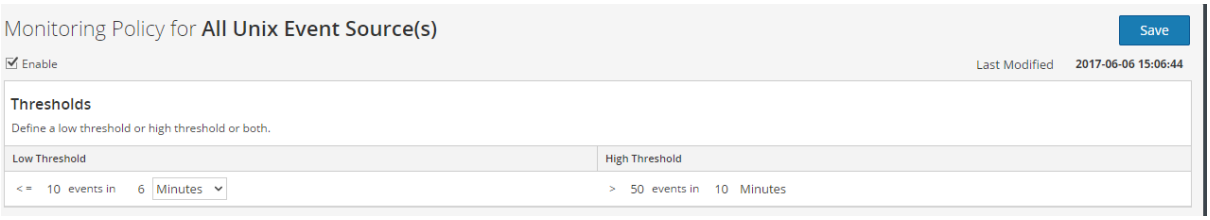


Set and View the Thresholds for an Alert Policy

Every event source group is also an alert policy. Thresholds are part of an alert policy. You can set thresholds for each alert policy. For each policy, you can set a low threshold, a high threshold, or both. Additionally, you can enable a policy without setting any thresholds; this allows you to receive notifications based on automatic alerts. Automatic alerts are generated when the baseline for an event source is out of normal bounds.

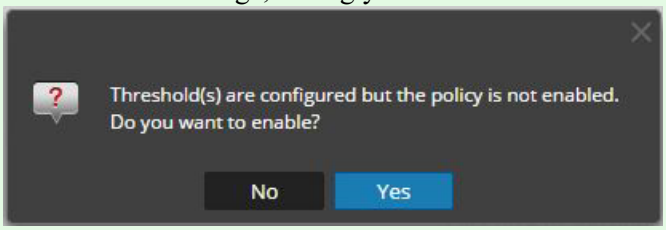
If you configure longer policy duration for low threshold policy, it may result in increased memory (SMS Heap) usage. In order to avoid increased memory usage (For more information, see 'Troubleshooting Health & Wellness' in the *System Maintenance Guide*), make sure you optimize the policy filter to match the required Event Sources and also NetWitness recommends having a policy duration of < 3 days if 30 low threshold policies match 20K Event Sources. If there are < 10 low threshold policies, you can have < 7 days policy duration with each policy matching unique Event Sources.

1. Go to  **(Admin) > Event Sources**.
2. Select the **Monitoring Policies** tab.
3. In the **Event Groups** panel, select a group.
   Any thresholds set for the selected group are displayed in the **Thresholds** panel.



4. Edit the values in either the Low or High Threshold as follows:
   a. Enter the number of events for the threshold.
   b. Enter the number of minutes or hours for the threshold. The minimum value is 5 minutes.
5. Select **Enable** to enable alarms when thresholds are not met.

Note: If you configure a threshold and attempt to save the page without enabling it, you receive a confirmation message, asking you whether or not to enable the policy: ddd



For example, suppose you enter 10 and 30 for the values for the low threshold: `10 events in 30 minutes`, and 20 and 30 for the values for the high threshold: `20 events in 30 minutes`. This means that you expect between 10 to 20 events are logged in 30 minutes (for the selected event source group). That is, anything between the low and high threshold is considered normal, and does not trigger an alarm.

After you add a threshold for a policy, you cannot delete it. You can disable the policy, or set the low or high threshold to 0 events in 5 minutes. Five minutes is the minimum duration for a threshold.

Setting Up Notifications

This topic describes how to configure notifications for event source groups. Notifications are sent when thresholds are not met.

Notifications go hand-in-hand with Thresholds. Before you configure notifications, you should set up Thresholds for an event source group.

Note: After configuring the thresholds for an event source group, if you do not set any notifications, then even if an alarm is triggered, users are not notified. However, all alarms are visible on the Alarms Tab.

Prerequisites

Before you set up notifications for an event source group, you should review the available notification items:

- **Notification Servers:** These are the servers that you want to receive notifications from the system. For more details, see the **Notification Servers Overview** topic in the *System Configuration Guide*.
- **Notification Templates:** These are the available templates for each type of notification. For Event Source Management, default templates are supplied for Email (SMTP), SNMP, and Syslog. You can use these templates as supplied, or customize them if necessary. For more details, see the **Templates Overview** topic in the *Systems Configuration* Guide.
- **Notification Output:** The outputs contain the parameters for the notification type. For example, an email notification type contains the email addresses and subject for the notification. For more details, see the **Notification Outputs Overview** topic in the *Systems Configuration* Guide.

Add Notifications for an event source group

To add notifications for an event source group:

1. Go to  **(Admin) > Event Sources**.
2. Select the **Monitoring Policies** tab.
3. In the **Event Groups** panel, select a group.

   Note: You should have already set a threshold for the group. If not, see Set and View the Thresholds for an Alert Policy to set a threshold, and then return to this procedure. Alternatively, if you have automatic alerting turned on, then you do not need to set thresholds for a policy. Automatic alarms generate notifications without the need to set thresholds.

4. In the Notifications panel, click , and from the drop-down menu, select the type of notification you want to add:
   - Email
   - SNMP
   - Syslog

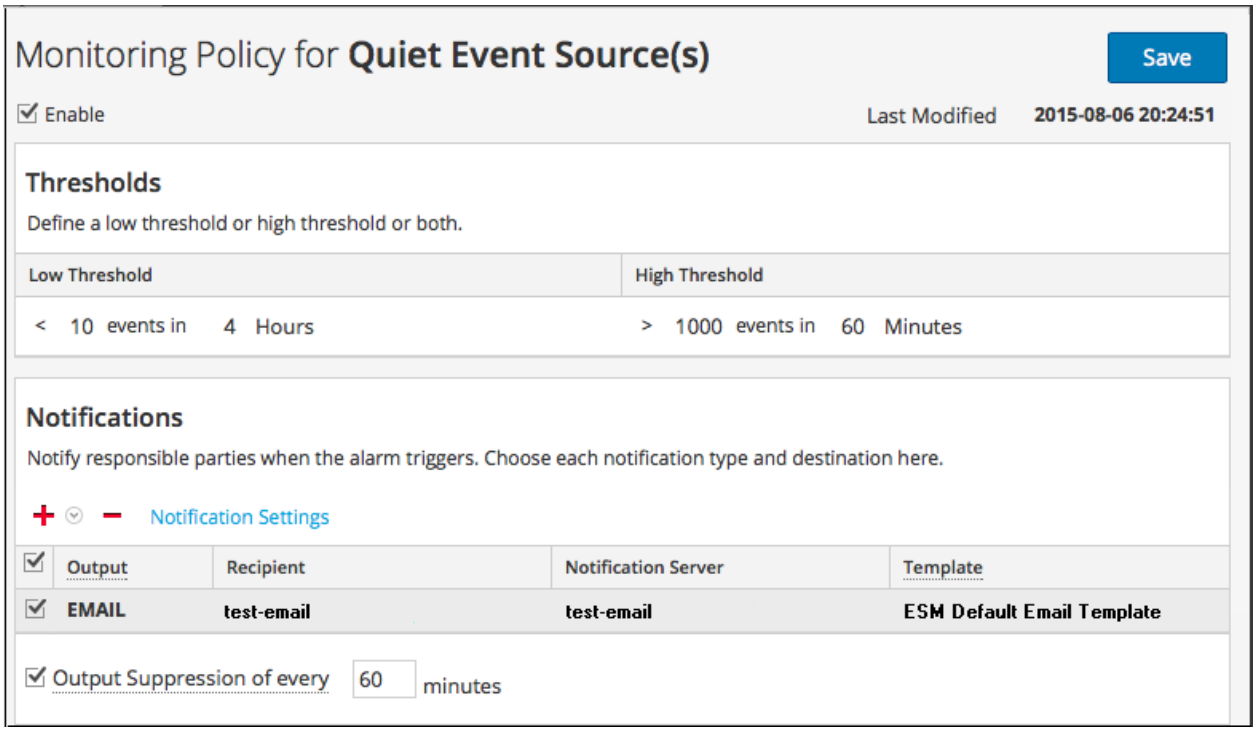
   Note: Default ESM (Event Source Monitoring) templates are provided for each type of notification.

5. Enter values for the Notification, Notification Server, and Template fields.
   a. For Notification, select from the list, or add a suitable notification type in **Notifications**, and then select it here.
   b. For the Server, select one from the list, or add a suitable server in **Notifications**, and then select it here.
   c. For Template, select an available template, or create a suitable template in **Notifications**, and then select it here.

   Note: If you need to add or edit one of these items, click **Notification Settings**. A new browser window opens on the **Administration > System > Global Notifications** page. Use this page to view or update the available Notification items.

6. Optionally, you can limit the rate of notifications for a policy.
   a. Select **Output Suppression** to enable setting a limit.
   b. Enter a value, in minutes, for the suppression rate. For example, if you enter **30**, notifications for this policy are limited to one notification every 30 minutes.
   c. Click **Save**.

Here is an example of a monitoring policy that contains a threshold and notification for an event source group.



Disabling Notifications

Notifications are sent when thresholds are not met. Additionally, automatic notifications are sent when baselines are not met. However, you may determine that you no longer require notifications for the event sources in a particular group. In this case, you can disable notifications for the event source group.

Note: Even if you disable all notifications, the details for alarms are still visible on the Alarms Tab.

Prerequisites

You must have configured thresholds and notifications for an event source group, and enabled them. For automatic notifications, you must have selected **Enable Notifications From Automatic Monitoring** on the Alarms Tab

Disable Notifications

To disable notifications (both manual and automatic) for an event source group:

1. Go to **(Admin) > Event Sources**.
2. Select the **Monitoring Policies** tab.
3. In the **Event Groups** panel, select a group.
4. Click **Enable** to remove the check mark. Clearing this option means that notifications are not sent for

this event source group, even if thresholds are not met or exceeded.

5. Additionally, you can remove all notifications. However, this is not required to stop the notifications.

Configuring Automatic Alerting

Note: Automatic alerting, and it settings, are currently in Beta testing.

Prerequisites

Before you set up notifications for an event source group, you should review the available notification items:

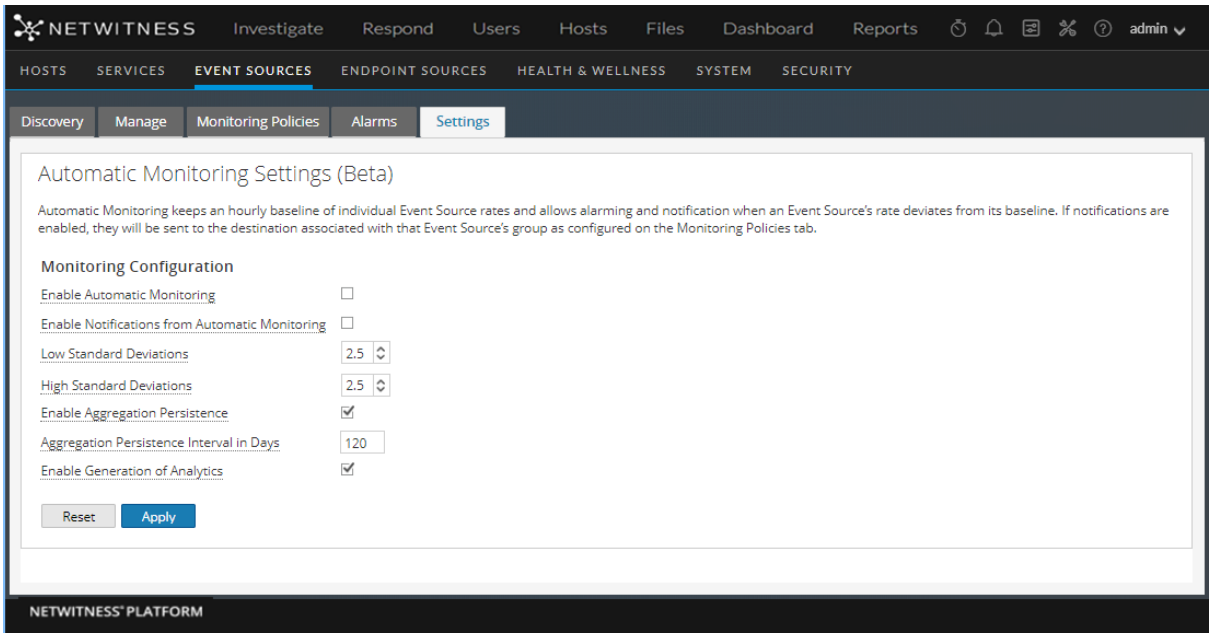
- **Notification Servers:** These are the servers that you want to receive notifications from the system. For more details, see the **Notification Servers Overview** topic in the *System Configuration Guide*.
- **Notification Templates:** These are the available templates for each type of notification. For Event Source Management, default templates are supplied for Email (SMTP), SNMP, and Syslog. You can use these templates as supplied, or customize them if necessary. For more details, see the **Templates Overview** topic in the *Systems Configuration* Guide.
- **Notification Output:** The outputs contain the parameters for the notification type. For example, an email notification type contains the email addresses and subject for the notification. For more details, see the **Notification Outputs Overview** topic in the *Systems Configuration* Guide.

Configure Automatic Alerting

To configure automatic alerting:

1. Go to  **(Admin) > Event Sources**.
2. Select the **Settings** tab.

   The Settings tab is displayed.



3. By default, automatic monitoring is turned on. To turn off automatic alerting, clear the **Enable Automatic Monitoring** option.
4. By default, notifications for automatic alerts is turned off. To turn on automatic notifications, select the **Enable Notifications From Automatic Monitoring** option.
5. Configure the parameters, based on your usage patterns:
   - **Low Standard Deviations**: standard deviations below which to receive alerts. Default is **2.5** (95% confidence).
   - **High Standard Deviations**: standard deviations above which to receive alerts. Default is **2.5** (95% confidence).

   Note: You can adjust the standard deviation settings in increments of 0.1 (one tenth) of a standard deviation.
6. Click **Save** to close the dialog and save your settings.

Viewing Event Source Alarms

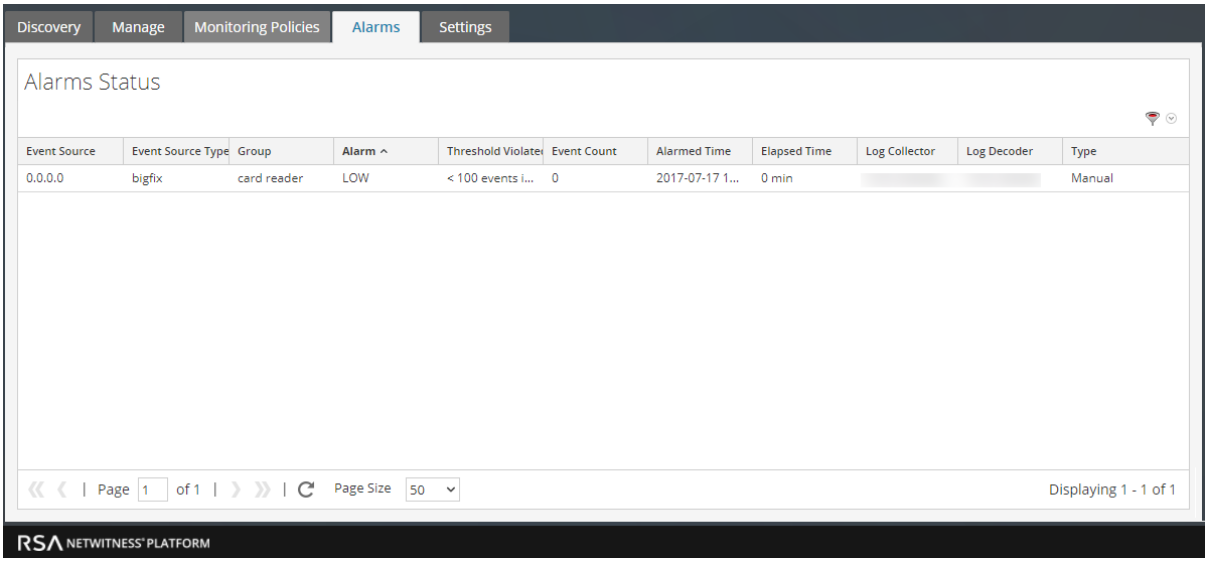
This topic describes how to view alarms for your event source groups. After you have configured and set alerts, you can view all of the generated alarms in the **Alarms** tab of the **Event Sources** view.

Sort the Alarms Information

When you first access this view, the data is sorted by most recent alarm (the Alarmed time column). You can sort by any column.

1. Go to  **(Admin) > Event Sources**.
2. Mouse over a column that you want to sort.
3. Click the Select the **Alarms** tab.
4. Mouse over the column that you want sorted, and click the  icon.

   This is an example when you mouse over the Alarm column.



5. Select either **Sort Ascending** or **Sort Descending** to sort the column in the way you wish.

   The data is sorted across all pages.

Note: You can also sort by two columns. To do this, first sort by the secondary column, then sort by the primary column. For example, if you want to see all the HIGH alarms by their group order, first sort on **Group**, then sort on **Alarm**.

Filter Alarms by Type

You can also filter the alarms by their type: you can display only the Manual or Automatic (baseline) alarms. To filter by alarm type, select the filter icon on the right side of screen, in the heading area: 

Select either Automatic or Manual:

- If you select Automatic, only the alerts based on baselines are displayed.
- If you select Manual, only the alarms for which you have set thresholds are displayed.

Event Source Management References

The following topics contain reference information for Event Source Management:

- Discovery Tab
- Manage Tab
- Manage Event Source Tab
- Event Sources View
- Create/Edit Group Form
- Details View
- Historical Graph for System Stats
- Manage Parser Mappings
- Alarms Tab
- Monitoring Policies Tab
- Settings Tab

Discovery Tab

To access the Discovery tab, go to NetWitness  (Admin) > **Event Sources**. The Discovery tab is displayed.

The Discovery tab lets you review the event source types that NetWitness has discovered for each address and the system's confidence of how likely it is that they were identified completely accurately. If the discovered event source types are correct, you can acknowledge to filter out that event source. If incorrect, you can set the allowed event source types for a particular address so that future logs will parse against the correct parsers.

Note: The following features apply to NetWitness version 11.1 and later:
- Acknowledging multiple event sources
- Filtering by event source type
- (for 11.2 and later) Mapping filter options include None, Auto, and Manual
- Mapping multiple event sources
- Searching for event sources on the Event Source Discovery page

NetWitness, version 11.2 and later, automatically maps incoming events to a type based on previous logs received from that address, reducing the mis-parsing of messages and reducing the number of items that need attention in the Discovery workflow. A value of **Auto** in the **Mapping Type** column indicates that an address has been auto-mapped.

Workflow

This workflow shows the overall process for configuring event sources.



What do you want to do?

Role	I want to...	Documentation
Administrator	Acknowledge and map event sources.*	Acknowledging and Mapping Event Sources
Administrator	Add and configure parser mappings for a Log Decoder.*	Manage Parser Mappings
Administrator	View event source alarms.	Viewing Event Source Alarms
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

*You can perform this task here.

Related Topics

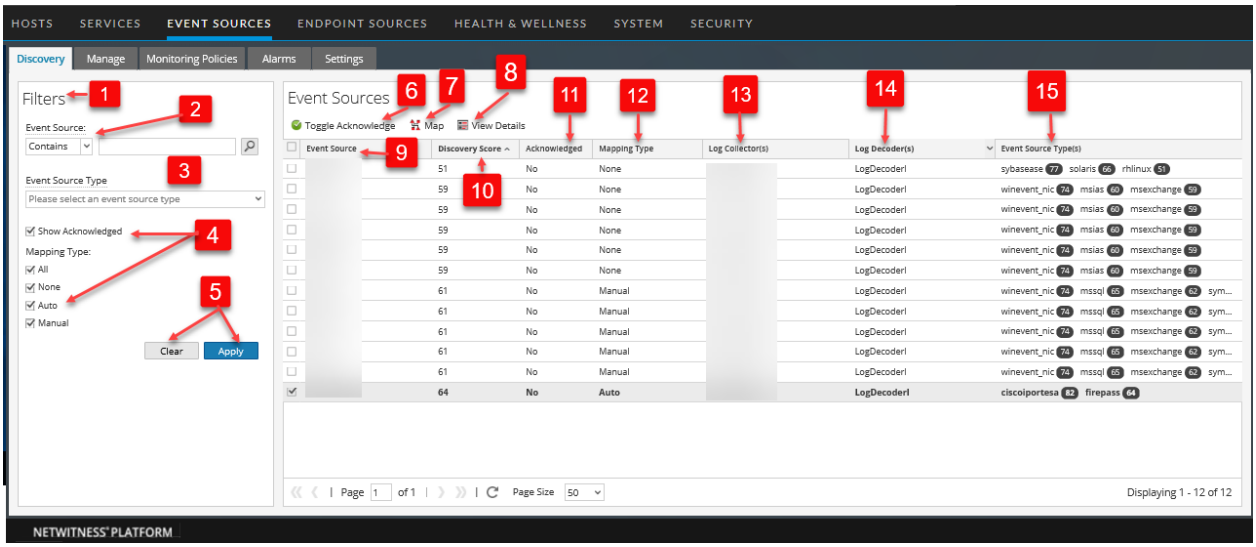
Manage Parser Mappings

Details View

Quick Look

The following example displays a list of addresses and their discovered Event Source types. The Event Source types display the Event Sources that have been discovered.

This is an example of the tab.



1 Displays the Filters and Event Sources panels with the Discovery tab open.

2 Displays the Event Source Filter field with a drop-down menu that offers the following options:

- Enter the full or partial address (IP, IPv6 or Hostname) of the source(s) you want to review. You can also enter multiple entries that are separated by commas.
  For example, **10.10.10.10,10.10.10.11,host1.company.com**
- **Exact:** Returns sources that completely match the search term.
  For example, **10.10.10.10** only returns **10.10.10.10**, not **10.10.10.101**.
- **Starts With:** Returns sources that start with the search term.
  For example,**10.10.10.** returns the whole **10.10.10.x subnet**.
- **Contains:** Returns sources that start with the search term.
  For example, **exch** returns all terms such as **us-exch-1.company.com**, or **lab21** returns all **hostx.lab21.company.com** terms.
- **Ends With:** Returns sources that end with the search term.
  For example, **lab21.company.com** returns all hosts.

Note: When specifying the search string, you can use **. - :** (period, dash, colon).

3 The Event Source Type drop-down menu filters for addresses containing all of the selected event source types.

4
- Select the **Show Acknowledged** checkbox to display acknowledged Event Sources.
- Mapping filter options can include just one of the mapping types listed in the Filter Panel, or multiple Mapping Types can be selected.

Note: If no mapping filter options are selected, the default is to display **All**, **None**, **Manual**, and **Auto** mapping types.

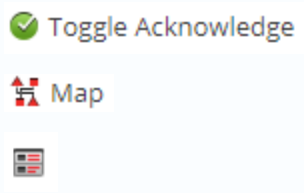
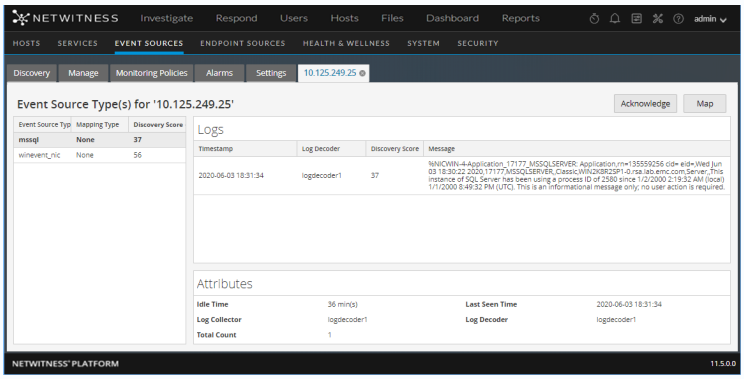
5
- The **Apply** button uses all criteria that is set in all filters.
- The **Clear** button clears all filters from the panel.

6 Toggles the event sources between acknowledged and not acknowledged states.

7	Maps the selected event sources.
8	View Details button to view details of the selected Event Source.
9	Displays the addresses of the selected Event Sources.
10	Displays the discovery scores of the selected Event Sources.
11	Displays whether or not the selected Event Sources have been acknowledged.
12	Displays the selected Event Source Mapping type as Auto, Manual, or None. Any changes to the mapping are only displayed here.
13	Displays the host names of the Log Collectors where the Event Sources are located.
14	Displays the host names of the Log Decoders where the Event sources are located.
15	Displays the discovered Event Source Types and their associated discovery scores.

Toolbar and Features

The Discovery tab contains the following features:

Field	Description
Tools 	The following items are available on the toolbar: • **Toggle Acknowledge**: Toggles the acknowledged state for the selected Event Source between **Yes** and **No**. • **Map**: Opens the Manage Parser Mappings dialog box, where you can map an event source to the correct log parser. • **View Details**: Provides details on the selected Event Source, as well as logs that have been received for this event source. For example: 
Event Source	The IP, IPv6, or Hostname of the Event Source.
Discovery Score	Displays the overall discovery score associated with that particular address. Higher scores indicate better confidence. Discovery scores range from 0 (least confident) to 100 (most confident).
Acknowledged	Selections are either **Yes** (you have acknowledged the Event Source) or **No** (you have not acknowledged the Event Source).

Field	Description
Mapping Type	Selections are **Manual** (you mapped the Event Source), **Auto** (the system automatically mapped the Event Source), or **None** (you have not mapped the Event Source). Auto mapping is content aware. When a log message is parsed to a high confidence header or message that has been tagged, an auto mapping will be set for that address and type. This auto-mapping is valid for 24 hours and will be renewed every time a log message matches a tagged header of a message. Log messages are first parsed against auto-mapped parsers, and only fall back to discovery if there is no match amongst the mapped parsers. Log messages that fall back to discovery can match the tagged headers or messages from other event sources: this results in multiple types being mapped. For example, an address could eventually be mapped to Windows, MS SQL, and Apache, and these parsers are evaluated first. If an event source is decommissioned, and its IP re-purposed, the 24-hour timer ages out the mappings for the decommissioned types. **Note:** This features applies to RSA NetWitness version 11.2 and later.
Log Collector(s)	Log Collectors that have received logs from this Event Source address.
Log Decoder(s)	Log Decoders that have received logs from this Event Source address.
Event Source Type(s)	The parsed type(s) of the Event Source address and the corresponding Discovery Score for each type.

Note: Discovery Scores are only available for 11.0 and above Log Decoders. Discovery Scores for pre-11.0 Log Decoders display as Unavailable.

The following table describes the sorting order for discovery scores. To access the Sorting Order drop-down menu, click on the down arrow in the Event Sources column.

Field	Description
Sort Ascending	Sort the column by discovery score in ascending order.
Sort Descending	Sort the column by discovery score in descending order.
Columns	Used to hide or show one or more columns.

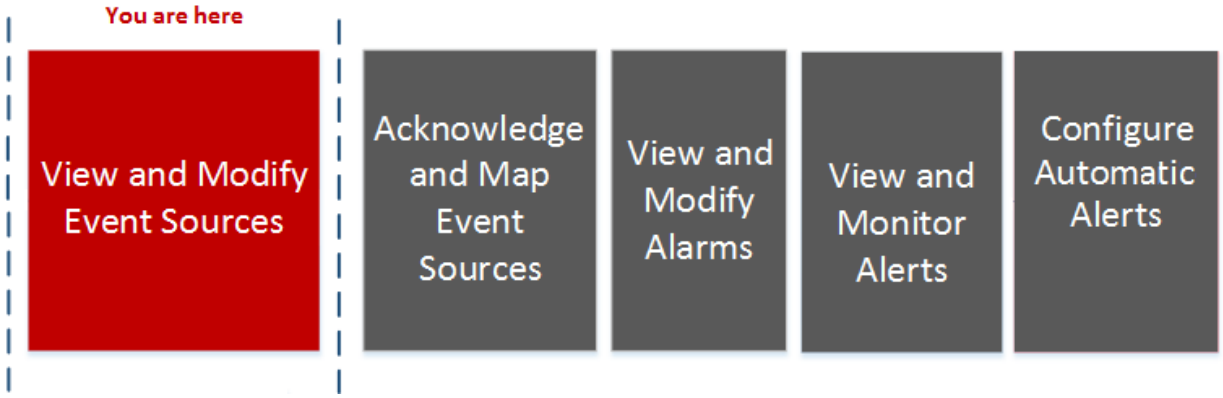
Manage Tab

The Manage tab organizes event sources into groups, and displays attributes for each event source.

To access this tab, go to (Admin) > **Event Sources** > **Manage**.

Workflow

This workflow shows the overall process for configuring event sources.



What do you want to do?

Role	I want to...	Documentation
Administrator	*View and modify event sources.	Managing Event Source Groups
Administrator	Acknowledge and map event sources.	Acknowledging and Mapping Event Sources
Administrator	Add and configure parser mappings for a Log Decoder	Manage Parser Mappings
Administrator	View event source alarms.	Viewing Event Source Alarms
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

*You can perform this task here.

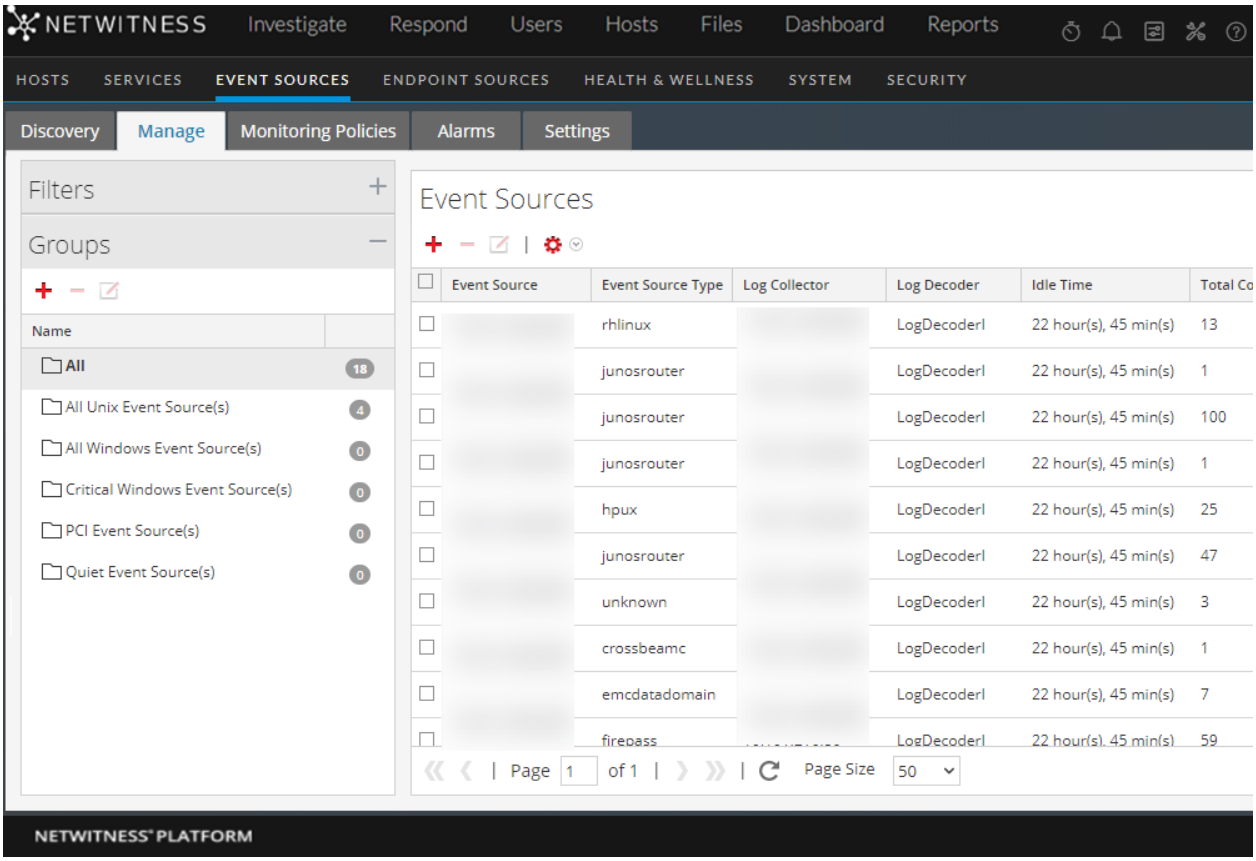
Related Topics

Creating Event Source Groups

Creating an Event Source and Editing Attributes

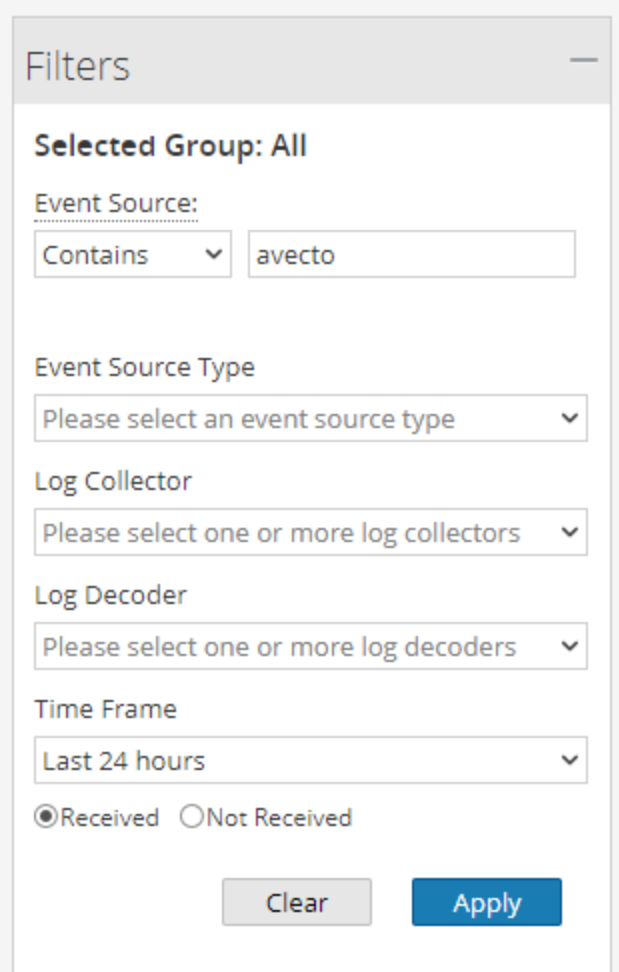
Quick Look

The Manage tab organizes event sources into groups, and displays attributes for each event source. The Manage tab consists of two panels, Groups and Event Sources.



Filter Panel

The Filter Panel provides options for filtering the set of event sources shown in the grid view. This is an example of the Filter Panel:



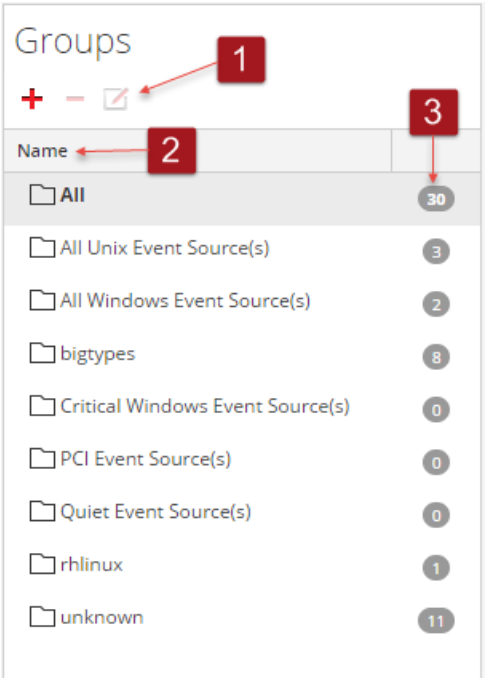
This panel provides the following options:

- You can filter on the event source name, using **Contains**, **Exact**, **Starts With**, or **Ends With**. Select one of these choices, then enter the corresponding string to match against.
- Select one or more Event Source Types to filter based on this value.
- To view event sources that send data to a particular Log Collector, select a Log Collector from the drop-down list.
- To view event sources that send data a particular Log Decoder, select a Log Decoder from the drop-down list.
- Select the time frame for when the events were collected. You can choose a value from 5 minutes to the previous 90 days, or all data that has been collected.
- Use the **Received** and **Not Received** radio buttons to filter the query results to contain only event sources that logs have been received from within the selected time, or query results to contain only event sources that logs have not been received from within the selected time.

After you complete the set of filters, click **Apply** to view the query results in the Event Sources grid.

Groups Panel

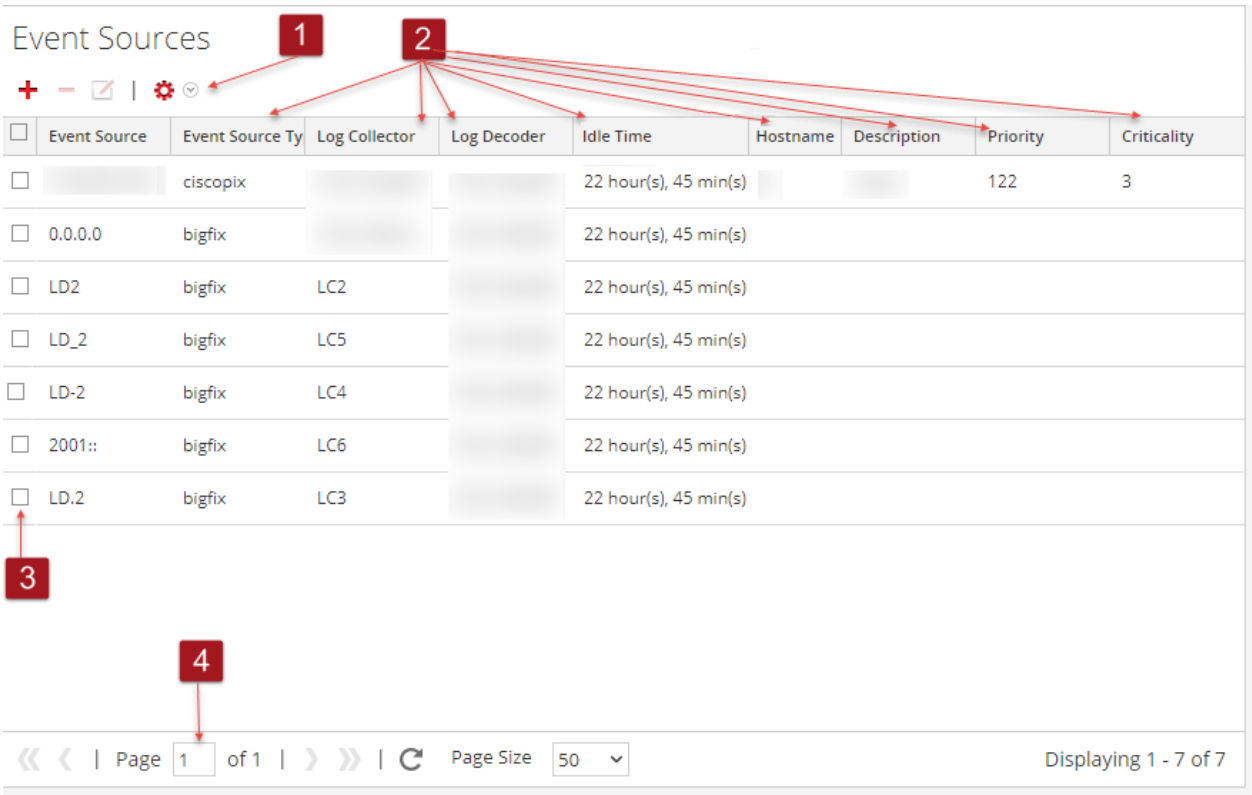
The Groups Panel lists the event source groups, as well as a count of the members for each group. To see all event sources, select **All** from the groups list. This is an example of the Groups panel.



1	Displays the standard NetWitness icons for adding, removing, or editing groups.
2	Lists the identifier for each group in the Name column. You can use the group names to quickly identify some of the criteria used to form the group. For example, if you create a group that consists of Windows event sources for the Sales organization, you could name the group **Windows Sales Sources**. **Note:** The event source group name is not editable. After you create a group, that name exists as long as the group itself.
3	The count for an event source group indicates the number of event sources in that group. That is, the number of event sources that match the criteria used to define the group. **Note:** The count is not dynamically updated when new event sources are added. Thus, you may need to refresh to see an updated group count.

Event Sources Panel

The Event Sources panel displays the attributes for the event sources in the selected group. Or, if All is selected in the Groups panel, the Event Sources panel lists all event sources.



1	The toolbar contains the following tools:
	• **Add**: manually add an event source
	• **Remove**: remove an event source
	• **Edit**: Update attributes for an existing event source
	• **Import / Export** menu: Displays a menu with the following options:
	  • **Import**: Import event sources from a Content Management Database (CMDB), spreadsheet, or other tool.
	  • **Export**: Export selected event sources and their attributes in CSV format.
	  • **Export Group**: Export the entire group that is currently selected.

2	Columnar display of attributes. You can choose which attributes to display.
3	Checkboxes: Select rows to use when performing tasks on multiple event sources, such as bulk editing.

4	Navigation Tools: At the bottom of the screen, there are items that help in navigating your group: • **Page *x* of *y***: indicates which page you are currently displaying, and how many total pages exist for this group. • **<<**, **<**, **>** and **>>**: click these icons to move between pages either one at a time (< and >) or to the first (<<) or last (>>) page. • **Page Size**: use this selector to choose your page size. • **Displaying *x* - *y* of *z***: quick check of which event sources are currently displayed out of the total number for the group.

Sorting

In the Event Sources panel, the list of items is presented in a sorted order. You can choose which column on which to sort. Note, however, that the sort order depends on capitalization.

For any string column, if the values contains a mix of lower case and upper case, the upper case appear in the list before the lower case values.

For example, assume the Event Source Type column contains the following entries: Netflow, APACHE, netwitnessspectrum, ciscoasa. The sort order would be as follows:

- APACHE
- Netflow
- ciscoasa
- netwitnessspectrum

Historical Graph for System Stats

Note: The information in this topic applies to NetWitness Version 11.4.1 and later.

Historical system statistics are available for each event source.

To access the Historical Graph for the System Stats:

1. Go to  **(Admin) > Event Sources**.
   The Event Sources view is displayed.
2. Click the **Manage** tab.
   The Event Sources Manage tab is displayed.
3. In the **Historical Graph** column, select the view chart icon () for any listed event source.
   The Historical graph for the selected event source is displayed. The figure displays event source

activity.



Parameters

You can customize the graph view as required. The table lists the various parameters used to customize the historical graph view.

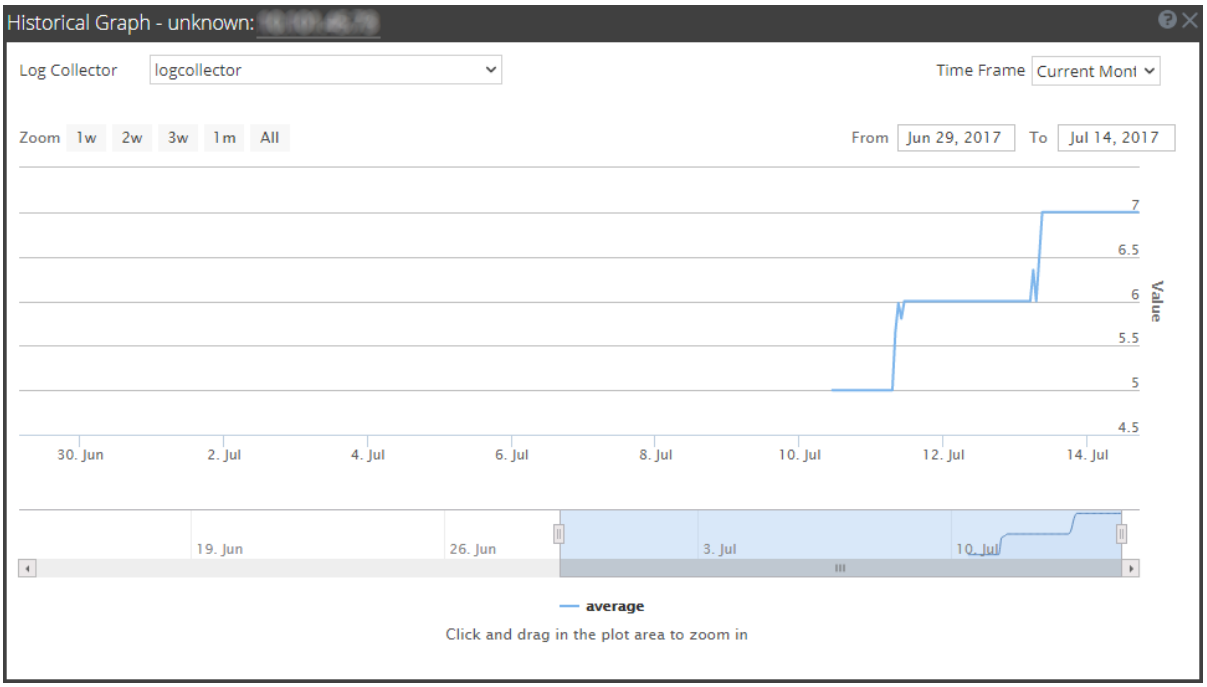
Parameter	Description
Log Collector	Select the Log Collector for which you want to view historical data. In some cases, the historical data for an event source may be spread across multiple Log Collectors.
Time Frame	Select the time frame for which you want to view the historical data. The available options are: **Current Day**, **Current Week**, **Current Month**, and **Current Year**.
From <date> To <date>	Select the date range for which you want to view the historical data.

You can zoom in for a detailed view of the data in the Historical graph.

Examples of Zooming In

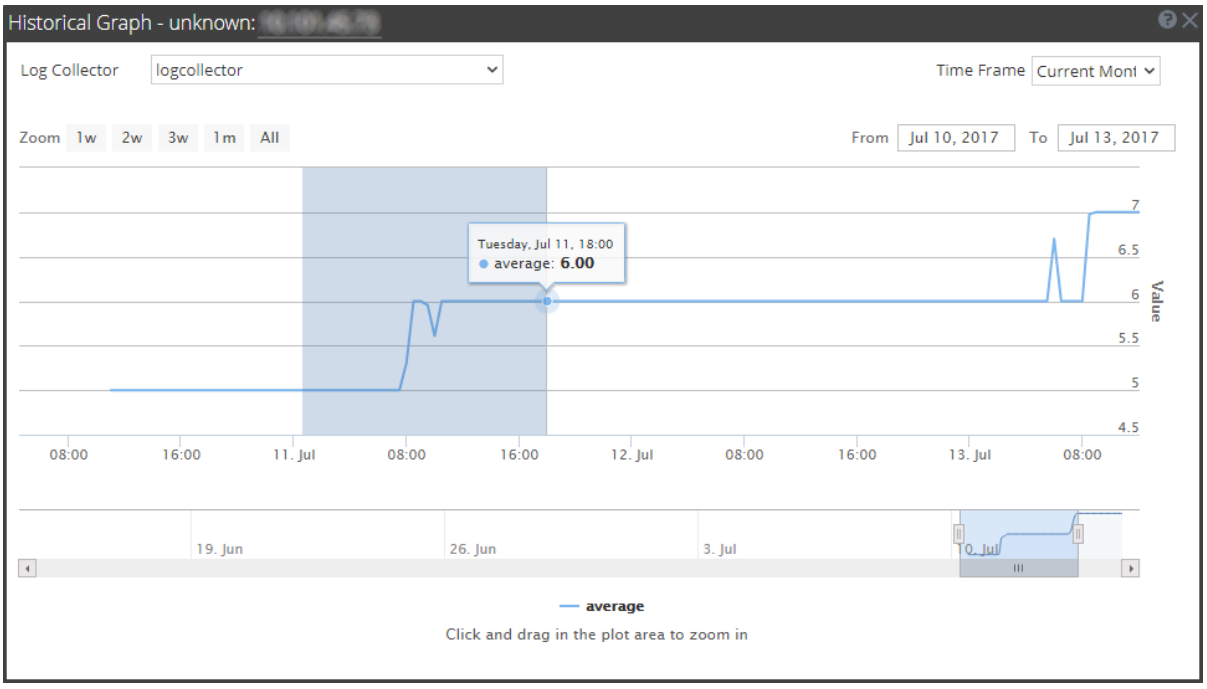
You can select one of the values to view the historical data for the selected value. The figure below displays an example for the 6 hour frame selected for zoom in. The slider bar at the right bottom corner is also changed to a 6 hour window.

Alternatively, you can slide the bar in the right hand corner to zoom in to a required frame.



You can click and drag in the plot area to zoom in for a required frame of time.

The figure below displays an example of how the graph appears while you click and drag.



Manage Event Source Tab

The Manage Event Source screen has several integrated components that present different perspectives of an event source.

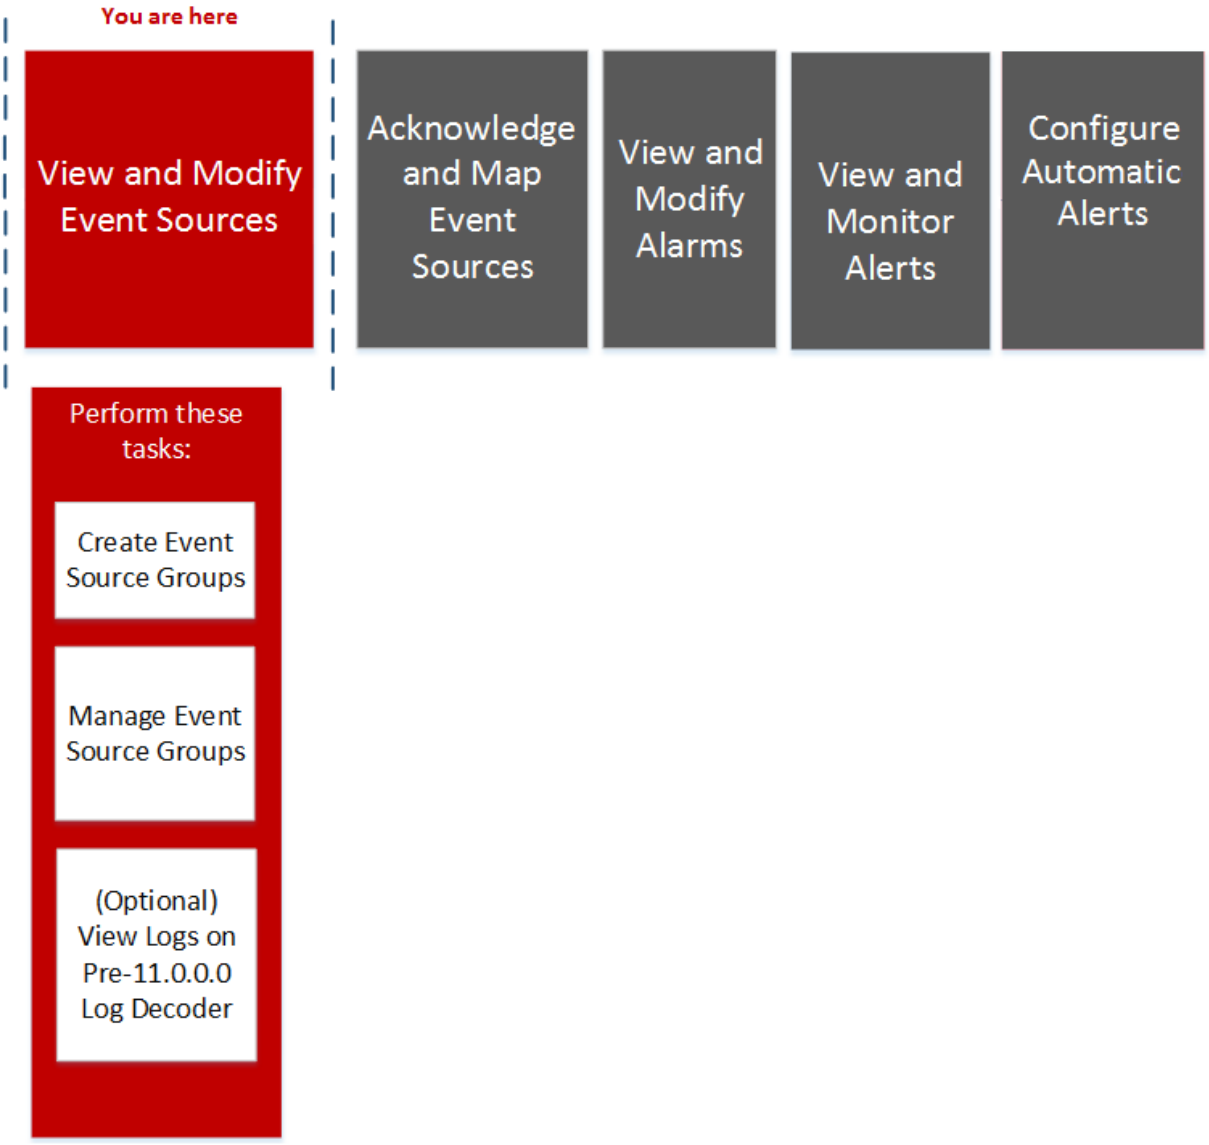
- Show Event Source Details
- Add attribute values to an event source
- Remove attribute values for an event source

To view the Manage Event Source screen for an event source:

1. Go to  **(Admin) > Event Sources**.
2. Select the **Manage** tab.
3. From the Event Sources pane, select an event source from the list and click + .

Workflow

This workflow shows the end-to-process for modifying, acknowledging, mapping, and configuring event sources, along with viewing and configuring event source alarms and alerts.



What do you want to do?

Role	I want to...	Documentation
Administrator	Create an event source group that contains all the high priority event sources.	Creating Event Source Groups
Administrator	Edit event source attributes.	Creating an Event Source and Editing Attributes

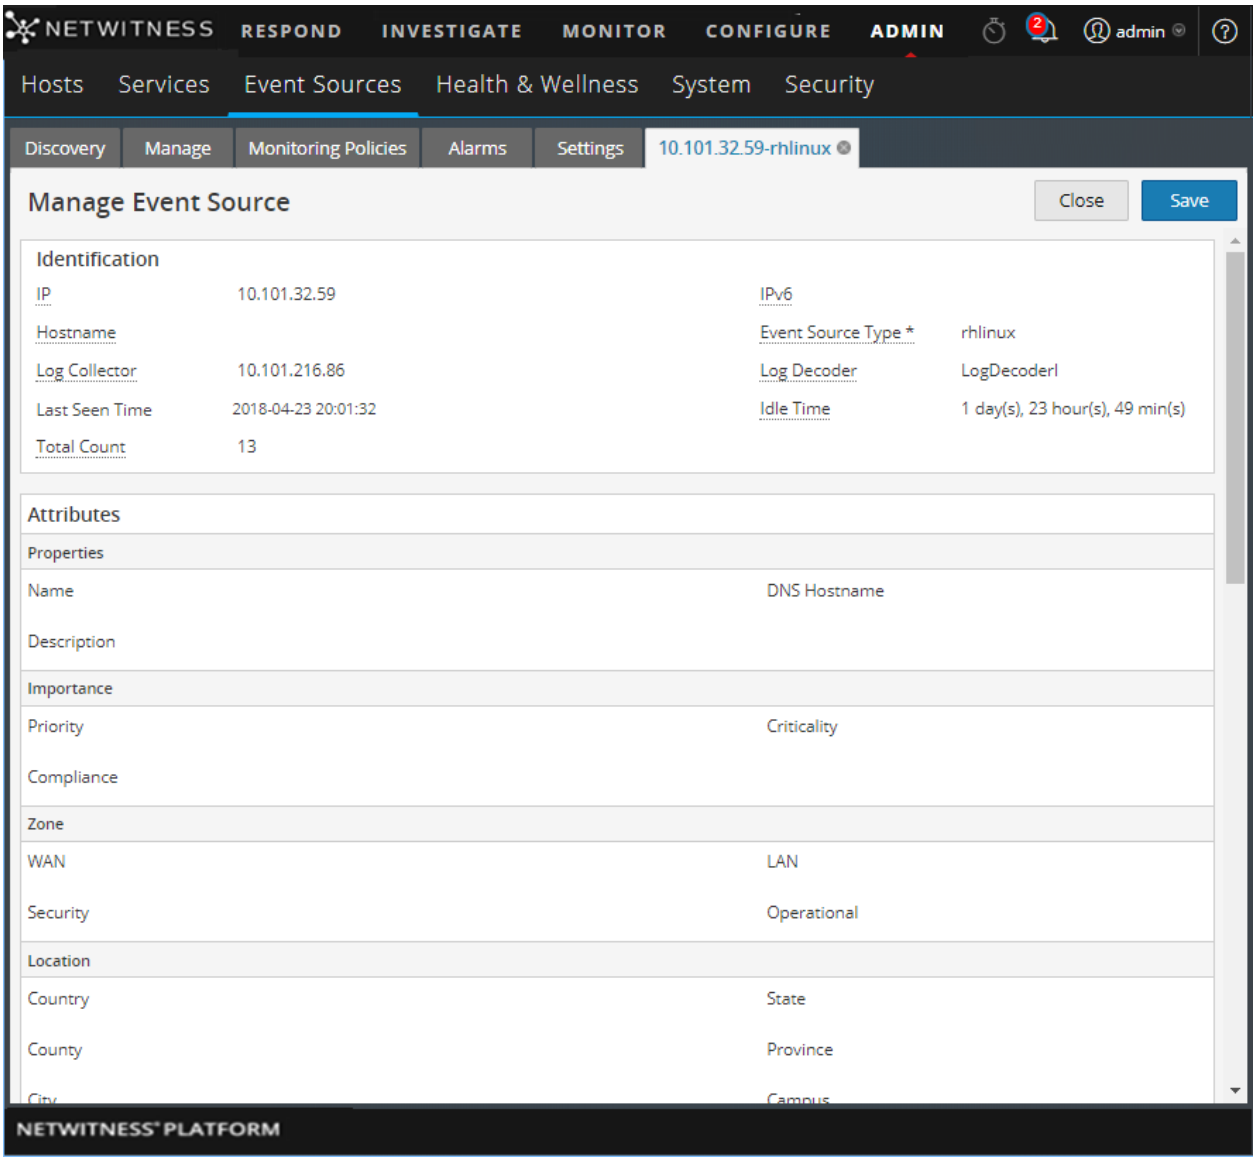
Related Topics

Creating Event Source Groups

Creating an Event Source and Editing Attributes

Quick Look

This is an example of the Event Source tab:



This table describes event source attribute categories.

Attribute Section	Description
Identification	These attributes are the main attributes that collectively identify an event source. You can only change these attributes when you are specifying the details for a new event source. For an existing event source, the attributes in this section are auto-populated, and cannot be changed while on this screen. Attributes available for a new event source: • IP • IPv6 • Hostname • Event Source Type • Log Collector • Log Decoder The following attributes are displayed when viewing the details for an existing event source: • Last Seen Time: this is the last time there was communication between NetWitness and the event source • Idle Time: this is the amount of time elapsed since the **Last Seen Time**. This time can be useful if you want to filter event sources that have been inactive for a certain duration. • Total Count: total count of all event sources for this **Event Source Type**.
Properties	These attributes provide the name and description. • Name • DNS Hostname • Description
Importance	These attributes can be used for grouping by priority. • Priority • Criticality • Compliance

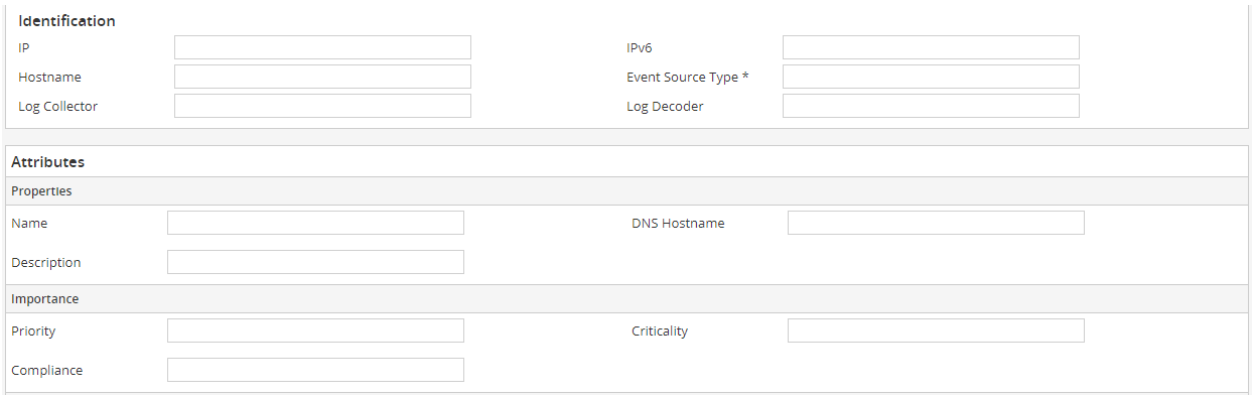
Attribute Section	Description
Zone	These attributes can be used for grouping by zone. • WAN (Wide Area Network) • LAN (Local Area Network) • Security • Operational
Location	These attributes can be used to group by the physical or geographical location. • Country • State • County • Province • City • Campus • Postal Code • Building • Floor • Room
Organization	These attributes can be used to group by organization, and also to provide contact information. • Company • Division • Business Unit • Department • Group • Contact • Contact Phone • Contact Email
Owner	These attributes specify those responsible for the event source. • Manager • Primary Administrator • Backup Administrator

Attribute Section	Description
Physical	These attributes specify the physical properties for the event source. • Vendor • Serial Number • Asset Tag • Voltage • UPS Protected • Rack Height • Depth • BTU Output • Color
Function	These attributes can be used to group by function. • Primary Role • Sub Role 1 • Sub Role 2
System Information	These attributes specify system information. • Domain Name • System Name • Identifier • System Description
Custom	This section provides eight custom attributes, for any other attributes that your organization might need.

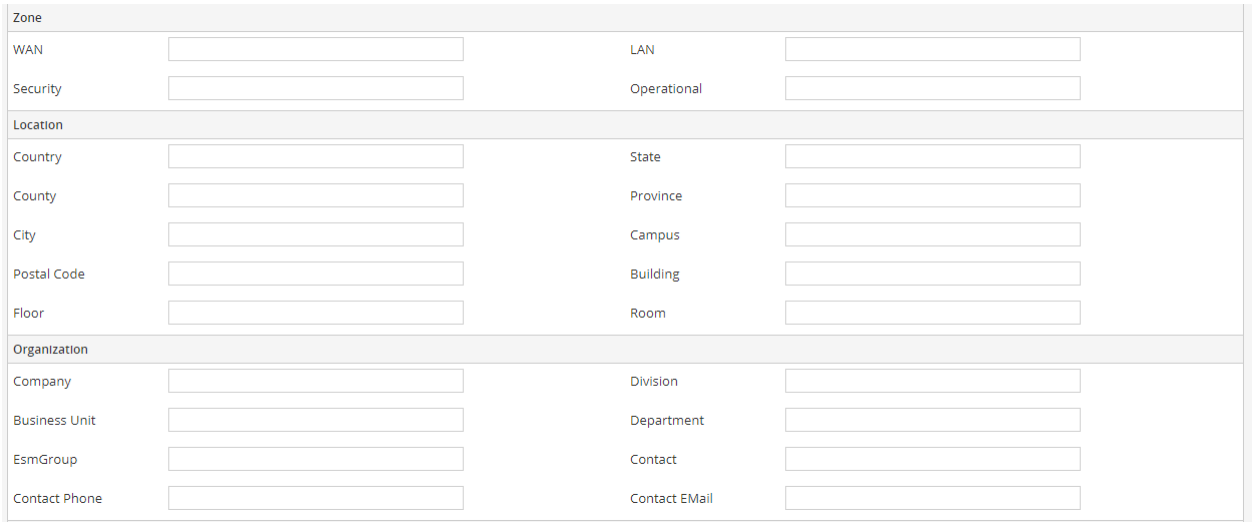
Features

The settings in the Manage Event Source tab are a combination of auto-populated and user-entered information. When an event source sends log information to NetWitness, it is added to the list of event sources, and some basic information is auto-populated. At any time after that, users can add or edit details for other event source attributes.

This figure shows an example of the **Identification**, **Properties**, and **Importance** sections.



This figure shows an example of the **Zone**, **Location**, and **Organization** sections.



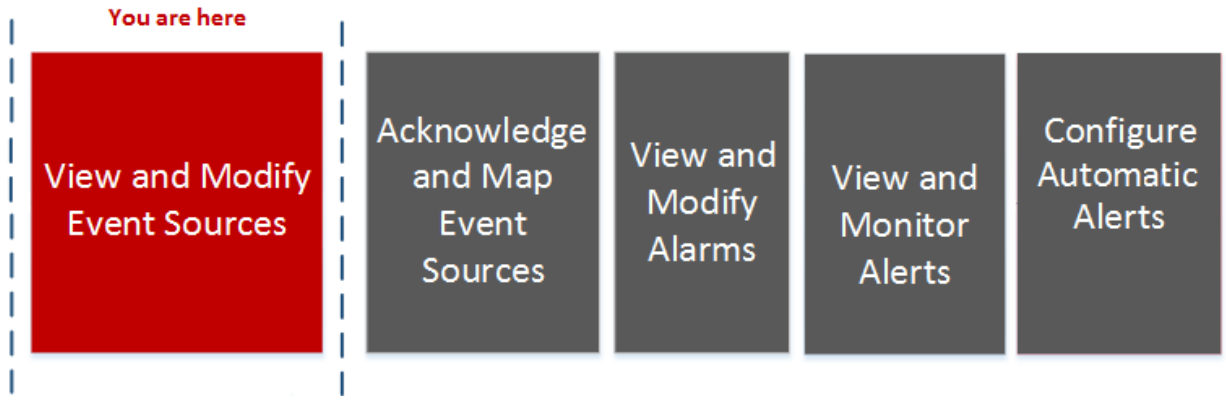
Event Sources View

The Event Source Attributes panel has the following tabs.

To access this panel, go to **(Admin) > Event Sources**.

Workflow

This workflow shows the end-to-process for modifying, acknowledging, mapping, and configuring event sources, along with viewing and configuring event source alarms and alerts.



What do you want to do?

Role	I want to...	Documentation
Administrator	Create an event source group.	Creating Event Source Groups
Administrator	Edit or delete an event source group.	Editing or Deleting Event Source Groups
Administrator	Edit event source attributes.	Creating an Event Source and Editing Attributes

Related Topics

Managing Event Source Groups

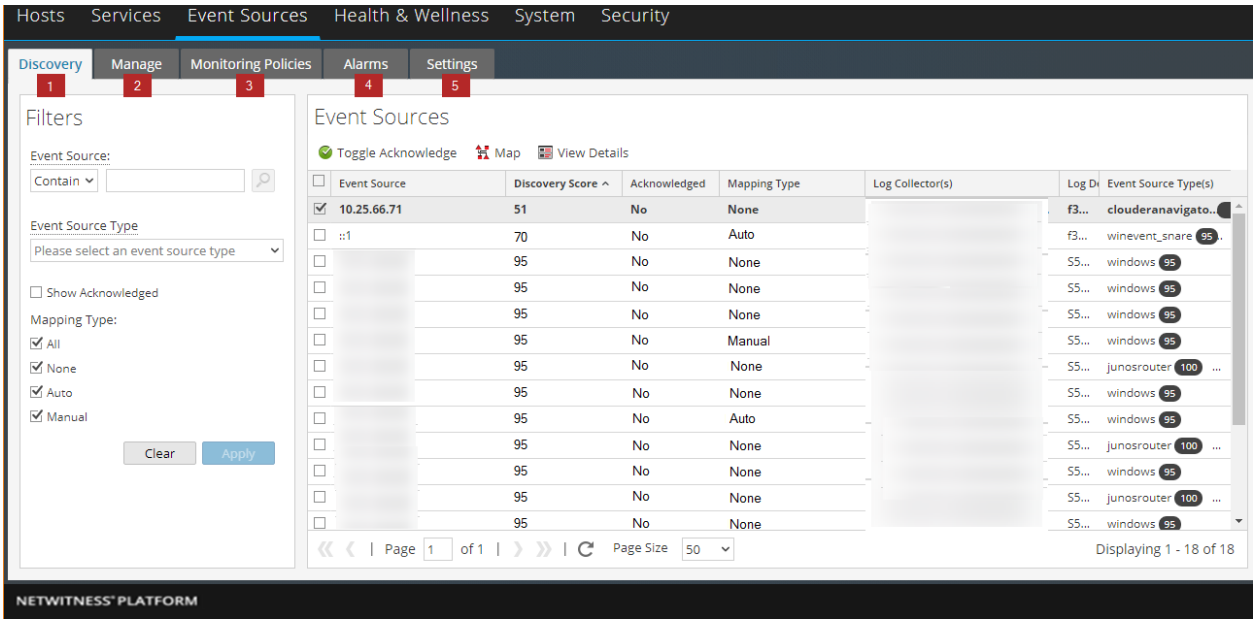
Creating Event Source Groups

Editing or Deleting Event Source Groups

Creating an Event Source and Editing Attributes

Quick Look

The Event Sources view presents the details for Event Sources that are discovered, acknowledged, or mapped by NetWitness.



1 [Discovery Tab]

Use this tab to review the event source types that NetWitness has discovered for each address and the system's confidence of how likely it is that they were identified accurately.

2 [Manage Tab]

Use this tab to create, edit, and delete Event Source Groups. It presents a customizable, searchable view of all of your event sources and groups.

3 [Monitoring Policies Tab]

Use this tab to manage alert configuration for event sources.

4 [Alarms Tab]

Use this tab to see the details of the alarms that have been generated.

5 [Settings Tab]

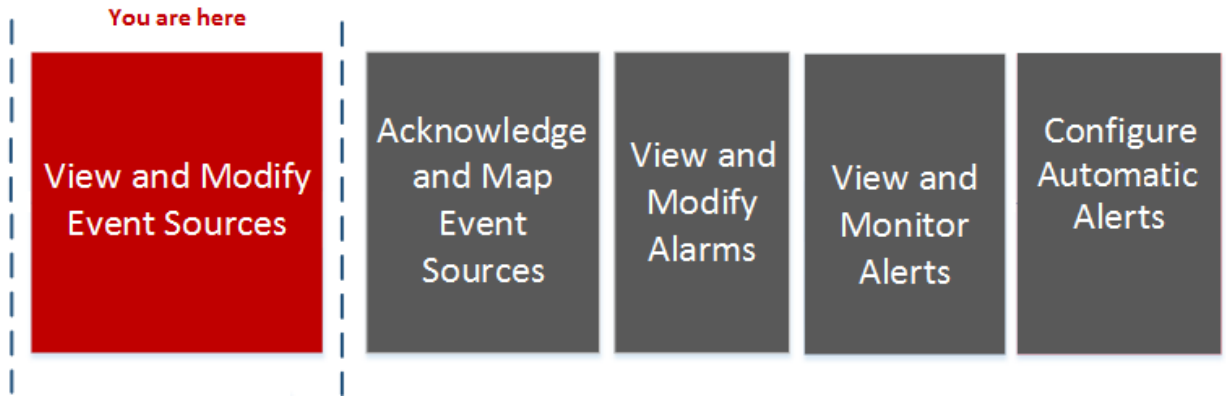
Use this tab to view or change the behavior for automatic (baseline) alerts.

Create/Edit Group Form

This Create Event Source Group form is displayed when you are creating or editing an Event Source Group.

Workflow

This workflow shows the overall process for configuring event sources.



What do you want to do?

Role	I want to...	Documentation
Administrator	*View and modify event sources.	Managing Event Source Groups
Administrator	Acknowledge and map events sources.	Acknowledging and Mapping Event Sources
Administrator	Add and configure parser mappings for a Log Decoder	Manage Parser Mappings
Administrator	View event source alarms.	Viewing Event Source Alarms
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

*You can perform this task here.

Related Topics

Creating Event Source Groups

Managing Event Source Groups

Details View

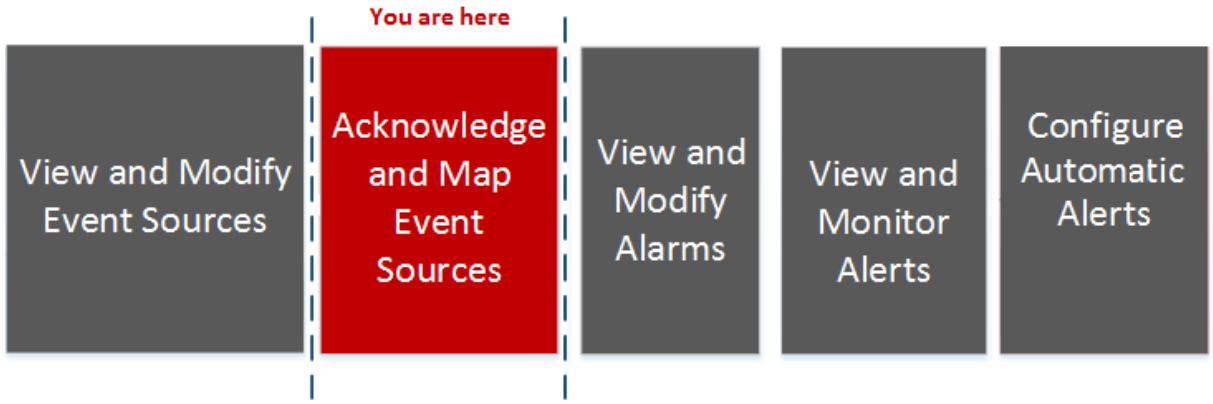
The **Details** view allows you to see details about the Event Source, as well as viewing a sample of the logs identified for each type in order to verify their accuracy.

You can access the **Details** view in a couple of ways.

- From the Toolbar, click the **View Details** button. Or, you can
- Double-click on the Event Source you selected.

Workflow

This workflow shows the overall process for configuring event sources.



What do you want to do?

Role	I want to...	Documentation
Administrator	View and modify event sources.	Managing Event Source Groups
Administrator	*Acknowledge and map events sources.	Acknowledging and Mapping Event Sources
Administrator	Add and configure parser mappings for a Log Decoder	Manage Parser Mappings
Administrator	View log parser details	Manage Parser Mappings
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

*You can perform this task here.

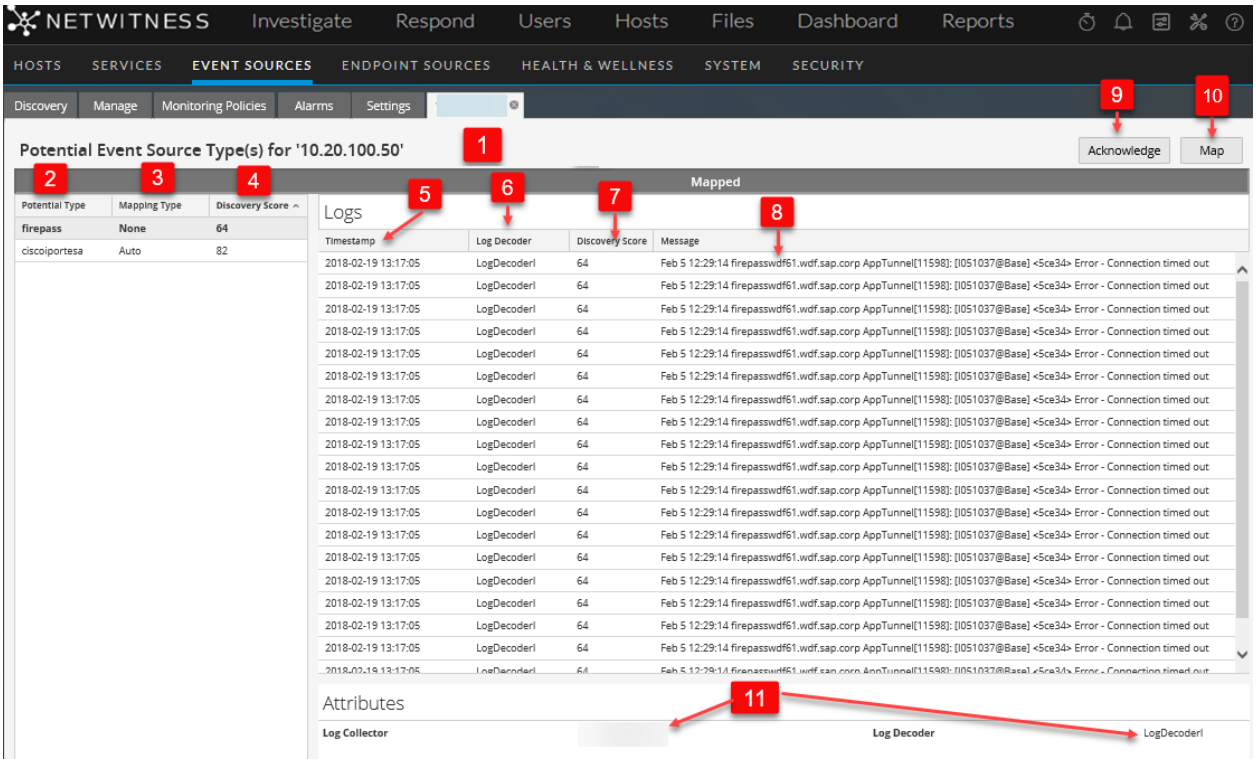
Related Topics

Viewing Logs from Pre-11.0 Log Decoder

Quick Look

The following example shows the discovery scores, event source types, logs, and attributes that correspond with the Event Source you selected in the Event Sources panel for a single Log Decoder.

Note: Device logs are only available for 11.0.0.0 and above Log Decoders.



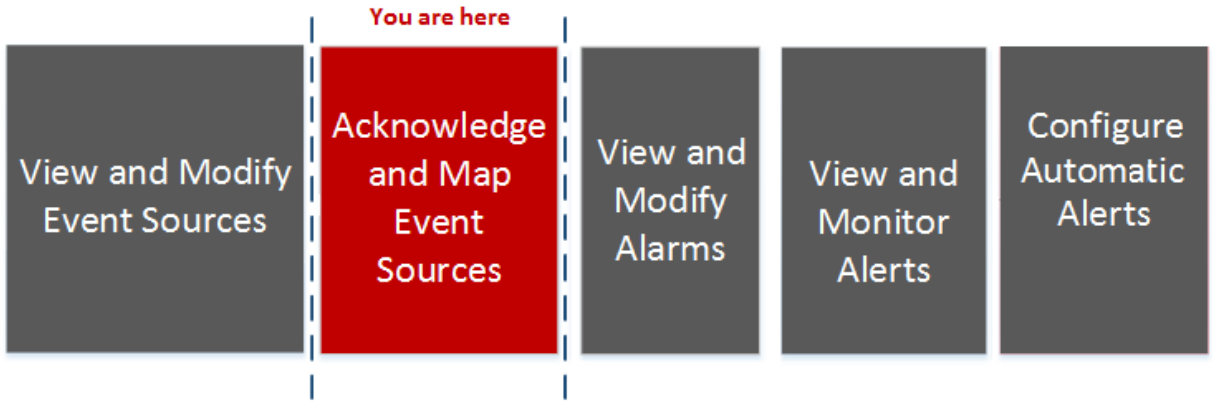
1	Displays the address of the selected Event Source.
2	Displays the potential type of the selected Event Source.
3	Displays the selected Event Source Mapping Type as Auto-Mapped, Manually Mapped, or None. Any changes to the Event Source Mapping are only displayed here.
4	Displays the discovery score for the selected Event Source type from least confident (0) to most confident (100).
5	Displays timestamps for the last few logs that have been parsed to the selected Event Source Type.
6	Displays the address of the Log Decoder that is parsing event sources.
7	Displays the discovery score of the corresponding log.
8	Displays logs for the selected Event Source type.
9	Allows you to acknowledge that all the discovered Event Source types are correct.
10	Allows you to set the appropriate parsers for selected Event Source addresses.
11	Displays the Event Source Management attributes for the selected Event Source Type.

Manage Parser Mappings

The **Manage Parser Mappings** dialog allows you to map the appropriate parsers for selected Event Source addresses. From the **Details** view, select the **Map** button.

Workflow

This workflow shows the overall process for configuring event sources.



What do you want to do?

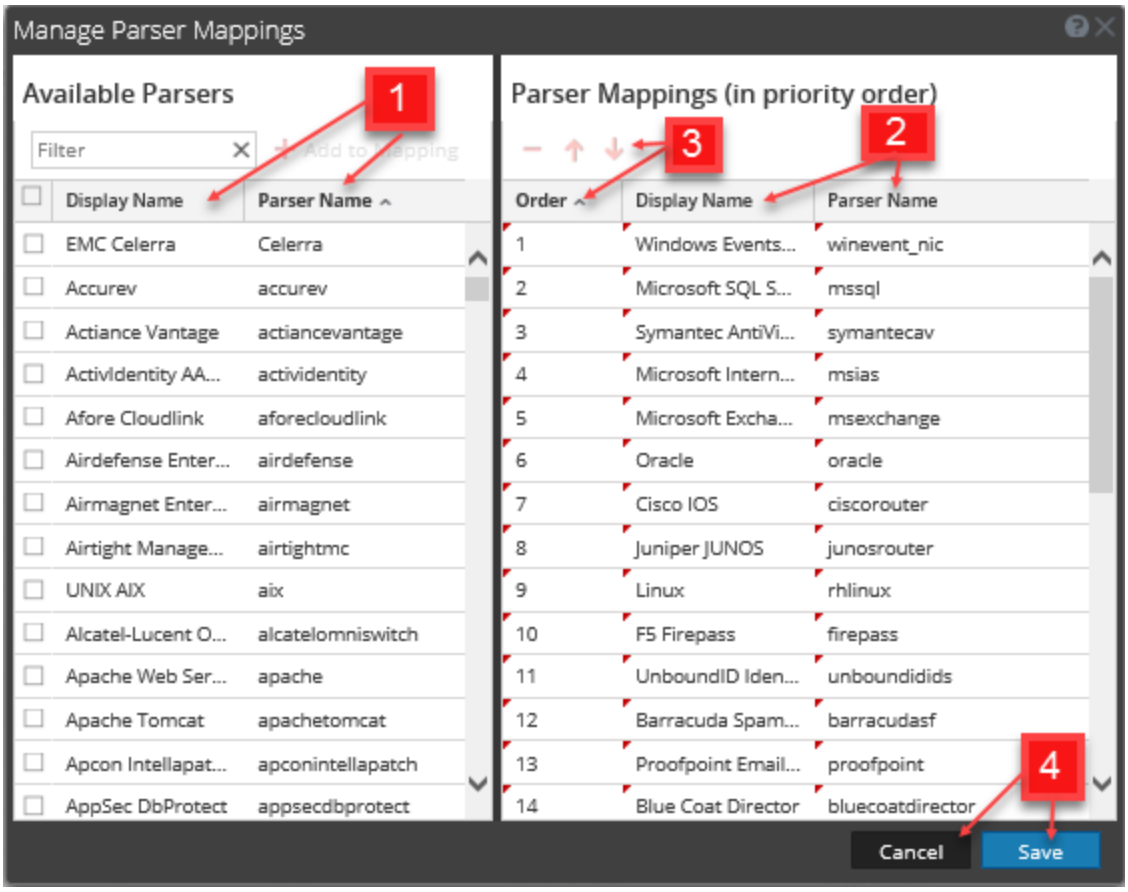
Role	I want to...	Documentation
Administrator	View and modify event sources.	Managing Event Source Groups
Administrator	Acknowledge and map events sources.	Acknowledging and Mapping Event Sources
Administrator	*Add and configure parser mappings for a Log Decoder	Manually Map Event Source Types
Administrator	View event source alarms.	Viewing Event Source Alarms
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

*You can perform this task here.

Related Topics

Viewing Logs from Pre-11.0 Log Decoder

Quick Look



Note: For event sources that were created manually, the Manage Parser Mappings window has an empty Display Name in the Log Parsers column. To view the missing display names, close the Manage Parser Mappings dialog box, and then reopen it.

1	Displays all the available parsers that you can map based on the event sources that you selected from the **Discovery** view. Also displays the mappings that are already present in the Log Decoders for the selected event source or the parsers that have been discovered. To filter your available parsers, type the first few letters of the parser name that you want to map. Click the **Add to Mapping** button to add the parser to the parser mappings listed in the right panel. You need to select parsers before the **Add toMapping** button is enabled. Add the selected parser by clicking the **Add to Mapping** button in the right panel. You can rearrange the order of the parser mappings using the up and down arrow keys and you can also drag and drop selected parser mappings. You can select multiple mappings by pressing the **Ctrl** key.
2	Displays the names of the selected parsers that you want to map.
3	Displays the order of the selected parser mappings. You can delete parser mappings by selecting the minus sign ( ). Press the **Ctrl** key to select multiple mappings to perform group operations on them.

4 Click **Save** to save your mappings to all the Log Decoders. A pop-up message informs you that your mappings are successfully saved. When the window is closed, the banner on the **Details** tab is updated to reflect the status. If mapped, the text displayed is **Mapped**.
Click **Cancel** to return to the **Details** tab.

Best Practices

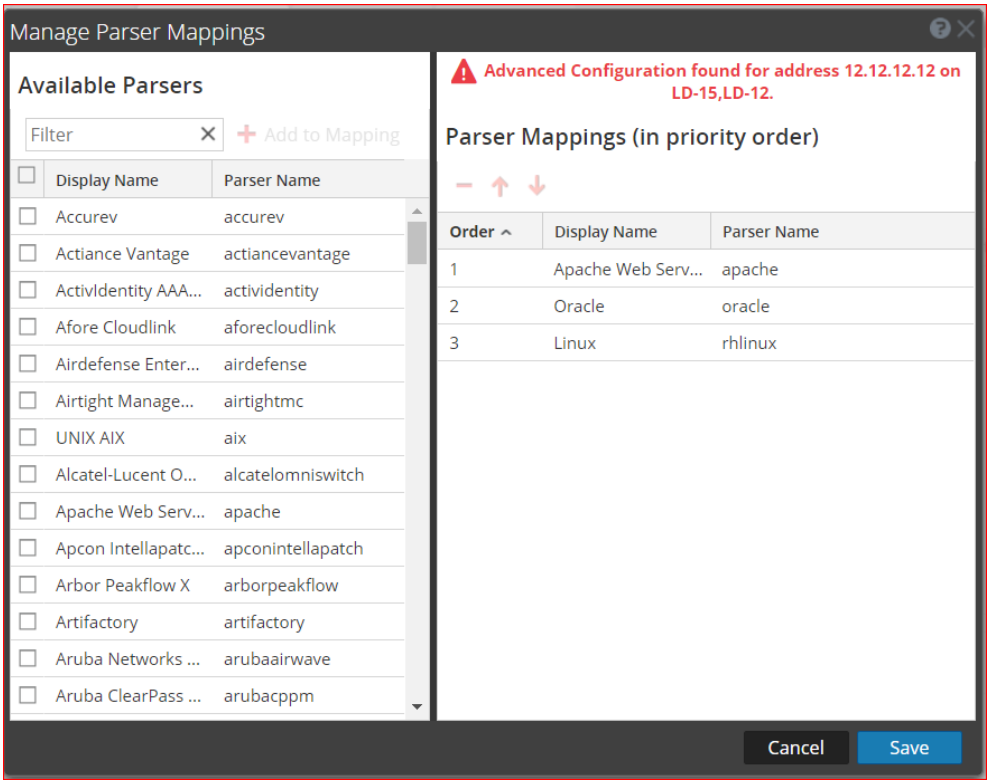
When mapping multiple device types from the same event source, assign the highest priority to the strictest log parser. Conversely, a log parser with generic headers should be lowest in priority. The CEF log parser is an example of a strict log parser.

Advanced Configuration

Mapping configurations with the Log Collector are not displayed in the Parser Mappings window. If the mapping is saved, it is saved for the corresponding IP address, not for the corresponding Log Collector entry. If no mappings are found for the corresponding IP address, the discovered event source types are displayed in the Parser Mappings window.

If advanced Log Decoder configurations are discovered, a message similar to the one below displays in the Manage Parser Mappings dialog.

Note: If you want to edit the advanced configuration, you need to navigate to the Log Decoder service's parser mappings configuration.



Alarms Tab

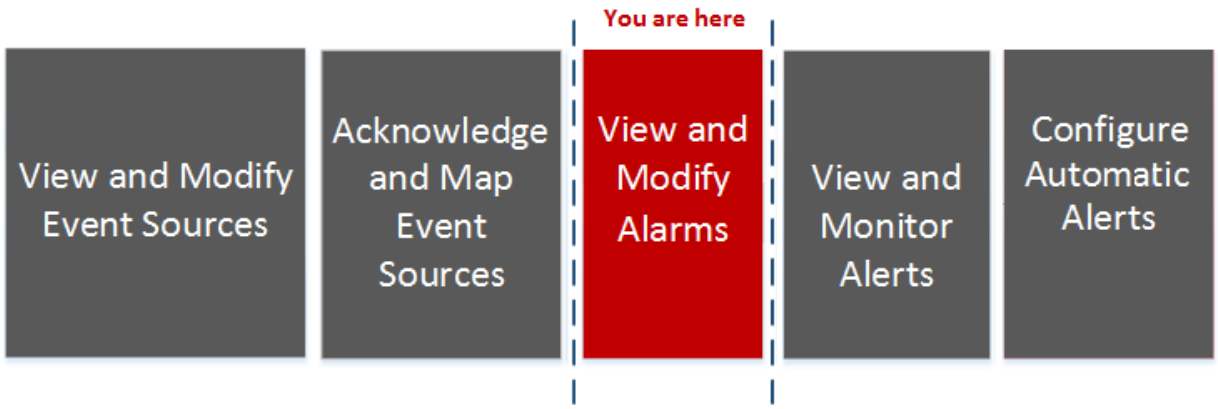
From the Alarms tab you can view details of the alarms that have been generated.

The Alarms tab has one panel that displays Alarm status.

To access this tab, go to (Admin) > **Event Sources** > **Alarms**.

Workflow

This workflow shows the overall process for configuring event sources. It also shows where configuring alarms and alerts settings are located in the process.



What do you want to do?

Role	I want to...	Documentation
Administrator	View and modify event sources.	Managing Event Source Groups
Administrator	Acknowledge and map events sources.	Acknowledging and Mapping Event Sources
Administrator	Add and configure parser mappings for a Log Decoder	Manage Parser Mappings
Administrator	*View event source alarms.	Viewing Event Source Alarms
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

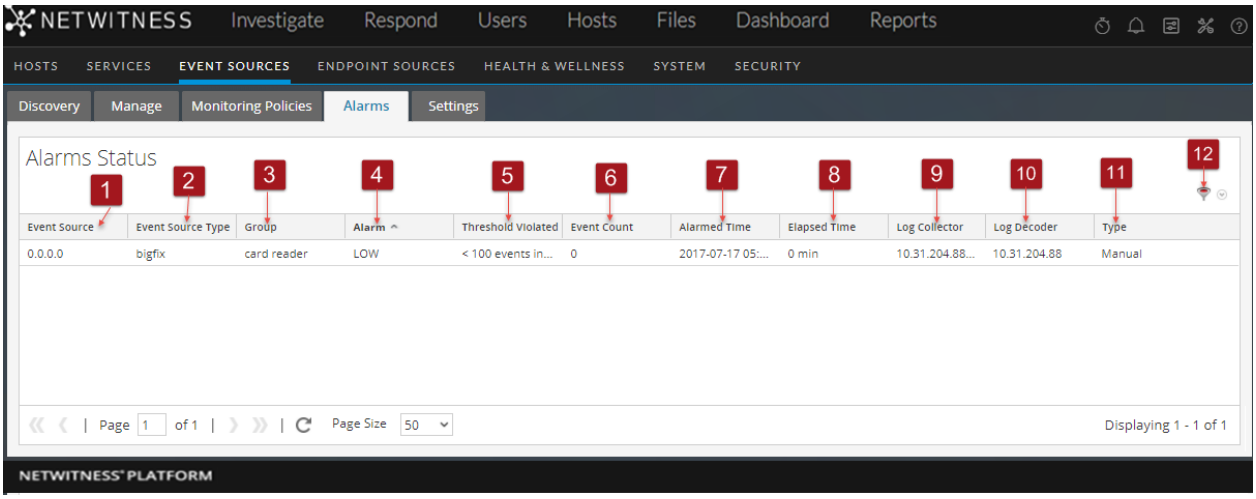
*You can perform this task here.

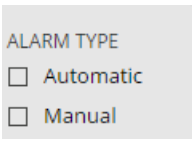
Related Topics

Configuring Automatic Alerting

Quick Look

The Alarms tab presents the details for Event Sources that are currently in violation of a policy and threshold. Only Event Sources in violation of a policy appear in the list. After the event source returns to a normal state, the corresponding alarm disappears from the list.



1	Displays the IP, IPv6, or Hostname of the event source that is alarmed.
2	Displays the type of the alarmed event source. For example, **winevent_nic** (for Microsoft Windows) or **rhlinux** (for Linux).
3	Displays the event source group that contains the event source for which the alarm has been triggered.
4	Displays the type of threshold that was triggered: **High** or **Low**
5	Displays the conditions of the threshold that was triggered. For example: `5,000,000 events in 5 minutes`
6	Displays the number of events in the threshold time period causing the alarm.
7	Displays the initial time the event source went into an alarmed state. **Note:** When you first access this view, the data is sorted by this column (most recent alarm first).
8	Displays the elapsed time since the event source entered an alarmed state.
9	Displays the Log Collector last collecting from this event source.
10	Displays the Log Decoder last receiving from this event source.
11	Displays the alarm type. Alarm type is either **Manual** or **Automatic**: • **Manual**: these are alarms that violate the configured threshold policy. • **Automatic**: these are alarms that deviate from the baseline for the alarmed event source.
12	Select the **Filter** icon to display the **Filter** menu:  Select either **Automatic** or **Manual**: • If you select **Automatic**, only the alerts that are based on baselines are displayed.

- If you select **Manual**, only the alarms for which you have set thresholds are displayed.

Note: You can hide or show columns by right-clicking in the table header and choosing **Columns** from the drop-down menu. Select a column to display it, or clear the column to hide it.

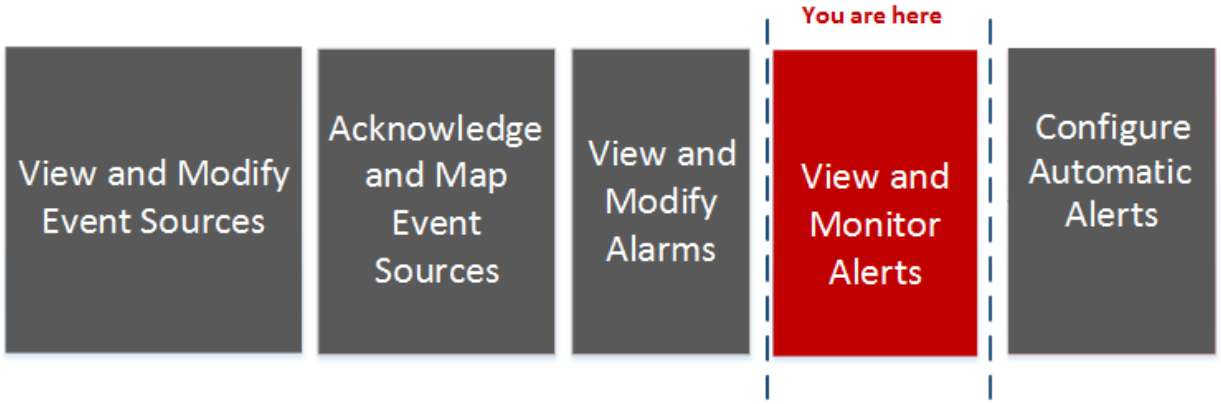
Monitoring Policies Tab

The Monitoring Policies tab organizes thresholds by event source group.

To access this tab, go to (Admin) > **Event Sources** > **Monitoring Policies**.

Workflow

This workflow shows the overall process for configuring event sources.



What do you want to do?

Role	I want to...	Documentation
Administrator	View and modify event sources.	Managing Event Source Groups
Administrator	Acknowledge and map events sources.	Acknowledging and Mapping Event Sources
Administrator	Add and configure parser mappings for a Log Decoder	Manage Parser Mappings
Administrator	View event source alarms.	Viewing Event Source Alarms
Administrator	*View Monitoring Policies.	Monitoring Policies

Role	I want to...	Documentation
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

*You can perform this task here.

Related Topics

Setting Up Notifications

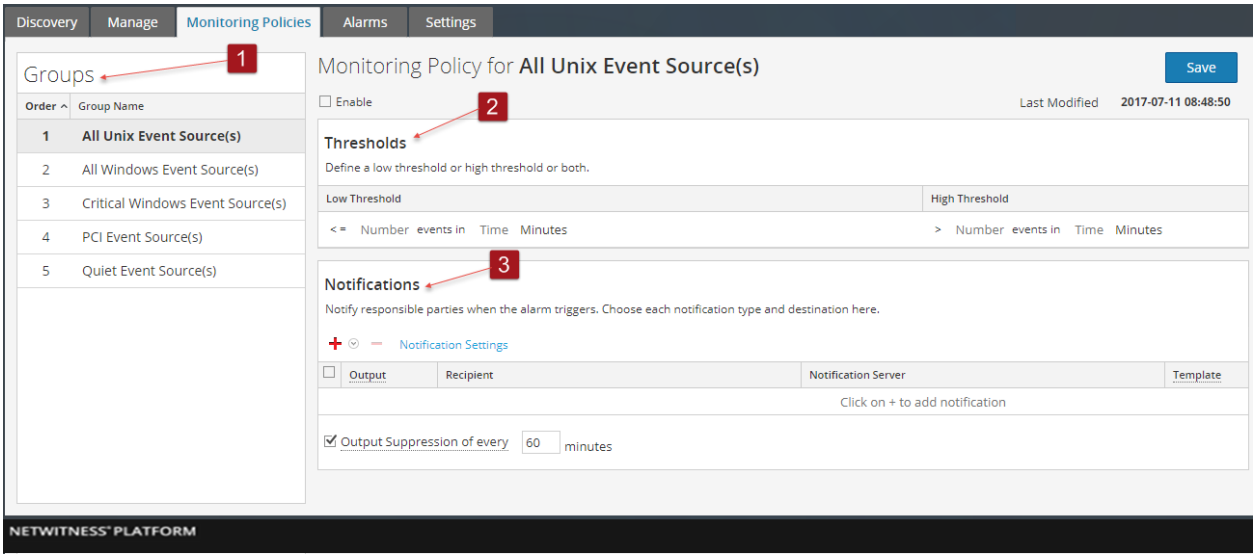
Disabling Notifications

Quick Look

The **Monitoring Policies** tab consists of three panels:

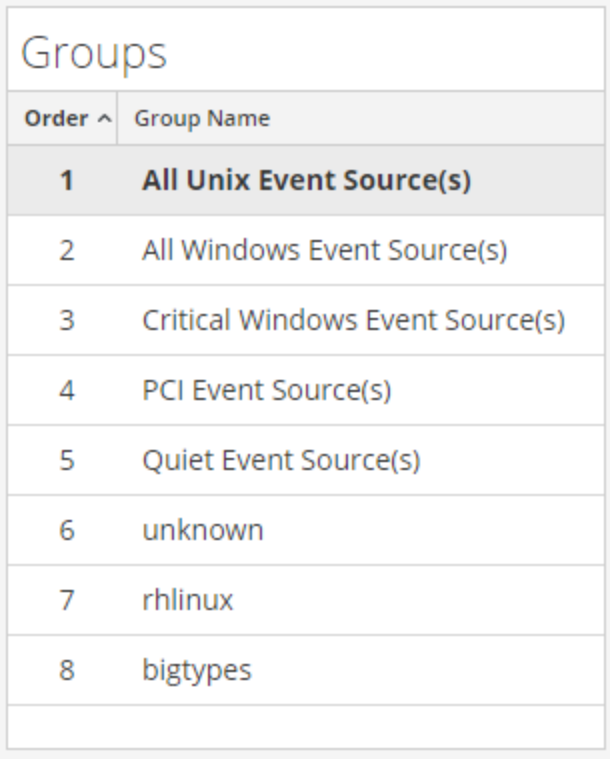
- Event Groups Panel
- Thresholds Panel
- Notifications Panel

This is an example of the **Monitoring Policies** tab.



1 Displays the Groups panel.
2 Displays the Thresholds panel.
3 Displays the Notifications panel.

Event Groups Panel



The group selected in this panel determines which thresholds appear in the Thresholds panel. You can define a set of thresholds for each event source group. Notice that the groups are listed in a specific order:

- Drag and drop groups to change the specified order.
- The higher a group is listed, the higher the precedence for that group's thresholds: RSA NetWitness checks the thresholds in the order provided in this panel. Thus, your highest priority groups should be at the top of this list

Thresholds Panel

This is an example of the Thresholds panel for an event source group.

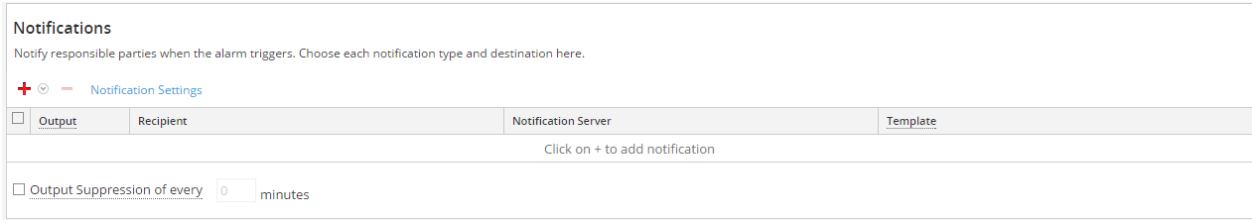


The Thresholds Panel contains the following features.

Feature	Description
Enable	The Enable checkbox designates whether or not the thresholds that you define for a group are enabled. If so, notifications are sent whenever the thresholds for that group are outside of the defined range. If not, then no monitoring of that event source group is occurring. **Note:** If you configure a threshold and attempt to save the page without enabling it, you receive a confirmation message, asking you whether or not to enable the policy. If you enable a policy, but do not have any thresholds set, then you can still receive automatic (baseline) notifications, as long as you have turned on automatic notifications. See below for more details on the look of notifications.
Low number of events Low number of minutes or hours	This is the low end of the threshold. Enter the fewest number of events and the time range. If the event source group receives fewer messages than specified here, the threshold is not met, and notifications are sent.
High number of events High number of minutes or hours	Works similarly as for the low values: If more messages than specified here are received, the threshold is not met, and notifications are sent.
Last Modified date and time	This field indicates the last time and date that the thresholds were changed.
Save	Saves the changes you have made to the thresholds.

Notifications Panel

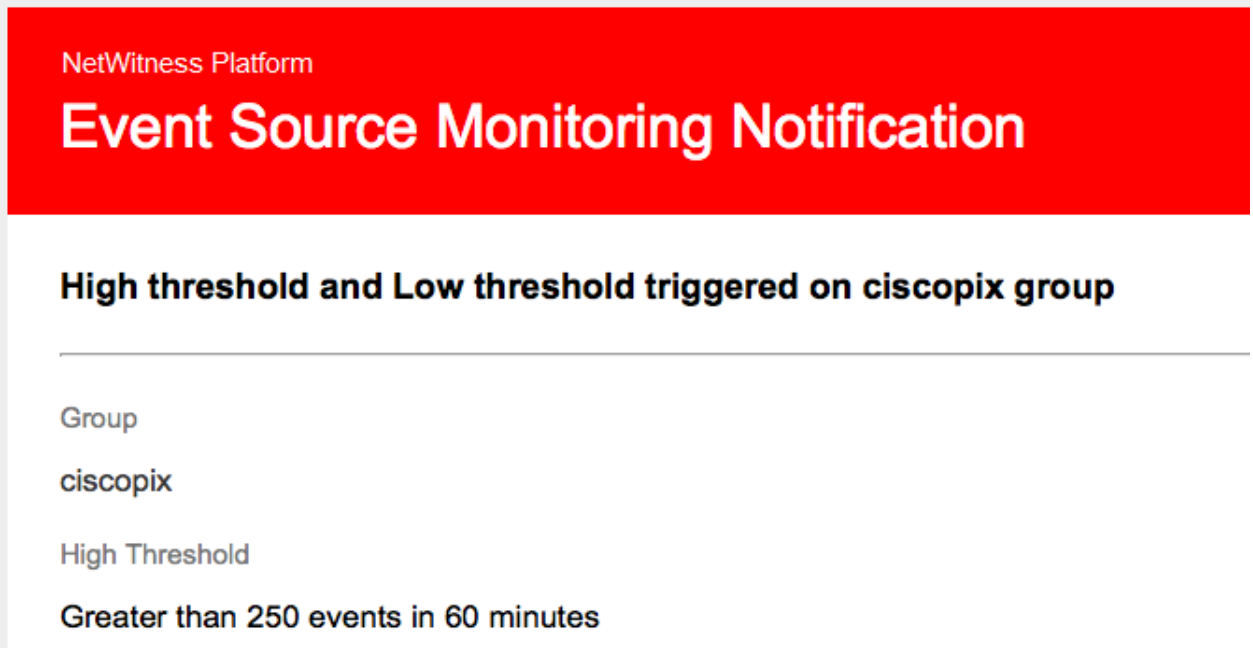
This is an example of the Notifications panel for an event source group.



The following table describes the fields on the Notifications panel

Field	Description
Tools + -	The following items are available on the toolbar: • **Add (+)**: clicking the **Add** presents a menu where you can choose the type of the notification • **Remove (-)**: removes the selected row from the list.
Notification Settings	Clicking this link opens a new browser tab, and takes you to the **Admin** > **System** > **Notifications** page in NetWitness.
Type	Displays the type of the notification that you have chosen. The available options are as follows: • Email • SNMP • Syslog
Notification	See the **Configure Notification Outputs** topic in the *System Configuration Guide* for more details.
Notification Server	See the **Configure Notification Servers** topic in the *System Configuration Guide* for more details
Template	For Event Source Management, RSA provides three out-of-the-box templates for notifications. You can use the following templates as delivered, or customize them based on the needs of your organization: • **Email template**: sends notifications to the specified email addresses. • **SNMP template**: sends notifications to the specified SNMP server • **Syslog** template: sends notifications to the specified Syslog server. See the **Configure /Templates for Notifications** topic in the *System Configuration Guide* for more details.
Output Suppression	Use this item to limit how often notifications are received for this policy, in case a lot of alerts are triggered in a short period of time.

The following are sample notifications, based on the supplied Templates.



- Email:

  For email notifications, the third column, **Alarm Type**, specifies whether the triggered alarm was based on a user threshold, or the baseline data being out of normal bounds. If you have automatic monitoring or notifications turned off, you will not receive any **Automatic** notifications. The same is true for Syslog and SNMP, except those notifications are formatted differently.

- SNMP trap:

```
11-11-2015 11:57:33 Local7.Debug 127.0.0.1 community=public,
enterprise=1.3.6.1.4.1.36807.1.20.1, uptime=104313, agent_
ip=10.251.37.92, version=Ver2, 1.3.6.1.4.1.36807.1.20.1="NetWitness
Event Source Monitoring Notification:
Group: PCI Event Source(s)
High Threshold:
Greater than 500 events in 5 minutes
10.17.0.10,ciscopix,Manual
10.17.0.13,ciscopix,Manual
10.17.0.8,ciscopix,Manual
10.17.0.8,ciscopix,Automatic
10.17.0.12,ciscopix,Manual
10.17.0.5,ciscopix,Manual
10.17.0.6,ciscopix,Manual
10.17.0.4,ciscopix,Manual
10.17.0.4,ciscopix,Automatic
10.17.0.3,ciscopix,Manual"
```

- Syslog sample:

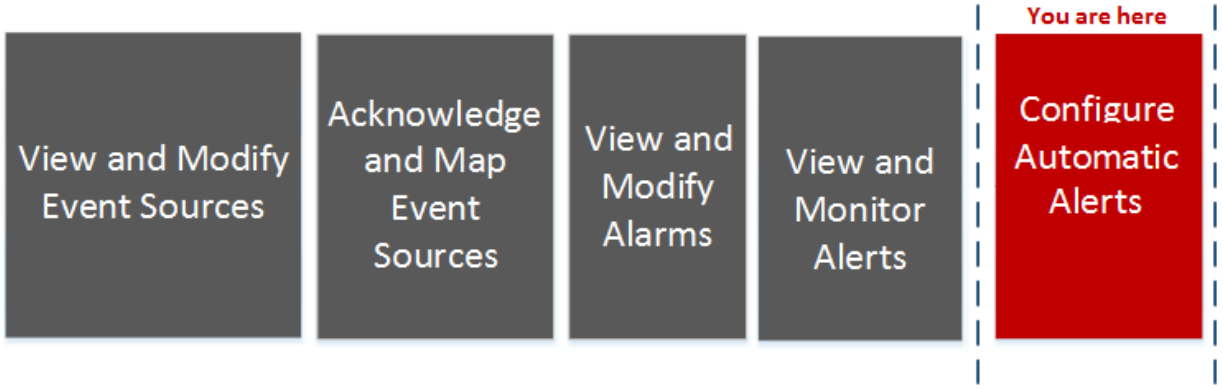
```
11-11-2015 11:57:33 User.Info 127.0.0.1 Nov 11 11:57:33 localhost
CEF:0|RSA|NetWitness Event Source Monitoring|10.6.0.0.0|
HighThresholdAlert|ThresholdExceeded|1|cat=PCI Event Source(s)|Devices|
src=10.17.0.10,ciscopix,Manual|src=10.17.0.13,ciscopix,Manual|src=10.1
7.0.8,ciscopix,Manual|src=10.17.0.8,ciscopix,Automatic|src=10.17.0.12,
ciscopix,Manual|src=10.17.0.5,ciscopix,Manual|src=10.17.0.6,ciscopix,M
anual|src=10.17.0.4,ciscopix,Manual|src=10.17.0.4,ciscopix,Automatic|s
rc=10.17.0.3,ciscopix,Manual|
```

Settings Tab

The Settings tab presents options for automatic monitoring (baseline alerting). To access this tab, go to  (Admin) > **Event Sources** > **Settings**.

Workflow

This workflow shows the overall process for configuring event sources.



What do you want to do?

Role	I want to...	Documentation
Administrator	View and modify event sources.	Managing Event Source Groups
Administrator	Acknowledge and map events sources.	Acknowledging and Mapping Event Sources
Administrator	Add and configure parser mappings for a Log Decoder	Manage Parser Mappings
Administrator	View event source alarms.	Viewing Event Source Alarms
Administrator	*Configure Automatic Alerts.	Automatic Alerting

Role	I want to...	Documentation
Administrator	Troubleshoot event source management.	ESM Troubleshooting & Appendix

*You can perform this task here.

Related Topics

Automatic Alerting

Disabling Notifications

Quick Look

You can set up policies and thresholds for your event source groups. You do this so that you can receive notifications when the thresholds are not met. NetWitness also provides an automatic way to receive alarms, if you do not want to set up thresholds to generate alarms.

About Automatic Alerting

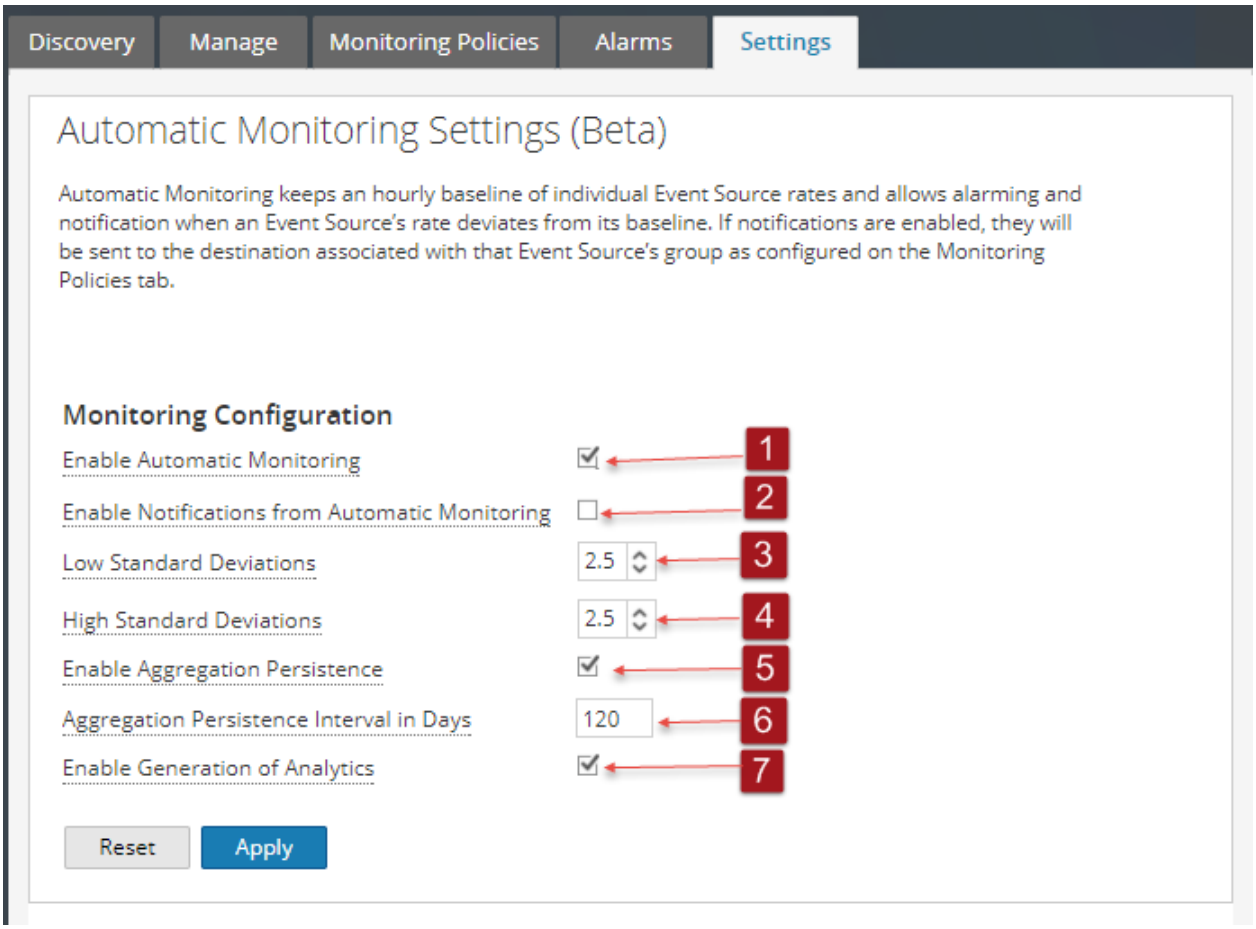
You can set up policies and thresholds for your event source groups. You do this so that you receive notifications when the thresholds are not met. NetWitness also provides an automatic way to receive alarms, if you do not want to set up thresholds to generate alarms.

To trigger automatic alerts, you can use baseline values. This way, you do not need to set up numerous group thresholds and policies in order to receive alerts. Any anomalous amount of messages trigger alerts, without needing to do any configuration (except for turning on automatic alerting).

Note the following:

- After you begin collecting messages from an event source, it takes the system approximately a week to store a baseline value for that event source. After this initial period, the system alerts you when the number of messages for a period are above or below the baseline by a set amount. By default, this amount is 2 standard deviations above or below the baseline.
- Base your high and low deviation settings on how "regular" your event sources behave. That is, if you expect little or no variance in the number of messages that arrive for a given time (for example, 8 to 9 am on a weekday), then you can set a low value for the Deviation. Conversely, if you often see peaks and valleys, set the Deviation value higher.
- If you enable a policy, but do not have any thresholds set, then you can still receive automatic (baseline) notifications, as long as you have turned on automatic alerting.

Note: Automatic alerting, and it settings, are currently in Beta testing.



1	Determines whether automatic alerting is on or off. By default, this option is selected (automatic alerting turned on)
2	Determines whether notifications for automatic alerts are on or off. By default, this option is cleared (automatic notifications are not sent when automatic alerts are triggered)
3	The standard deviations below which to receive alerts. Default is **2.5** (95% confidence)
4	The standard deviations above which to receive alerts. Default is **2.5** (95% confidence)
5	When selected, this option stores event source counts per one-hour interval. The data that is collected is used to form the baseline values for each event source. • **Enabled (default)**: one count per hour per event source is stored in the underlying database. These one-hour counts (or aggregations) form the historical basis for computing the normal range for each event source. • **Disabled**: when the SMS Server is restarted, Event Source Monitoring will have no historical data with which to compute the normal range and the user will have to wait until enough data (about a week's worth) is collected to form a new basis for each event source
6	Controls how much historical data (see **Enable Aggregation Persistence**) to maintain for each event source. The default value of 120 days means roughly 4 months of history is kept and used when reconstructing the basis for each event source
7	When enabled, data about the behavior of the automatic alerting is stored to disk. The default value is **Enabled**. The data retained includes baseline value over time and the alerting history for each event source.

Note, however, the event source address and type is anonymized, so only your event rate information is revealed.

Since automatic alerting is a beta feature, this data is important to measure the efficacy of the feature. This can be disabled without affecting the automatic alerting functionality

8 The **Reset** option discards any unsaved changes for all settings on the page.

9 Click **Apply** to save any changes you made to the values on the page.

Features

The Settings tab contains the following features.

Feature	Description
Enable Automatic Monitoring	Determines whether automatic alerting is on or off. By default, this option is selected (automatic alerting turned on)
Enable Notifications From Automatic Monitoring	Determines whether notifications for automatic alerts are on or off. By default, this option is cleared (automatic notifications are not sent when automatic alerts are triggered)
Low Standard Deviations	The standard deviations below which to receive alerts. Default is **2.5** (95% confidence)
High Standard Deviations	The standard deviations above which to receive alerts. Default is **2.5** (95% confidence)
Enable Aggregation Persistence	When selected, this option stores event source counts per one-hour interval. The data that is collected is used to form the baseline values for each event source. • Enabled (default): one count per hour per event source is stored in the underlying database. These one-hour counts (or aggregations) form the historical basis for computing the normal range for each event source. • Disabled: when the SMS Server is restarted, Event Source Monitoring will have no historical data with which to compute the normal range and the user will have to wait until enough data (about a week's worth) is collected to form a new basis for each event source
Aggregation Persistence Interval in Days	Controls how much historical data (see **Enable Aggregation Persistence**) to maintain for each event source. The default value of 120 days means roughly 4 months of history is kept and used when reconstructing the basis for each event source

Feature	Description
Enable Generation of Analytics	When enabled, data about the behavior of the automatic alerting is stored to disk. The default value is **Enabled**. The data retained includes baseline value over time and the alerting history for each event source. Note, however, the event source address and type is anonymized, so only your event rate information is revealed. Since automatic alerting is a beta feature, this data is important to measure the efficacy of the feature. This can be disabled without affecting the automatic alerting functionality
Reset	This option discards any unsaved changes for all settings on the page.
Apply	Click **Apply** to save any changes you made to the values on the page.

ESM Troubleshooting & Appendix

Troubleshooting Topics:

- Alarms and Notifications Issues
- Duplicate Log Messages
- Troubleshooting Feeds
- Import File Issues
- Negative Policy Numbering

Appendix: Viewing Logs from Pre-11.0 Log Decoder

Alarms and Notifications Issues

This topic describes how to address problems you may encounter with alarms or notifications.

Alarms

If you are not seeing alarms that you expect to see, make sure that you have configured all the necessary items, as discussed below.

Automatic Alarms

To see automatic alarms appear on the Alarms screen, the **Enable Automatic Monitoring** option must be selected.

This option is on the **Settings** tab ((Admin) > Event Sources > Settings), and is selected by default. However, at some point someone may have cleared this option.

Manual Alarms

To see manual alarms appear on the Alarms screen, all of the following conditions must be met:

- The event source must be part of a Group.
- The Group must have a policy with either a low or high (or both) threshold defined.
- The Group Policy must be enabled.

Notifications

If you are seeing alarms, but are not receiving the expected notifications, make sure that you have configured all the necessary items, as discussed below.

Also, make sure that you have correctly configured the Notification Servers and Notification Outputs.

Much of the preliminary configuration for Notifications is done from **(Admin) > System > Global Notifications**. For details, see the **Global Notifications Panel** topic in the *System Configuration Guide*.

Automatic Notifications

To have the system send automatic notifications, all of the following conditions must be met:

- The **Enable Automatic Monitoring** option must be selected (this option is selected by default).
- The **Enable Notifications From Automatic Monitoring** option must be selected. This option is cleared by default, so you or someone in your organization must select it. Navigate to  (Admin) **> Event Sources > Settings** to see this option.
- The event source that triggered the alarm must be in a group that has a policy enabled: note that no thresholds need to be set for automatic notifications.
- The policy must at least one notification configured (either email, SNMP or Syslog).

Manual Notifications

To have the system send manual notifications (that is, a notification which says that a manual alarm was triggered):

- The event source that triggered the alarm must be in a group that has a group policy enabled.
- There must be a threshold set for the policy.
- At least one notification has been configured for the policy.

Duplicate Log Messages

It is possible that you are collecting messages from the same event source on two or more Log Collectors. This topic describes the problem and ways to troubleshoot the issue.

Details

If the ESM aggregator detects the same events for the same event source on multiple Log Collectors, you receive a warning similar to the following:

```
2015-03-17 15:25:29,221 [pool-1-thread-6] WARN
com.rsa.smc.esm.groups.events.listeners.EsmStatEventListener -
 192.0.2.21-apache had a previous event only 0 seconds ago; likely because it
exists on multiple log collectors
```

This warning message means the `192.0.2.22-apache` event source is being collected by multiple hosts. You can see the list of hosts in the Log Collector column in the **Manage** tab in the Administration > Event Sources view.

Clean Up Duplicate Messages

1. Stop collectd on NetWitness and Log Decoders:

   ```
   service collectd stop
   ```

2. Remove the ESM Aggregator persisted file on NetWitness:

   ```
   rm /var/lib/netwitness/collectd/ESMAggregator
   ```

3. Reset the Log Decoder.
   a. Navigate to the Log Decoder REST, at `http://<LD_IP_Address>:50102`.
   b. Click **decoder(*)** to view the properties for the decoder.
   c. In the Properties drop-down menu, select **reset**, then click **Send**.
4. In the Event Sources panel from the Event Sources Manage tab, select all event sources and then click − to remove them.
5. Start the collectd service:
   ```
   service collectd start
   ```

Troubleshooting Feeds

The purpose of the feed generator is to generate a mapping of an event source to the list of groups to which it belongs.

If you have an event source from which you are collecting messages, and yet it is not displayed in the correct event source groups, then this topic provides background and information to help you track down the problem.

Details

The ESM Feed maps multiple keys to single value. It maps the DeviceAddress, Forwarder, and DeviceType attributes to groupName.

The purpose of the ESM feed is to enrich event source Meta with the groupName collected on the Log Decoder.

How it Works

The feed generator is scheduled to update every minute. However, it is triggered only if there are any changes (create, update, or delete) in event sources or groups.

It generates a single feed file with event source to group mapping, and pushes the same feed to all of the Log Decoders that are connected to NetWitness.

After the feed file is uploaded on the Log Decoders, for any new events, it enriches events Meta data with groupName, and appends this groupName to logstats.

After the groupName is in logstats, the ESM Aggregator groups information and sends it to ESM. At this point, you should see the **Group Name** column under the **Event Source Monitoring** tab.

The entire process can take some time. Therefore, you may need to wait for several seconds after you add a new group or event source, before the Group name is displayed.

Note: If the event source type attribute changes when the feed is updated, NetWitness adds a new logstats entry, rather than updating the existing one. Thus, there will be two different logstats entries in logdecoder. Previously existing messages would have been listed under the previous type, and all new messages are logged for the new event source type.

Feed File

The format of the feed file is as follows:

```
DeviceAddress, Forwarder, DeviceType, GroupName
```

The `DeviceAddress` is either `ipv4, ipv6`, or `hostname`, depending on which of these have been defined for the event source.

The following is a sample of the feed file:

```
"12.12.12.12","d6","NETFLOW","grp1"
"12.12.12.12","ld4","netflow","grp1"
"12.12.12.12","d6","netfow","grp1"
"0:E:507:E6:D4DB:E:59C:A","10.25.50.243","apache","Apachegrp"
"1.2.3.4","LCC","apache","Apachegrp"
"10.100.33.234","LC1","apache","Apachegrp"
"10.25.50.248","10.25.50.242","apache","Apachegrp"
"10.25.50.251","10.25.50.241","apache","Apachegrp"
"10.25.50.252","10.25.50.255","apache","Apachegrp"
"10.25.50.253","10.25.50.251","apache","Apachegrp"
"10.25.50.254","10.25.50.230","apache","Apachegrp"
"10.25.50.255","10.25.50.254","apache","Apachegrp"
"13.13.13.13","LC1","apache","Apachegrp"
"AB:F255:9:8:6C88:EEC:44CE:7",,"apache","Apachegrp"
"Appliance1234",,"apache","Apachegrp"
"CB:F255:9:8:6C88:EEC:44CE:7","10.25.50.253","apache","Apachegrp"
```

Troubleshooting Feeds

You can check the following items to narrow down where the problem is occurring.

10.5 Log Decoders

Are your NetWitness Log Decoders at version 10.5 or later? If not, you need to upgrade them. For NetWitness version 10.6, feeds are sent only to version 10.5 and later Log Decoders.

Feed File Existence

Verify that the feeds ZIP archive exists in the following location:

```
/opt/rsa/sms/esmfeed.zip
```

Do not modify this file.

Group Meta Populated on LD

Verify that the group meta is populated on the Log Decoder. Navigate to the Log Decoder REST and check logstats:

```
http://LogDecoderIP:50102/decoder?msg=logStats&force-content-type=text/plain
```

This is a sample logstats file with group information:

```
device=apache forwarder=NWAPPLIANCE10304 source=1.2.3.4 count=338
lastSeenTime=2015-Feb-04 22:30:19 lastUpdatedTime=2015-Feb-04
22:30:19
groups=IP1234Group,apacheGroup
device=apachetomcat forwarder=NWAPPLIANCE10304 source=5.6.7.8
count=1301 lastSeenTime=2015-Feb-04 22:30:19 lastUpdatedTime=2015-
Feb-04 22:30:19
groups=AllOtherGroup,ApacheTomcatGroup
```

In the above text, the group information is bolded.

Device Group Meta on Concentrator

Verify that the **Device Group** meta exists on the Concentrator, and that events have values for the `device.group` field.



SMS Log File

Check the SMS log file in the following location to view informational and error messages:
`/opt/rsa/sms/logs/sms.log`

The following are example *informational* messages:

```
Feed generator triggered...
Created CSV feed file.
Created zip feed file.
Pushed ESM Feed to LogDeocder : <logdecoder IP>
```

The following are example *error* messages:

```
Error creating CSV File : <reason>Unable to push the ESM Feed:
Unable to create feed zip archive.
Failed to add Group in CSV: GroupName: <groupName> : Error: <error>
Unable to push the ESM Feed: CSV file is empty, make sure you have
al-least on group with at least one eventsource.
Unable to push the ESM Feed: No LogDecoders found.
Unable to push the ESM Feed: Unable to push feed file on LogDecoder-
<logdecoderIP>Unable to push the ESM Feed:
admin@<logdecoderIP>:50002/decoder/parsers received error: The zip
archive "/etc/netwitness/ng/upload/<esmfeedfileName>.zip" could not
be opened
Unable to push the ESM Feed: <reason>
```

Verify Logstats data is getting Read & Published by ESMReader & ESMAggregator

These are the steps to verify that logstats are collected by **collectd** and published to Event Source Management.

ESMReader

1. On Log Decoders add **debug "true"** flag in in **/etc/collectd.d/logdecoder-esm.conf**:

```
#
# Copyright (c) 2014 RSA The Security Division of EMC
#
<Plugin generic_cpp>    PluginModulePath "/usr/lib64/collectd"
    debug "true"

    <Module "NgEsmReader" "all">
        port      "56002"
        ssl       "yes"
        keypath   "/etc/pki/nw/node/node-key.pem"
        certpath  "/etc/pki/nw/node/node-cert.pem"
        interval  "600"
        query     "all"
        <stats>
        </stats>
    </Module>
    <Module "NgEsmReader" "update">
        port      "56002"
        ssl       "yes"
        keypath   "/etc/pki/nw/node/node-key.pem"
        certpath  "/etc/pki/nw/node/node-cert.pem"
        interval  "60"
        query     "update"
        <stats>
        </stats>
    </Module>
</Plugin>
```

2. Run the command:

```
collectd service restart
```

3. Run the following command:

   ```
   tail –f /var/log/messages | grep collectd
   ```

   Verify that ESMReader is reading logstats and there are no errors. If there are any read issues, you will see errors similar to the following:

   ```
   Apr 29 18:47:45 NWAPPLIANCE15788 collectd[14569]: DEBUG: NgEsmReader_all:
   error getting ESM data for field "groups" from logstat device=checkpointfw1
   forwarder=PSRTEST source=1.11.51.212. Reason: <reason>Apr 29 18:58:36
   NWAPPLIANCE15788 collectd[14569]: DEBUG: NgEsmReader_update: error getting
   ESM data for field "forwarder" from logstat device=apachetomcat
   source=10.31.204.240. Reason: <reason>
   ```

ESMAggregator

1. On NetWitness, uncomment the verbose flag in **/etc/collectd.d/ESMAggregator.conf**:

   ```
   # ESMAggregator module collectd.conf configuration file
   #
   # Copyright (c) 2014 RSA The Security Divsion of EMC
   #
   <Plugin generic_cpp>
   PluginModulePath "/usr/lib64/collectd"
   <Module "ESMAggregator">
           verbose 1
           interval "60"
           cache_save_interval "600"
           persistence_dir "/var/lib/netwitness/collectd"
   </Module>
   </Plugin>
   ```

2. Run the following:

   ```
   collectd service restart.
   ```

3. Run the following command:

   ```
   run “tail –f /var/log/messages | grep ESMA
   ```

   Look for ESMAggregator data and make sure your logstat entry is available in logs.

Sample output:

```
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[0]
logdecoder[0] = d4c6dcd4-6737-4838-a2f7-ba7e9a165aae
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[1]
logdecoder_utcLastUpdate[0] = 1425174451
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[2]
groups = Cacheflowelff,Mixed
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[3]
```

```
logdecoders = d4c6dcd4-6737-4838-a2f7-ba7e9a165aae
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[4]
utcLastUpdate = 1425174451
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: Dispatching
ESM stat NWAPPLIANCE15788/esma_update-cacheflowelff/esm_counter-3.3.3.3 with a
value of 1752 for NWAPPLIANCE15788/cacheflowelff/esm_counter-3.3.3.3
aggregated from 1 log decoders
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[0]
logdecoder[0] = 767354a8-5e84-4317-bc6a-52e4f4d8bfff
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[1]
logdecoder_utcLastUpdate[0] = 1425174470
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[2]
groups = Cacheflowelff,Mixed
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[3]
logdecoders = 767354a8-5e84-4317-bc6a-52e4f4d8bfff
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[4]
utcLastUpdate = 1425174470
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: Dispatching
RRD stat NWAPPLIANCE15788/esma_rrd-cacheflowelff/esm_counter-3.3.3.3 with a
value of 1752 for NWAPPLIANCE15788/cacheflowelff/esm_counter-3.3.3.3
aggregated from 1 log
```

Configure JMX Feed Generator Job Interval

Although the feed generation job is scheduled to execute every minute by default, you can change this by using **jconsole**, if necessary.

To change the feed generator job interval:

1. Open **jconsole** for the SMS service.
2. On the MBeans tab, navigate to **com.rsa.netwitness.sms** > **API** > **esmConfiguration** > **Attributes**.
3. Modify the value for the property **FeedGeneratorJobIntervalInMinutes**.
4. Go to **Operations** under the same navigation tree, an click **commit()**. This persists the new value in the corresponding json file under **/opt/rsa/sms/conf**, and uses the value if SMS is restarted.

Setting a new value reschedules the feed generator job for the new interval.

Import File Issues

If your import file is not formatted correctly, or is missing required information, an error is displayed, and the file is not imported.

Check the following:

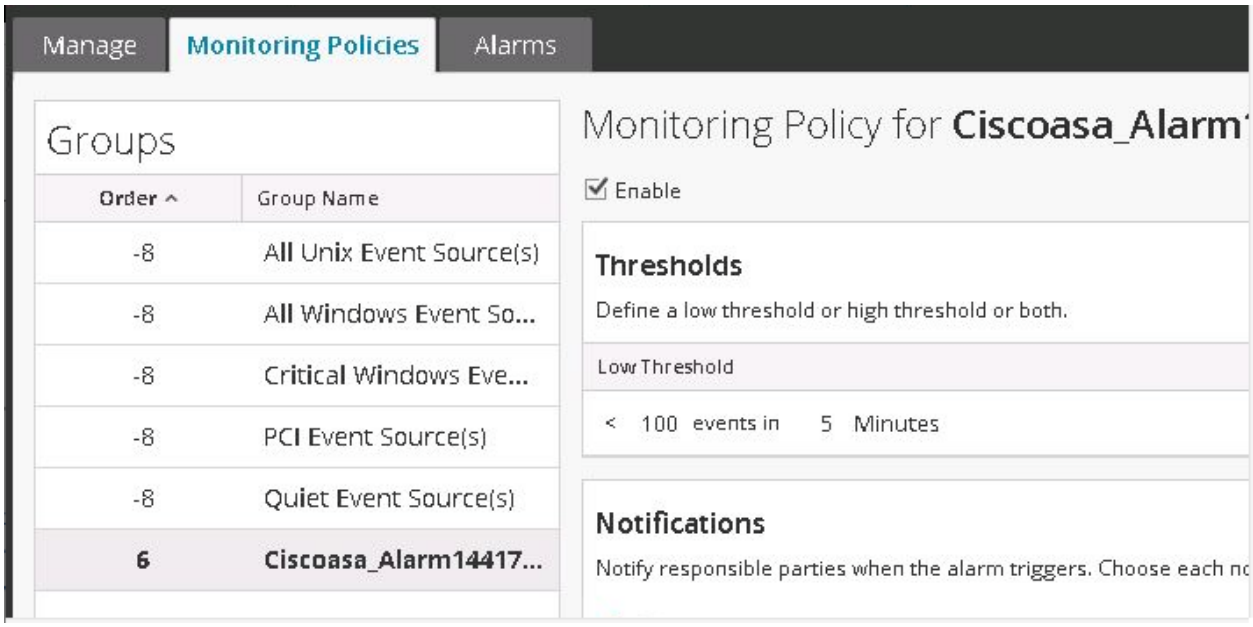
- If you are adding unknown sources, each line in the file must contain a combination of the required attributes:
  - IP or IPv6 or Hostname, and
  - Event Source Type
- The first line of the file must contain header names, and the names must match the names in NetWitness. To get a list of correct column names, you can export a single event source. Examine the exported CSV file: the first row of the file contains the correct set of attribute/column names.

Negative Policy Numbering

You may see negative numbers in the Order field in the Groups section of the Monitoring Polices tab. This topic describes a workaround to restore the correct numbering scheme for your policies.

Details

The following screen shows an example of the situation where the numbers of group policies become negative.



If you encounter this situation, drag and drop the top group (**All Unix Event Source(s)** in the above image) to after the last group (**Ciscoasa_Alarm14417**). This restores normal, ordinal numbering. You can then continue to drag and drop groups until you have them in their proper order for your organization.

Clean Up Duplicate Messages

1. Stop collectd on NetWitness and Log Decoders:

   ```
   Service collectd stop
   ```

2. Remove the ESM Aggregator persisted file on NetWitness:

   ```
   rm /var/lib/netwitness/collectd/ESMAggregator
   ```

3. Reset the Log Decoder.

   a. Navigate to the Log Decoder REST, at `http://<LD_IP_Address>:50102`

   b. Click **decoder(*)** to view the properties for the decoder.

   c. In the Properties drop-down menu, select **reset**, then click **Send**.

4. In the Event Sources panel from the Event Sources Manage tab, select all event sources and then click − to remove them.

Viewing Logs from Pre-11.0 Log Decoder

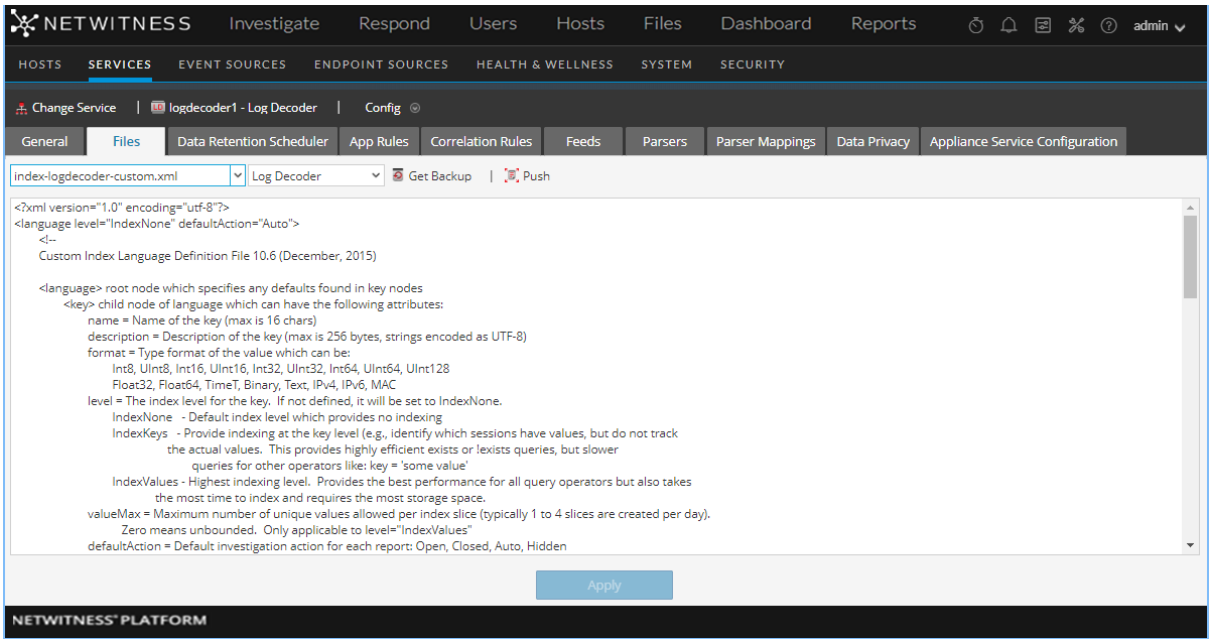
NetWitness 11.0 added the capability to view a small sampling of recent logs for specific devices through detail tabs of the Discovery View. By default, Log Decoders prior to 11.0 do not have the necessary configuration to enable this feature, but a few minor changes can make it available.

To enable logs preview for a pre-11.0.0.0 Log Decoder, follow these steps on the Log Decoder:

1. Go to  **(Admin) > Services >** select a Log Decoder, then select  **> View > Config**.
2. Click **Files** tab, then select **index-logdecoder-custom.xml** from the drop-down menu.
3. Add the following three lines at the end of the file (before the closing language tag):

   ```
   <key description="Device IP" level="IndexValues" name="device.ip" format="IPv4"
   valueMax="100000" defaultAction="Open"/>
   <key description="Device IPv6" level="IndexValues" name="device.ipv6" format="IPv6"
   valueMax="100000" defaultAction="Open"/>
   <key description="Device Host" level="IndexValues" name="device.host" format="Text"
   valueMax="100000" defaultAction="Open"/>
   ```

4. Click **Apply**.



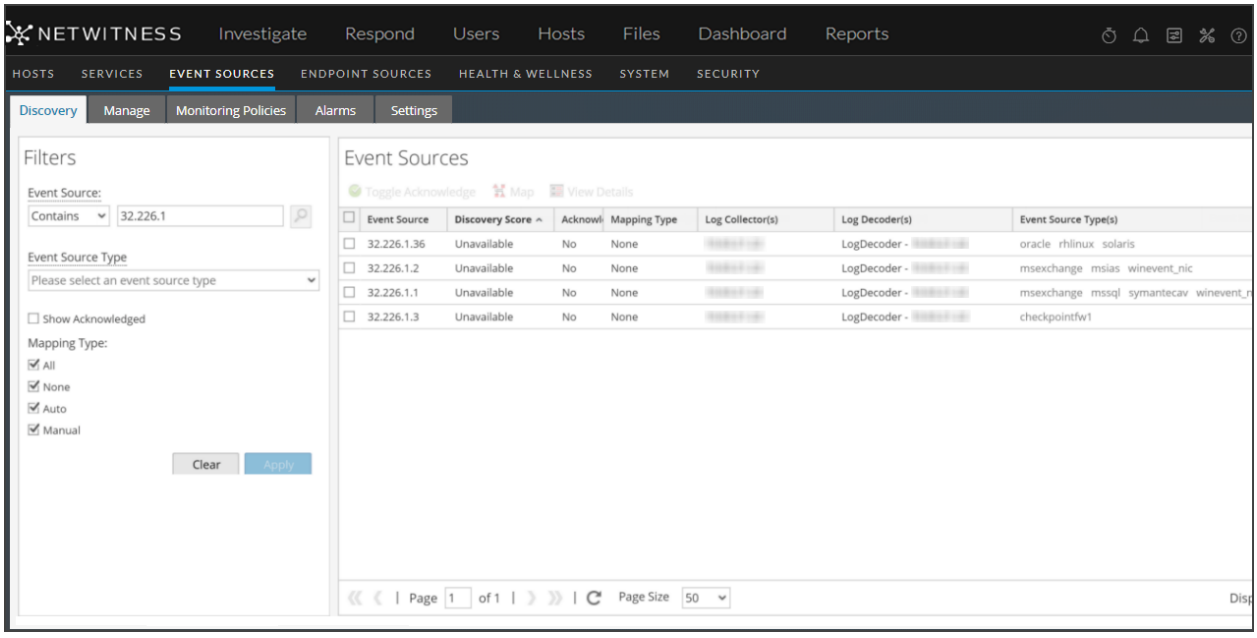
5. Restart the Log Decoder as follows.

   Select **Log Decoder > Explore > sys > Properties > shutdown**

This is an example of the **index-logdecoder-custom.xml** file.

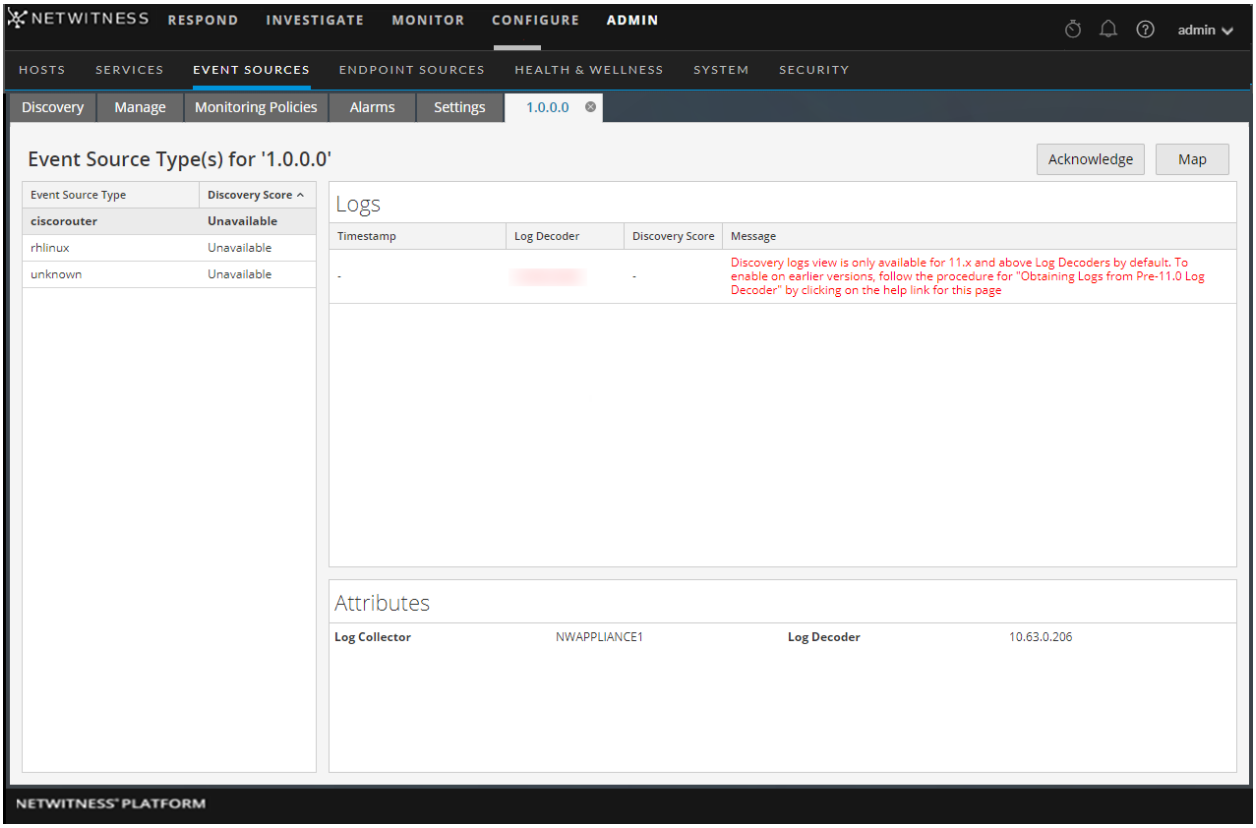
Note: Discovery Scores are only available for 11.0 and above Log Decoders. Discovery Scores for pre-11.0 Log Decoders display as Unavailable.

The following example displays the Discovery Score as **Unavailable** in the **Details** view for a pre-11.0 Log Decoder.



Note: Device logs are only available for 11.0 and above Log Decoders.

The following example shows the message that displays in the displays in the Logs panel for a pre-11.0 Log Decoder.



NetWitness® Platform XDR

Version 12.0

NetWitness Export Connector - Installation and Configuration Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

Overview **4**
Work Flow of NetWitness Export Connector 4
Configuration Process **6**
VM Sizing Recommendations **7**
Log Decoder 7
Network Decoder 8
Install Logstash **10**
Troubleshooting 10
Install NetWitness Export Connector **11**
Configure NetWitness Export Connector **12**
Position tracking and start session 16
Configure SSL 16
Certificate and Keystore 17
Health and Wellness 18
Download and Install Dashboard. 20
Create a User in New Health & Wellness (Kibana) 21
Enable the Logstash Plugin Metrics 21
Enable the Logstash Host Metrics 22
Configuring Custom Multi-valued Meta 24
Configure Logstash Filter Plugin (optional) **25**
Configure Logstash Output Plugin **26**
Performance tuning for Kafka and Kafka Output Plugin 26
Known Issues **27**

Overview

The NetWitness Export Connector is an input plugin for Logstash, used to export NetWitness Platform events and routes the data where you want, all in continuous, streaming fashion. Giving you the flexibility to unlock a variety of downstream use cases.

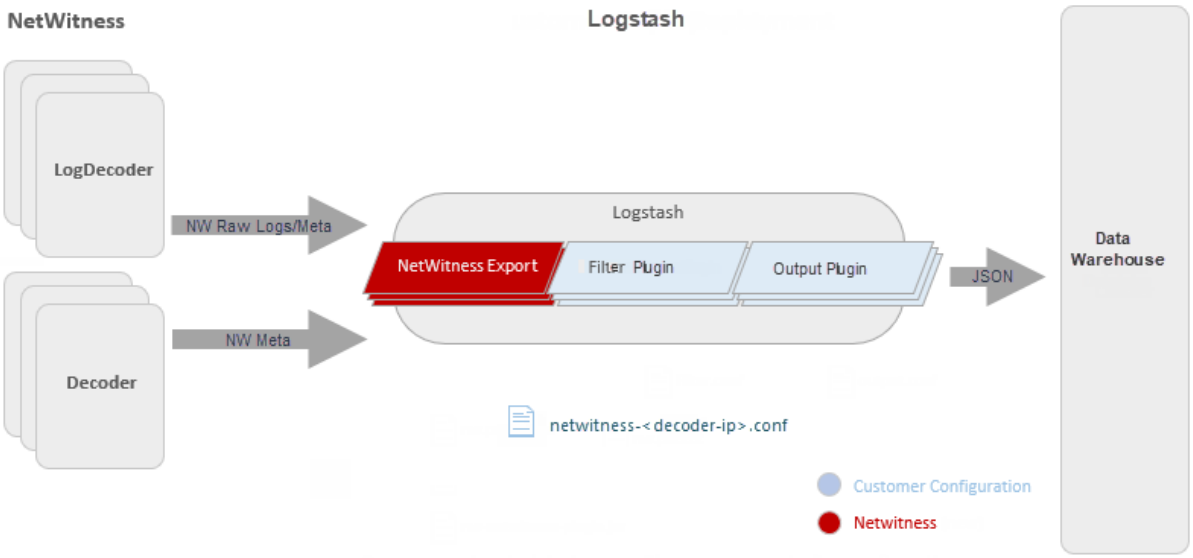
This plugin is installed on Logstash and integrates with NetWitness Decoders and Log Decoders. This plugin aggregates meta data and raw logs from the Decoder or Log Decoder and converts it to Logstash JSON object, which can easily integrate with numerous consumers such as Kafka, AWS S3, TCP, Elastic and others.

Install NetWitness Export Connector on the Logstash service. To activate the connector, restart the Logstash service.

Note: From 11.6 onwards, the Logstash server is packaged and supported along with the NetWitness Log Collector or Virtual Log Collector (VLC) service to provide easy access to Logstash. This is referred to as Managed Logstash and it eliminates the need for a separate Logstash server outside of the NetWitness Platform. For more information, see "Configure Logstash Event Sources in NetWitness" in the *Log Collection Configuration Guide*.

Work Flow of NetWitness Export Connector

Following diagram shows how NetWitness Export Connector works.



There are of three plugins available that helps with export.

- Input plugin
- Filter plugin (optional)
- Output plugin

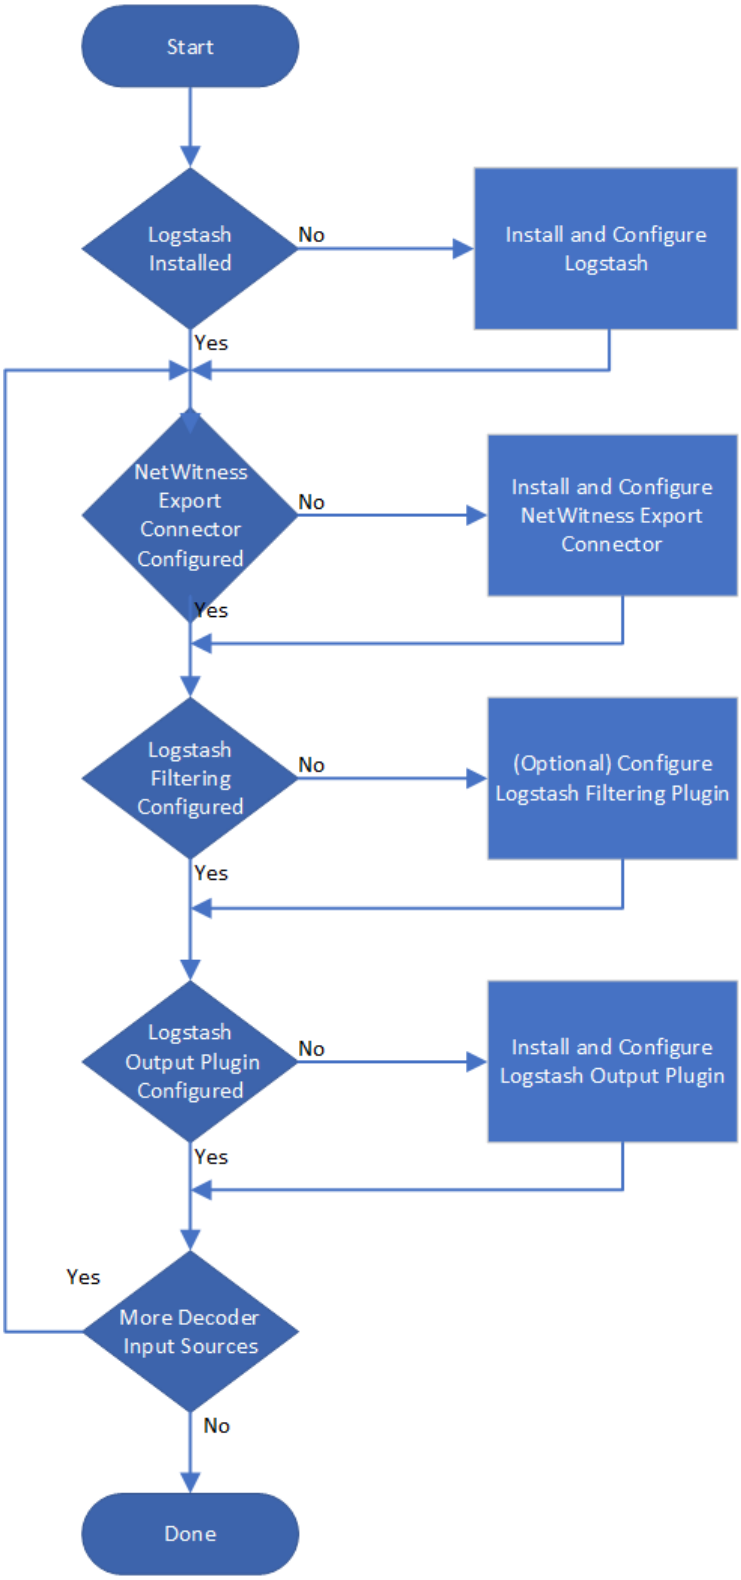
1. The Input plugin collects the events from the event sources. You must install the NetWitness Export Connector to collect events from Decoder or Log Decoder. The NetWitness Export Connector uses NetWitness API that collects the following data and forwards it as Logstash messages.
   - Meta data and raw log data from the Decoder
   - Meta data from Log Decoder

The data is then forwarded to the Filter plugin.

2. (Optional) The Filter plugin adds, removes, or modifies the received data and forwards it to the Output plugin. You can use the standard Logstash filter plugins to add, remove, or modify the data.
3. The Output plugin sends the processed event data to the data warehouse destinations. You can use the standard Logstash output plugins to send the data.

Configuration Process

The following flowchart describes the steps to configure NetWitness Export Connector.



VM Sizing Recommendations

It is recommended to install the Logstash and the NetWitness Export Connector in an independent VM (Virtual Machine) deployment. Following are the requirements for different variations of EPS (Events Per Second).

Total VM Memory: Set an additional 4GB memory for the JVM (Java Virtual Machine) memory mentioned in the following table for the respective EPS. For Example, for 5K EPS, the JVM memory required is 4GB, the total VM memory must be set to 8 GB.

JVM Memory: The Xms and Xmx for the Logstash service must be set to the same value as shown in the table. You can update the JVM settings for Logstash in the `/etc/logstash/jvm.options` file. For more information, see JVM Settings documentation.

For more info, https://www.elastic.co/guide/en/logstash/7.x/jvm-settings.html

vCPUs: The vCPUs values that are shown is required only for NetWitness Export Connector to be functional for respective EPS. It is recommended to add extra 4 CPUs for other OS (Operating System) services. The recommended vCPU specifications for all the EPS listed in the following tables are: Intel Xeon CPU @2.6 Ghz.

Note: The above recommendations is valid only if the VM instance is completely dedicated for Logstash deployment.
The resources consumption is dependent on cumulative EPS irrespective of number of sources. For Example, if you have three Decoders of 5000 EPS each, the cumulative EPS is 15000 as a result we can use the sizing option of 15000 EPS instead of 5000 EPS 3 times.

Log Decoder

- Metadata over SSL (~50 meta keys)

EPS	vCPUs	JVM Memory
5000	4	4 GB
10000	4	4 GB
15000	8	8 GB
20000	8	8 GB
25000	14	12 GB
30000	14	12 GB
60000	16	32 GB

- Metadata over SSL (~140 meta keys)

EPS	vCPUs	JVM Memory
5000	8	4 GB
10000	11	4 GB
15000	14	8 GB
20000	17	12 GB
25000	19	16 GB
30000	19	16 GB

- Meta data and raw logs over SSL (~50 meta keys)

EPS	vCPUs	JVM Memory
5000	4	4 GB
10000	8	4 GB
15000	10	8 GB
20000	12	12 GB
25000	15	24 GB
30000	18	24 GB
60000	23	48 GB

- Meta data and raw logs over SSL (~140 meta keys)

EPS	vCPUs	JVM Memory
5000	8	4 GB
10000	10	8 GB
15000	17	10 GB
20000	19	16 GB
25000	20	32 GB
30000	20	32 GB

Network Decoder

- Meta data over SSL

EPS	vCPUs	JVM Memory
500 Mbps	4	2 GB
1 Gbps	4	2 GB
1.5 Gbps	8	4 GB
2 Gbps	8	4 GB
3.4 Gbps	12	8 GB
6 Gbps (10G Decoder)	12	24 GB
8 Gbps	15	32 GB
9.4 Gbps	15	48 GB

Install Logstash

IMPORTANT: Ensure that you follow all the security-related best practices and guidelines outlined in the Logstash documentation to avoid any potential security risks.

The NetWitness Export Connector is an input plugin for Logstash. Hence, installation of Logstash is important for NetWitness Export Connector to work. For more information, see the Overview section.

You can install the open source version of Logstash (OSS) or the paid version (Elastic). The supported version is Logstash 7.6.2.

Information on released versions of Logstash is available at Logstash Reference.

Note: It is recommended to install the Logstash service in CentOS operating system for better results.

Do the following steps to Install the Logstash.

1. Install the service: Installing Logstash
2. After installation, set Logstash to run as a service: Running Logstash
3. Enable Logstash to start when the system boots up: https://www.unix.com/man-page/centos/1/systemctl/

If you are using CentOS, make a note of the following:

- Logstash logs are stored in `/var/log/logstash/logstash-plain.log`
- If you install Logstash using rpm install, make sure it installs as logstash user and folders get created with the same user and **not** the root user.

Troubleshooting

For Generic Troubleshooting Instructions for Logstash, see Logstash Troubleshooting.

Install NetWitness Export Connector

Note: From 11.6 onwards, the Logstash server is packaged and supported along with the NetWitness Log Collector or Virtual Log Collector (VLC) service to provide easy access to Logstash. This is referred to as Managed Logstash and it eliminates the need for a separate Logstash server outside of the NetWitness Platform. For more information, see "Configure Logstash Event Sources in NetWitness" in the *Log Collection Configuration Guide*.

Do the following steps to install NetWitness Export Connector .

1. Download the offline installer from RSA Link in the following location: NetWitness Export Connector Installer.
2. Copy the downloaded NetWitness ZIP archive to the system where Logstash runs.
3. Open a command prompt and run the following command to change directory to Logstash home.
   `cd /usr/share/logstash`
4. Check the status of the Logstash service by running the following command.
   `systemctl status logstash`
5. Stop the Logstash service by running the following command.
   `systemctl stop logstash`
6. Install the NetWitness Export Connector by running the following command.
   `bin/logstash-plugin install file:///<path-to-file>/netwitness-export-connector-x.x.x.zip`
7. Make sure that all the required configuration files (`netwitness-<decoder-ip>-input.conf`) are available in the following folder.
   `/etc/logstash/conf.d/`
8. Start the Logstash service by running the following command.
   `systemctl start logstash`

Configure NetWitness Export Connector

Note: Make sure you open the firewall of the Decoder or Log Decoder to establish connection with the Logstash. For more information, see "Network Architecture and Ports" in *Deployment Guide for RSA NetWitness*.

You must configure the Logstash configuration file to process the NetWitness events. Create a Logstash configuration file and add the NetWitness Export Connector plugin parameter settings for event processing. Save the file as `netwitness-<decoder-ip>-input.conf`. After adding the NetWitness Export Connector plugin parameter settings, place the configuration file in `/etc/logstash/conf.d/` location.

A Logstash configuration file can have three separate sections for each type of plugin that you want to add to the event processing pipeline. The first section is for Input plugin (NetWitness Export Connector), the second section is for Filter plugin (optional) and the third section is for Output plugin.

To configure the NetWitness Export Connector plugin, add the parameter settings in the first section the Logstash configuration file.

For multiple pipelines configuration, see Multiple Pipelines Configuration documentation.

The configuration of each NetWitness Export Connector plugin must consist of the plugin name followed by a block of parameter settings for that plugin. If the NetWitness Export Connector has multiple plugins with block of parameters, they are applied in the order of their appearance.

The following is an example of NetWitness Export Connector with one plugin instance with block of parameter settings which fetches data from a single decoder .

```
input {
 netwitness {
 host => "<host>" # Mandatory
 username => "<username>" # Mandatory
 password => "<password>" # Mandatory
 decoder_type => "logdecoder" # Mandatory
 }
}
```

The following is an example of NetWitness Export Connector with two plugin instances with block of parameter settings which fetches data from two different decoders. Each plugin in the configuration is applied in the order as shown.

```
input {
 netwitness {
 host => "<host>" # Mandatory
 username => "<username>" # Mandatory
 password => "<password>" # Mandatory
 decoder_type => "logdecoder" # Mandatory
 }
netwitness {
 host => "<host>" # Mandatory
 username => "<username>" # Mandatory
 password => "<password>" # Mandatory
 decoder_type => "logdecoder" # Mandatory
 }
}
```

Note: When configuring the Logstash, you may need to specify sensitive settings such as passwords. You can use the Logstash keystore to securely store secret values instead of file system permissions for using it in configuration settings. For more information, see Logstash keystore Documentation.

Following are the parameters accepted by NetWitness Export Connector.

Parameter	Description	Parameter Type	Default Value
host	IP address or hostname of the Decoder or Log Decoder (mandatory)	String	N/A
username	Username used to access the Decoder or Log Decoder (mandatory)	String	N/A
password	Password of the user (mandatory)	String	N/A
decoder_type	Accepts only 'decoder' or 'logdecoder' (mandatory)	String	N/A
ssl_enable	Enable SSL connection between the Decoders and the NetWitness Export Connector. For more information, see Configure SSL	Boolean	false
ssl_certificate_path	Path of the CA certificate that is used for SSL and Trusted Connections		/etc/pki/nw/trust/truststore.pem
ssl_certificate_password	Password of the certificate in use. Mandatory if SSL is enabled	String	N/A
ssl_client_certificate_path	Client's SSL certificate		/etc/pki/nw/node/node-cert.pem
ssl_version	Version of the SSL connection	String	TLSv1.2
plugin_metrics_enable	Enables metrics reporting to Elastic (New Health and Wellness), for more information see Health and Wellness	Boolean	false
elastic_host	IP address or hostname of the Elastic host. Mandatory if plugin metrics is enabled	String	N/A
elastic_port	Port number of the Elastic host	String	9200

Parameter	Description	Parameter Type	Default Value
elastic_username	The username that is used to access the Elastic host. Mandatory if plugin metrics is enabled	String	N/A
elastic_password	The password that is used to access the Elastic host. Mandatory if plugin metrics is enabled	String	N/A
meta_include	Aggregates only the meta keys that are added in this parameter setting. Accepts comma separated values (csv) format	String	nil
meta_exclude	Excludes the meta keys that are added in this parameter setting from aggregation. Accepts comma separated values (csv) format	String	nil
start_session	Session from which the aggregation starts. Setting the value to 0 starts the aggregation from `last.session.id` in the Decoder	Number	0
export_log	Includes the raw log with the meta in the session (applicable only for Log Decoder aggregation)	Boolean	false
aggregate_ sessions	Number of sessions ingested in a batch from the Decoders	Number	1000
aggregation_ interval	Time interval (in milliseconds) between two event cycles	Number	1000
prefetch_count	Controls the number of batches to be pulled into the buffer that is available for the plugin to collect	Number	2

Parameter	Description	Parameter Type	Default Value
compression	The number of bytes in each message before message is compressed. Setting the value to '0' does not allow compression. Ranges between 0 to 131071	Number	0
compression_ level	The level of compression. Ranges between 0 to 9 where 1 is fastest and 9 has the better compression. A value of 0 selects the best balance between speed and compression	Number	6
buffer_size	Controls the number of records that the stream in the buffer before it is ingested	Number	40000
position_ tracking_path	Path where the last consumed session id is stored. Default storage location is `/var/lib/logstash`. For more information, see Position tracking and start session	String	`/var/lib/logstash`
custom_meta_ config_path	Path to the custom multivalued meta configuration file. For more information, see Configuring Custom Multi-valued Meta	String	nil
Query	Takes any NetWitness Platform query as Input **Note:** Only Indexed meta key must be the part of the query. For example, select `* where device.type = 'rhlinux'`	String	Select *

Position tracking and start session

Position tracking or bookmarking is used to track the sessions that are aggregated by logstash and sent to the consumer. Position tracking initiates automatically and updates the tracking file every 60 seconds in the path mentioned in `position_tracking_path` parameter in the configuration file. The file consists of two parts `[sessionid,timestamp]`. Default location is `/var/lib/logstash`, if it is not mentioned in the configuration file.

The `start_session` parameter accepts a number (long: primitive datatype) and indicates which is the first `sessionid` the plugin should request from its corresponding source. if the `start_session` parameter is not mentioned in the configuration file or if the value is mentioned as 0, the first session requested by the plugin will be the `last.session.id + 1`, `last_session_id` as in the decoder's REST API `/database/stats/last.seesion.id`.

if the position tracking file exists for a source and `start_session` is not configured or the `start_session` is set as value 0, the plugin with initiate aggregation from the `sessionid` indicated in the position tracking file.

if the position tracking file exists for a source and the `start_session` is set as non-zero value, the `start_session` value will take precedence over the position tracking file. The plugin will request from the session mentioned in the `start_session` parameter onwards.

Configure SSL

Note: When configuring the Logstash, you may need to specify sensitive settings such as passwords. You can use the Logstash keystore to securely store secret values instead of file system permissions for using it in configuration settings. For more information, see Logstash keystore Documentation.

To support trusted connections, the Decoder or Log Decoder has two ports, an unencrypted non-SSL port and an encrypted SSL port. Trusted connections require the encrypted SSL port.

To establish trusted connection for the Decoder or Log Decoder with Logstash, add the following parameters in the Logstash Configuration file (`netwitness-<decoder-ip>-input.conf`) to enable the SSL mode.

Parameters	Settings	Parameter Type	Default Value
ssl_enable	Set the value to 'true' to enable SSL	Boolean	false
ssl_certificate_path	Enter the path of the Logstash keystore file. The keystore file must be in `.p12` format	String	N/A
ssl_certificate_password	Enter the password of the keystore	String	N/A

Certificate and Keystore

The NetWitness Export Connector must have a valid CA (Certificate Authority) certificate and a server certificate to establish a trusted connection with Decoders or Log Decoders.

Create a CA certificate

To create a CA certificate, do the following steps.

1. SSH to NW Logstash host.
2. Change the directory to /root by running the following command.
   `cd /root`
3. Create a private key of 2048 bits by running the following command.
   `openssl genrsa -out CA-key.pem 2048`
4. Create a CA certificate by running the following command.
   `openssl req -new -key CA-key.pem -x509 -days 1000 -out CA-cert.pem`

Create Certificate Signing Request (CSR) and Keystore

Do the following steps provided in this procedure to create a CSR for server and Keystore for the NetWitness Export Connector. You must submit the CSR to the Certificate Authority (CA) server to obtain a server certificate. Once the server certificate is created, do the following steps to package the private key and the signed certificate that must be uploaded to the Logstash keystore.

To create a CSR, do the following steps.

Note: You can skip till step 6 if you have PEM already available.

1. SSH to NW Logstash host.
2. Change the directory to /root by running the following command.
   `cd /root`
3. Create a private key of 2048 bits by running the following command.
   `openssl genrsa -out server-key.pem 2048`
4. Create a CSR by running the following command.
   `openssl req -new -key server-key.pem -out signingReq.csr`
5. Submit the CSR to the CA and get a signed server certificate by running the following command.
   `openssl x509 -req -days 1000 -in signingReq.csr -CA CA-cert.pem -CAkey CA-key.pem -CAcreateserial -out server-cert.pem`
6. Create a Keystore for the NetWitness Export Connector by running the following command.
   `- openssl pkcs12 -export -in server-cert.pem -inkey server-key.pem -certfile CA-cert.pem -out logstash-input-netwitness.p12`
7. Copy the `/etc/pki/nw/trust/truststore.pem` from the Decoder or Log Decoder to the Logstash host by running the following command.
   `scp /etc/pki/nw/trust/truststore.pem <logstash ssh username>@<logstash destination directory path>`
8. Import `truststore.pem` from the Decoder or Log Decoder to the Logstash keystore (`logstash-input-netwitness.p12`) by running the following command.

```
keytool -importcert -keystore logstash-input-netwitness.p12 -trustcacerts -
alias nw-inter -file truststore.pem -storetype PKCS12
```

Note: `truststore.pem` is same for all Decoders, if the Logstash is aggregating from same NetWitness setup. If the Logstash is aggregating from two different NetWitness setups, you must copy the `truststore.pem` for each Decoders.

9. Copy `logstash-input-netwitness.p12` file to the `/etc/logstash` directory by running the following command.
   ```
   cp logstash-input-netwitness.p12 /etc/logstash
   ```
10. Change the owner of `logstash-input-netwitness.p12` as Logstash by running the following command.
    ```
    chown logstash:logstash /etc/logstash/logstash-input-netwitness.p12
    ```

Note: You must use the same keystore "`logstash-input-netwitness.p12`" for all NetWitness hosts that is integrated with Logstash.

11. Connect to the Decoders to add the `server-cert.pem` to `trustpeer` and `caupload` APIs by running the following commands. You can also use to REST API port to connect.
    ```
    curl -X POST -d server-cert.pem http://<logdecoder or decoder
    IP>:<logdecoder or decoder REST port>/sys/trustpeer
    curl -X POST -d server-cert.pem http:// <logdecoder or decoder
    IP>:<logdecoder or decoder REST port>/sys/caupload
    ```

Note: If you are not able to use the REST API, follow the below steps to copy the certificates to `/sys peerCert` and `/sys caCert` APIs using NetWitness CLI .
1. SSH to Admin Server.
2. Type "help" for a list of commands or "man" for a list of manual pages.
3. Run the following command.
```
login localhost:50002 <username> <password>
```
3. Run the following commands.
```
send /sys peerCert op=add --file-data=/root/server-cert.pem
send /sys caCert op=add --file-data=/root/server-cert.pem
```

12. Restart the Logstash service running the following command.
    ```
    systemctl restart logstash
    ```

Note: If you are using an untrusted certificate, copy the `truststore.pem` file from the Decoder or Log Decoder and import it to `<JAVA_HOME>/lib/security/cacerts` for the Logstash service to trust the CA certificate of Decoder or Log Decoder. For example: `keytool -importcert -file </path/to/file/>truststore.pem -keystore /usr/java/jdk-X.Y.Z/lib/security/cacerts -alias nw-core-cert -storepass <password>`

Health and Wellness

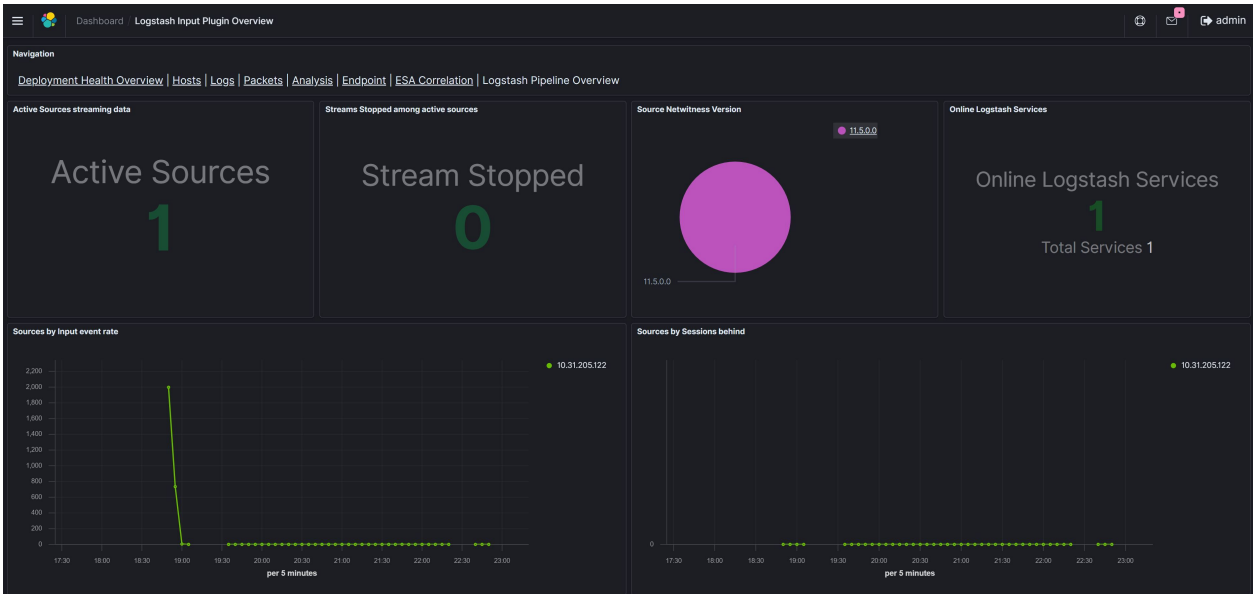
Note: This feature is available only from NetWitness version 11.5 and later. For more information, see "Monitor New Health and Wellness" in *System Maintenance Guide for RSA NetWitness*.

You can monitor the operational state of the Logstash service and details of sources configured in the New Health & Wellness tab available in NetWitness user interface. The New Health & Wellness is built on Elastic and Kibana. The metrics from the Logstash instance is sent to the Elastic service and visualized using Kibana. The tab displays the Logstash service status and the NetWitness Export Connector sources in a dashboard view. The New Health & Wellness service is referred as Elastic host in this document.

The Logstash service sends metrics for every 30 seconds. These metrics are used to view the operational status of the Logstash service and the NetWitness Export Connector. There are two types of metrics that are available for monitoring the Logstash status.

- Plugin metrics - For monitoring the NetWitness Export Connector status.
- Host metrics - For monitoring Logstash service instance (appliance or Virtual Machine) health status such as IO stats, CPU, memory usage and others.

By default, the plugin metrics are available in Logstash. To collect the host metrics, you must install Metricbeat, a third party software used for collecting the host metrics. It is recommended to enable both the plugin metrics and host metrics to view all the dashlets in the dashboard.





Download and Install Dashboard.

You need to download and deploy the dashboards from live to view the Logstash service and the NetWitness Export Connector stats in Kibana. Do the following steps.

1. Log in to NetWitness Platform UI.
2. Click **(Configure) > LIVE CONTENT**.
3. In the **Search Criteria** panel, select the **Resource Types** as:
   - Health and Wellness Dashboards
   - Health and Wellness Monitors
4. Click **Search**.
5. In the **Matching Resources** view, select the **Logstash Input Plugin Overview** dashboard to deploy.
6. In the **Matching Resources** toolbar, click Deploy.
7. In the **Deployment Wizard > Resources** tab, click **Next**.
8. In the **Services** tab, select the Metrics Server service.
9. Click **Next**.
10. Click **Deploy**.
    The **Deploy** page is displayed. The Progress bar turns green when you have successfully deployed the resources to the selected services.
11. Click **Close**.

For more information, see Advanced Configuration section in "Monitor New Heath and Wellness" topic in the *System Maintenance Guide for RSA NetWitness*.

Create a User in New Health & Wellness (Kibana)

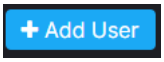
You must create user account in New Health & Wellness in NetWitness (Kibana), to send Logstash metrics to New Health & Wellness in NetWitness (Kibana). This user account is used in configuring the Logtstash plugin metrics and host metrics.

To create a user account in Kibana, do the following steps.

Note: This is applicable only for hosts or services that are not managed by NetWitness.

1. Go to  **(Admin) > HEALTH & WELLNESS > New Health & Wellness**.
   New Heath & Wellness panel view is displayed.
2. Click **Pivot to Dashboard**.
   Kibana Dashboard view is displayed in a new tab.
3. Click the Security icon  on the left.
   Kibana Security view is displayed.
4. In the Authentication Backends section, click **Internal User Database** and click the  icon to create new user.
5. Enter the username and the password in the respective columns.
6. In the Backend Roles section, click 

and add `nwservice-role`, and then click **Submit**.
   Username is displayed in the Internal Database section.
7. Click the 

icon to go back to Kibana Security view.
8. In the Permissions and Roles section, click **Role Mappings** and click **nwservice-role** to add the user name.
9. Click 

and add newly created username, and then click **Submit**.

Enable the Logstash Plugin Metrics

To enable the Logstash plugin metrics, add the following parameters in the Logstash Configuration file (`netwitness-<decoder-ip>-input.conf`).

Parameter	Setting	Parameter Type	Default Value
plugin_metrics_ enable	Set the value to 'true' to enable the plugin metrics	Boolean	false

Parameter	Setting	Parameter Type	Default Value
elastic_host	Enter the IP address or hostname of the Elastic host or IP address of the New Health & Wellness service	String	N/A
elastic_port	Enter the port number of the Elastic host	String	9200
elastic_username	Enter the username that is used to access the Elastic host (user account created in Kibana, see Create a User in New Health & Wellness (Kibana))	String	N/A
elastic_password	Enter the password that is used to access the Elastic host	String	N/A

Enable the Logstash Host Metrics

Note: This is applicable only for hosts or services that are not orchestrated by NetWitness.

To enable the Logstash host metrics, follow the below steps.

1. Download Metricbeat software. For more information refer to Metricbeat.
   ```
   curl -L -O https://artifacts.elastic.co/downloads/beats/metricbeat/
   metricbeat-oss-7.8.0-x86_64.rpm
   ```
2. Install Metricbeat on the Logstash service by running the following command.
   ```
   sudo rpm -ivh metricbeat-oss-7.8.0-x86_64.rpm
   ```
3. Disable the default collection of Logstash monitoring metrics in the configuration file `logstash.yml` by changing `xpack.monitoring.enabled` parameter to `false`.
   ```
   xpack.monitoring.enabled: false
   ```

Note: Generally, the `x.pack.monitoring.enabled` parameter is commented. Remove the (#) character to activate it.

4. Enable the `logstash-xpack` module in Metricbeat to collect host status by running the following command.
   ```
   metricbeat modules enable logstash-xpack
   ```
5. In `/etc/metricbeat/modules.d/system.yml` file, add the following lines at the last and save it.

```
# Processors to parse process names of Netwitness Logstash Service.
- module: system
  period: 1m
  metricsets:
    - process
  processors:
    - add_fields:
        fields: {process.given.name: "logstash-service"} # Adding process
name to a new field
        when.contains:
          system.process.cmdline: "logstash-core.jar" # Condition on which
process name will be updated
```

Note: Use proper indentation while adding the lines. Refer an existing module in the `system.yml` file. If indentation is incorrect, the service may fail to start.

6. Configure the monitoring data to send metrics to Metricbeat. Follow the below steps.
   a. Update the `/etc/metricbeat/metricbeat.yml`file.
      - Replace the content in 'Elasticsearch output' section in `metricbeat.yml` with the following lines.

```
output.elasticsearch:
  # Array of hosts to connect to.
  hosts: ["<ELASTIC_HOST>:9200"]
  index: "nw-logstash-metricbeats-%{+yyyy.MM.dd}"

  # Protocol - either `http` (default) or `https`.
  protocol: "https"
  ilm:
    enabled: false
    check_exists: false
  # Authentication credentials - either API key or username/password.
  username: "<ElASTIC_USERNAME>"
  password: "<ElASTIC_PASSWORD>"
  ssl.verification_mode: none

setup.ilm.enabled: false
setup.template:
   enabled: false # Disabled so that metricbeat auto-creates the
template everyday based on incoming metrics.
   name: "nw-logstash-metricbeats-%{+yyyy.MM.dd}"
   pattern: "nw-logstash-metricbeats-%{+yyyy.MM.dd}"
```

   b. Update the `ELASTIC_HOST`, `ELASTIC_USERNAME` and `ELASTIC_PASSWORD` details in `metricbeat.yml` file.

7. Start the Metricbeat service by running the following command.

```
sudo service metricbeat start
```

For more information visit the official documentation of Metricbeat.

Configuring Custom Multi-valued Meta

Default multi-valued meta keys are `action, alias.host, alias.ip, alias.ipv6, email` and `username`, if the custom multi-valued path parameter is not set in Logstash configuration file (`netwitness-<decoder-ip>-input.conf`), then only default values are considered for multi-valued meta. Custom valued meta can be setup in generic (applicable for all the Netwitness Platform host where the json file is referenced) or specific to the Netwitness Platform host and can have custom multi-valued meta for multiple host in single json file. Both the `custom_multi_valued_meta_generic field` and `custom_multi_valued_meta_specific field` can be used in single file or at-least one 1 field must be set.

Note: Provide absolute path to the json file in the configuration (`netwitness-<decoder-ip>-input.conf`) and make sure Logstash user has access to this file. Meta name should be the same as the Netwitness Platform default name format. Do not use underscore format of the meta key.

For example, below shown is the json file which has both generic and specific field set in `custom_meta.json`

```
{
  "custom_multi_valued_meta_generic": [
    "meta.1,meta.2,meta.3"
  ],
  "custom_multi_valued_meta_specific": {
    "0.0.0.0": [
      "meta.4,meta.5,meta.6"
  ],
  "1.1.1.1": [
    "meta.4,meta.5,meta.7"
  ]
 }
}
```

In above example, the meta keys `meta.1, meta.2 and meta.3` are applied to all the NetWitness Platform host in the configuration file that has custom_meta_config path is set. The meta keys `meta.1, meta.2, meta.3, meta.4, meta.5, meta.6` are set as custom multi-valued meta for hosts `0.0.0.0` and the meta keys `meta.1, meta.2, meta.3, meta.4, meta.5, meta.7` are set as custom multi-valued meta keys for the host `1.1.1.1` along with default mutli-valued meta.

Configure Logstash Filter Plugin (optional)

You can configure the Logstash Filter plugin to add, remove, or modify the specific input events from the Log Decoder or Decoder. To configure the Filter plugin, add the Filter plugin parameter settings in the second section of the Logstash configuration file (`netwitness-<decoder-ip>-input.conf`). This plugin modifies the events based on the parameter settings. You can use the existing standard Logstash filter plugins for adding the parameter settings to the configuration file. For more information on existing Logstash standard filter plugins, see Filter Plugin Documentation.

The configuration of the plugin must consist of the plugin name followed by a block of parameter settings for that plugin. The following is an example of Logstash mutate filter plugin configuration parameters that remove specific meta keys.

```
filter
 {
 mutate {
 remove_field => ["ip_dst", "ip_addr"]
 }
}
```

Configure Logstash Output Plugin

You must configure the Logstash Output plugin to send the input events to a data warehouse destination. To configure the Output plugin, in the third section of the Logstash configuration file (`netwitness-<decoder-ip>-input.conf`), add the Output plugin parameter settings. You can use the existing Logstash standard output plugins for adding the parameter settings to the configuration file. For more information on existing Logstash standard output plugins, see Output Plugin Documentation.

The configuration of the plugin must consist of the plugin name followed by a block of parameter settings for that plugin. The following is an example of Logstash Output plugin configuration parameters.

```
output
{
 kafka {
 codec => json
 topic_id => "logstash"
 }
}
```

Performance tuning for Kafka and Kafka Output Plugin

Following are the output plugin parameters for Kafka and Kafka Broker.

Kafka Output Plugin:
Following parameters require change in value for each parameters to add additional throughput in the deployment.
`batch_size, buffer_memory, compression_type, send_buffer_bytes`
For more information, see Kafka output plugin documentation.

Kafka Broker:
Following parameters require change in value for each parameters to add additional throughput in the deployment.
`num.network.threads, num.io.threads, socket.receive.buffer.bytes, socket.request.max.bytes, message.max.bytes, replica.fetch.max.bytes`
For more information, see Kafka Broker Performance documentation.

Known Issues

Com-ponents	Title, Problem & Workaround	Found In / Exists In	Fixed Version	Track-ing Number
NetWitness Export Connector	**Title:** Position tracking is not updated properly when query is used to filter the sessions. **Problem:** Last `session.id` aggregated is not updated to the latest aggregated ID but is updated only to the last filtered session retrieved. For Example, If a decoder has 1000 sessions, and only first 800 sessions match the filter, the `last.session.id` in the position tracking shows 800 even though all 1000 sessions are processed. This does not cause any data loss but if the Logstash service is restarted, aggregation will begin from 801 `session.id`, processing the 200 sessions again. **Workaround:** None	netwitness-export-connector-1.0.0		ASOC-101795
NetWitness Export Connector	**Title:** Consumer does not consume all the events after the Logstash service is restarted. **Problem:** While running at a high EPS or if there is a large session behind, the NetWitness Export Connector aggregates the sessions of the decoders. The Logstash pipeline sends to the output plugin which in turn sends it out to the consumer. At this point, if a Logstash service is stopped or restarted, the Logstash pipeline ensures the data is sent out of the output plugin to the consumer, however all the sessions are not consumed by the consumer. This leads to data loss on the consumer end even though Logstash has sent out all the data. **Workaround:** Update the `start_session` parameter in the Logstash configuration file (`netwitness-<decoder-ip>-input.conf`) by referring to the latest received `session.id` on Kafka or any other consumer used.	netwitness-export-connector-1.0.0	netwitness-export-connector-1.1.0	ASOC-102440

Com-ponents	Title, Problem & Workaround	Found In / Exists In	Fixed Version	Track-ing Number
NetWitness Export Connector	**Title:** Position tracking not saved for some Decoder sources after sometime in a multiple sources configuration. **Problem:** When aggregating data from multiple sources in a single Logstash instance, after a period of time, the position tracking thread fails for some or all of the sources gradually and stops tracking the position (bookmarking) for those sources. **Workaround:** None	netwitness-export-connector-1.0.0	netwitness-export-connector-1.1.0	ASOC-102411
NetWitness Export Connector	**Title:** Unable to start session aggregation from `first.session.id` of Decoder when it is not 1. **Problem:** If the Decoder has the `first.session.id` other than 1, which means the Decoder has data that has rolled over, the NetWitness Export Connector's `start_session` parameter will not start aggregation from that value, instead it would begin from the `last.session.id`. **Workaround:** Add 1 to the value to the `start_session` parameter and the aggregation would start as expected dropping the first session only. For example, if `first.session.id` on a Decoder is '1105000' add 1 to it and set the value of the `start_session` parameter as '1105001'.	netwitness-export-connector-1.0.0	netwitness-export-connector-1.1.0	ASOC-101951

NetWitness® Platform XDR

Version 12.0

Google Cloud Platform Installation Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

Google Cloud Platform Installation Overview **4**

GCP Deployment Scenarios **5**

NetWitness Full Stack VPC Visibility 5

Hybrid Deployment 5

GCP Deployment **6**

Checklist 6

Prerequisites 6

Find NetWitness Platform GCP Images 6

Establish gcloud Environment 7

Create an Instance using Google Cloud SDK Shell 7

Installation Tasks 8

Tasks - Install 12.0.0.0 on the NetWitness Server (NW Server) Host and Component Hosts 9

Storage Configurations 14

Create a Disk 14

Configure Hosts (Instances) in NetWitness Platform 18

GCP Instance Configuration Recommendations **19**

Google Cloud Platform Installation Overview

Note: Google Cloud Platform is supported from version 11.4 and later.

Google Cloud Platform (GCP) instances have the same functionality as the NetWitness hardware and virtual hosts. NetWitness recommends that you perform the following tasks when you set up your GCP environment.

Before you can deploy NetWitness in GCP, you need to:

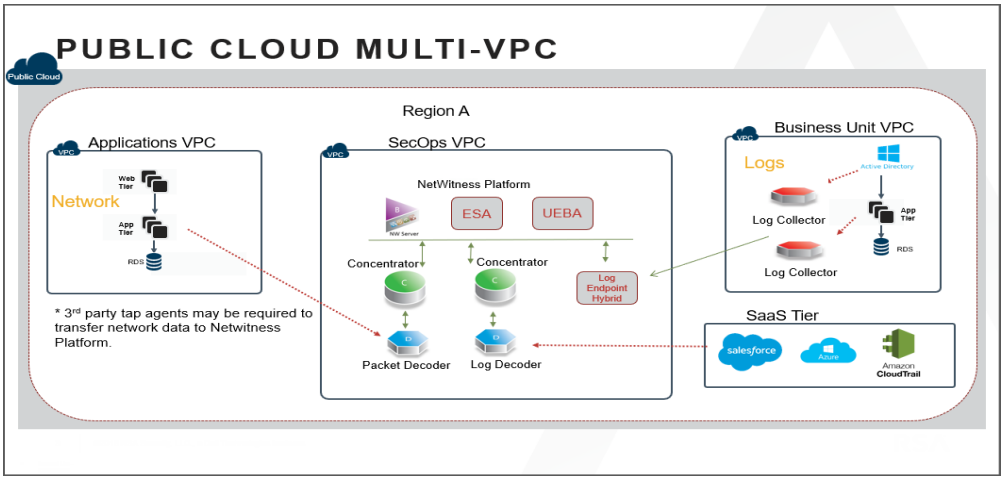
- Review the recommended compute and memory specifications needed for each NetWitness instance.
- Get familiar with the NetWitness Storage Guide to understand the types of drives and volumes needed to support NetWitness instances. For more information, see Storage Guide for NetWitness® Platform 12.0.
- Make sure that you have a NetWitness Throughput license.

GCP Deployment Scenarios

The following diagrams illustrate some common GCP deployment scenarios.

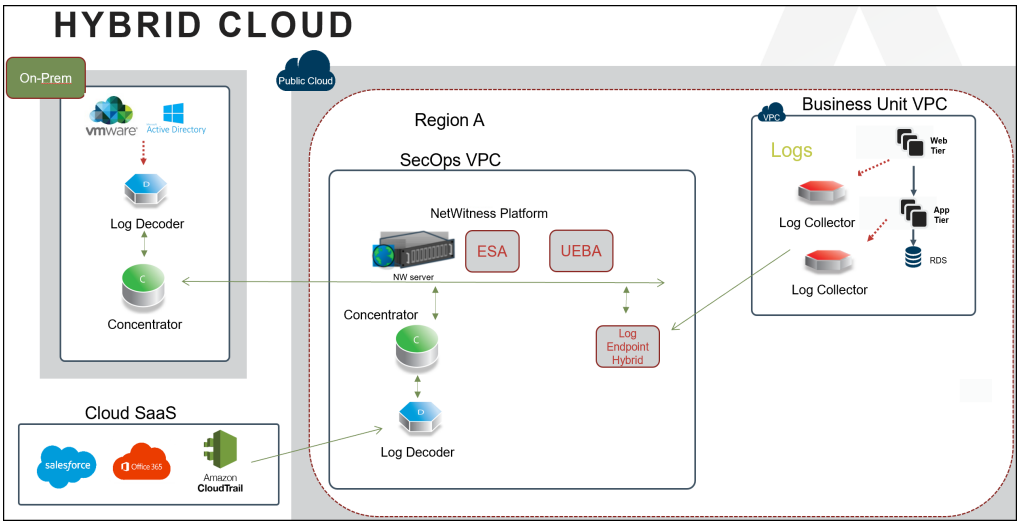
NetWitness Full Stack VPC Visibility

This diagram shows all NetWitness components (full stack) deployed in GCP.



Hybrid Deployment

This diagram shows the Log Decoder and Concentrator deployed in GCP with all other NetWitness components deployed on premises.



GCP Deployment

This topic contains the rules and high-level tasks you must follow to deploy NetWitness components in the GCP.

Checklist

Step	Description
1	Prerequisites
2	Find NetWitness Platform GCP Images
3	Establish gcloud Environment
4	Create an Instance using Google Cloud SDK Shell
5	Installation Tasks
6	Configure Hosts (Instances) in NetWitness Platform

Prerequisites

You need the following items before you begin the integration process:

- Access to GCP console.
- Google Cloud SDK Toolkit installed.
- Network routability (and proper GCP firewall rules) for the instances to transfer data to the NetWitness.

Find NetWitness Platform GCP Images

There are two types of NetWitness GCP images. The following description guides you on which image is appropriate for your deployment:

- **Lite Image**: rsa-nw-12-0-0-0-19340-lite
  If you have an active NW Server or NetWitness software repository, use the Lite image. This image has a small footprint and does not contain the NetWitness software packages. When you set up NetWitness using the Lite image, it will require the IP address of the NW Server or NetWitness software repository. This image is available publicly, and you can deploy it without needing to contact NetWitness Customer Support.
- **Full Image**: rsa-nw-12-0-0-0-19340-full
  For fresh or new deployments, where an NW Server is deployed for the first time, use the Full image. This image contains the entire NetWitness software library and other files required to complete the

installation.
To get access to the Full image, open an NetWitness Customer Support case (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897) to get granted the proper image permissions.

Note: For more information on GCP image deployment, see the blog post Running NetWitness in Google Cloud.

Establish gcloud Environment

1. Download and install the Google Cloud SDK Toolkit (https://cloud.google.com/sdk).
2. Run the following gcloud commands to log in and set the proper project:

```
gcloud auth login
gcloud config set project <project name>
```

Create an Instance using Google Cloud SDK Shell

1. Determine what machine type is appropriate for selection. For more information, see GCP Instance Configuration Recommendations.
2. Using the image names listed under Find NetWitness Platform GCP Images, create a GCP Instance by running the command:

```
gcloud compute instances create <instance name> --image <netwitness image name> --image-project <rsa project> --machine-type <machine type> --zone <gcp zone> --network <gcp project network> --subnet <gcp project subnet> --no-address
```

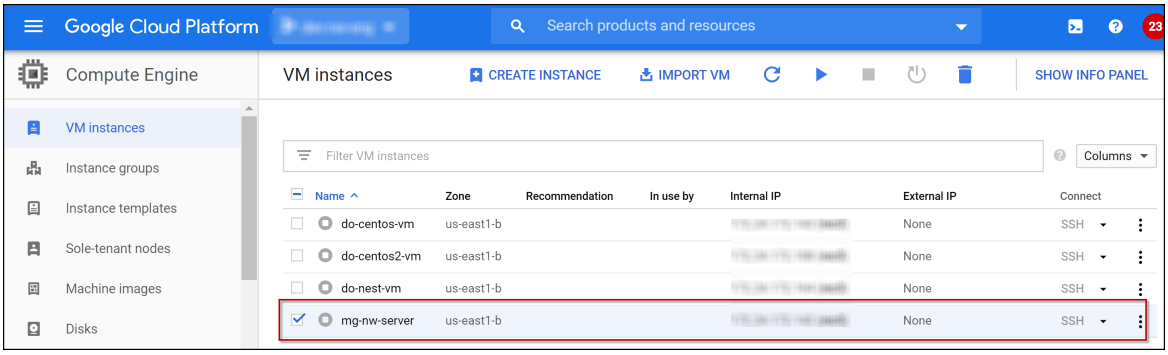
For example:

```
gcloud compute instances create nw-server --image rsa-nw-12-0-0-0-19340-full --image-project gcp-nw-prod-images --machine-type n1-standard-8 --zone us-west1-c --network rsa-network --subnet rsa-subnet --no-address
```

Note:
- The network and subnet values may vary based on the setup.
- The simple example is shown above, but there are many other options available. For more information, see the Quickstart: Creating a New Instance Using the Command Line section in the GCP documentation.

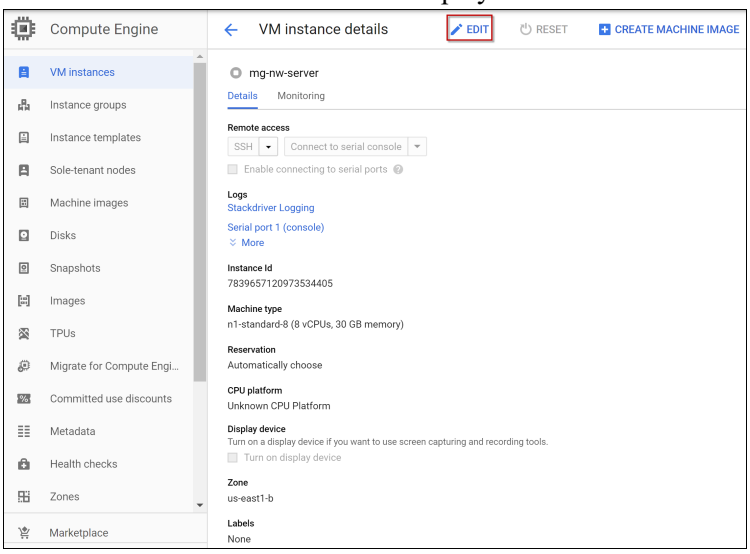
3. To modify the machine type, region, zone, or other configurations:
   a. Go to **Google Cloud Platform > Compute Engine > VM Instance** view to find the VM Instance.

b. Select the instance and click STOP.



c. Click the Instance name (**mg-nw-server**).

d. Click EDIT to modify the settings according to your preference and click **Save**. The **VM instance details** view is displayed.



e. SSH to the newly-created instance using the default NetWitness credentials.

Installation Tasks

Before you begin the installation tasks make sure you open the firewall ports. For more information on the lists of all the ports in a deployment, see the Network Architecture and Ports topic in the *Deployment Guide for NetWitness Platform 12.0.*

Caution: Do not proceed with the installation until the ports on your firewall are configured.

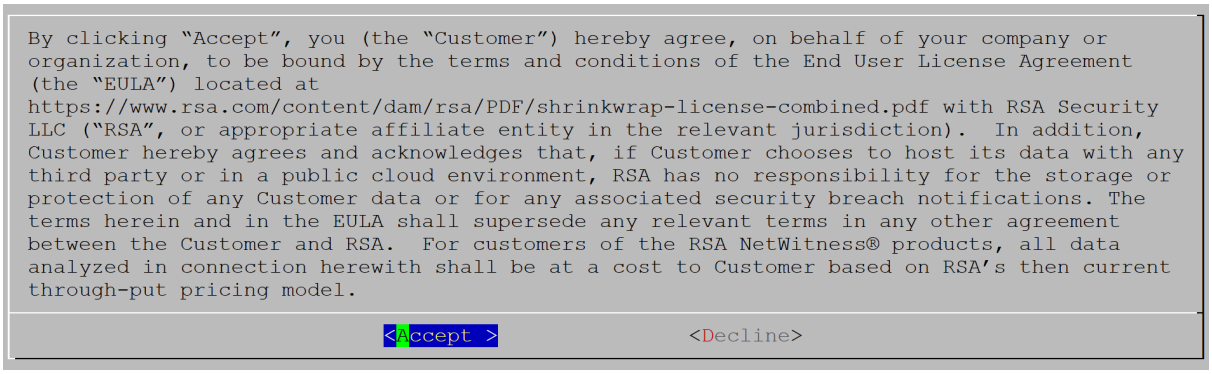
Tasks - Install 12.0.0.0 on the NetWitness Server (NW Server) Host and Component Hosts

Complete the following steps to install 12.0 on NW Server host and other component hosts. Steps that are specific to the NW Server host or to component hosts are noted.

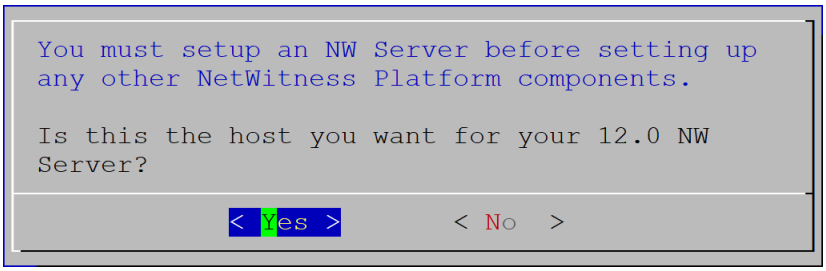
Caution: If you want to install the Endpoint Relay Server, do not run the nwsetup-tui script. Follow the instructions in (Optional) Installing and Configuring Relay Server in the *Endpoint Configuration Guide for NetWitness Platform 12.0*.

1. Log in to the host with the `root` credentials and run the `nwsetup-tui` command to set up the host. This initiates the `nwsetup-tui` (Setup program) and the EULA is displayed.

   Note: Use the following options to navigate the Setup prompts.
   1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, and use the Tab key to move to and from commands (such as **<Yes>**, **<No>**, **<OK>**, and **<Cancel>**). Press **Enter** to register your command response and move to the next prompt.
   2.) The Setup program adopts the color scheme of the desktop or console you use to access the host.
   3.) If you specify DNS servers during the Setup program (`nwsetup-tui`) execution, they MUST be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` script to proceed. Any misconfigured DNS servers cause the Setup program to fail. If you need to reach a DNS server after setup that is unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see "(Optional) Task 1 - Re-Configure DNS Servers Post 12.0" in the "Post Installation Tasks" section in this guide.
   If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Update Repository** prompt in step 10 (the DNS servers are not defined so the system cannot access the external repo).

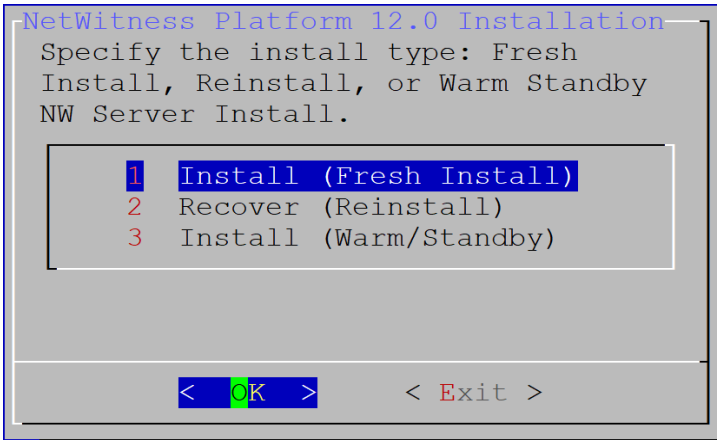
2. Tab to **Accept** and press **Enter**.
   The **Is this the host you want for your 12.0 NW Server** prompt is displayed.



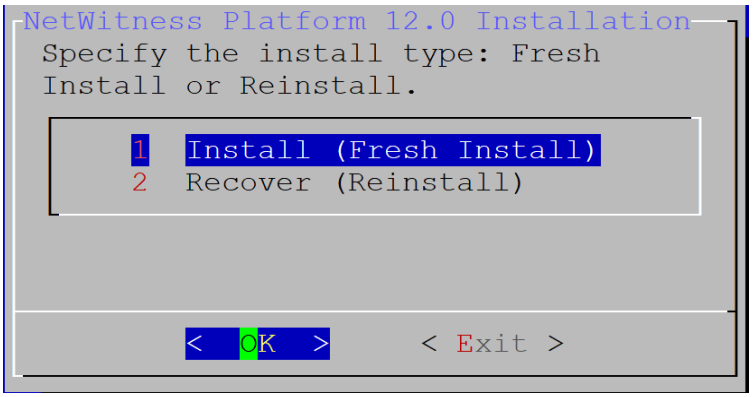
3. Tab to **Yes** and press **Enter** to install NW Server.
   Tab to **No** and press **Enter** to install Component Hosts.

   Caution: If you choose the wrong host for the NW Server and complete the setup, you must restart the setup program (step 2) and complete all the subsequent steps to correct this error.

4. The **Install** or **Recover** prompt is displayed.
   NetWitness Server Host prompt

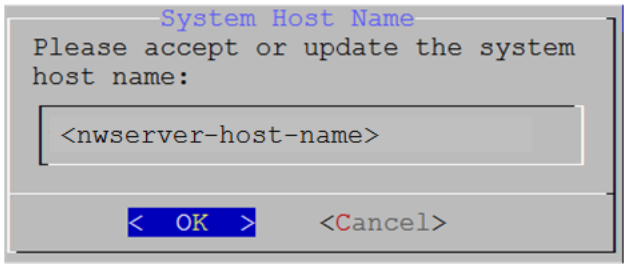


   Component Hosts prompt



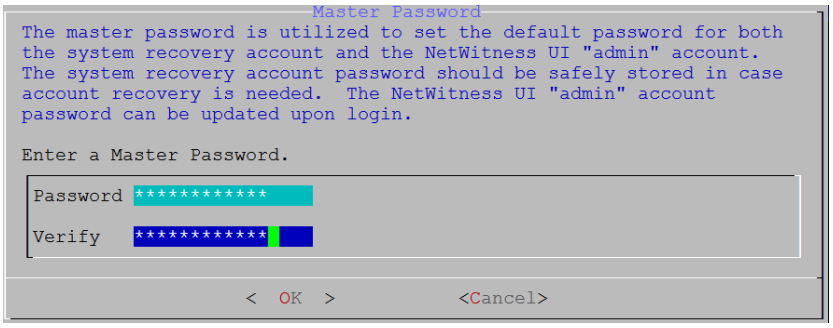
   Press **Enter**. By default, the **Install (Fresh Install)** option is selected.

5. The **System Host Name** prompt is displayed.



Press **Enter** if you want to keep this name. If not, edit the host name, tab to **OK**, and press **Enter** to change it.

6. The **Master Password prompt** is displayed.



Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

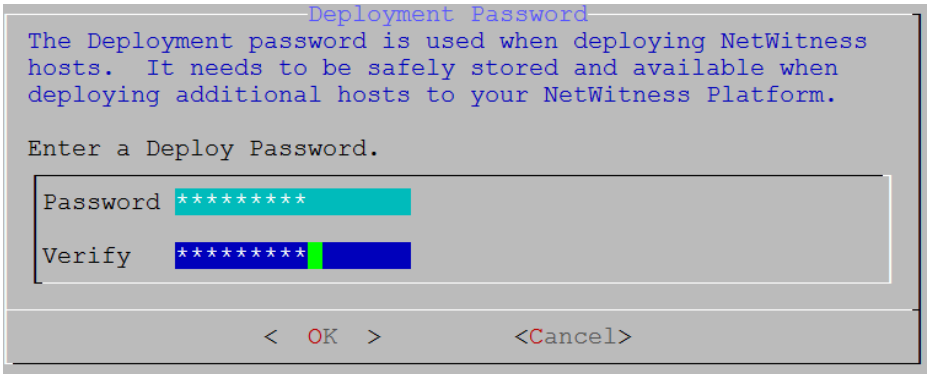
The following list of characters are supported for Master Password and Deployment Password:

- Symbols : ! @ # % ^ + ,
- Numbers :0-9
- Lowercase Characters : a-z
- Uppercase Characters : A-Z

No ambiguous characters are supported for Master Password and Deployment Password (for example: space { } [ ] ( ) / \ ' " ` ~ ; : . < > -.

Caution: This step 6 is applicable only for NW Server Host.

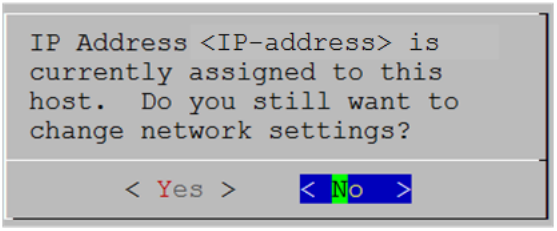
7. The **Deployment Password** prompt is displayed.



Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

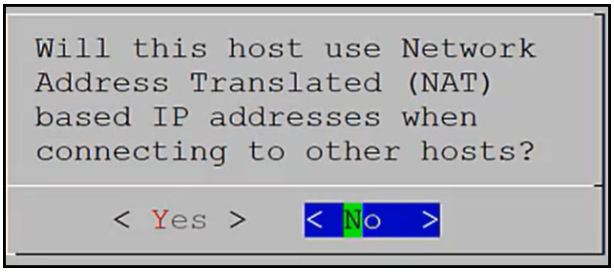
Caution: This step 7 is applicable for both NW Server and Component Hosts.

8. The setup program finds a valid IP address for this host and the following prompt is displayed.



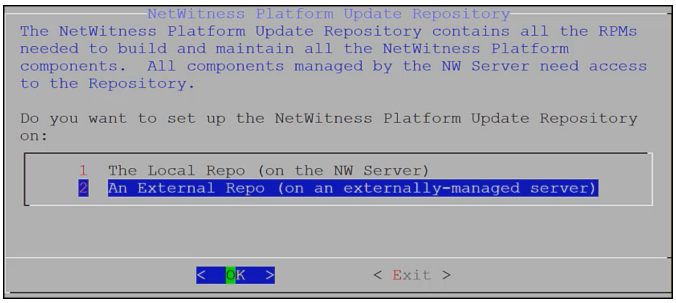
Tab to **No** and press **Enter**, if you want to use this IP and avoid changing your network settings.

9. The Use Network Address Translation (NAT) prompt is displayed.



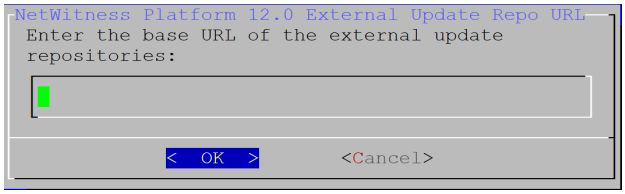
- For the NW Server, tab to **No** and press **Enter**.
- For Component Hosts, if this host requires the use of NAT-based addresses to communicate with the NW Server, tab to **Yes**. Otherwise, tab to **No** and press **Enter**.

10. The **NetWitness Platform Update Repository** prompt is displayed.



- For NW Server Host: Select **2 An External Repo (on an externally-managed server)**. Tab to **OK** and press **Enter**.
- For Component Hosts: Select **1 The Local Repo (on the NW Server)**. Tab to **OK** and press **Enter**.

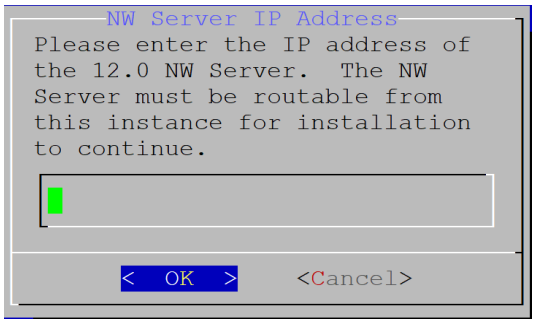
11. The **NetWitness Platform 12.0 External Update Repo URL** prompt is displayed.



Enter the base URL of the NetWitness Platform external repo, tab to **OK** and press **Enter**.

Caution: This step 11 is applicable only for NW Server Host.

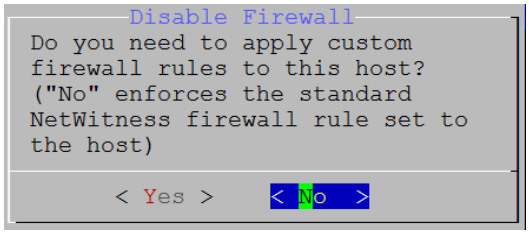
12. The **NW Server IP Address** prompt is displayed.



Enter the IP address of 12.0 NW Server host, tab to **OK** and press **Enter**.

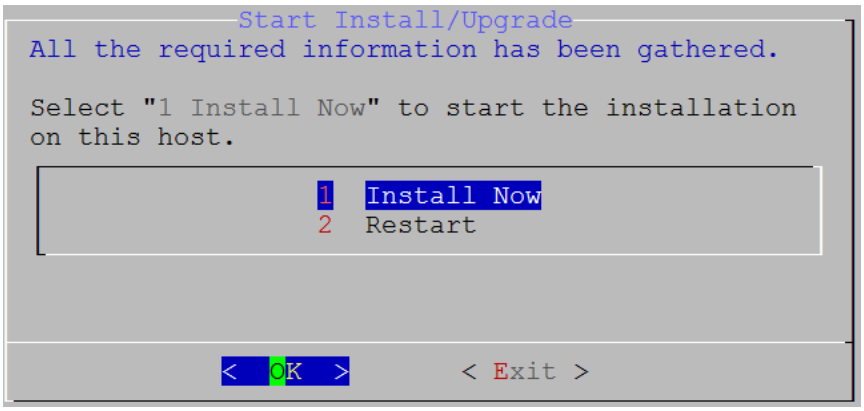
Caution: This step 12 is applicable only for Component Hosts.

13. The **Disable Firewall** prompt is displayed.



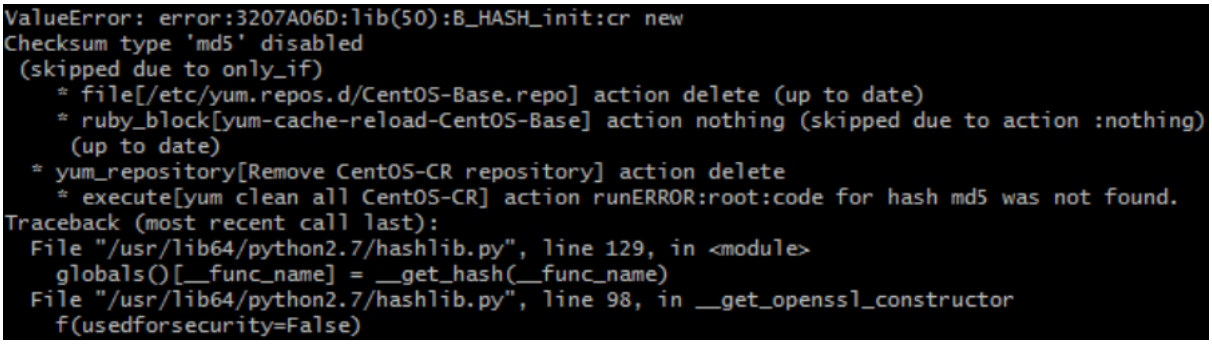
Tab to **No** and press **Enter**.

14. Press **Enter** to install 12.0 on the NW Server.
The **Start Install/Upgrade** prompt is displayed.



15. When **Installation complete** is displayed, you have installed the 12.0 NW Server on this host.

Note: Ignore the hash code errors similar to the errors shown in the following screen shot that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.



Note: If you want to perform a silent installation using CLI, see the Silent Installation Using CLI topic in the *Physical Host Installation Guide*.

Storage Configurations

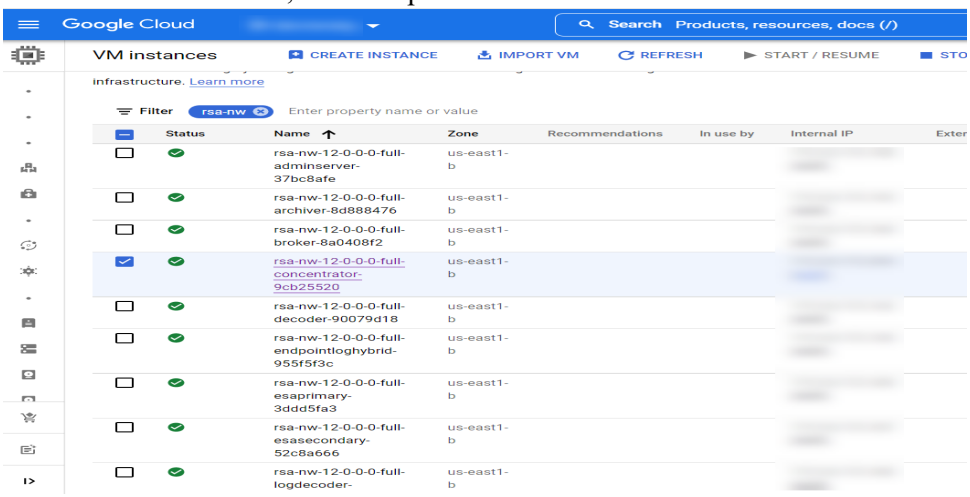
For storage allocations of all host types, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness® Platform 12.0*.

Create a Disk

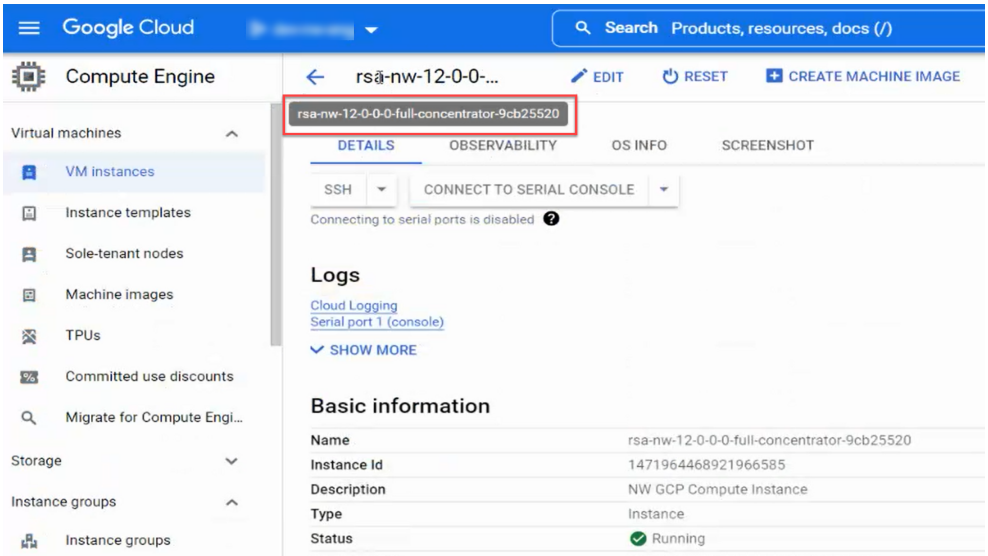
To understand which volumes are required to support the instance, see the Storage Requirements topic in the *Storage Guide for NetWitness® Platform 12.0*.

To create a disk:

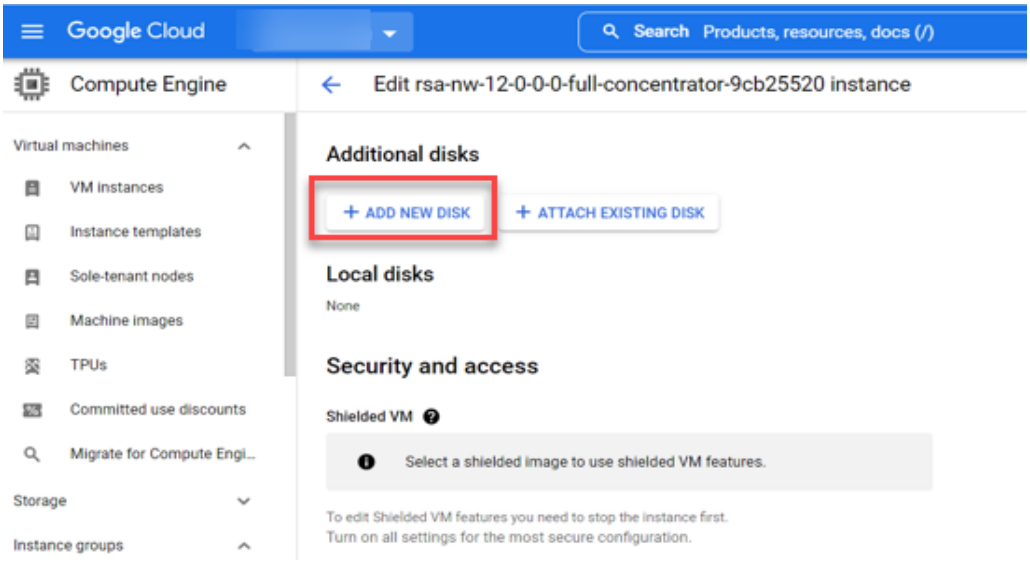
1. Go to **Google Cloud Platform > Compute Engine > VM Instance** to identify the VM instance.
2. Click the instance name, for example **rsa-nw-12-0-0-0-full-concentrator-9cb25520**.



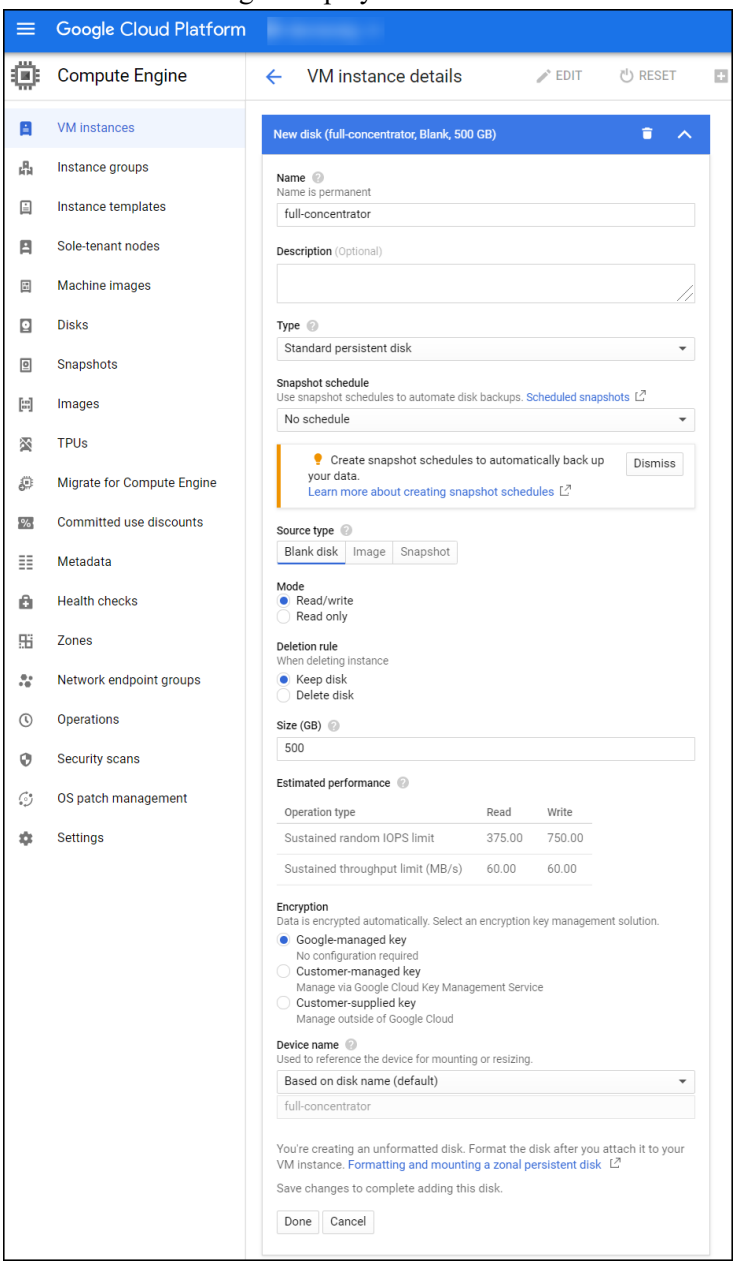
3. Click EDIT to make the required changes.
The **VM instance details** view is displayed.



4. Under **Additional disks** (Optional), click + Add new disk to add a new disk.



5. The **New disk** dialog is displayed.



Enter values in the following fields:

a. **Name**: Enter a unique name.

b. (Optional) **Description** : Enter the required additional information.

c. **Type**: Select **Standard persistent disk** from the drop-down menu.

d. **Snapshot schedule**: By default, **No schedule** option is selected.

e. Under **Source type > Blank disk**.
   - **Mode**: By default, **Read/write** option is selected.
   - **Deletion rule**: By default, **Keep disk** option is selected.

f. **Size (GB)**: Enter the required size.

g. Under **Encryption**: By default, **Google-managed key** is selected.

h. **Device name**: By default, **Based on disk name (default)** is selected.

i. Click **Done**.

6. Click **Save**.

Note: To add another disk to the instance type, repeat steps 1 to 6.

Configure Hosts (Instances) in NetWitness Platform

Configure individual hosts and services as described in NetWitness Host and Services Getting Started Guide. This guide also describes the procedures for applying updates and preparing for version upgrades.

GCP Instance Configuration Recommendations

Note: These recommendations can be used as a baseline for 12.0.0.0 and adjusted as needed.

Instance compute, and memory utilization will vary depending on content applied, ingestion rates, and the number of running queries.

This topic contains the minimum GCP instance configuration settings recommended for the NetWitness virtual stack components.

- Compute Engine Instance:
  - Minimum instance type - **n2-standard-32** is the minimum instance type required for any NetWitness component image so that it can function.
  - Machine type adjustments - You must adjust machine types according to your ingestion rate, content and parsers, dashboard reports, scheduled reports, investigations, and active users.
    - Ingestion rates of 15,000 EPS and 1.5 Gbps were used.
    - All the components were integrated.
    - The Log stream includes a Log Decoder, Concentrator, and Archiver.
    - The Endpoint Hybrid stream includes an Endpoint Server, Concentrator, and Log Decoder.
    - Respond receives alerts from the Reporting Engine, and Event Stream Analysis.
    - The background load includes reports, charts, alerts, investigation, and respond.
- Persistent Disk (Storage)

  For performance recommendations, recommended storage allocation per NetWitness host, and input/output operations per second, see the "Storage Requirements" topic in the *Storage Guide for NetWitness® Platform 12.0*.

  The following table displays the specification recommendations for NetWitness GCP instances.

Virtual Log Decoder (VLC)

Compute Engine Instance			
EPS	**Machine Type**	**Virtual CPU's**	**Memory**
5,000	n2-standard-4	4	16 GB
10,000	n2-standard-4	4	16 GB
15,000	n2-standard-8	8	32 GB

Archiver

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-4	4	16 GB
10,000	n2-standard-8	8	32 GB
15,000	n2-standard-16	16	64 GB

Broker

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-4	4	16 GB
10,000	n2-standard-4	4	16 GB
15,000	n2-standard-4	4	16 GB

Log Concentrator

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-8	8	32 GB
10,000	n2-standard-8	8	32 GB
15,000	n2-standard-16	16	64 GB

Event Stream Analysis (ESA)

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
9,000	n2-standard-8	8	32 GB
18,000	n2-standard-16	16	64 GB
30,000	n2-standard-32	32	128 GB

Log Decoder

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-8	8	32 GB
10,000	n2-standard-16	16	64 GB
15,000	n2-standard-32	32	128 GB

NetWitness Endpoint Hybrid

Compute Engine Instance			
Agents	Machine Type	Virtual CPU's	Memory
15,000 agents	n2-standard-48	48	192 GB

New Health and Wellness

Compute Engine Instance		
Machine Type	Virtual CPU's	Memory
n2-standard-4	4	16 GB

NetWitness Server and Co-Located Components

Compute Engine Instance		
Machine Type	**Virtual CPU's**	**Memory**
n2-standard-16	16	64 GB

Analyst UI

Compute Engine Instance		
Machine Type	**Virtual CPU's**	**Memory**
n2-standard-8	8	32 GB

UEBA

Compute Engine Instance		
Machine Type	**Virtual CPU's**	**Memory**
n2-standard-16	16	64 GB

NetWitness® Platform XDR

Version 12.0

Hosts and Services Getting Started Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Hosts and Services Basics **8**

What Is a Host? 8

What Is a Category? 8

What Is a Service? 9

Setting Up a Host 12

Maintaining Hosts 12

Update Version Naming Convention 12

Maintaining Services 13

Services Implemented with the NetWitness Server 13

Running in Mixed Mode 15

Functionality Gaps Encountered During in Staggered Updates 15

Examples of Staggered Updates 15

Example 1. Multiple Network Decoders and Concentrators, Alternative 1 16

Example 2. Multiple Network Decoders and Concentrators, Alternative 2 16

Example 3. Multiple Regions 17

Hosts and Services Set Up Procedures **18**

Step 1. Deploy a Host 18

Step 2. Install a Service on a Host 19

Step 3. Review SSL Ports for Trusted Connections 20

Encrypted SSL Ports 20

Step 4. Manage Access to a Service 21

Test a Trusted Connection 21

Hosts and Services Maintenance Procedures **25**

Apply Version Upgrades to a Host 27

Pre-Stage Host 27

Create and Manage Host Groups 28

Create a Group 29

Change the Name of a Group 29

Add a Host to a Group 29

View the Hosts in a Group 29

Remove a Host from a Group 30

Delete a Group 30

Search for Hosts 30

Search for a Host 31

Find the Host that Runs a Service 31

Execute a Task From the Host Task List ... 31
Add and Delete a Filesystem Monitor ... 34
Configure the Filesystem Monitor ... 34
Delete a Filesystem Monitor ... 35
Reboot a Host ... 35
Shut Down and Restart a Host from the Hosts View ... 36
Shut Down and Restart a Host from the Host Task List ... 36
Set Host Built-In Clock ... 36
Set the Time on the Local Clock ... 37
Set Network Time Source ... 37
Specify the Network Clock Source ... 38
Set SNMP ... 39
Toggle SNMP Service on the Host ... 39
Set Syslog Forwarding ... 40
Set Up and Start Syslog Forwarding ... 40
Show Network Port Status ... 42
Display the Network Port Status ... 42
Show Serial Number ... 42
Show the Serial Number ... 43
Shut Down Host ... 43
Shut Down the Host ... 44
Stop and Start a Service on a Host ... 44
Stop a Service on a Host ... 44
Start a Service on a Host ... 45
Add, Replicate, or Delete a Service User ... 46
Add a User Role to a Service ... 49
Change a Service User Password ... 51
Create and Manage Service Groups ... 53
Create a Group ... 53
Change the Name of a Group ... 54
Add a Service to a Group ... 54
View the Services in a Group ... 54
Remove a Service from a Group ... 55
Delete a Group ... 55
Duplicate or Replicate a Service Role ... 55
Duplicate a Service Role ... 56
Replicate a Role ... 57
Edit Core Service Configuration Files ... 57
Edit a Service Configuration File ... 58
Revert to a Backup Version of a Service Configuration File ... 58

Push a Configuration File to Other Services 59
Edit a Service Index File 59
Index and Custom Index Files 60
Configure the Task Scheduler 60
Scheduler File 60
Scheduler Task Syntax 60
Task Line Parameters 61
Messages 61
Sample Task Line 62
Enable the Crash Reporter Service 62
The crashreporter.cfg File 62
Configure the Crash Reporter Service 64
Start and Stop the Crash Reporter Service 65
Maintain the Table Map Files 65
Prerequisites 66
Edit or Delete a Service 67
Edit a Service 68
Delete a Service 69
Explore and Edit Service Property Tree 69
Display or Edit a Service Property 70
Send a Message to a Node 70
Terminate a Connection to a Service 71
Terminate a Session on a Service 71
Terminate an Active Query in a Session 72
Search for Services 72
Search for a Service 72
Filter Services by Type 73
Find the Services on a Host 74
Start, Stop, or Restart a Service 75
Start a Service 75
Stop a Service 75
Restart a Service 75
View Service Details 76
Purpose of Each Service View 76
Access a Service View 76
View Topology Details 78
Centralized Service Configuration via Policy 79
Creating Groups and Policies 80
Create a Group 80
Create a Policy 81

Managing Groups and Policies 82
View a Group 82
Delete a group 82
Edit a Group 83
Filter Groups 83
Delete a Policy 83
Revert a Policy 84
Clone a Policy 84
Edit a Policy 84
View a Policy 85
Filter Policies 85
Hosts and Services Views References 86
Hosts View 87
Services View 92
Edit Service Dialog 99
Services Config View 102
Services Config View - Appliance Service Configuration Tab 105
Services Config View - Data Retention Scheduler Tab 107
Services Config View - Files Tab 110
Services Explore View 113
Services Explore View - Properties Dialog 116
Services Logs View 119
Services Security View 122
Services Security View - Users Tab 124
Services Security View - Roles Tab 130
Services Security View - Service User Roles and Permissions 132
Services Security View - Aggregation Role 135
Services Security View - Settings Tab 136
Services Stats View 140
Services Stats View - Chart Stats Tray 145
Services Stats View - Gauges 148
Services Stats View - Timeline Charts 149
Services System View 152
Services Topology View 155
Services System View - Host Task List Dialog 157
Service Configuration Settings 160
Aggregation Configuration Parameters 160
Appliance Service Configuration Parameters 162
Archiver Service Configuration View 162
Broker Service Configuration Parameters 164

Concentrator Service Configuration Parameters ... 165
Core Service Logging Configuration Parameters ... 165
Core Service-to-Service Configuration Parameters ... 166
Core Service System Configuration Parameters ... 167
Decoder Configuration Parameters ... 168
Network Decoder Service Configuration Parameters ... 171
Log Decoder Service Configuration Parameters ... 171
REST Interface Configuration Parameters ... 174
NetWitness Platform Core Service system.roles Modes ... 174
Centralized Service Configuration via Policy Settings ... 176
Centralized Service Configuration – Groups Tab ... 176
Create Group Dialog ... 178
Define Group Dialog ... 179
Assign Policy Dailog ... 182
Centralized Service Configuration – Policies Tab ... 185
Create Policy ... 187
Define Policy Settings ... 189
Assign to Group ... 191
Troubleshooting Version Installations and Updates ... 193
deploy_admin User Password Has Expired Error ... 194
Downloading Error ... 195
Error Deploying Version <version-number> Missing Update Packages ... 196
Upgrade Failed Error ... 196
External Repo Update Error ... 198
Host Installation Failed Error ... 198
Host Update Failed Error ... 199
Missing Update Packages Error ... 200
OpenSSL 1.1.x ... 201
Patch Update to Non-NW Server Error ... 201
Reboot Host After Update from Command Line Error ... 202
Reporting Engine Restarts After Upgrade ... 202
Log Collector Service (nwlogcollector) ... 204
NW Server ... 205
Orchestration ... 206
Reporting Engine Service ... 206
Event Stream Analysis ... 206

Hosts and Services Basics

This guide gives administrators the standard procedures for adding and configuring hosts and services in NetWitness. After introducing you to the basic purpose of hosts and services and how they function within the NetWitness network, this guide covers:

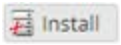
- Tasks you must complete to set up hosts and services in your network
- Additional procedures that you complete based on the long-term and daily, operational needs of your enterprise
- Reference topics that describe the user interface

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

What Is a Host?

A host is the machine on which a service runs and can be a physical or virtual machine. See the "NetWitness Detailed Host Deployment Diagram" in the *NetWitness Deployment Guide* for an illustration of how hosts are deployed.

What Is a Category?

A category assigns a service or services to a host when you install a host from the Hosts view. You choose a host **Category** in the **Install Services** dialog which is displayed when you select a host in the Hosts view and click 

. The following table lists each category and the services it installs. See the "NetWitness Detailed Host Deployment Diagram" in the *NetWitness Deployment Guide* for an illustration of how hosts are deployed.

Category	Services Installed
Analyst UI	Investigate Server, Broker, NetWitness UI, Reporting Engine, Respond Server
Archiver	Workbench and Archiver
Broker	Broker
Concentrator	Concentrator
Endpoint	Endpoint Server
Endpoint Broker	Endpoint Broker Server
Endpoint Log Hybrid	Log Collector, Log Decoder, Endpoint Server, and Concentrator
ESA Primary	Contexthub Server and ESA Correlation

Category	Services Installed
ESA Secondary	ESA Correlation
Log Collector	Log Collector
Log Decoder	Log Collector and Log Decoder
Log Hybrid	Log Collector, Log Decoder, and Concentrator
Log Hybrid - Retention	Log Collector and Log Decoder (deployed on NetWitness Series 6 Hybrid hardware with Log Hybrid-Retention Optimization)
Malware Analysis	Malware Analysis and Broker
Network Decoder	Decoder (Packets)
Network Hybrid	Concentrator and Network Decoder
New Health and Wellness	Metrics Server
UEBA	UEBA
Warehouse Connector	Warehouse Connector

What Is a Service?

A service performs a unique function, such as collecting logs or archiving data. Each service runs on a dedicated port and is modeled as a plug-in to enable or disable, according to the function of the host.

You must configure the following Core services first:

- Network Decoder
- Concentrator
- Broker
- Log Decoder

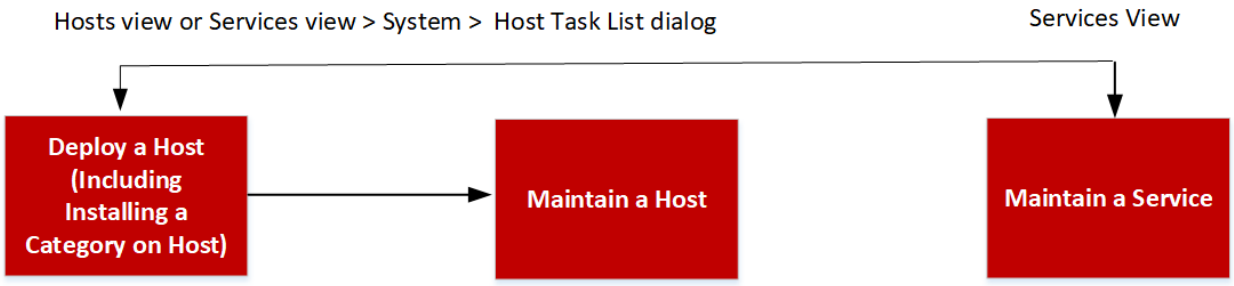
All the services are listed below and each service except the Log Collector has its own guide or shares a guide in the *Host and Services Configuration Guides* section of the NetWitness documentation page on NetWitness Community at https://community.netwitness.com/t5/netwitness-platform/ct-p/netwitness-documentation. The Log Collector has its own set of configuration guides to handle the configuration for all the supported event collection protocols. For Log Collector information, see *Log Collection Guides*.

Services	Notes
NW Server	

Services	Notes
Admin Config Content Integration Investigate License Orchestration Reporting Engine Respond Security	Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server
Analyst UI	
Broker Investigate Server NetWitness UI Reporting Engine Respond Server	Implemented with the Analyst UI Implemented with the Analyst UI Implemented with the Analyst UI Implemented with the Analyst UI Implemented with the Analyst UI
Archiver	
Archiver Workbench	Core Service
Broker	
Broker	Core Service
Concentrator	
Concentrator	Core Service
Endpoint	
Endpoint Server	
Endpoint Broker	
Endpoint Broker Server	
Endpoint Log Hybrid	
Log Collector Log Decoder Endpoint Server Concentrator	Core Service Core Service Core Service
ESA Primary	
Contexthub ESA Correlation	
ESA Secondary	

Services	Notes
ESA Correlation	
Log Collector	
Log Collector	Core Service
Log Decoder	
Log Collector Log Decoder	 Core Service
Log Hybrid	
Log Collector Log Decoder Concentrator	 Core Service Core Service
Log Hybrid - Retention	Deployed on Series 6 Hybrid hardware with Log Hybrid-Retention Optimization.
Log Collector Log Decoder	 Core Service
Malware Analysis	
Malware Analysis Broker	 Core Service
Network Decoder	
Decoder (Packets)	Core Service
Network Hybrid	
Concentrator Network Decoder	Core Service Core Service
New Health and Wellness	
Metrics Server	
UEBA	
UEBA	
Warehouse Connector	
Warehouse Connector	Command line installation

You must configure hosts and services to communicate with the network and each other so they can perform their functions such as storing or capturing data. For information about ports and a comprehensive list of ports for all services, see "Network Architecture and Ports" in the *Deployment Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.



Setting Up a Host

You use the Hosts view to add a host to NetWitness. See Step 1. Deploy a Host for detailed instructions.

Maintaining Hosts

You use the main Hosts view (**(Admin) > Hosts**) to add, edit, delete, and perform other maintenance tasks for the hosts in your deployment. You use the Host Task List dialog to perform tasks relating to a host and its communications with the network. See Hosts and Services Maintenance Procedures for detailed instructions.

After initial implementation of NetWitness, the major task you perform from the Hosts view is updating your NetWitness deployment to a new version.

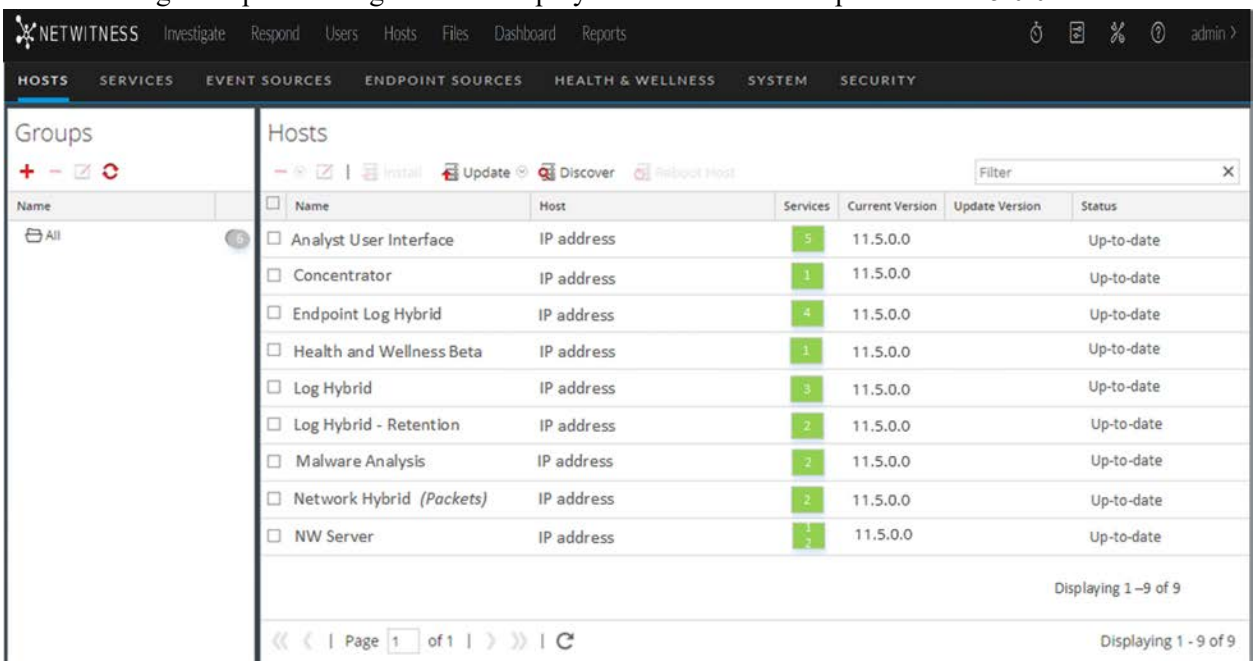
Update Version Naming Convention

You use the Hosts view to apply the latest version updates from your Hosts and Services Maintenance Procedures. You must understand the update version naming convention to know which version you want to apply to the host. The naming convention is ***major-release.minor-release.service-pack.patch***. For example, if you choose 11.6.1.2, you apply the following version to the host.

- 11 = major release
- 6 = minor release
- 1 = service pack
- 2 = patch

NetWitness supports multiple versions in your deployment. For more information, see Running in Mixed Mode. The NetWitness Server (NW Server Host) is updated first and all other hosts must have the same or earlier version as the NW Server Host.

The following example is a single version deployment with all hosts updated to 11.5.0.0.



Maintaining Services

You use the Services view (**(Admin) > Services**) to add, edit, delete, monitor, and perform other maintenance tasks for the services in your deployment. See Hosts and Services Procedures for detailed instructions.

Services Implemented with the NetWitness Server

The services in the following list are implemented when you deploy the NW Server to support:

- The expansion of physical and virtual deployment platforms and improvements to host and service maintenance.
- Content, Investigate, Respond, and Source functionality.

Caution: You do not need to configure these services to deploy NetWitness. recommends that you monitor the operating status of these services using Health-and-Wellness. Do not attempt to modify the parameters in the Explore view without contacting Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897).

Service	Purpose
Admin	The Administration (Admin) Server is the back-end service for administrative tasks in the NetWitness User Interface (UI). It abstracts authentication, global preferences management, and authorization support for the UI. The Admin server requires the Config server and the Security server to be online to perform its role.
Config	The Configuration (Config) Server stores and manages configuration sets. A configuration set is any logical configuration group that is managed independently. The Config server facilitates the sharing of properties among services, provides configuration backup and restore facilities, and tracks changes to properties.
Content	The Content server manages the NetWitness provided and user-created parser rules. For more information on parser management, search for "parsers" in NetWitness Community.
Integration	The Integration Server manages interactions with external systems. The service handles the following outbound or inbound channels. • REST API Gateway - gateway to external REST clients that assigns calls to the NetWitness Application Programming Interface (API). • Notifications Dispatcher - centralized dispatcher for all outbound notifications originating in the NetWitness deployment.
Investigate	The Investigate server supports Investigate and Malware Analysis functionality. For more information see the *NetWitness Investigate User Guide*.
Orchestration	The Orchestration server provisions, installs, and configures all services in your NetWitness deployment.
Respond	The Respond server supports Respond functionality. For more information see the *NetWitness Respond Configuration Guide*.

Service	Purpose
Security	The NetWitness Security Server (**Security server**) manages the security infrastructure of a NetWitness deployment. It handles the following security-related concerns. • Users and the authentication accounts • Role Based Access Control (RBAC) • Deployment PKI infrastructure A NetWitness deployment has users with authentication accounts. Independent of how you verify the identity of the analyst (for example, Active Directory), NetWitness must maintain the user state, which is not provided by all authentication providers (for example, last login time, failed login attempts, and roles). The concept of a user is separate from the identify associated with the user and the **Security server** maintains these as separate User and Account entities. In addition to the out-of-the-box local NetWitness accounts available to all NetWitness deployments, the server supports external authentication providers. The **Security server** also implements RBAC by managing Role and Permission entities. Permissions can be assigned to roles and roles to users. Together these enable a flexible authorization policy for the deployment. The server also manages generation of cryptographically secure tokens that encode the applicable authorization for a user. These tokens form the basis for deployment-wide authorization.
Source	The Source server is reserved for future use and will provide a centralized location to configure sources (for example, Endpoints and Log Sources).

Running in Mixed Mode

Mixed mode occurs when some services are updated to the latest version and some are still on older versions. This happens when you update the hosts in your deployment to the latest version in phases (or stagger the update).

Functionality Gaps Encountered During in Staggered Updates

If you stagger the update, you:

- May not have all the features operational until you update your entire deployment.
- Will not have service administrative features available until you update all the hosts in your deployment.
- May be without data capture for a period of time.

Examples of Staggered Updates

In the following examples, all the hosts are on 11.4.0.0 and you want to stagger the host updates to version 11.4.1.0.

Example 1. Multiple Network Decoders and Concentrators, Alternative 1

In this example, the 11.4.0.0 deployment includes one NW Server host, two Network Decoder hosts, two Concentrator hosts, one Archiver host, one Broker host, one Event Stream Analysis host, one Endpoint Log Hybrid host, and one Malware Analysis host.

You must complete Session 1 first and update the hosts in the order listed.

NetWitness recommends that you update the Sessions 2 and 3 hosts in the order listed.

Session 1: Update Essential Hosts

1. Update the NetWitness Server host.
2. Update the Event Stream Analysis host.
3. Update the Endpoint Log Hybrid host.
4. Update the Malware Analysis host.
5. Update the Broker host.

Session 2: Update Other Hosts

1. Update the two Network Decoder hosts.
2. Update the two Concentrator hosts and the Archiver host.

Session 3: Update Other Hosts

1. Update all other hosts.

Example 2. Multiple Network Decoders and Concentrators, Alternative 2

In this example, the 11.4.0.0 deployment includes one NW Server host, two Network Decoder hosts, two Concentrator hosts, one Broker host, one Event Stream Analysis host, one Endpoint Log Hybrid host and one Malware Analysis host.

You must complete Session 1 first and update the hosts in the order listed.

NetWitness recommends that you update the Sessions 2 and 3 hosts in the order listed.

Session 1: Update Essential Hosts

1. Update the NetWitness Server host.
2. Update the Event Stream Analysis host.
3. Update the Endpoint Log Hybrid host.
4. Update the Malware Analysis host.
5. Update the Broker host.

Session 2: Update Other Hosts

1. Update one Network Decoder host and one Concentrator host.

Note: It does not matter which of the Network Decoder hosts or which of the Concentrator hosts you update first.

Time elapses during which NetWitness processes a significant amount of data.

Session 3: Update Other Hosts

1. Update the second Network Decoder host and the second Concentrator host.
2. Update all Log Decoder hosts before you update Virtual Log Collectors.
3. Update all other hosts.

Example 3. Multiple Regions

In this example, the 11.4.0.0 deployment includes one NW Server host, one Event Stream Analysis host, one Endpoint Log Hybrid host, one Malware Analysis host. Additionally, there are two sites, each with two Network Decoders, two Concentrators, and one Broker, for a total of four Network Decoder hosts, four Concentrator hosts, and two Broker hosts.

Session 1: Update Essential Hosts and Site 1

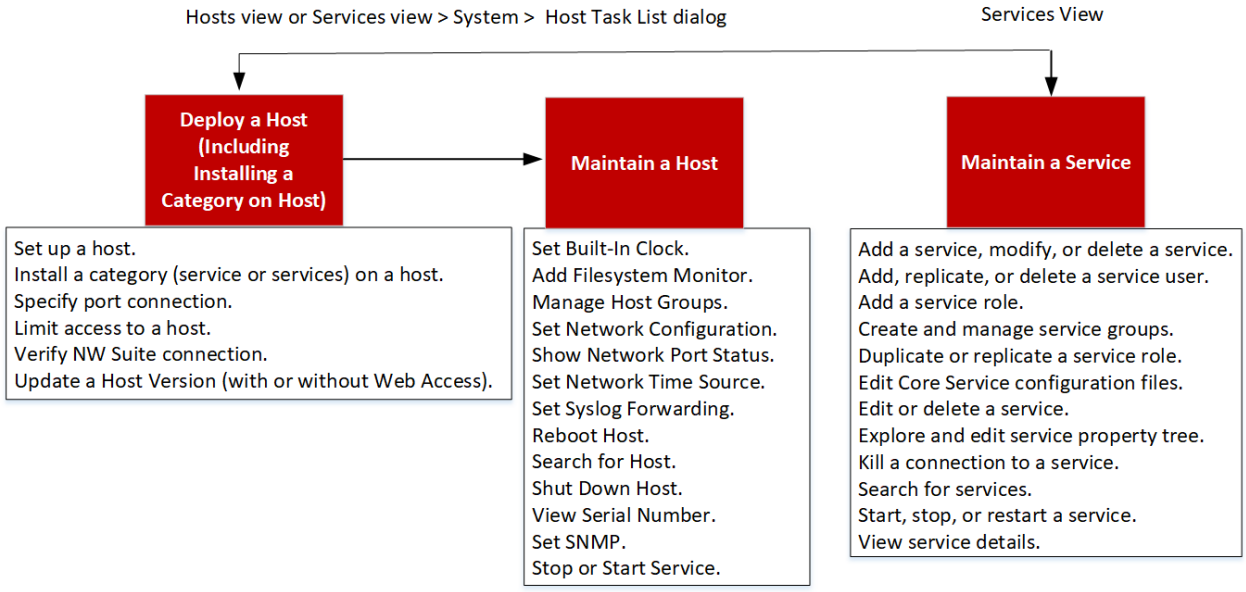
1. Update the NW Server host.
2. Update the Event Stream Analysis host.
3. Update the Endpoint Log Hybrid host.
4. Update the Malware Analysis host.
5. Update one Broker host, two Network Decoder hosts, and two Concentrator hosts.

Session 2: Update Other Hosts and Site 2

1. Update the second Broker host.
2. Update the two remaining Network Decoder hosts.
3. Update the two remaining Concentrator hosts.
4. Update all the other hosts.

Hosts and Services Set Up Procedures

Every service requires a host. After you set up a host, you can assign services to and from this host to other hosts in your NetWitness deployment. This topic contains information about basic procedures. For additional procedures, see Hosts and Services Maintenance Procedures.



High-Level Task	Description
Set Up a Host	Complete the following tasks in the order shown to set up a host. Step 1. Deploy a Host Step 2. Install a Service on a Host Step 3. Review SSL Ports for Trusted Connections Step 4. Manage Access to a Service

Step 1. Deploy a Host

Caution: If you include "." in a host name, the host name must also include a valid domain name.

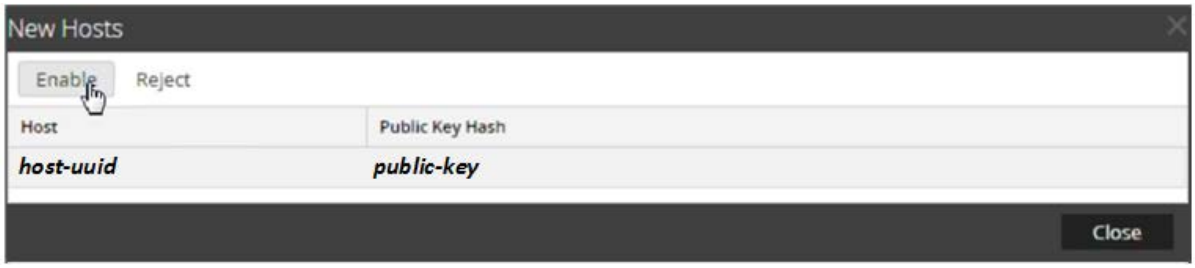
1. Deploy a host.

   You can deploy a physical host (NetWitness Appliance), virtual host on-prem, a virtual in AWS, a virtual host in Azure, or a virtual host on Google Cloud Platform. See the following guides for instructions on how to deploy hosts.

   - *Physical Host Installation Guide*
   - *Virtual Host Installation Guide*
   - *AWS Installation Guide*

- *Azure Installation Guide*
- *GCP Installation Guide*

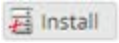
2. Go to  **(Admin) > Hosts**.
   The New Hosts dialog is displayed with the hosts that you deployed.
3. Select the hosts that you want to enable.
   The Enable menu option becomes active.
4. Click **Enable**.

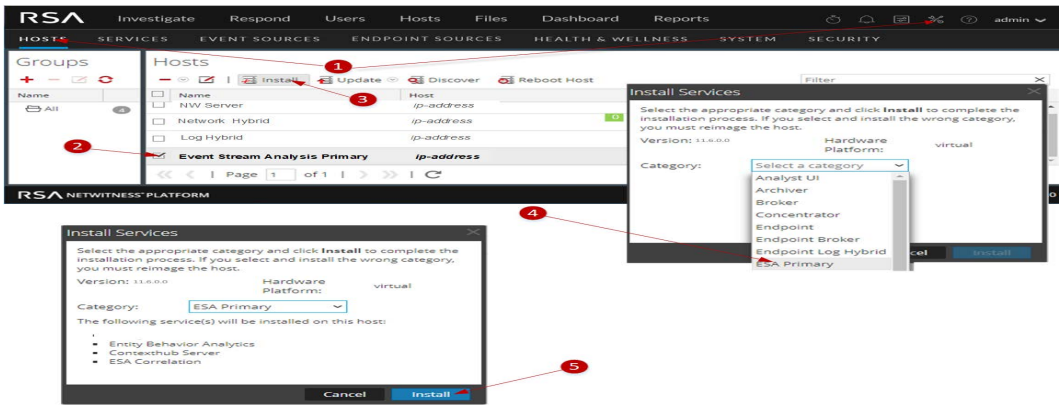


5. Select the host you enabled.
   The host is displayed in the Hosts view. At this point, you can install a service on the host.

Step 2. Install a Service on a Host

Perform the following steps to install a service on a host.

1. In NetWitness, go to  **(Admin) > Hosts**.
   The Hosts view is displayed.
2. Select the host on which you want to install the service (for example, **Event Stream Analysis**).
3. Click  in the toolbar.
   The Install Services dialog is displayed.
4. Select a service from the **Category** drop-down list (for example, **ESA Primary**).
5. Click  in the Install Services dialog.



6. A pop-up listing all the services already installed on this host is displayed. If there are no services installed, this pop-up will not be displayed.
7. Click **Yes** to install the new service.

Step 3. Review SSL Ports for Trusted Connections

To support trusted connections each core service has two ports, an unencrypted non-SSL port and an encrypted SSL port. Trusted connections require the encrypted SSL port.

Encrypted SSL Ports

By default, trusted connections are established with two settings:

- SSL is enabled.
- Core service is connected to an encrypted SSL port.

Each NetWitness Core service has two ports:

- Unencrypted non-SSL port
  Example: Archiver 50008
- Encrypted SSL port
  Example: Archiver 56008

The SSL port is the non-SSL port + 6000.

For information about ports and a comprehensive list of ports for all services, see "Network Architecture and Ports" in the *Deployment Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Step 4. Manage Access to a Service

In a trusted connection, a service explicitly trusts the NW Server to manage and authenticate users. With this trust, services in  **(Admin) > Services** no longer require credentials to be defined for every NetWitness Core service. Instead, users who have been authenticated by the server can access the service without entering another password.

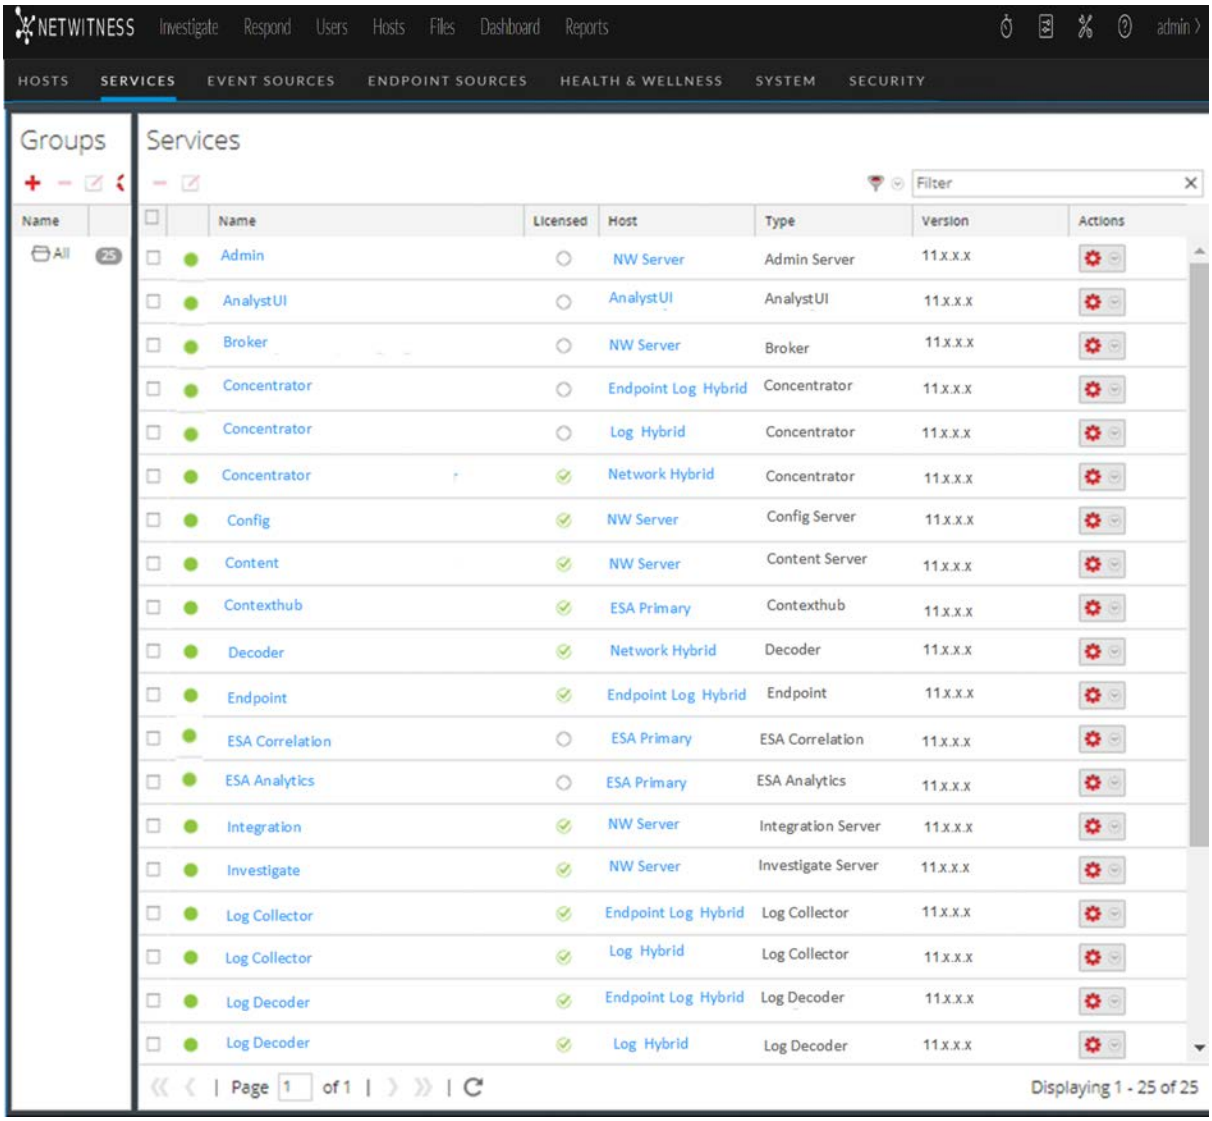
Test a Trusted Connection

Prerequisites

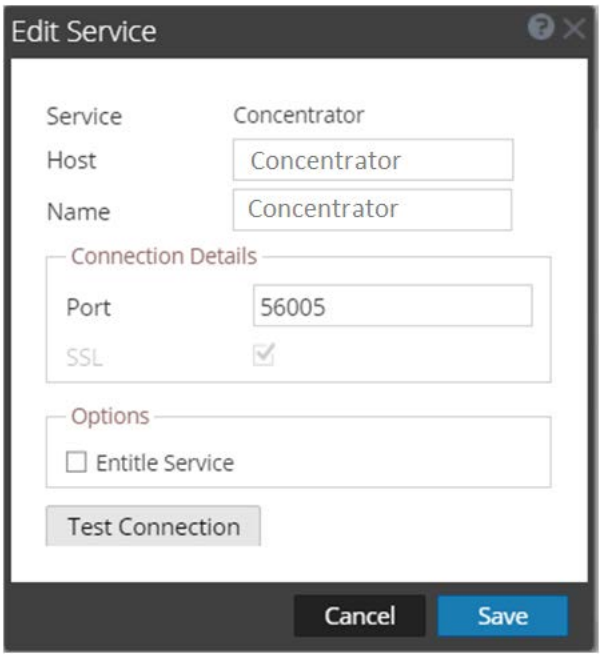
1. The administrator must assign a role to the user.
   For more information, see "Add a User and Assign a Role" in the *System Security and User Management Guide*.
2. The user must:
   - Log in to NetWitness for the server to authenticate the user.
   - Have access to the service.

Procedure

1. In NetWitness, go to  **(Admin) > Services**.
   The Services view is displayed.

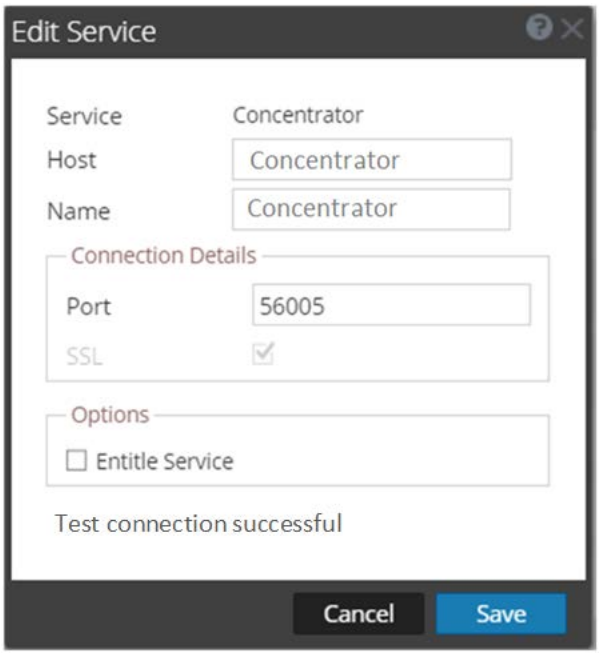

NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Groups
Services
Filter
Name
All
25
Name
Licensed
Host
Type
Version
Actions
Admin
NW Server
Admin Server
11.x.x.x
AnalystUI
AnalystUI
AnalystUI
11.x.x.x
Broker
NW Server
Broker
11.x.x.x
Concentrator
Endpoint Log Hybrid
Concentrator
11.x.x.x
Concentrator
Log Hybrid
Concentrator
11.x.x.x
Concentrator
Network Hybrid
Concentrator
11.x.x.x
Config
NW Server
Config Server
11.x.x.x
Content
NW Server
Content Server
11.x.x.x
Contexthub
ESA Primary
Contexthub
11.x.x.x
Decoder
Network Hybrid
Decoder
11.x.x.x
Endpoint
Endpoint Log Hybrid
Endpoint
11.x.x.x
ESA Correlation
ESA Primary
ESA Correlation
11.x.x.x
ESA Analytics
ESA Primary
ESA Analytics
11.x.x.x
Integration
NW Server
Integration Server
11.x.x.x
Investigate
NW Server
Investigate Server
11.x.x.x
Log Collector
Endpoint Log Hybrid
Log Collector
11.x.x.x
Log Collector
Log Hybrid
Log Collector
11.x.x.x
Log Decoder
Endpoint Log Hybrid
Log Decoder
11.x.x.x
Log Decoder
Log Hybrid
Log Decoder
11.x.x.x
Page
1
of 1
Displaying 1 - 25 of 25

2. Select the checkbox of the service (for example, a Concentrator) to test and click .
   The **Edit Service** dialog is displayed.



Note: The Options box will only display if the selected service is not licensed. A licensed service is denoted by a  in the Services view.

3. Remove the username to test the connection without credentials.
4. Click **Test Connection**.

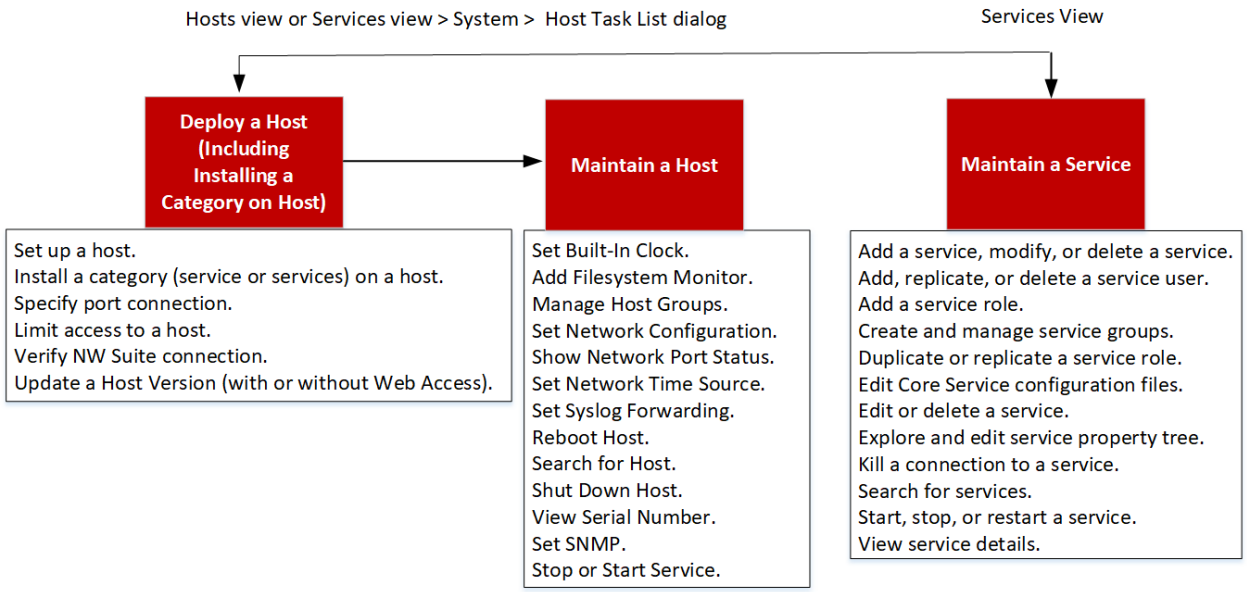


The message `Test connection successful` confirms the trusted connection is established. The previously authenticated user can access the service without typing a username and password on the service.

5. Click **Save**.

Hosts and Services Maintenance Procedures

Every service requires a host. After you set up a host, you can assign services to and from this host to other hosts in your NetWitness deployment.



High-Level Task	Description
Maintain a Host - Basics	The following maintenance tasks are shown in alphabetical order. • Apply Version Upgrades to a Host • Create and Manage Host Groups • Search for Hosts • Set Network Time Source • Show Network Port Status • Show Serial Number • Shut Down Host • Stop and Start a Service on a Host

High-Level Task	Description
Maintain a Host from the Host Task List Dialog	You use the Host Task List dialog to manage tasks that relate to a host and its communications with the network. Several service and host configuration options are available for Core hosts. • Execute a Task From the Host Task List • Add and Delete a Filesystem Monitor • Reboot a Host • Set Host Built-In Clock • Set Network Time Source • Set SNMP • Set Syslog Forwarding • Show Network Port Status • Show Serial Number • Shut Down Host • Stop and Start a Service on a Host
Maintain a Service	The following procedures describe how to maintain services. • Add, Replicate, or Delete a Service User • Add a User Role to a Service • Change a Service User Password • Create and Manage Service Groups • Duplicate or Replicate a Service Role • Edit Core Service Configuration Files • Edit or Delete a Service • Explore and Edit Service Property Tree • Terminate a Connection to a Service • Search for Services • Start, Stop, or Restart a Service • View Topology Details • Services System View • Centralized Service Configuration via Policy

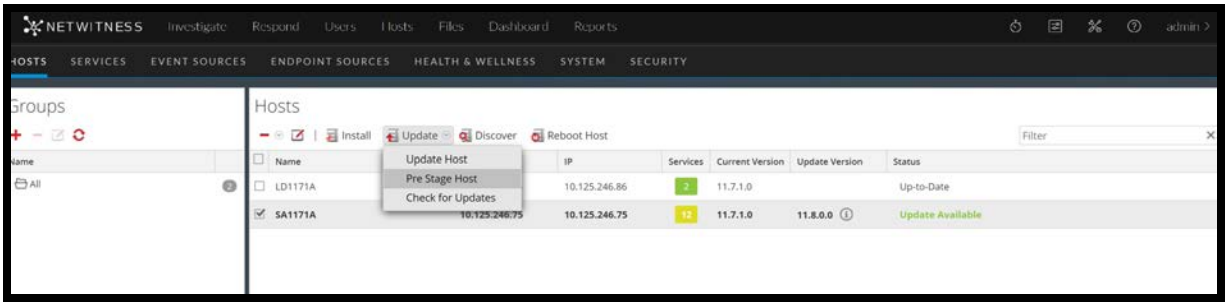
Apply Version Upgrades to a Host

Upgrade guides are available for all supported versions of NetWitness. Refer to the upgrade guide for the latest version in the *Install and Upgrade* section of the NetWitness documentation page on NetWitness Link.

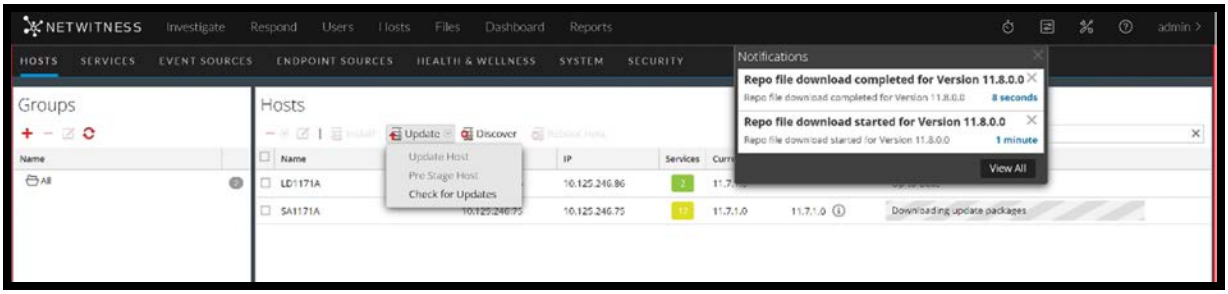
Pre-Stage Host

Note: The **Pre Stage Host** option is enabled for upgrade on the Admin Server. This option will be available after you upgrade to 11.7.1 and can be used to upgrade to versions later than 11.7.1.

1. Go to  > Hosts.
2. Click **Update** > **Check for Updates**.
   All possible update versions will displayed in the Versions drop-down list.
3. Click **Update** > **Pre Stage Host** and select the version in the update version column.
   A confirmation message for downloading the files is displayed.



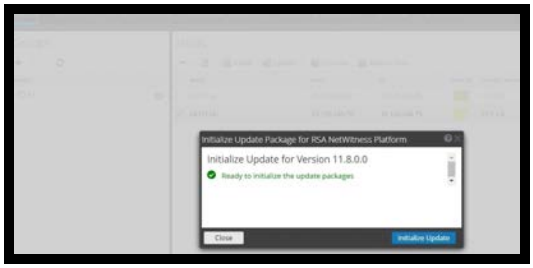
4. Click **Yes** to download the upgrade packages to the repo.
5. Verify the status of the download in the notifications tray as shown below.
   The **Pre Stage Host** and **Upgrade Host** will be disabled until pre stage is completed.



Note: The current version and the update version in the UI will be the same during the pre stage as it is not the actual update. This is because only the repo files are downloaded and no actual upgrade is done. The version will change only after upgrade.

6. If the download is successful, **Check for Updates** again to start the initialization.

7. Click on **Initialize Update**.
   The initialization of the package will take some time as the files are large and will need to be unzipped.



IMPORTANT: Pre Stage Repo preparation steps from 1 to 4 can be performed at any time. However, from steps 5 to 8 the upgrade process begins and you must NOT reboot the host or restart the jetty server during this time as it will corrupt the .ZIP files.

8. Check the status of initialization in the notifications tray.
9. After the initialization is completed successfully, click **Update** > **Update Host.**
   After the host is updated, you will be prompted to reboot the host.



10. Set up the host and reboot the host.

Create and Manage Host Groups

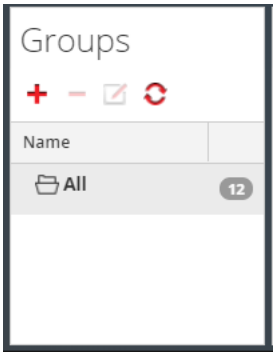
The Hosts view provides options for creating and managing groups of hosts. The Groups panel toolbar includes options for creating, editing, and deleting host groups. Once groups are created, you can drag individual hosts from the Hosts panel into a group.

Groups may reflect functional, geographical, project-oriented, or any other organization principle that is useful. A host may belong to more than one group. Here are some examples of possible groupings:

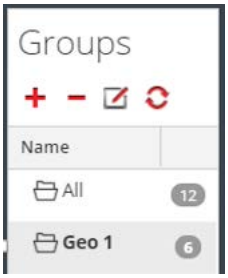
- Group different categories to make it easier to configure and monitor all Brokers, Network Decoders, or Concentrators.
- Group hosts that are part of the same data flow; for example, a Broker, and all associated Concentrators and Network Decoders.
- Group hosts according to their geographic region and location within the region. If a major power outage occurs in a location, potentially affected hosts are easily identifiable.

Create a Group

1. Select **(Admin) > Hosts**.
   The Hosts view is displayed.
2. In the **Groups** panel toolbar, click +.
   A field for the new group opens with a blinking cursor.



3. Type the name of the new group in the field (for example, `Geo 1`) and press **Enter**.
   The group is created as a folder in the tree. The number next to the group indicates the number of hosts in that group.



Change the Name of a Group

1. In the Hosts view **Groups** panel, double-click the group name, or select the group and click .
   The name field opens with a blinking cursor.
2. Type the new name of the group and press **Enter**.
   The name field closes and the new group name is displayed in the tree.

Add a Host to a Group

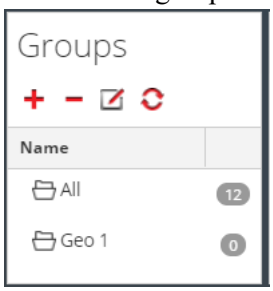
In the Hosts view **Hosts** panel, select a host and drag the host to a group folder in the Groups panel. The host is added to the group.

View the Hosts in a Group

To view the hosts in a group, click the group in the **Groups** panel.
The Hosts panel lists the hosts in that group.

Remove a Host from a Group

1. In the Hosts view **Groups** panel, select the group that contains the host that you want to remove. The hosts in that group appear in the Hosts panel.
2. In the **Hosts** panel, select one or more hosts that you want to remove from the group, and in the toolbar, select − > **Remove from Group**.
   The selected hosts are removed from the group, but are not removed from the NetWitness user interface. The number of hosts in the group, which is listed near the group name, decreases by the number of hosts removed from the group. The `All` group contains the hosts that were removed from the group.
   In the following example, the host group called `Geo 1` does not contain any hosts, because all the hosts in that group are removed.



Delete a Group

1. In the Hosts view **Groups** panel, select the group that you want to delete.
2. Click −.
   The selected group is removed from the Groups panel. The hosts that were in the group are not removed from the NetWitness user interface. The `All` group contains the hosts from the deleted group.

Search for Hosts

You can search for hosts from a list of hosts in the Hosts view. The Hosts view enables you quickly filter the list of hosts by Name and Host. It is possible to have numerous NetWitness hosts in use for various purposes. Instead of scrolling through the host list, you can quickly filter the host list to locate the hosts that you want to administer.

In the Services view, you can search for a service and quickly find the host that runs that service.

Search for a Host

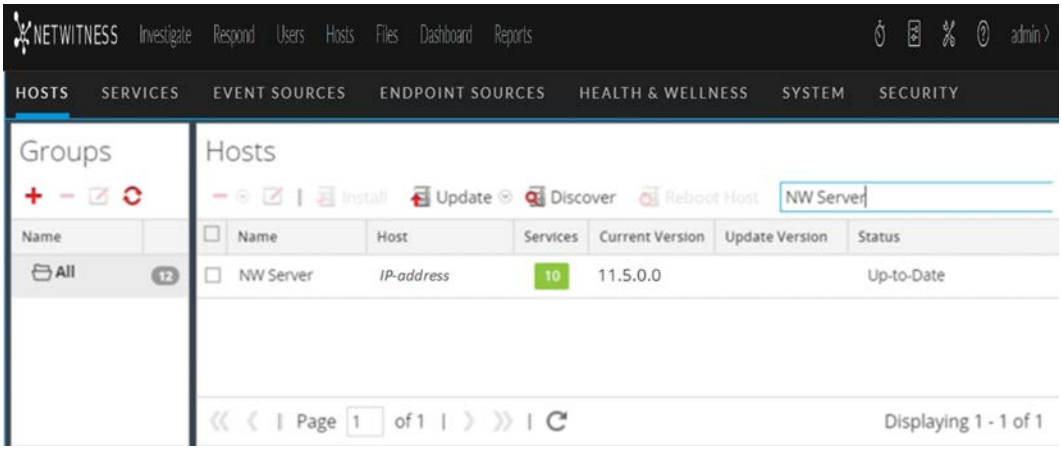
1. Select **(Admin) > Hosts**.
2. In the **Hosts** panel toolbar, type a host **Name** or **Hostname** in the **Filter** field.

   The Hosts panel lists the hosts that match the names entered in the Filter field.

Find the Host that Runs a Service

1. Select **(Admin) > Services**.
2. In the **Services** view, select a service. The associated host is listed in the Host column for that service.
3. To administer the host in the Hosts view, click the link in the **Host** column for that service. The host associated with the selected service is displayed in the Hosts view.

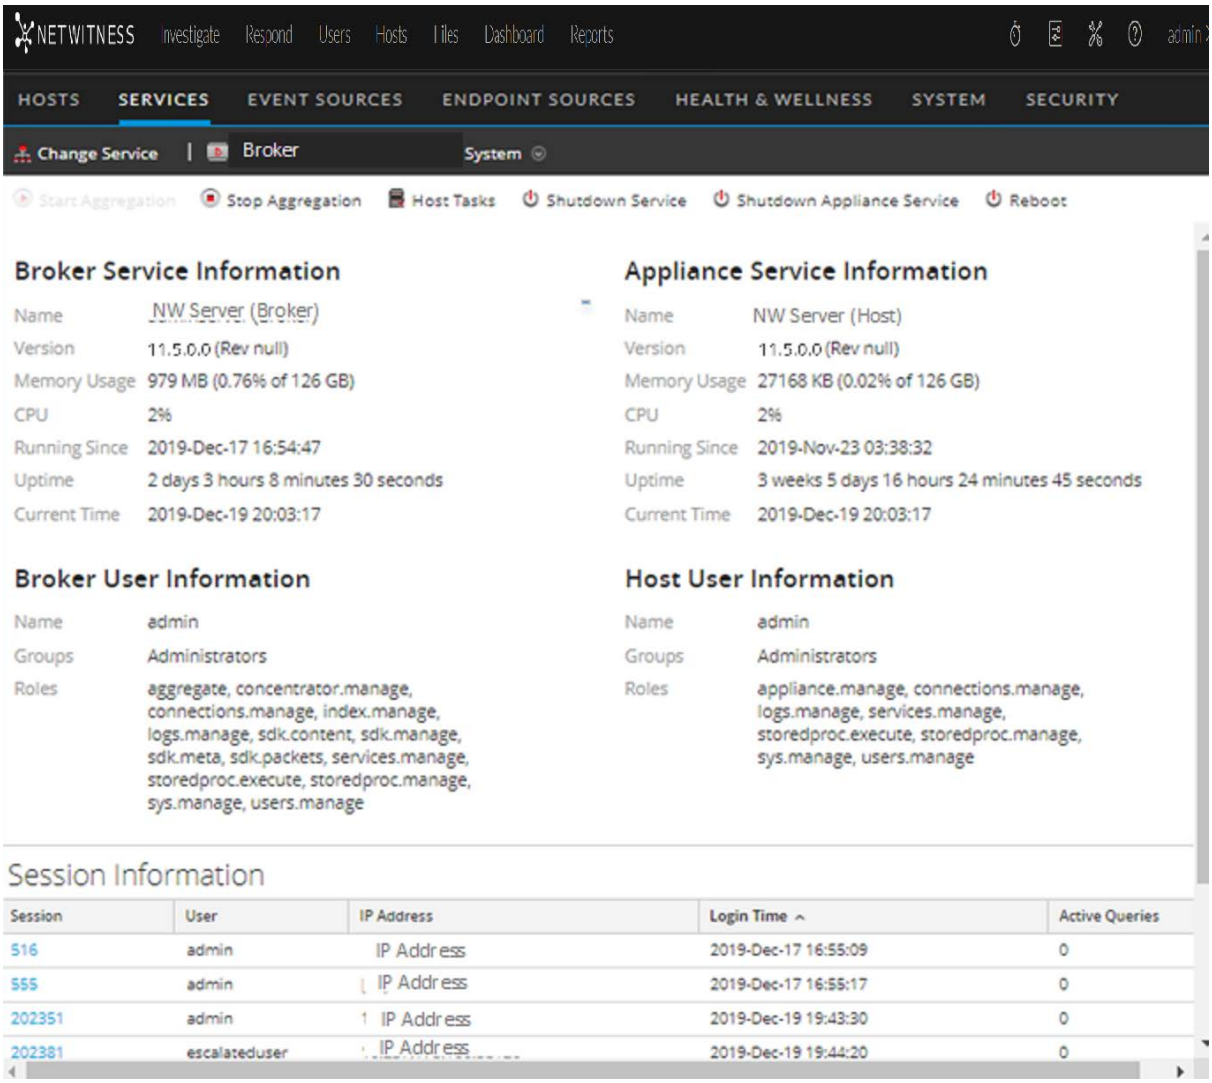


Execute a Task From the Host Task List

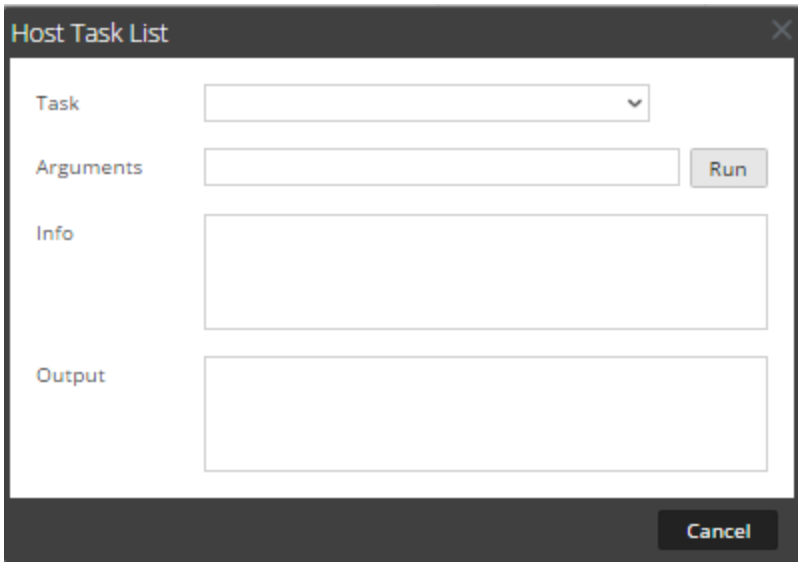
1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and click **> View > System**.

   Note: The Admin, Config, Orchestration, Security, Investigate, and Respond services do have access to the System view. They only have access to the Explore view.

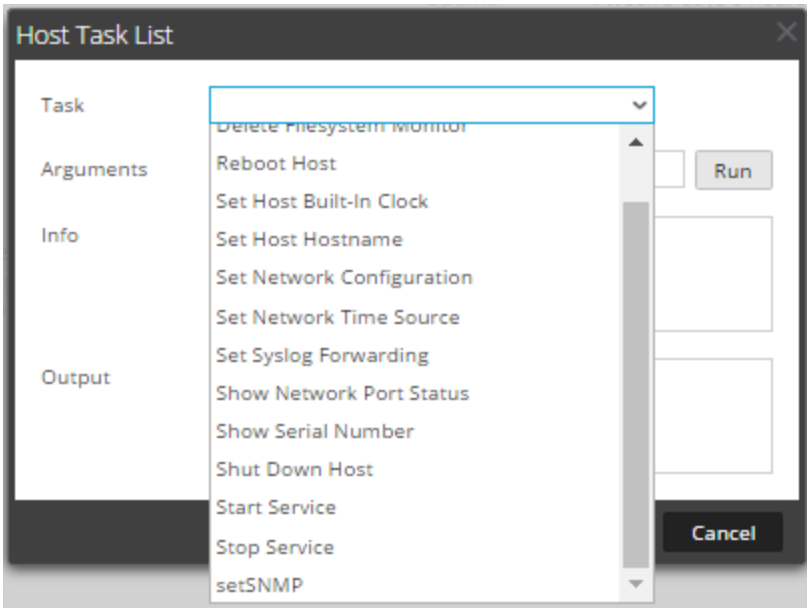
   The System view for the service is displayed below.



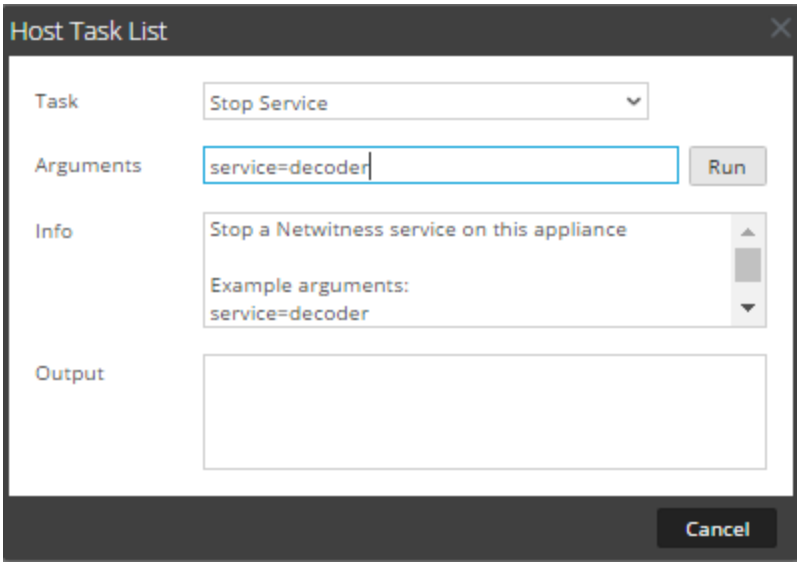
3. In the **Services System** view toolbar, click Host Tasks.



4. In the **Host Task List** dialog, click in the **Task** field to display a drop-down list of tasks that run on a host.



5. Select a task (for example, click **Stop Service**).
The task is displayed in the Task field. Task description, example arguments, security roles, and parameters are displayed in the Info area.



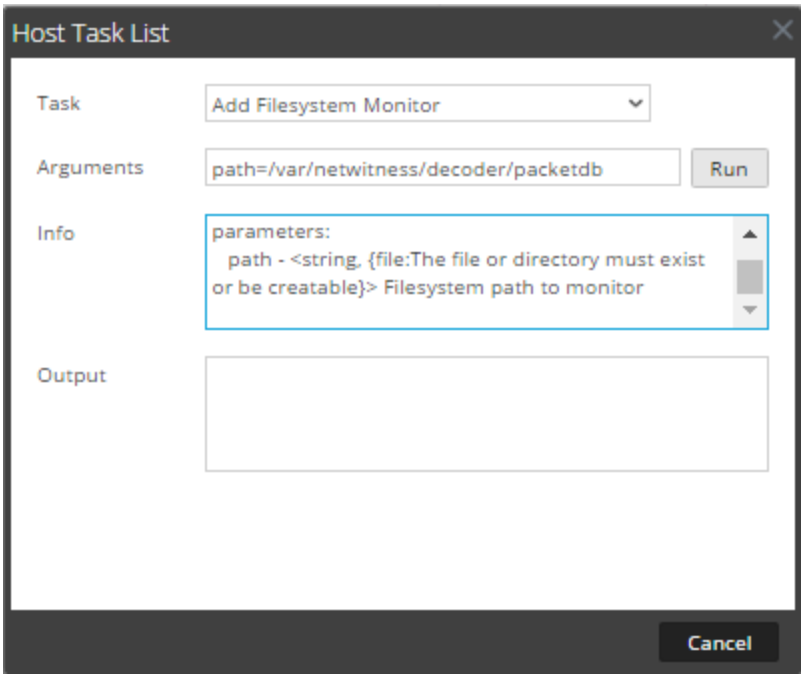
6. Type arguments if necessary and click **Run**.
The command executes and the result is displayed in the Output section.

Add and Delete a Filesystem Monitor

When you want a service to monitor traffic on a specific file system, you can select the service and then specify the path. NetWitness Platform adds a filesystem monitor. Once a file system monitor is added to a service, the service continues to monitor traffic on that path until the file system monitor is deleted.

Configure the Filesystem Monitor

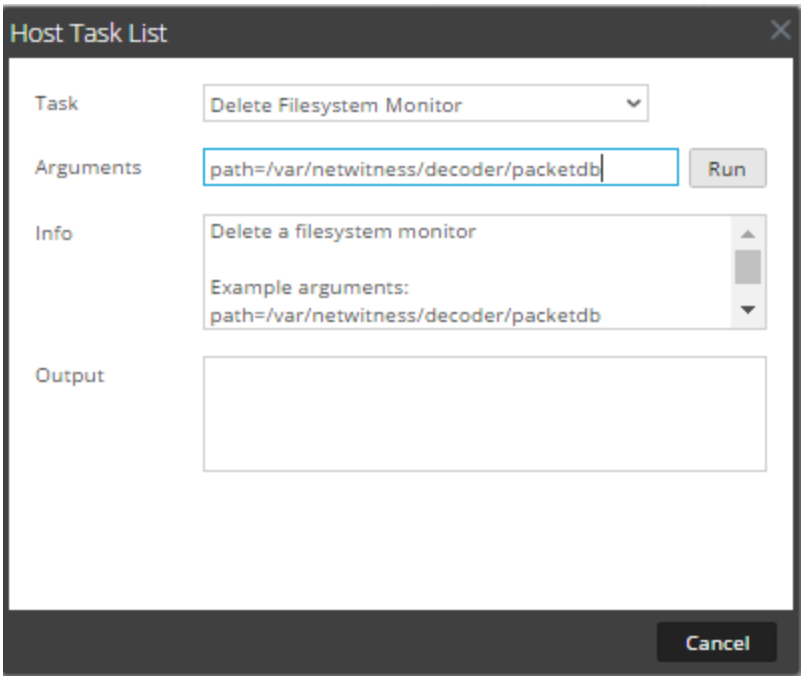
1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and click **> View > System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Add Filesystem Monitor**.
   In the Info area, a brief explanation of the task and the task arguments is displayed.
5. To identify the file system to monitor, type the path in the **Arguments** field. For example:
   `path=/var/netwitness/decoder/packetdb`



6. Click **Run**.
   The result is displayed in the Output area. The service begins to monitor the file system and continues to monitor it until you delete the filesystem monitor.

Delete a Filesystem Monitor

1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and click **> View > System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Delete Filesystem Monitor**.
   In the Info area, a brief explanation of the task and the task arguments is displayed.
5. To identify the filesystem to stop monitoring, type the path in the **Arguments** field. For example:
   `path=/var/netwitness/decoder/packetdb`



6. Click **Run**.
   The result is displayed in the Output area. The service stops monitoring the file system.

Reboot a Host

Under certain conditions, you must reboot a host; for example, after installing a software upgrade. This procedure uses a Host Task List message to shut down and restart a host.

NetWitness Platform also offers other options for shutting down a host:

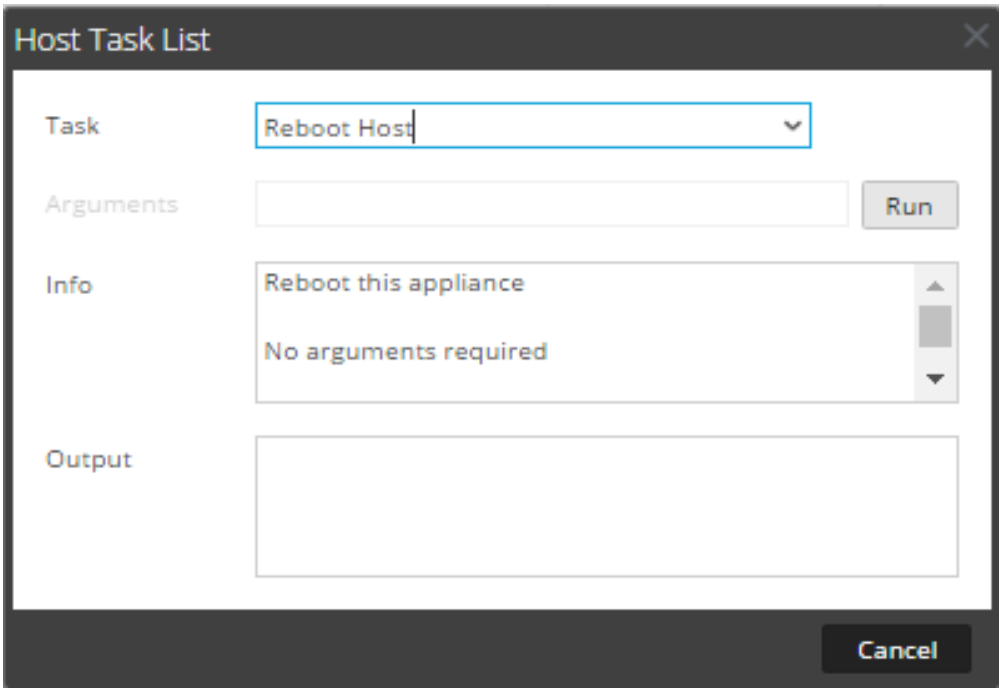
- To shut down and restart a host through an attached service, go to the Hosts view from a service in the Services view (see Search for Hosts) and then follow the Shut Down and Restart a Host from the Hosts View procedure below.
- To shut down the physical host without restarting, see Shut Down Host.

Shut Down and Restart a Host from the Hosts View

1. Select **(Admin) > Hosts**.
2. In the **Hosts** panel, select a host.
3. Select Reboot Host from the toolbar.

Shut Down and Restart a Host from the Host Task List

1. Select **(Admin) > Services**.
2. In the **Services** panel, select a service and click **> View > System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Reboot Host** in the **Task** field.
   No arguments are required.



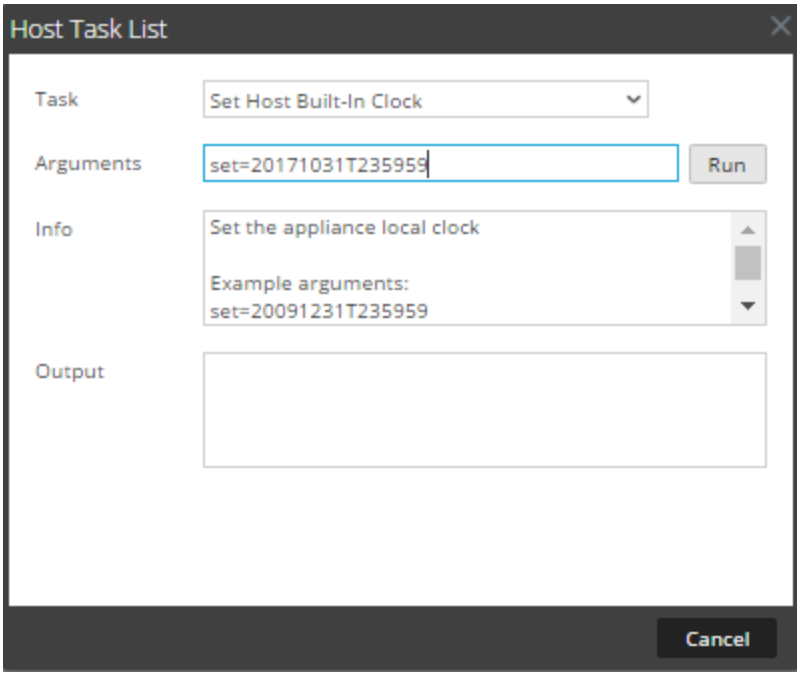
5. Click **Run**.
   The host is rebooted and the result is displayed in the Output area.

Set Host Built-In Clock

After a shutdown or battery failure, it may be necessary to set the local clock on a host. The Set Host Built-In Clock task resets the clock time.

Set the Time on the Local Clock

1. Select  **(Admin) > Services**.
2. In the **Services** list, select a service and  **> View > System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Set Host Built-In Clock**.
   Help for the task is displayed in the Info area.
5. Enter the date and time arguments in the **Arguments** field.
   For example, to specify October 31, 2017 at 11:59:59 PM, type:
   `set=20171031T235959`



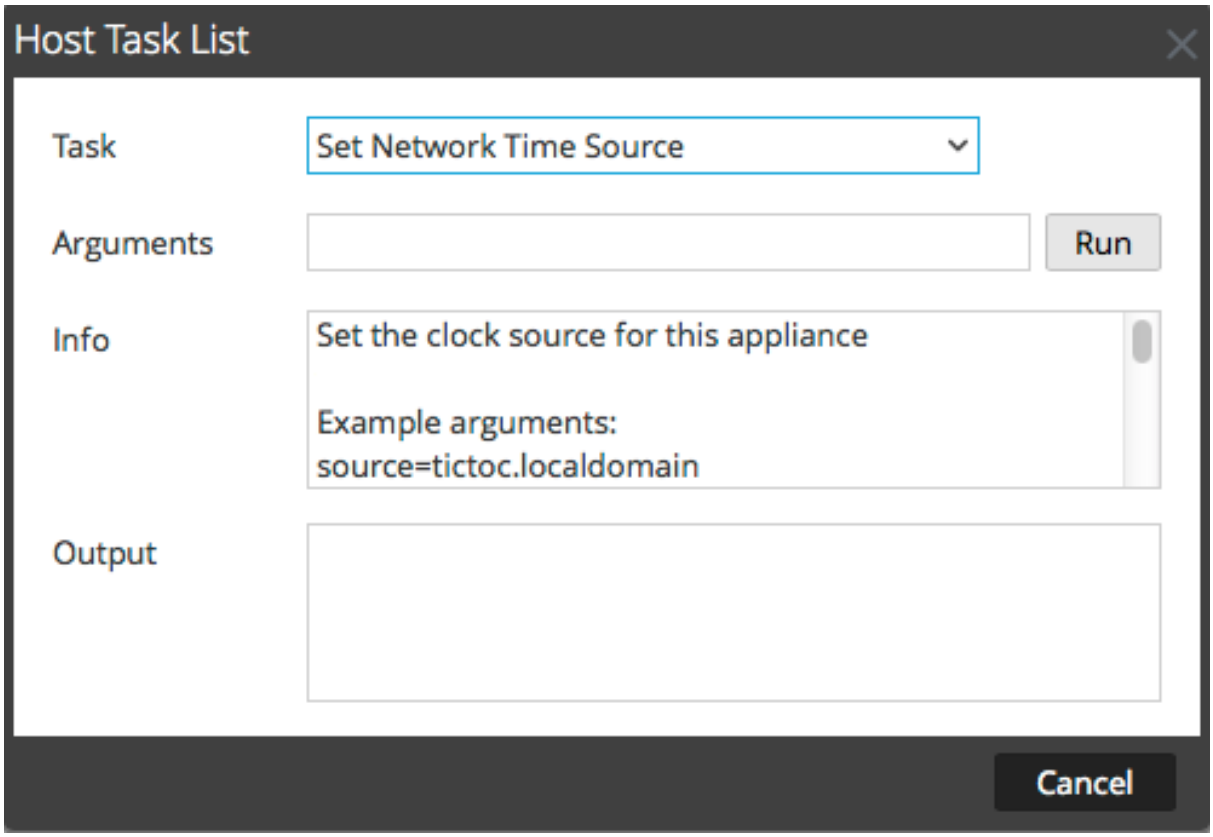
6. Click **Run**.
   The clock is set to the specified time and a message is displayed in the Output area.

Set Network Time Source

When setting the clock source for a host, set the hostname or address of an Network Time Protocol (NTP) server to be the network clock source for the host. If the host is using a local clock source, you must specify **local** here to allow **Set the Local Clock Source** to be effective.

Specify the Network Clock Source

1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and click **> View > System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Set Network Time Source**.



5. Do one of the following:
   - Type the hostname or address of the NTP server to serve as the clock source for this host; for example: **`source=tictoc.localdomain`**
   - If you want to use the host clock as a clock source, type: **`source=local`**
6. Click **Run**.
   The clock source is set and a message is displayed in the Output area.

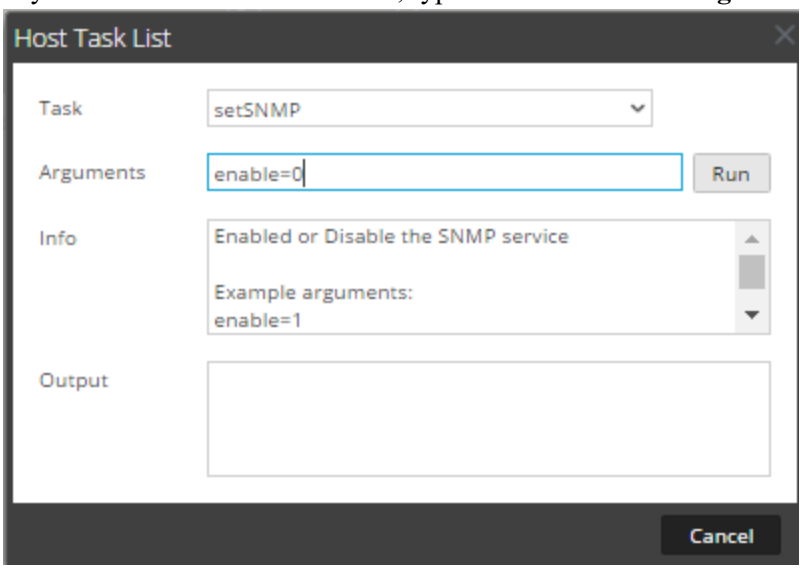
Note: If you specified a NTP clock source of **local**, the host clock serves as the clock source and the time is configured using Set Host Built-In Clock.

Set SNMP

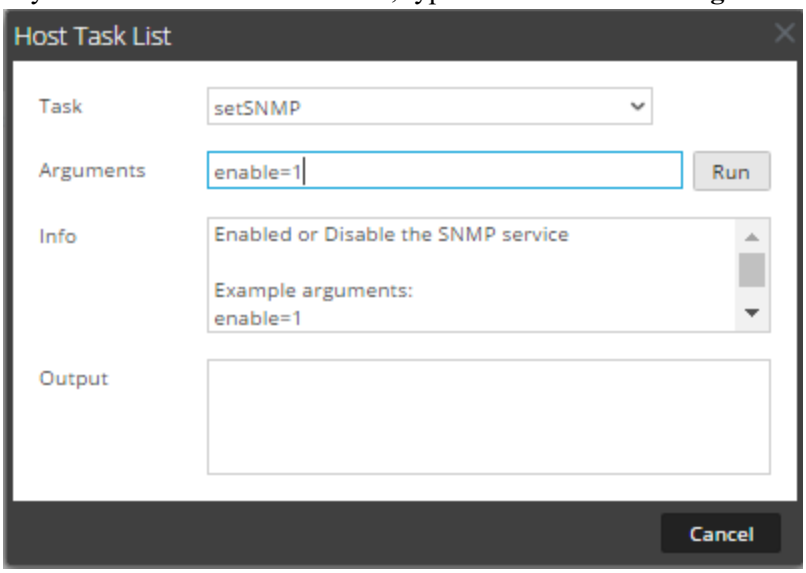
The Set SNMP task in the Host Task List enables or disables the SNMP service on a host. For a host to receive SNMP notifications, enable the SNMP service. If you are not using SNMP for NetWitness notifications, it is not necessary to enable the service.

Toggle SNMP Service on the Host

1. Select  **(Admin) > Services**.
2. In the **Services** list, select a service and click  **> View > System**.
   The System view for the service is displayed.
3. In the **Services System view** toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **setSNMP**.
   In the Info area, a brief explanation of the task and the task arguments is displayed.
5. Do one of the following:
   - If you want to disable the service, type **`enable=0`** in the **Arguments** field.



- If you want to enable the service, type **`enable=1`** in the **Arguments** field.



6. Click **Run**.
   The result is displayed in the Output area.

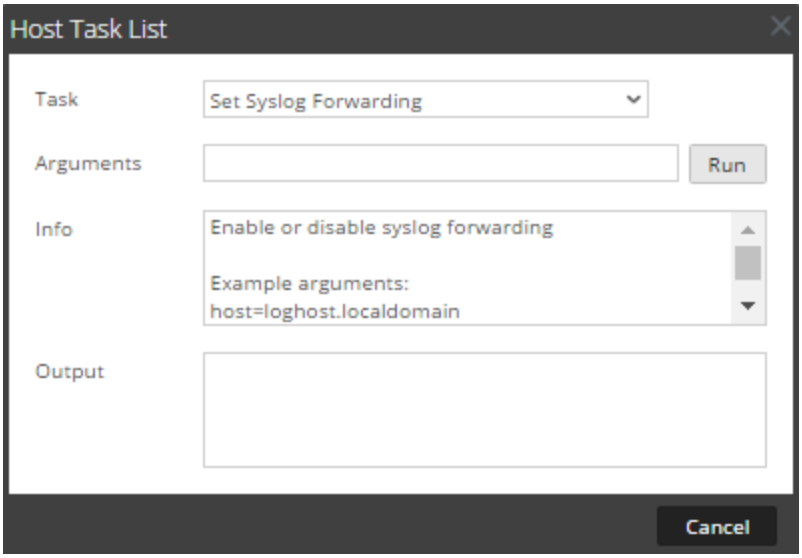
Set Syslog Forwarding

You can configure Syslog forwarding to forward the operating system logs of your NetWitness Hosts to a remote syslog server. You can use the Set Syslog Forwarding task in the Host Task List to enable or disable syslog forwarding.

Set Up and Start Syslog Forwarding

1. Select  **(Admin) > Services**.
2. In the **Services** list, select a service and click  **> View > System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.

4. In the **Host Task List**, select **Set Syslog Forwarding**.
   In the Info area, a brief explanation of the task and the task arguments is displayed.



5. In the **Arguments** field, do any one of the following.

   - To enable syslog forwarding, specify any one of the following formats:
     `host=<loghost>.<localdomain>` (for example, `host=syslogserver.local`).

     `host=<loghost>.<localdomain>:<port>` (for example, `host=syslogserver.local:514`).

     `host=<IP>` (for example, `host=10.31.244.244`).

     `host=<IP>:<port>` (for example, `host=10.31.244.244:514`).

     The following table lists the parameters used to enable syslog forwarding.

Parameter	Description
`loghost`	The host name of the remote syslog server.
`localdomain`	The domain of the remote syslog server.
`port`	IP address of the remote syslog server.
`IP`	The port number on which the remote syslog server receives a syslog messages.

   - To disable syslog forwarding, type **`host=disable`**.

6. Click **Run**.
   The result is displayed in the Output area.

Once syslog forwarding is enabled or disabled, the `/etc/rsyslog.conf` file is updated automatically to enable or disable syslog forwarding to the remote syslog destination and the syslog service is restarted.

If you enable syslog forwarding, the logs from the configured service are forwarded to the defined syslog server and continues forwarding until disabled.

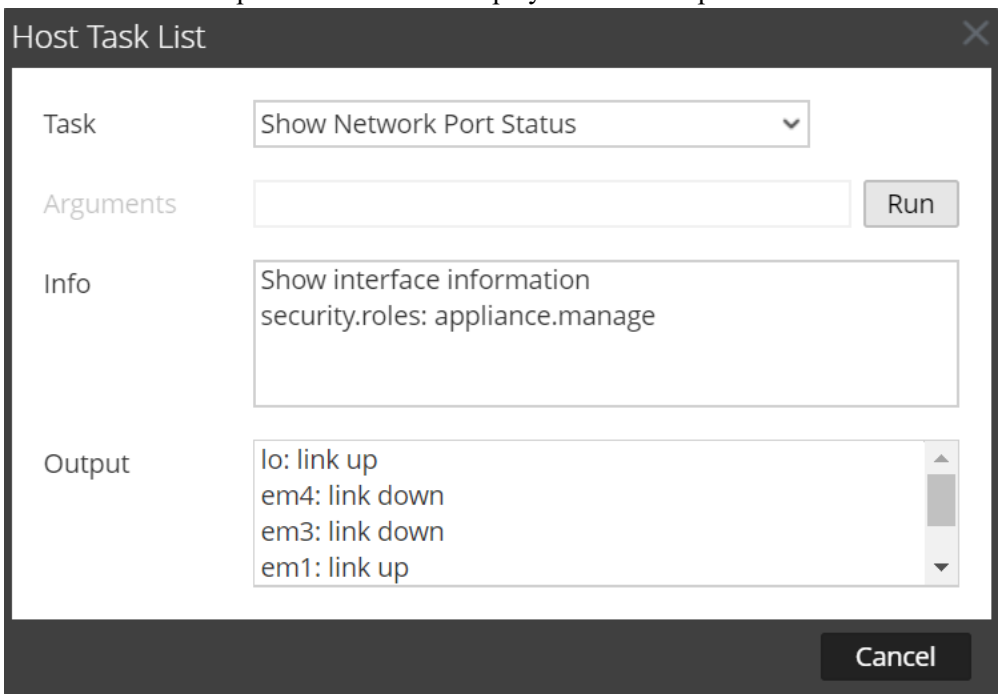
Note: You can now log in to the remote syslog server and verify if the messages are being received from the NetWitness services configured for syslog forwarding.

Show Network Port Status

The Show Network Port Status task in the Host Task List gives you the status of all configured ports on the host.

Display the Network Port Status

1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and **> View> System**.
   The System view for the selected service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, click **Show Network Port Status**.
   The task is displayed in the Task field, and information about the task is displayed in the Info area.
5. No arguments are required for this task. Click **Run**.
   The status for each port on the host is displayed in the Output area.

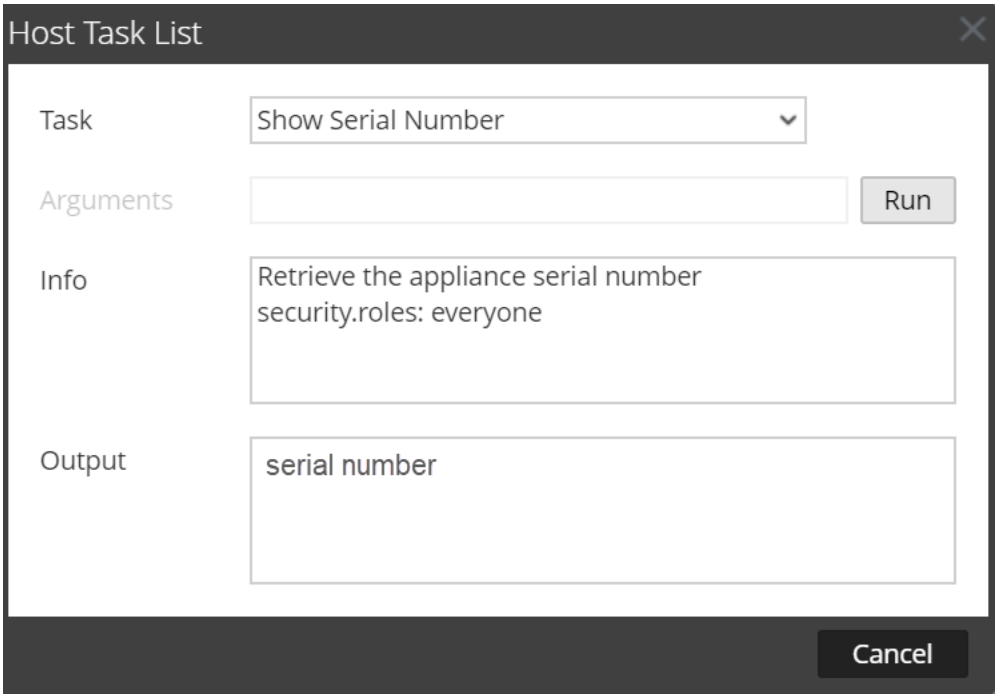


Show Serial Number

The Show Serial Number task in the Host Task List displays the serial number of a host.

Show the Serial Number

1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and click **> View> System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Show Serial Number**.
   In the Info area, a brief explanation of the task and the task arguments is displayed.
5. No arguments are required for this task. Click **Run**.
   The serial number of the selected host is displayed in the Output area.



Shut Down Host

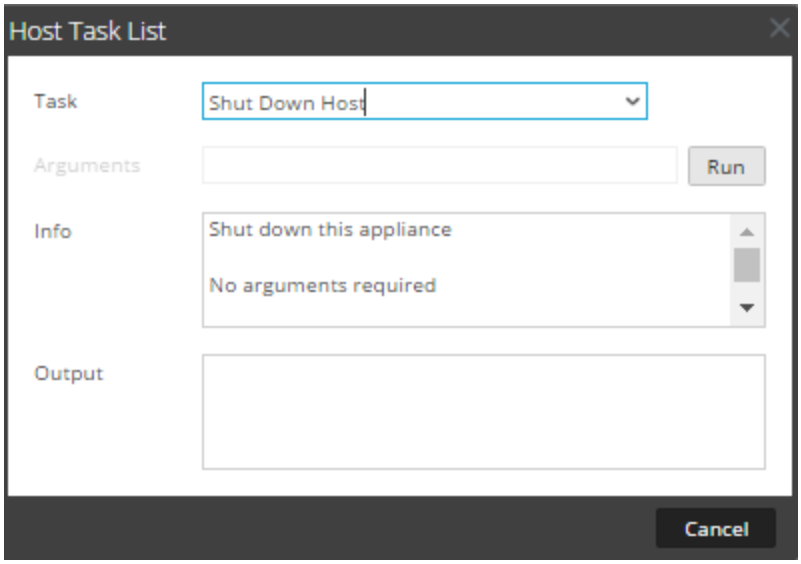
Under certain circumstances (for example, a hardware upgrade or an extended power outage that exceeds backup power capacity), it may be necessary to shut down a physical host. When you shut down a host, all services running on the host are stopped and the physical host turns off.

The physical host does not restart automatically. Use the power switch to restart the host. Once the physical host restarts, the host and services are configured to restart automatically.

See Reboot a Host for how to start and stop a host without shutting down the host.

Shut Down the Host

1. In the **Host Task List**, select **Shut Down Host**.



2. To execute the task, click **Run**.
   The host shuts down, and the host turns off.

Stop and Start a Service on a Host

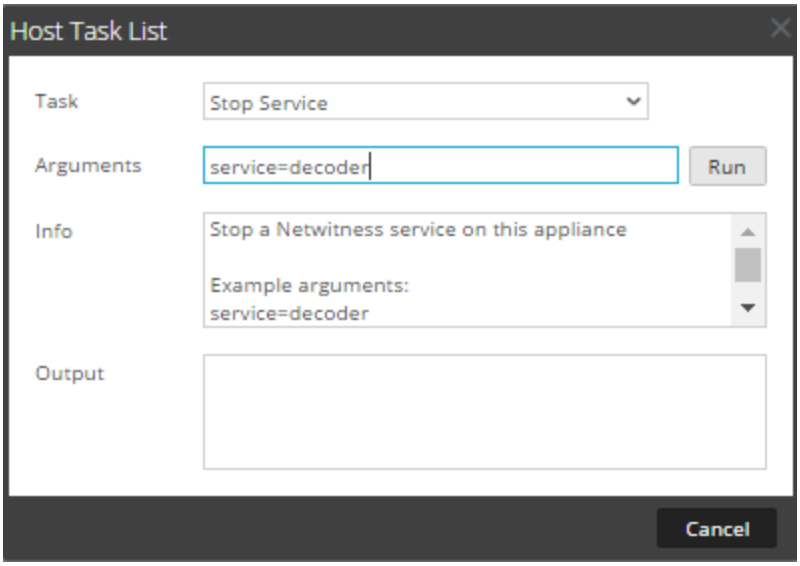
The Host Task List has two options for stopping and starting a service on a host. When you stop a service using the **Stop Service** message, all processes of the service are stopped and users connecting to the service are disconnected. Unless there is a problem with the service, it restarts automatically. This is the same as the **Shutdown Service** option in the Services System view.

If a service does not restart automatically after being stopped, you can restart it manually using the **Start Service** message.

Stop a Service on a Host

1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and click **> View> System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Stop Service**.
   The task is displayed in the Task field, and information about the task is displayed in the Info area.

5. Specify the service (`decoder`, `concentrator`, `broker`, `logdecoder`, `logcollector`) to stop in the **Arguments** field; for example, **`service=decoder`**.

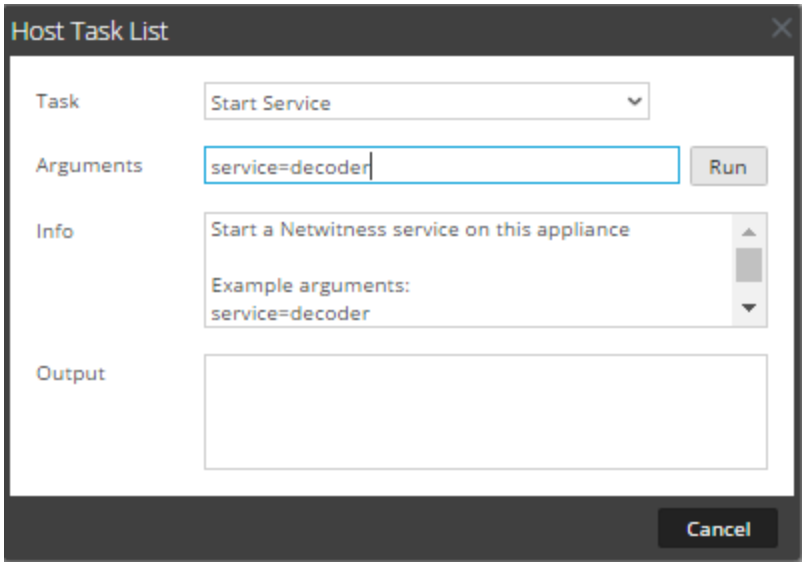


6. To execute the task, click **Run**.
   The service stops and the status is displayed in the Output area. All processes of the service are stopped and users connecting to the service are disconnected. Unless there is a problem with the service, it restarts automatically.

Start a Service on a Host

1. Select **(Admin) > Services**.
2. In the **Services** list, select a service and click **> View> System**.
   The System view for the service is displayed.
3. In the **Services System** view toolbar, click **Host Tasks**.
4. In the **Host Task List**, select **Start Service**.
   The task is displayed in the Task field, and information about the task is displayed in the Info area.

5. Specify the service (`decoder`, `concentrator`, `broker`, `logdecoder`, `logcollector`) to stop in the **Arguments** field; for example, **`service=decoder`**.



6. To execute the task, click **Run**.
   The service starts and the status is displayed on the Output area.

Add, Replicate, or Delete a Service User

You must add a user to a service for:

- Aggregation
- Accessing the service with the:
  - Thick client
  - REST API

Note: This topic does not apply to users who access services through the user interface on NetWitness Server. You must add those users to the system, not a service. For details, see the "Set Up a User" in *System Security and User Management Guide*.

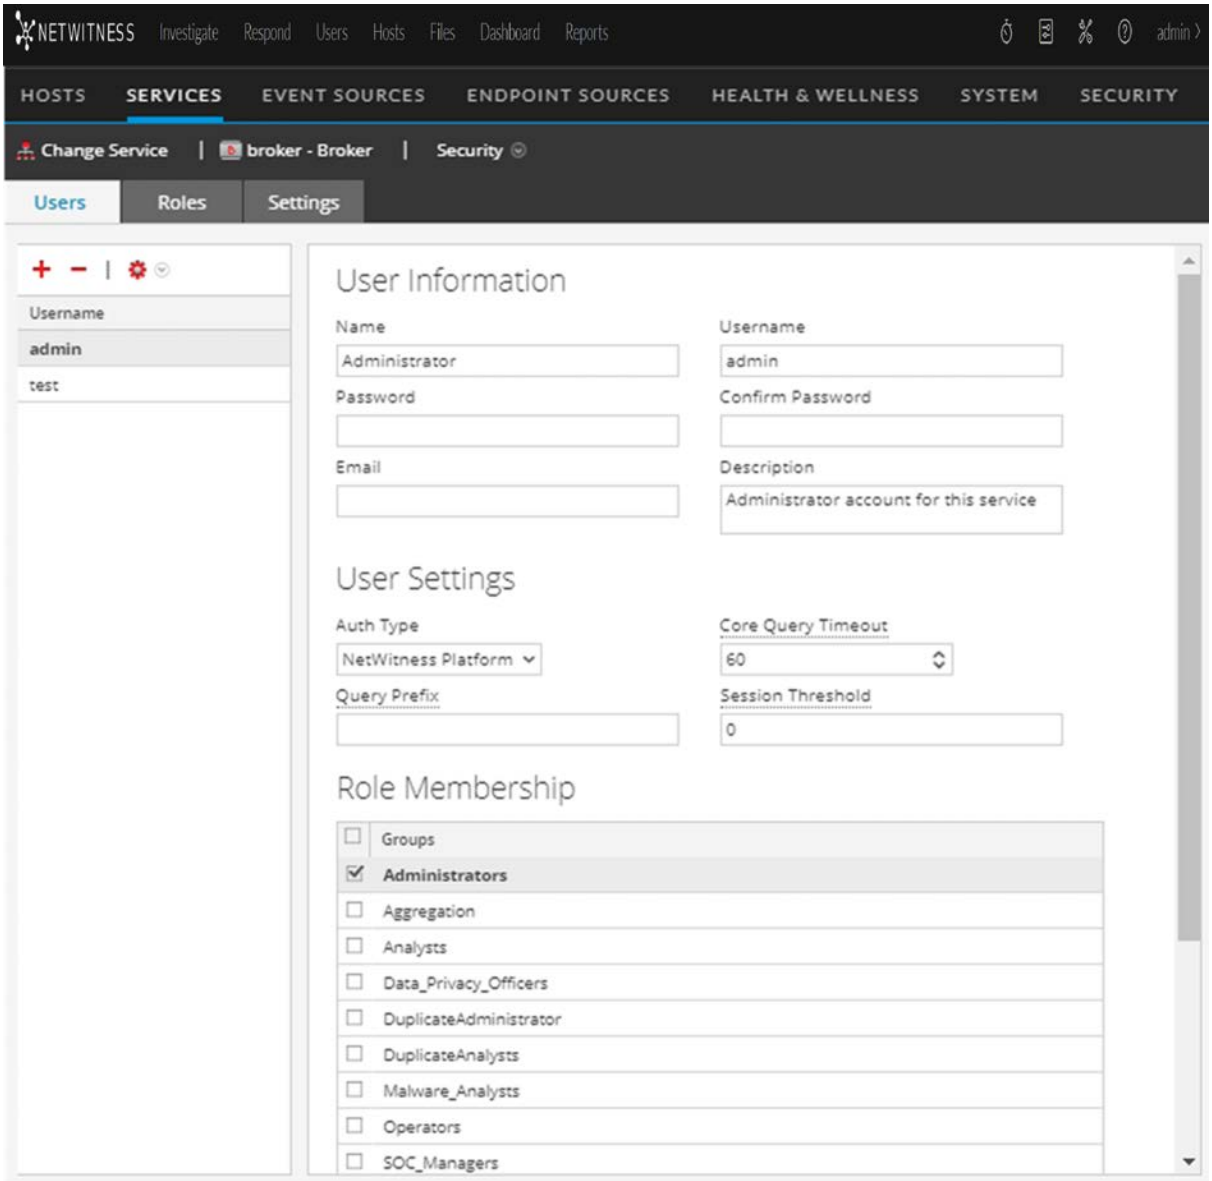
For each service user, you can:

- Configure user authentication and query handling properties for the service
- Make the user a member of a role, which has a set of permissions the user receives
- Replicate the user account to other services
- Change the service user password on selected services

Change a Service User Password provides instructions for changing the service user password across services.

To navigate to the Services Security view:

1. In NetWitness, go to **(Admin) > Services**.
2. Select a service, then click **> View > Security**.
   The Security view for the selected service is displayed with the Users tab open.



Add a Service User

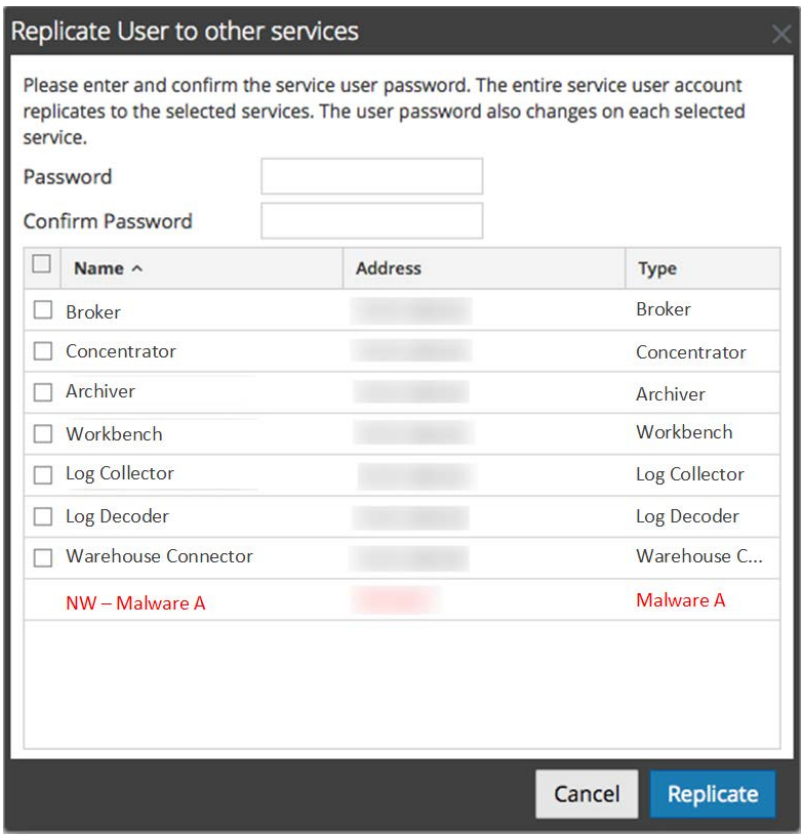
1. On the **Users** tab, click .
2. Type the user name to access the service, then press **Enter**.
   The User Information section displays the user name and the rest of the fields are available for editing.

3. Type the password for logging on to the service in the **Password** and **Confirm Password** fields.
4. (Optional) Provide additional information:
   - **Name** for logging on to NetWitness
   - **Email** address
   - **Description** of the user
5. In the User Settings section, select the following information:
   - **Authentication Type**
     - If NetWitness authenticates the user, select **NetWitness**.
     - If Active Directory or PAM is configured on NetWitness Server to authenticate the user, select **External**.
   - **Core Query Timeout** is the maximum number of minutes a user can run a query on the service. This field applies to NetWitness 10.5 and later service versions and does not appear for 10.4 and earlier versions.
6. (Optional) Specify additional query criteria:
   - **Query Prefix** filters queries. Type a prefix to restrict results the user sees.
   - **Session Threshold** controls how the service scans meta values to determine session counts. Any meta value with a session count that is above the threshold stops its determination of the true session count.
7. In the **Role Membership** section, select each role to assign to the user. When a user is a member of a role on a service, the user has the permissions assigned to the role.
8. To activate the new service user, click **Apply**.

Replicate a User to Other Services

Note: The **admin** user cannot be replicated to other services.

1. In the Users tab, select a user and click  **> Replicate**.
   The Replicate Users to Other Services dialog is displayed.



2. Enter and confirm the password.
3. Select each service to which you are replicating the user.
4. Click **Replicate**.

Delete a Service User

1. On the **Users** tab, select the **Username** and click −.
   NetWitness requests confirmation that you want to delete the selected user.
2. To confirm, click **Yes**.

Add a User Role to a Service

There are pre-configured roles in NetWitness that are installed on the server and on each service. You can also add custom roles. The following table lists the pre-configured user roles and their permissions.

Role	Permission
Administrators	Full system access
Operators	Access to configurations but not to metadata and session content
Analysts	Access to metadata and session content but not to configurations

Role	Permission
SOC_Managers	Same access as Analysts and additional permissions to handle incidents
Malware_ Analysts	Access to malware events and to metadata and session content
Data_Privacy_ Officers	Access to metadata and session content and configuration options that manage obfuscation and viewing of sensitive data within the system (see *Data Privacy Management Guide*).

You must add a service role when you have added a:

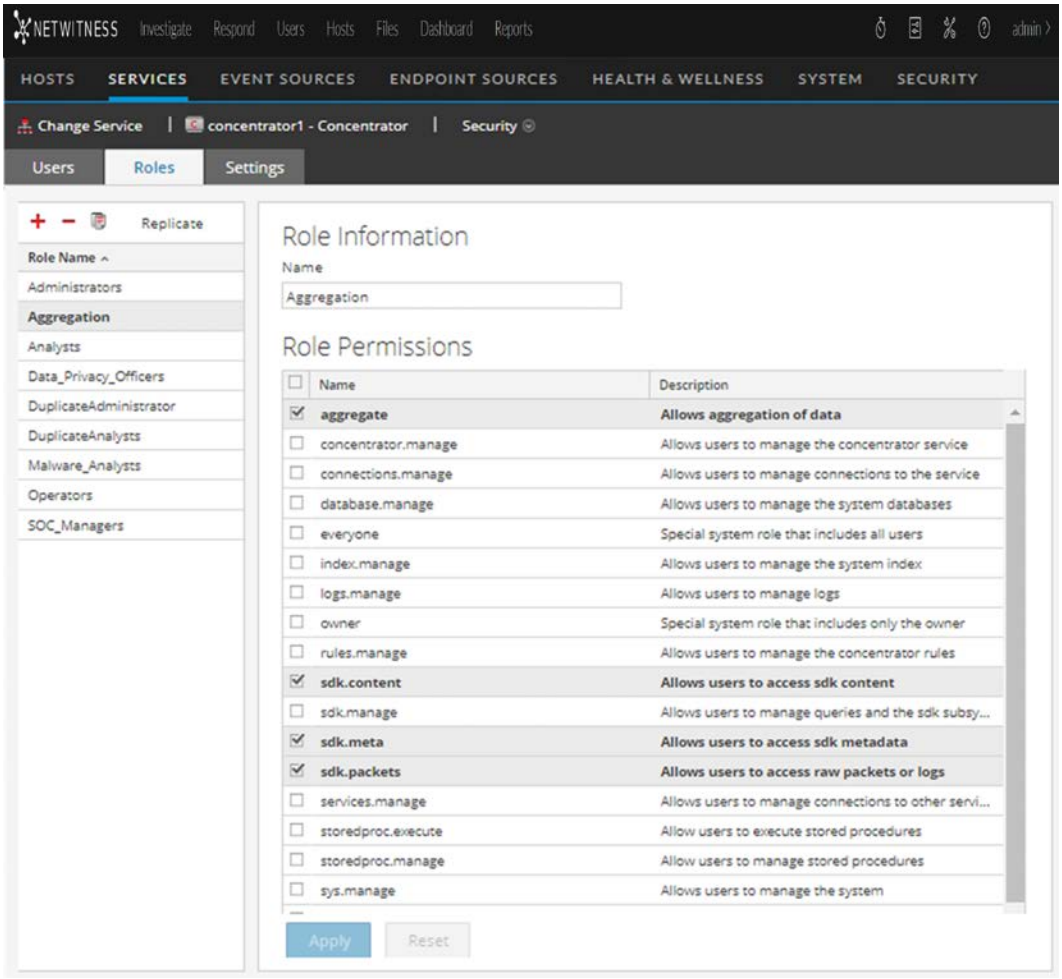
- **Service** user or users that requires a new set of permissions.
- **Custom role on NetWitness Server** because trusted connections require that the same custom role exists both on the server and on each service the custom role will access. The names must be identical. For example, if you add a Junior Analysts role on the server then you must add a Junior Analysts role on each service the role will access. For more information, see "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.

There is also a pre-configured **Aggregation** service role. Services Security View - Aggregation Role and Services Security View - Service User Roles and Permissions provide additional information.

To add a service user role and assign permissions to it:

1. In NetWitness, go to  **(Admin) > Services**.
2. Select a service, then  **> View > Security**.
   The Security view for the selected service is displayed with the Users tab open.

3. Select the **Roles** tab and click +.
   The Services Security view is displayed and five pre-configured roles are already listed.



4. Click +, type the **Role Name** and press **Enter**.
   The Role Name is displayed above a list of **Role Permissions**.
5. Select each permission the role will have on the service.
6. Click **Apply**.

You can add service users to the role in the **Users** tab.

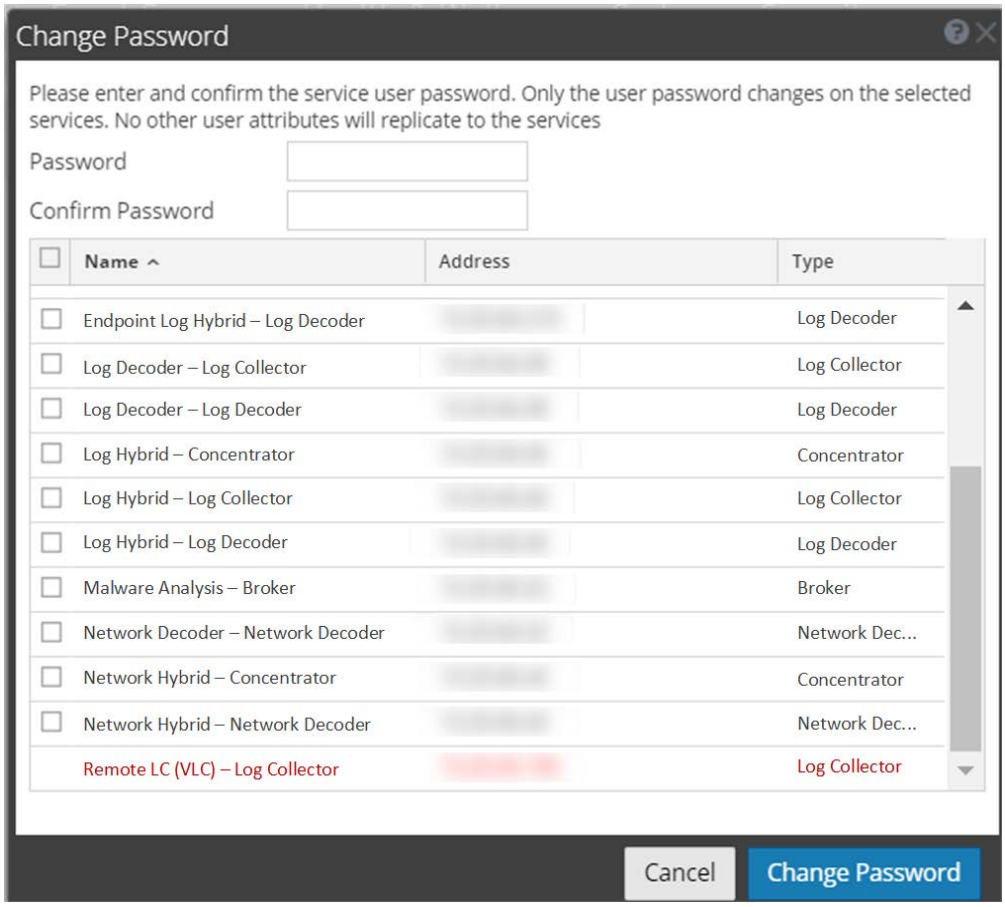
Change a Service User Password

This procedure allows administrators to change the password of a service user and replicate the new password to all Core services with that user account defined. It replicates only the password change to the Core services selected and does not replicate the entire user account. Administrators can also change the password of the **admin** account on the Core services.

Note: The Change Password option does not apply to external users.

To change the password of a service user:

1. In NetWitness, go to **(Admin) > Services**.
   The Admin Services view is displayed.
2. Select a service, then click **> View > Security**.
   The Security view for the selected services is displayed.
3. In the **Users** tab, select a user and select **Change Password** from .
   The **Change Password** dialog is displayed.



4. Type a new password for the user and confirm the password.
5. Select the services where you want the user password to change.
6. Click **Change Password**.
   The status of the password change on the selected services is displayed.

IMPORTANT: If you change the admin password on a NetWitness service that is used as a Reporting Engine data source, you must remove and then re-add the service as a data source. For details, see "Configure the Data Sources" topic in the *Reporting Engine Configuration Guide for RSA NetWitness Platform 11.x Guide*.

Create and Manage Service Groups

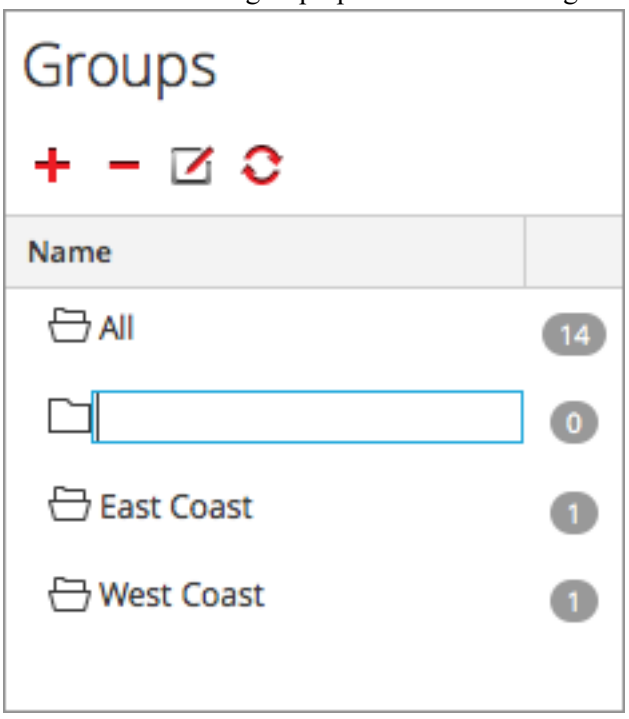
The Admin Services view provides options to create and manage groups of services. The Services list toolbar includes options to create, edit, and delete service groups. Once groups are created, you can drag individual services from the Services panel into a group.

Groups may reflect functional, geographical, project-oriented, or any other organization principle that is useful. A service may belong to more than one group. Here are some examples of possible groupings.

- Group different service types to make it easier to configure and monitor all Brokers, Network Decoders, or Concentrators.
- Group services that are part of the same data flow; for example, a Broker, and all associated Concentrators and Network Decoders.
- Group services according to their geographic region and location within the region. If a major power outage occurs in a location, potentially affected services are easily identifiable.

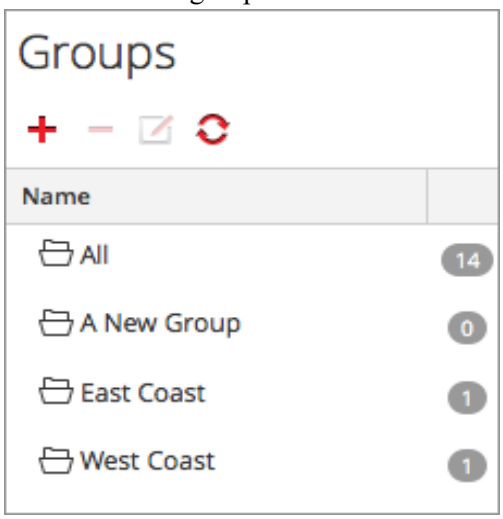
Create a Group

1. In NetWitness, go to **(Admin) > Services**.
   The Admin Services view is displayed.
2. In the **Groups** panel toolbar, click .
   A field for the new group opens with a blinking cursor.



3. Type the name of the new group in the field (for example, **`A New Group`**) and press **Enter**.
   The group is created as a folder in the tree. The number next to the group indicates the number of

services in that group.



Change the Name of a Group

1. In the Services view **Groups** panel, double-click the group name or select the group and click . The name field opens with a blinking cursor.
2. Type the new name of the group and press **Enter**.
   The name field closes and the new group name is displayed in the tree.

Add a Service to a Group

In the Services view **Services** panel, select a service and drag the service to a group folder in the groups panel.
The service is added to the group.

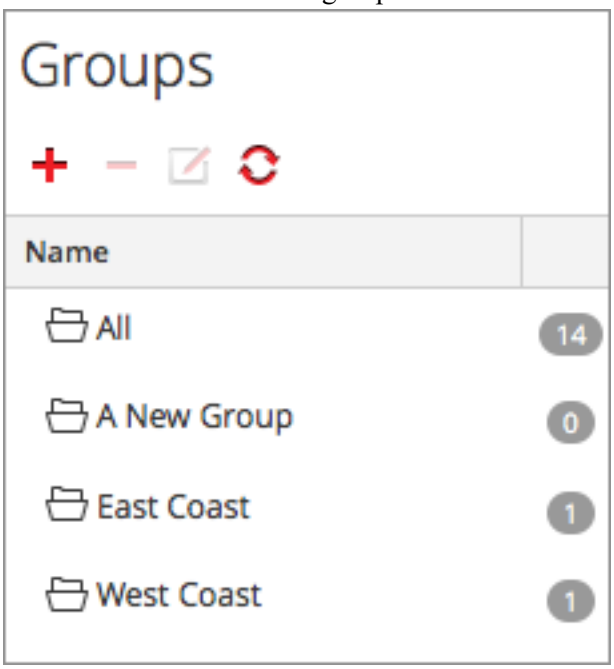
View the Services in a Group

To view the services in a group, click the group in the **Groups** panel.

The Services panel lists the services in that group.

Remove a Service from a Group

1. In the Services view **Groups** panel, select the group that contains the service that you want to remove. The services in that group appear in the Services panel.
2. In the **Services** panel, select one or more services that you want to remove from the group, and in the toolbar, select − > **Remove from Group**.
   The selected services are removed from the group, but are not removed from the NetWitness user interface. The number of services in the group, which is listed near the group name, decreases by the number of services removed from the group. The **All** group contains the services that are removed from the group.
   In the following example, the service group called **A New Group** does not contain any services, because the service in that group is removed.



Delete a Group

1. In the Services view **Groups** panel, select the group that you want to delete.
2. Click −.
   The selected group is removed from the Groups panel. The services that were in the group are not removed from the NetWitness user interface. The **All** group contains the services from the deleted group.

Duplicate or Replicate a Service Role

An efficient way to add a new service role is to duplicate a similar role, save it with a new name and revise the permissions that are already assigned. For example, you could duplicate the Analysts role. Then, save it as `JuniorAnalysts` and modify the permissions.

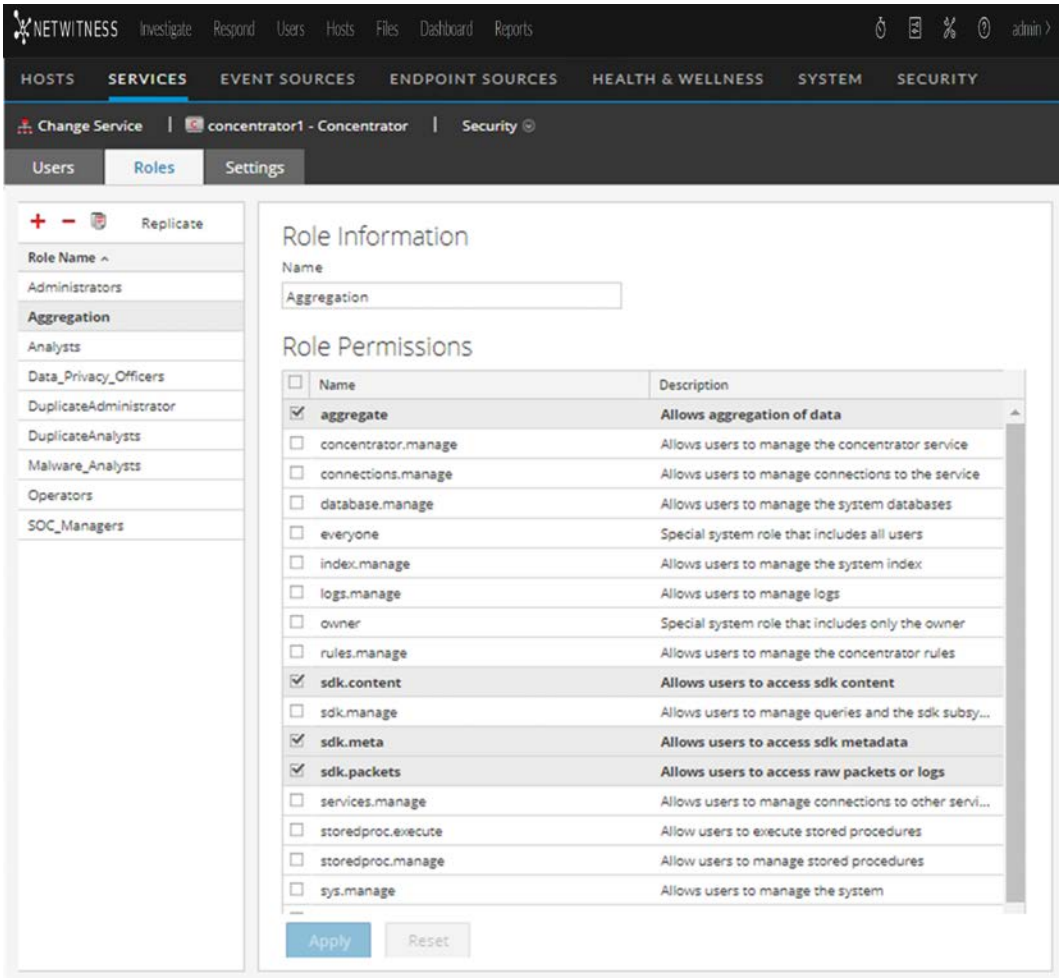
The quick way to add an existing role to other services is to replicate the role. For example, you could replicate the `JuniorAnalysts` role that exists on a Broker to a Concentrator and Log Decoder.

To navigate to the Services Security view:

1. In NetWitness, go to **(Admin) > Services**.
2. Select a service, then click **> View > Security**.
   The Security view for the selected service is displayed with the Users tab open.
3. Select the **Roles** tab.

Duplicate a Service Role

1. In the **Roles** tab, select the role you want to duplicate.



2. Click **> Duplicate Role.**
3. Type a new name and click **Apply**.

4. Select the new role.
5. In the **Role Permissions** section, select or deselect permissions to modify what the new role can do.

Replicate a Role

1. In the **Roles** tab, select the role you want to replicate and click **Replicate**.
2. In the **Replicate Role to Other Services** dialog, select each service on which to add the role.
3. Click **Replicate**.

Edit Core Service Configuration Files

The service configuration files for Network Decoder, Log Decoder, Broker, Concentrator, Archiver, and Workbench services are editable as text files. In the Services Config view > Files tab, you can:

- View and edit a service configuration file that the NetWitness system is currently using.
- Retrieve and restore the latest backup of the file you are editing.
- Push the open file to other services.
- Save changes made to a file.

The files available to edit vary depending upon the type of service being configured. The files that are common to all Core services are the:

- The NetWitness file (`netwitness`). This is preconfigured and does not require editing.
- The service index file (`index-<service>`). This is preconfigured and may require editing.
  See Edit a Service Index File for more information.
- The scheduler file (`scheduler`). The scheduler service is optional and requires editing.
  See Configure the Task Scheduler for more information.
- The crash reporter file (`crashreporter`). The crash reporter service is optional and requires editing.
  See Enable the Crash Reporter Service for more information.
- The feed definitions file (`feed-definitions`). This file is optional and may require editing.
  See "Feed Definitions File" in the *Decoder Configuration Guide* for more information.

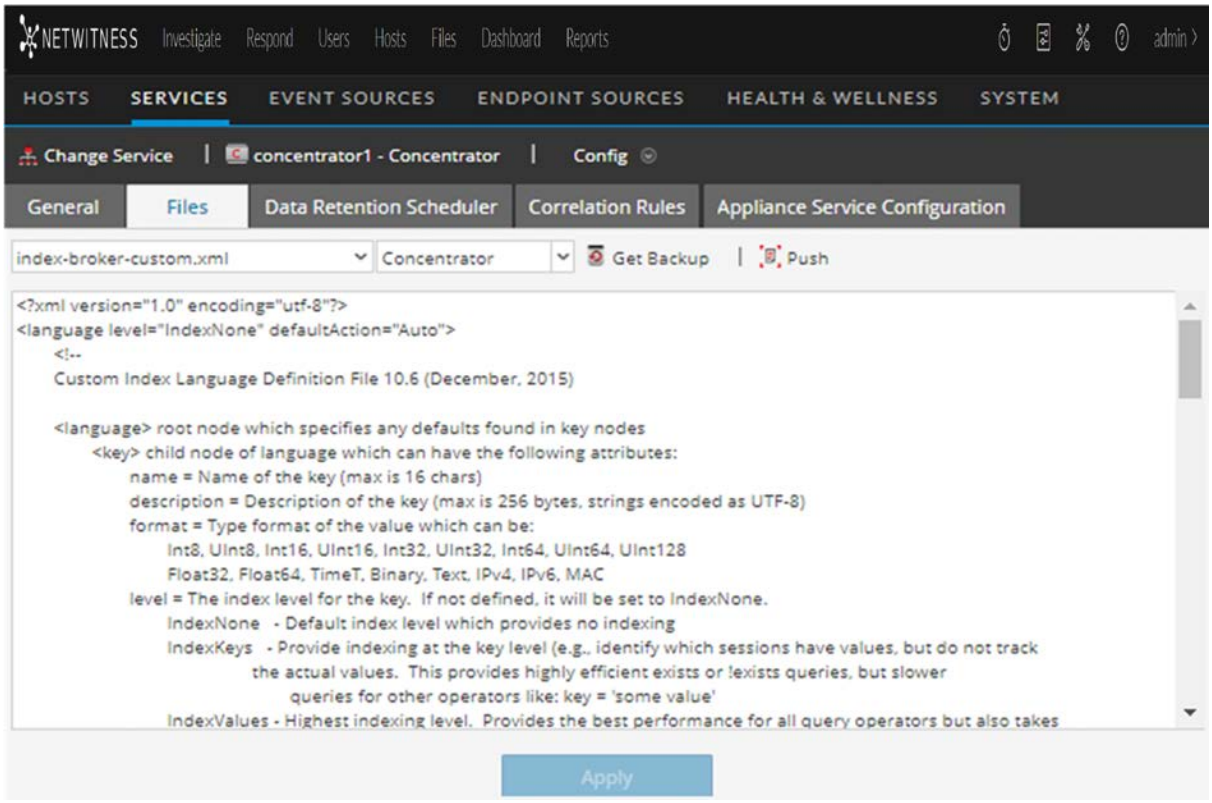
In addition, the Network Decoder has files that configure parsers, feed definitions, and a wireless LAN adapter. There is also the table mapping file provided by NetWitness, `table-map.xml`, which is an important part of the Log Decoder.

Note: The default values in these configuration files are good for the most common situations, however some editing is necessary for optional services, such as the crash reporter or scheduler. Only administrators with a good understanding of the networks and the factors that affect the way services collect and parse data should make changes to these files in the Files tab.

Edit a Service Configuration File

To edit a file:

1. In NetWitness, go to **(Admin) > Services**.
2. In the Services list, select a service.
3. Select **> View > Config**.
   The Service Config view is displayed with the General tab open.
4. Click the **Files** tab.
   cThe selected service, such as Concentrator, appears in the drop-down list on the right.
5. (Optional) To edit a file for the host instead of the service, select **Host** in the drop-down list.
6. Choose a file from the **Please Select A File To Edit** drop-down list.
   The file content is displayed in edit mode.



7. Edit the file and click **Apply**.

The current file is overwritten and a backup file is created. The changes go into effect after the service is restarted.

Revert to a Backup Version of a Service Configuration File

After you make changes to a configuration file, save the file, and restart the service, a backup file is available.

To revert to a backup of a configuration file:

1. Select a configuration file by completing steps 1-6 of Edit a Service Configuration File.
2. Click 

.
The backup file opens in the text editor.
3. To revert to the backup version, click **Save**.

The changes go into effect after the service is restarted.

Push a Configuration File to Other Services

Once you have edited a service configuration file, you can push the same configuration to other services of the same type.

1. Select a configuration file by completing steps 1-6 of Edit a Service Configuration File.
2. Click 

.
The Select Services dialog is displayed.
3. Select each service to push the configuration file on it. Each service must be the same type as the one you selected in the Services view.

   Caution: If you decide not to push the configuration file, click **Cancel**.
4. To push the configuration file to all selected services, click **OK**.

The configuration file is pushed to all selected services.

Edit a Service Index File

This topic provides important information and guidelines for configuring service custom index files, which are editable in the Service Config view > Files tab.

The index file, along with other configuration files, controls operation of each core service. Accessing the index file through the Service Config view in NetWitness opens the file in a text editor, where you can edit the file.

Note: Only administrators with a thorough and comprehensive understanding of Core service configuration are qualified to make changes to an index file, which is one of the central configuration files for the appliance service. Changes made should be consistent across all Core services. Invalid entries or a misconfigured file can prevent the system from starting and can require the assistance of NetWitness Support to bring the system back into a working state.

These are the index files:

- `index-broker.xml`, and `index-brokereustom.xml`
- `index-concentrator.xml`, and `index-concentrator-custom.xml`
- `index-decoder.xml`, and `index-decodereustom.xml`
- `index-logdecoder.xml`, and `index-logdecodereustom.xml`

- index-archiver.xml, and index-archiver-custom.xml
- index-workbench.xml, and index-workbench-custom.xml

Index and Custom Index Files

All customer-specific index changes are made in index-<service>-custom.xml. This file overrides any settings in index-<service>.xml, which is solely controlled by NetWitness.

The custom index file, index-<service>-custom.xml, allows creation of custom definitions or overrides of your own language keys that are not overwritten during the upgrade process.

- Keys that are defined in index-<service>-custom.xml replace the definitions found in index-<service>.xml.
- Keys that are added to index-<service>eustom.xml and not found in index-<service>.xml are added to the language as a new key.

Some common applications for editing the index file are:

- To add new custom meta keys to add new fields to the NetWitness user interface.
- To configure protected meta keys as part of a data privacy solution as described in the *Data Privacy Management Guide*.
- To adjust the NetWitness Core database query performance as described in the *NetWitness Core Database Tuning Guide*.

Caution: Never set the index level to IndexKeys or IndexValues on a Network Decoder if you have a Concentrator or Archiver aggregating from the Network Decoder. The index partition size is too small to support any indexing beyond the default time meta key.

Configure the Task Scheduler

Scheduler File

You can edit the scheduler file that in the Service Config view > Files tab. This file configures the built-in task scheduler for a service. The task scheduler can automatically send messages at predefined intervals or specific times of the day.

Scheduler Task Syntax

A task line in the scheduler file consists of the following syntax, where <Value> has no spaces:

```
<ParamName>=<Value>
```

If <Value> has any spaces, this is the syntax:

```
<ParamName>="<Value>"
```

In each task line, these guidelines apply:

- Parameter `time` or one of the interval parameters (`seconds`, `minutes` or `hours`) is required.
- Escape special characters with a `\` (backslash).

Task Line Parameters

The following task line parameters are accepted by the scheduler.

Syntax	Description
daysOfWeek: <string, optional, {enum-any:sun|mon|tue|wed|thu|fri|sat|all}>	The days of week to execute a task. The default value is all.
deleteOnFinish: <bool, optional>	Delete the task when it has successfully finished.
hours: <uint32, optional, {range:1 to 8760}>	The number of hours between executions.
logOutput: <string, optional>	Output the response to log using the specified module name.
minutes: <uint32, optional, {range:1 to 525948}>	The number of minutes between executions.
msg: <string>	The message to send the node.
params: <string, optional>	The parameters for the message.
pathname: <string>	The path of the node that receives the message.
seconds: <uint32, optional, {range:1 to 31556926}>	The number of seconds between executions.
time: <string>	The time of execution in HH::MM:SS format (local time of this server).
timesToRun: <uint32, optional>	How many times to run because service start, **0** = unlimited (default).

Messages

The following are the message strings to use in the Task Scheduler `msg` parameter.

Message	Description
addInter	Add a task to run at a defined interval. For example, this message runs the `/index save` command every 6 hours: `addInter hours=6 pathname=/index msg=save`

Message	Description
addMil	Add a task to run at a specific time of day or even day(s) of the week. For example, this message runs the `/index save` command at 1 AM every business day: `addMil time= 01:00:00 pathname=/index msg=save daysOfWeek=mon,tue,wed,thu,fri`
delSched	Deletes an existing scheduled task. The `id` parameter of the task must be retrieved from the print message.
print	Prints all scheduled tasks.
replace	Assign all scheduled tasks in one message, deleting any existing tasks.
save	Save node.

Sample Task Line

The following example task line in the `scheduler` file downloads the feeds package file (`feeds.zip`) to the selected Network Decoder every 120 minutes from the feeds host server:

```
minutes=120 pathname=/parsers msg=feed params="type\=wget
file\=http://feedshost/nwlive/feeds.zip"
```

Enable the Crash Reporter Service

The Crash Reporter is an optional service for NetWitness services. When activated for any of the Core services, the Crash Reporter automatically generates a package of information to be used for diagnosing and solving the problem that resulted in the service failure. The package is automatically sent to NetWitness for analysis. The results are forwarded to NetWitness Support for any further action.

The information package sent to NetWitness does not contain captured data. This information package consists of the following information:

- Stack trace
- Logs
- Configuration settings
- Software version
- CPU information
- Installed RPMs
- Disk geometry

The Crash Reporter crash analysis can be activated for any Core product.

The `crashreporter.cfg` File

One of the files available for editing in the Service Config view > Files tab is `crashreporter.cfg`, the Crash Reporter Client Server configuration file.

This file is used by the script that checks, updates, and builds crash reports on the host. The list of products to monitor can include Network Decoders, Concentrators, Brokers, and hosts.

This table lists the settings for the `crashreporter.cfg` file.

Setting	Description
`applicationlist=decoder, concentrator, host`	Define the list of products to monitor.
`sitedir=/var/crashreporter`	Location of the site directory for the report.
`webdir=/usr/share/crashreporter/Web`	Location of the web directory.
`devdir=/var/crashreporter/Dev`	Location of the development directory.
`datadir=/var/crashreporter/data`	Location of the directory storing data files.
`perldir=/usr/share/crashreporter/perl`	Location of the Perl files.
`bindir=/usr/share/crashreporter/bin`	Location of the binary executables.
`libdir=/usr/share/crashreporter/lib`	Location of the binary libraries.
`cfgdir=/etc/crashreporter`	Location of the configuration files.
`logdir=/var/log/crashreporter`	Location of the log files.
`scriptdir=/usr/share/crashreporter/scripts`	Location of the directory containing scripts.
`workdir=/var/crashreporter/work`	Location of the process work directory.
`sqldir=/var/crashreporter/sql`	Location where created SQL files are placed.
`reportdir=/var/crashreporter/reports`	Location where temporary reports are created.
`packagedir=/var/crashreporter/packages`	Location of the created package files.
`gdbconfig=/etc/crashreporter/crashreporter.gdb`	Location of the `gdb` configuration file.
`corewaittime=30`	Define the number of seconds to wait after finding a core to determine if the core is still being written.
`cyclewaittime=10`	Define the number of minutes to wait between search cycles

Setting	Description
`deletecores=1`	Specify if the Core files should be deleted after report. 0 = No 1 = Yes **Note:** Until the Core file is deleted, it is reported each time crashreporter is restarted.
`deletereportdir=1`	Specify if the report directory should be deleted after the report. Useful to view ore reports on box. 0 = No 1 = Yes **Note:** If not deleted, the directory will be included in each subsequent package.
`debug=1`	Specify whether debugging messages are turned on or off in the `crashreporter` logging output. 0 = No 1 = Yes
`posturl=https://www.netwitnesslive.com/crash...ter/submit.php`	Define the webserver post URL.
`postpackages=0`	Specify if the packages should be posted to the webserver. 0 = No 1 = Yes
`deletepackages=1`	Specify if packages should be deleted after they are posted to webserver. 0 = No 1 = Yes

Configure the Crash Reporter Service

To configure the Crash Reporter service:

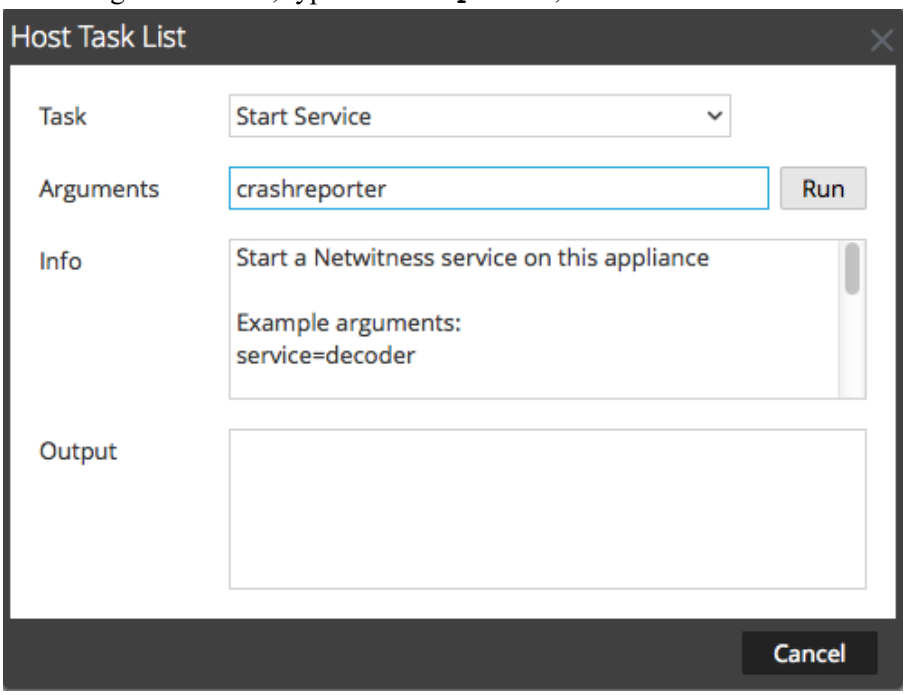
1. Select  **(Admin) > Services**.
2. Select a service and click  **> View > Config**.
3. Select the **Files** tab.
4. Edit `crashreporter.cfg`.

5. Click **Save**.
6. To display the Service System view, select **Config > System**.
7. To restart the service, click Shutdown Service.
   The service shuts down and restarts.

Start and Stop the Crash Reporter Service

To start the Crash Reporter Service:

1. Select **(Admin) > Services**.
2. Select a service and click **> View > System**.
3. In the toolbar, click Host Tasks.
   The Host Task List is displayed.
4. In the Task drop-down list, select **Start Service**.
5. In the Arguments field, type **`crashreporter`**, then click **Run**.



The Crash Reporter service is activated and remains active until you stop it.

To stop the Crash Reporter service, select **Stop Service** from the Task drop-down list.

Maintain the Table Map Files

The table mapping file provided by NetWitness, `table-map.xml`, is a very important part of the Log Decoder. It is a meta definition file which also maps the keys used in a log parser to the keys in the `metadb`.

Note: Do not edit the `table-map.xml` file. If you want to make changes to the table-map, make them in the `table-map-custom.xml` file. The latest `table-map.xml` file is available on Live Services, which NetWitness updates as required. If you make changes to the `table-map.xml` file, they can be overwritten during a content or service upgrade.

The table map and custom table map files have two purposes:

- To translate the variables used in the Log Parsers to NetWitness meta key names
- To tell the system which keys to move onto the Concentrator.

For example, look at the out-of-the-box Palo Alto log parser, and examine one of its meta keys: `stransaddr`. This key represents the source translated address. If we look in the `table-map.xml`file we can see that this variable is listed as `Transient`:

```
<mapping envisionName="stransaddr" nwName="stransaddr" flags="Transient"
format="Text" />
```

Because this variable is listed as, `Transient`, it never moved to the Concentrator. In fact, if you look at all the metadata that we parse from that log in the Concentrator, it is not listed as an available key.

Assume we change the value in the `table-map-custom.xml`file to the following:

```
<mapping envisionName="stransaddr" nwName="stransaddr" flags="None"
format="Text" />
```

In this case, the key-value pair would get copied to the Concentrator, and from there you can choose whether or not to index it.

In the`table-map.xml`file, some meta keys are set to `Transient` and some are set to `None`. To store and index a specific meta key, the key must be set to `None`. To make changes to the mapping, you need to create a copy of the file named `table-map-custom.xml` on the Log Decoder and set the meta keys to `None`.

For meta key indexing:

- When a key is marked as `None` in the `table-map.xml` file in the Log Decoder, it is indexed.
- When a key is marked as `Transient` in the `table-map.xml` file in the Log Decoder, it is not indexed. To index the key, copy the entry to the `table-map-custom.xml` file and change the keyword `flags="Transient"` to `flags="None"`.
- If a key does not exist in the `table-map.xml`file, add an entry to the `table-map-custom.xml`file in the Log Decoder.

IMPORTANT: Do not update the **`table-map.xml`** file because an upgrade can overwrite it. Add all of the changes that you want to make to the **`table-map-custom.xml`**file.

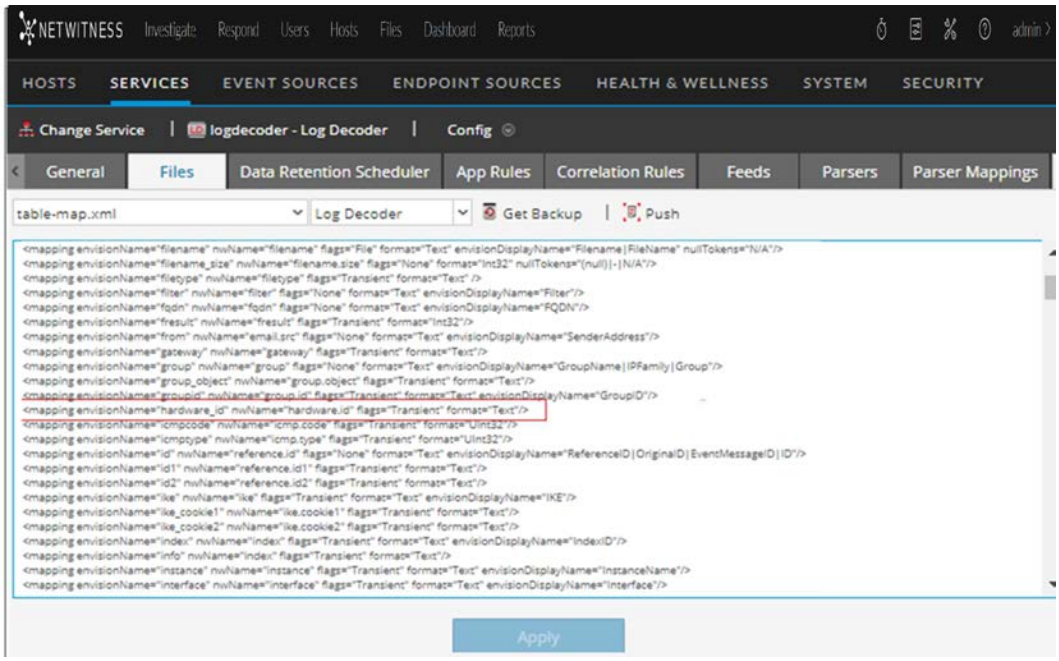
Prerequisites

If you do not have a `table-map-custom.xml` file on the Log Decoder, create a copy of `table-map.xml` and rename it to `table-map-custom.xml`.

To verify and update the table mapping file:

1. Go to  **(Admin) > Services**.
2. In the Services list, select a Log Decoder and click **> View > Config**.

3. Click the **Files** tab and select the **`table-map.xml`** file.



4. Verify that the `flags` keywords are set correctly to either `Transient` or `None`.
5. If you need to change an entry, do not change the `table-map.xml` file. Instead, copy the entry, select the `table-map-custom.xml` file, find the entry in the `table-map-custom.xml` file and change the `flags` keyword from `Transient` to `None`.
   For example, the following entry for the `hardware.id` meta key in the `table-map.xml` file is not indexed and the `flags` keyword shows as `Transient`:
   `<mapping envisionName="hardware_id" nwName="hardware.id"`
   `flags="Transient"``/>`
   To index the `hardware.id` meta key, change the flags keyword from `Transient` to `None` in the `table-map-custom.xml` file:
   `<mapping envisionName="hardware_id" nwName="hardware.id"` **`flags="None"`**`/>`
6. If an entry does not exist in the `table-map.xml` file, add an entry to the `table-map-custom.xml` file.
7. After making your changes to the `table-map-custom.xml` file, click **Apply**.

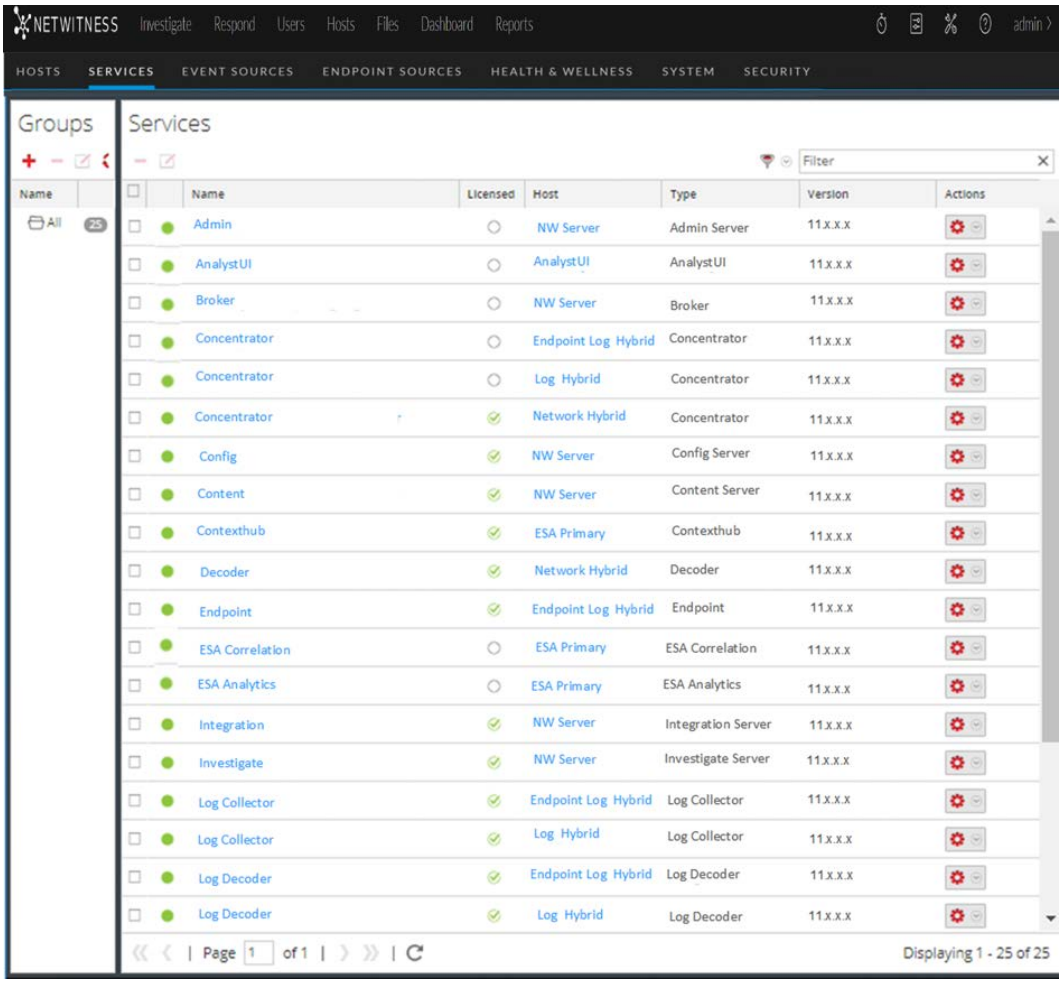
Caution: Before changing the table mapping files, carefully consider the effect of changing the index from `Transient` to `None` because it can impact the available storage and performance of the Log Decoder. For this reason, only certain meta keys are indexed out-of-the-box. Use the `table-map-custom.xml` file for different use cases.

Edit or Delete a Service

You can edit service settings, such as changing the host name or port number, or deleting a service that you no longer need.

Each of the following procedures starts in the Services view.

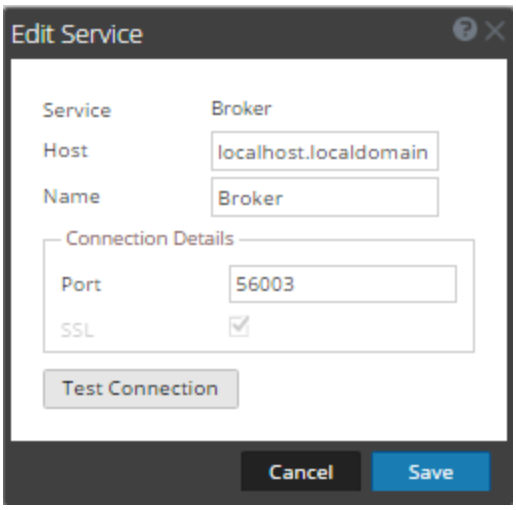
To navigate to the Services view, in NetWitness, go to **(Admin) > Services**.



Edit a Service

1. In the Services view, select a service and either click or > **Edit**.

   The **Edit Service** dialog is displayed. It shows only the fields that apply to the selected service.



2. Edit the service details by changing any of the following fields:
   - **Name**
   - **Port** - Each Core service has two ports, SSL and non-SSL.
   - **SSL** - For trusted connections, you must use SSL.
   - **Username** and **Password** - Use these credentials to test the connection to a service.
     a. If you use a trusted connection, delete the username.
     If you do not use a trusted connection, type a username and password.
     b. Click **Test Connection**.
3. Click **Save**.

Delete a Service

1. In the Services view, select one or more services and either click − or > **Delete**.
2. A dialog requests confirmation. To delete the service, click **Yes**.

The deleted service is no longer available to the NetWitness modules.

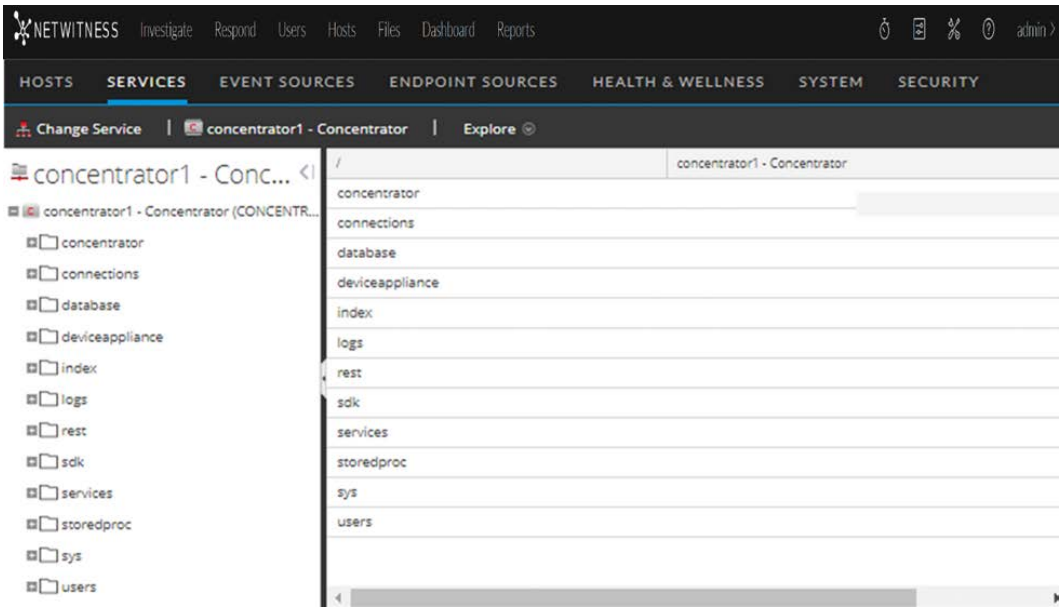
Explore and Edit Service Property Tree

You have advanced access and control of service functionality in the Services Explore view, which consists of two parts. The Node list displays service functionality in a tree structure of folders. The Monitor panel displays properties of the folder or file selected in the Nodes list.

Each of the following procedures starts in the Explore view.

To navigate to the Explore view:

1. In NetWitness, go to **(Admin) > Services**.
2. Select a service, then select **> View > Explore**.
   The Explore view is displayed. The Node list is on the left and the Monitor panel is on the right.



Display or Edit a Service Property

To display a service property:

1. Right-click a file in the Node list or Monitor panel.
2. Click **Properties**.

To edit the value of a service property:

1. In the **Monitor** panel, select an editable property value.
2. Type a new value.

Send a Message to a Node

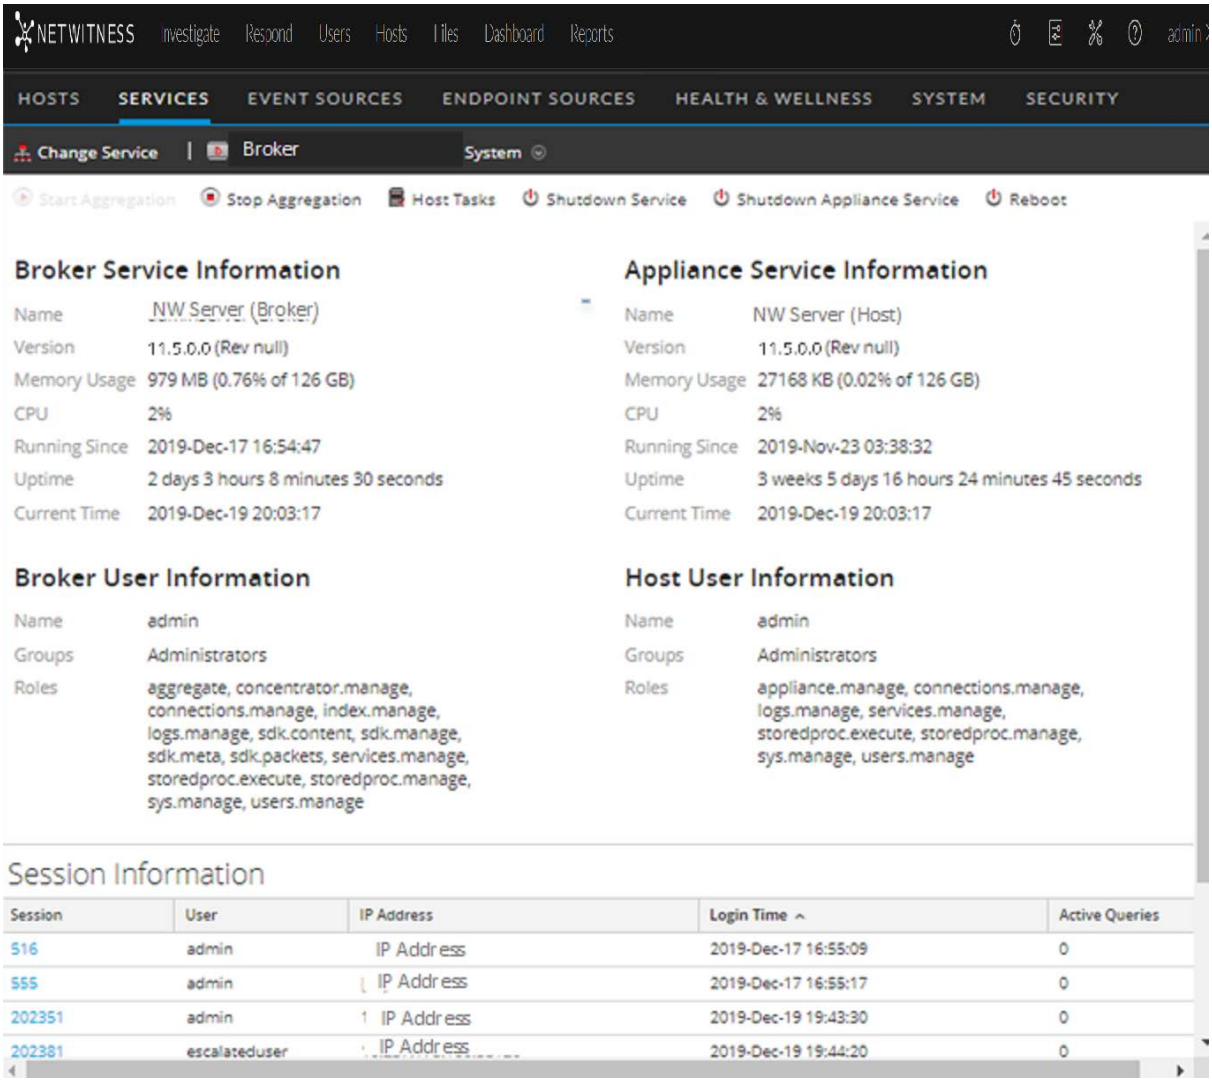
1. In the Properties dialog, select a message type from the drop-down list. Options vary according to the file selected in the Node list.
   A description of the selected message type is displayed in the **Message Help** field.
2. (Optional) If the message requires them, type the **Parameters**.
3. Click **Send**.
   The value or format is displayed in the Response Output field.

Terminate a Connection to a Service

You can view sessions that are running on a service in the Service System view. From within the list of sessions, you can terminate the session and the active queries in a session.

Terminate a Session on a Service

1. In NetWitness, go to **(Admin) > Services**.
   The Admin Services view is displayed.
2. Select a service, and select **> View > System**.
   The Services System view is displayed.



3. In the **Session Information** list at the bottom, click a session number from the Session column.

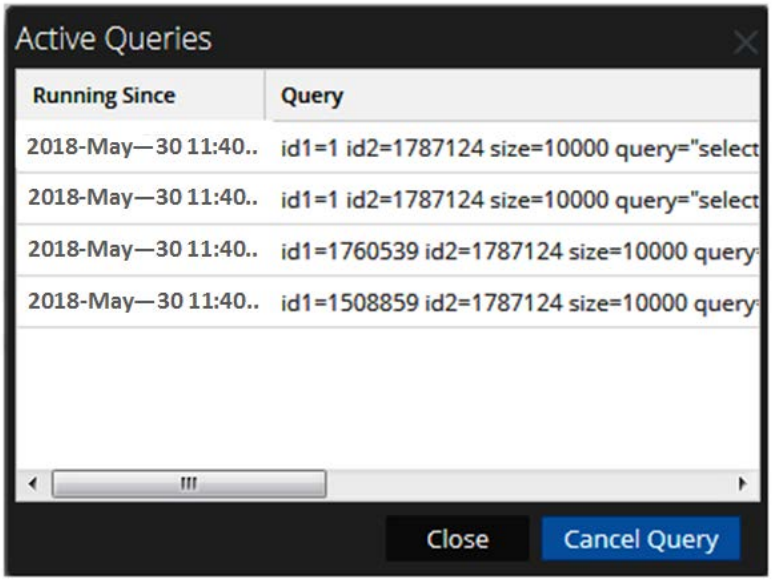
The confirmation dialog is displayed.

4. Click **Yes**.

Terminate an Active Query in a Session

1. Scroll down to the **Sessions** list.
2. In the **Active Queries** column, click a non-zero count of active queries for a session. You cannot click on it if there are 0 active queries.

   The Active Queries dialog is displayed.



3. Select a query and click **Cancel Query**.

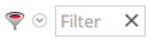
   The query stops and the Active Queries column is updated.

Search for Services

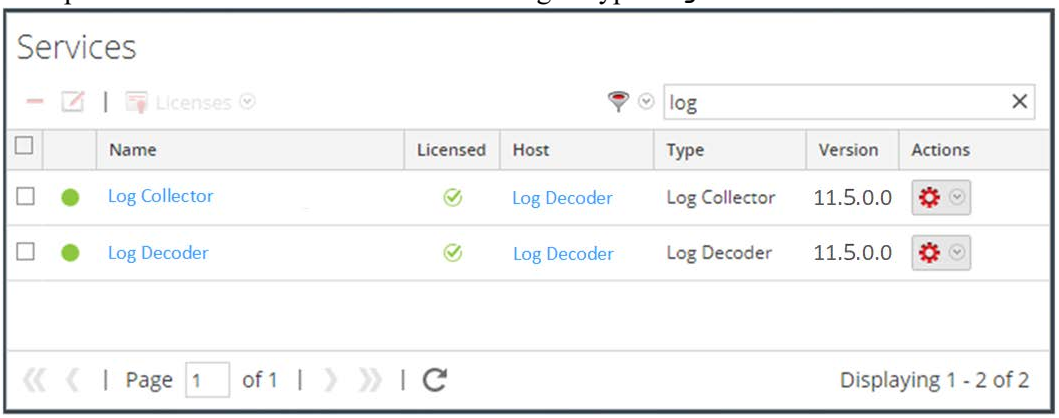
You can search for services from the list of services in the Services view. The Services view enables you to quickly filter the list of services by Name, Host, and Service Type. You can use the Filter drop-down menu and the Filter field separately or at the same time to filter the Services view.

Search for a Service

1. In NetWitness, go to  **(Admin) > Services**.
2. In the **Services** list toolbar, type a service **Name**, **Host**, or service **Type** in the **Filter** field.

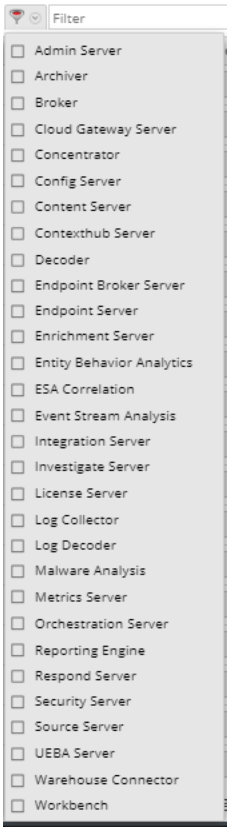


The Services panel lists the services that match the names entered in the Filter field. The following example shows the search results after starting to type **log** in the filter field.

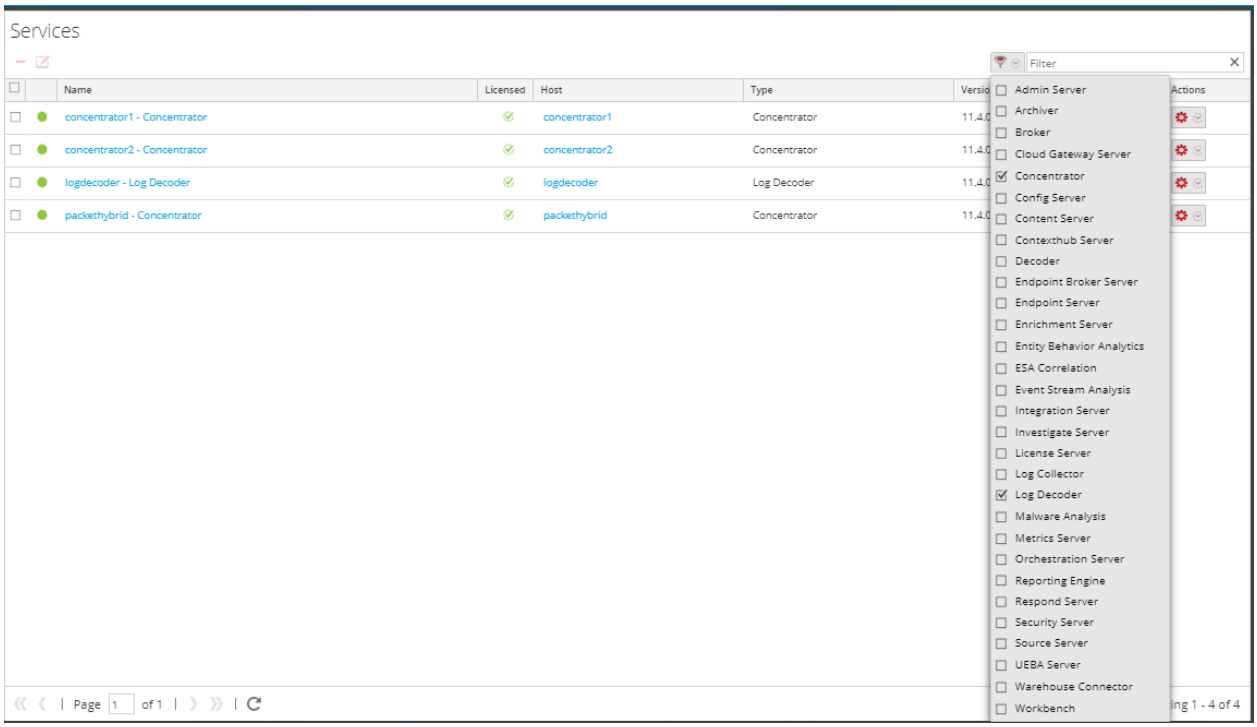


Filter Services by Type

1. In NetWitness, go to **(Admin) > Services**.
2. In the Services view, click and select the service types that you want to appear in the Services view.



The selected service types appear in the Services view. The following example shows the Services view filtered for Concentrator and Log Decoder.



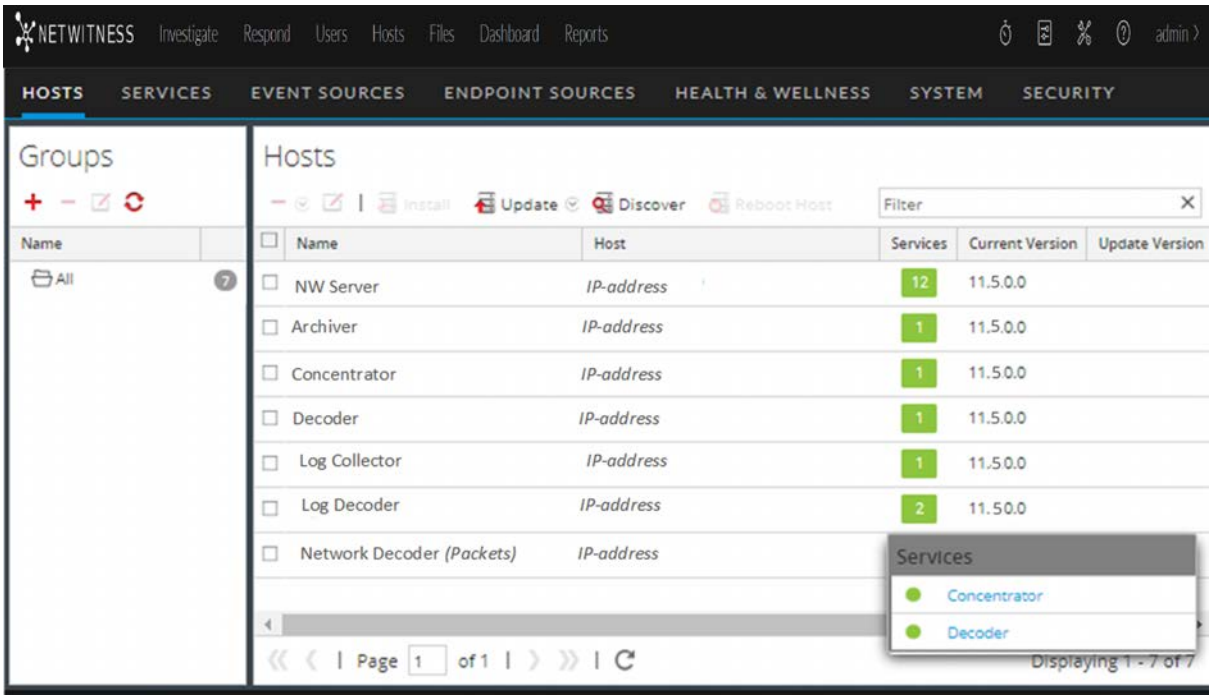
Find the Services on a Host

In addition to being able to locate the services for a host in the Services view, you can also quickly find the services that run on a host in the Hosts view.

1. In NetWitness, go to **(Admin) > Hosts**.
2. In the Hosts view, select a host and click the box that contains a number (the number of services) in the **Services** column.

   A list of the services on the selected host is displayed.

   In the following example, a list of two services on the selected host are listed after clicking the box containing the number 2.



3. You can click the service links to view the services in the Services view.

Start, Stop, or Restart a Service

These procedures apply to Core services only.

Each of the following procedures starts in the Services view. In NetWitness, go to **(Admin) > Services**.

Start a Service

1. Select a service and click > **Start**.

Stop a Service

When you stop a service, all of its processes stop and active users are disconnected from it.

To stop a service:

1. Select a service and click > **Stop**.
2. A dialog requests confirmation. To stop the service, click **Yes**.

Restart a Service

Occasionally, you have to restart a service for changes to take effect. When you change a parameter that requires a restart, NetWitness displays a message.

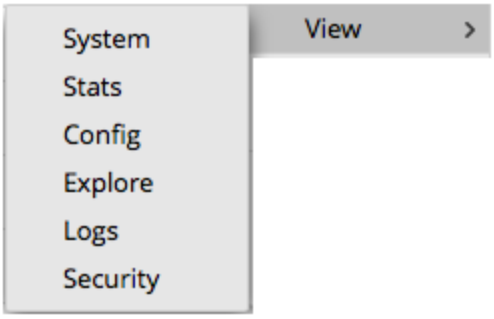
To restart a service:

1. Select a service and click > **Restart**.
2. A dialog requests confirmation. To stop the service, click **Yes**.

The service stops, then restarts automatically.

View Service Details

You can view and edit information about services using options in the View menu for a service.



Purpose of Each Service View

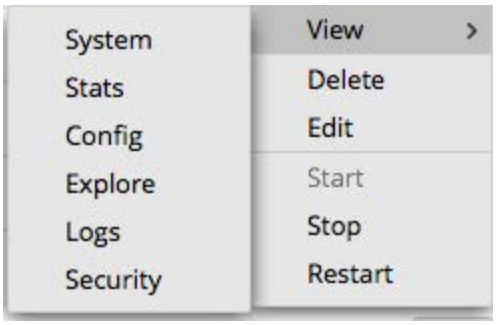
Each view displays a functional piece of a service and is described in detail in its own section:

- Services System View shows a summary of service, appliance service, service user, host user, and session information.
- Services Stats View provides a way to monitor service operations and status.
- Services Config View is for configuring all aspects of a service.
- Services Explore View is for viewing and editing host and service configurations.
- System Logging Panel shows service logs that you can search.
- Services Security View is a way to add NetWitness Platform Core user accounts for aggregation, thick client users, and REST API users.

Access a Service View

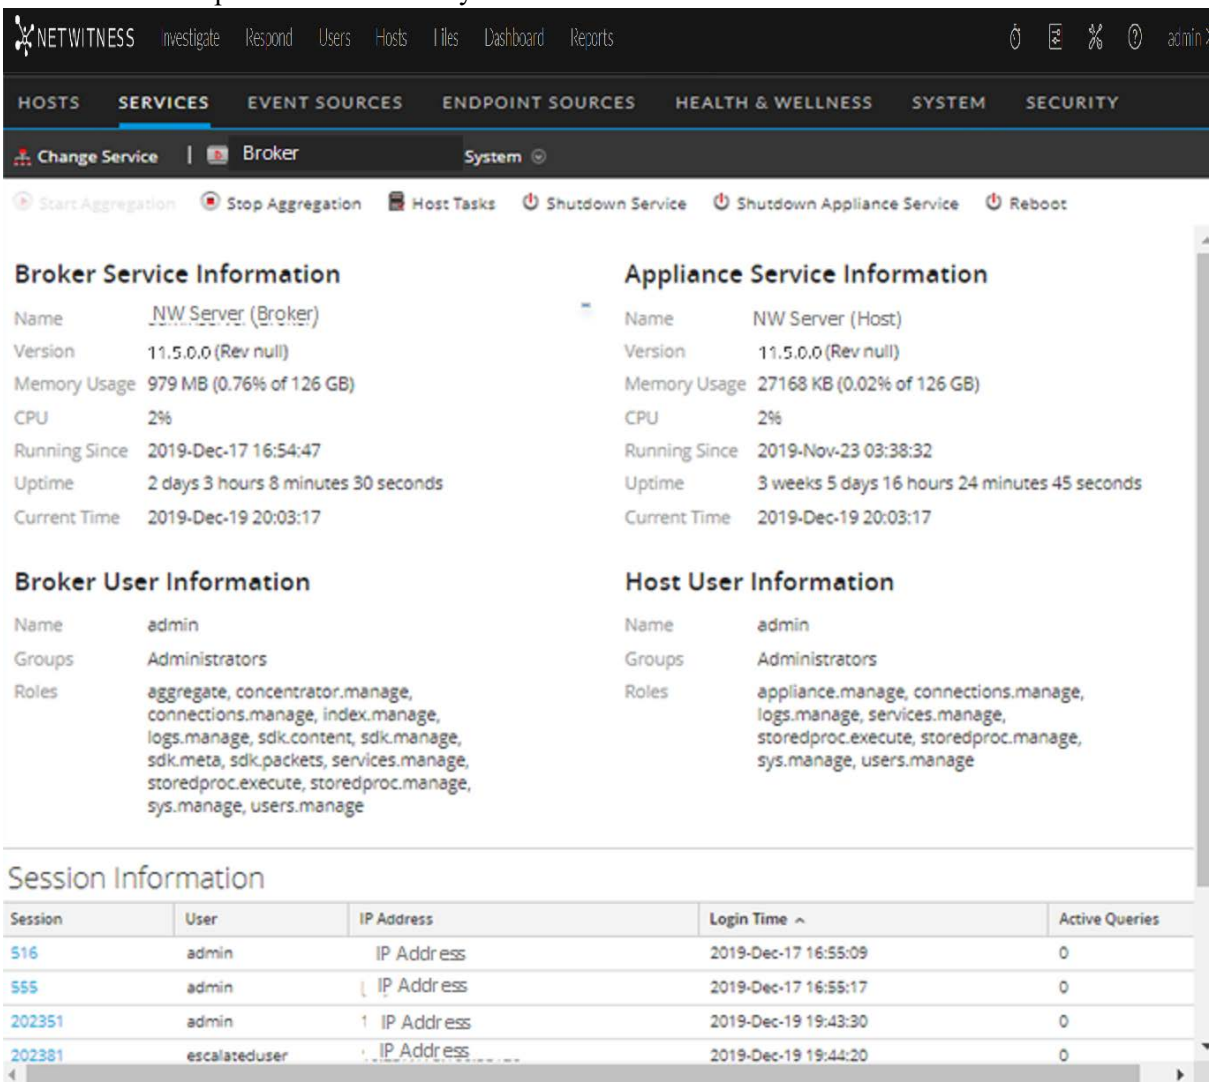
To access a view for a service:

1. In NetWitness, go to **(Admin) > Services**.
2. Select a service and click > **View**.
   The View menu is displayed.

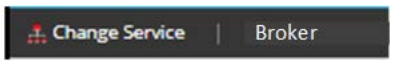


3. From the options on the left, select a view.

   Below is an example of the Services System view for a Broker.



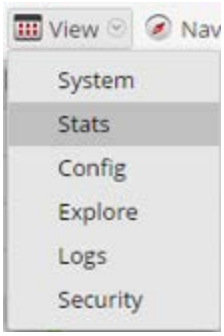
4. Use the toolbar to navigate:



a. Click **Change Service** to select another service.
The Administrate Service dialog is displayed.

b. Select the checkbox to the left of the service that you want.

c. Select the view that you want for the service you selected in the View drop-down list.



The new view (for example, Stats) is displayed for the service you selected.

View Topology Details

Note: This feature is supported in 11.7 and later.

Note: To view the **Service Topology Tab**, you must ensure this feature is enabled under **Admin Server** > **Explorer** > **Feature**> **service-topology-feature**.

Administrators and analysts can view all the NetWitness core services in a hierarchical layout depicting the collection and aggregation of the services in your deployment. This visualization displays the topology for Broker, Concentrator, Log Decoder, Packet Decoder, Hybrids, ESA and Log Collector and provides insights on which:

- Broker aggregates from the Concentrators.
- Concentrator aggregates from the Log Decoders.
- Log Decoder receives data from the Log Collectors.
- Log Collector sending data to other Log Collectors and vice-veNetWitness and the direction of the data flow.

Note: Services such as Reporting Engine, Malware Analysis, UEBA, Endpoint Server, Cloud Link Service, and Warehouse connector are not supported.

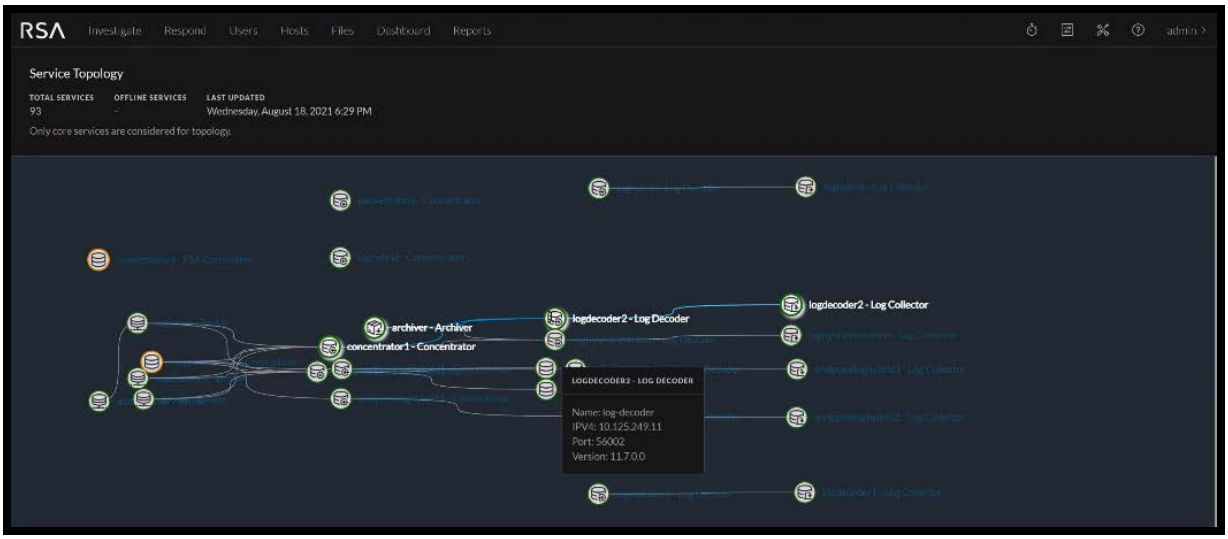
The topology is automatically refreshed once in 3 hours. You can click on any node to view the details.

To view the Topology:

1. Click > **Service Topology.**
2. Do one of the following:

- Scroll horizontally to view the complete topology.
- Click **Search** and enter the service name to locate a specific service. The specific service will be highlighted in blue color in the hierarchical layout.

3. Click on the node to view the details.
   For more information, see Services Topology View topic.



Centralized Service Configuration via Policy

Note: The information in this section applies to NetWitness Platform Version 11.7 and Later.

Centralized Service Configuration allows you to manage the configuration of services in your environment efficiently. The Decoder, Concentrator, and Log Decoder deployed in your environment may be large in number and geographically distributed. Managing common configurations across the services can be time consuming. With Centralized Service Configuration, you can centrally create a policy of common configuration settings and apply the policy to a set of groups.

Groups

A group is a set of services based on the service type Decoder, Concentrator, and Log Decoder. You can create a group and assign policy to it based on your requirement. For example, group all the 10G Decoder services or group services within a geographical area. For more information on creating groups, see Creating Groups and Policies.

Policies

A Policy is a set of service configurations that you can apply to a set of groups. This allows you to efficiently manage the service settings. Once you create a customized policy, you can assign it to a group.

Benefits of Centralized Service Configuration

- Apply customized settings in one step to any number of services

- Centrally restart all services within a group (when needed) to apply changes
- Indicates when an action is required, such as service restart, failed publication, or Out of Compliance services
- Clone policy to quickly create similar policies with minor changes
- Quickly revert changes to a policy in case of issues
- Groups of the same service type can be created based on similar hardware profiles or other criteria
- Add configuration items to policies in order to customize settings. Any settings which are not in the policy will be left as default
- (For version 11.7.1 and Later) Reconfigure 10G Network Decoders from the Policy UI. Administrators can quickly create 10G policies for each Decoder group based on the hardware profile
- (For version 11.7.1 and Later) Clone policy from an existing service, to save policy transition time for existing users

Creating Groups and Policies

Create a Group

You can create a group with one or more services and assign one policy to it. Atleast one service is required to publish a group.

To create a group:

1. Go to  (**CONFIGURE**) > **Policies**.
2. In the left panel, click **Groups**.
3. In the tool bar, click .
4. In the New Group panel, do the following:
   - Enter the name of the group.
   - Enter a description for the group.
   - Select the type of service for the group.
5. Click **Next**.
6. In the Define Group, click  to assign services to the group.

Note: A service is disabled if it is assigned to another group.

7. Click **Next**.
8. In the Assign Policies, click  to assign one policy to a group. Make sure that the policy and group is of same service type.
9. Do any one of the following:

- Click **Save and Publish** to save and publish the policy.

  The Publish and Restart Services dialog is displayed only if some settings require a service restart for the changes to take effect. In the Publish and Restart Services dialog, do any one of the following:

  - To publish and restart the services immediately, click **Publish & Restart**.
  - To publish and restart the services later, click **Publish & Restart Later**. In this case you must restart the services manually.
  - To cancel the restart dialog, click **Cancel**.

  The settings that do not require restart will be applied immediately.

- Click **Save and Close** to save the settings.

Create a Policy

You can create a policy and assign it to one or more groups.

To create a policy:

1. Go to  **(CONFIGURE)** > **Policies**.
2. Click **Policies**.
   The available policies are displayed.
3. Click 

to add a new policy.
4. In the New Policy panel, do the following:
   - Select the policy type from the drop-down list.
   - Enter a unique policy name.
   - Enter a description for the policy.
   - Select the service type for the policy from the drop-down list.
   - (For version 11.7.1 and Later) Select the service from which you want to clone the settings.
5. Click **Next**.
6. Do any one of the following:
   - Customize the settings for the policy based on your requirement, for example, in the Database setting, click  and change the Policy value to sha256 or change the byte range to 10 MB.

     Note: (For version 11.7.1 and Later) If you have preferred to clone from existing service, the setting values are populated automatically from the selected service on customize settings section. You can also add or modify the settings based on your requirement.

   - (For version 11.7.1 and Later) To configure 10g Settings:
     - Click **10g Settings**.

- Select the Decoder service to fetch the settings.
- Click **Done**.

7. In the Group List, click  to assign groups to the policy. The policy and group should be of same service type.

Note: A group is disabled if another policy of the same type is already assigned to this group.

8. Do any one of the following:
   - Click **Save and Publish** to save and publish the policy.

     The Publish and Restart Services dialog is displayed only if some settings require a service restart for the changes to take effect. In the Publish and Restart Services dialog, do any one of the following:
     - To publish and restart the services immediately, click **Publish & Restart Now**.
     - To publish and restart the services later, click **Publish & Restart Later**. In this case you must restart the services manually.
     - To cancel the restart dialog, click **Cancel**.

     The settings that do not require restart will be applied immediately.
   - Click **Save and Close** to save and return to the Policies view.

Managing Groups and Policies

View a Group

To view properties of the selected group:

1. Go to  (**CONFIGURE**) > **Policies**.
2. Click **Groups**. The available groups are displayed.
3. Click a row to view details about the selected group in the right panel.

Delete a group

You can delete one or more groups. Once a group is deleted, the services and policies in the group will no longer displays the deleted group details.

To delete a group:

1. Go to  (**CONFIGURE**) > **Policies**.
2. Click **Groups**. The available groups are displayed.
3. Select one or more groups and click **Delete**.

The confirmation message is displayed.

Edit a Group

You can edit the properties of the group at any point in time. The service type cannot be edited. The status of the updated group is unpublished if you change the service or policies in a group. If you just change the group name and description the status remains published (if it is already published).

To edit a group:

1. Go to  (**CONFIGURE**) > **Policies**.
2. Click **Groups**. The available groups are displayed.
3. Select a group and click **Edit**.

Note: You cannot edit the service type.

4. Make the required changes in the group.
5. Do any one of the following:
   - Click **Save and Publish** to save and publish the policy.
     The policy will be listed under the published category.
   - Click **Save and Close** to save the settings and return to the Groups view.

Filter Groups

The Filters Panel allows you to filter the list of displayed groups, based on the service type:

- Decoder
- Log Decoder
- Concentrator

Additionally, you can filter based on publication status or service status:

- **Published**: Groups that are published.
- **Unpublished**: Groups that are saved but not published.
- **Failed**: Groups that are failed to publish.
- **N/A**: Groups for which publication status is not applicable.
- **Service Require Restart**: The services associated with the group that require restart.

The Filters panel can be hidden or displayed:

- To display if hidden, click  in the toolbar.
- To hide, click  at the top-right of the panel.

Delete a Policy

To delete a policy:

1. Go to (**CONFIGURE**) > **Policies**.
2. Click **Policies**. The available policies are displayed.
3. Select one or more policies and in the **More Actions** drop-down list in the tool bar, click **Delete**.

The delete dialog is displayed.

Note: The services associated with this policy still require a restart if the restart is pending.

4. Click **Delete** to permanently delete the selected policies.

The deletion will take effect immediately.

Revert a Policy

You can revert a policy to the previously published version for a maximum of 5 times. Once a policy is reverted you cannot get back to the newer version and you must publish a policy.

Note: The revert is disabled if there is no previously published version.

To revert a policy:

1. Go to (**CONFIGURE**) > **Policies**.
2. Click **Policies**. The available policies are displayed.
3. Select a policy and in the **More Actions** drop-down list in the tool bar, click **Revert**.
   All changes made to policy will be discarded and reverted to the previously published version. You must publish the policy once it is reverted.

Note: If there are any missing settings in the previously published version, then a default value is set for those settings.

Clone a Policy

You can clone only one policy at a time. Once cloned, all the settings from the old policy is copied to the new policy.

To clone a policy:

1. Go to (**CONFIGURE**) > **Policies**.
2. Click **Policies**. The available policies are displayed.
3. Select a policy to clone and in the More actions drop-down list in the tool bar, click **Clone**.

The policy is cloned successfully.

Edit a Policy

You can edit the settings of the policies. The service type and policy type cannot be edited. Once the policy is edited the changes in the policy is reflected upon saving the policy. The updated settings are applied to the service if published.

After saving and before publishing the publication status of the changed policy is set to Unpublished if any settings is changed.

To edit a policy:

1. Go to  (**CONFIGURE**) > **Policies**.
2. Click **Policies**. The available policies are displayed.
3. Select a policy and click **Edit**.
4. Make the required changes in policy.

Note: You cannot edit service type and policy type.

5. Do any one of the following:
   - Click **Save and Publish** to save and publish the policy.
     The policy will be listed under the published category.
   - Click **Save and Close** to save the settings and return to the Policies view.

View a Policy

To view a policy:

1. Go to  (**CONFIGURE**) > **Policies**.
2. Click **Policies**. The available policies are displayed.
3. Click a row to view details about the selected policy in the right panel.

Filter Policies

The Filters Panel allows you to filter the list of displayed policies, based on the policy status and service type:

- **Published**: Policies that are published.
- **Unpublished**: Policies that are saved but not published.
- **Failed**: Policies that are failed to publish.
- **N/A**: Policies for which publication status is not applicable.

The Filters panel can be hidden or displayed:

- To display if hidden, click  in the toolbar.
- To hide, click  at the top-right of the panel.

Hosts and Services Views References

This topic is a reference for features in the NetWitness Admin user interface.

The Admin module pulls NetWitness Admin activities into a single view to monitor and manage hosts (appliances), services, tasks, and security.

Topics

- Hosts View
- Services View
- Services Config View
- Services Explore View
- Services Logs View
- Services Security View
- Services Stats View
- Services System View

Hosts View

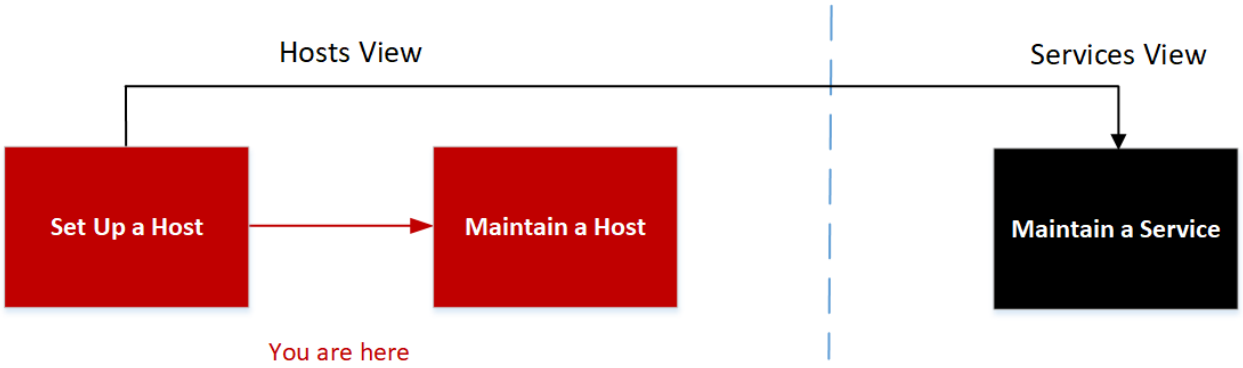
You set up and maintain the physical or virtual machine on which NetWitness services run in the Hosts view.

IMPORTANT: For help with resolving errors you receive during version installation and update, see Troubleshooting Version Installations and Updates .

A service performs a unique function, such as collecting logs or archiving data. Each service runs on a dedicated port and is modeled as a plug-in to enable or disable, according to the function of the host. You must configure the Core services first. You must configure hosts and services to communicate with the network and each other so they can perform their functions such as storing or capturing data.

Workflow

This workflow shows the procedures you complete to set up a host, maintain a host, and update the host with new NetWitness versions. Setting up a host is the first task in this workflow.The hosts with Core services are set up out-of-the-box. After that, you can set up additional hosts to enhance your NetWitness deployment. The other two tasks, maintaining a host and updating versions for a host, are performed when required and do not have a specific order of completion.



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.*	Setting Up a Host
Administrator	maintain a host.*	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	apply version updates to a host.*	Apply Version Upgrades to a Host

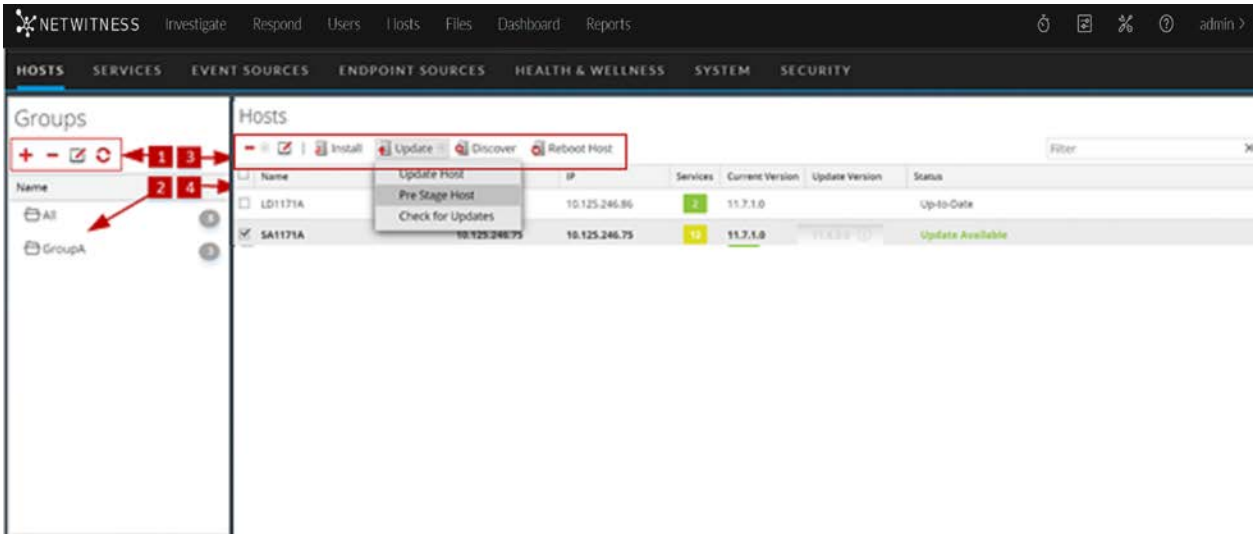
* You can perform these tasks in the current view.

Related Topics:

- Services View
- Hosts and Services Set Up Procedures
- Hosts and Services Maintenance Procedures

Quick Look

This is an example of the Hosts view.



1 Groups Panel Toolbar - Provides options to work with host groups in the list.
2 Groups Panel - Lists all host groups currently in your deployment.
3 Hosts List Toolbar - Provides options to work with the Hosts list.
4 Hosts List - Lists all hosts currently in your deployment.

Groups Panel Toolbar

Feature	Description
+	Displays a new row in the Groups panel in which you enter the name of a new group.
−	Asks for confirmation that you want to delete the group. You can confirm or cancel the deletion.
	Opens the field for renaming the selected preexisting group. You can also double click on the group name in the Groups panel to rename the group. Changes take effect immediately.
	Refreshes the Groups panel to reflect the changes and goes back to the All group view. Changes take effect immediately.

Groups Panel

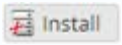
The Groups panel provides a logical way to organize hosts, such as by function, geography, or project. Once hosts are grouped, it is easier to perform operations on multiple hosts by interacting with each host in a group rather than individual hosts from an non-grouped list. You can drag a host from the Hosts list to add it to a group. A host may belong to more than one group.

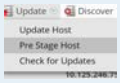
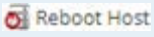
Note: In NetWitness Live, groups can subscribe to resources while individual hosts cannot.

Column Title	Description
Name	The host groups are displayed in the Groups panel. The number next to each group name displays the number of hosts that added to the group.

Hosts List Toolbar

The Hosts list toolbar contains the tools that you use to maintain the hosts in your NetWitness deployment.

Feature	Description
	This drop-down list displays two options. • **Remove Host**: Deletes the selected host from both the host group and the Hosts list altogether. All related services will be removed as well. • **Remove From Group**: Removes the selected host from the host group. The Host will still be available in the Hosts list. Changes take effect immediately.
	Opens the Edit Host dialog in which you edit a host or service identification and basic communication settings. See Step 1. Deploy a Host for more information.
	Opens the Install Services dialog from which you can install a service on a deployed host. See Step 2. Install a Service on a Host for more information.

Feature	Description
	This drop-down list displays three options. • **Update Host**: Updates the selected hosts with the version you select in the Update Version column. • **Pre-Stage Host:** Pre-stage the upgrade repository by downloading the required packages without affecting the system. • **Check for Updates**: Checks the Local Update Repo for the latest updates available from NetWitness. Changes take effect immediately. See Apply Version Upgrades to a Host for more information.
	Most of the time, the Discovery function completes automatically and you do not need to click  . For a fresh installation, click Discover to access the Provision dialog box so you can complete the provisioning phase. After the provisioning phase, NetWitness automatically discovers services running on the host and you do not need to click Discover.
	Restarts the host immediately.
	Entering a Name or Host here filters the list. This field allows you to quickly find a particular host.

Hosts List

Column	Description
☐	Select the host by clicking the corresponding checkbox in this column. To select all of the hosts, select the checkbox in the header.
Name	Displays the name of the host that was given when the host was installed. This column is organized in alphabetical order by default. Click the Name column title to view in reverse alphabetical order.
Host	Displays the IP address of each host.
Services	Displays the number of services added to the host.
Current Version	Displays the version that the host is currently on.
Update Version	Displays a drop-down list of versions that the user can upgrade to. See Apply Version Upgrades to a Host for more information.

Column	Description
Status	Displays whether or not the host is upgraded to the most current version. If the host is on the most current version, then the Status displays "Up-to-Date".

Services View

You set up and maintain the NetWitness services in the Services view. In the Services view, you can:

- Quickly search for and locate a specific service or type of service, such as Log Decoder or Warehouse Connector.
- Use shortcuts to get to administration tasks.
- Add, edit, and remove services.
- Sort services by name and host.
- Filter services by type, name, and host.
- Start, stop, and restart services.

A service performs a unique function, such as collecting logs or archiving data. Each service runs on a dedicated port and is modeled as a plug-in to enable or disable, according to the function of the host. You must configure the following Core services first.

Services	Notes
NW Server	
Admin Config Content Integration Investigate License Orchestration Reporting Engine Respond Security	Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server Implemented with the NW Server
Analyst UI	
Broker Investigate Server NetWitness UI Reporting Engine Respond Server	Implemented with the Analyst UI Implemented with the Analyst UI Implemented with the Analyst UI Implemented with the Analyst UI Implemented with the Analyst UI
Archiver	
Archiver Workbench	Core Service
Broker	
Broker	Core Service

Services	Notes
Concentrator	
Concentrator	Core Service
Endpoint	
Endpoint Server	
Endpoint Broker	
Endpoint Broker Server	
Endpoint Log Hybrid	
Log Collector Log Decoder Endpoint Server Concentrator	Core Service Core Service Core Service
ESA Primary	
Contexthub ESA Correlation	
ESA Secondary	
ESA Correlation	
Log Collector	
Log Collector	Core Service
Log Decoder	
Log Collector Log Decoder	 Core Service
Log Hybrid	
Log Collector Log Decoder Concentrator	 Core Service Core Service
Log Hybrid - Retention	Deployed on Series 6 Hybrid hardware with Log Hybrid-Retention Optimization.
Log Collector Log Decoder	 Core Service
Malware Analysis	
Malware Analysis Broker	 Core Service
Network Decoder	

Services	Notes
Decoder (Packets)	Core Service
Network Hybrid	
Concentrator Network Decoder	Core Service Core Service
New Health and Wellness	
Metrics Server	
UEBA	
UEBA	
Warehouse Connector	
Warehouse Connector	Command line installation

You must configure hosts and services to communicate with the network and each other so they can perform their functions such as storing or capturing data. For information about ports and a comprehensive list of ports for all services, see "Network Architecture and Ports" in the *Deployment Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Workflow

This workflow shows the procedures you complete to set up and maintain a service. Adding a service to a host is the first task in this workflow. The hosts with Core services are set up out-of-the-box. After that, you can set up additional services on hosts to enhance your NetWitness deployment.



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.*	Maintaining Services

* You can perform these tasks in the current view.

Related Topics

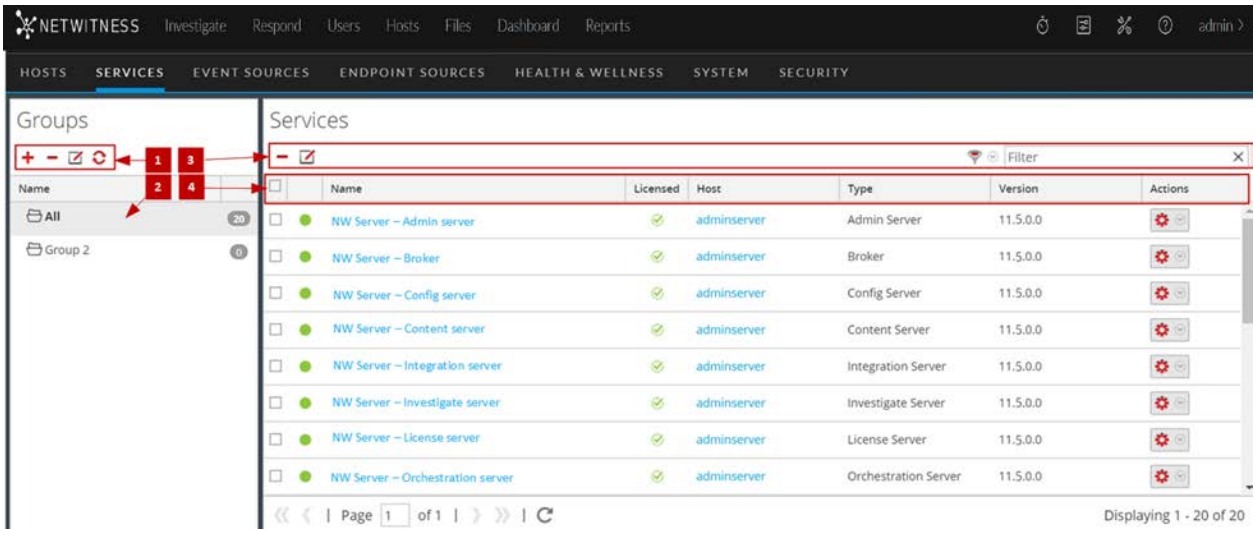
- View Topology Details
- Hosts View
- Hosts and Services Set Up Procedures
- Hosts and Services Maintenance Procedures

See the following NetWitness guides for detailed information on individual services. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Archiver Configuration Guide
Broker and Concentrator Configuration Guide
Context Hub Configuration Guide
Decoder Configuration Guide
Endpoint Configuration Guide
Event Stream Analysis (ESA) Configuration Guide
Malware Analysis Configuration Guide
Log Collection Configuration Guide
Malware Analysis Configuration Guide
Reporting Engine User Guide
NetWitness Respond Configuration Guide
NetWitness UEBA User Guide
Workbench Configuration Guide
Warehouse Connector Configuration Guide

Quick Look

This is an example of the Services view.



1 Groups Panel Toolbar - Provides options to work with service groups in the list.
2 Groups Panel - Lists all service groups currently in your deployment.
3 Services List Toolbar - Provides options to work with the Services list.
4 Services List - Lists all services currently in your deployment.

Groups Panel Toolbar

Feature	Description
+	Displays a new row in the Groups panel in which you enter the name of a new group.
−	Asks for confirmation that you want to delete the group. You can confirm or cancel the deletion.
	Opens the field for renaming the selected preexisting group. You can also double click on the group name in the Groups panel to rename the group. Changes take effect immediately.
	Refreshes the Groups panel to reflect the changes and goes back to the All group view. Changes take effect immediately.

Groups Panel

The Groups panel provides a logical way to manage groups of services, such as by function, geography, or project. After you create a group, you can drag individual services from the Services panel into the group. A service may belong to more than one group.

Column Title	Description
Name	The service groups are displayed in the Groups panel. The number next to each group name displays the number of hosts that added to the group.

Services List Toolbar

This topic introduces the options in Services list toolbar to add, remove, edit, and get a license for services. You can also filter the services listed in the Services list.

To access the Admin Services view, in NetWitness, go to **(Admin) > Services.** The Services list toolbar is at the top of the Services list in the Services view.

Feature	Description
+	Adds a service for your deployment of NetWitness to manage. See Step 2. Install a Service on a Host.
−	Deletes a service from your deployment of NetWitness. See Edit or Delete a Service.
	Edits service identification and basic communication settings.
Filter	Filters the services listed in Services view. In the Filter drop-down list, you can filter the services by one or more selected service types. In the Filter field, you can filter the services by Name and Host. You can use the Filter drop-down list and the Filter field at the same time to filter the services listed in the Services view.

Services List

Column	Description
☐	Select the service by clicking the corresponding checkbox in this column. To select all of the services, select the checkbox in the header.
Online/Offline Indicator	Displays ● if the service is online. Displays ● if the service is offline.
Name	Displays the name of the service that was given when the service was installed. This column is organized in alphabetical order by default. Click the Name column title to view in reverse alphabetical order.
Licensed	Displays ✓ if the service is licensed. Displays ● if the service is not licensed. If one or more services are not licensed, a red banner will appear at the top of the screen that will prompt you to fix this. One or more services are not licensed. For more information, see License Details
Host	Displays the host name that the service belongs to.
Type	Displays the service type.
Version	Displays the version that the service is currently on.

Column	Description
Actions	Use drop-down list to: • Navigate to the different service views (System, Stats, Config, Explore, Logs, Security) See View Topology Details for more information. • Delete, edit, start, stop, and restart a service. See Start, Stop, or Restart a Service for more information.

Topics

- Edit Service Dialog
- Services Config View
- Services Explore View
- Services Logs View
- Services Security View
- Services Stats View
- Services System View

Edit Service Dialog

NetWitness services are automatically discovered in NetWitness.

You can use the Edit Service dialog to modify services. To access the Edit Service dialog, go to **(Admin) > Services** and click in the **Services** list toolbar.

Workflow



What do you want to do?

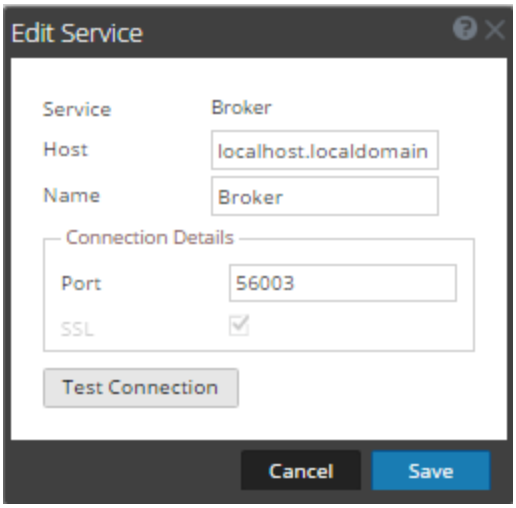
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	edit a service.*	Edit a Service

* You can perform these tasks in the current view.

Related Topics:

- Services View
- Hosts and Services Set Up Procedures
- Hosts and Services Maintenance Procedures

Quick Look



The following list describes the features of the Add Service or Edit Service dialogs.

Field or Option	Description
Service	Displays the service type. You can add the following services: Archiver, Broker, Concentrator, Network Decoder, Event Stream Analysis, Incident Management, IPDB Extractor, Log Collector, Log Decoder, Malware Analysis, Reporting Engine, Warehouse Connector, and Workbench.
Host	Specifies the host on which the service resides.
Name	Specifies the name used to identify the service, for example, **Broker**. An understandable naming convention can make administrative tasks easier. Some administrators find it convenient to use the hostname or IP address (specified in the **Host** field) for the **Name** as well.
Port	Specifies the port used to communicate with this service. The default port, based on the selected service type in the **Service** field, is autofilled here. If you select **SSL** below, this port becomes an SSL port. If you do not select **SSL**, it becomes a non-SSL port. You can customize this port by opening a firewall for the port that you add. For information about ports, see "Network Architecture and Ports" in the *Deployment Guide*.
SSL	Indicates that NetWitness uses SSL for communications with this service.
Username	Specifies the username used to log in to this service. The default username is `admin`.
Password	Specifies the password used to log in to this service. The default password is `netwitness`.
Test Connection	Tests the connection of a service that you are adding.
Cancel	Closes the Add Service or Edit Service dialog. If you do not save the service before closing the dialog, the service is not added or edited.

Field or Option	Description
Save	Adds a new service or saves changes to existing service.

Services Config View

The Services Config view is one of the views available from the **(Admin) > Services >** . It provides a user interface for configuring all aspects of a Core service or NetWitness service.

The configuration options in the Services Config view are organized as tabs, with each tab providing a view of a set of related parameters. Unlike the Services Explore view, which offers direct access to all configuration files for a service, these tabs present the most commonly modified parameters of service configuration in a user-friendly view.

Due to configuration requirements for different services; each type of service has variations in available tabs and configuration parameters in this view. Individual topics describe configuration parameters that are specific to a host (Brokers, Concentrators, Network Decoders, Log Decoders, and Archivers) or service (for example, Reporting Engine, IPDB Extractor, Log Collector, and Warehouse Connector).

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	configure an Archiver service.	See *Archiver Configuration Guide*.
Administrator	configure a Broker service.	See *Broker and Concentrator Configuration Guide*.
Administrator	configure a Concentrator service.	See *Broker and Concentrator Configuration Guide*.
Administrator	configure a Context Hub service.	See *Context Hub Configuration Guide*.
Administrator	configure an Endpoint Broker service.	See *Endpoint Configuration Guide*.

User Role	I want to...	Documentation
Administrator	configure an Endpoint Log Hybrid service.	See *Endpoint Configuration Guide*.
Administrator	configure an ESA Primary service.	See *ESA Configuration Guide*.
Administrator	configure an ESA Secondary service.	See *ESA Configuration Guide*.
Administrator	configure a Log Collector service.	See *Log Collection Configuration Guide*.
Administrator	configure a Log Decoder service.	See *Decoder Configuration Guide*
Administrator	configure a Malware Analysis service.	See *Malware Analysis User Guide*.
Administrator	configure a Network Decoder service.	See *Decoder Configuration Guide*.
Administrator	configure a Reporting Engine service.	See *Reporting Engine Configuration Guide*.
Administrator	configure a Respond service.	See *NetWitness Respond Configuration Guide*.
Administrator	configure a UEBA service.	See *NetWitness UEBA User Guide*.
Administrator	configure a Warehouse Connector service.	See *Warehouse Connector Configuration Guide*.
Administrator	configure a Workbench service.	See *Workbench Configuration Guide*.

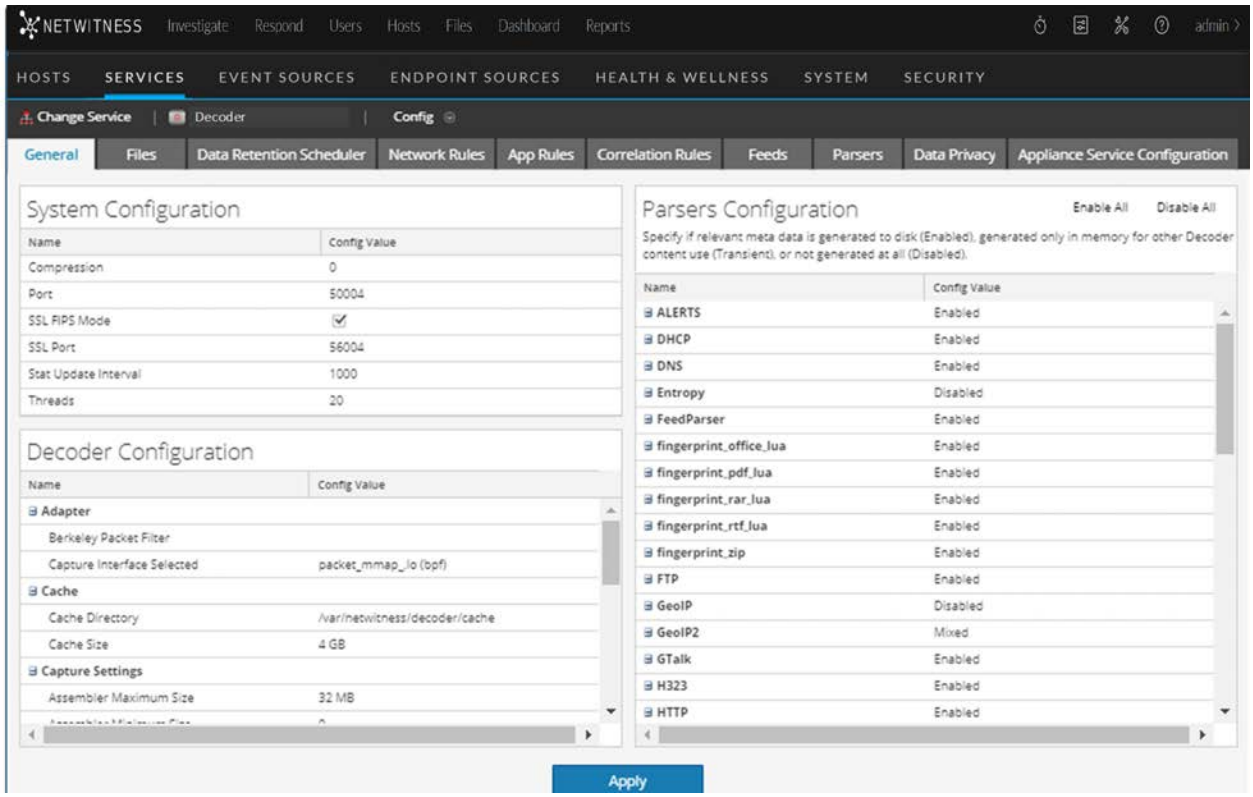
Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Related Topics

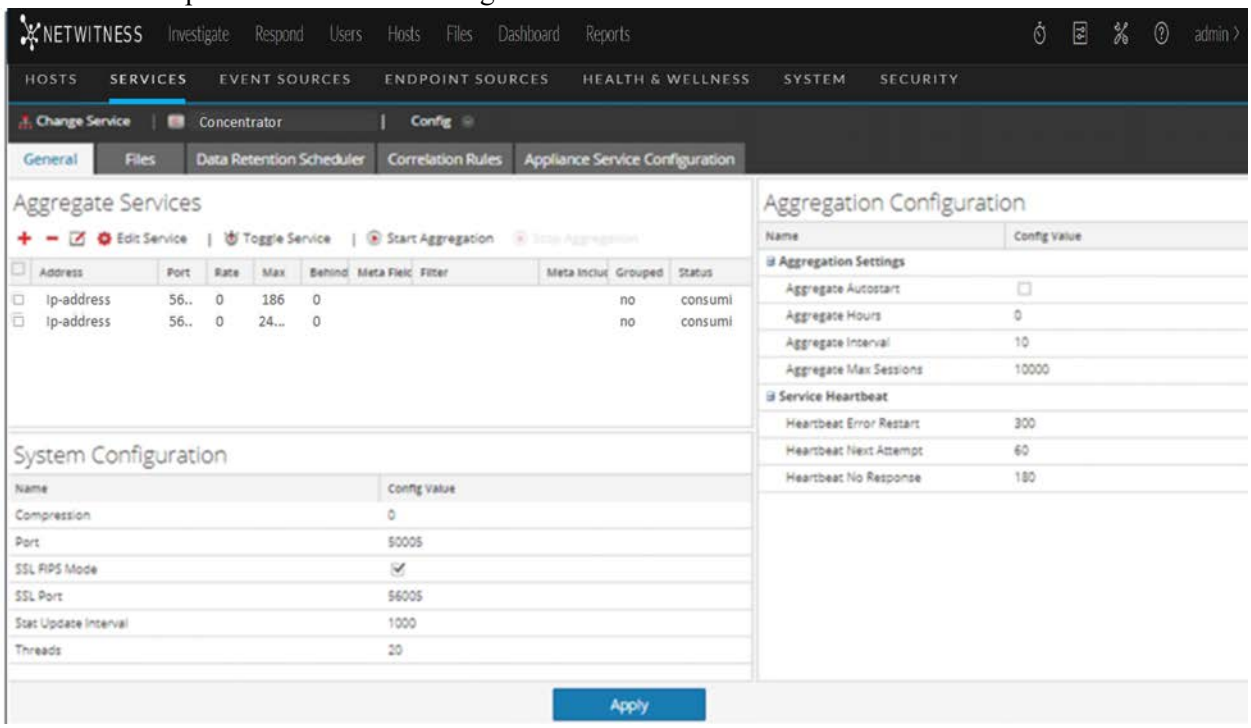
- Edit Service Dialog
- Services Explore View
- Services Logs View
- Services Security View
- Services Stats View
- Services System View
- Services View

Quick Look

This is an example of the Services Config view for a Network Decoder.



This is an example of the Services Config view for a Concentrator.



Note: Although both examples show the Services Config view, the tabs and panels are different for each. Refer to the respective guides for detailed instructions on how to configure a particular Services Config view.

Topics

- Services Config View - Appliance Service Configuration Tab
- Services Config View - Data Retention Scheduler Tab
- Services Config View - Files Tab

Services Config View - Appliance Service Configuration Tab

The Appliance Service Configuration tab appears in the Services Config view for the Archiver, Broker, Concentrator, IPDB Extractor, Network Decoder, Log Collector, and Log Decoder services.

This topic lists and describes the available configuration parameters for the NetWitness Core Appliance service. The NetWitness Core Appliance service provides hardware monitoring on legacy NetWitness hardware.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.*	Maintaining Services
Administrator	view or modify appliance service parameters.*	See "Services Config View - Appliance Service Configuration Tab" in the *Hosts and Services Getting Started Guide for Version 10.x*and prior.
Administrator	specify how long to retain database records.	See "Configure Data Retention" in the *Data Privacy Management Guide*. For information about the Data Retention tab for Archiver, see "Data Retention Tab - Archiver" in the *Archiver Configuration Guide*.

User Role	I want to...	Documentation
Administrator	edit `.xml` and `.lua` files.	Edit Core Service Configuration Files

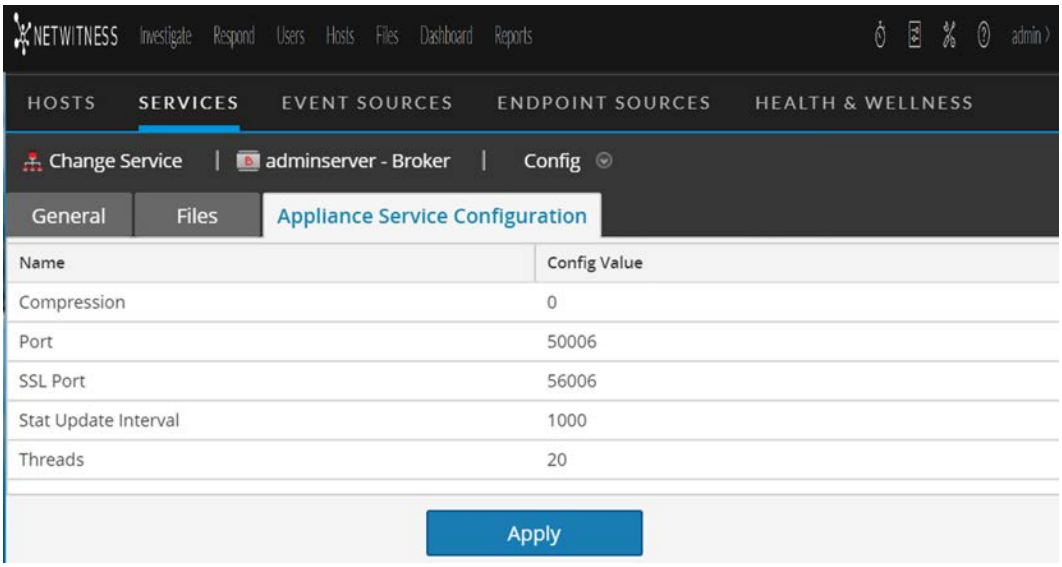
* You can perform these tasks in the current view.

Related Topics

- Services Config View - Data Retention Scheduler Tab
- Services Config View - Data Retention Scheduler Tab
- Services Config View - Files Tab
- Services Config View

Quick Look

This is an example of the Appliance Service Configuration tab for a Broker.



The following list describes the configuration values for this tab.

Name	Description of Configuration Value	When Changes Take Effect
Compression	Compresses a message when it reaches the positive number (in bytes) that you specify.	The next time you connect to this service.
Port	Unencrypted listening port. **0** indicates that the port is disabled.	Upon restart of the service.

Name	Description of Configuration Value	When Changes Take Effect
SSL FIPS Mode	One of the parameters you need to enable or disable Federal Information Processing Standards (FIPS). For detailed instructions, see "Activate or Deactivate FIPS" in the *System Maintenance Guide*.	Upon restart of the service.
SSL Port	SSL (Secure Sockets Layer) listening port. **0** indicates that the port is disabled. SSL is the standard security technology for establishing an encrypted link between a web server and a browser. This link ensures that all data passed between the web server and browsers remain private and integral.	Upon restart of the service.
Stat Update Interval	How often (in milliseconds) the system updates statistic nodes for monitoring Health and Wellness.	Immediately.
Threads	Threads in thread pool required to used to handle requests. The `Threads` parameter works with the `Polling Interval` parameter for event and log threads.	Immediately.

Services Config View - Data Retention Scheduler Tab

The Data Retention Scheduler tab appears in the Services Config view for the Network Decoder, Log Decoder, and Concentrator.

In the Data Retention Scheduler tab, you can define the criteria for removing database records from primary storage on Network Decoder, Log Decoder, and Concentrator services, and schedule the timing for checking the threshold.

Note: If additional customization is necessary, use the Scheduler under the Files tab in the Services Config view. For example, if you have storage available to save the RAW data versus the metadata, use `Capacity` as the threshold and to set different thresholds per database (metadata versus packet).

Workflow



What do you want to do?

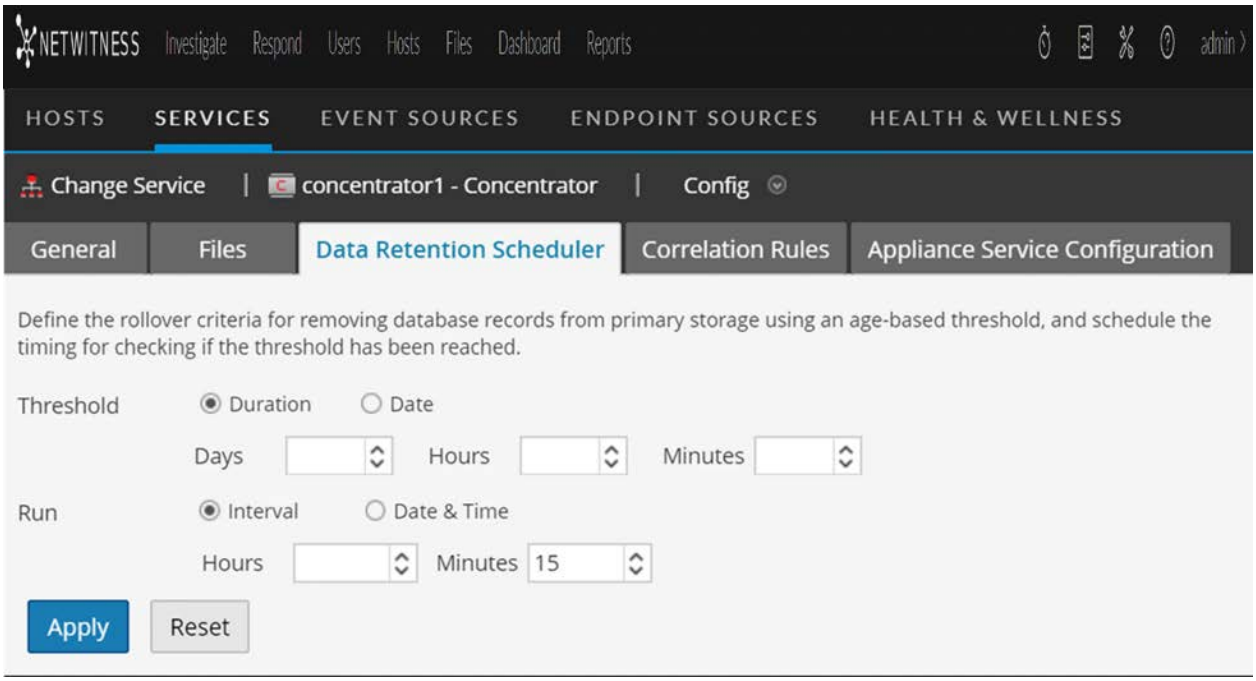
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.*	Maintaining Services
Administrator	view or modify appliance service parameters.	See "Services Config View - Appliance Service Configuration Tab" in the *Hosts and Services Getting Started Guide for Version 10.x* and prior.
Administrator	specify how long to retain database records.*	See "Configure Data Retention" in the *Data Privacy Management Guide*. For information about the Data Retention tab for Archiver, see "Data Retention Tab - Archiver" in the *Archiver Configuration Guide*.
Administrator	edit `.xml` and `.lua` files.	Edit Core Service Configuration Files

Related Topics

- Services Config View - Appliance Service Configuration Tab
- Services Config View - Files Tab
- Services Config View

Quick Look

The following figure illustrates the parameters in the Data Retention Scheduler tab for a Concentrator.



The Data Retention Scheduler tab has sections to specify Threshold settings and Run settings. The following table lists the parameters supported for data retention configuration.

Parameter	Description
Threshold	The threshold is based on the age of the data, the amount of time the data was stored or the date on which the data was stored. The date is from the database file, not from the actual session time. • **Duration**: The duration of time that data can be stored before removal. Specifies the number of days (365 maximum), hours (24 maximum), and minutes (60 maximum) that have elapsed since the time stamp on the data. • **Date**: The removal of data based on the date of the timestamp. Specifies the monthly date and time in the and fields.
Run	The schedule for running the job that checks rollover criteria. • **Interval**: Schedule the database check to occur at a regular interval. Specifies the hours and minutes between the scheduled checks. • **Date and Time**: Schedule the database check to occur at a regular day and time. Specifies the day from the drop-down list and the system clock time in `hh:mm:ss` format. Possible values for day are `Everyday`, `Weekdays`, `Weekends`, and `Custom`, where `Custom` allows you to select one or more specific days of the week.

Parameter	Description
Apply	Overwrites any previous schedule for this service and applies the new settings immediately. **Caution:** After you apply these settings, when the threshold is met the system deletes the old data from the database and you can no longer access it.
Reset	Resets the schedule to the last applied state.

Services Config View - Files Tab

The Files tab appears in the Services Config view for Archivers, Brokers, Concentrators, Log Decoders, and Network Decoders.

In the Files tab, you can edit service configuration files as text files. The files you can edit vary depending upon the type of service you are configuring. The following files are common to all Core services.

- The NetWitness file (`netwitness`)
- The service index file (`index-<service>`)
- The scheduler file (`scheduler`)
- The crash reporter file (`crashreporter`)
- The feed definitions file (`feed-definitions`)

In addition, the Decoder has files that configure parsers, feed definitions, and a wireless LAN adapter.

Note: The default values in the configuration files cover most common situations. You may need to edit configuration parameters and values for optional services, such as the crash reporter or scheduler. Do not change these values in the Files tab unless you understand networks and the factors that affect the way services collect and parse data.

More detail on the service configuration parameters is available in the Service Configuration Settings.

Workflow



What do you want to do?

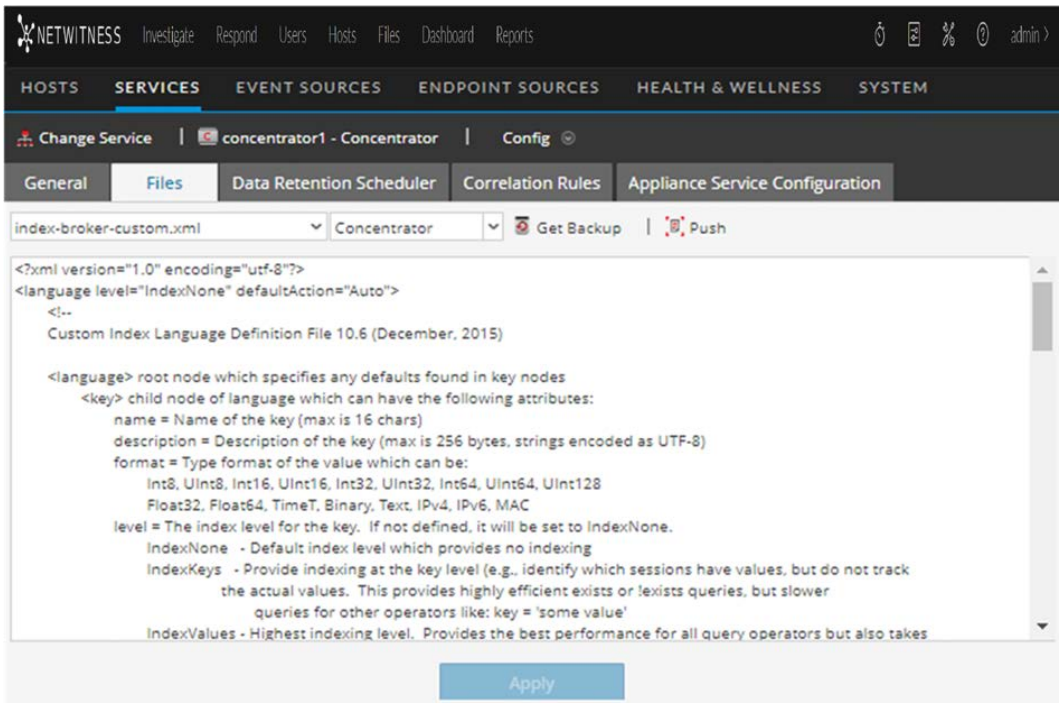
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.*	Maintaining Services
Administrator	view or modify appliance service parameters.	See "Services Config View - Appliance Service Configuration Tab" in the *Hosts and Services Getting Started Guide for Version 10.x*and prior.
Administrator	specify how long to retain database records.	See "Configure Data Retention" in the *Data Privacy Management Guide*. For information about the Data Retention tab for Archiver, see "Data Retention Tab - Archiver" in the *Archiver Configuration Guide*.
Administrator	edit `.xml` and `.lua` files.*	Edit Core Service Configuration Files

Related Topics

- Services Config View - Data Retention Scheduler Tab
- Services Config View - Data Retention Scheduler Tab
- Services Config View

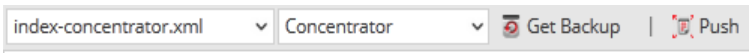
Quick Look

This is an example of the Files tab.



Files Tab Toolbar

The Files tab has a toolbar and an edit window. This is an example of the toolbar.



These are the features of the Files tab toolbar.

Feature	Description
File drop-down list	Displays a list of files that the system is currently using. When you select a file, the text of the file is displayed in the text edit window. In the text window, you can edit the file and save the changes, or create alternate files to use.
Service / Host drop-down list	Displays the service type and host. You can open a file from either the service or the host for editing.
Get Backup	Retrieves the latest backup of the current file, which can prove useful when you have made changes and want to go back to the previous version of the file. The backup does not replace the current file unless you click **Save**.
Push	Displays a dialog in which you can select services of the same type and push the currently viewed file to the services.
Apply	Overwrites the current file and creates a backup file.

Services Explore View

You can use the NetWitness Services Explore view (**(Admin) > Services**, **> View > Explore**) to display and edit both host and service configurations.

The Services Explore view offers advanced access and control of all NetWitness hosts and services. All services expose their functionality through a treelike series of nodes, similar to the Windows Explorer view of a Windows file system. Here you can:

- View a directory tree showing common files for all selected services.
- Navigate down through the directory to a file.
- Open the same file for each service, and display the contents side by side.
- Select an entry in the file and edit the value.
- Apply a property value from one service to other services.

The Services Explore view can also display a Properties dialog, a simple interface for viewing properties of any node in the system and sending messages to the node.

Caution: A good understanding of the nodes and parameters is required when editing in this view. Incorrect settings can cause performance problems.

Workflow



What do you want to do?

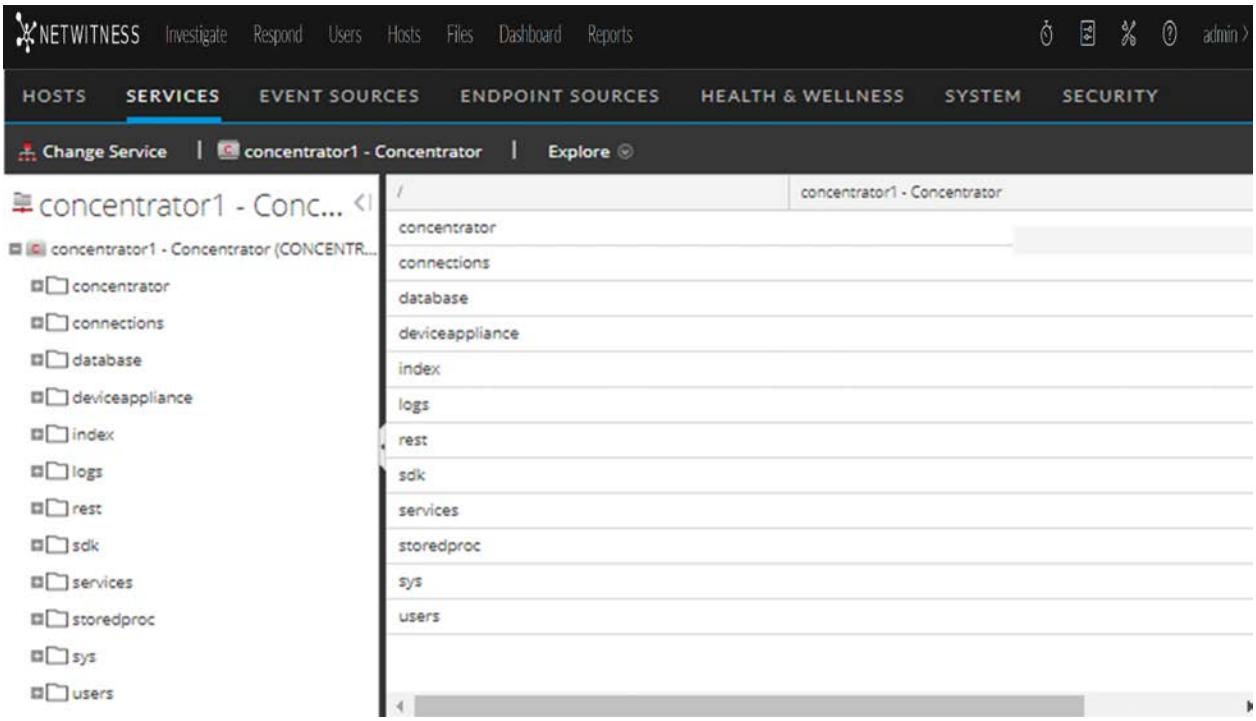
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	view and edit host and service configurations.*	Explore and Edit Service Property Tree

* You can perform these tasks in the current view.

Related Topics

- View Topology Details
- Edit Service Dialog
- Services Config View
- Services Logs View
- Services Security View
- Services Stats View
- Services System View
- Services View

Quick Look



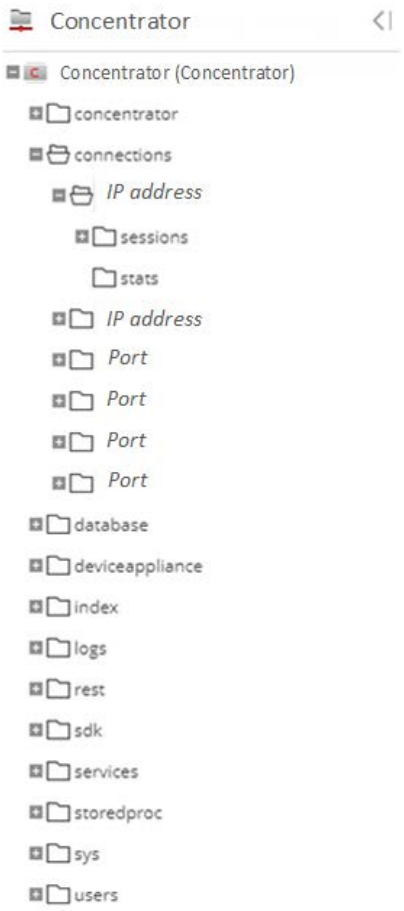
The Services Explore View has two main panels:

- The Node list
- The Monitor panel

You can right-click on a file to access the Properties for the file. See Services Explore View - Properties Dialog for more information.

The Node List

The Node list displays the services as a tree-like series of nodes and folders. The levels in the Node list expand and collapse to display the full hierarchy.



Each root folder is named based on the functionality it exposes. For instance, the `/connections` folder shows all connected IP addresses. Underneath each IP address or port in the list are two folders, `sessions` and `stats`.

- The `sessions` folder displays all authenticated user sessions originating from the IP/Port.
- The `stats` folder displays values, such as the number of messages sent or received, bytes sent or received, and other values set by the service. These are not editable.

Selecting any folder in the tree view displays its children in the Monitor panel. Every node in the tree is actively monitored, so when a statistic or configuration node changes value, it is immediately reflected in the tree and Monitor panel.

The Monitor Panel

The Monitor panel displays properties and values for a selected node (such as `index`) and a child folder (such as `config`). There are two ways to edit values:

- Click the value and type a new value
- Send a `set` message in the Properties dialog

/index/config	Concentrator
index.dir	/var/netwitness/concentrator/index=2.89 TB
index.dir.cold	
index.dir.warm	
index.slices.open	42
page.compression	huffhybrid
reindex.enable	true
save.session.count	auto

Topics

- View Topology Details

Services Explore View - Properties Dialog

You can use the Services Explore view > Properties dialog to perform the following tasks.

- Send messages to a system node
- Retrieve values for a property for multiple services
- Set values for a property for multiple services

Right-click any file in the Node list or Monitor panel to display the Properties context menu.



When you select `Properties` from the Node list or Monitor panel context menu, the Properties dialog opens below the Monitor panel. All nodes have support help that contains the following information.

- A description of the node
- The list of supported messages with a corresponding description
- Security roles needed to access the messages

The available messages vary according to the service and root folder. Many of these messages are also accessible as options with a NetWitness dashboard or view.

Workflow



What do you want to do?

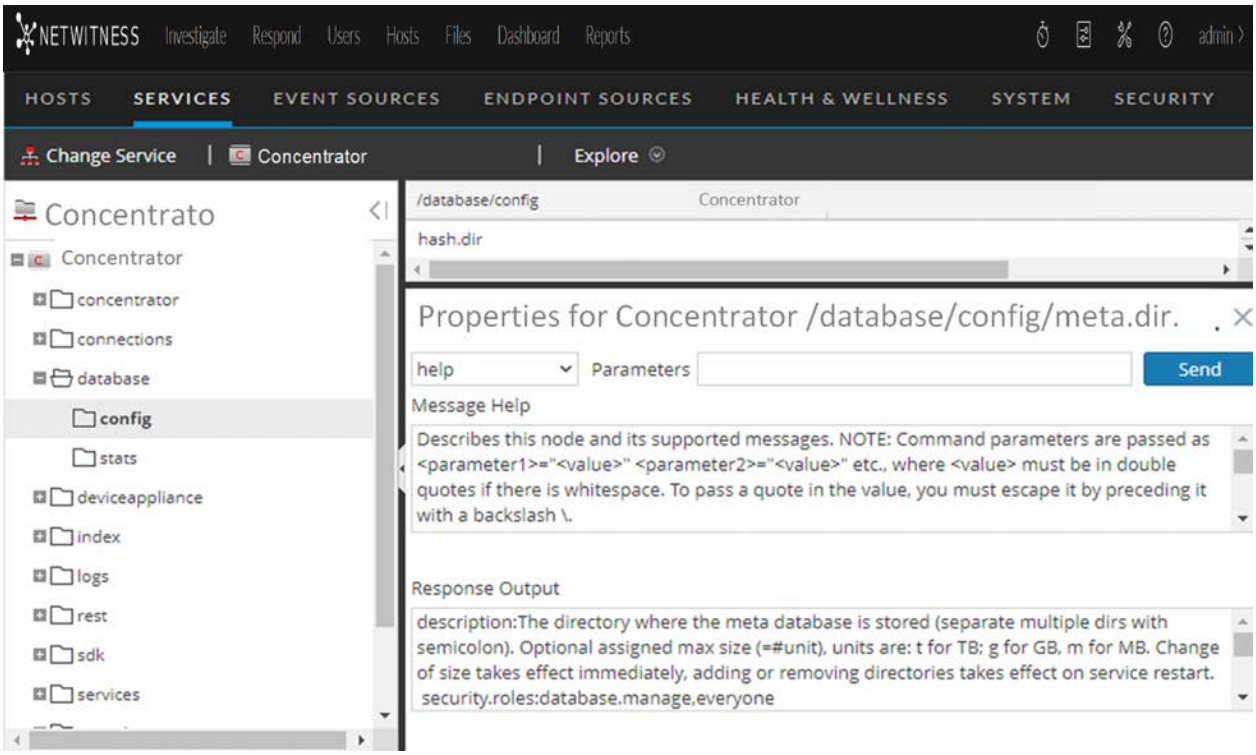
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	send messages to a system node.*	Explore and Edit Service Property Tree
Administrator	retrieve values for a property for multiple services.*	Explore and Edit Service Property Tree
Administrator	set values for a property for multiple services.*	Explore and Edit Service Property Tree

* You can perform these tasks in the current view.

Related Topics

- Services Explore View

The following example shows the Properties dialog with information in Message Help displayed.



The Properties dialog has the following features.

Feature	Description
Message drop-down list	Lists all available messages for the current node. Select a message from this drop-down list to send to the node.
Parameters input field	Type the message parameters in this field.
Send	Sends the message to the selected node.
Message Help	Displays help text for the current message.
Response Output	Displays the response to a message or output from a message.

Services Logs View

The Services Logs view provides the ability to view and search the logs for a specific service. The Services Logs view is identical to the System Logging panel (**(Admin) > System tab > System Logging**) with two exceptions:

- The Services Logs view has an additional filter to select messages for the service or host.
- The System Logging panel has an additional tab for Settings.

For a complete description of NetWitness logging features in the System Logging panel, see "Monitor Health and Wellness of NetWitness Platform" in the *System Maintenance Guide*.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	view and search the logs for a specific service.*	See "Monitor Health and Wellness of NetWitness Platform" in the *System Maintenance Guide*

* You can perform these tasks in the current view.

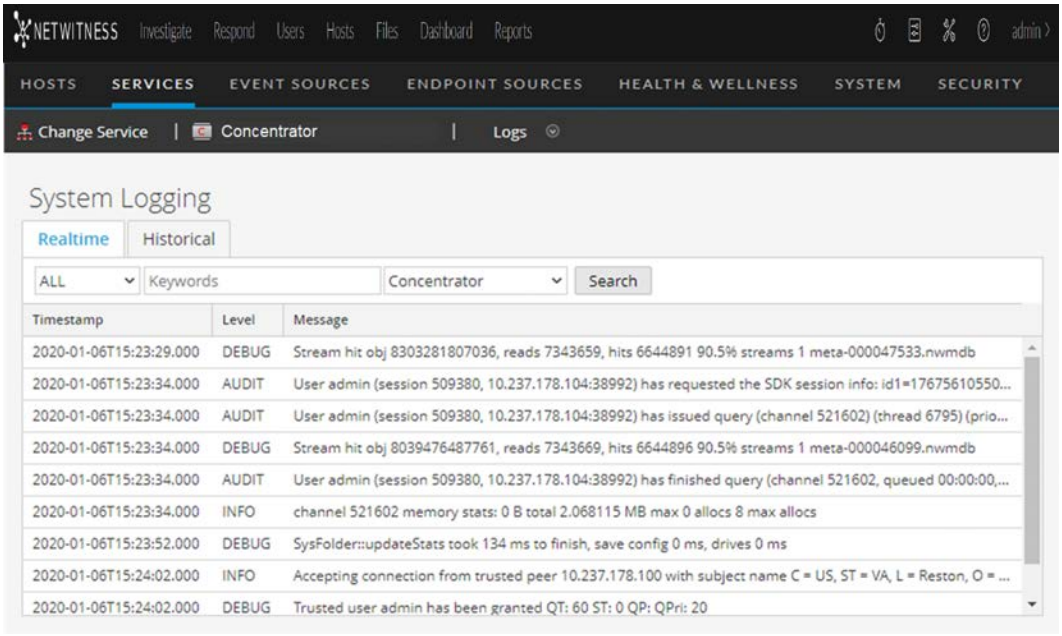
Related Topics

- Edit Service Dialog
- Services Config View
- Services Explore View
- Services Security View
- Services Stats View

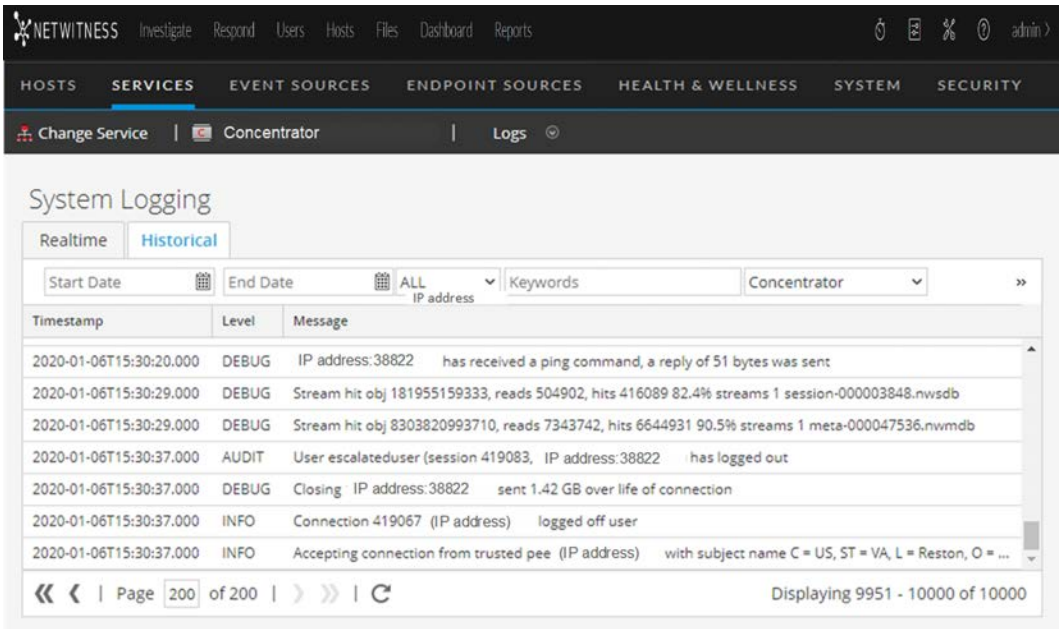
- Services System View
- Services View

Quick Look

The following figure shows the Services Logs view Realtime tab.



The following figure shows the Services Logs view Historical tab.



Feature	Description
Realtime tab	This is the monitor mode of the service log. For more information, see "System Logging - Realtime Tab" in the *System Maintenance Guide*.
Historical tab	This is a searchable view of the service log. For more information, see "System Logging - Historical Tab" in the *System Maintenance Guide*.

Services Security View

In NetWitness, each service has a separate configuration of users, roles, and role permissions, which are managed in the Services Security view.

To access service information and perform service operations through NetWitness, a user must belong to a role that has permissions on that service. For NetWitness version 10.4 or later Core services that utilize trusted connections, it is no longer necessary to create NetWitness Core user accounts for users that log on through the web client. You only need to create NetWitness Core user accounts for aggregation, thick client users, and REST API users.

Note: Only the default admin user in NetWitness is created by default on all services. As a prerequisite to managing service security, the default admin user account must be present in the Admin Services view. For every other user, you must configure access to each particular service through NetWitness.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	manage configuration of users, roles, and role permissions.*	See the *System Security and User Management Guide*.

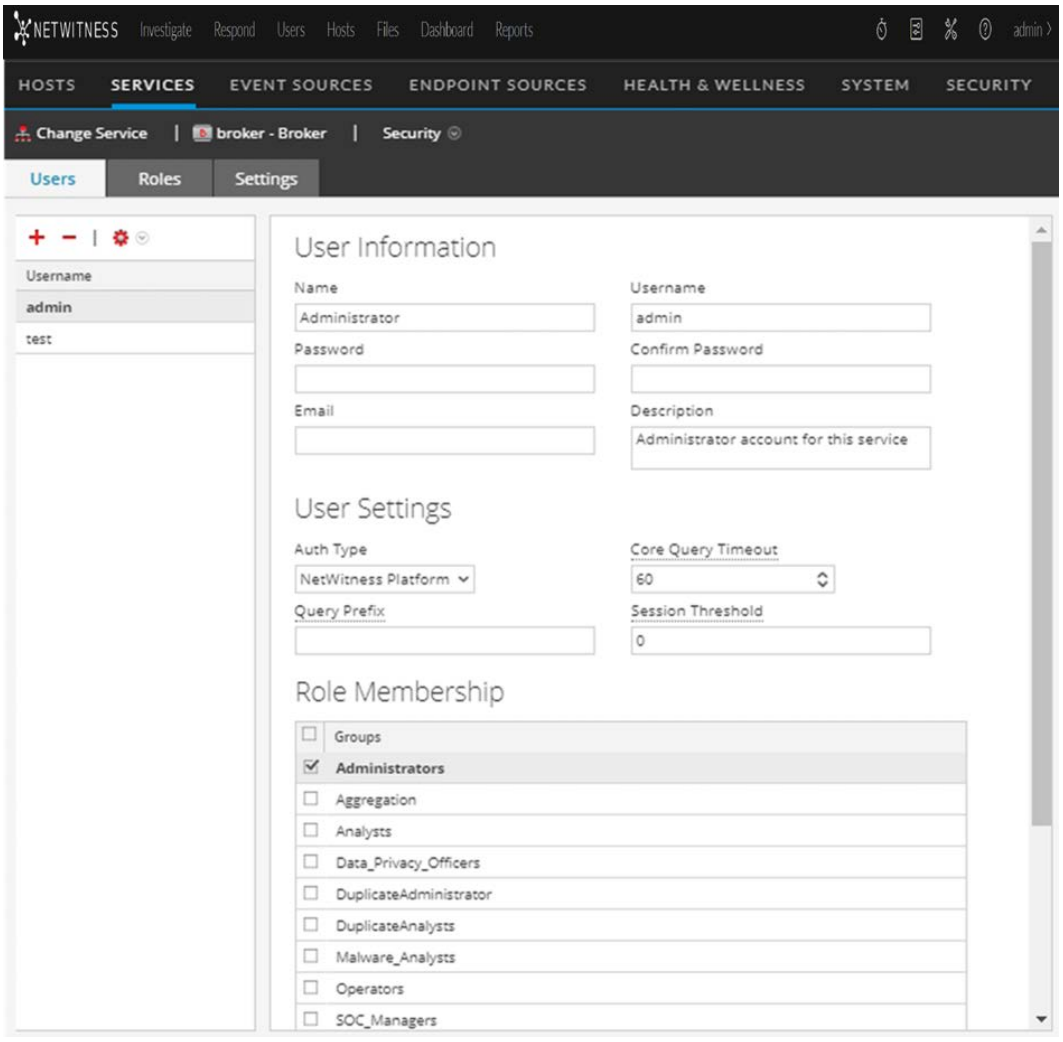
* You can perform these tasks in the current view.

Related Topics

- Add, Replicate, or Delete a Service User
- Add a User Role to a Service
- Change a Service User Password
- Duplicate or Replicate a Service Role

- Edit Service Dialog
- Services Config View
- Services Explore View
- Services Logs View
- Services Stats View
- Services System View
- Services View

Quick Look



The Services Security view has three tabs: Users, Roles, and Settings.

Roles and Service Access

Primary considerations in configuring service security are defining the roles and assigning users to the roles. The Service Security view separates these two functions into the Users tab and the Roles tab.

- In the Users tab, you can add a user, edit user settings, change the user password, and edit the role membership of the user for a selected service. Although you select a single service in the Services Security view, you can apply the settings for one service to other services.
- In the Roles tab, you can create roles and assign permissions to the roles for a selected service.

Topics

- Services Security View - Users Tab
- Services Security View - Roles Tab
- Services Security View - Settings Tab

Services Security View - Users Tab

In the Services Security view Users tab, you can configure the following for a service:

- Add user accounts.
- Change service user passwords.
- Configure user authentication properties and query handling properties for the service.
- Specify the user role membership, which specifies the roles that the user belongs to on the selected service.

Note: For version 10.4 or later NetWitness Core services that utilize trusted connections, it is no longer necessary to create NetWitness Core user accounts for users that log on through the web client. You only need to create NetWitness Core user accounts for aggregation, thick client users, and REST API users.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host

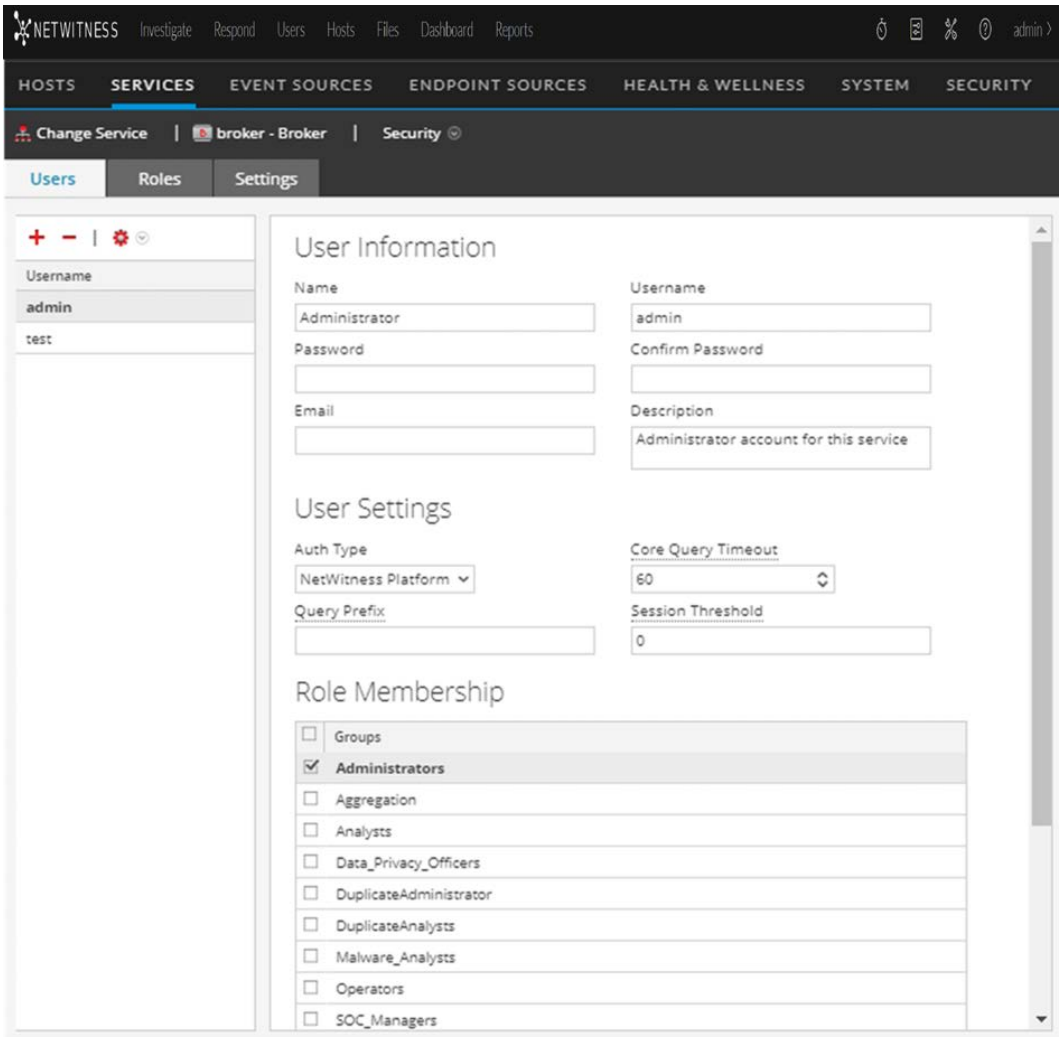
User Role	I want to...	Documentation
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	add user accounts.*	Add, Replicate, or Delete a Service User
Administrator	change service user passwords.*	Change a Service User Password
Administrator	configure user authentication properties and query handling properties for the service.*	See "Verify Query and Session Attributes per Role" in the *System Security and User Management Guide*.
Administrator	specify the user role membership (roles that the user belongs to on the selected service).*	See "Add a User and Assign a Role" in the *System Security and User Management Guide*.

* You can perform these tasks in the current view.

Related Topics

- Services Security View - Roles Tab
- Services Security View - Settings Tab
- Services Security View

Quick Look



The Users tab has a User List panel on the left. Selecting a username from the panel makes the User Definition panel on the right available.

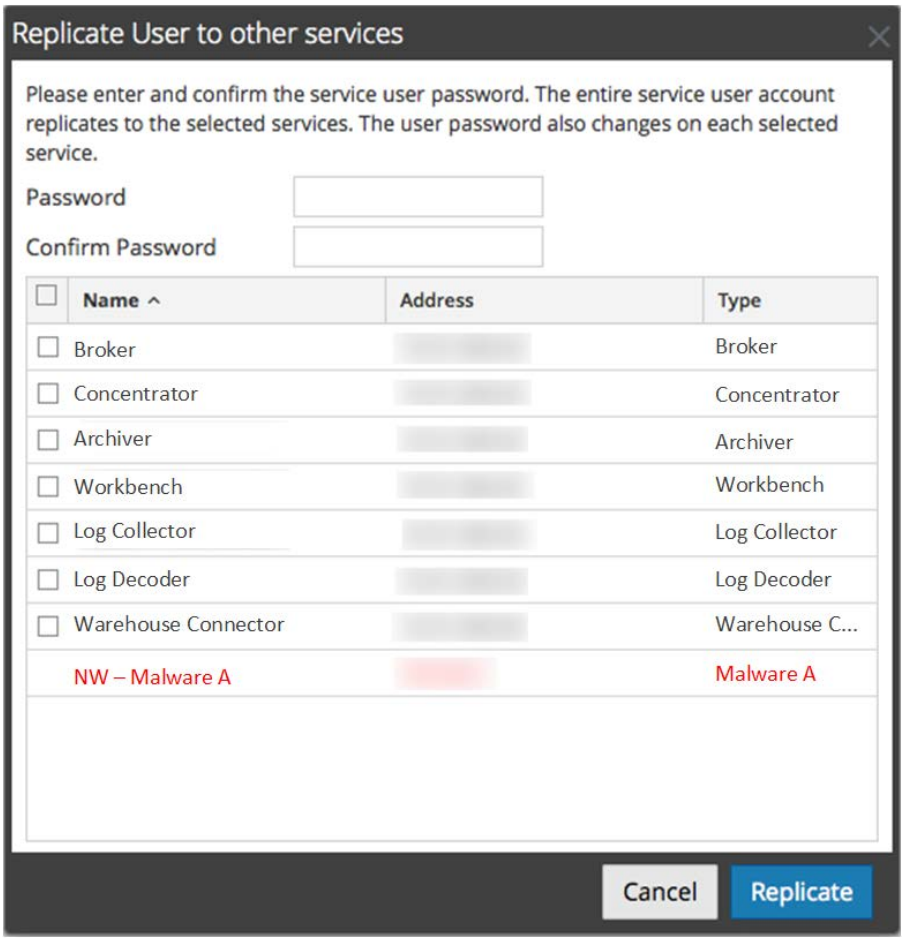
User List Panel

The User List panel has the following features.

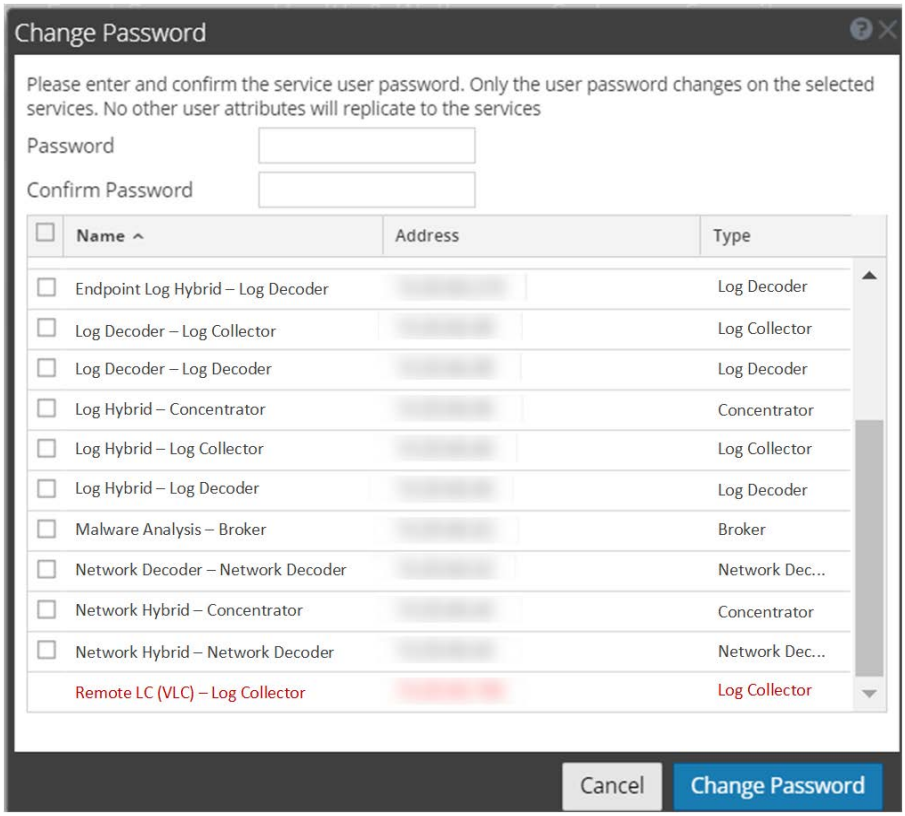
Feature	Description
+	Adds a new user to the current service.
−	Deletes the selected users from the service.

Feature	Description
	Performs one of the following actions on the selected service user account: • **Replicate:** Replicates the entire service user account to selected services. • **Change Password:** Changes the password of a service user and replicates the new password to Core services with that user account defined. The Change Password option replicates only the password change to the Core services selected and does not replicate the entire user account.
Username	The usernames for all user accounts that access the service. The username must be one used to log on to NetWitness.

The following figure shows the "Replicate User to other services" dialog.



The following figure shows the **Change Password** dialog.



User Definition Panel

The User Definition panel has three sections:

- User Information identifies the user as created in the Admin Services Security view.
- User Settings define parameters that apply to this user's access to the service.
- Role Membership defines user roles to which the user belongs.

There are two buttons at the bottom of the panel:

- The **Apply** button saves the changes made in the User Definition panel, and they become effective immediately.
- If you have not saved changes in the User Definition panel, the **Reset** button resets all fields and settings to their values before editing.

User Information

The User Information section has the following features.

Field	Description
Name	The name of the user.
Username	The username that this user enters to log in to the service. This is the NetWitness username generated when the administrator added the user and the associated credentials in the Admin Services Security view.

Field	Description
Password (and **Confirm Password**)	The password that the user enters to log on to the service. This is the NetWitness password generated when the administrator added the user and the associated credentials in the **Administration Security** view. The NetWitness account password and the service password must match in order to allow the user to connect to the service through NetWitness.
Email	(Optional) The user's email address.
Description	(Optional) A general description field to describe this user.

User Settings

The User Settings section has the following features.

Field	Description
Auth Type	The authentication scheme for this user. The product line supports internal and external authentication. • **NetWitness Platform** specifies internal authentication, and is enabled by default. In this mode, all users must authenticate with the user account and passwords that are generated when the administrator uses the NetWitness Admin Services Security view to create the user and their associated credentials. • **External** specifies that authentication is enabled through the host interface with PAM (Pluggable Authentication Modules). For more information, see "Configure PAM Login Capability" in the *System Security and User Management Guide*.
Core Query Timeout	**Note:** This field was previously known as "SA Core Query Timeout" and does not appear for 10.4 and earlier service versions. NetWitness version 10.4 and earlier services use "Query Level" instead. Specifies the maximum number of minutes a user can run a query on the service. If this value is set to **0**, the query timeout is not enforced for the user on the service.
Query Prefix	(Optional) Restricts query results seen by the user by appending the query syntax to every query. For example, adding the query prefix `email != 'ceo@company.com'` prevents those email results from showing up in the sessions.
Session Threshold	(Optional) Controls the behavior of the application when scanning meta values to determine session counts. If any meta value has a session count that is above the set threshold, the determination of the true session count stops when the threshold is reached. If a threshold is set for a session, the Navigate view (INVESTIGATE > Navigate) shows that the threshold was reached and the percentage of query time used to reach the threshold.

Role Membership

The Role Membership section shows a list of all roles. The checkbox next to a role is selected for the roles that a user is a member of for the selected service.

Services Security View - Roles Tab

The Services Security view Roles tab enables you to create roles and assign permissions. Each role can have different permissions for different services. For example, the Analysts role can have different role permissions based on the selected service.

Before you can add users to roles, you need to define user roles, usually by function, and assign permissions to the roles.

Procedures related to this tab are described in Add a User Role to a Service.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	define user roles.*	Add a User Role to a Service
Administrator	assign permissions to the roles.*	See "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.
Administrator	add a user role to a service.*	See "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.

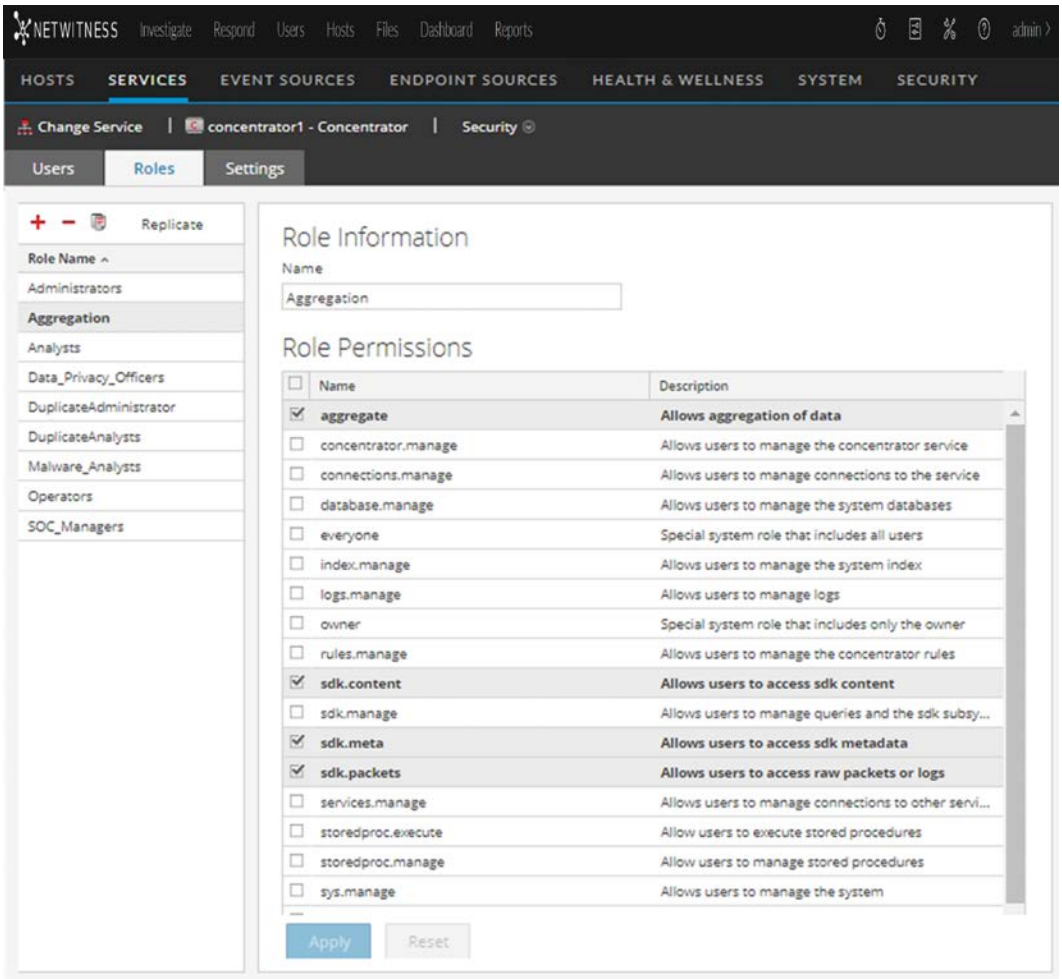
* You can perform these tasks in the current view.

Related Topics

- Duplicate or Replicate a Service Role
- Duplicate or Replicate a Service Role
- Services Security View - Users Tab

- Services Security View - Settings Tab
- Services View

Quick Look



The Roles tab has a **Role Name** panel on the left. Selecting a role name shows the **Role Information** panel for the selected role on the right.

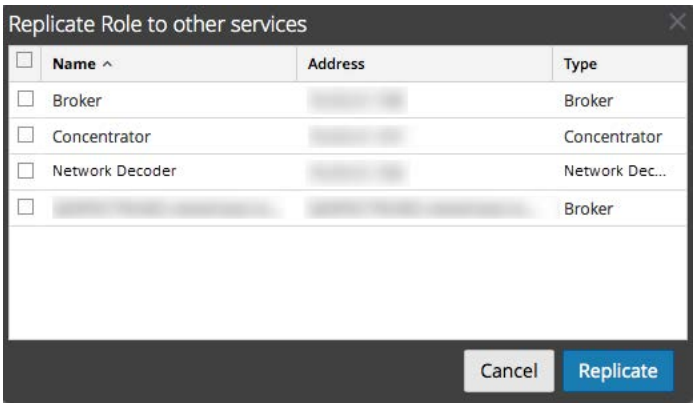
Role Name Panel

The **Role Name** panel has the following features.

Feature	Description
	Adds a new group to the current service.
	Deletes the selected group from the current service.
	Copies a role and its assigned permissions to a new role. The name of the new role must be unique. For example, you can copy the `Analysts` role and create another role with a new name, such as `Analyst_Managers`.

Feature	Description
Replicate	Pushes a role and its assigned permissions to other services. After you select a role and click **Replicate**, the Replicate Role to other services dialog is displayed. In the dialog, you can select the services where you want to replicate the role.

The following figure shows the Replicate Role to other services dialog.



Role Information and Permissions Panel

The Role Information and Permissions panel defines role permissions.

There are two buttons:

- The **Apply** button saves the changes made in the Role Permissions panel and they become effective immediately.
- If you have not saved changes in the Role Permissions panel, the **Reset** button resets all fields and settings to their values before editing.

Topics

- Services Security View - Service User Roles and Permissions
- Services Security View - Aggregation Role

Services Security View - Service User Roles and Permissions

The Services Security view Roles tab enables you to create service user roles and assign permissions. You can also use the pre-configured service user roles included with NetWitness to assign user permissions.

Related Topics

- Services Security View - Aggregation Role
- Services Security View - Roles Tab

Service User Roles

NetWitness has the following pre-configured service user roles.

Role	Assigned Permissions	Personnel/Account
Administrators	All permissions	NetWitness System Administrator

Aggregation	`aggregate`
	`sdk.content`
	`sdk.meta`
	`sdk.packets`
		You can use this role to create an Aggregation account.
		This role provides the minimum permissions necessary to perform aggregation of data. It is only available on NetWitness version 10.5 and later services.

Analysts, Malware_ Analysts, and SOC_ Managers	`sdk.meta`
	`sdk.content`
	`sdk.packets`
	`storedproc.execute`
		Users can use specific applications, run queries and view content for purposes of analysis.

Data_Privacy_ Officers	`sys.manage`
	`users.manage`
	`sdk.meta`
	`sdk.content`
	`sdk.packets`
	`sdk.manage`
	`logs.manage`
	`database.manage`
	`index.manage`
	`dpo.manage`
		Data Privacy Officer
		Data Privacy Officers have the `dpo.manage` permission on Network Decoders and Log Decoders.

Operators	`sys.manage`
	`services.manage`
	`connections.manage`
	`users.manage`
	`logs.manage`
	`parsers.manage`
	`rules.manage`
	`database.manage`
	`index.manage`
	`sdk.manage`
	`decoder.manage`
	`archiver.manage`
	`concentrator.manage`
	`storedproc.manage`
		Operators are responsible for the daily operation of the services.

Service User Permissions

There are many permissions that you can assign a service role in NetWitness. Users can have different permissions on each service, depending on their role assignments and the permissions selected for each role. This table describes the permissions that you can assign to a role.

Permission	Definition
sys.manage	Allows the user to edit the service configuration settings.
services.manage	Allows the user to manage connections to other services.
connections.manage	Allows the user to manage connections to the service.
users.manage	Allows the user to create individual users and user roles and specify user permissions.
aggregate	Allows the user to perform aggregation of data.
sdk.meta	Allows the user to run queries in the Investigation and Reporting applications and to view the metadata returned by the query.
sdk.content	Allows the user to access raw packets and logs from any client application (Investigations and Reporting).
sdk.packets	Allows users to access raw packets and logs from any client application.
appliance.manage	Allows the user to manage the appliance (host) tasks. This permission is required by the Appliance service.
decoder. manage	Allows the user to edit the configuration settings for the Network Decoder service.
concentrator.manage	Allows the user to edit the configuration settings for the Concentrator/Broker service.
logs.manage	Allows the user to view the service logs and edit the logging configuration settings for the specified service.
parsers.manage	Allows the user to manage all attributes under the parsers node.
rules.manage	Allows the user to add and delete all rules.
database.manage	Allows the user to set database locations, sizes, and the various configuration settings for the session, meta and/or packet/log databases.
index.manage	Allows the user to manage all index-related attributes.
sdk.manage	Allows the user to view and set all SDK configuration items.
storedproc.execute	Allows the user to execute a Lua stored procedure.
storedproc.manage	Allows the user to manage Lua stored procedures.
archiver.manage	Allows the user to modify the Archiver configuration.

Permission	Definition
dpo.manage	Allows the user to manage the transform configuration and the applicable keys.

Services Security View - Aggregation Role

This topic describes the Aggregation role and permissions that allow service users to perform aggregation.

Related Topics

- Services Security View - Service User Roles and Permissions
- Services Security View - Roles Tab

The Aggregation role is a service user role intended only for aggregation of data. It has the minimum role permissions required to do aggregation:

- aggregate
- sdk.meta
- sdk.packets
- sdk.content

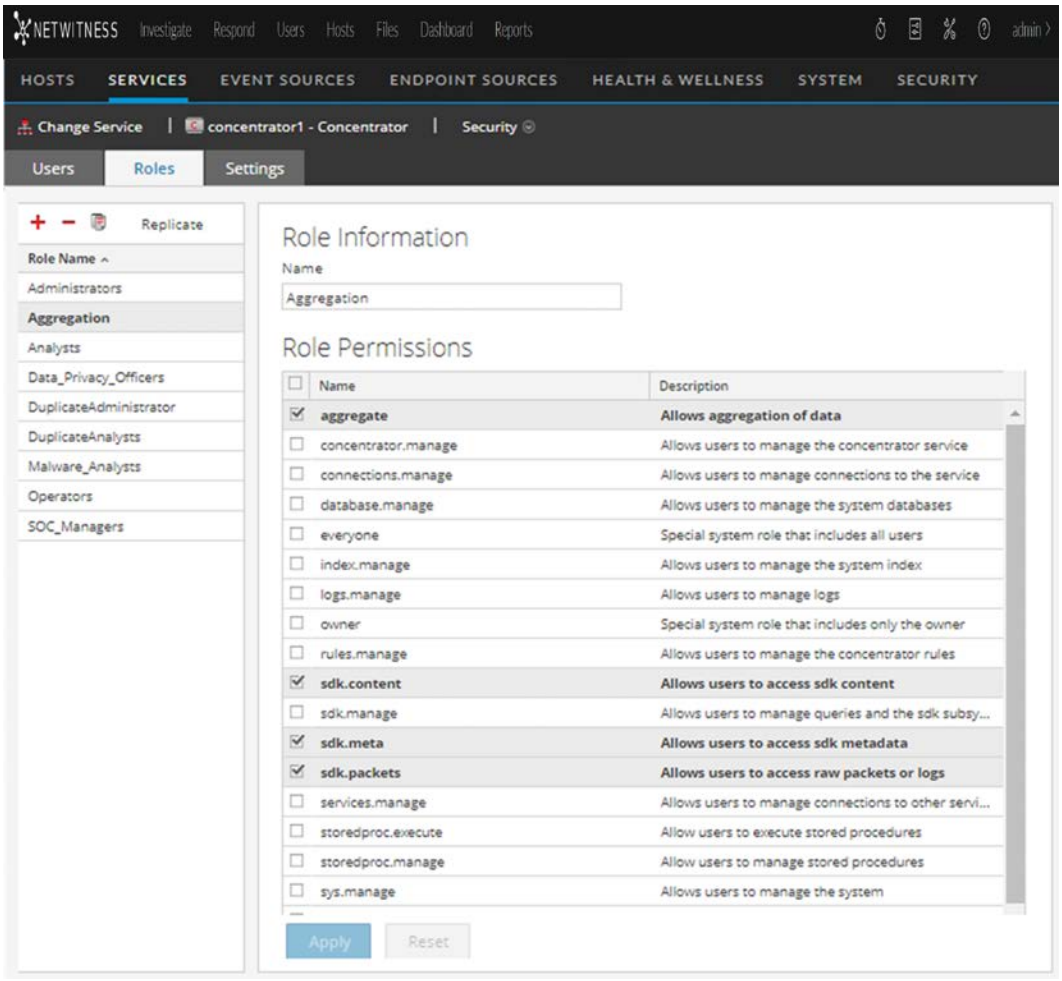
The Aggregation role is available only on NetWitness version 10.5 and later services and it can be used for an aggregation account. Members of this role or service users with these permissions can perform aggregation on Network Decoders, Concentrators, Archivers, and Brokers. The aggregate permission allows service users to perform aggregation of sessions and metadata along with raw packets and logs.

You can still use the decoder.manage, concentrator.manage, and archiver.manage permissions, but the Aggregation role permissions allow aggregation only and prevent the other available operations.

You access the service roles from the **(Admin) > Services** (select a service) > > **View** > **Security** > **Roles** tab.

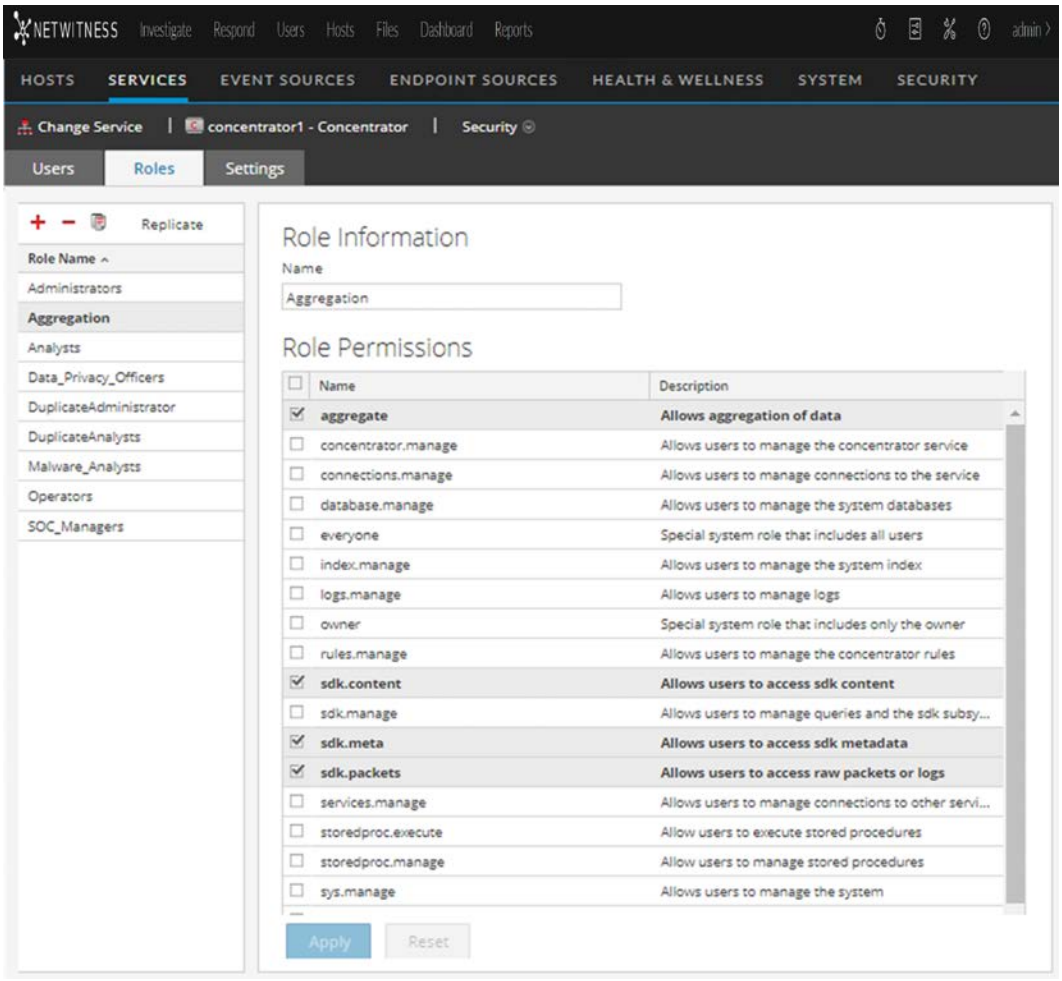
Procedures related to roles are described in Hosts and Services Maintenance Procedures. Services Security View - Service User Roles and Permissions provides detailed information on the pre-configured roles.

The following figure shows the permissions in the Aggregation role.



Services Security View - Settings Tab

In the Services Security view Settings tab, Administrators can enable and configure system roles that define permissions on a per meta key basis for individual Brokers, Concentrators, Network Decoders, and Log Decoders. Configuring this feature adds configurable meta keys to the Services Security view > Roles tab so that individual meta keys can be applied to specific roles on a specific service. The following figure illustrates this.



This configuration is generally part of a data privacy plan implemented to ensure that specific types of content consumed or aggregated by a service are kept secure by limiting visibility of the metadata and content to privileged users. See the *Data Privacy Management Guide* for more information.

Workflow



What do you want to do?

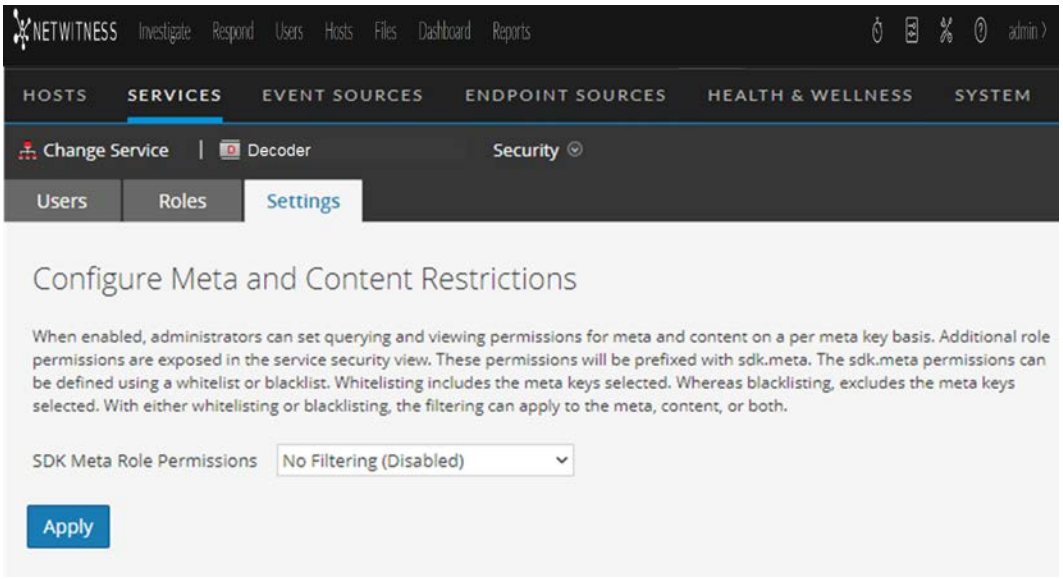
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	enable and configure system roles that define permissions on a per meta key basis for individual Brokers, Concentrators, Network Decoders, and Log Decoders.*	See the *System Security and User Management Guide* for more information.

* You can perform these tasks in the current view.

Related Topics

- Services Security View - Users Tab
- Services Security View - Roles Tab
- Services Security View

Quick Look



The Settings tab includes two features.

Feature	Description
SDK Meta Role Permissions field	Provides option for disabling or configuring meta key and content restrictions. The filtering options are described.
Apply button	Applies the selected configuration immediately. If not disabled, the meta keys are added to the Roles tab so they can be applied to specific roles.

SDK Meta Role Permissions Options

The following table lists the filtering options available in the SDK Meta Role Permissions selection list, and the numeric values used to disable (0) and the types of filtering (1 through 6).

Note: There is no need to know the numeric value unless configuring metadata and content visibility manually in the `system.roles` node.

`system.roles` Node Value	Settings Tab Option	Description
0	No Filtering (Disabled)	System roles that define permissions on a per meta key basis are disabled.
1	Whitelist metadata and content	Metadata and content for the specified SDK meta roles are white listed, or visible to users assigned the system role.
2	Whitelist only metadata	Metadata for the specified SDK meta roles is white listed, or visible to users assigned the system role.
3	Whitelist only content	Content for the specified SDK meta roles is white listed, or visible to users assigned the system role.
4	Blacklist metadata and content	Metadata and content for the specified SDK meta roles are black listed, or not visible to users assigned the system role.
5	Blacklist only metadata	Metadata for the specified SDK meta roles is black listed, or not visible to users assigned the system role.
6	Blacklist only content	Content for the specified SDK meta roles is black listed, or not visible to users assigned the system role.

Services Stats View

The Services Stats view provides a way to monitor the status and operations of a service. This view displays key statistics, service system information, and host system information for a service. In addition, more than 80 statistics are available for viewing as gauges and in timeline charts. In historical timeline charts, only statistics for session size, sessions, and packets are viewable.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	monitor the status and operations of a service.*	See the *System Maintenance Guide*.
Administrator	chart statistical information for a service over a user-specified period of time.*	See the *System Maintenance Guide*.

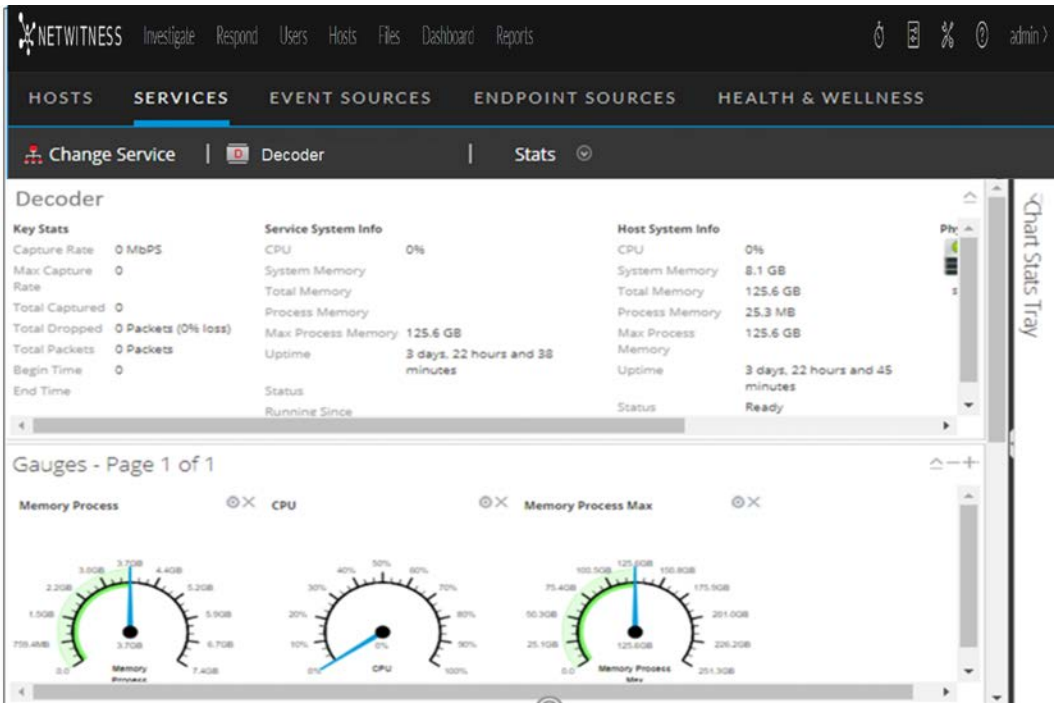
* You can perform these tasks in the current view.

Related Topics

- Edit Service Dialog
- Services Config View
- Services Explore View
- Services Logs View
- Services Security View
- Services System View
- Services View

Quick Look

The following figure shows an example of the Services Stats view for a Network Decoder.



Although different statistics are available for different types of services, certain sections are common to the Services Stats view for any Core service:

- Summary Stats
- Gauges
- Timeline Charts
- Historical Timeline Charts
- Chart Stats Tray

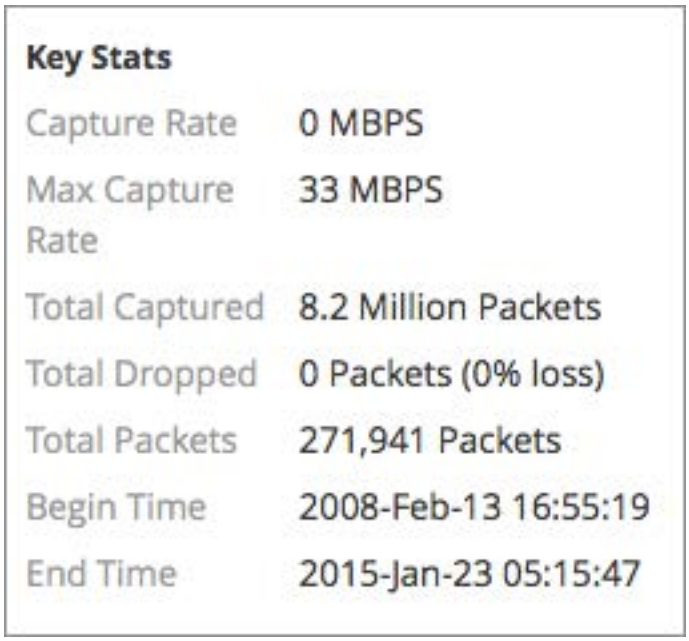
Summary Stats Section

The Summary Stats section is at the top of the default view and has no editable fields. There are five panels in the Summary Stats section: Key Stats, Service System Info, Host System Info, Logical Drives, and Physical Drives. The **Key Stats** panel displays different statistics for different types of services. The remaining four panels in the Summary Stats section are the same for all types of services.

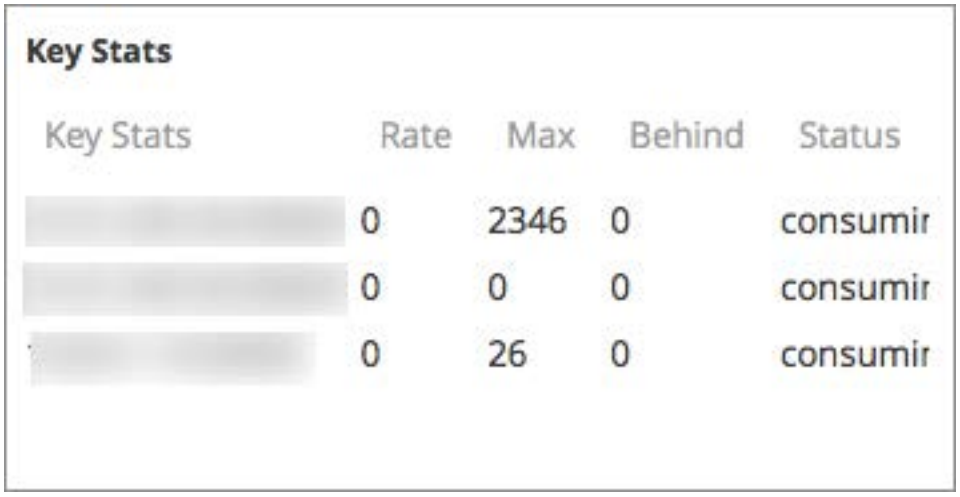
Key Stats

The Key Stats panel displays different statistics for different types of services.

- For a **Network Decoder** or **Log Decoder**, key statistics include capture statistics, such as capture rate, total packets or logs captured, total packets or logs dropped, the data capture begin time and end time.



- A **Broker** or **Concentrator** aggregates data from multiple services. Therefore, the key statistics for all aggregate services are presented in a list. The columns in the list provide the service name, the capture rate, the maximum capture rate, the number of sessions behind (that need to be aggregated), and the service status.



Service System Info

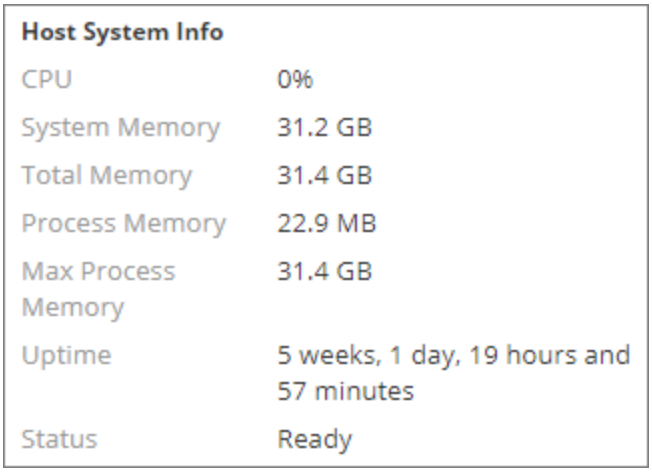
The Service System Info panel includes the percentage of CPU used by the service, the memory usage statistics (system, total, process, and maximum process), service uptime, status, running since time, and the current time.



Service System Info	
CPU	7%
System Memory	14.9 GB
Total Memory	15.6 GB
Process Memory	111.4 MB
Max Process Memory	15.6 GB
Uptime	1 week, 6 days, 3 hours and 25 minutes
Status	Ready
Running Since	2015-Jan-23 09:29:11

Host System Info

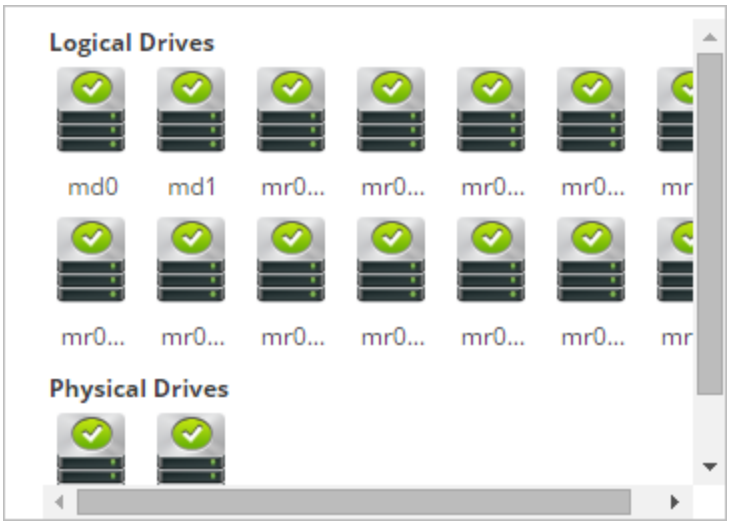
The Host System Info panel includes the percentage of CPU used by the host, the memory usage statistics (system, total, process, and maximum process), host uptime, status, running since time, and the current time.



Host System Info	
CPU	0%
System Memory	31.2 GB
Total Memory	31.4 GB
Process Memory	22.9 MB
Max Process Memory	31.4 GB
Uptime	5 weeks, 1 day, 19 hours and 57 minutes
Status	Ready

Logical Drives and Physical Drives

The Logical Drives panel and Physical Drives panel are shown with an icon for the drive name and state. Drive types are used in the names and the drive status options are listed below.



Drive Types and Status

Drive Type	Description	Comments	Status Options
`sd`	SCSI block device	Directly connected SAS, SATA MegaRAID volumes.	Green: OK Red: FAIL
`ld`	MegaRAID Logical Volume	Defined in BIOS or with MegaCLI tool.	Green: OK Yellow: DEGRADED/BUILDING Red: FAIL
`pd`	MegaRAID Physical Disks	Not directly exposed to Linux.	Green: OK Red: FAIL
`md`	Linux software RAID Volume		Green: OK Yellow: DEGRADED/BUILDING Red: FAIL

Gauges

The Gauges section in the Services Stats view presents statistics in the form of analog gauges. See Services Stats View - Gauges for details on configuring gauges.

Timeline Charts

Timeline charts display the selected statistics in a running timeline with focus on the current time. This is the same for all types of services, and only the display name of the timeline is editable. See Services Stats View - Timeline Charts for details on configuring timelines.

Historical Timeline Charts

Historical timeline charts display statistics for session size, sessions, and packets in a historical timeline. This is the same for all types of services. Only the display name, begin date, and end date. See Services Stats View - Timeline Charts for details on configuring timelines.

Note: Historical timeline charts is being deprecated for Log Collector, Virtual Log Collector (VLC) and Windows Legacy Collector services.

Chart Stats Tray

The Chart Stats Tray lists all available statistics for the selected service type. Different services have different statistics to monitor. See Services Stats View - Chart Stats Tray for a detailed description.

Topics

- Services Stats View - Chart Stats Tray
- Services Stats View - Gauges
- Services Stats View - Timeline Charts

Services Stats View - Chart Stats Tray

In the Services Stats view, the Chart Stats Tray provides a way to customize the monitored statistics for individual services. The Chart Stats Tray lists all available statistics for the service. The number of statistics varies according to the type of service being monitored. Any statistic in the Chart Stats Tray can be displayed in a gauge or a timeline chart. Only statistics for session size, sessions, and packets are viewable in historical timeline charts.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	customize the monitored statistics for individual services.*	See the *System Maintenance Guide*.

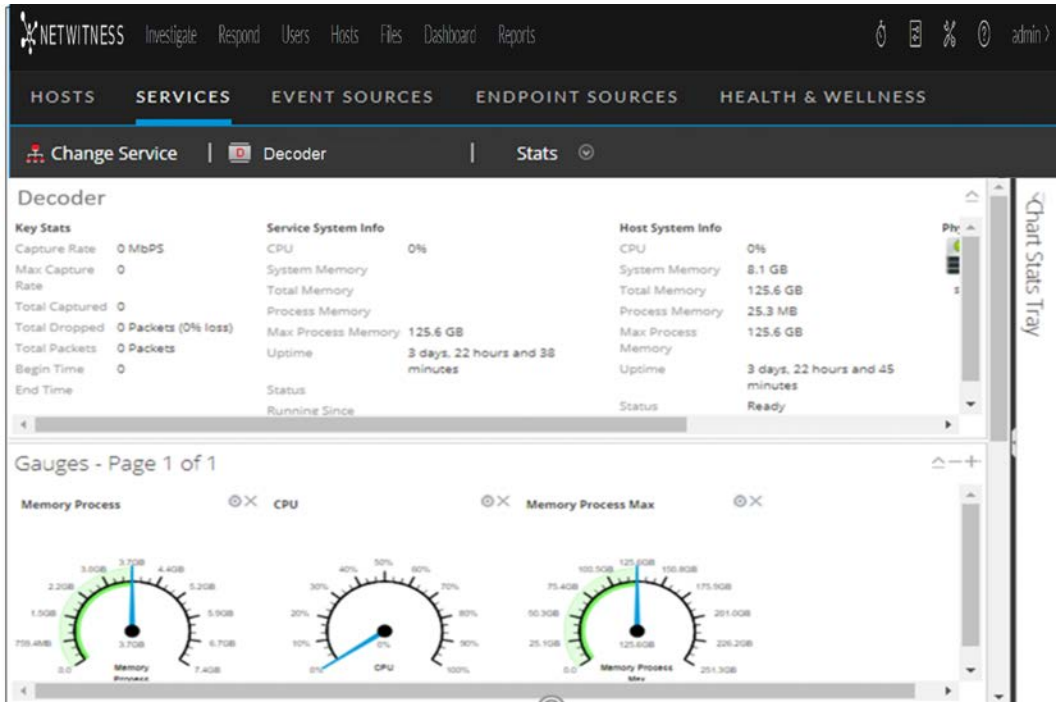
* You can perform these tasks in the current view.

Related Topics

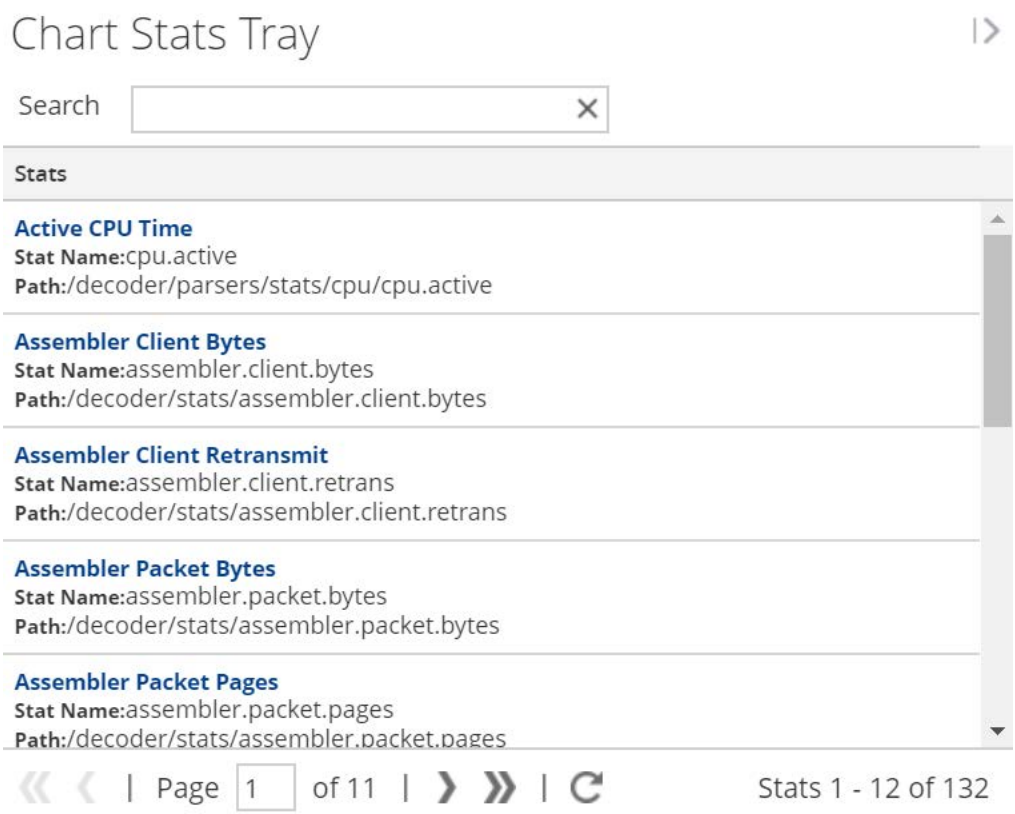
- Services Stats View - Gauges
- Services Stats View - Timeline Charts
- Services Config View

Quick Look

The following example shows the Services Stats view for a Network Decoder. The Chart Stats Tray is collapsed.



To view the Chart Stats Tray, click on the  to expand the Chart Stats Tray.



The Chart Stats Tray has different statistics for different types of services. In the example above, 132 statistics are available for the Network Decoder. The following table describes features of the Chart Stats Tray.

Feature	Description
	Click to expand the panel horizontally.
	Click to collapse the panel horizontally.
Search	Type a search term in the field and press **RETURN**. Statistics that match are displayed with the matching word highlighted.
	Click to go to the first page.
	Click to go to the previous page.
Page 1 of 11	Type a page number in the Page field.
	Click to go to the next page.
	Click to go to the last page.

Feature	Description
	Click to refresh the view.
Stats 1 - 12 of 132	Displays the range of statistics being displayed. The total number statistics varies by service type.

Services Stats View - Gauges

The Gauges section of the Services Stats view presents statistics in the form of an analog gauge. You can drag any statistic available in the Chart Stats Tray to the Gauges section. The properties of each individual gauge are editable; all gauges have an editable title and some have additional editable properties.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	view statistics in the form of an analog gauge.*	See the *System Maintenance Guide*.

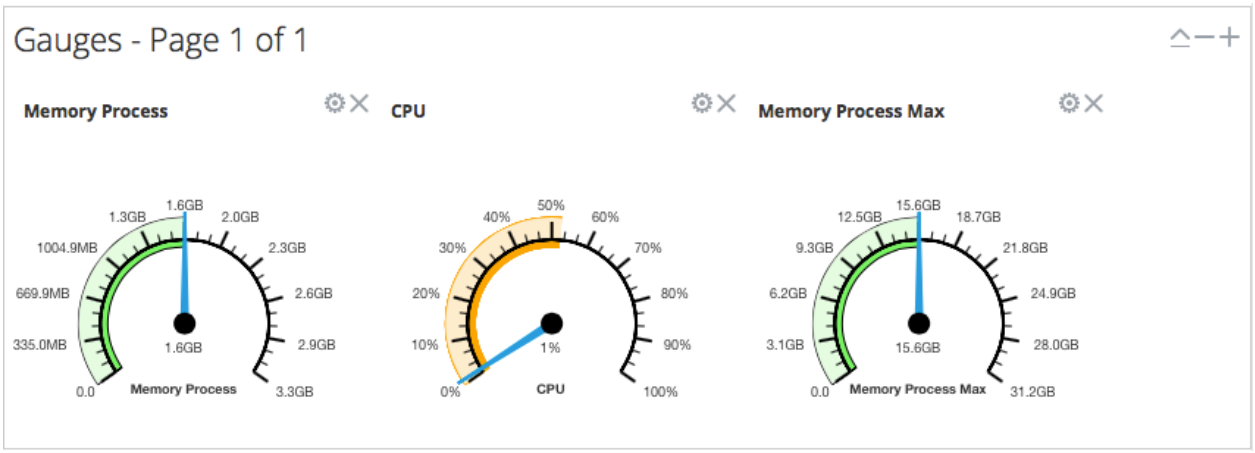
* You can perform these tasks in the current view.

Related Topics

- Services Stats View - Chart Stats Tray
- Services Stats View - Timeline Charts
- Services Config View

Quick Look

The following figure shows the default gauges in the Services Stats view for a Log Decoder.



The default gauges show these statistics:

- Process memory use
- CPU use
- Maximum process memory use

The controls in the Gauges title bar and in each gauge are the standard dashlet controls. Dashlets are the parts that make up a dashboard.



- In the Gauges title bar (from left to right), you can collapse/expand, delete a page, add a page, page backward, and page forward.
- In each gauge, you can edit properties ( ) and delete ( ) the gauge.

Services Stats View - Timeline Charts

The Services Stats view Timeline Charts sections display statistics in a running timeline. The Services Stats view includes two types of timelines: current time and historical. You can drag any statistic available in the Chart Stats Tray to the Timeline Charts section or Historical Timeline Charts section. Only statistics for session size, sessions, and packets are viewable in historical timeline charts. The properties of an individual timeline chart are editable; all timeline charts have an editable title and some have additional editable properties.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	view statistics in the form of a current or historical timeline.*	See the *System Maintenance Guide*.

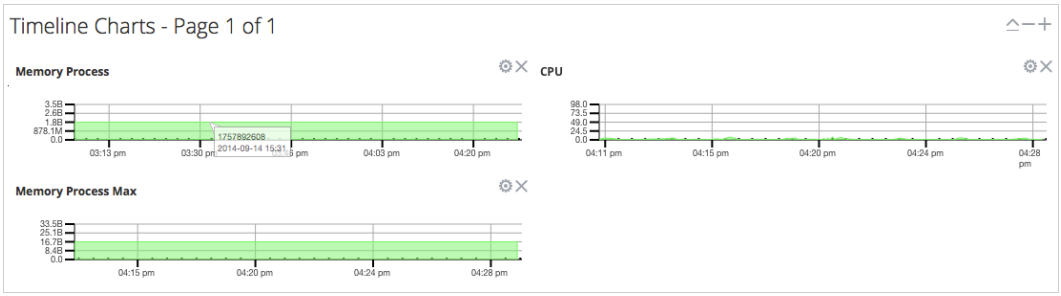
* You can perform these tasks in the current view.

Related Topics

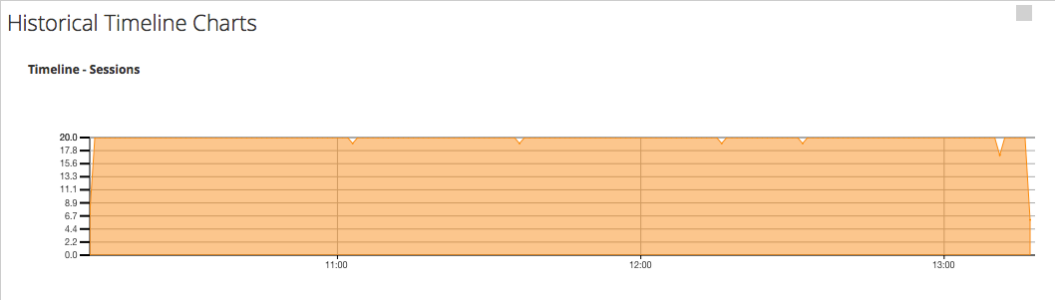
- Services Stats View - Chart Stats Tray
- Services Stats View - Gauges
- Services Config View

Quick Look

The following figure is an example of a current timeline showing the value and timestamp of a data point.



The following figure is an example of a historical timeline chart.



The default current timeline charts show these statistics:

- Memory Process
- CPU
- Memory Process Max

The historical time charts show these statistics:

- Sessions
- Packets
- Session Size

The controls in the Timeline Charts title bar and in each timeline are the standard dashlet controls. The Historical Timeline Charts title bar and timelines have the same controls.



- In the Timeline Charts title bar (from left to right), you can collapse/expand, delete a page, add a page, page backward, and page forward.
- In each timeline, you can edit Properties ( ) and delete ( ) the timeline.
- Hovering over a data point in the chart, displays the value and timestamp for the selected point.

Services System View

This topic introduces features in the Services System view using Decoders (Network Decoder and Log Decoder) as an example. See the appropriate configuration guides for other services (for example, the *NetWitness Broker and Concentrator Configuration Guide*) for details on their respective **(Admin) > Services > System** views.

A Log Decoder is a special type of Decoder, and is configured and managed in a similar way to a Network Decoder. Therefore, most of the information in this section refers to both types of Decoders. Differences for Log Decoders are noted. For more information on Decoders, see the *Decoder Configuration Guide*.

Workflow



What do you want to do?

User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	view system and session information about a service.*	See the *System Maintenance Guide*.

* You can perform these tasks in the current view.

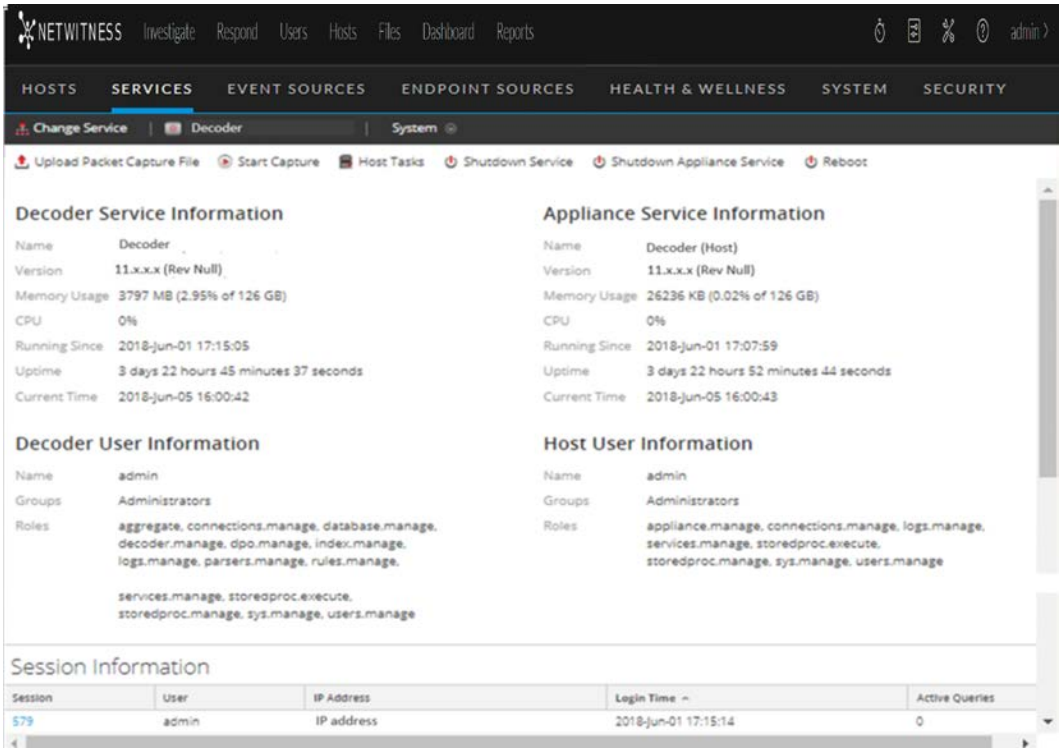
Related Topics

- Edit Service Dialog
- Services Config View
- Services Explore View
- Services Logs View
- Services Security View

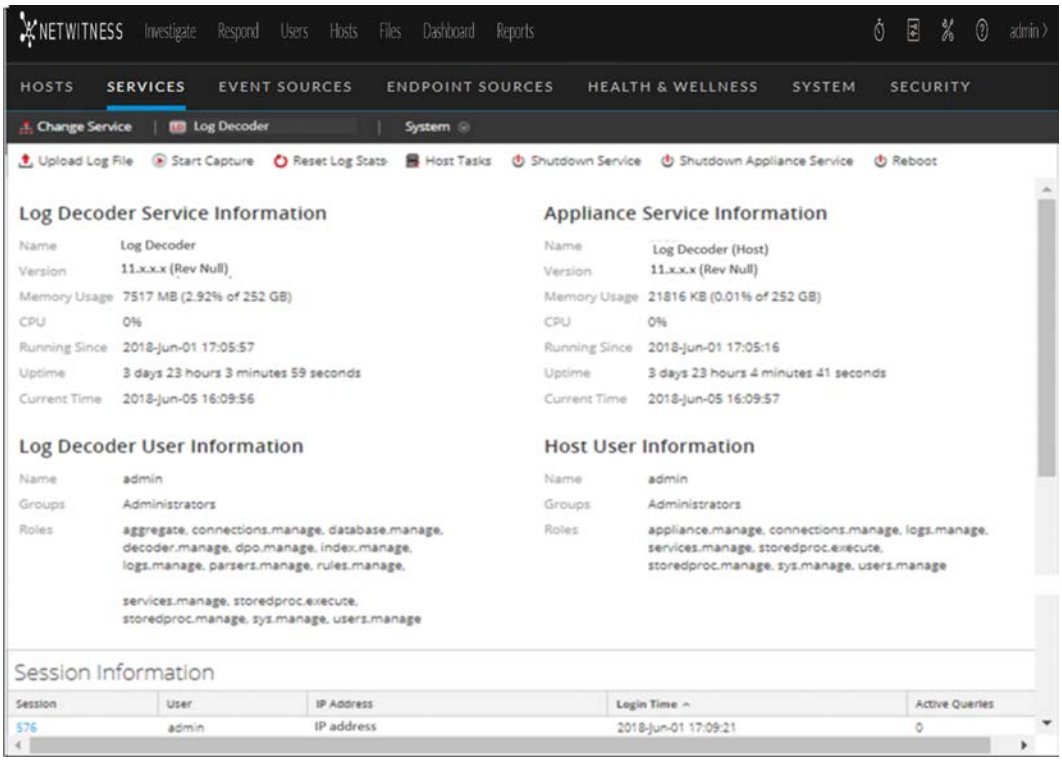
- Services Stats View
- Services View

Quick Look

The following figure shows an example of the Services System view for a Network Decoder.



The following figure shows the Services System view for a Log Decoder.



Services Info Toolbar

The following toolbars show the options specific to Log Decoders (top) and Network Decoders (bottom).





In addition to the common options in the Services System view toolbar, you can start and stop capture of packets or logs. The upload file options are different for the standard Network Decoder (packet capture file) and the Log Decoder (log file).

Action	Description
Upload Packet Capture File	Displays a dialog that provides a way to select a packet capture (`.pcap`) file for upload to the selected Network Decoder. For more information, see "Upload Packet Capture File" in the *Decoder Configuration Guide*. **Note:** This option does not apply to Log Decoders.
Upload Log File	Displays a dialog that provides a way to select a log (`.log`) file for upload to the selected Log Decoder. For more information, see "Upload Log File to a Log Decoder" in the *Decoder Configuration Guide*.
Start/Stop Capture	Starts packet capture on the selected Decoder. When packet capture is in progress, the option in the toolbar changes to Stop Capture, and the option to upload a file is unavailable.

Related Topics

- Services System View - Host Task List Dialog

Services Topology View

This topic introduces the services topology, a hierarchical view of the NetWitness core services depicting the collection and aggregation in your deployment.

This visualization displays the topology for the Broker, Concentrator, Log Decoder, Packet Decoder, Hybrids, ESA and Log Collector. You can view the services that are online and details of each node.

Workflow

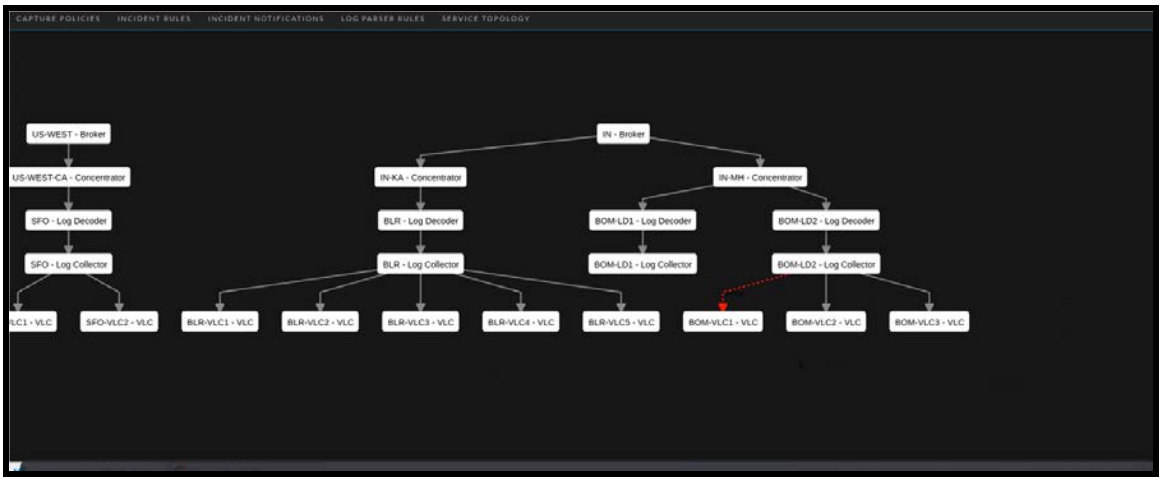


What do you want to do?

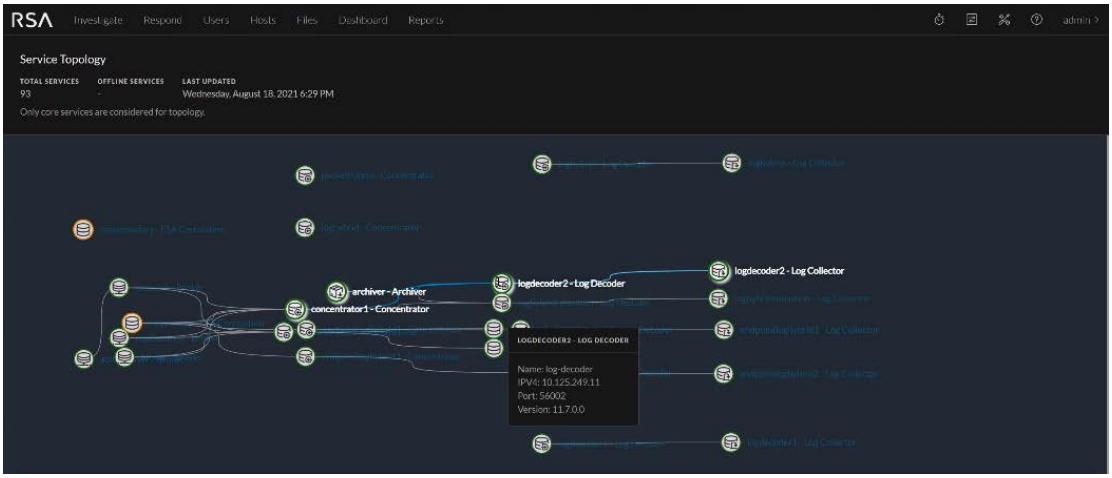
User Role	I want to...	Documentation
Administrator	maintain a service.	Maintaining Services

Quick Look

The following figure shows an example of how each service is aggregating data in a hierarchy along with which services are online or offline.



The following figure shows the details of a node. When you hover over any of the services, additional details such as name, IP version, Port number and Version number are displayed.



Action	Description
Node	Displays a dialog that provides name, IP version, Port number and Version number.
Search	Enter the service name to locate a specific service.

Services System View - Host Task List Dialog

In the NetWitness Services System view, you can use the Host Tasks option to manage tasks that relate to a host and its communications with the network. Several service and host configuration options are available for Core services.

Workflow



What do you want to do?

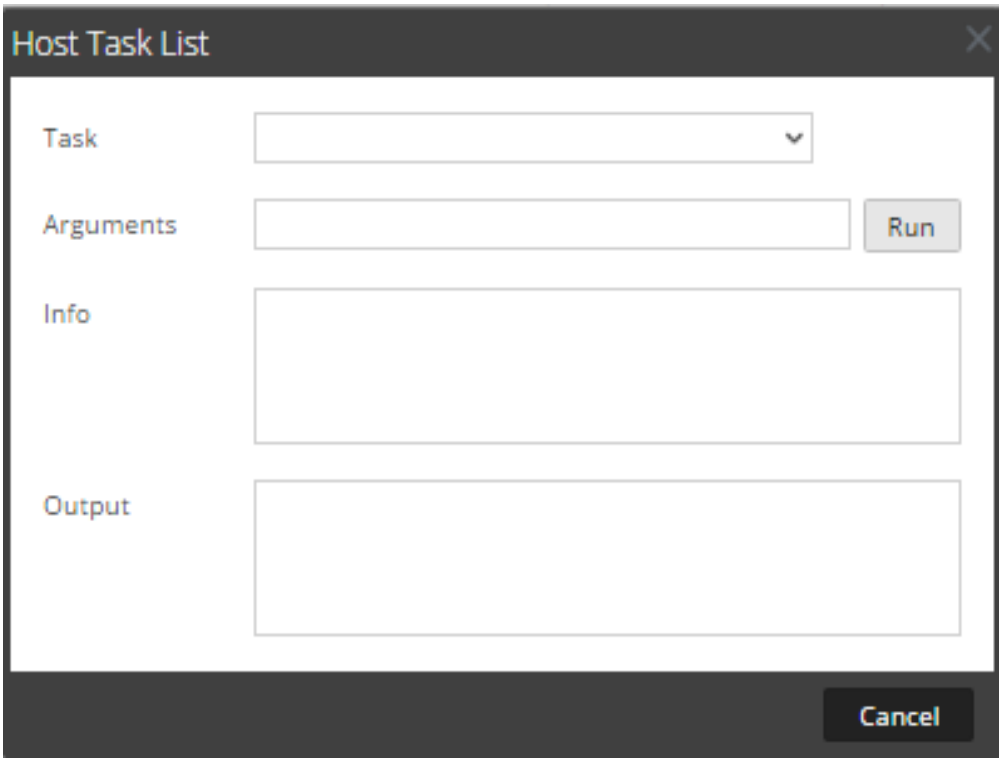
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	manage host-related tasks and host communications with the network.*	Hosts and Services Maintenance Procedures

* You can perform these tasks in the current view.

Related Topics

- Services System View

Quick Look



The table below describes the Host Task List dialog features.

Field	Description
Task	An entry field in which you type or select a message for a Core host. When you click in this field a drop-down list of available host tasks is displayed.
Arguments	An entry field in which you enter the arguments, if any, for the message.
Run	Executes the task and arguments in the entry fields.
Info	Information about the message purpose and syntax.
Output	The output or result of an executed task.
Cancel	Closes the Host Task list dialog.

Host Task Selection List

These tasks are displayed as a drop-down list in the Task field. The available options are regulated by the security role required to execute the option.

Task	Description
Add Filesystem Monitor	Starts monitoring the storage services attached to the specified filesystem. See Add and Delete a Filesystem Monitor.

Task	Description
Delete Filesystem Monitor	Stops monitoring the storage services attached to the specified filesystem. See Add and Delete a Filesystem Monitor.
Reboot Host	Shuts down and restarts the host. See Reboot a Host.
Set Host Built-in Clock	Sets the local host clock. See Set Host Built-In Clock.
Set Host Hostname	This method of changing the hostname is deprecated in NetWitness 10.6. To edit a hostname, see Edit or Delete a Service.
Set Network Time Source	Sets the clock source for this host. See Set Network Time Source.
Set Syslog Forwarding	Enables or disables syslog forwarding from a remote server to the selected service. See Set Syslog Forwarding.
Show Network Port Status	Shows the network interface information for a host. See Show Network Port Status.
Show Serial Number	Gets the host serial number. See Show Serial Number.
Shut Down Host	Shuts down the physical host and the host remains off. See Shut Down Host.
Start Service	Starts a service on this host. See Stop and Start a Service on a Host.
Stop Service	Stops a service on this host. See Stop and Start a Service on a Host.
setSNMP	Enables or disables the SNMP service on a host. See Set SNMP.

Service Configuration Settings

This topic introduces the available service configuration settings for NetWitness Core services.

NetWitness Core services include Brokers, Concentrators, Network Decoders, Log Decoders, and Archivers. The service configuration parameters listed below constitute all viewable and configurable parameters. Some parameters are configurable in various parts of the NetWitness user interface while other parameters are viewable or configurable only on the Services Explore view.

Aggregation Configuration Parameters

This table lists and describes the available configuration parameters that are common to services that perform aggregation, such as Concentrators and Archivers.

Configuration Path	/concentrator/config or /archiver/config
aggregate.autostart	Automatically restarts aggregation after a service restart, if enabled. Change takes effect immediately.
aggregate.buffer.size	Displays the size of the buffer (default unit is KB) used per round of aggregation. Larger buffers may improve aggregation performance but could impact query performance. Change takes effect after aggregation restart.
aggregate.crc	If enabled, all aggregation streams will be CRC validated. Change takes effect immediately.
aggregate.hours	Displays the maximum number of hours behind a service will be allowed to start aggregation. Change takes effect immediately.
aggregate.interval	Lists the minimum number of milliseconds before another round of aggregation is requested. Change takes effect immediately.
aggregate.meta.page.factor	Lists the allocated number meta pages per session used for aggregation. Change takes effect on service restart.
aggregate.meta.perpage	Lists the allocated number of meta stored on one page of data. Change takes effect on service restart.
aggregate.precache	Determines if the concentrator will precache the next round of aggregation for upstream services. Can improve aggregation performance but could impact query performance. Change takes effect immediately.
aggregate.sessions.max	Lists the number of sessions to aggregate on each round. Change takes effect after aggregation restart.
aggregate.sessions.perpage	Lists the number of sessions stored on one page of data. Change takes effect on service restart.

Configuration Path	/concentrator/config or /archiver/config
aggregate.time.window	Displays the maximum +/- time window, in seconds, that all services must be inside before another round of aggregation is requested. Zero turns off time window. Change takes effect immediately.
consume.mode	Determines if the concentrator can only aggregate locally or over a network, based on licensing restrictions. Change takes effect on service restart.
export.enabled	Allows export of session data, if enabled. Change takes effect on service restart.
export.expire.minutes	Lists the number of minutes before export cache files are expired and flushed. Change takes effect immediately.
export.format	Determines the file format used during data export. Change takes effect on service restart.
export.local.path	Displays the local location to cache exported data. Optional assigned max size (=#unit), units are: t for TB; g for GB, m for MB. Change takes effect on service restart.
export.meta.fields	Determines which meta fields are exported. Comma-separated list of fields. * means all fields. * and field list means all fields except the listed fields. Just the field list means only those fields are included. Change takes effect immediately.
export.remote.path	Displays the remote protocol (nfs://) and location to export data. Change takes effect on service restart.
export.rollup	Determines the rollup interval for export files. Change takes effect on service restart.
export.session.max	Displays the maximum sessions per exported file. For export file types that cache, this determines cached memory sizes. Zero is no limit. Change takes effect immediately.
export.size.max	Displays the maximum bytes per exported file. For export file types that cache, this determines cached memory sizes. Zero is no limit. Change takes effect immediately.
export.usage.max	Displays the maximum percentage of cache space used before stopping aggregation. Zero is no limit. Change takes effect immediately.
heartbeat.error	Lists the number of seconds to wait after a service error before attempting a service reconnect. Change takes effect immediately.
heartbeat.interval	Lists the number of milliseconds between heartbeat service checks. Change takes effect immediately.

Configuration Path	/concentrator/config or /archiver/config
heartbeat.next.attempt	Lists the number of seconds to wait before attempting a service reconnect. Change takes effect immediately.
heartbeat.no.response	Lists the number of seconds to wait before taking unresponsive service offline. Change takes effect immediately.

Appliance Service Configuration Parameters

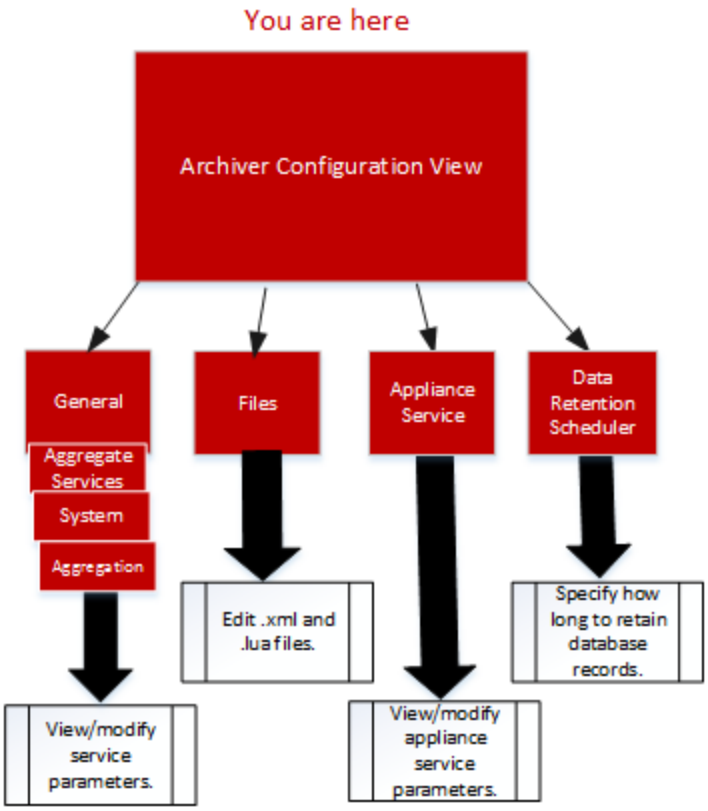
The NetWitness Core Appliance service provides hardware monitoring on legacy NetWitness hardware. The list describes the Appliance Configuration parameters.

Appliance Parameter Field	Description
Logs	/logs/config refer to Core Service Logging Configuration Parameters
REST	/rest/config refer to REST Interface Configuration Parameters
Services	/services/<service name>/config refer to Core Service-to-Service Configuration Parameters
System	/sys/config refer to Core Service System Configuration Parameters

Archiver Service Configuration View

This topic lists and describes the available configuration settings for NetWitness Archivers.

Workflow



Role	I want to ...
Administrator	Configure Meta Filters for Aggregation. Refer to "(Optional) Configure Meta Filters for Aggregation" in the *NetWitness Archiver Configuration Guide* for instructions.
Administrator	Configure Group Aggregation. Refer to "Configure Group Aggregation" in the *NetWitness Deployment Guide* for instructions.

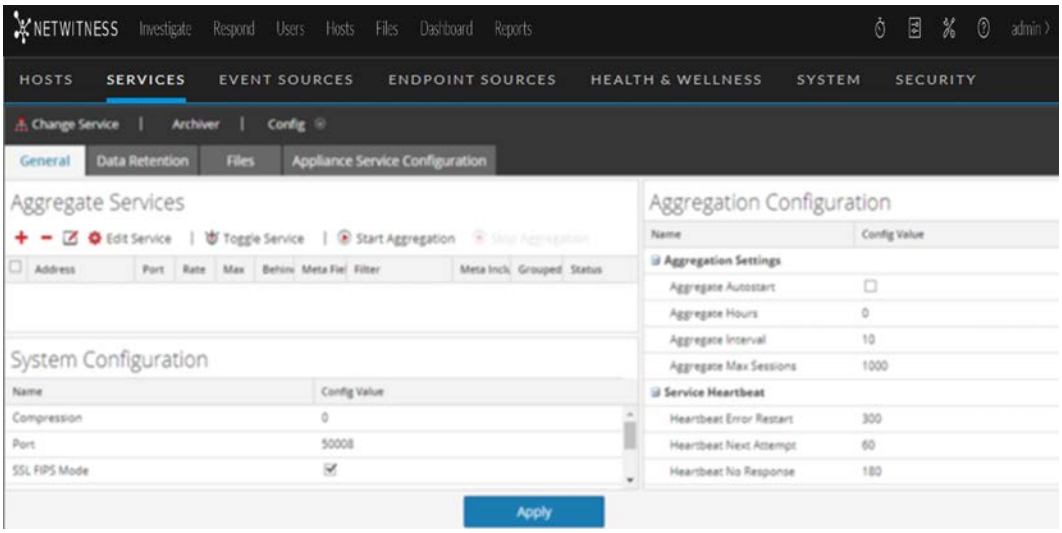
Quick Look

To access the Services Config view:

1. In **NetWitness**, select  **(Admin) > Services**.
   The Admin Services view is displayed.

2. Select an Archiver service and select  **>View > Config**.
   Services Config view for the Archiver service is displayed.

This is an example of the Services Config view for an Archiver.



Broker Service Configuration Parameters

The following list describes the Broker configuration parameters.

Broker Parameter Field	Description
Broker	`/broker/config` refer to Aggregation Configuration Parameters
`aggregate.interval.behind`	Minimum number of milliseconds before another round of aggregation is requested when the broker is behind. Change takes effect immediately.
Database	`/database/config` refer to "Database Configuration Nodes" in the *NetWitness Core Services Database Tuning Guide*.
Index	`/index/config`
`index.dir`	The directory where the broker device mapping files are stored. Change takes effect on service restart.
`language.filename`	The index language specification (XML) that is loaded on startup. Change requires service restart.
Logs	`/logs/config` refer to Core Service Logging Configuration Parameters
REST	`/rest/config` refer to REST Interface Configuration Parameters
SDK	`/sdk/config` refer to "SDK Configuration Nodes" in the *NetWitness Core Services Database Tuning Guide* and NetWitness Platform Core Service system.roles Modes
Services	`/services/<service name>/config` refer to Core Service-to-Service Configuration Parameters

Broker Parameter Field	Description
System	`/sys/config` refer to Core Service System Configuration Parameters

Concentrator Service Configuration Parameters

The following list describes the Concentrator configuration parameters.

Concentrator Parameter Field	Description
Concentrator	`/concentrator/config` refer to Aggregation Configuration Parameters
Database	`/database/config` refer to "Database Configuration Nodes" in the *NetWitness Core Database Tuning Guide*.
Index	`/index/config` refer to "Index Configuration Nodes" in the *NetWitness Core Database Tuning Guide*.
Logs	`/logs/config` refer to Core Service Logging Configuration Parameters
REST	`/rest/config` refer to REST Interface Configuration Parameters
SDK	`/sdk/config` refer to "SDK Configuration Nodes" in the *NetWitness Core Database Tuning Guide*and NetWitness Platform Core Service system.roles Modes
Services	`/services/<service name>/config` refer to Core Service-to-Service Configuration Parameters
System	`/sys/config` refer to Core Service System Configuration Parameters

Core Service Logging Configuration Parameters

The following table describes the logging configuration parameters for all NetWitness Core services. Logging configuration is the same for all Core services.

Logs Configuration Folder	/logs/config
`log.dir`	Displays the directory where the log database is stored. Optional assigned max size (=#) is in MBs. Change takes effect on service restart.
`log.levels`	Controls what types of log messages are stored (comma separated). Module specific settings are defined like this: `<Module>=[debug|info|audit|warning|failure|all|none]`. Change takes effect immediately.

Logs Configuration Folder	/logs/config
log.snmp.agent	Sets a remote SNMP Trap Receiving agent.
snmp.trap.version	Sets the SNMP version (2c or 3) to be used for gets and traps.
snmpv3.engine.boots	Displays the SNMPv3 engine boots count. This field auto-increments on startup and should not normally need to be set by the user.
snmpv3.engine.id	Sets the SNMPv3 engine ID, which is 10-64 hexadecimal digit number optionally preceded by 0x. You can add suffix values at the end of the engine ID for each of the SA Core services running on the same host. For example, if the generated Engine ID for the SA Core host is 0x1234512345, you can set the Engine ID for the Decoder service as 0x123451234501 and set 0x123451234504 for the Appliance service.
snmpv3.trap.auth.local.key	Sets the SNMPv3 Trap Authentication Local Key, which is a 16 or 20 hexadecimal digit number (depending on which authentication protocol is used) preceded by 0x. For MD5, the key is 16 hexadecimal digits, while SHA uses 20 hexadecimal digits. You can use any desired algorithm to generate the local keys. It is recommended that a generation method involving randomness be used as opposed to selecting key values manually.
snmpv3.trap.auth.protocol	Displays the SNMPv3 Trap Authentication Protocol (none, MD5 or SHA).
snmpv3.trap.priv.local.key	Sets the SNMPv3 Trap Privacy Local Key, which is a 16 hexadecimal digit number preceded by 0x.
snmpv3.trap.priv.protocol	Displays the SNMPv3 Trap Privacy Protocol (none or AES).
snmpv3.trap.security.level	Displays the SNMPv3 Trap Security Level, which indicates whether authentication and privacy are used or not. Possible values are noAuthNoPriv, authNoPriv, or authPriv.
snmpv3.trap.security.name	Sets the SNMPv3 Trap Security Name used during SNMPv3 trap authentication.
syslog.size.max	Displays the maximum size of a log sent to syslog (some syslog daemons have issues with very large messages). Zero means no limit. Change takes effect immediately.

Core Service-to-Service Configuration Parameters

This topic lists and describes the configuration parameters that control how a Core service connects to another Core service. For example, when a Concentrator connects to a Network Decoder, the connection parameters are controlled by these settings.

Whenever a Core service establishes a connection to another Core service, the service that acts as the **client** creates a new sub-folder in the `/services` folder of the configuration tree. The name of the sub-folder corresponds to the name of the service and has the form `host:port`. For example, the service connection folder for a Concentrator connection to a Network Decoder could be `/services/reston-va-decoder:50004`. Inside each service connection folder, there is a `config` sub-folder that holds configurable parameters.

The following list describes the Service Configuration parameters:

Services	/services/host:port/config
allow.nonssl.to.ssl	Allows a non-SSL connection to connect to a SSL service, when set to true. Otherwise, if `false`, non-secure to secure connections will be denied. Change takes effect immediately.
compression	Displays a config node that determines if data is compressed before sending. A positive value determines the number of bytes that need to be sent before it will be compressed. Zero means no compression.
crc.checksum	Displays a config node that determines if data streams are validated with a CRC checksum. A positive value determines the number of bytes that need to be sent before it will be CRC validated. Zero means no CRC validation.
ssl	Displays a config node that enables or disables SSL encryption on the connection.

Core Service System Configuration Parameters

The following list describes the System configuration parameters that are common to all NetWitness Core services.

System Configuration Folder	/sys/config
compression	Displays the minimum amount of bytes before a message is compressed, when set to a positive value. Zero means no compression for any message. Change takes effect on subsequent connections.
crc.checksum	Displays the minimum bytes before a message is sent over the network with a CRC checksum (to be validated by the client), when set to a positive value. Zero means no CRC checksum validation with any message. Change takes effect on subsequent connections.
drives	Displays drives to monitor for usage stats. Change takes effect on service restart.
port	Displays the port this service will listen on. Change takes effect on service restart.
scheduler	Displays the folder for scheduled tasks.

System Configuration Folder	/sys/config
service.name.override	Displays an optional service name used by upstream services for aggregation in lieu of hostname.
ssl	Encrypts all traffic using SSL, if enabled. Change takes effect on service restart.
stat.compression	Compresses stats as they are written to the database, if enabled. Change takes effect on service restart.
stat.dir	Displays the directory where the historical stats database is stored (separate multiple dirs with semicolon). Optional assigned max size (=#unit), units are: t for TB; g for GB, m for MB. Change takes effect on service restart.
stat.exclude	Lists stat pathnames to be excluded from the stat database. The following wildcards are permitted: ? match any single character * match zero or more characters to delimiter / ** match zero or more characters including delimiter. Change takes effect immediately.
stat.interval	Determines how often (in milliseconds) statistic nodes are updated in the system. Change takes effect immediately.
threads	Lists the number of threads in the thread pool to handle incoming requests. Change takes effect immediately.

Decoder Configuration Parameters

The following list describes the configuration parameters that are identical on both Network Decoder and Log Decoder services.

Configuration Path	<service>/config
aggregate.buffer.size	Displays the size of the buffer (default unit is KB) used per round of aggregation. Larger buffers may improve aggregation performance but could impact capture performance. Change takes effect after capture restart.
aggregate.precache	Determines if the decoder will pre-cache the next round of aggregation for upstream services. Can improve aggregation performance but could impact capture performance. Change takes effect immediately.
assembler.pool.ratio	Displays the percentage of pool pages that assembler manages and uses for the assembly process. Change takes effect on service restart.

Configuration Path	<service>/config
assembler.session.flush	Flushes sessions either when they are complete, or when they are parsed. Change takes effect on service restart.
assembler.session.pool	Lists the number of entries in the session pool. Change takes effect on service restart.
assembler.size.max	Lists the maximum size that a session will obtain. A setting of 0 removes the session size limit. Change takes effect immediately.
assembler.size.min	Lists the minimum size that a session must be before persisting. Change takes effect immediately.
assembler.timeout.packet	Lists the number of seconds before packets are timed out. Change takes effect immediately.
assembler.timeout.session	Lists the number of seconds before sessions are timed out. Change takes effect immediately.
assembler.voting.weights	Displays the weights used to determine which session stream is marked client and server. Change takes effect immediately.
capture.autostart	Determines if capture begins automatically when the service starts. Change takes effect on service restart.
capture.buffer.size	Displays capture memory buffer allocation size (default unit is MB). Change takes effect on service restart.
capture.device.params	Displays capture service specific parameters. Change takes effect on service restart. The parameters understood by this field are specific to the currently selected capture device. If any of the parameters are not recognized by the current capture device, they are ignored. On Log Decoders, there is only the Log Events capture device. It accepts some optional parameters. • use-envision-time: If this is set to 1, the time metadata for each event will be imported from the Log Collector stream. If this is 0 or not set, the imported event time will be stored in the event.time meta. • port: This parameter can be set to a numeric value to override the default syslog port listener, 514.
capture.selected	Displays current capture service and interface. Change takes effect immediately.
export.expire.minutes	Lists the number of minutes before export cache files are expired and flushed. Change takes effect immediately.
export.packet.enabled	Allows export of packet data, if enabled. Change takes effect on service restart.

Configuration Path	<service>/config
export.packet.local.path	Displays the local location to cache packet exported data. Optional assigned max size (=#unit), units are: t for TB; g for GB, m for MB. Change takes effect on service restart.
export.packet.max	Displays the maximum packets per exported file. For export file types that cache this determines cached memory sizes. Zero is no limit. Change takes effect immediately.
export.packet.remote.path	Lists the remote protocol (nfs://) and location to export data. Change takes effect on service restart.
export.packet.size.max	Displays the packet maximum bytes per exported file. For export file types that cache this determines cached memory sizes. Zero is no limit. Change takes effect immediately.
export.rollup	Determines the rollup interval for export files. Change takes effect on service restart.
export.session.enabled	Allows export of session data, if enabled. Change takes effect on service restart.
export.session.format	Determines the file format used during session export. Change takes effect on service restart.
export.session.local.path	Displays the local location to cache session exported data. Optional assigned max size (=#unit), units are: t for TB; g for GB, m for MB. Change takes effect on service restart.
export.session.max	Displays the maximum sessions per exported file. For export file types that cache, this determines cached memory sizes. Zero is no limit. Change takes effect immediately.
export.session.meta.fields	Determines which meta fields are exported. Comma-separated list of fields. * means all fields. * plus field list means all fields BUT listed fields. Just field list means only those fields are included. Change takes effect immediately.
export.session.remote.path	Displays the remote protocol (nfs://) and location to export data. Change takes effect on service restart.
export.session.size.max	Lists the session maximum bytes per exported file. For export file types that cache, this determines cached memory sizes. Zero is no limit. Change takes effect immediately.
export.usage.max	Lists the session maximum bytes per exported file. For export file types that cache, this determines cached memory sizes. Zero is no limit. Change takes effect immediately.
parse.threads	Lists the number of parse threads to use for session parsing. Zero means let server decide. Change takes effect on service restart.

Configuration Path	<service>/config
pool.packet.page.size	Displays the size of a packet page (default is KB). Change takes effect on service restart.
pool.packet.pages	Lists the number of packet pages decoder will allocate and use. Change takes effect on service restart.
pool.session.page.size	Displays the size of a session page (default is KB). Change takes effect on service restart.
pool.session.pages	Lists the number of session pages decoder will allocate and use. Change takes effect on service restart.

Network Decoder Service Configuration Parameters

This topic lists and describes the available configuration parameters for NetWitness Network Decoders.

Decoder Parameter Field	Description
Decoder	/decoder/config refer to Decoder Configuration Parameters.
Database	/database/config refer to "Database Configuration Nodes" in the *NetWitness Core Database Tuning Guide*.
Index	/index/config refer to "Index Configuration Nodes" in the *NetWitness Core Database Tuning Guide*.
Logs	/logs/config refer to Core Service Logging Configuration Parameters.
REST	/rest/config refer to REST Interface Configuration Parameters.
SDK	/sdk/config refer to "SDK Configuration Nodes" in the *NetWitness Core Database Tuning Guide* and NetWitness Platform Core Service system.roles Modes.
System	/sys/config refer to Core Service System Configuration Parameters.

Log Decoder Service Configuration Parameters

This topic lists and describes the available configuration parameters for Log Decoder configuration settings.

Log Decoder Setting Field	Description
Database	/database/config refer to "Database Configuration Nodes" in the *NetWitness Core Database Tuning Guide*.

Log Decoder Setting Field	Description
Decoder	`/decoder/config` refer to Decoder Configuration Parameters.
Index	`/index/config` refer to "Index Configuration Nodes" in the *NetWitness Core Database Tuning Guide.*
Logs	`/logs/config` refer to Core Service Logging Configuration Parameters.
REST	`/rest/config` refer to REST Interface Configuration Parameters.
SDK	`/sdk/config` refer to "SDK Configuration Nodes" in the *NetWitness Core Database Tuning Guide*and NetWitness Platform Core Service system.roles Modes.
System	`/sys/config` refer to Core Service System Configuration Parameters.

Log Tokenizer Configuration Settings

The Log Decoder has a set of configuration items that control how the automatic log tokenizer creates meta items from unparsed logs. The log tokenizer is implemented as a set of built-in parsers that each scan for a subset of recognizable tokens. The functionality of each of these native parsers is shown in the table below. These word items form a full-text index when they are fed to the indexing engine on the Concentrator and Archiver. By manipulating the parsers.disabled configuration entry, you can control which Log Tokenizers are enabled.

Parser Name	Description	Configuration Parameters
Log Tokens	Scans for runs of consecutive characters to produce 'word' meta items.	`token.device.types`, `token.char.classes`, `token.max.length`, `token.min.length`, `token.unicode`
IPSCAN	Scans for text that appears to be an IPv4 address to produce `ip.addr` meta items.	`token.device.types`
IPV6SCAN	Scans for text that appears to be an IPv6 address to produce `ipv6` meta items.	`token.device.types`
URLSCAN	Scans for text that appears to be a URL to produce `alias.host`, `filename`, `username`, and `password` meta items.	`token.device.types`
DOMAINSCAN	Scans for text that appears to be a domain name to produce `alias.host`, `tld`, `cctld`, and `sld` meta items.	`token.device.types`

Parser Name	Description	Configuration Parameters
EMAILSCAN	Scans for text that appears to be an email address to produce `email` and `username` meta items.	`token.device.types`
SYSLOGTIMESTAMPSCAN	Scans for text that appears to be syslog-format timestamps. Syslog is missing the year and time zone. When such text is located, it is normalized into UTC time to create `event.time` meta items.	`token.device.types`
INTERNETTIMESTAMPSCAN	Scans for text that appears to be RFC 3339-format timestamps to create `event.time` meta items.	`token.device.types`

Log Tokenizer Configuration Parameters.

Log Decoder Parser Setting Field	Description
`token.device.types`	The set of device types that will be scanned for raw text tokens. By default, this is set to `unknown`, which means only logs that were not parsed will be scanned for raw text. You can add additional log types here to enrich parsed logs with text token information. If this field is empty, then log tokenization is disabled.
`token.char.classes`	This field controls the type of tokens that are generated. It can be any combination of the values `alpha`, `digit`, `space`, and `punct`. The default value is `alpha`. • `alpha`: Tokens may contain alphabetic characters • `digit`: Tokens may contain numbers • `space`: Tokens may contain spaces and tabs • `punct`: Tokens may contain punctuation marks
`token.max.length`	This field puts a limit on the length of the tokens. The default value is **5** characters. The maximum length setting allows the Log Decoder to limit the space needed to store the word metadata. Using longer tokens requires more meta database space, but may provide slightly faster raw text searches. Using shorter tokens causes the text query resolver to have to perform more reads from the raw logs during searches, but it has the effect of using much less space in the metadb and index.
`token.min.length`	This is the minimum length of a searchable text token. The minimum token length will correspond to the minimum number of characters a user may type into the search box in order to locate results. The recommended value is the default, **3**.

Log Decoder Parser Setting Field	Description
token.unicode	This boolean setting controls whether unicode classification rules are applied when classifying characters according to the token.char.classes setting. When this is set to true, each log is treated as a sequence of UTF-8 encoded code points and then classification is performed after the UTF-8 decoding is performed. When this is set to false, each log is treated as ASCII characters and only ASCII character classification is done. Unicode character classification requires more CPU resources on the Log Decoder. If you do not need non-English text indexing, you can disable this setting to reduce CPU utilization on the Log Decoder. The default is enabled.

REST Interface Configuration Parameters

The following list describes the available configuration parameters for the REST interface built in to all NetWitness Core Services.

REST Configuration Path	/rest/config
cache.dir	Displays the host directory to use for temporarily creating and storing files. Change takes effect on service restart.
cache.size	Displays the total maximum size (default unit is MB) of all files in the cache directory before the oldest are deleted. Change takes effect on service restart.
enabled	Switches to enable or disable REST services, **1** is on, **0** is off. Change takes effect on service restart.
port	Displays the port the REST service will listen on. Change takes effect on service restart.
ssl	Encrypts all REST traffic using SSL, if enabled. The default system means use setting from /sys/config/ssl. Change takes effect on service restart.

NetWitness Platform Core Service system.roles Modes

All NetWitness Platform Core services offer role-based authorization modes. This topic describes the modes that are available, and how they are configured within every service.

The configuration node /sdk/config/system.roles sets querying and viewing permissions for metadata and content on a per key basis. This parameter supports the data privacy management function and when enabled using one of the non-zero values helps a data privacy officer to control access to specific meta keys and content. This parameter is configurable in the NetWitness Platform user interface (see "Data Privacy Tab" in the *Data Privacy Management Guide* for details). When the value is edited, change takes effect immediately.

Zero means that service permissions based on SDK meta keys are disabled.

- 0 - disabled

When one of the non-zero values is specified, the data privacy officer can select a meta key to whitelist or blacklist the display of the associated metadata, content, or both, for a specific user role on a service.

- 1 - whitelist meta and content filtered
- 2 - whitelist meta filtered
- 3 - whitelist content filtered
- 4 - blacklist meta and content filtered
- 5 - blacklist meta filtered
- 6 - blacklist content filtered

Centralized Service Configuration via Policy Settings

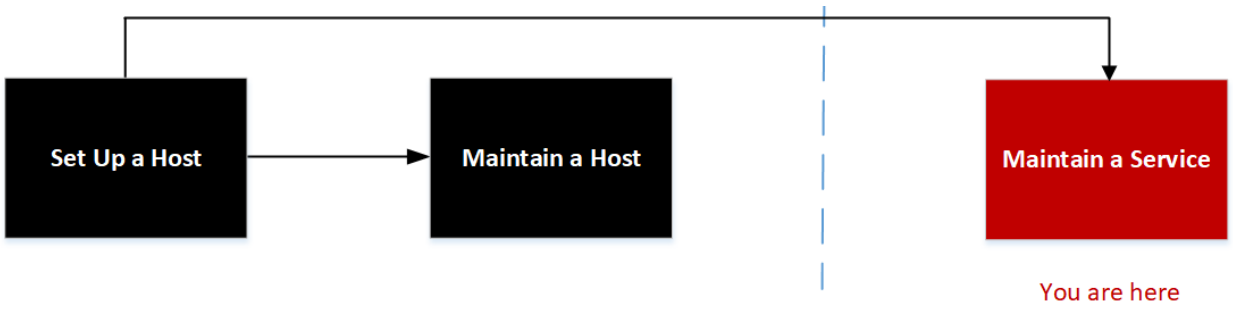
This topic introduces the available centralized service configuration settings for Concentrators, Decoders and Log Decoders services.

Centralized Service Configuration – Groups Tab

Note: The information in this topic applies to NetWitness Platform Version 11.7 and Later.

The  (**Configure**) > **Policies** view contains two tabs: **Groups** and **Policies**.

Workflow



What do you want to do?

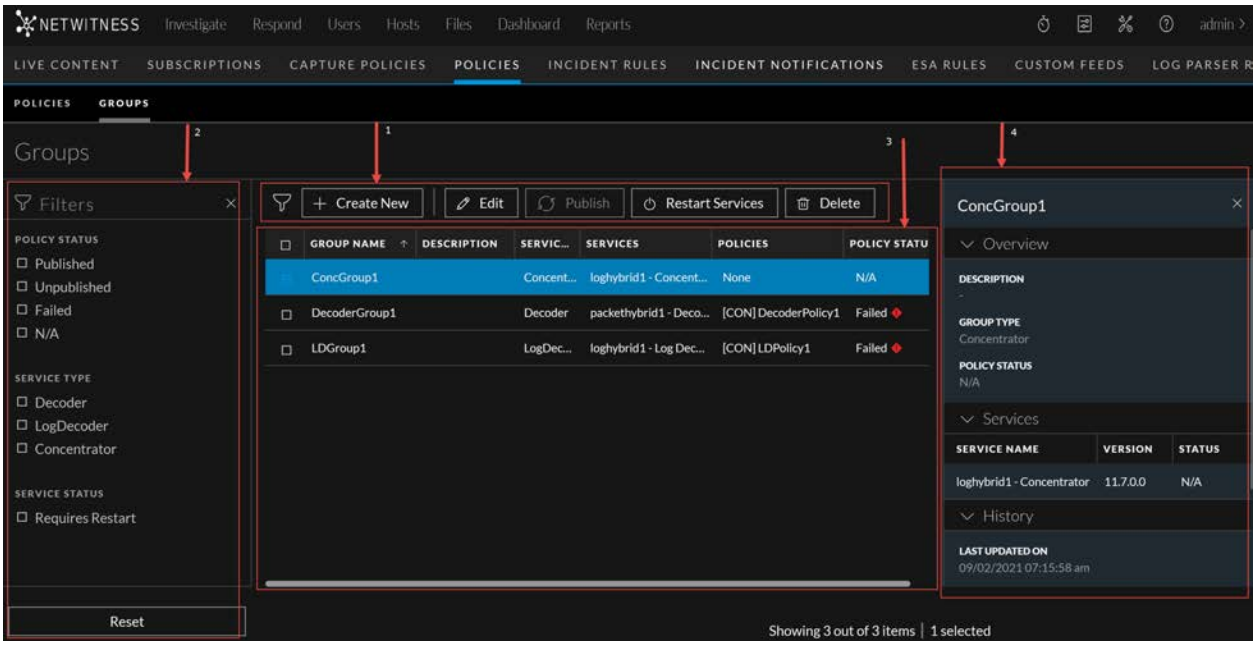
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services
Administrator	Centrally Create and Maintain Service Groups.*	Hosts and Services Maintenance Procedures

* You can perform these tasks in the current view.

Related Topics:

- Hosts and Services Set Up Procedures
- Hosts and Services Maintenance Procedures

Quick Look



Below is an example of the **Groups** tab:

The following table describes the Groups tab.

1	**Toolbar** • **Create New**- Lets you create a new group. For more information, see "Create a Group" in the Centralized Service Configuration via Policy. • **Edit** – Lets you edit the group. For more information, see "Managing Policies and Groups" in the Centralized Service Configuration via Policy. • **Publish** – Publishes selected groups. • **Restart Services** – Restarts the service associated with the selected group. This option is disable if: ◦ There are no services in the group. ◦ No service restart is required. • **Delete**- Deletes the selected groups.

2 **Filter Pane**

- **Filters**: You can filter groups based on Policy Type and Policy Status, Service Type and Service Status.

  To hide, click the  icon at the top-right of the panel. To display if hidden, click the  icon in the toolbar.

- **Reset**: Removes the currently applied filter criteria.

For more information, see "Filter Groups" in the Centralized Service Configuration via Policy.

3 **Group List Pane**

- **Group Name** – Name of the group.
- **Description** – Description of the group.
- **Service Type** – Displays the service type to which the group is applied.
- **Services** – Displays the service to the which the group is applied.
- **Policies** – Displays the policy to which the group is applied.
- **Policy Status** – Status of the policy. The values are: Published, Unpublished, Failed, N/A.

You can also sort on any column. If you mouse over a column header, a sort icon is displayed: . Click the icon to sort by the selected column.

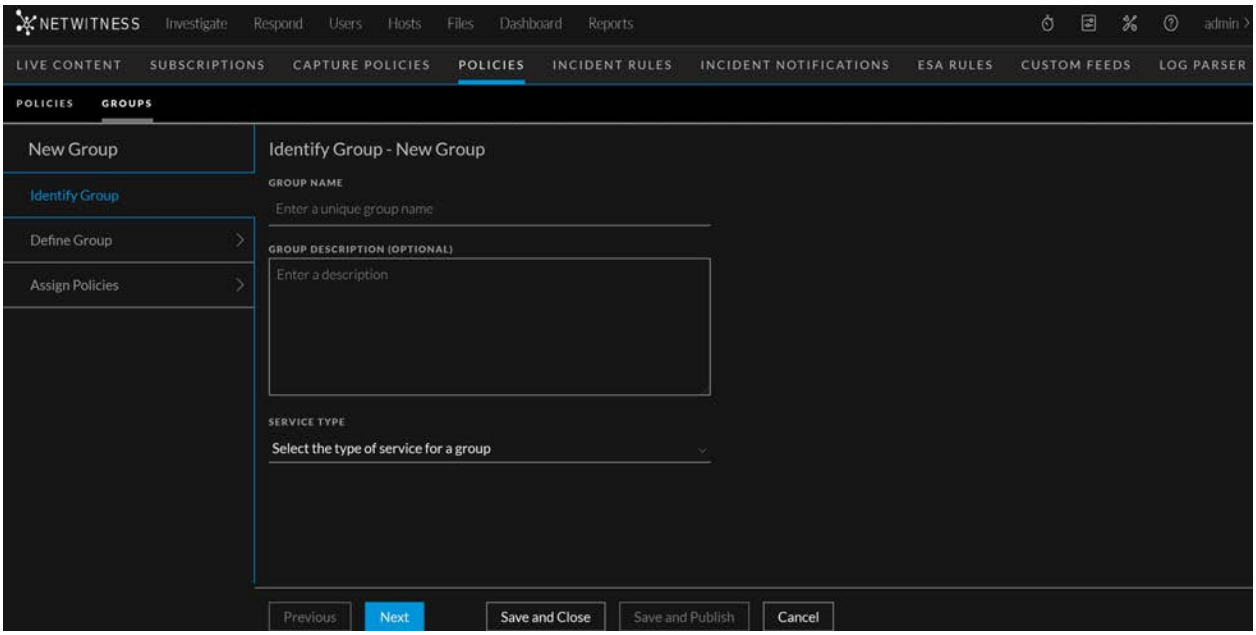
4 **Groups Details Panel**

Displays the properties of the selected group.

Note: Click the row to view the Properties panel for a group.

Create Group Dialog

Below is an example of the Create Group dialog.

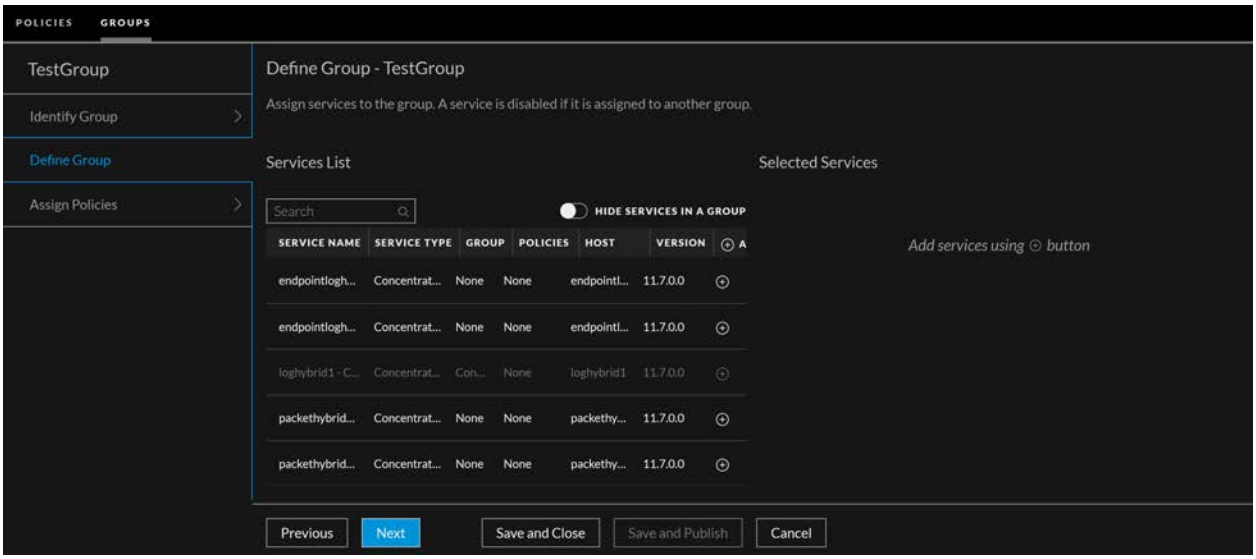


The table describes the information and options in the Create Group dialog.

Field	Description
Group Name	Name of the group. The name should be unique.
Group Description (Optional)	Description of the group. Description should not exceed 8000 characters.
Service Type	Select the type of service for a group. Available options are Log Decoder, Decoder, and Concentrator.
Save and Close	Saves the settings and closes the Create Group dialog.

Define Group Dialog

Below is an example of the define group dialog.



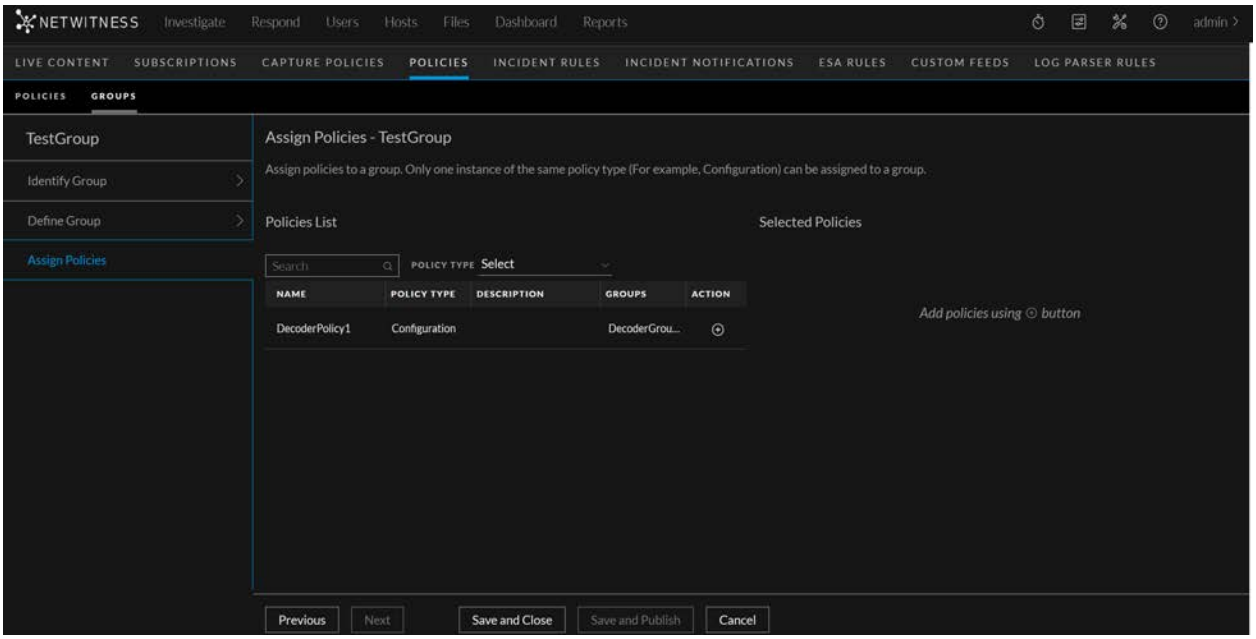
The table describes the information and options in the Define Group dialog.

Field	Description
Service List	Displays the list of services associated with the selected service type. The following describes services list: • **Service name** – Name of the service. • **Service type** – Displays the service type of the service. • **Group** – Name of the group. • **Policies** – Displays the policies associated with the service. • **Host** – Host name of the service. • **Version** – Service version. • **All** – Lets you to add services to the group. You can either click  to add all services or click  to add specific service.

Field	Description
Hide Services in a Group	Displays the services that is not assigned to any group. By default, this option is disabled.
Selected Services	Displays the list of selected services for the group.
Save and Close	Saves the setting and closes the Define Group dialog.
Save and Publish	Saves and publishes the created group. **Note:** This option is disabled if you have not: - Assigned services - Assigned policies

Assign Policy Dailog

Below is an example of Assign Policy dialog.



The following table describes assign policy dialog:

Field	Description
Policies List	Displays the list of policies associated with the group. The following describes policies list: • **Name** – Name of the policy. • **Policy Type** – Displays the policy type. • **Description** – Description of the policy. • **Groups** – Groups associated with the policy. • **Action** - Click  to add policies to the group.
Policy Type	Select the policy type from the drop-down list.
Selected Policies	Displays the list of selected policies for the group.
Save and Close	Saves the settings and closes the Create Group dialog.

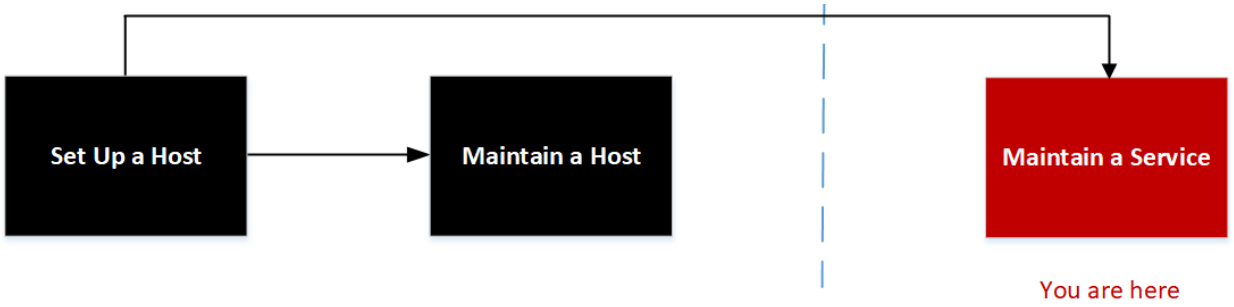
Field	Description
Save and Publish	Saves and publishes the created group. **Note:** This option is disabled if you have not: - Assigned services. - Assigned policies.

Centralized Service Configuration – Policies Tab

Note: The information in this topic applies to NetWitness Platform Version 11.7 and Later.

The  (**Configure**) > **Policies** view contains two tabs: **Groups** and **Policies**.

Workflow



Related Topics:

- Hosts and Services Set Up Procedures
- Hosts and Services Maintenance Procedures

What do you want to do?

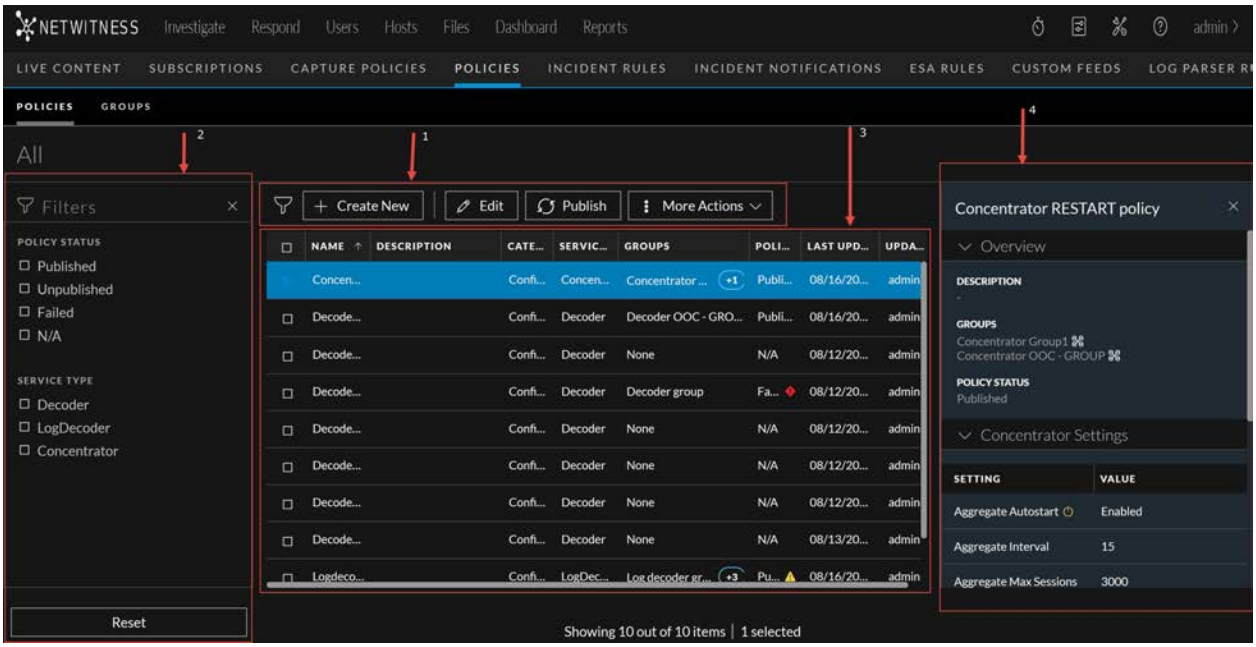
User Role	I want to...	Documentation
Administrator	set up a host.	Setting Up a Host
Administrator	maintain a host.	Maintaining Hosts
Administrator	maintain a service.	Maintaining Services

User Role	I want to...	Documentation
Administrator	Centrally Create and Maintain Service Policy.*	Hosts and Services Maintenance Procedures

* You can perform these tasks in the current view.

Quick Look

Below is an example of the **Policies** tab:



1 **Toolbar**

- **Create New**- Lets you create a new policy. For more information, see "Create a Policy" in the Centralized Service Configuration via Policy.
- **Edit** – Lets you edit the policy. For more information, see "Managing Policies and Groups" in the Centralized Service Configuration via Policy.
- **Publish** – Publishes the selected policy or policies.
- **More Actions**
  - **Assign to Group** -Lets you assign policy to a group.
  - **Clone** – Lets you clone a policy.
  - **Revert** – Lets you revert to previous policy settings. By default, the revert value is set to 5. You can customize the revert value in Explorer view.
  - **Delete** - Deletes the selected policy or policies permanently.

2 **Filter Pane**

- **Filters**: You can filter policies based on Policy Type and Policy Status and Service Type.

  To hide, click the  icon at the top-right of the panel. To display if hidden, click the  icon in the toolbar.

- **Reset**: Removes the currently applied filter criteria.

For more information, see "Filter Policies" in the Centralized Service Configuration via Policy.

3 **Policy List Pane**
Name – Name of the policy.

• **Description** – Description of the policy.

• **Category** – Type of policies applied.

• **Service Type** – Displays the service type to which the policy is applied.

• **Groups** - Lists the group to which this policy is applied.

• **Policy Status** – Status of the policy. The values are: Published, Unpublished, Failed, N/A.

• **Last Updated** – Displays the time when the policy is updated.

• **Updated By** – The user who updated the policy.

You can also sort on any column. If you mouse over a column header, a sort icon is displayed: . Click the icon to sort by the selected column.

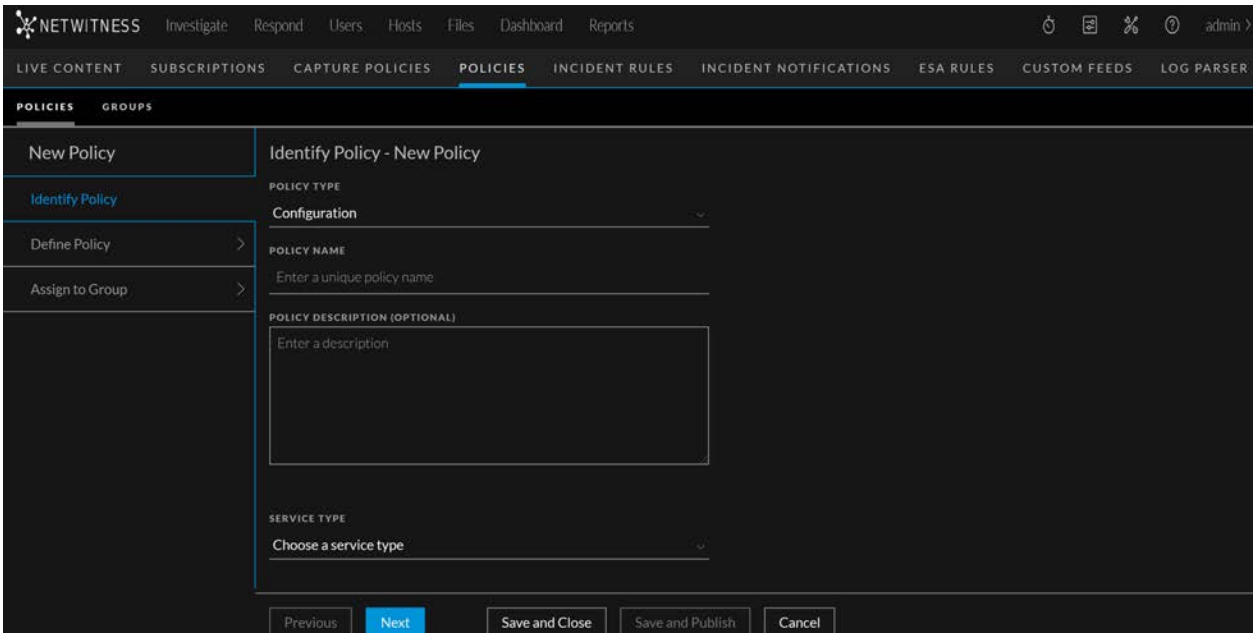
4 **Policy Details Pane**

Displays the properties of the selected policy.

Note: Click the row to view the Properties panel for a policy.

Create Policy

Below is an example of the Create Policy dialog.

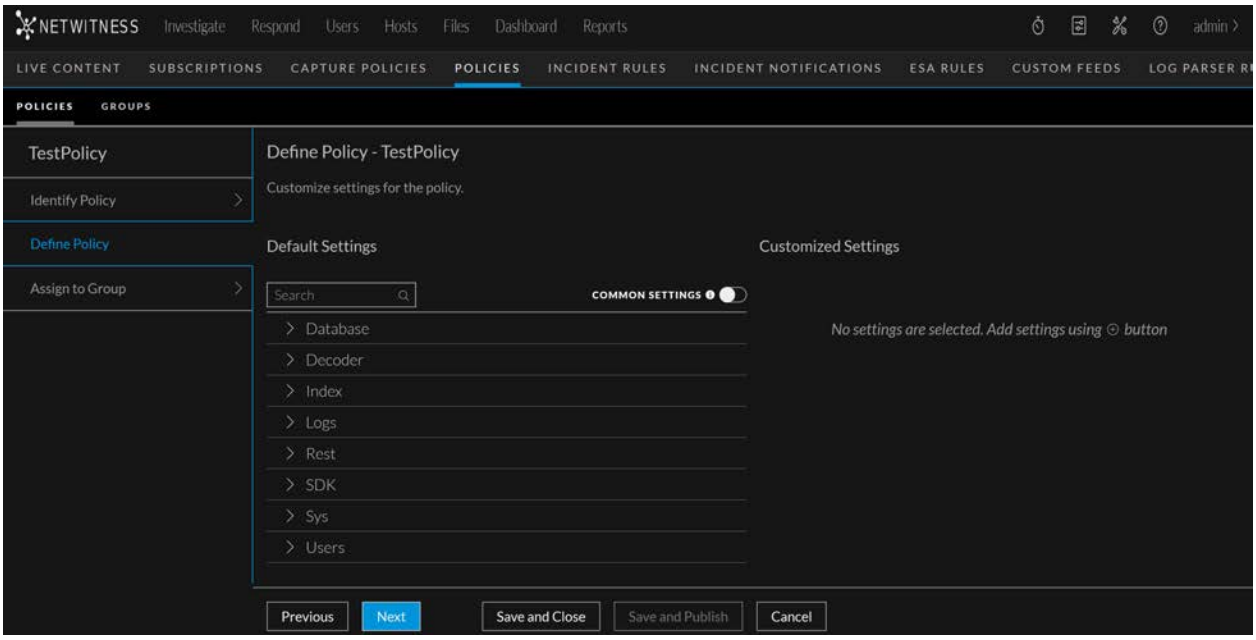


The table describes the information and options in the Create Policy dialog.

Field	Description
Policy Type	Displays the policy type.
Policy Name	Name of the policy. The name should be unique.
Policy Description (Optional)	Description of the policy. Description should not exceed 8000 characters.
Service Type	Displays the type of the service. Available options are Log Decoder, Decoder, and Concentrator.

Field	Description
(For version 11.7.1 and Later) Clone from Existing Service (Optional)	Displays the list of services based on the service type to clone those settings.
Save and Close	Saves the settings and closes the Create Policy dialog.

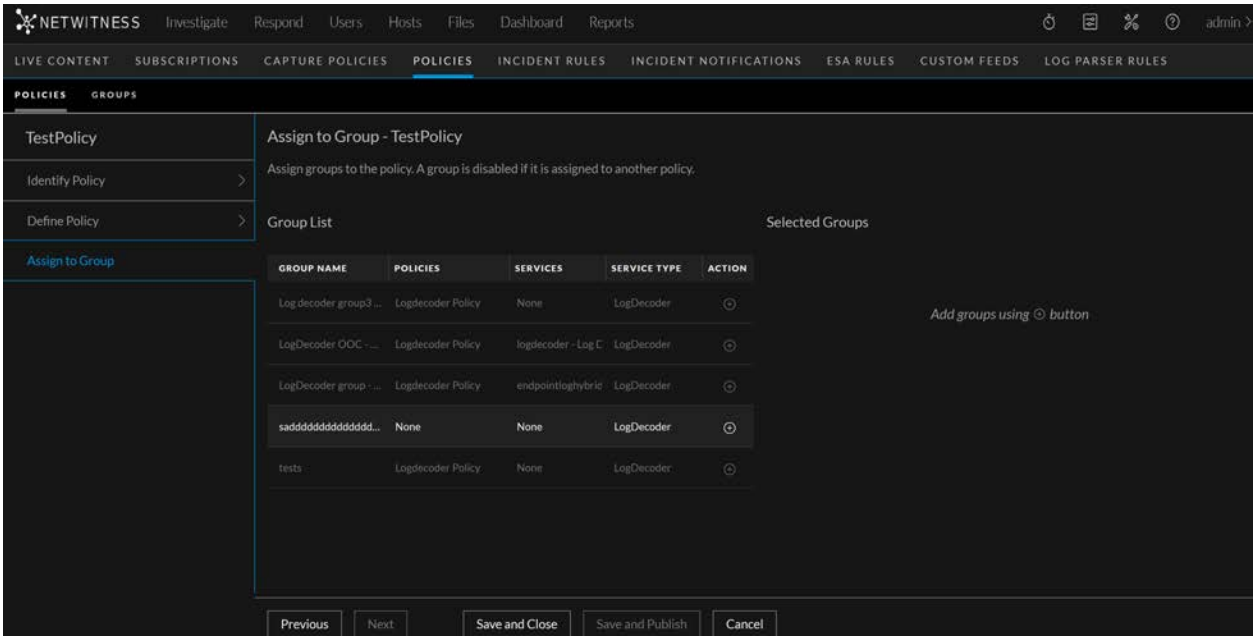
Define Policy Settings



Field	Description
Default Settings	Displays the settings of the service based on the category, for example, Database, Decoder and so on. Click to expand the setting category.  The following describes Default settings: • **Setting**- Name of the setting. • **Value**- The settings value. • **Description**- Description of the setting. • **All** – Lets you to select the settings for customization. You can either click  to customize all the settings value or click  on a specific setting
Common Settings	Displays the common settings of the service. By default, it is disabled.
(For version 11.7.1 and Later) 10g Settings	Fetches the 10g settings.

Field	Description
Customized Settings	Lists the selected settings for customization. Click ⊕ to add the setting.  Update the value based on your requirement. 

Assign to Group



Field	Description
Group List	Displays the list of group associated with the policy. A group is disabled if it is already assigned to another policy. • Group Name • Policies • Services • Service Type • Action
Selected Group	Lists the selected groups. Click  to add groups.
Save and Close	Saves the settings and closes the Assign to Group dialog.
Save and Publish	Saves and publishes the created policy. **Note:** This option is disabled if: - Policy settings are not customized. - Policy is not assigned to groups.

Troubleshooting Version Installations and Updates

This section describes the error messages displayed in the Hosts view when it encounters problems updating host versions and installing services on hosts in the Hosts view. If you cannot resolve an update or installation issue using the following troubleshooting solutions, contact Customer Support.

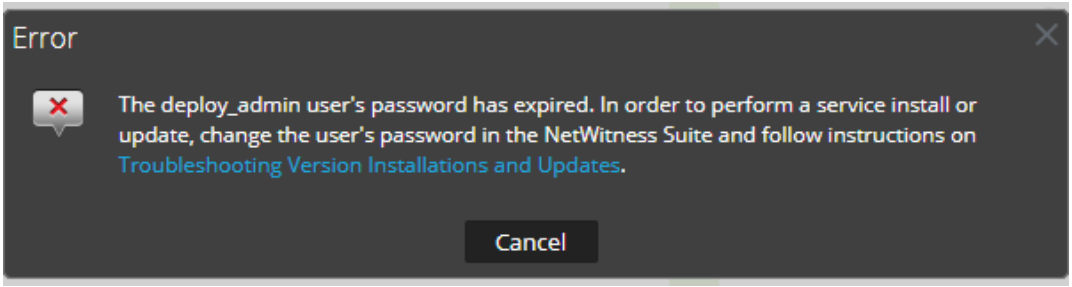
Troubleshooting instructions for the following errors that may occur during the upgrade are described in this section.

- deploy_admin Password Expired Error
- Downloading Error
- Error Deploying Version <version-number> Missing Update Packages
- Upgrade Failed Error
- External Repo Update Error
- Host Installation Failed Error
- Host Update Failed Error
- Missing Update Packages Error
- OpenSSL 1.1.x Error
- Patch Update to Non-NW Server Error
- Reboot Host After Update from Command Line Error
- Reporting Engine Restarts After Upgrade

Troubleshooting instructions are also provided for errors for the following hosts and services that may occur during or after an upgrade.

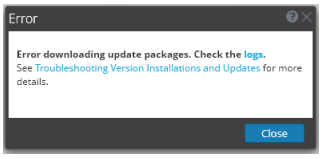
- Log Collector Service
- NW Server
- Orchestration
- Reporting Engine
- Event Stream Analysis

deploy_admin User Password Has Expired Error

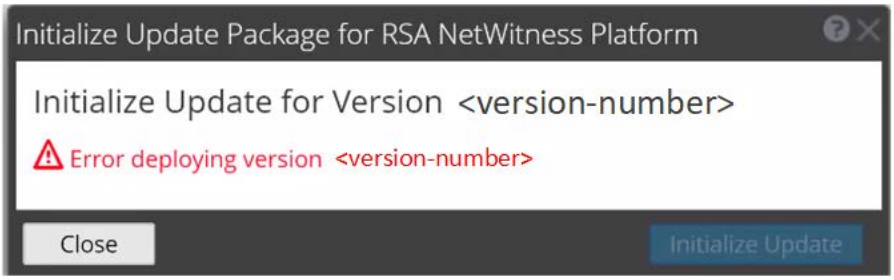
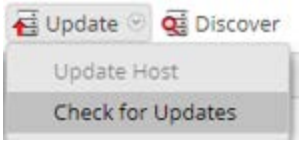
Error Message	
Cause	The `deploy_admin` user password has expired.

Solution	Reset your `deploy_admin` password password.
	1. On the NW Server host only, run the following command.
	`nw-manage --update-deploy-admin-pw`
	`Please enter the new deploy_admin account password: <new-deploy-admin-password>`
	`Please confirm the new deploy_admin account password: <new-deploy-admin-password>`
	2. Review the output of the `nw-manage --update-deploy-admin-pw` command to verify the `deploy_admin` password was successfully updated on all hosts. If an NW host is down or fails for any reason as displayed by the output of the `nw-manage --update-deploy-admin-pw` command, run `nw-manage --sync-deploy-admin-pw --host-key <host-identifier>` to synchronize the password between the NW Server and the host that failed once the communication failure is resolved.
	3. On the host that failed installation or orchestration, run the `nwsetup-tui` command and use the new **deploy_admin** password in response to the **Deployment Password** prompt.

Downloading Error

Error Message	
Problem	When you select an update version and click **Update** >**Update Host**, the download starts but fails to complete.
Cause	Version download files can be large and take a long time to download. If there are communication issues during the download it will fail.
Solution	1. Try to update again. 2. If it fails again with the same error, try to update using the offline methods as described in "Offline Method from Hosts View" or "Offline Method Using Command Line Interface" in the *Upgrade Guide for NetWitness Platform 11.6*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues. 3. If you are still not able to update, contact Customer Support.

Error Deploying Version <version-number> Missing Update Packages

Error Message	
Problem	Error deploying version <version-number> is displayed in the **Initialize Update Package for RSA NetWitness Platform** dialog after you click on **Initialize Update** if the update package is corrupted.
Solution	1. Click **Close** to close the dialog. 2. Remove the version folder from staging folder. 3. Make sure that the salt-master service is running. 4. Recopy the update package zip file to the staging folder. 5. In the Hosts view toolbar, select Check for Updates again.  6. Click **Initialize Update**. 7. Click **Update** > **Update Hosts** from the toolbar. 8. Click **Begin Update** from the **Update Available** dialog. After the host is updated, it prompts you to reboot the host. 9. Click **Reboot** from the toolbar.

Upgrade Failed Error

Error Message	While updating/installing a device to version 11.2 or above, the following error can occur and be found in /var/log/netwitness/config-management/chef-solo.log: 

Cause	The reason can be because the target host is unable to communicate to the Admin Server on port 53 as it is attempting to use the dnsmasq service on the Admin Server to resolve, in this case, 889e5752-6ae3-4286-a944-c182 33f4ccbc. This is the salt minion id of the admin server. You can see this by running "cat /etc/salt/minion" on the Admin Server to compare. Example output:
	`[root@S5-NWAdmin ~]# cat /etc/salt/minion`
	`master: localhost`
	`hash_type: sha256`
	`log_level: info`
	`id: 889e5752-6ae3-4286-a944-c18233f4ccbc`

Solution	If possible, configure any firewalls between the target host and the Admin Server host to be able to communicate on port 53. If this is not possible, the workaround is to include the minion id in the /etc/host file on the component hosts and starting in the 11.4 release, modify the chef recipe not to overwrite this workaround.
Workaround	Refer to Install/Upgrade fails in RSA NetWitness Platform because Resolv::ResolvError: no address for a particular host KB Article.

Error Message	Received an error in the error log similar to the following when trying to update to version 11.6 : `FATAL: Chef::Exceptions::Package: yum_package[rsa-protobufs-rt] (rsa-audit::packages line 11) had an error: Chef::Exceptions::Package: No version specified, and no candidate version available for rsa-protobufs-rt`
Cause	Custom builds/rpms installed for certain components installed on hosts, such as in the case of installing Hotfixes.

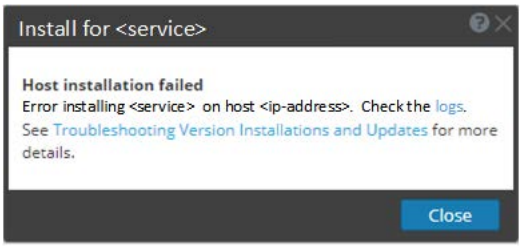
Solution	To resolve the issue, follow the below steps.
	1. SSH to Admin Server.
	2. Locate the component descriptor file by running the following command.
	`cd /etc/netwitness/component-descriptor/`
	3. Open the component descriptor file by running the following command.
	`vi nw-component-descriptor. json`
	4. Search for “packages” section for the component you have custom build/rpm. For example, below shown is the package details for “concentrator” host that has custom build/rpm.
	`“concentrator”: {`
	`“cookbook_name”: “rsa-concentrator”,`
	`“service_names”: [“rsa-nw-concentrator”],`
	`“family”: “launch”,`
	`“default_port”: xxxx, “description”: “Concentrator”,`
	`“packages”:[{ “name”: “rsa-nw-concentrator”,`
	`“version” : “11.6.0.0-2003001075220.5.cecf24b.e.17.centos”`
	`},`
	5. Delete the complete version details including (,) character in the packages section. For

	example, it should look like as shown below after you delete the version details.
	`"packages": [{`
	`"name": "rsa-nw-concentrator"`
	`},`
	Note: You must delete the version details for all the host that has custom builds/rpms in the component descriptor of the admin server.
	6. Run the upgrade process again.

External Repo Update Error

Error Message	Received an error similar to the following error when trying to update to a new version from the : `.Repository 'nw-rsa-base': Error parsing config: Error parsing "baseurl = 'https://nw-node-zero/nwrpmrepo /<version-number>/RSA'": URL must be http, ftp, file or https not ""`
Cause	There is an error the path you specified.
Solution	Make sure that: • the URL does exist on the NW Server host. • you used the correct path and remove any spaces from it.

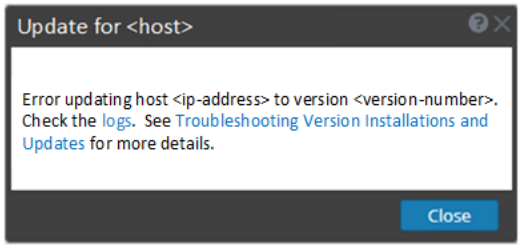
Host Installation Failed Error

Error Message	
Problem	When you select a host and click **Install** the install service process fails.

Solution	1. Try to install the service again.
	Often this is all you need to do.
	2. If you still cannot install the service:
	a. Monitor the following logs on NW Server as it progresses (for example, submit the `tail -f` command string from the command line'):
	`/var/netwitness/uax/logs/sa.log`
	`/var/log/netwitness/orchestration-server/orchestration-server.log`

	`/var/log/netwitness/deployment-upgrade/chef-solo.log`
	`/var/log/netwitness/config-management/chef-solo.log`
	`/var/lib/netwitness/config-management/cache/chef-stacktrace.out`
	The error appears in one or more of these logs.
	b. Try to resolve the issue and reinstall the service.
	• Cause 1 - Entered the wrong `deploy_admin` password in the nwsetup-tui.
	Solution - Reset your `deploy_admin` password password.
	1. On the NW Server host and all other hosts on 11.x, run the following command.
	`/opt/rsa/saTools/bin/set-deploy-admin-password`
	2. On the host that failed installation or orchestration, run the `nwsetup-tui` command and use the new **deploy_admin** password in response to the **Deployment Password** prompt.
	• Cause 2 -The `deploy_admin` password has expired.
	Solution - Reset your `deploy_admin` password password.
	1. On the NW Server host and all other hosts on 11.x, run the following command.
	`/opt/rsa/saTools/bin/set-deploy-admin-password`
	2. On the host that failed installation or orchestration, run the `nwsetup-tui` command and use the new **deploy_admin** password in response to the **Deployment Password** prompt.
	3. If you still cannot apply the update, gather the logs from step 2 and contact Customer Support.

Host Update Failed Error

Error Message	
Problem	When you select an update version and click **Update** > **Update Host**, the download process is successful, but the update process fails.
Solution	1. Try to apply the version update to the host again. Often this is all you need to do. 2. If you still cannot apply the new version update:

a. Monitor the following logs on NW Server as it progresses (for example, run the `tail -f` command from the command line):
`/var/netwitness/uax/logs/sa.log`
`/var/log/netwitness/orchestration-server/orchestration-server.log`
`/var/log/netwitness/deployment-upgrade/chef-solo.log`
`/var/log/netwitness/config-management/chef-solo.log`
`/var/lib/netwitness/config-management/cache/chef-stacktrace.out`
The error appears in one or more of these logs.

b. Try to resolve the issue and reapply the version update.

- Cause 1 - `deploy_admin` password has expired.
  Solution - Reset your `deploy_admin` password .
  Complete the following steps to resolve Cause 1.

  1. In the NetWitness Suite menu, select  **(Admin) > Security > Users** tab.
  2. Select the `deploy_admin` and click **Reset Password**.
  3. (Conitional) If NetWitness Suite does not allow you to expired `deploy_admin` password in the **Reset Password** dialog, complete the following steps.
     a. Reset `deploy_admin` to use a new password.
     b. On all non-NW Server hosts on 11.x , run the following command using the matching `deploy_admin` password from NW Server host.
     `/opt/rsa/saTools/bin/set-deploy-admin-password`

- Cause 2 -The `deploy_admin` password was changed on NW Server host but not changed on non-NW Server hosts.
  Complete the following step to resolve Cause 2.

  - On all non-NW Server hosts on 11.x , run the following command using the matching `deploy_admin` password from NW Server host.
  `/opt/rsa/saTools/bin/set-deploy-admin-password`

3. If you still cannot apply the update, gather the logs from step 2 and contact Customer Support.

Missing Update Packages Error

Error Message	**Initialize Update for Version xx.x.x.x** Missing the following update package(s) Download Packages from RSA Link

Problem	Missing the following update package(s) is displayed in the **Initialize Update Package for RSA NetWitness Platform** dialog when you are updating a host from the **Hosts** view offline and there are packages missing in the staging folder.
Solution	1. Click Download Packages from RSA Link in the **Initialize Update Package for RSA NetWitness Platform** dialog. The NetWitness Community page that contains the update files for the selected version is displayed. 2. Select the missing packages from the staging folder. The **Initialize Update Package for RSA NetWitness Platform** dialog is displayed telling you that it is ready to initialize the update packages.

OpenSSL 1.1.x

Error Message	The following example illustrates an ssh error that can occur when the ssh client is run from a host with OpenSSL 1.1.x installed: `$ ssh root@10.1.2.3` **`ssh_dispatch_run_fatal: Connection to 10.1.2.3 port 22: message authentication code incorrect`**

Problem	Advanced users who want to ssh to a NetWitness Platform host from a client that is using OpenSSL 1.1.x encounter this error because of incompatibility between CENTOS 7.x and OpenSSL 1.1.x. For example:
	`$ rpm -q openssl`
	`openssl-1.1.1-8.el8.x86_64`

Solution	Specify the compatible cipher list on the command line. For example:
	`$ ssh -oCiphers=aes128-ctr,aes192-ctr,aes256-ctr root@10.1.2.3`
	`I've read & consent to terms in IS user agreement.`
	`root@10.1.2.3's password:`
	`Last login: Mon Oct 21 19:03:23 2019`

Patch Update to Non-NW Server Error

Error Message	The `/var/log/netwitness/orchestration-server/orchestration-server.log` has an error similar to the following error: **`API|Failure /rsa/orchestration/task/update-config-management [counter=10 reason=IllegalArgumentException::Version '11.x.x.n' is not supported`**
Problem	After you update the NW Server host to a version, you must update all non-NW Server hosts to the same version. For example, if you update the NW Server from 11.4.0.0 to 11.6.0.0 , the only update path for the non-NW Server hosts is the same version (that is, 11.6.0.0). If you try to update any non-NW Server host to a different version (for example, from 11.4.0.0 to an 11.4.x.x) you will get this error.
Solution	You have two options:

	• Update the non-NW Server host to 11.6.0.0, or • Do not update the non-NW Server host (keep it at its current version)

Reboot Host After Update from Command Line Error

Error Message	You receive a message in the User Interface to reboot the host after you update and reboot the host offline. 
Cause	You cannot use CLI to reboot the host. You must use the User Interface.
Solution	Reboot the host in the Host View in the User Interface.

Reporting Engine Restarts After Upgrade

Problem	In some cases, after you upgrade to 11.6 from versions of 11.x, such as 11.4, the Reporting Engine service attempts to restart continuously without success.
Cause	The database files for live charts, alert status, or report status may not be loaded successfully as the files may be corrupted.
Solution	To resolve the issue, do the following: 1. Check which database files are corrupted: Navigate to the file located at `/var/netwitness/reserver/rsa/soc/reporting-engine/logs/reporting-engine.log` and check the following blocks: • If the live charts db file is corrupted, the following logs are displayed: Attempt Failed!!! Clearing pending acquires. While trying to acquire a needed new resource, we failed to succeed more than the maximum number of allowed acquisition attempts (30). Last acquisition attempt exception: org.h2.jdbc.JdbcSQLException: File corrupted while reading record: null. Possible solution: use the recovery tool [90030-196] at org.h2.message.DbException.getJdbcSQLException(DbException.java:345) at org.h2.message.DbException.get(DbException.java:168) org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'chartSummaryDAOImpl': Invocation of init method failed; nested exception is com.rsa.soc.re.exception.ReportingException: java.sql.SQLException: Connections could not be acquired from the underlying database! • If the alert status db file is corrupted, the following logs are displayed:

Attempt Failed!!! Clearing pending acquires. While trying to acquire a needed new resource, we failed to succeed more than the maximum number of allowed acquisition attempts (30). Last acquisition attempt exception:

org.h2.jdbc.JdbcSQLException: File corrupted while reading record: null. Possible solution: use the recovery tool [90030-196]

at org.h2.message.DbException.getJdbcSQLException(DbException.java:345)

at org.h2.message.DbException.get(DbException.java:168)

org.springframework.beans.factory.UnsatisfiedDependencyException: Error creating bean with name 'alertStatusHandler': Unsatisfied dependency expressed through field 'alertExecutionStatsDAO'; nested exception is org.springframework.beans.factory.UnsatisfiedDependencyException: Error creating bean with name 'alertExecutionStatsDAOImpl': Unsatisfied dependency expressed through field 'persistedAlertExecutionStatsDAO'; nested exception is org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'persistedAlertExecutionStatsDAOImpl'

- If the report status db file is corrupted, the following logs are displayed:

```
org.h2.jdbc.JdbcSQLException: File corrupted while reading
record: null. Possible solution: use the recovery tool
[90030-196]
```

2. To resolve the live charts database file corruption, perform the following steps:
   a. Stop the Reporting Engine service.
   b. Move the `livechart.mv.db` file from `/var/netwitness/reserver/rsa/soc/reporting-engine/livecharts` folder to a temporary location.
   c. Restart the Reporting Engine service.

   Note: Some live charts data may be lost on performing the above steps.

   To resolve the alert status or report status database file corruption, perform the following steps:
   a. Stop the Reporting Engine service.
   b. Replace the corrupted db file with the latest `alertstatusmanager.mv.db` or `reportstatusmanager.mv.db` file from `/var/netwitness/reserver/rsa/soc/reporting-engine/archives` folder.
   c. Restart the Reporting Engine service.

   For more information, see the Knowledge Base article Reporting Engine restarts After upgrade to RSA NetWitness Platform 11.4.

Problem	After you upgrade to version 11.6, the Reporting Engine service does not restart.
Cause	The Reporting Engine service may not start due to any of the following reasons. - workspace.xml not updated. - Time is not converted properly in livechart h2 database. - JCR (Jackrabbit repository) is corrupted with primary key violation.
Solution	To resolve the issue, run the Reporting Engine Migration Recovery tool (`rsa-nw-re-migration-recovery.sh`) on the Admin Server where the Reporting Engine service is installed.

	Note: You can find the Reporting Engine Migration Recovery tool in the below location.
	`/opt/rsa/soc/reporting-engine-11.6.0.0-<Tag>/nwtools`
	1. SSH to Admin Server.
	2. Untar the RE (Reporting Engine) tool, run the following command.
	`tar -xvf rsa-nw-re-recovery-tool-bundle.tar`
	3. (Optional) If you want to untar the RE tool file in some other directory, you can create a directory and untar the RE tool. Run the following commands.
	`mkdir <NAME OF THE DIRECTORY>`
	`tar -xvf rsa-nw-re-recovery-tool-bundle.tar --directory <PATH OF THE DIRECTORY>`
	4. Run the script, run the following command.
	`./<PATH OF THE DIRECTORY>/rsa-nw-re-recovery-tool.sh`
	For more information, see the Knowledge Base article **Reporting Engine Migration Recovery Tool**.

Log Collector Service (`nwlogcollector`)

Log Collector installation logs are posted to `/var/log/install/nwlogcollector_install.log` on the host running the `nwlogcollector` service.

Error Message	`<timestamp>.NwLogCollector_PostInstall: Lockbox Status : Failed to open lockbox: The lockbox stable value threshold was not met because the system fingerprint has changed. To reset the system fingerprint, open the lockbox using the passphrase.`
Cause	The Log Collector Lockbox failed to open after the update.
Solution	Log in to NetWitness and reset the system fingerprint by resetting the stable system value password for the Lockbox as described in the "Reset the Stable System Value" topic under "Configure Lockbox Security Settings" topic in the *Log Collection Configuration Guide*.

Error Message	`<timestamp> NwLogCollector_PostInstall: Lockbox Status : Not Found`
Cause	The Log Collector Lockbox is not configured after the update.
Solution	If you use a Log Collector Lockbox, log in to NetWitness and configure the Lockbox as described in the "Configure Lockbox Security Settings" topic in the *Log Collection Configuration Guide*.

Error Message	`<timestamp>: NwLogCollector_PostInstall: Lockbox Status : Lockbox maintenance required: The lockbox stable value threshold requires resetting. To reset the system fingerprint, select Reset Stable System Value on the settings page of the Log Collector.`
Cause	You need to reset the stable value threshold field for the Log Collector Lockbox.
Solution	Log in to NetWitness and reset the stable system value password for the Lockbox as described in "Reset the Stable System Value" topic under "Configure Lockbox Security Settings" topic in the *Log Collection Configuration Guide*.

Error Message	Decoder tries to start capture events but fails. `Aug 27 01:44:41 nwdecoder NwDecoder[8052]: [Decoder] [failure] The PF_RING driver is not installed.`

Solution	To resolve the issue, do the following steps,
	1. SSH to the Decoder host.
	2. Run the following commands.
	`yum reinstall pfring*`
	`systemctl restart nwdecoder`

NW Server

These logs are posted to `/var/netwitness/uax/logs/sa.log` on the NW Server Host.

Problem	After upgrade, you notice that Audit logs are not getting forwarded to the configured Global Audit Setup; or, The following message seen in the `sa.log`. `Syslog Configuration migration failed. Restart jetty service to fix this issue`
Cause	NW Server Global Audit setup migration failed to migrate from 11.4.x.x or 11.5.x.x. to 11.6.0.0.
Solution	1. SSH to the NW Server. 2. Submit the following command. `orchestration-cli-client --update-admin-node`

Orchestration

The orchestration server logs are posted to `/var/log/netwitness/orchestration-server/orchestration-server.log` on the NW Server Host.

Problem	1. Tried to upgrade a non-NW Server host and it failed. 2. Retried the upgrade for this host and it failed again. You will see the following message in the `orchestration-server.log`. `"'file' _virtual_ returned False: cannot import name HASHES""`
Cause	Salt minion may have been upgraded and never restarted on failed non-NW Server host

Solution	1. SSH to the non-NW Server host that failed to upgrade.
	2. Submit the following commands.
	`systemctl unmask salt-minion`
	`systemctl restart salt-minion`
	3. Retry the upgrade of the non-NW Server host.

Reporting Engine Service

Reporting Engine Update logs are posted to to`/var/log/re_install.log` file on the host running the Reporting Engine service.

Error Message	`<timestamp> : Available free space in /var/netwitness/re-server/rsa/soc/reporting-engine [ ><existing-GB ] is less than the required space [ <required-GB> ]`
Cause	Update of the Reporting Engine failed because you do not have enough disk space.
Solution	Free up the disk space to accommodate the required space shown in the log message. See the "Add Additional Space for Large Reports" topic in the *Reporting Engine Configuration Guide* for instructions on how to free up disk space.

Event Stream Analysis

Problem	After upgrading to version 11.6, the ESA correlation server does not aggregate events from the configured data sources.
Error Message	`Invalid username or password at com.rsa.netwitness.streams.base.RecordSourceSubscription.run (RecordSourceSubscription.java:173)`
Solution	To resolve the issue, do the following steps. In the NetWitness user interface,

1. Go to  **(Configure) > ESA Rules**.
   ESA Rules panel is displayed with **Rules** tab open.
2. In the Rules tab options panel, under Deployments, select a deployment.
3. In the **Data Sources** section, select the data source and click  in the toolbar.
4. In the **Edit Service** dialog, type the password for that data source.
5. Click the **Test Connection** button to make sure that it can communicate with the ESA service and then click **OK**.

Note: Do the above procedure for all the configured data sources.

6. After you finish making changes to the deployment, click **Deploy Now** to redeploy the ESA rule deployment.

RSA NetWitness

Version 12.0

NetWitness Investigate Quick Start Guide



Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA Conference Logo, RSA, and other trademarks, are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

JuneYY

What Is NetWitness® Investigate?

NetWitness audits and monitors all traffic on a network. One type of service--a Decoder--ingests, parses, and stores the packets, logs, and endpoint data traversing the network. The configured parsers and feeds on the Decoder create *metadata* that analysts can use to investigate the ingested logs and packets. Another type of service, called a Concentrator, indexes and stores the metadata. NetWitness Investigate provides the data analysis capabilities in NetWitness, so that analysts can analyze packet, log, and endpoint data, and identify possible internal or external threats to security and the IP infrastructure.

About This Guide

This guide provides end-to-end guidelines for all members of the SOC team to configure NetWitness Investigate and to investigate log and network events. End-to-end guidelines for investigating endpoints and user entity behavior using NetWitness Investigate are provided in separate documents:

- *NetWitness Endpoint Quick Start Guide*
- *NetWitness UEBA Quick Start Guide*

NetWitness Platform 11.7 Documentation in NetWitness Community

NetWitness Platform product documentation is organized along functional lines. If you are looking for a specific guide or version, go to the a.

Use these links to view the NetWitness Platform documentation. Both links provide the same documentation, in these two formats:

- HTML Guides include the latest information for currently supported versions: NetWitness Platform Documentation.
- PDF Guides provide the information for a specific version: NetWitness Platform 11.7 PDFs.

Use these links to access documentation that is not related to a particular version of the software:

- Hardware setup guides: https://community.netwitness.com/t5/netwitness-platform-hardware/tkb-p/netwitness-hardware-documentation
- Documentation for Content such as feeds, parsers, application rules, and reports: https://community.netwitness.com/t5/netwitness-platform-threat/tkb-p/netwitness-threat-intelligence.

Getting Started

The following tasks can be performed in any sequence and are for the entire SOC team.

Description	References
	
View information about product updates, improvements, and known issues	*Release Notes for NetWitness Platform 11.7*
Understand how NetWitness Investigate works	"How NetWitness Investigate Works" in the *NetWitness Investigate User Guide*

Setup, Installation, or Upgrade

No special setup, installation, or upgrade tasks are required for Investigate; it is part of NetWitness Platform for Logs and Network. However, setup is required for several components with which NetWitness Investigate works if you plan to do this type of analysis. These tasks are for the Administrator, and the SOC Manager may want to understand the setup.

Description	References
	
Install and set up the Malware Analysis (standalone or service)	Malware Analysis Configuration Guide
Install and set up NetWitness Endpoint (standalone or service)	NetWitness Endpoint Quick Start Guide
Install and set up NetWitness UEBA (standalone or service)	NetWitness UEBA Quick Start Guide

System-Level Configuration

Administrators configure system-level preferences for NetWitness Investigate.The following tasks are for the administrator, and the tasks can be performed in any sequence. SOC Managers should understand the possible configuration options.

Description	References
	
Configure role-based access control (RBAC) for analysts who will be using Investigate. These components have permissions related to investigate: investigate (Navigate view and Legacy Events view), investigate-server (Events view), Malware (Malware Analysis view), Endpoint-broker-server, and Endpoint-server.	"Role Permissions" in the System Security and User Management Guide
Configure Investigate to limit content available for different user roles (preQueries).	"Verify Query and Session Attributes per Role" in the System Security and User Management Guide
Configure default settings and limits for NetWitness Investigate on a system level.	"Configure Investigation Settings" in the System Configuration Guide

User Preference Configuration

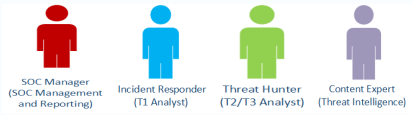
The following tasks are for Threat Hunters, Content Experts, and Incident Responders, and SOC Managers. The tasks can be performed in any sequence.

Description	References
	
Configure Navigate view and Events view preferences.	"Configure the Navigate and Legacy Events View" in the *NetWitness Investigate User Guide*
Configure Event Analysis view preferences.	"Configure the Events View" in the *NetWitness Investigate User Guide*
Configure the Malware Analysis view preferences.	"Configure Malware Analysis" in the *Malware Analysis User Guide*

Investigation

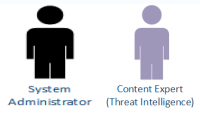
Different types of investigation may be handled by analysts with different skill levels and goals.

- Incident Responders (T1 Analysts) typically pivot to Investigate from NetWitness Respond to find detailed information about an incident so that they can respond to and remediate incidents.
- Threat Hunters (T2/T3 Analysts) typically peruse events, metadata, and raw content so that they can recommend issues for remediation and remediate issues.
- Content Experts (Threat Intelligence) typically peruse events, metadata, raw content, user and host data, and UEBA data so that they can investigate new threat intelligence, evaluate and create new feeds, and create correlation rules to flag indicators of compromise.
- SOC Managers need to understand the use cases.

Description	References
	
Learn about practical use cases	"Sample Use Cases for NetWitness Investigate" in the *NetWitness Investigate User Guide*
Investigate metadata and raw events in logs and network traffic	"Beginning an Investigation" in the *NetWitness Investigate User Guide*
Investigate possible malware	Malware Analysis User Guide
Investigate endpoints	NetWitness Endpoint User Guide
Perform user and entity behavior analysis	NetWitness UEBA User Guide

Maintenance

The administrator can perform the following tasks in any sequence.

Description	References
	
Maintain the list of queries and analyze the query patterns of other users of the NetWitness Platform system.	"Maintaining Queries Using URL Integration" in the System Maintenance Guide

Description	References
Fine tune system-level configuration settings to improve performance or limit access to data.	"Verify Query and Session Attributes per Role" in the System Security and User Management Guide "Configure Investigation Settings" in the System Configuration Guide

NetWitness® Platform XDR

Version 12.0

NetWitness Investigate User Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

How NetWitness Investigate Works **14**

Metadata, Meta Keys, Meta Values, and Meta Entities 14

Triggers for an Investigation 15

Workflow of an Investigation 15

Focus on Metadata, Query, and Time 21

Focus on Respond View Incidents and Alerts 23

NetWitness Investigate Views 23

Navigate View 23

Events View 24

Legacy Events View 26

Contextual Information for an Event 27

Reconstruction and Event Analysis 29

Configuring NetWitness Investigate Views and Preferences **31**

Configure the Navigate View and Legacy Events View 32

Configure Common Settings 32

Access the Navigate View and Legacy Events View Settings 33

Calibrate Navigate View Value Loading Parameters 34

Configure Navigate View and Legacy Events View Parameters 35

Configure the Default Log Export Format 36

Configure the Default Meta Value Export Format 36

Calibrate Legacy Events View Retrieval and Default Reconstruction 37

Enable or Disable Cascading Style Sheet Rendering in Web Content Reconstructions 37

Configure Search Options 38

Configure the Events View 39

Set the Default Investigate View 39

Set User Preferences for the Events View 40

Beginning an Investigation **44**

Focus on Metadata, Raw Events, and Event Analysis 44

Focus on Hosts and Files 45

Focus on Risky User and Entity Behavior 45

Focus on Scanning Files for Malware 45

Begin an Investigation in the Navigate or Legacy Events View 46

Begin an Investigation (No Default Service) 46

Set or Clear the Default Service 48

Begin an Investigation (Default Service Specified) 49

Change the Service or Collection to Investigate 50
Investigate Workbench Restoration Collections 51
Begin an Investigation in the Events View 53
Access the Events View 53
Refining the Results Set 56
Use Meta Groups to Focus on Relevant Meta Keys 57
Live Meta Groups 57
Default Meta Keys Group (Version 11.5 Events View) 58
Custom Meta Groups 59
Work with Meta Groups in the Events View (Version 11.5 and Later) 60
View the Meta Keys in a Meta Group 60
Select a Meta Group 62
Create a Custom Meta Group 62
Delete a Custom Meta Group 66
Edit a Custom Meta Group 67
Copy a Meta Group (Version 11.5 and Later) 68
Meta Group Folders 71
Create a Meta Group Folder 71
Edit and Move Meta Group Folder 72
Copy Meta Group Folder 73
Copying Meta Group Private or Shared Folders 73
Copy Meta Group Folder Deployed from Live 75
Delete Meta Group Folder 76
Work with Meta Groups in the Navigate View 77
Create a Meta Group and Add Meta Keys 77
Copy and Edit a Meta Group 81
Edit a Custom Meta Group 81
Delete a Meta Group 83
Export a Meta Group 83
Import a Meta Group 83
Use Columns and Column Groups in the Events List 85
Built-In Column Groups 92
Live Column Groups 93
Custom Column Groups 98
Filtering Folders 98
Dialogs for Managing Column Groups 99
Work with Columns and Column Groups in the 11.4 and Later Events View 101
Manually Select Columns to Display and Adjust Column Order and Width 101
Select a Column for Sorting Events in the Events Panel (Version 11.4) 102
Sorting by Column (Version 11.4.1 and Later) 102

Sorting by Column (Version 11.4) 104
View the Meta Keys Included in a Column Group 105
Select a Column Group 107
Create a Custom Column Group 109
Delete a Custom Column Group 113
Edit a Custom Column Group 116
Create a Copy of a Column Group (Version 11.5 and Later) 119
Create a Column Group Folder 123
Edit and Move Column Group Folder 123
Copy Column Group Folder 124
Copy Group Folder Deployed from Live 125
Delete Column Group Folder 126
Select a Column Group and Columns (11.3 and Earlier Event Analysis View) 127
Work with Column Groups in the Legacy Events View 128
Select a Column Group 128
Create a Custom Column Group in the Legacy Events View 129
Delete a Column Group (Legacy Events View) 131
Edit a Column Group (Events View) 132
Import and Export a Column Group (Legacy Events View) 135
Use Query Profiles to Encapsulate Common Areas for Investigation 137
Built-In Query Profiles 137
Live Query Profiles 138
Custom Query Profiles 139
Dialogs for Managing Query Profiles 139
View Query Profile Details (Events View) 142
Apply a Query Profile (Events View) 144
Create or Edit a Custom Query Profile (Events View) 145
Delete a Custom Query Profile (Events View) 148
Copy a Query Profile (Version 11.5 and Later) 150
Create a Query Profile Folder 153
Edit and Move Query Profile Folder 153
Copy Query Profile Folder 154
Copy Query Profiles Group Folder Deployed from Live 155
Delete Query Profile Folder 155
Add Springboard Panels from Events View 156
Navigate to the Manage Profiles Dialog (Navigate and Legacy Events Views) 158
Create, Edit, or Delete a Profile Group (Navigate or Legacy Events View) 159
Create and Edit Profiles (Navigate or Legacy Events View) 161
Delete a Profile (Navigate or Legacy Events View) 162
Change the Active Profile (Navigate or Legacy Events View) 162

Import Profiles (Navigate or Legacy Events View) ... 163
Download Profiles (Navigate or Legacy Events View) ... 163
Drill into Metadata in the Events View ... 164
Modes of Operation ... 165
View Metadata in the Filter Events Panel ... 166
Show Max Value of Meta Groups ... 168
View the Context Lookup Panel in the Filter Events Panel ... 169
Understand Visible Metadata ... 170
Stop and Resume Metadata Loading ... 171
Close All Except One Meta Key ... 172
Set the Ordering Method for Meta Values ... 173
Drill into Meta Values ... 175
Copy the Meta Values for a Meta Key ... 177
View a Selected Meta Value in RSA Live ... 179
Append and Refocus the Investigation of a Meta Value in Unified Panel ... 180
Filter Results in the Events View ... 188
Initial Filter Using the Query Bar ... 188
Find a Text String in the Events Panel (Version 11.4 and Later) ... 189
Refining the Results in the Events Panel ... 190
Query Builder Concepts ... 191
Guided Mode vs. Free-Form Mode ... 192
Concepts for Editing Multiple Filters ... 193
The Version 11.3 and Earlier Query Builder ... 195
The Version 11.4 Query Builder ... 195
Meta Keys Cached for Faster Loading ... 195
Text Filter ... 196
Pasting Text Instead of Typing ... 196
Select All Filters and Copy All Filters (Version 11.4.1) ... 196
Use of Recent Queries ... 196
Use of Advanced Operators ... 196
Easy Use of AND/OR Operators ... 197
Automatically Balanced Parentheses ... 197
Hints about Available Values ... 197
CIDR Notation and Shorthand ... 198
Ranges or Series of Values ... 198
No Separating Space Required After Meta Key and Operator (Version 11.4.1) ... 198
Select a Time Range ... 198
Submit a Query ... 202
Cancel Execution of a Query ... 203
View Status of a Query ... 203

Build a Query in Guided Mode ... 205
Keyboard Actions to Use in Guided Mode ... 206
Visual Feedback in Guided Mode ... 208
Add a Simple Filter in Guided Mode ... 210
Add a Free-Form Filter in Guided Mode (Version 11.3 and Later) ... 214
Add a Text Filter to Find a Value Anywhere in the Data Set (Version 11.4 and Later) ... 216
Select All Filters and Copy All Filters in the Query Bar (Version 11.4.1 and Later) ... 219
Paste Text in the Query Bar (Version 11.4 and Later) ... 219
Insert a Filter Based on a Recent Query (Version 11.4 and Later) ... 220
Edit a Filter in Guided Mode ... 222
Query Using Selected Filters in Guided Mode ... 222
Delete a Filter and Delete Text or Parentheses in a Filter in Guided Mode ... 223
Create a Query in the Free-Form Mode ... 224
Filter Results in the Navigate View ... 226
Set the Time Range ... 226
Set the Quantification Method and Sort Sequence of Meta Key Results ... 228
Manage and Apply Default Meta Keys in an Investigation ... 229
Drill into Data in the Navigate View Time Chart ... 231
Drill into Data in the Values Panel ... 232
Filter Results in the Legacy Events View ... 240
Filter Events Displayed in the Legacy Events View ... 240
Page Through Events in the Legacy Events View ... 241
Create a Query in the Navigate and Legacy Events Views ... 242
Create a Query Using the Basic Method ... 242
Create a Query Using the Advanced Method ... 243
Apply a Recent Query ... 245
Search for Text Patterns in the Navigate and Legacy Events Views ... 247
Keyword Text Search ... 247
Options Controlling Search Behavior ... 248
Regular Expression Search Syntax ... 250
Raw Text Keyword Search ... 250
Search Procedures ... 250
Search in the Navigate View ... 250
Search in the Legacy Events View ... 250
View and Modify Queries Using URL Integration ... 251
Service Id Known ... 251
Host and Port Known ... 251
Examples ... 252
Additional Notes ... 252

Reconstructing and Analyzing Events **253**
Examine Event Details in the Events View 256
Event Details for Each Event Type 256
Text Reconstruction 257
Packet Reconstruction 260
File Reconstruction 262
Host Information 263
Email Reconstruction 265
Analyze Events in the Events View 268
How Results Are Loaded and Sorted 268
Actions to Refine the Events List 269
Actions to Analyze Events 269
Select the Analysis Type for an Event 270
Adjust the Display of Requests and Responses 270
View Associated Metadata for an Event 270
Show or Hide the Event Header 274
Page Through Events in the Packet and Text Tabs 274
Expand Truncated Text Entries in the Text Tab 275
Perform URL and Base64 Encoding and Decoding in the Text Tab 275
View Decompressed Text in an HTTP Network Session in the Text Tab 277
View a JSON String in Tree Format in the Text Tab (Version 11.5.1) 278
Use the Payload Only Option in the Packet Tab 279
View Highlighted Bytes in the Packet Tab 280
Highlight Common File Types in the Packet Tab 281
Reconstruct an Event in the Legacy Events View 283
Reconstruct an Event Using an Event ID 283
Reconstruct an Event from a Drill Point in the Navigate View 284
View Side by Side or Top to Bottom 286
Select Event Information to View 286
Select Event Reconstruction Type 286
Open or Download an Email Attachment 287
Export an Event as a PCAP File 287
Extract Files from a Reconstructed Event 288
Look Up Additional Context for Results 289
Open the Context Lookup Panel 290
Add an Entity to a Whitelist 293
Create a List (Events View) 293
Pivot to Investigate > Navigate (Events View) 294
Pivot to Archer (Events View) 294
Pivot to NetWitness Endpoint Thick Client (Events View) 295

View the Context Lookup Panel in the Navigate View or Legacy Events View 295
Add Meta Values to an Existing List (Navigate and Legacy Events Views) 296
Remove a Meta Value from a Context Hub List (Navigate and Legacy Events Views) 297
Create a New List (Navigate and Legacy Events Views) 297
Launch a Lookup of a Meta Key 299
Launch an Endpoint Thick Client Lookup in the Events View 299
Launch an Endpoint Thick Client Lookup in the Navigate View 300
Perform Lookups of Meta Values in Events 302
Launch Other External Lookups from the Navigate View 304
Launch a Malware Analysis Scan from the Navigate View 305
Group Events from Split and Related Sessions in the Events and Legacy Events Views 307
Split Network Session 308
Session Size and Time Split 308
Transaction Handling Split 309
Session Fragments Highlighting 310
Related Network Session 310
Use Cases for Viewing Events from Split and Related Network Sessions 311
Show and Hide Relationships in the Events List 312
Session Fragments Highlighting (11.3 Events View) 313
Find and Combine Fragments in the Legacy Events View 314
Visualize Metadata as Parallel Coordinates 317
Best Practices for Effective Parallel Coordinates Charts 317
RSA Meta Groups for Parallel Coordinates Use Cases 318
View a Parallel Coordinates Visualization 318
Select Meta Keys for a Parallel Coordinates Visualization 319
Optimize a Parallel Coordinates Visualization 324
Sample Use Case 326
Sample Visualization of a Large Data Set 327
Visualize the Current Drill Point in Informer 329
Downloading and Acting Upon Results 330
Download Data in the Events View 331
Download Events or Metadata in the Events Panel 331
Download a Log in the Text Reconstruction 336
Download Network Event Data in the Text or Packet Reconstruction 338
Download Files from a Network Event in the File Reconstruction 340
Download Attachments from an Email Reconstruction 342
Export or Print a Drill Point in the Navigate View 345
Export Events in the Legacy Events View 347
Add Events to an Incident in the Events View 348
Add Events to an Incident in the Legacy Events View 350

Troubleshooting NetWitness Investigate **352**
Navigate View and Legacy Events View Issues 352
Events View Issues 353
Investigate Reference Materials **359**
Add Events to an Incident Dialog 360
Workflow 360
What do you want to do? 360
Related Topics 361
Quick Look 362
Add/Remove from List Dialog 364
Workflow 364
What do you want to do? 365
Related Topics 365
Quick Look in the Events View 366
Quick Look in the Navigate and Legacy Events Views 367
Column Groups Dialogs 369
Related Topics 370
Quick Look - Column Group Menu, Create Column Group Dialog, and Column Group Details Dialog 370
Quick Look - Manage Column Groups Dialog 372
Context Lookup Panel 375
Workflow 375
What do you want to do? 375
Related Topics 376
Quick Look (in the Navigate and Legacy Events Views) 377
Incidents 378
Alerts 378
Lists 378
Endpoint 379
Quick Look in the Events View 379
Lists Tab 382
Archer Tab 383
Active Directory Tab 384
NetWitness Endpoint Tab 385
Alerts Tab 387
Incidents Tab 388
File Reputation Tab 389
TI Tab 390
REST API Tab 391
Create an Incident Dialog 393

Workflow ... 393
What do you want to do? ... 393
Related Topics ... 394
Quick Look ... 394
Events View ... 396
Workflow ... 396
What do you want to do? ... 396
Related Topics ... 397
Quick Look ... 398
Filter Events Panel ... 401
Query Console ... 402
Timeline ... 403
Quick Look for Version 11.0.0.x (End-Of-Life) ... 404
Events View - Email Tab ... 406
Workflow ... 406
Related Topics ... 406
Quick Look ... 406
Events View - File Tab ... 408
Workflow ... 408
What do you want to do? ... 408
Related Topics ... 409
Quick Look ... 409
Events View - Host Tab ... 411
Workflow ... 411
What do you want to do? ... 411
Related Topics ... 412
Quick Look ... 412
Events View - Packet Tab ... 414
Workflow ... 414
What do you want to do? ... 414
Related Topics ... 415
Quick Look ... 415
Events View - Text Tab ... 417
Workflow ... 417
What do you want to do? ... 417
Related Topics ... 418
Quick Look ... 418
Investigate Dialog ... 420
Workflow ... 420
What do you want to do? ... 420

Related Topics ... 421
Quick Look ... 422
Investigation Tab - User Preferences Panel ... 424
Related Topics ... 424
Quick Look ... 424
Investigate View ... 428
Legacy Event Reconstruction View ... 429
What do you want to do? ... 429
Related Topics ... 430
Quick Look ... 430
Legacy Events View ... 433
What do you want to do? ... 433
Related Topics ... 434
Quick Look ... 434
Detailed Description ... 437
Manage Default Meta Keys Dialog ... 440
Related Topics ... 440
Quick Look ... 440
Meta Groups Dialogs ... 442
Related Topics ... 442
Quick Look - Meta Groups Menu, Create Meta Group Dialog, and Meta Group Details Dialog ... 442
Quick Look - Manage Meta Groups Dialog ... 445
Navigate View ... 448
Workflow ... 448
What do you want to do? ... 449
Related Topics ... 449
Quick Look ... 450
Toolbar ... 450
Pause/Reload Button and Breadcrumb ... 453
(Optional) Debug Information ... 454
Time Banner ... 454
Visualizations ... 454
Timeline Chart ... 454
Parallel Coordinates Chart ... 456
Values Panel ... 458
Values Panel Loading Behavior ... 459
Iterative results ... 460
Partial results ... 460
Debug Information ... 460
Load Complete ... 461

Query Dialog 462
What do you want to do? 462
Related Topics 463
Quick Look 463
Simple View 463
Advanced View 464
Recent View 464
Query Profiles Dialogs 466
Related Topics 466
Quick Look - Query Profile Menu, Create Query Profile Dialog, and Query Profile Details Dialog 467
Quick Look - Manage Profiles Dialog 469
Generate Springboard Panel Dialog 472
What do you want to do? 472
Related Topics 472
Quick Look - Generate Springboard Dialog 472
Settings Dialogs for Investigate Views 475
Related Topics 475
Quick Look 475
Navigate View Settings Dialog 476
Legacy Events View Settings Dialog 477
Events View Preferences Dialog 479

How NetWitness Investigate Works

NetWitness Investigate provides analysts the means to analyze events that have been captured by NetWitness. Using Investigate, analysts can examine packet, log, and endpoint data, and identify possible internal or external threats in their environment. There are several different views available to analysts to gain different perspectives into the data in their environment. A key element that all the views have in common is metadata.

Note: In Version 11.1 and later, the Hosts and Files views provide a view into endpoint data. Earlier versions offer access to endpoint data using a standalone NetWitness Endpoint server.

Metadata, Meta Keys, Meta Values, and Meta Entities

NetWitness audits and monitors all data communications in an environment. One type of service--a Decoder--ingests, parses, and stores the original packets captured on the network, logs forwarded by a device, and endpoint events seen by the endpoint agent. The configured rules, parsers, and feeds on the Decoder create *metadata* that analysts can use to investigate the ingested logs, packets, and endpoint data. Another type of service, called a Concentrator, indexes and stores the metadata, making it more efficient to search through all types of metadata.

The metadata is created to give analysts valuable points of reference associated with the original data. This allows analysts to quickly get a sense of what has transpired without being required to examine every detail of an event. The metadata is in the form of a *meta key* and *meta values* for the key. For example, `ip.src` is a meta key, and an IP address (192.168.1.1) that is the source of the traffic is a meta value tagged as `ip.src`. When you view data in Investigate, you see the meta key `ip.src` and all of the IP addresses (meta values) that are tagged with that key. Some meta keys are built-in and others may be custom keys specific to your environment and defined by the administrator. All metadata, no matter the source of the data, is normalized into the Unified Data Model for NetWitness Platform to keep similar metadata concepts grouped together into like meta keys (see https://community.netwitness.com/t5/netwitness-platform-unified-data/tkb-p/netwitness-udm).

Meta entities are available in Version 11.1 and later. A *meta entity* is an alias that groups together the results from other meta keys. Meta entities organize similar meta keys into a single, easier to use, meta type. For example, the default Core database language includes distinct meta keys for IP source and IP destination. One of the built-in meta entities named `ip.all` represents the combined set of all IP sources and destinations. Some meta entities are already included by default, and the administrator can create custom meta entities. Analysts can use a meta entity in a query, a meta group, a column group, and a query profile. Parallel coordinates visualizations do not support meta entities. Administrators can use meta entities to define a query prefix to apply to a user role and a user as described in the *System Security and User Management Guide*. The *Decoder Configuration Guide* provides additional information about creating meta entities and how they can be used in rules.

Note: Meta entities need to be configured on all upstream Concentrators. If any Concentrator does not have a meta entity configured, that meta entity will be empty when you query the Broker.

Analysts usually query the Broker or Concentrator to discover threats. The Concentrator handles queries, only going to the Decoder for raw logs or endpoint events or a full reconstruction of network events. ESA, Malware Analysis, and Reporting Engine also query the Concentrator, where they can quickly get all the pertinent metadata associated with an event and generate information about the event without having to query each Decoder. In some special cases, analysts may query a Decoder.

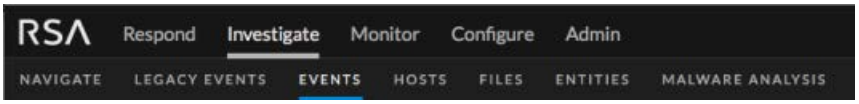
Triggers for an Investigation

These are a few examples of triggers for an investigation:

- You receive intelligence about a new active directory hack. Starting in the Events view, you use that intelligence to run a search across all of your raw Active Directory log data for the last 24 hours.
- You are asked by the SOC manager to find any Pokemon Go malware due to its popularity. Starting in the Navigate view, you craft a query to look for an HTTP session using a specific user agent related to the malware that your SOC manager found on a security blog.
- An incident responder escalates a ticket that shows some odd indicators related to a host. Starting in the Hosts view, you examine that host to find specific details.
- You are looking for the next zero day attack and start drilling into the network metadata in the Navigate view (or the Filter Events panel in the Events view) to find any abnormal automated sessions leaving the enterprise.
- You are asked by your SOC manager to find any information related to user `jarvis`, an employee who was just let go. Starting in the Users view you can filter for that username to make sure there is no longer any activity for that user and see if that user deviated from their typical behavior prior to being let go.
- A phishing attack detected has an associated attachment, and you want to know what devices in your environment have seen that file by searching for the file hash in the Files view.
- A malicious file has been automatically found in your environment, and you want to review the static and dynamic analysis done on that file along with how many systems it has been transmitted to or from. Starting in Investigate > Malware Analysis you can see the analysis results.

Workflow of an Investigation

Analysts can investigate data captured by NetWitness, and deep dive from information on a NetWitness dashboard, the Springboard (Version 11.5 and later), a NetWitness Respond incident or alert, a report created by the NetWitness Reporting Engine, or a third-party application. During the course of an investigation, analysts can move seamlessly between views: the Navigate view, the Events view, the Legacy Events view, the Hosts view, the Files view, the Users (Entities) view, and the Malware Analysis view. By default, the Navigate view and the Legacy Events view are disabled. This figure illustrates the NetWitness Investigate submenus in Version 11.4 and earlier.



This figure illustrates the Version 11.5 menu, which is optimized for the analyst workflow with Users, Hosts, and Files moved to the top-level menu.



This figure illustrates the Version 12.0 menu, which is optimized with new NetWitness logo.



Note:
- The Files and Hosts views are available in Version 11.1 and later. Before Version 11.5, they were a submenu of Investigate.
- The Users view is available in Version 11.2 and later; in Version 11.4 it is labeled Entities view. Before Version 11.5, it was submenu of Investigate.
- By default, the Legacy Events view is disabled in Version 11.4, but can be enabled by an administrator as described in the *System Configuration Guide*.
- By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.
- Specific user roles and permissions are required for a user to conduct investigations and malware analysis in NetWitness. If you cannot see a view, the administrator may need to adjust the roles and permissions configured for you.

The views are closely integrated to reduce the need to jump from one view to another. Each use case determines the starting point for your investigation; every situation is unique in terms of the types of information you are attempting to find. Many investigations start in one view, and end in a different view as you learn something and then need to follow that result to a different line of questioning. Experienced analysts frequently begin an investigation in the Navigate view or the Events view. Less experienced analysts can begin in the Dashboard, Respond view, or the Springboard (Version 11.5 and later), in which clickable incidents and alerts link to detailed information and analyses in the different views.

Go to...	Focus
Navigate view	The Navigate view (Version 11.5 and earlier) works well for providing a high-level vantage point of what has been seen in your environment for a specific time range by showing meta keys and meta values for log, endpoint, and network events. After drilling into meta values, go to the Events view to see the raw event (See Filter Results in the Navigate View.)

Go to...	Focus
Events view	The Events view (default Investigate view) is the workflow for analysts interacting with events, presenting different facets of the same data in adjacent panels. In Version 11.6, there is no need to go to the Navigate view to drill into meta values as the Filter Events Panel in the Events view provides this functionality. (See Refining the Results Set, Reconstructing and Analyzing Events, and Downloading and Acting Upon Results.)

Go to...	Focus
Legacy Events view	The Legacy Events view was the original workflow for looking at event details in Version 11.0 to 11.3.x.x. The Legacy Events is replaced by the 11.4 and later Events view and it is hidden unless the administrator enables it. (See Refining the Results Set, Reconstructing and Analyzing Events, and Downloading and Acting Upon Results.)

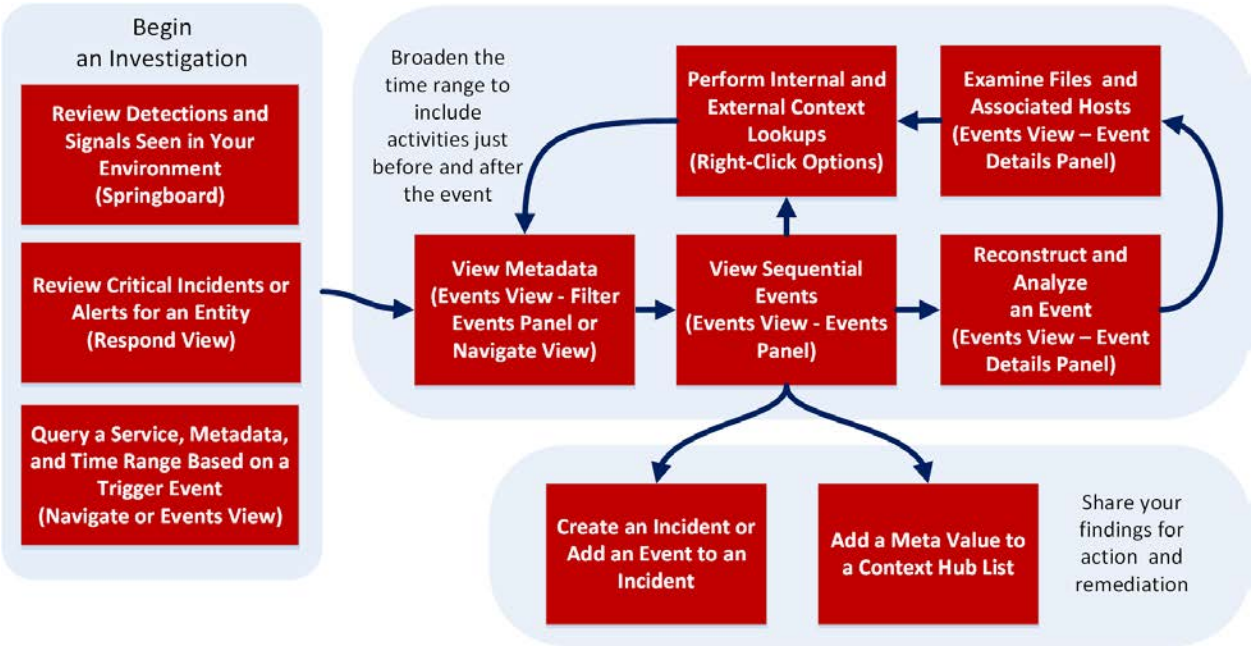
Go to...	Focus
Hosts view	The Hosts view option was moved to the main menu in Version 11.5. Hosts on which the NetWitness Endpoint agents are running are listed. For every host, you can view processes, drivers, DLLs, files (executables), services, anomalies, and autoruns that are running, and information related to logged-in users. (See the *NetWitness Endpoint User Guide*.)
Files view	The Files view option was moved to the main menu in Version 11.5. Unique files in your deployment, such as PE, Macho, and ELF, are listed. For each file, you can view details such as file name, reputation status, file status, risk score, signature, checksum, and others. (See the *NetWitness Endpoint User Guide*.)

Go to...	Focus
Malware Analysis view	If you are running a Malware Analysis appliance, you can scan files and see results of four types of analysis: network, static, community, and sandbox. If a file is malware, you can go to the Hosts view to see which hosts downloaded the file. (See the *Malware Analysis User Guide*.)

Go to...	Focus
Users view	The Users view (labeled as the Entities view in Version 11.4) was moved to the main menu in Version 11.5. It provides visibility into risky user behaviors across your enterprise using NetWitness UEBA. You can view a list of high-risk users and a summary of the top alerts for risky behavior for your environment, and then select a user or an alert and view details about the risky behavior, and a timeline during which the behaviors occurred. NetWitness Platform users assigned the Administrators or UEBA Analysts role have access to this view. (See the *NetWitness UEBA User Guide*.)

Focus on Metadata, Query, and Time

The following figure depicts the workflow for an investigation with focus on metadata, a query, and time range.



Analysts use Investigate to hunt for events that drive the incident response workflow and to do strategic analysis after another tool has generated an event. Beginning in the Navigate view, Events view, or Legacy Events view:

1. Start by executing a query on a service for a specific time range, filter results to get a subset of events, reconstruct or analyze an event, and repeat the process to reconstruct or analyze another event. Built-in query profiles, meta groups, and column groups provide a good starting point. For example, you can choose the RSA Email Analysis query profile to see only metadata that is useful when investigating email risks.
2. When an event bears a closer look, view the context around the event, and decide whether to create an incident or add the event to an incident. If you decide not to add the event to an incident, you can run another query to gain further insight, which starts again at the beginning of the workflow.
3. If you notice suspicious activity or files on a specific host in the network, gather additional information about the host and files found on the host in the Hosts and Files view, or in a standalone NetWitness Endpoint server.
4. If you find a file or event that potentially contains malware, do a Malware Analysis scan of the file or open Malware Analysis and start a scan of the service on which the event was seen.

Here is one simple use case: If there is a concern regarding suspicious traffic with certain countries, the Destination Country meta key reveals all destinations and the frequency of the contact. Drilling into those values yields the specifics of the traffic, such as the IP address of the originator and the recipient. Checking other metadata can expose the nature of attachments exchanged between the two IP addresses. When suspect IP addresses are identified, looking at the addresses in the Navigate view or Events view with a broader time range can provide clues about what happened before and after the event being investigated.

Another use case is to investigate an alert to discover a malicious insider in the network who is exfiltrating intellectual property or other sensitive data from a specific IP address. The investigation begins with this meta value: `Upload without change request followed by download alert`. Start in the Navigate view or Events view by filtering the values to the IP address during the time range in which the alert was generated. Alerts metadata shows risk indicators as meta values, and you can click on different meta values to filter the events list and then reconstruct the event. Next extract files and examine the files to understand what happened. With this information, you can filter on the same IP address and broaden the time range to see activities before and after the event.

Focus on Respond View Incidents and Alerts

An analyst who is working on an incident or an alert in Respond can open the incident in Investigate to do a deeper analysis of the event or alert.

- The workflow to respond to an incident typically begins in the Respond view, where the analyst who is investigating an incident needs to gather intelligence about the incident in Investigate. Hover over an underlined entity in an incident or alert, such as an IP address, and select the action **Pivot to Investigate**. The Events view opens and is filtered for the selected entity. Defined meta keys are queried and the captured packets, logs, and endpoint events are displayed in the Events view.
- If you find events that are relevant to the incident, add the events to the incident in Respond. You can also create a new incident in Respond based on one or more events found in Investigate.
- From the Incident Details view Indicators panel in Respond, open the Events view to get a better understanding of an indicator event.

NetWitness Investigate Views

This section provides a brief description and example of the Navigate view, Events view, and Legacy Events view, as well as the context information, event details and reconstructions available in those views. Refer to the *Malware Analysis User Guide* for information about features and functions of the Malware Analysis view.

Navigate View

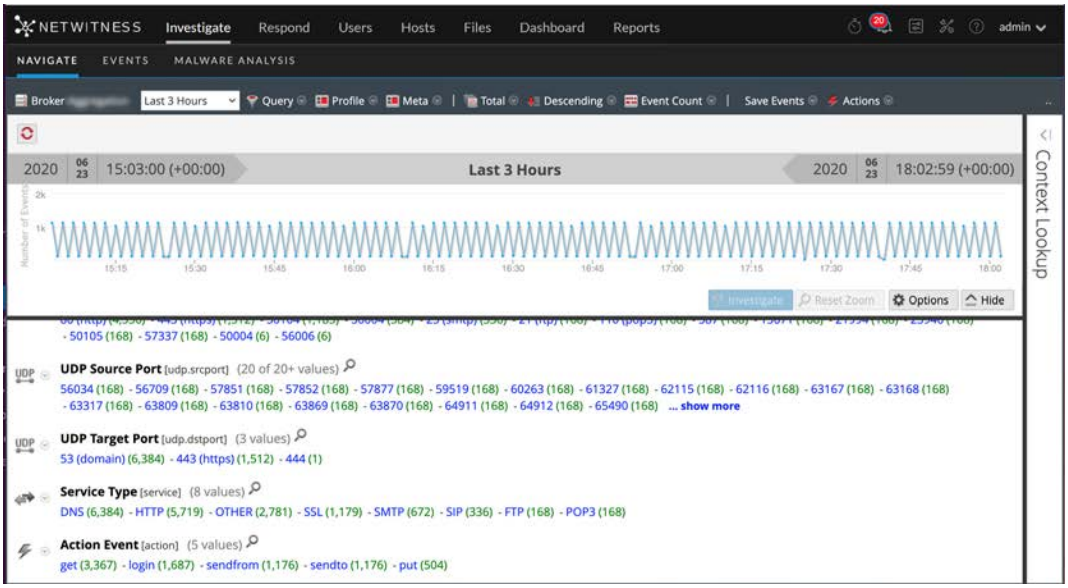
Note: By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.

The Navigate view provides the capability to drill into and query metadata for network, logs, and endpoint events on a Broker, Concentrator, or Decoder (though investigating a Decoder is not typical). For certain configured meta keys, such as IP address, or hostname, you can see additional context information around a value using the Context Hub. The Navigate view also provides a sequential visualization of the data in a timeline. This figure illustrates the Navigate view.

- Each meta key listed displays the top 20 values based on how many events have those values.
- Drilling into the meta values by successive left or right clicking of values applies each clicked value as an additional filter to your query. As you drill, the subset of metadata seen grows smaller based on

the filters you have applied. For example, if you filter to only show HTTP (service=80), all the remaining metadata presented will be what is contained inside those HTTP events.

- With a smaller set of refined results, you can go to the Events view to examine further event details, or perform other lookups inside or outside the platform to gain further insight.



Events View

The Events provides the ability to view events sequentially, analyze raw event data and metadata, and (in Version 11.5) drill into metatdata as you can in the Navigate view.

- **In the Events panel**, Network, endpoint, and log events are listed in order by time. You can view the raw event, filter, sort, search, look at details and reconstructions, and download events. Clicking an event opens the Event Details panel for the event. Different reconstructions are available in the Events view (packets, text, files, email, and web) with helpful cues to identify points of interest, such as interesting bytes, file types, and encoded data.

- **In the Event Meta panel**, you can view related metadata for an event that is open in the Event Details panel. Analysts reviewing the metadata can change the order of the metadata to better track down what they are looking for. The items in the list of metadata can be grouped by the sequence they were generated or alphabetically.

- **In the Filter Events panel** (Version 11.5 beta feature), you can drill into meta values for the listed events with your actions reflected in the Events panel. When the Filter Events panel is expanded to the full width of the browser, you can drill into meta values to hunt for specific information before listing the events in the Events panel (comparable functionality to drilling into data in the Navigate view).
  (Version 11.5.1) In the Filter Events panel, the meta values result threshold is 100000. If results are

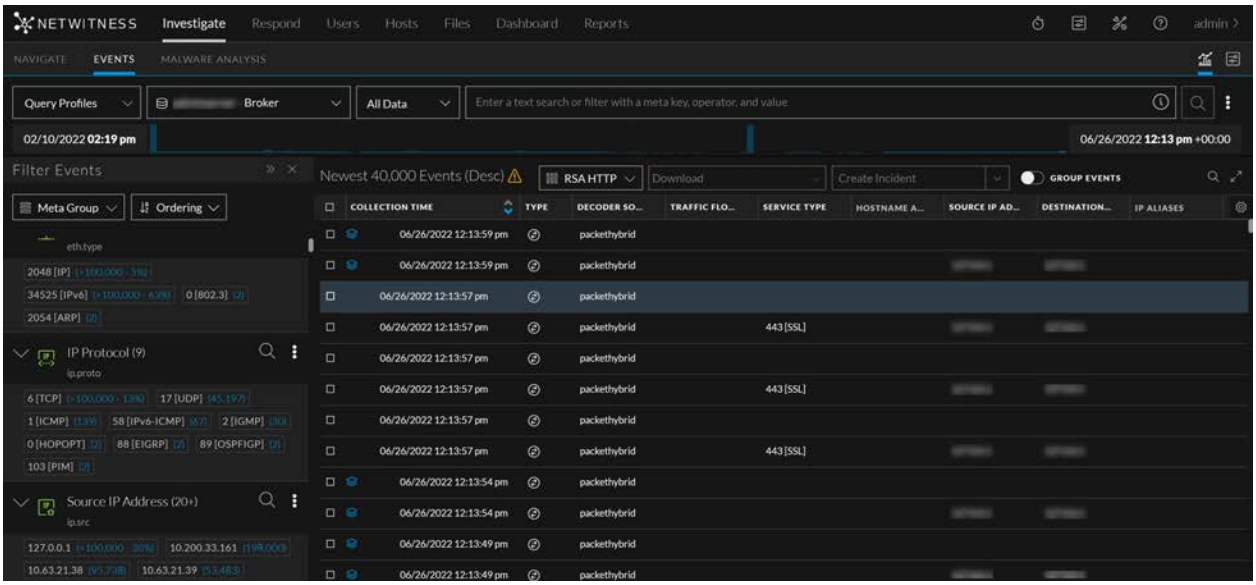
above the threshold, it is indicated using either ~ or >. For example, **(>100000)** indicates that the results are sorted based on count and are greater than the threshold. Similarly, **(~100000)** indicates that the results are sorted based on size and are greater than the threshold.

- **In the Event Details panel**, you can view details of a network, log, or endpoint event and safely reconstruct an event in a format similar to the original format. The tabs are Text, Packet, Files, Hosts, Users, Email, and Web.
- From various points in the Events view, you can pivot to standalone Endpoint, look up in Live, and do other internal lookups. External lookups allow you to search the internet for meta values with which you interacted, determine passive DNS information related to an IP address, check if a URL is blacklisted, and other third-party context integrations.
- (Version 11.5) If network data is enriched with endpoint data from an Endpoint deployment and the Endpoint Agent is configured for expanded network visibility, host information for network events is also displayed in the Events view in the header and in the Hosts tab.

  Note: For Expanded Network Visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Services Security View - Aggregation Role" in the *Hosts and Services Getting Started Guide for NetWitness Platform*.
- For certain configured meta keys, such as IP address, or hostname, you can search for additional context information around a value using the Context Hub. The additional context may include incidents, alerts, Threat Intelligence, and other sources where the value was mentioned.
- You can export different types of data. When viewing files, you can export files in a zip archive to your local file system. When viewing email reconstructions, you can download attachments. You can download logs from the text reconstruction, and export packets from the packet reconstruction. You can download multiple events from the Events list.
- (Version 11.6) To view the download jobs, click . Unlike the  icon in the **Legacy Events** and **Navigate** page, it does not open the Jobs tray. It opens the Jobs page where all jobs are listed. To display the Jobs tray, go to **Investigate** > **Navigate** (Version 11.5 and earlier) or **Investigate** > **Legacy Events** (Version 11.3 and earlier), and click the  (Jobs) icon.

This figure is an example of the Events view with a network event selected in the Events list, analyzed in the middle panel, and related metadata in the right panel.

This figure illustrates the Events view with the Filter Events panel open on the left side to filter the results set by drilling into meta values.

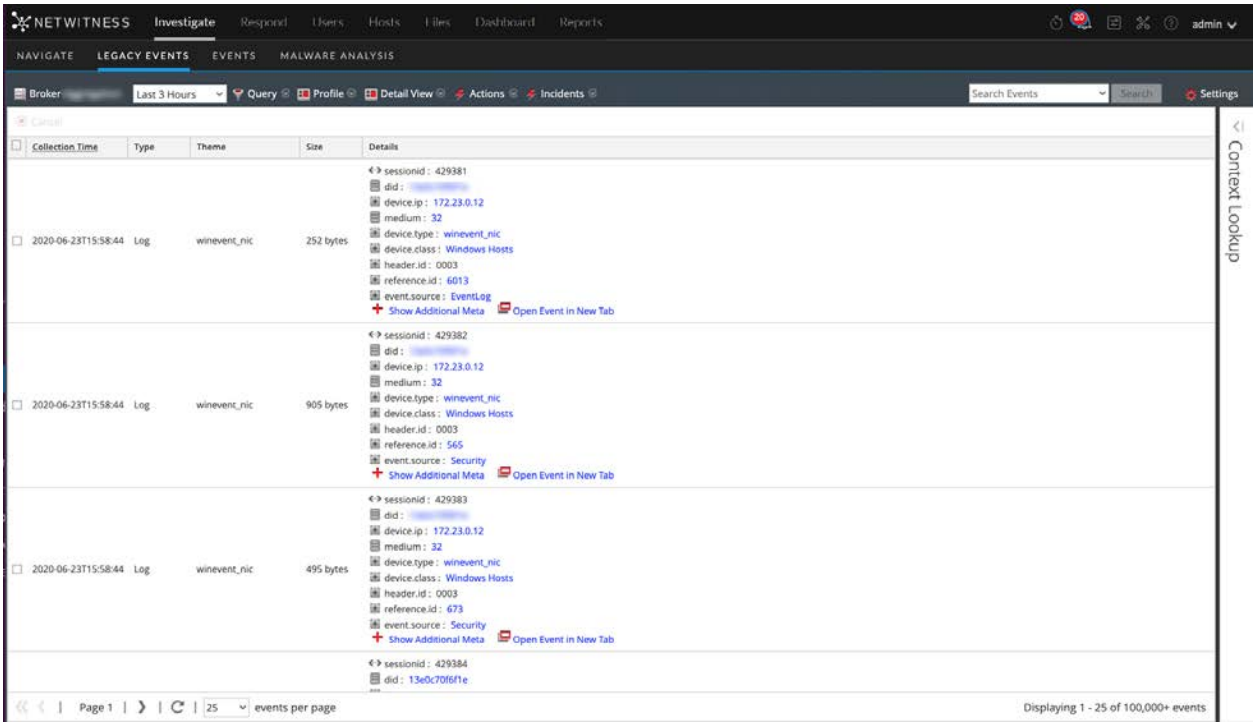


Legacy Events View

The Legacy Events view was the original user interface for analysts examining raw event data (11.0 to 11.3.x.x). The Legacy Events view is no longer needed in Version 11.4; it is hidden unless the administrator enables it as described under "Configure Investigation Settings" in the *System Configuration Guide*. When the Legacy Events view is enabled, both the Events view and the Legacy Events view are visible in the menu bar. The Legacy Events view provides a view of packet, log, and endpoint events in list form so that you can view events sequentially and reconstruct events safely.

- You can open the Legacy Events view for a meta value that you see in the Navigate view.
- For analysts without sufficient privilege to navigate a service, the Legacy Events view is a standalone investigation view in which analysts can access a list of network, log, and endpoint events from a NetWitness Core service without having to drill down through metadata first.
- The Legacy Events view presents event information in three standard forms, a simple list of events, a detailed listing of events, and a log view.
- For certain configured meta keys, such as IP address, or hostname, you can see additional context information around a value using the Context Hub. The additional context may include incidents, alerts, and other sources where the value was mentioned.
- You can export events and associated files, and create an incident from an event.

This figure illustrates the Legacy Events view.



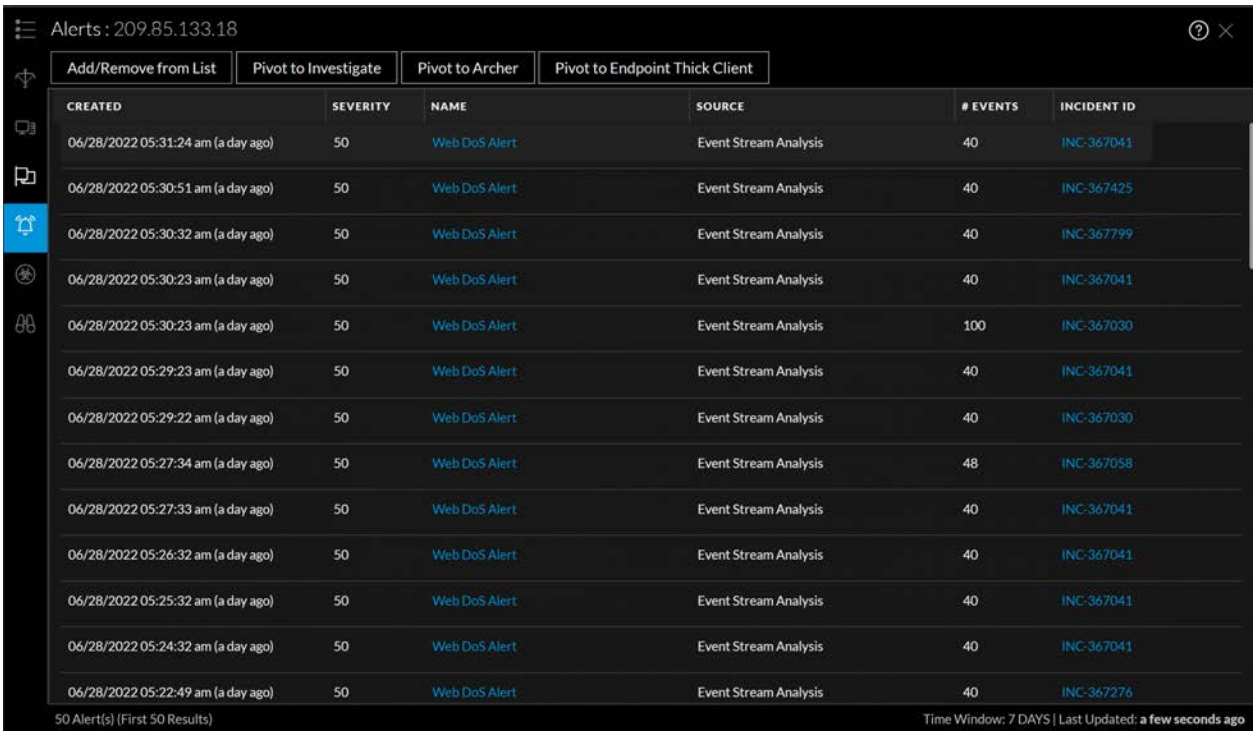
Contextual Information for an Event

In the Navigate view, Events view, and Legacy Events view, the Context Lookup panel shows details about elements associated with an event in the Context Hub for these meta types: IP Address, User, Host, Domain, MAC Address, Filename, and File hash. In addition, you can right-click all meta keys except time to see additional context.

You can interact with the elements of an event to get further insight including related incidents, alerts, custom lists, Archer assets, active directory details, NetWitness Endpoint IIOCs, and STIX data sources (namely File, TAXII Server, and REST Server). (See Look Up Additional Context for Results.)

Note: Archer assets and active directory details are available in the Events view context lookup. Endpoint context lookup is available for NetWitness Endpoint 4.4.0.2 or later hosts, but not for the NetWitness Endpoint 11.1 or later hosts.

The following figure shows the Context Lookup panel, which opens to the right of the Events panel in the Events view.



The following figure shows the Context Lookup panel, which opens to the right of the Events list in the Legacy Events view.

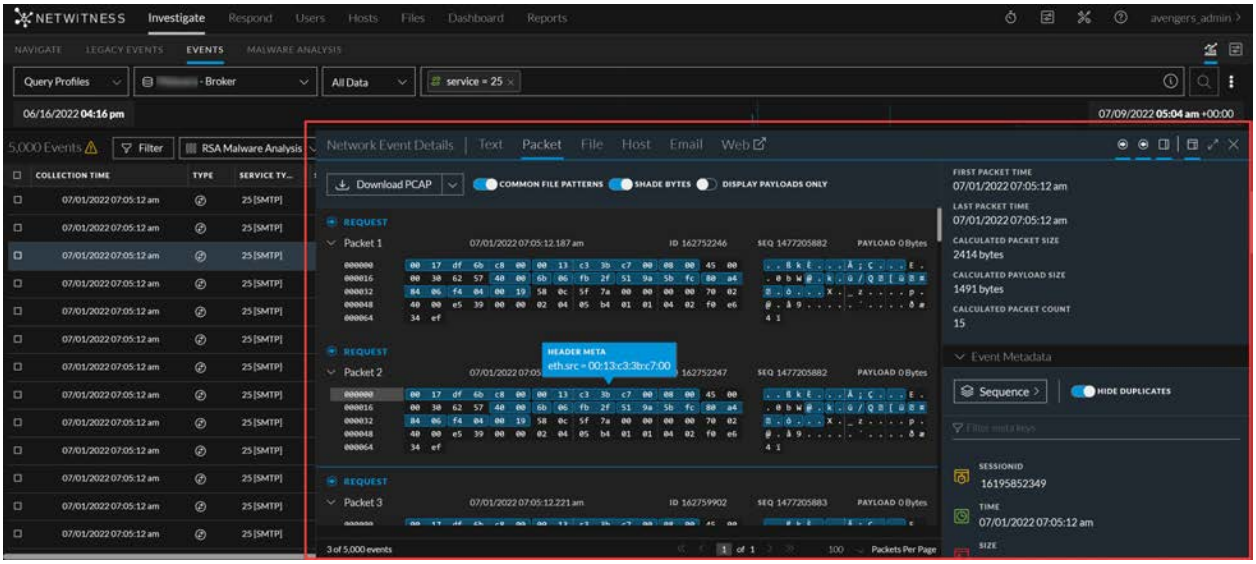


Reconstruction and Event Analysis

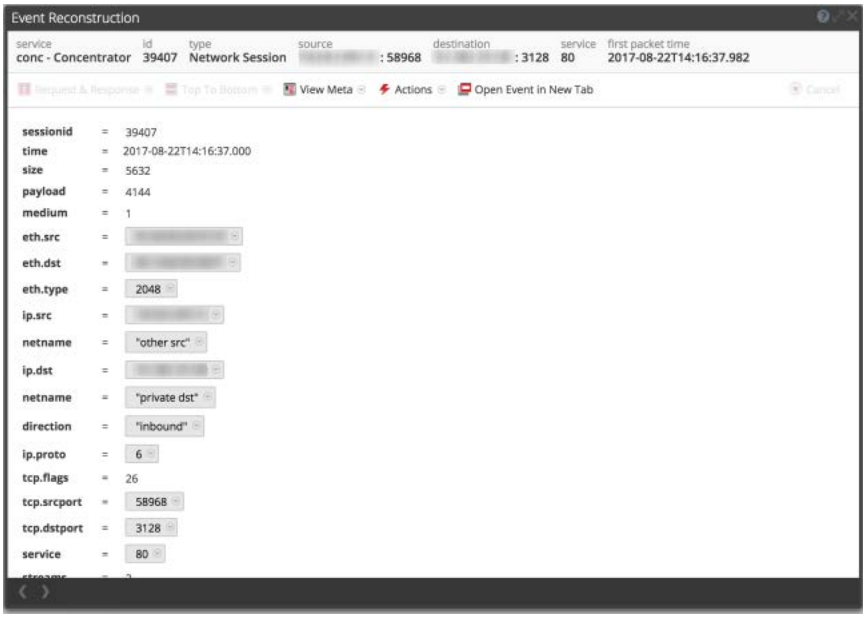
When you discover an event that merits additional investigation, you can analyze it in a form that best suits the content of event. Some forms of analysis involve safely reconstructing the event in a form similar to its native form, for example, packets, text, email, and web content. The rendering of events restricts the use of dynamic or active code that may be contained in the event to limit adverse outcome to your system or browser. Using cache improves performance when viewing previously viewed events. Each analyst has a separate cache of reconstruction data, and you can only access reconstructed events in your own cache.

Some network events are enriched with host data if you have an Endpoint deployment and the Endpoint Agent is configured for expanded network visibility. For such an event, you can view the host details.

The Events view gives you the ability to interactively analyze an event, looking at raw data, meta keys, meta values. This figure is an example of a network event rendered as packets in the Events view.



The Event Reconstruction in the Legacy Events view presents the raw data and the meta keys and meta values for an event in a list form. This figure is an example of the Event Reconstruction.



Configuring NetWitness Investigate Views and Preferences

Analysts can configure some aspects of NetWitness Investigate views and behavior. You can customize the way that Investigate views appear, the types of information displayed, and factors that affect performance in returning results and reconstructing events. All configurable settings have default values that are effective in most deployments; however, analysts have the option to adjust these if necessary.

Analysts using Investigate need to have the appropriate system roles and permissions set up for their user accounts. An administrator must configure roles and permissions as described in the *System Security and User Management Guide*.

These topics provide details:

- Configure the Navigate View and Legacy Events View
- Configure the Events View

Configure the Navigate View and Legacy Events View

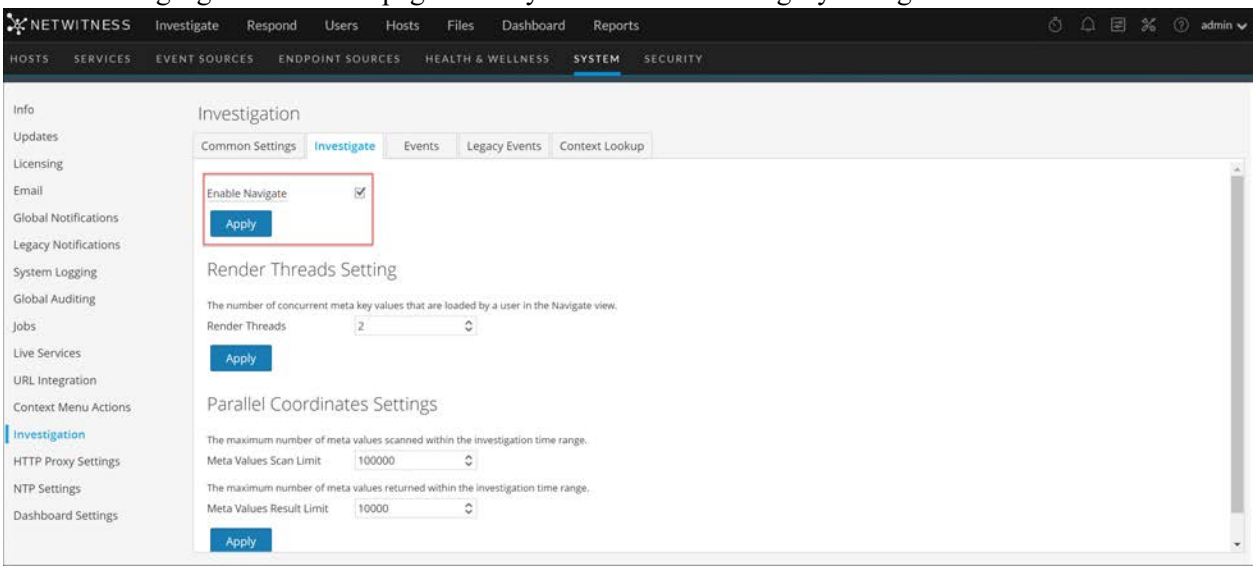
Analysts can set preferences that affect performance and behavior of NetWitness when using the Navigate view and Legacy Events view. Some of the same settings are available in two places in NetWitness, and changes made in either location are applied in the other view:

- Investigate view > Settings dialog for the Navigate view and the Legacy Events view.
- Profiles > Preferences panel > Investigation tab.
- Navigate view and Legacy Events view Search Options drop-down.

By default, the legacy Navigate view is disabled. To enable the Navigate tab in Investigate:

1. Go to (Admin) > **System** > **Investigation** > **Navigate**.
2. Select the **Enable Navigate** checkbox.
3. Click **Apply**.

The following figure shows the page where you can enable the legacy Navigate view.



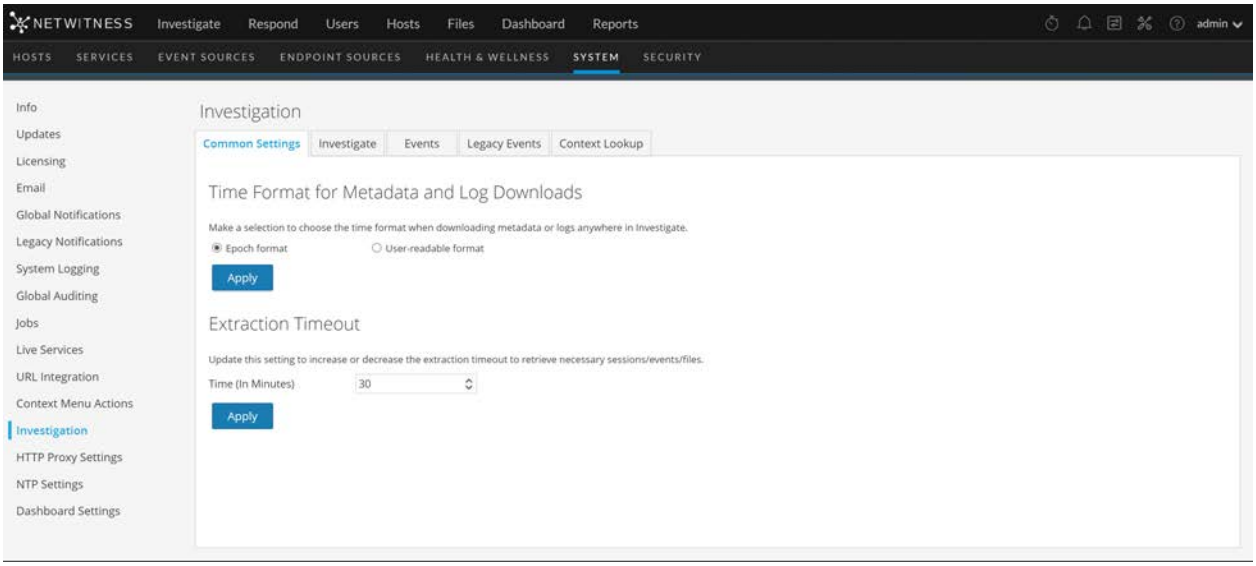
Configure Common Settings

In version 11.5 and later, the Common Settings tab allows you to configure settings that apply to the Navigate view, the Events view, and the Legacy Events view. You can set the time format used when downloading metadata and logs, and extraction timeout settings.

By default, the time format for downloads is Epoch format, which shows the time as a numerical value representing the number of seconds from the Unix epoch, January 1, 1970. The resulting number requires a conversion to be understood. You can change the setting to get a more understandable format that combines the user preference time zone, date format, and time format into an easily understood representation, which follows the industry standard ISO 8601 representation when possible.

This setting applies to all 11.5 Investigate views.

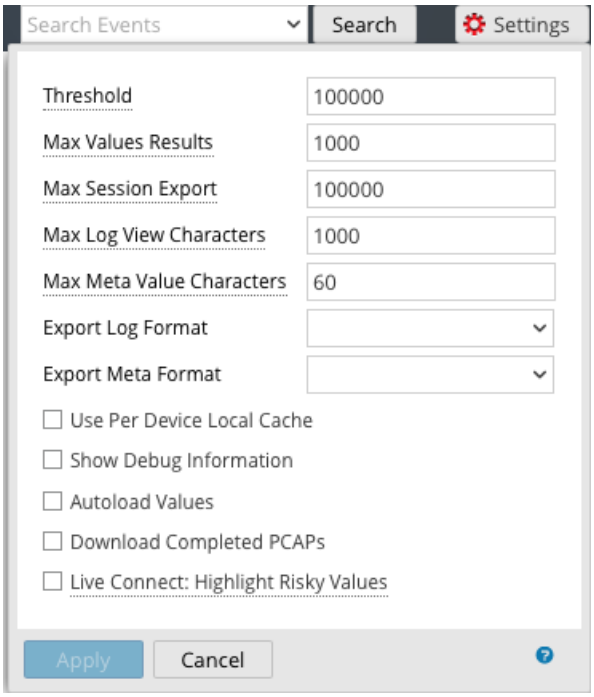
Go to  **(Admin) > System**, and in the options panel, select **Investigation**.
The Investigation Configuration panel is displayed.



Access the Navigate View and Legacy Events View Settings

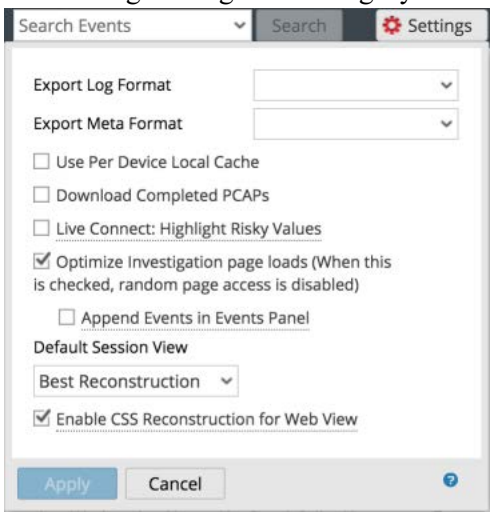
To access the settings, do one of the following:

- In the **Navigate** view toolbar, select the **Settings** option.
  The Settings dialog for the Navigate view is displayed.



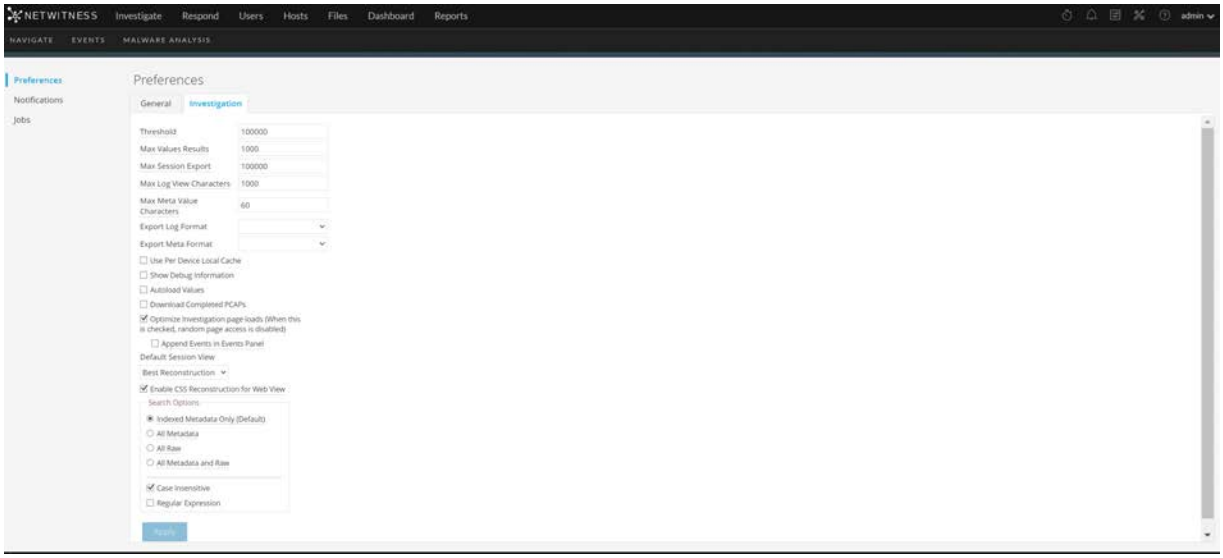
Note: Version 11.0 included a setting to Append Events in Events Panel, which was moved to the Events view settings panel in Version 11.1.

- In the **Legacy Events** view toolbar, select the **Settings** option.
  The Settings dialog for the Legacy Events view is displayed.



Note: Version 11.1 and later includes the Append Events in Events Panel setting.

- In the top right corner of NetWitness, go to > , Profile, and in the **Preferences** panel click the **Investigation** tab.
  The Investigation panel is displayed. The figure below illustrates the Investigation panel.



Calibrate Navigate View Value Loading Parameters

Several settings influence the performance of NetWitness when loading values in the Values panel. Default values are set based on common usage, and individual analysts can adjust these settings for their own investigations. To adjust these settings:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view.
2. Adjust the following parameters:
   - **Threshold:** Set the threshold for the maximum number of sessions loaded for a meta key value in the Values panel. A higher threshold allows accurate counts for a value, and also causes longer load times. The default value is **100000**.
   - **Max Values Results:** Set the maximum number of values to load in the Navigate view when the Max Results option is selected in the Meta Key Menu for an open Meta Key. The default value is **1000**.
   - **Max Session Export:** Specify the number of events that can be exported in a single PCAP or Log file.
   - **Max Log View Characters:** Set the maximum number of characters to be displayed on **Investigate > Events > Log Text**. The default value is **1000**.
   - **Max Meta Value characters:** Set the maximum number of characters in a meta value name displayed in the Navigate view Values panel. The default value is **60**.
   - **Show Debug Information:** If you want NetWitness to display the `where` clause beneath the breadcrumb in the Navigate view and the elapsed load time for each aggregated service on a Broker, check this option. The default value is **Off.**
   - **Append Events in Events Panel:** This option affects paging in the Events view and is described below under "Calibrate Events View Retrieval and Default Reconstruction."
   - **Autoload Values:** If you want NetWitness to automatically load values for the selected service in the Navigate view, check this option. When not selected, NetWitness displays a **Load Values** button, allowing the opportunity to modify options. The default value is **Off**.
3. Click **Apply**.
   The settings become effective immediately and are visible the next time you load values.

Configure Navigate View and Legacy Events View Parameters

Several settings influence the performance of NetWitness when loading values in the Navigate view and the Legacy Events view. Default values are set based on common usage, and individual analysts can adjust these settings for their own investigations. You can set these parameters separately in the Navigate view and the Legacy Events view. When configured in one view, the setting does not automatically apply to the other view. To adjust these settings:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view or the Legacy Events view.
2. Adjust the following parameters:
   - **Live Connect: Highlight Risky Values:** If you want NetWitness to highlight and display only IP addresses that are considered as risky by NetWitness community, check this option. When not selected, NetWitness displays all IP addresses. By default, this option is not selected (**Off**).

- **Use Per Device Local Cache:** You can specify the use of locally cached data from the selected service. By default, this option is not selected **(Off)**. When unchecked, Investigate sends a fresh query to the database rather than displaying cached data in the Investigate views after the initial load. If checked, Investigate uses the data from local cache.
- **Download Completed PCAPs:** You can automate the downloading of extracted PCAPs in the Navigate view and Legacy Events view so that the browser downloads the extracted PCAP and opens it in the default application for opening PCAP files, such as Wireshark. By default, this option is not selected (**Off**). If you are going to enable this option, ensure that an application that can open PCAPs is installed on your local file system and that the application is set as the default application to handle PCAP file formats.
- **Live Connect: Highlight Risky Values:** If this option is unchecked, all the meta values that have context available in Live Connect are highlighted in the Navigate view Values panel. If the option is checked, among the values that have context in Live Connect, only those values deemed Risky/Suspicious/Unsafe by the community are highlighted. By default this option is unchecked (**Off**).

3. Click **Apply**.
   The settings become effective immediately.

Configure the Default Log Export Format

You can export logs from the Navigate view and the Legacy Events view as Text, XML, comma-separated values (CSV), and JSON. There is no built-in default value for the log export format. If you do not select a format here, NetWitness displays a selection dialog when you invoke export of logs. To select the format for exported logs:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view or Legacy Events view.
2. Select one of the options from the **Export Log Format** drop-down menu.
3. Click **Apply**.
   The setting goes into effect immediately.

Configure the Default Meta Value Export Format

You can export meta values from the Navigate view and Legacy Events view as Text, CSV, tab-separated values (TSV), and JSON. There is no built-in default value for the meta value export format. If you do not select a format here, NetWitness displays a selection dialog when you invoke export of meta values. To select the format for exported meta values:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view or Legacy Eventsview.
2. Select one of the options from the **Export Meta Format** drop-down menu.
3. Click **Apply.**
   The setting goes into effect immediately.

Note: If you upgrade to version 11.5.2, the **Export Meta Format** preference is not retained and is reset to blank. You must re-configure this value after you upgrade to version 11.5.2.

Calibrate Legacy Events View Retrieval and Default Reconstruction

You can configure several parameters that control the how NetWitness retrieves events and reconstructs events in the Legacy Events view. To adjust these paramaters:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Legacy Events view.
2. Configure the following parameters.
   - **Optimize Investigation page loads**: Set a paging option. When optimized, results are returned as quickly as possible,and you cannot go to a specific page in the event list. Unchecking this box changes the Events list pagination to allow you to go to a specific page in the list (or to the last page). The default value is **enabled**.
   - **Default Session View:** Selects the default reconstruction type for the initial reconstruction in the Legacy Events view. The default value is **Best Reconstruction** in which events are reconstructed using the reconstruction method most appropriate to the event.
3. Go to the **Preferences** panel > **Investigation** tab, or to the **Settings** dialog for the Navigate view (11.1) or the Legacy Events view (11.2 and later), and set the **Append Events in Events Panel** option. When this option is selected, the events displayed in the Events Panel are added incrementally. For example, each time you click the next page icon, the next increment of events is added, at first you see 1 to 25, then 1 to 50, then 1 to 75 and so on. This option is available only if the **Optimize Investigation Page Loads** option is enabled.
4. To activate the changes immediately, click **Apply**.

Enable or Disable Cascading Style Sheet Rendering in Web Content Reconstructions

Analysts can enable the use of cascading style sheets (CSS) when reconstructing web content. If enabled, the web reconstruction includes CSS styles and images so that its appearance matches the original view in a web browser. This includes scanning and reconstructing related events, and searching for style sheets and images used in the target event. The option is enabled by default. Disable this option if there are problems viewing specific websites.

Note: The appearance of the reconstructed content may not match the original web page perfectly if related images and style sheets could not be found or were loaded from the web browser's cache. Also, any layout or styling that is performed dynamically through the client side javascript is not rendered in the reconstruction because all client side javascript is removed for security purposes.

To enable or disable this option:

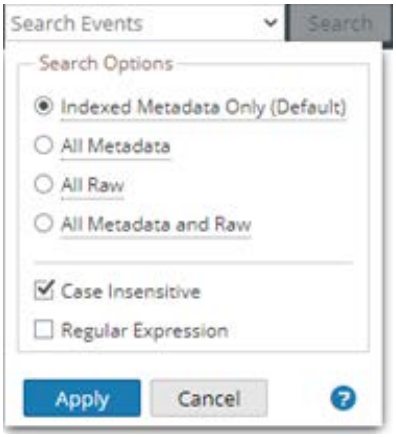
1. Go to the **Preferences** panel > **Investigation** tab.
2. Select the **Enable CSS Reconstruction for Web View** checkbox.

3. Click **Apply**.
   The setting becomes effective immediately and is visible in the next web content reconstruction.

Configure Search Options

You can configure search options to apply when you type a search string in the Search field. Edit the Search Options in the Profile > Preferences panel > Investigation tab or in the Navigate and Legacy Events view Search Options drop-down menu. To configure search options:

1. Go to the Search Options.
   The following figure illustrates the Search Options drop-down menu for Version 11.2 and later.



2. Select one or more search options to apply to the search. Search for Text Patterns in the Navigate and Legacy Events Views provides detailed information about each option.
3. To save the search settings, click **Apply**.
   The preferences are saved and effective immediately.

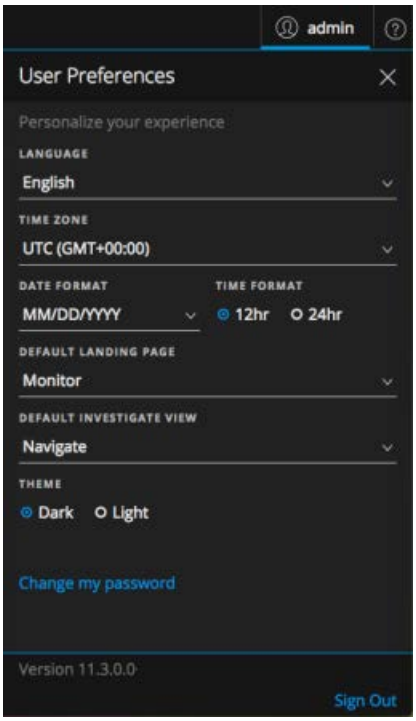
Configure the Events View

Note: The information in this topic applies to NetWitness Version 11.1 and later.

Analysts can set preferences that affect the behavior of NetWitness when using the Investigate > Events view. If the Events view is open; these two buttons give access to preferences dialogs: and . The User menu ( ) is focused on global user preferences such as time zone, while the Events preferences menu ( ) is focused on user preferences for behavior in the Events view. The rest of this section describes both sets of preferences.

Set the Default Investigate View

You can select the default view when you open Investigate: Navigate view, Events view, Hosts view, Files view, Entities view, or Malware Analysis view. The default Investigate view is set in the global User Preferences dialog (in the upper right corner of the NetWitness browser window, select). The global user preferences are described in detail in the *NetWitness Platform Getting Started Guide*.

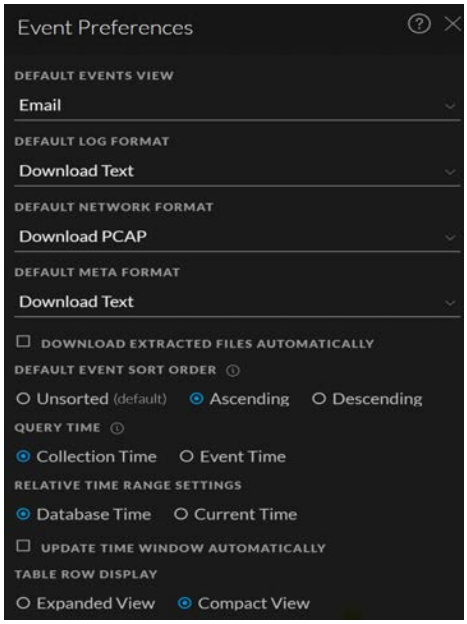


Set User Preferences for the Events View

You can set your own preferences relevant to the Events view. The preferences selected persist per user and are available whenever the specific user logs in to the application.

To set default values for working in the Events view:

1. In the **Events** view, click .
   The Event Preferences dialog is displayed. Different versions of the dialog have some differences in labeling and available options as shown in the following figure.



2. In the **Default Events View** field, select the default reconstruction type when you open an event in the Events panel: **Text**, **Packet**, **File**, (Version 11.5 and Later) **Host** or **Email** .
   If you have not selected a default analysis type, when you open an event, the default reconstruction type is the Packet analysis, except for log and endpoint events, which open to the Text analysis. If you select a default reconstruction type, the reconstruction type is the default reconstruction that you specified. In both cases, the default is the starting point, and if you change the type while you are working, the type you choose is used for the next reconstruction.
3. In the **Default Log Format** field, select the download format for exporting logs: **Download Log** (11.3) or **Download Text** (11.4), **Download XML**, **Download CSV**, or **Download JSON**. If you do not select a format here, the default download format is **Download Text**. These options are also available at the time of download in a drop-down menu.
4. In the **Default Packet Format** (11.3) or **Default Network Format** (11.4) field, select the default format for downloading packets. If you do not select a format here, the default download format is **Download PCAP**. These options are also available at the time of download in a drop-down menu:

- **Download PCAP** to download the entire event as a packet capture (*.pcap) file
- **Download All Payloads** (11.3) or **Download Payloads** (11.4) to download the payload as a *.payload file
- **Download Request Payload** to download the request payload as a *.payload1 file
- **Download Response Payload** to download the response payload as a *.payload2 file

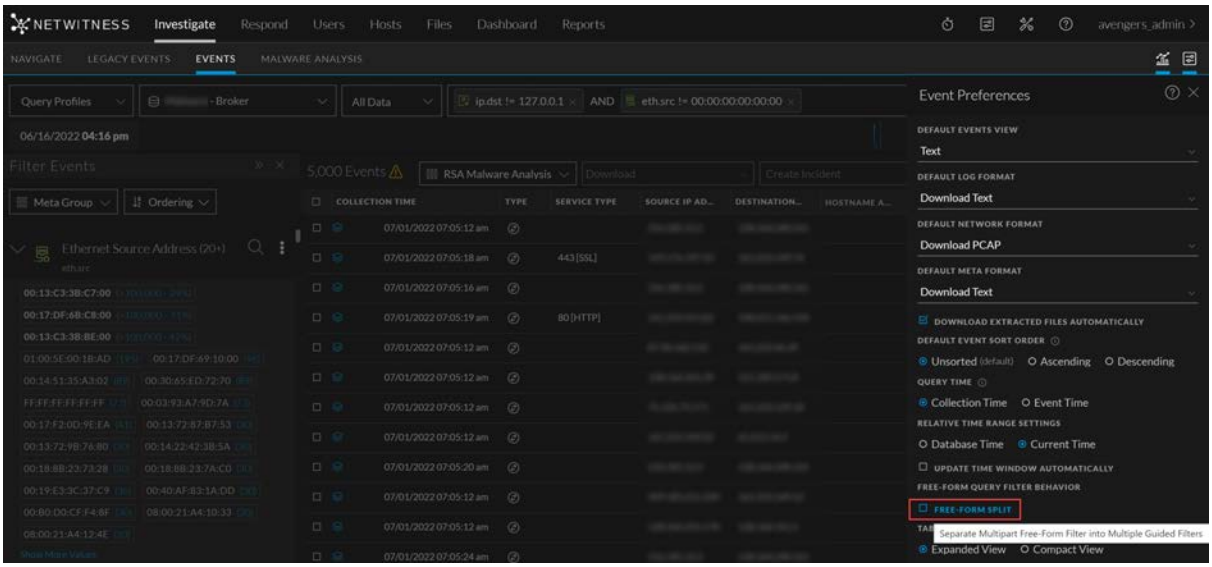
5. (Version 11.4 and later) In the **Default Meta Format** field, select the download format for exporting metadata: **Download Text**, **Download CSV**, **Download TSV**, or **Download JSON**. If you do not select a format here, the default download format is **Download Text**.
6. (Version 11.5.1 and later) To change the preference for the time that is matched when you submit queries, select an option in the **Query Time** field.
   - When **Collection Time** is selected, the time of an event reflects when it was received and stored into the system. This is the default setting.
   - When **Event Time** is selected, the time of an event reflects the time at which the event actually occurred. A good case for using the event time is when investigating logs or endpoints and looking for events that occurred around the same time. Using Event Time filters out all network events.

   Note: By default, the time range for a query is based on the time that the Decoder consumed the event, which is always not the same time that the event occured. To see the actual event time instead of the collection time in the Events view, select Event Time.

   Note: When you query with the event time preference enabled, you must use the collection time column and time zone column. This is to help you differentiate if the events are listed in a sequential order or not. This scenario occurs because there is no global standard to follow while logging the event times, which results in different events from different sources being in different time zones.
7. (Version 11.5.1 and later) Under **Relative Time Range Settings**, choose either **Database Time** or **Current Time**. The Events view can display results based on the relative time range, Last 2 hours or Last 30 Days, for example. The time range can be relative to the time when the event was received and stored into the system or the current time zone's clock time. When you set the time format, your individual user preference is saved until changed again. The default setting for this preference is **Database Time**, which is the same time format used to display query results in the Navigate view and Legacy Events view.
   - When **Database Time** is selected, the time range is relative to the time of the last stored event.
   - When **Current Time** (labeled **Wall Clock Time** in Version 11.3 and earlier) is selected, the time range is relative to the current time in the timezone set in user preferences.

   Note: (Version 11.6) **Current Time** is the default for **Relative Time Range Settings**. In previous versions, **Database Time** was the default value. Make a note that this may cause time range mismatch between **Events View** (using **Current Time** as default) and **Navigate View** (using **Database Time** as default). This change does not affect the existing users and is applicable only to the new users.

8. (Version 11.4 and 11.5) Under **Time Format for Query**, choose either **Database Time** or **Current Time**. The Events view can display results based on the database time or the current clock time. When you set the time format, your individual user preference is saved until changed again. The default setting for this preference is **Database Time**, which is the same time format used to display query results in the Navigate view and Legacy Events view.
   - When **Database Time** is selected, the end time for a query is based on the time that the event was stored.
   - When **Current Time** (labeled **Wall Clock Time** in Version 11.3 and earlier) is selected, the query is executed with the current time in accordance with the timezone set in user preferences.
9. (Version 11.4 and later) To set the sort sequence by collection time for the events listed in the Events panel, select one option under **Default Event Sort Order**. After you have selected a preference, you can also interact with the table column headers to request the results again sorted differently as described in Use Columns and Column Groups in the Events List.
   - **Unsorted** (default for Version 11.4.1): To list events as processed by the Core services. Unsorted is faster because it streams back the events as soon as a match is found versus waiting for all Core services to respond and then displaying them in the chosen order.
   - **Ascending** (default for Version 11.4 and earlier): To put the events with the earliest collection time first in the list. The oldest events are displayed first if in ascending order.
   - **Descending**: To put the events with the latest collection time first in the list. The newest events are displayed first if in descending order. When investigating logs, you may want to change the sort sequence to see the latest collection time first.
     If results exceed the events limit, not all events can be loaded. The portion of returned events loaded in the Events panel matches the sort order preference: the oldest portion of events is loaded when Ascending order is selected, and newest portion of events is loaded when Descending order is selected. When Unsorted is selected, the oldest portion of events is matched and then listed unsorted. if you changed the sort order preference after events were loaded, you must refresh the view to apply the new sort order.
10. If you want all extracted files to be downloaded automatically, select the **Download extracted files automatically** checkbox. You can go to the Jobs queue to view the extracted files.
11. (Version 11.3 and later) To automatically update the time range window in the query bar when the service is polled (at one minute intervals) and sends fresh results, select the **Update Time Window Automatically** checkbox. When the time range is updated, the  (Submit Query) button is activated and you can submit a query to get fresh results. To keep the time range window in the query bar synchronized with the current results, clear the checkbox (this is the default value
12. (Version 11.6 and later) If you have multiple rows of data and want to word wrap or unwrap the content of an event , choose either **Compact View** or **Expanded View**. Based on the selection you make the rows will appear compact or expanded.
13. (Version 11.7.1 and later) With the new preference, analysts can choose to split the free-form queries into multiple guided filters or a single free-form query. Analysts can switch the modes using the **Free Form Split** checkbox.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
avengers_admin
NAVIGATE
LEGACY EVENTS
EVENTS
MALWARE ANALYSIS
Query Profiles
- Broker
All Data
ip.dst != 127.0.0.1
AND
eth.src != 00:00:00:00:00:00
06/16/2022 04:16 pm
Filter Events
5,000 Events
RSA Malware Analysis
Download
Create Incident
Meta Group
Ordering
Ethernet Source Address (20+)
COLLECTION TIME
TYPE
SERVICE TYPE
SOURCE IP AD...
DESTINATION...
HOSTNAME A...
07/01/2022 07:05:12 am
07/01/2022 07:05:18 am
443 [SSL]
07/01/2022 07:05:16 am
07/01/2022 07:05:19 am
80 [HTTP]
07/01/2022 07:05:12 am
07/01/2022 07:05:12 am
07/01/2022 07:05:12 am
07/01/2022 07:05:12 am
07/01/2022 07:05:20 am
07/01/2022 07:05:12 am
07/01/2022 07:05:12 am
07/01/2022 07:05:24 am
Event Preferences
DEFAULT EVENTS VIEW
Text
DEFAULT LOG FORMAT
Download Text
DEFAULT NETWORK FORMAT
Download PCAP
DEFAULT META FORMAT
Download Text
DOWNLOAD EXTRACTED FILES AUTOMATICALLY
DEFAULT EVENT SORT ORDER
Unsorted (default)
Ascending
Descending
QUERY TIME
Collection Time
Event Time
RELATIVE TIME RANGE SETTINGS
Database Time
Current Time
UPDATE TIME WINDOW AUTOMATICALLY
FREE-FORM QUERY FILTER BEHAVIOR
FREE-FORM SPLIT
Separate Multipart Free-Form Filter into Multiple Guided Filters
Expanded View
Compact View

Beginning an Investigation

Based on the question you are attempting to answer, NetWitness Investigate offers different starting points: Navigate view (Version 11.5 and earlier), Events view, Legacy Events view (Version 11.3 and earlier), Hosts view, Files view, Users (Entities) view, and Malware Analysis view.

Specific user roles and permissions are required for a user to conduct investigations in NetWitness. If you cannot perform a task or see a view, the administrator may need to adjust the roles and permissions configured for you.

Note:
- The Files and Hosts views are available in Version 11.1 and later (refer to the *NetWitness Endpoint Quick Start Guide* and *NetWitness Endpoint User Guide* for details). Before Version 11.5, these views were submenu of Investigate.
- The Users view is available in Version 11.2 and later (refer to the *NetWitness UEBA Quick Start Guide* and the *NetWitness UEBA User Guide* for details); in Version 11.4 it is labeled Entities view. Before Version 11.5, it was submenu of Investigate.
- By default, the Legacy Events view is disabled in Version 11.4, but can be enabled by an administrator as described in the *System Configuration Guide*.
- By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.
- Specific user roles and permissions are required for a user to conduct investigations and malware analysis in NetWitness. If you cannot see a view, the administrator may need to adjust the roles and permissions configured for you.
- The 11.4 Events view is the default view for investigating events. The default workflow for analysts interacting with events is optimized to limit the need to transition from one view to another. By combining capabilities that were previously in two distinct workflows, known as Event Analysis and Events, the analyst has a single workflow for analyzing events. With the new functionality added to the Events view, the Legacy Events view is no longer needed. By default the previous workflow is not in the Investigate menu, but an administrator can re-enable it as described in "Configure Investigation Settings" in the *System Configuration Guide*.

Focus on Metadata, Raw Events, and Event Analysis

To hunt for events that drive the incident response workflow and to do strategic analysis after another tool has generated an event, go to **Investigate > Navigate**, **Investigate > Events**, or **Investigate > Legacy Events**. You can investigate the metadata and raw events for a single Broker or Concentrator. In each of these views, you can execute a query and filter the results by narrowing the time range and querying metadata. These topics provide details about beginning an investigation:

- Begin an Investigation in the Events View
- Begin an Investigation in the Navigate or Legacy Events View

Focus on Hosts and Files

To hunt for information on hosts that have the Endpoint agent running, go to **Hosts** (Version 11.5) or **Investigate > Hosts** (Version 11.4). For every host, you can see processes, drivers, DLLs, files (executables), services, and autoruns that are running, and information related to logged-in users. To begin an investigation by looking at files in your deployment, go to **Investigate > Files**. (See the *NetWitness Endpoint User Guide* for detailed information.)

Focus on Risky User and Entity Behavior

To discover, investigate, and monitor risky behaviors across all users and entities in your network environment, go to **Users** (Version 11.5), **Investigate > Entities** (Version 11.4), or NetWitness UEBA (User and Entity Behavior Analytics) . In versions 11.3 and earlier, the menu option is **Investigate > Users**. You can detect malicious and rogue users, pinpoint, high-risk behaviors, discover attacks, and investigate emerging security threats. (See the *NetWitness UEBA User Guide* for detailed information.)

Focus on Scanning Files for Malware

To scan files for potential malware, or set up a continuous scan of a service, go to **Investigate > Malware Analysis**. Scan results are expressed as four types of analysis: network, static, community, and sandbox with an indicator of compromise (IOC) rating. There are several other ways to begin working in Malware Analysis:

- You can begin Malware Analysis from the Malware Analysis dashlets in the Monitor view to quickly see the riskiest potential threats.
- You can right-click a meta key in the Navigate view, and select **Scan for Malware**.

See the *Malware Analysis User Guide* for detailed information.

Begin an Investigation in the Navigate or Legacy Events View

The Navigate view is the default view for Investigate unless you have selected a different view as your opening view. This user preference is set on the application level as described in Configuring NetWitness Investigate Views and Preferences. In the Navigate view and Legacy Events, you are hunting for events of interest based on a query. In the Navigate view you can also refine results by clicking on meta keys and meta values. When you find interesting events, you can take a closer look at the event in the other Investigate views.

To begin an investigation in the Navigate view or Legacy Events view, a service must be specified.

- Investigate opens the Navigate view or the Legacy Events view with the user-specified default service selected.
- If no default service is currently specified and the service id is not in the URL, Investigate presents a dialog for selecting the service or collection to investigate.
- When a service is selected manually or by default in the Navigate view or Legacy Events view, you can change the service or collection to investigate by selecting the service name in the toolbar. Investigate presents the dialog for selecting the service to investigate.

Note: The Archiver service does not appear in the Navigate view to minimize user experience of slow performance when performing investigations. The Archiver is available in the Legacy Events view for log exports and enhanced search capabilities.

With a service or collection selected, Investigate is ready to load data for the service or collection. It is recommended that you also select a time range so that results load faster. Several settings in the Navigate view and Legacy Events view Settings dialog or the Profiles > Preferences panel > Investigations tab affect the loading process: Threshold, Max Values Results, Show Debug Information, Autoload Values, and Optimize Investigation page loads (see Configuring NetWitness Investigate Views and Preferences).

Note: In the Legacy Events view data loads automatically. If you specified Autoload Values in the Navigate view preferences, Investigate populates the data automatically. Otherwise, you must select Load Values. Investigate populates the metadata in the Navigate view Values panel and results become visible almost immediately.

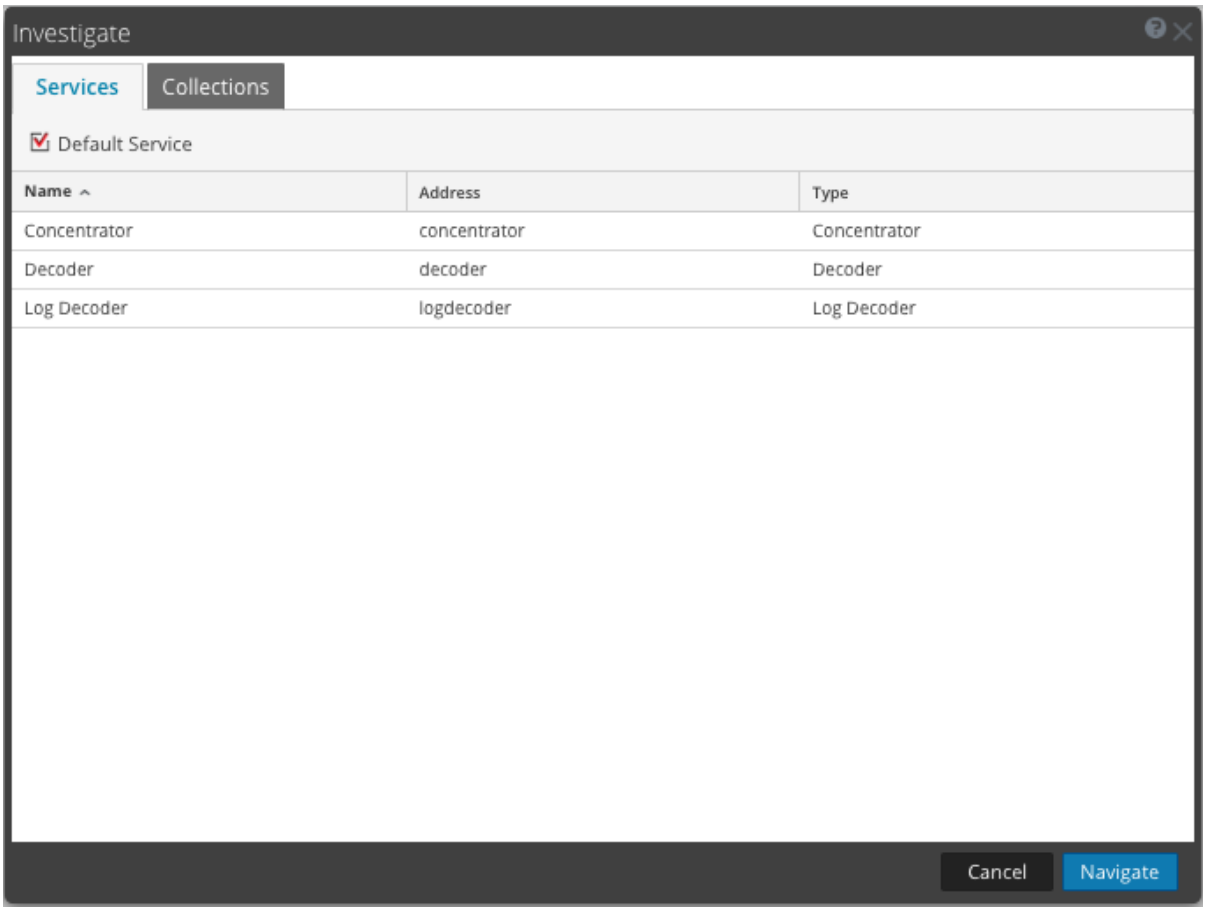
The rest of this topic provides instructions for beginning the investigation of data on a service.

Note: Only users with the administrator role can create a collection, and only the creator of the collection is able to investigate a collection.

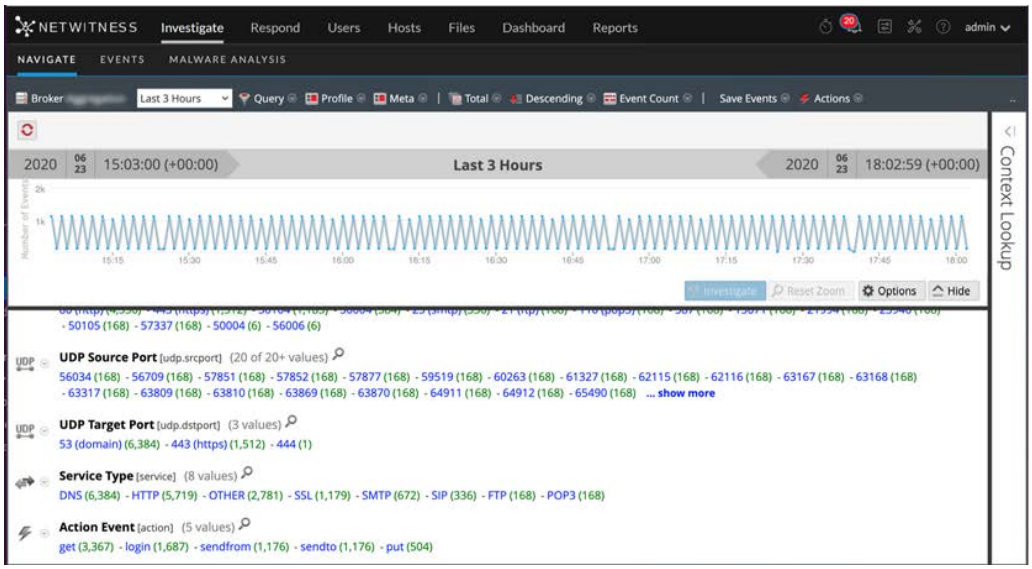
After loading data in the Navigate or the Legacy Events view, refine results, reconstruct and analyze events, then download and act upon results (see Refining the Results Set and Reconstructing and Analyzing Events and Downloading and Acting Upon Results).

Begin an Investigation (No Default Service)

1. Go to **Investigate > Navigate** or **Legacy Events**.
   The Investigate dialog is displayed.



2. Double-click a service or select a service, usually a Concentrator, and click **Navigate**.
   The data loads automatically in the Legacy Events view. If you are working in the Navigate view, the resulting panel displays the activity for the selected service, but the data is not loaded automatically.
3. (Recommended) Select a specific time range so that results load faster.
4. If you want to modify investigation options before loading, you can create or modify a custom profile, apply a different time range, create or apply a meta group, and perform a custom query as described Refining the Results Set. You can also modify options at any time during the investigation.
5. To load data in the Navigate view, click Load Values.
   The data for the selected service begins loading.

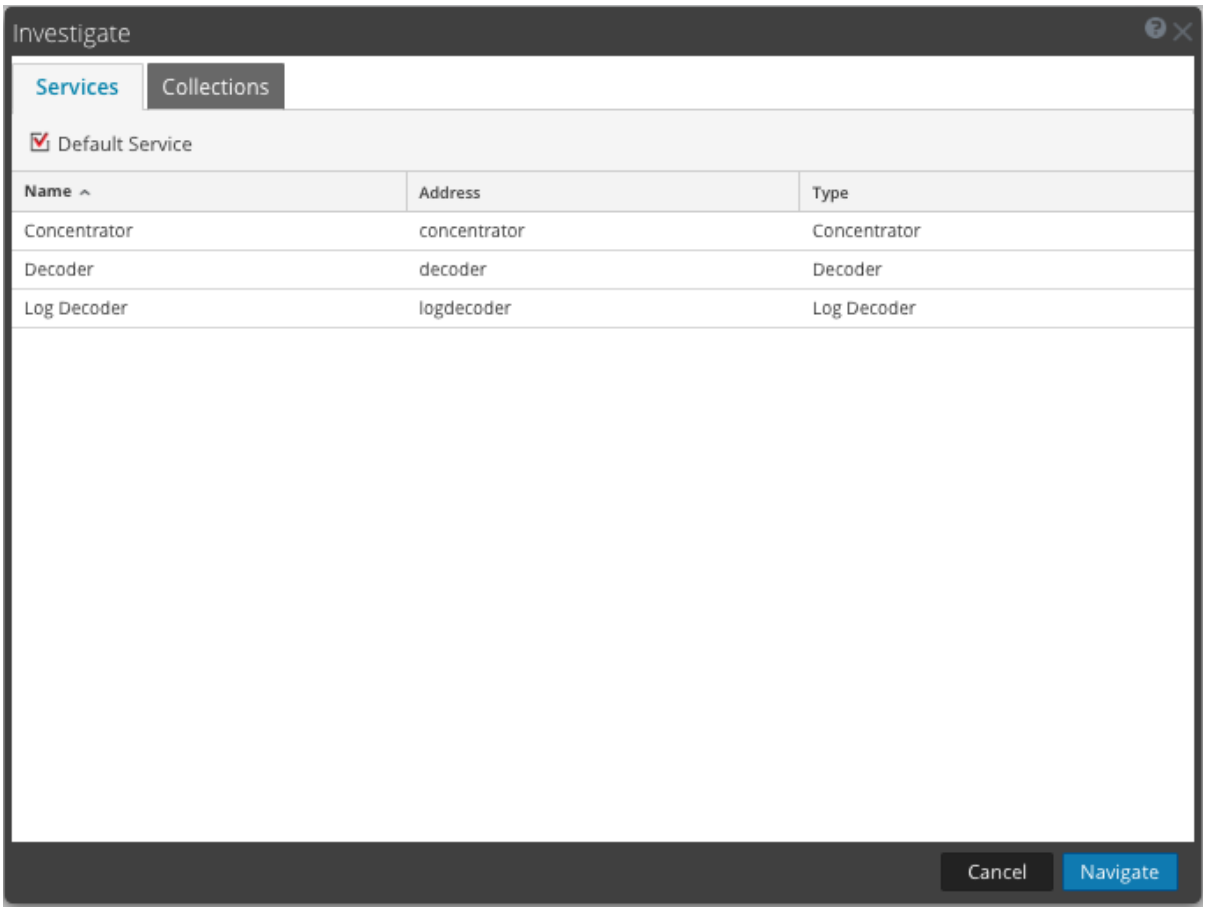


With the service selected and data loaded, you are ready to begin analyzing the data.

Set or Clear the Default Service

You can set the default service and clear the default service in the Investigate a Service dialog.

1. Click the service name in the toolbar.
   The Investigate dialog is displayed.



2. Select a service on the **Services** grid, and click Default Service.
   The service becomes the default, (indicated by **Default** in parentheses after the service name).

3. To clear the default service, select the default service in the grid, click Default Service, and click **Cancel** to close the dialog.
   No default service is set.

> **Note:** Clicking cancel does not cancel your selection of the default service. It closes the dialog without navigating to the currently selected service in the grid. Setting a default service that is different from the service currently being investigated, does not refresh the Navigate view. You must explicitly select and navigate to a different service.

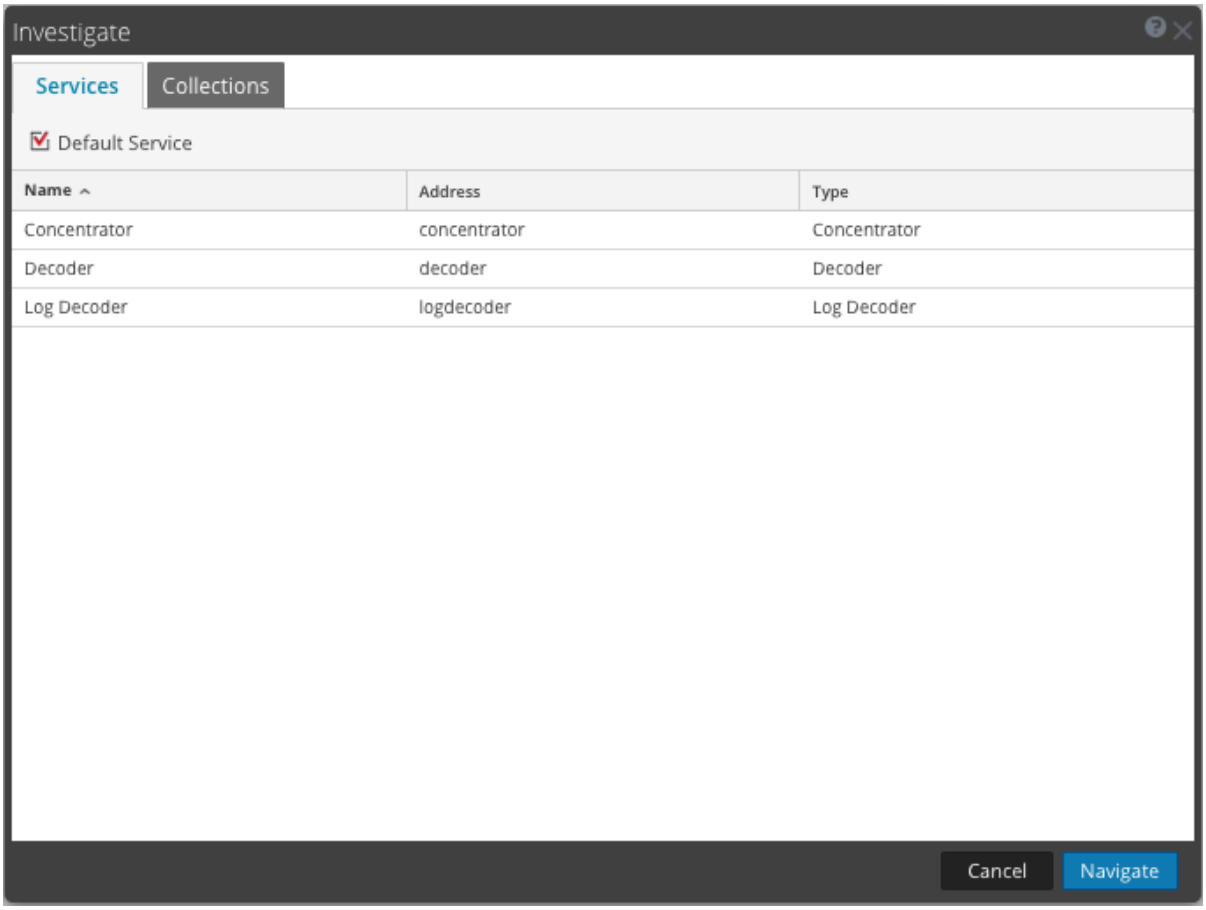
Begin an Investigation (Default Service Specified)

1. Go to **Investigate > Navigate** or **Legacy Events**.
   If the Autoload Values setting is set to off, the Navigate view is displayed with the default service selected, and ready to load data. If the Autoload Values setting is on, the values are loaded as shown in Step 3. In the Legacy Events view, the data is loaded automatically.

2. If you want to modify investigation options in the Navigate view before loading, you can create or modify a custom profile, apply a different time range, create or apply a meta group, and perform a custom query.
3. When ready, click Load Values.
   The values for the service are loaded in accordance with the selected options. With the service selected and data loaded you are ready to begin analyzing the data.

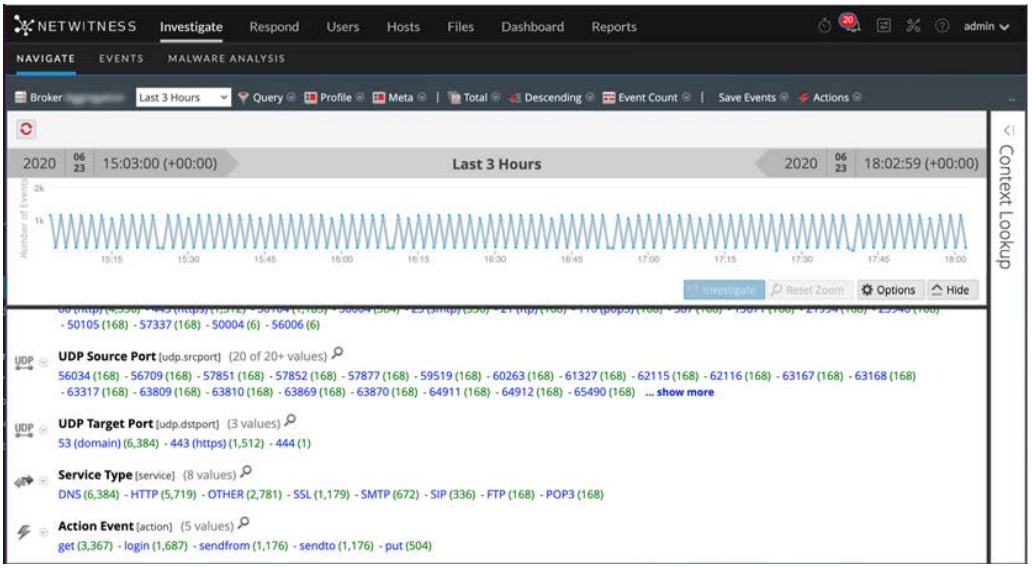
Change the Service or Collection to Investigate

1. In the Navigate view or the Legacy Events view, click the service name at the top of the options panel.
   The Investigate dialog is displayed.



2. Double-click a service or select a service and click **Navigate**. The resulting panel displays the activity for the selected service.
   If the Autoload Values setting is on, the values are loaded as shown in Step 3. Otherwise, the Navigate view is displayed with the default service selected, and data ready to load. In the Events view the data is loaded automatically.

3. When ready, click Load Values.
   The values for the service begin loading in accordance with the selected options.



   With the service selected and data loaded you are ready to begin analyzing the data.

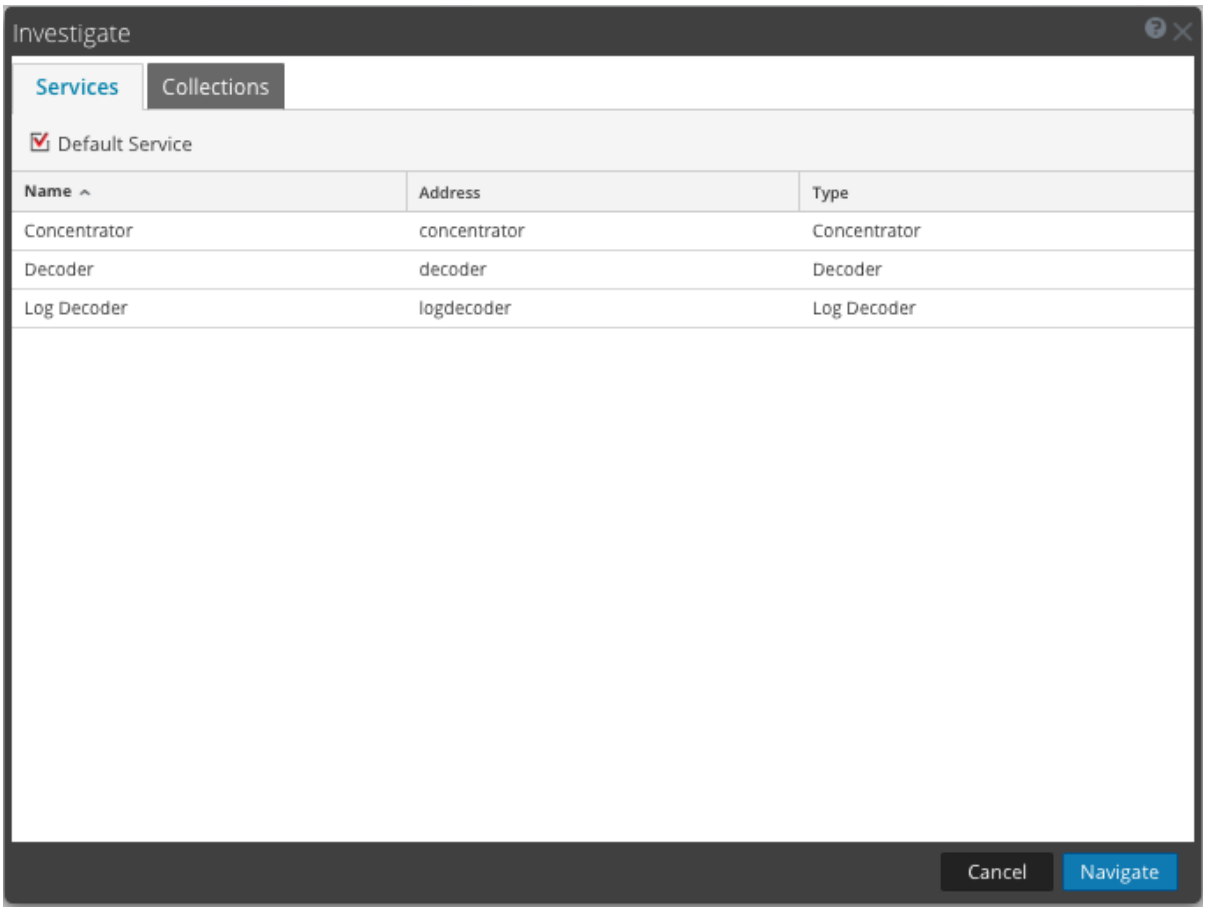
Investigate Workbench Restoration Collections

This procedure enables administrators to select content from an existing collection to reprocess for further investigation. This applies to Decoders that use Workbench services.

Note: Only a user with administrative privileges can create a collection, and you can view only those collections that you created.

To reprocess data for further investigation:

1. Go to **Investigate > Navigate** or **Legacy Events**.
   The Investigate dialog is displayed.

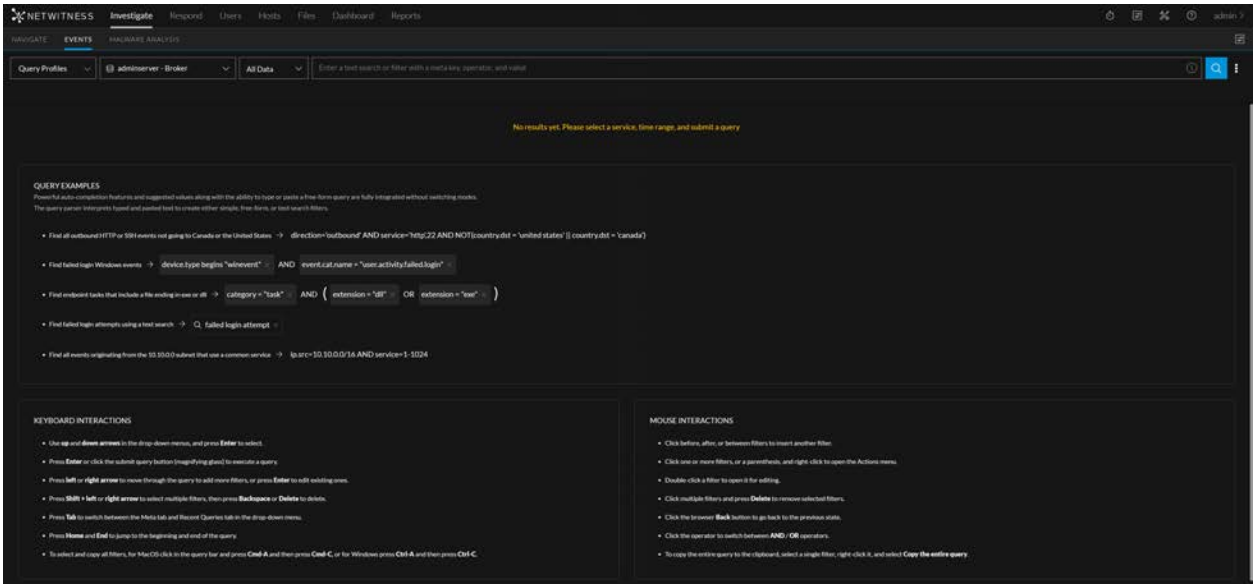


2. Select a workbench service and workbench name that you want to investigate.
3. Click **Navigate** to perform an investigation on the selected workbench service.
   Click **Cancel** to select a different workbench service to investigate.
   The Investigation view is displayed. With the collection selected and data loaded, you are ready to begin analyzing the data.

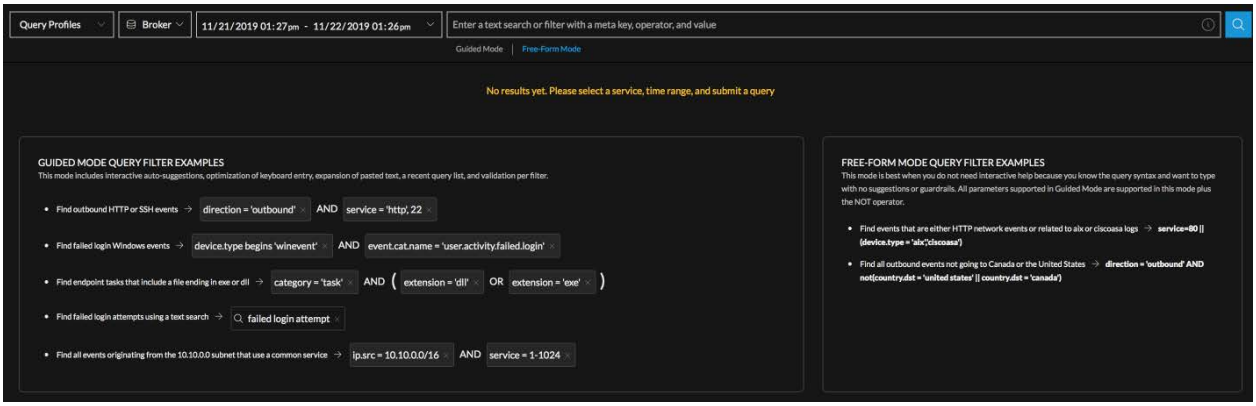
Begin an Investigation in the Events View

The Events view offers most of the features that are available in both the Navigate view and the Legacy Events view. Similar to the Navigate view, there is a view into meta keys and meta values for logs, endpoints, and packets. Like the Legacy Events view, an events list shows events listed in the order by time, and you can view the raw event, related metadata, and a reconstruction of an event. The Event reconstruction has some helpful cues to identify points of interest. See Reconstructing and Analyzing Events.

The following figure shows the initial Events view with some examples of queries and information about keyboard and mouse interaction. This figure depicts the initial view.



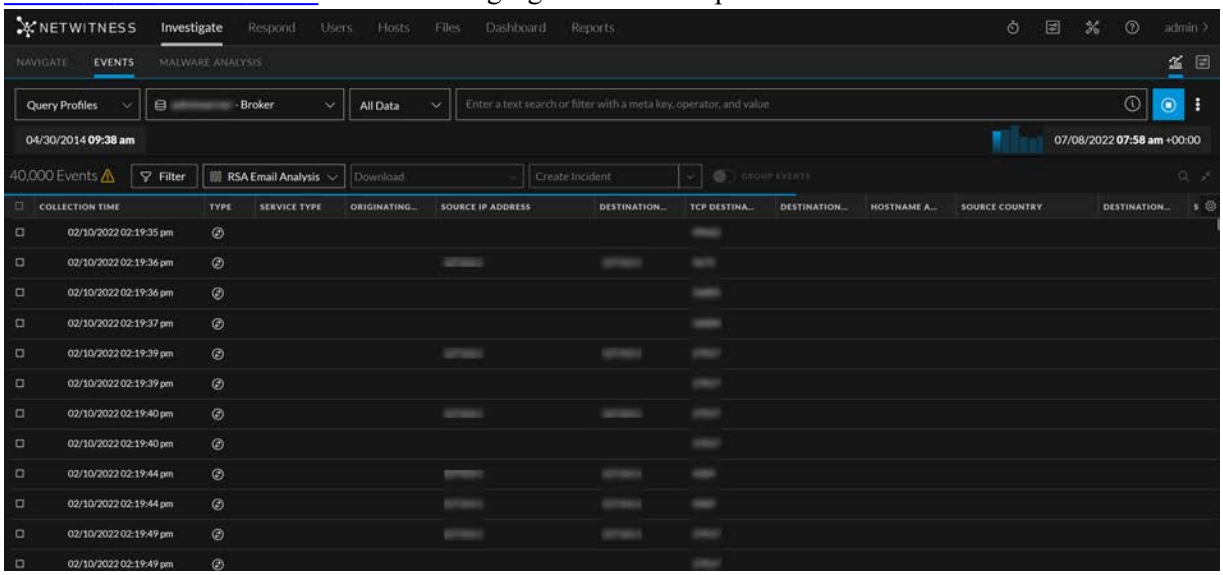
This figure show the examples provided in Version 11.4, which has two modes for creating queries.



Access the Events View

Several ways to access the Events view are available in Version 11.1 and later.

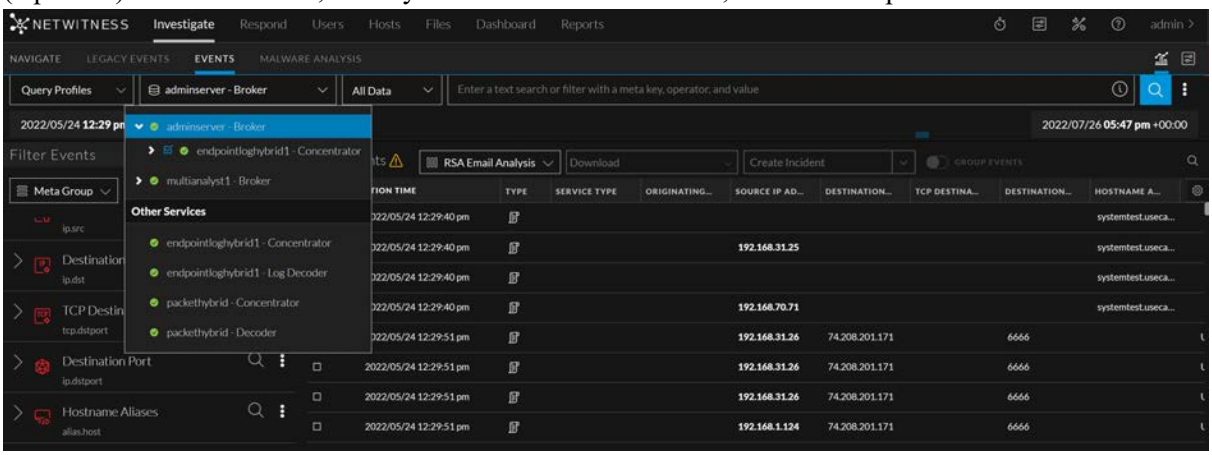
- Go to **Investigate > Events** or select the **Investigate** option in the main menu if you have made the Events view your default Investigate view. The following procedure provides detailed steps.
- Hover over and click a count (the green number after a meta value) in the Navigate view. The Events view opens with the list of events for the selected drill point, and you can begin working as described in Analyze Events in the Events View.
- Hover over a count and control-click **Open Events in new tab**. The Events view opens in a new tab with the list of events for the selected drill point, and you can begin working as described in Analyze Events in the Events View. The following figure is an example of the list of events.



To begin an investigation in the Events view using direct access:

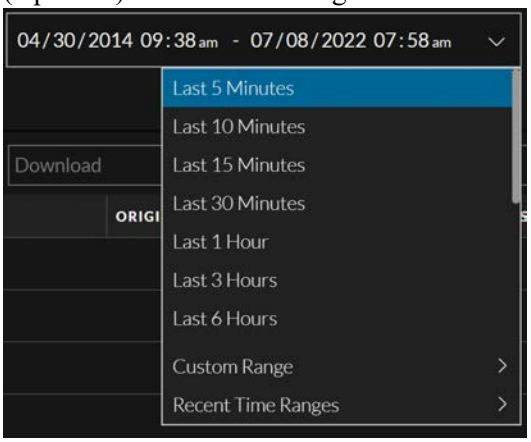
1. Go to **Investigate** > **Events**.
   The Events view opens with a service selected and no data displayed. A drop-down list offers a list of available services in alphabetical order. The **Select a service** field is populated with the first service in the list or the most recently selected service. By default the list of available services is retrieved every twelve hours and cached on the NetWitness server. If a service is added or removed from the NetWitness server before the next time to retrieve, the cache is updated with the latest list of services. An icon provides the status of the service.
   - and selected service name = The service is selected.
   - = Investigate is attempting to connect to the selected service.
   - = There was an error connecting to the selected service or there is no data in the selected service. In this state, the service selector control also turns red, and a tooltip explains why the connection attempt failed and advises you to choose another service.

2. (Optional) Select a service, usually a Broker or Concentrator, from the drop-down list.



The time range selector shows either the default time range of 24 hours, or the time range that you last selected for this service. The or Query Events button becomes active and you can create filters. If you launch a query without creating filters, the selected time is used.

3. (Optional) Edit the time range as described in Filter Results in the Events View.



The selected time range is stored in your browser for this service; you can set different time ranges for different services.

4. Create a query that consists of one or more filters that contain a meta key, operator, and optional value. See Filter Results in the Events View for details on creating queries.

5. When ready to submit the query, click or **Query Events**.
The Events view displays the data for the selected service, time range, and query, in accordance with permissions assigned to your role by the administrator. You are ready to begin analyzing the data. Refer to Examine Event Details in the Events View and Analyze Events in the Events View to learn how to work in the Events view.

Refining the Results Set

When conducting an investigation, results load faster and it is easier to find what you are looking for if you refine the results to get a smaller number of results. In addition, limiting the time range and submitting a good query gives you more relevant results to answer the question at hand. Use a combination of the methods described in the rest of this section to get the information you need quickly.

- Use Meta Groups to Focus on Relevant Meta Keys
- Use Columns and Column Groups in the Events List
- Use Query Profiles to Encapsulate Common Areas for Investigation
- Filter Results in the Events View
- Filter Results in the Navigate View
- Filter Results in the Legacy Events View
- Create a Query in the Navigate and Legacy Events Views
- View and Modify Queries Using URL Integration
- Search for Text Patterns in the Navigate and Legacy Events Views

Use Meta Groups to Focus on Relevant Meta Keys

A meta group combines selected meta keys and meta entities into a group to show only data in which the meta keys and meta entities were found. In the Navigate view and the Version 11.5 and later Events view, you can use meta groups to filter data displayed in the Navigate view (Values panel) and the Events view (Filter Events panel). The same shared meta groups are available for use in both views. Private meta groups created in the Events view are not available for use in the Navigate view or in query profiles in the Legacy Events view.

Note: In the Navigate view and Legacy Events view, you can manually add non-indexed meta keys (or keys that are not in the index at all) to a meta group or column group. The non-indexed meta keys are fully available (manageable and displayable) in the Navigate view and Legacy Events view, but only partially (displayable in the Filter Events panel) in the Events view. The Events view (Filter Events panel) can display data for non-indexed meta keys that are already included in a meta group, but you cannot add non-indexed meta keys while you are editing a meta group. The non-indexed meta keys in a column group do not display data in a column and new non-indexed meta keys cannot be added to a column group in Events view.

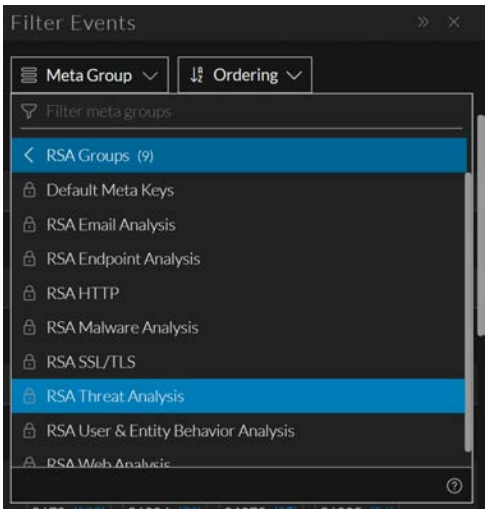
With a meta group in effect during an investigation, the information in the Values panel or the Filter Events panel shows only the meta keys in the selected group. When you open a Parallel Coordinates visualization in the Navigate view, the meta keys and meta entities in a group appear as axes from left to right. It may be useful to create two versions of each custom meta group; one for analysis of meta values and one for creating a parallel coordinates chart focusing on a smaller subset of the same use case.

A fresh installation of NetWitness includes built-in meta groups to help you find interesting data sets in Investigate. The built-in meta groups can be duplicated but cannot be edited or deleted. You can also create your own groups and edit a copy of a built-in group to create a custom group.

All groups in the Navigate view are shared and visible to all users of a service; you can export a group for import to any service, limited by the available meta keys for that service. In the Version 11.5 Events view Filter Events panel, you can create both shared and private custom meta groups; only the shared groups are visible and usable in the Navigate view.

Live Meta Groups

In 11.6 and later, NetWitness supports deploying the investigate content from live. The meta groups are categorized as RSA Groups (RSA Live content and RSA OOTB Groups), and Shared Groups. The content deployed from Live are marked by the live symbol (). The content is displayed in a folder structure. The groups are displayed as non-editable folders and sub-folders. The number inside **()** depicts the number of contents inside a folder and **>** symbol helps you to drill down inside the folder.



Built-In Meta Groups

NetWitness has built-in meta groups, prefixed with RSA, that are available immediately after installation. The built-in meta groups are useful to focus an investigation on common use cases and to support threat detection using the RSA Hunting Pack. You can copy these groups, give the copy a new name, then edit the copy. These are the built-in meta groups:

- RSA Email Analysis includes meta keys that outline email interactions.
- RSA Endpoint Analysis contains meta keys that provide insight on processes, files, users, and connections from NetWitness Endpoint (NWE) hosts.
- RSA Malware Analysis includes meta keys that mark indicators of compromise in files contained in events.
- RSA HTTP includes meta keys that provide insight into outbound web traffic.
- RSA SSL/TLS includes meta keys that focus on encrypted web traffic.
- RSA Threat Analysis includes meta keys that mark potential threats in the data set.
- RSA User & Entity Behavior Analysis includes meta keys that encompass all the meta keys to analyze user and entity behavior.
- RSA Web Analysis includes meta keys that mark anomalies in web traffic.

Default Meta Keys Group (Version 11.5 Events View)

The Default Meta Keys meta group is a special type of built-in meta group that consists of all the meta keys for the currently selected service, returned in the order of appearance in the index file for the service. Unlike the other built-in meta groups, you cannot copy this group and you cannot see which keys are included when you view information in the Meta Group Details dialog; instead, a message in the Details dialog explains that the group includes all meta keys for the selected service The Default Meta Keys group is always at the top of the list in the Meta Groups menu.

The Default Meta Keys group is used to select meta keys shown in the Filter Events panel when no meta group has been selected and none exists in local storage. You can also select this group as you would any other group. When using the Default Meta Keys group in the Filter Events panel, only the first 30 meta keys with values are open and the remaining are closed.

Custom Meta Groups

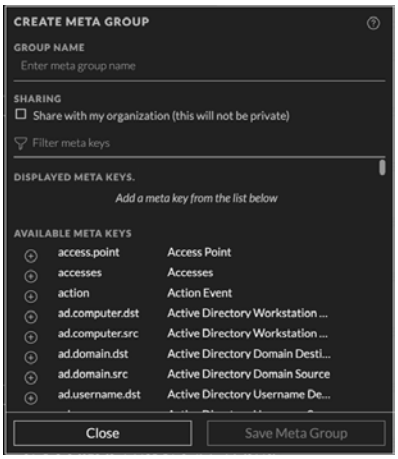
You can create custom meta groups to support scenarios that you use frequently while working in Investigate. When an administrator adds custom meta groups manually by editing the custom index file for a service, the new meta groups become available to use in meta groups after the service is restarted. Custom meta groups can be shared or private . Shared meta groups are available globally within your organization in the Navigate view and in the Filter Events panel. If you edit a shared custom meta group, your changes are applied globally. If you delete a shared custom meta group, the group is deleted and no longer available for all analysts. The Navigate view supports only shared groups. When you create custom meta group in the Events view, you can choose to share it or you can keep it private (default); you cannot change a shared group to private or a private group to shared.

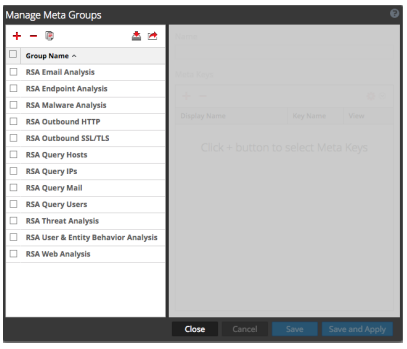
Note: Private custom meta groups created in the Events view are not visible or usable in the Navigate view.

Icons identify the group type in the Meta Groups menu. These are examples of each type of custom meta group with the edit icon displayed at the end of the row.



While the functionality of meta groups is similar in the Navigate view and the Events view, the user interface and some of the procedures are different. The following figures illustrate the (Events view) Create Meta Group dialog and the (Navigate view) Manage Meta Groups dialog.





Using options in the Events view Meta Groups menu (Version 11.5 and later) , you can:

- Select a meta group to apply.
- See the details of a meta group.
- Create, edit, and delete custom meta groups.
- Copy and edit the copy of a built-in or custom meta group.

Using options In the Navigate view Manage Meta Groups dialog, you can do all of the above as well as import and export a meta group.

The rest of this topic provides instructions for working with meta groups in the 11.5 Events view and the Navigate view.

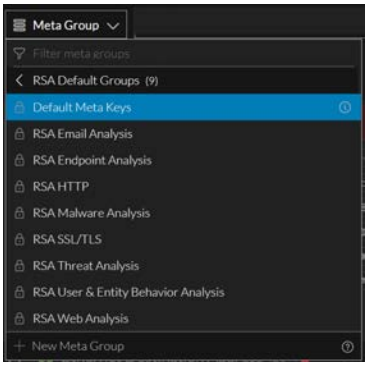
Work with Meta Groups in the Events View (Version 11.5 and Later)

After the upgrade to Version 11.5 or later, all of the existing meta groups -- both built-in and custom -- are available for filtering events in the Filter Events panel. The meta group selection persists between logins unless browser cache is cleared.

View the Meta Keys in a Meta Group

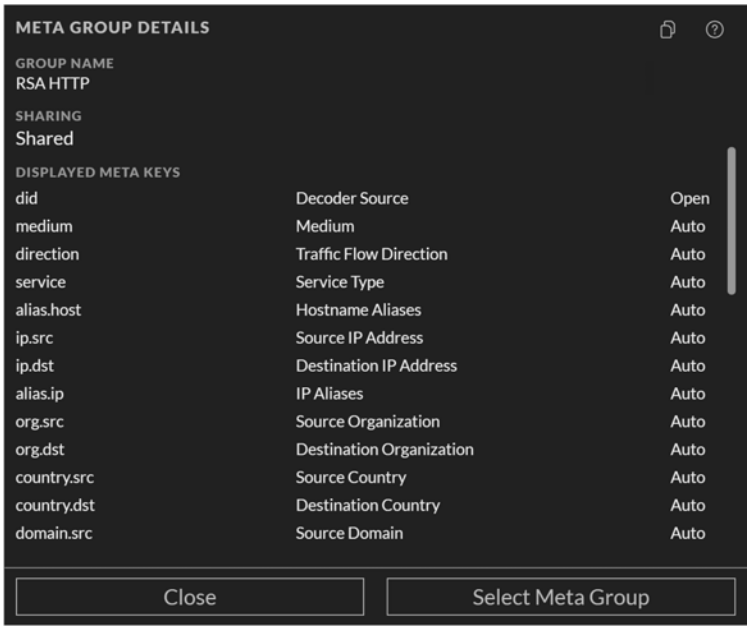
To view details of a meta group:

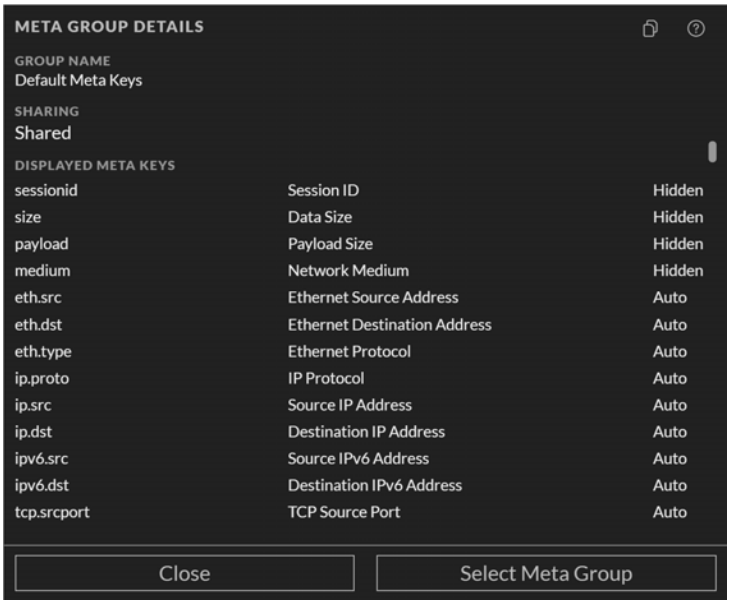
1. Go to **Investigate > Events** and click  to load events.
   The events for the default service and the default time range are loaded in the Events panel.
2. To display the Filter Events panel, click  above the Events panel.
   The Filter Events panel opens to the left of the Events panel.
3. To display the Meta Groups menu, click the Meta Groups menu title. The menu title is either Meta Group: Default Meta Keys or Meta Group: <currently selected meta group>. If this is your first visit after logging in, the Default Meta Key group is selected; any subsequent visits use the meta group selected in the previous session. If the selected meta group from the previous session is deleted, the Default Meta Keys group is selected when you log in. When opened, the menu displays a list of built-in meta groups (RSA), shared custom meta groups, and your private custom meta groups. Above the list, visibility options and a filter make it easier to find a particular meta group.



4. (Optional) To filter the listed meta groups by name, type some text in the **Filter Meta Groups** field. The list is updated to show only the group names that contain the exact text.

5. Hover over the meta group name and click the information icon ( ) to see which meta keys are included in the group.
   The figure on the left shows the columns for the RSA HTTP meta group. The figure on the right shows the columns for the Default Meta Keys meta group.

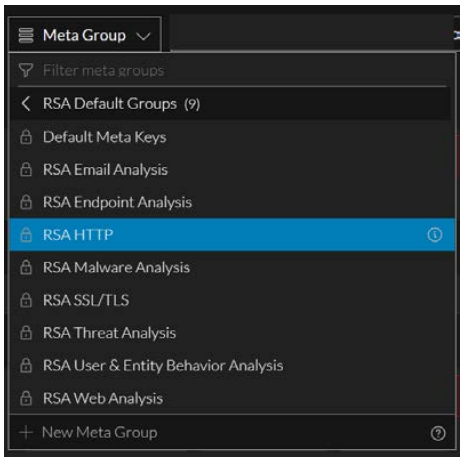




6. Do one of the following.
   a. To close the dialog, click **Close**.
   b. If you want to apply the meta group, click **Select Meta Group**.
      The dialog closes and the Filter Events panel is updated to reflect the meta keys in the selected meta group.

Select a Meta Group

1. With the Filter Events panel open in the Version 11.5 Events view, click the **Meta Groups** menu title.
   The menu drops down to display a list of meta groups and folders with a filtering option and a New Meta Group option. The list is sorted alphabetically and the name of the selected meta group is displayed in the menu label. This figure shows the menu after RSA HTTP was highlighted, but not selected.



2. Do one of the following:
   a. If the highlighted group is the one you want to apply, press **ENTER**.
   b. Begin typing text in the **Filter Meta Groups** field to search for a meta group name. As you type, the list is filtered to show only the meta group names that contain that string.
      When you see the group that you want to apply, click it or use the down or up arrow to highlight it, then press **ENTER**.
      The Filter Events panel is refreshed to include only meta keys in the selected meta group, and the menu title includes the selected group name. Your selection persists when you navigate away from the Events view.

Note: If a meta key in a meta group is not part of the selected service, it does not appear in the Filter Events panel or in the Events panel.

Create a Custom Meta Group

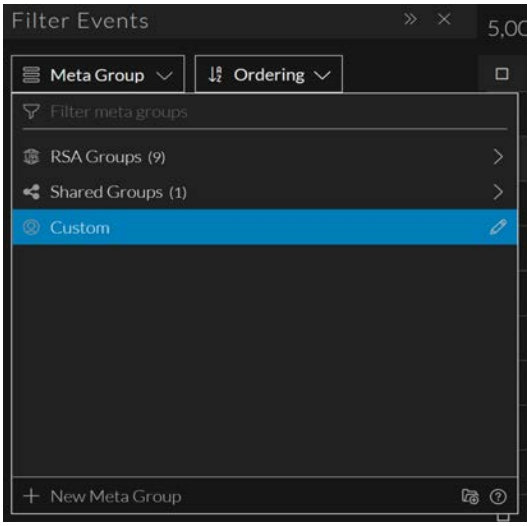
Custom meta groups must have a unique name up to 80 characters in length, and must have at least one meta key. If any other meta group has the name you type, whether shared or private, a message informs you that you need to use a different name. The Save Meta Group button is enabled when these criteria have been met. You can adjust the order of meta keys in a group by dragging keys in the Displayed Meta Keys list.

You can also set the initial view of each meta key: Open, Closed, Hidden, or Auto (the default setting).

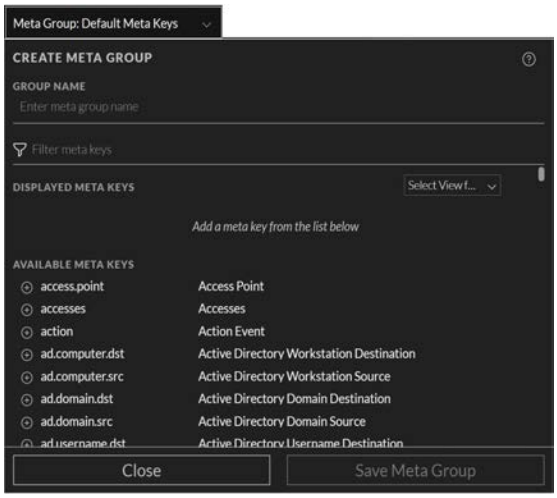
Note: You can also set the desired value for all meta keys at once. Make a note that changing the value of all meta keys might impact the performance.

- When set to Auto, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are Closed until opened manually. If you change the default view for a group of meta keys to Open and some of the meta keys are non-indexed, the non-indexed meta keys revert to Auto.
- Open meta keys are listed in the Filter Events panel, and the values are loaded.
- Closed meta keys are listed in the Filter Events panel, but the meta values are not loaded until you open the meta key.
- Hidden meta keys are not listed in the Filter Events panel at all. This is useful if you are using a single meta group for multiple purposes instead of creating several meta groups; you can turn off certain keys off without removing them from the meta group. You can also use the Hidden view when testing out some new keys or if you want to prepare a meta group with some new meta keys that are not yet available and would error out if in an Auto, Open, or Closed state.

1. With the Filter Events panel open in the 11.5 Events view, click the **Meta Groups** menu title.
   The menu drops down to display a list of meta groups and folders with the Filter Meta Groups field at the top and the + New Meta Group and Folder icon option at the bottom.

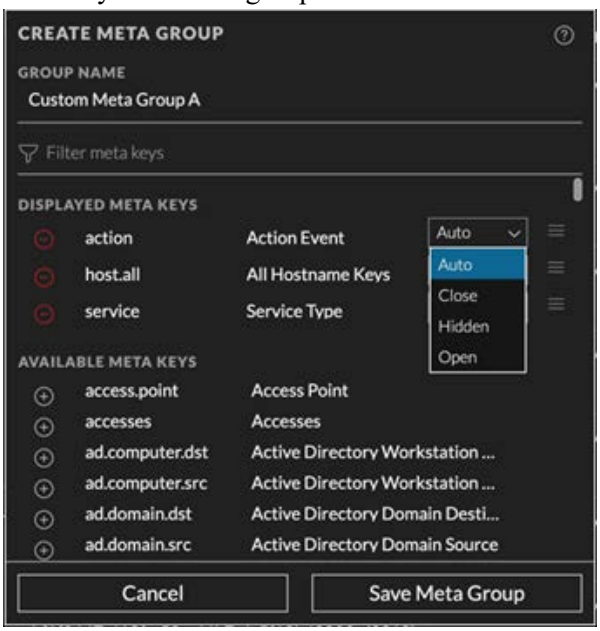


2. Select + **New Meta Group**.
   The Create Meta Group dialog is displayed.



3. In the **Group Name** field, type a unique name (maximum length of 80 characters) for the new meta group, for example, **Custom Meta Group A**.
4. If you want to share the new meta group with your organization, set the **Share with my organization** option.
5. To add a meta key to the meta group, select and add each meta key as follows:
   a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list.
   b. When you see the meta key that you want to add, click the add icon  that precedes the meta key name.
      The meta key is added to the end of the Displayed Meta Keys list. (This list is also filtered using the text you typed.) The maximum number of meta keys in a meta group is 500. If you attempt to add another meta key when 500 are already included in the Displayed Meta Keys list, a message

advises you that the group has the maximum number of meta keys.



6. (Optional) Next to each meta key, choose the initial view for the meta key: **Open**, **Close**, **Hidden**, or **Auto**.
7. (Optional) To find and remove a meta key from the meta group, type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Displayed Meta Keys** list. When you see the meta key that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
8. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
9. Do one of the following:
   a. To close the dialog without creating the custom meta group, click **Cancel**.
   b. To create the group, click **Save Meta Group**.
      The new meta group is saved. If the new group is shared, it becomes available for all analysts. If it is private, only you can use the meta group. The buttons change to Done and Select Meta Group.
10. Do one of the following:
    a. To close the dialog, click **Done**.
    b. To close the dialog and select the new meta group, click **Select Meta Group**.
       The new group is added to the Meta Groups menu (in alphabetical order), and if you clicked Select Meta Group, the Filter Events panel is updated to show the meta keys and values in the new meta group.

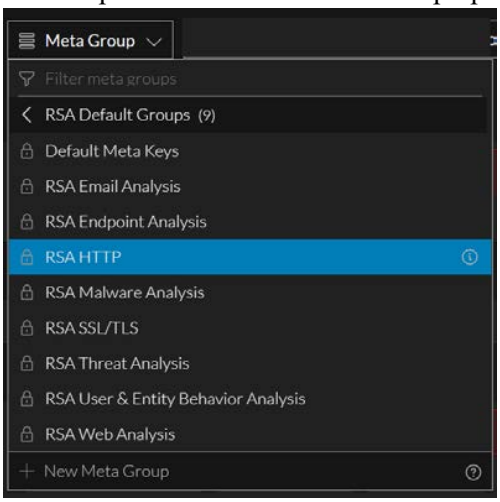
Delete a Custom Meta Group

You can delete any custom meta group, shared or private, that is not currently applied in the Events list and not used in a query profile. When you click the Delete button, a confirmation message allows you to confirm or cancel the deletion. If a meta group is being used in a query profile, the Delete button is disabled and a message identifies the query profile in which the meta group is used. The built-in meta groups are read only, and cannot be deleted.

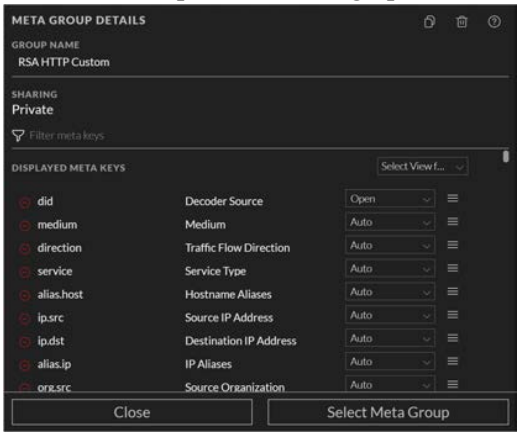
Caution: When you delete a shared meta group, the effect is global and the group is no longer available to any analyst.

To delete a custom meta group:

1. With the Filter Events panel open in the 11.6 Events view, click the **Meta Group** menu title.
   The menu drops down to display a list of meta groups and folders with the Filter Meta Groups field at the top and the + New Meta Group option at the bottom.



2. To delete a meta group, highlight a custom meta group and click the edit icon ( ) to the right of the name.
3. The Meta Group Details dialog opens with the details for the selected group displayed.



4. Click the delete group icon ( ).
   If the meta group is currently in effect, the following message is displayed: `This meta group`

`cannot be deleted because it is currently active.`

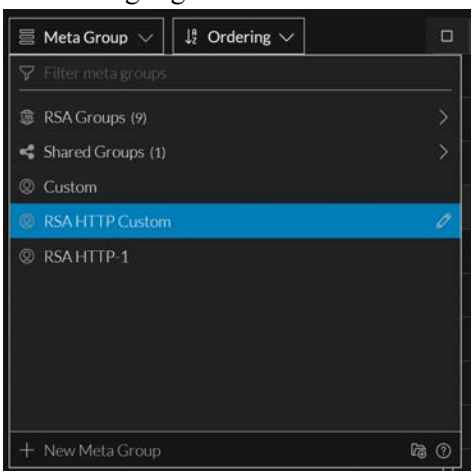
In Version 11.5, a confirmation message gives you the opportunity to confirm or cancel the deletion. Click **Cancel** or **Delete Meta Group**.

The group is deleted and removed from the Meta Group menu. The meta group no longer appears anywhere for any analyst working in Investigate.

Edit a Custom Meta Group

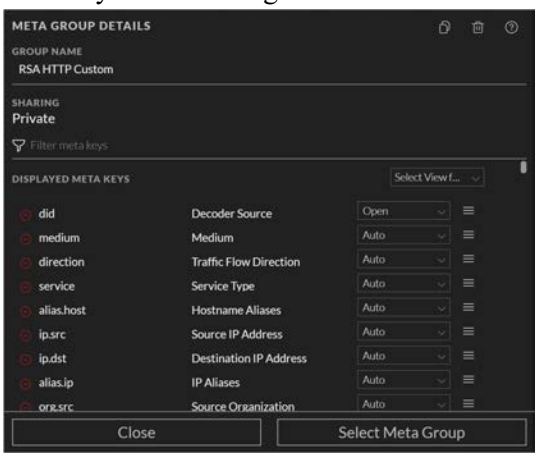
You can edit a shared custom meta group, your own private meta group, a copy of a built-in meta group or a copy of live meta groups.

1. With the Filter Events panel open in the 11.5 Events view, click the **Meta Group** menu title and highlight the meta group that you want edit. This figure shows private column group RSA HTTP Custom highlighted with the edit icon is displayed to the right.



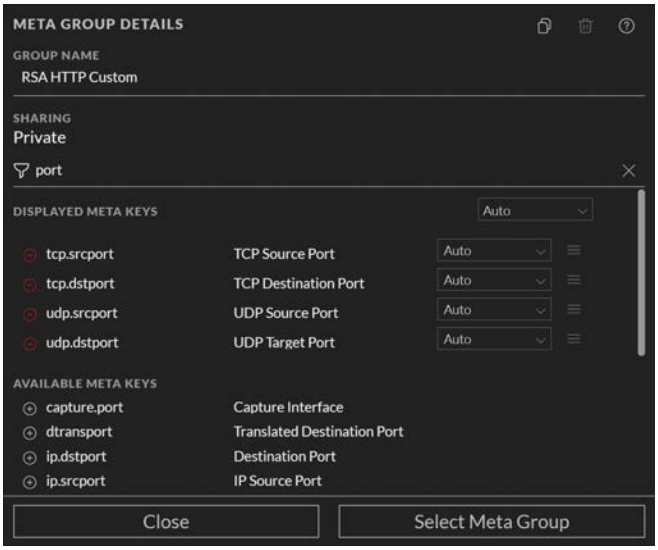
2. Click the edit icon ( ).

   The Meta Group Details dialog is displayed so that you can edit the location. You can add or delete meta keys and rearrange the order of the meta keys in the list.



3. (Optional) In the **Group Name** field, edit the name and location of the meta group.
4. (Optional) To add a meta key to the meta group, select and add each meta key as follows:

a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list. Or just scroll through the list to find the meta key. For example, type **port** in the **Filter meta keys** field.



b. When you see the meta key that you want to add, click the add icon that precedes the meta key name.

5. (Optional) To find and remove a meta key from the meta group, type a text string in the **Filter meta keys** field to look for meta keys that contain that text in the **Dispayed Meta Keys** list, or simply scroll through the list. When you see the meta key that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
6. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
7. Do one of the following:
   a. To close the dialog without saving the changes to the custom meta group, click **Reset**.
   b. To save the edits to the meta group, click **Update Meta Group**.
   The updated meta group is saved, and the dialog is closed.

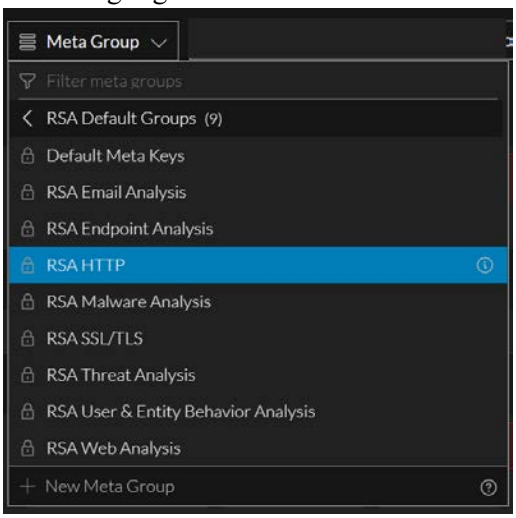
Copy a Meta Group (Version 11.5 and Later)

You can copy any meta group, built-in or custom, Live meta group, shared or private, as long as it does not have unsaved edits in progress. This is useful when you want a customized version of a built-in group. Also since you cannot change a custom group from private to shared or from shared to private, creating a copy allows you to select a different Sharing setting. When you copy a meta group, the same name is used with a number appended. For example, if you copy RSA HTTP twice, the first copy is named RSA HTTP-1, and a second copy is named RSA HTTP-2. After you copy the group, you can edit the copy to give it a new name and manage meta keys in the group.

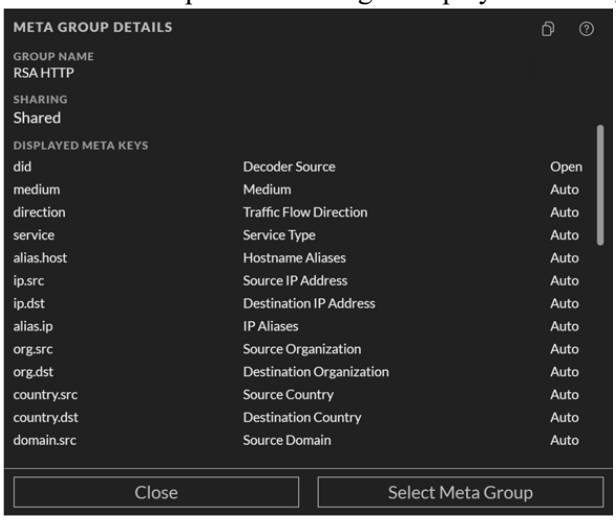
Note: Some meta groups created in the Legacy Events view may have more 500 meta keys, which is above the limit for meta groups in the Events view. If you copy a group with more than 500 meta keys, you must remove the excess meta keys when you edit the meta group.

To copy a meta group:

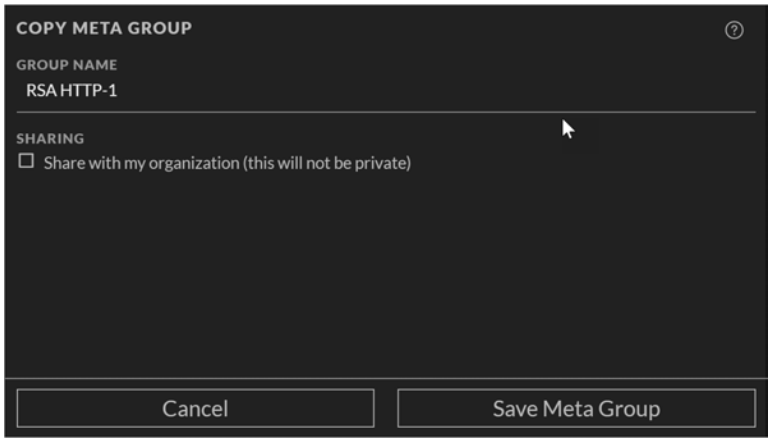
1. With the Filter Events panel open in the 11.6 Events view, click the **Meta Group** menu title.
   The menu drops down to display a list of meta groups and folders.
2. Highlight the meta group that you want copy.
   If you highlighted a built-in meta group, the information icon ( ) is displayed to the right. If you highlighted a custom meta group, the edit icon ( )is displayed to the right. This figure shows RSA HTTP highlighted.



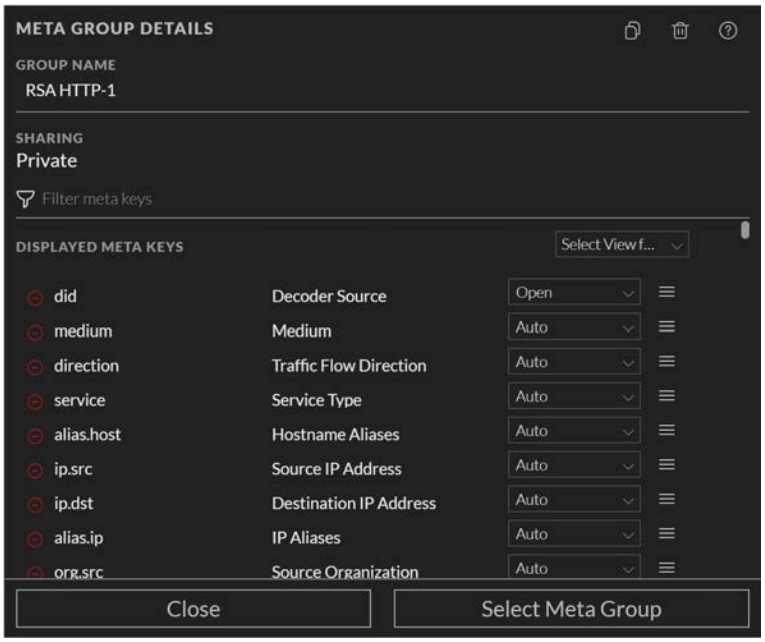
3. Do one of the following:
   a. Click the information icon ( ).
   b. Click the edit icon ( ).
      The Meta Group Details dialog is displayed. This figure shows the dialog for a built-in group.



4. Click the Copy icon ( ).
   The Copy Meta Group dialog is displayed with a `-n` appended to the original meta group name.

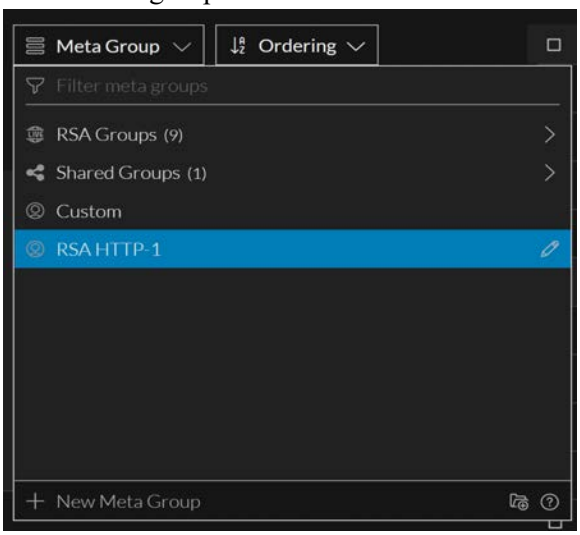


5. (Optional) In the **Group Name** field, edit the name and location of the meta group.
6. Do one of the following:
   a. To close the dialog without copying the group, click **Cancel**.
   b. To save the copy of the meta group, click **Save Meta Group**.
      The copy of the meta group is saved, and the Meta Group Details dialog for the copied group is displayed.



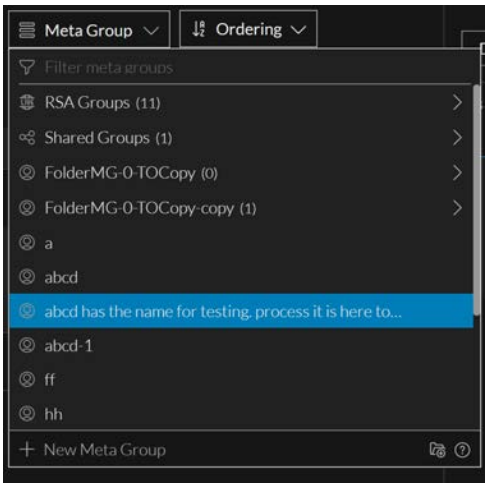
7. Do one of the following:
   a. To close the dialog without editing, click **Close**.
   b. To close the dialog and select the copy of the meta group, click **Select Meta Group**.
      The group is added to the Meta Group menu. The figure below has a private copy of the RSA

HTTP meta group.



Meta Group Folders

Users can create editable Shared and Private group folders. The contents of a private group folder and their contents are displayed outside RSA Groups and Shared group folders. For example, the below image shows private content below the Shared Groups folder.



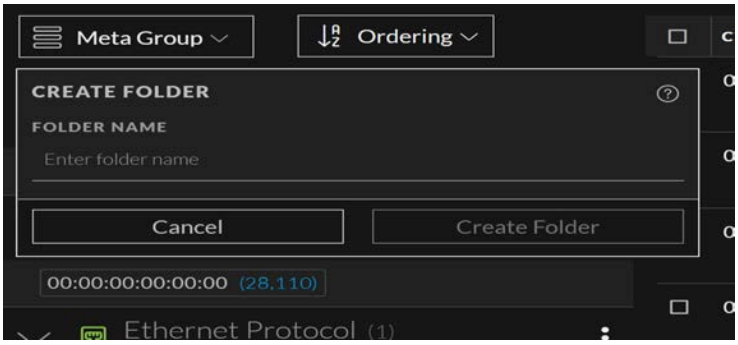
This section describes how to add, edit, import, export, copy, and delete custom meta groups and folders.

Create a Meta Group Folder

You can create meta group folders as shared and private folders. And, if the folder name already exists then you are prompted to provide a unique name.

1. With the Filter Events panel open in the Events view, click the Meta Groups menu title. The menu drops down to display a list of meta groups and folders.

2. Click .
   The Create Folder dialog is displayed.

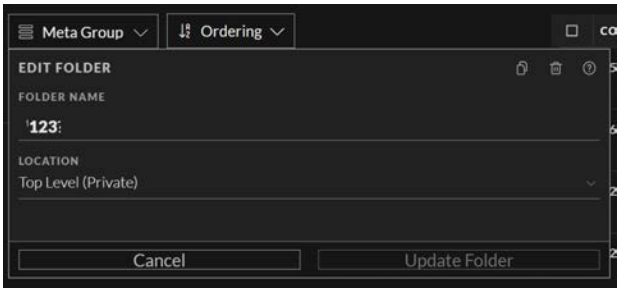


3. In the **Folder Name** field, type a unique name for the new meta group folder.
4. Click **Create Folder**.

Edit and Move Meta Group Folder

After you create a meta group folder you can edit or move it, however the folders inside RSA Groups (RSA Live content and RSA OOTB Groups) cannot be edited and moved. The folders inside private and shared folders can be edited and moved only within their respective groups. For example, you cannot move a shared folder into a private folder and vice-versa.

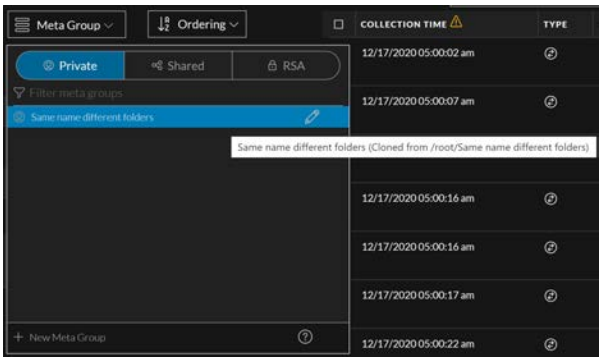
1. With the Filter Events panel open in the 11.6 Events view, click the Meta Group menu title and highlight the meta group that you want edit.
2. Click .
   The Edit Folder dialog is displayed.



3. In the **Folder Name** field, type a unique name for the meta group folder.
4. Select the location of the folder to be edited.
5. Click **Update Folder**.

Copy Meta Group Folder

Users can copy any type of meta group folders namely - RSA , Shared and Private. However, by default for RSA groups copy folder will create a copy in the private section (root level) but can change the location of the folder to a shared folder or any other private folder. The meta group can be copied by clicking the clone icon ( ). After copying, the meta group folders are displayed selected location (Shared or Private category). You can hover over on the copied item to view a tooltip that indicates the path from which the meta group has been copied. In case you need to search for a specific meta group, you can type the name of the meta group in the filter field ( ) at the folder level and the meta group will be filtered from the selected folder.

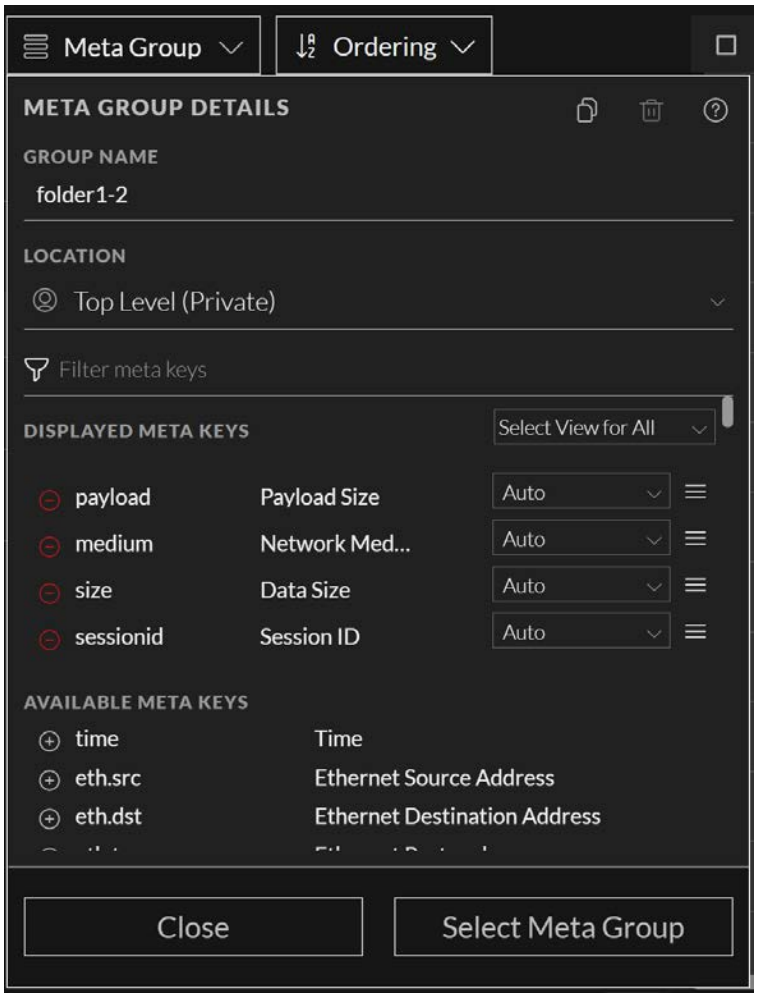


Copying Meta Group Private or Shared Folders

You can copy meta group folders from RSA groups to private and RSA groups to shared, private to shared, private to private, shared to shared and shared to private groups. When you copy a folder the content inside it gets copied except the sub-folders. When you copy a private folder into a shared folder, the folder and its content no longer remain private.

1. With the Filter Events panel open in the Events view, click the Meta Group menu title. The menu drops down to display a list of meta groups and folders.
2. Select a folder you want to copy.

3. Click edit and then click copy .



The Copy Folder dialog is displayed.

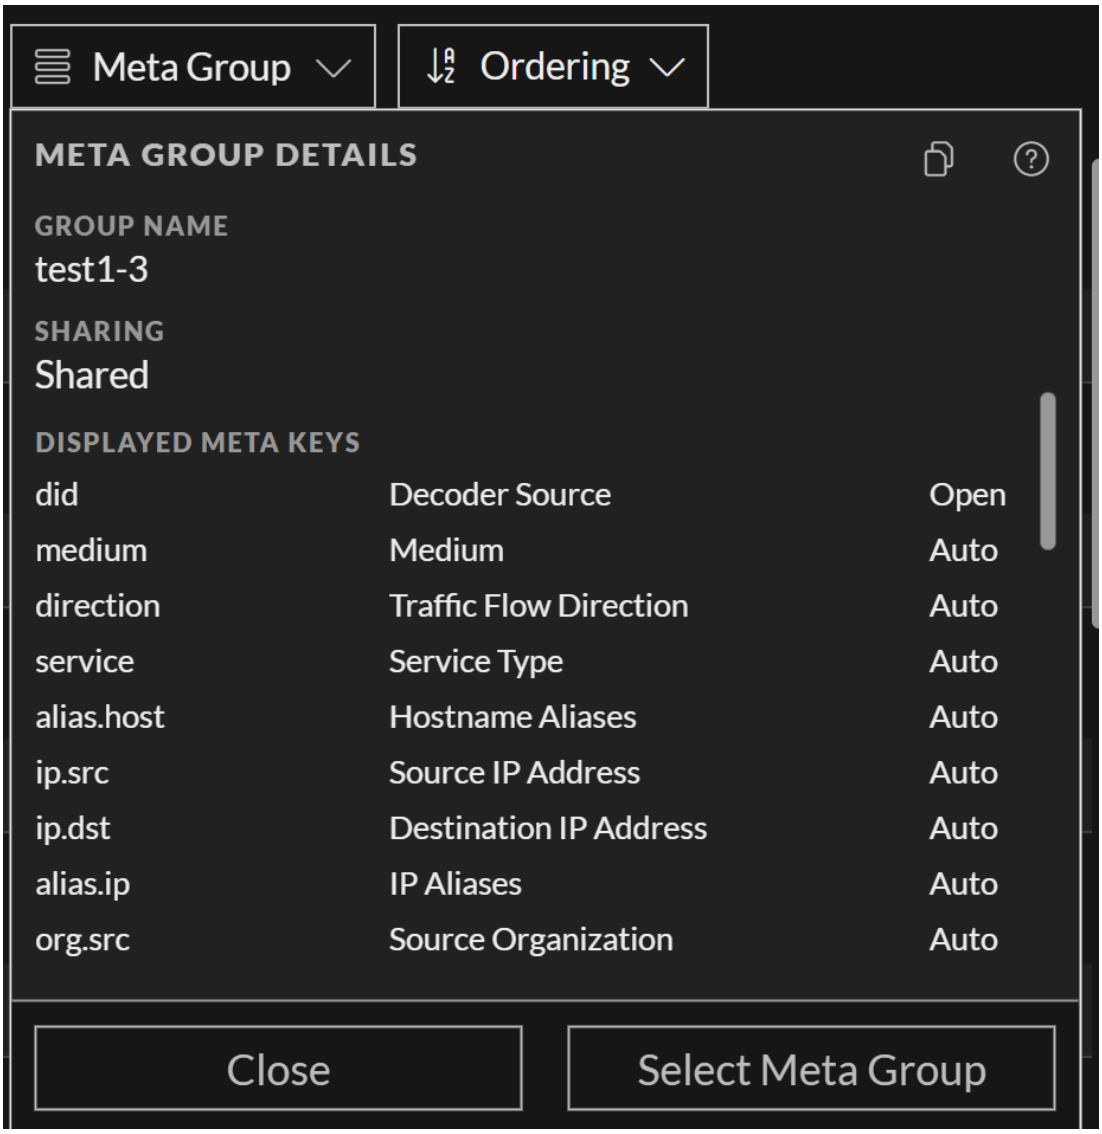


4. In the **Folder Name** field, type a unique name for the new meta group folder.
5. Select the location of the folder to be copied.
6. Click **Copy Folder**.

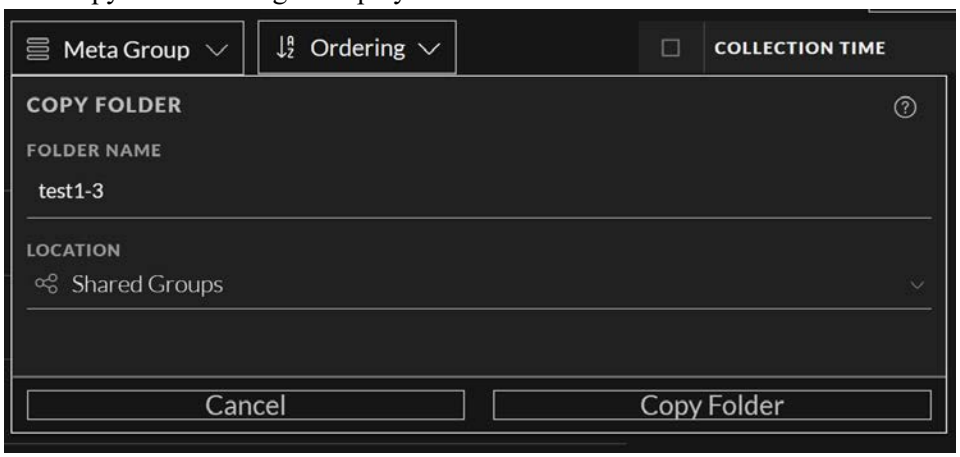
Copy Meta Group Folder Deployed from Live

You can copy meta group folder deployed from Live located under RSA Groups category to any other location like Shared groups or to a private folder.

1. With the Filter Events panel open in the Events view, click Meta Group menu title. The menu drops down to display a list of meta groups and folders.
2. Click on RSA groups and select a Live Meta Group folder you want to copy.
3. Click  and the click copy .



The Copy Folder dialog is displayed.

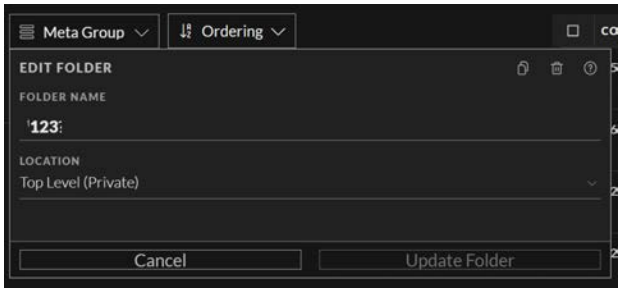


4. Select the location of the folder to be copied.
5. Click **Copy Folder**.
   The folder and the first level contents are copied, sub-Folders are not copied. The copied Meta Group Folder and its contents are displayed as the original meta group name appended with a -n .

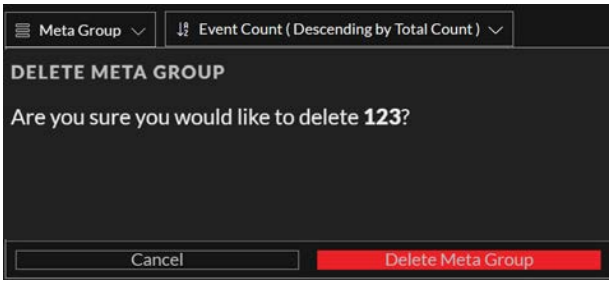
Delete Meta Group Folder

If you don't want to retain a folder you can delete it. However, once the folder is deleted it cannot be retrieved.

1. With the Filter Events panel open in the Events view, click the Meta Group menu title. The menu drops down to display a list of meta groups and folders.
2. Select a folder to be deleted.
3. Click edit .
   The Edit Folder dialog is displayed.



4. Click delete .
   A warning message is displayed to confirm the action.

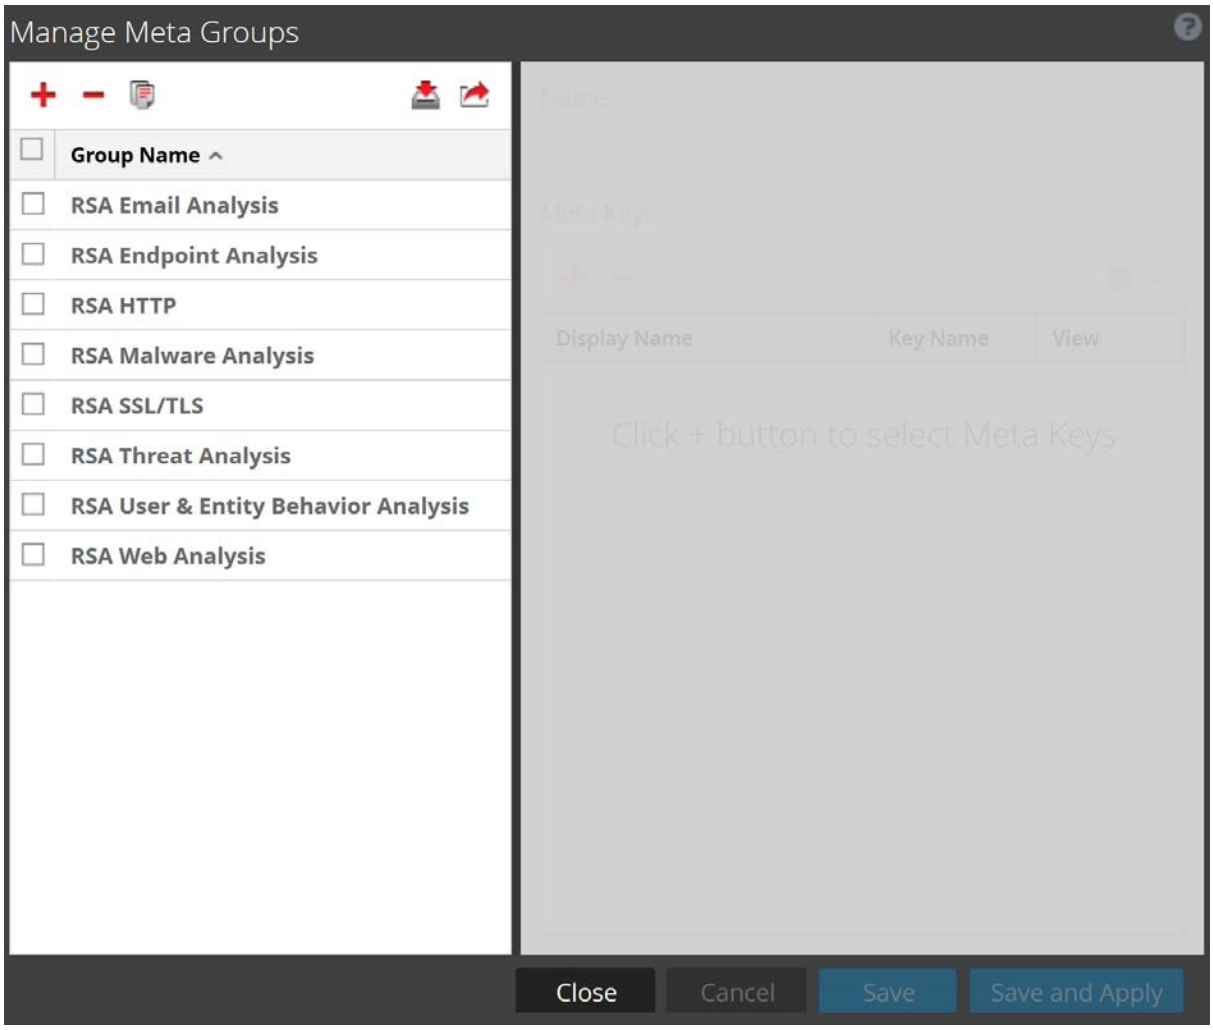


5. (Optional) Select the checkbox, if you want to delete the folder along with all the contents inside the selected folder.
   If you do not select the checkbox, then the content will be moved to the parent folder after the required folder is deleted.
6. Click **OK** to delete.

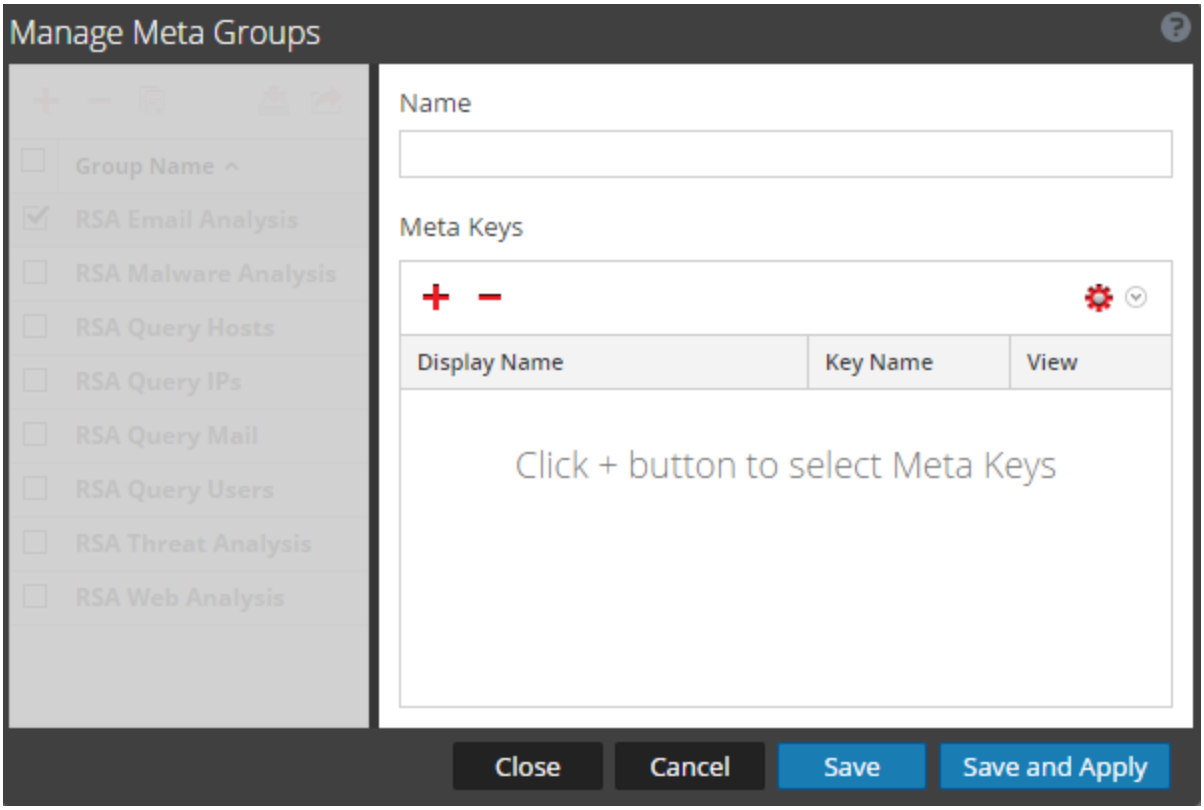
Work with Meta Groups in the Navigate View

Create a Meta Group and Add Meta Keys

1. While investigating a service in the **Navigate** view, select **Meta > Manage Meta Groups** in the toolbar.
   The Manage Meta Groups dialog is displayed. Initially only built-in groups are configured for a service and listed under Group Name. If other custom groups have been configured, they are also listed under Group Name.

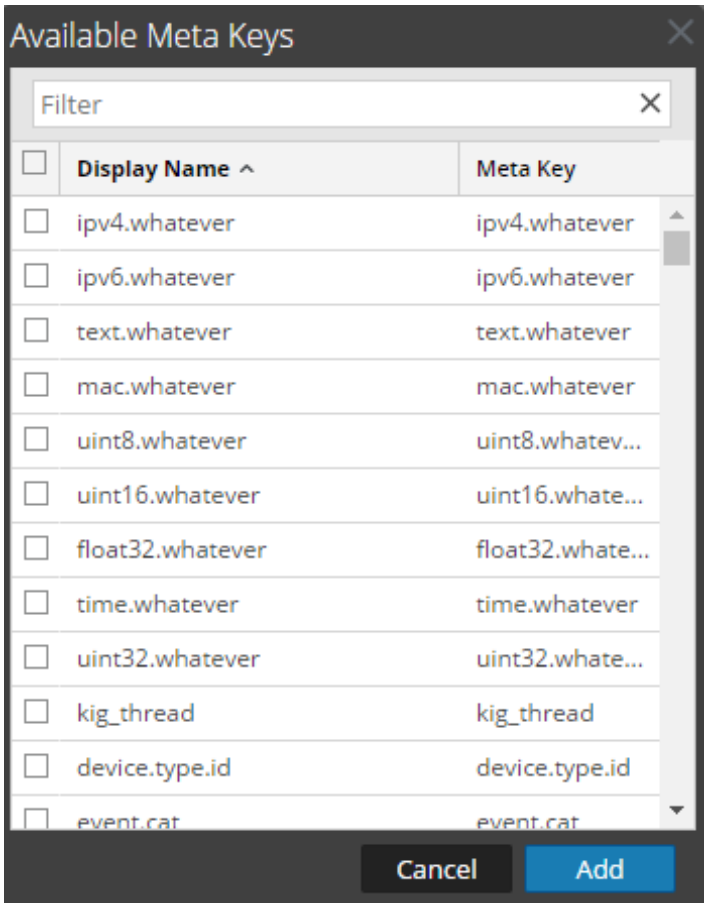

Manage Meta Groups
Group Name
RSA Email Analysis
RSA Endpoint Analysis
RSA HTTP
RSA Malware Analysis
RSA SSL/TLS
RSA Threat Analysis
RSA User & Entity Behavior Analysis
RSA Web Analysis
Name
Meta Keys
Display Name
Key Name
View
Click + button to select Meta Keys
Close
Cancel
Save
Save and Apply

2. In the toolbar at the top of the Meta Groups list, click +.
   The form to the right opens for editing.



3. Type a name for the new meta group in the **Name** field.

4. In the **Meta Keys** toolbar, click +.
   The Available Meta Keys dialog is displayed, with keys in alphabetical order.

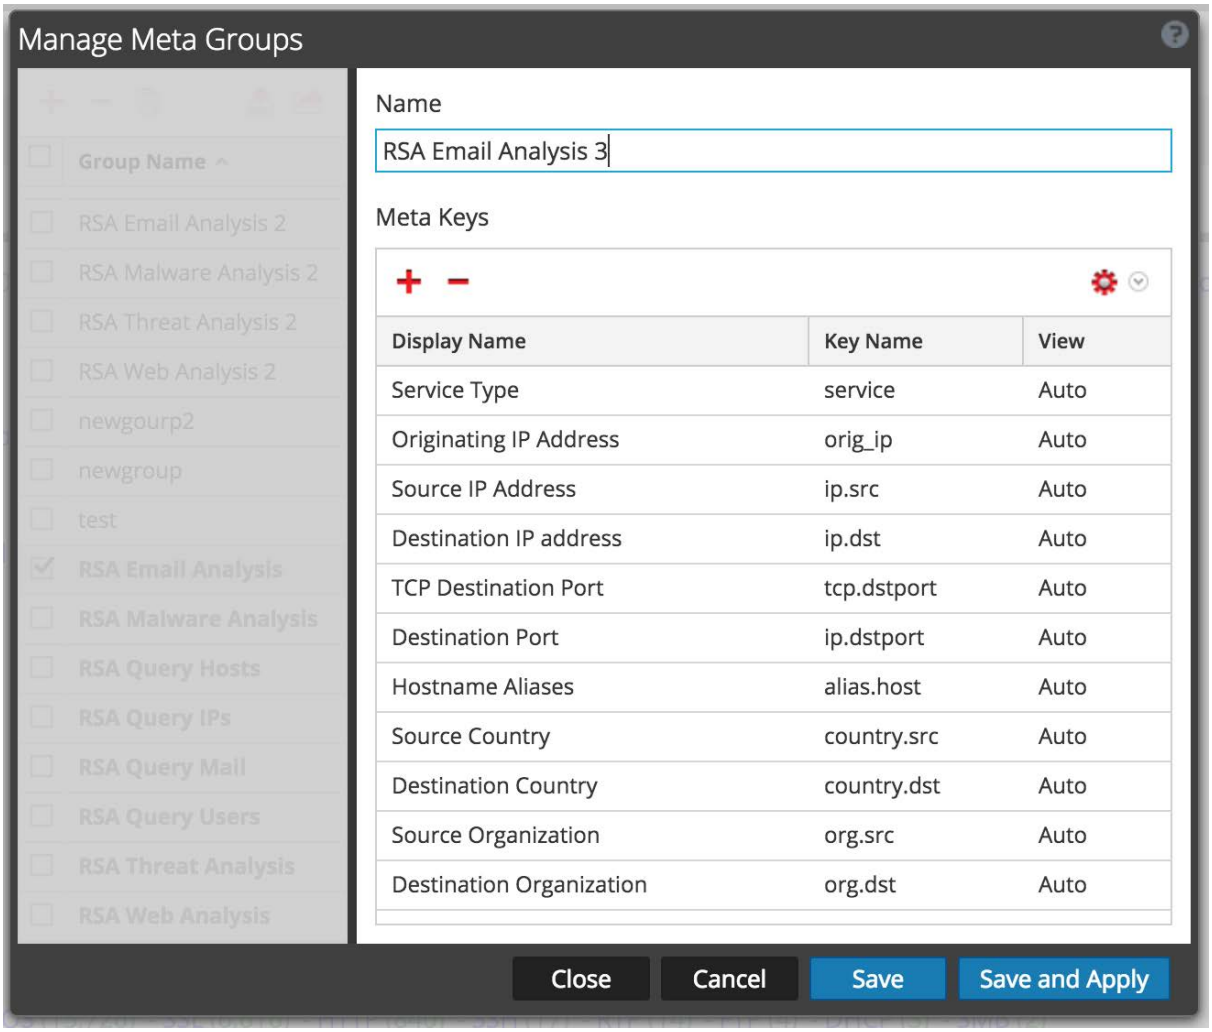


5. To filter the list of meta keys, type a word or phrase in the **Filter** field and press **Enter**.
   The list displays matching meta keys based on a case-insensitive search. Delete the filter text and press **Enter** to remove the filter.
6. To select individual meta keys to include in the meta group, select the checkboxes. To select all meta keys, select the checkbox in the title bar and click **Add**.
   The selected meta keys are added to the meta keys list.
7. (Optional) If you want to change the order in which the meta keys load and are listed in an investigation, click and drag one or more meta keys to a new position.
8. To finish creating the meta group do one of the following:
   a. To save the meta group, click **Save**.
      The group is created and available for use.
   b. To save and apply the meta group to the current Investigation view, click **Save and Apply**.
      The group is created and applied immediately to the current Investigation view.
9. Click **Close**.

Copy and Edit a Meta Group

If you want to customize a built-in meta group, you need to duplicate the group and then edit the duplicate.

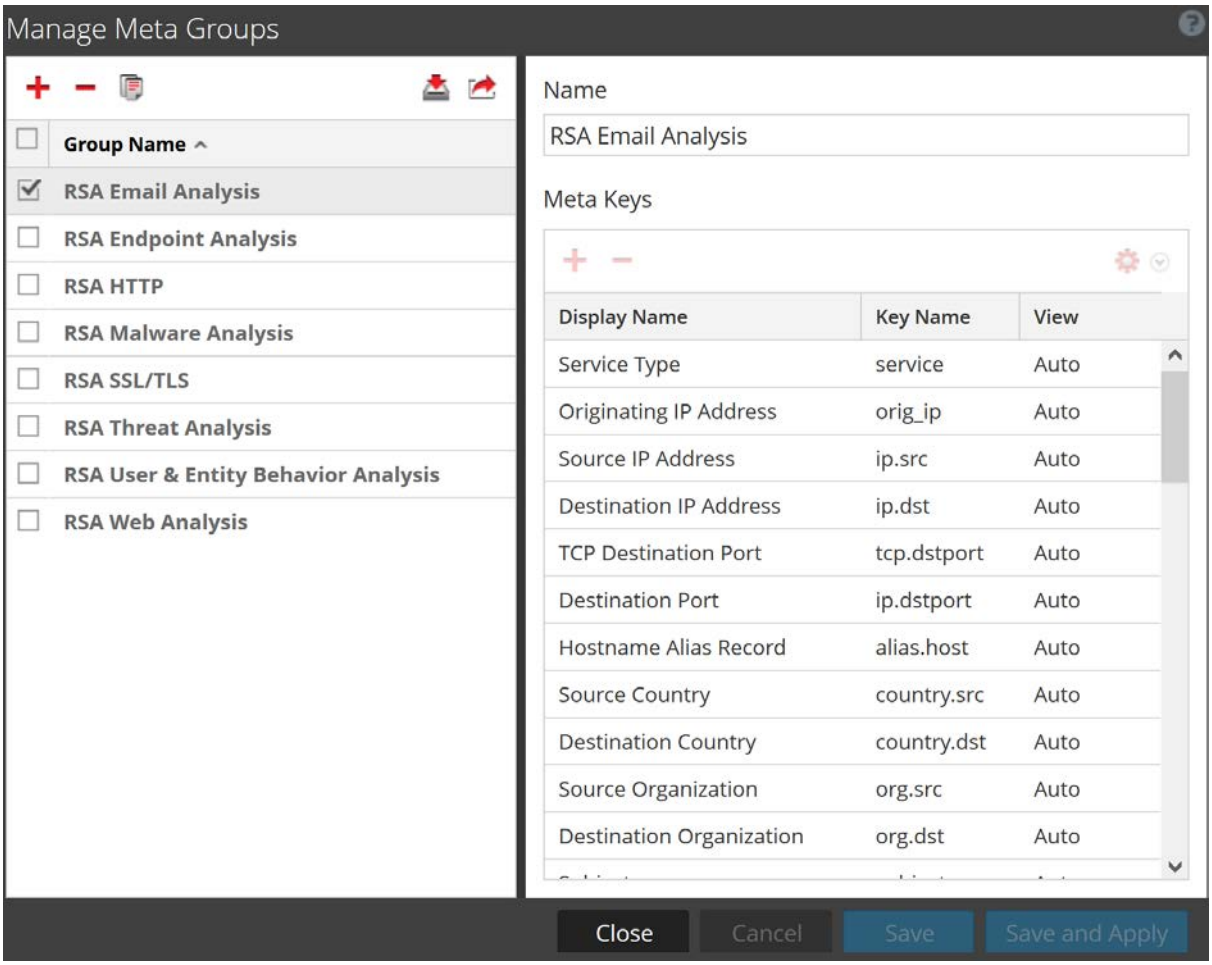
1. Select a built-in meta group from the Manage Meta Groups list and click .
   The form to the right opens for editing with all of the meta keys as they are in the built-in group.



2. Enter a name for the new group and continue editing as described in "Edit a Meta Group" below.

Edit a Custom Meta Group

1. Select a custom group from the **Meta Groups** list.
   The form to the right opens for editing.



2. (Optional) Edit the Name of the group.
3. (Optional) Add new meta keys, as described above in "Create a Meta Group and Add Meta Keys."
4. (Optional) To set the order for the keys, drag and drop one or more keys.
5. (Optional) To change the initial view of a meta key, click  and choose one of the possible views.
   When you modify the meta group, you cannot set the key to OPEN. If you change the default view for a group of meta keys to OPEN and some of the meta keys are non-indexed, the non-indexed meta keys revert to AUTO. As a result, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are CLOSED until opened manually.
   The value for the initial view is displayed in the View column.
6. To save, the changes, click **Save**.
7. To apply the changes to the current Navigate view, click **Save and Apply**.

Delete a Meta Group

1. In the **Meta Groups** list, select the group to be removed.
2. Click −.
   A confirmation dialog provides an opportunity to cancel or complete the request.
3. Click **Yes**.
   The meta group is deleted. When you close the window, if the deleted group was the currently applied meta group, it is removed and the default meta keys are used to build the view.

Export a Meta Group

User-defined meta groups are created on individual services. To make meta groups available to another service, you must export them to your local file system. To export one or more meta groups:

1. In the **Meta Groups** list, select one or more groups to be exported.
2. Click  .
   The selected groups are downloaded to your local file system as a **MetaGroups.jsn file**. Every download of meta groups has the same name with a numeral appended to avoid overwriting previous downloads.

Import a Meta Group

To make user-defined meta groups from another service available to the currently investigated service, you must import the `MetaGroups.jsn` file from the local file system. When you import meta groups, an error message is displayed if any of the groups are already present. To import a group that is a duplicate, you must first delete the existing group. If you want to delete a meta group, it cannot be in use by a profile.

To import meta groups:

1. In the **Meta Groups** list, select a file to import and click  .
   The selection dialog is displayed.

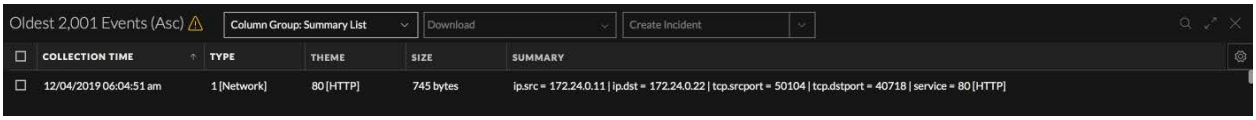


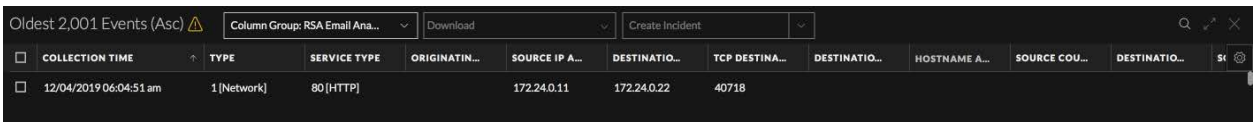
2. Click **Browse** and navigate to the directory on your local file system where the downloaded `MetaGroups.jsn` files are stored. Select a file and click **Open**.
   The filename is displayed in the Upload File field.

3. Click **Upload**.
   The upload process begins, and a message indicates that the upload was successful. The meta groups are added to Meta Group list. If the file is a duplicate of an existing meta group, a dialog tells you that the meta group already exists.

Use Columns and Column Groups in the Events List

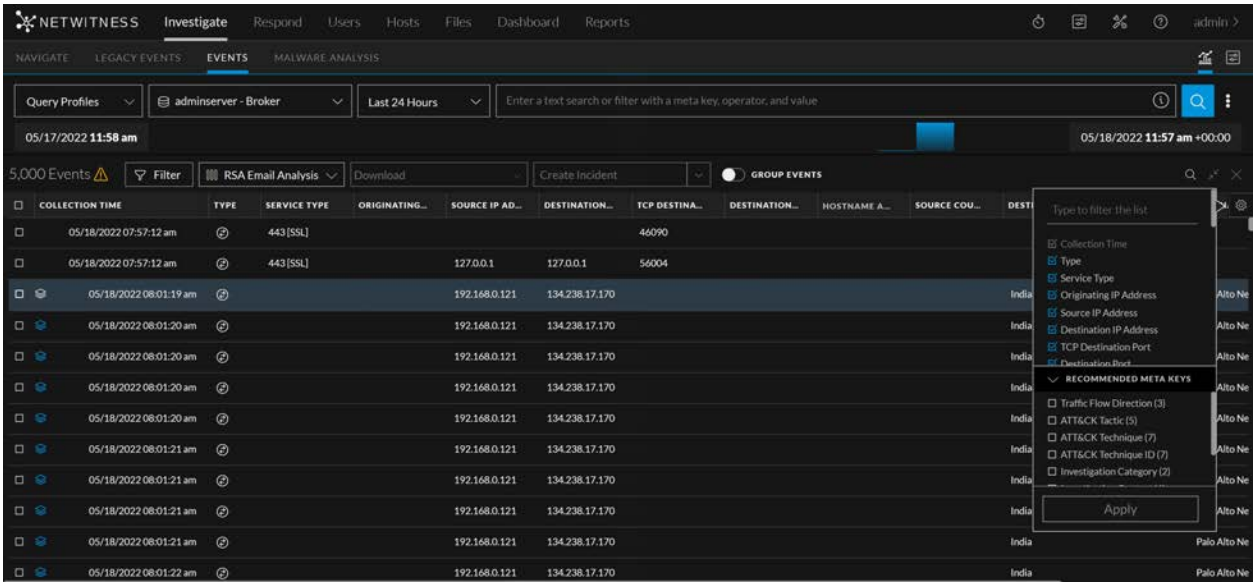
When the events list in Investigate is populated with events, each column lists the values returned for a meta key. Changing the meta keys displayed in the events list is a useful method of narrowing the focus of your investigation. For example, compare these two figures showing the same set of events with different columns. The first figure has five columns, **Collection Time**, **Type**, **Theme**, **Size**, and **Summary**. These are just the basic information, not specialized in any way. The second figure has many more columns that contain information useful when investigating email; you can scroll to the right to see the additional columns.





You can adjust the events list as you work, selecting different columns to be displayed, rearranging the order of the columns, changing the width of the columns, and choosing a column by which the list is sorted. Manual adjustments are easy to make if you know which meta keys are relevant. Manual adjustments apply only to the current session in Version 11.5; in Version 11.5.1, the column width is an exception. When you adjust the column width, it is preserved as a personal preference and is applied every time the column is used in the Events list, overriding any default column width.

In version 11.6, you can select additional columns with data for the meta keys that you are viewing. This will enable you to obtain all the relevant meta key information from the Filter Events panel however the recommendations can change based on the selected meta groups. The following figure displays the additional meta key information under the Recommended Meta Keys section.



To improve your ability to see relevant meta keys quickly when looking at events in the Legacy Events view and the Events view, you can change the set of meta keys displayed by applying a column group. A column group defines the meta keys or meta entities that are displayed as columns, the position of the column in the Events list, and the default width of the column. A column group must have at least one column. Column groups are useful in themselves, and they become even more useful when you combine them with meta groups and preQueries to define query profiles (see Use Query Profiles to Encapsulate Common Areas for Investigation).

The same column groups are shared between the Legacy Events view and the Events view. When importing a column group, the imported group is limited to the available meta keys for the service being investigated. Private column groups created in the Events view are not available in the Legacy Events view or for use in Query Profiles in the Navigate view.

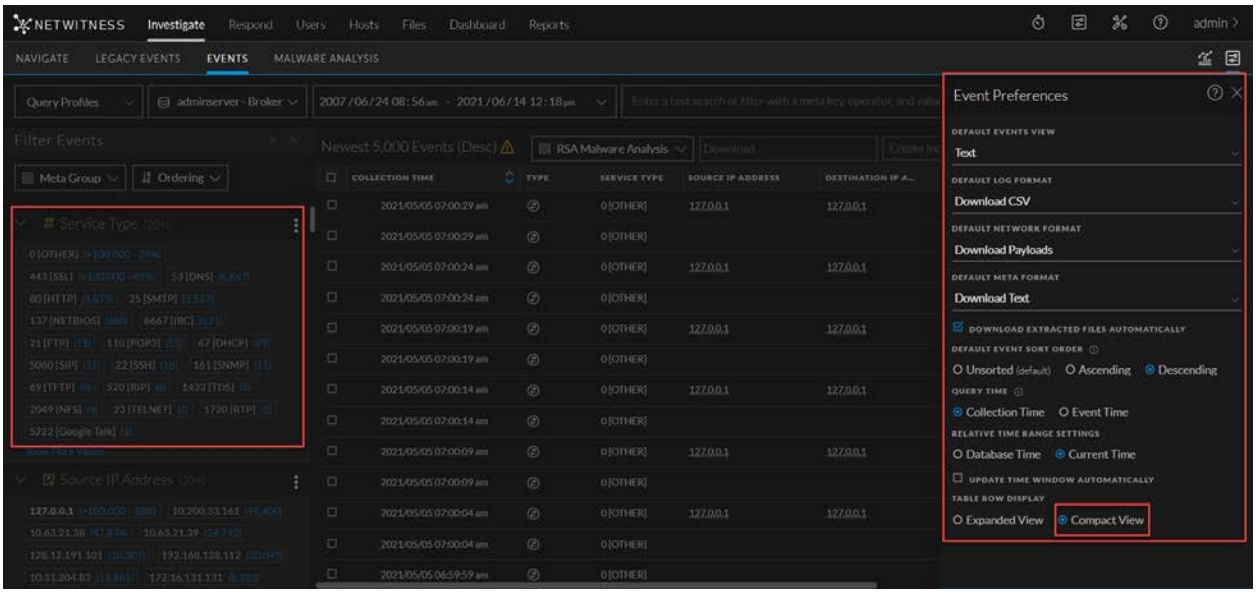
Note: In the Navigate view and Legacy Events view, you can manually add non-indexed meta keys (or keys that are not in the index at all) to a meta group or column group. The non-indexed meta keys are fully available (manageable and displayable) in the Navigate view and Legacy Events view, but only partially (displayable in the Filter Events panel) in the Events view. The Events view Filter Events panel can display data for non-indexed meta keys that are already included in a meta group, but you cannot add non-indexed meta keys while you are editing a meta group. The non-indexed meta keys in a column group do not display data in a column and new non-indexed meta keys cannot be added to a column group in Events view.

Large column groups can have a performance impact when loading data because the values for each meta key are loaded in the events list. To minimize impact on performance, the Events view has a fixed limit on the number of meta keys in a column group. The maximum number of meta keys in a column group is 40. (Because several default meta keys are included you may see a few more than 40 displayed on the screen.) Meta keys that are not in the selected column group are not loaded in the events list. By default we load all columns in the group, but only 15 are displayed by default.

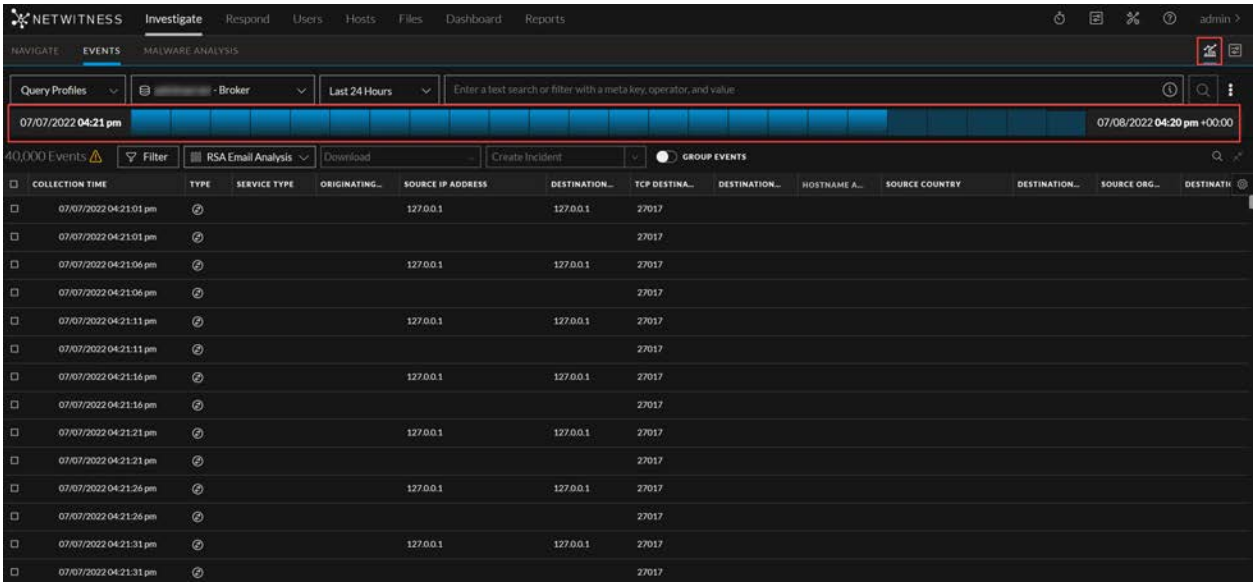
The Legacy Events view does not have a limit on the number of meta keys in a column group, and may have more than 40 meta keys in a column group. If you apply a column group with more than 40 meta keys that was created in the Legacy Events view, all columns are loaded in the Events view. If you copy a group with more than 40 columns, you must remove the excess columns when you edit the column group.

Note: In Version 11.3, column groups are created and managed in the Events view and available for use in the Event Analysis view. All existing column groups, both built-in and custom, are available in the 11.4 Events view. The complete column group management functionality is available in the Legacy Events view, and all functionality except cloning, importing, and exporting column groups is available in the 11.4 Events view. In Version 11.5, cloning is also available in the Events view, but importing and exporting are not.

In 11.6.1, the Investigate > Events Preferences view has been added to make optimum use of the space to enable analysts to view maximum details related to the events they are analyzing.

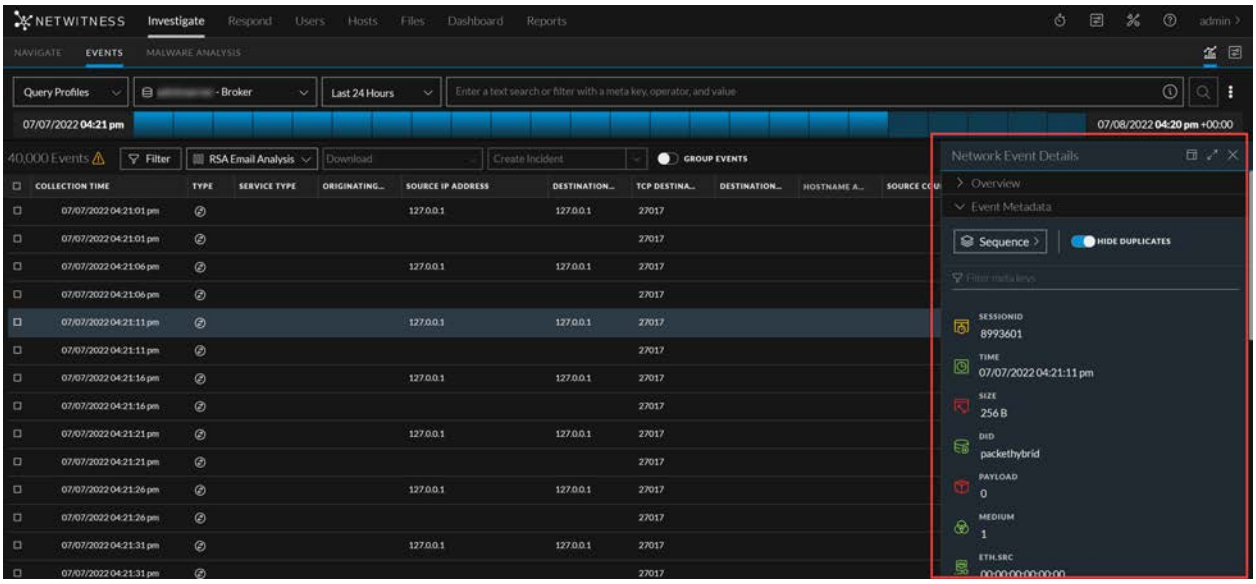


Additionally, an analyst can now view the timeline details of an event by clicking the ■ icon. When clicked the time line that displays the date and time range s displayed as shown in the image.

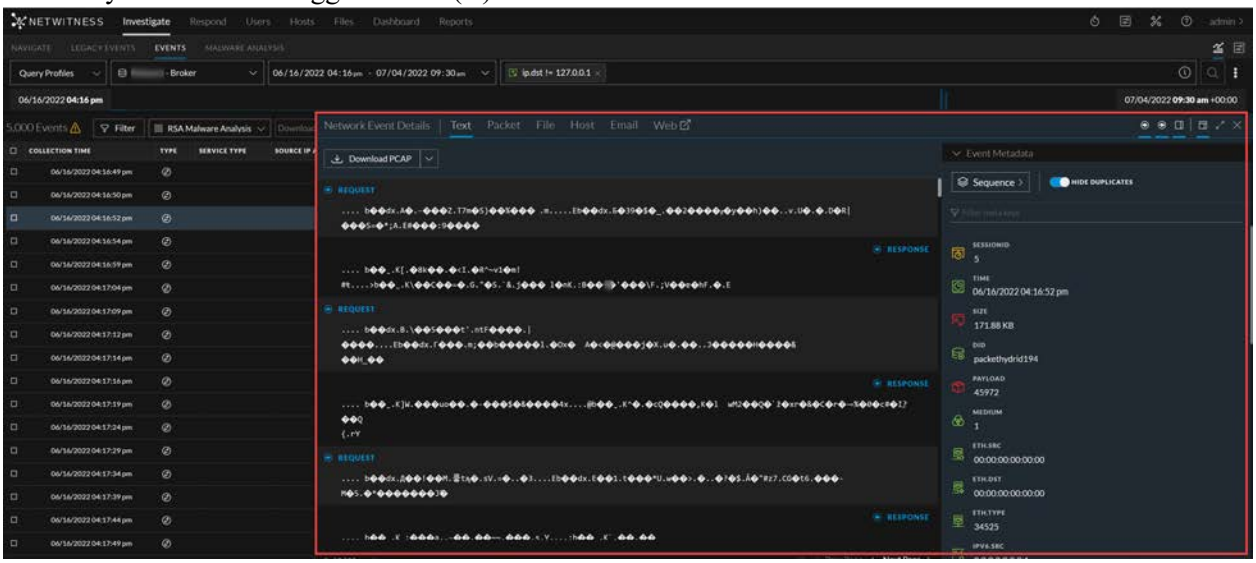


The Investigate > Events reconstruction panel has been modified to display an overlay that will contain an Overview tab and a Meta panel tab that can be expanded or collapsed. This will enable the analyst to view the headers and meta panel of the events optimally. The analysts can also toggle **Hide Duplicate Events** option and view only the relevant details of a selected event.

When an analyst navigates to this page the following view will be displayed.



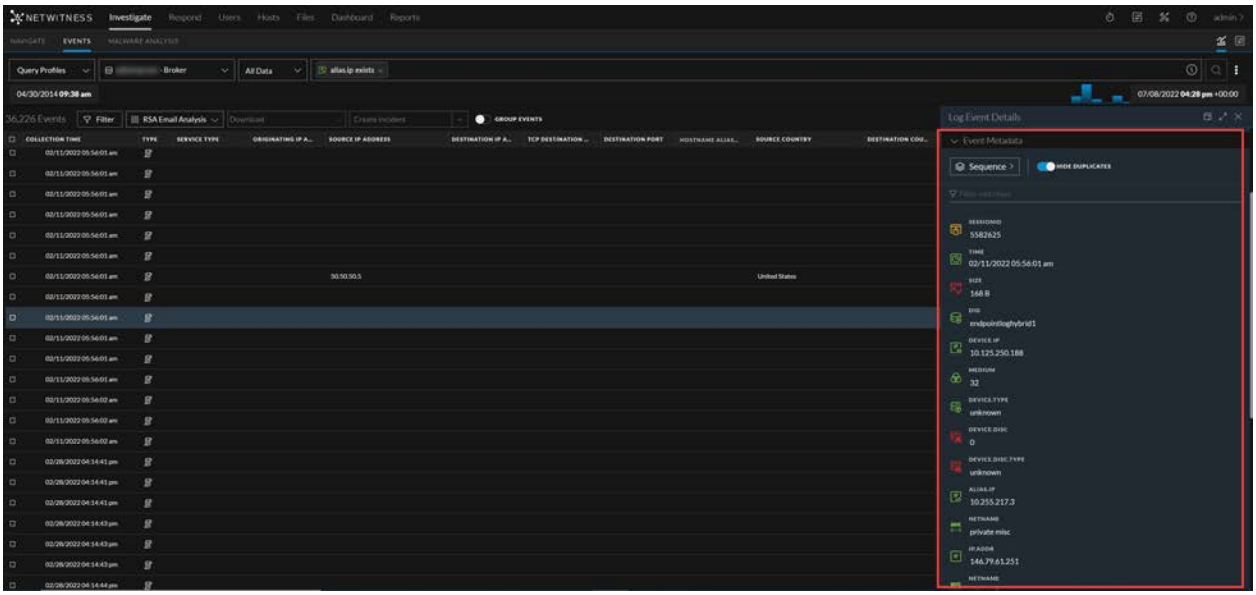
The analyst can use the toggle button ( ) to view the details related to the selected event.



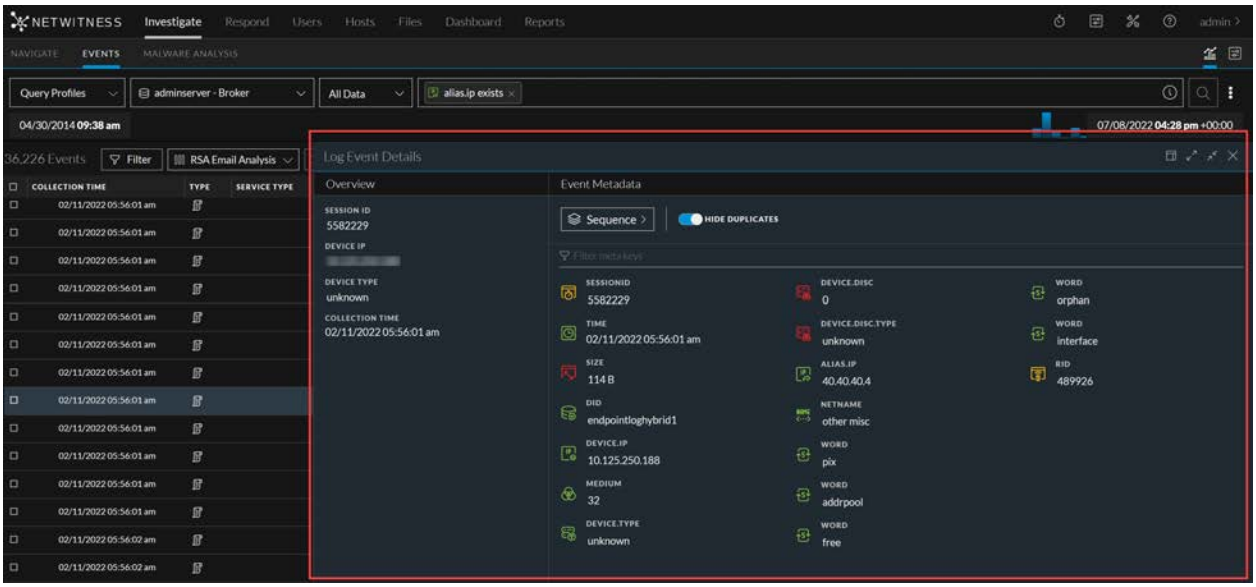
The binocular icon is enabled only if an event payload is open for a search related to a selected event.

The Overview Tab displays all the headers related to a specific event and the Event Metadata Panel displays all the metadata related to the selected event.

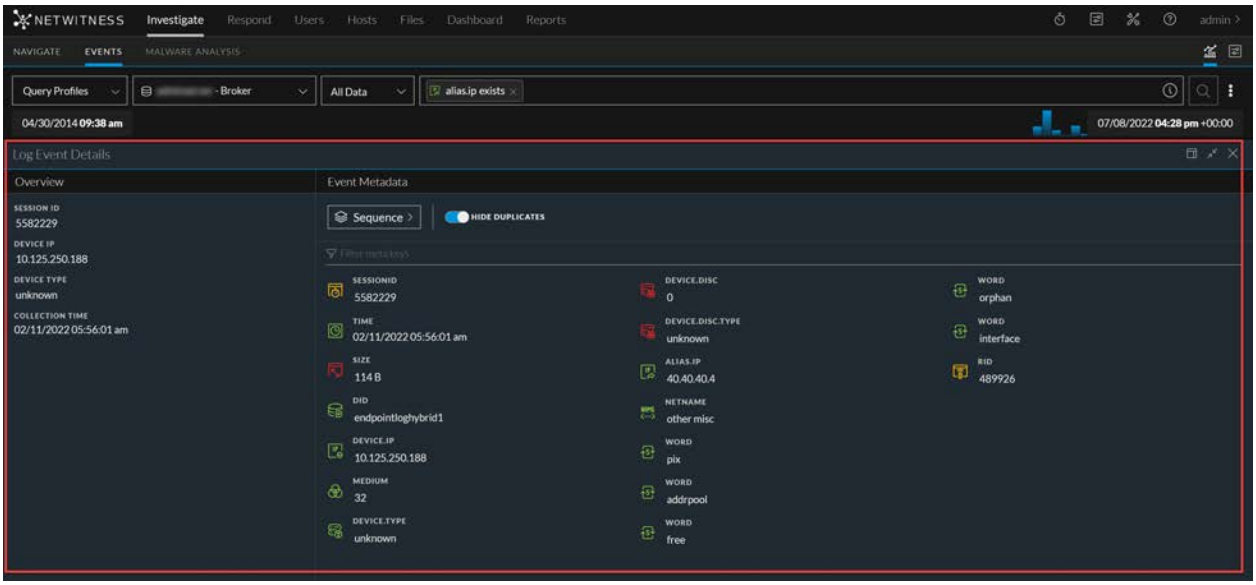
In 11.7, when you open the meta panel to view the selected event details, all the available headers will be displayed. The meta panel contains Expand buttons (). When you click on an event, the meta panel is displayed. In case, there are no additional headers to display, an error related to the header error is displayed.



The analyst can use the expand option to broaden the meta panels in three different views. When the outward facing arrow is clicked, it expands all the details displayed under the Overview and Event Metadata panel are displayed.



When clicked one more time, you can get an expanded view.



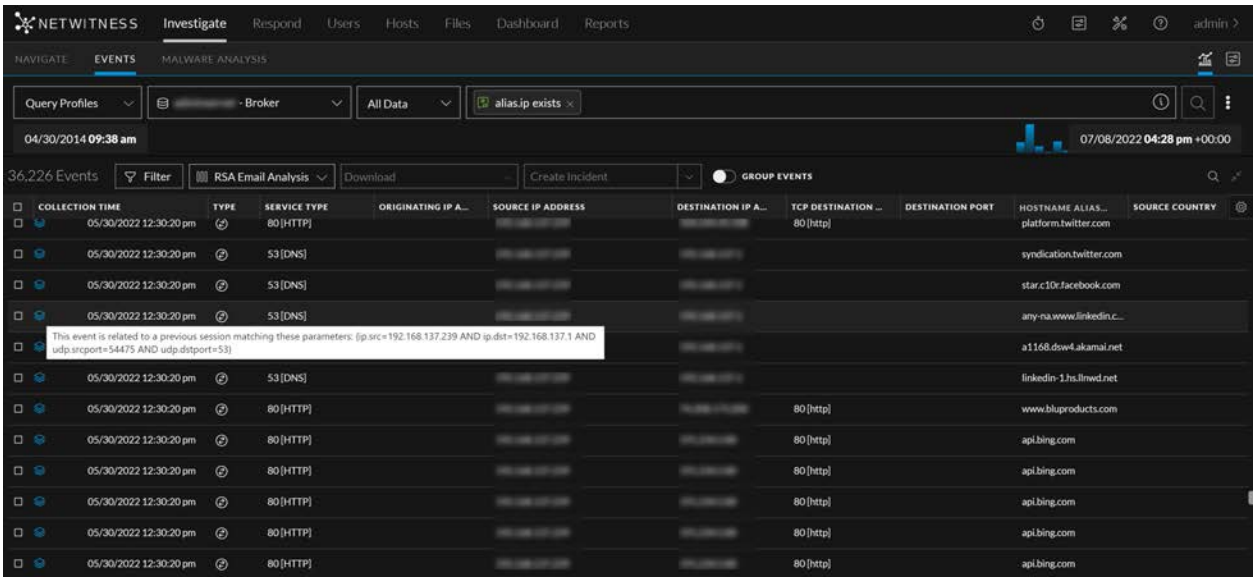
The analyst can revert the screen display by clicking the icon with the inward facing arrow , so that the window reverts to the earlier position. This helps the analyst to view the details of each event in an optimum manner.

The analysts can search for related sessions for a specific event as part of an investigation. The search for related sessions can be performed by navigating to the **Investigate** > **Events** page. You can click the icon and select either Find Related Sessions or Find Related Sessions in New Tab option from the drop-down.



When you select the **Find Related Sessions** option, all the events that are matched by selected event's query will be displayed in the current window. And if you select the **Find Related Sessions in New Tab** option the results are displayed in a new tab. The analyst can further investigate on each of the related session.

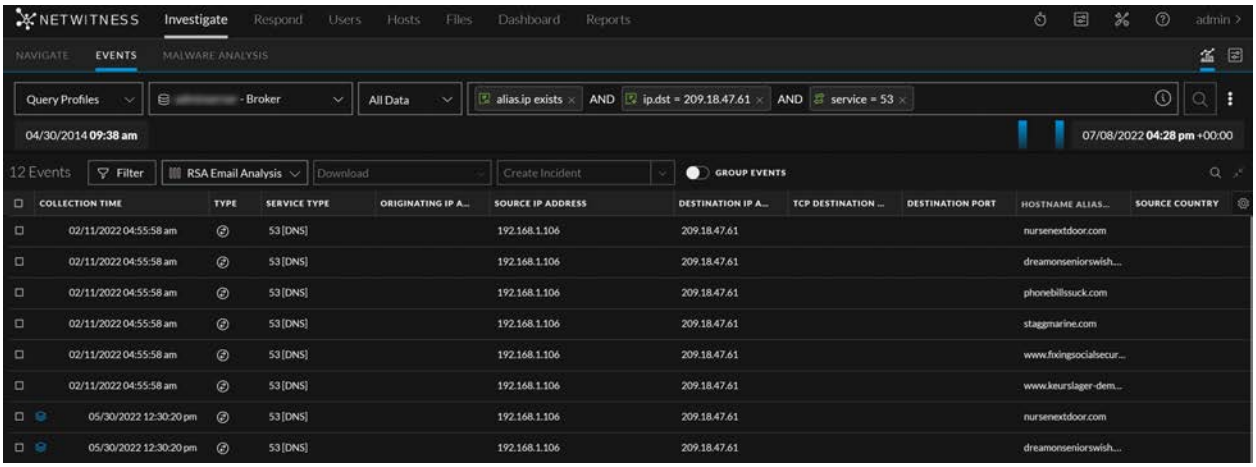
The query is based on the information displayed in the hover over text of an event. For example, in the below image the event has two split sessions where one event is split into one another session.



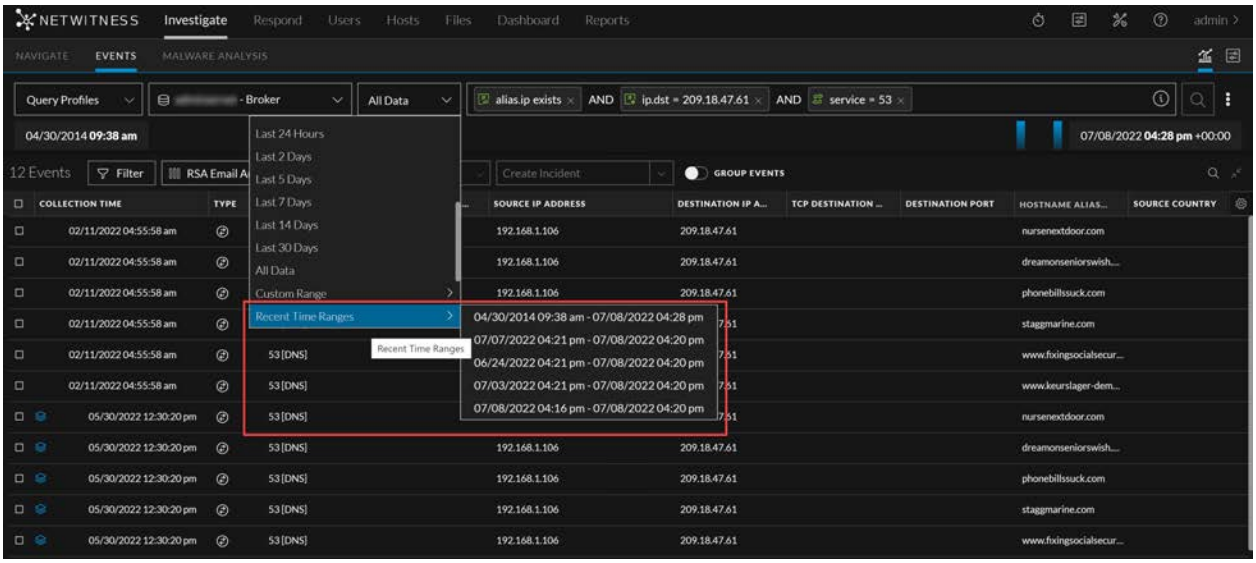
So, in this case when a search is done on these parameters, the related session for the following query is displayed.



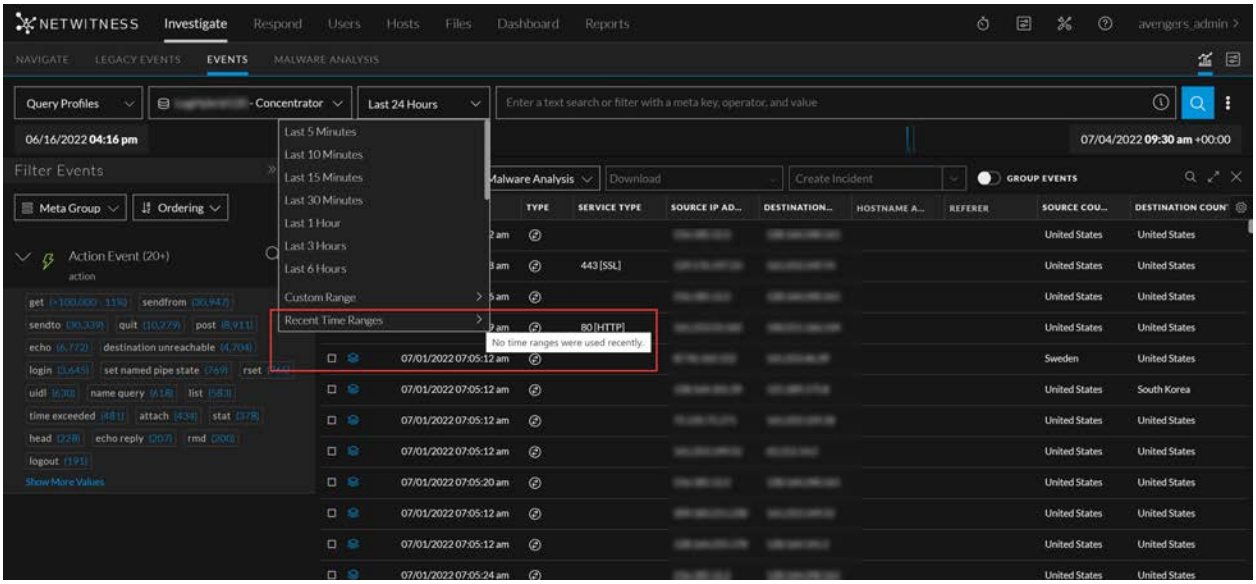
The results are displayed in the following format:



An analyst can view the last five time-ranges that were used recently, as the selection will be saved and displayed under the Recent Time Ranges section. The selection is saved separately for a user per service. For example, for Concentrator service, if you select Last 30 Days as the time-range, this will create an entry under the Recent Time Ranges section for Concentrator and when you select the same service in the next session, then the time range will be shown as below:

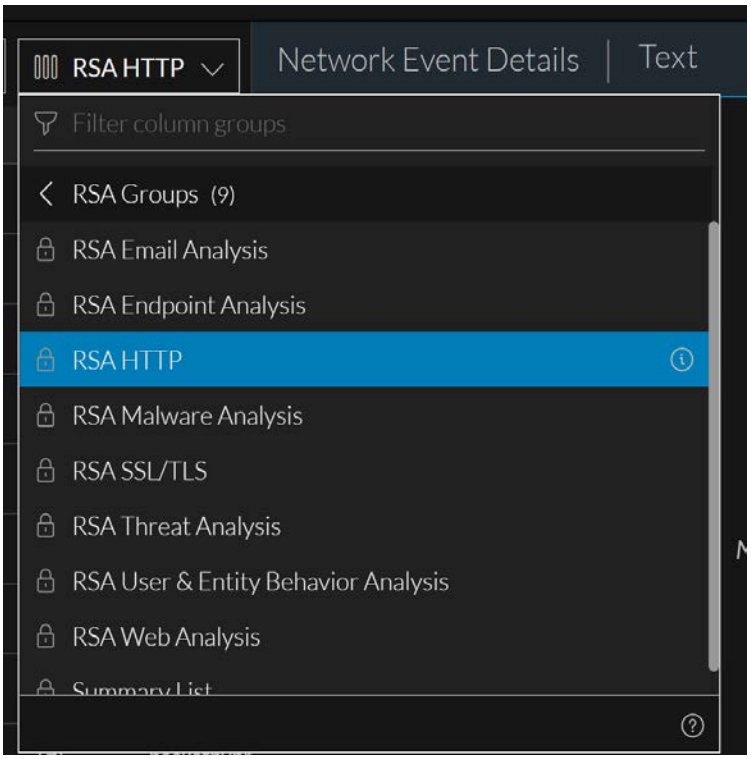


Now, if you select Concentrator service and if you have not viewed the details of the Concentrator recently, then the time range drop-down will not display any details as shown below:



Built-In Column Groups

NetWitness Platform has built-in column groups that include useful meta keys for specific types of investigation. The built-in groups cannot be edited or deleted, but you can create a copy of the group and edit the copy. The column groups are listed in alphabetical order in the Column Group menu in a way that makes built-in groups distinguishable from custom groups that you imported or created.

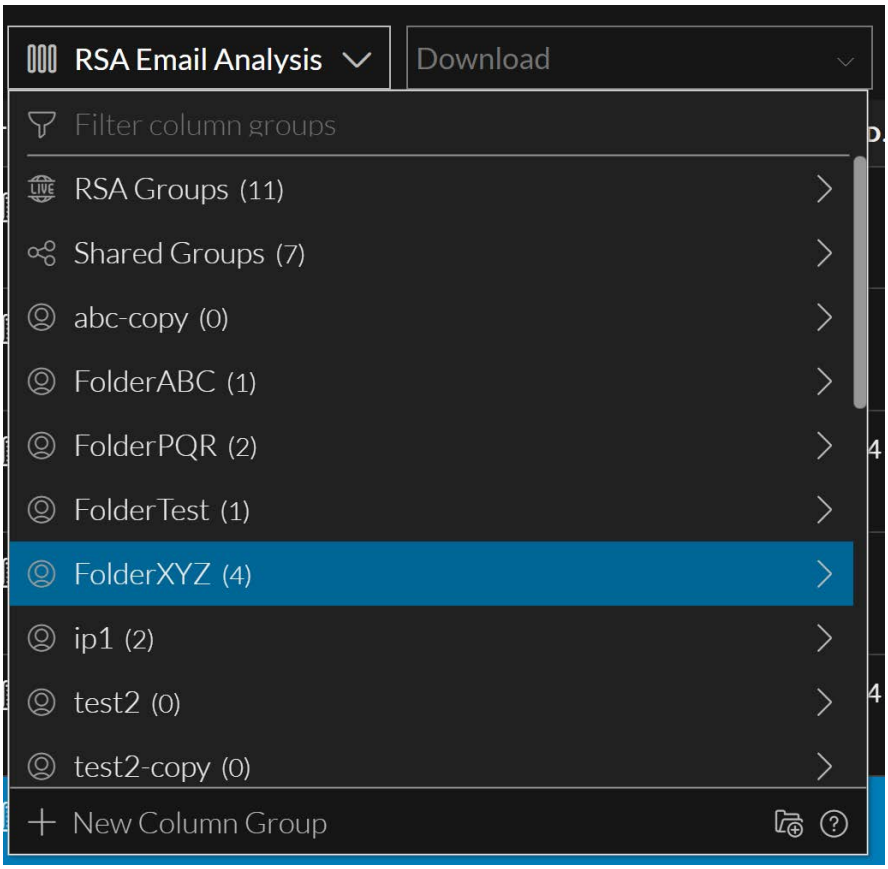


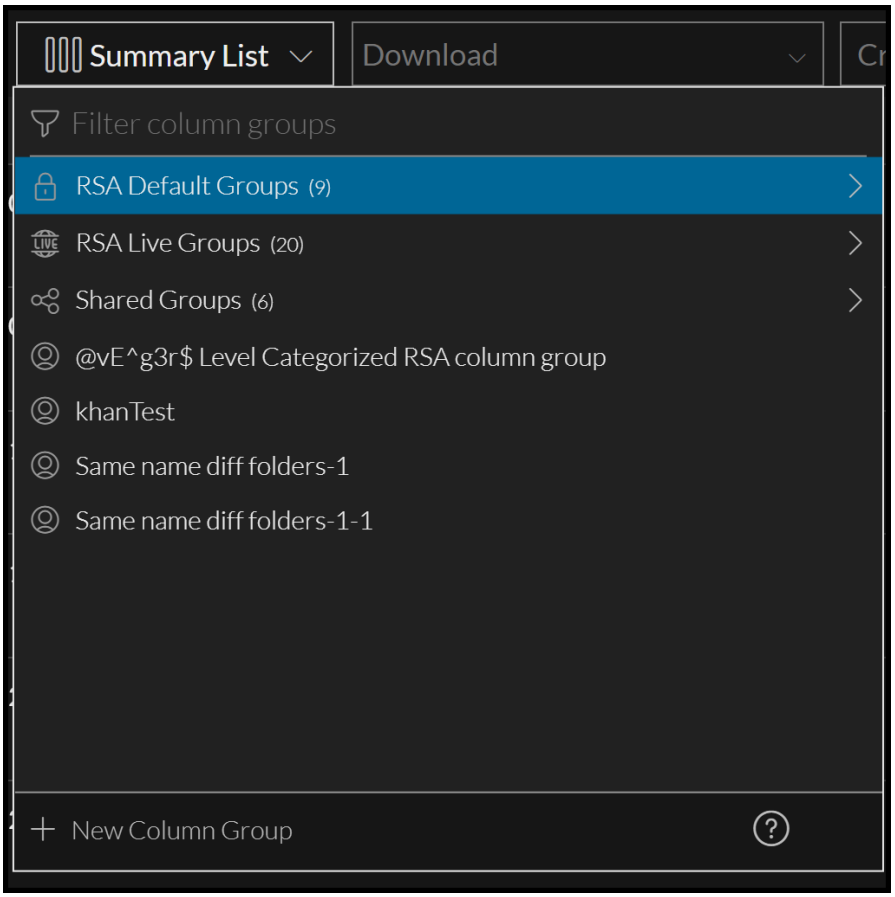
In the Legacy Events view, "RSA" precedes the name of built-in column groups. In the Events view (Version 11.4 and later), RSA precedes the name and the group is marked by the lock symbol ( ). This is an example of a selected built-in column group in the Column Groups menu. The information icon is displayed at the end of the row.

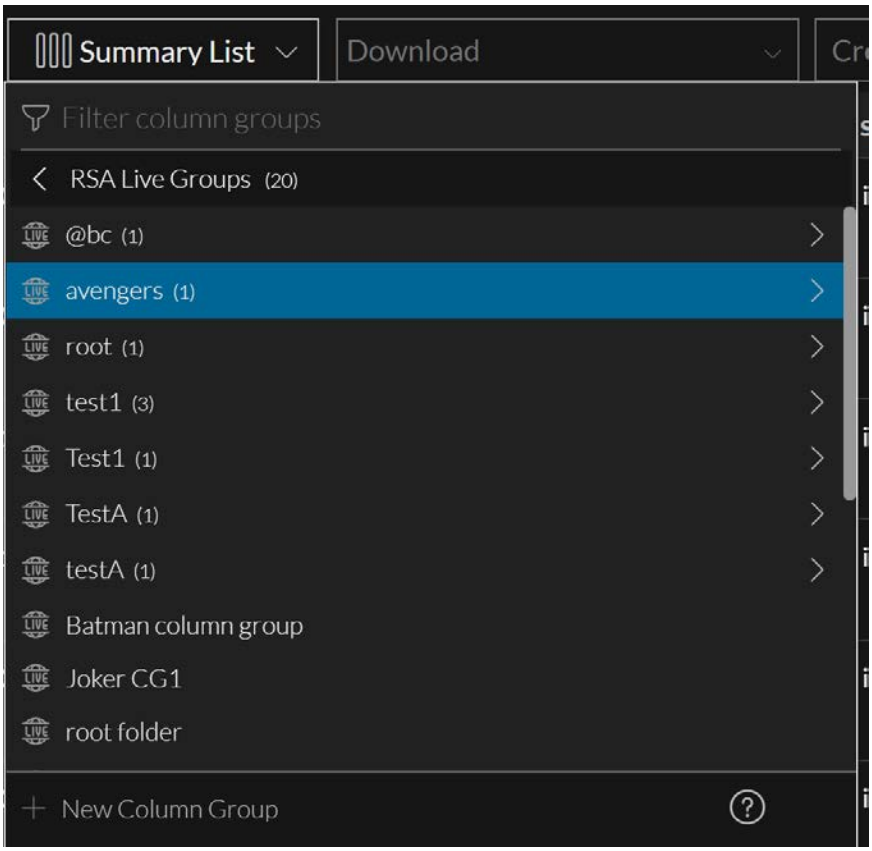


Live Column Groups

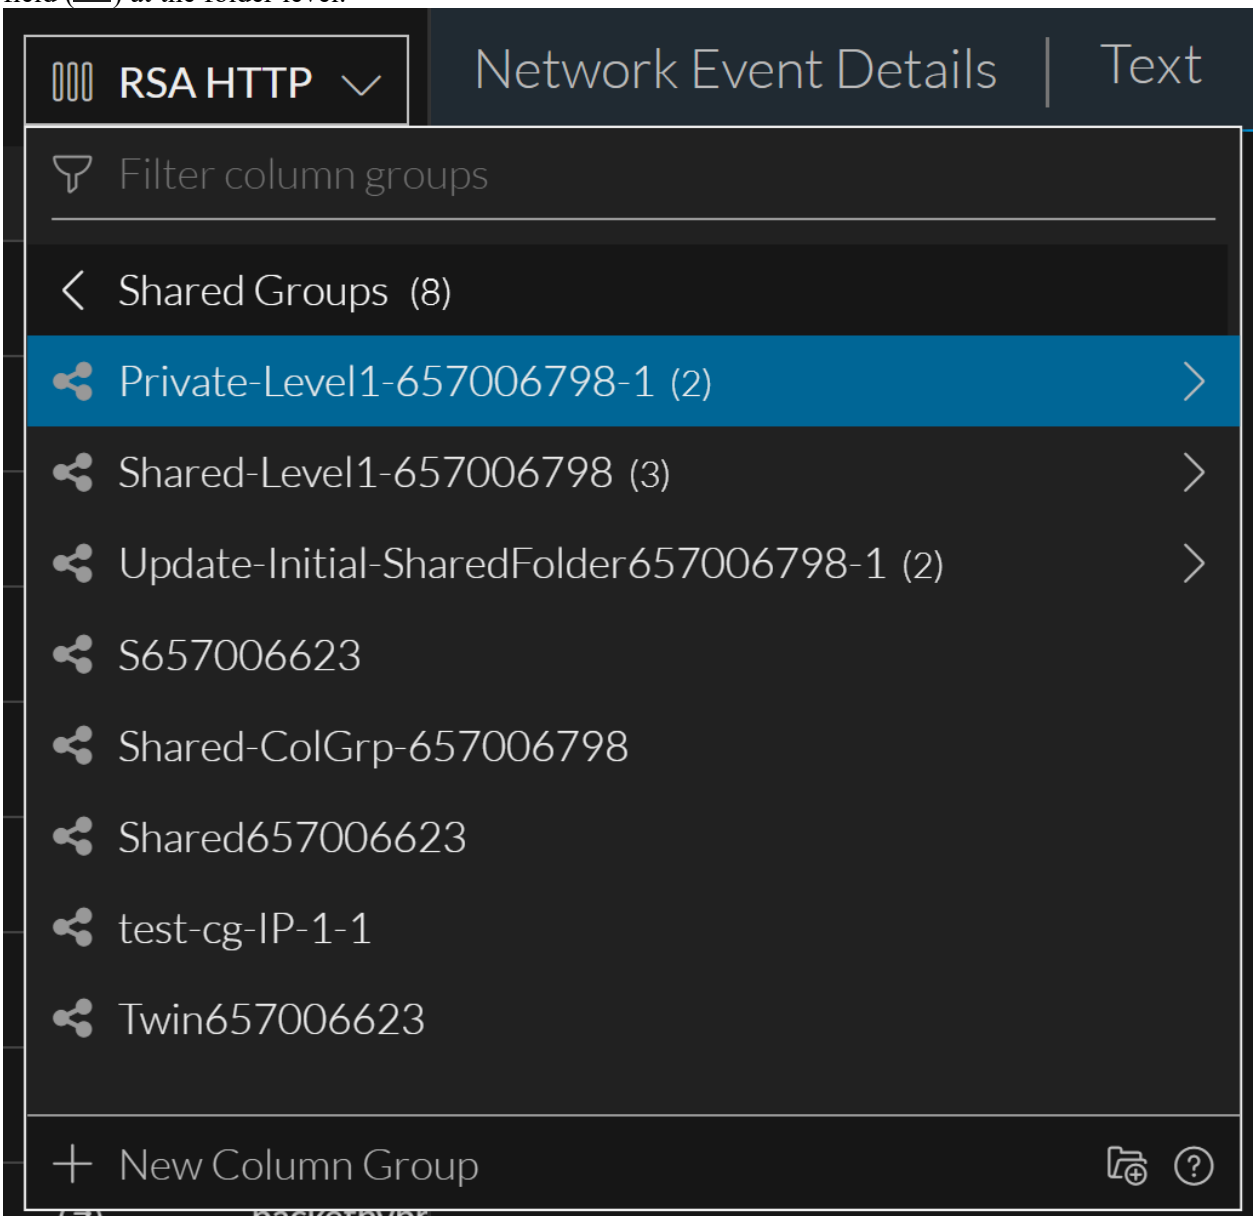
In 11.6 and later, NetWitness Platform supports deploying the investigate content from live and are marked by the live symbol () . The column groups are categorized as RSA Groups (RSA Live content and RSA OOTB Groups), and Shared Groups. The groups are displayed as non-editable folders and sub-folders except for Shared Groups that can be edited. All private content is displayed outside these groups. For example, the below image shows private content below the Shared Groups folder. The number inside () depicts the number of contents inside a folder and > symbol helps you to drill down inside the folder.


RSA Email Analysis
Download
Filter column groups
RSA Groups (11)
Shared Groups (7)
abc-copy (0)
FolderABC (1)
FolderPQR (2)
FolderTest (1)
FolderXYZ (4)
ip1 (2)
test2 (0)
test2-copy (0)
New Column Group


Summary List
Download
Filter column groups
RSA Default Groups (9)
RSA Live Groups (20)
Shared Groups (6)
@vE^g3r$ Level Categorized RSA column group
khanTest
Same name diff folders-1
Same name diff folders-1-1
New Column Group


Summary List
Download
Filter column groups
RSA Live Groups (20)
@bc (1)
avengers (1)
root (1)
test1 (3)
Test1 (1)
TestA (1)
testA (1)
Batman column group
Joker CG1
root folder
New Column Group

The column group can be copied by clicking the copy icon ( ). After copying, the copied column group is displayed under the selected location (Private folders or Shared groups). You can hover over on the cloned item to view a tool tip that displays the path from which the column group is cloned. In case you need to search for a specific column group, you can type the name of the column group in the filter field ( ) at the folder level.



These are the built-in column groups.

- **RSA Email Analysis**: Includes meta keys that are useful when investigating email-related metadata.
- **RSA Endpoint Analysis**: Includes meta keys that are useful when investigating endpoint-related metadata.
- **RSA Malware Analysis**: Includes meta keys that are useful when investigation potential malware.

- **RSA HTTP**: Includes meta keys that are useful when investigating HTTP related metadata.
- **RSA SSL/TLS**: Includes meta keys that are useful when investigating SSL/TTS analysis related metadata.
- **RSA Threat Analysis**: Includes meta keys that mark potential threats in the data set.
- **RSA User and Entity Behavior Analysis**: Includes meta keys that are useful when investigating UEBA data.
- **RSA Web Analysis**: Includes meta keys that mark anomalies in web traffic.
- **Summary List**: Includes meta keys that are useful in a general investigation. This is the default column group.

Custom Column Groups

You can create custom column groups to support scenarios that you use frequently while working in Investigate. When an administrator adds custom meta groups manually by editing the custom index file for a service, the new meta groups become available to use in column groups after the service is restarted.

Custom column groups are shared globally within your organization in Version 11.4. If you edit a shared custom column group, your changes are applied globally. If you delete a shared custom column group, the group is deleted and no longer available for all analysts. In Version 11.5 and later, you can create shared column groups as before, and can also create private column groups. When you create a group in Version 11.5, you can choose to share it or you can keep it private (default); you cannot change a shared group to private or a private group to shared.

Note: Private column groups created in the Events view are not visible or usable in the Legacy Events view.

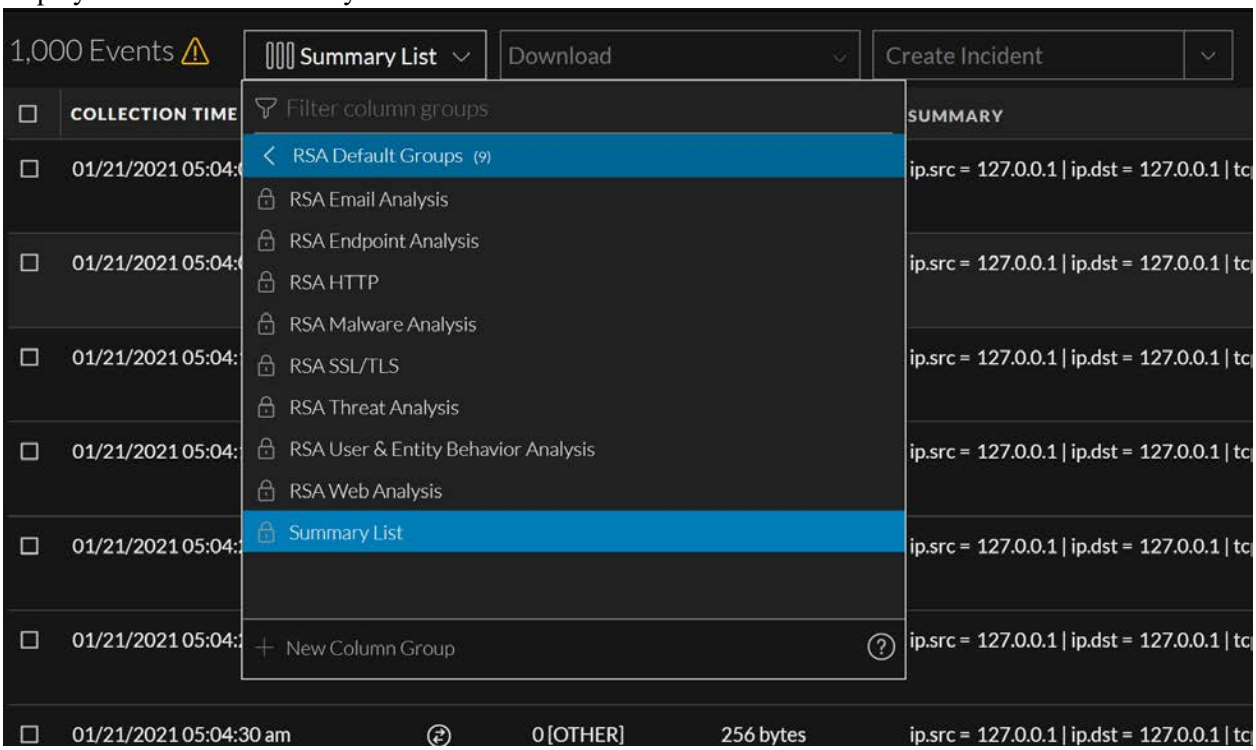
Icons identify the group type in the Column Group menu. These are examples of each type of custom column group with the edit icon displayed at the end of the row.

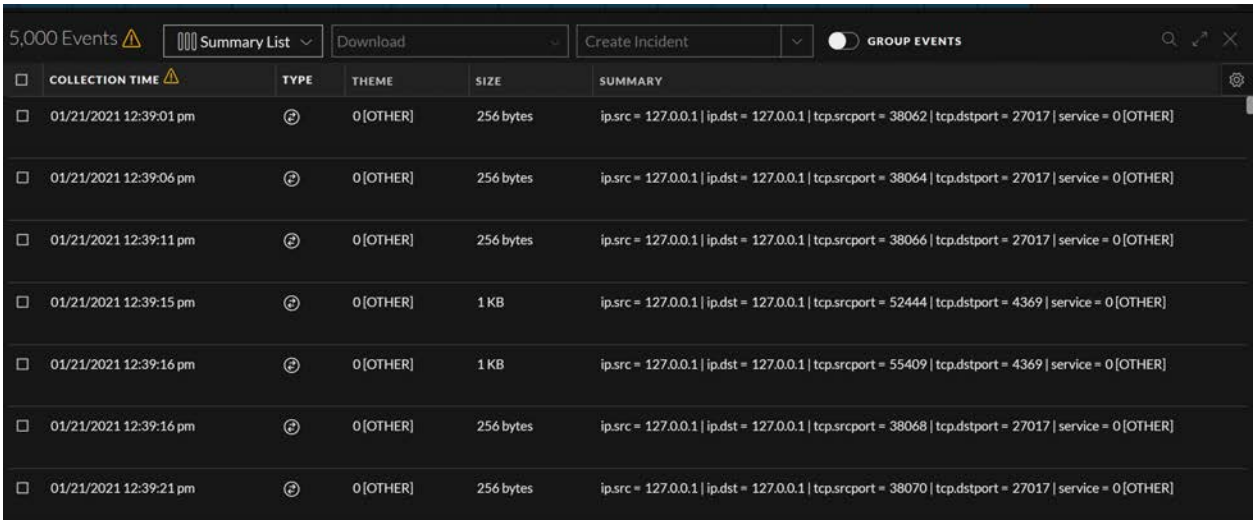


Filtering Folders

In case there are many folders, you can type the folder name and filter for a specific folder. The filtering is applicable to the current level folders and will not display folders available within a sub-folder. To search content within a sub-folder you need to navigate to the specific folder and filter.

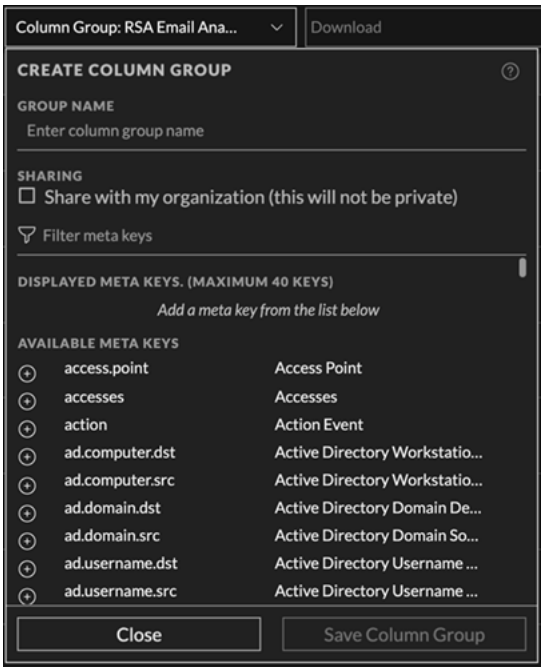
Also, when you select a specific folder, the content of the selected folder is displayed and the filter field becomes empty and when you navigate back the last selected folder is displayed. In the following example, the folder selected is RSA Groups with the its content and the column group drop-down displays the filtered Summary List folder.

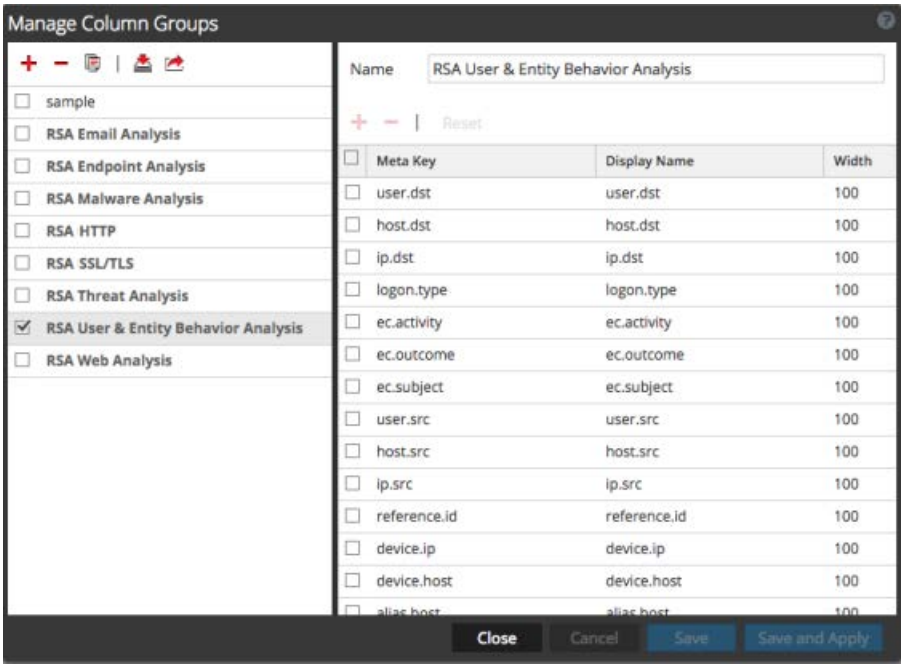




Dialogs for Managing Column Groups

While the functionality of column groups is similar in the Legacy Events view and the Events view, the user interface and some of the procedures are different. The following figures illustrate the (Events view) Create Column Group dialog and the (Legacy Events view) Manage Column Groups dialog. The Version 11.5 and later dialog includes a Sharing option.





Using options in the Create Column Group dialog and the Column Group Details dialog, you can:

- See the details of a column group.
- Create, edit, and delete custom column groups.

Using options In the Manage Column Groups dialog, you can do all of above and these additional functions:

- Clone and edit the clone of a built-in or custom column group.
- Import and export a column group.

The rest of this topic provides instructions for working with column groups in the Version 11.4 and later Events view, the 11.3 and earlier Event Analysis view, and the Legacy Events view.

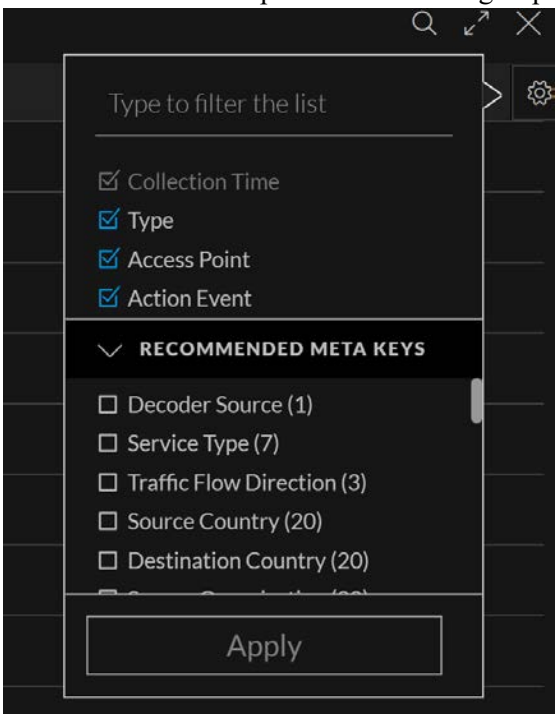
Work with Columns and Column Groups in the 11.4 and Later Events View

After the upgrade to Version 11.4, all of the existing column groups -- both built-in and custom -- are available for management in the Events view. Unless noted, the procedures in this section are for the Events view.

Manually Select Columns to Display and Adjust Column Order and Width

Note: The Column Selector was also available in the 11.3 Event Analysis view. If a column group includes a column for a meta key that your administrator has blacklisted (hidden), the data for that column cannot be displayed. The column is not available in the Column Selector and is not displayed in the Events panel.

1. With the Events list open and a column group applied, click to display the column selector.



2. Select the meta keys or enter the name of a meta key that you want to display in additional columns.
3. Deselect the meta keys that you do not want to display in a column.
   The data is redisplayed using the selected columns.

4. To change the width of the columns in the events list, hover the cursor over the column title and drag the column divider to the right or the left.
5. To rearrange the order of the columns across the top of the events list, hover the cursor over the column title and drag the column to the right or the left.
   The changes that you make in the events list are in effect during the current session and are not retained as part of the column group. The next time the column group is applied, the original composition and order of columns is applied.

Select a Column for Sorting Events in the Events Panel (Version 11.4)

Note: You can sort events in the Events panel after results have finished loading if all connected services are updated to 11.4. or later. Sorting by column is disabled when any connected service is running an earlier version of NetWitness Platform. Version 11.4.1 has more visible sorting toggles in the column heads and the ability to view results without sorting, but otherwise it functions the same as in Version 11.4.

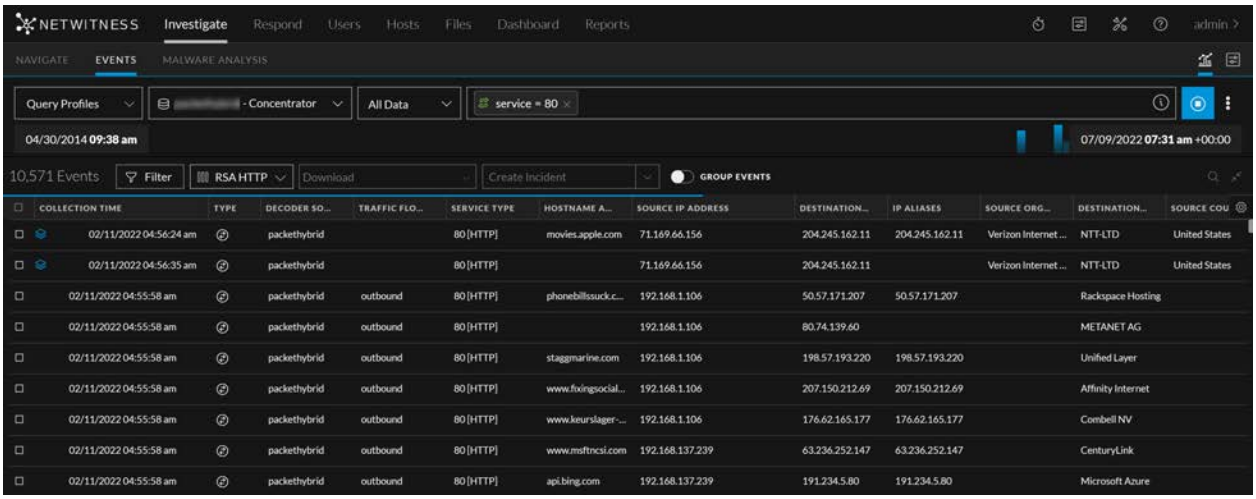
You can change the order of the events list in the Events panel based on the value for a meta key in the event. Each column title represents a meta key, and the column is populated by the values found for the meta key in the displayed events. In Version 11.4, the events in the Events panel are sorted using the method selected in the Event Preferences dialog: Ascending or Descending. If no sort method is selected, the default order is ascending (see Configure the Events View). In Version 11.4.1, the events in the Events panel are sorted only when the sort preference in the Event Preferences dialog is selected and is either Ascending or Descending. The events are not sorted if you do not have a sort preference selected under Events Preferences or if you selected Unsorted.

Sortability of a column is based on the definition of the meta key in the Broker and Concentrator index files. Columns for meta keys that are indexed by value are sortable. If the meta key is not indexed, is indexed by meta key, or has multiple values in the same event, it is not sortable.

- These are some examples of keys that are indexed by value and sortable: `time`, `eth.type`, `city.src`, `ip.src`, `ipv6.dst`, and `ipv6.src`.
- Meta entities are not sortable. For example, the meta entity `ipv6.all` is not sortable because it includes `ipv6.dst` and `ipv6.src`, and a single event has both `ipv6.dst` and `ipv6.src`.
- These are some examples of multiple value keys, which cannot be sorted: `filename`, `filetype`, and `attachment`. A single event can have more than one file and therefore more than one value for `filename`, `filetype`, and `attachment`.
- These are some examples of meta keys that cannot be sorted because they are not indexed or not indexed at the values level: `password`, `query`, and `size`.

Sorting by Column (Version 11.4.1 and Later)

The initial view of the Events list with the sorting preference set to Unsorted and no column sorting has an event count in the title, with no indication of a sorting method applied to a column. If the event sorting preference is set to Ascending, the count label is "Oldest 1,000 Events." If the event sorting preference is set to Descending, the count label is "Newest 1,000 Events." In the figure below, the Ascending method is in effect, more than 2001 events matched the query, and only the oldest 2001 are displayed. Clicking the amber warning triangle displays an explanation. Refer to Configure the Events View for more information about the sorting preference.



When you move the mouse over a column title, sortable columns have a pair of arrows after the column title, one pointing up for ascending and one pointing down for descending ( ). You can choose one sort column and the direction of the sort. A blue up arrow ( ) indicates that ascending sort order is in effect; which means the earliest events or the lowest numbers, or the text strings beginning with an 'A' appear first. A blue down arrow ( ) indicates that descending sort order is in effect; which means the latest events or the highest numbers, or the text strings beginning with a 'Z' appear first.

- When a column has a blue arrow, you can click the white arrow to change the sort order. When you change the sort order, a blue progress bar is displayed in the Events list title bar to show progress. As sorting begins, there is a short segment on the left side of the window; as sorting progresses the blue color extends to the right across the entire title bar. The directional arrow does not change until the events are re-sorted in the chosen sort order.
- To change the column to unsorted, you can click the blue arrow. Both arrows are white now to show that the column is unsorted. This figure shows the Type column sorted in ascending order.

- If a column is not sortable, no arrow is displayed when you hover the mouse over the column title. Instead a tooltip explains why it is not sortable.

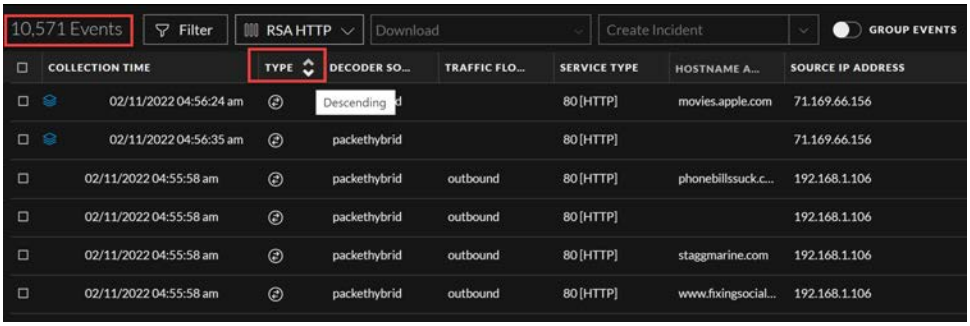
Sorting on a column is done on the client side without re-executing the query if the number of displayed results is less than the events limit set by the administrator. If there are more results that are not displayed because the number of results exceeded the events limit, a new query is submitted with the new sort order, and the same service, time range, and filters. The current results are removed, a spinner indicates progress, the Cancel button becomes available, the reconstruction closes, and progress is visible in the Query console.

Note: The re-sorting of events takes place in the browser when the number of results of the original query is less than the event display threshold.

To change the sort order or the sort column:

1. Move the mouse over the column titles to find a sortable column.
   If a column is not sortable, a tooltip that explains the reason is displayed.

2. To sort the list based on a column, move the mouse over a sortable column and click one of the arrows ( ).
   The arrow turns blue and the events are reloaded in the selected order. If both arrows are white, the column is not being used to sort the events list. If one arrow is blue, the column is being used to sort the events list, and the sort order (Asc or Desc) is appended to the events count in the title bar. This figure shows a column sorted in ascending order. When a column is descending order, (Desc) is appended to the event count.



   a. Click a white arrow to sort the events list in that order.

   b. Click a blue arrow to return to unsorted order.

Sorting by Column (Version 11.4)

When you move the mouse over a column title, sortable columns have an up or down arrow (or) after the column title. You can choose one sort column and the direction of the sort. An up arrow indicates that Ascending sort order is in effect; which means the earliest events or the lowest numbers, or the text strings beginning with an 'A' appear first. A down arrow indicates that Descending sort order is in effect; which means the latest events or the highest numbers, or the text strings beginning with a 'Z' appear first. When you select a sort column, it is sorted in descending order by default, with events having a null value for the meta key first.

- A column that is being used to sort the events list has a bright white arrow indicating the direction that you can choose for sorting: click to change to Ascending or to change to Descending order. When you click to change to Ascending sort order, the directional arrow does not change until the events are re-sorted in ascending order. The same behavior applies when you click the to change to Descending order.
- If a sortable column is not being used to sort the events list, the arrow is dimmed. If a column is not sortable, no arrow is displayed when you hover the mouse over the column title. Instead a tooltip explains why it is not sortable.
- If you click the arrow on a different column, the column is sorted in the same order as the previously active sort column. You can select a different sort order if desired.

Sorting on a column is done on the client side without re-executing the query if the number of displayed results is less than the events limit set by the administrator. If there are more results that are not displayed because the number of results exceeded the events limit, a new query is submitted with the new sort order, and the same service, time range, and filters. The current results are removed, a spinner indicates progress, the Cancel button becomes available, the reconstruction closes, and progress is visible in the Query console.

Note: The re-sorting of events takes place in the browser when the number of results of the original query is less than the event display threshold. If some of those events have the exact same time, they will not change order as you might expect when you reverse the sort order.

To change the sort order or the sort column:

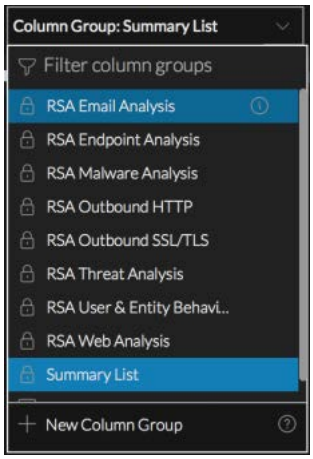
1. Move the mouse over the column titles to find a sortable column.
   If a column is not sortable, a tooltip that explains the reason is displayed.
2. To sort the list based on a column:
   a. Move the mouse over a sortable column and click the arrow ( or ).
   The events are sorted in the correct sort order. If you hover over the column title, you can see that the arrow is no longer dimmed. A column that is being used to sort the events list has a bright white arrow that you can click to change the sorting direction.
   b. To change the sort order, click  to change to Ascending or  to change to Descending order.
   The direction of the arrow changes and the events are reloaded in the selected order.

View the Meta Keys Included in a Column Group

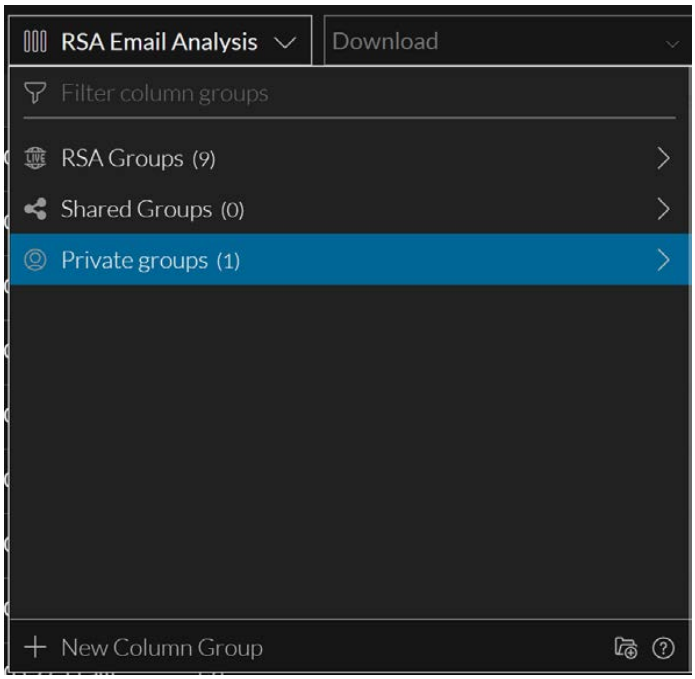
To view details of a column group:

1. Go to **Investigate > Events** and click  to load events.
   The events for the default service and the default time range are loaded in the Events panel. The Summary List column group or the column group from your last session is applied to the list.
2. To display the Column Groups menu, click the Column Groups menu title. The Column Group menu title includes the title of the currently selected column group. If this is your first visit after logging in, the Summary List group is selected; any subsequent visits use the column group selected in the previous session. When opened, the menu displays a list a list of built-in column groups (RSA), shared custom column groups, and your private custom column groups. The figure shows the Version 11.6 menu initially when Summary List is selected by default and all types of column groups are

visible: Private, Shared, and RSA.

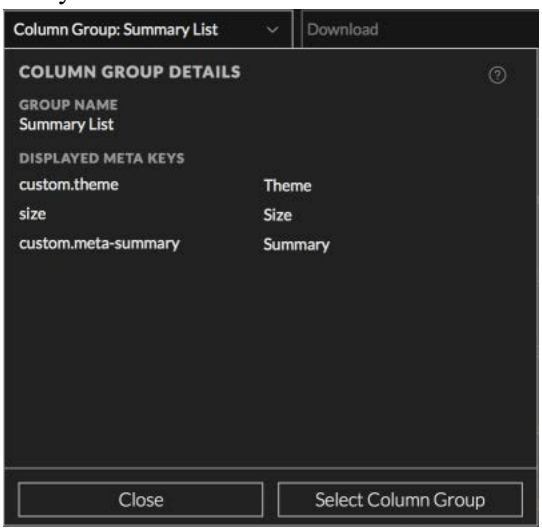


3. (Optional) To control the types of column groups that are visible in the list, use any combination of the visibility options (blue = selected, black = not selected):
   Private = display private groups that only you can manage
   Shared = display shared groups that anyone in your organization can manage
   RSA = display built-in groups that only RSA can manage
   The visibility options work together with the Filter Column Groups field. If the visibility option is hiding built-in groups (which include "RSA" in the group name) and you search for a name that contains "RSA," the list is empty. The figure below shows private and shared visibility options selected.



4. Hover over the **Summary List** group and click the information icon ( ) to see which columns are included in the group.
   This figure shows the columns for the Summary List. The Collection Time and Type column are

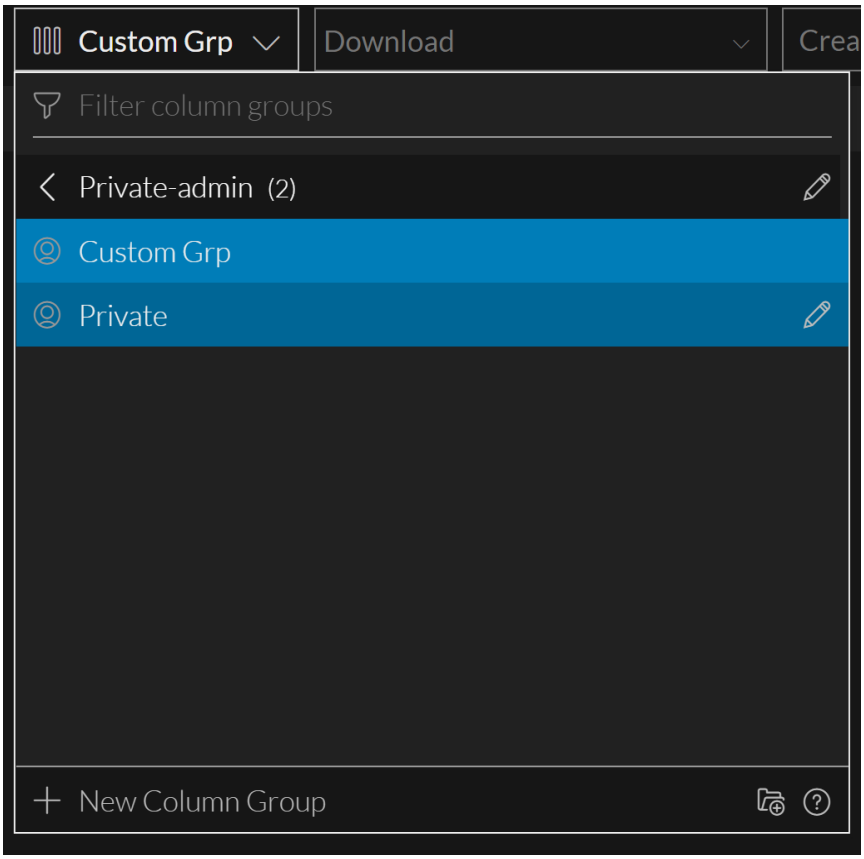
always the first two columns in the Events list, but are not listed in the Column Group Details dialog.



5. Do one of the following.
   a. To close the dialog, click **Close**.
   b. If you want to apply the column group, click **Select Column Group**.
      The dialog closes and the Events list is updated to reflect the selected column group.

Select a Column Group

1. With the Events panel open in the 11.4 or later Events view, click the **Column Group** menu title.
   The menu drops down to display a list of column groups with a filtering option and a New Column Group option. The list is sorted alphabetically and the selected column group is displayed in the menu label. The first option in the list is highlighted. The selected column group has a slightly different background color than the highlighted column group.
   The following figure shows the menu after RSA Endpoint Analysis was highlighted, but RSA Email Analysis is still selected.

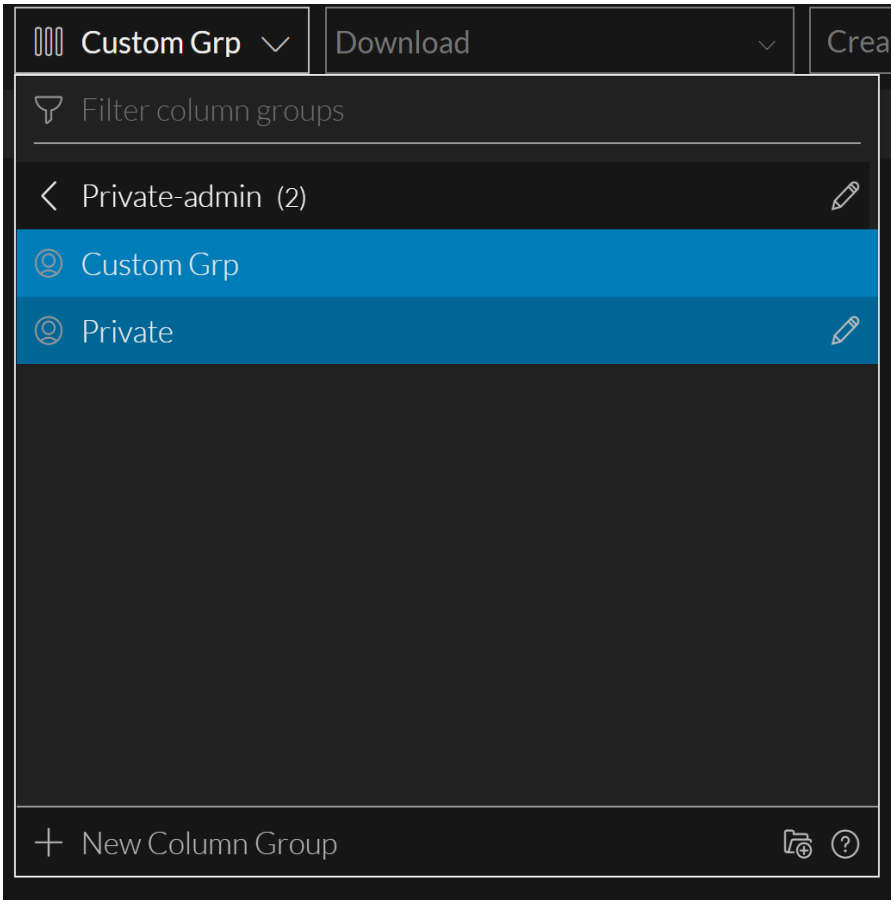


2. Do one of the following:
   a. If the highlighted group is the one you want to apply, press **ENTER**.
   b. (Version 11.5 and later) If you want to see only certain types of groups, use the visibility options (**Private**, **Shared**, and **RSA**) to hide one or two group types.
   c. Begin typing text in the **Filter column groups** field to search for a column group name. As you type, the list is filtered to show only the column group names that contain that string.
   When you see the group that you want to apply, click it or use the down or up arrow to highlight it, then press **ENTER**.
   The Events list is refreshed to include only columns in the selected column group, and the menu title includes the selected group name. Your selection persists when you navigate away from the Events view. The order of the columns in the Events list reflects the order of the meta keys in the column group. A column group may contain more columns that are only visible when you scroll to the right. For optimal viewing, the first 15 columns are displayed by default when you select a column group.

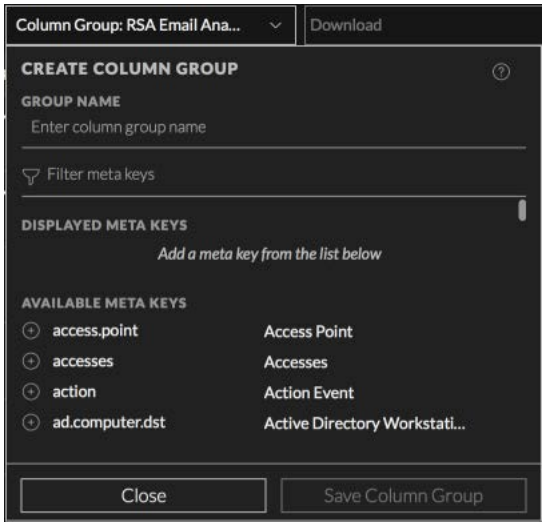
Note: If a meta key in a column group is not part of the selected service, it does not appear in the Filter Events panel or in the Events panel.

Create a Custom Column Group

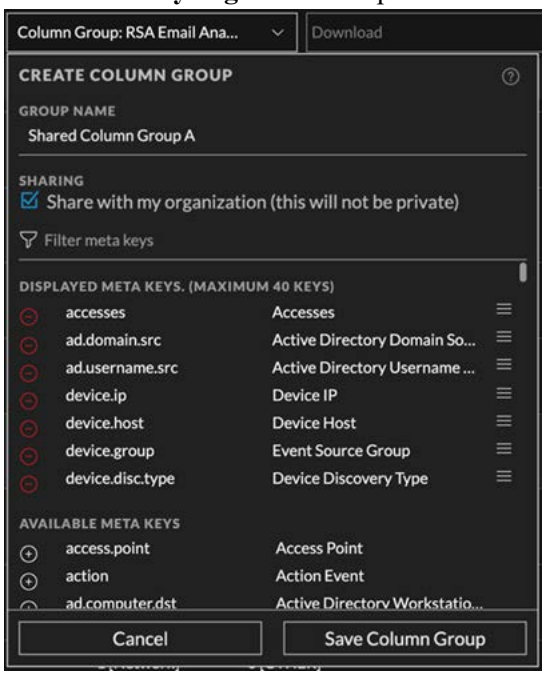
1. Go to **Investigate > Events** and submit a query to load data in the Events panel.
2. In the Events panel toolbar, click the **Column Group** menu title.
   The menu drops down to display a list of column groups with the Visibility Options and Filter Column Groups field at the top and the **+ New Column Group** option at the bottom.



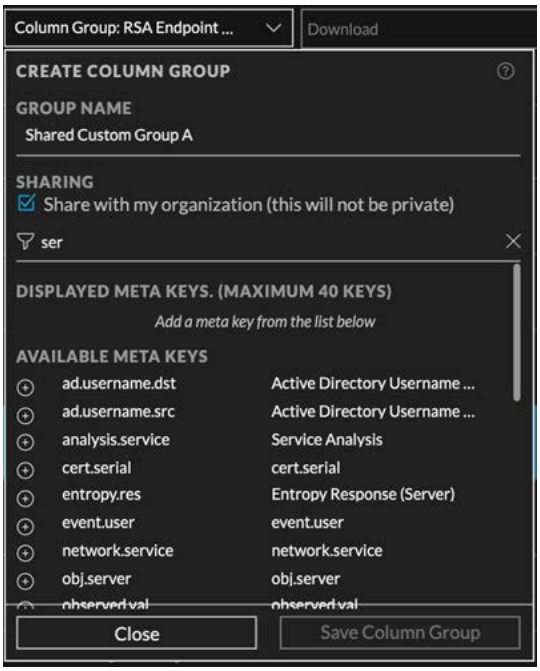
3. Select + **New Column Group**.
   The Create Column Group dialog is displayed. Version 11.5 includes the Sharing option.



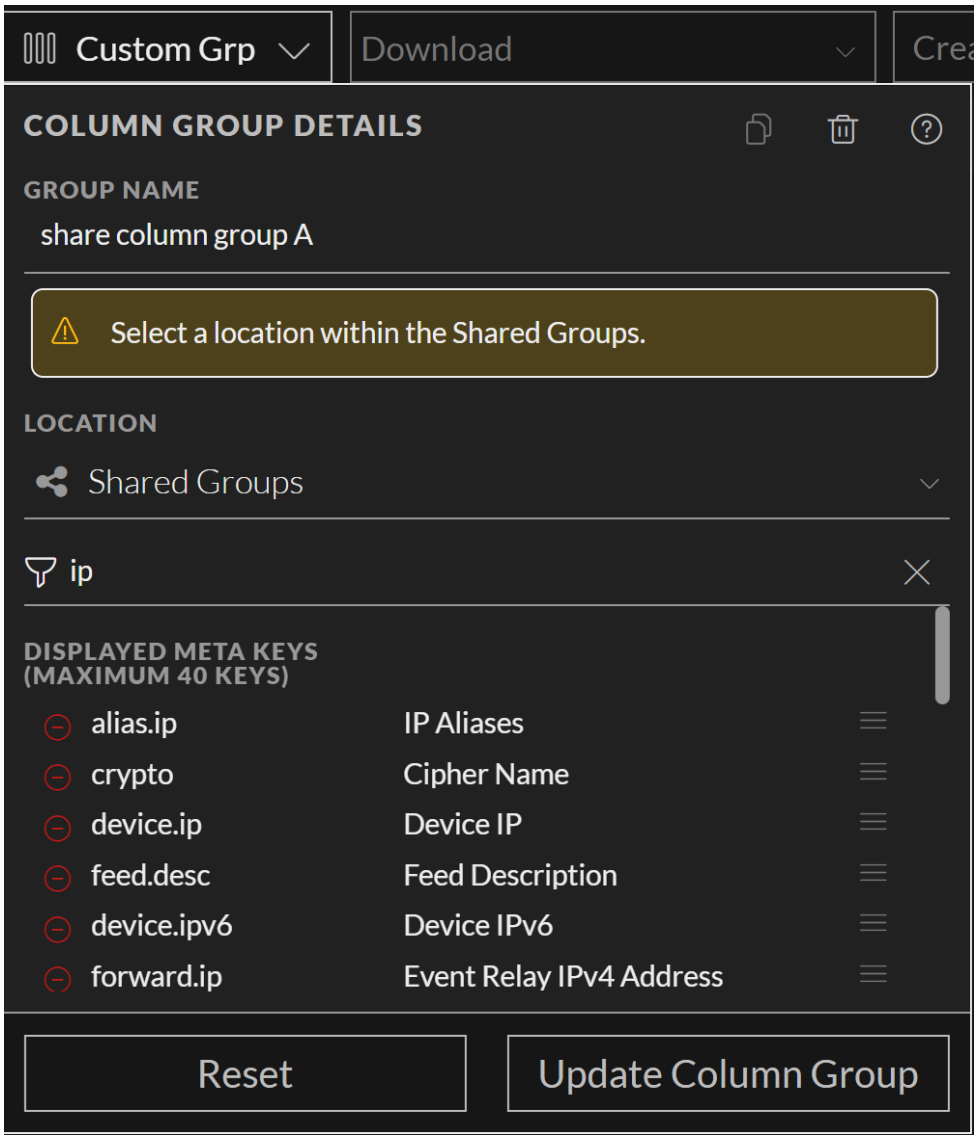
4. In the **Group Name** field, type a unique name (maximum length of 256 characters) for the new column group, for example, **Custom Column Group A**.
5. (Version 11.5 and later) If you want to share the new column group with your organization, set the **Share with my organization** option.



6. To add a meta key to the column group, select and add each meta key as follows:
   a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list.



b. When you see the meta key that you want to add, click the add icon  that precedes the meta key name.
The meta key is added to the end of the Displayed Meta Keys list. (This list is also filtered using the text you typed.) The maximum number of meta keys in a column group is 40. If you attempt to add another meta key when 40 are already included in the Displayed Meta Keys list, a message advises you that the group has the maximum number of meta keys.



7. (Optional) To find and remove a meta key from the column group, type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Displayed Meta Keys** list. When you see the column that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
8. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
9. Do one of the following:
   a. To close the dialog without creating the custom column group, click **Cancel**.

b. To create the group, click **Save Column Group**.
The new column group is saved and becomes available for all analysts. The buttons change to Done and Select Column Group.

10. Do one of the following:

a. To close the dialog, click **Done**.

b. To close the dialog and select the new column group, click **Select Column Group**.
The new group is added to the Column Groups menu (in alphabetical order), and if you clicked Select Column Group, the Events list is updated to show the columns in the new column group.

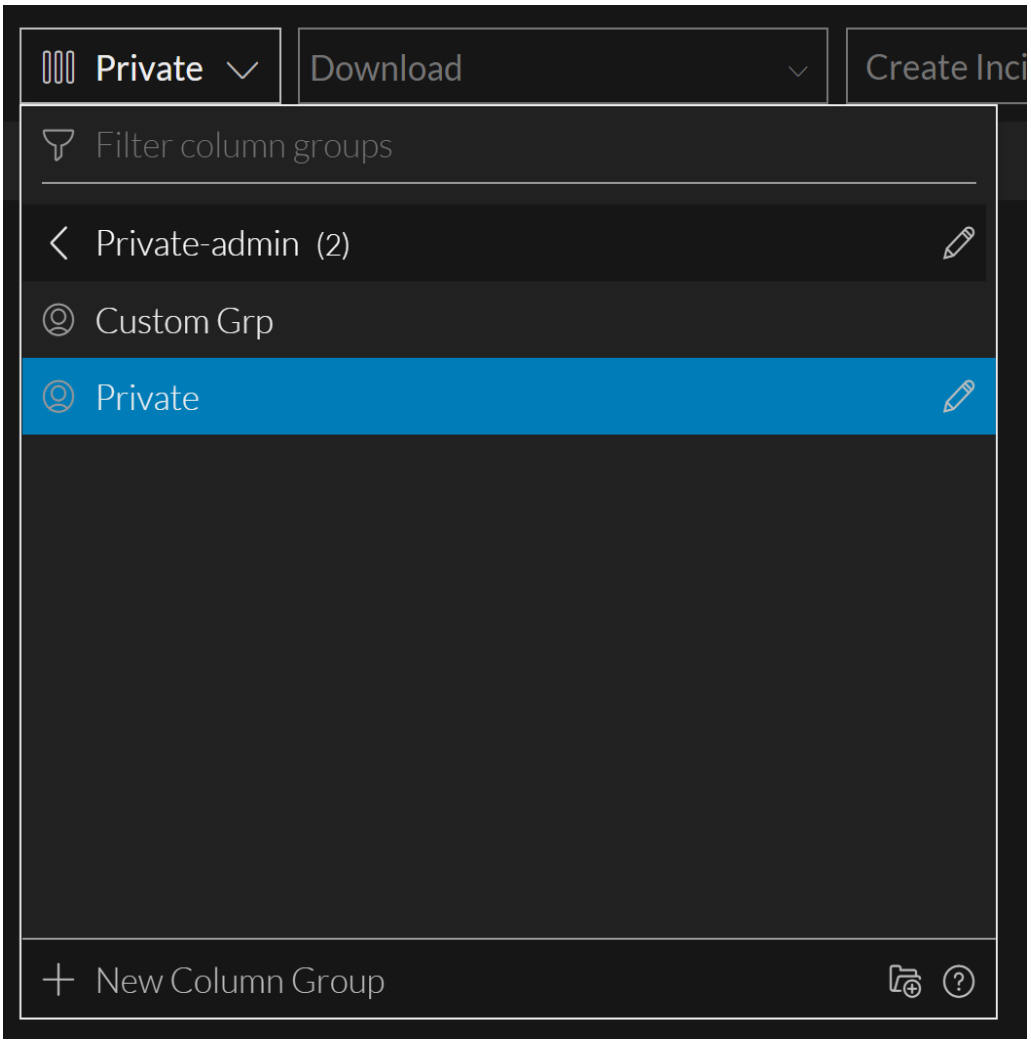
Delete a Custom Column Group

You can delete any custom column group that is not currently applied in the Events list and not part of a query profile. The built-in column groups are read only, and cannot be deleted. In Version 11.5 and later, a confirmation message allows you to confirm or cancel the deletion. When you delete a custom column group, it is removed from the Column Group menu.

Caution: When you delete a custom column group (Version 11.4) or a shared column group (Version 11.5), the effect is global and the group is no longer available to any analyst.

To delete a custom column group:

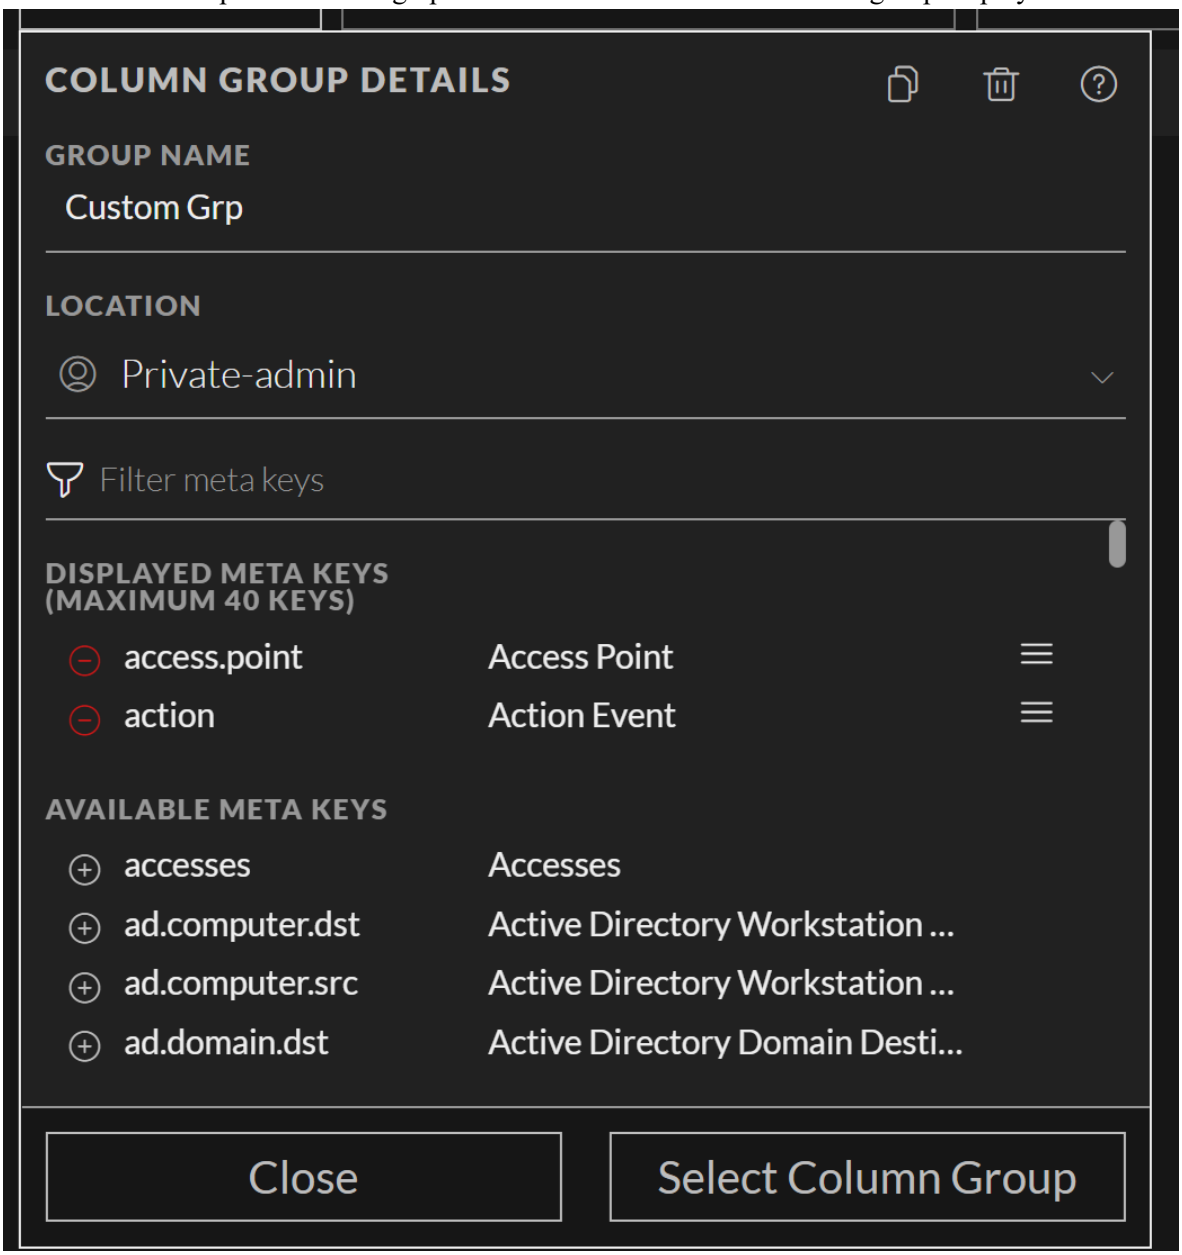
1. Go to **Investigate > Events** and click  to load events.
The events for the default service and the default time range are loaded in the Events panel. The Summary List column group or the column group from your last session is applied to the list. This figure shows the initial view with the Summary List column group selected. The label on the Column Group menu includes the name of the selected column group.



2. To delete a column group, highlight a custom column group as shown in the following figure and click the edit icon (✎) to the right of the name.



3. The Column Group Details dialog opens with the details for the selected group displayed.

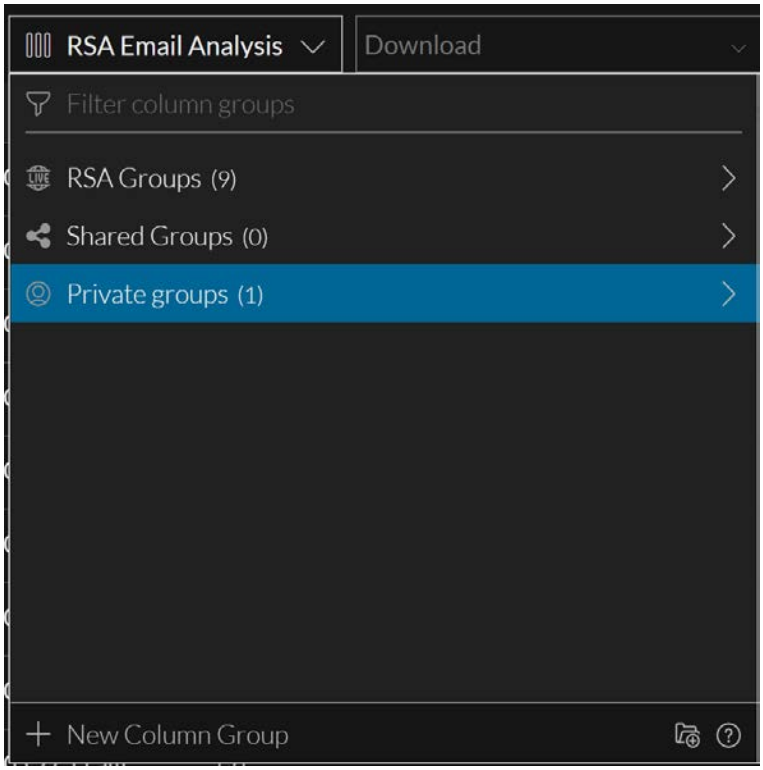


4. Click the delete group icon ( ).
   If the column group is currently in effect, the following message is displayed: `This column group cannot be deleted because it is currently active.`
   In Version 11.5, a confirmation message gives you the opportunity to confirm or cancel the deletion. Click **Cancel** or **Delete Column Group**.
   In Version 11.4, if the column group is not in effect and is not a built-in column group, there is no request for confirmation before the column is deleted.
   The group is deleted and removed from the Column Groups menu. The column group no longer appears anywhere for any analyst working in Investigate.

Edit a Custom Column Group

You can create a shared or private copy of any column group that is not open for editing. After you create the copy, you can edit the new group in the usual way.

1. Go to **Investigate> Events** and submit a query to load data in the Events panel.
2. In the Events panel toolbar, click the **Column Group** menu title.
   The menu drops down to display a list of column groups.

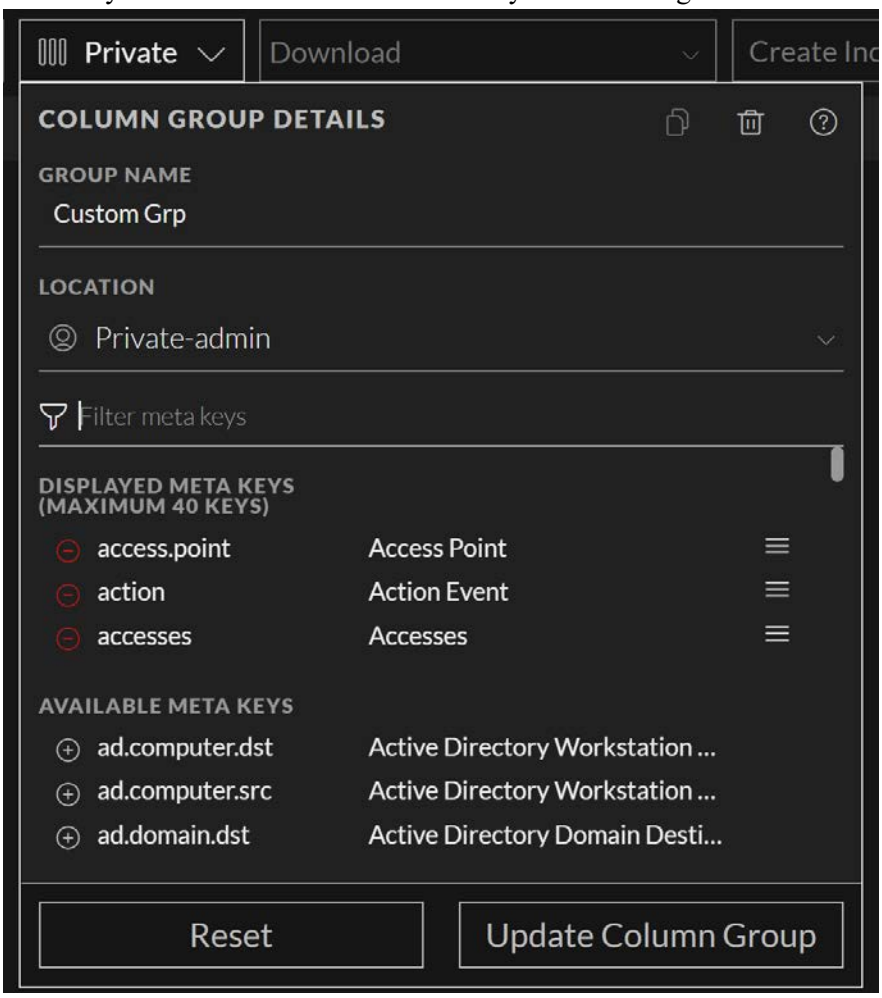


3. Highlight the column group that you want edit. This figure shows a custom column group highlighted, with the edit icon displayed to the right.



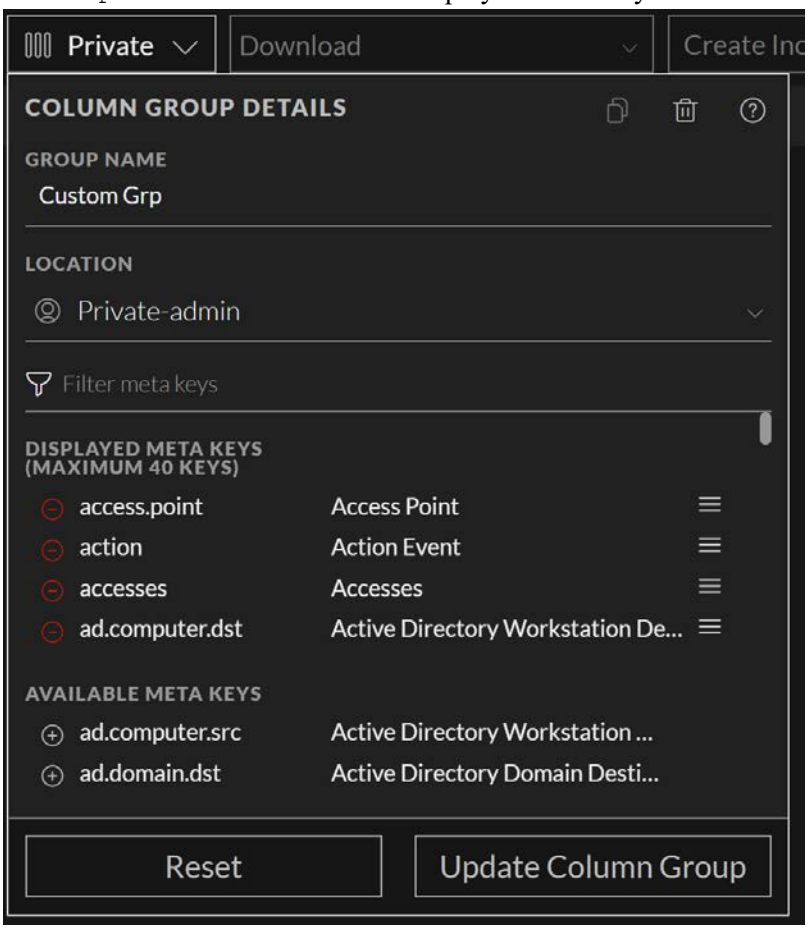
4. Click the edit icon ().
   The Column Group Details dialog is displayed so that you can edit the Group Name and Displayed

Meta Keys. You can add or delete meta keys and rearrange the order of the meta keys in the list.



5. (Optional) In the **Group Name** field, edit the name of the column group.
6. (Optional) To add a meta key to the column group, select and add each meta key as follows:
   a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list. Or just scroll through the list to find the meta key.
   b. When you see the meta key that you want to add, click the add icon  that precedes the meta key name.
   The meta key is added to the end of the Displayed Meta Keys list. (This list is also filtered using the text you typed.) This figure shows the group name changed to Column Group C and

`ad.computer.dst` added to the Displayed Meta Keys list.



7. (Optional) To find and remove a meta key from the column group, type a text string in the **Filter meta keys** field to look for meta keys that contain that text in the **Displayed Meta Keys** list, or simply scroll through the list. When you see the column that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
8. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
9. Do one of the following:
   a. To close the dialog without saving the changes to the custom column group, click **Reset**.
   b. To save the edits to the column group, click **Update Column Group**.
      The updated column group is saved globally for all analysts, and the buttons change to Done and Select Column Group.
10. Do one of the following:
    a. To close the dialog, click **Close**.

b. To close the dialog and select the updated column group, click **Select Column Group**.
The column group is updated, and if you clicked Select Column Group, the Events list is updated to show the columns in the new column group.

Create a Copy of a Column Group (Version 11.5 and Later)

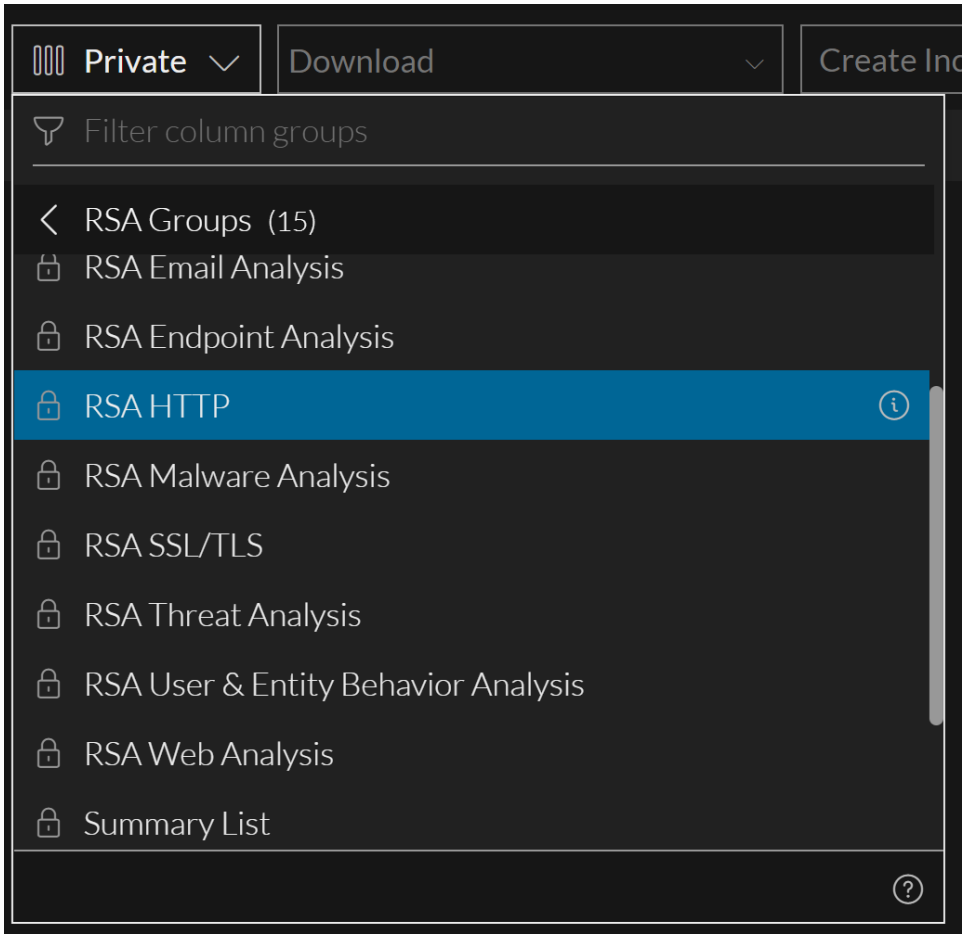
You can copy any column group, built-in or custom, shared or private, as long as it does not have unsaved edits in progress. This is useful when you want a customized version of a built-in group. Also since you cannot change a custom group from private to shared or from shared to private, creating a copy allows you to select a different Sharing setting. When you create a copy of a column group, the same name is used with a number appended. For example, if you copy RSA HTTP, the first copy is named RSA HTTP-1, and a second copy of the same group is named RSA HTTP-2. After you create the copy, you can edit the new group to give it a new name and manage meta keys in the group.

Note: Some column groups created in the Legacy Events view may have more the 40 columns, which is above the limit for column groups in the Events view. If you copy a group with more than 40 columns, you must remove the excess columns when you edit the column group.

To copy a column group:

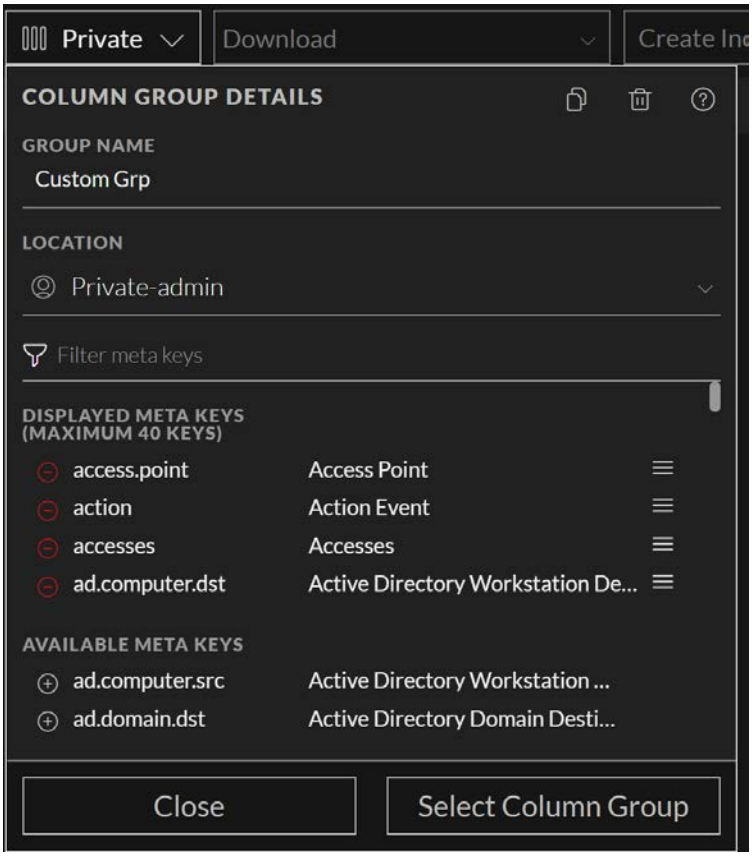
1. Go to **Investigate> Events** and submit a query to load data in the Events panel.
2. In the Events panel toolbar, click the **Column Group** menu title.
The menu drops down to display a list of column groups with the Filter Column Groups field at the top and the + New Column Group option at the bottom. The first group on the list is highlighted, and the selected group has a light blue background.

3. Highlight the column group that you want copy. This figure shows RSA HTTP highlighted. The information icon ( ) is displayed to the right.

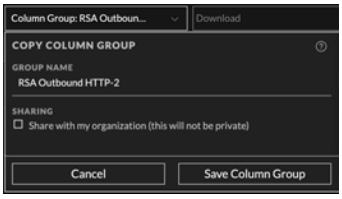


4. Do one of the following:
   a. Click the information icon ( ).

b. Click the edit icon ( ).
The Column Group Details dialog is displayed. This figure shows the dialog for a built-in group.

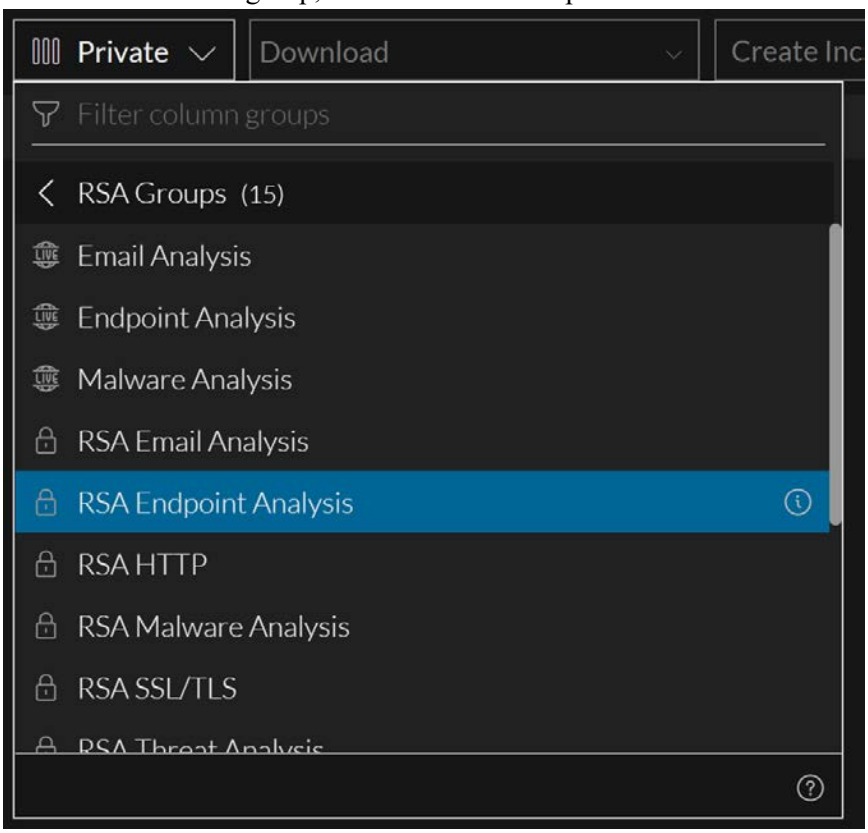


5. Click the Copy icon ( ).
The Copy Column Group dialog is displayed with a `-n` appended to the column group name. The following figure has -2 because it is the second copy of this column group.



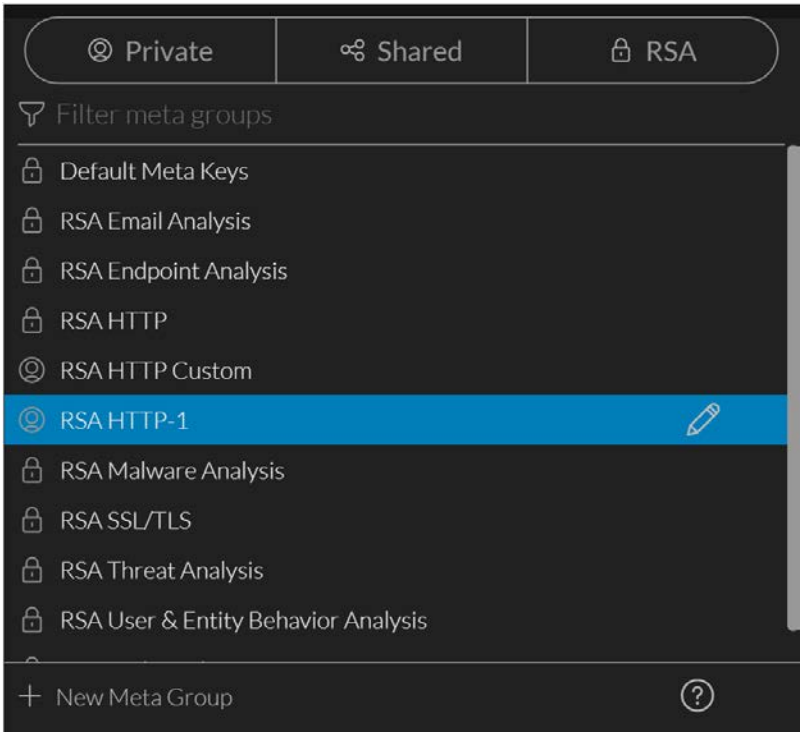
6. (Optional) In the **Group Name** field, edit the name of the column group.
7. If you want to share the new column group with your organization, set the **Share with my organization** option. By default the new group is private.
8. Do one of the following:
   a. To close the dialog without copying the group, click **Cancel**.
   b. To save the copy of the column group, click **Save Column Group**.
   The copy of the column group is saved, and the buttons change to Done and Select Column Group.

9. Do one of the following:
   a. To close the dialog, click **Close**.
   b. To close the dialog and select the copy of the column group, click **Select Column Group**.
   The column group is copied, and if you clicked Select Column Group, the Events list is updated to show the columns in the copy of the column group. The figure below has two copies of the RSA HTTP column group, one shared and one private.



1. Do one of the following:
   a. To close the dialog without editing, click **Close**.
   b. To close the dialog and select the copy of the meta group, click **Select Meta Group**.
   The group is added to the Meta Group menu. The figure below has a private copy of the RSA

HTTP meta group.



Create a Column Group Folder

You can create custom column group folders which reside at the current level and are be added as a private or shared folder. And, if the folder name already exists then you are prompted to provide a unique name.

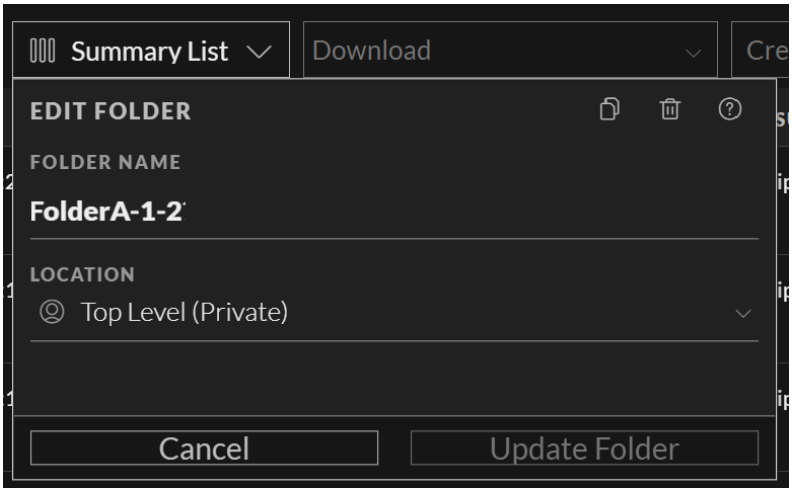
1. With the Filter Events panel open in the Events view, click the Column Groups menu title. The menu drops down to display a list of column groups and folders.
2. Click .
   The Create Folder dialog is displayed.
3. In the **Folder Name** field, type a unique name (maximum length of 255 characters) for the new meta group folder.
4. Click **Create Folder**.

Edit and Move Column Group Folder

After you create a column group folder you can edit or move it, however the folders inside RSA Groups (RSA Live content and RSA OOTB Groups) cannot be edited and moved. The folders inside private and shared folders can be edited and moved only within their respective groups. For example, you cannot move a shared folder into a private folder and vice-versa.

1. With the Filter Events panel open in the Events view, click the Column Group menu title and highlight the column group that you want edit.

2. Click .
   The Edit Folder dialog is displayed.

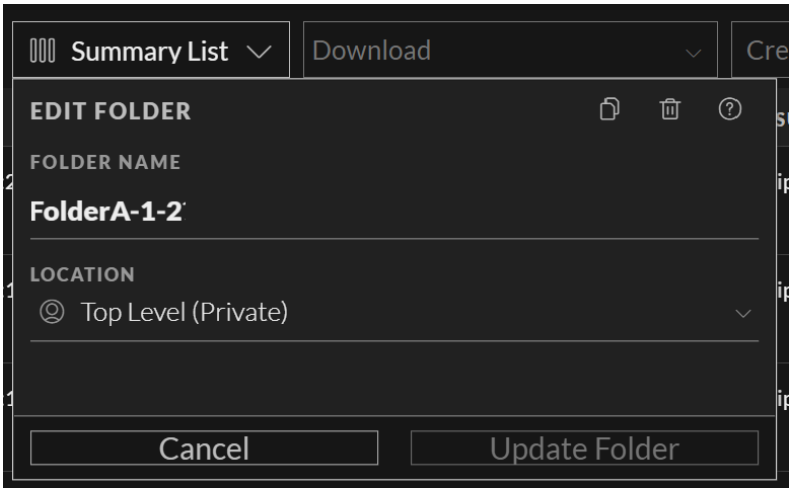


3. In the **Folder Name** field, type a unique name for the column group folder.
4. Select the location of the folder to be edited.
5. Click **Update Folder**.

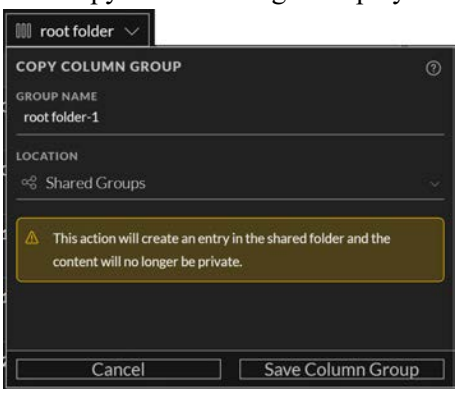
Copy Column Group Folder

You can copy column group folders from private to shared, private to private, shared to shared and shared to private groups. When you copy a folder the content inside it gets copied. When you copy a private folder into a shared folder, the folder and its content no longer remain private.

1. With the Filter Events panel open in the Events view, click the Column Group menu title. The menu drops down to display a list of column groups and folders.
2. Select a folder you want to copy.
3. Click edit and the click copy .



The Copy Folder dialog is displayed.

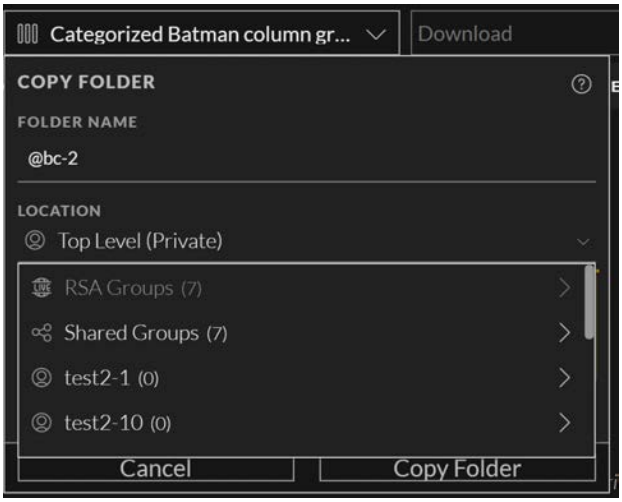


4. In the **Folder Name** field, type a unique name for the new meta group folder.
5. Select the location of the folder to be edited.
6. Click **Copy Folder**.

Copy Group Folder Deployed from Live

You can copy column group folder deployed from Live located under RSA Groups category to any other location like Shared groups or to a private folder.

1. With the Filter Events panel open in the Events view, click Column Group menu title. The menu drops down to display a list of column groups and folders.
2. Click on a Live Column Group folder you want to copy.
3. Click 
   The Copy Folder dialog is displayed.

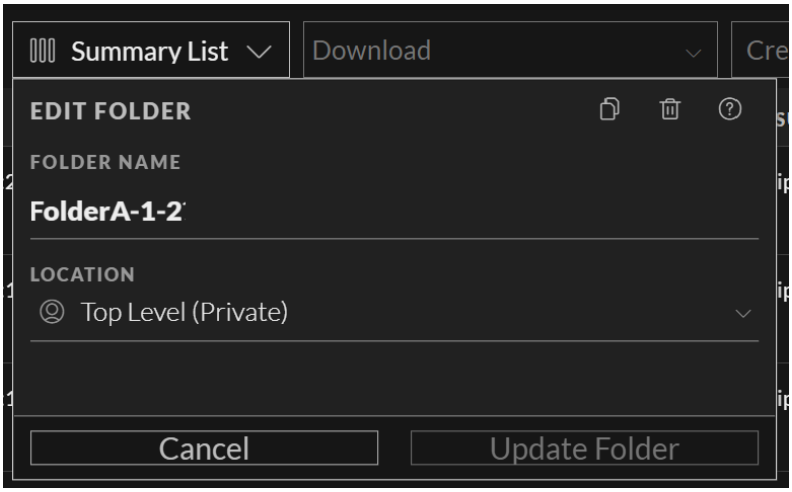


4. Select the location of the folder to be copied.
5. Click **Copy Folder**.
   The folder is created with the original name of the folder and appended with the 'copy' in the end.

Delete Column Group Folder

If you don't want to retain a folder you can delete it. However, once the folder is deleted it cannot be retrieved.

1. With the Filter Events panel open in the Events view, click the Column Group menu title. The menu drops down to display a list of column groups and folders.
2. Select a folder to be deleted.
3. Click edit .
   The Edit Folder dialog is displayed.



4. Click delete .
   A warning message is displayed to confirm the action.



5. (Optional) Select the checkbox, if you want to delete the folder along with all the contents inside the selected folder.
   If you do not select the checkbox, then the content will be moved to the parent folder after the required folder is deleted.
6. Click **OK** to delete.

Select a Column Group and Columns (11.3 and Earlier Event Analysis View)

In the 11.3 and earlier Event Analysis view, you can select a column group to apply to the Events list. These are the built-in groups and custom column groups created in the Legacy Events view.

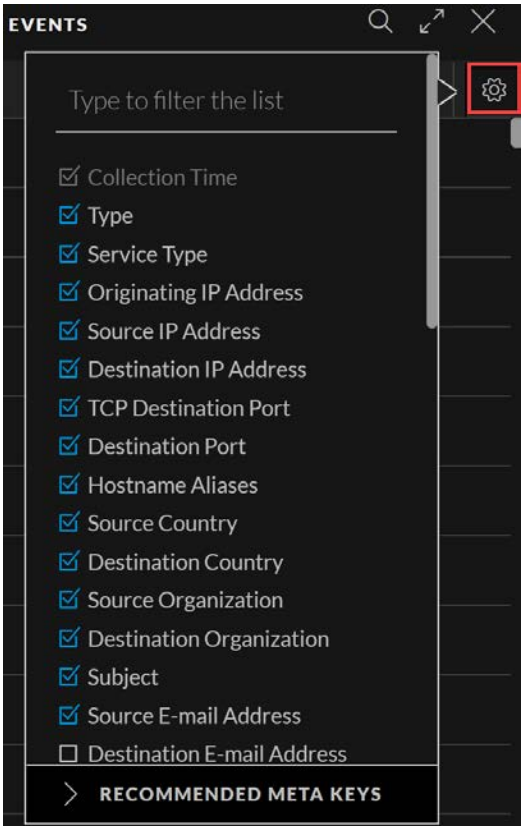
To select a column group:

Begin in the Column Group menu and do one of the following:

1. Select a column group (for example, **Summary List**).
2. To filter the list of column groups, start typing the name of the column group. Type one character and a list of column groups that contain that character is displayed; as you continue to type the list is filtered to match. When you see the column group that you want, click to select it. To clear the filter text, click **X** or delete the typed text.
   The Events panel displays data in the columns that belong to the selected column group.

To select columns to display:

1. With the Events list open and a column group selected, click ⚙ to display the column selector.



2. Select the meta keys or enter the name of a meta key that you want to display in additional columns.
3. If you do not want to see a meta key displayed in a column, deselect the meta key.
   The data is redisplayed using the selected columns.

Work with Column Groups in the Legacy Events View

This section includes procedures for working in the 11.4 Legacy Events view (and the 11.3 Events view). Three different forms of the events list with hard-coded columns are built in and labeled as follows: Detail View, List View, Log View. You can remove columns, rearrange the order, and change the width of a column. In addition the built-in or custom column groups are available; these give you more flexibility in choosing columns.

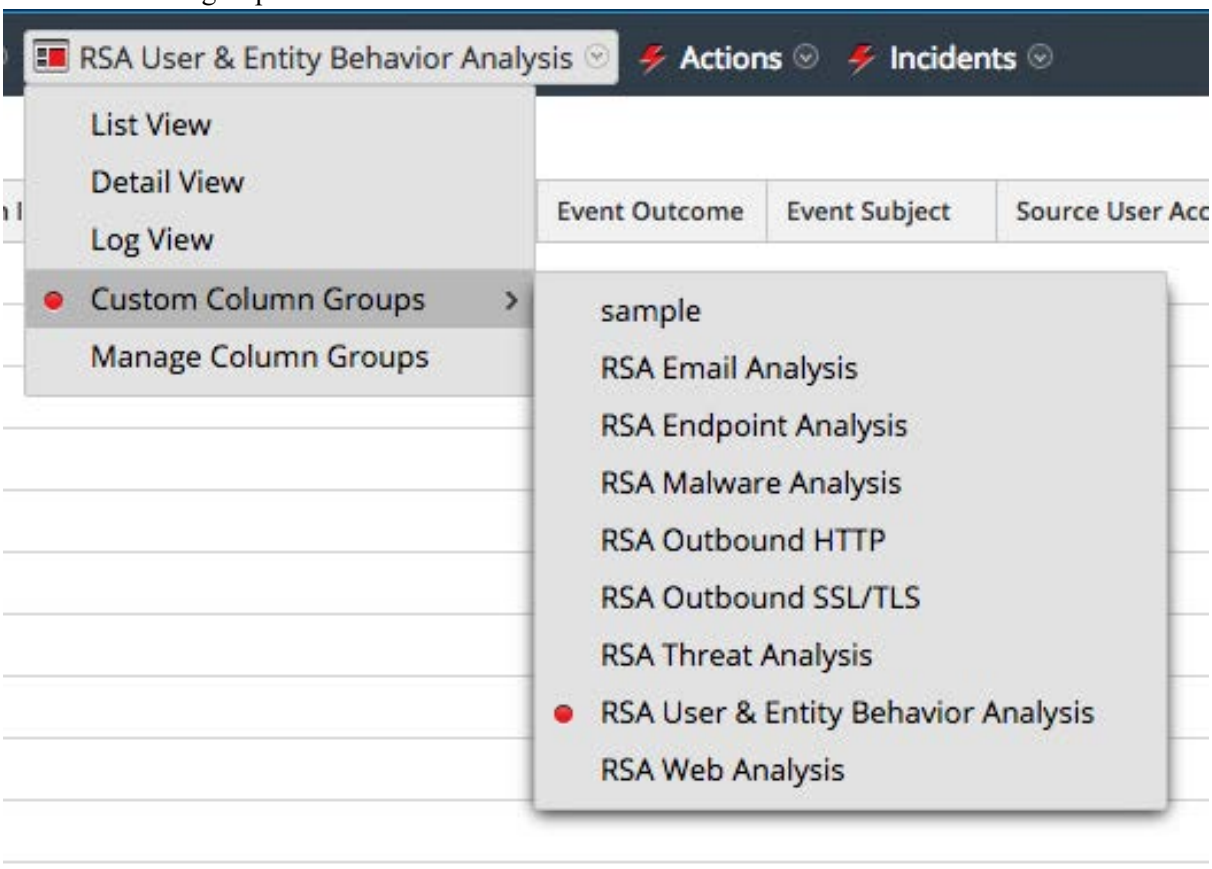
Column groups are shared globally, per service, across Investigate. Any changes you make to custom column groups are applied globally, affecting all analysts using the service. If you delete a column group, the column group is no longer available to anyone who is investigating the service.

Select a Column Group

Note: Investigate profiles can include custom column groups. If a custom column group is used in a profile and you are viewing events in the Legacy Events view using a custom column group, you cannot change the view type (Detail, List, or Log).

To select a column group:

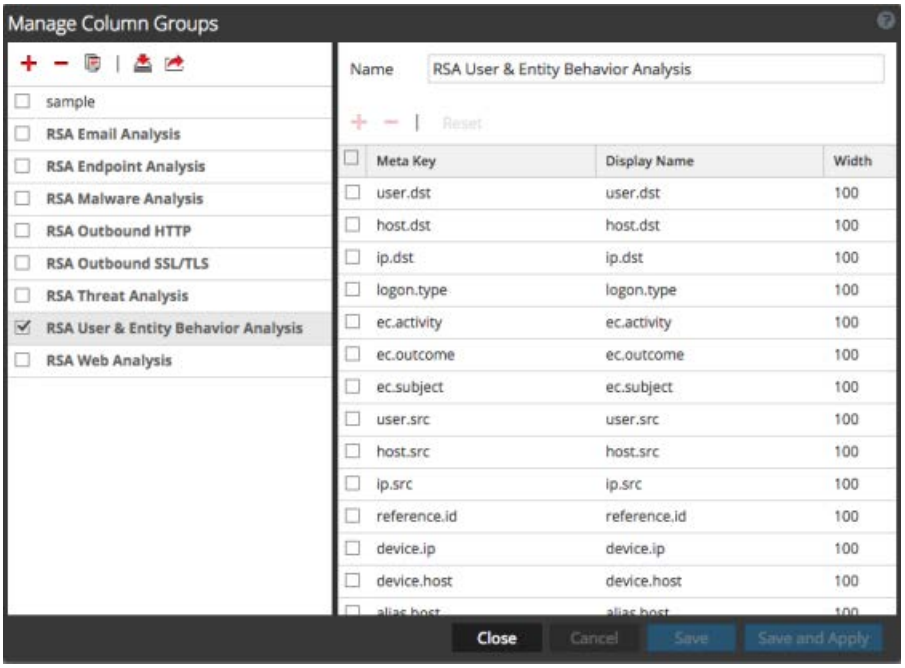
1. With the Legacy Events view open, select **Custom Column Groups** in the **View** drop-down menu. The menu label reflects the selected option: Detail View, List View, Log View, or the currently selected column group.



2. Select one of the column groups from the submenu.
   The Legacy Events view is refreshed to reflect the custom column group.

Create a Custom Column Group in the Legacy Events View

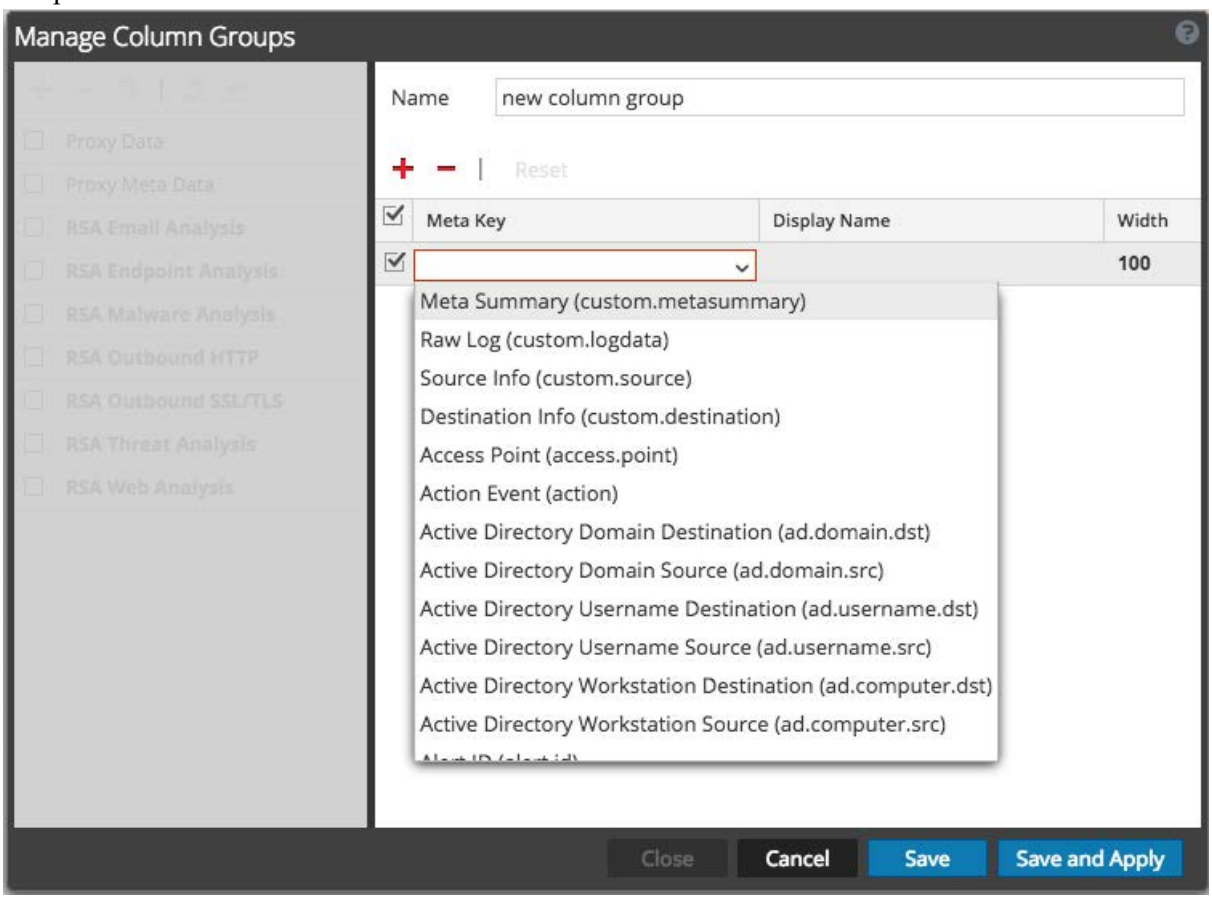
1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.
   The Manage Column Groups dialog is displayed.



3. To add a new column group in the column group panel, click + and type the name of the new group in the resulting field.
   The column definition panel opens on the right with the group name filled in. You can edit the group name.

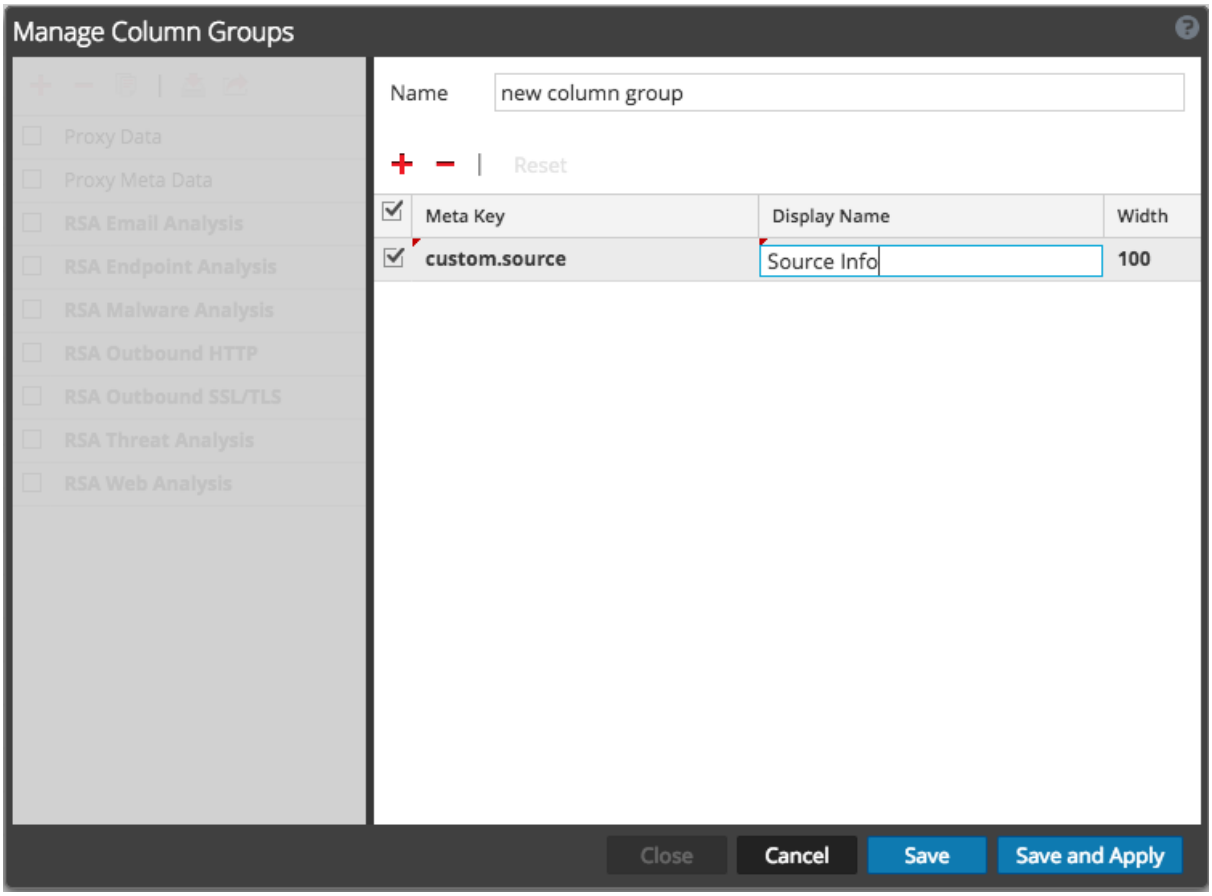
4. To add a column to the group, click +, and click in the empty **Meta Key** field to display the **Meta Key** drop-down list. Select a meta key field from the list, and repeat this step until the column set is

complete.



5. (Optional) To delete a meta key from the column group, click −.
6. (Optional) To rearrange the sequence in which the columns appear in the Events list, drag meta keys to the desired position.

7. (Optional) To set the default width for a column, click in the corresponding value in the **Width** column, and type a new column width.

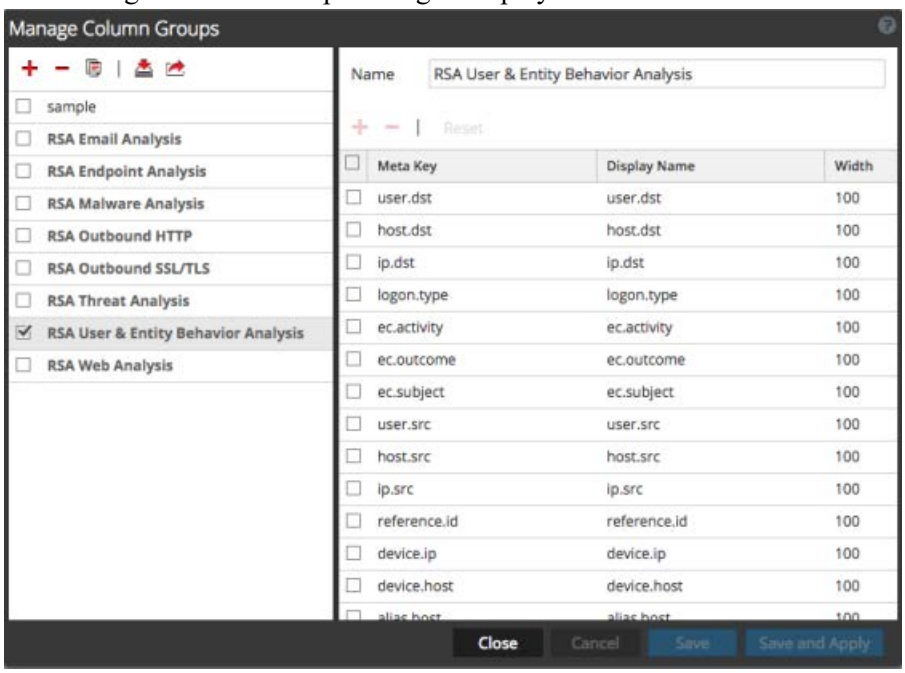


8. (Optional) To revert to the previous settings for the column group, and undo all of your changes, click **Cancel**.
9. When ready to save, do one of the following:
   a. To save the edited column group and refresh the Legacy Events view with the column group settings, click **Save and Apply**.
   b. To save the edited column group without refreshing the Legacy Events view, click **Save**.

Delete a Column Group (Legacy Events View)

1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.

The Manage Column Groups dialog is displayed.

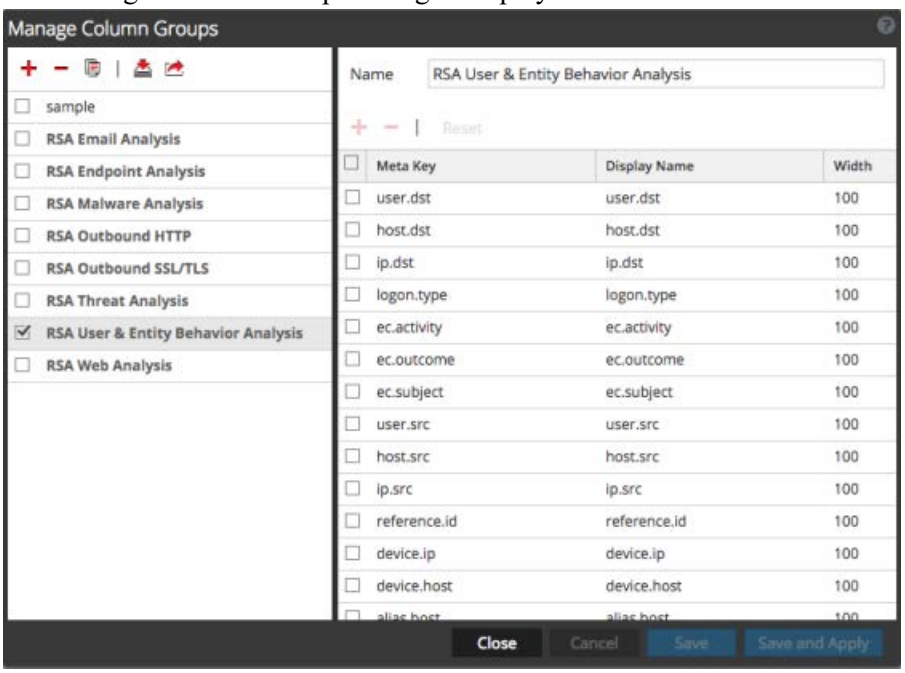


3. To delete a custom column group in the column group panel, select one or more custom column groups and click − in the toolbar.
   A confirmation request is displayed.
4. Do one of the following:
   a. To delete the column group and refresh the Legacy Events view, click **Yes**.
   b. If you decided not to delete the column group , click **No**.
      The selected column groups are deleted and no longer appear anywhere for this service in Investigate.

Edit a Column Group (Events View)

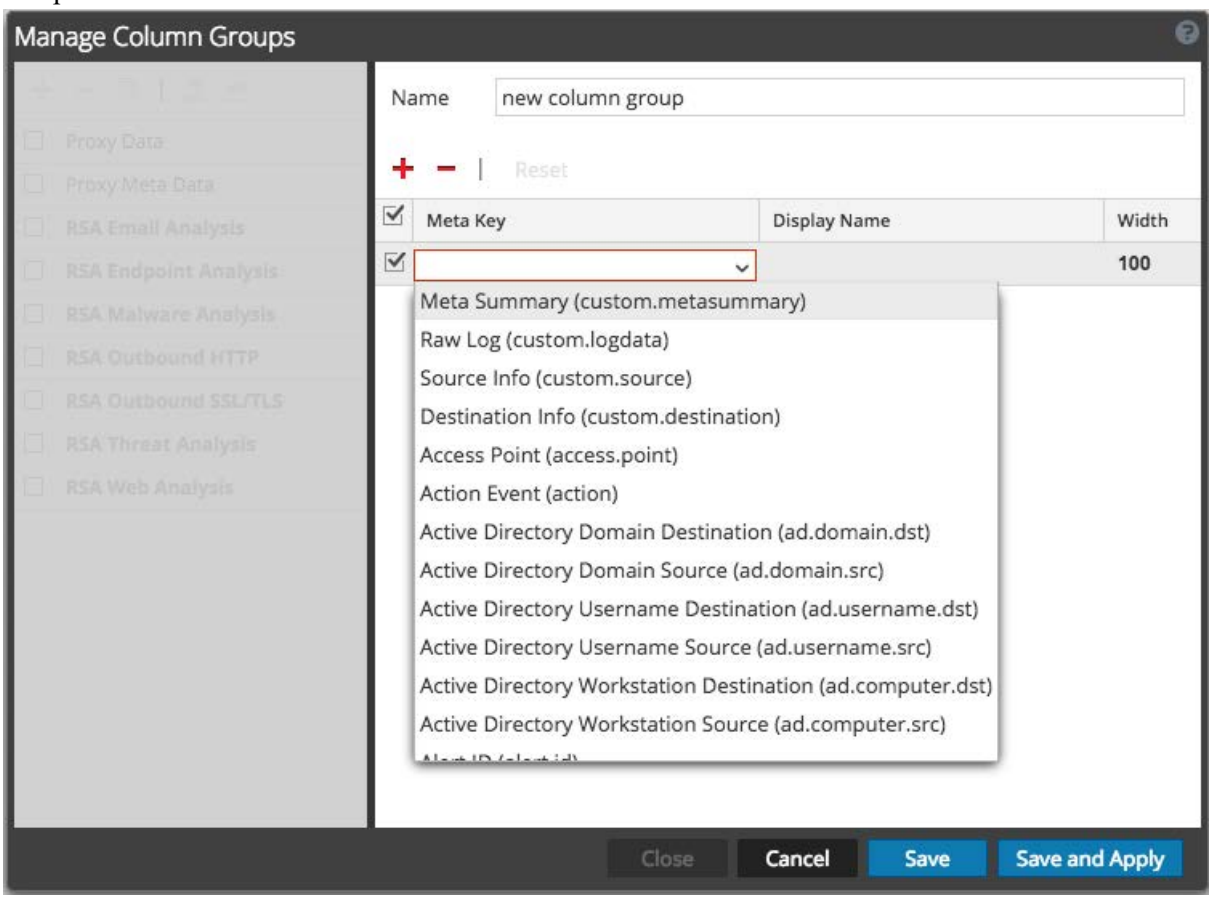
1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.

The Manage Column Groups dialog is displayed.



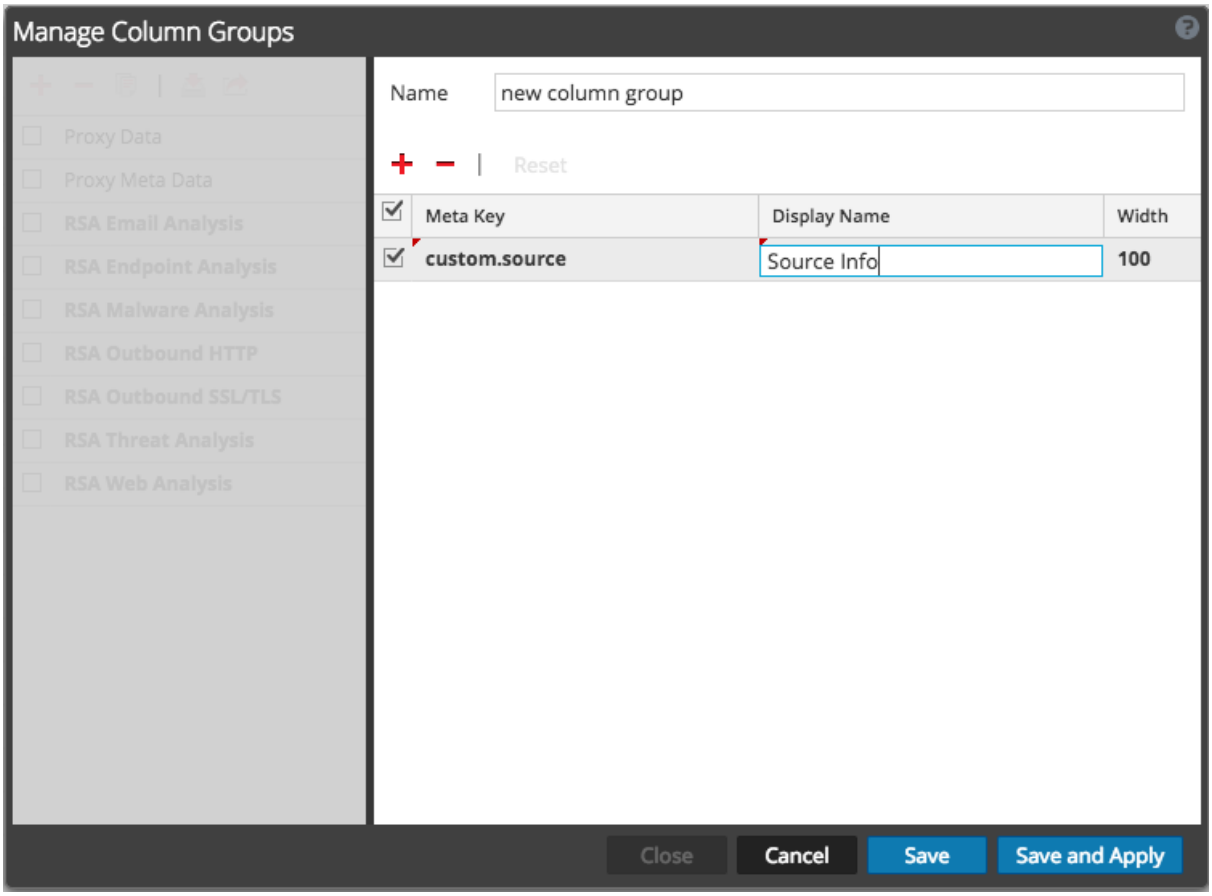
3. Do one of following:
   a. To edit a custom column group in the column group panel, select the checkbox before the name. The column definition panel opens on the right.
   b. To clone and edit a built-in column group or a custom column group, select the checkbox before the name and click the clone icon ( ).
   The column definition panel opens on the right.
4. (Conditional) If you are editing a clone of a group, type the new name of the group.
5. To add a column to the group, click , and click in the empty **Meta Key** field to display the **Meta Key** drop-down list. Select a meta key field from the list, and repeat this step until the column set is

complete.



6. (Optional) To delete a meta key from the column group, click −.
7. (Optional) To rearrange the sequence in which the columns appear in the Events list, drag meta keys to the desired position.

8. (Optional) To set the default width for a column, click in the corresponding value in the **Width** column, and type a new column width.



9. (Optional) To revert to the previous settings for the column group, and undo all of your changes, click **Cancel**.
10. When ready to save, do one of the following:
    a. To save the edited column group and refresh the Legacy Events view with the column group settings, click **Save and Apply**.
    b. To save the edited column group without refreshing the Legacy Events view, click **Save**.

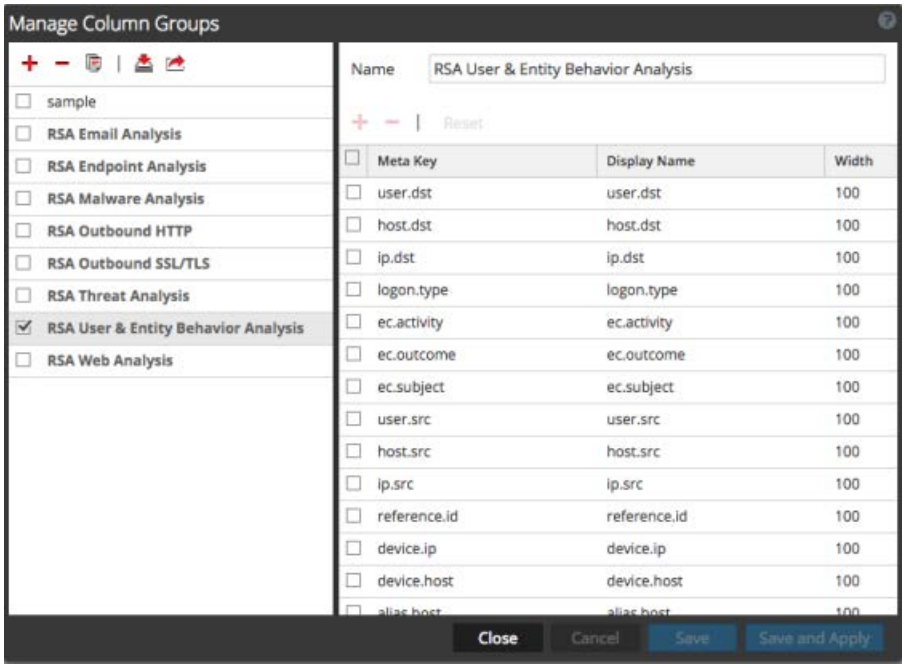
Import and Export a Column Group (Legacy Events View)

You can export custom column groups for use by other members of your team, and other analysts can import column groups if you give them a copy of the exported file.

To export a column group:

1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group. Each of these views is a differently formatted events list, and each column represents one

meta key.
The Manage Column Groups dialog is displayed.



3. To export a column group, select the checkbox before the name and click the Export option ( ).
   The column group is exported to your local file system as a .jsn file, for example, `CustomColumnGroupsExport.jsn`. If you export another group, the next file is named `CustomColumnGroupsExport-2.jsn` to differentiate.
4. To import a column group that you have available on your local file system, click the Import option ( ).
   The Import Column Groups dialog is displayed.
5. Browse your local drive to find the column group (jsn file), and click **Upload**.
   The column group is added to the list. If it has the same name as an existing column group, a message is displayed and the column group is not imported.

Use Query Profiles to Encapsulate Common Areas for Investigation

Query profiles offer a quick and easy way to define a meta group, column group, and a limiting filter (pre-query condition) that you can apply in the Navigate view, the Events view, and the Legacy Events view. The same query profiles are shared between all views, and they are available in the Springboard (Version 11.5) for use in panels. Private query profiles created in the Events view are only available in the Events view for the analyst who created them.

Each query profile specifies a meta group, column group, and sometimes includes a pre-query condition appropriate for the type of investigation.

In a query profile:

- The meta group defines the meta keys that are queried (see Use Meta Groups to Focus on Relevant Meta Keys).
- The column group defines which meta keys from the meta group are displayed as columns in the Events list. (see Use Columns and Column Groups in the Events List).
- When the query profile is in effect, the optional pre-query conditions add a limiting filter in the query bar. You can edit or delete the limiting filter and then create additional filters for your query (see Filter Results in the Events View).

Built-In Query Profiles

You cannot edit or delete built-in profiles, but you can copy an existing profile and edit the copy in the Navigate view, the Legacy Events view, or the Events view. In the Navigate view, the built-in profile names begin with the RSA prefix and are grouped under Default Profiles. The Events view does not support grouping of query profiles. This figure is an example of a built-in query profile as listed in the Query Profiles menu.



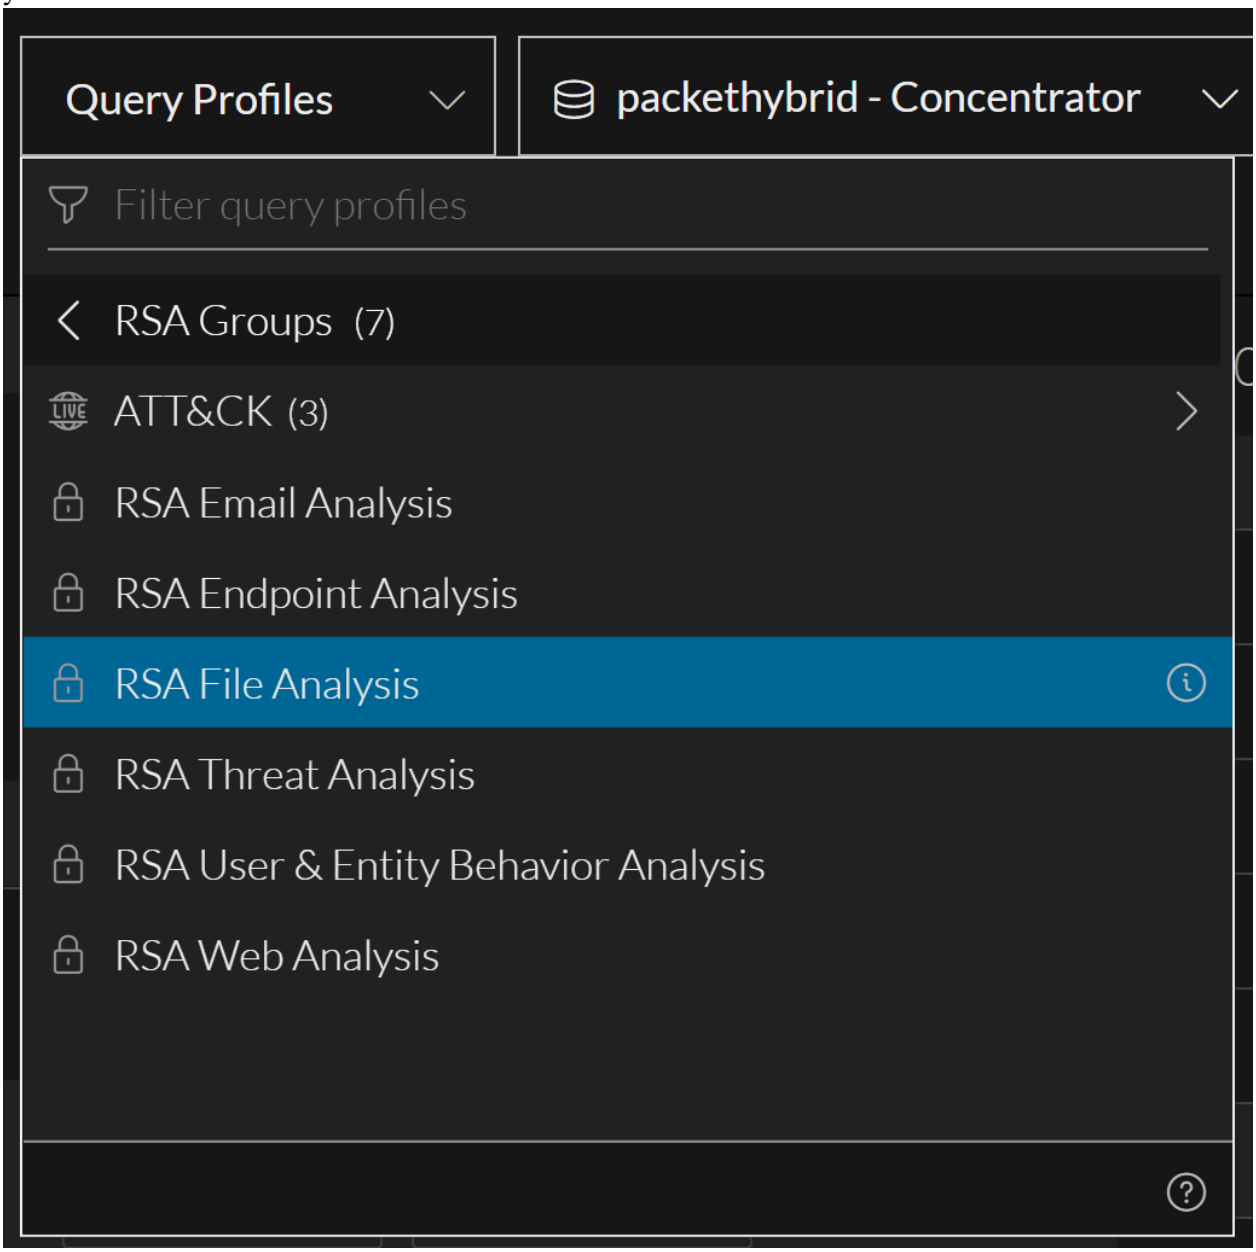
NetWitness Platform has these built-in profiles:

- RSA Email Analysis
- RSA Endpoint Analysis
- RSA File Analysis
- RSA Threat Analysis
- RSA User & Entity Behavior Analysis
- RSA Web Analysis

Built-in query profiles make it easy for you to query a specific area of interest; for example, selecting the built-in RSA Email Analysis query profile automatically specifies the meta group, and column group, and pre-query conditions that are most useful for investigating email activity. As you become familiar with the meta keys, you can create your own custom query profiles.

Live Query Profiles

In 11.6 and later, NetWitness supports deploying the investigate content from live and are marked by the live symbol () under the query profiles group drop down. The query profiles are categorized as RSA Groups (RSA Live content and RSA OOTB Groups), and Shared Groups. The groups are displayed as non-editable folders and sub-folders except for Shared Groups that can be edited. All private content is displayed outside these groups. For example, the below image shows private content below the Shared Groups folder. The number inside () depicts the number of contents inside a folder and > symbol helps you to drill down inside the folder.



Custom Query Profiles

Custom query profiles are shared globally within your organization in Version 11.4. In Version 11.5 and later, you can create shared query profiles as before, and can also create private query profiles. If you edit a shared custom query profile, your changes are applied globally. If you delete a shared custom query profile, the profile is deleted and no longer available for all analysts.

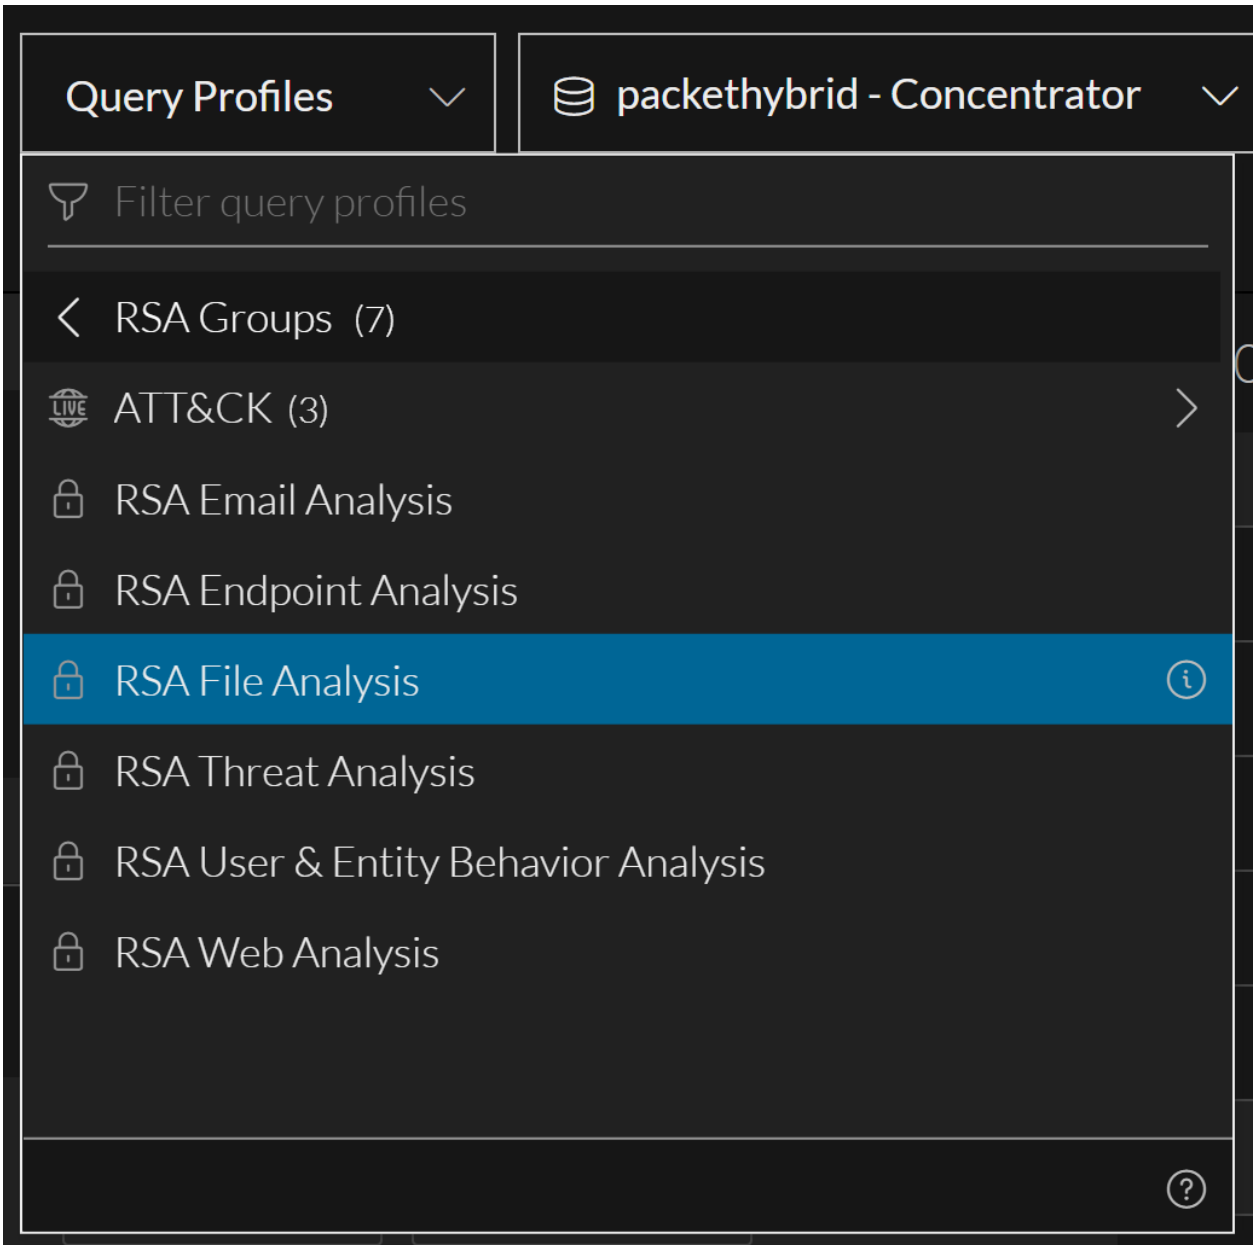
Note: If a Springboard panel is using a query profile as a filter, the profile can be edited, but cannot be deleted in the Events view. However, nothing prevents deletion of the profile in the Navigate view or the Legacy Events view. In this case, Springboard panels that use the deleted query profile as a filter continue to work, but the filter is removed and unexpected results may be displayed in the panel. Refer to "Managing the Springboard" in the *NetWitness Platform Getting Started Guide* for details.

When you create a query profile in Version 11.5, you can choose to share it or you can keep it private (default); you cannot change a shared profile to private or a private profile to shared. Private query profiles are not visible or usable in the Navigate view, the Legacy Events view, or the Springboard. Icons identify the profile type in the Query Profile menu. These are examples of a shared and a private custom query profile as listed in the Query Profile menu, with the edit icon displayed at the end of the row.

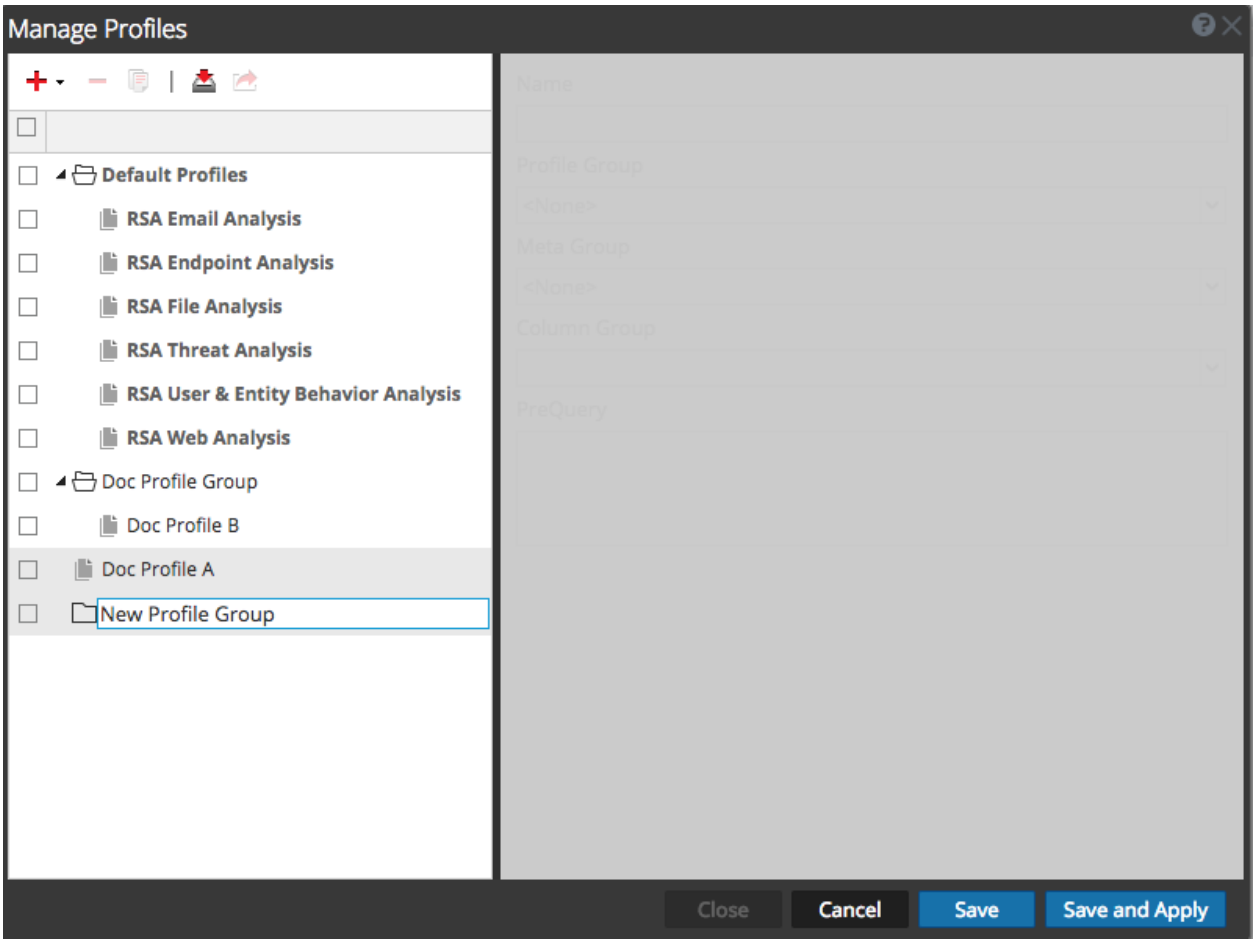


Dialogs for Managing Query Profiles

The profiles are listed in alphabetical order in the Query Profiles menu in a way that makes built-in profiles distinguishable from custom profiles that you imported or created. While the functionality for managing query profiles is similar in the Navigate view, the Legacy Events view, and the Events view, the dialogs are different. The following figure illustrates the Query Profiles menu in the Version 11.5 Events view. This menu lists the same profiles that are available in the Navigate view and the Legacy Events view. You can create, copy, edit, delete, and apply profiles.



This is an example of the Manage Profiles dialog in the Navigate and Legacy Events views.



Note: Query profiles are available in the Navigate view, the Legacy Events view, and the Events view; in Version 11.4.1 and earlier, they are shared globally across users. If one user modifies or deletes a custom query profile it has an effect on what is available to the other users. In the Events view, use the Query Profiles menu to work with profiles. In the **Navigate** or **Legacy Events** view toolbar, select **Profile > Manage Profiles** to open the Manage Profiles dialog. In Version 11.5, custom profiles can be shared globally, but private custom profiles created in the Events view are not available in the Navigate view or the Legacy Events view.

From the Query Profiles menu (11.4 and later Events view):

- You can apply a query profile and use options in the menu to create (Create Query Profile dialog), copy, edit, and delete (Query Profile Details dialog) custom query profiles.
- Selecting a profile applies the meta group, column group, and pre-query condition, and these are visible in the Meta Group menu title, Column Group menu title, and the query bar.
- In Version 11.4, the Events view does not use meta groups or profile groups defined in other views. Version 11.5 allows you to use meta groups and to create private custom query profiles, in addition to the previously available shared custom query profiles.
- If a query profile created in the Legacy Events view uses the Log View, Detail View, or List View instead of a column group, the same profile in the Events view uses the Summary List column group.

From the Manage Profiles dialog (Navigate view and Legacy Events view):

- You can configure, add, delete, import, and export profiles and profile groups.
- You can organize your custom query profiles in profile groups (Version 11.2 and later). When upgrading to Version 11.4 from an earlier version, only profile groups that contain profiles are imported. The built-in query profiles are in the Default Profiles group, which cannot be edited. Analysts can create new query profile groups, which anyone can use.
- After creating profiles, you can edit a profile group to add profiles, remove profiles, or move profiles from one group to another. When you create a profile, it is not added to any profile group by default.
- Selecting a profile applies the meta group, column group, and pre-query condition, and the label of the Profile menu is replaced with the query profile name. The following figure illustrates the RSA Email Analysis query profile selected in the Navigate view or Legacy Events view.

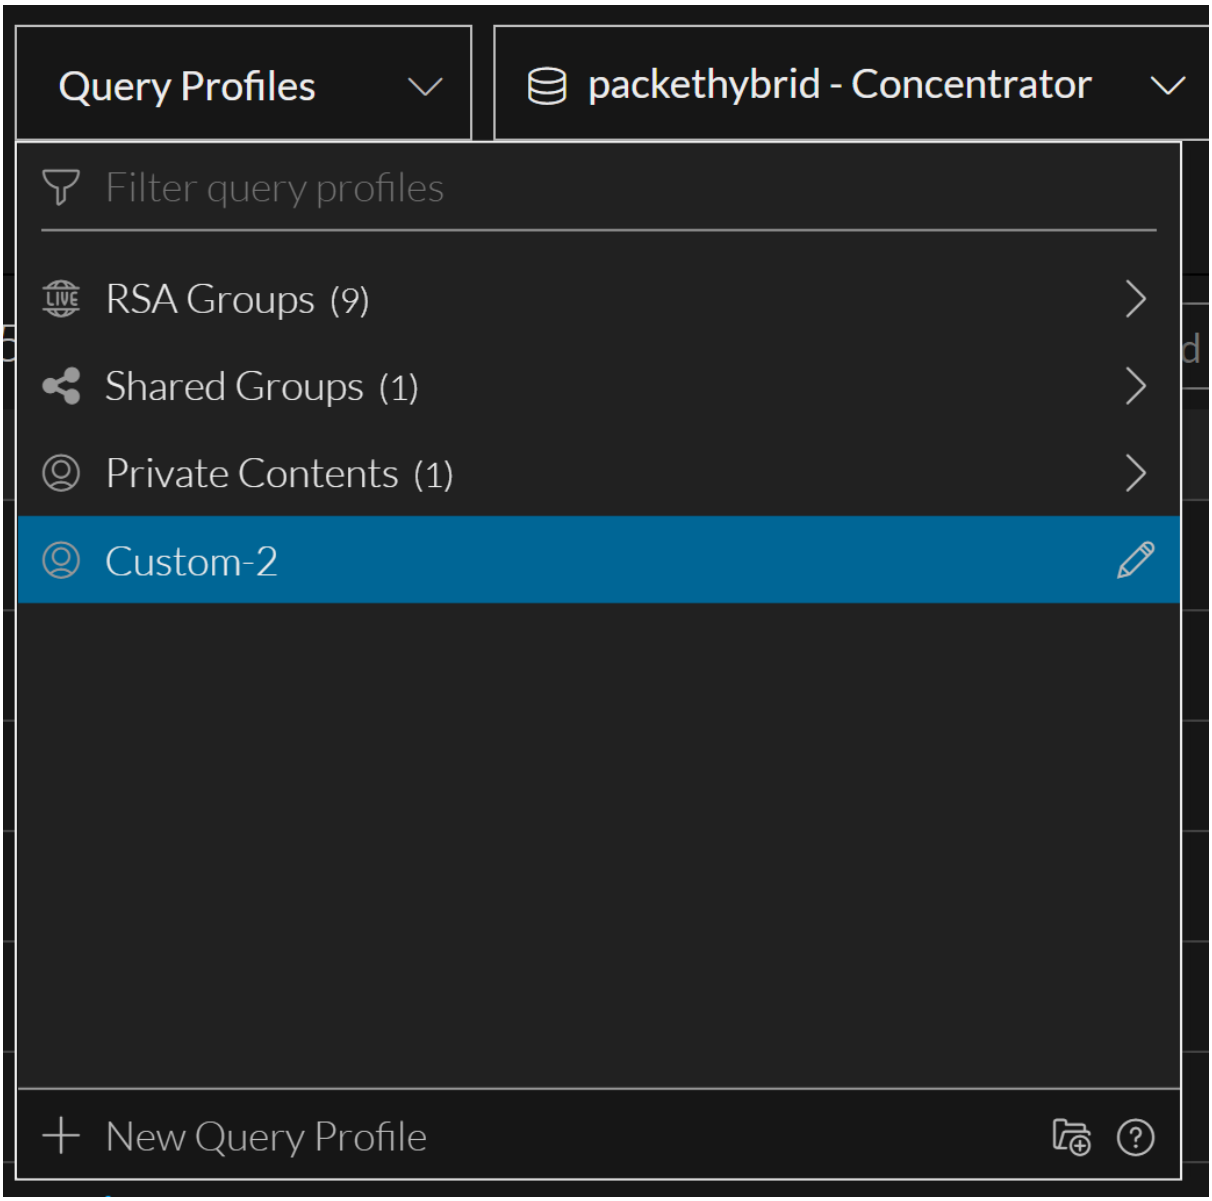


View Query Profile Details (Events View)

If you want to know which meta groups, column groups, and limiting filters (called pre-query conditions) define a query profile, you can view the details of the profile.

To view the details:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles. This menu displays a list a list of built-in query profiles (RSA), shared custom profiles, and your private custom profiles with visibility options and a filter field make it easier to find a particular query profile.



2. Hover over a query profile in the list and click the information icon ( ) to see the meta group, column group, and pre-query conditions configured for the profile.
   This figure shows the details for the RSA Email Analysis profile, one of the built-in profiles. In Version 11.5.1, an icon identifies the type of meta group and column group (shared, private, or RSA).

3. Do one of the following:

   a. To close the dialog, click **Close**.

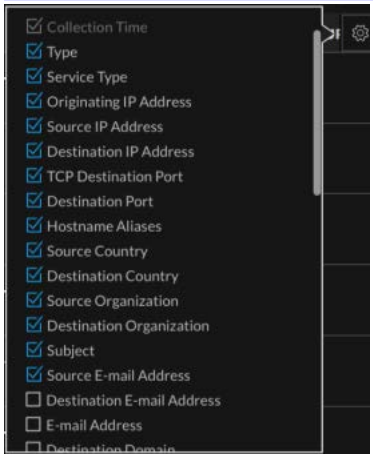
   b. If you want to apply the profile, click **Select Query Profile**.
      The dialog closes. The Events list is updated to reflect the selected query profile. If the profile uses a different column group, the query is re-executed with the pre-query conditions and column

group for the selected profile. If only the pre-query conditions are different, existing filters in the query bar are removed and the pre-query conditions (for example, this filter: `service=24,25,109,110,995,143,220,993`) is added in the query bar, but the query is not submitted. The first 15 columns in the associated column group are used in the Events list.

i. (Optional) Create additional filters in the query bar before executing the query (see Filter Results in the Events View).



ii. (Optional) If you want to select different columns from the associated column group before executing the query, click ⚙ above the Events list on the right.
The Column Selection list is displayed and you can choose up to 40 columns to display (see Use Columns and Column Groups in the Events List.



Apply a Query Profile (Events View)

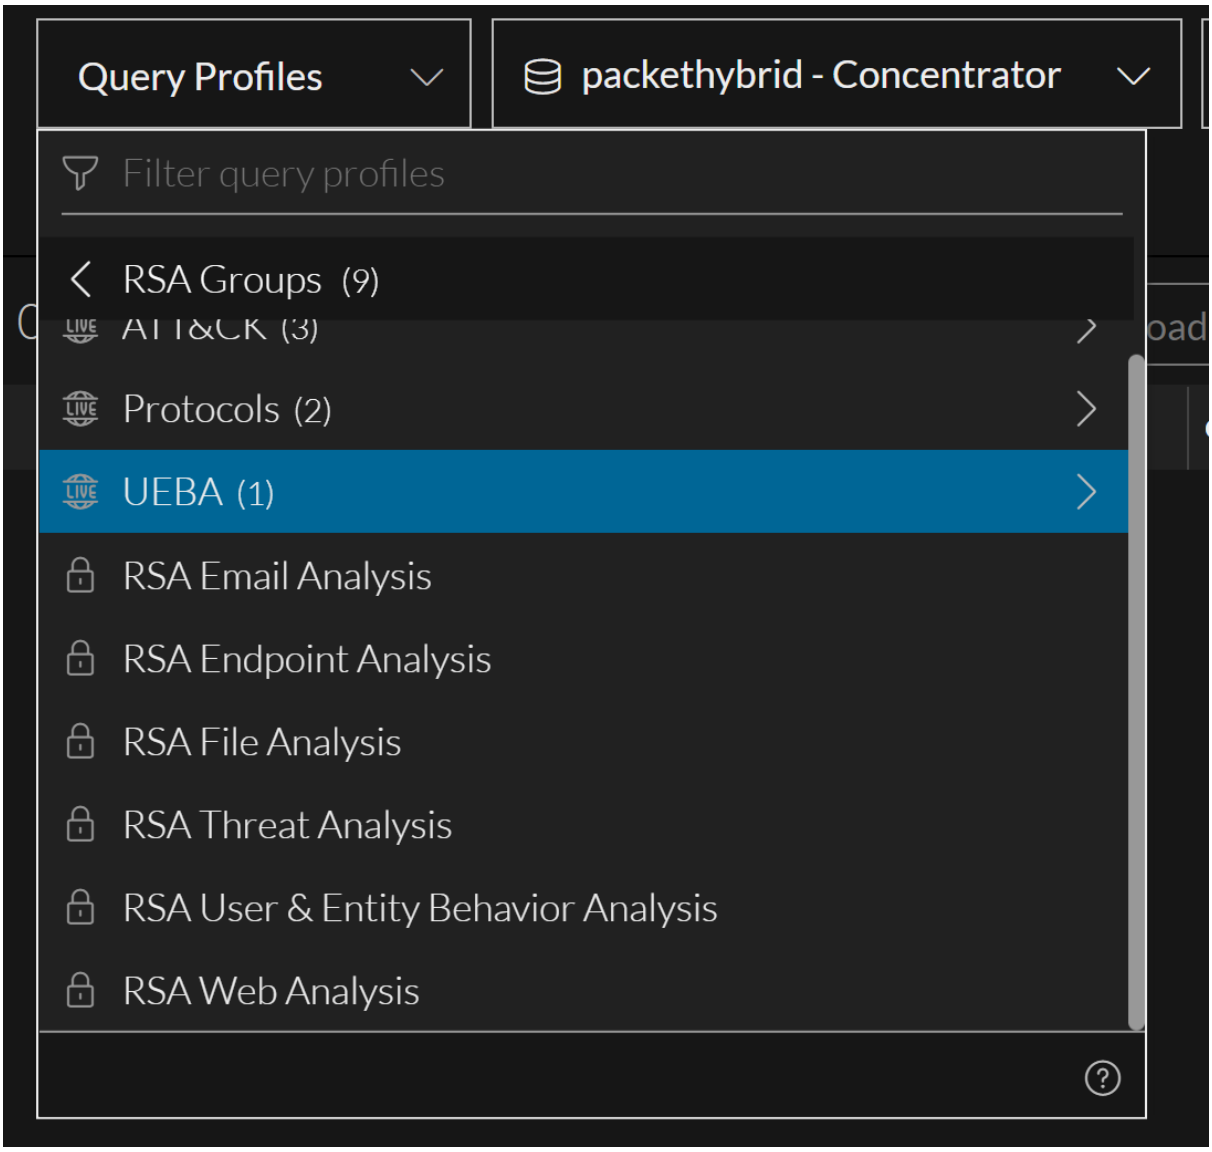
When a query profile is applied, there is no indication of it in the Query Profile menu, but you can see if a column group or meta group is in effect. If pre-query conditions are applied, the filters are visible at the beginning of the query bar as shown in this figure:



Note: If you do not see enough results or the right results in the Events view, an applied profile may be limiting results with pre-query conditions.

To apply a query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles.

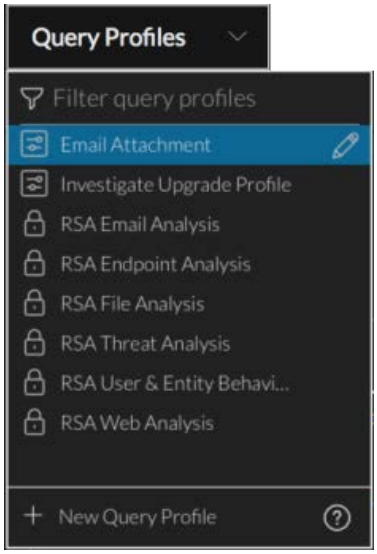


2. Use the Down and Up arrow keys or the mouse to highlight a profile.
3. Click the highlighted profile.
   The query profile settings are applied immediately. The Events list is updated to reflect the selected profile. If the profile uses a different column group the query is re-executed with the pre-query conditions and column group for the selected profile. If only the pre-query conditoins are different, existing filters in the query bar are removed and the pre-query conditions are added in the query bar. The button becomes active so that you can resubmit the query with the new pre-query conditions. You can add more filters as usual before or after resubmitting the query.

Create or Edit a Custom Query Profile (Events View)

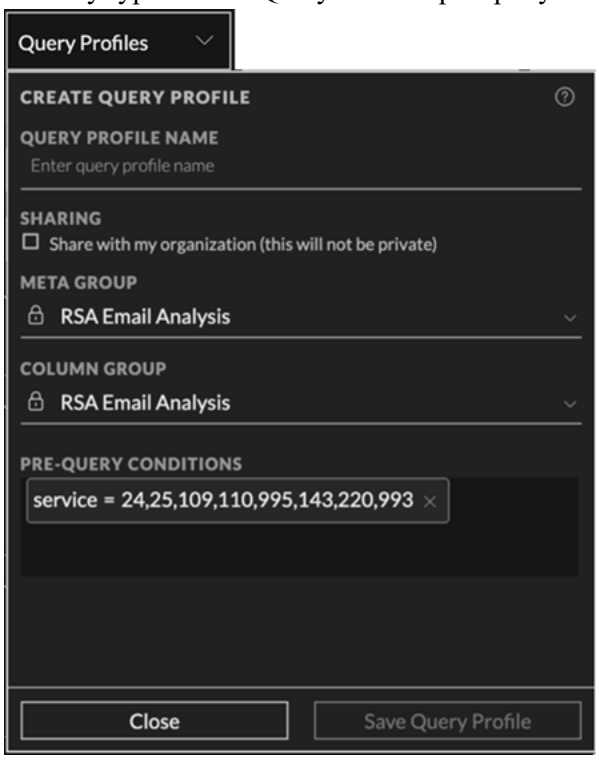
To create or edit a custom query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles.

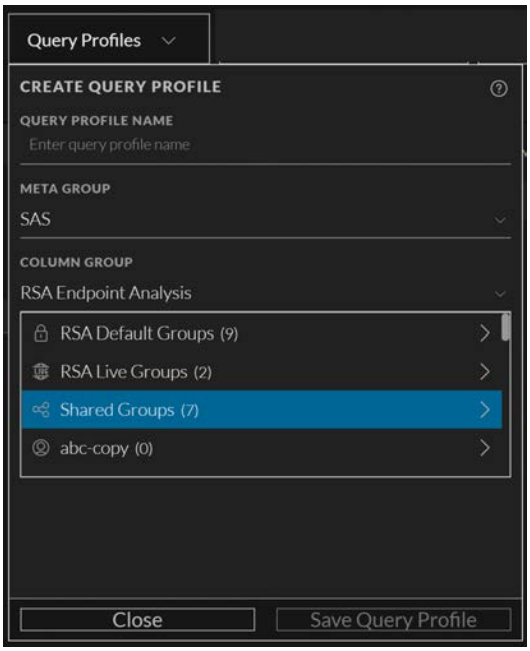


2. Do one of the following:

   a. To create a new query profile, click **+ New Query Profile**.
   The Create Query Profile dialog is displayed. The Create Query dialog shows a new empty profile that includes the currently selected meta group, column group, and filter that you have currently typed in the Query bar as a pre-query condition.

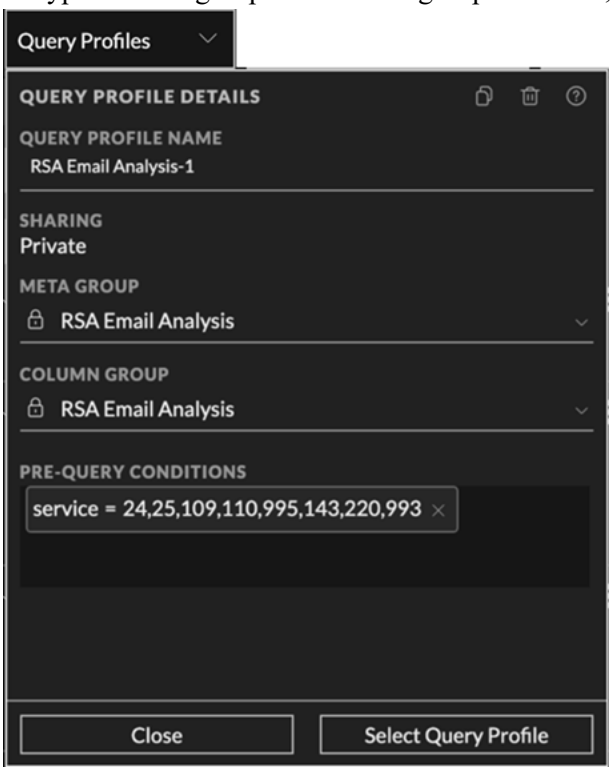


   In version 11.6, the meta groups and column groups are displayed in a drop-down in a folder structure.



b. To edit an existing query profile, highlight a custom query profile in the menu, and click the edit ( ) icon.

The Query Profile Details dialog is displayed. The Version 11.5.1 dialog (on the right) identifies the type of meta group and column group as shared, private, or RSA.



3. In the **Profile Name** field, type a unique profile name that has no more than 80 characters.
   In the Create Query dialog, the Save Query Profile button is activated. In the Query Profile Details dialog, the Select Query Profile button is relabeled as Update Query Profile.
4. (Version 11.5 and later), do one of the following
   a. If you want to share the new query profile with your organization, set the **Share with my organization** option. You cannot change a query profile from shared to private after it is created.
   b. If you want to create a private query profile that only you can see and manage, leave the **Share with my organization** checkbox empty. You cannot change a query profile from private to shared after it is created.
5. (Version 11.5 and later) Select a meta group from the **Meta Group** drop-down list. If a shared group and a private group have the same name, the private group is listed before the shared group.
6. Select a column group from the **Column Group** drop-down list. In Version 11.5, there can be shared or private groups and they can have the same name. In this case, the private group is listed before the shared group.
7. In the **Pre-Query Conditions** field, check the default filters from the query bar and add or remove filters if you wish.
8. Click **Save Query Profile** or **Update Query Profile**.
   The new profile is saved or the edited profile is updated with your changes.
9. To close the dialog, click **Close**.

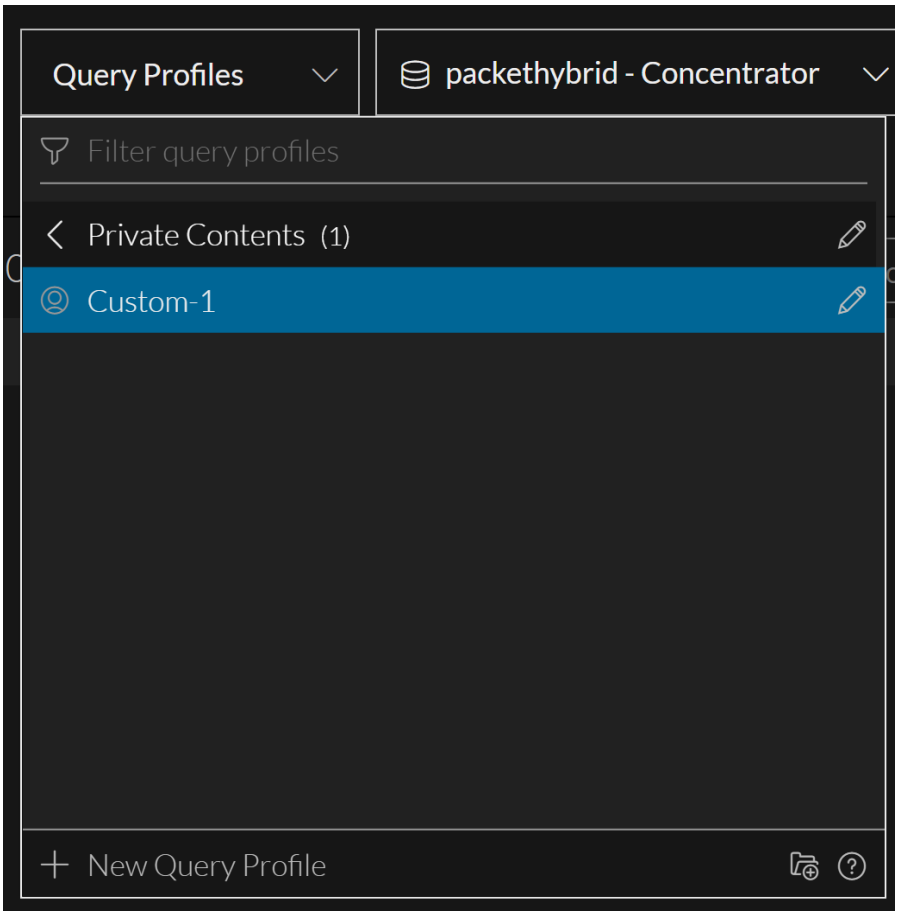
Delete a Custom Query Profile (Events View)

Built-in query profiles are read only, and cannot be deleted, but you can delete any custom query profile. A confirmation message allows you to confirm or cancel the deletion. When you delete a shared query profile, the effect is global and the profile is no longer available to any analyst.

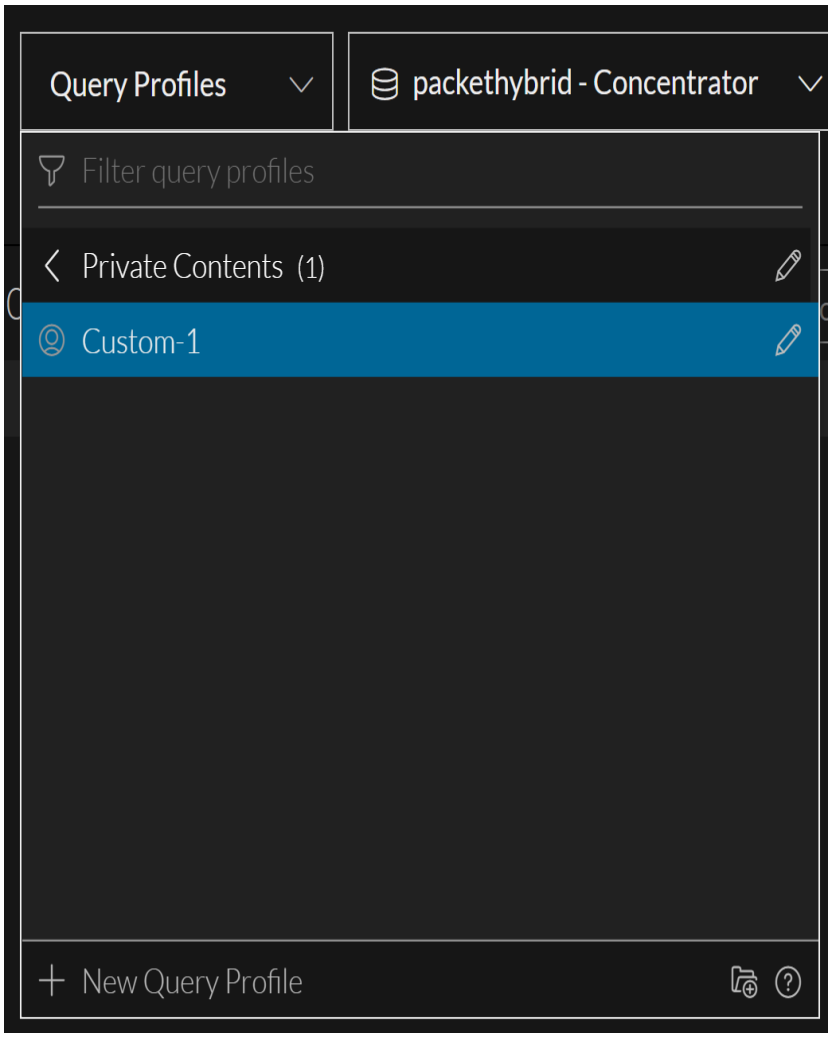
Note: If a Springboard panel is using a query profile as a filter, the profile can be edited, but cannot be deleted in the Events view. However, nothing prevents deletion of the profile in the Navigate view or the Legacy Events view. In this case, Springboard panels that use the deleted query profile as a filter continue to work, but the filter is removed and unexpected results may be displayed in the panel. Refer to "Managing the Springboard" in the *NetWitness Platform Getting Started Guide* for details.

To delete a custom query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles.



2. Highlight a custom query profile that you want to delete, and click the edit ( )icon.
   The Query Profile Details dialog is displayed.



3. Click the delete icon ( ).
   In Version 11.5, a confirmation message gives you the opportunity to confirm or cancel the deletion. Click **Cancel** or **Delete Query Profile**.
   In Version 11.4, if the query profile is not a built-in profile, there is no request for confirmation. The profile is deleted and removed from the Query Profiles menu. The profile no longer appears anywhere for any analyst working in Investigate.

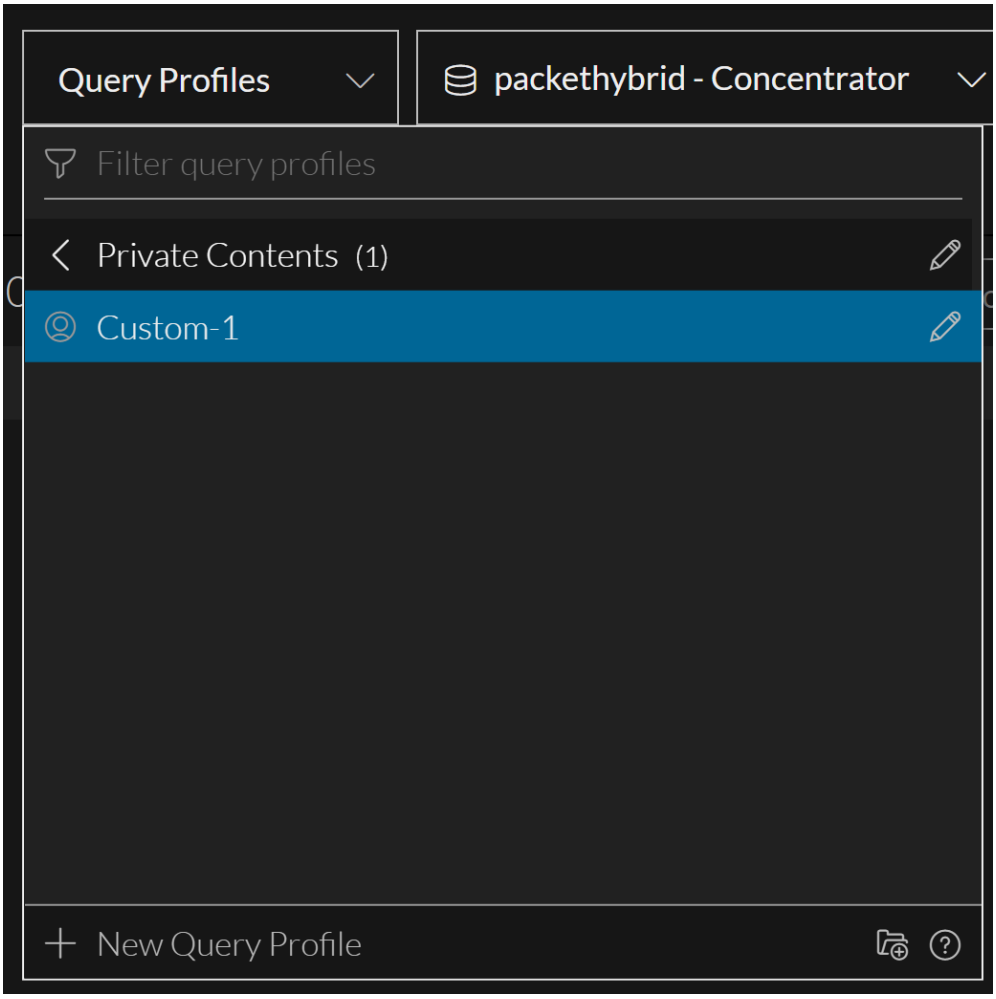
Copy a Query Profile (Version 11.5 and Later)

You can copy any query profile, built-in or custom, shared or private, as long as it does not have unsaved edits in progress. This is useful when you want a customized version of a built-in profile. Also since you cannot change a custom profile from private to shared or from shared to private, creating a copy allows you to select a different Sharing setting. When you copy a profile, the same name is used with a number appended. For example, if you copy RSA Email Analysis, the first copy is named RSA Email Analysis-1, and a second copy of the same profile is named RSA Email Analysis-2. After you create the copy, you can edit the new profile to give it a new name and edit the pre-query conditions, meta group, and column group in the profile.

Note: If you are making a shared copy of a private query profile that uses a private meta group or column group, a message notifies you that a shared copy of the meta group or column group is being created and used in the query profile. It may take a little longer to copy the query profile when a private meta group or column group has to be copied.

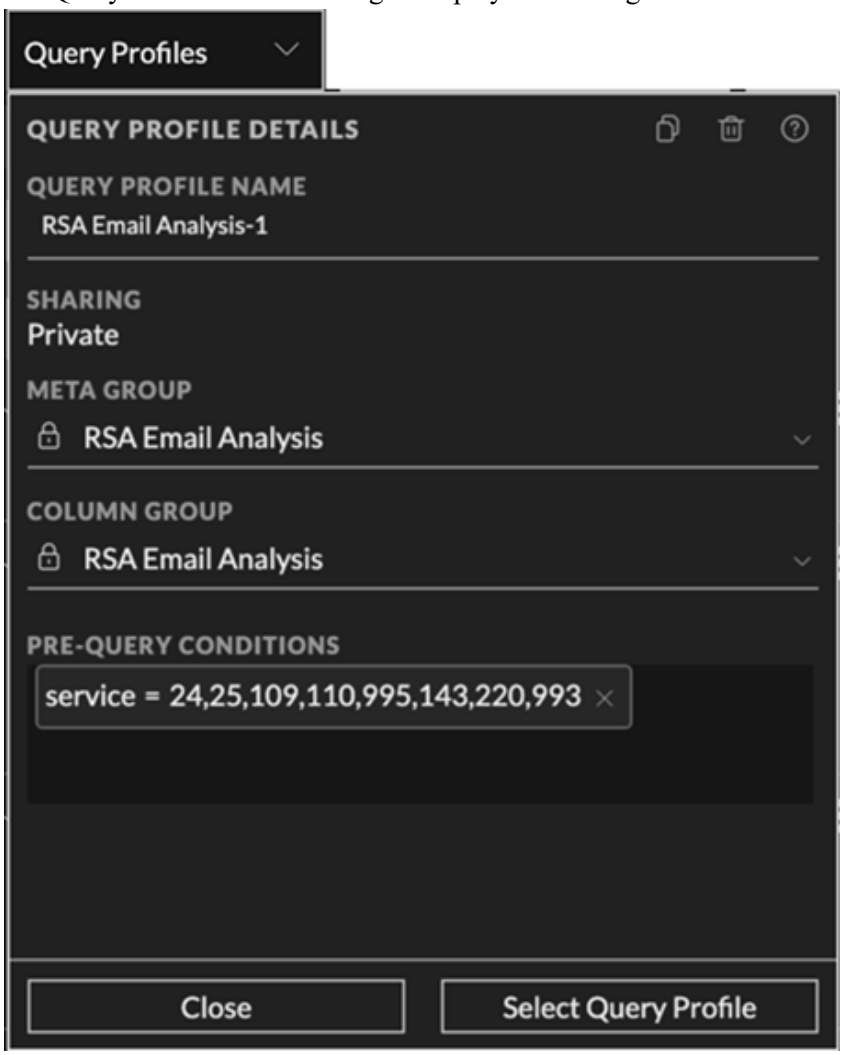
To copy a query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles.
2. Highlight the query profile that you want copy. This figure shows RSA Email Analysis highlighted. The information icon ( ) is displayed to the right.

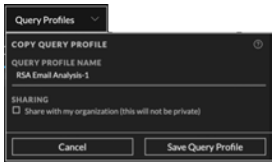


3. Do one of the following:
   a. Click the information icon ( ).
   b. For a custom profile, click the edit icon ( ).

The Query Profile Details dialog is displayed. This figure shows the dialog for a built-in profile.



4. Click the Copy icon ( ).
   The Copy Query Profile dialog is displayed with a number appended to the profile name to create a unique name among all query profiles.



5. (Optional) In the **Query Profile Name** field, edit the name of the query profile.
6. If you want to share the new profile with your organization, set the **Share with my organization** option. By default the new profile is private. If the profile being copied has a private column group or meta group, a shared copy is created and used in the copy of the profile.

7. Do one of the following:
   a. To close the dialog without copying the profile, click **Cancel**.
   b. To save the clone of the query profile, click **Save Query Profile**.
      The clone is saved, and the Query Profile Details dialog for the cloned profile is displayed.
8. Do one of the following:
   a. To close the dialog, click **Close**.
   b. To close the dialog and select the new profile, click **Select Query Profile**.
      The clone is added to the Query Profiles menu.

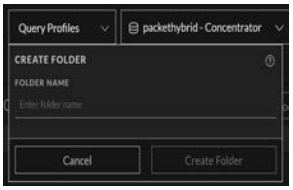
Create a Query Profile Folder

You can create query profiles folders which reside at the top level and are be added as a private or shared folders. And, if the folder name already exists then you are prompted to provide a unique name.

1. In the Events view, select the Query Profiles menu title. The menu drops down to display a list of meta groups and folders with the Filter Query Profiles field at the top and the  option at the bottom.



2. Click .
   The Create Folder dialog is displayed.



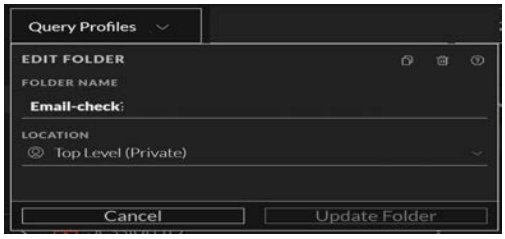
3. In the **Folder Name** field, type a unique name for the new query profile group folder.
4. Click **Create Folder**.

Edit and Move Query Profile Folder

After you create a query profile group folder you can edit or move it, however the folders inside RSA Groups (RSA Live content and RSA OOTB Groups) cannot be edited and moved. The folders inside private and shared folders can be edited and moved only within their respective groups. For example, you cannot move a shared folder into a private folder and vice-versa.

1. In the Events view, select the Query Profiles menu title that you want edit.

2. Click .
   The Edit Folder dialog is displayed.

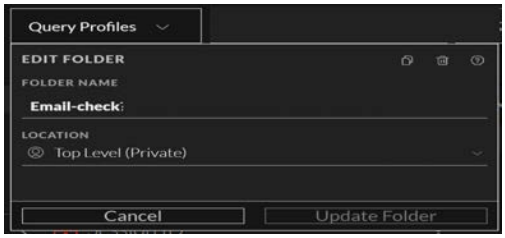


3. In the **Folder Name** field, type a unique name for the query profile folder.
4. Select the location of the folder to be edited.
5. Click **Update Folder**.

Copy Query Profile Folder

You can copy query profiles folder from private to shared, private to private, shared to shared and shared to private groups. When you copy a folder the content inside it gets copied except for the sub-folders. When you copy a private folder into a shared folder, the folder and its content no longer remain private.

1. In the Events view, click the Query Profiles menu title. The menu drops down to display a list of query profiles and folders.
2. Select a folder you want to copy.
3. Click edit and then click .



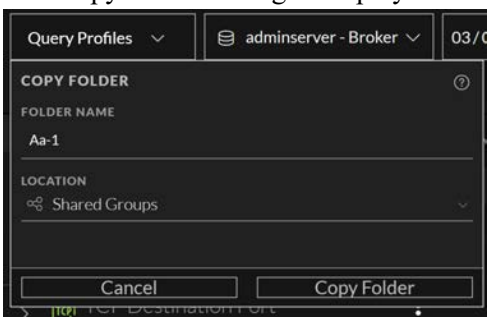
   The Copy Folder dialog is displayed.



4. In the **Folder Name** field, type a unique name (maximum length of 80 characters) for the new query profile group and folder.
5. Select the location of the folder to be edited.
6. Click **Copy Folder**.

Copy Query Profiles Group Folder Deployed from Live

You can copy query profiles group folder deployed from Live located under RSA Groups category to any other location like Shared groups or to a private folder.

1. In the Events view, click Query Profiles Group menu title. The menu drops down to display a list of query profiles groups and folders.
2. Click on a Live Query Profiles Group folder you want to copy.
3. Click 
   The Copy Folder dialog is displayed.
   

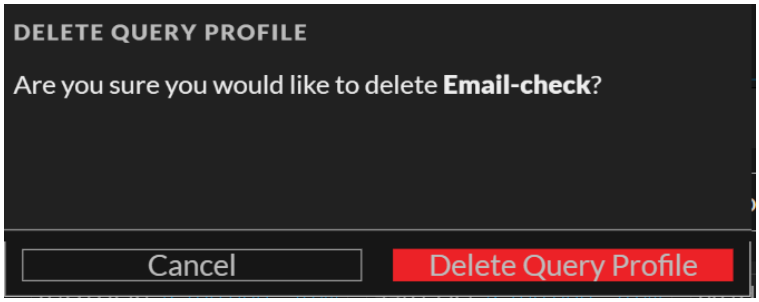
4. Select the location of the folder to be copied.
5. Click **Copy Folder**.
   The folder is created with the original name of the folder and its contents are displayed as the original meta group name appended with a -n.

Delete Query Profile Folder

If you don't want to retain a folder you can delete it. However, once the folder is deleted it cannot be retrieved.

1. In the Events view, click the Query Profile menu title. The menu drops down to display a list of query profile groups and folders.
2. Select a folder to be deleted.
3. Click edit .
   The Edit Folder dialog is displayed.
   

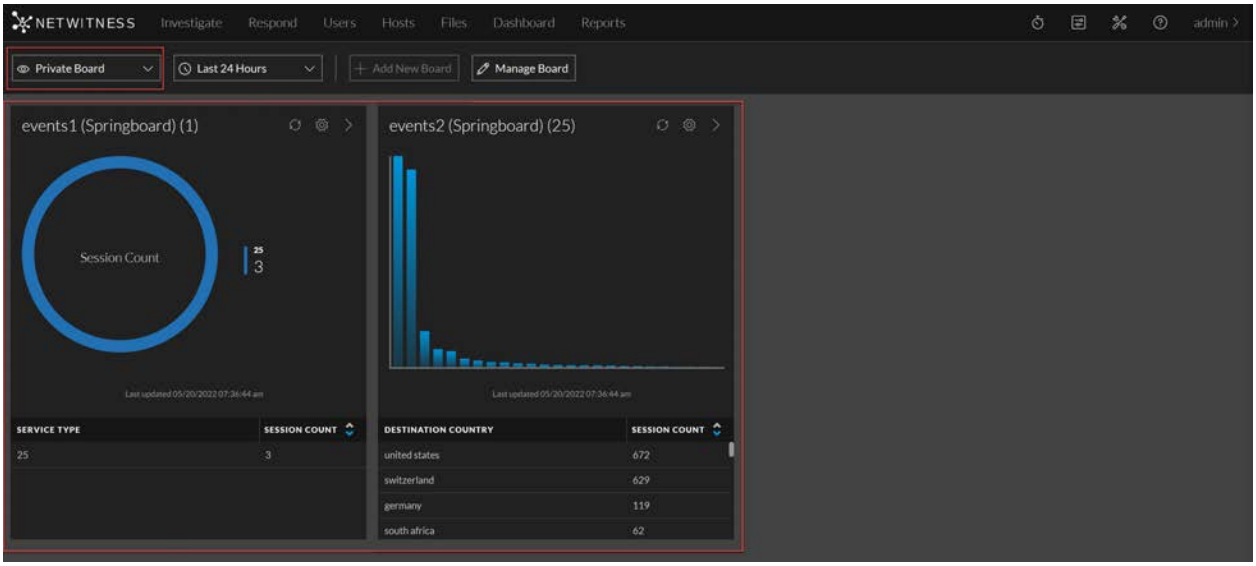
4. Click delete .
   A warning message is displayed to confirm the action.



5. (Optional) Select the checkbox, if you want to delete the folder along with all the contents inside the selected folder.
   If you do not select the checkbox, then the content will be moved to the parent folder after the required folder is deleted.
6. Click **OK** to delete.

Add Springboard Panels from Events View

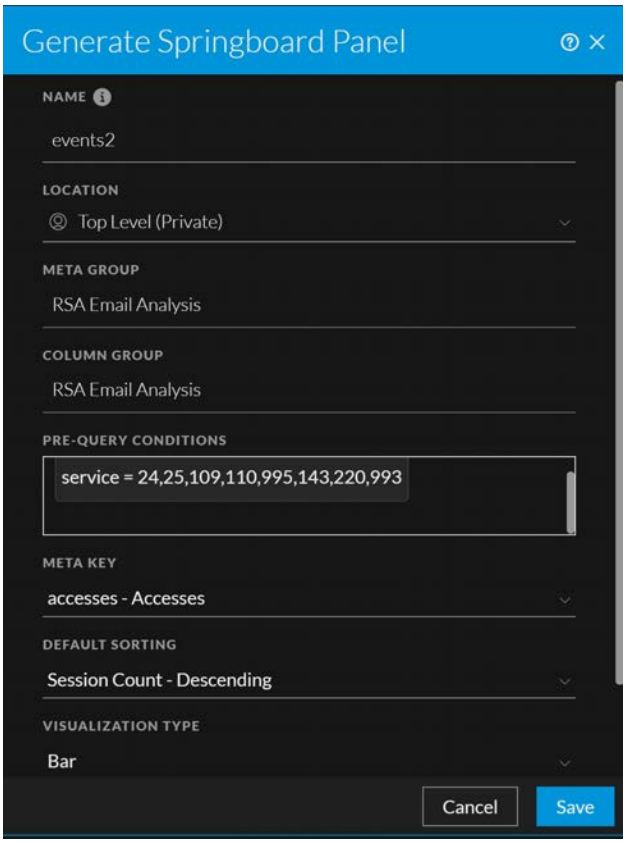
(From 12.0 and later) Administrators and Analysts can now create a Springboard panel from **Investigate** > **Events** view. Analysts can add any number of filters on the query bar and convert them into Springboard panels with important system indicators for threat hunting and investigation.



IMPORTANT: Ensure that you create a custom private board first in order to add the Springboard panel.

To add a Springboard panel from Events view

1. Go to **Investigate** > **Events**.
2. Create a query that consists of one or more filters that contain a meta key, operator, and optional value.
3. Click  > **Generate Springboard Panel**.
   The Generate Springboard Panel dialog is displayed.



4. Enter the following details:
   - **Name**: Enter a unique name for the panel. The name can include letters, numbers, spaces, and special characters, such as _ - ( ) [ ].

     Note: The query profile will be created with the same name as the Springboard panel.
   - **Meta Group**: It is selected by default.
   - **Column Group**: It is selected by default.
   - **Location**: It is the location where the query profile will be saved.
   - **Pre-Query Conditions**: Displayed based on the input criteria entered in the search query panel.
   - **Meta Key**: Select the appropriate meta key value from the drop-down list.

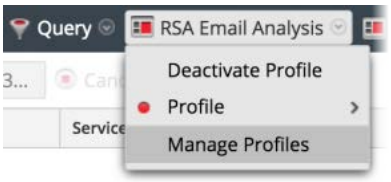
- **Default Sorting**: Select the appropriate sorting from the drop-down list.
- **Visualization Type**: Select the appropriate visualization type from the drop-down list.
- **Visualization Metric**: Select the appropriate visualization metric from the drop-down list.

5. Click **Save**.

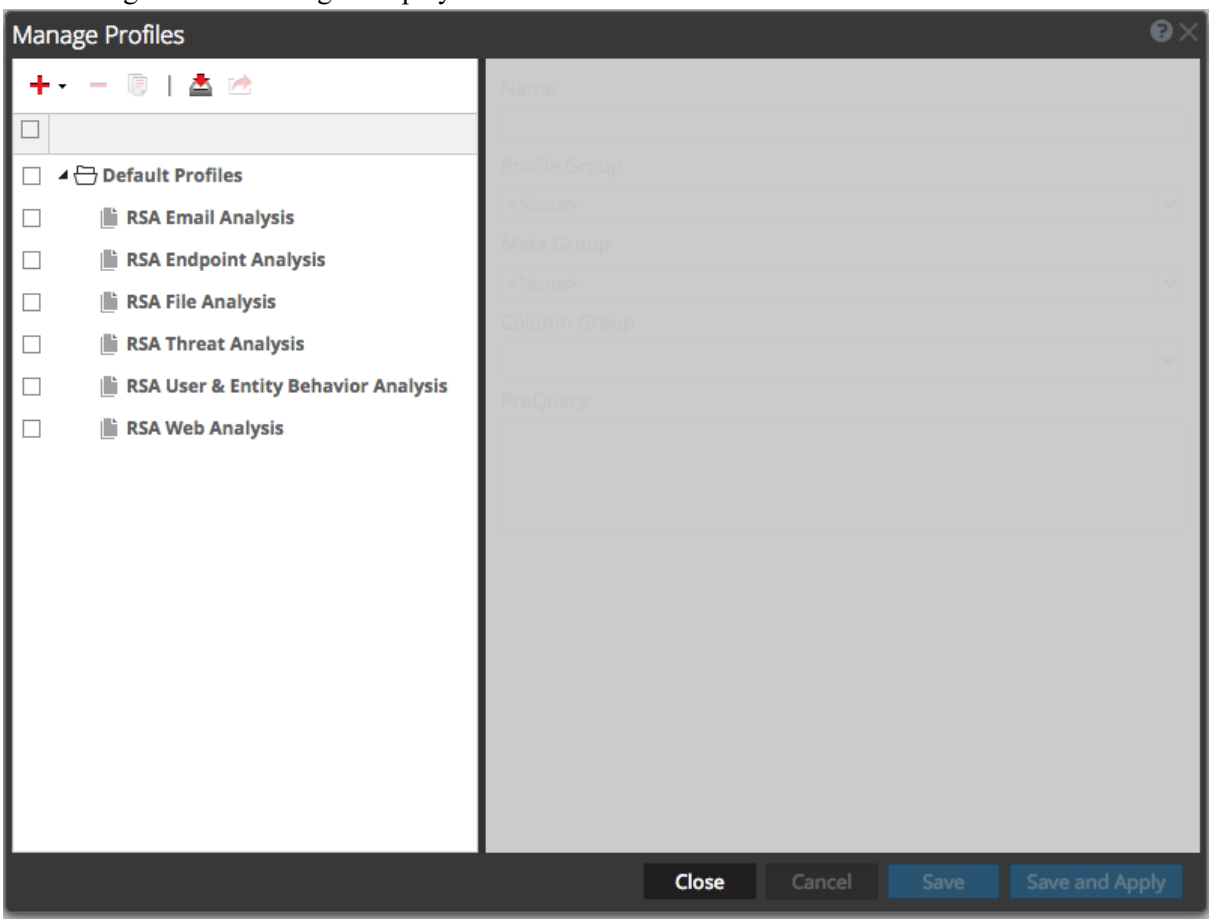
   The panel will be added successfully to the custom private board in the Springboard.

Navigate to the Manage Profiles Dialog (Navigate and Legacy Events Views)

1. Go to **Investigate > Navigate** or **Legacy Events**. (If the **Investigate** dialog is displayed, select a service and click **Navigate**.)
2. In the toolbar, select **Profile > Manage Profiles**.



The Manage Profiles dialog is displayed.



Create, Edit, or Delete a Profile Group (Navigate or Legacy Events View)

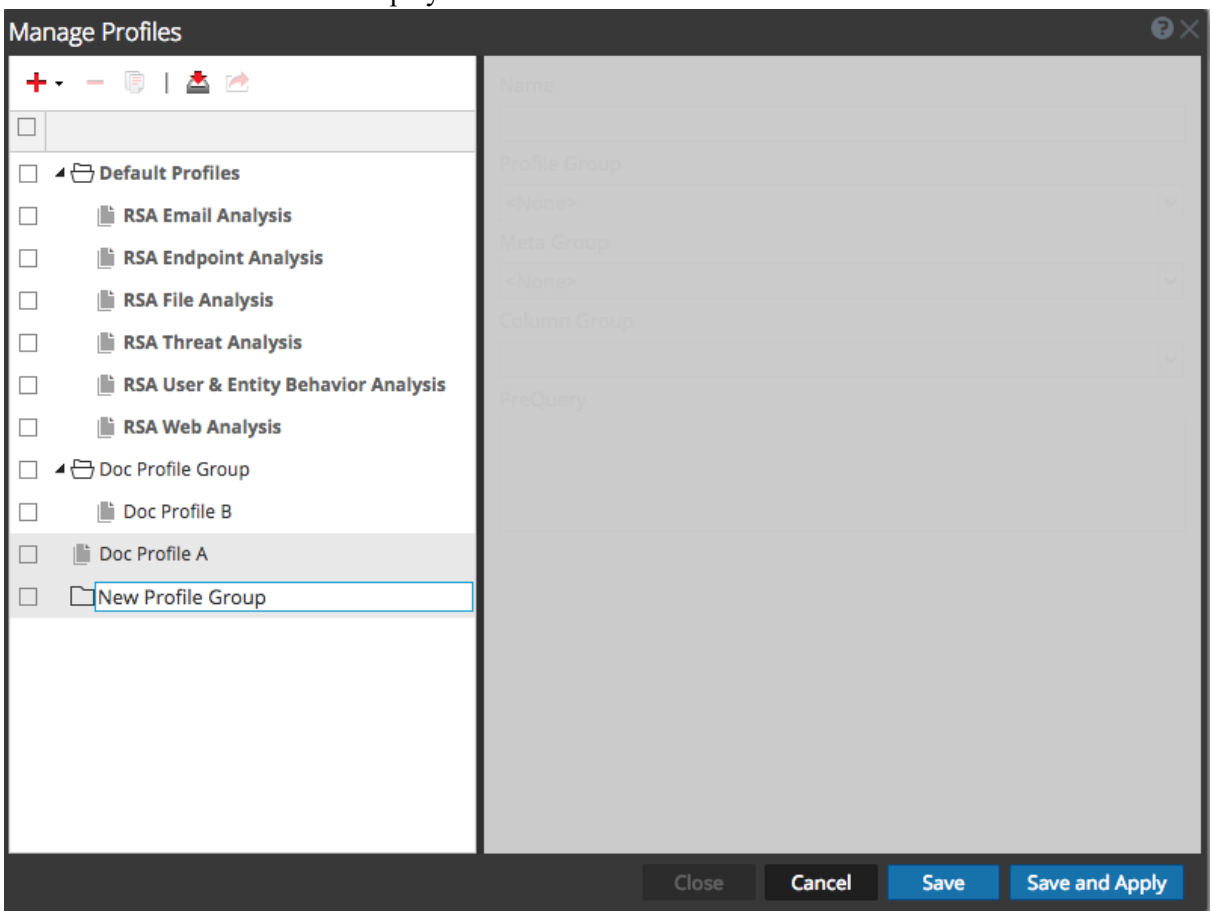
You can create a custom profile group to organize different profiles. Once created, the only edit you can make directly to a profile group is to edit the name of the profile group. To add or remove a profile in a group, edit the profile and assign it to a different profile group as described in Create and Edit Profiles (Navigate or Legacy Events View).

Note: If you migrated profile groups from Version 11.3, empty groups were not migrated.

1. In the **Manage Profiles** dialog, do one of the following:
   - To select an existing profile group to edit, double-click the profile group.
   - To add a new profile group, click  and select **Add New Profile Group**.

Note: If you want to edit one of the built-in profile groups, click  to make an editable copy.

A folder with a blank field is displayed at the bottom of the Profiles list in the left column.



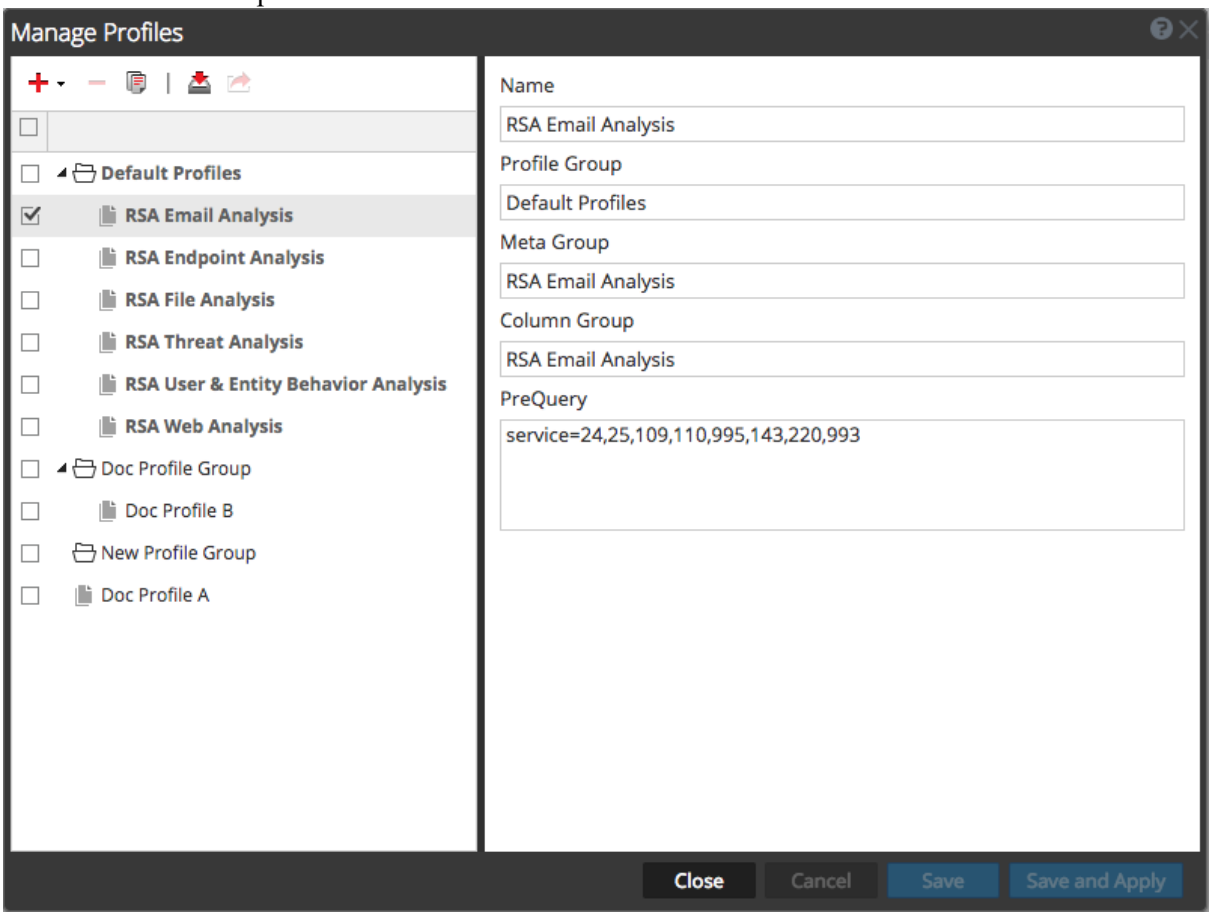
2. To edit or enter the name of the profile group, double-click the Profile Group and type in the entry field. The name must be between 2 and 80 characters.
   The profile group name is applied to a new profile group or to the profile group you edited. The profile group is now available when configuring a profile.
3. To delete a profile group do one of the following:
   - If you want to delete a profile group but keep the profiles, click the checkbox to select the group, uncheck the profiles in the group, and click delete.
   - If you want to delete a profile group and the profiles that the group contains, click the checkbox to select the group, and leave the profiles that you want to delete checked.
     A dialog asks for confirmation that you want to delete the group. If you left the mark in the checkbox next to the profiles, the group and the profiles in the group are deleted. If you unchecked the profiles, only the profile group is deleted and the profiles are moved out of the group and available to add to another profile group.

Create and Edit Profiles (Navigate or Legacy Events View)

1. In the **Manage Profiles** dialog, do one of the following:
   - To select an existing profile to edit, click the checkbox beside the name.
   - To add a new profile in Version 11.2 and later, click + or click the down arrow next to + and select **Add New Profile**.
   - To create a new profile in versions prior to 11.2, click +.

   Note: If you want to edit one of the built-in profiles, click to create a copy, and edit the copy.

   The definition of the profile is available to edit in the right panel. This figure illustrates the definition of one of the built-in profiles.



2. Edit or enter the profile name by typing in the **Name** field. The name must be between 2 and 80 characters.
3. (Optional for Version 11.2 and later) If you want to add the profile to a profile group, select a profile group from the **Profile Group** drop-down list.

If you select a profile group, the profile is added to the group when you save the changes. If you do not select a profile group, the profile is not part of a group.

4. Select a meta group from the **Meta Group** drop-down list. You can add custom meta groups as described in Use Meta Groups to Focus on Relevant Meta Keys. Private meta groups created in the Events view are not available in the Navigate view.
5. Select a column group for the **Column Group** drop-down list. You can add custom column groups as described in Use Columns and Column Groups in the Events List. Private column groups created in the Events view are not available in the Navigate view.
6. Type queries to filter results in the **PreQuery** field. PreQuery follows the same syntax as the Query builder. The PreQuery in the figure uses a meta group called **service = 24,25,109,110,995,143,220,993**.
7. Click **Save** to save the profile without using it, or click **Save and Apply** to save the profile and use it immediately.
   If you click **Save and Apply**, a confirmation dialog is displayed before applying the selected profile. For Version 11.2 and later, the PreQuery that you entered in the Manage Profiles dialog is displayed in the breadcrumb.



Delete a Profile (Navigate or Legacy Events View)

1. In the **Manage Profiles** dialog, select a profile by clicking the checkbox beside the name.

   Note: You cannot delete any of the built-in profiles.

2. Click −.
   A prompt requests confirmation that you want to delete the profile, and the profile is deleted. The option name in the toolbar reverts to **Profile** to show that no profile is in effect.

Change the Active Profile (Navigate or Legacy Events View)

If you do not see enough results or the right results in the Navigate or Events views, you may have an active profile that is applying a PreQuery. If you do not want to use any profiles, you can click **Deactivate Profile** in the **Profile** drop-down menu.

To use a different profile:

1. In the **Navigate** or **Legacy Events** view toolbar, open the **Profiles** drop-down menu.
2. Hover over the **Profile** option to display a drop-down list of available profiles.
3. Select the profile you want to use.
   The profile settings are applied immediately.

If you want to change the active profile from the Manage Profile dialog:

1. In the **Navigate** or **Legacy Events** view toolbar, select **Profiles > Manage Profiles**.
   The Manage Profiles dialog is displayed.
2. Select a profile from the left panel and click **Save and Apply**.
   A confirmation dialog is displayed.
3. Click **Yes**.
   The profile settings are applied immediately.

Import Profiles (Navigate or Legacy Events View)

In the Navigate view and the Legacy Events view, you can upload or import .jsn files that have been downloaded from another service. When profile groups are exported and then imported, the grouping of profiles is maintained.

1. In the **Manage Profiles** dialog, click  in the left panel toolbar.
   The Profile Import dialog is displayed.
2. Click **Browse** or the **Upload File** field to select a file from your computer.
3. When the file is selected, click **Upload**.
   The profile is displayed in the left panel.

Download Profiles (Navigate or Legacy Events View)

In the Navigate view and the Legacy Events view, profiles are downloaded as .jsn files.

1. In the **Manage Profiles** dialog, select one or more profiles from the left panel.
2. In the left panel toolbar, click .
   The download begins immediately.

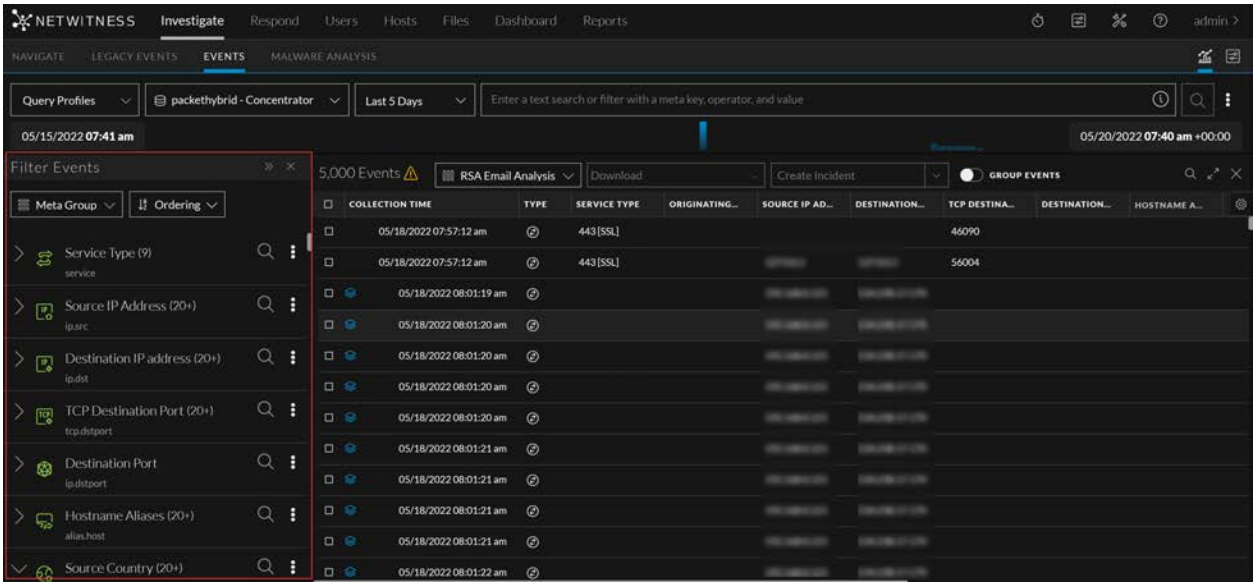
Drill into Metadata in the Events View

Note: This section applies to Version 11.5 and later. The feature is a beta feature that is enabled by default, and can be disabled by the system administrator as described in the *System Security and User Management Guide*.

When working in the Events view, the focus of an investigation is the smallest possible set of relevant events in sequential order. You can reduce the number of visible events loaded in the Events view using query profiles, column groups, meta groups, and queries. However, it is more efficient to limit the data set using the metadata indexed on the Concentrator before looking at the actual events stored on the Decoder or Log Decoder.

In Version 11.4.x and earlier, it is best to start by looking at the meta keys and meta values indexed on the Concentrator and drill into the metadata in the Navigate view to find a relevant set of events, with each drill or query further limiting the data set. When you have a meaningful data set, or drill point, you can examine the details of the related events in sequential order in the Events view.

Beginning with Version 11.5, you can drill into the metadata in the Filter Events panel, without leaving the Events view. The list of meta keys and meta values shown is related to all events seen in the environment for the time range in the query. When you find the drill point of interest in the Filter Events panel, you can open the Events panel to see the sequential events. The set of events loaded in the Events view is smaller and loads faster. The flow of an investigation is smoother with less hopping between views. The following figure illustrates the panel, open to the left of the Events panel.



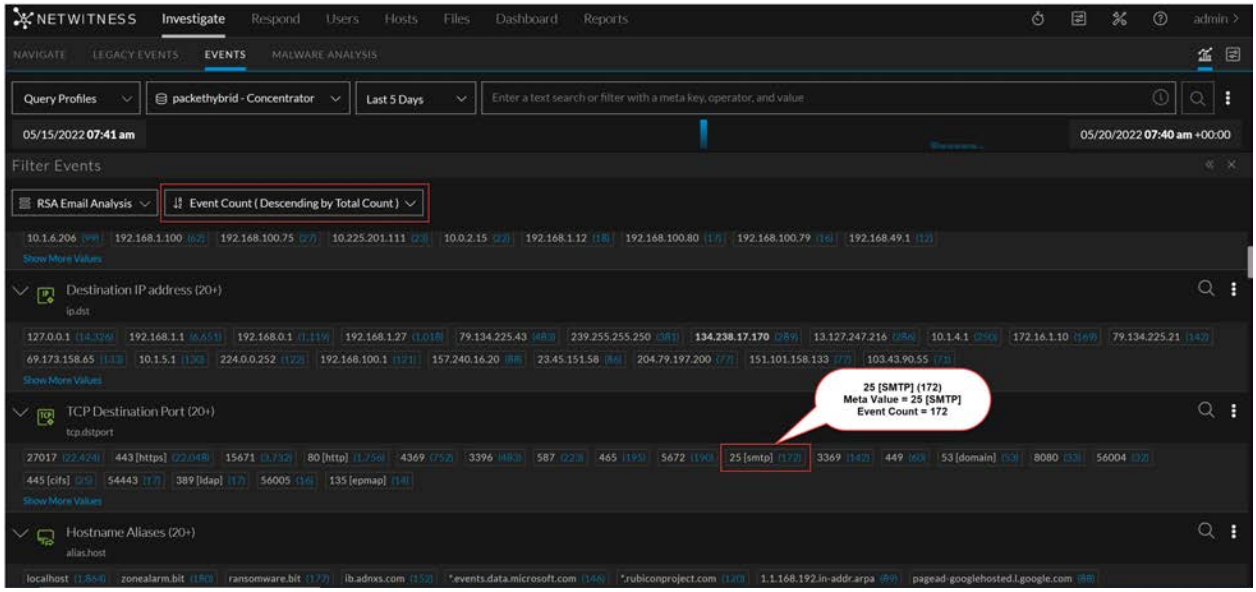
Note: There are two situations in which results in the Filter Events panel may not be as expected:
-In a mixed-mode environment with a Version 11.5 Broker and some Core services at NetWitness Platform Version 11.4 or earlier, a text filter is not supported in the Filter Events panel. If the query in the Events panel includes a text filter, the result set in the Events panel and Filter Events panel may be different.
-If the query in the Events view query builder has a logical `OR` or `&&`, the results in the Events view may be different from results for the same query in the Navigate view and Legacy Events view. In this situation, a set of parentheses automatically encloses the logical `OR` expression in the Navigate view and Legacy Events view, while parentheses have to be manually added in the Events view. If this occurs, you need to enclose the logical OR expression in an additional set of parentheses; select the two filters in the query bar, right-click one of them, and select **Wrap in parentheses** in the menu.

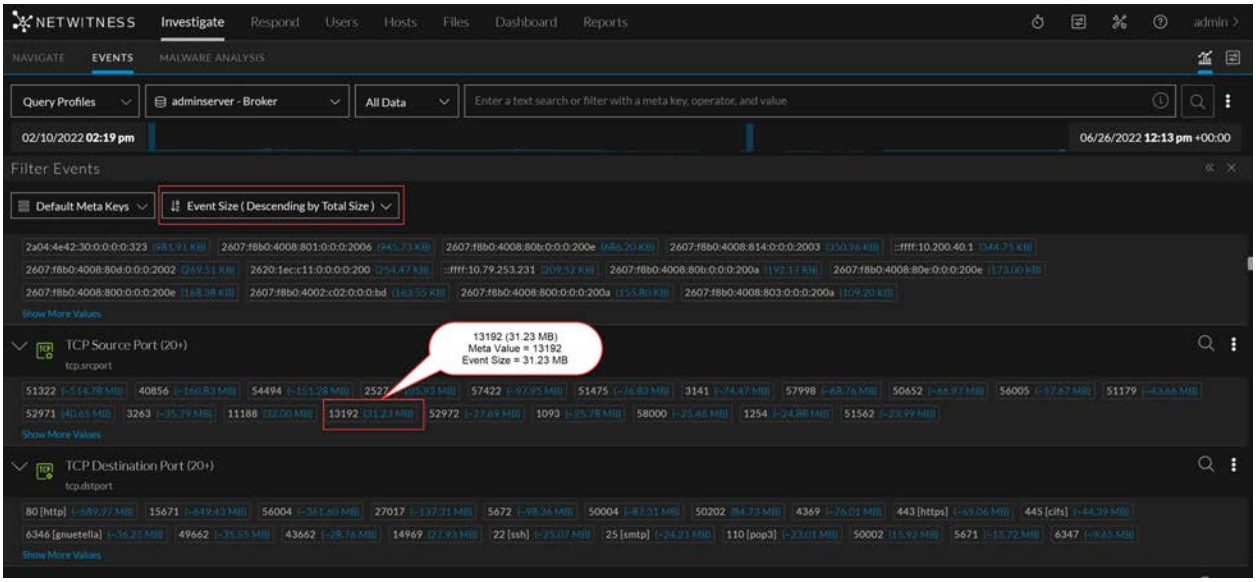
Note: (Version 11.5.1) In the Filter Events panel, the meta values result threshold is 100000. If results are above the threshold, it is indicated using either ~ or >. For example, **(>100000)** indicates that the results are sorted based on count and are greater than the threshold. Similarly, **(~100000)** indicates that the results are sorted based on size and are greater than the threshold.

Modes of Operation

The Filter Events panel has two modes of operation.

- The **narrowFilter Events panel** is part of a faceted search view into the data (shown above). Left- or right-clicking a meta value adds a new filter, automatically executes a new query, and displays matching events in the sequential list of events. When both panels are open, you can drill into the data in both the Filter Events panel and the Events panel . Each time you left-click a meta value in the Filter Events panel, an expression is appended to the query bar, and the query is executed by default. The query results show new metadata to filter by in the Filter Events panel and the resulting events that match the query in the Events panel. If you change the service or other query elements in the Events panel, you need to execute the query to reload the Filter Events panel.
- The **fully expanded Filter Events panel** uses the full width of the browser window to provide ample real estate to hunt through the metadata without the performance load of immediately submitting a query or viewing the sequential events. As you click a new meta value and drill into the meta values, each meta value is added to the query filter and executed in the Filter Events panel, so that the number of events seen is reduced. Because the Events panel is closed, the query in the Events panel is not updated and the query is not executed. When you collapse the Filter Events panel back to original size, the Events list opens and the query is executed. This is an example of the fully expanded panel.



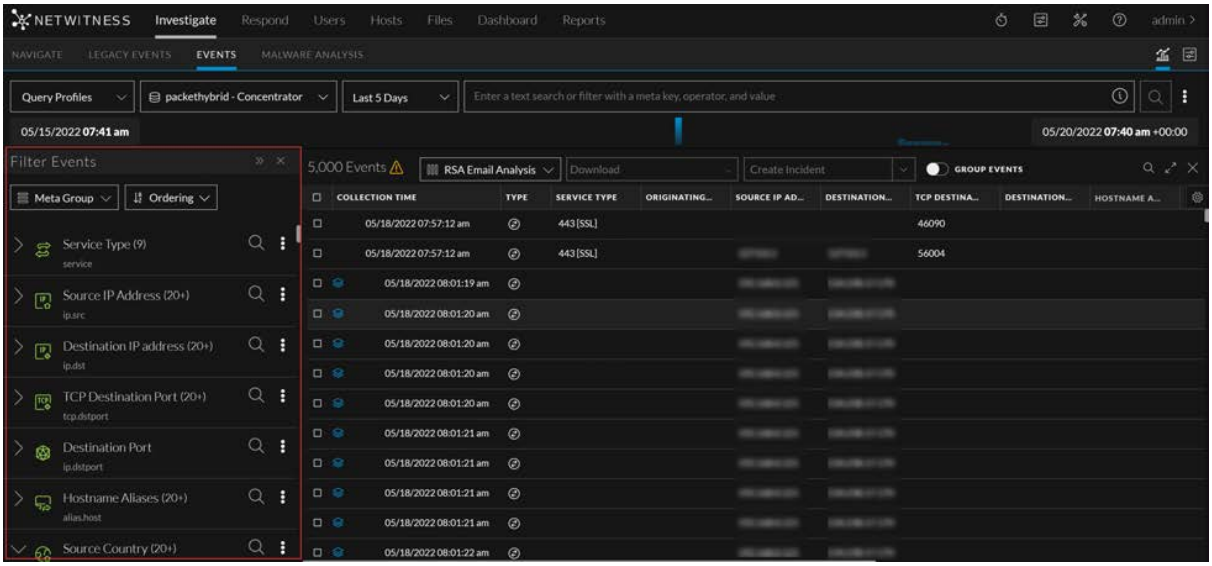


View Metadata in the Filter Events Panel

To view metadata in the Filter Events panel:

1. Go to **Investigate > Events**, select a service to investigate, and select a time range.
2. (Optional) Select a column group or a query profile.
3. Click  to load events in the Events panel.
   A query is executed in the Events panel and matching events are listed,

4. Click the Filter button (Filter) in the Events panel.
   The Filter Events panel opens to the left of the Events panel.



Note: (Version 11.6) By default, the Filter Events panel is open in the Events view. The last used state of the panel (narrow or fully expanded) is saved throughout the session and across logins. Also, the Filter Events panel provides additional contrast between meta keys, meta values, and meta counts to improve readability.

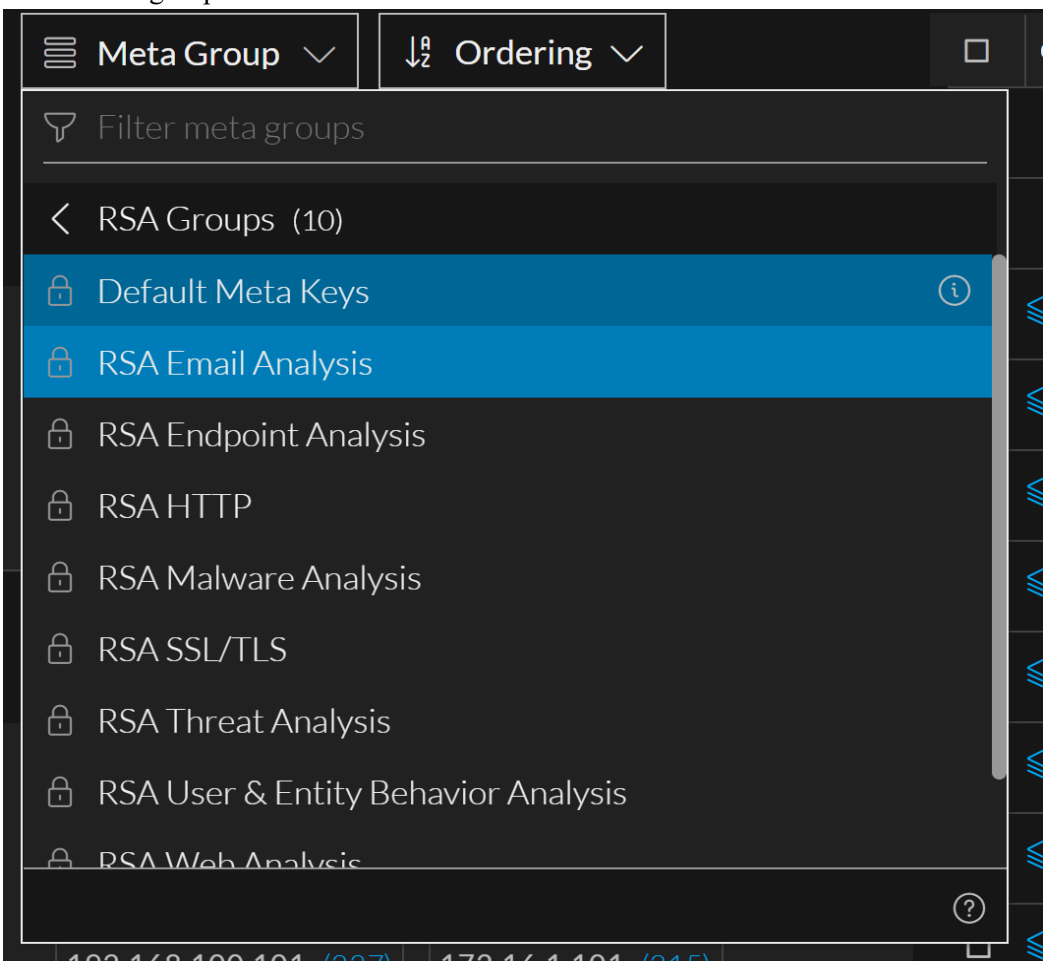
The Default Meta Keys meta group is in effect the first time you log in. If you selected a different meta group the last time you logged in, it remains in effect until browser cache is cleared. In Version 11.5.1, the meta group you selected previously is not stored in browser cache so it remains in effect until you change it. See Use Meta Groups to Focus on Relevant Meta Keys for details about meta groups. Based on the contents of the index file for the service, the Filter Events panel is populated with the first 25 meta keys that have at least one meta value and are open. When using the Default Meta Keys group in the Filter Events panel, only the first 30 meta keys with values are open and the remaining are closed. Closed meta keys may be listed, but they do not count toward the 25 or 30 meta keys total. Meta keys with no values are listed at the bottom of the panel. You can expand, collapse, and close the panel using the standard panel controls (, , and).

1. Do one of the following:

   a. To close the dialog without editing, click **Close**.

   b. To close the dialog and select the copy of the meta group, click **Select Meta Group**.
      The group is added to the Meta Group menu. The figure below has a private copy of the RSA

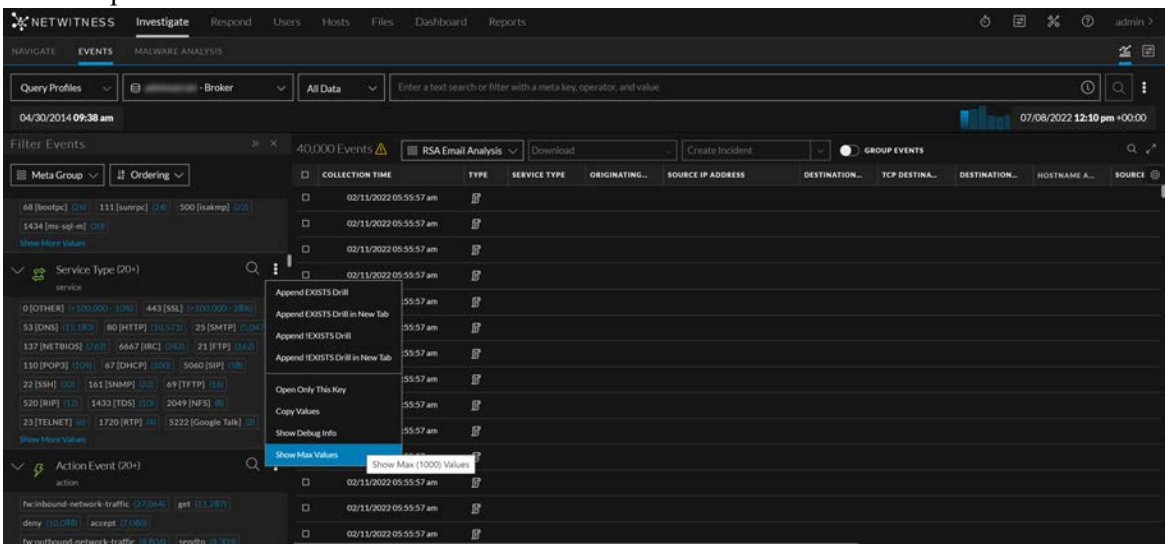
HTTP meta group.



Show Max Value of Meta Groups

In case all the values have not rendered and displayed, you can click Show Max Value to view all the values at once.

1. With the Filter Events panel open in the 11.6 Events view, click  and select the **Show Max Values** option.



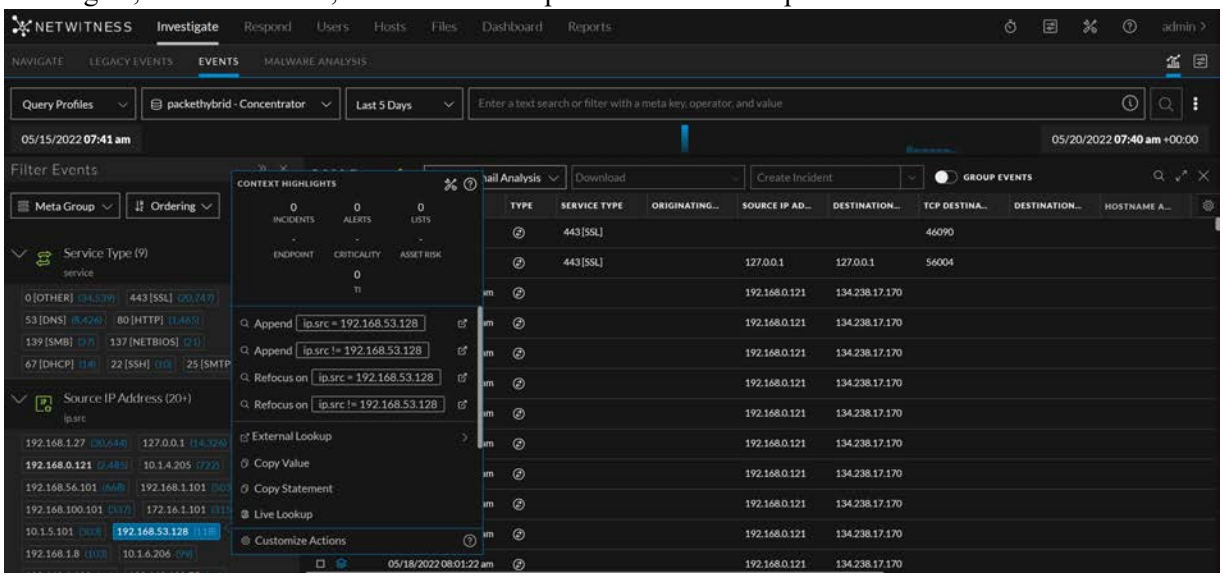
2. The values that were not rendered earlier will begin to load and a maximum of 1000 results are displayed.

View the Context Lookup Panel in the Filter Events Panel

In the Filter Events panel, you can click a meta entity to open the context tooltip. The context tooltip is available only for the meta keys that are defined as an entity the Context Hub supports. The Context Hub service is pre-configured with default meta types and meta keys mapping. For information about mapping of the context hub meta values with investigation meta keys, see "Configure Meta Type Mapping for Context Hub" in the *Context Hub Configuration Guide*. The context tooltip includes the following two sections.

- **Context Highlights** - The information in this section helps you to determine the actions that you would like to take. It can show related data for Incidents, Alerts, Lists, Endpoint, Criticality, Asset Risk, and Threat Intelligence (TI). Depending on your data, you may be able to click these items for more information.

- **Actions** - This section lists the available actions. In the example, the Add/Remove from List, Pivot to Investigate, Pivot to Archer, and Pivot to Endpoint Thick Client options are available.



Understand Visible Metadata

Each meta key has a list of meta values, with up to 20 values displayed by default. You can click **Show More Values** to incrementally add 20 meta values, up to a total of 1,000 meta values, which is a hard-coded limit to optimize performance. The meta key name and plain English name of each meta key found in the service, both populated and non-populated, are listed. For each meta value, you can see the number of events in the current results that contain the value (count) or the size of the events in the current results (size). For example, the following might be listed:

```
Action Event [action] (3)
get(3016) login (1346) put (501)
```

In this example, the meta key name is `action`, the English name is Action Event, and three meta values were found for this meta key. There were 3016 events containing `get`, 1346 events containing `login`, and 501 events containing `put`. The values are ordered so that the value with the largest count is listed first.

In the following example, the same meta key has the values ordered based on the event size in bytes. The smallest size is listed first:

```
Action Event [action] (3)
login (13,034,588) put (21,848,760) get (1,409,079,256)
```

An icon before each meta key name identifies the indexing method for the key. The indexing method determines the types of interactions and queries possible using that meta key.

- This meta key is indexed by value: Action Event [action] (40+). The green color indicates that the all available interactions and queries are supported. You can see the available interactions in the context menu by right-clicking the meta value.
- This meta key is indexed by meta key: Bytes Sent [bytes.src]. The yellow color is a clue that a subset of available interactions is supported, and queries on this meta key may take longer than meta keys that

are indexed by value. You can see the available interactions in the context menu by right-clicking the meta value.

- This meta key is not indexed: 

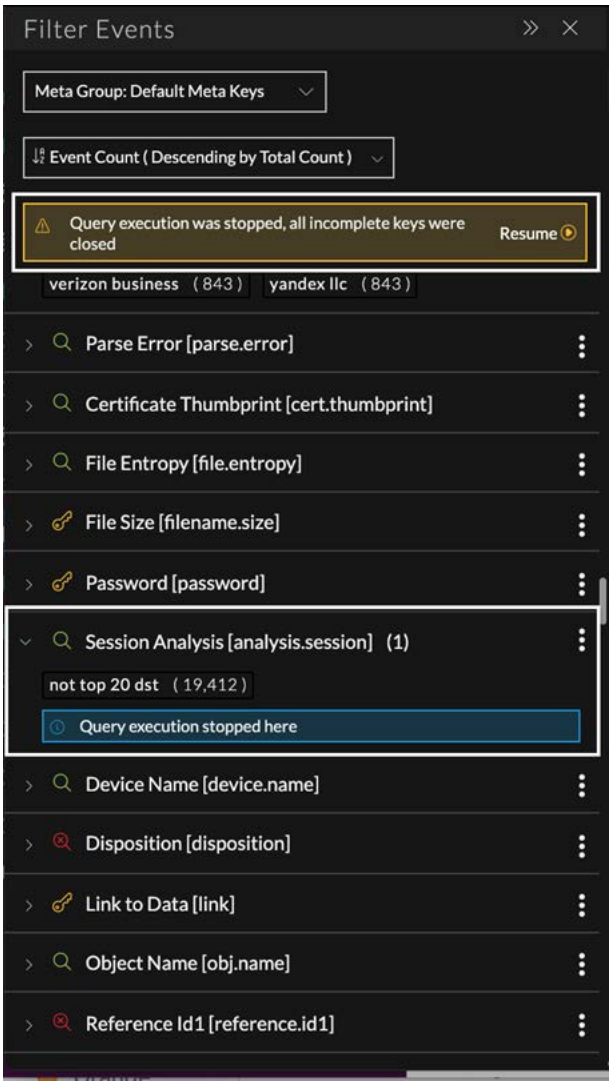
. Values for non-indexed meta keys cannot be used to query. If you want to query a meta key that is not indexed, your administrator needs to edit the index file for the service to index the meta key by value or meta key.
- For Version 11.5.1, a set of more than 200 meta key symbols replaces the three indexing method symbols to provide a visual indicator of the purpose of the meta key. The color of the meta key symbol identifies the indexing method using the same colors as before: green, yellow, and red. A tooltip also identifies the indexing method and provides a description of the icon. The icons are defined based on categories outlined in the Unified Data Model (https://community.netwitness.com/t5/netwitness-platform-unified-data/tkb-p/netwitness-udm). There is a generic icon for most categories that do not have specific meta keys and a default meta key icon to use when a new custom meta key is added.

If an error occurs while loading a meta key, the other meta keys load as usual and an error message is displayed in the meta key that did not load. When you execute a new query, some error messages disappear. Meta keys that have no values in the set of events are listed at the bottom of the panel.

Stop and Resume Metadata Loading

You can stop and resume loading of metadata in the Version 11.5.1 Filter Events panel as the meta keys and values are loading. When loading a lot of metadata this can save time because you do not have to wait for all data to load. If you stop loading and still need to see more metadata, you can resume loading, then stop again when you see the data you want.

1. While data is loading in the Filter Events panel, click the Stop button () in the query bar.
   Meta keys stop loading, and keys that did not finish loading are closed. A message above the meta keys list informs about the status, and you can scroll down to find the last meta key to finish loading. In this example, Session Analysis finished loading; all the meta keys below remain closed.



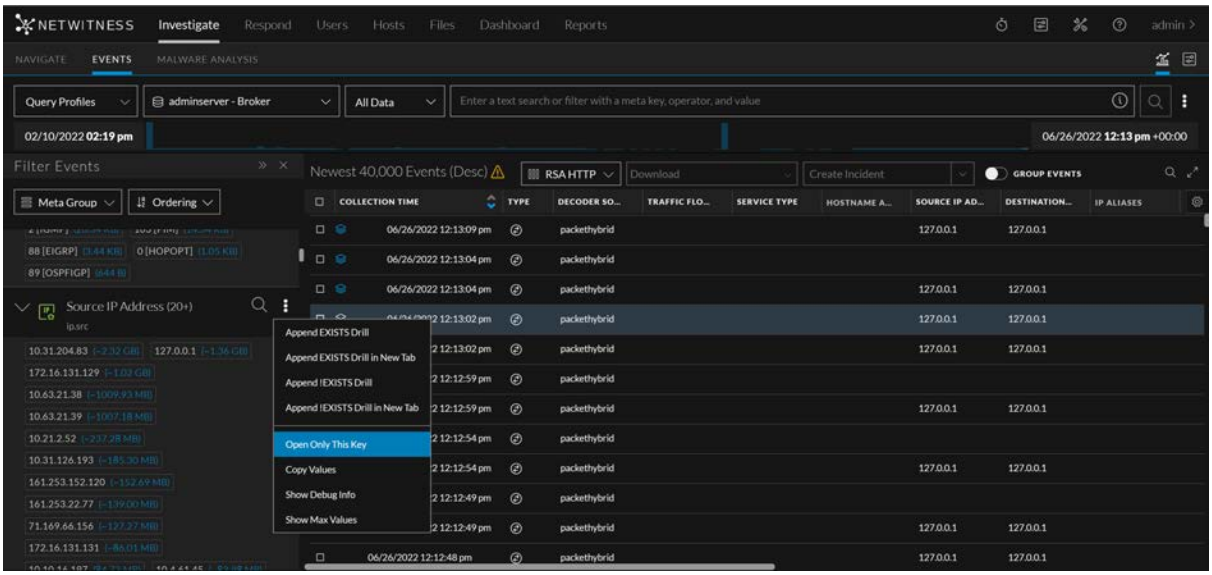
2. Do one of the following:
   a. To resume loading at the meta key where loading stopped, click **Resume**.
   b. If you want to review the values of specific meta keys that were not loaded, without resuming the query, click the meta key name to open any key.

Close All Except One Meta Key

(Version 11.5.1 and Later) When the Filter Events panel is open, seeing many meta keys at the same time can be distracting.

To close all except one meta key:

1. In the meta key row of an entry, click the Meta Key options button (⋮).
   The Meta Key options are displayed.



2. Select Open Only This Key.
   All except the current meta key close. If the selected key is closed, it opens and data is loaded, while all other keys close.

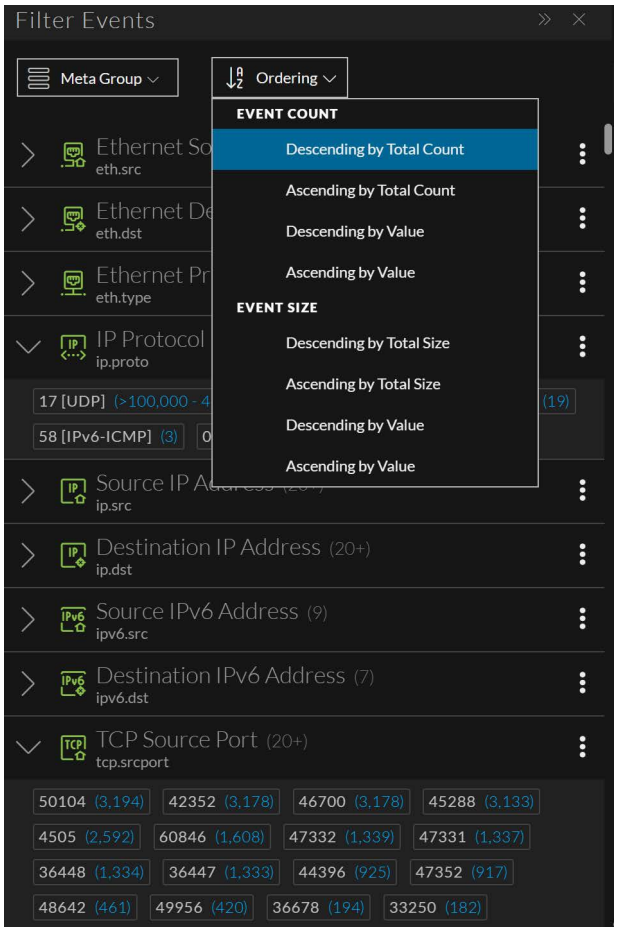
Set the Ordering Method for Meta Values

With the Filter Events panel open, you can look at two parameters for each value: the event count or the event size. Each meta key entry includes either the event count or the event size in parentheses after the value. In both cases, there are four options for ordering.

To use the ordering options:

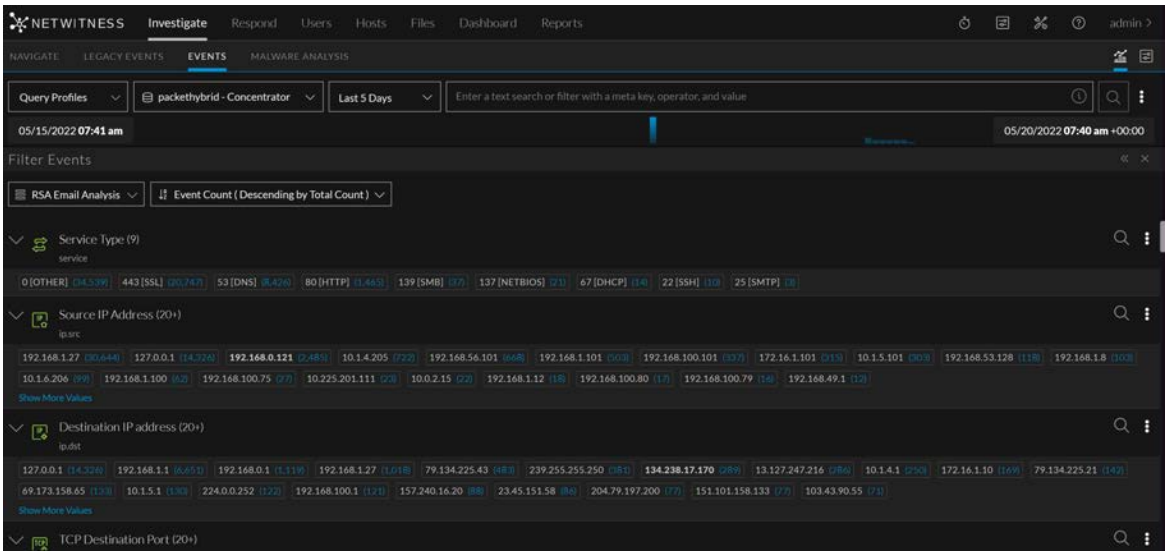
1. With the Filter Events panel open, click the ordering menu label, which is named according to the selected ordering option. This is an example of the menu label when ordering by event count in ascending order by total count: 

   .
   The Ordering menu is displayed. This figure shows the narrow version of the menu.



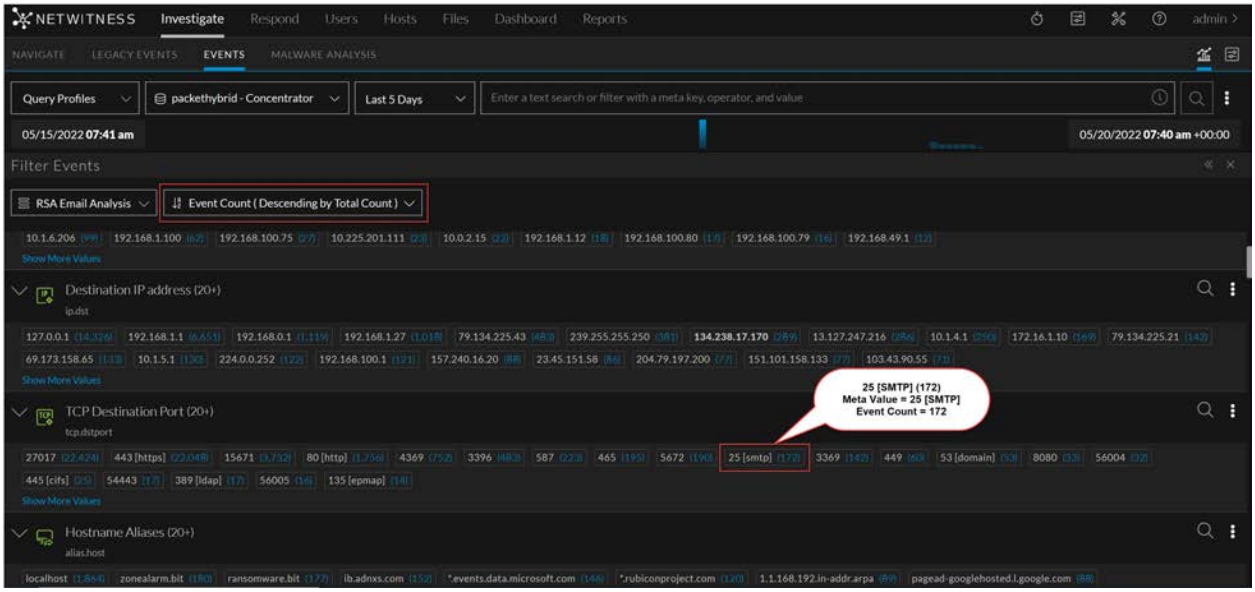
2. If you want to see the event count in parentheses after each value, select one of the following options. By default, the meta keys are displayed using the Event Count > Descending by Total Count method.
   a. To order by total count of events in which the value was found, select either **Descending by Total Count** or **Ascending by Total Count**.
   b. To order by the name of the value, select either **Ascending by Value** or **Descending by Value**.
3. If you want to see the size in bytes of the events in which the value was found, select one of the following options.
   a. To order by total size of events in which the value was found, select either **Descending by Total Size** or **Ascending by Total Size**.
   b. To order by the name of the value, select either **Ascending by Total Size** or **Descending by Total Size**.
   Under each meta key in the Filter Events panel, the values are ordered according to your

selection.

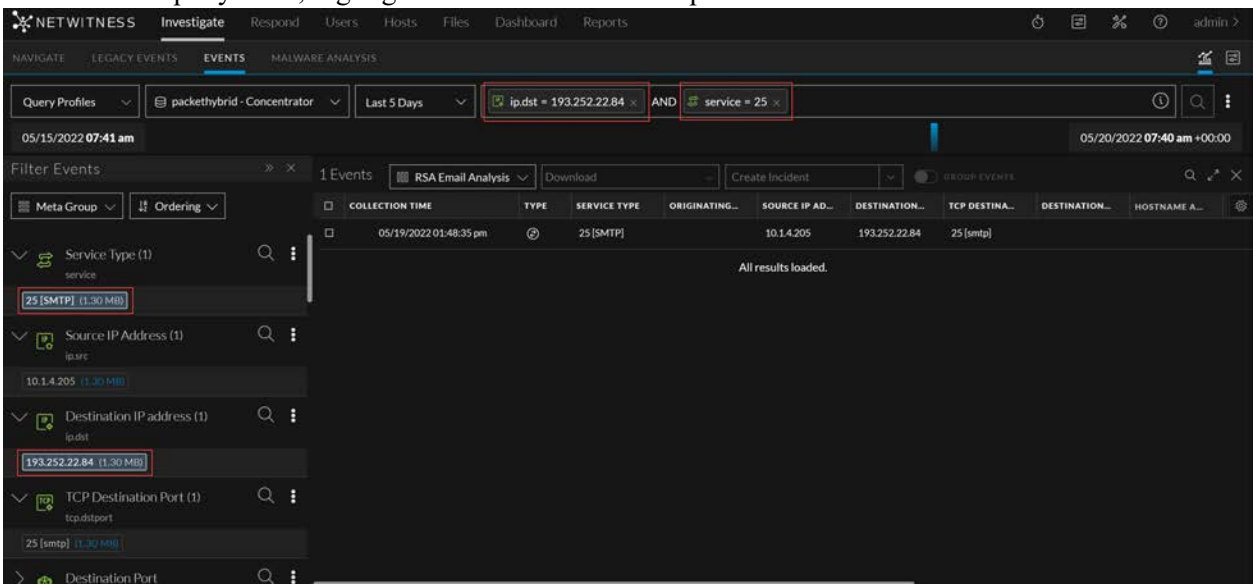


Drill into Meta Values

With the Filter Events panel open, you can drill into meta values to focus an investigation down to the smallest possible set of relevant events. Drilling in the fully expanded Filter Events panel adds filters to the query bar and refines the displayed metadata in the Filter Events panel, but does not execute the query in the Events panel. Drilling in the narrow panel, side by side with the Events panel, adds the filter to the query bar and executes the query in the Events panel and the Filter Events panel. This figure is an example of the fully expanded panel with some metadata loaded.



You can drill into metadata in the Filter Events panel to find relevant meta values. A simple query using the (=) operator highlights the meta value used in the Filter Events panel. This helps to associate the metadata with the filter added to the query. For example, the following figure shows the meta key value, related to the query filter, highlighted in the Filter Events panel.



To drill into meta values in the fully expanded Filter Events panel:

1. Look for a meta value that is of interest, and click the value. Using the figure above as an example, to investigate the SMTP service type as opposed to other service types, click **25[SMTP]**.
   The other service types are filtered out of the metadata in the Filter Events panel, but the query is not executed in the Events panel.
2. Look for a meta value that is of interest, and do one of the following:
   a. Click the value. Using the figure above as an example, to investigate the SMTP service type as opposed to other service types, click **25[SMTP]**.
      The filter is added as the last filter in the query bar, and other service types are filtered out of the metadata in the Filter Events panel. With the Events panel closed, no query is executed there.
   b. (Version 11.5.1) Right-click the value and select **Add Filter - Do Not Run Query** in the drop-down menu.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata in the Filter Events panel. With the Events panel closed, no query is executed there.
   c. (Version 11.5.1) Press **CTRL (Windows)** or **CMD (MacOS)** and click the value.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata in the Filter Events panel. With the Events panel closed, no query is executed there.
3. Repeat step 1 with another meta value, for example, **writetoexecutable** in the **Action Event [action]** meta key. Continue drilling into values until you find a set of events (drill point) that you want to see in sequential order.

4. To view the sequential events for the drill point, click  to shrink the Filter Events panel.
   The Events panel opens to the right, and the query is executed in the Events panel so that you can see the raw events in sequential order.

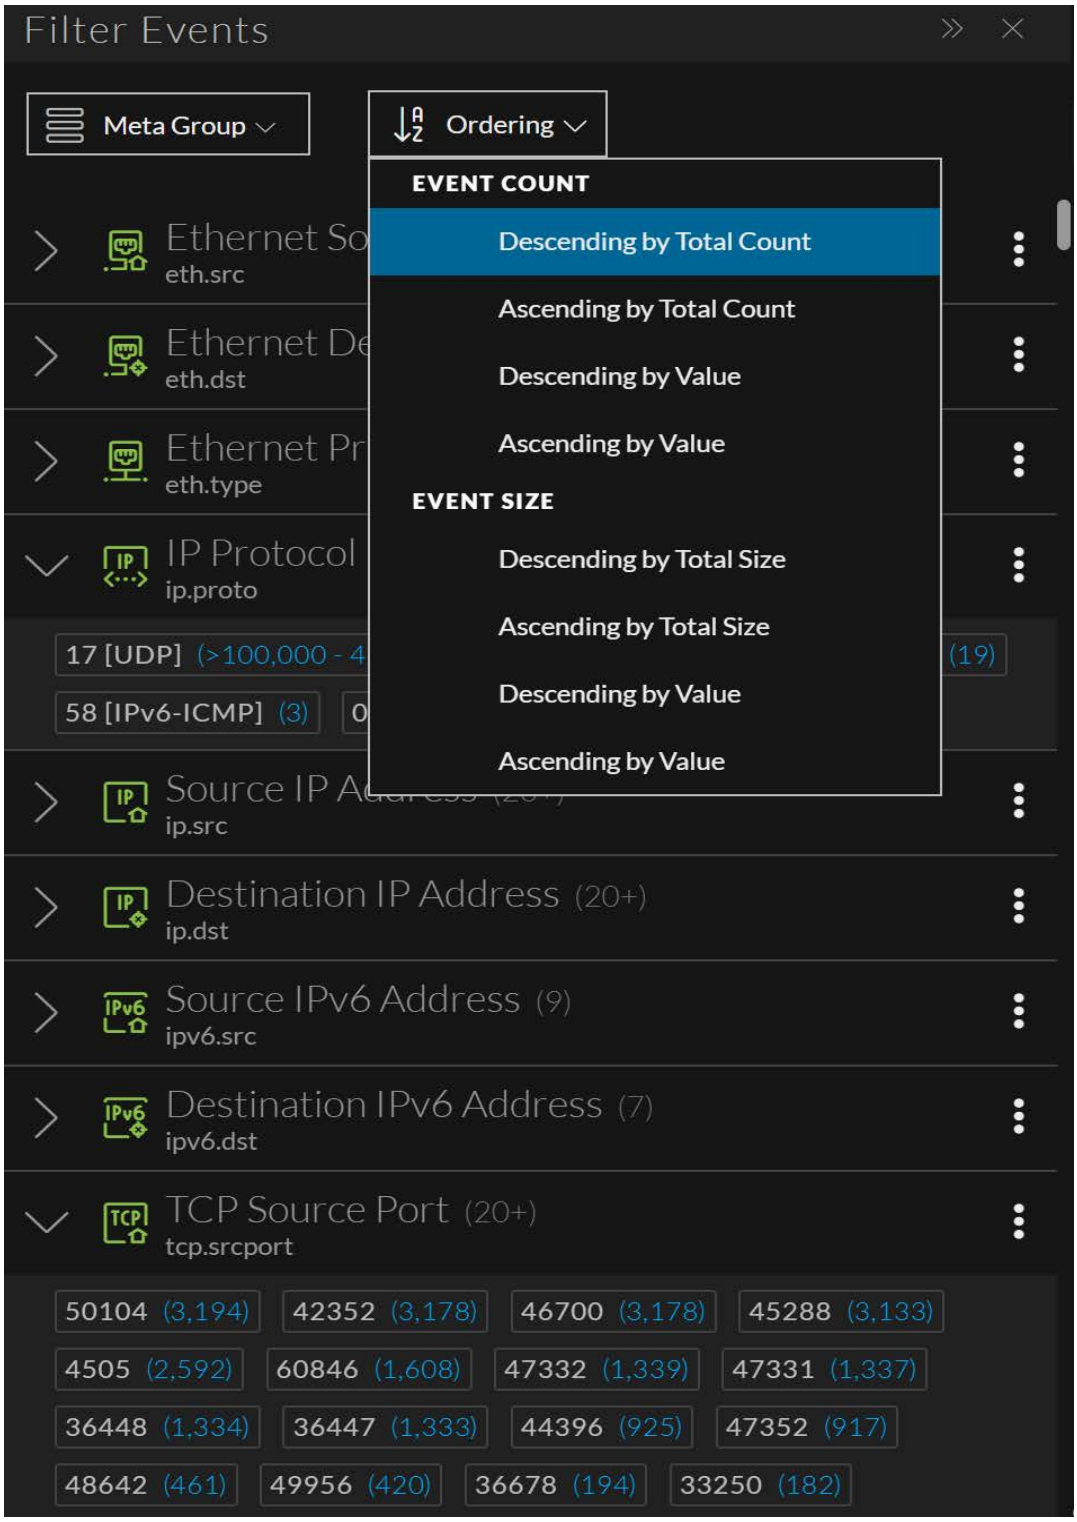
To drill into meta values in the narrow Filter Events panel:

1. Look for a meta value that is of interest, and click the value. Using the figure above as an example, to investigate the SMTP service type as opposed to other service types, click **25[SMTP]**.
   The filter is added as the last filter in the query bar, other service types are filtered out of the metadata in the Filter Events panel, and the query is executed in the Events panel.
2. Look for a meta value that is of interest, and do one of the following:
   a. Click the value. Using the figure above as an example, to investigate the SMTP service type as opposed to other service types, click **25[SMTP]**.
      The filter is added as the last filter in the query bar, and other service types are filtered out of the metadata in the Filter Events panel and the data set showing in the Events panel.
   b. Right-click the value and select **Add Filter - Do Not Run Query** in the drop-down menu.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata in the Filter Events panel, and the query is not executed in the Events panel until you click the query button.
   c. Press **CTRL (Windows)** or **CMD (MacOS)** and click the value.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata, and the query is not executed in the Events panel until you click the query button.
3. Continue clicking values to refine the set of events (drill point). As you refine the set of events, examine and reconstruct the raw events for the same set in the Events panel.

Copy the Meta Values for a Meta Key

To copy all of the visible meta values for a meta key:

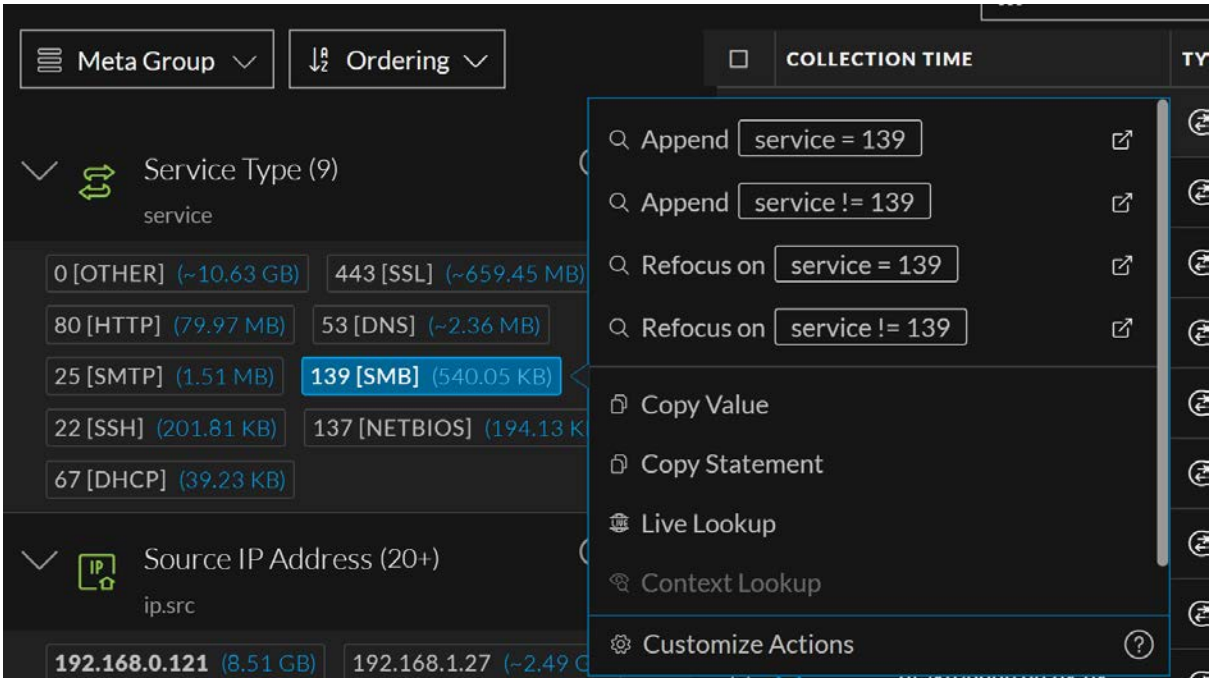
1. In the meta key row of an entry, click the Meta Key options button (⋮).
   The Meta Key options are displayed. Currently the only option is Copy Values.



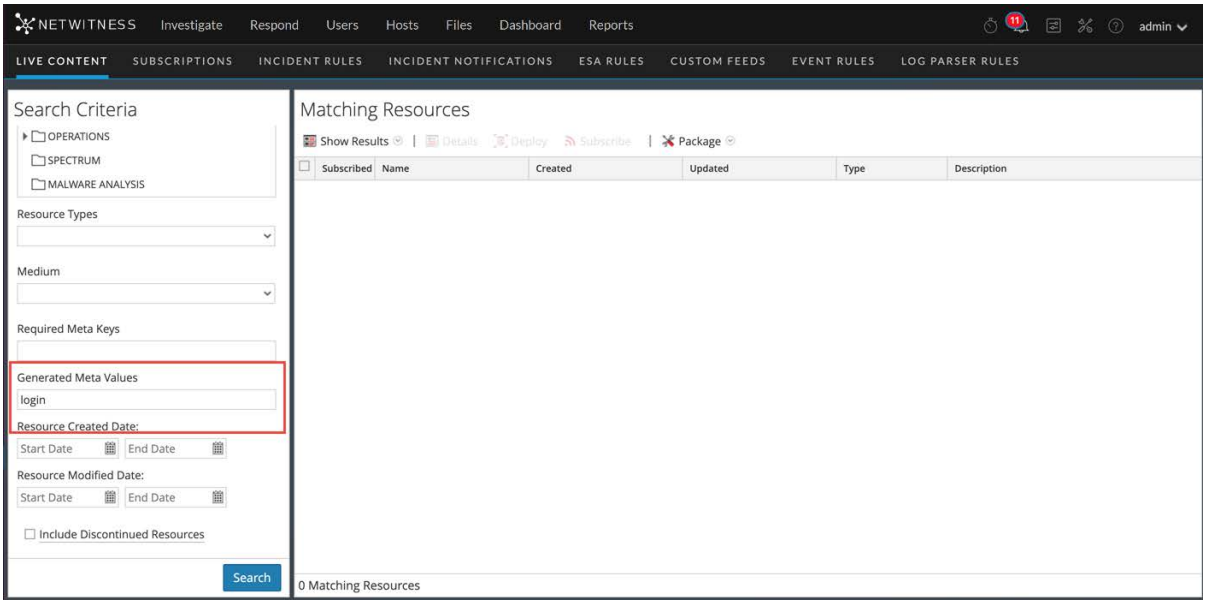
2. Click **Copy Values**.
   A comma-separated list of the values is copied to your local clipboard. This is an example of the clipboard contents: `"get", "login", "put"`.

View a Selected Meta Value in RSA Live

1. Left or right-click a meta value, for example **SMB**.
   The Meta Value drop-down menu is displayed.



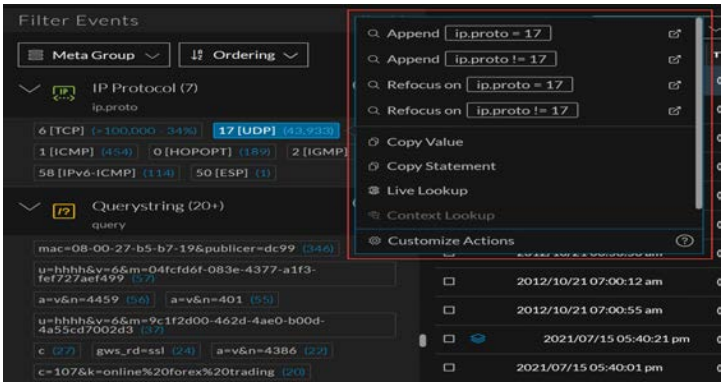
2. To look up the meta value, for example **success**, in RSA Live, select **Live Lookup**.
   The Live Search view is displayed with the meta value entered in the Generated Meta Values field, and ready for a search.



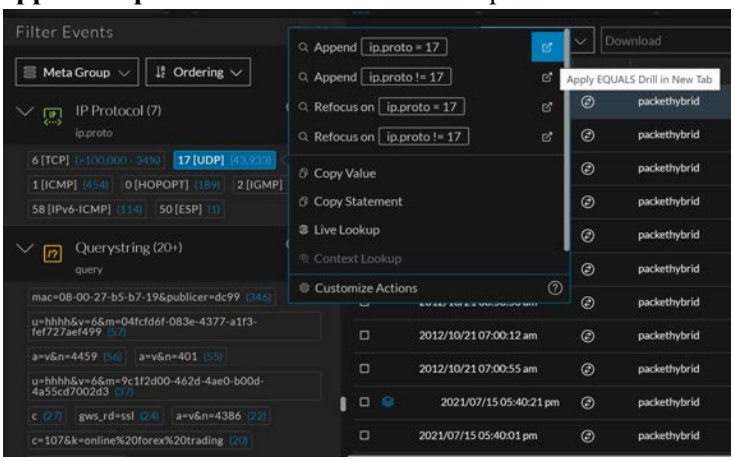
Append and Refocus the Investigation of a Meta Value in Unified Panel

For each value listed under a meta key, the focus is `<meta key> = <meta value>`. When you right-click a meta value, a context menu with different Append and refocus options is displayed. All of the append and refocus actions update the drill point in the Events panel, Filter Events panel and Meta Event panel.

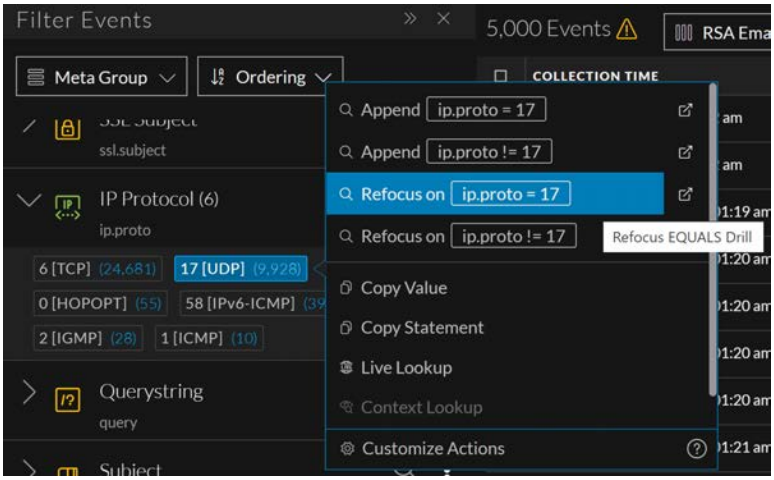
1. To append the key-value pair to the query with different operators (=, `!=`, `contains`), right-click a meta value (for example **UDP** in the figure below) and select one of the **Apply <operator> Drill** options.



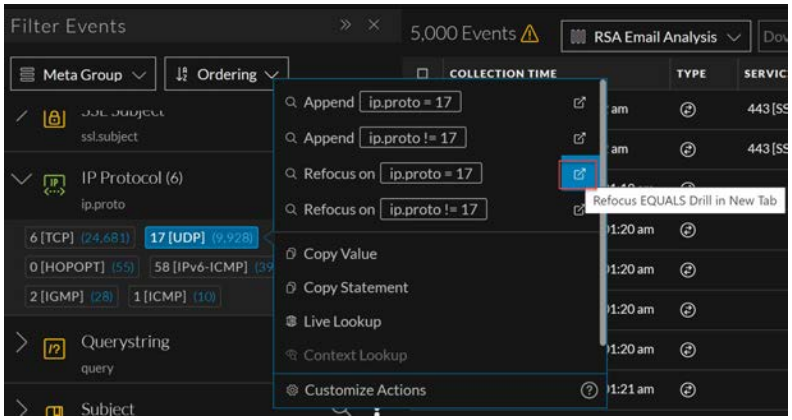
2. To append the key-value pair to the query or start the key-value pair over in a new browser tab, right-click a value and select one of the **Append New Tab > Append <operator> Drill in New Tab** or **Append <operator> Drill in New Tab** options.



3. To start the query over with the key-value pair and a different operator (=, `!=`, `contains`), right-click a value and select one of the **Refocus <operator> Drill** options.

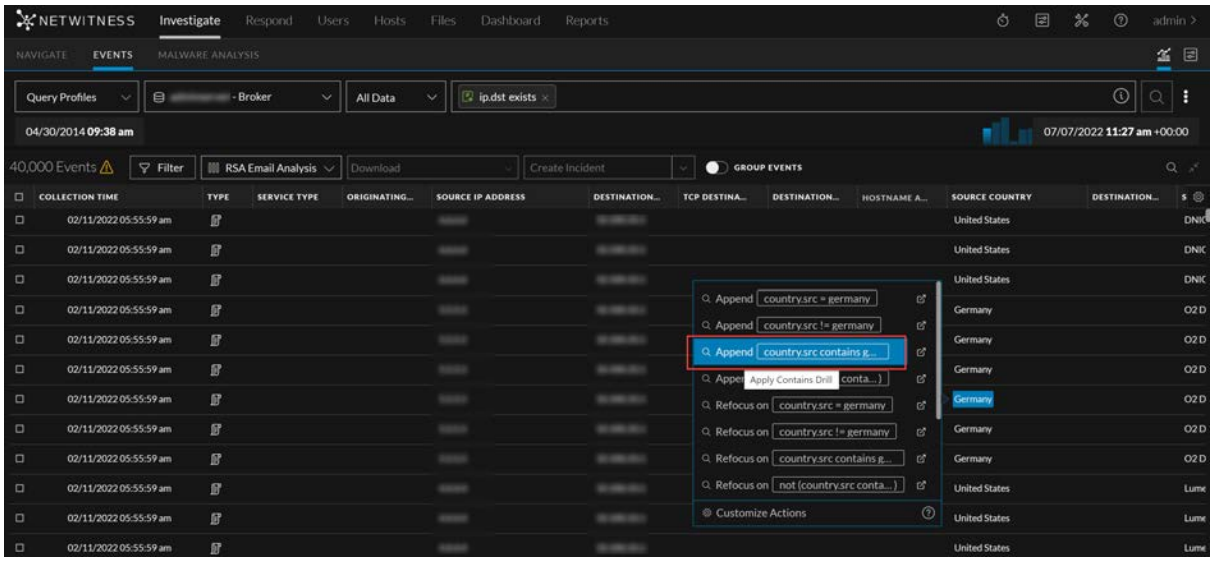


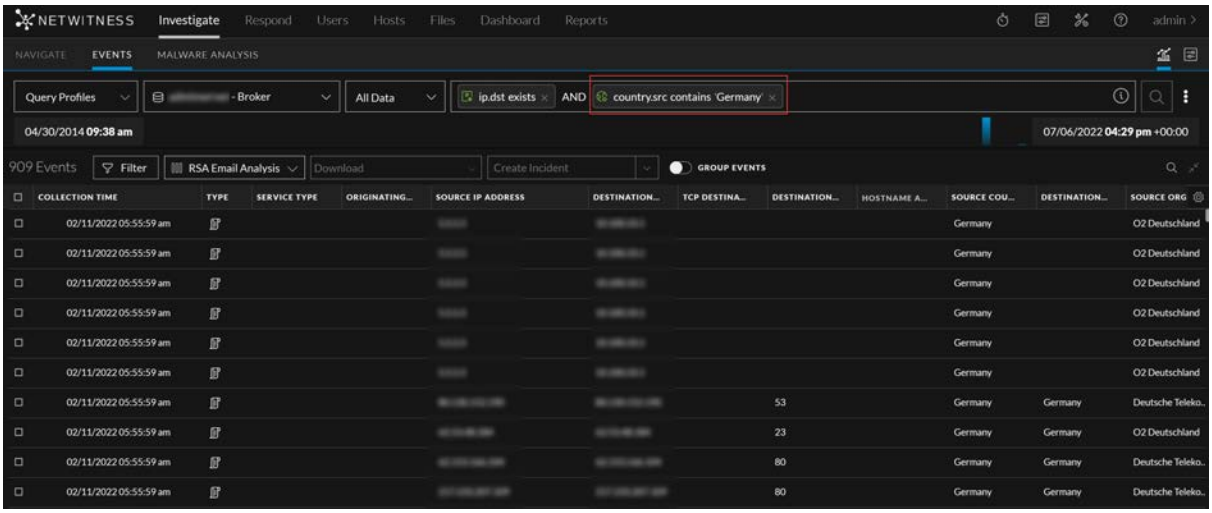
4. To append the key-value pair to the query or start the key-value pair over in a new browser tab, right-click a value and select one of the **Refocus New Tab > Refocus <operator> Drill in New Tab** or **Refocus <operator> Drill in New Tab** options.



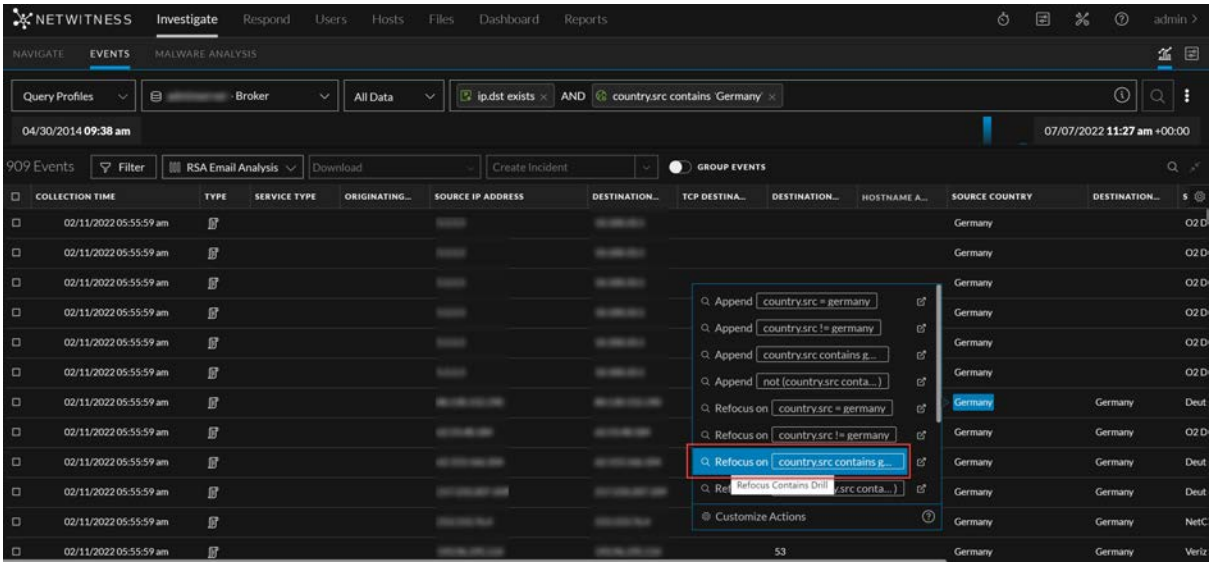
The drill is refocused according to your choice, and the new query is executed in the Events panel.

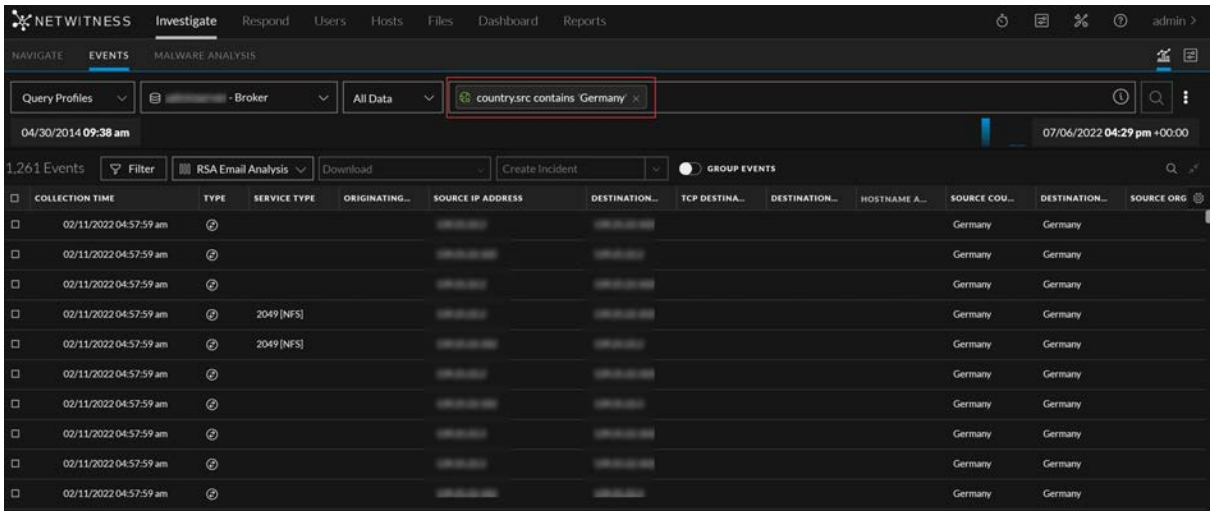
5. When you select a meta value using **Append** (*Meta* contains Meta value), the existing query is updated along with new query filters in the displayed search results.



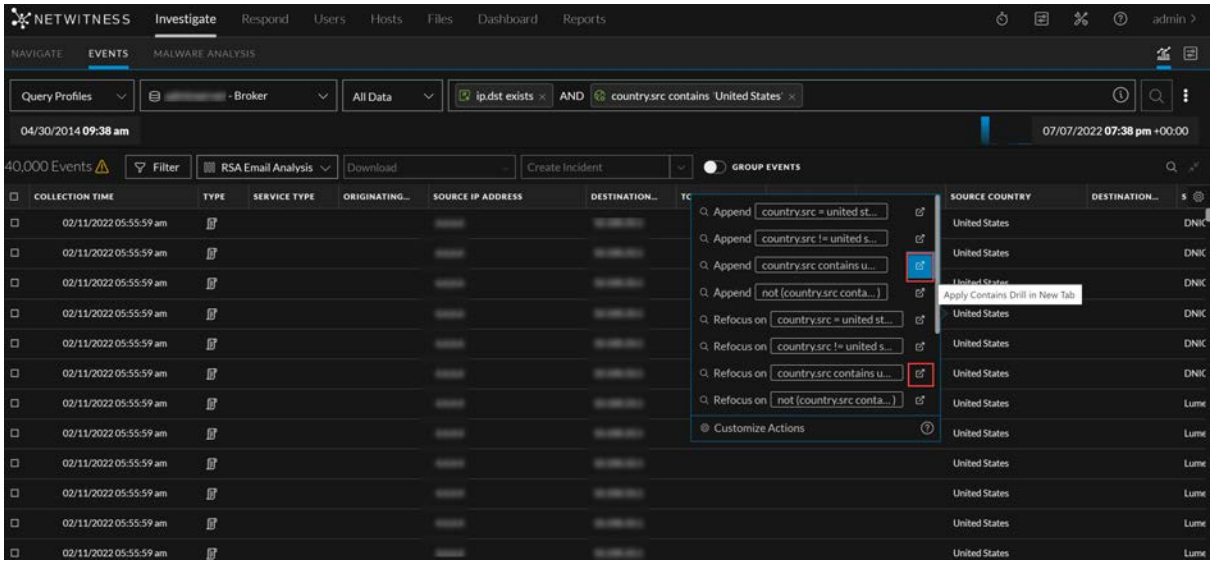


6. When you select a meta value using **Refocus** (*Meta* contains Meta value), the existing query is removed from the search result and the specified meta value results are displayed.





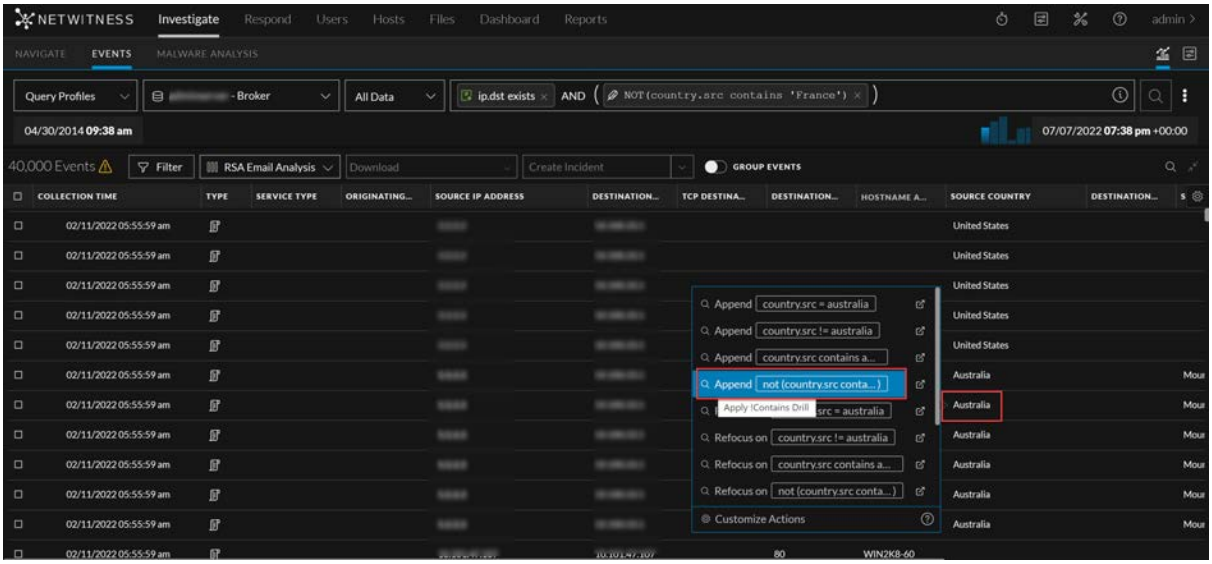
7. You can also Append and Refocus the meta values with Contains option using new tab option as shown below:

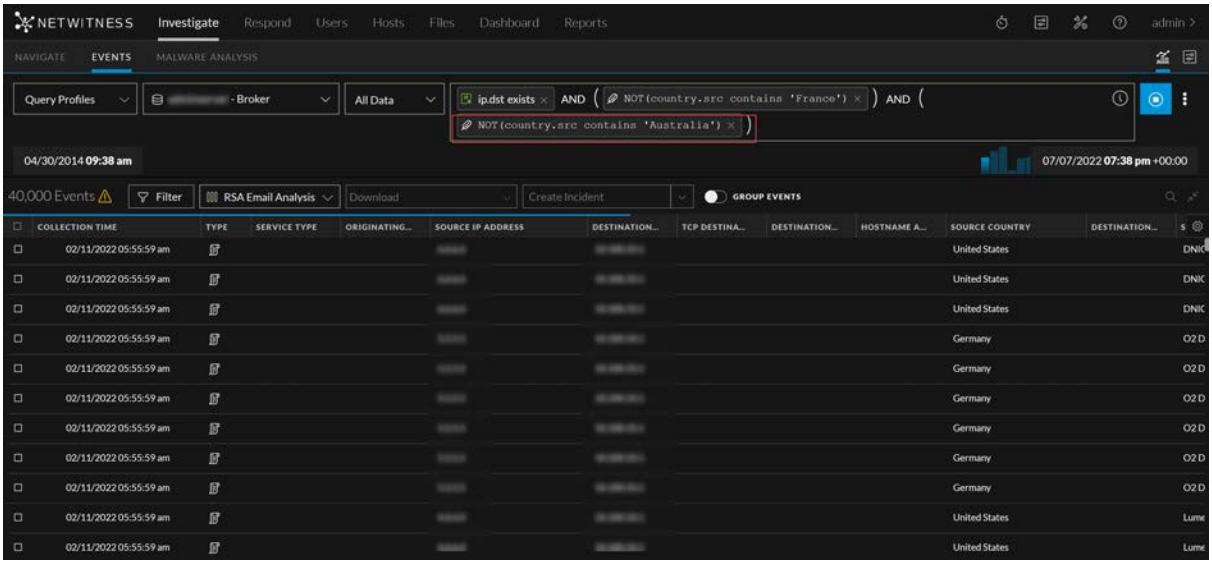


From 12.0 and later, Analysts can now exclude particular meta values while querying using the NOT (*meta* contains 'meta value') option available in the investigate unified panel. The specified meta value is removed from the query results when you use NOT(*meta* contains 'meta value') with **Append** or **Refocus** option on a specific meta value. This enhancement helps the analysts to view only the required data results in an optimized manner and conduct further investigation efficiently.

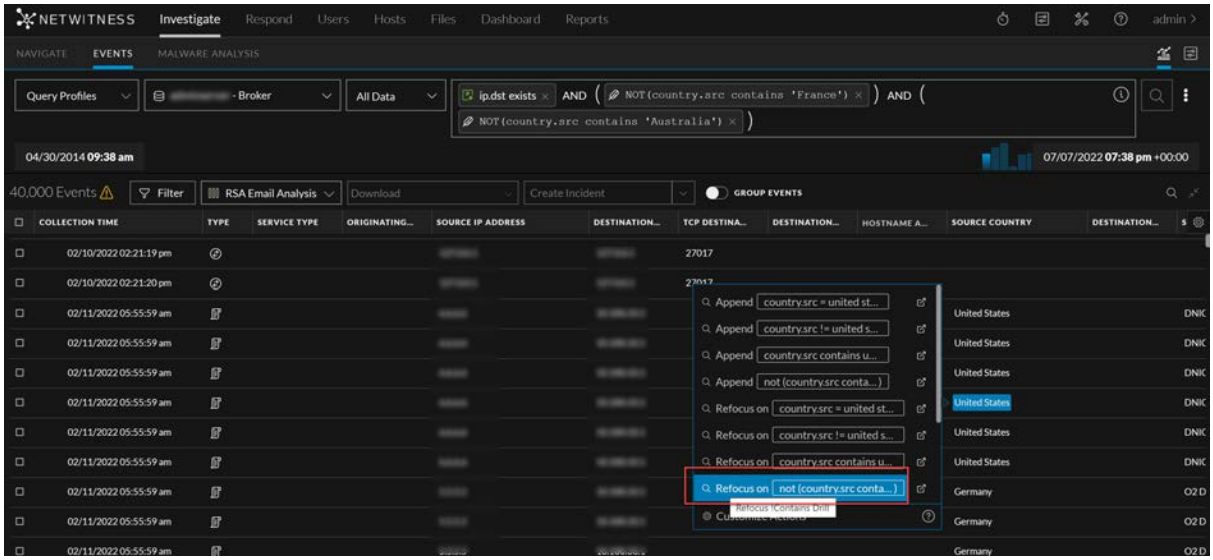
In the Events view, you can further investigate meta values in an event by left or right-clicking certain meta values and use the options in the drop-down menu.

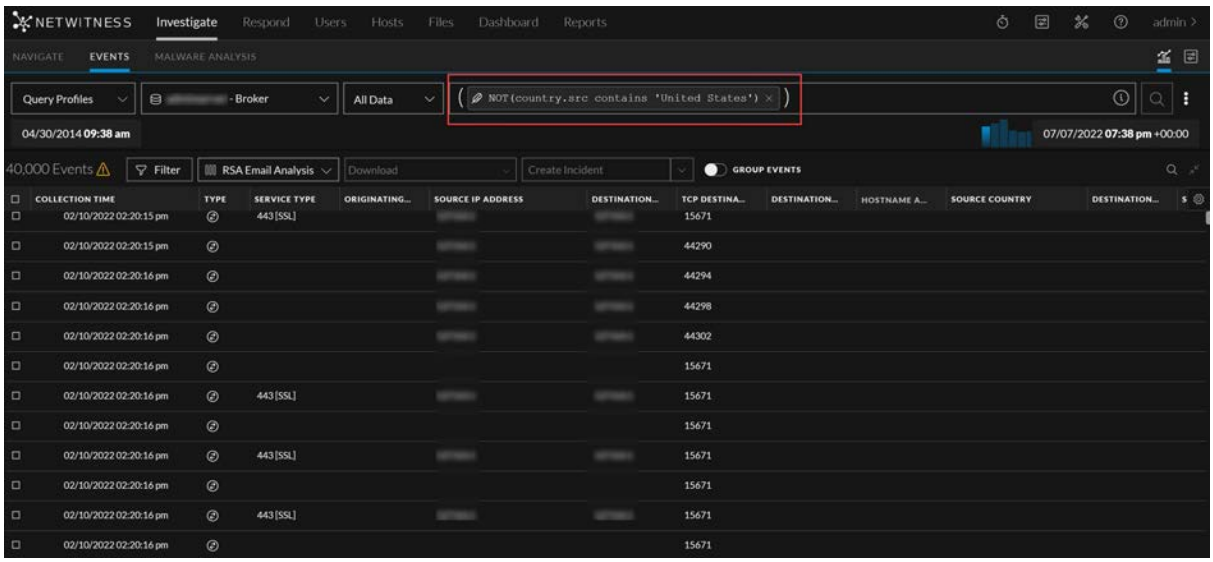
1. When you select a meta value using Append NOT(*Meta* contains Meta value), the particular query is removed, and the new query filters will be appended, resulting in NOT containing meta values.



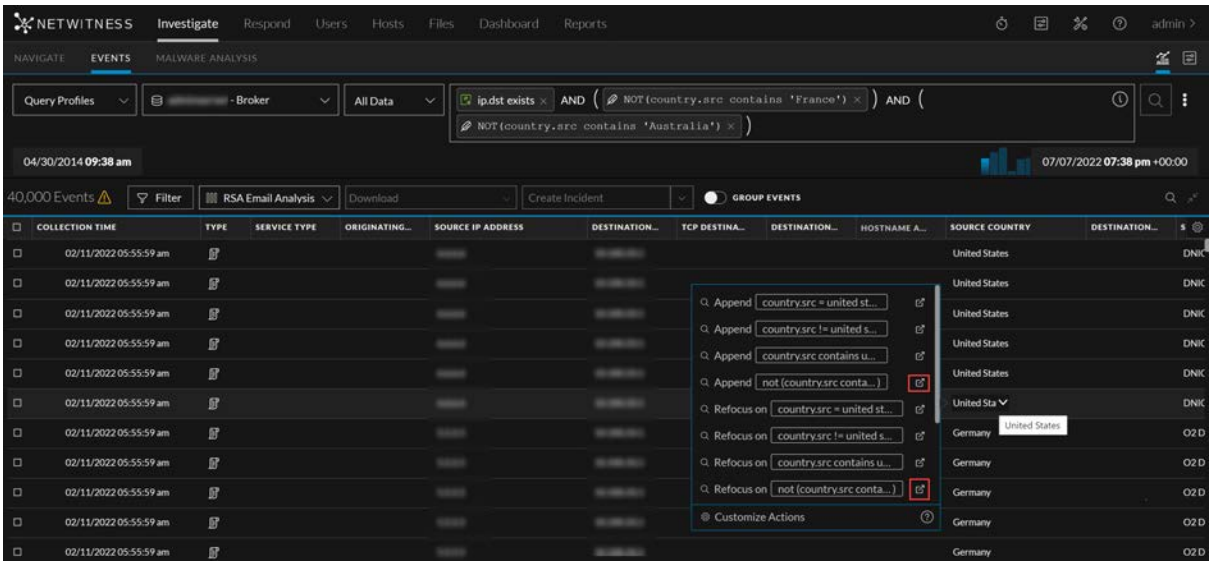


2. When you select a meta value using **Refocus** NOT(*Meta* contains 'Meta value'), the particular meta value is removed from the search query results.





3. You can also Append and Refocus the meta values with Not Contains option using new tab option as shown below:


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
NAVIGATE
EVENTS
MALWARE ANALYSIS
Query Profiles
- Broker
All Data
ip.dst exists
AND
NOT(country.src contains 'France')
AND
NOT(country.src contains 'Australia')
04/30/2014 09:38 am
07/07/2022 07:38 pm +00:00
40,000 Events
Filter
RSA Email Analysis
Download
Create Incident
GROUP EVENTS
COLLECTION TIME
TYPE
SERVICE TYPE
ORIGINATING...
SOURCE IP ADDRESS
DESTINATION...
TCP DESTINA...
DESTINATION...
HOSTNAME A...
SOURCE COUNTRY
DESTINATION...
02/11/2022 05:55:59 am
United States
Germany
Append country.src = united st...
Append country.src != united s...
Append country.src contains u...
Append not (country.src conta...)
Refocus on country.src = united st...
Refocus on country.src != united s...
Refocus on country.src contains u...
Refocus on not (country.src conta...)
Customize Actions

Filter Results in the Events View

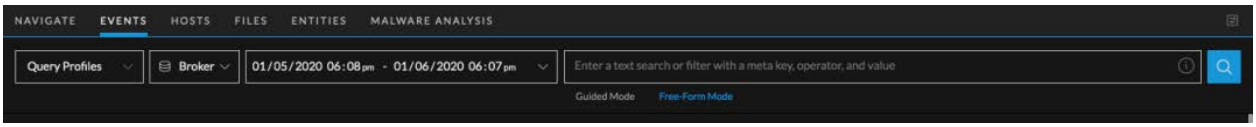
Note: This section applies to version 11.1 and later.

Filtering events in the Events view helps to narrow the focus of an investigation to a smaller, relevant set of events. You can filter events in the Events view using the Events Filter panel (Version 11.5 and later), the options in the query bar, and the options in the Events panel.

Initial Filter Using the Query Bar

When initially opening the Events view, the most basic filtering is to select a service, and time range and then query the service in the query bar. This returns a list of matching events in the Events panel. You can also select a query profile (Version 11.4 and later) and build a query to look for events that contain certain meta keys, meta values, and text in the query bar.

This figure illustrates the Version 11.4 and earlier query bar with the options to select a query profile, service, and time range to filter events and load them in the Events panel. Two modes are available Guided Mode and Free-Form Mode.

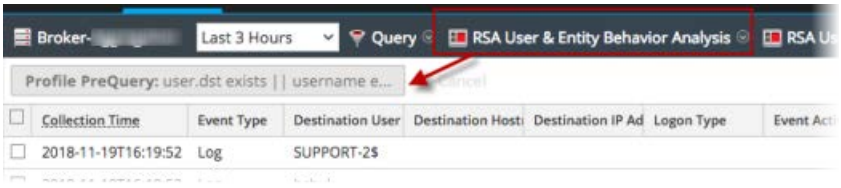


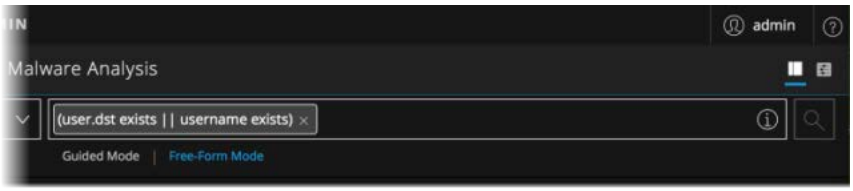
This figure illustrates the Version 11.4.1 and later query bar in which the Guided and Free-From Modes are no longer necessary. The simplified filter entry form allows you to use advanced auto-suggest options and also to enter a free-form query.



- The Query Profiles menu is available in Version 11.4 and later. You can encapsulate a query and a column group in a profile so that a useful combination of attributes is easily recalled and applied to a set of events in the Events panel (see Use Query Profiles to Encapsulate Common Areas for Investigation.)
- By default, the first service is automatically selected (unless you previously selected a service and the selected service is in browser cache). You can select a service as described in Begin an Investigation in the Events View
- If you do not select a time range, the default time range (24 hours) is used.
- The query builder field is an empty field to the right of the time range selector. This is where you build a query by creating filters. Clicking submits the query and sends a request to the selected service to load the data. In Version 11.3 and later, clicking the (console icon) opens the query console, where detailed status of the query is provided.
- When you go to the Events view from the Legacy Events view or the Navigate view, the service, time range, and any filters that were selected in the Legacy Events view or Navigate view are displayed in the query bar. The service, time range, and individual filters can be modified.

- If a profile is selected in the Legacy Events view when you right-click or double-click an event and go to the Events view, the filters from the profile (preQuery) are added to the query builder field as an editable filter. The following figures show a preQuery in the Legacy Events view, and the same query added as the first filter in the Events view.

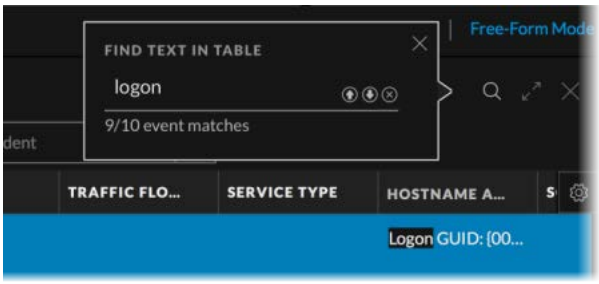




Find a Text String in the Events Panel (Version 11.4 and Later)

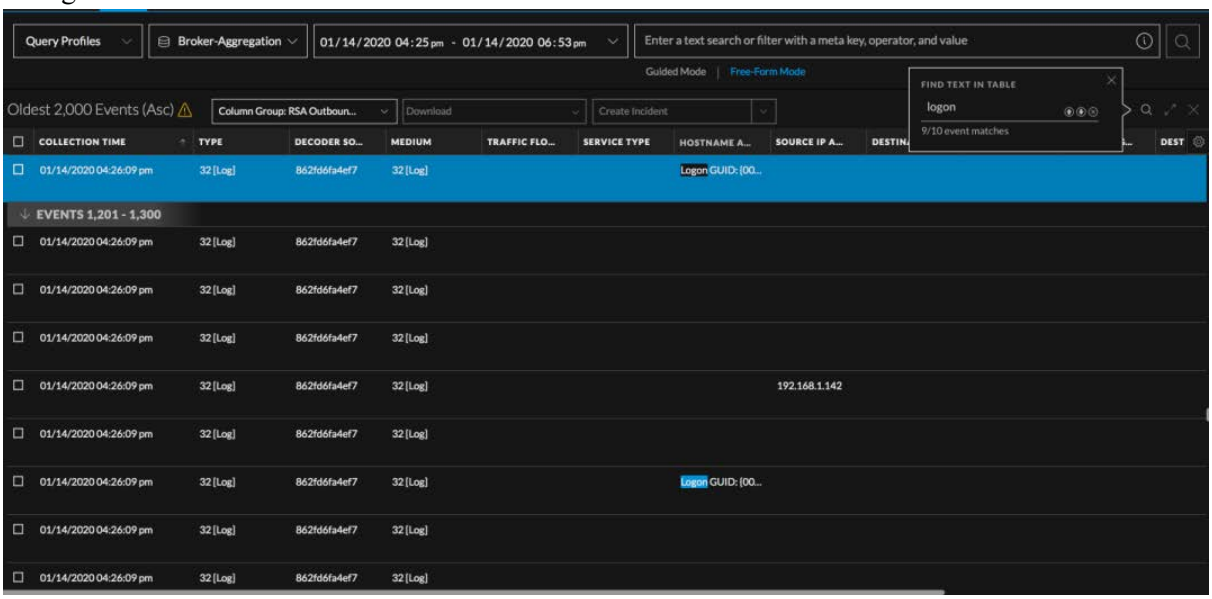
With the Events panel open, you can search for a text string in the list of events. This search is similar to the CTRL-F search in a browser window. The search scans all text in every row of the table, visible columns only, to find matching text and highlights the matches. Columns that are not displayed are not searched. The search function is disabled if the Summary column is part of the table.

1. With events loaded in the Events panel, click  on the right side of the toolbar.



2. In the **Find Text in Table** dialog, start typing a text string.
   After you type two characters, exact matches of the text string without regard to case are highlighted in the Events panel. As you type more text, highlighted events are further refined. The following figure is an example of the results found after entering "192.168" in the Find Text in Table dialog. The text string was found in 10 events. The first event is highlighted in blue with the text string within the event also highlighted. Icons are available for navigating the search results and closing the

dialog.



3. To navigate through the search results, click the up and down arrows.
   - To view the next event that contains the text string and navigate downward through the search results, click the down arrow. If you click the down arrow when viewing the last result, the first result is highlighted.
   - To view the immediately prior event that contains the text string and navigate upward through the search results, click the up arrow. If you click the up arrow while viewing the first result, the last result is highlighted.
4. To close the search dialog, click X or press the ESCAPE key. The dialog also closes if you open a reconstruction, select a new column group, or execute a new query.

Refining the Results in the Events Panel

After the initial filter and query submission, you can continue to use the options in the query bar to refine results, with two added methods of filtering the results.

- In Version 11.4, you can use column groups to optimize the number or attributes (meta keys, meta groups, meta entities) you look at for a given event (see Use Columns and Column Groups in the Events List).
- In Version 11.5 and later, you can filter events by exploring meta keys and meta values in the results in the Filter Events panel, which is a Beta release feature. This allows you to pivot through the metadata as you can in the Navigate view, and offers the added convenience of immediately seeing the matching events in sequence in the Events panel based on your drill point. The administrator can enable or disable this feature as describe in the *System Configuration Guide*.

Query Builder Concepts

In the query builder, you can reduce the number of events to an interesting set by creating three types of filters: simple, free-form, or text.

The basic syntax for each filter is as follows: `<meta key><operator><meta value>`. Here is an example: `direction = 'outbound'`.

In Version 11.4, when you type or paste a query in the query bar, the text is parsed into individual filters separated by the AND operator if the parsing engine determines that AND is needed. Earlier versions use only the AND operator between filters, and the logical operator is not visible.

- If you type `action = 'get' action = 'put'`, the result is two filters separated by AND.
- If you type `action = 'get' OR action = 'put'`, the result is two filters separated by OR.

When typing or pasting a filter for `event.time`, use one of the following formats:

- `event.time = '2020-DEC-02 23:00:00'`
- `event.time = '2000-12-20 21:00:00.000'`
- `event.time = '2000-12-20 21:00:00'`

In Version 11.4, the parsing engine converts a longer string of text that you type or paste in the query bar into individual filters. Parts of the filter that are not parsable are converted to a free-form filter. In earlier versions, a long text string is added to the query bar as a single filter. Further enhancement in Version 11.4.1 provides the ability to keep typing text for any query, where you type a meta key and an operator or an operator and a value, as a free-from query. The free-from query is parsed as usual.

- If you type, `action = 'GET' OR action is 20 || action = 'PUT'` in the query bar the Free-Form option is used. Part of this text cannot be parsed so the result is three filters separated by OR. The following figures show the query bars of version 11.4 and version 11.6 respectively.



- In Version 11.4.1, If you type a meta key-operator-value sequence and you continue typing without pressing Enter, the Free-Form option is automatically used so that you can continue typing the query. For example, you can type `medium = 1 OR medium = 2` without pressing the Enter key before `OR`. The Free-Form option is highlighted while you type and when you press Enter at the end, a free-form filter is created in the query bar.
- Text filters (Version 11.4 and later) are text strings that do not contain spaces. You can search the data set for any exact match of indexed meta keys, not all meta keys. Here are some examples: `failed`, `login`, or `attempt`.

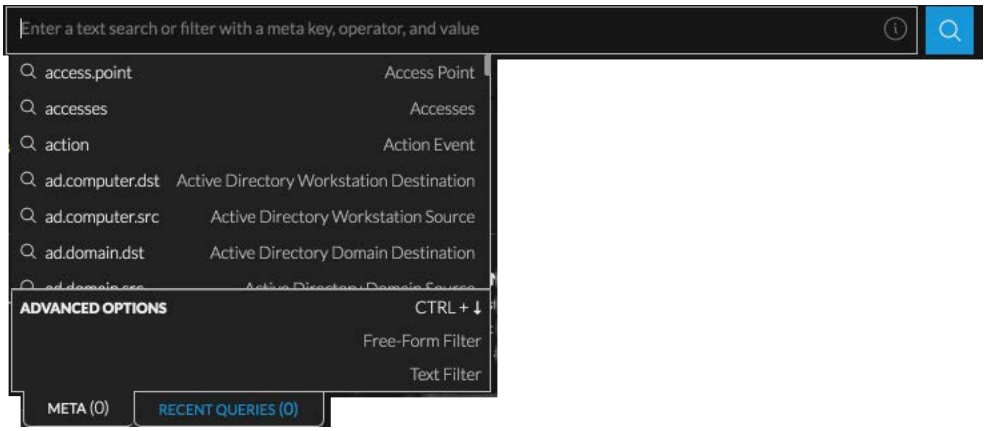
Note: In some cases, when you are typing a text filter that is close to matching a meta key and operator statement, the auto-suggest feature erroneously suggests a filter using the meta key and operator. The workaround is to begin typing the text and select Text Filter at the point where the auto-suggest feature turns the text into a meta key and operator. For example, there is meta key named `crypto`, and an operator named `contains`, and you want to create a text filter to search for `cryptocurrency`. As you type `c-r-y-p-t-o`, the next `c` in "currency" triggers the `contains` operator instead of continuing to type as a single word. To complete the text filter, right before typing that `c` in `currency`, which would trigger the `contain` operator, highlight the Text Filter option to let the system read the input as text filter.

In the query builder, each filter becomes an editable field. Filters line up from left to right, representing the sequence in which the filters were created. As more filters are added and exceed the length of a single line, they wrap to another line and the input area expands vertically so that all filters are visible without scrolling to the right.

Guided Mode vs. Free-Form Mode

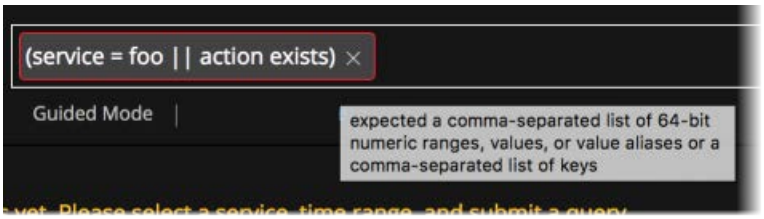
Note: Version 11.4 offered two modes for entering queries in the filter entry form: Guided Mode and Free-Form Mode. Beginning with Version 11.4.1, the powerful auto-completion features and suggested values of Guided Mode and the ability to type or paste a free-form query are fully integrated. References in this document that differentiate between Guided Mode and Free-Form Mode are for analysts using Version 11.4.0.x and earlier.

In Guided Mode, you are guided with suggestions for auto-completion that show valid meta keys and operators, and suggested values in the filter entry form. In version 11.4, you can type, paste, choose a recent query, or select from the drop-down menu. Earlier versions do not support pasting text and recent queries. This is an example of the 11.4 filter entry form.



As you create filters, the syntax of each filter is validated and invalid filters are marked by a red outline. If you hover the mouse over the filter, a message that explains the error is displayed.

In Version 11.3 and later, free-form filters are validated on the server side, which may take additional time. If you submit the query before the server has returned filter validation results, the is replaced by a spinner. When server validation returns, a query with no invalid filters begins execution. If the query contains an invalid filter, execution is terminated and the invalid filter is outlined in red. This is an example of an invalid query.



In Free-Form Mode, you can type or paste a long text string. There is no auto-suggestion, and validation is performed on the server side when you submit the query. If an error is found, the query does not execute.

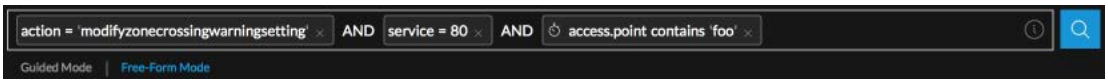
Note: The button has a different label in versions earlier than Version 11.3. It was previously named **Query Events**.

Clicking Guided Mode or Free-From Mode toggles between modes. If you selected Free-Form Mode the last time you logged in, this choice is stored in browser cache and is used until the browser cache is cleared.

- When you switch from Guided Mode to Free-Form Mode, filters that you created in Guided Mode are transformed to a text query in the Free-Form field.
- When you switch from Free-Form Mode to Guided Mode, the query you were typing is added to the query bar as individual simple filters, but it does not include auto-suggest options.

  Note: Before Version 11.3, a Free-Form filter could not be edited in Guided Mode.

The following figure is an example of the query bar with the Guided Mode query builder with several filters.



The following figure is an example of the Free-Form query builder in use.



Concepts for Editing Multiple Filters

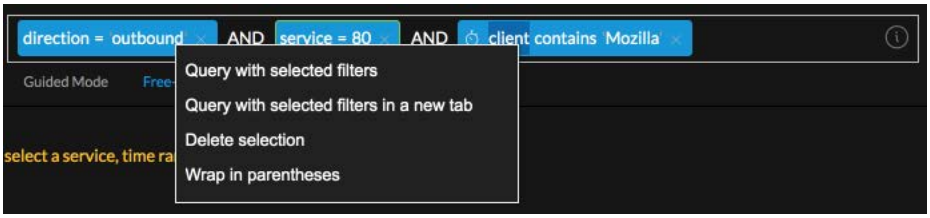
As you work in the query builder, you can see when a filter has focus for editing (a green outline) and which filters are selected (blue background). This is useful because you can have multiple filters selected for right-click actions, but only one can be edited at a time. The figure below shows the green outline marking a filter that has focus and the blue background indicating that two filters are selected.



This figure illustrates the same set of filters with all filters selected (blue background) and one filter that has focus (blue background and green outline).

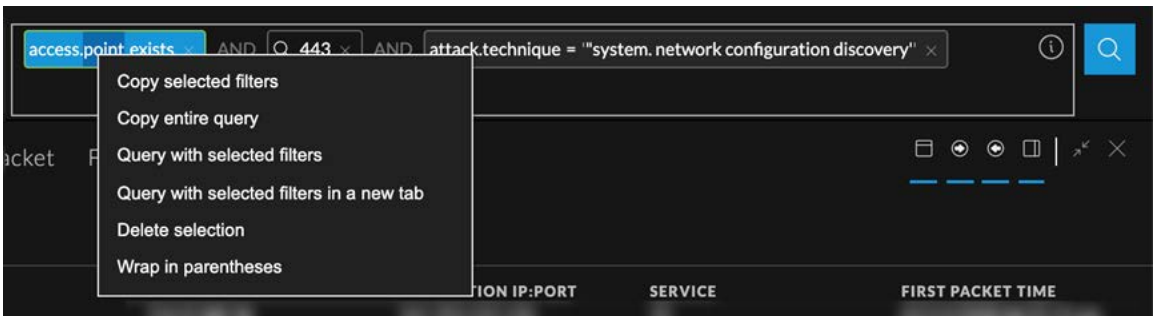


A right-click action from the drop-down menu applies to all selected filters as shown in this figure showing Version 11.4 options.



In Version 11.4.1, the menu has two new copy options as shown in the following figure. The options allow you to share the clipboard contents with other analysts or paste the contents to the query bar. You can:

- Select a single filter, right-click it, and then copy the entire query to the local clipboard.
- Select multiple filters, right-click one of them, then copy the selected filters.



These are a few basic concepts that explain how to work in the query builder:

- You can select multiple filters, but only one can have focus and the last selected filter is the one with active focus at any point in time.
- To select a filter and give it focus, click the filter. To deselect the filter and remove focus, click the filter again, press **Esc**, or click anywhere else on the page.
- To add a filter, click before or after an existing filter. To create a new filter before or after the filter in focus, press the right or left arrow key.
- To open a filter for editing, double-click the filter or click it and press **Enter**. To exit without saving changes and leave the filter in focus, press **Esc**.
- To delete a filter, click the filter and press **Delete** or click **X** on the filter.
- (Version 11.6) If a filter is awaiting user's input, it is highlighted with blue color .



- (Version 11.6) An invalid filter is highlighted with red color.



- You can hover over the filters in the query bar to display tooltip messages.

The Version 11.3 and Earlier Query Builder

In Version 11.1, the user interface guides you as you create and edit each simple filter: `<meta key> <operator> <meta value>` filter. The user interface supports only simple filters. If you open an event from the Legacy Events view or Navigate view and the filter has more than one operator, `||`,`&&`, `()`, `REGEX`, or `LENGTH`, the filter is added, but editing is not supported in the Events view. Refer to the Investigate User Guide for NetWitness Platform 11.3 for details; it is available here in PDF form: Master Table of Contents for NetWitness Platform Version 11.

In Version 11.2, the user interface has two modes: Guided Mode and Free-Form Mode. Guided Mode guides you as you create and edit each simple `<meta key> <operator> <meta value>` filter. The default mode is Guided Mode, which includes auto-suggest and validation options. You can enter a long text string in Free-Form Mode. In Free-Form Mode validation is performed when you run the query. Refer to the Investigate User Guide for NetWitness Platform 11.2 for details; it is available here in PDF form: Master Table of Contents for NetWitness Platform Version 11.

In Version 11.3, the user interface has these added features:

- When you go to the Events view from the Navigate or Events view, with a preQuery in effect for a profile, the preQuery that was displayed in the breadcrumb is available in the Events view query builder as an editable filter.
- The auto-suggest options in the user interface are augmented with an Advanced Options section that allows for the creation of a free-form filter. Free-Form Mode is still useful for pasting an entire long text string.
- You can cancel a query while it is executing.
- You can view detailed status information about the query being executed in the query console.

The Version 11.4 Query Builder

You can type; select meta keys, operators, and values from the drop-down menus; or paste a filter in the query bar. Added 11.4 features in the Guided Mode filter entry form are described in detail below.

Meta Keys Cached for Faster Loading

When the Events view opens, meta keys from all connected services are cached for faster data loading. These meta keys are available in user interface elements that have auto-suggest meta keys. (When you are building a column group or profile, if you are expecting to see a meta key and it is not displayed, select the service where the key was added to force a cache update. This usually occurs only when a meta key is not added to all concentrators.)

Text Filter

You can create a text filter to find a text string in the data set which is indicated with . You can use a text filter with no knowledge of meta keys that would contain the values. One text filter per query is supported. A text filter looks through indexed meta keys, not all meta keys.

Pasting Text Instead of Typing

When creating a filter, you can paste a meta key or value in the filter entry form. When you paste text into the filter entry form instead of typing the text, the text is parsed appropriately to create one or more filters. Any portion of the text that cannot be parsed is converted to a free-form filter.

Select All Filters and Copy All Filters (Version 11.4.1)

While creating a filter in the Events view query bar, you can use keyboard commands to select all filters (Cmd-A for MacOS, Ctrl-A for Windows) and then copy the selection to the clipboard (Cmd-C for MacOS, Ctrl-C for Windows). The clipboard text is available to share with other analysts or to paste in the query bar using Cmd-V or Ctrl-V).

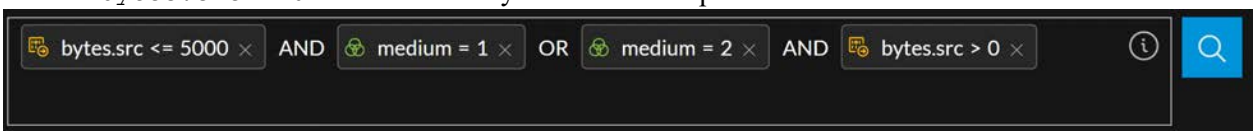
Use of Recent Queries

The filter entry form offers two methods of entering meta keys, operators, and values: the Meta tab and Recent Queries tab. The Meta tab is the same as the filter entry form for prior versions except that a count of matching results is given in the tab label and icons mark meta keys that are indexed by key, indexed by value, and not indexed. In the Recent Queries tab, up to 100 recent queries are displayed. The list is filtered as you type to show only queries that contain the typed text, and you can select a query from the list.

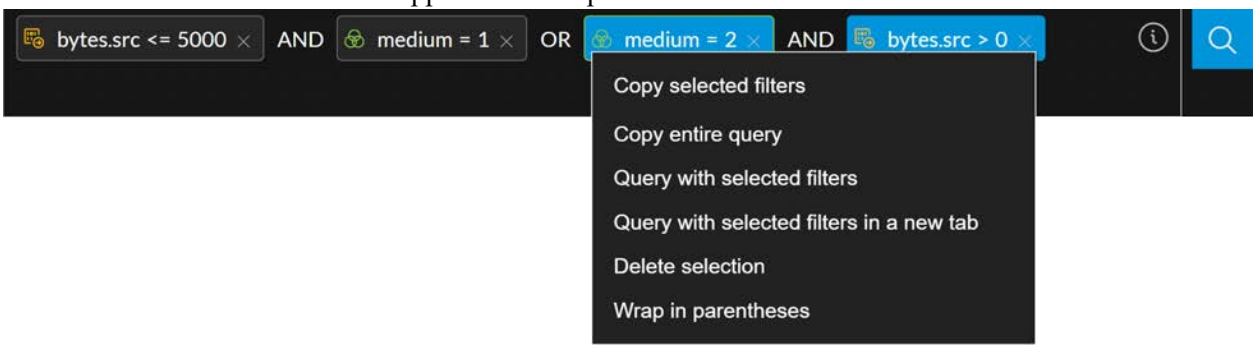
Use of Advanced Operators

Auto-suggest can parse the following advanced operators that you paste or type into the filter entry form: `<`, `>`, `<=`, `>=`, `OR`, `||`, `AND`, `&&`, `()`, `regex`, and `length`. The text is parsed as multiple filters. For example, if you type or paste `medium > 0 && medium <= 100`, the text is parsed as two simple filters with an explicit AND operator: `medium > 0 AND medium <= 100`. If you type or paste `bytes.src <= 5000 && medium = 1 || medium = 2 && bytes.src > 0`, four simple filters are created with `AND` and `OR` operators separating them: `bytes.src <= 5000 AND medium = 1 OR medium =`

`2 AND bytes.src > 0` to make as many valid filters as possible.



This filter is an example of a filter in which it would be useful to add parentheses. You can select `medium = 2` and `bytes.src > 0`, then right-click and select **Wrap in Parentheses** from the drop-down menu. Text filters are not supported inside parentheses.



The resulting query is `bytes.src <= 5000 AND medium = 1 OR (medium = 2 AND bytes.src > 0)`.



If you encounter errors while creating filters, look for tooltip messages and check the documentation.

Easy Use of AND/OR Operators

When you type `||` and `&&`, they are displayed as `OR` and `AND` in the query bar. You can change `OR` to `AND` and `AND` to `OR` by clicking the word. When you insert the cursor to add a filter, the `AND` operator is added before the cursor. When you delete a filter, orphaned `OR` and `AND` operators are removed. The operator for a text filter must be AND because text filters are always ANDed to a query.

Automatically Balanced Parentheses

When creating and editing filters in the query builder, parentheses pairs are automatically balanced as you type. If you type an open parenthesis in a filter that is open for editing or before a selected filter, the close parenthesis is added at the end of the filter. This works intuitively as you type so that you can add new filters on either side of the parenthesis and between parentheses when there are nested parentheses. Orphaned parentheses are automatically removed. If adding parentheses would create an invalid filter, the parentheses are not added. You can also right-click selected filters, and add parentheses using the Wrap in parentheses option. This option is only available when the result would be a valid filter.

Hints about Available Values

For properly indexed meta keys, the user interface provides hints about available values related to the time range of the query. Up to 100 suggested values are returned and, when you type text, the list of 100 values is filtered to include only relevant values. If no matching values are returned, a message advises "No suggestions found." (The suggested values are based solely on the time range; filters in the query do not filter the list of 100.)

CIDR Notation and Shorthand

When entering a value for an IP address in a filter, you can use CIDR notation to filter for addresses within a range.

The IPv4 CIDR block range is 0 to 32. For example `10.20.30.0/24` specifies `10.20.30.0` with a subnet mask of `255.255.255.0`, which will match an IP in the range `10.20.30.0` through `10.20.30.255`.

The IPv6 CIDR block range is 0 to 128, for example,
`1203.0fe1:fe82:b896:89b0:8a7c:99bf:323d/32` specifies
`1203:0fe1:0000:0000:0000:0000:0000:0000` through
`1203:0fe1:ffff:ffff:ffff:ffff:ffff:ffff`.

You can also use shorthand to remove groups of zeros or leading zeros in a group in IPv6 addresses, for example,

`1203:fe1::`

There must be no spaces between the IP address and the CIDR mask that you are using.

Ranges or Series of Values

For meta keys that have numerical data, you can use a range of values, a series of values, or both to filter data. For example, this query has a comma-separated list, and two of the values in the series are ranges `src.port = 0-1023, 1024-1050, 65535`. If a comma is a part of a value, the value must be wrapped in quotes. For example, `get,post` is interpreted as two separate values, while `'get,post'` is interpreted as one value. A range of values must be a valid range of positive integers, separated by a dash (with our without a space before and after). The first number in the range must be smaller than the second. For example, `0-1023` and `0 - 1023` are valid ranges, but these are not valid ranges: `-10 - 50`, `50 - 10`, `50.8 - 60.2`, `50 - 70x`.

No Separating Space Required After Meta Key and Operator (Version 11.4.1)

Filters in the query bar need a space between the meta key and the operator, and between the operator and the value. Operators must be typed with a separating space in the filter entry form in order to use the auto-suggest functions for operators and values. To improve the user experience when typing a query, the filter entry form accepts operators typed with no separating space after the meta key. When you type an operator with no separating space, a value is auto-suggested as usual and a space is added between the meta key and the operator. When you type an operator and a value without a separating space, the space between them is automatically added.

Select a Time Range

The Time Range selector limits the events returned in the Events view to a specific time range. The time range is displayed in the format `Start Time - End Time`, showing the date, hours, and minutes in your current timezone based on timezone settings configured for your profile. In Version 11.3 and later, you can choose a time range relative to the current collection time or create a custom time range. The time and date format is based on preferences set for the Events view in the User Preferences dialog (select  > **Profile**).

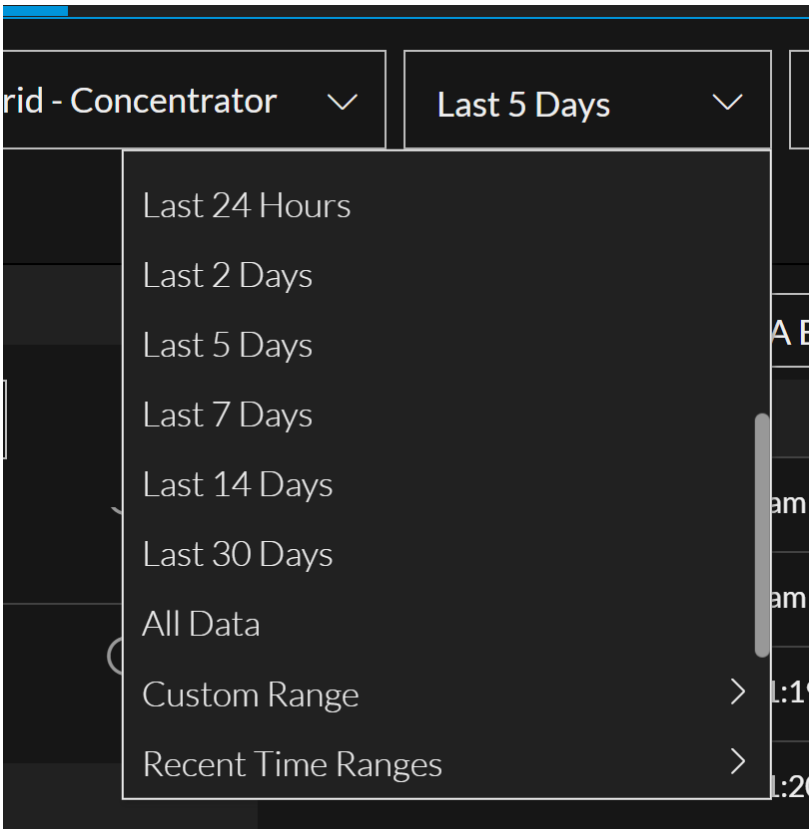
- By default, the date format is MM/DD/YYYY. You can change the format to DD/MM/YYYY, or YYYY/MM/DD in the User Preferences dialog.
- Start time and end time are in the format of HH:MM. Although seconds are not displayed, the value for start time always defaults to HH:MM:00 seconds, and the value for end time always defaults to HH:MM:59 seconds. As an example, a time range of 6:45 pm - 7:45 pm is interpreted as 06:45:00 - 07:45:59 pm.
- The default time range is the 24-hour clock; you can change it to 12-hour periods.

Note: By default, the time format for downloads is Epoch format, which shows the time as a numerical value representing the number of seconds from the Unix epoch, January 1, 1970. The resulting number requires a conversion to be understood. Your administrator can change the setting for time format in downloads to combine your user preference time zone, date format, and time format into an easily understood representation, which follows the industry standard ISO 8601 representation when possible. Given this time: `04/13/2020 09:17:36 am` with timezone US/Pacific `(GMT-7:00`, this is an example of the time on the 12-hour clock as it appears in the user interface: `04/13/2020 09:17:36 am`. In the download, Epoch format would represent the time as `61547519856000`. If your administrator set the time format for downloads to the easily readable representation, the same time would be represented as follows: `04-13-2020T09:17:36AM-07:00`.

The time format for a query is based on preferences set for the Events view in the Event Preferences dialog (select). The time format can be either database time or wall clock time. When Database Time is selected, the start and end time for a query is based on the time that the event was captured (collection time). When Wall Clock Time is selected, the query is executed using the end time based on the current browser time; the start time is calculated based on that end time and the time range. This and other Events view preferences are described in Configure the Events View.

To edit the time range, do one of the following:

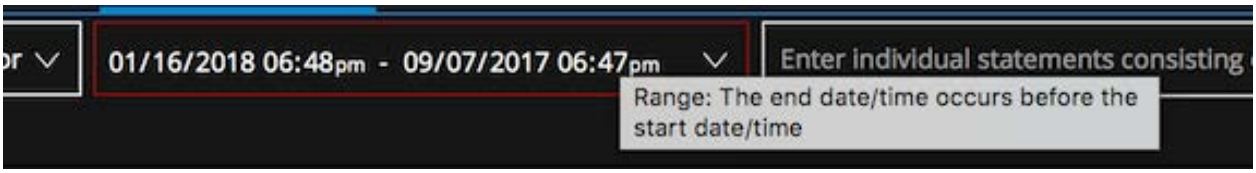
1. Click the drop-down arrow inside the **Time Range** selector and select a time range from the list. Options are in minutes, hours, days, or all data.



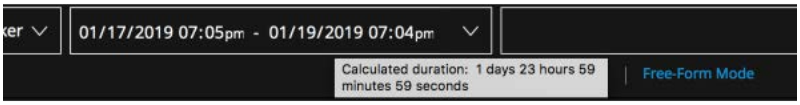
(Version 11.3 and later) Edit the time range directly by clicking the year, month, day, hour, or minute displayed in the query bar. When a value is highlighted, type a new value for either the start or end time. If your time format preferences are set to 12-hour periods, click **am** or **pm** to toggle between the two options.



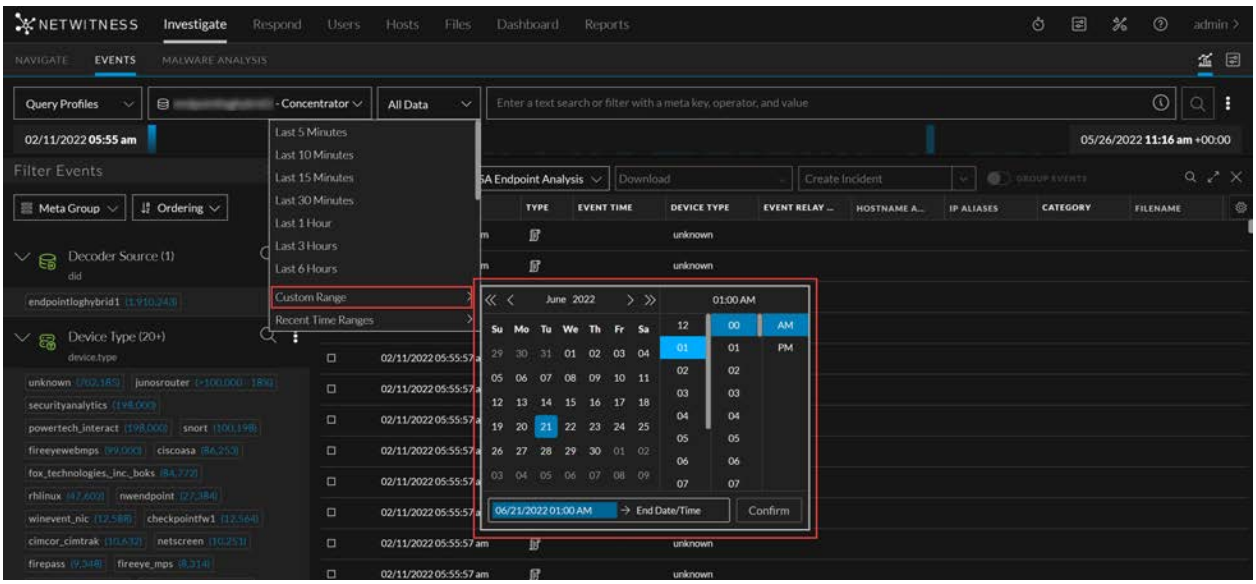
If the time range is invalid (for example, the start time is later than the end time), a red border appears around the Time Range selector. The button is disabled because the query is no longer possible, and a tool tip shows an error message explaining what you need to change. The following figure shows an invalid time range.



The selected time range is stored in your browser for the service being queried; you can set different time ranges for different services. A tool tip shows the calculated duration of the query. The following figure is an example of the tool tip.

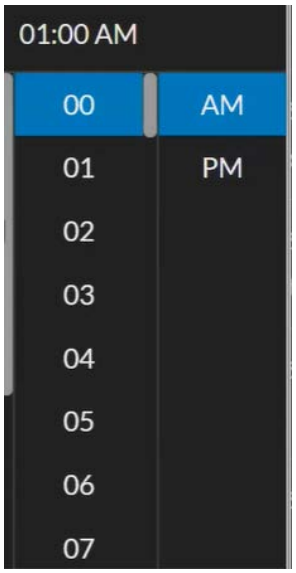


In 12.0 and later, in addition to the existing options, the **Custom Range** option in Investigate Events view allows analyst to select a specific time, date, month, and year or a date range to run a query and filter events. On clicking the **Custom Range** option, a calendar view is displayed with a current day, time, and date details. This enhancement helps the analysts to select date and time quickly and avoid manual intervention therefore avoiding human errors (typos).



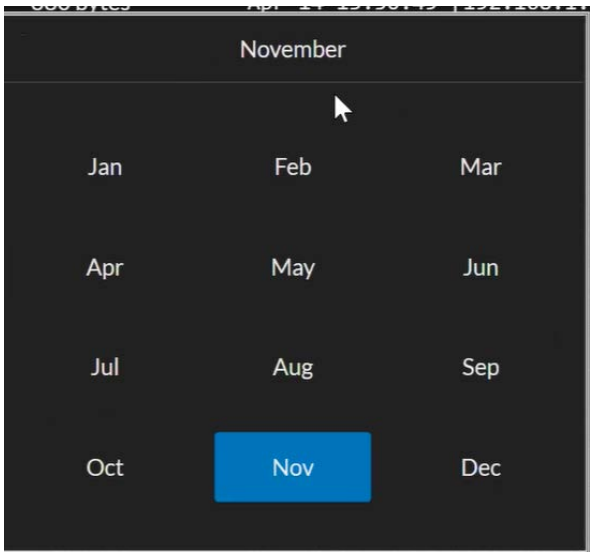
The analyst can use the following to navigate within the calendar:

- « and » to toggle between months.
- « and » to toggle between years.

- 

to select a specific time.
- 

to select the start date and time.
- 

to select the end date and time.
- 

to select an year and the year range.
- Click **Confirm** to save the selection.

Submit a Query

The button on the right side of the query bar is active as needed to submit a query. In version 11.3 and earlier, when you click , all of the filters are ANDed to generate results and the button becomes inactive. In Version 11.4, because the query may contain other operators besides AND, the query is submitted as is. The button becomes active again in these conditions:

- If you change the service in the query bar or change the column group in the Events panel, a network call for data for a reconstruction in the Events panel continues to use the previous service, time range, and metadata filters until you submit the new query. The button becomes active as an indicator that the data in the view is stale.
- If more than a minute has passed and the original query's time range would no longer generate the same result set, the button becomes active as an indicator that results may be stale. In Version 11.3 and later, a setting in the Events view preferences determines this behavior by enabling or disabling the Update Relative Time Window Automatically option (see Configure the Events View.)

Cancel Execution of a Query

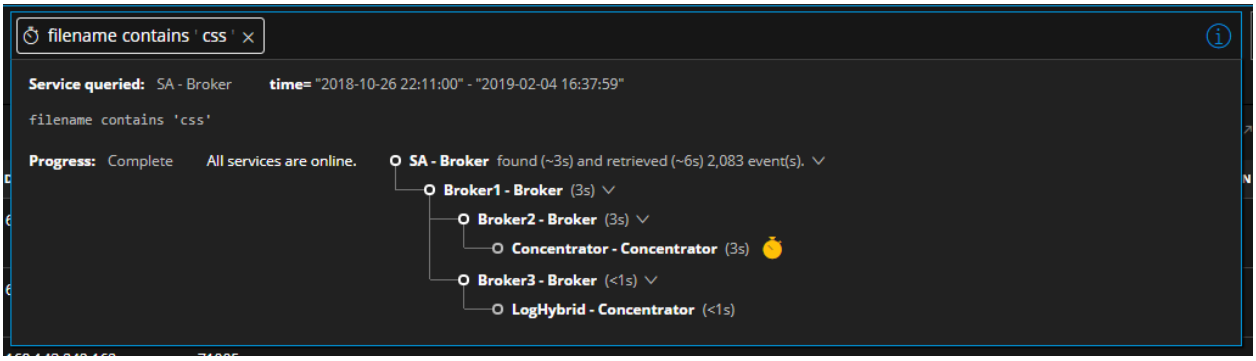
After you click to submit a query, the button changes to (the stop query option). The stop query option remains until all the events are loaded in the Events panel. To cancel the query, click .

If the query is canceled before all results have been returned, the following message is displayed at the end of results in the Events list: "Because the query was canceled, only partial results are displayed."

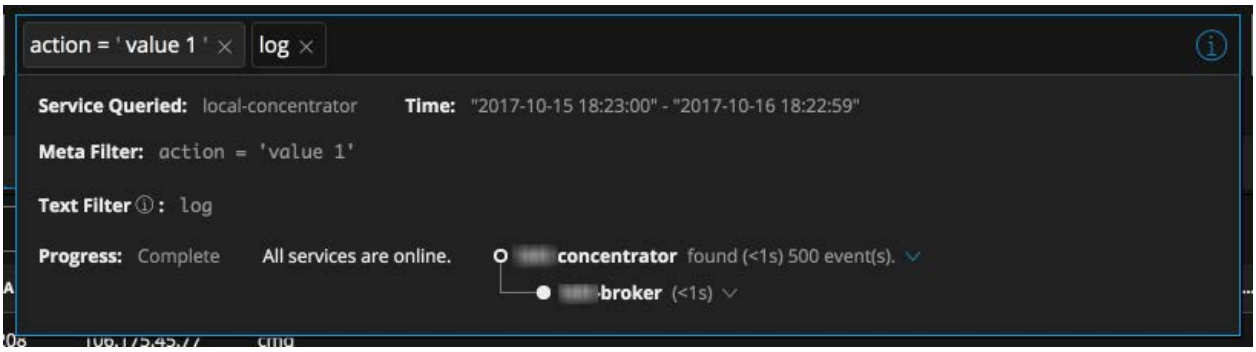
View Status of a Query

After a query is submitted, you can click the information icon ( ) in the query bar to open the query console. In the query console, you can see which service, time range, and metadata was queried as well as real-time information about the status of the query and the services being queried. The time range displayed in the Query console always shows each date as YYYY-MM-DD. Here is an example range from the Query Console: `"2014-09-20 20:57:00"-"2018-11-02 18:57:59"`.)

The following figure is an example of the Query Console for Version 11.3 when a query executes successfully and the slowest service is marked by an amber stopwatch.



The following figure is an example of the information in the Query Console for Version 11.4 after a query that includes a text filter executes. Notice that the query is shown in two fields, Meta Filter and Text Filter.



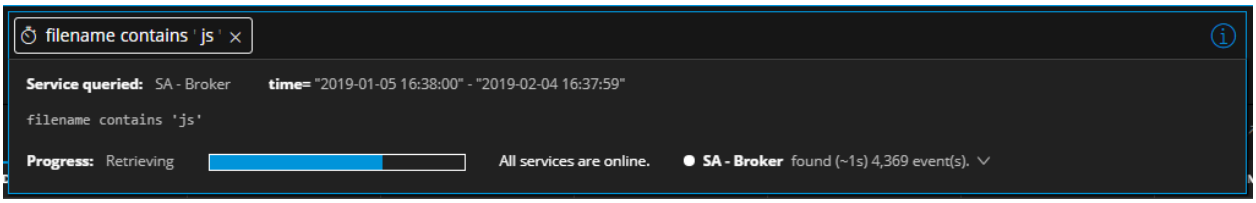
While a query is executing, a progress bar indicates the query's completion percentage at the bottom of the console. The status lets you know details about what is happening; for example, you can tell when the query is executing, queued, reading the index file for the queried service, retrieving events, and complete. All statuses and non-fatal messages are displayed as they come in, and the border color of the query bar changes to amber if a non-fatal error occurs.

Icons provide additional information about individual services.

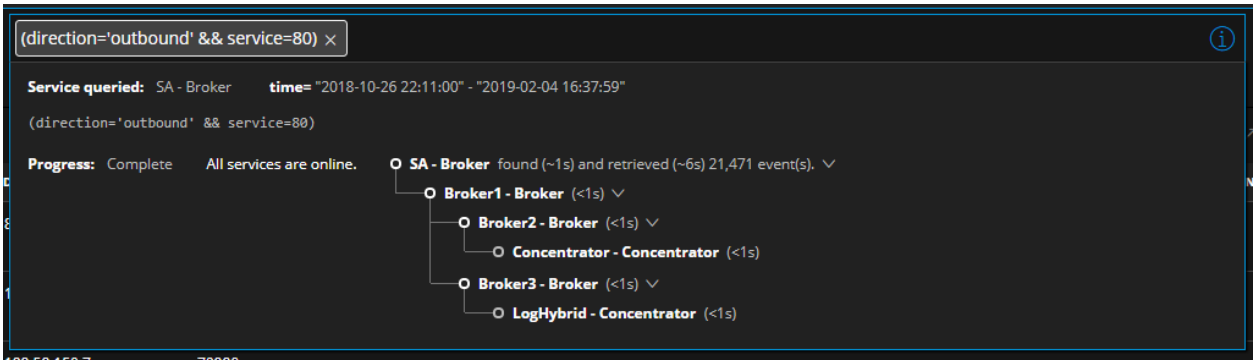
- An amber stopwatch marks the slowest service.
- An amber triangle indicates a warning was received.
- A red triangle shows that an error was received when trying to query the service.

Executing and Reading the Index File to Find Events. The first stage of a query is complete when the queried services have found results. The query console provides a nested hierarchical listing of all the services being queried with indicators showing which are online or offline, and the time in seconds that the service took to find results.

Retrieving Events and Loading in the Events Panel. While the found events are being retrieved and loading in the Events panel, the progress bar shows a visual indicator and text description of what is happening. In the figure below, results were found and are being retrieved.

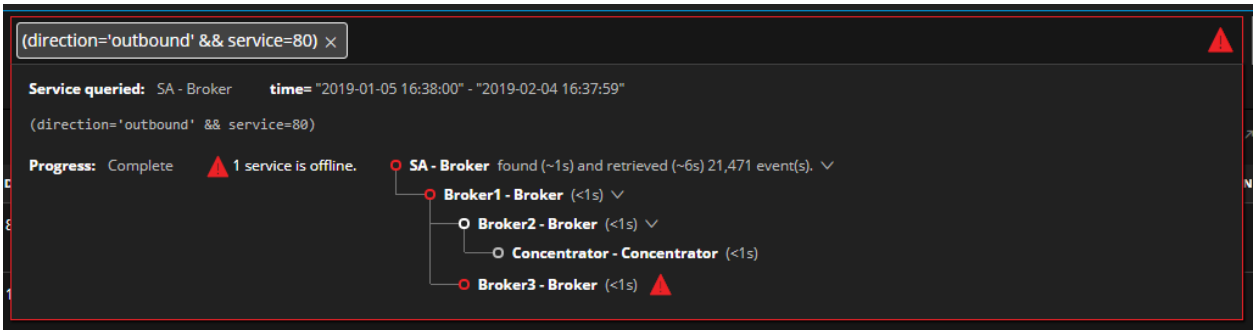


Request Complete. If there are no errors or warnings when loading is complete, the query console is outlined in blue and the button is disabled as an indicator that the data in the view is fresh. The following figure is an example of the query console for a completed query with no errors or warnings.



Errors and Warnings. A fatal error such as a syntax error in the query, or the queried service being offline, stops execution of the query. A red triangle is displayed in the upper right corner of the query console, and the console is outlined in red to indicate that the query failed. If the queried service is offline, only the queried service with no hierarchy of services is listed in the query console and marked by a red triangle.

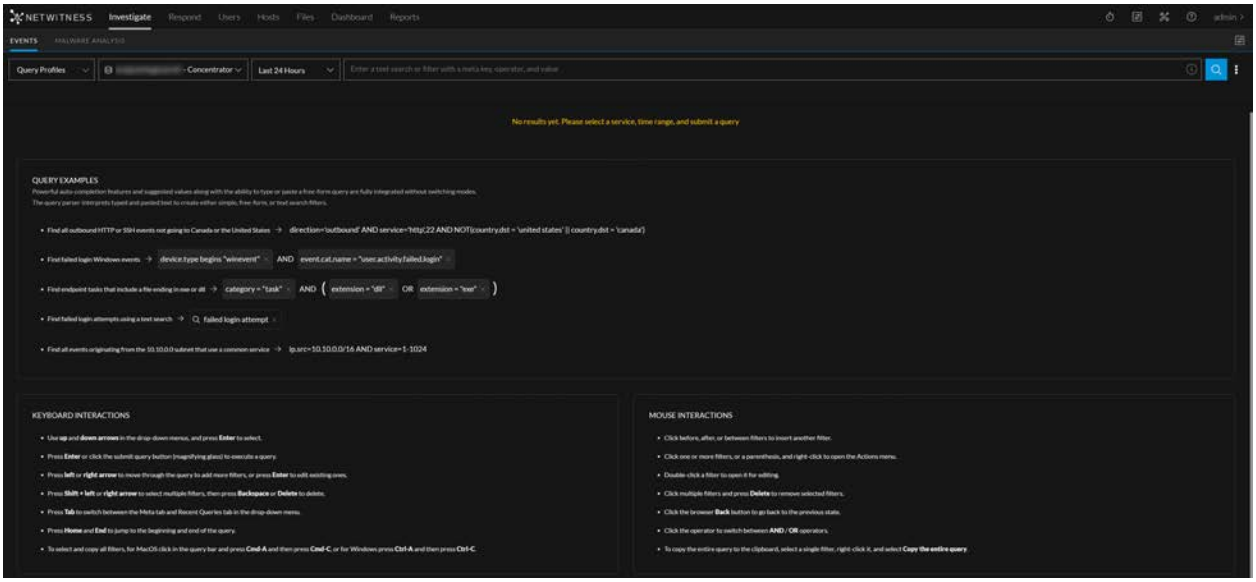
A non-fatal error does not prevent a query from executing. The query is executed and events are loaded, but a red triangle is displayed in the upper right corner of the query console, and the console is outlined in red as a warning. The following figure shows the appearance of the query console when the queried service proxies to another service that is offline.



A warning does not prevent a query from executing. The query is executed and events are loaded, but an amber triangle is displayed in the upper right corner of the query console, and the console is outlined in amber.

Build a Query in Guided Mode

Guided Mode is the easiest way to create a query with features to help analysts enter valid queries. The following figure illustrates the initial Events view with Guided Mode in effect in the query bar.



There is no longer a need to select Guided Mode because all Guided Mode features and more improvements are built in to Version 11.4.1. This figure depicts the query bar in Version 11.4.1.





Keyboard Actions to Use in Guided Mode

In Guided Mode, the query builder allows entry, editing, and deletion of filters using the key strokes without having to use a pointer. Although you can use the pointer, you have the option to keep your fingers on the keyboard. This table identifies the available keyboard actions in Guided Mode when the cursor is located in the query bar; these do not apply to the service selector and time range.

Action	Keyboard Entry
Copy all filters (Version 11.4.1 and later)	With the cursor in the query bar, but not in a filter being edited, and all filters selected, press **Ctrl-C** (Windows OS) or **Cmd-C** (MacOS).
Delete characters in a filter	Selected characters: With characters selected in the query bar, press **Delete** or **Backspace**. Previous character (Version 11.4 and later): With the cursor next to a character in the query bar, press **Backspace** (Windows OS) or **Delete** (MacOS). All characters (Version 11.4 and later): With the cursor in a filter, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).
Delete filters	Selected filters: With one or more filters selected do one of the following: • **Right-click > Delete selected filters** or **Delete selection** (11.4 and later). • Press **Delete**. • Press **Backspace**. Filter that has focus (Version 11.4 and later): With the cursor in a filter that has focus, press **Backspace** (Windows OS) or **Delete** (MacOS). The focused filter is deleted and focus moves to the left. Filter that has focus (Version 11.4 and later): With the cursor in a filter that has focus, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS). The focused filter is deleted and the focus moves to the right.
Delete parentheses in a filter, do not delete the contents (Version 11.4 and later)	With a set of parentheses, but not the contents selected, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS). The selected parentheses are deleted, but the contents of the parentheses remain.

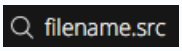
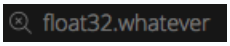
Action	Keyboard Entry
Delete parentheses and their contents in a filter (11.4 and later)	Selected parentheses: With a set of parentheses selected, do one of the following: • **Right-click > Delete selection**. • Press **Backspace** (Windows OS) or **Delete** (MacOS). The selected parentheses and contents are deleted and the focus moves to the left. • Press **Delete** (Windows OS) or **Fn + Delete** (MacOS). The selected parentheses and contents are deleted and the focus moves to the right.
Deselect all filters	With a filter selected, press **Esc**.
Edit a selected filter	With a single filter selected, press **Enter**.
Insert a new filter at the beginning of the query bar, and open for editing (Version 11.4 and later)	With a filter selected, press **Home** (Windows OS) or **Fn + Left Arrow** (MacOS).
Insert a new filter at the end of the query bar, and open for editing (Version 11.4 and later)	With a filter selected, press **End** (Windows OS) or **Fn + Right Arrow** (MacOS).
Insert a new filter to the immediate left of the selected filter, and open for editing	With a filter selected, press **Shift + Left Arrow**.
Insert a new filter to the immediate right of the selected filter, and open for editing.	With a filter selected, press **Shift + Right Arrow**.
Insert a new filter to the immediate left of the selected filter	With a filter selected, press the **Left Arrow**.
Insert a new filter to the immediate right of the selected filter	With a filter selected, press the **Right Arrow**.
Open a new tab with the selected filters	With filters selected, **right-click > Query with selected filters in a new tab**.
Query with the selected filters	With filters selected, **right-click > Query with selected filters**.

Action	Keyboard Entry
Query with content of parentheses (Version 11.4 and later)	With parentheses selected: • To query with the selected parentheses contents, select one side of a parentheses set and **right-click > Query with selected filters**. • To query with the selected parentheses contents in a new browser tab, select one side of a parentheses set and **right-click > Query with selected filters in new tab**.
Select all filters in the query bar (Version 11.4.1 and later)	With the cursor in the query bar, but not in a filter being edited, press **Ctrl-A** (Windows OS) or **Cmd-A** (MacOS).
Select all filters to the left of the current filter	(Version 11.3.x and earlier) With a filter selected, press **Shift + Up Arrow**. (Version 11.4 and later) With a filter selected, press **Shift + Right Arrow** twice.
Select all filters to the right of the current filter	(Version 11.3.x and earlier) With a filter selected, press **Shift + Down Arrow**. (Version 11.4 and later) With a filter selected, press **Shift + Right Arrow** twice.
Select the filter to the immediate left if one exists	With no filter selected, press the **Left Arrow** key.
Select the filter to the immediate right if one exists	With no filter selected, press the **Right Arrow** key.
Submit a query.	With focus on the query bar and no pending filters, press **Enter**.

Visual Feedback in Guided Mode

Guided Mode provides visual feedback during query construction. This table identifies and describes the possible feedback.

Feedback	Icon	Description
Blue background on a Filter		Indicates that a filter is selected.
Green circle between two filters		(Version 11.3 and earlier) A green circle indicates the location of the cursor between two existing filters. Clicking inserts a new filter at this location. (Version 11.4) A bold cursor indicates the insertion point.

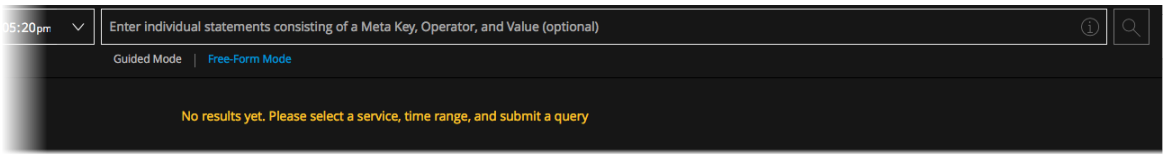
Feedback	Icon	Description
Green filter outline		Marks the single filter that has focus and ready to edit. This is combined with the blue background, when multiple filters are selected and this filter has focus.
Red filter outline		Indicates that the filter is invalid. A tool tip that explains the error is displayed.
Index indicators in the Meta tab		(Version 11.4 and later) Indicate the index level of the meta keys in the Meta tab, which determines if you can use it in a filter:  This meta key is indexed by meta value and can be used in a filter.  This meta key is indexed by meta key, and can be used in a filter.  This meta key is not indexed, and not selectable for a filter. The `sessionID` meta key is a special case. Unlike other non-indexed meta keys, it is not configurable, but you can use it in a filter so it is marked by the key symbol. Supported operators are `exists`, `!exists`, `=`, and `!=`.

Feedback	Icon	Description
Query Events button		Used to submit a query, show the status of the query, and cancel a query. The button has three possible states:  Ready to submit a query using filters in the Query Builder.  Waiting for server validation to complete before executing the query.  The query is executing, click to cancel execution.
Slow Service icon		In the query console, marks the service that took the longest time to load results from the query.
Spinner in the Events list		Indicates that the query is currently being processed. The **Query Events** button is disabled while this occurs.
Stopwatch		(Version 11.5 or earlier) Indicates that the meta key/operator combination requires extra time to process. While the query is still executable, a more efficient meta key or operator is recommended.

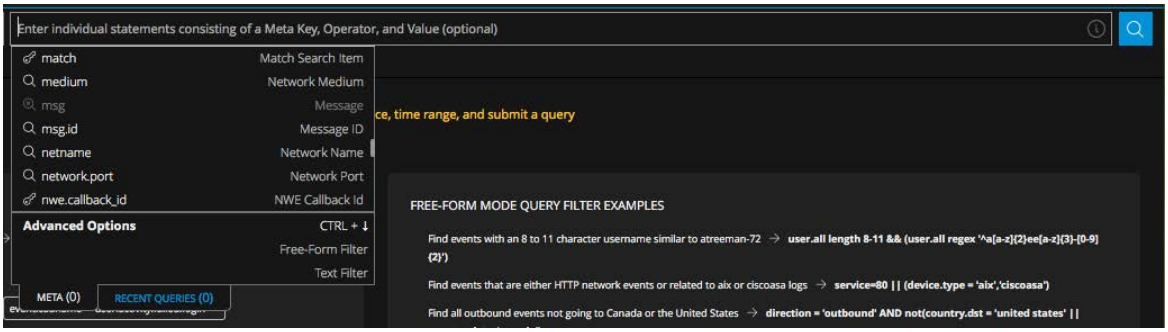
Add a Simple Filter in Guided Mode

To create a simple filter in Guided Mode:

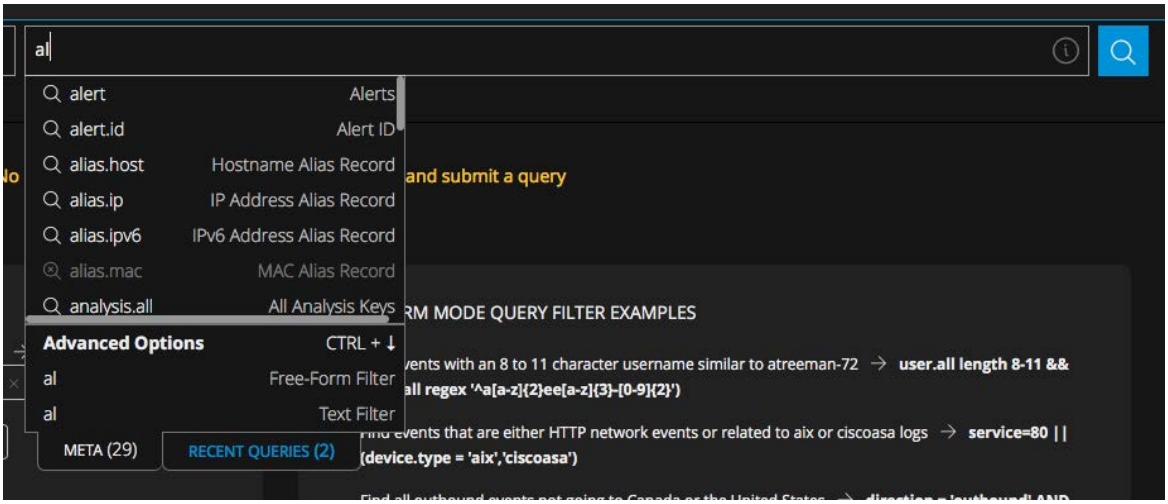
1. Go to the **Events** view (Event Analysis view in Version 11.3 and earlier) and do one of the following:
   a. (Version 11.4.1 and later) Click in the query bar and when the filter entry form is displayed, select the Meta tab if it is not already selected.
   b. (Version 11.4 and later) Select **Guided Mode**, click in the query bar and when the filter entry form is displayed, select the Meta tab if it is not already selected.
   c. (Version 11.2 and later) Select **Guided Mode** and click in the query bar.
   d. (Version 11.1) Click in the empty query bar, or before or after an existing filter. This is an example of the empty query bar in Guided Mode before you begin entering a filter.



If the insertion point is between two filters, a green circle (Version 11.3 and earlier) or a bold cursor (Version 11.4 and later) marks the insertion point. If the insertion point is at the end of the query bar, the filter entry form opens with a blinking cursor at the entry point. A drop-down list displays the available meta keys passed from the service being investigated in alphabetical order. This figure shows the filter entry form from Version 11.4.



2. To select a meta key do one of the following:
   a. If there is only one option in the drop-down list, press **Enter**.
   b. If there are two or more options in the drop-down list, click a meta key or select a meta key using the up/down arrows, then press **Enter**.
   c. Start typing the meta key. As you type the meta key, the list is filtered to include only meta keys that contain the text you typed. The count next to the label on the Meta (0) tab increments to enumerate the indexed meta keys that match the typed text. Keys that are not indexed are disabled and not selectable and are not included in the count, for example, `alias.mac` in the figure below is not indexed and is dimmed. Click a meta key or select a meta key using the up/down arrow, then press **Enter**.



d. To select a highlighted meta key, press **Enter**.
The count on the Meta label changes to 1.

> **Note:** If no meta key in the drop-down list is selected, and the list has no meta keys to select, either the Free-Form Filter or the Text Filter option is highlighted based on the content already typed in the query bar.
> --If the text typed in the query bar includes some form of query syntax and other operators not yet supported by the user interface, the Free-Form Filter option is highlighted and you can create a free-form filter. In Version 11.3 and earlier, the `**`, `&&`, `||`, `()`, `AND`, `OR`, `comma`, `-`, `length`, and `regex` operators are not supported by the user interface. The Version 11.4 user interface supports these operators. If the Free-Form Filter is not highlighted, and the query bar has no text filter, the Text Filter is highlighted so you can create one.
> --If the first condition is true, and there is already one text filter, the Free-Form Filter option is highlighted so you can create a free-form filter.

e. If you want to edit or delete the meta key, press **Backspace** or **Delete**.
As you backspace and delete characters, the meta key drop-down list is filtered to include meta keys that contain those characters. To select a meta key, press **Enter**.
The meta key is added to the filter entry form, and a list of valid operators for the selected meta key is displayed. Operations that require more time to process are marked by a (stopwatch icon). This figure shows the stopwatch icon marking the `contains` operator.



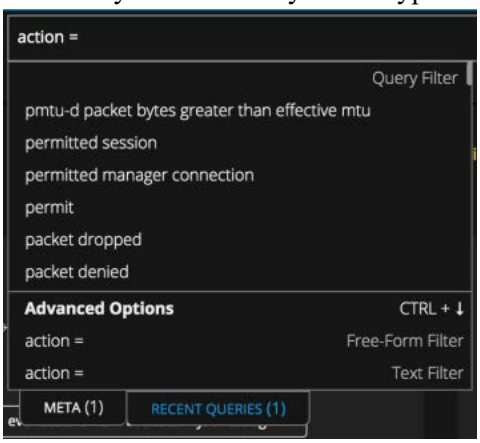
3. To select an operator, do one of the following:

a. If there is only one option in the operators drop-down list, press **Enter** to select it.

b. If there are two or more options in the operators drop-down list, click an operator or select one using the up/down arrows, then press **Enter**.

c. Type the operator and press **Enter**. As you type, the operators drop-down list is filtered to show only operators that contain the typed text. Click an operator or select one using the up/down arrows, then press **Enter**.
The operator is added to the filter entry form. In Version 11.4 and later, if the operator accepts a value, the suggested values drop-down list is displayed. Earlier versions leave the cursor in the

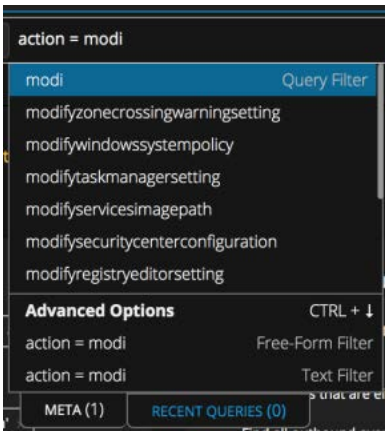
filter entry form so that you can type a value.



4. (Optional) If the selected operator in the filter entry form accepts a value, do one of the following:
   a. In Version 11.3 and earlier, type the value and press **Enter**.
   b. In Version 11.4 and later, paste a value that you have copied from somewhere and press **Enter**.
   c. In Version 11.4 and later, begin typing in the **Query Filter** field.
   As you type, the meta value drop-down list is filtered to return up to 100 properly indexed values that begin with the typed text. The suggested values are based solely on the time range; filters in the query do not filter the list of 100. The auto-suggest function looks for matches in all events in the current data set, not just the (up to 10,000) downloaded events. If nothing in the list matches exactly, the text you typed in the Query Filter field is highlighted and this message tells you that no suggestions were found. Some values, such as the integers for the `service` meta key, also display the definition of the service type.



   If there is an exact match, that value is highlighted. In the following example, there is no exact match for the typed text, modi.

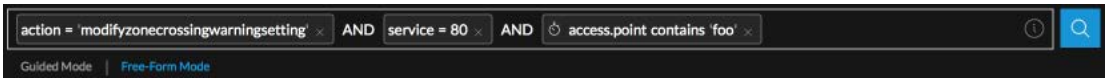


i. If the typed text is the value you want to use in the filter, press **Enter**.

ii. If you see the value that you want to query in the list and it is not highlighted, click the value or use the up/down arrows to highlight the value. Then press **Enter**.

iii. If you want to edit or delete the value, press **Backspace** or **Delete**.
As you backspace and delete a character, the meta value drop-down list is filtered to include values that begin with the remaining characters. To select a value, press **Enter**.
The value is added to the filter entry form.

5. To create the filter, press **Enter**. If you click anywhere outside the box before pressing **Enter**, the filter is not created.
The new filter is inserted, and the blinking cursor is refocused after the last filter, the meta keys drop-down list is displayed. If there is an error in the filter, it is outlined in red. You can hover over the filter to see a tool tip explaining the error. This figure shows a query being created with no errors.



6. If the filters have no errors, you are ready to execute the query in the query bar. Click .
The results are returned and loaded in the Events panel. The first 10,000 events that match the query begin loading in the Events panel. As the events are loaded, a status bar at the top tracks progress and you can scroll to the bottom of the list to see the completion status.

7. (Optional in Version 11.3 and later) If you want to see detailed status in the Query Console, click the information icon .

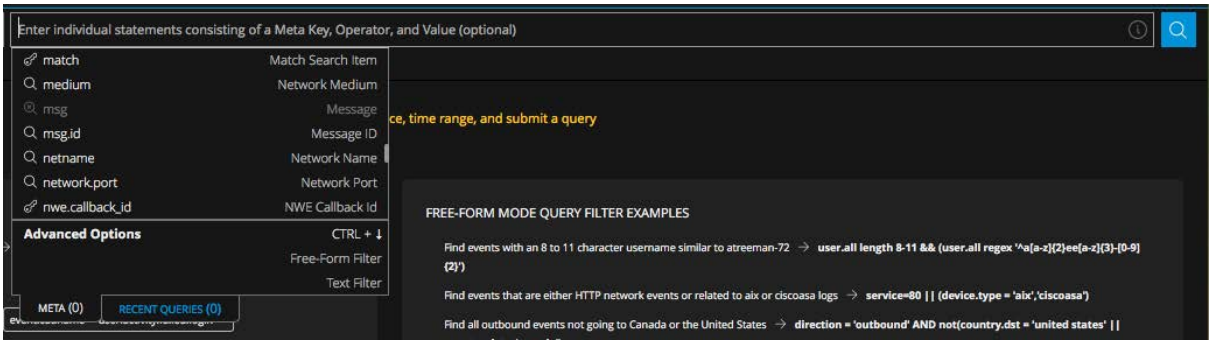
8. (Optional in Version 11.3 and later) If you want to cancel the query before it finishes executing, click .
The query stops executing and a notification that the query has been canceled is displayed.

Add a Free-Form Filter in Guided Mode (Version 11.3 and Later)

To filter the data displayed in the Events view using a free-form filter in Guided Mode:

1. Go to the **Events** view, select **Guided Mode** below the query bar, and click in the query builder field. (For Version 11.4.1, simply click in the query builder field.)
If the insertion point is between two filters, a green circle or a bold cursor marks the insertion point.

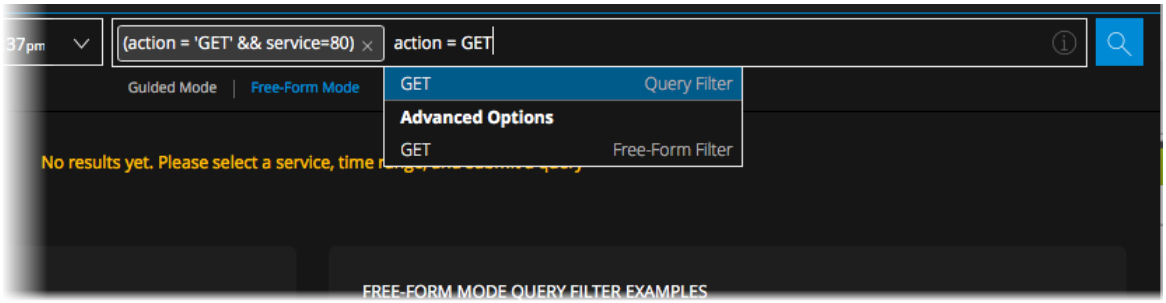
If the insertion point is at the end of the query bar, the filter entry field opens with a blinking cursor at the entry point. A drop-down menu lists available meta keys passed from the service being investigated in alphabetical order.



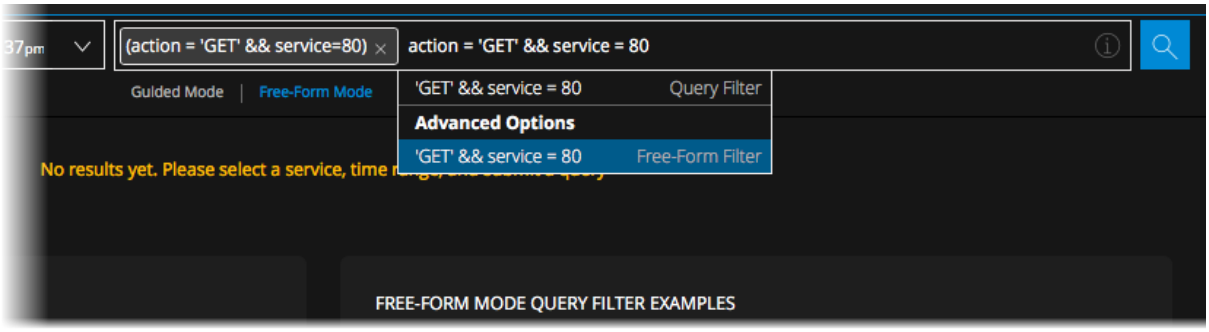
2. Do one of the following:

   a. Place the cursor in the **Free-Form Filter** field and begin typing the query.

   b. Begin typing the filter beginning with a meta key or with an open parenthesis. When entering and editing filters in the query builder, parentheses pairs are automatically balanced. If you type an open parenthesis, the other part of the pair is added to the filter.
   When no matching meta keys or operators are available in the drop-down menu, the Free-Form Filter option becomes available, and the text you typed is available in the **Free-Form Filter** field.

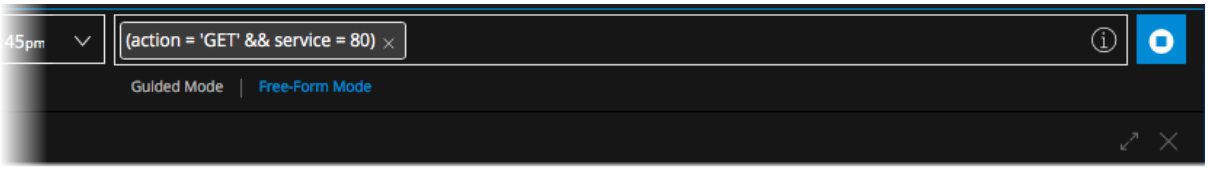


3. Continue typing the entire expression and press **Enter**.
   (If you click anywhere outside the box before pressing **Enter**, the filter is not created.) This figure shows a free-form expression created by continuing to type after the value GET.



   The new filter is inserted, and the blinking cursor is refocused after the last filter, a new filter entry form is displayed. If there is an error in the filter, it is outlined in red. You can hover over the filter to see a tool tip explaining the error.

4. To execute the query, click . While the query is executing, the button changes to .



5. If you want to cancel the query before it finishes executing, click .
   If you do not cancel the query, you can click to view the status of query execution. When the query is finished executing, the Events panel displays appropriate results for the query.

Add a Text Filter to Find a Value Anywhere in the Data Set (Version 11.4 and Later)

In Version 11.4 and later, the text filter allows you to find a specific value in the current data set (endpoint, logs, and network events). The text filter initiates a case-insensitive search against all the data for meta keys that are indexed by value. The text filter does not search for values that are indexed by meta key or not indexed so you not see all results. A message advises that `Results may be limited by a text filter, which matches only indexed meta keys. If you want to conduct a more exhaustive search against raw events, click here and choose the appropriate options in the Search Events drop-down menu.` Icons in the drop-down list indicate the index level of each meta key:

- filename.size - indexed by meta key
- filename.src - indexed by meta value
- float32.whatever - not indexed

Note: All services in the hierarchy being queried (Broker, Concentrators, and Decoders) must be at Version 11.3 or later. The text filter is not available in the drop-down menu when there are services below Version 11.3 in the hierarchy.

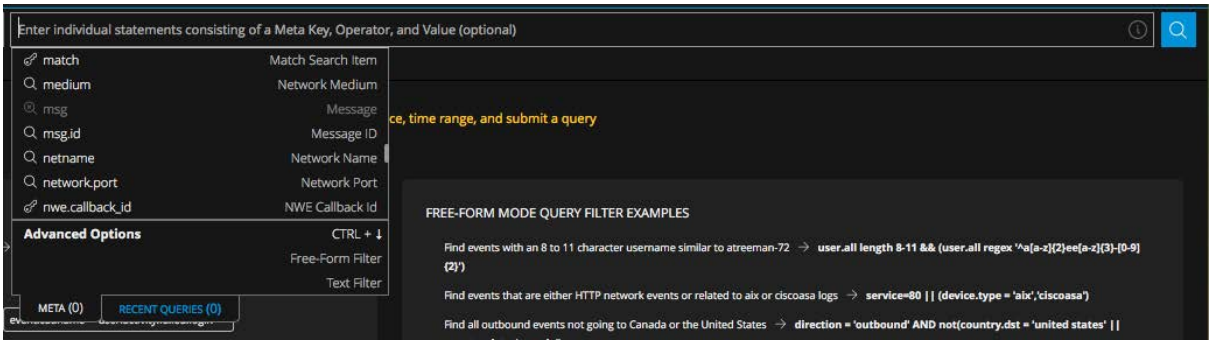
The text filter is useful when you have some idea of what you are looking for, but are not sure where to look (which meta key or service). As an example, if you are interested in looking for a file name, click in the query bar, type the complete text string, and click **Text Filter**. The text filter initiates a search against all the data in the index, within the services and time range being investigated, and returns exact matches to the text string.

A query can include one text filter and any combination of simple and free-form filters. The operator for a text filter must be AND because it acts as a filter over the results of all the other filters in the entire query. If one text filter already exists in the query bar, the Text Filter option is disabled as shown in the figure below. Text filters are not supported inside parentheses.

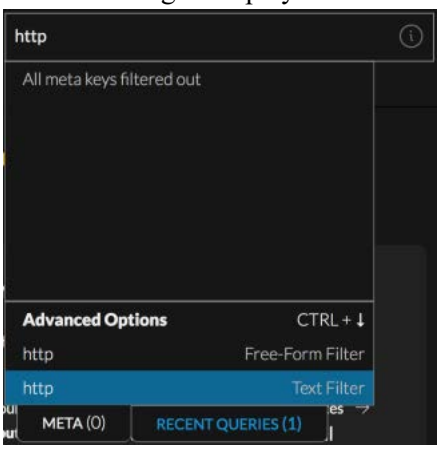


To create a text filter:

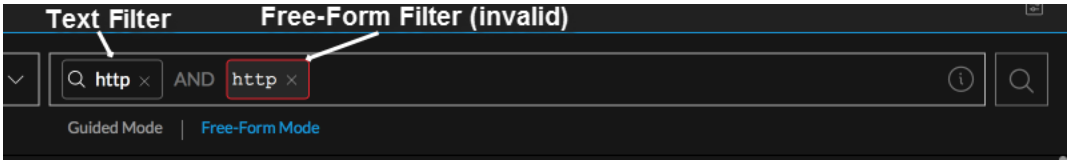
1. Go to the **Events** view and click in the query bar.
   The query entry form is displayed.



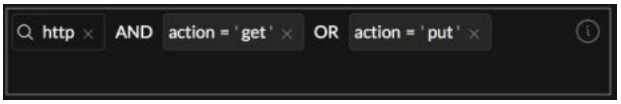
2. Type the text string that you want to find, for example, `http`.
   The text string is displayed in the meta key drop-down list under Advanced Options.



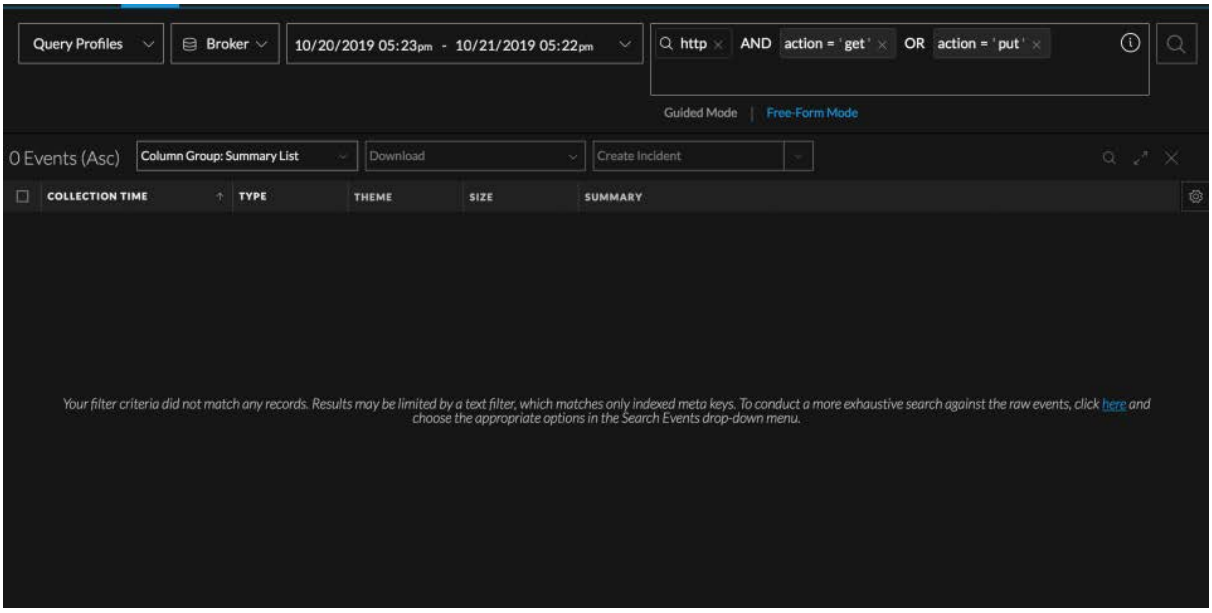
3. Click **Text Filter** under **Advanced Options**.
   The text filter is created in the query bar. The following figure illustrates the different appearance of a text search filter versus a free-form filter. The free-form filter is in a fixed-space font and outlined in red. The red outline indicates a syntax error because a valid expression is expected in a free-form filter. The text filter is marked by the search icon. No syntax requirements are applied to text search filters.



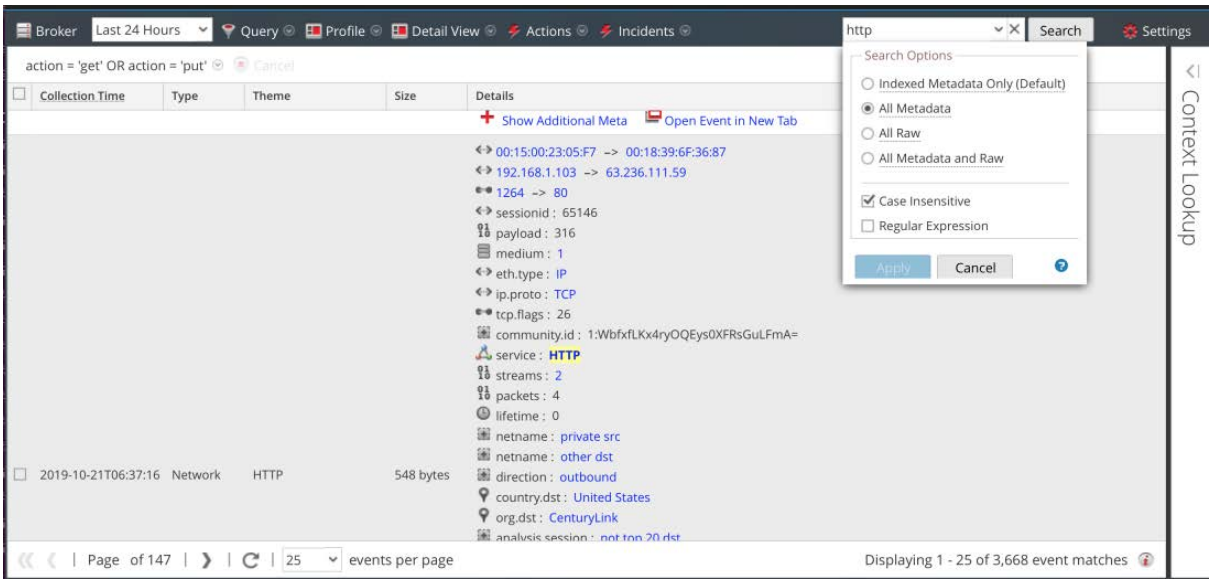
4. (Optional) Create additional simple or free-form filters in the query bar. There can be only one text filter in the query. This example was created by typing `http` as a text filter and then adding two more filters - `action = 'get' OR action = 'put'`



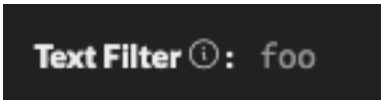
5. To submit the query, click .
   The results are displayed in the Events panel. This figure illustrates the Events panel with no results displayed and a message with instructions for improving results. Every time you use a text filter, this message is at the bottom of the results offering a link to expand your search.



6. Click the **here** link in the message.
   A new browser tab opens with the query results displayed in the Legacy Events view, where you have additional options to improve the search. This figure shows the results for the same query when non-indexed metadata is included.



7. Click the Information icon in the query console to view the status of the query. This figure shows a text filter in the query console.



Select All Filters and Copy All Filters in the Query Bar (Version 11.4.1 and Later)

While creating a filter in the Events view query bar, you can use keyboard commands to select all filters (Ctrl-A for Windows OS, Cmd-A for MacOS) and then copy the selection to the local clipboard (Ctrl-C for Windows OS, Cmd-C for MacOS).

To select all filters and copy them to the clipboard:

1. In the **Events** view > **Events** panel, click on a focused pill or in the query entry form, and press **Ctrl-A** for Windows OS or **Cmd-A** for MacOS.
   All filters in the query bar are selected.
2. To copy the selected filters to the clipboard, type **Ctrl-C** for Windows OS or **Cmd-C** for MacOS.
   You can share the clipboard with other analysts or paste the contents in the query bar.

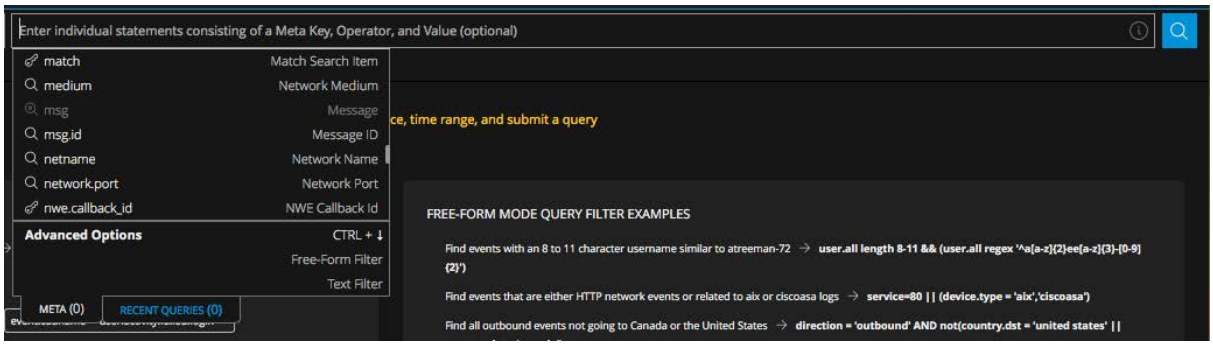
Paste Text in the Query Bar (Version 11.4 and Later)

While creating a filter in the Events view query bar, you can paste instead of typing the complete text of a filter that you have copied from somewhere else. You can paste the text into an empty query bar or next to an existing filter in the query bar. Depending on the text you typed, the query parsing engine parses the information that you pasted and creates a new filter, which can be a simple filter, a free-form filter, or a text filter.

- A text string of this form is added as a new simple filter in the query bar: `<valid meta key> <valid operator> <optional value>`. This is an example: `alias.host contains 's'`.
- A text string of this form is added as two simple filters in the query bar: `<valid meta key> <valid operator> <optional value> && <valid meta key> <valid operator> <optional value>`. This is an example: `alias.host contains 's' && action exists`, which is converted to `alias.host contains 's' AND action exists`.
- A text string that contains unparsable text may be converted to a free-form filter. For example, using `NOT (device.ip = 10.10.10.10)` is unsupported for creation of a filter in Guided Mode, so this would be converted to a free-form filter. Free-form filters are validated by the server when they are submitted.
- Text that does not conform to the filter syntax is added as a free-form filter.

To create a filter by pasting text:

1. Go to the **Events** view > **Events** panel, select **Guided Mode** under the query bar, and click in the query bar. (For Version 11.4.1, simply click in the query bar.)
   The query entry form is displayed.



2. Type **Ctrl-V** (Windows OS), **Cmd-V** (MacOS), or **right-click** and **Paste** to paste text that you have copied into the clipboard from somewhere else. Do one of the following:

   a. If the text you pasted is a statement that can be parsed, one or more simple filters is created.
   If the text you pasted is a statement that cannot be parsed, a new free-form filter is created.
   If the text you pasted is not a statement and not a valid meta key, an invalid syntax error is displayed.
   If you pasted a valid meta key for a new filter you are building, the meta key is highlighted in the drop-down list, and you can continue creating a filter as usual by entering an operator and a value.
   After you select a valid meta key and a valid operator (for example, `city.dst =`) any text that you paste is treated as a text string if the meta key supports a text value, and one filter is created.
   If the meta key does not support a text value all of the text in the query bar is parsed as described in step a above.

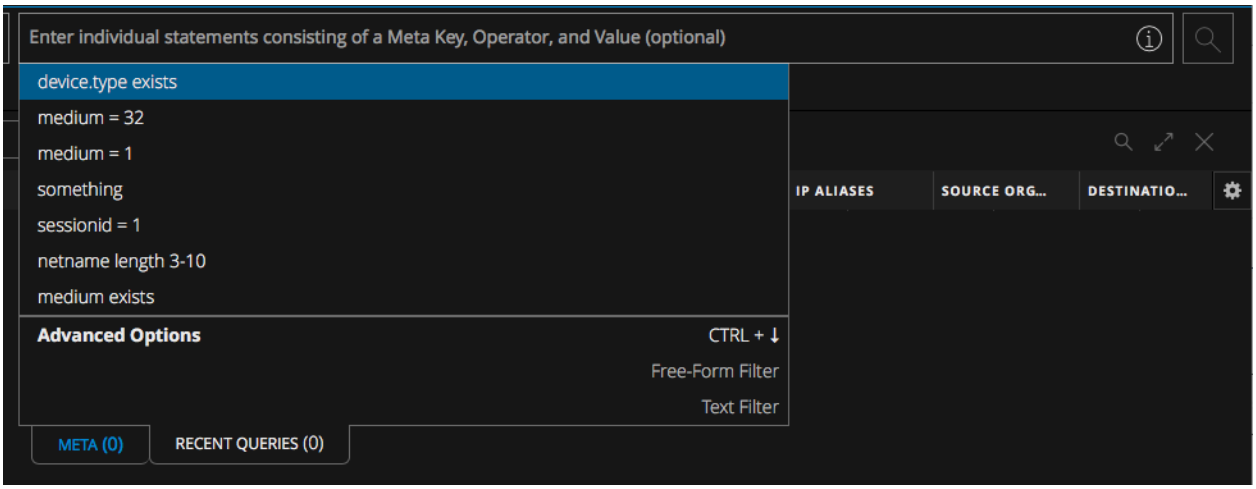
3. Add more filters in the query bar if you wish, and then submit the query.
   The query is executed.

Insert a Filter Based on a Recent Query (Version 11.4 and Later)

In the Guided Mode query bar, you can insert a filter based on a recent query. When the Recent Queries tab is opened and nothing has been typed in the query bar, up to 100 of your most recently executed queries are displayed in a scrollable list. The list is sorted to show the most recent at the top, and the Recent Query count is set to 0. When you begin typing, the list is filtered to display up to 100 queries from the query history database that contain matching text, even if the matches are not in the most recent 100 queries. The Recent Query count changes to reflect the number of matching queries as you type.

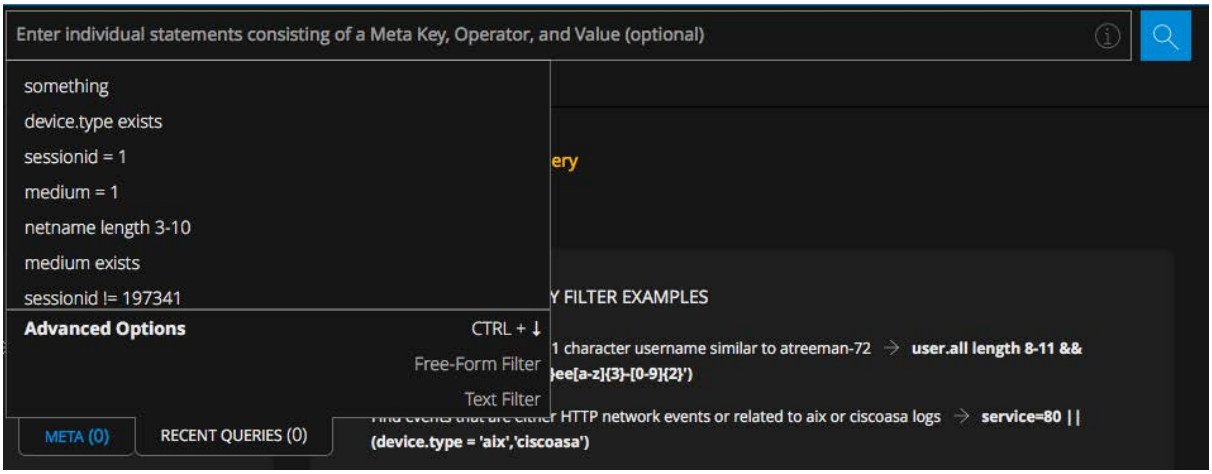
The top entry in the list is highlighted by default. To select a recent query, you can move the highlighting up and down in the list using the up and down arrow or by mouse-over of a recent query. As you type the list is filtered and the highlighting moves back to the top of the list. Clicking a query, or pressing Enter while a query is highlighted, creates a new filter with the text of the selected query.

Whenever you submit a query, the list is sorted to add that query, now the most recent, at the top.

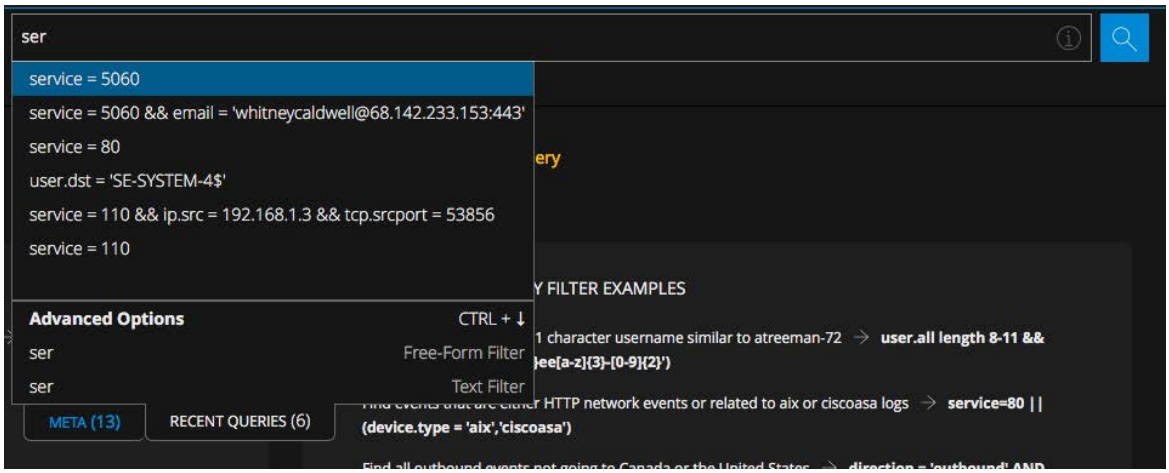


To create a filter based on a recent query:

1. Go to the **Events** view, select **Guided Mode** under the query bar, and click in the query bar. (For Version 11.4.1, simply click in the query bar.)
   The Meta Key drop-down list is displayed in the Meta tab.
2. Select the **Recent Queries** tab.
   The Recent Queries drop-down list is displayed with a count of 0.



3. To search for a recent query, do one of the following:
   a. Begin typing some text.
      As you type more characters or backspace to delete characters, the list is filtered to show recent queries that contain the text you typed. The count in the Recent Queries label increments to show the number of matching queries as you type.



b. To select a query and add a new filter, continue to type and use the up and down arrows until the query you want to use as a new filter is highlighted.

c. With a query highlighted, press **Enter** or simply click a query that you see in the list.
The filter is added in the query bar.

4. Add more filters in the query bar if you wish, and then submit the query.
The query is executed and the list is sorted to add that query, now the most recent, at the top.

Edit a Filter in Guided Mode

With a query in the Guided Mode query bar, you can edit a filter. To edit a filter:

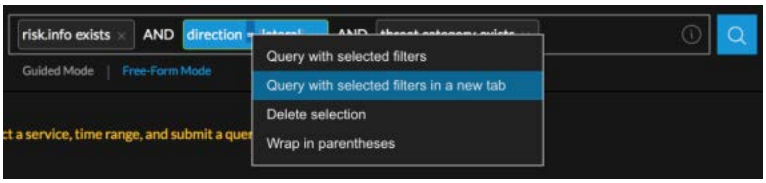
1. Double-click the filter, or click the filter and press **Enter**.
2. Edit the filter. When finished editing, press **Enter** to update the filter.
3. If you want to execute the query again, click .
The Events panel displays results for the updated filter.

Query Using Selected Filters in Guided Mode

When you have one or more filters in the query bar in Guided Mode, you can refocus the query to include only selected filters, displaying results in the current browser tab or a new browser tab. Some filters include expressions with nested parentheses in Version 11.4, and you can refocus part of a filter that includes nested parentheses. To update the query using only selected filters, do one of the following:

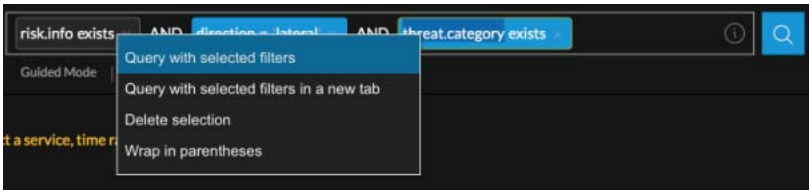
1. Using a query that includes one or more simple filters, for example a query has three filters: `risk.info exists`, `direction ='lateral'`, and `threat.category exists`.

   a. Select `direction = 'lateral'`, right-click the filter and select **Query with selected filters in a new tab** in the drop-down menu.



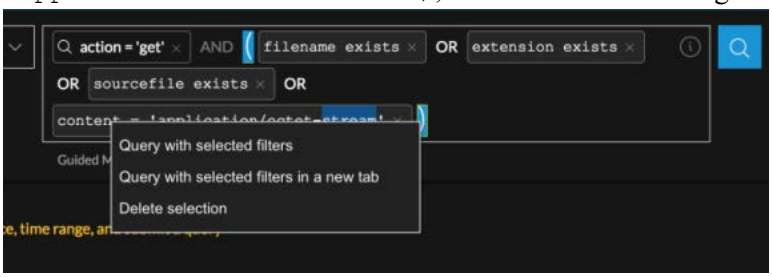
A new tab opens with the results for the selected filter and the original query is left intact on the previous tab.

b. To query the selected filters in the same tab, select `direction = "lateral"` and `threat.category exists`. Then right-click and select **Query with selected filters** in the drop-down menu.



A query with only the selected filters is submitted and all remaining filters are removed.

2. (Version 11.4) For a query that includes a filter containing nested parentheses, for example: `action = 'get' AND (filename exists OR sourcefile exists OR content = 'application/octet-stream')`, do one of the following:



a. Select the close parenthesis after `'application/octet-stream'`, right-click, and select **Query with selected filters in a new tab**.
A new tab opens with results for `(filename exists OR sourcefile exists OR content = 'application/octet-stream')`.

b. Select the same, right-click, and select **Query with selected filters**.
The results for `(filename exists OR sourcefile exists OR content = 'application/octet-stream')` are displayed in the current tab.

Delete a Filter and Delete Text or Parentheses in a Filter in Guided Mode

Some keystroke editing features became available in Version 11.4; these are labeled in the steps.

1. To delete a filter, do any of the following:
   a. Click **X** in a filter.
   b. Select the filter and press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).
   c. (Version 11.4 and later) Select the filter and press **Backspace** (Windows OS) or **Delete** (MacOS).

d. Right-click one or more filters and select **Delete selected filters** or **Delete selection** (Version 11.4 and later) in the drop-down menu.
The filter and the operator to the right or left of the filter is deleted, ensuring that no extraneous operators remain in the query bar.

2. (Version 11.4 and later) To delete characters in a filter or parentheses and contents in a filter, do any of the following:

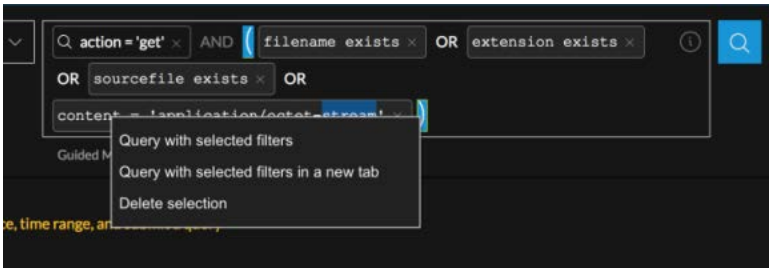
a. To delete the previous character: With the cursor next to a character in the query bar, press **Backspace** (Windows OS) or **Delete** (MacOS).

b. To delete all characters: With the cursor in a filter, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).

c. To delete the selected characters: With characters selected in the query bar, press **Delete** or **Backspace**.

d. To delete parentheses, but not the characters inside the parentheses, select one of the parentheses and press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).

e. To delete a set of parentheses and the contents, for example, `(filename exists OR sourcefile exists OR content = 'application/octet-stream')`, select the parenthesis after `get`, right-click, and select **Delete selection**.

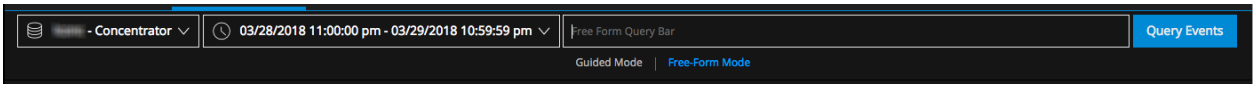


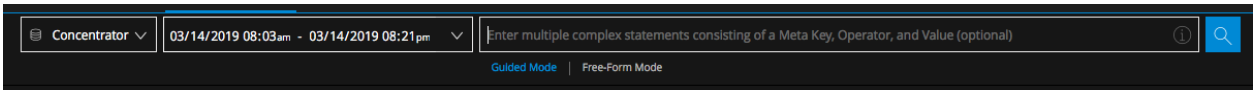
Everything except `action = 'get'` is deleted.

Create a Query in the Free-Form Mode

Free-Form Mode is used in Version 11.2, 11.3, and 11.4, but no longer available in Version 11.4.1.

Free-form queries are most useful when you have a long text string saved that you want to paste, or if you have one in mind that you want to enter quickly, and you know the meta keys, valid operators, and valid syntax for entering values. The following figure illustrates the initial Events view with the empty Free-Form query builder field. The first example is Version 11.2 and the second example is Version 11.3.





The blinking cursor indicates that you can enter a query. You can enter free text here. As more expressions are added and they cannot be displayed in a single line, they wrap to another line and the input area expands vertically so that all filters are visible without scrolling to the right.

These are some examples of queries that you can enter in Free-Form mode:

To find events with an 8- to 11- character username similar to atreeman-72:
`user.all length 8-11 && (user.all regex '^a[a-z]{2}ee[a-z]{3}-[0-9]{2}')`

To find events that are either HTTP network events or related to aix or ciscoasa logs:
`service=80 || (device.type = 'aix','ciscoasa')`

To find all outbound events not going to Canada or the United States:
`direction = 'outbound' AND not(country.dst = 'united states' || country.dst = 'canada')`

If you have a submitted query in Guided Mode, the query is transformed into text when you click switch to Free-Form mode. This is an example of a query submitted in Guided Mode as two filters, `service = 80` and `direction = 'outbound'`, and then viewed in Free-Form mode.



The button on the right side of the query builder is visible as needed to input a query. The query is applied when you click . At that time the query is validated to show syntax and logic errors.

Operations that require more processing time are not highlighted as they are in Guided Mode, but this table provides a summary of expensive operations for reference.

Index Method	Non-Text Value	Text Value	Regular Operations	Expensive Operations
By Key	✓		`exists`, `!exists`	`eq`, `!eq`
By Key		✓	`exists`, `!exists`	`eq`, `!eq`, `begins`, `ends`, `contains`
By Value	✓		`exists`, `!exists`, `eq`, `!eq`	no expensive operators
By Value		✓	`exists`, `!exists`, `eq`, `!eq`, `begins`	`ends`, `contains`
By None	special case for `sessionid`		`exist`, `!exits`, `eq`, `!eq`	no expensive operators

Filter Results in the Navigate View

When conducting an investigation in the Navigate view, there are several methods available to refine the results displayed when meta key values are loaded in the Navigate view. The rest of this topic is focused on the basic methods of filtering data:

Note: By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.

- Set the Time Range
- Set the Quantification Method and Sort Sequence of Meta Key Results
- Manage and Apply Default Meta Keys in an Investigation
- Drill into Data in the Navigate View Time Chart
- Drill into Data in the Values Panel

Set the Time Range

When conducting an investigation in the Navigate view, the time range options limit the results returned. You can select:

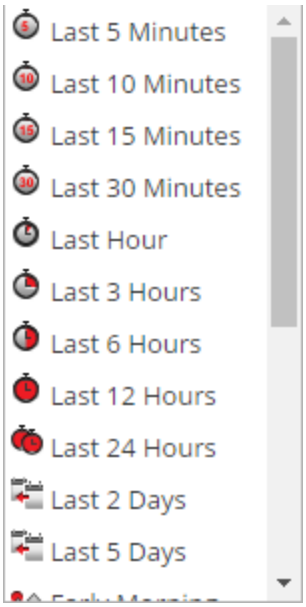
- A time range relative to the collection- ranges relative to the collection are based on the last collection time for data.
- A time range relative to the calendar.
- A custom date range.
- All data.

The selected Date Range is shown in the Navigate view tool bar as the Time Range label. By default the label is **Last 3 Hours**. The Time Range displayed in the timeline banner shows the first and last timestamp for the date range being used for the metadata.

Note: Time range is based on the Time Zone configured in the Profile Preferences panel as described in "Setting User Preferences" in the *NetWitnessGetting Started Guide*.

To select a built-in time range:

1. Click the **Time Range** option in the Navigate view toolbar. The default time range is for the **Last 3 Hours**, but a different value from the selection list, for example, **All Data** or **Last Hour,** may already be selected and used as the label in the options panel.
   The Time Range selection list is displayed.



2. Do one of the following:
   - If you want to see all data, select **All Data**.
   - If you want to set a time range in minutes, hours, or days that is relative to the collection, select a value such as **Last 10 minutes**, **Last 3 Hours**, or **Last 5 days**.
   - If you want to set a time range relative to today, select **Yesterday**, **This Week**(Version 11.1), **Last Week** (Version 11.1), **All Day**, or a part of the day such as **Early Morning**, **Morning**, **Afternoon**, or **Evening**.
   - If you want to set a unique date range, select **Custom** in the **Time Range** menu and follow the procedure below.
     The selected time range is applied to the current results in the Values panel.

To specify a custom time range:

1. Select **Custom** in the **Time Range** menu.
   Date selection options are displayed in the toolbar.



2. Within the time **Start Date** and **End Date** fields, do the following to specify the date and time:
   a. Click a date from the calendar.
   b. (Optional) Select the time from the Hour and minute fields or click **Now**. The time selection defaults to the current time of day.

Note: The value for start time in seconds always defaults to :00, and the value for end time in seconds always defaults to :59. For example, if you are using time to drill down into an issue, the drill time is interpreted as "HH:MM:00 - HH:MM:59."

3. To apply the range, click **Go**.
   The selected time range is applied to the current results in the Values panel.

Set the Quantification Method and Sort Sequence of Meta Key Results

You can select the way results for each meta key are quantified and sequenced in the Navigate view.

Note: If meta entities (Version 11.1 and later) are used in meta groups, the results will show the top 20 values that matched any of the meta keys contained in the meta entity.

Each meta key section in the Navigate view contains an ordered list of values showing each meta key value (Value) and its count (Total). You can specify whether:

- The results in each meta key section are sorted based on Value or Total.
- The results are sorted in ascending or descending order.
- The values shown for each meta key are quantified by number of packets (Packet Count), number of sessions or logs (Quantify by Event Count) or by the size of events (Quantify by Event Size).

Note: If you have both a log decoder and a packet decoder for which you are viewing the metadata, the calculation of what is actually being counted is dependent on the type of key. If you select to Quantify by Packet Count and are looking at logs, the Navigate view output is the same output as if you had selected Quantify by Event Count (see Navigate View for details).

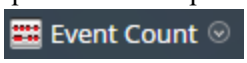
This image shows the `Event Type` meta key presented in order by **Total** in **Descending** order. The value with the greatest count of matches is presented first. The value `failure audit` has 71 matches and is listed first. The value `logon` has only one match and is presented last. The quantification method is **Event Count**.



This image shows the `Event Type` meta keys presented in order by **Value** in **Descending** order. The value names are presented in alphabetical order starting at the end of the alphabet. The value `success audit` is listed first. The value `connect` is presented last. The quantification method is **Event Count**.

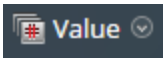


To select the quantification method of meta key count and ordering of meta key results displayed in the Navigate view:

1. In the toolbar, select **Event Count**, **Event Size**, or **Packet Count** and choose one of the quantification options in the drop-down menu. The label for the menu displays the selected option.
   

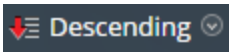
   The current view is reloaded according to your selection.

2. In the toolbar, select **Total** or **Value** and choose one of the ordering methods in the drop-down menu. The label for the menu displays the selected option.



   The current view is reloaded according to your selection.

3. In the toolbar, select **Ascending** or **Descending** and choose one of the sort order options in the drop-down menu. The label for the menu displays the selected option.
   The current view is reloaded according to your selection.



Manage and Apply Default Meta Keys in an Investigation

When analysts are conducting an investigation of captured data in Investigate, a default set of meta keys is loaded and displayed in a default sequence in the Navigate view > Values panel. The default content and sequence is based on the meta keys for the service being investigated. Analysts can specify the meta keys to display during navigation by selecting the default meta keys or by selecting a user-defined group of meta keys, which provides great flexibility to define meta keys. This can help to drill down more directly to the desired data and to reduce the load time by preventing the loading of meta that is not of interest in the current investigation.

Note: In Version 11.1 and later, wherever meta keys are used, you can also use configured meta entities.

If no custom meta groups are in effect, the Navigate view is displayed with the meta key visibility specified in the Default Meta Keys dialog. To optimize loading of meta keys in the Navigate view > Values panel, NetWitness does not open non-indexed meta keys by default. When you open a non-indexed meta key in the Values view, NetWitness begins loading values for that meta key. If the load time is excessive, the load of the meta key times out with a message. Title, values, and counts for non-indexed meta keys are not drillable in the Values panel. Additional labeling in Investigation identifies the non-indexed meta keys.

To select the meta keys to apply to your investigation, you can:

- Select the default meta keys.
- Select a set of meta keys, called a meta group.

Note: Investigate has built-in meta groups and user-defined meta groups. Once created, user-defined meta groups can be edited, deleted, exported for use on other services, and imported to the service you are investigating. All of these procedures are provided in a separate topic: Use Meta Groups to Focus on Relevant Meta Keys.

The Default Meta Keys dialog allows you to specify the default view and display sequence for meta keys during navigation in the Investigate > Navigate view for a specific service. For each key or for all keys, you can set the default view to:

- Hidden: Results for default meta key are hidden and are not available to load.
- Open: Results for default meta key are open with all values and counts displayed.
- Close: Results for default meta key are closed with only the meta name visible.

- Auto: The loading of default meta keys is controlled by the index level, which must be Indexed By Value.

When using the default meta keys, be aware that these can be modified for different services, and you may not be seeing the same set of default meta keys when navigating to a drill point on different services. If you do not see the expected data, you may need to change the initial view of the default meta keys.

When you change the initial state of default meta keys from within the Navigate view, the change persists for that service. When new keys are added to the custom index file for a Core service (for example, `concentrator-custom-index.xml` or `decoder-custom-index.xml`), the new keys are added to the default meta keys list. The changes made in the Navigate view apply only to the current service.

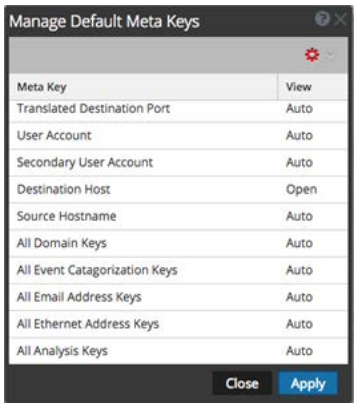
To specify that the initial Navigate view opens using default meta keys:

1. Go to **Investigate** > **Navigate**.
2. Select a service and select **Navigate**.
3. In the **Meta** menu, select **Use Default Meta Keys**.
   If an investigation is already in progress, the data is reloaded in the current view and an icon highlights the selected option. If no data is loaded yet, the default meta keys are used for the next load.

Configure Default Meta Keys

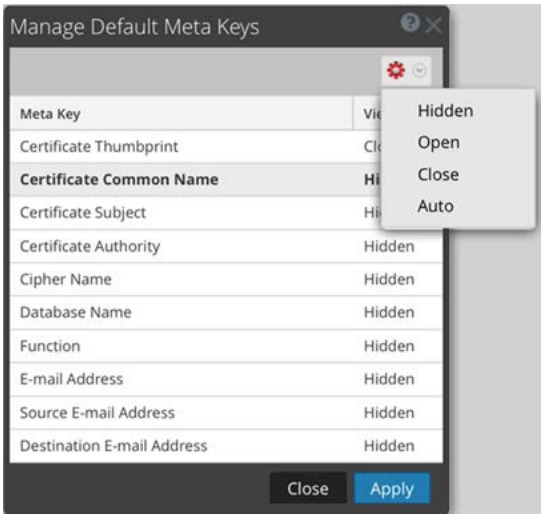
To configure the default view of default meta keys in the Navigate view:

1. In the **Navigate** view toolbar, select **Meta** > **Manage Default Meta Keys**.
   The Manage Default Meta Keys dialog is displayed with the list of available meta keys for the service.



2. (Optional) To change the order of the keys, select one or more keys, and drag the values up or down through the list of keys.
3. Do one of the following:
   - (Optional) To change the default view for all meta keys, make sure that no keys are selected and in the toolbar, select .

- (Optional) To change the default view for one or more keys, select the keys and in the toolbar, select .
  A drop-down menu of possible initial views for all default meta keys is displayed.



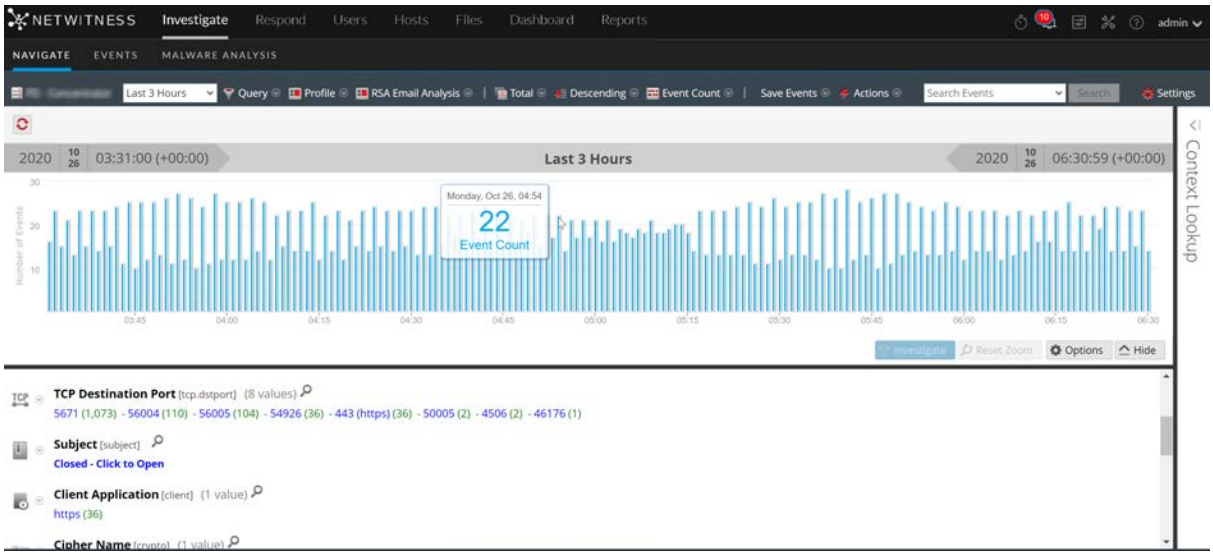
- (Optional) To revert to the default view for meta keys as specified in the service index file, make sure that no keys are selected and in the toolbar, select > **Auto**.
  When you modify the default view for a non-indexed meta key, you cannot set the key to OPEN. If you change the default view for a group of meta keys to OPEN and some of the meta keys are non-indexed, the non-indexed meta keys revert to AUTO. As a result, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are CLOSED until opened manually.

4. Select one of the views.
5. To save the changes, click **Apply**.
   The meta keys displayed in the Navigate view are set to your specifications. If the default meta keys are hidden, values for the meta keys are not shown in the investigation at all. If the default meta keys are closed, the values for the meta keys are not loaded by default, but you can load individual meta keys manually in the Navigate view.

Drill into Data in the Navigate View Time Chart

The Time Chart visualization allows analysts to visualize activity over time. You can zoom into the data by selecting a time window then selecting the Investigate option. You can then reset the navigation to the time range that was in effect before zooming.

1. Go to **Investigate > Navigate**.
   The Time Chart for the current drill point and selected time range is displayed. You can hover over the time chart to display total number of events occurred at a specific time.



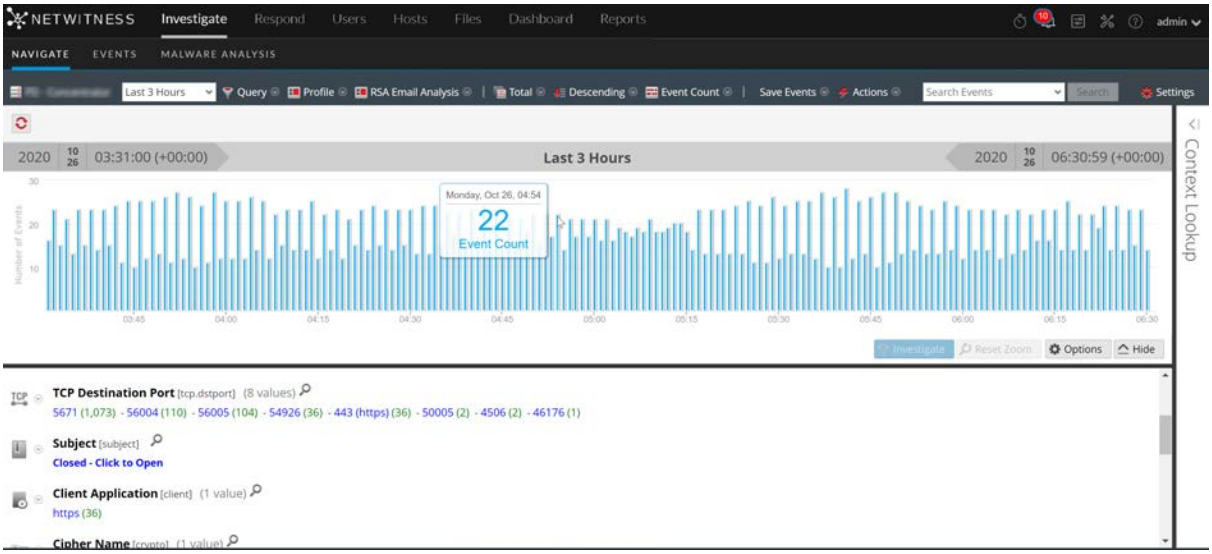
2. To highlight a period of time on the Time Chart, click over the desired time period and drag the mouse.
   The Time Chart is redrawn for the selected time range; however, the meta values are unchanged.
3. To drill into the data for the selected time range, click **Investigate**.
   The URL is updated to reflect the time range override, and the Investigation options panel is updated to reflect the custom time range. The Time Chart is redrawn and the meta values are loaded for the selected time range.
4. To reset the Time Chart to the original time range, click **Reset Zoom**.
   The URL is updated to reflect the original URL prior to zooming into the data, and the Investigation options panel is updated to reflect the time range selected before zoom. The Time Chart is redrawn for the selected time range and the meta values are loaded for that time range.

Drill into Data in the Values Panel

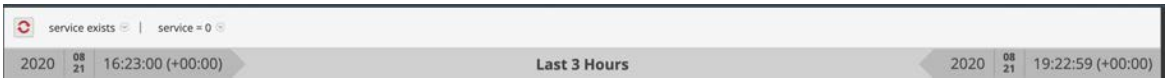
NetWitness displays the activity and values for the selected service in the Investigation > Navigate view. To investigate data, analysts drill into data by clicking on a meta key or a meta value, which is treaty as a query. In the Values panel, each query is added to the breadcrumb data in the Values panel. This results in a breadcrumb at the top with a crumb for each query. You can edit the breadcrumb to insert or remove a query.

To drill into a subset of the metadata:

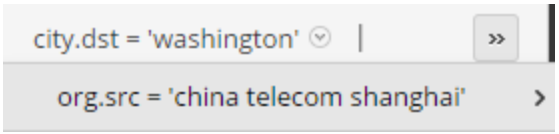
1. Begin an investigation so that metadata is displayed in the Navigate view.



2. To drill down into the metadata, do any combination of the following:
   a. Click a **meta key**, for example, **Service Type**.
   b. Click a **meta value**, the blue text in the results. For example, **OTHER**.
      Each time you click a meta key or meta value, the investigation query pivots to a narrowed focal point, or drill point, in the data. At each drill point, the Values panel is updated and the new drill point is displayed in the breadcrumb. Below is an example of the first breadcrumb.



      This is an example of a long breadcrumb that does not fit in the toolbar. The last query that fits is followed by a drop-down menu that lists additional queries. To select a drill point within the overflow, click the overflow icon and a query in the drop-down list.



To add a query in the breadcrumb:

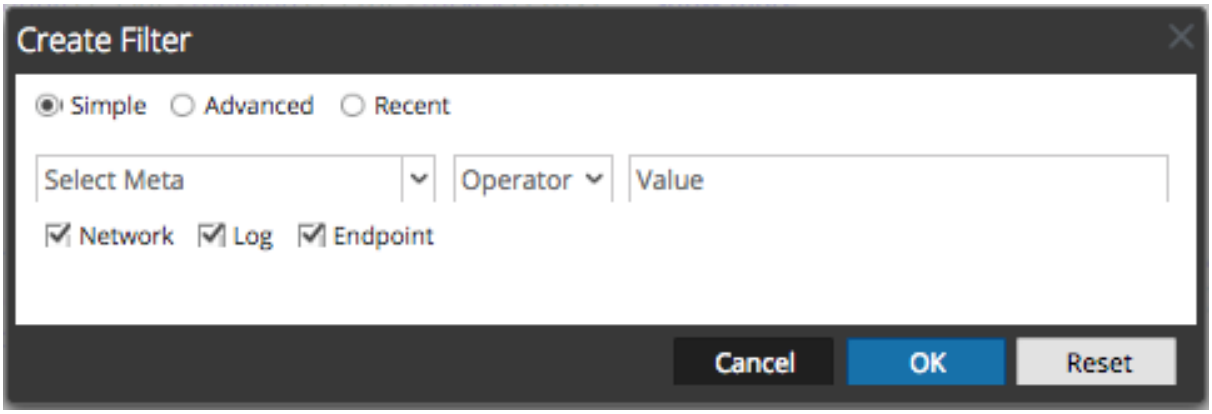
In the breadcrumb, you can click any of the crumbs to display the Query menu. You can insert a new query before a crumb, and append a new query to the end of breadcrumb. After each edit in the breadcrumb, NetWitness refreshes the results.

To add a query in the breadcrumb:

1. Click a crumb.
   The Breadcrumb menu is displayed.



2. To add a query in the breadcrumb, select **Append** or **Insert Before**.
   The Create Filter dialog is displayed.



3. Create the Query as described in Create a Query in the Navigate and Legacy Events Views.

To edit a query in the breadcrumb:

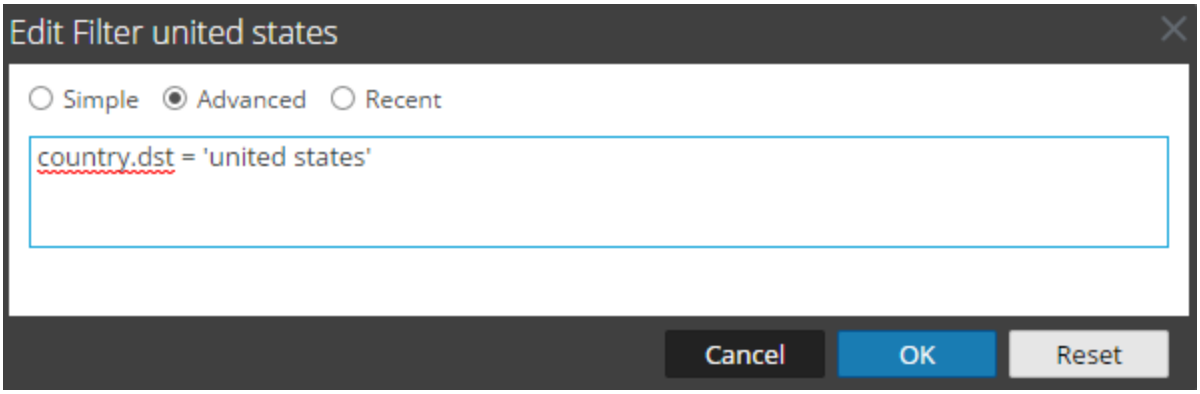
In the breadcrumb, you can click any of the crumbs to display the Query menu. You can delete a crumb and edit a query in a crumb. After each edit in the breadcrumb, NetWitness refreshes the results.

To work with queries in the breadcrumb:

1. Click a crumb.
   The Breadcrumb menu is displayed.



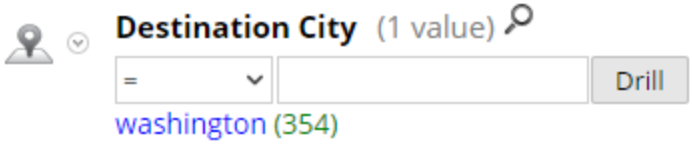
2. To edit a query in the breadcrumb, select **Edit**.
   The Create dialog is displayed with the selected query open for editing.



3. Edit the fields as described in Create a Query in the Navigate and Legacy Events Views.

To quick search within a meta key:

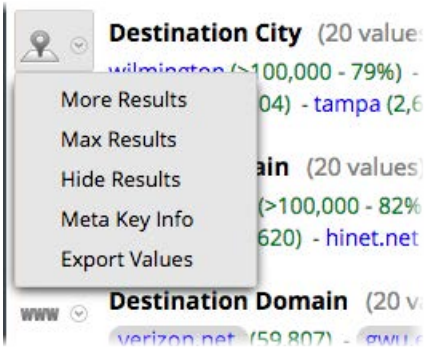
1. Move the mouse over a meta key section and click the magnifying glass.
   The Quick Search form, which contains a comparator and an optional operand for the search, is displayed.



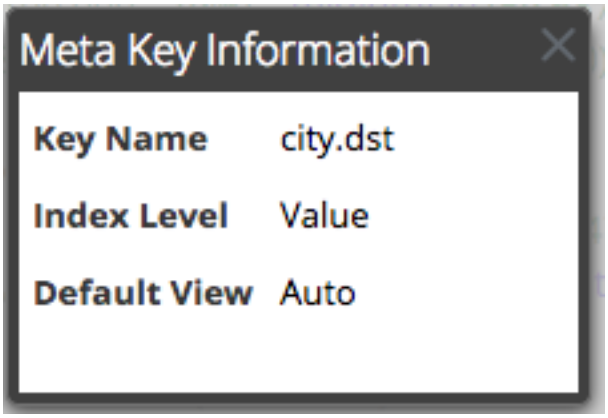
2. (Optional) If you want to close the search form, click the magnifying glass again.
3. Select the operation from the drop-down list on the left and type the text value to search for. Then click **Drill** to perform the execution.
   The metadata for that meta key is used to drill down in the current metadata.

To view meta key information and copy meta values for a meta key:

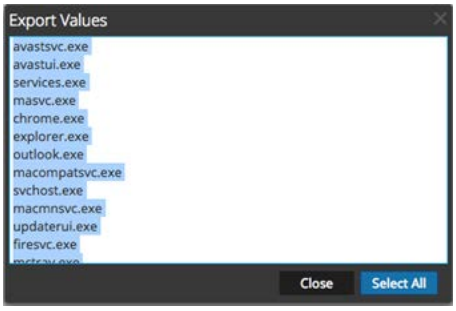
1. To view the key name, index level set for displaying the meta key, and the default view set for the meta key, click the drop-down menu next to the meta key. This figure shows the drop-down menu for Version 11.1 and later.



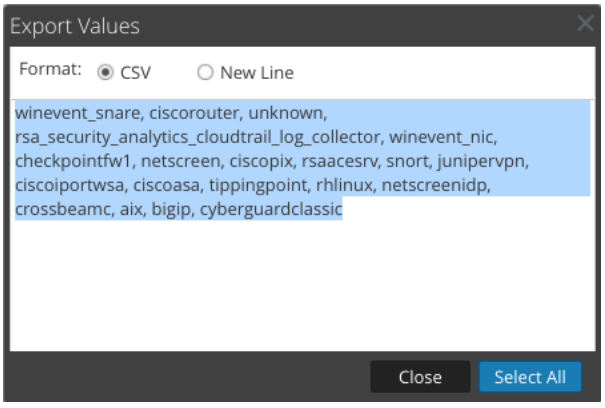
2. Select **Meta Key Info**.
   The Meta Key Info dialog is displayed.



3. When finished viewing, click ☒.
4. (Optional for Version 11.1 and later) To view meta values found for a meta key in a simple list that you can copy, click the drop-down menu next to the meta key.
   The Export Values dialog is displayed.
   The Version 11.1 dialog displays a list of values with one value per line.



   The Version 11.3 dialog allows you to select the method of separating values: either New Line or CSV.



5. Select the values that you want to copy, and click **Export Values**.
   The values are copied to the local clipboard and you can paste them into a file to save or share them.
6. To close the dialog, click **Close**.

7. (Optional) If you want to hide the results for the meta key in the current drill point, click the drop-down menu next to the meta key and click **Hide Results**.

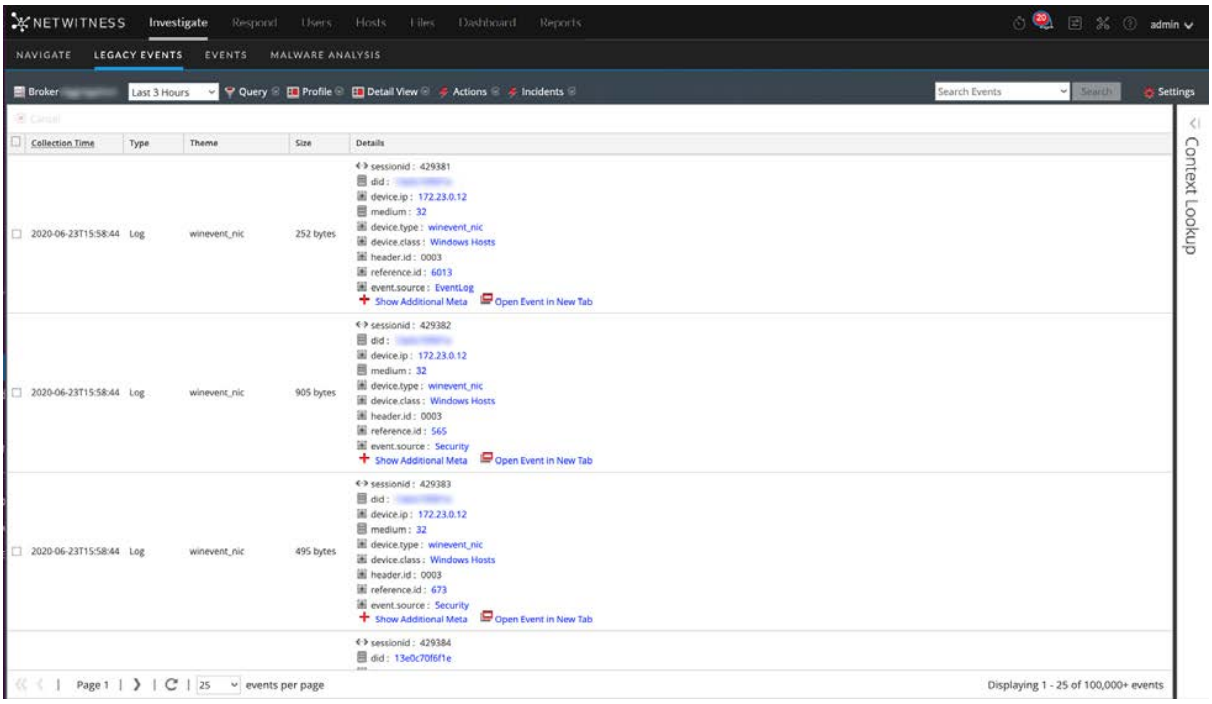
To display events associated with a meta value:

The Legacy Events view provides additional details for an event in two different views: Events List and Detail View.

1. In the Navigate view, drill into metadata that is the focus of your investigation.
2. Click the count (the number in green) next to a blue meta value.
   The Events view corresponding to the current drill point is displayed.
   The operations that you can perform in the events view are described in Reconstructing and Analyzing Events.

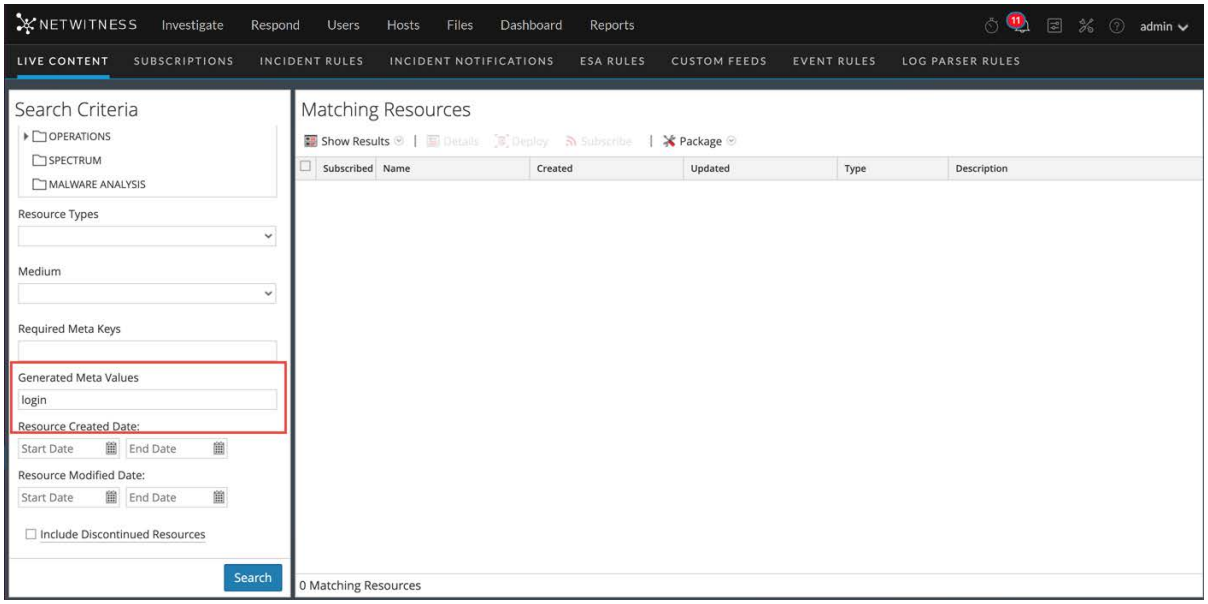
To search for specific events associated with a meta value:

1. In the **Navigate** view, drill into metadata that is the focus of your investigation (click a meta value or add a query).
2. Type a search string in the Search box and press **Enter** or click **Search**.
   You can also select and set search mode preferences. See Search for Text Patterns in the Navigate and Legacy Events Views for detailed search information.
   The Events view opens in a new tab and shows the search results. If you do not see the search term highlighted, click **Show Additional Meta**. Your time range selection and drills (queries) carry forward to the Events view.



To view a selected meta value in RSA Live:

1. In the Navigate view, drill into metadata that is the focus of your investigation.
2. Right-click a meta value (the text in blue).
   The Meta Value drop-down menu is displayed.
3. To look up the meta value in RSA Live, select **Live Lookup**.
   The Live Search view is displayed with the meta value entered in the Generated Meta Value(s) field, and ready for a search.



To refocus the investigation in a drill point:

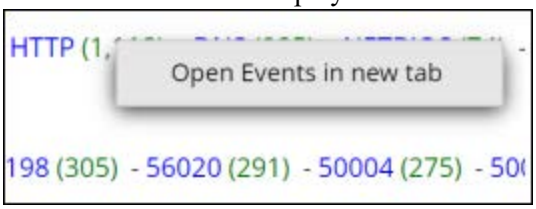
1. Right-click a meta value (the text in blue).
   The Meta Value drop-down menu is displayed.



2. Choose one of the refocus options.
   The drill is refocused according to your choice.

To look at a specific count in a new tab:

To view a count for a meta value in the Legacy Events view or the Events view, right-click a count for a meta value (the green number following the blue meta value).
The context menu is displayed.



Filter Results in the Legacy Events View

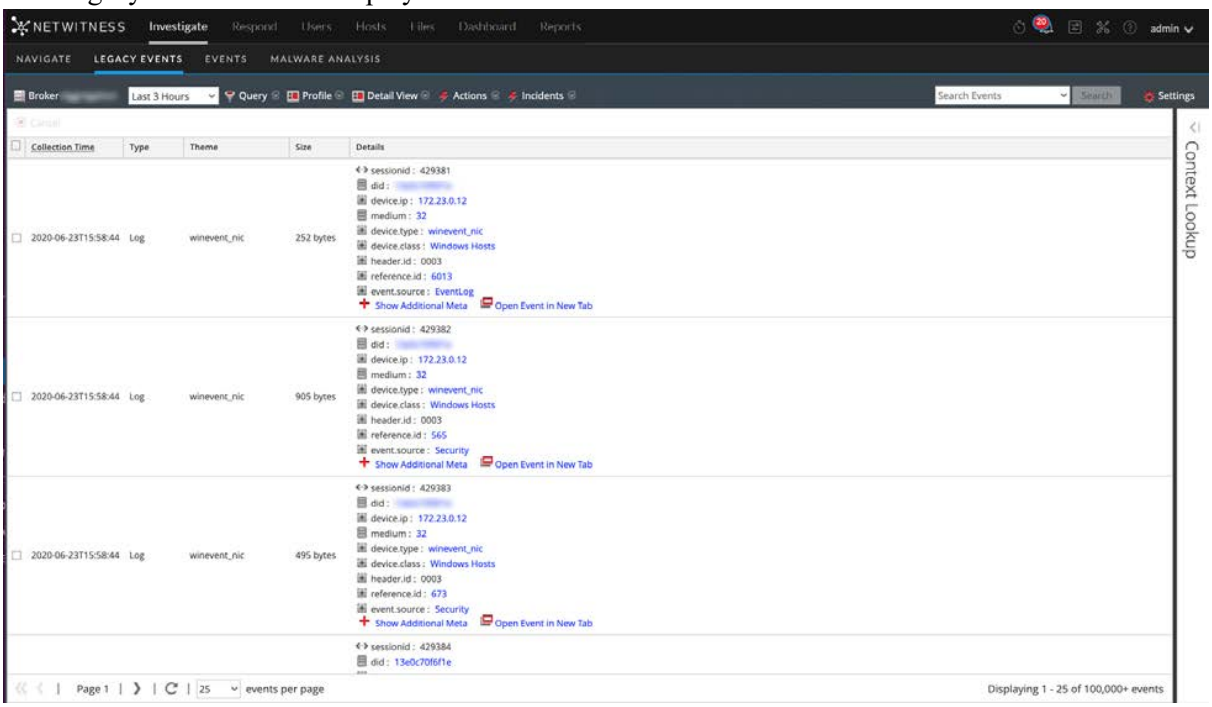
Analysts can filter events in the Legacy Events view by searching for events or selecting the service, setting the time range, and querying the metadata. If you opened the Legacy Events view from a Navigate view drill point, the view opens to the Detail view of events by default. Analysts who do not have permissions to use the Navigate view can query services directly from the Legacy Events view.

Note: When an Archiver is the currently selected service in the Legacy Events view and you are searching against a Broker or Concentrator, the search is slower than if searching against a Broker or Concentrator because the data on the Archiver is compressed and there is typically more data.

Filter Events Displayed in the Legacy Events View

To filter the data displayed in the Legacy Events view:

1. Go to **Investigate** > **Legacy Events**.
   The Legacy Events view is displayed.



2. To select a time range other than the default (**Last 3 Hours**), in the toolbar, click the time range field and select a value. For example, **Last Hour**.
   The Legacy Events view is refreshed with the selected time range.
3. Create a query as described in Create a Query in the Navigate and Legacy Events Views.
   The matching results for the query are displayed in the Detail View in the Legacy Events view. The breadcrumb reflects the query. In the breadcrumb, you can click any of the crumbs to display the Query menu. You can insert a new query before a crumb, and append a new query to the end of breadcrumb. After each edit in the breadcrumb, the results are refreshed.

Page Through Events in the Legacy Events View

Pagination controls allow more flexibility in paging through a list of Events in the List View, Logs View, or Details View. You can select the number of events to display per page, and your selection is saved across logins to the NetWitness application. When a control is unavailable, the control is dimmed; for example, when you are viewing page 1, the ❮ and ❮❮ controls are dimmed.

To use pagination controls:

1. With results displayed in the Legacy Events view, click the current number of events per page (**10**, **25**, **50**, **100**, or **200**), and select the new number of events per page from the drop-down menu.
2. To page forward or back, use the page control icons:
   Click ❯ to go to the next page.
   Click ❯❯ to go to the last page.
   Click ❮ to go the previous page.
   Click❮❮ to go to the first page.
3. To go to a specific page, type a page number in the page number field | 3 | Page 3 |.

Create a Query in the Navigate and Legacy Events Views

In the Navigate view or the Legacy Events view, you can create a query using dialogs for that offer syntax help with drop-down lists of applicable meta keys or meta entities and operators.

When viewing the drop-down list, you can expand and collapse each meta group to view or hide the individual meta keys in that group. When you select a meta group, NetWitness generates the complex query equal to a query with all of the meta keys in that group ORed together. So if a meta group contains `ip.src` and `ip.dst`, the query generated is `ip.src = <value> OR ip.dst = <value>`. If the meta group contains meta keys that have different meta value types, the value input is disabled and the query uses `exists` statements. For example, a meta group that contains `ip.src`, `ip.dst`, and `alias.host` includes meta keys that have different value types; `ip.src` and `ip.dst` are ip addresses and `alias.host` is text. The generated query is `ip.src exists OR ip.dst exists OR alias.host exists`.

A basic query is in the following form:

```
<metakey> <operator> [<metavalue>]
```

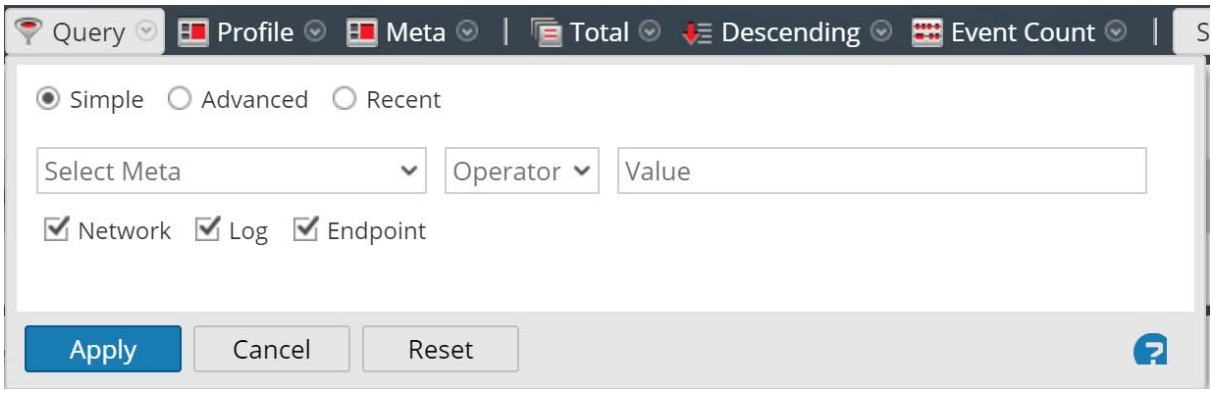
These are a few examples:

```
action exists
action = 'get'
alias.host = '10.25.55.115'
extension = 'exe'
orig_ip != "10.0.0.0" - "10.255.255.255"
```

Create a Query Using the Basic Method

When you create a query using the basic method, drop-down lists of meta keys and operators are displayed.

1. In the **Navigate** view or the **Legacy Events** view toolbar, select **Query.**
   The Query dialog is displayed, with the Simple option selected.

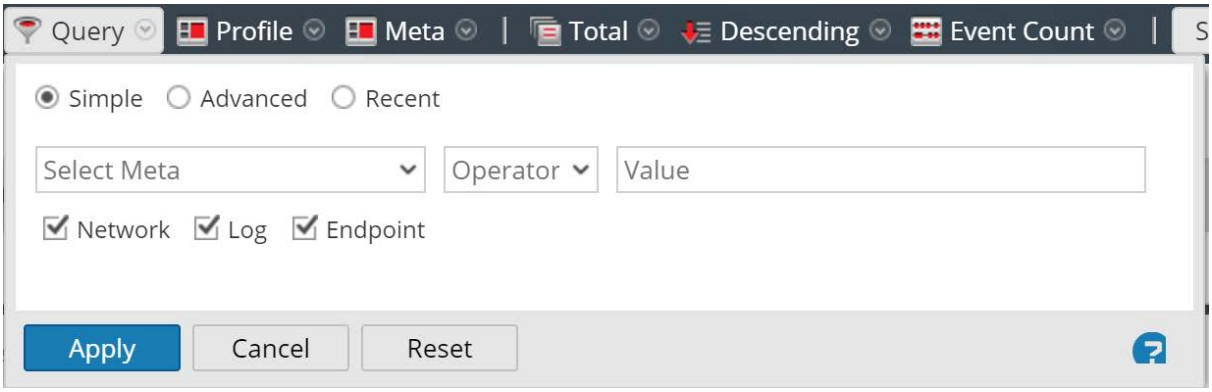


2. In the **Select Meta** field, click to display the drop-down list. The drop-down list has two sections: Meta Groups and All Meta.
3. Select a single meta key under **All Meta** or select a meta group under **Meta Groups**. You can also type in a meta key or meta group in the field.

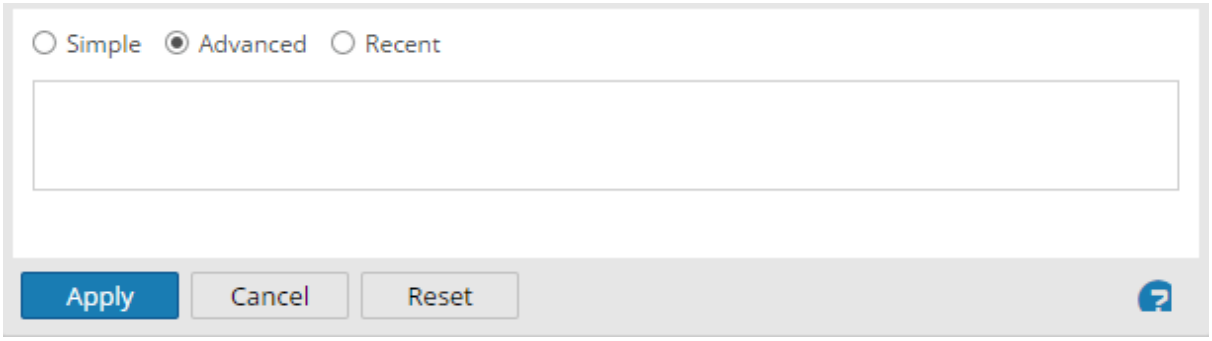
4. In the **Operator** field, type an operator or click on the drop-down list to select a valid operator.
5. (Optional) If you selected an operator that requires a value, for example, =, in the third field type the value for the meta key.
6. In the Network, Log, and Endpoint checkboxes, choose the type of data to query. Do one of the following:
   a. To limit the query to packets select **Network** and de-select **Log** and **Endpoint**.
   b. To limit the query to logs, select **Log** and de-select **Network** and **Endpoint**.
   c. To limit the query to endpoint events, select **Endpoint** and de-select **Network** and **Log**.
   d. To apply the query to packets, logs, and endpoints, select **Network**, **Log**, and **Endpoint**.
7. Do one of the following:
   a. Click **Apply**.
   The window is closed and the view is updated with the results of the new query. The query is displayed in the breadcrumb.
   b. Click **Cancel**.
   The window is closed and no changes are made to the view or current query.

Create a Query Using the Advanced Method

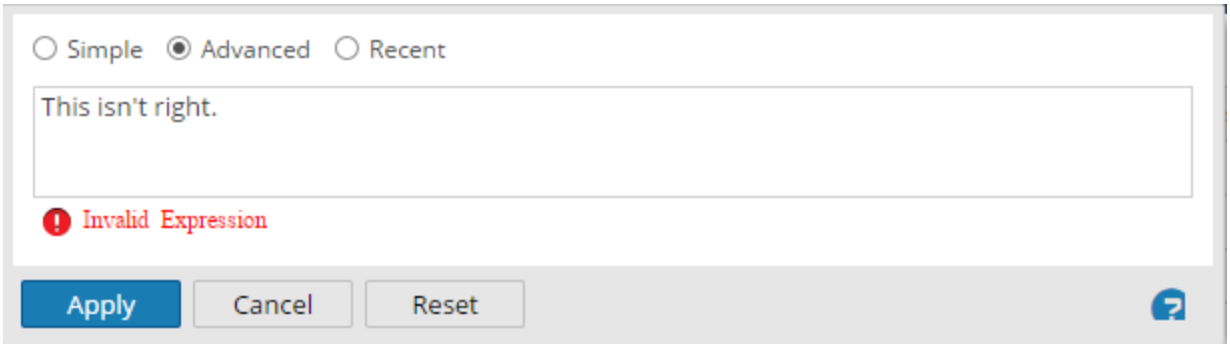
1. In the **Navigate** view or the **Legacy Events** view toolbar, select **Query.**
   The Query dialog is displayed.



2. Select **Advanced**.
   The advanced query field is displayed.



3. In the field, create a query, which can include the meta key, operator, and value. When you begin typing a meta key in the field a drop-down list of available meta keys for the selected service is displayed.
4. Select the meta key for your query.
   The display is updated. If the expression is not yet complete, the status indicates that the query is invalid.
5. Continue with an operator, from the drop-down list, then a value if necessary. The display is updated as you continue to enter the query. If you enter an operator, such as **exists** or **!exists**, which does not use the value field, the value field is disabled and the invalid status is cleared. If you enter an operator, such as =, which requires the value field, the invalid status remains until you enter a value. When the query is valid the invalid status is no longer displayed.

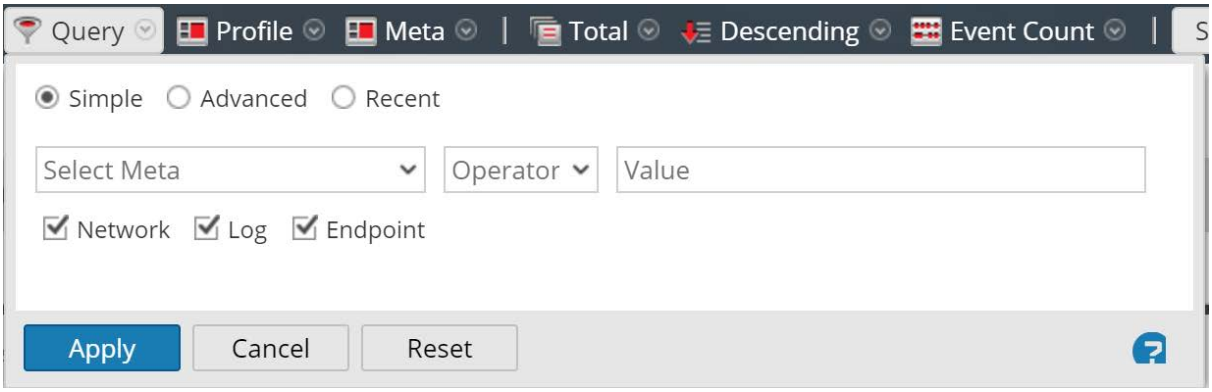


6. Do one of the following:
   - Click **Apply**.
     The window is closed and the view is updated with the results of the new query. The query is displayed in the breadcrumb.
   - Click **Cancel**.
     The window is closed and no changes are made to the view or current query.

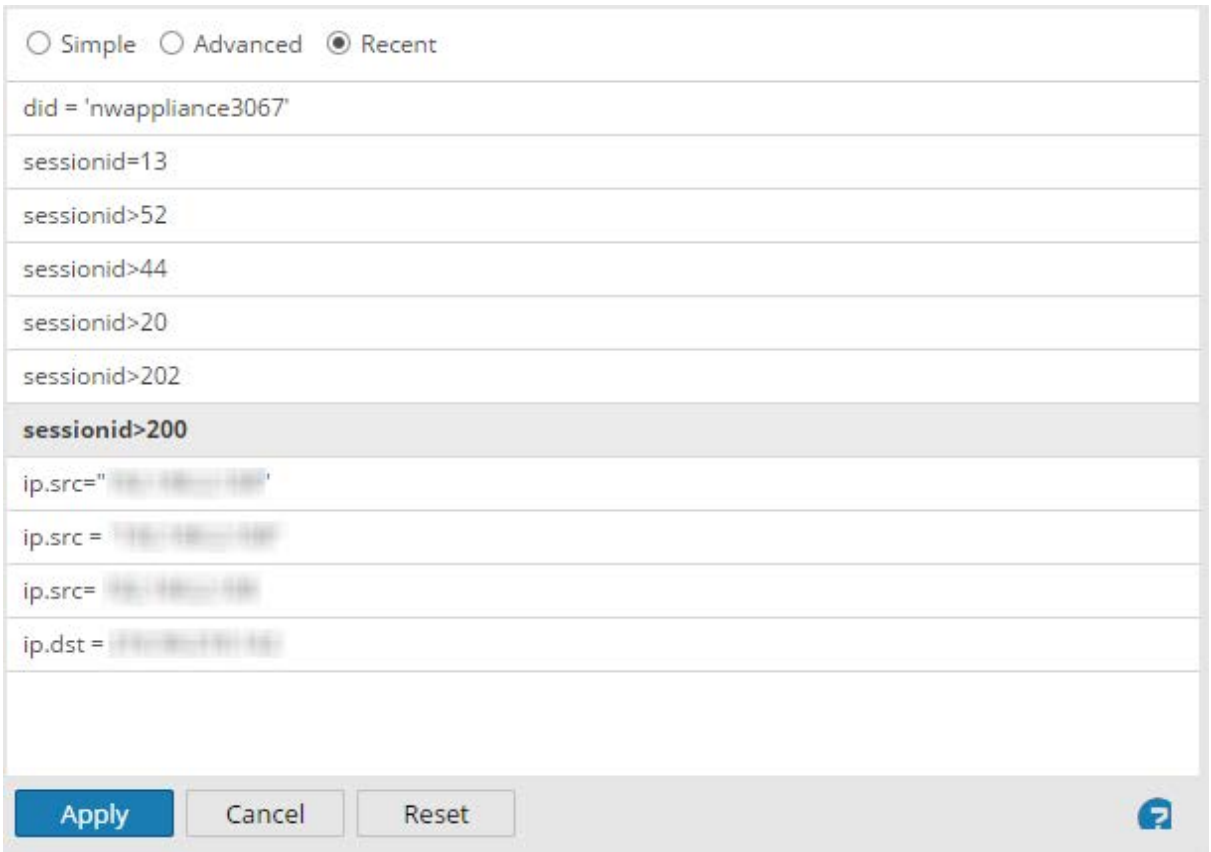
Apply a Recent Query

You can view recent queries and select one to apply to the current service being investigated. To select a recent query:

1. In the **Navigate view** or the Events view toolbar, select **Query.**
   The Query dialog is displayed, with the Simple option selected.



2. Select the **Recent** option.
   The list of recent queries is displayed in the bottom portion of the dialog.



3. In the list of recent queries, click to select a query.
4. Do one of the following:
   - Double-click a query.
   - Select a query and click **Apply**.
     The window is closed and the view is updated with the results of the new query. The query is displayed in the breadcrumb.
   - Click **Cancel**.
     The window is closed and no changes are made to the view or current query.

Search for Text Patterns in the Navigate and Legacy Events Views

You can search for text patterns within the current set of events in the Navigate view, the Events view, and the Legacy Events view. This section provides information about searching in the Navigate view and the Legacy Events view.

You can perform a keyword text search or do regex (Regular Expression) matching. In the Navigate view, you can click a meta value, such as HTTP, to drill into the data and then enter a search string in the Search field to search for events within that subset of data. The search opens a tab in the Legacy Events view, brings your drill and time range forward, and shows your search results. You can also drill into the data using queries before starting a search. To execute the search, enter a search string in the Search box, and press **Enter** or click **Search**.

Note: By default search results are only for exact matches found in indexed data. Only meta values shown as blue links in the Events Detail view are indexed. The regex option must be selected if the value contains a space. To broaden the search change the options in the Search Events drop-down menu.

Keyword Text Search

The text search provides these capabilities:

- Each white space delimited word is ANDed, so that every word must be found, but the order or location position in relation to the other words is irrelevant. For example, if you search on `Mark Albert`, both Mark and Albert must be found in the session, but they need not be together or in any specific order.
- The word OR is special. If you search `Mark OR Albert`, either Mark or Albert must be found in the session to match; both are not required.
- You can mix or match implicit ANDs and ORs together in the search string. The explicit OR has higher precedence than the implicit (whitespace) AND. The following examples make the same logical statement, which requires that both the terms cheese and dumplings be present in a match and one of toast or bread:

  ```
  cheese toast OR bread dumplings
  cheese AND (toast OR bread) AND dumplings
  ```

- You can exclude words from search results using the `-` operator. For example, searching for `cheese -toast` would return any result that has the word cheese, unless the word toast is also present.
- The keyword search can match metadata stored in the following patterns:
  - **IPv4 and IPv6 addresses**. Any term that can be recognized as an IP address is converted to the native metadata format so that it can be found in indexed metadata.
  - **IPv4 CIDR ranges**. You can use CIDR notation to locate IPv4 addresses within a range.

- **Timestamps**. Timestamps are matched against the native time metadata, and any additional time meta fields stored with the Time type.
- **Numbers**. The search function will attempt to automatically identify decimal search terms and match them against numeric meta data fields.

Options Controlling Search Behavior

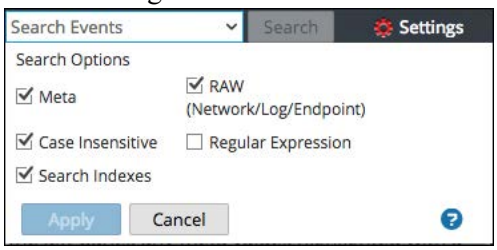
To access the Search box and search options in the Navigate or Legacy Events view:

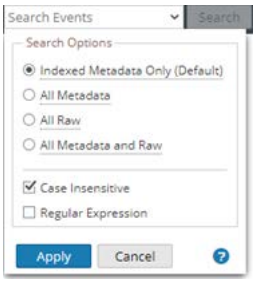
1. You can see the Search Events field in the toolbar.



Note: If you cannot see the Search Events field in the toolbar, click ⋯ on the right side of the toolbar.

2. Click in the **Search Events** field to view the Search Options drop-down menu. In Version 11.2 and later, the menu options are slightly different. The first figure illustrates the menu for 11.1 and below; the second figure illustrates the menu for Version 11.2 and above.





The options selected in this box change how the search is executed. The default search mode is to search indexed metadata and raw data only.

Note: Because the Index or Indexed Metadata Only (default) checkbox is selected by default, the search returns results based on data that is indexed. If you want to search for a complete set of metadata or raw data, select those checkboxes and clear the Index or Indexed Metadata Only (default) checkbox. This type of search takes longer, but it contains a more complete set of data.

The following table describes the Investigation search options.

Feature	Description
Indexed Metadata Only (default) checkbox (Version 11.2 and later) **Index** radio button (Version 11.1)	This search only returns results on indexed data. Searching the index is the fastest way to locate keywords within a large data set. The index search uses any relevant indexes present within your data collection. **Caution:** Substring matches are not located by index searches. If you require substring matches, clear this checkbox and use a non-index search mode.
All Metadata radio button (Version 11.2) **Meta** checkbox (Version 11.1)	Searches the metadata. Your keyword or regex pattern is matched against any parsed metadata.
All Raw radio button (Version 11.2 and later) **RAW** (Network/Log/Endpoint) checkbox (Version 11.1)	Searches the network, log, and endpoint event text. Every event is decoded and content is searched for matches on the keyword or regex pattern. If you select all data with no filters on an Archiver, execution time may be excessive and a warning may be displayed. **Caution:** Searching raw network sessions causes sessions to be decoded, which is very time intensive. You may want to disable raw searches when looking at network-only collections.
All Metadata and Raw radio button (Version 11.2)	Searches the metadata and the log or event text. This option is a combination of two options in Version 11.1: Meta and RAW (Network/Log/Endpoint), which you could select together. In Version 11.2, you can select only one radio button.
Case Insensitive	Ignores case when searching.
Regular Expression	Searches using a Perl regular expression, rather than text. By default executes a text search. To execute a regular expression search, select the Regular Expression option. **Caution:** - Regular expression searches can be very slow. - When combining regular expressions and index search options, the regular expression pattern is matched against unique index values instead of meta values. This produces results faster, but it is not an exhaustive search of all the metadata or raw data.
Apply	Sets the default search options to apply to a search in the Navigate and Legacy Events views. This also updates your Investigation preferences in your Profile (Profile > Preferences > Investigation tab). The preferences are saved and effective immediately. You can select search options to use for a particular search without changing your default search preferences.

Regular Expression Search Syntax

A regular expression search uses Perl regular expression syntax, which is documented in detail in http://perldoc.perl.org/perlre.html.

Raw Text Keyword Search

The Log Decoder has the capability to create a raw text index for unparsed log events. This functionality creates metadata items that form a full-text index on downstream services such as Concentrators and Archivers. When you enable the Search Indexes option in your search preferences, your search automatically uses the text index. Note that the text index produces meta items that have a coarse granularity. For example, the default text indexer configuration truncates text terms. By comparing the index matches against raw data, the search engine will find accurate results for your search. However, you can improve search times by disabling the raw search checkbox. If you do so, results will be returned faster, but you may see false positive hits in your search results.

Search Procedures

Search in the Navigate View

To search within the currently displayed data in the Navigate view:

1. Type a search string in the Search field and press **Enter** or click **Search**.
2. To clear the search box and return to the previous Navigate view with results unfiltered by the search, click the **X** in the search box.

Search in the Legacy Events View

To search within the currently displayed data in the Legacy Events view:

1. Type a search string in the Search box, and press **Enter** or click **Search**.
   The search results are displayed. Events that match the search criteria are displayed in the events list. In the Details view and List view, matches are highlighted in the Details column. In addition, when searching RAW, matches are highlighted in the Log view Logs column.
2. If you want to narrow the search, change the query and time.
3. If you want to stop the search and return to the Legacy Events view, click **Cancel**.
   Any results that are displayed remain.
4. To clear the search box and return to the normal Events view, click **X** in the search box.

View and Modify Queries Using URL Integration

NetWitness Investigate includes an External URL Integration that facilitates integration with third-party products by allowing a search against the NetWitness architecture. By using a query in a URI, you can pivot directly from any product that allows custom links, into a specific drill point in the Investigate view. This integration provides an internal presentation of the user's query.

URL Integration allows the user to identify the service either by the host id or by the service and port, as defined in NetWitness. If NetWitness is unable to resolve the service, the analyst is redirected to the Navigate view, showing the Service selection dialog. Once the service is selected, the Navigate view is loaded with the drill point, defined by the query.

Service Id Known

When the ID of the service to use for an investigation is known, the format for entering a URI using a URL-encoded query is:

```
http://<sa host:port>/investigation/<deviceId>/navigate/query/<encoded query>/date/<start date>/<enddate>
```

where

- `<sa host: port>` is the IP address or DNS, with or without a port, as appropriate (ssl or not). This designation is needed only if access is configured over a non-standard port through a proxy.
- `<deviceId>` is the internal Service ID in the NetWitness instance for the service to query against. The service ID can be represented only as an integer. You can see the relevant service ID from the URL when accessing the Investigate view within NetWitness. This value changes based on the service being connected to for analysis.
- `<encoded query>` is the URL-encoded NetWitness query. The length of query is limited by the HTML URL limitations.
- `<start date>` and `<end date>` define the date range for the query. The format is `<yyyy-mm-dd>T<hh:mm:ss>Z`. The start and end dates are required. If no date is provided then the user defaults for that service are used. Relative ranges (for example, Last Hour) are not supported. All times are run as UTC.
  For example:
  `http://localhost:9191/investigation/12/navigate/query/alias%20exists/date/2012-09-01T00:00:00Z/2012-10-31T00:00:00Z`

Host and Port Known

When the host and port of the service to use for investigation is known, the format for entering a URI using a URL-encoded query is:

```
http://<sa host:port>/investigation/<device host:port>/navigate/query/<encoded query>/date/<start date>/<enddate>
```

where

- `<sa host: port>` is the IP address or DNS, with or without a port, as appropriate (ssl or not). This designation is needed only if access is configured over a non-standard port through a proxy.
- `<device host:port>` is the host and port of a service defined in the NetWitness instance for the service to query against. NetWitness attempts to resolve the host and port as a service ID defined in NetWitness.
- `<encoded query>` is the URL-encoded NetWitness query. The length of query is limited by the HTML URL limitations.
- `<start date> and <end date>` define the date range for the query. The format is `<yyyy-mm-dd>T<hh:mm:ss>Z`. The start and end dates are required. If no date is provided then the user defaults for that service are used. Relative ranges (for example, Last Hour) are not supported in this version. All times are run as UTC.
  For example:
  `http://localhost:9191/investigation/concentrator:50105/navigate/query/alias%20exists/date/2012-09-01T00:00:00Z/2012-10-31T00:00:00Z`

Examples

These are query examples where the NetWitness Server is 192.168.1.10 and the deviceID is identified as 2.

All activity on 03/12/2013 between 5:00 and 6:00 AM with a hostname registered

- Custom Pivot: `alias.host exists`
- `https://192.168.1.10/investigation/2/navigate/query/alias%2Ehost%20exists/date/2013-03-12T05:00:00Z/2013-03-12T06:00:00Z`

All activity on 3/12/2013 between 5:00 and 5:10 PM with http traffic to and from IP address 10.10.10.3

- Custom Pivot: `service=80 && (ip.src=10.10.10.3 || ip.dst=10.0.3.3)`
- Encoded Pivot Dissected:
  - `service=80 => service&3D80`
  - `ip.src=10.10.10.3 => ip%2Esrc%3D10%2E10%2E10%2E3`
  - `ip.dst=10.10.10.3 => ip%2Esrc%3D10%2E10%2E10%2E3`
  - `https://192.168.1.10/investigation/2/navigate/query/service%3D80%20%26%26%20%28ip%2Esrc%3D10%2E10%2E10%2E3%20%7C%7C%20ip%2Edst%3D10%2E10%2E10%2E3%29/date/2013-03-12T17:00:00Z/2013-03-12T17:10:00Z`

Additional Notes

Some values may not need to be encoded as part of the query. For example, commonly the IP src and dst is used for this integration point. If leveraging a third-party application for integration of this feature, it is possible to reference those without encoding applied.

Reconstructing and Analyzing Events

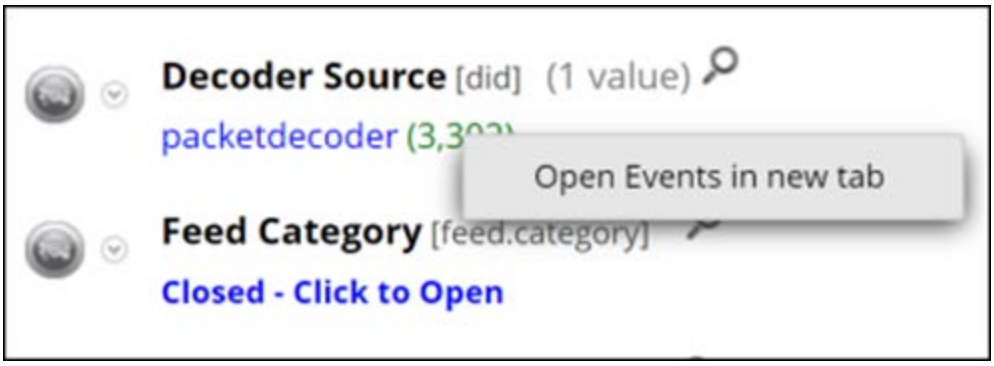
Having refined events in the Navigate view or in the Events list as described in Refining the Results Set, your next step is to learn more about the events by reconstructing them, looking at attachments, and viewing additional context in third-party lookups or internal lookups.

Reconstructions are done in the Events view or the Legacy Events view. If you are starting from the Navigate view, you need to go to the Events view or the Legacy Events view to see a reconstruction.

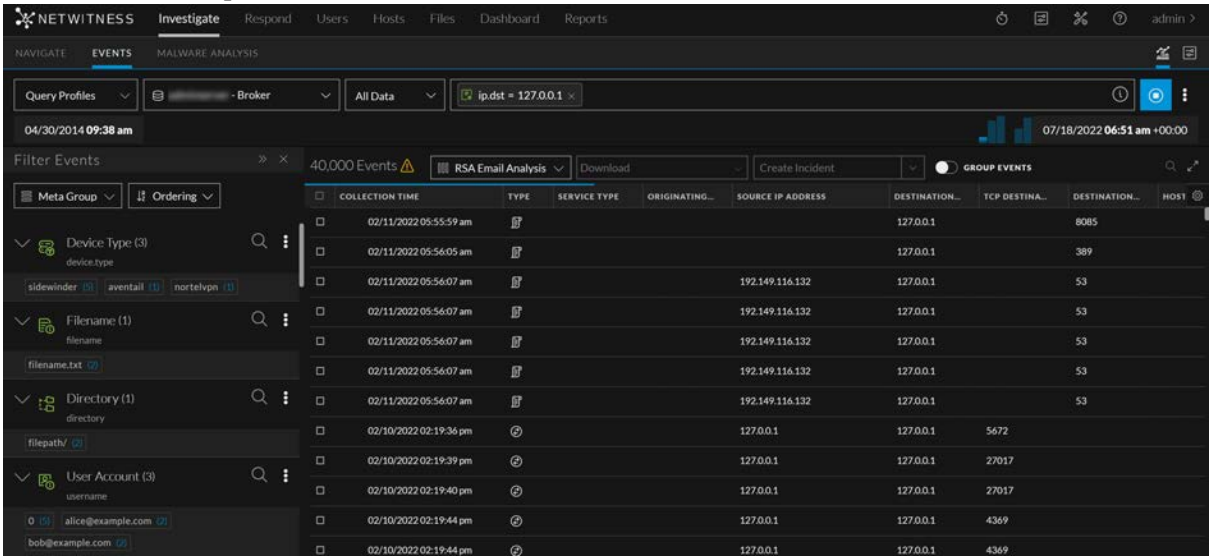
Note: The Legacy Events view is disabled by default. The administrator can enable the view as described in "Configure Investigation Settings" in the *System Configuration Guide*.

To display events in the Events view, do one of the following:

1. Go to **Investigate** > **Events**.
2. Go to **Investigate > Navigate** (Version 11.5 and earlier), right-click the meta count for a meta value (the meta count is in green text). When the context menu is displayed, select **Open Events in new tab**.



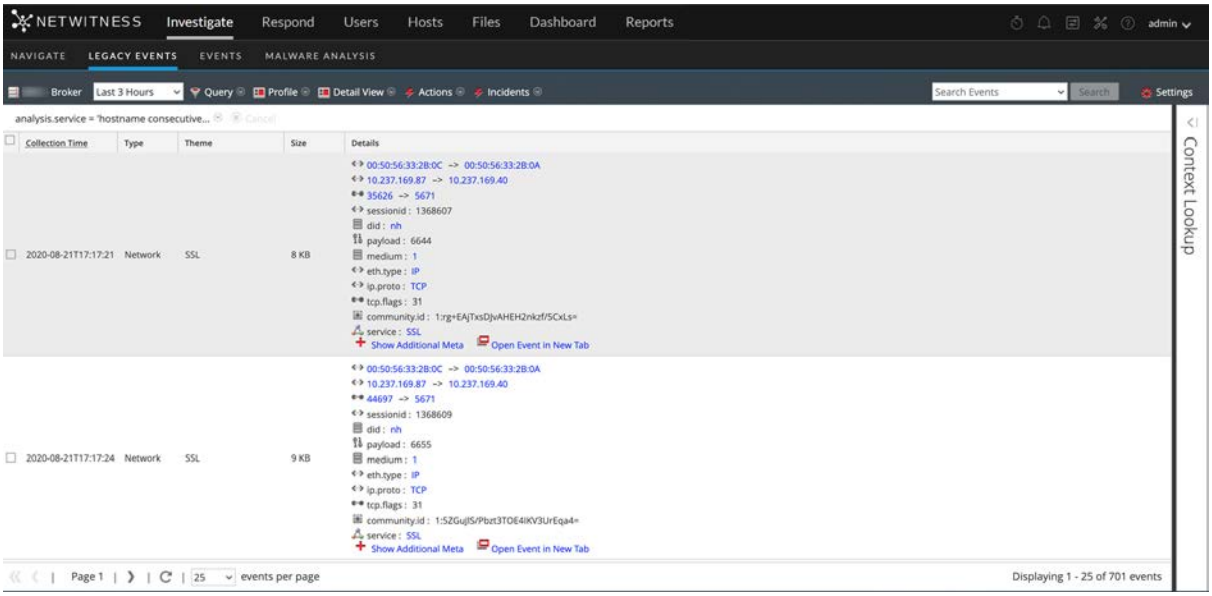
The Events view opens with a list of events for the selected meta value.



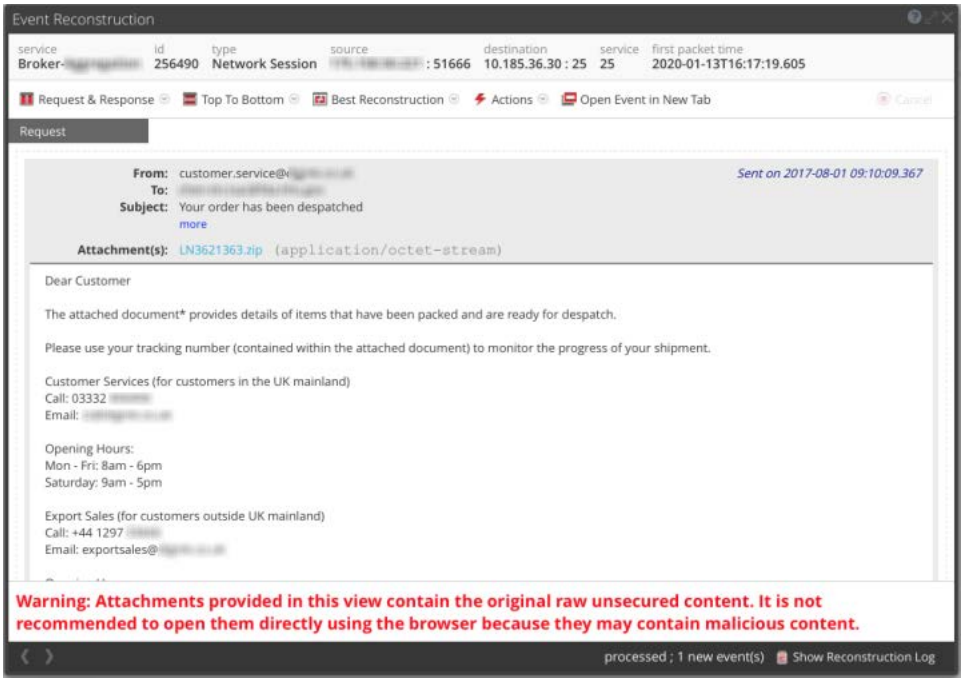
For detailed information about the types of reconstruction and analysis that you can use in this view, see Examine Event Details in the Events View.

To display an event in the Legacy Events view, do one of the following:

1. To open the Legacy Events view using the default query for the default service, go to **Investigate** > **Legacy Events**.(This option is available only if the administrator has enabled the view.)
2. To view events for a specific meta value in the Legacy Events view, go to **Investigate > Navigate** and when events are loaded in the Values panel, click a meta count (the meta count is in green text). You can also right-click the meta count for a meta value. When the context menu is displayed, click **Open Legacy Events in new tab**.
   The Legacy Events view displays the events for the selected meta value. The Legacy Events view provides three built-in presentations of event data: the Detail view, the List view, and the Log view. This figure is an example of the Detail view. You can use queries, the time range setting, and profiles to filter the events listed in the Legacy Events view. You can extract files, export events, export logs, and open the Event Reconstruction panel by double-clicking an event. See Downloading and Acting Upon Results for detailed information about these capabilities.
   NetWitness runs a default query on the last three hours for the default service (if one is set) or displays a dialog in which you can select a service and then runs the default query. The default query selects all events and the Events view displays events on the selected service, with the oldest events first.



3. To view a reconstruction of the first event in the list, double-click the event.
   The reconstruction opens in a pop-up window in front of the Events list.



Examine Event Details in the Events View

When you find an interesting session in the Navigate view or the Events view > Filter Events panel, you can see the list of sequential events for the session in the Events view > Events panel. Clicking an event in the list opens the Network Event Details panel for that type of event: Network Event Details, Log Event Details, or Endpoint Event Details. Within the Event Details panel, you can select a tab that shows an event reconstruction (text, packet, file, email, and web) or (Version 11.5 or later) the tab that shows host information for network events that are enriched with endpoint data.

Note: (Version 11.5 or Later) For expanded network visibility of existing network events in your network (packet) deployment. Network events are enriched with endpoint data namely the host and process that triggered the network event and other details such as user name, risk score, reputation, and so on.
You can view endpoint data in the following ways:
- (Quick View) **Investigate** > **Events** - Event Summary Header
- (Detailed View) **Investigate** > **Events** > **Host**

For more information to enable expanded network visibility, see "Creating Groups and Policies" in the *Endpoint Configuration Guide*.

Note: For expanded network visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.

Event Details for Each Event Type

Within the Event Details panel, different tabs are available per event type as shown in the following table. Procedures for working in the Event Details panel are provided in Analyze Events in the Events View.

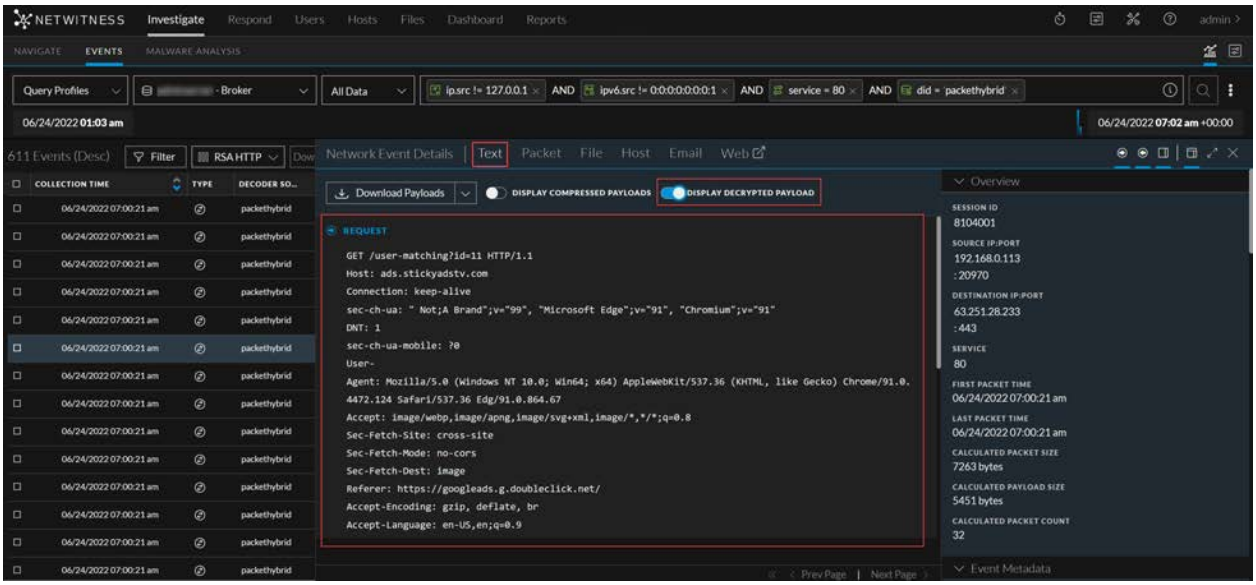
Action	Network Event	Log Event	Endpoint Event
View the text reconstruction (default unless last selected overrides)	✓	✓	✓
View the file reconstruction	✓		
(Version 11.5 or Later) View the host Information for an Endpoint Agent configured with expanded network visibility (see Host Information.	✓		
View the packet reconstruction	✓		
View the email reconstruction	✓		
View the web reconstruction in the Legacy Events view (see Reconstruct an Event in the Legacy Events View)	✓		

Each tab has settings to enhance your analysis. If you change a setting, the setting is preserved between browser refreshes and logins within the same browser. These are the preserved settings:

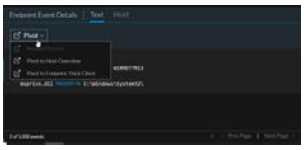
- The currently selected reconstruction: Text, Packet, File, (Version 11.5 and later) **Host**, or Email.
- Whether the Event Meta panel is open or closed.
- Whether the Event header is open or closed.
- Whether the Request or Response, or both are displayed.
- Whether packet payloads are displayed without the headers in the packet reconstruction.
- Whether shaded bytes are displayed in the packet reconstruction.
- Whether other common file types are highlighted in the packet reconstruction.
- The number of packets per page in the packet reconstruction.
- Whether compressed or uncompressed text is displayed in the text reconstruction.

Text Reconstruction

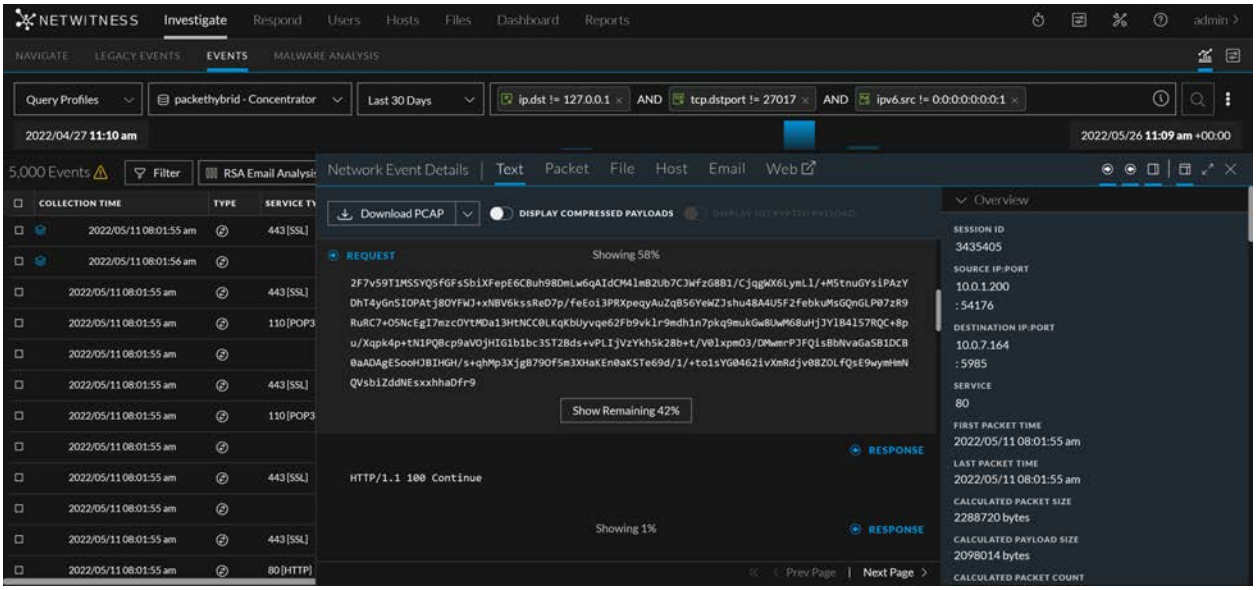
In 12.0 and later, analysts can directly view an encrypted data in a decrypted format by switching on a toggle option **Display Decrypted Payload**. However, the data can be displayed in the decrypted format only if a Decoder service in the **Events** > **Text** tab, has a TLS key available. This feature enables the analyst to focus on the important data in less time and perform investigations on the selected events with optimum accuracy and quality.

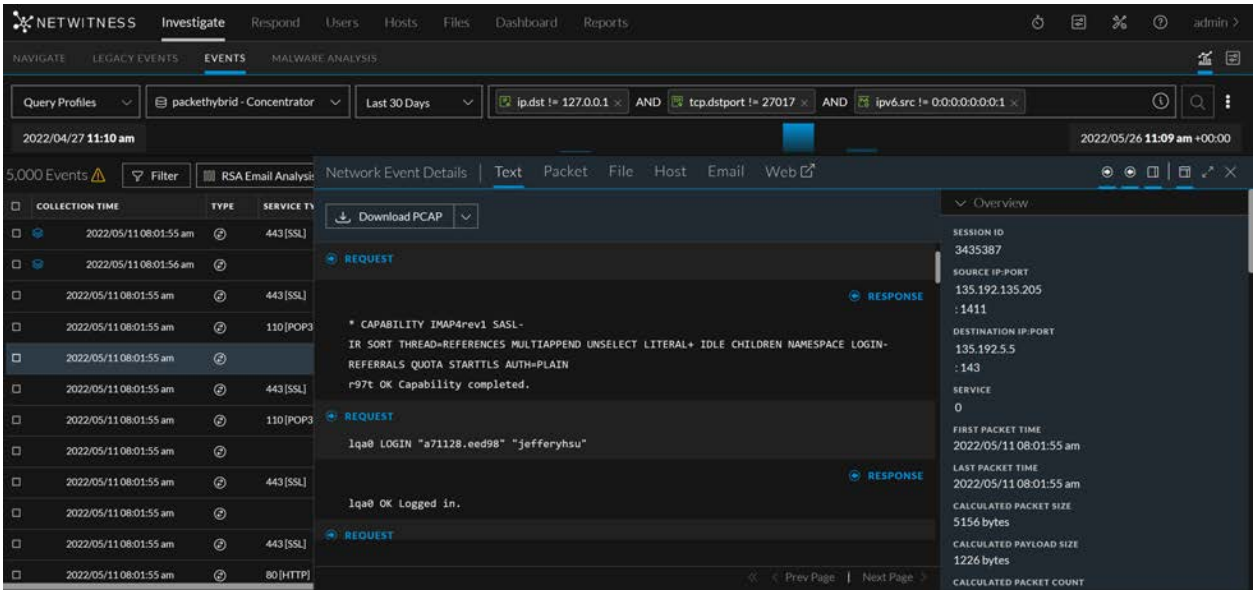


For Endpoint Events Details, the Pivot options available under Text analysis page has been replaced with a Pivot dropdown menu that contains the three options to perform further investigation.



You can view all types of events (network events, log events, and endpoint events) in their original text format in the Text tab. The text reconstruction for some network events can be quite large. To ensure the best rendering, an excessively large payload is truncated to fit. If a single reconstructed request or response in the reconstructed event exceeds the maximum number of bytes, the header indicates what percentage of bytes is shown. Pagination controls add flexibility when paging through the reconstructed text of an event. This figure illustrates a single response that has been truncated because it exceeds the maximum number of bytes.



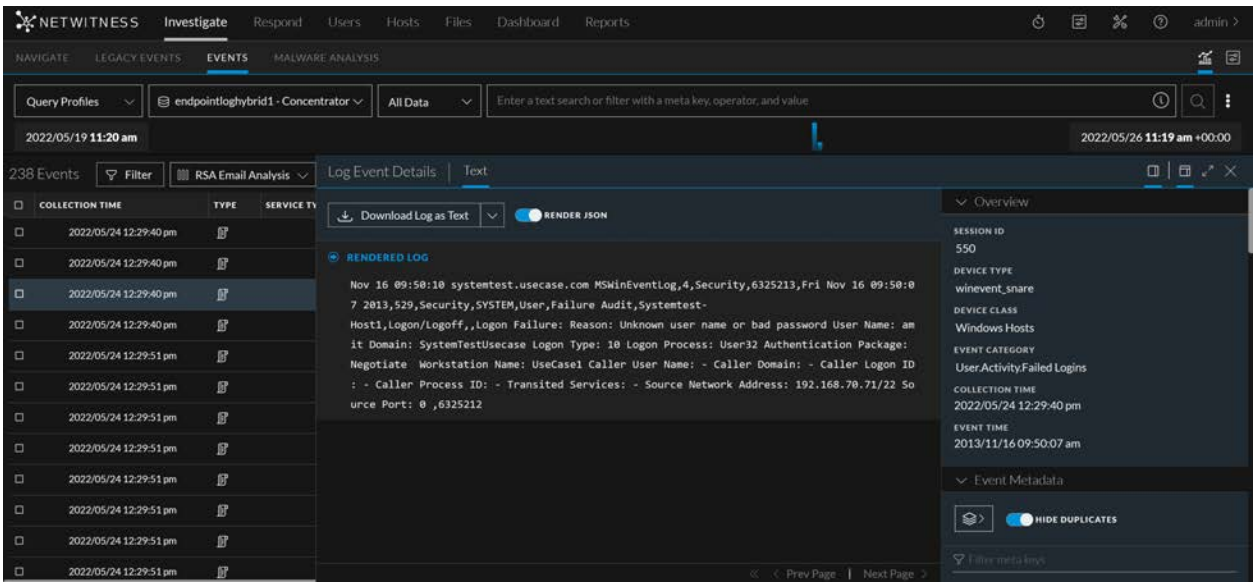


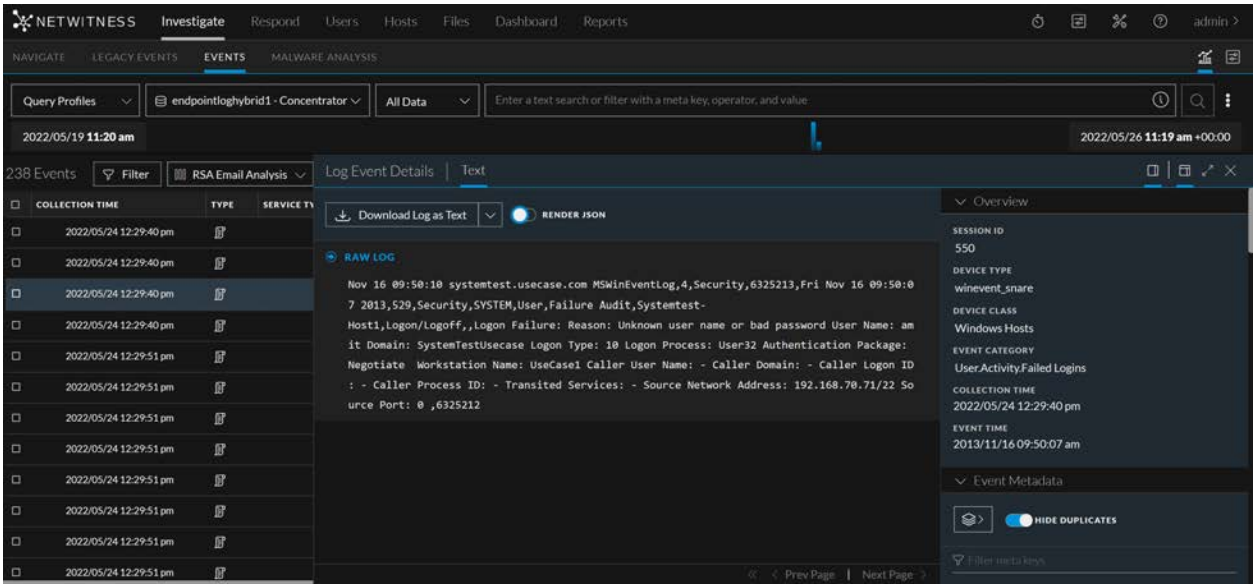
Note: Version 11.1 handles large payloads differently; the payload for a single event is limited to 2500 packets. When the packet limit is reached, a warning in the footer advises the limit has been reached and provides the total number of packets in the event. For Version 11.1, the Show More option is still available for messages that are truncated; however, the entire text of the message is not visible without downloading the raw payload.

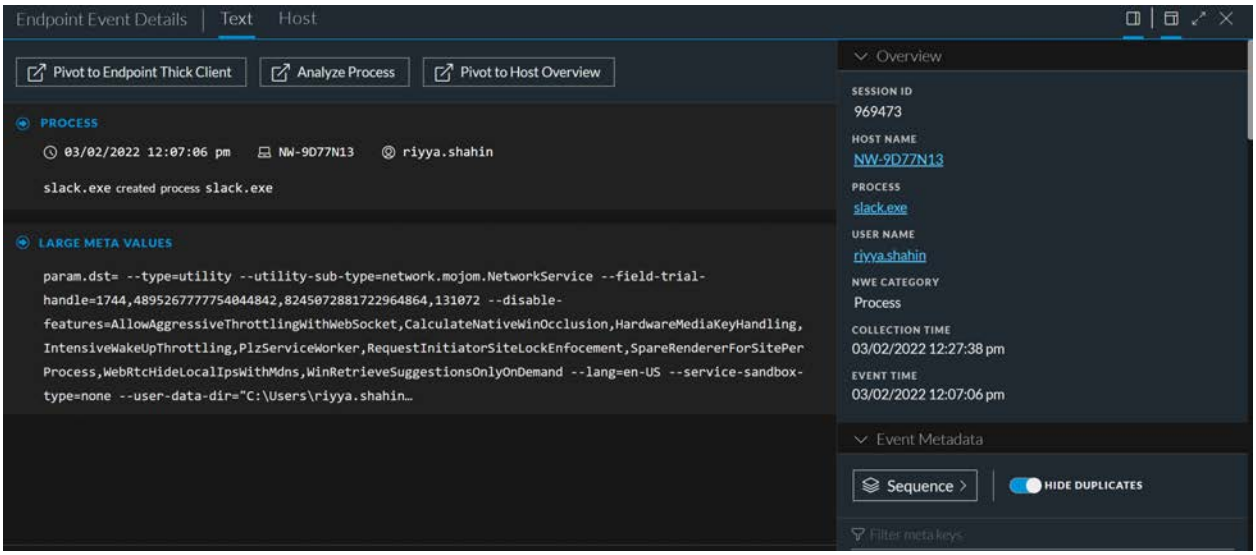
In the text reconstruction, network events, log events, and endpoint events are presented differently.

- For network events, the reconstruction provides the direction of the packet (Request or Response) and contents of each packet in text format. If you are reconstructing a network event, the text reconstruction is scrollable. When you scroll, the text identification information and the Request and Response labels remain visible rather than scrolling out of view.
- Log events and endpoint events have no request or response; only the raw event is displayed in the Text tab. Endpoint events include additional information relevant to an endpoint event.
- (Version 11.5.1) Log events that contain JSON snippets are rendered in an easily readable JSON tree view with nested indentations if the RenderJSON option is enabled.

For each type of event (network, log, or endpoint), there are differences in the Overview panel and the options for downloading the event. Below is an example of the text reconstruction for each type of event: a network event, a log event, and an endpoint event.





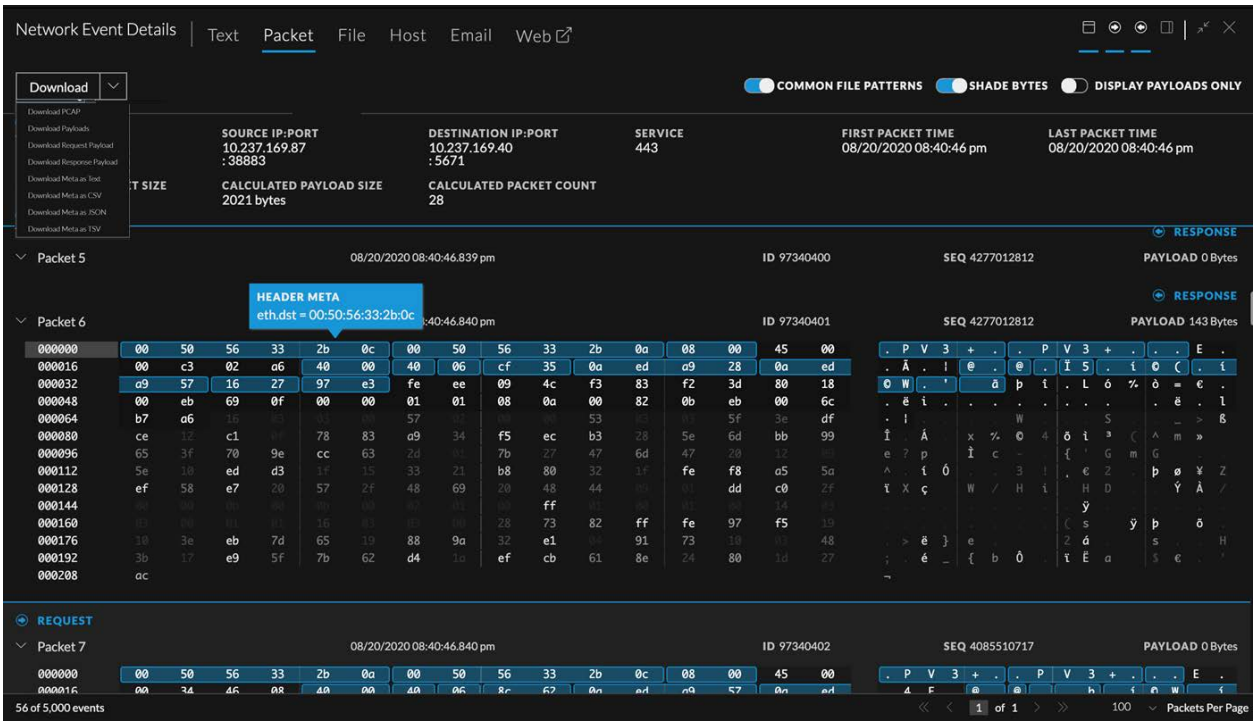


Note: The calculated packet count, calculated packet size, and calculated payload size in the Event header may be different than the same statistics in the Event Meta panel because the metadata is sometimes written before event parsing completes and may include packet duplicates.

Packet Reconstruction

The packet reconstruction is for network events. The panel is scrollable, and the packet identification information and the Request and Response labels remain visible rather than scrolling out of view. In the Packet tab, the headings provide the direction of the packet (Request or Response), the packet number, the packet start time, the packet ID and the sequence, and the payload size. All packets begin with a header, and some packets have a footer. Pagination controls add flexibility when paging through packets.

The metadata in the hexadecimal and ASCII data is highlighted in blue; when you place the cursor over the highlighted metadata, the meta key/meta value information is displayed in a hover box.

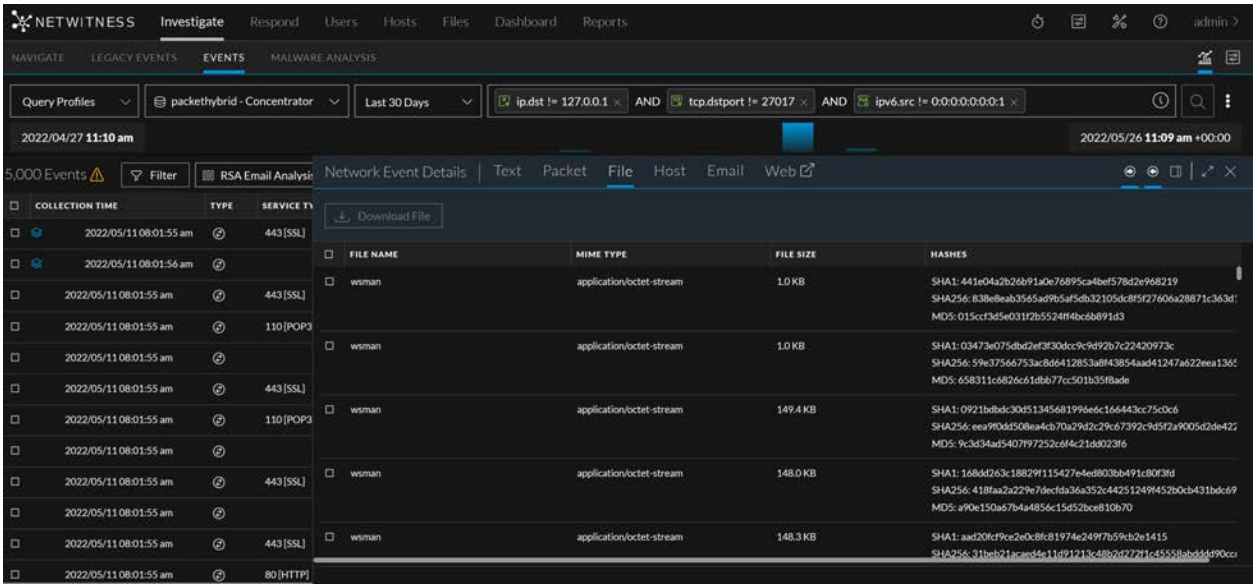


Common file signatures are highlighted with an orange background. When you place the cursor over the highlighted text, the description of the file type is displayed in a hover box.

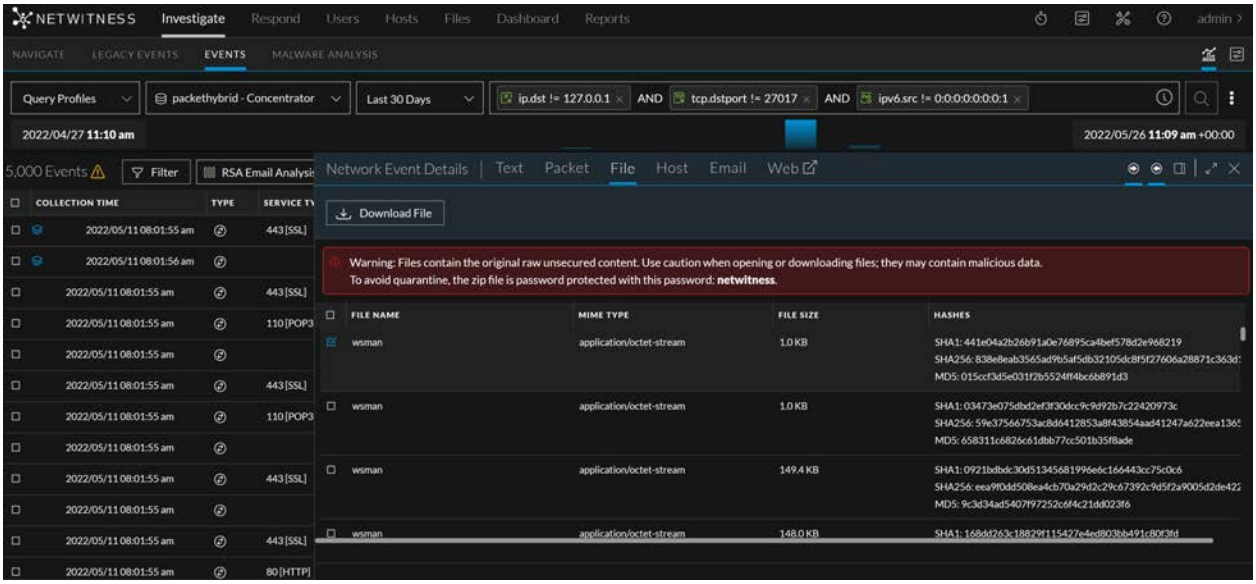


File Reconstruction

The file reconstruction shows a list of files associated with the selected network event. This is an example of the file reconstruction.



You can select one or more files, or all files, to export to your local file system. When files are selected, The Download File options becomes active and reflects the number of files selected.



Caution: Caution is advised when unzipping and opening files that are associated with a default application; for example, an Excel spreadsheet may automatically open in Excel before you have a chance to verify it is safe.

Host Information

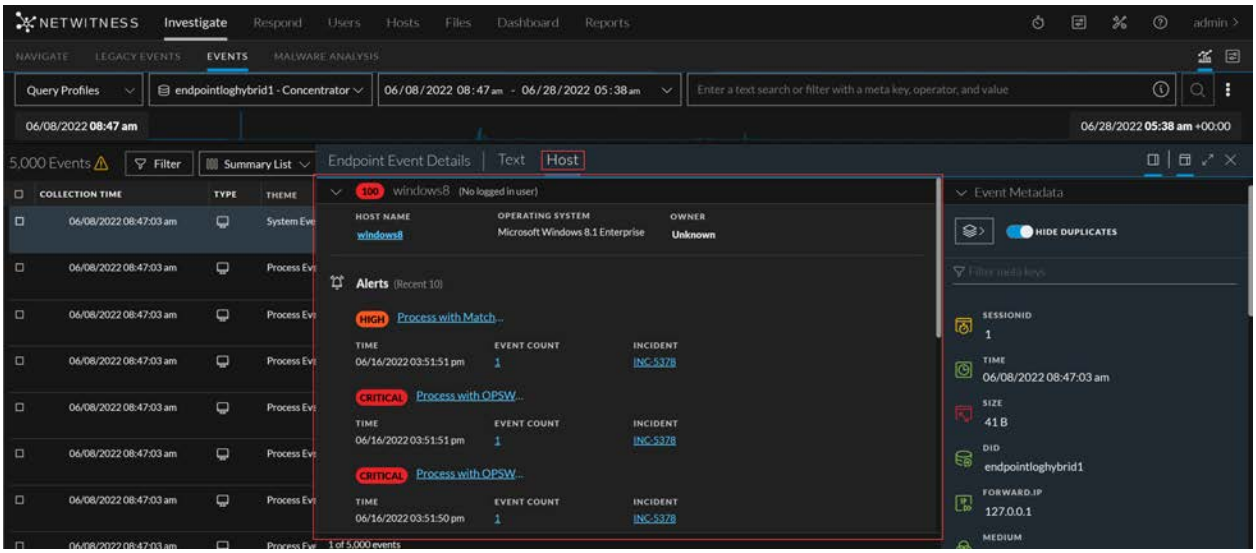
Host information is displayed for network and endpoint events with the endpoint data.

You can filter the events in the Events view using the Events Filter panel (version 11.5 and later). For more information, see Filter Results in the Events View.

Note: Endpoint data is displayed only if you have an Endpoint deployment, and the Endpoint agents are configured for expanded network visibility. For more information to enable expanded network visibility, see "Creating Groups and Policies" in the *Endpoint Configuration Guide*.

Note: For expanded network visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.

Below is the example of the host information.



- The hosts with the closest matching process are listed in the order of event time.
- By default, the first host is expanded, and you can view additional information such as:
  - Host details – This provides details on the host's operating system and the owner (logged in user) associated with the host.
    - To investigate on the host name, click the **Host name** link highlighted in blue. For more information, see "Investigating Hosts" in the *NetWitness Endpoint User Guide*.
    - To investigate alerts associated with the user, click the **owner** link highlighted in blue. For more information, see "Investigate High-Risk Entities" in the *NetWitness UEBA User Guide*
  - Process details – This provides details like risk score, process name, reputation, event time, on hosts, signed status, process ID, signer, user, launch arguments, SHA256, and path.

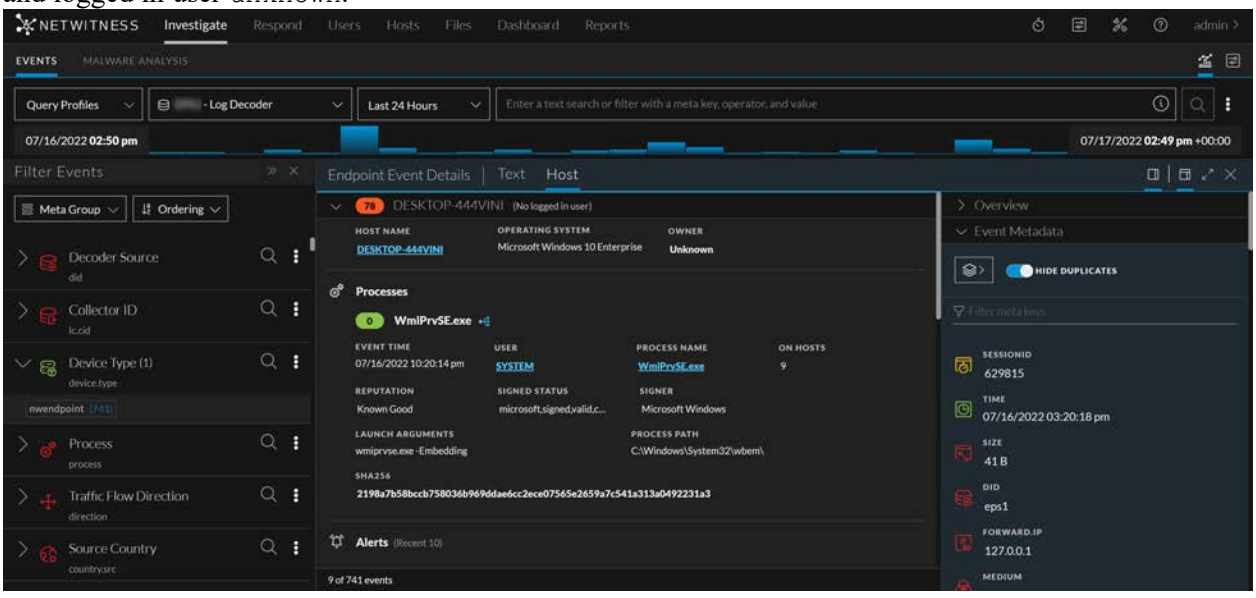
- Click to open the process tree. By default, the process tree will open the process details of last 14 days. The icon to open the process tree does not appear when the process tree is not available.
- To investigate on the process, click the **process** link highlighted in blue. For more information, see "Investigating Files" in the *NetWitness Endpoint User Guide*.
- To investigate alerts associated with the user, click the **user** link highlighted in blue. For more information, see "Investigate High-Risk Entities" in the *NetWitness UEBA User Guide*.

◦ Alert details – It displays the recent ten alerts associated with the host. These alerts can be from endpoint, network, and log events. You can click **View All** to open the host details page. The host details page lists all the alerts that contribute to the risk score. You can click on the alert name to open the alert details. For information on how to review an alert, see “Reviewing Alerts” in the *Net Witness Respond User Guide*. This section provides the following details.

- Severity – Displays the severity of the alert.
- Time – Date and time when the alert was triggered.
- Event Count – It displays the number of events that triggered the alert. To view the events associated with the alert, click the **EVENT COUNT** link highlighted in blue. The **EVENT COUNT** link is available only when the events are from the same source.
- Incident – Lists the incidents associated with the respective alerts. To view the details and respond to an incident, click the **INCIDENT** link highlighted in blue. For more information, see "Responding to Incidents" in the *NetWitness Respond User Guide*.

You can hover over the meta values of the host name, process, user, owner, and SHA256 to view additional information about the specific metadata. For more information on context look up, see Look Up Additional Context for Results.

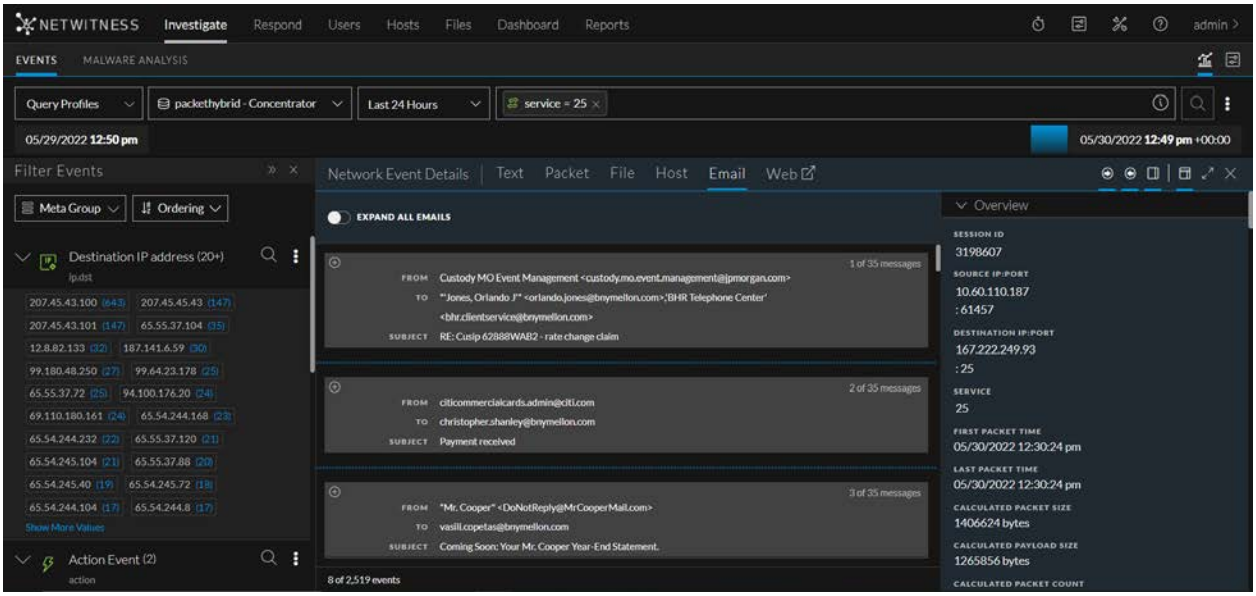
Below is an example of the Host Information tab with a single host, process, and user associated with the selected network event. The `WmiPreSV.exe` is the process associated with the host `DESKTOP-44VINI` and logged in user `unknown`.



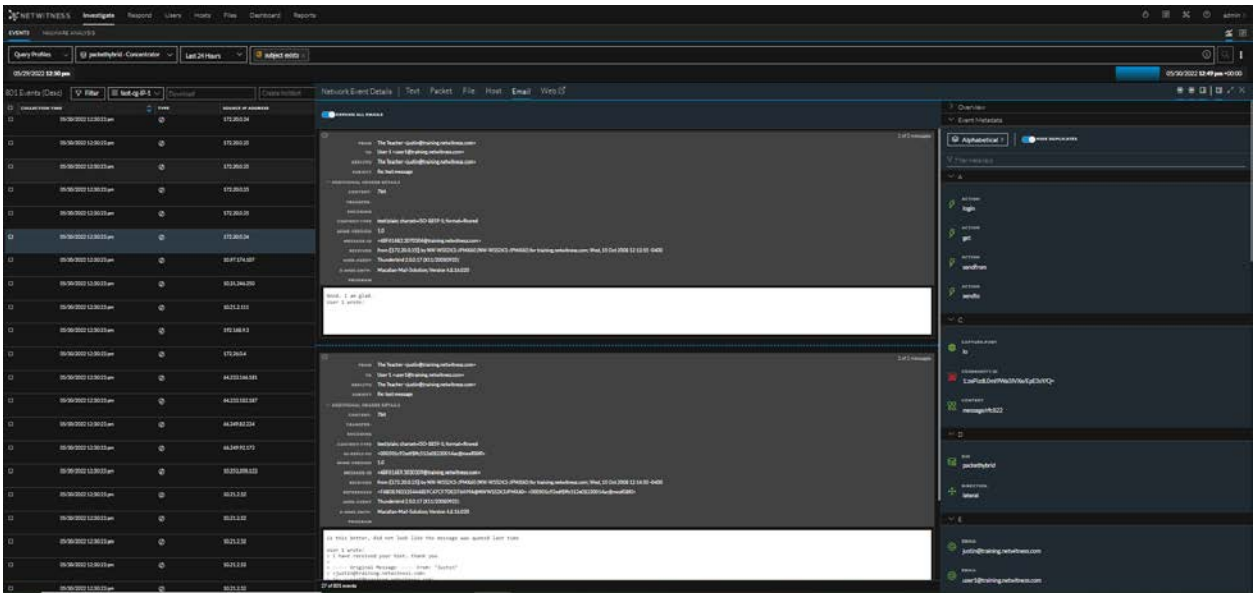
Note: You may see multiple hosts and processes triggered for the selected network event; in such cases, the host from where the event is triggered first is listed first and then the other hosts where a similar event is triggered.
For example, if `10.63.0.240` IP address is assigned to Host1, and User1 is logged in to the machine and accessed `www.nyu.edu/` using Chrome. Meanwhile Host1 is powered off (within a span of 30 minutes), and the same IP address is assigned to Host2. The user logged in is User2 and accesses `www.nyu.edu/` using Internet Explorer. In this case, network events for the endpoint data are as follows:
- Hostname - `Host1`, `Host2`
- Process - `chrome.exe`, `iexplore.exe`
- User - `User1`, `User2`

Email Reconstruction

In 11.7 and later, if the analysts need to review all email contents in a single session, then they can click on the Expand All Emails toggle button by navigating to **Investigate** > **Events** > **Email** view.

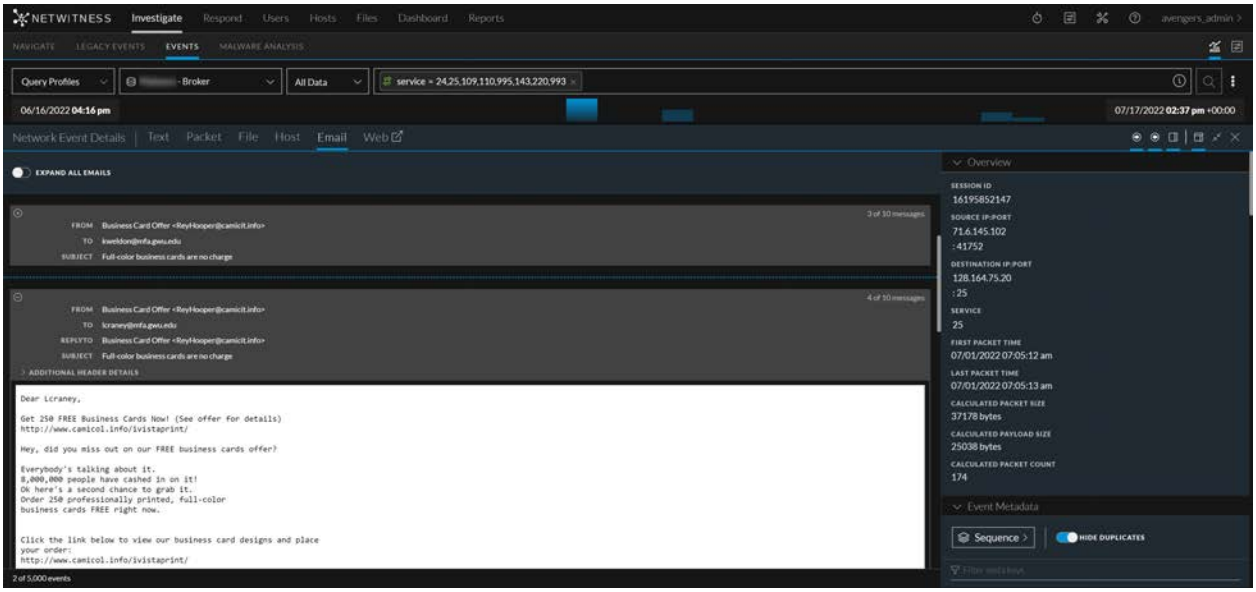


When the Expand All Emails toggle button switched on, the email content is displayed in an expanded form.



When the Expand All Emails toggle button is switched off, the email content is displayed in a collapsed form. If there are no emails to display the toggle button is disabled.

The email reconstruction shows a list of emails associated with the selected network event. This is an example of the email reconstruction.



- By default, a single email is expanded and multiple emails are collapsed.
- If an email contains attachments, you can download attachments as described in Download Data in the Events View.
- **Caution:** When you download and open attachments from an email, they may contain malicious data.

  An external link in an email cannot be accessed. Clicking an external link displays a **Link Address**

popup window that provides the actual link.

- When an email body is too long, **Showing %** is displayed in the beginning of the email. To view the remaining content, click **Show Remaining %** at the bottom of the email.
- If an event contains a web email supported by the `alias.host` metadata of `mail.google.com`, `mail.live.com`, or `mail.yahoo.com`, a message is displayed with a link to view the reconstruction for the associated session in the Event Reconstruction page. If not, a "*No Email reconstruction is available for this event*" message is displayed.

Analyze Events in the Events View

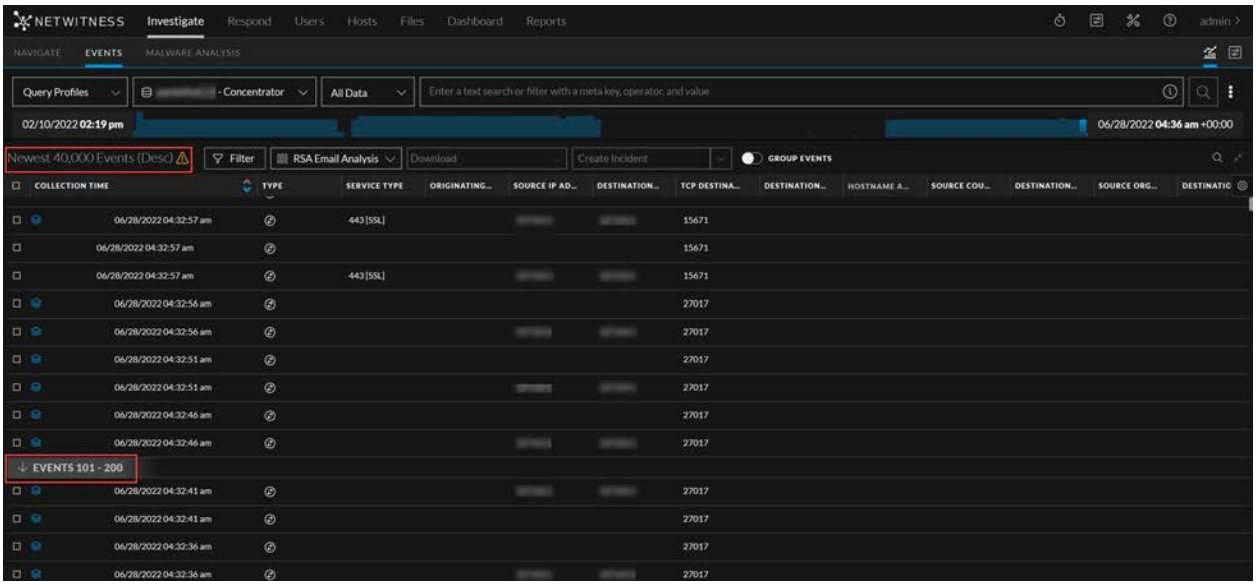
Note: In Version 11.4, the Event Analysis view was renamed as the Events view, replacing the Legacy Events view as the default view analyzing events. Information regarding Events view features prior to version 11.4 also applies to the 11.3 and earlier Event Analysis view.
The Legacy Events view is disabled by default, but the administrator can enable it as described in "Configure Investigation Settings" in the *System Configuration Guide*.

After a query is submitted in the Events view, the Events panel opens with a list of sequential events. The events listed here meet two conditions:

- They match the submitted query.
- They include a value for one or more meta keys required by the selected column group. If you change the column group while viewing the Events list, the original query with the new column group is resubmitted. Unsubmitted query changes made to service, time range, or filter, are ignored.

How Results Are Loaded and Sorted

There is a configurable limit on the number of events that can be loaded; the default value is 5,000. Administrators can configure the limit as described in the *System Configuration Guide*. The events begin loading into the Events panel; a progress bar at the top of the list tracks progress while events are loaded. Events with the earliest collection time are loaded first and a row number indicator of the form " EVENTS xxx - xxx" is inserted in the list after every group of 100 events as shown in the following figure.



A spinner is displayed while the events are loaded. If that count is greater than or equal to the threshold, a message under the spinner advises that fact and directs you to the Query Console for more details. As data begins loading, the message is removed and the spinner remains until all events are loaded. When all events have loaded, one of these messages is added to the bottom of the list:

- "All events loaded."
- "Reached the 5,000 event limit. Consider refining your query."
- "Retrieved 4,000 of 5,000 events prior to query cancellation."

A message at the top of the list indicates the total number of events loaded, if the 5,000 event limit has been reached, and the sorting method in effect:

- The message when fewer than 5,000 events are listed is: "xx,xxx Events"
- The message when more than 5,000 events are listed is: "Oldest 10,000 Events (Asc)"

If the number of events that match the query exceeds the limit of 5,000 events, the newest or oldest 5,000 events in the time window are loaded in ascending order. The portion of events loaded is based on the sort order. For example, if 300,000 events match your query and the sort sequence is set to Ascending, the oldest 5,000 events are loaded by default. You can change this by changing the sort order to Descending and the newest 5,000 events are loaded. Ascending sort, which loads the oldest events first is usually the best setting for investigating network events. If you want to view the newest 5,000 events in the time window, you can change the Default Event Sort Order to Descending in the Event Preferences dialog.

The sorting method for the list is configured in the Event Preferences dialog (see Configure the Events View). Any change in the setting goes into effect the next time you submit a query. The Default Event Sort Order from the Events Preferences dialog is saved in the database and persists after logging out and logging back in.

- **Unsorted** (default for Version 11.4.1): To list events as processed by the Core services. Unsorted is faster because it streams back the events as soon as a match is found versus waiting for all Core services to respond and then displaying them in the chosen order.
- **Ascending** (default for Version 11.4 and earlier): To list the events with the earliest collection time first. The earliest collection time first is well suited to most investigations. When investigating logs, you may want to change the sort sequence to latest collection time first.
- **Descending**: To list the events with the latest collection time first. The latest collection time first is often useful for investigating logs.

Actions to Refine the Events List

When the results are loaded in the Events panel, you can take the actions to refine the list:

- Select a column for sorting events (Use Columns and Column Groups in the Events List).
- Select a set of meta keys (column group) that help with a specific type of investigation (Use Columns and Column Groups in the Events List).
- Apply a query profile (Use Query Profiles to Encapsulate Common Areas for Investigation).
- (Version 11.5) Filter events by pivoting through the metadata Drill into Metadata in the Events View.

Actions to Analyze Events

The rest of this section has procedures for working in the Events view and adjust reconstructions to bring interesting data into focus.

- You can download events and create an incident in Respond.
- Clicking an event in the Events panel opens the Event Details panel to a tab that shows an event reconstruction (text, packet, file, email, and web) or the tab that shows host information for network events that are enriched with endpoint data (Version 11.5).
- The Events panel and the Event Details panel can be open at the same time.
- Within the Packet tab and the Text tab, you can use additional features to adjust the way the reconstruction is displayed and bring interesting data into focus.

Select the Analysis Type for an Event

To select the analysis type for an event, with an event open in the Event Details panel, click one of the tabs: **Text**, **File**, **Host** , **Packet**, **Email**, or **Web**.

- If you chose **Hosts**, the host information from extended endpoint data is displayed.
- If you chose **File**, **Text**, **Packet**, or **Email**, the reconstruction is displayed.
- If you chose **Web**, the reconstruction of the single event opens in a new tab. This is the same reconstruction of a session used in the Legacy Events view (see Reconstruct an Event in the Legacy Events View).

Note: The packet reconstruction is only available for network events.

Adjust the Display of Requests and Responses

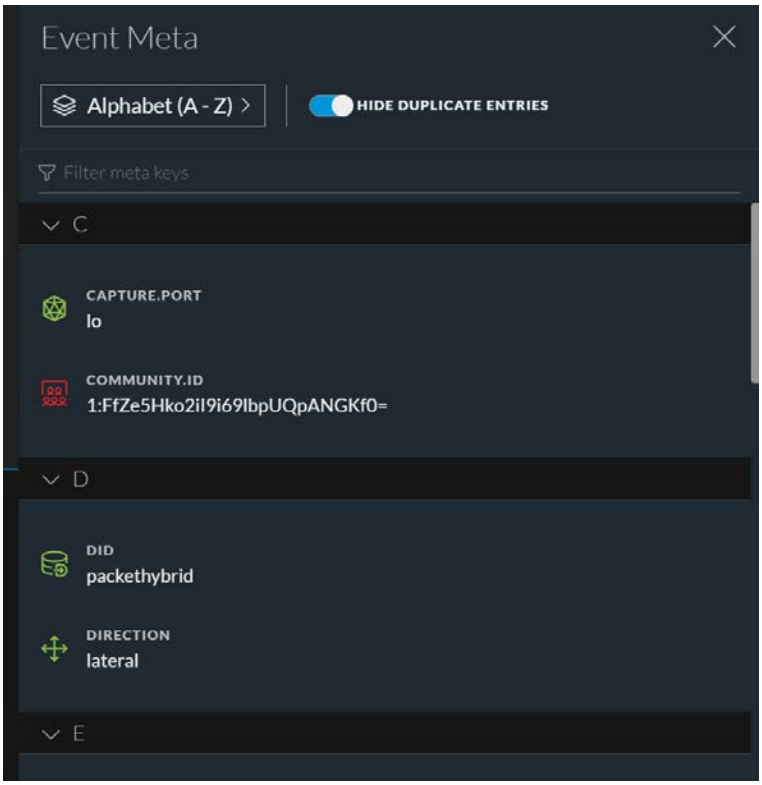
For analysis types that have requests and responses, you can select which side of the conversation to show, Request , Response , or both, click one or both of the direction icons. The reconstruction is refreshed with the selected information.

Note: If you do not see any data, you may have deselected both Request and Response. You must select one of the two to see data displayed.

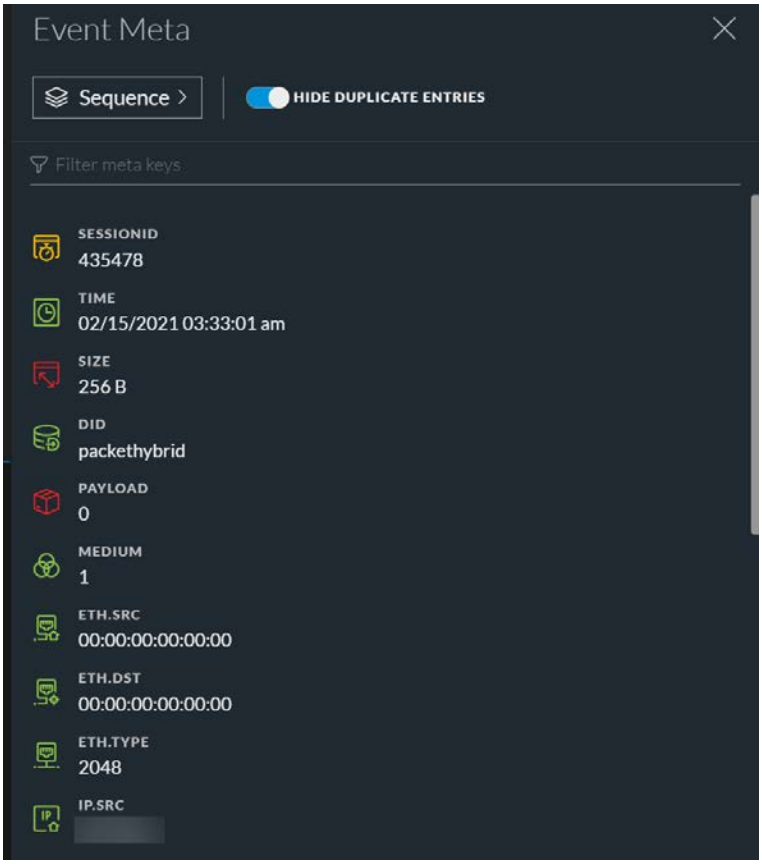
View Associated Metadata for an Event

When examining events in the Text tab, Packet tab, or Filetab, you can click  to show the associated metadata in an adjacent panel, the Event Meta panel.You can change the order of the metadata listed in the Event Meta panel to better find what you are looking for. Metadata can be organized by the sequence in which they were generated or alphabetically by meta key. This figure illustrates metadata organized as meta keys sorted alphabetically."

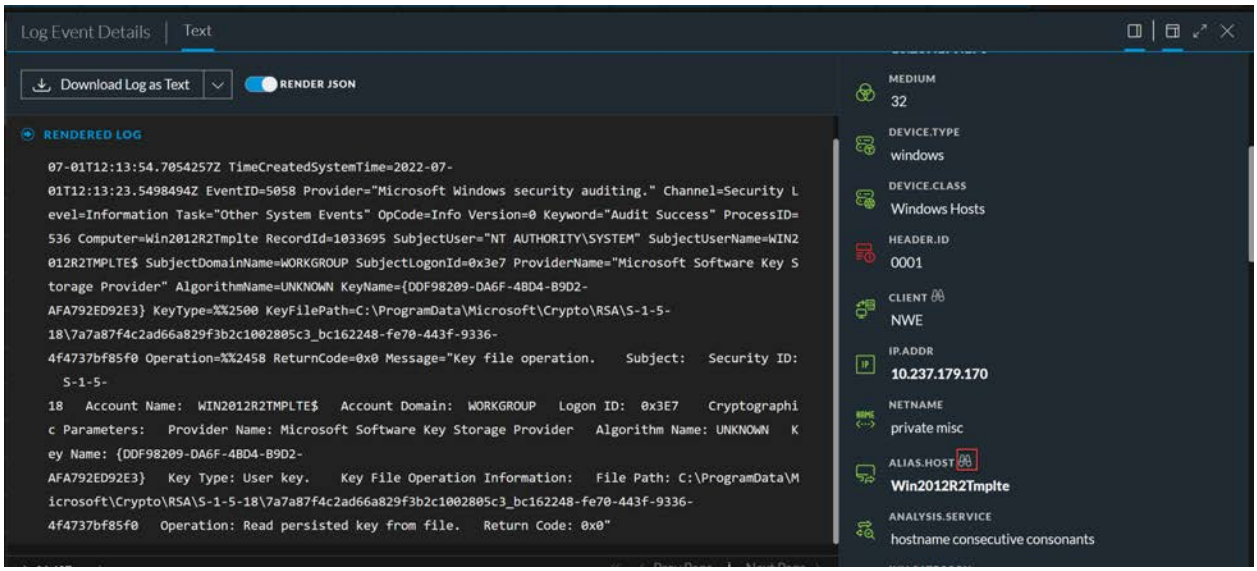
You can change the order of the metadata listed in the Event Meta panel to better find what you are looking for. Metadata can be organized by the sequence in which they were generated or alphabetically by meta key. This figure illustrates metadata organized as meta keys sorted alphabetically.



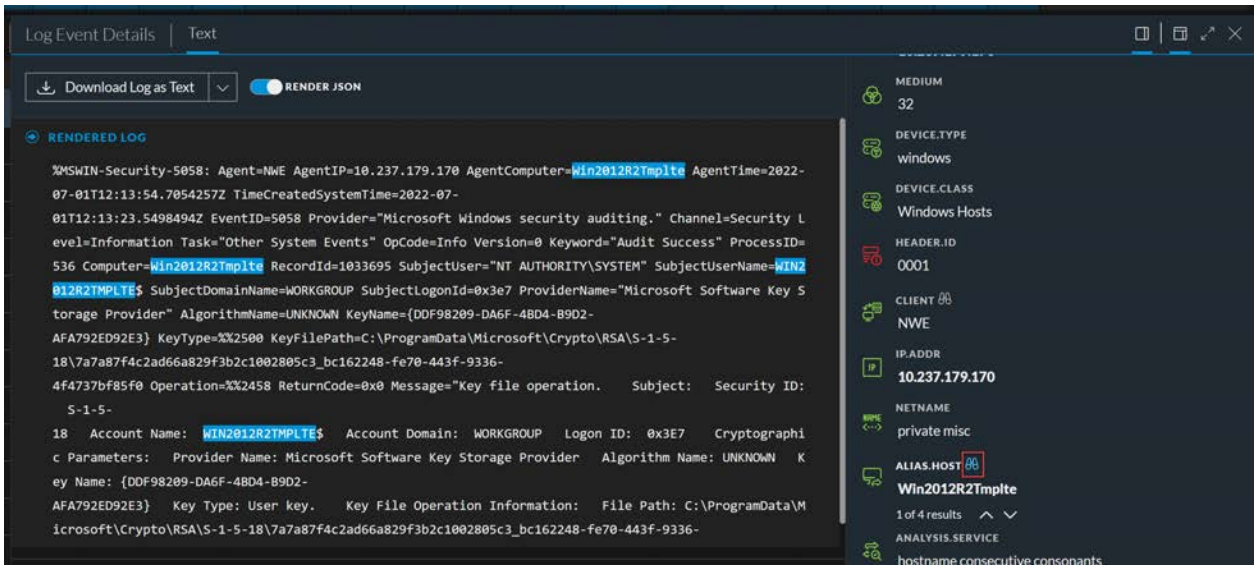
This figure illustrates the same metadata presented in the sequence in which it was generated.



From 12.0 and later, the meta key and meta value pairings now display a binocular icon while viewing a text reconstruction in the Event Meta panel, indicating the search option. This enhancement helps the analysts to visually see directly the indication (binocular icon) rather than going through the list of all metadata to figure out which ones may be searched. This figure is an example of binocular icon marking a searchable meta key.



Clicking the icon triggers a search for the meta key/meta value pair (case-insensitive) in the Text tab and each instance is highlighted. This is an example of binoculars with a blue background after clicking a searchable meta key/ meta value combination.

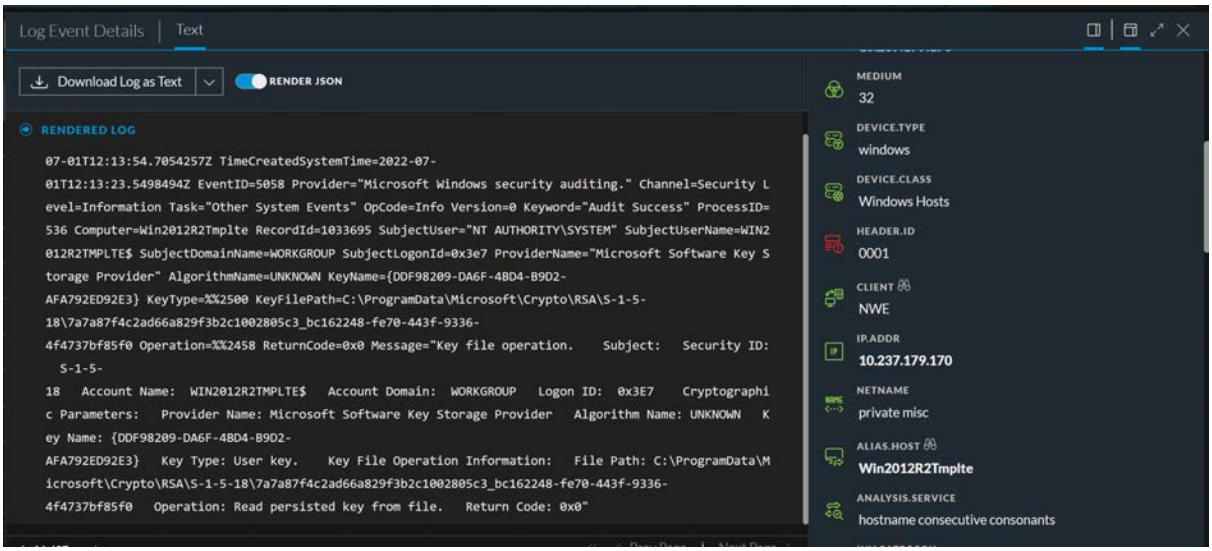


In the Event Meta panel, the highlighted row has a count of the results and up and down arrows that you can use to quickly find each result in the Text tab. You can view each highlighted location of the data that triggered generation of the meta key, going forward to view the next, and back to view the previous.

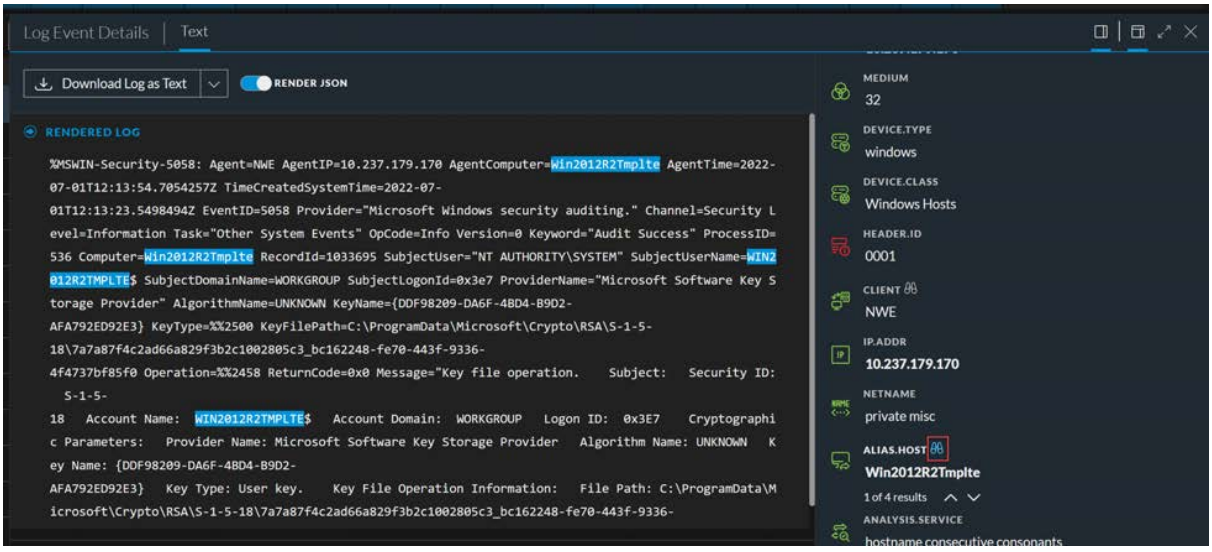
Only meta keys that have relevant values inside the RAW text are searchable. You can search only one meta key at a time. If the value is currently hidden due to truncation of a text entry with more than 3000 characters, the text entry is expanded to reveal the found meta value.

To search the raw text for meta values that triggered a meta key:

1. Open a network event in the Text tab and click  to open the Event Meta panel.



2. Search for the meta key/meta value pairs in the list until you see a binoculars icon next to a meta key.
3. To search for the value in the raw text, click a row that has the binoculars icon, indicating it is searchable.
   If no relevant occurrence of the value is in the text, the value that you are searching for is highlighted in the Event Meta panel and nothing is highlighted in the Text tab.
   If one or more relevant instances of the value is found in the Text tab, each occurrence is highlighted. The value that you are searching for is highlighted in the Event Meta panel and the up/down arrows for scrolling are visible.



4. To remove the highlighting, click the same meta key/meta value pair in the Event Meta panel, click a different meta key/meta value pair in the Event Meta panel, or close the Event Meta panel.
   The highlighting is removed from the raw text.

Note: When a meta value is more than 255 characters, you can hover over that meta key to view the complete value.

Show or Hide the Event Header

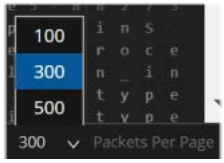
To hide the Event Header in the Packet tab, Text tab, or File tab, providing more vertical space for the data, click . Clicking the icon again shows the Event Header.

Page Through Events in the Packet and Text Tabs

Pagination controls allow more flexibility in paging through a list of packets or text. In the Packet tab, you can select the number of packets to display per page, and your selection is saved across logins to the NetWitness Platform application. When a control is unavailable, the control is dimmed; for example, when you are viewing page 1, the  and  controls are dimmed.

Note: Pagination controls are available in Version 11.2 and later of the Text tab. When in the Text tab, you must navigate manually to the last page before the last page control icon is available.

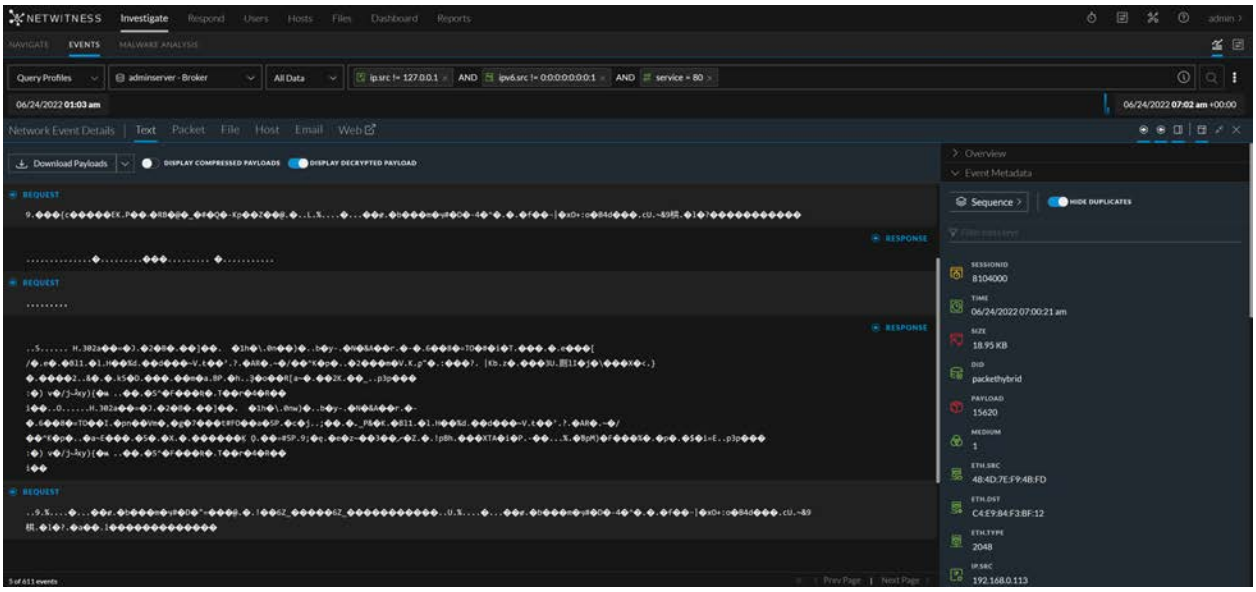
To use pagination controls:

1. With an event open in the Events view, click the current number of packets per page (**50, 100**, **300**, or **500**), and select the new number of packets per page from the drop-down menu.
   

2. To page forward or back, use the page control icons:
   Click  to go to the next page.
   Click  to go to the last page.
   Click to go the previous page.
   Click to go to the first page.
3. To go to a specific page, type a page number in the page number field .

Expand Truncated Text Entries in the Text Tab

A reconstruction of a network event in the Text tab may include requests and responses of many hundred thousands of characters, and scrolling through a long entry that is not of interest can waste time. To save time, text entries that have more than 6000 characters are truncated to show only the first 2000 characters. This example shows an entry that has more than 2000 characters, and a message in the header indicates the percentage of total characters that is being displayed.



You can see that 46% of the characters (the first 2000) are displayed, and click **Show Remaining 54%** to reveal the rest of the entry.

If you search for metadata seen in the Event Meta panel while text is truncated in the Text tab, the truncated text is searched. If the metadata exists inside hidden text, the text entry expands to reveal the text with the found metadata.

Perform URL and Base64 Encoding and Decoding in the Text Tab

If a network session being reconstructed in the Text tab contains Base64 or URL encoded strings, you can decode a string to better understand the session. If the session contains decoded strings for Base64 or URL, you can view a string in its encoded form in order to search for additional instances of the encoded text in other sessions.

When viewing any network session that contains encoded text in the Text tab, you can select a subset of the text within a single Request or Response to view in either encoded or decoded form. Depending on the content loaded on the Decoder, there may be additional metadata outlining that Base64 or URL encoded data is contained within the session.

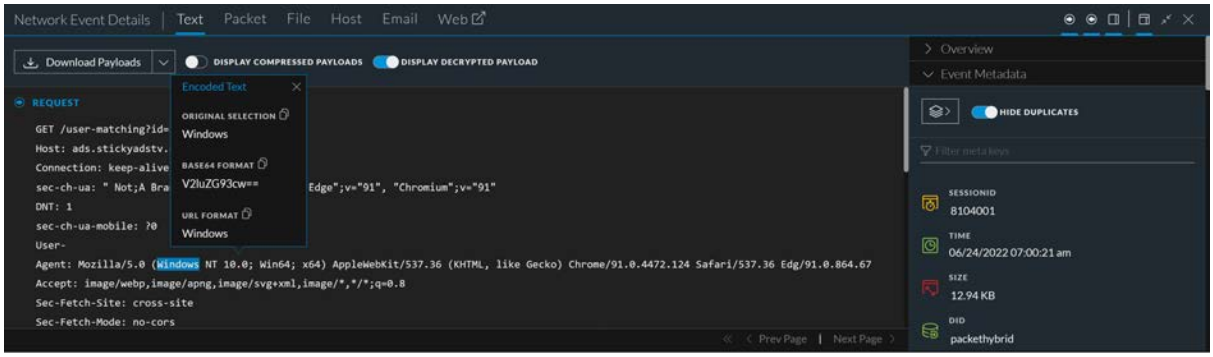
To perform encoding and decoding in the Text tab:

1. In the **Events view**, go to the Text reconstruction of a session that contains encoded or decoded content.

2. To view some decoded text in encoded form, drag to select the text within a Request or Response.
   A menu offers options to encode and decode.



3. Click **Encode Selected Text**.
   The encoded text is displayed in a hover box, which remains in place until you click ✕, select different text in the Text tab, close the Events panel, select another event for reconstruction, or switch to a different reconstruction view.

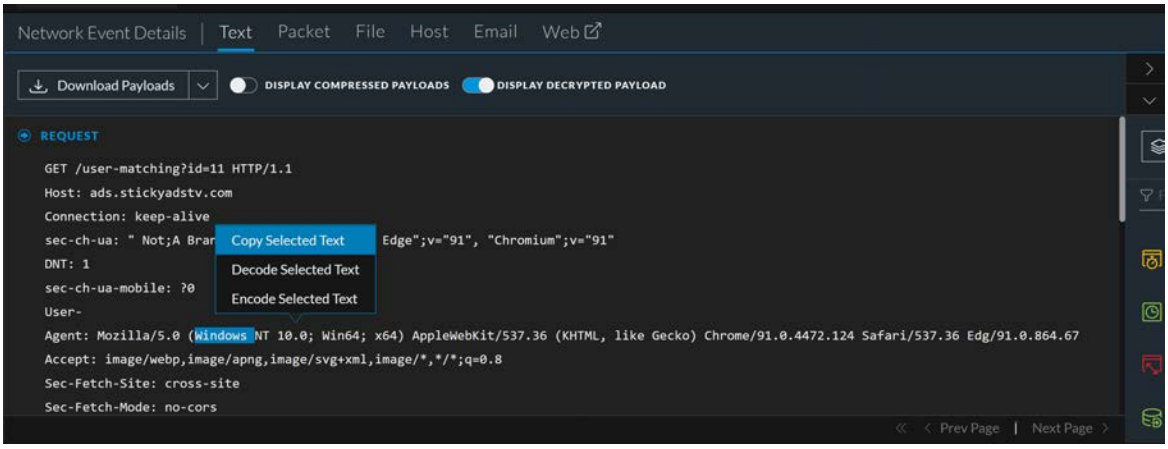


   When a longer text is selected, the hover box is scrollable and large enough to fit the entire selected text and the decoded text.

4. If the session contains encoded text that you want to see in decoded form, drag to select the text within a Request or Response.
   A menu offers options to encode and decode.

5. Click **Decode Selected Text.**
   The decoded text is displayed in a hover box, which remains in place until you click ✕, select different text in the Text tab, close the Events panel, select another event for reconstruction, or switch to a different tab in the Event Details panel.

6. If you want to copy some text from the text reconstruction, do one of the following:
   a. Drag to select some text, right-click, and select **Copy Selected Text** from the pop-up menu.



   b. Drag to select some text, then select either **Decode Selected Text** or **Encode Selected Text**.
   Select the desired text and type **Control-C**.
   The selected text is copied to the clipboard and available to paste in a query.

7. When finished, click ✕ to close the hover box.

View Decompressed Text in an HTTP Network Session in the Text Tab

When the content of an HTTP network session is compressed and you are viewing the Text tab, NetWitness displays decompressed content by default. This helps you to determine if there are any patterns and view the readable characters. You can switch between a compressed and decompressed view of compressed text.

The toggle for changing between compressed and decompressed text is only displayed in the Text tab, and is enabled only if there is compressed text content.

1. Open the Text tab of an HTTP session that contains compressed content.
   By default the session is reconstructed with the text decompressed, and above the reconstruction, is the **Display Compressed Payloads** toggle switch.



2. To view the same text in its compressed form, click the toggle switch.
   The view changes so that the compressed text is no longer readable, and the switch indicates the Display Compressed Packets is on.

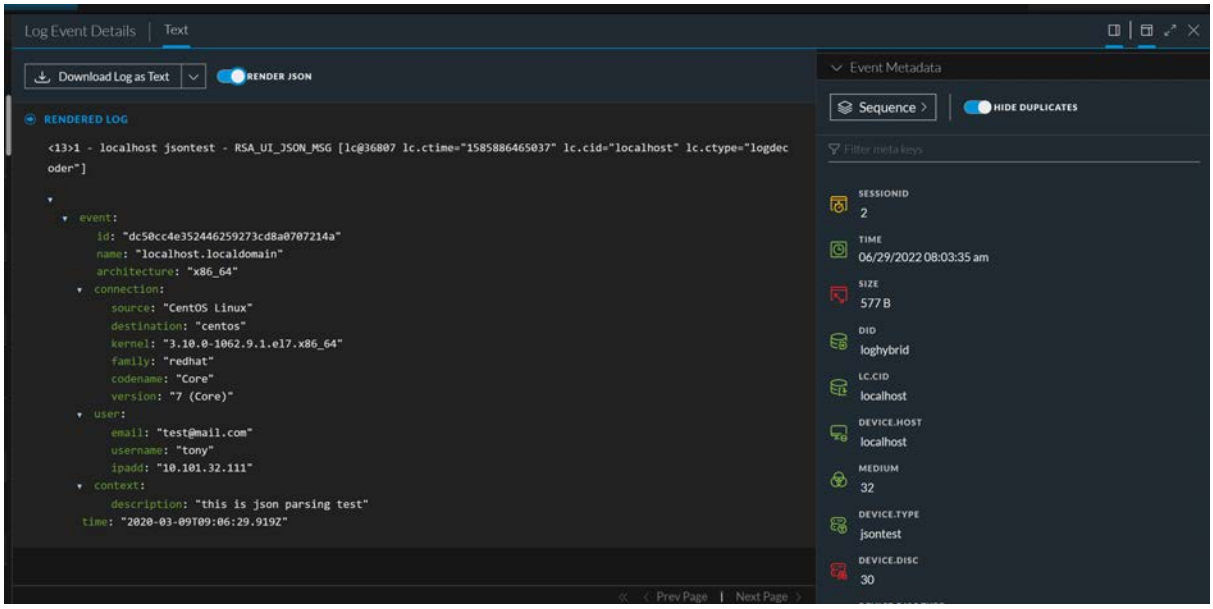


3. To return to the view of decompressed text, click the switch again.

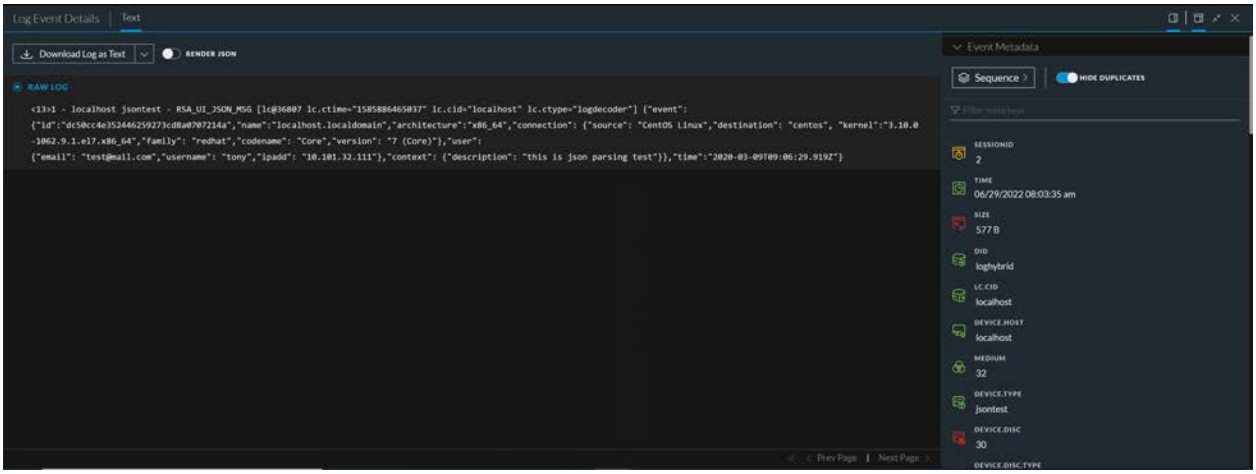
View a JSON String in Tree Format in the Text Tab (Version 11.5.1)

You can view a text reconstruction of a log event in easy-to-read JSON format instead of the raw block format using the Render JSON toggle switch. By default, the switch is enabled, and JSON snippets in a log event are detected and displayed in fully expanded tree format. Invalid JSON snippets are presented as raw text. If you change the setting of the switch, your setting persists in local storage.

1. In the **Events** view, open a log event in Text tab.
   If the raw log contains JSON strings and the Render JSON switch is enabled, all JSON strings found are rendered in tree format.



2. If you prefer to see the log as raw text, click the Render JSON switch.
   The log is rendered as a single block of text without the nested indentation. The setting persists until you change it and will be in the same state when you log in again.

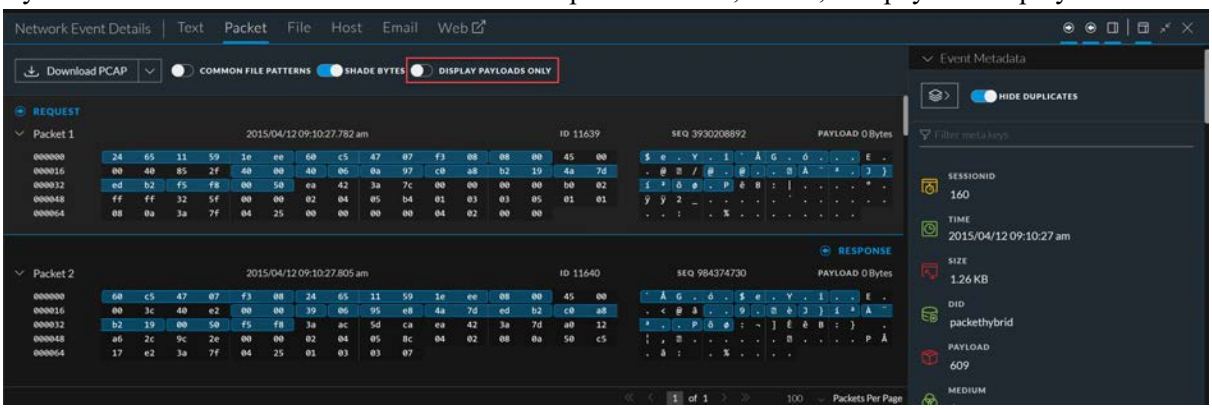


Use the Payload Only Option in the Packet Tab

When viewing a reconstruction of a network session in the Packet panel, you can hide the header and footer bytes. The view changes so that only the payload is visible and contiguous same-side packets are concatenated together to make the payload more readable and understandable. This setting persists until you change it or refresh the browser.

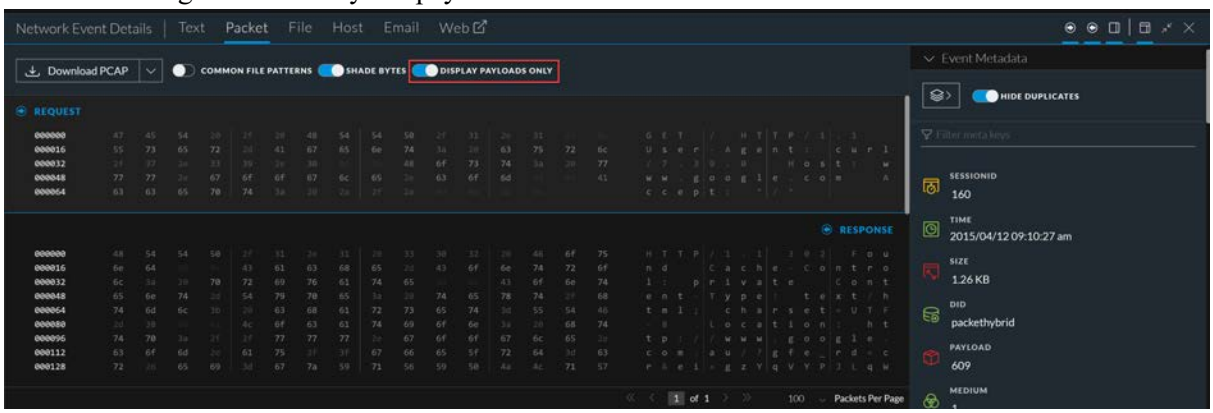
- With the Display Payloads Only option off, the number of packets, packet header, packet footer, and payload are displayed.
- With the Display Payloads Only option on, no packet header and footer bytes are displayed. Only the packet content of 16 hexadecimal bytes per line and the corresponding ASCII per line is displayed.

1. In the **Events** view, go to the Packet tab of a network session.
   By default the session is reconstructed with the packet header, footer, and payload displayed.



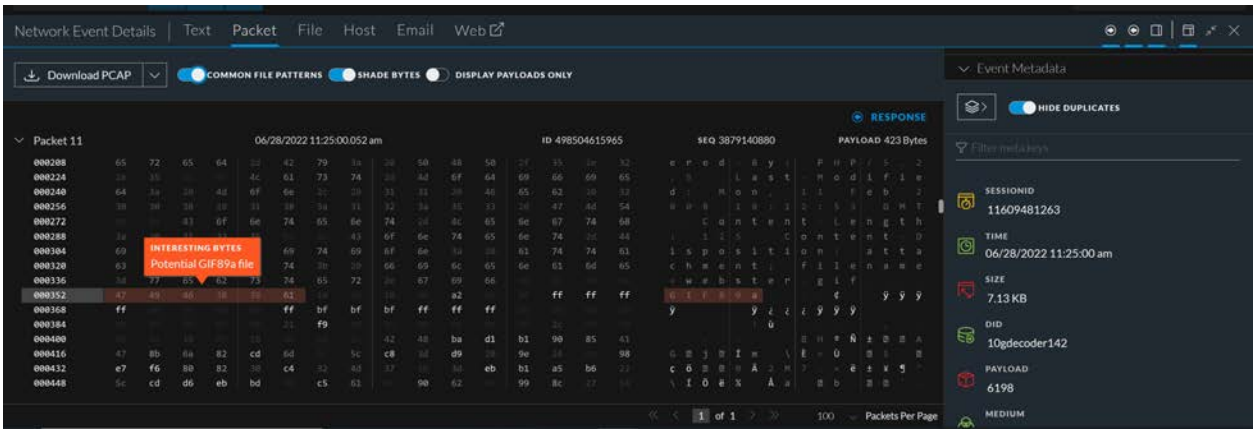
2. To change the view to show only the payload for each packet, click the **Display Payloads Only** toggle switch.

The view changes so that only the payload is visible.

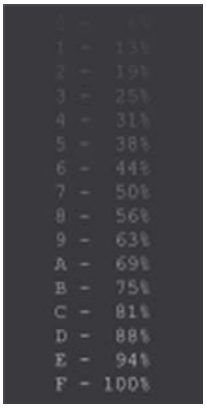


View Highlighted Bytes in the Packet Tab

When you open a reconstruction in the Packet tab, the significant header bytes in each packet are highlighted in blue, and the payload bytes are distinguished using shading to help you understand the contents of the packet. This figure shows the default appearance of the packet reconstruction with highlighting and byte shading.



The Shade Bytes option adds shading to identify the different hexadecimal bytes (00 to FF) using degrees of highlighting. Bytes near the lower range are more transparent, and bytes near 255 are more opaque. Both hexadecimal and ASCII bytes are shaded. This is an example of the shading applied to each hexadecimal byte.



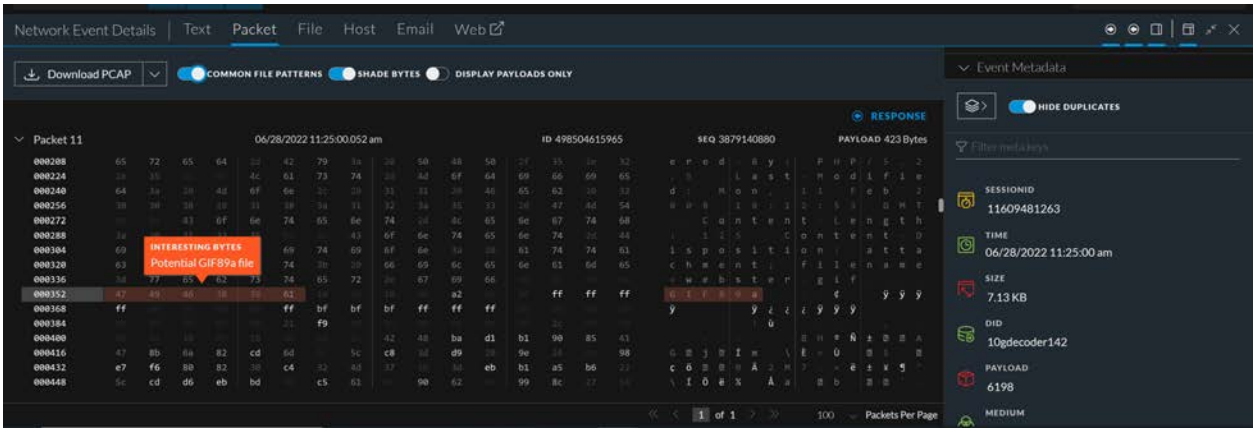
The Shade Bytes switch controls the shading of bytes. When you set Shade Bytes on or off, your setting persists until you change it or refresh the browser.

Highlight Common File Types in the Packet Tab

In the Packet tab, analysts can show or hide highlighting of certain common file types based on the file signature. When the Common File Patterns feature is turned on, the magic number bytes in the file signature are highlighted in the payload and you can hover over the highlighting to see the potential type of file. In this example, `42 4d` is highlighted in the hexadecimal payload and `BM` is highlighted in the ASCII payload. When you hover over the highlighted bytes, the potential file type associated with the magic number is provided in a hover box.

To view common file signatures in the Packet tab:

1. With a reconstruction open in the Packet tab, turn on the **Common File Patterns** option.
   If there is more then one highlight in view, all are shown.
2. To view the hover box, place the cursor over the highlighting.



These are the file types and corresponding magic numbers that are highlighted if present in the payload:

File Type	Hexadecimal Signature	ASCII Encoding
DOS Executable / Windows PE	4D 5A	MZ
Portable Network Graphics (PNG)	89 50 4E 47 0D 0A 1A 0A	PNG
JPEG	FF D8 FF	JPEG
JPEG/JFIF	4A 46 49 46	JFIF
JPEG/Exif	45 78 69 66	Exif
GIF	47 49 46 38 37 61	GIF87a
GIF	47 49 46 38 39 61	GIF89a
Non-portable Executable	5A 4D	ZM

File Type	Hexadecimal Signature	ASCII Encoding
BMP	42 4D	BM
PDF	25 50 44 46	%PDF
Old Office Document (doc, xls, ppt, msg, and other)	D0 CF 11 E0 A1 B1 1A E1	ÐÏ.à¡±.á
ZIP file formats and formats based on it, such as JAR, ODF, OOXML	50 4B	PK..
7-Zip File Format (7z)	37 7A BC AF 27 1C	7z¼¯'
Java Class File, Mach-O Fat Binary	CA FE BA BE	Êþº¾
Postscript	25 21 50 53	%!PS
Unix/Linux Shell script	23 21	#!
Executable and Linkable Format (ELF) executables	7F 45 4C 46	.ELF

Reconstruct an Event in the Legacy Events View

When viewing a list of events in Legacy Events view, you can safely create a reconstruction of the event in a readable form that matches the original. By default, the initial view of a reconstructed event is the most suitable format (Best Reconstruction); for example, web content is reconstructed as a web page; an IM conversation is displayed with both parts of the conversation. Each user can select a different default reconstruction in the Profile > Preferences view.

You can also open a reconstruction from the Navigate view if you know the Event ID of the event.

In the reconstruction, you can:

- Select event information to view. Possible values are: request data, response data, both request and response data.
- Select the reconstruction type: details, text, hex, packets, web, mail, or IM.
- Export raw logs.
- Export the event as a PCAP file.
- Extract any files available in the event.
- Extract all the meta data associated with the event.

  Caution: Be careful when clicking a link to a file in the Reconstruction. If your system has an application associated with the file, or the browser is capable of opening them, and the attachments are malicious, they can negatively affect your system.
- Display the event in a separate window or tab (depending on your browser configuration).
- If you are viewing the reconstruction as a preview in the current view, you can page forward to the next event and back to the previous using the navigation buttons in the bottom left corner.

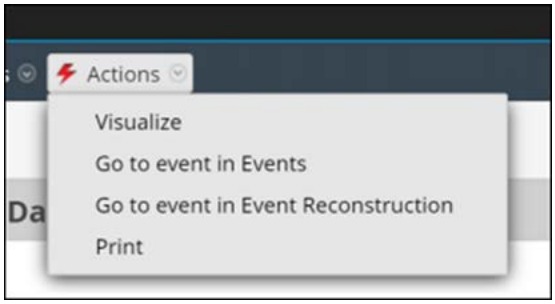
Note: Reconstruction Settings and Reconstruction Cache Settings allow an administrator to manage application performance for Investigation (as described in the *System Configuration Guide*). When analysts reconstruct sessions, two situations can affect performance and results.
- Some large events contain many thousands of source packets. Reconstructing these sessions can degrade application performance.
- In some cases, the reconstruction cache can present incorrect content; for this reason, NetWitness cleans cache that is older than a day every 24 hours. Between the daily cache cleanings, certain actions my result in stale cache being used for a reconstruction, and if the need arises, administrators can manually clear cache for one or more services that are connected to the current NetWitness Server.

Reconstruct an Event Using an Event ID

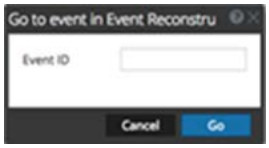
You can reconstruct an event directly from the Navigate view given a known Event ID. You can use this option without executing a query as you usually do when beginning an investigation. A service and time range must be selected to be able to jump directly to an event using just its `eventid`.

To view a reconstruction or event analysis directly from the Navigate view:

1. Go to **Investigate > Navigate** and select **Actions > Go to event in Events** or **Go to event in Event Reconstruction**.

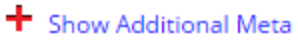


The Go to event dialog is displayed. There are two dialogs, one for Events and one for Legacy Event reconstruction. Both ask for the Event ID.

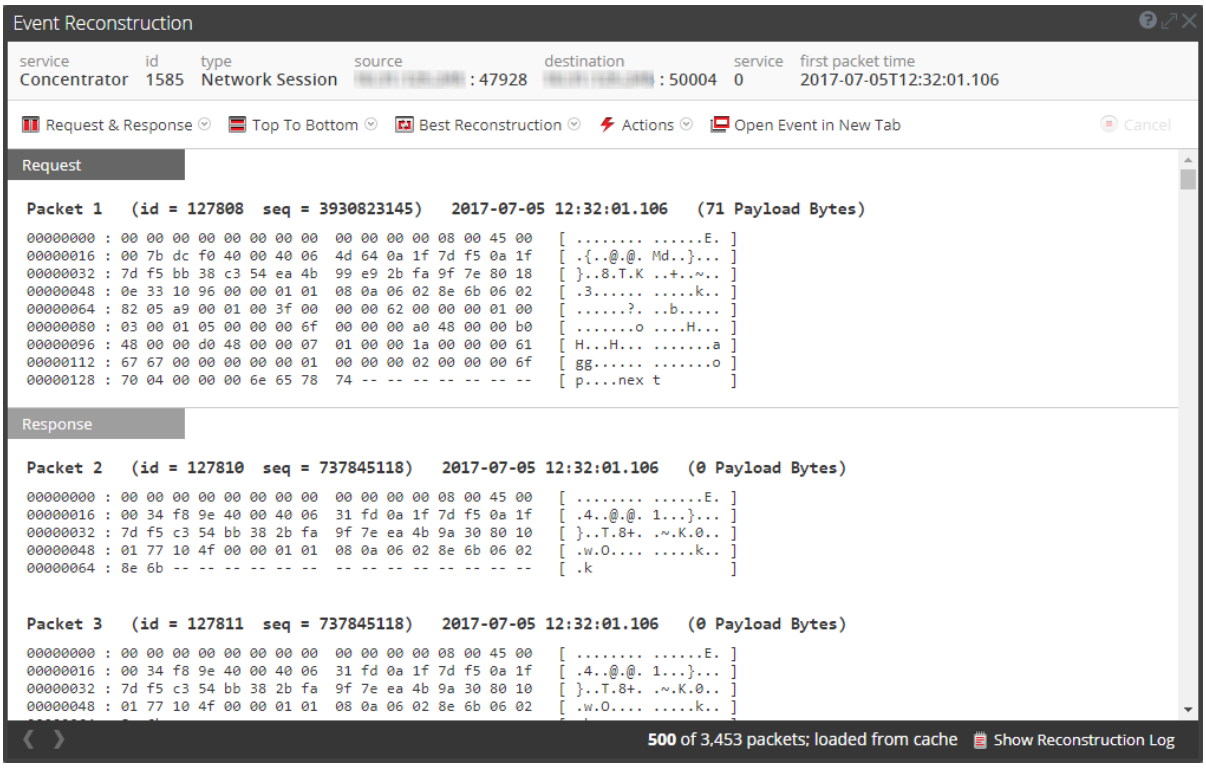


2. In the **Event ID** field, type the ID and click **Go**.
   The specified event is reconstructed in the legacy Event Reconstruction view or the Events view.

Reconstruct an Event from a Drill Point in the Navigate View

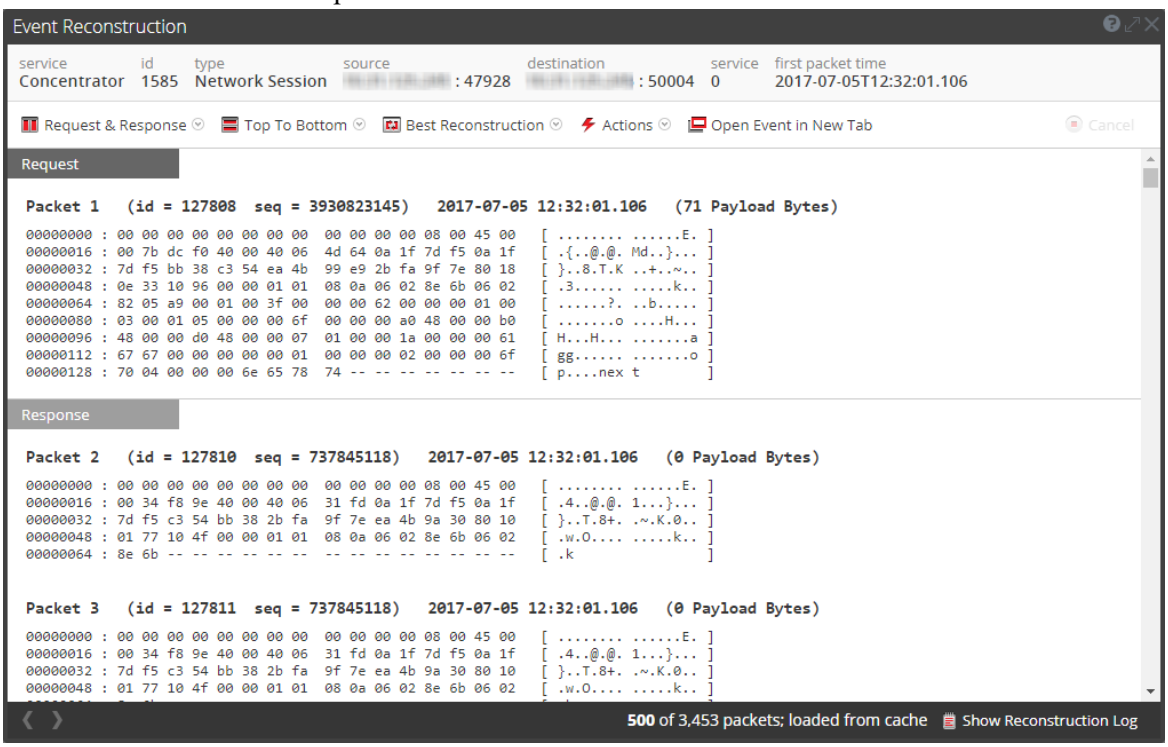
1. Click the count (the green number following a value) for a value in the Navigate view to open a drill point in the **Events** view.
2. To show all meta data, click 

.
3. To open an event reconstruction in the Legacy Events view, select an event to reconstruct and select **Actions > View Event > Preview Inline**.
   The Event Reconstruction opens in a popup window in the same view. By default, NetWitness displays the best reconstruction for the event determined by the event content or the reconstruction that you have selected in the Default Session View setting for Investigation. You can use the options in the Event Reconstruction toolbar to change the reconstruction method, view side-by-side results, export an event, open an email attachment, extract files, and open the event in a new tab. The toolbar options vary depending on the type of event being reconstructed (network event, log event, or endpoint event). This is an example of the reconstruction for a network event.



4. To preview a reconstruction of the next event, click  in the lower left corner of the reconstruction or to preview a reconstruction of the previous event, click .
5. To open an event reconstruction in a new tab, do one of the following:
   a. In the **Legacy Events** view, select an event to reconstruct and select **Actions > View Event > Open in New Tab**.
   b. In the **Event Reconstruction** toolbar of previewed reconstruction, click **Open Event in New Tab** in the toolbar.

The Event Reconstruction opens in a new tab.



View Side by Side or Top to Bottom

To select the way requests and responses for an event are displayed:

1. In the **Event Reconstruction** toolbar, click **Top to Bottom** or **Side by Side**.
2. In the drop-down menu, select the information you want to see in the event: **Side by Side** or **Top to Bottom.**
   The reconstruction is refreshed with the selected information.

Select Event Information to View

To select what event information to view:

1. In the **Event Reconstruction** toolbar, click **Request & Response**.
2. In the drop-down menu, select the information you want to see in the event: **Request & Response**, **Request**, or **Response**.
   The reconstruction is refreshed with the selected information.

Select Event Reconstruction Type

To select the reconstruction type for an event:

1. In the **Event Reconstruction** toolbar, click **Best Reconstruction**.
2. In the drop-down menu, select the reconstruction type to view: **meta**, **text**, **hex**, **packets**, **web**, **mail**, or **files**.
   The reconstruction is refreshed with the selected reconstruction type.

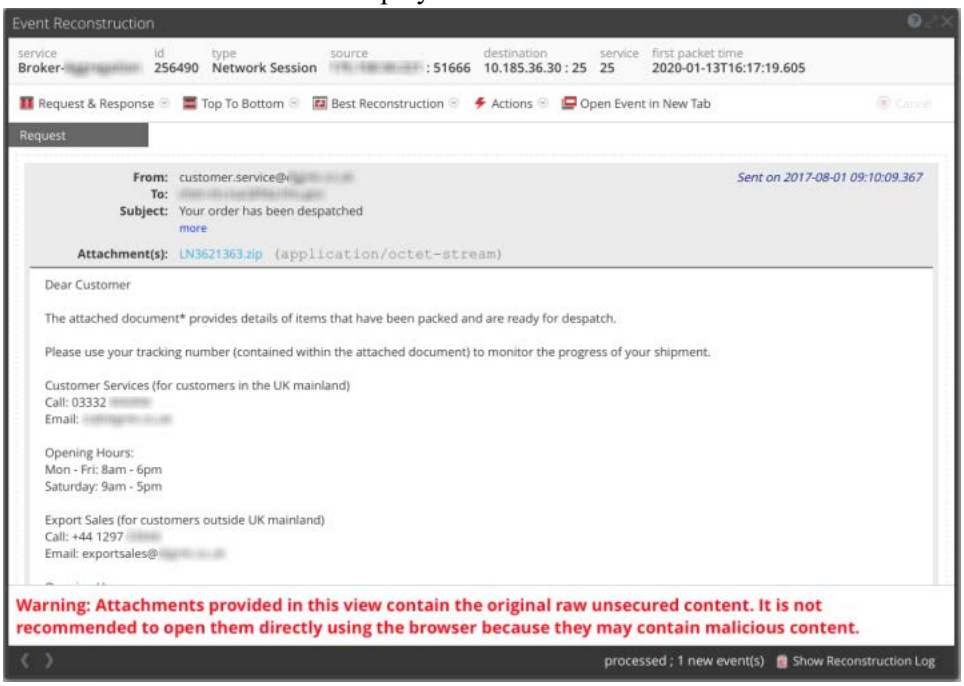
Open or Download an Email Attachment

When viewing a reconstruction of an email that has attachments, you can open supported file types or download the files to the local system.

Caution: Be careful when selecting file attachments. If your system has an application associated with the file attachments, or the browser is capable of opening them, and the attachments are malicious, they can negatively affect your system.

To open or download email attachments:

1. In the **Event Reconstruction** toolbar, select the **View** drop-down and select **View Mail**.
   The Event Reconstruction is displayed.



2. In the **Event Reconstruction** section of the email, click the Attachment.
   If the file type is supported by the browser, the attachment will open in a new tab.
   If the file type is not supported, the Download dialog is displayed so that you can download the attachment.

Export an Event as a PCAP File

The PCAP export option downloads the sessions for the current time range and drill point to a PCAP file. To export an event as a pcap file:

1. In the **Event Reconstruction** toolbar, click **Actions**.
2. Click **Export PCAP**.
3. A confirmation dialog is displayed.
4. Click **OK**.
   The job is scheduled and when complete the PCAP is downloaded to the local file system. In the Profile > Jobs tab, you can download the PCAP.

Extract Files from a Reconstructed Event

The Extract Files option extracts and downloads the files associated with the event. To extract files:

1. In the **Event Reconstruction** toolbar, click **Actions**.
2. Click **Extract Files**.
   The File Extraction dialog is displayed.
3. Select the types of files to extract, and click **OK**.
4. The job is scheduled and when complete the selected file types are downloaded to the local file system. In the Profile > Jobs tab, you can download the files.

Look Up Additional Context for Results

The Context Hub is a centralized service that aggregates data about entities from multiple configurable data sources. This data can extend your investigation with additional context beyond the immediate results of a specific query. For example, the Context Hub can tell you if a given entity has been mentioned in any incidents, alerts, feeds, or community intelligence publications.

To enable viewing of contextual information, your administrator must add the Context Hub service in NetWitness Platform and configure data sources for the Context Hub service as described in the *Context Hub Configuration Guide*. Analysts also need a role with the permission `Context Lookup` as described in "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

When the Context Hub service is enabled and configured, NetWitness provides enrichment data from NetWitness Respond, custom lists, and NetWitness Endpoint directly in the Navigate view, Events view, and Legacy Events view. A visual cue highlights meta values for which enrichment data is available in the Investigate views, and you can click on the highlighted value to look up the contextual information and intelligence. You can look up details and intelligence about elements associated with an event in the Context Hub. These elements, or entities, are identifiers, such as an IP address, a user name, a host name, a domain name, a file name, or a file hash. The data from configured sources, such as NetWitness Endpoint, can help you understand what is happening. In Version 11.5 and later, you can add STIX data sources and view related data using context lookup; associated elements are IP address, file name, file hash, domain name, and URL.

In addition, you can add lists and list values for Context Hub enrichment; you can view lists, edit meta values in an existing list, or create a new list. When you add meta values to a list, you can investigate the meta values using the context lookup option.

Note: In Version 11.2 and earlier, you can look up additional context in the Navigate view or the Legacy Events view, but not the Event Analysis view.

For an analyst to manage lists in Investigate, the administrator must:

- Enable the Context Hub service.
- Assign an analyst role with permission `Manage List from Investigation` to the user who will perform Context Lookup from Investigation views.
- Configure appropriate roles and permissions as described in "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

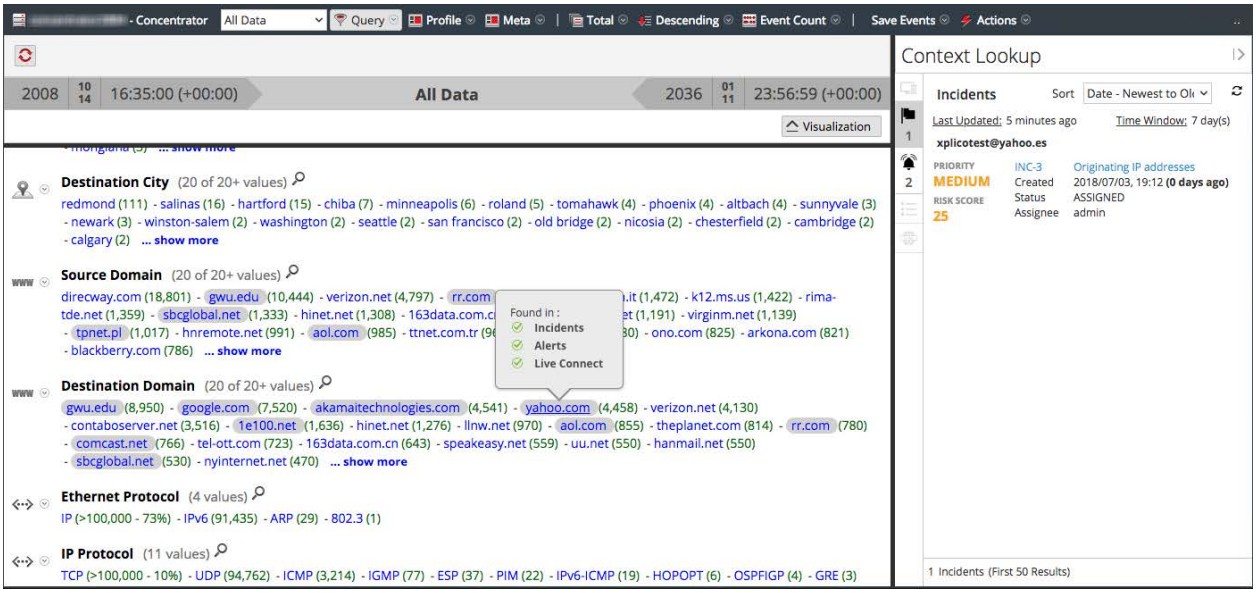
In Version 11.6 and later, you can add REST API data sources and view related data using context lookup. Also, Administrators can now configure specific Context Hub sources (For example, specific lists, Respond, Endpoint, and so on) for context highlighting during investigation. If context highlighting is disabled for a Context Hub source, analysts will see results from all the data sources when opening the Context Panel for a meta value, but the values are not highlighted in the Investigation views. In the View Context:

- No data is shown if the meta values is not highlighted.
- If there are entities common across different data sources, the meta values for those entities are underlined for all the data sources, but the data is shown only for the data source on which context highlighting is enabled.

Open the Context Lookup Panel

In the Context Lookup panel, you can view and explore individual data sources for further investigation. For a detailed description of the information displayed for each data source, see Context Lookup Panel.

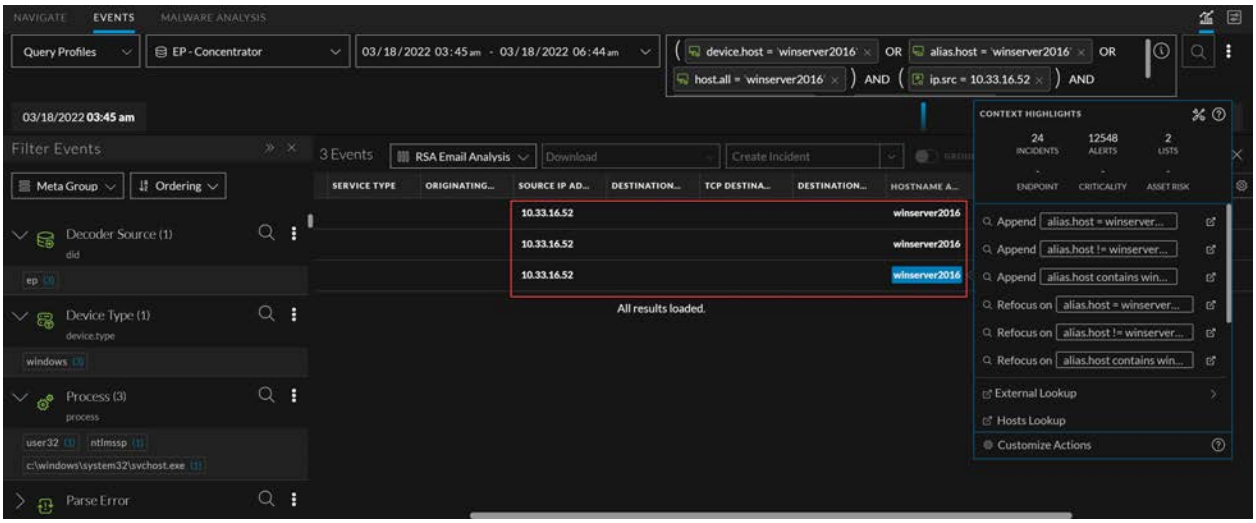
In the Navigate view and Legacy Events view, entities that have associated context data available are highlighted with a gray background; hovering over an entity displays a hover box giving a summary of the available data. When you right-click the entity, the Context Hub queries the configured data sources for relevant information, and the Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available. You can perform another lookup by right-clicking on another entity, and the Context Lookup panel is updated with that entity's information.



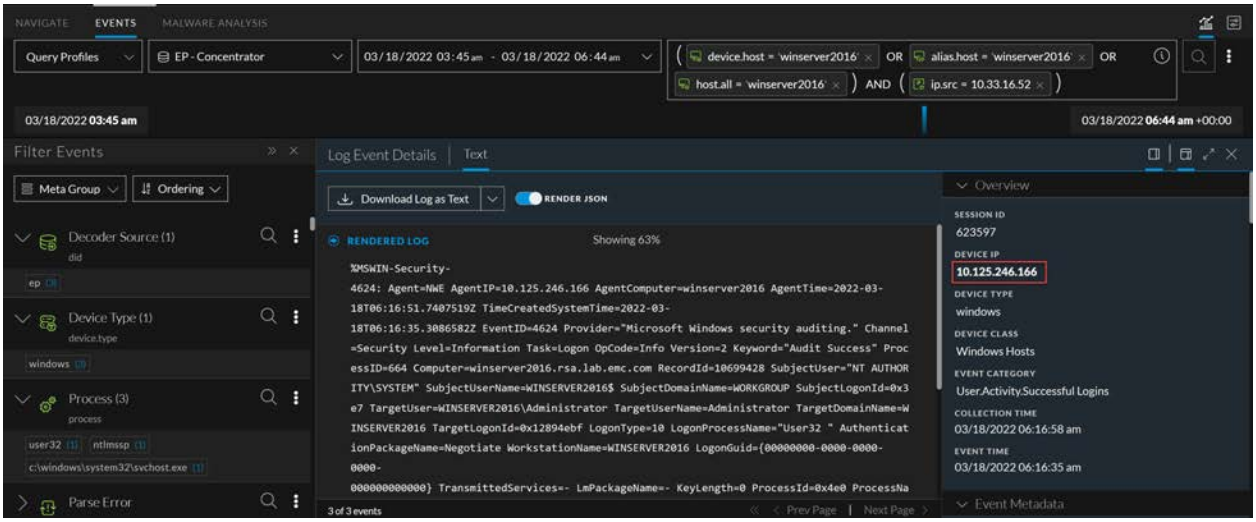
In the Events view, you can see underlined entities in the Events panel, the Event Header, or the Event Meta panel. If an entity is underlined, NetWitness is populating information about that entity type in the Context Hub. There may be additional information available about that entity in the Context Hub.

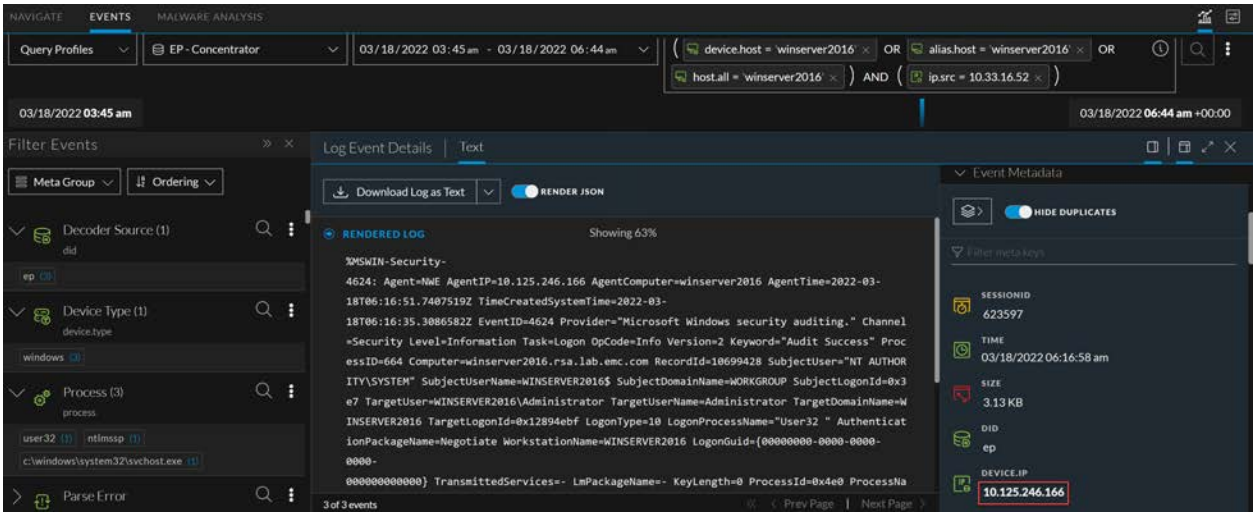
The following figure shows underlined entities in the Events panel with the context tooltip open. The context tooltip has two sections: Context Highlights and Actions.

- The information in the Context Highlights section helps you to determine the actions that you would like to take. It can show related data for Incidents, Alerts, Lists, Endpoint, Criticality, Asset Risk, and STIX. Depending on your data, you may be able to click these items for more information.
- The **Actions** section lists the available actions. In the example, the Add/Remove from List, Pivot to Investigate > Navigate, Pivot to Archer, and Pivot to Endpoint Thick Client options are available.



The following figures shows bolded entities in the Overview and the Event Meta panel.



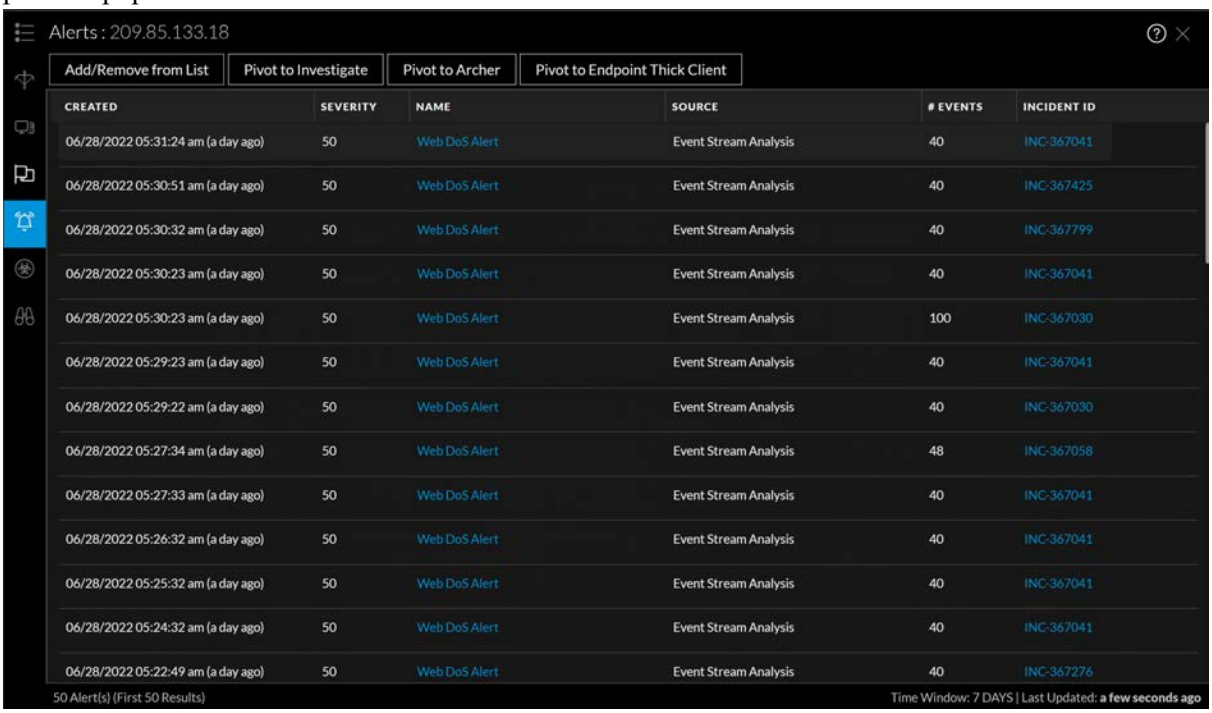


When you click View Context in the context tooltip, the Context Hub queries the configured data sources for relevant information, and the Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available. You can perform another lookup by using the View Context option on another entity, and the Context Lookup panel is updated with that entity's information.

You can also take any available action in the Actions section.

To view information in the Context Lookup panel in the Events view:

1. Hover over different meta values to see the data sources for which data is available.
   A context tooltip displays a list of the context data available for the selected meta value.
2. Click **View Context** in the context tooltip to open the Context Lookup panel.
   The Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available.



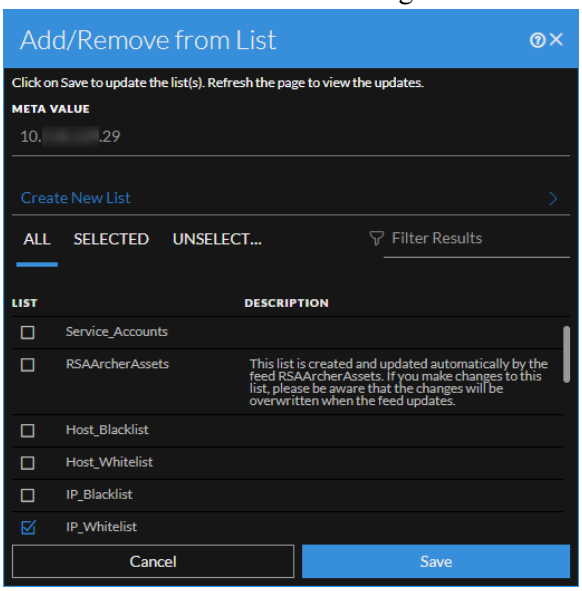
3. To perform actions on an entity, select one of the available actions in the context tooltip: Add /Remove from List, Pivot to Investigate > Navigate, Pivot to Archer, Pivot to Endpoint Thick Client. For more information, see Pivot to Investigate > Navigate (Events View), Pivot to Archer (Events View), Pivot to NetWitness Endpoint Thick Client (Events View), and Add an Entity to a Whitelist.

Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer data source is not responding. Check that the Archer configuration is enabled and configured properly.

Add an Entity to a Whitelist

You can add any underlined entity to a list, such as a Whitelist or Blacklist, from a context tooltip. For example, to reduce false positives, you may want to whitelist an underlined domain to exclude it from the related entities.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over the underlined entity that you would like to add to a Context Hub list.
   A context tooltip showing the available actions is displayed.
2. In the **ACTIONS** section of the tooltip, click **Add/Remove from List**.
   The Add/Remove from List dialog shows the available lists.



3. Select one or more lists and click **Save**.
   The entity is added to the selected lists.

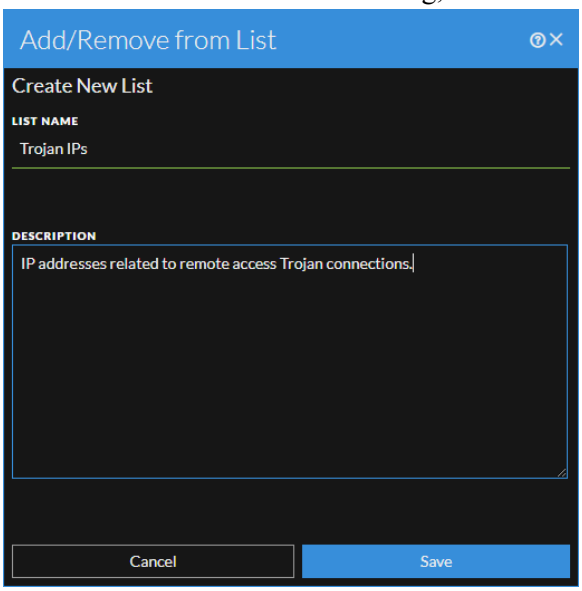
Create a List (Events View)

You can create lists in Context Hub from the Events view. In addition to using lists to whitelist and blacklist entities, you can use lists to monitor entities for abnormal behavior. For example, to improve the visibility of a suspicious IP address and Domain under investigation, you may want to include them in two separate lists. One list could be for domains suspected of being related to command and control connections, and another list could be for IP addresses related to remote access Trojan connections. You can then identify indicators of compromise using these lists.

To create a list in the Context Hub:

1. In the Events panel, the Event Header, or the Event Meta panel, hover over the underlined entity that you would like to add to a Context Hub list.
   A context tooltip showing the available actions is displayed.
2. In the **ACTIONS** section of the tooltip, click **Add/Remove from List**.

3. In the Add/Remove from List dialog, click **Create New List**.



4. Type a unique **List NAME** for the list. The list name is not case sensitive.
5. (Optional) Type a **DESCRIPTION** for the list.
   Analysts with the appropriate permissions can also export lists in CSV format to send to other analysts for further tracking and analysis. The *Context Hub Configuration Guide* provides additional information.

Pivot to Investigate > Navigate (Events View)

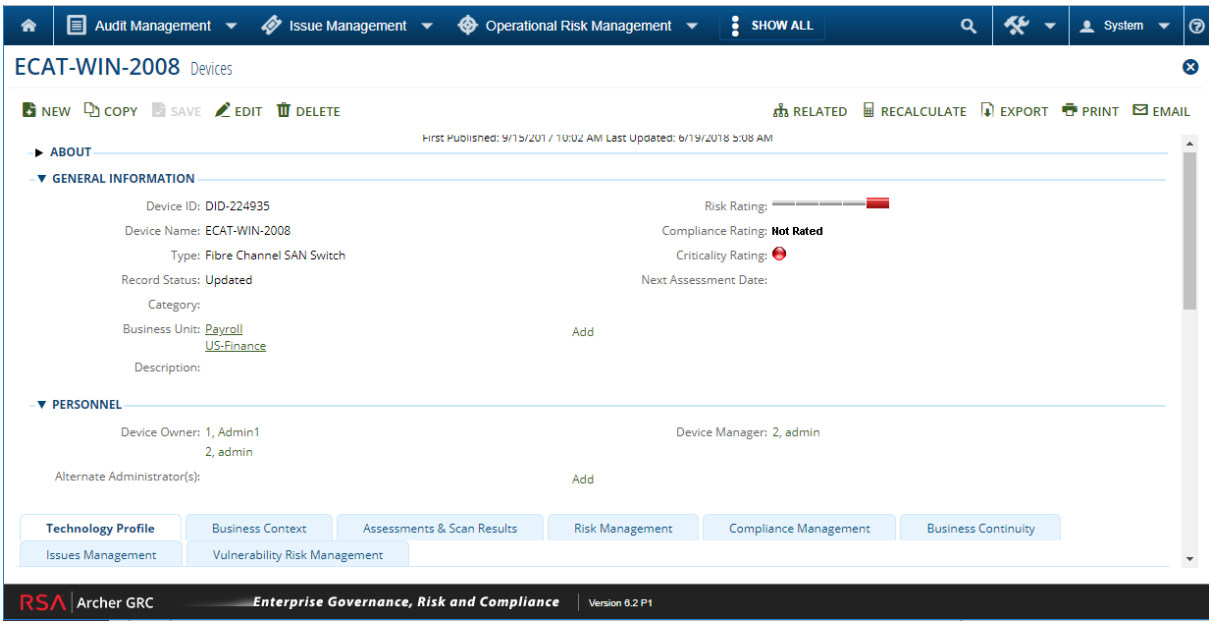
For a more thorough investigation of an entity, you can open the the Navigate view.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over any underlined entity.
2. In the **ACTIONS** section of the tooltip, select **Pivot to Investigate > Navigate**.
   The Navigate view opens, enabling you to perform a deeper dive investigation. For more information, see Begin an Investigation in the Navigate or Legacy Events View.

Pivot to Archer (Events View)

For viewing more details about the device in Archer Cyber Incident & Breach Response, you can pivot to the device details page. This information is displayed only for IP address, host, and Mac address.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over any underlined entity (IP address, host, and Mac address).
2. In the **ACTIONS** section of the context tooltip, select **Pivot to Archer**.
3. The device details page in **Archer Cyber Incident & Breach Response** opens if you are logged in to the application, otherwise the login screen is displayed.



Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer Datasource is not responding. Check that the Archer configuration is enabled and configured properly.

For more information, see the *Archer Integration Guide*.

Pivot to NetWitness Endpoint Thick Client (Events View)

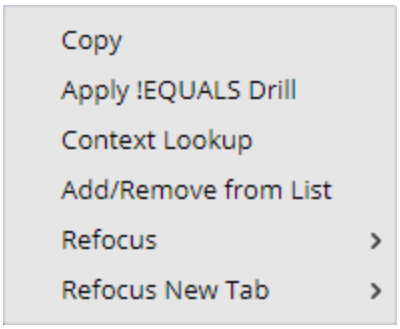
If you have the NetWitness Endpoint thick client application installed, you can launch it through the context tooltip. From there, you can further investigate a suspicious IP address, Host, or MAC address.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over any underlined entity.
2. In the **ACTIONS** section of the tooltip, select **Pivot to Endpoint Thick Client**.
   The NetWitness Endpoint thick client application opens outside of your web browser.

For more information on the thick client, see the *NetWitness Endpoint User Guide*.

View the Context Lookup Panel in the Navigate View or Legacy Events View

1. Hover over different meta values to see the data sources for which data is available.
   A hover box displays a list of the data sources that have context data available for meta value. These are the possible data sources: NetWitness Endpoint, Incidents, Alerts, Hosts, Files, Feeds, and Live Connect.
2. Right-click a meta value, and click **Context Lookup** in the drop-down menu to open the Context Lookup panel.



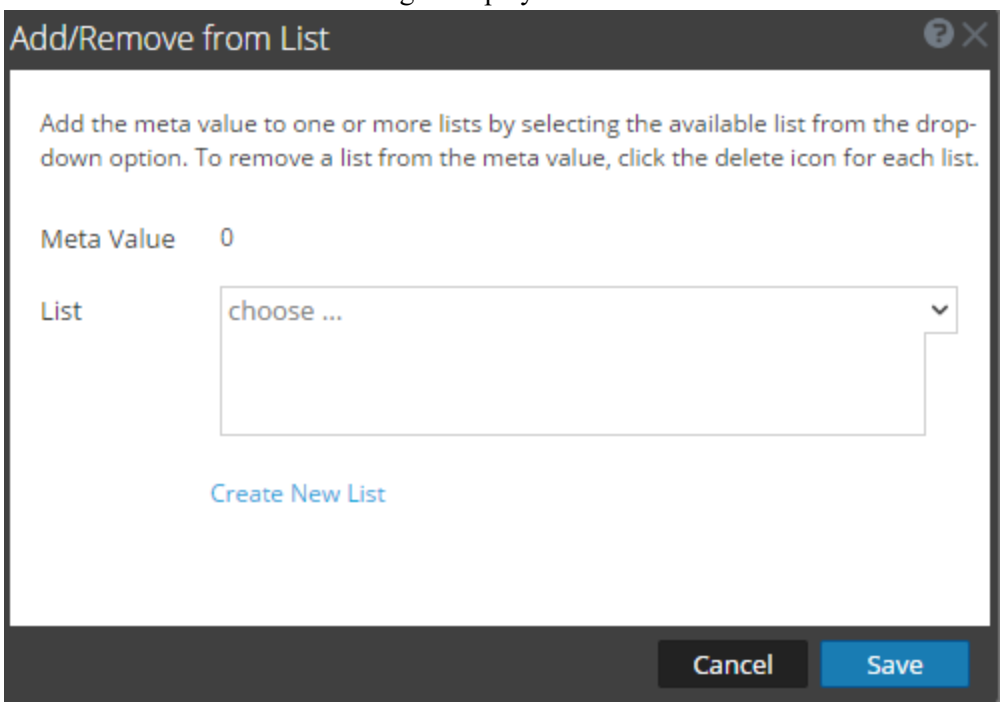
The Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available.

3. To perform actions from the Context Lookup panel, right-click an entity such as IP address.
   The following options are available: Open Link in New Tab, Query in Investigate, Copy Link, Paste, Google Lookup, Virus Total Lookup, and Query in Endpoint.
4. To close the Context Lookup panel, click ☒ in the panel.

Add Meta Values to an Existing List (Navigate and Legacy Events Views)

To add a meta value to an existing list in Context Hub:

1. While investigating a service in the **Navigate** view or the **Legacy Events** view, right-click a meta value (for example, values under Source IP, Destination IP, or Username) and select **Add/Remove from List** in the context menu.
   The Add/Remove from List dialog is displayed.



2. In the **List** field, select one or more lists from the drop-down option to which the meta value must be added.
3. Click **Save**.
   The meta value is added to the selected lists.

Remove a Meta Value from a Context Hub List (Navigate and Legacy Events Views)

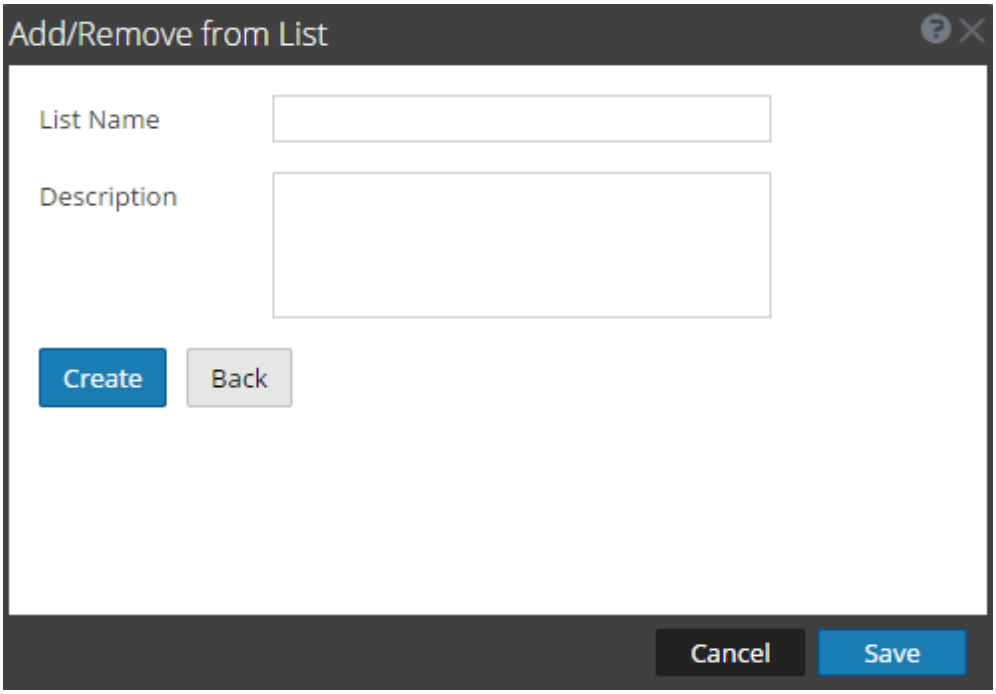
To remove a meta value from list:

1. In the **Add/Remove from List** dialog, in the **List** field, view the lists which include the meta value.
2. Click the delete icon (x) for each list that should not include the meta value.
3. Click **Save**.
   The meta value is removed from the deleted list.

Create a New List (Navigate and Legacy Events Views)

To create a Context Hub list in Investigate:

1. In the **Add/Remove from List** dialog, click **Create New List**.



2. In the **List Name** field, enter an unique name for the list.
3. In the **Description** field, enter the description of the list.
4. Click **Create** to create the list.

5. Click **Save** to add the meta value to the created list.
   These lists are considered as data sources for retrieving context information.

Launch a Lookup of a Meta Key

When you have found data of interest in the Navigate view, the Events view, or the Legacy Events view, you can do internal lookups to NetWitness Endpoint and RSA Live, as well as external lookups of meta values in community resources such as SANS IP History and ThreatExpert Search.

Analysts can use the external lookups to save time during investigations. The external lookups are available by right-clicking one of the these meta keys: IP address (`ip-src, ip-dst, ipv6-src, ipv6-dst, orig_ip`), `host (alias-host, domain.dst)`, `client`, and `file-hash`.

For all `ip` and `host` meta keys, the following lookups are built in to NetWitness:

- Google Malware: Opens a Google Malware search in a new tab.
- SANS IP History: Opens a SANS IP History search in a new tab.
- McAfee SiteAdvisor: Opens a McAfee SiteAdvisor search in a new tab.
- Endpoint Thick Client Lookup: Opens a search in the NetWitness Endpoint Thick Client in a new tab.
- BFK Passive DNS Collection: Opens a BFK Passive DNS collection search in a new tab.
- CentralOps Whois for IPs and Hostnames: Opens a CentralOps Whois search for IPs and hostnames in a new tab.
- Malwaredomainlist.com Search: Opens a Malwaredomainlist.com search in a new tab
- Robtex IP Search: Opens a RobtexIP search in a new tab.
- ThreatExpert Search: Opens a ThreatExpert search in a new tab
- IPVoidSearch: Opens a UrlVoid Search in a new tab n a new tab

For the `file-hash` and `alias-host` meta keys, the Google lookup opens a Google search in a new tab.

For the `client` meta key, the NetWitness Endpoint Lookup option opens an Endpoint Thick Client in a new tab if the client is installed on the same system on which the browser is being used.

Administrators can add additional external lookups and other custom actions as described in "Add Custom Context Menu Actions" in the *System Configuration Guide*.

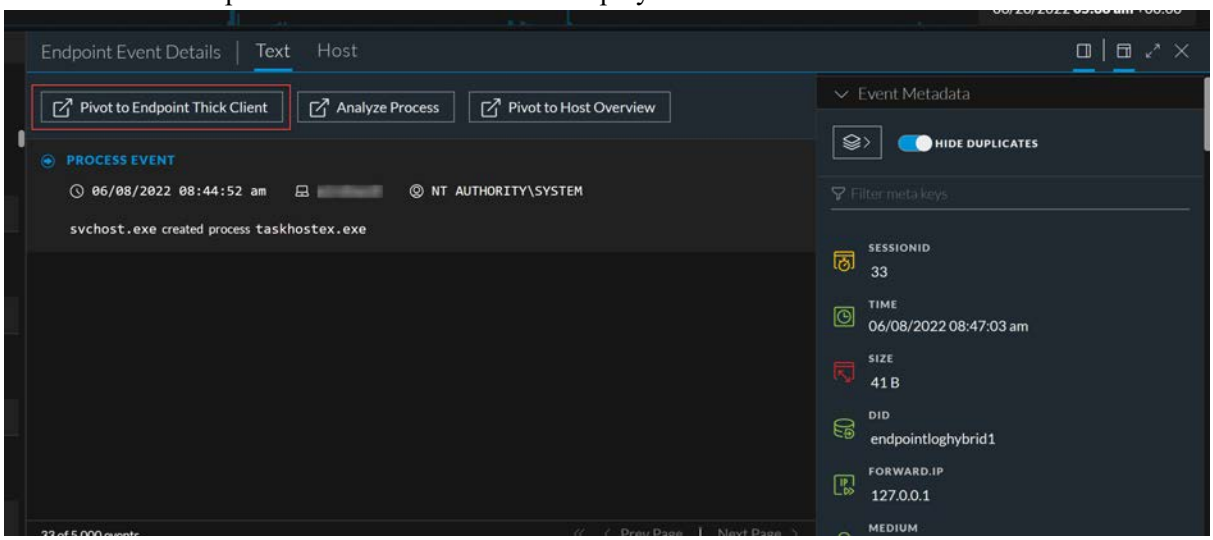
Launch an Endpoint Thick Client Lookup in the Events View

When viewing an endpoint event in the Text panel, you can pivot to analyze the same event in NetWitness Endpoint.

Note: Version 4.4.0.x of the NetWitness Endpoint (NWE) thick client must be installed on the same server, the NWE meta keys must exist in the `table-map.xml` file on the Log Decoder, and the NWE meta keys must exist in the `index-concentrator-custom.xml` file. The NWE thick client is a Windows only application. Complete setup instructions are provided in the *NetWitness Endpoint User Guide* for Version 4.4.

To open an event in NetWitness Endpoint:

1. Starting from the Navigate view:
   a. In the **Query** drop-down, select **Advanced**, and enter one of the following queries: `nwe.callback_id exists` or `device.type='nwendpoint'`
   Endpoint data is displayed in the Values panel.
   b. Right-click an event, and select **Events** in the menu.
2. (Version 11.1 and later) Go to **INVESTIGATE > Events**. In the **Query** drop-down, select **Advanced**, and enter one of the following queries: `nwe.callback_id exists` or `device.type='nwendpoint'`
   Endpoint data is displayed in the Events panel.
3. Select an event.
   The Events view opens with the selected event displayed in the Text view.



4. In the Event Header click **Pivot to Endpoint.**
   A new browser tab with the url `ecatui://<id>` opens and the NWE Thick Client is launched. If the NetWitness Endpoint Thick Client is not installed, no data is displayed and the following message is displayed: `Applicable for hosts with 4.x Endpoint agents installed, please install the NetWitness Endpoint Thick Client.`

Launch an Endpoint Thick Client Lookup in the Navigate View

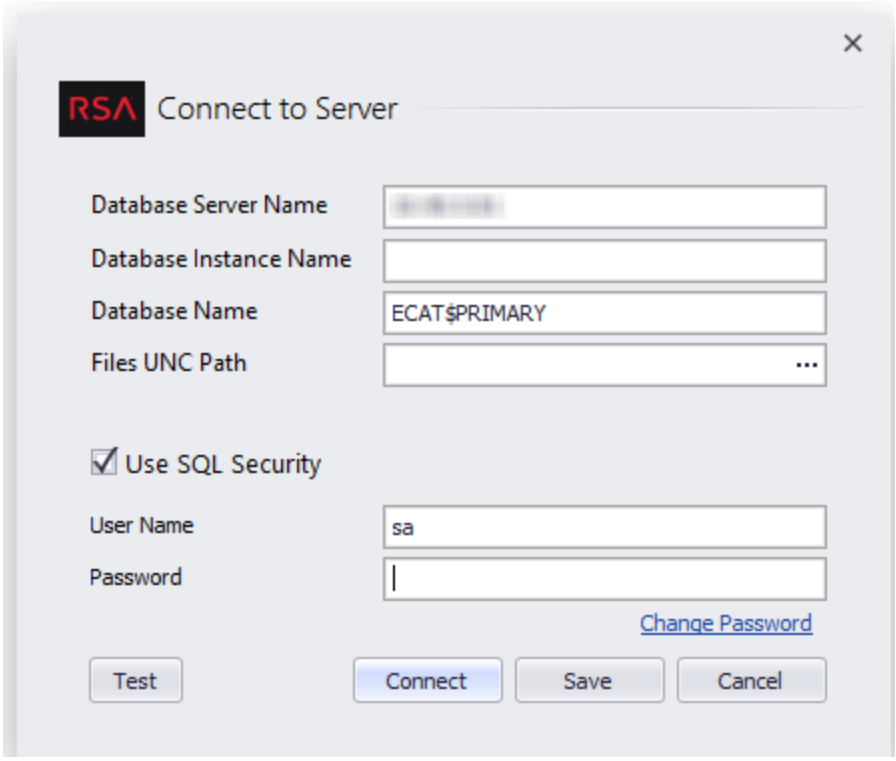
To launch an Endpoint Thick Client lookup of data from the Navigate view:

1. Right-click a meta value for one of the following meta keys: `ip-src`, `ip-dst`, `ipv6-src`, `ipv6-dst`, `orig_ip`, `alias-host`, `domain.dst`, or `client`.

2. Select **External Lookup** in the context menu.
   A submenu of external lookup options is displayed.

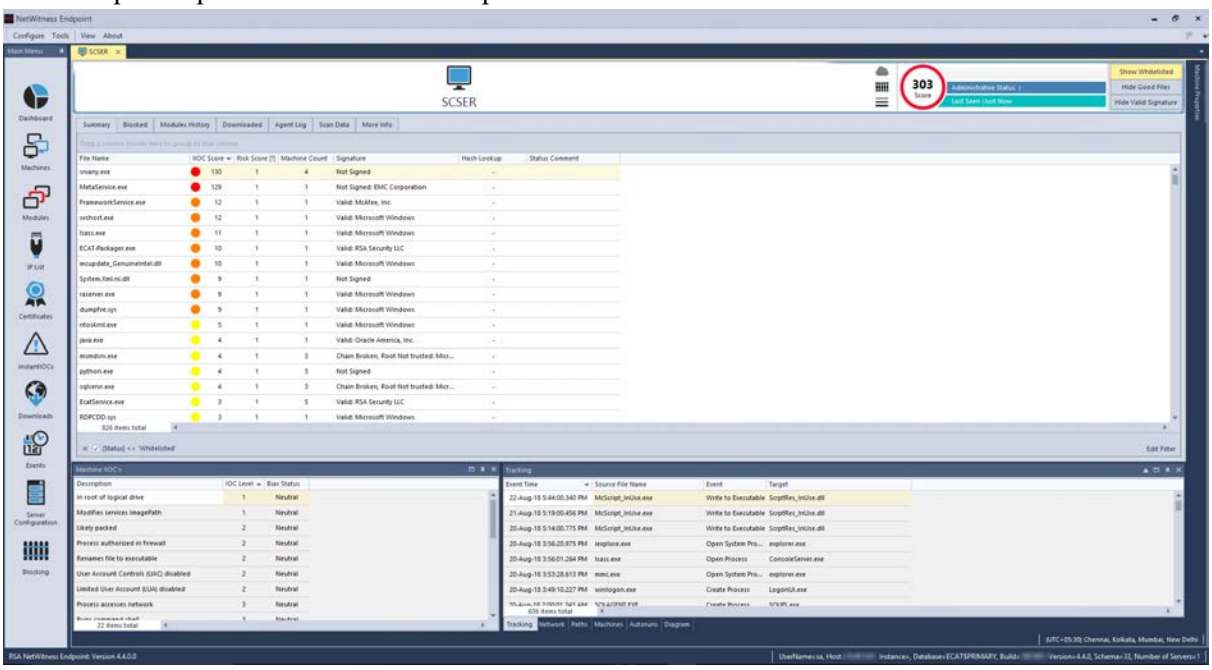


3. Select **Endpoint Thick Client Lookup**.
   The Connect to Server dialog is displayed.



4. Enter the user name and password required to log in to the Endpoint Thick Client, and click **Connect.**

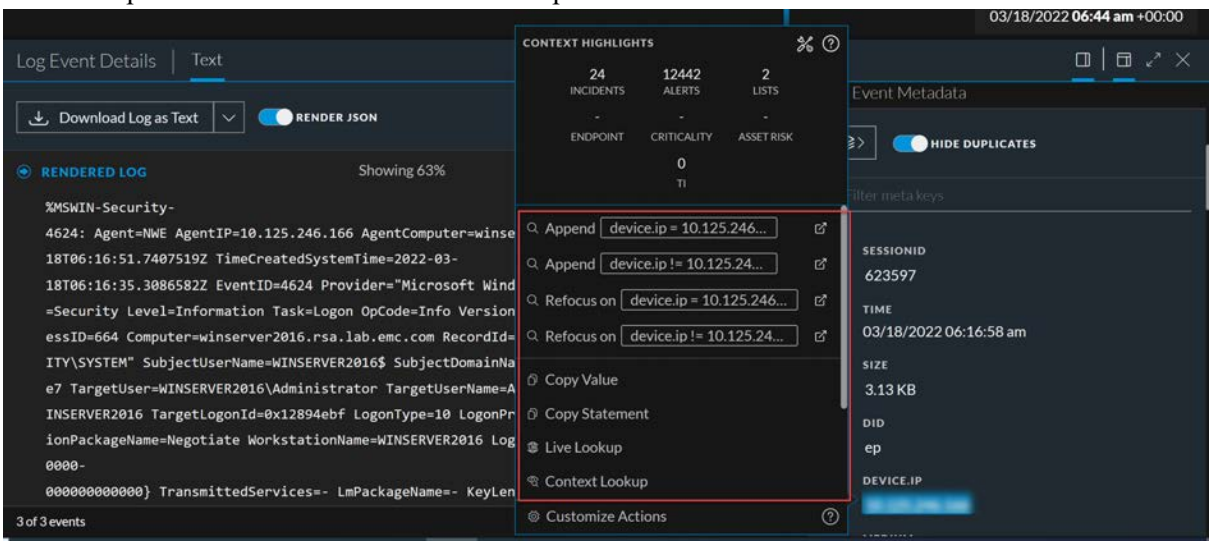
The drill point opens in NetWitness Endpoint.



Perform Lookups of Meta Values in Events

In the Events view, you can further investigate meta values in an event by lift or right-clicking certain meta values and using the options in a drop-down menu. To perform internal and external lookups:

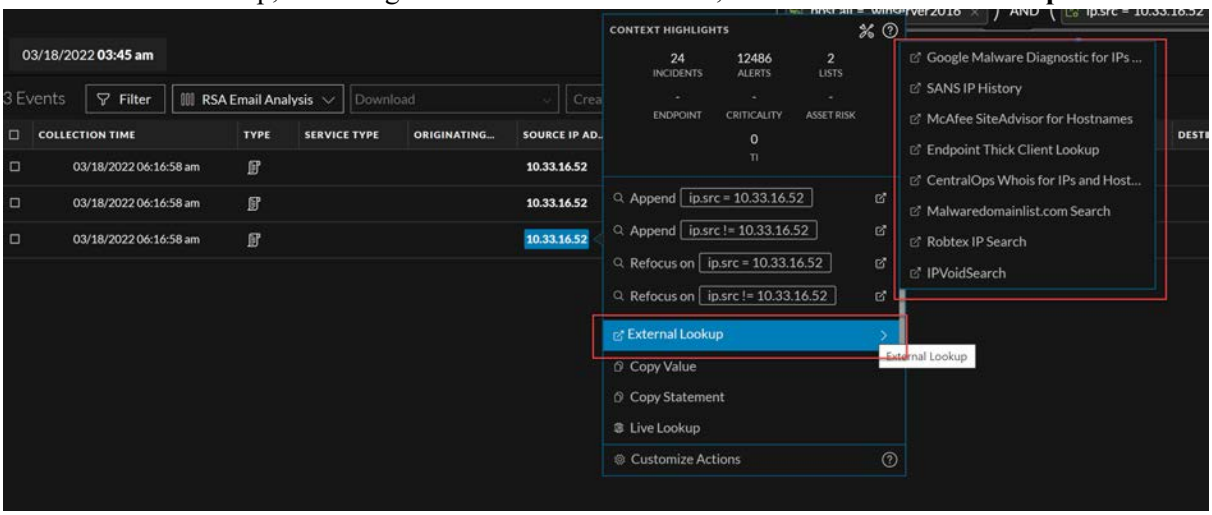
1. In the Events view, left or right-click a meta value in the Events List, the Event Meta panel, or the Overview panel. Some meta values have a drop-down menu.



2. Select one of the following internal lookups:

- **Copy Value**: Copies the meta value to the clipboard.
- **Refocus Investigation in New tab**: Launches the another investigation in a new tab with the focus on the selected meta value.
- **Apply Drill in New Tab**: Applies the drill and launches it in a new tab to drill the data in Navigate view.
- **Apply !EQUALS Drill in New Tab**: Applies (!EQUALS) to the meta and launches a new tab, effectively excluding the meta value from the results.
- **Hosts Lookup**: Looks up the value in the Investigate > Hosts view.
- **Endpoint Thick Client Lookup**: Analyzes the meta value in the Endpoint Thick Client (for clients which have Endpoint Agent).
- **Live Lookup**: Looks up a meta value on Live for further analysis.

3. For an external lookup, left or right-click on a selected meta, and click **External Lookup**.



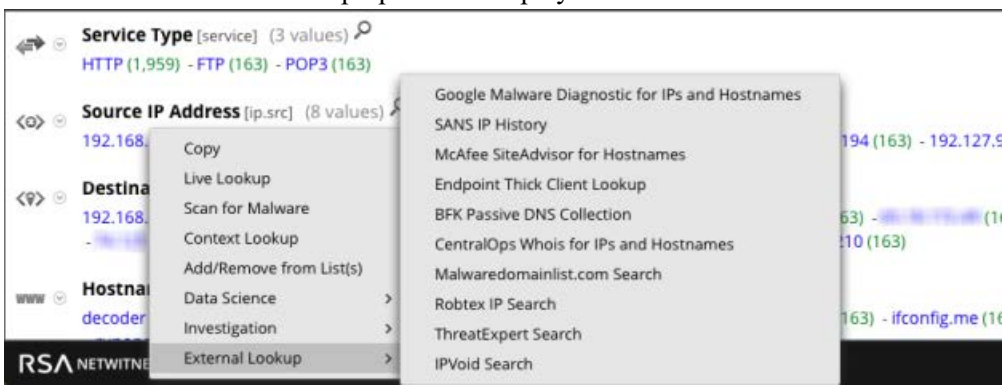
4. In the submenu select one of the available external lookups:

- **Google**: Looks up a meta value on Google.com
- **SANS IP History**: Looks up a meta value on SANS IP History, domain = http://isc.sans.org/ipinfo.html?ip=ipaddress
- **CentralOps Whois for IPs and Hostnames**: Looks up a meta value on CentralOps Whois for IPs and Hostnames, domain = http://centralops.net/co/DomainDossier.aspx?addr=domain&dom_whois=true&dom_dns=true&net_whois=true
- **Robtex IP Search**: Looks up a meta value on Robtext IP Search, domain = https://www.robtex.com/cidr/domain.ipaddress
- **IPVoid**: Looks up a meta value on IPVoid, domain = http://www.ipvoid.com/scan/domain/
- **URLVoid**: Looks up a meta value on URLVoid, domain = http://www.urlvoid.com/scan/ipaddress/
- **ThreatExpert Search**: Looks up an IP meta value on ThreatExpert Search, domain = http://www.threatexpert.com/reports.aspx?find=IP address

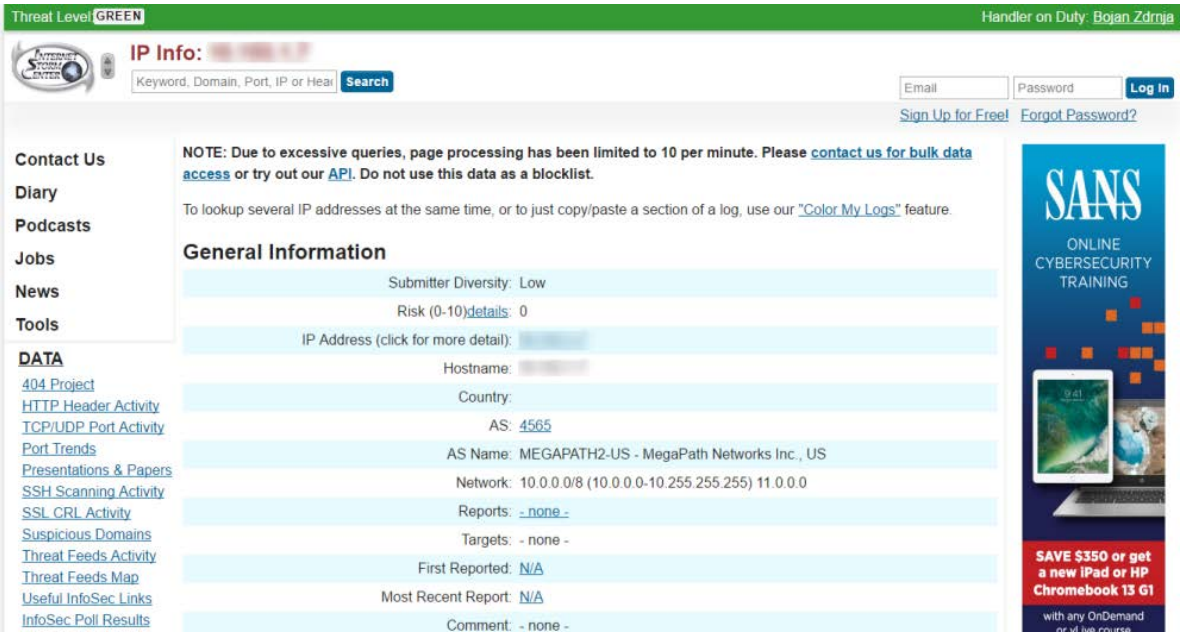
Launch Other External Lookups from the Navigate View

To launch an external lookup of data from the Navigate view (other than NetWitness Endpoint Thick Client Lookup):

1. Right-click a meta value for one of the following meta keys: `ip-src`, `ip-dst`, `ipv6-src`, `ipv6-dst`, `orig_ip`, `alias-host`, `domain.dst`, or `client`.
2. Select **External Lookup** in the context menu.
   A submenu of external lookup options is displayed.



3. Select one of the lookup options.
   The selected meta value opens in the selected lookup, for example, if you selected SANS IP History, the drill point information is displayed in SANS Internet Storm Center.



Launch a Malware Analysis Scan from the Navigate View

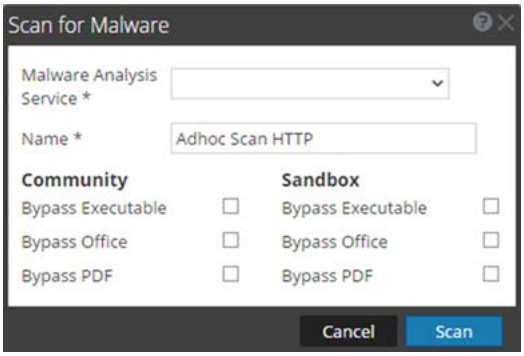
From within Investigate, analysts can launch an on-demand Malware Analysis scan by selecting a service and meta value, and choosing an option from the context menu. When polling is complete, the scanned data is available for malware analysis.

To launch a Malware Analysis scan of data from the INVESTIGATE > Navigate view:

1. Right-click a meta value (for example, OTHER, DNS, or FTP) and select **Scan for Malware** in the context menu.

   The Scan for Malware dialog is displayed with a suggested name for the on-demand scan and no service selected.

2. In the Scan for Malware dialog, select a service to perform the scan, edit the name, and select the types of files to bypass under community and sandbox.

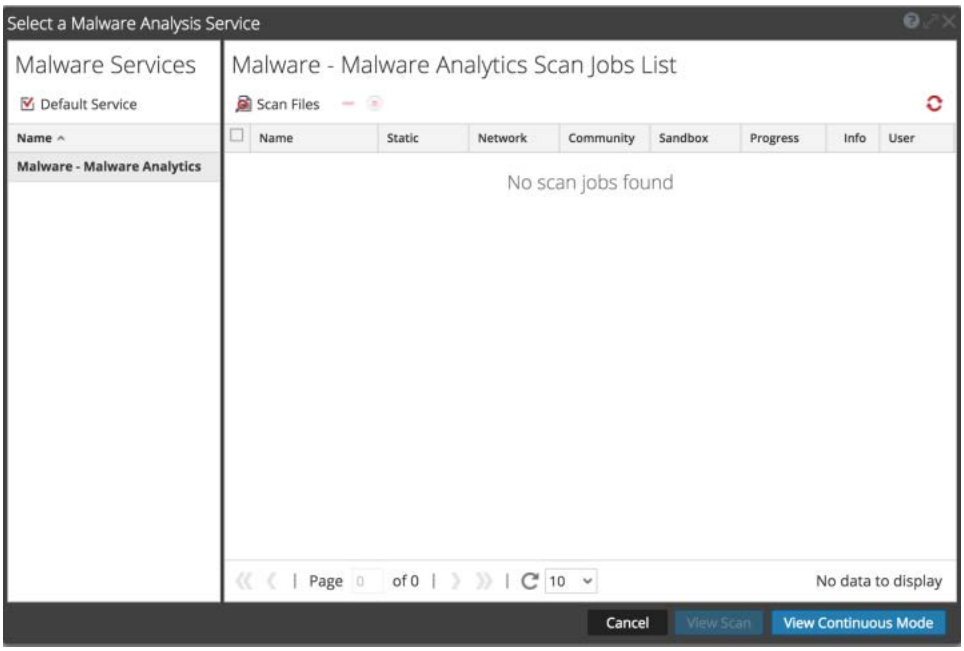


3. Click **Scan**.

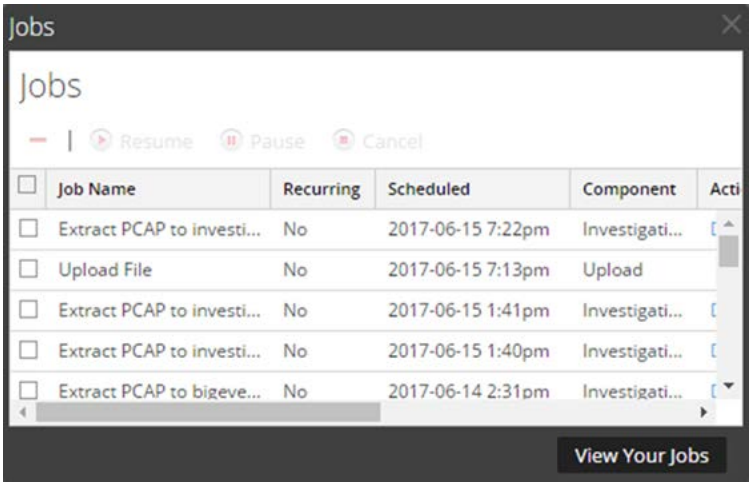
   The scan request is added to the Scan Jobs List dashlet and the Jobs Tray. The bypass settings in this dialog override the default settings in the basic Malware Analysis configuration settings.

4. To view the jobs, do one of the following:

   a. Go to the Scan Jobs List in the Malware Analysis view or in the Unified dashboard. Double-click a scan to view the scan.



b. To view the job in the Jobs tray, click  in the NetWitness toolbar. When the job is complete, scroll to the left and click **View**.



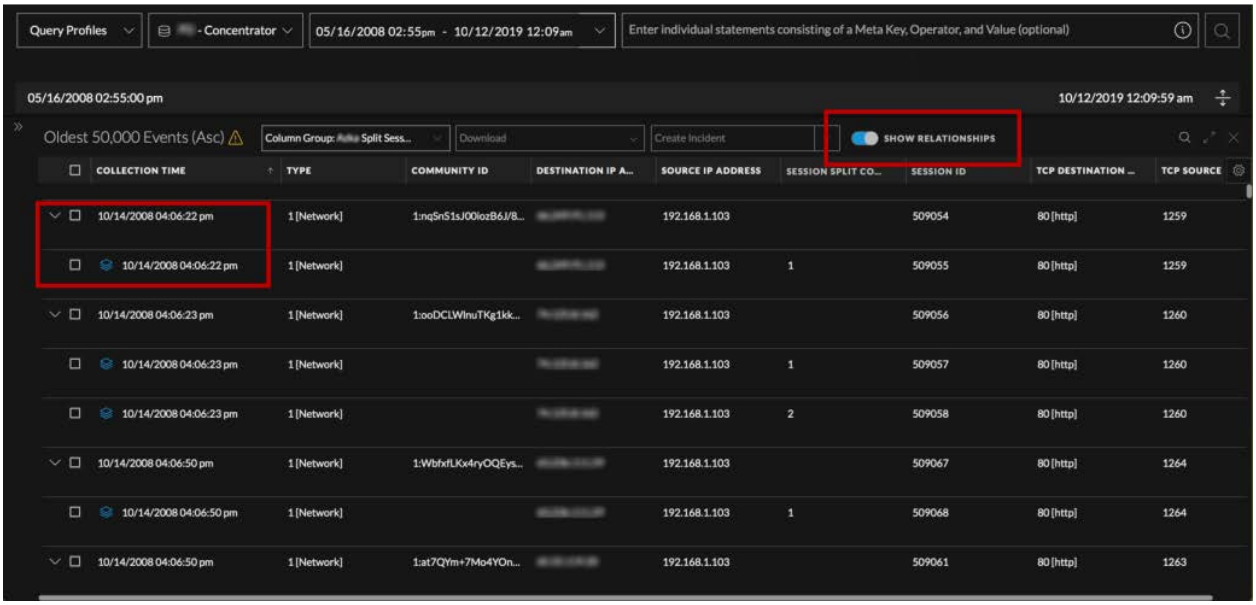
The Malware Summary of Events for the selected scan is displayed. The scan is also added to the list of available scans in the dialog for selecting scans in the Investigation > Malware tab.

Group Events from Split and Related Sessions in the Events and Legacy Events Views

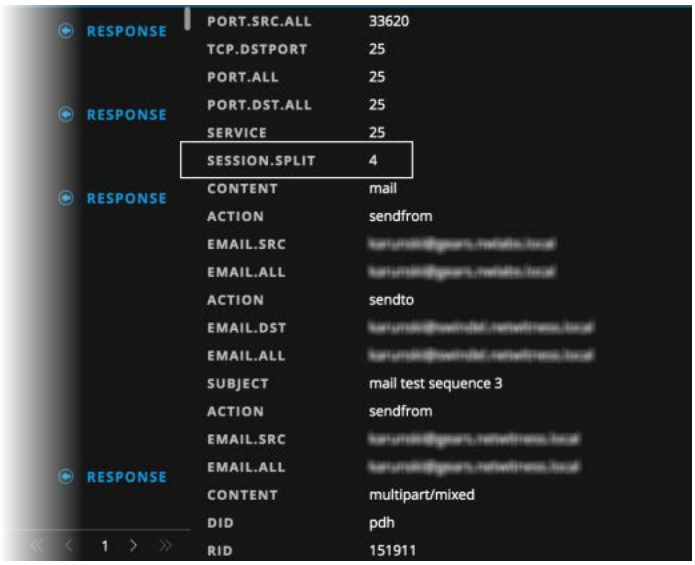
In the Events view listing of events, events from split and related sessions are listed in the order they were parsed; for this reason, they are not always listed together in the Events list. In Version 11.4.1 and later, the Group Events option allows you to change the order in which the events are listed so that you can more easily detect relationships in the captured data. When events are grouped, we refer to the first event as the leading event.

The user interface is designed to help you identify grouped events. Solid row lines delineate different groups of related events while dotted lines depict the events that are part of the same related group. In a group of events, the leading event is first and subsequent events are nested under the leading event with indentation and relationship icons for subsequent events. The numbers next to the relationship icon distinguish the session split count.

If the leading event is not included in the current data set, the subsequent events are still grouped together under the first subsequent event. Only the leading event or the first subsequent event, in the case of no leading event, is sorted; indented events are not sorted. Hovering over the subsequent event marker () displays a tool tip explaining the relationship. The following figure is an example of related events as they are displayed in the Events list.



If events are related based on session fragments, when you select a subsequent event and open the reconstruction, you can see the `session.split` meta key in the Event Meta panel.



Split Network Session

If you see a tool tip like this, the listed events are part of a split network session:
`The event is part of a split session (session.split: #) matching these parameters: ip.src=127.0.0.1 AND ip.dst=127.0.0.1 AND tcp.srcport=25 AND tcp.dstport=1234.`

The reason for the split is one of the following:

- The original event was split into sub-parts by creating additional events for each transaction inside the original event.
- The original session was split as it was ingested into the Network Decoder because the size was larger than the Assembler Maximum Size (default=32 MB).
- The original session was split as it was ingested into the Network Decoder because the time was longer than the Assembler Timeout Session (default=60 sec).

Session Size and Time Split

Network Decoders are configured with a default session size (`assembler.size.max`) and a session timeout (`assembler.timeout.session`). The configuration is described in "Configure Session Split Timeouts" in the *Decoder Configuration Guide*. When a session exceeds either limit, the Network Decoder splits the session, and all subsequent packets become part of a new session, fragmenting the actual network session across multiple Network Decoder sessions. To process session fragments with the context that they are fragments of the larger network session and to improve association of source and destination addresses, ports, and application protocols, Network Decoders parse contextual fragments and highlight session fragments.

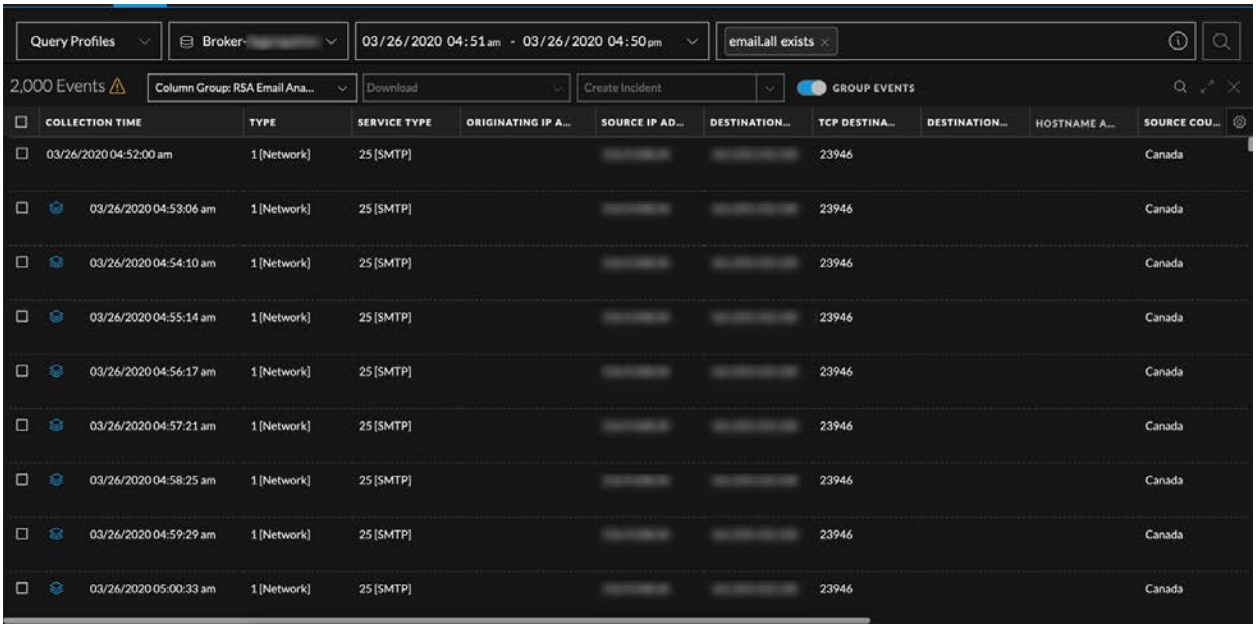
Note: In the Legacy Events view, you can find session fragments and export all packets viewed in the Events list to a single PCAP as described in Find and Combine Fragments in the Legacy Events View.

The Network Decoder completes session parsing before splitting the session based on the configured maximum session size (default = 32 MB) or the configured timeout (default = 60 seconds). When parsing is complete, the parsed results include the proper address directionality and application protocol, which are propagated to each subsequent session fragment to ensure consistency with the logical network session they represent.

Transaction Handling Split

Administrators can configure a Network Decoder to subdivide incoming sessions into smaller transaction sessions when using Lua parsers that are designed to create transactions. The configuration is described in "Configure Transaction Handling on a Decoder" in the *Decoder Configuration Guide*. The Decoder service configuration node has a parameter that controls the behavior of the Network Decoder when a parser defines a transaction within a network session: `/decoder/parsers/config/parser.transaction.mode`. If the mode is set to `split`, a large session with multiple application-level transactions is split when a parser generates an application-level transaction such as an email. An example of this is a large session containing multiple emails. For each email (transaction), a new session item (split session) is created, network meta items are copied into the new session, and meta items marked in the transaction are copied from the original session to the new session.

Transactions require a parser update to function, and initially they only support SMTP and HTTP pipelining use cases. This is an example of the reconstruction of an email that is separated based on the individual emails inside the original event. Each transaction highlights a single email and the metadata associated with the transaction is only related to that email. To provide this functionality, the original packets are stored on the Network Decoder as usual for a network event, but the new related transaction events are created on the Concentrator. As a result, analysts see visual queues in the user interface, and they also have the ability to query to find only specific emails or attributes of emails that were previously all bundled together. To eliminate the original event from the query results, the `session.split` meta key has been indexed. When there are transaction splits, the original event does not have that meta key associated with it while all the related transactional events do.



Session Fragments Highlighting

Both types of session fragments have an additional meta key: `session.split`. The earliest fragment is session 0 and sessions with a later time stamp are incrementally numbered 1, 2, 3, and so on. The `session.split` meta key indicates the number of preceding sessions fragments; however, it does not always indicate that there are subsequent session fragments, even with a value of 0. It is also possible for the first fragment of the session to not have `session.split` metadata if the session is parsed before exceeding the maximum session size.

Transaction splits start with a `session.split` value of 1. When viewing sessions, the `session.split` meta key clearly identifies sessions that are fragments in the Events view and the Legacy Events view (Events List view and the Events Detail view).

If this was a session size and timeout split, you can view the session fragments and determine the maximum session size or session timeout necessary for parsing to combine the split sessions into one again. For example, if you have four fragments at 32 MB, you need to configure your test Decoder (usually a virtual machine set up separately from the production service) with a maximum session size greater than 128 MB. The steps are the same to find all fragments based on a session timeout.

Related Network Session

If you see a tool tip like this, separate events processed by the Network Decoder share a set of four values that identify the IP source, IP destination, source port, and destination port:

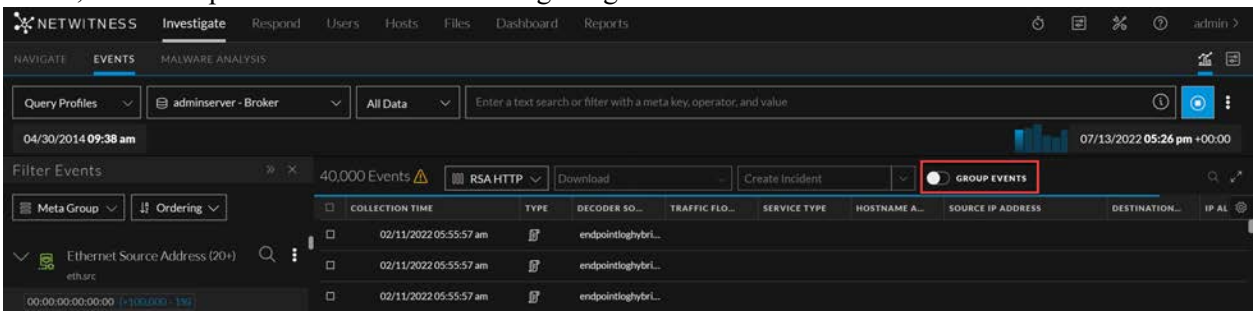
```
 The event is related to a previous session matching these parameters:
ip.src=127.0.0.1 AND ip.dst=127.0.0.1 AND tcp.srcport=25 AND
tcp.dstport=1234"Second category: Related Network Session
```

In this case, the Network Decoder has not inserted a split, and there is no `session.split` metadata associated to any of the events. The reason these events grouped together is to highlight the events that are worthy of scrutiny based on a pattern. Each event has the same source IP address, the same destination IP address, the same source port, and the same destination port. The grouping of related events does not occur if any of the four meta keys is obfuscated for data privacy reasons.

These are the combinations of meta keys that must match to categorize an event as a Related Network Session:

- `ip.dst`, `ip.src`, `tcp.dstport`, `tcp.srcport`
- `ip.dst`, `ip.src`, `upd.dstport`, `udp.srcport`
- `ipv6.dst`, `ipv6.src`, `tcp.dstport`, `tcp.srcport`
- `ipv6.dst`, `ipv6.src`, `upd.dstport`, `udp.srcport`

In Version 11.6, the group events option that retrieves events from logs, network and endpoints is disabled when data privacy is enabled. For example, if the `'ip.src'` is blacklisted or restricted by the admin, then the option shown in the following image is disabled.



This option is disabled if one or more of the following 12 metas are restricted:

`'sessionid'`

`'nwe.callback_id'`

`'medium'`

`'session.split'`

`'ip.dst'`

`'ip.src'`

`'ipv6.src'`

`'ipv6.dst'`

`'tcp.dstport'`

`'tcp.srcport'`

`'udp.dstport',`

`'udp.srcport'`

Use Cases for Viewing Events from Split and Related Network Sessions

These are examples of practical use cases for viewing events from split sessions:

- A Network Decoder that is inline with a proxy server receives a lot of email connections that get bundled into a single session based on event time as NetWitness sees them. The single session has multiple meta values for `subject`, `email.src`, `email.dst`, and other meta keys relevant to email, which are difficult to map together correctly. Seeing the session organized as leading and subsequent events, gives the analyst a clear idea of the details of each email.
- An analyst is attempting to understand what IP address out of all the metadata associated with a session resulted in the generated metadata or an alert, but the IP address is not in the output. For example, a feed that is parsing for an indicator of compromise may have many triggers in a session that has many IP addresses. By viewing the complete event organized as leading and subsequent events, the analyst can understand which IP triggered the alert.
- An analyst wants to know which file was deleted from which directory versus which file was read out in which directory, but the session encompasses multiple files and directories. For example, an HTTP

connection with commands: `directory /keep/`, `directory /temp/`, `filename foo.txt`, `filename me.doc`, `action delete`, and `action read`. Viewing the leading and subsequent events tells the analyst that `/temp/me.doc` was deleted and `/keep/foo.txt` was read. Now the analyst or analytics can form an opinion about how bad these actions really were.

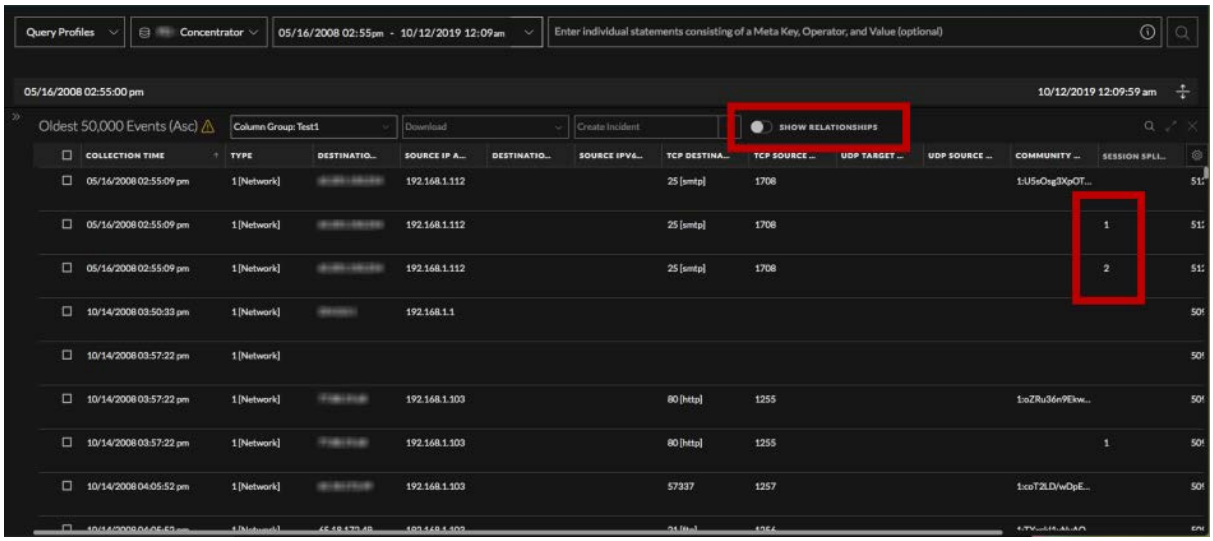
- An analyst is attempting to retrieve a large file related to an event that triggered a suspicous alert. However, the file that was transferred was so large that the network decoder split it into 100 separate sessions. When viewing this group related split sessions, the analyst can download a PCAP of the sessions, then extract the original file by either running it through a Decoder with larger assembler settings or a third-party tool.

Show and Hide Relationships in the Events List

For both types of related events, you can see the relationship of events in the Events view Events list. When the Events list is first displayed, you can tell if set of results includes related events by looking at the Group Events switch at the top of the Events list. If the results do not include related events, the switch is dimmed as shown in the following figure.

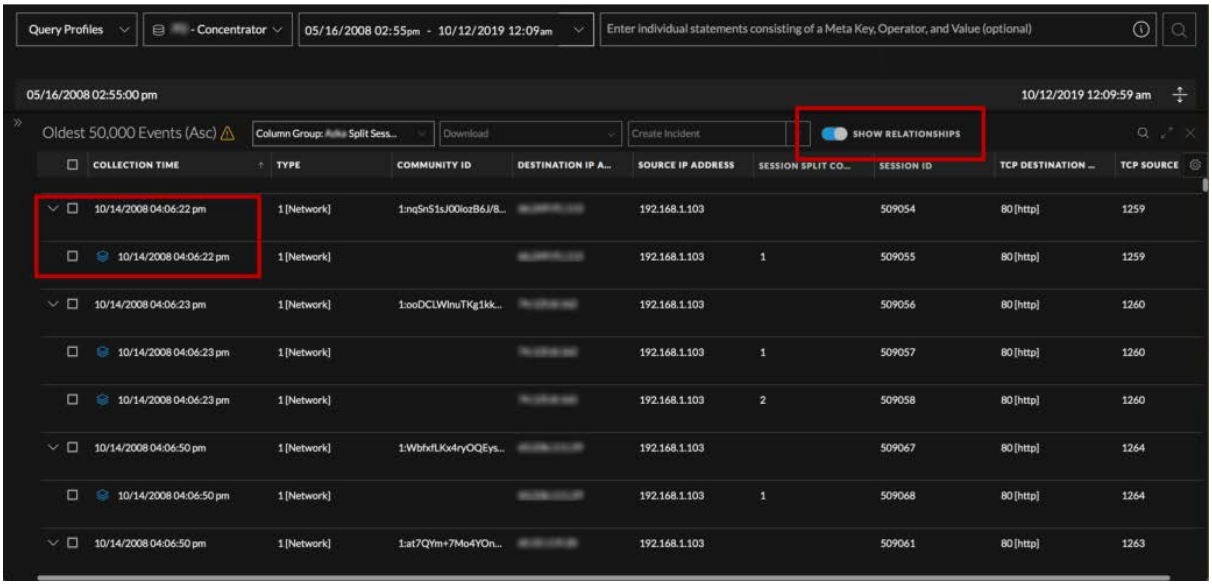
To look for related events in the Events list:

1. Go to **Investigate > Events** and submit a query.
   If the results include related events, the Group Events switch is active, but not enabled. The figure below illustrates a set of results that include split sessions, and the Group Events switch is disabled. The related events are not nested.



2. Click the **Group Events** switch.
   Related subsequent events are nested below the leading event. The subsequent events are indented

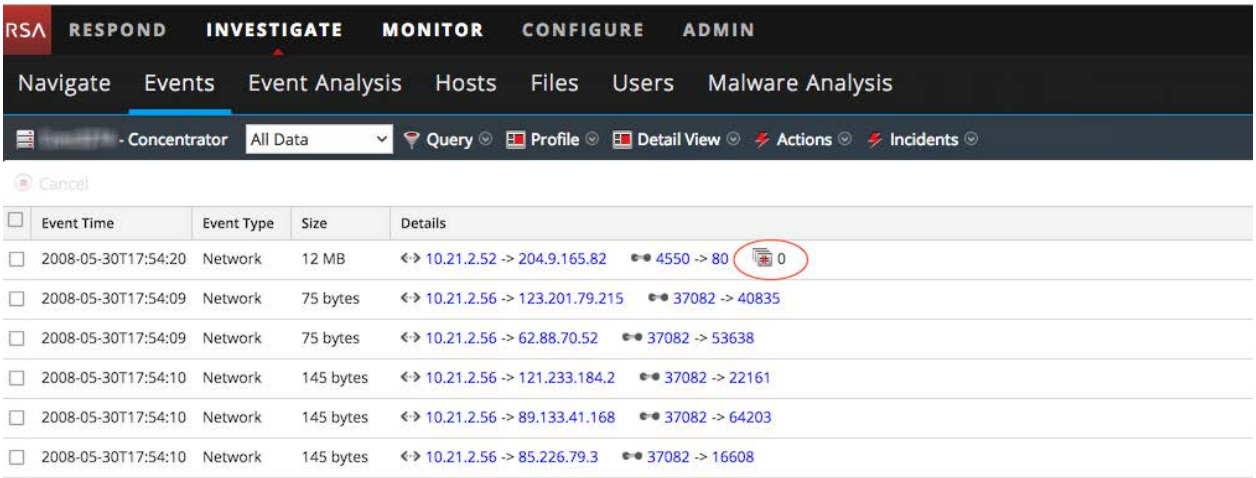
and marked by an icon. Clicking on the icon explains why the event is grouped.

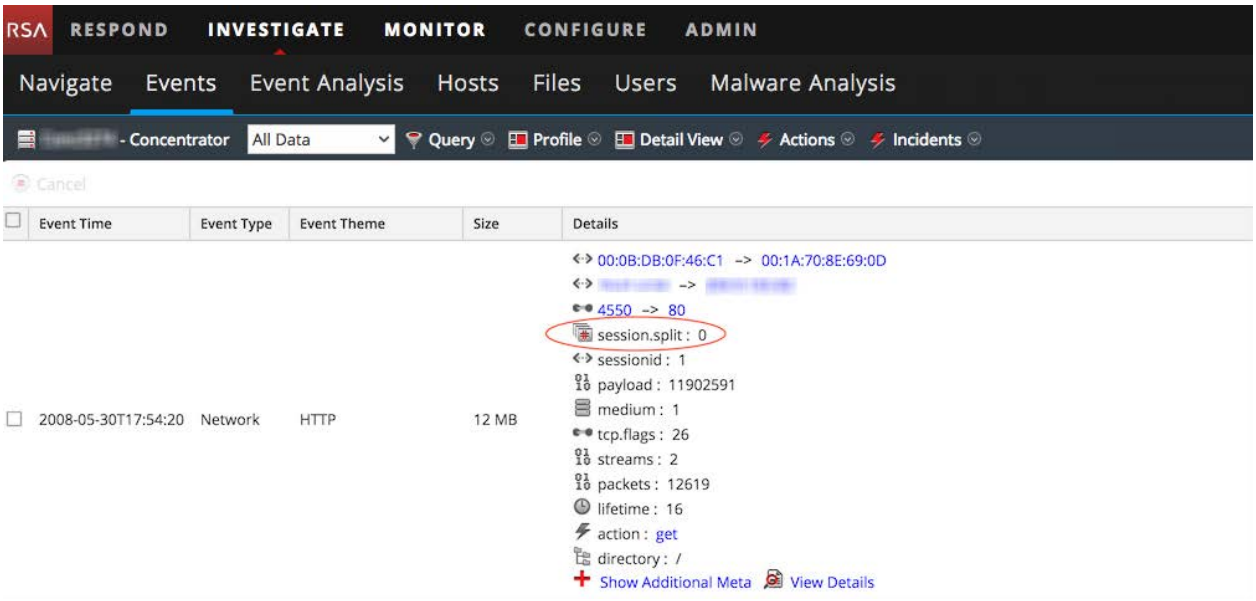


Session Fragments Highlighting (11.3 Events View)

The figures below show the Events List view and the Events Detail view with fragmented session information highlighted.

Note: A maximum session size of 12 MB was configured at the time the screen captures below were created.





Note: The `session.split` meta key must be indexed to help you find the split sessions in the results.

The `session.split` metadata is displayed immediately following the address and port metadata in the details view (Legacy Events view). These enhancements make it possible to quickly:

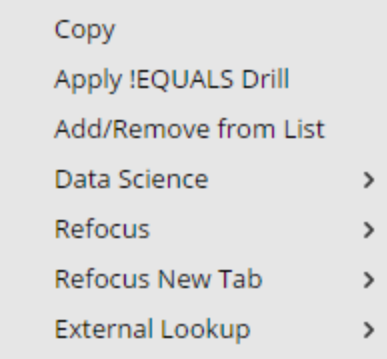
- Identify sessions that are fragments of a network sessions.
- View all of the session fragments of a network session given a single session fragment.
- Export the packets for the entire network session as a single PCAP file.

Find and Combine Fragments in the Legacy Events View

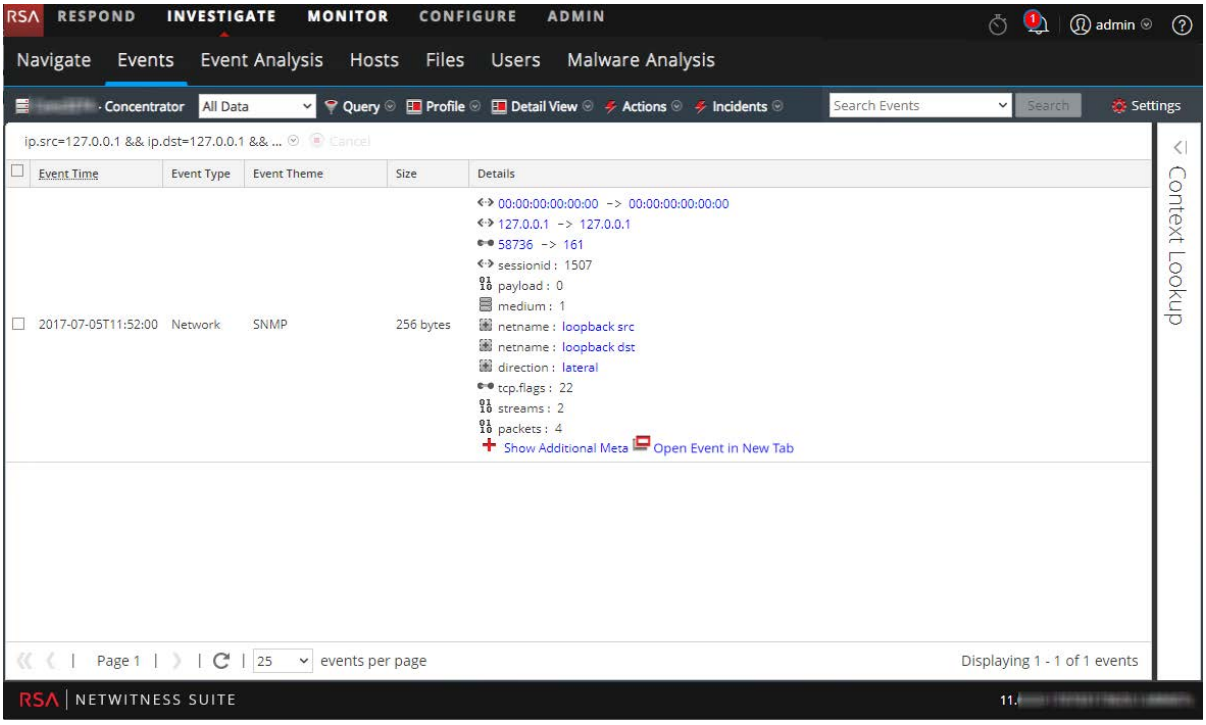
From within the Legacy Events view, you can find fragments of a session using the Refocus > Find Session Fragments context menu option. NetWitness composes a query using the source and destination addresses and ports of the selected session and displays all sessions that match that query within the current time window.

To find session fragments:

1. In the **Legacy Events** view, right-click any of the source and destination address and port values: `ip.src`, `ip.dst`, `ipv6.src`, `ipv6.dst`, `tcp.srcport`, `tcp.dstport`, `udp.srcport`, and `udp.dstport`) as well as `session.split` values.
   The context menu is displayed.



2. Select **Refocus > Find Session Fragments** or **Refocus New Tab > Find Session Fragments**. NetWitness repopulates the Events list with session fragments for a single session within the current time range. Depending on the option you selected, the refocus replaces the current view or opens in a new tab. (All data is used in these examples but is not recommended on production systems).



3. If necessary, adjust the time range to include any session fragments that may precede or follow the current time window. You can tell that the time range needs to be expanded if the fragments occur near the time boundary, especially if the first visible fragment does not have a split value of 0 (or none). Alternately, inspecting the packets of the last visible session may lead you to believe that the session continues. Here is an example:

   a. If you are looking at fragments that are obviously not the first fragment, for example, 1, 2, 3, and 4 in time range 10:30 to 10:35, there should be a fragment 0. You can increase the time range to start earlier (for this example, 10:25) to find the additional fragment.

b. If the session size of last fragment is close to maximum session size (12 MB in this example), look for additional fragments by increasing the time window to include a later time (for this example, 10:40).
When all of the session fragments of a network session are included within a single Events list, the list can span multiple pages.

4. (Optional) To export the packets for every session fragment to a single PCAP file, select **Actions > Export All PCAP**.
A message informs you that the PCAP is being downloaded. When download is complete, PCAP file includes the entire network session that was fragmented.

Visualize Metadata as Parallel Coordinates

Analysts can use the parallel coordinates visualization in the Navigate view to focus the investigation on combinations of meta keys and meta entities, and values that may indicate events are abnormal and worth investigation. The parallel coordinates chart is a way of visualizing the current drill point in Investigate to examine more than two meta keys simultaneously. Visualizing multiple meta keys simultaneously can help in identifying security issues associated with multivariate patterns and comparisons, such as when individual meta keys and values may not be of concern, but combining them together may bring an abnormal pattern or relationship to light. Meta groups (see Use Meta Groups to Focus on Relevant Meta Keys) can be used effectively to define a collection of meta keys that you want to visualize as parallel coordinates.

Best Practices for Effective Parallel Coordinates Charts

To create effective parallel coordinates charts, follow these recommendations:

- Use the RSA built-in meta groups that are included in a new installation.
- Start from a drill point rather than attempting to visualize all data.
- Limit the time range if necessary.
- Choose the smallest useful set of meta keys to display as axes.
- Specify the sequence of axes to highlight anomalies between the meta values as you follow a line across the chart.
- When you can identify a useful set of meta keys and sequence, create a custom meta group to use for future investigations. For example, you can create a custom meta group for Windows executable file types.
- Reuse and share custom meta groups by importing and exporting groups as `.jsn` files.
- It may be useful to create two versions of each custom meta group. One for analysis of meta values and one for creating a parallel coordinates chart focusing on a smaller subset of the same use case.

Note: When you import meta groups, an error message is displayed if any group is already present. To import a group that is a duplicate, you must first delete the existing group. If you want to delete a meta group, it cannot be in use by a profile.

To help build better parallel coordinates charts, several optimizations are included in NetWitness.

- Analysts can specify that only sessions in which all meta keys exist are rendered in the chart.
- The administrator can increase the number of meta values rendered in the Parallel Coordinates Settings in the ADMIN > System > Investigation panel > Navigate tab.

RSA Meta Groups for Parallel Coordinates Use Cases

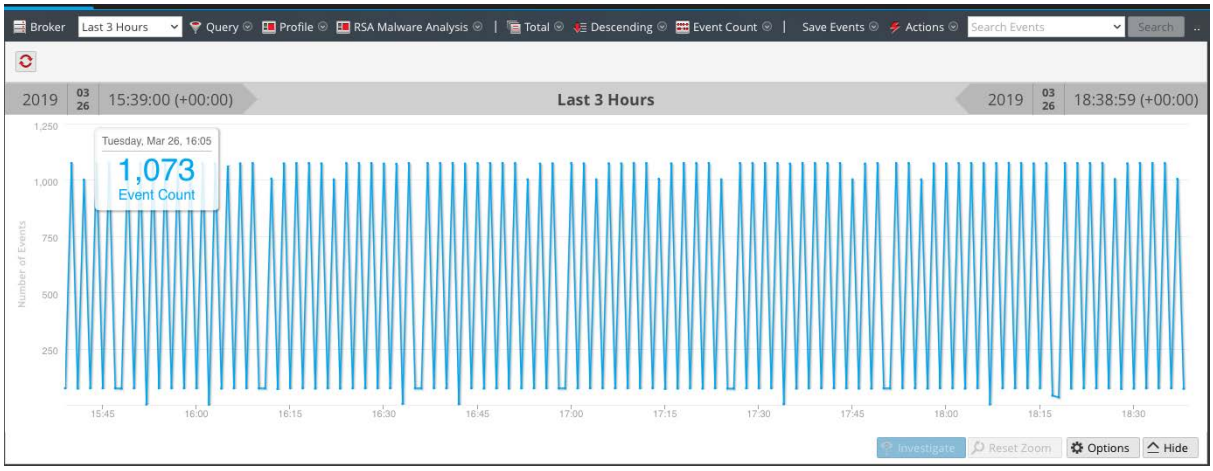
A set of predefined meta groups is included with NetWitness. If you want to get the latest version, you can import the meta groups file, `MetaGroups_ootb_w_query.jsn`, in the Manage Meta Groups dialog. Some of the targeted activities that lend themselves well to Parallel Coordinates visualizations are:

- Botnet Beaconing
- Covert Channels
- Email Analysis
- Encrypted Sessions
- Endpoint Analysis
- File Analysis
- Malware Analysis
- HTTP
- SSL/TLS
- SQL Injection Attacks
- Threat Analysis
- Web Analysis

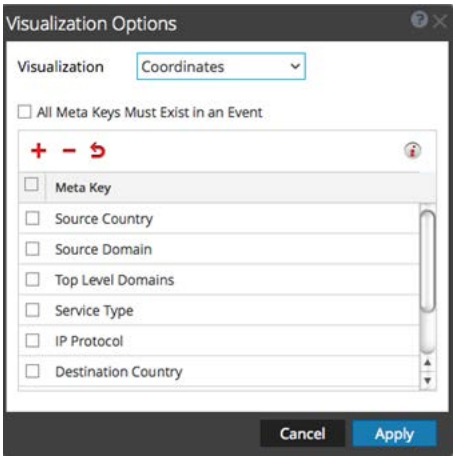
View a Parallel Coordinates Visualization

From an investigation in the Investigate > Navigate view:

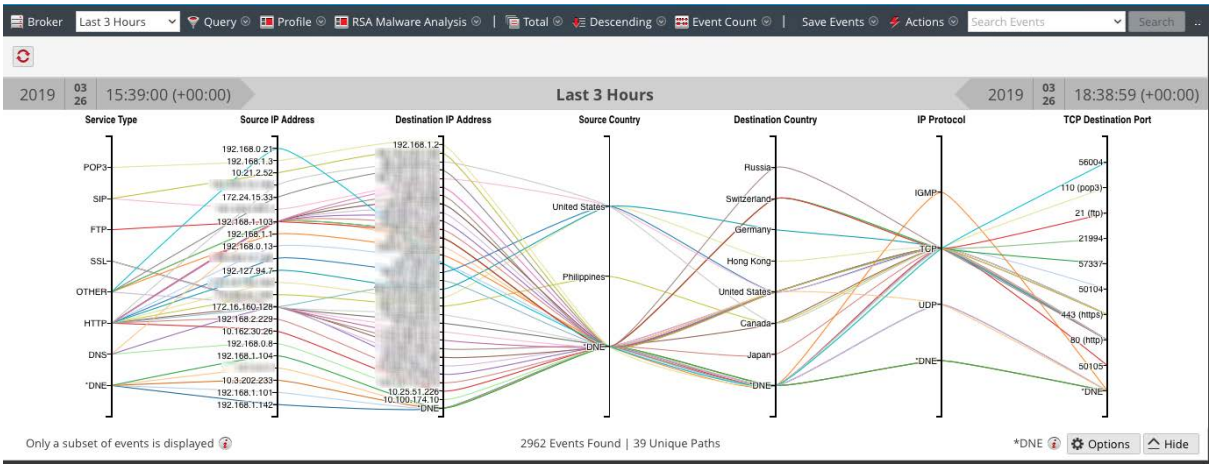
1. If the Visualization panel above the Values panel is closed, select **Visualization**.
2. In the toolbar, select **Meta** > **Use Meta Group > RSA Malware Analysis**.
3. A default visualization for the current drill point is displayed as a timeline.



4. In the **Visualization** panel, select **Options**.
   The Visualization Options dialog is displayed.
5. In the **Visualization** drop-down list, select **Coordinates** and click **Apply**.



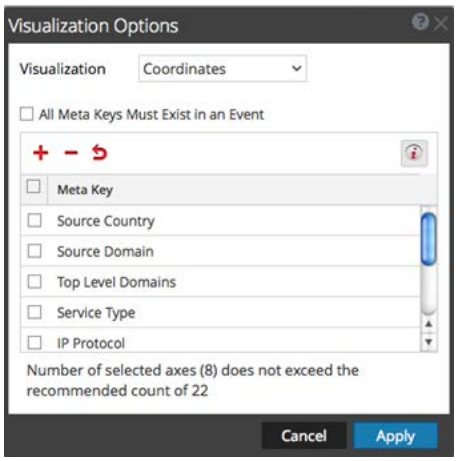
The visualization is loaded. In this example, 2962 events are found and 39 unique paths are visualized.



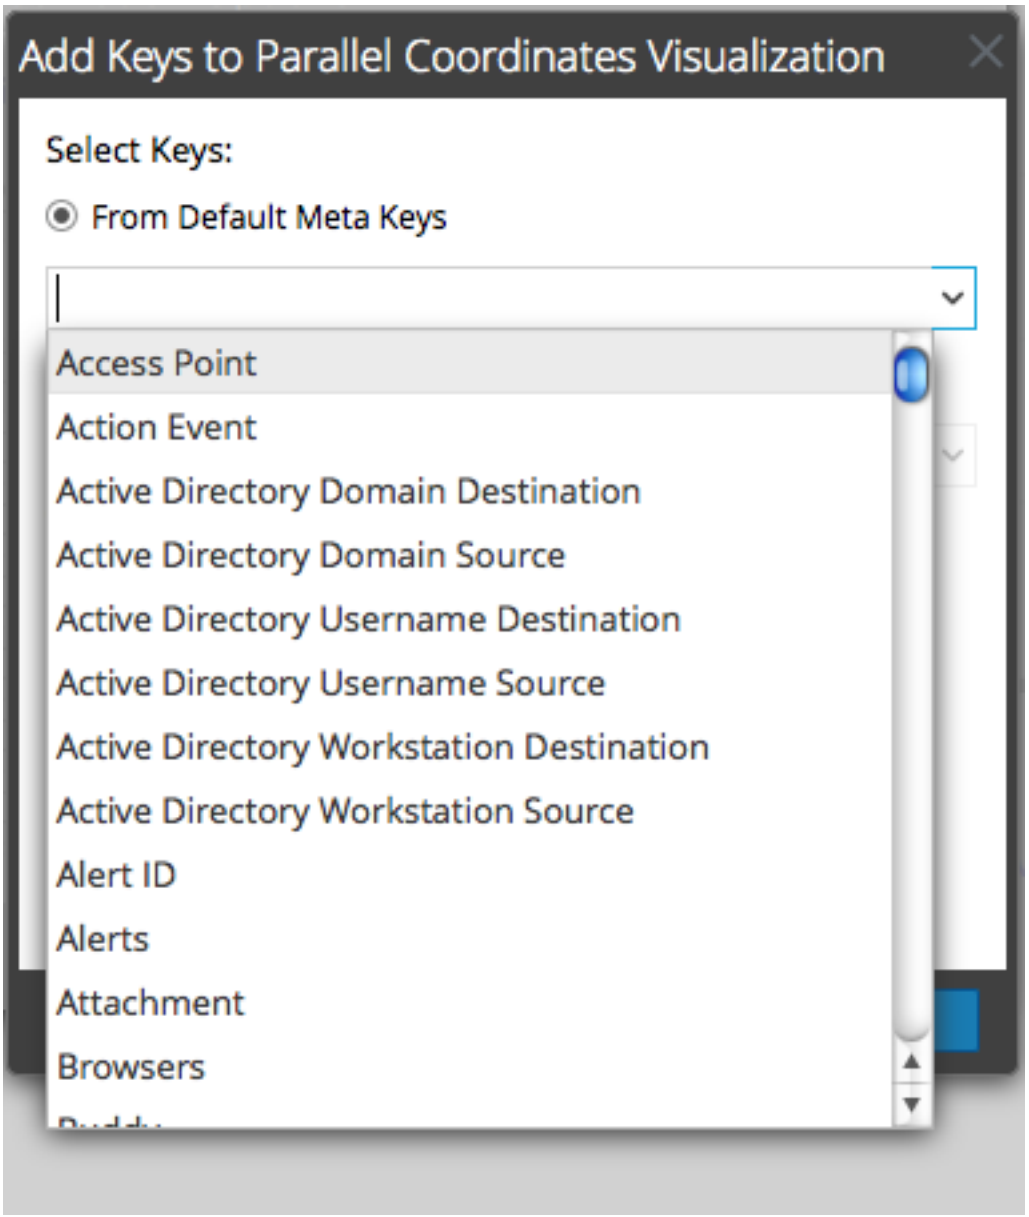
Select Meta Keys for a Parallel Coordinates Visualization

With a Parallel Coordinates visualization open, do the following:

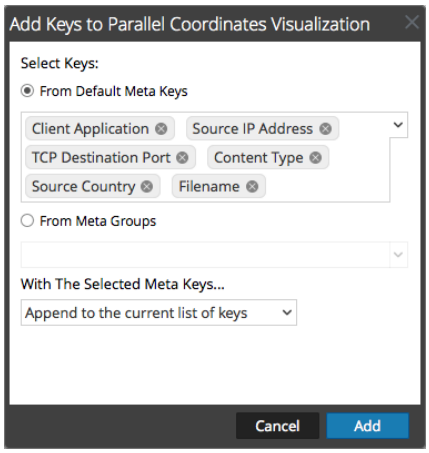
1. In the Visualization panel, select **Options**.
   The Visualization Options dialog is displayed. In the toolbar, click  to display the recommended number of axes for a readable visualization. When a recommended count of keys is displayed, the count changes based on the browser size. If you make the browser window larger, the recommended count is increased.



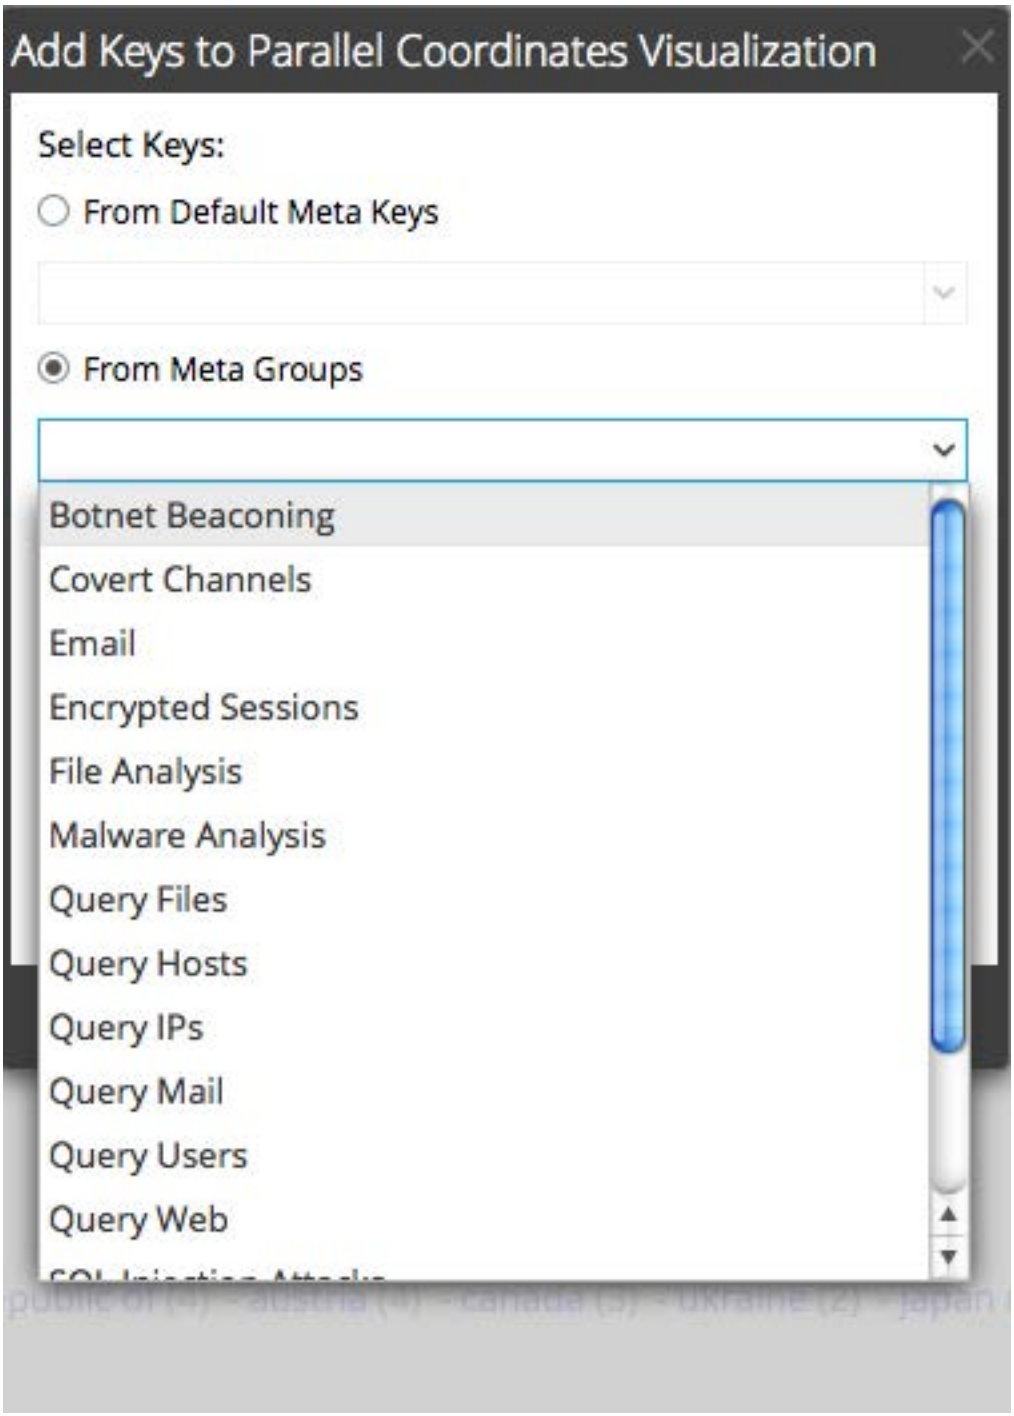
2. If you want to change the sequence of the meta keys, drag meta keys up or down to the desired sequence.
3. If you want to delete any meta keys, click in the selection box, and click .
   The meta keys are removed, but the change has not been applied.
4. If you want to revert to the previous state, click .
   Any meta keys you have deleted are restored and any changes that you made are removed.
5. If you want to select individual meta keys, click , select **From Default keys**, and in the drop-down list, select the meta keys.



The selected keys are listed.

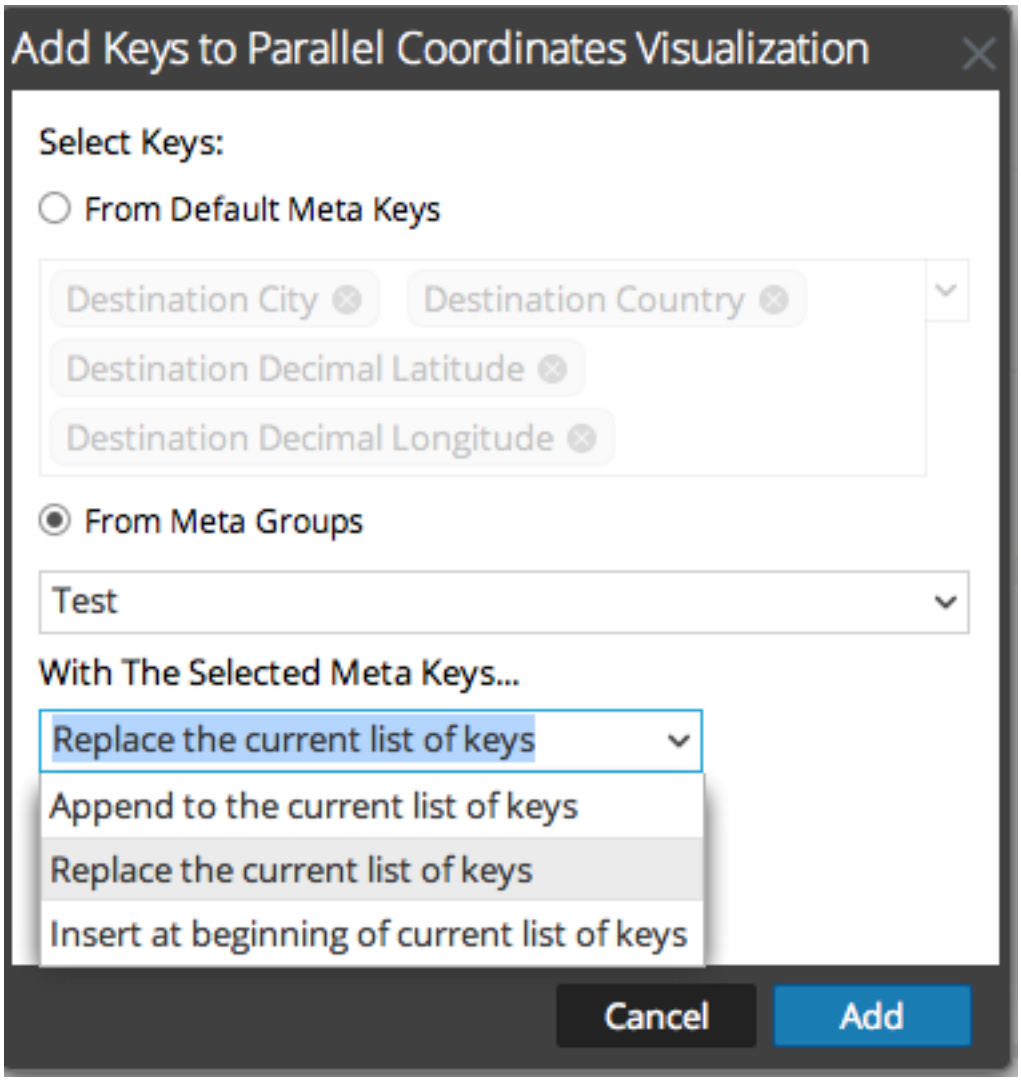


6. If you want to add all the keys in a meta group, you cannot add individual meta keys. Select **From Meta Groups**, and select a group from the drop-down list.



The selected meta groups are listed in the field.

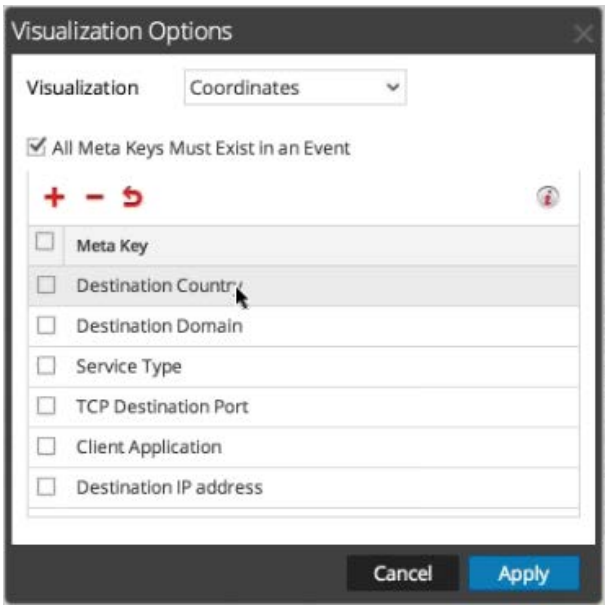
7. Select the method of adding the keys or groups: **Replace the current list of keys**, **Append to the current list of keys** (at the end), or **Insert at the beginning of current list of keys.**



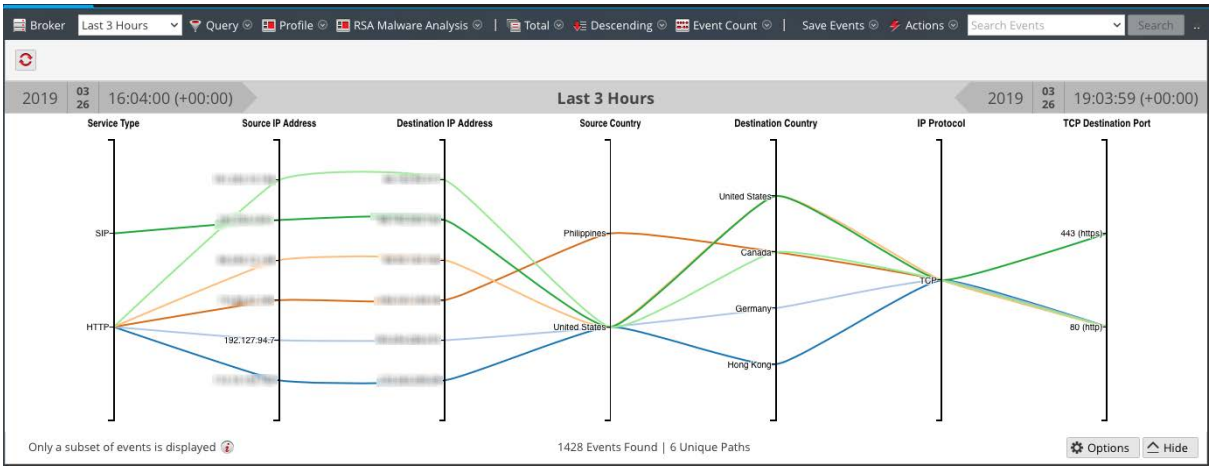
8. To complete the procedure, click **Add**.
   The Visualization Options dialog is displayed with the meta keys or groups you selected.
9. To display the new visualization chart, click **Apply**

Optimize a Parallel Coordinates Visualization

1. To optimize the visualization by removing events in which not all meta keys exist, select **Options.**



2. In the Visualization Options dialog, select **All Meta Keys Must Exist in an Event**. Click **Apply**. The resulting graph is more readable and useful and has fewer unique paths.



3. If you want to highlight a small set of points to see the path of the line from right to left, click on an axis. The cursor changes to cross hairs, which you can drag to select one or more values. When you

let go of the mouse, the lines are highlighted.



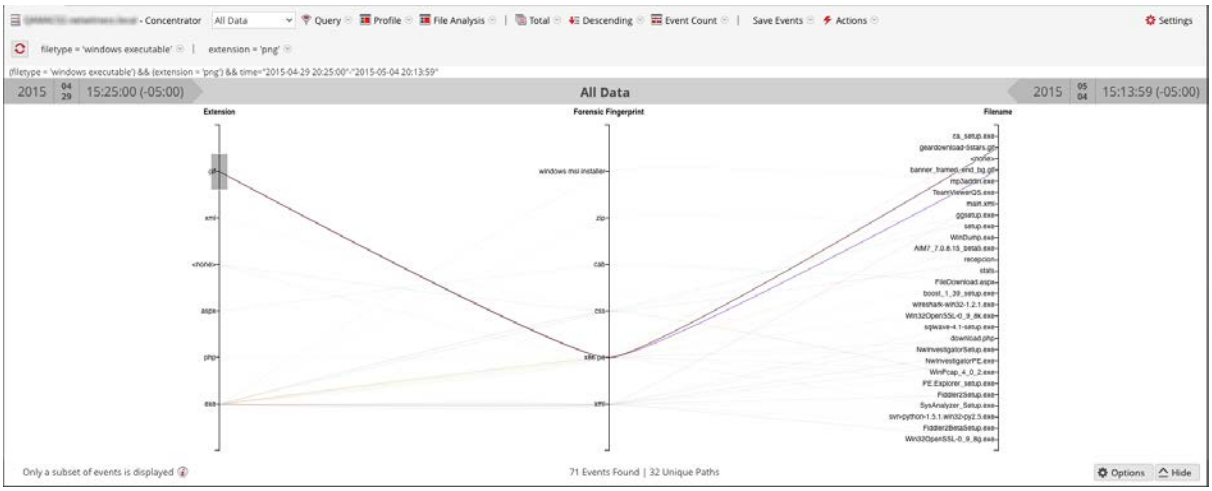
4. If you want to enlarge the visualization, drag the bottom edge of the panel down and drag the right edge of the browser window wider.

Sample Use Case

Below is an example of a parallel coordinates visualization of meta keys representing file metadata in a session. There are three meta keys or axes from left to right: Extensions, Forensic Fingerprint, and Filename with values listed along each axis. Values on the Extension axis show the file extension, and values on the Forensic fingerprint axis are windows executables. Normally the file type matches the expected forensics fingerprint; however, it is abnormal for a `gif` file type to be combined with the Windows executable fingerprint. The `gif` file type is selected to highlight the correlations of that file type, `x86pe` , and two filenames in the third axis so that an analyst can quickly identify the files that merit investigation.

To reach this view:

1. Order by Value and Sort in Ascending order.
2. Apply two filters (file type = 'windows executable' and extension = 'gif') in the Navigate view to limit the amount of data.
3. Configure a parallel coordinates chart by choosing three axes: `file extension`, `forensic fingerprint`, and `filename`.

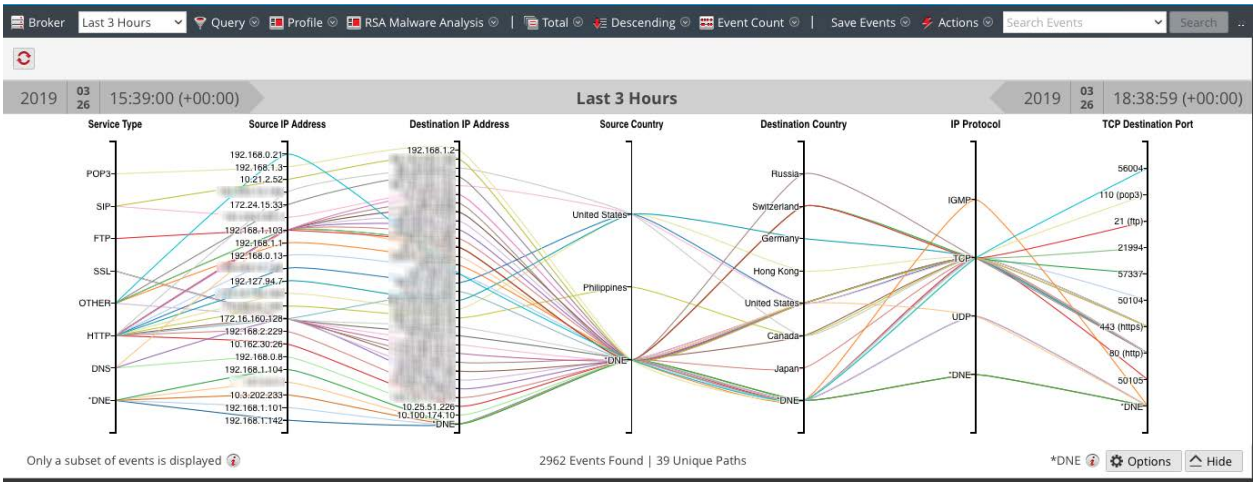


Sample Visualization of a Large Data Set

This example of a parallel coordinates visualization applied to a larger set of data illustrates several messages that help analysts to understand what has been charted.

- To create a chart, NetWitness begins scanning meta values and returning results. A typical time range could have up to 10,000,000 meta values. When the number of meta values returned reaches the Meta Values Result Limit, the chart is rendered even if NetWitness has not scanned a number of meta values equal to the Meta Values Scan Limit.
- There is a fixed limit on the amount of data that can be rendered as a parallel coordinates chart.The administrator configures parallel coordinates limits as part of the Investigation settings In the ADMIN > System view.

With a larger set of data, the parallel coordinates chart takes longer to process than the smaller set of data and meta keys. To preserve performance, NetWitness renders the meta values from the Values panel below until the limits set by the Administrator are reached. An informational message tells you: **Only a subset of events is displayed.**



Of all the data visualized for 2962 events, there were only 39 unique parallel coordinates paths. Some events are included though they do not include some of the meta keys; these are labeled **DNE** because the metadata does not exist in the event.

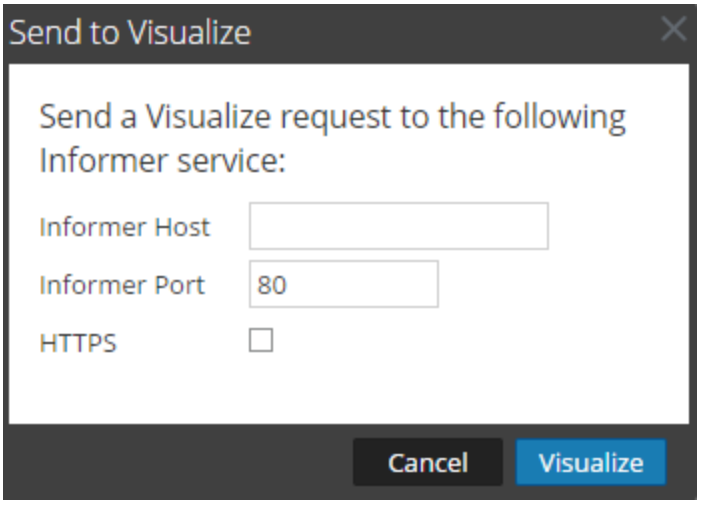
Visualize the Current Drill Point in Informer

This topic provides instructions for sending a drill point in the Navigate view to an Informer visualization.

Informer must be installed in your network and accessible by the service being investigated. You need to supply the host name and the port used on the Informer host to communicate with NetWitness.

To display a visualization in Informer of the current drill point:

1. With a drill point open in the Navigate view, click **Actions > Visualize**.
   The Send to Visualize dialog is displayed.



2. Type the Informer hostname or IP address, and verify the NetWitness server port used to communicate with the Informer host.
3. (Optional) Select the HTTPS option if the Informer host uses secured communications.
4. Click **Visualize**.
   The visualization is displayed in a new tab.

Downloading and Acting Upon Results

When working in Investigate, you may want to extract and shared data with other analysts, incident responders, SOC managers, and others. The topics in this section provide instructions for downloading results and creating incidents that appear in the Respond view:

- Download Data in the Events View
- Export or Print a Drill Point in the Navigate View
- Export Events in the Legacy Events View
- Add Events to an Incident in the Events View
- Add Events to an Incident in the Legacy Events View

Download Data in the Events View

In the Events view, you can download data from the Events panel and from a reconstruction. The Events panel download available in Version 11.4 and later is a bulk download of log and network events for all event types.

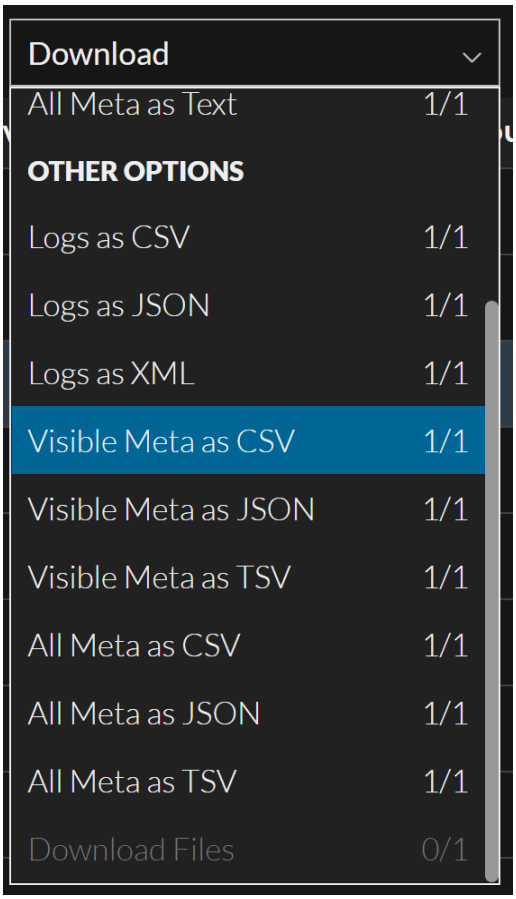
- Version 11.4.1 has the added capability to download visible metadata for all event types. From within a reconstruction, you can download events, logs, and files.
- Version 11.5 has the added capability to download metadata for all event types in the Events panel and in the event reconstructions.

Note: The information that you can view and download is managed by Role-Based Access Controls (RBAC) that your administrator has implemented. When RBAC is configured to prevent downloads of certain data, events for which you do not have download permission may appear to download successfully, but they are 0 byte in size. When RBAC is configured to prevent reconstruction of certain events, the reconstruction is disabled from the Events panel, but the bulk download button is still enabled.

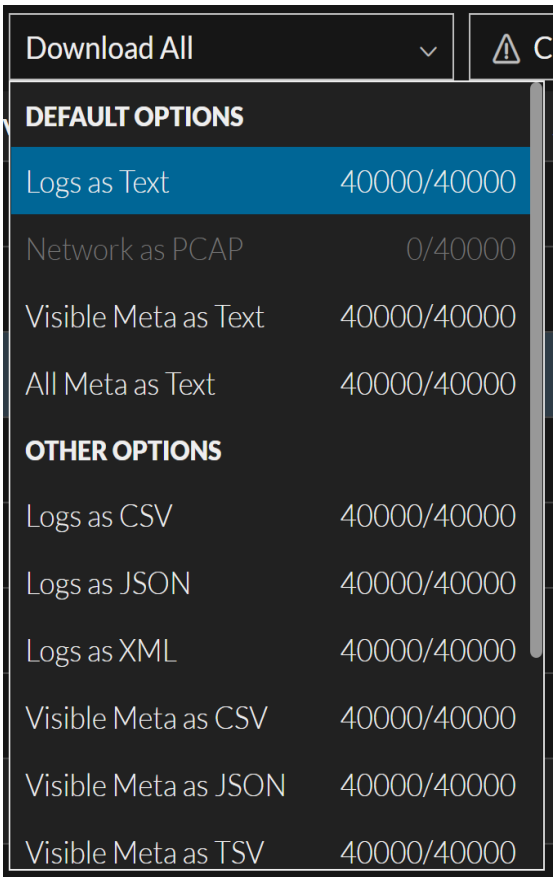
Download Events or Metadata in the Events Panel

After submitting a query, you can download logs, network events, visible metadata (Version 11.4.1), or all metadata (Version 11.5) for events in your preferred format, directly from the Events panel. The preferences are set in the Event Preferences dialog and any changes made there are reflected inside the **Download** menu. See Configure the Events View for more information about preferences.

In the Events panel, you can select events individually or select all events returned by your search. The selection checkboxes appear only if you have permission to download the events. All checkboxes are deselected when a new query is submitted. When you select events and click Download, the Download menu is displayed. The number of events selected for each event type is displayed next to each option in the format `Events of this type selected/ Total number of events selected`. If an event type has no events selected, the corresponding download option is disabled and the number of events selected is displayed as `0 / Total number of events selected` as shown in the following figure.



If the Select All checkbox is selected in the Events list, the Download All options are available. In addition to the options that download all logs or network events.



The difference between the All Meta options and the Visible Meta options are as follows:

- In Version 11.4.1 and later, visible metadata for the selected events is downloaded in the format that you selected in the Events Preference menu (**Visible Meta as Text**, **Visible Meta as CSV**, **Visible Meta as JSON**, or **Visible Meta as TSV**) or the format that you select under **Other Options** in the **Download** menu at the time of download. The downloaded metadata for each event corresponds to the columns visible when the metadata is downloaded. The visible columns are determined by the selected column group and the Column Selector. For additional information about selecting columns, see Use Columns and Column Groups in the Events List. If the Summary column group is selected in the Events panel, all metadata for the events is downloaded. When you use one of the Download Visible Meta options; the downloaded metadata is sorted in order of collection time rather than the current sort order in the Events panel.
- In Version 11.5, all metadata for the selected events is downloaded in the default format that you selected in the Events Preference menu (**All Meta as Text**, **All Meta as CSV**, **All Meta as JSON**, or **All Meta as TSV**) or the format that you select under **Other Options** in the **Download All** menu at the time of download. The resulting download includes all metadata for the events selected, regardless of what columns are visible in the Events list. For example, if an event has 40 meta keys in the meta database, even if the column group in the Events list has 20 columns with 10 columns visible, all 40 meta keys for that event are included in the downloaded file.

- In Version 11.5.1, all metadata for the selected events is downloaded in the default format that you selected in the Events Preference menu (**All Meta as Text**, **All Meta as CSV**, **All Meta as JSON**, **All Meta as TSV**, or **Download Files**) or the format that you select under **Other Options** in the **Download All** menu at the time of download. The resulting download includes all metadata for the events selected, regardless of what columns are visible in the Events list. For example, if an event has 40 meta keys in the meta database, even if the column group in the Events list has 20 columns with 10 columns visible, all 40 meta keys for that event are included in the downloaded file.

Note: When you select all events for download, only the events in the current result set are downloaded. If you canceled the query before all results were returned, only the events that were loaded are downloaded.

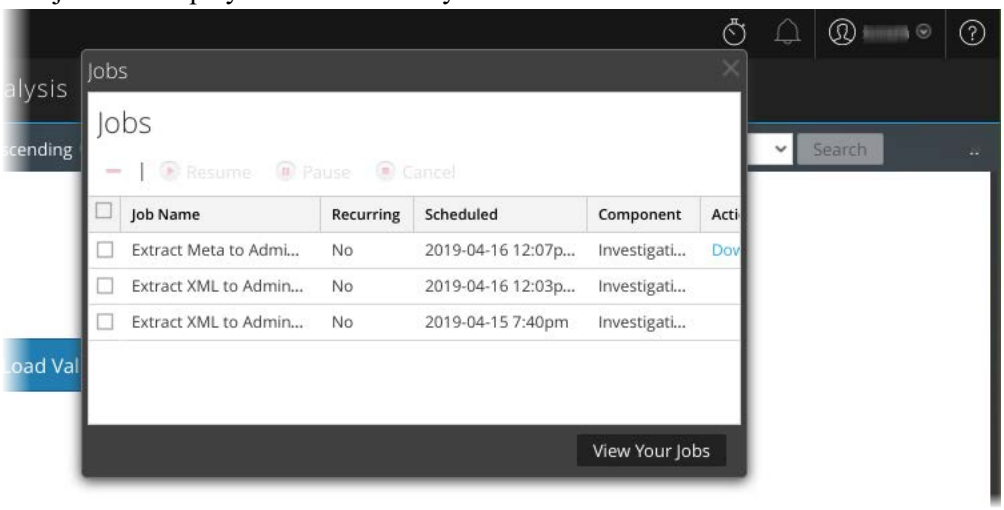
To download event data for one event, multiple events, or all events in the Events panel:

1. Do one of the following:
   a. To select events individually, select the checkbox next to each event you want to download, and click the downward arrow on the **Download** menu button to see the options.

   b. To select all events displayed in the Events panel, select the checkbox at the top of the Events panel and click the **Download All** menu button.

2. Review the **Default Options** in effect in the top section of the menu. If you do not want to use the default format, you can choose a different format from the **Other Options** section of the menu.
   - Logs are downloaded in the preferred format that you selected in the Events Preference menu (**Logs asText**, **Logs as CSV**, **Logs as JSON**, or **Logs as XML**). If you want to choose a different format for this download, select one of the formats from **Other Options**.
   - Network events are downloaded as a PCAP. When downloading multiple network events in the Events panel, the format is always PCAP. The preferred format that you specified in the Events Preference menu (**Network as PCAP**, **Network as Payloads**, **Network as Request Payload**, or **Network as Response Payload**) is ignored in this menu. Your preferred format applies only to downloading a single network event in the network reconstruction panel.
   - Visible metadata is downloaded in the format that you selected in the Events Preference menu (**Visible Meta as Text**, **Visible Meta as CSV**, **Visible Meta as JSON**, or **Visible Meta as TSV**). If you want to choose a different format for this download, select one of the formats from **Other Options**. The downloaded metadata for each event corresponds to the columns visible when the metadata is downloaded. If the Summary column group is selected in the Events panel, all metadata for the events is downloaded.
   - All metadata is downloaded in the format that you selected in the Events Preference menu (**All Meta as Text**, **All Meta as CSV**, **All Meta as JSON**, **All Meta as TSV**, or **Download Files**). If you want to choose a different format for this download, select one of the formats from **Other Options**. The downloaded metadata for each event includes all metadata, not just the visible columns.
3. Click the menu label: **Download** or **Download All**.
   The download begins immediately within the browser window if the **Download extracted files automatically** preference is set (**Events** view >). If the preference is not set, the download job for the selected events is added to the Jobs tray, where you can download the events.
   If the download fails, a message provides feedback regarding why the download failed. The download button is re-enabled and any selected events remain selected. These are examples of reasons for a failed download: timeout after X minutes, connection failed, event limit reached, and permission denied.
4. To display the Jobs tray, go to **Investigate > Navigate** or **Investigate > Legacy Events**, and click the  Jobs icon, which looks like a stop watch.

The jobs are displayed in the Jobs tray.



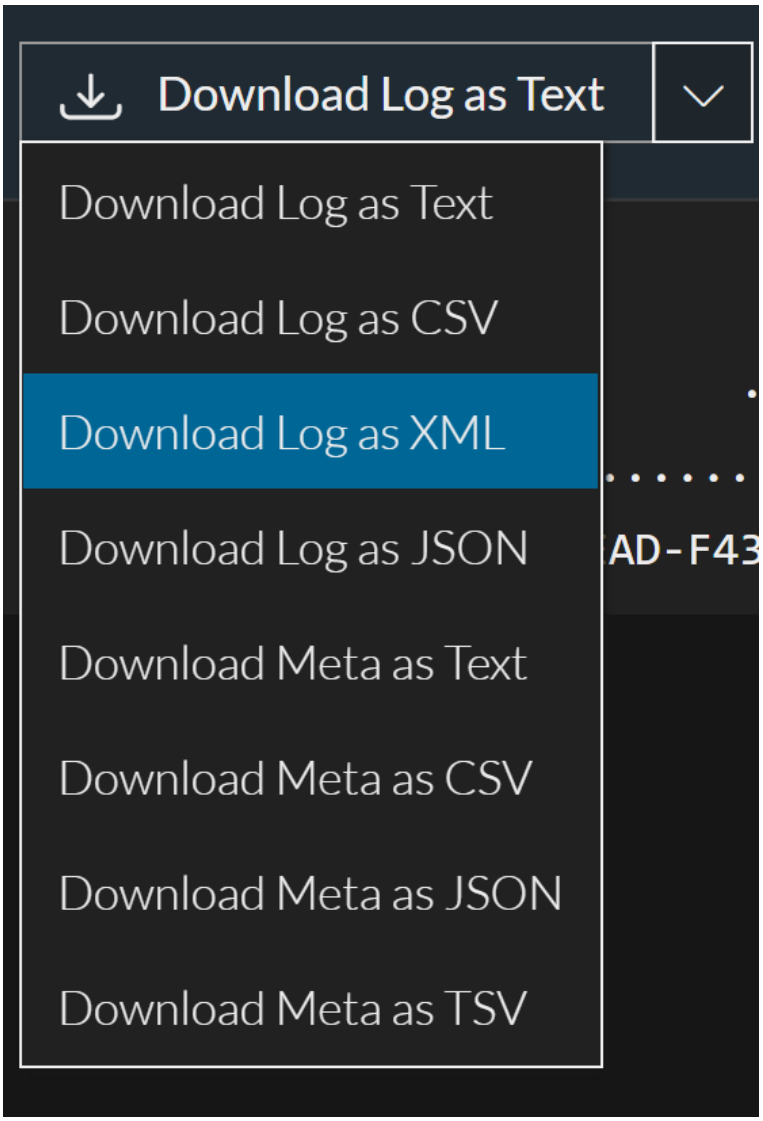
Download a Log in the Text Reconstruction

When viewing a text reconstruction of a log event, you can download a log file in the following formats using options in the Download Log menu:

- Raw log (log) using the **Download Log** (11.3) or **Download Text** (11.4 and later) or **Download Log as Text** (11.5 and later) option.
- Comma-separated values (CSV) using the **Download CSV** or **Download Log as CSV** (11.5 and later) option.
- Extensible Markup Language (XML) using the **Download XML** or **Download Log as XML** (11.5 and later) option.
- JavaScript Object Notation (JSON) using the **Download JSON** of **Download Log as JSON** (11.5 and later) option.

In Version 11.5 and later, you can also download the metadata for the log using one of these options:

Download Meta as Text, **Download Meta as CSV**, **Download Meta as JSON**, or **Download Meta as TSV**.



Note: For endpoint events, the **Download Log**, **Download Text**, or **Download Log as Text** option applies only to events that have at least one meta value exceeding 256 characters. For an endpoint event, the raw log is populated only when the meta value exceeds 256 characters. Long-running or historically downloaded files are not downloadable. For example, meta values like launch arguments can exceed 256 characters. In this case, 256 characters are available as a meta value while the full value is available in the raw log to view.

The downloaded log file contains the log and is named to help identify the service on which the log was collected, the session ID, and the file type. This is an example of the filename for a raw log: **Concentrator_SID2.log**. The exported log file is named using the following convention:

```
<service-ID or host name>_SID<n>.<filetype>
```

where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `<filetype>` identifies the format of the downloaded log. These are the possible log types: raw log, CSV, XML, and JSON. By default, the format is a raw log.

Note: Some formats do not have time stamps or the device IP where the event was generated, so a log downloaded in CSV, XML, or JSON format has an extra value called `timestamp` along with the raw log content. The additional information inside the log is in this form: `Log timestamp="1490824512" source="10.12.35.65".`

To download the log or the metadata for a log:

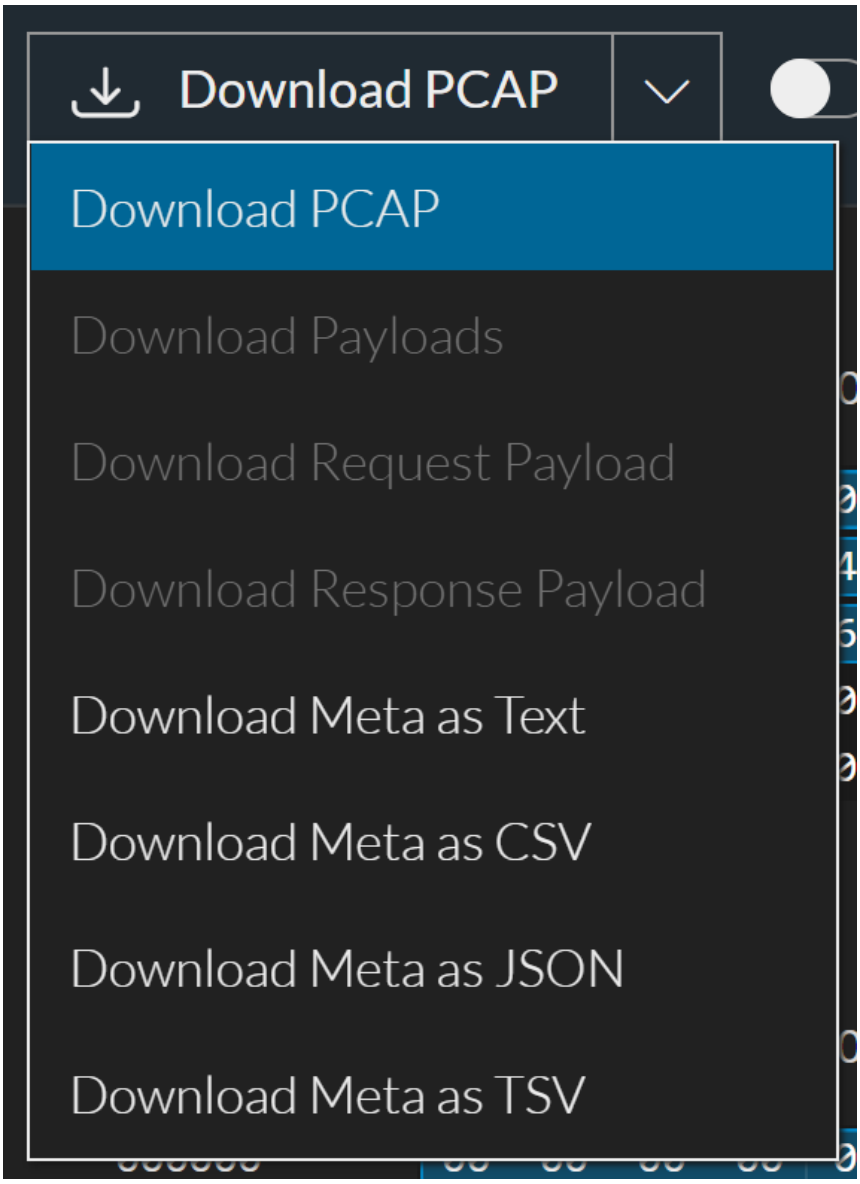
In the text reconstruction of a log event, do one of the following:

1. To download the log as a raw log (the default format), click **Download Log**, **(Download Text)** or **(Download Log as Text)**.
2. To download the log in one of the other formats, click the downward arrow on the **Download Log**, **Download Text**, or **Download Log as Text** button, then select one of the file formats for the downloaded log.
3. To download the metadata for the log, click the downward arrow on the **Download Log**, **Download Text**, or **Download Log as Text** label, then select **Download Meta as Text**, **Download Meta as CSV**, **Download Meta as JSON**, or **Download Meta as TSV**).

The log file or the metadata for the log is downloaded to your local file system in the format specified. If you initiate a download and move away from the view while the log is being extracted and before the log starts to download, the log is not downloaded in your browser. A message notifies you that you can find the downloaded log in the job queue.

Download Network Event Data in the Text or Packet Reconstruction

When viewing a packet reconstruction or a text reconstruction of a network event, you can export network data files for further analysis. In Version 11.5 and later, you can also download metadata for the reconstructed event.



The download includes events for the current time range and drill point. You can download the data in these formats:

- The entire event as a packet capture (`*.pcap`) file using the **Download PCAP** option.
- The payload as a `*.payload` file using the **Download All Payloads** (11.3) or **Download Payloads** (11.4) option.
- The request payload as a `*.payload1` file using the **Download Request Payload** option.
- The response payload as a `*.payload2` file using the **Download Response Payload** option.
- (Version 11.5) The metadata for the event using one of these options: **Download Meta as Text**, **Download Meta as CSV**, **Download Meta as JSON**, or **Download Meta as TSV**.

The label on the download menu button is one of these formats, based on the setting selected in the Event Preferences dialog. If the event does not have that type of data, the menu button is dimmed. You can click the downward arrow on the menu button to see which options are available. For example, if an event has a request payload, but no response payload, the Download Response Payload label is dimmed. You can click the downward arrow on the button and select **Download Request Payload** for this download. After selecting a valid format, clicking the button executes the download.

This is an example of the filename for a PCAP file: `C01 - Concentrator_SID1697309.pcap`. The exported network data file is named using the following convention:

`<service-ID or host name>_SID<n>.<filetype>`

where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `<filetype>` is pcap, payload, payload1, or payload2.

The network data is downloaded directly into your browser if the download is quick. If the download takes longer due to network factors or file size, the file is downloaded in the background and the task is tracked in the Jobs queue. In this case, you can check your jobs in the queue and get the file when the download is complete.

Note: If you initiate a download and move away from the view while the file is being extracted and before the file starts to download, the file is not downloaded in your browser. A message notifies you that you can find the downloaded document in the job queue.

To export an event as a network data file or to download the metadata for the event:

Go to the packet reconstruction of a network event and do one of the following:

1. To download the event as a PCAP file (the system-defined default format) or in the user-defined default format, click the **Download <format>** button. The label is the same as the download option set in the Events Preferences dialog.
2. To download the event in one of the other formats, click the downward arrow on the button, and select one of the file formats for the downloaded event data.
3. To download the metadata for the event, click the downward arrow on the button, and select one of the file formats for the downloaded metadata.

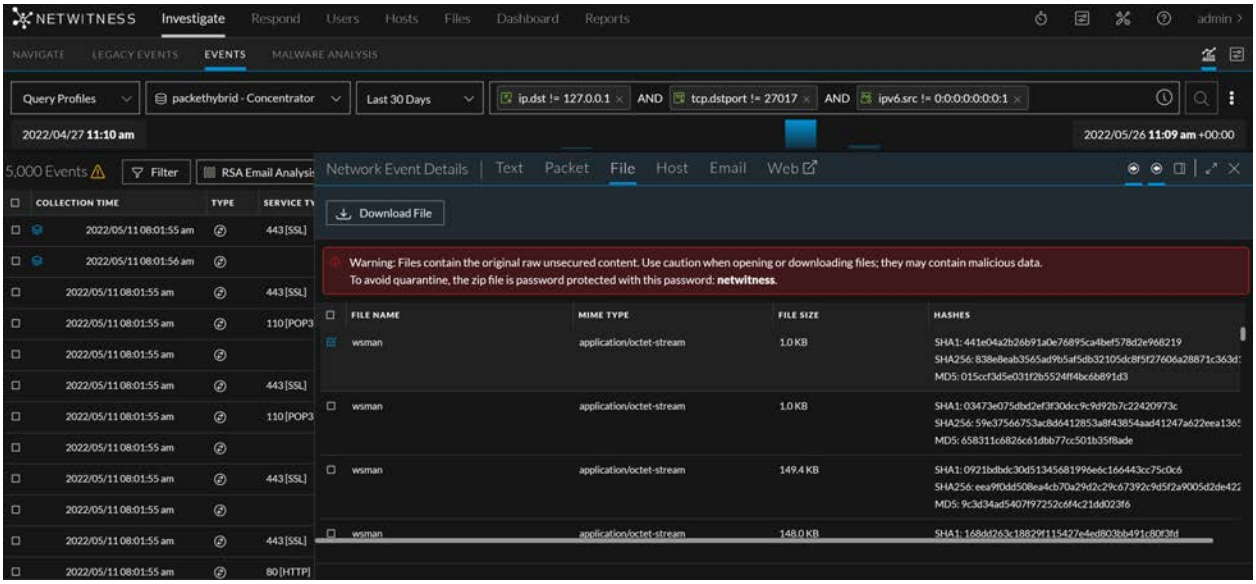
The network data file is downloaded to your local file system in the format specified or the metadata for the event is downloaded in the format specified.

Download Files from a Network Event in the File Reconstruction

When viewing reconstructed network events that contain files in the file reconstruction, you can select one or more files, or all files, to download to your local file system.

Note: If you initiate a download and move away from the view while the file is being extracted and before the file starts to download, the file is not downloaded in your browser. A message notifies you that you can find the downloaded file in the job queue.

When files are selected, the Download Files button becomes active and reflects the number of files selected.



Clicking Download Files exports the selected files as a password-protected zip archive. The password to open the exported archive is `netwitness`. Exporting the files in this form ensures that:

- The archive is not quarantined by antivirus software.
- Potentially malicious files are not automatically opened by the default application and executed.

When downloading files form the file reconstruction, the exported archive is of the form `<service-name>SID<service ID><file-count>_FILES_FILES`, for example, `Broker_SID8_1_FILES_FILES.zip`. This is the password to open the zip archive: netwitness.

`<service-ID or host name>_SID<n>_<file-count>FILES_FILES.zip`

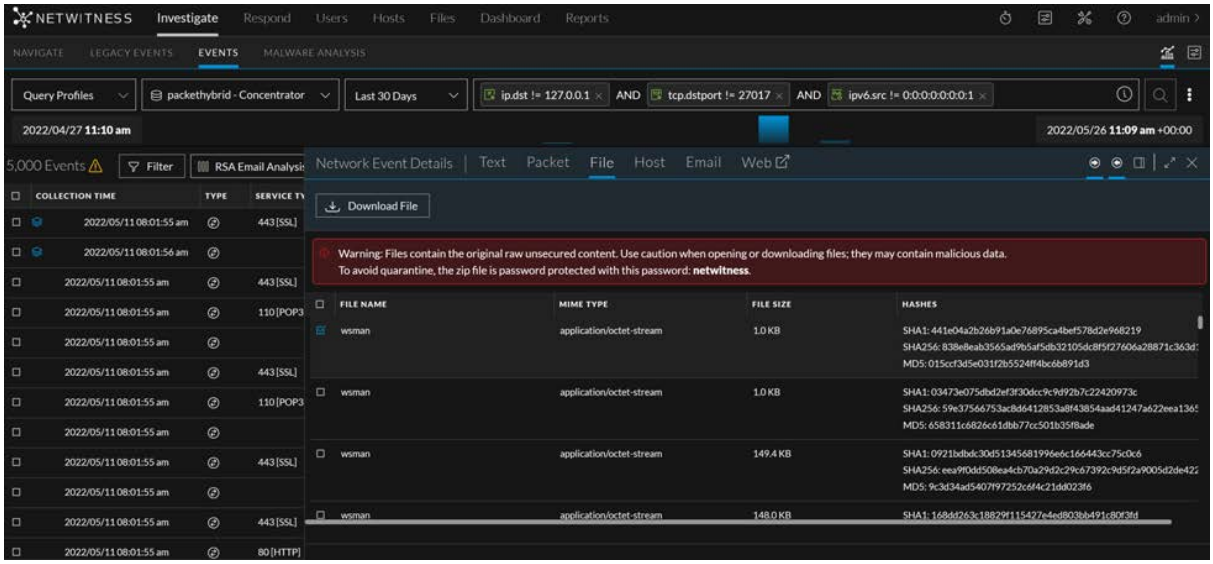
where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `<file-count> FILES` is the number of files in the archive.
- `FILES` identifies the reconstruction type from which the files were downloaded.

Caution: Caution is advised when unzipping and opening files that are associated with a default application; for example, an Excel spreadsheet may automatically open in Excel before you have a chance to verify it is safe.

To export files in a reconstructed event:

1. In the **Events** view, go to the file reconsruction of an event that contains files.



2. Click one or more files that you want to extract, and click **Download File** or **Download Files**. The job is scheduled and when complete the selected file are downloaded, in the form of a password-protected zip archive, to the local file system.
3. To open the archive on your local file system, enter the following password when prompted: `netwitness`.

Download Attachments from an Email Reconstruction

When viewing an email reconstruction that contains attachments, you can select one or more attachments, or all attachments (Version 11.4.1.x), to download to your local file system. This feature exports the selected files as a password-protected zip archive. The password to open the exported archive is `netwitness`. Exporting the files in this form ensures that:

- The archive is not quarantined by antivirus software.
- Potentially malicious files are not automatically opened by the default application and executed.

When downloading files from an email reconstruction, the filename is of the form `<service-name>_SID<n>_EMAIL`, for example, `Broker-_SID34_EMAIL.zip`. This is the password to open the zip archive: netwitness. The exported archive is named using the following convention:

`<service-ID or host name>_SID<n>_EMAIL.zip`

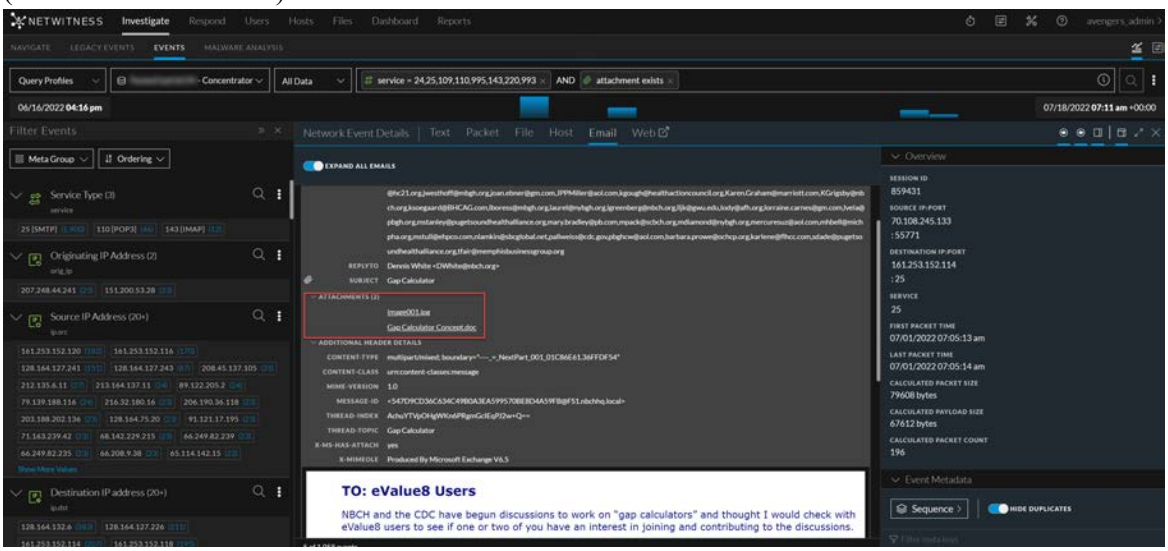
where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `EMAIL` is the type of reconstruction from which the files were downloaded.

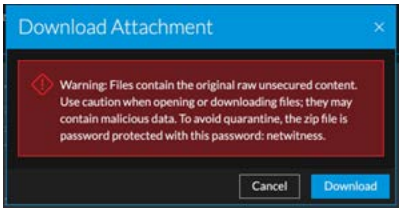
Caution: Caution is advised when unzipping and opening files that are associated with a default application; for example, an Excel spreadsheet may automatically open in Excel before you have a chance to verify it is safe.

To download email attachments:

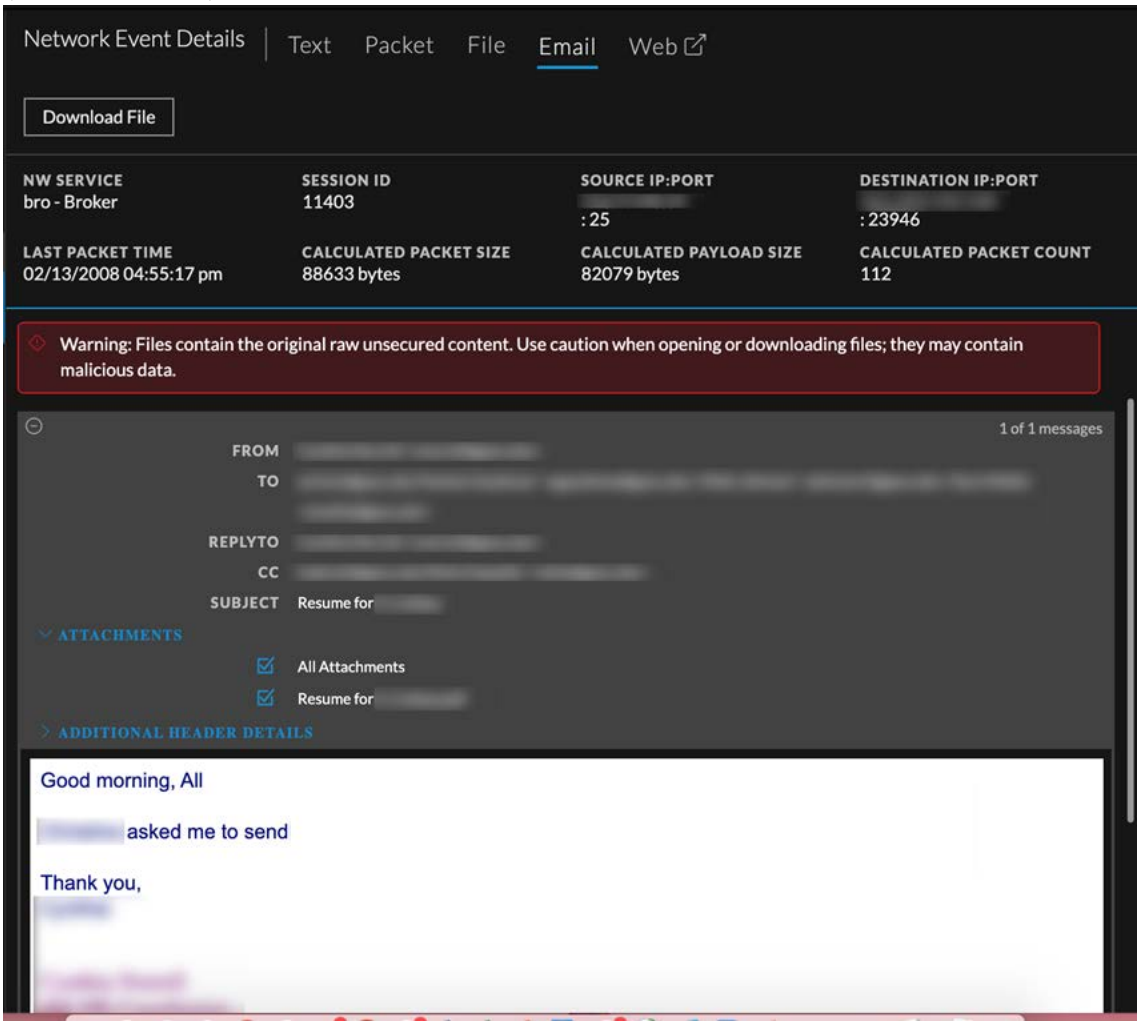
1. Go to the **Events** view and open an email reconstruction by clicking an event that contains an email with attachments.
2. Expand the **Attachments** drop-down list and do one of the following:
   a. (Version 11.5 and later) Click a link to an attachment.



A dialog warns you that downloaded email attachments may contain malicious data and asks you to cancel or confirm the download. If you want to complete the download, click **Download**. Otherwise, click **Cancel** to cancel the download.



b. (Version 11.4.1.x) Select one or more attachments or **All Attachments**.



A warning message is displayed in the reconstruction. Click the **Download File** or **Download Files** button. The attachments are downloaded with no additional opportunity to cancel.

Export or Print a Drill Point in the Navigate View

In NetWitness Investigate, when the data for a drill point is displayed in the Navigate view, you can:

- Extract files from a session and choose the type of files to extract: archives, audio BitTorrent, documents, executable, images, other, video, and web.
- Export the drill point as a packet capture (PCAP) file, a log file, or a metadata file.
- Print the drill point.

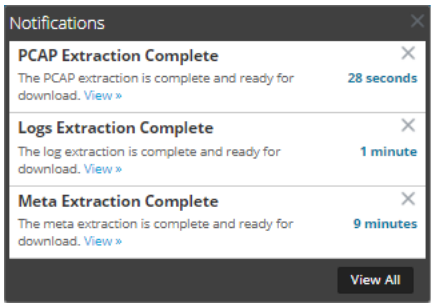
The details being exported are affected by both the time range and drill point at the time of exporting.

Note: When you export the drill point as a log file, only the log sessions are exported. The job queue message refers to the total number of sessions in the drill point rather than the number of logs. For example, if the drill point has 505 sessions and only five log sessions, the job queue message states that NetWitness is extracting logs for 505 sessions.

To export a drill point from the Navigate view:

1. Conduct an investigation until you reach the desired drill point.
2. For Version 11.0, In the toolbar, select **Actions > Export** and select one of the export options: **PCAP**, **Logs**, or **Meta**.
   The drill point is extracted, and a message advises that the job is scheduled. You can check the jobs page for the status.
3. For Version 11.1, in the toolbar, select **Save Events** > and select one of the export options: **PCAP**, **Logs**, **Files**, or **Meta**.
   A dialog gives you an opportunity to edit the default filename for the file. The default filename is in the form `investigation-Feb-21-15-44-33`. When you are exporting a PCAP, the file is exported with no choice of formats. If you are using one of the other export options, a dialog is displayed.
4. In the dialog, select:
   - The export log format: **Text**, **XML**, **CSV**, or **JSON**.
   - The file types to export: Archives, Audio, BitTorrent, Documents, Executables, Images, Other, Video, and Web.
   - The Meta format: **Text**, **CSV**, **TSV**, **JSON**.

5. When the scheduled file extraction is complete, it is displayed in the Job Notifications tray.



6. Click the **View** link in the Jobs tray and download the specific extraction file requested.

To print the current drill point:

In the Navigate view, you can display the contents of the current drill point in printer friendly format in the browser window.

To display the current drill point in a print view:

1. With a drill point open in the **Navigate** view, select **Actions** > **Print** in the toolbar.
   A new tab is created with the print view of the current drill point.



2. Use the print option in your browser to send the printable view to the printer.

Export Events in the Legacy Events View

In the Legacy Events view, the Actions menu has an option to export events from the event being viewed to an archive.

Note: You can only export files that you have permission to view or access.

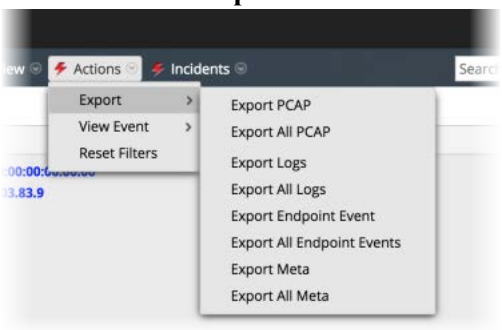
The export function queries the service for all sessions inside the selected time range and drill point to extract the content of each session. The details being exported are affected by both the time range and drill point at the time of exporting. In the File Extraction dialog, you can choose to export:

- PCAPs
- Logs
- NetWitness Endpoint events
- Meta values

The format of the exported archive: ZIP or GZIP file. After you send the request, a job is scheduled and you can track the job in in the Jobs tray. If there is an error retrieving the log or PCAP from the service, an error notification is displayed.

To extract files from an event:

1. While in the **Event view**, click an event.
2. Click **Actions > Export.**.



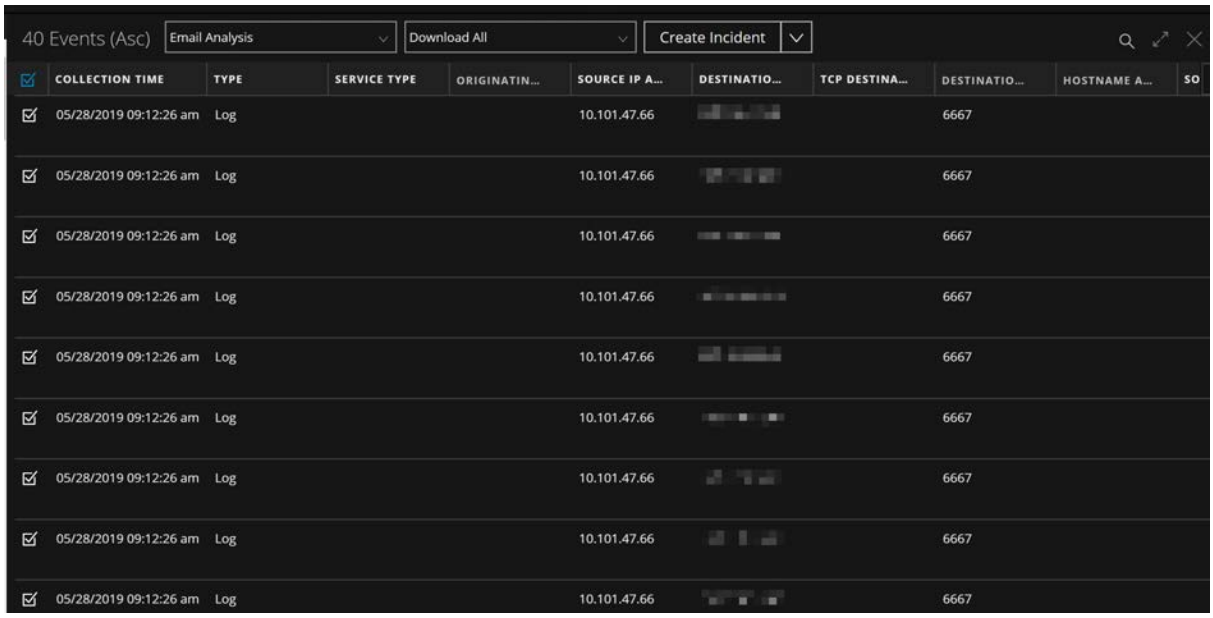
3. Select the export option.
   A message informs you that the PCAP is being downloaded.

Add Events to an Incident in the Events View

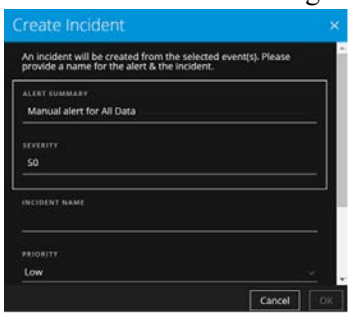
When conducting an investigation in the Events view, you can select one or more events and create an incident that is available for incident responders in Respond. When you create an incident, if access restrictions are in effect, you can view only incidents to which you have access. For example, when creating incidents from the Investigate view, analysts must assign the incidents to themselves to view them in the Respond view. You can also add events to an existing incident in Respond to which you have access.

Note: An administrator must configure the `respond-server.incident.manage` and `investigate-server.incident.manage` roles and permissions. For more information, see "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

1. Go to **Investigate > Events**.
2. In the Events view, select one or more events.



3. Click **Create Incident.**
   The Create Incident dialog is displayed. Complete the information in the Create Incident dialog.



a. Select the severity. The alert summary field is a pre-defined value which is auto-populated but can be edited if required.

b. Type a name for the incident in the **Incident Name** field.

c. From the **Priority** drop-down list, select a priority for the incident. For example, an incident may be critical, high, medium, or low priority.

d. Select an assignee for the incident from the drop-down list. This list includes the built-in users that have access to Investigate as well as any custom users that have been added to your system. For example, this list might include users for admin, analyst, dpo, operator, and users for incident responders.

e. From the Categories drop-down list, select one or more categories of events that apply to this incident.

f. Click **OK**.
An incident is created with the selected event in Investigate.

4. To add one or more events to an existing incident, select one or more events, and then click **Add to Incident**.

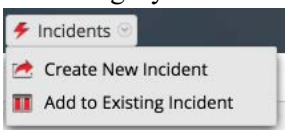
5. In the Add to Incident dialog, select the alert summary and severity, and select one or more open and existing incidents to which the incidents will be added. You can Search for an existing incident by Incident-ID or Incident Name. When ready, click **OK**. The event is added to the selected incidents and updated in Respond.

Add Events to an Incident in the Legacy Events View

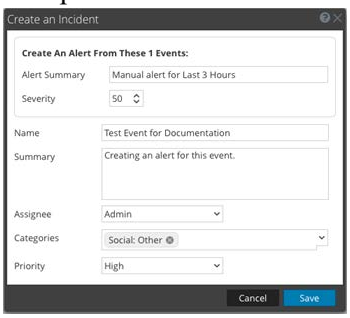
When conducting an investigation in the Legacy Events, you can select one or more events and create an incident that is available for incident responders in Respond. When you create an incident, if access restrictions are in effect, you can view only incidents to which you have access. For example, when creating incidents from the Investigate view, analysts must assign the incidents to themselves to view them in the Respond view. You can also add events to an existing incident in Respond to which you have access.

Note: An administrator must configure the required roles and permissions as described in "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

1. Go to **Investigate > Legacy Events**.
2. In the Legacy Events view, select one or more events, and then **Incidents > Create New Incident**.



3. Complete the information in the Create an Incident dialog.



   a. Select the severity, an integer between 1 and 100, with 100 being the most severe.
   b. Type a name for the incident and describe the incident in the **Summary** field.
   c. Select an assignee for the incident from the drop-down list. This list includes the built-in roles that have access to Respond as well as any custom roles that have been added to your system. For example, this list might include roles for admin, analyst, dpo, operator and roles for incident responsers.
   d. From the **Categories** drop-down list, select one of more categories of alerts that apply to this incident.
   e. From the **Priorities** drop-down list, select a category for the incident. For example, an incident may be critical, high, medium, or low priority.
   f. Click **Save**.
   The new incident is created and is available immediately in the incident queues for the selected role in Respond.

4. To add one or more events to an incident, select one or more events, and then **Incidents > Add to Existing Incident**.
5. In the Add Events to an Incident dialog, select the severity, and select one or more incidents to which the events will be added. You can Search for an existing incident by Incident-ID or Incident Name. When ready, click **Add to Incident**.
   The events are added to the selected incidents and updated in Respond.

Troubleshooting NetWitness Investigate

This section provides information about possible issues when using NetWitness Investigate.

Navigate View and Legacy Events View Issues

Behavior	A meta key that normally returns values in the Navigate view returns values, but has a `Not Indexed` message following the meta key name. For example, the in this figure the Service Type meta key is followed by the message: `Service Type[service] Not Indexed`. 
Issue	When you first set up the environment or very rarely after performing a data reset on the broker due to other issues, you see meta keys as Not Indexed when they are indexed at meta key or meta values level.
Explanation	To fix the issue on a Broker, log out of NetWitness Platform and then log in again. Valid sessions will be displayed.

Message	`Not indexed; will experience longer than usual load times.` in the Manage Meta Groups dialog.
Issue	Meta keys in the Manage Meta Groups dialog are marked by a red exclamation point, and the error message is displayed. This can occur when investigating a Broker or Decoder while adding a meta group with meta keys that are not indexed in the index file or the custom index file for the service. For a Broker, it could mean that the Broker has not begun aggregating data from a Concentrator. In this case the Broker will not have the contents of the custom index file from the aggregate services and the keys will not be indexed. For a Decoder, it means that the meta keys are not indexed in the Decoder index or custom index file.
Explanation	To fix the issue on a Broker, log out, log in, and restart the Broker service so that it can aggregate the meta key information from connected Concentrators. To fix the issue on a Decoder, edit the custom index file to index the meta keys, log out, log in, and restart the Decoder service.

Behavior	When downloaded from the Event Reconstruction view, logs and metadata are always in text format irrespective of the format selected in the Legacy Events view.

Issue	When you download metadata or a log in the Event Reconstruction view, the format that you selected in the Legacy Events view is not used. The exported data is always in text format.
Explanation	Download metadata and logs from the Legacy Events view if you want to use a format other than text format.

Events View Issues

Message	`Applicable for hosts with 4.x Endpoint agents installed, please install the NetWitness Endpoint Thick Client.`
Issue	When you click **Pivot to Endpoint** in the Events view, no data is displayed and the message is displayed.
Explanation	Version 4.4 of the NetWitness Endpoint Thick Client must be installed on the same server, the NWE meta keys must exist in the `table-map.xml` file on the Log Decoder, and the `index-concentrator-custom.xml` file on the Concentrator. The NWE Thick Client is a Windows only application. Complete setup instructions are provided in the *NetWitness Endpoint User Guide* for Version 4.4.

Behavior	Download jobs are in a Waiting state or Failed state in the Jobs tray during and after upgrading the software to Version 11.4.
Issue	If you had download jobs running while your administrator was upgrading the software, you may see a job in a Waiting state while the upgrade is in progress and then in a Failed state after the upgrade is complete. You cannot resume or cancel the failed job.
Explanation	To delete the failed jobs, select the failed jobs in the Jobs tray and click .

Message	Event counts in the Filter Events panel and the Events panel may sometimes differ when showing results for the same query.
Issue	The Filter Events panel uses only index data to produce counts of events, which is less accurate than the Events panel. The Events panel results are filtered for exact matches on data retrieved from the meta database, which takes longer to process.
Explanation	At worst the difference is in false positives in the Filter Events panel, not false negatives; so you will not miss an event.

Message	`Event Analysis requires all core services to be NetWitness 11.1. Connecting prior versions of services to the 11.1 NetWitness Server results in limited functionality (see "Investigate in Mixed Mode" in the *Physical Host Upgrade Guide*).`
Issue	When attempting to investigate a service that has not been updated to Version 11.1 in the Event Analysis view, the informational message is displayed.
Explanation	When an analyst opens the Event Analysis view in mixed mode (that is, some services are upgraded to 11.1 and later, and some are still on 11.0.0.x or 10.6.x), Role-Based Access (RBAC) is not applied uniformly. This affects viewing and downloading content, and validation of filters in the interactive breadcrumb. You will see this informational message when you open Events. As you select a service, services that are not up to date are displayed in a red box, with the message that the service is not up to date. When your administrator has upgraded all connected services to 11.1 and later, these features work as expected.

Message	`Forbidden. You cannot access the requested page.`
Issue	When attempting to access the Events view, the view opens with the message.
Explanation	Your administrator has prevented access to the Events view using role and permissions.

Behavior	If you can download an event in the Events view, but get a 0-byte file, the administrator may have restricted access to the content.
Issue	Role-Based Access Controls applied by your administrator allowed you to download an event for which you did not have permission; therefore, the file download was empty.
Explanation	If you believe you should have access to the event, contact your administrator.

Message	`Insufficient permissions for the requested data.`
Issue	While attempting to access an event in the Events view, the message is displayed.
Explanation	You have entered an event ID for an event that you do not have permission to view. The administrator may have placed some restrictions to limit access by role and permissions.

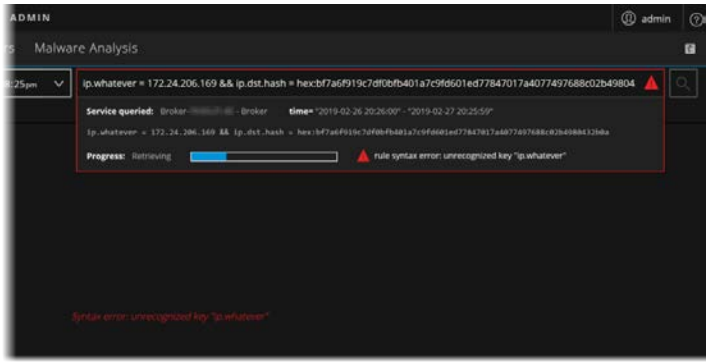
Message	`Invalid session ID: <<eventId>>`
Issue	No `sessionId` matches the `sessionId` that you queried.
Explanation	The reason for an invalid session ID can vary. Perhaps you edited the session ID manually, and no such session exists. Another case may be when you query a Broker, and the aggregated data has not been refreshed, you may see this error for a session that no longer exists.

Behavior	Investigation Profiles and built-in column groups are not present in 11.1 Event Analysis.
Issue	Post upgrade to NetWitness v11.1, the default column groups - Endpoint Analysis, Outbound SSL and Outbound HTTP are not added under column groups. Also, a few of the Investigation Profiles are missing post upgrade.
Explanation	It is observed that this issue occurs only when you have created a custom column group with the name which is same as one of the new 11.1 OOTB custom column group name. For example, if you create a custom column group in 11.0 with name **RSA Endpoint Analysis** then after upgrade to 11.1. Due to the same name already existing in 11.1, OOTB column groups and built-in profiles will not be available in the UI. To fix this, change the name of custom column group to something other than one of the OOTB column groups and restart the jetty server by using the following command on the NetWitness server: `systemctl restart jetty`

Message	`Memory limit of <XXXXXX> GB reached, controlled by setting max.query.memory`
Issue	The query that you submitted failed because the result set was too large, and the memory limit set by `max.query.memory` was reached.
Explanation	To avoid this error, try to further limit results by narrowing the time range, adding filters, and decreasing the number of columns in the column group. You can also ask an administrator to limit the number of events returned.

Behavior	No text data was generated during content reconstruction. This could mean that the event data was corrupt or invalid, or that an administrator has disabled the transmission of raw endpoint events in the Endpoint server configuration. Check the other reconstruction views.
Issue	When you reconstruct an event as text in the Events view, no data is displayed and the message is displayed.
Explanation	If you do not see the raw text in other Events views or Legacy Events view reconstructions, and you believe the data is not corrupted or invalid, your administrator has likely disabled transmission of raw endpoint events on the NetWitness Endpoint server. Contact your administrator for additional information.

Message	`Rule Syntax error: Unrecognized key "<meta key or meta entity name>"` `Syntax error: Unrecognized key "<meta key or meta entity name>"`
Issue	While querying a service, the matching events are not listed and the message is displayed in the query console and the Events view.

	
Explanation	The query you entered is querying a meta entity that is not configured properly. All upstream devices connected to the Broker being queries should have the same entity configuration. This error indicates that the Broker is operating with mismatched entity definitions. Ask your administrator to review the configuration described in "Index Customization" in the *Core Database Tuning Guide*.

Message	`Selected Column Group is no longer available. The default summary column group has been selected instead.`
Issue	If you had set a preferred column group before the 11.4 upgrade, on your first visit to the Events view, the flash message is displayed even when the column group is available or is the default group (summary). This issue was resolved in Version 11.4.1.
Explanation	This is a one-time occurrence. If you reload the Events view, the message is not displayed.

Message	`Session is unavailable for viewing.`
Issue	While querying an event ID, the reconstruction is not displayed and the message is displayed.
Explanation	The query you entered is trying to look at restricted data, for example, if you are allowed to see only log data and you are using a link to network data .

Message	`The query on channel <channel-number> was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`
Issue	If you continually get this timeout message, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or other warnings that may need to be addressed to increase query response time.
Explanation	If there are no messages indicating any specific warnings, ask your administrator to increase the Core Query Timeout from 5 minutes to 10 minutes as described in the *System Security and User Management Guide*.

Message	`The session id is too large to be handled:<<eventId>`
Issue	The session id that you typed in, or got from the Legacy Events view or Navigate view is too large.
Explanation	If you manually typed the sessionId or edited a sessionId in the Events view, you may have created an integer that is too large for Events to process.

Behavior	When reconstructing network events with a large number of packets (>250) in the Events view > Packets panel, with the option to display only payloads enabled and the packets per page setting higher than the default (100), the current browser tab may become unresponsive for up to 45 seconds as it is working to render the payloads.
Issue	Depending on the amount of resources (memory and CPU) on the client machine and the number of packets in the event there may be a performance lag when displaying only payloads in packet reconstruction.
Explanation	To limit the amount of data processed in a reconstruction of a single event, change the Packets per Page setting in the footer to a lower value. 

Behavior	When working in the Version 11.4 Events view, the Query Profile drop-down menu and Column Group drop-down menu do not function.
Issue	You do not have permission to read columns groups and profiles. The default column group , Summary List, is applied to the Events list, and you cannot change the column group, create a column group, or delete a column group.
Explanation	This occurs only when the administrator has created a custom role for you instead of assigning the default Analyst role. Ask your administrator to enable column group read and profile read permission for your role.

Issue	No matching Endpoint data available on Investigate > Events view > Host tab.
Explanation	The Endpoint data may not be available due to any of the following: • **No Endpoint Deployment** – You must install Endpoint Log Hybrid, see "NetWitness Endpoint" in the *Physical Host Installation Guide*. • **Endpoint data is not captured for the host associated with the selected network event-** Make sure that NetWitness Endpoint Agent is installed, and expanded network visibility is configured to track the network events. To enable expanded network visibility, see "Creating Groups and Policies" in the *NetWitness Endpoint Configuration Guide*. **Note:** For Expanded Network Visibility to work, ensure the service user account

used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Services Security View - Aggregation Role" in the *Hosts and Services Getting Started Guide for NetWitness Platform*.

- **Concentrators or Endpoint services are offline or very slow** – You must check the status (online or offline) of the services on Health and Wellness. If the service is online, you must check the Endpoint server logs and (Endpoint) Concentrator logs for details.
- **Endpoint data is rolled over for the host associated with the selected network event** - The Endpoint data may be rolled over due to data retention period configured. You must configure the data retention period to retain the endpoint data for a longer period. For more information, see “Configure Data Retention” in the *Data Privacy Management Guide*.

Investigate Reference Materials

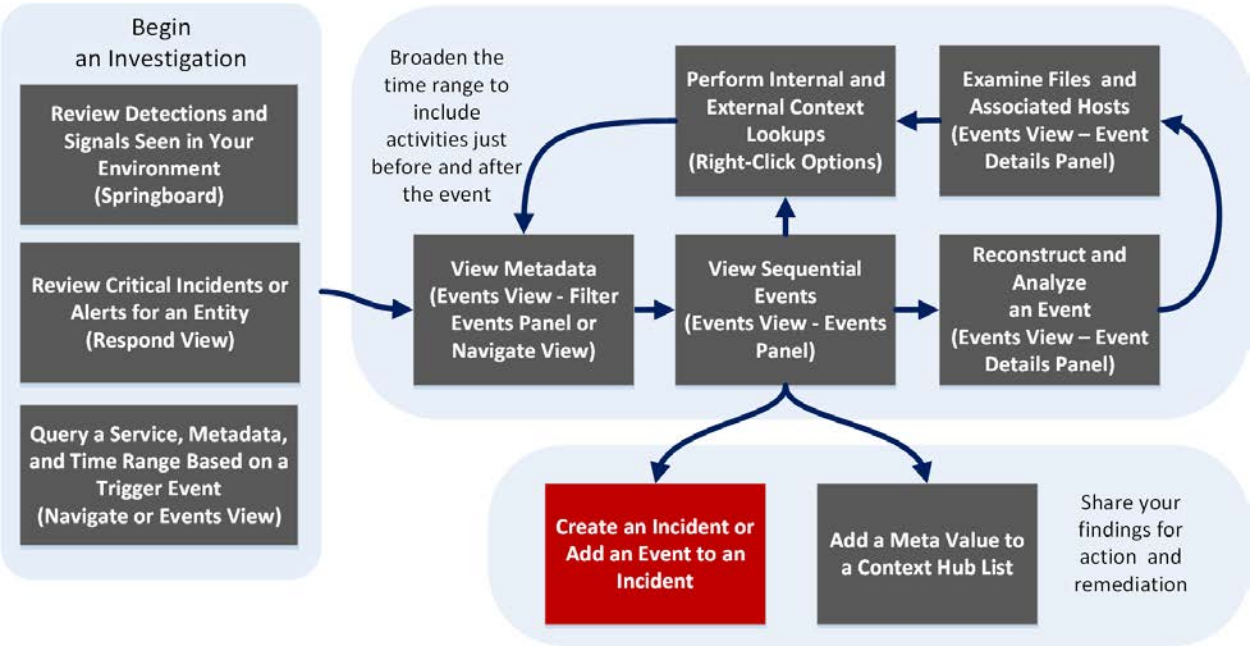
This section provides is intended to help you understand the purpose and application of NetWitness Investigate views. For each view, there is a brief introduction and a What Do You Want To Do table with links to related procedures. In addition some of the reference materials include workflows and Quick Looks to highlight important features in the user interface.

- Investigate View
- Navigate View
- Legacy Events View
- Events View
- Add/Remove from List Dialog
- Add Events to an Incident Dialog
- Column Groups Dialogs
- Context Lookup Panel
- Create an Incident Dialog
- Events View - Email Tab
- Events View - Text Tab
- Events View - Packet Tab
- Events View - File Tab
- Investigate Dialog
- Investigation Tab - User Preferences Panel
- Legacy Event Reconstruction View
- Manage Default Meta Keys Dialog
- Meta Groups Dialogs
- Navigate View
- Query Dialog
- Query Profiles Dialogs
- Generate Springboard Panel Dialog
- Settings Dialogs for Investigate Views

Add Events to an Incident Dialog

In the Add Events to an Incident dialog, analysts can add alerts to an existing incident so that incident responders look at the associated events as part of an incident response. To access this dialog while investigating a service in the Events view and the Legacy Events view, see Add Events to an Incident in the Events View and Add Events to an Incident in the Legacy Events View.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View

User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

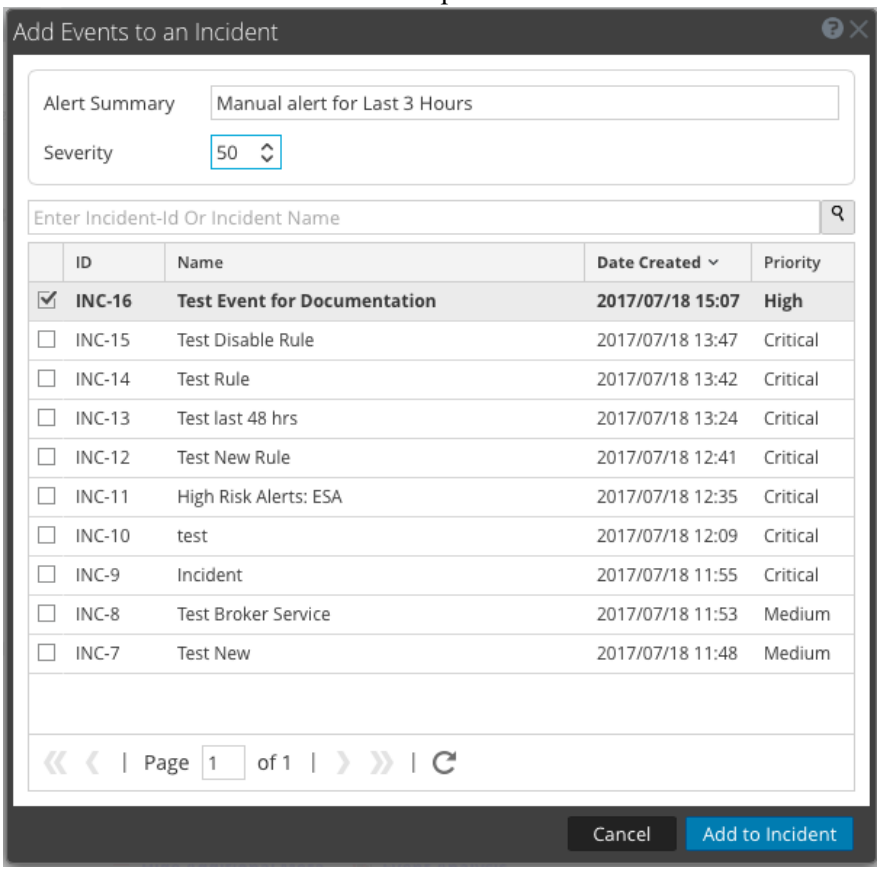
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View
- Legacy Events View

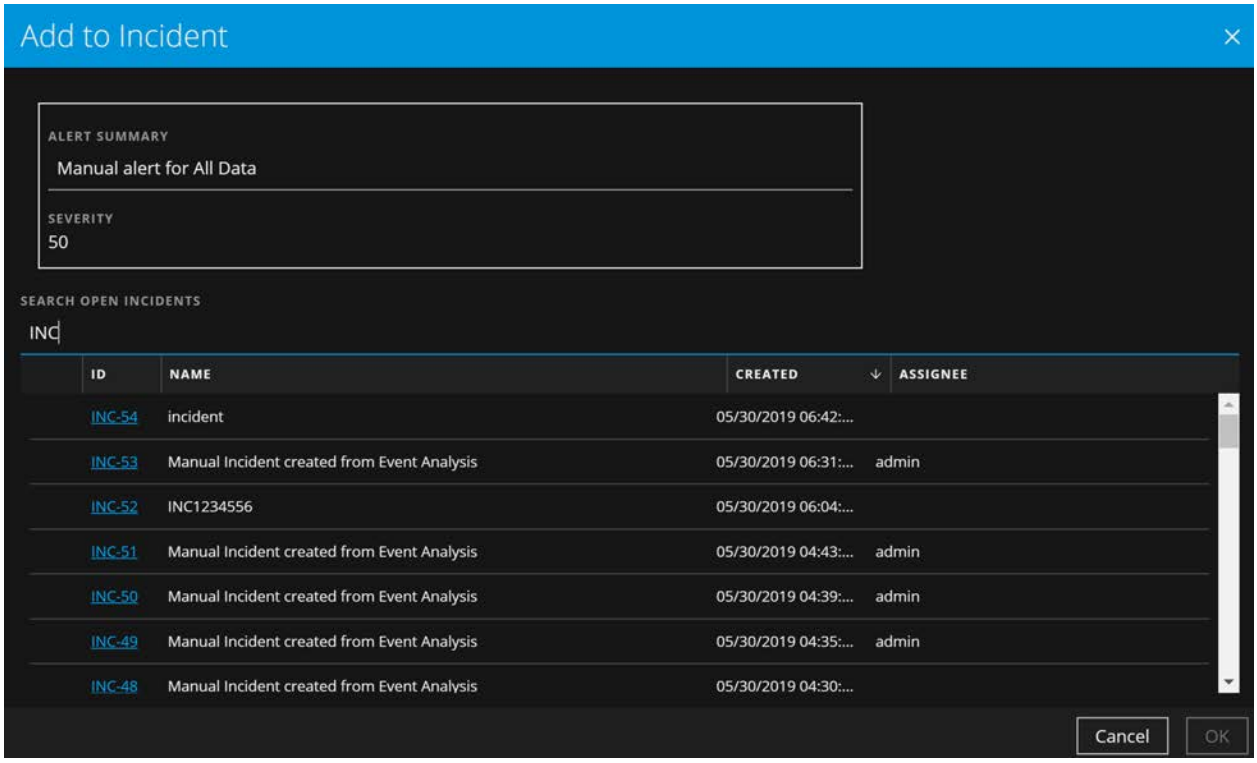
Quick Look

The following figure is an example of the Add Events to an Incident dialog in the Legacy Events. The table describes the information and options in the Add Events to an Incident dialog .



Feature	Description
Alert Summary	The Alert Summary field is filled by the query that produced the select alerts, which you selected to create this incident. The Severity field reflects the Severity of the selected alert, an integer between 1 and 100.
Search	Allows you to search for an existing event.
ID	The ID of the incident. You can sort IDs in ascending or descending order.
Name	The incident name. You can sort the Name in ascending or descending order.
Date Created	Displays the date and time the incident was created. You can sort the dates in ascending or descending order.
Priority	Displays the priority of the incident: either low or critical.
Cancel	Closes the dialog without saving changes.
Add to Incident	Adds the alerts to the incident. A dialog confirms that alerts are successfully added

The following figure is an example of the Add to Incident dialog in the Events view. The table describes the information and options in the Add to Incident dialog.



Feature	Description
Alert Summary	The Alert Summary field is filled by the query that produced the select alerts, which you selected to create this incident.
Severity	The Severity field reflects the Severity of the selected alert, an integer between 1 and 100.
Search Open Incidents	Allows you to search for an existing incidents.
ID	The ID of the incident.
Name	The incident name.
Created	Displays the date and time the incident was created.
Assignee	Displays the team member currently assigned to the incident
Cancel	Closes the dialog without saving changes.
OK	Adds the alerts to the incident. A confirmation message is displayed after the incident is successfully added

Add/Remove from List Dialog

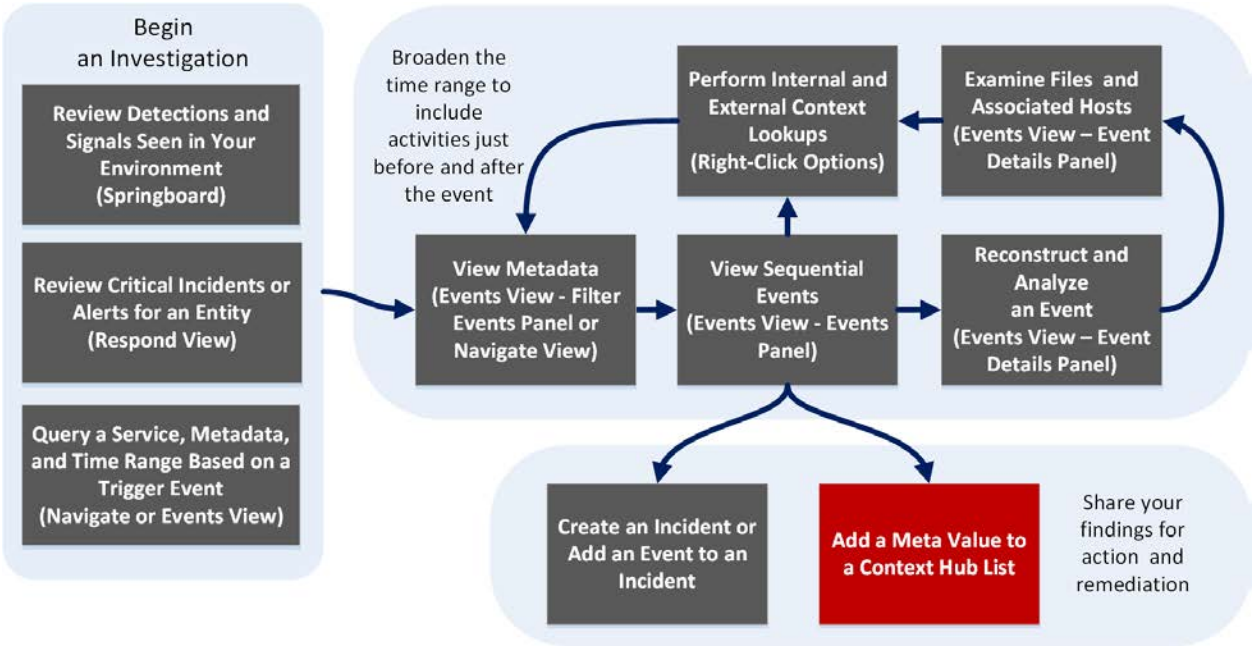
The Add/Remove from List dialog allows you to add an entity or meta value to an existing Context Hub list, remove an entity or meta value, or create a new Context Hub list containing the entity or meta value. When you look up an IP address or other entity and you find it suspicious or interesting, you can add it to a list that has been added as a data source. An example of a commonly used list is a white list or black list. This improves the visibility of the suspicious IP addresses and reduces false positives that do not need further investigation.

You can add entities or meta values to more than one list. For example, you can add them to one list for suspected domains related to command and control connections and to another list for Trojan connection IP addresses related to remote access. If a list is not available, you can create a list.

The dialog is available in NetWitness Investigate and in NetWitness Respond. When working in Investigate, in the Navigate view, Legacy Events view, or Events view, you can add meta values for the `Source IP`, `Destination IP`, or `Username` meta keys to an existing context hub list or you can create a new list containing the meta values. When you add meta values to a list, you can look up additional context on those meta values.

- To display the dialog in the Navigate view or the Legacy Events view, right-click a meta value under `Source IP`, `Destination IP`, or `Username`) and select **Add/Remove from List** in the context menu.
- To display the dialog in the Events view, hover over a value and select **Add/Remove from List** in the Actions section of the context tooltip.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

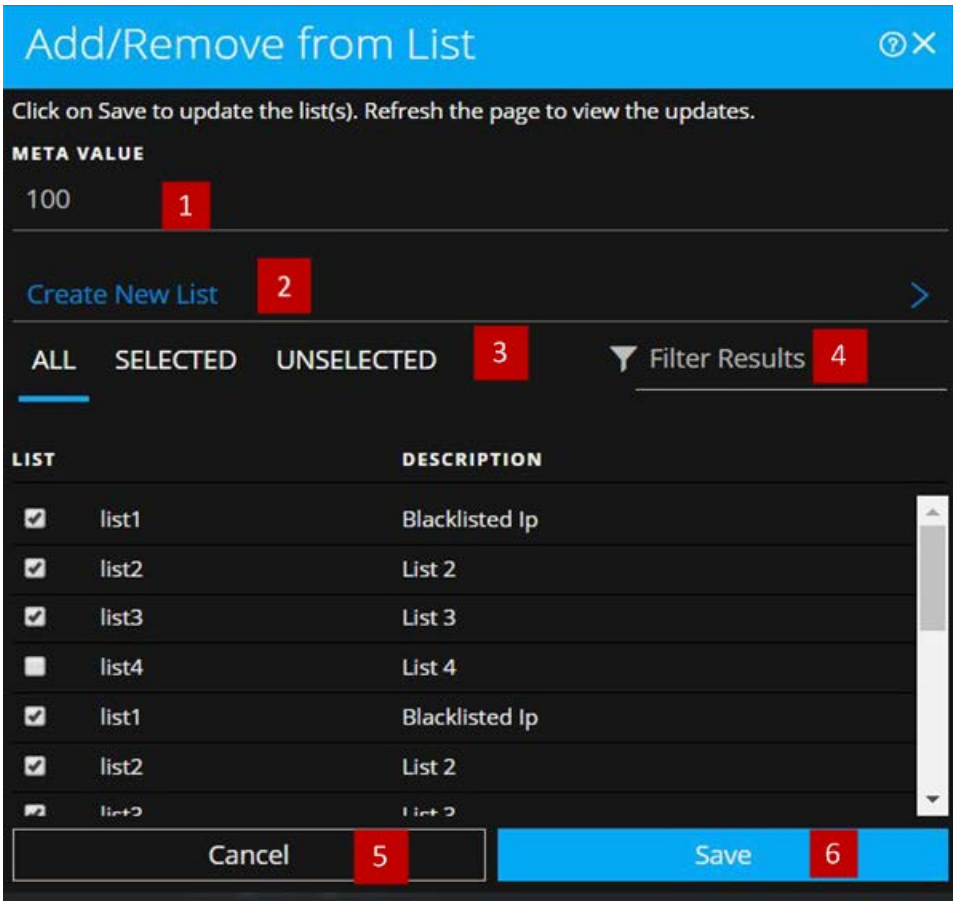
Related Topics

- Look Up Additional Context for Results
- Navigate View

- Legacy Events View
- Events View

Quick Look in the Events View

The following is an example of the **Add/Remove from List** dialog in the Events view.



1 Entities or meta values to be added or removed.
2 Create a new list using the selected meta.
3 Select any of the tabs: All, Selected, or Unselected.
4 Search using the list name or description.
5 Cancel the action.
6 Save to update lists or create a new list.

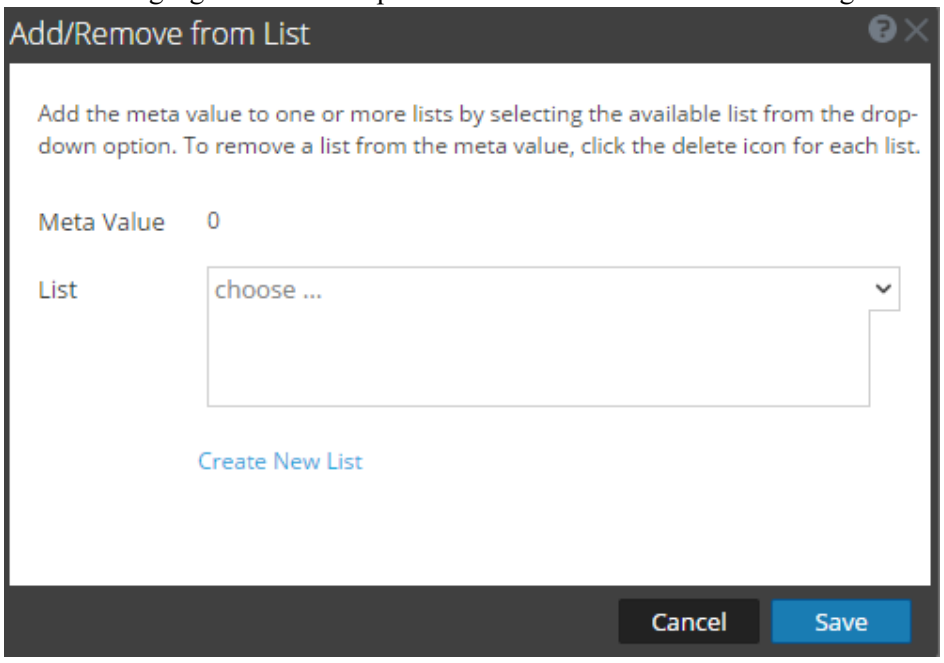
The following table shows the options in the Add/Remove from List dialog.

Option	Description
META VALUE	Displays the selected entity or meta value that needs to be added to or removed from one or more lists. You can also a create a new list using the selected value.

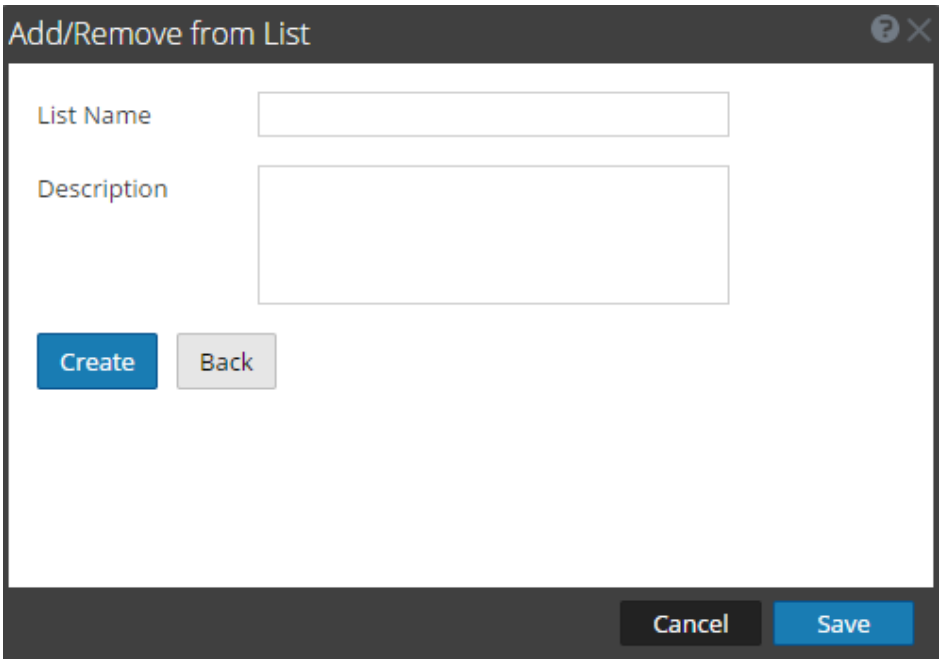
Option	Description
Create New List	Displays a dialog to create a new list using the selected meta value.
ALL	Shows all of the available Context Hub lists. The lists that contain the selected entity or meta value are selected. Select a checkbox to add an entity or meta value to a list. Clear a checkbox to remove it from the list.
SELECTED	Shows only the lists that contain the selected entity or meta value. (All lists are selected.)
UNSELECTED	Shows only the lists that do not contain the selected entity or meta value. (All lists are unselected.)
Filter Results	Enter the name or description of a specific list to search from multiple lists.
LIST	Displays the name of all the lists.
DESCRIPTION	Displays information about the selected list. The description that you provide when creating a list appears in this dialog. For example: This list contains all of the blacklisted IP addresses.
Cancel	Cancels the operation.
Save	Saves the changes.

Quick Look in the Navigate and Legacy Events Views

The following figure is an example of the Add/Remove from List dialog when initially opened.



The following figure shows the dialog when you select Create New List.



The following table describes the features of the Add/Remove from List and Create New List dialogs.

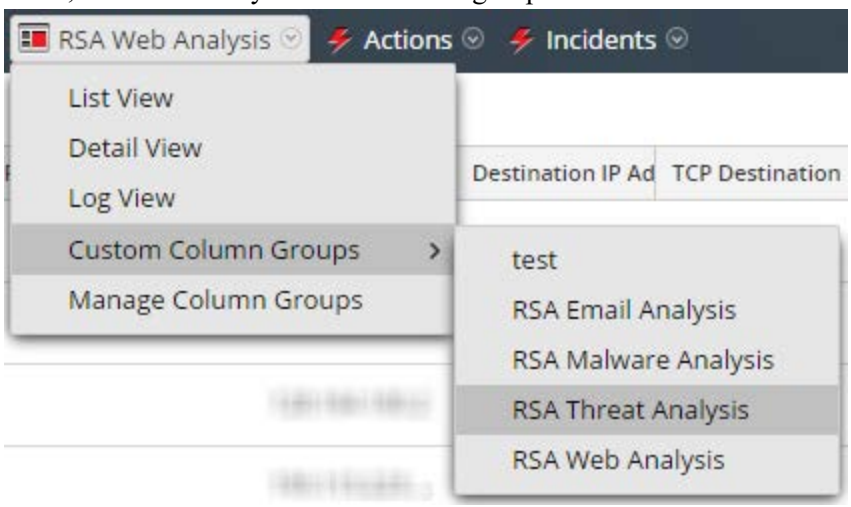
Feature	Description
Meta Value	The selected meta value to be added to the existing or new list.
List	The list to which the selected meta value must be added. A drop-down menu provides a list of available lists to which you can add the meta value.
Create New List	Opens a new dialog in which you can create a new list for the selected meta value.
List Name	The name of the new list.
Description	The description of the new list.
Create	Creates a new list after entering the required fields.
Back	In the new list mode, cancels the new list creation and returns to the original dialog.
Cancel	Cancels the addition of the meta value to a list and closes the dialog.
Save	Saves the changes made to the lists and closes the dialog.

Column Groups Dialogs

Column groups allow you to format the events list to include only the relevant meta keys in the Events view and Legacy Events view (see Use Columns and Column Groups in the Events List). When the events list in Investigate is populated with events, each column lists the values returned for a meta key. Changing the meta keys displayed in the events list is a useful method of narrowing the focus of your investigation. For example, the default column group includes columns for **Collection Time**, **Type**, **Theme**, **Size**, and **Summary**. These are just the basic information, not specialized in any way. The NetWitness Email Analysis list has only that contain information useful when investigating email.

The column group definition includes the meta keys to use as column titles, the position of the column in the list, and the default width of the column. You can add, delete, import, export, and edit column groups. At fresh installation, built-in column groups are available. The built-in column groups are prefixed with NetWitness and can be duplicated but cannot be edited or deleted. You can also create custom column groups.

- The Create Column Group dialog is for the 11.4 and later Events view. To access this dialog, select **Column Group > New Column Group** in the **Events** view toolbar.
- The Column Group Details dialog is for the 11.4 and later Events view. To access this dialog, select **Column Group** in the **Events** view toolbar, then click the edit icon ( ) next to a custom column group name.
- The Manage Column Groups dialog is for the Legacy Events view (Version 11.4) , and the Events view (Version 11.4 and earlier). The Manage Column Groups dialog has these features that are not yet available in the Create Column Group dialog: set column width, import, and export. To access this dialog, go to **Investigate > Legacy Events** and in the **View** drop-down list select **Manage Column Groups**. The **View** option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.



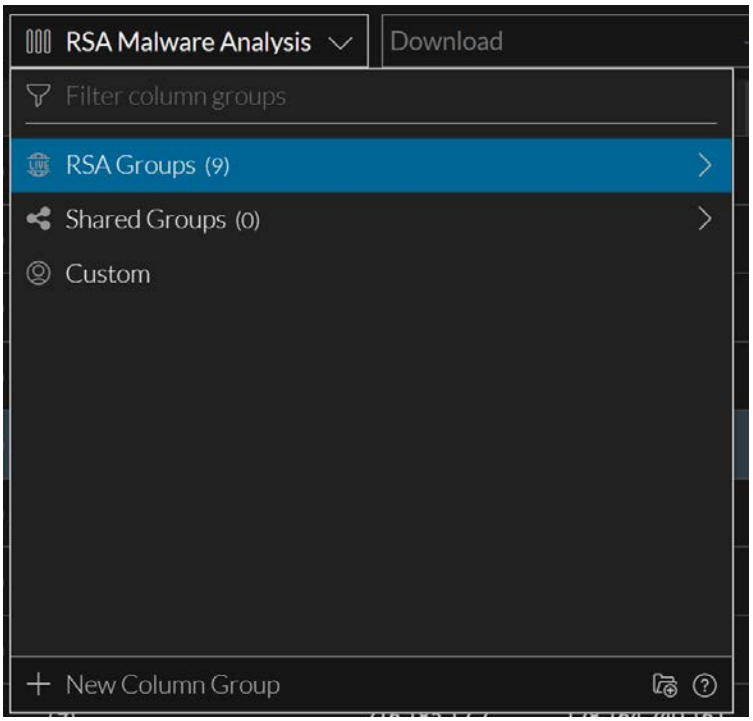
After column groups are defined, you can use them in other Investigate views. In the Navigate view, Query Profiles allow you to select a column group to use when the profile is applied. In the Events view and the Legacy Events view, you can select a column group to apply to the Events panel.

Related Topics

- How NetWitness Investigate Works
- Events View
- Legacy Events View

Quick Look - Column Group Menu, Create Column Group Dialog, and Column Group Details Dialog

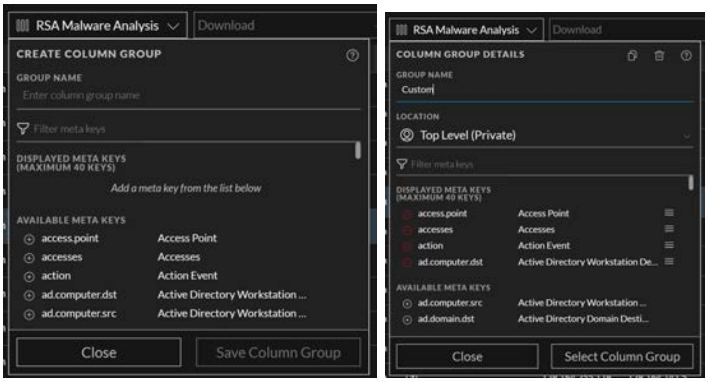
This section introduces the Column Group Menu, Create Column dialog, and the Column Group Details dialog. The following figure is example of the Column Group menu. The table describes the options.



Feature	Description
Visibility Options	(Version 11.5 and later) Control the types of column groups that are visible in the list, using any combination of the visibility options (blue = selected, black = not selected): **Private** = display private groups that only you can manage **Shared** = display shared groups that anyone in your organization can manage **RSA** = display built-in groups that only RSA can manage The visibility options work together with the Filter Column Groups field. If the visibility option is hiding built-in groups (which include "RSA" in the group name) and you search for a name that contains "RSA," the list is empty.

Feature	Description
Filter Column Groups	Filters the list of column groups as you type text so that only group names that contain that text are displayed.
Column Group List	The list of column groups consists of custom and built-in groups, which are distinguished by the icons that precede the name. In Version 11.5 and later, custom groups can be shared or private. The RSA column groups are built-in column groups. The icons distinguish the private custom groups, shared custom groups, and built-in groups.
New Column Group	Displays the Create Column Group dialog, where you can create a custom column group.

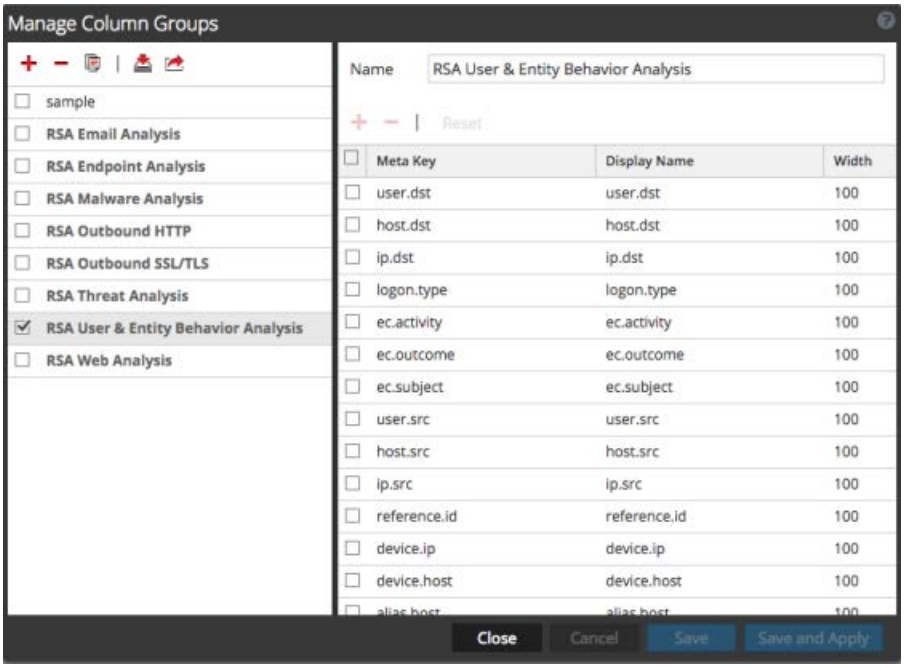
The Create Column Group dialog, shown in the figure on the left, allows you to define a custom column group. The figure on the right illustrates the Column Group Details dialog, in which you can edit a custom column group. The table describes the fields and options in the dialogs.



Feature	Description
	Deletes the custom column group in the Column Group Details dialog. This action is irreversible and applies globally; the column group is no longer available to anyone who is using it on this service.
Group Name	Displays the name of the column group. The name must be unique and contain fewer than 64 characters. You can type in this field to edit the name in a custom column group.
Sharing	In Version 11.5 and later, you can create column groups that are shared or private. This setting is available when you first create the group. After it is created, you cannot change a shared column group to private, or a private column group to shared.
Filter Meta Keys	Filters the Displayed Meta Keys and Available Meta Keys based on the text that you type. Only meta keys that contain the typed text are displayed.
Displayed Meta Keys	Displays a scrollable list of meta keys that are selected for use in the custom column group. You can add meta keys in the Available Meta Keys list to this list, remove meta keys from this list , and drag meta keys up or down to change the order in this list ().

Feature	Description
Available Meta Keys	Displays a scrollable list of meta keys that are available (on the service) for use in the custom column group. You can add them to the Available Meta Keys list. Clicking ⊕ next to the meta key name adds it to the Displayed Meta Keys list.
Close button	Closes the dialog.
Save Column Group	For the Create Column Group dialog only, saves the new column group.
Reset	For the Column Group Details dialog only, reverts the edited column group to the last saved state.
Update Column Group	For the Column Group Details dialog only, applies changes to an edited column group.
Select Column Group	Applies the column group.

Quick Look - Manage Column Groups Dialog



The Manage Column Groups dialog has two panels: Groups and Settings. At the bottom of this dialog are four buttons: Close, Cancel, Save, and Save and Apply.

The left panel is the Groups panel. This is where you can add, delete, import, or export column groups. At the top of the panel is a toolbar. Below the toolbar is a list of added column groups, where you can select one or more groups.

The following table lists the actions in the toolbar.

Action	Description
	Adds a column group. Clicking this button highlights the Settings panel on the right, where you can name the column group and add or delete meta keys. At least one meta key is required to add a group.
	Deletes a column group. A confirmation dialog is displayed before the selected group is deleted. OOTB column groups cannot be deleted.
	Creates a copy of the selected column group.
	Displays the Import Column Groups dialog, where you can select a file to upload.
	Exports one or more selected groups to your local file system.

The right panel is the Settings panel. This is where you can create and edit column groups. This panel contains the Name field, a toolbar, and a list. The following table describes the features of the Settings panel.

Feature	Description
Name	The name of the selected column group.
	Adds a new row to the list of meta keys, where you can open a drop-down menu to select a new meta key.
	Deletes one or more selected meta keys. Displays a confirmation dialog before deleting.
Reset	Returns the column group to its most recently saved settings.
Meta Key	Lists the meta keys added to the selected column group.
Display Name	Lists the names of the meta keys as they are displayed in the Navigate, Events, and Event Analysis views.
Width	Specifies the width of each meta key's column. The width can be set between **10** and **1000**. The default width is **100**.

The following table provides descriptions of the action buttons.

Feature	Description
Close	Closes the dialog without saving.
Cancel	Cancels all unsaved changes.

Feature	Description
Save	Saves all changes without closing the dialog.
Save and Apply	Saves and applies all changes immediately, closing the dialog.

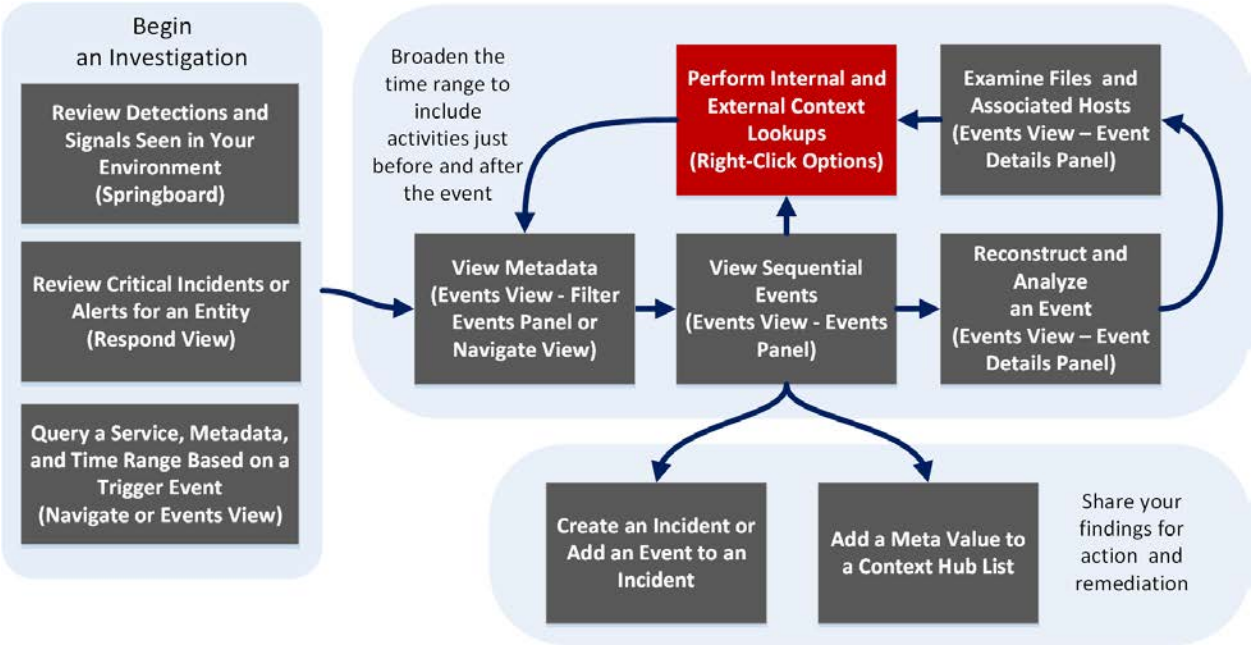
Context Lookup Panel

After an administrator configures the Context Hub service, you can view the contextual information for the meta values in the Navigate view, Legacy Events view, and Events view (Version 11.2 and later). The Context Hub service is pre-configured with a default meta type and meta key mapping. For information about mapping of the context hub meta value with investigation meta key, see "Manage Meta Type and Meta Key Mapping" in the *Context Hub Configuration Guide*.

The Context Lookup panel is displayed on the right side of the Navigate view, Legacy Events view, or Events view. Meta values that have been added to a Context Hub list are highlighted in gray in the Navigate view or Legacy Events view results. In the Events view, they are marked by an underscore. When you right-click a highlighted value and select **Context Lookup** in the resulting context menu, the lookup results are displayed in the Context Lookup panel for configured sources for the selected meta value. You can select a source in the Context Lookup panel icon bar to view the contextual information.

There are some differences between the appearance and contents of the Context Lookup panel when open in the Navigate view or Events view and when open in the Events view.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*

User Role	I want to ...	Show me how
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

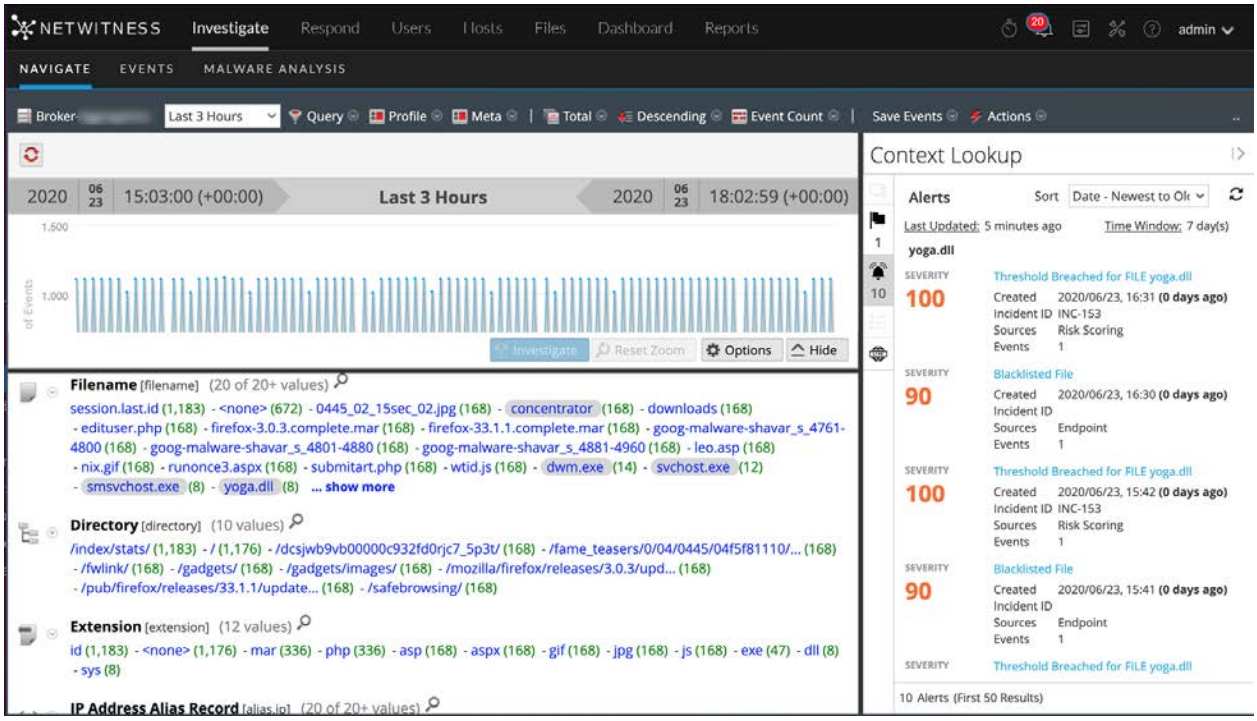
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Legacy Events View
- Navigate View
- Events View
- "NetWitness Feedback and Data Sharing" in the *Live Services Management Guide*

Quick Look (in the Navigate and Legacy Events Views)

The following figure is an example of the Context Lookup panel as it appears in the Navigate view. Controls and features are described in the table.



Feature	Description
Source Options Bar	Displays the icons for the available sources: Endpoint, Incidents, Alerts, and Lists.
Source Name	Displays the source name based on the selected icon: • Endpoint • Incidents • Alerts • Lists
Sort	Provides a drop-down of sort options for the listed context information. Possible sort options are Severity - High to Low, Severity Low to High, Date - Oldest to Newest. and Date - Newest to Oldest. The sorting options vary by source type.
	Refreshes the lookup results.
<n items> (First <n> Results)	The footer provides a count of results currently displayed and the total number of results. For example, 5 Alerts (First 50 Results).

Incidents

Incidents are displayed based on time first (Newest to Oldest) and then priority status. The following information is displayed for incident lookups:

- Incident Name and ID
- Priority status of the incidents
- Risk Score value of the incidents
- Date when the incident was created
- Status of the incident
- Assignee for the incident
- Last Updated: Indicates when contextual data was last fetched from data source and updated to cache.
- Time window: This is based on the value that is set for the "Query Last (Days)" field in the Configure Respond window. For details, see the "Configure Respond as a Data Source" topic in the *Context Hub Configuration Guide*.
- Sort: This drop-down field provides options to change the sorting of result based on time or priority.

Alerts

Alerts are displayed based on the Severity. ;The following information is displayed for alert lookups:

- Alert Name
- Severity value of the alerts
- Date when the alert was created
- Incident ID: This is the ID of the incident that the alert is associated with (If any).
- Sources: Event source name
- Number of events associated with the alert.
- Last Updated: Indicates when contextual data was last fetched from data source and updated to cache.
- Time window: This is based on the value that is set for the "Query Last (Days)" field in the Configure Respond window. For details, see the "Configure Respond as a Data Source" topic in the *Context Hub Configuration Guide*
- Sort: This drop-down field provides option to change the sorting of result based on time or priority.

Lists

The following information is displayed for list lookups.

- List Name
- Owner who created the list

- Created Date
- Last Updated Date
- Description of the list

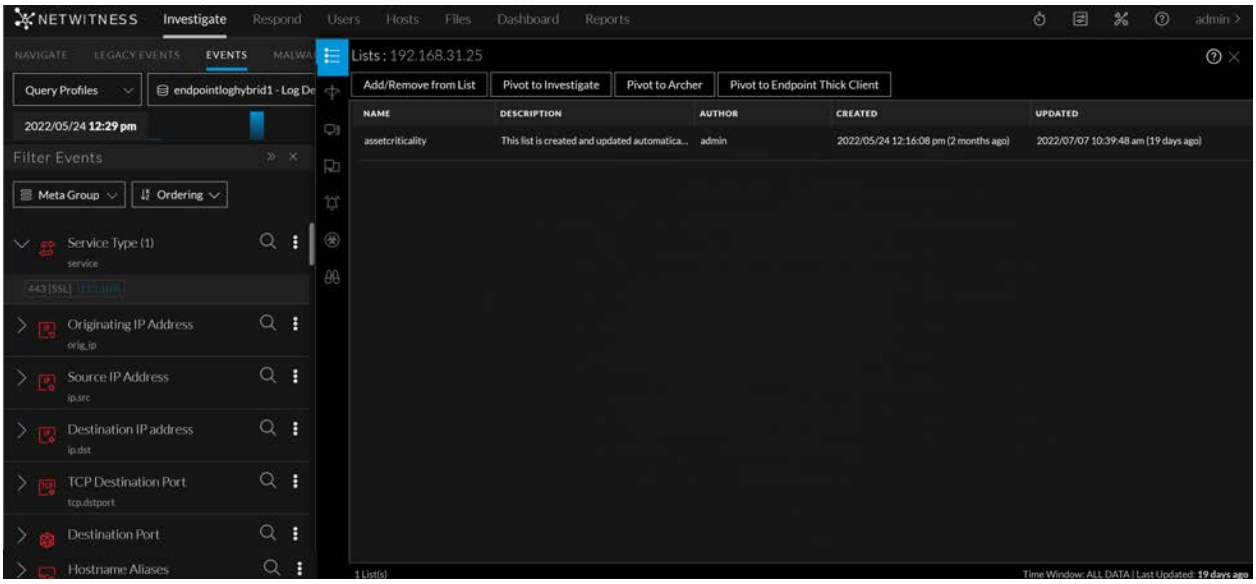
Endpoint

The following information is displayed for Endpoint lookups.

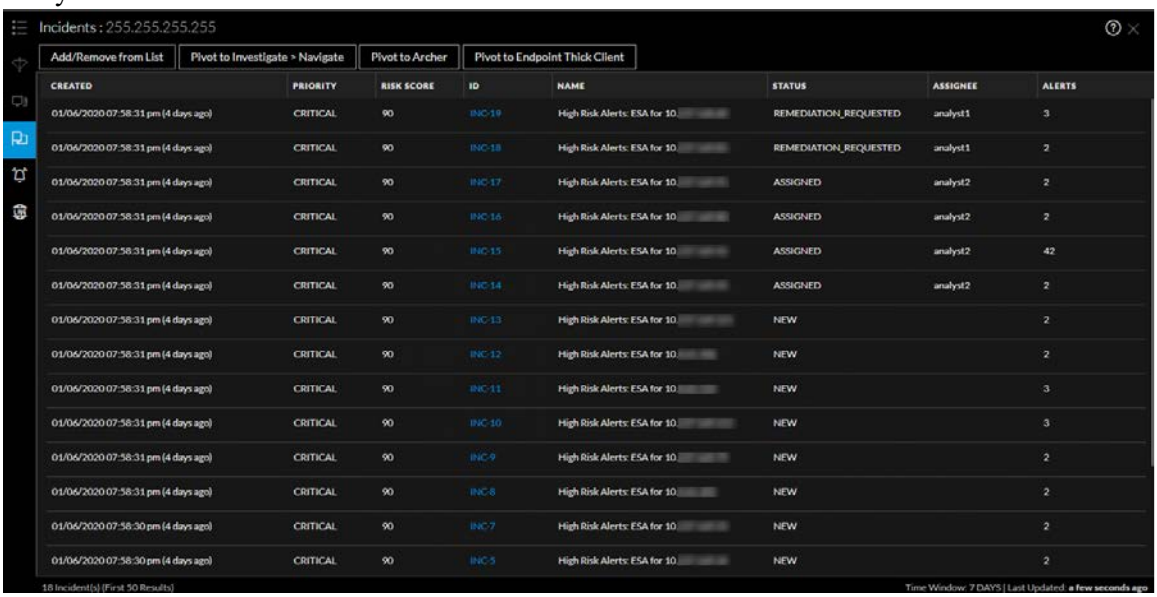
- Machine name and IP address of the machine.
  By clicking on the IP or Endpoint machine name, you will be navigated to Endpoint UI to perform further investigation.
- Last Updated: Indicates when contextual data was last fetched from data source and updated to cache.
- Machine Score: A machine IIOC score is aggregated based on the module scores.
- Number of modules: Number of active files for the selected machine.
- Last Updated: Indicates when the scan results were last updated in Endpoint database.
- Last Login User
- Machine MAC Address
- Operating System Version
- Admin Notes (if any)
- Admin Status (if any)
- Top Suspicious Modules (Modules that have an IIOC score > 500). This is based on the value set for "Minimum IIOC Score" field in the Configure Endpoint window. The default value for "Minimum IIOC Score" is 500.
- Machine IIOC Levels

Quick Look in the Events View

The following figure is an example of the Context Lookup panel as it appears in the Events view.



The contextual information or query results displayed in the Context Lookup panel depends on the selected entity and the associated data sources. The Context Lookup panel has separate tabs for each of the data sources. The tabs are: List data source, Archer, Active Directory, Endpoint, Incidents, Alerts, Live Connect, and REST API. The following figure shows the Context Lookup panel for a selected entity in the Incident Details view.



The following table describes the data available on each tab and the supported entities.

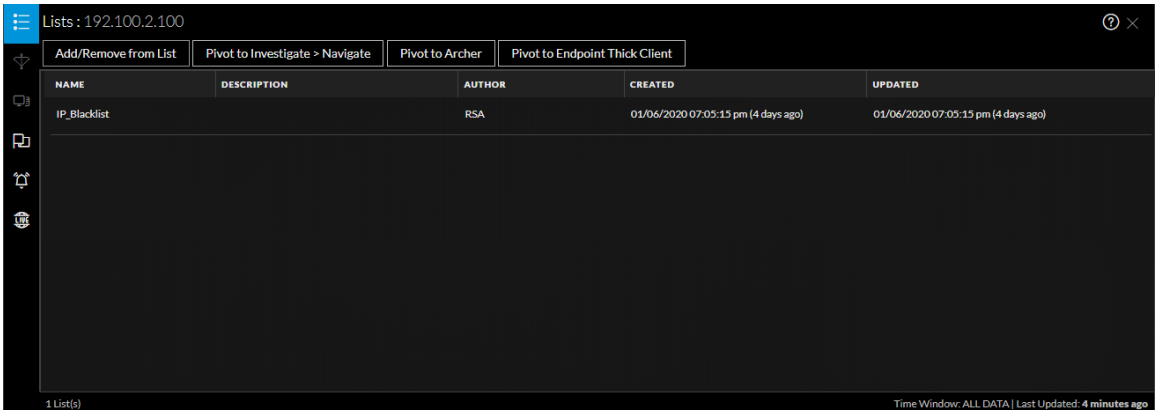
Tab	Description	Supported Entities
 (Lists)	Displays all of the list data associated with the selected entity or meta value. The result is sorted by the last updated list.	All entities

Tab	Description	Supported Entities
 (Archer)	Displays asset information along with criticality ratings using the Archer data source.	IP, Host, and Mac
 (Active Directory)	Displays all user information for the selected user.	User
 (NetWitness Endpoint)	Displays the NetWitness Endpoint data source information for the selected entity or meta value, which includes the Machines, Modules, and IIOC levels. Modules are by highest IOC score to lowest IIOC score and IIOC levels are sorted by highest IOC levels to lowest IOC levels.	IP, MAC address, and Host
 (Incidents)	Displays the list of incidents associated with the selected entity or meta value. The result is sorted by newest incidents to oldest incidents.	All entities
 (Alerts)	Displays the list of alerts associated with the selected entity or meta value. The result is sorted by newest alerts to oldest alerts.	All entities
 (Live Connect)	Displays information related to Live Connect.	IP, Domain, and Filehash
 (File Reputation)	Displays file reputation status for Filehash entities.	Filehash entities
 TI	Displays information for STIX data sources.	IP address, email address, domain, filename, URL's, and file hash. **Note:** The context lookup for email address and URL will be displayed only if these metas are mapped. Navigate to  (Admin) > System > Investigation > Context Lookup.

Tab	Description	Supported Entities
REST API	Displays the list of REST APIs (enabled in Context Hub) associated with selected the entity.	All entities

Lists Tab

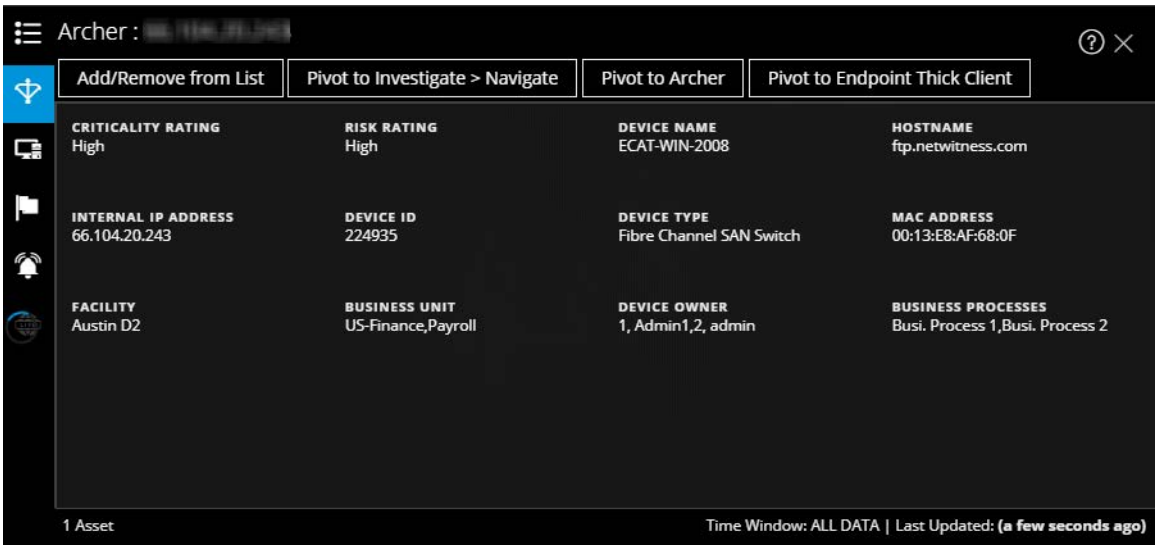
The Context Lookup panel for Lists shows one or more lists associated with the selected entity or meta value. The following figure is an example of the Context Panel for Lists, and the table describes the fields.



Field	Description
Name	The name of the list (defined while creating the list).
Description	The description of the list (defined while creating the list).
Author	The owner who created the list.
Created	The date when the list was created.
Updated	The date when the list was last updated or modifed.
Count	The number of lists in which the selected entity or meta value is available.
Time Window	The time window based on the value set for the "Query Last" field in the Configure Responses dialog. By default, all Lists data is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

Archer Tab

The Context Lookup panel for Archer displays asset information along with criticality ratings using the Archer data source for IP, Host, and Mac entities. The following figure is an example of the Context Lookup panel for Archer, and the table describes each field.



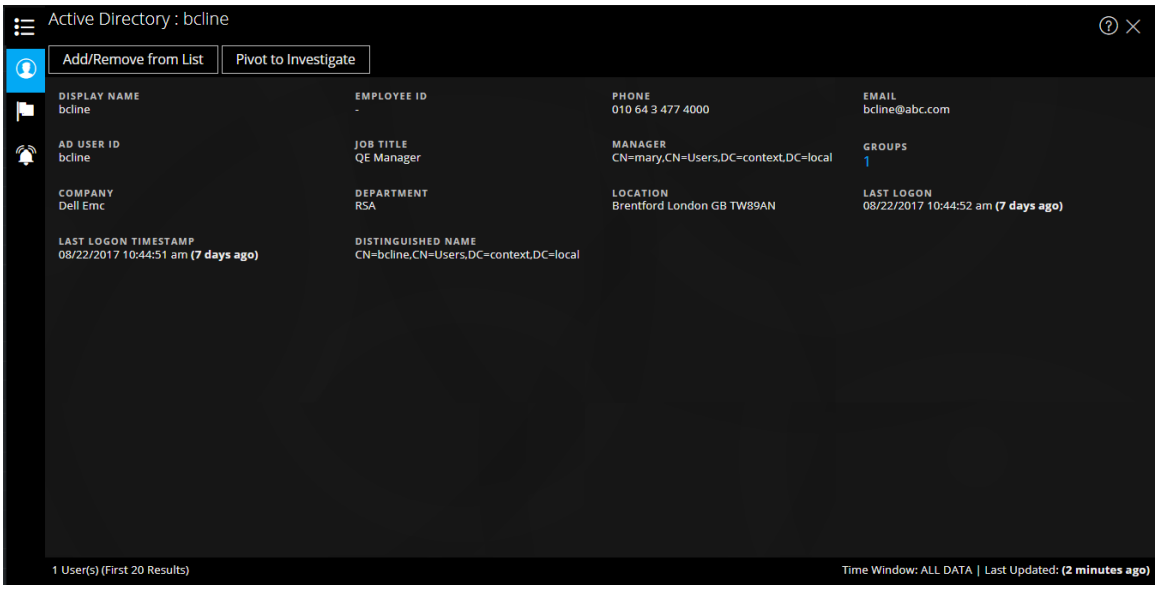
Field	Description
Criticality Rating	The device operational criticality based on the applications it supports. The criticality ratings can be set as Not Rated, Low, Medium-Low, Medium, Medium-High, or High.
Risk Rating	The calculated risk rating for the device based on the most recent assessment and the average risk rating of facilities using the device. The risk rating can be set as Severe, High, Medium, Low, or Minimal.
Device Name	The unique name of the device.
Host Name	The host name of the device.
IP Address	The primary internal IP address of the device.
Device ID	The automatically populated value that uniquely identifies the record across all applications within the system.
Type	The device type, for example, server, laptop, desktop, and others.
Facilities	Links to records in the Facilities application that are related to this device.
Business Unit	Links to records in the Business Unit application that are related to this device. For more than three business unit values, you can hover over the field to view the values.
Device Owner	The person who is responsible for the device and receives read and update rights of the record.

Field	Description
Count	The number of assets available.
Time Window	The time window based on the value that is set for the "Query Last" field in the Configure Responses dialog. By default, all data for Archer is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

Note: In the localized versions, only these twelve fields are displayed: Criticality Rating, Risk Rating, Device Owner, Business Unit, Host Name, MAC Address, Facilities, IP Address, Type, Device ID, Device Name, and Business Processes.

Active Directory Tab

The following figure is an example of a Context Lookup panel for Active Directory.



The Context Lookup panel for Active Directory displays all the related information, incidents, and alerts for a user. You can perform a look up using the following formats:

- userPrincipalName
- Domain\UserName
- sAMAccountName

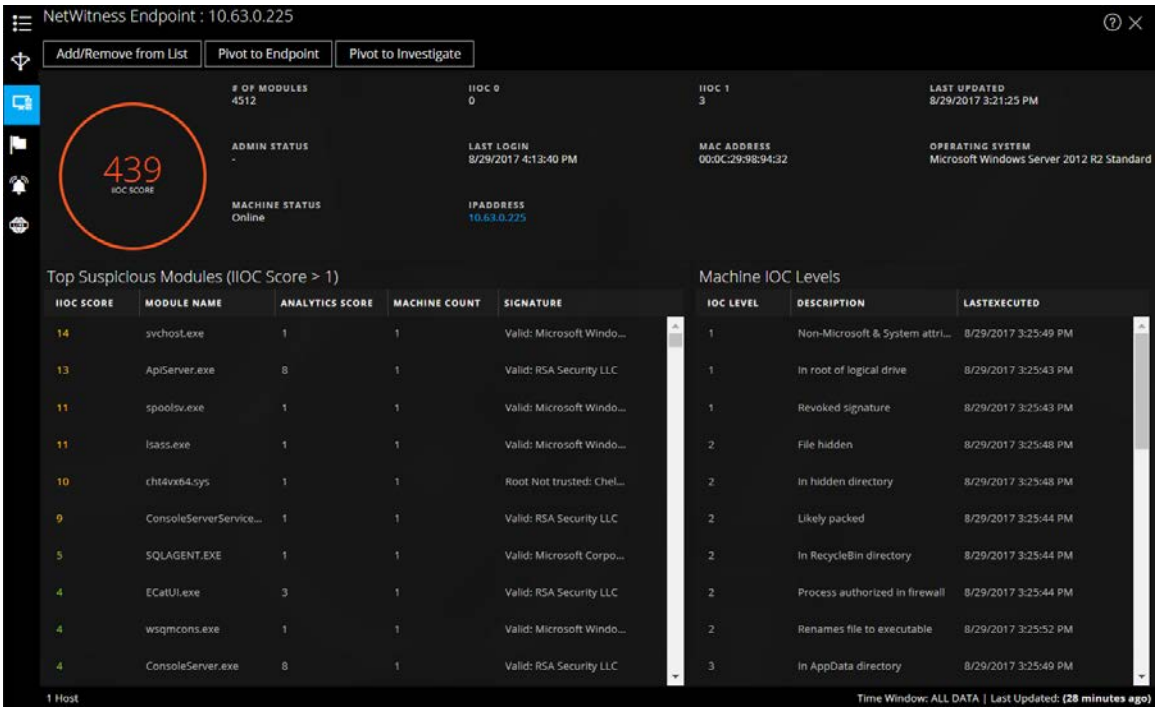
The following information is displayed for Active Directory.

Field	Description
Display Name	The name of the user.
Employee ID	The employee ID of the user.

Field	Description
Phone	The phone number of the user.
Email	The email ID of the user.
AD User ID	The unique identification of the user within an organization.
Job Title	The designation of the user.
Manager	The name of the user's manager.
Groups	The list of groups of which the user is a member.
Company	The name of the user's company.
Department	The department name to which the user belongs within the organization.
Location	The location of the user.
Last Logon	The time when the user logged into the system, only if the Global Catalogue is defined.
Last Logon TimeStamp	The time when the user logged into the system.
Distinguished Name	The unique name assigned to the user.
Count	The number of users.
Time Window	The time window based on the value that is set for the "Query Last" field in the Configure Data Source Settings dialog. By default, all data for Active Directory is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

NetWitness Endpoint Tab

The following figure is an example of the Context Lookup panel for NetWitness Endpoint.



The following information displayed for IIOCs.

Field	Description
# Of Modules	The number modules that are looked up.
Admin Status	The admin status (if any).
Last Updated	The time when the data was last refreshed.
Last Login	The time when the user last logged in.
MAC Address	The Machine MAC Address.
Operating System	The Version of the Operating System used by the NetWitness Endpoint machine.
Machine Status	The state of the module being viewed: Online, Offline, Active, or Inactive.
IP Address	The IP address of the specific module.

The following information is displayed for modules.

Field	Description
IIOC Score	A machine IIOC score is an aggregated score based on the module scores. This is based on the value set for Minimum IIOC Score field in the Context Hub Data Source Settings dialog. The default value for Minimum IIOC Score is 500. See "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.

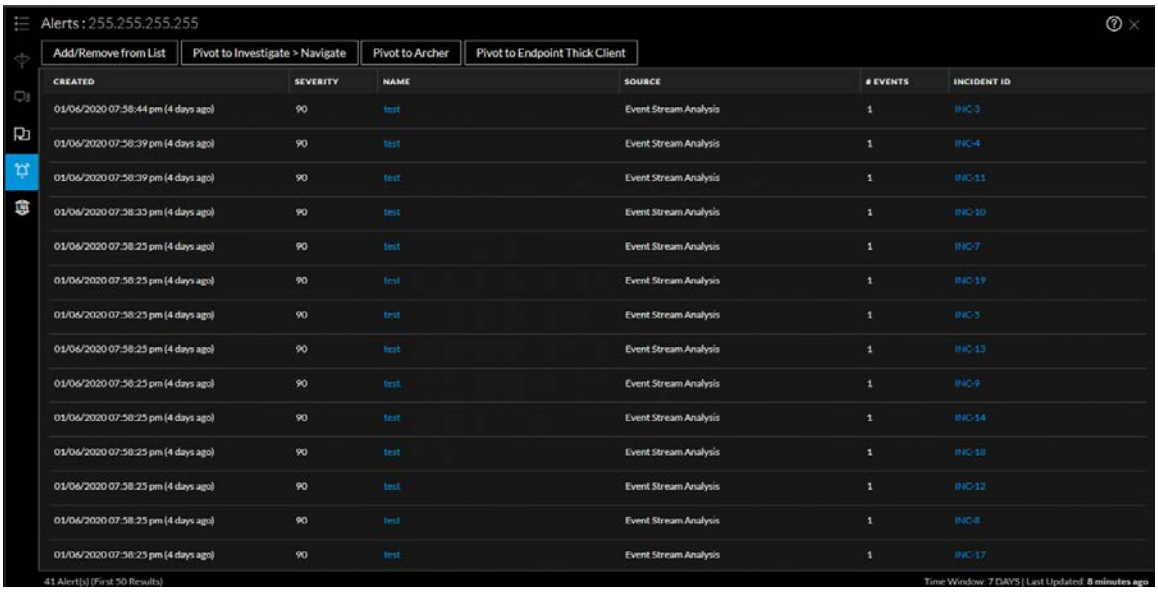
Field	Description
Module Name	The name of the module that is being looked up.
Analytic Score	The number of active files for the selected machine.
Machine Count	The number of machines on which that particular IOC got triggered.
Signature	Indicator of whether the file is signed or unsigned, valid or invalid, and signatory information. For example, Google, Apple, and so on.

The following information is displayed for machines.

Field	Description
IOC Levels	The IOC levels.
Description	The description for he IOC level if available.
Last executed	The time when the action was executed.
Count	The number of hosts that are being looked up.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, all data for NetWitness Endpoint is fetched.
Last Updated	The time when scan results were last updated in NetWitness Endpoint database.

Alerts Tab

The following figure is an example of Context Panel for Alerts that is displayed based on time first (Newest to Oldest) and then severity.

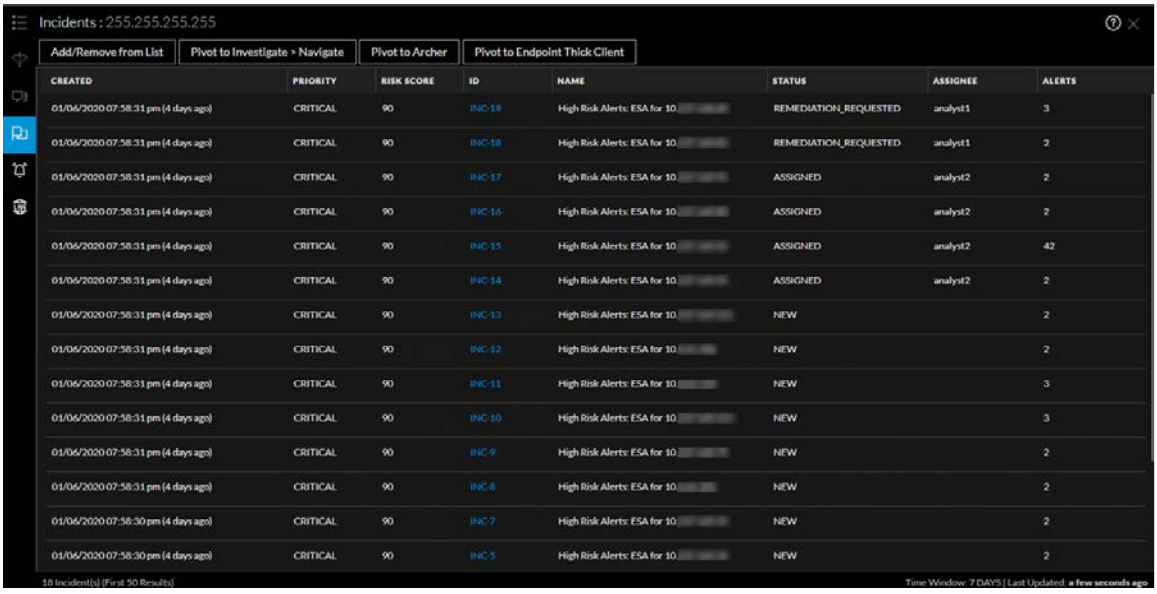


The following information is displayed in the Context Lookup panel for Alerts.

Field	Description
Created	The date and time when the alert was created.
Severity	The severity value of the alerts.
Name	The name of the alert. You can click the name to view the details of a specific alert.
Source	The alert source name from which the alert is triggered.
#Events	The number of events associated with the alert.
Incident ID	The ID of the incident (if any) with which the alert is associated. You can click the ID to view the details of a specific alert.
Count	The number of alerts. By default only the first 100 alerts are displayed. For more information on how to configure the settings, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, the alert data for last 7 days is fetched.
Last Updated	The time when contextual data was last fetched from data source.

Incidents Tab

The following figure is an example of the Context Panel for Incidents, which is based on time first (Newest to Oldest) and then priority status.

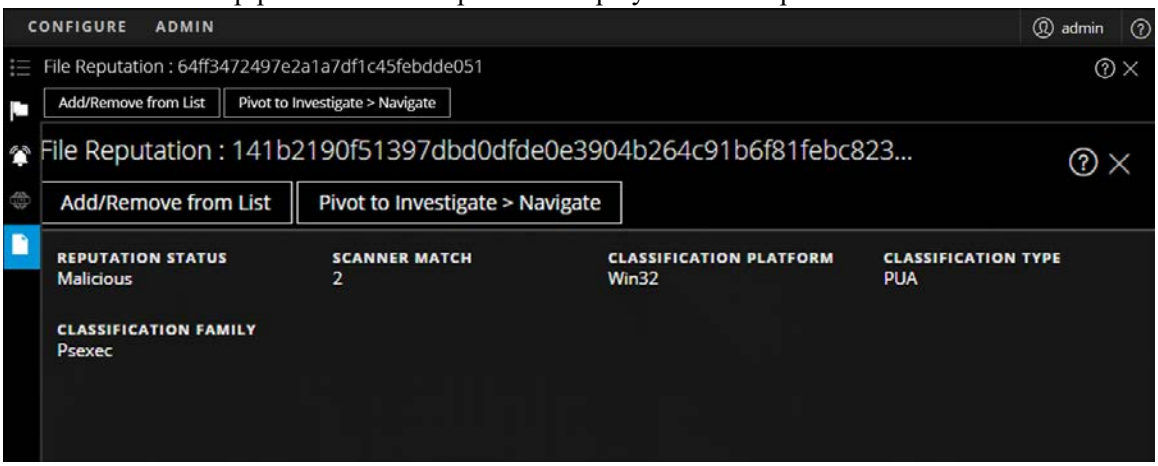


The following information is displayed in the Context Lookup panel for Incidents.

Field	Description
Created	The date when the incident was created.
Priority	The priority status of the incidents.
Risk Score	The risk score of the incidents.
ID	The Incident ID of the incident. You can click on the ID to display further details about the incident.
Name	The incident name.
Status	The status of the incident
Assignee	The current owner of the incident.
Alerts	The number of alerts associated with the incident.
Count	The number of incidents. By default only the first 100 incidents are displayed. For more information on how configure the settings, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, the alert data for last 7 days is fetched.
Last Updated	The time when contextual data was last fetched from data source.

File Reputation Tab

The Context Lookup panel for File Reputation displays the file reputation status of a file.

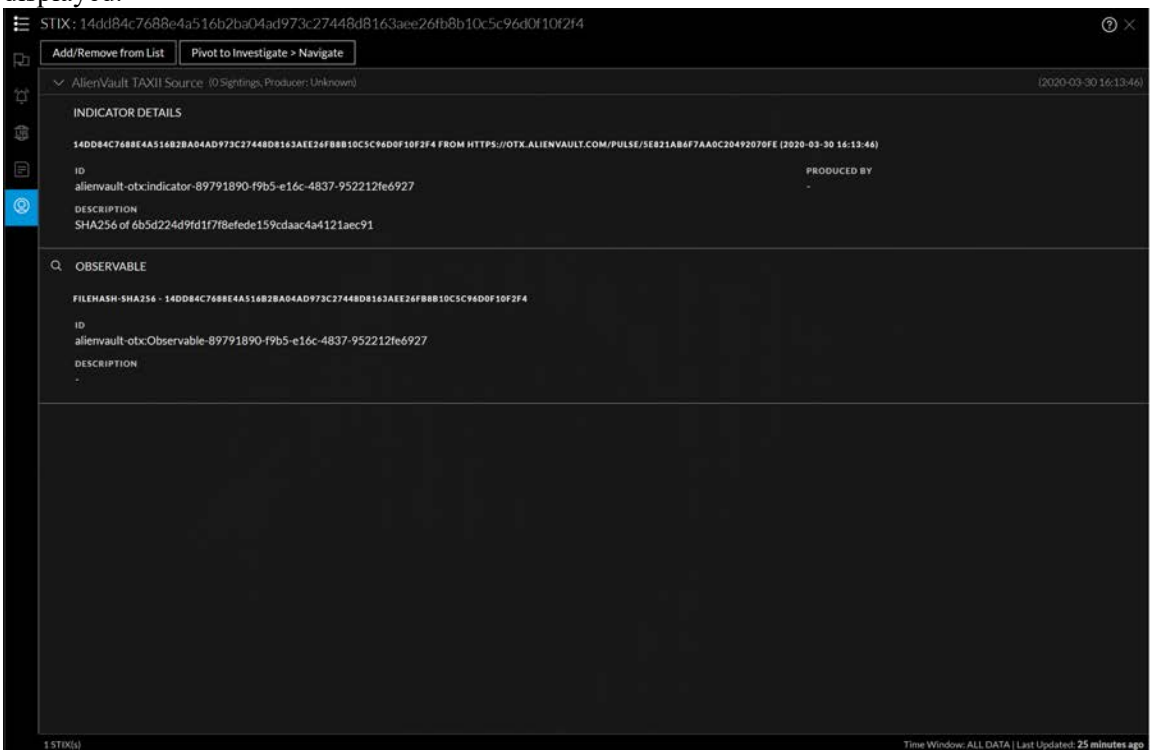


Field	Description
Reputation Status	Reputation Status of filehash. For more information about reputation status, see "View Reputation of files" in the *UEBA User Guide*.

Field	Description
Scanner Match	Number of scanners that detected malware or suspicious activity in the last scan.
Classification Platform	Classification for the queried filehash based on the platform. For example, the platform can be Win 32.
Classification Type	Classification for the queried filehash based on the type.
Classification Family	Classification for the queried filehash based on the malware family name.

TI Tab

The following figure is an example of a Context Panel for TI, and the table describes the information displayed.



Field	Description
Data Source name	Displays the STIX data source name from where the data is retrieved.
Timestamp	The time when the event was created.

Field	Description
Indicator Details	Indicator Title: Displays the details that contains a pattern that can be used to detect suspicious or malicious cyber activity. ID: Displays the ID of the selected indicator. Produced by: Displays the user role who requested for the STIX data. Description: Displays details about the selected IP address which are being watch listed.
Observable	Observable Title: Displays and conveys information about cyber security related entities such as files, systems, and networks using the STIX Cyber-observable Objects (SCOs). ID: Displays the ID of the selected observable.
(Optional) SightingsREST	Sightings Title: Displays the name of the sighting source. Confidence: Displays the criticality of the sighting. Reference: Displays the reference URL of the sighting source.

REST API Tab

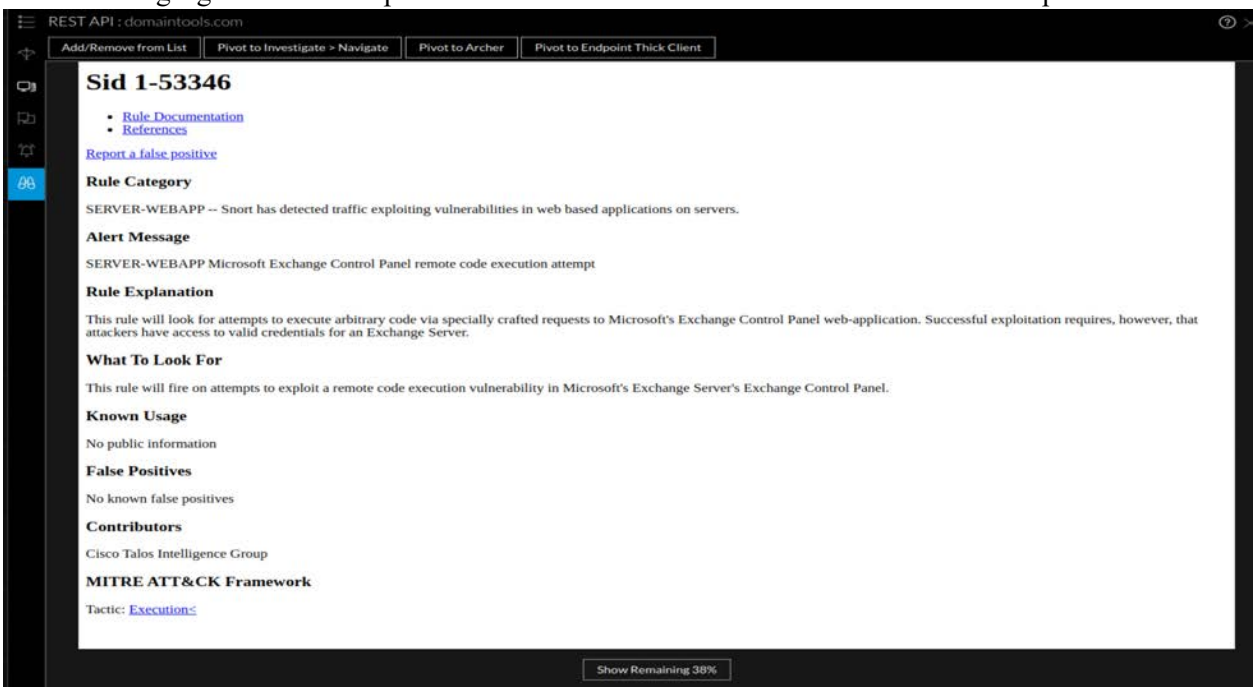
The Context Lookup panel for REST API shows HTML or JSON response (based on the response type configured) associated with the selected entity or meta value.

Note: For JSON response type, the fields that are mapped with friendly names (during REST API configuration) are only displayed for context Lookup. If you have not mapped any fields, all fields are displayed for context lookup.

The following figure is an example of the Context Panel for REST API with JSON response:



The following figure is an example of the Context Panel for REST API with HTML response:

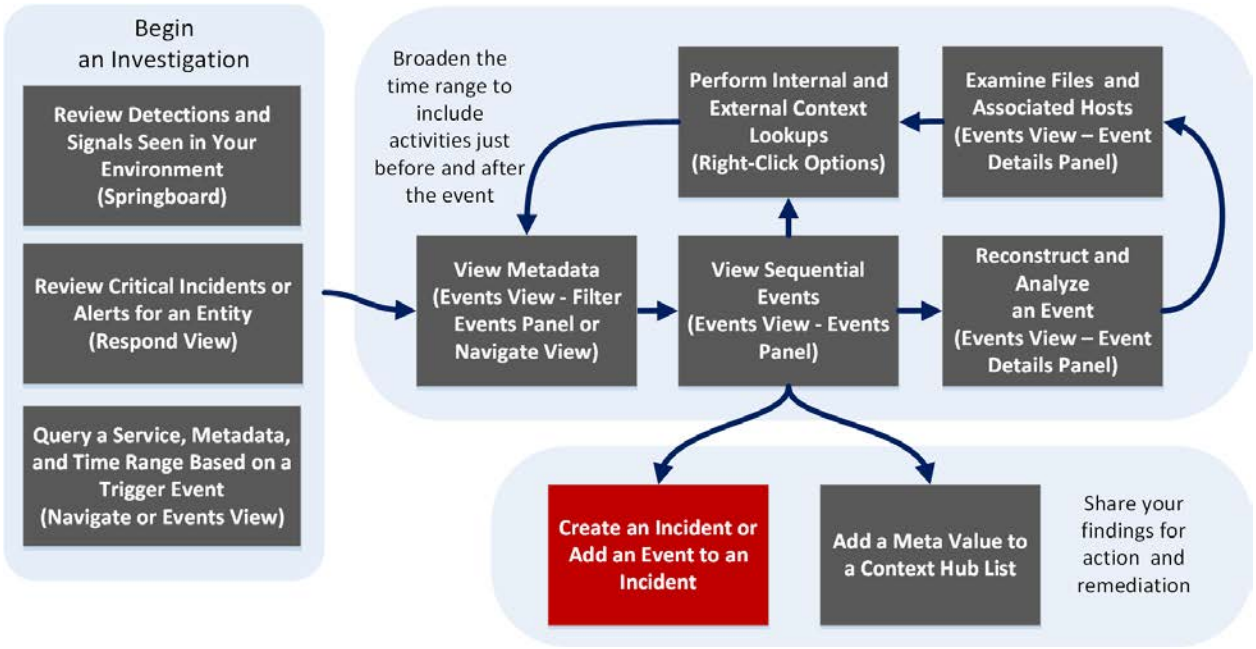


Create an Incident Dialog

In the Create an Incident dialog, analysts can create an incident from selected events in the Events view. The incident is then available to incident responders working in Respond.

To access this dialog, while investigating a service in the Investigation > Events view, select **Incidents > Create New Incident** from the toolbar.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View

User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

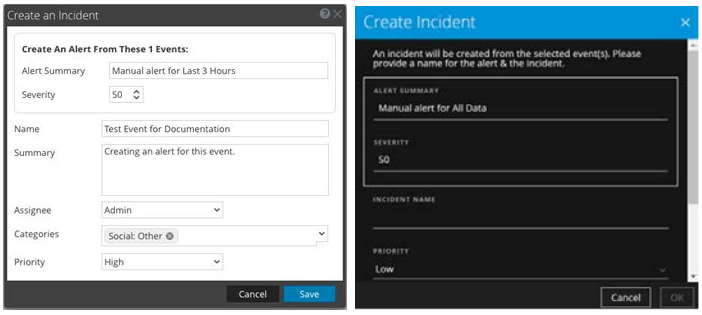
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look

The following figure is an example of the Create an Incident Dialog, and the features are described in the table.



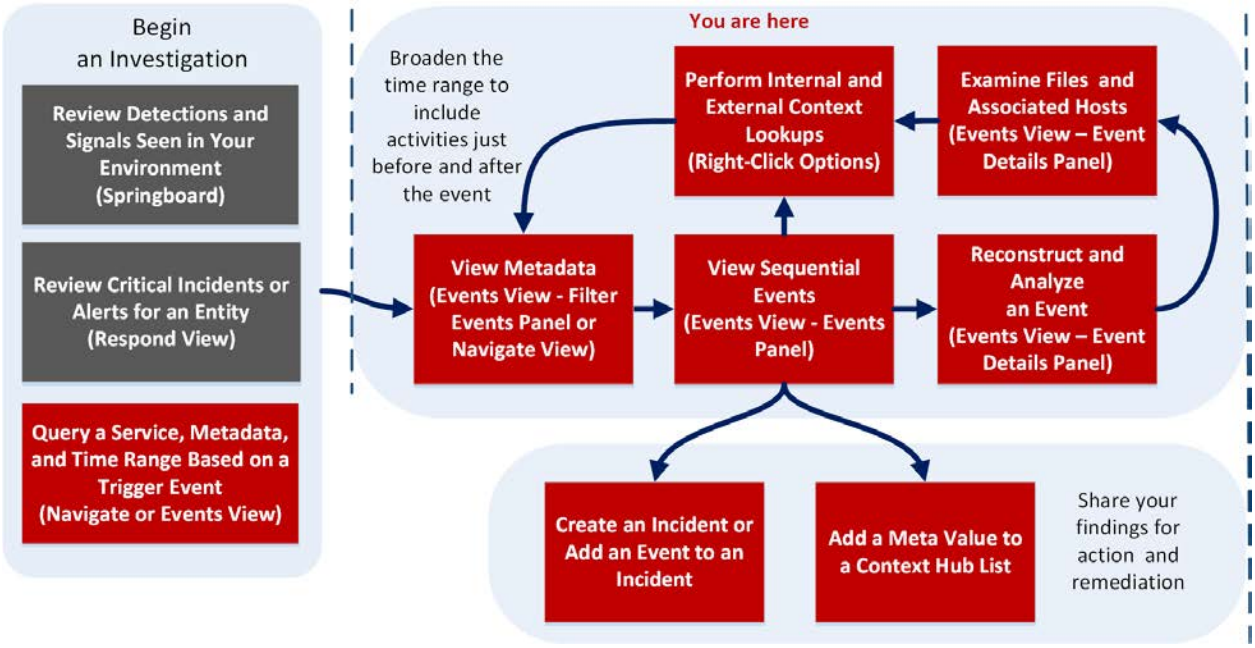
Feature	Description
Create Summary from These Events	The Alert Summary field is filled by the query that produced the select alerts, which you selected to create this incident. The Severity field reflects the Severity of the selected alert, an integer between 1 and 100.
Name	(Required) Specifies a name to identify the incident. In the example, the name is Sample Incident. You can provide a name that clearly identifies the nature of events that will be added to this incident
Summary	(Optional) Specifies a description for the incident. A good summary clearly identifies the incident for other analysts and responders.
Assignee	(Optional) Assigns the incident to a user in the SOC. Clicking Assignee opens a drop-down list showing the user names of SOC personnel who respond to incidents.
Categories	(Optional) Identifies categories of incidents. Clicking Categories, opens a drop-down list of Incident categories and subcategories. You can select one or more categories to which the incident belongs. Categories fall into these major groups: Environmental, Error, Hacking, Malware, Misuse, and Social.
Priority	Identifies the priority for the incident. Clicking Priority opens a drop-down list of priorities: Critical, High, Medium, or Low displayed in the drop-down list.
Cancel	Closes the dialog without saving changes.
Save	Saves the incident and closes the dialog. A message confirms that the incident was created successfully.

Events View

In the Events view analysts can view a sequential list of network, log and endpoint events, select an event for reconstruction and analysis, and view the raw event and metadata with interactive features that enhance the ability to see meaningful patterns in the data. In Version 11.5 and later, you can drill into metadata for the listed events. The Events view offers packet, file, host, text, log, and email reconstruction. When you open a web reconstruction of an event, the same web reconstruction used in the Legacy Events view is displayed.

Workflow

The following figure is a high-level workflow illustrating the tasks you can do in NetWitness Investigate, with the Events view tasks highlighted in red.



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*

User Role	I want to ...	Show me how
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

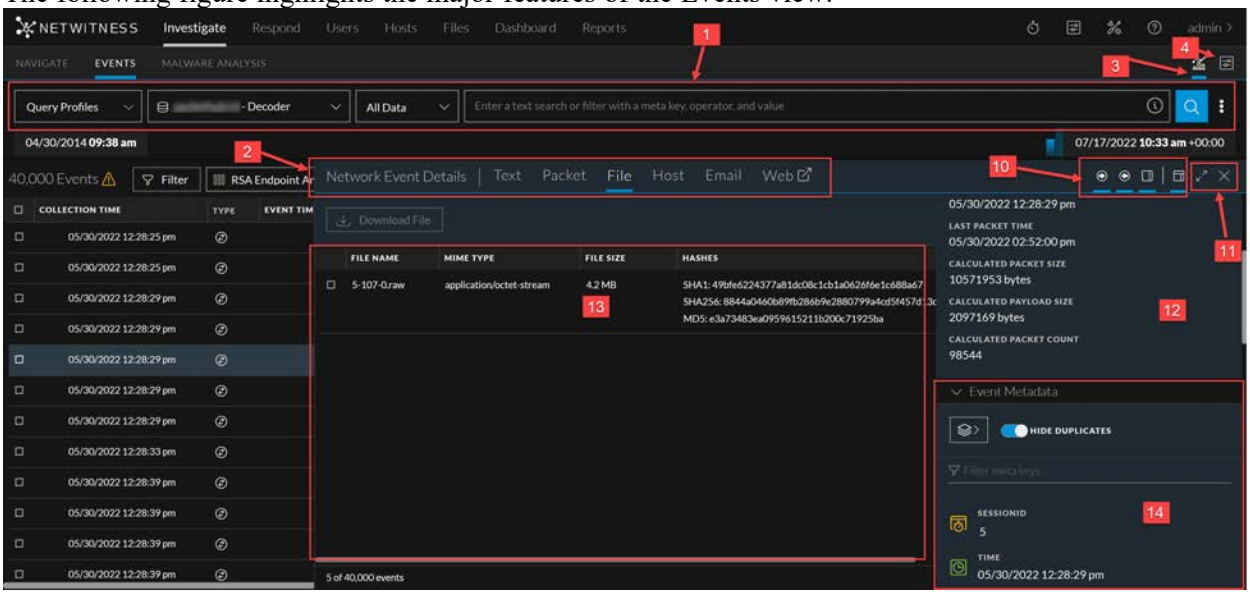
There are multiple access points to this view, which are described in Begin an Investigation in the Events View. If you access the Events view from the Respond view, you can see the analysis for a selected event in an incident. The options are a subset of the options available when you open an event from within the Investigate view. To get complete functionality and examine other events, you can go to the Event view directly (INVESTIGATE > Event).

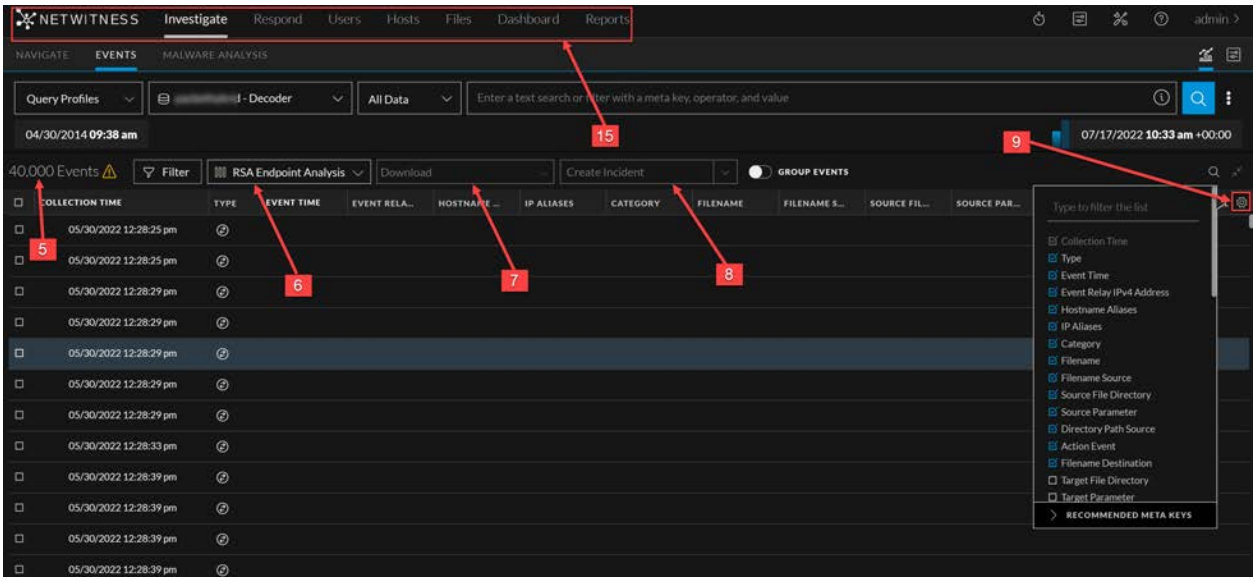
The Events view lists events in ascending order by time in the Events panel. The events displayed can be results for the drill point in the Navigate view or Legacy Events view, or results for a query entered in the Events view query bar.

Input fields for a query are displayed so that you can select a service and time range, and type an optional query. When you submit a query, the service being investigated counts the results up to a limit of 10,000 events, and 10,000 network, log, and endpoint events are loaded in the Events panel. Different columns are displayed, depending on the selected column group. You can rearrange and resize the columns, choose a built-in or custom column group, and choose individual columns that you want to see. When you find an event of interest, clicking the event opens the reconstruction in a new panel (Packet, Text , or File).

Note: For versions earlier than 11.3, the first 100 events are loaded. You can scroll through the list and click **Show Next 100 Events** at the bottom of the list. If the next page contains fewer than 100 events, the button changes to reflect the number of remaining events.

The following figure highlights the major features of the Events view.





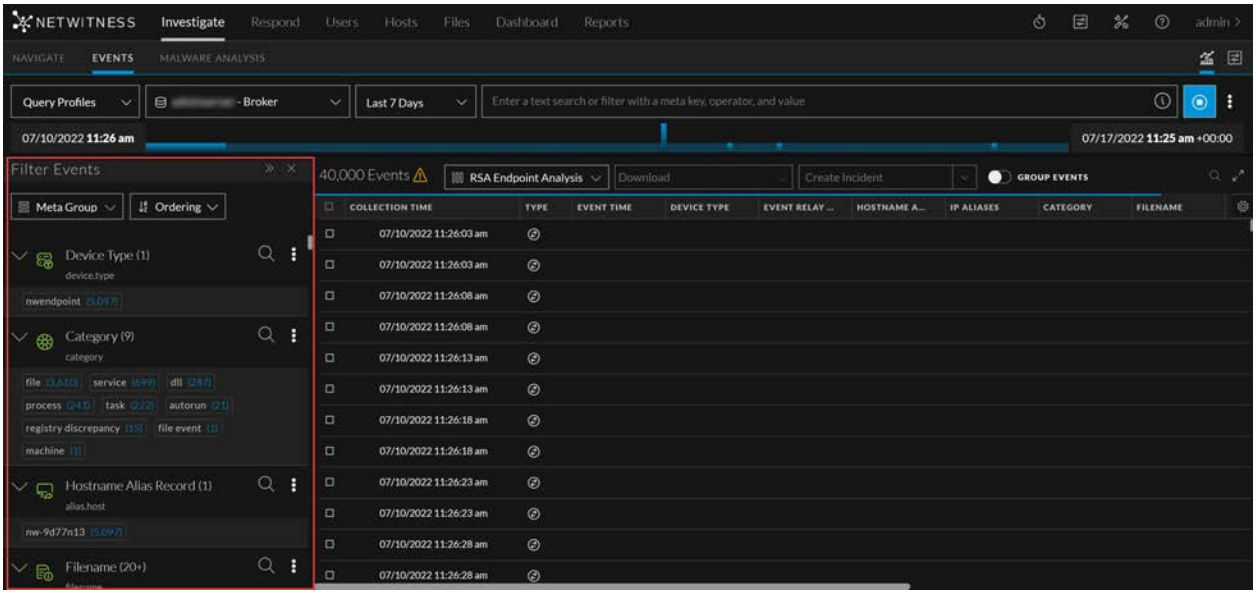
1	**Query Bar**: When a service is selected, displays the service selector, time range selector, and the queries you have entered. You can select a service as described in Begin an Investigation in the Events View and refine the query as described in Filter Results in the Events View. Clicking submits the query and sends a request to the selected service to load the data. In Version 11.3 and later, clicking the (console icon) opens the query console, where detailed status of the query is provided (see Events View below).
2	The type of event being analyzed and the type of reconstruction are reflected in the heading. • These are the event types: **Network Event Details**, **Log Event Details**, or **Endpoint Event Details**. • The types of analysis available for the event type are Text, Packet, File, Host, Email, and Web. Network events can use all types of analysis: text, packet, file, and email (Version 11.4.1 and later). Log and endpoint events use only text analysis. The email (Version 11.4.0.x and earlier) type and web type open the current event as an email or web reconstruction in the Events view. For details, see Examine Event Details in the Events View.
3	Reopens the Events panel if you have closed it. For details, see Analyze Events in the Events View.
4	Sets preferences for the Event view (see Configure the Events View).
5	The Events panel title. • In Version 11.3 and later, the Events panel title is slightly different than the title in prior versions, and a row number indicator has been added. The title lists the number of events and sort order; for example, **40,000 Events (Asc)** means that 40,000 events were found and they are listed in ascending order by time. If more than 10,000 events are found, only the oldest 10,000 events are displayed in ascending order, and an amber triangle highlights the fact that not all events were loaded. This may indicate that you need to refine the query. For more information about refining the events listed here, see Filter Results in the Events View.

	• Versions prior to 11.3 simply list the number of events found, and you can load 100 of them at a time. In Version 11.4 and later, clicking  opens the Find Text in Table dialog.
6	The Column Group drop-down lists built-in and custom column groups that you can apply to the Events panel. Built-in column groups are sometimes updated between one version and the next. Some examples of built-in column groups are Email, Endpoint Analysis, Malware Analysis, Outbound HTTP, Outbound SSL/TLS, and Summary List. Summary List is the default column group. For details, see Use Columns and Column Groups in the Events List.
7	The Download drop-down menu lists the available options for downloading event data. The options are Log, Visible Meta, and Network (see Downloading and Acting Upon Results. You can change the preferred format of the event type data in the Event Preferences dialog (see Configure the Events View).
8	The Create Incident button enables you to create incidents from events. The Add to Incident button enables you to add selected events to an open and existing incident (see Add Events to an Incident in the Events View and Add Events to an Incident in the Legacy Events View).
9	Displays the column selection settings to select the individual columns displayed in the Events panel. For details, see Use Columns and Column Groups in the Events List.
10	Controls to show or hide the Overview panel, show or hide requests and responses, and open the Event Meta panel. For details, see Analyze Events in the Events View.
11	Controls to change the size of the panel and close the panel. For details, see Analyze Events in the Events View.
12	The Overview panel provides summary information about the event you are currently analyzing. The selected event is highlighted in the Events panel with a blue background. The summary information is different for the different event types (packet, log, and endpoint). In Version 11.5, the redundant NW Service is removed.
13	The event data for the event you are currently analyzing.
14	The Event Meta panel is redesigned in Version 11.5, but has the same functions as in Version 11.4. The Event Meta panel lists the meta keys and values found in the data. This data can be sorted in two ways - Alphabets or Sequence. Some metadata are searchable; they have a binoculars icon, which you can click to see the associated data highlighted in the event data (see Analyze Events in the Events View). • For a packet, the data is called a payload and is displayed in the form of a request and response. • For a log event, the data is a line of text from the raw log. • For an endpoint event, the event data is relevant to data from the NetWitness Endpoint agents running on hosts in the network. It may be a single process, driver, DLL, file (executable), service, or autorun, and information related to logged-in users. (See the *NetWitness Endpoint User Guide* for complete information about endpoint event data.)
15	The Version 11.5 main menu for NetWitness Platform has relocated Hosts, Files, and Users (Entities) options for easier access.

Filter Events Panel

The Filter Events panel is a beta feature added from Version 11.5. Clicking the Filter button (Filter) in the Events panel, opens the panel to provide a view of meta keys and meta values found in the data set. (Version 11.6) By default, the Filter Events panel is open in the Events view. The user preference (open, closed, or fully expanded) is saved across sessions and logins. See Drill into Metadata in the Events View for more information about drilling into metadata.

Note: (Version 11.6) By default, the Filter Events panel is open in the Events view. The last used state of the panel (narrow or fully expanded) is saved throughout the session and across logins. Also, the Filter Events panel provides additional contrast between meta keys, meta values, and meta counts to improve readability.

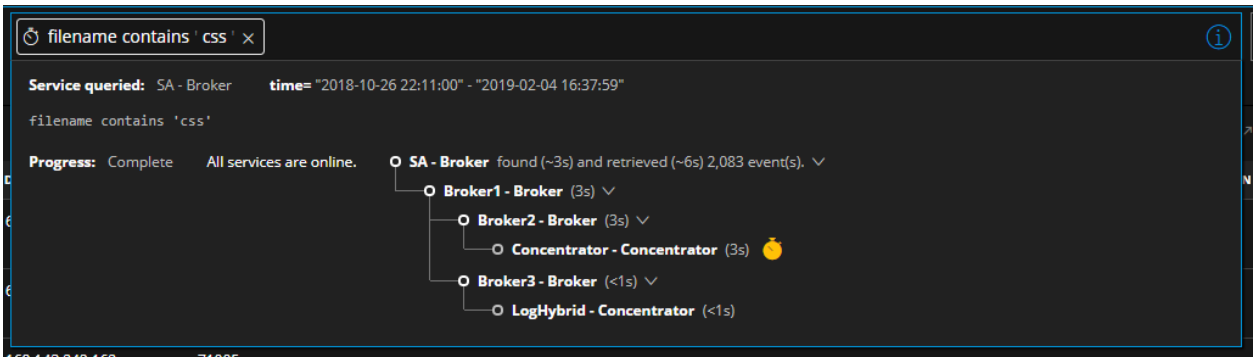


Meta Groups Menu	With the Filter Events panel open, you can select a meta group to define the meta keys displayed in the Filter Events panel.The Default Meta Keys meta group is in effect the first time you log in. If you selected a different meta group the last time you logged in, it remains in effect until browser cache is cleared. See Use Meta Groups to Focus on Relevant Meta Keys for details about meta groups.

Ordering Menu	With the Filter Events panel open, you can look at two parameters for each value: the event count or the event size. Each meta key entry includes either the event count or the event size in parentheses after the value. In both cases, there are four options for ordering: • By default, the meta keys are displayed using the Event Count > Descending by Total Count method. When showing the event count for each value, you can order by Descending by Total Count, Ascending by Total Count, Ascending by Value, and Descending by Value. • When you prefer to see the size of the event that contains the value, you can use one of the four Event Size ordering options: Descending by Total Size, Ascending by Total Size, Ascending by Value, and Descending by Value.
Meta Key options button (E)	The Meta Key options button offers actions that you can take on an individual meta key. In Version 11.5, the only action is to copy all of the visible meta values for a meta key.
Meta Key List	An icon before each meta key name identifies the indexing method for the key. The indexing method determines the types of interactions and queries possible using that meta key. • This meta key is indexed by value: Action Event [action] (40+). The green color indicates that the all available interactions and queries are supported. You can see the available interactions in the context menu by right-clicking the meta value. • This meta key is indexed by meta key: Bytes Sent [bytes.src]. The yellow color is a clue that a subset of available interactions is supported, and queries on this meta key may take longer than meta keys that are indexed by value. You can see the available interactions in the context menu by right-clicking the meta value. • This meta key is not indexed: MAC Alias Record [alias.mac]. Values for non-indexed meta keys cannot be used to query. If you want to query a meta key that is not indexed, your administrator needs to edit the index file for the service to index the meta key by value or meta key.

Query Console

Clicking (the console icon) opens the query console, where detailed status of the query is provided.



In the query console, you can see which service, time range, and metadata was queried as well as real-time information about the status of the query and the services being queried. A progress bar indicates the query's completion percentage at the bottom of the console. The statuses let you know details about what is happening; for example, you can tell when the query is executing, queued, reading the index file for the queried service, retrieving events, and complete. All statuses and non-fatal messages are displayed as they come in, and the border color changes if a non-fatal error occurs. View Status of a Query provides additional details on this subject.

Several messages that may be displayed in the query console require additional explanation.

Message: Maximum value limit (`valueMax`) of %1% reached on meta key %2% in index slice %3%

Explanation: The `valueMax` property on the specified meta key has been reached in the index being queried. An administrator configures this inside the index files available in `ADMIN > Services > [Service Name] > Files > index-[service type].xml` or `index-[service type]-custom.xml`. As an example, the statement below from the index file states the meta key called `client` has a limit of 250,000 values by default.

```
<key description="Client Application" level="IndexValues" name="client"
format="Text" valueMax="250000" />
```

Message: The query on channel %2% was auto-canceled by the system for exceeding time usage limits. Check timeout values.

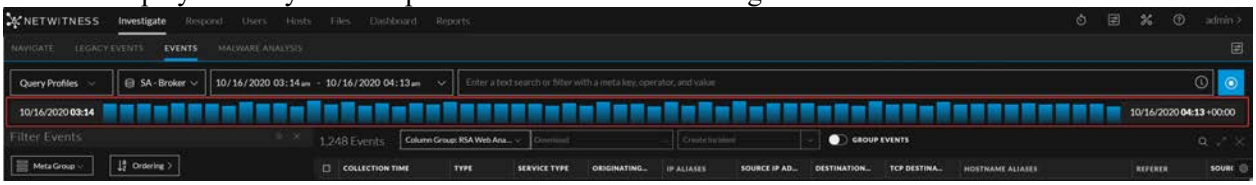
The server has a per-operation limit on execution time, and the requested operation exceeded the limit. To avoid this error, split the operation into smaller pieces, such as smaller time ranges.

Memory limit of %1% reached, controlled by setting max.query.memory

The server has a per-operation limit on memory utilization, and the requested operation exceeded the limit. The limit is related to the amount of memory in the server, which an administrator can adjust in **ADMIN > Services > [Service Name] > sdk > config**. To avoid this error, split the operation into smaller pieces, such as smaller time ranges.

Timeline

The timeline visualizes the events count that occurs at a specific instance. The timeline provides event counts so that you can see if the number of events increases drastically at a given point in time. The timeline displays activity for the specified service and time range as a bar chart.



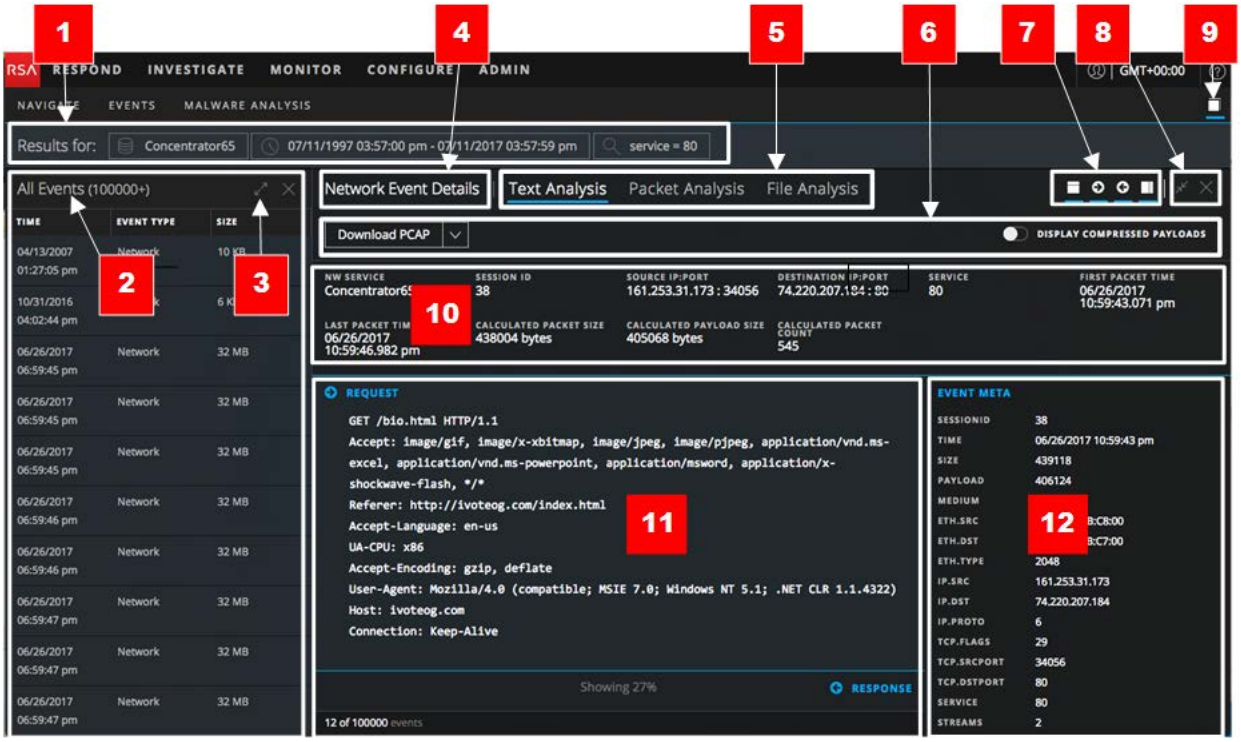
The following are some key features of the timeline:

- The start time and end time are displayed on either side of the timeline to depict the matching events over the duration of the query. The timezone is also displayed on the right side of the timeline by showing the offset from UTC.

- The section of the timeline that matches the event data is highlighted. The data is highlighted up to the specified threshold. If the data is not sorted using time, the entire timeline is highlighted. Also, if the data is not sorted in a specific order, it might highlight the individual bars in the timeline as events are not ordered using time.
- The query results that exceed the threshold are displayed in grey color. If events are sorted using time, the threshold highlight shifts from left (ascending) to right (descending) depending on your sort preference.
- Hover over the timeline to display the total number of events that occurred in the queried time duration. You can hover over the individual bars in the timeline to get the total number of events that occurred at a specific time.
- When you sort the events, the timeline does not refresh. You must run the query again to refresh the timeline.
- All capabilities of the timeline work when user is querying using the collection time (`time`) or when the preference is set to use the event time (`event.time`).

Quick Look for Version 11.0.0.x (End-Of-Life)

The following figure highlights the major features of the Event Analysis view for Version 11.0.0.x.



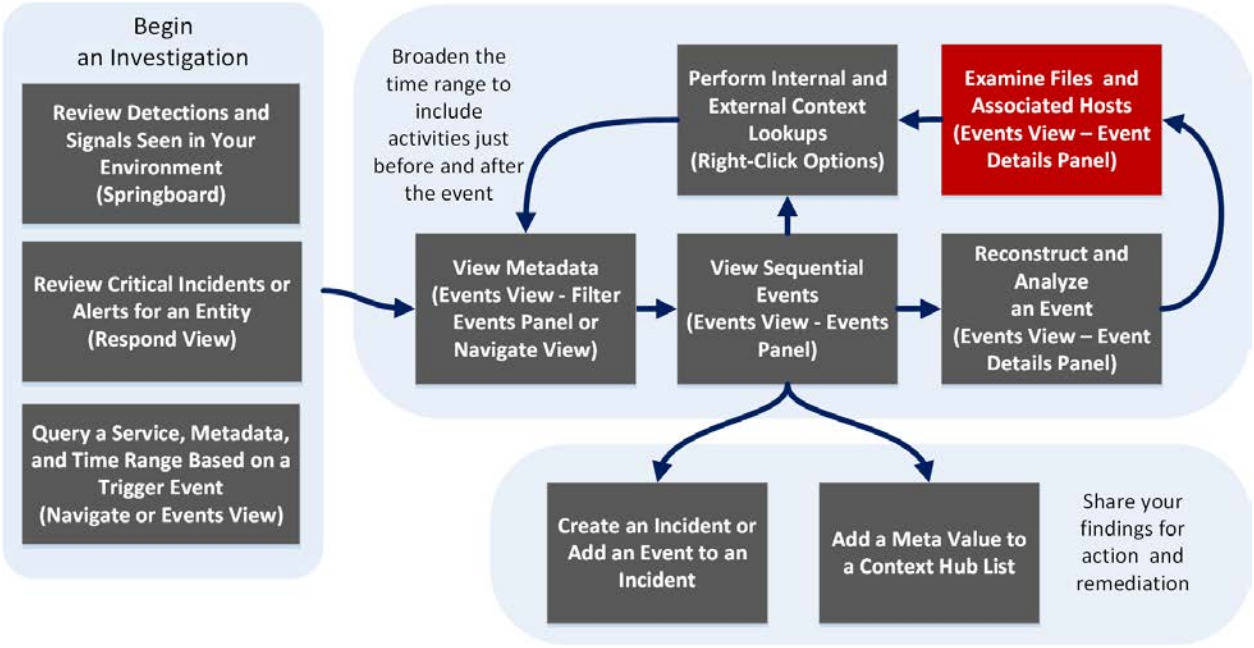
1	The read-only breadcrumb displays the selected service, time range, and query entered in the Navigate view or Events view.
2	This is a read-only list of events based on the query made in the Navigate or Events view. The Events panel includes a count of the events. You can rearrange and resize columns. You can scroll to the bottom of the list, and load more events (see Analyze Events in the Events View).
3, 8	Controls to change the size of the panel and close the panel.
4	The type of event being analyzed is reflected in the heading: Network Event Details, Log Event

	Details, or Endpoint Event Details. Each view is discussed in detail in Analyze Events in the Events View.
5	The types of analysis available for the event type. Network events can use all three types of analysis: text, packet, and file. Log and endpoint events use only text analysis.
6	These options vary for the different types of analysis. They are discussed in detail in Analyze Events in the Events View.
7	Controls to show or hide the Event Header, show or hide requests and responses, and open the Event Meta panel (12). These controls are described in Analyze Events in the Events View.
9	Reopens the Events panel or the Event Meta panel if you have closed it.
10	The Event Header provides summary information about the event. This information is different for the different event types (packet, log, and endpoint).
11	The event data (sometimes called a payload for packets). The event data for a log event or endpoint event is typically a line of text from the raw log rather than request and response shown for a packet.
12	The Event Meta panel lists the meta keys and values found in the data. Some meta data are searchable; they have a binoculars icon, which you can click to see the associated data highlighted in the event data (see Analyze Events in the Events View.

Events View - Email Tab

The Email tab is in the Event Details panel. Here you can view a list of email received and associated attachments for an event.

Workflow

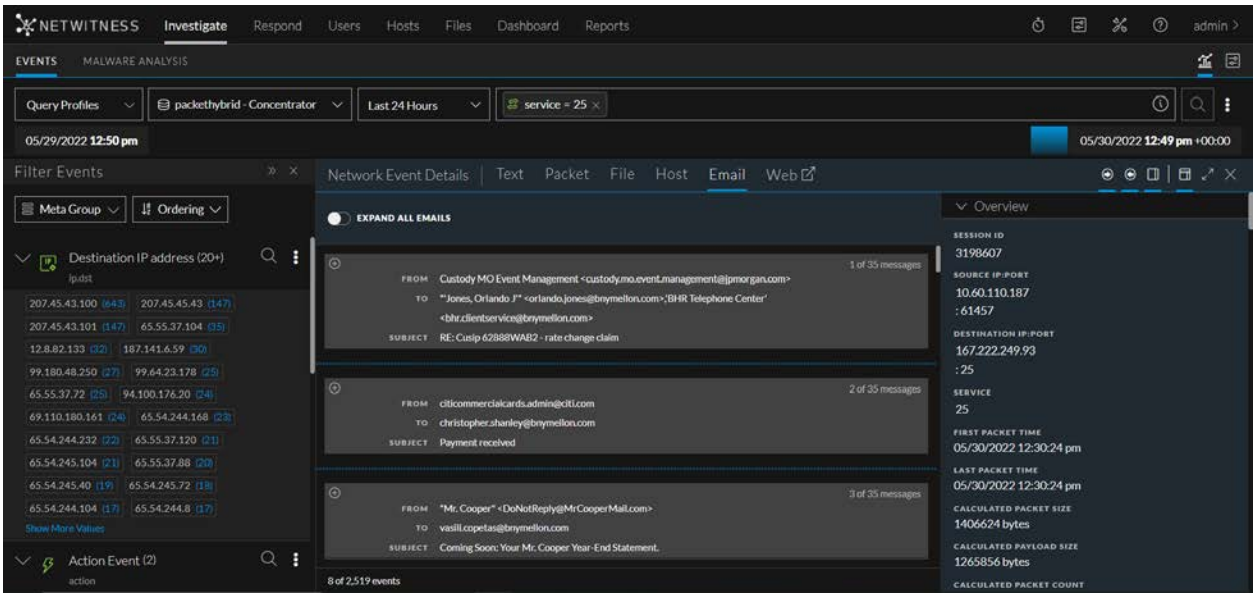


Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

The Email panel displays a list of emails associated with a network event. When an analyst opens the email, the email reconstruction is displayed along with the associated attachments and additional header details, if any.
The following figure is an example of an email reconstruction.



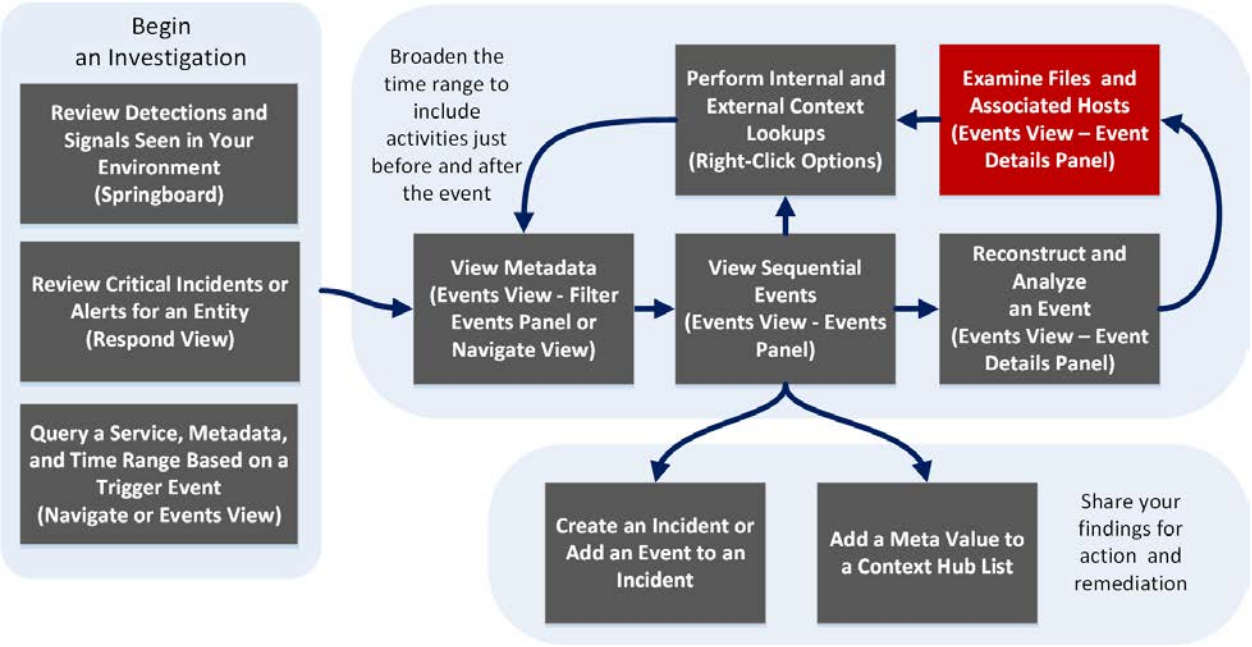
The following table describes all the fields within an email.

Field	Description
From	Displays the email address of the sender of the email.
To	Displays the email addresses of the recipients of the email.
CC (Carbon Copy)	Displays email addresses of additional recipients of the email. The field is displayed only if the sent email has any value and the email addresses are visible to the recipient.
BCC	Displays email addresses of additional recipients privately. This field is displayed only if the sent email has any value and the email addresses are not visible to the recipient.
Reply to	Displays the address designated to receive replies, the sender address.
Subject	Displays the subject of the email.
Attachments	Displays any files shared by the sender that can be downloaded by the recipient. This field is displayed only if email contains attachments. See Download Data in the Events View for details about downloading email attachments.
Additional Header Details	Provides additional details of the email event such as Received, Sender, Message-ID and others.

Events View - File Tab

The File tab is in the Event Details panel. Here you can safely view a list of files and download one or more files in an event.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

*You can perform this task in the current view.

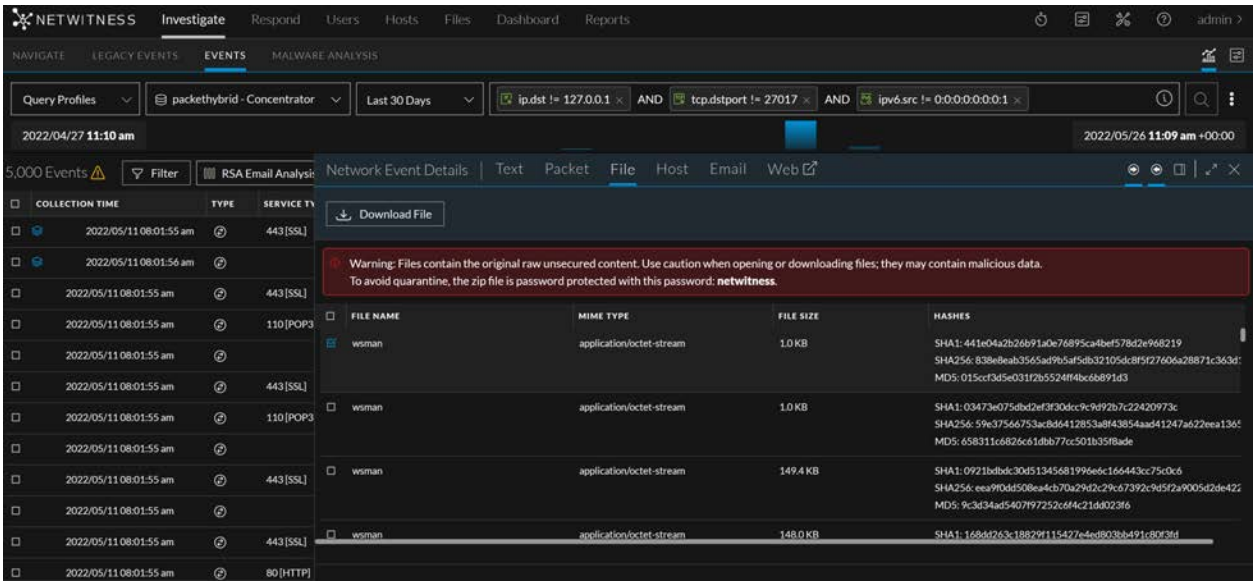
Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

The File panel displays a list of files associated with a network event. You can download files in this view.

Below is an example of the File panel.

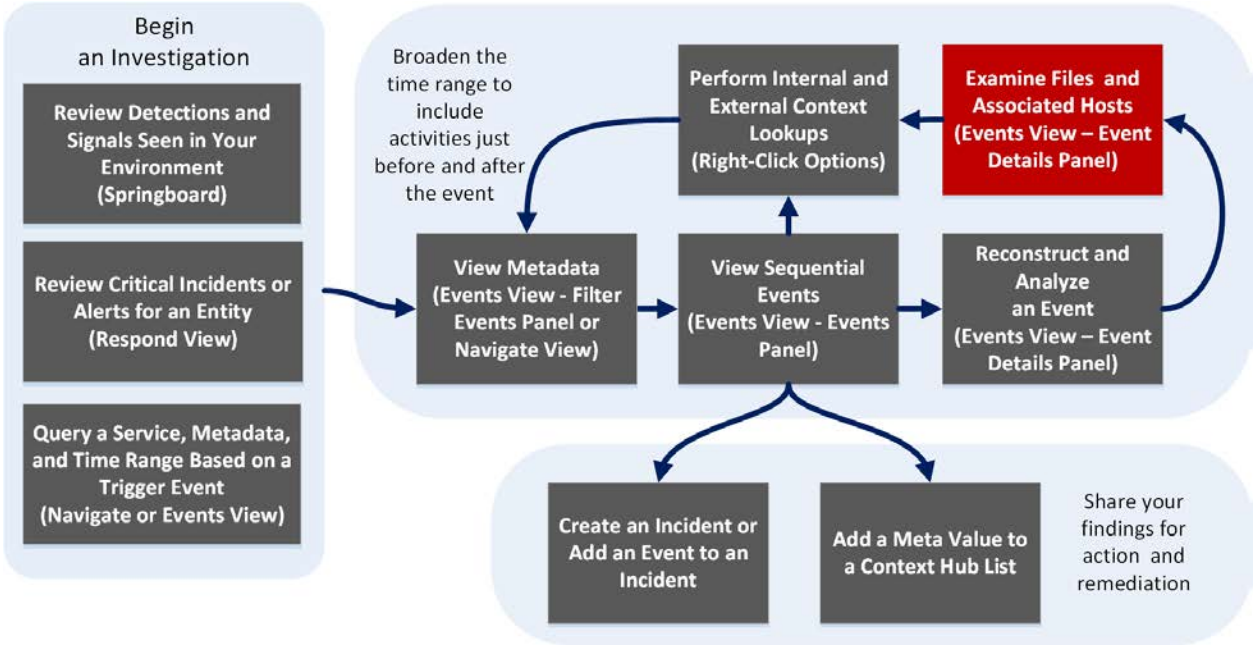


Feature	Description
Download Files button	Click to download one or more selected files.
Event Header	The Event Header displays summary information for the network event that contains the files.
Files List	Scrollable list of associated files that you can select and download.

Events View - Host Tab

The Host tab is in the Event Details panel. Here you can view network events enriched with endpoint data such as host and process triggered for the selected network event and other details such as risk score, reputation, and logged in user. The host panel is available for network events with endpoint data only.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View

User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

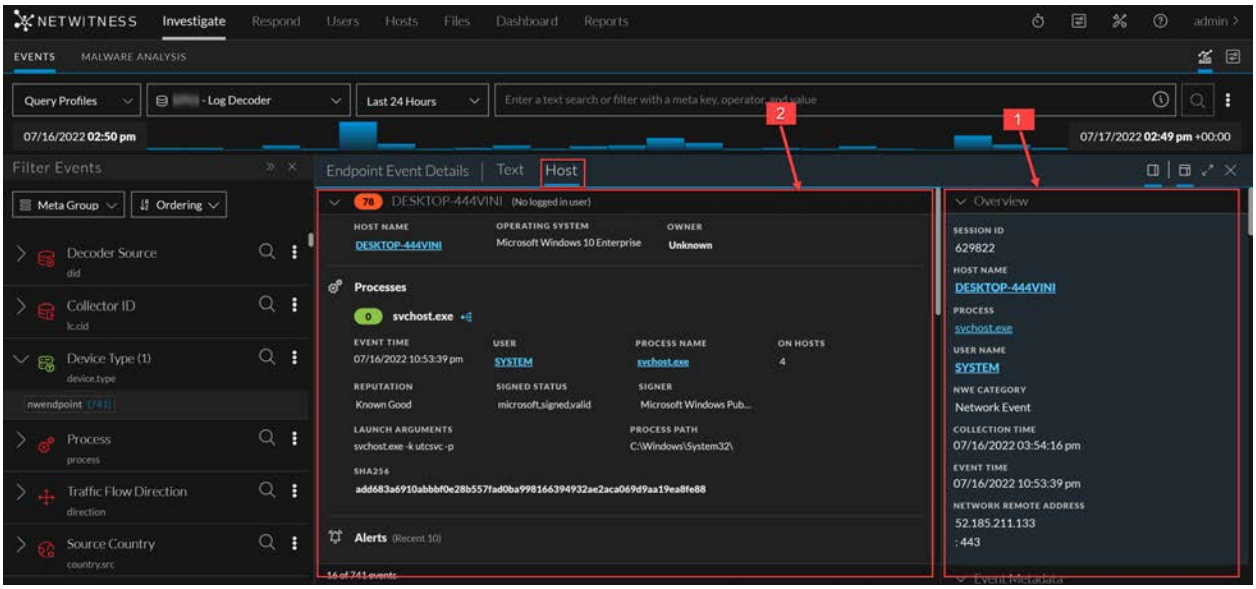
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

Below is an example of the Host panel with labeled features.



1 The event header displays the summary of network events enriched with endpoint data. It includes:

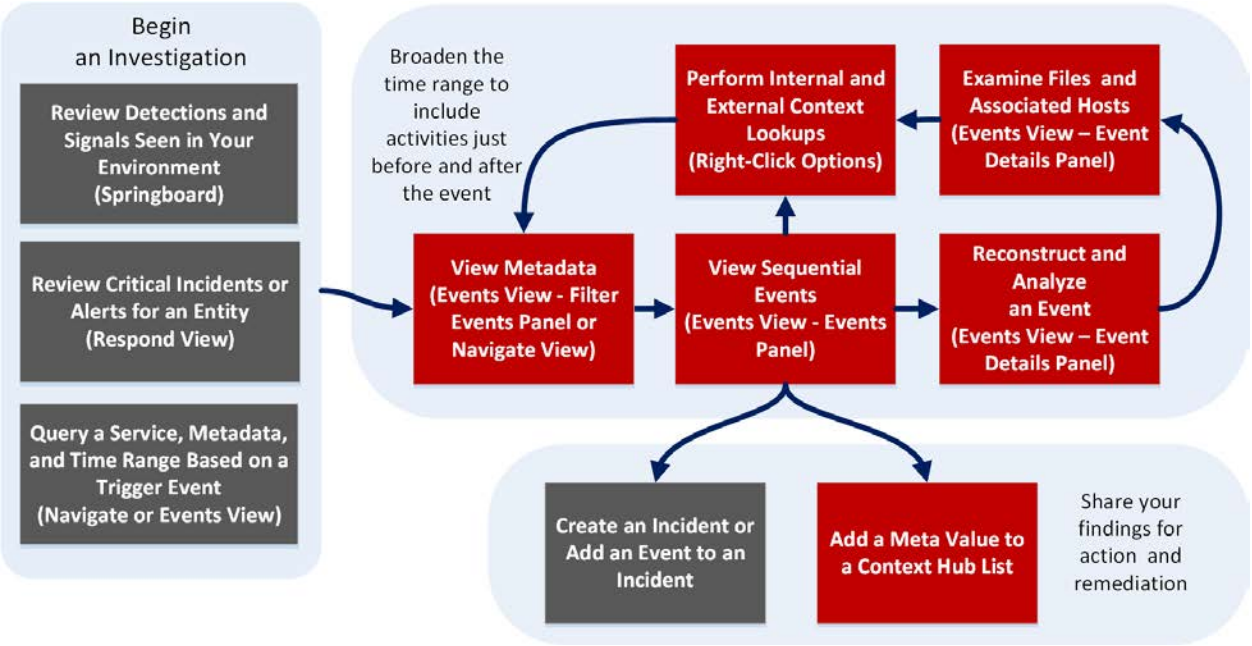
- Host - The host from where the event originated.
- Process - The source process which triggered the event.
- User - The user associated with the triggered process.

2 You can view additional details about the host and process. For more information, see Host Information.

Events View - Packet Tab

The Packet tab is in the Event Details panel. Here you can safely view and interactively analyze the packets and payload of an event.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View

User Role	I want to ...	Show me how
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

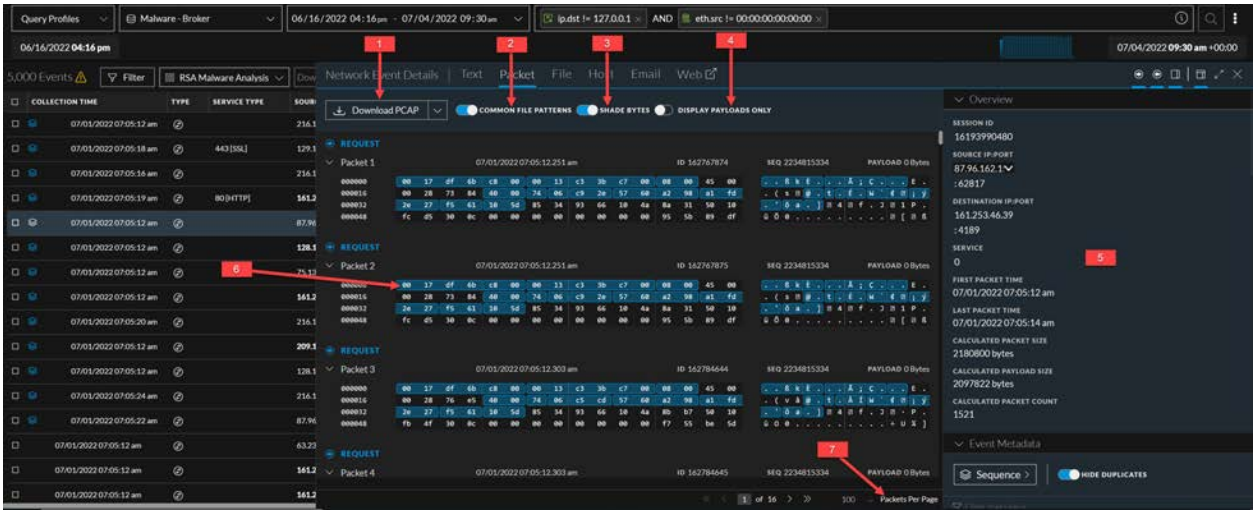
Quick Look

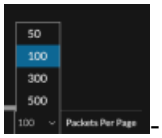
Only network events can be analyzed in the Packet panel. The Packet panel lists each packet in the event. The list of packets is scrollable. When you scroll, the packet or text identification information as well as the Request and Response labels remain visible rather than scrolling out of view.

In Version 11.1 and later, you can use pagination controls to go backward and forward through the pages, go to a specific page, and select the number of packets to display per page (50, 100, 300, or 500).

Each packet is displayed with shading and highlighting to help identify common file patterns: significant header and payload bytes, hexadecimal and ascii bytes, and common file signatures. In addition, you can adjust the request/response display, and display or hide the packet summary.

Below is an example of the Packet panel (formerly known as Packet Analysis) with labels to identify features. For details and examples of each feature, see Analyze Events in the Events View.

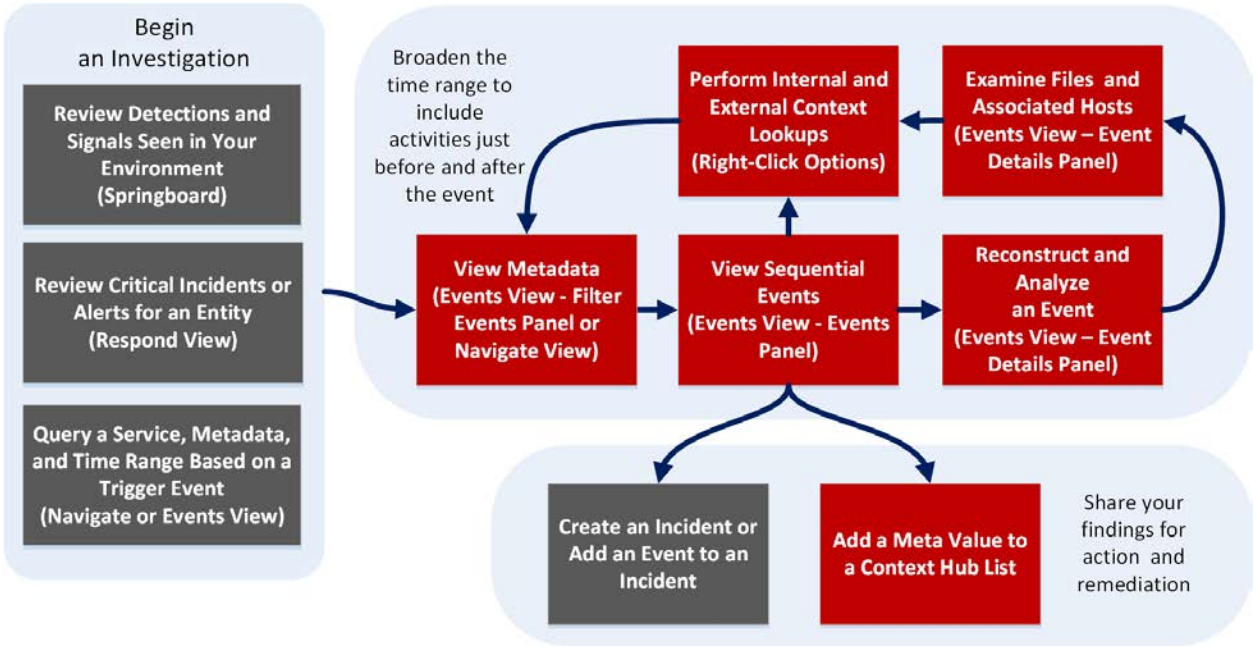


1	Options for exporting a network event. You can export a PCAP, all payloads, request payloads, or response payloads for deeper analysis and to share with others.
2	The option to identify common file signatures is activated by default. Common file signatures are highlighted in orange; hovering over the highlight reveals the file type.
3	The Shade Bytes option adds shading to identify the different hexadecimal bytes (00 to FF) using degrees of highlighting.
4	The option to display payloads only hides the packet headers, leaving more space for the payload.
5	The Overview panel information.
6	Significant bytes are highlighted in a blue background; as you move the cursor over the highlighting the meta data is displayed in a hover box.
7	(Version 11.1 and later) Packet pagination controls allow more flexibility in paging through a list of packets. When a control is unavailable, the image is dimmed; for example, when you are viewing page 1, the ⟪ and ⟨ controls are dimmed. ⟪ - Go to the first page ⟨ - Go to the previous page 1 of 206 - Go to a specific page ⟩ - Go to the next page ⟫ - Go to the last page  - Select the number of packets per page. If you are reconstructing large packets, lowering this limit can improve performance.

Events View - Text Tab

The Text tab is in the Event Details panel. Here you can safely view and analyze the raw text payload of an event. The Text reconstruction includes features that can show decompressed or compressed text, expand truncated entries, perform URL and Base64 encoding and decoding, and download network events, logs, and endpoint events. The text reconstruction is available for all types of events: network, log, and endpoint.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

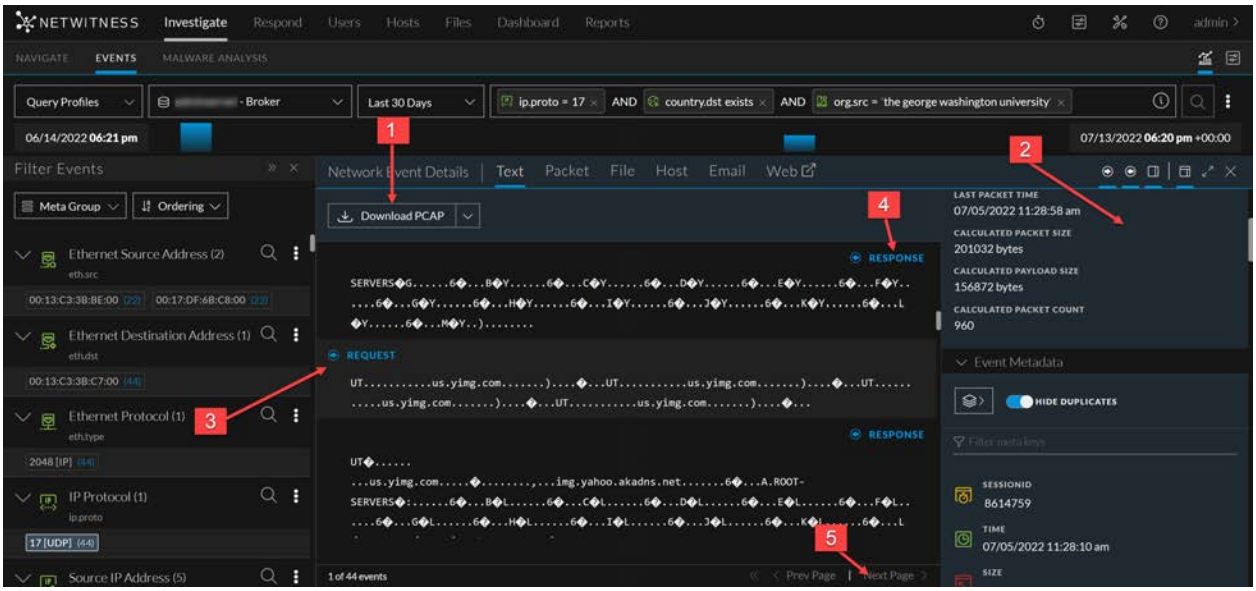
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

The Events view displays the text of a single event in the Text panel (formerly known as Text Analysis). When you click an event in the Event list panel, the adjacent panel shows the text reconstruction. Only the raw log for log events and endpoint events is shown in the Text panel. For network events, the direction of the packet (Request or Response) and contents of each packet are provided in text format. For more examples of the Text, see Reconstructing and Analyzing Events. For detailed procedures, see Analyze Events in the Events View.

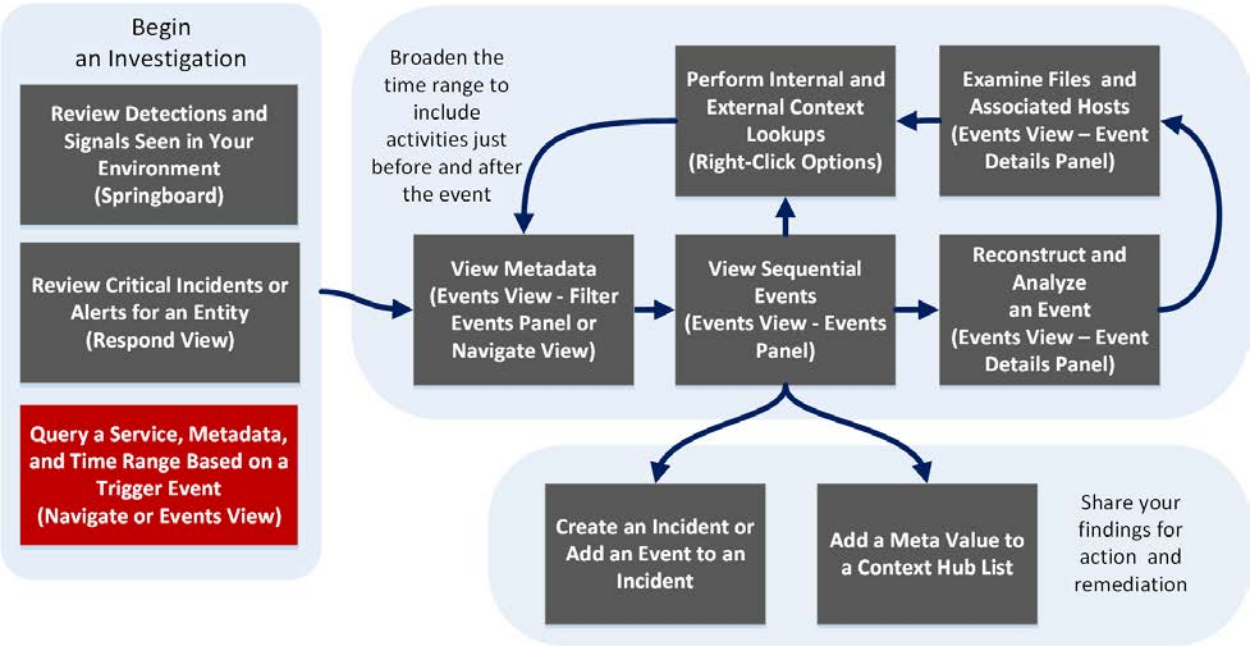


1	Options for exporting a log, a PCAP, or files for deeper analysis and to share with others. This download menu is for network data.
2	The Overview panel information.
3	The payload for a network event includes requests and responses. This is the request side of the packet.
4	This is the response side of the packet.
5	(Version 11.2 and later) Event pagination controls allow more flexibility in paging through a list of events. When a control is unavailable, the image is dimmed; for example, when you are viewing page 1, the ⏮ and ◀ controls are dimmed. ⏮ - Go to the first page ◀ - Go to the previous page ▶ - Go to the next page ⏭ - Go to last page (Only available after last page has already been navigated to)

Investigate Dialog

In the Investigate dialog, analysts can select a service or a collection to investigate. The dialog is automatically displayed when you first go to the Navigate view or Legacy Events view and have not selected a default service to investigate. To access the dialog from a current investigation, select the current service name in the toolbar.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View

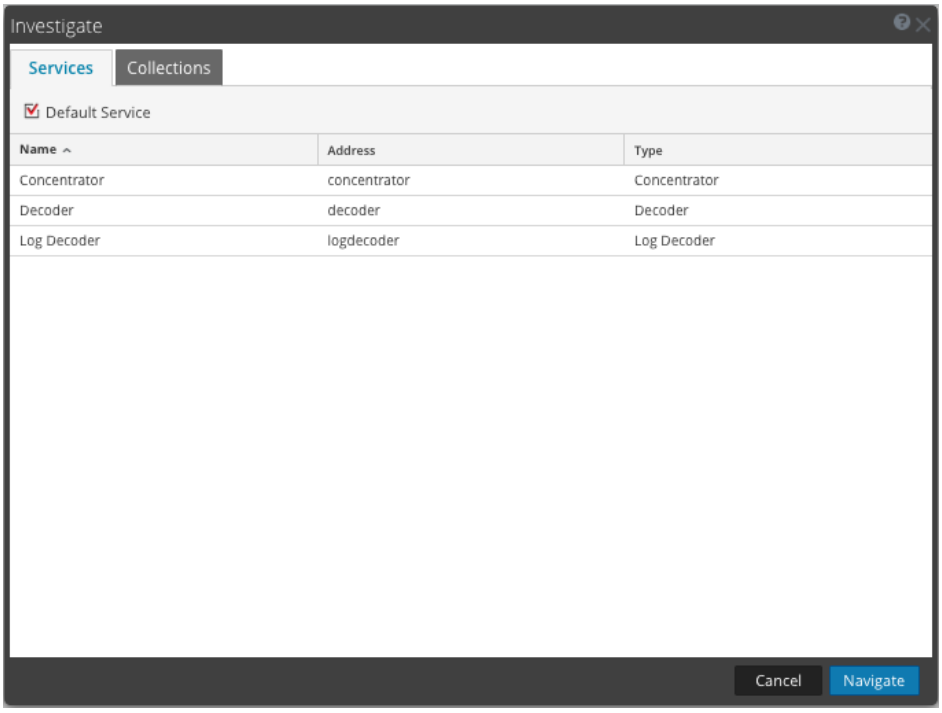
User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look



The Investigate dialog has two tabs: Services and Collections.

Note: Collections are also known as workbench collections. You can only view workbench collections that you have created, and only administrators can create a workbench collection.

The Services tab includes a list of services available for investigation, and three buttons. All features are described in the following table.

Feature	Description
Default Service	Clicking this button sets or clears the default service to investigate. When a service has been set as the default service, the word (Default) is appended to the service name.
Name	The name of the service.
Address	The IP address of the service.
Type	The type of service.
Cancel	Closes the dialog.
Navigate	Opens the selected service in the Navigate or Legacy Events view.

The Collections tab has two buttons and two panels: Workbench and Collections.

The Workbench panel lists available Workbench services by name. After a Workbench service is selected, you can select a collection from the Collections panel.

The Collections panel lists available collections to investigate. After a collection is selected, you can click Navigate to view the collection.

The following table describes the features of the Collections panel.

Feature	Description
Name	The name of the collection.
Type	The type of collection.
Size	The size of the collection.
Data Type	The type of data within the collection.
Date Created	The date the collection was created.

Investigation Tab - User Preferences Panel

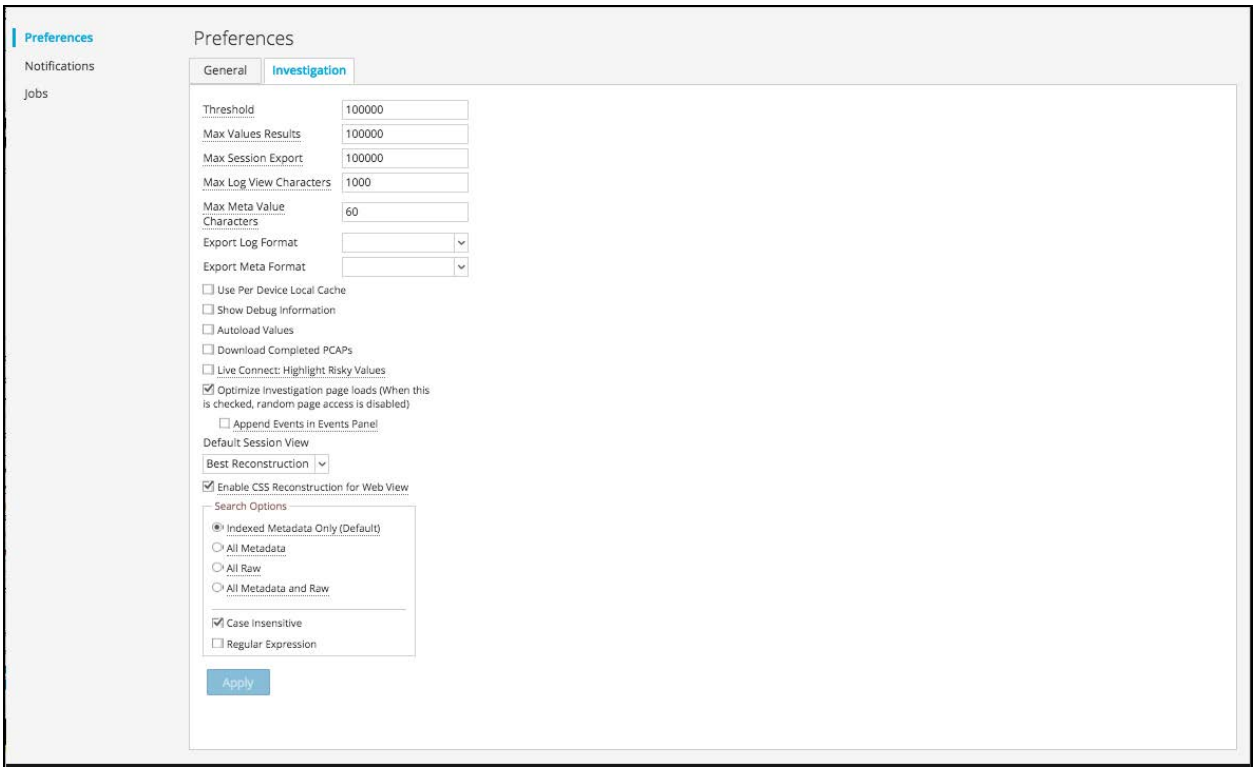
In the Profile view > Preferences panel > Investigation tab, users can set several preferences that affect the performance and behavior of NetWitness when analyzing data, viewing events, and reconstructing events in NetWitness Investigate. To access this tab, select > Profile from the Navigate view or the Legacy Events view. When the Profile view is displayed, select **Preferences > Investigation**. You can change user preferences at any time when you are working in NetWitness.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look

This figure is an example of the Investigation tab, and the following table describes the preferences that affect Investigate. There are slight differences between the 11.1 search settings and later versions of the search settings and these are explained in Search for Text Patterns in the Navigate and Legacy Events Views.



Feature	Description
Threshold	This setting controls the count shown for a meta key value in the Navigate view during the load. A higher threshold allows more accurate counts for a value. However, a higher threshold causes longer load times. When the threshold is reached, NetWitness displays the count and the percentage of time used to reach the count in comparison to the time necessary to load all sessions with that value. For example, **(>100000 - 18%)** indicates that the threshold was set at 100000 and this load took only 18% of the time it would have taken with no threshold set. The default value is **100000**.
Max Values Results	This setting controls the maximum number of values to load in the Navigate View when the Max Results option is selected in the Meta Key Menu for an open meta key. The default value is **1000**.
Max Session Export	This setting controls the maximum number of sessions that can be exported. The default value is **100000**.
Max Log View Characters	This setting controls the maximum number of characters to be displayed on **Investigate > Legacy Events > Log Text**. The default value is **1000**.
Export Log Format	This setting specifies the default format for exporting logs from Investigate. Available options are **Text**, **XML**, **CSV**, and **JSON**. There is no default value for the log export format. If you do not select a format for logs here, NetWitness displays a selection dialog when you invoke export of logs. When you select one of the options from the Export Log Format drop-down menu and click Apply, the setting goes into effect immediately.
Export Meta Format	This setting specifies the default format for exporting meta values from Investigate. Available options are Text, XML, CSV, and JSON. There is no default value for the meta export format. If you do not select a format for exporting meta values here, NetWitness displays a selection dialog when you invoke export of meta values. When you select one of the options from the Export Meta Format drop-down menu and click Apply, the setting goes into effect immediately. **Note:** If you upgrade to version 11.5.2, the **Export Meta Format** preference is not retained and is reset to blank. You must re-configure this value after you upgrade to version 11.5.2.
Use Per Device Local Cache	Allows you to specify the use of locally cached data from the selected service. By default, this checkbox is cleared (Off), which means that Investigate sends a fresh query to the database rather than displaying cached data in the Investigate views after the initial load. If the option is set, Investigate uses the data from local cache.
Show Debug Information	When this option is set, NetWitness displays the `where` clause beneath the breadcrumb in the Navigate view. For each meta value load, the load time is displayed. If the service is a Broker, the elapsed time for each aggregated service is reported. The default value is **Off.**

Feature	Description
Append Events in Events Panel	When this option is set, the events displayed in the **Events Panel** are added incrementally rather than overwriting the currently displayed events. Each time you click the next page icon, the additional events are appended to the previous events; 1 -25, then 1 -50, then 1 -75 and so on. **Note:** This option is available only if the Optimize Investigation Page Loads option is enabled.
Autoload Values	When this option is set, the service values are automatically loaded in the Navigate view. When not set, NetWitness displays a **Load Values** button, allowing the user the opportunity to modify the options. The default value is **Off**.
Download Completed PCAPs	This setting automates the downloading of extracted PCAPs in the Investigate so that you do not have to manually download and open extracted PCAP files in an application, such as Wireshark, that can handle viewing data in a PCAP format.
Live Connect: Highlight Risky Values	If you want NetWitness Platform to highlight and display only IP addresses that are considered to be risky by the NetWitness community, set this option. When not enabled, NetWitness Platform displays all IP addresses. By default, this option is cleared (Off).
Optimize Investigation Page Loads	This option is enabled by default (checked) and controls how the Legacy Events view retrieves events. When enabled, results are returned as quickly as possible, but you cannot go to a specific page in the event list. Clearing the checkbox changes the Events list pagination to allow you to go to a specific page in the list (or to the last page). Being able to go to any page in the list costs additional overhead to determine the events in advance.
Default Session View	This setting selects the default reconstruction type for the initial reconstruction view. By default events are reconstructed using the reconstruction type most appropriate to the event.
Enable CSS Reconstruction for Web View	This setting controls how web content reconstruction is performed. If enabled, the web reconstruction includes cascaded style sheet (CSS) styles and images so that its appearance matches the original view in a web browser. This includes scanning and reconstructing related events, and searching for stylesheets and images used in the target event. The option is enabled by default. Clear the checkbox if there are problems viewing specific websites. **Note:** The appearance of the reconstructed content may not match the original web page perfectly if related images and stylesheets could not be found or were loaded from the web browser's cache. Also, any layout or styling that is performed dynamically via client side javascript does not render in the reconstruction because all client side javascript is removed for security purposes.

Feature	Description
Search Options	This setting specifies the default search options to apply to a search in the Navigate and Legacy Events views. Search for Text Patterns in the Navigate and Legacy Events Views provides detailed information.
Apply	Saves your preferences and puts them into effect immediately.

Investigate View

The Investigate view is the primary entry point to NetWitness Investigate. In Version 11.5, several of the Investigate submenus are moved to the main menu for easier access. Prior to Version 11.5, the Investigate view had six submenus, which opened different views that allow you to analyze events from different perspectives. The submenus that remain under Investigate are: Navigate, Legacy Events, Events (formerly Event Analysis), and Malware Analysis. The Hosts, Files, and Users views (formerly Entitites) are available from the main menus to improve the workflow for analysts.

Note: The Legacy Events view was the original Events view (11.0 to 11.3.x.x). In Version 11.4 and later, the Legacy Events is no longer needed and it is hidden unless the administrator enables it. By default only the Events view appears in the menu, but when the Legacy Events view is enabled, both the Events view and the Legacy Events view are visible in the menu bar.

How NetWitness Investigate Works provides an introduction to all of the capabilities available in the Investigate view.

Legacy Event Reconstruction View

The Event Reconstruction view is deprecated in favor of the Events view. The Legacy Events view provides a reconstruction of a selected event from the Legacy Events view. By default, NetWitness displays the best reconstruction for the event determined by the event content, or the default reconstruction that you have selected in the Default Session View setting for Investigate. You can use the options in the Event Reconstruction toolbar to change the reconstruction method, view top-to-bottom or side-by-side results, select request and response views, export an event, export meta values, extract files, open an email attachment, and open the event in a new tab.

To access this view, do one of the following:

- In any Legacy Events view, double-click an event.
- In the Legacy Events view with Detail View selected, right-click **Events** at the end of the event, and select **Event Reconstruction.**
- In the Event Reconstruction toolbar of previewed reconstruction, click **Open Event in New Tab**.
- In the Navigate view, select **Actions > Go to event in Event Reconstruction**, and enter an event ID.

What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

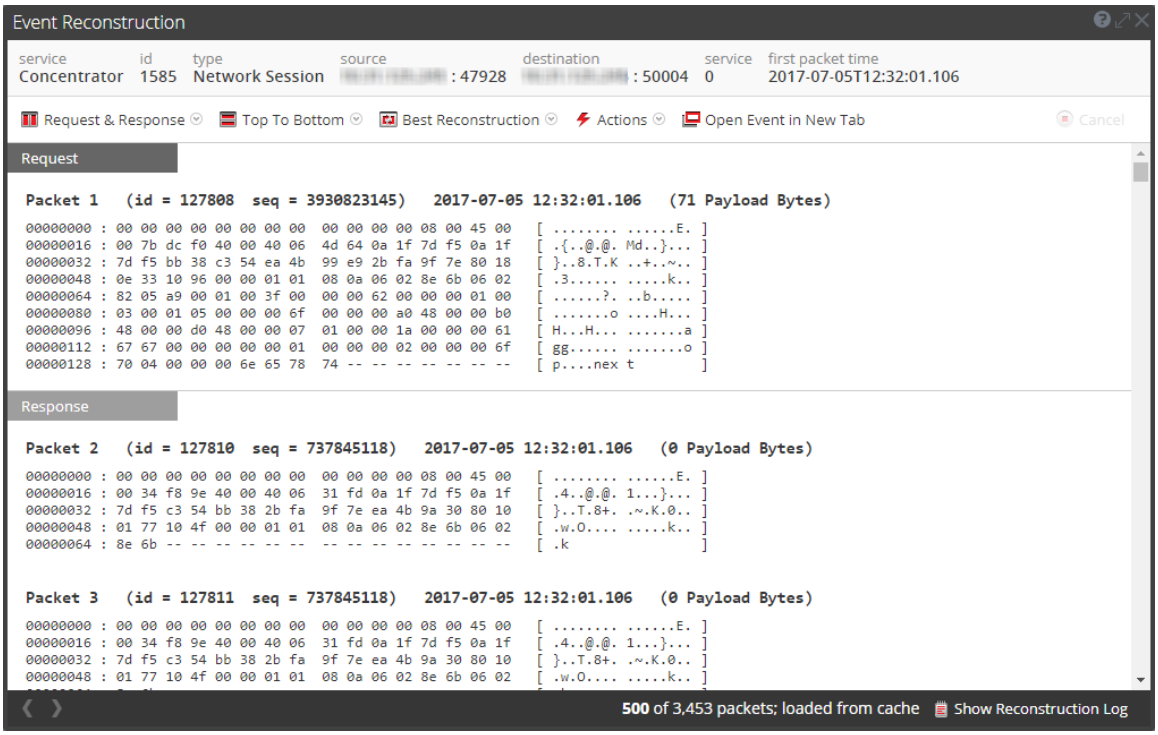
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View

Quick Look

This figure is an example of the Event Reconstruction view. The following table describes the toolbar options.



Feature	Description
Request & Response	Displays a drop-down menu for selecting whether the view displays: • Request & Response • Request • Response
Organization	Displays a drop-down menu for selecting whether the information is displayed top to bottom or side by side.
Reconstruction View	Displays a drop-down menu for selecting what information is displayed. By default, Best Reconstruction is selected. Other options are: • View Meta • View Text • View Hex • View Packets • View Web • View Mail • View Files
Actions	Displays a drop-down menu with the actions available in the Event Reconstruction view (Export PCAP, Extract Files, and Export Meta).

Feature	Description
Open Event in New Tab	Opens the event in a new browser tab.
Event Analysis	Open the event in the Event Analysis view.

Beneath the toolbar is a list of meta keys and values. Some of the keys offer a drop-down menu with available actions.

The bar at the bottom of the view offers several options.

Feature	Description
	Displays the previous event.
	Displays the next event.
Show Reconstruction Log	Displays the reconstruction log at the bottom of the view. Once you click this button, it changes to Hide Reconstruction Log.

Legacy Events View

The Legacy Events view is deprecated in favor of the Events view. In the Legacy Events view a list of events associated with a session is available; this view is optimized for viewing raw events in sequence by time. You can display the events list in several forms, filter events, search for events, and open a reconstruction of an event.

There are two ways to display the Legacy Events view:

- Go to **Investigate > Legacy Events.** NetWitness runs a default query on the last three hours for the default service (if one is set) or displays a dialog in which you can select a service and then runs the default query. The default query selects all events and the Legacy Events view displays events on the selected service, with the oldest events first.
- From within the **Navigate** view, double-click an event. The Legacy Events view displays the events on the selected service based on the drill point in the Navigate view.

Note: The Legacy Events view was the original Events view (11.0 to 11.3.x.x). The Legacy Events is no longer needed and it is hidden unless the administrator enables it. By default only the Events view appears in the menu, but when the Legacy Events view is enabled, both the Events view and the Legacy Events view are visible in the menu bar.

What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Filter Results in the Legacy Events View
- Downloading and Acting Upon Results

Quick Look

The Legacy Events view provides three built-in presentations of event data: the Detail view, the List view, and the Log view. The List view and Detail view provide more information for each event including the timestamp, event type, event theme, and size.

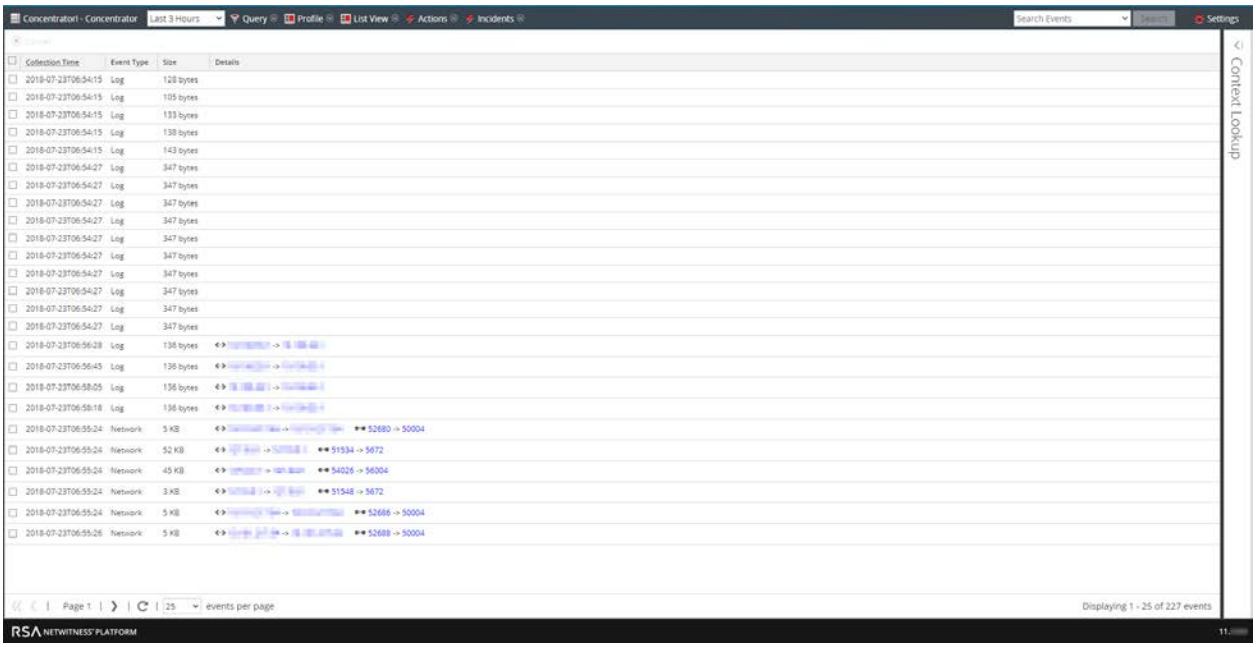
- The List View shows corresponding source and destination address and port information for events in summary form in a grid.
- The Detail View shows all metadata collected for the event in a paged view.
- The Log View is optimized for viewing log and endpoint information, and provides more information for each log including the timestamp, event type, service type, service class, and the logs.

You can use queries, the time range setting, and profiles to filter the events listed in the Legacy Events view. From any view type in Legacy Events view, you can extract files; export network events, endpoint events, logs, and meta values, and open the Event Reconstruction panel. In the Detail View you can also open the event in the Events view.

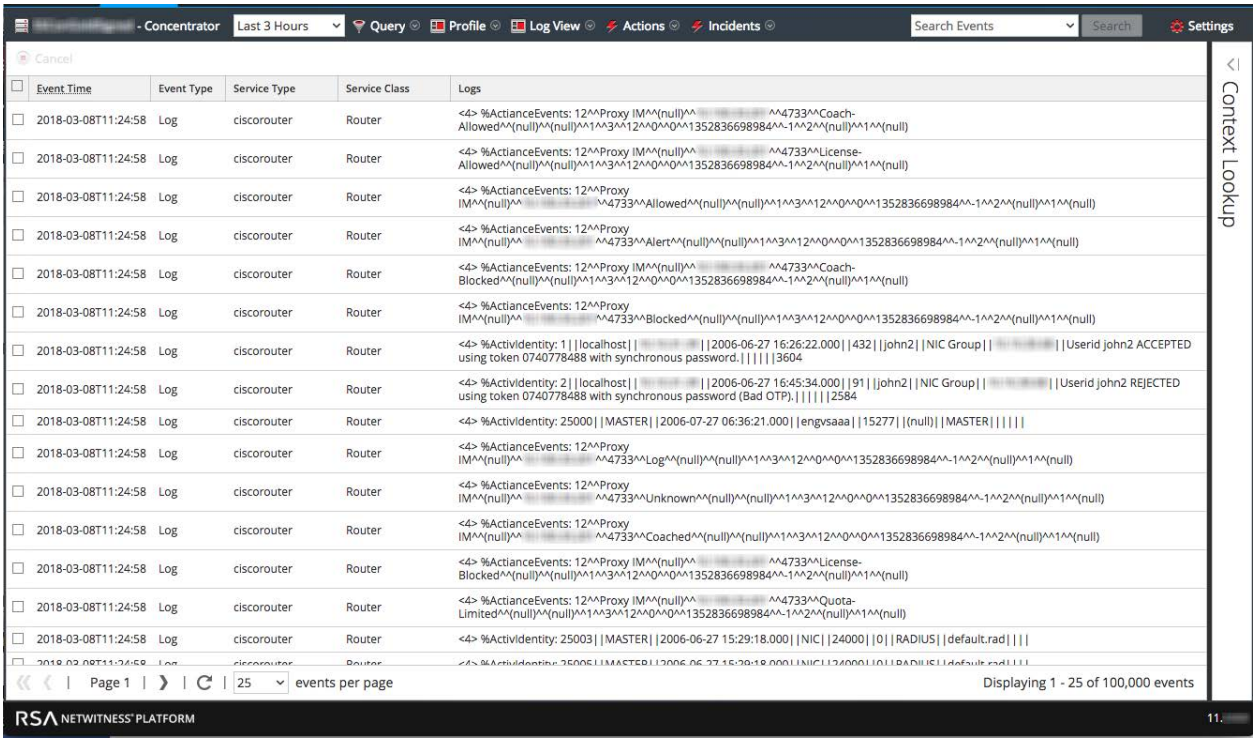
The following figure is an example of events in the Detail View. The Context Lookup panel is visible only if the Context Hub service is configured.



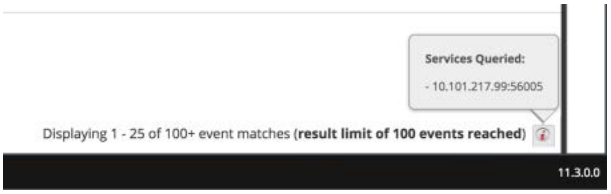
The following figure is an example of events in the List View.



The following figure is an example of the Log View.



The following figure shows the information added to the footer for Version 11.3 and later.



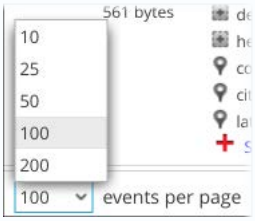
Detailed Description

The Legacy Events view has a toolbar at the top with the following options.

Feature	Description
Select Service	Displays the selected service name next to the icon. Opens the Investigate dialog, in which you can select a service for which the event list is displayed.
Time Range	Displays a drop-down menu for selecting the time range to apply to the event list. You can choose one of the standard options or specify a custom time range.
Query	Displays the Query dialog, in which you can enter a custom query directly instead of drilling down the data (see Create a Query in the Navigate and Legacy Events Views).
Profile	Displays the Profile menu; the currently selected profile is displayed in the toolbar. The menu options include built-in (Default) profiles and custom profiles, as well as an option to manage profiles. Each profile can include a meta group, a column group, and a beginning query that is applied to the Navigate view (meta groups and queries) and the Legacy Events view (column groups and queries) as you investigate events. (see Use Query Profiles to Encapsulate Common Areas for Investigation).
View Type Drop-down	Displays a drop-down menu for selecting the event view type. • Detail View shows events in a paged format with detailed information for each event. • List view shows the events in table form with a summary of each event in a separate row. • Log View shows a log-oriented events grid with a summary of each log in a separate row. • Custom Column Groups displays the event list using a column group selected from a drop-down list of custom column groups. • Manage Column Groups displays the dialog for creating and editing custom column groups.
Actions	Displays a drop-down menu with actions in the Legacy Events view: • Export an events as PCAP files, export logs, export endpoint events, or export meta values. • View an event reconstruction in a popup window or in a new tab. • Reset all filters in the Legacy Events view.
Incidents	Create a new incident in Respond and add the selected events, or add selected events to an existing incident in Respond.

Feature	Description
Search	Displays the Search Events options, which allow you to specify the export log and export meta value format with additional options explained in Search for Text Patterns in the Navigate and Legacy Events Views.
Settings	Displays the Investigation settings for the Legacy Events view (which are also available in the Profile view) so that you can change Investigation settings without navigating away from the Legacy Events view. When you change a setting In the Legacy Events view, the setting is also changed in the Profile view (see Configure the Navigate View and Legacy Events View).

Other features of the Legacy Events view are described in this table.

Feature	Description
+ Show Additional Meta (in the Detail View of an event)	Displays the rest of the metadata for the event.
Event Analysis (in the Detail View of an event)	Opens the selected event in the Events view.
 (in the footer)	Pagination controls allow more flexibility in paging through a list of events. When a control is unavailable, the image is dimmed; for example, when you are viewing page 1, the « and ‹ controls are dimmed. « - Go to the first page ‹ - Go to the previous page 3 Page 3 - Go to a specific page › - Go to the next page » - Go to the last page  - Select the number of packets per page When you select a number of events per page, the setting is saved in browser cache so that you do not have to select your preferred number of events each time you log in. The setting applies to all views: Log View, List View, and Details View.

Feature	Description
Displaying 1-100 of 100,000 events (in the footer) Displaying 1-25 of 100+ event matches (result limit of 100 events reached) (in the footer)	Displays the count of events displayed versus the total number of matching events. In Version 11.3 and later, the footer includes a notification if the results limit configured by the administrator has been reached to let you know that more results are available but not viewable. To view the additional results, you need to refine the filter to get fewer results. Clicking the information icon  in the footer displays the IP address and connecting port number for all services queried.

Manage Default Meta Keys Dialog

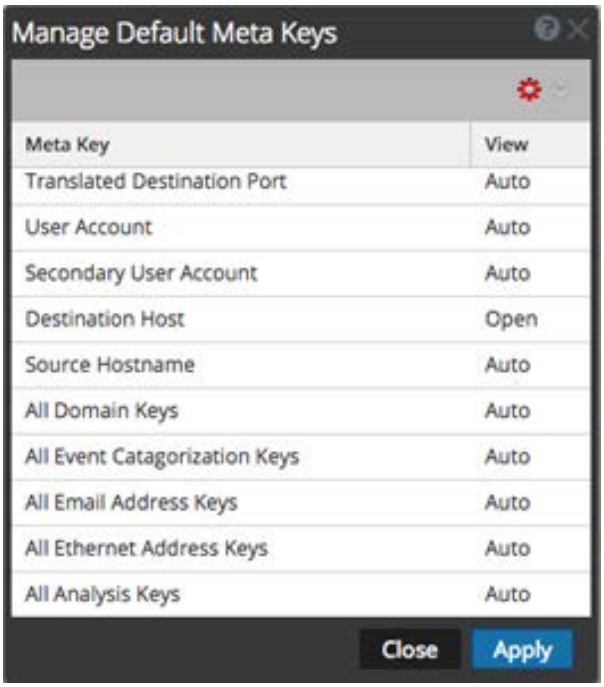
In the Manage Default Meta Keys dialog, analysts can specify the meta keys to be displayed in the Navigate view > Values panel (see Manage and Apply Default Meta Keys in an Investigation. This can help you find the desired data more quickly and prevents the loading of meta keys that are not of interest. To access this dialog, in the **Navigate View** toolbar, select **Meta > Manage Default Meta Keys**.

Related Topics

- How NetWitness Investigate Works
- Use Meta Groups to Focus on Relevant Meta Keys

Quick Look

The following figure illustrates the Manage Default Meta Keys dialog, which has a list of meta keys, toolbar, Close button, and Apply button. In the list, you can view, sort, and manage default meta keys. If you click and drag meta keys, you can rearrange their order. The following table describes columns in the list.



Column	Description
Meta Key	This column displays the meta keys available for the service. In Version 11.1 and later, default meta entities are also included, for example All Domain Keys and All Email Address Keys.

Column	Description
View	This column displays the type of view assigned to each meta key. By clicking on the view in each row, you can assign the meta key a different default view. There are four views: • Auto: Reverts to the default view for meta keys as specified in the service index file. • Close: The values of this meta key are closed by default, and can be opened manually. • Hidden: These meta keys are hidden by default, and are not shown in Investigation at all. • Open: The values of this meta key are displayed by default. When you modify the default meta keys for a non-indexed meta key, you cannot set the key to **Open**. If you change the default view for a group of meta keys to **Open** and some of the meta keys are non-indexed, the non-indexed meta keys revert to **Auto**. As a result, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are **Closed** until opened manually.

The following table describes the toolbar options and buttons.

Feature	Description
	Displays a drop-down menu that allows you change the default view of all the meta keys. There are four views: • Auto: Reverts to the default view for meta keys as specified in the service index file. • Close: The values of this meta key are closed by default. • Hidden: The values of this meta key are hidden by default. • Open: The values of this meta key are displayed by default.
Close	Closes the dialog. Any unsaved changes are lost.
Apply	Applies the changes, and they become effective immediately.

Meta Groups Dialogs

You can use meta groups to filter data displayed in an investigation. A fresh installation of NetWitness includes built-in meta groups to help you find interesting data sets in Investigate. The built-in meta groups are prefixed with RSA for identification and can be duplicated but cannot be edited or deleted. You can create your own groups and you can duplicate and edit a built-in group to create a custom group. With a meta group in effect during an investigation, the information in the Navigate view and Events view includes only the meta keys in the selected group.

While the functionality of meta groups is similar in the Navigate view and the Events view, the user interface and some of the procedures are different.

Using options in the Events view Meta Groups menu (Version 11.5 and later) , you can:

- Select a meta group to apply.
- See the details of a meta group.
- Create, edit, and delete custom meta groups.
- Clone and edit the clone of a built-in or custom meta group.

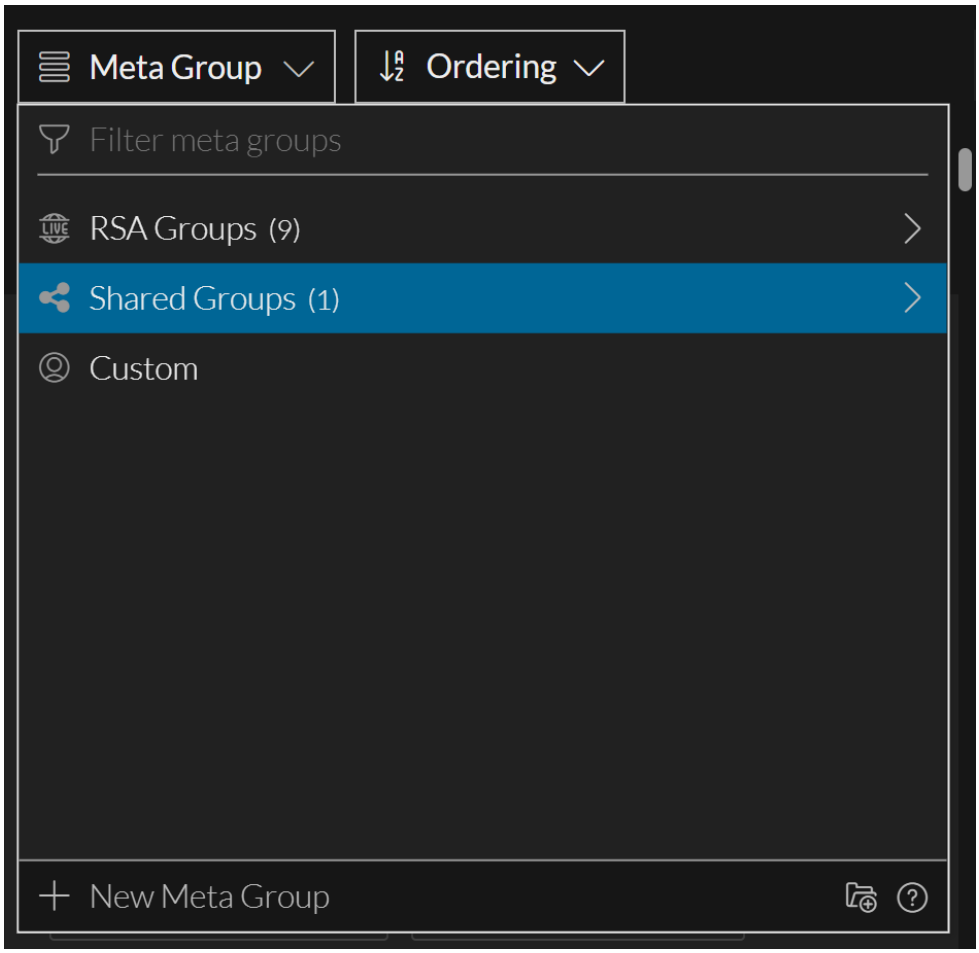
Using options In the Navigate view Manage Meta Groups dialog, you can do all of the above as well as import and export a meta group. Refer to Use Meta Groups to Focus on Relevant Meta Keys for detailed information.

Related Topics

- How NetWitness Investigate Works
- Use Meta Groups to Focus on Relevant Meta Keys
- Filter Results in the Navigate View

Quick Look - Meta Groups Menu, Create Meta Group Dialog, and Meta Group Details Dialog

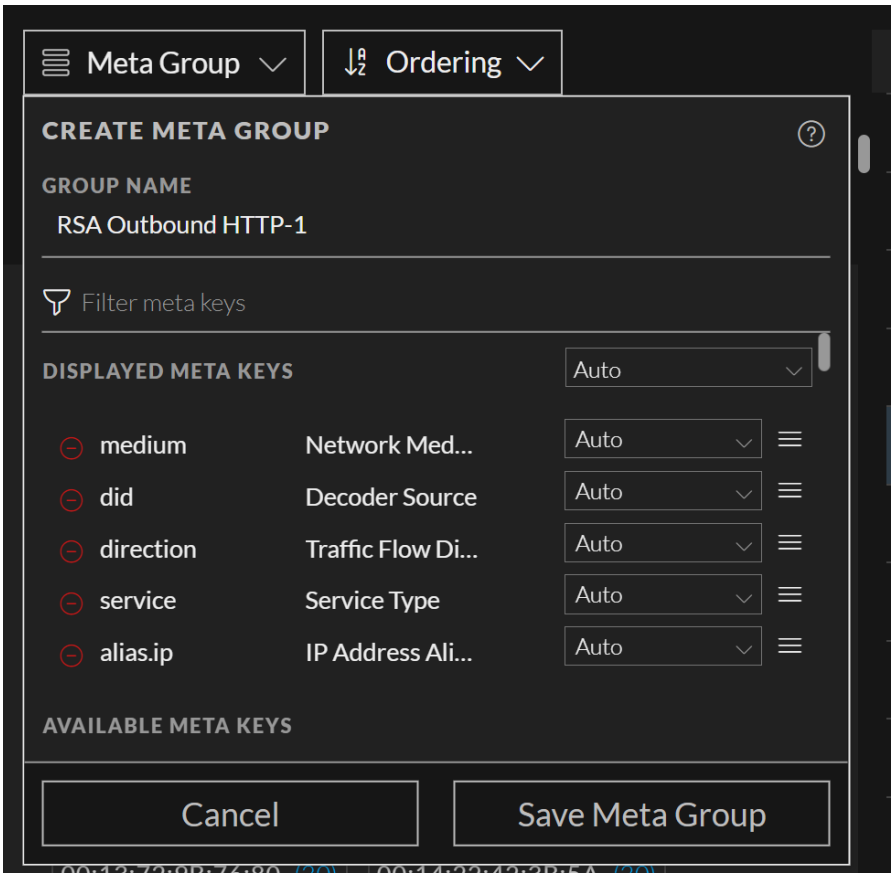
This section introduces the Meta Groups Menu, Create Meta Group dialog, and the Meta Group Details dialog. The following figure is an example of the Meta Group menu. The table describes the options.



Feature	Description
Visibility Options	Control the types of meta groups that are visible in the list. You can use any combination of the visibility options: **Private**, **Shared**, or **RSA** (blue = selected, black = not selected). Initially none of the buttons are selected and all meta group types are visible. This is the same result as if all three buttons are selected. The visibility options work together with text in the Filter Meta Groups field. If the visibility option is hiding built-in groups (which include "RSA" in the group name) and you search for a name that contains "RSA," the list is empty. Private = display private groups that only you can manage Shared = display shared groups that anyone in your organization can manage RSA = display built-in groups that only RSA can manage
Filter Meta Groups	Filters the list of meta groups as you type text so that only group names that contain the typed text are displayed.
Meta Group List	The list of meta groups consists of custom and built-in groups. Custom meta groups can be shared or private. The RSA meta groups are built-in meta groups; you cannot edit or delete these, but you can make a copy and edit the copy. Icons preceding the meta group name distinguish the private groups, shared groups, and built-in groups.

Feature	Description
New Meta Group	Displays the Create Meta Group dialog, where you can create a custom meta group.

The Create Meta Group dialog, shown in the figure on the left, allows you to define a custom meta group. The figure on the right illustrates the Meta Group Details dialog, in which you can edit a custom meta group. The table describes the fields and options in the dialogs.

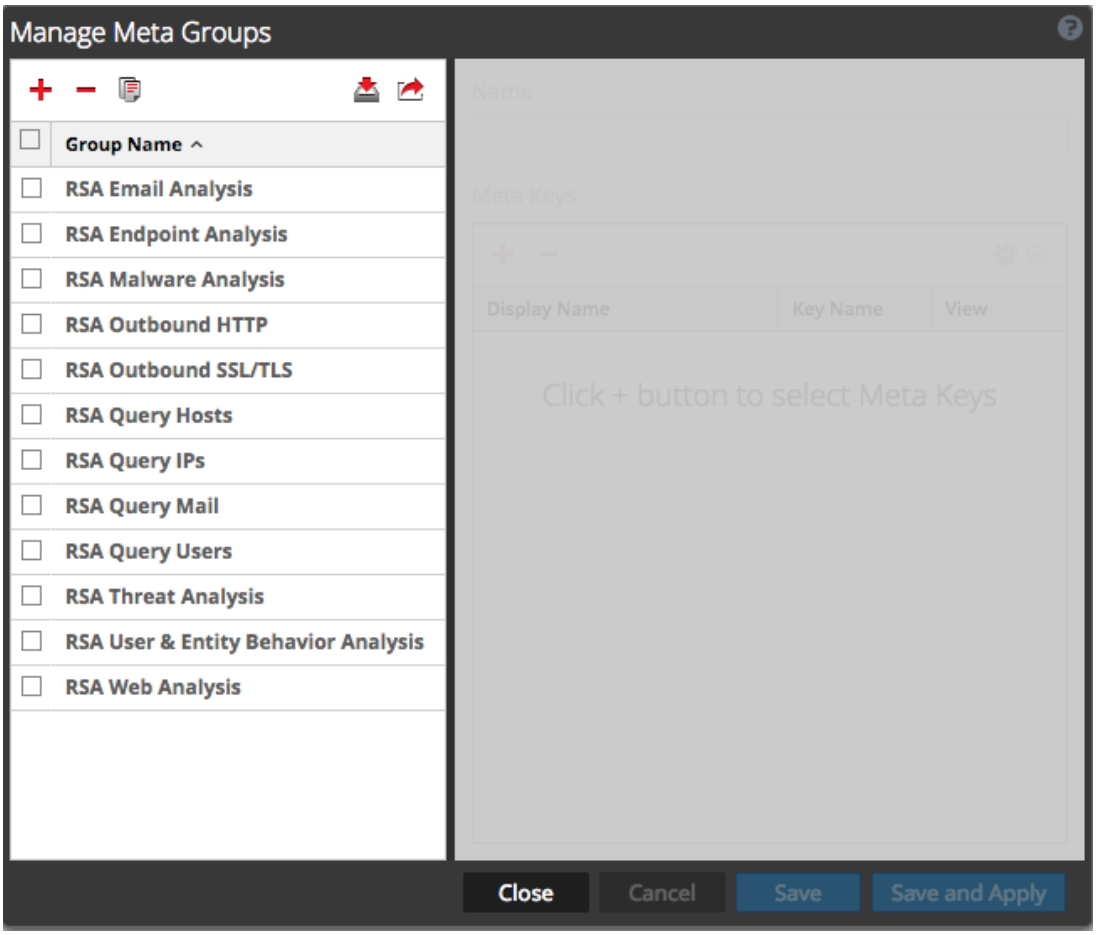


Feature	Description
	Creates a copy of the meta group so that you can edit a copy. This is useful if you want your own copy of a built-in group, a shared copy of a private group, or a private copy of a shared group.
	Deletes the custom meta group that you are currently editing. This action is irreversible and applies globally. If the meta group is a shared group, it is no longer available to anyone.
Group Name	Displays the name of the meta group. The name must be unique and contain fewer than 64 characters. You can type in this field to edit the name in a custom meta group.
Sharing	Specifies whether the meta group is shared or private. This setting is available when you first create the group. After it is created, you cannot change a shared column group to private, or a private column group to shared.

Feature	Description
Filter Meta Keys	Filters the Displayed Meta Keys and Available Meta Keys based on the text that you type. Only meta keys that contain the typed text are displayed.
Displayed Meta Keys	Displays a scrollable list of meta keys that are selected for use in the custom meta group. You can add meta keys in the Available Meta Keys list to this list, remove meta keys from this list , and drag meta keys up or down to change the order in this list ( ). Drag and Drop is disabled when text is typed in the Filter Meta Keys field. For each displayed meta key you can choose:
Available Meta Keys	Displays a scrollable list of meta keys that are available (on the service) for use in the custom column group. You can add them to the Available Meta Keys list. Clicking next to the meta key name adds it to the Displayed Meta Keys list. You can also set the initial view of each meta key: Open, Closed, Hidden, or Auto (the default setting).
Initial View Option	For each meta key, you can set the initial view option: -When set to Auto, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are Closed until opened manually. If you change the default view for a group of meta keys to Open and some of the meta keys are non-indexed, the non-indexed meta keys revert to Auto. -Open meta keys are listed in the Filter Events panel, and the values are loaded. -Closed meta keys are listed in the Filter Events panel, but the meta values are not loaded until you open the meta key. -Hidden meta keys are not listed in the Filter Events panel at all. This is useful if you are using a single meta group for multiple purposes instead of creating several meta groups; you can turn off certain keys off without removing them from the meta group. You can also use the Hidden view when testing out some new keys or if you want to prepare a meta group with some new meta keys that are not yet available and would error out if in an Auto, Open, or Closed state.
	Allows you to drag and drop meta keys in the Displayed Meta Keys list so that you can see the data in the order you prefer.
Close button	Closes the dialog.
Save Meta Group	For the Create Meta Group dialog only, saves the new meta group.
Reset	For the Meta Group Details dialog only reverts the edited meta group to the last saved state.
Update Meta Group	For the Meta Group Details dialog only, applies changes to an edited meta group.
Select Meta Group	Applies the meta group. The Filter Events panel is refreshed to display only the meta keys in the selected meta group.

Quick Look - Manage Meta Groups Dialog

The following figure is an example of the Manage Meta Groups dialog.



The Meta Groups panel is on the left side of the Manage Meta Groups dialog. This is where you can add, delete, import, and export meta groups. The following table describes the features of the Meta Groups panel.

Feature	Description
	Adds a meta group using the Settings panel on the right side of the Manage Meta Groups dialog.
	Deletes the selected meta group. A confirmation dialog is displayed before the meta group is deleted.
	Creates a copy of the selected meta group.
	Displays the Meta Group Import dialog, where you can upload a file.
	Exports the selected meta group to your computer.
Group Name	Lists all meta group names.

The Settings panel is on the right side of the Manage Meta Groups dialog. This is where you create and edit meta groups. Below the Name field is the Meta Keys list. The following table describes the features of the Settings panel.

Feature	Description
Name	Displays the name of the selected meta group.
	Displays the Available Meta Keys dialog, where you can select meta keys to add to the group.
	Deletes the selected meta keys.
	Displays a drop-down menu, where you can select the view for all meta keys. There are four options based on the possible values for the `defaultAction` property used to define a key in the custom index file for the service: • Hidden: These meta keys are hidden by default, and are not shown in Investigation at all. • Open: The values of this meta key are displayed by default. • Close: The values of this meta key are closed by default, and can be opened manually. • Auto: Reverts to the default view for meta keys as specified in the service index file.
Display Name	Indicates the name that is displayed for the key in Investigate views, and is defined by the `description` property for the key in the custom index file for the service..
Key Name	Indicates the `name` of the meta key as defined in the custom index file for the service.
View	Indicates the view to which the meta key is set. You can change: • Change the view for all meta keys by clicking  in the View column header, then selecting a view from the drop-down menu. • Change the view for a single meta key by clicking a single meta key in the View column, then  and selecting a view from the drop-down menu.

The following table describes the buttons at the bottom of the dialog.

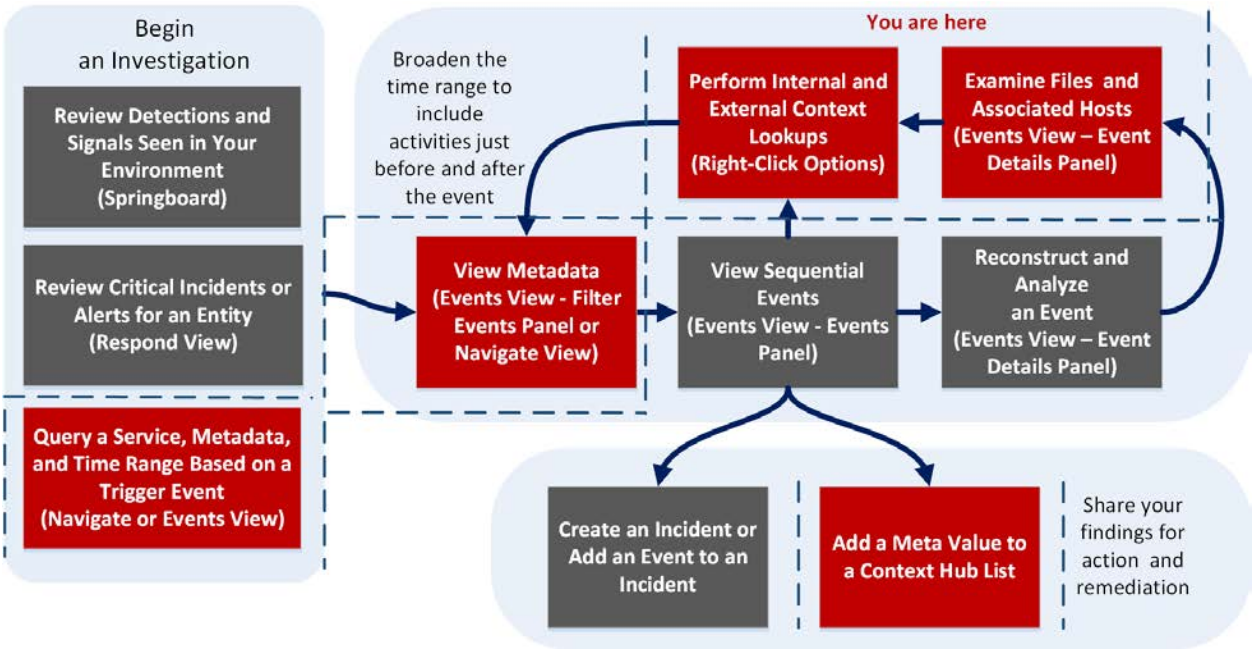
Feature	Description
Close	Closes the dialog.
Cancel	Cancels all changes.
Save	Saves all changes.
Save and Apply	Saves and immediately applies all changes.

Navigate View

The Navigate view (**Investigate** > Navigate) displays event metadata--the meta keys and meta values--that were found in captured data for the selected service. The data is filtered and displayed in accordance with the options you set for profile, time range, meta group, and query. You can also drill into the data by clicking meta keys and meta values.

Note: By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.

Workflow



In the Navigate view, you can:

- View metadata for events in the Values panel.
- Visualize events in a timeline or parallel coordinates chart.
- Save events, go to an event using the event ID, visualize an event, and print an event.
- View additional contextual data for meta keys and values.
- Open a drill point or an event in the Legacy Events or the Events view.

What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

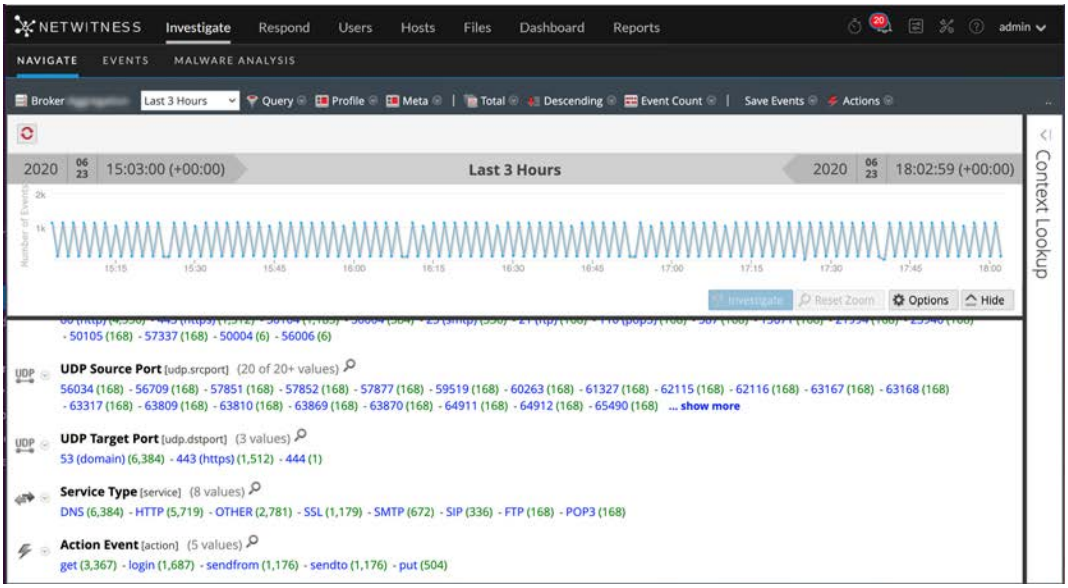
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Legacy Events View
- Events View

Quick Look

This figure illustrates the Version 11.5 Navigate view.



The Navigate view consists of these features:

- Toolbar
- Pause/reload button and breadcrumb
- Time banner
- Optional debug information.
- Collapsible Visualization panel
- Values panel
- Context Lookup panel
- Context menus

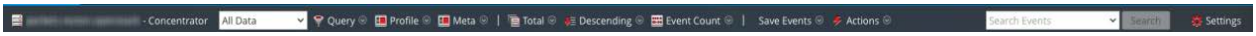
Toolbar

The following figure is an example of the toolbar. The toolbar provides a way to:

- Change the service being investigated.
- Control the range of data displayed: You can select use profiles, set a time range, use meta groups, and create queries to apply to the data.
- Set the quantification method and sorting method for data in the Values panel.
- Perform actions on the results. You can export and print results, open an event for which you have an

event ID in the Legacy Events view or Events view, and pass a query to Informer.

- Configure Investigate settings without navigating away from the Investigate views.



Some of the toolbar options are labeled with the default value or the selected value rather than displaying the name of the option. For example, the time range option in the example above is labeled **Last 5 Minutes** to reflect the currently selected value. These are the toolbar options.

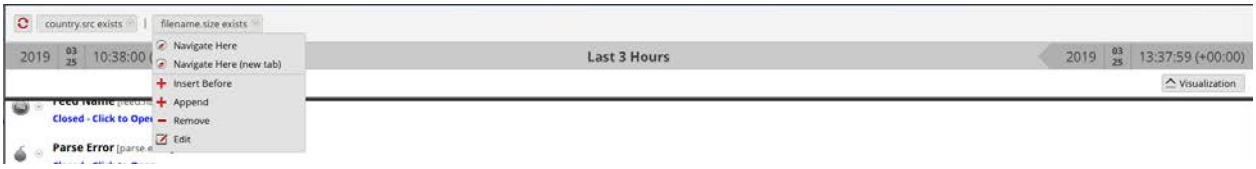
Option	Description
	Displays the selected service name next to the icon. Clicking the icon opens the Investigate a Service dialog, in which you can select a service to investigate and set the default service to investigate (see Begin an Investigation in the Navigate or Legacy Events View). Changing the service does not cause a reload of the data.
Time Range	Displays the Time Range options; the currently selected option is displayed in the toolbar (see Filter Results in the Navigate View). Possible choices are: • All Data • Last 5, 10, 15, or 30 Minutes • Last Hour, Last 3, 6, 12, or 24 Hours • Last 2 or 5 Days • Early Morning • Morning • Afternoon • Evening • All Day • Yesterday • This Week • Last Week • Custom **Note:** If you specify custom start or end times in seconds, the value for start time in seconds always defaults to :00, and the value for end time in seconds always defaults to :59. For example, if you are using time to drill down into an issue, the drill time will be interpreted as `HH:MM:00 - HH:MM:59`. Seconds display in this format in Investigate functions.
Query	Displays the Query dialog, in which you can enter a custom query directly instead of drilling down the data. See Query Dialog for a description of the dialog.

Option	Description
Profile	Displays the Profile menu; the currently selected profile is displayed in the toolbar. A profile allows you to manage and use profiles that can include custom meta groups, a default column group, and a beginning query. The Profiles apply to the Navigate view (meta groups and queries), the Legacy Events view, and the Events view (column groups and queries). See Use Query Profiles to Encapsulate Common Areas for Investigation for more information.
Meta	Displays the Meta Group menu. You can use Default Meta Keys or a custom Meta Group. You also have the option to make changes to both group types (see Use Meta Groups to Focus on Relevant Meta Keys).
Sort Field	Displays the Sort Field menu; the currently selected option is displayed in the toolbar. The menu has two options: Order by Total and Order by Value. The Sort Field is a complement to the Sort Order option; the data for each meta key is ordered based on the total (green number) or the meta value (blue text) (see Filter Results in the Navigate View).
Sort Order	Displays the Sort Order menu; the currently selected option is displayed in the toolbar. The menu has two options: Sort in Ascending Order and Sort in Descending. The Sort Order is a complement to the Sort Field option; the selected sort field for each meta key is ordered in ascending or descending order (see Filter Results in the Navigate View).
Quantification Method	Displays the Quantification Method menu; the currently selected option is displayed in the toolbar. The Quantification Method only applies to the meta key results in the Values panel. It does not apply to the timeline. The drop-down menu contains three options for calculating the quantity (green number in parentheses) for a meta value: Quantify by Event Count, Quantify by Event Size, and Quantify by Packet Count (see Filter Results in the Navigate View). These are applied differently depending on the type of data in view. For packet data: • Quantify by Event Count shows the number of sessions. • Quantify by Event Size shows the size in bytes. • Quantify by Packet Count shows the number of packets. For log data: • Quantify by Event Count shows the number of logs. • Quantify by Event Size shows the size in bytes. • Quantify by Packet Count shows the number of logs.
Save Events	Displays the Save Events menu, in which you can use options to: extract files associated with an event, export the current drill point as a PCAP file, and export the current drill point as a log file (see Export a Drill Point).

Option	Description
Actions	The Actions menu includes actions that you can perform in the Navigate view (see Refining the Results Set). In Version 11.1 and later, the options are Visualize, Go to event in Event Reconstruction, Go to event in Events view, and Print).
Search Events	Enables you to search for text patterns within the current set of events. If you click in the Search field, it shows a drop-down menu with search options. If you click Apply, it saves the selected options and also updates the search options in the Legacy Events view and the Investigations profile (see Search for Text Patterns in the Navigate and Legacy Events Views).
Settings	Displays the settings for the Navigate view (which are also editable in the Profile view) so that you can change Investigate settings without navigating away from the Navigate view. When you change a setting In the Navigate view the setting is also changed in the Profile view (see Configure the Navigate View and Legacy Events View).

Pause/Reload Button and Breadcrumb

The breadcrumb tracks each query as you drill down through the metadata for the service. The following figure is an example of the breadcrumb.



Each query is listed with a drop-down menu in a pipe separated string. The last point is the current point, also called the tip. The icon in front of the breadcrumb allows you to pause the loading of meta values and to reload meta values. The breadcrumb does not include the service name and appears only if a query is in effect. If too many drill points exist for display, the overflow is shown as double angle brackets, >>, at the end of the breadcrumb. Each drop-down menu in the breadcrumb is the same, with slight variation based on the position of the crumb.

The following table describes the controls and menu options in the breadcrumb.

Feature	Description
Pause	Pause and Reload button. Controls the loading of data in the view. It has three possible functions: pause loading, continue loading, and reload.
Navigate Here	Opens the selected drill point in the current Values panel.
Navigate Here (new tab)	Opens the selected drill point in a new tab.
Insert Before	Inserts a query before the current drill point. The Create Filter dialog opens and you can define a custom query to insert in the breadcrumb (see Create a Query in the Navigate and Legacy Events Views).

Feature	Description
Append	Appends a query after the current drill point. The Create Filter dialog opens and you can define a custom query to append to the end of the breadcrumb (see Create a Query in the Navigate and Legacy Events Views).
Remove	Removes the selected drill point from the breadcrumb.
Edit	Opens the selected drill point in the Create Filter dialog so that you can edit the query.
>>	Clicking the angle brackets displays a drop-down menu of the breadcrumb overflow.

(Optional) Debug Information

If you have activated the Show Debug Information setting and the service you are navigating is a Broker, NetWitness, displays the debug information beneath the breadcrumb.

The debug information is the `where` clause from the current query. The only time there is no `where` clause is when the time range is all data and there are no drill points. If the Broker has at least one aggregate service that is offline, the debug information also lists the offline service.

For example:

```
(attachment exists)&&(tcp.dstport = '80')&&(risk.info exists)$$time='2014-05-
04 18:50:00"-"2014-05-09 18:59:59(attachment exists) && (tcp.dstport = '80')
&& (risk.info exists) && time="2014-05-04 18:50:00"-"2014-05-09 18:50:59"
```

In addition, the time taken to load is displayed at the end of each meta key in the Values panel.

Time Banner

Just below the breadcrumb and debug information (if present), the time banner shows the time range used to create the chart. The following figure is an example of the time banner.

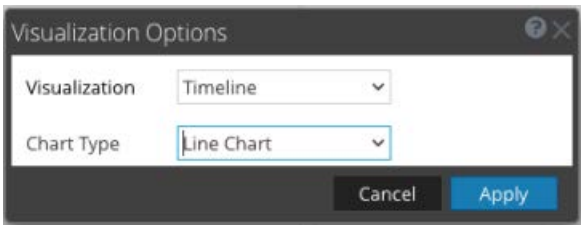


Visualizations

At the top of the Navigate view is a visualization of the current drill point. You can use this to drill into data from the Visualization panel (see Filter Results in the Navigate View). You can show or hide the visualization, and choose one of thevisualization options: Timeline or Coordinates. The Visualization opens initially to the last saved Visualization.

Timeline Chart

The timeline is the count of the number of events that occur at a specific instance. The timeline provides event counts so that you can see if the number of events increases drastically at a given point in time. The timeline displays activity for the specified service and time range as a line chart or a bar chart based on your choice in the Options menu. The second figure illustrates a line chart and third figure illustrates a bar chart.





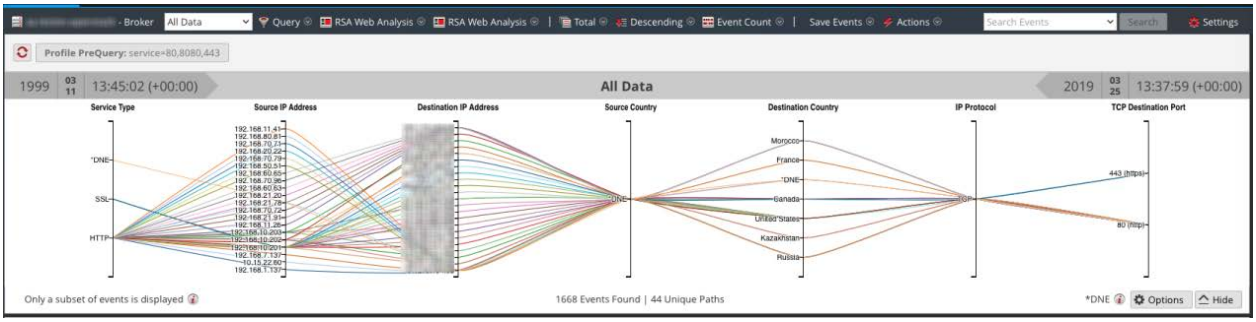
The timeline displays activity for the specified service and time range, as a line chart or a bar chart based on your choice in the Options menu.

Feature	Description
Number of Events (Timeline)	The Y axis of the chart based on thousands of events.
Time Line (Timeline)	The X axis of the chart based on the time the events occurred.
Event point (Timeline)	If you want to explore a specific section, simply select the range from the chart. The new time range will be reflected in the chart.
Investigate (Timeline)	Displays the meta values for the selected subset.
Reset Zoom (Timeline)	To return to the original time range, click Reset Zoom.

Feature	Description
Options	Displays the Visualization Options dialog. Data points can be displayed as a Line chart (default), a Bar chart, or Coordinates chart. When a chart type is select, the relevant options are displayed.
Hide	Collapses the chart.

Parallel Coordinates Chart

The Parallel Coordinates chart is one of the choices in the Options menu for visualizing the current drill point. With Coordinates selected in the Visualization Options dialog, you can select the meta data to be displayed (see Visualize Metadata as Parallel Coordinates). An easy way to view a useful Parallel Coordinates chart is to choose a profile group as shown in the following figure.



Feature	Description
Axes	Each axis is a meta key. The number of meta keys affects the load time for the chart. All meta keys are loaded, but it there the number of events per meta key is limited.
Lines	Lines represent events and they connect values on the axes to show the correlation between multiple meta keys.
Options	Displays the Visualization Options dialog. Data points can be displayed as a Line chart (default), a Bar chart, or Coordinates chart. When a chart type is select, the relevant options are displayed.
Only a subset of events is displayed.	This message is a notification that not all events in the values panel are drawn in the chart. Removing axes or filtering the data in the Values panel can help to display all events.
Events Found | Unique Paths	Displays the total number of events charted versus the number of unique paths charted. Setting the All Meta Keys Must Exist in an Event option redraws the chart so that it is more targeted and legible.
DNE	Indicates that there is no values for this meta key in the event.

In the Visualization Options dialog for Coordinates, you can select the meta keys to chart.

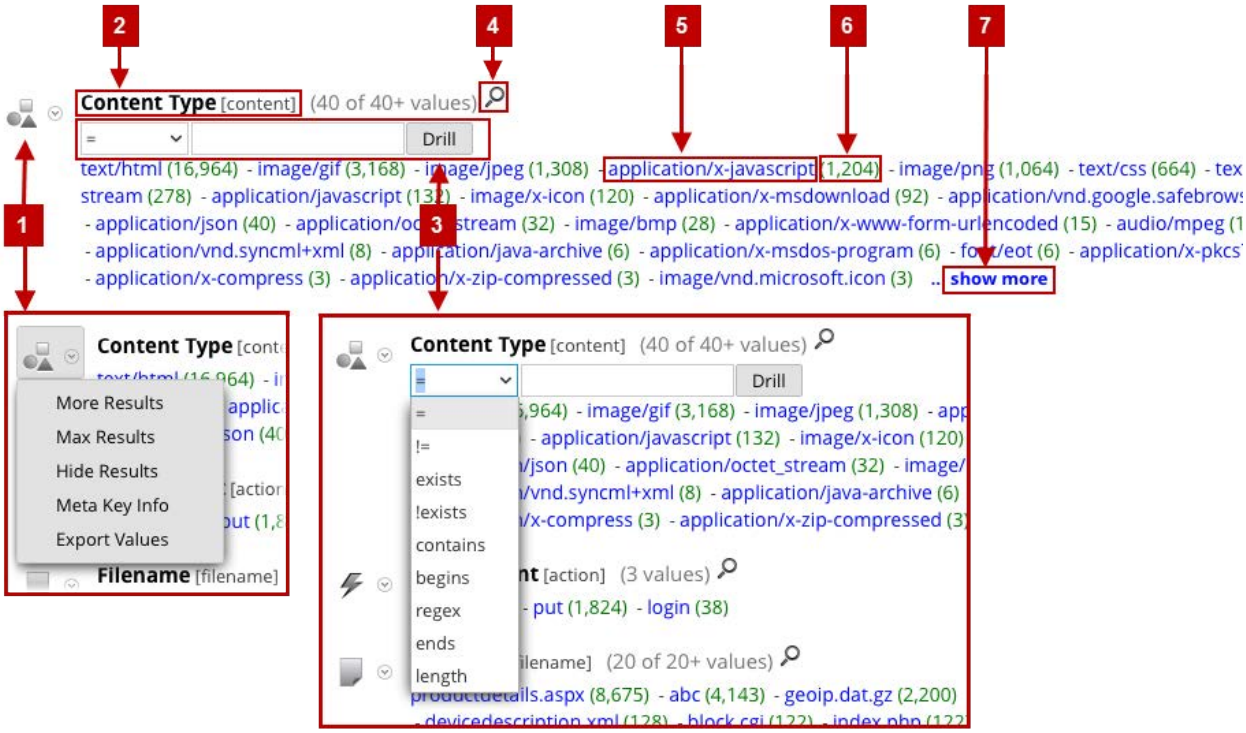
Feature	Description
Visualization selection	Displays a drop-down list of visualization types: Timeline and Coordinates
All Meta Keys Must Exist in an Event	Limits the data represented in the visualization to only those events that include all selected meta keys. This can result in a cleaner, more targeted visualization.
	Displays the Add Keys to Parallel Coordinates Visualization dialog so that you can add axes to the visualization. This is useful if you are looking for relationships between the default meta keys and some additional ones.
	Deletes the selected keys so that they do not appear as axes in the visualization. This can help to make the visualization less cluttered and allow for more data points to be included in the visualization.
	Reverts to the default meta keys for visualization, which consist of all meta keys in the current drill point.
	Controls the display of additional information about the number of selected axes versus the recommended count. This helps to make you aware of possible performance improvements by removing axes.
Axes	Lists the meta keys selected as axes in the visualization.
Cancel	Cancels any changes made to the visualization options.
Apply	Saves the changes made to the visualization options and applies to the current visualization.

In the Add Keys to Parallel Coordinates Visualization dialog, you can select the meta keys or meta groups to use as axes the Parallel Coordinates visualization.

Feature	Description
Visualization selection	Select Keys: Two options for selecting meta keys are: • From Default Meta Keys • From Meta Groups Each option offers a drop-down list from which to select.
With the Selected Meta Keys...	The options for the method of adding meta keys allow you to: • Replace the current list of keys • Append to the current list of keys • Insert at beginning of the current list of keys
Cancel	Closes the dialog and does not add any keys.
Add	Closes the dialog and adds the selected keys as specified.

Values Panel

The major feature of the Navigate view is the Values panel, which presents meta keys and meta values found in the service being investigated. Procedures for analyzing data in the Values panel are provided in Filter Results in the Navigate View.



Note: Title, values, and counts for non-indexed meta keys are not drillable; the Values and counts are shown in black.

1 The meta keys in the Values panel have drop-down menus, which offer actions that can apply to that meta key. You can use these to change the way the results for the meta key are displayed in the current view. Changes made to meta keys are displayed in the current view and persist until you refresh the page or select a new service in the Navigate view toolbar. See Drill into Data in the Values Panel

Refresh reverts to the current view of meta keys as defined in the Manage Default Meta Keys dialog (see Manage and Apply Default Meta Keys in an Investigation). If you have never made modifications in the Manage Default Meta Keys dialog, NetWitness, a refresh restores the default meta keys from the core service.

- More Results
- Max Results
- Hide Results
- Meta Key Info

	• Export Value
2	The name of the meta key for which values are listed. In Version 11.3 and later, the user friendly name of the meta key is displayed with the index file name of the meta key following in brackets. For example **Content Type [content]** gives the user friendly name of the `content` meta key with the index file name in parentheses. For meta groups, the name of the group is given in plain English with the meta group name following in parentheses. This is an example of a meta group name as it would appear in the Values panel: **All User Keys [users.all]**.
3 and 4	Clicking on an indexed meta key opens the Search dialog in which you can enter a filter for the current meta key. The search function is not available for non-indexed meta keys, and is based on the actual meta value rather than the alias. Drilling in the Search dialog using aliases is not supported. NOTE: Check with your administrator to obtain a list of aliases used for a meta key in Investigation. When an alias is used, this search dialog does not provide results. Instead, you must query the meta key using the Right-click query capability or the Query dialog.
5	The meta value associated with the found meta key. These are listed in order by meta value name or by the count of events in which the meta value was found, according to your preference.
6	The number of events that include the meta value.
7	The number or values rendered is specified by the Render Threads value in the Investigation Preference settings. In the example above, the meta key is **Content Type**, and 40 of 40+ values are currently displayed. You can display additional values by clicking **...show more**. The number of instances found for a particular meta in the session.

Values Panel Loading Behavior

The default view is for the last 3 hours of collection, using the default meta keys and non-indexed meta keys closed. The meta keys within the meta groups are displayed in the order that NetWitness queries the keys. As the data loads into the Values panel, NetWitness is optimized to show partial results, loading progress, and service status as the data loads.

The loading behavior is determined by several configuration settings. The highest level settings are configured by the administrator for each user. These are:

- The maximum amount of time allowed for this user to run a query (Query Timeout).
- The limit at which NetWitness stops counting the number of meta values in a session (Session Threshold). If a threshold is set for a session, the Navigate view shows that the threshold was reached and the percentage of results loaded. Any session that does not show a percentage is accurate and was processed to completion. If there is a percentage, that reflects how much processing was completed. The percentage displayed is estimated by extrapolating from the value at the time processing finished, considering the amount of work remaining. Larger percentages are generally more accurate because they require less extrapolating
- The limit at which NetWitness stops counting the number of meta values in a session (Session Threshold). If a threshold is set for a session, the Navigate view shows that the threshold was reached and the percentage of query time used to reach the threshold.

Note: The values for non-indexed meta keys take longer to load in the Values panel. To optimize loading, NetWitness does not open non-indexed meta keys by default. Refer to Manage and Apply Default Meta Keys in an Investigation for a detailed description of non-indexed meta keys in Investigation.

When you have launched an investigation of a service, NetWitness displays results in the Values panel.

1. NetWitness loads meta keys and meta values in the Values panel. For each meta key load, the stages of load are:
   a. **Waiting to Be Loaded** or **Closed**. If Closed, no data for that key is loaded.
   b. **Loading**
      i. **Loading progress**: NetWitness is receiving and displaying progress messages.
      ii. **Partial results**: NetWitness is receiving values messages and partial results are displayed in the Values panel.
   c. **Load Complete**: All results are finished loading.
2. As each meta key load is completed, and final values are displayed, the next meta key is started. The number or values rendered for each meta key is specified by the Render Threads value in the Investigation Preference settings. Loading continues until all keys to be loaded have finished.
3. If **Show Debug Information** is active and the service you are navigating is a 10.4 or later Broker, NetWitness displays load time information beneath the values for each meta key and displays additional load details for the aggregated services. NetWitness also displays the debug information beneath the breadcrumb.

Iterative results

Iterative results provide feedback on the status of queries within the interfaces to provide additional context for how long the data load will take and if any service data is missing. For example, if you are querying a Broker that is aggregating from two Concentrators, NetWitness starts displaying the results from the first Concentrator as soon as it is available, even if the second Concentrator is still waiting for results.

Iterative results also include a notification that service data is missing because the service is unreachable.

Partial results

When partial values from the Core service are returned but not completed, a message at the end of the meta key listing shows the progress of values loaded. For example, Currently looking at 38 ip.src values 71% indicates that loading of values for the meta key is 71% complete.

Debug Information

If the Show Debug Information setting is in effect, a field at the end of the values displays the status for the different systems against which you are querying within NetWitness. For example, when you are querying against a 10.4 broker pulling from multiple concentrators, NetWitness displays the status of the query on each of the Concentrators, which provides insight into the relative speed of data loading from each of the Concentrators. Each service that participated in the query is listed with the total elapsed time for the query.

Each service that participated in the query is listed with the total elasped time for the query. In the example above, two services returned in 3.207 seconds, localhost:50005 took 2 seconds to return the results. In addition, the where clause of the query is displayed below the breadcrumb. You can copy this syntax directly into an application rule or Reporting where clause of a rule.

Load Complete

For each meta key, there is a list of values (blue text) and counts (green text) found in the current drill point. When you click a value to drill down into a subset of the currently selected data, the display is updated and the new drill point is recorded in the breadcrumb. You can specify the sorting and quantification methods for the values list using the option in the toolbar.

Query Dialog

In the Navigate view or Legacy Events view, you can create a query rather than clicking through the meta keys and values to drill down into the meta data. The dialogs for creating a query offer syntax help with drop-down lists of applicable meta keys and operators. To access this dialog in the **Navigate** or **Legacy Events** view toolbar, select **Query**.

What do you want to do?

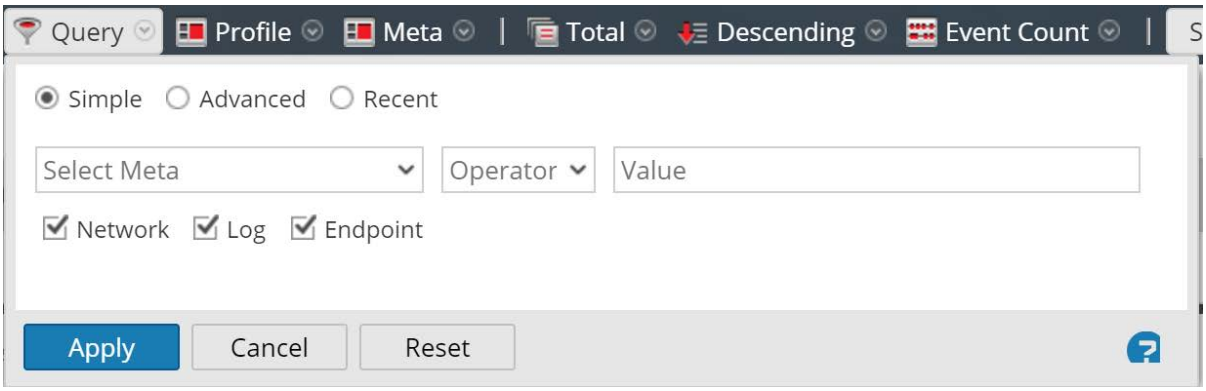
User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look

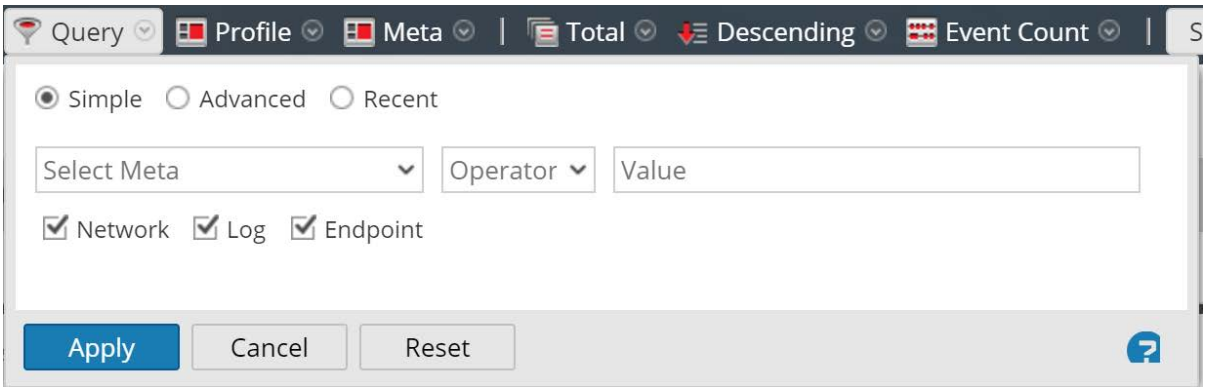


The Query dialog has three views:

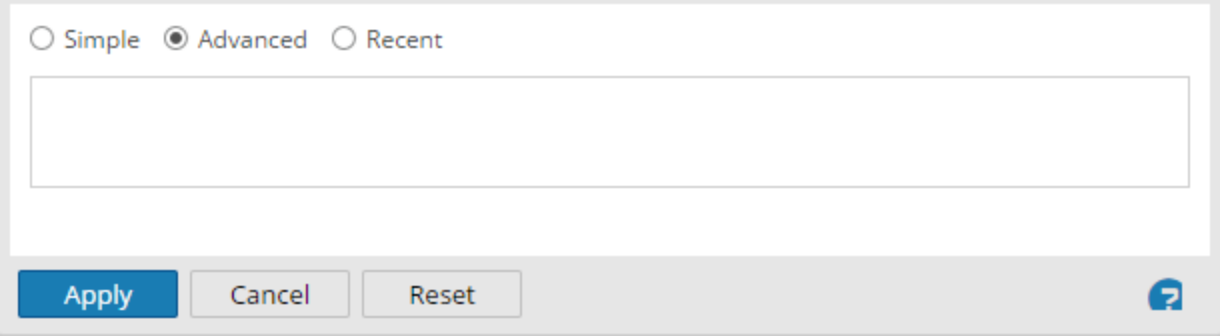
- Simple
- Advanced
- Recent

In the Simple view, you can create a query using the options displayed in the dialog. In the Advanced view, you can create a query without guidance. In the Recent view, you can select a query from a drop-down list of recent queries.

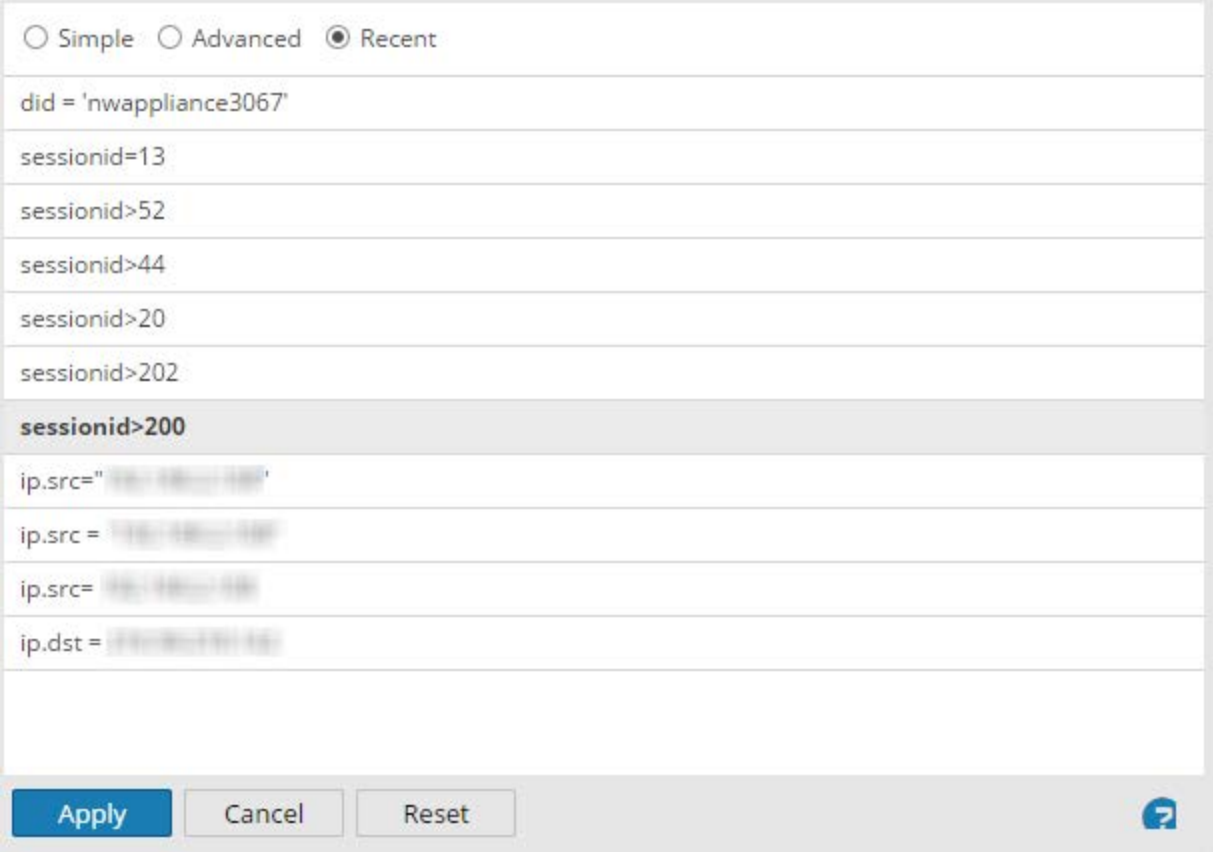
Simple View



Advanced View



Recent View



The following table describes features of the Query dialogs.

Feature	Description
Select Meta	Displays a drop-down list of meta groups.
Operator	Displays a drop-down list of operators (=,NetWitness!=,NetWitnessexists,NetWitness!exists)

Feature	Description
Value	Allows you to enter a value to complete the query.
Network	Limits the query to packets if Log is not selected.
Log	Limits the query to logs if Network is not selected.
Query box	Allows you to enter a query in the Advanced view. When you begin typing, a drop-down list of available meta keys for the service is displayed, then a drop-down of operators is displayed as you type. If the expression currently entered in the query box is invalid, a warning appears near the box. When the query is valid, the warning is removed.
Query list	Allows you to select a query from a list of recent queries in the Recent view. Double-clicking a query automatically applies it.
Apply	Applies the new query to the current Investigation view.
Cancel	Closes the dialog without applying changes.
Reset	Resets all fields.

Query Profiles Dialogs

Query profiles offer a quick and easy way to define a meta group, column group, and a limiting filter (pre-query condition) that you can apply in the Navigate view, the Events view, and the Legacy Events view (see Use Query Profiles to Encapsulate Common Areas for Investigation). The same query profiles are shared between all views, and they are available in the Springboard (Version 11.5) for use in panels. Private query profiles created in the Events view are only available in the Events view for the analyst who created them.

Each query profile specifies a meta group, column group, and sometimes includes a pre-query condition appropriate for the type of investigation.

In a query profile:

- The meta group defines the meta keys that are queried (see Use Meta Groups to Focus on Relevant Meta Keys).
- The column group defines which meta keys from the meta group are displayed as columns in the Events list. (see Use Columns and Column Groups in the Events List).
- When the query profile is in effect, the optional pre-query conditions add a limiting filter in the query bar. You can edit or delete the limiting filter and then create additional filters for your query (see Filter Results in the Events View)

You can manage profiles in the Manage Profiles dialog, the Create Query Profile dialog and the Query Profile Details dialog.

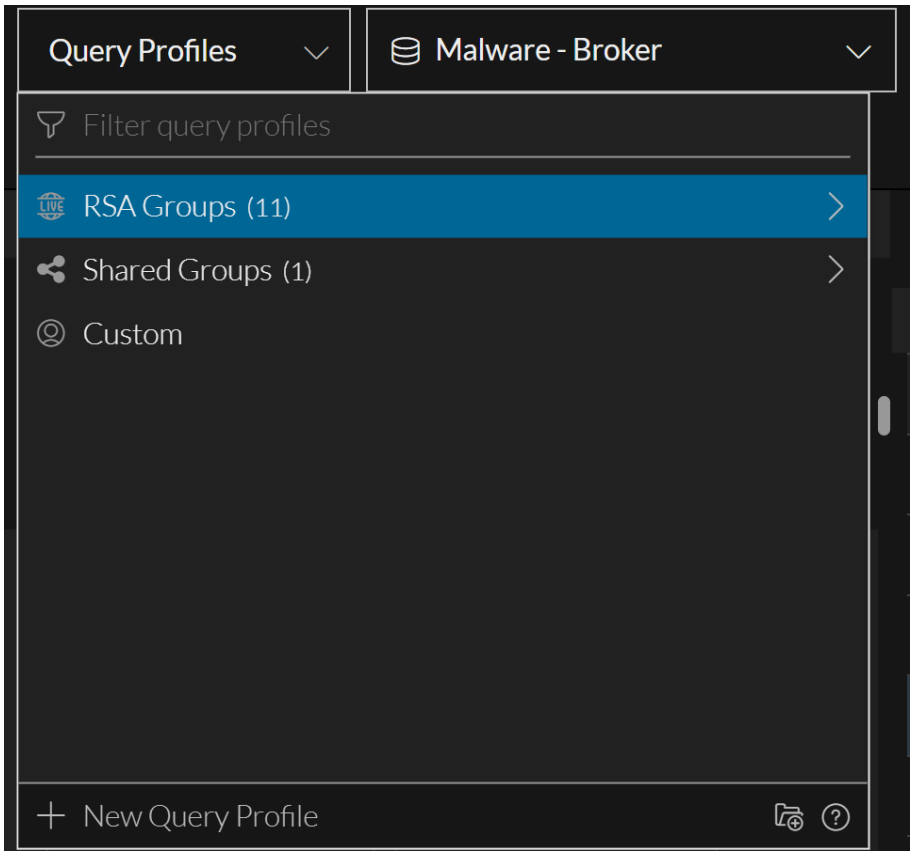
- The Manage Profiles dialog is for the Navigate view, the Legacy Events view (Version 11.4 and later) , and the Events view (Version 11.3 and earlier). To access this dialog, select **Profile > Manage Profiles** in the **Navigate** or **Legacy Events** view toolbar.
- The Create Query Profile dialog is for the 11.4 and later Events view. To access this dialog, select **Query Profiles > New Query Profile** in the **Events** view query bar.
- The Query Profile Details dialog is for the 11.4 and later Events view. To access this dialog, select **Query Profiles** in the **Events** view query bar, then click the edit icon () next to a custom profile name.

Related Topics

- How NetWitness Investigate Works
- Use Query Profiles to Encapsulate Common Areas for Investigation
- Navigate View
- Events View
- Legacy Events View

Quick Look - Query Profile Menu, Create Query Profile Dialog, and Query Profile Details Dialog

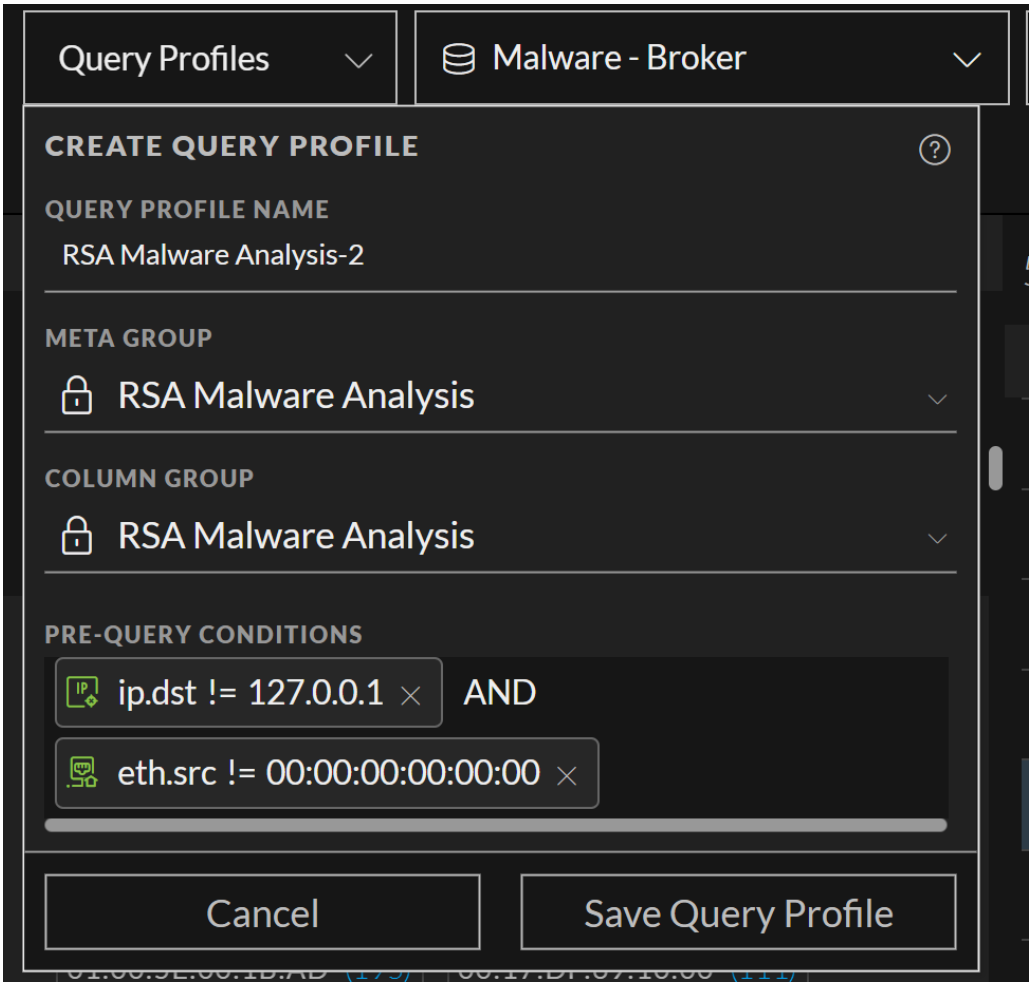
This section introduces the Query Profile menu, Query Profile dialog, and the Query Profile Details dialog. The following figure is an example of the Query Profiles menu and the table describes the options. The example on the left has built-in profile highlighted so that the information icon is visible. The Version 11.4 menu is on the left and the Version 11.5 menu is on the right.



Feature	Description
Visibility Options	Control the types of query profiles that are visible in the list. You can use any combination of the visibility options: **Private**, **Shared**, or **RSA** (blue = selected, black = not selected). Initially none of the buttons are selected and all profile types are visible. This is the same result as if all three buttons are selected. The visibility options work together with text in the Filter Query Profiles field. If the visibility option is hiding built-in profiles (which include "RSA" in the name) and you search for a name that contains "RSA," the list is empty. Private = display private groups that only you can manage Shared = display shared groups that anyone in your organization can manage RSA = display built-in groups that only RSA can manage

Feature	Description
Filter Query Profiles	Filters the list of query profiles as you type text so that only profile names that contain that text are displayed.
Query Profile List	The list of profiles consists of custom and built-in profiles, which are distinguished by the icons that precede the name. In the example, RSA Email Analysis-1 and RSA Email Analysis-2 are custom profiles. The RSA Email Analysis is a built-in profile.
New Query Profile	Displays the Create Query Profile dialog, where you can create a custom profile.

The Create Query Profile dialog, shown in the figure on the left, allows you to define a custom profile. The figure on the right illustrates the Query Profile Details dialog, in which you can edit a custom profile. The table describes the fields and options in the dialogs.

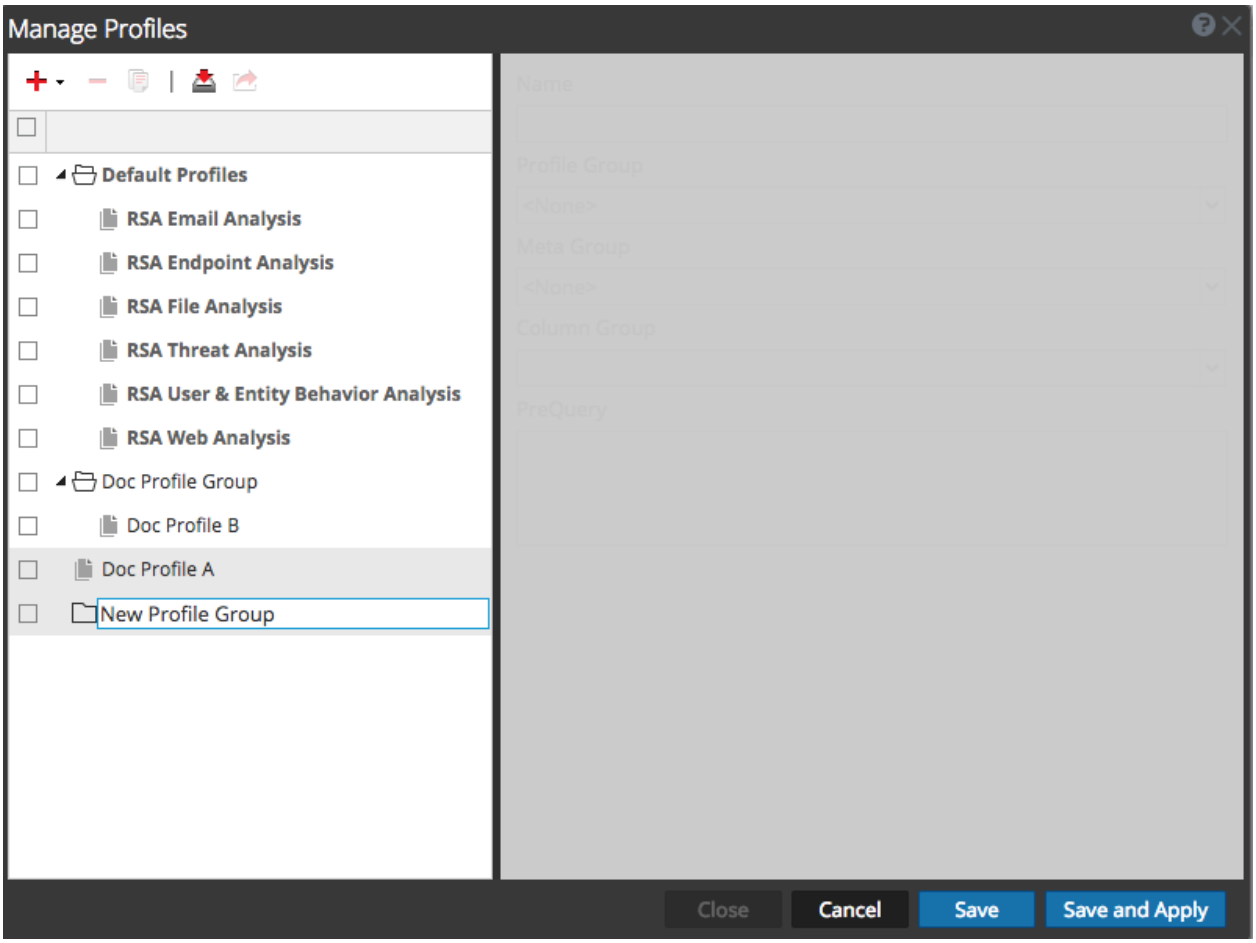


Feature	Description

Feature	Description
	Creates a clone of the meta group so that you can edit a copy. This is useful if you want your own copy of a built-in group, a shared copy of a private group, or a private copy of a shared group.
	Deletes the custom profile in the Query Profile Details dialog. This action is irreversible and applies globally; the profile is no longer available to anyone who is using the profiles on this service.
Query Profile Name	Displays the name of the profile. The name must be unique and contain fewer than 64 characters. You can edit the name in a custom profile.
Column Group	Displays a drop-down menu listing available column groups, with the currently selected column group from the Events list already selected. You can change the column group in a custom profile.
Pre-Query Conditions	Defines a limiting filter for results in the Events view. If you had a query active in the query bar when you began to create the new profile, the active query is added to the pre-Query field. In a custom profile, you can delete the prepopulated pre-query condition and type additional text for a text search or additional filters in the Pre-Query Conditions field. This is an example of a pre-query condition: `'service=80,25,110'`.
Close button	Closes the dialog.
Save Query Profile	For the Create Query Profile dialog only, saves the new profile.
Reset	For the Query Profile Details dialog only, reverts the edited profile to the last saved state.
Update Query Profile	For the Query Profile Details dialog only, applies changes to an edited profile.
Select Query Profile	Applies the query profile.

Quick Look - Manage Profiles Dialog

This is an example of the Manage Profiles dialog showing several profile groups.



The Profile panel on the left side of the dialog displays available profiles and allows you to add, delete, import, and export profiles. The following table describes the fields in the Profile panel.

Field	Description
	Adds a new profile using the Settings panel on the right side of the Manage Profiles dialog.
	Deletes the selected profile. A confirmation dialog is displayed before the profile is deleted.
	Creates a copy of the selected profile.
	Displays the Profile Import dialog, where you can upload a file.
	Exports the selected profile to your computer.
Profile Name	Lists all profile names.

The Settings panel on the right side of the dialog offers options to configure profiles. It can only be used when one profile is selected. The following table describes the fields in the Settings panel.

Feature	Description
Name	Displays the name of the profile.
Meta Group	Displays a drop-down menu listing available meta groups.
Column Group	Displays a drop-down menu listing available column groups. The OOTB column groups and these three groups are available by default: • List View • Detail View • Log View
PreQuery	Defines a limiting query for filtering Investigate results. This query is used when the associated profile is activated and the preQuery applies to any queries used in the Navigate and Events views. This is an example of a preQuery: `'service=80,25,110'`.

The following table describes the buttons.

Field	Description
Close	Closes the dialog.
Cancel	Cancels all changes.
Save	Saves all changes.
Save and Apply	Saves and applies all changes immediately.

Generate Springboard Panel Dialog

In the Generate Springboard Panel dialog, you can create a Springboard panel from selected query in the Events view. You can add any number of filters in the query search bar and convert them into Springboard panels. The Springboard panels are then available to analysts for detection and monitor the results.

To access this dialog, while investigating a service in the **Investigate** > **Events** view, add a query on the query search bar >  > **Generate Springboard Panel** from the toolbar.

IMPORTANT: Ensure that you first create a custom private board in Springboard.

What do you want to do?

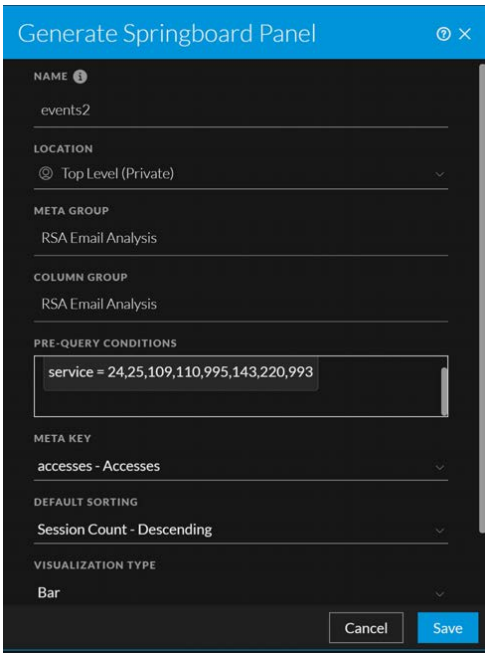
User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View

Related Topics

- Use Query Profiles to Encapsulate Common Areas for Investigation
- Managing the Springboard

Quick Look - Generate Springboard Dialog

This is an example of the Generate Springboard dialog.



The following table describes the fields in the Generate Springboard Dialog view.

Feature	Description
Name	(Required) Specifies a name to identify the panel. In the example, the name is Sample Incident. You can provide a name that clearly identifies the nature of events that will be added to this incident
Meta Group	Displays the currently selected meta group from the filter query already selected.
Column Group	Displays the currently selected column group from the filter query already selected.
Location	Displays the location where the defined query is saved.
Pre-Query Conditions	Defines a limiting filter for results in the Events view. If you had a query active in the query bar when you began to create the new Springboard panel, the active query is added to the pre-Query field. This is an example of a pre-query condition: `'service=80,25,110'`.
Meta Key	Displays a drop-down listing meta keys available for the service.
Default Sorting	Displays a drop-down listing sorting options.
Visualization Type	Displays a drop-down listing Visualization type options: • **Donut** • **Bar**

Feature	Description
Visualization Metric	Displays a drop-down listing available Visualization metric options.
Cancel	Closes the dialog without applying changes.
Save	Saves the changes.

Settings Dialogs for Investigate Views

NetWitnessVersion 11.0 has two settings dialogs, one for the Navigate view and one for the Legacy Events view. With the addition of the settings dialog for the Events view in Version 11.1, Investigate has three settings dialogs.

The settings in these dialogs are a subset of the Investigation settings made in the Profiles > Preferences panel > Investigation. Analysts can save time by editing these settings within the Investigate view. If you change a setting here, the same setting is changed in the Profiles view, and if you change a setting in the Profiles view, the same setting is changed here.

To access this dialog, go to the **Navigate** or **Legacy Events** view, and select the **Settings** option in the toolbar.

The settings in the Events view have no corresponding settings in the Profiles > Preferences panel.

Related Topics

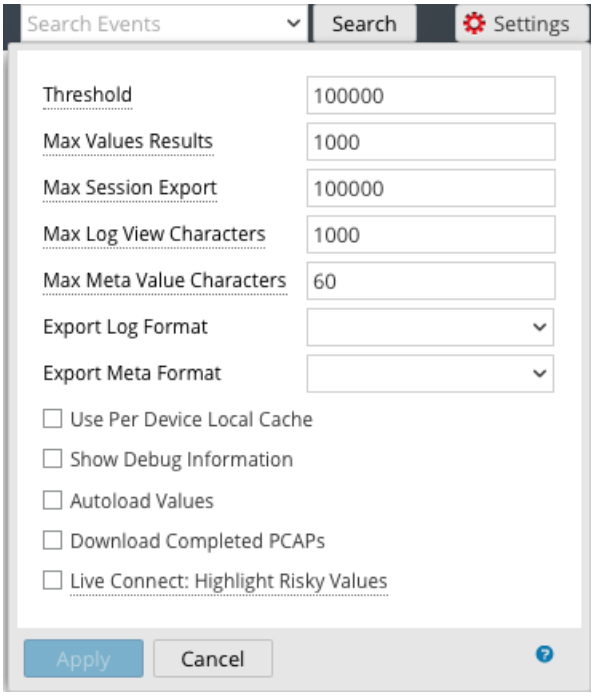
- How NetWitness Investigate Works

Quick Look

This is a quick look at the settings dialog for the Navigate view, Legacy Events view, and Events view.

Navigate View Settings Dialog

The following figure illustrates the Navigate view Settings dialog. The settings that influence performance when loading values in the Values panel have default values based on common usage, and individual analysts can adjust these settings for their own investigations. The following table describes the features.

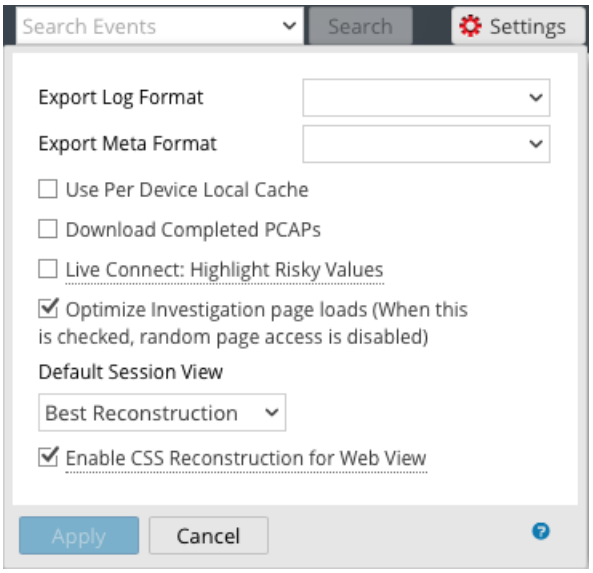


Feature	Description
Threshold	Sets the threshold for the maximum number of sessions loaded for a meta key value in the Values panel. A higher threshold allows accurate counts for a value, and also causes longer load times. The default value is **100000**.
Max Values Results	Sets the maximum number of values to load in the Navigate view when the Max Results option is selected in the Meta Key menu for an open meta key. The default value is **1000**.
Max Session Export	Sets the maximum number of sessions that can be exported. The default value is **100000**.
Export Log Format	Sets the file format of exported logs. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format.

Feature	Description
Export Meta Format	Sets the file format of exported meta values. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format. **Note:** If you upgrade to version 11.5.2, the **Export Meta Format** preference is not retained and is reset to blank. You must re-configure this value after you upgrade to version 11.5.2.
Use Per Device Local Cache	When the checkbox is cleared, Investigate sends a fresh query to the database rather than displaying cached data in the Investigate views after the initial load. If the checkbox is selected, Investigate uses the data from local cache.
Show Debug Information	This option controls the display of the `where` clause beneath the breadcrumb in the Navigate view and the elapsed load time for each aggregated service on a Broker. When the checkbox is selected, the debug information is displayed. The default value is **Off** (checkbox cleared).
Autoload Values	This option controls automatic loading of values for the selected service in the Navigate view. When the checkbox is selected, values are automatically loaded when you select a service to investigate. When the checkbox is cleared, Investigate displays a **Load Values** button, allowing the opportunity to modify options. The default value is **Off**.
Download Completed PCAPs	This setting automates the downloading of extracted PCAPs in Investigate so that you do not have to manually download and open extracted PCAP files in an application, such as Wireshark, that can handle viewing data in a PCAP form. When the checkbox is selected, the option is enabled. The default setting is disabled (checkbox is cleared).
Live Connect: Highlight Risky IPs	If the checkbox for this option is cleared, all the meta values that have context available in Live Connect are highlighted in the Navigate view Values panel. If the checkbox is selected, among the values that have context in Live Connect, only those values deemed Risky/Suspicious/Unsafe by the community are highlighted. By default this option is disabled (checkbox cleared).
Apply	Applies the settings immediately and they are visible the next time you load values. The same changes are also applied in the Profiles view.
Cancel	Cancels the editing operation and closes the dialog, leaving the settings unchanged.

Legacy Events View Settings Dialog

The following figure is an example of the Settings dialog for the Legacy Events view, and the table describes the features.

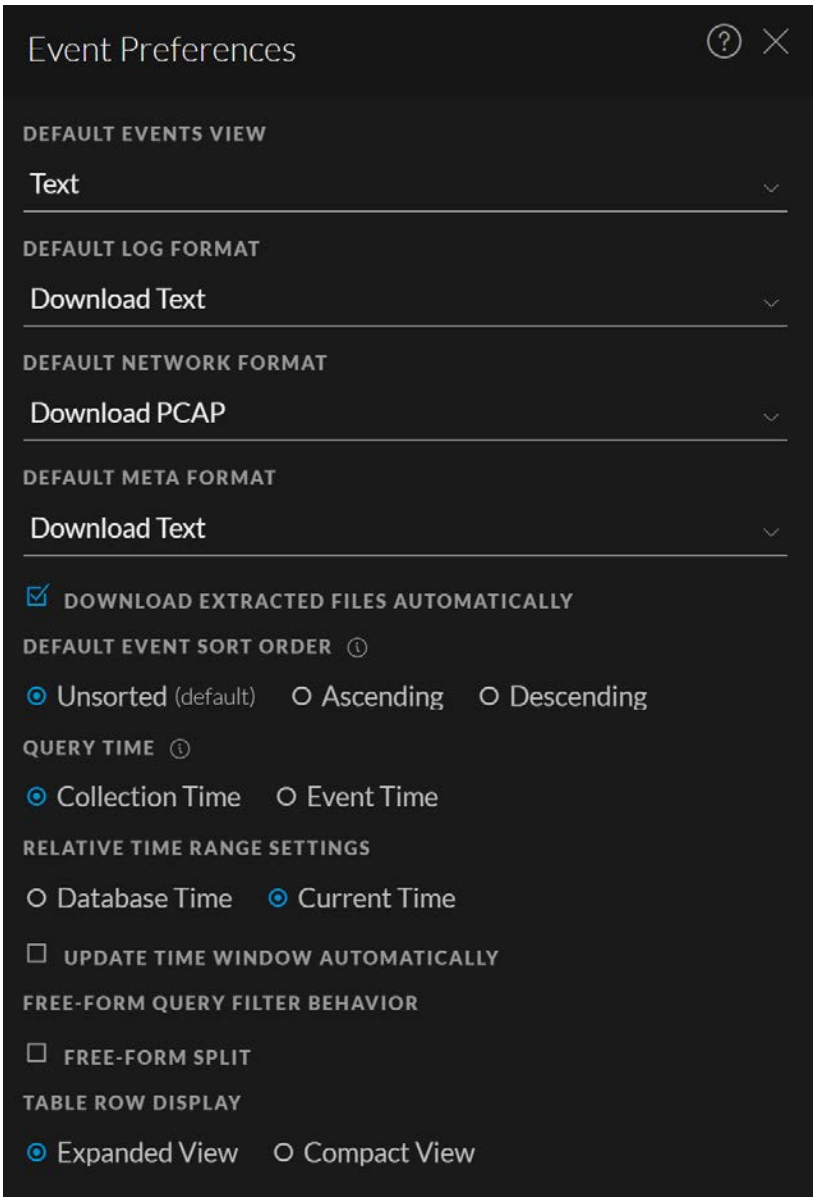


Feature	Description
Export Log Format	Sets the file format of exported logs. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format.
Export Meta Format	Sets the file format of exported meta values. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format.
Download Completed PCAPs	This setting automates the downloading of extracted PCAPs in Investigate so that you do not have to manually download and open extracted PCAP files in an application, such as Wireshark, that can handle viewing data in a PCAP form.
Live Connect: Highlight Risky IPs	When the checkbox is selected, Investigate uses a filter to fetch only IP addresses that are considered as risky by RSA community. When the checkbox is cleared, NetWitness displays all IP addresses. By default, this option is disabled (checkbox cleared).

Feature	Description
Optimize Investigation page loads	Sets a paging option. When optimized, results are returned as quickly as possible, sacrificing the original ability to go to a specific page in the event list. Clearing this checkbox changes the Events list pagination to allow you to go to a specific page in the list (or to the last page). The default value is enabled (checkbox selected).
Append Events in Event Panel	This option affects paging in the Legacy Events panel and in prior releases was located in the Navigate view settings dialog. When the checkbox is selected, the next group of events is appended to the already displayed events. When cleared, the previous page of events is replaced by the next page. The default value is Off (checkbox cleared).
Default Session View	Selects the default reconstruction type for the initial reconstruction in the Events view. The default value is **Best Reconstruction** in which events are reconstructed using the reconstruction method most appropriate to the event.
Enable CSS Reconstruction for Web View	This setting controls how web content reconstruction is performed. If enabled, the web reconstruction includes cascaded style sheet (CSS) styles and images so that its appearance matches the original view in a web browser. This includes scanning and reconstructing related events, and searching for style sheets and images used in the target event. The option is enabled by default. Clear the checkbox to disable this option if there are problems viewing specific websites.
Apply	Applies the settings immediately and they are visible the next time you view events. The same changes are also applied in the Profiles view.
Cancel	Cancels the editing operation and closes the dialog, leaving the settings unchanged.

Events View Preferences Dialog

Beginning with Version 11.1, the Events view has user preferences that you can configure in the Events view > Event Preferences dialog. These settings persist so that they are applied each time you log in and go to the Events view. The following figures are examples of the dialog for Version 11.3 and Version 11.6. The table below describes the options.



Feature	Description
Default Events View	Selects the default event analysis view that is displayed every time you open the Events view. For example, if you select File, the File Analysis panel is highlighted and displayed every time you investigate an event in the Events view. These are following options: • **Text** : View and analyze the raw text payload of an event. • **Packet** : View and interactively analyze the packets and payload of an event. • **File** : View a list of files and download one or more files in an event.

Feature	Description
Default Log Format	Selects the default format for downloading logs: • **Download Log** or **Download Text**: Raw log (log) format. • **Download CSV**: Comma-separated values (CSV) format. • **Download XML**: The Extensible Markup Language (XML) format. • **Download JSON**: The JavaScript Object Notation (JSON) format.
Default Packet Format or Default Network Format	Selects the default format for downloading packets. • **Download PCAP**: To download the entire event as a packet capture (*.pcap) file. • **Download All Payloads** or **Download Payloads**: To download the payload as a *.payload file. • **Download Request Payload**: To download the request payload as a *.payload1 file. • **Download Response Payload**: To download the response payload as a *.payload2 file.
Default Meta Format	Selects the default format for downloading metadata: • **Download CSV**: Comma-separated values (CSV) format. • **Download JSON**: The JavaScript Object Notation (JSON) format. • **Download Text**: Plain text format. • **Download TSV**: Tab-separated values (TSV) format.
Time Format for Query	The Events view can display results based on the database time or the current clock time. **Note:** (Version 11.6) **Current Time** is the default for **Relative Time Range Settings**. In previous versions, **Database Time** was the default value. Make a note that this may cause time range mismatch between **Events View** (using **Current Time** as default) and **Navigate View** (using **Database Time** as default). This change does not affect the existing users and is applicable only to the new users. When **Database Time** is selected, the start and end time for a query is based on the time that the event was captured (collection time). When **Current Time** (labeled **Wall Clock Time** in Version 11.3 and earlier) is selected, the query is executed using the end time based on the current browser time; the start time is calculated based on that end time and the time range.

Feature	Description
Event Sort Order (Version 11.4 and Later)	Sets the sort sequence by collection time for the events listed in the Events panel. If results exceed the events limit, not all events can be loaded. The portion of returned events loaded in the Events panel matches the sort order preference: the oldest portion of events is loaded when Ascending order is selected, and newest portion of events is loaded when Descending order is selected. A change to this setting becomes effective the next time you submit a query. **Unsorted**: Default sorting method for Version 11.4.1. To list events as processed by the Core services. Unsorted is faster because it streams back the events as soon as a match is found versus waiting for all Core services to respond and then displaying them in the chosen order. **Ascending**: Default sorting method for Version 11.4. To put the events with the earliest collection time first in the list. **Descending**: To put the events with the latest collection time first in the list. When investigating logs, you may want to change the sort sequence to latest collection time first.
Download Extracted Files Automatically	Enables the automatic download of files if they are in the selected default format in the **Default Log Format** and **Default Packet format** fields set in the Event Preferences dialog. Select the checkbox to enable downloading the selected format automatically to local file system. Otherwise, the download job goes to the job queue, and you can download it manually.
Update Time Window Automatically	(Version 11.3 and later) Enables automatic update of the time range window in the query bar when the service is polled (at one minute intervals) so that fresh results are sent. The default setting is disabled. When the checkbox is selected, as the time range is updated, the  (Submit Query) button is activated so you can click to get the fresh results. When the checkbox is cleared, the automatic update is disabled keeping the time range window in the breadcrumb synchronized with the current results.

NetWitness® Platform XDR

Version 12.0

Licensing Management Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

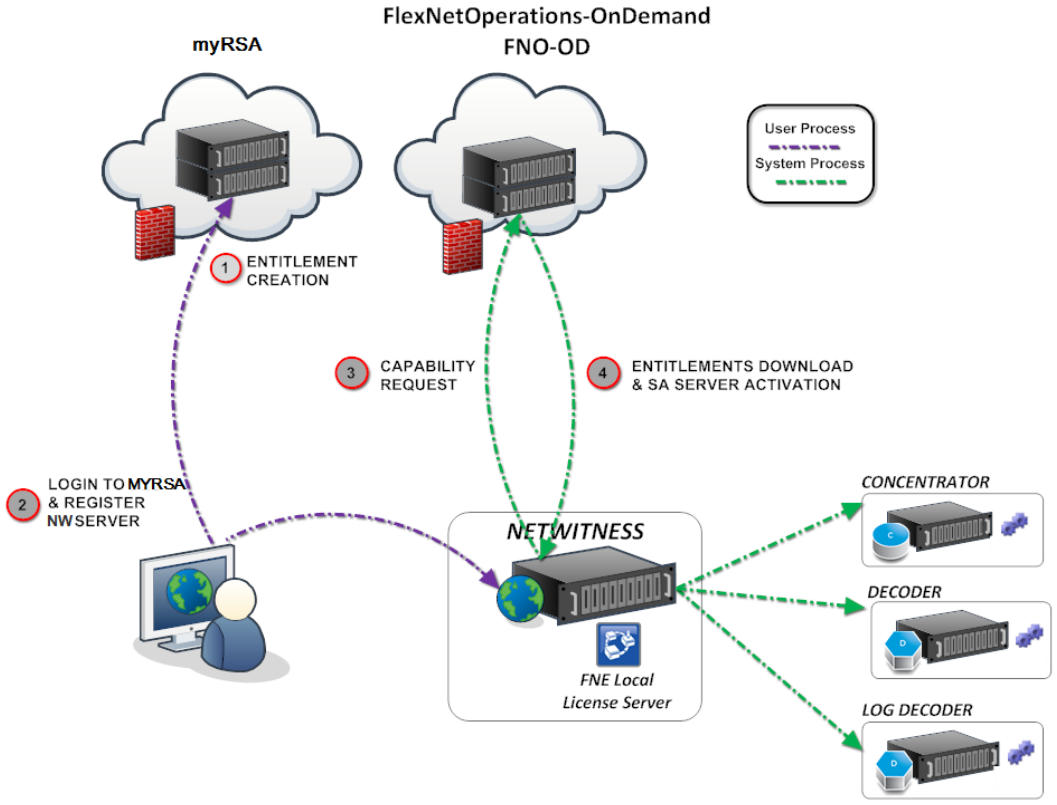
Contents

Entitlement Capability Implementation **5**
NetWitness Platform Licensing Overview **8**
Obtain License Server ID from NetWitness Platform User Interface **10**
Access myRSA **11**
License for NetWitness Endpoint 4.4.0.x Agents **13**
Register the Server (Online Registration) **16**
Map Entitlements **21**
View Current Licenses 23
View and Manage Licenses 23
Register the Server (Offline Capability Request) **25**
Prerequisites 25
Download an Offline Request for Submission to myRSA 25
Upload an Offline Capability Response to NetWitness Platform 28
Refresh Licenses 28
License Types **30**
Choosing a License Type 30
Throughput Licenses 31
Appliance Licenses 32
User and Entity Behavior Analytics Licenses 32
Endpoint Licenses 32
NetWitness Endpoint 4.4.0.x License 33
Out-of-the-Box Trial License 33
Configure License Support For Multiple NetWitness Servers **34**
Configure License Notifications **35**
About Out-of-Compliance Banners **37**
Troubleshoot Licensing **39**
Out-of-Compliance Banners 39
Common Log and Configuration Files 39
NetWitness Server Problems 39
Wrong License Mapping Issue 40
License Details **42**
Related Topics 43
Usage Trends 45

Reassign Service Licenses **50**
View and Export Usage Stats **53**
Settings **55**
Out-of-Compliance Reference **58**
Out-of-Compliance State 59
License Approaching Out-of-Compliance 60

Entitlement Capability Implementation

This topic describes how licensing of appliances and services is implemented in NetWitness. The entitlement capability leverages myRSA (https://my.rsa.com/) as the mechanism for entitlement delivery.



Key	Description
1	**Entitlements Created and Available to Customer.** After a customer order is processed, the entitlements (licenses) become available in myRSA. The entitlements are tied to an individual account.
2	**Register NetWitness Server on myRSA and Map Entitlements to the Local License Server (LLS).** • Customers log on to myRSA and view the entitlements to which they have access within their account. • Customers map entitlements to their Local License Server using the License Server ID (displayed in the NetWitness (Admin) **> System > Info** panel). The License Server ID is used only for mapping entitlements to a Local License Server and does not pertain to appliance activation.

Key	Description
3	**Synchronize the Server and Download Mapped Entitlements.** There are two methods for customers to synchronize with FlexNet Operations-On Demand (FNO-OD) and download the mapped entitlements to their LLS. • **Sites with Internet connectivity.** If the LLS has Internet connectivity, the LLS attempts to synchronize with FNO-OD every 24 hours over HTTP (TCP-80). Customers with Internet connectivity can also perform on-demand synchronization, using the **Refresh** option in the  (Admin) **> System > Licensing** panel on the NetWitness Server. For more information, see Register the Server (Online Registration). After the synchronization, endpoint entitlements that were mapped to the Local License Server on the NetWitness Platform appliance are synchronized on hourly basis by port 443. • **Sites in closed environments.** Customers can synchronize the mapped entitlements by downloading a capability request and importing it on the NetWitness Server. For more information, see Register the Server (Offline Capability Request). After the synchronization, entitlements that were mapped to the Local License Server on the NetWitness appliance are synchronized, but the entitlements have not been used in any way. For example, if you have purchased 10 Decoders and 10 Concentrators, 10 of 10 Decoder entitlements and 10 of 10 Concentrator entitlements would be available on the NetWitness Server.

Note: FNO-OD is the license server in the cloud on myRSA. The URL is rsasecurity.subscribenet.com. The customer's firewall must allow communications between this URL on port 80 (whatever it resolves to when using lookup or whois) and the NetWitness IP address.

NetWitness Platform Licensing Overview

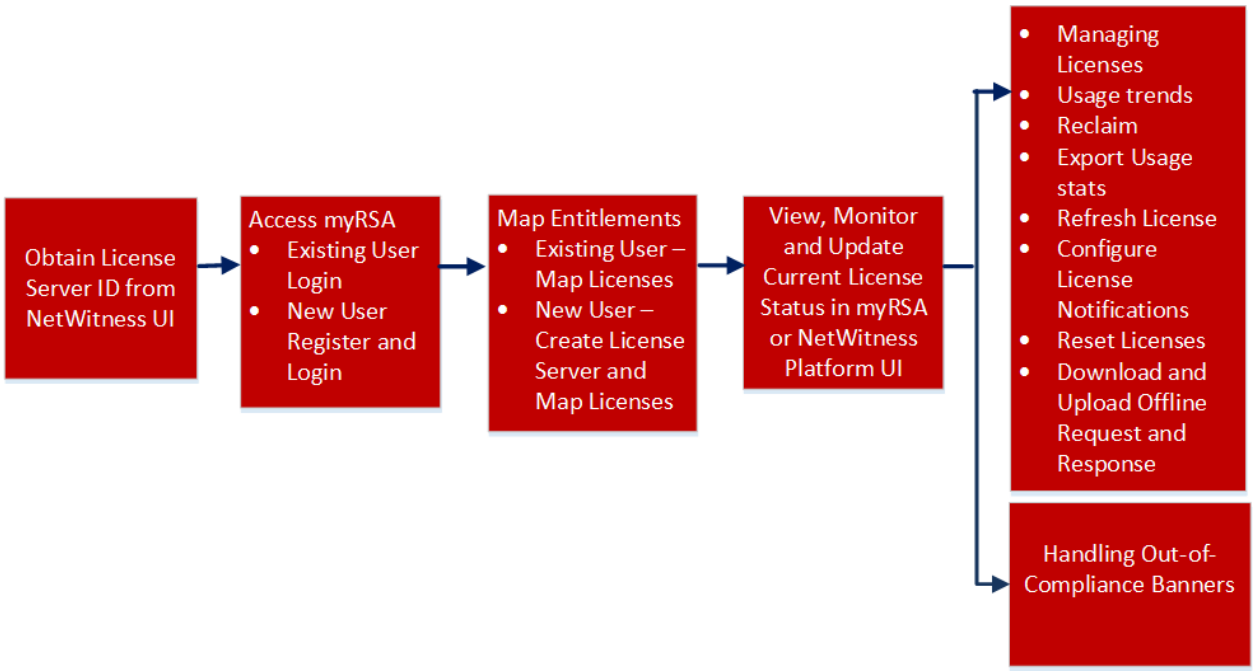
After you have purchased the licenses and have received all the relevant email communications, perform the steps required to map entitlements in NetWitness. You must perform each step in the proper sequence.

This document is divided into three parts:

- How to obtain license information from myRSA and view in the NetWitness Platform UI
- Understanding License Types and its functionality
- Troubleshooting or maintenance information

Workflow

The following workflow illustrates the end-to-end licensing process after you have the NetWitness product installed .



Configuration Step	Description
Obtain License Server ID from NetWitness Platform User Interface	Before you begin the licensing process, you must ensure that you obtain the License Server ID displayed on the NetWitness User Interface.
Access myRSA	Your myRSA Welcome e-mail message contains system log in instructions to myRSA. Instructions for downloading your product licenses can be found in this document, as well as the myRSA website.

Configuration Step	Description
Register the Server (Online Registration)	Your NetWitness Server must be registered to myRSA and entitlements must be mapped. There are two methods of synchronizing NetWitness with myRSA: online and offline.

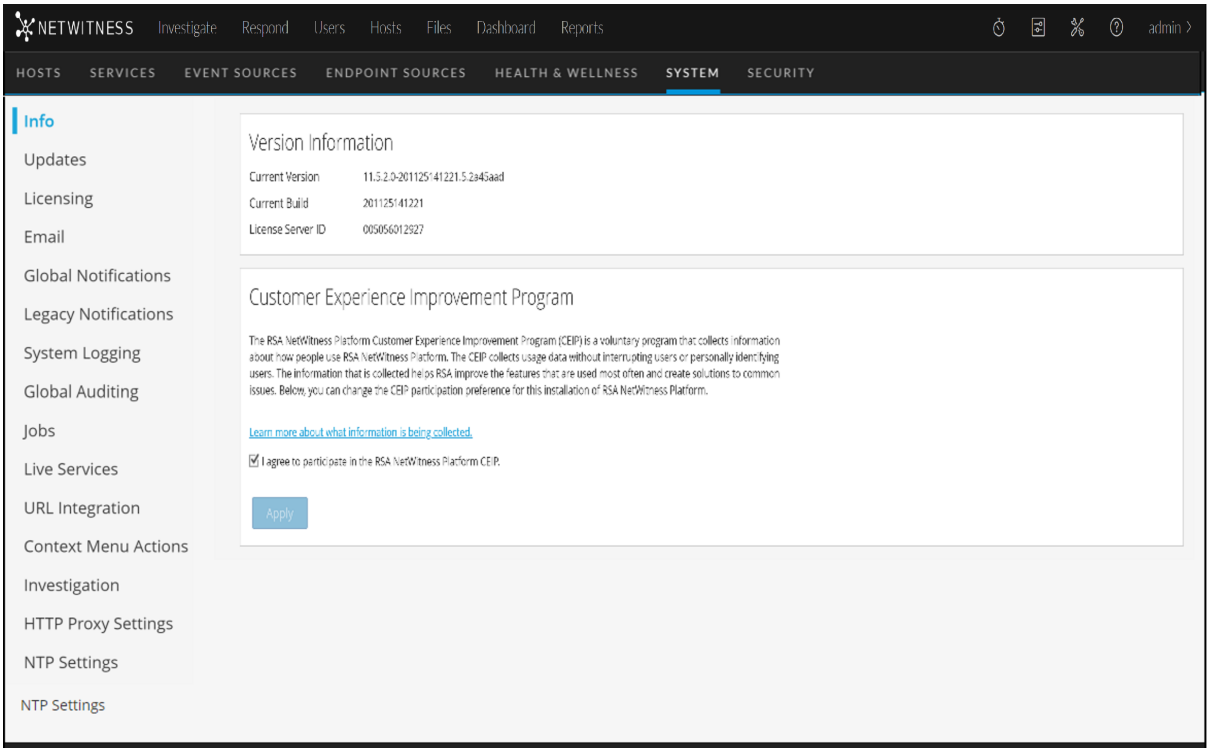
Obtain License Server ID from NetWitness Platform User Interface

Before you begin the licensing process, you must make sure that you obtain the License Server ID displayed in the NetWitness User Interface which is generated on the successful installation of the product.

To obtain the license service ID:

1. Log in to the NetWitness user interface.
2. Go to  **(Admin) > System**.

   The Admin System view opens to display the Version Information in the **Info** panel.

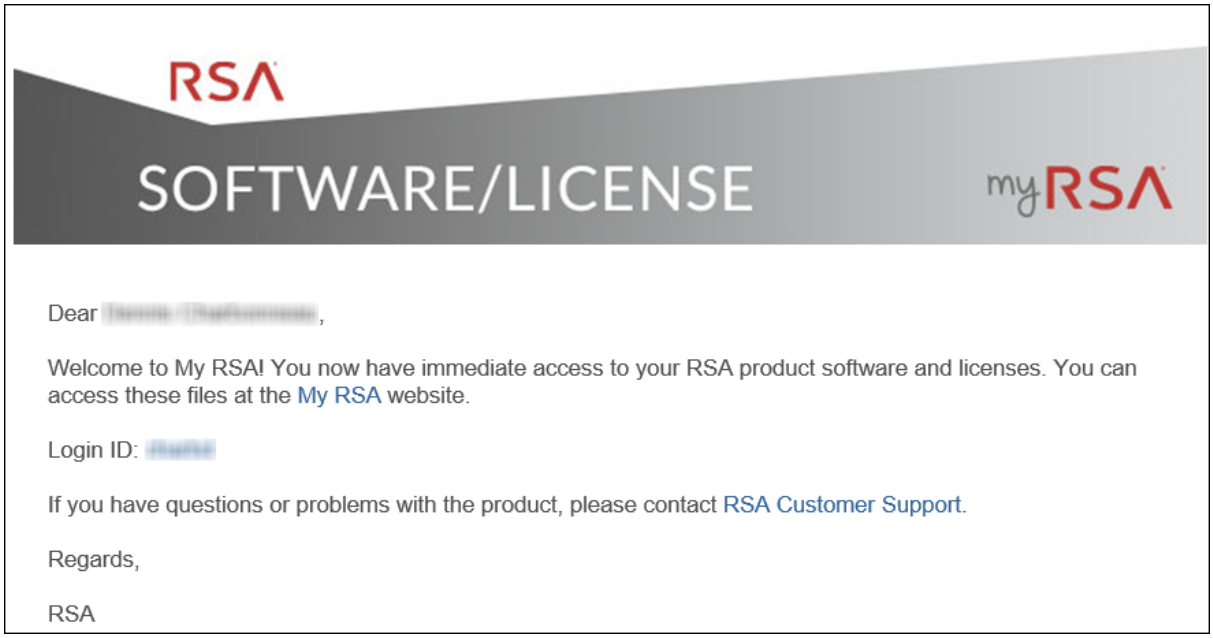


3. Under **Version Information**, locate the **License Server ID** and make a note.
   Log in to **myRSA** and enter the License Server ID to view your entitled license information. For more information, see Access myRSA.

Access myRSA

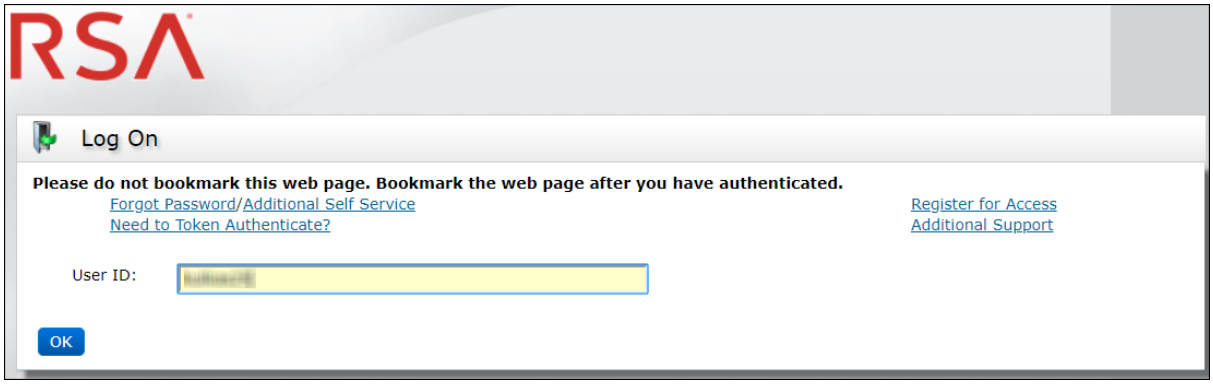
You will receive an order confirmation e-mail. Refer to your myRSA email for login information. If you are a new myRSA user, you also receive an e-mail containing instructions on how to create a new account. Your myRSA Welcome e-mail message contains system log in instructions to myRSA. This document provide you instructions on how to map your product licenses on the myRSA website.

1. The new users instruction e-mail contains a **Click Here** link. This link opens an Enrollment Portal, where the -Based Authentication (RBA) method must be configured for the account.
2. After the RBA method is enabled, a Confirmation e-mail message containing the User ID (which is your e-mail address), along with a temporary password is sent. After the first login, you are prompted to change the password. Once the password is changed, you are logged into myRSA.



Note: If you have a pre-existing account for the RSA Link or RSA Online websites, you receive only one e-mail on how to use your existing login credentials. You must log into myRSA with your existing User ID, password, and RSA method.

3. When you navigate to https://my.rsa.com/, the **RSA Secure Logon** screen is displayed.



4. Enter your User ID and click **OK**, which displays the **Password** field. When you enter the password, you are logged into myRSA.

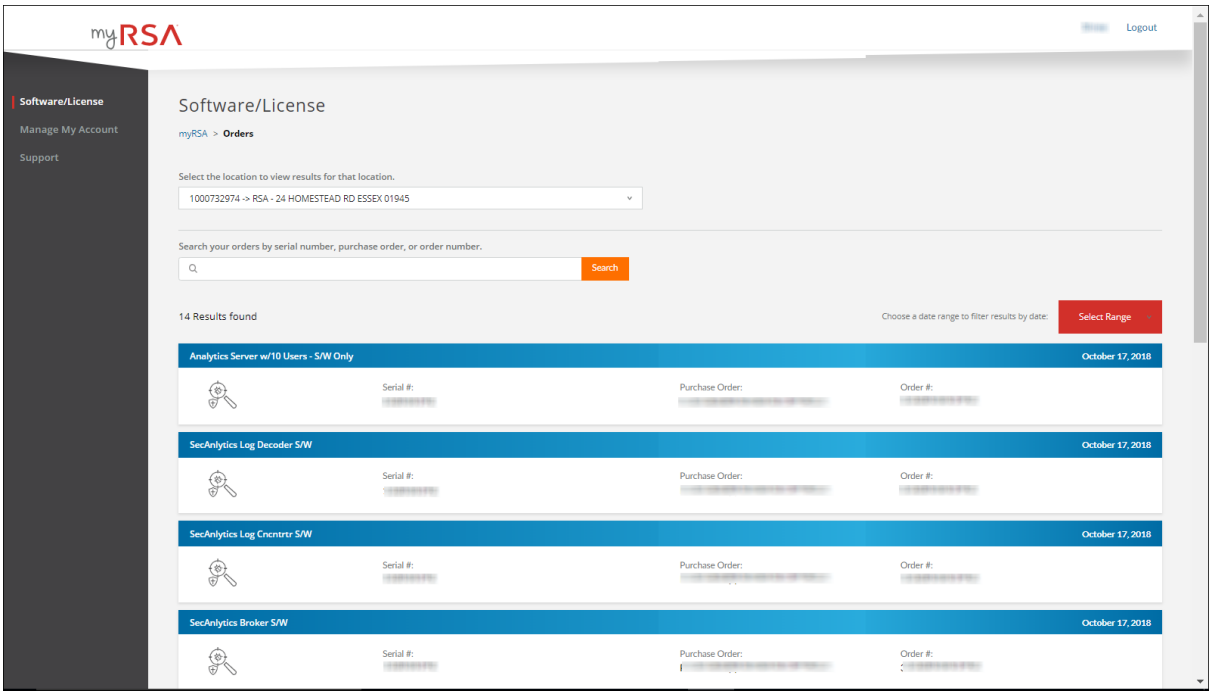
   Your contact e-mail address is used to authenticate your User ID. If the Customer Authentication process is successful, the myRSA Software/License page displays a list of all the following products available which are associated to your site location:

   - RSA Products
   - Serial Numbers
   - Purchase Orders
   - Sales Orders

   Note: You may be prompted to verify your identity through your RBA method, if multiple log in failures occur in a row, or if you have not logged into myRSA within the past several months.

5. In the Sofware/License page, select the Order Location from the drop-down menu. The list of products, sales orders, and purchase orders are filtered and displayed.

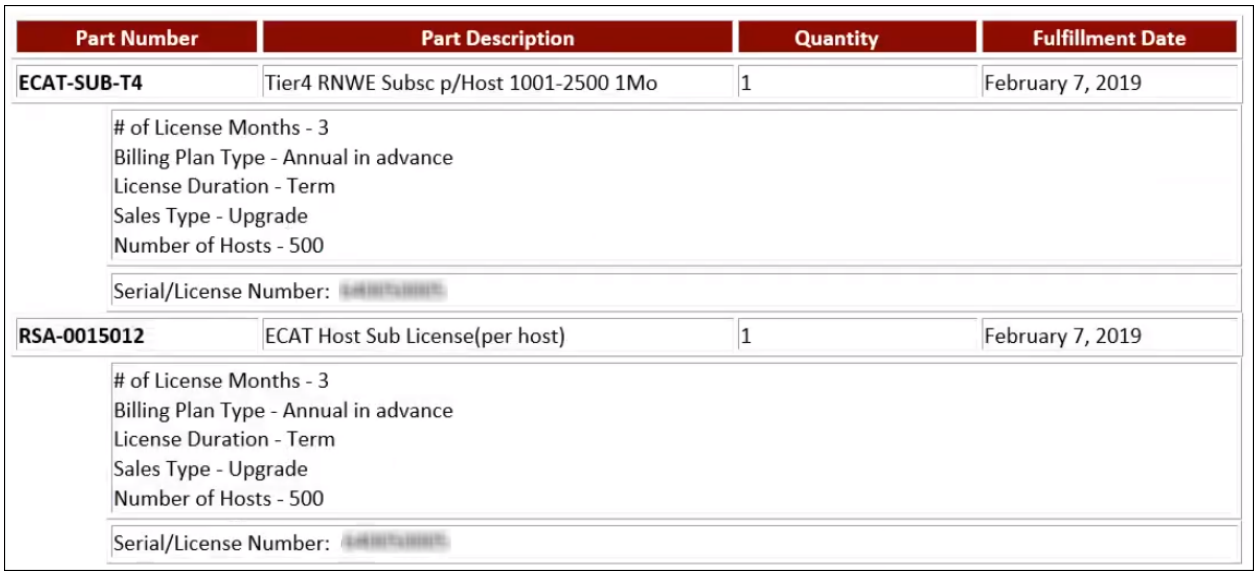
   The **Order Detail** screen is displayed.



License for NetWitness Endpoint 4.4.0.x Agents

If you are an existing NetWitness Endpoint 4.4.0.x customer, who also has NetWitness Server, Log Hybrid and an ESA box, you receive licenses for NetWitness Endpoint 4.4 and 11.6 agents. You receive the following RSA Order Fulfillment Confirmation email which contains the license details for NetWitness Endpoint 4.4.0.x and NetWitness 11.6.

Note: If you are a new NetWitness 11.6 customer, the RSA Order Fulfillment Confirmation email contains the license details for the current 11.6 version only.

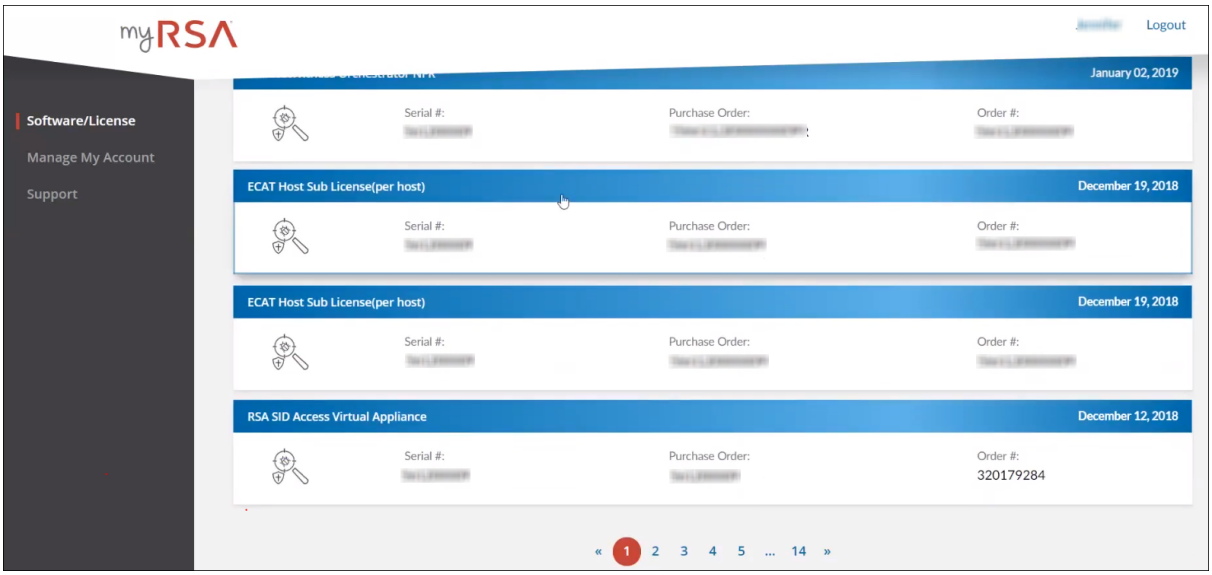


Part Number	Part Description	Quantity	Fulfillment Date
ECAT-SUB-T4	Tier4 RNWE Subsc p/Host 1001-2500 1Mo	1	February 7, 2019
	# of License Months - 3 Billing Plan Type - Annual in advance License Duration - Term Sales Type - Upgrade Number of Hosts - 500		
	Serial/License Number: [redacted]		
RSA-0015012	ECAT Host Sub License(per host)	1	February 7, 2019
	# of License Months - 3 Billing Plan Type - Annual in advance License Duration - Term Sales Type - Upgrade Number of Hosts - 500		
	Serial/License Number: [redacted]		

In the above screenshot, the part number with **RSA-0015012** indicates that it is aNetWitness Endpoint 4.4 license and the part number with **ECAT-SUB-T4** indicates that it is a NetWitness 11.6 license. You will get NetWitness 11.6 license and if you have NetWitness Endpoint 4.4 license, myRSA will upgrade the number of host for both licenses.

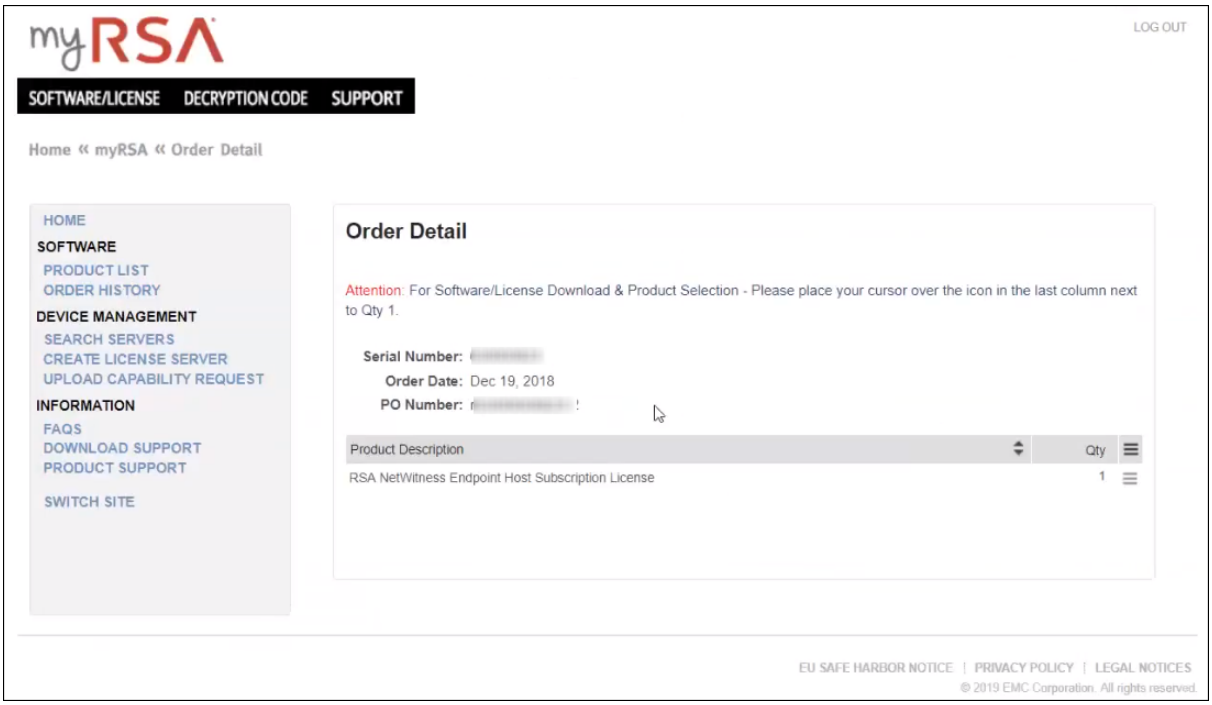
To download an Endpoint (ECAT) 4.4.x license, perform the following steps:

1. After you log in to myRSA, an ECAT Subscription Licenses option is displayed.



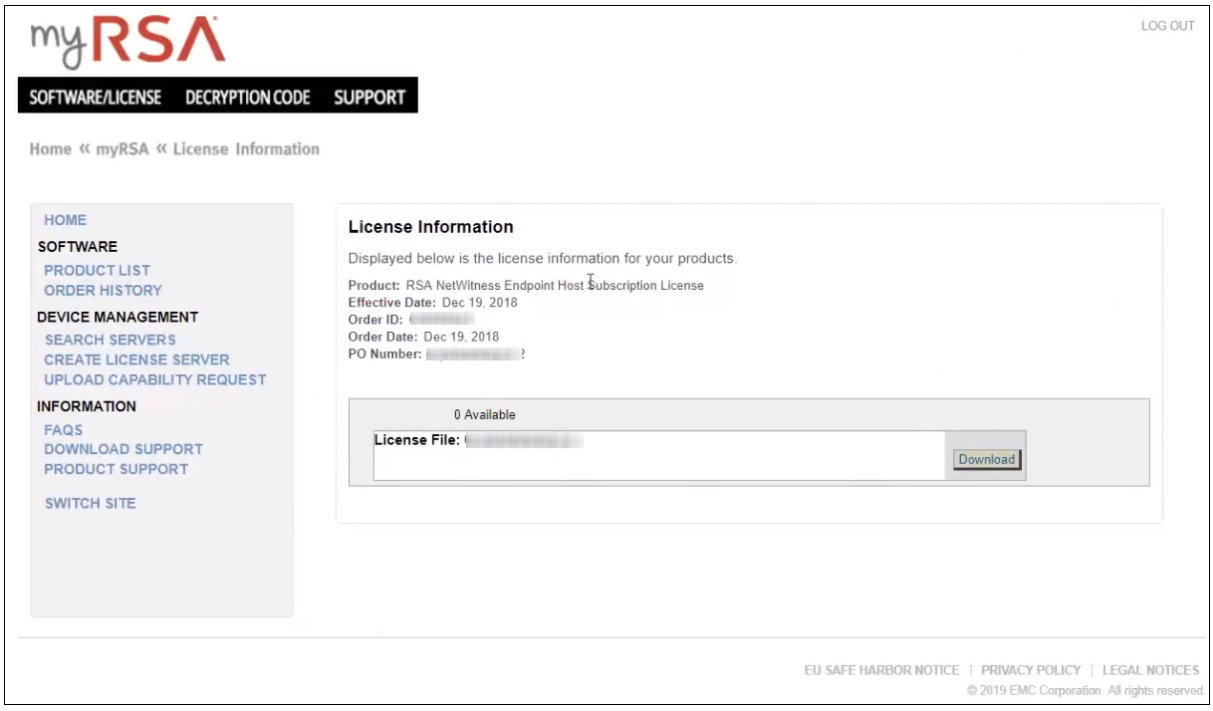
2. Click on ECAT Host Sub License (per host).

   The Order Detail page is displayed.



3. Click on **Product Description**.

   The License information page is displayed.



4. Click **Download** to download NetWitness Endpoint 4.4.0.x license.

If you have an existing NetWitness Endpoint 4.4.0.x and do not have RSA NetWitness Platform, then you must register to the license server. For more information. see Register the Server (Online Registration).

If you have an existing NetWitness Endpoint 4.4.0.x customer, and also NetWitness, then you can use the same license in NetWitness.

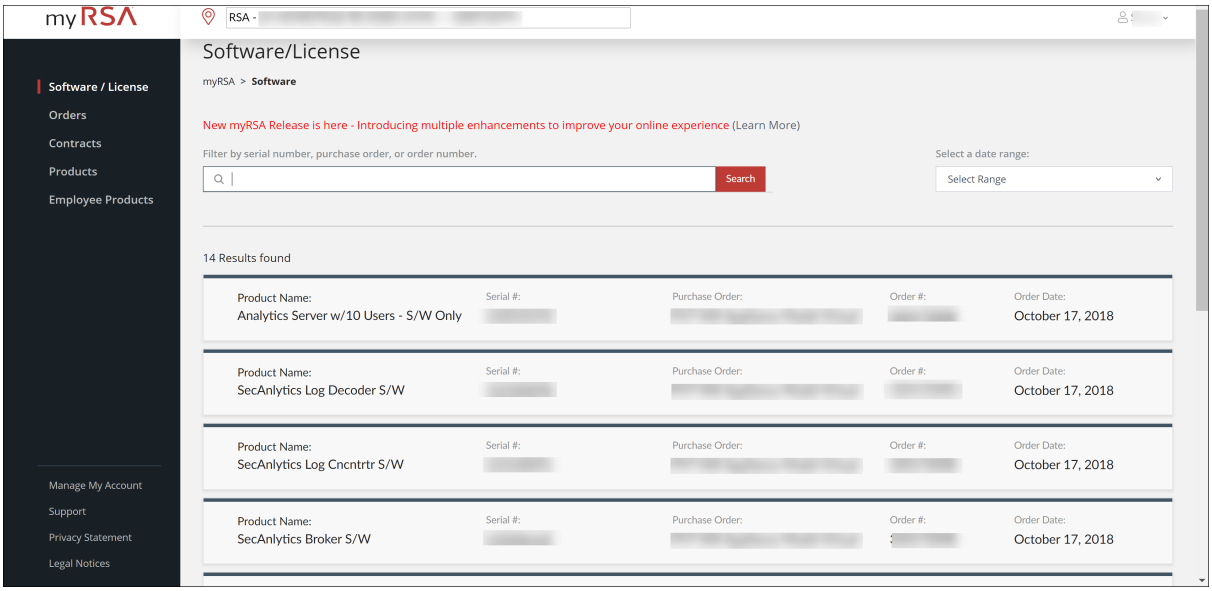
If you are new to NetWitness 11.6 , and have NetWitness Endpoint, then you can view the license only for the current version in NetWitness.

Register the Server (Online Registration)

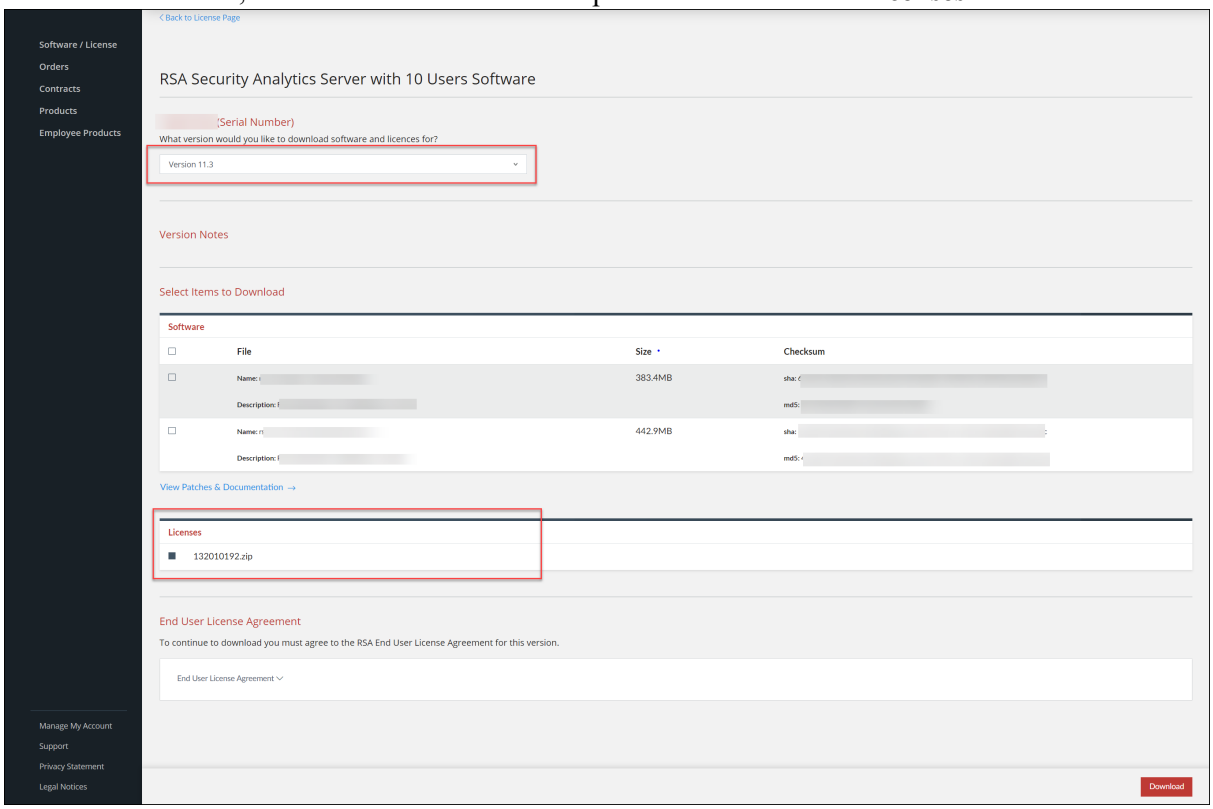
In the NetWitness entitlement process, you need to register the NetWitness Server and map the entitlements to the Local License Server (LLS). There are two methods of synchronizing NetWitness with myRSA: online and offline.

Note: By default NetWitness is configured to synchronize with myRSA at regular intervals hence manual synchronization is not required.

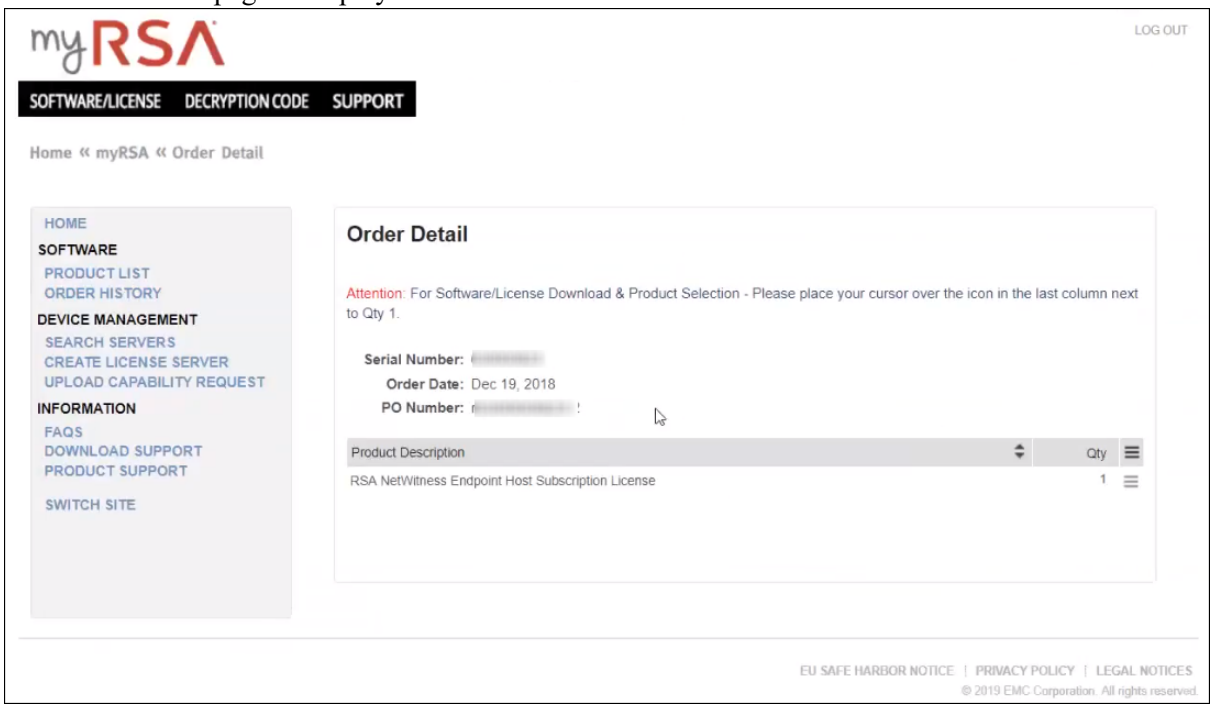
1. On the Sofware/License page, select the Order Location from the drop-down menu. The list of products, sales orders, and purchase orders are filtered and displayed.



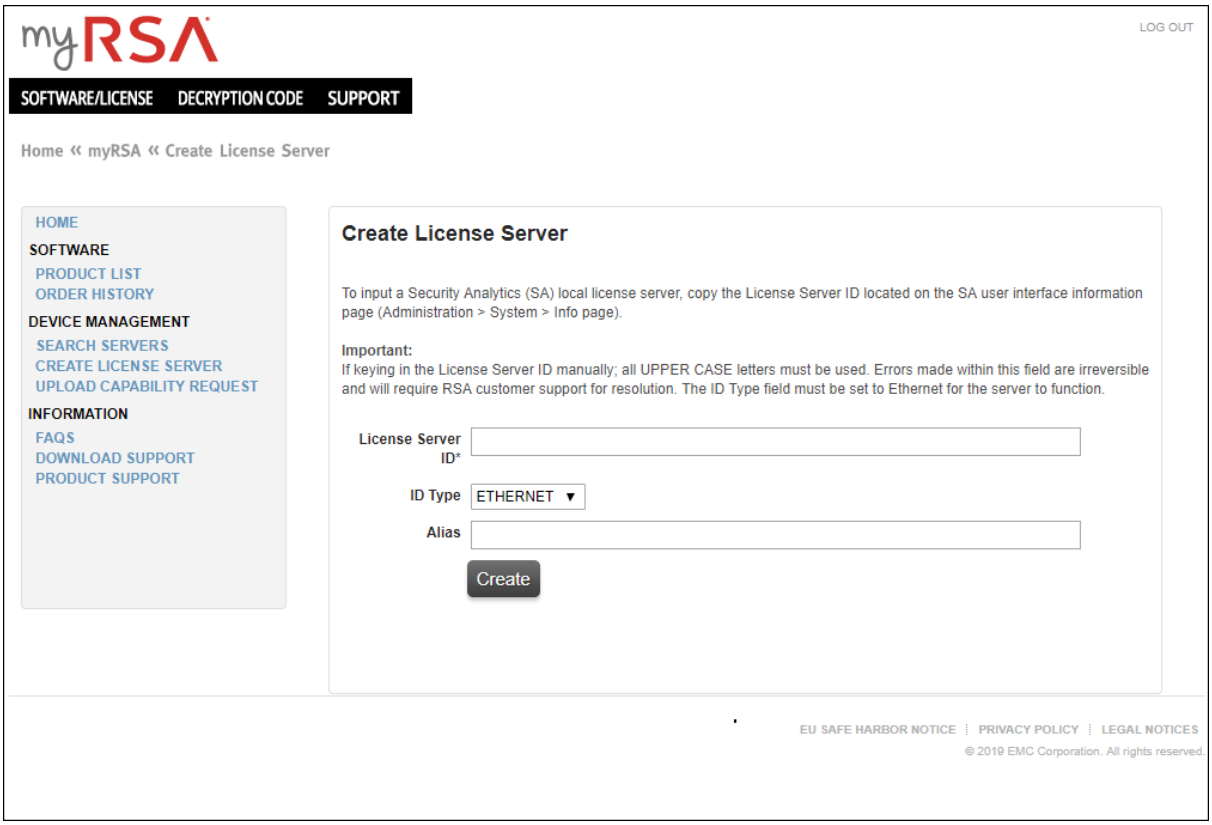
2. Click on a line item, select a version from the drop-down and select the **Licenses** checkbox.



3. Click **Download**.
The Order Detail page is displayed.

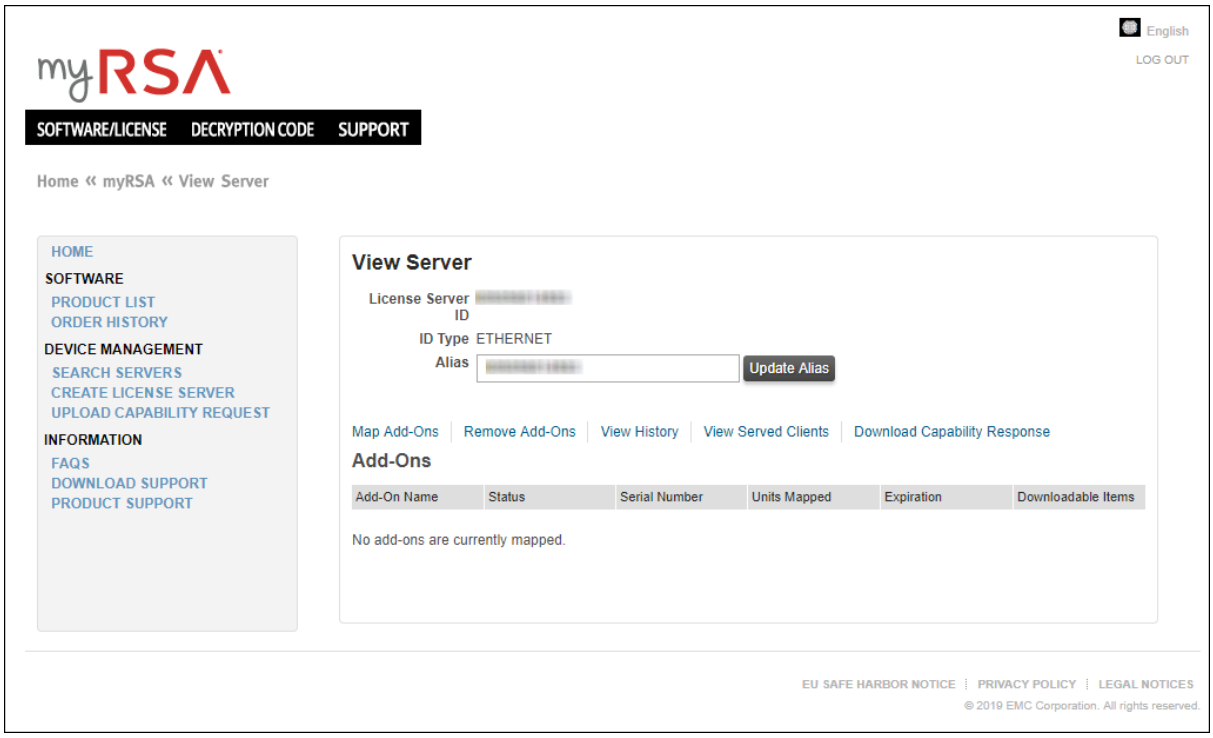


4. Click **Create License Server**.



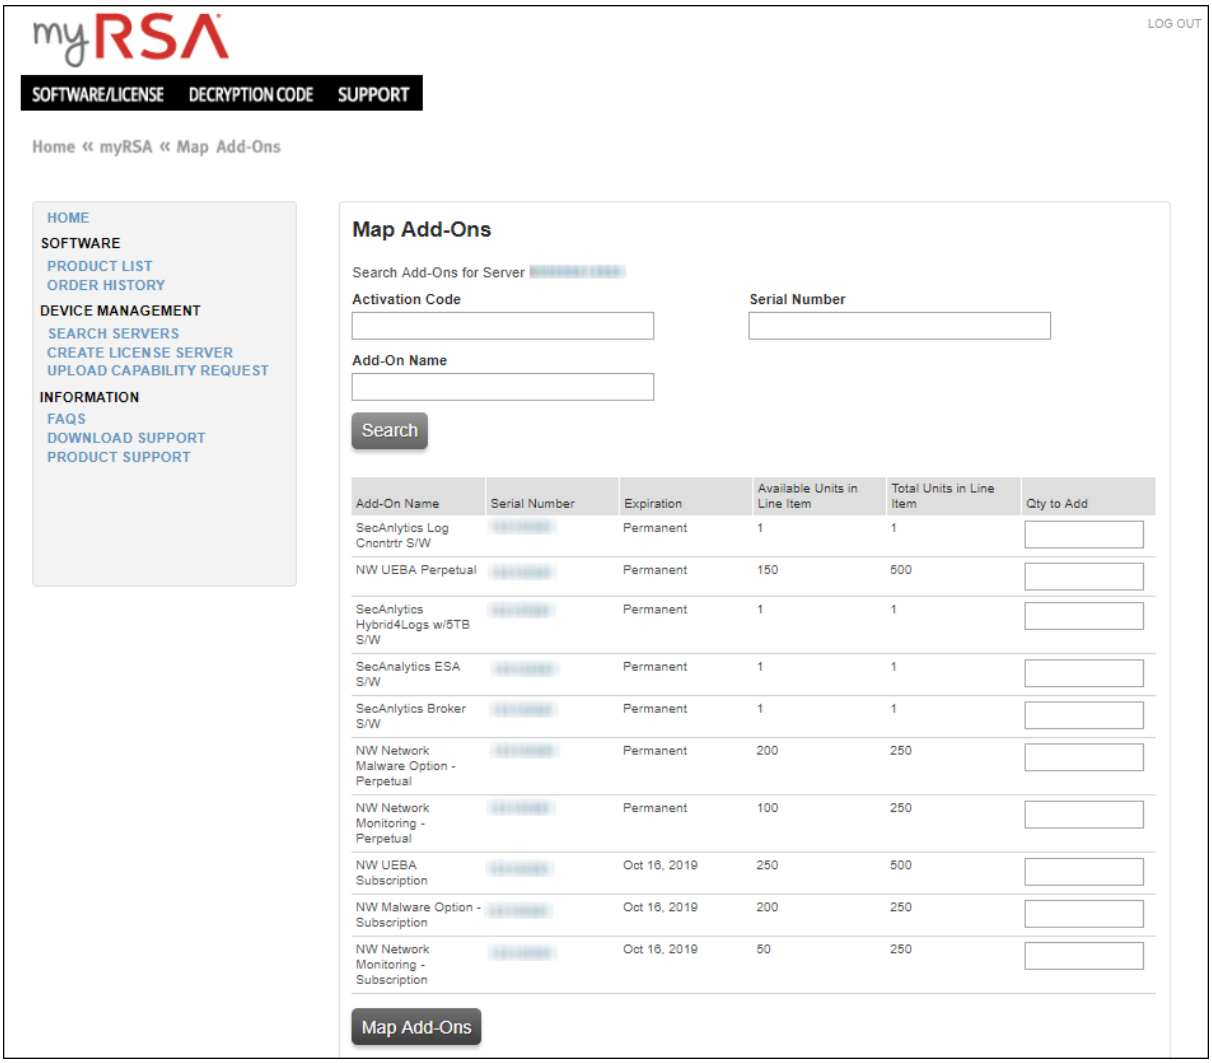
5. Complete the following fields:
   - Copy or enter (in uppercase letters) the License Server ID in the License Server ID field.
   - In the **ID Type** drop-down, select **ETHERNET** (the default value).
   - In the **Type** drop-down, select **Ethernet** (the default value).
   - (Optional) In the **Alias** field, type an alias to your Appliance ID.
6. Click **Create**.

   The View Server page is displayed.



7. Click Map Add-Ons.

   The Map Add-Ons page is displayed.



8. After you have registered the license server and viewed the license details, you must map the entitlements. For more information, see Map Entitlements .

Note: By default NetWitness host is configured to synchronize with myRSA at regular intervals and also a designated nameserver (DNS). No action is required.

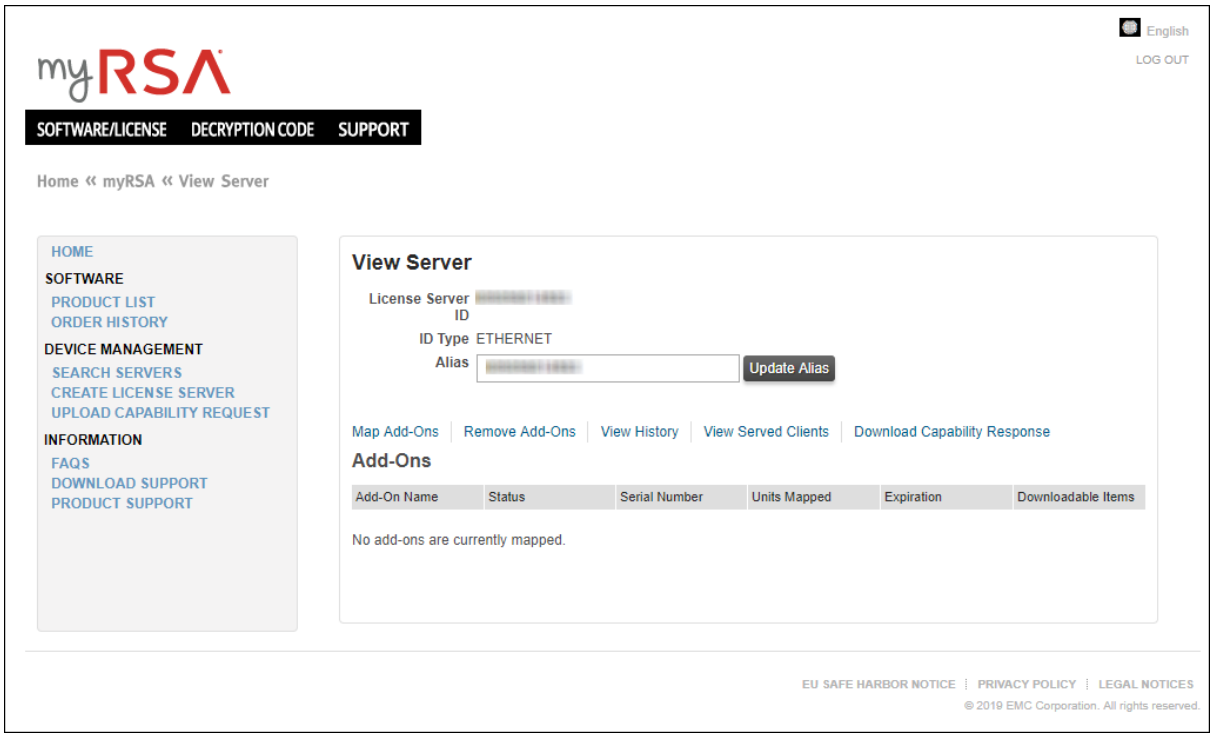
Note: In a multiple NetWitness deployment where the services are connected to both primary and secondary NetWitness hosts and the services are licensed only with the primary NetWitness, a license expiry message is shown for the same services on the secondary NetWitness. You can ignore the message and continue using the product.

Map Entitlements

Mapping entitlements involves choosing the quantity of available entitlements to pull to the NetWitness Server during synchronization.

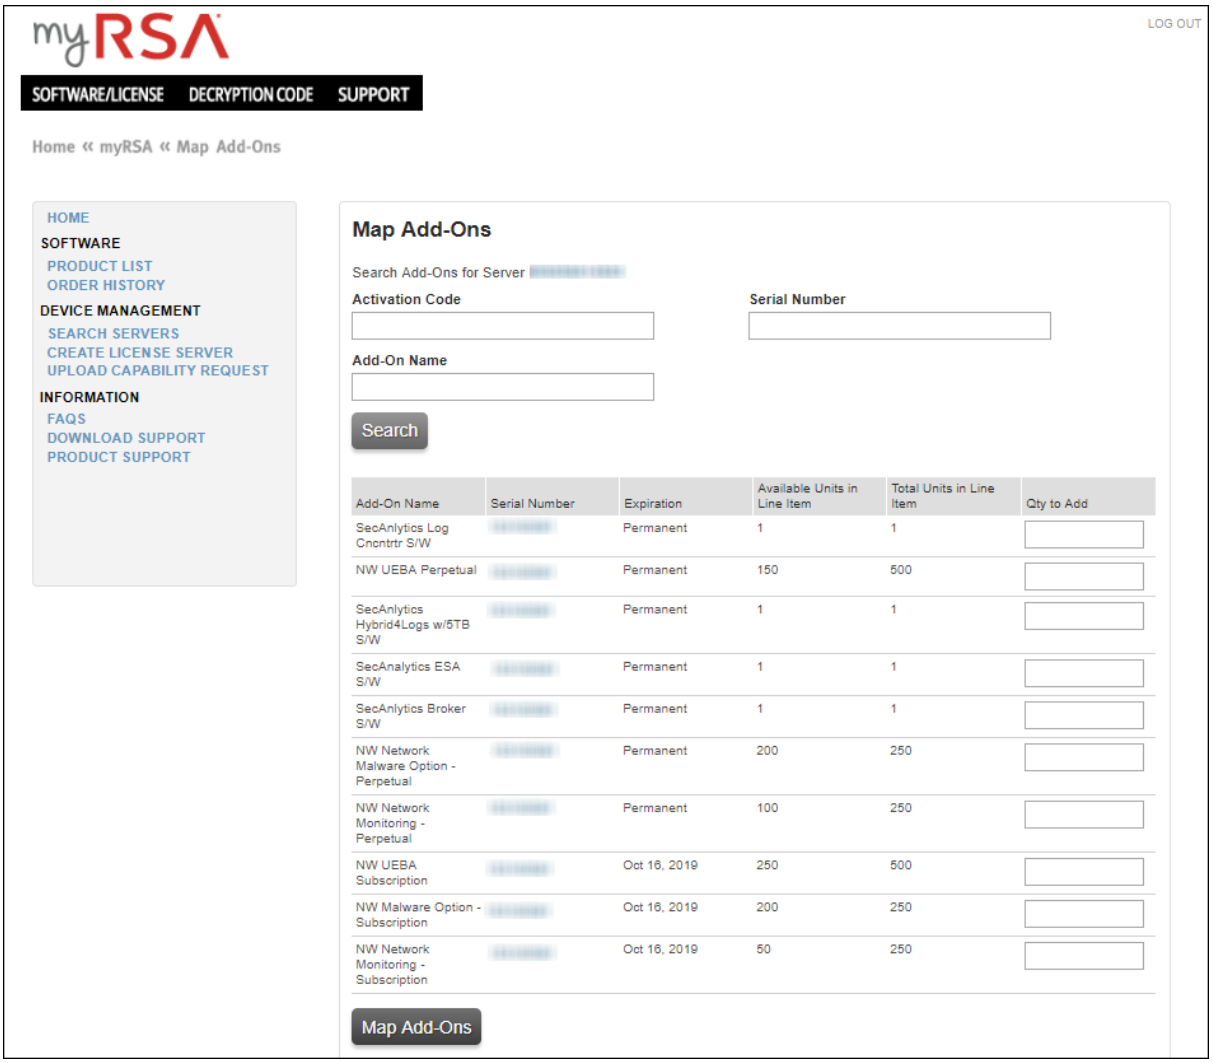
To map appliance entitlements to the server:

1. Log in to myRSA.
2. Click on the **SEARCH SERVERS** and then click on the License Server ID to view the entitled Server Details.



3. Click **Map Add-Onss**.

   The Map Add-Ons page is displayed.



The Add-On table lists all entitlements that are available for your account. The table has a row for each NetWitness product entitlement, with the following information:

- **Add-On Name**: The name of the entitlement; for example, SMC Concentrator or SMC Decoder.
- **Serial Number**: The serial number associated with a product.
- **Expiration**: For licenses that are not permanent, the expiration information is displayed. The value in this field is a specific date (for example, Oct 6, 2019) or a time range (for example, 90 days). If the value is a time range, the expiration period begins when the add-on is mapped to a server.
- **Available Units in Line Item**: The quantity of entitlements currently available for that license. This quantity is the difference between the Total Units and the entitlements that have been mapped to a NetWitness Server for licensing.
- **Total Units in Line Item**: The total quantity of entitlements tied to a specific license.
- **Quantity to Add**: The number of entitlements for a specific license.

4. To designate the quantity of entitlements to map to the NetWitness Server from a license, type a quantity in the **Units to Configure** column.
5. Click **Map Add-Ons**.

   The View Server page displays a message indicating that the entitlements were successfully mapped to the NetWitness Server.

   Entitlements are now dedicated and set aside from an accounts pool. The message **Waiting to add to appliance** is displayed in the **Status** for each entitlement. The entitlements are not yet synchronized with the server.
6. (Optional) If you want to add more entitlements, use the **Map Add-Ons** option.
7. (Optional) If you want to remove entitlements, use the **Remove Add-Ons** option.

Now you can synchronize and pull down the mapped entitlements to the NetWitness Server

View Current Licenses

After you have completed the license process, you can view the current licensing status in the NetWitness UI.

Each NetWitness Server is a license server providing capabilities to entitle services connected to it. To make entitlements available for licensing services, the entitlements must be mapped to the Local License Server (LLS) on the NetWitness Server.

Note: If licensing a hybrid system, which has a Concentrator and Decoder on the same appliance, each component must be licensed separately.

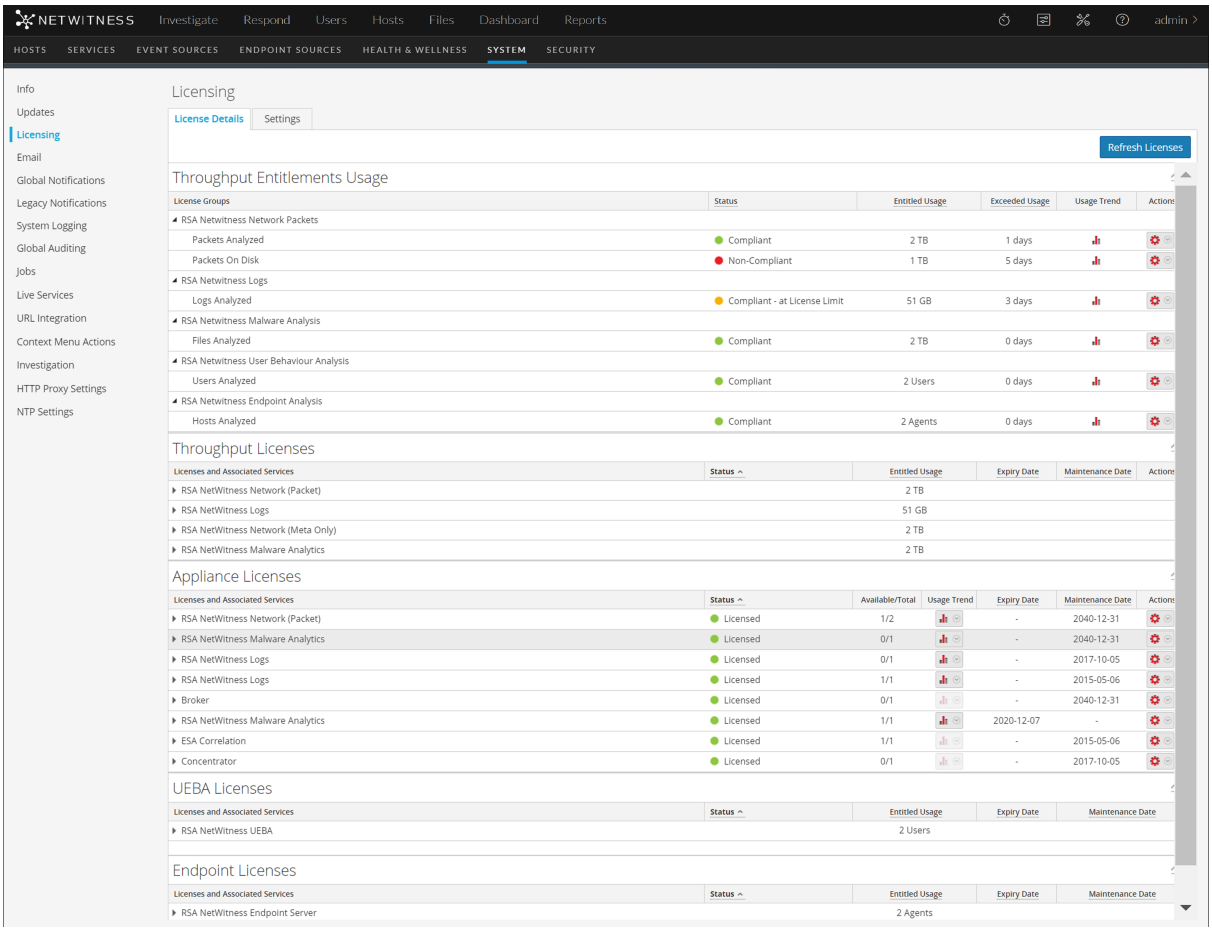
View and Manage Licenses

In NetWitness, you can view and manage available licenses.

To view the licenses that are available on this instance of NetWitness:

1. Go to  **(Admin) > System.**
2. In the options panel, select **Licensing**.

   The **License Details** tab is displayed.



Each license is listed in the grid by license type. Information includes the status of the license indicated using color-coded circles and the related information.

Register the Server (Offline Capability Request)

NetWitness manages licensing through a Local License Server (LLS). Each client appliance is shipped with an installed LLS. This topic provides instructions for offline synchronization of the Local License Server (LLS) with the online repository. For more information on the functional description of the LLS, see Entitlement Capability Implementation.

If you are unable to register the NetWitness Server online, you can download an offline capability request in NetWitness and upload that binary request to the myRSA Portal. If the NetWitness Server is not connected to the Internet, you can perform offline synchronization of entitlements through the View Server page in myRSA.

Prerequisites

Before implementing the NetWitness Platform entitlements capability offline, ensure the following

- The NetWitness Server is registered to myRSA (https://my.rsa.com/) and entitlements are mapped. Internet access is not required for offline synchronization.
- An Offline Capability Request has been downloaded from the NetWitness UI to be uploaded to myRSA.
- An Offline Response has been uploaded to NetWitness that was received from myRSA within 24 hours.

Here is a workflow that describes how to use the offline capability to acquire the licenses from myRSA and view them.

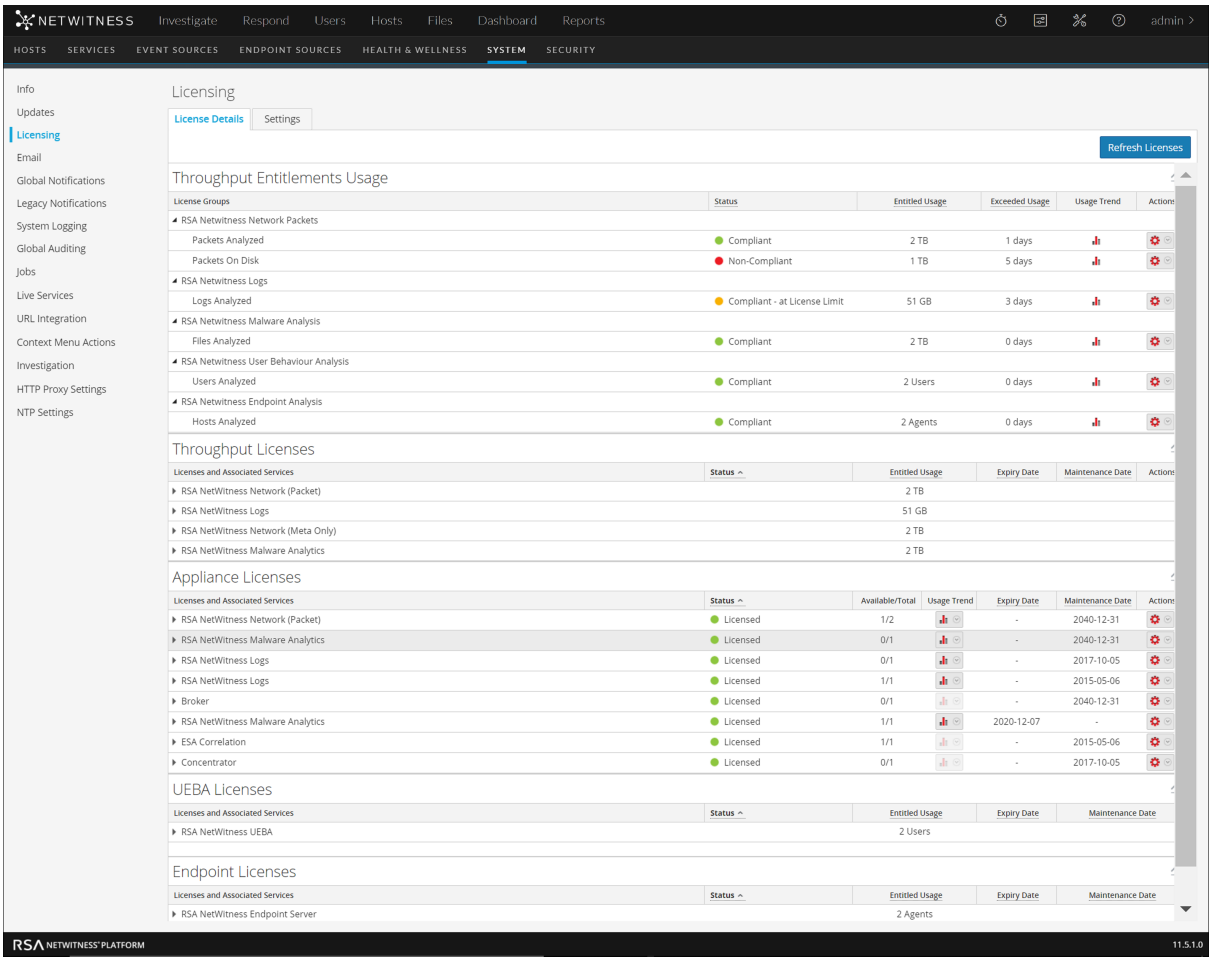


Download an Offline Request for Submission to myRSA

To register the server using an offline capability request:

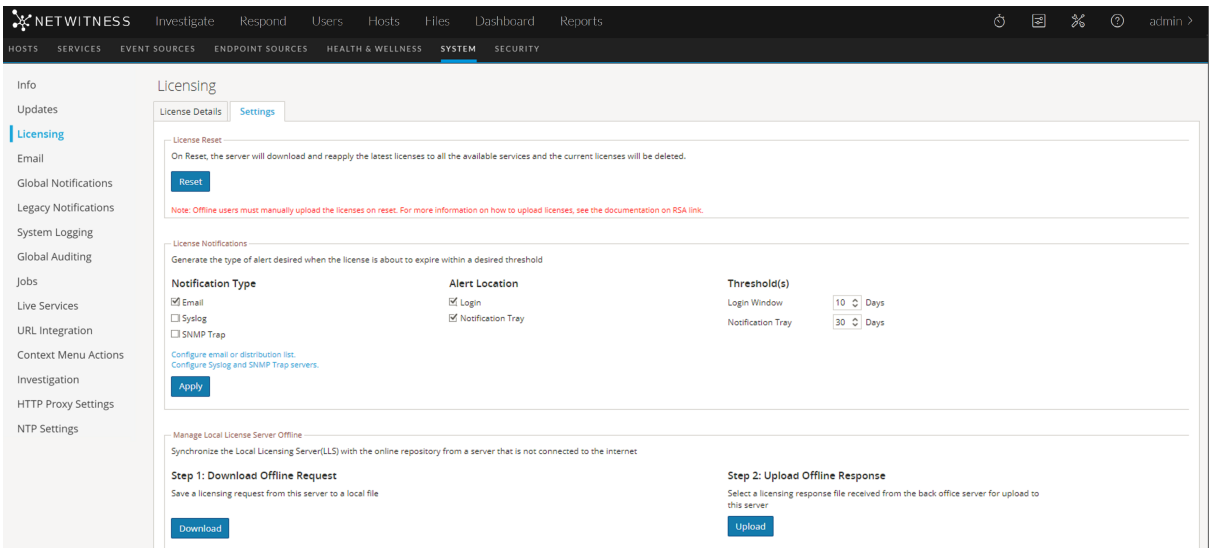
1. Go to  **(Admin) > System**.
2. Click **Licensing**.

   The License Details page is displayed.



3. Select the **Settings** tab.

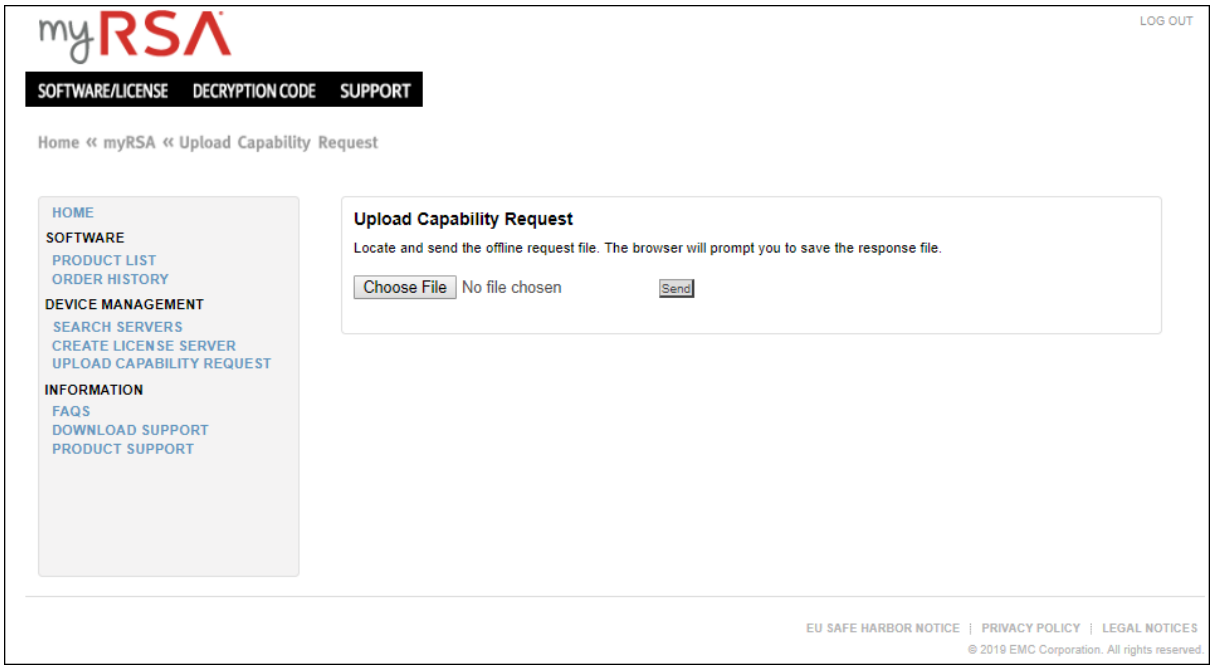
   The Settings page is displayed.



4. In the **Step 1. Download Offline Request** panel, click **Download**.
   A file called **OfflineCapabilityRequest.bin** is downloaded to the local system.
5. Login to the myRSA Portal at https://my.rsa.com/ with your user credentials.
6. Click on a service which needs to be licensed.

   Under **Device Management**, click **Upload Capability Request**.

   The **Upload Capability Request** dialog is displayed.



7. Click **Choose File** and browse the local file system to find the file downloaded from the NetWitness Server. Select **OfflineCapabilityRequest.bin**.

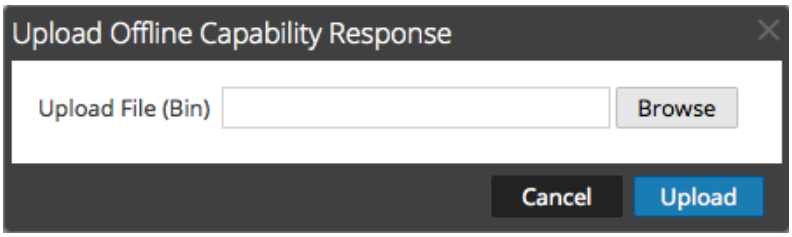
   The filename is displayed next to the **Choose File** button.
8. Click **Send**.
   The license data in a .bin file format is downloaded. The .bin file contains the data you entered as well as information about any entitlements that have been added to the NetWitness Server. If the server has just been added, there are no entries under **Add-Ons**. The server is registered and you can now map entitlements as described in the following sections.
9. After you download the .bin file from myRSA, login to NetWitness UI to upload the .bin file on the Settings page. For more information, see Upload an Offline Capability Response to NetWitness Platform.

Upload an Offline Capability Response to NetWitness Platform

If the NetWitness Server is not connected to the Internet, you can perform offline synchronization of entitlements through the View Server page in myRSA. To upload an offline capability response (**response.bin**) file saved to the local file system from myRSA:

1. Follow the steps 1 to 3 mentioned in the Download a Capability Request for Submission to myRSA procedure.
2. Login to the NetWitness user interface.
3. Go to **(Admin) > System** > **Licensing** > **Settings** tab.
4. In the **Upload Offline Response** section, click **Upload Response**.

   A dialog to upload the file is displayed.



5. Browse and select the **response.bin** file so that it is displayed in the Upload File (bin) field.
6. Click **Upload**.

   The entitlements are uploaded to NetWitness and the licenses added to the grid in the **Licensing Details** tab. They are available for licensing appliances.

   After you have uploaded the entitlements, you can verify the synchronization by performing any one of the following:

   - To view results in NetWitness, go to **(Admin) > System**>**Licensing** > **Licensing Details** tab. The individual product entitlements that have been pulled down to NetWitness are displayed in the **Available/Total** column.
   - Within the myRSA interface, you can see the status for entitlements changed to **In Sync**.

Refresh Licenses

When a new license is added, to map the view with the new license, click **Refresh Licenses.**

Refreshing your licenses performs the following behind-the-scenes tasks:

- Restarts the LLS server to ensure the latest licenses are pulled down from the central Flexera server.
- Associates any unlicensed service with a valid license (if available).
- Replaces expired or Out-of-the-Box license with valid licenses (if available).

License Types

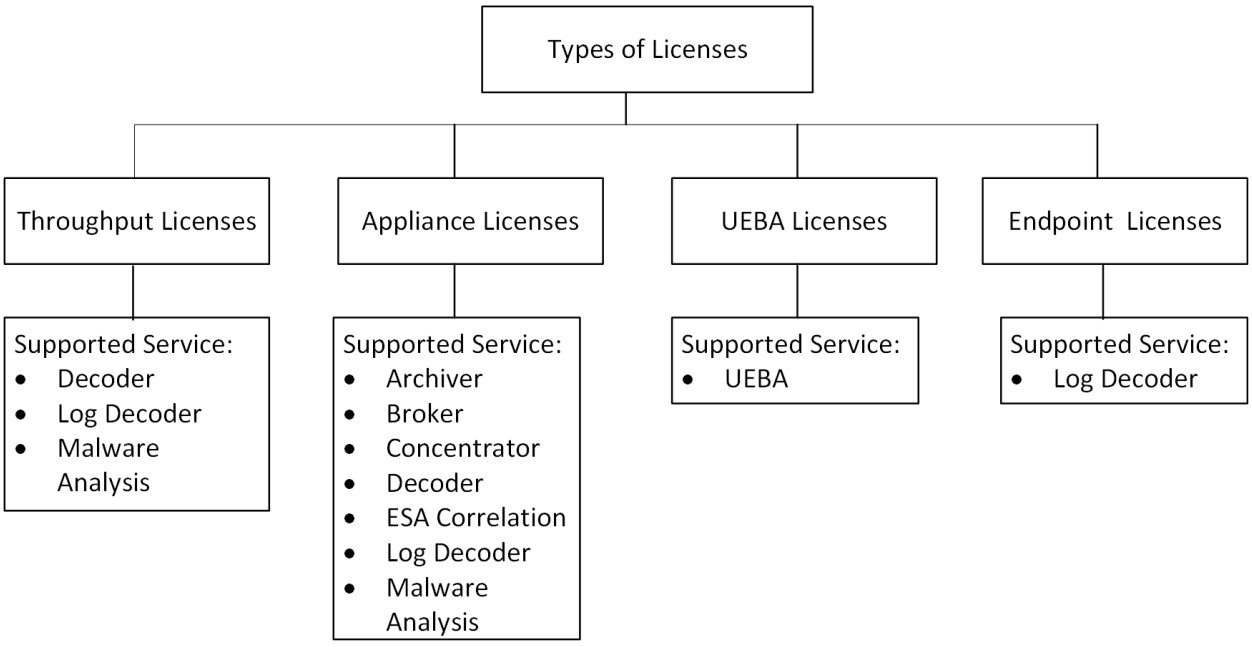
After you have installed the NetWitness software and the required services, you need to acquire the relevant licenses for the each of the services or a group of services based on your requirements. NetWitnessentitlement uses a trust-based licensing model. Appliances continue to function as usual even when the license is out-of-compliance.

Choosing a License Type

The type of license you choose is based on your network requirements. The following types of licenses are available in RSA NetWitness:

- Throughput License
- Appliance License
- UEBA (User and Entity Behavior Analytics)
- Endpoint License

Here is a chart, followed by a description of each license type available for the NetWitness Platform products and services.



Throughput Licenses

Throughput license are based on the amount of data used per day for logs (SIEM), or network (packets) or malware.

The throughput per day is measured in Gigabytes per day for logs, and in Terabytes per day for packets. The total amount of throughput is selected based on the total amount of throughput per day that is being licensed across your entire enterprise deployment of NetWitness Platform. This license is measured as follows:

Service	Unit of Measurement	Increments	License Duration
Network Decoder	Terabytes (TB) per day	1 TB	Subscription or perpetual
Log Decoder	Gigabytes (GB) per day	50 GB	Subscription or perpetual
Malware	Terabytes (TB) per day of Files Analyzed by Malware Analysis	1 TB	Subscription or perpetual

Note: In version 11.5 and later, NetWitness Platform introduced Network Meta-Only license in addition to the Network Full Packet license.

Network Full Packet License captures, analyzes, and stores packet payloads and metadata and allows users to retain the packet payload for analysis or for session reconstruction. Network Full Packet License measures the bytes analyzed and the bytes written to the disk for Network Packets.

Network Meta-Only License captures and analyzes packet payloads and discards the packet payload data after analysis. Using this license, NetWitness Platform can be deployed in an environment where full packet capture is not required. This helps to optimally manage the storage space, and easily detect threats without the need to retain the full payload of the sessions. Network Meta-Only License measures the bytes analyzed for Network Packets and can be used with or without the Network Full Packet License.

When customers have both - Network Meta-Only License and Network Full Packet Licenses, the entitlement is aggregated and measured against the bytes analyzed for Network Packets. However, the entitlement for bytes written to disk for Network Packets will continue to be measured against the Network Full Packet Licenses. The following table is an example of how the license is used and analyzed.

Available Licenses	Network Full Packet License	Network Meta-Only License	Network Packet Bytes Analyzed	Network Packet Bytes Written to Disk
10TB of Network Full Packet	10TB	0TB	10TB	10TB
10TB of Network Meta-Only	0TB	10TB	10TB	0TB
10TB of Network Full Packet 5TB of Network Meta-Only	10TB	5TB	15TB	10TB
5TB of Network Full Packet 10TB of Network Meta-Only	5TB	10TB	15TB	5TB

For more information on the status of your license, usage and out of compliance licenses, see About Out-of-Compliance Banners.

Appliance Licenses

NetWitness Platform supports the Appliance license, which is applicable to all hosts that require a license. Other services do not require a license. Appliance licenses are measured as follows:

- Services are licensed automatically if you have a valid appliance based license for a specific service to be licensed.
- Appliance licenses can be purchased as a perpetual license that does not expire and will have a maintenance contract. If you purchase a subscription license, then it will expire if you do not renew the contract.

User and Entity Behavior Analytics Licenses

NetWitness Platform supports the User and Entity Behavior Analytics License (UEBA). This license is based on the number active of users.

NetWitness Platform supports the User and Entity Behavior Analytics License (UEBA). UEBA monitors the number of active users from the previous day and sends it to the licensing server. The entitlement is measured for logs and endpoint events for the number of active users and is checked against a user id. If a threshold is identified in a calendar month an appropriate banner is displayed. For more information on banners, see About Out-of-Compliance Banners.

Endpoint Licenses

Endpoint license are entitled based on the number of active agents deployed.

There are two types of agents:

- **Advanced Agents**: The license for these agents is based on the number of advanced agents in your deployment. A 90-days trail license period is provided. After the 90-days trial period, a zero MB and zero Agent license is applied to the Log Decoder service and Endpoint service in the NetWitness Endpoint Log Hybrid. Once an Endpoint license is applied any Archivers, Brokers, Concentrators, and ESA are automatically licensed as a result. An usage exceeded banner is displayed when the license goes out-of-compliance in the following scenarios:

  a. If the number of active agents exceeds the number of licensed agents

  b. If the Endpoint Subscription is about to expire in near future or has already expired.

  For example, if you have purchased a license for 50,000 agents and if the number of agents exceeds more than 50,000, the banner is displayed.

  Or, if you have purchased a license for 50k agents but have mapped the entitlements for only 10k agents on myRSA, an out-of-compliance banner is displayed when your usage exceeds these 10k active agents.

- **Insights Agents** - There is no license required for these agents if they are used to collect only endpoint data.
- **Log Collection with Endpoint Agents** - All Endpoint agents (Advanced or Insights) can forward Windows Log data only to a licensed Log Decoder (or Hybrid). Windows Logs sent to a licensed Log Decoder (or Hybrid) will count against either the applied Throughput or Appliance license. Logs may be retained in the NetWitness Endpoint Log Hybrid as long as a Log (SIEM) Throughput license with available capacity is available. In either case, a license for Logs is required, irrespective of the Insights or Advanced agent.

NetWitness Endpoint **4.4.0.x License**

If you have a **NetWitness Endpoint** 4.4.0.x license, you can use the same amount of license on NetWitness 11.6. For example, if you have purchased a 50,000 license for NetWitness Endpoint 4.4.0.x, you will get a 50,000 license on NetWitness Endpoint 4.4.0.x as well as on NetWitness 11.6. For more information on how to get a license for **NetWitness Endpoint 4.4.0.x**, see License for NetWitness Endpoint 4.4.0.x Agents.

Out-of-the-Box Trial License

RSA NetWitness Platform comes with an OOTB 90-days trial license.

For Endpoint metered license, you will be provided with 0MB Log Decoder license. The 0MB license from the Log Decoder will be utilized after the 90-days trial period is complete. If you are still within the 90-days trial period and if you have an Endpoint Metered license, then the Log Decoder uses the remaining trial license. The Log Decoder will apply the 0MB license only after the trial license completed.

Configure License Support For Multiple NetWitness Servers

If you have multiple NetWitness servers deployed in your environment, and require license support, you must configure multiple Mac addresses.

To configure online:

1. Log in to the NetWitness Server console.
2. Navigate to /opt/netwitness/flexnetls.
3. Edit the local-configuration.yaml file and add the mac address of the back up or secondary host.
   For example, `backup-hostid: 685B3596A5F3` or `backup-hostid: "685B3596A5F3:785B3596A5F3"`.
4. Restart `systemctl restart flexnetls-RSALM`.
5. Obtain the instance details from `wget http://localhost:3333/api/1.0/instances` and set the `"failOverRole"` to `MAIN` to verify the failover.
   For example, `"failOverRole" : "MAIN"`.

To configure offline:

1. Perform the step 1 to step 6 from the online procedure .
2. Download the capability request from NetWitness UI.
3. Upload the request to FNO (myRSA).
4. Upload the response from FNO (myRSA) to the NetWitness UI.

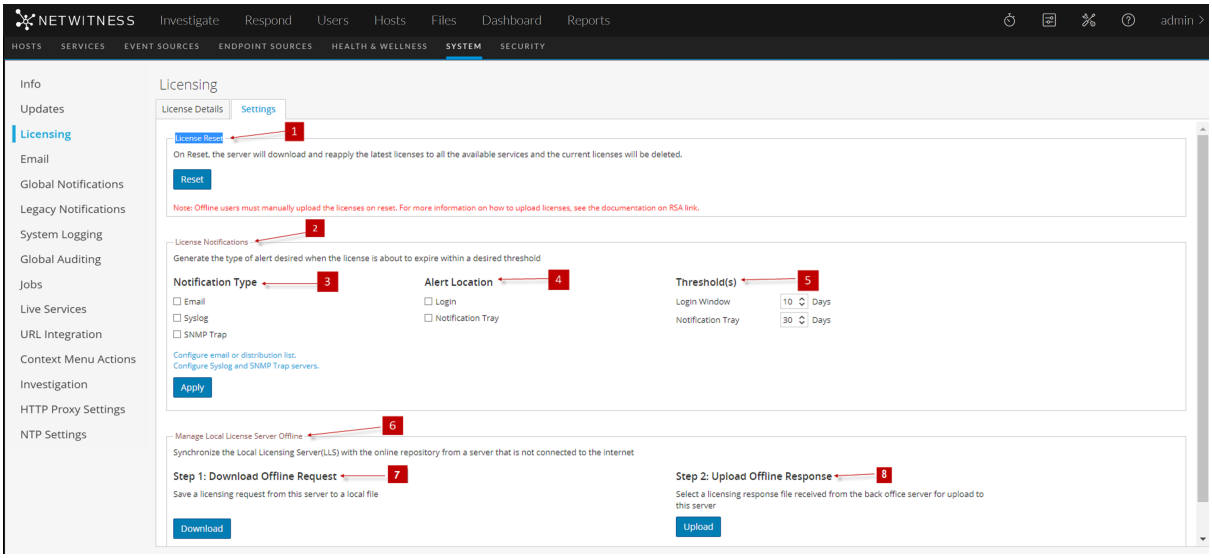
Configure License Notifications

This topic provides instructions for configuring notification settings for the Local License Server (LLS). If you wish to receive alerts about the approaching license expiration date, you can configure NetWitness to send notifications. You can receive notification by email, syslog and SNMP. You can also view the notification during system log on and in the Notification Tray. You can also specify the number of days before expiration as a threshold for notification.

Note: In the Analyst UI, the license notifications will not be displayed in the notification tray or login window when the license goes out of compliance or when the license expires. This is displayed only on the Admin UI.

To configure the notification:

1. Go to  **(Admin) > System**.
2. Select **Licensing**.
3. Select the **Settings** tab.



4. Select each of the methods for NetWitness to use when sending a notification about the license status.
   a. To receive a notification at log on, select **Login** and specify the number of days before the license expires that you want to receive notification in the **Login Window Threshold** field.
   b. To receive a notification in the Notifications tray, select **Notification Tray** and specify the number of days before the license expires that you want to receive notification in the **Notification Tray Threshold** field.
   c. To receive an Email notification to a configured distribution list, select **Email** and select **Configure email or distribution list**. The Email panel is displayed in a separate tab, and you can

configure notifications in the Email Server Settings section. Refer to the *System Configuration Guide* for further details.

d. To receive syslog notifications, select **Syslog** and click **Configure Syslog and SNMP Trap servers**. The System Auditing panel opens in another tab and you can configure the system settings.

e. To receive notifications through SNMP Trap, select **SNMP Trap** and select **Configure Syslog and SNMP Trap servers.** The System Auditing panel opens in another tab and you can configure the SNMP auditing settings.

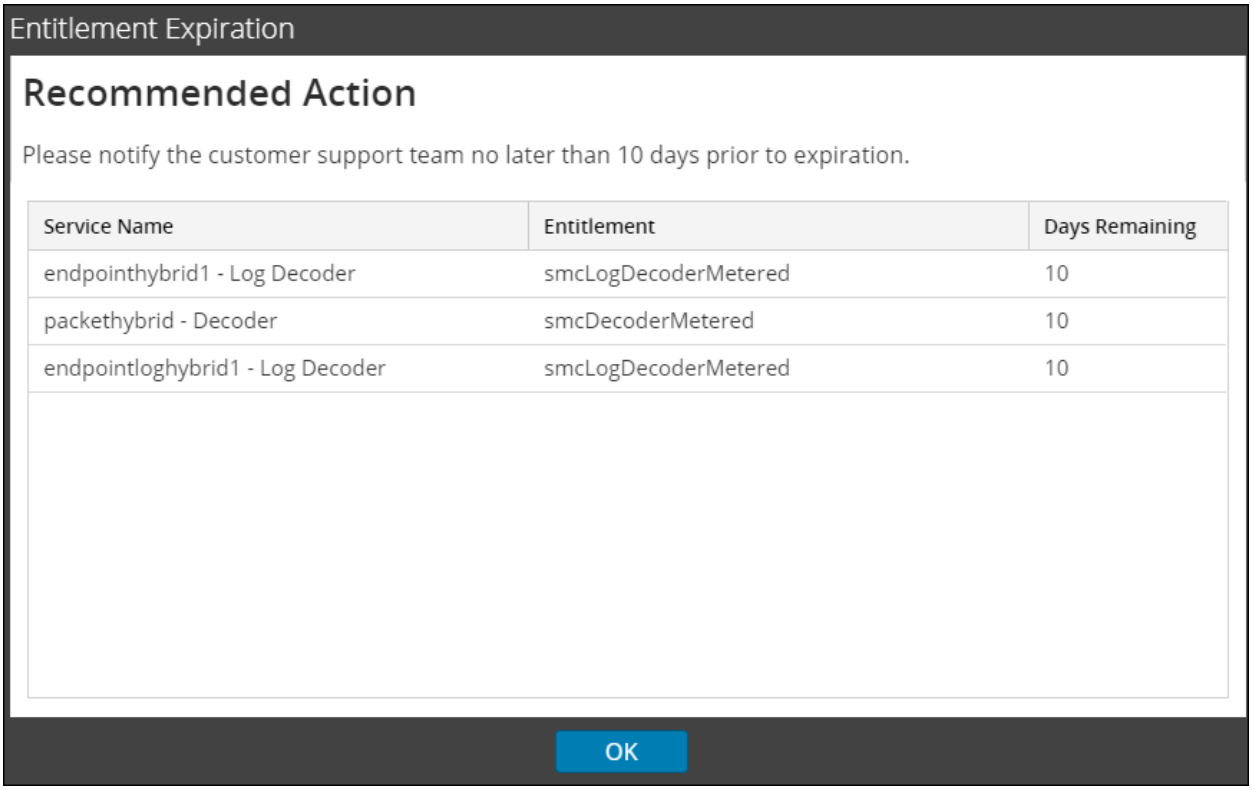
5. Click **Apply Notifications**.
   The settings are saved and go into effect immediately.

About Out-of-Compliance Banners

Banners let you know the status of your license and usage compliance.

Note: To help customers transition seamlessly to Throughput license for their network (packet) usage, NetWitness Platform provides a 45 days grace period. So, on upgrading to 11.6, you will not see the Out-of-Compliance notification or banner for 45 days from the day of upgrade for any network license usage. If you upgrade from 11.5.1 or later to 11.6, the available balance of the unused grace period will be applicable after the upgrade. For example
• If the grace period of 20 of the 45 days is used post upgrade to 11.5.1 or later then customers upgrading to 11.6 will get the remaining 25 days of the grace period.
• If the grace period of 45 days is used post upgrade to 11.5.1 or later the customers upgrading to 11.6 will not have any grace period left.

When a license for a specific service is about to expire, an **Entitlement Expiration** message is displayed with the list of services about to expire.



After you log on, either a red or yellow banner may be displayed, based on the license issue:

- A red banner is displayed if your license has expired, or exceeded the allotted usage, or if the license has internal errors.



Note: A red banner cannot be dismissed and the license issue must be resolved.

- A yellow banner is displayed if the license is approaching expiration or the license is nearing the allotted usage. The banner can be dismissed by clicking the **Reject** button.



Here is an of example of how the license usage is calculated and a way that you can resolve the license issues:

- Contracted daily usage can be exceeded three times in a calendar month. The fourth spike puts the system in an out-of-compliance state. If you are able to keep your usage within compliance for seven consecutive days until the end of calendar month, the Out-of-Compliance Red banner disappears.
- For example, if the fourth spike occurs on November 23, 2019, the Grace Period ends on December 31, 2019 and the Out-of-Compliance Red banner disappears.
- The breach period starts immediately after Grace Period ends. However, if your deployment is in a breach state, you can bring the state back to normal by keeping the usage in a compliant state for seven consecutive days.

Note: Even when the Red banner is displayed, there is no loss of functionality, all NetWitness appliances continue to work with full functionality. All other functionality is included in the license (ESA, storage, and so on).

Troubleshoot Licensing

This topic provides information about possible issues that NetWitness users may encounter when setting up licensing in NetWitness. You can look up explanations of issues and their solutions. NetWitness notifies users of issues using the popup notifications and the system log as described in the "Troubleshoot NetWitness" topic in the *System Maintenance Guide*.

Out-of-Compliance Banners

Problem	Possible Causes	Solutions
Yellow and Red Out-of-Compliance Banners	• A service is not licensed. • A license has expired, or is due to expire within the next two weeks. • Usage exceeds entitled limit. • Usage is approaching entitled limit.	• Contact Customer Support to buy or renew your license. • Reduce usage or • Adjust contracted usage amount

Common Log and Configuration Files

When troubleshooting licensing, the following files contain information that may help to diagnose the problem. Specific conditions for searching the files are described in the troubleshooting tables.

On the NetWitness Server, run `wget` for the following files after you ssh onto the NetWitness Server:

- `http://localhost:3333/api/1.0/features`
- `http://localhost:3333/api/1.0/instances`
- `http://localhost:3333//api/1.0/hostids`
- `http://localhost:3333/api/1.0/capability_request`

NetWitness Server Problems

This table lists possible problems with the NetWitness Server errors that can affect entitlements.

Problem	Possible Causes	Solutions
Some features have been mapped in the central Flexera server, but the NetWitness Server does not display them.	Ensure that the NetWitness Server is connected to the Internet.	To resolve the error execute a License Refresh as follows: 1. GSo to **(Admin) > Services > Licensing.** 2. In the License Details tab, click **Refresh Licenses**. **Note:** If the NetWitness Server is not connected to the Internet, try to synchronize offline. For more information, see Register the Server (Offline Capability Request)
NetWitness Platform Licensing page not showing any license information although there are services available.	Ensure that you have the required entitlements pulled down from the Central Flexera server.	To resolve the error execute a License Refresh as follows: 1. Go to **(Admin) > Services > Licensing**. 2. In the License Details tab, click **Refresh Licenses**.

Wrong License Mapping Issue

Problem	Possible Causes	Solution
Decoder license not available due to core appliances being removed from the NetWitness Server without releasing the license. Several core appliance licenses were not available for use.	Various possible causes.	Reset the license on NetWitness Server.

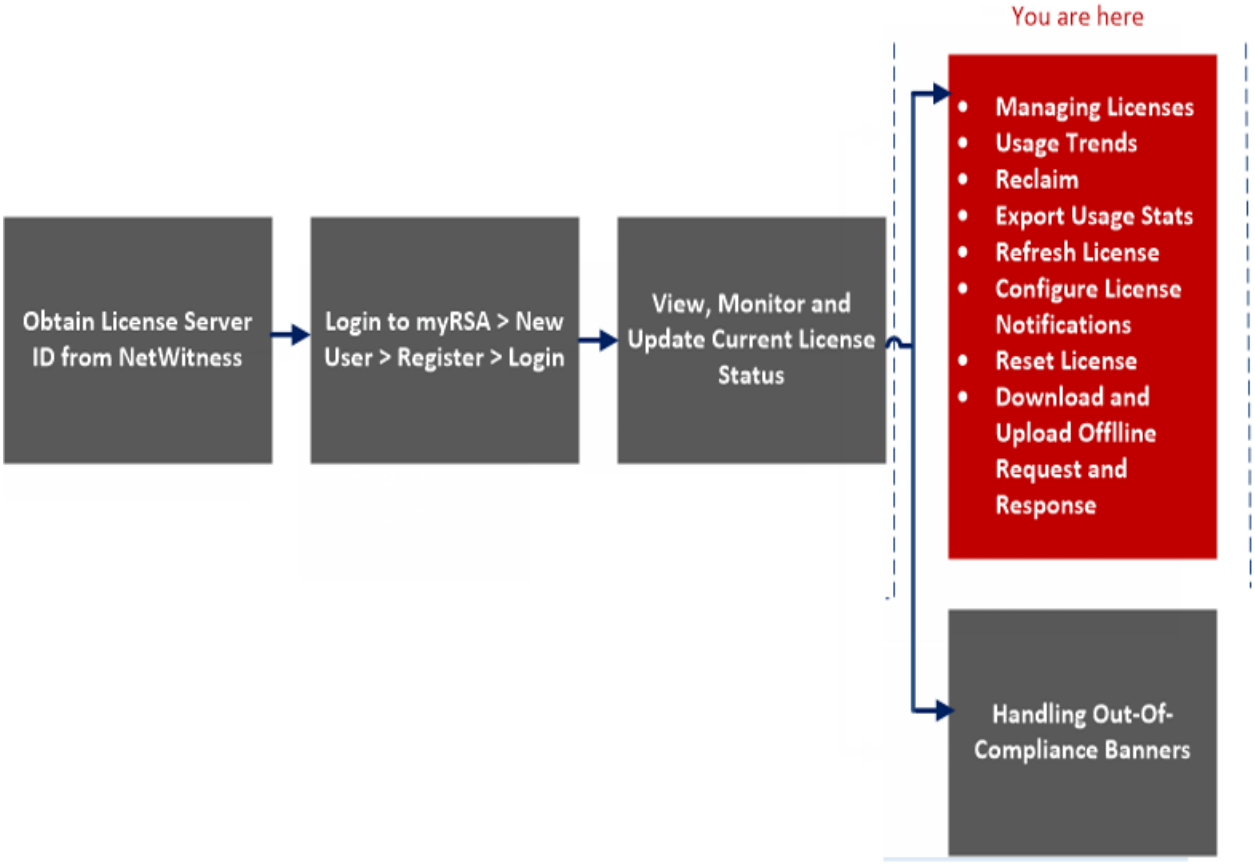
Problem	Possible Causes	Solution
Archiver DACs are not mapped to the license server with all other appliances' licenses.	Various possible causes.	1. Enter **1** in Quantity field to add for each license. 2. Select **Map Add-ons** at the bottom of the screen. 3. Click **Download Capability Request** and upload license to the Offline Capability Request in the User Interface under the **License** tab.
Two new appliances were installed: one Log Hybrid and one Log Archiver. Able to license the Log Hybrid, but the following error occurred when attempting to license the Archiver: "There is an issue with registering your product, please contact RSA Customer Support." Also, one of the Concentrators showed as a Trial license, and a separate Log Decoder showed as a Trial license when they should be licensed.	The new entitlements are not mapped to the license server.	Map add-ons to myRSA and upload the `.bin` file into the NetWitness User Interface.

License Details

This topic introduces the features of the System Licensing panel. NetWitness manages licensing through a Local License Server (LLS). Each client appliance is shipped with an installed LLS.

Workflow

This workflow shows the end-to-end licensing process.



What do you want to do?

Role	I want to...	Show me how...
Administrator	Register NetWitness Server	Obtain License Server ID from NetWitness Platform User Interface
Administrator	Synchronize NetWitness Server	Register the Server (Online Registration)
Administrator	Install product licenses from DLC	Access myRSA
Administrator	*Monitor and update current licenses.	View Current Licenses

Role	I want to...	Show me how...
Administrator	Configure licensing notifications.	Configure License Notifications
Administrator	View Out-of-Compliance banners	About Out-of-Compliance Banners

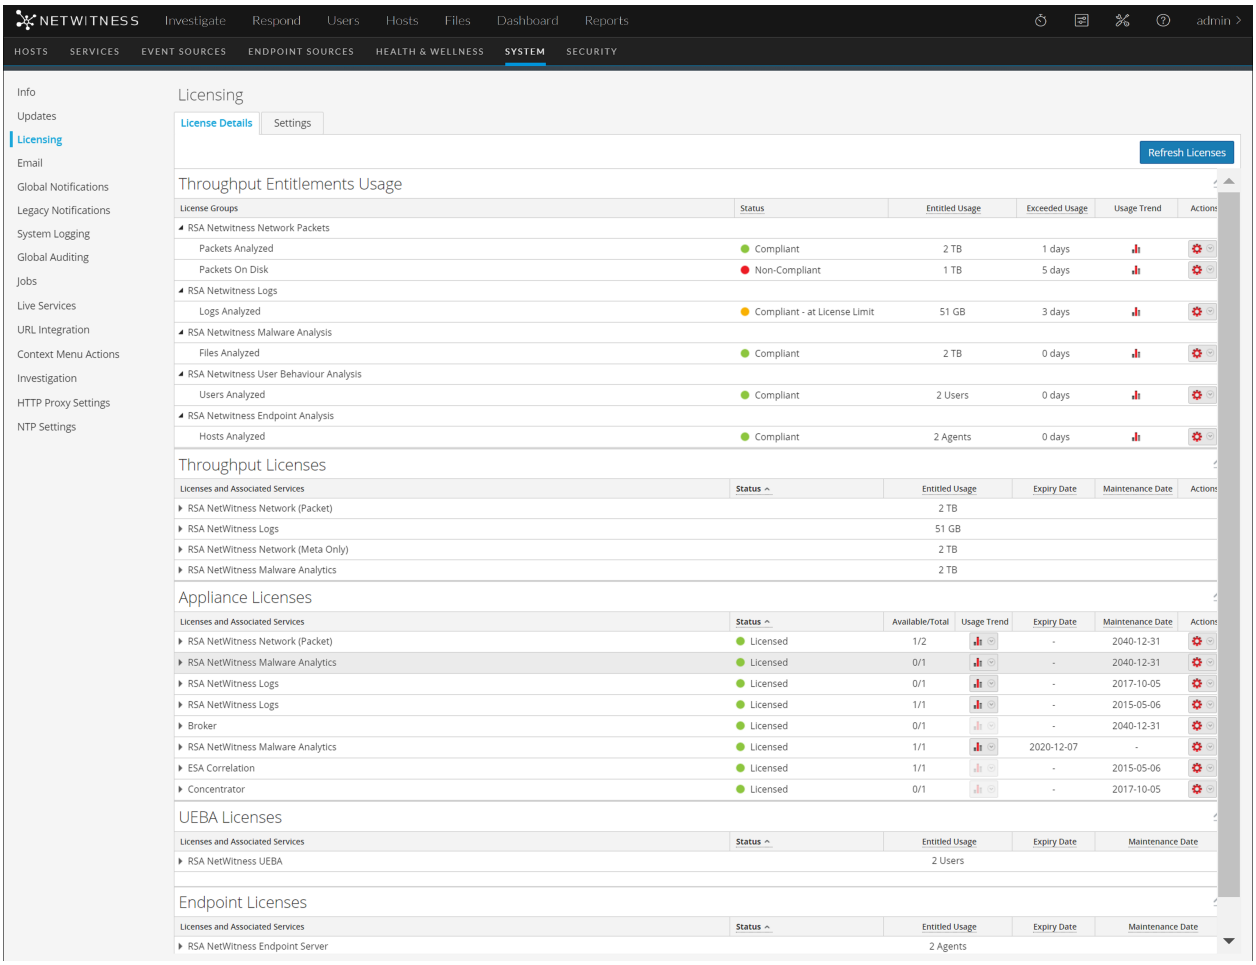
*You can perform this task here.

Related Topics

- Settings
- View and Export Usage Stats

Note: On initial start up, the usage shown in the Licensing page displays zero usage for the initial one hour.

Quick Look



The following table describes the features of the License Details tab.

1 **Refresh Licenses**:Refreshes and maps the new licenses.

2 Displays the following details of the Throughput license or licenses.

- **License and Associated Service** - Displays the license and the services assigned to it.
- **Status** - Displays the status of the license such as Expired License, Over Usage Limit, Near Usage Limit, and Within Usage limit.
- **Entitled Usage** - Displays the entitled usage for aggregate as well as individual services .
- **Exceeded Usage** - Displays the number of days the usage exceeded the entitled usage in the last 30 days for each individual service.
- **Usage Trend** - Displays the trend of how the license usage has been for a period of time for aggregate as well as individual services.
- **Expiry Date** - Displays the expiry date of the customer subscription contract.
- **Maintenance Date** - Displays the maintenance expiration date for the permanent license or licenses and the date on which the license or licenses expire.
- **Actions** - Displays the Licensing Actions button that offers the following options:
  - **Export Usage Stats**: Exports license usage statistics for the selected service or license.
  - **Reassign to Another License**: Reassigns an extensively used Throughput license to another unused Throughput license. This is applicable only for Throughput and Appliance license.

3 Displays the following details of the Appliance license or licenses.

- **License and Associated Service** - Displays the license and the services assigned to it.
- **Status** - Displays the status of the license such as Expired License, Over Usage Limit, Near Usage Limit, and Within Usage limit.
- **Available/Total** - Displays the number of available license and the total number of licenses.
- **Usage Trend** - Displays the trend of how the license usage has been for a period of time for aggregate as well as individual services
- **Expiry Date** - Displays the renewal date of the customer subscription contract.
- **Maintenance Date** - Displays the maintenance expiration date for the permanent license or licenses and the date on which the license or licenses expire.
- **Actions** - Displays the Licensing Actions button that offers the following options:
  - **Export Usage Stats**: Exports license usage statistics for the selected service or license.
  - **Reassign to Another License**: Move the license from Appliance to Throughput license.
  - **Reclaim**: This option gets activated when any appliance service is down. When you click **Reclaim** the license becomes available in the pool of the specific license service. This is applicable only for Appliance license.

4 Displays the following details of UEBA license or licenses.

- **License and Associated Service** - Displays the license and the services assigned to it.
- **Status** - Displays the status of the license such as Expired License, Over Usage Limit, Near

Usage Limit, and Within Usage limit.

- **Entitled Usage** - Displays the number of entitled users.
- **Actual Usage** - Displays the daily active users.
- **Exceeded Usage** - Displays the number of days the usage exceeded the entitled usage in the last 30 days for each individual service.
- **Usage Trend** - Displays the trend of how the license usage has been for a period of time for aggregate as well as individual services
- **Expiry Date** - Displays the renewal date of the customer subscription contract.
- **Maintenance Date** - Displays the maintenance expiration date for the permanent license or licenses and the date on which the license or licenses expire.
- **Actions** - Displays the Licensing Actions button that offers the following options:
  - **Sort Ascending**
  - **Sort Descending**

5 Displays the following details of Endpoint license or licenses.

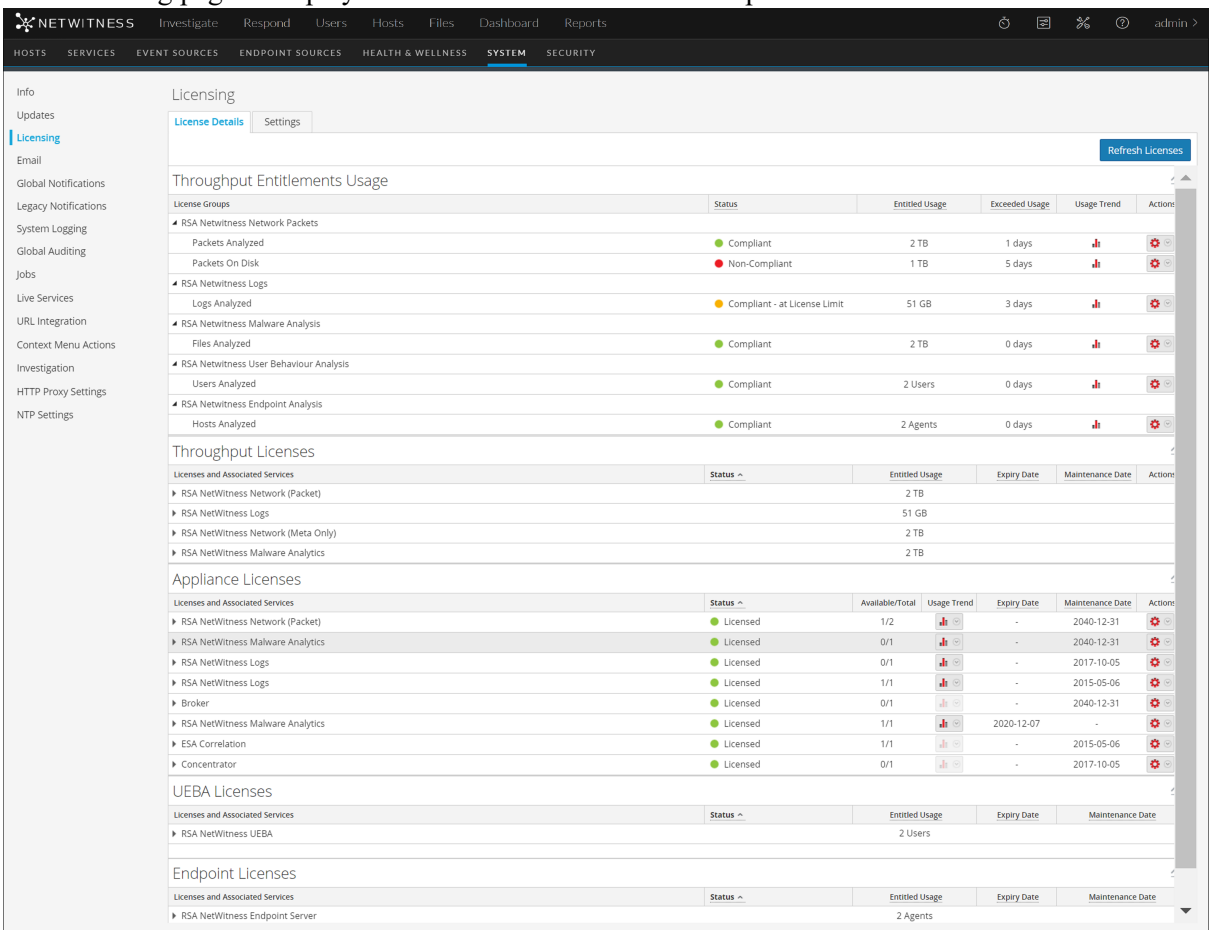
- **License and Associated Service** - Displays the license and the services assigned to it. For example, Endpoint Hybrid or Endpoint Log Hybrid.
- **Status** - Displays the status of the license such as Expired License, Over Usage Limit, Near Usage Limit, and Within Usage limit.
- **Entitled Usage** - Displays the users of the entitled agents .
- **Actual Usage** - Displays the number of daily active agents .
- **Exceeded Usage** - Displays the number of days on which the actual number of agents exceeded the entitled threshold.
- **Usage Trend** - Displays the trend of how the license usage has been for a period of time.
- **Expiry Date** - Displays the renewal date of the customer subscription contract.
- **Maintenance Date** - Displays the maintenance expiration date for the permanent license or licenses and the date on which the license or licenses expire.
- **Actions** - Displays the Licensing Actions button that offers the following options:
  - **Sort Ascending**
  - **Sort Descending**

Usage Trends

You can view the usage trend for aggregate licenses as well as individual licenses. The data is displayed for each day. When you click the usage trend icon, it opens a new window and displays the usage data as a chart. The user can use this data to analyze the license usage in the deployment.

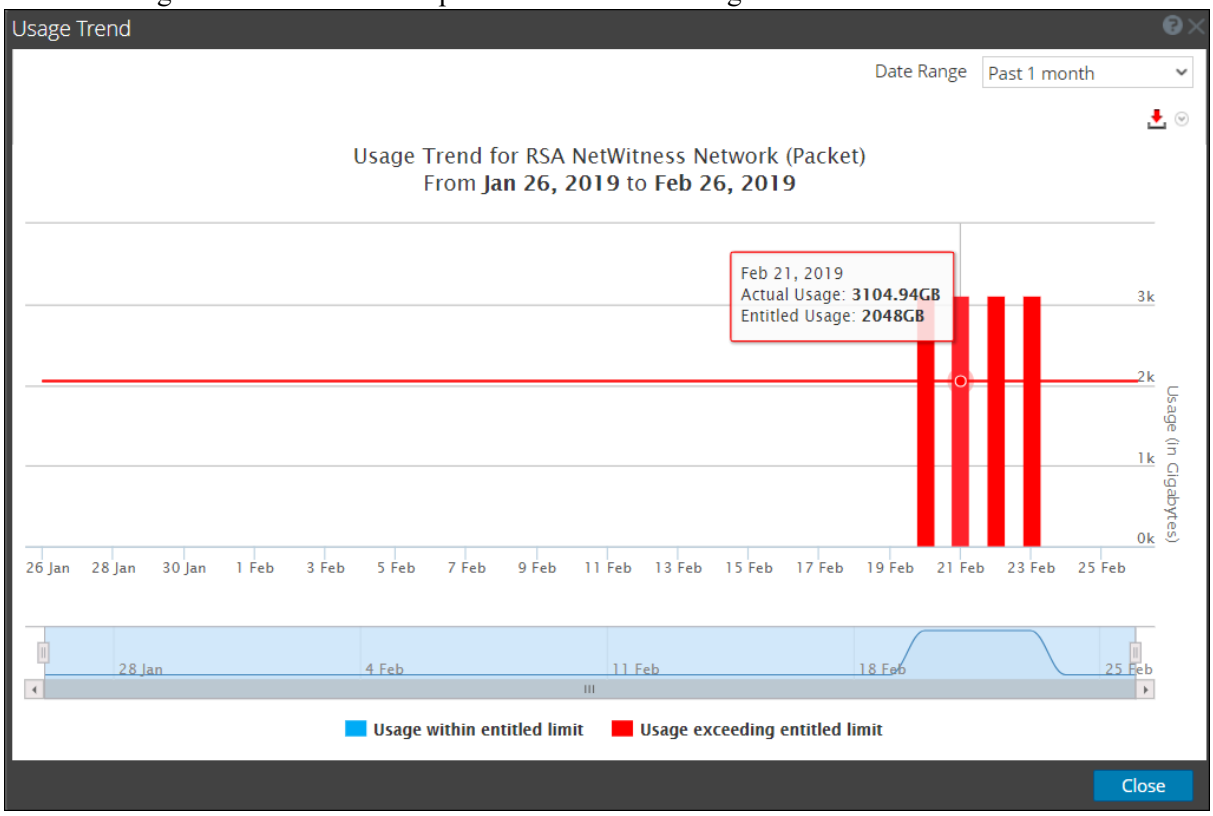
To access this view:

1. Go  (Admin) > **System** and select **Licensing**.
   The Licensing page is displayed with the License Details tab opens.

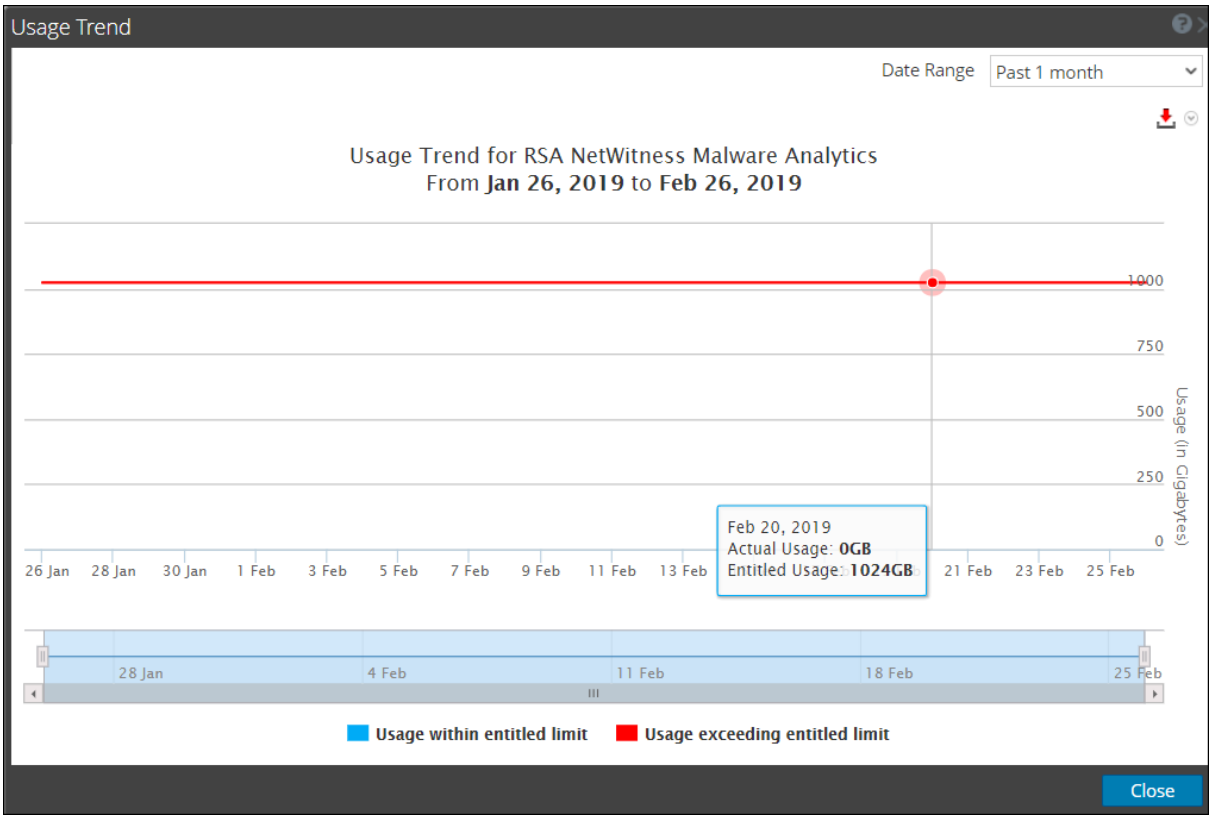


2. Select a service or license and click the  icon under the Usage Trend column.
   The Usage Trend window is displayed.

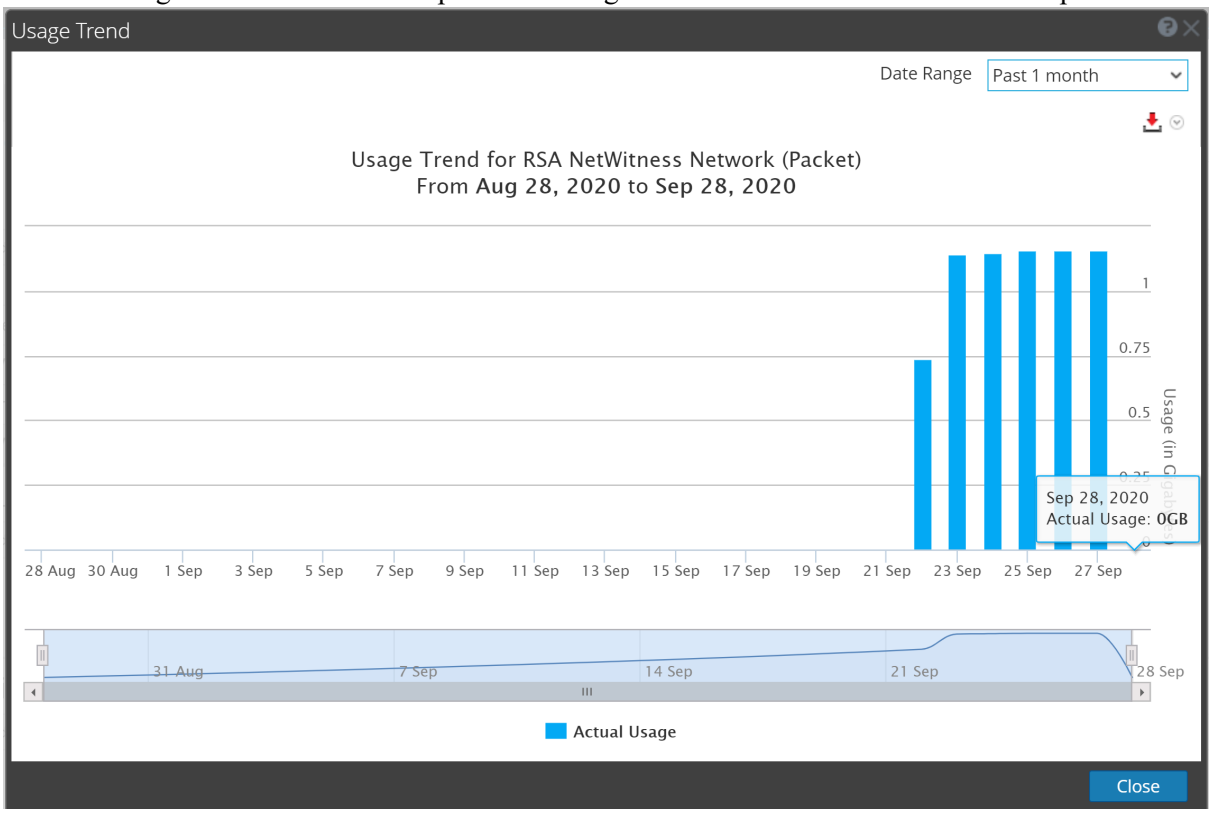
The following screenshot is an example of the Exceeded Usage Trend.



The following screenshot is an example of Within Limits Exceeded Usage Trend.


Usage Trend
Date Range
Past 1 month
Usage Trend for RSA NetWitness Malware Analytics
From Jan 26, 2019 to Feb 26, 2019
1000
750
500
250
0
Usage (in Gigabytes)
Feb 20, 2019
Actual Usage: 0GB
Entitled Usage: 1024GB
26 Jan
28 Jan
30 Jan
1 Feb
3 Feb
5 Feb
7 Feb
9 Feb
11 Feb
13 Feb
21 Feb
23 Feb
25 Feb
28 Jan
4 Feb
11 Feb
18 Feb
25 Feb
Usage within entitled limit
Usage exceeding entitled limit
Close

The following screenshot is an example of the Usage Trend chart for a license with multiple services.



The threshold limit is indicated as a horizontal red-dotted line across the chart. When actual data usage exceeds the entitled daily usage, those days are indicated by red colored bars.
The chart can be generated to collect data for 1 month, 3 months, 6 months, 12 months or custom date range.



These charts can be exported in CSV, and PDF formats by clicking the icon.

Note: The maximum range for which trend can be viewed is 12 months. The pan-zoom bar at the bottom of the chart can be used to narrow the chart to a smaller time range for better visibility in that range.

If the entitled limit exceeds the usage limit then an appropriate out-of-compliance banner is displayed. For more information, see About Out-of-Compliance Banners.

Reassign Service Licenses

You can move service between licenses only if a similar service license is available.

You can move the services between the following licenses.

- Throughput License to Throughput License
- Throughput License to Appliance License
- Appliance License to Throughput License
- Appliance License to Appliance License

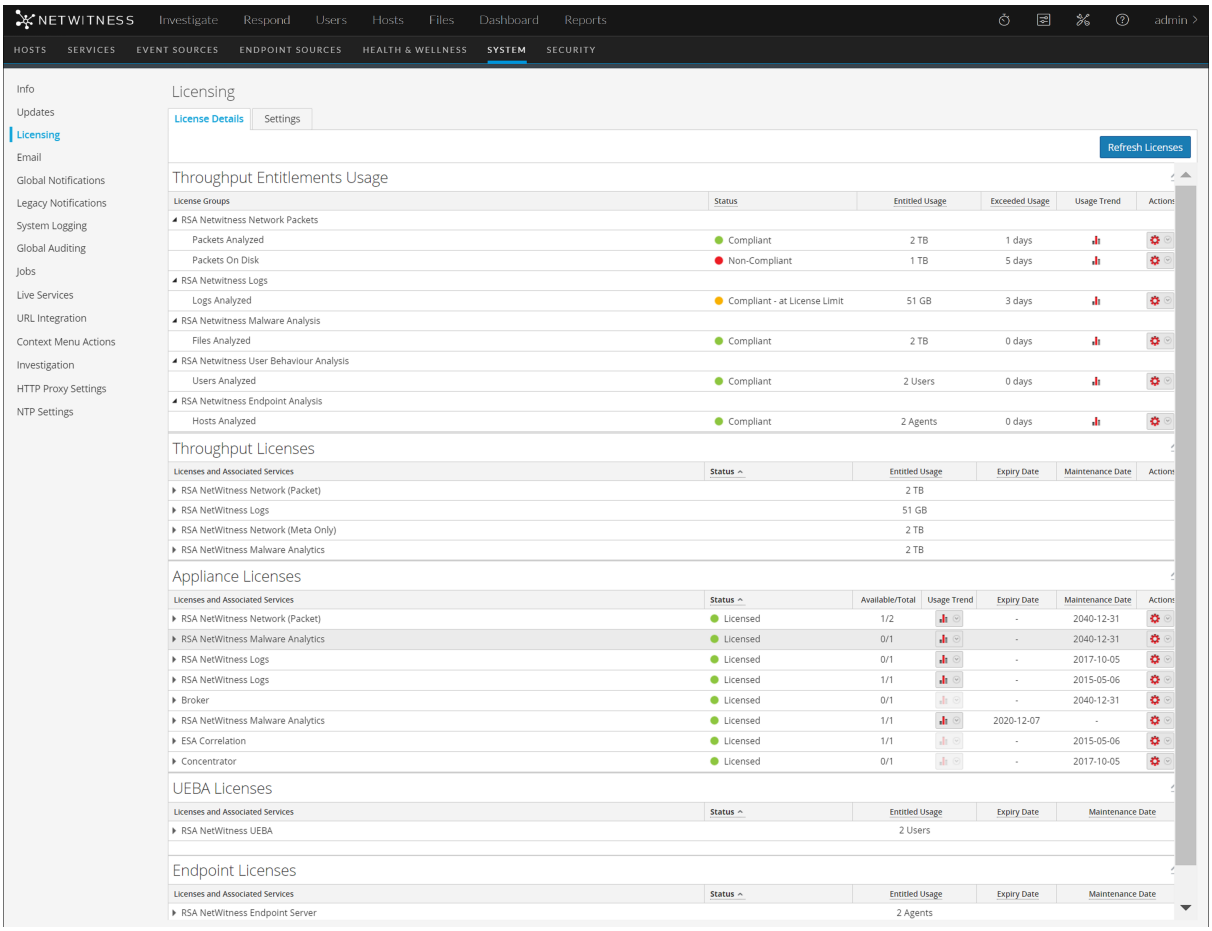
Note: If you want to change the licenses by moving between throughput and appliance, you can do this by selecting the license under the actions of each license. Trail licenses cannot be moved.

You can move the licenses for following reasons:

1. If the subscription based license has expired
2. If there are any unused available license for any service
3. If you want to reduce the usage of the service

To access this view:

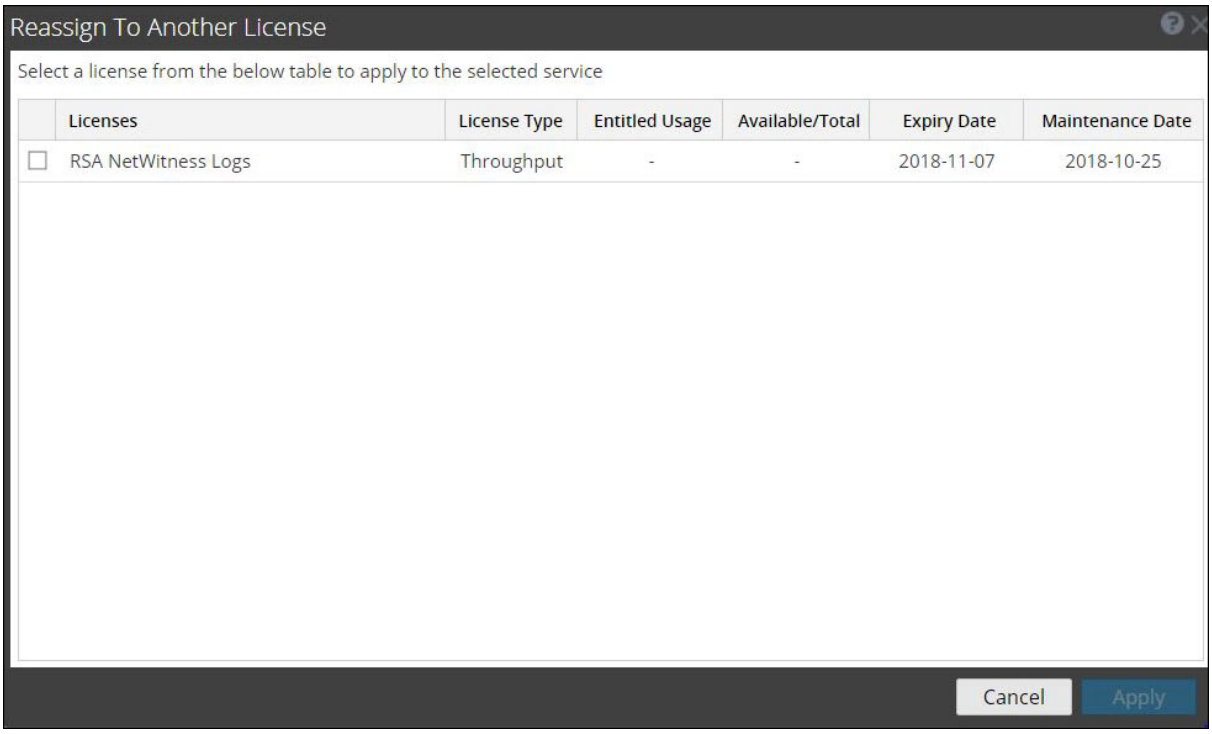
1. Go  **(Admin) > System** and select **Licensing**.
   The License Details tab is displayed.



2. Select a service of which you want to move the license.

3. Click  , and select the **Reassign to Another License** option.

   The Reassign To Another License dialog is displayed with a list of the available licenses that can be moved.



4. Select a license to be applied for the selected service.
5. Click **Apply**.

Note: The above procedure can be used to move both Throughput Licenses and Appliance Licenses. However, the UEBA and Endpoint licenses cannot be moved.

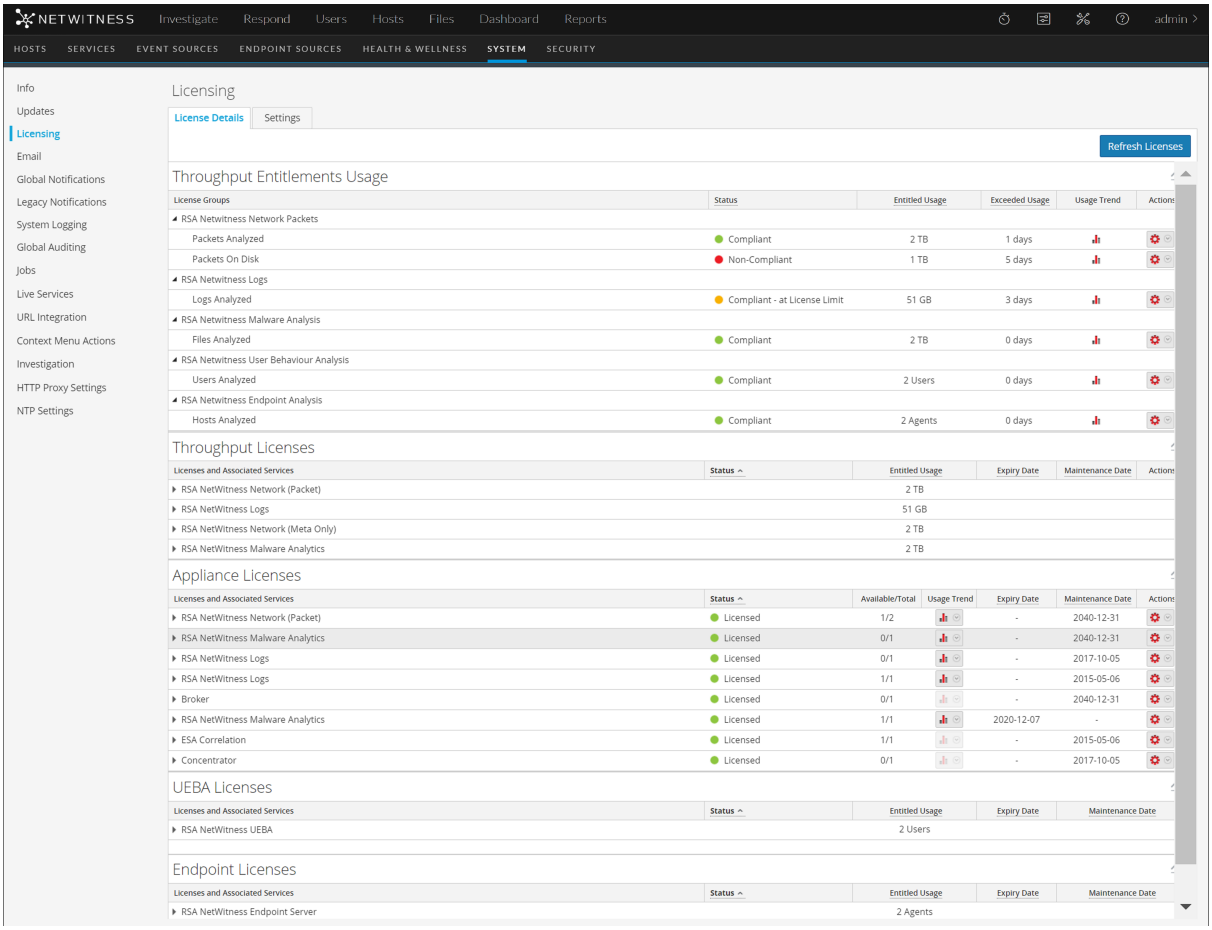
View and Export Usage Stats

NetWitness provides the ability for Administrators to view usage statistics of device types that are eligible for a Throughput , Appliance Licenses such as Log Decoder, Decoder, Malware, UEBA, and Endpoint Server licenses. An hourly usage statistics is captured for all supported services connected to the NetWitness Server and are made available in CSV, and PDF formats. Metrics can be tracked securely, allowing Administrators to save data locally on their systems to use in reporting usage compliance.

To access Export Usage Stats:

1. Go **(Admin) > System** and select **Licensing**.

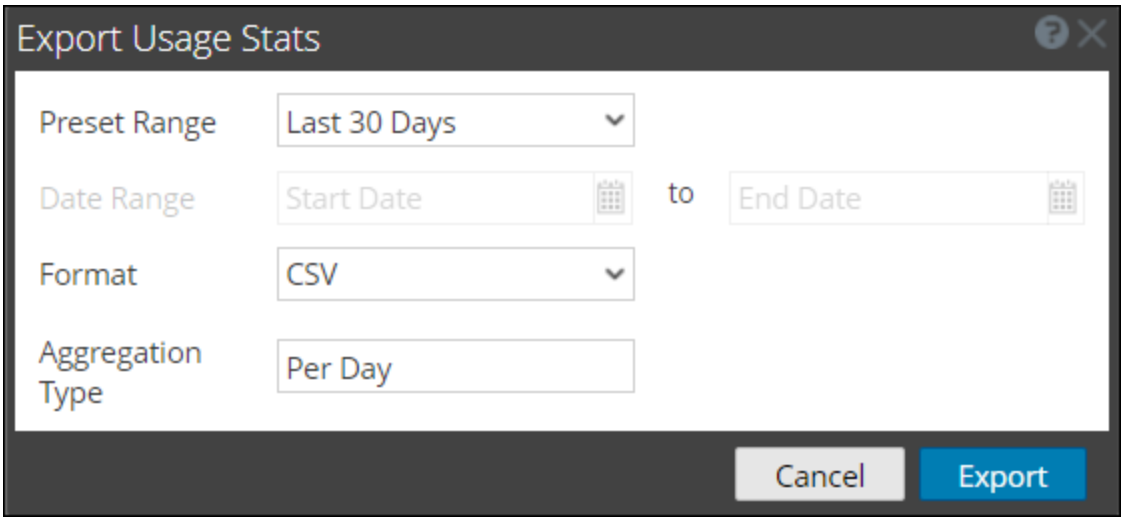
   The License Details tab is displayed.



2. Select a service for which you want to export usage stats.

3. Click , and select the **Export Usage Stats** option.

   The **Export Usage Stats** dialog is displayed.



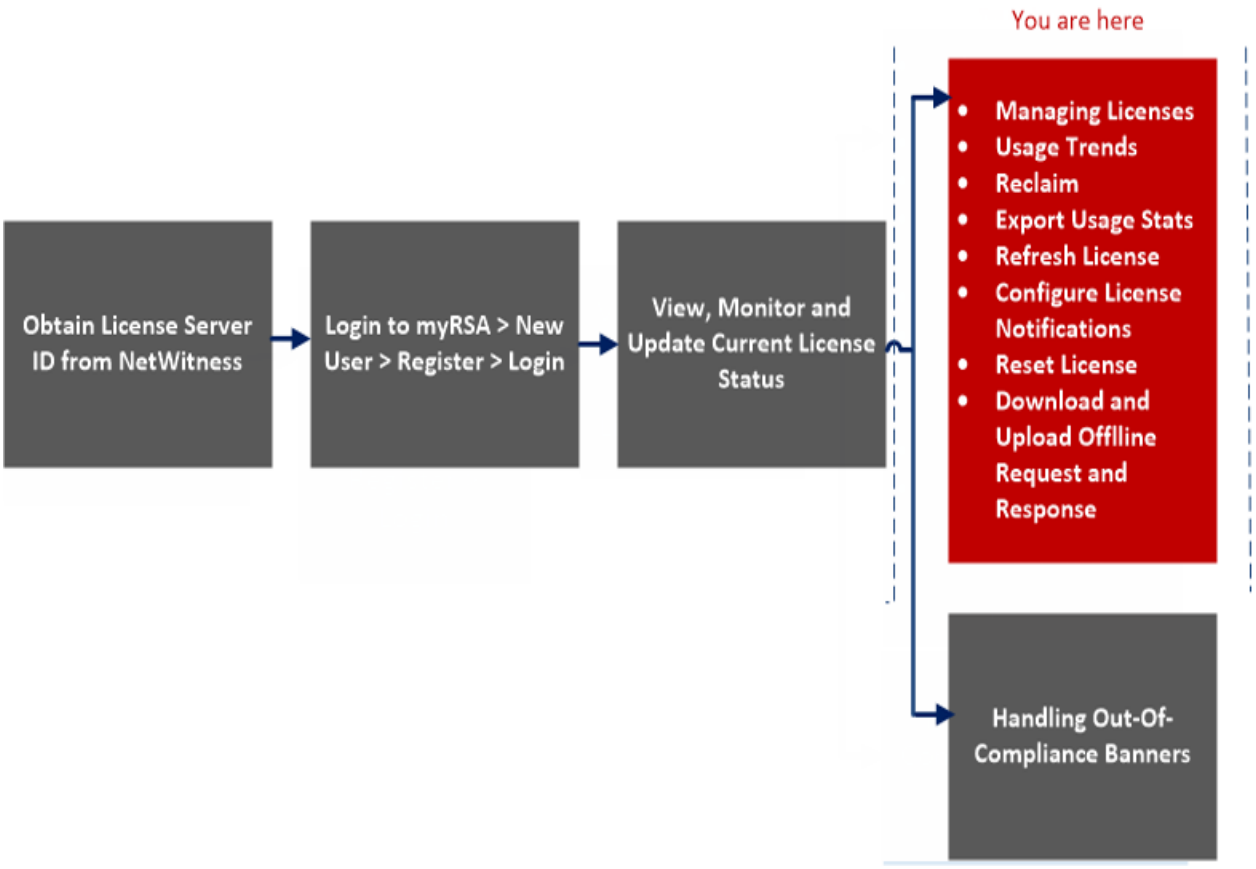
4. Select a **Preset Range**, **Date Range** (in case of custom range), **Format** and **Aggregation Type** to generate the report.
5. Click **Export** to export the report or click **Cancel** to return to the **License Details** tab.

Settings

This topic describes the notification settings for the NetWitness in the **Licensing panel** > **Settings** tab.

Workflow

This workflow illustrates the end-to-end licensing process.



What do you want to do?

Role	I want to...	Show me how...
Administrator	Register NetWitness Server	Obtain License Server ID from NetWitness Platform User Interface
Administrator	Synchronize NetWitness Server	Register the Server (Online Registration)
Administrator	Install product licenses from DLC.	Access myRSA
Administrator	Monitor and update current licenses.	View Current Licenses
Administrator	*Configure licensing notifications.	Configure License Notifications

Role	I want to...	Show me how...
Administrator	View Out-of-Compliance banners	About Out-of-Compliance Banners

*You can perform this task here.

Related Topics

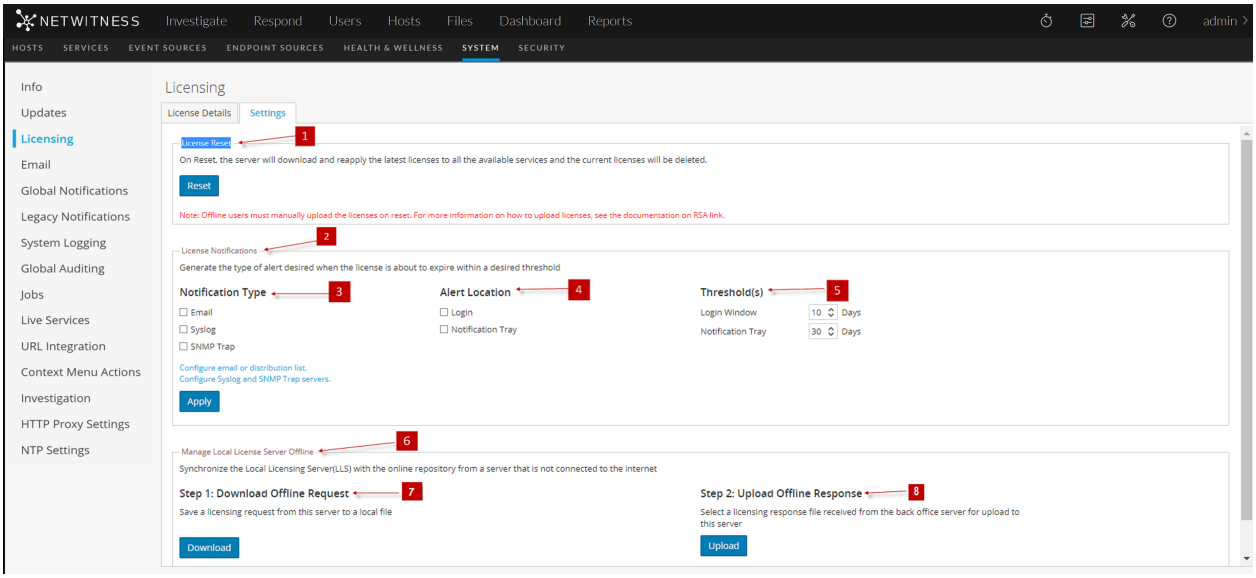
Obtain License Server ID from NetWitness Platform User Interface

Configure License Notifications

Quick Look

From the **Settings** tab you can:

- Configure licensing notifications.
- Download an Offline Capability Request in NetWitness for submission to DLC.
- Within 72 hours, upload to NetWitness an Offline Response that was received from DLC.



The following table describes the **Settings** tab features.

1	Displays the **License Reset** panel which applies the default license settings.
2	Displays the **Licensing Notifications** panel.
3	Displays the **Notification Type**. There are three types of notifications: • **Email:** Checkbox to receive a notification of approaching license expiration in an email message. The email is sent to the configured email or distribution list. • **Syslog:** Checkbox to receive a notification of approaching license expiration in an syslog message. The syslog is generated in accordance with the settings in the Syslog Auditing Settings. • **SNMP Trap:** Checkbox to receive a notification of approaching license expiration in an SNMP trap. The trap is generated in accordance with the settings in the SNMP Auditing

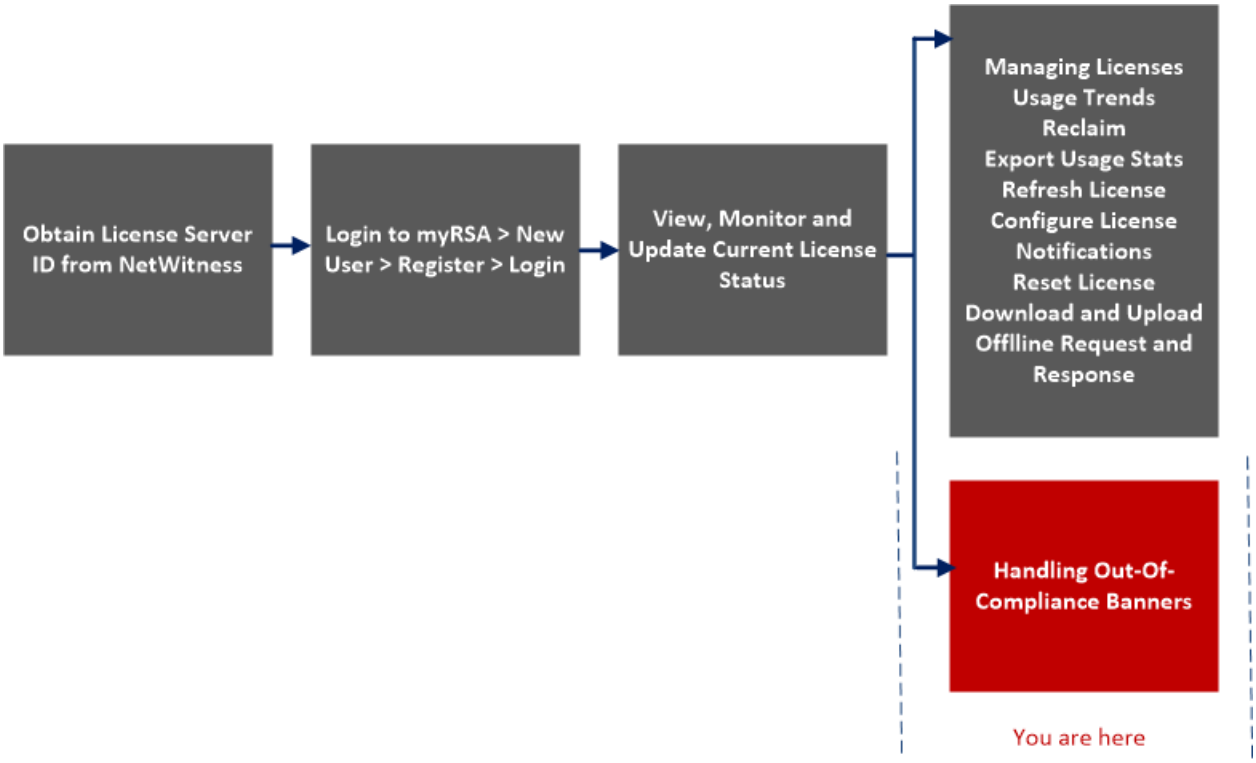
	Settings.
4	Displays the type of **Alert Notification**. • **Login:** Select this checkbox to receive a notification of your approaching license expiration when you log on to NetWitness. The **Login Window Threshold** field specifies the number of days before the license expires to display the notification at log on. • **Notification Tray:** Select this checkbox to receive a notification of approaching license expiration in the Notifications tray.
5	Displays the **Threshold** field, which specifies the number of days before the license expires to send a notification to the Notifications tray.
6	Displays the **Manage Local License Server Offline** panel.
7	Displays the **Download Offline Request** button. This button enables you to download a request from the NetWitness Platform LLS into a local file for processing by a back-office server. The downloaded bin file should be uploaded to DLC to generate the offline response.
8	Displays the **Upload Offline Request** button. This button enables you to browse for an offline response that you received from the back-office server, and uploads the selected response to NetWitness. The file must be uploaded within 72 hours after receiving the file.

Out-of-Compliance Reference

This topic describes the workflow when your license is out of compliance.

Workflow

This workflow illustrates the end-to-end licensing process.



What do you want to do?

Role	I want to...	Show me how...
Administrator	Register NetWitness Server	Obtain License Server ID from NetWitness Platform User Interface
Administrator	Synchronize NetWitness Server	Register the Server (Online Registration)
Administrator	Install product licenses from DLC.	Access myRSA
Administrator	Monitor and update current licenses.	View Current Licenses
Administrator	Configure licensing notifications.	Configure License Notifications
Administrator	*View Out-of-Compliance banners.	About Out-of-Compliance Banners

*You can perform this task here.

Note: When throughput devices are under trial period, warning messages will not be displayed unless usage is observed on the corresponding device.

Out-of-Compliance State

The following sample banner is displayed when a license expires:



If your license has internal errors, the following banner is displayed:



In addition to a red banner being displayed during system log on, an Out of Compliance Acknowledgment dialog is also displayed. Click **Accept** to continue using your NetWitness product.

Licenses may become out-of-compliance due to reasons provided in the following table:

Red Banner Message	Possible Causes	Solutions
One or more services is not licensed.	Trial license period has expired. There are pre-11.0.0.0 services in the deployment that are not licensed.	Contact Sales team to procure a NetWitness license. Upgrade the services to NetWitness version 11.0.0.0 or later.
One or more licenses is expired.	Log ingestion usage has been observed after the date of renewal. The license is not valid anymore for the corresponding usage.	Contact Sales team to renew or resolve the license.
You have exceeded license usage limits.	If the allotted daily usage is exceeded on four or more occasions, the Grace Period begins. The Grace Period begins on the day of the fourth occurrence and ends at the end of the following calendar month. Seven continuous days of standard usage will end the Grace Period. If the daily allotted usage is still being exceeded at the end of the Grace Period, the 30-day Breach Period begins. Seven continuous days of standard usage will end the Breach Period.	Contact Sales to extend or increase your allotted usage by purchasing a NetWitness license.

Red Banner Message	Possible Causes	Solutions
Your Trial license has internal errors.	An internal licensing issue was reported during your Out-of-the-Box Trial period.	Contact Technical Support to resolve this issue.

Note: If a license has not been installed within 90 days, you must contact Sales to purchase a NetWitness license.

License Approaching Out-of-Compliance

When your license is approaching expiration, or it is nearing its allotted usage, a yellow banner with a brief description is displayed. A yellow banner is displayed 14 days before your license is due to expire. You will also see a yellow banner if you are approaching your allotted license usage. You can get rid of the yellow banner by clicking the **Dismiss** button.

The following sample banner is displayed in the NetWitness screen if your license is approaching its allotted usage:



The following table explains the messages that are displayed when you see a yellow banner.

Yellow Banner Message	Possible Causes	Solutions
You are nearing license usage limits.	One or more Throughput licenses has exceeded your allotted usage for three times during the current calendar month. The fourth time that you exceed your allotted usage during the current month will push the deployment into an Out-of-Compliance state.	Contact RSA Sales if your allotted usage spikes four times within a calendar month.
One or more licenses is expiring.	One or more licenses is due to expire within 14 days. Or Log ingestion usage has been observed. The license is not valid anymore for the corresponding usage.	Contact RSA Sales to purchase a new license.

NetWitness® Platform XDR

Version 12.0

Live Services Management Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

November, 2022

Contents

Live Services Management **7**

NetWitness Live 7

NetWitness Feedback and Data Sharing 7

For Debian Linux and NetWitness Endpoint Users 7

Deploy Content **8**

Create Live Account 9

Reset the Password for Your Live Account 12

Set Up Live Services on NetWitness 13

Deploy Content using Live Content UI 15

Live Services Required Procedures 16

Find and Deploy Live Resources 17

Find Resources in Live 17

Deploy Resources in Live 18

Manage Live Resources 23

Manage Subscription and Deployment 23

Remove a Deployed Resource 24

Deploy a Resource Bundle 24

Download Resources 24

Set Up Data Feeds 24

Search and Download Content from NetWitness XDR Cloud Services Live 25

Quick Search for Content 25

Advanced Search for Content 26

Download Content 27

Additional Procedures 29

Export Data to RSA 30

About Live Feedback 30

Download Live Feedback Historical Data 30

Share Telemetry Data to NetWitness 30

Packaging Resources 32

Create and Deploy Resource Package Use Case 32

Prerequisites to Create a Resource Package 32

Creating a Resource Package 32

Creating Threat Package 33

Deploying a Threat Package 34

Manage Custom Feeds 36

Custom Feed Creation 36

Sample Feed Definition File 36

Feed Definition Equivalents for Custom Feed Wizard Parameters 37

Creating a Custom Feed 40

Import Certificates for HTTPS Service 46

Create a STIX Custom Feed ... 48
MetaCallback Feeds using CIDR Index Range for IPv4 and IPv6 ... 53
Creating and Managing an Identity Feed ... 54
Import the SSL Certificate ... 61
Cannot Verify Identity Feed URL ... 61
Investigating an Identity Feed ... 62
Editing a Feed ... 64
Removing a Feed ... 66
Subscribing to Live Resources ... 68
Subscription Updates ... 68
Adding Subscribed Resources for Deployment to Services ... 69
Deleting a Subscription ... 69
Removing Subscribed Resources from the Deployments Subscriptions Grid ... 70
Subscribe and Unsubscribe to a Resource ... 70
Viewing Subscribed Resources Selected to Deploy on Services ... 72
Miscellaneous Live Services Procedures ... 73
Displaying Resource Details in Live Resource View ... 73
Downloading a Resource ... 74
Locating and Removing a Deployed Resource from Services ... 74
Showing Results as a List or in Detail ... 75
Viewing Resource Details ... 76
References ... 78
Live Configure View ... 79
Deployments Tab ... 80
Groups Panel ... 81
Subscriptions Panel ... 81
Subscriptions Tab ... 82
Toolbar ... 82
Grid ... 83
Discontinued Resources Tab ... 84
Groups Panel ... 84
Discontinued Resources on Service Panel ... 85
Live Feeds View ... 86
Toolbar ... 86
Feeds Grid ... 87
Live Resource View ... 88
Resource Details ... 88
Resource View Toolbar ... 89
Live Search View ... 91
Search Criteria Panel ... 91
Matching Resources Panel ... 93
Detailed Results ... 93

Grid Results 94
See Also 95
Live Search Content View 96
Search Content Panel 97
Search Results Panel 99
Content Details Panel 100
Resource Package Deployment Wizard 103
Features 103
Package Tab 103
Resources Tab 104
Services Tab 104
Review Tab 105
Deploy Tab 106
NetWitness Live Registration Portal 108
NetWitness Feedback and Data Sharing 110
Additional Live Services 110
Live Feedback 110
File Reputation 110
Troubleshooting Live Services 112
OutOfMemoryError on Context Hub Server 112
Troubleshooting Live Connect Threat Data Sharing 112
Query Log Retrieval Sample 113
System Logging: Debug 113
Policy-based Centralized Content Management 115
About Policy-based Centralized Content Management 117
Migrate Content from Core Services to Content Library 119
About Content Library 121
Import Content to Content Library 121
Create an Application Rule 122
Edit Application Rule 122
Delete Application Rule 123
View Application Rule Details 123
Create a Network Rule 123
Edit Network Rule 124
Delete Network Rule 124
View Network Rule Details 125
About Groups 125
Create a Group 125
View a Group 126
Delete a Group 126
Edit a Group 127
About Policies 127

Create and Publish Policies 128
Clone a Policy 129
Delete a Policy 130
Edit a Policy 130
View a Policy 131
Enable Content for a Policy 132
Disable Content for a Policy 132
References 132
Content Library Tab 132
Groups Tab 134
Policies Tab 138
Appendix A: Endpoint Risk Scoring Rules 142

Live Services Management

RSA NetWitness Live is the gateway to a rich environment that offers access to feeds, tools, and other resources.

NetWitness Live

Live is the component of NetWitness that manages communication and synchronization between NetWitness services and a library of Live content available to NetWitness customers. Live provides a simple interface for browsing, selecting, and deploying content from the NetWitness Live Content Management System to NetWitness services and software. In addition to managing feeds from the CMS Library, Live allows users to deploy custom feeds and packages.

Note: Any customer with valid maintenance can access RSA Live.

The content management system (CMS) library (known as *Live*) is a valuable source of the latest internet security resources for NetWitness customers. It provides a view into the collective intelligence and analytical skills of the worldwide security community to ensure that users have the most current visibility into attack vectors.

Live gathers the best advanced threat intelligence and content in the global security community - the ideas, research, ongoing tracking, and analysis - and brings it directly into the user's security operations center to definitively classify computers associated with botnets, malware, and other malicious exploits. Live aggregates, consolidates, and illuminates only the most pertinent information relevant to an organization on a real-time basis.

NetWitness Feedback and Data Sharing

Live Feedback is intended to help improve NetWitness. Once you set up and configure a Live account, usage data is shared with RSA.

For more details, see NetWitness Feedback and Data Sharing.

For Debian Linux and NetWitness Endpoint Users

If you are upgrading to NetWitness 11.5 or later, and you are using NetWitness Endpoint and also have any Debian Linux endpoint systems, NetWitness recommends that you go to Live and download the following application rules:

- autorun debian package mismatch
- autorun file path not part of debian package
- debian package hash mismatch in important system directory
- debian package hash mismatch
- file path not part of debian package in important system directory
- file path not part of debian package

Deploy Content

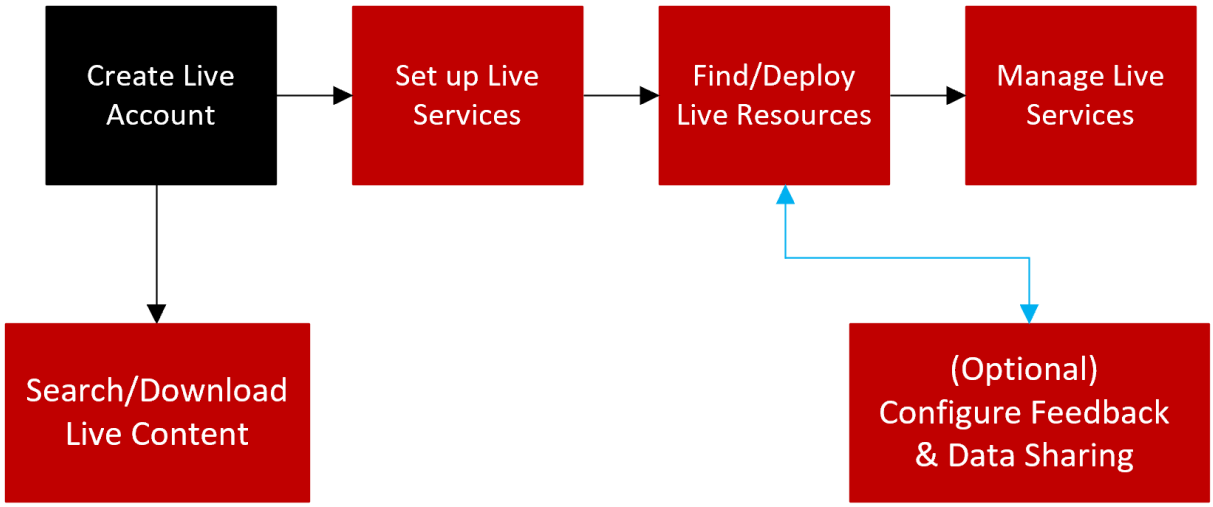
This section explains the different ways available to deploy content:

- Deploy Content using Live Content UI
- Policy-based Centralized Content Management

Create Live Account

Note: The NetWitness Live Registration Portal now has a new user interface and supports email verification.

You must create a Live account using the NetWitness Live Registration Portal (https://live.netwitness.com/registration) on the Live server. Live Account is required to access all Live services including CMS. The CMS Library provides access to all NetWitness content in one place where you can view, search, deploy, and subscribe to NetWitness content.



Make sure the following are available to set up a NetWitness Live account:

- Active internet connection to access the portal.
- A valid and registered NetWitness License Server on the Flexera Server, before you can register for a Live account. You can view the License ID on the **(Admin) > System > Info** panel.

Note: If the License Server is not set up, contact NetWitness Customer Support.

To create a Live Account:

1. Access the NetWitness XDR Cloud Services Live Registration Portal using the URL: https://live.netwitness.com/registration

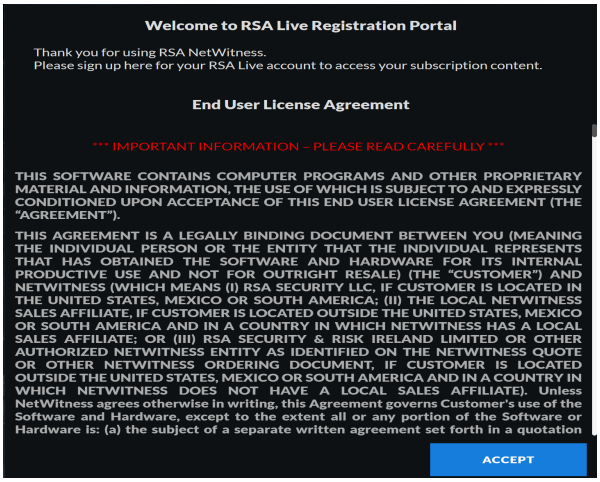
   The NetWitness XDR Cloud Services Live sign up page is displayed.



2. Click **Sign Up For Live**.

   The End User License Agreement page is displayed.

   Read the Terms and Conditions carefully and click **Accept**.



3. In the **Sign Up for NetWitness Live Account** page, enter all the fields:
   - The **First Name** and **Last Name** of the user.
   - The **Company** for which the Live Account is being created.
   - The **Email address** you enter will be used to receive the verification code for your new Live account and other notifications related to the Live account.
   - The **License ID** can be viewed on **(Admin) > System > Info** panel.
   - The **Username** and **Password** for the Live Account.



4. Click **Create Account**.
   You will be directed to **Confirm Sign up** page.
5. Enter the **Confirmation Code** sent to your registered email address.
   Click **Confirm**.
   You can see the confirmation message for your NetWitness Live Account registration.

   Note: You cannot create more than one Live account for the same License ID. For additional license, contact NetWitness Customer Support .
6. Once the account is created, enter your credentials and click **Sign In** to access the NetWitness XDR Cloud Services Live.
7. After you sign in, you can perform the following:
   - Search and Download Content from NetWitness XDR Cloud Services Live
   - Share Telemetry Data to NetWitness

Reset the Password for Your Live Account

If you want to reset the password for your Live Account, do the following:

1. Access the NetWitness XDR Cloud Services Live Registration Portal using the URL:

   https://live.netwitness.com/registration

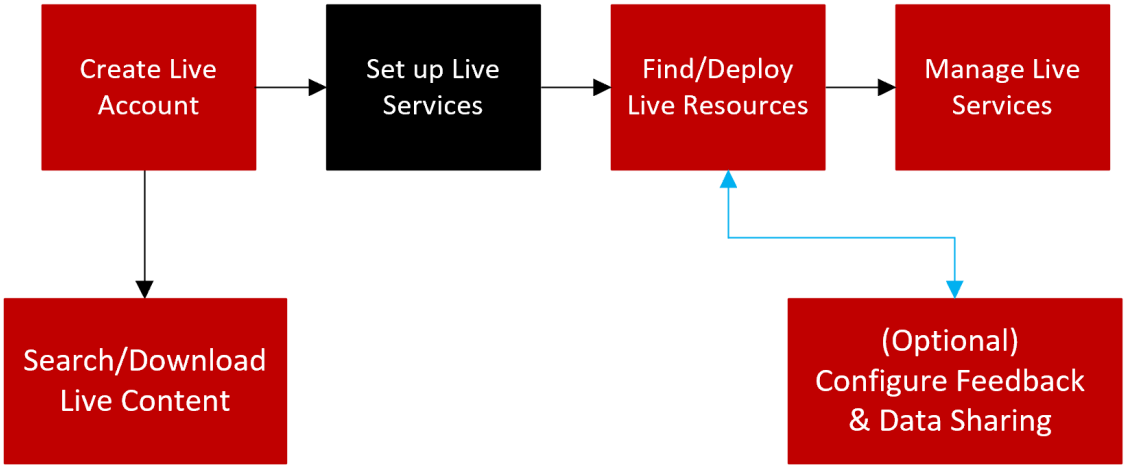
   The NetWitness XDR Cloud Services Live sign up page is displayed.



2. On the Sign Up page, click **Forgot Password?**.
3. Enter your **Username** and click **Send Code**.
   A verification code will be sent to your registered email address.
4. Enter the **Verification Code** and **New Password** on the Reset Password page and click **Submit**.

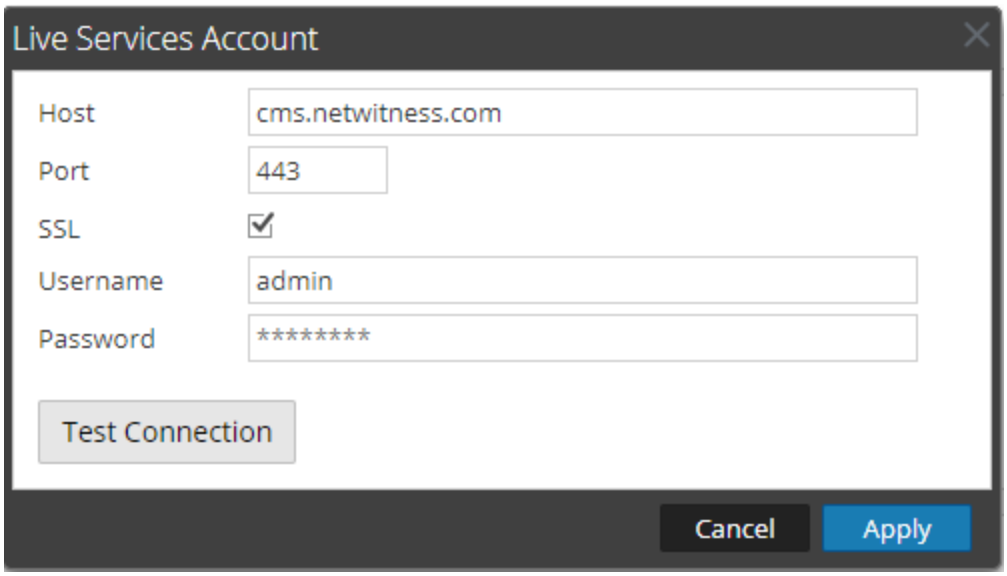
Set Up Live Services on NetWitness

To set up Live on NetWitness, you configure the connection and synchronization between the CMS server and NetWitness. The user interface for this setup is the **(Admin) > System > Live Services** Configuration panel.



To configure the connection to the CMS Server:

1. Navigate to **(Admin) > System > Live Services**.
2. Click Sign In and enter your credentials in the Live Services Account dialog box.



3. Click Test Connection to make sure your connection is working.

4. If the test is successful, click **Apply**. If not, contact NetWitness Customer Support for help connecting to the Live server.
5. Configure the timing for synchronization of NetWitness with updates from Cloud Services Live.

For more details, see the "Configure Live Services Settings" topic in the *System Configuration Guide*.

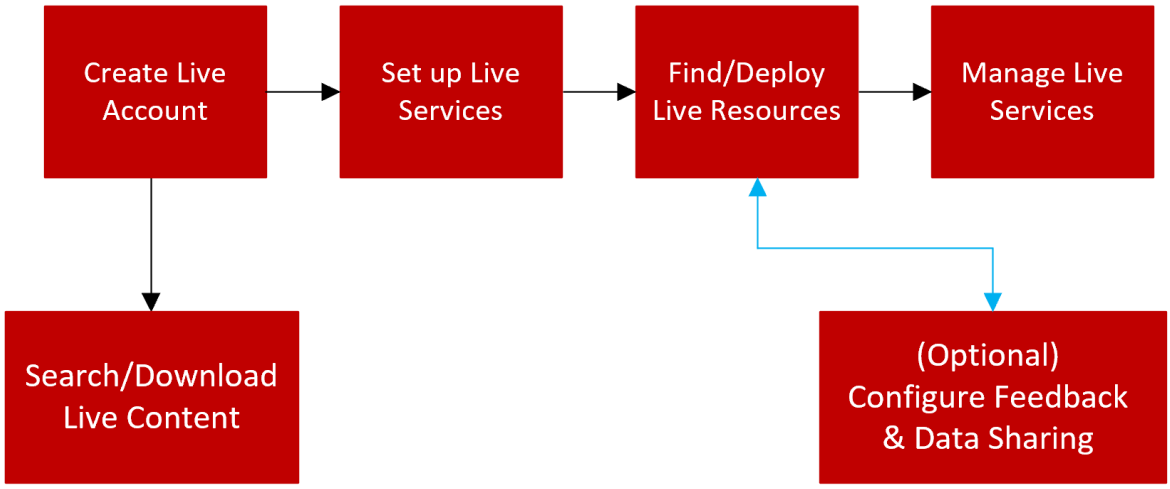
Deploy Content using Live Content UI

This topic explains the process of deploying the content using Live Content UI.

- Live Services Required Procedures
- Additional Procedures
- References
- Troubleshooting Live Services

Live Services Required Procedures

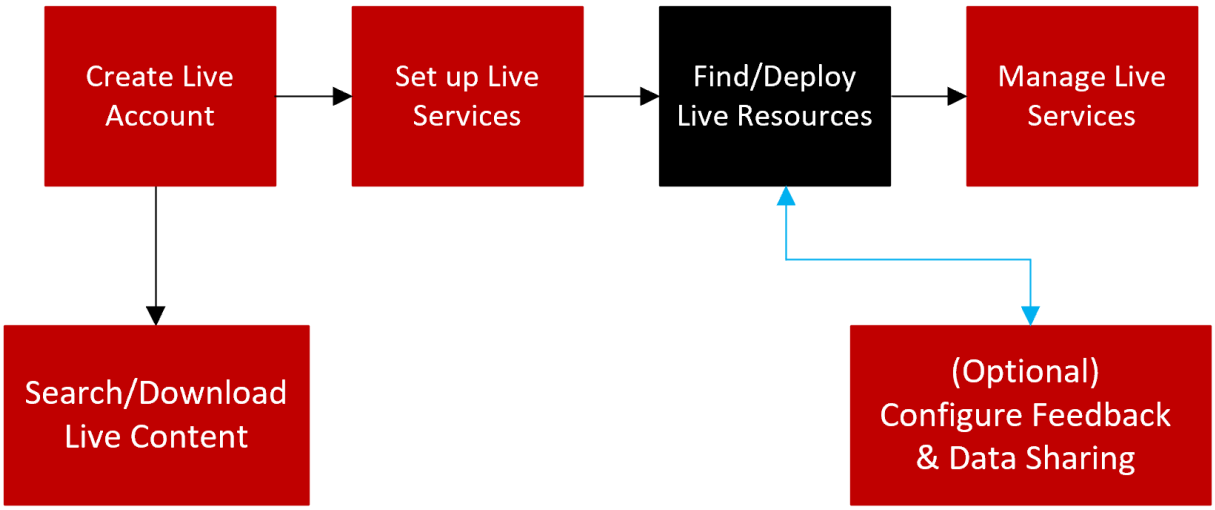
The following workflow describes the basic setup into four steps, which you perform individually.



Configuration Step	Description
Create Live Account	Create a Live Account on the Cloud Services Live Registration portal URL: https://live.netwitness.com/registration. If you have an existing account, you can manage your account using this portal.
Set Up Live Services on NetWitness	Set Up Live Services on NetWitness by configuring a connection with the CMS server.
Find and Deploy Live Resources	Search and browse for resources in the Live Search view, and then, deploy the selected resources.
Manage Live Resources	Procedures for administrators to search for, subscribe to, and deploy resources from Live.
Search and Download Content from NetWitness XDR Cloud Services Live	Search and browse for content in the Cloud Services Live, and then, download the selected content.
NetWitness Feedback and Data Sharing	Describes the feedback and data sharing features provided in NetWitness, from Live Services. Participation is optional, but can help to provide useful threat intelligence for the community.

Find and Deploy Live Resources

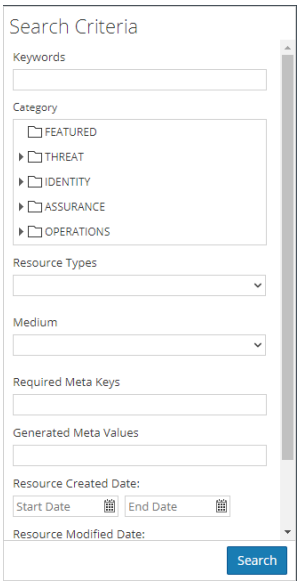
Administrators can search for resources in the Live Search view, which is also the same as browsing the Live CMS for resources using the Search Criteria panel of the Live Search View.



Find Resources in Live

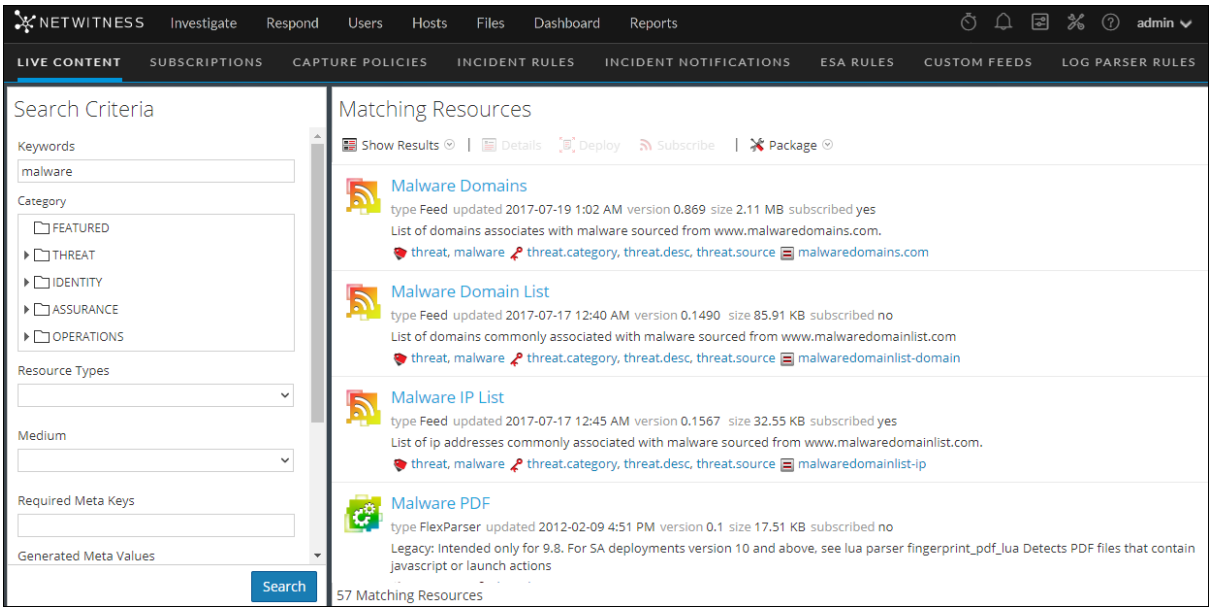
To find resources:

1. Navigate to **(Configure) > Live Content**.
2. In the **Search Criteria** panel, specify search criteria. Enter any or all of these: keyword, category, type of resource, medium, meta keys, meta values, date resource was created, and date resource was modified.



3. Click **Search**.

   The Matching Resources panel displays detailed results.



4. (Optional) To further narrow the results In the Matching Resources panel, click on a tag, meta key, medium or resource meta value in a result.

Deploy Resources in Live

In NetWitness, you can deploy selected resources manually, using the Deployment Wizard, or you can subscribe to a group of resources.

- When you have results from browsing resources in NetWitness Live, you can deploy resources manually to a service or a service group without subscribing to the resources. To deploy resources, select one or more from the list.
- Deploying resources manually deploys to services without taking advantage of the powerful resource management capabilities of NetWitness. If you want to receive notification and updates for updated resources and be able to easily remove resources from a service, you must subscribe to resources in the Live Search view and deploy them in the Live Configure View.
- If you have previously created and saved a resource package, you can deploy the package to services. Please refer to Resource Package Deployment Wizard for instructions on how to create a package.

Caution: For NetWitness 11.3, there is a new Content bundle for Endpoint, which contains approximately 400 application rules. Do not deploy this bundle (or the Endpoint application rules) onto any Log Decoder that is running an earlier version of NetWitness. The rules are only useful for 11.3 and newer, and would have major performance implications if deployed on Log Decoders that cannot process them.

To deploy resources manually:

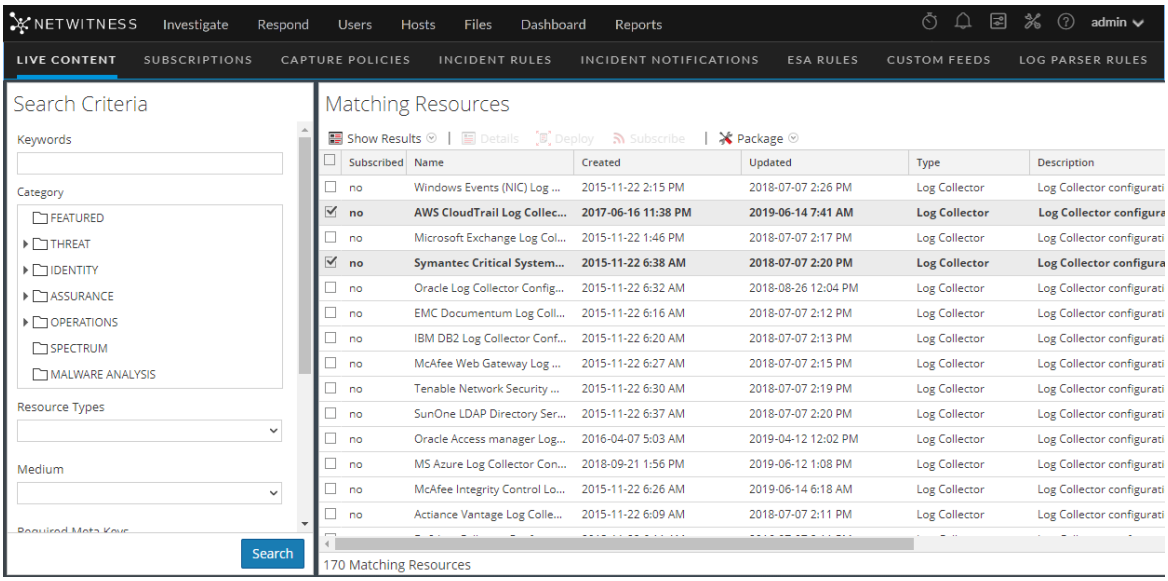
1. Go to  **(Configure) > Live Content**.
2. Select a group of resources, or a previously created resource package.

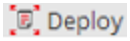
   To select a resource or group of resources:

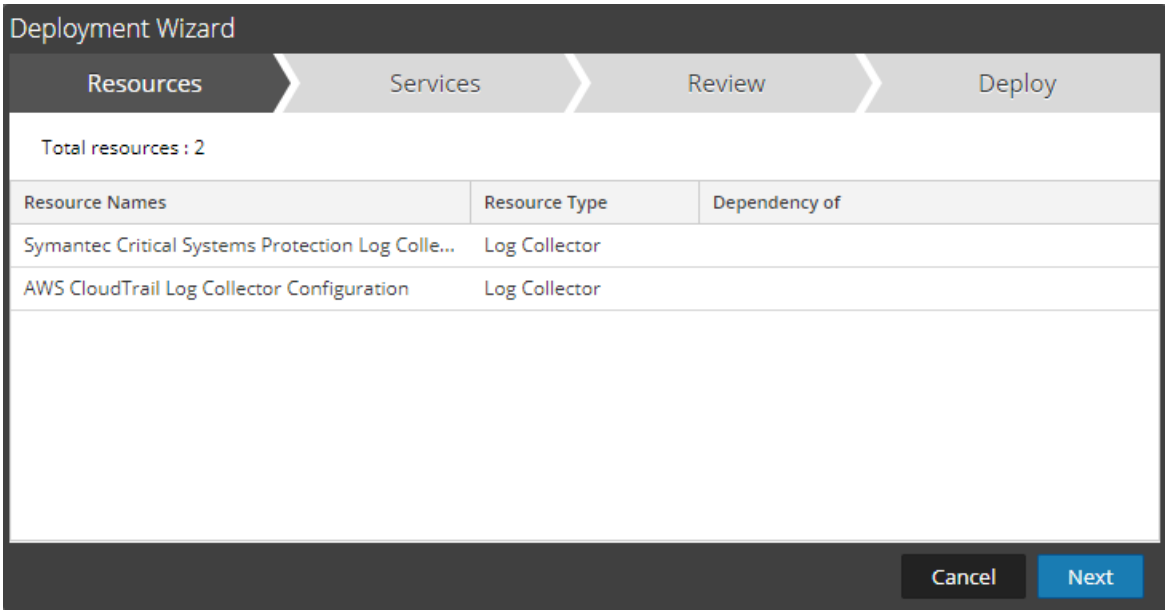
   a. In the **Live Search View,** browse Live resources (for example, search for the **Log Collector** resource Type).

   b. In the **Matching Resources** panel, select **Show Results > Grid**.

   c. Select the checkbox to the left of the resources that you want to deploy.

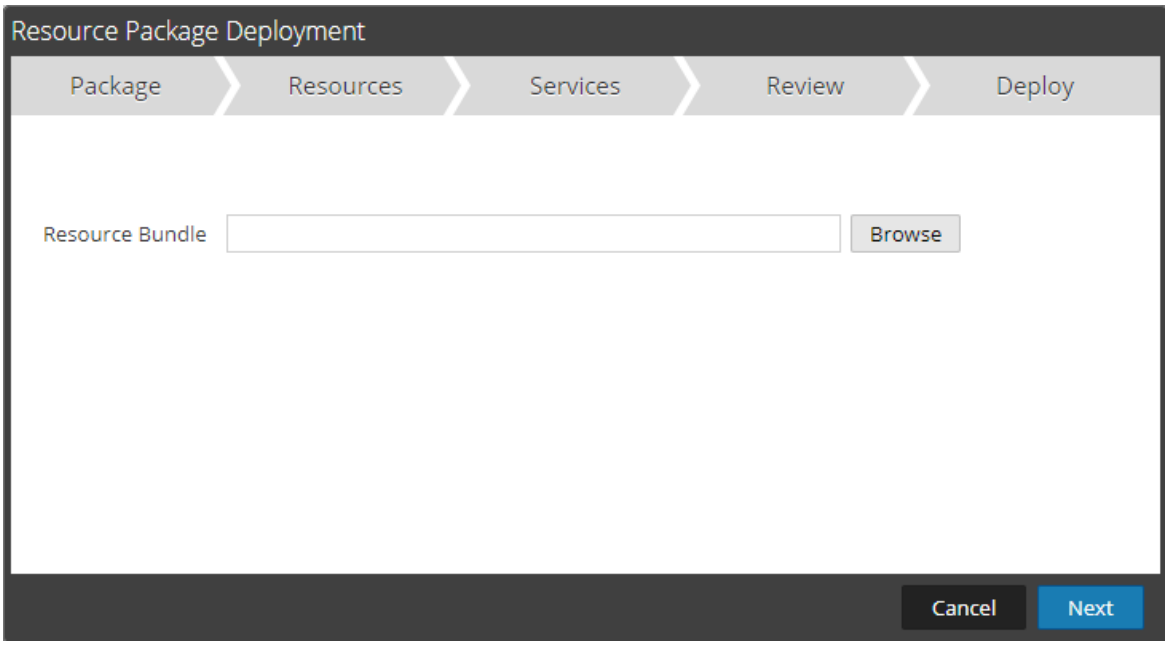


   d. In the Matching Resources toolbar, click .



3. To select a resource package to deploy:
   a. In the **Live Search** view **- Matching Resources** toolbar, select **Package** > **Deploy** .

      The Package page of the Resource Package Deployment wizard is displayed.

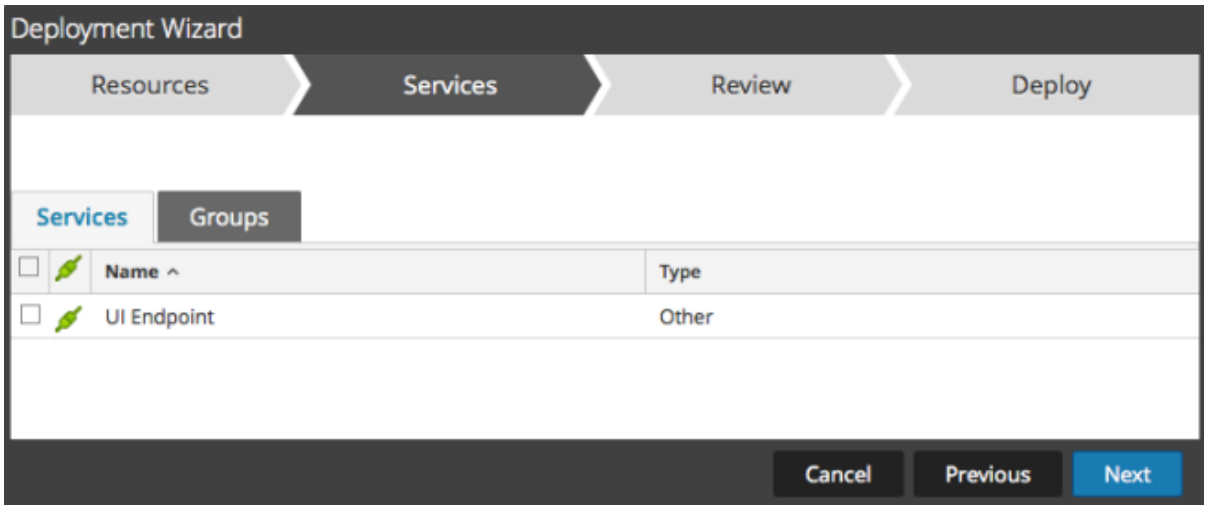


   b. Click Browse and select a package from your network (for example **resourceBundle-FeedsParsersContent.zip**).
   c. Click **Open**.

      At this point, whether you are deploying a package or a group of resources, the **Deployment Wizard** opens, and the **Resources** page is displayed.
4. Click **Next**.

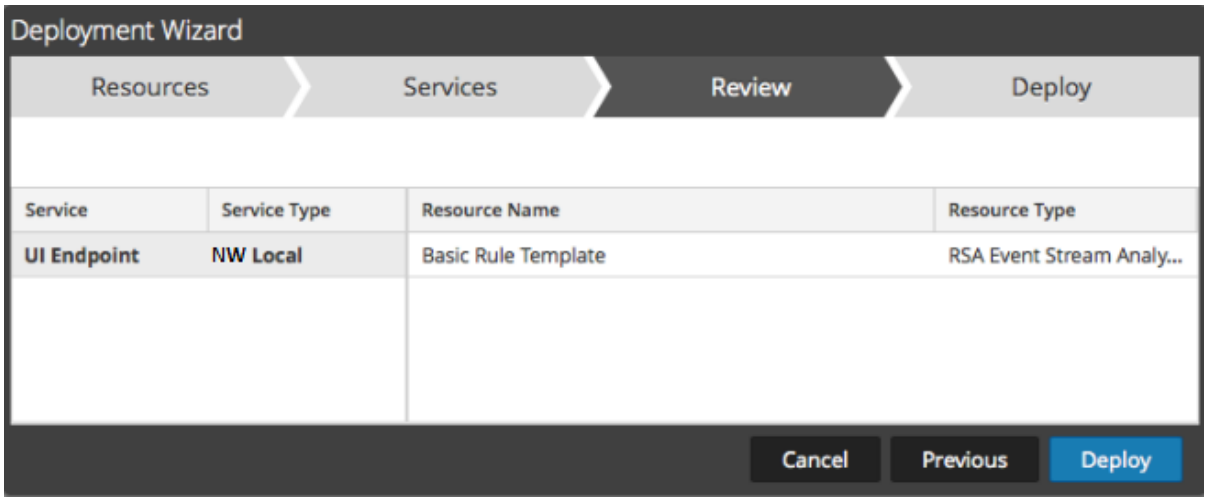
   The **Services** page displayed has two tabs, **Services** and **Groups**, which provide a list of services and service groups that are configured in the **(Admin) > Services** view. The columns are a subset of the columns available in the Services view.

   Note: The Live server is "smart" about deploying resources to Services. For example, it does not deploy resources that have a Medium of packets to any Log Decoders. This means that only applicable content resources are deployed to each Service.
5. Select the services on which you want to deploy the content. You can select any combination of services and service groups.
   - Use the **Services** tab to select individual services, list of services, and service groups that are configured in the **(Admin) > Services** view.
   - Use the **Groups** tab to select groups of services.



6. Click **Next**.

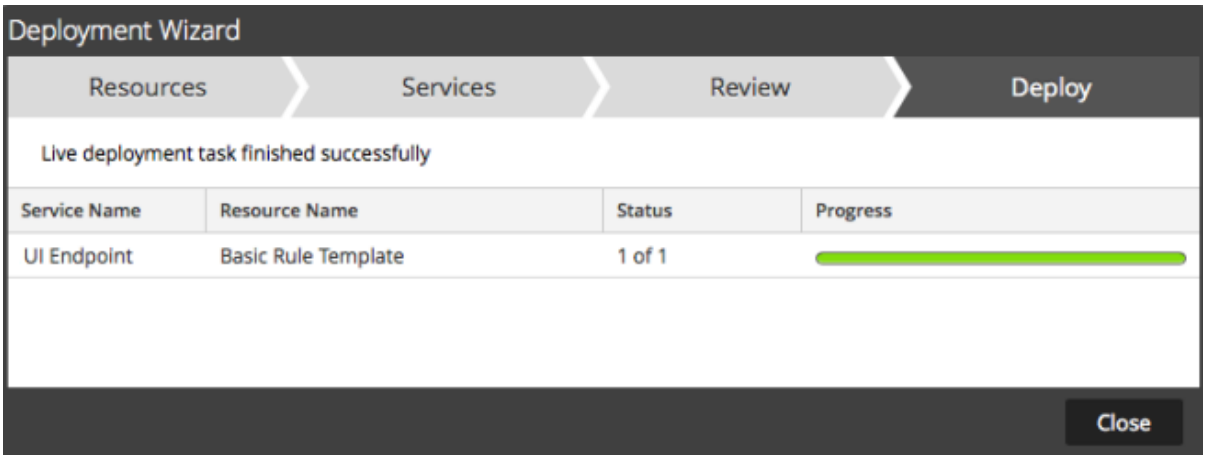
   The **Review** page is displayed.



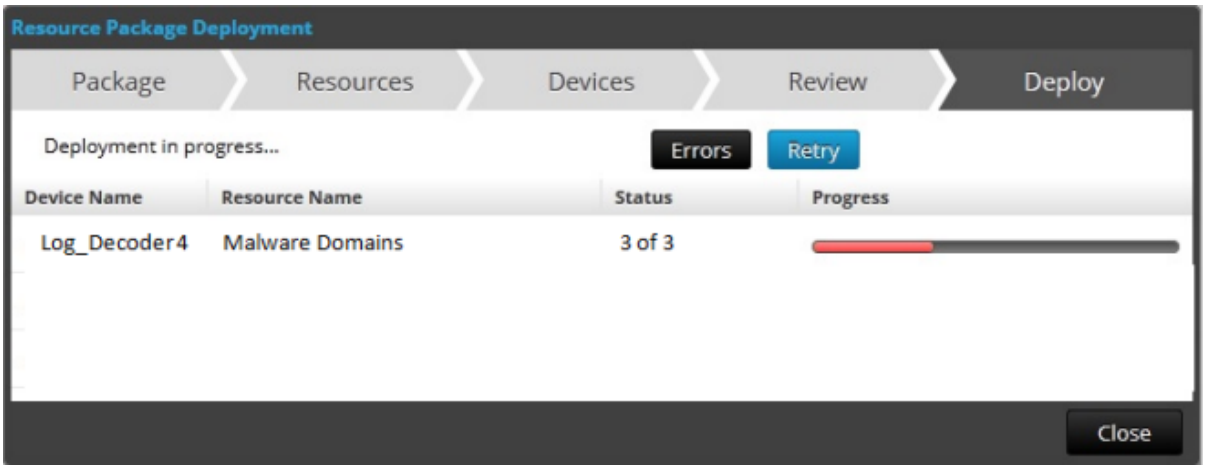
   Make sure that you have selected correct resources and the services on which you want to deploy them.

7. Click **Deploy**.

   The **Deploy** page is displayed. The Progress bar turns green when you have successfully deployed the resources to the selected services.



If you try to deploy resources and services that are not compatible, NetWitness displays the Errors and Retry buttons, which you can click to review the errors and re-attempt the deployment.



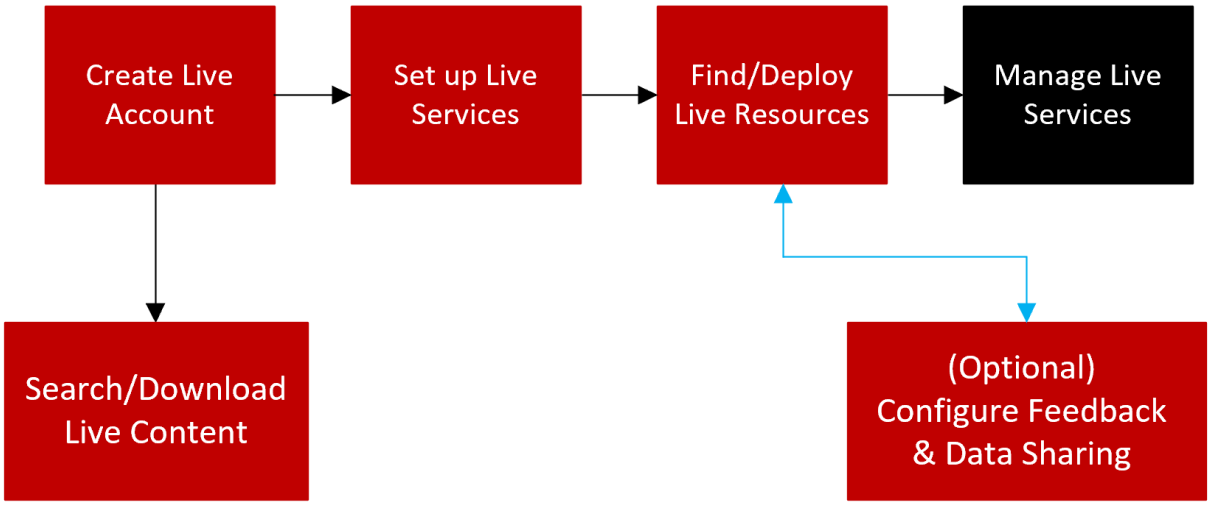
8. Click **Close**.

Next steps

After deploying parsers to Decoders and Log Decoders, you must enable parsers on the individual services. For more information, see the *Decoder and Log Decoder Configuration Guide*.

Manage Live Resources

With a connection to the CMS server, you can search for, subscribe to, and deploy resources from Live in accordance with your subscription level. Once you have found resources, you deploy them to services and service groups that have been configured in the the **(Admin) > Services** view.



There are several workflows for deploying resources to services and managing those deployments. These include:

- Subscribe and deploy resources
- Deploy a resource bundle
- Remove deployments of resources
- Download resources
- Set up data feeds

Manage Subscription and Deployment

The subscription and deployment workflow takes advantage of the resource management tools available in Live. By subscribing to resources, you agree to receive updated resources in accordance with the synchronization configured in the **(Admin) > System > Live Services panel**.

By adding subscribed resources to the deployments list, you configure NetWitness to automatically push those resources to the selected services at the configured synchronization intervals. This method requires some planning of service groups and services where resources are deployed. In addition:

- You can remove a resource from the deployments list in the Deployments Tab.
- You can unsubscribe from a resource in the Subscriptions Tab and the Live Resource View.

To manage subscriptions and deployment:

1. In the **(Admin) > System > Live Services** panel, specify an interval at which NetWitness checks for updates to subscribed resources in Live and specify the email addresses of people to receive an email listing subscribed resources that have been updated.
2. In the **(Configure) > Live Content** search view, search for and subscribe to Live resources.
3. In the **(Configure) > Subscriptions > Deployments** tab, select subscribed resources and add them to the deployment list for services groups.
4. (Optional) In the **(Configure) > Subscriptions > Deployments** tab, click to deploy the resources listed in the Deployments tab immediately.
5. In the **(Configure) > Subscriptions > Deployments** tab, select deployed resources from a Group, and remove them from services.
6. In the **(Configure) > Subscriptions** tab, unsubscribe from resources.

Remove a Deployed Resource

Once deployed to a service, Live resources remain on the service until removed. It is a good practice to remove unused resources from services on which they are deployed.

To remove deployed resources:

1. Go to the Live Resource View
2. Unsubscribe from a resource, and remove it from deployed services.

Deploy a Resource Bundle

To deploy a content package, use the Resource Package Deployment Wizard. You can deploy a content package created in Live to one or more services. NetWitness accepts packages in **.nwp** files or **.zip** files.

Download Resources

To download resources to your local file system, use the **Download** button in the Live Resource view.

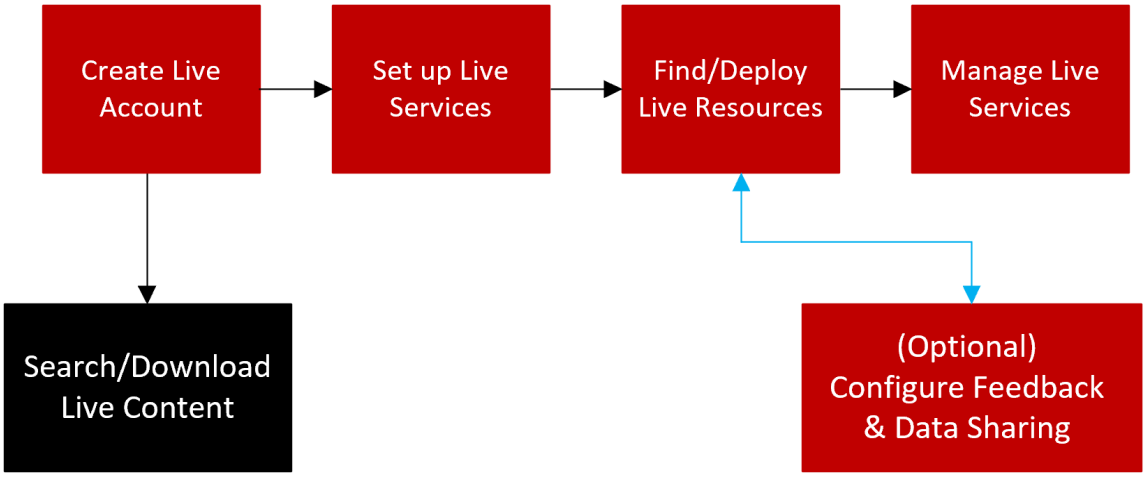
Set Up Data Feeds

In the **Live** > **Feeds** view, you can set up and maintain Custom and Identify feeds.

Search and Download Content from NetWitness XDR Cloud Services Live

Administrators can search for live content using the Search Content panel in the NetWitness XDR Cloud Services Live, which is similar to browsing the live CMS content in the (Configure) > Live Content page on the NetWitness Platform.

Note: If Admin server is not connected to the Live Services, you can use the NetWitness XDR Cloud Services Live to search and download the required content.



Prerequisites

- Ensure that you have created the Live account. For more information, see Create Live Account.

Quick Search for Content

You can now select and view the content based on the Sources available in the Cloud Services Live. You can select either NetWitness or Community from the Source drop-down list.

- **NetWitness**: Displays all the content provided by NetWitness.
- **Community**: Displays the content collected and retrieved from third party and open source communities.

  Note: The **Only Opensource** option will appear under the Search Content panel only when the community is selected as the source. You can use this option to select and search open-source related content.

You can also quickly select and view the available content types under Content section.

Clicking expands the **Content** section and displays the following options:

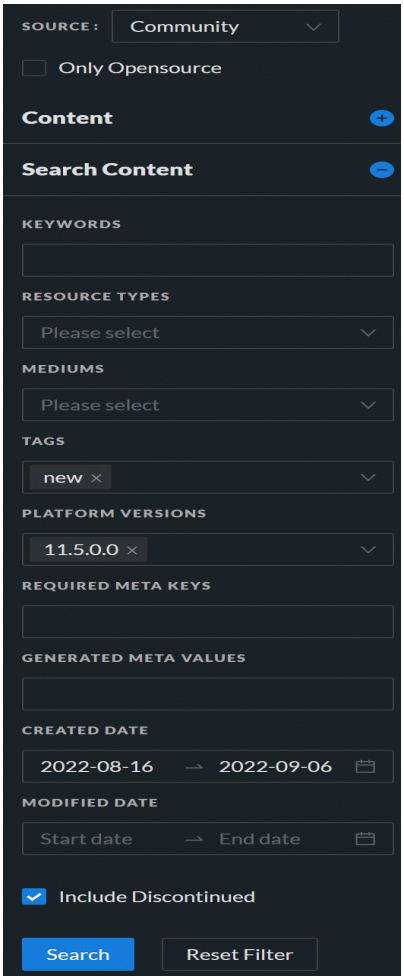
- **New**: Displays the content which is created in the last 21 days.
- **Recently Updated**: Displays the content which is created or updated in the last 21 days.

Advanced Search for Content

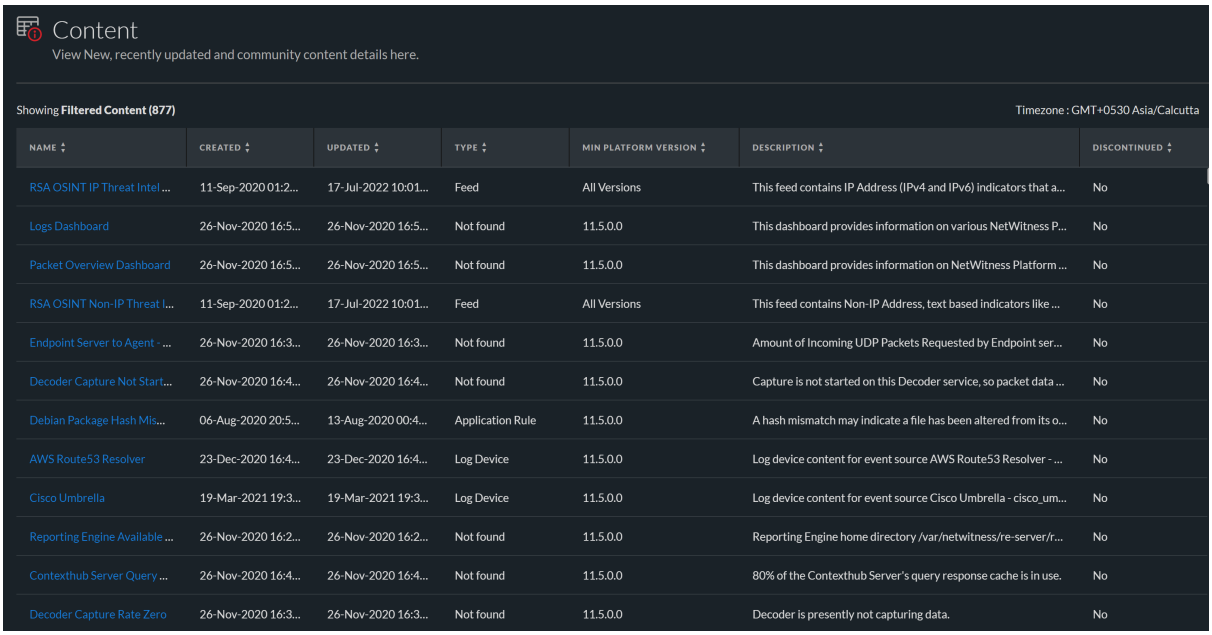
You can search for the specific content in the Search Content view. For more information, see Live Search Content View.

To search the content:

1. Click  to expand the **Search Content** section.
2. In the **Search Content**, specify the search criteria. Enter any or all of these: keyword, type of resource, medium, tag, platform versions, meta keys, meta values, date when content was created, date when content was modified, and (optional) discontinued content.



3. Click **Search**.
   The matching results are displayed on the right panel.



You can sort the content using the name, created, updated, type, or any of the column.

Note: Clicking **Reset Filter** removes the existing filters applied from the **Search Content**, and displays all the available content on the right panel.

Download Content

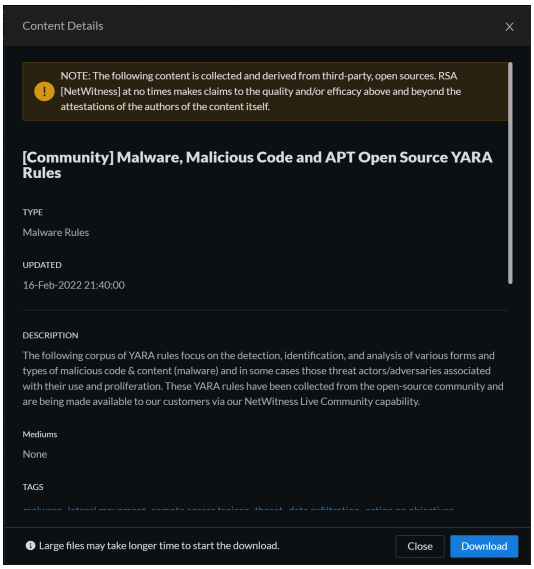
You can download the content from the results displayed in the right panel by performing the following steps:

Note:
- You cannot download a discontinued content.
- NetWitness provides no assurance related to the quality and accuracy of the content provided by the third parties and open source communities.

To download the content:

1. Click name of the content that you want to download.

   The **Content Details** dialog is displayed.



2. Click **Download**.

   The content file is downloaded.

Additional Procedures

This topic explains the additional procedures an administrator could choose to follow which are not essential for the configuration or use of Live Services.

- Export Data to RSA
- Packaging Resources
- Manage Custom Feeds
  - Creating a Custom Feed
  - Create a STIX Custom Feed
  - Creating and Managing an Identity Feed
  - Editing a Feed
  - Removing a Feed
- Miscellaneous Live Services Procedures

Export Data to RSA

A NetWitness administrator can export the metrics in NetWitness for Live Feedback.

About Live Feedback

In the Live Services Configuration panel, there is a Live Feedback Activity Log which enables you to download the usage data required for Live Feedback. This is active regardless of the Live Account configuration.

If the Live Account is not configured, you can manually upload the usage data to RSA. For more information, see the "Configure Live Services Panel" topic in the *System Configuration Guide*.

You must first download the Live Feedback historical data, and then upload it to share with RSA.

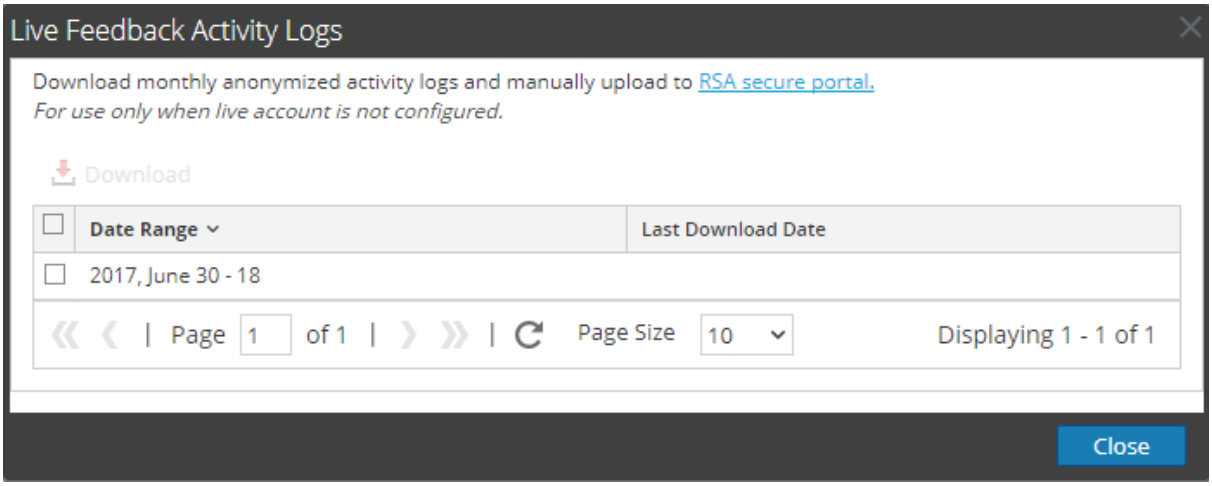
Download Live Feedback Historical Data

To download the Live Feedback historical data:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Live Services**.

   The **Live Account** screen is displayed which consists of the **RSA Live Status** and Download **Live Feedback Activity Log**.
3. Click **Live Feedback Activity Log**.

   The **Live Feedback Activity Log** window opens which allows you to download the required Live Feedback historical data.



4. Select one or multiple entries by selecting the checkboxes and click **Download**.

   Note: If you select multiple entries in the history table, the Live Feedback data is downloaded into a ZIP archive, consisting of individual JSON files for each month.

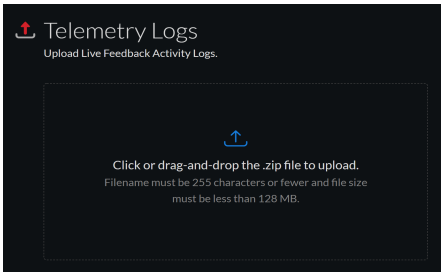
Share Telemetry Data to NetWitness

After you download the Live Feedback data, you can then upload it using the following procedure.

Note:
- To download the Live Feedback data, see topic Download Live Feedback Historical Data.
- You can share data NetWitness XDR Cloud Services Live portal. For more information, Create Live Account.

To share the data to RSA:

1. Log in to the NetWitness XDR Cloud Services using your Live account credentials.
2. Click  on the left panel.
   The **Telemetry Logs** dialog is displayed.



Note:
- You can upload only .zip files.
- Filename must be 255 characters or less and file size must be less than 128 MB.

3. Click or drag-and-drop a file onto this area to upload.

Packaging Resources

The primary use for creating and subsequently deploying a resource package is for customers using an air gap network environment. In this case, you create a resource package on the network that is connected to the internet, and then deploy the resource package on a more secure network.

Create and Deploy Resource Package Use Case

The basic steps are as follows:

1. Access NetWitness Live Services using an instance that is connected to the internet.
2. Create a Resource package as described below, adding whichever content items you need.
3. Copy the ZIP archive of the packages to your secure NetWitness instance, by using a thumb drive or other manual copying process.
4. On the secure NetWitness instance, deploy the resource package. For more information, see Resource Package Deployment Wizard.

Prerequisites to Create a Resource Package

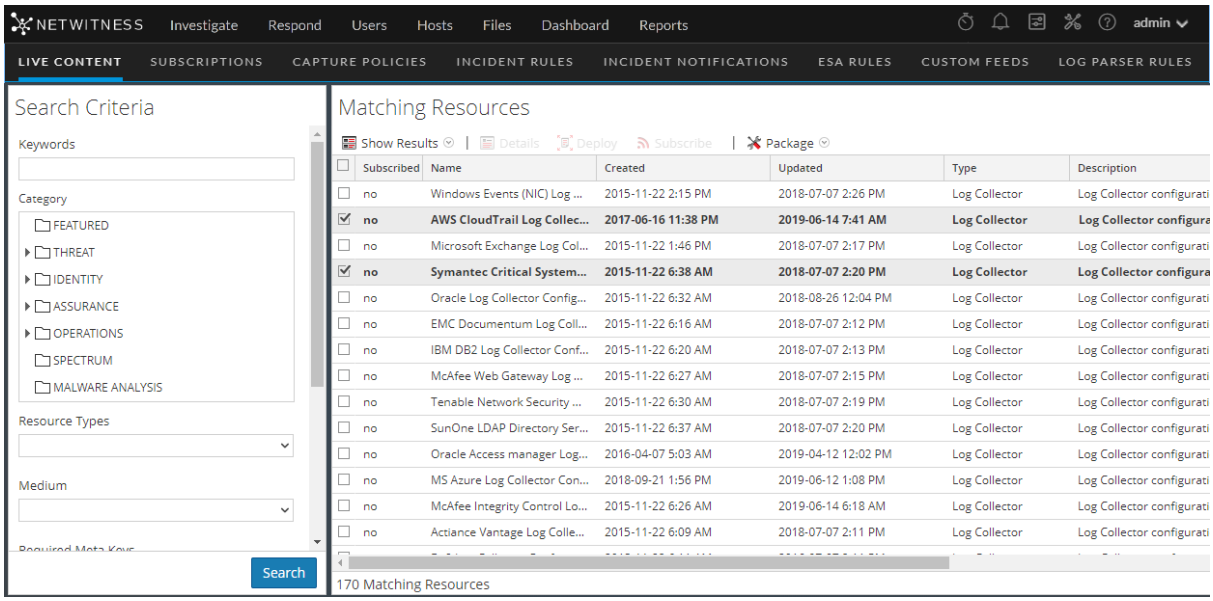
A prerequisite for creating resource packages is configuration of the connection and synchronization between the CMS server and NetWitness and the ability to search for resources in the User Interface.

Creating a Resource Package

The following procedure describes how to create a resource package, as a ZIP archive and save it to your local file system.

To create a resource package:

1. Go to  **(Configure) > Live Content** from the NetWitness UI.
2. Select the resources that you want to package in the Matching Resources grid.



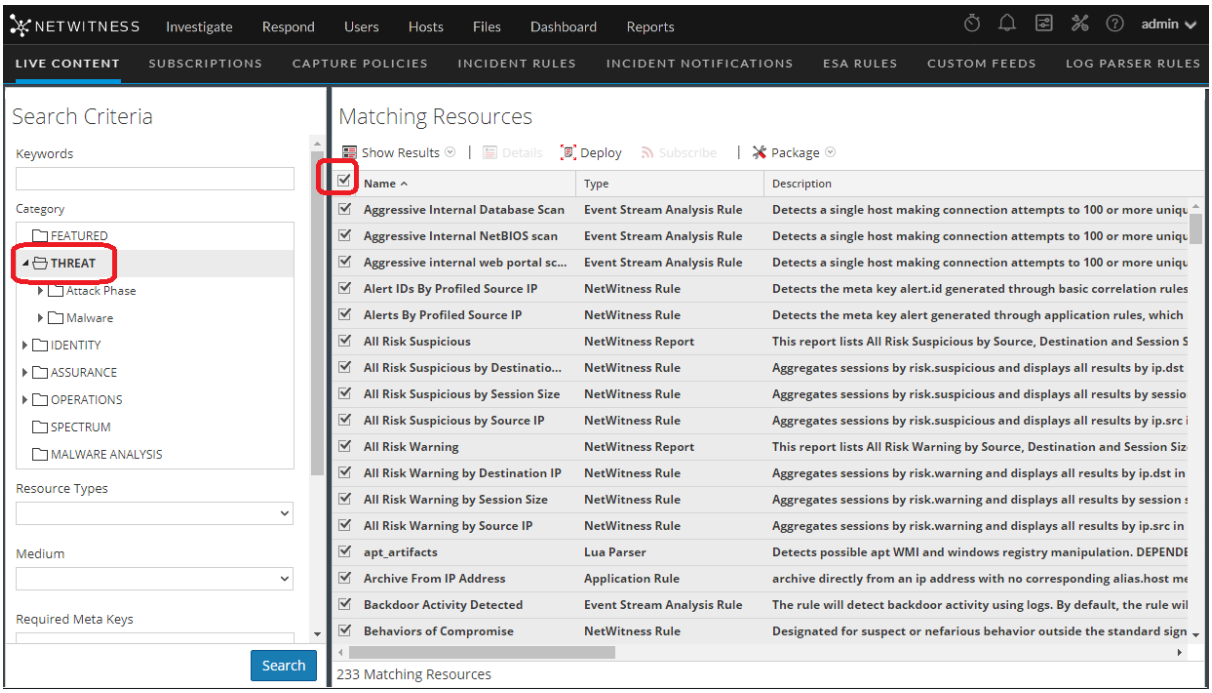
3. Select some or all the resources that are listed in the Matches Resources pane.
4. Select Package > **Create**.

NetWitness creates a **.zip** archive that contains the selected resources and downloads it to your default download folder. NetWitness gives the package a generic name. You should rename it when you save it so that it identifies the resources contained in the package.

Creating Threat Package

The following procedure describes how to create a resource package that contains all the content that is categorized as **Threat**. Then we rename it, using the type of content and date.

1. Go to **(Configure) > Live Content**.
2. From the **Category** section, select **Threat**.
3. Select all items returned by clicking on the checkbox in the column header row of the **Matching Resources** pane.



4. Select Package > **Create**.

   A ZIP archive is saved to your Downloads folder. For example, **resourceBundle8740753704980701969.zip**.

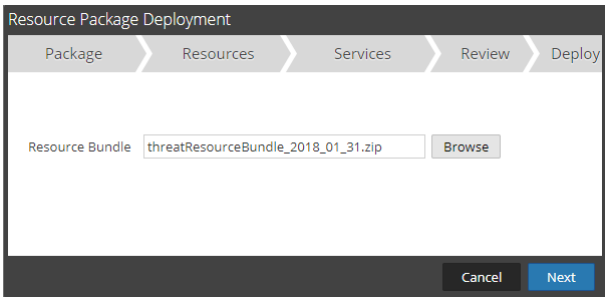
5. Rename the package to something meaningful. For example, in this case, you cold change the package name to **threatResourceBundle_2018_01_31.zip** (assuming today's date is January 31, 2018).

The resource package is now available for later deployment.

Deploying a Threat Package

This procedure assumes that you saved a package named **threatResourceBundle_2018_01_31.zip**, as described in the previous section. It describes how to deploy a saved resource package

1. Go to **(Configure) > Live Content**.
2. In the **Matching Resources** pane, select Package > **Deploy**.
3. Click **Browse** and navigate to the **threatResourceBundle_2018_01_31.zip** file that were created earlier.



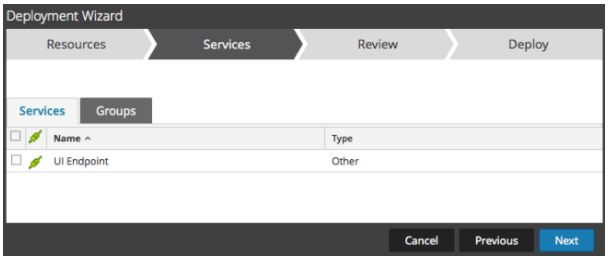
4. Click **Next**.

   The **Resources** page displays details for the resources in the package.

5. Click **Next**.

   The **Services** page displays two tabs, **Services** and **Groups**, which provide a list of services and service groups that are configured in the **(Admin) > Services** view. The columns are a subset of the columns available in the Services view.

6. Select the services on which you want to deploy the content. You can select any combination of services and service groups.



7. Click **Next**.

   The **Review** page is displayed.

   Note: Make sure that you have selected correct resources and the services to which you want to deploy them.

8. Click **Deploy** to complete the deployment process. Alternatively, you can choose **Cancel** or **Previous** to either cancel the deployment or go back to the previous screen.

Manage Custom Feeds

The custom feed capability is implemented using the Custom Feed Wizard in NetWitness, allowing you to quickly populate Decoders with custom and identity feeds.

Custom Feed Creation

You can use the **Live > Custom Feeds > Setup Feed > Configure a Custom Feed** wizard to create and deploy Decoder feeds based on deterministic logic that offers the meta keys specific to the selected Decoders and Log Decoders. Although the wizard guides you through the process to create both on-demand and recurring feeds, you should understand the form and content of a feed file when you create a feed.

Feed file names in NetWitness are in the form `<filename>.feed`. To create a feed, NetWitness requires a feed **data** file in `.csv` or `.xml` (for STIX) format and a feed **definition** file in `.xml` format, which describes the structure of a feed data file. The Configure a Custom Feed wizard can create the feed definition file based on a feed data file, or based on a feed data file and the corresponding feed definition file.

The files that you use to create an on-demand feed must be stored on your local file system. The files used to create a recurring feed must be stored at an accessible URL, whence NetWitness can fetch the most current version of the file for each recurrence. After a NetWitness feed is created, you can download the feed to your local file system, edit the feed files, and edit the NetWitness feed to use the updated feed files.

Sample Feed Definition File

This is an example of a feed definition file named `dynamic_dns.xml`, which NetWitness creates based on your entries in the Feed wizards. It defines the structure of the feed data file named `dynamic_dns.csv`.

Note: The feed file path should be `.csv` regardless of the Feed Type (Default or STIX).

Sample Feed Definition File

```
<?xml version="1.0" encoding="utf-8"?>
<FDF xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="feed-definitions.xsd">
    <FlatFileFeed name="Dynamic DNS Domain Feed"
        path="dynamic_dns.csv"
        separator=","
        comment="#"
        version="1">

        <MetaCallback
            name="alias.host"
            valuetype="Text"
            apptype="0"
            truncdomain="true"/>

        <LanguageKeys>
            <LanguageKey name="threat.source" valuetype="Text" />
```

```
            <LanguageKey name="threat.category" valuetype="Text" />
            <LanguageKey name="threat.desc" valuetype="Text" />
        </LanguageKeys>

        <Fields>
            <Field index="1" type="index" key="alias.host" />
            <Field index="4" type="value" key="threat.desc" />
            <Field index="2" type="value" key="threat.source" />
            <Field index="3" type="value" key="threat.category" />
        </Fields>

    </FlatFileFeed>
</FDF>
```

Feed Definition Equivalents for Custom Feed Wizard Parameters

The NetWitness Feeds wizard provide options to define the structure of the data feed file. These correspond directly to attributes in the feed definition (.xml) file.

NetWitness Parameter	Feed Definition File Equivalent
	Define Feed tab
Feed Type	Select: **Default** - to define a feed based on a `.csv` formatted feed data file. **STIX** - to define a feed based on STIX formatted `.xml` file.
Feed Task Type	Select: **Adhoc** - to create an on-demand feed. **Recurring** - to create a feed that recurs automatically.
Name	Enter a custom feed name in the feed data file that corresponds to the `flatfeedfile name` attribute in the feed definition file; for example, `Dynamic DNS Test Feed`.
File/ Browse	Enter a name of the feed data file that corresponds to the `flatfeedfile path` attribute in the feed definition file; for example, `dynamic_dns.csv`.
(STIX, Recurring) **Trust All Certificate**	Select **Trust All Certificate**, if you do not want to validate the REST server certificate. This option is enabled by default (checked).
(STIX, Recurring) **Certificate/Browse**	For client authentication with the REST URL, in the **Certificate** field, click **Browse** and select the self signed certificate. The supported certificate formats are .cer, .crt with Base64 & DER encoded files.
	Define Feed tab - Advanced Options
XML Feed File	Enter a name of the feed definition file, for example, `dynamic_dns.xml`.
Separator	The separator character used to separate attributes in the feed data file. It corresponds to the **`flatfeedfile separator`** in the feed definition file; for example, a comma.

NetWitness Parameter	Feed Definition File Equivalent
Comment	The character used to identify a comment in the feed data file. It corresponds to the `flatfeedfile comment` attribute in the feed definition file; for example, #.
Remove STIX data older than	The number of days for which the STIX packages downloaded from TAXII server have to be stored. The STIX packages older than the specified number of days are deleted automatically. The default value is 180 days, which is also the maximum.
Select Services tab	Select the services to which you want to send the data feed.
(Define Columns tab, Define Index) **Type**	The type of lookup value in the index position of the feed data file. **IP** means that each row in the feed data file contains an IP address in the lookup value position. The IP value is in dotted-decimal format (for example, 10.5.187.42). **IP Range** means that each row in the feed data file contains a range of IP addresses in the lookup value position. The IP range is in CIDR format (for example, 192.168.2.0/24). **Non IP** means that the each row in the feed data file contains a metadata value other than IP address in the lookup value position. The Service Type and Truncate Domain, and Callback Keys fields become active for a Non IP index.
(Define Columns tab, Define Index) **CIDR**	Specifies that the IP value in the lookup position is in CIDR format. The **CIDR** attribute sets the IP address format in the field to Classless Inter-Domain Routing (CIDR) notation.
(Define Columns tab, Define Index) **Service Type**	For a Non IP index, the integer service type to filter meta lookups. It corresponds to **MetaCallback apptype** attribute in the feed definition file. A value of **0** indicates no filtering by service type.
(Define Columns tab, Define Index) **Truncate Domain**	For a Non IP index, for meta values that contain domain names (for example, hostnames), the system can strip off the host specific element in the data. Truncate Domain corresponds to the **MetaCallback truncdomain** attribute. If the value is www.example.com, it is truncated to example.com. A value of **False** selects no truncation, and **True** selects truncation.
(Define Columns tab, Define Index) **Ignore Case**	If this option checked, the feed will ignore the case.
(Define Columns tab, Define Index) **Callback Keys**	For a Non IP index, the available meta keys to match on instead of ip.src/ip.dst (the defaults for IP index type) are selectable from the drop-down list. The Callback Key corresponds to the **MetaCallback name** attribute, and the index column of the csv file must contain data that can match the chosen meta key. For example, if the username meta key is chosen, the index column of the csv file needs to be populated with users to be matched.

NetWitness Parameter	Feed Definition File Equivalent
(Define Columns tab, Define Index) **Index Column**	Identifies the column in the feed data file that provides the lookup value for the row. Each position in each row of the feed data file is identified by a **Field index** attribute in the feed definition file. A field with an index of **1** is the first entry in a row, the second field has an index of **2**, the third field has an index of **3**, and so on. You can select multiple index columns, if the **Feed Type** is **STIX** and **Index Type** is **Non IP**. When you select multiple index columns the values from all the selected columns are merged in the first index column that you selected.
(DEFINE VALUES) **Key**	The name of the **LanguageKey**, as defined in the feed definition file, for which meta is created from this row of the feed data file. It corresponds to the **Field key** attribute in the feed definition file. A key applies only to a field whose type is set to **value**. In the feed definition file, there is a list of LanguageKeys from **index.xml**, or a summary name if Source Name and Destination Name are used. For example, **reputation** is a summary name for **reputation.src** and **reputation.dst**). This value is referenced by the Field key attribute.

Creating a Custom Feed

This topic provides instructions for creating a custom feed using a .csv or STIX formatted feed data file in NetWitness. For more information about STIX and creating a STIX custom feed, see Create a STIX Custom Feed.

You can easily create a custom feed using the Custom Feed wizard. To complete this procedure, you need a feed data file in `.csv`or `.xml` format. If you also have an associated feed definition file in `.xml` format, which describes the structure of the feed data file, you can use the feed definition file to create a feed. The Custom Feed wizard can create the feed based on a feed data file, or based on a feed data file and corresponding feed definition file.

After completing this procedure, you will have created a custom feed.

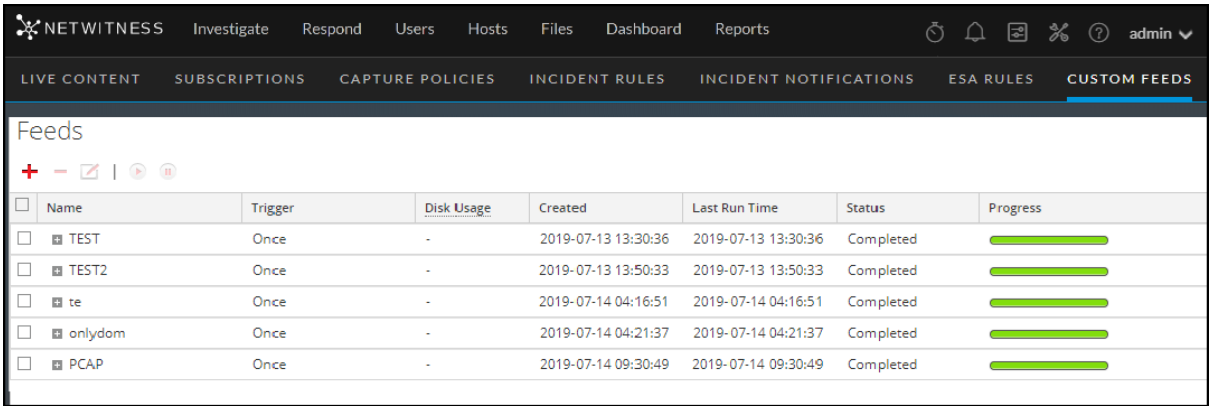
The feed data file (`.csv` or STIX (`.xml`)) and optionally the feed definition file (`.xml`) must be available on the local file system for an on-demand custom feed. For a recurring custom feed, the files must be available at a URL that is accessible to the NetWitness server.

Note: Any feeds that are created in 11.2 release or prior will be automatically pushed to Context Hub as Lists. The lists can be looked up in the context lookup panel of the Respond and Investigate pages. If Context Hub is not configured or the service is down, then the feeds will be pushed to Context Hub the next time the server is available.

To create a custom feed:

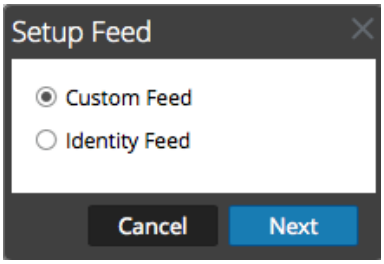
1. Go to **(Configure) > CUSTOM FEEDS**.

   The Custom Feeds view is displayed.



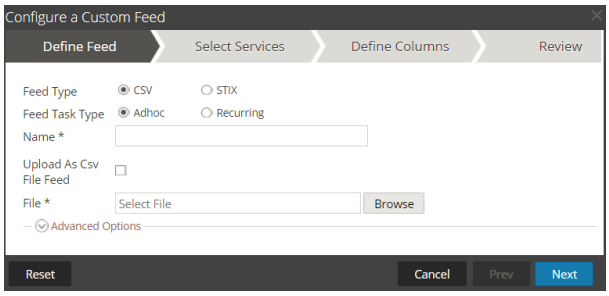
2. In the toolbar, click .

   The Setup Feed dialog is displayed.



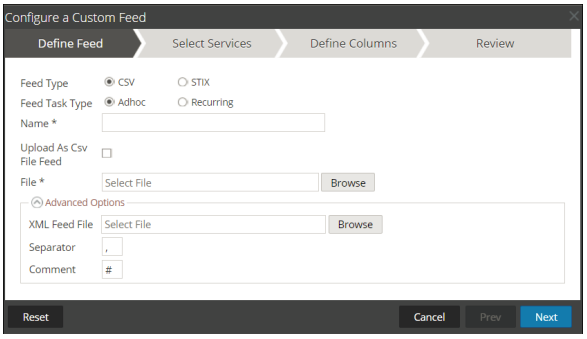
3. To select the feed type, click **Custom Feed** and **Next**.

   The Configure a Custom Feed wizard is displayed, with the Define Feed form open.



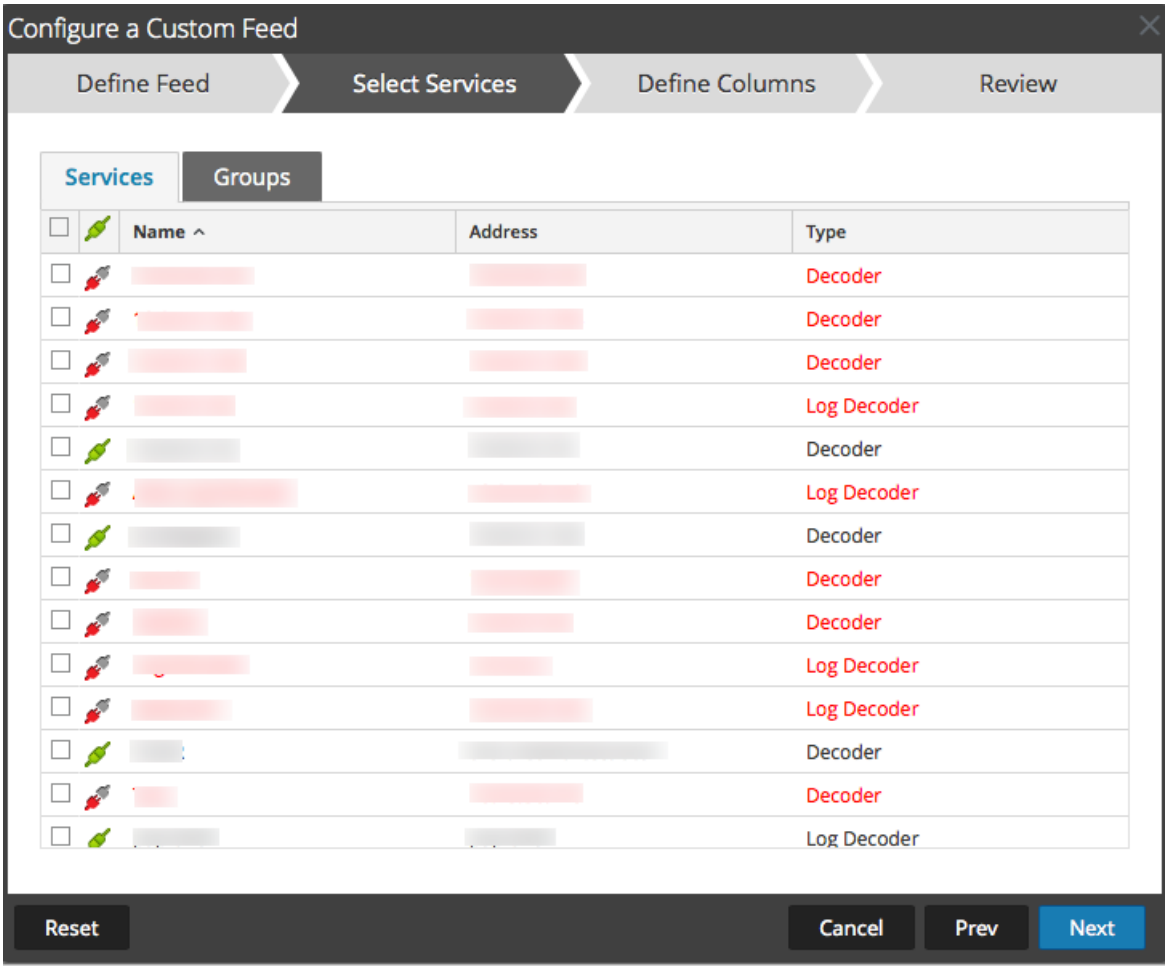
4. To define a feed based on a `.csv` formatted feed data file, select **CSV** in the **Feed Type** field.
5. To define an on-demand feed task that executes once, select **Adhoc** in the **Feed Task Type** field and do one of the following:
   a. (Conditional) To define a feed based on a .csv formatted feed data file, type the feed **Name**.
   b. Select the checkbox **Upload As CSV File Feed**, if required.
   c. Select a .csv content **File** from the local file system, and click **Next**.
   d. (Conditional) To define a feed based on an XML feed file, select **Advanced Options**.

      The Advanced Options are displayed:



   e. Select an XML feed file from the local file system, choose the **Separator** (default is comma), and specify the **Comment** characters used in the feed data file (default is #), and click **Next**.
   f. The Select Services form is displayed. This is an example of the form for a feed based on a feed data file with no feed definition file. If you are defining a feed based on a feed definition file, the

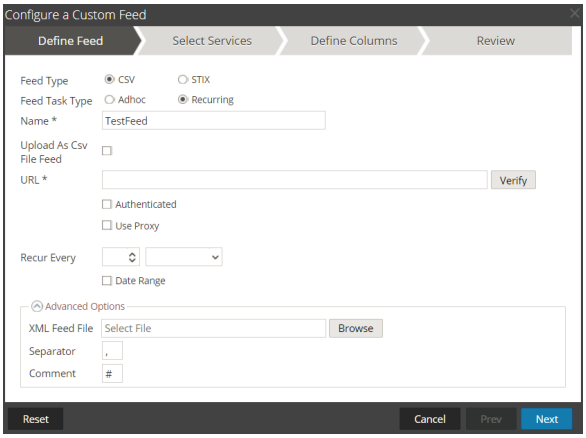
Define Columns tab is not needed.



6. To define a recurring feed task that executes repeatedly at specified intervals, during a specified date range.

   a. Select **Recurring** in the **Feed Task Type** field.

   The Define Feed dialog includes the fields for a recurring feed.



b. In the **URL** field, enter the URL where the feed data file is located, for example, `http://<hostname>/<feeddatafile>.csv`, and click **Verify**.

   NetWitness verifies the location where the file is stored, so that NetWitness can check for the latest file automatically before each recurrence.

c. (Optional) If the URL has restricted access and requires authentication using your username and password, select **Authenticated**.

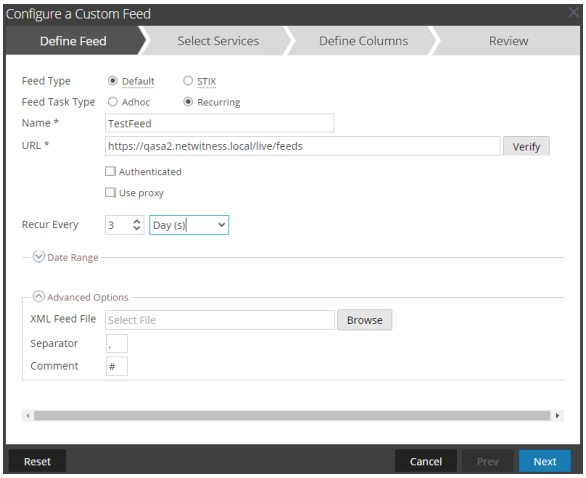
   NetWitness provides your user name and password for authentication to the URL.

d. If you want the NetWitness server to access the Feed URL through a proxy, select **Use Proxy**. For more information on configuring a proxy, see the **Configure Proxy for NetWitness** topic in the *System Configuration Guide*. By default, the **Use Proxy** checkbox is not selected.
   NOTE:
   If you are using an HTTPS based feed server, ensure that you import and install the certificates. For more information, see Import Certificates for HTTPS Service

e. To define the interval for recurrence, do one of the following:

   - Specify the number of minutes, hours, or days between recurrences of the feed.
   - Specify recurrence every week, and select the days of the week.

f. To define the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.



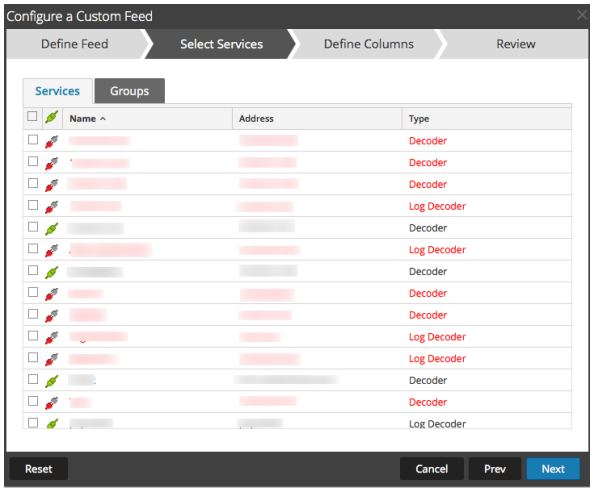
7. (Conditional) If you want to define a feed based on an XML feed file:

   - Type the feed **Name**, select **Advanced Options**.

     The Advanced Options fields are displayed.

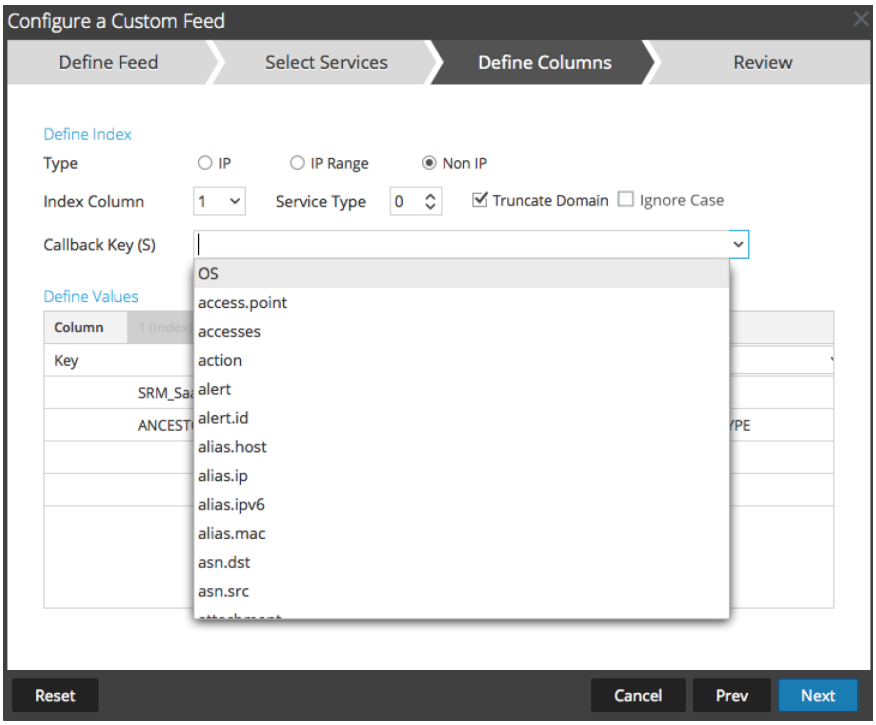
   - Select an XML feed file from the local file system, choose the **Separator** (default is comma), specify the **Comment** characters used in the feed data file (default is #) and click **Next**.

The Select Services dialog is displayed.

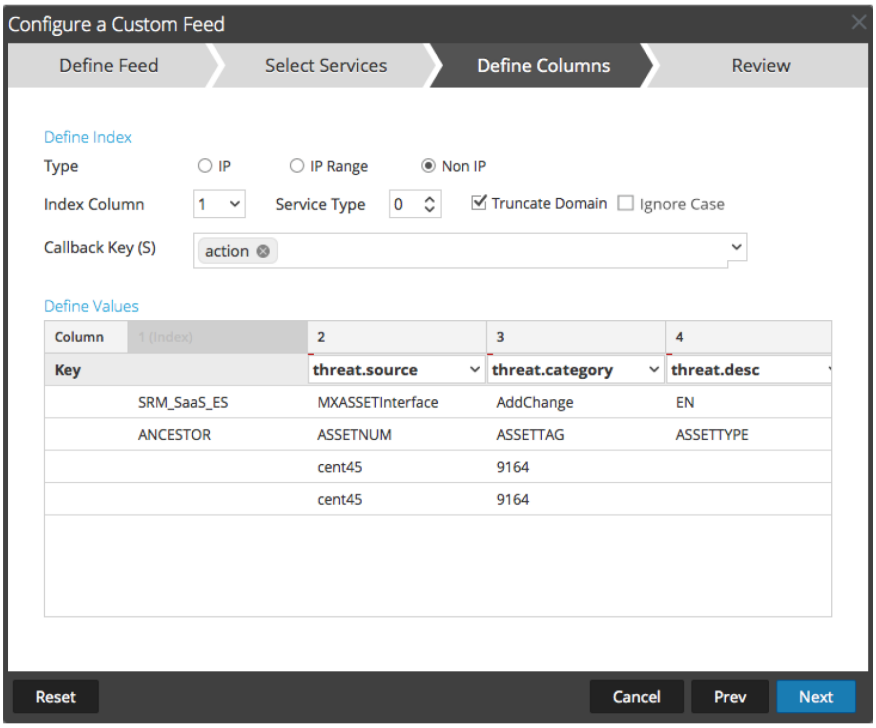


8. To identify services on which to deploy the feed, do one of the following:
   a. Select one or more Decoders and Log Decoders, and click **Next**.
   b. Click the **Groups** tab and select a group. Click **Next**.

      The Define Columns dialog is displayed.
9. To map columns in the Define Columns form:
   a. Define the Index type: **IP**, **IP Range**, or **Non IP**, and select the index column.
   b. (Conditional) If the index type is **IP** or **IP Range** and the IP address is in CIDR notation, select **CIDR.**
   c. (Conditional) If the index type is **Non IP**, additional settings are displayed. Select the service type and **Callback Keys,** and optionally select the **Truncate Domain** and **Ignore Case** option.



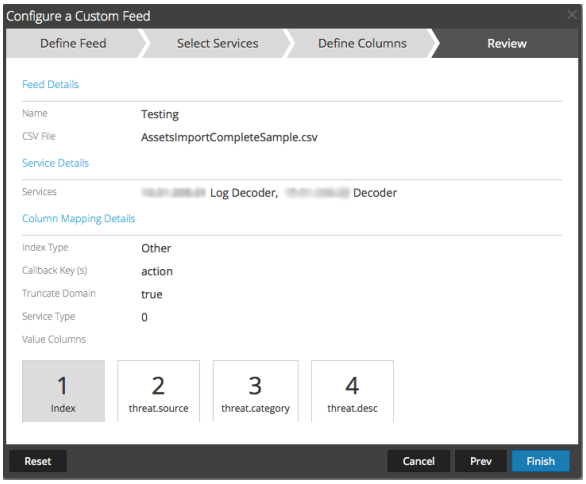
d. Select the language key to apply to the data in each column from the drop-down list. The meta displayed in the drop-down list is based on the meta available for the service define values. You can also add other meta based on advanced expertise.



Note: When a custom feed gets converted into a context hub list, you must map at least one meta key with one or more meta types by mapping a column header with a meta. However, you can add or edit the entity mapping of a list by clicking ⚙ in the Lists tab. For more information, see the *Context Hub Configuration Guide*.

e. Click **Next**.

The Review dialog is displayed.



10. Anytime before you click **Finish**, you can:
    - Click **Cancel** to close the wizard without saving your feed definition.
    - Click **Reset** to clear the data in the wizard.
    - Click **Next** to display the next form (if not viewing the last form).
    - Click **Prev** to display the previous form (if not viewing the first form)
11. Review the feed information, and if correct, click **Finish**.
12. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file are listed in the Feed grid and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.

Note: When you create a feed, and if there is no entity mapping done such as in case of custom meta, then those columns in the List will not have entity mappings in Context Hub. You have to manually map the entities from the List page.

Import Certificates for HTTPS Service

Import certificates to communicate with the HTTPS services:

1. SSH to the NW node and copy the CA certificate located in the following directory:
   /etc/pki/ca-trust/source/
2. Execute the following command to update the certificates:
   `update-ca-trust`

3. Execute the following command to add the certificate to the java keystore:
   ```
   keytool -list -keystore /etc/pki/java/cacerts -storepass changeit |& head
   ```
4. Restart the service on the NW node.

Note: Perform the procedure for all the HTTPS servers.
Example: HTTPS proxy server and HTTPS feed server.

Create a STIX Custom Feed

You can create a custom feed using a .csv or STIX formatted feed data file in NetWitness.

Note: NetWitness supports Structured Threat Information Expression (STIX) 1.0, 1.1 and 1.2 versions only.

Structured Threat Information Expression (STIX™) is a structured language for describing cyber threat information so it can be shared, stored, and analyzed in a consistent manner. For more information about STIX, see https://oasis-open.github.io/cti-documentation/.

Caution: If STIX recurring feed is configured and you update Security Analytics from 10.6.x to NetWitness 11.0, you must re-configure the STIX recurring feed.

In NetWitness Platform, STIX feeds are supported. STIX content (with version 11.x) can be uploaded in an ".xml" format. The constructs such as Indicator Title and Description, Observable Title and Description, and Indicator Sightings information are parsed from STIX and pushed to the decoders or log decoders that are selected during feed configuration. Information such as IP addresses, File hashes, Domain names, URIs, and Email addresses are extracted from the STIX observable to be included in the feed.

Make sure the following criteria are met before you upload the STIX file:

1. Only STIX Observables with property values in the "Equals" operator
2. The uploaded STIX xml file must have only one STIX_Package

TAXII (Trusted Automated eXchange of Indicator Information) is the main transport mechanism for cyber threat information represented in STIX. Using the TAXII services, organizations can share cyber threat information in a secure and automated manner.

The STIX and TAXII communities work closely together to ensure that they continue to provide a full stack for sharing threat intelligence.

Apart from TAXII server, STIX data can also reside on REST server and you can fetch STIX file from the REST server by providing the URL of the REST server. For example, http://stixrestserver.internal.com.

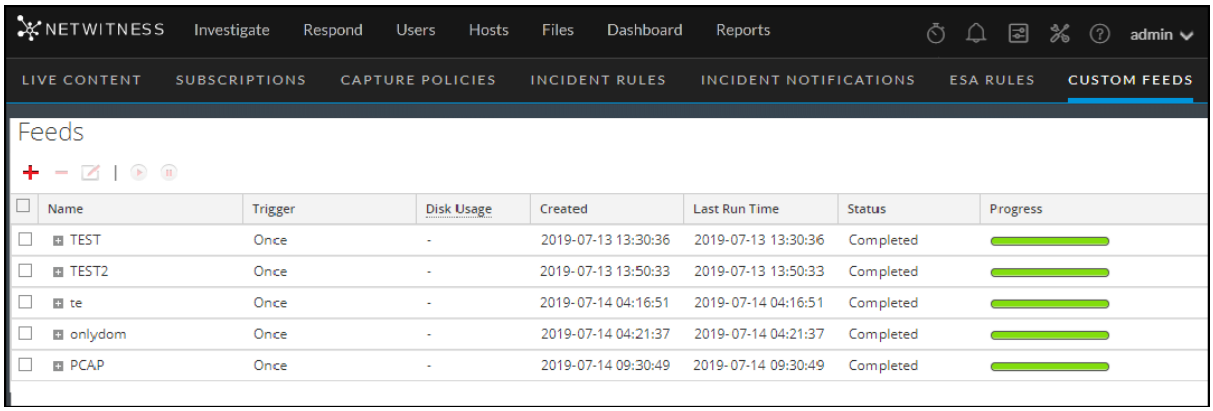
The feed data file (`.csv` or STIX (`.xml`)) and optionally the feed definition file (`.xml`) must be available on the local file system for an on-demand custom feed. For a recurring custom feed, the files must be available at a URL that is accessible to the NetWitness server.

In NetWitness Platform, STIX (.xml) feed of type Indicators or Observable which contains the properties such as the IP addresses, File hashes, Domain names, URIs and Email addresses are supported. The properties values in the Equals operator is only supported. The STIX constructs that are parsed are Indicator Title and Description, Observable Title and description and Indicator Sightings information. The STIX (.xml) with a single STIX_Package is only supported."

To create a STIX custom feed:

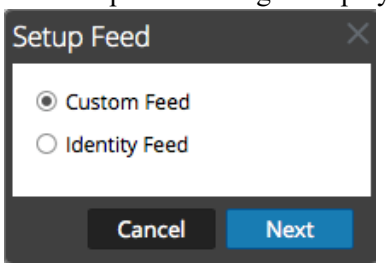
1. Go to  **(Configure) > CUSTOM FEEDS**.

   The Custom Feeds view is displayed.



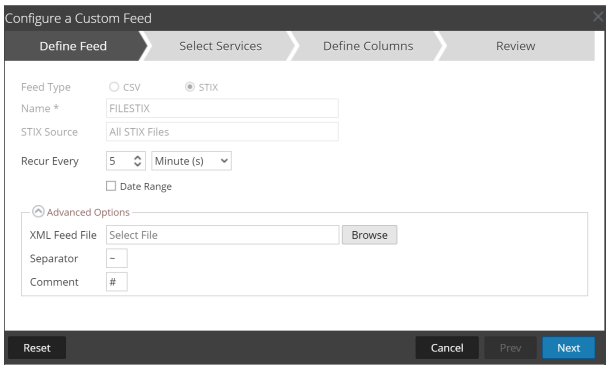
2. In the toolbar, click .

   The Setup Feed dialog is displayed.



3. To select the feed type, click **Custom Feed** and **Next**.

   The Configure a Custom Feed wizard is displayed, with the Define Feed dialog open.



4. Enter the following details:

   a. **Feed Type**: Select **STIX**, to define a feed based on a STIX formatted `.xml` file.

   b. **Name**: type the feed name, to define a feed based on STIX formatted `.xml` file.

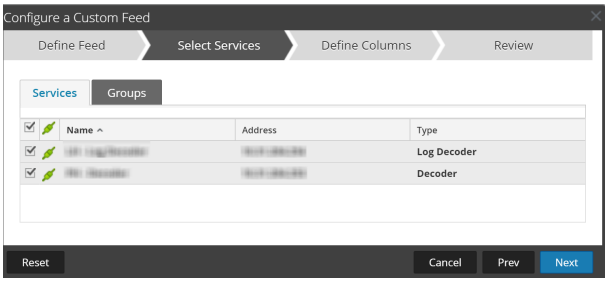
   c. **STIX Source**:Select a STIX data source from the drop-down which is added in Context Hub.

   d. **Recur Every**: Specify a recurring feed task that executes repeatedly at specified intervals.

   Note: NetWitness verifies the connection to the server, so that NetWitness can check for the latest file automatically before each recurrence.

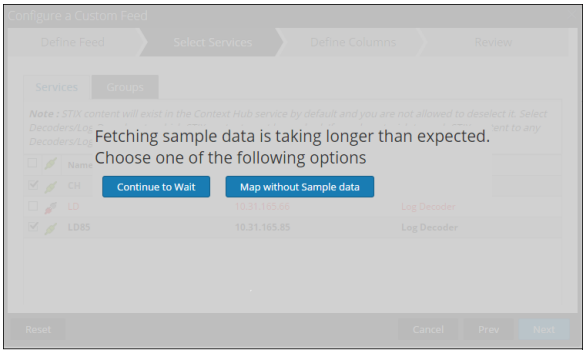
   e. **Date Range**: Select the checkbox and specify the date range for the feed task to recur.

5. (Optional) Select **Advanced Options**,to define a feed based on an XML feed file.
   a. **XML Feed file**: Browse and select an XML feed file from the local file system.
   b. **Separator**: Choose a separator (default is comma).
   c. **Comment**: Specify the comment characters used in the feed data file (default is #).
6. Click **Next**.
7. The Select Services dialog is displayed. This is an example of the form for a feed based on a feed data file with no feed definition file. If you are defining a feed based on a feed definition file, the Define Columns tab is not needed.



8. To identify services on which to deploy the feed, do one of the following:
   a. Select one or more Decoders and Log Decoders, and click **Next**.
   b. In case of STIX feed, Context Hub will be selected by default and you are not allowed to deselect it. In addition, you can select one or more Decoders and Log Decoders and click **Next** or Click the **Groups** tab and select a group. Click **Next**.

      If the data from the STIX server is large, the following message is displayed:



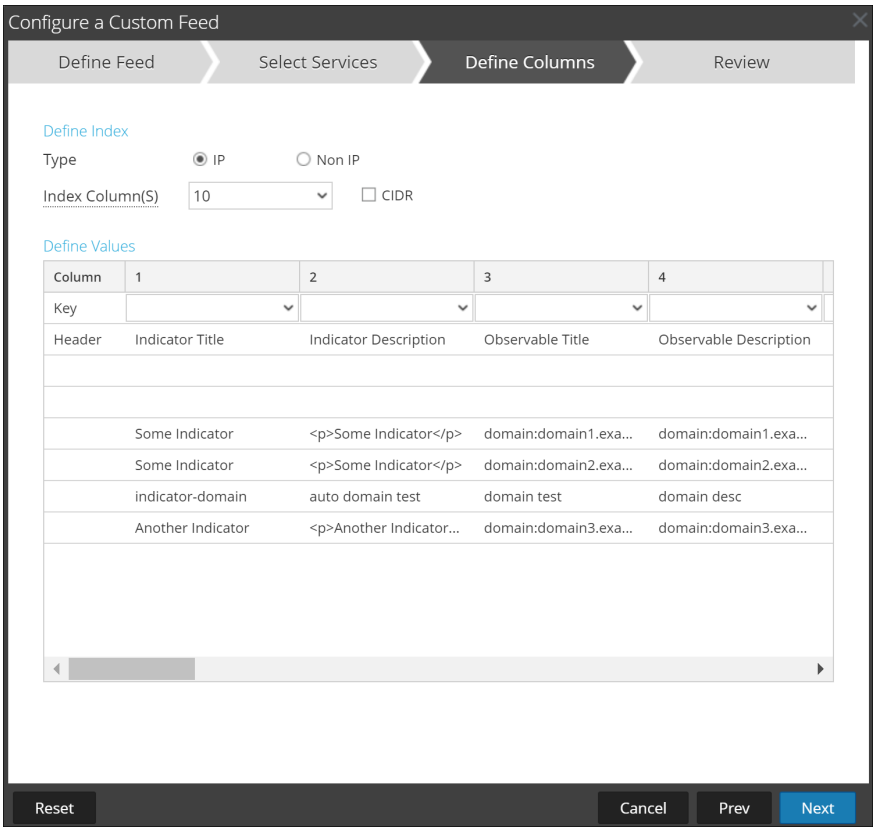
      - If you click **Continue to Wait**, it continues to wait till the sample data is fetched or timeout (10 minutes) whichever is sooner. In case of timeout no sample data is retrieved even after 10 minutes.
      - If you click **Map Without Sample data**, the mapping column is displayed without any sample data.

      The Define Columns dialog is displayed.
9. To map columns in the Define Columns form:

a. Define the Index type: **IP**, **IP Range**, or **Non IP**, and select the index column.

b. (Optional) If the index type is **IP** or **IP Range** and the IP address is in CIDR notation, select **CIDR**.

c. (Optional) If the index type is **Non IP**, additional settings are displayed. Select the service type and **Callback Keys,** and optionally select the **Truncate Domain** option.

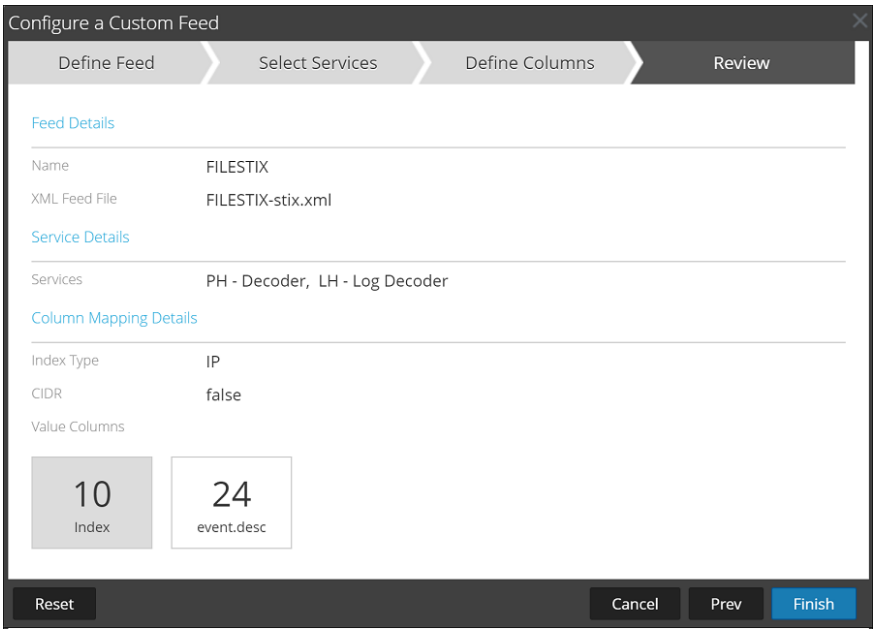


Note:
- If the **Index Type** is Non IP, you can select multiple index columns in the **Index Column(S)**. The values from all the selected columns are merged in the first index column that you selected and the merged values are pushed to the Log Decoder for parsing. For example, in the **Index Column(S)** if you select 2,4,7 as index columns the values from the 2,4 and 7 columns are merged in the column 2 and the values are pushed to Log Decoder for parsing.
- Indexing cannot be done for the columns such as Indicator Title, Indicator Description, Observable Title, Observable Description, as the look up cannot be performed for those columns.

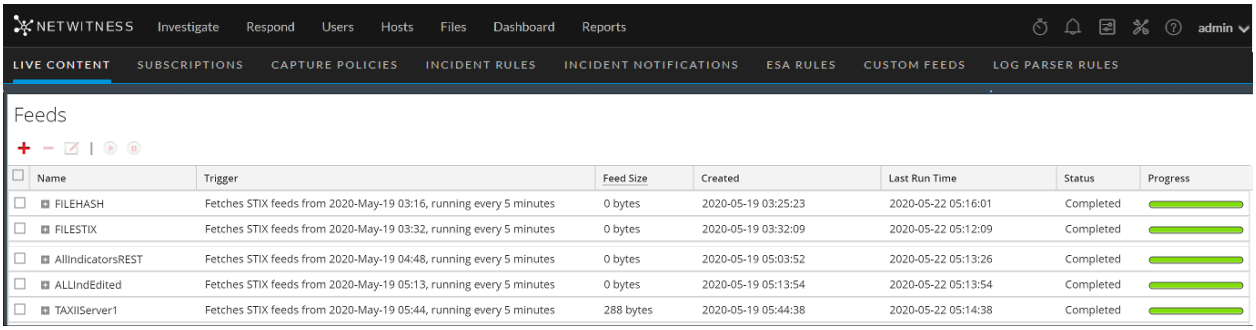
d. Select the language key to apply to the data in each column from the drop-down list. The meta displayed in the drop-down list is based on the meta available for the service define values. You can also add other meta based on advanced expertise.

e. Click **Next**.

The Review dialog is displayed.



10. Anytime before you click **Finish**, you can:
    - Click **Cancel** to close the wizard without saving your feed definition.
    - Click **Reset** to clear the data in the wizard.
    - Click **Next** to display the next dialog (if not viewing the last form).
    - Click **Prev** to display the previous dialog (if not viewing the first form)
11. Review the feed information, and if correct, click **Finish**.
12. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file are listed in the Feed grid and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.



Note: Health and Wellness raises alerts when the available heap memory of Context Hub server is critically low and the status displays as unhealthy due to low memory. For more information on how to troubleshoot the `OutOfMemoryError` on the Context Hub Server, see "Troubleshooting" in the *Live Services Management Guide*.

MetaCallback Feeds using CIDR Index Range for IPv4 and IPv6

You can use CIDR index ranges for IPv4 and IPv6 in custom MetaCallback feeds. As with other custom feeds, you must create feed data file in .csv format, and a feed definition file in .xml format.

Note: Using MetaCallback feeds with CIDR index ranges is supported only through the Advanced Configuration wizard or the REST interface.

The following example shows the content of both a .csv file and an .xml file for a MetaCallback feed using CIDR index ranges for IPv4 or IPv6.

CSV File Content

```
192.168.0.0/24, Sydney
192.168.1.0/24, Melbourne
```

XML File Content

```
<?xml version="1.0" encoding="UTF-8"?><FDF
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="feed-definitions.xsd">
<FlatFileFeed name="ip_test" path="ip_test.csv" separator="," comment="#">
<MetaCallback name="DstIP" valuetype="IPv4" apptype="0" truncdomain="false">
<Meta name="ip.dst"/>
</MetaCallback>
<LanguageKeys>
<LanguageKey name="alert" valuetype="Text" />
</LanguageKeys>
<Fields>
<Field index="1" type="index" range="cidr"/>
<Field index="2" type="value" key="alert" />
</Fields>
</FlatFileFeed>
</FDF>
```

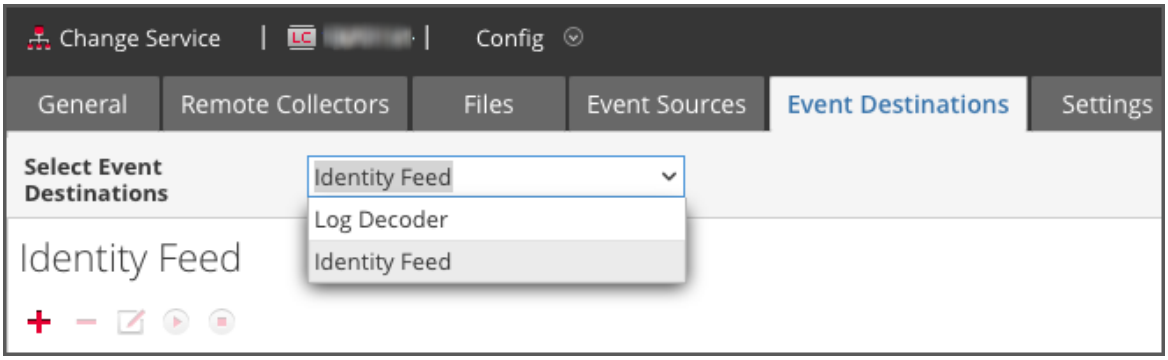
Note: To configure a CIDR index range for feeds with single or multiple MetaCallbacks of value type IPv4 or IPv6, the field of type index MUST contain a range attribute with range="cidr". Also, configuring "cidr" index ranges for feeds with MetaCallbacks of multiple different value types is not supported.

Creating and Managing an Identity Feed

You can easily create an Identity feed and populate it to selected Decoders and Log Decoders. After completing this procedure, you will have created an Identity feed.

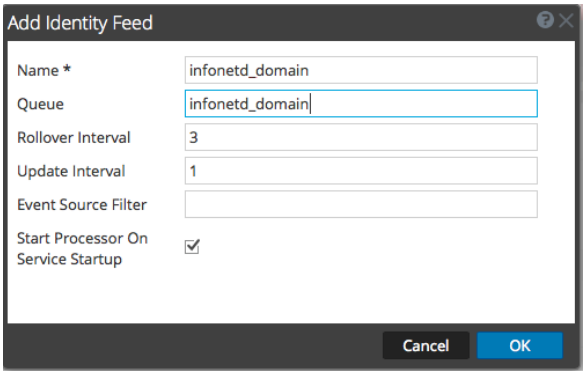
To create an identity feed:

1. Add a destination for the feed.
   a. Go to **(Admin) > Services** and in the **Services**.
   b. In the list of services, select a **Log Collector** service, and select **View > Config**.
   c. Select the **Event Destinations** tab.
   d. In the Select **Event Destinations** field, select **Identity Feed**.



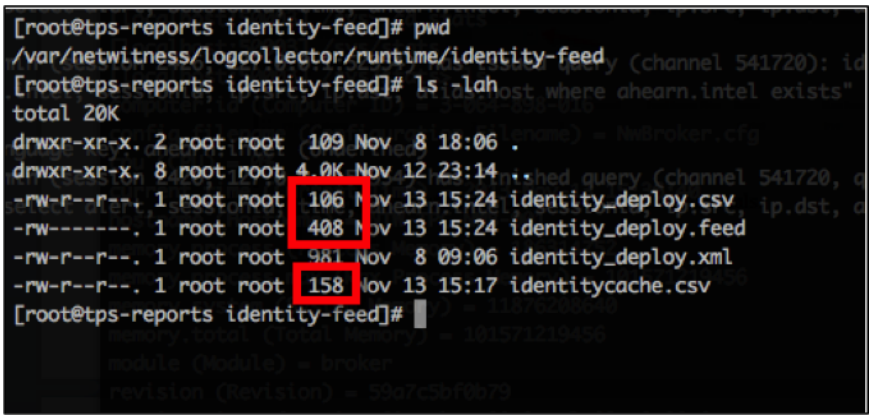
   e. Click and enter a unique name for the feed.

   The Queue name identifies the feed within the Log Collector. Use the name of the feed for the Queue.



   f. Click **OK**.
2. Test generation of messages.
   a. Have users log into Windows boxes on the domain to generate the appropriate log messages on the domain controllers for testing.

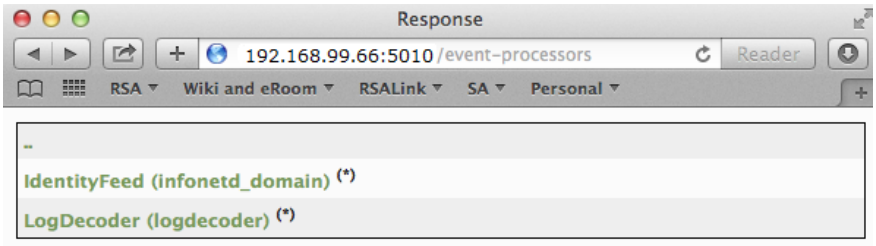
b. Verify that data is written to the feed files. SSH to the Log Decoder/Collector or Virtual Log Collector being configured. Navigate to `/var/netwitness/logcollector/runtime/identity-feed` and verify that the `Identity_deploy` files are getting populated with data.



c. Open up a web browser (Non-Internet Explorer browsers preferred) and log in to the REST interface of the Log Collector. Use administrative credentials when logging in. For example, if the IP address of your Log Collector is 192.168.99.66, the URL would be:

- SSL not enabled: **http://192.168.99.66:50101/event-processors**
- SSL enabled: **https://192.168.99.66:50101/event-processors**

The browser screen should look like this:

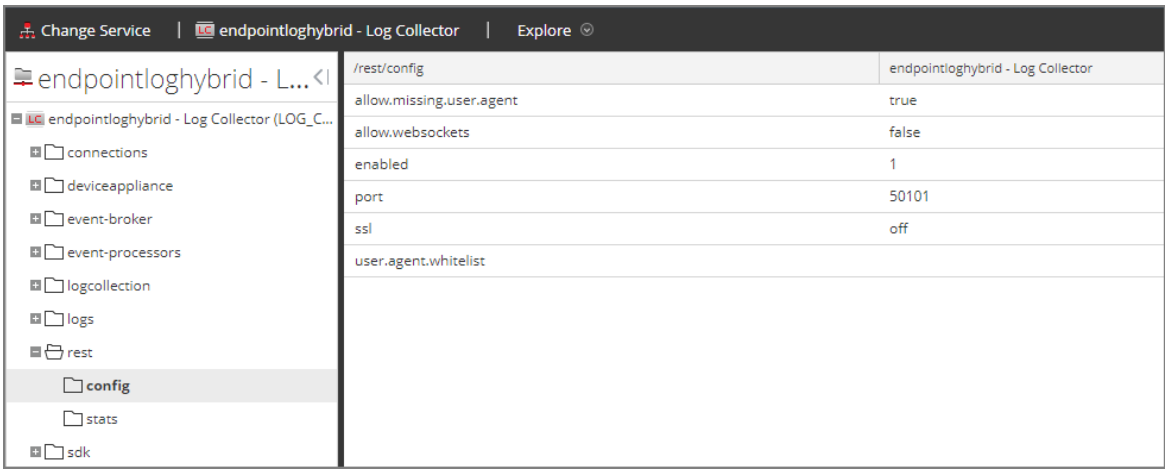


The screen contains the name of the identity feed you created earlier (`infonetd_domain`, in this example).

For the identity feed to function correctly, port 50101 must be active on the Log Collector, and you must determine whether SSL encryption is active.

d. Go to  **(Admin) > Services > <*Log Collector being setup*>**  **> View > Explore**.

e. In the left pane, expand **rest > config**.



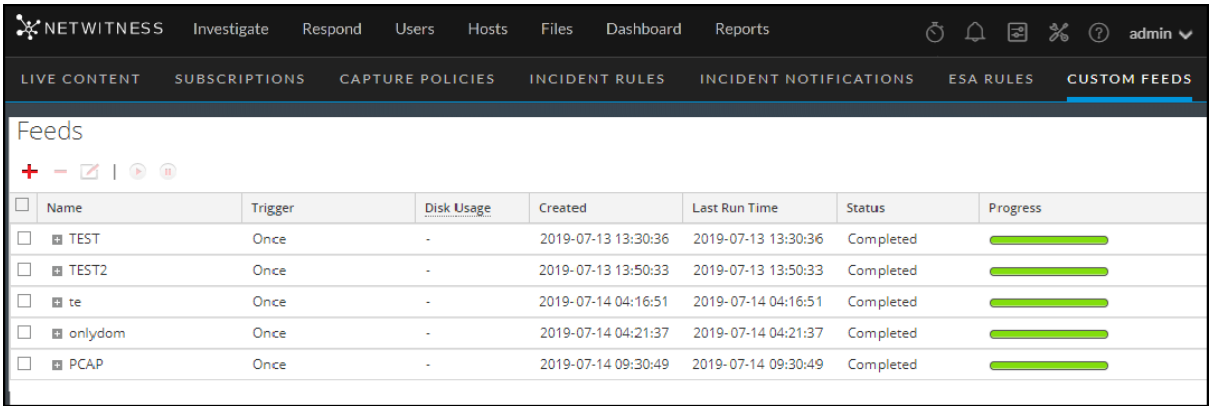
For REST to be active, **enabled** must be set to **1**.

f. Note the value for **ssl**. If SSL should be enabled for your environment, this must be set to **on**.

Note: If you changed the setting for either the **enabled** or **ssl** option you must restart the Log Collector service before moving forward.

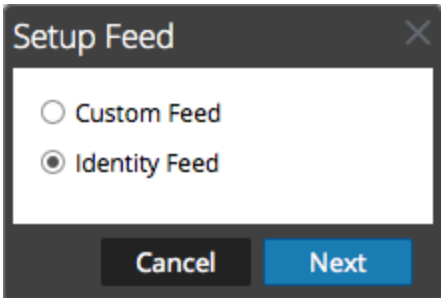
3. Go to **(Configure) > Custom Feeds**.

The Feeds dialog is displayed.



4. In the toolbar, click .

The Setup Feed dialog is displayed.



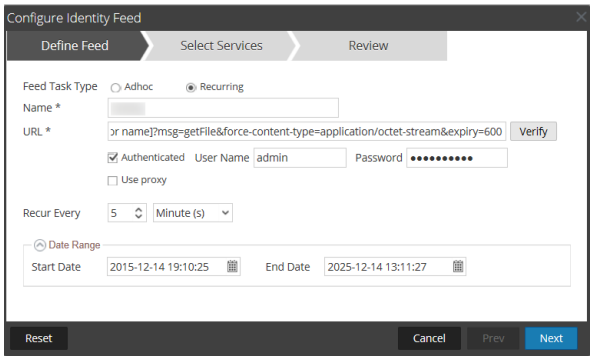
5. Make sure **Identity Feed** is selected and click **Next**.

The Configure Identity Feed panel opens with the **Define Feed** tab displayed.

6. (Conditional) You can create an on-demand or recurring feed.

   - To define an on-demand Identity feed task that executes once, select **Adhoc** in the **Feed Task Type** field, type the feed **Name**, and browse for and open the feed.
   - To define a recurring Identity Feed task that executes on a recurring basis, select **Recurring** in the **Feed Task Type** field.

     The **Define Feed** dialog includes the fields for a recurring feed.



Note: NetWitness verifies the location where the file is stored, so that NetWitness can check for the latest file automatically before each recurrence.

7. Enter a value and verify the URL field.

   a. In the **URL** field, enter the URL where the feed data file is located. This is the REST API interface that was setup earlier. Make sure you have the following information to construct the URL:

      - The IP address of the Log Collector being used to construct the Identity Feed file.
      - The identity queue name, as set in step 2c.
      - Whether or not SSL is enabled on the Log Collector REST port, as set in step 2f.

      You can construct this value as follows:

      - SSL enabled: `https://<LogCollector>:50101/event-processors/<ID Event processor name>?msg=getFile&force-content-type=application/octet-stream&expiry=600`
      - SSL not enabled: `http://<LogCollector>:50101/event-processors/<ID Event processor name>?msg=getFile&force-content-type=application/octet-stream&expiry=600`

      So, using the example from earlier, the complete value that you would enter into this field is as follows:

      ```
      http://192.168.99.66:50101/event-processors/infonetd_
      domain?msg=getFile&force-content-type=application/octet-
      stream&expiry=600?msg=getFile&force-content-type=application/octet-
      stream&expiry=600
      ```

b. For the URL verification to work correctly, it is important that the NetWitness UI server can access the Log Collector's REST API port (50101). This can be tested by going to the NetWitness UI server via SSH. Once there, run the following command:

- SSL enabled: `curl -vk https://<ip of log collector>:50101`
- SSL not enabled: `curl -v http://<ip of log collector>:50101`

If the `curl` command does not connect then there may be a network firewall or routing issue between the NetWitness UI server and the Log Collector.

Example of a bad connection:

```
* About to connect() to 192.168.99.66 port 50105 (#0)
* Trying 192.168.99.66... No route to host
* couldn't connect to host
* Closing connection #0
curl: (7) couldn't connect to host
```

Example of a good connection:

```
* About to connect() to 192.168.99.66 port 50105 (#0)
* Trying 192.168.99.66... connected
* Connected to 192.168.99.66 (192.168.99.66) port 50105 (#0)
> GET / HTTP/1.1
> User-Agent: curl/7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7
NSS/3.19.1 Basic ECC zlib/1.2.3 libidn/1.18 libssh2/1.4.2
> Host: 192.168.99.66:50105
> Accept: */*
>
< HTTP/1.1 401 Unauthorized
< Content-Length: 71
< Connection: Keep-Alive
< Pragma: no-cache
< Expires: -1
< Cache-Control: no-cache, no-store, must-revalidate
< WWW-Authenticate: Basic realm="NetWitness"
< Content-Type: text/xml; charset=utf-8
<
<?xml version="1.0" encoding="utf-8"?>
<error>401 Unauthorized</error>
* Connection #0 to host 192.168.99.66 left intact
* Closing connection #0
```

8. The REST API requires a username and password when attempting to pull the `identity_deploy.csv` file from the Log Collector. This can be any username and password that is available on the service itself. For more information, see the "Services Security View" topic in the *Hosts and Services Guide*.

   To see which accounts are available, go to  **(Admin) > Services > <log collector being setup> > Actions > View > Security**.

   Under the Users table, you see all the users that can be used in this step. It is suggested that a separate user account is created specifically for this setup, and is used nowhere else in the environment, for added security. For details, see "Add a User and Assign a Role" in the *System Security and User Management Guide*. (Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.)

9. To define the recurrence interval, do one of the following:

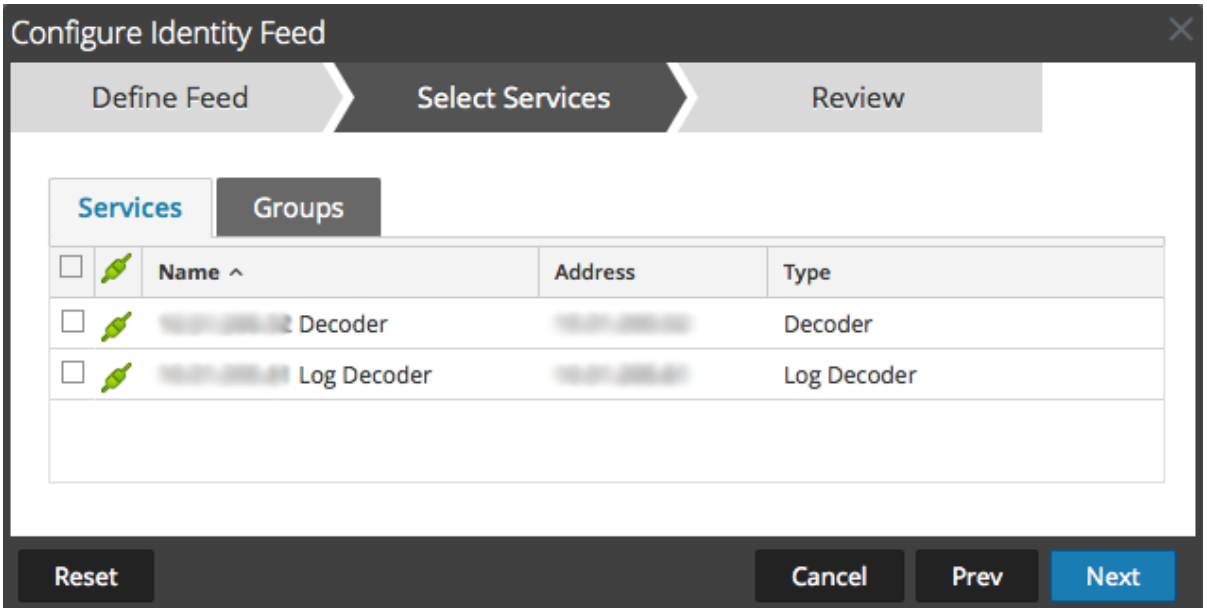
- Specify the number of minutes, hours, or days between recurrences of the feed.
- Enter the date range for the execution of the feed to recur, specify the **Start Date** and time and the **End Date** and time.

10. If using SSL encryption, you need to install the REST API SSL certificate for the Log Collector into the NetWitness UI server. For more information, see Import the SSL Certificate.

    If, after importing the SSL certificate, the verification of the URL still fails, see Cannot Verify Identity Feed URL.

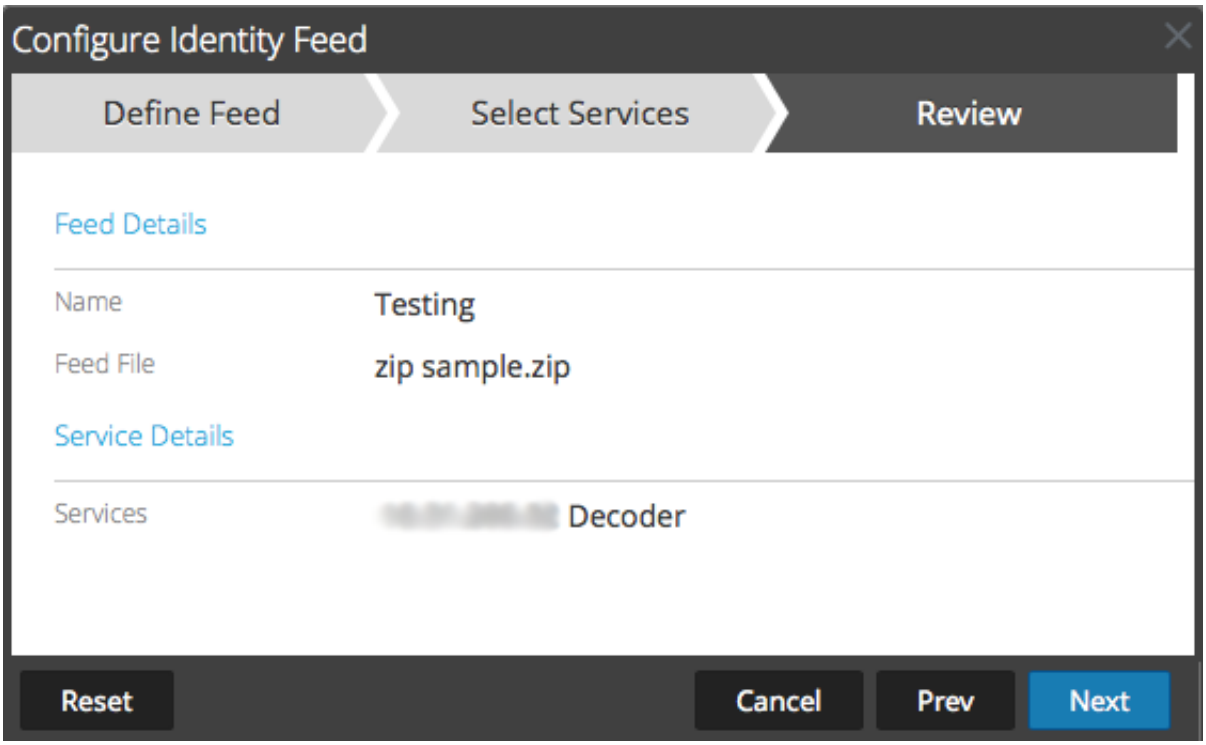
11. Click **Verify** to verify your identity feed configuration before you proceed to the Select Services dialog.
12. Click **Next**.

    The Select Services dialog is displayed.



13. To identify services on which to deploy the feed, select one or more Decoders and Log Decoders and click **Next**.
14. Click the **Groups** tab, select a group, and click **Next**.

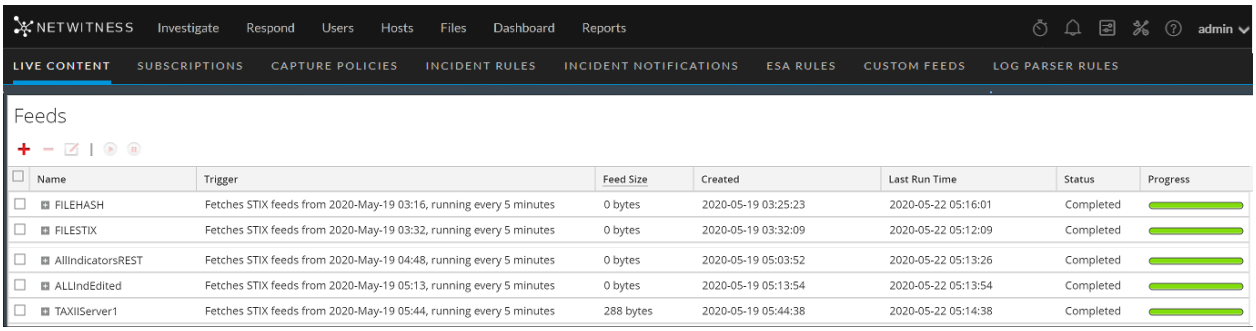
    The Review dialog is displayed.



Note: If a group of devices with Decoders and Log Decoders is used to create recurring or custom feeds and this group is deleted, you can edit the feed and add a new group to the feed.

15. Anytime before you click **Finish**, you can:
    - Click **Cancel** to close the wizard without saving your feed definition.
    - Click **Reset** to clear the data in the wizard.
    - Click **Next** to display the next form (if not viewing the last form).
    - Click **Prev** to display the previous form (if not viewing the first form).
16. Review the feed information, and if correct, click **Finish**.

Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file are listed in the Feed grid and progress bar tracks completion. You can expand or collapse the entry to see how many services are included, and which services were successful.



Import the SSL Certificate

If SSL is configured on the Identity feed's Log Collector, follow these steps to import the Log Collector's SSL certificate into the NetWitness UI server key store. If this certificate is not imported, the NetWitness UI server will be unable to pull the Identify feed file from the Log Collector.

1. To pull the SSL certificate off the Log Collector, SSH into the NetWitness UI server and run the following command:

   ```
   echo -n | openssl s_client -connect <HOST>:<PORT> | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/<SERVERNAME>.cert
   ```

   This command saves the SSL certificate to `/tmp/<SERVERNAME>.cert`. For example:

   ```
   echo -n | openssl s_client -connect 192.168.99.66:50101 | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/logcollector.cert
   ```

2. To import the SSL certificate into the NetWitness UI server, SSH into the UI server and run the following command:

   ```
   keytool -importcert -alias <name an alias for the cert> -file <the cert file pathname> -keystore /etc/pki/java/cacerts
   ```

   For example:

   ```
   keytool -importcert -alias logcollector01 -file /tmp/logcollector.cert -keystore /etc/pki/java/cacerts
   ```

3. The system requests a password. Enter the password for the keystore on the NetWitness UI server, not for the jetty keystore. The default password is **changeit**.

4. Restart **jettysrv** to allow jetty to read the new certificate in the store.

Cannot Verify Identity Feed URL

If the Identity feed URL cannot be verified, and you are using SSL, make sure you followed the steps in Import the SSL Certificate.

If there are issues, it is possible that the internal name of the certificate does not match the hostname of the Log Collector. The following procedure checks this.

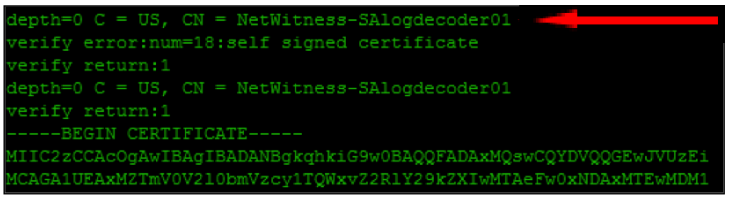
1. SSH to the NetWitness UI server.
2. Run the following command to output the CN name of the SSL cert:

   ```
   echo -n | openssl s_client -connect <log decoder>:50101 | sed -ne '/BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
   ```

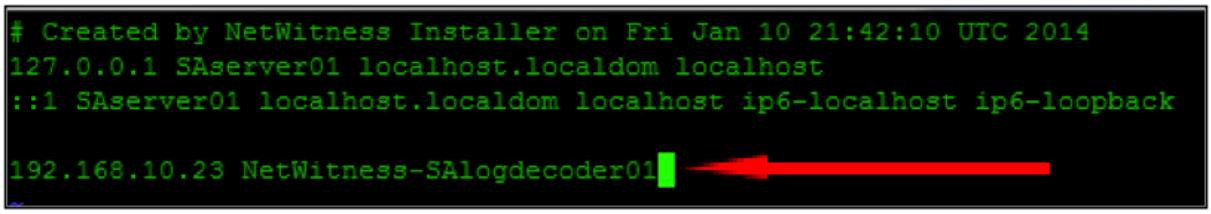
   For example:

   ```
   echo -n | openssl s_client -connect salogdecoder01:50101 | sed -ne '/BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
   ```

3. Retrieve the CN name of the SSL certificate.



4. Edit the `/etc/hosts` file and add the IP address and CN name to the file.



5. Restart the network service on the appliance.
6. Confirm that the name placed in the **/etc/hosts** file is used instead of the FQDN or IP address in the Identity feed URL.
7. Re-verify the Identity feed URL.

Investigating an Identity Feed

An identity feed tracks interactive log on events from the Windows operating system. Identity feeds do not track interactive log off events.

In order for an identity feed to process events and tag them, the events need to be collected using a Windows Log Collection module where an Active Domain Controller or non-Domain Controller is configured. Note that identity feeds can only be processed via an Identity Feed Event Processor.

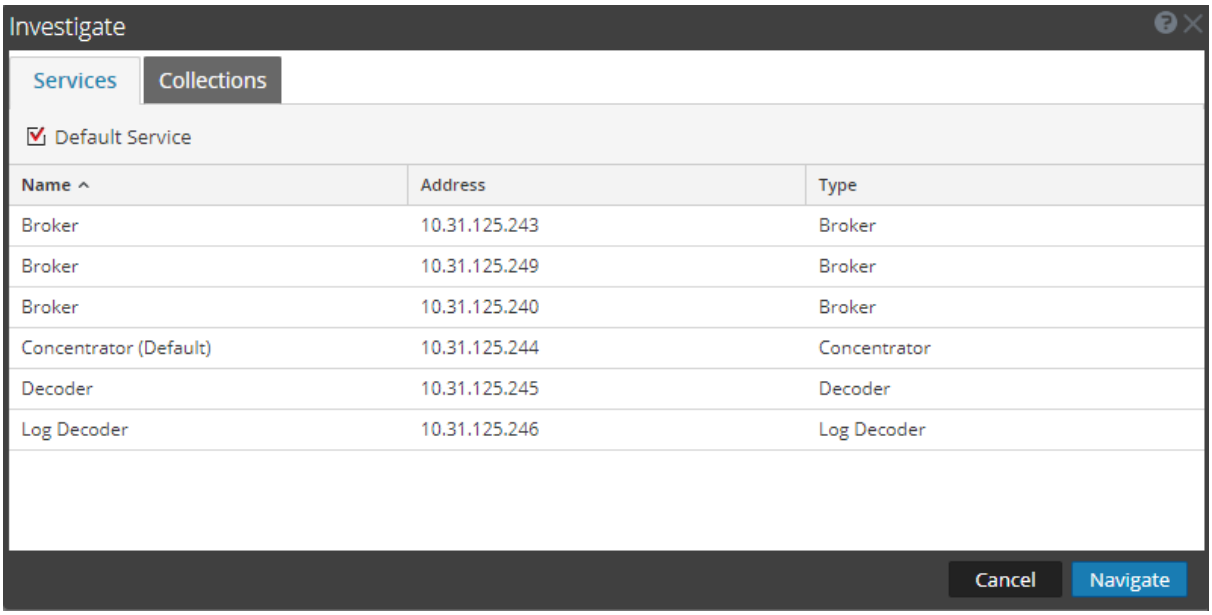
Note: An identity feed only tracks one log in at a time. If two users log in to a system at the same time, the second user will overwrite the first user's data in the identity feed.

Once you have created an identity feed, you can view the results by investigating the feed.

To investigate a configured identity feed:

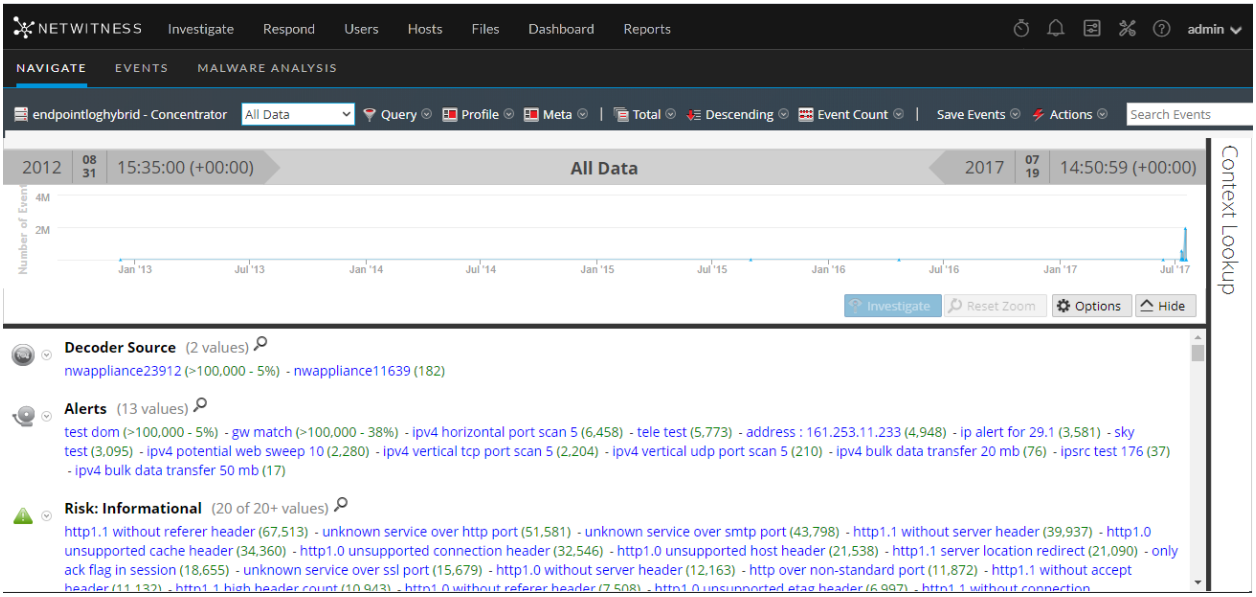
1. Go to **INVESTIGATE > Navigate**.

   If no default service is selected, the Investigate dialog is displayed.



2. Select a service, usually a Concentrator, and click **Navigate**.
3. Select **Load Values** to retrieve meta data.

In the Values panel, scroll down to find the Meta Keys:



The identity feed provides information to selected Decoders and Log Decoders. It associates the Host IP data from the Windows operating system to the user logging into that Host in order to tag all logs associated with that IP and investigate.

Editing a Feed

This topic provides instructions for editing a custom feed using the Custom Feed Wizard.

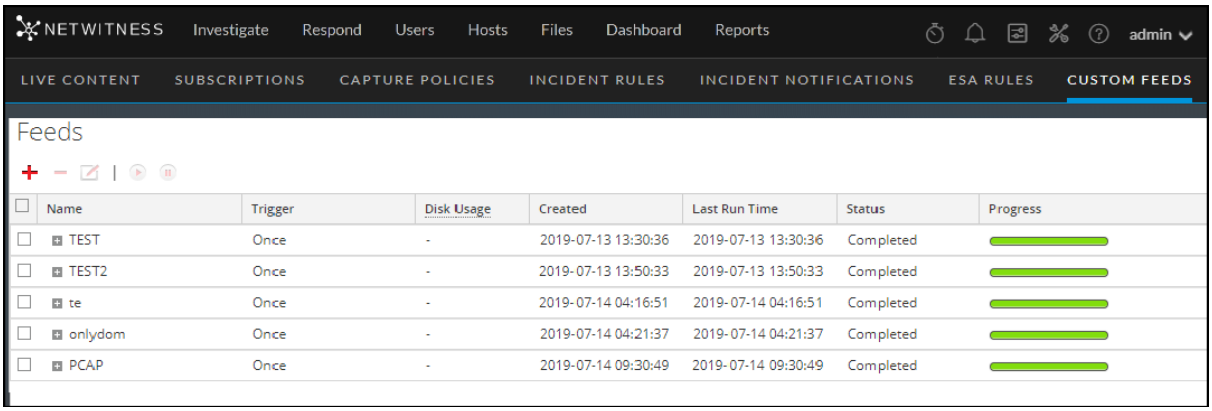
After you edit a feed:

- The feed (**.zip** format) or the file used to create the feed (**.csv** or **.xml**) has been downloaded and edited.
- The feed has been recreated with the updated file and new feed specifications.

To edit an existing feed:

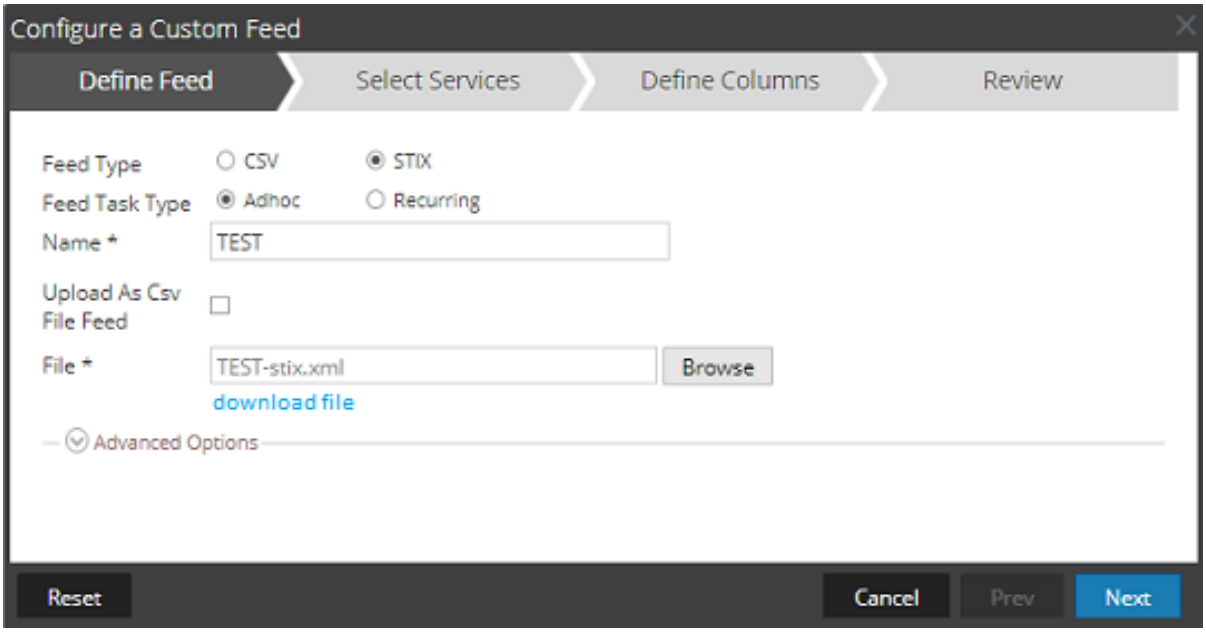
1. Go to **(Configure) > CUSTOM FEEDS**.

   The Custom Feeds dialog is displayed.



2. In the toolbar, select a feed and click .

   The Configure Custom Feed or Configure Identity Feed panel opens in the Custom Feed wizard.



3. If you want to edit the feed file:
   a. Click **download file**.

      For an Identity feed, the .zip file is downloaded. For a custom feed, the .csv or .xml file is downloaded to your local file system.

   b. Edit and save the file.

   c. In the **Define Feed** tab, browse for and open the edited file.
4. Edit any other parameters in the **Define Feed** tab, **Select Services** tab, and **Define Columns** tab that apply to the type of feed.
5. Anytime before you click **Finish**, you can:
   - Click **Cancel** to close the wizard without saving your changes.
   - Click **Reset** to clear the data in the wizard.
   - Click **Next** to display the next form (if not viewing the last form).
   - Click **Prev** to display the previous dialog (if not viewing the first form).
6. In the **Review** tab, review the feed information, and if correct, click **Finish**.

The feed is added to the feeds list and progress bar tracks completion. Upon successful creation of the feed definition file, the Create Feed wizard closes, and the feed and corresponding token file is listed in the Feeds list. You can expand or collapse the entry to see how many services are included, and which services are successful.

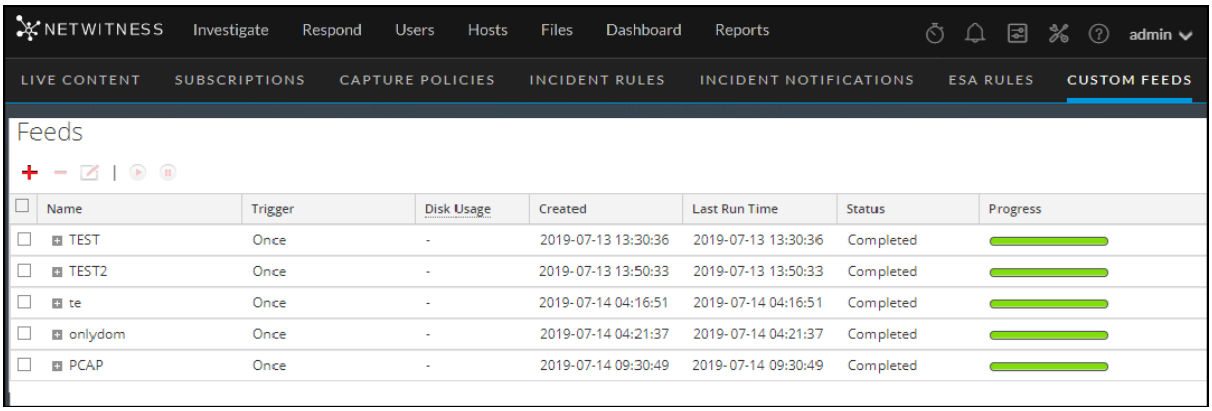
Removing a Feed

This topic provides instructions for removing a feed. You might want to remove a feed when some or all of the information in the feed is no longer useful for your organization.

To remove a feed:

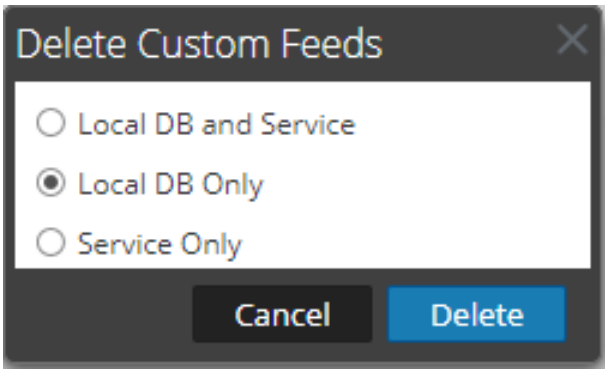
1. Go to **(Configure) > CUSTOM FEEDS**.

   The Custom Feeds dialog is displayed.



2. In the toolbar, select a feed and click .

   The Delete Custom Feeds dialog is displayed.



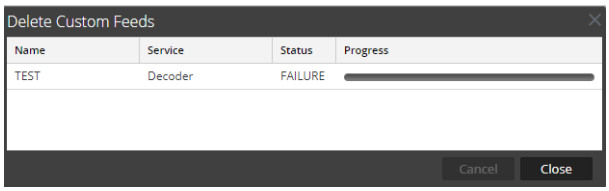
   You can select one of the following options to delete the feed:

   - If you choose to delete the feed from **Local DB and Service**, the feed is deleted from both the service and the local NetWitness box. The deleted feed will no longer seen on the NetWitness user interface.
   - If you choose to delete the feed from **Local DB Only**, the feed is deleted from the local NetWitness box. The deleted feed will not be seen on the NetWitness user interface; however, the last deployed version of the feeds will be present on the service. The undeployed feeds will be deleted forever.
   - If you choose to delete the feed from **Service Only**, the feed is deleted from the service. The deleted feed will appear on the NetWitness user interface and can be deployed again

3. Select which feed you want to delete and click **Delete**.

   A warning dialog is displayed.

4. Click **yes** to confirm that you want to delete the feed from the selected areas.



Subscribing to Live Resources

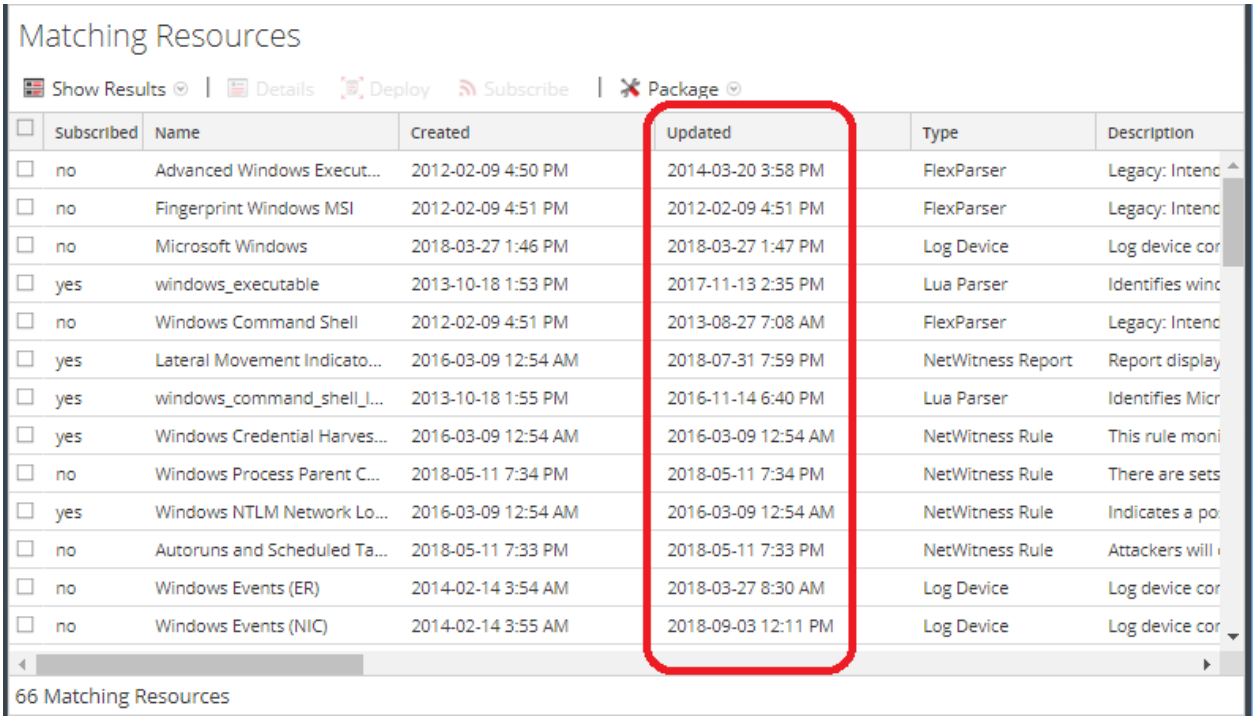
This section describes subscriptions in Live.

Threats and the corporate landscape change over time. NetWitness periodically reviews existing content to determine whether it needs to be updated based upon current campaigns, or has become irrelevant due to changes in technology or attack techniques and tools.

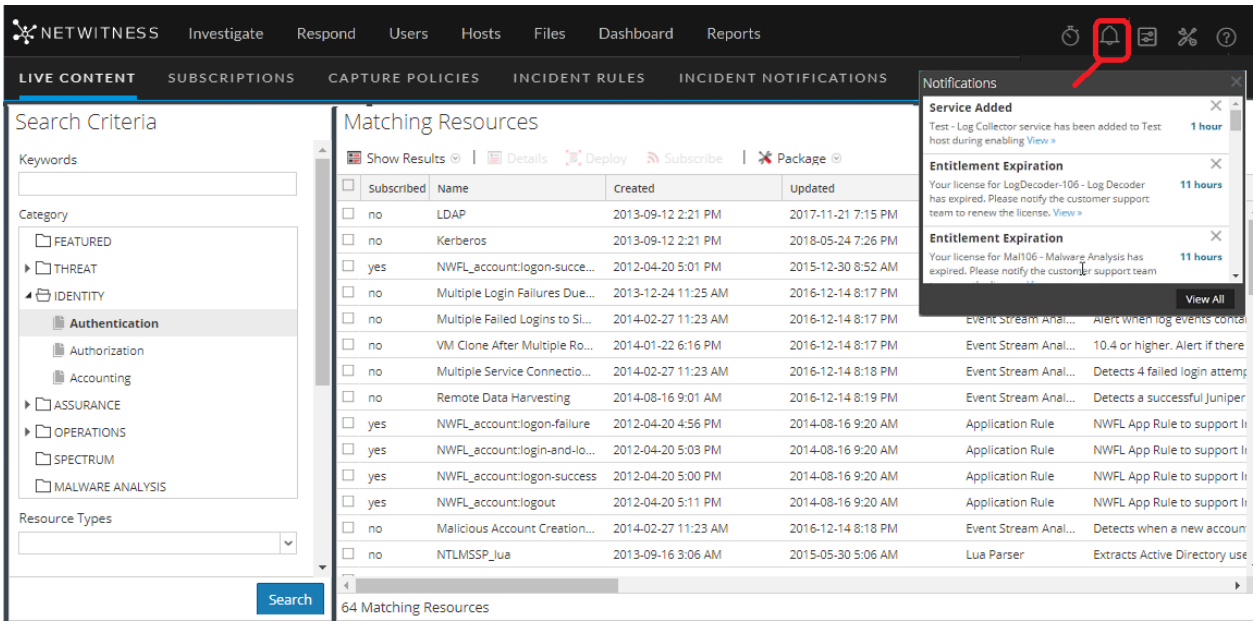
You can discover new content by using the What's New dashlet within the Default Dashboard, or by searching through RSA Live by data range since last deployed. Be sure to subscribe to any content for which you want to receive update notifications.

Subscription Updates

When you view resources in the Matching Resources panel of the Live Content view, there is a column named **Updated**:



This value is also displayed when you select the detailed view for a resource. Every time a resource changes, its **Updated** value changes to match the specific update. If you are subscribed to a resource, and it gets updated, your system is automatically updated with the latest version, and you receive a notification. You can view your notifications by clicking the Notification icon, , from anywhere in the NetWitness UI.



You can also get email notifications when subscribed resources are updated. System Administrators can add email addresses in the **(Admin) > SYSTEM > Live Services** view. For more information, see the "Live Services Configuration Panel" topic in the *System Configuration Guide*.

Adding Subscribed Resources for Deployment to Services

1. Go to **(Configure) > Subscriptions > Deployments**.
2. In the **Groups** panel, select a group.
   Subscribed resources, if any, are listed in the Deployments tab Subscriptions panel.
3. In the **Subscriptions** panel, click .
   The Add Subscription dialog, which lists subscriptions available for deployment, is displayed.
4. Select the subscribed resources that you want to deploy to the services group.
5. Click **Save**.
   The dialog closes and the subscriptions are added to the listing in the Deployments tab, Subscriptions panel. This stages the resources for deployment at the next synchronization.
6. You can click the Synchronize icon, , to immediately synchronize your changes.

Deleting a Subscription

When you delete a subscription to a resource, deployed instances of the resource are not deleted. The deployed resource remains on services until explicitly removed, but the resource is no longer synchronized with the resource in NetWitness Live.

To delete a subscription:

1. Go to **(Configure) > Subscriptions**.
2. In the **Subscriptions** tab, select the subscriptions you want to delete.
3. Click , then choose **Delete** to delete your selected resources or **Delete All** delete all subscriptions.

   A dialog asks for confirmation that you want to delete the subscription.
4. To confirm removal, click **Yes**.

   Your selected subscriptions are deleted from the subscriptions list, but any deployed instances of the subscribed resource remain on the services.

Removing Subscribed Resources from the Deployments Subscriptions Grid

Subscriptions that are selected for deployment to a service group are deployed during synchronization.

You can remove subscriptions from the Live (Configure) > Subscriptions > Deployments panel, but any that have actually been deployed to services remain deployed until someone removes them.

To remove resources from the Deployments tab Subscriptions panel:

1. Go to **(Configure) > Subscriptions > Deployments**
2. In the **Groups** panel, select a group.

   Subscribed resources, if any, are listed in the Subscriptions panel.
3. In the Subscriptions panel, click .

   A dialog requests confirmation that you want to delete the resource from the service group. The resource is removed from the Deployments tab Subscriptions panel, but is not removed from services on which it is deployed.

Subscribe and Unsubscribe to a Resource

When you subscribe to resources, you will receive notification when new versions of the resources are available.

To subscribe to a resource:

1. Go to **(Configure) > Live Content**.
2. In the **Search Criteria** panel, specify search criteria and click **Search**.
3. Select one or more resources and click Subscribe.

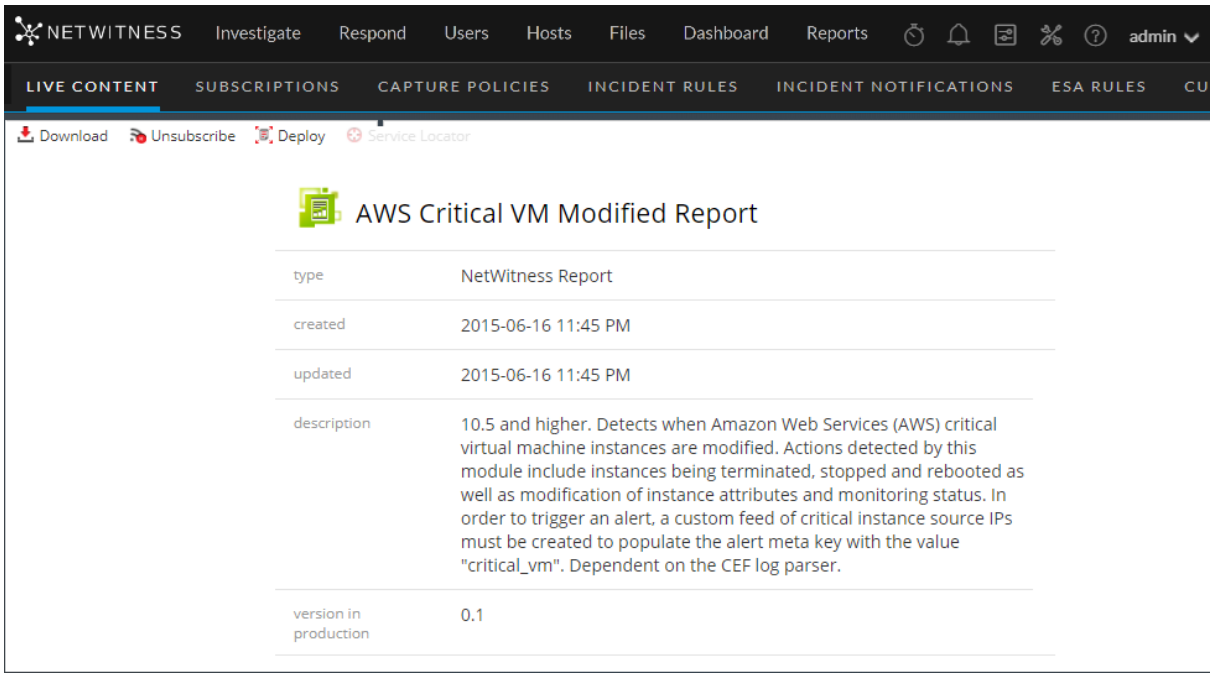
   A confirmation dialog is displayed: **By subscribing to these resources, you are indicating that you wish to receive notification when new versions are available.**
4. To confirm that you want to subscribe to the resource, click **OK**.

   The resource is added to the subscriptions managed in the Subscriptions tab and is available for deployment in the Deployments tab.

When unsubscribing from a resource, you have the option to leave the resource on services on which it is deployed or to remove it from services.

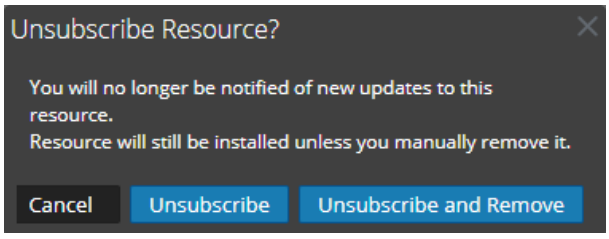
To unsubscribe from a resource:

1. Open a detailed view of a resource in one of the following ways:
   - Perform a search, **(Configure) > Live Content** > enter search criteria, then select the resource in the Matching Resources panel, then click Details.
   - View subscriptions, **(Configure) > Subscriptions**, select the resource from the Subscriptions list, then click Details.



2. With the detailed view of a resource displayed. click Unsubscribe.

   A confirmation dialog is displayed.



3. Do one of the following:
   - To confirm that you want to unsubscribe from the resource and leave it on the services where it is deployed, click **Unsubscribe**.

- To confirm that you want to unsubscribe from the resource and remove it from the services where it is deployed, click **Unsubscribe and Remove from Services**.
- To close the dialog without unsubscribing, click **Cancel**.

The selected action is applied.

Viewing Subscribed Resources Selected to Deploy on Services

In the  **(Configure) > Subscriptions > Deployments** tab you can view subscribed resources that have been selected for deployment on services.

To view subscribed resources that have been selected for deployment on services:

In the **Groups** panel, select a group, and expand it to view services in the group. The resource subscriptions selected for deployment are listed in the Deployments tab Subscriptions panel.

Miscellaneous Live Services Procedures

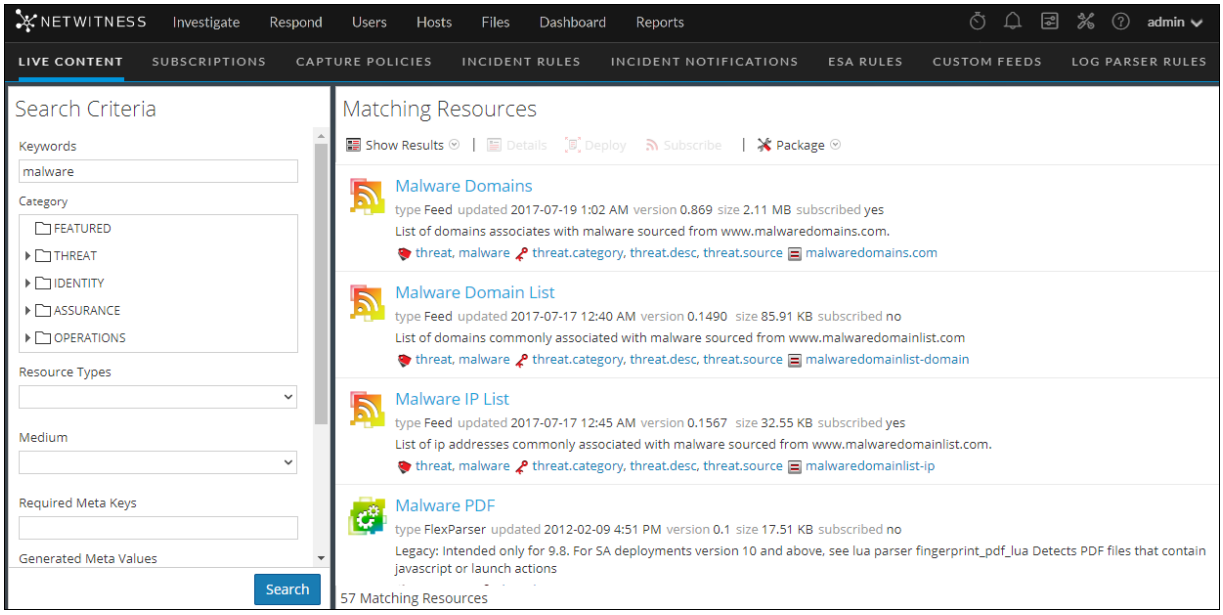
This section describes several other procedures.

Displaying Resource Details in Live Resource View

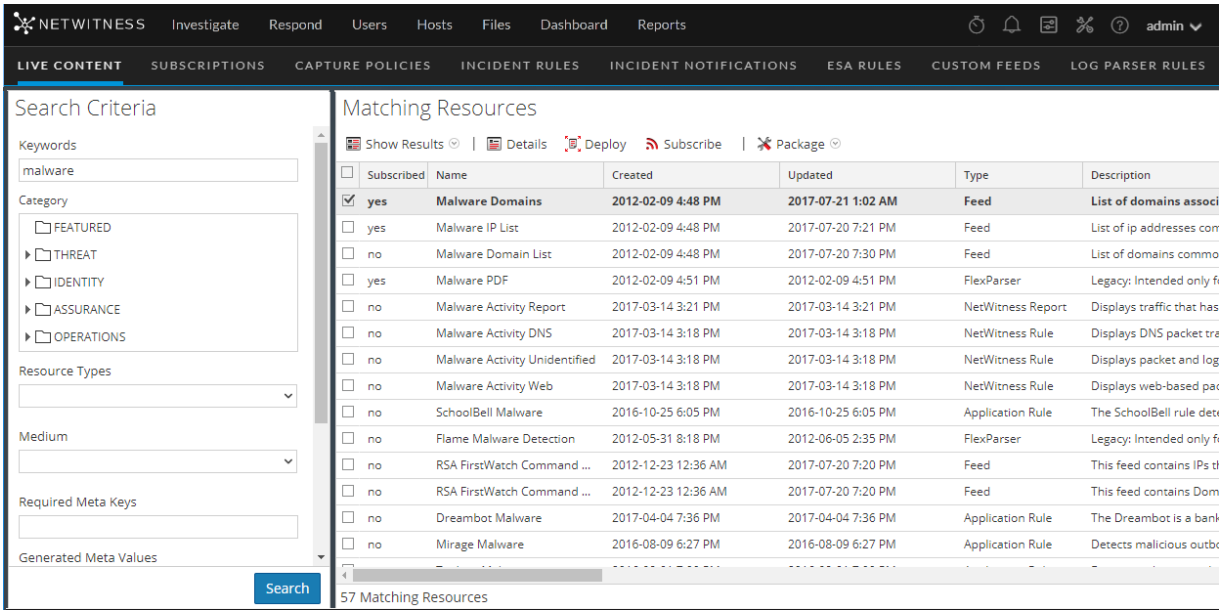
After you select a resource (in the Live Resource View), you can view its detailed information.

To open a separate tab in the Live Resource view with details of a selected resource, do one of the following:

- If you are viewing the **Detailed Results**, click the resource type icon or the resource name.



- If you are viewing the list results, double-click a resource or select a resource and click **Details**.



Downloading a Resource

You can download a single resource from the Live Resource View.

To download a resource:

1. Go to **(Configure) > Live Content**.
2. In the **Search Criteria** panel, enter the criteria needed to return the resource you want to download.
3. Select a single resource, then click Details.
4. Click Download.

The resource is saved as a ZIP archive to your local Downloads folder.

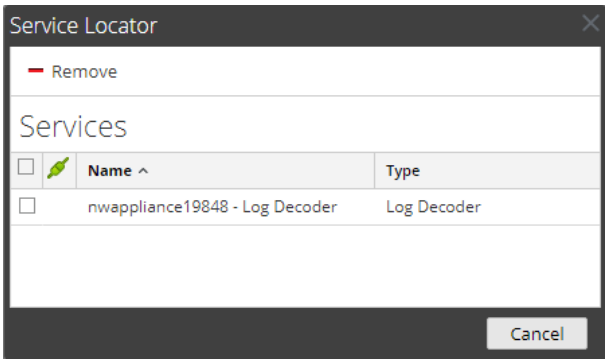
Locating and Removing a Deployed Resource from Services

You can locate and remove a deployed resource from services from the Live Resource View.

To view a list of services on which a resource is deployed:

1. With a resource displayed in the **Resource View**, click Service Locator.

   The Service Locator dialog is displayed.

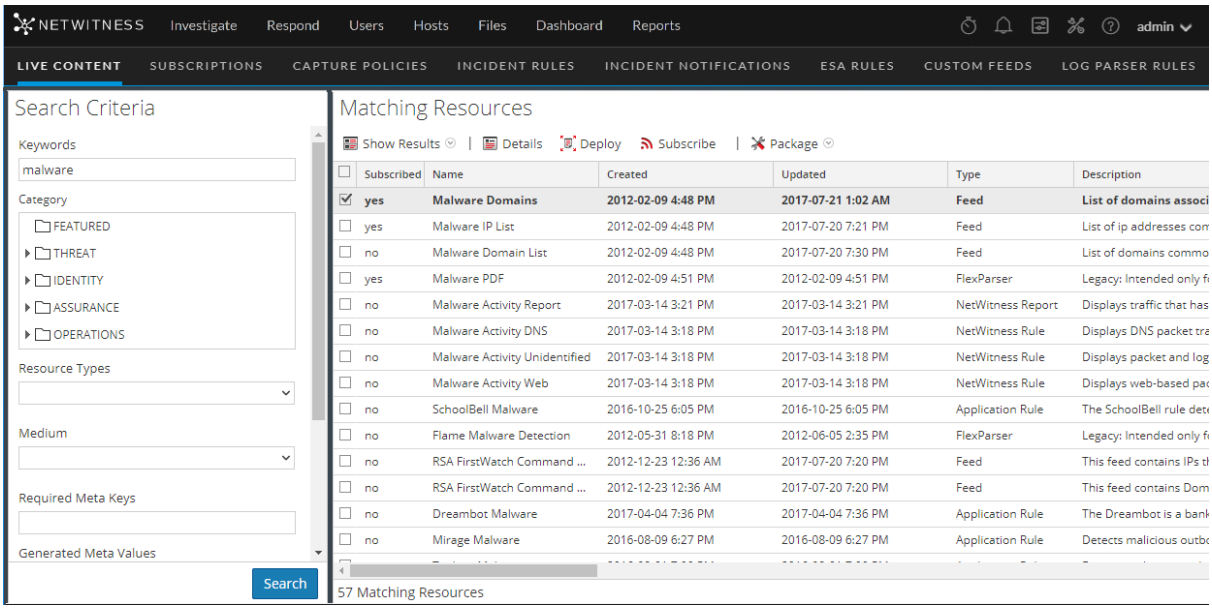


2. Select one or more services in the **Services** list.
3. Click −.

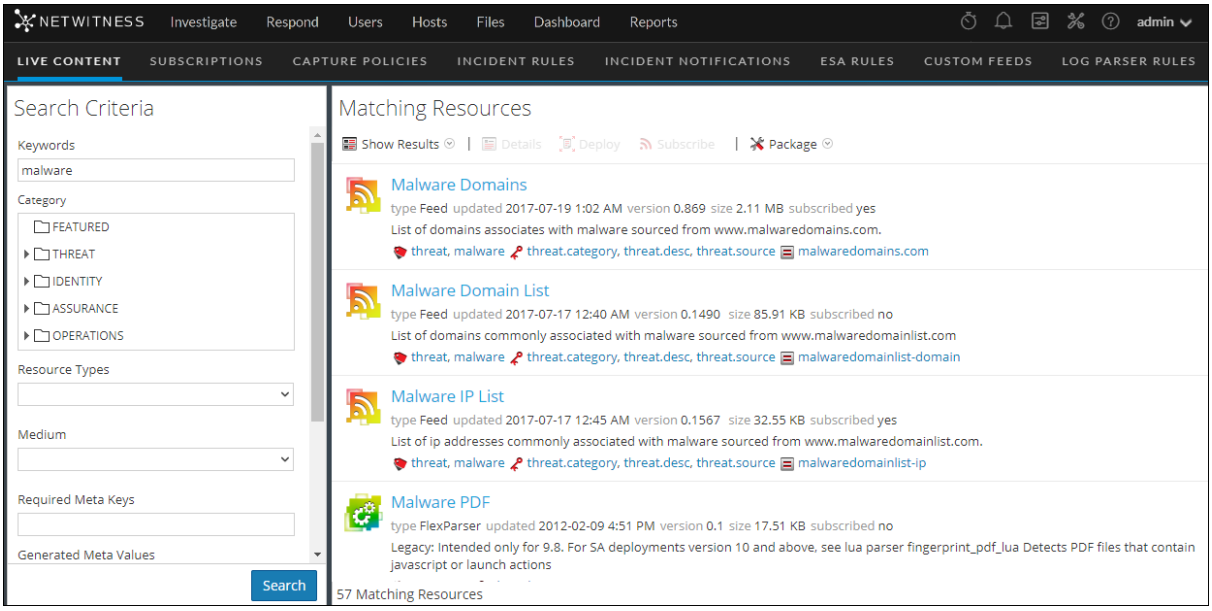
   The resource is removed from the selected services.

Showing Results as a List or in Detail

1. Select **Show Results > Grid** to change to grid results when viewing detailed results.



2. Select **Show Results > Detailed** to change to detailed results when viewing grid results.



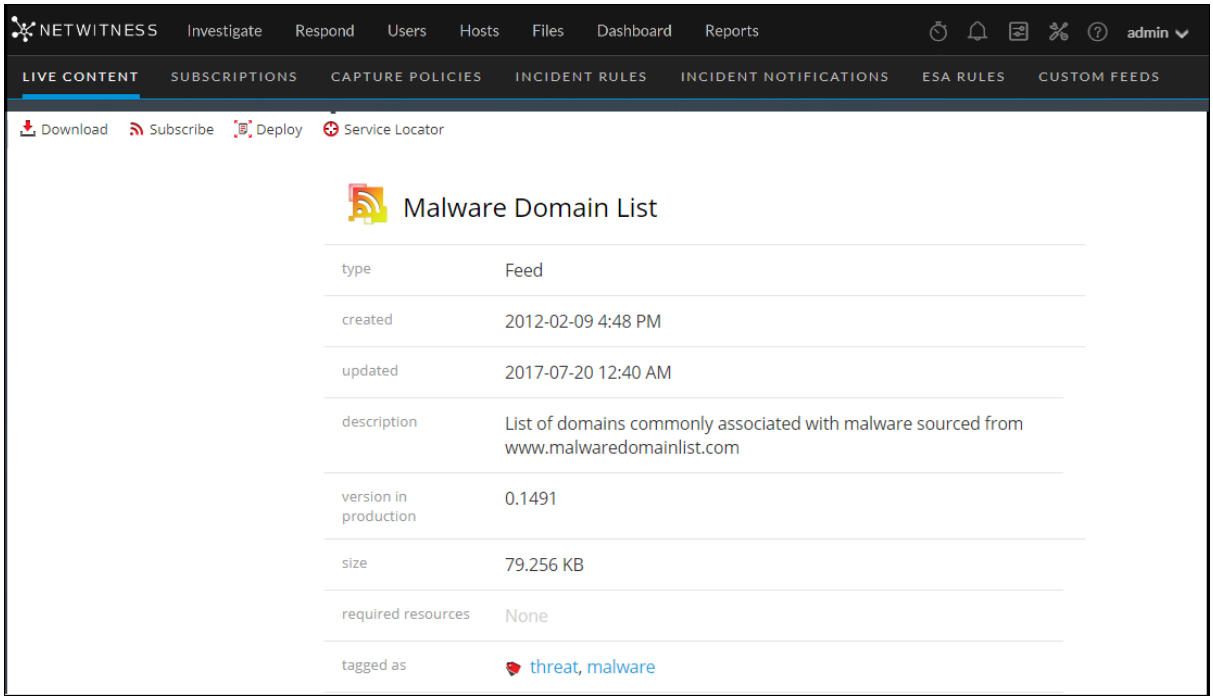
Viewing Resource Details

You can display detailed information about a subscribed resource in the Resource View.

To view details:

1. In the **Subscriptions** tab, select a single subscription.
2. Click Details .

The details of the resource are displayed in the Resource View.



References

This topic is a collection of references, which describe the user interface and more detailed information about how Live works in NetWitness. These topics are presented in alphabetical order.

Live Configure View

In the Live Configure view, NetWitness provides integrated tools for managing Live resources. You can manage resource subscriptions, deployments to services and discontinued resources. The required role to access this view is **Configure Live Resources**. For a high-level description of how to use the different views in NetWitness Live, please read Live Services Management.

To access this view, navigate to  **(Configure) > Subscriptions**. The view has three tabs: Deployments Tab, Subscriptions Tab, and Discontinued Resources Tab.

Deployments Tab

The Deployments tab provides a user interface in the Live Configure view for:

- Viewing subscribed resources that are selected for deployment on services in a service group.
- Selecting subscribed resources to deploy to services in a service group.
- Removing resources that are selected for deployment on services in a service group.

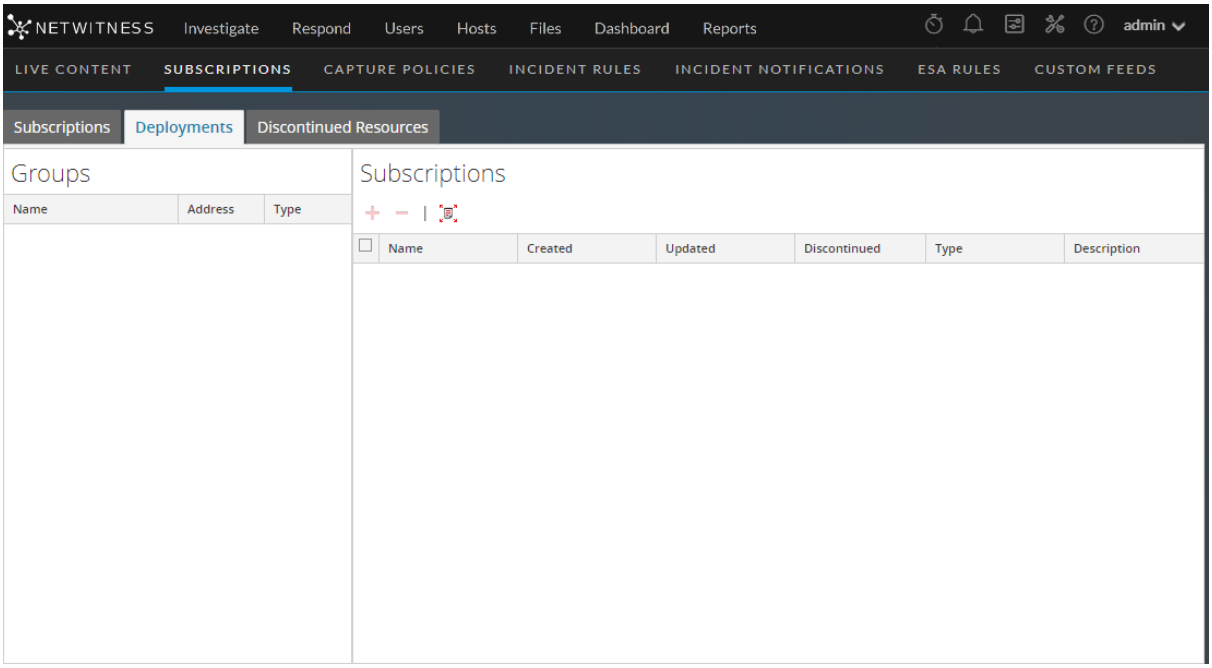
The resources listed here are not deployed immediately after adding to a service group. Instead the subscribed resources are pushed to the services when NetWitness synchronizes with NetWitness Live. The synchronization schedule is configured in the Live Configuration panel. Additionally, you can synchronize immediately in the **(Configure) > Subscriptions > Deployments** tab.

Likewise, resources deleted from the Deployments panel are not deleted from service where they have been deployed. To delete resources from services, delete them in the Live Resource View.

The required permission to access this view is **Manage Live Resources**.

To access this view:

1. Go to **(Configure) > Subscriptions**.
   The **Subscriptions** tab is displayed.
2. Click the **Deployments** tab.



The Deployments tab has two panels: **Groups** and **Subscriptions**.

Groups Panel

The Groups panel is a static display of configured service groups that were created in the Administration Services view. Selecting a group in the Groups panel populates the Subscriptions panel with a list of subscriptions that are selected for deployment on the services in the service group.

Feature	Description
Name	Displays the service group name. Clicking the plus sign displays a nested list of services in the group.
Address	Displays the IP address of each service in the group.
Type	Displays the type of service.

Subscriptions Panel

The following table describes the features in the Subscriptions panel.

Feature	Description
	Click  to open a dialog that lists subscriptions that were added in the Live Search view or in the Live Resource view and are available for deployment.
	Click to delete the selected subscriptions from the deployment list for service group.
	Click  to synchronize your resources to the latest versions available on Live.
Name	Displays name of the resource.
Created	Displays date and time that the resource was created.
Updated	Displays date and time that the resource was last updated.
Type	Displays type of resource.
Description	Displays description of the resource.

Subscriptions Tab

Subscriptions are NetWitness Live resources to which you subscribed in the Live Search view or Live Resource view. When you subscribe to a resource, you agree to receive updates on a regular basis from RSA NetWitness Live. The choices made in the Live Configuration panel determine the synchronization frequency and also whether you receive update notifications through email. In addition, if you don't want to wait for the next update, you can force an immediate synchronization.

The Subscriptions tab provides a way to manage subscriptions. Each resource to which NetWitness is subscribed is listed in this tab.

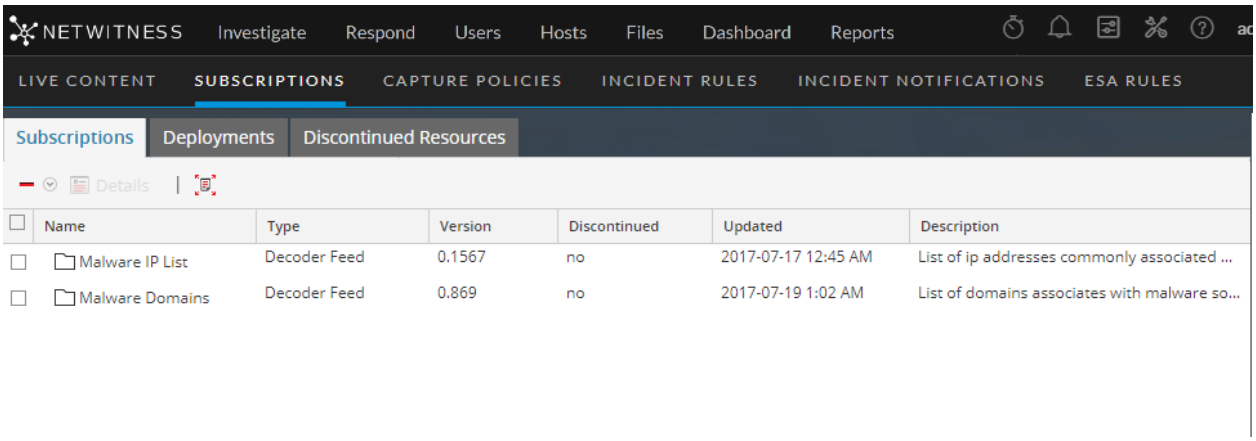
In the Subscriptions tab, you can:

- View all resources to which this NetWitness instance is subscribed.
- Open a detailed view of a subscription in the Live Resource View.
- Delete a subscription.

Note: Subscribing to a resource does not deploy the resource to any services. To deploy one or more subscribed resources, go to the Deployments tab. To deploy a single resource manually, use the Deploy option in the Resource View.

The required permission to access this view is **Manage Live Resources**.

To access this view, in the main menu, select **(Configure) > Subscriptions**. The Subscriptions tab is displayed.



The **Subscriptions** tab has a toolbar and a grid.

Toolbar

This table describes the options available in the toolbar.

Feature	Description
	Deletes the selected subscriptions.
	Displays the details of a single subscribed resource in the Resource View.

Feature	Description
	Check the Live Server for the latest discontinued resources.

Grid

Column	Description
	Selects subscribed resources to view in detail or delete. You can view details for a single resource. You can delete one or more resources from the subscribed resources, in effect unsubscribing.
Name	Displays name of the subscribed resource.
Type	Displays type of subscribed resource.
Version	Displays version of the subscribed resource.
Discontinued	Indicates the status of the discontinued resources for the subscribed resource. **Yes** - Resource is discontinued. **No** - Resource is not discontinued. **--** - The Live Server is not checked for the discontinued resources.
Updated	Displays date and time when the subscribed resource was last updated.
Description	Displays description of the subscribed resource.

Discontinued Resources Tab

The Discontinued Resources tab provides a user interface in the Live Configure view:

- Scan the services for the discontinued resources.
- Remove the discontinued resources from any service or service group.

Note: Discontinued content still appears. With discontinued content there just won't be any updates, and users won't see these items when they search in Live, unless they check the **Include Discontinued Resources** box while searching.

In the RSA Content space on NetWitness Community, you can view the complete, up-to-date list of discontinued resources (Discontinued Content). For each resource, there is a description of why it was discontinued. Use these details to determine whether or not to remove a discontinued resource from your installation. .

The required permission to access this view is **Manage Live Resources**.

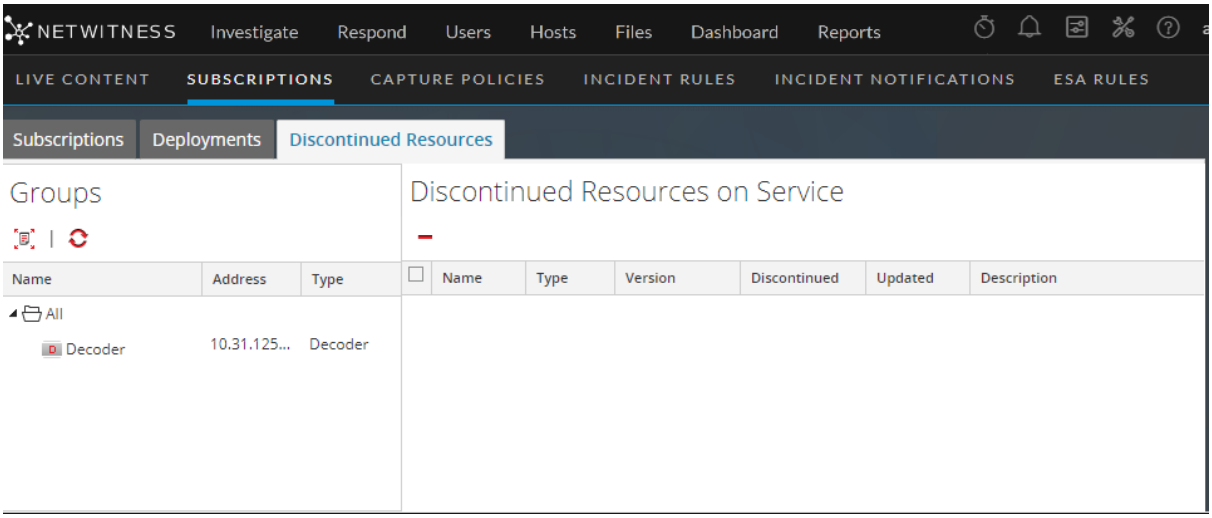
To access this view:

1. Go to **(Configure) > Subscriptions**.

   The **Subscriptions** tab is displayed.

2. Click the **Discontinued Resources** tab.

   The Discontinued Resources tab is displayed.



The Discontinued tab has two panels: Groups and Discontinued Resources on Service.

Groups Panel

The Groups panel is a static display of configured service groups that were created in the Admin Services view. Selecting a group in the Groups panel populates the Discontinued Resources panel with a list of discontinued resources which are deployed on the selected service or service group.

Feature	Description
	Click the button to scan the services for a discontinued resource.
	Displays the current status of the discontinued resources on a service. **Note:** The status of a service may change while the services are being scanned.
Name	Displays service group name. Clicking the plus sign displays a nested list of services in the group.
Address	Displays IP address of each service in the group.
Type	Displays type of service.

Discontinued Resources on Service Panel

The following table describes the features in the Discontinued Resources on Service panel.

Feature	Description
	Click the button to delete the selected resources from the service or service group.
Name	This is the name of the resource.
Type	This is the type of resource.
Version	Version of the discontinued resource.
Discontinued	Indicates the status of the discontinued resources for the subscribed resource. **Yes** - The resource is discontinued. **No** - The resource is not discontinued. **--** - The Live Server is not checked for the discontinued resources.
Updated	Displays date and time that the resource was last updated.
Description	Displays description of the resource.

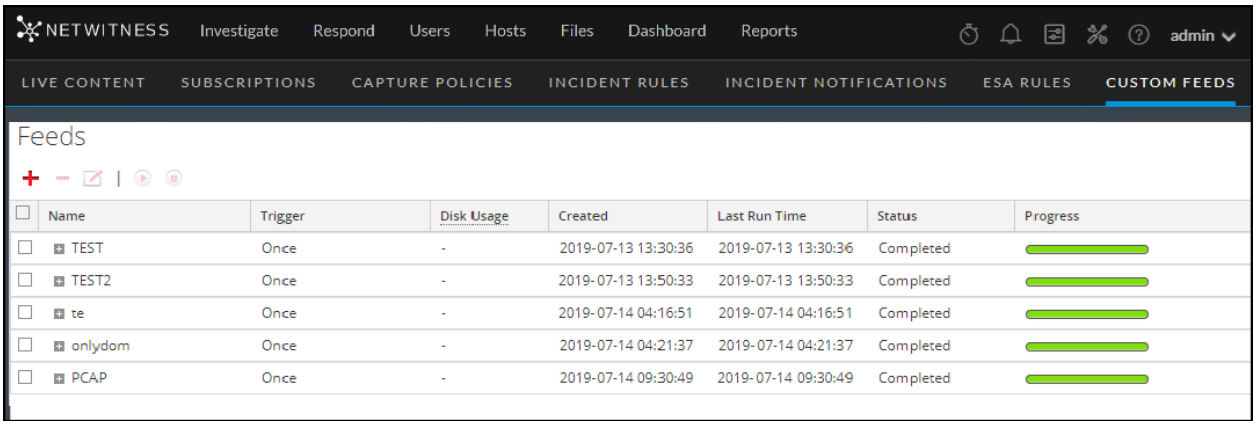
Live Feeds View

Use the Live Feeds View to:

- Create custom feeds.
- Create identity feeds.
- Edit feeds.

The required role to access this view is **Manage Devices**.

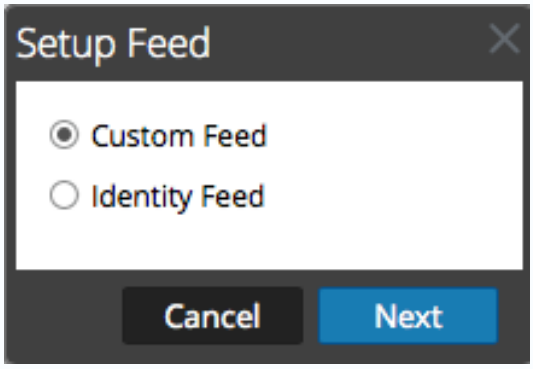
To access this view, navigate to **(Configure) > Custom Feeds**.
This is an example of the Feeds view.



The **Feeds** tab has a toolbar and a grid.

Toolbar

This table describes the options in the toolbar.

Feature	Description
	Initiates the creation of a custom or identify feed by displaying the **Setup Feed** dialog is displayed.  • Custom Feed opens the **Configure a Custom Feed** wizard. • Identity Feed opens the **Configure Identity Feeds** wizard.
	Deletes the feed that you selected.
	Opens the Configure Custom Feed or Configure Identity Feed wizard for the feed that you selected (see Editing a Feed.
	Start or resume data feed.
	Stop or pause data feed.

Feeds Grid

This table describes the columns in the grid.

Column	Description
	Selects a feed.
Name	Name of the feed. **Note:** You can now use special characters to define the name of the custom feed.
Trigger	Displays how often the feed runs which is determined by what you defined in **Feed Task Type** when the feed was created.
Created	Displays date and time when the feed was created.
Disk Usage	Displays the MongoDB storage size used by the TAXII feed.
Last Run Time	Displays date and time when the feed was last run.
Status	The status of the feed.
Progress	Progress bar.

Live Resource View

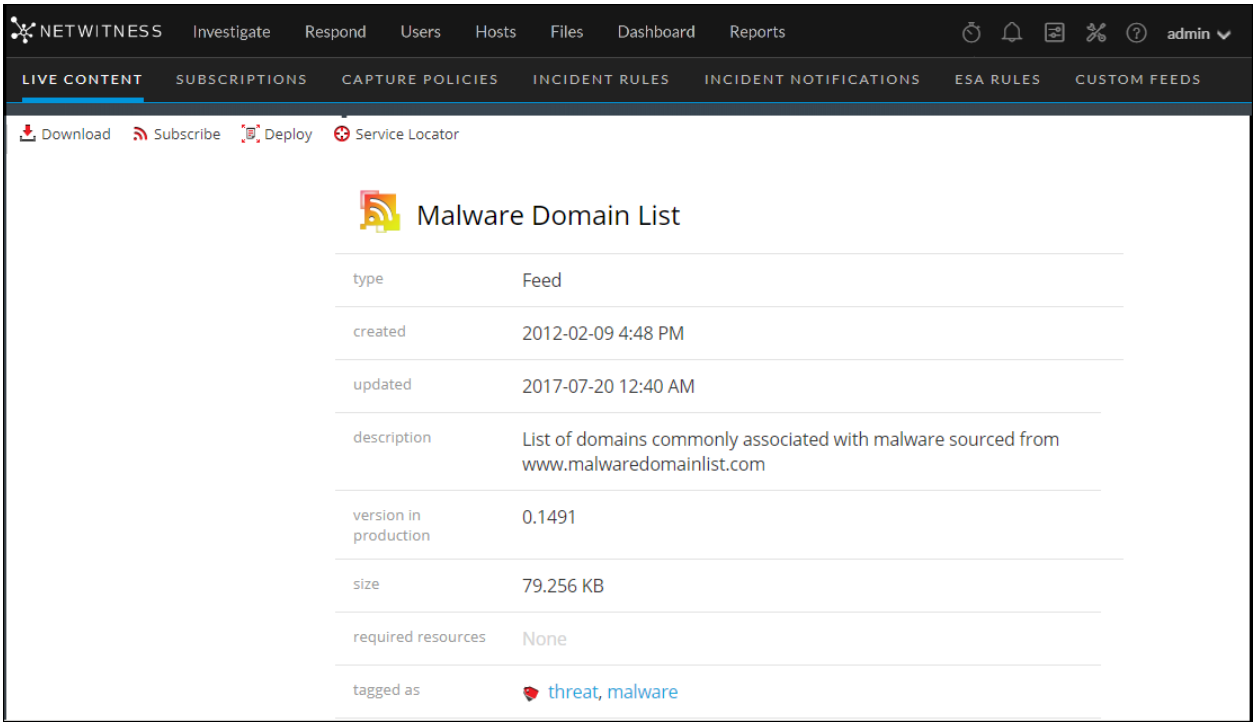
The Live Resource View shows a detailed view of a selected resource, and has the following options:

- Download the resource.
- Subscribe or unsubscribe the resource.
- Deploy the resource to services.
- Locate services on which the resource is deployed and remove the resource from services.

The required permission to access this view is View Live Resource Details.
To access this view, do one of the following:

1. Go to **(Configure) > LIVE CONTENT > Search Criteria**.
2. In the Live Search view, **Detailed Results**, click the resource type icon or the resource name.
3. In the Live Search view, **Grid Results**, double-click a resource or select a resource and click **Details**.

This is an example of the Resource view.



The Live Resource View has a detailed view of a single resource and a toolbar.

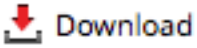
Resource Details

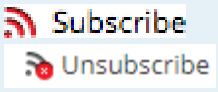
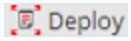
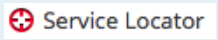
The following table describes the elements in the Resource Details section.

Feature	Description
Resource Type Icon	A graphic representation of the resource type, for example  .
Name	The name of the resource, for example, **fingerprint_office_lua**.
Type	The type of resource, for example, **RSA Lua Parser**.
Created	The date the resource was created, for example, **2013-09-15 02:16 PM**.
Updated	The date the resource was last updated, for example, **2013-09-15 02:16 PM**
Description	The description of the resource, for example, **Identifies Microsoft Office 95, 2007 Word, Excel, and PowerPoint documents**.
Version in production	The version of the resource, for example, **0.1**.
Size	The size of the resource, for example, **9.079 KB**.
Required Resources	A list of resources on which this resource depends, for example, **NetWitness Lua Library**. Clicking a resource replaces the currently displayed details with the details of the one you clicked.
Tagged as 	The tags that apply to the resource. In the example, the tags are **featured, informational**. Clicking a tag opens the Live Search View with the search narrowed to match resources with that tag.
Required Meta Keys	The meta keys  that apply to the resource. In the example, there are no meta keys required. Clicking a meta key opens the Live Search View with the search narrowed to match resources with that meta key.
Generates Meta Values	The meta values  that the resource generates. In the example, there are no meta values generated. Clicking a meta value opens the Live Search View with the search narrowed to match resources with that meta value.
Permissions	The permissions required for the resource.

Resource View Toolbar

This table describes the Live Resource view toolbar options.

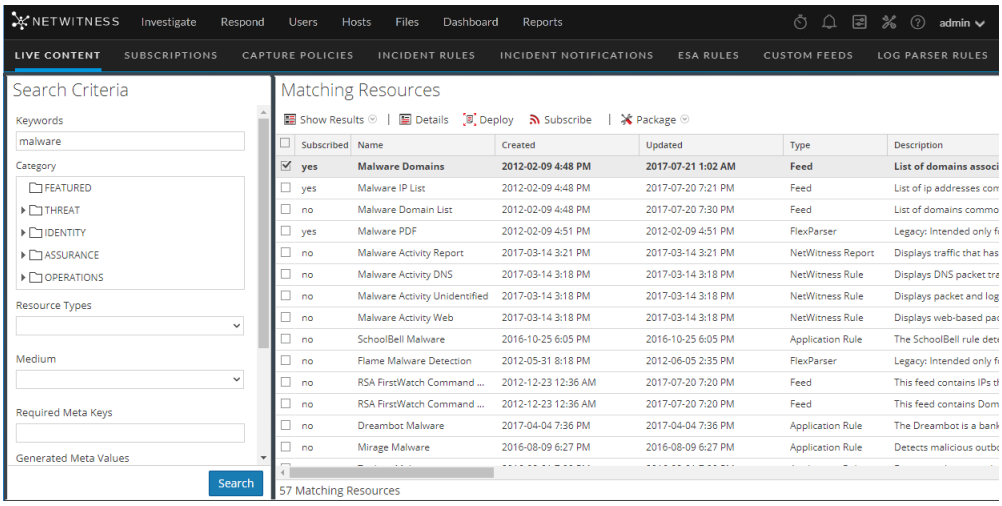
Feature	Icon	Description
Download		This option downloads the resource currently displayed in the Resource View.

Feature	Icon	Description
Subscribe or Unsubscribe		This option subscribes to or unsubscribes from the resource currently displayed in the Resource View. • Clicking **Subscribe** opens a dialog notifying that you are agreeing to receive notification when the selected resources are updated. You can cancel or click **OK**. • Clicking **Unsubscribe** asks for confirmation that you want to stop receiving notification when the selected resources are updated. You can then choose to cancel or you can click **Unsubscribe** or **Unsubscribe and Remove**, which also removes the resource from services on which it is deployed.
Deploy		This option provides a way to deploy the resource currently displayed in the Resource View. Clicking **Deploy** opens the Manual Resource Deployment dialog.
Service Locator		This option displays a list of services on which the currently displayed resource is deployed. You can remove the resource from all services or selected services.

Live Search View

The Live Search view provides the ability to browse the configured Live CMS for resources. Once matching resources are found, you can view details, subscribe to resources, and deploy resources to services and service groups.

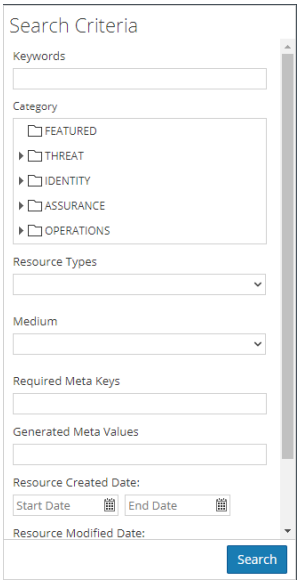
This is an example of the Search view.



The Live Search view has a panel for specifying search criteria and a panel that displays matching resources. The Search Criteria panel is collapsible to provide more width for viewing the Matching Resources panel.

Search Criteria Panel

This is an example of the Search Criteria panel.



The following table provides descriptions of the Search Criteria panel features.

Feature	Description
Keyword(s)	Enter a keyword or keywords to browse for resources that have the keyword in the resource name or the resource description. You can use wildcards when you enter a keyword.
Category	The categories mirror the hierarchical Investigation Model that NetWitness uses to organize resources. The purpose of the Investigation model is to deliver an accurate path to information security incident response. For more information, see the Investigation Model topic in the NetWitness Content space on NetWitness Community.
Resource Types	Select resources types from the drop-down list to filter resources by type of resource. Possible values are: • Advanced Analytics (Warehouse) • Application Rule • Bundle • Correlation Rule • Event Stream Analysis Rule • Feed • FlexParser • Investigation Column Group • Investigation Meta Group • Investigation Profile • Log Collector • Log Device • Lua Parser • Malware Rules • NetWitness List • NetWitness Report • NetWitness Rule • (Version 11.5 and later) Health and Wellness Dashboards • (Version 11.5 and later) Health and Wellness Monitors **Note:** Some rules that have been deployed to an earlier version of NetWitness may not deploy or execute on NetWitness 11.x. For more information, see the Troubleshooting Live Services.

Feature	Description
Medium	Select one or more mediums from the drop-down list to search for content based on the meta data source. Available values for medium are as follows: • **endpoint**: for NetWitness 11.3 and later): applied to content that uses meta derived from endpoint agent and endpoint server data • **log**: applied to content that uses meta derived from log data • **packet**: applied to content that uses meta derived from network packets • **log and packet**: applied to content that correlates meta derived across log and packet data
Tags	Select meta tags from the drop-down list to browse based on how the meta is tagged. For example, to browse resources for a Log Decoder, select the **netwitness for logs** tag. Alternatively, you can click a tag in the Matching Resources panel to insert that tag in this field.
Required Meta Key(s)	Enter a specific meta key; for example, **threat.source**. Alternatively, you can click a meta key in the Matching Resources panel to insert that tag in this field.
Generated Meta Value(s)	Enter a generated meta value; for example, **netwitness**. Alternatively, you can click a generated meta key in the Matching Resources panel to insert that tag in this field.
Research Created Date	Specify a date range during which resources were created. For example, to browse resources that were created between January 1 and January 4, you select January 1 as the start date and January 4 as the end date. You must enter dates in mm/dd/yyyy format or you click  and pick dates from a calendar.
Research Modified Date	Specify a date range during which resources were modified. For example, to browse resources that were modified between January 1 and January 4, you select January 1 as the start date and January 4 as the end date. You must enter dates in mm/dd/yyyy format or you click  and pick dates from a calendar.
Search	Click **Search** to send the search request to the Live server. More specific search criteria return matching resources more quickly.
Cancel	Click **Cancel** to cancel the search in progress.
Include Discontinued Resources	Check **Include Discontinued Resources** to include the discontinued resources in the search result. For an up-to-date list of resources that have been discontinued, see the Discontinued Content topic.

Matching Resources Panel

The Matching Resources panel displays search results based on the selections made in the Search Criteria panel. Results are initially displayed in a grid, but you can switch between two Show Results options: Detailed or Grid.

Detailed Results

In the detailed results, you can click a tag, meta key, or resource meta value to auto fill the Search Criteria panel and pivot the search results.

The following table describes the elements in the detailed results.

Feature	Description
Resource Type Icon	A graphic representation of the resource type. For example  .
Name	The name of the resource, for example, **Group Management**. **Note:** (Discontinued) is displayed next to the resource name if a resource is discontinued.
Type	The type of the resource, for example, **Rule**.
Updated	The date when the resource was last updated, for example, **2015-09-15 4:27 PM**.
Version	The version of the resource, for example, **0.1**.
Size	The size of the resource, for example, **153 B**.
Subscribed	Subscription status: • **yes**: This NetWitness instance is subscribed to this content resource. • **no**: This NetWitness instance has not subscribed to this content resource.
Description	The description of the resource, for example, **Compliance Rule-Group Management**.
Tags	The tags that apply to the resource. Clicking a tag narrows the search to resources with that tag. For example,  .
Meta Keys	The meta keys that apply to the resource. Clicking a meta key narrows the search to resources with that meta key. For example,  .
Resource Meta Values	The meta values generated by the resource. Clicking a meta value narrows the search to resources that generated the meta value. For example,  .

Grid Results

In the grid view, you can select one or more resources and use additional options in the toolbar to view the details of a single resource, subscribe to resources, and deploy resources.

The following table describes the elements in the grid results.

Feature	Description
Subscribed	Subscription status: • **yes**: This NetWitness instance is subscribed to this content resource. • **no**: This NetWitness instance has not subscribed to this content resource.
Name	The name of the resource, for example, **Group Management**. **Note:** The resource name is displayed in red color if it is discontinued.
Created	The date when the resource was created, for example, **2015-08-12 3:11 PM**.
Updated	The date when the resource was last updated, for example, **2015-09-15 4:27 PM**.

Feature	Description
Type	The type of the resource, for example, **Rule**.
Discontinued	The status of the discontinued resources: • **yes**: The resource that matches the search criteria is discontinued • **no**: The resource is not discontinued • **--**: The Live Server is not checked for the discontinued resources
Description	The description of the resource, for example, **Compliance Rule-Group Management**.
Toolbar	
Show Result	This menu offers two ways to view search results: **Detailed** and **Grid**.
Details	This option applies to a single selected resource. Clicking **Details** opens the selected resource in the Live Resource view.
Deploy	This option applies to one or more selected resources.
Subscribe	This option applies to one or more selected resources. Clicking **Subscribe** opens a dialog that asks for confirmation that you want to receive notification when the selected resources are updated.

Package	This menu offers two packaging functions for the selected resources:
	• **Create**: creates a **resourceBundle.zip** file that contains the selected resources and opens a dialog in which you can either:
	  • open the file, or
	  • save the file for subsequent deployment.
	• **Deploy**: opens the Deployment Wizard, in which you can choose a **resourceBundle.zip** file and deploy it.

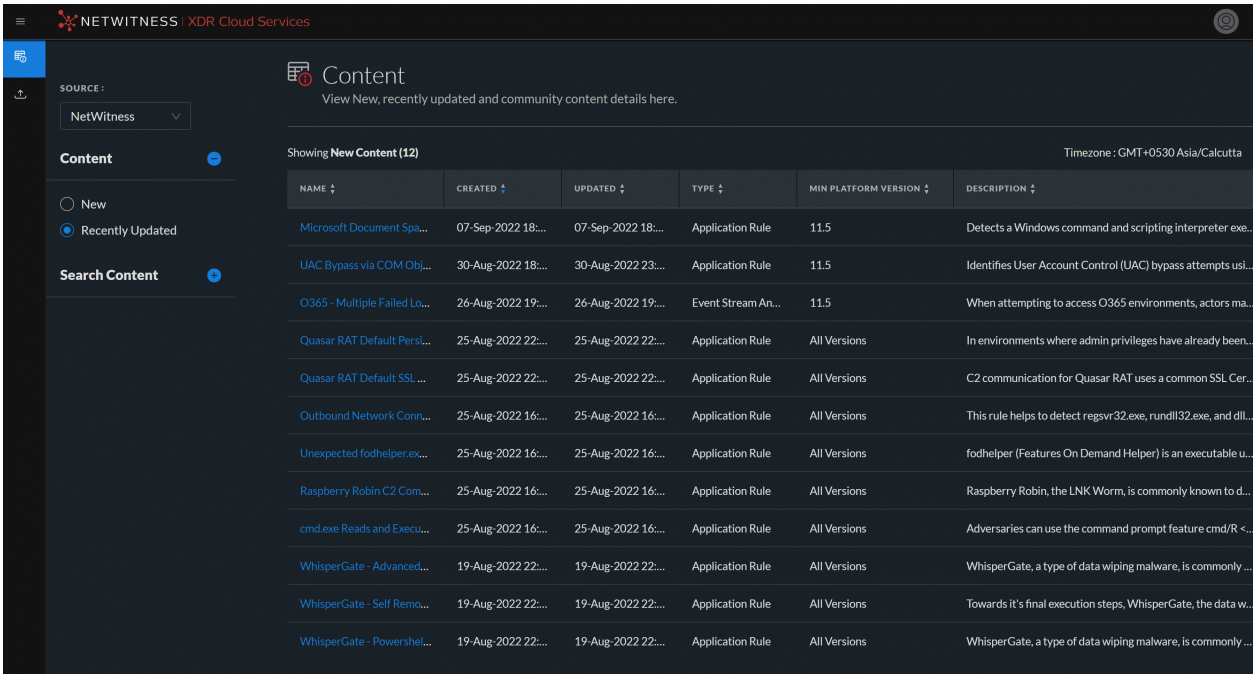
See Also

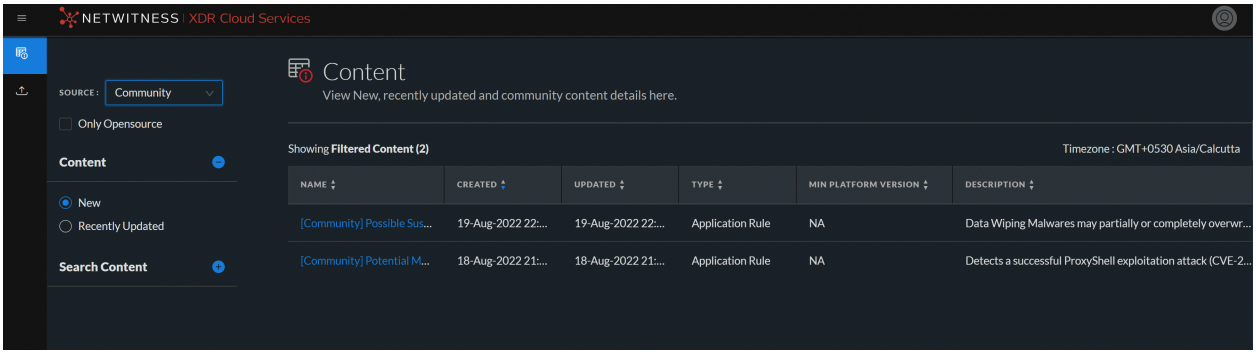
- For more information on Deployment (Deploy), see Find and Deploy Live Resources.
- For more information on Deploying a Package (Package), see the Resource Package Deployment Wizard,.

Live Search Content View

The Live Search Content view provides the ability to search the configured Live CMS for content. Once matching content are found, you can view the details, and download the content.

This is an example of the Search Content view.





The Live Search Content view has a panel for selecting the source and specifying search content. The matching content are displayed on the right panel.

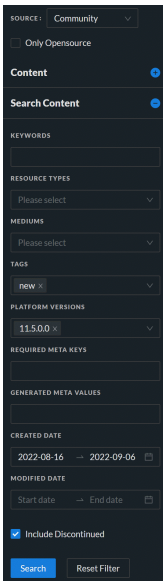
The following table provides descriptions of the Live Search Content panel features.

Feature	Description
NetWitness	Select NetWitness from the Source drop-down menu to search for in-built content that NetWitness Platform Live provides.
Community	Select Community from the Source drop-down menu to search for the content collected and retrieved from third party and open source communities.

Feature	Description
Only Opensource	Select the Only Opensource checkbox to retrieve the content from the open-source communities. **Note:** When the community is selected as the source, the Only Opensource option will be displayed under the Search Content Panel to select and search for open source-related content.
New	Select New to retrieve the content which is created in the last 21 days.
Recently Updated	Select Recently Updated to retrieve the content which is updated in the last 21 days.

Search Content Panel

This is an example of the Search Content panel.



The following table provides descriptions of the Search Content panel features.

Feature	Description
Keywords	Enter a keyword or keywords to browse for content that have the keyword in the resource name or the resource description. You can use wildcards when you enter a keyword.

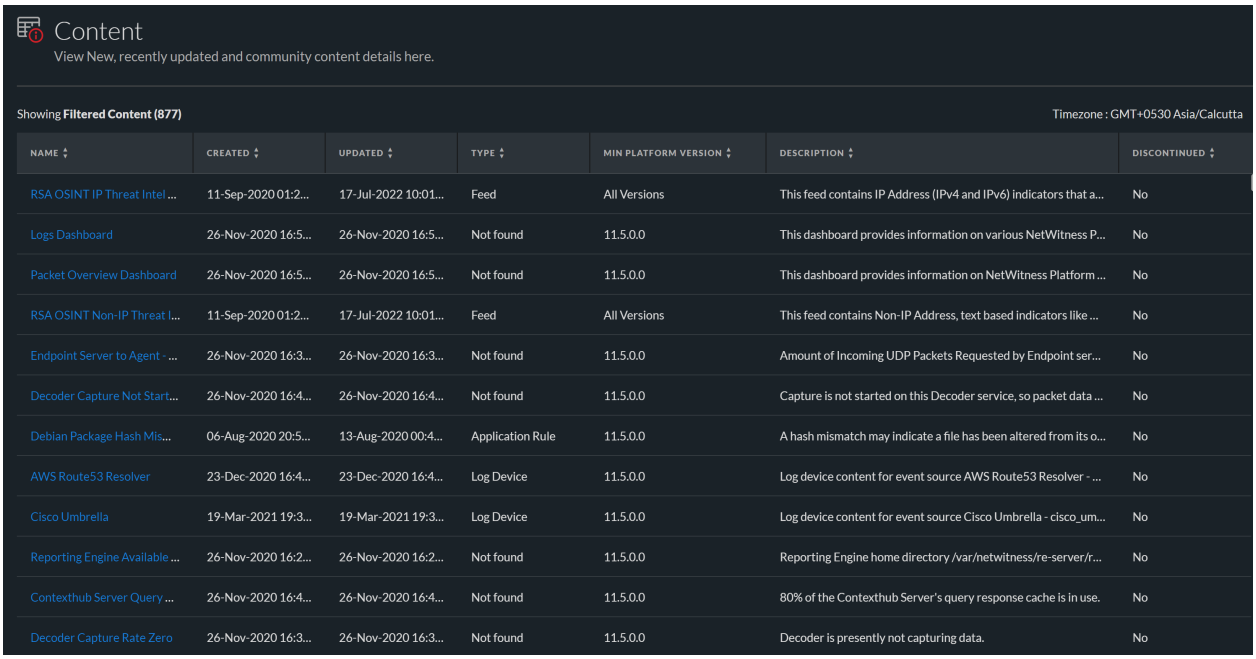
Feature	Description
Resource Types	Select resources types from the drop-down list to filter resources by type of resource. Possible values are: • Application Rule • Feed • Log Device • Correlation Rule • NetWitness Rule • NetWitness Report • Lua Parser • Log Collector • NetWitness List • Malware Rules • Event Stream Analysis Rule • Advanced Analytics (Warehouse) • Bundle • Health and Wellness Dashboards • Health and Wellness Monitors • Investigate Profile • Investigate Column Group • Investigate Meta Group
Mediums	Select one or more mediums from the drop-down list to search for content based on the meta data source. Available values for medium are as follows: • endpoint: for 11.3 and later): applied to content that uses meta derived from endpoint agent and endpoint server data • log: applied to content that uses meta derived from log data • packet: applied to content that uses meta derived from network packets • log and packet: applied to content that correlates meta derived across log and packet data.
Tags	Select meta tags from the drop-down list to browse based on how the meta is tagged. For example, to browse content for a Log Decoder, select the **netwitness for logs** tag.

Feature	Description
Platform Versions	Select one or more platform versions from the drop-down list to search for content based on the versions. For example, **11.5**.
Required Meta Keys	Enter a specific meta key. For example, **threat.source**.
Generated Meta Values	Enter a generated meta value. For example, **rsa-firstwatch**.
Created Date	Specify a date range during which content were created. For example, to browse content that were created between January 1 and January 4, you select January 1 as the start date and January 4 as the end date. You must enter dates in yyyy/mm/dd format or you click and pick dates from a calendar.
Modified Date	Specify a date range during which content were modified. For example, to browse content that were modified between January 1 and January 4, you select January 1 as the start date and January 4 as the end date. You must enter dates in yyyy/mm/dd format or you click and pick dates from a calendar.
Search	Click **Search** to send the search request to the Live server. More specific search criteria return matching content more quickly.
Reset Filter	Click **Reset Filter** to reset the existing search results and displays all the content on the right panel.
Include Discontinued	Check **Include Discontinued** to include the discontinued content in the search result. For an up-to-date list of content that have been discontinued, see the Discontinued Content topic.

Search Results Panel

The Search Results panel displays search results based on the selections made in the Search Content panel.

This is an example of the Search Results panel.



The following table describes the elements in the search results panel.

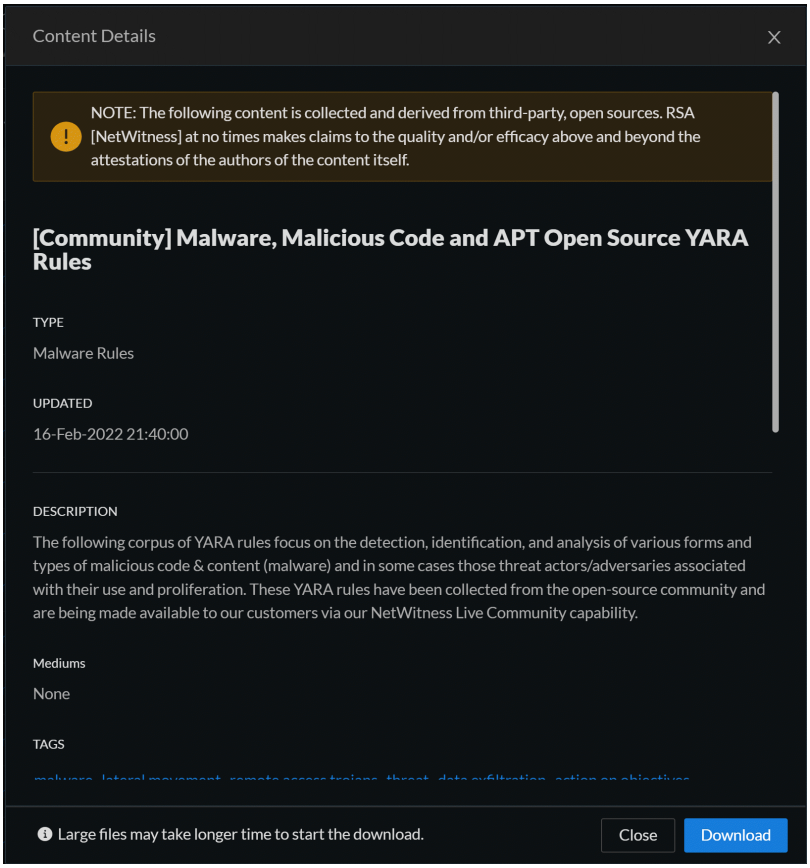
Feature	Description
Name	The name of the content. For example, **Log Parser Pack**.
Created	The date when the content was created. For example, **04-Aug-2017 15:19:06**.
Updated	The date when the content was last updated. For example, **29-Sep-2020 20:27:14**.
Type	The type of the content. For example, **Bundle**.
Min Platform Version	Platform version that the content supports. For example, 11.5 and higher. **Note:** Min Platform Version is not applicable for Community content.
Description	The description of the content. For example, **Contains all parser files and log collection files**.
Discontinued	The status of the discontinued content: • **Yes**: The content that matches the search criteria is discontinued • **No**: The content is not discontinued

Content Details Panel

In the Search Results panel, you can select any content titles to view the details in the pop-up window and download the content.

Note: NetWitness provides no assurance related to the quality and accuracy of the content provided by the third parties and open source communities.

This is an example of the Content Details panel.



The following table describes the elements in the Content Details section.

Feature	Description
Name	The name of the content. For example, **Log Parser Pack**.
Type	The type of the content. For example, **Bundle**.
Created	The date when the content was created. For example, **04-Aug-2017 15:19:06**.
Updated	The date when the content was last updated. For example, **29-Sep-2020 20:27:14**.
Description	The description of the content. For example, **Contains all parser files and log collection files**.
Version on Production	The version of the content. For example, **0.5**.
Size	The size of the content. For example, **14.96 KB**.
Required Resources	A list of resources on which this resource depends. For example, **NetWitness Lua Library**. Clicking a resource replaces the currently displayed details with the details of the one you clicked in the pop-up window.

Feature	Description
Tags	The tags that apply to the content. For example, **threat**. Clicking a tag opens the Live Search Content view with the search narrowed to match content with that tag.
Required Meta Keys	The meta keys that apply to the content. For example, **Threat Category**. Clicking a meta key opens the Live Search Content view with the search narrowed to match content with that meta key.
Generated Meta Values	The meta values that the content generates. For example, **rsa-firstwatch**. Clicking a meta value opens the Live Search Content view with the search narrowed to match content with that meta value.
Discontinued	The status of the discontinued content: • **Yes**: The content that matches the search criteria is discontinued • **No**: The content is not discontinued

Resource Package Deployment Wizard

If you have created a package of resources and saved it on a network drive, you can use the Resource Package Deployment Wizard to deploy the resources manually to a service or a service group without subscribing to the resources. NetWitness accepts packages in **.nwp** files or **.zip** files.

Deploying resources manually deploys them directly to the services without taking advantage of the powerful resource management capabilities of NetWitness.

If you want to receive notification and updates for updated resources and be able to easily remove resources from a service, you must subscribe to resources in the Live Search view and deploy the resources in the **Live Configure** view.

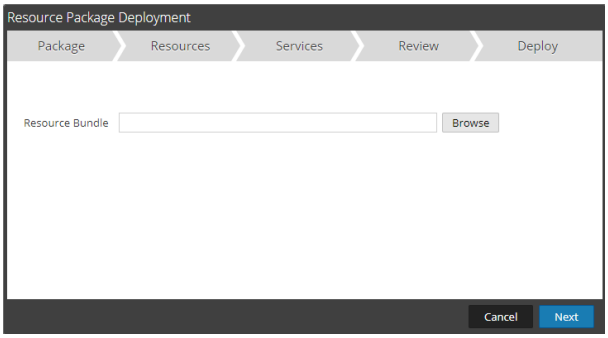
Note: Use NetWitness Live to create resource bundles; this is a different application that is not part of NetWitness. Selecting **Package** > **Create** in the **Live Search - Matching Resources** toolbar displays the Content Package Tool window. You can choose resources to include in a package and save the package as a NetWitness Package File.

The required permission to access this view is **Deploy Live Resources**.

To access this view:

1. Go to **(Configure) > Live Content**.
2. In the **Live Search - Matching Resources** toolbar, select Package > **Deploy**.

   The Resource Package Deployment wizard is displayed.



Features

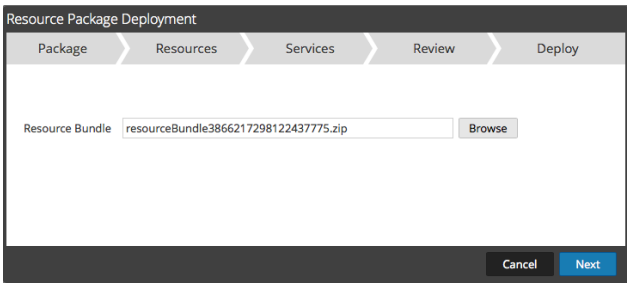
The Deployment Wizard has five tabs: **Package, Resources**, **Services**, **Review** and **Deploy**.
Use **Close** to exit before you complete the wizard.

When you complete the wizard, NetWitness returns to the Live Resources View.

Package Tab

You use this tab to select a resource bundle from your network in this page.

This is an example of the Package tab, with a resource bundle already selected.



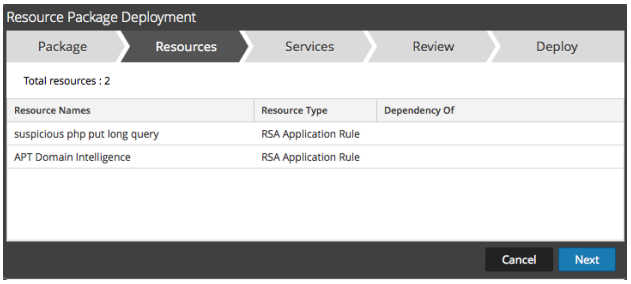
The following table describes the elements in the Package tab.

Column	Description
Resource Bundle	The input field to specify a resource bundle. You can type a path in this field or search using the Browse button.
	Command Buttons
Browse	This button opens a File Upload dialog in which you can browse the local file system and select a bundle.
Cancel	Cancels the deployment and closes the wizard.
Next	Displays the next tab of the wizard.

Resources Tab

This tab displays the resources contained in the bundle.

The following figure shows an example of the Resources tab.



The following table describes elements in the Resources tab.

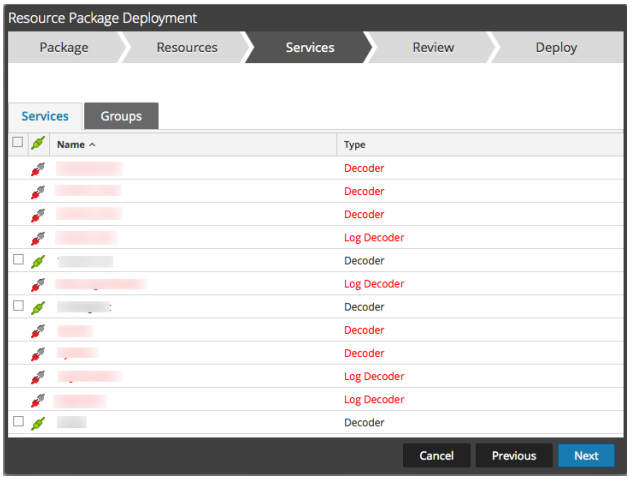
Column	Description
Resource Name	Displays the name of the resources in the bundle (for example, **NetWitness Lua Library**).
Resource Type	Displays the resource types for the resources in the bundle (for example, **RSA Lua Parser**)
Dependency Of	Displays Resources on which the selected resource depends (for example, **AIM lua**).

Services Tab

You select the services on which you want to deploy the resources in the bundle.

The Services tab has two tabs, **Services** and **Groups**. These provide a list of services and service groups that are configured in the **(Admin)** > **Services** view. The columns are a subset of the columns available in the Services view. You can select the services or the service groups on which you want to deploy the resources in the bundle.

This is an example of the Services tab.



The following table describes the elements in the Services tab.

Column	Description
	Services
	Selects services on which you want to deploy the content. You can select any combination of services and service groups.
Name	Displays the services in your environment on which you can deploy the content.
Host	Displays the name of the resource host.
Type	Displays the type of NetWitness service.
	Groups
	Selects service groups (if you have service groups defined in your environment).
Name	Displays the names of the service groups.

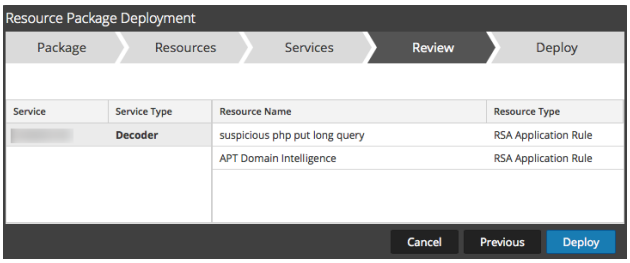
Review Tab

Displays the resources and services on which the resources will be deployed.

In this tab, you can do the following:

- Review the content and services before you deploy.
- Initiate the deployment of the resources.

The following figure shows an example of the Review tab.



The following table describes the elements in the Review tab.

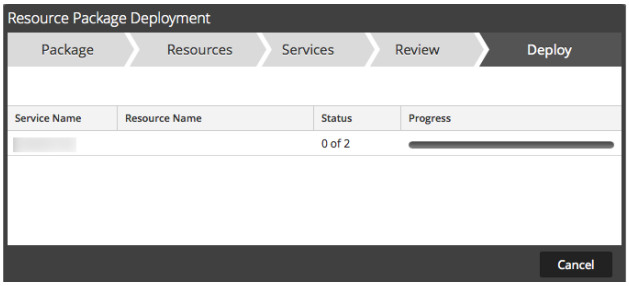
Column	Description
	Service Information
Service	Displays the services in your environment on which you can deploy the content.
Service Type	Displays the type of each NetWitness service (type of host or service).
	Resource Information
Resource Name	Displays the name of the resources you have selected (for example, **NetWitness Lua Library**).
Resource Type	Displays the resource types for the resources you have selected (for example, **RSA Lua Parser**).
Deploy	Initiates the deployment of the resources and displays the **Deploy** page (final page of the wizard).

Deploy Tab

This tab allows you to do the following:

- View the progress of the job
- Cancel the job

This is an example of the Deploy tab.



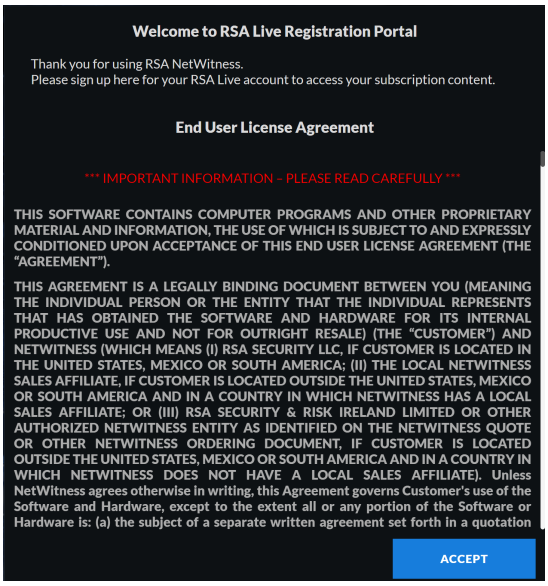
The following table describes the elements in the Deploy tab.

Feature	Description
Service Name	Name of the services to which resources are deployed.
Resource Name	Name of the resources.

Feature	Description
Status	Status of the manual deployment.
Progress	Progress of the manual deployment in a progress bar. When complete, the bar turns solid green.
Command Buttons	
Close	Closes the wizard.
Errors	Only displays if NetWitness encountered any errors. Click to display the errors.
Retry	Only displays if NetWitness encountered any errors. Click this button to try to deploy the resources again using the wizard.

NetWitness Live Registration Portal

The NetWitness Live Registration Portal is a self-service wizard in which customers can set up a Live account and change or reset the password. A Live account is required to get access to the feeds, parsers, rules, and other content in NetWitness Live library. To access the portal, go to the following URL: https://live.netwitness.com/registration.





Click **Sign Up For Live**. The License Agreement page is displayed, once you agree to the Terms and Conditions, click **Accept**: the fields for setting up an account are displayed. These include Contact Information, and License ID.

The following table lists the contact information section fields and its descriptions:

Parameter	Description
First Name	Your first name.
Last Name	Your last name.
Company	The name of your company.
Email	The email address where you want to receive notifications related to the Live account.

Parameter	Description
License ID	This is the License ID on the **(missing or bad snippet) > System > Info** page. **Caution:** The license ID on the NetWitness must be valid and must be registered on the Flexera Server. If not, contact NetWitness Customer Support.
Username	The username used to sign in to Cloud Services Live account. The username must contain a minimum of four characters and a maximum of 16 characters.
Password	The password for the Cloud Services Live account. The password must contain minimum of eight characters and the maximum length is 24, with at least one uppercase, one lowercase, one number, and one special character.

NetWitness Feedback and Data Sharing

The Live Feedback Activity Log enables you to download the usage data required for Live Feedback. After you download the Live Feedback data, you can then upload it to share with NetWitness.

The settings for these features are available in **(Admin) > System > Live Services** view, in the Additional Live Services section.

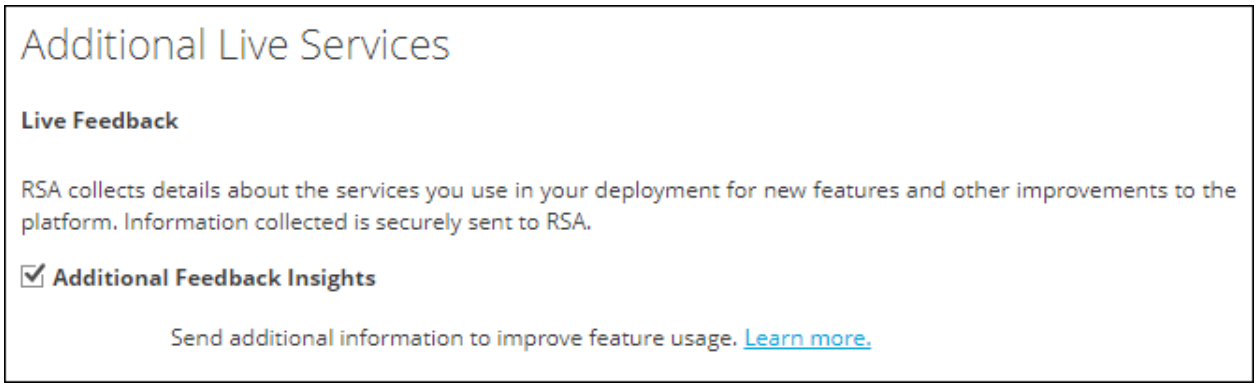
Additional Live Services

Participation in the Additional Live Services is configured in the **(Admin) > System > Live Services** view.

Live Feedback

Note: For NetWitness 11.4.1 and later, this section in the UI has been removed. As of 11.4.1, NetWitness has created the Customer Experience Improvement Program. For details, see "Configure the Customer Experience Improvement Program" in the *NetWitness System Configuration Guide*.

Live Feedback is intended to help improve NetWitness.



Once you set up and configure a Live account, usage data is automatically shared with NetWitness and is protected in accordance with the applicable license agreement. All such data shall be anonymized and shall not have any Personally Identifiable Information.

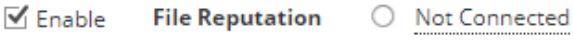
Before data is sent to NetWitness, all Personally Identifiable Information is removed. Thus, only anonymous usage data gets transferred to NetWitness.

For more information, see the "Live Feedback Overview" topic in the *System Configuration Guide*.

File Reputation

File Reputation service provides instant access to the latest signatures using the RSA Live feed so data is more relevant, with fewer false positives. With this service, users always have reliable data about the reputation of files in their NetWitness Endpoint system. In addition to the whitelisting service, it provides blacklisting information as well.

File Reputation



This option is used to view reputation status of files. The File Hash information from NetWitness Platform is sent to RSA Live to get the reputation status. Reputation status is leveraged by analysts during investigation of files.Learn more.

Troubleshooting Live Services

This section provides troubleshooting instructions for issues faced when using the Live Services module in NetWitness.

Some Rules Are Invalid for Version 11.x

The rules "NetWitness Incident Management - Alert Details" and "NetWitness Incident Management - Incident Summary" are not valid for NetWitness version 11.x. Do not deploy these rules to an 11.x system.

Note: Rules are updated frequently, and the documentation for them is available in the Content space on NetWitness Community. For the latest information on Rules, see RSA NetWitness Rules.

OutOfMemoryError on Context Hub Server

You may encounter an OutOfMemoryError on Context Hub server, and the service becomes unresponsive.

If there are any TAXII feeds configured, Health and Wellness raises alerts when the available heap memory of Context Hub server is critically low. If the status of Context Hub server is Unhealthy because of low memory, perform the following steps:

1. Make sure that the feeds **Start Date** is within 180 days.
2. Check if any TAXII feed is consuming too much disk space. A TAXII feed can consume maximum of 300 MB. If it consumes more disk space, you must reduce the value in the **Remove STIX data older than** field under **Advanced Options** in the **Custom Feed Creation Wizard** when you edit a TAXII feeds.

   Note: If the issue still persists, you must execute step 3.

3. To decrease the number of parallel threads available for processing STIX:

   a. Go to  **(Admin) > Services > Context Hub service > View > Explore**.

   b. In the tree panel, navigate to **enrichment/stix/ config**.

   c. In the right panel, set the **stix-query-scheduler-pool-size** field value to 2. By default the value is 5. This setting controls how many number of threads are allowed to process queries for STIX data at the same time.

   d. Set the **taxii-poll-scheduler-pool-size** field value to 2. By default the value is 5. This setting controls how many number of threads are allowed to poll TAXII servers at the same time.

   e. Restart the Context Hub server.

Troubleshooting Live Connect Threat Data Sharing

This section discusses troubleshooting Live Connect Threat Data Sharing.

Query Log Retrieval Sample

To retrieve a sample of threat intelligence data sent to Live Connect, you must construct a URL by setting the following parameters:

- **sendReport**: value is **true** or **false**: true to send this report to the Live Connect server. False to just create the report for viewing. The value defaults to false.
- **hashValues**: value is **true** or **false**: true to hash the values as md5/sha256. False to show values in clear text – should use only for manual viewing. Defaults to false.
- **startDate / endDate**: Dates for time boundaries for log entries. Format: YYYY-MM-DD HH:mm:ss

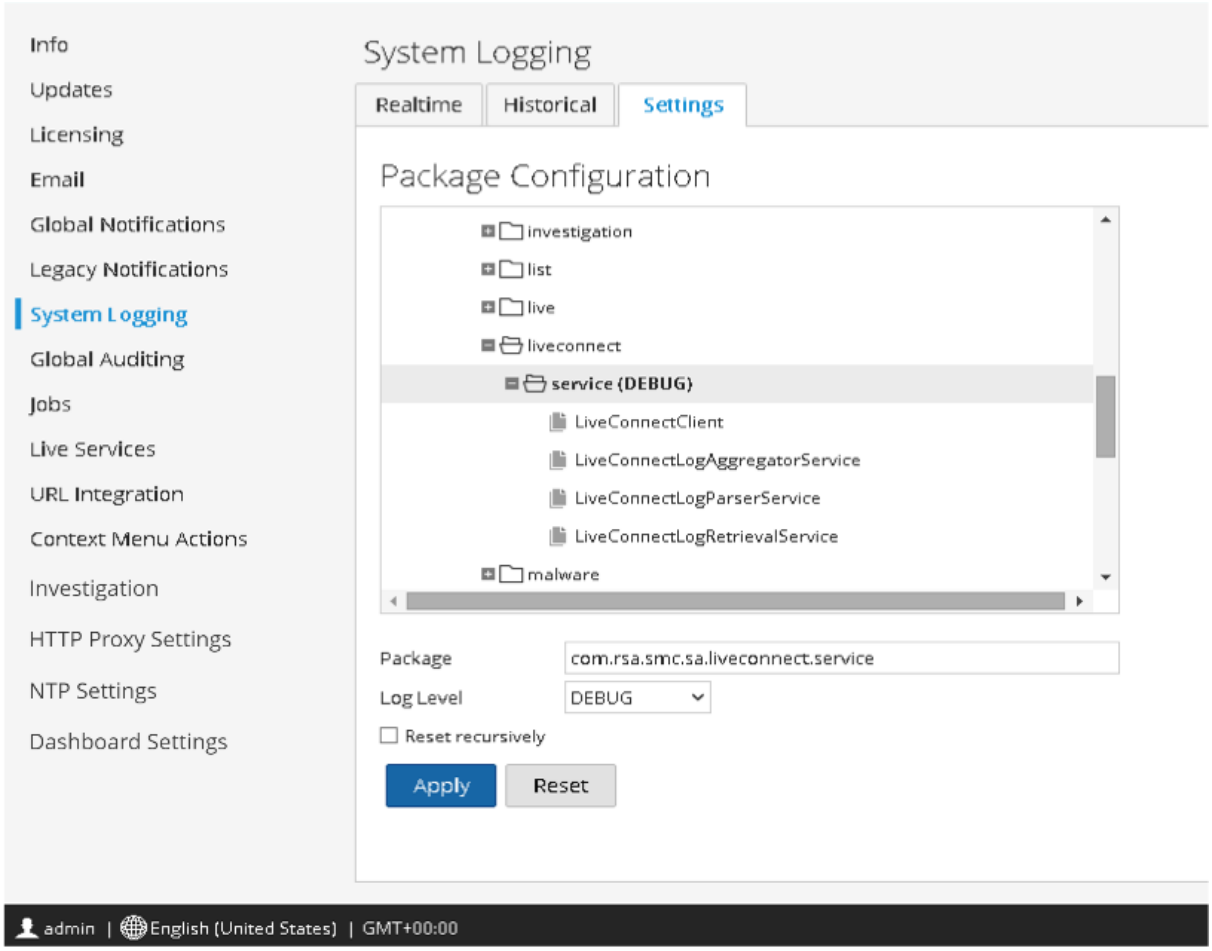
The following is an example of the URL used to retrieve query logs:

```
https://<server>/admin/liveconnect/force_aggregation?startDate=2016-01-18%2000:00:00&endDate=2016-01-19%2010:10:00&sendReport=false&hashValues=true
```

System Logging: Debug

To access debug information:

1. Go to **(missing or bad snippet) > System > System Logging**.
2. Select the **Settings** tab.
3. In the Package Configuration section, select **com > netwitness > platform > server > liveconnect > service (DEBUG)**.


Info
Updates
Licensing
Email
Global Notifications
Legacy Notifications
System Logging
Global Auditing
Jobs
Live Services
URL Integration
Context Menu Actions
Investigation
HTTP Proxy Settings
NTP Settings
Dashboard Settings
System Logging
Realtime
Historical
Settings
Package Configuration
investigation
list
live
liveconnect
service (DEBUG)
LiveConnectClient
LiveConnectLogAggregatorService
LiveConnectLogParserService
LiveConnectLogRetrievalService
malware
Package
com.rsa.smc.sa.liveconnect.service
Log Level
DEBUG
Reset recursively
Apply
Reset
admin | English (United States) | GMT+00:00

Policy-based Centralized Content Management

This chapter covers different topics that lets you configure Policy-based Centralized Content Management.

- About Policy-based Centralized Content Management
- About Content Library
  - Migrate Content from Core Services to Content Library
  - Import Content to Content Library
  - Create an Application Rule
  - Edit Application Rule
  - Delete Application Rule
  - View Application Rule Details
  - Create a Network Rule
  - Edit Network Rule
  - Delete Network Rule
  - View Network Rule Details
- About Groups
  - Create a Group
  - View a Group
  - Delete a Group
  - Edit a Group
- About Policies
  - Create and Publish Policies
  - Clone a Policy
  - Delete a Policy
  - Edit a Policy
  - View a Policy
  - Enable Content for a Policy
  - Disable Content for a Policy
- References

- Content Library Tab
- Groups Tab
- Policies Tab

About Policy-based Centralized Content Management

Legacy content management involves deploying and managing content in multiple places in the UI.

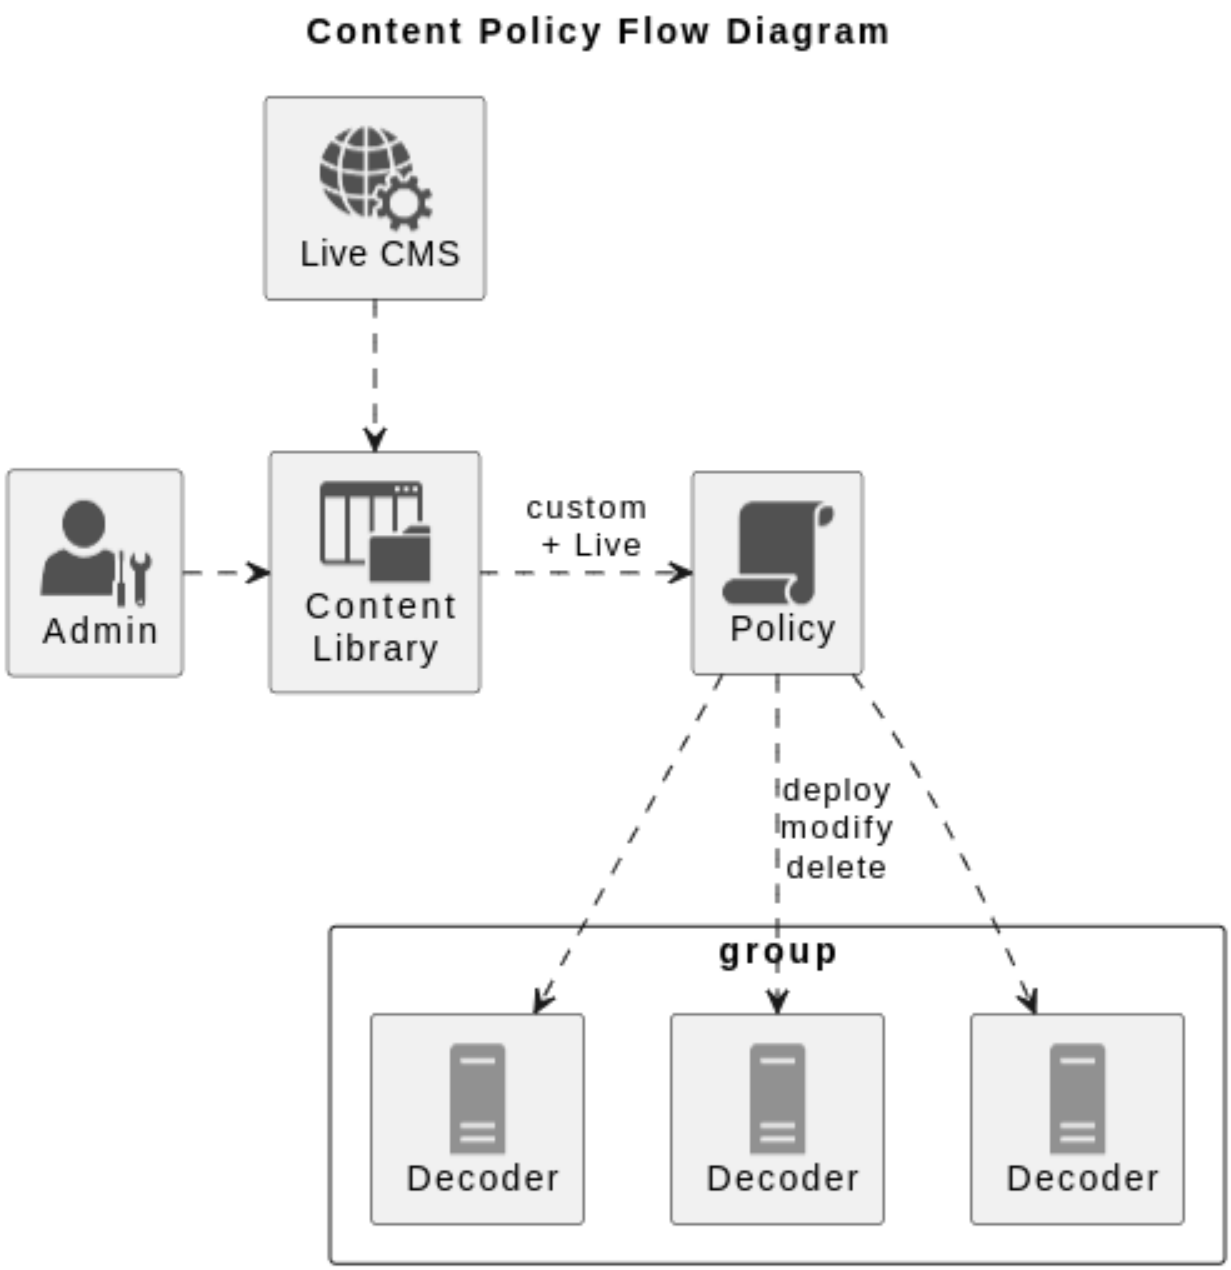
- **Live Content UI**: Located under the Configuration interface, this allows a “push” deployment of Live content to one or more services, but does not provide any management of content once it is deployed
- **Service Config UI**: Located under **Admin > Services > View Config**, this UI enables you to view, edit or delete content on individual services.

Policy-based Centralized Content Management is a unified approach to find, deploy, and manage content through the entire life cycle based on policies that can be assigned to groups of devices. It is a single location to view, modify and manage the content deployed across all services in the environment.

This approach consists of three elements:

- Groups: A collection of NetWitness services (such as Decoders, Log Decoders, etc.) to assign and manage content.
- Content Policies: A container of content and subscription settings used to assign and manage content within a Group.
- Content Library: A local repository of content which resides on the Admin Server and is used to assign content to policies. This includes both Live and Custom content.

The Content Library contains Live content (synchronized with the Live CMS) and any custom content you create or import. To deploy, remove or manage content on your services, content is assigned from the Content Library to a Content Policy. Once that content policy is assigned to a group and Published, the content changes are put into effect on the services within the group.



Benefits of Policy-based Centralized Content Management:

- Add or remove content without repeating the process on each service.
- Add content from RSA Live or add your custom content into a single content repository. You can add content from this repository to a policy.
- Add a new service to an existing group to automatically deploy all necessary content.
- One-click management of subscriptions and automatic updates

- Provides highly responsive and updated UI for browsing RSA Live content that can help you with the following:
  - View Live content along with your content policies and click  to add content from Live.
  - Seamlessly view Live content along with your custom content.

IMPORTANT: It is recommended not to use the Centralized Content Management and Service Config page or Live Content page simultaneously for managing the content. Using the Service Config UI to add or modify content can cause the content to become out-of-sync with the Content Policy.

Migrate Content from Core Services to Content Library

The customers who want to use Centralized Content Management, and if their content is already deployed, a migration process is required.

Note: Existing Live content does not need to be exported or imported. All Live content will be available in the Content Library and will only need to be added to one or more policies and published as needed.

This process includes the following steps:

- Make sure that the necessary Live content has been selected and applied to one or more Policies.
- Enable subscriptions for Live content as desired.
- Export any custom content, including Application Rules, Network Rules, Lua Parsers and Log Parsers.
- Import custom content into the Content Library.
- Apply custom content to one or more Policies.
- Create Groups to which Policies will be assigned.
- Publish Policies to their assigned Groups.

Warning: Initially, when a Policy is published to a Group , all the content which are not included in the policy will be removed from the services in that Group.

To migrate Application Rules or Network Rules

1. Go to  **(Admin) > Services**.
2. Go to Config view of the service where application rule or network rule is deployed.
3. Click either the **Application Rule** or the **Network Rule** tab.

   Note: The **Network Rule** tab is only available for **Network Decoder** services.
4. Select the content to migrate.
5. Click **Export** to export the selected content or click **All** to export all the content.

6. Type a file name which contains exported content and import the content to Content Library.

   For details on importing content to content library, view Import Content to Content Library topic.

The following table lists the supported file types and file extensions for Application Rules and Network Rules:

Content	Supported File Types	Supported File Extensions
Application Rules	**.NWR**	NA
Network Rules	**.NWR**	NA

To migrate Feeds, LUA Parsers, or Log Devices

The content file locations are as given below:

- Feeds content file location: /etc/netwitness/ng/feeds
- Lua Parsers content file location: /etc/netwitness/ng/parsers
- Log Devices content file location: /etc/netwitness/ng/envision/etc/devices

You can upload the files which are copied locally from these locations and import these files to Content Library.

For details on importing content to content library, view Import Content to Content Library topic.

The following table lists the supported file types and file extensions for Log Devices, LUA Parsers and Feeds:

Content	Supported File Types	Supported File Extensions
Feeds	**.zip**	**.feed** and **.token**
Log Devices	**.envision**	NA
LUA Parsers	**.zip**	**.luax**, **.lua** and **.flextoken**

Note: Any imported content will be treated as custom content. If imported content has the same name as existing Live content, then it must be renamed upon import. Custom content with the same name can be overwritten.

To create .envision files

1. Keep all the Log Devices in a root folder in your local drive. For example, "logDevices".
2. From the command prompt, run the python script specified in the NetWitness Community portal with input argument as the path of the above folder.

   Note: The command to run the python script is "python3 pythonscriptname.py inputArg".

3. Once you run the script, a new zip named "nw_content_logDevices.zip" is created. This zip file will contain all the envision files.

About Content Library

This section contains:

- Migrate Content from Core Services to Content Library
- Import Content to Content Library
- Create an Application Rule
- Edit Application Rule
- Delete Application Rule
- View Application Rule Details
- Create a Network Rule
- Edit Network Rule
- Delete Network Rule
- View Network Rule Details

Import Content to Content Library

Before the custom content can be used in policies, it must be imported to the Content Library.

To view the list of supported file types and file extensions for different content types, refer Migrate Content from Core Services to Content Library topic.

To import content to Content Library

1. Go to  **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.
4. Depending upon the type of content to be imported, click the following tabs:
   - **Application Rule**
   - **Network Rule**
   - **LUA Parser**
   - **Feeds**
   - **Log Devices**

   Note: Log Devices content should be converted to .envision files before importing.
5. In the respective content panel, click **Import**.
6. In the **Import** panel, click or drag the file to upload.

7. Click **Overwrite** to overwrite content. This is applicable only in case of overriding an already imported content.

   Note: You can overwrite the content if the content name is the same as the custom content. However, overwriting is not supported if the content name is the same as existing content of the same type from the live server.

8. Select the medium types.
9. Click **Import** to complete the import process.

Create an Application Rule

This topic describes the steps to create an application rule.

To create a new Application Rule

1. Go to  **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.

   The available rules are displayed.
4. In the application rule panel, click **+ Create Rule** to add an application rule.
5. In the **New Create Rule** panel, do the following:
   - Enter a unique rule name. If the name of that application rule is the same as an existing rule, an error message is displayed.
   - Enter the condition for the rule.
   - Select the medium to be applied for the rule.
   - Select the session data to be applied for the rule.
   - Select the session options to be applied for the rule.
   - Enter the meta value for the alert on. This is a mandatory field.
   - Click **Save** to save the new application rule.

Edit Application Rule

When you edit the application rule, follow these guidelines:

- You can only edit the custom rules.
- The rule name cannot be edited if the custom rule is assigned to a policy.
- If the custom rule assigned to a policy is edited, then the customer must republish the policy for the changes to take effect in the service.

To edit an Application Rule

1. Go to  **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.
4. Select an application rule to edit.
5. Click **Edit Rule** to edit the application rule.

Delete Application Rule

When you delete the application rule, follow these guidelines:

- You can delete only the custom application rules.
- You cannot delete the application rule if it is associated to a policy. You should first disassociate the application rule from the policy and then delete it.

To delete an Application Rule

1. Go to  **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.
4. Select an application rule to delete.
5. Click **Delete** to permanently delete the selected application rule.

View Application Rule Details

This topic describes the steps to view the application rule details.

To view Application Rule details

1. Go to  **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.

   The list of application rules is displayed.
4. Click a row to view details about the selected application rule in the right panel.

   The various details of the application rule are displayed.

Create a Network Rule

This topic describes the steps to create a network rule.

To create a Network Rule

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.
4. Click the **Network Rule** tab.
5. In the network rule panel, click **+ Create Rule** to add a network rule.
6. In the **New Create Rule** panel, do the following:
   - Enter a unique rule name. If the name of that network rule is the same as an existing rule, an error message is displayed.
   - Enter the condition for the rule.

     Note: The medium is selected as **Packet** by default, and it cannot be modified.
   - Select the session data to be applied for the rule.
   - Select the session options to be applied for the rule.
   - Click **Cancel** to cancel the operation.
   - Click **Reset** to reset the data.
   - Click **Save** to save the new network rule.

Edit Network Rule

When you edit the network rule, follow these guidelines:

- You can only edit the custom rules.
- The rule name cannot be edited if the custom rule is assigned to a policy.
- If the custom rule assigned to a policy is edited, then you must republish the policy for the changes to take effect in the service.

To edit a Network Rule

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.
4. Click the **Network Rule** tab.
5. Select the network rule to edit.
6. Click **Edit Rule** to edit the network rule.

Delete Network Rule

When you delete the network rule, follow these guidelines:

- You can delete only the custom network rules.
- You cannot delete the network rule if it is associated to a policy. You should first disassociate the network rule from the policy and then delete it.

To delete a Network Rule

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.
4. Click **Network Rule** tab.
5. Select a network rule to delete.
6. Click **Delete** to permanently delete the selected network rule.

View Network Rule Details

This topic describes the steps to view the network rule details.

To view Network Rule details

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Content Library**.
4. Click the **Network Rule** tab.
5. The list of network rules is displayed.
6. Click a row to view details about the selected network rule in the right panel.

   The various details of the network rule are displayed.

About Groups

This section contains:

- Create a Group
- View a Group
- Delete a Group
- Edit a Group

Create a Group

You can create a group with one or more services and assign one policy to it. Groups may be created without any assigned policy; however, a policy must be assigned to a group and Published in order for any content changes to take effect.

To create a Group

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Groups**.
4. In the tool bar, click **+ Create New**.
5. In the **New Group** panel, do the following:
   - Enter the name of the group.
   - Enter the description for the group.
6. Click **Next**.
7. In the **Define Group**, click + to assign services to the group.

   Note: A service is disabled if it is assigned to another group.
8. Click **Next**.
9. In the **Assign Policies**, click + to assign policies to a group. You can assign only one policy to any particular group.
10. Do any one of the following:
    - Click **Save and Publish** to save and publish the settings.
      - To publish all the content, click **Publish All**.
      - To publish only the content that is not published on the service, click **Publish Changes**.
      - To cancel the publish content dialog, click **Cancel**.
    - Click **Save and Close** to save the settings.

View a Group

This topic describes the steps to view the properties of Group.

To view the properties of the selected Group

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. Click **Groups**. The available groups are displayed.
4. Click a row to view details about the selected group in the right panel.

Delete a Group

You can delete one or more groups. Once the group is deleted, all services will be removed from the group and all the policy content will be deleted from the services.

To delete a Group

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. Click **Groups**. The available groups are displayed.
4. Select one or more Groups and click **Delete**.

   The confirmation message is displayed.

Edit a Group

You can edit the properties of the group at any point in time. The status of the updated group is unpublished if you change the service or policies in a group. If you just change the group name and description, then the status remains published (if it is already published).

To edit the selected Group

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Group**. The available groups are displayed.
4. Select a group to edit and click **Edit**.
5. Make the required changes in the group.
6. Do any one of the following:
   - Click **Save and Publish** to save and publish the policy.
   - The policy will be listed under the Unpublished category.
   - Click **Save and Close** to save the settings and return to the Policies view.

Note: All the content within a service will be deleted when a service is removed from a group.

About Policies

This section contains:

- Create and Publish Policies
- Clone a Policy
- Delete a Policy
- Edit a Policy
- View a Policy
- Enable Content for a Policy
- Disable Content for a Policy

Create and Publish Policies

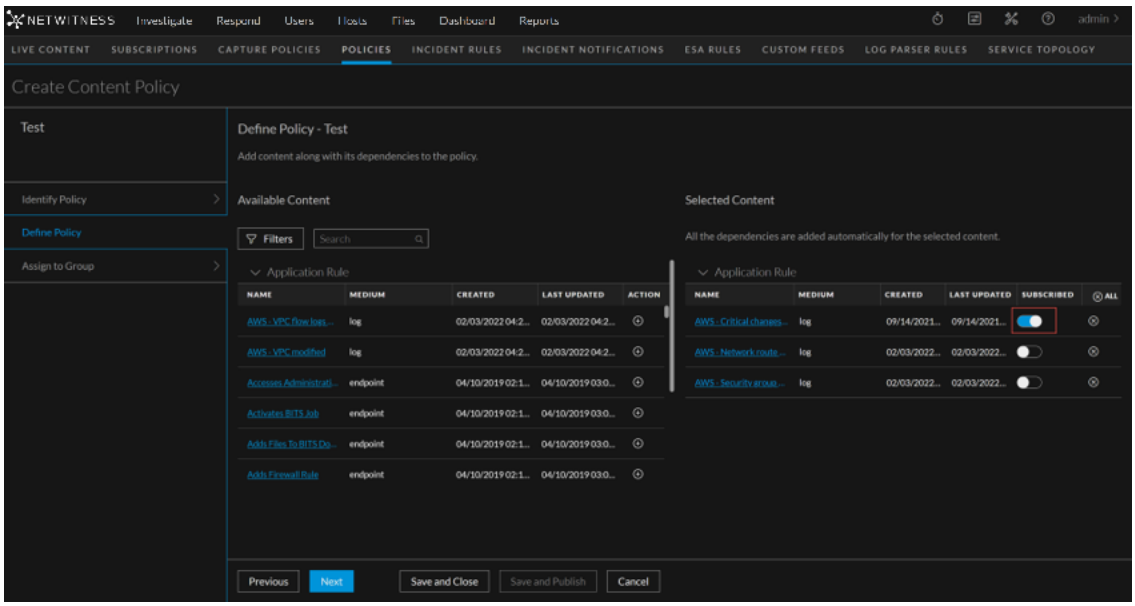
You can create a policy and assign it to one or more groups.

To create a Policy

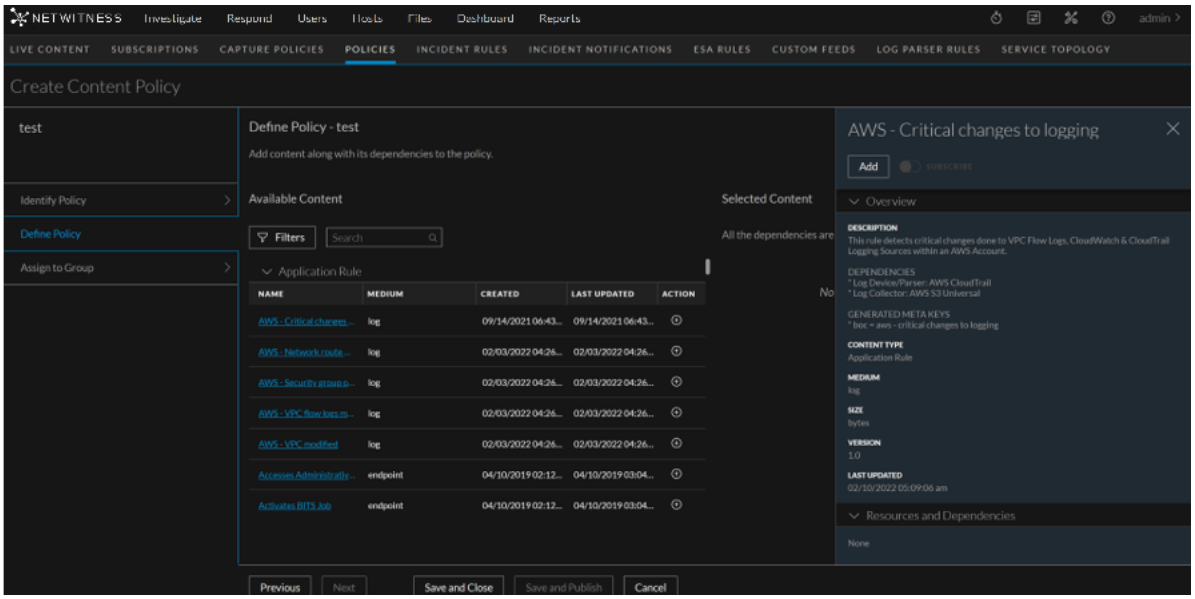
1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. Click **Policies**.

   The available policies are displayed.
4. Click **+ Create New** to add a new policy.
5. In the **New Policy** panel, do the following:
   - Enter a unique policy name.
   - Enter a description for the policy.
6. Click **Next**.
7. In the **Available Content**, select the content type and click + to add the content to the policy. After you add the content, you can enable subscription (if required) by clicking subscribed toggle. Once the content is subscribed the updates are pushed automatically.

Note: Subscription is not allowed for custom content.



Note:
- All the dependencies are added automatically for the selected content. You can click on the content name highlighted in blue and look for details such as content description, content type, resources and dependencies and so on. You can also add and subscribe the resource from the details view.



8. In the Group List, click + to assign groups to the policy.

   Note: A group is disabled if another policy of the same type is already assigned to this group.

9. Do any one of the following:

   - Click **Save and Publish** to save and publish the settings.
     - To publish all the content, click **Publish All**.
     - To publish only the content that is not published on the service, click **Publish Changes**.
     - To cancel the publish content dialog, click **Cancel**.

       Note: While publishing the first policy to a service, all previous content except custom feeds, will be deleted. Ensure that all custom content are migrated to content library before publishing the first policy.

   - Click **Save and Close** to save the settings.

   Note: You can also publish a policy from **Policy Details** screen. For more information on publishing a policy from **Policy Details** screen, refer View a Policy feature.

IMPORTANT:
When first applying a Content Policy to manage content, the existing Live and custom content on the service (excluding Custom Feeds) will be removed and replaced with the Policy content. You should compare the existing service content with the policy before applying to ensure required content is added to the policy. Endpoint risk scoring requires certain application rules. Refer Endpoint Risk Scoring Rules to view the list of these application rules.

Clone a Policy

When you clone a policy, all the content from the old policy is copied to the new policy. The cloned policy can be assigned to a new group. You can clone only one policy at a time.

To clone a Policy

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. Click **Policies**. The available policies are displayed.
4. Select a policy to clone and in the More actions drop-down list in the tool bar, click **Clone**.

   The policy is cloned successfully.

Delete a Policy

Deleting a policy removes all content from the associated group.

To delete a Policy

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. Click **Policies**. The available policies are displayed.
4. Select one or more policies and in the More actions drop-down list in the tool bar, click **Delete**.

   The delete dialog is displayed.
5. Click **Delete** to permanently delete the selected policy.

   Deletion will take immediate effect and the policy will no longer be available in any group.

> **Note:**
> - The services associated with this policy still require a restart if the restart is pending.
> - You can also delete a policy from **Policy Details** screen. For more information on deleting a policy from **Policy Details** screen, refer View a Policy feature.

Edit a Policy

You can edit the content and settings of the policies. Once the policy is edited, the changes in the policy are reflected upon saving the policy. The changes are applied to the service once published. After saving and before publishing, the publication status of the changed policy is set to **Unpublished** if any settings are changed.

To edit a Policy

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. Click **Policies**. The available policies are displayed.
4. Select a policy to edit and click **Edit**.

5. Make the required changes in policy.
6. Do any one of the following:
   - Click **Save and Publish** to save and publish the policy. The policy will be listed under the Unpublished category.
   - Click **Save and Close** to save the settings and return to the Policies view.

Note: You can also edit a policy from **Policy Details** screen. For more information on editing a policy from **Policy Details** screen, refer View a Policy feature.

View a Policy

This topic describes the steps to view the properties of a Policy.

To view properties of the selected Policy

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. Click **Policies**. The available policies are displayed.
4. Click a row to view details about the selected policy.
5. To change the order of application rule assigned to the policy, do the following:
   1. To move the application rule or network down the order, click in the **Order** column.
   2. To move the application rule up the order, click in the **Order** column.
   3. You can also manually enter the order number in the **Order** column.
6. To change the order of network rule assigned to the policy, do the following:
   1. Click **Network Rules** tab.
   2. To move the network rule or network down the order, click in the **Order** column.
   3. To move the network rule up the order, click in the **Order** column.
   4. You can also manually enter the order number in the **Order** column.

   IMPORTANT: It is recommended not to order application rules or network rules deployed on the service from Service Config page if the service is part of Centralized Content Management.
7. To edit the policy, click **Edit Policy**. For more information on editing a policy, refer **Edit a Policy** feature.
8. To delete the policy, click **Delete Policy**. For more information on deleting a policy, refer **Delete a Policy** feature.
9. To publish the policy, click **Publish Policy**. For more information on creating and publishing a policy, refer **Create Policies** feature.
10. To enable or disable subscription, click **Subscribe** or **Unsubscribe** respectively.

Note:
- Subscription is not allowed for custom content.
- The **Subscribe** and **Unsubscribe** button is disabled if any one of the content selected is custom.

Enable Content for a Policy

This topic describes the steps to enable the content for a Policy.

To enable content

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Policies**.
4. Click the policy name to view the policy details.
5. In the **Application Rule**, **Network Rule**, **Feed**, **Log Device** or **LUA Parser** panel, click the row to select the content to be enabled. You can either select all content or select any specific content to be enabled.
6. Click **Enable**.

Disable Content for a Policy

This topic describes the steps to disable the content for a Policy.

To disable content

1. Go to **(CONFIGURE) > Policies**.
2. In the policies panel, click **Content**.
3. In the left panel, click **Policies**.
4. Click the policy name to view the policy details.
5. In the **Application Rule**, **Network Rule**, **Feed**, **Log Device** or **LUA Parser** panel, click the row to select the content to be disabled. You can either select all content or select any specific content to be disabled.
6. Click **Disable**.

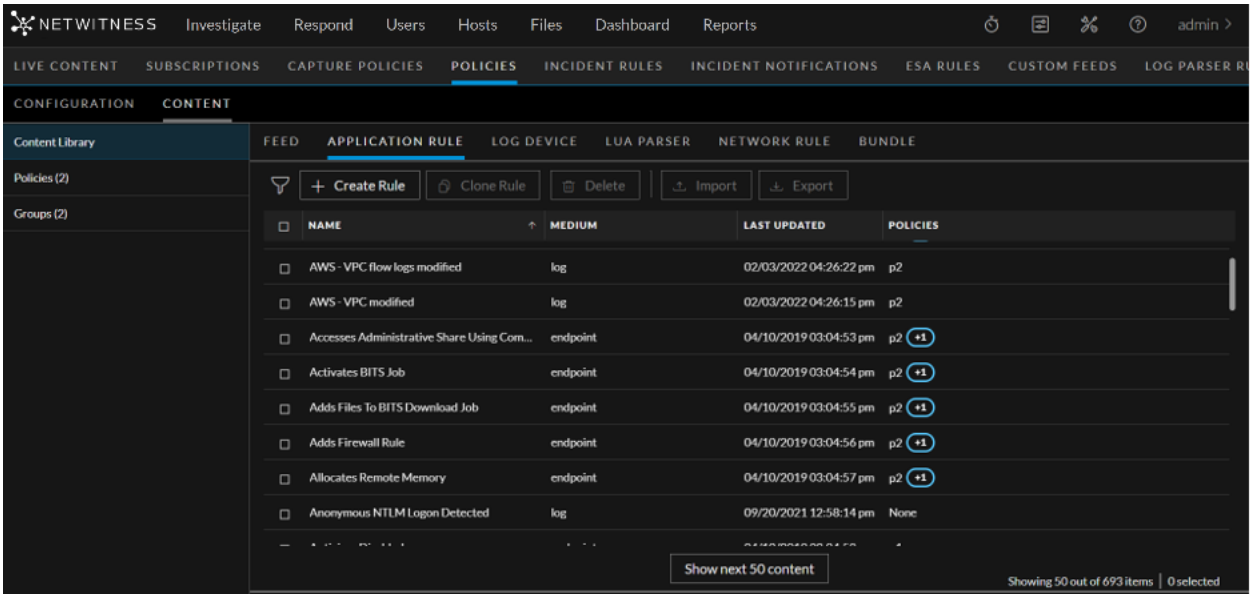
References

This section is a collection of references, which describe the user interface and more detailed information about how Policy-based Centralized Content Management works in NetWitness. The topics are presented in alphabetical order.

Content Library Tab

The **(CONFIGURE) > Policies** view contains two tabs: **Configuration** and **Content**.
The **CONTENT** tab has **Content Library**, **Policies** and **Groups** on the left panel.

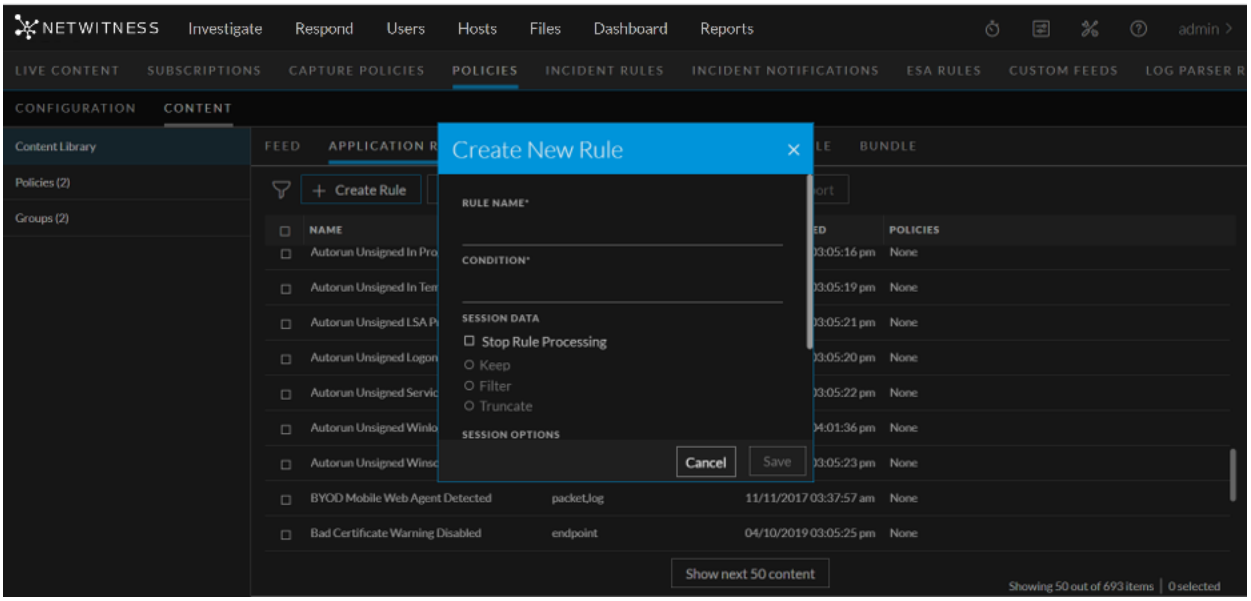
Below is an example of the Content > Content Library tab:



1	Toolbar • Create Rule - Lets you create a rule.
2	Rule List Pane • Name - Name of the rule. • Medium - Medium through which the rule is created. • Last Updated - Displays the time when the rule is updated. • Policies - Policies to which the rule is applied. You can also sort on any column. If you mouse over a column header, a sort icon is displayed: . Click the icon to sort by the selected column.

Create New Rule dialog:

Below is an example of the Create new rule dialog:

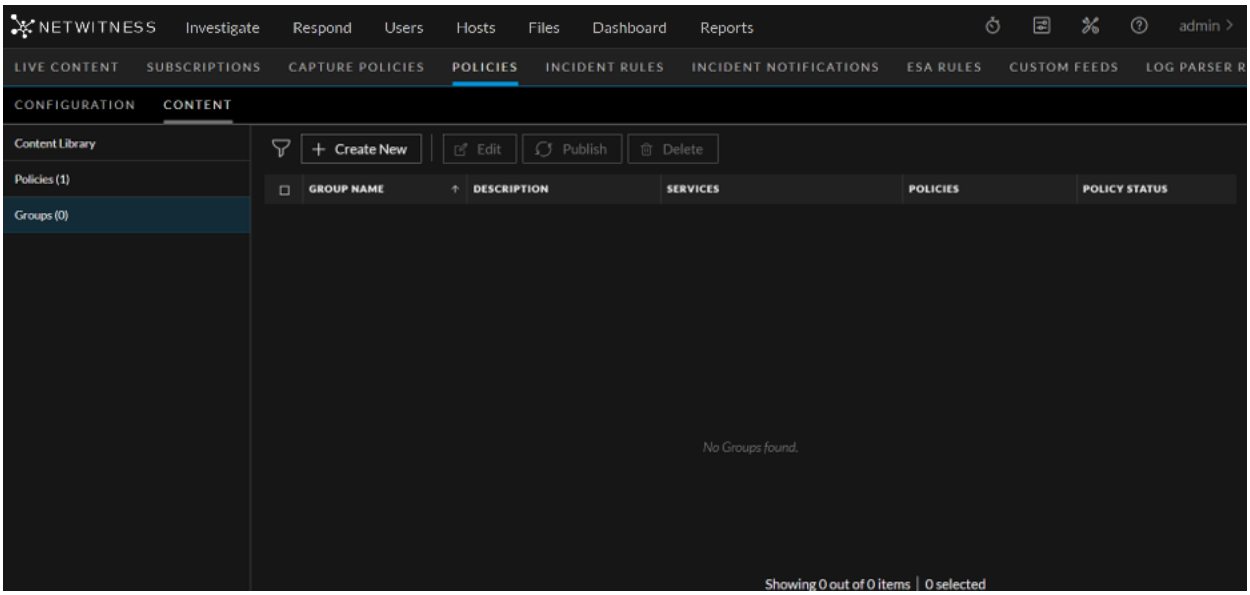


The table describes the information and options in the Create New Rule dialog:

Field	Description
Rule Name	Name of the new rule. The name should be unique.
Condition	Condition for the new rule.
Medium	Medium through which the rule is created.
Session Data	Session data for the new rule. Indicates if the rule processing should stop, keep, filter or truncate when the session data is running.
Session Options	Session options for the new rule. Indicates if the session options should be alert, forward or transient.
Alert On	Conditions for which the alert should be turned on.
Save	Saves the settings and closes the Create New Rule dialog.
Cancel	Cancels the operations.

Groups Tab

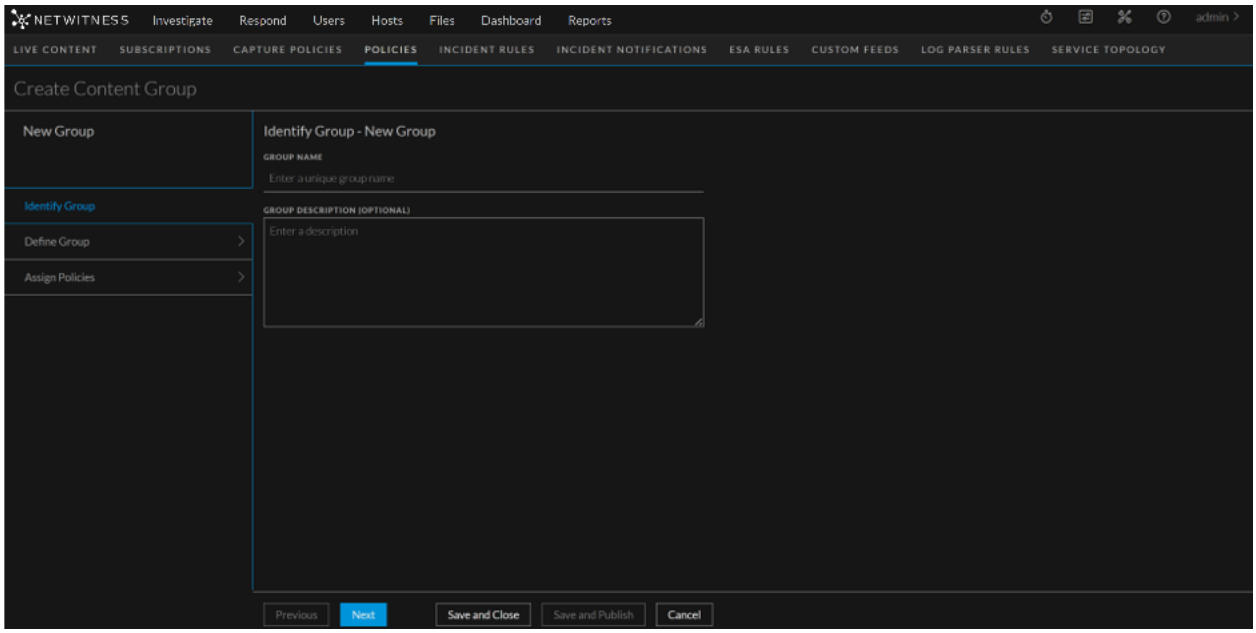
Below is an example of the Content > Groups tab:



The following table describes the Groups tab.

1	Toolbar • Create New - Lets you create a new group. For more information, see Create a group. • Edit - Lets you edit the group. For more information, see Managing Groups. • Publish - Publishes selected groups. • Delete - Deletes the selected group.
2	Group List Pane • Group Name - Name of the group. • Description - Description of the group. • Services - Displays the service to the which the group is applied. • Policies - Displays the policy to which the group is applied. • Policy Status - Status of the policy. The values are: Published, Unpublished, Failed, N/A. You can also sort on any column. If you mouse over a column header, a sort icon is displayed:  . Click  to sort by the selected column.
3	Groups Details Panel Displays the properties of the selected group.

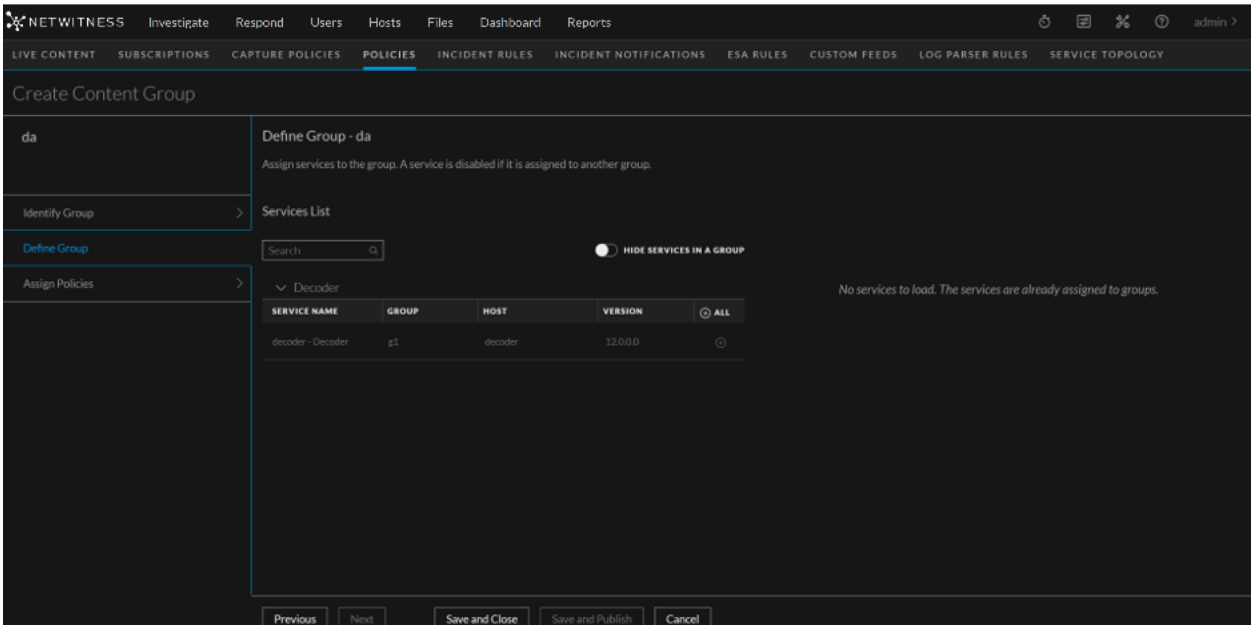
Below is an example of the Create group dialog:



The table describes the information and options in the Create Group dialog:

Field	Description
Group Name	Name of the group. The name should be unique.
Group Description (Optional)	Description of the group. Description should not exceed 8000 characters.
Save and Close	Saves the settings and closes the Create Group dialog.

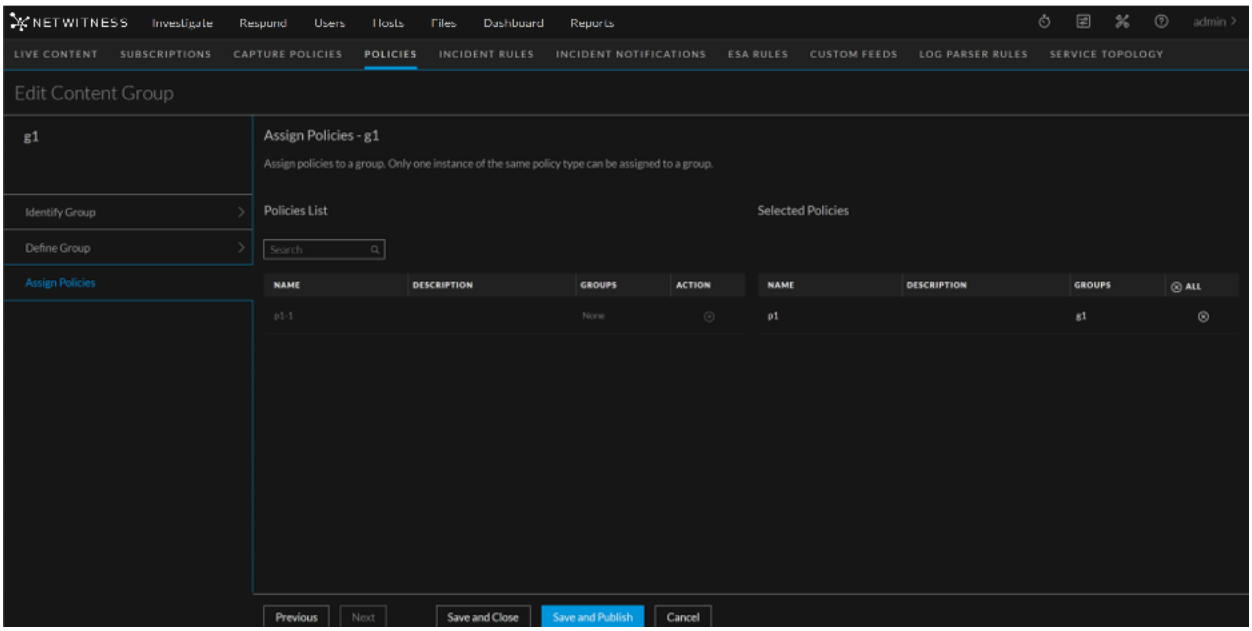
Below is an example of the define group dialog:



The table describes the information and options in the Define Group dialog:

Field	Description
Services List	Displays the list of services. The following describes services list: Service name – Name of the service. • Group - Name of the group. • Host - Host name of the service. • Version - Service version. • All - Lets you to add services to the group. You can either click ALL to add all services or click to add specific service.
Hide Services in a Group	Displays the services that is not assigned to any group. By default, this option is disabled.
Selected Services	Displays the list of selected services for the group.
Save and Close	Saves the setting and closed the create group dialog.
Save and Publish	Saves and publishes the created group. **Note:** This option is disabled if you have not: - Assigned services. - Assigned policies.

Below is an example of Assign policy dialog:



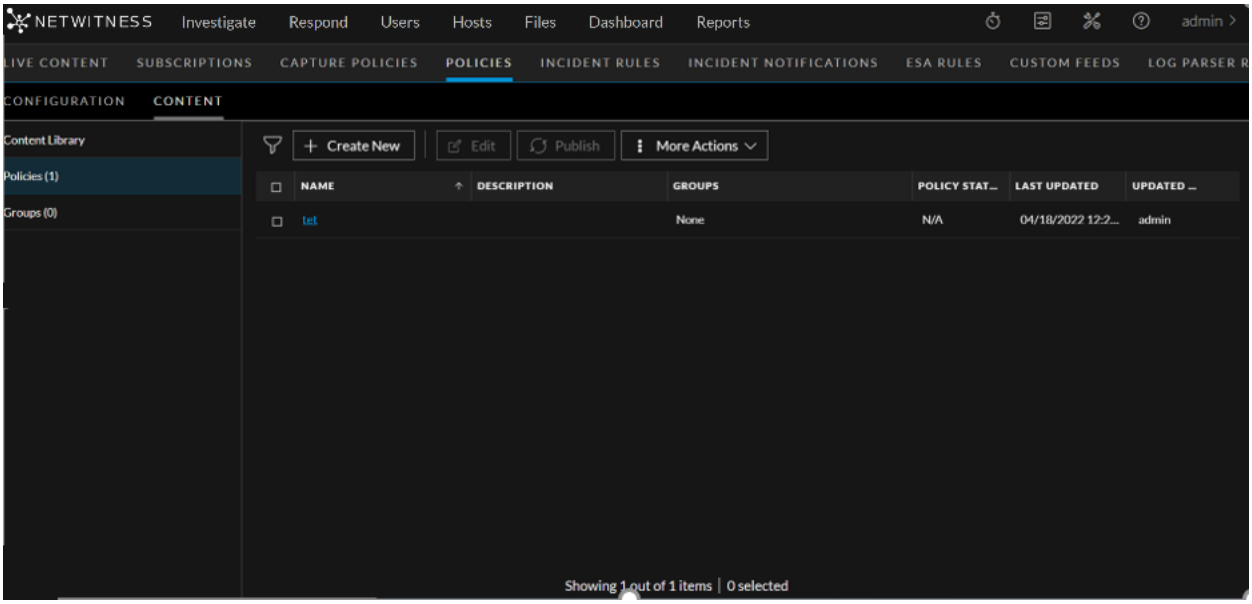
The following table describes assign policy dialog:

Field	Description

Policies List	Displays the list of policies associated with the group. The following describes policies list: • Name - Name of the policy. • Description - Description of the policy. • Groups - Groups associated with the policy. • Action - Click to add policies to the group.
Selected Policies	Displays the list of selected policies for the group.
Save and Close	Saves the setting and closed the create group dialog.
Save and Publish	Saves and publishes the created group. **Note:** This option is disabled if you have not: - Assigned services. - Assigned policies.

Policies Tab

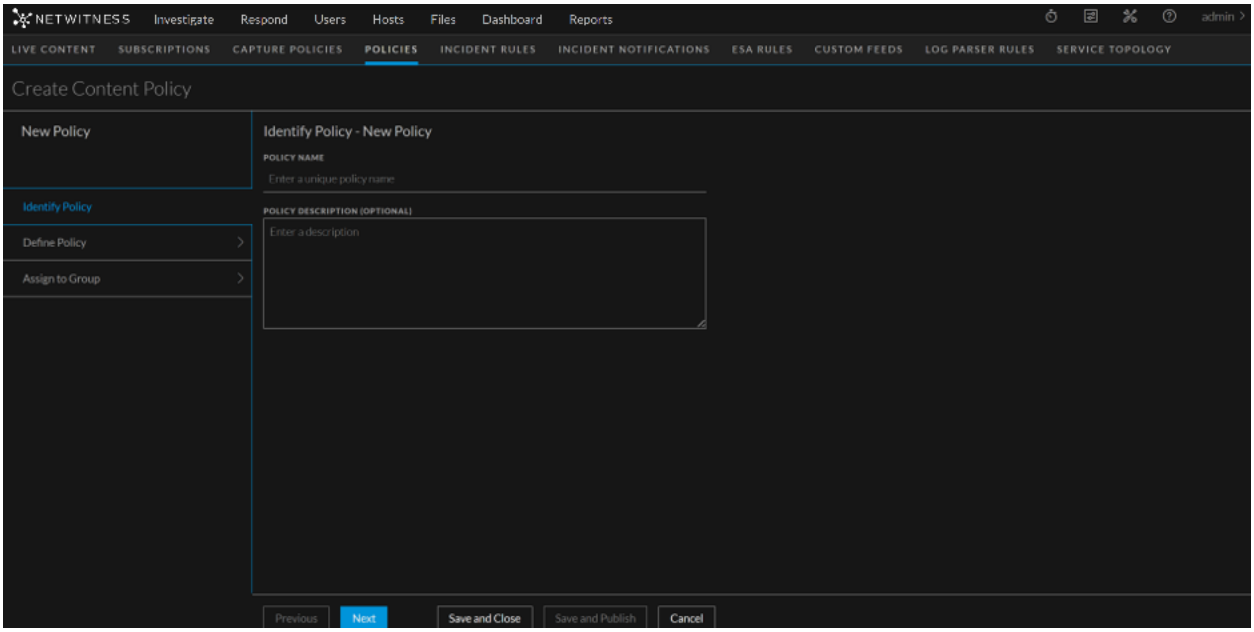
The **(CONFIGURE) > Policies** view contains two tabs: **Configuration** and **Content**.
Below is an example of the **Content > Policies** tab:



1	Toolbar:
	• Create New - Lets you create a new policy. For more information, see Create a policy.
	• Edit - Lets you edit the policy. For more information, see Managing Groups.
	• Publish - Publishes selected policy or policies.
	• More Actions:
	  • Assign to Group --Lets you assign policy to a group.
	  • Clone - Lets you clone a policy.
	  • Delete - Deletes the selected group or groups permanently.

2	Policy List Pane: • Name - Name of the policy. • Description - Description of the policy. • Groups - Lists the group to which this policy is applied. • Policy Status - Status of the policy. The values are: Published, Unpublished, Failed, N/A. • Last Updated - Displays the time when the policy is updated. • Updated By - The user who updated the policy. You can also sort on any column. If you mouse over a column header, a sort icon is displayed: . Click the icon to sort by the selected column.
3	Policy Details Panel: Displays the properties of the selected policy.

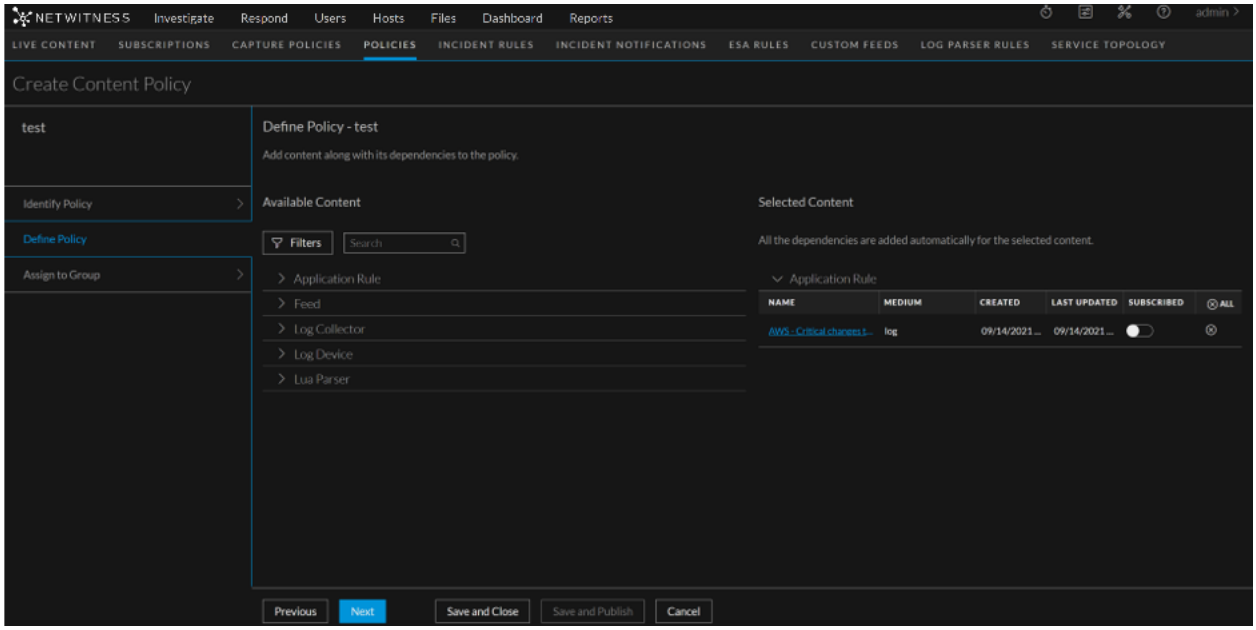
Below is an example of the Create Content Policy dialog.



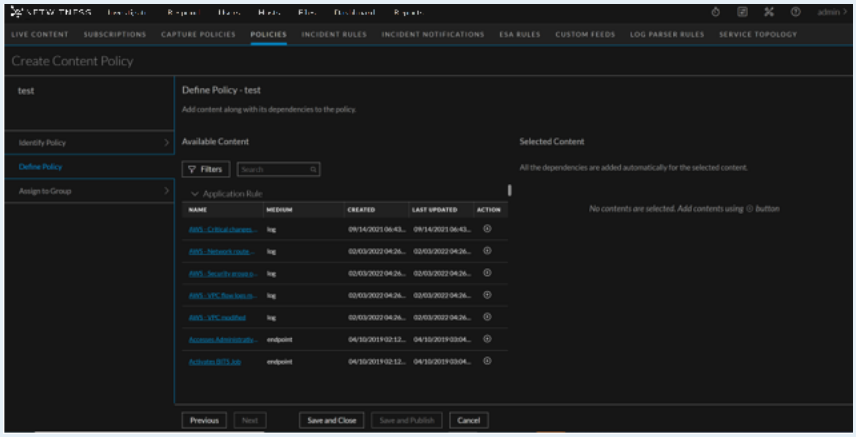
The table describes the information and options in the Create Policy dialog:

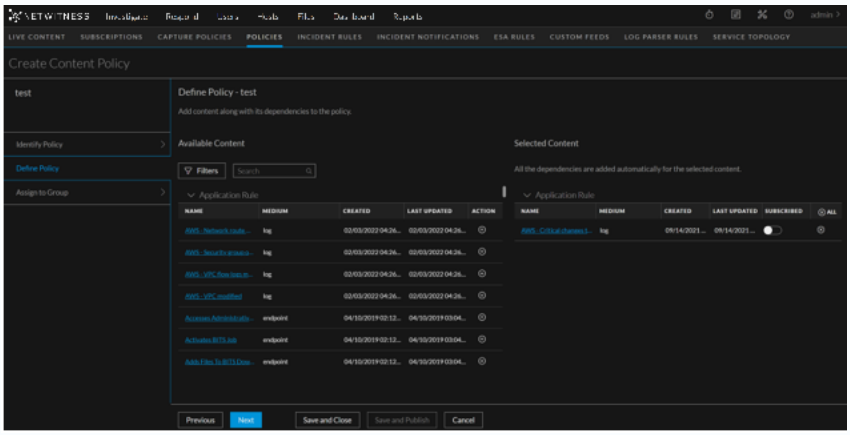
Field	Description
Policy Name	Name of the policy. The name should be unique.
Policy Description (Optional)	Description of the policy. Description should not exceed 8000 characters.

Define Policy Settings:

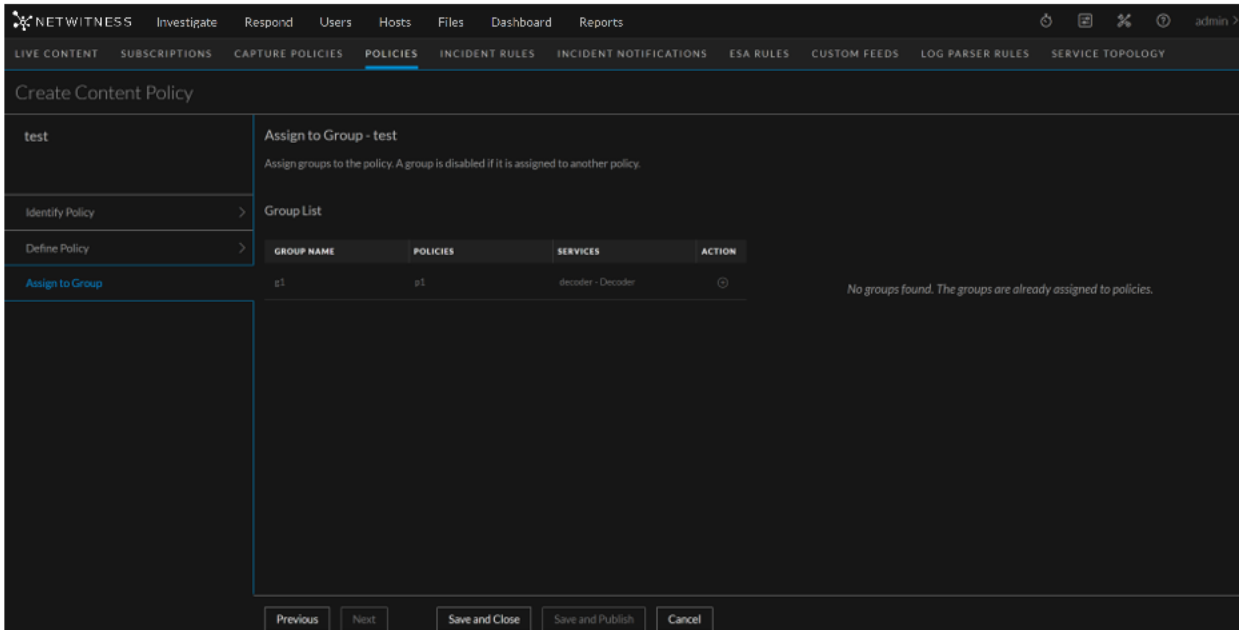


Field	Description

Available Content	Displays the available content resources in your deployment. Click expand the resource type.
	
	The following describes resource type:
	• Name - Name of the resource.
	• Medium - Meta data source medium. Available values for medium are as follows:
	  • **Endpoint**: applied to content that uses meta derived from endpoint agent and endpoint server data
	  • **Log**: applied to content that uses meta derived from log data
	  • **Packet**: applied to content that uses meta derived from network packets
	  • **Log and packet**: applied to content that correlates meta derived across log and packet data
	• Created - Displays the time when the resource is created
	• Last Updated - Displays the time when the resource is updated last.
	• Action- Click + to add the resource and its dependencies to your deployment.

Selected Content	Lists the selected resource. Additionally, you can subscribe the content. Once the content is subscribed, the content resources are updated automatically in case of any changes. 

Assign to Group:



Group List	Displays the list of groups associated with the policy. A group is disabled if it is already assigned to another policy. • Group Name • Policies • Services • Action
Selected Group	Lists the selected groups. Click ⊕ to add groups.
Save and Close	Saves the settings and closes the Create Policy dialog.
Save and Publish	Saves and publishes the created policy. **Note:** This option is disabled if: - Policy settings are not customized. - Policy is not assigned to groups.

Appendix A: Endpoint Risk Scoring Rules

Endpoint risk scoring requires the following content:

- "accesses administrative share using command shell"
- "activates bits job"

- "adds files to bits download job"
- "adds firewall rule"
- "allocates remote memory"
- "antivirus disabled"
- "archiving software reads multiple documents"
- "autorun debian package mismatch"
- "autorun file path not part of debian package"
- "autorun file path not part of rpm"
- "autorun key contains non-printable characters"
- "autorun"
- "autorun rpm mismatch"
- "autorun unsigned active setup"
- "autorun unsigned appinit_dlls"
- "autorun unsigned bho"
- "autorun unsigned bootexecute registry startup method"
- "autorun unsigned explorer registry startup method"
- "autorun unsigned hidden"
- "autorun unsigned hidden only executable in directory"
- "autorun unsigned ie toolbar"
- "autorun unsigned in appdatalocal directory"
- "autorun unsigned in appdataroaming directory"
- "autorun unsigned in programdata directory"
- "autorun unsigned in temp directory"
- "autorun unsigned logontype registry startup method"
- "autorun unsigned lsa provider"
- "autorun unsigned servicedll"
- "autorun unsigned winlogon helper dll"
- "autorun unsigned winsock lsp"
- "bad certificate warning disabled"
- "blacklisted file"

- "browser runs command prompt"
- "browser runs mshta"
- "browser runs powershell"
- "builds script incrementally"
- "clears application event log"
- "clears event logs using powershell"
- "clears security event log"
- "clears setup event log"
- "clears system event log"
- "combines binaries using command prompt"
- "command line usage of archiving software"
- "command line writes script files"
- "command prompt obfuscation"
- "command prompt obfuscation using value extraction"
- "command shell runs rundll32"
- "completes bits download job"
- "configures image hijacking"
- "configures port redirection"
- "copies binary over administrative share"
- "created in last month"
- "creates browser extension"
- "creates domain user account"
- "creates executable in startup directory"
- "creates local driver service"
- "creates local service"
- "creates local task"
- "creates local user account"
- "creates password-protected archive"
- "creates recursive archive"
- "creates remote process using wmi command-line tool"

- "creates remote service"
- "creates remote task"
- "creates run key"
- "creates shadow volume for logical drive"
- "creates suspicious service running command prompt"
- "debian package hash mismatch in important system directory"
- "debian package hash mismatch"
- "deletes backup catalog"
- "deletes firewall rule"
- "deletes shadow volume copies"
- "deletes shadow volume copies using powershell"
- "deletes usn change journal"
- "disables event logging service"
- "disables firewall"
- "disables safe mode"
- "disables security service"
- "disables startup repair"
- "disables uac"
- "disables uac remote restrictions"
- "disables windows audit policy"
- "disables windows defender using powershell"
- "downloads binary using certutil"
- "drops credential dumping tools"
- "dumps dns cache"
- "dyld inserted"
- "enables cleartext credential storage"
- "enables login bypass"
- "enables rdp from command-line"
- "enables safe mode"
- "enumerates arp table"

- "enumerates available systems on network"
- "enumerates domain account policy"
- "enumerates domain administrators"
- "enumerates domain computers"
- "enumerates domain controllers"
- "enumerates domain groups"
- "enumerates domain users"
- "enumerates enterprise administrators"
- "enumerates exchange domain servers"
- "enumerates exchange servers"
- "enumerates ip configuration"
- "enumerates local account policy"
- "enumerates local administrators"
- "enumerates local administrators on domain controller"
- "enumerates local groups"
- "enumerates local services"
- "enumerates local users"
- "enumerates logical disk"
- "enumerates mapped resources"
- "enumerates network connections"
- "enumerates primary domain controller"
- "enumerates processes on local system"
- "enumerates processes on remote system"
- "enumerates remote netbios name table"
- "enumerates remote resources"
- "enumerates route table"
- "enumerates services hosted in processes"
- "enumerates system info"
- "enumerates trusted domains"
- "evades scanning within windows defender"

- "evasive powershell used over network"
- "event viewer executes uncommon binary"
- "execute dll through rundll32"
- "exports sensitive registry hive"
- "extracts password-protected archive"
- "file encrypted"
- "file hidden"
- "file path not part of debian package in important system directory"
- "file path not part of debian package"
- "file path not part of rpm in important system directory"
- "file path not part of rpm"
- "file vault disabled"
- "floating module and hooking"
- "floating module in browser process"
- "floating module in os process"
- "floating module"
- "gatekeeper disabled"
- "gets current user as system"
- "gets current username and group information"
- "gets current username"
- "gets hostname"
- "gets remote time"
- "gina replacement"
- "graylisted file"
- "hidden and hooking"
- "hidden in appdata"
- "hidden plist and autorun"
- "hidden running as root"
- "hooks audio output function"
- "hooks authentication function"

- "hooks crypto function"
- "hooks dnsquery function"
- "hooks gui function"
- "hooks network http function"
- "hooks network io function"
- "hooks ntldr function"
- "hooks registry access function"
- "hooks registry enumeration function"
- "http daemon runs command prompt"
- "http daemon runs powershell"
- "http daemon runs reconnaissance tool"
- "http daemon writes executable"
- "ie dep disabled"
- "ie enhanced security disabled"
- "in appdata directory"
- "in hidden directory"
- "in recycle bin directory"
- "in root of appdatalocal directory"
- "in root of appdataroaming directory"
- "in root of logical drive"
- "in root of program directory"
- "in root of users directory"
- "installs root certificate"
- "in system volume information directory"
- "in temporary directory"
- "in uncommon directory"
- "invalid signature"
- "kext signature validation disabled"
- "lateral movement with credentials using net utility"
- "ld preload"

- "library preferences directory"
- "lists anti-spyware products"
- "lists antivirus products"
- "lists firewall products"
- "login bypass configured"
- "lua disabled"
- "mac firewall disabled"
- "malicious file by reputation service"
- "maps administrative share"
- "maps ipc$ share"
- "misleading file extension"
- "modifies file associations"
- "modifies image file execution for persistence"
- "modifies registry using command-line registry tool"
- "modifies run key"
- "modifies shell-open-command file association"
- "modifies startup folder location"
- "modifies winlogon dll for persistence"
- "modifies winlogon registry settings"
- "mshta runs command prompt"
- "mshta runs powershell"
- "mshta runs scripting engine"
- "mshta writes executable"
- "network access"
- "no antivirus notification disabled"
- "no firewall notification disabled"
- "non-microsoft modifies bad certificate warning setting"
- "non-microsoft modifies firewall policy"
- "non-microsoft modifies internet zone setting"
- "non-microsoft modifies lua setting"

- "non-microsoft modifies registry editor setting"
- "non-microsoft modifies security center config"
- "non-microsoft modifies services imagepath"
- "non-microsoft modifies task manager setting"
- "non-microsoft modifies windows system policy"
- "non-microsoft modifies zone crossing warning setting"
- "no uac notification disabled"
- "no windows update notification disabled"
- "office application crashed"
- "office application injects remote process"
- "office application runs bits"
- "office application runs command prompt"
- "office application runs powershell"
- "office application runs scripted ftp"
- "office application runs scripting engine"
- "office application runs task scheduler"
- "office application runs wmi scripting engine"
- "office application writes executable"
- "opens browser process"
- "opens os process"
- "opens process"
- "opswat reported infected"
- "opswat reported suspicious"
- "os process runs command shell"
- "packed and autorun"
- "packed and network access"
- "packed"
- "performs scripted file transfer"
- "possible login bypass"
- "possible mimikatz activity"

- "possible rdp session hijacking"
- "possibly configures uac bypass"
- "possibly renamed net.exe detected"
- "potential abuse of odbcconf"
- "potential outlook exploit"
- "powershell command using string manipulation"
- "powershell injects remote process"
- "powershell opens lsass process"
- "powershell runs command prompt"
- "powershell runs scripting engine"
- "process authorized in firewall"
- "process redirects to stdout or stderr"
- "process with matched yara rule"
- "process with opswat reported infected"
- "process with opswat reported suspicious"
- "psexesvc runs powershell"
- "psexesvc runs scripting engine"
- "psexesvc runs shell commands"
- "pubprn detection"
- "queries cached kerberos tickets"
- "queries processes on local system"
- "queries processes on remote system"
- "queries registry using command-line registry tool"
- "queries terminal sessions"
- "queries users logged on local system"
- "queries users logged on remote system"
- "record screen captures using psr tool"
- "registers always install elevated policy"
- "registers appcert dll"
- "registers appinit dll"

- "registers boot execute"
- "registers lsa authentication package"
- "registers lsa notification package"
- "registers lsa security package"
- "registers netsh helper dll"
- "registers port monitor dll"
- "registers shim database"
- "registers startup during safe mode boot"
- "registers time provider dll"
- "registry tools disabled"
- "regsvr32 creates windows task"
- "regsvr32 runs powershell"
- "regsvr32 runs rundll32"
- "regsvr32 writes executable"
- "remote directory traversal"
- "removes windows defender definitions"
- "rpm hash mismatch in important system directory"
- "rpm hash mismatch"
- "rpm ownership changed"
- "rpm permissions changed"
- "rundll32 creates windows task"
- "rundll32 runs powershell"
- "runkey persistence"
- "runs acl management tool"
- "runs active directory service query tool"
- "runs binary located in recycle bin directory"
- "runs binary located in root of logical drive"
- "runs binary located in root of program directory"
- "runs binary located in root of users directory"
- "runs binary located in system volume information directory"

- "runs blacklisted file"
- "runs certutil with decode arguments"
- "runs certutil with encode arguments"
- "runs certutil with hashfile arguments"
- "runs chained command shell"
- "runs chmod"
- "runs credential dumping tools"
- "runs curl"
- "runs ditto"
- "runs dns lookup tool for txt record"
- "runs dns lookup tool"
- "runs file attributes modification tool"
- "runs file transfer tool"
- "runs forfiles.exe"
- "runs graylisted file"
- "runs ifconfig"
- "runs kextload"
- "runs kextstat"
- "runs launchctl"
- "runs malicious file by reputation service"
- "runs mshta with http argument"
- "runs mshta with script argument"
- "runs msiexec with http argument"
- "runs netstat"
- "runs network configuration tool"
- "runs network connectivity tool"
- "runs one letter executable"
- "runs one letter script"
- "runs ping"
- "runs powershell bypassing execution policy"

- "runs powershell decoding base64 string"
- "runs powershell defining function"
- "runs powershell downloading content"
- "runs powershell invoke-mimikatz function"
- "runs powershell memory stream function"
- "runs powershell"
- "runs powershell shellexecute function"
- "runs powershell using encoded command"
- "runs powershell using environment variables"
- "runs powershell with hidden window"
- "runs powershell with http argument"
- "runs powershell with long arguments"
- "runs psexec on remote system and silently accepts user license"
- "runs psexec on remote system as system user"
- "runs ps"
- "runs registry tool"
- "runs regsvr32 com scriplets"
- "runs regsvr32 using one letter dll"
- "runs regsvr32 with http argument"
- "runs regsvr32 without arguments"
- "runs remote execution tool"
- "runs remote powershell command"
- "runs robocopy.exe"
- "runs rundll32 using one letter dll"
- "runs rundll32 with http argument"
- "runs rundll32 with javascript argument"
- "runs rundll32 without arguments"
- "runs scripting engine in batch mode using execution engine argument"
- "runs scripting engine"
- "runs service control tool"

- "runs shim database installer"
- "runs sh"
- "runs suspicious file by reputation service"
- "runs tar"
- "runs tasks management tool"
- "runs unzip"
- "runs waitfor.exe"
- "runs wmi command-line tool"
- "runs wmi scripting engine"
- "runs xcopy.exe"
- "safari fraud website warning disabled"
- "scripting addition in process"
- "scripting engine injects remote process"
- "scripting engine runs powershell"
- "scripting engine runs regsvr32"
- "scripting engine runs rundll32"
- "self signed"
- "services in programdata directory"
- "services runs command shell"
- "smartscreen filter disabled"
- "starts local service"
- "starts rdp service"
- "starts remote service"
- "stops diagtrack service"
- "stops error reporting service"
- "stops security service"
- "stops windows update service"
- "sudo no password prompt"
- "suspicious file by reputation service"
- "suspicious regsvr32.exe task"

- "system integrity protection disabled"
- "system restore disabled"
- "tampers with windows defender registry"
- "task manager disabled"
- "tasks in programdata directory"
- "terminates process"
- "transfers file using bits"
- "uac disabled"
- "unexpected csrss.exe parent"
- "unexpected explorer.exe destination location"
- "unexpected explorer.exe parent"
- "unexpected explorer.exe source location"
- "unexpected lsass.exe parent"
- "unexpected lsm.exe parent"
- "unexpected msdtc.exe parent"
- "unexpected os process destination location"
- "unexpected os process source location"
- "unexpected runtimebroker.exe parent"
- "unexpected services.exe parent"
- "unexpected smss.exe parent"
- "unexpected svchost arguments"
- "unexpected svchost.exe parent"
- "unexpected taskhostw.exe parent"
- "unexpected wininit.exe parent"
- "unexpected winlogon.exe parent"
- "unknown segment"
- "unsigned copies self"
- "unsigned creates remote thread and file hidden"
- "unsigned creates remote thread"
- "unsigned cron job"

- "unsigned deletes self"
- "unsigned kext"
- "unsigned library in suspicious daemon"
- "unsigned module in signed process"
- "unsigned reserved name"
- "unsigned runs python"
- "unsigned writes executable"
- "unsigned writes executable to appdatalocal directory"
- "unsigned writes executable to appdataroaming directory"
- "unsigned writes executable to library application support directory"
- "unsigned writes executable to library directory"
- "unsigned writes executable to library preferences directory"
- "unsigned writes executable to scripting additions directory"
- "unsigned writes executable to system directory"
- "unsigned writes executable to var directory"
- "unsigned writes executable to windows directory"
- "unsigned writes to autorun"
- "uses libnss"
- "uses libpcap"
- "uses mach injection"
- "uses mach override"
- "warning on post redirect disabled"
- "windows firewall disabled"
- "windows task runs powershell"
- "windows update disabled"
- "wmic remote node activity"
- "wmiprvse runs command shell"
- "wmiprvse runs powershell"
- "wmiprvse runs scripting engine"
- "writes blacklisted file"

- "writes executable to recycle bin directory"
- "writes executable to root of logical drive"
- "writes executable to root of program directory"
- "writes executable to root of users directory"
- "writes executable to system volume information directory"
- "writes graylisted file"
- "writes malicious file by reputation service"
- "writes suspicious file by reputation service"
- "yara rule matched"
- "executable in ads"
- "explorer public folder dll load"
- "powershell double base64"
- "outbound from windows directory"
- "outbound from unsigned temporary directory"
- "unsigned opens lsass"
- "outbound from unsigned appdata directory"
- "rdp launching loopback address"
- "autorun invalid signature windows directory"
- "command shell copy items"

NetWitness® Platform XDR

Version 12.0

Log Parser Customization Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Log Parser Customization **5**
- Dynamic Rules 5
- JSON Mapping 5

Add or Delete a Log Parser **7**
- Add a Log Parser 7
- Delete a Log Parser 8
- Add Dynamic Log Parser Parameters 8
  - Import or Export a Log Parser 8

JSON Mappings **10**
- View JSON Mappings 10
- Add a JSON Mapping 11
- Auto Discover JSON Mappings 12
- Deploy JSON Parser 14

Add or Delete a Log Parser Rule **15**
- About Log Parser Rules 15
- Custom Log Parser Rules 15
- Guidelines for Custom Rules 16

Default Log Parser and Log Parser Rules **17**
- Default Log Parser 17
- Highlight Matching Patterns 18
- Highlight Overlapping Patterns 20

Use Cases **22**
- Use Case 1: On Board a New Event Source 22
- Use Case 2: Modify an Existing Parser 22

Extend an Existing Log Parser Example **23**
- Task Overview 23
- Notes 23
- Add the Log Parser 23
- About Custom Rules 25
- Add Rules and Deploy 25
- Regex Values 28

Appendix A: Select the Reference Log Decoder **30**

Appendix B: Move Log Parsers to Production **31**

Appendix C: Troubleshooting and Limitations **32**
- Troubleshooting 32
- Delete a Log Parser Manually 32

NwLogPlayer 33
Limitations 34
Parser Rules Tab 35
Log Parsers Panel 36
Dynamic Rules 38
Details Pane 38
Rules Pane 41
JSON Mappings 42
Sample JSON Message 42
Auto Discover JSON Mappings 44
Remove Unmapped Entries 46
Meta Mappings 47
Mapping Details 47
Disable log Parser Rules 48
Add VARTYPE Support to CEF Parser 49

Log Parser Customization

Note: The JSON Mapping information in this guide applies to NetWitness Version 11.5 and later.

You use the Log Parser Rules view (available from the (Configure) view) to customize rules and to map meta for your log parsers.

The *default log parser* parses logs that do not match any installed log parsers. The information contained in such a log is processed against the default log parser's rules, and metadata is then extracted by those rules and is available for Enrichment, Investigation, Reporting, and Alerting. This provides immediate visibility into logs from custom or unsupported sources.

You can also add or extend a log parser. For example, you may need to parse certain fields differently than in the manner provided by the log parser for a particular event source. You can add rules that change the way meta information is extracted from the logs for the event source. And you can then map the extracted information to NetWitness meta keys.

Finally, you can view and test sample log messages and rules for your log parsers, including the default log parser.

To access this tab, go to **(Configure) > Log Parser Rules**. For more details, see Parser Rules Tab.

Dynamic Rules

The **Dynamic Rules** entry for a log parser displays the following information:

- The default log parser that parses logs that are not associated with a particular log parser
- Native XML-defined device parsers that have been extended with dynamic log parser rules, and
- User-created custom device parsers used to parse unsupported custom event sources

The dynamic parse rules are used to parse arbitrary values triggered by a literal token on the left hand side of the value in unstructured logs. Currently, this is used to parse name-value pairs.

Following are some examples (the token is in bold red):

```
src: 1.2.3.4
src = 1.2.3.4
src=1.2.3.4
Source address is 1.2.3.4
```

JSON Mapping

This allows you to parse JSON nodes from Logstash as well as Log Collection Plugins by mapping JSON nodes to an appropriate meta on which the value should be saved.

The **JSON Mappings** for a log parser displays the following information:

- Sample JSON Message - This allows you paste the log messages.
- The list of JSON Mappings: these are the names that represent the meta information.
- Details of each mapping: for each mapping, the display name, path, NetWitness meta key, and a text description.

The JSON mapping functionality is for strictly paring structured JSON logs, and mapping values from the log to meta or fine parsing. The parsing is not applied to arbitrary logs; only logs where we know the exact structure of the data.

For example:

```
{ "event": { "source": { "address": "1.2.3.4", "port":8080 } } }
```

NetWitness knows the structure of logs when it knows the event source type, or when you add specific JSON mappings.

Add or Delete a Log Parser

Note: The information in this topic applies to NetWitness Version 11.2 and later.

For version 11.2, NetWitness has added the ability to add log parsers through the UI. You can also delete log parsers, as long as they have never been deployed to a Decoder. You can create a new log parser definition from scratch, or extend an existing one.

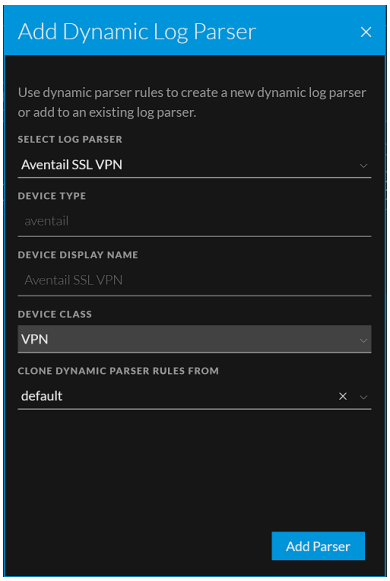
You can add a log parser to extend the functionality for an existing parser. For example, if you have some unknown messages for the Cisco Pix parser, you could add rules to match your unknowns.

IMPORTANT: If you are adding a new log parser, for example when onboarding an event source, you must map the event source IP to the new log parser in order for messages to be parsed. For details, see "Acknowledging and Mapping Event Sources" in the *Event Source Management User Guide*.

Add a Log Parser

1. In the NetWitness UI, navigate to **(Configure) > Log Parser Rules**.
2. From the **Log Parsers** pane, click **Add Parser**.

   The Add Dynamic Log Parser dialog box is displayed.



3. Fill in details for this dialog box. For details, see Add Dynamic Log Parser Parameters below.
4. Click **Save** to save the new log parser.

   This updates the definition file in the file system. It *does not* deploy the changes.
5. To deploy your changes to all of your Decoders, click **Deploy**.

Delete a Log Parser

You can use the UI to delete a log parser.

To delete a log parser:

1. In the NetWitness UI, navigate to **(Configure) > Log Parser Rules**.
2. From the **Log Parsers** pane, select a log parser, then click **Delete**.

   Delete Parser dialog box is displayed.
3. Click **Delete Parser** to remove the log parser from the system.

Note: If you have encounter any issues when you attempt to delete a parser, see the Troubleshooting section, Delete a Log Parser Manually.

Add Dynamic Log Parser Parameters

When you are adding a log parser, the following parameters are available.

Field	Details
SELECT LOG PARSER	Select NEW, or choose an existing log parser. By choosing an existing log parser, you can add rules to that parser, essentially extending its parsing capabilities. **Note:** If you select an existing log parser, the remaining fields are auto-filled based on the values for selected log parser.
DEVICE TYPE	Enter a string to define the device type. The name must be between 3 and 30 alphanumeric characters (including underscores), and must not match the name of any existing log parsers.
DEVICE DISPLAY NAME	Enter the display name for the log parser. **Note:** The display name must be 64 characters or fewer, and must not match the name of any other device display name.
DEVICE CLASS	Select a device class.
CLONE DYNAMIC PARSER RULES FROM	Leave blank to start with no rules, or select one of the existing log parsers to clone its rules.

Import or Export a Log Parser

To import the log parser:

1. In the NetWitness UI, navigate to **(Configure) > Log Parser Rules**.
2. From the **Log Parsers** panel, click **Import**.

3. Select a log parser to import.
   The log parser is imported successfully. This will import the selected log parser rules along with the JSON mapping (if any).

To export the log parser:

1. In the NetWitness UI, navigate to  **(Configure) > Log Parser Rules**.
2. Select the log parser to export.
3. In the **Log Parsers** panel, click **Export**.
   The selected parser is exported successfully in a .json format. This will export the selected log parser rules along with the JSON mapping (if any).

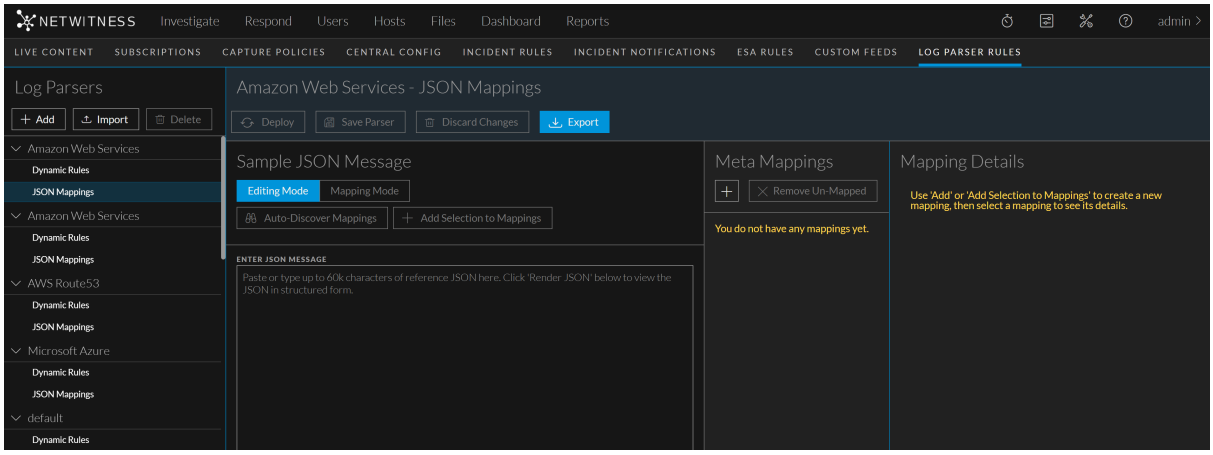
JSON Mappings

JSON mappings is applicable from version NetWitness Platform 11.5 and Later.

View JSON Mappings

1. In the NetWitness UI, go to **(Configure) > Log Parser Rules**.
2. From the Log Parsers pane, select a parser, then click JSON Mappings.

   The JSON Mappings and Mapping Details are shown for the parser you selected.



The Mapping Details pane displays the following information.

Field	Details
DISPLAY NAME	This name corresponds to the name displayed in the JSON Mappings panel.
PATH	The path to where the values for this portion of the log are stored.
DESCRIPTION	Optionally, you can enter a text description for this mapping.
META	Select a meta key to which this value from the log is mapped. Select a value from the drop-down menu. Optional if you choose a Value Format.

Field	Details
VALUE FORMAT	Choose a value format parser onto which to pass this JSON value. You can either select an existing meta or **Custom Regex Type**. If you select custom regex type, you must define the regex and capture to *fine* parse the value in the meta. Optional if you choose a Meta.
CUSTOM REGEX TYPE	Select Custom Regex Type from the Value Format drop-down, which allows you to add new custom regex type.
REGEX PATTERN	Specify a regex to identify different pieces of data contained within a JSON node value.
FIRST CAPTURE	Select a meta key that should be captured first based on the value defined in the Regex pattern.
ADD A CAPTURE	New capture field is added. By default, it is loaded with meta keys in the drop-down. You can add maximum of 20 captures and this option will be disabled once it reaches maximum.

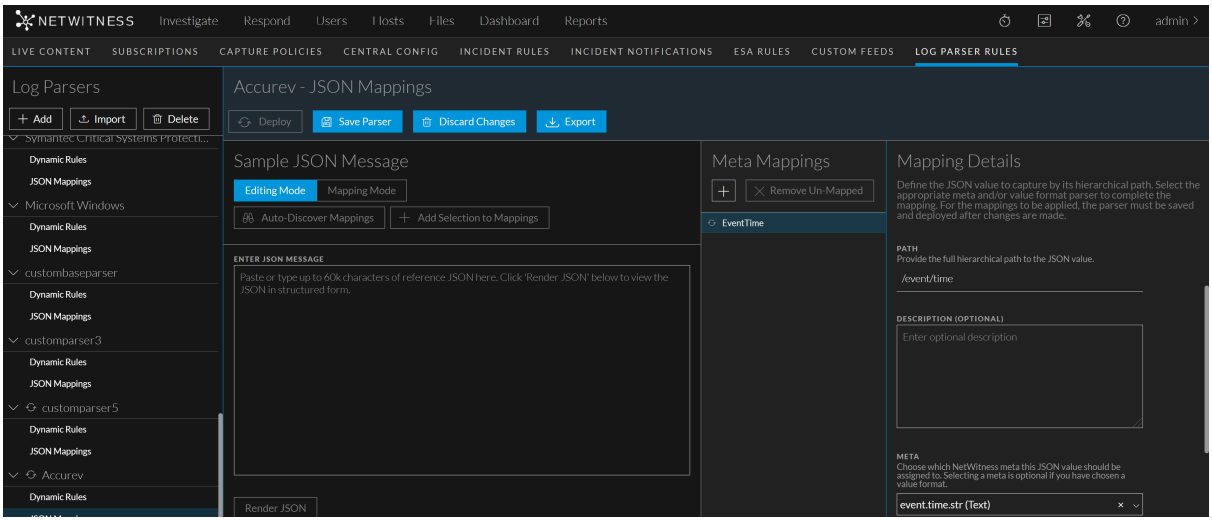
Note: You need to select a meta or enter a Value Format, but you do not need to fill in values for both settings.

Add a JSON Mapping

After you add a parser, as described in Add a Log Parser, you can then add JSON mappings.

1. Follow the procedure to add a parser.
2. Select the **JSON Mappings** entry for the newly-added parser.

   The following screen shows an example where an Accurev is added:



3. Click **Add New** to begin adding a mapping.
4. Enter values for display name, path, meta or value format (or both), and (optionally) a description.
5. Click **Save Parser** to save your new mapping.

Auto Discover JSON Mappings

Beginning with NetWitness version 11.5.1, you can automatically create the mappings without the need to manually enter the name and path of the mapping.

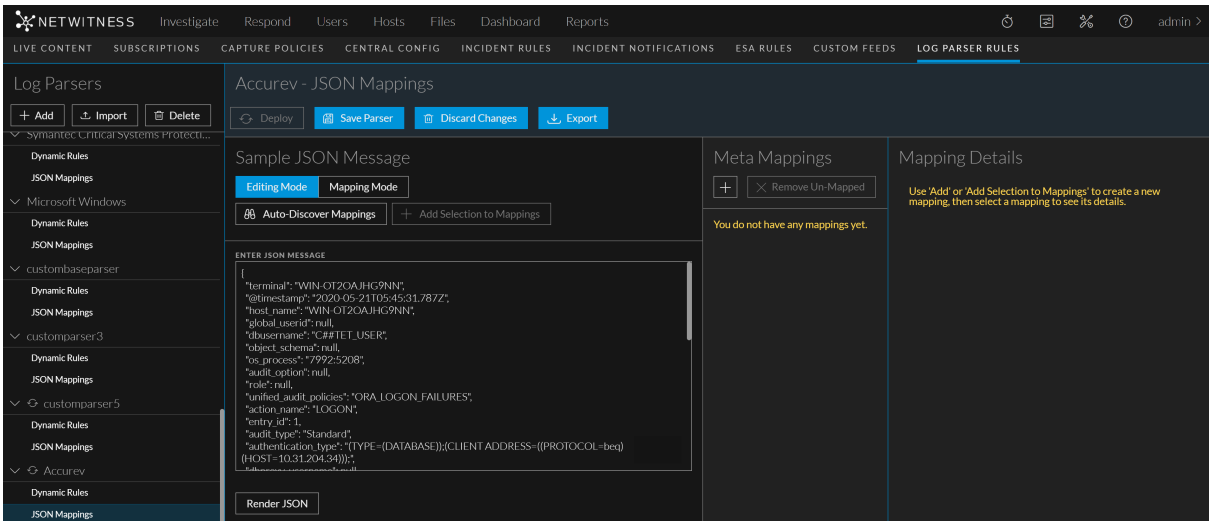
As an example, we are using the following JSON log:

```
{"terminal":"WIN-OT2OAJHG9NN","@timestamp":"2020-05-21T05:45:31.787Z","host_
name":"WIN-OT2OAJHG9NN","global_userid":null,"dbusername":"C##TET_
USER","object_schema":null,"os_process":"7992:5208","audit_
option":null,"role":null,"unified_audit_policies":"ORA_LOGON_
FAILURES","action_name":"LOGON","entry_id":1,"audit_
type":"Standard","authentication_type":"(TYPE=(DATABASE));(CLIENT ADDRESS=
((PROTOCOL=beq)(HOST=10.31.204.34)));","dbproxy_username":null,"external_
userid":null,"@version":"1","new_schema":null,"new_name":null,"statement_
id":1,"proxy_sessionid":0,"os_username":"WIN-
OT2OAJHG9NN\\Administrator","system_privilege":null,"sql_
binds":null,"timestamp":"2020-05-21 10:22:12","client_program_
name":"sqlplus.exe","sessionid":4125005309,"userhost":"WORKGROUP\\WIN-
OT2OAJHG9NN","rman_device_type":null,"object_name":null,"event_timestamp_
utc":"2020-05-20T23:22:12.452Z","system_privilege_used":null,"return_
code":1017,"version":"19.0.0.0.0","instance_name":"orcl","sql_
text":null,"target_user":null,"fga_policy_name":null,"rman_object_
type":null,"dbid":1566661212,"rman_operation":null}
```

To auto-discover JSON mappings:

1. Select the **JSON Mappings** entry for the appropriate parser.
2. Paste JSON formatted log message in to the **Sample JSON message** text box, and click **Render JSON**.

The screen looks like this:



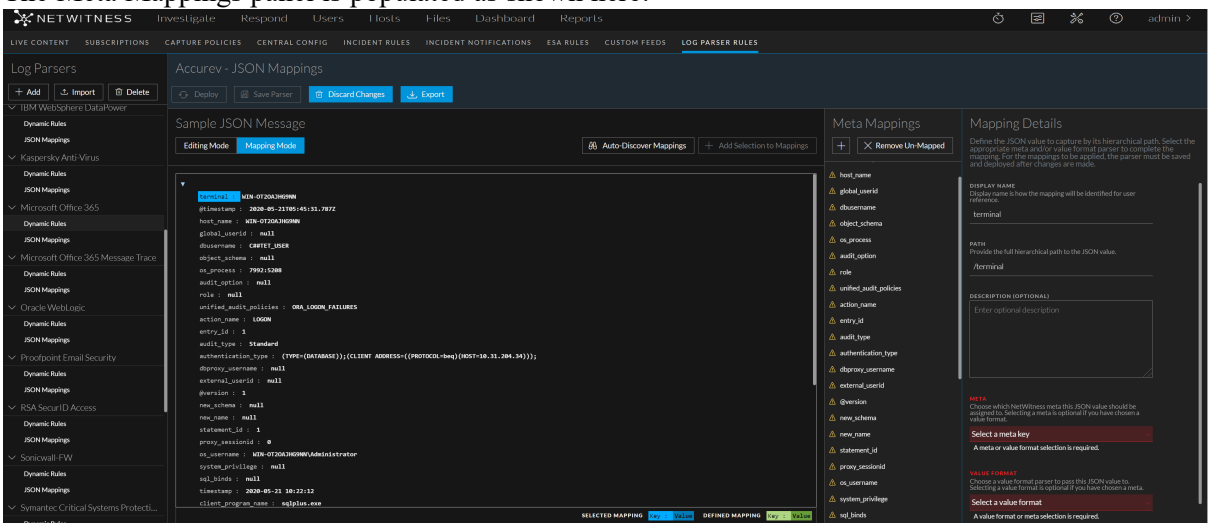
Rendering JSON in Editing Mode allows you to view and edit (if needed) the logs in a pretty format. Additionally, if your text is not valid JSON, the text field is displayed in red.

3. Click **Mapping Mode**, to view the JSON in a collapsable tree format which also highlights the mapping.

Note: In **Mapping mode**, you cannot edit the sample Log.

4. Click **Auto-Discover Mappings** to discover the JSON nodes and create mappings.

The Meta Mappings panel is populated as shown here:



5. After you auto-discover, note that all the mappings are invalid (preceded by this icon: ⚠). You cannot save your changes until all the mappings are valid (mapping is preceded by this icon: 🔄) or removed.

Note:
The meta value is highlighted in blue if it matches the selected mapping including the regex (if used).
The meta value is highlighted in green if it matches any other mappings that are not selected.

6. To make a mapping valid, you need to select a **Meta Key** or **Value Format** for all the mappings that you want to parse and save.

Note: In the **Value Format** you can either select an existing meta or **Custom Regex Type**.

7. If you want to *fine* parse the value in the meta in the **Value Format** drop-down, select **Custom Regex Type** and enter the regex in the **Regex Pattern** field. For example,`\s*(\b(?:[0-9]{1,3}\.){3}[0-9]{1,3}\b):?(\d*).`

Note: The custom regex will be added to the database, when you **Save Parser**.

8. For a value in the Regex, you can select a meta key and store the capture as below:
   - **First capture**: The first portion of the string, up to the period character is stored to the meta key. For example, `ip.src`
   - To add more capture, click **Add Capture** and select the meta key.
     For more information on Regex, see Regex Values.
9. If there are mappings that you do not want to save, you can select the mapping and click **Delete**. Alternatively, after you complete all of the mappings that you want to keep, you can click **Remove Unmapped** to remove all mappings that you have not yet validated.
10. After you have either completed or removed all of your mappings, click **Save Parser** to save your new mappings. Note that the icon preceding each mapping is removed.

Deploy JSON Parser

You need to deploy a JSON parser so that logs coming in to any decoder are parsed appropriately and meta is generated and stored correctly.

To deploy a parser, select it from the list and click **Deploy**. The parser, its dynamic rules, and its mappings are sent to all Log Decoders.

Note: A JSON parser must have at least one rule or mapping to enable deployment.

Add or Delete a Log Parser Rule

Note: The information in this topic applies to NetWitness Version 11.2 and later.

For version 11.2, NetWitness has added the ability to create custom rules for log parsers. You can create rules to change how meta values are parsed for a particular log parser. Prior to version 11.2, you could only view the out-of-the-box log parser rules.

About Log Parser Rules

Parsers are described within their XML files. Each log parser has an XML file that contains rules on how to parse messages for that parser. The out-of-the-box rules are contained within these XML files. For details, see the Log Parser Customization topic in the Link space for NetWitness Content.

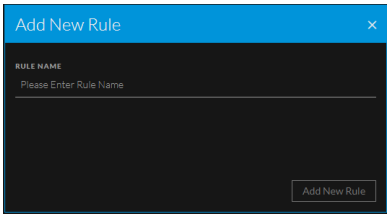
Custom Log Parser Rules

When you create a new log parser rule, it is saved to another XML definition file for the parser. These files are known as token files. This is important, since the out-of-the-box rules are overwritten if you update the parser through Live, but any custom log parser rules are not overwritten, since Live does not update the token files for log parsers.

To create a custom log parser rule:

1. In the NetWitness UI, navigate to (Configure) > **Log Parser Rules**.
2. From the **Log Parsers** pane, select a log parser, then **Dynamic Rules**.
3. From the **Rules** pane, click **Add Rule**.

   The Add Rules dialog box is displayed.



IMPORTANT: If you click outside of the Add Rule dialog box before you save your rule, your changes will be lost.

4. Enter a name for the new rule, and click **Add New Rule**.
5. Add at least one meta key and a value to match, in order to create a valid rule.
6. Click **Save** to save your new rule.

   This updates the definition file in the file system. It *does not* deploy the changes.
7. To deploy your changes to all of your Decoders, click **Deploy**.

Guidelines for Custom Rules

When you are creating a custom rule, keep in mind the following:

- For the list of tokens that match strings from the log file, very short tokens are not useful. For example, a one- or two-character string can match more items than desired.
- Remember to add the delimiter (especially if it is a space) as part of the token. For example "domain=" or "email ".
- When constructing regular expressions, the more complexity you add, the more performance overhead added to the system to compare against the rule.
- To see examples of good tokens and regular expressions, examine the rules that are provided for the default log parser.

Default Log Parser and Log Parser Rules

Note: The information in this topic applies to NetWitness Version 11.1 and later.

This tab displays information about pattern matching and rules for the parsers in your system. The features on this tab apply to all log parsers, including the Default Log Parser.

Default Log Parser

The NetWitness default log parser is used to parse logs coming from the Log Decoder that do not match any of the configured log parsers. This default parser parses these logs by using a default set of rules and tokens.

You can view the default log parser and its details by going to **Admin > Event Sources > Log Parser Rules** and selecting **default** from the Log Parsers panel.

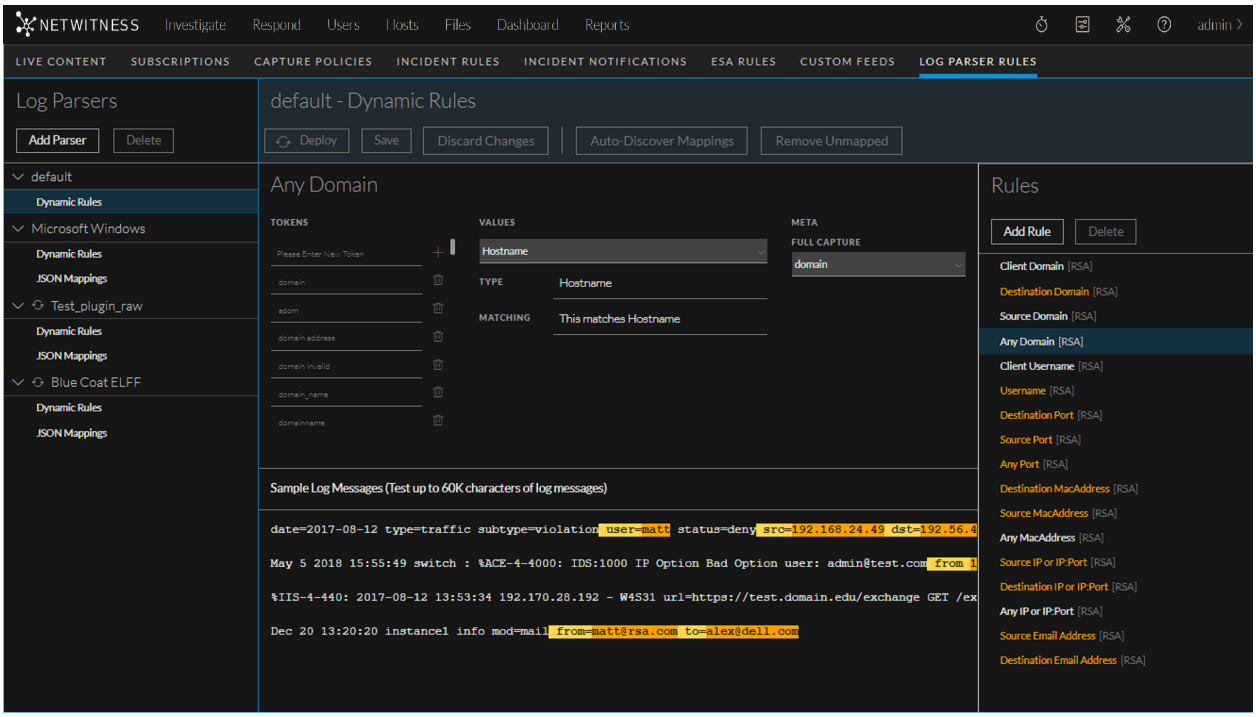
Note: If you do not see the default log parser and its rules, you might need to go to Live and deploy the Content to your log decoders. Additionally, you must have at least one Log Decoder at version 11.2 to view the default log parser.

You can view the default log parser and its details, depending on your version:

- For NetWitness version 11.1, go to **Admin > Event Sources > Log Parser Rules**, then select **default** from the Log Parsers panel.
- For NetWitness version 11.2 and later, go to **Configure > Log Parser Rules**, then select **default** from the Log Parsers panel.

Note: The list of log parsers is based on the first Log Decoder that is installed or registered by the Orchestration Server. If you have more than one Log Decoder, this tab only lists log parsers that have been configured on the first one.

This is a view of the Log Parser Rules tab, showing the **Default Log Parser** and **Any Domain** rule selected:



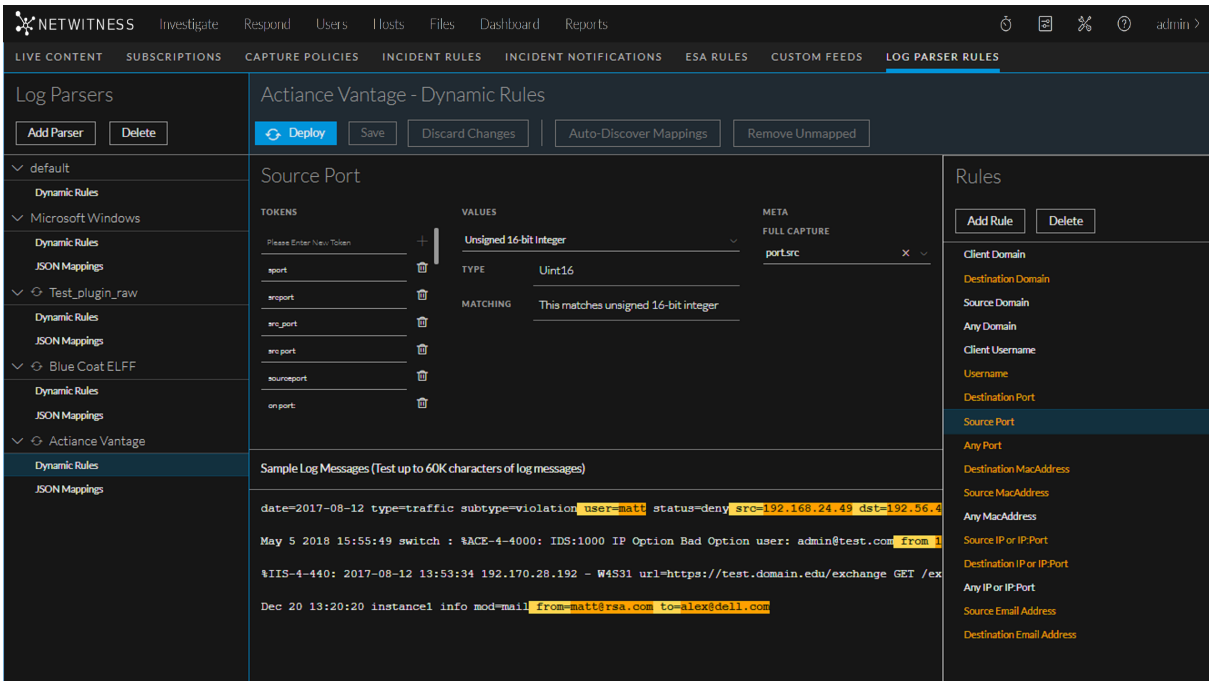
The Parser Rules Tab topic describes the items available for the Log Parsers tab.

Highlight Matching Patterns

You can paste logs into the Log Messages text box, and the system highlights the matching literals and patterns for the rules for the selected event source type. Use this feature to confirm that the parser is behaving as expected.

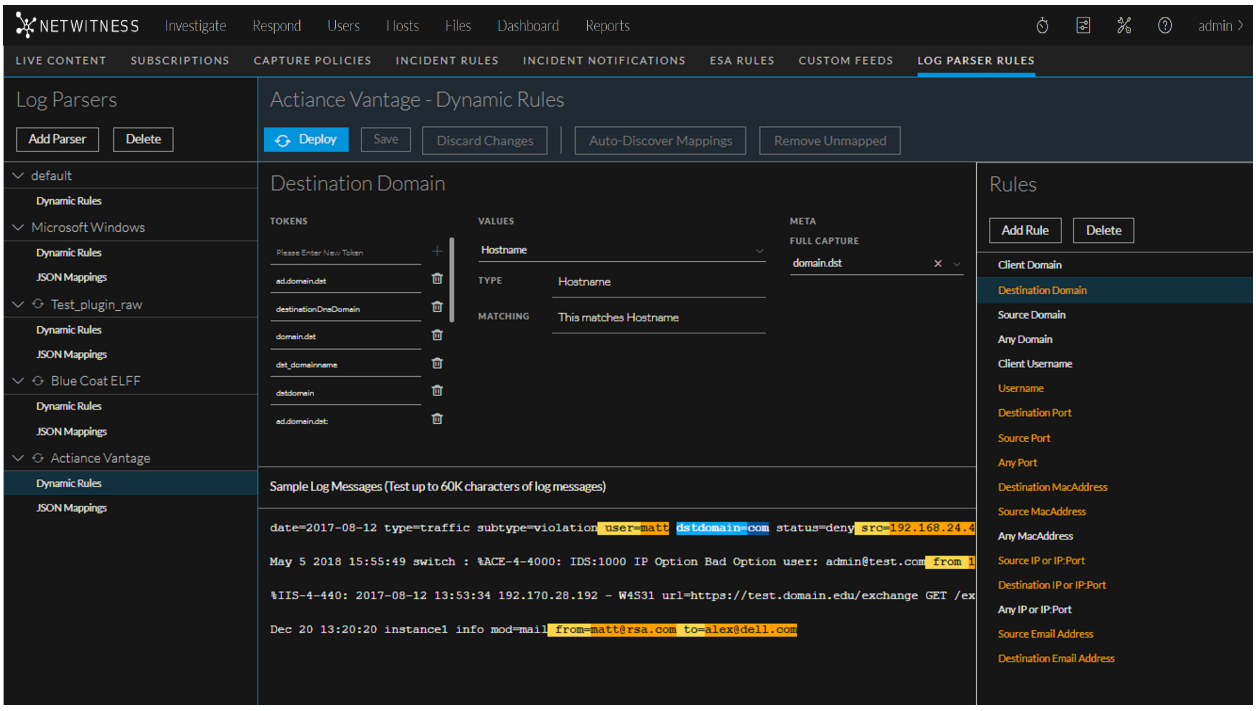
1. In the NetWitness UI, navigate as follows, depending on your version:
   - For NetWitness version 11.1, go to **Admin > Event Sources > Log Parser Rules**.
   - For NetWitness version 11.2 and later, go to  **(Configure) > Log Parser Rules**.
2. From the **Log Parsers** pane, select the **Dynamic Rules** entry for a log parser.
3. From the **Rules** pane, select a rule.

   For example, this screen shows the **Source Port** rule for the **Actiance Vantage** log parser:



4. Add text or paste in a sample log message.

Strings that match tokens for the selected rule are highlighted in blue. Strings that match other rules for the parser (and the rules themselves) are highlighted in orange.



For example, in the previous screen, note:

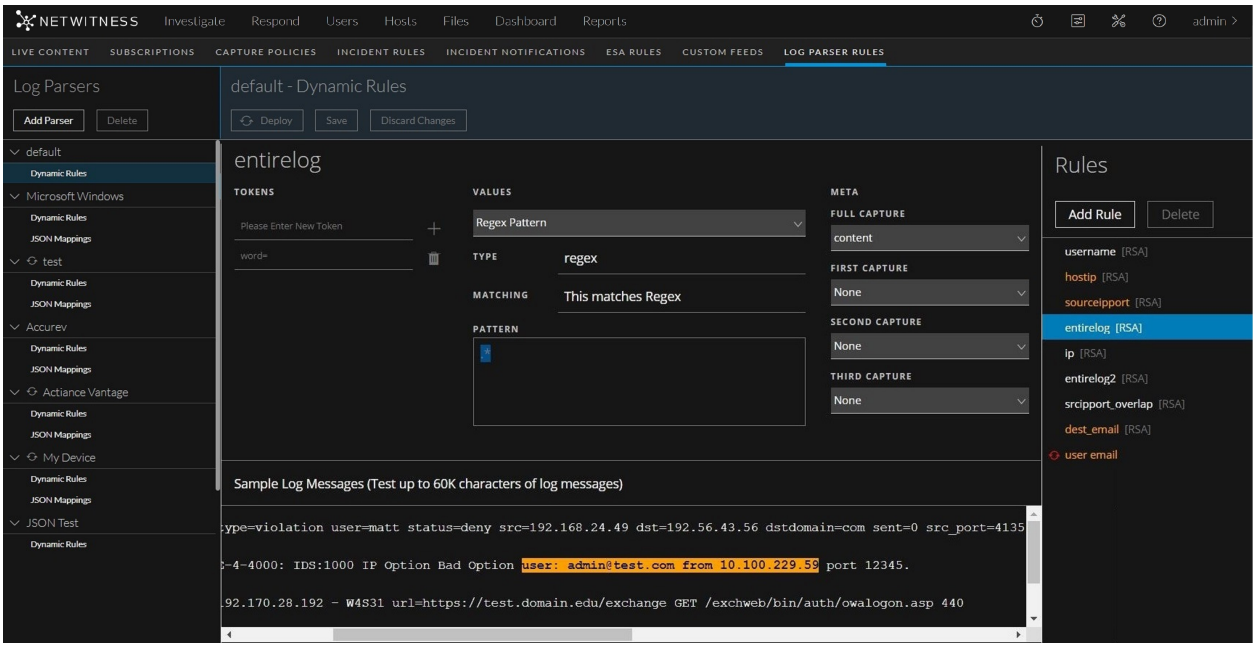
- The destination domain address, matching the **dstdomain** token, is highlighted in blue. The token is in dark blue, and the matching string is highlighted in light blue. This is because the **Destination Domain** is the currently selected Rule.
- The strings highlighted in orange match tokens for rules for **Username**, **Source IP or IP:Port**, **Destination IP or IP:Port**, **Source Port**, **Source Email Address**, and **Destination Email Address**. This is because they are in rules for the default parser that are not currently selected.

Highlight Overlapping Patterns

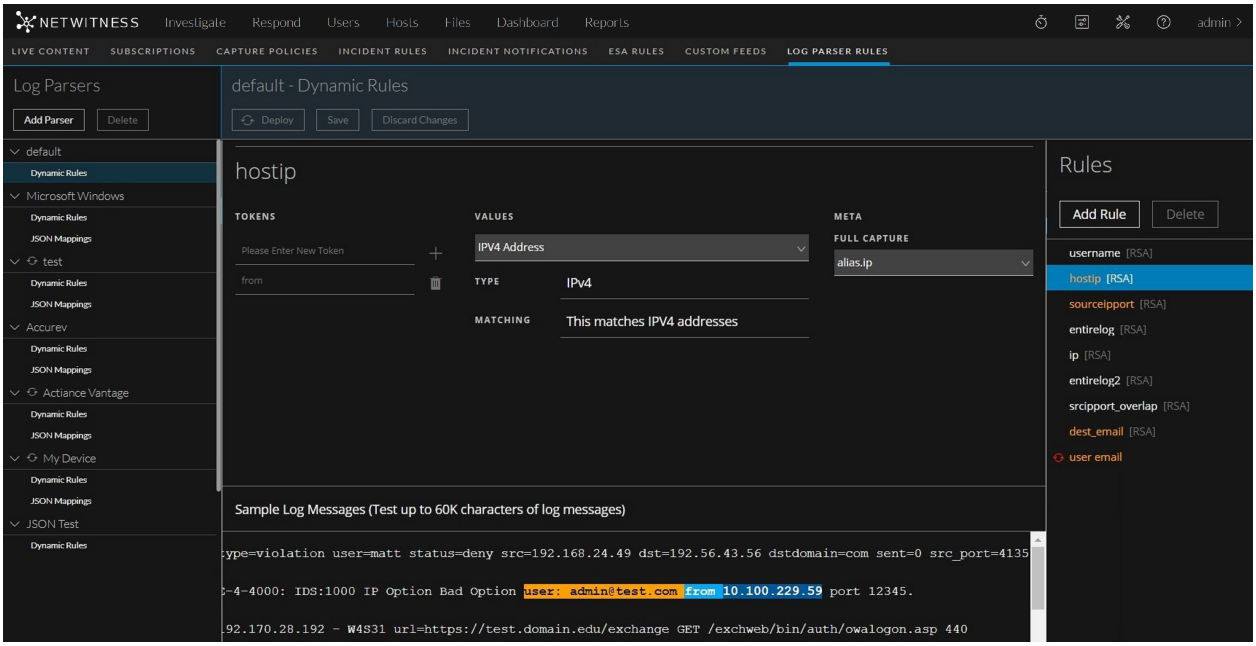
When you have patterns that overlap rules (that is, one pattern matches more than one rule), the behavior is as follows:

- The pattern is displayed in a single color (yellow)
- When you select one of the matching rules, the exactly-matched pattern is displayed in light and dark blue

For example, the pattern `user: admin@test.com from 10.100.229.59` matches several rules.



When you select the **hostip** rule, the highlighting that matches only this rule is shown in dark and light blue.



Use Cases

This topic describes the procedures you use to either on board a new event source, or to extend the parsing capabilities for an existing log parser.

Use Case 1: On Board a New Event Source

In this case, a customer has an event source and wants to add it into the NetWitness. Perform the following tasks:

I. For your event source, get examples of the logs.

II. In the **CONFIGURE > Log Parser Rules** view, add the Log Parser.

III. From your sample logs, paste applicable sections into the Sample Log Messages section of the **Log Parser Rules** screen.

IV. Use the sample area to understand which items are being parsed by the current parser, and note the items that are not being parsed.

V. For anything that is not currently being parsed, add rules.

- If the new rules apply to all parsers, you can add them to the Default parser.
- If not, add them only to the new log parser you are creating.

VI. Save the new rules, and deploy them to all Log Decoders.

VII. Map the IP address for the newly added event source to the newly-created log parser. For details, see "Acknowledging and Mapping Event Sources" in the *Event Source Management User Guide*.

Use Case 2: Modify an Existing Parser

In this case, a customer wants to parse some items from the logs that are not currently being parsed by the existing log parser. Perform the following tasks:

I. For your event source, get examples of the logs.

II. In the **CONFIGURE > Log Parser Rules** view, add the Log Parser.

III. From your sample logs, paste applicable sections into the Sample Log Messages section of the **Log Parser Rules** screen.

IV. Use the sample area to understand which items are being parsed by the current parser, and note the items that are not being parsed.

V. For anything that is not currently being parsed, add rules.

VI. Save the new rules, and deploy them to all Log Decoders.

For a detailed walk through of some of the steps in these use cases, see Extend an Existing Log Parser Example.

Extend an Existing Log Parser Example

Note: The information in this topic applies to NetWitness Version 11.2 and later.

Parse Rules can be used to parse unknown logs from existing devices. If a log is identified as a particular type (**device.type** is populated), and is not already being parsed (**msg.id** is not populated), then Parse Rules can be added to pull out relevant data from these logs.

If the neither **device.type** nor **msg.id** are populated for the logs from an existing device, then you need to map the device before Parse Rules can be processed against these logs.

Note: If a log message is already being parsed (**msg.id** is populated) then Parse Rules will not be processed against that log.

Task Overview

In this example, a customer wants to parse some items from the logs that are not currently being parsed by the existing log parser. Perform the following tasks:

I. For your event source, get examples of the logs.

II. In the CONFIGURE > Log Parser Rules view, Add the Log Parser

III. From your sample logs, paste applicable sections into the Sample Log Messages section of the **Log Parser Rules** screen.

IV. Use the sample area to understand which items are being parsed by the current parser, and note the items that are not being parsed.

V. For anything that is not currently being parsed, add rules as described in Add Rules and Deploy.

VI. Save the new rules, and deploy them to all Log Decoders.

Notes

Note: All the procedures in the topic use the CONFIGURE > Log Parser Rules view.

In the Log Parser Rules tab, you may see the Refresh icon () next to an item. This indicates that the item has undeployed changes.

Add the Log Parser

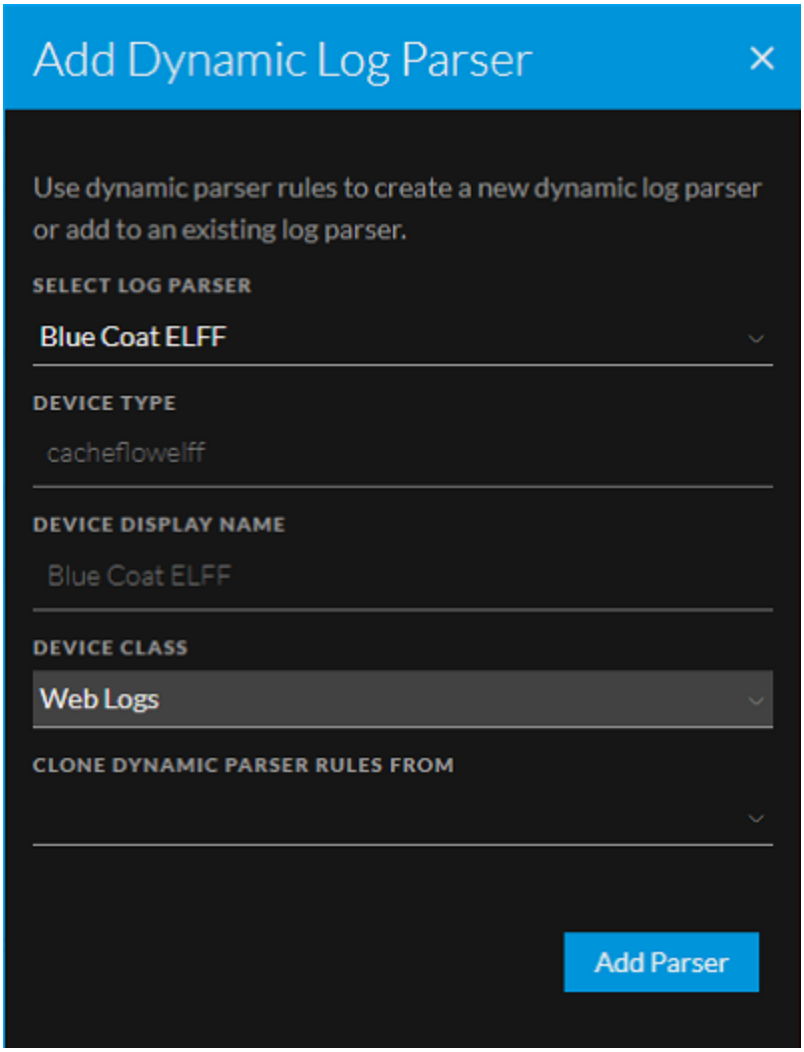
The first step in the process is to add a log parser, based on an existing log parser that you want to customize.

To add a log parser

1. In the NetWitness menu, navigate to **(Configure) > Log Parser Rules**.
2. In the Log Parsers panel, click **Add Parser**.

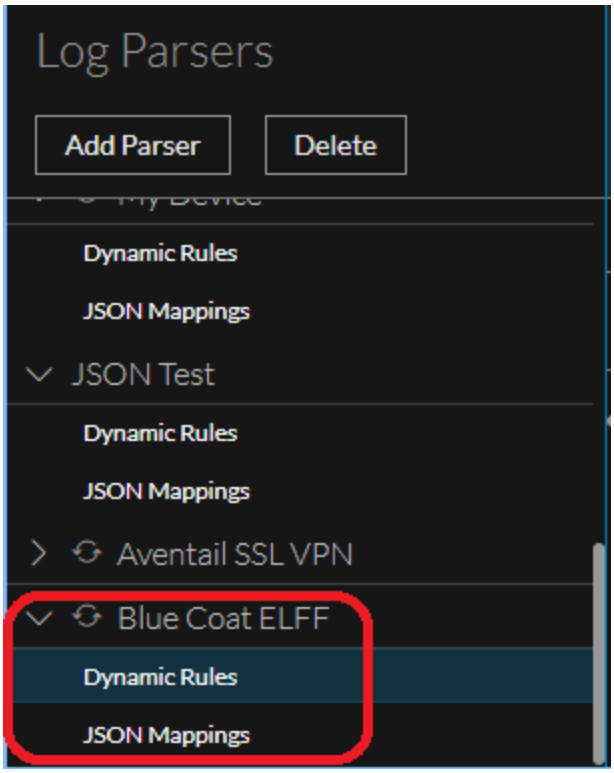
   The Add Dynamic Log Parser dialog box is displayed.
3. In the **SELECT LOG PARSER** field, select the existing parser to extend. In this example, we use Blue coat ELFF.

   Note: Values for Device Type, Device Display Name, and Device Class are filled in automatically, based on the log parser you selected.



   You can clone the rules from any of your existing parsers, including the **default** parser. For simplicity, in this example we leave this field blank: thus, only the rules we create are added to the new parser.
4. Click **Add Parser** to create the new parser.

The new parser is listed in the Log Parsers panel. Note the symbol next to the new parser—this indicates that your changes have not yet been saved.



About Custom Rules

When you create a new log parser rule, it is saved to an XML definition file for the parser. These files are known as token files. This is important, because the built-in rules are overwritten if you update the parser through Live, but any custom log parser rules are not overwritten, since Live does not update the token files for log parsers.

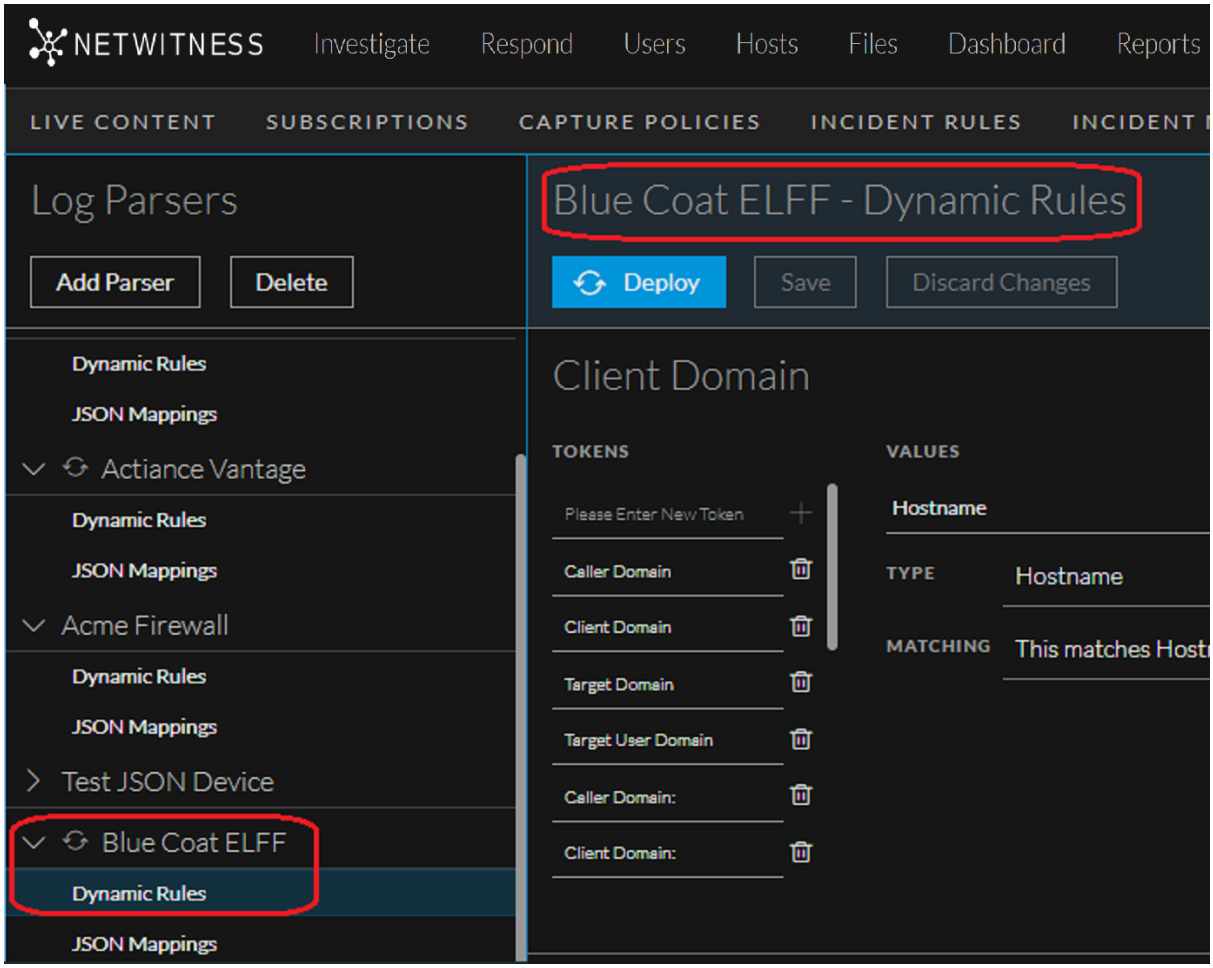
Add Rules and Deploy

After you have added the parser, the next step is to add one or more rules.

For example, you know that your log messages have some email addresses that follow a "source_mail" string. You could add the following rule to parse these strings:

IMPORTANT: If you click on another parser in the **Log Parsers** panel, before you save your rule, your changes are lost.

1. Make sure the Dynamic Rules entry for the Blue Coat ELFF parser is selected.

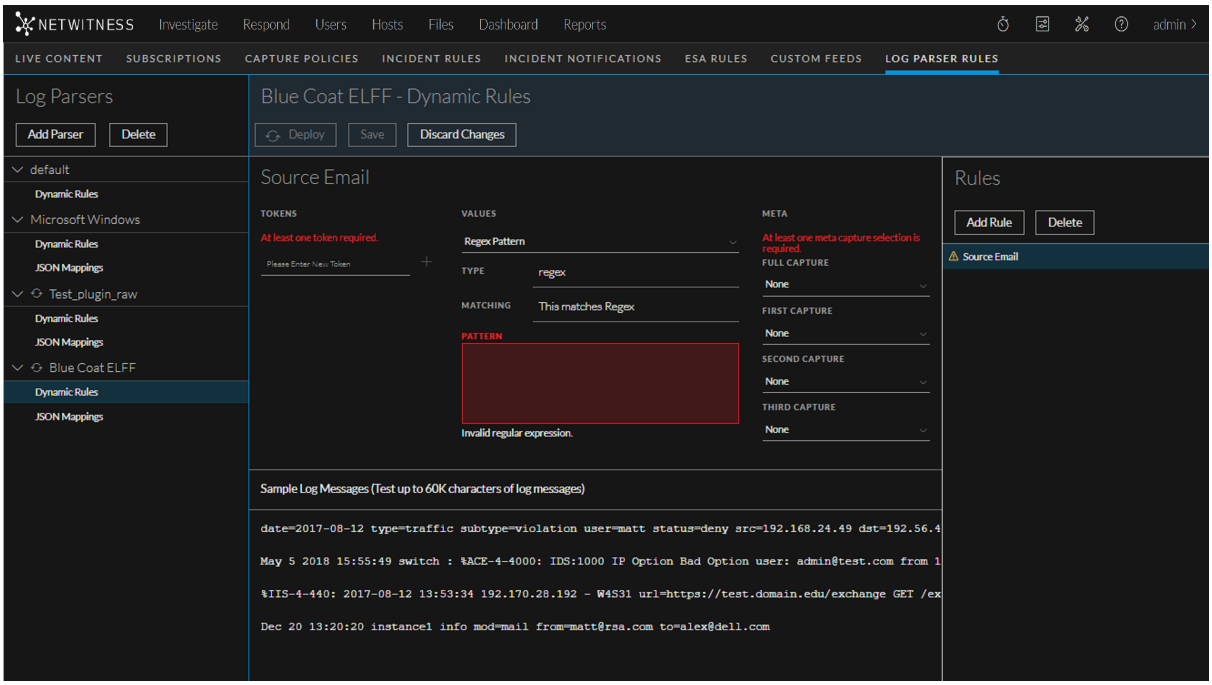


2. In the Rules panel, click **Add Rule**.

   The Add New Rule dialog box is displayed.

3. Enter a name for the rule, and click **Add New Rule**.

   The center panel is updated to reflect that you are working on a new rule.



In the TOKENS section, enter a string for the token that you want to match, then click +.

In this example, we entered **email** .

Note: Make sure to include a delimiter for your token. For example, in this case, the token consists of 6 characters: the string "email," and then a space. Some tokens might use a colon, semicolon, or some other character as the delimiter, but it can be easy to forget to add the space character when that is the delimiter.

4. You can enter more tokens, or continue to add values.
5. In the VALUES section, choose the value for the rule. If you choose to match a Regex Pattern, you need to enter the pattern in the PATTERN field. Other values do not require any options.

   In this example, we selected **Email Address**.

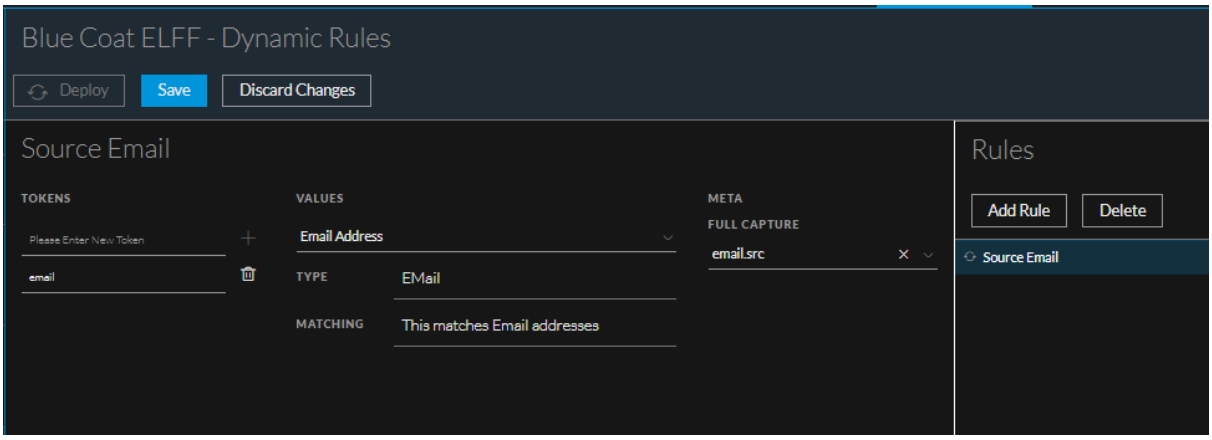
6. In the META section, click  to select a meta key to which the rule stores its information. Some notes:

   - Enter characters to filter the list of available meta keys.
   - For Regex values, you can select "pieces" of the value, and store each piece to its own meta key.

   Note: If any new meta keys are added to the Log Decoder, they do not appear in the list of Meta immediately. They appear automatically after 24 hours, or you can restart the **content server** service to view them.

   In this example, we selected the **email.src** meta key.

   The following image shows an example rule:



7. Click **Save** to save the rule. Repeat this procedure to continue adding rules.
8. After you have added all of your rules, click **Deploy** to deploy the new parser to your Log Decoders. Some notes about deploying rules:
   - You deploy an entire set of rules for a parser. That is, you can continue adding rules for a specific parser until you have all of your rules, and then you can deploy them all at once.
   - After you deploy a custom parser, you can no longer delete it. You can only delete parsers that you have not yet deployed.

Note: In this example, we extended an existing log parser. However, if you are creating a new log parser for a new event source, make sure to map the new log parser to the IP address of the event source, as described in "Acknowledging and Mapping Event Sources" in the *Event Source Management User Guide*.

Regex Values

Custom Log Parser Rules can match regular expression patterns. If you select a Regex pattern for your Value, you can capture the entire matched token, or sections of it:

- Full capture: the entire matched string is stored to your selected meta key.
- First capture: the first portion of the string, up to the period character, is stored to the meta key.
- Second capture: the second portion of the string, starting after the first period character, is stored to the meta key.
- Third capture: the third portion of the string, starting after the second period character, is stored to the meta key.

You can choose any or all four of these captures, depending on the token you are matching.

For example, we examine the **Source IP or IP:Port** rule:

- Regex Pattern: \s*(\b(?:[0-9]{1,3}\.){3}[0-9]{1,3}\b):?(\d*)
- Full capture: none

- First capture: **ip.src**
- Second capture: **port.src**
- Third capture: none
- Assume example string of "src=192.168.24.4:8080", where **src** is one of the tokens defined for this rule:
  - **192.168.24.4** is saved to the **ip.src** meta key.
  - **8080** is saved to the **port.src** meta key.

For more details, see any online reference that describes PERL regular expressions. There are many tutorials available online.

IMPORTANT: Be careful when constructing regular expressions in your custom rules. Badly constructed regular expressions could impact your performance.

Appendix A: Select the Reference Log Decoder

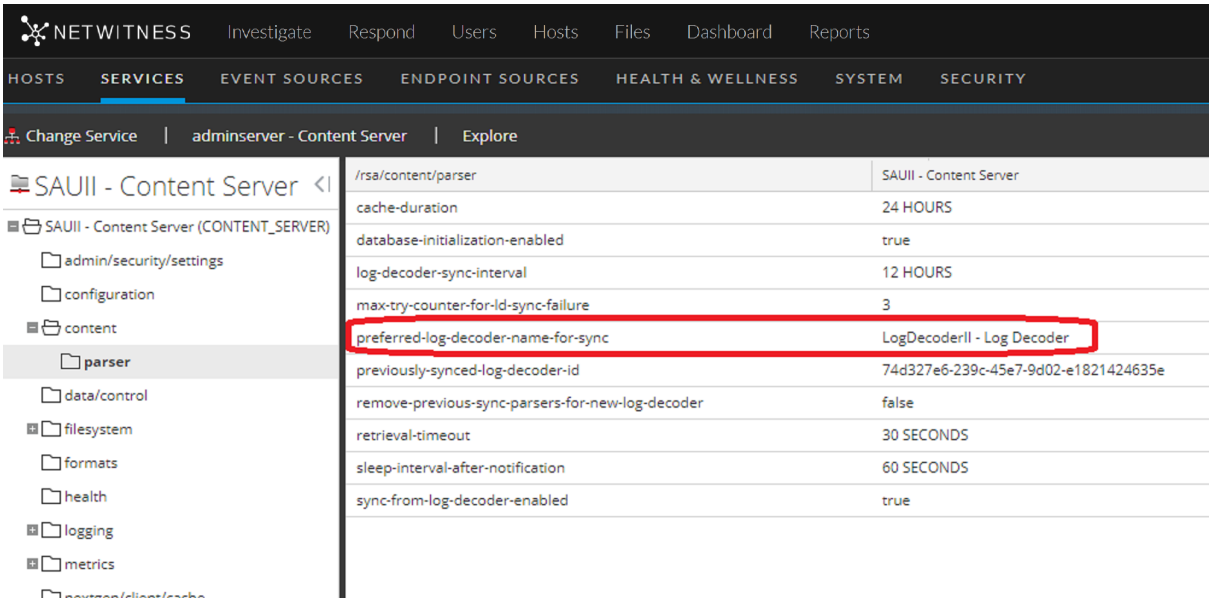
For version 11.2, NetWitness has added the ability to add log parsers and log parsing rules through the UI, using the Log Parsers view. The Log Parsers tab is populated based on your reference Log Decoder. If you have more than one Log Decoder, you can select which acts as the reference one for populating the tab in the UI. This topic describes the procedure to do so.

Note: If you have previously set a reference log decoder, make sure that it is at NetWitness version 11.5 or later to get full functionality.

To change the reference log decoder:

1. In the NetWitness UI, navigate to **(Admin) > Services**.
2. For the **Content Server**, select **View > Explore**.
3. From the left navigation panel, expand **content > parser**.
4. To set the reference log decoder, enter a value for `preferred-log-decoder-name-for-sync`.

   Enter the name listed in the **Name** column on the **ADMIN > Services** screen for your preferred log decoder.



5. The change takes effect during the next system sync, based on the `log-decoder-sync-interval`. To sync sooner, you can do either of the following:

   - To sync immediately, restart the Content Sever: in the **(Admin) > Services** view, from the **Actions** menu for the Content Server, select **> Restart**.
   - Change the `log-decoder-sync-interval` parameter from its default of 12 hours to your preferred interval. Note that the minimum value for this parameter is **1 HOUR**.

Appendix B: Move Log Parsers to Production

You may have a development or test environment where you work on new and updated log parsers and log parser rules. In this case, at some point you need to move your new and updated log parsers into your production environment. This topic describes how to do this.

To move custom log parsers and log parser rules from development to production environment:

1. On the development system, do the following:
   a. SSH to the NetWitness Server
   b. Export the log parser information by running the following command:

      ```
      mongodump --host localhost --port 27017 --db "content-server" --username
      "deploy_admin" --password "netwitness" --authenticationDatabase admin
      ```

   c. Copy the "dump" folder to your production NetWitness Server.
2. On the production system, do the following:
   a. SSH to the NetWitness Server
   b. Drop the content-server table from Mongo by running below commands in the order listed:

      ```
      mongo --username deploy_admin --password netwitness --
      authenticationDatabase admin
      use content-server
      db.logDeviceParser.drop()
      db.patternFormatType.drop()
      exit
      ```

   c. Run the following restore command:

      ```
      mongorestore --host localhost --port 27017 --db "content-server" --
      username "deploy_admin" --password "netwitness" --authenticationDatabase
      admin PATH_TO_DUMP_FOLDER
      ```

      Make sure to replace *PATH_TO_DUMP_FOLDER* with the actual path to the "dump" folder.
   d. Restart the content-server by running the following command:

      ```
      systemctl restart rsa-nw-content-server
      ```

Appendix C: Troubleshooting and Limitations

This section describes some common issues that can occur when you customize log parsers and log parser rules.

Troubleshooting

You do not see any log parsing against a newly created parser.	You may have forgotten to map the new parser. To map a parser, go to **(Admin) > Event Sources > Discovery** tab. See the "Discovery Tab" topic in the *Event Source Management Guide* for details.
Deployment fails	If you click Deploy to deploy a new or updated log parser, and it fails, you should check the log for your reference log decoder. You access this log in the following location on the NetWitness Server: `/var/log/netwitness/content-server/content-server.log`

Delete a Log Parser Manually

If you have any issues when you attempt to remove a Log Parser through the UI, you can manually delete a log parser by using **NwConsole**.

To delete a log parser that has been deployed:

1. Access the NetWitness Console, using the **NwConsole** command. For details, see "Access NwConsole and Help" in the *NwConsole User Guide*.
2. Run the following command:

   ```
   [localhost:50002] /decoder/parsers> send . delete file=filename.xml
   type=device
   ```

   where ***filename*** is the name of the XML file for the log parser. For example, to delete the log parser for Oracle Access Manager, run the following command:

   ```
   [localhost:50002] /decoder/parsers> send . delete file=oracleam.xml
   type=device
   ```

Notes about the log parser filename:

- Log parser files are located on the Log Decoder in the following path:

  `/etc/netwitness/ng/envision/etc/devices`

- Each log parser has its own sub-folder. For example, the Cisco ASA parser files are in the following folder:

  `/etc/netwitness/ng/envision/etc/devices/ciscoasa`

- Some log parser file names begin with **v20_**, while others do not—the only way to tell is by examining the `devices` folders. For Cisco ASA, the log parser file name is **v20_ ciscoasamsg.xml**. However, in the previous command, when you specify the filename, do ***not*** use the **v20_** prefix.

NwLogPlayer

NwLogPlayer is a troubleshooting tool that simulates syslog traffic. `NwLogPlayer.exe` is a command line utility located on the Log Decoder host in `/usr/bin`.

At the command line, type `nwlogplayer.exe -h` to list the available options, as reproduced here:

Option	Description
`--priority arg`	set log priority level
`-h [ --help ]`	show this message
`-f [ --file ] arg (=stdin)`	input message; defaults to **stdin**
`-d [dir ] arg`	input directory
`-s [ --server ] arg (=localhost)`	remote server; defaults to **localhost**
`-p [ --port ] arg (=514)`	remote port; defaults to **514**
`-r [ --raw ] arg (=0)`	Determines raw mode. • 0 = add priority mark (default) • 1= File contents will be copied line by line to the server. • 3 = auto detect • 4 = enVision stream • 5 = binary object
`-m [ --memory ] arg`	Speed test mode. Read up to 1 Megabyte of messages from the file content and replays.
`--rate arg`	Number of events per second. This argument has no effect if **rate** > eps that the program can achieve in continuous mode.
`--maxcnt arg`	maximum number of messages to be sent
`-c [ --multiconn ]`	multiple connection
`-t [ --time ] arg`	simulate time stamp time; format is `yyyy-m-d-hh:mm:ss`
`-v [ --verbose ]`	If **true**, output is verbose
`--ip arg`	simulate an IP tag
`--ssl`	use SSL to connect
`--certdir arg`	OpenSSL certificate authority directory
`--clientcert arg`	use this PEM-encoded SSL client certificate
`--udp`	send in UDP

Limitations

Please note the following limitations when using the Log Parser Rules tab:

- **Log Decoder must be at version 11.2:** For the functionality in the Log Parser Rules tab to work, your installation must have at least one Log Decoder running NetWitness version 11.2 or later.
- **JSON Mapping: Log Decoder must be at version 11.5:** For the JSON Mapping Beta functionality to work, your installation must have at least one Log Decoder running NetWitness version 11.5 or later.
- **Mixed Mode:** If any Log Decoders are at version 11.2 or later, and the NetWitness Server is at version 11.2 or later, the Log Decoders will have parseall rules enabled by default, and thus will begin to parse logs accordingly. However, the 11.2 NetWitness Server does not support Log Decoders with versions less than 11.2, so the Log Parser Rules tab in the UI stays blank.
- **Meta key fields list refresh:** If any new meta keys are added to the Log Decoder, they do not appear in the list of Meta in the Log Parser Rules tab immediately. They appear automatically after 24 hours, or you can restart the **content server** service to view them.
- **Field Restrictions:** Note the following field restrictions:
  - **Rule name** must be 64 characters or fewer.
  - **Parser Name** must be between 3 and 30 alphanumeric characters (including underscores), and must not match the name of any existing log parsers.
  - **Parser Display Name** must be 64 characters or fewer, and cannot match any other parser display name.
  - **Regex Expression** must be 1-255 characters, and a valid regex (closed capture list allowed).
  - **Tags** cannot be duplicates.
- **Deploy only to 11.2 Log Decoders:** The Deploy operation only deploys log parsers to version 11.2 or later Log Decoders.
- **Cannot Remove Deployed Parsers:** Once deployed, you cannot delete a log parser using the UI.
- **See log for errors:** Refer to content-server logs for more details on deploy failure details and log decoder names.

Parser Rules Tab

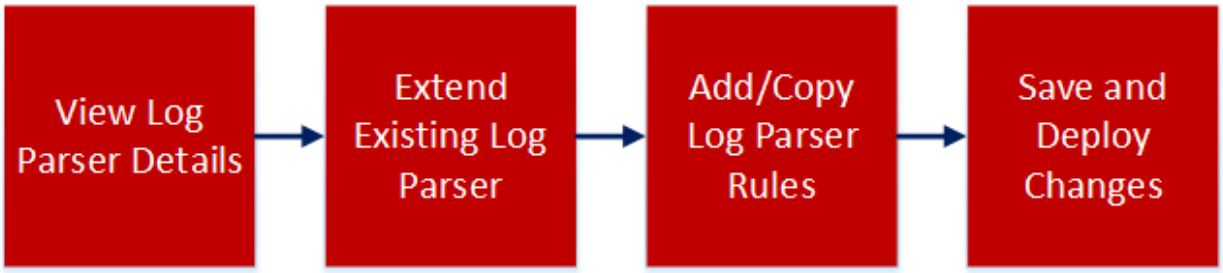
Note: The information in this topic applies to NetWitness Version 11.1 and later.

This tab contains details about the rules for the default log parser, as well as any other custom rules and log parsers that have been defined.

To access this tab, go to  **(Configure) > Log Parser Rules**.

Workflow

This workflow shows processes available from the Log Parser Rules view.



What do you want to do?

Role	I want to...	Documentation
Administrator	*View log parser rules.	Default Log Parser and Log Parser Rules
Administrator	*Add, edit or delete a log parser rule (version 11.2 and later)	Add or Delete a Log Parser Rule
Administrator	*Add or remove a log parser (version 11.2 and later)	Add or Delete a Log Parser

*You can perform this task here.

Related Topics

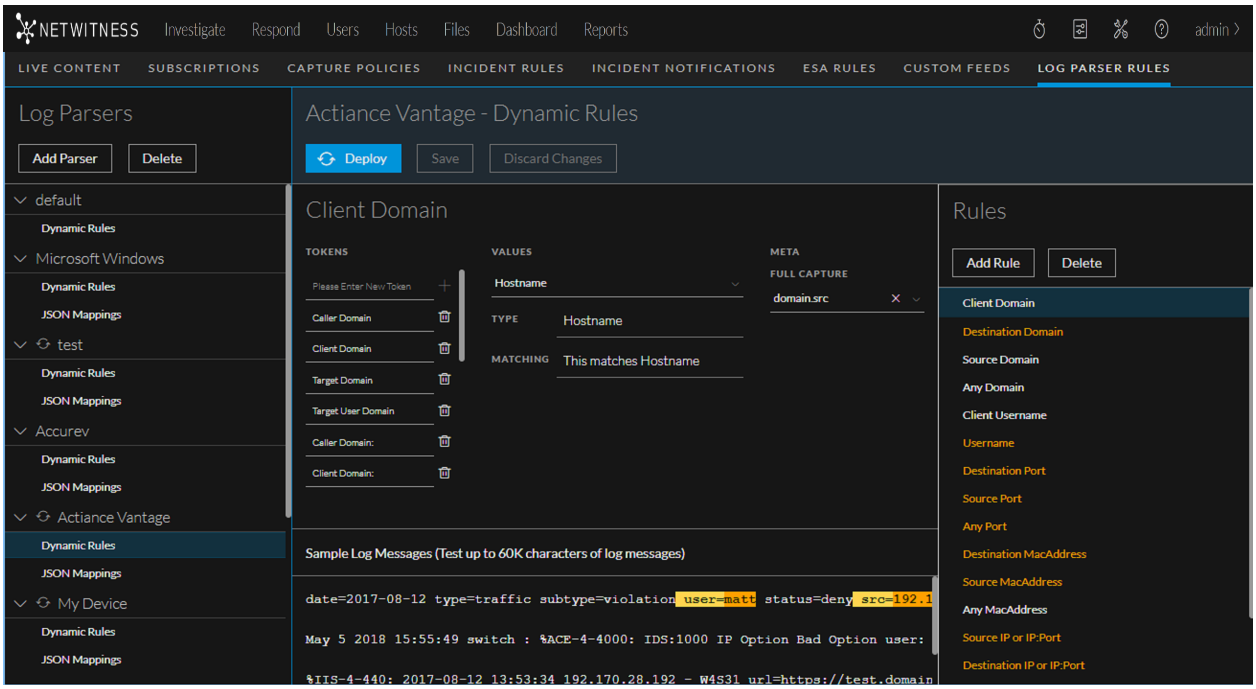
Default Log Parser and Log Parser Rules

Quick Look

Note: The list of log parsers is based on the first Log Decoder that is installed or registered by the Orchestration Server. If you have more than one Log Decoder, this tab only lists log parsers that have been configured on the first one.

The Log Parser Rules tab organizes and displays information about the configured log parsers in your system.

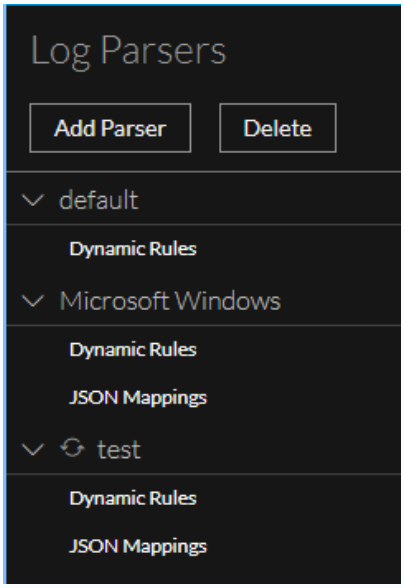
If you select the Dynamic Rules entry for a parser, the screen displays the dynamic rules:



Note: The icon indicates that there is some uncompleted or unsaved work. For example, next to a parser name, it means that the parser has not yet been deployed.

If you select JSON Mappings, the screen displays JSON mappings for the parser.

Log Parsers Panel

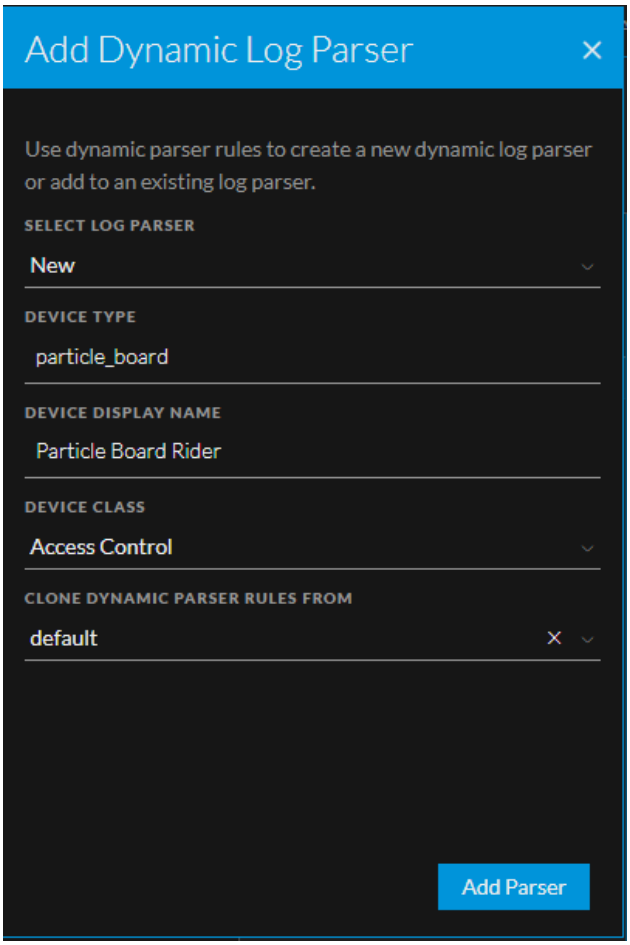


The Log Parsers Panel lists the configured log parsers.

- Until you add rules to existing XML parsers on your reference Log Decoder, (or add a new, custom log parser) only the **default** parser is listed here.
- Select a specific log parser to view its details in the Details and Rules panels.
- Click **Add Parser** to open the Add Dynamic Log Parser dialog box.
- Click **Delete** to delete a log parser.

IMPORTANT: Once you deploy a log parser, you can no longer delete it through this interface. The **Delete** button is not available for deployed parsers. To manually delete a log parser, see Add or Delete a Log Parser.

The Add Dynamic Log Parser dialog box allows you to add a custom log parser.



When you are adding a log parser, the following parameters are available.

Field	Details
SELECT LOG PARSER	Select NEW, or choose an existing log parser. By choosing an existing log parser, you can add rules to that parser, essentially extending its parsing capabilities. **Note:** If you select an existing log parser, the remaining fields are auto-filled based on the values for selected log parser.
DEVICE TYPE	Enter a string to define the device type. The name must be between 3 and 30 alphanumeric characters (including underscores), and must not match the name of any existing log parsers.
DEVICE DISPLAY NAME	Enter the display name for the log parser. **Note:** The display name must be 64 characters or fewer, and must not match the name of any other device display name.
DEVICE CLASS	Select a device class.
CLONE DYNAMIC PARSER RULES FROM	Leave blank to start with no rules, or select one of the existing log parsers to clone its rules.

Dynamic Rules

If you select the Dynamic Rules entry for a parser, you see the following panes:

- Details
- Rules

Details Pane

The details pane shows the three pieces for the selected rule:

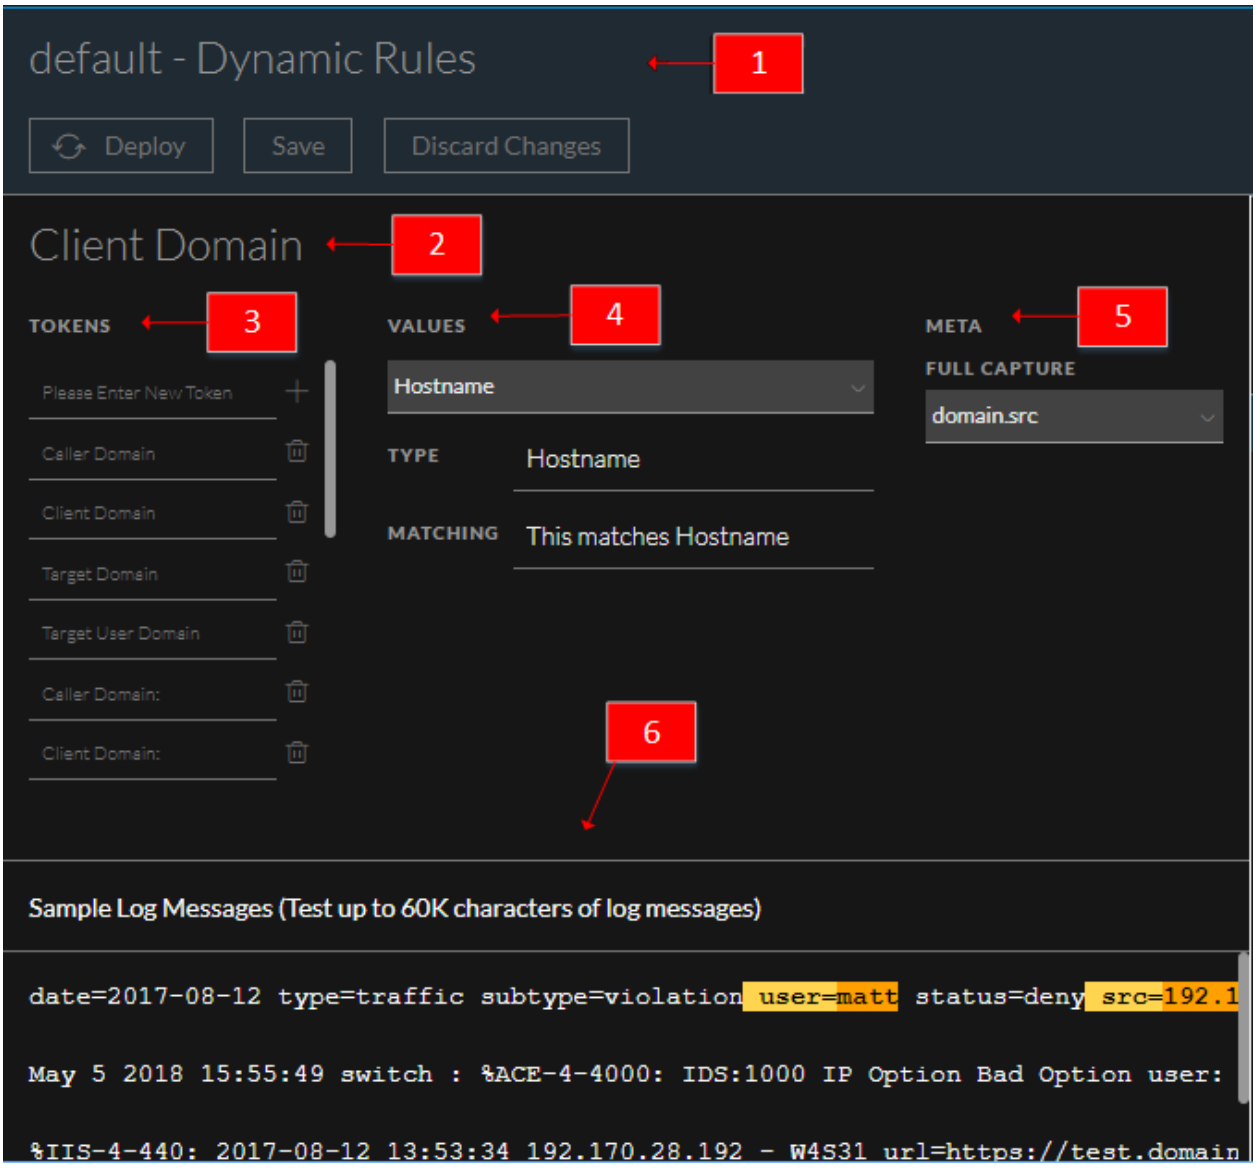
- **Tokens:** one or more tokens to match in the message. For example, the Any Port rule looks for the following strings to match against: **port** , **port:**, **port=**, and others.
- **Values:** the value that follows the token. This is a string that is captured as meta. For example, assume a log contains the following string:

  ```
  port 12345
  ```

  The Any Port rule has a token that matches "port ". When it encounters that string, it assigns the token value, "12345" to a meta key.
- **Meta:** the meta keys to which the value is mapped. For example, the Any Port rule maps the port value to the **port** meta key.

Essentially, a rule says, "when you are parsing a message, if you match one of my tokens, assign the value that follows the token to the meta key that I want it stored as."

The bottom section of the Details panel contains sample log messages, and how they would be parsed for the selected log parser.



1	Displays the name of the selected log parser, and the buttons for deploying, saving, and discarding changes. This value changes when you select a different parser.
2	Displays the name of the selected rule. This value changes when you select a different rule for this parser.
3	Displays the list of tokens defined for the selected rule.
4	Displays the type and pattern of the value matching for the selected parser. The values here are determined by the type of the selected value. You can also use the Regex option to define a custom regular expression.
5	Displays the NetWitness meta to which the selected rule maps any matched tokens. The values here are determined by the selected Rule.
6	Displays a sample log message, and highlights strings that match tokens in the selected log parser. You can edit this field, and add in your own logs to preview how the selected parser will parse your logs. **Note:** The sample section refreshes whenever a rule is changed or updated, as well as when you

paste in samples from your logs.

For example, consider the following scenario:

- The **default** parser is selected.
- The **Any Domain** rule is selected.
- The Tokens matching list displays all of the tokens that are matched when found in a log message: **Domain**, **Domain Name**, **domain**, **ADMIN_DOMAIN**, and so on.
- The Meta list displays the NetWitness meta to which the value for the token is mapped: **domain**.

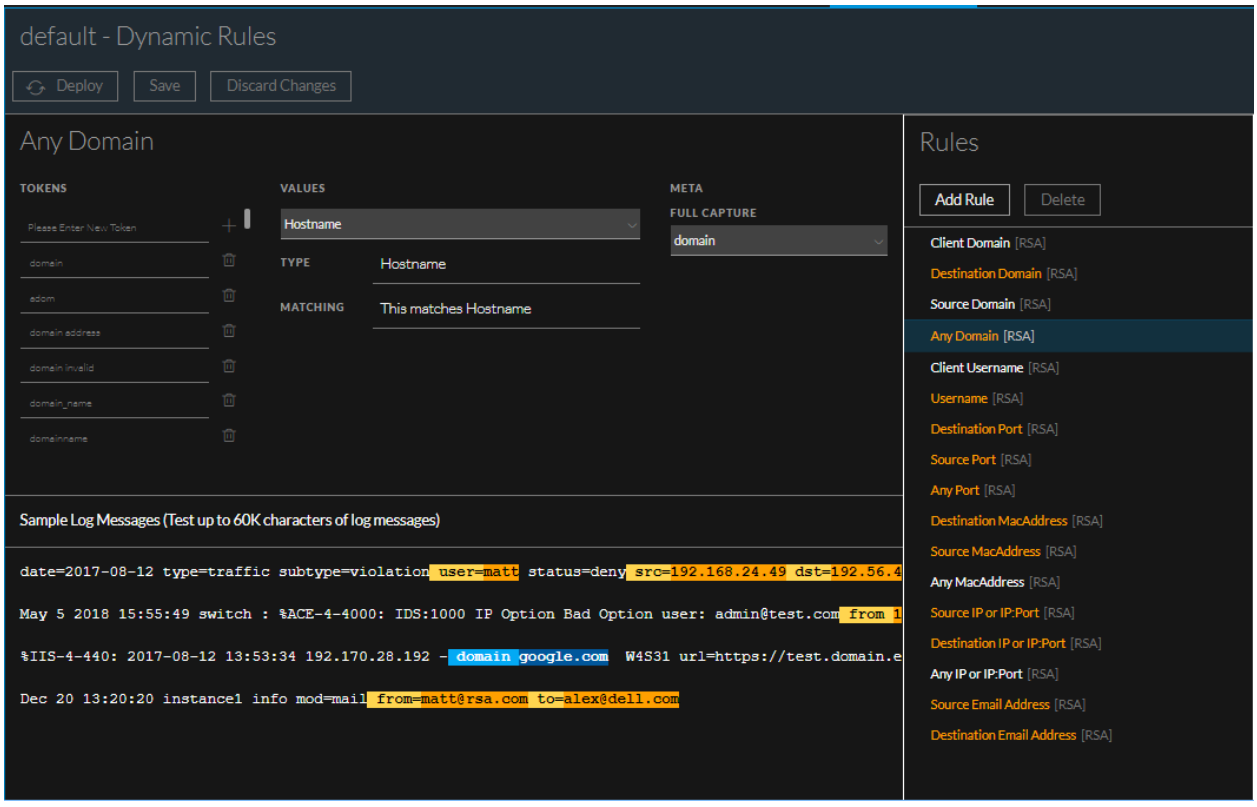
So, let's say the sample log message area has the following text:

```
Below are sample log messages:

May 5 2010 15:55:49 switch : %ACE-4-400000: IDS:1000 IP Option Bad Option
List by user admin@test.com from 10.100.229.59 to 224.0.0.22 on port 12345.

Apr 29 2010 03:15:34 pvg1-ace02: %ACE-3-251008: Health probe failed for
server 218.83.175.75:81, connectivity error: server open timeout (no SYN ACK)
domain google.com with mac 06-00-00-00-00-00.
```

In this case, the Sample Log Message area looks like this:



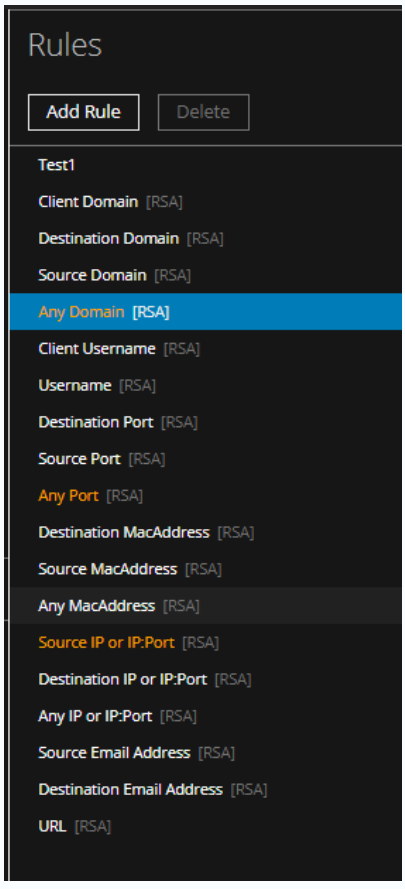
Note that some strings are highlighted, and that there are two "pairs" of highlight colors:

- Dark blue and light blue highlighting is applied to the strings that match the currently selected rule.
  - Dark Blue highlighted strings match a token in the selected rule. In this case, **domain** is the token that is matched for the **Any Domain** rule.
  - Light Blue highlighted strings are the values that correspond to the tokens in dark blue. For example, **google.com** is highlighted in light blue, because it corresponds to the **domain** token.
- Orange and yellow highlighting is applied to the strings that match rules for the current parser that are *not* currently selected.
  - Orange highlighted strings match a token in a rule that is not currently selected.
  - Yellow highlighted strings are the values that correspond to the tokens in orange. For example, the **user** token matches the **Username** rule (which is not currently selected).

In this example, the **domain** meta would be assigned a value of **google.com** for this log message, if it was parsed using the default log parser.

Rules Pane

The Rules pane displays the list of rules used by the selected log parser. When you select a rule, you change the values that are displayed in both the **Tokens** and **Values** areas of the panel.



Note the highlighted rules:

- The currently selected rule is highlighted in blue.
- Other rules that match tokens in the sample log message area are highlighted in orange.

Other notes for the Rules panel:

- RSA rules (the rules provided out-of-the-box for each log parser) are identified by **[RSA]** following the rule name.

  You can copy these rules when adding a new log parser, and then change them as needed.
- The **Delete** button is only available for custom rules; for RSA rules, it is greyed out.
- Use the **Add Rule** button to add a custom rule.

JSON Mappings

If you select the JSON Mappings entry for a parser, you see the following panes:

- Sample JSON Message
- Meta Mappings
- Mapping Details

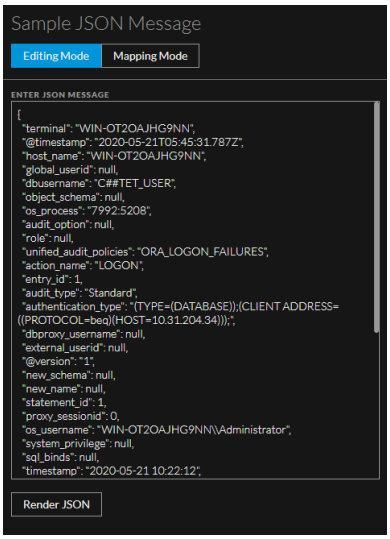
Sample JSON Message

You can enter or paste sample JSON text. Click the **Render JSON** button to automatically format the text into JSON code. If the text is not valid JSON, you receive a message and the text is not formatted.

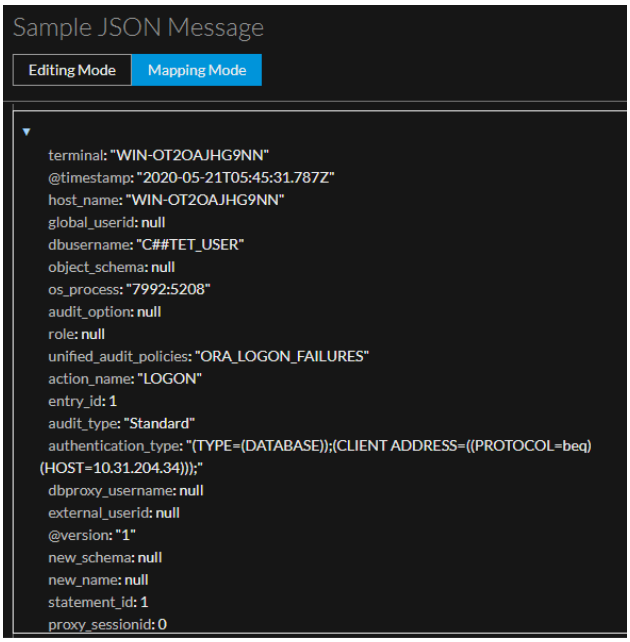
The following screen shows the Sample JSON Message area with some JSON that has been pasted in:



Since this is valid JSON, clicking **Render JSON** produces the following:

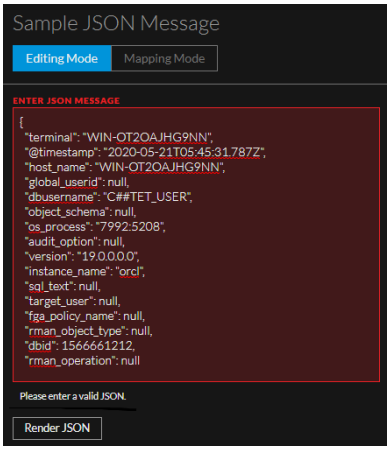


Note that you can see the tree mode of the sample JSON if you click **Mapping Mode**:



Note: You can edit the sample JSON in Editing mode, but not in Mapping mode. Mapping mode is read only.

On the other hand, if you enter text that is not valid JSON, the screen looks as follows:



Auto Discover JSON Mappings

Beginning with NetWitness version 11.5.1, you can automatically create the mappings without the need to manually enter the name and path of the mapping. For details, see .Auto Discover JSON Mappings.

When the system auto discovers mappings, the path is filled in automatically, based on the structure of the mappings. For simple name-value pairs, this is straightforward. For example, for this pair, **"host_ name": "WIN-OT2OAJHG9NN"** , the path is set to `/host_name`.

However, the rules for nested mappings and arrays are as follows:

- For a nested structure, names are separated with a forward slash (/).

```
{
  "parent": {
    "child": "value"
  }
}
```

The path is set to `/parent/child`.

- Arrays are accessed by omitting the index.

```
{
  "array": [
    "x",
    "y",
    "z"
  ]
}
```

The path is set to `/array/`.

```
{
  "array": [
    {
      "name": "value"
    }
  ]
}
```

The path is set to `/array//name`.

- Variable names are accessed by omitting them from the path.

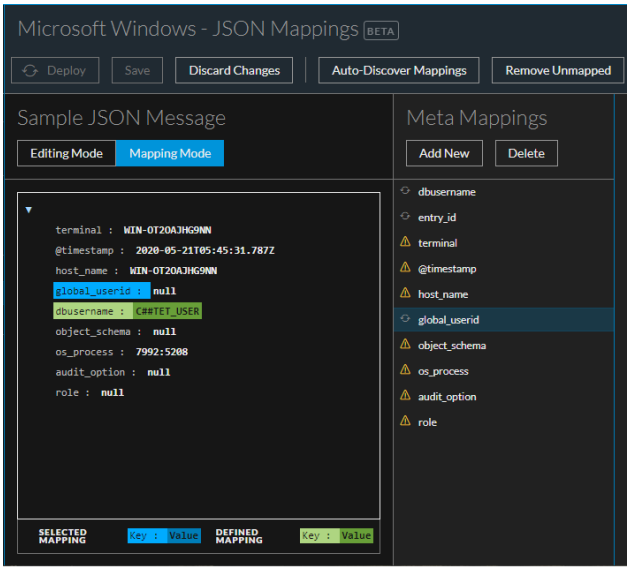
```
{
  "root": {
    "x": {
      "name": "value"
    },
    "y": {
      "name": "value"
    }
  }
}
```

The paths are set to `/root/x/name` and `/root/y/name`.

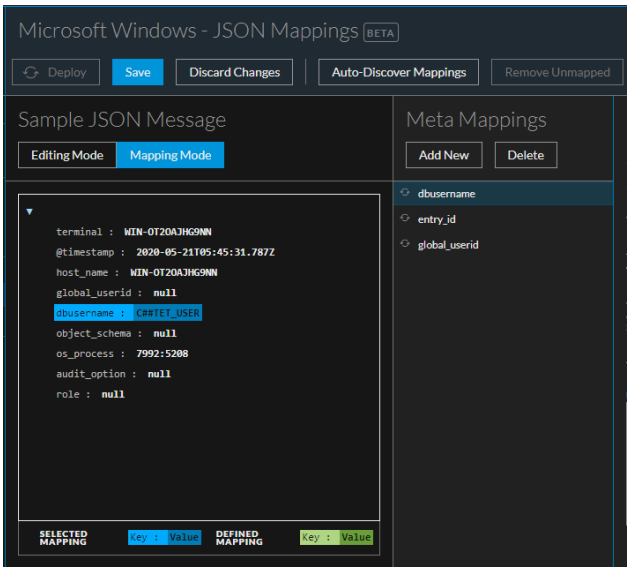
Remove Unmapped Entries

If there are mappings that you do not want to save, you can remove them. After you validate all of the mappings that you want to keep, you can click **Remove Unmapped** to remove all mappings that you have not yet validated.

For example, assume you have auto-discovered some mappings as shown here:



After you click **Remove Unmapped**, you can see that only mapped entries remain:



Note: You cannot save your work until all of the entries have either been mapped or removed from the list.

Meta Mappings

This panel lists the mappings that exist for this parser. You can add a mapping by using the **Add New** button, or delete an existing mapping by selecting it and clicking **Delete**.

Mapping Details

The Mapping Details pane displays the following information.

Field	Details
DISPLAY NAME	This name corresponds to the name displayed in the JSON Mappings panel.
PATH	The path to where the values for this portion of the log are stored.
DESCRIPTION	Optionally, you can enter a text description for this mapping.
META	Select a meta key to which this value from the log is mapped. Select a value from the drop-down menu. Optional if you choose a Value Format.
VALUE FORMAT	Choose a value format parser onto which to pass this JSON value. You can either select an existing meta or **Custom Regex Type**. If you select custom regex type, you must define the regex and capture to *fine* parse the value in the meta. Optional if you choose a Meta.
CUSTOM REGEX TYPE	Select Custom Regex Type from the Value Format drop-down, which allows you to add new custom regex type.
REGEX PATTERN	Specify a regex to identify different pieces of data contained within a JSON node value.
FIRST CAPTURE	Select a meta key that should be captured first based on the value defined in the Regex pattern.

Field	Details
ADD A CAPTURE	New capture field is added. By default, it is loaded with meta keys in the drop-down. You can add maximum of 20 captures and this option will be disabled once it reaches maximum.

Note: You need to select a meta or enter a Value Format, but you do not need to fill in values for both settings.

Disable log Parser Rules

You can disable log parser rules, so that none of them are processed by the Log Decoder. You might have your log parsers working as you like, and do not want any extra processing that you do not need.

You disable them from the reference Log Decoder.

1. Go to  **(Admin) > Services**.
2. In the **Administration Services view**, select the Decoder and  **> View > Config.**

   The Services Config view is displayed with the General tab open.
3. Under **Parsers Configuration**, look at the Config Value for **PARSERULESCAN**.

   If it is **Enabled**, log parser rules are processed. If it is **Disabled**, they are not processed.
4. If the rules are Enabled, click Enabled and select Disabled to disable the log parser rules.

   To save the changes, click **Apply**.

Add VARTYPE Support to CEF Parser

VARTYPE added in 11.0 and later versions provide type and pattern validation for variables. While the variables are parsed, NetWitness has added the ability to condition them to match a certain format.

CEF Parser format logs with **CEF:0** and will use CEF parsers without any further identification. Some CEF sources have structured data embedded in variables. This data requires further processing. To direct parse rules scanning of the data, support for scanned variables is required for the CEF parser. Scanned variables are supported during the VARTYPE implementation.

To add VARTYPE support to CEF parser:

1. Load VARTYPE from CEF parser (cef.xml and cef-custom.xml) successfully.
2. Load all existing CEF parsers and parse the data successfully.
3. Enable VARTYPE to validate the header information.
4. Enable VARTYPE to fine parse CEF variable (MetaName).
5. Ensure that VARTYPE is working fine with custom CEF parser.

NetWitness® Platform XDR

Version 12.0

Logstash and NetWitness Integration Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED

WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Contents **4**
Overview **6**
Configuration Process **8**
Install Logstash **10**
Install and Configure the NetWitness Codec **11**
Configure Logstash Output Plugins **13**
Logstash TCP Output 13
Logstash TLS Output 13
TLS with Log Decoder (or Virtual Log Collector) Verification 13
Configure the Event Source **15**
Collect Apache File Logs 15
Collect CentOS Audit Logs 15
Configure Parameters for Filebeat or Auditbeat 15
Configure Logstash Filters to Add NetWitness Meta **18**
Advanced NetWitness Configuration **19**
Grok Filter Plugin 19
Logstash Input and Filter plugins 19
Filter out unwanted logs 19
Configure heartbeat plugin to send test logs to NetWitness 19
Configure logstash to persist events in case of failure (Recommended) 20
NetWitness Codec Advanced Configuration 20
Troubleshoot Installation Issues 21
Configure NetWitness to Collect Events **22**
Linux Event Source Example **23**
Input Plugin 23
Output Plugin 23
Filter Plugin 24
Create a Pipeline 25
Build Custom JSON Parser **26**

Sample JSON Log Received on Log Decoder 26
Create the JSON Parser for a Linux Device 27
Initial Parser to Match Message ID and Device Type 27
Map Payload Contents to Datatypes 28
Parse the Message String 28
Parse an Array in JSON 29
Parse a Nested JSON Object 30
Capture Data That Has Varying Parent Key 30
The Parsed Example Event on the Log Decoder 31
Example Parser Listing 33
Deploy JSON parser 35
Reload Parsers from REST 35
Reload Parsers from NetWitness UI 35

Overview

This document is intended to provide a general overview of Logstash and NetWitness integration. The intention is to provide enough implementation detail that users can have comfort using and troubleshooting these integrations on their own.

To describe Logstash, here is some introductory text from Logstash reference documentation:

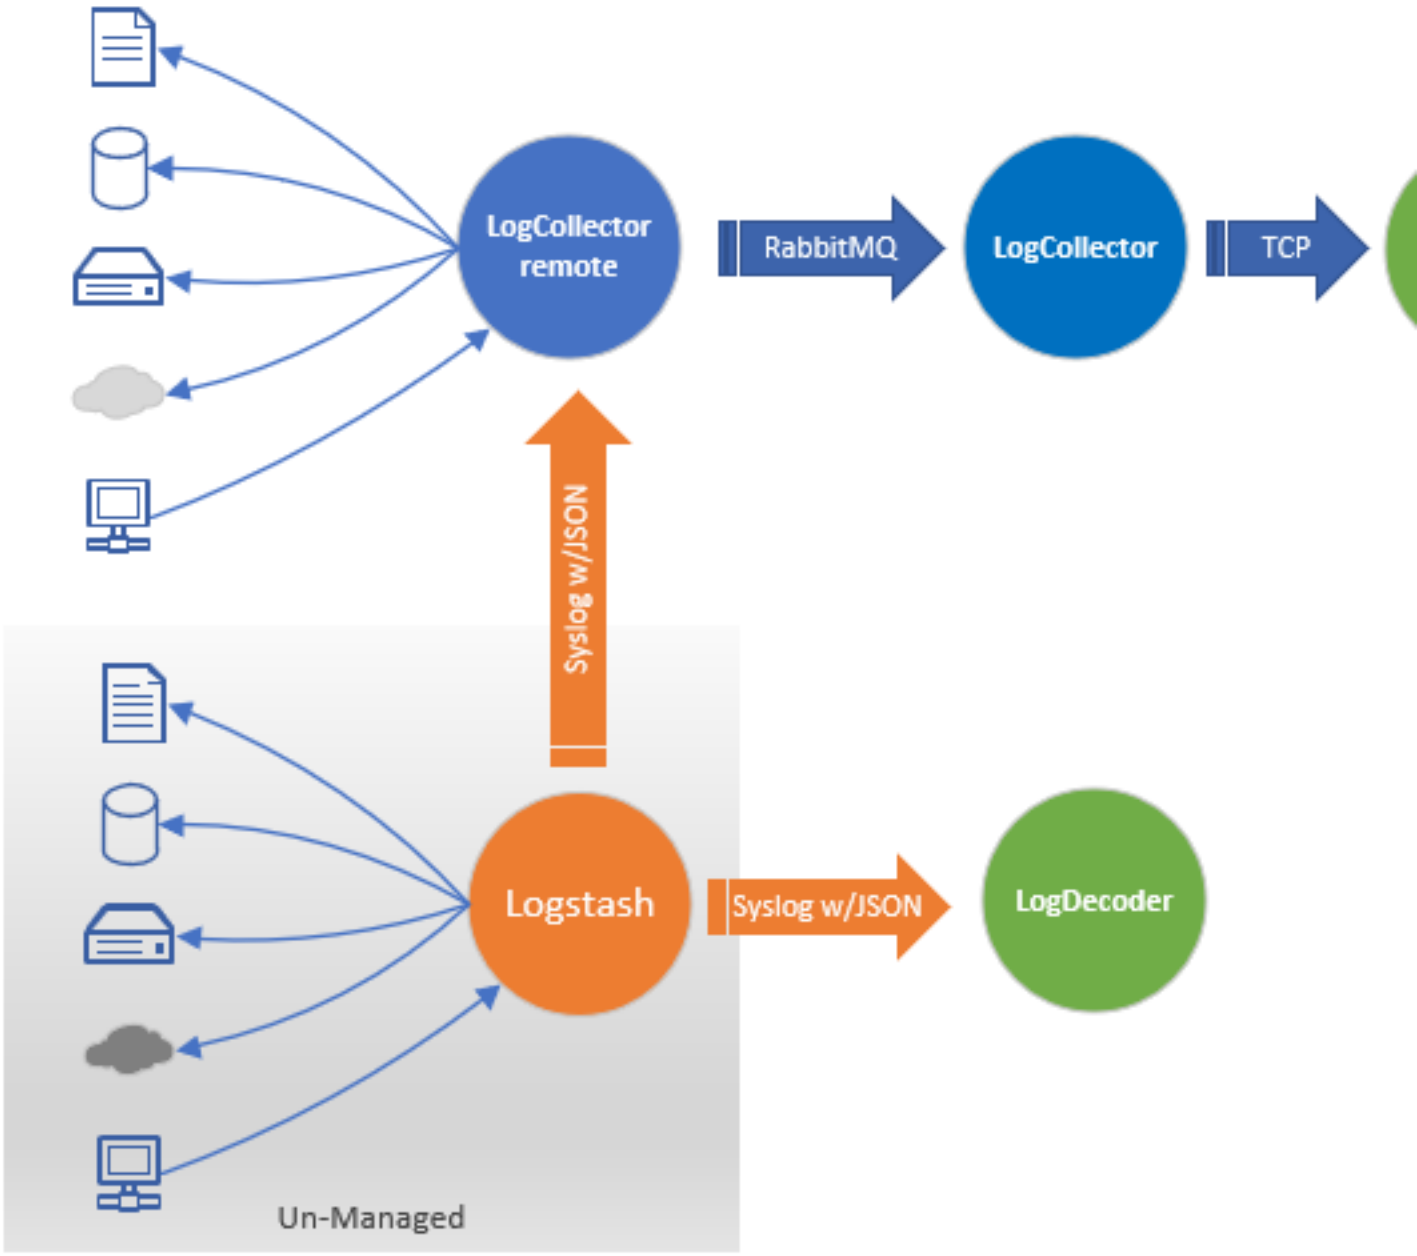
> *Logstash is an open source data collection engine with real-time pipelining capabilities. Logstash can dynamically unify data from disparate sources and normalize the data into destinations of your choice…*

From a NetWitness standpoint, there are two basic use cases:

- For customers that have an event source for which NetWitness does not already provide an integration, or if you want a customized integration that is different from the one provided by NetWitness.
- For customers that already have an existing Logstash configuration, you can use Logstash to integrate as many of your event sources as you like. Integrating your event sources should be a matter of updating the destination for where you currently send the log information: either adding NetWitness as a destination, or changing your current output destination to NetWitness.

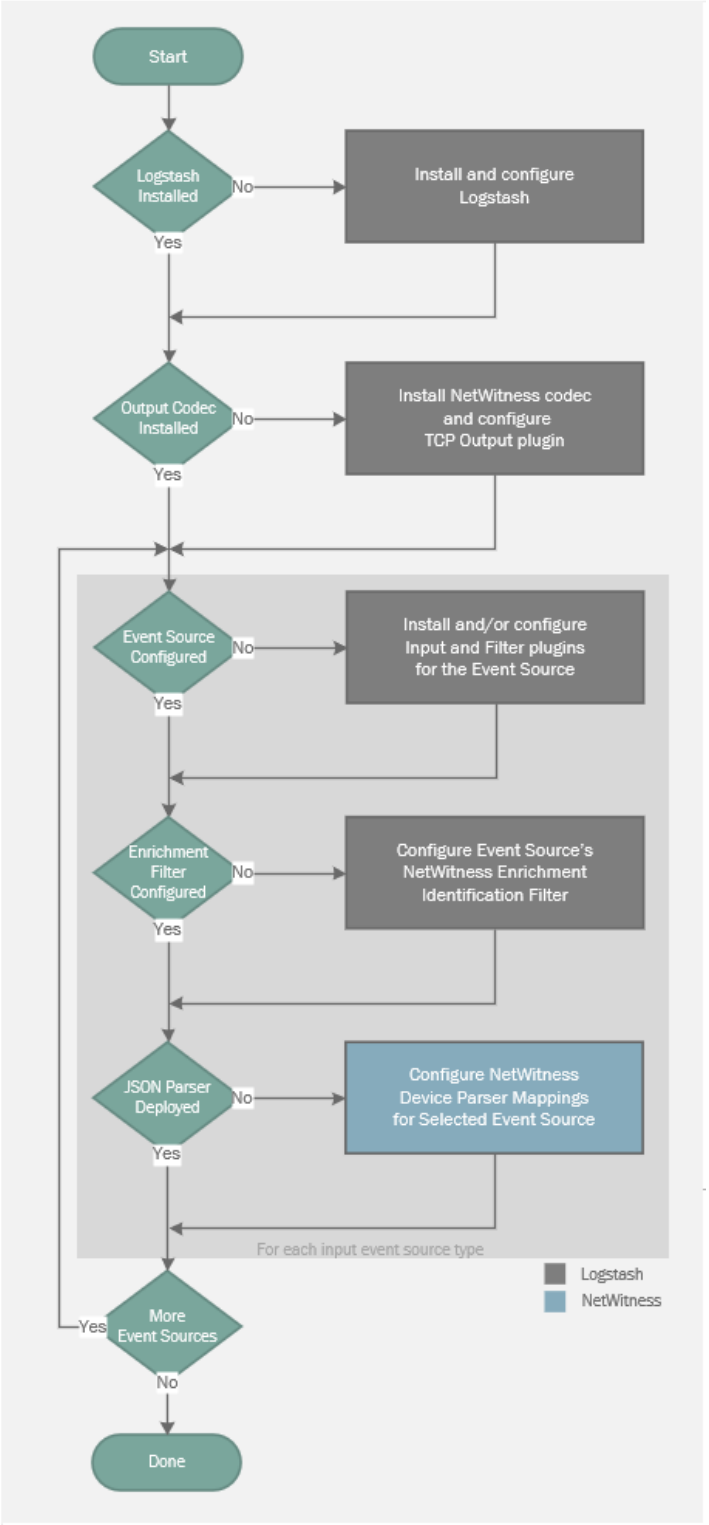
From 11.6 onwards, the Logstash server is packaged and supported along with the NetWitness Log Collector or Virtual Log Collector (VLC) service to provide easy access to Logstash. This is referred to as Managed Logstash and it eliminates the need for a separate Logstash server outside of the NetWitness Platform. For more information, see Configure Logstash Event Sources in NetWitness in the *Log Collection Configuration Guide*.

The following diagram displays a view of how Logstash integrates with the NetWitness.


LogCollector
remote
RabbitMQ
LogCollector
TCP
Syslog w/JSON
Logstash
Syslog w/JSON
LogDecoder
Un-Managed

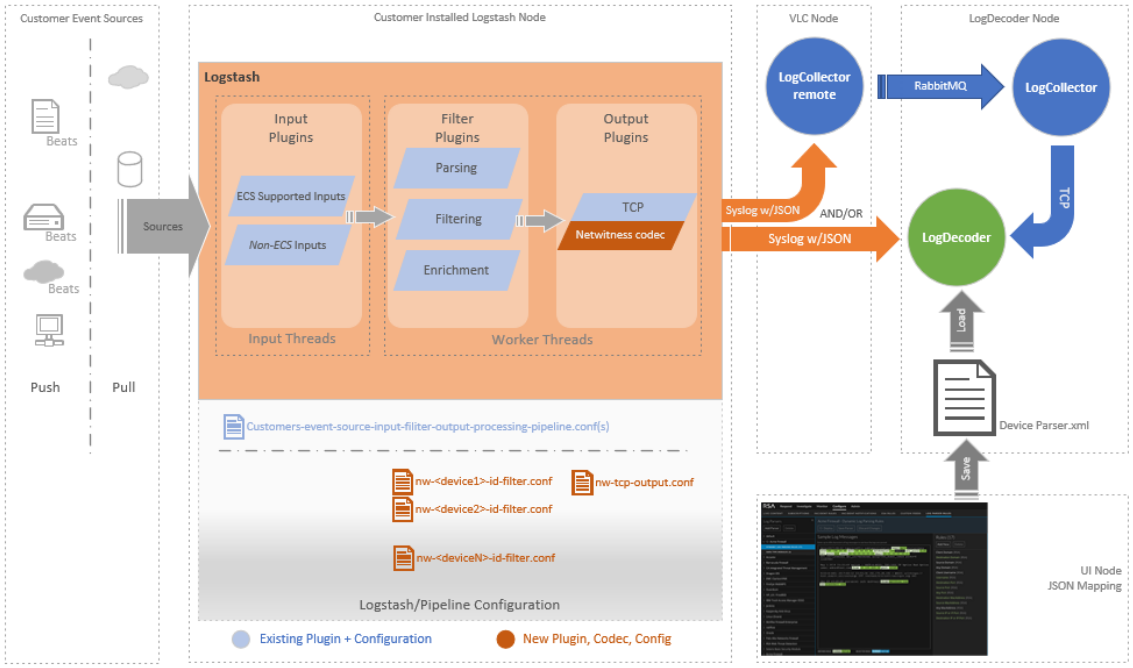
Configuration Process

The following flowchart describes the steps customers take to integrate Logstash with NetWitness, depending on their prior familiarity with and use of Logstash.



The following sequence describes the data flow from an event until it becomes NetWitness meta in a Log Decoder.

1. An event source generates events.
2. The collection plugin (for example a Beats plugin) collects events from the event source.
3. Logstash processes the data from the events.
4. A NetWitness codec encodes the Logstash-processed data into a format that can be consumed by NetWitness.
5. An output plugin sends the processed event data to the NetWitness.
6. A JSON parser populates meta from the processed event data.



Install Logstash

Skip this section if you already have Logstash installed and configured.

Please ensure that you follow all the security-related best practices and guidelines outlined in the Logstash documentation to avoid any potential security risks.

You can install either the free, open source version of Logstash (OSS) or the paid version (Elastic).

Information on released versions of Logstash is available at Logstash Reference. Links in the following steps show an example of installing current version of Logstash on Linux.

1. Install the service: Installing Logstash (OSS free version)
2. Based on your OS, after installation do one of the following:
    - Linux: set Logstash to run as a service: Running Logstash
    - Windows: see Running Logstash on Windows
3. Next, enable Logstash to start when the system boots up:
    - For CentOS, see https://www.unix.com/man-page/centos/1/systemctl/
    - For Debian, see https://wiki.debian.org/systemd/documentation

    Those are 2 examples: adjust instructions according to your particular Operating System.

For Generic Troubleshooting Instructions for Logstash, follow this link: Logstash Troubleshooting

If you are using CentOS, note the following:

- Logstash logs are stored in /var/log/logstash/logstash-plain.log
- If you install logstash using rpm install, make sure it installs as logstash user and folders get created with the same user: **not** the root user.

Install and Configure the NetWitness Codec

To forward Logstash events to the NetWitness in RFC-5424 format, you need to install the NetWitness codec on your system and refer to it in your output plugin configuration.

To install or update the codec:

The following procedure can be performed on either Linux or Windows: instructions that are specific to an OS are noted.

1. Download offline installer from NetWitness Link in the following location: NetWitness Codec Installer
2. Copy the downloaded NetWitness ZIP archive to the system where Logstash runs.
3. Open a command prompt and change directory to Logstash home:
    - On Linux: cd /usr/share/logstash.
    - On Windows: cd: ***logstash_directory***

      For example: cd: c:\Logstash\
4. Stop the logstash service, if it is running.
    - On Linux, run the following command:

      systemctl stop logstash
    - On Windows, open the Services window (you can search for it or find it from the Start menu), then locate the Logstash service in the list and click **Stop the service**.
5. Run the following command and check to see if **logstash-codec-netwitness** is listed.
    - On Linux:

      bin/logstash-plugin list
    - On Windows:

      .\bin\logstash-plugin list
6. If it is listed, delete it by running the following command:
    - On Linux:

      bin/logstash-plugin remove logstash-codec-netwitness
    - On Windows:

      .\bin\logstash-plugin remove logstash-codec-netwitess
7. To install the latest package, run the following command:

- On Linux:

  bin/logstash-plugin install file:///***<path-to-file>***/logstash-codec-netwitness-offline-**<version>**.zip

- On Windows:

  .\bin\logstash-plugin install file:///***<path-to-file>***/logstash-codec-netwitness-offline-**<version>**.zip

  Make sure to use forward slashes (/) as a separator between Windows subfolders for the path to the logstash installation package.

If the codec is installed correctly, you receive a confirmation message. The following is an example:

Installing file: /usr/share/logstash/logstash-codec-netwitness-offline-1.0.0.zip

Install successful

8. Below are default paths for logstash configurations. All output, filter, and input configuration files are under these folders.

   we would go over those configuration files in later sections.

   - For Linux: /etc/logstash/conf.d/
   - For Windows: ***logstash_directory***/config/

   For details on output files, see "Configure Logstash Output Plugins" on page 13. For details on input and filter plugins, see "Logstash Input and Filter plugins" on page 19.

9. Start the logstash service:

   - On Linux, run the following command:

     systemctl start logstash

   - On Windows, open the Services window (you can search for it or find it from the Start menu), then locate the Logstash service in the list and click **Start the service**.

Configure Logstash Output Plugins

Logstash TCP Output

In order to send the events from Logstash to NetWitness, we use the TCP output plugin: https://www.elastic.co/guide/en/logstash/current/plugins-outputs-tcp.html

The TCP output is configured with the NetWitness codec, which formats the outgoing events to be consumable by a NetWitness Log Decoder or Virtual Log Collector (VLC).

The following is an example of a properly configured output block using TCP & the NetWitness codec:

Output Block

```
output {
  tcp {
    id => "nw-output-tcp"
    host => "10.10.1.2" #IP or Hostname of destination Log Decoder or VLC
    port => 514
    codec => netwitness
  }
}
```

Logstash TLS Output

The output block can be further configured to allow for TLS communication between Logstash and NetWitness. An example of a properly configured output block using TLS and the NetWitness codec:

TLS Output Block

```
output {
  tcp {
    id => "nw-output-tcp"
    host => "10.10.1.2" #IP or Hostname of destination Log Decoder or VLC
    port => 6514
    ssl_enable => true
    codec => netwitness
  }
}
```

TLS with Log Decoder (or Virtual Log Collector) Verification

TLS can also be set up to verify the Log Decoder or Virtual Log Collector (VLC) to which it will be communicating. To do this, the Root and Intermediate CA certificates need to be obtained and stored in a truststore for Logstash.

1. On the Log Decoder (or VLC) to which you will be sending events, run the following command:

   cat /etc/pki/nw/ca/nwca-cert.pem /etc/pki/nw/ca/ssca-cert.pem > nw-truststore.pem
2. Copy the nw-truststore.pem file to the Logstash machine and store it in a known location.
3. Create a certificate for the Logstash machine using a self-signed CA or your own CA.
4. Store the cert and private key files in a location of your choosing.

   You need to specify the locations of these files in your TLS output block.

The following code snippet shows an example of a properly configured output block using TLS and the NetWitness codec

Output Block with Verification

```
output {
  tcp {
    id => "nw-output-tcp"
    host => "10.10.1.2" #IP or Hostname of destination Log Decoder or VLC
    port => 6514
    ssl_enable => true
    ssl_verify => true
    ssl_cacert => "/path/to/certs/nw-truststore.pem"
    ssl_key => "/path/to/certs/privkey.pem"
    ssl_cert => "/path/to/certs/cert.pem"
    codec => netwitness
  }
}
```

Configure the Event Source

If you have logstash already install and configured to collect events from desired event-sources you can skip this section.

As a starting point, view the list of input plugins here: Input Plugins. This is not meant to be an exhaustive list, just as a starting point to view some of the available input plugins.

The remainder of this section describes some examples.

Collect Apache File Logs

Typically, Apache logs are written to files on the disk, so to collect events from file you could use the **Filebeat** plugin.

To install and configure Filebeat, see the following websites:

1. Install the Filebeat service on Linux. Refer to the following link: Filebeat Installation
2. Configure Filebeat to collect from specific logs. Refer to the following link: Filebeat Configuration
3. Configure Filebeat to send the output to Logstash. Refer to the following link: Filebeat Logstash Output

Collect CentOS Audit Logs

To collect audit events from an operating system (for example CentOS), you could use the **Auditbeat** plugin.

To install and configure Auditbeat, see the following websites:

1. Install the Auditbeat service on Linux. Refer to the following link: Auditbeat Installation
2. Configure Auditbeat to collect from specific logs. Refer to the following link: Auditbeat Configuration
3. Configure Auditbeat to send the output to Logstash. Refer to the following link: Auditbeat Logstash Output

Configure Parameters for Filebeat or Auditbeat

For Filebeat or Auditbeat plugin, make sure to configure input and output parameters. Modify /etc/filebeats/filebeat.yml as shown below.

To enable file collection, modify the **Filebeat inputs** section as shown in the following image:



```
#========================== Filebeat inputs =============================

filebeat.inputs:

# Each - is an input. Most options can be set at the input level, so
# you can use different inputs for various configurations.
# Below are the input specific configurations.

- type: log

  # Change to true to enable this input configuration.
  enabled: true

  # Paths that should be crawled and fetched. Glob based paths.
  paths:
    - /var/log/*.log
    #- c:\programdata\elasticsearch\logs\*
```

Modify the Outputs as follows:

- Comment out settings in the **Elasticsearch Output** section
- Uncomment **Logstash output** settings and provide logstash IP and Port.

Below is an example of the sections after changes are made:

```
#================================ Outputs =====================================

# Configure what output to use when sending the data collected by the beat.

#-------------------------- Elasticsearch output ------------------------------
#output.elasticsearch:
  # Array of hosts to connect to.
  # hosts: ["localhost:9200"]

  # Optional protocol and basic auth credentials.
  #protocol: "https"
  #username: "elastic"
  #password: "changeme"

#----------------------------- Logstash output --------------------------------
output.logstash:
  # The Logstash hosts
  hosts: ["localhost:5044"]

  # Optional SSL. By default is off.
  # List of root certificates for HTTPS server verifications
  #ssl.certificate_authorities: ["/etc/pki/root/ca.pem"]

  # Certificate for SSL client authentication
  #ssl.certificate: "/etc/pki/client/cert.pem"

  # Client Certificate Key
  #ssl.key: "/etc/pki/client/cert.key"
```

Configure Logstash Filters to Add NetWitness Meta

In order for an event to be processed in the Log Decoder as a specific data type, you need to add some meta key information to the event in Logstash.

- [@metadata][nw_type] — NetWitness device parser content name
- [@metadata][nw_msgid] — NetWitness message pattern to parse message
- [@metadata][nw_source_host] — Originating event source's IP or host value

The value for nw_type *must* match the device parser name. It should be composed of lowercase characters, numbers, or underscore and be less than 29 characters in length.

Optionally, you can add the following meta key:

[@metadata][nw_collection_host] — Collection system identifier (**lc.cid**)

By default, the NetWitness codec sends the complete JSON event as payload to the NetWitness Log Decoder. If the NetWitness nw_type device parser type has a custom payload format and failover payload format, the NetWitness codec plugin must be configured to use them. Please see the "Configure the Event Source" on page 15 section for more details.

The following code snippet contains an example of adding the required meta:

Code to Populate NetWitness Meta

```
filter {
  if ![@metadata][nw_type] {
    if [agent][type] == "filebeat" {
      mutate {
        add_field => {
          "[@metadata][nw_type]" => "linux"
          "[@metadata][nw_msgid]" => "LOGSTASH001"
          "[@metadata][nw_source_host]" => "%{[host][hostname]}"
        }
      }
    }
  }
}
```

Advanced NetWitness Configuration

Grok Filter Plugin

You can use Grok to parse incoming logs and extract valuable meta information. For details, see Grok filter plugin: match. The meta extracted using grok is part of the full event package sent to NetWitness, where it can be mapped to NetWitness meta.

Resources:

- List of Grok patterns: https://github.com/logstash-plugins/logstash-patterns-core/tree/master/patterns
- Grok debugger: https://grokdebug.herokuapp.com/

Logstash Input and Filter plugins

See the following URLs:

- Input Plugins
- Filter Plugins

You can use the beats input plugin for Logstash to receive events from beats sources, including Filebeat & Auditbeat: Beats Input Plugin

You can use the drop filter plugin for Logstash to filter out unwanted logs being passed through Logstash: Drop Filter Plugin

Filter out unwanted logs

You can also filter out unwanted logs using the drop-event processor for Filebeat & Auditbeat.

- Examples for Filebeat:

  Filebeat Reference: Filtering and Enhancing Data

  Filebeat Reference: Drop Events

- Examples for Auditbeat:

  Auditbeat Reference: Filter and Enhance the Exported Data

  Auditbeat Reference: Drop Events

Configure heartbeat plugin to send test logs to NetWitness

The **Heartbeat** plugin can be used to send a test message to verify connectivity between logstash and NetWitness. For details, see the Heartbeat Plugin configuration guide: Heartbeat Input Plugin.

Example of Heartbeat Plugin

```
...

heartbeat {
  id => "sample_plugin"
  #interval => 60
  count => 3
  message => "sequence"
  add_field => {
  "[@metadata][nw_type]" => "logstash_testlog"
  "msg" => "This is test log from some eventsource"
}

...
```

Configure logstash to persist events in case of failure (Recommended)

By default, Logstash uses in-memory queues to buffer events. The size of these in-memory queues is not configurable. If there is a machine failure, or if the service is forcibly stopped, the contents of these queues are lost. To protect against data loss in these situations, Logstash supports persistent queues that are stored on disk and thus can survive failures. For details, see Logstash Persistent Queues.

NetWitness Codec Advanced Configuration

By default, the NetWitness codec sends the complete JSON event as payload to the NetWitness Log Decoder. If the NetWitness nw_type device parser type has a custom payload format, you must configure the NetWitness codec plugin to recognize this custom format.

The payload_format and payload_format_failover mappings use nw_type as the key. The payload_format mapping is searched first for the device type (nw_type). If the device type is not set, or no format is specified for nw_type, or the system fails to make all configured variable substitutions, the complete JSON output is used as the payload. If the primary format from the payload_format mapping fails, the payload_format_failover mapping is tried. If that also fails, the complete JSON output is used as the payload. You can use The format can use Logstash event field syntax for this custom configuration.

The following code snippet shows example of adding meta with custom payload formats.

Add meta with custom payload formats

```
output {
   if [@metadata][nw_type] { # Only targeted NetWitness items
   tcp {
```

```
        id => "netwitness-tcp-output-conf-output"
        host => "127.0.0.1"
        port => 514
        ssl_enable => false
        codec => netwitness {
           # Payload format mapping by nw_type.
           # If nw_type is absent or formatting fails,
           # JSON event is used as the payload
           payload_format => {
              "apache" => "%APACHE-4-%{verb}: %{message}"
           }
           # Failover format, if above format fails
           # If nw_type is absent or formatting fails,
           # JSON event is used as the payload
           payload_format_failover => {
              "apache" => "%APACHE-4: %{message}"  # When verb is missing
              }
           }
        }
     }
  }
```

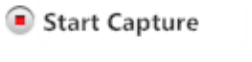
Troubleshoot Installation Issues

If you encounter any issues during installation of Logstash, see Logstash Installation and Setup.

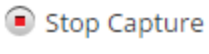
Configure NetWitness to Collect Events

You need to start capture on the Log Decoder to which your are sending your Logstash data.

To start or restart network capture on a Log Decoder:

1. Log in to **NetWitness** and click the **ADMIN** icon > **Services**.
   The Services view is displayed.
2. Select a **Log Decoder** service.
3. Under  (actions), select **View** > **System**.
4. In the toolbar, click 

.

   If the toolbar is displaying the Stop Capture (

) icon, then capture has already been started.

By default, Log Decoders support events that are up to 32 KB in size. If your events are getting truncated on the Log Decoder, use the following procedure to change the event size:

1. Change LogDecoder REST config at **http://*LogDecoder_IP*:50102/-decoder/config**, where ***LogDecoder_IP*** is the IP address of your Log Decoder.
2. Set pool.packet.page.size to 64 KB.
3. Restart the Log Decoder: this is required after you change the pool.packet.page value.

If you are collecting events larger than 64 KB in size, follow instructions above in the "Filter out unwanted logs" on page 19 section. You can drop unwanted logs or fields for a specific event source, to reduce the size of the incoming data.

Linux Event Source Example

This section shows sample input, filters and output configuration to collect system and audit events from CentOS.

Input Plugin

An input plugin enables a specific source of events to be read by Logstash. The following code represents an example input plugin.

input-beats.conf

```
# Below input block collects events using beats plugins (e.g filebeats, auditbeats)
# Skip this block if it's already defined in another pipeline.
input {
    beats {
        port => 5044
    }
}
```

Make sure that port 5044 is open on the Logstash machine. As an example, if Logstash is on a CentOS system, run the following commands to open port 5044:

firewall-cmd --add-port=5044/tcp
firewall-cmd --add-port=5044/tcp --permanent
firewall-cmd --reload

Output Plugin

An output plugin sends event data to a particular destination. Outputs are the final stage in the event pipeline.

output-netwitness-tcp.conf

```
# Below is tcp output plugin with netwitness codec to tranform events in syslog and send
it to LogDecoder
# Only one of these configurations can be within the same pipeline.
output {
  #if [@metadata][nw_type] { # Only targeted Netwitness items
    tcp {
      id => "netwitness-tcp-output-conf-output"
      host => "10.10.100.100"  ## LogDecoder IP
      port => 514
      ssl_enable => false
      #ssl_verify => true
      #ssl_cacert => "/path/to/certs/nw-truststore.pem"
      #ssl_key => "/path/to/certs/privkey.pem"
      #ssl_cert => "/path/to/certs/cert.pem"
```

```
      codec => netwitness {
        # Payload format mapping by nw_type.
        # If nw_type is absent or formatting fails, JSON event is used as the payload
        payload_format => {
          "apache" => "%APACHE-4-%{verb}: %{message}"
        }
        # Failover format, if above format fails
        # If nw_type is absent or formatting fails, JSON event is used as the payload
        payload_format_failover => {
          "apache" => "%APACHE-4: %{message}"  # When verb is missing
        }
      }
    }
  #}
}
```

Filter Plugin

A filter plugin performs intermediary processing on an event. Below is a filter plugin configuration for system events collected from linux using the Filebeat plugin.

linux-system.conf

```
# Filters are often applied conditionally depending on the characteristics of the events.
# Requires these additional configurations within the same pipeline:
#   input-beats.conf
#   output-netwitness-tcp.conf

filter {
if ![@metadata][nw_type] {
  if [ecs][version] and [host][hostname] and [agent][type] == "filebeat" {
    if [event][module] == "system" {
      mutate {
        add_field => {
          "[@metadata][nw_type]" => "linux"
          "[@metadata][nw_msgid]" => "LOGSTASH001"
          "[@metadata][nw_source_host]" => "%{[host][hostname]}"
        }
      }
    }
  }
 }
}
```

Below is filter plugin configuration for audit events collected from linux using the Auditbeat plugin.

linux-audit.conf

```
filter {
    if ![@metadata][nw_type] { # Update Once
```

```
      if [ecs][version] and [host][hostname] and [agent][type] == "auditbeat" {
        if [event][module] == "audit" {
          mutate {
            add_field => {
              "[@metadata][nw_type]" => "linux"
              "[@metadata][nw_msgid]" => "LOGSTASH002"
              "[@metadata][nw_source_host]" => "%{[host][hostname]}"
            }
          }
        }
      }
    }
  }
```

Create a Pipeline

It is recommended to have one pipeline for each input type. For example, all beats collection should be in the same pipeline. To run collection as separate pipeline, create a directory and add above input, filters, and output configuration files to it.

Example Pipeline for Beats

```
/etc/logstash/pipeline1/
/etc/logstash/pipeline1/input-beats.conf
/etc/logstash/pipeline1/output-netwitness-tcp.conf
/etc/logstash/pipeline1/linux-system.conf
/etc/logstash/pipeline1/linux-audit.conf
```

Modify /etc/logstash/pipeline.yml and add the following entries:

Add to pipeline.yml

```
- pipeline.id: my-sample-pipeline-1
  path.config: "/etc/logstash/pipeline1/*.conf"
```

Build Custom JSON Parser

This section is intended for advanced programmers who want to build their own JSON parser. It describes how to build a Logstash parser for a sample device. We use the Linux device as an example throughout.

This chapter is optional: you do not need to build a custom JSON parser from scratch to input logs from Logstash to NetWitness.

Major sections in this document:

- Configure a filter by defining several required pieces of metadata.
- Examine a sample log message from the Linux device
- Walk through creating the parser, based on the sample log message
- View the parsed meta from the sample log message, as it appears on the Log Decoder

Sample JSON Log Received on Log Decoder

Let's examine a sample log and discuss its contents.

<13>1 - Centos7 linux - LOGSTASH001 [lc@36807 lc.ctime="1585886465037" lc.cid="Centos7" lc.ctype="logstash"] {"message": "msg='op=PAM:accounting grantors=pam_access,pam_unix,pam_localuser acct=root exe=/usr/sbin/crond hostname=? addr=? terminal=cron res=success'", "user":{ "email":"john.deaux@test.com", "username":"CORP\\deauxj" }, "host": { "name": "Centos7", "hostname": "Centos7", "containerized": false, "architecture": "x86_64", "id": "d1059ac783b24eb7bbde70a41fa572c9", "os": { "name": "CentOS Linux", "kernel": "3.10.0-1062.el7.x86_64", "version": "7 (Core)", "codename": "Core", "platform": "centos", "family": "redhat" } }, "@timestamp": "2020-04-03T04:01:05.037Z", "files": ["test1.log", "test2.log", "test3.log"],"machine_details" : { "1" : { "hostname" : "USXXLinux" }, "2" : { "hostname" : "USXXWindows" }}}

The first portion of the log is the RFC-5424 header:

<13>1 - **Centos7 linux** - **LOGSTASH001** [lc@36807 lc.ctime="1585886465037" lc.cid="Centos7" lc.ctype="logstash"]

This header contains the information that we used in setting our fields above:

- **nw_source_host**: Centos7 (Hostname)
- **nw_type**: linux (Device Type)
- **nw_msgid**: LOGSTASH001 (Message ID)

The remainder of the log is the JSON Payload.

JSON Payload

```
{
```

```
    "message": "msg='op=PAM:accounting grantors=pam_access,pam_unix,pam_localuser acct-
t=root exe=/usr/sbin/crond hostname=? addr=? terminal=cron res=success'",

    "user": {
        "email": "john.deaux@test.com",
        "username": "CORP\\deauxj"
    },

    "host": {
        "name": "Centos7",
        "hostname": "Centos7",
        "containerized": false,
        "architecture": "x86_64",
        "id": "d1059ac783b24eb7bbde70a41fa572c9",
        "os": {
            "name": "CentOS Linux",
            "kernel": "3.10.0-1062.el7.x86_64",
            "version": "7 (Core)",
            "codename": "Core",
            "platform": "centos",
            "family": "redhat"
        }
    },

    "@timestamp": "2020-04-03T04:01:05.037Z",

    "files": [
        "test1.log",
        "test2.log",
        "test3.log"
    ],

    "machine_details": {
        "1": { "hostname": "USXXLinux"},
        "2": { "hostname": "USXXWindows"}
    }

}
```


Create the JSON Parser for a Linux Device

Now that we have the sample log from the Linux device, we can construct a filter plugin for this device.

Initial Parser to Match Message ID and Device Type

The parser name should match the device type. We call this initial parser v20_linuxmsg.xml, which matches the message ID from the event. We set content to a variable, logstash_json_payload, which represents the JSON payload. We will parse the payload later in the process.

Message ID and Device Type Parsing

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<DEVICEMESSAGES
    name="linux"
    displayname="Linux"
    group="Unix">
    <VERSION device="2.0"/>

    <MESSAGE
        id1="LOGSTASH001"
        id2="LOGSTASH001"
        content="&lt;logstash_json_payload&gt;" />

    <!-- Additional logic to parse JSON payload -->

</DEVICEMESSAGES>
```

Map Payload Contents to Datatypes

We create datatypes to map each element from the payload to meta that can be saved to the NetWitness database.

The entire payload is assigned to the **FileBeatsEvent** datatype.

```
<VARTYPE name="logstash_json_payload" dataType="FileBeatsEvent"/>
```

The timestamp is parsed an assigned to the **InternetTime** datatype.

```
<DataType name="InternetTime"dateTime="%W-%M-%DT%H:%T:%S.%V%E" />
```

Parse the Message String

Using the parser above, we parse (using the FineParse type defined above) the message string from the log file.

"message": "msg='**op**=PAM:accounting grantors=pam_access,pam_unix,pam_localuser **acct**=\"root\"
exe=\"/usr/sbin/crond\" hostname=? addr=? terminal=cron **res**=success'",

The following code extracts the values from the string and saves them to meta keys:

- **op** is saved to **operation.id**
- **acct** is saved to **service.account**
- **res** is saved to **result**
- **exe** is saved to **process.src**

Note that he **search** flag is set to **true** to parse key/value pars regardless of their order in the string.

Extract Values: save to Meta

```
<DataType name="TagValParse" regex="(?: |^)(?:exe=(\S+)|acct=(\S+)|res=(\S+)|op=(\S+))"
```

```
search="true">
    <Capture index="1" meta="process.src" />
    <Capture index="2" meta="service.account" />
    <Capture index="3" meta="result" />
    <Capture index="4" meta="operation.id" />
</DataType>
```

The following code:

- Assigns the whole key as **message** and maps it to the **FineParse** type.
- The **FineParse** type is then mapped to **TagValParse** type.

```
<DataType name="FineParse" regex="msg='(.*)'">
    <Capture index="1" type="TagValParse" />
</DataType>

<DataType name="FileBeatsEvent" type="ElasticCommonSchemaSubset">
    <Capture key="/message" type="FineParse" meta="message"/>
    <Capture key="/user/email" meta="email"    />
    <Capture key="/user/username" type="DomainUser"/>
    <Capture key="/files/" meta="sourcefile" />
    <Capture key="/machine_details//hostname"   meta="host.dst"/>
</DataType>
```

Give the previous string and the code above, the output on the Log Decoder is as follows:

message: msg='op=PAM:accounting grantors=pam_access,pam_unix,pam_localuser acct=root exe=/usr/sbin/crond
  hostname=? addr=? terminal=cron res=success'
service.account: root
process.src: /usr/sbin/crond
result: success
operation.id: PAM:accounting

Parse an Array in JSON

In our sample JSON log file from earlier, one section contained an array object:

```
"files": [
  "test1.log",
  "test2.log",
  "test3.log"
]
```

To fetch all the values of an array, you need to define a capture key enclosed in forward slashes to fetch all values, for example **/files/**.

```
<DataType name="FileBeatsEvent" type="ElasticCommonSchemaSubset">
   <Capture key="/files/" meta="sourcefile" />
```

```
  </DataType>
```

Using the code above on the sample array, the following would be the output on the Log Decoder:

sourcefile: test1.log
sourcefile: test2.log
sourcefile: test3.log

Parse a Nested JSON Object

Let's look at an example nested object from our sample log file from earlier:

```
"host": {
  "name": "Centos7",
  "hostname": "Centos7",
  "containerized": false,
  "architecture": "x86_64",
  "id": "d1059ac783b24eb7bbde70a41fa572c9",
  "os": {
    "name": "CentOS Linux",
    "kernel": "3.10.0-1062.el7.x86_64",
    "version": "7 (Core)",
    "codename": "Core",
    "platform": "centos",
    "family": "redhat"
  }
}
```

To fetch nested values, you need to build a path that contains the keys from each nested level. For example, to fetch the OS name from our example, you use the following code:

```
  <DataType name="ElasticCommonSchemaSubset">
     <Capture key="/host/os/name"> meta="OS" />
  </DataType>
```

Using the code above on the sample nested object, the following would be the output on the Log Decoder:

OS: CentOS Linux

Capture Data That Has Varying Parent Key

When capturing structured data types like JSON, instead of a numbered capture index, you can provide a field name path that uses the key attribute. For example, assume we want to capture the **hostname** from **machine_details** and ignore the indexed key:

```
"machine_details": {
  "1": { "hostname": "USXXLinux"},
  "2": { "hostname": "USXXWindows"}
}
```

To fetch the required values, which have a varying parent key name, we leave the parent key empty in the path:

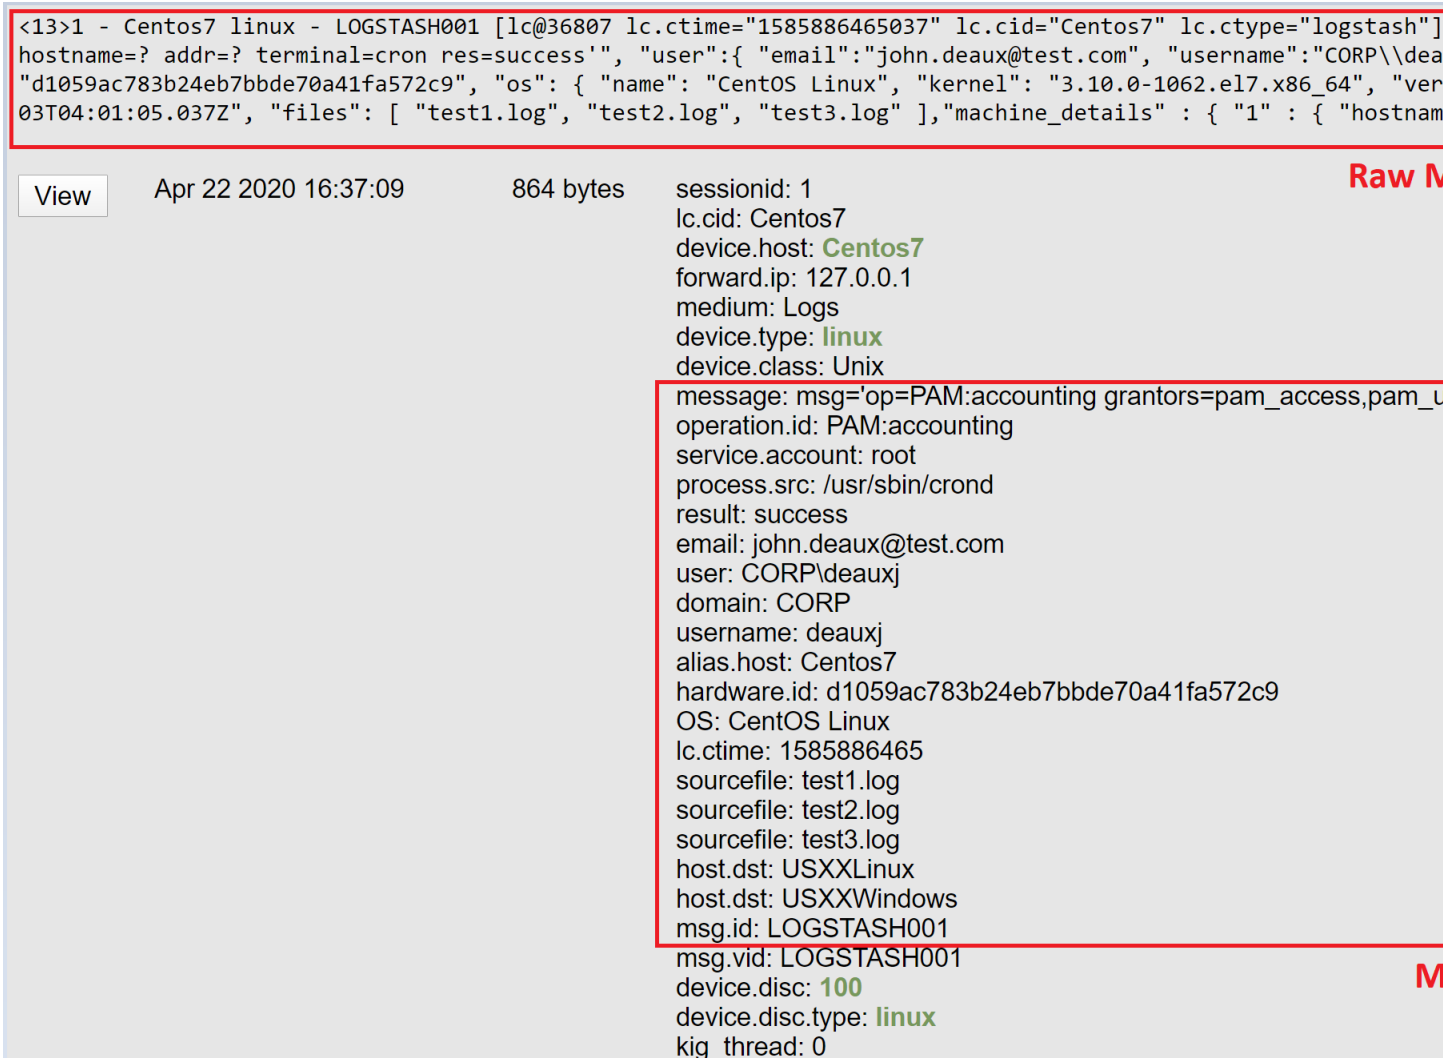
```
<DataType name="FileBeatsEvent" type="ElasticCommonSchemaSubset">
   <Capture key="/machine_details//hostname">  meta="host.dst" />
</DataType>
```

Using the code above on the sample, the following would be the output on the Log Decoder:

host.dst: USXXLinux
host.dst: USXXWindows

The Parsed Example Event on the Log Decoder

Assuming the sample log message from the beginning of this document, and using the parser that we have built, the image below details the event as it would appear on the Log Decoder.



The following representation of the sample log has meta values highlighted.

<13>1 - Centos7 linux - **LOGSTASH001** [lc@36807 **lc.ctime="1585886465037"** lc.-cid="Centos7" lc.ctype="logstash"] {**"message":** "msg='**op=PAM:accounting** grantors=pam_access,pam_unix,pam_localuser **acct=root exe=/usr/sbin/crond** hostname=? addr=? terminal=cron **res=success**'", "user":{ **"email":"-john.deaux@test.com"**, **"username":"CORP\\deauxj"** }, "host": { "name": "Centos7", **"hostname": "Centos7"**, "containerized": false, "architecture": "x86_64", **"id": "d1059ac783b24eb7bbde70a41fa572c9"**, "os": { **"name": "CentOS Linux"**, "kernel": "3.10.0-1062.el7.x86_64", "version": "7 (Core)", "codename": "Core", "platform": "centos", "family": "redhat" } }, "@timestamp": "2020-04-03T04:01:05.037Z", "files": [**"test1.log", "test2.log", "test3.log"**],"machine_details" : { "1" : { **"hostname" : "USXXLinux"** }, "2" : { **"hostname" : "USXXWindows"** }} }

Example Parser Listing

The following code represents the complete parser, including the components we built earlier in this document.

Example Parser Listing

```
<?xml version="1.0" encoding="ISO-8859-1"?>

<DEVICEMESSAGES
   name="linux"
   displayname="Linux"
   group="Unix">
   <VERSION device="2.0" />

<MESSAGE
      id1="LOGSTASH001"
      id2="LOGSTASH001"
      content="&lt;logstash_json_payload&gt;" />

<VARTYPE name="logstash_json_payload" dataType=FileBeatsEvent"/>

<DataType name="InternetTime" dateTime="%W-%M-%DT%H:%T:%S.%V%E" />

<DataType name="CollectionTime" type="InternetTime" meta="lc.ctime"/>

<DataType name="ElasticCommonSchemaSubset" format="JSON">
   <Capture key="/@timestamp">     type="CollectionTime" />
   <Capture key="/host/hostname">  meta="alias.host" />
   <Capture key="/host/id">        meta="hardware.id" />
   <Capture key="/host/os/name">   meta="OS" />
</DataType>

<DataType name="DomainUser" regex="(?:(\w+)\\)?(\w+)">
   <Capture index="0" meta="user" />
   <Capture index="1" meta="domain" />
   <Capture index="2" meta="username" />
</DataType>

<DataType name="TagValParse" regex="(?: |^)(?:exe=(\S+)|acct=(\S+)|res=(\S+)|op=(\S+))"
search="true">
   <Capture index="1" meta="process.src" />
   <Capture index="2" meta="service.account" />
   <Capture index="3" meta="result" />
   <Capture index="4" meta="operation.id" />
</DataType>

<DataType name="FineParse" regex="msg='(.*)'">
   <Capture index="1" type="TagValParse" />
</DataType>

<DataType name="FileBeatsEvent" type="ElasticCommonSchemaSubset">
   <Capture key="/message" type="FineParse" meta="message"/>
   <Capture key="/user/email" meta="email" />
   <Capture key="/user/username" type="DomainUser"/>
   <Capture key="/files/" meta="sourcefile" />
   <Capture key="/machine_details//hostname" meta="host.dst" />
```

```
</DataType>

</DEVICEMESSAGES>
```

Deploy JSON parser

After you have built or changed a JSON parser, you need to upload it to the NetWitness Log Decoder.

1. SSH to the Log Decoder system.
2. Copy the custom parser file to the following folder:

   /etc/netwitness/ng/envision/etc/devices/***eventsource***

   where ***eventsource*** is the name of the event source. You may need to create the folder if it doesn't already exist.

   For example, we need to create **linux** folder under /etc/net-witness/ng/envision/etc/devices directory and copy the **v20_linuxmsg.xml** parser file to /etc/netwitness/ng/envision/etc/devices/linux directory.
3. To get the new parser loaded into memory, you need to reload the parsers on the Log Decoder.

Reload Parsers from REST

From a browser, run the REST reload command by entering the following URL:

http://***<logdecoder_ip>***:50102/decoder/parsers?msg=reload

For example, if your Log Decoder IP address is **10.10.100.101**, use the following string:

http://10.10.100.101:50102/decoder/parsers?msg=reload

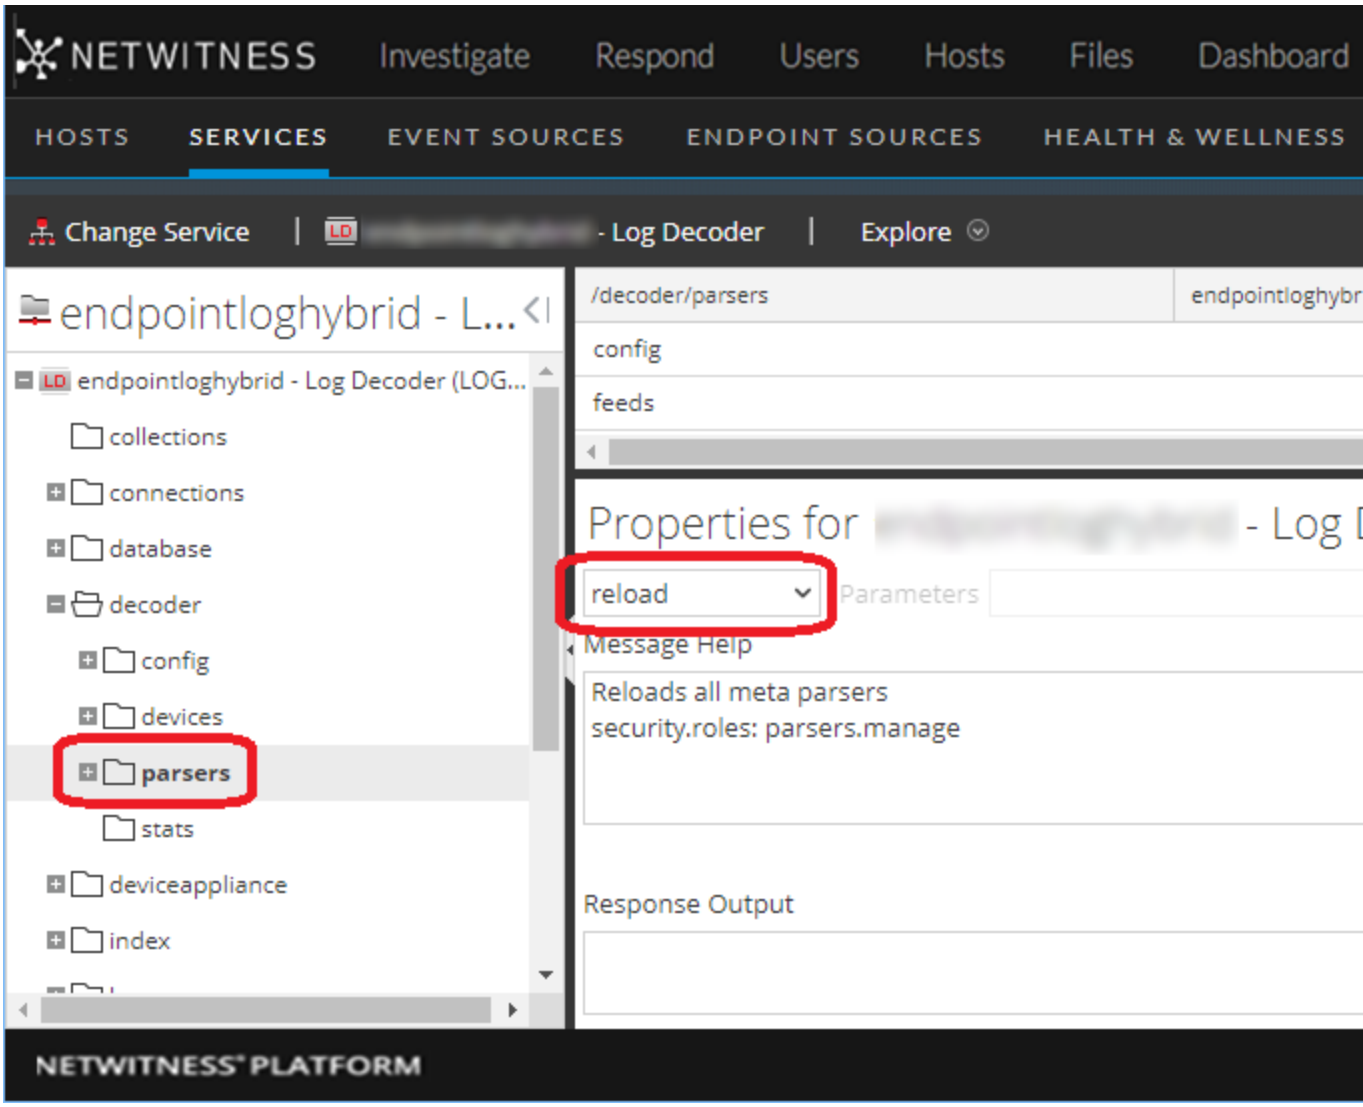
If the call is successful, you should see a REST response, "The parsers have been reloaded."

Reload Parsers from NetWitness UI

You can also reload your parsers from the UI as follows.

1. In the NetWitness UI, navigate to  (Admin) **> Services**.

   The Services view is displayed.
2. Select the Log Decoder to which your want to reload the parsers, and click **View > Explore**.
3. In the left pane, navigate to **decoder > parsers**.
4. Right-click **parsers** and select **Properties**.
5. From the drop-down menu in the Properties panel, select **reload**.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
HOSTS
SERVICES
EVENT SOURCES
ENDPOINT SOURCES
HEALTH & WELLNESS
Change Service
- Log Decoder
Explore
endpointloghybrid - L...
endpointloghybrid - Log Decoder (LOG...
collections
connections
database
decoder
config
devices
parsers
stats
deviceappliance
index
/decoder/parsers
endpointloghybr
config
feeds
Properties for
reload
Parameters
Message Help
Reloads all meta parsers
security.roles: parsers.manage
Response Output
NETWITNESS PLATFORM

6. Click **Send**.

NetWitness® Platform XDR

Version 12.0

Malware Analysis Configuration Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

How Malware Analysis Works **1**

Functional Description 1

Analysis Method 3

NetWitness Server Access to the Malware Analysis Service 3

Scoring Method 3

Deployment 4

Scoring Modules **5**

Network 5

Static Analysis 5

Community 6

Sandbox 6

Roles and Permissions for Analysts **7**

Required Roles and Permissions 7

Basic Setup **9**

Basic Configuration Checklist 9

(Optional) Configure Dedicated Appliance 10

Network Connections 11

Inbound Connections 11

Outbound Connections 11

Configure General Malware Analysis Settings 12

View the Basic Settings 12

Configure Continuous Polling 13

Configure Manual File Upload Settings 15

Configure the Data Repository 15

Calibrate Scoring Modules 16

Configure Static Analysis Scoring 17

Configure Community Analysis Scoring 17

Configure Sandbox Analysis Scoring 18

GFI Sandbox Settings 19

ThreatGRID Sandbox Settings 19

Configure Indicators of Compromise 20

Filter Displayed IOCs by Module 22

Filter Displayed Modules to Show Only Modified Modules 22

Enable and Disable IOCs for a Scoring Module 23

Adjust the Score Weight for an IOC 23

Set the High Confidence Flag for an IOC 24
Reset IOCs to Default Settings 24
Configure Installed Antivirus Vendors 25
Identify Installed AV Software 26
Enable Community Analysis 27
(Optional) Configure Auditing on Malware Analysis Host 27
Configure the Auditing Threshold 28
Configure Incident Management Alerting 28
Configure SNMP Auditing 29
Configure File Auditing Settings 29
Configure Syslog Auditing Settings 30
(Optional) Configure Hash Filter 30
View the Hash List 31
Add a File Hash to the Hash Filter 31
Mark a Hash as Trusted or Untrusted 31
Delete a Hash from the Hash Filter 31
Search for a File Hash 32
Import a Hash List Using the Watched Folder 32
(Optional) Configure Malware Analysis Proxy Settings 34
Configure the Web Proxy 34
(Optional) Register for a ThreatGRID API Key 35
Additional Procedures for Configuring Malware Analysis 37
Create Custom Alert in CEF Format 37
The CEF Template 37
Understand a Syslog Auditing File Entry 37
First Line 38
Audit Common Event Format (CEF) Header 38
Audit CEF Extension 39
Analysis Scores 39
File Information 40
Event Meta Data Retrieved by NextGen 40
Edit the Configuration File 41
Example 42
Enable Custom YARA Content 53
Prerequisites 54
Install Libraries and Applications Required to Build YARA on a CentOs-Based Appliance 54
Set Up YARA 55
Supported Antivirus Vendors 57
Malware Analysis References 59
Services Config View - General Tab 60

Workflow 60
What do you want to do? 60
Quick Look 61
Continuous Scan Configuration Section 62
Repository Configuration Section 65
Miscellaneous Configuration Section (10.3 SP2 and Later) 66
Modules Configuration Section 66
Static Analysis Configuration 66
Community Analysis Configuration 67
Sandbox Analysis Configuration 68
GFI Sandbox Settings 69
ThreatGRID Sandbox Settings 70
Services Config View - Indicators of Compromise Tab 71
Workflow 71
What do you want to do? 71
Related topic 71
Quick Look 72
Services Config View - IOC Summary Tab 74
Workflow 74
What do you want to do? 74
Related Topic 75
Quick Look 75
Features 75
Services Config View - Auditing Tab 77
Workflow 77
What do you want to do? 77
Related Topics 78
Quick Look 78
Audit Thresholds 79
SNMP Auditing 80
Respond Alerting 81
File Auditing 81
Syslog Auditing 82
Services Config View - Hash Tab 84
Workflow 84
What do you want to do? 84
Related Topic 84
Quick Look 85
Services Config View - AV Tab 87
Workflow 87

What do you want to do? 87
Related Topic 87
Quick Look 88
Features 88
Service Config View - Proxy Tab 89
Workflow 89
What do you want to do? 89
Related topic 89
Quick Look 90
Services Config View - ThreatGRID Tab 91
Workflow 91
What do you want to do? 91
Related Topic 92
Quick Look 92
Features 92
Services Config View - Integration Tab 93
Workflow 93
What do you want to do? 93
Related Topic 93
Quick Look 94

How Malware Analysis Works

NetWitness Malware Analysis is an automated malware analysis processor designed to analyze certain types of file objects (for example, Windows portable executable (PE), PDF, and MS Office) to assess the likelihood that a file is malicious.

Malware Analysis detects indicators of compromise using four distinct analysis methodologies:

- Network Session Analysis (network)
- Static File Analysis (static)
- Dynamic File Analysis (sandbox)
- Security Community Analysis (community)

Each of the four distinct analysis methodologies is designed to compensate for inherent weaknesses in the others. For example, Dynamic File Analysis can compensate for Zero-Day attacks that are not detected during the Security Community Analysis phase. By avoiding malware analysis that strictly focuses on one methodology, the analyst is more likely to be shielded from false negative results.

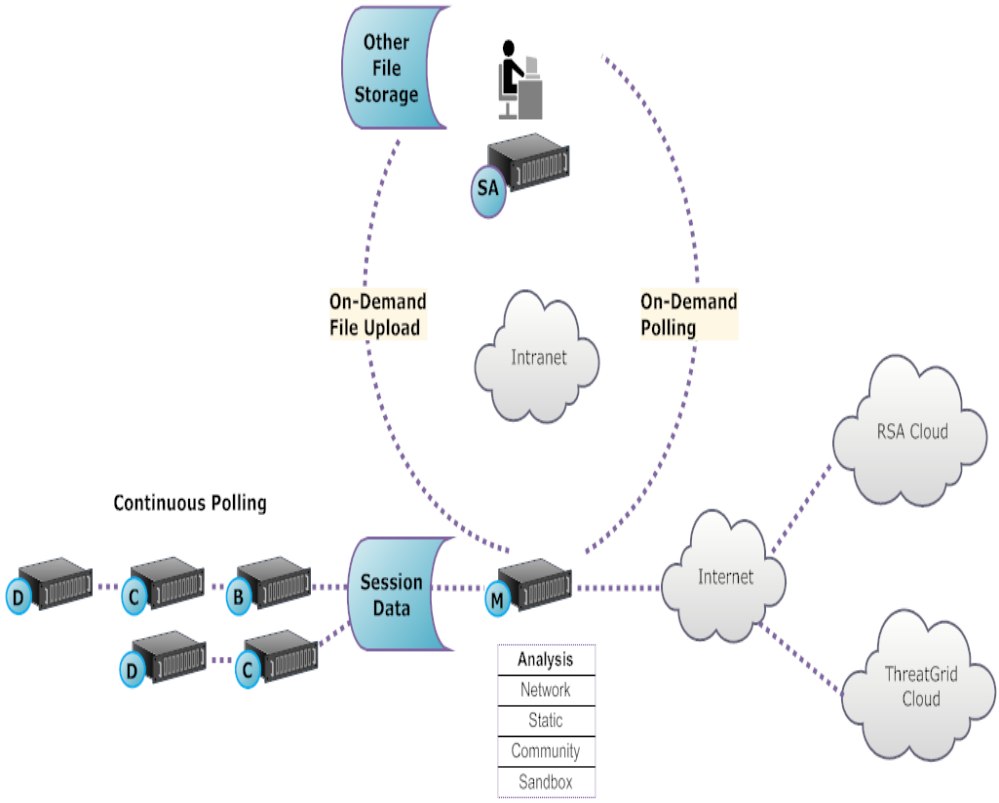
In addition to the built-in indicators of compromise, Malware Analysis supports indicators of compromise written in YARA. YARA is a rule language, which allows malware researchers to identify and classify malware samples. This allows Indicators of Compromise (IOC) authors to add detection capabilities to Netwitness Malware Analysis by authoring YARA rules and publishing them in Netwitness Live. These YARA-based IOCs in Netwitness Live will automatically be downloaded and activated on the subscribed host, to supplement the existing analysis that is performed in each analyzed file.

Malware Analysis also has features that support alerts for NetWitness Respond.

Functional Description

This figure depicts the functional relationship between the Core services (the Decoder, Concentrator, and Broker), the Malware Analysis service, and the NetWitness Server.

Daily Quota (Number of Files)	Free	Standard	Enterprise
Malware Analysis	100	unlimited	unlimited
ThreatGrid Analysis	5	1000	5000



The Malware Analysis service analyzes file objects using any combination of the following methods:

- **Continuous automatic polling of a Concentrator or Broker** to extract sessions identified by a parser as potentially carrying malware content.
- **On-demand polling of a Concentrator or Broker** to extract sessions identified by a malware analyst as potentially carrying malware content.
- **On-demand upload of files** from a user-specified folder.

When automatic polling of a Concentrator or Broker is enabled, the Malware Analysis service continuously extracts and prioritizes executable content, PDF documents, and Microsoft Office documents on your network, directly from data captured and analyzed by your Core service. Because the Malware Analysis service connects to a Concentrator or Broker to extract only those executable files that are flagged as possible malware, the process is both rapid and efficient. This process is continuous and does not require monitoring.

When on-demand polling of a Concentrator or Broker is chosen, the malware analyst uses Investigation to drill into captured data and choose sessions to be analyzed. The Malware Analysis service uses this information to automatically poll the Concentrator or Broker and to download the specified sessions for analysis.

On-demand upload of files provides a method for the analyst to review files captured external to the Core infrastructure. The malware chooses a folder location and identify one or more files to be uploaded and analyzed by Malware Analysis. These files are analyzed using the same methodology as files automatically extracted from network sessions.

Analysis Method

For the Network analysis, the Malware Analysis service looks for characteristics that seem to deviate from the norm, much as an analyst does. By looking at hundreds to thousands of characteristics and combining the results into a weighted scoring system, legitimate sessions that coincidentally have a few abnormal traits are dismissed, while the actual bad ones are highlighted. A user can learn patterns that indicate anomalous activity in the sessions as indicators that warrant further investigation, Indicators of Compromise.

The Malware Analysis service can perform Static analysis against suspicious objects it finds on the network and determine whether those objects contain malicious code. For Community analysis, new malware detected on the network is pushed to the Netwitness Cloud for checking against Netwitness's own malware analysis data and feeds from the SANS Internet Storm Center, SRI International, the Department of the Treasury and VeriSign. For Sandbox analysis, the services can also push data into major security, information and event management (SIEM) hosts (the ThreatGRID Cloud).

Malware Analysis has a unique method for analysis that is partnered with industry leaders and experts, so their technologies can enrich the Malware Analysis scoring system.

NetWitness Server Access to the Malware Analysis Service

The NetWitness Server is configured to connect to the Malware Analysis service and import tagged data for deeper analysis in Investigation. Access is based on three subscription levels.

- Free subscription: All NetWitness customers have a free subscription, with a free trial key for ThreatGRID analysis. The Malware Analysis service is rate-limited to 100 file samples per day. The number of samples (within the set of files from above) submitted to the ThreatGRID Cloud for sandbox analysis is limited to 5 per day. If one network session had 100 files in it, customers would hit the rate limit after processing the one network session. If 100 files were manually uploaded, that would cause the rate limit to be reached.
- Standard subscription tier: The number of submissions to the Malware Analysis service is unlimited. The number of samples submitted to the ThreatGRID Cloud for sandbox analysis is 1000 per day.
- Enterprise subscription tier: The number of submissions to the Malware Analysis service is unlimited. The number of samples submitted to the ThreatGRID Cloud for sandbox analysis is 5000 per day.

Scoring Method

By default, the Indicators of Compromise (IOC) are tuned to reflect industry best practices. During analysis, the IOCs that trigger cause the score to move upward or downward to indicate the likelihood that the sample is malicious. The tuning of IOCs is exposed in NetWitness so that the malware analyst can choose to override the assigned score or to disable an IOC from being evaluated. The analyst has the flexibility to either use the default tuning, or to completely customize the tuning to specific needs.

YARA-based IOCs are interleaved with the built-in IOCs within each built-in category and are not distinguished from native IOCs. When viewing IOCs in the Service Configuration view, administrators can select YARA from the Module selection list to see a list of YARA rules.

After a session is imported into NetWitness, all of the viewing and analysis capabilities in Investigation are available to further analyze Indicators of Compromise. When viewed in Investigation, YARA IOCs are distinguished from the built-in native IOCs by the tag `Yara rule`.

Deployment

The Malware Analysis service is deployed as a separate Netwitness Malware Analysis host. The dedicated Malware Analysis host has an onboard Broker which connects to the Core infrastructure (either another Broker or a Concentrator). Prior to this connection, a collection of parsers and feeds must be added to the Decoders that are connected to the Concentrators and Brokers from which the Malware Analysis service pulls data. This allows suspicious data files to be marked for extraction. These files are `malware analysis` tagged content available through the Netwitness Live content management system.

Scoring Modules

Netwitness NetWitness Malware Analysis analyzes and scores sessions and the embedded files within these sessions by scoring four categories: Network, Static Analysis, Community, and Sandbox. Each category comprises many individual rules and checks that are used to calculate a score between -100 and 100. The higher the score, the more likely the session is to be malicious and worthy of more in-depth follow-on investigation.

Malware Analysis can facilitate a historical investigation into events leading up to a network alarm or incident. If you know that a certain type of activity is taking place on your network, you can select only the reports of interest to examine the content of data collections. You can also modify behavior for each scoring category based on the scoring category or the file type (Windows PE, PDF, and Microsoft Office).

Once you become familiar with data navigation methods, you can explore the data more completely through:

- Searching for specific types of information
- Reviewing specific content in detail.

Category scores for Network, Static Analysis, Community, and Sandbox are maintained and reported independently. When events are viewed based on the independent scores, as long as one category detects malware, it is evident in the Analysis section.

Network

The first category examines each core network session to determine if the delivery of the malware candidates was suspicious. For example, benign software being downloaded from a well-known safe site, using proper ports and protocols, is considered less suspicious than downloading software known to be malicious from a known dubious download site. Sample factors used in the scoring of this criteria set may include sessions that:

- Contain threat feed information
- Connect to well-known bad sites
- Connect to high-risk domains/countries (for example, .cc domain)
- Use well-known protocols on non-standard ports
- Contain obfuscated JavaScript

Static Analysis

The second category analyzes each file in the session for signs of obfuscation in order to predict the likelihood of the file behaving maliciously if allowed to run. For example, software that links to networking libraries is more likely to perform suspicious network activity. Sample factors used in the scoring of this criteria set may include:

- Files found to be XOR encoded
- Files found embedded within non-EXE formats (for example, PE file found embedded in a GIF format)
- Files linking to higher risk import libraries
- Files highly deviating from the PE Format

Community

The third category scores the session and files based on the collective knowledge of the security community. For example, files whose fingerprint/hash is already known to be good or bad by respected anti-virus (AV) vendors is scored accordingly. Files are also scored based on knowledge that a file was delivered from a site known to be good or bad by the security community.

Community scoring also indicates whether the AV on your network flagged the files as malicious. It does not indicate that the resident AV product acted to protect your system.

Sandbox

The fourth category examines the behavior of the software by actually running it in a sandbox environment. By running the software to watch its behavior, a score can be calculated by identifying well-known malicious activity. For example, software that configures itself to autostart on each reboot and make IRC connections would score higher than a file with no known bad behavior.

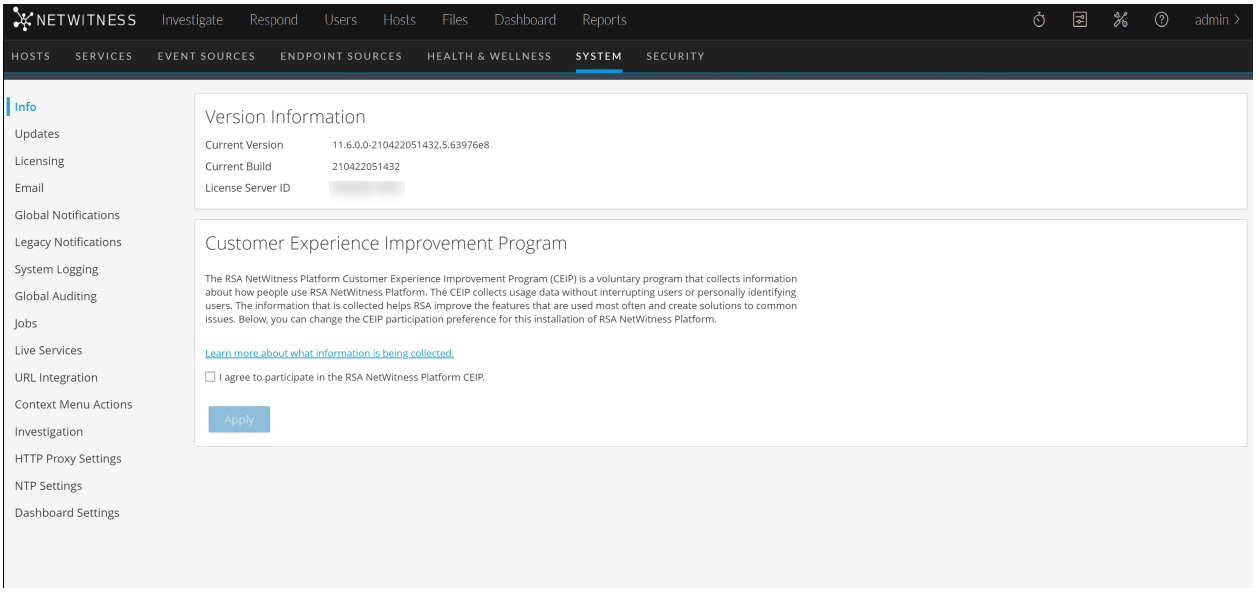
Roles and Permissions for Analysts

This topic identifies the user roles and permissions required for a user to conduct malware analysis in NetWitness. If you cannot perform an analysis task or see a view, the administrator may need to adjust the roles and permissions configured for you.

Required Roles and Permissions

NetWitness manages security by providing access to views and functions using both system permissions and permissions on individual services.

On the system level, the user needs to be assigned a system role, in the (Admin) > System view, that provides access to specific views and functions.



The default `Malware_Analysts` role in NetWitness12.0 is assigned all of the permissions listed below. If necessary, an Administrator can create a custom role with some combination of the following permissions:

- Access Investigation Module (required)
- Investigation - Navigate Events
- Investigation - Navigate Values
- Access Incident Module
- View and Manage Incidents
- View Malware Events (to view events)
- File Download (to download files from the Malware Analysis service)
- Initiate Malware Scan (to initiate a one-time service scan or one-time file upload)

- Dashlet permissions for convenience: Dashlet - Investigate Top Values Dashlet, Dashlet - Investigate Service List Dashlet, Dashlet - Investigate Jobs Dashlet, Dashlet - Investage Shortcuts Dashlet.

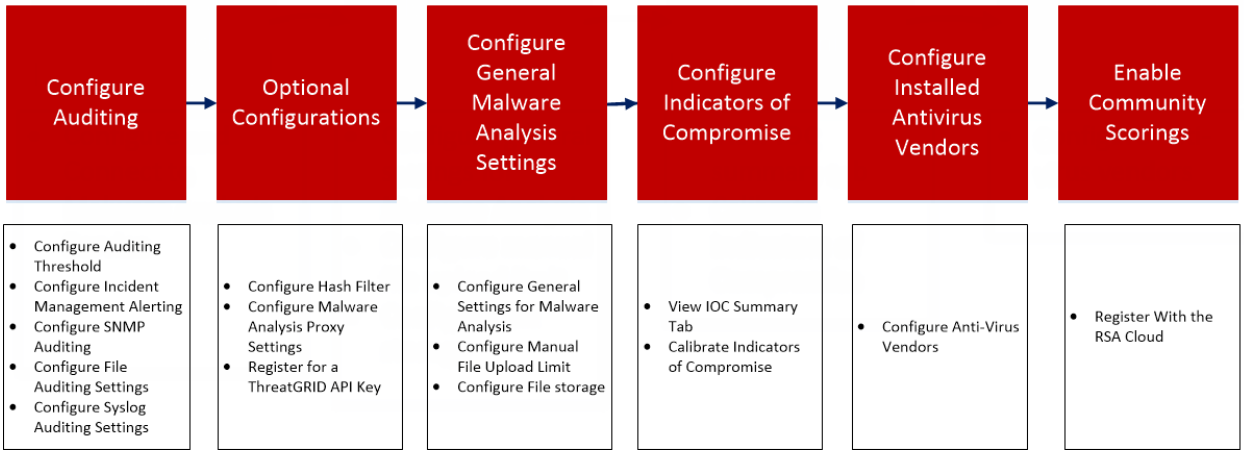
A use case for creating a custom role would be a Junior Malware Analyst role, with limited permissions that do not include the File Download permission.

On specific services, a malware analyst needs to be a member of the **Analysts** group, or to a group that has the two default permissions assigned to the Analyst group: **sdk.meta** and **sdk.content**. Users who have these permissions can use specific applications, run queries, and view content for purpose of analysis on the service.

Basic Setup

Malware Analysis can operate as a service on a Decoder or as a service on a dedicated appliance. This guide includes instructions for setting up the operating environment and then configuring the Malware Analysis service. After this configuration is complete, analysts can conduct malware analysis.

These are the configuration steps for Malware Analysis, and also for editing the configuration. Perform the steps in the section in the sequence they are given.



Basic Configuration Checklist

The following checklist provides the sequence for tasks that are required to configure Malware Analysis that has been added to NetWitness in accordance with the *Hosts and Services Guide*.

Step	High-Level Task
Step 1 - (Optional) Configure Dedicated Appliance	(Optional) Configure Dedicated Appliance This topic describes the procedures for configuring the environment to connect to the Malware Analysis service.
Step 2 - Configure General Malware Analysis Settings	Configure General Malware Analysis Settings • Enable continuous polling. • Configure manual file upload limit. • Configure the file storage repository and database. • Calibrate the Static, Network, Community, and Sandbox scoring modules.

Step	High-Level Task
Step 3 - Configure Indicators of Compromise	Configure Indicators of Compromise Calibrate Indicators of Compromise that are applied for each scoring module (Static, Network, Community, Sandbox) and for YARA-based IOCs.
Step 4 - Configure Installed Antivirus Vendors	Configure Installed Antivirus Vendors
Step 5 - Enable Community Scoring	Enable Community Analysis Register with the Netwitness cloud and test connections to enable Community scoring.
Step 6 - Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host Configure auditing thresholds and enable Syslog, SNMP, and file auditing.
Step 7 - Configure Hash Filter	(Optional) Configure Hash Filter Configure hash filtering to fine tune Malware Analysis event analysis based on known good or bad file hashes.
Step 8 - Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings (Optional) Configure Malware Analysis to communicate with the Netwitness Cloud through a web proxy instead of directly.
Step 9 - Register for a ThreatGRID API key	(Optional) Register for a ThreatGRID API Key

(Optional) Configure Dedicated Appliance

You can configure the NetWitness operating environment to connect to a NetWitness Malware Analysis service.

Malware Analysis operates as a service on a dedicated Malware Analysis appliance. If your site is using a dedicated appliance, do one of the following:

- If your site is adding a new dedicated NetWitness Malware Analysis appliance, install the physical appliance in your network and configure the operating environment.
- If your site is upgrading a dedicated Spectrum appliance to a dedicated NetWitness Malware Analysis appliance, re-image the Spectrum appliance as a Malware Analysis appliance.

Malware Analysis is dependent on the Core infrastructure to operate. The following steps are necessary before Malware Analysis can successfully analyze data.

1. Configure the onboard Broker on the Malware Analysis appliance to connect another Broker or Concentrator in the existing Core infrastructure.

   Note: If no Core infrastructure exists, only manually uploaded files can be analyzed.

2. Use NetWitness Live to find all Live resources with the `malware analysis` tag and deploy these resources to each Decoder service that will be capturing traffic for Malware Analysis to analyze. NetWitness uses this proprietary set of parsers and feeds to find events that are likely to be malware.
3. Configure communications ports. Malware Analysis requires a number of different communications ports to be open, including TCP/443 for HTTPS. These are described below in Network Connections.
4. Configure the NextGen source to which Malware Analysis will connect. This is the Broker or the Concentrator.
   Malware Analysis is now ready to begin analyzing network traffic.

Network Connections

The inbound and outbound network connections must be configured for the Malware Analysis appliance to properly communicate with services, Netwitness sources for software updates, and other critical information.

Your network firewall must be configured to allow the Malware Analysis access to the internet. Proxy servers may be used to facilitate these connections, if necessary.

Inbound Connections

TCP/22 - Secure Shell access to the Malware Analysis server to review log files and troubleshoot. Access can be limited to IP addresses that will be managing Malware Analysis.

- TCP/443 - HTTPS web-based connection to access the Malware Analysis user interface.
- TCP/50008 - JMX port for performance troubleshooting, using an application such as JVisualVM. This is optional and access can be limited to IP addresses that will be managing Malware Analysis.

Outbound Connections

- TCP/443 - HTTPS connections to SSL-based web servers. Some features include Malware Analysis sending files or documents to servers for analysis, which require a secure connection. Use of a web proxy server is supported.
- (TCP/443 - SSL connection from Malware Analysis to the Netwitness Cloud. Use of a SOCKS proxy server is supported. Customer infrastructure changes may be required to ensure that 443 is open to cloud.netwitness.com.)
- TCP/50103 - REST API port used to communicate with a Broker. (NetWitness 10.3.x and earlier)
- TCP/50105 - REST API port used to communicate with a Concentrator. (NetWitness 10.3.x and earlier)
- TCP/50003 TCP/56003 - Ports used to communicate with a Broker. (NetWitness 10.4 and later)
- TCP/50005 TCP/56005 - Ports used to communicate with a Concentrator. (NetWitness 10.4 and later)

- ICMP - JMS connection from NetWitness to the Malware Analysis service to verify if the hostname and ip address entered is valid for a successful test connection.

Configure General Malware Analysis Settings

You can configure several basic settings required to enable and calibrate the consumption of sessions, manual file upload, and the different scoring modules that Malware Analysis uses to analyze data.

You can also set up file sharing with the data repository. Malware Analysis has three modes of consuming sessions and files. Any combination of the three choices may be used to initiate analysis in Malware Analysis. The choices are:

- **Continuous Polling of the Core service:** You can enable and configure continuous polling of the Core service. When enabled and configured, Malware Analysis continuously polls the Core service for sessions tagged for analysis. By default, continuous polling is disabled. You can enable Denial of Service (DOS) attack prevention for use during continuous polling. You can test the connection to the Malware Analysis service that is being continuously polled using an option in the Integration tab.

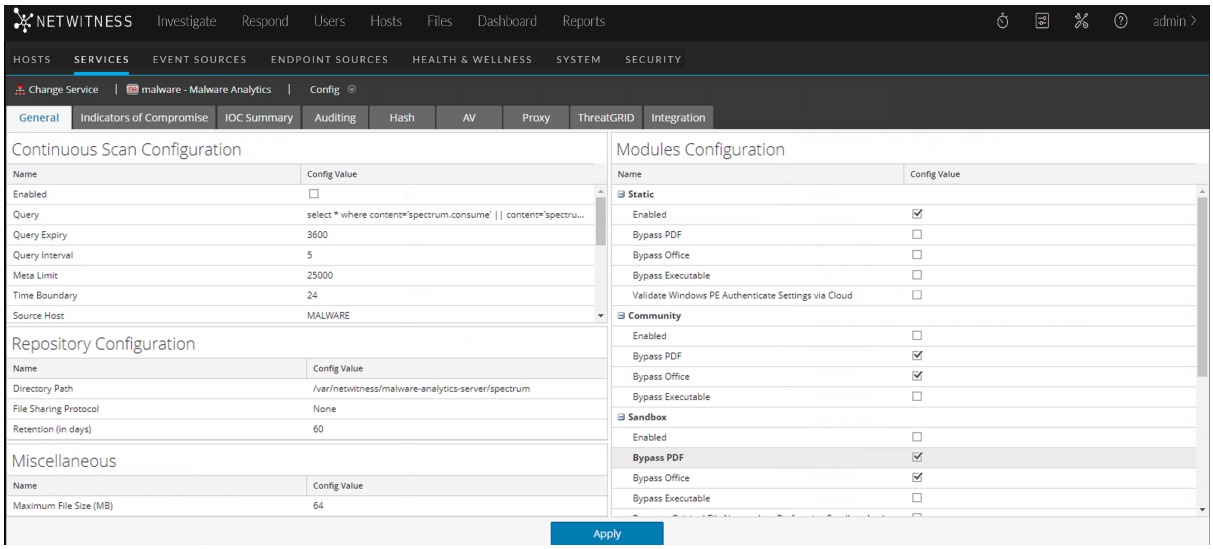
  Note: When adding a Core service as a service for continuous polling on 10.3.5 and earlier Malware Analysis, use the REST port; for example, add a Concentrator to 10.3.5 Malware Analysis with REST port (50105) instead of the native NexGen port (50005).

- **On-Demand Analysis of the Core service:** You can analyze sessions based on Investigations initiated directly in NetWitness. This method allows manually controlled consumption of Core sessions and allows tighter control over how files in those sessions are processed (for example, send to sandbox for processing). Document types can bypass the default restrictions and be sent to community or sandbox processing regardless of the configured setting.
- **Manual File Upload:** You can manually upload one or more files for analysis by navigating to a visible folder on your computer and selecting files to be uploaded. The maximum size for the uploaded files is configurable.

View the Basic Settings

To view the basic settings:

1. Go to  **(Admin) > Services**.
2. In the **Services** grid, select a Malware Analysis service and click  **> View > Config**.
   The Service Config for the service is displayed with the **General** tab open.



Configure Continuous Polling

Malware Analysis is rate limited so that 1,000 files per day may be submitted to ThreatGRID's Cloud for sandbox processing. To optimize your use of the sandbox, Malware Analysis configuration allows you to choose which of several methods of consumption Malware Analysis uses; you can enable or disable continuous polling.

An important consideration when configuring continuous polling is the Denial of Service (DOS) Prevention parameters. By default this feature is disabled because you need to carefully consider the settings for your environment before enabling the feature.

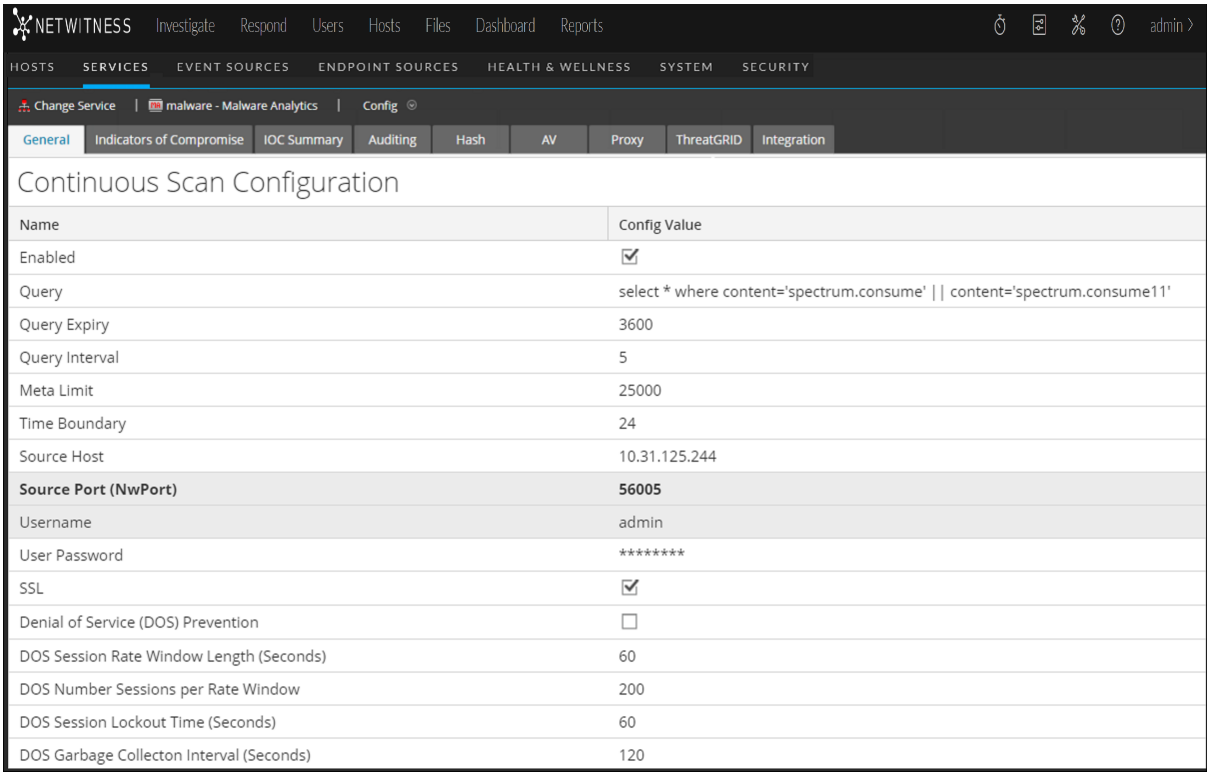
When DOS Prevention is disabled, Malware Analysis analyzes the queued sessions in first-in first-out order. A DOS attack may rapidly fill the queue so that Malware Analysis is busy handling those sessions, while a malware attack is occurring in a later session. The later session with the actual attack may not reach the beginning of the queue and undergo analysis until after the attack has begun.

When DOS Prevention is enabled, Malware Analysis treats too many sessions from a single IP address as a DOS attack. If an IP address exceeds the Number of Sessions per Rate Window, Malware Analysis begins to disregard sessions from that address until the Session Lockout time is reached. Then Malware Analysis resumes analysis of the sessions from that IP address. The disregarded sessions from the IP address are not analyzed at all, so a malware attack may slip through during the Session Lockout period.

Using the DOS Garbage Collection Interval setting, Malware Analysis clears in-memory storage of an IP source after a specified number of seconds. IP addresses with little activity during this interval are cleared from memory. If an IP address is active at intervals that exceed the DOS Garbage Collection Interval, Malware Analysis may not identify it as a DOS attack.

To configure Malware Analysis for continuous polling, in the Continuous Scan Configuration section:

1. In the **General** tab, under **Continuous Scan Configuration** you can configure continuous polling.



2. To enable continuous polling, click **Enabled**.
3. (Optional) If you want to change the default values for querying, enter new values for the **Query Expiry**, **Query Interval**, **Meta Limit**, and **Time Boundary**.
4. To configure the Malware Analysis appliance that Malware Analysis queries to retrieve data for analysis, specify the **Source Host** and **Source Port (NwPort)**.
5. (Optional) If you want to change the default logon credentials for the Malware Analysis appliance, specify the **Username** and **User Password**.
6. If you want to use SSL for communication between the Malware Analysis appliance and the Core service, enable **SSL**.
7. (Optional) If you want to configure Denial of Service (DOS) prevention:
   a. Enable the **Denial of Service (DOS) Prevention** parameter.
   b. Set up the DOS prevention session limitations:
      - Specify the number of seconds of the time window during which Malware Analysis counts sessions for a single IP address (**DOS Session Rate Window Length**). The window is called a Rate Window and a counter is set when the first session is received from that IP source. The default value is 60 seconds.
      - Specify the number of sessions allowed per Rate Window in the **DOS Number Session per Rate Window**. The default value is 200 sessions. When the number of sessions is reached within the Rate Window; Malware Analysis begins disregarding sessions from the IP address

and the disregarded sessions from that IP are not analyzed at all. Malware Analysis continues to disregard sessions until the lockout time is reached.

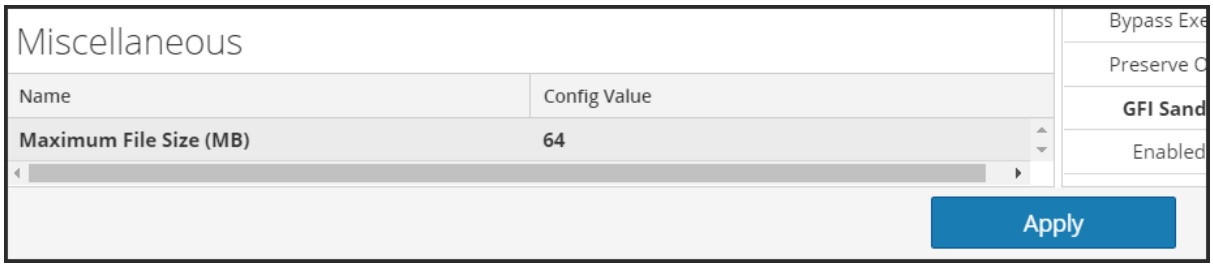
- Specify the length of lockout time (during which sessions from the IP address are disregarded and not analyzed) in the **DOS Session Lockout Time (Seconds)**. The default value is 60 seconds. When the lockout duration has elapsed, Malware Analysis resumes analysis of sessions from that IP address.
- Specify the interval of inactivity for an IP address before NetWitness removes the in-memory object for the IP source in **DOS Garbage Collection Interval (Seconds)**. The default value is 120 seconds.

8. Click **Apply** to apply the changes.
   The applied changes become immediately effective as Malware Analysis receives new packets.
9. Test the connection of the Malware Analysis service to the Core service selected in the **Integration** tab by clicking the **Test Connection** button in the **Continuous Scan Connection Test** section.

Configure Manual File Upload Settings

To configure the maximum file size for manual file upload:

1. In the Miscellaneous section, type the maximum file size in Megabytes allowed for files uploaded manually for Malware Analysis scanning.



2. Click **Apply**.
   The changes become immediately effective.

Configure the Data Repository

Malware Analysis can store a finite number of files on the appliance. The data repository configuration has a file system retention period of 60 days. This setting determines how long files are retained in the Malware Analysis appliance. When old files are deleted, they cannot be recovered. Every day, Malware Analysis deletes files that exceed the file system retention period to ensure that there is no wasted disk space.

The File System Retention Period is the only setting that governs when files are deleted. Files are not deleted based on the amount of disk space being used. If the setting needs to be changed, the administrator must configure the retention period based on the anticipated space usage during the number of retention days specified.

The visible data repository parameters in the NetWitness user interface are:

- The location of the repository is **/var/lib/netwitness/malware-analytics-server/spectrum**. Do not edit this value.
- The file sharing protocol, which allows access through one of the File Sharing Protocols to copy files from the Malware Analysis service.
- The file retention period in number of days.

To configure file sharing, in the Repository Configuration section:

1. Click on the File Sharing Protocol and select **FTP** or **SAMBA** or **None**.
2. Select the number of days that files are maintained in the repository before deletion.
3. Click **Apply.**

The changes become immediately effective.

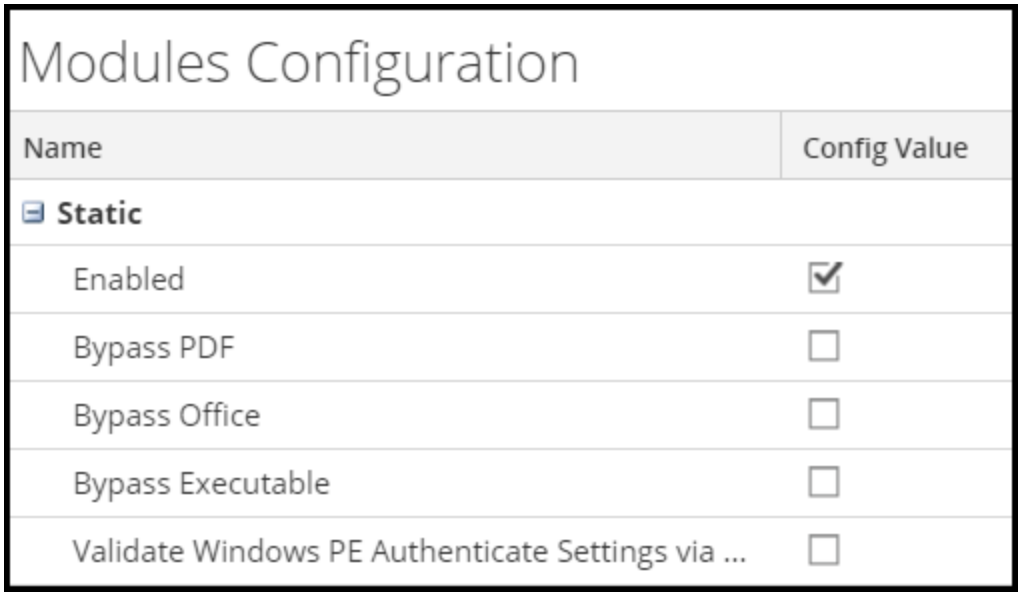
Calibrate Scoring Modules

The Modules Configuration section helps configure the following components of Malware Analysis to:

- Completely disable any or all of three scoring modules (Static, Community, and Sandbox). Before disabling or enabling any scoring module, ensure that you understand what each scoring module detects.
- Malware Analysis tags sessions containing Microsoft Office, Windows PE, and PDF files for consumption by the Malware Analysis service. You can configure Malware Analysis to ignore Windows PE, Microsoft Office, and PDF documents entirely. If this is the case, a better option is to adjust your Core settings to ignore these files so they are not tagged for Malware Analysis consumption.

A sample application for using scoring module calibration is this: when setting up rule groups or analyzing system performance, you can test various scenarios in which PDF documents are not analyzed, but Microsoft Office and Windows PE documents are. You can test the scenario in each of the three scoring modules. If you see a measurable improvement in system performance, you can apply this knowledge on a broader scale.

Configure Static Analysis Scoring



To configure Static analysis scoring, in the **Modules Configuration** section:

1. By default the Static module is enabled. To enable or disable Static analysis entirely, click the **Enabled** checkbox.
2. To configure handling of PDF, Microsoft Office, and Windows PE files in a session, select any of the checkboxes **Bypass PDF**, **Bypass Office**, and **Bypass Executable**.
3. To configure your preference for Authenticode validation of digitally signed Windows PE files, click the **Validate Windows PE Authenticate Settings via Cloud** checkbox. If you want to prevent Windows PE files that are digitally signed from being transmitted to the Netwitness Cloud for validation, remove the check.
   When disabled, ALL static analysis is performed locally (skipping Authenticode validation). Regardless of this setting, PDF and MS Office documents are not subject to Authenticode validation and are not transmitted over the network during static analysis.
4. Click **Apply.**The changes become immediately effective as Malware Analysis receives new packets.

Configure Community Analysis Scoring

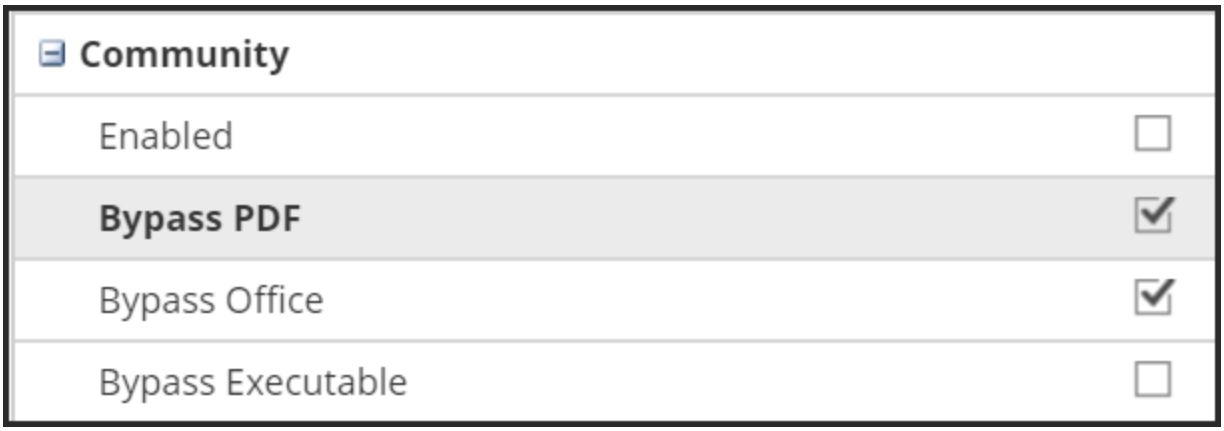
Once the Community module is enabled, the security community analyzes all documents not prevented from processing. This is achieved by sending network session and file attributes to the Netwitness Cloud for processing. The Netwitness Cloud then may make external connection to security community partners as needed to process the information.

The file content is never sent to the community for analysis. Instead, the MD5/SHA-1 hash of the file is sent for Anti-Virus detection and Blacklisting. Similarly, session Meta is harvested and analyzed as part of this process. Meta elements such as URL and Domain Name are examined and transmitted to the Netwitness Cloud to identify known bad URLs/Domains.

You can enable Community analysis and limit which document types are processed. There is no risk for the file content (except for a hash) being sent outside of your network.

Note: To gain access to the Netwitness Cloud where processing occurs, you must register your Malware Analysis service with Netwitness customer service. There are two methods: register the service using the options in the Integration tab or contact Netwitness Customer Care.

To configure Community analysis scoring, in the Modules Configuration section:



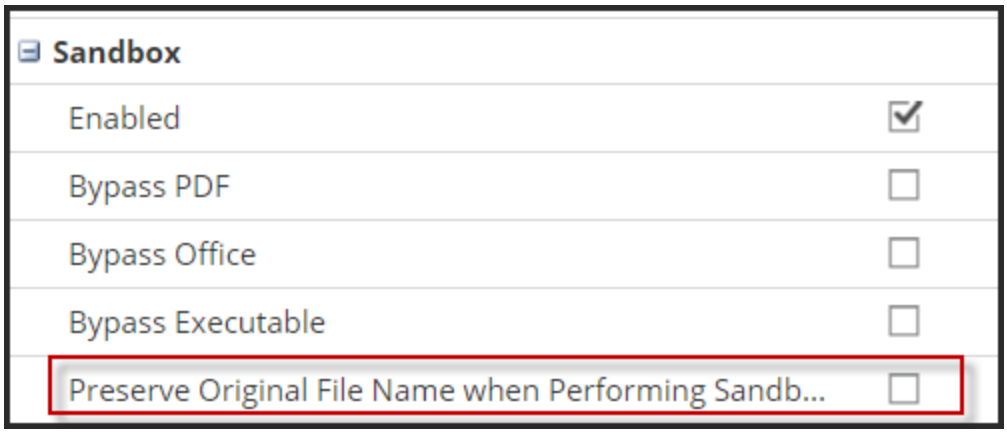
1. To enable or disable Community analysis entirely, click the **Enabled** checkbox. The default value is **Disabled**.
2. To configure handling of PDF, Microsoft Office, and Windows PE files in a session, select the specific checkboxes - **bypasspdf**, **bypass office, bypass executable.**
3. Click **Apply** to save the changes and put them into effect immediately as Malware Analysis receives new packets.

Configure Sandbox Analysis Scoring

By default, the sandbox module is disabled and MS Office and PDF files are prevented from being processed. The intent is to set to the most restrictive settings to force the user to specify whether or not potentially sensitive information is sent outside of the network for processing. If a document type is not prevented from being processed, the entire file (not just the hash) is sent to the destination sandbox server.

In addition, you can choose to preserve the original file name when performing sandbox analysis.

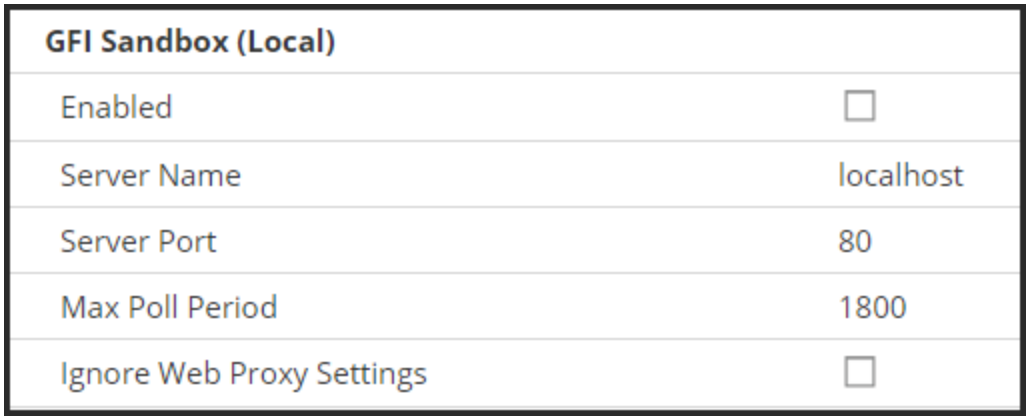
Note: If you do not specify the **Preserve Original File Name when Performing sandbox Analysis** parameter, NetWitness hashes the files.



When you enable the sandbox module, you must specify whether or not the sandbox processing is performed using a local GFI sandbox, a local ThreatGRID sandbox, or a cloud version of the ThreatGRID sandbox. The cloud version of the ThreatGRID sandbox is provided directly by ThreatGRID and requires an activation key to be obtained from ThreatGRID and configured in the ThreatGRID tab.

GFI Sandbox Settings

To use a locally installed GFI sandbox, you must enable GFI and supply the Server Name and Server Port of the GFI sandbox Server. The Max Poll Period and Polling Interval determine how long to wait for a submitted sample to finish processing and how often to check the status (in seconds). The Ignore Web Proxy Settings option allows you to indicate that you want Malware Analysis to bypass a web proxy when making this connection. If no Web Proxy has been configured in Malware Analysis, the setting is ignored.



ThreatGRID Sandbox Settings

Note: Before enabling ThreatGRID scoring, a ThreatGRID-supplied Service Key must be configured so that ThreatGRID can recognize that samples submitted from this site are legitimate. Use NetWitness to register for a ThreatGRID API key, then you can enable and configure a locally installed ThreatGRID sandbox or the ThreatGRID Cloud sandbox. Refer to the following detailed task: Register for a ThreatGRID API Key.

The Ignore Web Proxy Settings allows you to indicate that you want Malware Analysis to bypass a web proxy when making this connection. If no Web Proxy has been configured in Malware Analysis, the setting is ignored.

To configure sandbox scoring, in the Modules Configuration section:

1. To enable or disable sandbox analysis entirely, click the **Enabled** checkbox. The default value is **Disabled**.
2. To configure handling of PDF, Microsoft Office, and Windows PE files in a session, select any of the three checkboxes **Bypass PDF**, **Bypass Office, Bypass Executable.**
3. Configure the active sandbox vendor. You have three options:
   a. To use a locally installed instance of the GFI sandbox, provide the Server Name and Server Port of the GFI sandbox Server, the Max Poll Period and Polling Interval, and optionally, select the **Ignore Web Proxy** checkbox.
   b. To use a locally installed instance of ThreatGRID, enable ThreatGRID scoring, provide the ThreatGRID Service Key and optionally, select the **Ignore Web Proxy** checkbox.
   c. To use the ThreatGRID Cloud, you must first register for a ThreatGRID API key. Then enable ThreatGRID scoring, provide the ThreatGRID Service Key, enter the URL for the ThreatGRID server (https://panacea.threatgrid.com), and optionally, select the **Ignore Web Proxy** checkbox.
4. Click **Apply**.
   The changes become immediately effective.

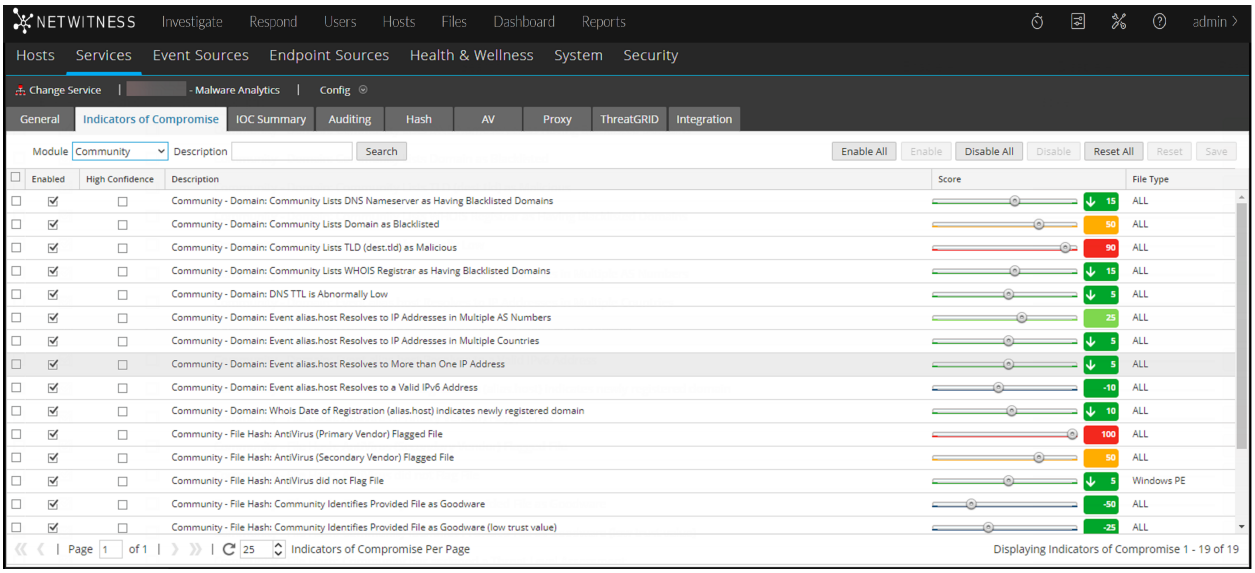
Configure Indicators of Compromise

The Indicators of Compromise (IOC) for the Malware Analysis scoring modules are configured since, each Malware Analysis scoring module -- Network, Static, Community, Sandbox, and YARA -- has a default set of Indicators of Compromise (IOCs) that it uses to evaluate the file and session data in order to assess the likelihood of malware being present.

Each IOC is assigned a numeric score weighting between -100 (good) and 100 (bad). When an IOC triggers, the numeric score weighting is factored into the total score for the session or file being analyzed. The individual score weightings for all matched IOCs are aggregated to produce the resulting score for each session or file. The aggregated score is adjusted to ensure that it does not exceed the valid score range (-100 through 100).

Note: The score weighting assigned to an IOC is not always the explicit score value that is aggregated (it is not a simple addition of score weights for each IOC that triggers). Instead, the IOC's score is a weighting or indicator of importance that is factored into calculating an overall score.

The Indicators of Compromise (IOC) configuration settings for Malware Analysis are in the Service Config view > Indicators of Compromise tab. Below is an example of the tab.



Using the **Community - File Hash: AntiVirus (Primary Vendor) Flagged File** IOC as an example, the IOC's score weighting could be set to 100. However, Malware Analysis dilutes this value based on the percentage of primary AV vendors that agree if the sample is malicious. The closer to 100% of the vendors who agree that the sample is malicious, the closer to the full 100 points are used in aggregating a score. As the percentage drops closer to 0%, the proportion of the full 100 points used in the aggregated score drops.

IOCs use logic implemented natively in Malware Analysis. You cannot adjust the logic. Instead, you can only adjust the IOC to increase or decrease its impact on scoring, to indicate a confidence setting, or to turn the IOC on or off. The typical scenario is to adjust a limited set of IOC score weighting values downward for IOCs that are inflating the final score and causing false positive analysis results. An extreme version of tuning would be to disable the IOCs entirely if they consistently contribute to false positive results. Additionally, the flexibility exists to allow you to disable all IOCs and to choose a select few to leave enabled. For example, all IOCs can be disabled with the exception of a select few IOCs that detect AntiVirus matches. Using Malware Analysis in this extremely limited configuration, you can reduce results in Malware Analysis such that only known A/V matches generate results.

You can configure this functionality in several ways:

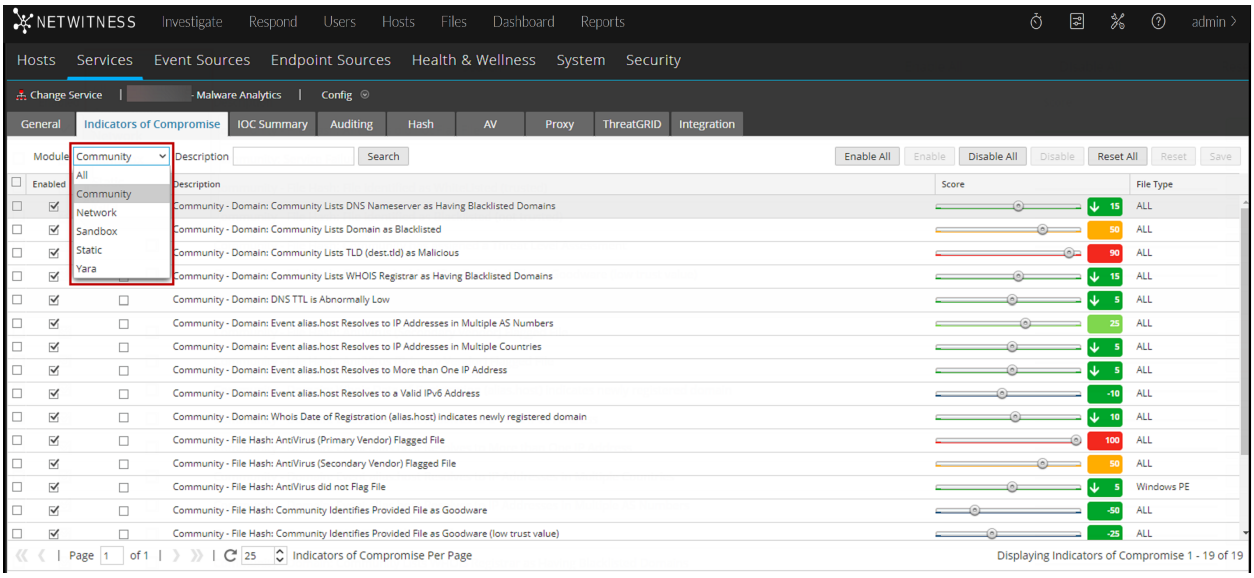
- Disable IOCs so that they are not evaluated as part of the scoring module to which they are assigned.
- Adjust the score weight for an IOC such that its impact on the aggregated score is increased or decreased.
- Mark IOCs that you expect to be strong indicators of malware and display a high-confidence (HC) flag on sessions that triggered these IOCs in the Malware Analysis results.
- Customize score and confidence settings uniquely to the file type being analyzed. Each IOC is pre-assigned a file type to which it is applied. Possible values are **ALL**, **PDF**, **MS Office**, and **Windows PE**. The IOC with the most applicable file type is used during file-based analysis. For example, if a PDF is analyzed, an IOC with a file type set to **PDF** will be chosen rather than the same IOC with a file type set to **ALL.** If no file-type specific match is found, the IOC with a file type set to **ALL** is selected.
- Search for rules to display in the grid based on a match to the rule description.

Filter Displayed IOCs by Module

You can filter the displayed IOCs by scoring module: one of the four built-in modules or YARA. YARA-based IOCs are interleaved with the native IOCs with each category. Although the YARA IOCs are not identified as such in the other views, you can select YARA from the Module selection list to see a list of YARA rules.

To view the IOCs for one or the four scoring modules or for YARA:

1. Go to **(Admin) > Services**.
2. Select a Malware Analysis service.
3. In the row, select **> View > Config**.
4. Click the **Indicators of Compromise** tab.
5. In the **Module** selection list, select **All**, **NextGen**, **Static**, **Community**, **Sandbox**, or **YARA**.
   The configured rules and settings for the module are displayed.



Filter Displayed Modules to Show Only Modified Modules

The **Indicators of Compromise** tab visually identifies IOCs that are locally modified. When an IOC has been modified, for example, the score weight has been changed, and the name is displayed in red and includes a modification indicator appended to the IOC name. The modification indicator is ++ and can be used as a filtering mechanism when searching for IOCs.
To limit the display to locally modified IOCs:

1. In the **Description** field, enter ++.
2. Click **Search**.
   The view is filtered to show only modified IOCs.

Enable and Disable IOCs for a Scoring Module

When an IOC is disabled, it no longer impacts the aggregate score for the scoring module to which it belongs. If the IOC has multiple instances (differentiated only by file type), disabling a more file-type specific IOC results in use of the more file-type agnostic version of the IOC in scoring.
For example, if the same IOC exists as file type **ALL** and file type **Windows PE**, disabling the **Windows PE** instance of the IOC causes the **ALL** version to be used in scoring. In order to disable the IOC entirely for **Windows PE**, while leaving the IOC enabled for other file types, set the score weighting of the **Windows PE** instance of the IOC to a value of zero as described below. This leaves the IOC enabled for Windows PE files (although it has a zero weighting and is suppressed from being displayed in analysis results), while not affecting the other file types. The remaining file types will continue to use the **ALL** instance of the IOC.

To enable or disable an IOC so that it no longer factors into a scoring module:

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service, and in the row select  > **View > Config**.
3. Click the **Indicators of Compromise** tab.
4. In the **Module** selection list, select a scoring module: **All**, **Community**, **Network**, **Sandbox**, **Static**, or **YARA**.
   The configured rules and settings for the module are displayed.
5. Do one of the following:
   a. Click the **Enabled** checkbox in the column next to a rule that you want to enable.
   b. Select one or more rules, and click **Enable** or **Disable** in the toolbar.
   c. To toggle between Enabled and Disabled for all rules displayed on the page, click the **Enabled** checkbox in the column title.
   d. To enable or disable all rules for the scoring module, click **Enable All** or **Disable All** in the toolbar.
6. To save the changes to the page, click **Save** in the toolbar.

> **Note:** Rules that have changed settings are displayed with a red corner. If you navigate to another page of rules before saving, all changes to this page are lost.

Adjust the Score Weight for an IOC

Adjusting the score weight for an IOC increases or decreases the IOC's overall impact on the aggregate score for the module in which it is configured. To raise or lower the overall impact of the IOC, reduce the current value to a new setting.

- Values ranging from -100 to -1 indicate that the session or file being analyzed is not likely to be malware (-100 being the least likelihood).
- Values ranging from 1 to 100 indicate a likelihood that the file or session being analyzed is malware (100 being the highest likelihood).
- Setting the value to zero leaves the IOC enabled, but causes the IOC to no longer impact the aggregate score and suppresses the IOC from being displayed in analysis results. Setting the value to zero is a method of disabling a file-type specific instance of an IOC while leaving the original file-type agnostic instance of the rule intact for scoring of the remaining file types.

To adjust the score weight:

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service.
3. In the row, select  **> View > Config**.
4. Click the **Indicators of Compromise** tab.
5. In the **Module** selection list, select a scoring module: **All**, **Network**, **Static**, **Community**, **Sandbox** or **YARA**.
   The configured rules and settings for the module are displayed.
6. Do one of the following:
   a. Drag the score slider left or right to decrease or increase the score weight.
   b. Click directly on the displayed score weight and enter a new score weight.
7. To save the changes to the page, click **Save** in the toolbar.

Note: Rules that have changed settings are displayed with a red corner. If you navigate to another page of rules before saving, all changes to this page are lost.

Set the High Confidence Flag for an IOC

The High Confidence setting is used as a method of flagging specific IOCs as high confidence indicators that malware is present. As an example, the **Community - File Hash: AntiVirus (Primary Vendor) Flagged File** IOC has a low probability of being a false positive, combined with a high probability of being an accurate measurement of malware being present. By flagging this IOC (and others) as High Confidence, you can use a filter in the Malware Analysis results to limit display to only those sessions that include one or more high confidence rules. By doing so, the display is limited to a smaller subset of results whose accuracy is accorded a higher degree of confidence. Displaying results not limited to high confidence IOCs still allows you to review results that are more grey in nature. This provides for results that are less prone to false negative results. Choosing to filter or to not filter results based on confidence level has a valid use case in the NetWitness workflow.
To set the High Confidence flag:

1. In the **Indicators of Compromise** tab, select a scoring module from the **Module** selection list: **All**, **Network**, **Static**, **Community**, **Sandbox**, or **YARA**.
   The configured rules and settings for the module are displayed.
2. Click the **High Confidence** checkbox in the column next to a rule that you want to flag or unflag as highly likely to indicate the presence of malware in a session when matched.
3. To save the changes to the page, click **Save** in the toolbar.

Note: Rules that have changed settings are displayed with a red corner. If you navigate to another page of rules before saving, all changes to this page are lost.

Reset IOCs to Default Settings

1. In the **Indicators of Compromise** tab, select a scoring module from the Module selection list: **All**, **Network**, **Static**, **Community**, **Sandbox**, or **YARA**.
   The configured rules and settings for the module are displayed.
2. If you want to reset all rules on the current page to their default settings, click **Reset** in the toolbar.
3. If you want to reset all rules for the selected scoring module to default settings, click **Reset All** in the

toolbar.

4. To save changes to the page, click **Save** in the toolbar.

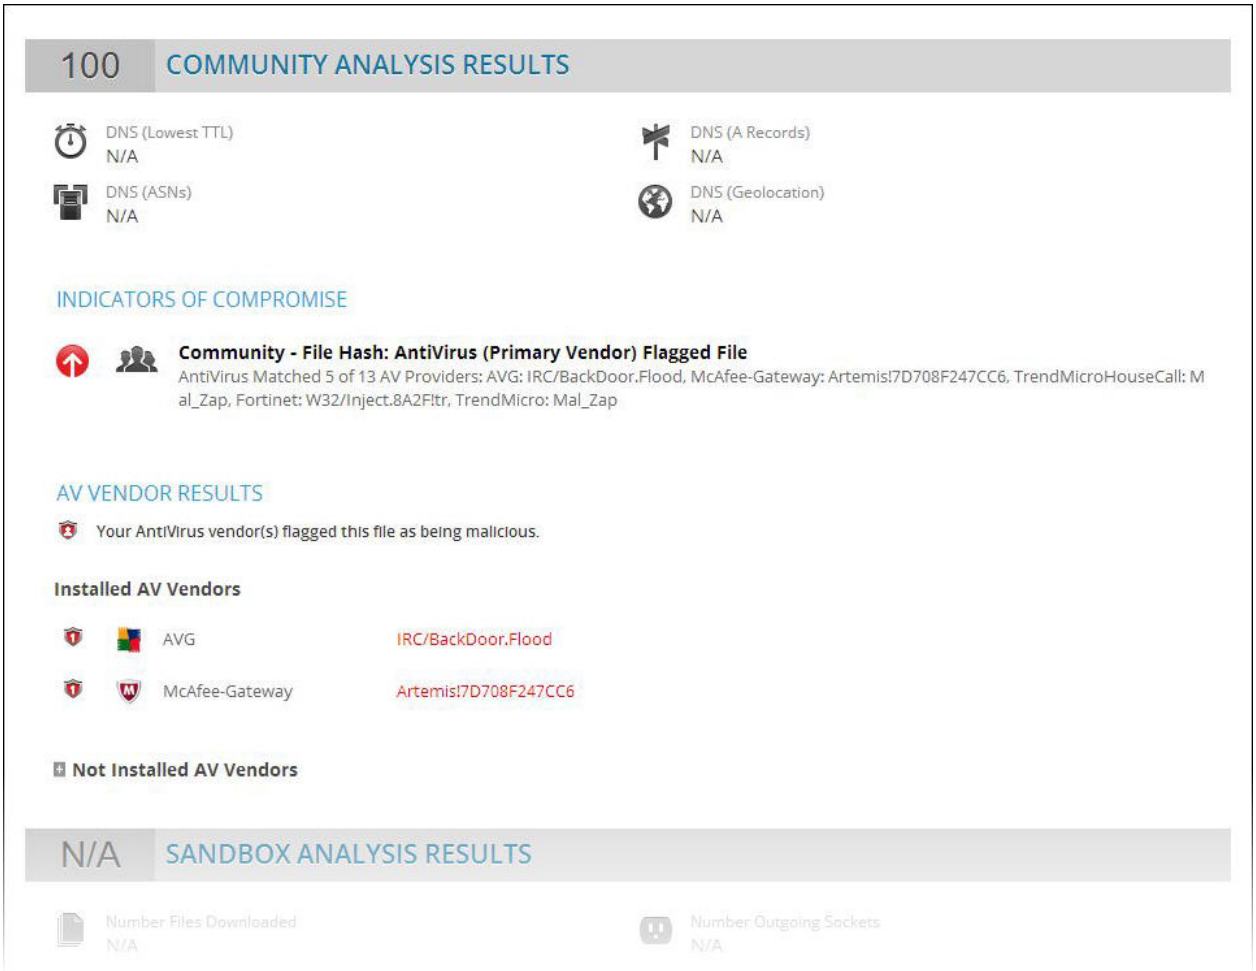
Configure Installed Antivirus Vendors

You can compare file analysis results from your installed antivirus (AV) vendors versus community results from the Malware Analysis knowledge base. While a file is being analyzed by community analysis, Malware Analysis checks an antivirus knowledge base to determine if the sample is already known to be malicious. If the file is known to be malicious, NetWitness flags the file to indicate whether a primary antivirus vendor or a secondary antivirus vendor identified the sample. NetWitness classifies vendors as primary and secondary to indicate the level of reputation the vendors have in the industry, and Indicators of Compromise factor the reputation into scoring. For example, detection made solely by secondary antivirus vendors may score less than detection by primary vendors.

Note: When choosing AV vendor software to install on your network, it is highly recommended that you include at least one from NetWitness Primary Vendors list. For more information, please see the Supported Antivirus Vendors.

You can identify the antivirus vendors installed on your network to NetWitness. NetWitness compares the antivirus results during community analysis against the results from the installed vendors selected in the AV tab. If a match is detected, the file being analyzed is flagged to indicate that your locally installed primary or secondary antivirus software detected the sample.

The example below shows the community analysis results for a file that had a score of 100. Under **Indicators of Compromise**, you can see that the file was flagged by the listed AV vendors in the Community. Under **AV Vendor Results**, NetWitness indicates whether the AV vendors installed in your environment flagged the file as malicious. If your installed AV vendors detected the virus, the name of the malware is displayed. If your installed AV vendors did not detected the virus, **--Not detected--** is displayed next the AV vendor name. Under **Not Installed Vendors**, you can click + to expand the section and see if other vendors not installed on your system detected the virus.



Identify Installed AV Software

To identify Antivirus software installed on your network:

1. Go to **(Admin) > Services.**
2. Select a Malware Analysis service, and in the row select **> View > Config**.
3. In the **Service Config View**, select the **AV** tab.
4. Select the checkbox next to each antivirus vendor (primary and other) whose software is installed on your network.
5. To save the changes, click **Apply**.
   The Community Analysis results will indicate whether your software flagged an event.
6. (Optional) If you want to reset the list of installed AV software to the default value (none), click **Reset**.
   All selections are removed.
7. To save changes, click **Apply**.

Enable Community Analysis

An Administrator can enable community analysis. For Community analysis, new malware detected on the network is pushed to the Netwitness Cloud for checking against Netwitness's own malware analysis data and feeds from the SANS Internet Storm Center, SRI International, the Department of the Treasury and VeriSign. To enable Community analysis, you must register with the Netwitness cloud and test connection to the cloud, then to test the connection between the Netwitness cloud and the service you have configured for continuous scanning.

A complete description of analysis methods is provided in How Malware Analysis Works.

1. Go to  **(Admin) > Services.**
2. Select a Malware Analysis service, and in the row select  **> View > Config**.
3. In the **Service Config View**, select the **Integration** tab.
4. Scroll down to **Netwitness Cloud Connection Test and Registration**.

   NetWitness tests communications with the site at `https://cloud.netwitness.com`. If your company uses a proxy for outbound traffic, please check your Proxy settings. A valid connection is required in order to register with the Netwitness Community Service.
5. Enter your company name and contact email. Click **Register**.

   If all required fields are complete, your registration is completed. The label on the button used to register changes to Update.
6. To verify that the Malware Analysis Service can connect to the Core service selected for continuous scanning, click **Test Connection**.

   NetWitness initiates a check based on the Source Host, Source Port, Username, and User Password specified in the General tab. When the test executes successfully, analysts are able to see Community Scoring in Malware Analysis.

(Optional) Configure Auditing on Malware Analysis Host

This topic introduces the configurable features of the Malware Analysis auditing log and the procedures for configuring the features. Malware Analysis is capable of generating auditing alerts based on configured score module thresholds. Once the analysis score for a file in an analysis session meets or exceeds the configured threshold(s), an auditing alert is generated. Thresholding allows sessions and files that score high enough to be likely malware candidates to automatically generate an alert.

Alerts can be configured to be formatted as SNMP, Syslog or File entries. Supporting various audit formats provides a method for external systems to ingest auditing events based on their capability of parsing the supported formats.

In addition to auditing analysis sessions, the following events will trigger an audit alert:

- User login successes and failures
- Changes to system configuration settings

- Server restart
- Server version upgrade and install

The Auditing configuration settings for Malware Analysis are in the Service Config view > Auditing tab.

Configure the Auditing Threshold

The sole purpose of the thresholds is to specify the criteria that must be reached prior to an alert being generated for an analyzed session/file. If auditing is enabled, each scored file/session is examined to determine if the score in each score module meets or exceeds the configured auditing threshold. If so, an alert is generated using the configured audit alert format (e.g., SNMP, Syslog or File). For example, by configuring SNMP and setting the Community Threshold to 90, all sessions/files that score 90 or higher in the Community Score module generate an SNMP trap. If all of the thresholds are set to 90, then an alert is not generated unless a session or file scores 90 or higher in the Network, Static, Community and sandbox score modules.

To configure the auditing threshold:

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service, and select  **> View > Config**.
3. In the **Services Config** view, click the **Auditing** tab.
4. In the **Auditing Thresholds** section:
   a. Set the threshold for the **Community**, **Static**, **Network**, and **Sandbox** by doing one of the following for each scoring module:
      - In the slider, click and drag the handle in either direction.
      - In the value field, type a number between 0 and 100, inclusive.
   b. (Optional for 10.3 SP2) Select one or more triggers to record a message and deliver it through all enabled auditing methods.
   c. Click **Apply**.
      - The threshold setting becomes effective immediately for all enabled auditing methods: SNMP, File, and Syslog.
      - The recorded messages are sent through all enabled auditing methods: SNMP, File, and Syslog.

Configure Incident Management Alerting

When enabled, Malware Analysis alerts feed into the NetWitness Respond workflow.

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service, and select select  **> View > Config**.
3. In the **Services Config** view, select the **Auditing** tab.

4. In the **Respond Alerting** section, select the **Enabled** checkbox and click **Apply**.
   Alerting becomes effective immediately.

Configure SNMP Auditing

The Simple Network Management Protocol (SNMP) is an Internet-standard protocol for managing services on IP networks. When SNMP auditing is enabled, Malware Analysis can send an audit event as an SNMP trap to a configured SNMP trap host. In addition to the score and event ID, the alert includes all session meta as well as generated meta data. This is useful for users who want to feed event data to third-party systems.

To configure SNMP auditing:

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service, and select  **> View > Config**.
3. In the **Services Config** view, select the **Auditing** tab.
4. In the **SNMP Auditing** section, click the checkbox to enable SNMP auditing.
5. Configure the SNMP server name and port.
6. Configure the SNMP version and trap OID for sending traps.
7. Configure the Malware Analysis community, and retry and timeout parameters when sending traps.
8. Click **Apply**.
   The SNMP auditing settings become effective immediately.

Configure File Auditing Settings

When file auditing is enabled, the audit log file is kept in the Malware Analysis Home Directory. The default location for this log file is:

`/var/lib/netwitness/malware-analytics-server/spectrum/logs/audit/audit.log`.

As each log reaches the maximum file size, it is archived and a new log is created. The size of these audit logs and their number are both configurable.

Caution: Avoid setting the max file size and archive file count too high, because it may have an adverse effect on the available disk space on the Malware Analysis appliance.

To configure the file auditing settings:

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service, and select  **> View > Config**.
3. In the **Services Config** view, select the **Auditing** tab.
4. In the **File Auditing** section, click the checkbox to enable file auditing.
5. (Optional) Set the Archive File Count and Max File Size.
6. Click **Apply**.
   The file auditing settings become effective immediately.

Configure Syslog Auditing Settings

When enabled, Syslog provides auditing through the use of the RFC 5424 syslog protocol. Regulations, such as SOX, PCI DSS, HIPAA, and many others are requiring organizations to implement comprehensive security measures, which often include collecting and analyzing logs from many different sources. Syslog has proven to be an effective format to consolidate logs, as there are many open source and proprietary tools for reporting and analysis.

In addition to the score and event ID, the syslog includes all session meta as well as generated meta data. This is useful for users who want to feed event data to third-party systems.

To configure the syslog auditing settings:

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service, and select  **> View > Config**.
3. In the **Services Config** view, select the **Auditing** tab.
4. In the **Syslog Auditing** section, click the checkbox to enable syslog auditing.
5. Configure the host where the target syslog process is running and the port on the host where the syslog process is listening.
6. Configure the facility, encoding, format, max length, and timestamp for outgoing syslog messages.

   Note: (Optional) Configure Identity String to prepend to syslog alerts.
   For CEF format, please refer to Create Custom Alert in CEF Format for additional considerations.

7. Click **Apply**.

   The syslog auditing settings become effective immediately.

(Optional) Configure Hash Filter

This topic introduces hash filters as a method of marking files in Malware Analysis that are known to be good or known to be bad. Hash filtering allows you to maintain a list of known good or known bad file hashes. In the Hash tab, you can fine tune Malware Analysis event analysis based on file hashes. When a file hash is marked as Good, Malware Analysis does not analyze the file the next time it is seen. When a file hash is marked as Bad, Malware Analysis automatically raises the file's community score by a large number of points. Malware Analysis still analyzes the file, just in case new information becomes available.

Note: If an event contains a single file and that file's hash is marked as Good, Malware Analysis filters the entire event and you do not see it in Malware Analysis results.

To add hash filters to the hash list, you can use either of these manual methods:

1. Context menu add in the Event Detail view: Right-click on a file, and a context menu allows marking of the hash for the selected file as Good (Normal) or Bad (Malicious).
2. Hash tab toolbar: Click **Add** in the Hash tab to add a file hash, file size, and optionally, mark the hash as trusted.

There is also an automated method to add hash filters to Malware Analysis by importing a hash list in bulk from the watched folder. Hashes imported through the watched folder do not appear in the hash list. With bulk importing and the watched directory (/var/netwitness/malware-analytics-server/spectrum/hashWatch) on the Malware Analysis server set up, copy a hash list into the watched folder to be automatically imported into the system. Hashes imported using the bulk import method overwrite hashes that were previously imported through the watched folder.

View the Hash List

To view the Hash List:

1. Go to  **(Admin) > Services**.
2. In the Services view, select a Malware Analysis service, and select  **> View > Config**.
3. Select the **Hash** tab.

   The hash list is displayed in the Hash tab. Only file hashes that have been added using one of the methods are displayed.

Add a File Hash to the Hash Filter

To add a file hash to the hash filter:

1. In the **Hash** tab, in the toolbar, click **Add**.

   The Add Hash dialog is displayed.
2. If the hash is trusted, select **Trusted**.
3. Enter the MD5 hash and the file size in bytes.
4. Click **Save**.

   The file hash is added to the hashes and used to perform hash filtering in Malware Analysis.

Mark a Hash as Trusted or Untrusted

To mark a file hash as trusted or untrusted:

1. In the **Hash** tab, to toggle between trusted and untrusted, click in the **Trusted** column for the hash.
2. In the toolbar, click **Save Edit**.

Delete a Hash from the Hash Filter

To delete a hash from the hash filter:

1. In the **Hash** tab, select one or more hashes that you want to remove from the hash filter.
2. In the toolbar, click **Delete**.

   A dialog requests confirmation and offers an opportunity to cancel.
3. To confirm the deletion, click **Yes**.

The file hash is deleted from the grid and no longer used to perform hash filtering in Malware Analysis.

Search for a File Hash

In the Hash tab, you can search for a file hash that is displayed in the grid. In the MD5 field, type the file hash for which you are searching, and click **Search**. The list of files that contain the hash is displayed in the grid.

Import a Hash List Using the Watched Folder

To import a hash list from the watched directory, the hash list must be in the specified format and must be sorted on md5. You can drop a file formatted as described below into a folder (/var/netwitness/malware-analytics-server/spectrum/hashWatch) on the Malware Analysis appliance, and it is automatically imported into the local hash database. This is the only way to import file hashes into. An additional use case is to allow a system administrator to expose the watched directory to some process that would push a file to this directory. This is a bulk import method designed to handle a high volume of hash imports.

This is a csv-formatted file with no spaces between the data in each row. The assumption with the data in the hash list is that there are no duplicates. Duplicates are ignored during processing. If duplicate hashes are encountered, the log file will display the following message to indicate the number of duplicate hashes contained in the file:

2013-08-09 09:46:00,674 [jobExecutor-2(HashFileWatch)] INFO com.netwitness.malware.core.scheduler.jobs.HashFileWatch - Processing - /var/lib/rsa>malware/hashWatch/test.csv
2013-08-09 09:47:56,619 [jobExecutor-2(HashFileWatch)] INFO com.netwitness.malware.core.services.file.hash.HashServiceImpl - Skipped 21 Duplicate Hashes Already on File
2013-08-09 09:48:06,638 [jobExecutor-2(HashFileWatch)] INFO com.netwitness.malware.core.scheduler.jobs.HashFileWatch - Processed - /var/lib/rsa>malware/hashWatch/test.csv

Below is an example of a hash list in the default file format.

```
[BeginFileExample]
392126E756571EBF112CB1C1cdEDF926,98865,True
0E53C14A3E48D94FF596A2824307B492,2226,True
176308F27DD52890F013A3FD80F92E51,42748,False
9B3702B0E788C6D62996392FE3C9786A,32768,False
937ADE76A75712B7FF339403B4FCB5A6,4821,False
B47139415F735A98069ACE824A114399,1723,False
E6CAF205E602CFA9A65663DB1A087874,704,False
680CA0BCE1FC7BC4136ADF4E210869C5,2075,False
[EndFileExample]
```

A NetWitness configuration file (**/var/netwitness/malware-analytics-server/spectrum/conf/hashFileWatchConfig.xml**) specifies the format and options in the hash list import process. Below is a listing of the configuration file.

```
<config>
 <enabled>true</enabled>
 <distributedCacheEnabled>true</distributedCacheEnabled>
 <watchDirectory>//var/lib/rsamalware/hashWatch</watchDirectory>

<processedDirectory>/
var/lib/rsamalware/hashWatch/processed</processedDirectory>

<erroredDirectory>/
var/lib/rsamalware/hashWatch/error</erroredDirectory>
 <md5Col>0</md5Col>
 <fileSizeCol>-1</fileSizeCol>
 <isTrustedCol>1</isTrustedCol>
 <isTrust>false</isTrust>
 <ignoreFirstLine>false</ignoreFirstLine>
 <frequencyInMinutes>1</frequencyInMinutes>
 <isGzipCompressed>false</isGzipCompressed>
 </config>
```

Line	Description
<md5Col>0</md5Col>	The location of the md5 hash in each entry. The default value is position **0**, or the first position.
<fileSizeCol>1</fileSizeCol>	The location of the hash size in each entry. The default value is position **1**, or the second position. If the hash size is not included in the csv file, the value must be **-1**.
<isTrustedCol>2</isTrustedCol>	The location of the Trusted Column in each entry. The default value is position **2**. If the Trusted parameter is not included in the csv file, the value must be **-1**.
<isTrust>false</isTrust>	The default assumption for **Trusted** in each entry is **false**.
<ignoreFirstLine>false</ignoreFirstLine>	The presence or absence of a header in the hash. The default value is **false**. If the hash has a header, the value must be set to **true**.
<frequencyInMinutes>1</frequencyInMinutes>	The interval between checks by NetWitness in the watched directory. The default value is **1** minute.
<isGzipCompressed>false</isGzipCompressed>	The hash is compressed using Gzip. The default value is **false**. If the hash is Gzip compressed, the value must be set to **true** here.

When the hash list has been imported, the system log has entries similar to this:

```
2013-04-11 03:22:00,597 [jobExecutor-9(HashFileWatch)] INFO
com.netwitness.malware.core.scheduler.jobs.HashFileWatch -
Processing - /var/lib/rsamalware/spectrum/hashWatch/simpleHash.csv
 2013-04-11 03:22:00,600 [jobExecutor-9(HashFileWatch)] INFO
com.netwitness.malware.core.scheduler.jobs.HashFileWatch - Processed
- /var/lib/rsamalware/spectrum/hashWatch/simpleHash.csv
```

If there is a problem loading the file, the system log has entries similar to this:

```
2013-04-11 03:17:00,597 [jobExecutor-4(HashFileWatch)] INFO
com.netwitness.malware.core.scheduler.jobs.HashFileWatch -
Processing - /var/lib/rsamalware/spectrum/hashWatch/simpleHash.csv
 ... Verbose log
 2013-04-11 03:17:00,632 [jobExecutor-4(HashFileWatch)] INFO
com.netwitness.malware.core.scheduler.jobs.HashFileWatch - Error
Processing - /var/lib/rsamalware/spectrum/hashWatch/simpleHash.csv
```

To import a hash list using the watched folder method:

1. Copy the hash lists that you want to import into the **/var/netwitness/malware-analytics-sever/spectrum/hashWatch** directory.
   Malware Analysis automatically watches this folder and processes files placed there.
   Malware Analysis adds every hash found in the hash lists to the hash filter.
   If there are processing errors, they are logged in **/var/netwitness/malware-analytics-sever/spectrum/hashWatch/error**
   Processed files are cataloged in **/var/netwitness/malware-analytics-sever/spectrum/hashWatch/processed**
   Processed files are not removed from the hashWatch directory.
2. After importing hashes in bulk, the System Administrator can use a cronjob to clean up old processed files.

(Optional) Configure Malware Analysis Proxy Settings

This topic describes the configuration of a web proxy for communicating with the Netwitness Cloud service and local ThreatGRID or GFI service. The settings in the Service Configuration view > Proxy tab set up communication by web proxy, which Malware Analysis can use to communicate with Netwitness Cloud for community analysis and sandbox analysis. Once the proxy is configured:

- Malware Analysis communicates via web proxy with the Netwitness Cloud for community analysis.
- Malware Analysis communicates via web proxy with the configured ThreatGRID or GFI sandbox service. Using a web proxy may negatively affect performance. ThreatGRID and GFI configuration sections in the General tab have an option to ignore the web proxy and communicate directly with the sandbox to improve performance.

Configure the Web Proxy

To configure the web proxy for Malware Analysis:

1. Go to  **(Admin) > Services** view.
2. Select a Malware Analysis service, and select  **> View > Config**.
3. In the **Services Config** view, select the **Proxy** tab.
4. To enable the proxy, select the **Enabled** checkbox.
5. (Optional) To automatically detect proxy settings for the NetWitness Server, select the checkbox.
   The proxy host and proxy port fields are autofilled.

6. If you want to use a different proxy, enter the **Proxy Host** and **Proxy Port**.
7. Enter the username and password used to log on to the proxy host.
8. (Optional) Select **SSL**, if the proxy host communicates over SSL.
9. Click **Apply**.
10. Restart the Malware service.

Note: Malware Analysis does not support NTML web proxy authentication.

(Optional) Register for a ThreatGRID API Key

This topic provides the procedure for obtaining a trial ThreatGRID API key for use in the ThreatGRID Cloud sandbox. Before enabling ThreatGRID as the sandbox service in the sandbox module, a ThreatGRID-supplied Service Key must be configured so that ThreatGRID can recognize that samples submitted from this site are legitimate.

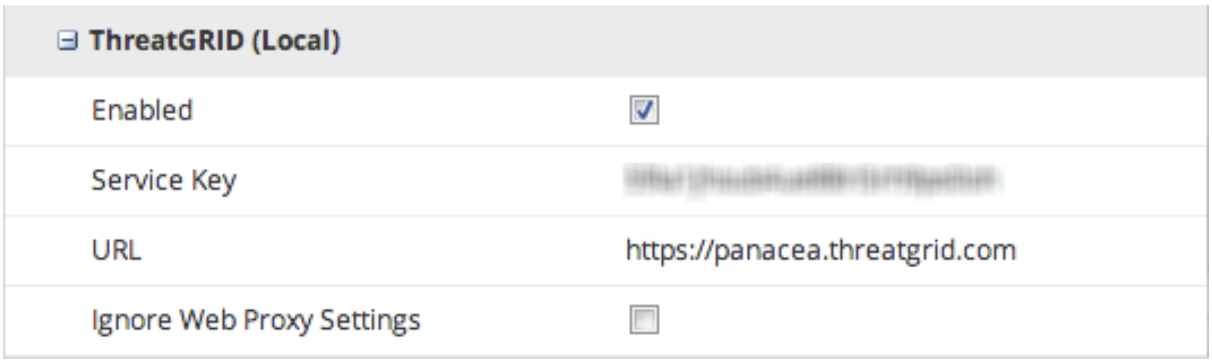
If you do not have a ThreatGRID-supplied Service Key, you can obtain a key using this tab. The key is provided on a trial basis.

When you fill in your user information and click **Register**, a key is displayed in this tab, and automatically added to the ThreatGRID configuration in the **General** tab. In a few minutes, you will receive an email from ThreatGRID containing a link to their page where you can log on. After you agree to the license terms on the ThreatGRID page, you can submit files for analysis, and ThreatGRID will recognize files that Malware Analysis submits for sandbox analysis.

To obtain a Trial ThreatGRID API key:

1. Go to  **(Admin) > Services**.
2. Select a Malware Analysis service, and select  **> View > Config**.
3. In the **Services Config** view, select the **ThreatGRID** tab.
4. Enter your full name, job title, organization name, and email address.
5. In the User Id and Password field, create a user ID and password for logging on to ThreatGRID.
6. Click **Register**.

   Your registration is sent to ThreatGRID and an API key is displayed below the Register button. The key is automatically filled in the **General** tab.
7. Select the **General** tab to confirm that the ThreatGRID configuration now includes the API key.



8. When you receive an email from ThreatGRID with a link where you can log on, log on and accept the terms of the agreement.

   Your trial of ThreatGRID begins and Malware Analysis can send five files per day to the ThreatGRID Cloud for sandbox analysis.

Additional Procedures for Configuring Malware Analysis

This topic provides procedures that an Administrator may perform to accomplish an objective that is not part of basic Malware Analysis setup. After Malware Analysis is configured, administrators may want to fine-tune the service and implement advanced customization; an example of this is implementing custom YARA content.

- Create Custom Alert in CEF Format
- Enable Custom YARA Content

Create Custom Alert in CEF Format

This topic provides instructions for creating custom alerts in Common Event Format (CEF) to send to a service that ingests events as CEF. This is an advanced configuration task, which requires sufficient knowledge to manually edit the configuration file: `/var/netwitness/malware-analytics-server/spectrum/conf/malwareCEFDictionaryConfiguration.xml`. Before editing the file, you must stop the Malware Analysis service in the operating system. The CEF Alert becomes active when you restart the Malware Analysis service.

The CEF Template

To send events to a service ingesting events as CEF, NetWitness runs them through a configuration file that serves as a CEF template before feeding the events to a correlation technology. You can tune the configuration file, which specifies the sequence and mapping of syslog fields in each alert.

The following example syslog message shows the CEF fields in the extensions section of the alert (following the last '|' in the alert). Each field can be configured to indicate the sequence (described in the Example section below). Fields can be excluded entirely from the alert via a configuration setting.

```
CEF:0|NetWitness|Spectrum|10.3.0.7995.1.0|Suspicious Event|Detected
suspicious network event ID 4 session ID n/a|2|static=100.0 nextgen=25.0
community=100.0 sandbox=25.0 file.name=myFile.exe file.size=1234556
file.md5.hash=DEADBEEFBABECAFEDEADBEEFBABECAFE
event.source=spectrum://admin@0:0:0:0:0:0:0:1:64563 event.type=MANUAL_
UPLOAD event.id=0 country.dst.code=-- country.dst=Unavailable
ip.src=0:0:0:0:0:0:0:1 ip.dst=0:0:0:0:0:0:0:1 event.uuid=f7a6155a-31de-
4fa6-ba16-41fb9a8e5f26 ...
```

Understand a Syslog Auditing File Entry

The description of the file structure is based on the following sample.

```
Feb 6 10:02:28 10.10.10.125 SpectrumServer125

CEF: 0|NetWitness|Spectrum|1.2.1.130|Suspicious Event|Detected suspicious
 network event ID 857 session ID 73|2|

static=100.0 network=29.0 community=8.0 sandbox=N/R

file.name=-CVE-00_DOC_2010-05-13_attachment.doc file.size=0
file.md5.hash=20a29259c0e5958afb2f50c4177bb307
```

```
com.netwitness.event.internal.id=73
com.netwitness.event.internal.uuid=37d2bad7-06bc-4b34-88e1-df43d9710204
alias.ip=10.25.50.149 client=Wget/1.11.4 Red Hat modified payload=108872
packets=136 country.dst=Private time=Fri Jan 27 10:09:25 EST 2012
threat.source=netwitness tcp.srcport=43580 action=get
com.netwitness.event.internal.source=http://QASpectrum2:50104/sdk filetype=rtf
alias.host=qa-fc12-149 eth.src=00:25:90:18:76:E2 ip.proto=6 tcp.flags=27
ip.src=10.25.50.61 tcp.dstport=80 threat.category=spectrum
eth.dst=00:0C:29:F8:50:2D lifetime=0 alert.id=nw32535 sessionid=73 medium=1
size=117864 content=spectrum.consume11 extension=doc
directory=/files/MALWAREMALWARE/OfficeDocs/DOC/ eth.type=2048
ip.dst=10.25.50.149 service=80 filename=-CVE-00_DOC_2010-05-13_attachment.doc
server=Apache/2.2.13 (Fedora) streams=2 referer=http://qa-fc12-
149/files/MALWAREMALW...fficeDocs/DOC/ risk.info=http client server version
mismatch
```

First Line

```
Feb 6 10:02:28 10.10.10.125 SpectrumServer125
```

Log Information	Description
Feb 6 10:02:28	The timestamp for the entry.
10.10.10.125	The source IP address for the event.
SpectrumServer125	The source hostname for the event.

Audit Common Event Format (CEF) Header

```
0|NetWitness|Spectrum|1.2.1.130|Suspicious Event|Detected suspicious network
event ID 857 session ID 73|2|
```

The audit CEF header is a pipe-separated listing of the following fields:

Log Information	Description
0	The ArcSight Common Event Format (CEF) version used for the audit syslog.
NetWitness	The service that created the syslog message.
Spectrum	Malware Analysis is the logger for the event.
1.2.1.130	Malware Analysis version.
event ID 857	Unique network event id for this event.
session ID 73	Core unique session id for the session that included this event.

Log Information	Description
2	Severity, an integer between 1 and 6 indicates the level of severity for the message. • 1 = INFORMATION_LEVEL • 2 = WARNING_LEVEL • 3 = ERROR_LEVEL • 4 = SUCCESS_LEVEL • 5 = FAILURE_LEVEL • 6 = AUDIT_FAILURE_LEVEL

Audit CEF Extension

```
static=100.0 network=29.0 community=8.0 sandbox=N/R
file.name=-CVE-00_DOC_2010-05-13_attachment.doc  file.size=0
file.md5.hash=20a29259c0e5958afb2f50c4177bb307
com.netwitness.event.internal.id=73
com.netwitness.event.internal.uuid=37d2bad7-06bc-4b34-88e1-df43d9710204
alias.ip=10.25.50.149 client=Wget/1.11.4 Red Hat modified payload=108872
packets=136 country.dst=Private time=Fri Jan 27 10:09:25 EST 2012
threat.source=netwitness tcp.srcport=43580 action=get
com.netwitness.event.internal.source=http://QASpectrum2:50104/sdk filetype=rtf
alias.host=qa-fc12-149 eth.src=00:25:90:18:76:E2 ip.proto=6 tcp.flags=27
ip.src=10.25.50.61 tcp.dstport=80 threat.category=spectrum
eth.dst=00:0C:29:F8:50:2D lifetime=0 alert.id=nw32535 sessionid=73 medium=1
size=117864 content=spectrum.consume11 extension=doc
directory=/files/MALWAREMALWARE/OfficeDocs/DOC/ eth.type=2048
ip.dst=10.25.50.149 service=80 filename=-CVE-00_DOC_2010-05-13_attachment.doc
server=Apache/2.2.13 (Fedora) streams=2 referer=http://qa-fc12-
149/files/MALWAREMALW...fficeDocs/DOC/ risk.info=http client server version
mismatch
```

Analysis Scores

The first entry in the audit CEF extension provides the four Malware Analysis scores for the event: Static, Network, Community, and Sandbox.

Log Information	Sample Value
static	100.0
network	29.0
community	8.0 A score of 0.0 can be a community score for the event or can indicate that no community services were enabled.

Log Information	Sample Value
sandbox	N/R N/R means not run. This indicates that the GFI sandbox was not enabled.

File Information

The next three entries provide file information: file name, size, and hash.

Log Information	Sample Value
file.name	-CVE-00_DOC_2010-05-13_attachment.doc
file.size	0
file.md5.hash	20a29259c0e5958afb2f50c4177bb307

Event Meta Data Retrieved by NextGen

The record continues with the Core meta data for the event. The meta data in the message depends on the event. The amount of data in the message is truncated to the maximum length in bytes configured in the Syslog Settings. The default value is 1024.

Log Information	Sample Value
com.netwitness.event.internal.id	73
com.netwitness.event.internal.uuid	37d2bad7-06bc-4b34-88e1-df43d9710204
alias.ip	10.25.50.149
client	Wget/1.11.4 Red Hat modified
payload	108872
packets	136
country.dst	Private
time	Fri Jan 27 10:09:25 EST 2012
threat.source	netwitness
tcp.srcport	43580
action	get
com.netwitness.event.internal.source	http://QASpectrum2:50104/sdk
filetype	rtf

Log Information	Sample Value
alias.host	qa-fc12-149
eth.src	00:25:90:18:76:E2
ip.proto	6
tcp.flags	27
ip.src	10.25.50.61
tcp.dstport	80
threat.category	spectrum
eth.dst	00:0C:29:F8:50:2D
lifetime	0
alert.id	nw32535
sessionid	73
medium	1
size	117864
content	spectrum.consume11
extension	doc
directory	/files/MALWAREMALWARE/OfficeDocs/DOC/
eth.type	2048
ip.dst	10.25.50.149
service	80
filename	-CVE-00_DOC_2010-05-13_attachment.doc
server	Apache/2.2.13 (Fedora)
streams	2
referer	http://qa-fc12-149/files/MALWAREMALWARE/OfficeDocs/DOC/
risk.info	http client server version mismatch

Edit the Configuration File

1. Stop the Malware Analysis service.
2. Edit the configuration file as described in the Example.

3. Start the Malware Analysis service.

   The Malware Analysis service begins processing alerts through the configuration file and sending CEF alerts to designated services.

Example

The configuration file can be used to dictate which fields appear in the resulting alert as well as the label associated with each field and the order in which the data fields appear. The configuration file is composed of one or more XML `MalwareCefExtension` blocks as shown in the example below. The ordering of these blocks in the configuration file implies the order of the data fields in the CEF alert.

In the example below, the CEF alert would include two data fields, `ip.src` followed by `ip.dst`. The `customKey` is used to indicate the labeling of the data field in the alert. This allows the user to choose a custom label in order to force the alerting format to better match the expectations of the alert consumer. In other words, the format can be tuned to prevent unwanted changes to an existing alert parser. Lastly, the `isDisplay` setting determines if the field is included in the alert output. This allows the user to turn off data fields without having to physically delete the `MalwareCefExtension` block from the configuration.

```
 <config>
 <malwareExtensionList>
<com.netwitness.malware.core.cef.MalwareCefExtension>
   <customKey>ip.src</customKey>
   <malwareKey>ip.src</malwareKey>
   <isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
   <customKey>ip.dst</customKey>
   <malwareKey>ip.dst</malwareKey>
   <isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
 </malwareExtensionList>
</config>
```

At the end of the configuration file are three additional settings that can be used to further tune the alert format. They are as follows:

Setting	Description
includesUnknownMeta	This true or false setting indicates if unknown data elements are included in the resulting alert. Any NextGen session meta can be considered for inclusion into a CEF alert. Because additional session meta can be introduced via authoring new NextGen parsers, meta that is not contained in the default configuration may be encountered. You can set includesUnknownMeta to true to include the unknown meta in the alert and label it using the NextGen meta key name. To force a custom key for the unknown meta, you must edit this file and add a new MalwareCefExtension to the dictionary. To omit unknown meta from the alert, set includesUnknownMeta to false.
displayNulls	This true or false setting indicates if values that are set to null are included in the alert. If displayNulls is set to false, the null value fields are omitted even if their MalwareCefExtension isDisplay property is turned on. This allows dynamic formatting of alerts to exclude null fields.
valueIfNull	This true or false setting allows you to specify a string placeholder (n/a by default) to be used as the value for any null valued fields. If displayNulls is set to true, then null valued fields are included in the alerts. Their value is set to the value specified in valueIfNull.

The following represents the default CEF configuration file. The default configuration file includes all default NextGen session meta.

```
<config>
  <malwareExtensionList>
   <com.netwitness.malware.core.cef.MalwareCefExtension>
   <customKey>static</customKey>
   <malwareKey>static</malwareKey>
   <isDisplay>true</isDisplay>
   </com.netwitness.malware.core.cef.MalwareCefExtension>
   <com.netwitness.malware.core.cef.MalwareCefExtension>
   <customKey>nextgen</customKey>
   <malwareKey>nextgen</malwareKey>
   <isDisplay>true</isDisplay>
   </com.netwitness.malware.core.cef.MalwareCefExtension>
   <com.netwitness.malware.core.cef.MalwareCefExtension>
   <customKey>community</customKey>
   <malwareKey>community</malwareKey>
   <isDisplay>true</isDisplay>
   </com.netwitness.malware.core.cef.MalwareCefExtension>
   <com.netwitness.malware.core.cef.MalwareCefExtension>
   <customKey>sandbox</customKey>
```

```
<malwareKey>sandbox</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>file.name</customKey>
<malwareKey>file.name</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>file.size</customKey>
<malwareKey>file.size</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>file.md5.hash</customKey>
<malwareKey>file.md5.hash</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>event.source</customKey>
<malwareKey>event.source</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>event.type</customKey>
<malwareKey>event.type</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>event.id</customKey>
<malwareKey>event.id</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>event.uuid</customKey>
<malwareKey>event.uuid</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>antivirus.primary.detected</customKey>
<malwareKey>antivirus.primary.detected</malwareKey>
```

```
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>antivirus.secondary.detected</customKey>
<malwareKey>antivirus.secondary.detected</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>antivirus.other.detected</customKey>
<malwareKey>antivirus.other.detected</malwareKey>
<isDisplay>true</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>country.dst.code</customKey>
<malwareKey>country.dst.code</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>city.dst</customKey>
<malwareKey>city.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>org.dst</customKey>
<malwareKey>org.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>payload</customKey>
<malwareKey>payload</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>packets</customKey>
<malwareKey>packets</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>country.dst</customKey>
<malwareKey>country.dst</malwareKey>
<isDisplay>false</isDisplay>
```

```
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>time</customKey>
<malwareKey>time</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>threat.source</customKey>
<malwareKey>threat.source</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>tcp.srcpport</customKey>
<malwareKey>tcp.srcpport</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>filetype</customKey>
<malwareKey>filetype</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>latdec.dst</customKey>
<malwareKey>latdec.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>eth.src</customKey>
<malwareKey>eth.src</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>agency.dst</customKey>
<malwareKey>agency.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ip.proto</customKey>
<malwareKey>ip.proto</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
```

```
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>tcp.flags</customKey>
<malwareKey>tcp.flags</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ip.src</customKey>
<malwareKey>ip.src</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>tcp.dstport</customKey>
<malwareKey>tcp.dstport</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>threat.category</customKey>
<malwareKey>threat.category</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>eth.dst</customKey>
<malwareKey>eth.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>lifetime</customKey>
<malwareKey>lifetime</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>latdec.src</customKey>
<malwareKey>latdec.src</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>did</customKey>
<malwareKey>did</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
```

```
<customKey>alert.id</customKey>
<malwareKey>alert.id</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>country.src</customKey>
<malwareKey>country.src</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>sessionid</customKey>
<malwareKey>sessionid</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>longdec.src</customKey>
<malwareKey>longdec.src</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>medium</customKey>
<malwareKey>medium</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>size</customKey>
<malwareKey>size</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ad.domain.dst</customKey>
<malwareKey>ad.computer.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ad.computer.dst</customKey>
<malwareKey>ad.computer.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ad.username.src</customKey>
```

```
<malwareKey>ad.username.src</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>rpackets</customKey>
<malwareKey>rpackets</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>action</customKey>
<malwareKey>action</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ad.domain.src</customKey>
<malwareKey>ad.domain.src</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>eth.src.vendor</customKey>
<malwareKey>eth.src.vendor</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>rpayload</customKey>
<malwareKey>rpayload</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ad.username.dst</customKey>
<malwareKey>ad.username.dst</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>content</customKey>
<malwareKey>content</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>extension</customKey>
<malwareKey>extension</malwareKey>
```

```
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>eth.dst.vendor</customKey>
<malwareKey>eth.dst.vendor</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>rid</customKey>
<malwareKey>rid</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>directory</customKey>
<malwareKey>directory</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>risk.suspicious</customKey>
<malwareKey>risk.suspicious</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>eth.type</customKey>
<malwareKey>eth.type</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>ip.dst</customKey>
<malwareKey>ip.dst</malwareKey>
<isDisplay>false</isDisplay>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>service</customKey>
<malwareKey>service</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>filename</customKey>
<malwareKey>filename</malwareKey>
```

```
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>streams</customKey>
<malwareKey>streams</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>risk.info</customKey>
<malwareKey>risk.info</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>dest.tld</customKey>
<malwareKey>dest.tld</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>alias.host</customKey>
<malwareKey>alias.host</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>tcp.srcport</customKey>
<malwareKey>tcp.srcport</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>udp.srcport</customKey>
<malwareKey>udp.srcport</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>udp.dstport</customKey>
<malwareKey>udp.dstport</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>domain.dst</customKey>
<malwareKey>domain.dst</malwareKey>
<isDisplay>false</isDisplay>
```

```
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>feed.name</customKey>
<malwareKey>feed.name</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>feed.description</customKey>
<malwareKey>feed.description</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>threat.description</customKey>
<malwareKey>threat.description</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>referer</customKey>
<malwareKey>referer</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>client</customKey>
<malwareKey>client</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>server</customKey>
<malwareKey>server</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>risk.warning</customKey>
<malwareKey>risk.warning</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
<com.netwitness.malware.core.cef.MalwareCefExtension>
<customKey>attachment</customKey>
<malwareKey>attachment</malwareKey>
<isDisplay>false</isDisplay>
</com.netwitness.malware.core.cef.MalwareCefExtension>
```

```
    <com.netwitness.malware.core.cef.MalwareCefExtension>
    <customKey>whois.registrar</customKey>
    <malwareKey>whois.registrar</malwareKey>
    <isDisplay>false</isDisplay>
    </com.netwitness.malware.core.cef.MalwareCefExtension>
    <com.netwitness.malware.core.cef.MalwareCefExtension>
    <customKey>whois.registrant</customKey>
    <malwareKey>whois.registrant</malwareKey>
    <isDisplay>false</isDisplay>
    </com.netwitness.malware.core.cef.MalwareCefExtension>
    <com.netwitness.malware.core.cef.MalwareCefExtension>
    <customKey>whois.date.creation</customKey>
    <malwareKey>whois.date.creation</malwareKey>
    <isDisplay>false</isDisplay>
    </com.netwitness.malware.core.cef.MalwareCefExtension>
    <com.netwitness.malware.core.cef.MalwareCefExtension>
    <customKey>whois.server</customKey>
    <malwareKey>whois.server</malwareKey>
    <isDisplay>false</isDisplay>
    </com.netwitness.malware.core.cef.MalwareCefExtension>
  </malwareExtensionList>
  <includesUnknownMeta>false</includesUnknownMeta>
  <displayNulls>false</displayNulls>
  <valueIfNull>n/a</valueIfNull>
</config>
```

Enable Custom YARA Content

This topic provides instructions for enabling custom YARA content on the NetWitness host on which the Malware Analysis service is installed. In addition to the built-in Indicators of Compromise, Malware Analysis supports indicators of compromise written in YARA. YARA is a rule language that allows malware researchers to identify and classify malware samples. Netwitness makes built-in YARA-based Indicators of Compromise (IOCs) available in Netwitness Live; these are automatically downloaded and activated on subscribed appliances.

Customers with advanced skills and knowledge can add detection capabilities to Netwitness Malware Analysis by authoring YARA rules and publishing them in RSA Live or placing YARA rules in a watched folder for the appliance to consume. This section provides instructions for the Administrator who configures appliances to enable the creation of custom YARA content.

Prerequisites

This is an advanced configuration task, which requires sufficient privilege and knowledge to set up a GNU Compiler Collection (GCC) and C++ Python development library to build YARA. In addition, you must be thoroughly familiar with the standard YARA documentation. The following components are required:

- The Perl-Compatible Regular Expression (PCRE) library: pcre-8.33.tar.bz2
- The yara 1.7 (rev:167) stand-alone YARA command line: yara-1.7.tar
- The YARA extension for Python: yara-python-1.7.tar.gz
- YARA rules documentation: YARA User's Manual 1.6.pdf

The components are available for download here: https://code.google.com/p/yara-project/downloads/list

Note: As of writing, YARA 2.0 is available but not supported for Malware Analysis 10.5.

Install Libraries and Applications Required to Build YARA on a CentOs-Based Appliance

As a prerequisite to building YARA on a host that is running CentOS, you must install `make`, the GNU Compiler Collection, and C++ Python Development Library on the appliance. To install the applications and libraries required to build YARA:

1. To ensure the standard YUM repo and no other repo files are in the /etc/yum.repos.d folder, enter the following command:
   `ls -al /etc/yum.repos.d`
   The results should be similar to the following:
   ```
   -rw-r-r-. 1 root root 1926 Jun 26  2012 CentOS-Base.repo
   -rw-r-r-. 1 root root  637 Jun 26  2012 CentOS-Debuginfo.repo
   -rw-r-r-. 1 root root  626 Jun 26  2012 CentOS-Media.repo
   -rw-r-r-. 1 root root 2593 Jun 26  2012 CentOS-Vault.repo
   ```
2. To install `make` on the appliance, enter the following commands:
   a. **`yum search make`**
      The following message is returned: `make.x86_64 : A GNU tool which simplifies the build process for user`
   b. **`yum install make.x86_64`**
3. To install and test GCC on the host, enter the following commands:
   a. **`yum search gcc`**
      The following messages are displayed:
      ```
      gcc-c+.x86_64 : C+ support for GCC
       gcc.x86_64 : Various compilers (C, C++, Objective-C, Java, ...)
      ```
   b. Enter the following commands:
      `yum install gcc.x86_64`
      `yum install gcc-c++.x86_64`

c. To test the gcc commands, enter the following commands:

```
gcc -v
cc -v
```

4. To install the C++ Python development library on the appliance, enter the following commands:

a. **yum search python dev**
The following message is returned:
`python-devel.x86_64 : The libraries and header files needed for Python development`

b. **yum install python-devel.x86_64**

Set Up YARA

To create a GCC and C++ Python development library in which you can build YARA on the NetWitness host that is running Malware Analysis:

1. Do one of the following:

a. If the host on which you are installing is running Mac OS, install xCode for Mac OS.

b. If the host on which you are installing is running CentOS, install make, GCC and C++ Python development library using the YUM command line.

2. To Install the PCRE library on the host, open a terminal window and enter the following commands:

```
tar -xvf pcre-8.33.tar.bz2
cd pcre-8.33
./configure
make
sudo make install
```

3. To install the stand-alone YARA command line, enter the following commands:

```
tar -xvf yara-1.7.tar
cd yara-1.7
./configure
make
sudo make install
```

4. To test the stand-alone YARA command line:

a. Enter the following command:

```
yara
```

b. If the command succeeds, continue with Step 7. If the command fails and returns the `yara: error while loading shared libraries: libpcre.so.1: cannot open shared object file: No such file or directory` error, enter the following command to check the `/etc/ld.so.conf` file or `LD_LIBRARY_PATH` environment variable.

```
ldconfig -v
```

5. To install the YARA extension for Python, enter the following commands:

```
tar -xvf yara-python-1.7.tar.gz
cd yara-python-1.7
```

```
python setup.py build
sudo python setup.py install
```

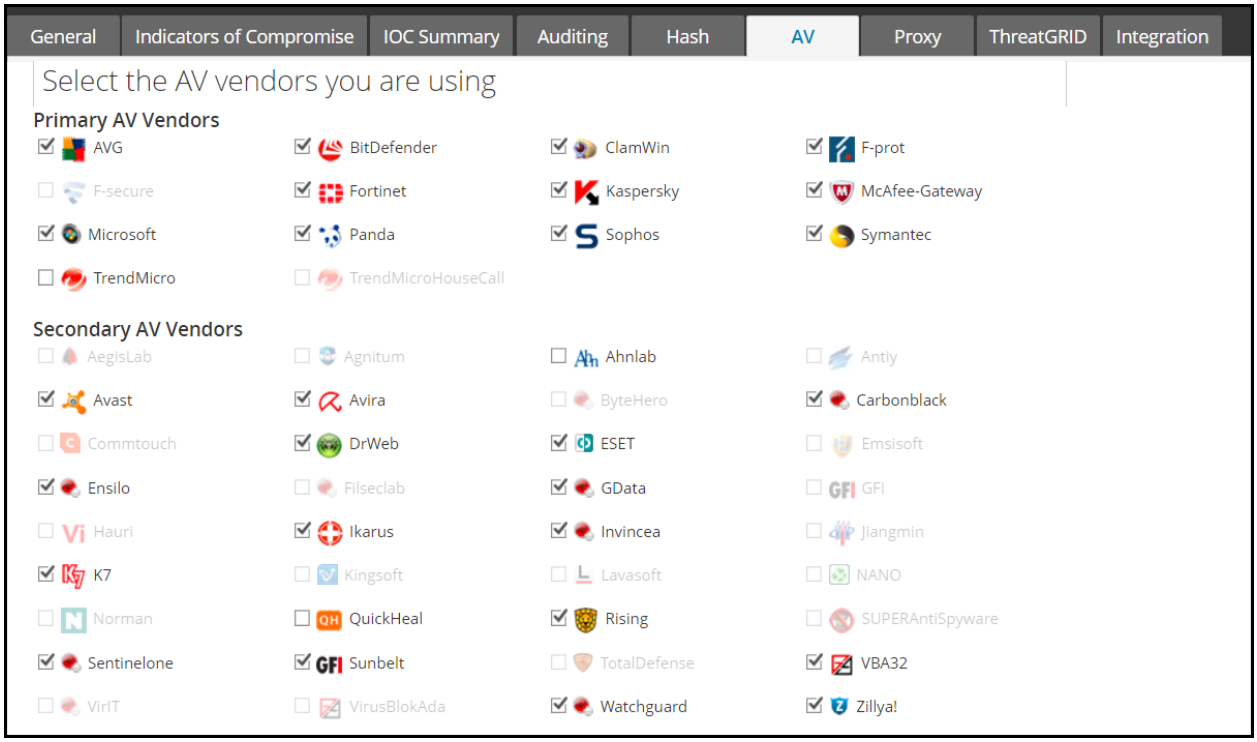
6. To test the YARA extension:
   a. Enter the following command: **python**
   b. At the Python prompt (>>>), enter the following commands:

```
import yara
exit()
```

When this configuration is complete, analysts can create custom YARA IOCs for consumption on a Malware Analysis host as described in "Implement Custom YARA Content" in the *Malware Analysis User Guide*

Supported Antivirus Vendors

This section provides detailed information about the NetWitness supported antivirus (AV) vendors.



The following are the new AV vendors supported in NetWitness version 11.4.0.1 or later.

- Panda
- Gdata
- Avast
- Carbonblack
- Dr.Web
- enSilo
- Invincea
- Rising
- SentinelOne
- Sunbelt
- VBA32
- Watchguard

The following AV vendors are not supported in NetWitness version 11.4.0.1 or later.

Note: AV vendors that are not supported by NetWitness Platform are disabled in the user interface.

- F-secure
- TrendMicroHouseCall
- AegisLab
- Agnitum
- Antiy
- ByteHero
- Commtouch
- Emsisoft
- Filseclab
- GFI
- Hauri
- Jiangmin
- Kingsoft
- Lavasoft
- NANO
- Norman
- nProtect
- SUPERAntiSpyware
- TotalDefense
- VirIT
- VirusBlokAda
- Zillya
- Zoner

Malware Analysis References

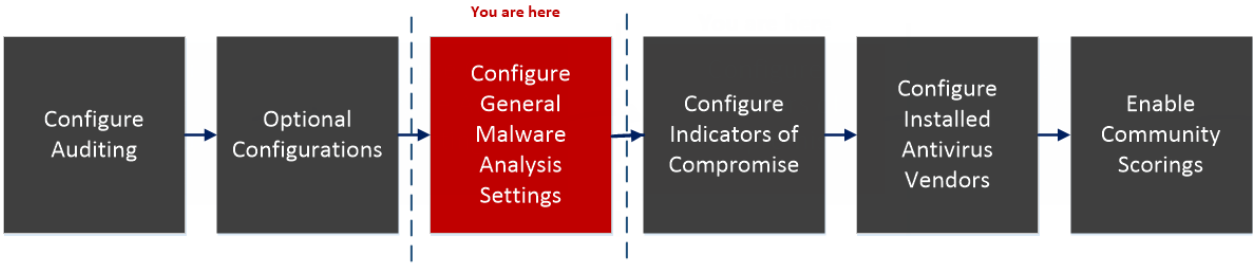
- Services Config View - General Tab
- Services Config View - Indicators of Compromise Tab
- Services Config View - IOC Summary Tab
- Services Config View - Auditing Tab
- Services Config View - Hash Tab
- Services Config View - AV Tab
- Service Config View - Proxy Tab
- Services Config View - ThreatGRID Tab
- Services Config View - Integration Tab

Services Config View - General Tab

This topic introduces the configuration settings in the Service Config view > General tab for Malware Analysis, which has parameters specific to the Malware Analysis service. In this tab, you can configure:

- The processing parameters for Core services that are capturing data.
- The repository for captured data.
- The static, community, and sandbox scoring categories used to analyze the data.

Workflow



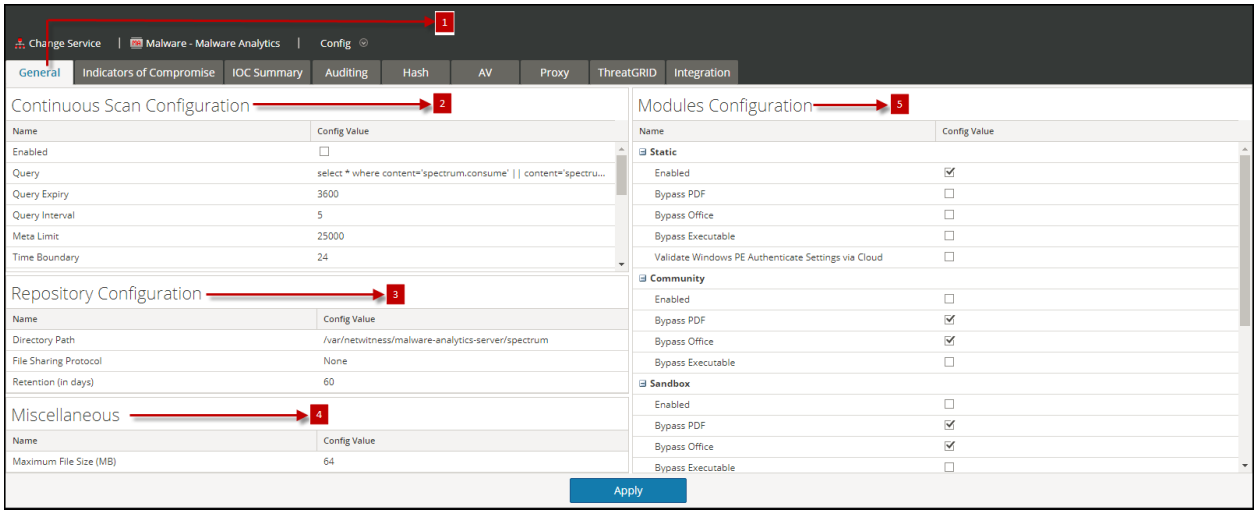
What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings*	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

*You can perform this task in the current view

Quick Look

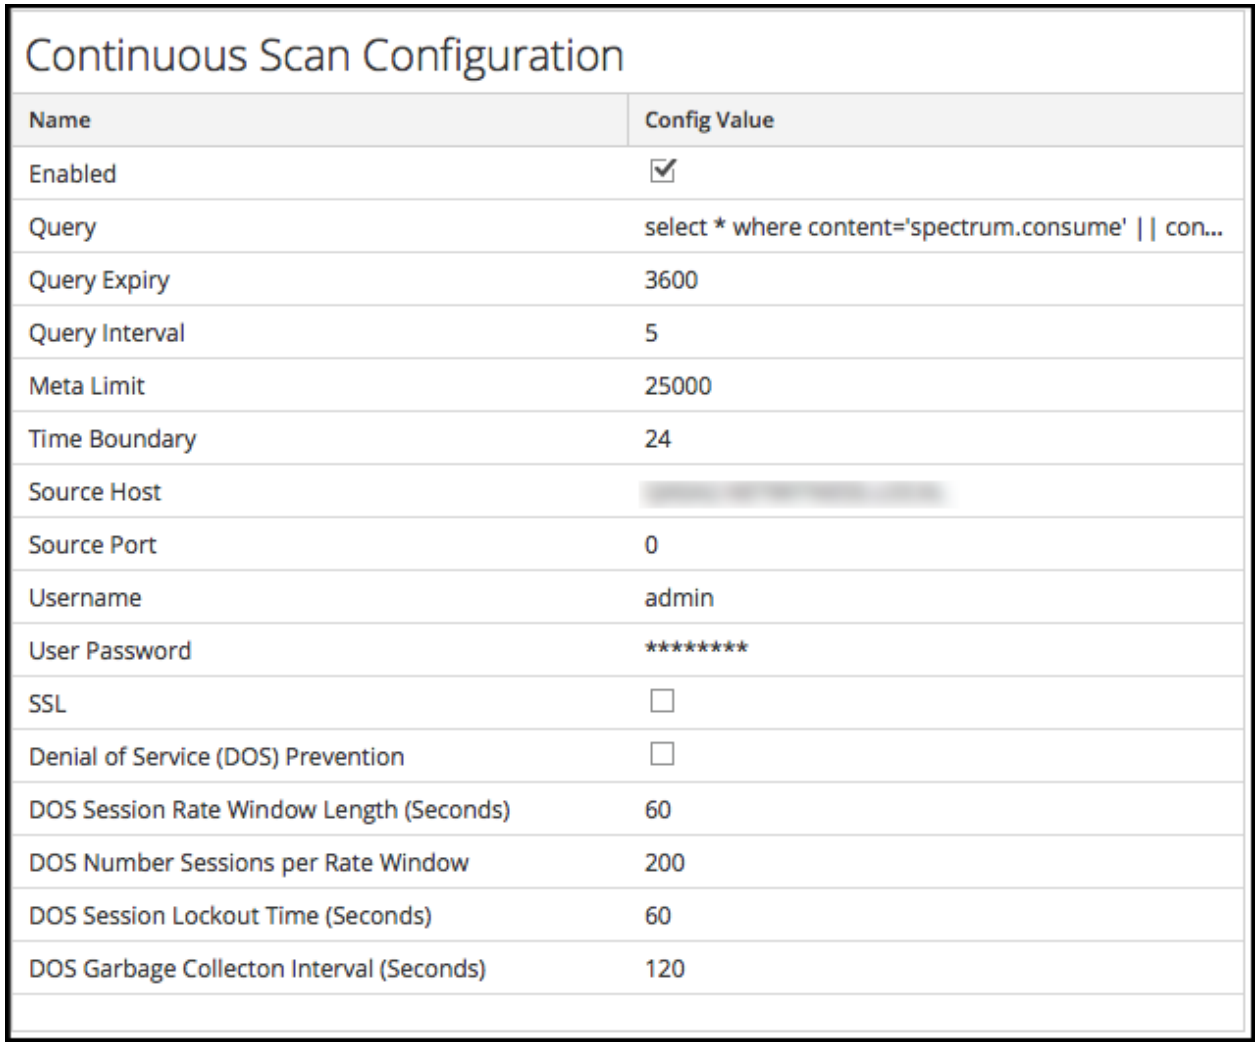
This is an example of the General tab.



1 Displays the General Tab.
2 Allows you to Configure Continuous Scan.
3 Allows you to Configure Repository.
4 Displays Miscellaneous Settings.
5 Allows you to Configure Modules.

This tab has four sections: Continuous Scan Configuration, Repository Configuration, Miscellaneous, and Modules Configuration.

Continuous Scan Configuration Section



Continuous Scan Configuration

Name	Config Value
Enabled	☑
Query	select * where content='spectrum.consume' || con...
Query Expiry	3600
Query Interval	5
Meta Limit	25000
Time Boundary	24
Source Host	[redacted]
Source Port	0
Username	admin
User Password	********
SSL	☐
Denial of Service (DOS) Prevention	☐
DOS Session Rate Window Length (Seconds)	60
DOS Number Sessions per Rate Window	200
DOS Session Lockout Time (Seconds)	60
DOS Garbage Collecton Interval (Seconds)	120

This table describes the features of the Continuous Scan Configuration section.

Parameter	Description
Enabled	Completely disable or enable continuous polling of the Core service. By default this is not selected (**disabled**).
Query	While the Decoder is analyzing network traffic, it creates a meta field called content with a value of **spectrum.consume** in sessions that are likely to contain malware. By default, Malware Analysis only performs analysis on events that have this particular meta value. By changing this query, Malware Analysis can be configured to analyze different types of events. Making this query too broad may force Malware Analysis to analyze too many events, causing it to fall behind or perform poorly. The default query is **select * where content='spectrum.consume'**

Parameter	Description
Query Expiry	When Malware Analysis queries the Core service for meta, it gets a result back within a few seconds. If there is a problem, such as a network connectivity issue, Malware Analysis abandons the query after this configured amount of time. The default value is **3600 seconds**.
Query Interval	How often, in minutes, to query for new session meta and files.
Meta Limit	Each time Malware Analysis queries the Core service, it pulls an amount of meta, up to this meta limit. Using this setting, in conjunction with the query interval, you can tune the performance of Malware Analysis in the Core infrastructure. The default value is **25000**.
Time Boundary	Malware Analysis analyzes sessions that occurred after the Time Boundary. This setting is most important when installing a new Malware Analysis appliance, because it determines how far back in time to begin analysis. Setting the boundary too many hours in the past may cause Malware Analysis to analyze too many past events, causing a large delay before you see any traffic happening in real time. The default value is **24 hours**.
Source Host	Hostname of the Malware Analysis appliance. This is the IP address, or the hostname, of the service that Malware Analysis queries to retrieve its data for analysis. Do not use localhost as the source host. Depending on the model of the appliance and the configuration of the NetWitness infrastructure, this source host can vary. **Note:** When you change the host name or the host IP address, ensure that you re add the Source Host in the Malware Service config page, and restart the service to take the source host field changes into effect.
Source Port	Malware Analysis communicates with the NetWitness infrastructure using the REST service listening on this port. This port number is specific to the type of the Core service that is being used as the Source host. This corresponds to the outbound connections for your Core service.
Username	Username. The default value is **admin**. Malware Analysis must authenticate to the Source host each time it queries for data. In most cases, the account used by Malware Analysis is the same account used to access the Core service through NetWitness. However, it is recommended to create a new account on the Core service dedicated to Malware Analysis.
User Password	User password. The default value is **netwitness**.
SSL	Use SSL when communicating with Core. If Malware Analysis is using an SSL connection to communicate with a Core service, check this option. The default value is unchecked.

Parameter	Description
Denial of Service (DOS) Prevention	The Denial of Service Prevention feature provides safeguards against malware that intentionally generates high volumes of network connections between two endpoints containing Windows PE content. Generating a high volume of connections artificially inflates the amount of traffic that security services monitoring the network must consume and analyze resulting in a denial of service. This feature helps identify these sessions so that you can have the analysis processing disregard them. The default value is unchecked.
DOS Session Rate Window Length (Seconds)	Malware Analysis uses this parameter with the **DOS Number Sessions per Rate Window** and **DOS Session Lockout Time (Seconds)** parameters to identify a Denial of Service Attack and determine how long to disregard sessions from a single IP address. To identify a Denial of Service Attack, Malware Analysis monitors the number of sessions established by a single IP address during a specific time frame. The **DOS Session Rate Window Length (Seconds)** defines this time frame. If the number of sessions exceeds the **DOS Number Sessions per Rate Window** setting within the number of seconds defined in **DOS Session Rate Window Length**, Malware Analysis identifies the activity as a Denial of Service attempt. In this case, traffic from the IP address is disregarded for the length of time specified in **DOS Session Lockout Time (Seconds)**. The default value is: **60** seconds.
DOS Number Sessions per Rate Window	Malware Analysis uses this parameter with the **DOS Session Rate Window Length (Seconds)** and **DOS Session Lockout Time (Seconds)** parameters to identify a Denial of Service Attack and determine how long to disregard sessions from the IP address. To identify a Denial of Service Attack, Malware Analysis monitors the number of sessions established by a single IP source during a specific time frame. The **DOS Session Rate Window Length (Seconds)** defines this time frame. If the number of sessions exceeds the **DOS Number Sessions per Rate Window** setting within the number of seconds defined in **DOS Session Rate Window Length**, Malware Analysis identifies the activity as a Denial of Service attempt. In this case, traffic is disregarded for the length of time specified in **DOS Session Lockout Time (Seconds)**. The default value is: **200** sessions
DOS Session Lockout Time (Seconds)	Malware Analysis uses this parameter with the **DOS Session Rate Window Length (Seconds)** and **DOS Number Sessions per Rate Window** parameters to identify a Denial of Service Attack and determine how long to disregard such an attack. To identify a Denial of Service Attack, Malware Analysis monitors the number of sessions established by a single IP address during a specific time frame. The **DOS Session Rate Window Length (Seconds)** defines this time frame. If the number of sessions exceeds the **DOS Number Sessions per Rate Window** setting within the number of seconds defined in **DOS Session Rate Window Length**, Malware Analysis identifies the activity as a Denial of Service attempt. In this case, traffic is disregarded for the length of time specified in **DOS Session Lockout Time (Seconds)**. The default value is: **60** seconds

Parameter	Description
DOS Garbage Collection Interval (Seconds)	Performs garbage collection on the internal memory structure used to track Denial of Service attempts. If memory usage is abnormally high, you can decrease this setting to free unused memory more often. If CPU usage is abnormally high, you can increase this setting to eliminate processing overhead (at the expense of memory usage). The default value is: **120** seconds

Repository Configuration Section

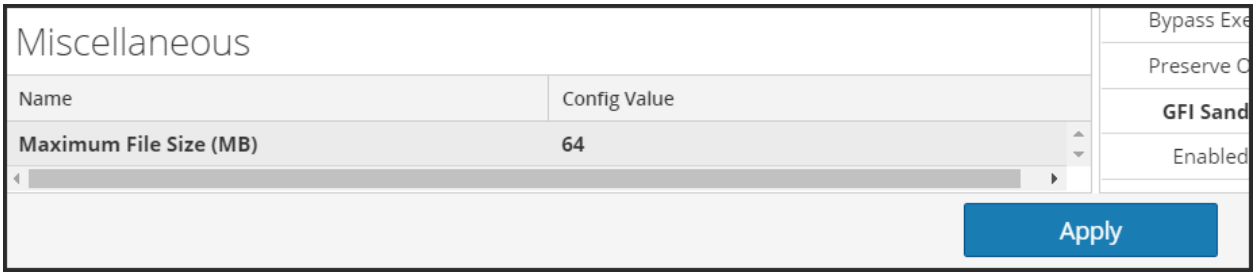


Malware Analysis stores all of the files that are analyzed for future use. These files can be downloaded through the user interface or accessed via one of the file sharing protocols.

This table describes the features of the Repository Configuration section.

Parameter	Description
Directory Path	All files are stored in the following directory on the Malware Analysis appliance: **/var/lib/netwitness/spectrum**
File Sharing Protocol	Possible values for the file sharing protocol are FTP, SAMBA, and None. You can enable FTP access and SAMBA file sharing to allow a user access to the stored files on the Malware Analysis from a remote location. No credentials are required to access these files. The port required for FTP access is TCP/21. The default file sharing protocol is **None**.
Retention (in days)	Malware Analysis maintains files stored in the repository for a specified number of days. You can set the number of days that files are retained before being deleted. The default value is **60** days.

Miscellaneous Configuration Section (10.3 SP2 and Later)



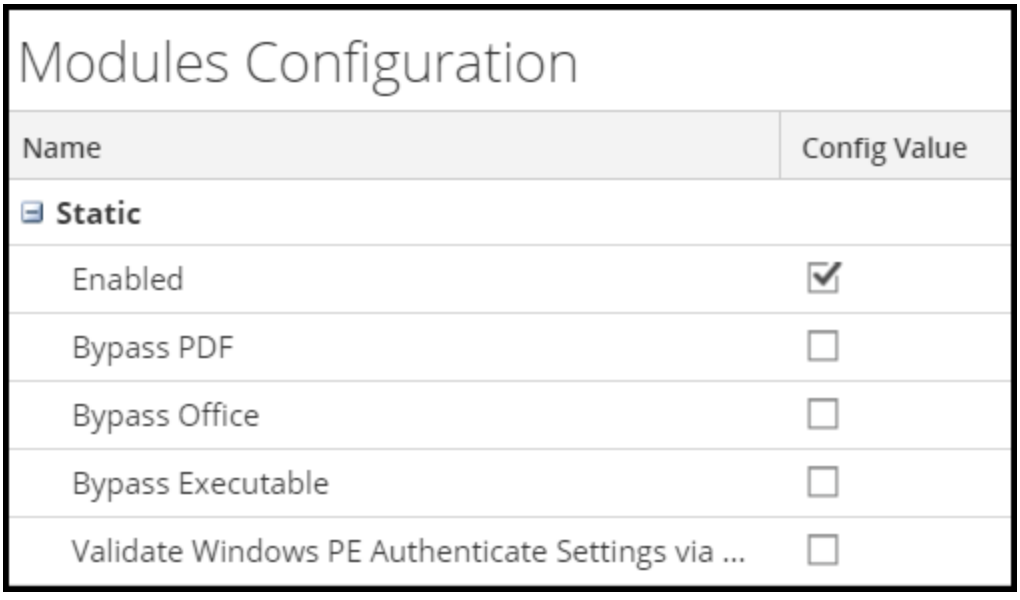
This table describes the features of the Miscellaneous Configuration section.

Parameter	Description
Maximum File Size	Limits the size of each file that you can scan for manually. This parameter applies to the feature described in "Upload Files for Malware Scanning" in the Investigation and Malware Analysis Configuration Guide. The default value is **64 MB**. If the file size limit is exceeded, prevents you from scanning the file.

Modules Configuration Section

The Modules Configuration section allows configuration of the static, community, and sandbox scoring categories.

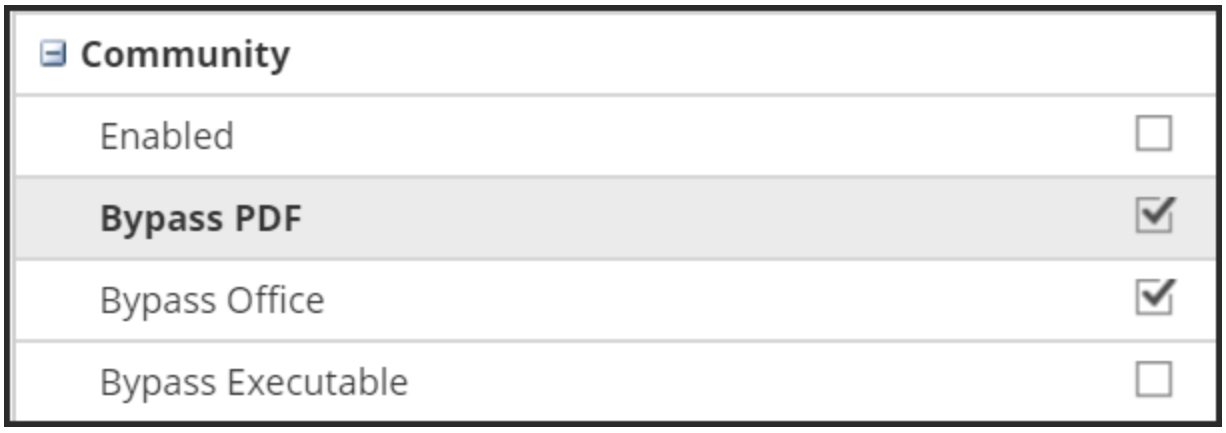
Static Analysis Configuration



The static module is the only scoring category that is enabled by default. This table describes the parameters for configuring static analysis.

Feature	Description
Enabled	Completely disable or enable static analysis. By default this is selected (**enabled**).
Bypass PDF	Disable analysis of PDF documents. By default this is not selected; all PDF files undergo static analysis.
Bypass Office	Disable analysis of Office documents. By default this is not selected; all MS Office files undergo static analysis.
Bypass Executable	Disable analysis of Windows PE documents. By default this is not selected; all Windows PE files undergo static analysis.
Validate Windows PE Authenticate Settings via Cloud	Specify whether or not Windows PE files are sent to the RSA-Netwitness Cloud for Authenticode validation. The default value is selected. • When selected, any Windows PE file that is digitally signed is transmitted over the network (in its entirety) to the RSA-Netwitness Cloud for validation. If the intent is to prevent Windows PE files from leaving the customer network, you should disable this option. • When not selected, ALL static analysis is performed locally (skipping Authenticode validation). Regardless of this setting, PDF and M/S Office documents are not subject to Authenticode validation and are not transmitted over the network during static analysis.

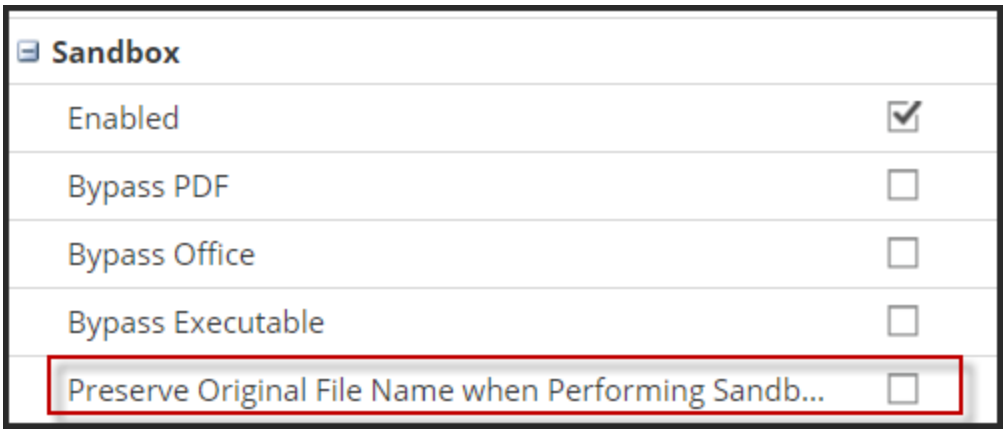
Community Analysis Configuration



By default, the community module is disabled and the options are selected to prevent PDFs and MS Office documents from being processed. The intent is to default the settings to the most restrictive choices so that no sensitive documents leave the network unless the user chooses. This table describes the parameters for configuring Community analysis.

Feature	Description
Enabled	Completely disable or enable community analysis. By default this is not selected (**disabled**). **Note:** Before you enable community, you must log in to live account. For more information about live account, see Live Services Management Guide.
Bypass PDF	Disable analysis of PDF documents. By default this is selected; PDF files are not processed.
Bypass Office	Disable analysis of Office documents. By default this is selected; Microsoft Office documents are not processed.
Bypass Executable	Disable analysis of Windows PE documents. By default this is selected; Windows PE documents are not processed

Sandbox Analysis Configuration



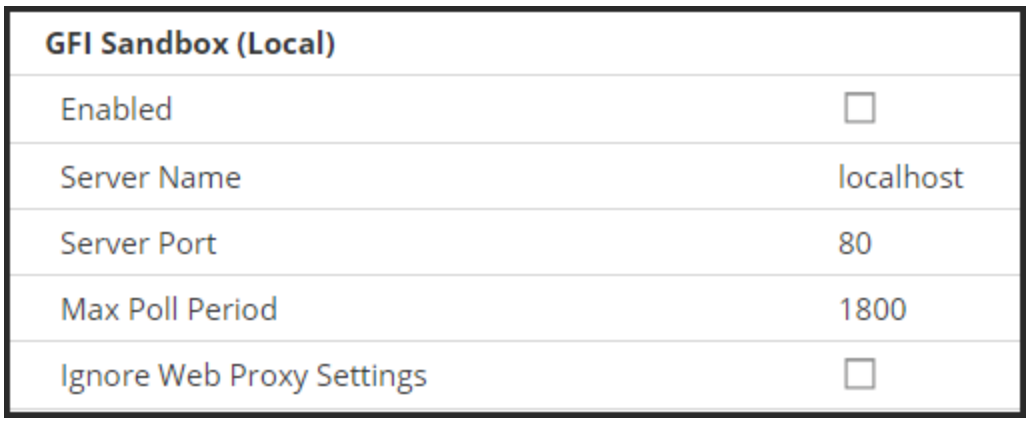
By default, the sandbox module is disabled and MS Office and PDF files are prevented from being processed. The intent is to set the most restrictive settings to force the user to specifically choose whether or not potentially sensitive information is sent outside of the network for processing. If the document type is not prevented from being processed, the file is sent to the destination sandbox server in its entirety (not limited to a hash of the file contents).

This table describes the parameters for configuring Sandbox analysis.

Feature	Description
Enabled	Completely disable or enable sandbox analysis. By default this is not selected **(disabled)**.
Bypass PDF	Disable analysis of PDF documents. By default this is selected; PDF files are not processed. When not selected, all PDF files are submitted in their entirety to the Sandbox for analysis.

Feature	Description
Bypass Office	Disable analysis of Office documents. By default this is selected; Microsoft Office documents are not processed. When not selected, all MS Office files are submitted in their entirety to the Sandbox for analysis.
Bypass Executable	Disable analysis of Windows PE documents. By default this is selected; Windows PE documents are not processed. When not selected, all Windows PE documents are submitted in their entirety to the Sandbox for analysis.
Preserve Original File Name when Performing Sandbox Analysis	In 10.3 SP2 and later, enable the ability to hash for filenames when they are sent to a local sandbox. By default this is not selected. **Note:** If you do not select this parameter, NetWitness hashes the files.

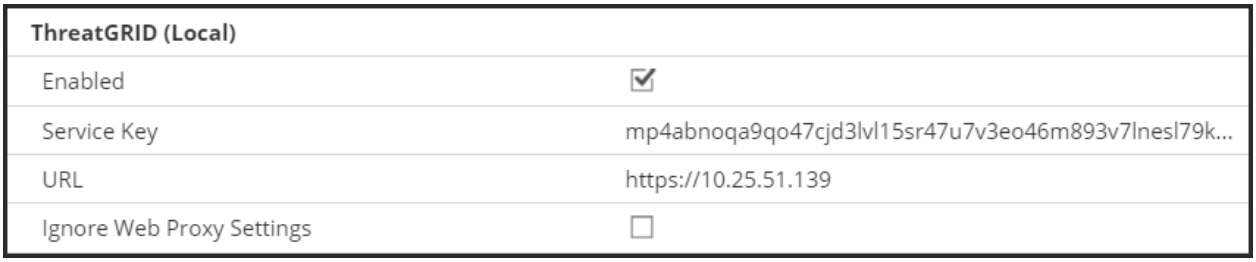
GFI Sandbox Settings



In the GFI Sandbox section, you can enable sandbox processing by GFI and configure the locally installed GFI sandbox. The table describes the parameters for configuring the GFI sandbox.

Feature	Description
Enabled	When enabled, sandbox processing is performed by a local copy of GFI. The default value is **disabled**. If you enable GFI, you need to configure the remaining parameters.
Server Name	The GFI Sandbox server name. No default value.
Server Port	The GFI Sandbox server port. Default value is **80**.
Max Poll Period	Determines how long to wait for a submitted sample to finish processing. Default value is **600 seconds**.
Ignore Web Proxy Settings	Tells Malware Analysis to bypass the web proxy, if a web proxy is configured, when making this connection. If no web proxy has been configured in Malware Analysis, the setting is ignored.

ThreatGRID Sandbox Settings



In the ThreatGRID Sandbox section, you can enable sandbox processing by ThreatGRID and choose whether to use the locally installed ThreatGRID or the ThreatGRID Cloud for sandbox analysis.

- If you have a local copy of ThreatGRID, configure sandbox processing to use the local copy.
- If no local instance of ThreatGRID has been purchased and installed, configure the ThreatGRID Cloud.

The table describes the parameters for configuring the ThreatGRID sandbox.

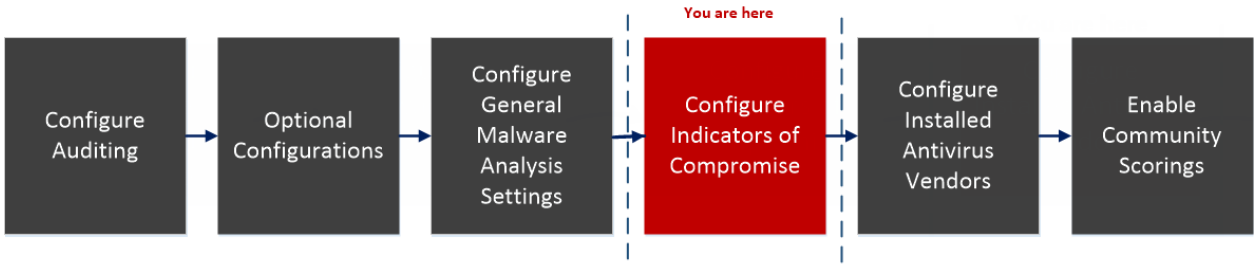
Note: Before enabling this service, you must configure a ThreatGRID-supplied Service Key. The service key allows ThreatGRID to recognize that samples submitted from this site are legitimate.

Feature	Description
Enabled	When enabled, sandbox processing is performed by ThreatGRID, either a local copy or the ThreatGRID Cloud. The default value is **disabled**.
Service Key	Before enabling the sandbox module, a ThreatGRID-supplied Service Key must be configured. The service key allows ThreatGRID to recognize that samples submitted from this site are legitimate.
URL	The URL for the ThreatGRID server to be used (if you are not using a locally installed ThreatGRID). The ThreatGRID Cloud is reachable via https://panacea.threatgrid.com
Ignore Web Proxy Settings	Tells Malware Analysis to bypass the web proxy, if a web proxy is configured, when making this connection. If no Web Proxy has been configured in Malware Analysis, the setting is ignored.

Services Config View - Indicators of Compromise Tab

This topic introduces the features and functions available in the Service Config view > Indicators of Compromise tab, which applies to the Malware Analysis service. This tab provides a way to configure the way each of the four scoring modules uses the available rules to score data.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise*	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

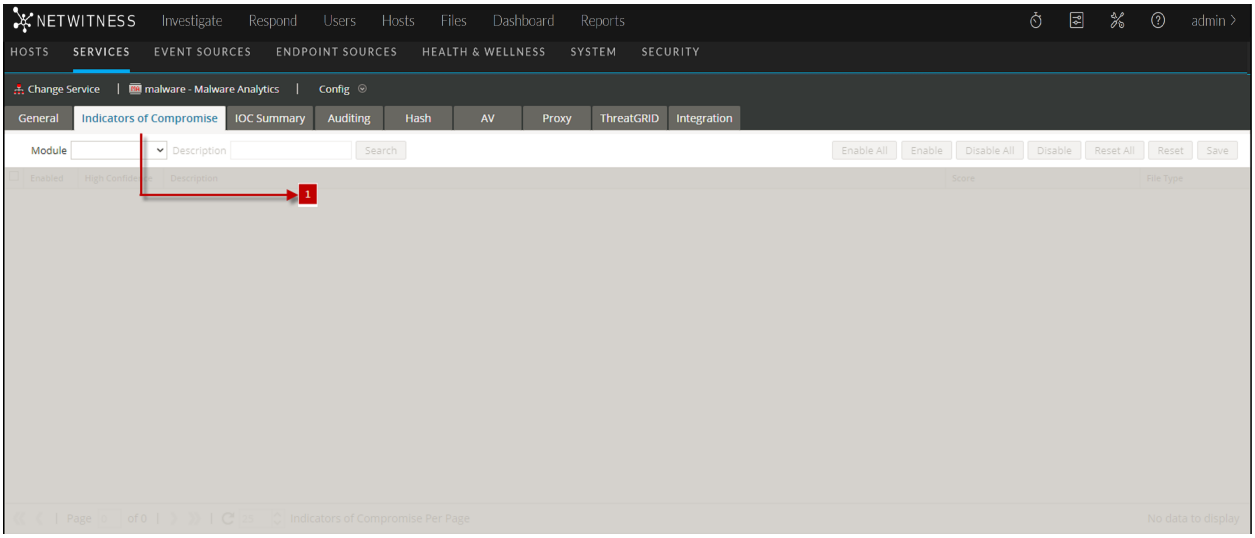
*You can perform this task in the current view

Related topic

Basic Setup

Quick Look

This is an example of the Indicators of Compromise tab.



1 Displays the Indicators of Compromise Tab.

Features

The Indicators of Compromise tab consists of a toolbar and page-able grid.

This table describes the features of the grid.

Feature	Description
Module selection list	Selects the scoring module for which you want to view the Indicators of Compromise: All, Network, Static, Community, Sandbox, or Yara.
Search field	Type text for which you are searching in the Description field.
Search option	Filters the grid to display only Descriptions that match the Description search term.
Enable All option	Click to enable all rules for the scoring module, as opposed to enabling all rules on the page using the checkbox.
Enable option	Click to enable selected rules.
Disable All option	Click to disable all rules for the scoring module, as opposed to disabling all rules on the page using the checkbox.
Disable option	Click to disable selected rules.

Feature	Description
Reset All option	Click to reset all rows on the page to their default values.
Reset option	Click to reset selected rows to their default values.
Save option	Click to save changes you made on this page. If you leave the page without saving, the changes are lost. The description of each row with unsaved changes has a red corner.

This table describes the features of the toolbar.

Column	Description
Selection checkbox	Checkboxes for selecting individual rows or all rows on the page.
Enabled checkbox	If the indicator of compromise is enabled, Malware Analysis uses the rule for scoring session data.
High Confidence checkbox	If checked, Malware Analysis treats the rule as one very likely to indicate the presence of malware, and an event that triggers that rule is marked in the results grid.
Description	Describes the Indicator of Compromise.
Score	Specifies the score that you want to factor in to the total score for any event that triggers the rule. The default score is displayed and you can raise or lower the score by dragging the slider or typing a number in the score box.
File Type	Displays the file types to which the rule applies. Possible values are **ALL**, **PDF**, **MS Office**, and **Windows PE**.

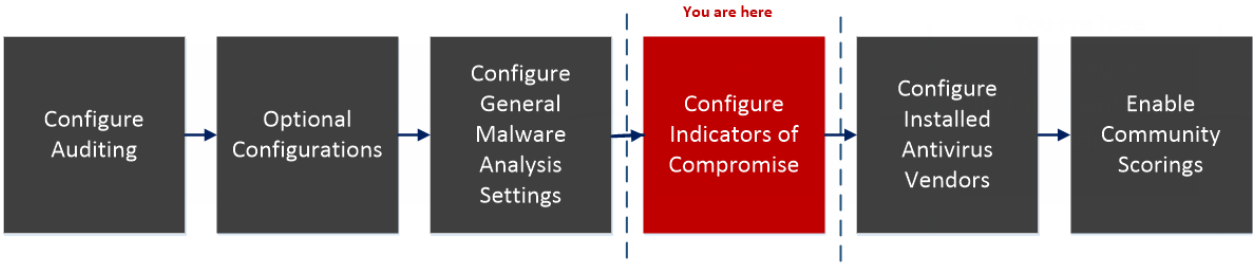
Services Config View - IOC Summary Tab

This topic introduces the features and functions available in the Service Config view > IOC Summary tab. This tab provides a way to view summary information for any IOC. A grid for each scoring module lists the configured IOCs along with statistics associated with that IOC of a specific range of time. The statistics include:

- The number of events for a network session or the number of files for a static, community, or sandbox event that were flagged with the IOC.
- The current score configured for the IOC in the Indicators of Compromise tab.
- The scores returned by each of the scoring modules.

When you select an event, you can show the Malware Events view or Malware Files view for the IOC. You can also open the selected IOC in the Indicators of Compromise tab to edit the Current Score.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise*	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings

Role	I Want to...	Show me how
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

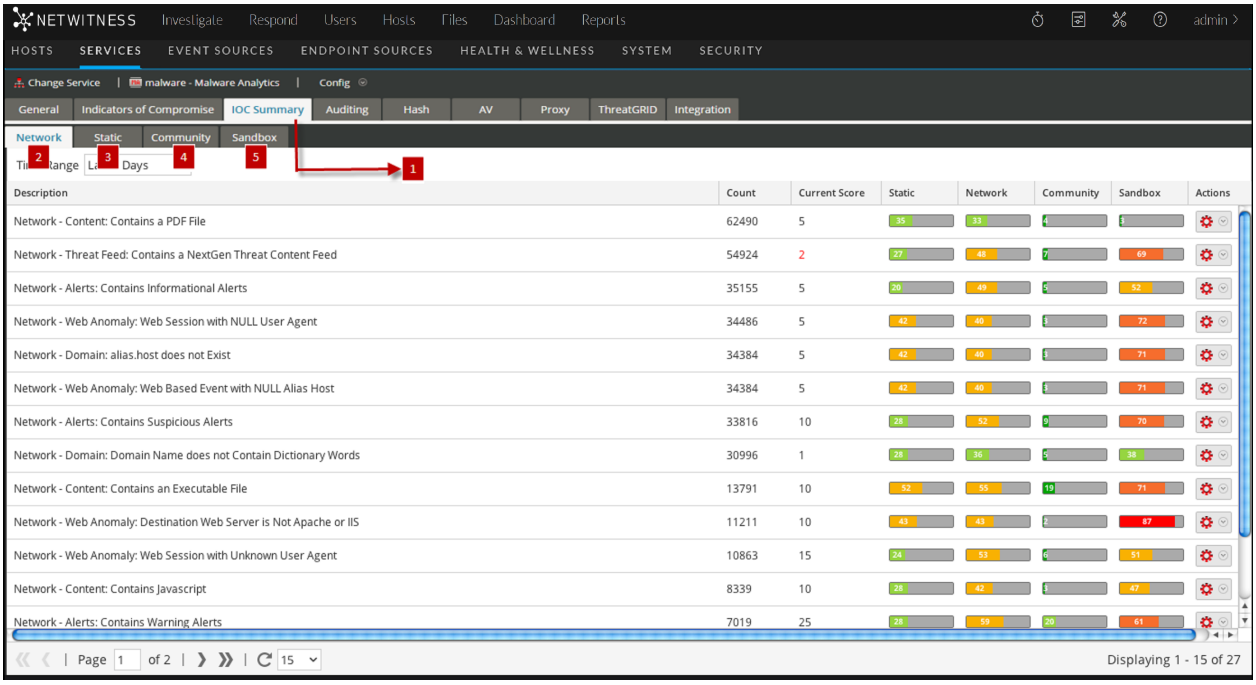
*You can perform this task in the current view

Related Topic

Basic Setup

Quick Look

This is an example of the IOC Summary tab for the Network scoring module.



1 Displays the IOC Summary Tab.
2 Displays the Network View.
3 Displays the Static View.
4 Displays the Community View.
5 Displays the Sandbox View.

Features

The IOC Summary consists of four tabs, one for each scoring module: Network, Static, Community, and Sandbox. Each tab has the same form and same information with a toolbar and page-able grid.

This table describes the features of each tab.

Feature	Description
Time Range	Selects the time range for the IOC Summary. Possible values are: Last 5 Minutes, Last 15 Minutes, Last 30 Minutes, Last Hour, Last 3 Hours, Last 6 Hours, Last 12 Hours, Last 24 Hours, Last 2 Days, Last 5 Days, Early Morning, Morning, Afternoon, Evening, All Day, Yesterday, This Week, Last Week, or Custom.
Description column	Lists the descriptions for the IOCs.
Count column	Lists the number of occurrences of the IOCs. In the Network tab, the count is the number of events in which the IOC was found. In the other tabs, the count is the number of files in which the IOC was found.
Current Score column	Lists the current score for the IOCs as configured in the Indicators of Compromise tab.
Static, Network, Community, and Sandbox columns	List the scores that each of the scoring modules gave the IOCs.
Actions drop-down	The Actions drop-down menu has two options: • Show Events/Files: opens the IOC in the Investigation Events view or Files view. This view can also be opened by double-clicking on the IOC. • Edit: opens the IOC in the Indicators of Compromise tab to edit the Current Score.

Services Config View - Auditing Tab

This topic introduces the features and functions of the Auditing tab in the Services Config view for Malware Analysis. The Auditing tab in the Services Config view for Malware Analysis provides a way to configure the auditing feature. Malware Analysis has an automated auditing system capable of sending alerts (syslog, snmp, audit log file entries) as Malware Analysis exceeds configured score value thresholds for each scoring module (Network, Static, Community, Sandbox). Malware Analysis can automatically feed any external system capable of ingesting the supported audit formats. One alert is generated for each file in an analyzed session that meets or exceeds the configure threshold.

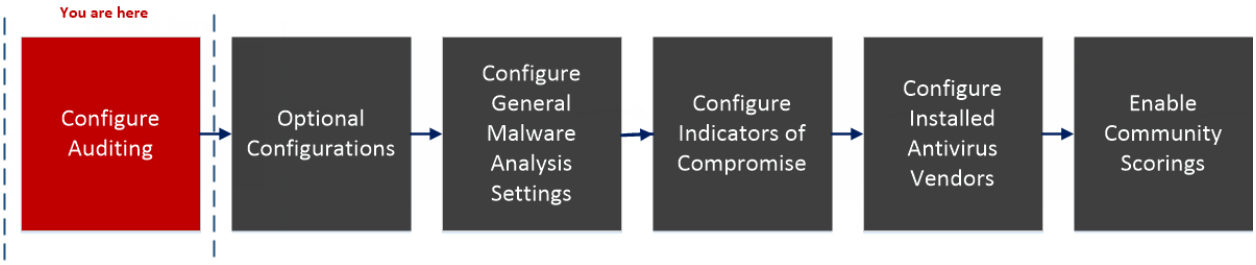
The audit log is a log file maintained on the Malware Analysis appliance for every significant event or action. Audit logs are rolled out and archived over time as they become large so an audit history is maintained. The size of these audit logs and their number are both configurable.

Some examples of events that are logged are:

- User login successes and failures
- Changes to system configuration settings
- Server restart
- Server version upgrade and install
- Suspicious events that exceed the Audit Thresholds

Malware Analysis can send audit events as an SNMP trap to a configured SNMP trap host, and consolidate logs in syslog format. Refer to the following task topic for detailed procedures: (Optional) Configure Auditing on Malware Analysis Host.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host*	(Optional) Configure Auditing on Malware Analysis Host

Role	I Want to...	Show me how
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

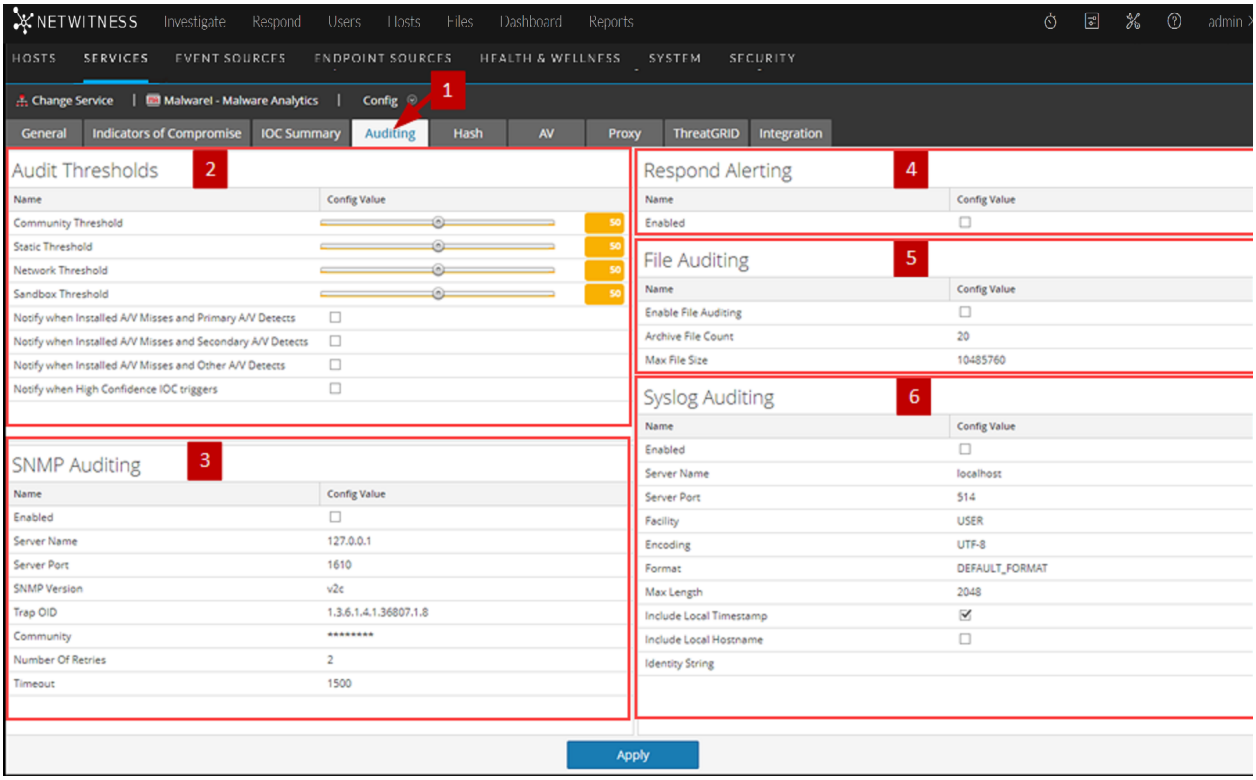
*You can perform this task in the current view

Related Topics

Basic Setup

Quick Look

This is an example of the Auditing tab.



1 Displays the Auditing Tab.
2 Displays the Audit Thresholds section.
3 Displays the SNMP Auditing section.

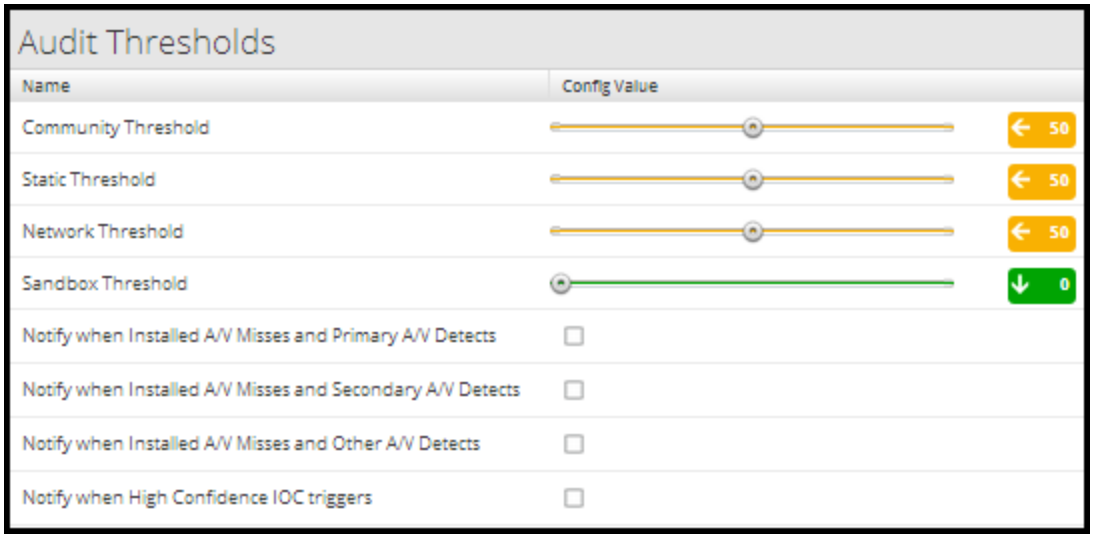
4 Displays the Respond Alerting section.
5 Displays the File Auditing section.
6 Displays the Syslog Auditing section.

Features

The Auditing tab includes five sections and an Apply button used to save changes made in this tab and put them into effect.

- Auditing Thresholds
- SNMP Auditing
- Respond Alerting
- File Auditing
- Syslog Auditing

Audit Thresholds



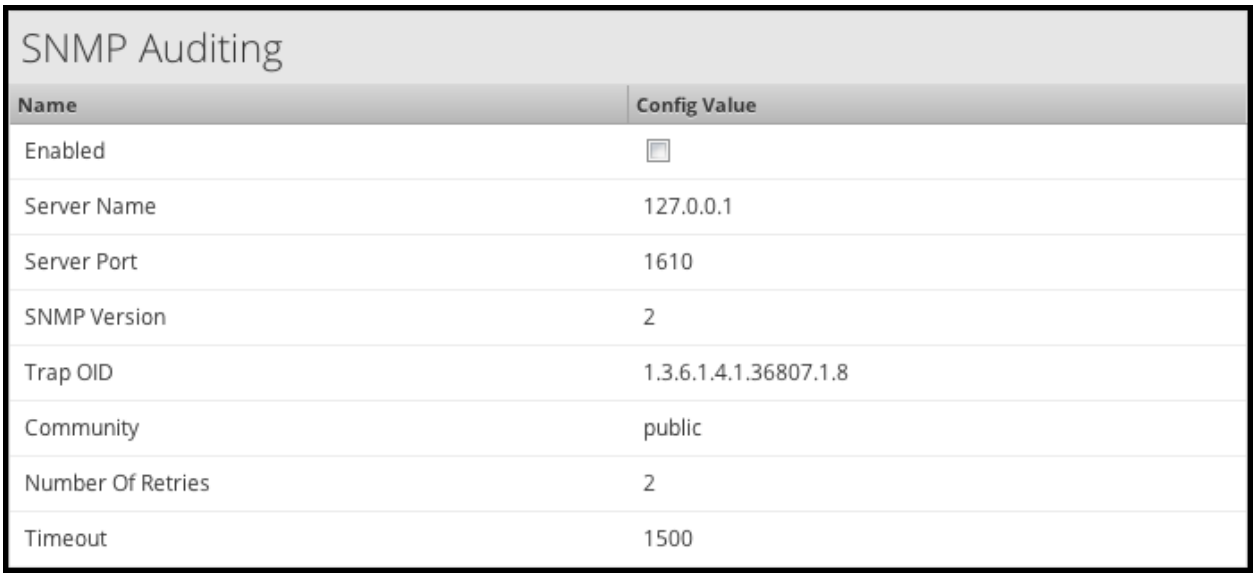
This table describes the features in the Audit Thresholds section.

Name	Config Value
Community, **Static**, **Network**, and **Sandbox Thresholds**	Malware Analysis scoring module thresholds for recording event information in a log file. Malware Analysis records the event information in a log file if the event scored high enough to satisfy all of the auditing thresholds. Each scoring category that completed analysis (for example, not all sessions invoke sandbox analysis) is compared against the configured audit threshold for that category. Any one of the category must exceed the threshold in order for an audit event to be triggered. An integer between 0 and 100 is a valid value. Setting these thresholds too low may cause a very large volume of audit events and notifications.

Name	Config Value
Notify when Installed A/V Misses and Primary A/V Detects	Records a message in a log file when installed antivirus software misses a virus and the primary antivirus software detects that virus. The recorded message is sent through all enabled auditing methods: SNMP, File, and Syslog. The default value is unchecked.
Notify when Installed A/V Misses and Secondary A/V Detects	Records a message in a log file when installed antivirus software misses a virus and the secondary antivirus software detects that virus. The recorded message is sent through all enabled auditing methods: SNMP, File, and Syslog. The default value is unchecked.
Notify when Installed A/V Misses and Other A/V Detects	Records a message in a log file when installed antivirus software misses a virus and the other antivirus software detects that virus. The recorded message is sent through all enabled auditing methods: SNMP, File, and Syslog. The default value is unchecked.
Notify when High Confidence IOC triggers	Records a message in a log file when a high confidence IOC (Indicators of Compromise) triggers. The recorded message is sent through all enabled auditing methods: SNMP, File, and Syslog. The default value is unchecked.

SNMP Auditing

The Simple Network Management Protocol (SNMP) is an Internet-standard protocol for managing services on IP networks. When SNMP auditing is enabled, Malware Analysis can send an audit event as an SNMP trap to a configured SNMP trap host.

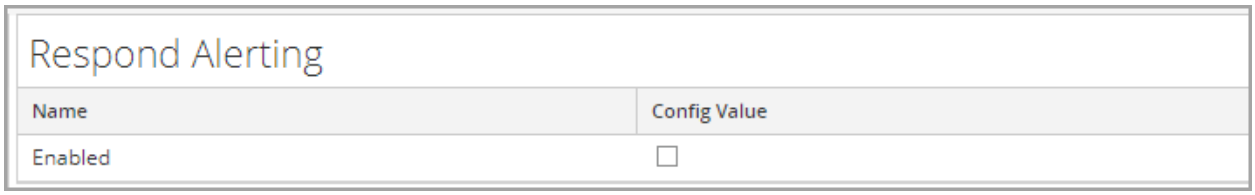


This table describes the features in the SNMP Auditing section.

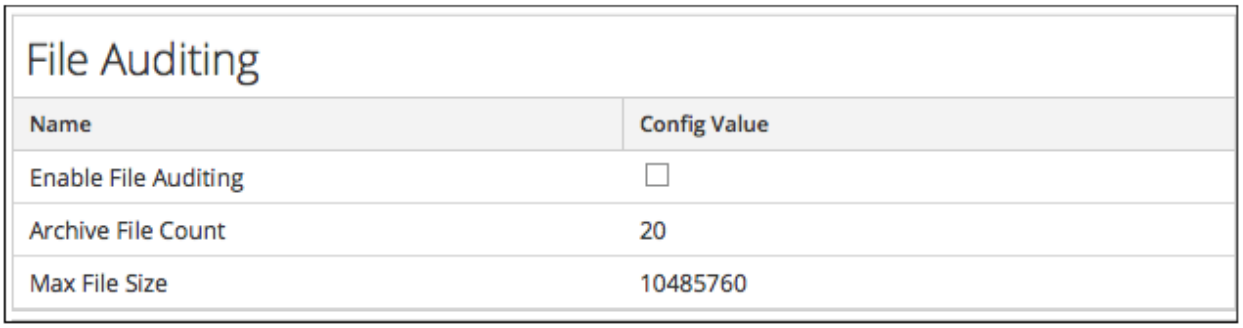
Name	Config Value
Enabled	Click to enable or disable SNMP auditing.
Server Name	The host where the target SNMP server is running.
Server Port	The port used where the SNMP trap receiver is listening.
SNMP Version	The version of the SNMP protocol to use when sending traps.
Trap OID	The object ID to use to identify the type of trap to send.
Community	The SNMP group to which Malware Analysis belongs.
Number Of Retries	The number of retries for sending a trap.
Timeout	The timeout period to wait for acknowledgment.

Respond Alerting

The Respond Alerting section enables NetWitness Respond to receive alerts from Malware Analysis. Select Enabled to forward alerts to the Respond view.



File Auditing

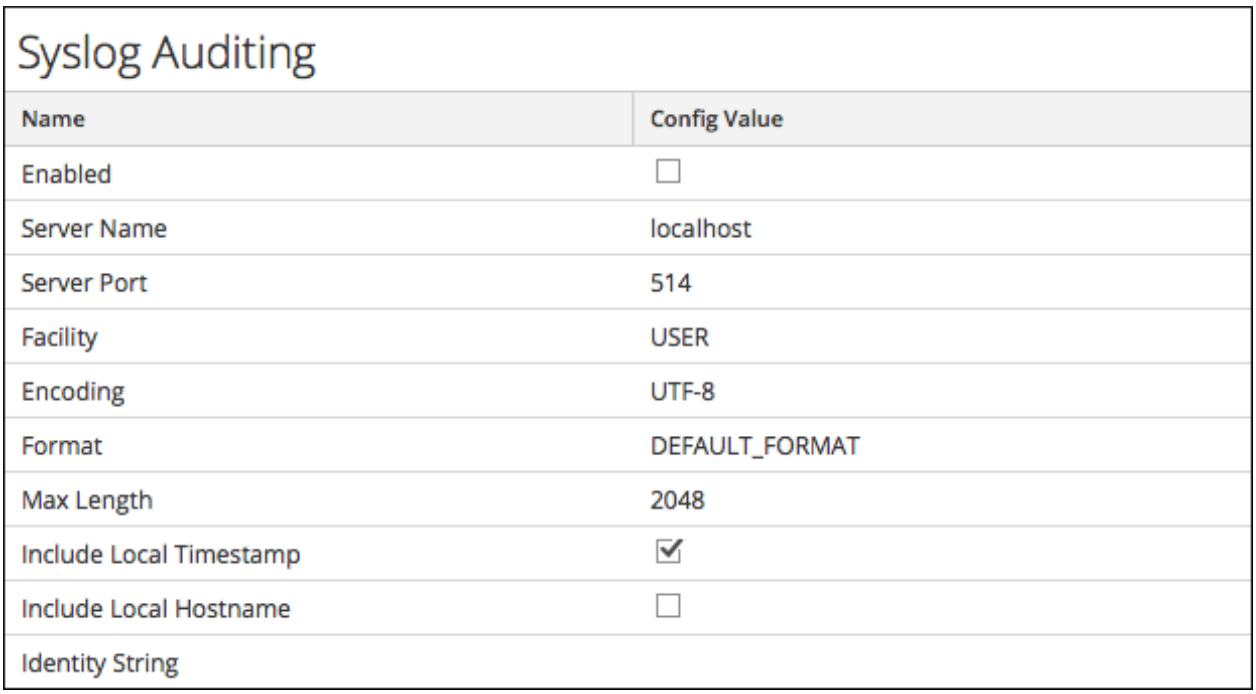


This table describes the features in the File Auditing section. Avoid setting the max file size and archive file count too high because it may have an adverse effect on the available disk space on the Malware Analysis appliance.

Name	Config Value
Enable File Auditing	Click to enable or disable file auditing.

Name	Config Value
Archive File Count	Malware Analysis keeps only as many log files as defined by this setting. When the maximum number is reached, the oldest log files are deleted and cannot be recovered. The default value is 20. Valid value: Integer between 1 and 50, inclusive.
Max File Size	The maximum file size for a single auditing log before it is archived. The default value is 10485760 bytes.

Syslog Auditing



This table describes the features in the Audit Thresholds section.

Feature	Description
Enabled	Click to enable or disable syslog auditing.
Server Name	This is the host where the target syslog process is running.
Server Port	This is the port where the target syslog process is listening.
Facility	This is the designated syslog facility to use for all outgoing messages. Possible values are KERN, USER, MAIL, DAEMON, AUTH, SYSLOG, LPR, NEWS, UUCP, CRON, AUTHPRIV, and LOCAL1 through LOCAL7.
Encoding	This is the encoding to use for text in syslog messages; for example, UTF-8.
Format	This is the desired message format. Possible values are: Default, PCI DSS, or SEC.

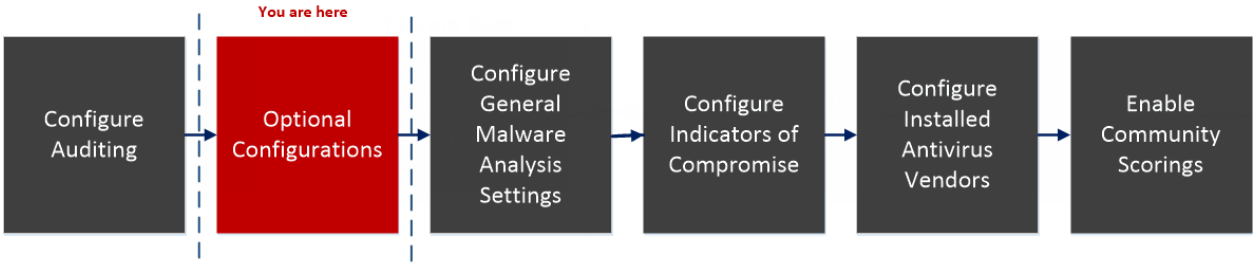
Feature	Description
Max Length	This is the maximum length in bytes that any syslog message can be. Default is 1024. Messages that exceed the maximum length are truncated.
Include Local Timestamp	Check this box to include the local timestamp in messages.
Include Local Hostname	Check this box to include the local hostname.
Identity String	This is an identity string to be prepended to each syslog alert. If the string is blank, no identity string is prepended to the outgoing syslog alerts. You can use this to identify the source of the alert. Users conventionally set it to the name of the program that will submit the messages to a syslog auditing.

Services Config View - Hash Tab

This topic introduces the features and functions available in the Service Config view > Hash tab for Malware Analysis.

In this tab, you can manage hash filtering in Malware Analysis. The hash grid is initially empty; the grid lists filters that have been added to Malware Analysis. In this view, you can add a hash filter, delete a hash filter, mark a hash filter as trusted or untrusted, and save changes.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter*	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

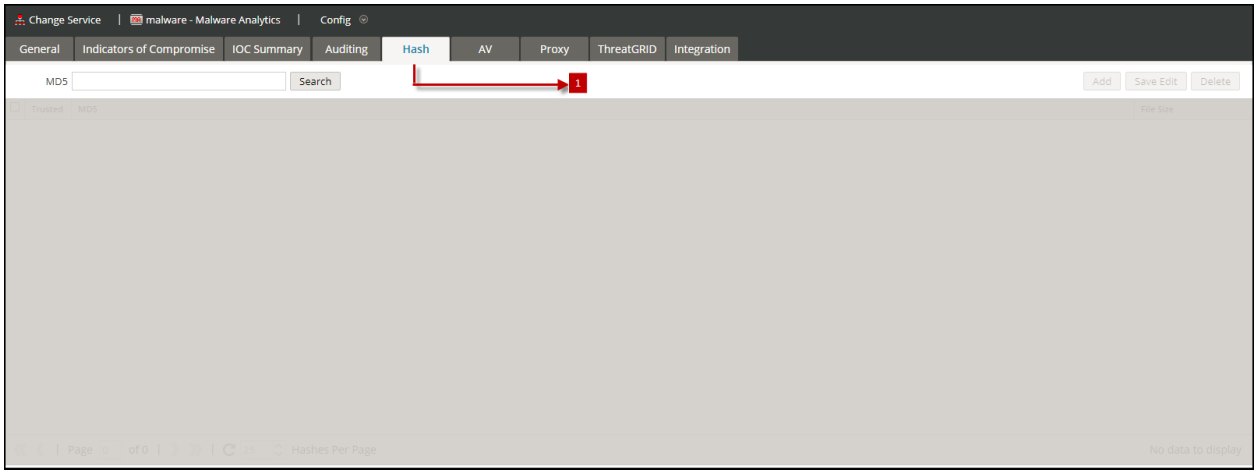
*You can perform this task in the current view

Related Topic

Basic Setup

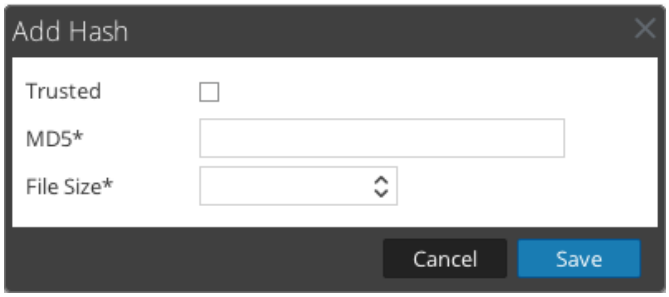
Quick Look

This is an example of the Hash tab.



1 Displays the Hash Tab.

This is an example of the Add Hash dialog.



Features

The **Hash** tab consists of a toolbar and a pageable hash grid.

This table describes the Hash tab toolbar.

Feature	Description
MD5 Search	Enter an MD5 hash for which you want to search the results in the grid. The search function is case-insensitive.
Add	Displays the Add Hash dialog in which you can add a new hash to the hash grid, specify whether the hash is trusted or not, and provide the hash file size.
Save Edit	Saves any additions or edits to hashes in the grid.
Delete	Deletes selected hashes from the grid.

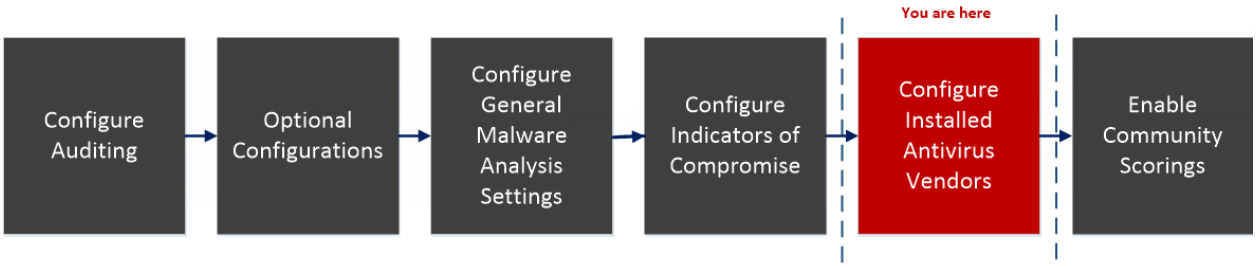
This table describes the Hash grid columns.

Feature	Description
Select Checkbox	Click to select a row. Click in the column header to select a header.
Trusted	Marks a hash as trusted or untrusted.
MD5	Identifies the MD5 hash.
File Size	Identifies the hash file size in kilobytes.

Services Config View - AV Tab

This topic introduces the features and functions of the AV tab in the Service Config view for a Malware Analysis service. The AV tab provides a way to identify the anti-virus software vendors whose software is in use on your network. NetWitness can include the results from these vendors in the detailed results view of an event that has been analyzed using Malware Analysis.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor*	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

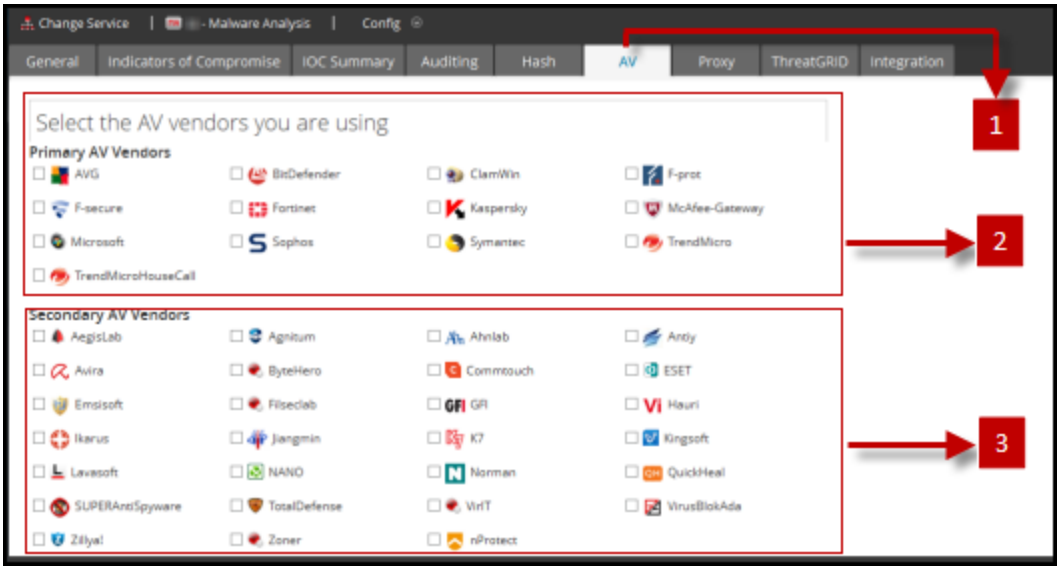
*You can perform this task in the current view

Related Topic

Configure Installed Antivirus Vendors

Quick Look

This is an example of the AV tab.



1 Displays the Av Tab.
2 Allows you to select the AV vendor that you are using.
3 Displays the Secondary AV vendors.

Features

The AV tab lists anti-virus vendors whose software may be installed in your network. There are two categories for vendors: Primary, which are the most trusted, and Secondary, which are less known. Each vendor name has a checkbox and an icon. Checking a vendor name indicates that you have installed the selected AV software from that vendor in your environment.

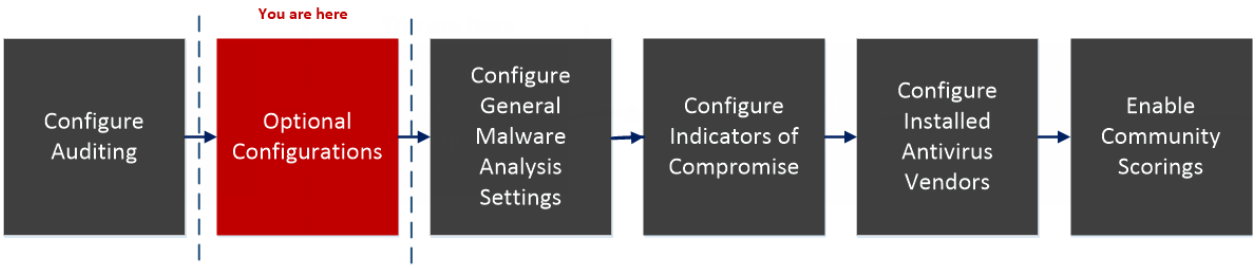
This table describes the options in the AV tab.

Feature	Description
Vendor Checkbox	Choose one or more Anti Virus vendors from the supplied list to indicate which products have been installed in the local organization.
Apply	Saves changes made in the AV tab.
Reset	Resets the AV list to the default state, which has no vendors selected.

Service Config View - Proxy Tab

This topic introduces the parameters configured in the Proxy tab in the Service Config view for a Malware Analysis service. This tab configures Malware Analysis communication via web proxy with the RSA Cloud for community analysis and with the sandbox service for sandbox analysis to preserve anonymity. If you are using a local sandbox service, communications via web proxy are unnecessary and may slow performance. When configuring the sandbox module in the **General** tab, you can choose to bypass the configured web proxy.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings*	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

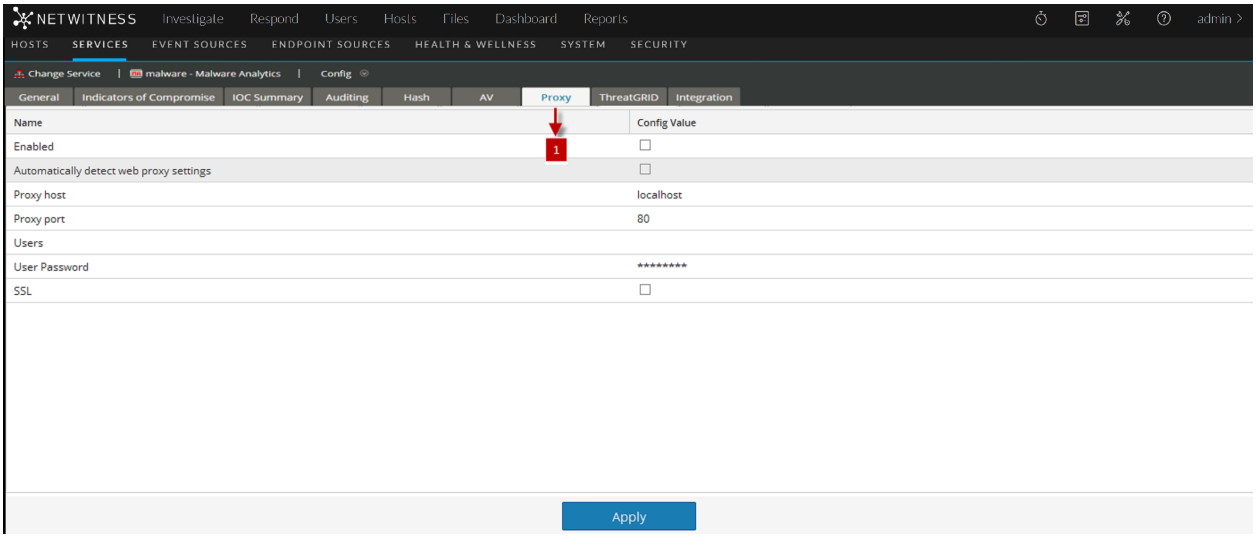
*You can perform this task in the current view

Related topic

Basic Setup

Quick Look

This is an example of the Proxy tab.



1 Displays the Proxy Tab.

Features

This table describes the features in the Proxy tab.

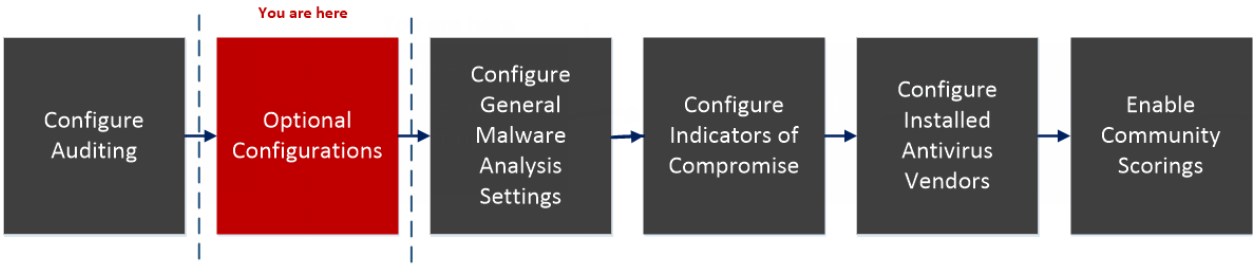
Feature	Description
Enabled	Select the checkbox to enable communication via web proxy with the RSA Cloud for community analysis and with the sandbox service for sandbox analysis to preserve anonymity.
Automatically detect web proxy settings	Select the checkbox to use settings configured in the System settings.
Proxy host	Enter the hostname for the proxy host.
Proxy port	Enter the port used for communication on the proxy host
Users	Enter the username used to log on to the proxy host.
User Password	Enter the user password used to log on to the proxy host.
SSL	(Optional) Select the checkbox to enable communication using SSL.
Apply button	Click the **Apply** button to submit chosen settings.

Services Config View - ThreatGRID Tab

This topic introduces the parameters required to obtain a trial ThreatGRID API key in the Malware Analysis **ThreatGRID** tab, which provides a method of obtaining a trial ThreatGRID API key for use in the ThreatGRID Cloud sandbox. Before enabling ThreatGRID as the sandbox service in the Sandbox module, a ThreatGRID-supplied Service Key must be configured so that ThreatGRID can recognize that samples submitted from this site are legitimate.

If you do not have a ThreatGRID-supplied Service Key, you can obtain a key using this tab. The key is provided on a trial basis.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key*	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis	Enable Community Analysis

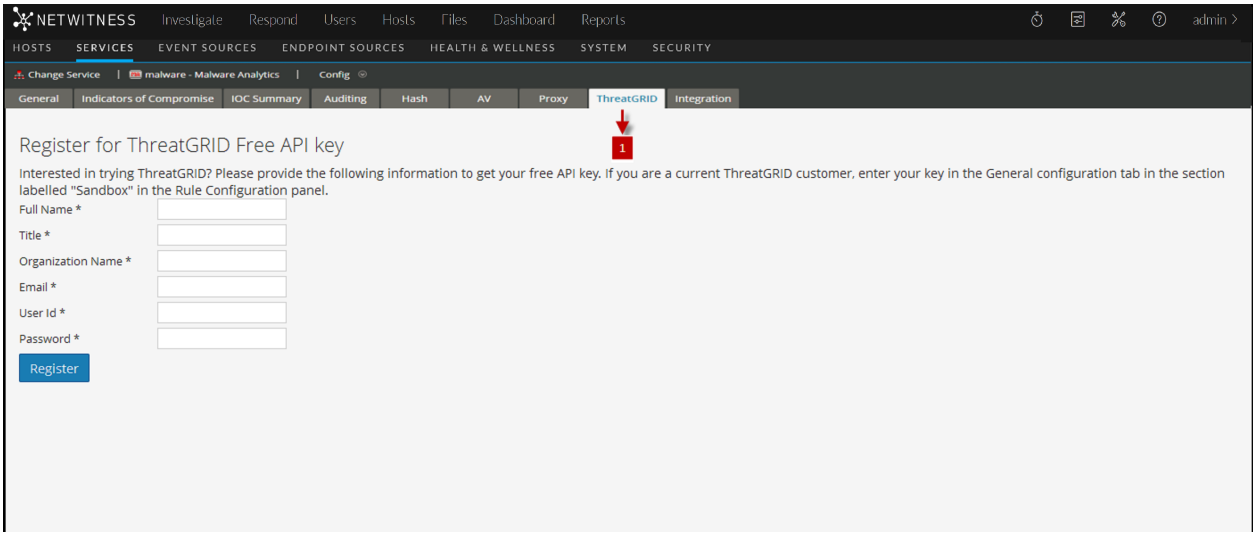
*You can perform this task in the current view

Related Topic

Basic Setup

Quick Look

This is an example of the ThreatGRID tab.



1 Displays the ThreadGRID Tab.

Features

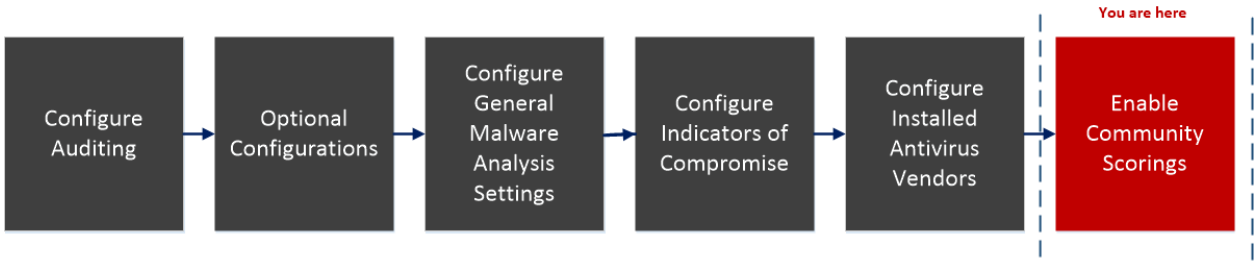
This table describes the features of the **ThreatGRID** tab.

Feature	Description
Full Name	Your first and last name.
Title	Your job title.
Organization Name	The name of your organization.
Email	Your email address.
User Id	Your user ID for ThreatGRID access.
Password	Your password for ThreatGRID access.
Register button	Click the **Register** button to submit the request.

Services Config View - Integration Tab

This topic introduces the features and functions of the Integration tab in the Administration Services Config view for Malware analysis. This tab provides a way to test connections and enable Community scoring by registering the Malware Analysis service. An administrator can test the connection to `cloud.netwitness.com` and to a core service that was configured for continuous scan.

Workflow



What do you want to do?

Role	I Want to...	Show me how
Administrator	Configure General Malware Analysis Settings	Configure General Malware Analysis Settings
Administrator	Configure Indicators of Compromise	Configure Indicators of Compromise
Administrator	Configure Auditing on Malware Analysis Host	(Optional) Configure Auditing on Malware Analysis Host
Administrator	Configure Hash Filter	(Optional) Configure Hash Filter
Administrator	Configure Installed Anti virus Vendor	Configure Installed Antivirus Vendors
Administrator	Configure Malware Analysis Proxy Settings	(Optional) Configure Malware Analysis Proxy Settings
Administrator	Register a TreadGRID API Key	(Optional) Register for a ThreatGRID API Key
Administrator	Enable Community Analysis*	Enable Community Analysis

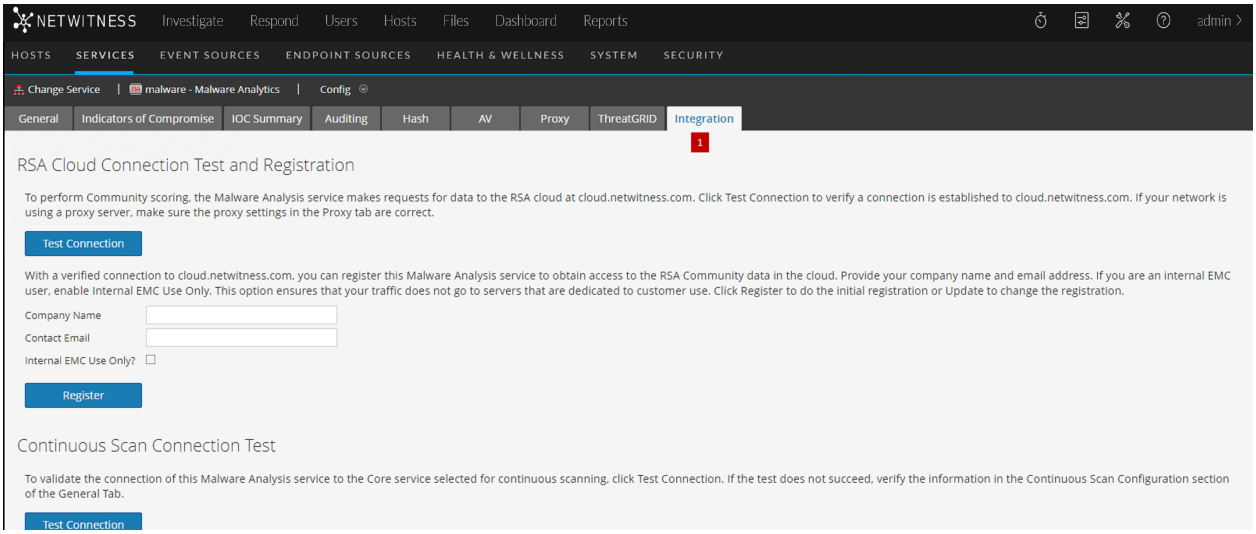
*You can perform this task in the current view

Related Topic

Basic Setup

Quick Look

The following figure is an example of the Integration tab.



1 Displays the Integration Tab.

Features

This tab has two sections: RSA Cloud Connection Test and Registration and Continuous Scan Connection Test. The following table describes the features.

Feature	Description
RSA Cloud Connection Test and Registration button	Clicking this button tests for an active connection to cloud.netwitness.com. NetWitness tests communications with the site and checks Proxy settings. A valid connection is required in order to register with the RSA Community Service.
Company Name	This is the name of your company. This is a required field.
Contact Email	This is the contact email. This is a required field.
Internal EMC Use Only Check box	This is an optional field. EMC customers, salespersons, or demo users should check this option to ensure that their requests do not use bandwidth on the production server. When the box is checked the following warning is displayed: `Checking this box may cause a less robust performance because the production server isn't being used.`
Register button	Clicking the Register button completes registration if all required fields are filled in. The Register button becomes the Update button after registration is complete.
Update button	The Update button is displayed after registration is complete.

Feature	Description
Continuous Scan Connection Test button	Clicking this button initiates a check to verify that the Malware Analysis service can connect to the Core service selected for continuous scanning (the Source Host, Source Port, Username, and User Password as specified in the General tab).

NetWitness® Platform XDR

Version 12.0

Malware Analysis User Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Malware Analysis Functions **6**

Functional Description 6

Analysis Method 8

NetWitness Server Access to the Malware Analysis Service 8

Scoring Method 9

Deployment 9

Malware Scoring Modules **10**

Network 10

Static Analysis 10

Community 11

Sandbox 11

Conducting Malware Analysis **12**

Begin a Malware Analysis Investigation 13

Fastest: Instant Launch from Malware Analysis Dashlets 13

On-Demand Polling from a Meta Value in the Navigate View 13

Investigate a Specific Service 13

Launch a Malware Investigation from a Malware Analysis Dashlet 13

Begin a Malware Analysis Investigation (No Default Service) 14

Set or Clear the Default Service 16

Upload and Scan Files 17

Begin an Investigation (Default Service Specified) 17

Apply Time Parameters Filter for Results 18

Apply a Threshold Filter to Continuous Mode Results 19

Delete or Resubmit an On-Demand Scan with New Bypass Settings 20

View the Files List 21

View the Events List 22

Implement Custom YARA Content 24

Prerequisites 24

YARA Version and Resources 24

Meta Keys in YARA Rules 24

YARA Content 25

Add Custom YARA Rules 26

Examine Scan Files and Events in List Form 28

Sort the Files List or Events List 29

Filter the List by Filename or MD5 File Hash 29

Download Files from the Files List ... 29
Delete Events from the Scan ... 30
Return to the Summary of Events ... 30
Open the Detailed Analysis for an Event ... 30
Configure the Malware Analysis Summary of Events View ... 31
Add a Dashlet ... 32
Modify or Delete a Dashlet Using Toolbar Options ... 32
Apply Threshold Filter to Multiple Dashlets ... 32
Set Title and Category Options for a Dashlet ... 33
Order Dashlets ... 34
Restore Default Dashlets ... 35
Filter Dashlet Data in the Summary of Events View ... 36
Configure the Score Wheel Dashlet ... 36
Zero-Day Candidates Example ... 37
Malicious Sessions Example ... 37
Arrange the Ring Sequence by Scoring Module ... 37
Configure the Meta Treemap Dashlet ... 38
Configure the Meta Breakdowns Dashlet ... 39
Configure the Events Timeline Dashlet ... 40
Open All Events in the Events List ... 40
Configure the Top Listing of Highly Suspicious Malware Dashlet ... 40
Configure the Malware with High Confidence IOCs and High Scores Dashlet ... 41
Configure the Top Listing of Possible Zero Day Malware Dashlet ... 41
Upload Files for Malware Analysis Scanning ... 43
Upload Files Manually ... 43
Upload Files from a Watched Folder ... 45
Import a Hash List ... 45
Import YARA rules to the IOC List ... 45
Import Files into the Scan Jobs List ... 46
View Detailed Malware Analysis of an Event ... 47
View Malware Analysis Details for an Event ... 47
Pivot Network Analysis Results ... 48
Use File Actions in the Static Analysis Results ... 48
View Community Analysis Results Details ... 49
View Sandbox Analysis Results in the ThreatGRID User Interface ... 50
Malware Analysis Reference Materials ... 52
Malware Analysis View ... 53
Workflow ... 53
What do you want to do? ... 53
Related Topics ... 54

Quick Look 54
Summary of Events Panel 55
Options Dialog 56
Meta Breakdowns 56
Meta Treemap 57
Score Wheel 58
Event Timeline 59
Malware Analysis Events List and Files List 61
Workflow 61
What do you want to do? 61
Related Topics 62
Quick Look 62
Scan For Malware Dialog 66
Workflow 66
What do you want to do? 66
Related Topics 67
Quick Look 67
Select a Malware Analysis Service Dialog 69
Workflow 69
What do you want to do? 69
Related Topics 70
Quick Look 70

Malware Analysis Functions

NetWitness Malware Analysis is an automated malware analysis processor designed to analyze certain types of file objects (for example, Windows portable executable (PE), PDF, and MS Office) to assess the likelihood that a file is malicious.

Malware Analysis detects indicators of compromise using four distinct analysis methodologies:

- Network Session Analysis (network)
- Static File Analysis (static)
- Dynamic File Analysis (sandbox)
- Security Community Analysis (community)

Each of the four distinct analysis methodologies is designed to compensate for inherent weaknesses in the others. For example, Dynamic File Analysis can compensate for Zero-Day attacks that are not detected during the Security Community Analysis phase. By avoiding malware analysis that strictly focuses on one methodology, the analyst is more likely to be shielded from false negative results.

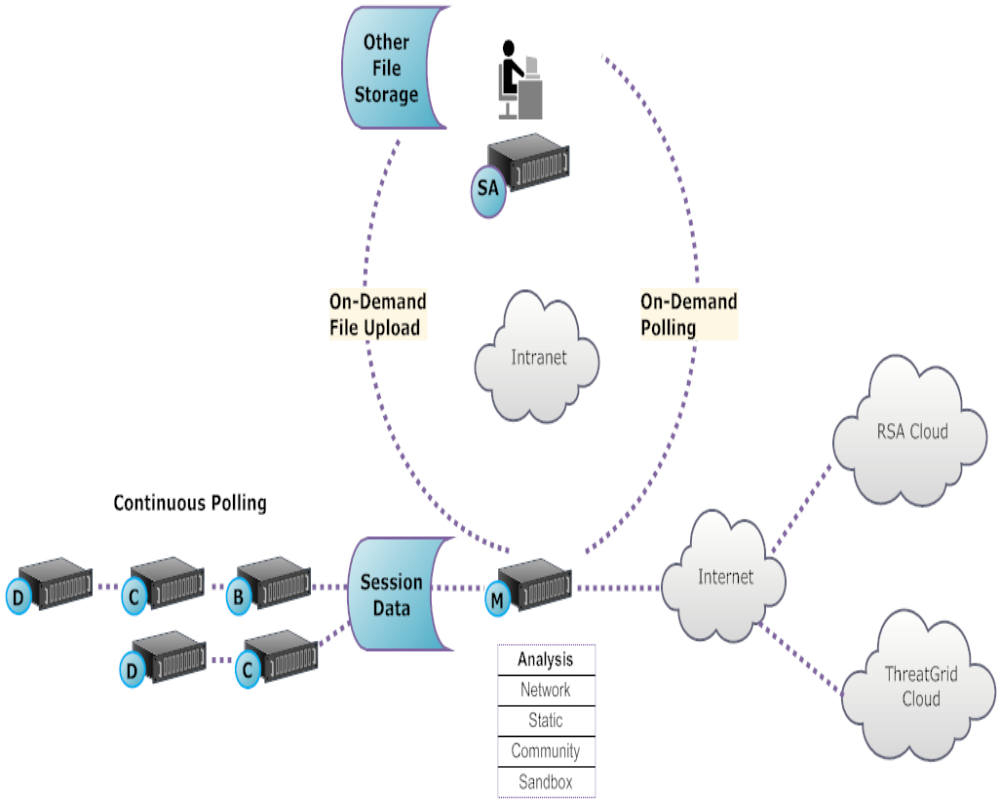
In addition to the built-in indicators of compromise, Malware Analysis supports indicators of compromise written in YARA. YARA is a rule language, which allows malware researchers to identify and classify malware samples. This allows IOC authors to add detection capabilities to RSA Malware Analysis by authoring YARA rules and publishing them in RSA Live. These YARA-based IOCs in RSA Live will automatically be downloaded and activated on the subscribed host, to supplement the existing analysis that is performed in each analyzed file.

Malware Analysis also has features that support alerts for Incident Management.

Functional Description

This figure depicts the functional relationship between the Core services (the Decoder, Concentrator, and Broker), the Malware Analysis service, and the NetWitness Server.

Daily Quota (Number of Files)	Free	Standard	Enterprise
Malware Analysis	100	unlimited	unlimited
ThreatGrid Analysis	5	1000	5000



The Malware Analysis service analyzes file objects using any combination of the following methods:

- **Continuous automatic polling of a Concentrator or Broker** to extract sessions identified by a parser as potentially carrying malware content.
- **On-demand polling of a Concentrator or Broker** to extract sessions identified by a malware analyst as potentially carrying malware content.
- **On-demand upload of files** from a user-specified folder.

When automatic polling of a Concentrator or Broker is enabled, the Malware Analysis service continuously extracts and prioritizes executable content, PDF documents, and Microsoft Office documents on your network, directly from data captured and analyzed by your Core service. Because the Malware Analysis service connects to a Concentrator or Broker to extract only those executable files that are flagged as possible malware, the process is both rapid and efficient. This process is continuous and does not require monitoring.

When on-demand polling of a Concentrator or Broker is chosen, the malware analyst uses Investigation to drill into captured data and choose sessions to be analyzed. The Malware Analysis service uses this information to automatically poll the Concentrator or Broker and to download the specified sessions for analysis.

On-demand upload of files provides a method for the analyst to review files captured external to the Core infrastructure. The malware chooses a folder location and identify one or more files to be uploaded and analyzed by Malware Analysis. These files are analyzed using the same methodology as files automatically extracted from network sessions.

Analysis Method

For the Network analysis, the Malware Analysis service looks for characteristics that seem to deviate from the norm, much as an analyst does. By looking at hundreds to thousands of characteristics and combining the results into a weighted scoring system, legitimate sessions that coincidentally have a few abnormal traits are dismissed, while the actual bad ones are highlighted. A user can learn patterns that indicate anomalous activity in the sessions as indicators that warrant further investigation, Indicators of Compromise.

The Malware Analysis service can perform Static analysis against suspicious objects it finds on the network and determine whether those objects contain malicious code. For Community analysis, new malware detected on the network is pushed to the RSA Cloud for checking against RSA's own malware analysis data and feeds from the SANS Internet Storm Center, SRI International, the Department of the Treasury and VeriSign. For Sandbox analysis, the services can also push data into major security, information and event management (SIEM) hosts (the ThreatGrid Cloud).

Malware Analysis has a unique method for analysis that is partnered with industry leaders and experts, so their technologies can enrich the Malware Analysis scoring system.

NetWitness Server Access to the Malware Analysis Service

The NetWitness Server is configured to connect to the Malware Analysis service and import tagged data for deeper analysis in Investigation. Access is based on three subscription levels.

- Free subscription: All NetWitness customers have a free subscription, with a free key for ThreatGrid analysis. The Malware Analysis service is rate-limited to 100 file samples per day. The number of samples (within the set of files from above) submitted to the ThreatGrid Cloud for sandbox analysis is limited to 5 per day. If one network session had 100 files in it, customers would hit the rate limit after processing the one network session. If 100 files were manually uploaded, that would cause the rate limit to be reached.
- Standard subscription tier: The number of submissions to the Malware Analysis service is unlimited. The number of samples submitted to the ThreatGrid Cloud for sandbox analysis is 1000 per day.
- Enterprise subscription tier: The number of submissions to the Malware Analysis service is unlimited. The number of samples submitted to the ThreatGrid Cloud for sandbox analysis is 5000 per day.

Scoring Method

By default, the Indicators of Compromise (IOC) are tuned to reflect industry best practices. During analysis, the IOCs that trigger cause the score to move upward or downward to indicate the likelihood that the sample is malicious. The tuning of IOCs is exposed in NetWitness so that the malware analyst can choose to override the assigned score or to disable an IOC from being evaluated. The analyst has the flexibility to either use the default tuning, or to completely customize the tuning to specific needs.

YARA-based IOCs are interleaved with the built-in IOCs within each built-in category and are not distinguished from native IOCs. When viewing IOCs in the Service Configuration view, administrators can select YARA from the Module selection list to see a list of YARA rules.

After a session is imported into NetWitness, all of the viewing and analysis capabilities in Investigation are available to further analyze Indicators of Compromise. When viewed in Investigation, YARA IOCs are distinguished from the built-in native IOCs by the tag `Yara rule`.

Deployment

The Malware Analysis service is deployed as a separate RSA Malware Analysis host. The dedicated Malware Analysis host has an on-board Broker which connects to the Core infrastructure (either another Broker or a Concentrator). Prior to this connection, a collection of parsers and feeds must be added to the Decoders that are connected to the Concentrators and Brokers from which the Malware Analysis service pulls data. This allows suspicious data files to be marked for extraction. These files are `malware analysis` tagged content available through the RSA Live content management system.

Malware Scoring Modules

NetWitness Malware Analysis analyzes and scores sessions and the embedded files within these sessions by scoring four categories: Network, Static Analysis, Community, and Sandbox. Each category comprises many individual rules and checks that are used to calculate a score between 1-100. The higher the score, the more likely the session is to be malicious and worthy of more in-depth follow-on investigation.

Malware Analysis can facilitate a historical investigation into events leading up to a network alarm or incident. If you know that a certain type of activity is taking place on your network, you can select only the reports of interest to examine the content of data collections. You can also modify behavior for each scoring category based on the scoring category or the file type (Windows PE, PDF, and Microsoft Office).

Once you become familiar with data navigation methods, you can explore the data more completely through:

- Searching for specific types of information
- Reviewing specific content in detail.

Category scores for Network, Static Analysis, Community, and Sandbox are maintained and reported independently. When events are viewed based on the independent scores, as long as one category detects malware, it is evident in the Analysis section.

Network

The first category examines each core network session to determine if the delivery of the malware candidates was suspicious. For example, benign software being downloaded from a well-known safe site, using proper ports and protocols, is considered less suspicious than downloading software known to be malicious from a known dubious download site. Sample factors used in the scoring of this criteria set may include sessions that:

- Contain threat feed information
- Connect to well-known bad sites
- Connect to high-risk domains/countries (for example, .cc domain)
- Use well-known protocols on non-standard ports
- Contain obfuscated JavaScript

Static Analysis

The second category analyzes each file in the session for signs of obfuscation in order to predict the likelihood of the file behaving maliciously if allowed to run. For example, software that links to networking libraries is more likely to perform suspicious network activity. Sample factors used in the scoring of this criteria set may include:

- Files found to be XOR encoded
- Files found embedded within non-EXE formats (for example, PE file found embedded in a GIF format)
- Files linking to higher risk import libraries
- Files highly deviating from the PE Format

Community

The third category scores the session and files based on the collective knowledge of the security community. For example, files whose fingerprint/hash is already known to be good or bad by respected anti-virus (AV) vendors is scored accordingly. Files are also scored based on knowledge that a file was delivered from a site known to be good or bad by the security community.

Community scoring also indicates whether the AV on your network flagged the files as malicious. It does not indicate that the resident AV product acted to protect your system.

Sandbox

The fourth category examines the behavior of the software by actually running it in a sandbox environment. By running the software to watch its behavior, a score can be calculated by identifying well-known malicious activity. For example, software that configures itself to autostart on each reboot and make IRC connections would score higher than a file with no known bad behavior.

Conducting Malware Analysis

Analysts can use theNetWitness Malware Analysis service to detect malware in selected data and files.

Analysts who conduct analysis using NetWitness Malware Analysis need to have the appropriate system roles and permissions set up for their user accounts.

The following procedures provide instructions for using Malware Analysis:

- Begin a Malware Analysis Investigation
- Upload Files for Malware Analysis Scanning
- Implement Custom YARA Content
- Configure the Malware Analysis Summary of Events View
- Filter Dashlet Data in the Summary of Events View
- Examine Scan Files and Events in List Form
- View Detailed Malware Analysis of an Event

Begin a Malware Analysis Investigation

You can investigate data that has been scanned, flagged, and rated by Malware Analysis as containing Indicators of Compromise. This includes all types of Malware Analysis scans: continuous mode polling, on-demand polling, and on-demand uploaded files. Continuous mode polling must be enabled when the administrator configures basic settings for the Malware Analysis service.

NetWitness provides several methods of launching a Malware Analysis investigation.

Fastest: Instant Launch from Malware Analysis Dashlets

The fastest way to begin a Malware Analysis investigation is an Instant launch from the NetWitness Dashboard using one of the Malware Analysis dashlets that lists events or files that are likely to contain malware. The dashlets are described as part of the NetWitness Content in Dashlets. From one of these dashlets, you can go directly to the Analysis Results for a specific event that has been listed as worthy of investigation:

- Top Listing of Highly Suspicious Malware
- Top Listing of Possible Zero Day Malware
- Malware with High Confidence IOCs and High Scores Dashlet

On-Demand Polling from a Meta Value in the Navigate View

You can initiate on-demand polling from within an investigation by right-clicking a meta value in the Navigate view, and choosing an option from the context menu. When polling is complete, the scanned data is available for malware analysis (see "Launch a Malware Analysis Scan from the Navigate View" in the *NetWitness Investigate User Guide*).

Investigate a Specific Service

You can also begin a Malware Analysis investigation of a service in the Investigate > Malware Analysis view. For Malware Analysis investigation on a service basis, a service must be specified in the Investigate > Malware Analysis view.

1. Investigate opens the Malware Analysis view with the user-specified default service selected.
2. If no default service is currently specified, a dialog allows you to select the Malware Analysis service to investigate.
3. When a service has been selected in the Malware Analysis view, the Summary of Events for the selected service and continuous scan data for the service is displayed.

This topic provides instructions for all methods of launching a Malware Analysis investigation.

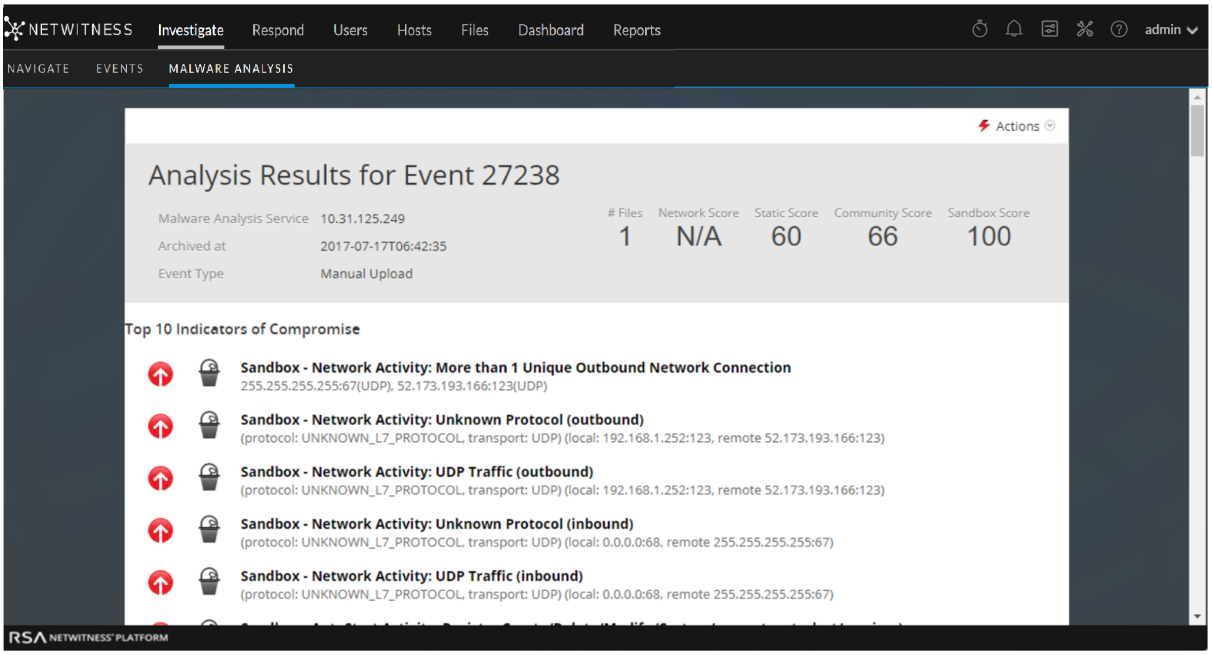
Launch a Malware Investigation from a Malware Analysis Dashlet

A prerequisite for this procedure is that one of the following dashlets must be visible in the NetWitness dashboard or in the Malware Analysis view, and must be populated with listed events or files. If you do not see the dashlets, add them and configure the dashlets.

- Top Listing of Highly Suspicious Malware
- Top Listing of Possible Zero Day Malware
- Malware with High Confidence IOCs and High Scores Dashlet

To launch a Malware Analysis investigation from a dashlet:

1. Log in to NetWitness and look for one of the above dashlets in the Monitor view or in the Malware Analysis view
2. In the dashlet, double-click an event or file for deeper analysis. A detailed analysis of the event in the Events List or the event with which the file in the File List is associated is displayed in the Malware Analysis view.



To learn more about configuring the Malware Analysis dashlets in the Monitor dashboard, see "Dashlets" in the *NetWitness Platform Getting Started Guide*.

To learn about the ways you can configure and filter information in dashlets in the Malware Analysis view, refer to Filter Dashlet Data in the Summary of Events View.

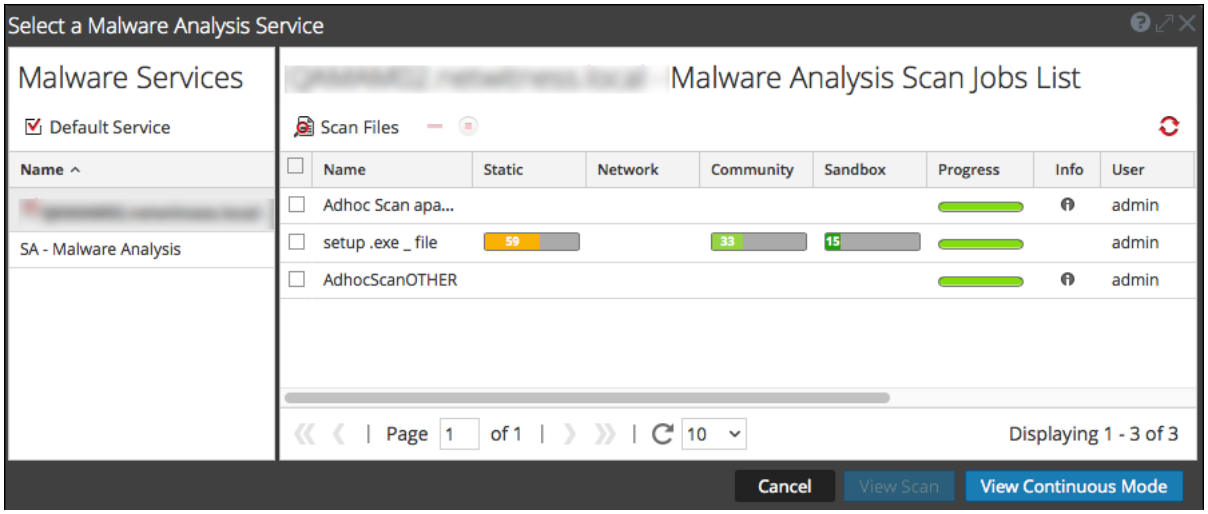
To learn about the actions you can perform in the Analysis Results, refer to View Detailed Malware Analysis of an Event.

Begin a Malware Analysis Investigation (No Default Service)

To begin an investigation with no default service specified:

1. Select **Investigate > Malware Analysis**.

   The Select a Malware Analysis Service dialog is displayed, with available Malware Analysis hosts and services for the current user in the left panel and available scan jobs in the right panel. This scan jobs panel contains the same columns as the Malware Scan Jobs dashlet in the Unified dashboard. In addition, it has a toolbar and View options, which are described in Select a Malware Analysis Service Dialog.



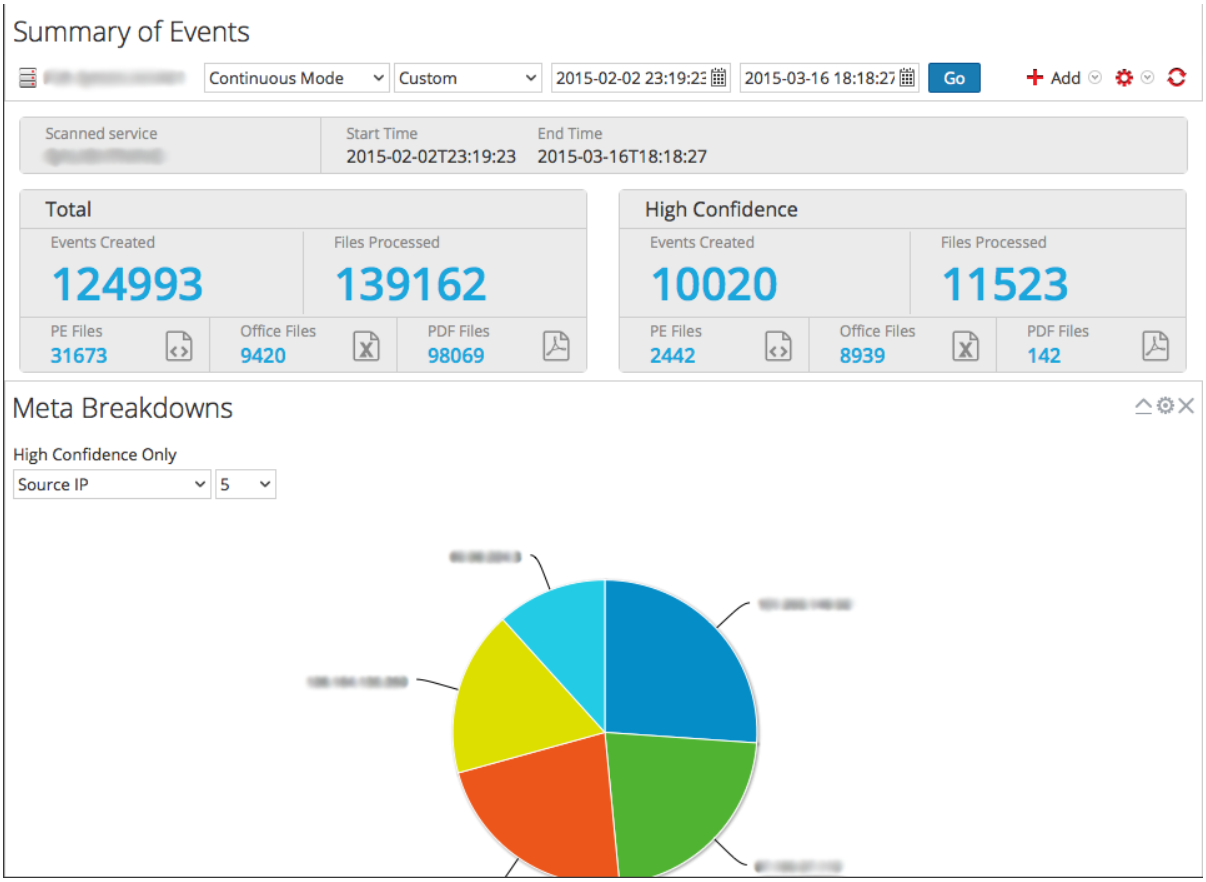
2. In the list of Malware Analysis hosts, select a host and a list of scan jobs is displayed in the right panel. These jobs are created when you scan an event or a file (see Upload Files for Malware Analysis Scanning and "Launch a Malware Analysis Scan from the Navigate View" in the *NetWitness Investigate User Guide*.
3. To begin analyzing a scan, do one of the following:

   a. Select a scan and click **View Scan**.

   b. Click **View Continuous Mode**.

      The Summary of Events for the selected scan is displayed with the default dashlets open. Each user can add, modify, and delete default dashlets, which persist through different scan investigations. Users can also restore default dashlets as described in Filter Dashlet Data in the

Summary of Events View.



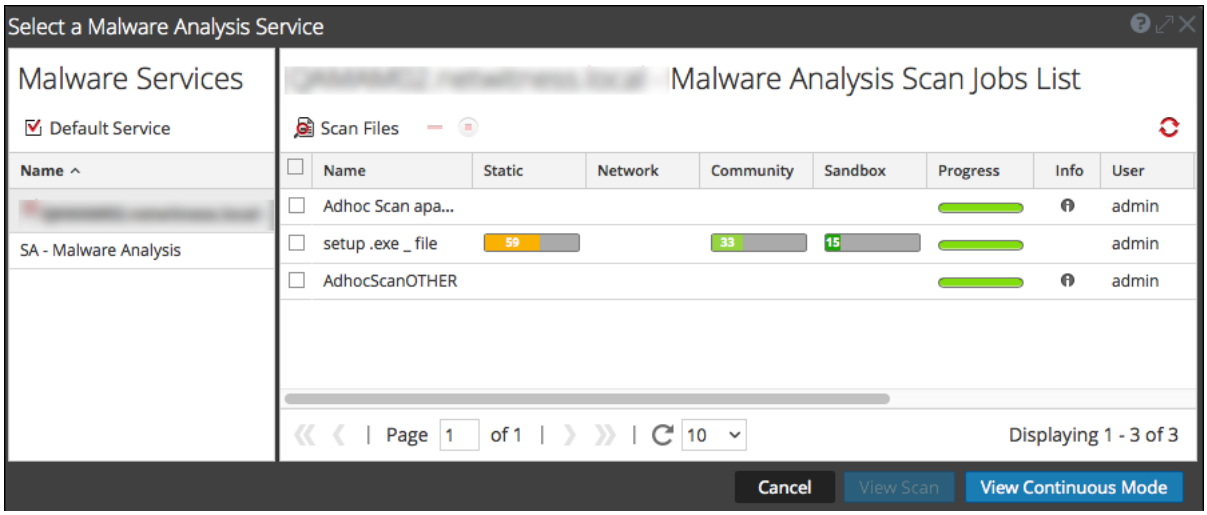
Set or Clear the Default Service

You can set the default service and clear the default service in the Select a Malware Analysis Service dialog.

To set a default service:

1. Click the service name in the Summary of Events toolbar.

   The Select a Malware Analysis Service dialog is displayed.



2. Select a service on the list of available Malware services, and click Default Service.

   The service becomes the default, (indicated by ☑ in front of the host name).

3. To clear the default service, select the default service in the grid, and click Default Service.

   No default service is set.

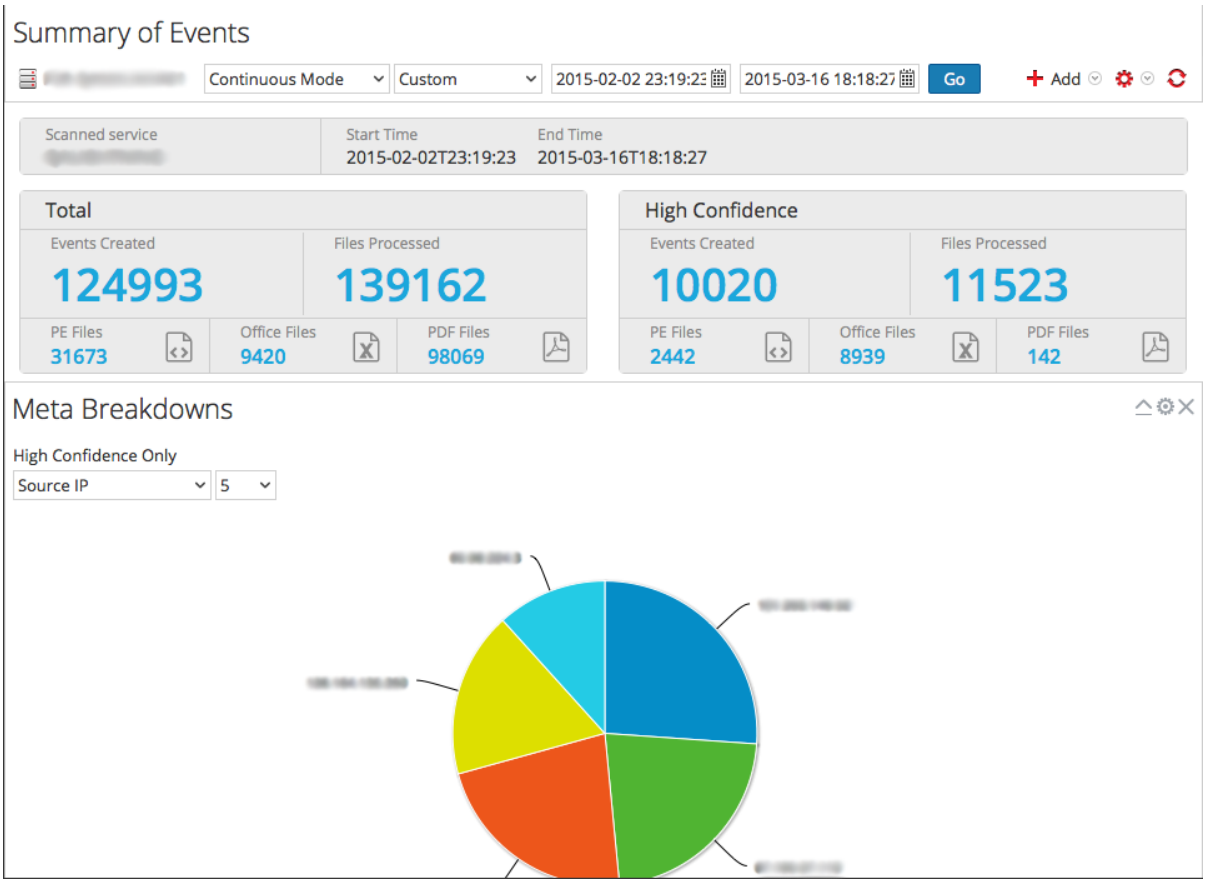
Upload and Scan Files

A Malware Analyst with permission to `Initiate Malware Analysis Scan` can upload files to scan using the Scan Files option in the Select a Malware Analysis Service dialog (see Upload Files for Malware Analysis Scanning. An administrator can upload packet capture files to a Decoder for Malware Analysis in the Services System view as described in "Upload Packet Capture File" in the *Decoder and Log Decoder Configuration Guide*.

Begin an Investigation (Default Service Specified)

To begin an investigation with a default service specified, select **Investigate > Malware Analysis**.

The Summary of Events for a continuous scan of the selected service is displayed with the default dashlets open. Each user can add, modify, and delete default dashlets, which persist through different scan investigations. Users can also restore default dashlets as described in Filter Dashlet Data in the Summary of Events View.



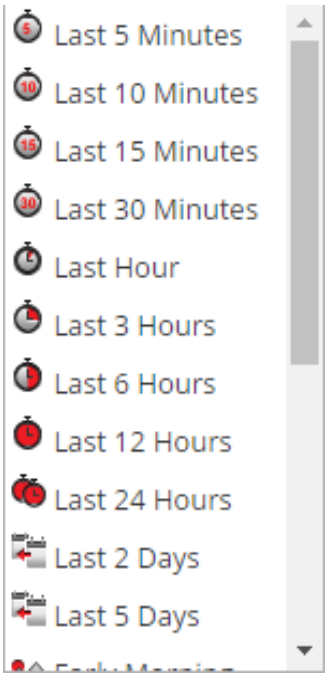
Apply Time Parameters Filter for Results

You can apply a Threshold filter to refresh the results of the chosen dashlets.

1. To select a different time range, select either **Continuous Mode** or a different scan from the toolbar.

   The Malware Summary of Events for the selected scan is displayed.

2. To select a new time range for the scan, click in the range selection list in the toolbar. Ranges available are: Last 5 minutes, Last 10 minutes, Last 15 minutes, Last 30 minutes, Last Hour, Last 3 Hours, Last 6 Hours, Last 12 Hours, Last 24 Hours, Last 2 Days, Last 5 Days, Early Morning, Morning, Afternoon, Evening, All Day, Yesterday, This Week, Last Week, or Custom.



The results are updated immediately.

3. To refresh a continuous mode scan with new data, click .

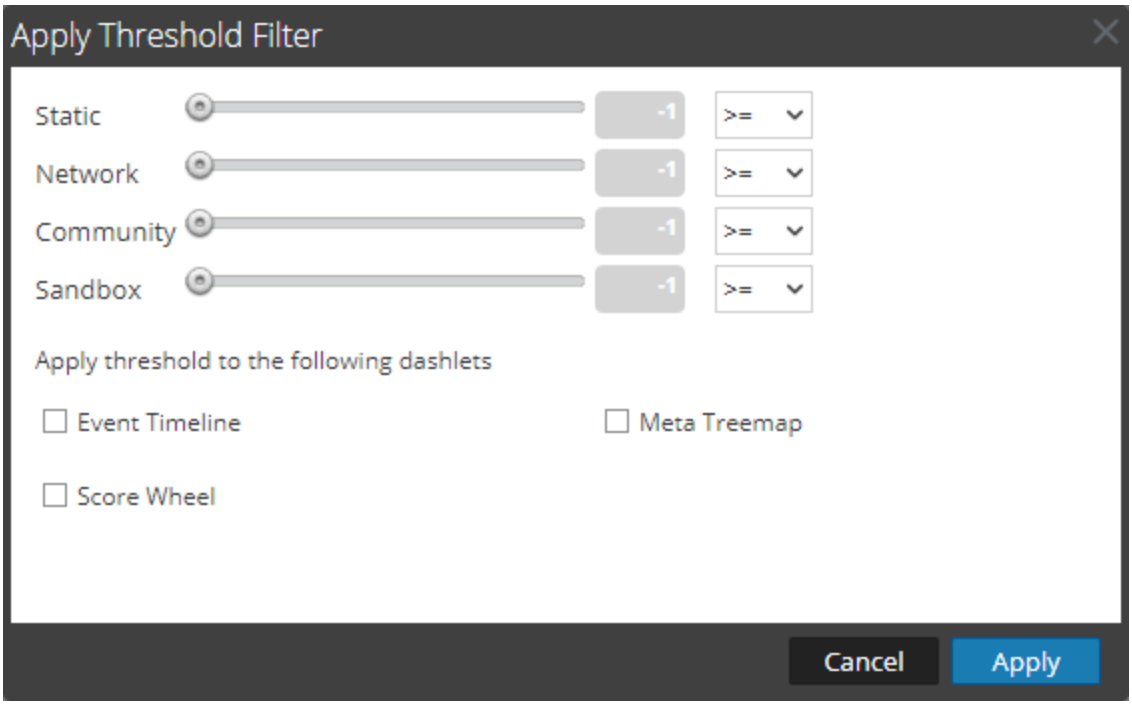
Apply a Threshold Filter to Continuous Mode Results

You can apply a new threshold filter to an instance of the Malware with High Confidence IOCs and High Scores dashlet, the Meta Treemap dashlet, the Score Wheel dashlet, and the Event Timeline dashlet.

To customize the scoring applied to the scan, in the toolbar, do the following:

1. Select > **Apply Threshold Filter**.

   The Apply Threshold Filter dialog is displayed.



2. If you want to limit the number of events displayed to events that were given a score above a certain number, do the following:
   a. Drag the slider in the Static, Network, Community, and Sandbox slider bars.
   b. To select the dashlets in which the thresholds apply, select the appropriate checkboxes.
   c. Click **Apply**.

Delete or Resubmit an On-Demand Scan with New Bypass Settings

You can delete an on-demand scan or resubmit an on-demand scan with different bypass settings than those specified in the Service Configuration view for a Malware Analysis service.

To delete a scan while viewing an on-demand scan, do the following:

1. Select **Actions > Delete Scan**.

   A dialog asks for confirmation that you want to delete the scan.

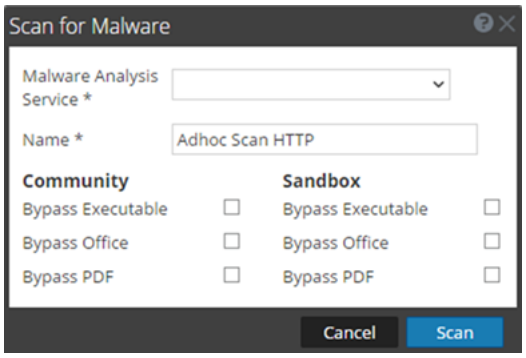
2. Click **Yes**.

   The selected scan is deleted.

To apply different bypass settings to the current scan:

1. Select **Actions > Resubmit Scan**.

   The Scan for Malware dialog is displayed.



2. Select the bypass settings that you want to use on the new scan, and click **Scan**.

   Malware Analysis resets cache and resubmits the file for a new scan, and the scan jobs are added to the jobs queue.

3. When the job is complete, scroll to the left and select **View**.

   The Malware Summary of Events for the selected scan is displayed.

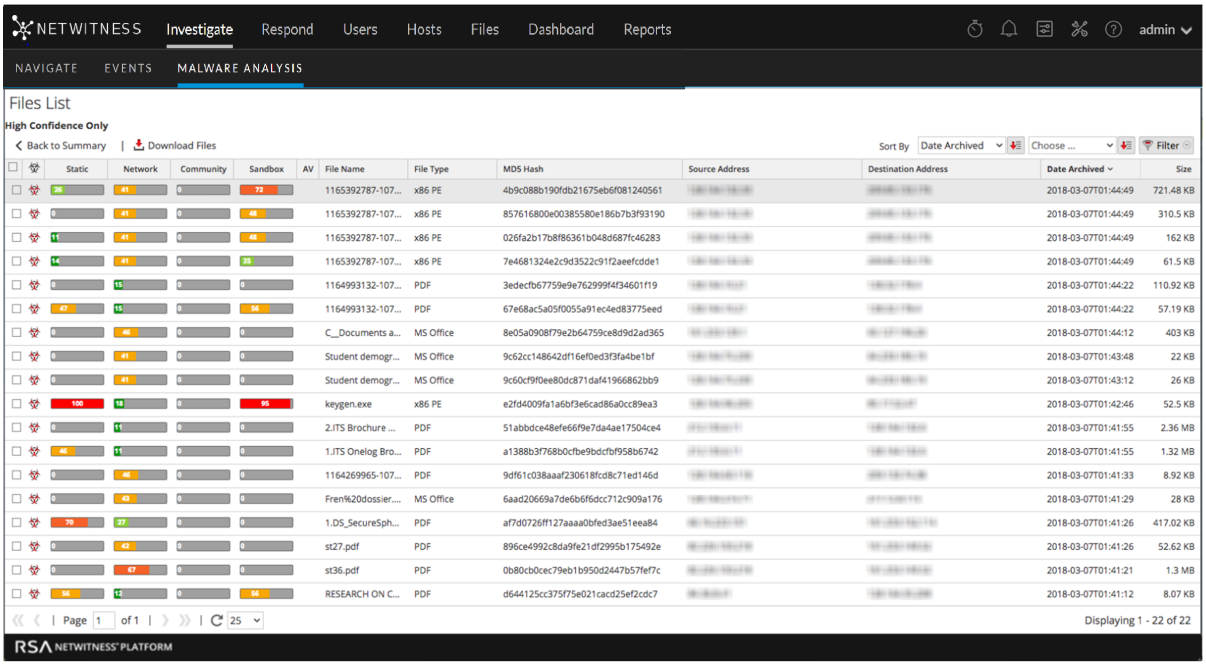
View the Files List

You can view a list of files for an event from the Malware Analysis Summary of Events and from each of the Visualization charts: Event Timeline, Meta Breakdowns, Meta Treemap, and Score Wheel.

To view the Files List, do one of the following:

- In the Summary of Events, click on the number of files in the **Total** row or the **High Confidence** row under **Files Processed**, **PE Files**, **Office Files**, or **PDF Files**. The Files List is displayed.
- In any visualization dashlet, click the number next to the **Files** field in the top right corner of the dashlet.

The Files List for the selected drill point is displayed.



From the Files List, you can search for a file by filename or MD5 file hash, sort the list using two criteria and ascending or descending order, and download files as described in Examine Scan Files and Events in List Form.

To return to the Summary of Events, click **Back to Summary**.

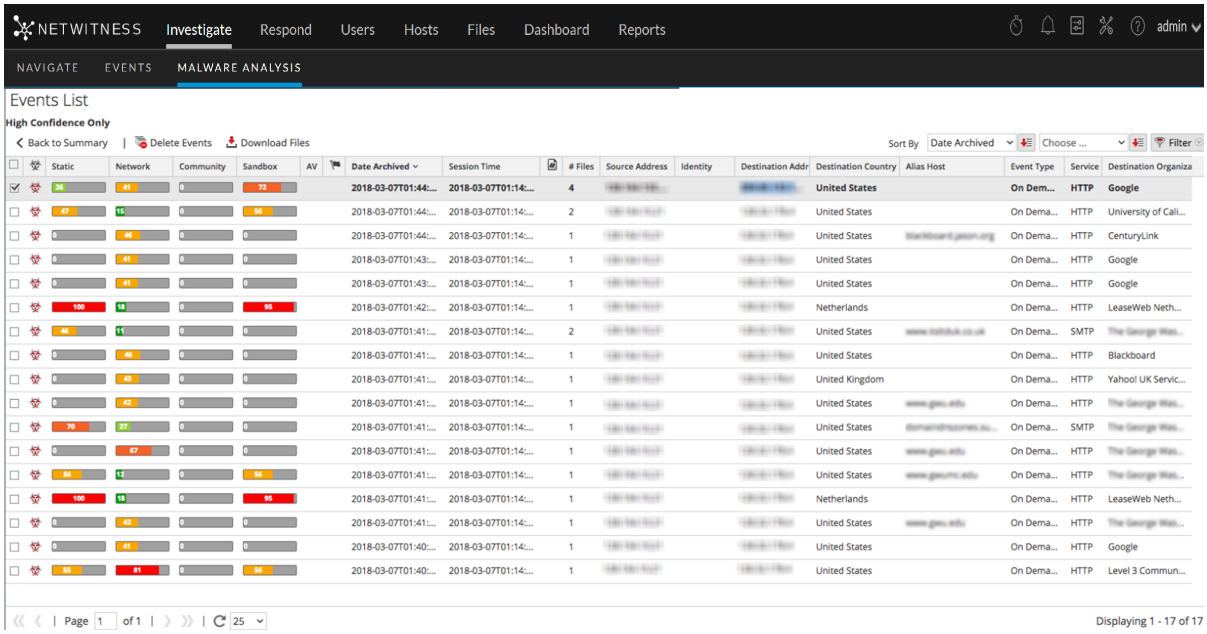
View the Events List

From the Malware Analysis Summary of Events and from each of the visualization charts (Event Timeline, Meta Breakdowns, Meta Treemap, and Score Wheel), you can select events to view in the Events grid.

To view the Events List, do one of the following:

- In the Summary of Events, click the number of Events Created in the **Total** row or the **High Confidence** row. The Events List is displayed.
- In any visualization dashlet, click the number next to the Events field in the top right corner of the dashlet.

The Events List for the selected time is displayed.



Implement Custom YARA Content

In addition to the built-in indicators of compromise, Malware Analysis supports indicators of compromise written in YARA. YARA is a rule language that allows malware researchers to identify and classify malware samples. It makes built-in YARA-based Indicators of Compromise (IOCs) available in Live; these are automatically downloaded and activated on subscribed hosts.

Customers with advanced skills and knowledge can add detection capabilities to Malware Analysis by authoring YARA rules and publishing them in Live or placing YARA rules in a watched folder for the host to consume.

As malware and the threat landscape evolve, it is important to review and examine existing custom rules. Updates are often necessary to incorporate new detection methods. It also updates YARA rules in Live from time to time. To receive updates, you can subscribe to the Blog and Live at http://blogs.rsa.com/feed.

This document provides information to help customers implement custom YARA rules in Malware Analysis.

Prerequisites

The host on which you are adding custom rules must be configured to support authoring of YARA rules as described in "Enable Custom YARA Content" in the *Malware Analysis Configuration Guide*.

YARA Version and Resources

Malware Analysis is packaged with YARA version 1.7 (rev:167). To find out the exact version, you can run `yara -v` on the Malware Analysis host as shown in this example:

```
[root@TESTHOST yara] # yara -v
 yara 1.7 (rev:167)
```

Meta Keys in YARA Rules

Malware Analysis is compliant with other sources of YARA rules, and it also consumes additional meta keys that are specific to Malware Analysis. Each YARA rule is equivalent to an Indicator of Compromise (IOC) within Malware Analysis. The example below illustrates the meta definitions in a rule:

```
meta:
 iocName = "FW.ecodedGenericCLSID"
        fileType = "WINDOWS_PE"
        score = 25
        ceiling = 100
        highConfidence = false
```

Meta Key	Description
iocName	(Required) This is the name that MA uses as the rule name. It is specific to Malware Analysis and is required to add the rule to the IOC list.

Meta Key	Description
fileType	Specifies the files type. Possible values are: WINDOWS_PE, MS_OFFICE, and PDF. If not specified, the default value is WINDOWS_PE.
score	This value that is added to the static score if the YARA rule is triggered. If not specified, the default value is 10.
ceiling	This is the maximum amount that is added to the static scores when a rule is triggered multiple times in one session. For example, if each time a rule is triggered, 20 points are added to the static, and you do not want more that 40 points added when the rule is triggered more than two times, you can specify a ceiling of 40. If not specified, the default value is 100.
highConfidence	This sets the High Confidence flag, which is set on IOCs when there are high confidence indicators that malware is present. If not specified, the default file value is false.

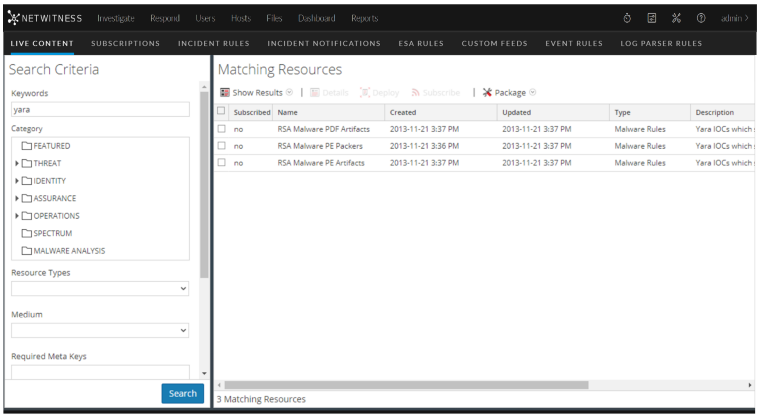
Note: Refer to the following URL for YARA resources: https://code.google.com/p/yara-project/downloads/list. NetWitness uses YARA 1.7, not YARA 2.0.

YARA Content

Live contains 3 sets of Yara rules:

- PE Packers
- PDF Artifacts
- PE Artifacts

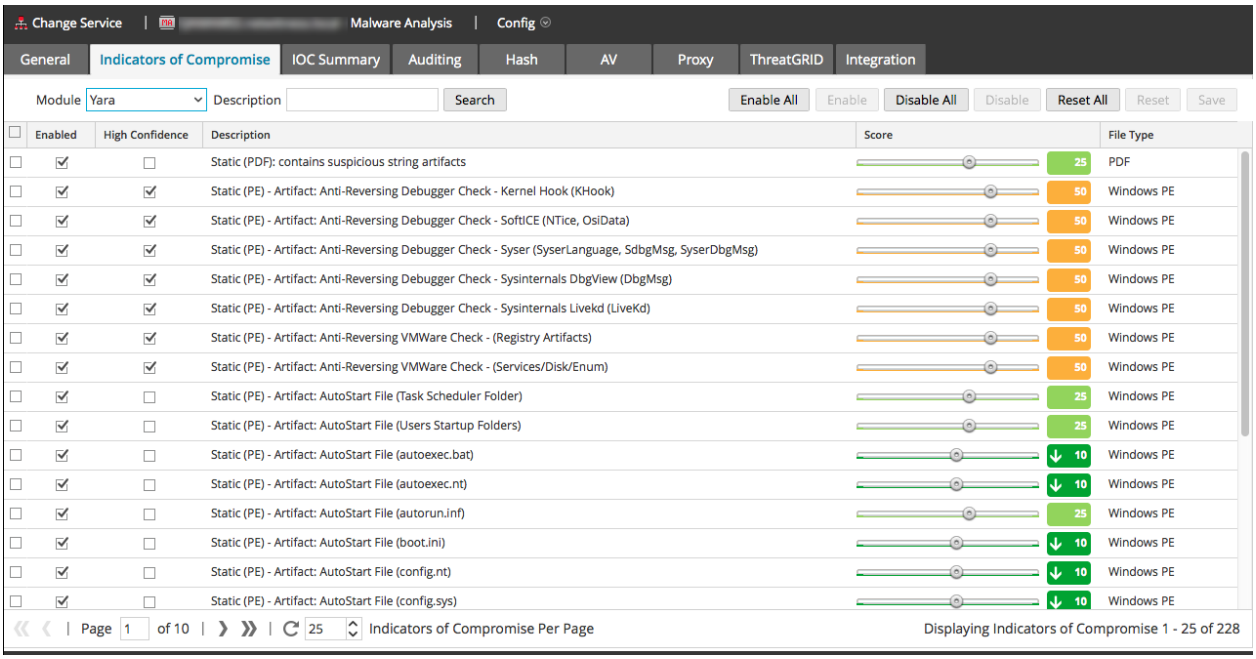
The following figure illustrates YARA content available as YARA rules in NetWitness Live.



On the Malware Analysis host, the YARA rules reside in `/var/lib/netwitness/malware-analytics-server/spectrum/yara`, as shown in the example below.

```
[root@TESTHOST yara]# pwd
 /var/lib/netwitness/malware-analytics-server/spectrum/yara
 [root@TESTHOST yara]# ls *.yara
 rsa_mw_pdf_artifacts.yara rsa_mw_pe_artifacts.yara rsa_mw_pe_packers.yara
```

The individual rules are listed as IOCs in the Malware Analysis Service Config view > Indicators of Compromise tab. To view them, use the Yara module as the filter. You can adjust the configuration of an individual in the same way that you configure other IOCs.



Add Custom YARA Rules

To introduce custom YARA rules from other sources:

1. To ensure that the YARA rules follows the correct format and syntax, use the YARA command to compile the YARA rule as shown in the following example. If the rule compiles with no errors, this indicates that the YARA rule has the correct syntax.

   ```
   [root@TESTHOST yara]# yara rsa_mw_pe_packers.yara dummy.txt
   [root@TESTHOST yara]#
   ```

2. Ensure that custom rules do not duplicate existing YARA rules from other sources. All YARA rules are in `/var/lib/netwitness/malware-analytics-server/spectrum/yara`

3. Ensure that the meta keys supports are included to organize the YARA rules as part of the configurable IOCs, and name the file with the yara extension (<filename>.yara). For better organization, make sure that the `iocName` meta is included in the meta section as shown in the following example.

   Example:

   ```
   rule HEX_EXAMPLE
   {
       meta:
   ```

```
        author = "RSA"
        info = "HEX Detection"
      iocName = "Hex Example"
    strings:
        $hex1 = { E2 34 A1 C8 23 FB }
        $wide_string = "Ausov" wide ascii
    condition:
        $hex1 or $wide_string
}
```

4. When ready, place the custom YARA file in the folder that the Malware Analysis service watches:
   `/var/lib/netwitness/malware-analytics-server/spectrum/yara/watch`
   The file is consumed within one minute.
   Once consumed, NetWitness moves the file to the `processed` folder, and the new rule is added to the Malware Analysis Services Config view > Indicators of Compromise tab.

Note: If you fail in adding the custom YARA rule, after step 4, do the following:
1. Check the custom YARA file is in `/var/lib/netwitness/malware-analytics-server/spectrum/yara/watch/error`.
2. Check the logs `/var/lib/netwitness/malware-analytics-server/spectrum/logs/spectrum.log`.

Examine Scan Files and Events in List Form

When viewing the Summary of Events in a Malware Analysis scan, you can click a file count or an event count to view the Files List or the Events List for the scan (see Begin a Malware Analysis Investigation). In the Files List and Events List, you can search for a file by filename or MD5 file hash, sort the list using two criteria and ascending or descending order, and download files. When you find an event or file of interest in the Events List or Files List, you can view many details about the event in the Event Details view.

For each event in the Events List, NetWitness provides the following information:

- Flagged as a High Confidence event, which is considered likely to contain Indicators of Compromise.
- The numeric score for each scoring module: Static, Network, Community, and Sandbox.
- Antivirus vendor scores.
- The Influenced by customized rule flag.
- The date the event was archived.
- The session time.
- The MD5 hash filter.
- The number of files in the event.
- The source IP address of the event.
- The Identity.
- The destination IP address.
- The destination country.
- The name of the alias host.
- The event type, for example, Network.
- The service used by the event.
- The destination organization

For each file in the Files List, NetWitness provides the following information:

- Flagged as a High Confidence event, which is considered likely to contain Indicators of Compromise.
- The numeric score for each scoring module: Static, Network, Community, and Sandbox.
- Antivirus vendor scores.
- The filename.
- The file type.
- The MD5 hash filter.
- The source IP address of the event that contained the file.

- The destination IP address.
- The date the event that contained the file was archived.
- The file size.

Sort the Files List or Events List

You can sort the Files List and Events List by column name in ascending and descending order. You can choose one or two columns.

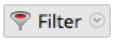
To sort the list:

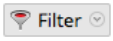
1. In the first **Sort By** drop-down list, choose a column name and sort direction:  for descending order or  for ascending order.
2. (Optional) In the second **Sort By** drop-down list, choose a column name, and sort direction,  for descending order or  for ascending order.
   The column titles reflect the selected sort order.

Filter the List by Filename or MD5 File Hash

You can filter the Files List and Events List by filename or file hash. With this feature, you can specify a limited subset of the original data based on the search criteria.

Note: When you perform a search, you search the scan that you are currently displaying, not all scans.

1. Click 

   .
   The Filter dialog is displayed.
2. Enter a value in **File Name** or **MD5 Hash** and click **Filter**. The File Name and Hash field are not case sensitive. Wild card or regular expressions are not supported. The filter is based on exact matches. You can drag across a filename or hash to select from the Files list or Events list, then copy and paste it in the dialog.
3. Click **Filter**.
   Malware Analysis filters the list to display only files or events with the selected hash
4. To revert to the unfiltered list, click 

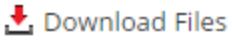
   . When the Filter dialog is displayed, click **Reset**.

Download Files from the Files List

NetWitness lets you select and download files from the Files List or the Events List.

Caution: Use caution when downloading files from Malware Analysis; some files may contain harmful code. File Download is a specific permission that can be configured, refer to "Define Roles and Permissions for Malware Analysts" in the *Malware Analysis Configuration Guide* for more details.

To download files from the Files List or Events List:

1. In the **Files List** or **Events List**, select the checkbox next to one or more rows.
2. In the toolbar, select 

.
   The Malware File Download dialog is displayed.
3. Do one of the following:
   a. If you decide not to download the file, click **Cancel**.
   b. If you want to download the file, select click the **Download** button.
      The file or files selected are downloaded in a zip archive with the name `Malware_Files.zip`.

Delete Events from the Scan

In the Events List, you select one or more events and delete them from the scan. This is useful for removing events that are not of interest.

To remove an event from the scan being viewed:

1. In the **Events List**, select one or more events.
2. In the toolbar, click 

.
   NetWitness asks for confirmation that you want to delete the events.
3. In the confirmation dialog, click **Yes**.
   The selected events are deleted.

Return to the Summary of Events

To leave the Files List or Events List and return to the Summary of Events, click **Back to Summary**.

Open the Detailed Analysis for an Event

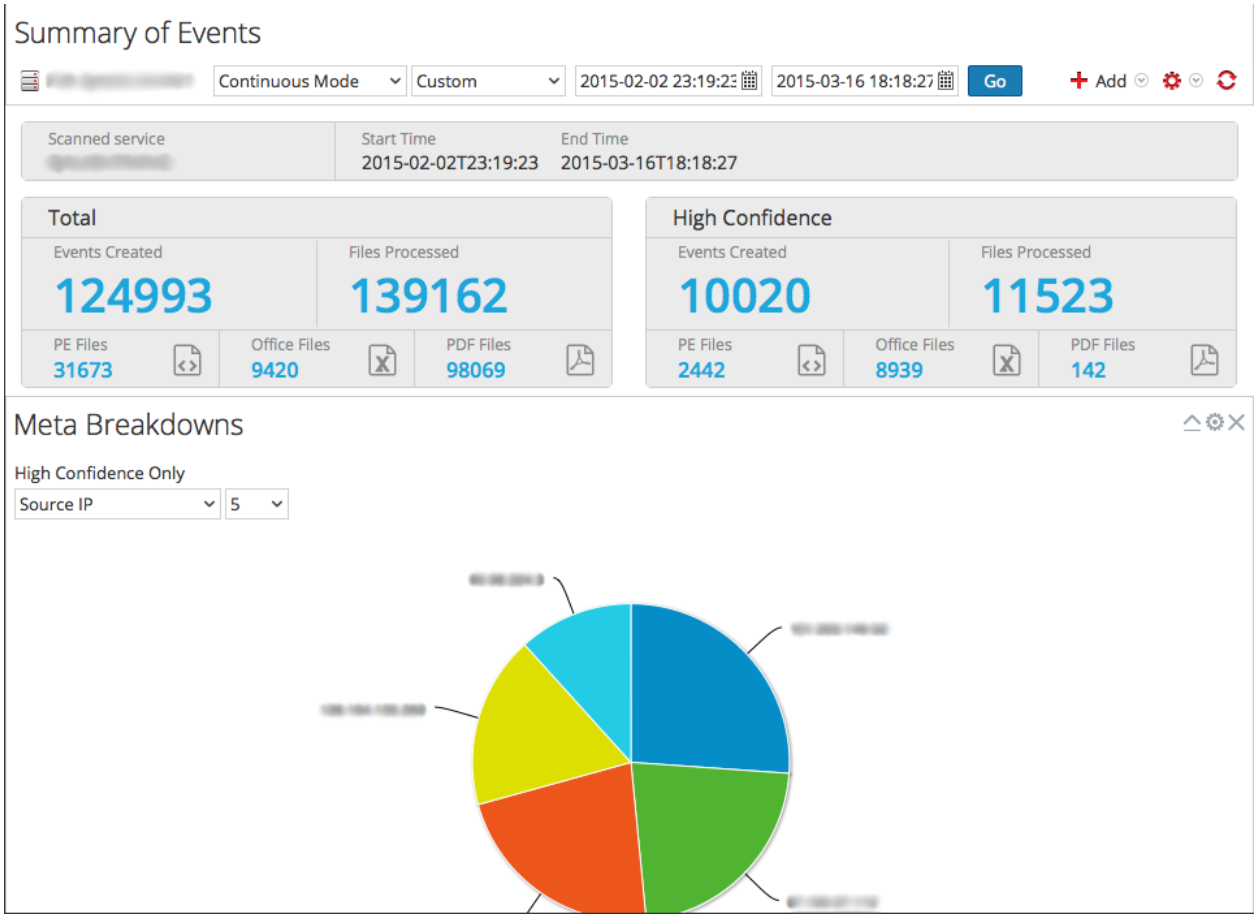
While you examine events or files in the Files List or Events List, you can double-click any event or file to open a detailed analysis of the event in the Events List or the event with which the file in the Files List is associated (see View Detailed Malware Analysis of an Event).

Configure the Malware Analysis Summary of Events View

The Summary of Events provides a summary of the scan being investigated, and below the summary are configurable dashlets such as visualization charts and listings. By default, the Summary of Events for a scan opens with the default dashlets displayed. You can customize the view by adding, modifying, and deleting default dashlets. The configured customization of dashlets persists through different scan investigations, and you can restore default dashlets at any time. The default dashlets are:

- Summary of Events (Fixed)
- Event Timeline
- Top Listing of Highly Suspicious Malware
- Meta Treemap
- Score Wheel
- Meta Breakdowns

The following figure is an example of the default Summary of Events.



The rest of this topic provides instructions for managing and configuring dashlets.

Add a Dashlet

You can add multiple copies of dashlets in the Malware Analysis Summary of Events. To add a dashlet:

1. In the toolbar, select **Add**.
   The drop-down list of dashlets is displayed. There are four visualization options: Score Wheel, Meta Treemap, Meta Breakdowns, and Event Timeline. The other three dashlets are the same dashlets available in the NetWitness dashboard: Malware with high Confidence IOCs and High Scores, Top Listing of Highly Suspicious Malware, Top Listing of Possible Zero Day Malware. Details for these common dashlets are provided in "Dashlets" in NetWitness Content.
2. Select a dashlet.
   The new dashlet is added as the last dashlet below the existing dashlets.
3. If the dashlet is a duplicate of an existing dashlet, change the name of the new dashlet so that it is unique.

Modify or Delete a Dashlet Using Toolbar Options

Each dashlet has a toolbar that offers options for modifying the dashlet. The visualization charts have the same configuration settings, while some of the other dashlets have different additional settings.



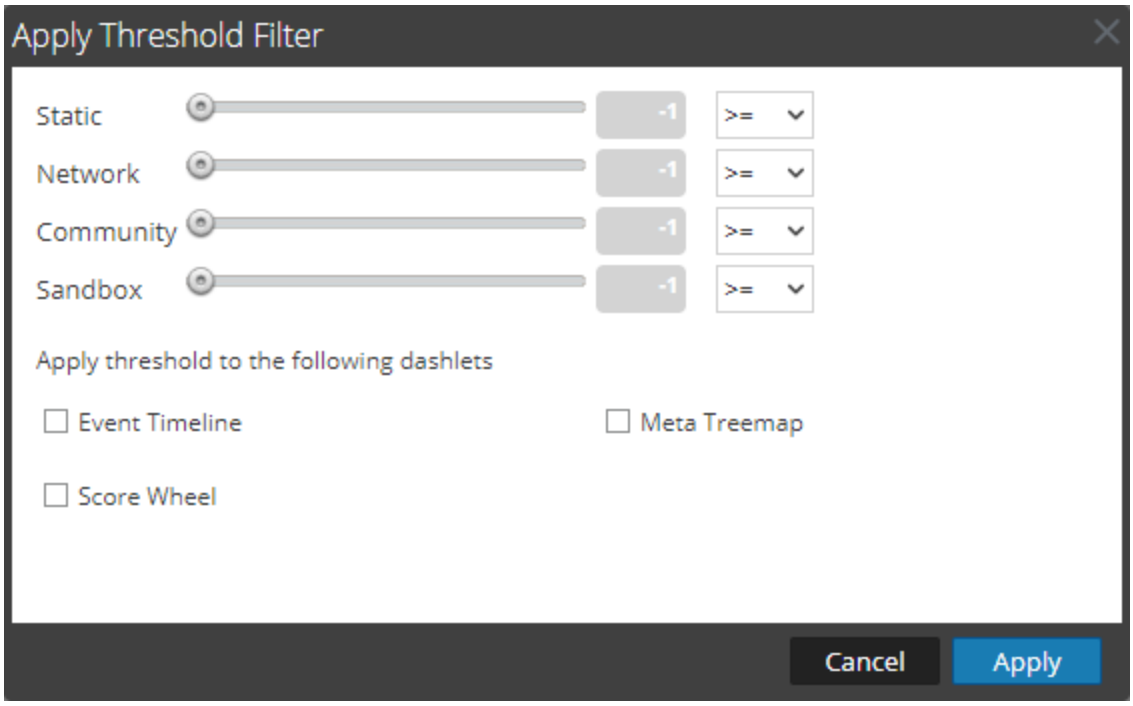
To use the toolbar options:

- To close a dashlet so that only the title bar is displayed, click .
- To open a dashlet that is closed, click .
- To display the configurable settings for a dashlet, click .
  The settings dialog for the dashlet is displayed.
- To delete a dashlet, click .

Apply Threshold Filter to Multiple Dashlets

Within dashlets, you can set a threshold to show only events equal to, above, or below a certain score in the four categories (Static, Network, Community, and Sandbox). This procedure sets the thresholds by dashlet type for these dashlets: Event Timeline, Score Wheel, and Meta Treemap. You can also set the threshold for individual dashlets.

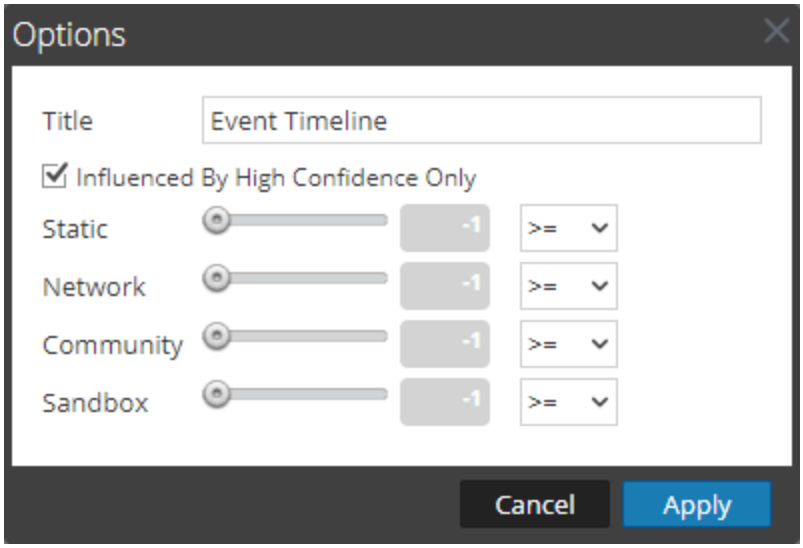
1. In the toolbar, select > **Apply Threshold Filter**.
   The Apply Threshold Filter dialog is displayed.



2. Select one or more dashlet types: Event Timeline, Score Wheel, and Meta Treemap.
3. Drag the corresponding slider or enter a numeric value, then select an operator in the drop-down list: =, >=, or <=.
4. Click **Apply**.
   The threshold filters are applied to the selected dashlet types in the Summary of Events.

Set Title and Category Options for a Dashlet

1. To display the configurable settings for a dashlet, click ⚙.
   The Options dialog for the dashlet is displayed.

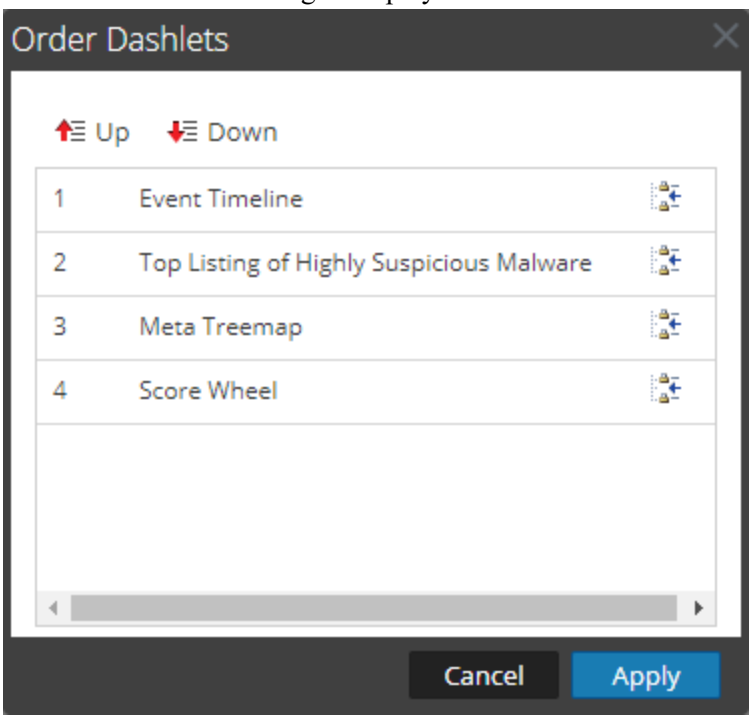


2. Type a new title for the dashlet in the **Title** field.
3. If you want to see only events that are influenced by a High Confidence tag, which means there is high confidence that the event contains harmful code, check the **Influenced By High Confidence Only** option.
4. If you want to see only events that were given a score above a certain score in the four categories (Static, Network, Community, and Sandbox), drag the corresponding slider or enter a numeric value, then select an operator in the drop-down list: =, >=, or <=.
5. Click **Apply**.
   The title and filters are applied to the dashlet.

Order Dashlets

To change the order of dashlets as they appear beneath the Summary of Events:

1. In the toolbar, select > **Order Dashlets**.
   The Order Dashlets dialog is displayed.



2. Select a dashlet that you want to move up or down, and click Up or Down.
3. When you are satisfied with the order, click **Apply**.
   The dialog closes and the order of dashlets below the Summary of Events is changed to match your choices.

Restore Default Dashlets

After you have added, modified, and arranged dashlets, you can revert to the default settings for dashlet display. To restore the default dashlets:

1. In the toolbar, select  > **Restore Default Configuration**.
   A dialog requests confirmation that you want to restore the configuration.
2. Do one of the following:
   a. If you decide to keep the dashlet arrangement you have configured, click **No**.
   b. If you are sure that you want to restore the defaults, click **Yes**,
      The dashlet display reverts to the default display.

Filter Dashlet Data in the Summary of Events View

The Summary of Events provides a summary of the scan being investigated with selectable dashlets. The Summary of Events is fixed, but Analysts can configure each dashlet to filter out information and drill into the data.

The rest of this topic provides instructions for managing and configuring dashlets.

Configure the Score Wheel Dashlet

The Score Wheel is a high-level visualization of analyzed sessions that scored high, medium, or low in each of the scoring categories: Static, Network, Community, and Sandbox. The Score Wheel is a quick way to drill into sessions to review them. Each ring represents a different scoring category so that you can visually compare results by category.

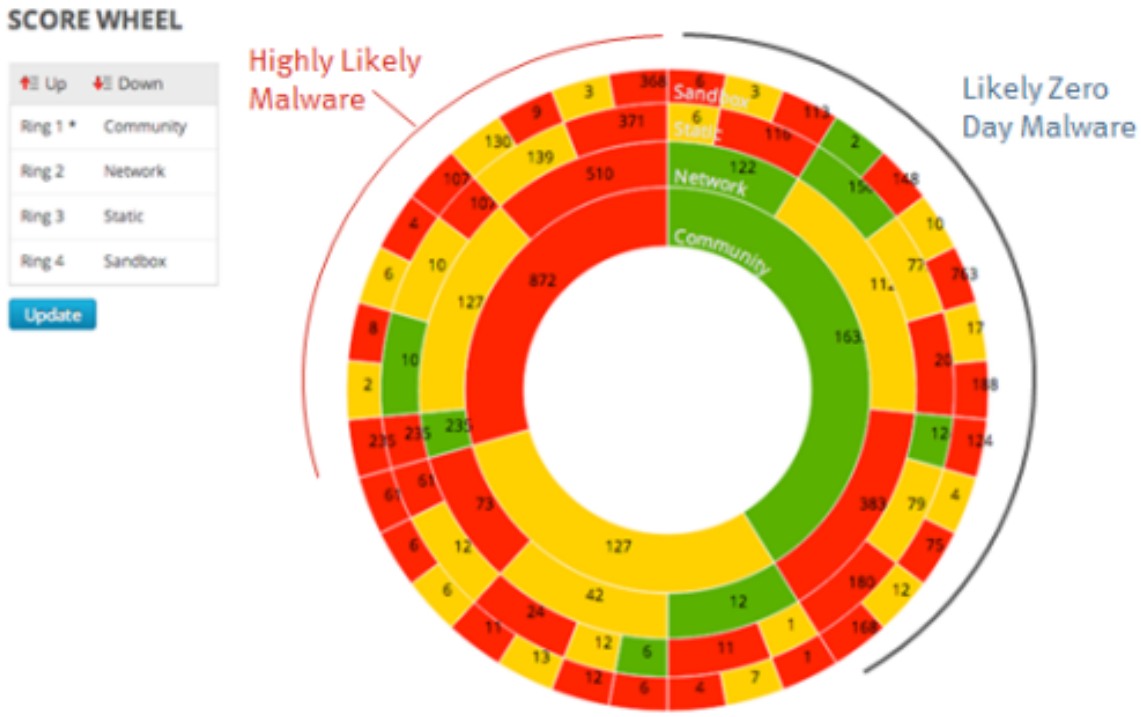


You can change the order of the rings to highlight indicators of compromise that were flagged in one category but not in another category. Comparing the same results in a different sequence of the rings provides visibility into additional vulnerabilities in a session, and you can drill into sessions of interest. The following examples show two possible use cases.

Zero-Day Candidates Example

This example shows how to drill into sessions that the Community did not flag as malicious, but all other scoring categories did. The resulting list of sessions highlights zero-day candidates.

1. Configure the Score Wheel rings in the following sequence:
   Community (innermost) > **Static** > **Network** > **Sandbox** (outermost)
2. Click the red slice in the outermost (Sandbox) ring that aligns with a green slice on the innermost ring (Community): green (innermost) -> **Static**: red -> **Network**: red -> **Sandbox**: red (outermost).



Malicious Sessions Example

This example shows how to drill into sessions in which all scoring categories identify the resulting list of sessions as malicious, indicating Malware Analysis has the most confidence that they are malware.

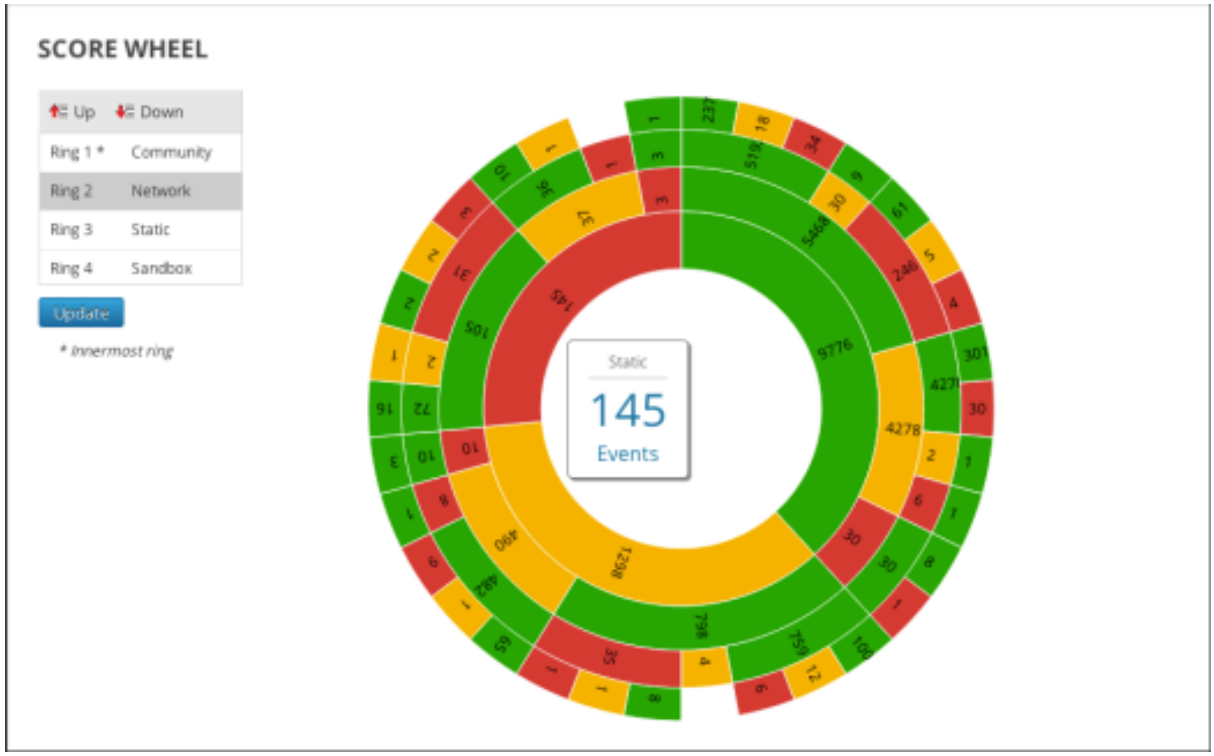
1. Configure the Score Wheel rings in the following sequence:
   Community (innermost) > **Static** > **Network** > **Sandbox** (outermost)
2. Click the red slice of the outermost (Sandbox) ring that aligns within a red slice on the innermost ring (Community): red (innermost) -> Static: red -> Network: red -> Sandbox: red (outermost).

Arrange the Ring Sequence by Scoring Module

In the Score Wheel, you can arrange the sequence of the rings by scoring module. Initially, the sequence of rings from inside to outside is Static, Network, Community, and Sandbox.

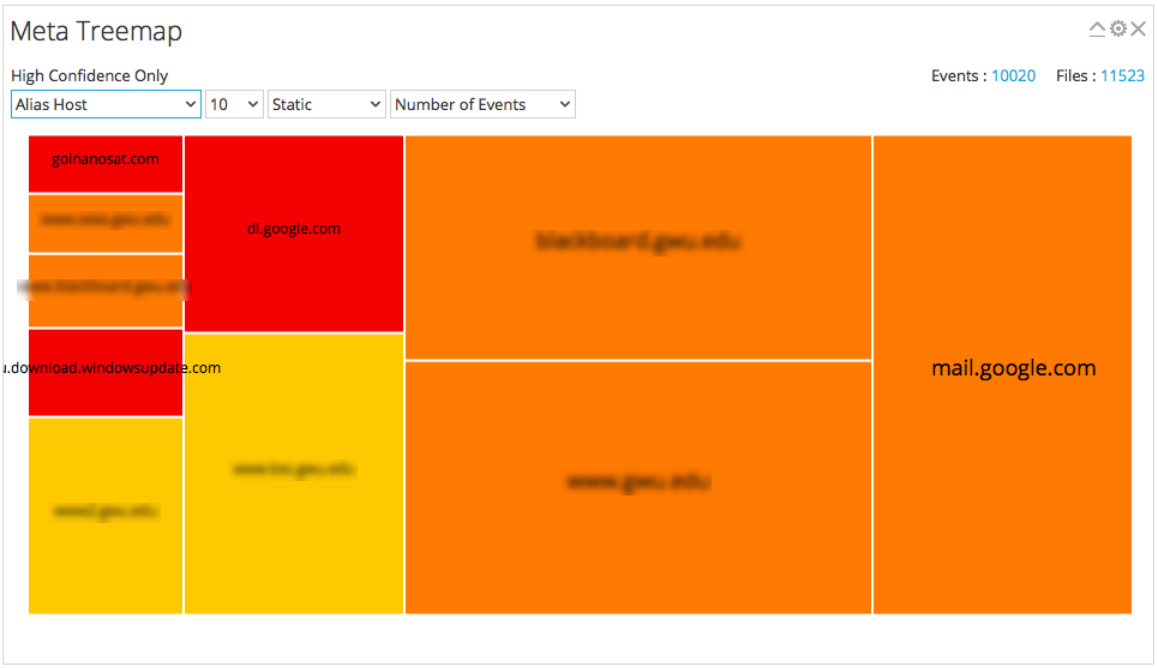
To change the ring sequence:

1. Do one of the following:
   a. Click and drag each scoring module up or down.
   b. Select each scoring module and use the Up and Down buttons to move it.
2. When the ring sequence is the way you want it, click the **Update** button.
   The Score Wheel is refreshed with the new sequence.



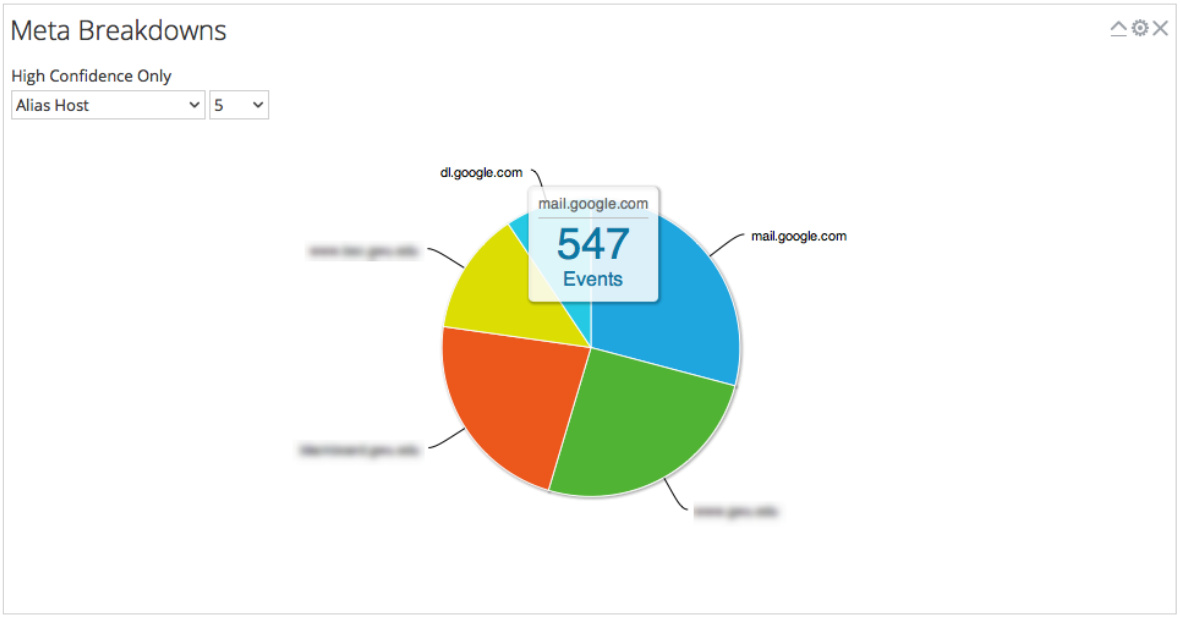
Configure the Meta Treemap Dashlet

In the Meta Treemap chart, you can visualize and filter meta breakdowns by meta type, count, and analysis type. Use the three selection lists to set the filter, and the Meta Treemap chart is refreshed immediately.



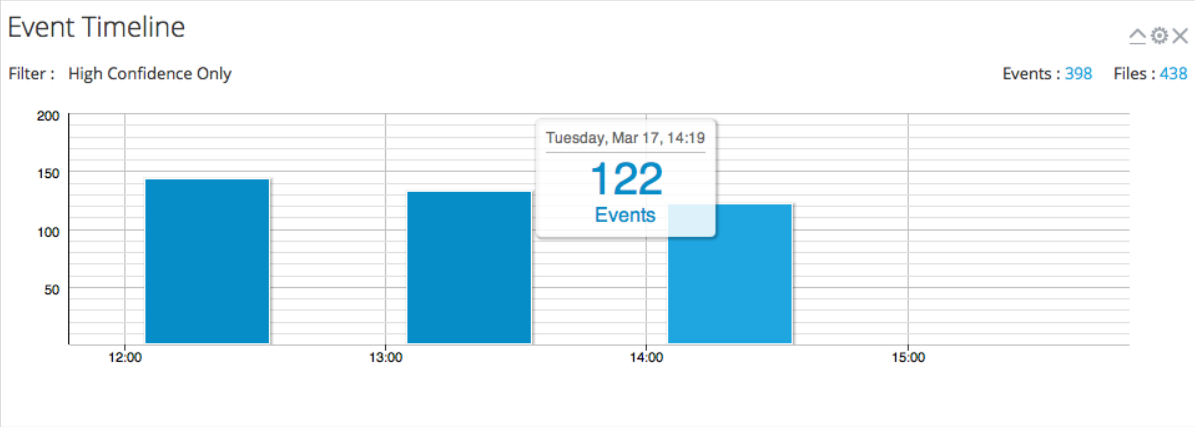
Configure the Meta Breakdowns Dashlet

The Meta Breakdowns dashlet is a visualization of values for a specific meta key in a pie chart. In the Meta Breakdowns chart, you can filter meta breakdowns by meta type and count. Use the two selection lists to set the filter, and the Meta Breakdowns chart is refreshed immediately.



Configure the Events Timeline Dashlet

The Events Timeline dashlet is a visualization of the events along a timeline. No additional filters are available for the Event Timeline.

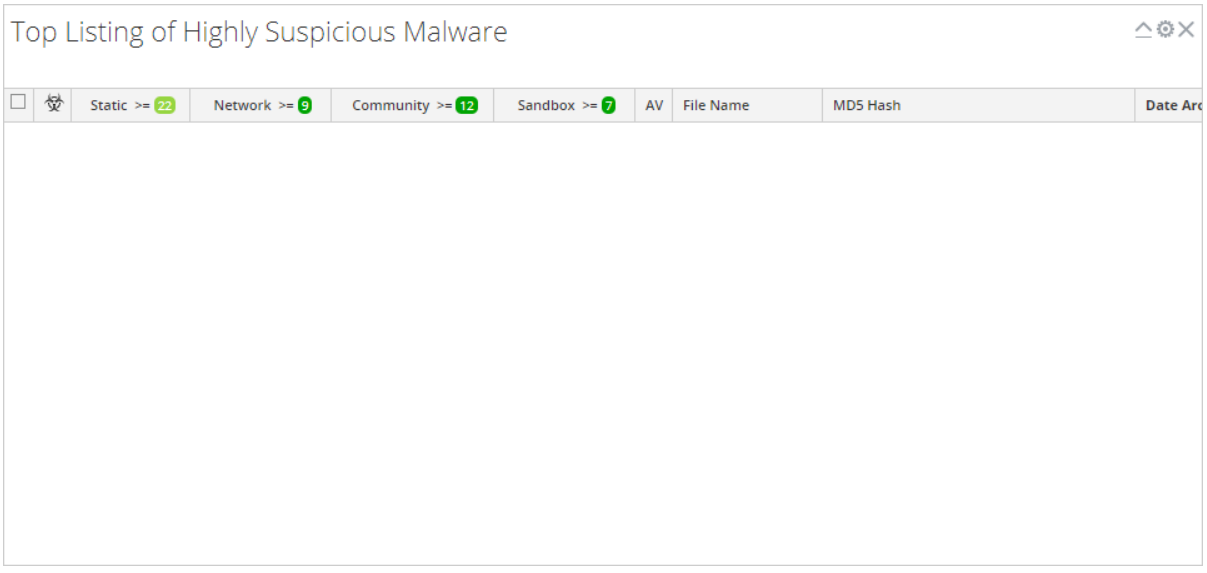


Open All Events in the Events List

From within the Event Timeline, you can open the entire list of events in the Events List. To do so, click View Events. This option is not the same as clicking the count next to Events, which is the same for all visualization charts and opens the current drill point in the Events List.

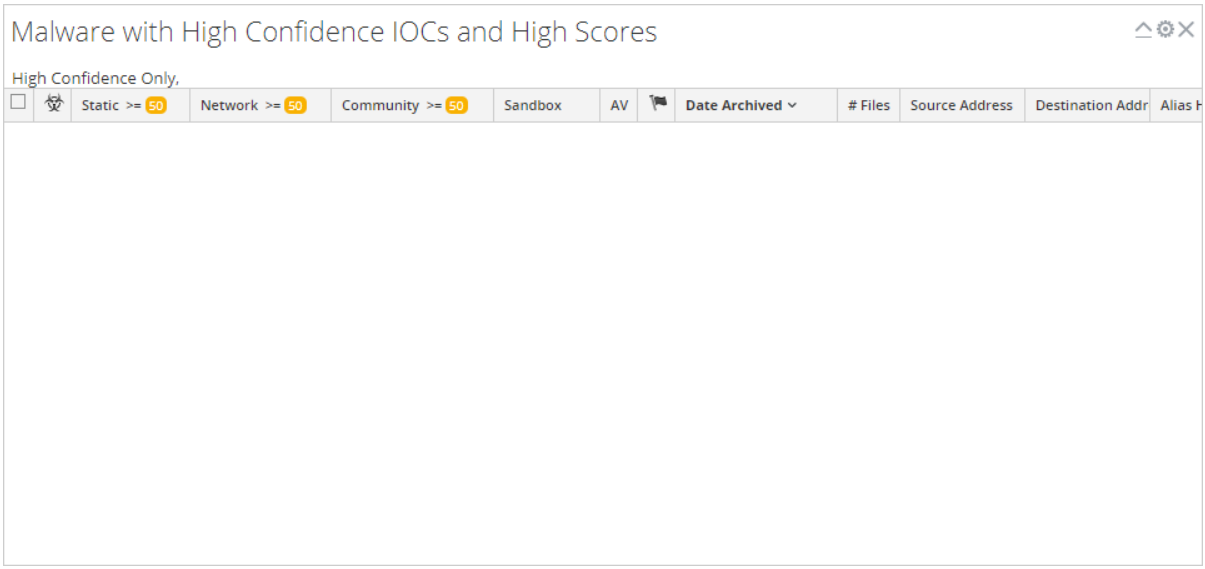
Configure the Top Listing of Highly Suspicious Malware Dashlet

The Top Listing of Highly Suspicious Malware Dashlet presents the Top 10 most suspicious events in the Events List or the Files List. This dashlet is also available in the Monitor dashboard, and the configuration options are described as part of the NetWitness Content in Dashlets.



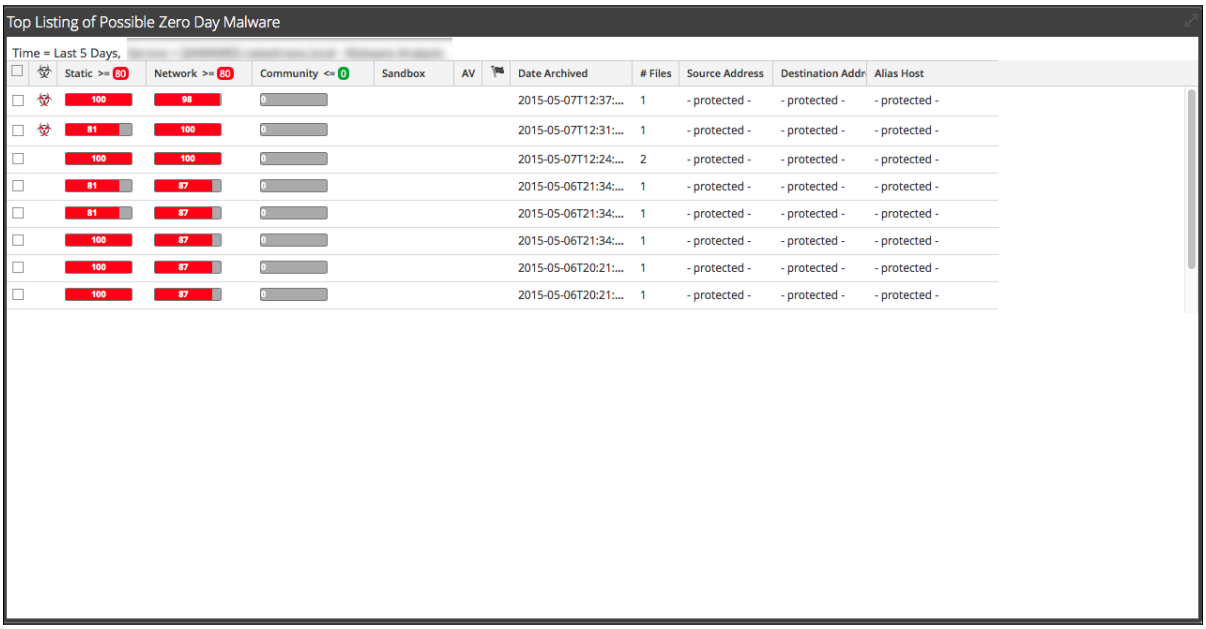
Configure the Malware with High Confidence IOCs and High Scores Dashlet

The Malware with High Confidence IOCs and High Scores dashlet presents Indicators of Compromise that have both high scores and high confidence that the events are likely to contain malware. The dashlet is also available in the Unified dashboard, and the configuration options are described as part of the NetWitness Content in Dashlets.



Configure the Top Listing of Possible Zero Day Malware Dashlet

The Top Listing of Possible Zero Day Malware dashlet presents potential zero day events in the Events List or the Files List. The dashlet is also available in the Unified dashboard, and the configuration options are described as part of the NetWitness Content in Dashlets.


Top Listing of Possible Zero Day Malware
Time = Last 5 Days,
Static >= 80
Network >= 80
Community <= 0
Sandbox
AV
Date Archived
Files
Source Address
Destination Addr
Alias Host
100 98 0 2015-05-07T12:37:... 1 - protected - - protected - - protected -
81 100 0 2015-05-07T12:31:... 1 - protected - - protected - - protected -
100 100 0 2015-05-07T12:24:... 2 - protected - - protected - - protected -
81 87 0 2015-05-06T21:34:... 1 - protected - - protected - - protected -
81 87 0 2015-05-06T21:34:... 1 - protected - - protected - - protected -
100 87 0 2015-05-06T21:34:... 1 - protected - - protected - - protected -
100 87 0 2015-05-06T20:21:... 1 - protected - - protected - - protected -
100 87 0 2015-05-06T20:21:... 1 - protected - - protected - - protected -

Upload Files for Malware Analysis Scanning

There are two methods for analysts to upload files for Malware Analysis scanning.

A Malware Analyst with permission to Initiate Malware Analysis Scan can upload files to scan using the Scan Files option in the Select a Malware Analysis Service dialog.

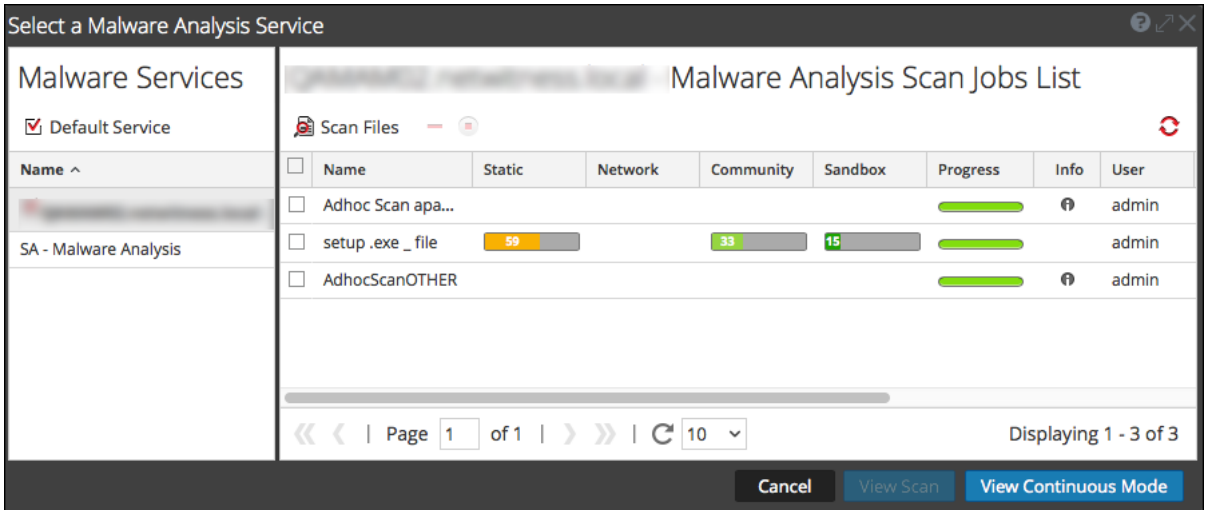
It is also possible to upload a file for scanning using a watched file share.

Upload Files Manually

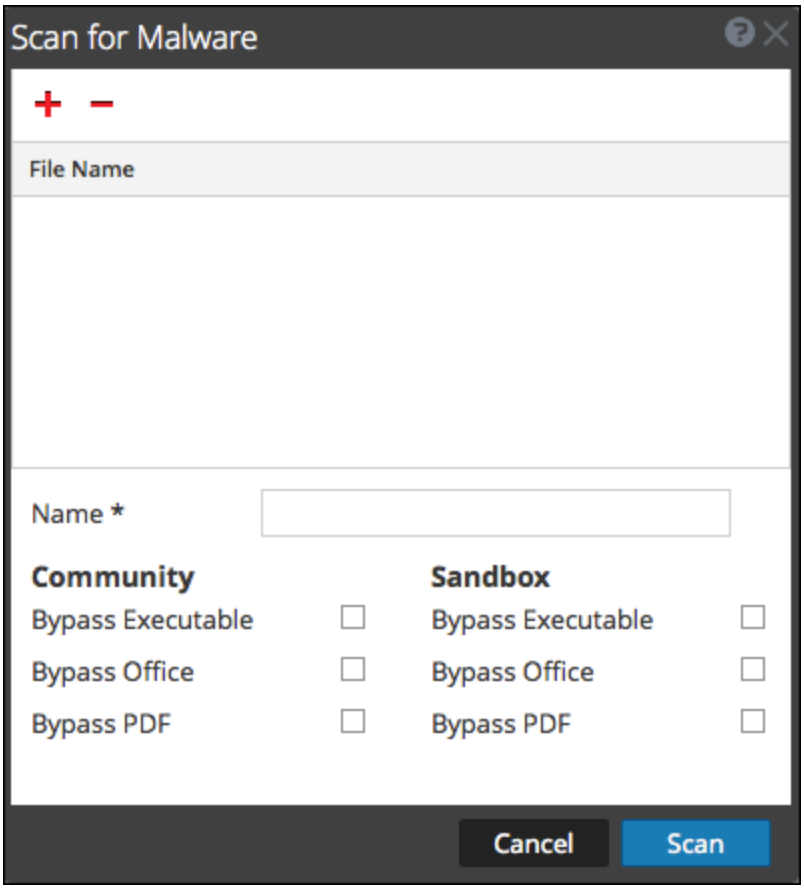
This topic provides instructions for initiating on-demand scanning of an uploaded file. When you upload a file for scanning, NetWitness starts the upload job and adds it to the jobs queue. When the job is complete, you can view the scan in Malware Analysis.

To upload a file to scan:

1. Go to **Investigate > Malware Analysis**.
   The Select a Malware Analysis Service dialog is displayed, with available Malware Analysis hosts and services for the current user in the left panel.



2. Click **View Scan**.
   The Scan for Malware dialog is displayed.



3. Click +
   A view of the files system is displayed so that you can choose files to upload.
4. Select one or more files from the list and click **Open**.
   The file names are added. Malware Analysis escapes the filename characters before processing a file. The maximum number of filename characters after escaping is 200. If the filename is greater than 200 characters, Malware Analysis truncates the filename characters and displays the truncated filename in the NetWitness user interface.
5. Continue adding and deleting files until you have a list of the files that you want to upload.
6. Name the scan and select the types of files to bypass. This is useful for a zip archive that contains different types of files, and overrides the default bypass settings.
7. Click **Scan**.
   The scan job is submitted and NetWitness displays a confirmation message for successful submission. The scan request is added to the Scan Jobs List dashlet. The bypass settings in this dialog override the default settings in the basic Malware Analysis configuration settings.
8. The job is added to the Scan Jobs List in the Select a Malware Analysis Service dialog and in the Unified dashboard Scan Jobs List dashlet.
9. To view the scan when complete, double-click the scan.
   The Malware Summary of Events for the selected scan is displayed.

Upload Files from a Watched Folder

To upload files from a watched folder, you can drop files into a watched file share for Malware Analysis. Analysts can share YARA rules, hash files, and infected zip archives with Malware Analysis.

Malware Analysis watches a file share and automatically consumes files placed in specific folders in the file share. This feature is useful for:

- Bulk import of hash files from `/var/lib/rsamalware/spectrum/hashWatch`.
- Addition of custom-YARA rules to the Indicators of Compromise (IOC) list on the host from `/var/lib/netwitness/malware-analytics-server/spectrum/yara/watch`.
- Creation of on-demand scan jobs from a zip archive of infected zip files from `/var/lib/rsamalware/spectrum/infectedZipWatch/watch`.

Analysts need to prepare the files for consumption in accordance with requirements, the file extension must be correct, and the file must be copied to the correct watched folder in the file share.

Import a Hash List

To import a hash list from the watched directory, the hash list must be in the specified format and must be sorted on md5. You can drop a file formatted into a folder (`/var/lib/rsamalware/spectrum/hashWatch`) on the Malware Analysis host, and it is automatically imported into the local hash database. This is described in "Configure Hash Filter" in the *Malware Analysis Configuration Guide.*

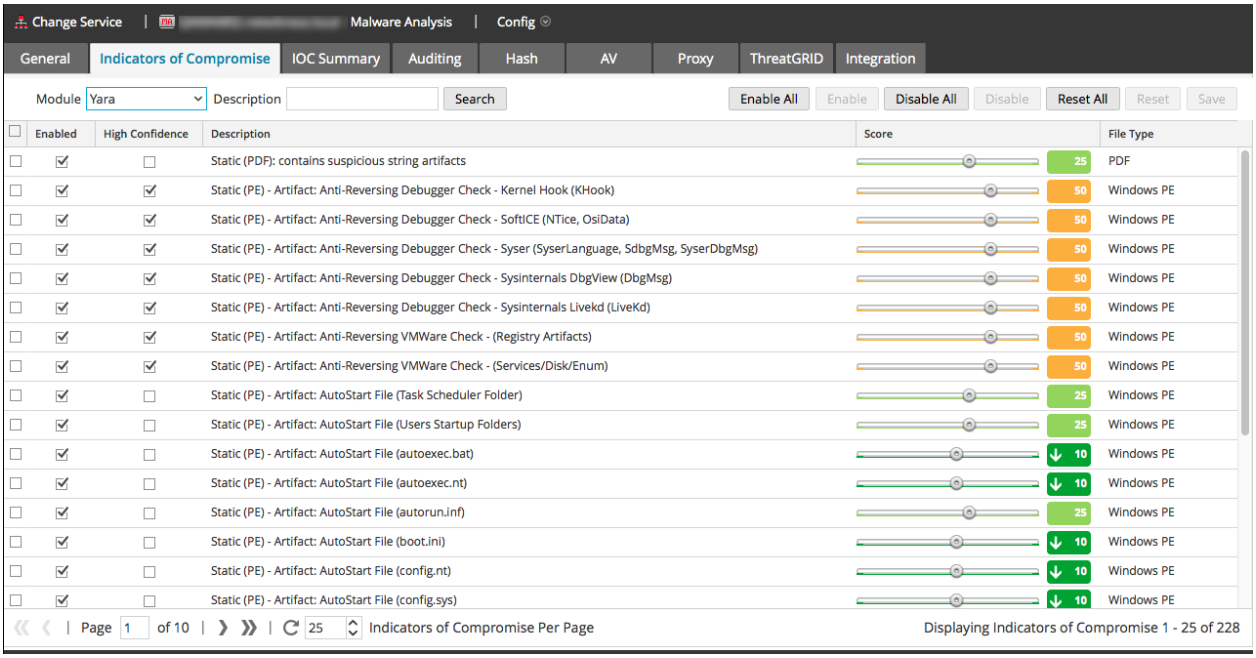
To import a hash list using the watched folder method:

1. Copy the hash lists that you want to import into the **`/var/lib/rsamalware/spectrum/hashWatch`** directory.
   NetWitness Malware Analysis automatically watches this folder and processes files placed there.
   a. Malware Analysis adds every hash found in the hash lists to the hash filter.
   b. If there are processing errors, they are logged in: `/var/lib/rsamalware/spectrum/hashWatch/error`
   c. Processed files are cataloged here: `/var/lib/rsamalware/spectrum/hashWatch/processed`
   d. Processed files are not removed from the hashWatch directory.
2. After importing hashes in bulk, the System Administrator can use a cronjob to clean up old processed files.

Import YARA rules to the IOC List

Customers with advanced skills and knowledge can add detection capabilities to Malware Analysis by authoring YARA rules and publishing them in Live or placing YARA rules in a watched folder for the host to consume. Implement Custom YARA Content provides complete information on the prerequisites for using custom YARA content and authoring rules.

When the rules are ready, place the custom YARA files in the folder that the Malware Analysis service watches:
`/var/lib/netwitness/malware-analytics-server/spectrum/yara/watch`
The file is consumed within one minute.
Once consumed, NetWitness moves the file to the `processed` folder, and the new rule is added to the Malware Analysis Service Config view > Indicators of Compromise tab.



Import Files into the Scan Jobs List

When you obtain samples from perimeter security solutions and would like to perform further analysis on the files, you can zip the files and password protect the archive with `infected`, then add to the watched folder for consumption by Malware Analysis. This zipped archive is ready to be placed in the watched folder: `/var/lib/rsamalware/spectrum/infectedZipWatch/watch`.

Note: The maximum size of the archive is 100 MB.

To analyze infected, password-protected zip files, Malware Analysis consumes archives place in a watched folder and creates an on-demand job that is added to the Scan Jobs List.

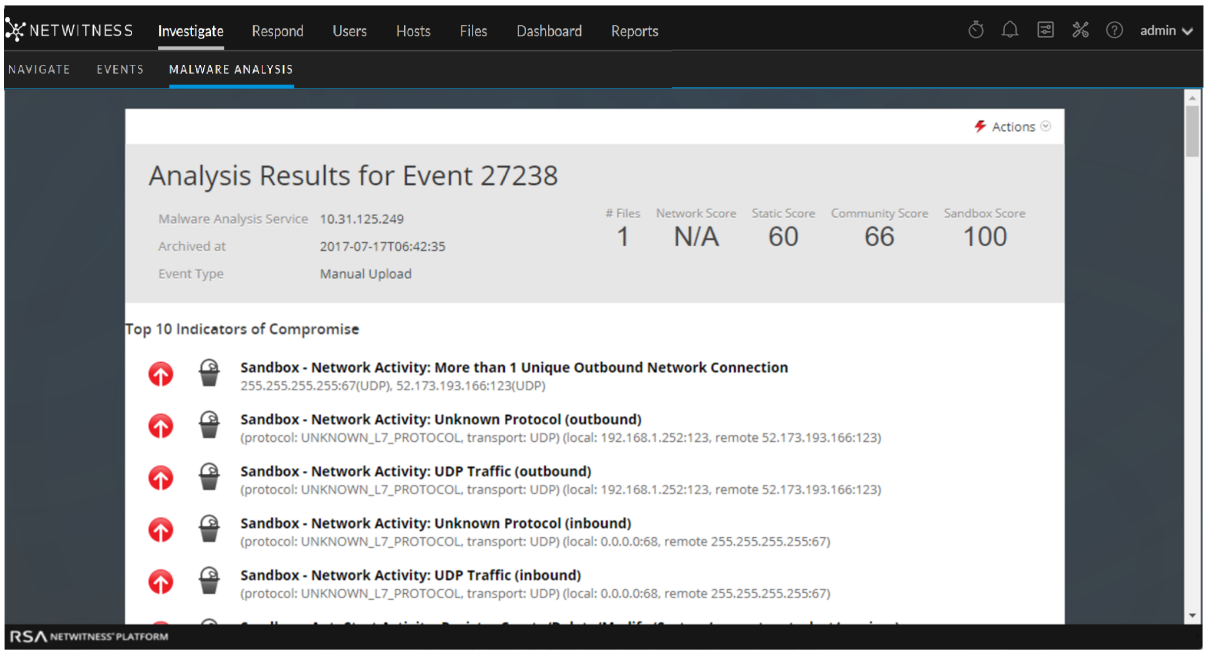
1. While logged on as administrator, place the files to be processed in a zip file with password `infected` at `/var/lib/rsamalware/spectrum/infectedZipWatch/watch`
   In a minute or two Malware Analysis consumes the archive and creates an on-demand job in the Scan Jobs List. The scan job name is the name of the file, the user is **file share**, and the Event Type is 1. The archive is moved to `/var/lib/rsamalware/spectrum/infectedZipWatch/processed`
2. After the job is added to the Scan Job List, run a script or cronjob to clean up the zip file in `/var/lib/rsamalware/spectrum/infectedZipWatch/processed`.

View Detailed Malware Analysis of an Event

When viewing the list of individual events in a Malware Analysis scan in the Malware Analysis Events grid, you can double-click an event to view the detailed analysis results for the event.

View Malware Analysis Details for an Event

1. Start an investigation in the **Malware Analysis** view.
   The Malware Summary of Events is displayed, and includes four charts, including the Event Timeline.
2. Do one of the following:
   a. To view all events in the Event Timeline, click the **View Events** button.
   b. Double-click data in the **Meta Breakdown**, **Meta Treemap Chart**, or **Score Wheel**.
   The Events List is displayed.
3. Double-click an event.
   The Analysis Results for the event are displayed.

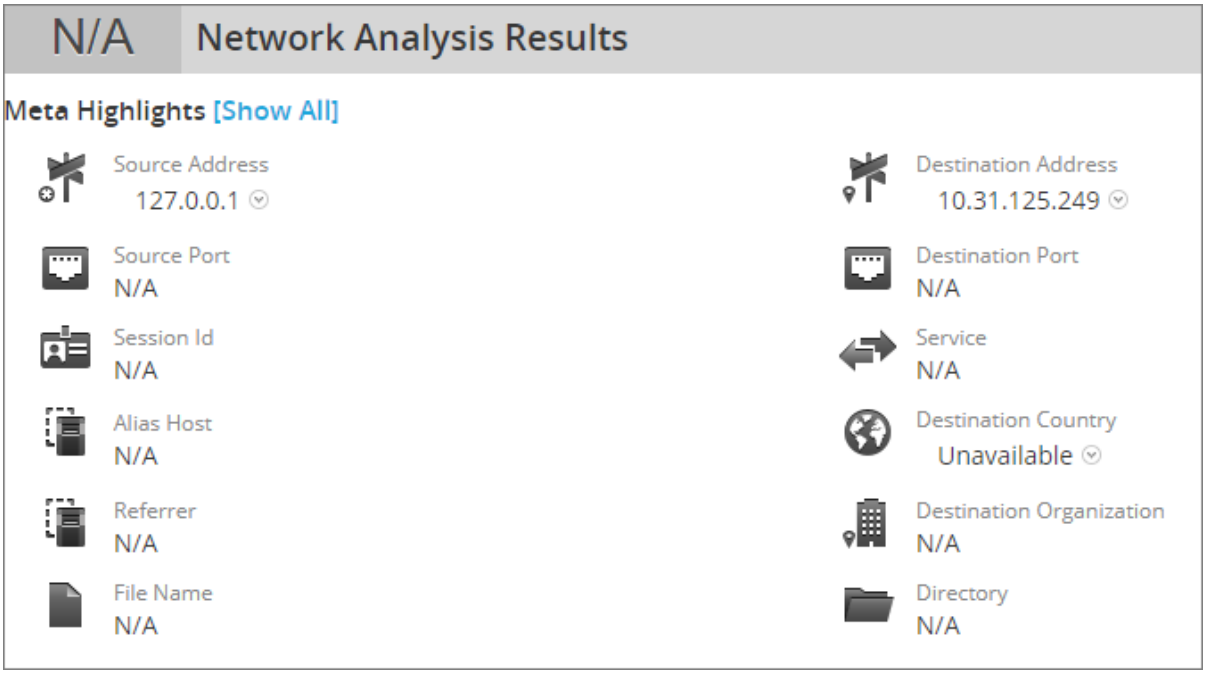


4. (Optional) If you want to delete an event, select **Actions > Delete Event**.
5. If you want to view a reconstruction of the network session, select **Actions > View Network Session**.
   The session opens in the Navigate view > Event Reconstruction.

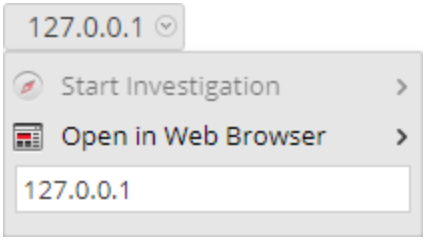
Pivot Network Analysis Results

You can pivot the Network Analysis Results in several ways:

1. Scroll down to the Network Analysis Results.



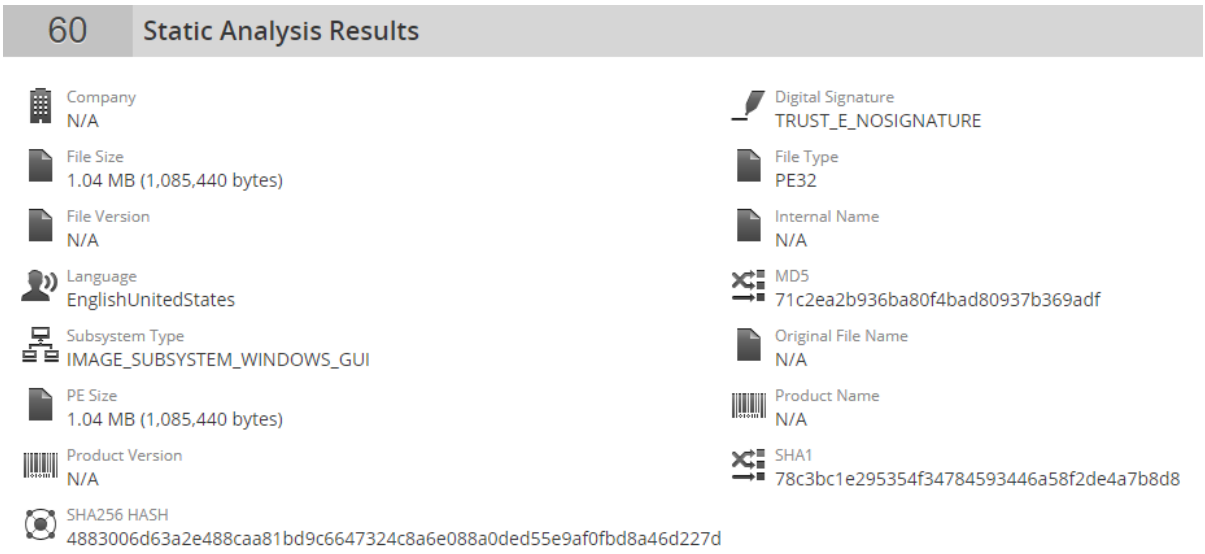
2. Hover over a meta value and left-click.
   The context menu is displayed.



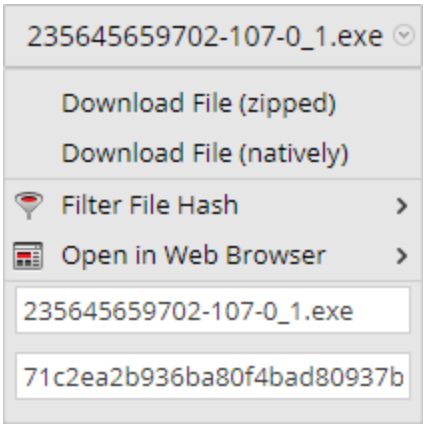
3. To view the selected meta value in the **Navigate** view, select **Start Investigation** and a time option.
4. To view the selected meta value in a browser, select **Open in Web Browser > Open in Google**.

Use File Actions in the Static Analysis Results

1. Scroll down to the Static Analysis Results.



2. If you want to download a file, select the file name and either **Download File (zipped)** or **Download File (natively)** in the drop-down menu. It is safer to download a file in zipped format.



3. If you want to mark the file as safe or unsafe in the hash list, select **Filter File Hash** and **Mark hash as good** or **Mark hash as bad**.

View Community Analysis Results Details

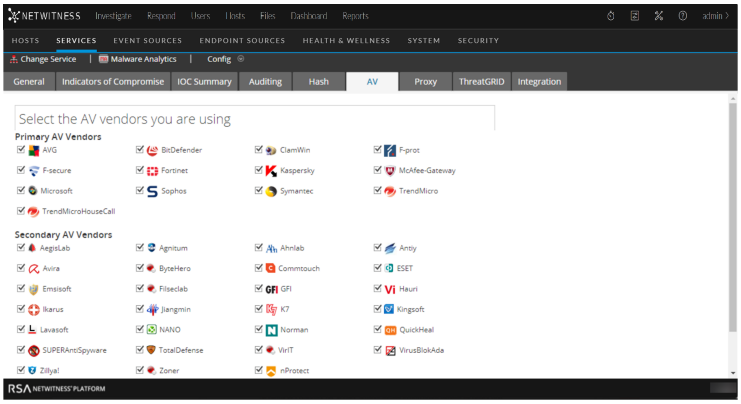
The Community Analysis Results summarizes results from the community, identifying Indicators of Compromise that were flagged as a risk or identified as good.

In addition, this view lists the results from Installed AV Vendors and Not Installed AV Vendors. You can compare results of the installed AV vendors that were configured for the current Malware Analysis service versus Community results. You can also see results from a list of AV vendors that are not configured as installed for the current Malware Analysis service.

Each row of AV vendor results includes the shield icon to show whether the IOC was discovered by a Primary ( ) or Secondary AV ( ) vendor in the community, the name of the Installed or Not Installed vendor, and the name of malware or risk detected by the community and AV vendor. If the AV vendor did not detect a risk, **-- Not detected --** is displayed instead of the name of the risk.

The Not Installed AV Vendors section is expandable to view all entries, but is collapsed by default to minimize the need to scroll. Clicking the + expands the list.

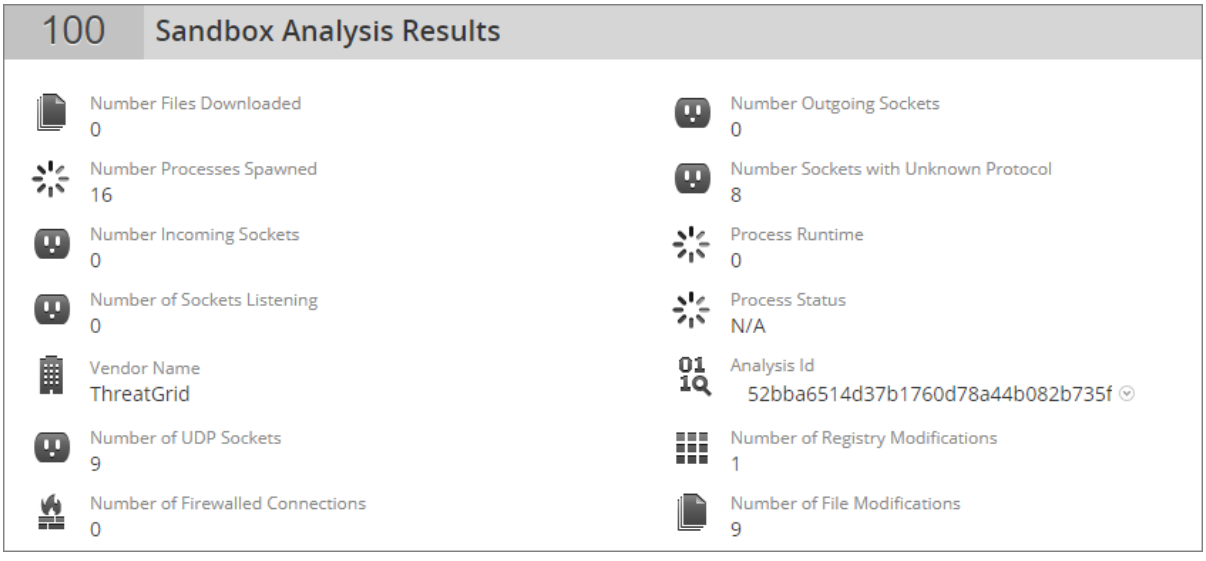
If no installed AV vendors have been configured for the current Malware Analysis service, the following message is displayed: No AV vendors were marked as installed. Please go to the Malware Analysis Service configuration page to identify installed AV vendors.



View Sandbox Analysis Results in the ThreatGRID User Interface

If you have registered with ThreatGRID, you can view the Sandbox results directly in ThreatGRID.

1. Scroll down to the Sandbox Analysis Results.



2. Click the **Analysis ID** and select **Open In ThreatGRID**.
   The analysis report in ThreatGRID is displayed.

Malware Analysis Reference Materials

This section provides is intended to help you understand the purpose and application of Malware Analysis views in NetWitness Investigate. For each view, there is a brief introduction and a What Do You Want To Do table with links to related procedures. In addition some of the reference materials include workflows and Quick Looks to highlight important features in the user interface.

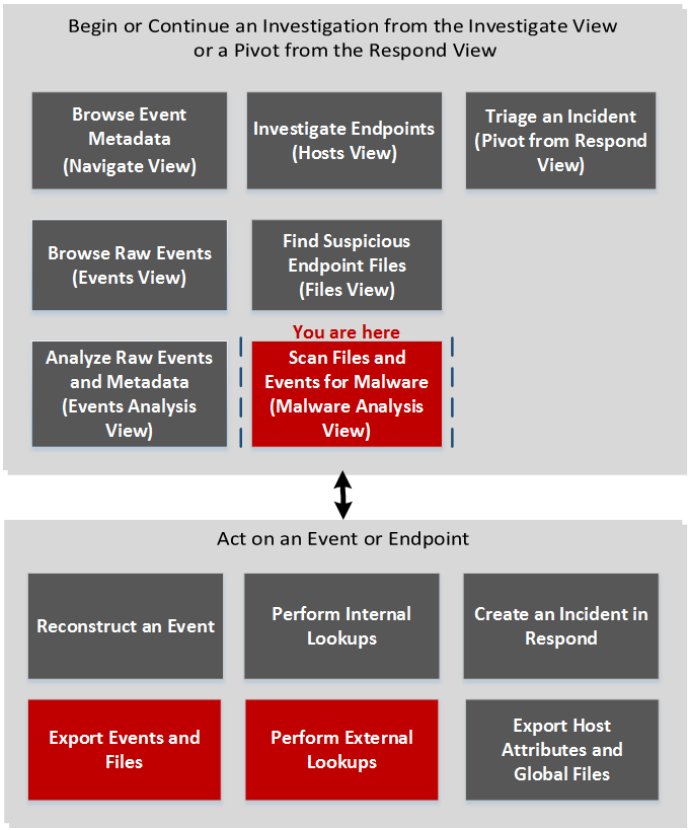
- Malware Analysis View
- Malware Analysis Events List and Files List
- Scan For Malware Dialog
- Select a Malware Analysis Service Dialog

Malware Analysis View

In NetWitness Investigate, the Malware Analysis view provides the user interface for conducting a malware analysis. The Malware Analysis view is in the form of a customizable dashboard, in which default dashlets in the initial view are based on the user role (Administration or Analyst) and user customizations. Initially, the Summary of Events dashlet is displayed in the Malware Analysis view. Additional dashlets present different visualizations of the events being viewed, and each representation is configurable to further refine your view as you search for Indicators of Compromise. The Malware Analysis dashlets available in the Dashboard are also available in the Malware view.

To access this view, select **Investigate > Malware Analysis**. If a default service has not been selected, the Select a Malware Analysis Service dialog is displayed. Select a service, then click **View Continuous Mode**.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	browse event metadata	*NetWitness Investigate User Guide*

User Role	I want to ...	Show me how
Threat Hunter	browse raw events	*NetWitness Investigate User Guide*
Threat Hunter	analyze raw events and metadata	*NetWitness Investigate User Guide*
Threat Hunter	investigate endpoints (Version 11.1)	*NetWitness Investigate User Guide*
Threat Hunter	find suspicious endpoint files (Version 11.1)	*NetWitness Investigate User Guide*
Threat Hunter	scan files and events for malware*	Conducting Malware Analysis
Incident Responder	triage an incident in Investigate	*NetWitness Investigate User Guide*
Threat Hunter	export events and Files*	Examine Scan Files and Events in List Form
Threat Hunter	perform external lookups*	View Detailed Malware Analysis of an Event

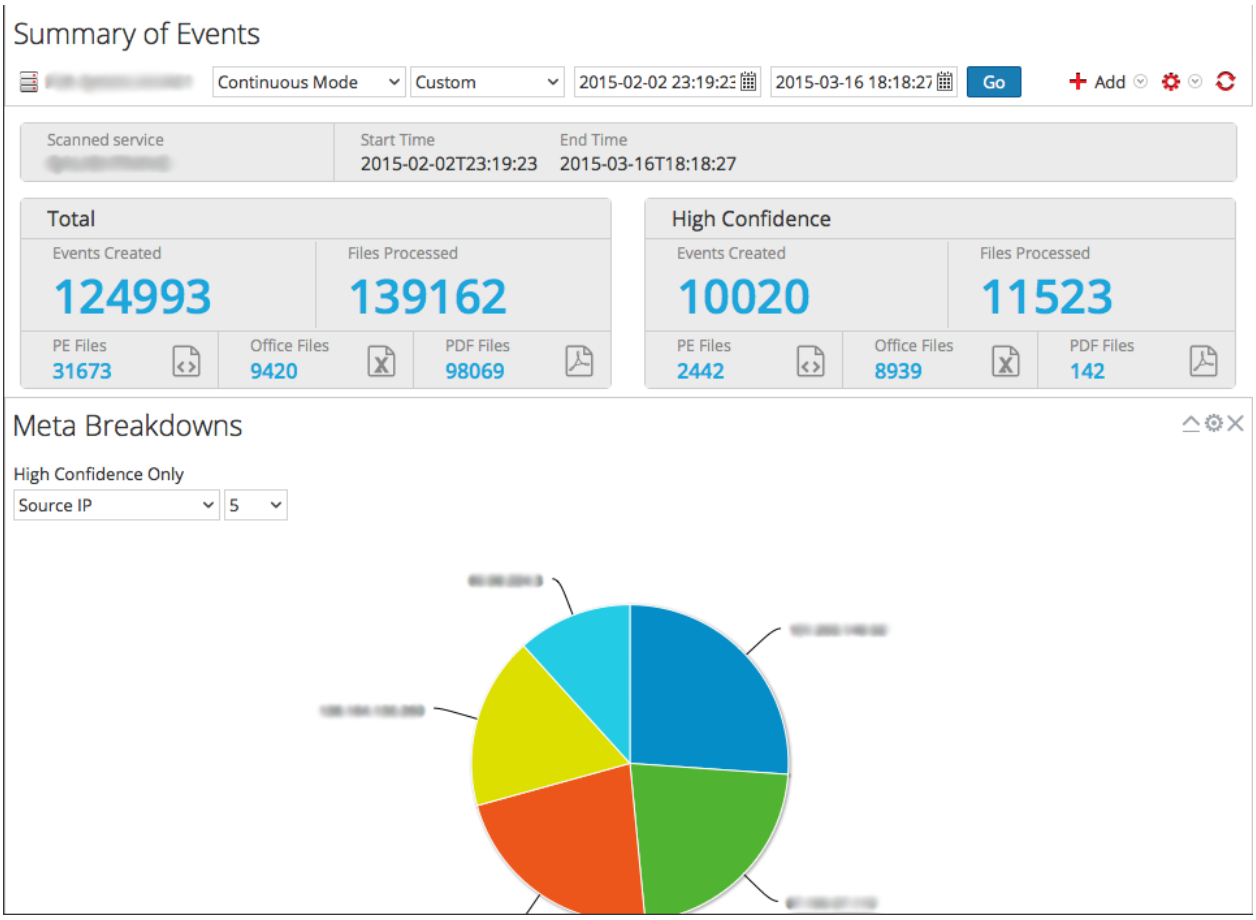
*You can perform this task in the current view.

Related Topics

- "How NetWitness Investigate Works" in the *NetWitness Investigate User Guide*
- "Launch a Malware Analysis Scan from the Navigate View" in the *NetWitness Investigate User Guide*

Quick Look

Below is an example of the Malware Analysis view.



The Malware Analysis view consists of the Summary of Events panel and four dashlets unique to this view. Each of the unique dashlets have identical Options dialogs. The Malware Analysis dashlets in the MONITOR view are also available, and are described in the Dashlets topic in the NetWitness Content space.

Summary of Events Panel

In the Summary of Events panel, you can select the service, the scan mode, and the time range. In addition, you can select a data point and view the events associated with the event.

The following table describes all features in the Summary of Events panel.

Feature	Description
	Selects a service to display.
Scan Mode	Displays a drop-down list of available scan modes.
Time Range	Displays a drop-down list of time ranges to view events.
Start Date	When Time Range is set to custom, offers a calendar from which to choose the start date of the time range.

Feature	Description
End Date	When Time Range is set to custom, offers a calendar from which to choose the end date of the time range.
	Displays a drop-down list of dashlets you can add to the view.
	Displays a drop-down list of actions you can perform in this view: • Restore Default Configuration • Order Dashlets • Apply Threshold Filter
	Refreshes the Malware Analysis view.

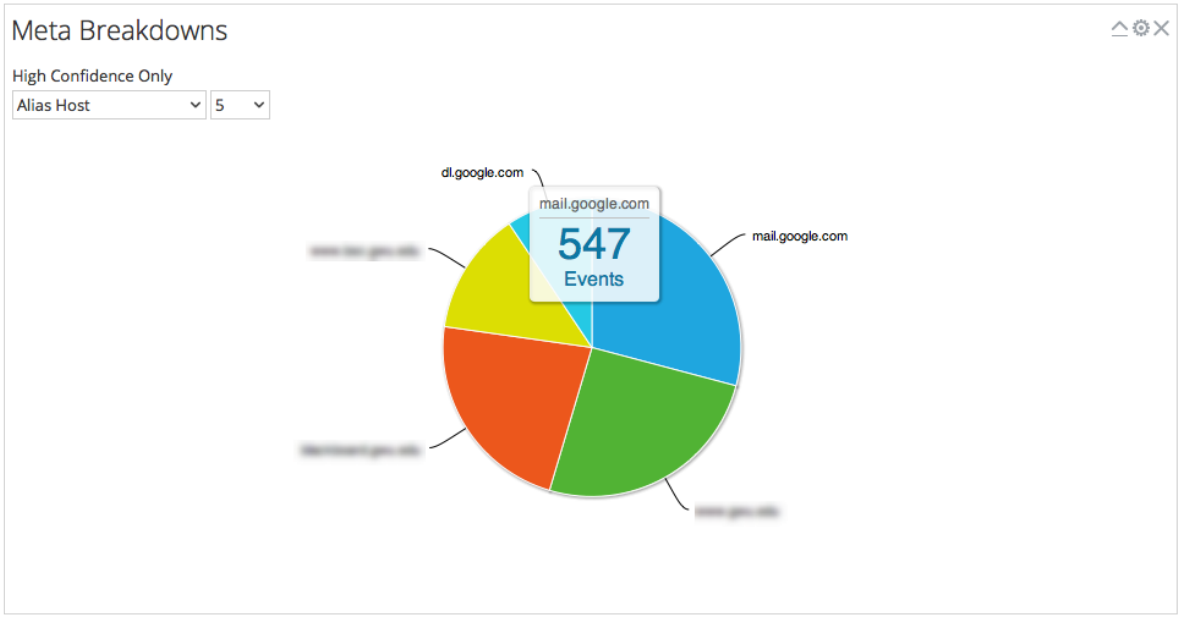
Options Dialog

In the Options dialog, you can customize the results displayed in the dashlet. This dialog can be accessed by clicking the  icon in the top right corner of each dashlet. The following table describes the features of the Options dialog.

Feature	Description
Title	Indicates whether the data shown is restricted to events flagged as high confidence or not. If the data is not restricted, this line will not be displayed.
Influenced By High Confidence Only	Indicates whether the data shown is restricted to events flagged as high confidence.
Static, Network, Community, Sandbox	Allows you to filter results based on the scores in the scoring modules.
Cancel	Closes the dialog without saving any changes.
Apply	Applies changes to the dashlet immediately and closes the dialog.

Meta Breakdowns

Meta Breakdowns presents events in the form of a pie chart, with each slice representing a meta value for the specified meta key. You can select the meta key and the count of meta values for that key to render in the chart, starting with the meta value having the most events. Hovering over an event displays the count.

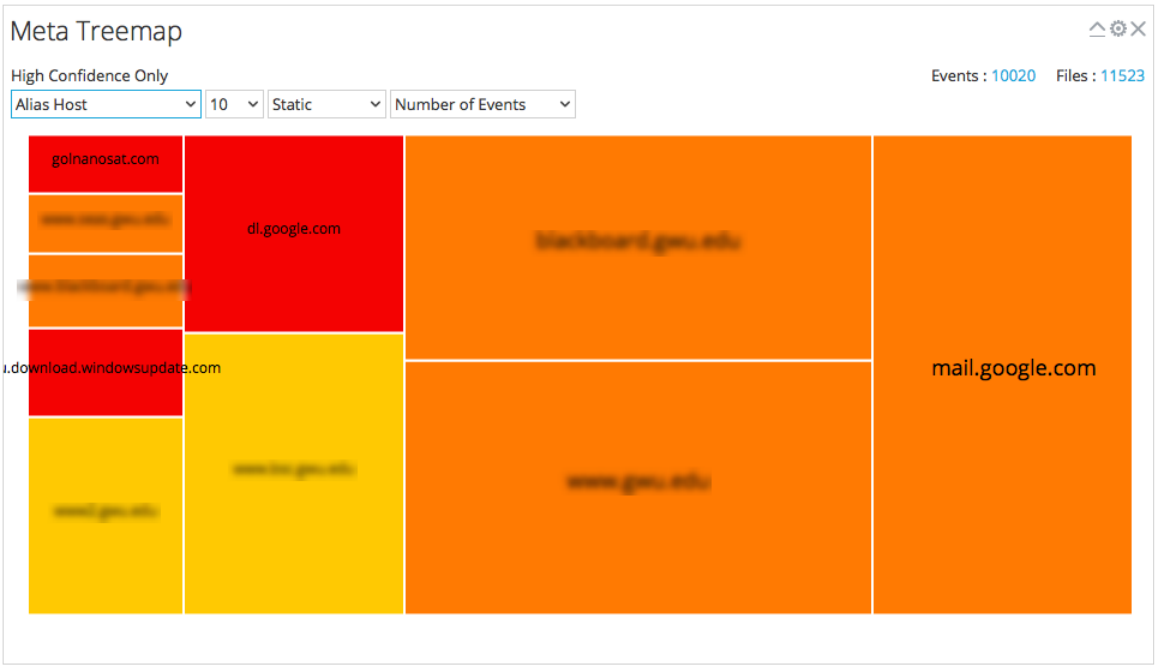


The following table describes the options in the Meta Breakdowns dashlet.

Feature	Description
High Confidence Only	Indicates whether the data shown is restricted to events flagged as high confidence or not. If the data is not restricted, this line will not be displayed.
Meta Key	Drop-down list of available meta keys.
Count	Drop-down list specifying how many of the top results are displayed.

Meta Treemap

Meta Treemap presents events in the form of a heat map. You can select the meta key and the count of meta values for that key to render in the chart, starting with the meta values having the most events. In addition, you can select the module that detected the meta value in the events: static, network community, or sandbox.



The following table describes the options in the Meta Treemap dashlet.

Feature	Description
High Confidence Only	Indicates whether or not the results are restricted to events flagged as high confidence or not. If the results are not restricted, this line will not be displayed.
Meta Key	Drop-down list of available meta keys to select as a filter.
Count	Drop-down list specifying how many of the top results are displayed.
Module	Drop-down list specifying which module results will be pulled from.
Value	Drop-down list specifying what information will be displayed when the mouse is hovering over a result (for example, Average Score).

Score Wheel

The Score Wheel offers a view of events as concentric rings with colors representing scores for events based on Indicators of Compromise and the scoring module. You can arrange the position of the rings using the Up and Down arrows to obtain a view that highlights events that were detected by one scoring module (red) and not detected by other scoring modules.

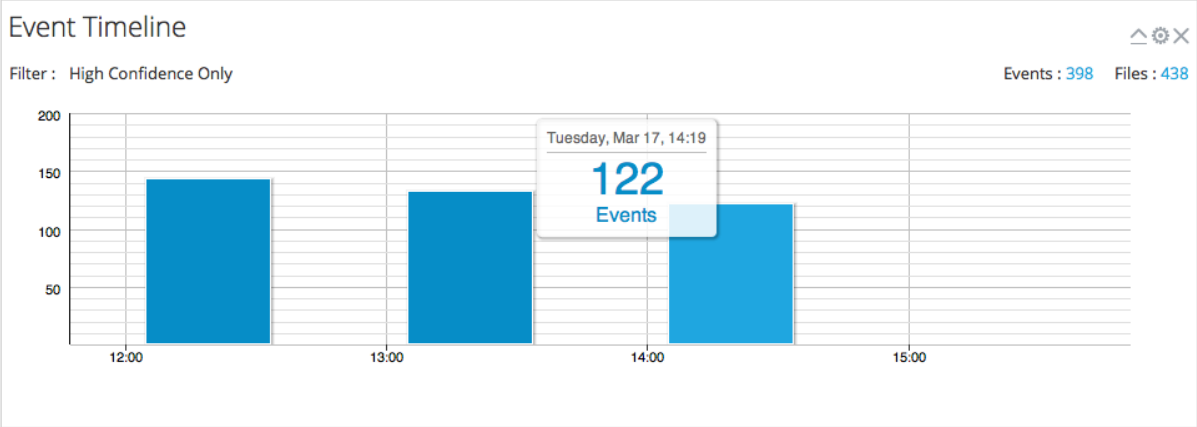


The following table describes the features of the Score Wheel dashlet.

Feature	Description
High Confidence Only	Indicates whether or not the results are restricted to events flagged as high confidence or not. If the results are not restricted, this line will not be displayed.
Module Order grid	Displays the order of the rings in Score Wheel, Ring 1 being the innermost ring and Ring 4 being the outermost ring. You can click the **Up** and **Down** buttons to reorder the modules, then click **Update** to apply the changes.

Event Timeline

The Event Timeline offers a view of events organized by the time of occurrence in a bar graph. Clicking and dragging to select a time range within the chart zooms in on the selected time.



The following table describes the features of the Event Timeline dashlet.

Feature	Description
High Confidence Only	Indicates whether or not the results are restricted to events flagged as high confidence or not. If the results are not restricted, this line will not be displayed.
View Events	Displays the Investigate > Events view.

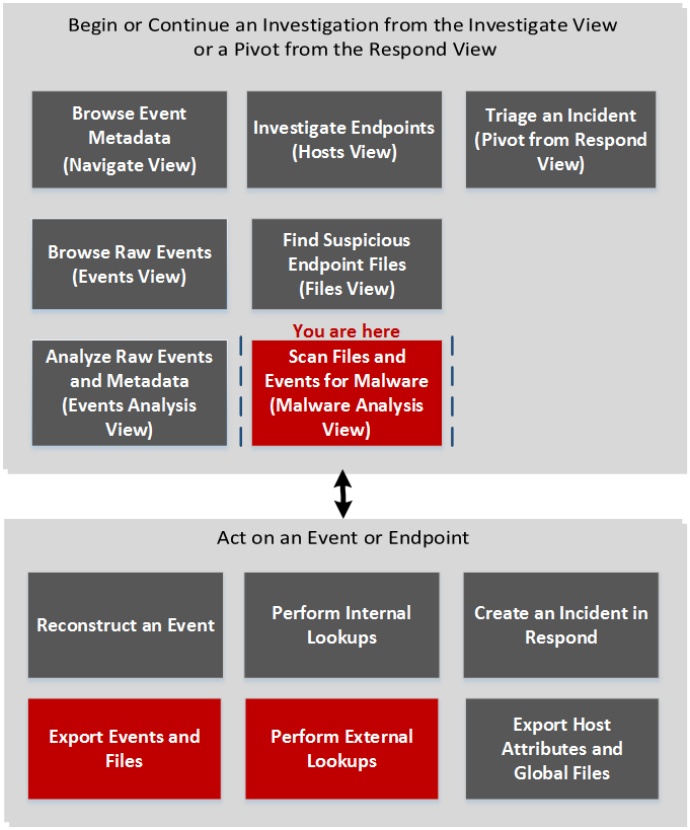
Malware Analysis Events List and Files List

The Malware Analysis Events List and Files List provide a detailed view of events or files. You can double-click on an event or file in either of the lists to display the Analysis Results view in a new browser tab.

To access this view, go to **Investigate > Malware Analysis** > **Select a Malware Analysis Service** dialog. Select a service from the left panel, then select a job from the right panel, and click **View Scan**. In the Summary of Events view do one of the following:

- In either the **Total** panel or the **High Confidence** panel, click the number in the **Events Created** section.
- If you want to view the Files List, click the number in the **Files Processed** section.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	browse event metadata	*NetWitness Investigate User Guide*

User Role	I want to ...	Show me how
Threat Hunter	browse raw events	*NetWitness Investigate User Guide*
Threat Hunter	analyze raw events and metadata	*NetWitness Investigate User Guide*
Threat Hunter	investigate endpoints (Version 11.1)	*NetWitness Endpoint User Guide*
Threat Hunter	find suspicious endpoint files (Version 11.1)	*NetWitness Endpoint User Guide*
Threat Hunter	scan files and events for malware*	Conducting Malware Analysis
Incident Responder	triage an incident in Investigate	*NetWitness Respond User Guide*
Threat Hunter	export events and Files*	Examine Scan Files and Events in List Form
Threat Hunter	perform external lookups*	View Detailed Malware Analysis of an Event

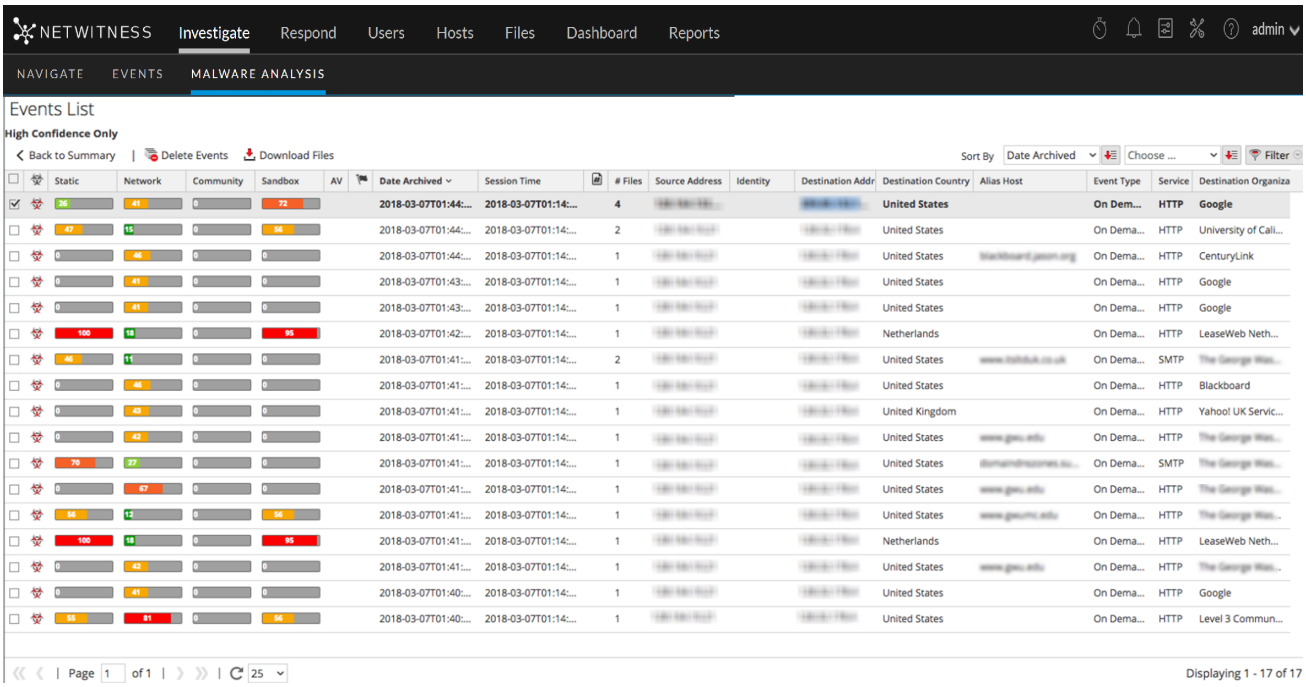
*You can perform this task in the current view.

Related Topics

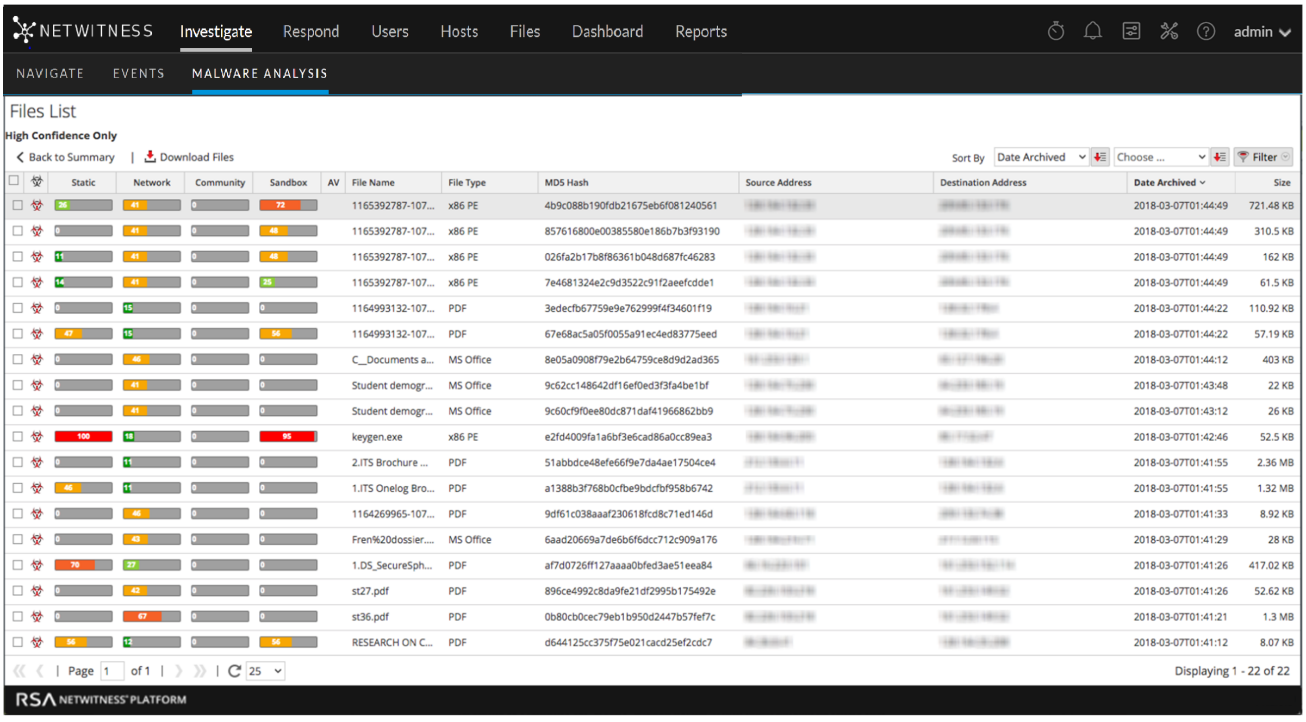
- "How NetWitness Investigate Works" in the *NetWitness Investigate User Guide*

Quick Look

This is an example of the Events List view.



This is an example of the Files List view.



These are the features in the Events List toolbar, and the Files List toolbar is the same, except it has no option to delete events.



Feature	Description
Back to Summary	Returns to the Summary of Events view.
Delete Events	Removes the selected events from the current events list.
Download Files	Displays the Malware File Download dialog, which allows you to download available files.
Choose ...	Displays a drop-down menu from which you can decide how to sort the list. These are the options for sorting: • High Confidence • Static • Network • Community • Sandbox • AV • File Name • File Type • Hash • Date Archived • Size The button directly to the right of this drop-down indicates whether the list will be sorted by ascending or descending values.
Choose ...	Displays a drop-down menu from which you can select a secondary sorting order. This menu includes an option forNetWitness**None**, so selecting a secondary sorting order is not necessary.
Filter	Displays a drop-down window in which you can filter the list by filename or MD5 Hash.

These are the features in the Events List.

Feature	Description
	Indicates whether the event is influenced by the high confidence flag.
Static, Network, Community, Sandbox	Displays the scores for each scoring module.
AV	Indicates whether the AV flagged this event as suspicious.

Feature	Description
	Indicates whether the event is influenced by a customized rule.
Date Archived	Displays the date and time the event was archived.
Session Time	Displays the time of the event's session.
	Indicates whether the hash value is marked as trusted.
# Files	Displays the number of files included in the event.
Source Address	Displays the address of the event source.
Identity	Displays the identity of the event source.
Destination Address	Displays the address of the event destination.
Destination Country	Displays the country of the event destination.
Alias Host	Displays the hostname of the alias.
Event Type	Displays the type of event. For example, Manual Upload.
Service	Displays the service on which the event occurred.
Destination Organization	Displays the organization of the destination.

These are the features in the Files List grid.

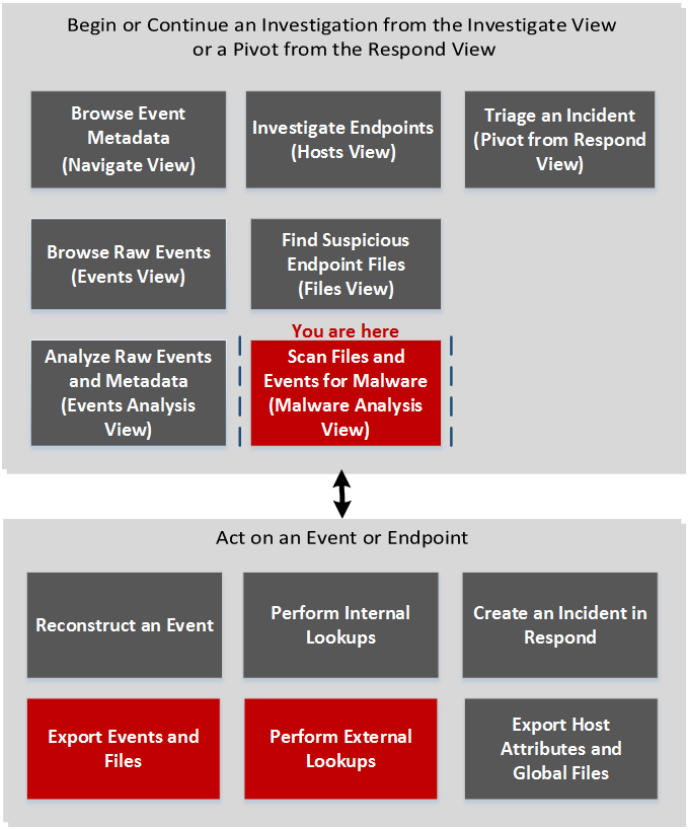
Feature	Description
	Indicates whether the event is influenced by high confidence flag.
Static, Network, Community, Sandbox	Displays the scores for each scoring module.
AV	Indicates whether the AV flagged this event as suspicious.
File Name	Displays the name of the file.
File Type	Displays the type of the file (for example, PDF or x86 PE)
MD5 Hash	Displays the MD5 hash.
Source Address	Displays the address of the file source.
Destination Address	Displays the address of the file destination.
Date Archived	Displays the date and time the file was archived.
Size	Indicates the size of the file.

Scan For Malware Dialog

In the Scan for Malware dialog, Malware Analysis analysts can upload files to investigate in Malware Analysis.

To access this dialog go to the **Malware Analysis** view. In the **Select a Malware Analysis Service** dialog, select a service in the left panel, then click Scan Files in the right panel.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	browse event metadata	*NetWitness Investigate User Guide*
Threat Hunter	browse raw events	*NetWitness Investigate User Guide*
Threat Hunter	analyze raw events and metadata	*NetWitness Investigate User Guide*

User Role	I want to ...	Show me how
Threat Hunter	investigate endpoints (Version 11.1)	*NetWitness Endpoint User Guide*
Threat Hunter	find suspicious endpoint files (Version 11.1)	*NetWitness Endpoint User Guide*
Threat Hunter	scan files and events for malware*	Conducting Malware Analysis
Incident Responder	triage an incident in Investigate	*NetWitness Respond User Guide*

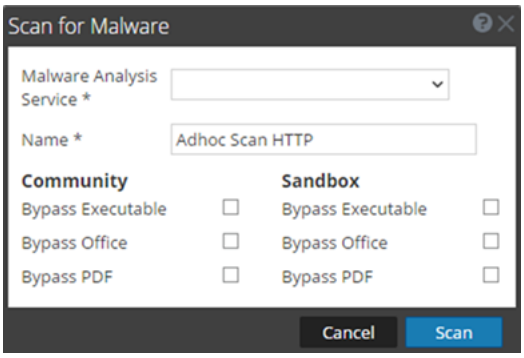
*You can perform this task in the current view.

Related Topics

- "How NetWitness Investigate Works" in the *NetWitness Investigate User Guide*
- Begin a Malware Analysis Investigation
- "Launch a Malware Analysis Scan from the Navigate View" in the *NetWitness Investigate User Guide*

Quick Look

The figure below illustrates the Scan for Malware dialog, and The following table describes the features available in the dialog.



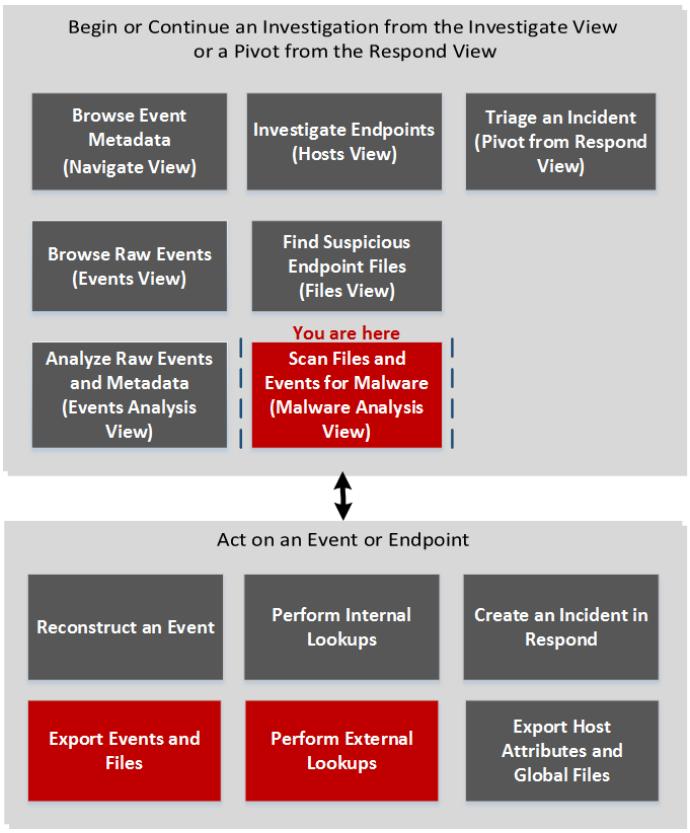
Feature	Description
+	Uploads a file from your computer.
−	Deletes a file from the list.
File Name	Displays the names of the files added to the list.
Name	Allows you to name the scan job.

Feature	Description
Community	Displays options for Community to bypass or ignore certain types of files: • Bypass Executable • Bypass Office • Bypass PDF
Sandbox	Displays options for Sandbox to bypass or ignore certain types of files: • Bypass Executable • Bypass Office • Bypass PDF
Cancel	Closes the dialog without performing any actions.
Scan	Scans the uploaded files.

Select a Malware Analysis Service Dialog

The Select a Malware Analysis Service dialog is accessible in the Malware Analysis view. In this dialog, Malware Analysis analysts can select a service to investigate, choose a scan on that service to investigate, upload a file to scan, and begin a continuous scan of the service.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Threat Hunter	browse event metadata	*NetWitness Investigate User Guide*
Threat Hunter	browse raw events	*NetWitness Investigate User Guide*
Threat Hunter	analyze raw events and metadata	*NetWitness Investigate User Guide*
Threat Hunter	investigate endpoints (Version 11.1)	*NetWitness Endpoint User Guide*

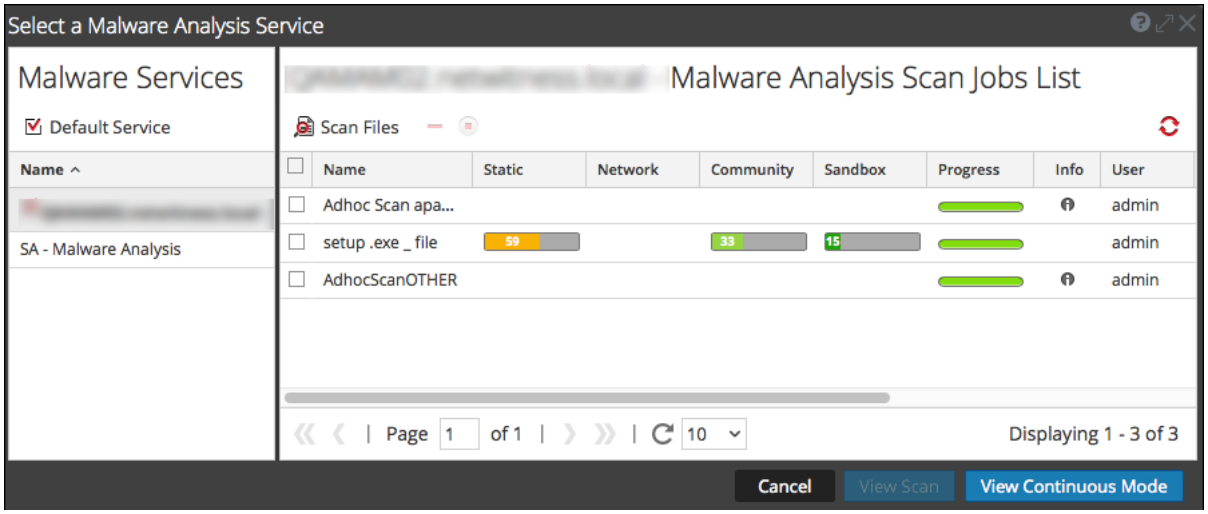
User Role	I want to ...	Show me how
Threat Hunter	find suspicious endpoint files (Version 11.1)	*NetWitness Endpoint User Guide*
Threat Hunter	scan files and events for malware*	Conducting Malware Analysis
Incident Responder	triage an incident in Investigate	*NetWitness Respond User Guide*

*You can perform this task in the current view.

Related Topics

- "How NetWitness Investigate Works" in the *NetWitness Investigate User Guide*
- Begin a Malware Analysis Investigation
- "Launch a Malware Analysis Scan from the Navigate View" in the *NetWitness Investigate User Guide*

Quick Look



The Select a Malware Analysis Service dialog has a Malware Services panel on the left and a Scan Jobs List on the right. The Scan Jobs List panel has a toolbar, list, and buttons to view scans.

The Malware Services panel is a list of services available for malware analysis. In this panel, you can select the service to investigate and you set a default service using the Default Service icon. When you select a service, the available scan jobs for that service are listed in the Scan Jobs list.

These are the features in the Scan Jobs List toolbar.

Feature	Description
Scan Files	Displays the Scan for Malware dialog, in which you can upload a file to the service for scanning.

Feature	Description
Delete scan job ( )	Deletes one or more selected scan jobs, NetWitness displays a confirmation dialog before deleting scan jobs.
Cancel scan job ( )	Pauses or continues one or more scan jobs.
Refresh ( )	Refreshes the list of scan jobs.

These are the columns in the Scan Jobs list. This list is also available in the Malware Scan Jobs dashlet.

Feature	Description
Name	Displays the name of the job.
Static, Network, Community, Sandbox	Filters the results based on the scores for each scoring module.
Progress	Displays the current progress made on the job. • Green: The job is finished. • Black: The job is in progress. • Red: An error occurred.
Info	Provides additional information. Displays the query for the job. If the job is not complete, it also displays more detailed description of the status.
User	Displays the name of the user who created the job.
Events	Counts the number of events for the job.
Dropped	Counts the number of files or events in the job that were dropped because the scores are below their configured threshold.
Event Type	Displays the type of job: Manual Upload, On Demand, or Resubmit.
Scheduled	Displays the date and time when the job was executed.

These are the available actions in the dialog.

Feature	Description
Cancel button	Cancels the selected scan job.
View Scan button	Displays the Summary of Events for the selected scan with the default dashlets displayed.
View Continuous Mode button	Displays the Summary of Events for the selected scan with the default dashlets displayed.

NetWitness® Platform XDR

Version 12.0

NwConsole User Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Access NwConsole and Help **4**
- Prerequisites 4
- Access NwConsole 4
- View Help 4
  - View a List of Commands 5
  - View Detailed Help on a Command 6
  - View a List of Help Topics 7
  - View a Specific Help Topic 7
- Quit NwConsole 8

Basic Command Line Parameters and Editing **9**
- Basic Command Line Parameters 9
- Line Editing 10

Connecting to a Service **11**

Monitoring Stats **15**

Useful Commands **16**
- Feeds 16
  - create 16
  - stats 17
  - dump 17
- Converting Packet DB Files to PCAP 17
- Packets 18
- Verifying Database Hashes 19

SDK Content Command **20**

SDK Content Command Examples **22**

Commands Used for Troubleshooting **27**
- whatIsWrong 27
- dbcheck 27
- topQuery 28
- netbytes 28
- netspeed 28

Access NwConsole and Help

NetWitness Console, also known as NwConsole, is a multi-platform terminal application that provides powerful tools and command line access to Core services, such as Decoder, Log Decoder, Concentrator, Broker, and Archiver. While most users complete their tasks and investigations through the NetWitness user interface, some advanced users, such as administrators and developers, require direct access to the services without going through the user interface. NwConsole enables you to enter commands from the command line or run multiple commands from a file.

This topic describes how to access NwConsole and view the internal help within NwConsole.

Extensive help information is available within the NetWitness console, also known as NwConsole. You can access this help from the NetWitness command line.

Prerequisites

All NetWitness appliances have the NwConsole application installed. You can also install it on Windows, Mac, and CentOS to connect and interact with a Core service.

NwConsole is available from the command line on a NetWitness appliance. If you are accessing a Core appliance remotely, you need to have the NetWitness Console application installed on a Windows, Mac, or CentOS machine. To obtain the NetWitness Console application installer, contact NetWitness Customer Care.

Access NwConsole

To run NwConsole from the command line on a NetWitness appliance or on a terminal emulator, at the <$> prompt, type `NwConsole` (Linux) or `nwconsole` (Windows). The actual command is `NwConsole`, but Windows is not case sensitive. NetWitness Console is displayed as shown in the following example.

```
Last login: Thu Sep 24 14:00:42 on console
usxx<username>m1:~ <username>$ NwConsole
RSA NetWitness Platform Console 11.2.0.0.6105
Copyright 2001-2020, RSA Security Inc. All Rights Reserved.

Type "help" for a list of commands or "man" for a list of manual
pages.
>
```

View Help

NwConsole provides help on individual commands as well as help on specific topics.

Caution: To get the latest information, view the command and help topics within NwConsole.

View a List of Commands

To view a list of available commands and their descriptions, at the (>) prompt, type `help`. The following example shows a list of available commands.

```
> help
Local commands:
    avro2nwd      - Convert AVRO files to NWD files
    avrodump      - Display schema and contents of AVRO file (for debugging)
    blockspeed    - Tests various write block sizes to determine best setting
    compileflex   - Compile all flex parsers in a directory
    createflex    - Create a flex parser that matches tokens read from a file
    dbcheck       - Perform a database integrity check over one or more
                    session, meta, packet, log or stat db files
    diskspeed     - Measures the speed of the disk(s) mounted at a specified
                    directory
    echo          - Echos the passed in text to the terminal
    encryptparser - Encrypt all parsers in a directory
    feed          - Create and work with feed files
    fmanip        - Manipulate a file with XOR and check for embedded PEs
    hash          - Creates or verifies hashes of database files
    help          - Provides help information for recognized console commands
    history       - Displays, erases or executes a command in the command
                    history
    httpAggStats  - Tests HTTP aggregation and reports statistics as it
                    continues
    log           - Perform operations on a log database
    logParse      - Parse line delimited logs on stdin and post results to
                    stdout
    logfake       - Create a fake log pcap file
    lua           - Execute a lua script
    makec3        - Generate C3 Test Data
    makepcap      - Convert packet database files to pcap or log files
    man           - Displays a list of topics or opens a specific manual page
                    on a topic
    metaspeed     - Tests read performance over an existing meta db
    netbytes      - Display statistics on network interface utilization
    nwdstrip      - Convert full NWD file into just session and meta file
    pause         - Wait for user input when running a script file
    reindex       - reindex a collection
```

```
    sdk             - Execute SDK commands based on the C SDK library, type "sdk
                      help" for more information
    sleep           - Sleeps for the specified milliseconds
    timeout         - Globally change the timeout for waiting for a response from
                      a service
    tlogin          - Open a trusted SSL connection to an existing service
    topQuery        - Returns the top N longest running queries from the audit
                      log (either a file or from the log API)
    vslice          - Validate index slices

Remote commands (executed on the connected service, see "login"):
    login           - Connect to a remote service.  Once connected, type help to
                      see commands available for remote execution.

For detailed help, type "help <command>"
>
```

View Detailed Help on a Command

To view detailed information about a command, type `help <command>`. The following example shows help for the `logParse` command after typing `help logParse`.

```
For detailed help, type "help <command>"
> help logParse
Usage: logParse {in=<pathname>} {indir=<pathname>} [out=<pathname>]
                [content=<c2|c3>] [device=<device,[device...]>]
                [path=<log-parsers-config-path>] [metaonly] [srcaddr=<src
                address>] [srcaddrfile=<filename,IP Address>]
Parse line delimited logs on stdin and post results to stdout

    in          - The input source file. "in=stdin" means interactive typing of
                  log.
    indir       - The input source files parent directory
    out         - The output file or output file parent directory if input is
                  set by indir. If not specified, use stdout as output.
    content     - Content version, either c2 or c3. Default is c2.
    device      - Comma delimited device list specifying devices that is
                  enabled. Default enable all devices.
    path        - The logparsers configuration path. Default will find
                  configuration file like logdecoder.
    metaonly    - The output will only contains parsed meta, otherwise will
```

```
                print log message after metas.
    srcaddr     - The source address of the all the logs
    srcaddrfile - The source address for logs in one input file, in the format
                filename,ipaddress
>
```

View a List of Help Topics

To view a list of help topics, type `man`. The following example shows a list of help topics.

```
> man
List of topics:

    Introduction
    Connecting to a Service
    Monitoring Stats
    Feeds
    Converting Packet DB Files to PCAP
    Packets
    Verifying Database Hashes
    SDK Content
    SDK Content Examples
    Troubleshooting

Type "man <topic>" for help on a specific topic, partial matches are acceptable
>
```

View a Specific Help Topic

To view help about a specific topic, type `man <topic>`. The following example shows the Packets help topic after typing `man Packets`.

```
Type "man <topic>" for help on a specific topic, partial matches are acceptable
> man Packets
                                    Packets
                                    =======

The *packets* command can be used to generate a pcap or log file based on a
list of Session IDs, a time period or a where clause.  The command is quite
flexible and can be used on any running service that has access to the raw
data from a downstream component.  Before running the command, you must first
```

```
*login* to a service and then change directory to the appropriate sdk node,
(e.g., "cd /sdk").  Unlike the *makepcap* command, which only works on the
local file system, this command is meant to be used on a remote service.

    login ...
    cd /sdk
    packets where="service=80 && time='2018-03-01 15:00:00'-'2018-03-01
    15:10:00'" pathname="/tmp/march-1.pcap"
Write 10 minutes of HTTP only packets from March 1st, to the file
/tmp/march-1.pcap.  All times are in UTC.

    packets time1="2018-04-01 12:30:00" time2="2018-04-01 12:35:00"
    pathname=/media/sdd1/packets.pcap.gz
Write all packets between the two times to a gzip compressed file at
/media/sdd1/packets.pcap.gz

    packets time1="2018-04-01 12:30:00" time2="2018-04-01 12:35:00"
    pathname=/media/sdd1/mylogs.log
Write all logs between the two times to a plaintext file at
/media/sdd1/mylogs.log.  Any pathname ending with .log indicates that the
format of the output file should be plaintext line-delimited logs.
>
```

Caution: To get the latest information, view the command and help topics within NwConsole.

Quit NwConsole

To exit the NwConsole application, type `quit` at the command line.

Basic Command Line Parameters and Editing

NwConsole is like a Swiss army knife; it contains many tools buried underneath its command line interface. NwConsole is multi-platform. Executables are available for CentOS (which ships on appliances), Windows, and Mac. NwConsole is included on all hosts.

Basic Command Line Parameters

Here are some basic command line parameters:

- `-f` To run a set of commands from a file, use the `-f` attribute as shown here:
  `NwConsole -f /tmp/<somefile.script>`
- `-e` You can use the `-e` attribute (which is similar to the -f attribute) to run a set of commands from a file and allow environment variable substitution within the file using `$ENV_VAR` syntax, as shown here:
  `NwConsole -e /tmp/<somefile.script>`
  Use `\$` to escape a dollar sign and `\\` to escape a backslash.
- `-c` To pass in a list of commands from the command line, use the `-c` attribute as shown here:

  `NwConsole -c <command1> -c <command2> -c <command3>`

  This is not recommended except for very simple scripts. The Bash interpreter can jumble quoted strings if you do not escape properly. If you have non-obvious errors passing through the command line, switch to reading from a file to see if that fixes the issues.
- `-i` Normally, the NwConsole exits after running commands passed by a file or command line. If you want to keep the interactive prompt open after the commands are executed, include `-i` in the command line.
- `-q` To suppress command information messages and only see command output, use `-q` for Quiet mode. This makes it easy to pipe command output (`stdout`) to other commands.

  You can also run NwConsole and type the commands in the console window.

When you use the `-c` option, you can use additional quotes and escape quotes to include embedded spaces for the `where` clause. For example:

```
where="\"tcp.dstport=443 && time='2020-09-07 17:25:00'-'2020-09-08
05:24:59'\""
```

The following example shows how to use the `-c` option for pcap extraction.

```
[root@nwadmin1 ~]# NwConsole -c login 127.0.0.1:56003:ssl <user> <passowrd> -
c cd sdk -c packets
where="\"tcp.dstport=443 && time='2020-09-07 17:25:00'-'2020-09-08
05:24:59'\""
pathname="/var/netwitness/ny_sdwan_tcp_443_20200930.pcap"
RSA NetWitness NextGen Console 11.3.1.0
```

```
Copyright 2001-2020, RSA Security Inc. All Rights Reserved.

>login ...
Successfully logged in to 127.0.0.1:56003 as session 168886
>cd sdk
[127.0.0.1:56003] /sdk
>packets where="tcp.dstport=443 && time='2020-09-07 17:25:00'-'2020-09-08
05:24:59'"
pathname=/var/netwitness/ny_sdwan_tcp_443_20200930.pcap
Writing packets to /var/netwitness/ny_sdwan_tcp_443_20200930.pcap (100%)
```

Line Editing

You can use the keys in the following table when editing a command.

Key	Descripton
Ctrl-U	Clears the current line
Ctrl-W	Deletes the word that the cursor is on
Ctrl-A	Moves the cursor to the beginning of the line
Ctrl-E	Moves the cursor to the end of the line
Ctrl-F	Moves the cursor forward to the next word
Ctrl-B	Moves the cursor backwards to the previous word
Up arrow	Displays the previously executed command
Down arrow	Displays the command executed after the current command (only valid if the up arrow has been pressed)
Left arrow	Moves the cursor to the previous character
Right arrow	Moves the cursor to the next character
Tab	Provides context sensitive completion of most commands and their parameters. The Tab key is very helpful for editing. For example, to view the *Connecting to a Service* help topic, at the command line, you can type `man con` and then press the `Tab` key. NwConsole completes the command for you: `man Connecting to a Service` Press **Enter** to run the command and view the topic.
`history`	Displays a numbered list of previous commands
`history execute=#`	Executes a previous command, which is also equivalent to typing `!#` For example, `!1` executes the previous command.
`history clear`	Clears all command history
`history erase=#`	Erases a specific command from the history buffer. History is automatically stored from one session to the next.

Connecting to a Service

To connect and interact with a NetWitness Core service (Decoder, Concentrator, Broker, Archiver, and so on), you must first issue the `login` command. You must have an account on that service. You can type `help login` at any time for more information. Here is the syntax of the `login` command:

`login <hostname>:<port>[:ssl] <username> [password]`

For example: `login 10.10.1.15:56005:ssl someuser`

If you do not include the password, the NwConsole prompts you.

If you have set up proper trust between NwConsole and the endpoint, you can use the `tlogin` command and avoid having to enter a password. Setting up trust is beyond the scope of this documentation, but it involves adding NwConsole's SSL cert to the endpoint with the `send /sys peerCert op=add --file-data=<pathname of cert>` command. You must first use a normal login with the proper permissions before you can add a peer cert for subsequent trusted logins.

Once connected, you can interact with the endpoint service through a virtual file system. Instead of files, what you are looking at are the nodes of that service. Some nodes are folders and have child nodes, forming a hierarchical structure. Each node serves a purpose and all of them support a subset of commands like `info` and `help`. The `help` message returns information about the commands that each node supports. When you first log on, you are on the root node, which is the path `/`, just like a Linux or Mac system. To see a list of nodes under `/`, type the `ls` command.

All services have nodes like `sys` and `logs`. To interact with the **/logs** API, you can first send the `help` command to the **/logs** node. To do this, you must use the `send` message, which has this syntax:

```
Usage: send {node pathname} {message name} [name=value [name=value]]
            [--file-data=<pathname>] [--string-data=<text>] [--binary-data=<text>]
            [--file-format={binary,params,params-list,string,params-binary}]
            [--output-pathname=<pathname>] [--output-append-pathname=<pathname>]
            [--output-format={text,json,xml,html}]
Sends a command to a remote pathname.  For remote help, use "send <pathname>help" for
details.
    pathname                - The node pathname to retrieve information on
    message                 - The command (message) to send
    parameters              - Zero or more name=value parameters for the command
    --file-data             - Loads data from a file and send as either a BINARY
                              message or as a PARAMS_BINARY message if other
                              parameters exist. Providing the --file-format parameter
will
                              override automatic detection.
   --file format            - Treat file input as the provided format
  --string-data            - Sends text as a STRING message type
   --binary-data            - Send text as either a BINARY message type or as a
                              PARAMS_BINARY message type if other parameters
                              exist
    --output-pathname       - Writes the response output to the given pathname,
```

```
                              overwriting any existing file
  --output-append-pathname - Writes the response output to the given pathname,
                              will append output to an existing file
  --output-format          - Writes the response in one of the given formats,
                              the default is text
```

Parameters are formatted in `<name>=<value>` pairs, with whitespace between each pair. If a `<value>` has whitespace, you should place each value in double quotes. If the value itself should have double quotes, you can escape them by preceding with a backslash. For example: `send /sdk query query="select * where time=\"12-31-2019 14:30:00\"-u" size=100`

You must also escape backslashes; a double backslash is treated as a single backslash. It is also possible to pass special characters by giving their hex value using `\x##`. For example, `\x0a` would pass a line feed in the parameter value.

To send a `help` message, you would send this:

```
send /logs help
```

And your response would look something like this:

```
description: A container node for other node types
  security.roles: everyone,logs.manage
  message.list: The list of supported messages for this node
  ls:  [depth:<uint32>] [options:<string>] [exclude:<string>]
  mon:  [depth:<uint32>] [options:<uint32>]
  pull:  [id1:<uint64>] [id2:<uint64>] [count:<uint32>] [timeFormat:<string>]
  info:
  help:  [msg:<string>] [op:<string>] [format:<string>]
  count:
  stopMon:
  download:  [id1:<uint64>] [id2:<uint64>] [time1:<date-time>] [time2:<date-time>]
    op:<string> [logTypes:<string>] [match:<string>] [regex:<string>]
    [timeFormat:<string>] [batchSize:<uint32>]
  timeRoll:  [timeCalc:<string>] [minutes:<uint32>] [hours:<uint32>] [days:<uint32>]
    [date:<string>]
```

To get more information about a specific message or command, you can specify the `msg=<message name>` on the help command as a parameter. For example, look at the `pull` message help:

```
send /logs help msg=pull
```

```
 pull: Downloads N log entries
  security.roles: logs.manage
  parameters:
  id1 - <uint64, optional> The first log id number to retrieve, this is
     mutually exclusive with id2
  id2 - <uint64, optional> The last log id number that will be sent, defaults to
```

```
       most recent log message when id1 or id2 is not sent
    count - <uint32, optional, {range:1 to 10000}> The number of logs to pull
    timeFormat - <string, optional, {enum-one:posix|simple}> The time format used in
each
       log message, default is posix time (seconds since 1970)
```

The built-in message help says that this command grabs the last N log entries if you omit ID1 and ID2. To look at the last 10 log entries:

```
send /logs pull count=10 timeFormat=simple
```

Almost all of the commands on the service follow this simple format. The only commands that do not are the ones that require more complicated handshaking, like importing a PCAP to a Decoder. To import a PCAP, use the NwConsole `import` command, which takes care of the complicated communication channel handshaking.

Some parameters are specific to NwConsole's `send` command and are not actually sent to the service. You can use these parameters to change the output format of the response, write the response to a file, or read a file from the local machine and send it to the service. The local parameters to NwConsole's `send` command all start with two dashes --.

- `--output-format` — This parameter changes the normal output of the command from plain text to one of these types: JSON, XML, or HTML. The format `values` is a text output that only writes values (query output, stat/config values, and so on) without any other decorations.
- `--output-pathname` — Instead of writing the output to the terminal, the output is written to the specified pathname (truncates any existing file).
- `--output-append-pathname` — This is the same as `--output-pathname` except that it appends the output to an existing file (or creates the file if it does not exist).
- `--file-data` — Reads in a file and uses it as the command payload. This is useful for commands like `/sys fileEdit`. The following example shows how you can send an updated **index-concentrator-custom.xml** file using NwConsole:

  ```
  send /sys fileEdit op=put filename=index-concentrator-
  custom.xml --file-data="/Users/user/Documents/index-
  concentrator-custom.xml"
  ```

- `--file-format` — When reading an input file with `--file-data`, this parameter forces NwConsole to interpret the file as a specific type of input. The allowed enumerations are: `binary`, `params`, `params-list`, `string` and `params-binary`. As an example, to send a file of application rules (`*.nwr`) to a Decoder, you can use this command:

  ```
  send /decoder/config/rules/application replace --file-
  data=/path/rules.nwr --file-format=params-list
  ```

- `--string-data` — Sends the command payload as a string instead of a list of parameters.
- `--binary-data` — Sends the command payload as binary instead of a list of parameters.

Example Streaming Query to JSON file (could be a large result set):

```
send /sdk query size=0 query="select * where service=80 && time='2018-
03-05 13:00:00'-'2018-03-05 13:59:59'" --output-format=json --output-
pathname=/tmp/query.json
```

One thing to note about the `send` command is the fact that, by default, there is a timeout of 30 seconds waiting for a response. Some commands (like the query above) may take longer to receive results. To avoid a premature client-side timeout, you can use the `timeout [secs]` command to increase the wait. For instance, `timeout 600` would wait 10 minutes for a response before timing out. Once enacted, it takes effect for all subsequent commands.

To navigate around the virtual node hierarchy of the service, you can use the `cd` command like you would on any command shell. This covers the basics of connecting and interacting with a service. Once you are connected, the `help` command lists all the commands that you can use to interact with the endpoint. These commands do not display when you are not connected to an endpoint.

Monitoring Stats

You can use NwConsole to watch statistics (stats) change on a service in real time. However, be warned that this can result in a LOT of output. If you are not careful and monitor too many nodes, the screen scrolls by too quickly to be useful.

As a simple example, if you log on to a Decoder, you can monitor the capture rate in real time. To do this, issue these commands after connecting to a Decoder:

```
cd /decoder/stats
```

```
mon capture.rate
```

That is all you need to do! Now, any time the capture rate changes, it outputs into the console window.

You can add another monitor:

```
mon capture.avg.size
```

Now it watches those two stats and outputs those values when they change. You may have noticed that as you tried to type the second command, the output from the original monitor was messing up your display. This is the problem with monitoring stats. It is not really meant for doing more than just watching the stats after the first command is entered.

However, you can stop the monitoring by typing `delmons` and pressing **Enter**. Just ignore the output while you type and it returns you to a proper command prompt. If you want to monitor many stats at once, you can give the path of the parent stat folder and it monitors all of the stats underneath it. For example, typing `mon /decoder/stats` or `mon .` (they are equivalent) monitors everything. Be prepared for a lot of output! Remember to enter `delmons` if it is scrolling too fast.

Useful Commands

The following NwConsole commands are useful when interacting with NetWitness Core services:

- **feed**: Enables you to create and work with feed files.
- **makepcap**: Converts Packet database (DB) files to PCAP.
- **packets**: Retrieves packets or logs from the logged in service.
- **hash**: Creates or verifies hashes of database files.

The following sections as well as the NwConsole help and topic information (man) pages, provide additional information.

Feeds

The `feed` command provides several utilities for creating and examining feed files. A feed file contains the definition and data of a single feed in a format that has been precompiled for efficient loading by a Decoder or Log Decoder. For a complete reference on feed definitions, see the "Feed Definitions File" topic in the *Decoder Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

create

```
feed create <definitionfile> [-x <password>]
```

The `feed create` command generates feed files for each feed defined in a feed definition file. A definition file is an XML document that contains one or more definitions. Each feed definition specifies a data file and the structure of that data file. The resulting feed files will be created in the same directory as the definition file with the same name as the data file, but with the extension changed to **.feed** (for example, **datafile.csv** results in **datafile.feed**). Any existing files with the target name will be overwritten without a prompt.

```
$ ls
 example-definition.xml    example-data.csv
 $ NwConsole
 RSA NetWitness Console 11.3.0.0.0
 Copyright 2001-2020, RSA Security Inc.  All Rights Reserved.

 Type "help" for a list of commands or "man" for a list of manual pages.
 > feed create example-definition.xml
 Creating feed Example Feed...
 done.  2 entries, 0 invalid records
 All feeds complete.
 > quit
 $ ls
 example-definition.xml    example-data.feed    example-data.csv
 $
```

Optionally, feed files can be obfuscated using the option `-x` followed by a password of at least 16 characters (no spaces). This will be applied to all feeds defined in the definition file. In addition to the feed file, a token file will be generated for each feed file. The token file must be deployed with the corresponding feed file.

```
feed create example-definition.xml -x 0123456789abcdef
```

stats

```
feed stats <feedfile>
```

The `feed stats` command provides summary information for an existing, un-obfuscated feed file. Specifying an obfuscated feed file will result in an error.

```
> feed stats example.feed
 Example Feed stats:
 version  : 0
 keys count  : 1
 values count: 2
 record count: 2
 meta key  : ip.src/ip.dst
 language keys:
    alert    Text
```

dump

```
feed dump <feedfile> <outfile>
```

The `feed dump` command generates a normalized, key-value pair listing of an un-obfuscated feed file. You can use the resulting file to validate a feed file or assist in determining which records were considered invalid when the feed was created. Specifying an obfuscated feed file will result in an error. If `outfile` exists, the command will abort without overwriting the existing file.

```
feed dump example.feed example-dump.txt
```

Converting Packet DB Files to PCAP

You can use the `makepcap` command to quickly convert any Packet DB file to a generic PCAP file, preserving the capture time order. This command offers many options (see `help makepcap`), but is easy to use. All it really needs is the Packet DB directory (with the `source=<pathname>` parameter) to get started.

Note: You must stop the Decoder or Archiver service before running this command. If you want to generate a PCAP while the service is running, see the `packets` command.

- `makepcap source=/var/lib/netwitness/decoder/packetdb`
  This command converts every Packet DB file into a corresponding PCAP file in the same directory. If the disk is almost full, see the next command.

- `makepcap source=/var/lib/netwitness/decoder/packetdb dest=/media/usb/sde1`
  This command writes all of the output PCAPs to the directory at **/media/usb/sde1**.

- `makepcap source=/var/lib/netwitness/decoder/packetdb dest=/media/usb/sde1 filenum=4-6`
  This command only converts the files numbered 4 through 6 and skips all other files. In other words, it converts the Packet DB files: **packet-000000004.nwpdb**, **packet-000000005.nwpdb**, and **packet-000000006.nwpdb**.

- `makepcap source=/var/lib/netwitness/decoder/packetdb time1="2020-03-01 14:00:00" time2="2020-03-02 07:30:00" fileType=pcapng`
  This command only extracts packets with a timestamp between March 1st, 2020 at 2 PM and March 2nd, 2020 before or on 7:30 AM. It writes the file as `pcapng` in the same directory as the source. All timestamps are UTC.

Packets

You can use the `packets` command to generate a PCAP or log file based on a list of Session IDs, a time period, or a `where` clause. This command is very flexible, and you can use it on any running service that has access to the raw data from a downstream component. Before running the command, you must first `login` to a service and then change directory to the appropriate SDK node (for example, `cd /sdk`). Unlike the `makepcap` command, which only works on the local file system, you use this command for a remote service.

```
login ...
cd /sdk
packets where="service=80 && time='2020-03-01 15:00:00'-'2020-03-01 15:10:00'"
pathname="/tmp/march-1.pcap"
```

This command writes 10 minutes of HTTP-only packets from March 1st to the file **/tmp/march-1.pcap**. All times are in UTC.

- `packets time1="2020-04-01 12:30:00" time2="2020-04-01 12:35:00" pathname=/media/sdd1/packets.pcap.gz`
  This command writes all packets between the two times to a GZIP compressed file at **/media/sdd1/packets.pcap.gz**.

- `packets time1="2020-04-01 12:30:00" time2="2020-04-01 12:35:00" pathname=/media/sdd1/mylogs.log`
  This command writes all logs between the two times to a plaintext file at **/media/sdd1/mylogs.log**. Any pathname ending with **.log** indicates that the format of the output file should be plaintext line-delimited logs.

Verifying Database Hashes

By default, Archiver writes an XML file for every DB file that is written. This XML file ends with the extension **.hash** and contains a hash of the file along with other pertinent information. You can use the `hash` command to verify that the DB file has not been tampered with by reading the hash stored in the XML file and then rehashing the DB file to verify that the hash is valid. The command does accept wildcards, so something like `/var/netwitness/archiver/database0/default/packetdb/*.hash` will verify all hashes in that directory. You can also use the `--output-pathname` and `--output-format` parameters to write out the verification in json, xml or text formats for scripting purposes. If the DB files are not found in the location that the XML file says it should be in, then the command will attempt to find the DB files in the current directory or the directory where the hash file is found.

```
hash op=verify
hashfile=/var/lib/netwitness/archiver/database0/alldata/packetdb/packe
t-000004880.nwpdb.hash
```

This command verifies that the Packet DB file **packet-000004880.nwpdb** still matches the hash in the XML file **packet-000004880.nwpdb.hash**. For proper security, the hash file should be stored somewhere else to prevent the XML file from being tampered with (such as write-once only media), but the hash command itself is not affected by where it is stored.

SDK Content Command

One of the powerful commands in NwConsole is `sdk content`. It contains numerous options to do just about anything, at least as far as extracting content from the NetWitness Core stack. You can use it to create PCAP files, log files, or extract files out of network sessions (for example, to grab all of the pictures from email sessions). It can append files, have a maximum size assigned before creating a new file, and automatically clean up files when the directory grows too large. It can run queries in the background to find new sessions. It breaks queries into manageable groups and performs those operations automatically. When the group is exhausted, it does a requery to get a new set of data for further operations. The list of options for the SDK content command is very extensive.

Because the command has so many options, this document provides examples of commands for different use cases.

Before you can run `sdk content`, there are a few commands (like logging into a service) that you need to run first. Here are some examples:

- First connect to a service:
  `sdk open nw://admin:netwitness@10.10.25.50:50005`
- If you need to connect over SSL, use the nws protocol:
  `sdk open nws://admin:netwitness@10.10.25.50:56005`
- Keep in mind that you are passing a URL and must URL encode it properly. If the password is `p@ssword`, the URL looks like this:
  `sdk open nw://admin:p%40ssword@10.10.25.50:50005`
  This also applies to username.
- Once you log in, you can set an output directory for the commands: `sdk output <some pathname>`
- For command line help, type: `sdk content`

Before you try the example commands, it is important to understand the `sessions` parameter. This parameter is very important and controls how much or how little data you want to grab (the `where` clause is also important). The sessions parameter is either a single session ID or a range of session IDs. All NetWitness Core services work with session IDs, which start at 1 and increase by 1 for every new session added to the service (network or log session). Session IDs are 64-bit integers, so they can get quite large.

For example, to keep it simple, assume we have a Log Decoder that has ingested and parsed 1000 logs. On the service, you now have 1000 sessions with session IDs from 1 to 1000 (session ID 0 is never valid). If you want to operate over all 1000 sessions, you pass `sessions=1-1000`. If you only want to operate over the last 100 sessions, you pass `sessions=901-1000`. Once the command finishes processing session 1000, it exits back to the console prompt.

Many times, however, we do not care about specific session ranges. We just want to run a query over all of them and process the sessions that match a query. Here are some shortcuts that simplify this:

- The letter `l` (lowercase L) means lower bound or the lowest session ID.
- The letter `u` means the highest session ID. In fact, it actually means the highest session ID for future sessions as well. In other words, if you pass `sessions=l-u`, this special range means operate over all the current sessions in the system, but also do not quit processing, and as new sessions enter the system, process those, too. The command pauses and waits for new sessions once it reaches the last session on the service. To summarize, the command never exits and goes into continuous processing mode. It runs for days, months, or years, unless it is killed.
- If you do not want the command to run forever, you can pass `now` for the upper limit. This determines the last session ID on the service at the time the command starts and processes all sessions until it reaches that session ID. Once it reaches that session ID, the command exits, regardless of how many sessions may have been added to the service since the command started. So, for the example Log Decoder, `sessions=200-now` starts processing at session 200 and goes all the way to session 1000 and quits. Even if another 1000 logs were added to the Log Decoder after the command started, it still exits after processing session 1000.
- The parameter `sessions=now-u` means start at the very last session and continue processing all new sessions that come in. It does not process any existing sessions (except the last one), only new sessions.
- If you want to run this continually and have it remember the last session it processed between invocations, use the `sessionPersist={pathname}` parameter. The pathname should be a valid filename where the command will write the last successful session id. Upon restart, it will read the last session id and pick up where it left off.

For example commands and what they do, type `man sdk content examples` or see SDK Content Command Examples.

SDK Content Command Examples

The first NwConsole SDK content command example below is simple and shows all of the commands that you need to enter. After that, the examples show only the `sdk content` commands. The first example creates a log file and grabs the first 1000 logs out of a Concentrator aggregating from a Log Decoder:

- ```
  sdk open nw://admin:netwitness@myconcentrator.local:50005
  sdk output /tmp
  sdk content sessions=1-1000 render=logs append=mylogs.log
  fileExt=.log
  ```
  This script outputs 1000 logs (assuming sessions 1 through 1000 exist on the service) to the file **/tmp/mylogs.log**. The logs are in a plain text format. The parameter `fileExt=.log` is necessary to indicate to the command that we want to output a log file.

- ```
  sdk content sessions=1-1000 render=logs append=mylogs.log
  fileExt=.log includeHeaders=true separator=","
  ```
  This command grabs the same 1000 logs as above, but it parses the log header and extracts the log timestamp, forwarder, and other information, and puts them in a CSV formatted file.
  Example CSV: `1422401778,10.250.142.64,10.25.50.66,hop04b-LC1,%MSISA-4:81.136.243.248...`
  The timestamp is in Epoch time. The `includeHeaders` and `separator` parameters can only be used on NetWitness installs 10.4.0.2 and later.

- ```
  sdk content sessions=1-now render=logs append=mylogs.log
  fileExt=.log includeHeaders=true separator=","
  where="risk.info='nw35120'"
  ```
  This command writes a log file across the current session range, but only logs that match `risk.info='nw35120'`. Keep in mind that when you add a `where` clause, it performs a query in the background to gather the session IDs for export. The query should be run on a service with the proper fields indexed (which is typically a Broker or Concentrator). In this case, since you are querying the field `risk.info`, double-check the service where you run the command to make sure it is indexed at the value level (`IndexValues`, see **index-concentrator.xml** for examples). By default, most Decoders only have time indexed. If you use any field but time in the `where` clause, you need to move the query from the Decoder to a Concentrator, Broker, or Archiver with the proper index levels for the query. You can find more information on indexing and writing queries in the *NetWitness Core Database Tuning Guide*.

- ```
  sdk content sessions=1-now render=logs append=mylogs.log
  fileExt=.log includeHeaders=true separator=","
  where="threat.category exists && time='2020-01-05 15:00:00'-'2020-01-05 16:00:00'"
  ```
  This command is the same as above, but it only searches for matching logs between 3 PM and 4 PM (UTC) on Jan 5, 2020 that have a meta key `threat.category`. Again, because this query has a field other than time in the `where` clause (`threat.category`), it should be run on a service with

`threat.category` indexed at least at the `IndexKeys` level (the operators `exists` and `!exists` only require an index at the key level, although values work fine, too).

- `sdk content sessions=1-now render=logs append=mylogs fileExt=.log where="event.source begins 'microsoft'" maxFileSize=1gb`
  This command creates multiple log files, each one no larger than 1 GB in size. It prepends the filenames with **mylogs** and appends them with the date-time of the first packet/log timestamp in the file. Some example filenames: **mylogs-1-2020-Jan-28T11_08_14.log**, **mylogs-2-2020-Jan-28T11_40_08.log** and **mylogs-3-2020-Jan-28T12_05_47.log**. On versions older than 10.5, the T separator between date and time is a space.

- `sdk content sessions=1-now render=pcap append=mypackets where="service=80,21 && time='2020-01-28 10:00:00'-'2020-01-28 15:00:00'" splitMinutes=5 fileExt=.pcap`
  This command grabs all the packets between the five-hour time period for service types 80 and 21 and writes a PCAP file. Every five minutes, it starts a new PCAP file.

- `sdk content time1="2020-01-28 14:00:00" time2="2020-01-28 14:15:00" render=pcap append=mydecoder fileExt=.pcap maxFileSize=512mb sessions=1-now`
  Pay attention to this command. Why? It works for both packets and logs and is *extremely fast*. The downside is that you get everything between the two time ranges and you cannot use a `where` clause. Again, it starts streaming everything back almost immediately and does not require a query to run first on the backend. Because everything is read using sequential I/O, it can completely saturate the network link between the server and client. It starts creating files prepended with **mydecoder** and splits to a new file once it reaches 512 MBs in size.

- `sdk tailLogs`
  or (the equivalent command):
  `sdk content render=pcap console=true sessions=now-u`
  This is a fun little command. It actually uses `sdk content` behind the scenes. The purpose of this command is to view all incoming logs on a Log Decoder. As logs come into the Log Decoder (you can also run it on a Broker or Concentrator), they are output on the console screen. It is a great way to see if the Log Decoder is capturing and what exactly is coming into the Log Decoder. This command runs in continuous mode. Do not use it if the Log Decoder is capturing at a high ingest rate (this command cannot keep up with it). However, it is helpful for verification or troubleshooting purposes.

- `sdk tailLogs where="device.id='ciscoasa'" pathname=/mydir/anotherdir/mylogs`
  This command is the same as above, except it only outputs logs that match the `where` clause and instead of outputting to the console, it writes them to a set of log files under `/mydir/anotherdir` that do not grow larger than 1 GB. Obviously, you can accomplish this with the `sdk content` command as well, but this command requires a little less typing if you like the default behavior.

- `sdk content sessions=now-u render=pcap where="service=80" append=web-traffic fileExt=.pcap maxFileSize=2gb maxDirSize=100gb`
  This command starts writing PCAPs of all web traffic from the most recent session and all new incoming sessions that match `service=80`. It writes out PCAPs no larger than 2 GBs and if all the PCAPs in the directory grow larger than 100 GBs, it deletes the oldest PCAPs until the directory is 10% smaller than the maximum size. Keep in mind that the directory size checking is not exact and it only checks every 15 minutes by default. You can adjust the number of minutes between checks by passing `cacheMinutes` as a parameter, but this only works with versions 10.5 and later.

- `sdk content sessions=79000-79999 render=nwd append=content-%1%.nwd metaFormatFilename=did`
  This is a basic backup command. It grabs 1000 sessions and outputs the full content (sessions, meta, packets, or logs) to the NWD (NetWitness Data Format) format. NWD is a special format that can be re-imported to a Packet or Log Decoder without reparsing. So essentially, the original parsed session imports without changes. The timestamp does not change as well, so if it was originally parsed six months ago, the timestamp upon import will be retained as six months ago.

  Note: Do not expect great performance with this command, especially with packets. Gathering the packets for a session involves a lot of random I/O and can drastically slow down the export. Logs do not suffer as much from this problem (only one log per session), but behind the scenes, this command uses the /sdk content API, which is not a performance-minded streaming API like /sdk packets.

  The `metaFormatFilename` parameter is very helpful in this command. If this command is run on a Concentrator with more than one service, the NWD filenames will be created with the `did` meta for each session (the `%1%` in the append parameter is substituted with the value of `did`). Each filename will indicate exactly which Decoder the data came from.

- `sdk content session=1-u where="service=80,139,25,110" render=files cacheMinutes=10 dir=/tmp/content-files maxDirSize=200mb`
  This is another fun little command. It works very similar to our old Visualize product if you pair the output directory with something like Windows Explorer in Icon mode. It extracts files from all web, email, and SMB traffic. This includes all kinds of files, such as images, zip files, videos, PDFs, office documents, text files, executables, and audio files. If it extracts malware, your virus scanner will flag it. Nothing will be executed by the command, so it does not infect the machine (unless you try to execute it yourself). However, it can be useful because if you do find malware, the filename indicates the session ID where it was extracted. You can then query that session ID and see which host the malware possibly infected and take action. You can filter what gets extracted with the parameters `includeFileTypes` or `excludeFileTypes` (see the command help). For example, adding `excludeFileTypes=".exe;.dmg;.msi"` prevents executables and installers from being extracted. This command just runs nonstop extracting files from all existing and any new sessions. After the directory gets littered with more than 200 MBs of files, it automatically starts cleaning up the files every 10 minutes.

  Note: This command only makes sense for packet sessions, not logs.

- `sdk content session=1-now render=files where="time='2015-01-27 12:00:00'-'2015-01-27 13:00:00' && (service=25,110,80)" subdirFileTypes="audio=.wav;.mp3;.aac; video=.wmv;.flv;.mp4;.mpg;.swf; documents=.doc;.xls;.pdf;.txt;.htm;.html images=.png;.gif;.jpg;.jpeg;.bmp;.tif;.tiff archive=.zip;.rar; other=*" renameFileTypes=".download|.octet-stream|.program|.exe;.jpeg|.jpg" maxDirSize=500mb`
  This command extracts files from HTTP and email sessions from a one-hour period and then groups the extracted files into directories specified by the `subdirFileTypes` parameter. For example, any extracted audio file with the extension `.wav`, `.mp3`, or `.aac` will be placed into the subdirectory `audio`, which will be created under the specified output directory. The same goes for all the other groups specified in that parameter. Some files will also be automatically renamed based on their file extension, which is handled by `renameFileTypes`. Any file with an extension `.download`, `.octet-stream` or `.program` will be renamed to `.exe`. Files with the extension `.jpeg` will be renamed `.jpg`. Once the top-level directory exceeds 500 MBs, the oldest files are cleaned. This command stops at the last session at the time the command started.

- `sdk content session=1-u render=111 where="filename exists" maxDirSize=20mb fileHash=md5,sha1,sha256 linkMeta=sessionid,time,did,service_uuid deviceType=filehash logDecoder="<hostname>:<port>[:ssl]" sessionPersist=/tmp/last-session.txt`
  If you want to extract file hashes from network protocols like HTTP, SMB, POP3 and SMTP, this is the command you want to use. Not only will it give you the file hash, but it will send an event to a Log Decoder with this information, which will be processed into a session for querying. For example, use `sdk open` to talk to a Broker or Concentrator that is aggregating from one or more Network Decoders. Set the `where` clause to find sessions of interest that have embedded files in the protocols we support for file extraction (use Live to install Malware Analysis content for example). Next, set up a Log Decoder to receive the file hash information. This information will go into the meta keys `checksum`, `checksum.algo`, `filename`, `filename.size` and `link`. The meta key link is used to indicate the session (with the file that has been hashed) on the device used by `sdk open`. When hashes of interest are discovered on the Log Decoder, the link meta can be used to find the original packet session that contains the file. The parameters `fileHash`, `linkMeta` and `deviceType` can be used to customize how the log is generated that is sent to the Log Decoder. The ports for sending Syslog to a Log Decoder are usually 514 or 6514 (for SSL). You may need to add a `table-map-custom.xml` file on the Log Decoder to properly save the meta keys sent in the file hash log. For the `linkMeta` parameter, `service_uuid` is a special meta key that inserts the UUID of the service the session ID belongs to. This can be used to link back to the original session from the Log Decoder. `sessionPersist` is used to write out our session IDs to pick up where we left off from the last run.

- `sdk search session=1-now where="service=80,25,110" search="keyword='party' sp ci"`
  This command searches all packets and logs (the `sp` parameter) for the keyword `party`. If party is

found anywhere in the packets or logs, it outputs the session ID along with the text it found and the surrounding text for context. The `where` clause indicates that it only searches web and email traffic. The `ci` parameter means that it is a case-insensitive search. You can substitute `regex` for `keyword` and it performs a regex search.

- `sdk search session=1-now search="keyword='checkpoint' sp ci" render=log append=checkpoint-logs.log fileExt=.log`
  This is an interesting command example. It searches all logs (or it could be packets) for the keyword `checkpoint` and if that keyword is seen, it extracts the log to a file **checkpoint-logs.log**. There are all kinds of possibilities with this command. Essentially, when a hit is detected, it hands off the session to the content call. So any parameters you pass to `sdk search` that it does not recognize are just passed along to the content call. This allows the full capabilities of the `sdk content` call, but they only work on those sessions with content search hits.

Commands Used for Troubleshooting

NwConsole provides the following commands that are helpful when troubleshooting NetWitness:

- **whatIsWrong**: Provides a snapshot of a service's configuration, stats, and failure and warning logs for a specified period of time in the past.
- **dbcheck**: Performs consistency checking of database files.
- **topQuery**: Helps pinpoint queries that are taking an excessively long time to run.
- **netbytes**: Troubleshoots the network connections on the current host.
- **netspeed**: Troubleshoots the connection between the host computer running NwConsole and the remote computer connected to it using the `login` command.

The following sections, as well as the NwConsole help and topic information (man) pages, provide additional information.

whatIsWrong

When a service is not working correctly, the reason is usually somewhere in the logs that the service has written. You can use the `whatIsWrong` console command to obtain a snapshot of a service's configuration, stats, and failure and warning logs (with surrounding context logs) for a specified past period of time, which defaults to the previous seven days. You can save the results of running `whatIsWrong` into a specified plain-text file. The output of this command can be a useful starting point to help determine what is currently wrong with a service.

To use the `whatIsWrong` console command, log on to the service to troubleshoot using the `login` command, and run the `whatIsWrong` command.

Note: Use `help whatIsWrong` to see all of the available parameters, including the number of days or hours to look back for events, the pathname to store results, whether or not to append or overwrite the results file, and the delimiter to use for log fields. You can also limit the number of most recent logs used to find context, and you can specify how many context logs per warning or failure log to retrieve.

Whenever you receive a request for logs for a Core service, you should run the `whatIsWrong` command first and use the results collected as a starting point.

dbcheck

The `dbcheck` command is used to perform consistency checking of database files (session, meta, packets, logs, stats, and so on). This might be necessary when a service cannot start because of errors in the consistency of the database files. Normally a service would automatically recover and correct any consistency issues on startup, but there are times when this does not occur. When a service starts (like Decoder), it typically does not read or open most database files in order to start quickly. It assumes most files are in a consistent state and only does a cursory check of the most recently written files. If there are problems, `dbcheck` can perform those consistency checks, but ONLY if the service is not running.

Caution: Do not attempt to run this command while a service is running.

For example, you can check a single file:

```
dbcheck /var/lib/netwitness/decoder/packetdb/packet-000000001.nwpdb
```

You can also use wildcards to check multiple files:

```
dbcheck /var/lib/netwitness/decoder/metadb/meta-00000002*.nwmdb
```

topQuery

The `topQuery` command can help pinpoint queries that are taking an excessively long time to run. This command parses the audit logs of a service and returns the top N longest running queries for the specified time period.

The easiest way to run it is to log on to the service (usually a Broker or Concentrator) and type `topQuery`. The default behavior is to return the top 100 longest running queries for the last seven days.

Type `help topQuery` for the list of parameters. Here are some additional examples with explanations:

- `topQuery hours=12 top=10`
  This command returns the top 10 queries for the last 12 hours.

- `topQuery time1="2020-03-01 00:00:00" time2="2020-03-14 00:00:00"`
  This command returns the top 100 queries between March 1, 2020 and March 14, 2020. Times are in UTC, not local.

- `topQuery input=/var/log/messages output=/tmp/top20.txt top=20 user=sauser1`
  Instead of connecting to a service, it parses the Syslog audit messages for the top 20 queries in the last seven days, but only for queries executed by user sauser1. It writes the top 20 queries to **/tmp/top20.txt** instead of the console screen. The parameter user is a regex, so you can specify multiple usernames by writing something like `user="(sauser1|sauser2)"`.

netbytes

The `netbytes` command is very useful for troubleshooting the network connections on the current host. It displays continuous send and receive statistics for all network interfaces. Once executed, you must press **Ctrl-C** to exit this command, which also exits NwConsole.

netspeed

The `netspeed` command is used to troubleshoot the connection between the host computer running NwConsole and the remote computer connected to it through the `login` command. You must supply the amount of bytes to transfer and it will time the speed of the connection. The netspeed command is very useful for troubleshooting aggregation performance issues that might be network related.

```
login somedecoder:50004 admin ...
netspeed transfer=4g
```

To troubleshoot the connection between a Concentrator and a Decoder, SSH into the Concentrator, run NwConsole, and then log on to the Decoder and run `netspeed`. The output from the command gives you an indication of the maximum network throughput. If it is much less than the standard 1 Gbps interface, it could indicate a network issue.

NetWitness® Platform XDR

Version 12.0

Physical Host Installation Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Introduction **5**
Supported Hardware 5
Endpoint Log Hybrid Host Hardware Specifications 5
NetWitness UEBA Host Hardware Specifications 5
External Attached Storage 6
Physical Host Installation Workflow 6
Self-Help Resources 7
Contact NetWitness Support 7
Installation Tasks **8**
Checklist 8
Install 12.0 on the NetWitness Server (NW Server) Host and Other Component Host 8
Create a Base Image on the RSA Appliance 8
Create a Base Image on the Third Party Server Hardware 10
Install NetWitness Platform 12
Set Up ESA Hosts 20
Install Component Services on Hosts 20
Complete Licensing Requirements 20
(Optional) Install Warm Standby NW Server 20
Update or Install Windows Legacy Collection **21**
Post Installation Tasks **22**
Event Stream Analysis (ESA) 22
Configure Meta Keys on New ESA Hosts to Match Upgraded ESA Hosts in the Same NetWitness Platform Network 22
NetWitness Endpoint 23
Step 1: Install additional Endpoint Log Hybrid 23
Step 2: Setup the Endpoint Log Hybrid 24
Step 3: Switch to the NetWitness UI and add Hosts 24
Do You Need to Install an Endpoint Service onto Separate Hardware 25
Install an Endpoint Service Category on an Existing Log Decoder 25
NetWitness UEBA 26
Deployment Options 30
Appendix A. Troubleshooting **31**
Command Line Interface (CLI) 32
Event Stream Analysis 33

Appendix B. Create an External Repository **34**
Appendix C. Silent Installation Using CLI **36**
Appendix D. Third Party Server System Requirement **39**
Hardware Requirements 39

Introduction

The instructions in this guide apply to physical hosts exclusively. See the *Virtual Host Installation Guide for NetWitness Platform 12.0* for instructions on how to set up virtual hosts in 12.0.

Note: Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Supported Hardware

Series 5, and Series 6.

Refer to the *NetWitness* Hardware Setup Guides for detailed information on each series type (https://community.netwitness.com/t5/netwitness-platform-hardware/tkb-p/netwitness-hardware-documentation).

Endpoint Log Hybrid Host Hardware Specifications

Series 5 (Dell R730) hardware or Series 6 (Dell R740 hardware. See "(Optional) Task 2 - Install Endpoint Log Hybrid" in Post Installation Tasks for instructions on how to install the Endpoint Log Hybrid.

Note: If you have RSA NetWitness Endpoint 4.x hardware, you can re-purpose it for NetWitness Endpoint Log Hybrid 12.0.

NetWitness UEBA Host Hardware Specifications

S5 (Dell R630 appliance) or S6 (Dell R640) hardware. See "(Optional) Task 3 - Install NetWitness UEBA" in Post Installation Tasks for instructions on how to install NetWitness UEBA.

SERIES 5 (DELL R630) SPECIFICATIONS

Specification	Capacity
Model	Dell PowerEdge R630xl
Processor Type	Intel Xeon E5 -2680v3
Processor Speed	2.5 GHz
Cache	30MB
Number of Cores	12
Number of Processors	2
Number of Threads	24
Total Memory	256GB
Internal Disk Controller	Dell PERC H730
External Disk Controller	Dell PERC H830

Specification	Capacity
SAN Connectivity (HBA) - Optional	N/A
Remote Management Card	iDRAC8 Enterprise
Drives	Total - 6 Drives 2 x 1TB, 2.5" HDD 4 x 2TB, 2.5" HDD
Chassis	1U
Weight	18.4 kg (40.5 lbs)
NIC Card*	On Board 2 x 10 Gb Copper 2 x 10 Gb & 2 x 1Gb Copper (Other options are available)
Dimensions	H: 4.28 cm (1.68 in.) x W: 48.23 cm (18.98 in.) x D: 75.51 cm (29.72 in.)
Power	1100W Redundant
BTU/hr	4100 BTU/hr (max)
Amps (Spec)	1100W / 220VAC = 5A
Actual Amp Draw (Post Startup)	2.1 Amps
Events Per Second (EPS)	100K EPS
Throughput	N/A

Note: Network Interface Controller (NIC) card options are available for swap with on-board daughter cards or add ons.

External Attached Storage

If you have external storage devices (for example, DACs or PowerVaults) attached to physical hosts, refer to the Hardware Setup Guides on NetWitness Community (https://community.netwitness.com/t5/netwitness-platform-hardware/tkb-p/netwitness-hardware-documentation) for information on how to configure this storage.

Physical Host Installation Workflow

The following diagram illustrates the NetWitness 12.0 Physical Host Installation workflow.



Self-Help Resources

There are several options that provide you with help as you need it for installing and using NetWitness:

- See the documentation for all aspects of NetWitness here: https://community.netwitness.com/t5/netwitness-platform/ct-p/netwitness-documentation
- Use the **Search** and **Ask it** fields in NetWitness Community to find specific information here: https://community.netwitness.com/
- See the NetWitness Knowledge Base: https://community.netwitness.com/t5/netwitness-knowledge-base/tkb-p/netwitness-knowledge-base
- See Troubleshooting section in the Guides.
- See also NetWitness® Platform Blog Posts.
- If you need further assistance, contact NetWitness Support.

Contact NetWitness Support

If you contact NetWitness Support, you should be at your computer. Be prepared to provide the following information:

- The version number of the NetWitness Platform product or application you are using.
- The type of hardware you are using.

Use the following contact information if you have any questions or need assistance.

NetWitness Community	https://community.netwitness.com/ In the main menu, click **Support** > **Case Portal** > **View My Cases**.
International Contacts (How to Contact NetWitness Support)	https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897
Community	https://community.netwitness.com/t5/support/ct-p/support

Installation Tasks

This topic contains the tasks you must complete to install NetWitness 12.0 on physical hosts.

Checklist

Complete the installation tasks in the following order.

Step	Description	Instructions
1	Install 12.0 on NetWitness hosts.	Install 12.0 on the NetWitness Server (NW Server) Host and Other Component Host
2	Set up ESA hosts.	Set Up ESA Hosts
3	Install component services on your hosts.	Install Component Services on Hosts
4	Complete licensing requirements for services.	Complete Licensing Requirements
5	(Optional) Install warm standby NW Server host.	(Optional) Install Warm Standby NW Server

Caution: Before you begin the installation process, open all your firewall ports. The"Network Architecture and Ports" topic in the *Deployment Guide for NetWitness Platform 12.0* lists all the ports in a deployment. Do not proceed with the installation until the ports on your firewall are configured.

Install 12.0 on the NetWitness Server (NW Server) Host and Other Component Host

Complete the following steps to install 12.0 on NW Server host and other component hosts. Steps that are specific to the NW Server host or to component hosts are noted.

Create a Base Image on the RSA Appliance

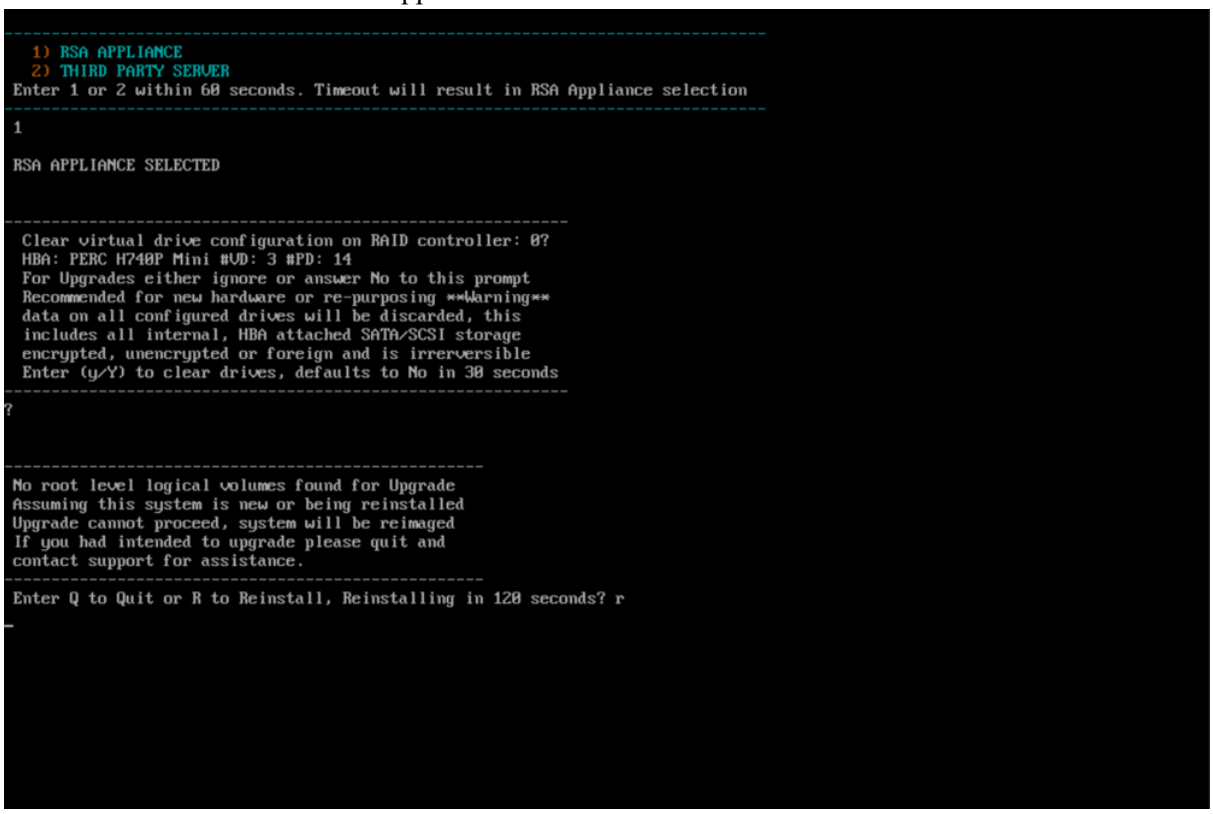
1. Attach media (ISO) to the host.
   See the *USB Build Stick Instructions for NetWitness* for more information. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
   - Hypervisor installations - see the *Virtual Host Installation Guide for NetWitness Platform 12.0*
   - Physical media - use the ISO image to create bootable flash drive media. You can use Rufus or another suitable imaging tool to create a Linux file system on the USB drive. Rufus is available at https://rufus.ie

- iDRAC installations - the virtual media type is:
  - **Virtual Floppy** for mapped flash drives
  - **Virtual CD** for mapped optical media devices or ISO file.

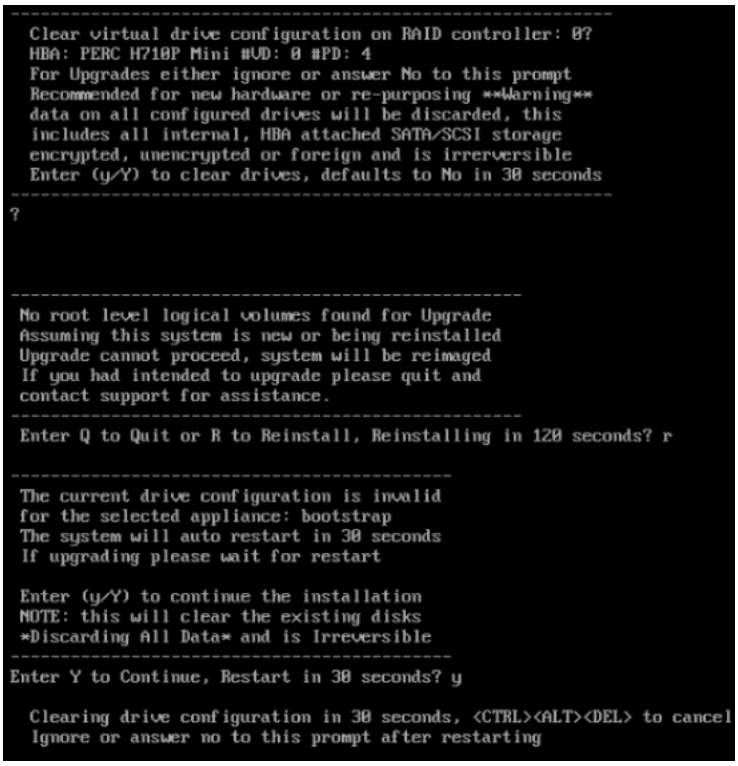
2. Log in to the host and reboot it.

```
login: root
Password:
Last login: Tue Sep 19 13:27:15 on tty1
[root@saserver ~]# reboot
```

3. Select **F11** (boot menu) during reboot to select a boot device and boot to the connected media. After system checks, the **Welcome to NetWitness Platform 12.0** installation menu is displayed.
4. Select **Install NetWitness Platform 12.0** (default selection) and press **Enter**. The Appliance Type selection menu is displayed.
5. You must enter **1** to select RSA appliance.

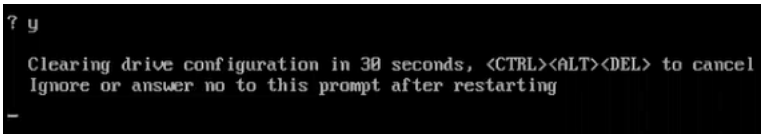


6. The Installation program runs and stops at the Enter (**y**/**Y**) to clear drives prompt that asks you to format the drives.



Caution: You must respond **y** or **Y** to this prompt even if the host does not have an internal RAID configuration or the installation will fail.

7. Type **y** to continue. The default action is No, so if you ignore the prompt, it will select No in 30 seconds, and will not clear the drives.



The system displays all the installation tasks it is performing. This can take a minute or so. After it completes the tasks, the installation program reboots the host.

Caution: Do not reboot with the attached media attached (media that contains the ISO file, for example a build stick).



Create a Base Image on the Third Party Server Hardware

Prerequisites

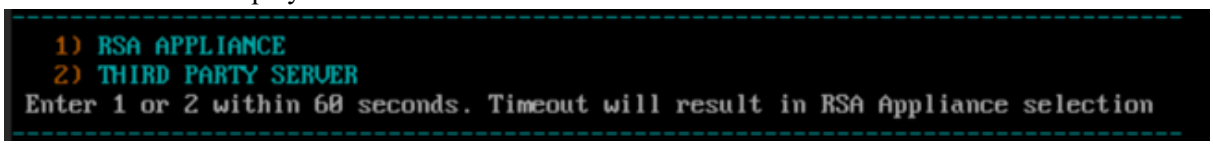
NetWitness recommends that the Third party Server Hardware meets the criteria defined in Appendix D. Third Party Server System Requirement.

1. Attach media (ISO) to the host.
   See the *USB Build Stick Instructions for NetWitness* for more information. Go to the NetWitness All

Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

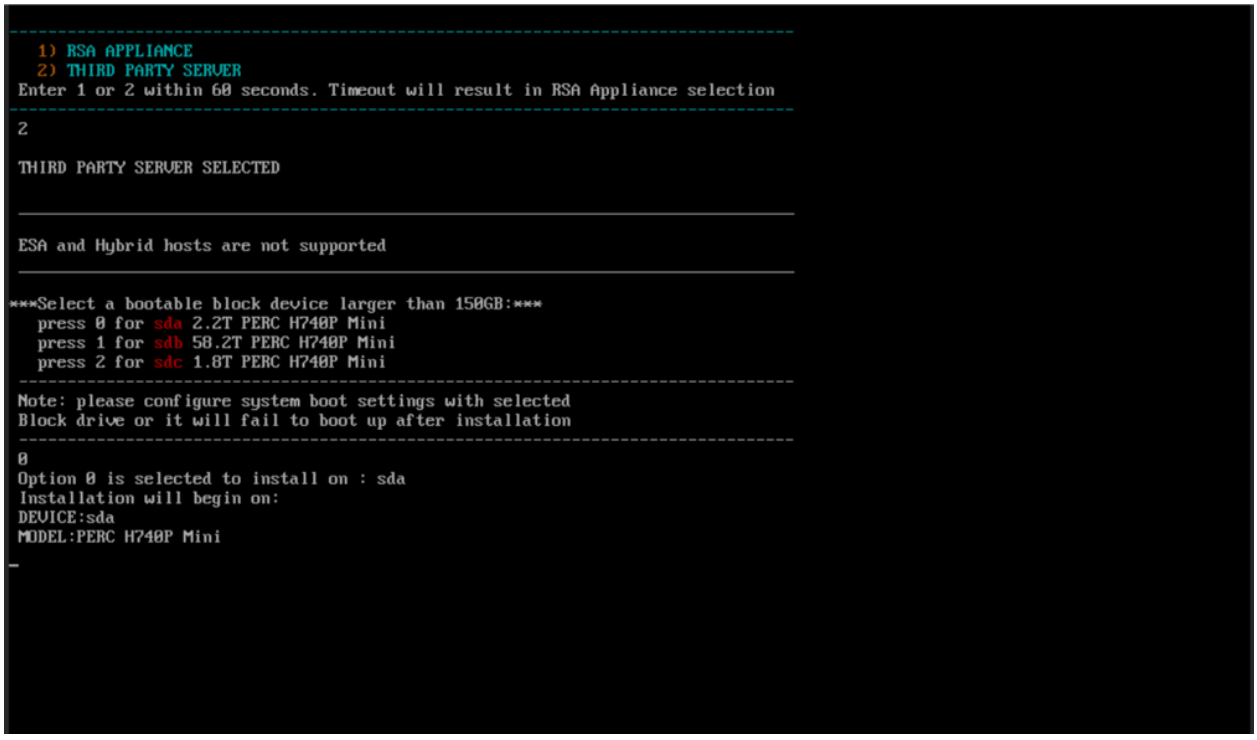
- Physical media - use the ISO image to create bootable flash drive media. You can use Rufus or another suitable imaging tool to create a Linux file system on the USB drive. Rufus is available at https://rufus.ie.

2. Log in to the host and reboot it.
3. Select **F11** (boot menu) during reboot to select a boot device and boot to the connected media. After system checks, the **Welcome to NetWitness Platform 12.0** installation menu is displayed.
4. Select **Install Netwitness Platform 12.0** (default selection) and press **Enter**. The Appliance Type selection menu is displayed.



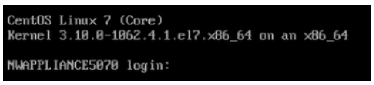
5. Enter **2** to select the Third Party Server.
6. All the available block devices are displayed. Select a block device larger than 150 GB to install the Netwitness Platform.

Note: You must configure the system boot settings with selected block device else the system will not boot after imaging.



7. The system displays all the installation tasks running, and it may take few minutes to complete the installation. Once the installation is complete, the installation program reboots the host.

Caution: Do not reboot with the attached media that contains the ISO file, for example, build stick.



Install NetWitness Platform

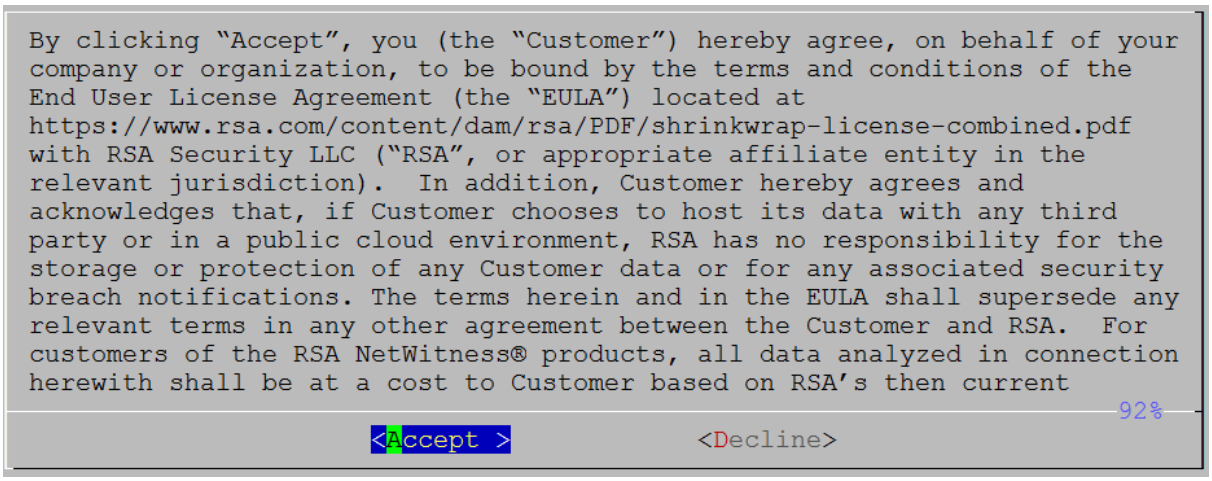
Caution: If you want to install the Endpoint Relay Server, do not run the `nwsetup-tui` script. Follow the instructions in "(Optional) Installing and Configuring Relay Server" in the *NetWitness Endpoint Configuration Guide*.

IMPORTANT: In NetWitness Platform version 11.6 or Later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`.
If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

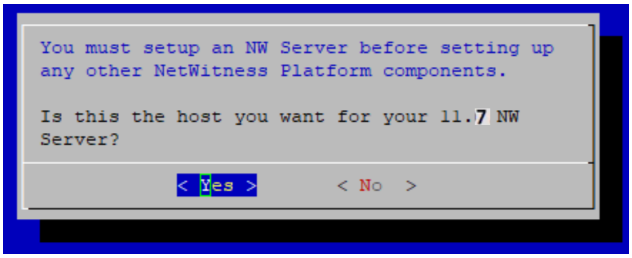
1. Log in to the host with the `root` credentials and run the `nwsetup-tui` command to set up the host.

   This initiates the `nwsetup-tui` (Setup program) and the EULA is displayed.

   Note: Use the following options to navigate the Setup prompts.
   1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, and use the Tab key to move to and from commands (such as **<Yes>**, **<No>**, **<OK>**, and **<Cancel>**). Press **Enter** to register your command response and move to the next prompt.
   2.) The Setup program adopts the color scheme of the desktop or console you use to access the host.
   3.) If you specify DNS servers during the Setup program (`nwsetup-tui`) execution, they MUST be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` script to proceed. Any misconfigured DNS servers cause the Setup program to fail. If you need to reach a DNS server after setup that is unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see "Change Host Network Configuration" topic in the System Maintenance Guide.
   If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Update Repository** prompt in step 12 (the DNS servers are not defined so the system cannot access the external repo).



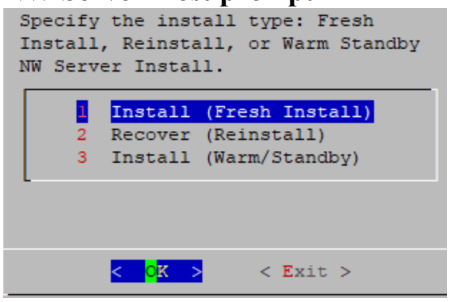
2. Tab to **Accept** and press **Enter**.
   The **Is this the host you want for your 12.0 NW Server** prompt is displayed.



3. Tab to **Yes** and press **Enter** to install 12.0 on the NW Server.
   Tab to **No** and press **Enter** to install 12.0 on other component hosts.

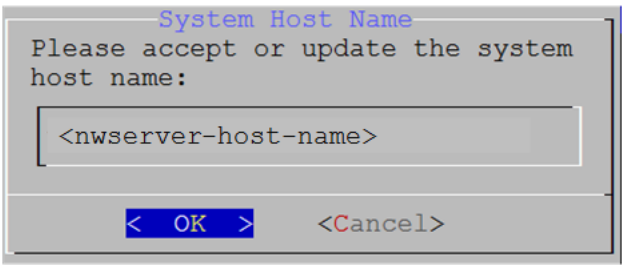
   Caution: If you choose the wrong host for the NW Server and complete the Setup, you must restart the Setup Program (step 2) and complete steps all the subsequent steps to correct this error.

4. The **Install** prompt is displayed (**Recover** does not apply to the installation.).
   NW Server Host prompt:



   Other Component Hosts, the prompt is the same, but does not include option 3 Install (Warm/Standby)

5. Press **Enter**. **Install (Fresh Install)** is selected by default.
   The **System Host Name** prompt is displayed.
   NW Server prompt:

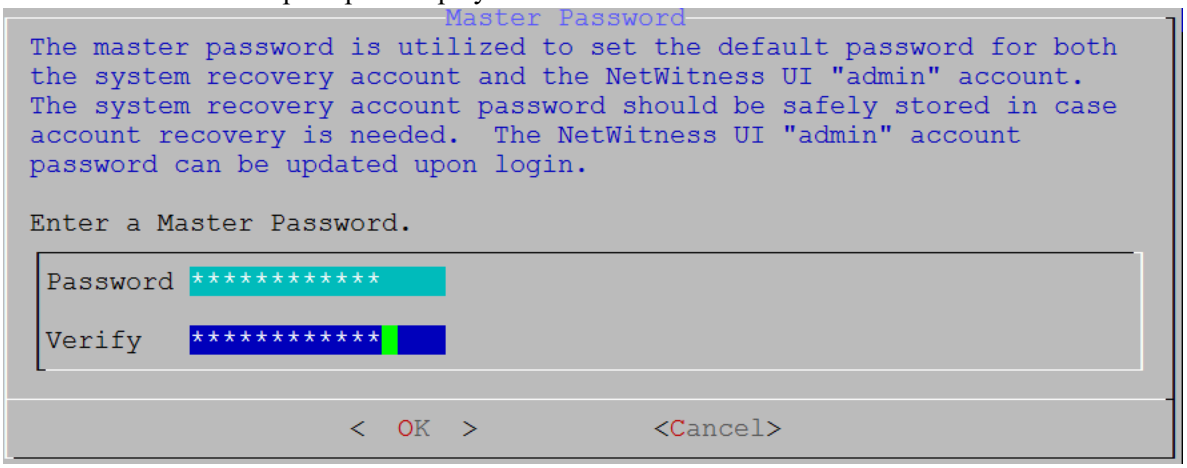


Other Component Hosts prompt says <non-nwserver-host-name>

Caution: If you include "." in a host name, the host name must also include a valid domain name.

Press **Enter** if want to keep this name. If not, edit the host name, tab to **OK**, and press **Enter** to change it.

6. **This step applies only to NW Server hosts**.
   The **Master Password** prompt is displayed.



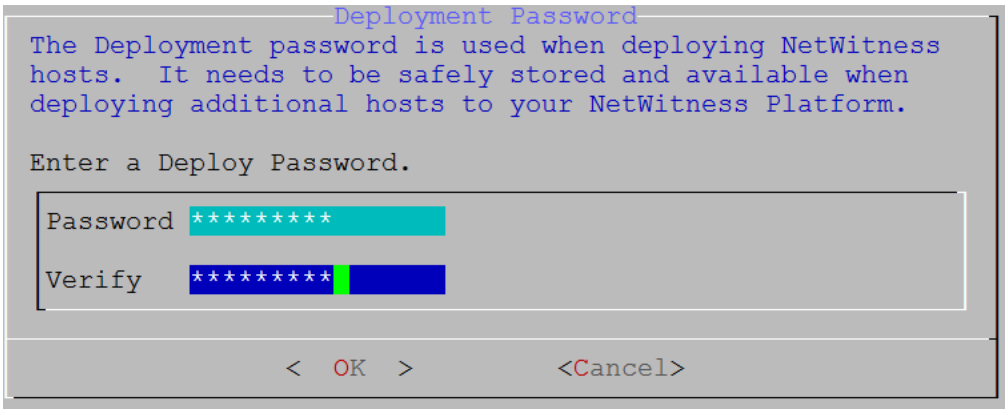
   The following list of characters are supported for Master Password and Deployment Password:

   - Symbols: **! @ # % ^ +**
   - Numbers: **0-9**
   - Lowercase Characters: **a-z**
   - Uppercase Characters: **A-Z**

   No ambiguous characters are supported for Master Password and Deployment Password. For example:
   space { } [ ] ( ) / \ ' " ` ~ ; : . < > -

   Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

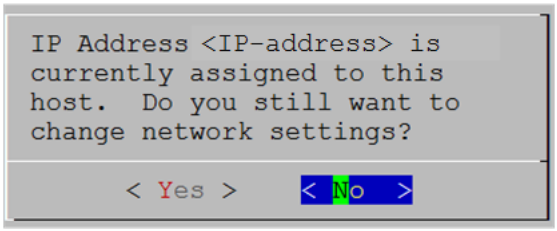
7. **This step applies to both NW Server hosts and component hosts.**
   The **Deployment Password** prompt is displayed.



Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

8. One of the following conditional prompts is displayed.

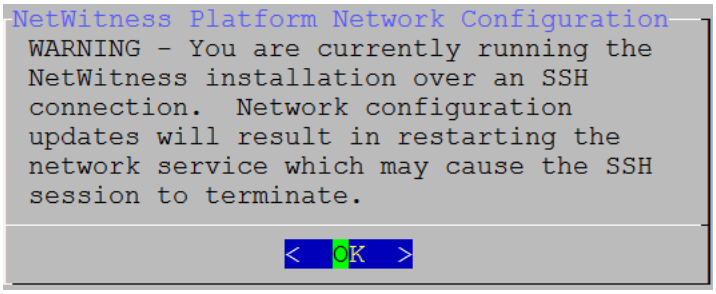
- If the Setup program finds a valid IP address for this host, the following prompt is displayed.



  Press **Enter** if you want to use this IP and avoid changing your network settings. Tab to **Yes** and press **Enter** if you want to change the IP configuration on the host.

- If you are using an SSH connection, the following warning is displayed.

  Note: If you connect directly from the host console, the following warning is not displayed.

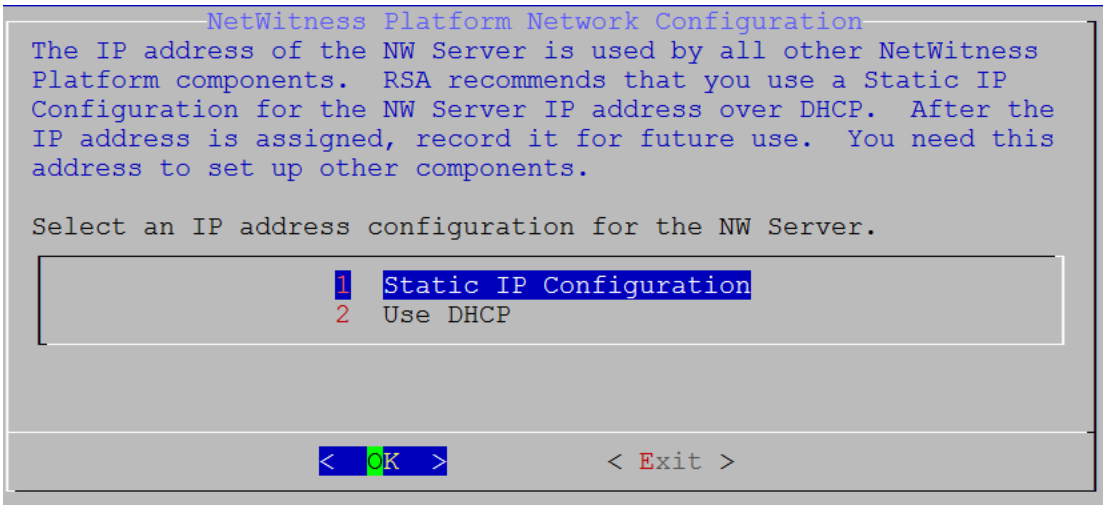


  Press **Enter** to close warning prompt.

- If the Setup Program finds an IP configuration and you choose to use it, the **Update Repository** prompt is displayed. Go to step 12 and complete the installation.

- If the Setup Program did not find an IP configuration or if you choose to change the existing IP configuration, the **Network Configuration** prompt is displayed.

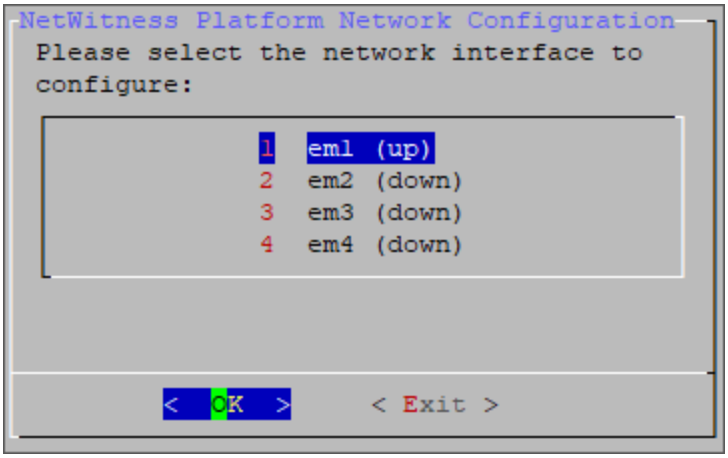
  Caution: Only select "**Use DHCP**" as an IP address configuration for the NW Server if DHCP issues static IP addresses.



Tab to **OK** and press **Enter** to use **Static IP**.

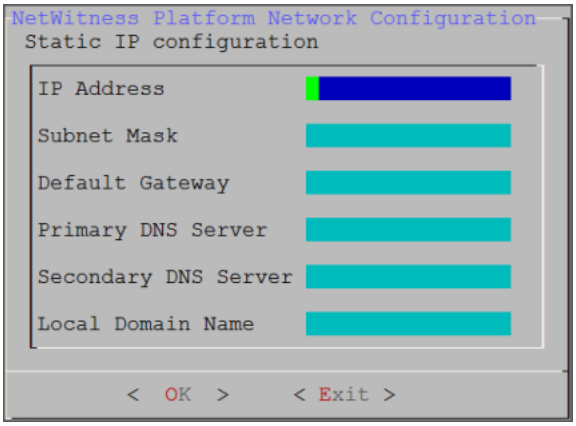
If you want to use DHCP, down arrow to **2 Use DHCP** and press **Enter**.

The **Network Configuration** prompt is displayed.



9. Down arrow to the network interface you want, tab to **OK**, and press **Enter**. If you do not want to continue, tab to **Exit.**

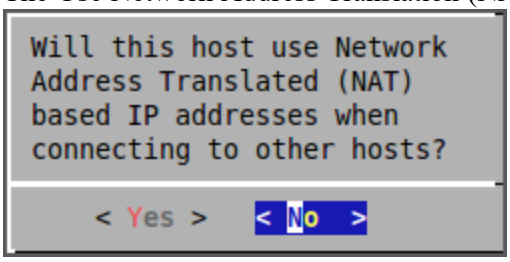
   The following **Static IP Configuration** prompt is displayed.



10. Type the configuration values, tab to **OK**, and press **Enter**. If you do not complete all the required fields, an `All fields are required` error message is displayed (**Secondary DNS Server** and **Local Domain Name** fields are not required). If you use the wrong syntax or character length for any of the fields, an `Invalid <field-name>` error message is displayed.

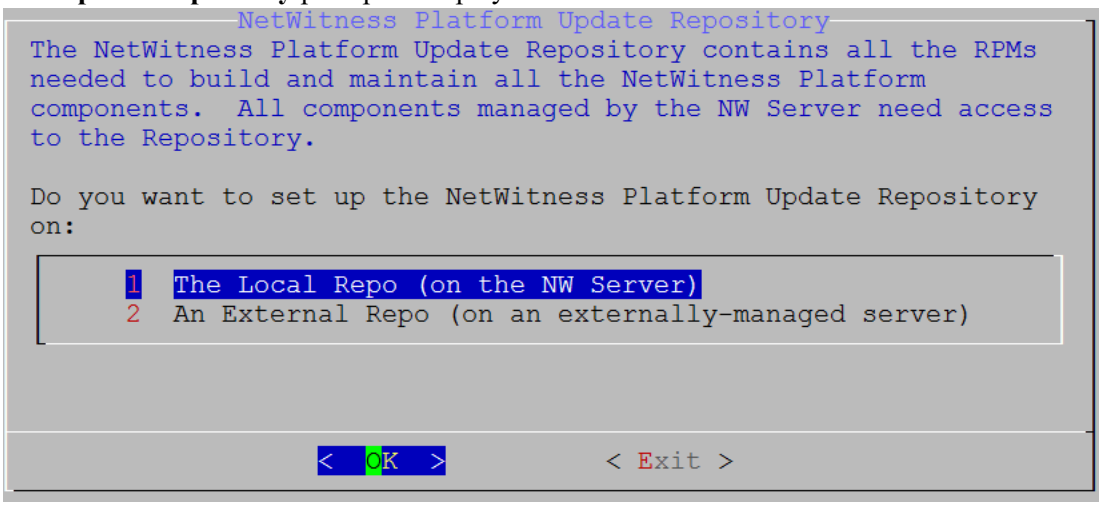
    Caution: If you select **DNS Server**, make sure that the DNS Server is correct and the host can access it before proceeding with the installation.

11. The Use Network Address Translation (NAT) prompt is displayed.



    For the NW Server, tab to **No** and press **Enter**.
    For component hosts, if this host requires the use of NAT-based addresses to communicate with the NW Server, tab to **Yes**. Otherwise, tab to **No** and press **Enter**.

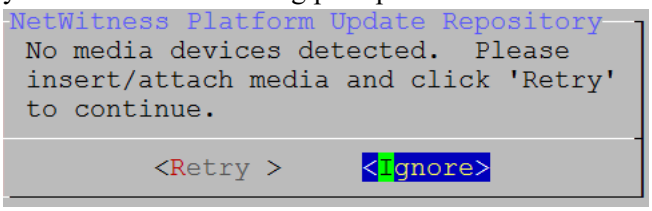
12. The **Update Repository** prompt is displayed.



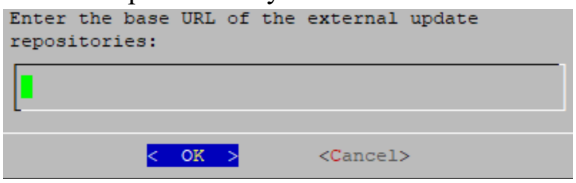
    For the NW Server:

    - Press **Enter** to choose the **Local Repo**.
    - If you want to use an external repo, down arrow to **External Repo**, tab to **OK**, and press **Enter**. If you select **1 The Local Repo (on the NW Server)** in the Setup program, make sure that you have the appropriate media attached to the host (media that contains the ISO file, for example a build stick) from which it can install NetWitness 12.0. If the program cannot find the attached media,

you receive the following prompt.



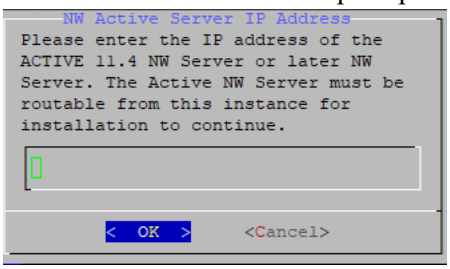
- If you select **2 An External Repo (on an externally-managed server)**, the UI prompts you for a URL. The repositories give you access to RSA updates and CentOS updates. Refer to "Appendix B. Create an External Repo" in this guide for instructions on how to create this repo and its external repo URL so you can enter it in the following prompt.



Enter the base URL of the NetWitness external repo and click **OK.** The **Start Install** prompt is displayed.

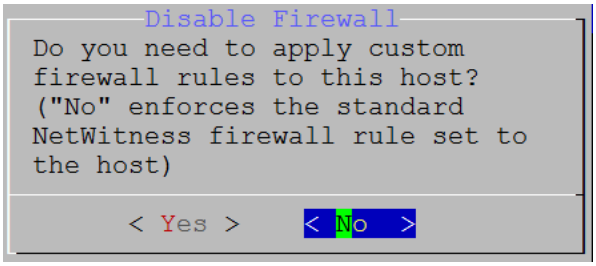
For component hosts:

- Select the same repo that you selected when you installed the NW Server host and follow the steps above.
- The NW Server IP Address prompt is displayed.



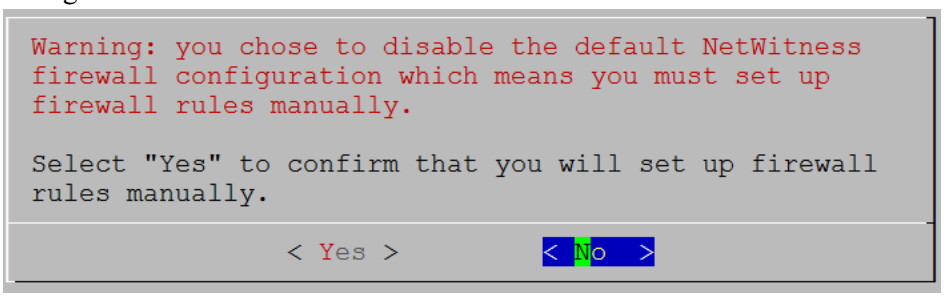
Type the NW Server IP address. Tab to **OK** and press **Enter**.

13. The Disable firewall prompt is displayed.

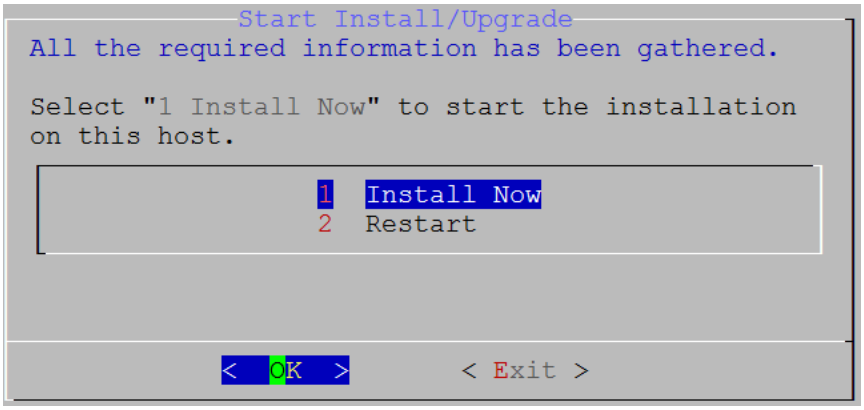


Tab to **No** (default), and press **Enter** to use the standard firewall configuration.
To disable the standard firewall configuration, tab to **Yes**, and press **Enter**.
If you select **Yes**, confirm your selection(select **Yes** again) or select **No** to use the standard firewall

configuration.

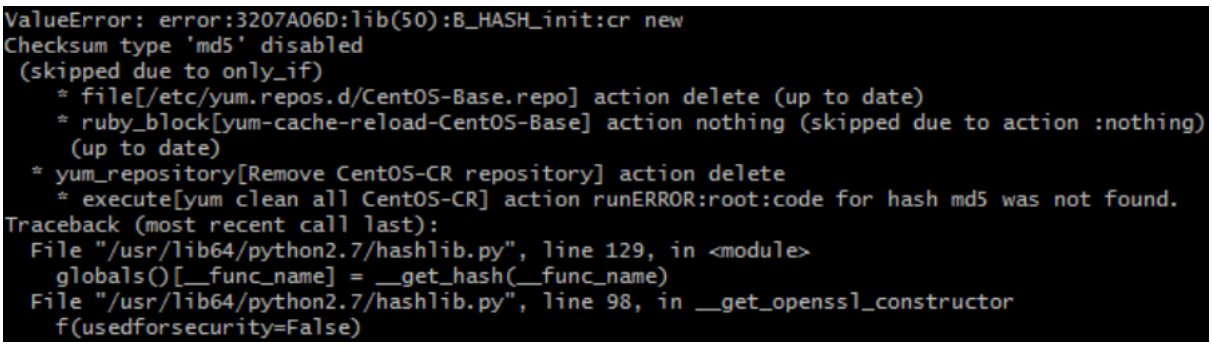


14. The **Start Install** prompt is displayed.



15. Press **Enter** to install 12.0.
When **Installation complete** is displayed, you have installed 12.0 on this host.

Note: Ignore the hash code errors similar to the errors shown in the following figure that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.



16. (Optional) If your system configuration requires that a component host must use a NAT IP address to reach the NW Server host, you must configure the NAT IP address of the NW Server by running the following command:
`nw-manage --update-host --host-id <NW Server Host UUID> --ipv4-public <NAT IP address>`

Set Up ESA Hosts

After you install your NW Server and component hosts, follow these steps to set up your ESA hosts.

- Install your primary ESA host following the instructions in "Install 12.0 on the NetWitness Server (NW Server) Host and Other Component Hosts" in this guide, and install the **ESA Primary** service on it after you finish the Set Up program in the UI in  (Admin) > **Hosts** > :
- (Conditional) If you have a secondary ESA host, install it and install the **ESA Secondary** service on it after you finish the Set Up program in the UI in  (Admin) > **Hosts** > :

Install Component Services on Hosts

After you have installed NW Server and component hosts, and set up your ESA hosts, follow these steps to install component services, such as Decoders and Concentrators, on your host systems.

1. Install a component service on the host.
   a. Log into NetWitness and go to (missing or bad snippet)> **Hosts**.
   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background.

   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the host in the **New Hosts** dialog and click **Enable**.
   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the **Hosts** view and click .
   The **Install Services** dialog is displayed.

   d. Select the appropriate host type (for example, **Concentrator**) in **Category** and click **Install**.

Complete Licensing Requirements

Complete licensing requirements for installed services. See the *NetWitness Platform 12.0 Licensing Management Guide* for more information. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

(Optional) Install Warm Standby NW Server

Refer to "Warm Standby NW Server Host" under "Deployment Option Setup Procedures" in the *Deployment Guide for NetWitness Platform 12.0* for instructions on how to set up a Warm Standby NW Server.

Update or Install Windows Legacy Collection

Refer to the Windows Legacy Collection Guide for NetWitness

Note: After you update or install Windows Legacy Collection, reboot the system to ensure that Log Collection functions correctly.

Post Installation Tasks

This topic contains the tasks you complete after you install 12.0.

- Event Stream Analysis (ESA)
- NetWitness Endpoint
- NetWitness UEBA

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Event Stream Analysis (ESA)

Configure Meta Keys on New ESA Hosts to Match Upgraded ESA Hosts in the Same NetWitness Platform Network

If you have one or more ESA hosts in a NetWitness Platform network, which were upgraded from a version before 11.3.0.2 to 11.7, and you add a new ESA host, you must configure the meta keys on the new ESA host to match the other ESA hosts. All ESA Correlation services on the same NetWitness Platform network must have the same Meta Key configurations.

1. For each ESA Correlation service on an upgraded ESA host and for the ESA Correlation service on the newly installed ESA host:
   a. Open a new tab, go to  (Admin) > **Services**, and in the Services view, select the ESA Correlation service and then select  > **View** > **Explore**.
   b. In the Explore view node list for the ESA Correlation service, select **correlation** > **stream**.
2. Ensure that the **`multi-valued`** and **`single-valued`** meta key values are the same on each of the upgraded ESA Correlation services.
3. Ensure that the **`multi-valued`** and **`single-valued`** meta key values on the newly installed ESA host are the same as those on the upgraded services.
4. To apply any changes on the ESA Correlation services, go to  **(Configure)** > **ESA Rules** and click the **Settings** tab. In the Meta Key References, click the **Meta Re-Sync (Refresh)** icon ().
5. If you updated the ESA Correlation services, redeploy the ESA rule deployments.

For more information, see "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*.

NetWitness Endpoint

The tasks in this section only apply to customers that use the NetWitness Endpoint component of NetWitness Platform.

Install Endpoint Log Hybrid

Depending on the number of agents and the location of the agents, you can choose to deploy a single Endpoint Log Hybrid host or multiple Endpoint Log Hybrid hosts. To deploy a host, you provision it and install a category on it.

- **Single Endpoint Log Hybrid host** - Deploy NetWitness Server host, Endpoint Log Hybrid host, and ESA host or hosts.
- **Multiple Endpoint Log Hybrid hosts** - Deploy NetWitness Server host, ESA host or hosts, Endpoint Log Hybrid hosts. You can deploy up to 6 Endpoint Log Hybrid hosts. For a consolidated view of all endpoint data from multiple Endpoint Log Hybrid hosts, install the Endpoint Broker. You can add only one broker in a NetWitness platform deployment which serves upto 6 Endpoint Log Hybrid hosts.

  Note: NetWitness recommends that you co-locate the Endpoint Broker on the NetWitness Broker host. However, you can deploy the Endpoint Broker on a separate host or co-locate it on the Endpoint Log Hybrid.

  Note: You must plan to scale your ESA deployment to support multiple Endpoint Log Hybrid hosts.

Follow these steps to deploy an Endpoint Log Hybrid host.

Complete the following steps first:

- For a physical host, complete steps 1 - 16 in "Install NetWitness Platform" under Installation Tasks in the *Physical Host Installation Guide for NetWitness Platform*
- For a virtual host, complete steps 1 - 16 in "Step 4. Install NetWitness Platform" under Install NetWitness Platform Virtual Host in Virtual Environment in the *Virtual Host Installation Guide for NetWitness Platform 12.0*

Configuring Multiple Endpoint Log Hybrids

Follow these steps to install another Endpoint Log Hybrid.

Step 1: Install additional Endpoint Log Hybrid

- To install a physical host, complete steps 1 - 16 in "Install NetWitness Platform" under Installation Tasks in the *Physical Host Installation Guide for NetWitness Platform 12.0*

- To install a virtual host, complete steps 1 - 16 in "Step 4. Install NetWitness Platform" under Install NetWitness Platform Virtual Host in Virtual Environment in the *Virtual Host Installation Guide for NetWitness Platform 12.0*

Step 2: Setup the Endpoint Log Hybrid

1. Create a directory `mkdir -p /etc/pki/nw/nwe-ca`.
2. Copy the following certificates from the first Endpoint Log Hybrid to the secondary Endpoint Log Hybrid:

   `/etc/pki/nw/nwe-ca/nwerootca-cert.pem`

   `/etc/pki/nw/nwe-ca/nwerootca-key.pem`

   Note: NetWitness recommends that you copy certificates from Endpoint Log Hybrid to secondary Endpoint Log Hybrid CentOS to Windows using the `SCP` command to avoid any corruption caused by Antivirus or third-party tools.

Step 3: Switch to the NetWitness UI and add Hosts

- See Add Hosts to the Endpoint Log Hybrid: for more information.

Add Hosts to the Endpoint Log Hybrid:

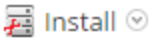
1. Log into NetWitness Platform and click  **(Admin) > Hosts**.

   The New Hosts dialog is displayed with the Hosts view grayed out in the background.

   Note: If the New Hosts dialog is not displayed, click **Discover** in the **Hosts** view toolbar.

2. Select the host in the **New Hosts** dialog and click **Enable**.

   The New Hosts dialog closes and the host is displayed in the Hosts view.

3. Select that host in the **Hosts** view (for example, **Endpoint**) and click 

.

   The Install Services dialog is displayed.

4. Select the **Endpoint Log Hybrid** category and click **Install**.
5. Make sure that the Endpoint Log Hybrid service is running.
6. Configure Endpoint Meta forwarding.

   See the *Endpoint Configuration Guide* for instructions on how to configure Endpoint Meta forwarding.

7. Deploy the ESA Rules from the Endpoint Rule Bundle. For more information, see "Deploy Endpoint Risk Scoring Rules on ESA" section in the ESA Configuration Guide.

   Note: The Endpoint IIOCs are available as OOTB Endpoint Application rules.

8. Review the default policies and create groups to manage your agents. See *Endpoint Configuration Guide*.

   Note: In 11.3 or later, agents can operate in Insights or Advanced mode depending on the policy configuration. The default policy enables the agent in an advanced mode. If you want to continue to use the Insights agent, before updating, review the policy, and make sure that the Agent mode is set to Insights.

9. Install the Endpoint Agent. You can install an Insights (free version) or an Advanced agent (licensed). See *Endpoint Agent Installation Guide* for detailed instructions on how to install the agent.

   Note: You can migrate the Endpoint Agent from 4.4.0.x to 11.7. For more information, see the *NetWitness Endpoint 4.4.0.x to NetWitness Platform 12.0 Migration Guide*.

(Optional) Configure an Endpoint Service on an Existing Log Decoder Host

You can install an Endpoint service category on an existing Log Decoder host. For an overview of installing service categories on hosts, see "Hosts and Services Set Up Procedures" in the *Host and Services Getting Started Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

- If you have an existing Endpoint Log Hybrid, you must copy certificates from that Endpoint Hybrid host to the Log Decoder before you install the Endpoint service category on the Log Decoder.
- If you do not have an Endpoint Log Hybrid host, you do not need to copy over the certificates before you install the Endpoint service category on the Log Decoder.

Do You Need to Install an Endpoint Service onto Separate Hardware

If you are only using NW Platform for collecting and analyzing logs, you can co-locate your Endpoint Server on the same physical hardware as your Log Decoder. For more information, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness Platform*.

If you exceed these guidelines, the amount of disk space usage and CPU might become so high as to create alarms for your Endpoint Server in Health and Wellness. If you notice this, and are running both log collection and EDR scans, you can use Throttling to control the amount of data coming into the Log Decoder.

If that doesn't help, NetWitness recommends that you move your Endpoint Server onto separate hardware from that used by your Log Decoder.

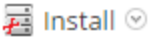
Install an Endpoint Service Category on an Existing Log Decoder

To install an Endpoint service category on an existing Log Decoder if you have an existing Endpoint Log Hybrid:

1. Create a directory `mkdir -p /etc/pki/nw/nwe-ca`.
2. Copy the following certificates from the first Endpoint Log Hybrid to the Log Decoder on which you are going to install the additional **Endpoint** service category.

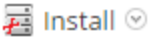
   Note: NetWitness recommends that you copy certificates from Endpoint Log Hybrid to secondary Endpoint Log Hybrid using the `SCP` command to avoid any corruption caused by Antivirus or third-party tools.

   `/etc/pki/nw/nwe-ca/nwerootca-cert.pem`

   `/etc/pki/nw/nwe-ca/nwerootca-key.pem`
3. Log into NetWitness Platform and click  **(Admin) > Hosts**.
4. Select the Log Decoder host in the **Hosts** view and click .

   The Install Services dialog is displayed.
5. Select the **Endpoint** category and click **Install**.

To install an Endpoint service category on an existing Log Decoder if you do not have an existing Endpoint Log Hybrid:

1. Log into NetWitness Platform and click  **(Admin) > Hosts**.
2. Select the Log Decoder host in the **Hosts** view and click .

   The Install Services dialog is displayed.
3. Select the **Endpoint** category and click **Install**.

NetWitness UEBA

The tasks in this section only apply to customers that use the UEBA component of NetWitness Platform.

Install UEBA

To set up NetWitness UEBA in NetWitness Platform 11.7, you must install and configure the NetWitness UEBA service.

The following procedure shows you how to install the NetWitness UEBA service on a NetWitness UEBA Host Type and configure the service.

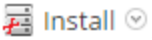
1. For:
   - A physical host, complete steps 1 - 16 in "Install NetWitness Platform" under "Installation Tasks" in the *Physical Host Installation Guide for NetWitness Platform 11.7*.
   - A virtual host, complete steps 1 - 16 in "Step 4. Install NetWitness Platform" in the *Virtual Host Installation Guide for NetWitness Platform 11.7*.

   Note: The Kibana and Airflow webserver User Interface password is the same as the deploy admin password. Make sure that you record this password and store it in a safe location.

2. Log into NetWitness Platform and go to  **(Admin) > Hosts**.
   The New Hosts dialog is displayed with the Hosts view grayed out in the background.

   Note: If the New Hosts dialog is not displayed, click **Discover** in the **Hosts** view toolbar.

3. Select the host in the **New Hosts** dialog and click **Enable**.
   The New Hosts dialog closes and the host is displayed in the Hosts view.

4. Select that host in the **Hosts** view (for example, **UEBA**) and click .
   The Install Services dialog is displayed.

5. Select the **UEBA** Host Type and click **Install**.

6. Make sure that the UEBA service is running.

7. Complete licensing requirements for NetWitness UEBA.
   See the *Licensing Management Guide* for more information.

   Note: NetWitness Platform supports the User and Entity Behavior Analytics License (UEBA). This license is used based on the number of users. The Out-of-the-Box Trial License is a 90-day trial license. In case of UEBA licenses, the 90-day trial period begins from the time the UEBA service deployed on the NetWitness Platform product.

Configure NetWitness UEBA

To start running UEBA:

1. Define the following parameters: data schemas, data source (NetWitness Broker or Concentrator) and start date.

   a. Define UEBA schemas:
      Choose schemas from the following list:

      AUTHENTICATION, FILE, ACTIVE_DIRECTORY, PROCESS, REGISTRY and TLS.

      Note: The TLS packet requires adding the hunting package and enabling the JA3 feature. For more information regarding events that each schema contains, see the *NetWitness UEBA Configuration Guide*.

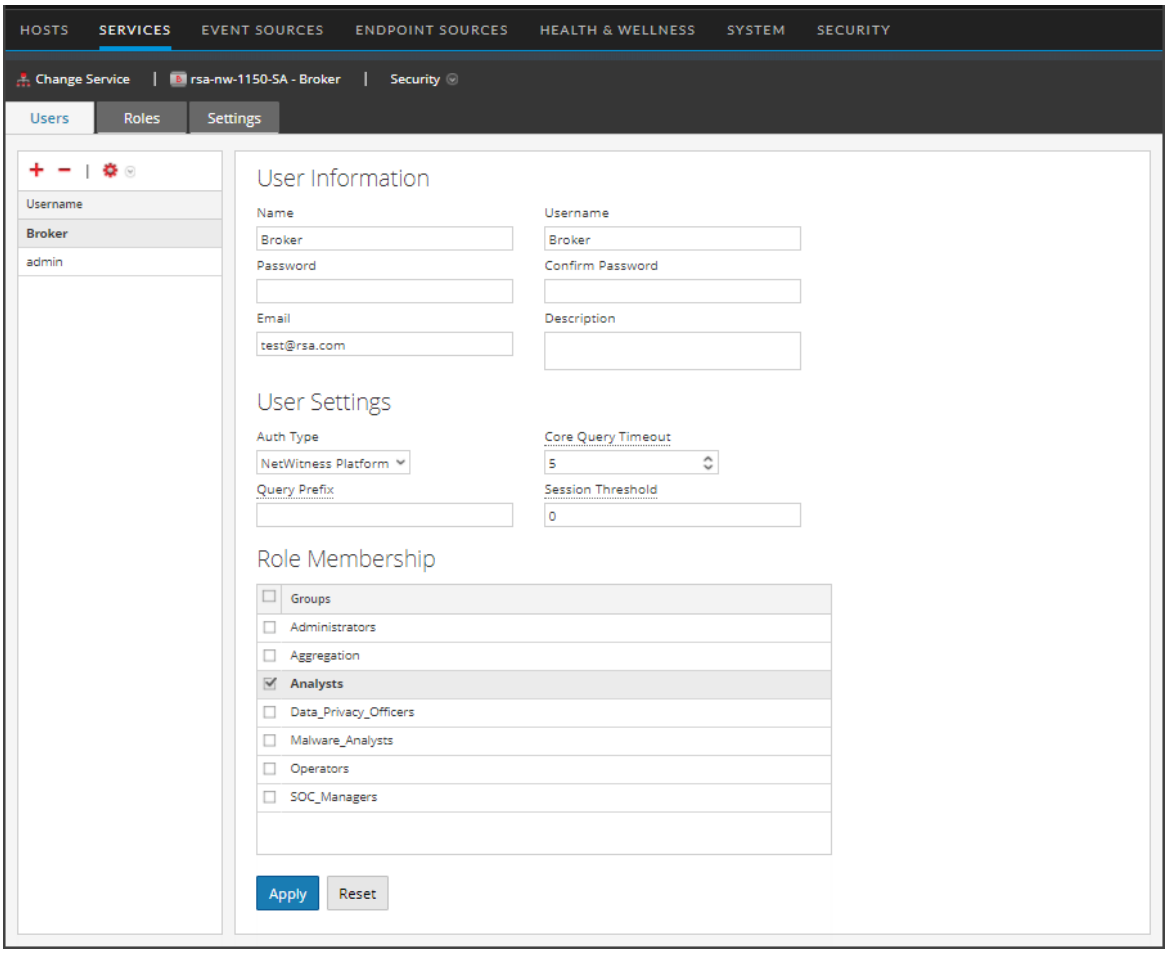
   b. Define the data source:
      If your deployment has multiple Concentrators, we recommend that you assign a Broker at the top of your deployment hierarchy for the NetWitness UEBA data source.

   c. Define the UEBA start-date:

      Note: The selected start date must contain events from all configured schemas.

      NetWitness recommends that the UEBA start date is set to 28 days earlier than the current date. For UEBA systems that intend to process TLS data, you must make sure that the start date is set to no later than 14 days earlier than the current date.

2. . Create a user account for the data source (Broker or Concentrator) to authenticate to the data.
   a. Log into NetWitness Platform.
   b. Go to (Admin) **> Services**.
   c. Locate the data source service (Broker or Concentrator).
      Select that service, and select (Actions) > **View** > **Security**.
   d. Create a new user and assign the "Analysts" role to that user.
      The following example shows a user account created for a Broker.



3. SSH to the NetWitness UEBA server host.
4. Submit the following commands with the above parameters that you already defined.
   `/opt/rsa/saTools/bin/ueba-server-config -u <user> -p <password> -h <host> -o <type> -t <startTime> -s <schemas> -v -e <argument>`
   Where:

Argument	Variable	Description
-u	<user>	User name of the credentials for the Broker or Concentrator instance that you are using as a data source.
-p	<password>	Password of the credentials for the Broker or Concentrator instance that you are using as a data source. The following special characters are supported in a password. !"#$%&()*+,-:;<=>?@[\]^_`\{\|} If you want to include a special character or special characters, you must delimit the password with an apostrophe sign, for example: sh /opt/rsa/saTools/bin/ueba-server-config -u brokeruser -p **'!"UHfz?@ExMn#$'** -h 10.64.153.104 -t 2018-08-01T00:00:00Z -s 'AUTHENTICATION FILE ACTIVE_DIRECTORY TLS PROCESS REGISTRY' -o broker -v
-h	<host>	IP address of the Broker or Concentrator used as the data source. Currently, only one data source is supported.
-o	<type>	Data source host type (broker or concentrator).
-t	<startTime>	Historical start time as of which you start collecting data from the data source in YYYY-MM-DDTHH-MM-SSZ format (for example, 2018-08-15T00:00:00Z). **Note:** The script interprets the time you enter as UTC (Coordinated Universal Time) and it does not adjust the time to your local time zone.
-s	<schemas>	Array of data schemas. If you want to specify multiple schemas, use a space to separate each schema (for example, AUTHENTICATION FILE ACTIVE_ DIRECTORY PROCESS REGISTRY TLS).
-v		**verbose mode**.
-e	<argument>	Boolean Argument. This enables the UEBA indicator forwarder to Respond. **Note:** If your NetWitness deployment includes an active Respond server, you can transfer NetWitness UEBA indicators to the Respond server and create incidents by enabling the indicator forwarder, from this data. For more information on how to enable the NetWitness UEBA incidents aggregation, see Enable User Entity Behavior Analytics Incident Rule.

5. If you are deploying a hot fix on 11.x.x.x version, you must do the following:
   a. Run the presidio-upgrade DAG.
   b. Press the play sign next to the DAG and then click the trigger button.

Enable Access Permission for the NetWitness UEBA User Interface

After you install NetWitness UEBA 11.7, you need to assign the UEBA_Analysts and Analysts roles to the UEBA users. For more information, see 'Assign User Access to UEBA' topic in the *NetWitness UEBA Configuration Guide*. After this configuration, UEBA users can access the **Investigate** > **Users** view.

Note: To complete NetWitness UEBA configuration according to the needs of your organization, See the *NetWitness UEBA Configuration Guide*.

Deployment Options

NetWitness Platform has the following deployment options. See the *NetWitness Deployment Guide* for detailed instructions on how to deploy these options.

- **Analyst User Interface** - gives you access to a subset of features in the NetWitness Platform UI that you can set up in individual locations when you deploy NetWitness Platform in multiple locations. It is designed to reduce latency and improve the performance that can occur when accessing all functionality from the Primary User Interface on the NW Server Host (Primary UI).
- **Group Aggregation** - configures multiple Archiver or Concentrator services as a group and share the aggregation tasks between them.
- **New Health and Wellness Search** - New Health and Wellness is an advanced monitoring and alerting system that provides insights on the operational state of the host and services in your deployment, and helps identify potential issues.
- **Hybrid Categories on Series 6 (R640) Hardware** - installs Hybrid Categories such as Log Hybrid and Network (Packet) Hybrid service categories on a Series 6 (R640) Physical host. This gives you the ability to attach multiple PowerVault external storage devices to the Series 6 (R640) Physical host.
- **NW Server Deployment on ESA Hardware** - installs the NW Server host on NetWitness Series 5 and Series 6 Analytics hardware. The Series 6 Analytics Hardware has more memory and storage capacity than the standard Core appliance on which NW Server has typically been deployed. This results in better overall responsiveness and larger retention capacity for Report Engine.
- **Second Endpoint Server** - deploys a second Endpoint Server.
- **Warm Standby NW Server** - duplicates the critical components and configurations of your active NW Server Host to increase reliability.

Appendix A. Troubleshooting

This section describes solutions to problems that you may encounter during installations and upgrades. In most cases, NetWitness creates log messages when it encounters these problems.

Note: If you cannot resolve an upgrade issue using the following troubleshooting solutions, contact Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897).

This section has troubleshooting documentation for the following services, features, and processes.

- Command Line Interface (CLI)
- Event Stream Analysis

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Command Line Interface (CLI)

Error Message	Command Line Interface (CLI) displays: "Orchestration failed." `Mixlib::ShellOut::ShellCommandFailed: Command execution failed. STDOUT/STDERR suppressed for sensitive resource in/var/log/netwitness/config-management/chef-solo.log`
Cause	Entered the wrong `deploy_admin` password in `nwsetup-tui`.
Solution	Retrieve your `deploy_admin` password. 1. SSH to the NW Server host. `security-cli-client --get-config-prop --prop-hierarchy nw.security-client --prop-name deployment.password` SSH to the host that failed. 2. Run the `nwsetup-tui` again using correct `deploy_admin` password.

Error Message	`ERROR com.rsa.smc.sa.admin.web.controller.ajax.health. AlarmsController - Cannot connect to System Management Service`
Cause	NetWitness sees the Service Management Service (SMS) as down after successful upgrade even though the service is running.
Solution	Restart SMS service. `systemctl restart rsa-sms`

Error Message	You receive a message in the User Interface to reboot the host after you update and reboot the host offline. 
Cause	You cannot use CLI to reboot the host. You must use the User Interface.
Solution	Reboot the host in the Host View in the User Interface.

Event Stream Analysis

For ESA Correlation troubleshooting information, see the *Alerting with ESA Correlation Rules User Guide*.

Appendix B. Create an External Repository

Complete the following procedure to set up an external repository (Repo).

Note: 1.) You need an unzip utility installed on the host to complete this procedure. 2.) You must know how to create a web server before you complete the following procedure.

1. (Conditional) Complete this step if you have an external repo and you want to override it.
   - Case 1: You bootstrapped the host from an external repo and you want to upgrade using a local repo on the Admin Server.
     a. Create the `/etc/netwitness/platform/repobase` file.
        `vi /etc/netwitness/platform/repobase`
     b. Edit the `repobase` file so that the only information in the file is the following URL.
        `https://nw-node-zero/nwrpmrepo`
     c. Complete the instructions on how to run the upgrade using the `upgrade-cli-client` tool.
   - Case 2: You bootstrapped the host from local repo on the Admin server (NW Server host) and you want to use an external repo for the upgrade.
     a. Create the `/etc/netwitness/platform/repobase` file.
        `vi /etc/netwitness/platform/repobase`
     b. Edit the `repobase` file so that the only information in the file is the following URL.
        `https://<webserver-ip>/<alias-for-repo>`
     c. Complete the instructions on how to run the upgrade using the `upgrade-cli-client` tool. The instructions are in "Appendix A. Offline Method (No Connectivity to Live Services) - Command Line Interface" in the *Upgrade Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
2. Set up the external repo.
   a. Log in to the web server host.
   b. Create directory to host the NW repository (`netwitness-11.7.0.0.zip`), for example `ziprepo` under `web-root` of the web server. For example, /var/netwitness is the web-root, run the following command string.
      `mkdir -p /var/netwitness/<your-zip-file-repo>`
   c. Create the 11.7.0.0 directory under `/var/netwitness/<your-zip-file-repo>`.
      `mkdir -p /var/netwitness/<your-zip-file-repo>/11.7.0.0`
   d. Create the `OS` and `RSA` directories under `/var/netwitness/<your-zip-file-repo>/11.7.0.0`.
      `mkdir -p /var/netwitness/<your-zip-file-repo>/11.7.0.0/OS`
      `mkdir -p /var/netwitness/<your-zip-file-repo>/11.7.0.0/RSA`

e. Unzip the `netwitness-11.7.0.0.zip` file into the `/var/netwitness/<your-zip-file-repo>/11.7.0.0` directory.
`unzip netwitness-11.7.0.0.zip -d /var/netwitness/<your-zip-file-repo>/11.7.0.0`
Unzipping `netwitness-11.7.0.0.zip` results in two zip files (`OS-11.7.0.0.zip` and `RSA-11.7.0.0.zip`) and some other files.

f. Unzip the:
`OS-11.7.0.0.zip` into the `/var/netwitness/<your-zip-file-repo>/11.7.0.0/OS` directory.
`unzip /var/netwitness/<your-zip-file-repo>/11.7.0.0/OS-11.7.0.0.zip -d /var/netwitness/<your-zip-file-repo>/11.7.0.0/OS`

The external url for the repo is `http://<web server IP address>/<your-zip-file-repo>`.

g. Unzip the:
`RSA-11.7.0.0.zip` into the `/var/netwitness/<your-zip-file-repo>/11.7.0.0/RSA` directory.
`unzip /var/netwitness/<your-zip-file-repo>/11.7.0.0/RSA-11.7.0.0.zip -d /var/netwitness/<your-zip-file-repo>/11.7.0.0/RSA`

h. (Conditional - For Azure) Follow these steps for Azure update.

   i. `mkdir -p /var/netwitness/<your-zip-file-repo>/11.7.0.0/OS/other`

   ii. `unzip nw-azure-11.3-extras.zip -d /var/netwitness/<your-zip-file-repo>/11.7.0.0/OS/other`

   iii. `cd /var/netwitness/<your-zip-file-repo>/11.7.0.0/OS`

   iv. `createrepo`

i. Use the `http://<web server IP address>/<your-zip-file-repo>` in response to **Enter the base URL of the external update repositories** prompt from NW 11.7.0.0 Setup program (`nwsetup-tui`) prompt.

Appendix C. Silent Installation Using CLI

You can use the following Command Line Interface commands to run the installation script (`nwsetup-tui`) without getting prompted for inputs. This enables you to automate the installation of a host by supplying response to the scripts prompts through the command line.

1. After you have created a base image on the host, log in to the host with the `root` credentials.
2. Submit the `nwsetup-tui` script with the `--silent` command and the arguments that you want to apply.

   The following command string is an example of how you would install a basic NW Server host.

   ```
   nwsetup-tui --silent --is-head=true --host-name=new-host --master-
   pass=netwitness --deploy-pass=netwitness --repo-type=1 --custom-
   firewall=false --ip-override=false --eula=true
   ```

Note: In NetWitness Platform version 11.6 or later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`.
If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

3. (Conditional - For Component Hosts Only) Install the appropriate service **Category** on the newly provisioned host in the NetWitness Platform Hosts view.

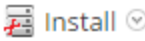
   a. Log into NetWitness and go to  **(Admin)** > **Hosts**.

   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background

   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the host in the **New Hosts** dialog and click **Enable**.

   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the **Hosts** view (for example, **Event Stream Analysis**) and click 

.

   The **Install Services** dialog is displayed.

   d. Select the appropriate host type in **Category** and click **Install**.

Arguments

Argument	Description
`--help-install-opts`	Display all the arguments in this table.

Argument	Description
`--eula`	Accept or decline the End User License Agreement (EULA). Specify: • `true` (default) to accept the agreement • `false` to decline it and cancel the installation. For example: `--eula=true`
`--is-head`	Designate the host as the NW Server host or a component host. Specify: • `true` for NW Server host. • `false` for Component host. For example: `--is-head=true`
`--host-name`	Specify new hostname. If you do not specify this argument, NetWitness Platform retains the existing hostname. For example: `--host-name=<hostname>`
`--master-pass`	Enter master password. For example: `--master-pass=<password>`
`--deploy-pass`	Enter deployment password. For example: `--deploy-pass=<password>`
`--iface-name`	Specify network interface. For example: `--iface-name=eth0`
`--ip-override`	Accept or override IP address found for this host or change the IP configuration found on the host. Specify: • `true` provide IP address. • `false` use IP address found on the host. For example: `--ip-override=false`
`--ip-type`	Select ip address configuration type. Specify: • `1` Static IP Configuration) • `2` DCHP For example: `--ip-type=1`
`--ip-addr`	For Static IP configuration, enter IP Address for static address. For example: `--ip-addr=<ip-address>`
`--ip-netmask`	For Static IP configuration, enter Subnet Mask for static address. For example: `--ip-gateway=<subnet-mask>`

Argument	Description
`--ip-gateway`	For Static IP configuration, enter default gateway for static address. For example: `--ip-gateway=<default-gateway>`
`--ip-nameserver`	IP address assigned to DNS server. `--ip-nameserver=<ip-address>`
`--ip-nameserver-secondary`	Optional - IP address assigned to a secondary DNS server. For example: `--ip-nameserver-secondary=<ip-address>`
`--ip-domain`	For Static IP configuration, enter Local Domain Name for static address. For example: `--ip-domain=<default-gateway>`
`--repo-type`	Select type of update repository. Specify: • `1` Local repository • `2` External repository For example: `--repo-type=1`
`--repo-url`	For an external update repository, specify the url of the repository. For example: `--repo-url=<url>`
`--head-ip`	For a component host, specify IP Address of the NW Server. For example: `--head-ip=<ip-address>`
`--custom-firewall`	Disable default firewall configuration and use your custom configuration. Specify: • `true` use custom firewall configuration. • `false` use default firewall configuration. For example: `--custom-firewall=true`
`--use-nat`	Configure the host to use Network Address Translation (NAT) based IP addresses: • `true` use NAT IPs to connect to other hosts • `false` do not use NAT IPs to connect to other hosts (default) For example: `--use-nat=false`

Appendix D. Third Party Server System Requirement

This section contains all the hardware requirements and configuration needed to successfully deploy NetWitness Platform on Third Party Server Hardware. It contains the required compute, memory, drive types and recommendations.

Third Party Server Deployments only support the following NetWitness Platform components:

- Core Services (Broker , Decoder, Log Decoder, Archiver , Concentrator)
- Analyst UI
- New Health & Wellness
- Log Collector
- Malware Analysis
- Warehouse connector

Hardware Requirements

Administrators must configure single bootable block device (RAID volume/group) and ensure it is bootable. After installation is complete, See core service storage configuration the NetWitness Storage Guide.

Item	Core
Host Type	• NW Server • Warm standby • Analyst UI • Health & Wellness • Core Services • Log Collector • Malware Analytics • Warehouse Connector
Memory	128 GB
Processor	
Processor Speed	3.2 Ghz
# of Processors	2
# of Cores	8 Cores Per Processor
# of Threads	16 Threads Per Processor
Storage Configuration	

Item	Core
Volume	Single block device. 150 GB or greater.
Drive Types	10K SAS or SSD
RAID Configuration	1, 5 or 6
Network	
NIC	Supported by Centos 7.9 (1G, 10G or 40G)
Capture Speed	3G

NetWitness® Platform XDR

Version 12.0

Getting Started Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

Getting Started with NetWitness Platform **6**
- Overview 6
- Architecture 6
- Core Versus Downstream Components 8

Logging in to NetWitness Platform **9**
- Log Off NetWitness Platform 10

Changing Your Password **11**

Identifying Your Role **12**

NetWitness Platform Basic Navigation **13**
- Menu changes 14
- Accessing Main Views 14
- Secondary Menus 14
- Additional Options 15
- Main Views 15
- Springboard 16
- Investigate 17
- Respond 21
- Users 25
- Hosts 26
- Files 27
- Dashboard 27
- Reports 29
- Configure 29
- Admin 31

Setting Up Your Default View by SOC Role **34**
- Set Your Default View 36

Managing the Springboard **38**
- Working with the Springboard 40
  - Add a Custom Private Board 40
  - Add a Panel 41
  - Edit a Panel 43
  - Rearrange Panels 43
  - Delete Panels 43
  - Restore System Default Settings 44
  - Refresh a Panel 44

Managing Dashboards **45**
Dashboard Basics 45
Dashboard Title 45
Dashboard Selection List 45
Dashboard Toolbar 46
The Default Dashboard 47
Selecting a Preconfigured Dashboard 47
Enabling or Disabling Dashboards 48
Enable a Dashboard 49
Disable a Dashboard 51
Setting a Dashboard as a Favorite 51
Creating Custom Dashboards 52
Working with Dashlets 53
Add a Dashlet 55
Edit Dashlet Properties 56
Rearrange a Dashlet 58
Maximize a Single Dashlet 59
Delete a Dashlet 60
Importing and Exporting Dashboards 60
Import a Dashboard 60
Export a Dashboard 61
Copying a Dashboard 61
Sharing a Dashboard 62
Removing Unwanted Dashboards 62
Using Dashboards in the Analyst User Interface 63
Setting User Preferences **65**
Preferences 65
View your Preferences 66
Set the Language and Time Zone 66
Enable or Disable System Notifications for Your User Account 66
Enable or Disable Context Menus for Your User Account 67
User Preferences 67
View Your User Preferences 67
Set the Language, Time Zone, and Date and Time Format 68
Select the Default NetWitness Platform Starting Location 69
Select the Default Investigate View 69
Choose the Appearance of NetWitness Platform 70
Managing Jobs **72**
Display the Jobs Tray 72
View All of Your Jobs 73

Pause and Resume Scheduled Execution of a Recurring Job 73
Cancel a Job 73
Delete a Job 74
Download a Job 74
Viewing and Deleting Notifications 75
View Recent Notifications 75
View All Your Notifications 76
Delete Notification Records 76
Viewing Help in the Application 77
View Inline Help 77
View Tooltips 77
View Online Help 77
Finding Documents on NetWitness Community 78
Locate NetWitness Documentation 78
Locate NetWitness Content 78
Locate NetWitness Supported Event Sources 78
Locate Hardware Setup Guides 79
Follow Content for Updates 79
Send Your Feedback to NetWitness 79
Troubleshooting the User Interface 80
Basic Troubleshooting Tips for User Setup 80
Analyst User Interface Dashlet Issue 81
Springboard Issue 81
Springboard Fails to Load the Panel Issue 82
Inconsistent Event Panel Count Issue 82
NetWitness Platform Getting Started References 83
User Preferences 84
Notifications Panel and Notifications Tray 89
Jobs Panel and Jobs Tray 92

Getting Started with NetWitness Platform

Overview

NetWitness is a powerful threat detection suite that enables Security Operation Centers (SOCs) to quickly locate, prioritize, and triage threats. NetWitness helps you to isolate and remediate known threats as well as those that were previously unknown. It provides deep insight into packets, logs, and endpoints that provide you with an unparalleled view into your enterprise or business.

NetWitness is powerful, but it is easier for Tier 1 Analysts to use because it automates the process of identifying and prioritizing suspicious threats. Tier 2 and Tier 3 Analysts can hunt for and locate threats by searching and filtering events and then examining events using reconstruction and analysis tools.

Architecture

NetWitness is a distributed and modular system that enables highly flexible deployment architectures that scale with the needs of the organization. NetWitness allows administrators to collect three types of data from the network infrastructure: packet data, log data, and endpoint data. The key aspects of the architecture are:

- **Distributed Data Collection**: The **Decoder** ingests packet data while the **Log Decoder** ingests log data. Decoders parse and reconstruct all collected network traffic from Layers 2 - 7, or log and event data from hundreds of devices and event sources, including NetWitness Endpoint data (if installed and configured). The **Concentrator** indexes metadata extracted from network or log data and makes it available for enterprise-wide querying and real-time analytics while also facilitating reporting and alerting. The **Broker** aggregates data captured by other devices and event sources. Brokers aggregate data from configured Concentrators; Concentrators aggregate data from Decoders. Therefore, a Broker bridges the multiple real-time data stores held in the various Decoder or Concentrator pairs throughout the infrastructure.
- **Real-time Alerting**: The NetWitness **Event Stream Analysis (ESA)** service provides advanced stream analytics such as correlation and complex event processing at high throughputs and low latency. It can process large volumes of disparate event data from Concentrators. ESA uses an advanced Event Processing Language (EPL) that allows analysts to express filtering, aggregation, joins, pattern recognition, and correlation across multiple disparate event streams. Event Stream Analysis helps to perform powerful incident detection and alerting.
- **Real-time Analytics** (automatic analysis of events): The Automated Threat Detection functionality includes preconfigured ESA analytics module for detecting Command and Control traffic.
- **NetWitness Server**: The NetWitness Server provides reporting, investigation, administration, and other aspects of the user interface.
- **Capacity**: NetWitness has a modular-capacity architecture enabled with direct-attached capacity (DACs) or storage area networks (SANs), that adapts to the organization's short-term investigation and long-term analytic and data-retention needs.

NetWitness provides large deployment flexibility. You can design its architecture using as many as multiple dozens of physical hosts or a single physical host, based on the particulars of the customer's performance and security-related requirements. In addition, the entire NetWitness system has been optimized to run on virtualized infrastructure.

The System Architecture comprises of these major components- Decoders, Brokers, Concentrators, Archivers, ESA, and Warehouse Connectors. NetWitness components can be used together as a system or can be used individually.

- In a security information and event management (SIEM) implementation, the base configuration requires these components- Log Decoder, Concentrator, Broker, Event Stream Analysis (ESA), and the NetWitness Server.
- In a forensics implementation, the base configuration requires these components- Decoder, Concentrator, Broker, ESA, Malware Analysis, and Endpoint Log Hybrid. The Respond Server service is also required and is used to prioritize alerts.

The table provides a synopsis of each major component:

System Component	Description
Decoder / Log Decoder	• NetWitness collects packet, log, and endpoint data. • Packet data, that is, network packets, are collected using the Decoder through the network tap or span port, which is typically determined to be an egress point on an organization's network. • A Log Decoder can collect four different log types - Syslog, ODBC, Windows eventing, and flat files. • Windows eventing refers to the Windows 2008 collection methodology and flat files can be obtained via SFTP. • Both types of Decoders ingest raw transactional data that is enriched, closed out, and aggregated to other NetWitness components. • The process for ingesting and parsing transactional data is a dynamic and open framework.
Endpoint Log Hybrid	• Collects and manages endpoint (host) data from Windows, Mac, or Linux hosts. • Records data about every critical action, such as process, file, registry modification, network connections, and user console interactions. • Collects Windows logs and file logs from Windows host, if configured. • Generates metadata to correlate endpoint data with sessions from other events sources, such as logs and network. • Performs on-demand memory analysis and suspicious user behavior detection.

System Component	Description
Concentrator	• Provides index and query capability to NetWitness Collections. • Can optionally forward data to ESA.
Broker	• Distributes NetWitness Collection access across many Concentrators or Archivers, making the entire NetWitness enterprise appear as a single collection.
Archiver	• The Archiver service enables long-term log archiving by indexing and compressing log data and sending it to archiving storage. • The archiving storage is optimized for long-term data retention, and compliance reporting. • Archiver stores raw logs and log metadata from Log Decoders for long-term retention, and it uses Direct-Attached Capacity (DAC) for storage. **Note:** Raw packets and packet metadata are not stored in the Archiver.
Event Stream Analysis (ESA)	• ESA provides event stream analytics such as correlation and complex event processing at high throughputs and low latency. It can process large volumes of disparate event data from Concentrators. • ESA uses advanced Event Processing Language that allows users to express filtering, aggregation, joins, pattern recognition, and correlation across multiple disparate event streams. • ESA helps to perform powerful incident detection and alerting.

Core Versus Downstream Components

In NetWitness, the Core services ingest and parse data, generate metadata, and aggregate generated metadata with the raw data. The Core services are Decoder, Log Decoder, Concentrator, and Broker. Downstream systems use data stored on Core services for analytics; therefore, the operations of downstream services are dependent on Core services. The downstream systems are Archiver, ESA, Malware Analysis, Investigate, and Reporting.

Although the Core services can operate and provide a good analytics solution without the downstream systems, the downstream components provide additional analytics. ESA provides real-time correlation across sessions and events as well as between different types of events, such as log, packet, and endpoint data. Investigate provides the ability to drill into data, examine events and files, and reconstruct events in a safe environment. The Malware Analysis service provides real-time, automated inspection for malicious activity in network sessions and associated files.

Logging in to NetWitness Platform

Note: NetWitness supports modern (or current) versions of Google Chrome, Mozilla Firefox, Microsoft Edge, and Apple Safari. It is possible to use a different browser, but some features may not function as expected. Internet Explorer is no longer supported.

Logging in to NetWitness can vary based on your environment. You may have an internal user account or an external user account. Internal user accounts are local to the NetWitness and internal users can log in to NetWitness and receive role-based permissions. External user accounts authenticate outside of the NetWitness and are mapped to NetWitness roles. If you are an external user and you cannot access NetWitness or view the information that you need, contact your System Administrator. Your Administrator can assign the appropriate roles to your account.

NetWitness Platform 11.4 or later also supports Single Sign-On authentication (SSO) using Security Assertion Markup Language 2.0 (SAML 2.0) protocol with Active Directory Federation Services (ADFS) as the Identity Provider.

If SSO authentication is enabled by your administrator, you will be redirected to the Identity Provider User Interface instead of the default NetWitness login page. After you enter the username and password you will be securely logged into NetWitness Platform.

Note: In 11.4 or later, Single Sign-On (SSO) authentication can be used to access the NetWitness Platform UI, and Analyst UI Deployment.

1. Use the icon provided by your Administrator, or type the following in your web browser:
   `https://<hostname or IP address>/login`
   Where `<hostname or IP address>` is the hostname or the IP address of your NetWitness server.
   If Single Sign-On authentication is enabled, this redirects you to the ADFS login screen.
2. Enter the username and password, and then click **Sign in**.
   If the login is successful, you will be logged into the landing page specified in the user preferences.

Note: If you had previously authenticated to any other application configured to the same IDP, then you may be redirected to the requested NetWitness Platform UI without being prompted for the credentials.

If you are locked out:

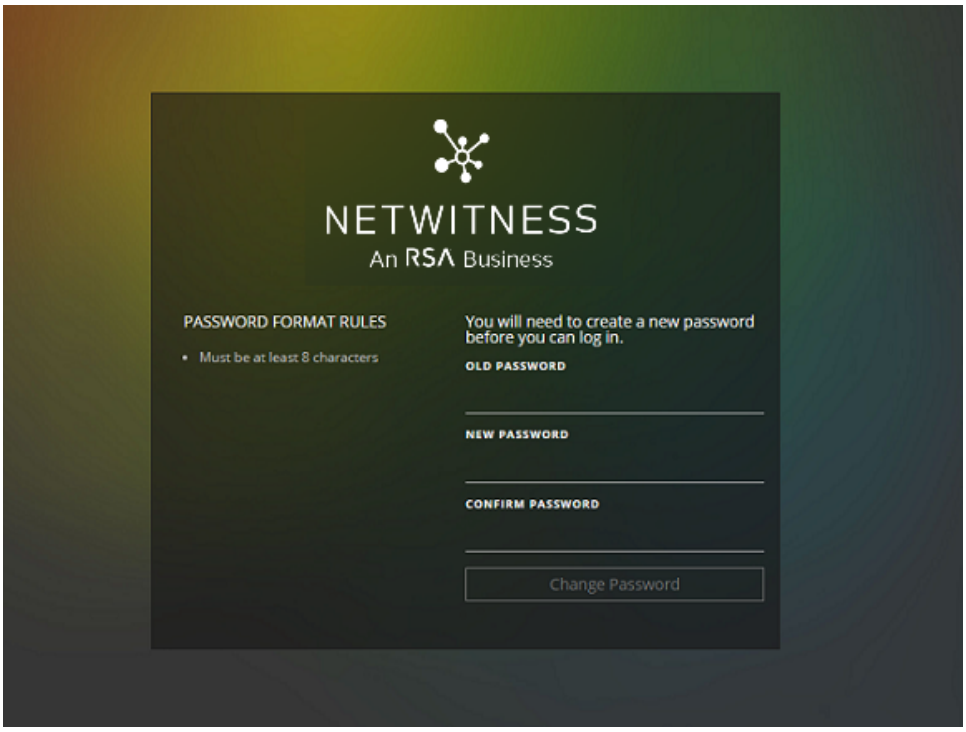
Note: This information applies to internal accounts only. It does not apply to Active Directory or PAM accounts.

If you try too many times to log in with an incorrect username or password, your account will be locked. Contact your Administrator to unlock your account.

If you have a new account or your account is expired:

Note: This procedure applies to internal accounts only. It does not apply to Active Directory or PAM accounts.

1. In the dialog to create a new password, enter your old password, type a new password, and confirm it. Password format rules (as defined by your system administrator) are provided on the left and your new password must conform to the indicated format rules.



2. Click **Change Password**.

If you do not have the appropriate access to NetWitness:

If you are able to log in successfully, but you are not able to view the information that you need, it is possible that you need a user role assigned to your user account. Contact your Administrator for assistance.

Log Off NetWitness Platform

To log off from the Respond and some Investigate views:

1. In the main menu bar, select your username, for example, admin.
2. In the User Preferences, click **Sign Out**.

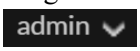
To log off from the other views:

In the main menu bar, select your username and then select **Sign Out**.

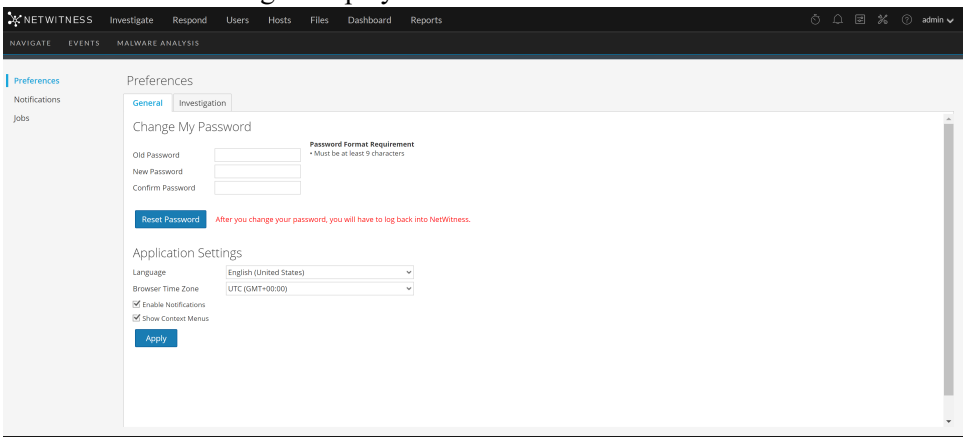
Changing Your Password

You can change the password that you use for NetWitness authentication at any time in your user preferences. Your administrator defines the appropriate password strength requirements for your NetWitness password, such as minimum password length and minimum number of uppercase, lowercase, decimal, non-Latin alphabetic, and special characters. These requirements are then displayed when changing your password.

To change your password:

1. Do one of the following:
   - For most views, such as Investigate, Dashboard, Reports,  (Configure) or  (Admin) , select your username, for example , and then select **Profile**.
   - In the Springboard, Investigate view (Events), Respond, Users, Hosts, and Files, select your username, for example , and in the User Preferences dialog click **Change my password**.

   The Preferences dialog is displayed.



2. In the **Change My Password** section, enter the password that you used to authenticate to NetWitness in the **Old Password** field.
3. In the **New Password** field, enter the password that you want to use for the next login.
4. In the **Confirm Password** field, retype the new password.
5. Click **Reset Password**.
   You will be logged out of NetWitness for the changes to take effect. The new password becomes effective the next time you log in to NetWitness.

For more information on user preferences, see Setting User Preferences.

Identifying Your Role

The roles listed here are the typical roles or functions of a Security Operations Center (SOC). Determine the role or roles that you perform in the SOC. You can use these functions as a guide to decide how to set up and navigate NetWitness so that you can efficiently perform your job tasks.



SOC Team



SOC Manager
(SOC Management and Reporting)

- Manage SOC readiness
- Respond to incidents
- Respond to data breaches



Data Privacy Officer

- Monitor and protect privacy and sensitive information



Incident Responder
(T1 Analyst)

- Respond to incidents
- Remediate incidents



Threat Hunter
(T2/T3 Analyst)

- Hunt for threats
- Conduct forensic analysis
- Recommend issues for remediation
- Remediate issues



Content Expert
(Threat Intelligence)

- Investigate new threat intelligence
- Evaluate and create new feeds
- Create correlation rules to flag indicators of compromise



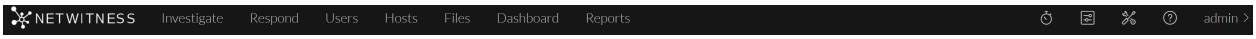
System Administrator

- Install and configure equipment and software
- Manage user access
- Monitor and fine tune performance
- Backup and restore data
- Manage storage and archives
- Update software
- Create reports for regulatory compliance

NetWitness Platform Basic Navigation

The NetWitness application is divided into ten main functional areas, known as views, that are based on typical Security Operation Center (SOC) roles.

Note: On upgrade to version 11.5 or later, by default the Springboard is displayed if you have not configured the default landing page in previous versions.



- **Springboard**: Springboard presents Analysts with the platform-wide detections and signals in a single view to hunt and investigate faster than ever before. System Administrators set up and maintain the Springboard. You can view the Springboard at any time by clicking NetWitness in the main menu. For more information, see Managing the Springboard.
- **Investigate**: This view is primarily for Threat Hunters, who prefer to manually hunt for threats using NetWitness metadata, raw event data, and event reconstruction and analysis. Incident Responders also use this view to get details about events associated with an incident being investigated. Both Threat Hunters and Incident Responders can use the forensic event reconstruction and event analysis features in this view.
- **Respond**: This view is for Incident Responders, who can view a list of prioritized incidents to triage. These incidents come from sources such as ESA rules, NetWitness Endpoint, or ESA Analytics modules for Automated Threat Detection. You can also view all of the alerts received by NetWitness here.
- **Users**: This view is for SOC Managers and Analysts to discover, investigate, and monitor risky behaviors across entities namely Users and Network in your environment.
- **Hosts**: This view is for Analysts, who can investigate or perform analysis on hosts using attributes such as IP address, host name, Mac address, risk score, and so on.
- **Files**: This view is for Analysts, who can investigate or perform analysis on files using attributes such as IP address, host name, Mac address, risk score, and so on
- **Dashboard**: This view is for all users. You can view dashboards on different areas of interest depending on your user permissions.
- **Reports**: This view is for all users. You can view reports on different areas of interest depending on your user permissions.
-  **Configure**: This view is for Threat Intel personnel (Content Experts), who configure data sources and inputs to NetWitness. Content Experts use this area to download and manage Live content. They can also create and manage incident and ESA rules.
-  **Admin**: This view is for System Administrators, who set up and maintain the overall application.

Menu changes

The following table illustrates the top-level menu changes in the 11.5 version.

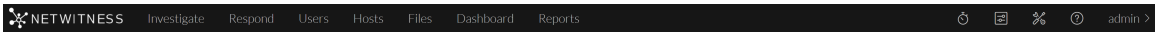
Previous Version - 11.4 and earlier	11.5 Version
N/A	RSA Click the RSA logo at the top left corner to view the Springboard.
Monitor > Dashboard Monitor > Reports	Dashboard Reports
Investigate > Hosts Investigate > Files Investigate > Entities	Hosts Files Users
Configure	
Admin	

The following table illustrates the menu changes in the 12.0 version.

Previous Version - 11.5 to 11.7.x.x	12.0 Version
RSA logo	NETWITNESS Click the NetWitness logo at the top left corner to view the Springboard.

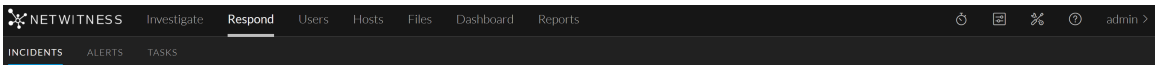
Accessing Main Views

The options that open each of the main views are listed at the top of the browser window. With the appropriate permissions, you can access any of these views at the top of every UI at any time.



Secondary Menus

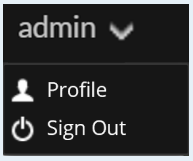
The main views have secondary menus with additional views that you can select, which vary according to the tasks that you can complete. The following example shows the Respond menu.



Additional Options

In addition to the main views, there are additional options at the top of the UI that are common to the application.

The following table describes the common options.

Common Option	Name	Description
	Jobs	In the Investigate, Dashboard, Reports,  (Configure) , and  (Admin) views, click this icon to view and manage your jobs in the Jobs tray. Jobs are on-demand or scheduled tasks that take some time to complete in the NetWitness application.
	Notifications	Click this icon to view notifications from the application.
	User Preferences	Click this icon to view your available user preference options. You can manage your user preferences and log out of NetWitness.
	User Profile	Click your user profile to view the available options. You can manage your user preferences, change your password, and log out of NetWitness UI.
	Help	Click this icon to view NetWitness help topics.

Main Views

The following sections explain the main views:

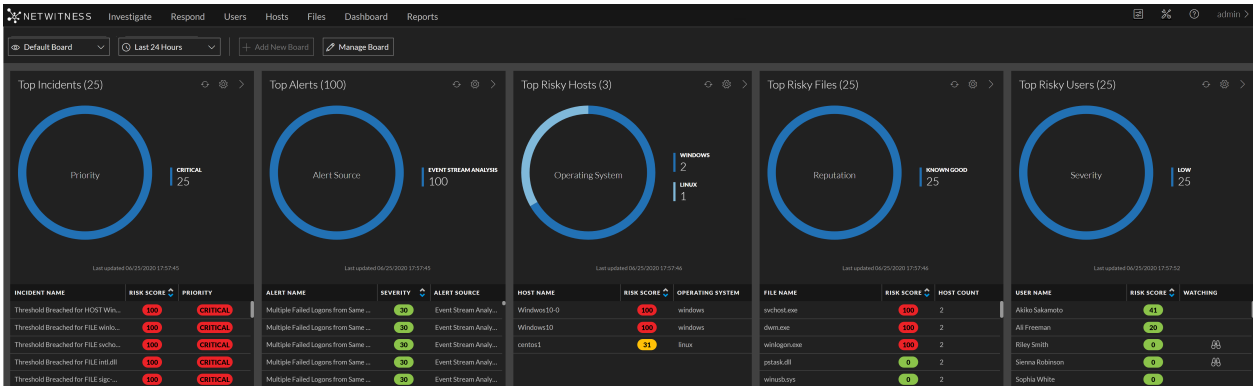
- Springboard
- Investigate
- Respond
- Users
- Hosts

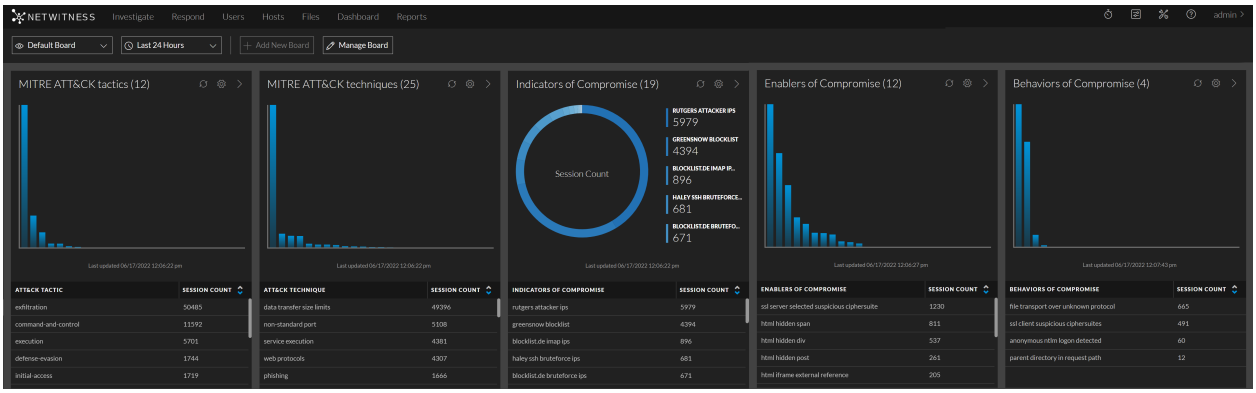
- Files
- Dashboard
- Reports
- Configure
- Admin

Springboard

(From 11.5 and later) NetWitness Platform Springboard is an easy-to-use landing page that presents platform-wide detections and signals in a single view to help analysts hunt and investigate faster than ever before.

Click the NetWitness logo at the top left corner to view the Springboard.





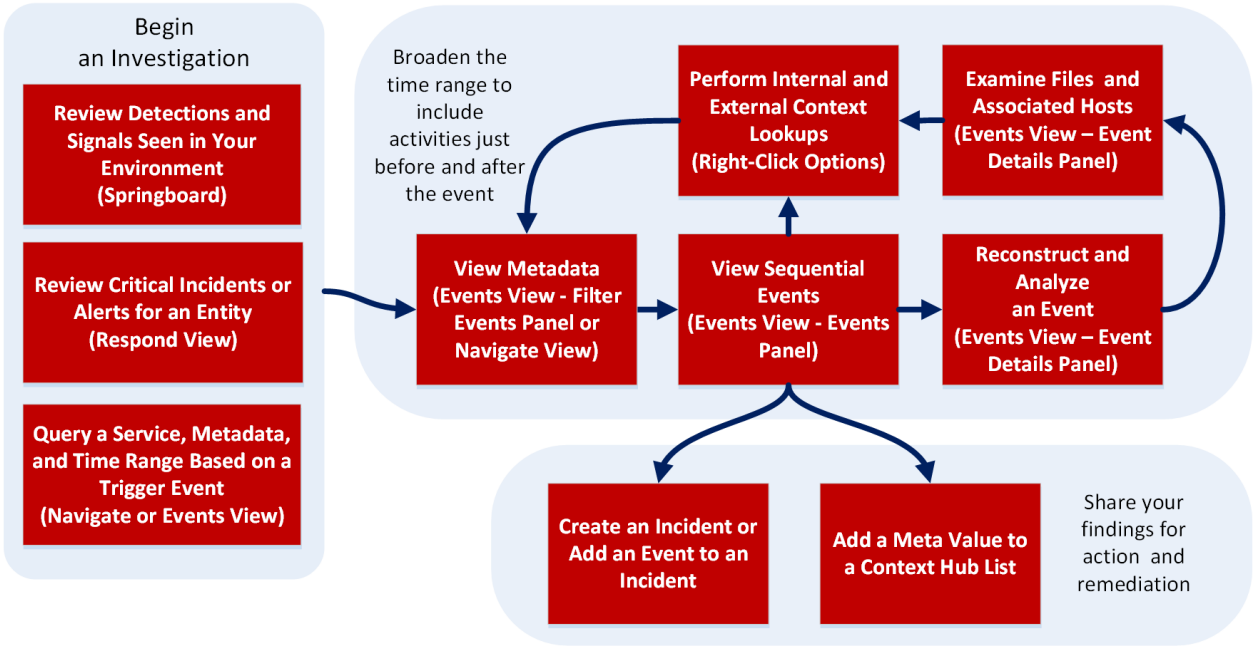
What can I do here?	Path	Show me how
View out-of-the-box panels Edit a panel Refresh a panel Select time range View all incidents, alerts, users, files, and hosts View details of selected incident, alert, user, file, and host Manage Board (add, rearrange, and delete panels) Add a new custom private board	Springboard view	See Managing the Springboard.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

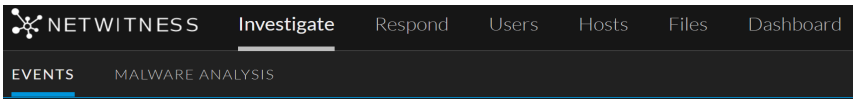
Investigate

The Investigate view is the tool for SIEM, network, and endpoint data investigation, presenting different views into a set of data. Analysts can see metadata and raw data for endpoints, logs, and events, as well as potential indicators of compromise. In addition to investigating data on a specific service, you can pivot into Investigate from Respond, the Dashboard view, an entry in a report generated by the Reporting Engine, or a properly configured third-party application.

You can begin your investigation in any Investigate view, then continue the investigation seamlessly in another Investigate view. The manner in which you proceed is determined by the question that needs to be answered. If you find an event that needs a response, you can create an incident in Respond where an incident responder will take further action. The following figure depicts the high-level flow of an investigation. The *NetWitness Investigate User Guide* provides detailed information.

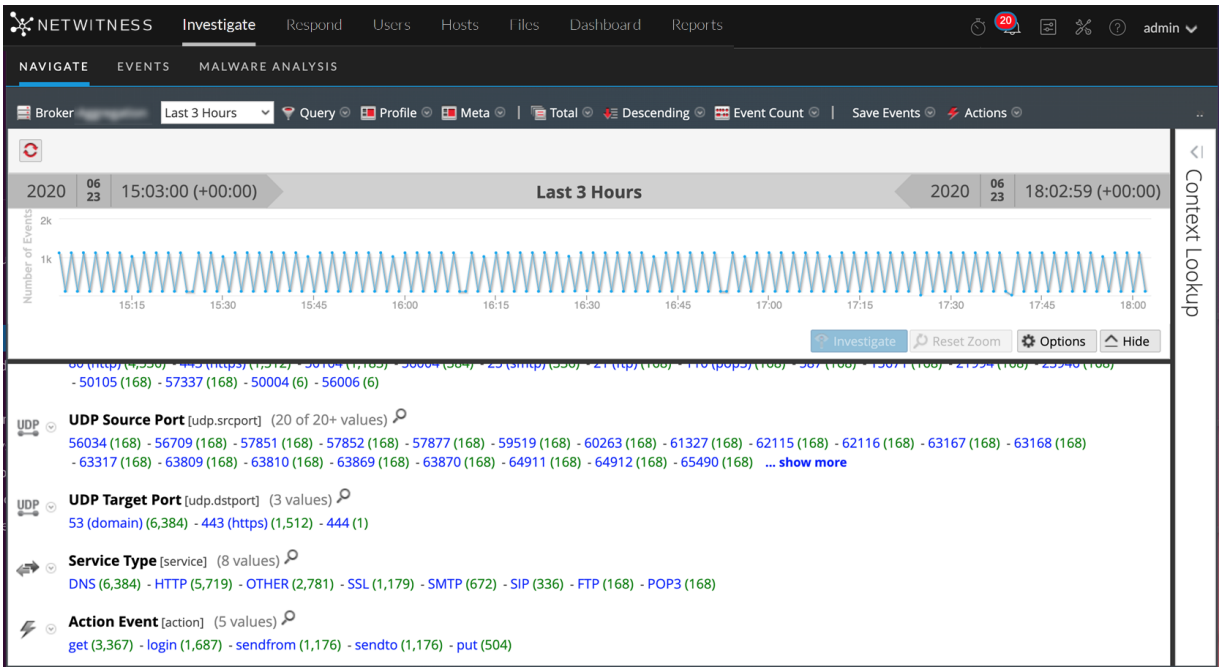


Investigate Menu

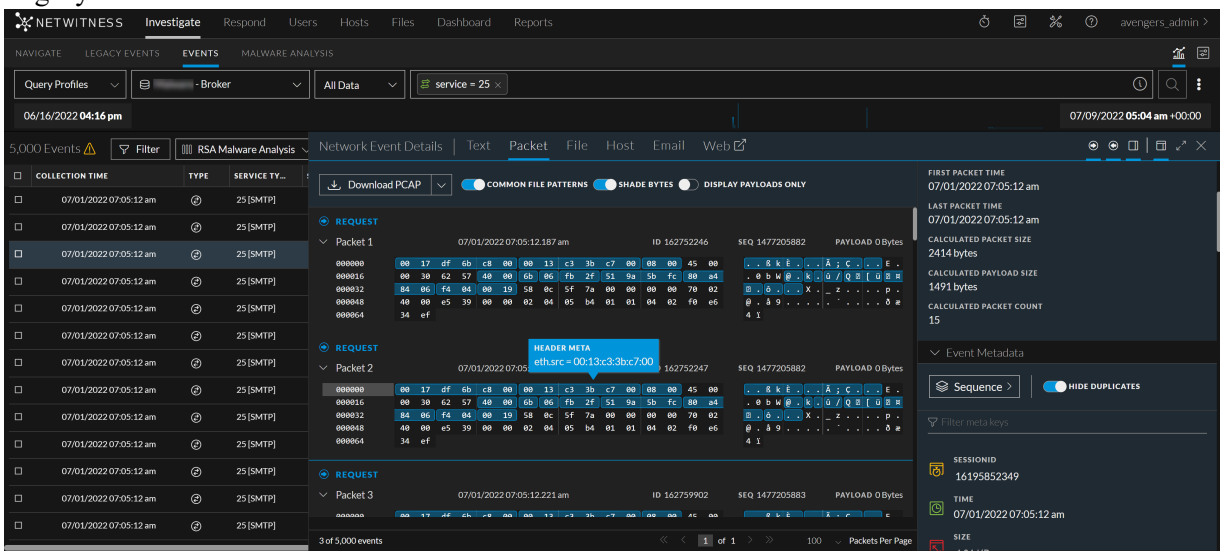


The Investigate menu has the following options:

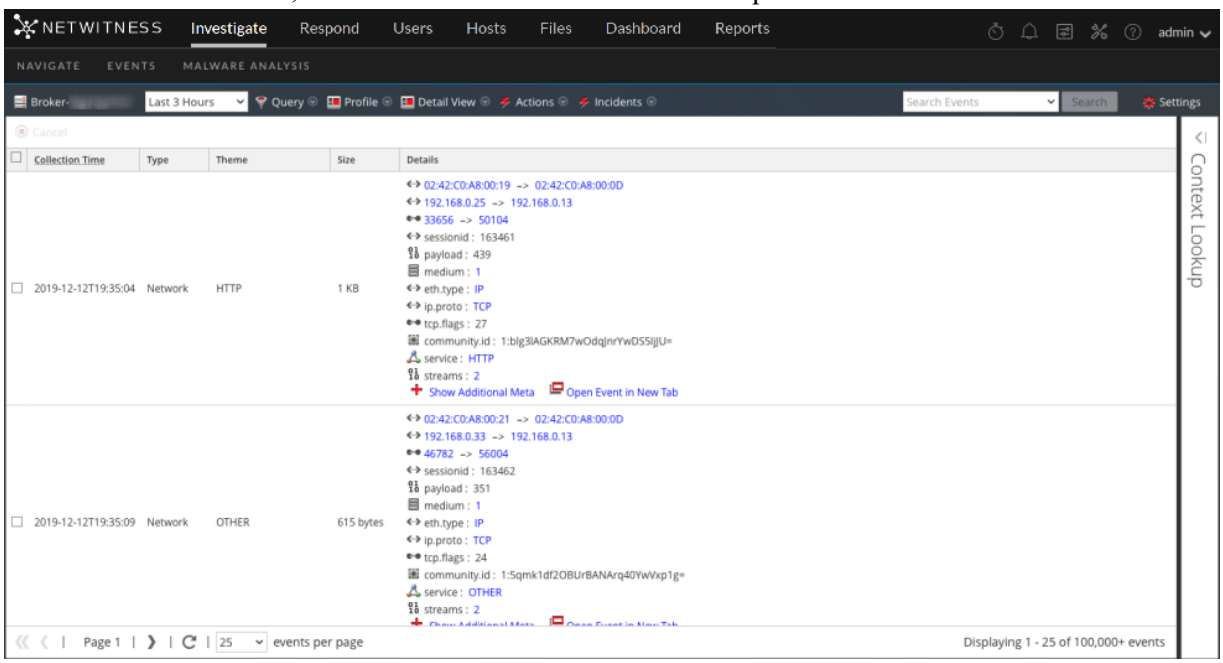
- **Navigate**: The Navigate view provides a list of meta keys and meta values with a focus on metadata. You can drill into the data, search for events, open a selected event in the Events view, and look up additional context from the Context Hub service.



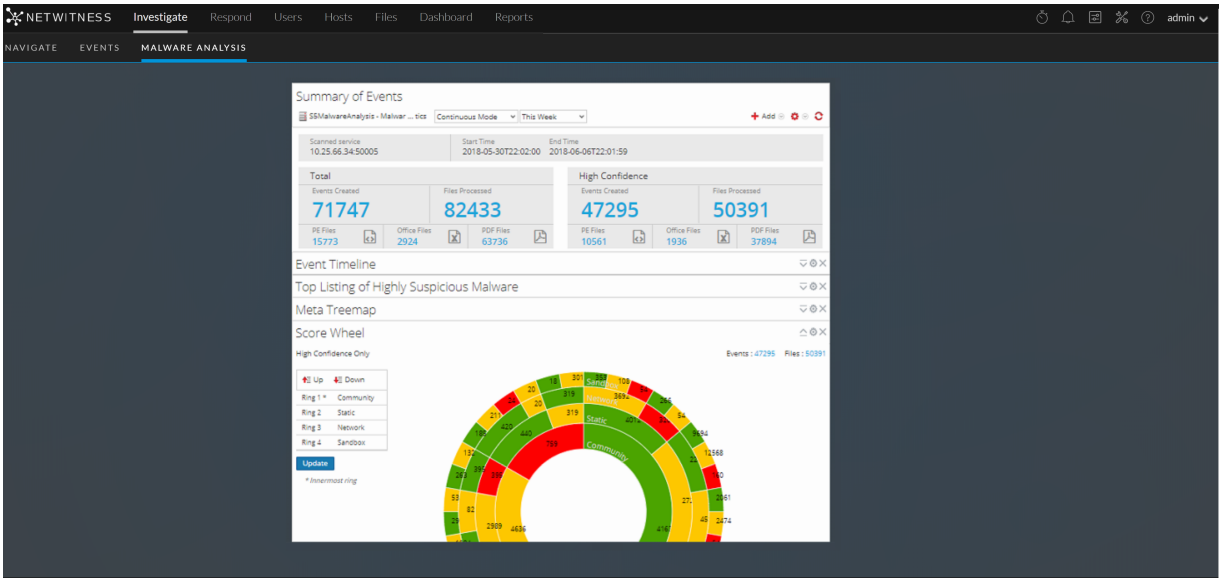
- **Events**: The Events view (formerly Event Analysis view) is the default user interface for interacting with events. It provides a sortable list of events with focus on metadata and raw data. You can search for events, view a reconstruction that offers helpful cues to identify points of interest, pivot to standalone Endpoint, look up additional context from the Context Hub service, look up data in Live, do external lookups, and create an incident for incident responders. By default only the Events view appears in the menu, but when the Legacy Events view is enabled, both the Events view and the Legacy Events view are visible in the menu bar.



- **Legacy Events**: With major functionality added to the 11.3 Events view, the Legacy Events is no longer needed and it is hidden unless the administrator enables it. The Legacy Events view provides a list of events with a focus on raw data. You can browse a simple list of events, a detailed list, and a log list. You can search for events, view a reconstruction of an event, look up additional context from the Context Hub service, and create an incident for incident responders.



- **Malware Analysis**: Malware Analysis is an automated malware analysis processor designed to analyze certain types of file objects (for example, Windows PE, PDF, and MS Office) to assess the likelihood that a file is malicious. Using Malware Analysis, you can prioritize the massive number of files captured in order to focus analysis efforts on the files that are most likely to be malicious.



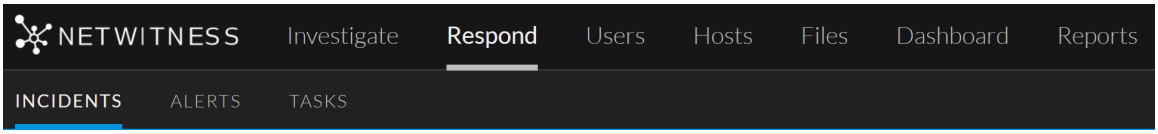
What can I do here?	Path	Show me how
Configure Investigate Views and Preferences	Investigate view	See "Configuring Investigate Views and Preferences" in the *NetWitness Investigate User Guide*.
Browse Event Metadata	Navigate view	See "Refining the Results Set" in the *NetWitness Investigate User Guide*.
Browse Raw Events	Events view	See "Refining the Results Set" in the *NetWitness Investigate User Guide*.
Analyze Raw Events and Metadata	Events view	See "Reconstructing and Analyzing Events" in the *NetWitness Investigate User Guide*.
Scan Files and Events for Malware	Malware Analysis view	See the *Malware Analysis User Guide*.
Triage an Incident	Pivot from the Respond view	See the *NetWitness Respond User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Respond

The Respond view presents analysts with a queue of incidents in severity order. When you take an incident from the queue, you receive relevant supporting data to help you investigate the incident. From there, you can determine the incident scope and escalate or remediate it as appropriate.

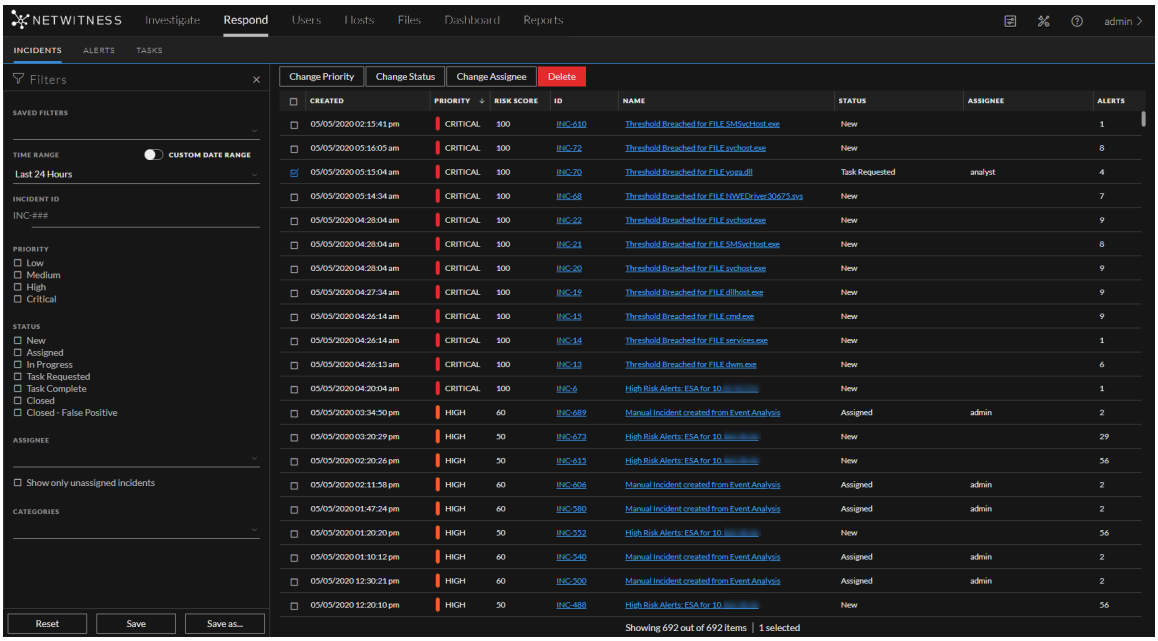
Respond Menu



The Respond menu has the following options:

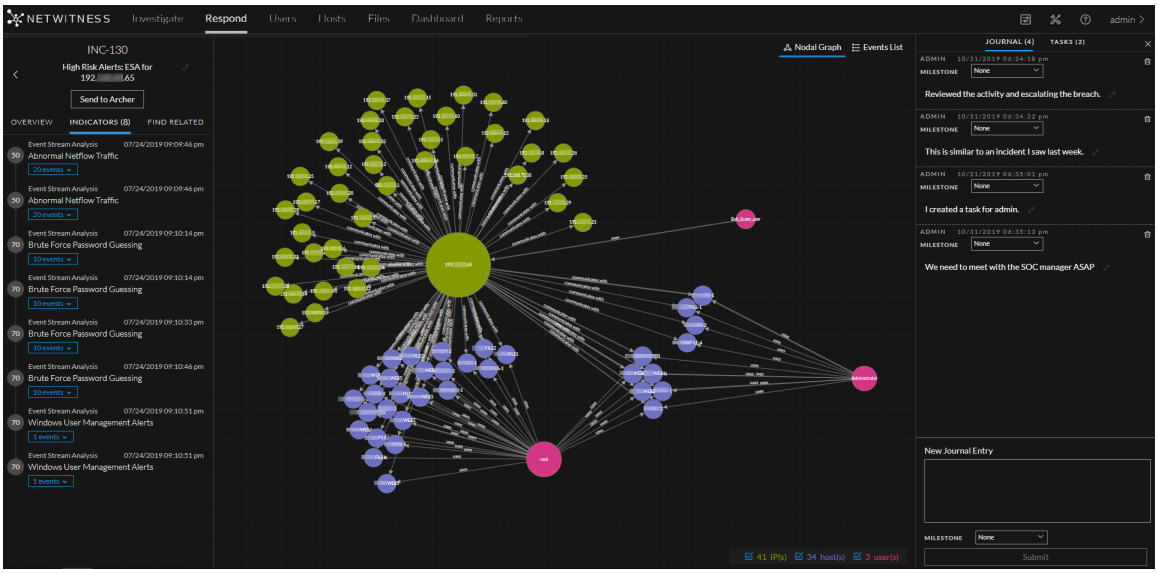
- **Incidents**: The Incidents List view contains a list of all incidents with basic information. The Incident Details view provides extensive details about the incident.
- **Alerts**: The Alerts List and Alert Details views provide information about all of the threat alerts and indicators received by NetWitness in one location.
- **Tasks**: The Tasks List view enables you to create tasks and track them to completion.

The following figure shows the Respond view - Incidents List view, which shows a list of prioritized incidents.

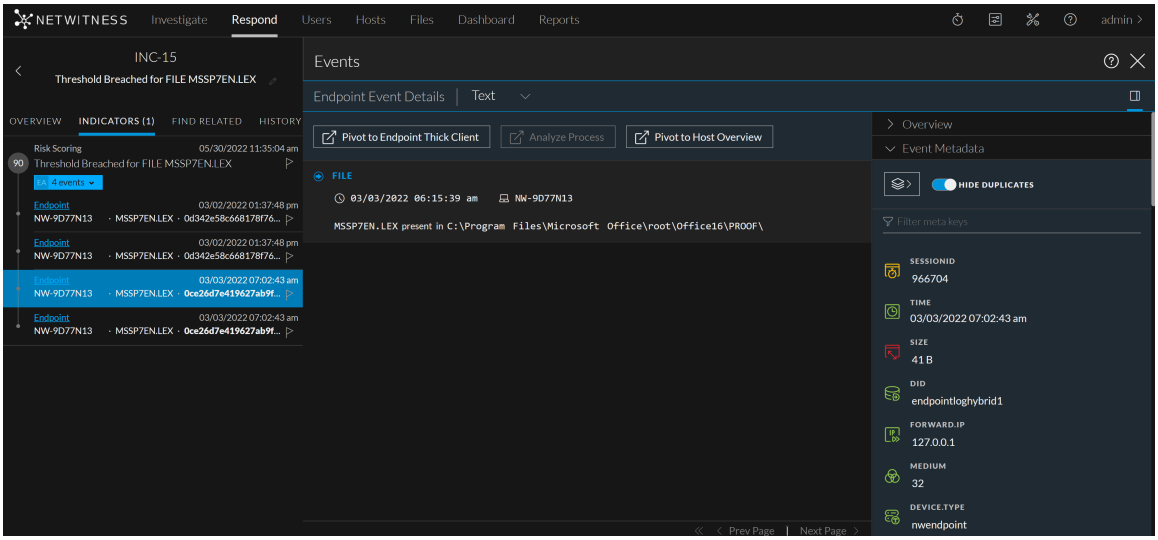


When using NetWitness as your case management tool, you can also manage incidents from this view. New incidents appear at the top of the incident queue.

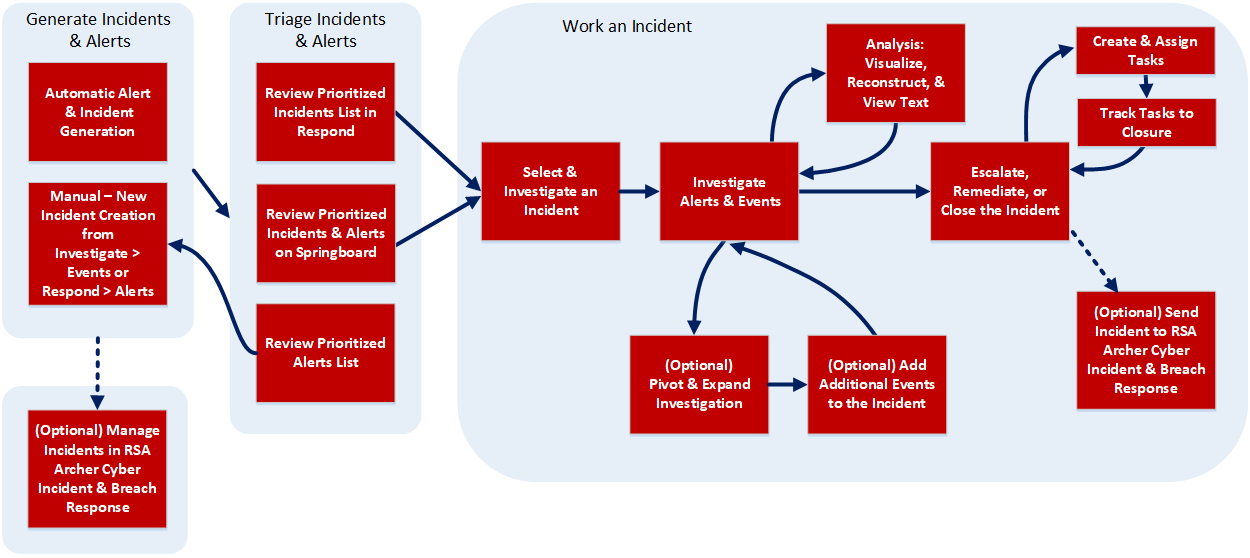
The following figure shows an example of the Respond view - Incident Details view, which shows details for a selected incident.



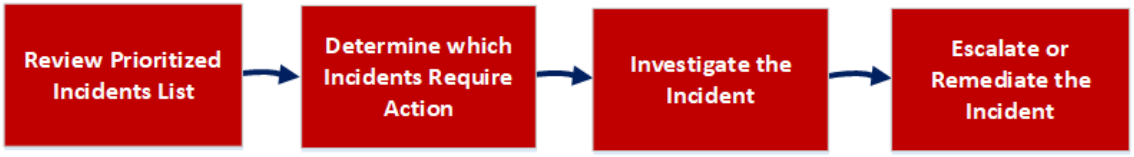
The Respond view is designed to make it easy to evaluate incidents, contextualize that data, collaborate with other analysts, and pivot to a deep-dive investigation as needed. The following figure shows an example of an event analysis in the Incident Details view.



The following figure shows the high-level Respond workflow process.



The following figure shows the high-level process that Incident Responders use to respond to incidents in the Respond view.



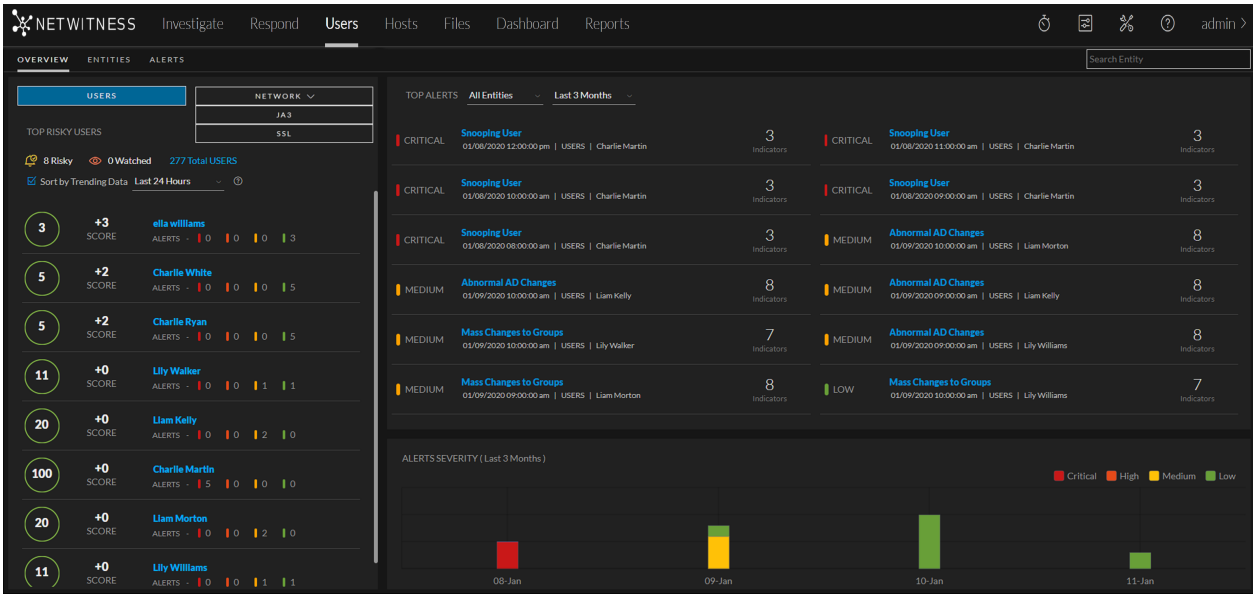
In the Respond view, analysts look at the prioritized list of incidents and determine which incidents require action. They click an incident for a clear picture of the incident with supporting details and they can investigate the incident further. Analysts can then determine how to respond to the threat, by escalating or remediating it.

What can I do here?	Path	Show me how
View prioritized incident lists	Respond > Incidents (Incidents List view)	See the *NetWitness Respond User Guide*.
Determine which incidents require action (Triage an incident)	Respond > Incidents (Incident Details view)	See the *NetWitness Respond User Guide*.
Investigate the incident	Respond > Incidents (Incident Details view)	See the *NetWitness Respond User Guide*. (You can also pivot to the Investigate view.)
Escalate or Remediate the Incident	Respond > Incidents (Incident Details view) and Respond > Tasks (Tasks List view)	See the *NetWitness Respond User Guide*.
Review Alerts	Respond > Alerts (Alerts List and Alert Details views)	See the *NetWitness Respond User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Users

The Users view provides visibility into risky user behaviors across your enterprise with NetWitness UEBA. You can view a list of high-risk users and a summary of the top alerts for risky behavior for your environment. Then you can select a user or an alert and view details about the risky behavior and a timeline during which the behaviors occurred.



The Users menu has the following options:

- **Overview**: It provides an initial view into the recent and most important user or network entity activities in the environment. Each panel shows either prioritized incidents for investigation or consolidated metrics reflecting potential risks to the enterprise.
- **Entities**: It is a proactive threat hunting console. You can use behavioral filters to build use case driven target lists, and to continuously monitor the environment for specific risky behavior patterns.

Note: The Entities view is only available if you are assigned the role of Administrator or UEBA Analyst.

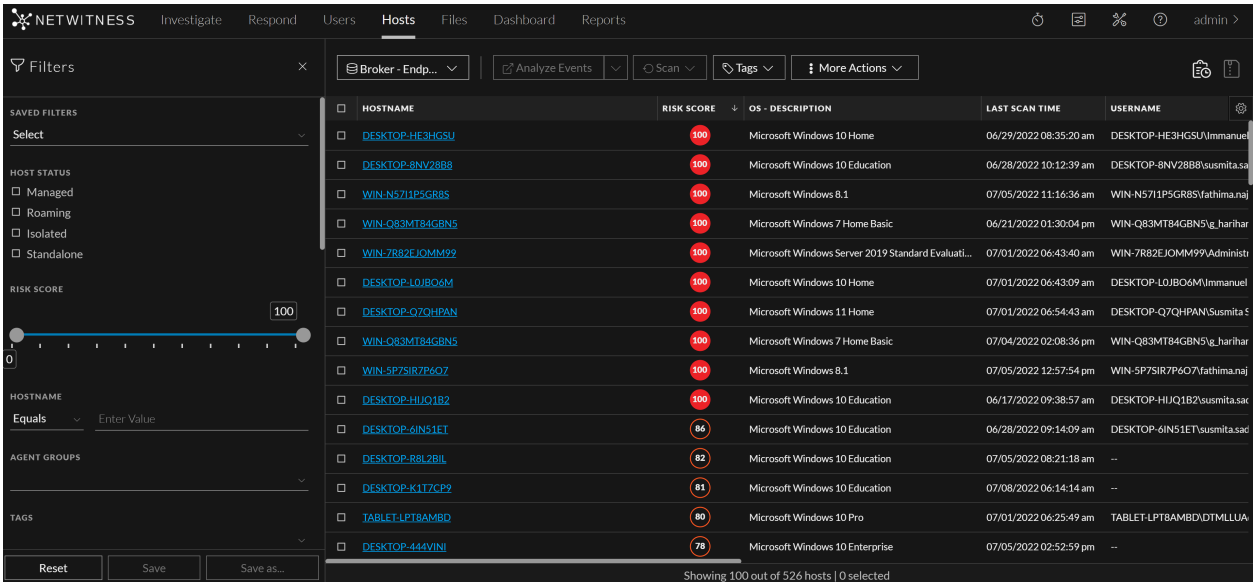
- **Alerts**: It displays details about all the alerts in your environment. You can view forensic information about suspicious activity in your environment that is based on a specific timeframe.

What can I do here?	Path	Show me how
Find Risky User Behavior	Users view	See the *NetWitness UEBA User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Hosts

The Hosts view lists all hosts that have a NetWitness Endpoint agent running. You can filter hosts based on operating system, agent last seen, last scan time, risk score, and other factors. You can open a specific host to view events related to alerts, anomalies, process details, and information related to logged-in users.

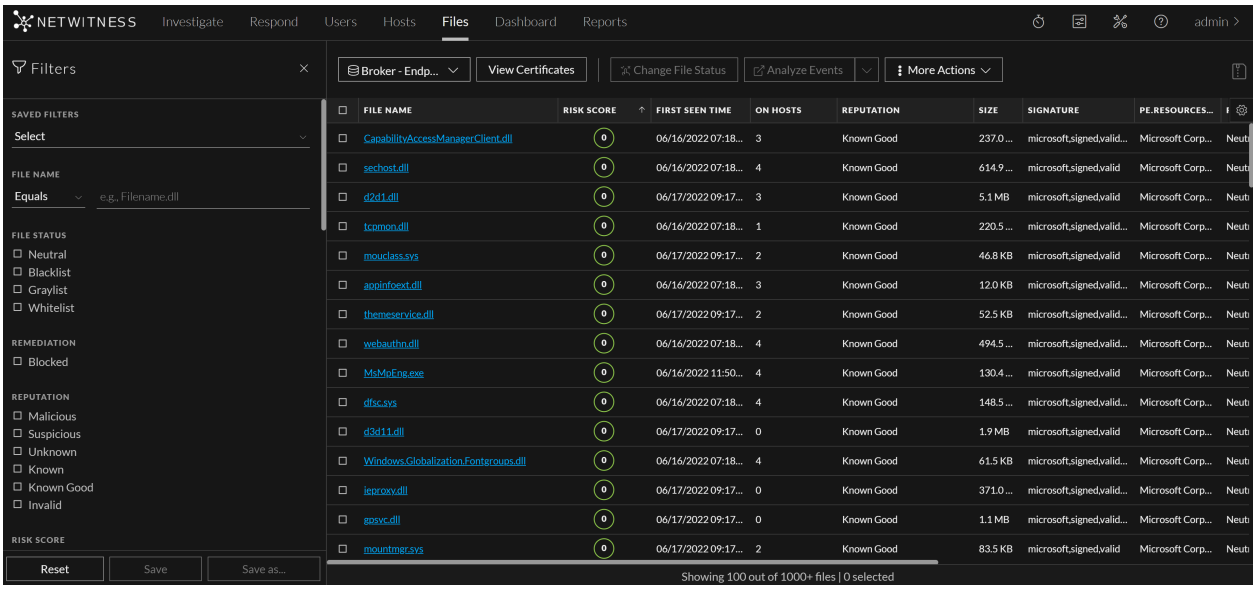


What can I do here?	Path	Show me how
Investigate Endpoints	Hosts view	See the *NetWitness Endpoint User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Files

The Files view provides a holistic view of all files in your deployment. You can apply filters, sort, and categorize files by status to reduce the number of files for analysis, and identify suspicious or malicious files.

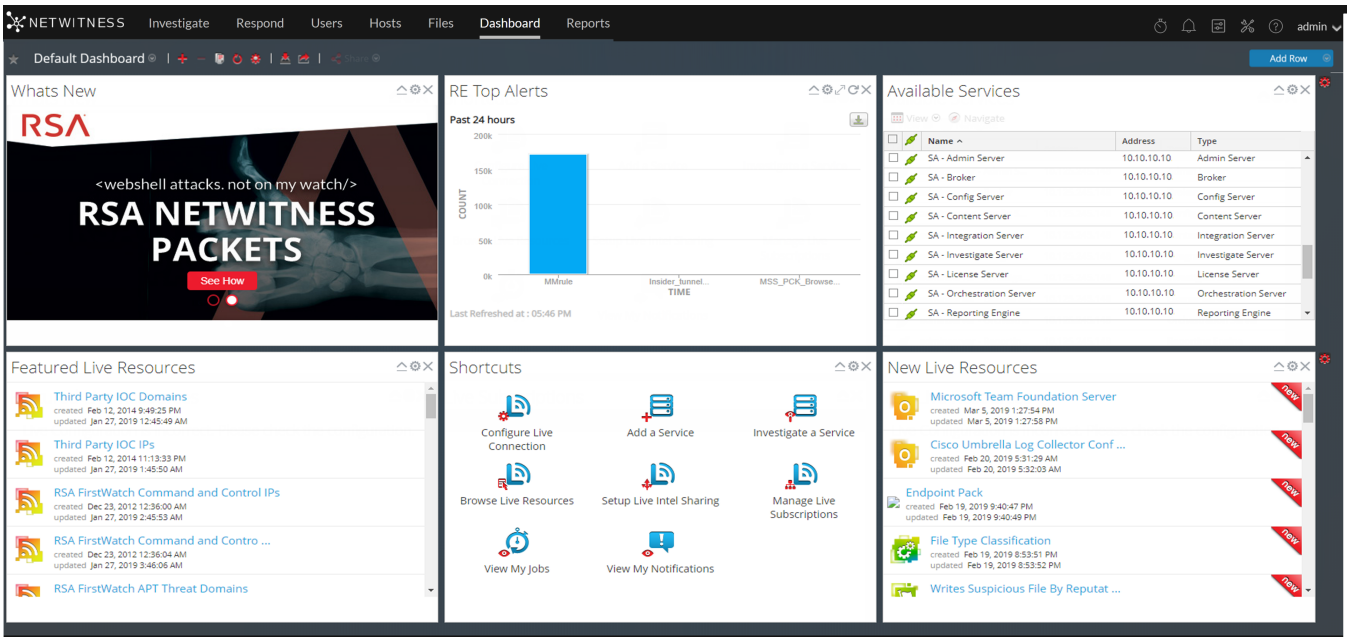


What can I do here?	Path	Show me how
Find Suspicious Endpoint Files	Files view	See the *NetWitness Endpoint User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Dashboard

A dashboard is a group of dashlets that give you the ability to view data in one space, the key snapshots of the various components that you consider important. In NetWitness® Platform, you can compose dashboards to obtain high-level information and metrics that portray the overall picture of a NetWitness Platform deployment, displaying only the information that is most relevant to the day-to-day operations.



NetWitness Platform has predefined dashboards that you can select in the Dashboard view depending on the tasks you perform:

You can select the following preconfigured dashboards:

- Default
- Identity
- Investigation
- Operations - File Analysis
- Operations - Logs
- Operations - Network
- Operations - Protocol Analysis
- Overview
- RSA SecurID
- Threat - Hunting
- Threat - Intrusion
- Threat - Malware Indicators

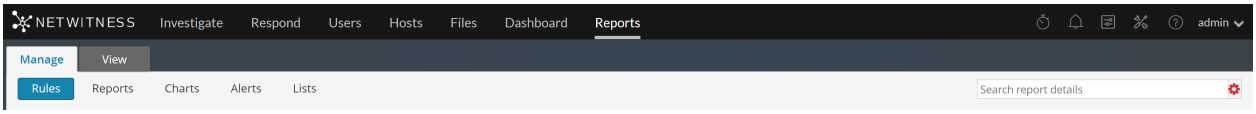
What can I do here?	Path	Show me how
Select a Dashboard	Dashboard view	See Managing Dashboards.
Create a Dashboard	Dashboard view	See Managing Dashboards.
Manage Dashboards	Dashboard view	See Managing Dashboards.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Reports

The Reports view enables you to view and manage reports relevant to your SOC role according to your assigned permissions.

Reports Menu



The Reports menu has the following options:

- **Manage**: This panel allows you to create or modify an rules, reports, charts, alerts, and lists as per the requirement.
- **View**: You can view a report or list of all reports. You can also view the scheduled reports to know the state of the scheduled report. If the scheduled report is in a stop or disable state, you can start or enable the scheduled report.

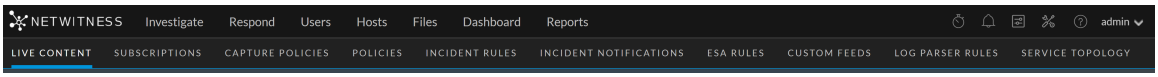
What can I do here?	Path	Show me how
View a Report	Reports > View	See the *Reporting User Guide*.
Manage Reports	Reports > Manage	See the *Reporting User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Configure

The Configure view enables Threat Intel personnel (Content Experts) to configure data sources and inputs to NetWitness in one convenient location.

Configure Menu



The Configure menu has the following options:

- **Live Content** (Live Services): The Live Content view enables you to search for and subscribe to Live Services resources. Live Services is the component of the NetWitness that manages communication and synchronization between NetWitness services and a library of Live content available to NetWitness customers. You can view, search, deploy, and subscribe to content from the RSA Live Content Management System (CMS) to NetWitness services and software. When you subscribe to a

resource, you agree to receive updates on a regular basis from RSA Live Services.

- **Subscriptions** (Live Services): The Subscriptions view enables you manage the Live content that you subscribed to, in the Live Content view. To set up Live Services on NetWitness, you configure the connection and synchronize between the CMS server and NetWitness.
- **Capture Policies**: The Capture Policies view enables you to set up selective network data collection, which gives you the ability to apply centrally managed capture policies across your Network Decoders. This results in better use of service resources, including hard drive space, which leads to more predictable costs and lessens the burden of managing multiple services. You can determine which traffic is stored and how it is stored by using policies. Each policy contains a list of supported base protocols and definitions for handling any other protocols that are detected.
- **Incident Rules**: The Incident Rules view enables you to create incident rules with various criteria to automatically create incidents. You can view prioritized incidents in the Respond view.
- **Incident Notifications**: The Incident Notifications view enables you to automatically send email notifications to SOC Managers and the Analysts assigned to the incidents when incidents are created or updated.
- **ESA Rules**: The ESA Rules view enables you to manage the Event Stream Analysis (ESA) rules that specify criteria for problematic behavior or threatening events in your network. When ESA detects a threat that matches the rule criteria, it generates an alert.
  You can create ESA rules yourself or download them from Live Services. The Rule Library shows all ESA rules created or downloaded. To activate rules, you have to add them to a deployment. Deployments map rules from your rule library to the appropriate ESA services.
- **Custom Feeds** (Live Services): The Custom Feeds view streamlines the task of creating and managing custom feeds, as well as populating the feeds to selected Decoders and Log Decoders. You can set up and maintain custom and identity feeds.
  NetWitness uses feeds to create metadata based on externally defined metadata values. A feed is a list of data that is compared to sessions as they are captured or processed. For each match, additional metadata is created.
  You can create custom feeds to provide extra metadata extraction, for example, to accommodate custom network applications.
- **Log Parser Rules**: The Log Parser Rules tab displays information about individual log parsers, as well as the default, "parse all" parser that can parse logs that are not associated with a particular log parser. This tab contains the following information:
  - You can view the rules for a particular event source type, including the default parser.
  - You can view the names, literals, patterns, and metadata for each configured log parser.
  - You can add log parsers.
  - You can add, edit, and delete custom rules for log parsers.

  Note: The Log Parser Rules tab is available in the Configure menu in versions 11.2 and later. For earlier versions, it is located in Admin > Event Sources.

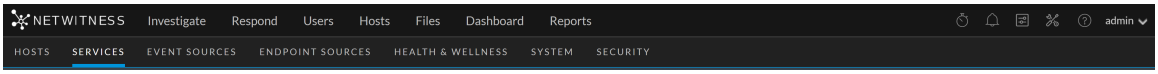
What can I do here?	Path	Show me how
Create a Live Services account.	RSA Live Registration Portal: https://cms.netwitness.com/registration/	See the *Live Services Management Guide*.
Find and deploy Live Services resources.	(Configure) > Live Content	See the *Live Services Management Guide*.
Set up selective network data collection.	(Configure) > Capture Policies	See the *Decoder Configuration Guide*.
Set up Live Services Services on NetWitness.	(Configure) > Subscriptions	See the *Live Services Management Guide*.
Create incidents automatically.	(Configure) > Incident Rules	See the *NetWitness Respond Configuration Guide*.
Configure incident notifications.	(Configure) > Incident Notifications	See the *NetWitness Respond Configuration Guide*.
Configure alerts.	(Configure) > ESA Rules	See the *Alerting with ESA Correlation Rules User Guide*.
Set up and maintain custom and identity feeds.	(Configure) > Custom Feeds	See the *Live Services Management Guide*.
View and edit log parsers and log parser rules.	(Configure) > Log Parser Rules	See the *Log Parser Customization Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Admin

In the Admin view, administrators can manage network hosts and services; monitor the health and wellness of NetWitness; and manage system-level security. They can also configure global system resources and manage event sources.

Admin Menu



The (Admin) menu has the following options:

- **Hosts**: The Hosts view is where you set up and maintain hosts. A host is the machine on which services run and a host can be a physical or virtual machine.

- **Services**: The Services view enables you to manage services, manage service users and roles, maintain service configuration files, and explore and edit service properties. A service performs a unique function, such as a Decoder service, which captures network data in packet form.
- **Event Sources**: The Event Sources view enables you to manage event sources and configure alerting policies for them. Organizations typically monitor event sources in groups based on the criticality of the event sources. You can create monitoring policies for each event source group and order them based on priority.
- **Endpoint Sources**: The Endpoint Sources view enables you to manage and update endpoint agent configurations through groups and manage the agents behavior using policies. You can either use the default policies or customize these policies.
- **Health & Wellness**: The Health & Wellness view enables you to monitor the health of the NetWitness hosts and services in your network environment.
- **System**: The System view enables you to set global NetWitness configurations. You can configure global audit logging, email, system logging, jobs, RSA Live Services, URL integration, Investigation, Event Stream Analysis (ESA), ESA Analytics, and advanced performance settings. In addition, you can manage NetWitness versions and configure the local licensing server.
- **Security**: The Admin Security view provides the capability to manage user accounts, manage user roles, map external groups to NetWitness roles, and modify other security-related system parameters. These apply to the NetWitness system and are used in conjunction with the security settings for individual services.

Note: For versions 11.2 and later, the Event Sources > Log Parser Rules tab can be found in the Configure view.

What can I do here?	Path	Show me how
Manage hosts.	(Admin) > Hosts	See the *Host and Services Getting Started Guide*.
Manage services including managing service user access and security.	(Admin) > Services	See the *Host and Services Getting Started Guide*.
Manage event sources and configure alerting policies for them.	(Admin) > Event Sources	See the *Event Source Management Guide*.
Manage endpoint sources and configure alerting policies for them.	(Admin) > Endpoint Sources	See the *Event Source Management Guide*.
Set up and monitor alarms for the hosts and services in your NetWitness domain.	(Admin) > Health & Wellness > Alarm	See the *System Maintenance Guide*.

What can I do here?	Path	Show me how
Monitor statistics for the NetWitness hosts and the services running on the hosts.	(Admin) > Health & Wellness > Monitoring	See the *System Maintenance Guide*.
Create and apply policies to your hosts and services to help you maintain the health and wellness of your NetWitness domain.	(Admin) > Health & Wellness > Policies	See the *System Maintenance Guide*.
Set global configurations for NetWitness.	(Admin) > System	See the *System Configuration Guide*.
Configure Global Audit Logging.	(Admin) > System > Global Auditing	See the *System Configuration Guide*.
Set up system security.	(Admin) > Security	See the *System Security and User Management Guide*.
Manage system users with roles and permissions.	(Admin) > Security	See the *System Security and User Management Guide*.
Set up Public Key Infrastructure (PKI) authentication. PKI is available in NetWitness 11.3 and later.	(Admin) > Security	See the *System Security and User Management Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Setting Up Your Default View by SOC Role

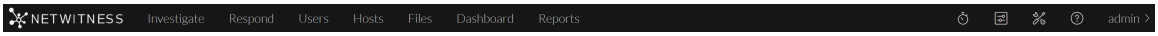
After logging in to NetWitness, you can navigate into the application easier by setting up your default view based on your Security Operations (SOC) role. You can set your default view, also known as a landing page, in your user preferences.

On upgrade to version 11.5 or later, by default the Springboard is displayed if you have not configured the default landing page in previous versions.

In version 11.4 and prior versions, if you configured the default landing page as Respond or Investigate in the User Preferences, then on upgrade to version 11.5, the default landing page will be Respond or Investigate view.

On fresh install of NetWitness 11.5 version, when you log in, by default Springboard is the landing page.

The following figure shows the main NetWitness views.



- **Springboard**: This view is for Analysts, who can see panels for prioritized alerts, incidents, risky hosts, risky users, risky files, and focused event data to help hunt and investigate faster than ever before.
- **Investigate**: This view is for Threat Hunters, who investigate and hunt for advanced threats. Other analysts such as Incident Responders may pivot into this view for deeper analysis of an incident.
- **Respond**: This view is for Incident Responders, who can view a list of incidents to triage and alerts.
- **Users**: This view is for SOC Managers and Analysts to discover, investigate, and monitor risky behaviors across entities namely Users and Network in your environment.
- **Hosts**: This view is for Analysts, who can investigate or perform analysis on hosts using attributes such as IP address, host name, Mac address, risk score, and so on.
- **Files**: This view is for Analysts, who can investigate or perform analysis on files using attributes such as IP address, host name, Mac address, risk score, and so on.
- **Dashboard**: This view is for all users. You can view dashboards on different areas of interest depending on your user permissions. You have the option to select a preconfigured dashboard, import a dashboard, or create your own custom dashboard.
- **Reports**: This view is for all users. You can view reports on different areas of interest depending on your user permissions.
- **Configure**: This view is for Threat Intel personnel (Content Experts), who configure data sources and inputs to NetWitness. Content Experts use this area to download and manage Live content. They can also create and manage incident and ESA rules.
- **Admin**: This view is for System Administrators, who set up and maintain the overall application.

You can select any of the main NetWitness views as your default view. In addition to the main views, NetWitness has predefined dashboards that you can select in the Dashboards view depending on the tasks you perform:

- Default Dashboard
- Identity Dashboard
- Operations - Logs Dashboard
- Operations - Network Dashboard
- Overview Dashboard
- Threat - Indicators Dashboard
- Threat - Intrusion Dashboard

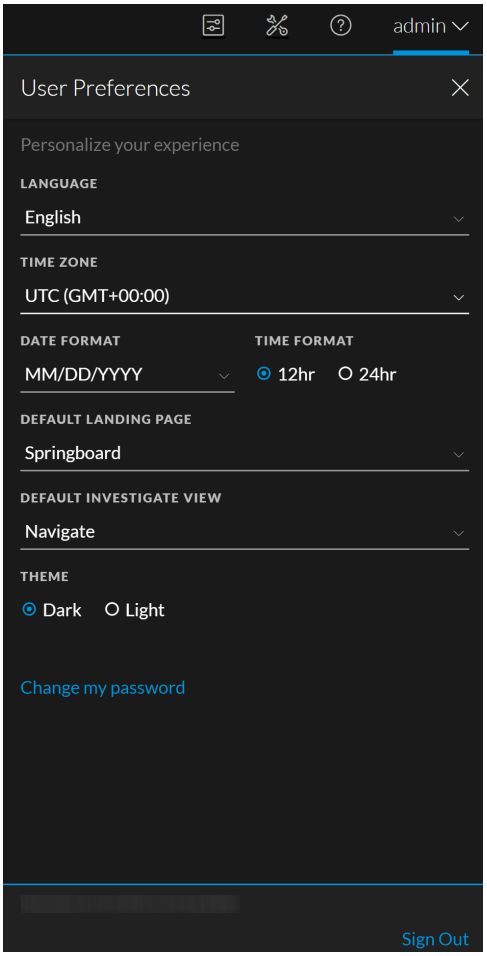
The following table shows typical SOC roles and the available views you can select as your landing page in your user preferences based on your SOC role. If you have more than one role, select the view that is most appropriate for you to start with when you log in to NetWitness.

SOC Roles	Role Description	Consider this Default Landing Page
Incident Responder (Tier 1 Analyst)	Addresses incidents and alerts queued for them to review and mitigate.	**Springboard** or **Respond**
Threat Hunter (Tier 2/Tier 3 Analyst)	Investigates and hunts for advanced threats.	**Springboard**, **Investigate**, **Users**, **Hosts**, or **Files** For information on selecting the default Investigate view, see the *NetWitness Investigate User Guide*.
SOC Manager (SOC Management and Reporting)	Manages SOC readiness and responds to incidents and data breaches.	**Springboard** or **Dashboard** When you log in, select the appropriate predefined dashboard for your SOC role. You can also import a dashboard or create your own dashboard.
Content Expert (Threat Intelligence)	Configures data sources and inputs to NetWitness.	**Dashboard** or  (Configure) When you log in, select the appropriate predefined dashboard for your SOC role. You can also import a dashboard or create your own dashboard. If you choose Dashboard as your default view, you can navigate to the  (Configure) view from the main menu.
Data Privacy Officer (DPO)	Similar to an Administrator, but a DPO monitors and protects privacy-sensitive information.	**Dashboard** When you log in, select the appropriate predefined dashboard for your SOC role. You can also import a dashboard or create your own dashboard.

SOC Roles	Role Description	Consider this Default Landing Page
System Administrator	Focuses on the configuration and stability of the overall application. Manages user access.	**Springboard**

Set Your Default View

1. On the main menu bar, select your username, for example, admin.
   The User Preferences dialog shows your current preferences.



2. In the **Default Landing Page** field, select the default view that you would like to see when you log in to NetWitness. Use the above table to make your selection based on your SOC role. For example, if you are an Analyst, you can select **Springboard**; if you are an Incident Responder you can select **Respond**; and if you are a Threat Hunter, you can select **Investigate**.

   Your preferences become effective immediately. You can change your default landing page at any time. For information on other preferences, see Setting User Preferences.

3. To verify that you can see the correct default view, click **Sign Out** to log out and then log back in to NetWitness.

Managing the Springboard

(From 11.5 and later) NetWitness Platform Springboard presents platform-wide detections and signals in this view so analysts hunt and investigate faster than ever before.

From 12.0 and later, NetWitness Platform Springboard introduces five more out-of-the-box panels based on the events processed with a specific query and presented on Springboard view. This helps analysts with further investigation.

The Springboard congregates the following information for analysts to view:

- Critical incidents and high severity alerts that require attention.
- Hosts and files with high risk scores that may be potential threats.
- Risky users that are potential leads for investigation.
- Events with specific query and high severity that require immediate attention.



The Springboard displays important information for the last 24 hours in the following out-of-the-box panels:

- **Top Incidents**
- **Top Alerts**
- **Top Risky Hosts**
- **Top Risky Files**

- **Top Risky Users**
- **MITRE ATT&CK tactics**
- **MITRE ATT&CK techniques**
- **Indicators of Compromise**
- **Enablers of Compromise**
- **Behaviors of Compromise**

For example, the Top Risky Hosts displays the top 25 risky hosts based on the highest risk score and Operating system (Windows, Linux, and Mac). The result displays hosts of all Endpoint Servers if the Endpoint Broker is available. Otherwise, it displays the result of the first Endpoint Server.

You can perform the following actions on the Springboard:

- Change the time range for some panels namely Incidents and Alerts panels. To change the time range, select the time range selection box from the drop-down menu in the top left corner of the Springboard view.

  IMPORTANT: If the selected filter has a time range selected, it is given priority, otherwise the Springboard time range for the specified panel is considered.

- Increase the display of the results in the table to view more than 25 results. Click  on the panel, the Edit Panel dialog is displayed. Edit the number of results field and click **Save Panel**.
- Click a row in the table to view details or to investigate.
- Click  at the top of the panel to view all the results. For example, in the Top Incidents panel, click  to view all incidents in the **Incidents** list view.
- Click a row name in the events panel to view or investigate the event details with relevant filters applied in the **Events** view.
- Scroll to view the different panels using the  scroll bar available below the panels.

Administrators can customize the Springboard by performing the following:

- Create own custom private board and add panels on the board. For more information, see Add a Custom Private Board.
- Edit the out-of-the-box panels. For more information, see Edit a Panel.
- Refresh the out-of-the-box panels. For more information, see Refresh a Panel.
- Create new panels with important system indicators. For example, a new panel showing focused event metadata based on pre-defined query conditions can be created. For more information, see Add a Panel.

Working with the Springboard

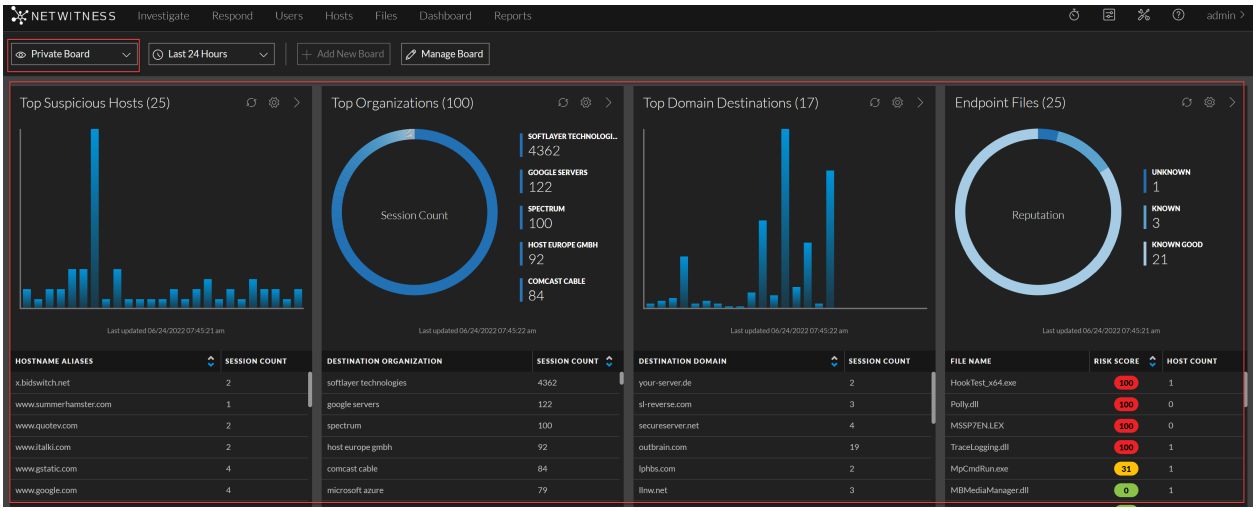
Note: An administrator must provide the appropriate permissions to allow users to edit the springboard panels. For more information see the the Springboard section in the "Role Permissions" topic in the *System Security and User Management Guide*.

You can customize the information on the out-of-the-box Springboard by adding, editing, copying, moving, and deleting panels.

From 12.0 and later, the data sources and query filters are automatically added for the new out-of-the-box panels.

Add a Custom Private Board

(From 12.0 and later), Administrators and Analysts can create their own custom private board in the springboard and add panels with important system indicators, which helps in threat hunting and investigation. The users can also add, edit, rearrange, and delete panels in the custom private board view. The board allows users to organize and manage information in an easy manner.

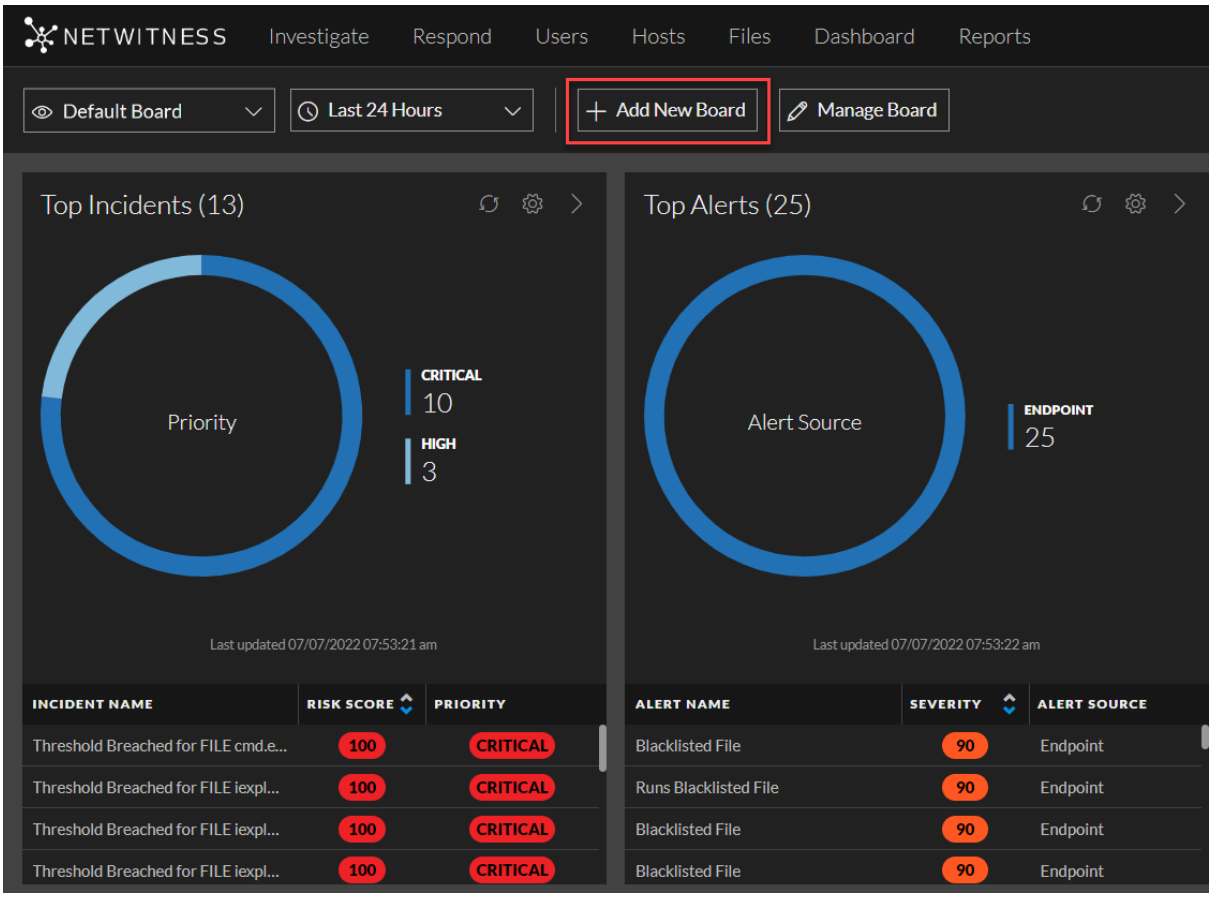


IMPORTANT:
- The board will be saved as a custom private board, and other users will not be able to view the board.
- Only one custom private board can be created.

Note: The maximum number of panels on the custom private board must not exceed 20 panels.

To add a custom Board:

1. Click + **Add New Board**.



   It navigates to a custom private board view to add panels.
2. Add the panels. For more information, see Add a Panel.
3. To edit the custom board's name, click at the top left corner and enter a unique name.
4. Click **Save Board**.

Add a Panel

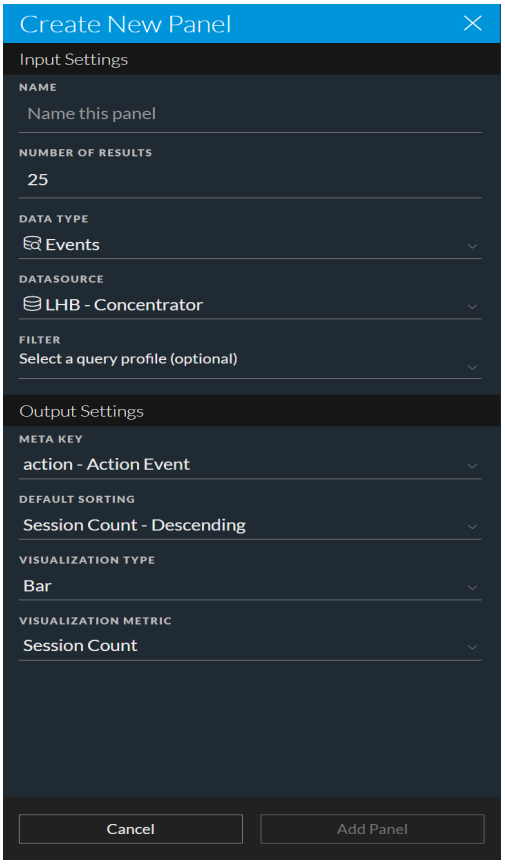
You can add a panel to the Springboard according to the analyst preferences. For example, an analyst can watch top risky users or top risky hosts for a particular region in a panel.

Note: The maximum number of panels on the Springboard should not exceed 20 panels.

To add a panel:

1. Click **Manage Board**.
2. Click + New Panel either on the top or on the right side of the view or click ⊞ at the bottom of the view to add a panel.

The Create New Panel dialog is displayed. The following figure is an example of the events panel configuration.



3. In the Input Settings section:
   - **Name**: Enter a unique name for the panel. The name can include letters, numbers, spaces, and special characters, such as `_ - ( ) [ ]`.
   - **Number of Results**: By default, the number of results is 25. Specify the number of results that range from 25 to 100.
   - **Data Type**: Select the type of data to use for the panel:
     - Alerts
     - Incidents
     - Events
     - Files
     - Hosts
     - Users
   - **Data Source**: Select the source of the data to use for the panel. This field is enabled when the data type is Events, Files, or Hosts.

    - Events: Select either Broker or Concentrator.
    - Files: Select either Endpoint Broker Server or Endpoint Server.
    - Hosts: Select either Endpoint Broker Server or Endpoint Server.
  - (Optional) **Filter** : Filter the data as required from the drop-down for each data type from the saved filters list.
4. In the Output Settings section, select the appropriate settings based on the data type.
5. Click **Add Panel**.
6. Click **Save Board** once you have added all the panels.

Edit a Panel

You can edit the out-of-the-box or newly added panels on the Springboard.

To edit a panel:

1. Click  on the panel that you want to edit.
   The Edit Panel dialog is displayed.
2. Edit and click **Save Panel**.

Rearrange Panels

You can arrange the panels by dragging and dropping them into a different order on the Springboard.

To rearrange panels:

1. Click **Manage Board**.
2. To move a panel, click anywhere on the panel, drag and drop the panel to the desired location.
3. Click **Save Board**.

Delete Panels

You can delete panels permanently in the following situations:

- Services are not installed. For example, if you do not have Endpoint Log Hybrid installed, then you can delete the panels for Top Risky Hosts and Files.
- The maximum number of panels have exceeded the limit, that is 20, and you want to add a new panel.

To delete existing panels:

1. Click **Manage Board**.
2. Select the panels that you want to delete.
3. Click **Remove Panel**.
4. Click **Save Board**.

Restore System Default Settings

Note: This is enabled only if any changes are made to the out-of-the-box Springboard panels.

To restore the out-of-the-box panels:

1. Click **Manage Board**.
2. Click **Restore System Default**.

   A confirmation pop-up is displayed to confirm if you want to restore the out-of-the-box panels or not.
3. Click **Restore System Default**.

Refresh a Panel

To refresh a panel:

Click  on the panel that you want to refresh, it loads the latest data in the panel.

Managing Dashboards

A dashboard is a group of dashlets that give you the ability to view data in one space, the key snapshots of the various components that you consider important. In NetWitness, you can compose dashboards to obtain high-level information and metrics that portray the overall picture of a NetWitness deployment, displaying only the information that is most relevant to the day-to-day operations.

The dashboards for all NetWitness components are available to add to the default NetWitness dashboard or a custom NetWitness dashboard.

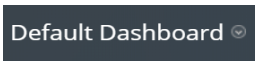
You can view dashboards on different areas of interest depending on your user permissions. You have the option to select a preconfigured dashboard, import a dashboard, or create your own custom dashboard. The dashboards help you to quickly and easily view reports. You can configure your dashboards to display the information that supports your workflow. This topic explains the high-level tasks that can be done when you are setting up a dashboard.

Dashboard Basics

If the Dashboard view is your default landing page following logging in to NetWitness, you always see either the default dashboard or the currently configured dashboard immediately after completing the login process. To return to the dashboard from another NetWitness component, click **Dashboard**.

Dashboard Title

The dashboard title reflects the currently active dashboard; for example, Default Dashboard.



Dashboard Selection List

You can access preconfigured and custom dashboards on the dashboard selection list. When you select a dashboard, its title is displayed below the NetWitness toolbar.



A dashboard has:

- The dashboard toolbar
- The dashboard title and the dashboard selection list.

Dashboard Toolbar

The dashboard toolbar is available next to the title of the selected dashboard. The dashboard toolbar allows various operations on dashboards and dashlets.



Note: The Copy, Delete, Import, Export, Share, and Add Row options are disabled for preconfigured dashboards.

Option	Description
★	Sets the selected dashboard as the Favorite.
Default Dashboard	Displays the list of available dashboards from which you can make a selection.
+	Displays the Create a Dashboard dialog, where you define or add a custom dashboard.
−	Deletes a custom dashboard. The default dashboard cannot be deleted.

Option	Description
	Allows you to copy a dashboard.
	Displays the Manage Dashlet dialog.
	Exports a dashboard as a .zip file.
	Imports a dashboard as .zip or.cfg file.
	Allows you to share a dashboard with another user.
	Enables user to add rows and columns to the dashboard based on the requirement. Click the  icon in a row to add a dashlet.

The Default Dashboard

The default dashboard is configured to display specific dashlets in specific positions. The default dashboard serves as an example of dashboard composition and a starting point for customization.

- You can customize the information on the default dashboard by editing, adding, moving, maximizing, and deleting dashlets.
- After modifying the default dashboard, you can restore the default dashboard () to its original layout.
- The default dashboard cannot be deleted or shared.

Selecting a Preconfigured Dashboard

On installation of NetWitness, the following preconfigured dashboards are automatically activated and are available to you:

- Default
- Identity
- Investigation
- Operations - File
- Operations - Logs
- Operations - Network
- Operations - Protocol Analysis
- Overview

- RSA SecurID
- Threat - Hunting
- Threat - Intrusion
- Threat - Malware Indicators

You cannot perform the following actions on a preconfigured dashboard:

- Edit a dashboard
- Export a dashboard
- Share a dashboard
- Delete a dashboard

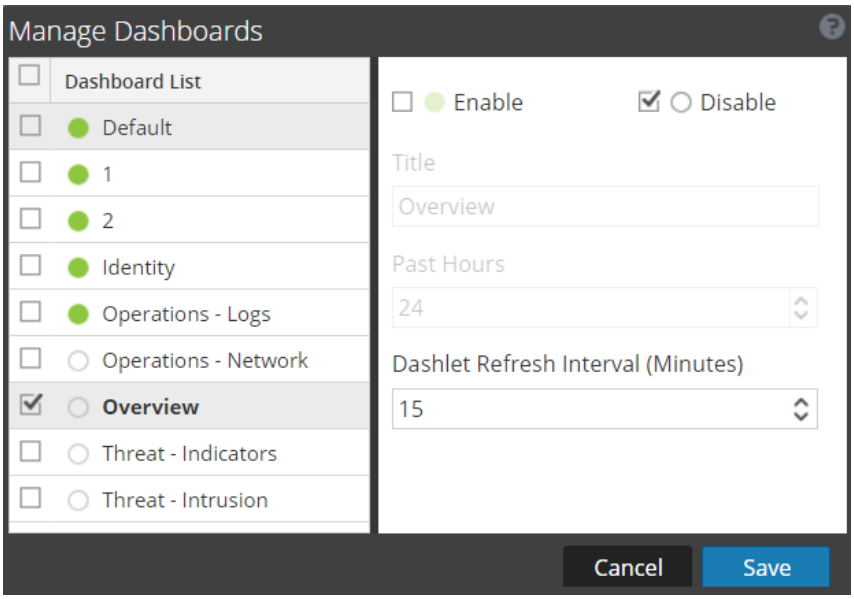
For more information on each preconfigured dashboard, see the Dashboards Catalog in the NetWitness Content space on NetWitness Community.

Note: If you are logged into an Analyst UI, you can view but cannot enable, disable, or edit preconfigured Dashboards.

Enabling or Disabling Dashboards

When you enable or disable a dashboard, all the dashlets within the dashboard are enabled or disabled along with the associated charts, unless they are used in any other dashboard.

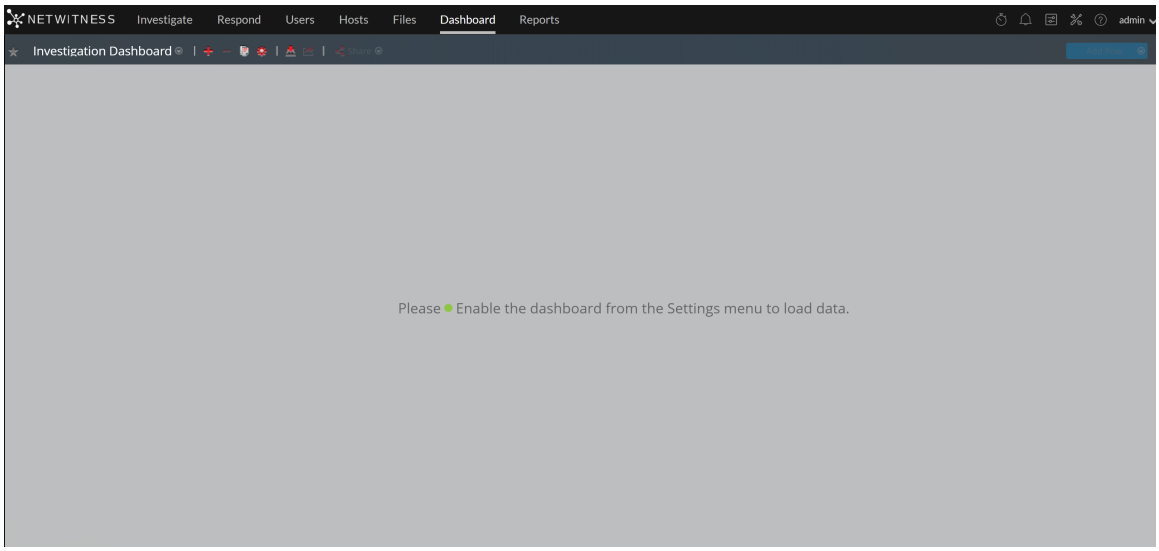
NetWitness modules can display only those dashlets presented in the Manage Dashboard dialog. The main dashboard offers all NetWitness dashlets. This is an example of currently available dashlets.



Name	Description
Dashboard List	Displays a list of the default, preconfigured, and custom dashboards.
Enable	Indicates if the selected dashlet is enabled.
Disable	Indicates if the selected dashlet is disabled.
Title	Displays the title of the selected dashlet and you can also rename the dashboard.
Past Hours	Displays the time for which the data is collected.
Dashlet Refresh Intervals (Minutes)	Displays the refresh interval time of a dashlet.

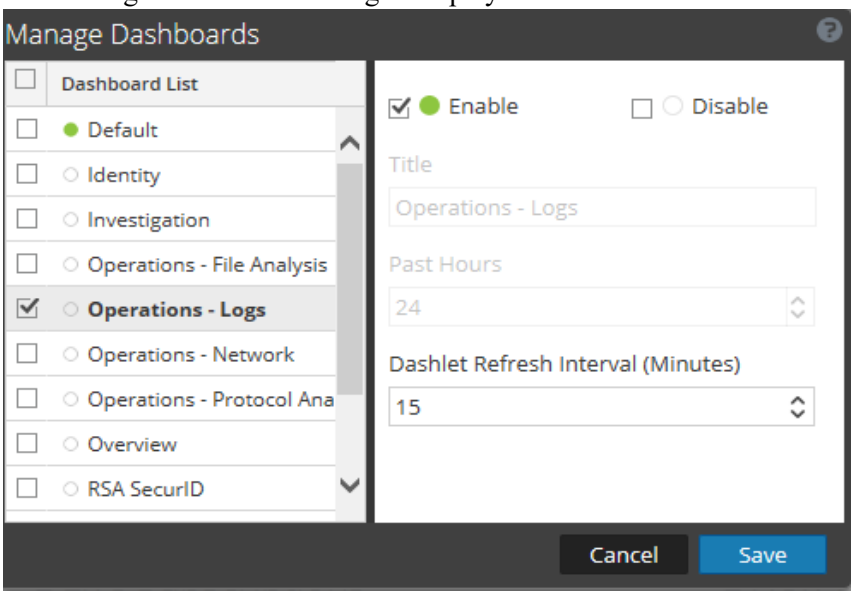
Enable a Dashboard

If you select a dashboard that is not enabled, a masked screen is displayed.



To enable one or more dashboards:

1. Navigate to the dashboard to be enabled.
2. In the dashboard toolbar, click  (Manage Dashboards).
   The Manage Dashboards dialog is displayed.

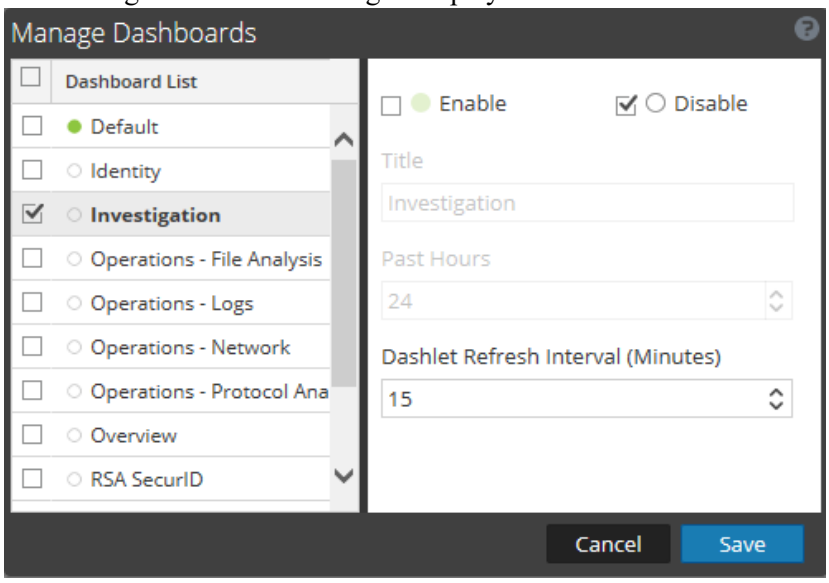


3. From the dashboard list, select the dashboards to be enabled.
4. Select the **Enable** checkbox.
5. Click **Save**.

Disable a Dashboard

To disable one or more dashboards:

1. Navigate to the dashboard to be disabled.
2. In the dashboard toolbar, click  (Manage Dashboards).
   The Manage Dashboards dialog is displayed.



3. From the dashboard list, select the dashboards to be disabled.
4. Select the **Disable** checkbox.
5. Click **Save**.

Setting a Dashboard as a Favorite

To customize the views in NetWitness, you can set a preconfigured or custom dashboard as a Favorite. The NetWitness dashboard offers all NetWitness dashlets. The Favorite dialog sets a specific dashboard as your favorite dashboard and is listed as favorite every time you log in to NetWitness.

1. Navigate to any dashboard.
2. In the dashboard toolbar, click 

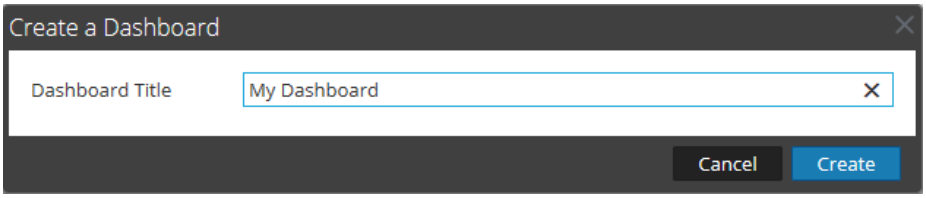
   If the favorite icon is red in color, it indicates that selected dashboard is set as a Favorite and is listed on top above the line.

Creating Custom Dashboards

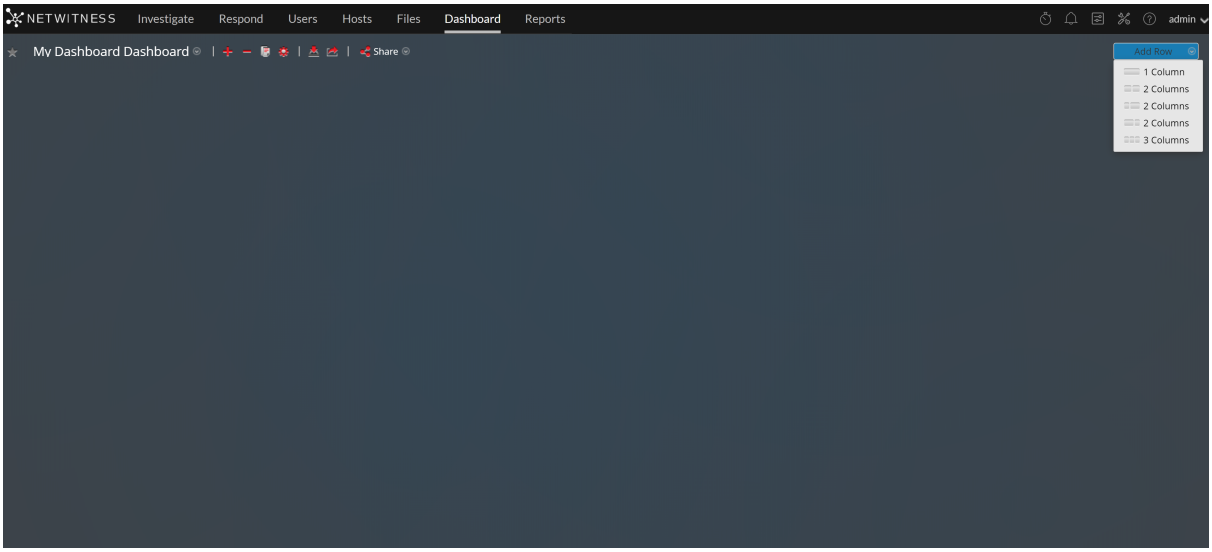
You can create custom dashboards to serve a particular purpose; for example, to represent a specific geographical or functional area of the network. Each custom dashboard is appended to the dashboard selection list.

To create a custom dashboard:

1. In the dashboard toolbar, click .
   The Create a Dashboard dialog is displayed.

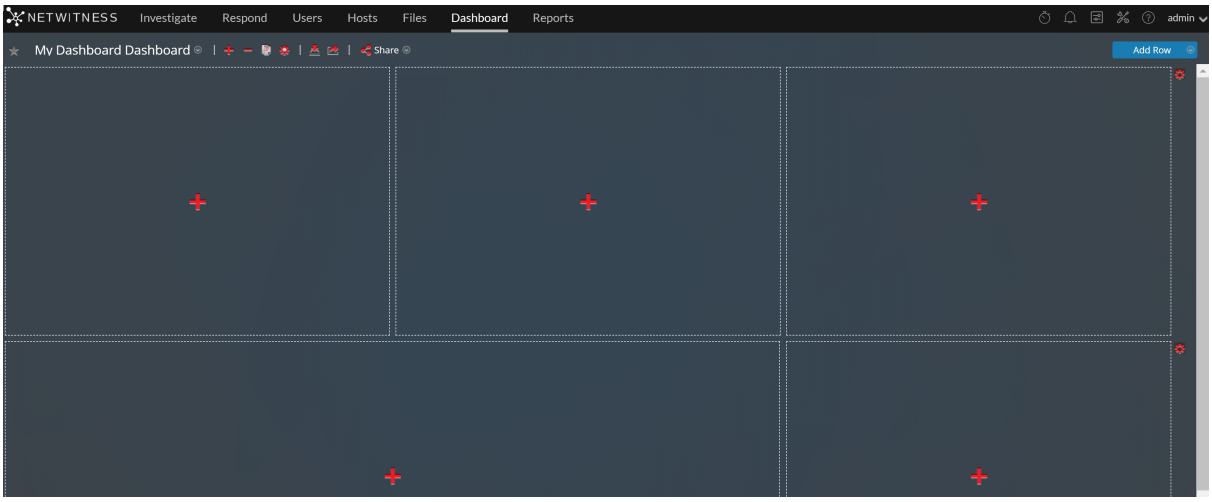


2. Enter a title for the new dashboard and click **Create**.
   The new dashboard is displayed as a blank screen.



3. Add rows to the dashboard, which can contain one or more columns, using the **Add Row** option on the right side of the screen (). Click the desired column configuration in the drop-down list to add one row to the dashboard with the selected number of columns. Repeat the process

to add more rows.



4. You can add any desired dashlets to the dashboard by clicking  in an empty placeholder in a row. For complete details on adding and managing dashlets, see Working with Dashlets.

After custom dashboards are created, you can:

- Switch between dashboards by selecting an option from the dashboard selection list.
- Delete any custom dashboard.
- Import or export a dashboard.

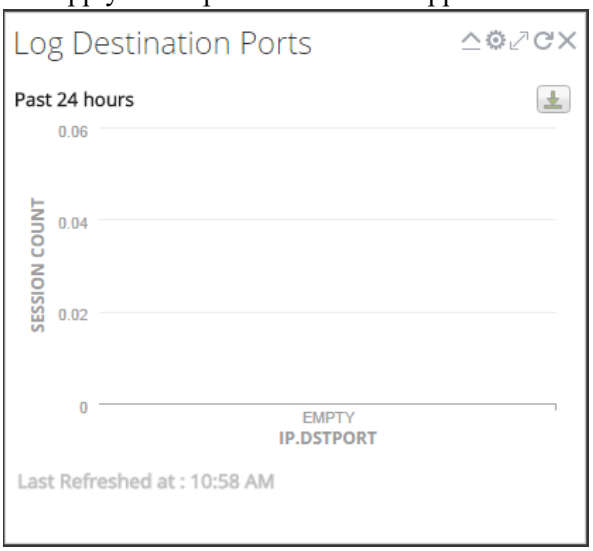
Each dashboard has:

- The dashboard toolbar.
- The dashboard title and the dashboard selection list.
- Zero or more dashlets.

Working with Dashlets

Dashlets are the parts that make up a dashboard. NetWitness uses dashlets to display focused subsets of system information, services, jobs, resources, subscriptions, rules, and other information.

The controls for a dashlet are in the title bar. All dashlets use a common set of controls, and only those that apply to the particular dashlet appear in the title bar of the dashlet.



The following table displays the description of each icon on the dashlet.

Icon	Name	Description
	Collapse vertically	Collapses the dashlet vertically so that only the title is visible.
	Expand vertically	Expands the dashlet to its original size.
	Reload	Reloads the dashlet.
	Settings	Displays configurable settings for the dashlet.
	Maximize	In some dashlets with content that does not fit horizontally within the width of the dashlet, maximizes a chart or a dashlet to full screen.
	Delete	Deletes the dashlet from the dashboard.
Last Refreshed at		Displays the time at which the data is polled from the related chart.
View More		When clicked, navigates to the corresponding dashboard which is linked to the main dashlet and displays more details. If you have not linked the dashboard to an existing dashlet, this link with not be available on the dashlet. To configure this option, click , and in the Dashboard Link field select a related dashboard view more details of the specific dashlet.

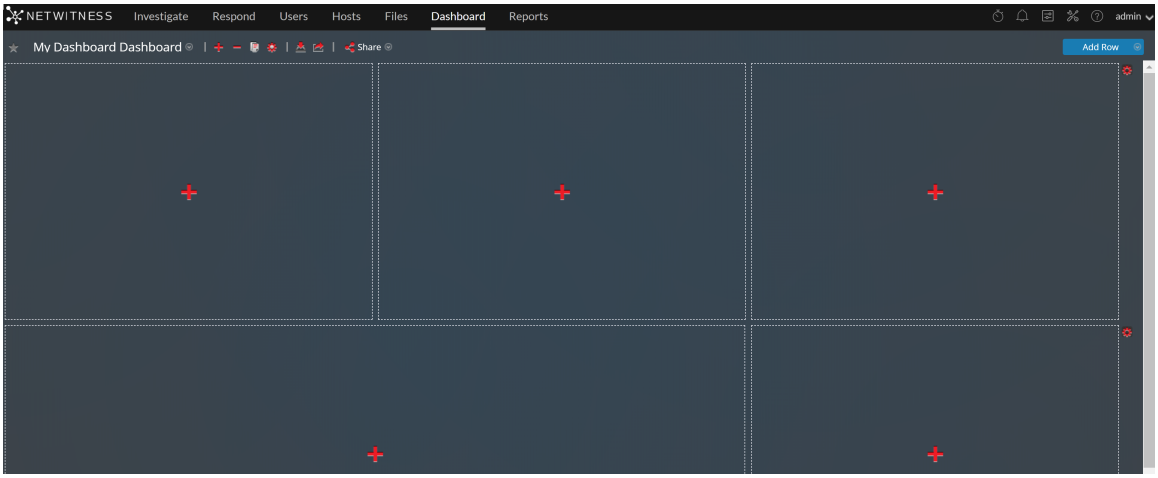
You can add dashlets to the default dashboard or construct a custom dashboard with your own useful set of dashlets to make your workflow more efficient.

Add a Dashlet

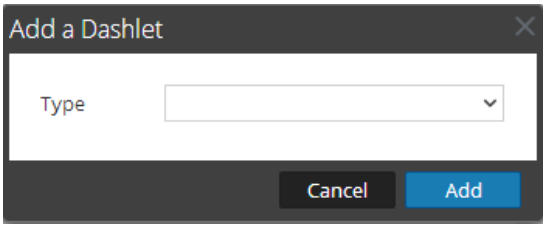
To customize the views in NetWitness, you can add dashlets to a default dashboard or create custom dashboards. However, you cannot add dashlets to preconfigured dashboards.

To add a dashlet:

1. Navigate to any dashboard or create a new dashboard.

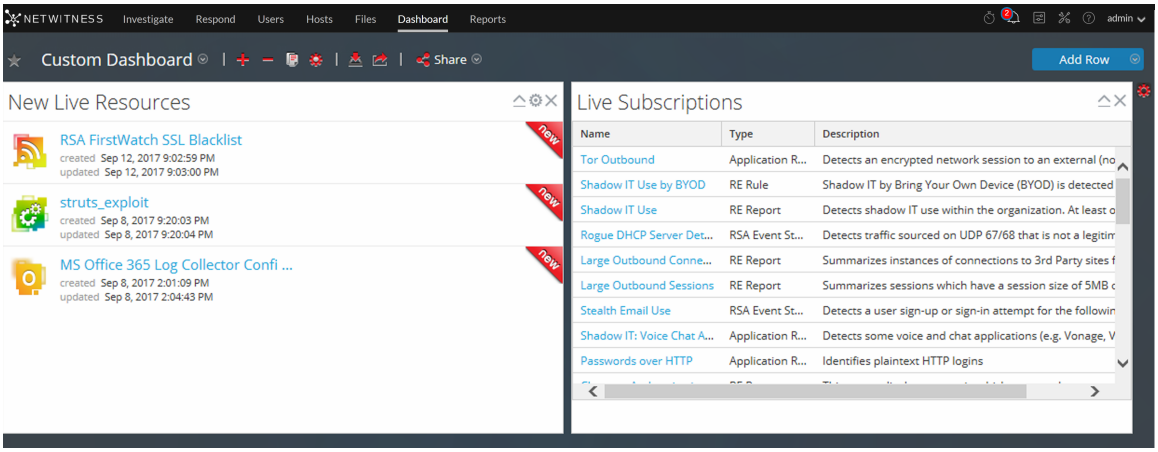


2. Click  on the placeholder where you want to add the dashlet.
   The Add a Dashlet dialog is displayed.



3. Click the **Type** selection list to view the available dashlets, and select the type of dashlet you want to add. Depending on the type of dashlet you are adding, some configurable fields appear in the **Add a Dashlet** dialog.
4. Type a title for the dashlet. The title can include letters, numbers, special characters, and spaces.
5. If there are additional configurable fields for the dashlet, set appropriate values.

6. When all required fields have been configured, click **Add**.
   The dashlet is added to the dashboard in the selected placeholder and is automatically saved.



Edit Dashlet Properties

All preconfigured dashlets are read-only and their properties cannot be edited. Other dashlets are editable and allow users to customize some aspect of the data displayed in the dashlet. A dashlet with editable properties has a settings ( ) option that displays all the editing options.

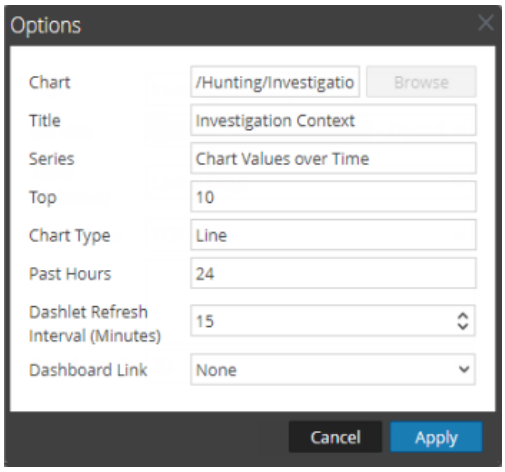
After the dashlets are added, you can drag and drop them and they can be swapped.

A dashlet without editable properties, such as the Live Subscriptions dashlet, does not display the settings option in the title bar. Many dashlets have an editable title where you can edit the following properties:

- Dashlet display title.
- Type of services to monitor; for example, you can monitor only Decoders, or you can monitor Decoders and Concentrators.

Other dashlets have parameters that you define to specify the kind and amount of information you want to see in the dashlet. For example, a Realtime Chart Dashlet has the settings option.

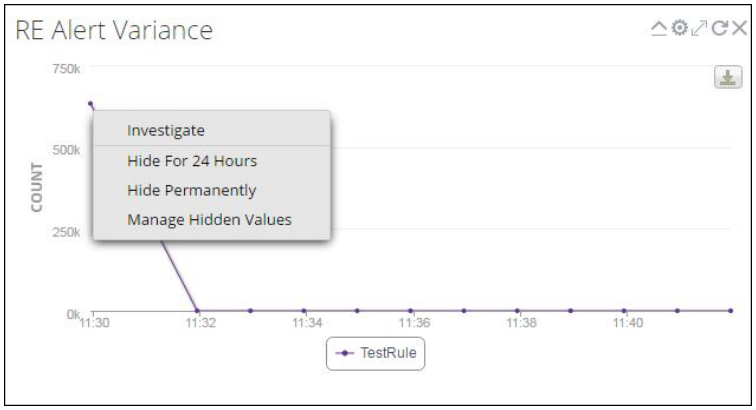
1. To display and modify the options for a dashlet, click settings ( ) in the dashlet title bar.
   The Options dialog is displayed.



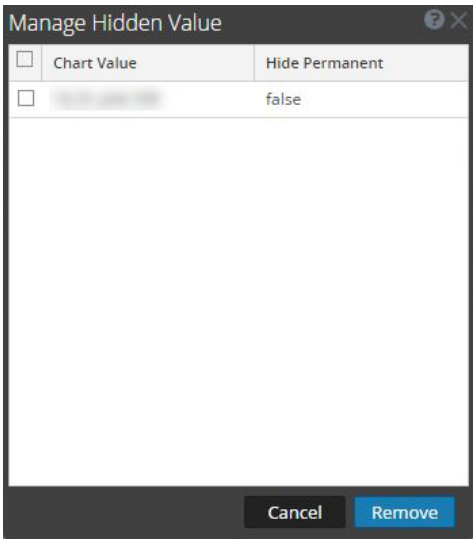
2. Edit any of the displayed properties. For example, in an in Operations - File Analysis Dashboard, edit Dashlet Refresh Interval from 15 to 20.
3. Click **Apply**.

Some dashlets have configuration options to tailor the appearance or the contents of the dashlet. The following options are available for RE Top Alerts, RE Alert Variance, and RE Realtime Charts dashlets on left-click:

- **Hide For 24 Hours**: This option allows you to hide the selected value for the next 24 hours. After 24 hours, the data is automatically displayed on the dashlet, if the value is configured and listed on top.
- **Hide Permanently**: This option allows you to hide the selected value permanently until you add it back using the Manage Hidden Values option.



- **Manage Hidden Values**: This option displays a list of all the hidden values. You can select the checkbox for a value and click **Remove** to view the data back on the chart.



The options to Hide for 24 Hours, Hide Permanently, and Manage Hidden Values are not available for Geomap charts.

Note: When you edit a value in a preconfigured dashboard, it is a user-specific change. The changes made to a preconfigured dashboard are applicable only to your dashboard and cannot be viewed by other users who use the same preconfigured dashboard. For example, if you hide a value in an overview dashboard, the change is applicable only to your dashboard. If another user views the same overview dashboard, the value is still displayed. The same applies to a custom dashboard. When you hide a value in the custom dashboard and share the same dashboard with another user, the values are still displayed even though the dashboard is shared.

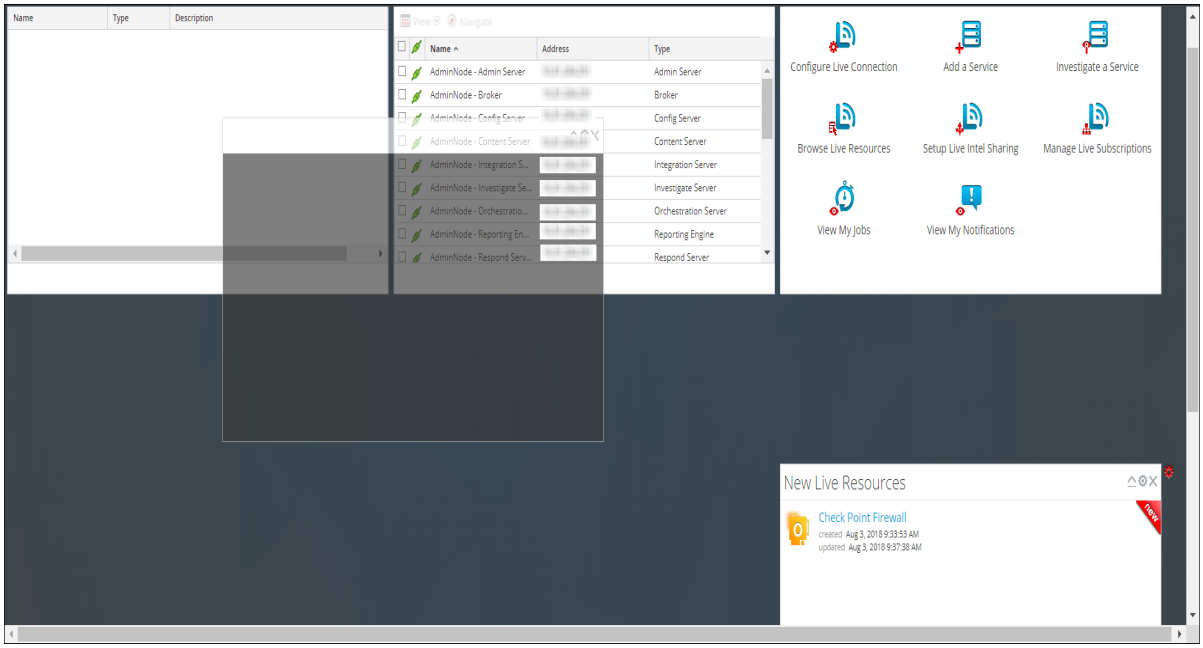
For more information on available dashlets, see the Dashboards Catalog in the NetWitness Content space on NetWitness Community.

Rearrange a Dashlet

You can arrange dashlets according to your preference by dragging and dropping them into a different order on the dashboard.

1. To move a dashlet, hover in the header of the dashlet that you want to move.
   The directional cursor appears over the dashlet. Click and hold in the header of the dashlet that you want to move.

2. Continue to hold the left mouse button and drag the window toward the new location.
   The figure below shows a dashlet as it is being re-arranged.



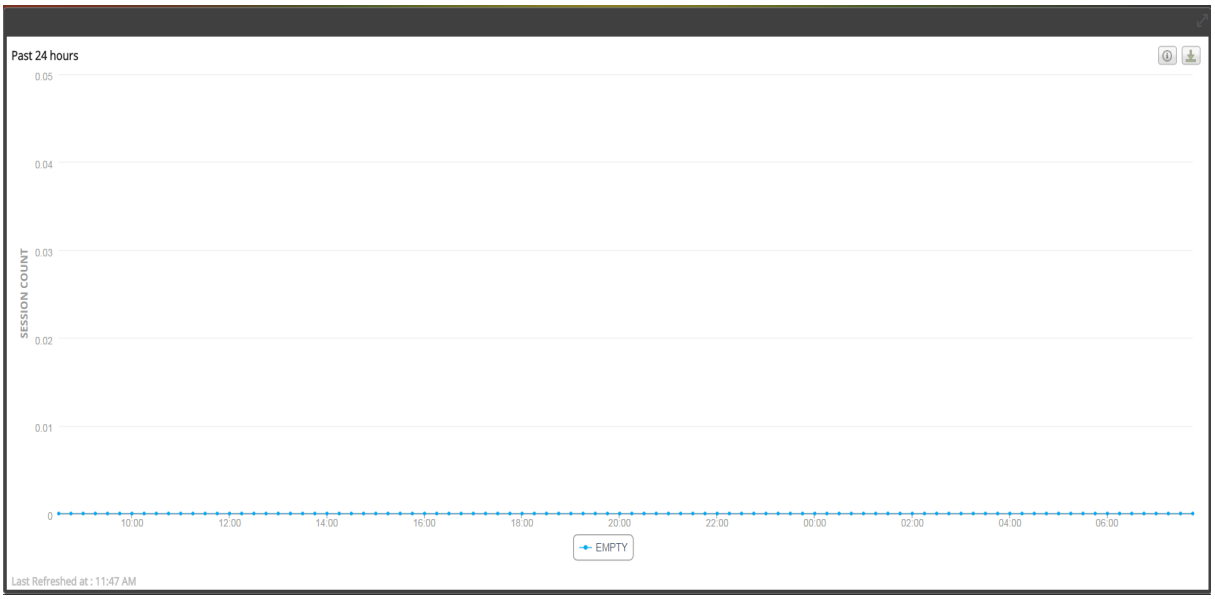
3. Release the mouse button when the dashlet is in the desired location.
   The dashlet that currently occupies that position moves down.

Maximize a Single Dashlet

This section explains how to open a dashlet on the entire area of the main NetWitness dashboard with the same dashlet title. Dashlets that have a lot of columns or charts, for example some Reporting dashlets, are easier to view when maximized so that the entire contents is visible without scrolling.

To maximize a dashlet, click the maximize control icon in the dashlet title bar: ⤢. The dashlet is displayed on full screen.

To minimize a dashlet, click the same control icon in the dashlet title bar: ⤢. The dashlet is restored to previous size.



Delete a Dashlet

1. Click ✕ in the dashlet title bar:
   A confirmation pop-up is displayed to confirm if you want to delete the dashlet.
2. Click **Yes**, if you want to delete. The dashlet is removed from the dashboard.
   Click **No**, if you do not want to delete.

Note: After you remove the dashlet, the empty space is replaced by a placeholder where you can add another dashlet using the above Add a Dashlet procedure.

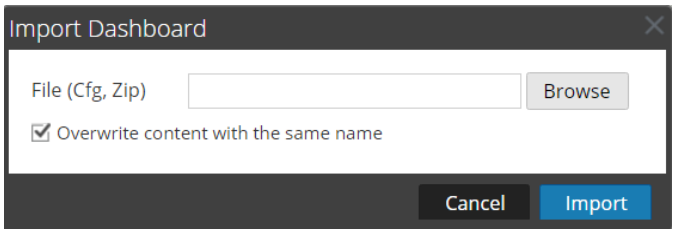
Importing and Exporting Dashboards

The ability to customize dashboards to changing circumstances and conditions could result in a large number of dashboards that are not needed on a daily basis. Rather than reinvent the wheel each time you want to recreate a particular custom dashboard, you can export your dashboards that are not currently in use. When you are ready to use a previously exported dashboard, import the dashboard into NetWitness.

Import a Dashboard

Note: You can import the Reporter Realtime Charts dashboard and its related charts in different instances of the NetWitness server and Reporting Engine from which it was exported.

1. In the dashboard toolbar, click (Import Dashboard).
   The Import Dashboard dialog is displayed.



2. Browse to the dashboard file in the **Import Dashboard** dialog. You can import .cfg and .zip files.
3. Click **Import**.
   The dashboard is displayed in NetWitness

Note: If you import a dashboard from Security Analytics 10.6. x into NetWitness11.x, the dashboard and the associated rules and charts must be imported separately. But when you import a dashboard from NetWitness11.x into NetWitness, the dashboard and all the rules and charts associated with it are imported in .zip format.

Export a Dashboard

Note: When you export a Reporter Realtime dashboard, the corresponding Reporting Engine contents is also exported.

Exported dashboards are designed to work within the same NetWitness instance. It is also possible to share your custom dashboards with other users in your organization, provided they have equivalent permissions.

To export a dashboard, you must have the dashboard open to access the Export Dashboard option under the Edit drop-down menu in the dashboard toolbar.

1. Navigate to the dashboard that you want to export. All existing dashboards appear in the drop-down **Dashboard Selection List** in the currently displayed dashboard.
2. In the dashboard toolbar, click  (Export Dashboard).
   The exported file is saved in .zip format.

Note: The Export feature is not applicable for preconfigured dashboards.

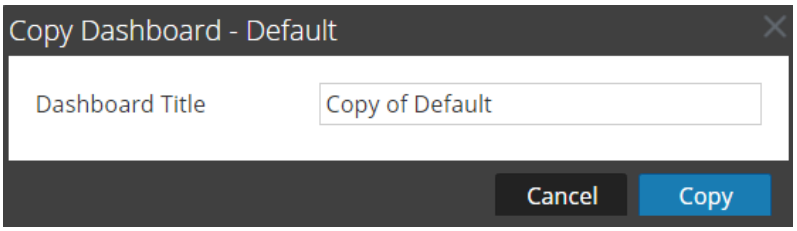
Copying a Dashboard

To customize the views in NetWitness, you can copy dashboards to the NetWitness dashboard or a custom dashboard. The NetWitness dashboard, as the name suggests, offers all NetWitness dashlets. The Copy Dashboard dialog creates a duplicate dashboard, which can be customized. When you copy a dashboard, the default name is prefixed with `Copy of`. For example, if the name of the original dashboard is `XYZ`, the default title of the copied dashboard will be `Copy of XYZ`.

To copy a dashboard:

1. Navigate to any dashboard.
2. In the dashboard toolbar, click .
   The Copy Dashboard dialog is displayed. The following screenshot is an example of copying a

dashboard.



3. Enter the Dashboard Title.
4. Click **Copy**.

Sharing a Dashboard

In NetWitness, as an administrator you can share dashboards for viewing purposes with other roles such as Administrators, Analysts, Operators and so on. When you share a dashlet, the users can only view the dashboard, make dashboard as favorite, copy the dashboard, and export the dashboard. In case of other roles such as Analysts, Operators, and so on, you can share the dashboard only with similar roles. For example, an analyst can share a dashboard with other analysts only.

1. Navigate to any dashboard.
2. In the dashboard toolbar, click 

and select the checkbox of the role with whom you want to share the dashboard.

Note: If you do not want to share the dashboard, clear the checkbox of the role. If a dashboard is shared, the dashboard name will be displayed with the share icon .

Removing Unwanted Dashboards

You can remove the following dashboards by enabling the dashboard cleaning job.

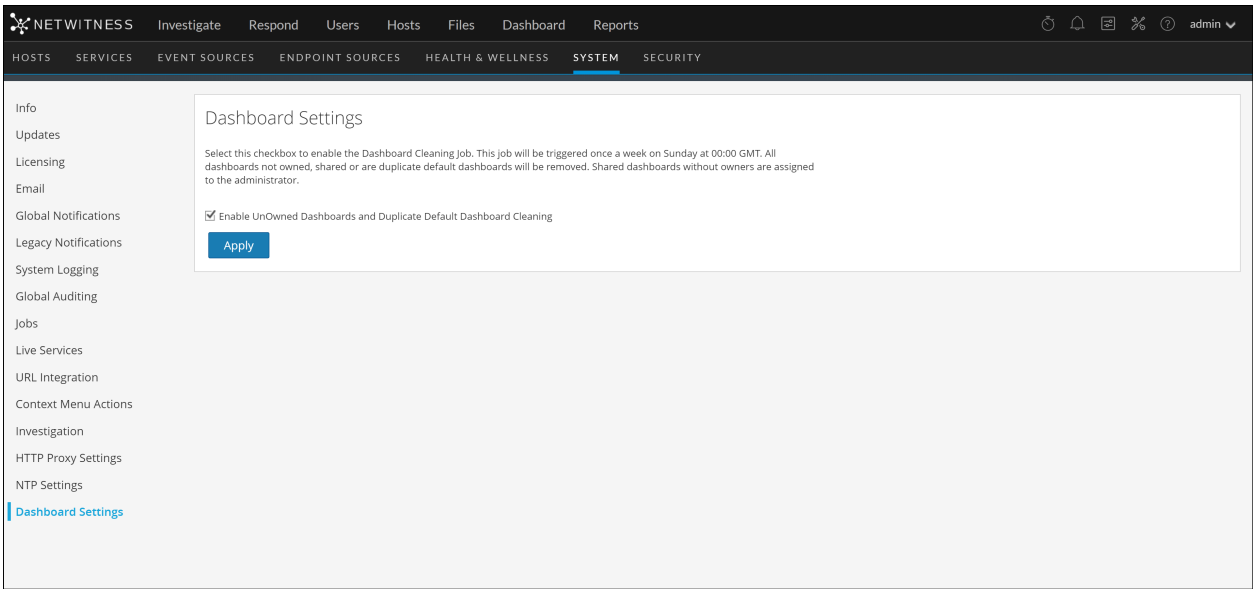
- Dashboards not owned by anyone
- Dashboards not shared
- Duplicates default dashboards

When the dashboard cleaning job is enabled, a scheduled job is triggered once a week on Sunday at 00:00.

Shared dashboards without owners are assigned to the administrator.

To enable Dashboard Cleaning Job:

1. Go to **(Admin) > System > Dashboard Settings**.
2. Select the **Enable UnOwned Dashboards and Duplicate Default Dashboard Cleaning** checkbox.
3. Click **Apply**.

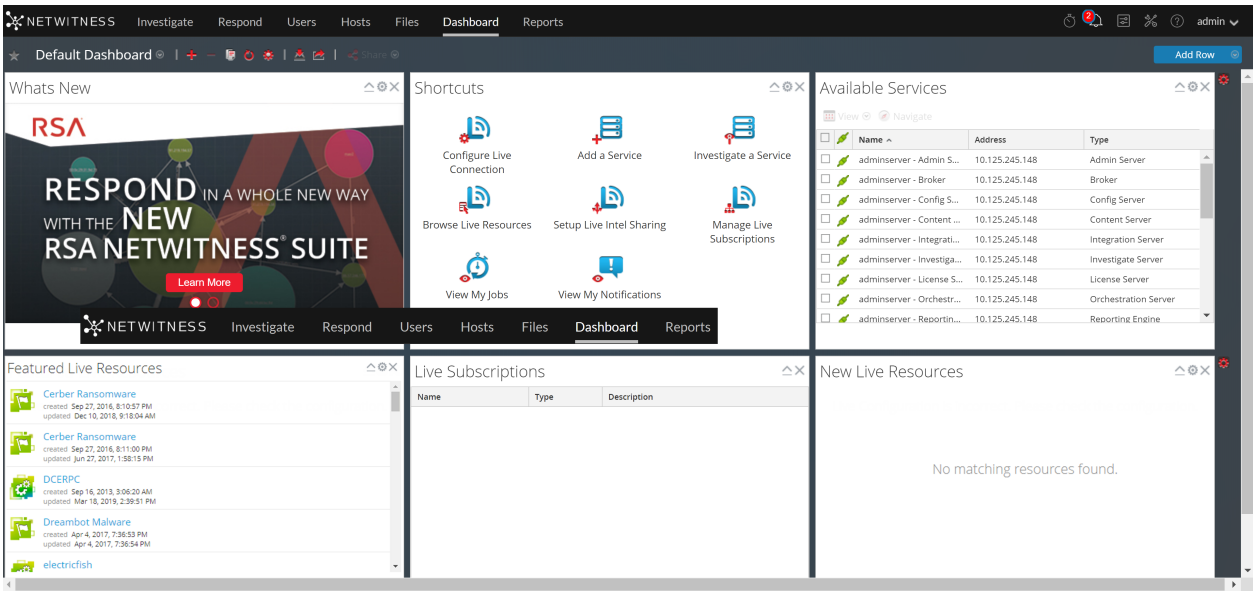


Using Dashboards in the Analyst User Interface

Large environments that include geographical distribution with a single data center and multiple NetWitness Servers require Analyst UI instances in all their NetWitness locations or managed entities. For example, if an Analyst UI is deployed for the EMEA SOC team, analysts can query their EMEA NetWitness Platform hosts directly. If the EMEA team has Broker hosts and Concentrator hosts within the region, the Analyst UI can connect and query them instead of connecting back to Primary User Interface (Primary UI).

When using the dashboards in the Analyst UI, these features apply and affect the dashboards available to analysts:

- Each Analyst UI instance is connected to its local Reporting Engine service. Every Reporting Engine has its own copy of the built-in content such as rules and charts and the local Reporting Engine runs them on the default data source (if configured in Reporting Engine) when the dashboard is enabled.
- The built-in Reporting Engine content that is modified on local Analyst UI is available only to that specific Analyst UI. This behavior is same for an Admin UI as well.
- Dashboards that are shared across different Analyst deployments display the Reporting Engine chart data from the shared instance. To edit the shared dashboard, the analyst must create a copy or contact the Admin to customize it.
- When built-in preconfigured dashboards are enabled from an Admin UI, the dashboard is enabled for all the roles and for all the Analyst UIs. However, the data displayed in every Analyst UI is specific to the associated Reporting Engine.
- By default, all the built-in preconfigured dashlets are disabled and can be enabled only from an Admin UI. Before you enable a preconfigured dashlet, you must set up the Live Services Account, see the *Set Up Live Services* topic in the Live Services Management Guide.



Note: The Analyst UI functions are similar to the Admin UI functions, except all configurations must be performed from the Admin UI.

Setting User Preferences

You can view and manage your NetWitness global application preferences from your user profile. There are two global user preference dialogs that have different options. The user Preferences dialog is accessible from the Springboard, Investigate (Events), Respond, Users, Hosts, Files and Configure views.

The Preferences dialog is accessible from the Investigate, Dashboard, Reports,  (Configure) , and  (Admin) views. The dialog that you see depends on where you access the user preferences.

You can:

- Change the application language
- Set the application time zone
- Set the application date and time format*
- Select your default NetWitness starting location*
- Select your default Investigate view*
- Choose a dark or light theme for the application*
- Change your password (See Changing Your Password for more information.)
- Enable or disable notifications**
- Enable or disable context menus**

* You can make this change from the **User Preferences** dialog accessible from these views:

Springboard, Investigate > Events (formerly Event Analysis), Respond, Users, Hosts, Files, and  (Configure) > (Capture Policies, Incident Rules, Incident Notifications, and Log Parser Rules). See User Preferences.

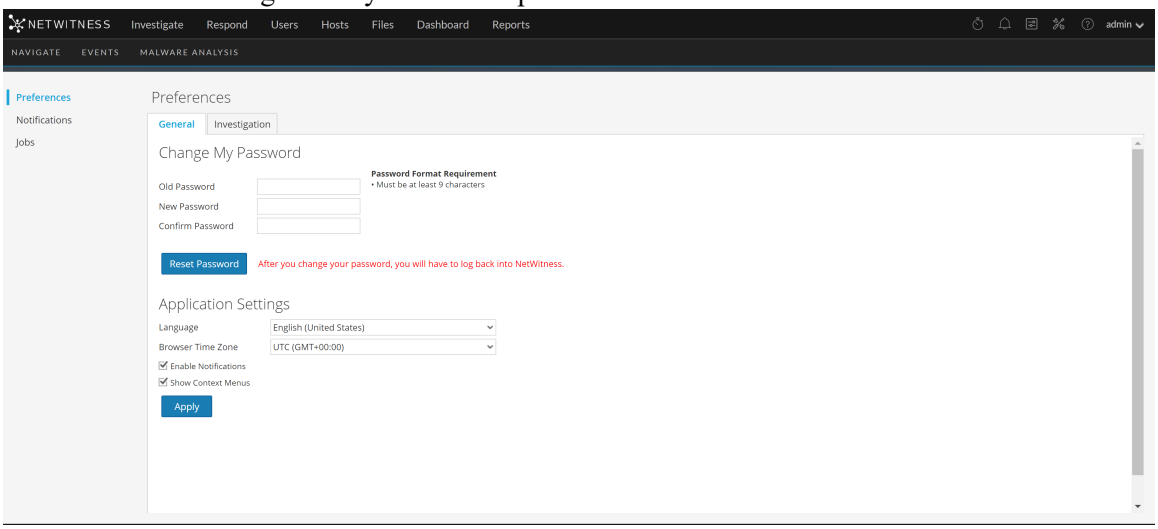
** You can make this change from the **Preferences** dialog accessible from these views: Investigate, Dashboard, Reports,  (Configure) > (Live Content, Subscriptions, ESA Rules, and Custom Feeds), and  (Admin) . See Preferences.

Preferences

This section gives instructions for various tasks that can be performed in the Preferences dialog that is accessible in Investigate, Dashboard, Reports,  (Configure) , and  (Admin) .

View your Preferences

In the upper right corner of the NetWitness browser window, select your username, for example admin, and then select **Profile**.
The Preferences dialog shows your current preferences.



Set the Language and Time Zone

Note: The Language preference option applies to NetWitness 11.2 and later.

You can change your preferred language for the entire NetWitness Platform. The default language is English (United States).

1. In the Preferences dialog, select your localization preferences:
   a. **Language**: Select your preferred language for NetWitness.
   b. **Browser Time Zone**: Set the time zone to use in the NetWitness.
2. Click **Apply**.
   Your preferences become effective immediately.

Note: When Daylight Saving Time (DST) starts or ends, if the selected time zone for the currently logged in user observes DST, the user interface automatically updates to reflect the correct time.

Enable or Disable System Notifications for Your User Account

By default, NetWitness system notifications are enabled when a new user account is created. You can disable and enable these notifications at any time.

1. In the Preferences dialog Application Settings section:
   - To enable notifications for your user account, select the **Enable Notifications** checkbox.
   - To disable notifications, clear the **Enable Notifications** checkbox.

2. Click **Apply**.
   Your preference becomes effective immediately.

Enable or Disable Context Menus for Your User Account

By default, context menus are enabled when a new user account is created. Context menus provide additional functions for specific views when you right-click a view.

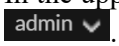
1. In the Preferences dialog Application Settings section:
   - To enable context menus for your user account, select the **Show Context Menus** checkbox.
   - To disable context menus, clear the **Show Context Menus** checkbox.
2. Click **Apply**.
   Your preference becomes effective immediately.

Note: Settings available on the Investigate tab in the Preferences dialog are documented in the *NetWitness Investigate User Guide*.

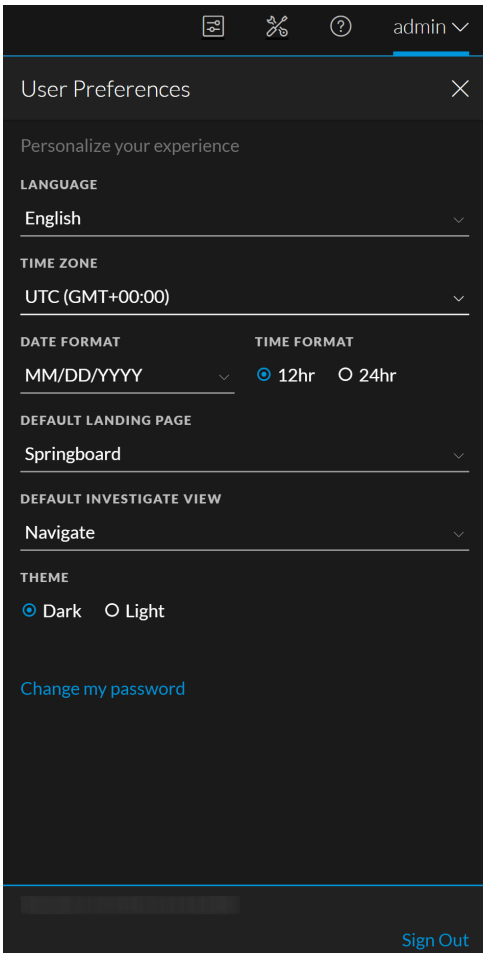
User Preferences

This section gives instructions for various tasks that can be performed in the User Preferences dialog that is accessible in the Springboard, Events, Respond, Users, Hosts, Files and  (Configure) views.

View Your User Preferences

In the upper right corner of the NetWitness browser window, select your username, for example, 

.
The User Preferences dialog shows your current preferences when accessed through the Springboard, Events, Respond, Users, Hosts, Files and  (Configure) views.



Any selections that you make become effective immediately.

Set the Language, Time Zone, and Date and Time Format

Note: The Language preference option applies to NetWitness 11.2 and later.

You can change your preferred language for the entire NetWitness Platform UI. The default language is English (United States). You can also change the time zone and the format of the date and time for your location.

1. Open the User Preferences dialog.
2. In the User Preferences dialog, select your localization preferences:
   a. **Language**: Select your preferred language for NetWitness.
   b. **Time Zone**: Set the time zone to use in the NetWitness.
   c. **Date Format**: Set the format for the order of the display of the month (MM), day (DD), and year (YYYY). For example, the MM/DD/YYYY format shows the date as 05/11/2019.
   d. **Time Format**: Set the time as 12-hour or 24-hour time. For example, 2:00 PM in 12-hour time is 14:00 in 24-hour time.

Changes in the Investigate and Respond views become effective immediately.

Note: When Daylight Saving Time (DST) starts or ends, if the selected time zone for the currently logged in user observes DST, the user interface automatically updates to reflect the correct time.

Select the Default NetWitness Platform Starting Location

1. Open the User Preferences dialog.
2. In the **Default Landing Page** field, select the opening view that you would like to see when you log in to NetWitness. You can choose Springboard, Investigate, Respond, Users, Hosts, Files, Dashboard, Reports, Configure, and Admin, and according to your user role. For example, you can choose Respond to go directly to the relevant section of the application for Incident Responders. For more information, see Setting Up Your Default View by SOC Role to help you select the appropriate default view.
   This selection sets the default view for the entire application. The changes take effect immediately.

Select the Default Investigate View

1. Open the User Preferences dialog.
2. In the **Default Investigate View** field, select the default landing page when you log in to NetWitness and navigate to Investigate. You can choose Navigate, Legacy Events (if enabled), Events (formerly Event Analysis), or Malware Analysis as the default Investigate view. For example, you can choose Events for the default Investigate view to go directly to the Events view to see the events generated for a service. See Setting Up Your Default View by SOC Role to help you select the appropriate default view. For more information, see the *NetWitness Investigate User Guide*.

Note: After you have applied the change in the drop-down, sometimes it takes few seconds for the changes to take effect.

Choose the Appearance of NetWitness Platform

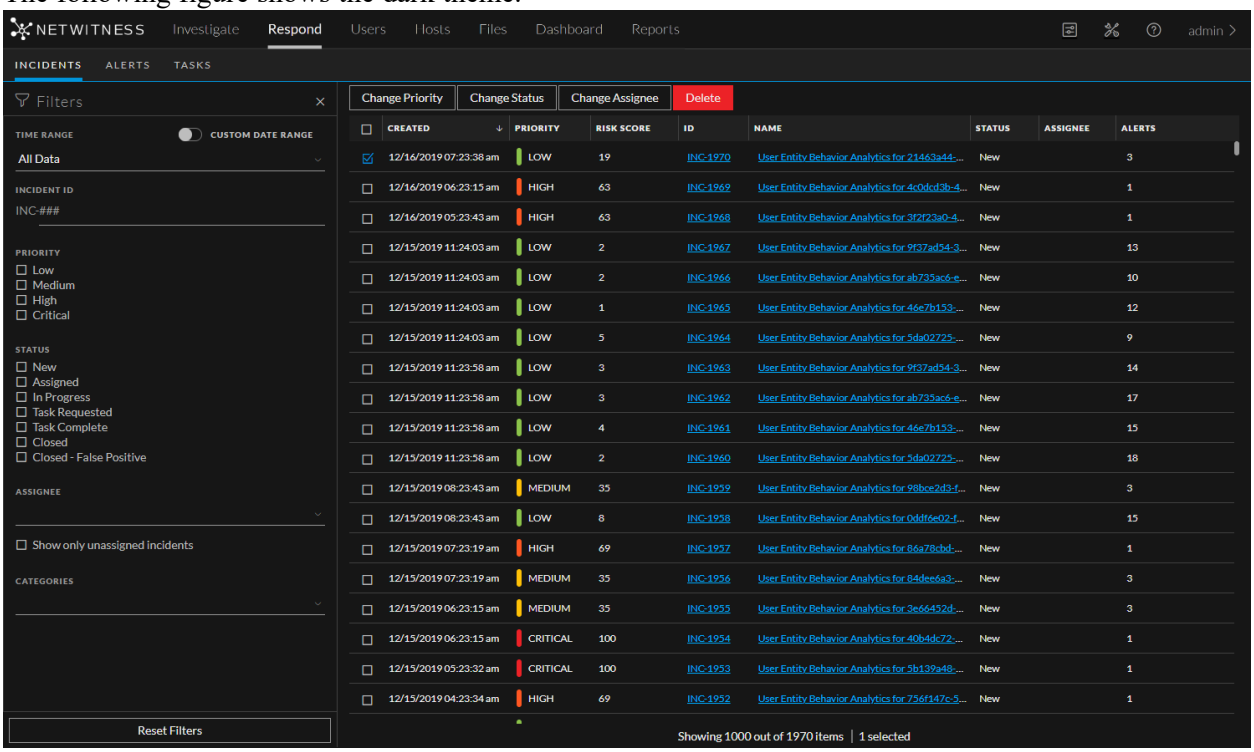
Note: This option is only available for NetWitness versions 11.1 and later.

You can choose a dark theme or a light theme for your application, depending on your personal preference. When you change the theme, the Respond view and some Investigate views change to the light or dark theme. Your selection only changes how NetWitness appears to you, not other users.

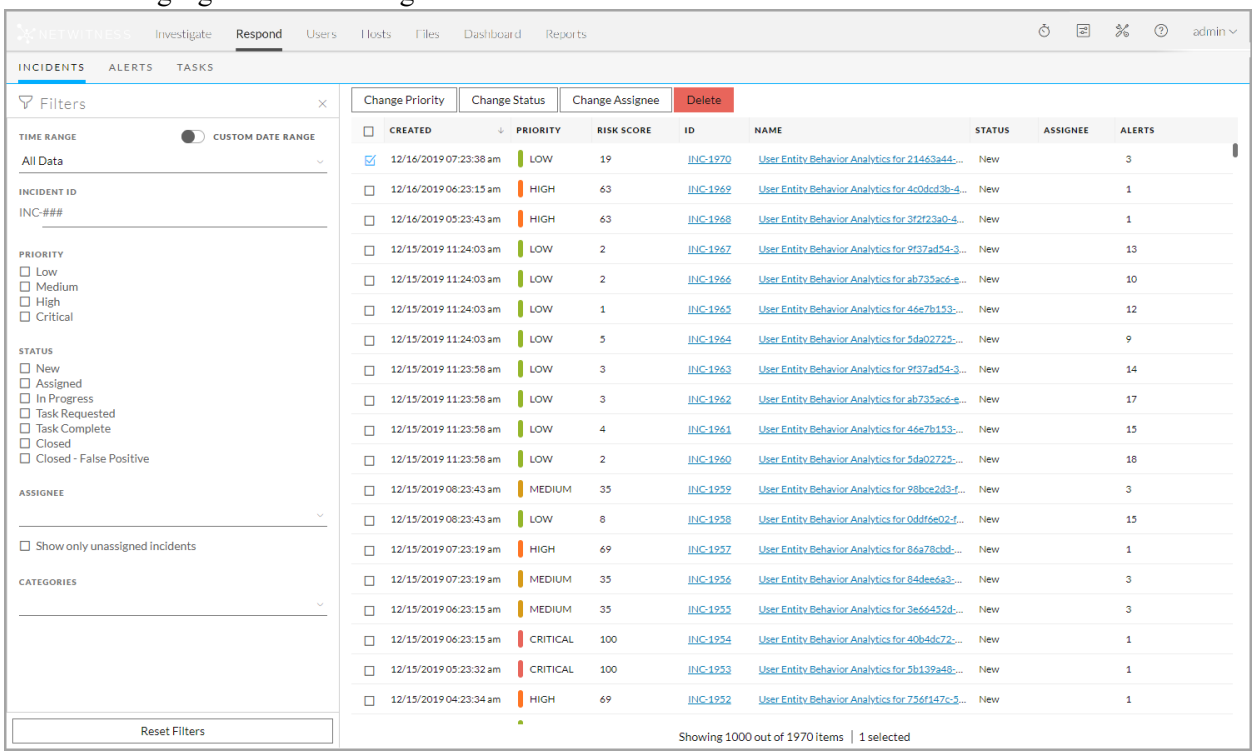
1. Open the User Preferences dialog.
2. Under **THEME**, select one of the following options:
   - **Dark**: The dark theme is best for darker environments or when you do not need as much contrast.
   - **Light**: The light theme is best for lighter environments, when you need more contrast, or when you are projecting the application for others to view. Since some views are not affected by the theme changes, you may want to choose the light theme for a more cohesive viewing experience.

Changes become effective immediately.

The following figure shows the dark theme.



The following figure shows the light theme.



Managing Jobs

Inevitably, there are on-demand or scheduled tasks in NetWitness that take a few minutes to be completed. The NetWitness jobs system lets you begin a long-running task and continue using other parts of NetWitness while the job is running. Not only can you monitor the progress of the task, but you can also receive notifications when the task has completed and whether the result was a success or failure.

While you are working in NetWitness, you can open a quick view of your jobs from the toolbar. You can look anytime, but when a job status has changed, the Jobs icon ( ) is flagged with the number of running jobs. Once all jobs are completed, that number disappears.

You can also see the jobs in these two views:

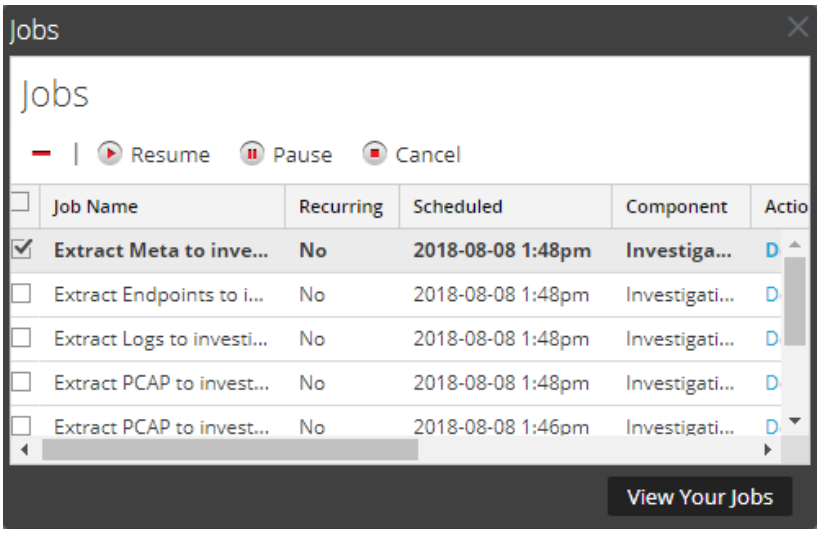
- In the user Profile Jobs panel, you see the same jobs in a full panel. These are only your jobs.
- In the System view, users with administrative privileges can view and manage all jobs for all users in a single jobs panel.

The structure of the jobs panel is the same in all views.

Display the Jobs Tray

In the NetWitness toolbar, click the Jobs icon ( ).

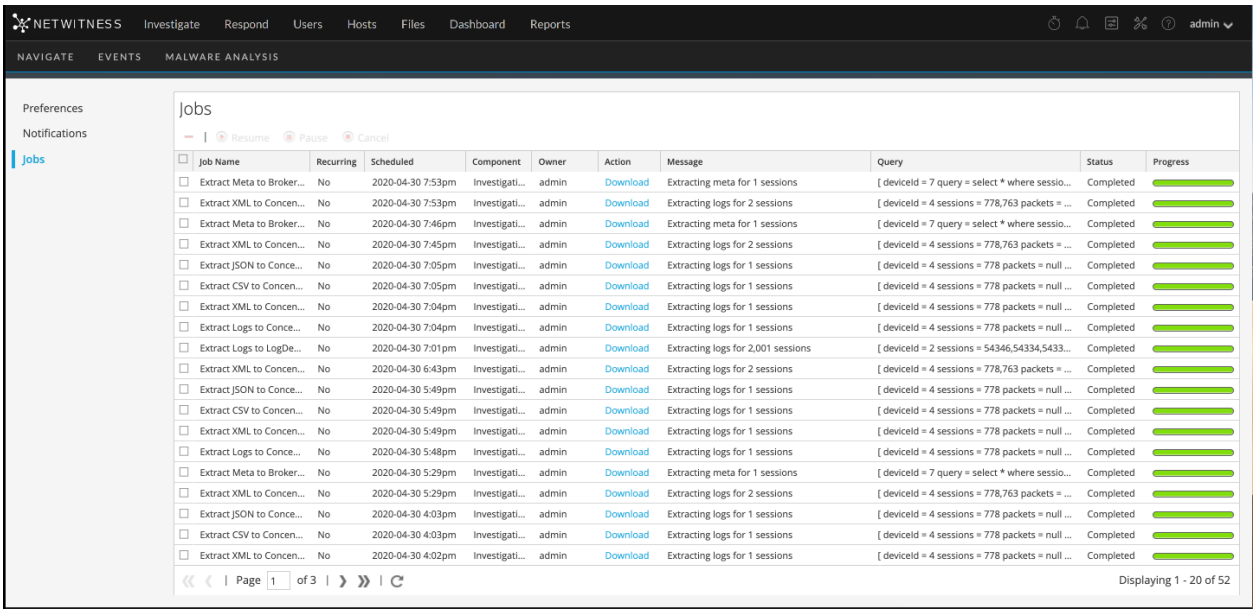
The Jobs Tray is displayed.



The Jobs Tray lists all recurring and non-recurring jobs that you own, using a subset of the columns available in the Jobs panel. Otherwise the Jobs Tray and the user Profile view Jobs panel are the same. In the Admin System view, the Jobs panel lists information about all NetWitness jobs for all users.

View All of Your Jobs

To see a complete view of your jobs, in the Jobs Tray, click **View Your Jobs**.
The Jobs panel is displayed. The Query column is available in Version 11.5 and later.



Pause and Resume Scheduled Execution of a Recurring Job

The Pause and Resume options apply only to recurring jobs. You can pause a recurring job that is running; however, it has no effect on that execution. The next execution (assuming the job is still paused) is skipped.

1. To stop the next execution of a recurring job, in any **Jobs panel**, select the job, and click **Pause**.
   The next execution of the job is skipped, and the schedule is paused until you click Resume.
2. To restart execution of paused recurring jobs, select the job and click **Resume**.
   The next execution of the job occurs as scheduled, and the schedule for the job resumes.

Cancel a Job

To cancel jobs that are executing or in the queue to execute:

1. In the **Jobs Tray** or either **Jobs** panel, select one or more jobs.
2. Click **Cancel**.
   A confirmation dialog is displayed.
3. Click **Yes**.
   The jobs are canceled, and the entries remain in the list with a status of **canceled**.

   If you cancel a recurring job, it cancels that execution of the job. The next time the job is scheduled to run, it executes normally.

Delete a Job

Caution: When you delete a job, the job is instantly deleted from the list. No confirmation dialog is offered. If you delete a recurring job, all future executions are removed as well.

Users can delete their own jobs before, during, or after execution. Administrators can delete any job. To delete jobs:

1. Select one or more jobs.
2. Click −.
   The jobs are deleted from the list.

Download a Job

When a job has the Download status in the Action column, you can download the result of the job. If you are working in the Investigate view and extract the packet data for a session as a PCAP file or extract the payload files (for example, Word documents and images) from a session, a file is created. To download the file to your local system, click **Download**.

Viewing and Deleting Notifications

While you are working in NetWitness, you can view recent system notifications without leaving the area where you are working. You can open a quick view of notifications from the NetWitness toolbar. You can look anytime, but when a new notification is received, the Notifications icon is flagged ().
Examples of notifications include:

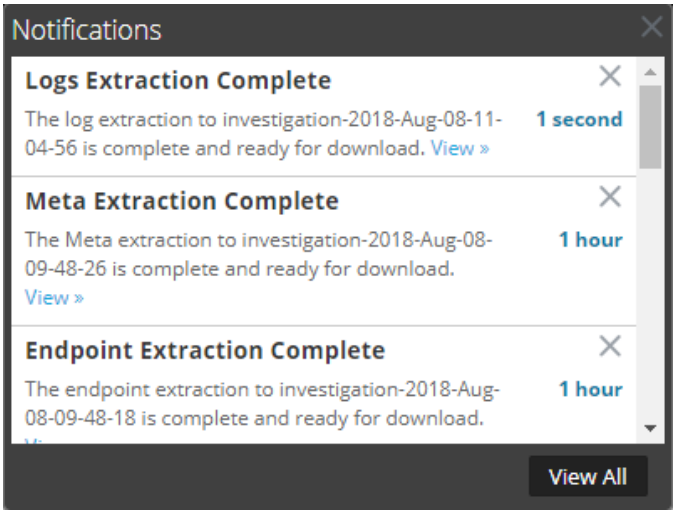
- A host upgrade completed.
- A parser push to decoders completed.
- A newer software version is available.

You can see notifications in these two views:

- In the Notifications tray, you can see your recent notifications.
- In the user Profile Notifications panel, you can view all of your notifications.

View Recent Notifications

To display recent notifications, click the Notifications icon ().
The Notifications tray is displayed.

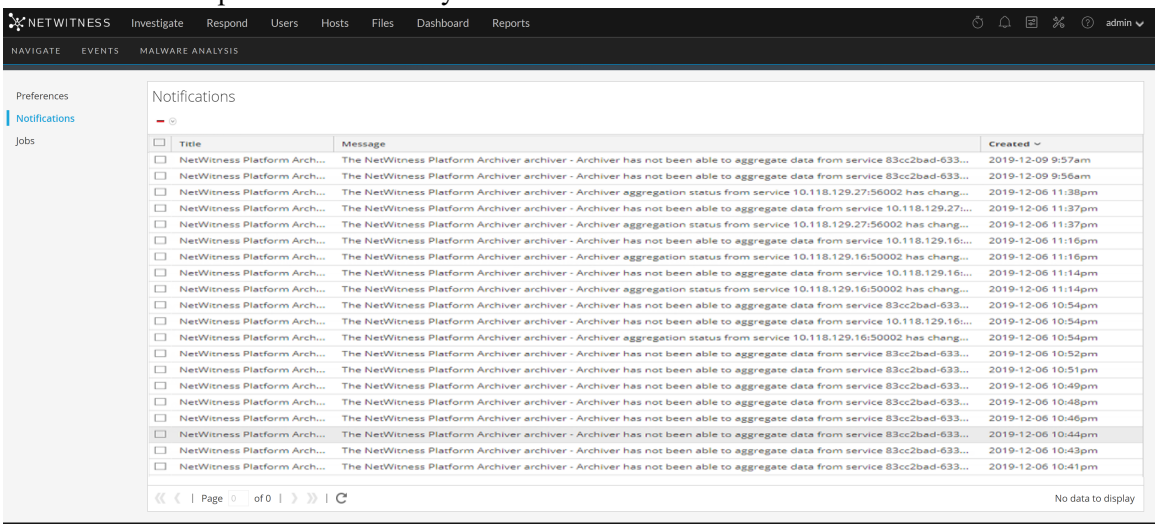


View All Your Notifications

To view all of your notifications, do one of the following:

- Click  to open the Notifications tray and then click **View All** in the Notifications tray.
- In the upper right corner of the NetWitness browser window, select your username and then select **Profile**. In the options panel of the Preferences dialog, select **Notifications**.

The Notifications panel shows all of your notifications.



Delete Notification Records

To delete notification records:

1. In the **Profile Notifications** list, select the notifications that you want to delete.
2. Click −.

   The selected notifications are deleted from this list and from the Notifications tray.

Viewing Help in the Application

There are different ways available to get help while using NetWitness. You can use inline help, tooltips, and online help links.

View Inline Help

Inline help provides additional information about what to do in sections or fields that you are currently viewing in the NetWitness user interface. To display inline help, hover over . The inline help shows a brief description of the element.

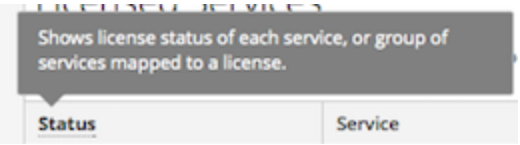
Inline help example:



View Tooltips

Tooltips are a quick way for you to see a description of the text or additional information about an action, field, or parameter. Tooltips appear as underlined text. To display the tooltip and see a brief description of the term, hover over the underlined text.

Tooltip example:

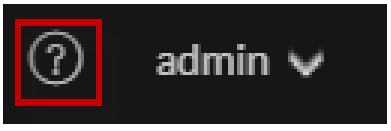


View Online Help

Online help links take you outside of NetWitness to the NetWitness Community online documentation. This site has a complete documentation set for NetWitness, and the links take you directly to the topic that describes the part of the user interface currently in view.

To view the online help topic for the current location, click in the NetWitness toobar or in a dialog. The relevant help topic is displayed in a separate browser window. The topic describes the features and functions of the current view or dialog. From that topic, you can quickly navigate to the related procedures.

The following figure is an example of the online help icon in the NetWitness toobar.



Finding Documents on NetWitness Community

The NetWitness documentation is located on NetWitness Community. NetWitness Community brings all of your NetWitness resources together in one place. It includes advisories, product documentation, knowledge base articles, downloads, and training. To view a *Guided Tour of NetWitness Community*, see https://community.netwitness.com/t5/netwitness-platform/ct-p/netwitness.

Locate NetWitness Documentation

NetWitness Logs, Networks, Endpoints, and UEBA documentation is at the following link: https://community.netwitness.com/t5/netwitness-platform/ct-p/netwitness-documentation

To navigate to NetWitness Logs, Networks, Endpoints, and UEBA documentation:

1. On the NetWitness Community homepage (https://community.netwitness.com), under **Products**, click **NetWitness Platform**.
2. On the NetWitness Platform page, click **Documentation**.

Locate NetWitness Content

NetWitness Content contains feeds, parsers, reports, and rules. It is located at the following link: https://community.netwitness.com/t5/threat-intelligence/ct-p/threat-intelligence.

To navigate to NetWitness Content:

1. On the NetWitness Community homepage (https://community.netwitness.com), under **Products**, click **NetWitness Platform**.
2. On the NetWitness Platform page, click **Documentation** > **Threat Intelligence**.

Locate NetWitness Supported Event Sources

The NetWitness Integrations Catalog is located at the following link: https://community.netwitness.com/t5/netwitness-platform-integrations/tkb-p/netwitness-integrations.

To navigate to the NetWitness Integrations Catalog:

1. On the NetWitness Community homepage (https://community.netwitness.com), under **Products**, click **NetWitness Platform**.
2. On the NetWitness Platform page, click **Integrations**

Locate Hardware Setup Guides

The Hardware Setup Guides are at the following link: https://community.netwitness.com/t5/netwitness-platform-hardware/tkb-p/netwitness-hardware-documentation.

1. On the NetWitness Community homepage (https://community.netwitness.com), under **Products**, click **NetWitness Platform**.
2. On the NetWitness Platform page, click **Documentation** > **Hardware Setup Guides**.

Follow Content for Updates

You can follow pages or documents to be notified of changes.

1. Log in to NetWitness Community.
2. Navigate to a page or a document and next to the Search bar, click  > **Subscribe**.

Send Your Feedback to NetWitness

Your feedback is very important to us and helps us to provide a better experience for our customers. Please send your suggestions to nwdocsfeedback@netwitness.com

Troubleshooting the User Interface

This section describes common issues that users may face during setup and provides basic troubleshooting information.

Basic Troubleshooting Tips for User Setup

The following table provides basic troubleshooting tips that may be helpful for user setup in NetWitness.

Problem	Troubleshooting Tip
When I log in to NetWitness, I see the wrong default view.	Verify that the correct default view is set in the Default Landing Page field in your user preferences. If you select the Dashboard view, you can select the predefined dashboard that is most appropriate for your SOC role. You can also import or create your own dashboard.
I see the correct view, but the metadata does not load.	Make sure that you are using the latest version of the browser. If that does not work, try using another browser. For example, if you are using Safari, try using Firefox or Chrome.
I am using Internet Explorer 10 and I get the following error: `The page can't be displayed.`	NetWitness supports current versions of Firefox, Chrome, and Safari. Internet Explorer is no longer supported.
When I log in, I cannot see anything.	See your administrator, you may need a user role assigned to your account or additional troubleshooting.
I can't see where to change my default landing page.	Go to the User Preferences in the Respond view or Events view (formerly Event Analysis view) or see your administrator.

Analyst User Interface Dashlet Issue

Problem	Cause	Solution
I see the following error message on my dashboard: `The underlying chart for this dashlet is unavailable.`	This scenario occurs on an Analyst UI deployment when an Analyst creates a dashboard on one NetWitness Server, logs in to another NetWitness Server, and tries to view a dashboard that the Analyst did not share.	Share the dashboard to view the data. For more information, see Sharing a Dashboard.

Springboard Issue

Problem	Cause	Solution
I see the following error message on a Springboard panel: `Filter definition is missing! Edit the panel to remove the filter or update with a new filter.`	A query profile that is used to filter events in the panel may have been deleted in the Investigate view. Springboard panels that use a deleted query profile as a filter do not work.	Edit the Springboard panel to remove the filter or use a different filter. For more information, see Managing the Springboard.

Springboard Fails to Load the Panel Issue

Problem	Cause	Solution

I see the following error messages on a Springboard panel: `Request timed out! Contact your administrator or refer to the Getting Started Guide on NetWitness Community for details.`
 	
`Failed to fetch the results. Refresh the panel. If the issue persists, contact your administrator or refer to the Getting Started Guide on NetWitness Community for details.`
	• The Springboard panels request time out in 60 seconds.
	• The associated data sources are offline or not reachable.
	• Allocated memory is insufficient to execute the query.
		• Limit the results by narrowing the time range for the Springboard using the time range field above the panels.
		• Go to the Investigate > Events view and refine the Query Profile used in the panel as described in "Use Query Profile to Encapsulate Common Areas of Investigation" in the *Investigate User Guide*.
		• Make sure that the associated data sources are online.
		• Increase the memory by modifying the `max.query.memory` parameter setting.
		• If the above solution does not work and if the issue still persists, check the service logs in the given order:
		  ▪ Admin server
		  ▪ Investigate server
		  ▪ Respective core services
		
		For more information, see “Display System and Service Logs” section in the *System Maintenance Guide*.

Inconsistent Event Panel Count Issue

Problem	Cause	Solution
Event count on the panel is inconsistent.	The underlying data sources are offline or not reachable.	Make sure that all underlying data sources are online and refresh the panel. For more information, see Managing the Springboard.

NetWitness Platform Getting Started References

The following section contains user interface reference information related to getting started with the NetWitness application.

- User Preferences
- Notifications Panel and Notifications Tray
- Jobs Panel and Jobs Tray

User Preferences

To adjust NetWitness to best fit your environment and work practices, you can set your own global application preferences. You can:

- Change the application language
- Set the application time zone
- Set the date and time formats
- Select your default NetWitness starting location
- Select your default Investigate view
- Choose a dark or light theme for the application
- Change your password
- Enable notifications
- Enable context menus
- Change Investigate preferences - Described in the *NetWitness Investigate User Guide*.

Your global preference options vary depending on whether you access them, such as Springboard, Investigate, Respond, Users, Hosts, Files, Dashboard, Reports, Configure, and Admin. There are two global user preferences dialogs accessible from the main menu bar:

- **User Preferences** dialog - Accessible from these views: Springboard, Investigate > Events (formerly Event Analysis), Respond, Users, Hosts, Files, and  (Configure) > (Capture Policies, Incident Rules, Incident Notifications, and Log Parser Rules).
- **Preferences** dialog - Accessible from these views: Investigate, Dashboard, Reports,  (Configure) > (Live Content, Subscriptions, ESA Rules, and Custom Feeds), and  (Admin) .

What do you want to do?

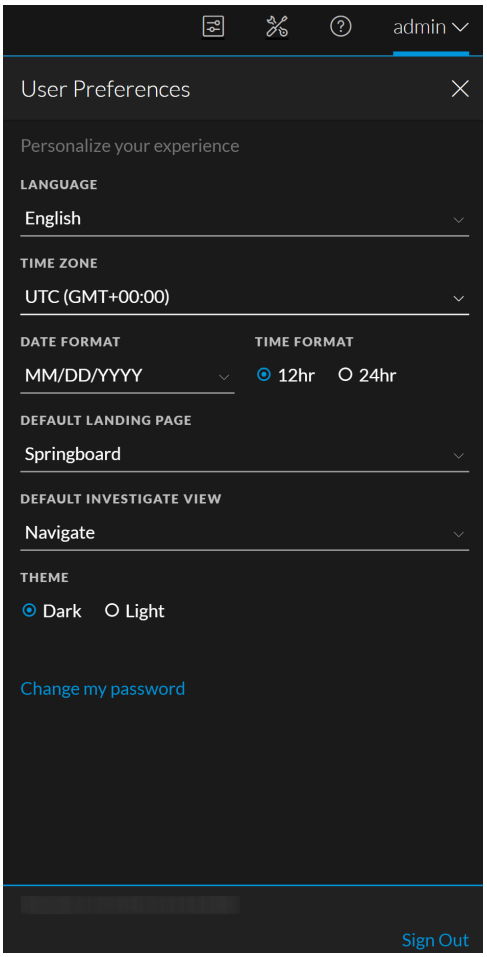
Role	I want to ...	Show me how
All	Change my Password	Change My Password Section
All	Choose my Default Landing Page	Setting Up Your Default View by SOC Role
All	Set my User Preferences	Setting User Preferences

Related Topics

- NetWitness Platform Basic Navigation

User Preferences

To access your user preferences, click your username, for example, admin.
The User Preferences dialog shows your current preferences and the NetWitness version.



The following table describes the global application preference options that you can access from the User Preferences dialog.

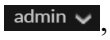
Option	Description
Language	(This option applies to NetWitness 11.2 and later.) Sets the preferred language for the entire NetWitness Platform. The default language is English (United States).
Time Zone	Sets the time zone to use in NetWitness.
Date Format	Sets the format for the order of the display of the month (MM), day (DD), and year (YYYY). For example, the MM/DD/YYYY format shows the date as 05/11/2017.
Time Format	Sets the time as 12-hour or 24-hour time. For example, 2:00 PM in 12-hour time is 14:00 in 24-hour time.

Option	Description
Default Landing Page	Enables you to select the default view when you log in to NetWitness. You can choose Springboard, Investigate, Respond, Users, Hosts, Files, Dashboard, Reports,  (Configure) , and  (Admin) , according to your user role. For example, you can choose Respond to go directly to the relevant section of the application for Incident Responders. This selection sets the default view for the entire application.
Default Investigate View	(This option applies to NetWitness 11.1 and later.) Select the default landing page for the Investigate view. You can choose Navigate, Legacy Events (if enabled), Events (formerly Event Analysis), or Malware Analysis as the default Investigate view. For example, you can choose Events for the default Investigate view to go directly to the Events page to view the events generated for a service.
Theme	(This option applies to NetWitness 11.1 and later.) Changes the appearance of the Respond view and some Investigate views that you see in the application. You can choose between light and dark themes: • **Dark**: The dark theme is best for darker environments or when you do not need as much contrast. • **Light**: The light theme is best for lighter environments, when you need more contrast, or when you are projecting the application for others to view. Since some views are not affected by the theme changes, you may want to choose the light theme for a more cohesive viewing experience. Your selection only changes how NetWitness appears to you, not other users.
Change my password	Opens the Preferences dialog where you can change your password.
Version	Shows the NetWitness version.
Sign Out	Enables you to log out of NetWitness.

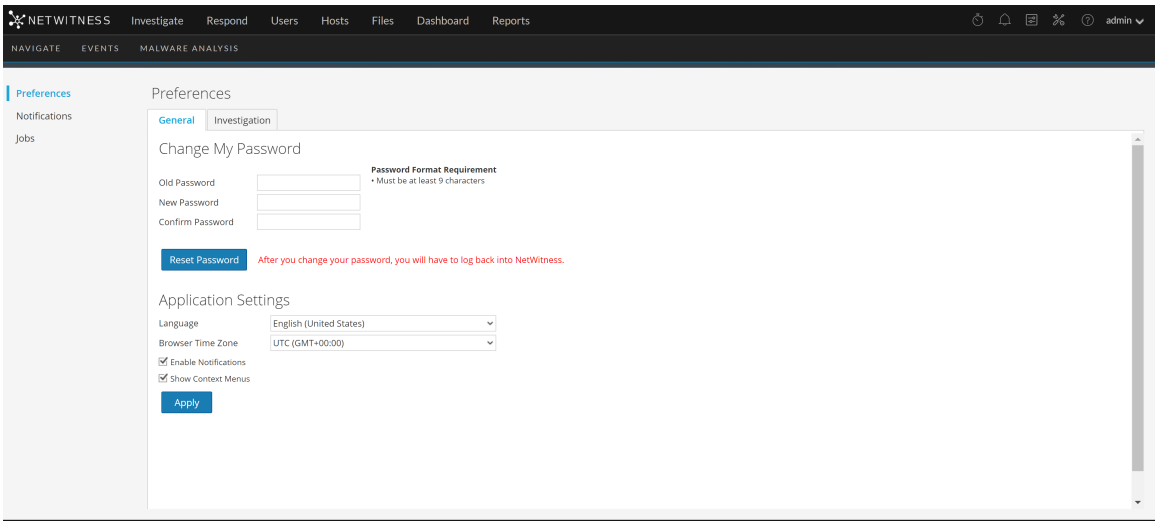
Any selections that you make become effective immediately.

Preferences

To access additional global user preferences, do one of the following:

- For most views, such as Investigate, Dashboard, Reports,  (Configure) , or  (Admin) , select your username and then select **Profile**.
- In the Springboard, Investigate view [Events (formerly Event Analysis)], Respond, Users, Hosts, Files, and some  (Configure) views, select your username, for example , and in the User Preferences dialog click **Change my password**.

The Preferences dialog shows your current preferences.



The following tables describe the global application preference options that you can access from the Preferences dialog.

Change My Password Section

This section enables you to change your password. Your administrator defines the appropriate password strength requirements for your NetWitness password, such as minimum password length and minimum number of uppercase, lowercase, decimal, non-Latin alphabetic, and special characters. These requirements are then displayed when changing your password.

The following tables describes the options in the Change My Password section.

Option	Description
Old Password	Enter the password that you used to log in to NetWitness.
New Password	Enter the password that you want to use for the next login.
Confirm Password	Retype the new password.
Reset Password	Updates your user profile with the new password. You will be logged out of NetWitness for the changes to take effect. The new password becomes effective the next time you log in to NetWitness. The password change is applied to your system login and to all NetWitness services on which your account has been added.

If you changed your password, you will be logged out of NetWitness for the changes to take effect. The new password becomes effective the next time you log in to NetWitness.

Application Settings Section

The following tables describes the options in the Application Settings section.

Option	Description
Language	(This option applies to NetWitness 11.2 and later.) Sets the preferred language for the entire NetWitness Platform. The default language is English (United States).
Browser Time Zone	Sets the time zone to use in NetWitness. Your time zone preference is displayed on the toolbar.
Enable Notifications	This checkbox enables and disables notifications for your user account. By default, NetWitness system notifications are enabled when a new user account is created.
Enable Context Menus	This checkbox enables and disables context menus for your user account. By default, context menus are enabled when a new user account is created. Context menus provide additional functions for specific views when you right-click in a view.
Apply	Updates your preferences and applies the changes immediately.

Notifications Panel and Notifications Tray

NetWitness provides system notifications to advise users about certain actions or conditions:

- A host upgrade completed.
- A parser push to decoders completed.
- A service went down (critical log of a certain type).
- A visualization completed.
- A report completed.
- A newer software version is available.

While you are working in NetWitness, you can view recent system notifications without leaving the area where you are working. You can open a quick view of notifications from the NetWitness toolbar. You can look anytime, but when a new notification is received, the Notifications icon is flagged ().

When you are viewing notifications in the Notifications tray, only recent notifications are displayed. You can access all of your notifications from your user Profile and from the Notifications tray by selecting the View All option. Procedures for viewing notifications are provided in Viewing and Deleting Notifications.

Note: In the Analyst UI, the license notifications are not displayed in the notification tray or login window when the license goes out of compliance or when the license expires. This is displayed only on the Admin UI.

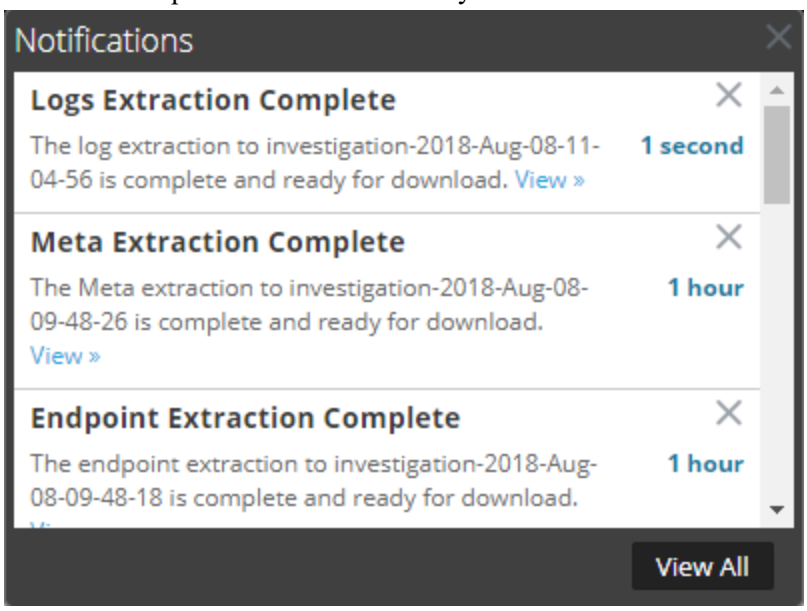
What do you want to do?

Role	I want to ...	Show me how
All	View all notifications	Viewing and Deleting Notifications
All	Delete notifications	Viewing and Deleting Notifications

Quick Look

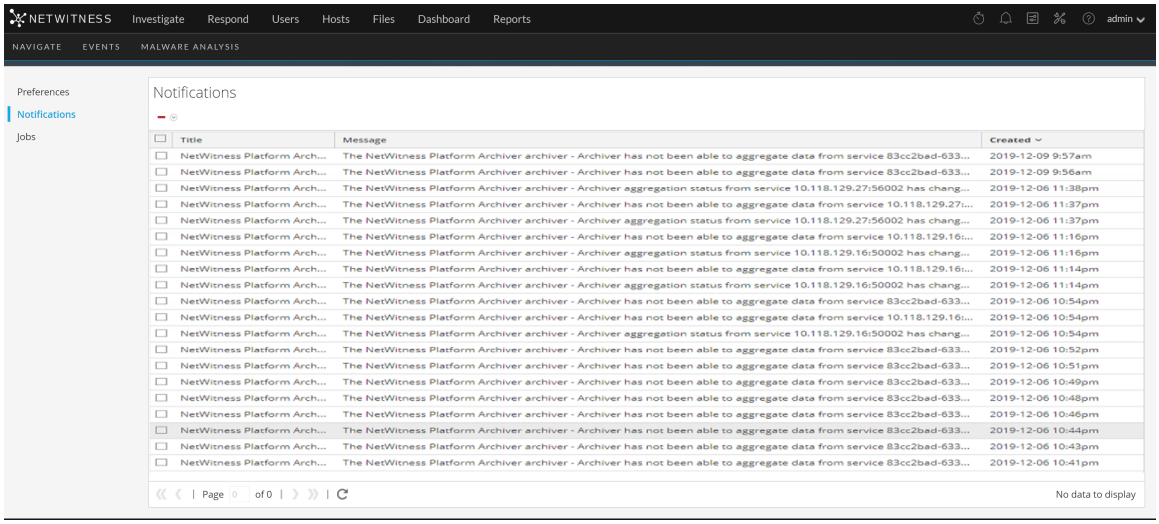
To access the Notifications panel, do one of the following:

- Click  to open the Notifications tray and then click **View All** in the Notifications tray.



- In the upper right corner of the NetWitness browser window, select your username and then select > **Profile**. In the options panel of the Preferences dialog, select **Notifications**.

  The Notifications panel is displayed.



The Notifications tray shows your recent notifications. It contains a subset of the information in the Notifications panel. The Notifications panel shows all of your notifications. The following table describes the Notifications panel and Notifications tray features.

Feature	Description
	(Notifications panel only) Displays a drop-down menu where you can delete the selected notification or all of your notifications in the Notifications panel and in the Notifications tray.
Title	The title of the notification, for example, **Logs Extraction Complete**.
Message	The entire message, for example, **The log extraction to Investigation is complete and ready for download.**
View	Some messages include a **View** link that displays a view where you can take action. For example, if there is a file to download, clicking this link opens the Jobs panel, the view where you can download the file.
Created	The date and time the notification was created. In the Notifications tray, it shows the number of hours or days since the notification was created.
View All	(Notification tray only) Opens the Notifications panel, which lists all of your notifications.

Jobs Panel and Jobs Tray

Jobs are started by various NetWitness components; for example, downloading Content Management System (CMS) resources from Live Services and extracting logs, meta, and PCAP files from NetWitness Investigate.

In the  (Admin) > System view, Administrators can manage all NetWitness jobs in the Jobs panel. Other non-administrative users can view their own jobs in the user Profile Jobs panel.

In addition, while working in NetWitness, you can open a quick view of your jobs from the NetWitness toolbar. When a job status has changed, the Jobs icon () is flagged with the number of running jobs. Once all jobs are completed, that number disappears.

In the Jobs panel, you can:

- View and sort the jobs
- Pause or resume a job
- Cancel a job
- Delete a job
- Download a job

The structure of the jobs panel is the same in all views.

What do you want to do?

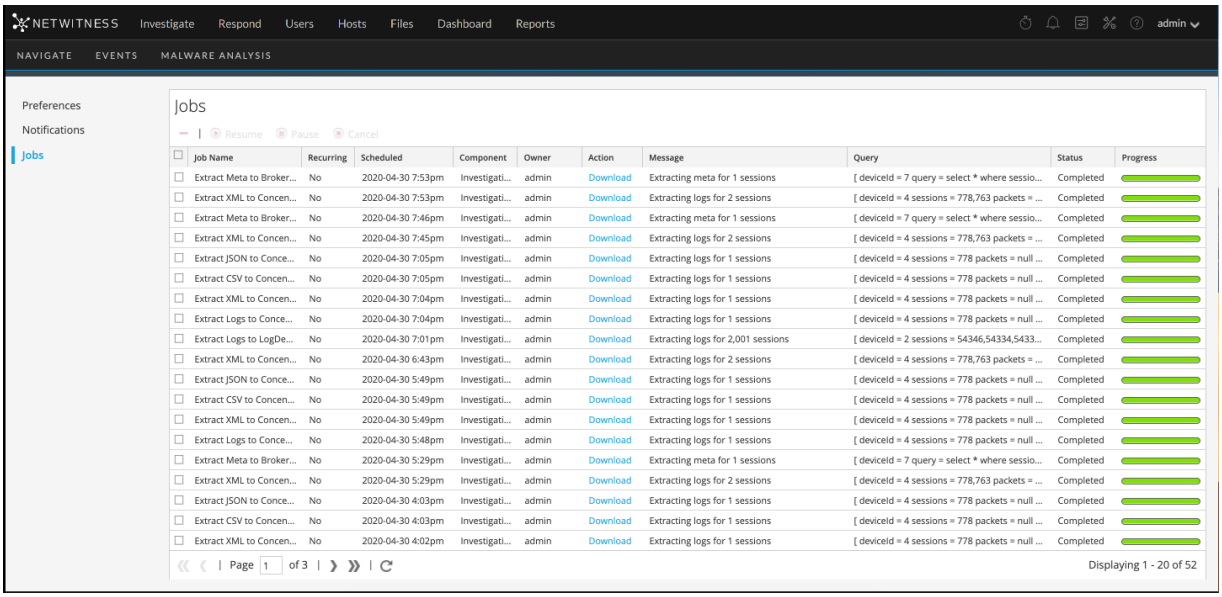
Role	I want to ...	Show me how
All	Pause and Resume a Scheduled Job	Managing Jobs
All	Cancel or Delete a Job	Managing Jobs
All	Download a Job	Managing Jobs

Your actions may be limited to your own jobs depending on your permissions.

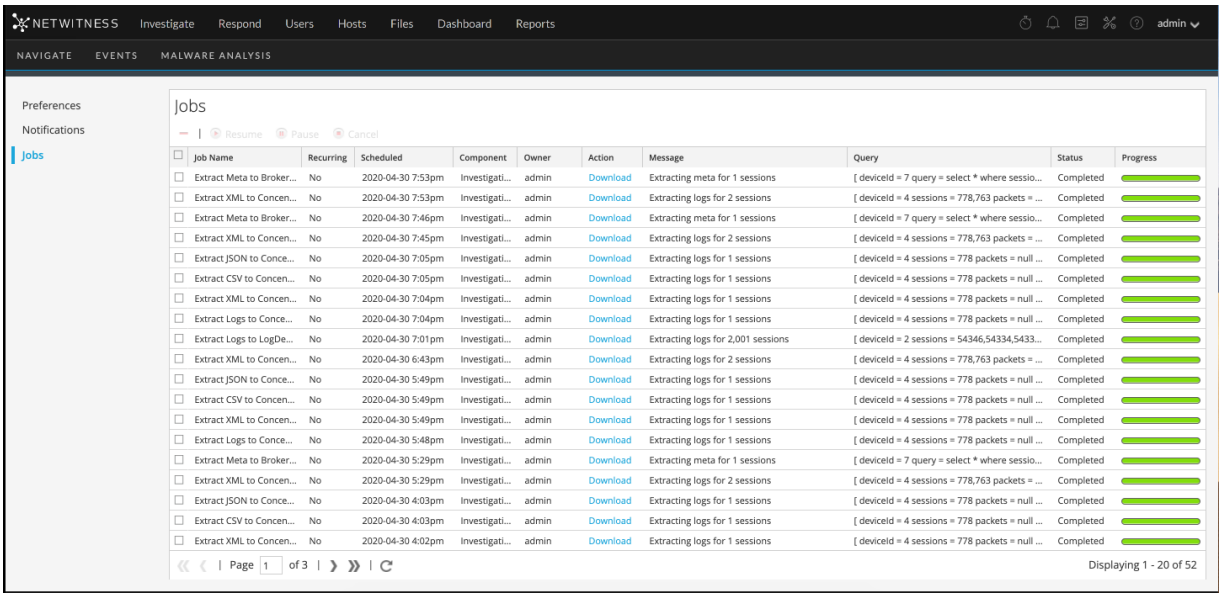
Quick Look

To access the Jobs panel, do one of the following:

- In the upper right corner of the NetWitness browser window, select your username and then select **Profile**. In the options panel of the Preferences dialog, select **Jobs**.
  The Jobs panel is displayed. It shows the jobs of a particular user.

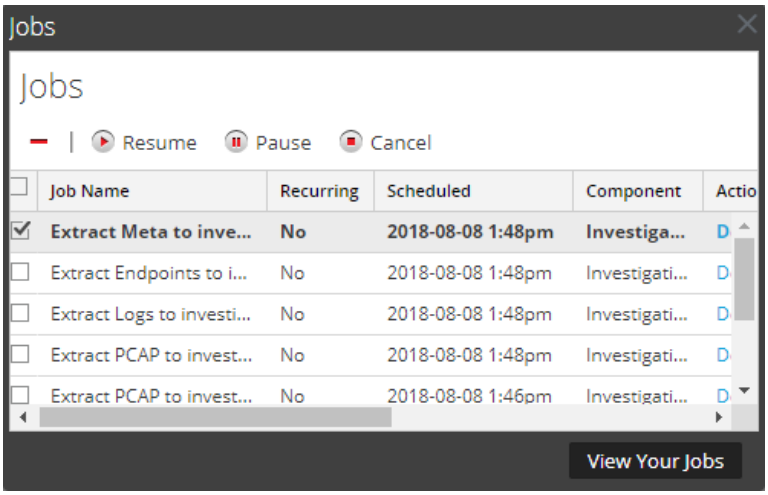


- Go to **(Admin) > System**, and in the options panel, select **Jobs**.
  The Jobs panel in the Admin System view is displayed. It shows the jobs for all users.



The Jobs panel organizes information about jobs into a list. The columns present a job progress bar, the job name, an indication that the job is recurring or not recurring, the NetWitness component that is controlling the job, the owner of the job, the status, any associated message, and a download button to allow downloading of a job's packet capture files or payload files.

To display the Jobs tray, click the **Jobs** icon .



The Jobs tray lists all jobs that you own, recurring and non-recurring, using a subset of the columns available in the **Jobs** panel. Otherwise the Jobs tray and the user Profile Jobs panel are the same. In the Admin System view, the Jobs panel lists information about all NetWitness jobs for all users.

The following table describes the available options in the Jobs panel.

Option	Description
Resume	The Resume option applies only to recurring jobs that have been paused. When you resume a paused job, the next execution of the job executes as scheduled.
Pause	The Pause option applies only to recurring jobs. When you pause a recurring job that is running, it has no effect on that execution. The next execution (assuming the job is still paused) is skipped.
Cancel	Cancels a recurring or non-recurring job. You can cancel a job while it is running. If you cancel a recurring job, it cancels that execution of the job. The next time the job is scheduled to run, it executes normally.
–	Deletes a recurring or non-recurring job from the Jobs panel. When you delete a job, the job is instantly deleted from the Jobs panel. No confirmation dialog is offered. If you delete a recurring job, all future executions are removed as well.

The following table describes the Jobs tray and Jobs panel columns.

Column	Description
Selection box	Enables you to select one or more jobs.
Job Name	Displays the name of the job; for example, **Extract Files** or **Upgrade Service**.
Recurring	Indicates whether the job is recurring or non-recurring. **Yes** = recurring, **No** = non-recurring.
Scheduled	Indicates the date and time at which the job was scheduled to begin.
Component	Indicates the component in which the job originated; for example, **Investigation** or **Administration**.
Owner	Indicates the owner of the job. The owner of the job is not included in the default **Jobs Tray**, because only the current user's jobs are displayed here. The column is available to add.
Action	Views the job in another view or downloads job files for the job to the default **Downloads** directory on the local system. Only successfully completed jobs have the **View** link in the **Action** column. Only jobs that create a file have the **Download** link in the **Action** column.
Message	Displays additional information about the job; for example, **Extracting files** or **No sessions found.**
Query	Displays the query associated with the job to help you understand any issues with the result. This example provides the deviceid and the query that was submitted: `[Deviceid = 7 query = select * where sessionid=35667 size = 0 flags = 0]`
Status	Indicates the status of the job. Common values for status are **Paused**, **Running**, **Canceled**, **Failed**, **Completed**, and other status values are possible.
Progress	Shows the percentage complete for a job.
View Your Jobs	(Jobs tray only) Displays your jobs in the **Jobs panel**.

NetWitness® Platform XDR

Version 12.0

Recovery Tool User Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Disaster Recovery (Backup and Restore Instructions) 4

(Recommended) NetWitness Recovery Wrapper Tool 4

Basic Usage of the NetWitness Recovery Wrapper Tool 5

Required Conditions 6

Status Check 10

Troubleshooting 10

NetWitness Recovery Tool (NRT) 11

Basic Usage of the NetWitness Recovery Tool 13

Required Conditions 14

Disaster Recovery Workflow 15

Back Up and Restore Data for Hosts 15

Back Up and Restore Data on the NetWitness Server 16

Back Up Data on a NetWitness Server Host 16

Restore Data on a NetWitness Server Host 17

Back Up and Restore Data on Other Component Hosts 19

Back Up Data on a Component Host 19

Restore Data on a Component Host 20

Hardware Refresh Only - Use Additional Space in New Hardware Hosts 23

Disaster Recovery in Azure Deployment 24

Task 1 - Backup and Export Data 24

Task 2 - Restore and Import Data 24

Disaster Recovery in AWS Deployment 26

Task 1 - Backup and Export Data 26

Task 2 - Restore and Import Data 26

Appendix A. Modify fstab for Series 5 and 6 Hybrid Storage After Recovery 28

Sample etc/fstab File Before Disaster 28

Sample etc/fstab File After Recovery - Before Modification 29

Sample etc/fstab File After Recovery - After Modification 30

Disaster Recovery (Backup and Restore Instructions)

You can take a backup and restore of NetWitness Hosts using any of the following:

- (Recommended) NetWitness Recovery Wrapper Tool
- NetWitness Recovery Tool (NRT)

(Recommended) NetWitness Recovery Wrapper Tool

Note: NetWitness Recovery Wrapper tool is supported from NetWitness 11.6.1.4 and Later. In case of host with large volume of data (>500GB), NetWitness recommends to use NetWitness Recovery Tool (`nw-recovery-tool`) for backup.

NetWitness Recovery Wrapper Tool (NRWT) provides centralized backup and restore that makes it easy for you to take a backup of all the supported installation options (Physical Host, Virtual Host, AWS, and Azure). With NRWT you can:

- Backup (export) an individual, a specific, or all hosts at a time.
- Restore (import) an individual host at a time.
- Customize files or folders during backup and restore.
- Copy back up data to/from remote host location from/to Netwitness hosts provided:
  - Remote host is reachable via SSH from each NetWitness hosts.
  - The credentials are correct.
  - The location specified has sufficient space to take a backup in case of export.
  - The location specified should contain valid backup data in case of import.
- (For version 11.7.1 and Later) Back up Mongo databases for Endpoint and ESA instances
- (For version 11.7.1 and Later) Include Broker index for NetWitness node in which Broker service is running
- (For version 11.7.1 and Later) Back up custom files and folders provided by user

For details on previous run, check NRWT logs at `/var/log/netwitness/recovery-tool/nw-recovery-wrapper.log` on Admin Server.

Basic Usage of the NetWitness Recovery Wrapper Tool

You can use the NRWT to back up data by using the `export` option. To restore data, use the `import` option. The basic usage of the tool is to run the following command from the root directory level:

```
nw-recovery-wrapper [command] [option]
```

The commands and options that you can use with this tool are described in the following tables.

Commands and Options	Description
`-h --help`	Display help on commands and option. For example, specify: `nw-recovery-wrapper --help` to get a list of supported operations and arguments.
`-e, --export`	Export data or configuration.
`-i, --import`	Import data or configuration.
`-d, --dump-dir <path>`	Path for the where data will be exported or imported from (for example, `/var/netwitness/backup`).
`--host-key HOST_KEY [HOST_KEY ...]`	Host IP, ID or display name.
`--host-all`	specify for all hosts - supported only for export
`--include CUSTOM_PATH [CUSTOM_PATH ...]`	Custom path or file.
`--remote-location REMOTE_LOCATION`	Remote host path for remote host configuration.
`--remote-ip REMOTE_IP`	Remote host IP for remote host configuration.
`--remote-password REMOTE_PASSWORD`	Remote host password for remote host configuration.
`--remote-user REMOTE_USER`	User for remote host configuration.

Required Conditions

- Make sure that there is adequate disk space on dump directory to take the backup on each NetWitness Hosts.
- Valid Host key is entered. Host key can be Host ID, IP address or display name.

Back Up using NRT wrapper:

1. Backup NetWitness Hosts and store on local dump directory of each hosts:
   ```
   nw-recovery-wrapper export --dump-dir <dir> --host-key <Host 1 IP/ID/Name> <Host 2 IP/ID/Name>......<Host N IP/ID/Name>

   nw-recovery-wrapper export --dump-dir <dir> --host-all
   ```
2. (Optional) Add custom files or folders during backup and restore other than what is predefined in recovery tool:

   Note: Make sure the custom files or directories are available on NetWitness Hosts, if not, the files or directories will be ignored.

   ```
   nw-recovery-wrapper export --dump-dir <dir> --include-file <custom files>/--include-dir --host-key <Host 1 IP/ID/Name> <Host 2 IP/ID/Name>......<Host N IP/ID/Name>
   nw-recovery-wrapper export --dump-dir <dir> --include-file <custom files>/--include-dir --host-all
   ```
3. (Optional) Copy backup data to remote Location:

   Note: Make sure that:
   - You specify valid values for `--remote-ip`, `--remote-location`, `--remote-password` arguments for remote copy operation.
   - Remote Host IP is valid and reachable via SSH from all NetWitness Hosts.
   - Remote Host location (`--remote-location`) has adequate space to take backup.

   ```
   nw-recovery-wrapper export --dump-dir <dir> --host-key <Host 1 IP/ID/Name> <Host 2 IP/ID/Name>......<Host N IP/ID/Name> --remote-ip <IP ADDRESS of remote host> --remote-password <ssh-password> --remote-location <remote-location-where-backups-should-be-copied-to>
   nw-recovery-wrapper export --dump-dir <dir> --host-all --remote-ip <IP ADDRESS of remote host> --remote-password <ssh-password> --remote-location <remote-location-where-backups-should-be-copied-to>
   ```

   Note: Optional argument `--remote-user` defaults to root if you do not specify any value.

   Example:

   For adminserver, the backup folder name will be adminserver-backup-2021-09-08-12:48:13

4. Backup (export) including custom files or folders and copy to remote location:

> **Note:** Make sure that:
> - the custom files or directories are available on NetWitness Hosts, if not, the files or directories will be ignored.
> - You specify valid values for `--remote-ip`, `--remote-location`, `--remote-password` arguments for remote copy operation.
> - Remote Host IP is valid and reachable via SSH from all NetWitness Hosts.
> - Remote Host location (`--remote-location`) has adequate space to take backup.

```
nw-recovery-wrapper export --dump-dir <dir> --include <custom files/folder>
--host-key <Host 1 IP/ID/Name> <Host 2 IP/ID/Name>......<Host N IP/ID/Name>
--remote-ip <IP ADDRESS of remote host> --remote-password <ssh-password> --
remote-location <remote-location-where-backups-should-be-copied-to>
```

```
nw-recovery-wrapper export --dump-dir <dir> --include <custom files/folder>
--host-all --remote-ip <IP ADDRESS of remote host> --remote-password <ssh-
password> --remote-location <remote-location-where-backups-should-be-
copied-to>
```

```
optional argument: --remote-user defaults to root if argument is not
specified.
```

Example:

For Admin server, the backup folder name will be adminserver-backup-2021-09-08-12:48:13

5. (For version 11.7.1 and Later) (Optional) Include Mongo service.

> **Note:** Make sure that:
> - Mongo service is running on the NetWitness host.
> - `--host-all and --host-key` with multiple values are not supported for include Mongo operation.

```
nw-recovery-wrapper export --dump-dir <dir> --host-key <Host 1 IP/ID/Name>
--include-mongo
```

6. (For version 11.7.1 and Later) (Optional) Include Broker index.

> **Note:** Make sure that:
> - Broker service is running on the NetWitness host.

```
nw-recovery-wrapper export --dump-dir <dir> --host-key <Host 1 IP/ID/Name>
<Host 2 IP/ID/Name>......<Host N IP/ID/Name> --include-broker-index
```

```
nw-recovery-wrapper export --dump-dir <dir> --host-all --include-broker-
index
```

7. (For version 11.7.1 and Later) (Optional) Backup (export) including Mongo and Broker index.

> **Note:** Make sure that:
> - Mongo service is running on the NetWitness host.
> - Broker service is running on the NetWitness host.
> - `--host-all and --host-key` with multiple values are not supported for include Mongo operation.

```
nw-recovery-wrapper export --dump-dir <dir> --host-key <Host 1 IP/ID/Name>
--include-mongo --include-broker-index
```

8. (For version 11.7.1 and Later) (Optional) Backup (export) including custom files or folders, copying to remote location, Broker index and Mongo.

> **Note:** Make sure that:
> - Custom files or directories are present on NetWitness hosts to be backedup, if it is not present it skips the files or directory.
> - Fields such as `--remote-ip`, `--remote-location`, `--remote-password` are mandatory for remote copy operation.
> - Remote host IP credentials should be valid and reachable via SSH from all NetWitness hosts.
> - Remote host location (`--remote-location`) should have sufficient space to contain backups.
> - Mongo service is running on the NetWitness host.
> - Broker service is running on the NetWitness host.
> - `--host-all and --host-key` with multiple values are not supported for include Mongo operation.

```
nw-recovery-wrapper export --dump-dir <dir> --include-broker-index --
include-mongo ---include-file <custom files>/--include-dir <custom folders>
--host-key <Host 1 IP/ID/Name> --remote-ip <IP ADDRESS of remote host> --
remote-password <ssh-password> --remote-location <remote-location-where-
backups-should-be-copied-to>
```

Restore (import) options supported in NRT Wrapper

> **Caution:** Use import commands carefully as it performs system level changes.

1. Restore (import) single host at a time (using IP address, Host name, or Host ID).
   ```
   nw-recovery-wrapper import --dump-dir <dir> --host-key <Host IP/ID/Name>
   ```
2. Restore custom files or folders (if any).

   > **Note:** Make sure the custom files or directories are available on NetWitness Hosts, if not, the files or directories will be ignored.

   ```
   nw-recovery-wrapper import --dump-dir <dir> --include-file <custom files>/-
   -include-dir --host-key <Host IP/ID/Name>
   ```
3. Restore from a remote location.

   > **Note:** Make sure that:
   > - `--remote-location` contains remote host location in which data is backedup.
   > - Remote Host IP is valid and reachable via SSH from all NetWitness Hosts.
   > - Remote Host location (`--remote-location`) has adequate space to take backup.

   ```
   nw-recovery-wrapper import --remote-ip <IP address of remote host> --
   remote-password <SSH password of remote host> --remote-location <location-
   of-backup-on-remote-host> --dump-dir <dir> --host-key <Host IP/ID/Name>
   ```
   Optional argument: `--remote-user` defaults to root if argument is not specified.
   Example, for adminserver, the backup folder name should be adminserver-backup-2021-09-08-12:48:13
   ```
   nw-recovery-wrapper import --dump-directory <dir> --host-key <host-1> --
   remote-ip <remote-ip> --remote-password <password> --remote-location
   /home/adminserver-backup-2021-09-08-12:48:13
   ```

4. Restore data from remote location including custom files or folders.

   Note: Make sure that:
   - The custom files or directories are available on NetWitness Hosts, if not, the files or directories will be ignored.
   - `--remote-location` contains remote host location in which data is backedup.
   - Remote Host IP is valid and reachable via SSH from all NetWitness Hosts.
   - Remote Host location (`--remote-location`) has adequate space to take backup.

   ```
   nw-recovery-wrapper import --dump-dir <dir> --include <custom files/folder>
   --host-key <host1> --remote-ip <IP ADDRESS of remote host> --remote-
   password <ssh-password> --remote-location <remote-location-where-backups-
   should-be-copied-to>
   ```

   Optional argument: `--remote-user` defaults to root if argument is not specified.

   Example, for Admin Server, the backup folder name will be adminserver-backup-2021-09-08-12:48:13

5. (For version 11.7.1 and Later) (Optional) Restore Mongo service.

   Note: Make sure that:
   - Mongo service is running on the NetWitness host.
   - `--host-all and --host-key` with multiple values are not supported for include Mongo operation.

   ```
   nw-recovery-wrapper import --dump-dir <dir> --host-key <Host 1 IP/ID/Name>
   --include-mongo
   ```

6. (For version 11.7.1 and Later) (Optional) Restore Broker index.

   Note: Make sure that:
   - Broker service is running on the NetWitness host.
   - `--host-all option is not support for include broker index operation.`

   ```
   nw-recovery-wrapper import --dump-dir <dir> --host-key <Host 1 IP/ID/Name>
   --include-broker-index
   ```

7. (For version 11.7.1 and Later) (Optional) Restore Mongo and Broker index.

   Note: Make sure that:
   - Mongo service is running on the NetWitness host.
   - Broker service is running on the NetWitness host.
   - `--host-all and --host-key` with multiple values are not supported for include Mongo operation.

   ```
   nw-recovery-wrapper import --dump-dir <dir> --host-key <Host 1 IP/ID/Name>
   --include-mongo --include-broker-index
   ```

8. (For version 11.7.1 and Later) (Optional) Restore custom files or folders, copying to remote location, Broker index, and Mongo.

Note: Make sure that:
- Custom files or directories are present on NetWitness hosts to be backedup, if it is not present the files or directory is skipped for backup.
- Fields such as `--remote-ip`, `--remote-location`, `--remote-password` are mandatory for remote copy operation.
- Remote host IP credentials should be valid and reachable via SSH from all NetWitness hosts.
- Remote host location (`--remote-location`) should have sufficient space to contain backups.
- Mongo service is running on the NetWitness host.
- Broker service is running on the NetWitness host.
- `--host-all and --host-key` with multiple values are not supported for include Mongo operation.

```
nw-recovery-wrapper import --dump-dir <dir> --include-file <custom files>/-
-include-dir <custom folders> --include-mongo --include-broker-index --
host-key <host1> --remote-ip <IP ADDRESS of remote host> --remote-password
<ssh-password> --remote-location <remote-location-where-backups-should-be-
copied-to>
```

Status Check

You can check the Backup or Restore status using the below command.

`/var/log/netwitness/recovery-tool/recovery.log`

Troubleshooting

Error Message	NRT Wrapper fails during backup or restore.
Solution	Do any one of the following: • Log in to host where backup is failing and check `/var/log/netwitness/recovery-tool/recovery.log`. • Run in debug mode (`nw-recovery-wrapper -l debug`) on Node 0 to get recovery logs of each host.

Error Message	NRT Wrapper fails due to incorrect password for remote copy operation (`--remote-password`).
Cause	NRWT

	fails if you enter wrong password multiple times during remote copy. Since SFTP uses SSH, it locks the system SSH for a while.
Solution	You must retry after sometime.

Error Message	NRT Wrapper fails after running for long hours for a particular host but the backup is still in progress. For example, Endpoint or ESA node.
Cause	NRT architecture is designed to take a backup of configuration and not user data. It uses salt to communicate between admin server(SA) and nodex. Due to large volume of data in the particular host, this issue is caused when salt communication times out.
Solution	SSH to the particular host and check the backup status at `/var/log/netwitness/recovery-tool/recovery.log`. If the data to be backed up is huge, recommend using "nw-recovery-tool" by logging into particular host.

NetWitness Recovery Tool (NRT)

You can use the NetWitness Recovery Tool (NRT) to back up and restore data from the NetWitness Server and component host systems. The NRT is a script that you run from the command line to back up and restore data on hosts for RMAs, hardware refreshes, and general backup and restore requirements. Refer to Disaster Recovery in Azure Deployment for specific steps on how to perform disaster recovery for hosts deployed in Azure VMs.

Note: You must run the NRT on each host system locally. You cannot run it from remote hosts or an external host.

The following types of hosts can be backed up and restored.

Note: In the NRT script, the following terms in bold are referred to as categories.

- **NetWitness Admin Server** (may include Broker, Investigate, Respond, Health and Wellness, and Reporting Engine)
- **AnalystUI** (may include Broker, Investigate, Respond, Reporting Engine)
- **Archiver** Log Archiver (Workbench and Archiver)
- **Broker** Stand-alone Broker
- **Concentrator** Network or Log Concentrator
- **Decoder** Network Decoder (Packets)
- **Endpoint** Endpoint Agents
- **Endpoint Broker** Endpoint Broker
- **Endpoint Log Hybrid** Log Collector, Log Decoder, Endpoint Server, and Concentrator
- **Event Stream Analysis (ESA) Primary** Contexthub, ESA Correlation, and Incident Management database
- **ESA Secondary** ESA Correlation
- **Gateway** Cloud Gateway
- **Log Hybrid Retention** Log Hybrid-Retention Optimized (for RSA Series 6 Hybrid hardware with Log Hybrid-Retention Optimization)
- **Log Collector** Log Collector including Virtual Log Collector if installed
- **Log Decoder** Log Decoder including Local Log Collector and Warehouse Connector, if installed
- **Log Hybrid** Log Collector, Log Decoder, and Concentrator
- **Malware** Malware Analysis and Broker
- **Network Hybrid** Concentrator and Decoder
- **Search** (for Health & Wellness Beta Host)
- **UEBA** User Entity and Behavior Analytics
- **Warehouse** Warehouse Connector

Basic Usage of the NetWitness Recovery Tool

You can use the NRT to back up data by using the `export` option. To restore data, use the `import` option. The basic usage of the tool is to run the following command from the root directory level:

```
nw-recovery-tool [command] [option]
```

The commands and options that you can use with this tool are described in the following tables.

Commands and Options	Description
`-h, --help`	Display help on commands and option. For example, specify: `nw-recovery-tool --help-categories` to get a list of all the valid category names.
`-e, --export`	Export data or configuration.
`-i, --import`	Import data or configuration.
`-d, --dump-dir <path>`	Path for the where data will be exported or imported from (for example, `/var/netwitness/backup`).
`-C, --category <name>`	Select components by category. Valid category names are `AdminServer`, `AnalystUI`, `Archiver`, `Broker`, `Concentrator`, `Decoder`, `Endpoint`, `EndPointBroker`, `EndpointLogHybrid`, `ESAPrimary`, `ESASecondary`, `Gateway`, `LogHybridRetention`, `LogCollector`, `LogDecoder`, `LogHybrid`, `Malware`, `NetworkHybrid`, `Search`, `UEBA`, and `Warehouse`. You can specify a single category or multiple categories if multiple categories are co-located on the same host. For example: • `--category AdminServer` for the Admin Server exclusively. `--category AdminServer --category Gateway` for the Admin Server and the Cloud Gateway. • `--category ESAPrimary` for the ESA Primary exclusively. • `--category Broker` for the Broker exclusively. `--category Broker --category EndpointBroker` for the Broker and the Endpoint Broker.

Required Conditions

Make sure that the following conditions are met:

- Read the entire document before backing up any data. The document covers all deployment scenarios, so you want to make sure you have all the information required to back up and restore your implementation of NetWitness Platform before going through this process.
- Run the NRT for both backup and recovery locally, on each system being backed up or restored. You cannot run the NRT on an external host, or back up or restore several hosts simultaneously. However, you can back up several components on the same host system simultaneously.
- Export and import data on the same host. If a host fails and you need to build a new system, the new system must have the same identity parameters (i.e., the same IP address), and must be on the same version of NetWitness Suite
- Make sure that there is adequate disk space in the backup location (`/var/netwitness/backup` is the recommended directory) before the export command in the nw-recovery tool is executed. Do not use a `tmp` directory because it fills up quickly and may cause the system to crash.
- Check the sizing of the Malware disks and adjust them before you back them up. The following table shows you the maximum size of Malware databases that you can back up by hardware type with the actions you can take to reduce them to the maximum size.

Host	Source Hardware	Target Hardware	Database	Maximum Size for Backup	Actions to Reduce Size to Backup Maximum
Malware	Series 4S Hybrid	Series 6 Core	`/var/netwitness`	2.5TB	Configure a rollover. Purge data that you do not need from the database.

- Restore to the exact ISO Image that each host had at the time of backup.
- If you have multiple services co-located on a single host, include all the services in a single command string for the `import` and `export` commands in the nw-recovery tool.

Note: 1.) When you run the NRT, the Malware , Reporting Engine, and Postgresql services are stopped and restarted during both the backup (export) and restore (import) processes.

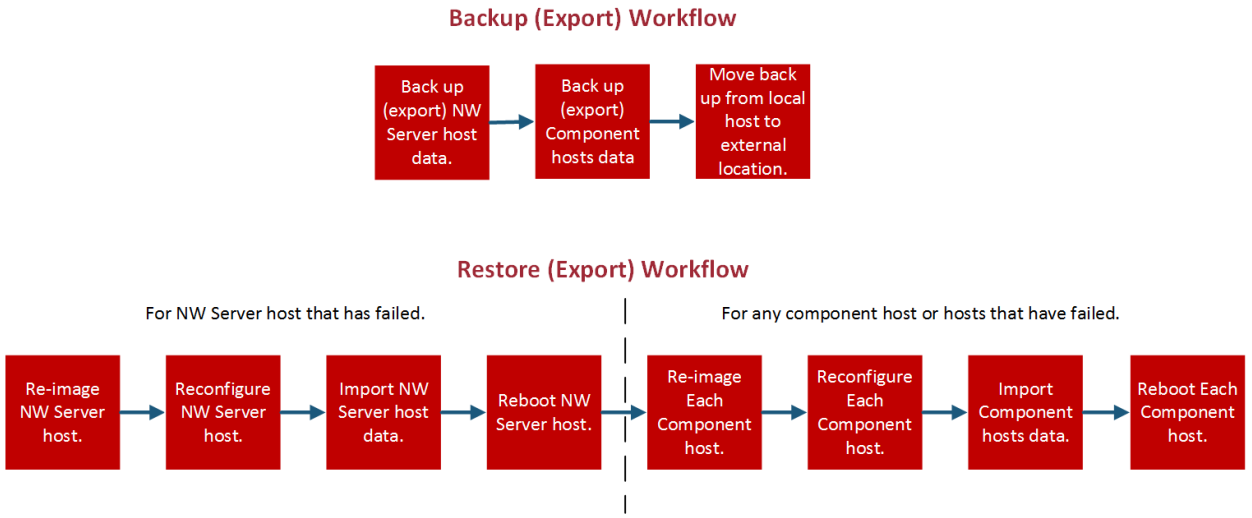
Disaster Recovery Workflow

The following diagram shows the high-level Disaster Recovery tasks.

Note: You only need to recover a host if it failed. This means that you can recover a single host, or any combination of hosts depending on which host or hosts failed.

The following diagram shows the tasks for:

- Backup (perform as soon as possible and as frequently as possible).
- Restore (only required if you need to restore your data).



Back Up and Restore Data for Hosts

The procedures for backing up and restoring data are different for NetWitness Server host systems and for component systems.

Caution: 1.) Do not remove component hosts (that is any host other than the NW Server host) from the Hosts View (**Admin** > **Hosts**) from the user interface when you are performing the following disaster recovery procedure. 2.) You must retain (restore) the 'Host name' that existed prior to performing the disaster recovery procedure.

Back Up and Restore Data on the NetWitness Server

Note: If you are using shared storage to export data from multiple hosts (for example, a shared mount or drive), use host-specific subfolders for the path to the location of the exported files for each host, to avoid overwriting one host's exported data with another. For example, you could use a path similar to `--dump-dir /mnt/storage/<host-specific-name>` for the path to the location of the exported files.

Back Up Data on a NetWitness Server Host

Perform this procedure on an existing, functional NetWitness Server host system.

1. At the root level, type the following command:

   ```
   nw-recovery-tool --export --dump-dir /var/netwitness/backup --category
   AdminServer
   ```

   Note: If a service is co-located with another category on the same host rather than on its own, dedicated host, you must include it in the command string. The `Gateway` and `EndpointBroker` can be co-located as show in the following examples:

   ```
   nw-recovery-tool--export --dump-dir /var/netwitness/backup --category
   AdminServer --category Gateway
   nw-recovery-tool--export --dump-dir /var/netwitness/backup --category
   Broker --category EndpointBroker
   ```

2. Replace `/var/netwitness/backup` with the path to the location to which the data should be exported.

   a. Ensure that this location has sufficient space to store the backup data.

   b. The backup directory path should be located on the local host. However, the backup files could be located on a network mount or an external device.

The data is backed up on the NetWitness Server host in the location you set up in step 2.

3. Move the backed up data from the local host to an external server or a USB stick.

Restore Data on a NetWitness Server Host

1. Re-image the NetWitness Server host using the same network configuration settings of the original host. For information about re-imaging the NetWitness Server host, see "Task 1 - Install 12.0 on the NetWitness Server Host" in the *Physical Host Installation Guide for Version 12.0 Guide*

   a. **Optional** If you need to establish network connectivity before you can fetch backup data, for example, if it is on a remote host, run the following script using the same IP address, subnet, gateway, DNS and domain information as the original host:

   ```
   netconfig --static --interface <name> --ip <address> --netmask <netmask>
   --gateway <gateway>
   ```

   For example:

   ```
   netconfig --static --interface eth0 --ip 192.168.1.100 --netmask
   255.255.255.0 --gateway 192.168.1.1
   ```

   Optional: To specify DNS server(s), include the following additional parameter:

   ```
   --dns <address>
   ```

   Optional: To set the local domain name, include the following additional parameter:

   ```
   --domain <name>
   ```

   b. (**Optional**) If you are using DHCP, run the following script:

   ```
   netconfig --dhcp --interface <name>
   ```

   For example:

   ```
   netconfig --dhcp --interface eth0
   ```

   c. Add the backup data to the backup directory path on the local host, for example:

   ```
   /var/netwitness/backup
   ```

2. Run the `nwsetup-tui` command. This initiates the Setup program.

   Note: During the Setup program, when you are prompted for the network configuration of the host, be sure to specify the same identical network configuration that was used for the original installation of the host.

3. When you are prompted, select install type option **2: Recover (Reinstall)**, click **OK**, and then enter the path to the backup directory containing the backup data.
4. After the installation completes successfully, ensure that the host is running the exact same release and patch version of the data that was backed up:

   - If the data was on an system that was updated to a later patch release, update the host by following the instructions for updating systems offline in the update guide for the same patch version as what was previously running on the host (the exact release or patch version for which data was backed up).
   - If the data was on a major release version (for example, 12.0) that had not been updated to a later patch version, you do not need to update the host system.

5. When the host is running at the correct version, run the following command on the NetWitness Server to restore data:

```
nw-recovery-tool --import --dump-dir /var/netwitness/backup --category
AdminServer
```

Note: If a service is co-located with another category on the same host rather than on its own, dedicated host, you must include it in the command string. The `Gateway` and `EndpointBroker` can be co-located as show in the following examples:

```
nw-recovery-tool--import --dump-dir /var/netwitness/backup --category
AdminServer --category Gateway
nw-recovery-tool--import --dump-dir /var/netwitness/backup --category
Broker --category EndpointBroker
```

6. (Conditional) If you use custom firewall rules or custom entries in `/etc/hosts`:
   a. (Conditional) If you use custom firewall rules (that is, replied "Yes" to the "Disable Firewall" nwsetup-tui prompt during installation), restore the `/etc/sysconfig/iptables` file from the backup copy located in the `<dump-dir>/unmanaged/etc/sysconfig/iptables` file.
   b. (Conditional) If you use custom entries in `/etc/hosts`, restore the `/etc/hosts.users` file from the backup copy located at `<dump-dir>/unmanaged/etc/hosts.user` to `/etc` on the host.
   c. If you performed step 6a or 6b, refresh the host by running the following command:
   `nw-manage --refresh-host --host-key <ID, IP, hostname or display name of host>`
7. Reboot the NetWitness Server host.

Note: If you want to add any more custom entries to `/etc/host`, you must add them to the `/etc/hosts.users` file and then refresh the host as described in step 6c.

Back Up and Restore Data on Other Component Hosts

Perform these procedures on each existing, functional component host system.

Back Up Data on a Component Host

1. At the root level, type the following command:
   `nw-recovery-tool --export --dump-dir /var/netwitness/backup --category <category name>`
   where the category name is one of the following:
   `AdminServer`, `AnalystUI`, `Archiver`, `Broker`, `Concentrator`, `Decoder`, `Endpoint`, `EndPointBroker`, `EndpointLogHybrid`, `ESAPrimary`, `ESASecondary`, `Gateway`, `LogHybridRetention`, `LogCollector`, `LogDecoder`, `LogHybrid`, `Malware`, `NetworkHybrid`, `Search`, `UEBA`,or `Warehouse`

   Note: 1.) Use the category that matches the host type. 2.) If services are co-located on a Component Host rather than on its own dedicated host, you must include it in the command string. For example, a Warehouse Connector resides on a Log Decoder host. The following is an example of this command string.
   `nw-recovery-tool--export --dump-dir /var/netwitness/backup --category LogDecoder --category Warehouse`

2. **(Optional)** Replace `/var/netwitness/backup` with the path to the location to which the data should be exported
   a. Ensure that this location has sufficient space to store the backup data.
   b. The backup directory path should be located on the local host. However, the backup files could be located on a network mount or an external device.

3. For **Endpoint Log Hybrid** and **ESA Primary** hosts, you can export application data that is stored in the database by running the following command:
   `nw-recovery-tool --export --dump-dir /var/netwitness/backup --component mongo`
   You can replace `/var/netwitness/backup` with the path to the location to which the data should be exported.

   Note: 1.) Make sure that there is enough space in the export location for the files from the Mongo database. 2.) You can back up the **Endpoint Log Hybrid** or **ESA Primary** host data and Mongo database in a single command string. For example, `nw-recovery-tool --export --dump-dir /var/netwitness/backup --category EndpointLogHybrid --component mongo`

4. For **Malware**, you can export application data from the Malware application database by running the following command:
   `nw-recovery-tool --export --dump-dir /var/netwitness/backup --component postgresql`
   You can replace `/var/netwitness/backup` with the path to the location to which the data should be exported.

Note: Ensure that there is enough space in the export location for the files from the Malware database.

5. Move the backed up data from the local host to an external server or a USB stick.

Restore Data on a Component Host

1. Re-image the component host using the same network configuration settings of the original host. For information about re-imaging a component host, see "Task 2 - Install 12.0 on Other Component Hosts" in the *Physical Host Installation Guide for Version 12.0 Guide*
2. **Optional** If you need to establish network connectivity before you can fetch backup data, for example, if it is on a remote host, run the following script using the same IP address, subnet, gateway, DNS and domain information as the original host:
   `netconfig --static --interface <name> --ip <address> --netmask <netmask> --gateway <gateway>`
   For example:
   `netconfig --static --interface eth0 --ip 192.168.1.100 --netmask 255.255.255.0 --gateway 192.168.1.1`
   Optional: To specify DNS server(s), include the following additional parameter:
   `--dns <address>`
   Optional: To set the local domain name, include the following additional parameter:
   `--domain <name>`
   a. (**Optional**) If you are using DHCP, run the following script:
      `netconfig --dhcp --interface <name>`
      For example:
      `netconfig --dhcp --interface eth0`
   b. Add the backup data to the backup directory path on the local host, for example, `/var/netwitness/backup`.
3. Run the `nwsetup-tui` command. This initiates the Setup program.

   Note: During the Setup program, when you are prompted for the network configuration of the host, be sure to specify the same identical network configuration that was used for the original installation of the host.
4. When you are prompted, select install type option **2: Recover (Reinstall)**, click **OK**, and then enter the path to the directory containing the backup data.

5. After you complete the `nwsetup-tui` command setup, you must re-install the appropriate services on the host using the **Install** command from the Hosts View in the NetWitness Platform User Interface.
6. After the service installation completes, ensure that the host is running the exact same release and patch version of the data that was backed up:
   - If the data was on system that was updated to a later patch release, update the host by following the instructions for updating systems offline for the same patch version as what was previously running on the host (the exact release or patch version for which data was backed up).
   - If the data was on a major release version (for example, 12.0) that had not been updated to a later patch version, you do not need to update the host system.
7. When the host is running at the correct version, return to the root level of the component host and run the following command to restore data:
   ```
   nw-recovery-tool --import --dump-dir /var/netwitness/backup --category
   <category name>
   ```

   Note: If services are co-located on a Component Host rather than on its own dedicated host, you must include it in the command string. For example, a Warehouse Connector resides on a Log Decoder host. The following is an example of this command string.
   ```
   nw-recovery-tool--import --dump-dir /var/netwitness/backup --category
   LogDecoder --category Warehouse
   ```
8. For **EnpointLogHybrid** and **ESAPrimary** systems, you can import application data to be restored by running the following command:
   ```
   nw-recovery-tool --import --dump-dir /var/netwitness/backup --component
   mongo
   ```
9. For **Malware**, you can import application data from the Malware application database to be restored by running the following command:
   ```
   nw-recovery-tool --import --dump-dir /var/netwitness/backup --component
   postgresql
   ```

10. For a Decoder, Log Decoder , Concentrator, Archiver, Network Hybrid, or Log Hybrid configured with external storage (that is, DAC, SAN, Unity or Powervault):

    a. Scan the `<dump-dir>/unmanaged/etc/fstab` file for devices with mount points that do not exist in the system `/etc/fstab` file.

    IMPORTANT: If you are migrating to new host hardware (that is a new Decoder, Log Decoder, Concentrator, Archiver, Network Hybrid, or Log Hybrid host), before you proceed to the next step you must:
    1. Power off the old hardware host and the external storage device attached to it.
    2. Attach the external storage device to the new host hardware.
    3. Power on both the new host hardware and the external storage device attached to it.

    a. Complete the following steps for each device in the backup copy of `<dump-dir>/unmanaged/etc/fstab`.

        i. Verify that the corresponding device is present and attached. If it not attached, attach it. If the device is no longer applicable, skip it and go to the next device.

        ii. Verify that the mount point directory exists on the file system. If it does not exist, create the directory with the `mkdir <path>` command.

        iii. Add the `fstab` entry from the backup copy to the system `/etc/fstab` file.

        Caution: For a Series 5 or 6 Hybrid, you must restore backed up data to the `/etc/fstab` directory according to the instructions in Appendix A. Modify fstab for Series 5 and 6 Hybrid Storage After Recovery.

    b. Run the following command on each host.

    ```
    mount -a
    ```

11. (Conditional) If you use custom firewall rules or custom entries in `/etc/hosts`:

    a. (Conditional) If you use custom firewall rules (that is, replied "Yes" to the "Disable Firewall" nwsetup-tui prompt during installation), restore the `/etc/sysconfig/iptables` file from the backup copy located in the `<dump-dir>/unmanaged/etc/sysconfig/iptables` file.

    b. (Conditional) If you use custom entries in `/etc/hosts`, restore the `/etc/hosts.users` file from the backup copy located at `<dump-dir>/unmanaged/etc/hosts.user` to `/etc` on the host.

    c. If you performed step 11a or 11b, refresh the host by running the following command:

    ```
    nw-manage --refresh-host --host-key <ID, IP, hostname or display name of host>
    ```

12. Reboot the component host.

Hardware Refresh Only - Use Additional Space in New Hardware Hosts

Refer to the *Core Database Tuning Guide for NetWitness Platform* for instructions on how to use all the space you have available on your new hardware. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Disaster Recovery in Azure Deployment

This section tells you how to back up and restore NetWitness Platform deployed on Azure virtual hosts (also referred to as VMs in this section). The two major tasks to back up and restore data in an Azure deployment are:

- Task 1 - Backup and Export Data
- Task 2 - Restore and Import Data

Task 1 - Backup and Export Data

1. Export the data by running the `nw-recovery-tool --export` commands as described in the Disaster Recovery (Backup and Restore Instructions) section of this document.

Task 2 - Restore and Import Data

You need to refer to the *10.6.6.x to 11.3 Azure Upgrade Guide* to complete this task. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

1. Delete the VM.

   Caution: Do not delete the resources (for example, do not delete Disks, Network Interface, and so on).

2. Complete the following steps for the NW Server host, Broker host, ESA host, Endpoint Log Hybrid host, and Log Collector host (where host = `--category`).

   a. Delete all the resources except the network interface card of the older 12.0 VM.

   b. Deploy the fresh 12.0 VM with the same disk and resources and power it off.
   For detailed instructions on how to deploy a virtual host in Azure, see the *Azure Installation Guide*.

   c. Run the `azure-mac-retention.ps1` from the local machine.
   See the *10.6.6 to 11.3 Azure Upgrade Guide* for instructions on how to run this script.

   d. Follow the procedure for the NRT restoration for the respective host as described in Disaster Recovery (Backup and Restore Instructions).

   e. After you restore NRT the component host, restore the following files.

   - `/etc/fstab`
   - `/etc/hosts` (if hostname is not changed)
   - `/etc/waagent.conf`
   - `/etc/logrotate.d/waagent.logrotate`
   - `/etc/krb5.conf` from the `<dump-dir>/unmanaged` folder

3. Complete the following steps for the Log Decoder host, Concentrator host, and Archiver host (where host = `--category`.

   a. Delete all the resources except the disks that are named **external** and the network interface card of the older 12.0 VM.

   b. Deploy the fresh 12.0 VM with the same disk and resources listed in the *Azure Installation Guide* and power it off.

   > **Note:** Do not create the **external** disk. Only create the **nwhome** disks.

   c. Run the `azure-mac-retention.ps1` from the local machine.
   See the *10.6.6 to 11.3 Azure Upgrade Guide* for instructions on how to run this script.

   d. Follow the procedure for the NRT restoration for the respective hosts as described in Restore Data on a Component Host.

   e. After you restore NRT the component host, restore the following files.

   - `etc/fstab`
   - `/etc/hosts` (if hostname is not changed)
   - `/etc/waagent.conf`
   - `etc/logrotate.d/waagent.logrotate`
   - `/etc/krb5.conf`

Disaster Recovery in AWS Deployment

This section tells you how to back up and restore NetWitness Platform deployed on AWS virtual hosts (also referred to as VMs in this section). The two major tasks to back up and restore data in an AWS deployment are:

- Task 1 - Backup and Export Data
- Task 2 - Restore and Import Data

Task 1 - Backup and Export Data

1. Export the data by running the `nw-recovery-tool --export` commands as described in the Disaster Recovery (Backup and Restore Instructions) section of this document.
2. Record the IP addresses. You need to refer to them later in the Disaster Recovery process. Refer to the *10.6.6 to 11.3 AWS Upgrade Guide* for instructions on how retain the IP addresses. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Task 2 - Restore and Import Data

You need to refer to the *10.6.6 to 11.3 AWS Upgrade Guide* to complete this task.

1. Delete the VM.

   Caution: Do not delete the resources (for example, do not delete Disks).

2. Complete the following steps for the NW Server host, Broker host, ESA (Primary/Seconday) host, Endpoint Log Hybrid host, and Log Collector host (where host = `--category`).
   a. Delete all the resources of the older 12.0 VM.
   b. Deploy the fresh 12.0 VM with the same IP address, disk and resources and power it off. For detailed instructions on how to deploy a virtual host in AWS, see the *AWS Installation Guide*.
   c. Follow the procedure for the NRT restoration for the respective host as described in Restore Data on a Component Host.
   d. After you restore NRT the component host, restore the following files.
      - `/etc/fstab`
      - `/etc/hosts` (if hostname is not changed)
3. Complete the following steps for the Log Decoder host, Decoder (Network Decoder) host, Concentrator host, and Archiver host (where host = `--category`.
   a. Delete all the resources except the **external disks** of the older 12.0 VM.
   b. Deploy the fresh 12.0 VM with the same IP address, disk, and resources listed in the *AWS Installation Guide* and power it off.

Note: Do not create the **external** disk. Only create the **nwhome** disks.

c. Follow the procedure for the NRT restoration for the respective hosts as described in Restore Data on a Component Host.

d. After you restore NRT the component host, restore the following files.

- `etc/fstab`
- `/etc/hosts` (if hostname is not changed)
- `/etc/krb5.conf`

Appendix A. Modify fstab for Series 5 and 6 Hybrid Storage After Recovery

Note: The procedures in this Appendix do not apply if the hybrid was an 11.4 fresh installation and you had a disaster.

You must modify the `/etc/fstab` file for a Series5 or Series 6 Hybrid if a Series 5 or Series 6 Network Hybrid was upgraded from 11.2.x.x or 11.3.x.x to 11.4 and a disaster occurred.

These are the tasks to recovery from this disaster scenario.

1. Image a new Series 5 or Series 6 Hybrid as Network Hybrid using an 11.4 ISO.
2. Import the backed up data or configuration (`nw-recovery-tool --import`).
3. Modify the recovered `/etc/fstab` file.

Sample etc/fstab File Before Disaster

The following data is an example of the external storage configuration backup for a Series 5 or 6 Hybrid that was upgraded to 11.4.

The data highlighted in yellow corresponds to internal storage on an upgraded system. This configuration is carried over from prior releases during upgrade process. This lay out changed in 11.4 (Fresh Install). You must only copy the entries corresponding to external storage (highlighted in green) to the new `etc/fstab` file as part of a disaster recovery.

When you export data or configuration using the `nw-recovery-tool --export` command, the storage configuration details are saved to `<back-location>/unmanaged/etc/fstab`. The `fstab` file contains both internal storage (highlighted in yellow) and external storage configuration (highlighted in green). The contents on a upgraded (10.6 or 11.x -> 11.4) Series 5 or Series 6 Network Hybrid look similar to the following storage configuration.

```
/dev/mapper/netwitness_vg00-root / xfs defaults 0 0 UUID=906e2a3d-3b59-46d1-
975d-fa2b8467d009
/boot xfs defaults 0 0 /dev/mapper/netwitness_vg00-usrhome
/home xfs nosuid 0 0
/dev/mapper/netwitness_vg00-varlog /var/log xfs defaults 0 0
/dev/mapper/netwitness_vg00-nwhome /var/netwitness xfs nosuid,noatime 0 0
/dev/mapper/concentrator-vlnwc /var/netwitness/concentrator xfs noatime,nosuid
0 0
/dev/mapper/index-vlnwci /var/netwitness/concentrator/index xfs noatime,nosuid
0 0
/dev/mapper/concentrator-vlnwcm /var/netwitness/concentrator/metadb xfs
noatime,nosuid 0 0
/dev/mapper/concentrator-vlnwcs /var/netwitness/concentrator/sessiondb xfs
noatime,nosuid 0 0
/dev/mapper/decoderpacket-vlnwd /var/netwitness/decoder xfs noatime,nosuid 0 0
```

```
/dev/mapper/decoderpacket-vlnwdi /var/netwitness/decoder/index xfs
noatime,nosuid 0 0
/dev/mapper/decodermeta-vlnwdm /var/netwitness/decoder/metadb xfs
noatime,nosuid 0 0
/dev/mapper/decoderpacket-vlnwdp /var/netwitness/decoder/packetdb xfs
noatime,nosuid 0 0
/dev/mapper/decoderpacket-vlnwds /var/netwitness/decoder/sessiondb xfs
noatime,nosuid 0 0
/dev/mapper/netwitness_vg00-swap swap swap defaults 0 0
/var/netwitness/decoder /var/netwitness/logdecoder none defaults,rbind 0 0
/dev/concentrator0/sessiondb /var/netwitness/concentrator/sessiondb0 xfs
noatime,nosuid 1 2
/dev/concentrator0/metadb /var/netwitness/concentrator/metadb0 xfs
noatime,nosuid 1 2
/dev/decoder0/packetdb /var/netwitness/decoder/packetdb0 xfs noatime,nosuid 1
```

Sample etc/fstab File After Recovery - Before Modification

After you install 11.4 with the 11.4 ISO and execute the recovery tool execution to restore all previous configurations, the `/etc/fstab` file appears looks like the following example.

```
#
# /etc/fstab
# Created by anaconda on Thu Dec 5 17:31:26 2019
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/netwitness_vg00-root / xfs defaults 0 0
UUID=d84db66c-fce6-4fec-9f84-b3449861f664 /boot xfs defaults 0 0
/dev/mapper/netwitness_vg00-usrhome /home xfs nosuid 0 0
/dev/mapper/netwitness_vg00-varlog /var/log xfs defaults 0 0
/dev/mapper/netwitness_vg00-nwhome /var/netwitness xfs nosuid,noatime 0 0
/dev/mapper/netwitness_vg00-swap swap swap defaults 0 0
/dev/hybrid-decoder-meta/decoroot /var/netwitness/decoder xfs noatime,nosuid 1
2
/dev/packet/packetdb /var/netwitness/decoder/packetdb xfs noatime,nosuid 1 2
/dev/hybrid-concentrator/root /var/netwitness/concentrator xfs noatime,nosuid
1 2
/dev/index/index /var/netwitness/concentrator/index xfs noatime,nosuid 1 2
```

Note: You can see that it is missing the external storage configurations. You must add the external storage configurations highlighted in green above) to the `/etc/fstab` file on the newly built Hybrid.

Sample etc/fstab File After Recovery - After Modification

After you make this updater, the `/etc/fstab` looks like the following example.

```
#
# /etc/fstab
# Created by anaconda on Thu Dec 5 17:31:26 2019
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/netwitness_vg00-root / xfs defaults 0 0
UUID=d84db66c-fce6-4fec-9f84-b3449861f664 /boot xfs defaults 0 0
/dev/mapper/netwitness_vg00-usrhome /home xfs nosuid 0 0
/dev/mapper/netwitness_vg00-varlog /var/log xfs defaults 0 0
/dev/mapper/netwitness_vg00-nwhome /var/netwitness xfs nosuid,noatime 0 0
/dev/mapper/netwitness_vg00-swap swap swap defaults 0 0
/dev/hybrid-decoder-meta/decoroot /var/netwitness/decoder xfs noatime,nosuid 1
2
/dev/packet/packetdb /var/netwitness/decoder/packetdb xfs noatime,nosuid 1 2
/dev/hybrid-concentrator/root /var/netwitness/concentrator xfs noatime,nosuid
1 2
/dev/index/index /var/netwitness/concentrator/index xfs noatime,nosuid 1 2
/dev/concentrator0/sessiondb /var/netwitness/concentrator/sessiondb0 xfs
noatime,nosuid 1 2
/dev/concentrator0/metadb /var/netwitness/concentrator/metadb0 xfs
noatime,nosuid 1 2
/dev/decoder0/packetdb /var/netwitness/decoder/packetdb0 xfs noatime,nosuid 1
```

NetWitness® Platform XDR

Version 12.0

Reporting Engine Configuration Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

How Reporting Engine Works **5**
Configure the Reporting Engine **6**
Configure the Data Sources **7**
Configure a NWDB Data Source 7
Configure a Warehouse Data Source 9
Enable Kerberos Authentication 12
Set a Data Source as the Default Source 14
(Optional) Add Workbench as Data Source 15
(Optional) Add Archiver as Data Source 18
(Optional) Integrate Endpoint Information Into Reports 20
(Optional) Add Collection as Data Source to Reporting Engine 21
Configure Data Privacy for the Reporting Engine **24**
Configure Data Source Permissions **28**
Configure Reporting Engine Settings **30**
Enable LDAP Authentication 30
Add Additional Space for Large Reports 30
Accessing Reporting Engine Log Files 32
Configure Task Scheduler for a Reporting Engine 32
Specify the Pools and Queues 33
Define Reports, Charts and Alerts **34**
Define Reports 34
Define Charts 34
Define Alerts 34
Configure Reporting Engine General Settings **36**
Access the General Tab 36
Troubleshooting Reporting Engine Configuration **37**
References **38**
Reporting Engine General Tab 39
System Configuration 40
Logging Configuration 43
Warehouse Kerberos Configuration 44
Reporting Engine Sources Tab 45
Reporting Engine Output Actions Tab 48
NetWitness Configuration 50
SMTP 51
SNMP 52

Syslog 53
SFTP 55
URL 56
Network Share 57
Reporting Engine Manage Logos Tab 59

How Reporting Engine Works

NetWitness Reporting Engine is a service on the NetWitness Admin Server. It facilitates the data extraction from different data sources to generate reports for compliance and analysis. Reporting Engine stores the definitions of the charts, rules, reports and alerts that are used to generate reports, charts and alerts.

Reporting Engine configuration includes configuring the data sources, definitions of outputs or notifications and parameters to improve the performance of data extraction and report, chart, and alert generation.

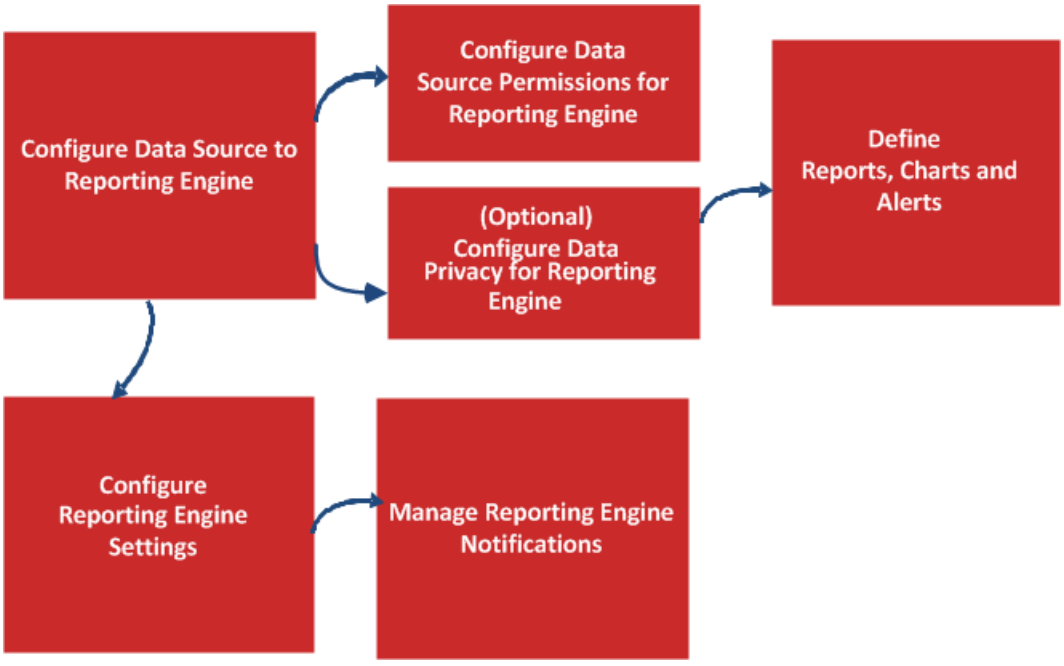
When you install the NetWitness, Reporting Engine is automatically installed as a service. This enables the Reports, Charts, and Alerts to be maintained in the NetWitness and be available to view, download reports as PDF or CSV format, download charts as PDF and be added as dashlets.

For the Reporting Engine to run reports and alerts based on the data drawn from a data source, you must associate a data source, or multiple data sources to a Reporting Engine. There are three types of data sources:

- NWDB - The NetWitness Database (NWDB) data sources are Decoders, Log Decoders, Brokers, Concentrators, Archiver, and Collection. Reporting Engine supports the generation of reports, alerts, and charts.
- Warehouse - The Warehouse data sources are Horton Works and MapR which collects information from the Warehouse Connector and generates reports and alerts. This data source generates Reports only.
- Respond - Respond is used to generate reports on alerts and incidents. This data source generates Reports only.

Workflow

The following workflow shows an overview of the Reporting Engine configuration which enables the user to generate Reports, Charts, and Alerts.



Configure the Reporting Engine

On installation of the NetWitness Server, the Reporting Engine service is automatically available and some parameters are pre-populated with default values to achieve optimal results.

Make sure that the data sources are deployed and configured in the NetWitness. For more information, see "Add Service or Edit Service Dialog" topic in the *Host and Service Configuration Guide*.

You can perform the following tasks:

- Check Live for the latest data source content and deploy it on a regular basis. (For more information, see "Manage Live Resources" topic in the in *Live Services Guide*).
- (Optional) Add Additional Space for Large Reports.

Configure the Data Sources

You must configure NWDB, Warehouse, and Respond to generate Reports, Charts, and Alerts. Optionally, you can also configure Archiver, Collection, and Workbench data sources.

IMPORTANT: If you change the admin password on a NetWitness service that is used as a Reporting Engine data source, you must remove and then re-add the service as a data source.

Note: To execute Reports and Charts on an Analyst UI, make sure the admin adds the data sources to each Reporting Engine instance from the admin node using the relevant procedure described in this topic.

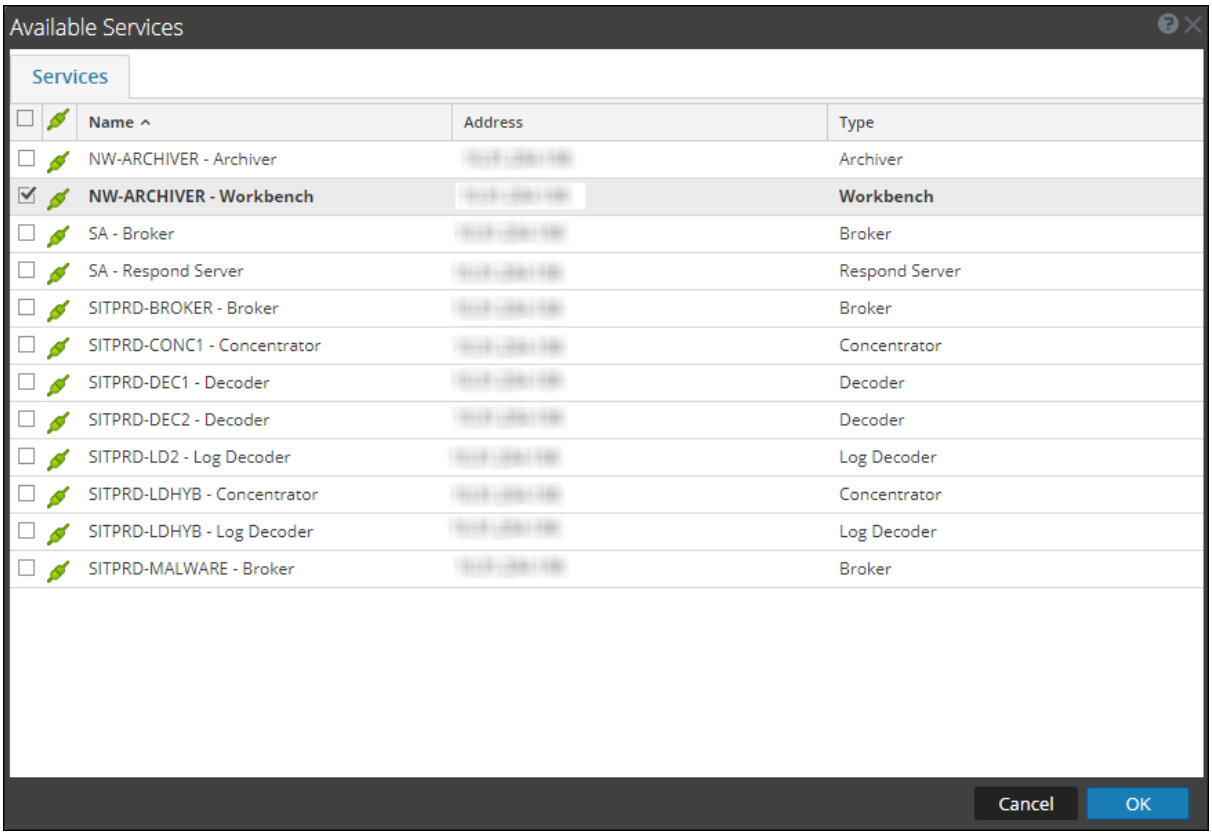
Configure a NWDB Data Source

To add a NWDB data source:

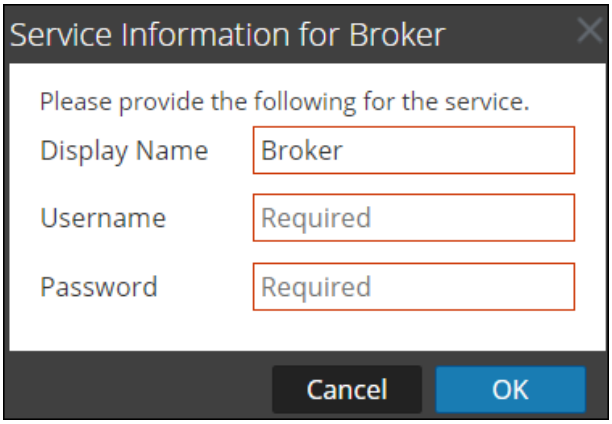
1. Go to  **(Admin) > Services**.
2. In the **Services**, select **Reporting Engine** service.
3. Click  > **View** > **Config**

   The Services Config View of Reporting Engine is displayed.
4. On the **Sources** tab, click  > **Available Services**.

   The **Available Services** dialog is displayed.



5. Select a NWDB service you want to add and click **OK**.
6. In the Service Information for Broker dialog, enter the service information for the service and click **OK**. In this example, we are adding a Broker service.



7. The service is displayed in the Sources tab when it is successfully added.

Note: The services with the Trust Model enabled must be added individually. You are prompted to provide a username and password for the selected service.

Configure a Warehouse Data Source

You can add the warehouse data source to Reporting Engine, so that you can extract the data from the required services, store them in MapR or Horton works and generate Reports and Alerts. The procedure to configure Warehouse as a data source differs. To extract data from a Warehouse data source, you must configure it using the following procedure.

Note: Warehouse Connector is still supported, as it is reporting against a warehouse in NetWitness 11.x.

Prerequisite

Make sure you:

- Add a Warehouse Data Source to Reporting Engine.
- Set Warehouse Data Source as the Default Source.
- HIVE server is in running state on all the Warehouse nodes. Use the following command to check the status of the HIVE server:

  ```
  status hive2 (MapR deployments)
  service hive-server2 status (Horton Works deployments)
  ```

- Warehouse Connector is configured to write data to the warehouse deployments.
- If Kerberos authentication is enabled for HiveServer2, make sure that the keytab file is copied to the `/var/netwitness/re-server/rsa/soc/reporting-engine/conf/` directory in the Reporting Engine Host.

  Note: The `rsasoc` user should have read permissions for the keytab file. For more information, see Configure Data Source Permissions.

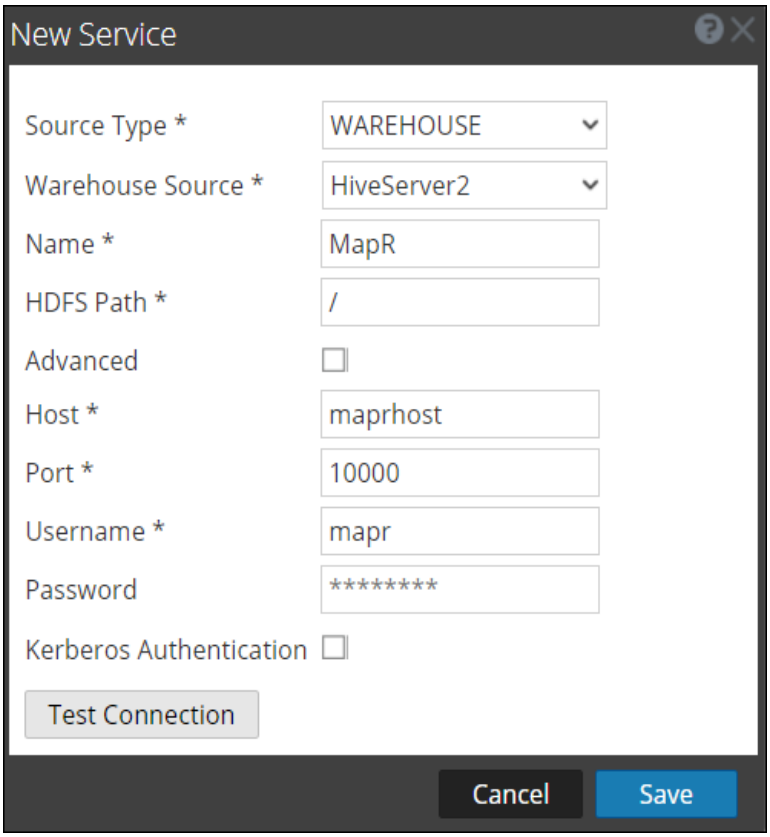
  Also, make sure that you update the keytab file location in the **Kerberos Keytab File** parameter in the Reporting Engine Service Config View. For more information, see Reporting Engine General Tab.

To add Warehouse data source for MapR:

1. Go to  **(Admin) > Services**.
2. In the **Services** list, select the **Reporting Engine** service.
3. Click  > **View** > **Config**.

   The **Service Config** view is displayed with the **General** tab open.
4. Click the **Sources** tab.
5. In the **Sources** tab, click  and select **New Service**.

   The New Service dialog is displayed.



6. In the **Source Type** drop-down menu, select **WAREHOUSE**.
7. In the **Warehouse Source** drop-down menu, select the warehouse data source.
8. In the **Name** field, enter the host name of the Warehouse data source.
9. In the **HDFS Path** field, enter the HDFS root path to which the Warehouse Connector writes the data.

   For example:

   If `/saw` is the local mount point for HDFS that you have configured while mounting NFS on the device. And if you have installed the Warehouse Connector service to write to SAW. For more information, see "Mount the Warehouse on the Warehouse Connector" topic in the *Warehouse (MapR) Configuration Guide*.

   If you have created a directory named `Ionsaw01` under `/saw` and provided the corresponding Local Mount Path as `/saw/Ionsaw01`, then the corresponding HDFS root path would be `/Ionsaw01`.

   The `/saw` mount point implies to `/`as the root path for HDFS. The Warehouse Connector writes the data `/Ionsaw01` in HDFS. If there is no data available in this path, the following error is displayed:

   `"No data available. Check HDFS path"`

   Make sure that `/lonsaw01/rsasoc/v1/sessions/meta` contains avro files of the meta data before performing test connection.
10. Select the **Advanced** checkbox to use the advanced settings, and fill in the **Database URL** with the complete JDBC URL to connect to the HiveServer2.

    For example:

If kerberos is enabled in HIVE then the JDBC url will be:

```
jdbc:hive2://<host>:<port>/<db>;principal=<Kerberos serverprincipal>
```

If SSL is enabled in HIVE then the JDBC url will be:

```
jdbc:hive2://<host>:<port>/<db>;ssl=true;sslTrustStore=<trust_store_
path>;trustStorePassword=<trust_store_password>
```

For more information on HIVE server clients, see https://cwiki.apache.org/confluence/display/Hive/HiveServer2+Clients.

11. If not using the advanced settings, enter the values for the **Host** and **Port**.
    - In the **Host** field, enter the IP address of the host on which HiveServer2 is hosted.

      Note: You can use the virtual IP address of MapR only if HiveServer2 is running on all the nodes in the cluster.

    - In the **Port** field, enter the HiveServer2 port of the Warehouse data source. By default, the port number is **10000**.
12. In the **Username** and **Password** field, enter the JDBC credentials used to access HiveServer2.

    Note: You can also use LDAP mode of authentication using Active Directory. For instructions to enable LDAP authentication mode, see Enable LDAP Authentication.
13. Enable Kerberos authentication: see Enable Kerberos Authentication.
14. If you want set the added Warehouse data source as default source for the Reporting Engine, select the added Warehouse data source and click Set default.

To add Warehouse data source for Horton Works (HDP):

Note: Make sure you download the `hive-jdbc-1.2.1-with-full-dependencies.jar`. This jar contains the driver file of HIVE 1.2.1 which connects to Reporting Engine for Hive 1.2.1 Hiveserver2.

1. SSH to the NetWitness server.
2. In the `/opt/rsa/soc/reporting-engine/plugins/` folder, take a backup of the following jar:

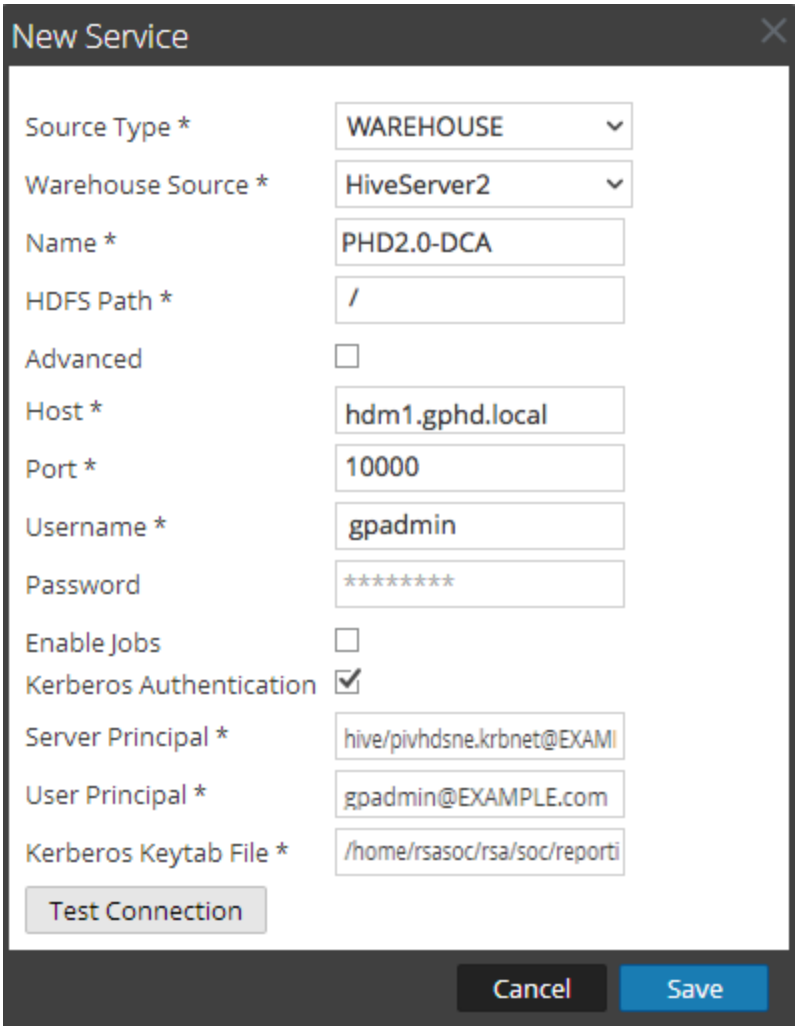
   `hive-jdbc-0.12.0-with-full-dependencies.jar` or `hive-jdbc-1.0.0-mapr-1508-standalone.jar`
3. Remove the following jar:

   `hive-jdbc-0.12.0-with-full-dependencies.jar` or `hive-jdbc-1.0.0-mapr-1508-standalone.jar`
4. In the /opt/rsa/soc/reporting-engine/plugins folder, copy the following jar using WinSCP:

   `hive-jdbc-1.2.1-with-full-dependencies.jar`
5. Restart the Reporting Engine service.
6. Log in to NetWitness UI.
7. Select the **Reporting Engine** service and select > **View** > **Explore**.
8. In the **hiveConfig**, set **EnableSmallSplitBasedSchemaLiteralCreation** parameter to **true**.

Enable Kerberos Authentication

1. Select **Kerberos Authentication** checkbox, if the Warehouse is Kerberos enabled HIVE server.



2. Fill in the fields as follows:

Field	Description
Server Principal	Enter the Principal used by the HIVE server to authenticate with the Kerberos Key Distribution Center (KDC) Server.
User Principal	Enter the Principal that HIVE JDBC client uses to authenticate with the KDC server for connecting the HIVE server. For example, gpadmin@EXAMPLE.COM.

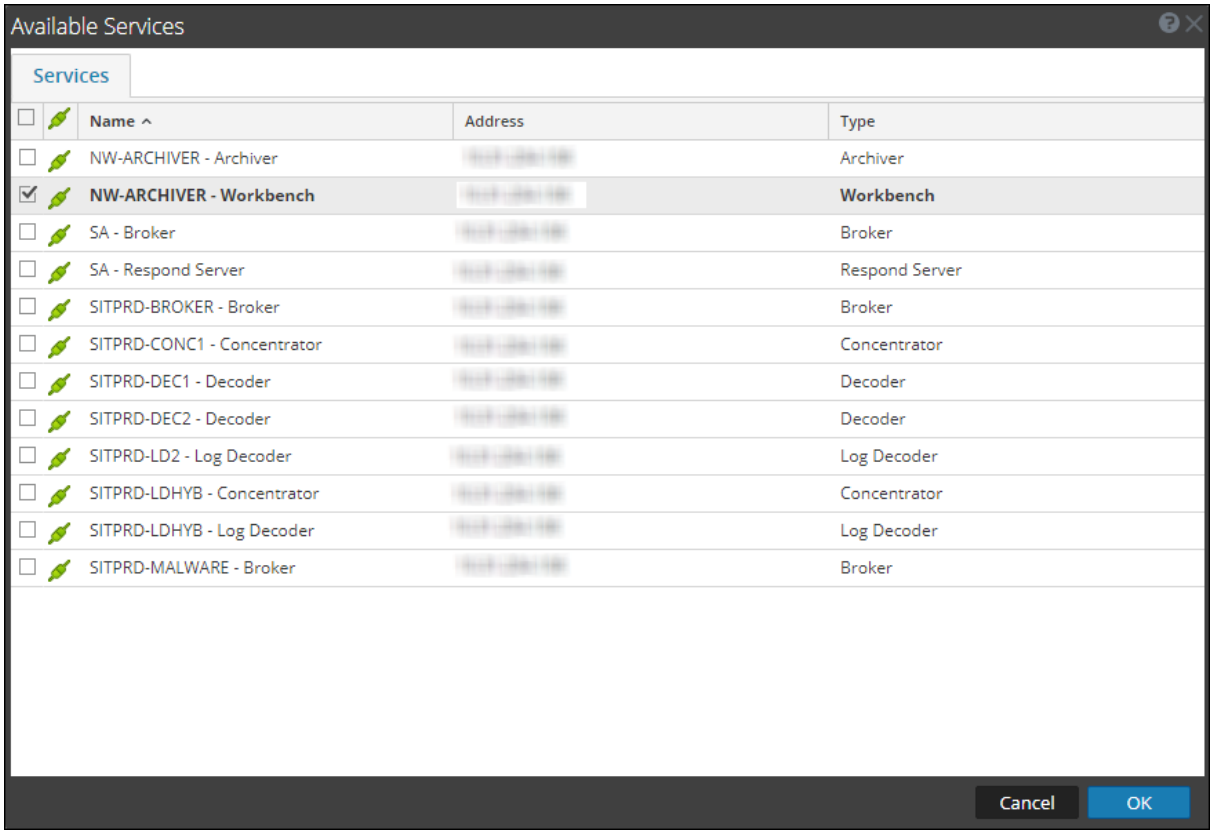
Field	Description
Kerberos Keytab File	View the Kerberos keytab file location configured in the HIVE Configuration panel on the Reporting Engine General Tab. **Note:** Reporting Engine supports only the data sources configured with the same Kerberos credentials, like, User Principal and key tab file.

3. Click **Test Connection** to test the connection with the values entered.
4. Click **Save**.

   The added Warehouse data source is displayed in the Reporting Engine Sources tab.

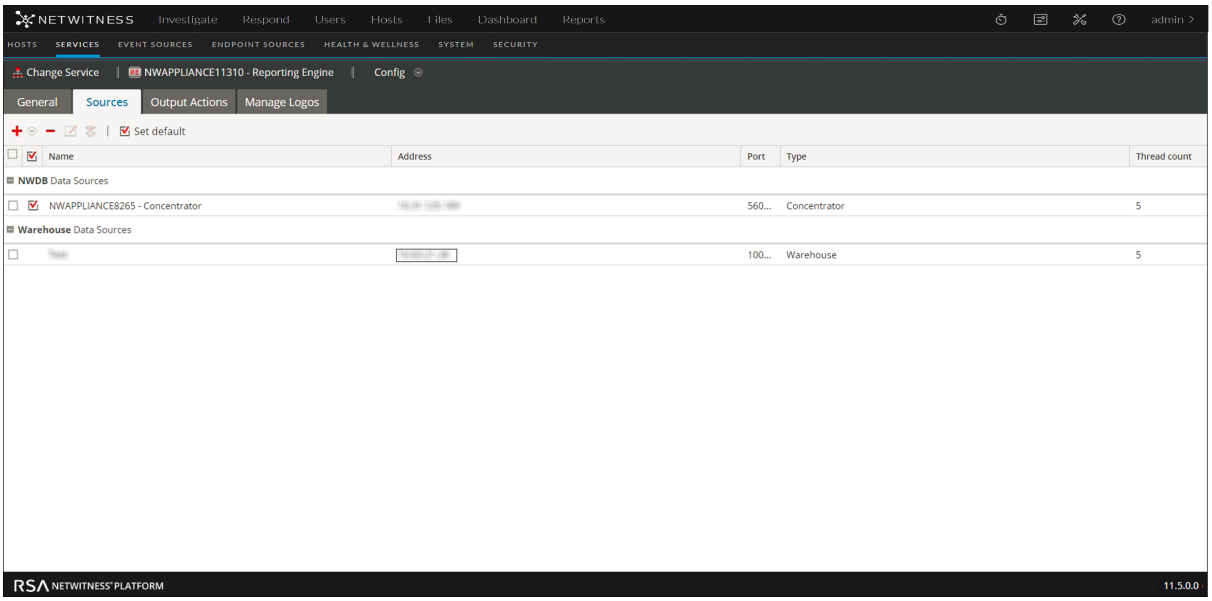
5. Click  > **Available Services**.

   The Available Services dialog box is displayed.



6. In the Available Services dialog box, select the service that you want to add as data source to the Reporting Engine and click **OK**.

   NetWitness adds this as a data source available to reports and alerts against this Reporting Engine.



Note: This step is relevant only for an Untrusted model.

Set a Data Source as the Default Source

To set a data source to be the default source when you create reports and alerts:

1. Go to  **(Admin) > Services**.
2. In the **Services** list, select a **Reporting Engine** service.
3. Select  > **View**> **Config**.

   The Services Config View of Reporting Engine is displayed.
4. Select the **Sources** tab.

   The **Services Config View** is displayed with the Reporting Engine Sources tab open.
5. Select the source that you want to be the default source (for example, Broker).
6. Click the  checkbox.

   NetWitness defaults to this data source when you create reports and alerts against this Reporting Engine.

(Optional) Add Workbench as Data Source

You must configure Workbench,to be able to use data from Workbench data source to generate Reports and Alerts. The following instructions describe how to add Workbench service as a data source to Reporting Engine to generate report for the data collected by Workbench.

Prerequisites

Make sure you have:

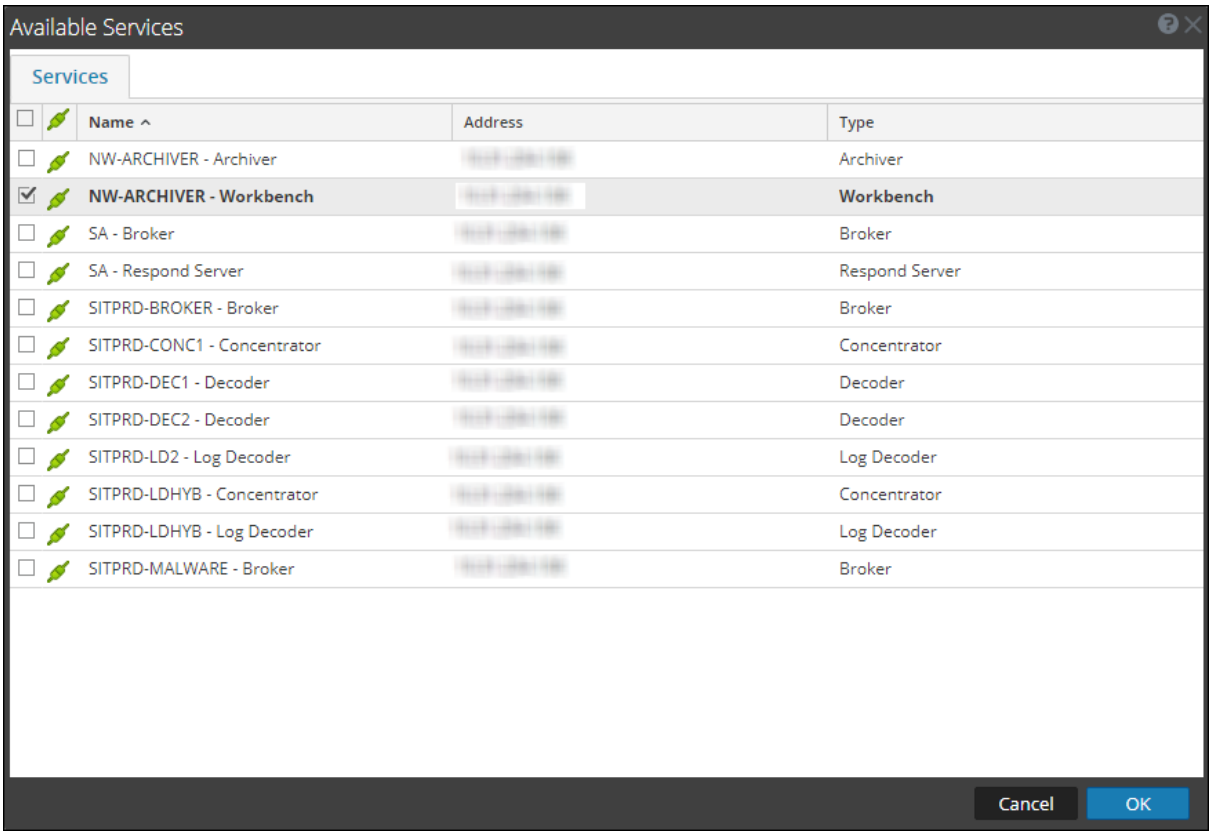
1. Added Workbench as a service to your NetWitness deployment. For more information, see the *Archiver Configuration Guide*.
2. Added a Collection on the Workbench service.

To add Workbench as a data source to Reporting Engine:

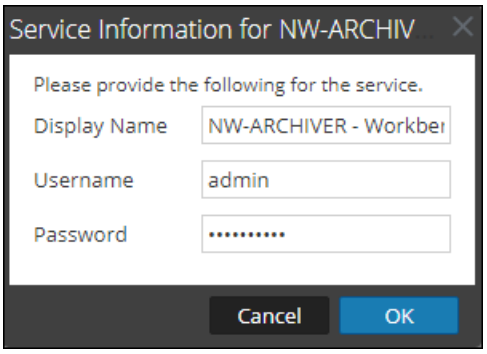
1. Go to  **(Admin) > Services**.
2. In the **Services** list, select a **Reporting Engine** service.
3. Select  > **View** > **Config**.

   The Services Config View of Reporting Engine is displayed.
4. Select the **Sources** tab.
5. Click  and select **Available Services**.

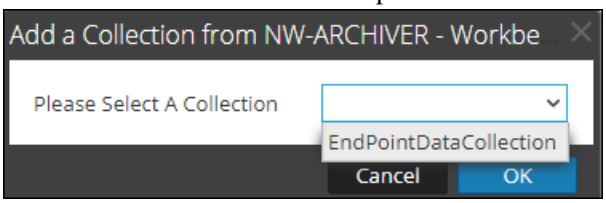
   The Available Services dialog is displayed:



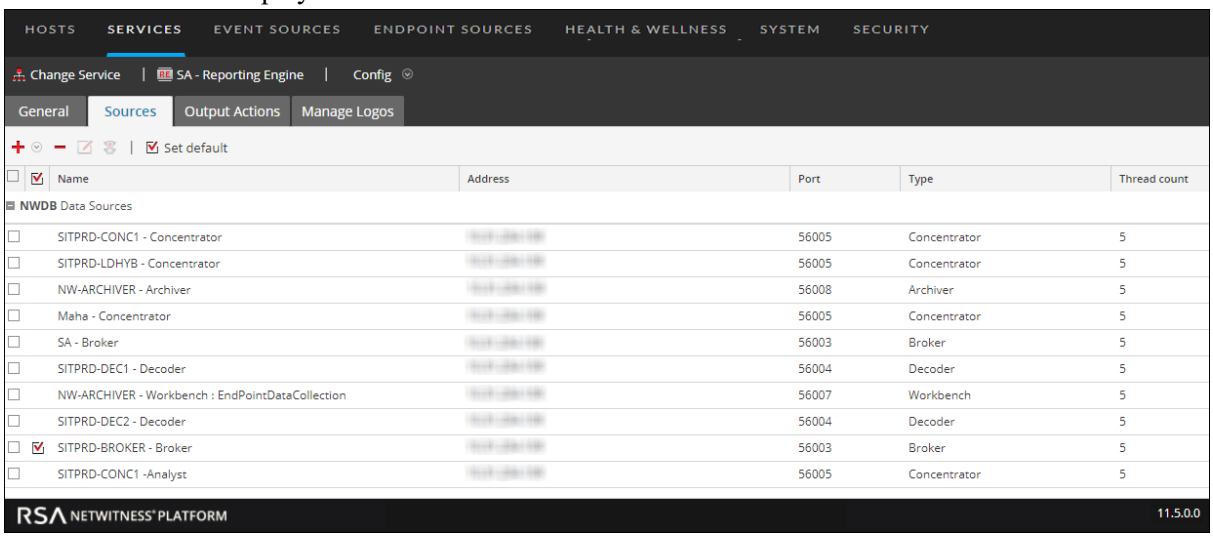
6. Select the Workbench service and click **OK**.
   A list of collections are displayed.
7. Enter the service information, and click **OK**



8. Select a collection from the dropdown.



9. The data source is displayed in the Sources tab.



The workbench service is now added as a data source to the Reporting Engine.

Note: The services with the Trust Model enabled must be added individually. You are prompted to provide a username and password for the selected service.

(Optional) Add Archiver as Data Source

You must configure Archiver, to be able to use data from Archiver data source to generate Reports and Alerts:

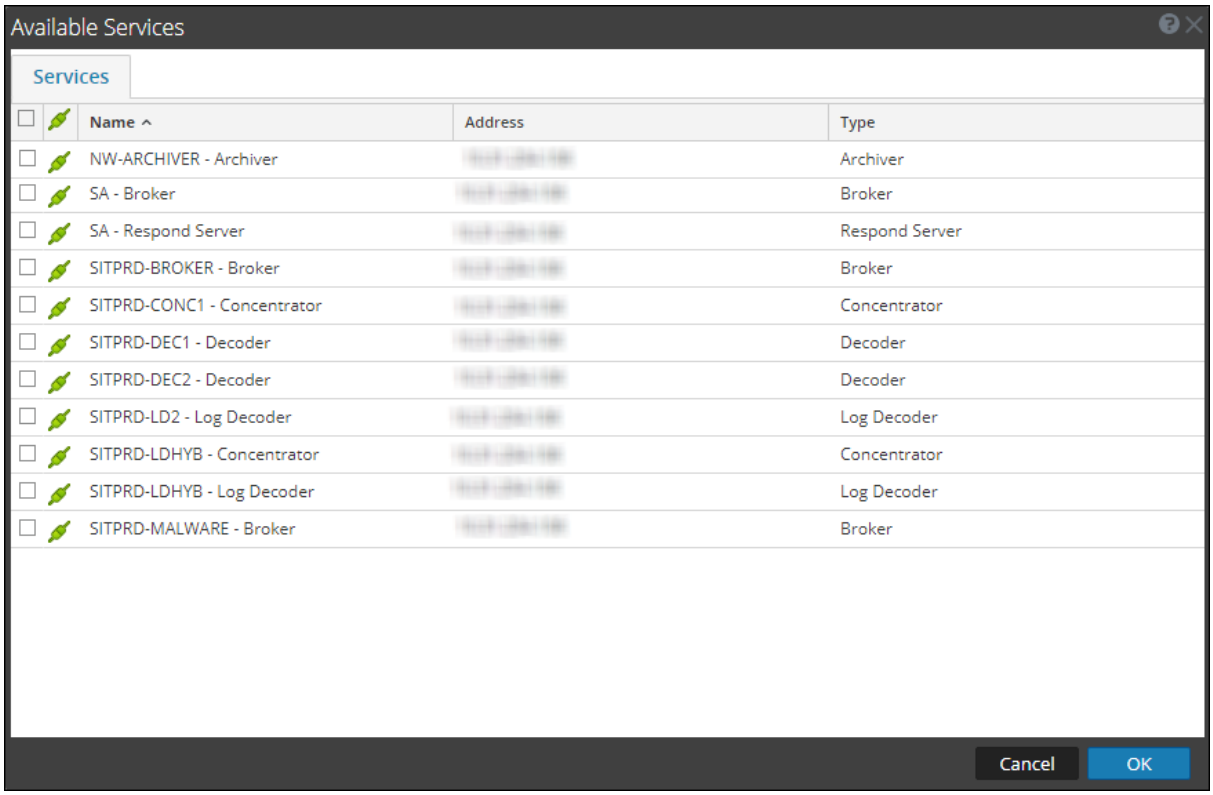
Prerequisites

Make sure that you have:

1. Installed the NetWitness Archiver host in your network environment. For more information, see the *Hosts and Services Getting Started Guide*.
2. Installed and configured Log Decoder in your network environment. For more information, see "Add Log Decoder as a Data Source to Archiver" in the *Archiver Configuration Guide*.
3. Reporting Engine service is available in your NetWitness deployment.
4. Added Archiver service to your NetWitness deployment. For more information, see "Add the Archiver Service" in the *Archiver Configuration Guide*.
5. Applied license to the Archiver service.

To add Archiver Data Source to Reporting Engine:

1. Go to  **(Admin) > Services**.
2. In the **Services** list, select the **Reporting Engine** service.
3. Click  > **View** > **Config**.
   The Services Config View of Reporting Engine is displayed.
4. Select the **Sources** tab.
5. Click  and select **Available Services**.
   The Available Services dialog is displayed.



6. Select the Archiver service and click **OK**.

   The service authentication dialog box is displayed.

   Note: The services with the Trust Model enabled must be added individually. You are prompted to provide a username and password for the selected service.

7. Enter the Username and Password for the Archiver.
8. Click **OK**.

   The selected Archiver is listed in the Aggregate Services pane.

(Optional) Integrate Endpoint Information Into Reports

You can use the Endpoint data by using the following instructions to add the Endpoint information into Reports. The *Endpoint Integration Guide* provides an overview of Endpoint integration into NetWitness.

Prerequisites

Make sure that:

- You have configured the Endpoint alerts via syslog into a Log Decoder. For more information see, "Configure Endpoint Alerts Via Syslog into a Log Decoder" topic in *Endpoint Integration Guide*).

To integrate Endpoint information into Reports:

1. In **Reporting Engine**> **View**> **Config**> **Sources**.
2. Add the Concentrator that is consuming data from the Log Decoder as a data source.
   Endpoint meta is populated in Reporting Engine.
3. Run reports by selecting the appropriate meta.

(Optional) Add Collection as Data Source to Reporting Engine

You must configure Collection, to be able to use data from Collection data source to generate Reports, Charts, and Alerts:

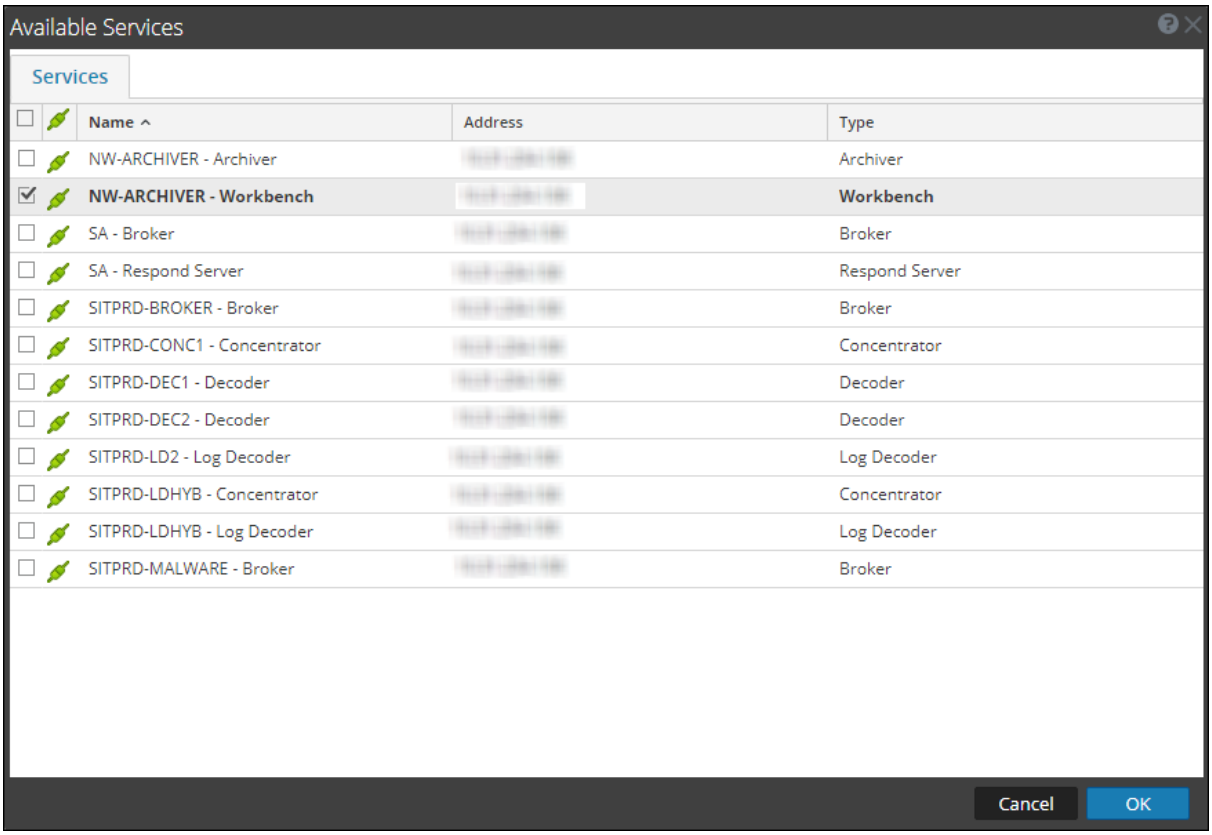
Prerequisites

Make sure that you have:

- Installed a Workbench service on a Reporting Engine host.
- Backed up data in a known location on your local host, if you are adding a collection using the data restored from the backed up data.

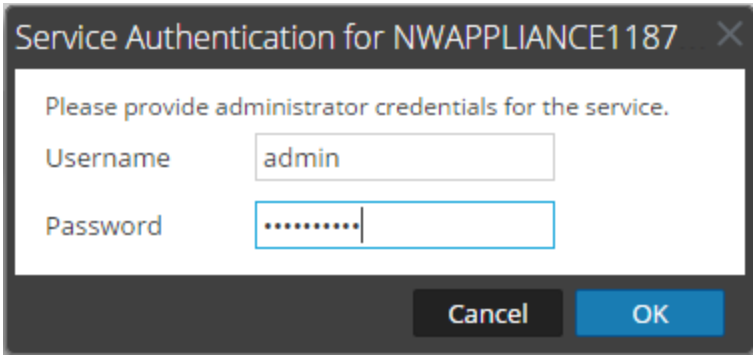
To associate a Collection as a data source with Reporting Engine:

1. Go to  **(Admin) > Services**.
2. In the **Services** list , select a **Reporting Engine** service.
3. Click  > **View** > **Config**.
   The Services Config View of Reporting Engine is displayed.
4. Select the **Sources** tab.
5. Click  and select **Available Services**.
   The Available Services dialog is displayed.



6. Select the Workbench service and click **OK**.

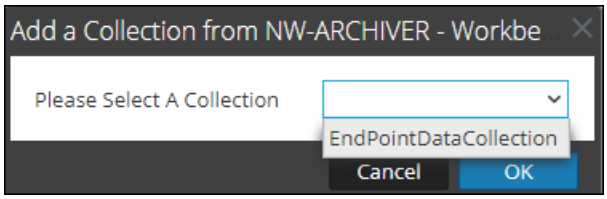
   The Service Authentication dialog for the selected service is displayed.



Note: The services with the Trust Model enabled must be added individually. You are prompted to provide a username and password for the selected service.

7. Enter the username and password for admin credentials for the service.
8. Click **OK**.

   The add collection dialog is displayed.



9. Select a collection from the drop-down list and click **OK**.

   The workbench service is now added as a data source to the Reporting Engine.

Configure Data Privacy for the Reporting Engine

You can configure the data privacy for all data sources of Reporting Engine using the Sources tab of the **Services** > **View** > **Config** view.

With the addition of the Data Privacy feature to NetWitness 11.0 and above, access to sensitive meta in NetWitness Core services can be restricted by configuring separate data sources for Data Privacy Officer (DPO) users and non-DPO users, and limiting access to those data sources by assigning appropriate permissions.

In the **Services** > **Config** view, you can add each Core service as two separate data sources: one with a service account having privileges equivalent to a DPO and the other with a service account having privileges equivalent to any other user. Then, to limit access to those data sources based on roles, you can assign read access or no access to those data sources for individual roles. To limit access to Warehouse data sources, you can do the same.or more information, see Configure Data Source Permissions.

Note: A user assigned to the `Data_Privacy_Officers` role (or an equivalent custom role), can create a report, chart and alert. Also, configure a report or alert output actions in the Reporting module. In an environment where data privacy features of NetWitness are enabled and one or more meta keys are configured as protected, these actions can result in the following:
- When an alert is created by a DPO user, any protected or sensitive meta involved in the alert is automatically available in Respond. This may inadvertently provide all the users of Respond module access to the sensitive meta values, regardless of their roles. One option to prevent this is to disable publishing into Respond from Reporting.
- When an Output Action is configured by a DPO user, either sensitive meta values, reports with sensitive meta values or both, may become available to target users or destinations of that Output Action, regardless of the role assigned to the target user.

It is strongly recommended that DPO users completely avoid creating alerts or configuring output actions for a report or alert in the Reporting module. If they do such configuration, the above implications must be carefully considered.

NetWitness Core services (for example, Concentrator, Broker, or Archiver) support the ability to restrict meta data based on the configured user role. To make use of the data privacy feature for Reporting Engine, you can configure two separate service accounts against Core services. One service account for general purpose reporting that does not include any sensitive data and the other account for privileged users with access to all data including sensitive data. The access to restricted meta data for the two service accounts is configured as part of the data privacy plan on each Core service.

In Reporting Engine, you can add each Core service as two separate data sources (one being the regular data source and the other a privileged data source) using the two separate service accounts. You can configure Reporting Engine to allow only users with privileged roles to access the sensitive data source. Hence, Reporting Engine can connect to a NWDB Data source in two ways:

- Using a service account with DPO role.
- Using a service account without a DPO role.

Note: You can also add two or multiple data sources for the same Core service.

After adding two data sources with different service accounts for the same Core service, you can configure data source permissions to manage access to these data sources. For more information, see Configure Data Source Permissions.

Note: If the content is changed to utilize the transformed meta key, the hash value of the original meta is displayed in its place when viewing reports, charts and alerts.

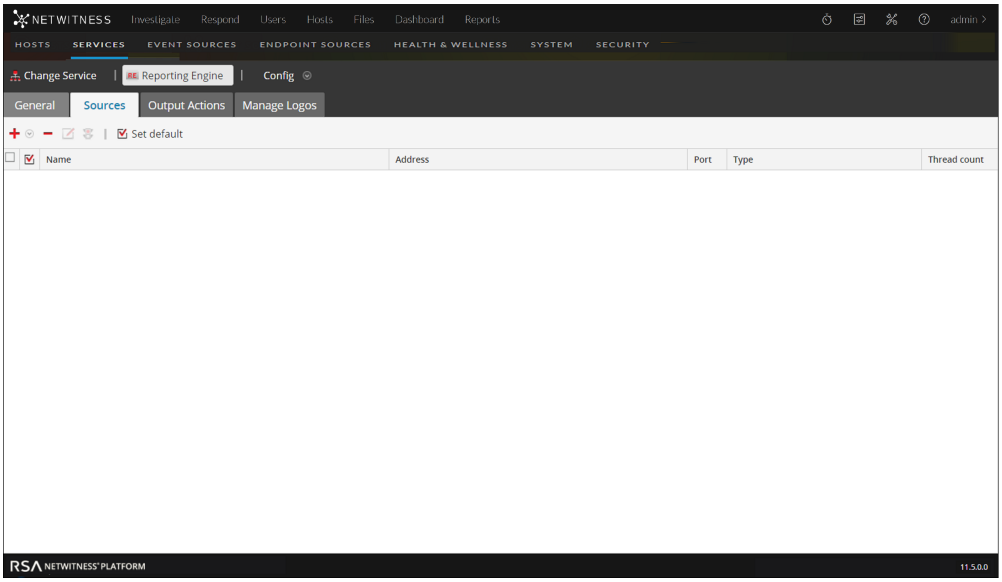
Add a NWDB Data Source with Different Service Accounts

To add a NWDB data source:

1. Go to **(Admin) > Services**.
2. In the **Services** list, select a **Reporting Engine** service.
3. Click **View** > **Config**.

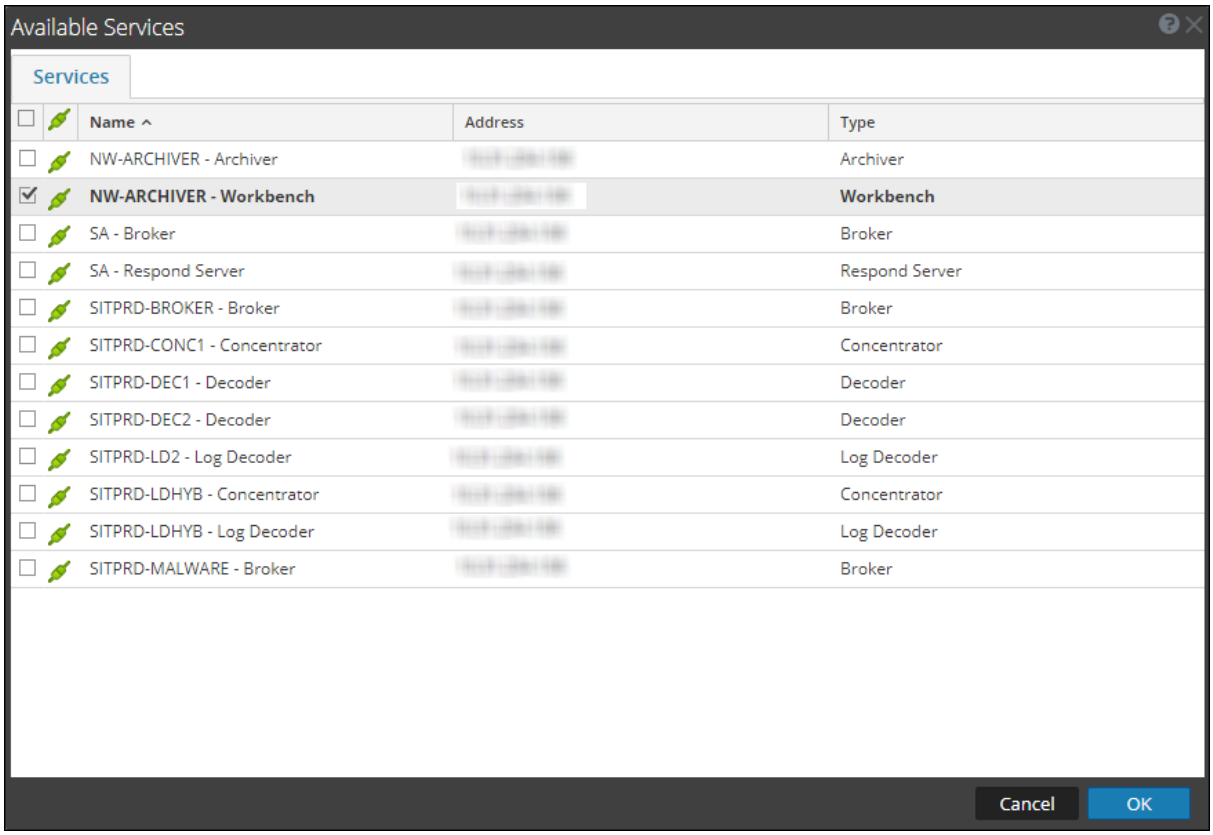
   The Services Config view of Reporting Engine is displayed.
4. Select the **Sources** tab.

   The Services Config View is displayed.



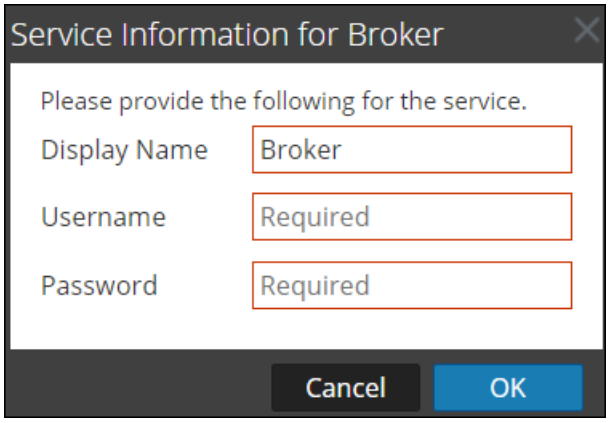
5. Click and select Available Services.

   The Available Services dialog is displayed. All services are listed, including those that have already been added to the Reporting Engine.



6. Select the checkbox next to the service and click **OK**.

   The Service Information dialog for the selected service is displayed.

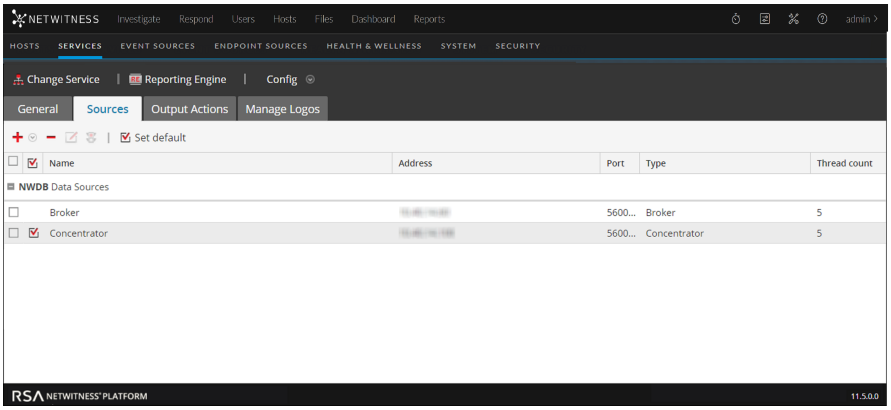


Note: NetWitness prompts you to provide a username and password for the selected service. To limit access to sensitive data, DPO users must use their credentials while adding the source instead of using the admin credentials. These credentials need to be applied to the host even if using trusted connections between the NetWitness server and NetWitness Core hosts.

   Repeat the step for Non-DPO data source.

7. Enter the username and password for the required service account.
8. Click **OK**.

The required service is added as a data source to the Reporting Engine. Two data sources are added to Reporting Engine for the same Core device.



Configure Data Source Permissions

You can configure data source permissions for Reporting Engine using the **Sources** tab of the **Services** > **Config** view. This helps manage access control to the data sources by setting the data source permissions. Now, with the ability to add more than one data source for the same Core service, you can configure different permissions for each data source of the same Core service. For example, data privacy officers (DPO) can create a Warehouse source using their credentials, and that allows them to execute reports against the Warehouse while restricting everyone else from being able to use that source.

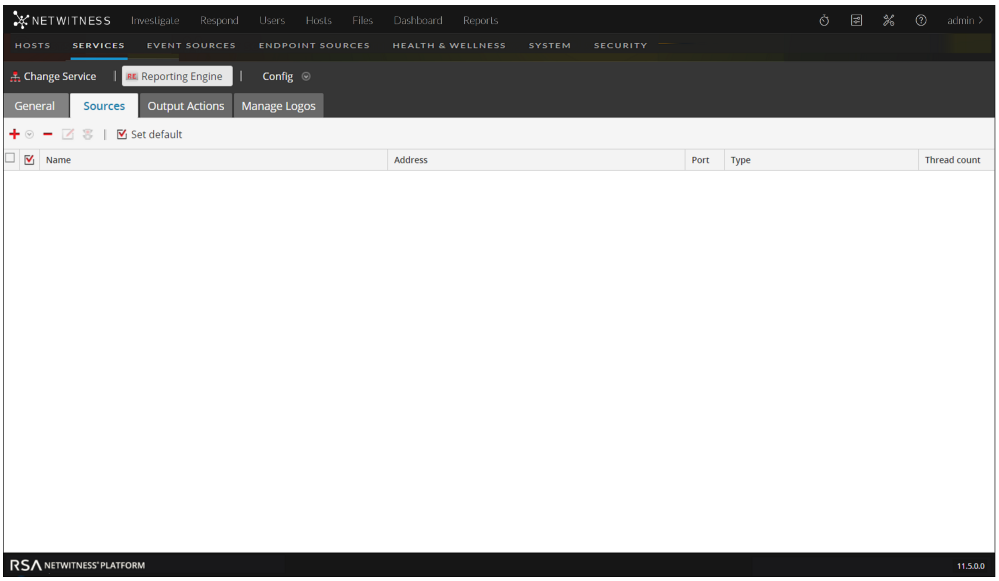
Note: In 11.x, the permissions for NWDB and Warehouse data sources are automatically set based on the permissions of the reporting objects. For example, if the role had the permissions set as **Read Only**/**Read & Write** for any reporting object in 10.5, then that role is automatically assigned read only permission for all the data sources that existed in 10.5. If no permission is set for the role, then the data source permission is automatically set to No Access.

To configure permissions to data sources:

1. Go to  **(Admin) > Services**.
2. In the **Services** list, select a **Reporting Engine** service.
3. Click  > **View** > **Config**.

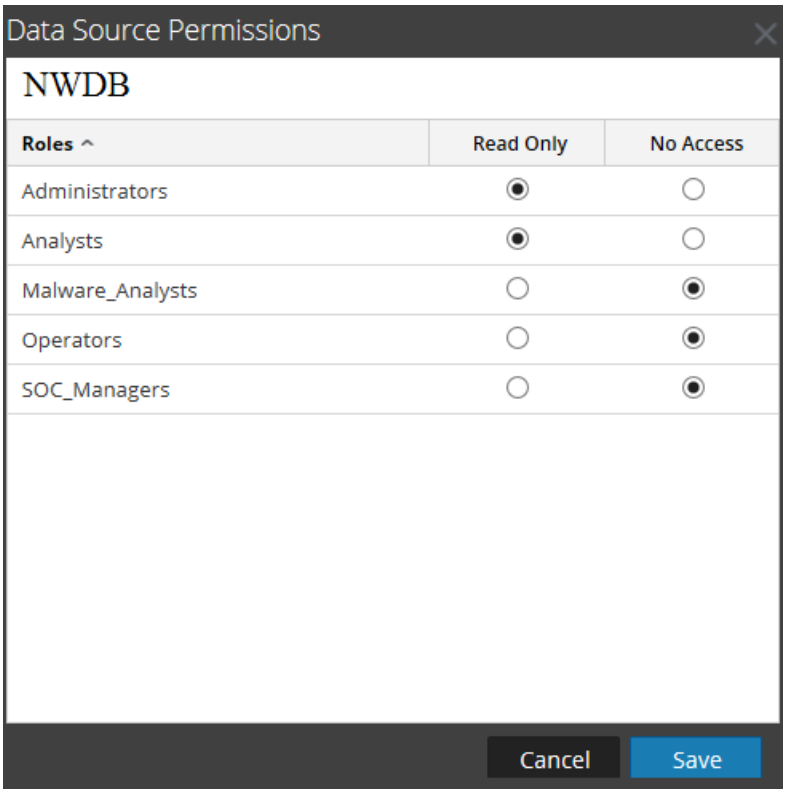
   The Services Config view of Reporting Engine is displayed.
4. Select the **Sources** tab.

   The Service Config View displays the Sources tab.



5. Select the data source for which you want to configure permissions by selecting the checkbox.
6. Click .

   The Data Source Permissions dialog is displayed.



7. Modify the access permission for different users based on the type of service account of the data source. The permission can be either **Read Only** or **No Access**.
8. Click **Save**.

   The required permissions are configured for the data source. For more information, see the *Reporting Guide*.

Configure Reporting Engine Settings

After you configure the Reporting Engine and required data sources based on your requirements, you can modify some of the configurations to customize your Reports, Charts, and Alerts.

To configure the settings:

1. Go to  **(Admin) > Services**.
2. In the **Services** list, select a **Reporting Engine** service.
3. Click  > **View** > **Config**.

   The Services Config View of Reporting Engine is displayed with the General tab highlighted. For more information on Reporting Engine General tab, see Reporting Engine General Tab.
4. Edit the Reporting Engine service settings and click **Apply**.

   The service settings are configured on Reporting Engine.

Enable LDAP Authentication

To enable LDAP mode of authentication using Active Directory for HiveServer2 for Warehouse data source, follow these steps.

1. Log on to the Analytics Warehouse appliance as root user.
2. Navigate to `/opt/mapr/hive/hive-0.11/conf.new/` directory. Type the following command and press ENTER:

   ```
   cd /opt/mapr/hive/hive-0.11/conf.new/
   ```
3. Edit the file `hive-site.xml`. Type the following command and press ENTER:

   ```
   vi hive-site.xml
   ```
4. Add the following properties under `<Configuration>` tag:

   ```
   <property> <name>hive.server2.authentication</name> <value>LDAP</value>
   </property> <property> <name>hive.server2.authentication.ldap.url</name>
   <value>LDAP_URL</value> </property>
   ```

   Where `LDAP_URL` is the URL of the LDAP Server.
5. Restart HiveServer2.

Add Additional Space for Large Reports

To add additional disk space to the Reporting Engine for large reports, follow the below steps. If large compliance reports have to be generated for Warehouse, the Reporting Engine disk space might get consumed quicker than expected. In such cases, you can mount any external storage such as SAN or NAS for storing reports.

The directories that tend to fill up disk space are `resultstore` and `formattedReports` under the Reporting Engine home directory. It is recommended to move only these two directories to SAN or NAS and replace the original locations with soft links pointing to the new locations. It is also recommended to leave the remaining directories in the local disk itself for reliable and high I/O performance.

Note: The following steps assume that the Reporting Engine home directory is located at `/var/netwitness/re-server/rsa/soc/reporting-engine/` and the external storage is mounted under `/externalStorage/`. Also, the 'rsasoc' user must have read-write access to the specified external storage path.

To move disk space for the Reporting Engine to external storage:

1. Stop Reporting Engine service as a root user.

   ```
   service rsasoc_re stop
   ```

2. Switch to `rsasoc` user.

   ```
   su rsasoc
   ```

3. Change to RE home directory.

   ```
   cd /var/netwitness/re-server/rsa/soc/reporting-engine/
   ```

4. Move the resultstore directory to a mounted external storage. Type the following command and press ENTER:

   ```
   mv resultstore /externalStorage
   ```

5. Move the formatted Reports directory to a mounted external storage. Type the following command and press ENTER:

   ```
   mv formattedReports /externalStorage
   ```

6. Create a soft link for resultstore. Type the following command and press ENTER:

   ```
   ln -s /externalStorage/resultstore /var/netwitness/re-
   server/rsa/soc/reporting-engine/resultstore
   ```

7. Create a softlink for formattedReports. Type the following command and press ENTER:

   ```
   ln -s /externalStorage/formattedReports /var/netwitness/re-
   server/rsa/soc/reporting-engine/formattedReports
   ```

8. Exit the `rsasoc` user.

   ```
   exit
   ```

9. Start Reporting Engine service as a root user.

   ```
   service rsasoc_re start
   ```

   Note: If the external storage is offline, you cannot perform the following tasks:
   1) Execute Reports or Reporting Alerts
   2) View existing Reports or Reporting Alerts
   However, you can create new Reporting objects such as Reports and Charts, and access Charts and Live Dashboard created for charts. Therefore, you must ensure that the external storage is reliable and has the required space.

   Additionally, if you want to store reports beyond 100 days, change the retention configuration appropriately for the service that you are using as a data source.

Accessing Reporting Engine Log Files

You can access the Reporting Engine log files which are stored in the following logs directory `/var/netwitness/re-server/rsa/soc/reporting-engine/logs/`

- Current logs in the `reporting-engine.log` file.
- Backup copies of previous logs in the `reporting-engine.log.*` file.
- All UNIX script logs in the files that have the following syntax: `reporting-engine.sh_timestamp.log` (for example, `reporting-engine.sh_20120921.log`)

The Reporting Engine rarely writes command line error messages to the `rsasoc/nohup.out` file.

The Reporting Engine appends the log messages and output written by systemd system and the commands used to start the reporting-engine to the directory `/var/log/messages`.
A `/var/log/messages` log file is a system log file so only the root user can read it.

Configure Task Scheduler for a Reporting Engine

You can configure queues and pools in the Reporting Engine to schedule NWDB or Warehouse reports. For more information on Task Schedulers, see "Task Scheduler for Warehouse Reporting" in the *Reporting Guide*

Prerequisites

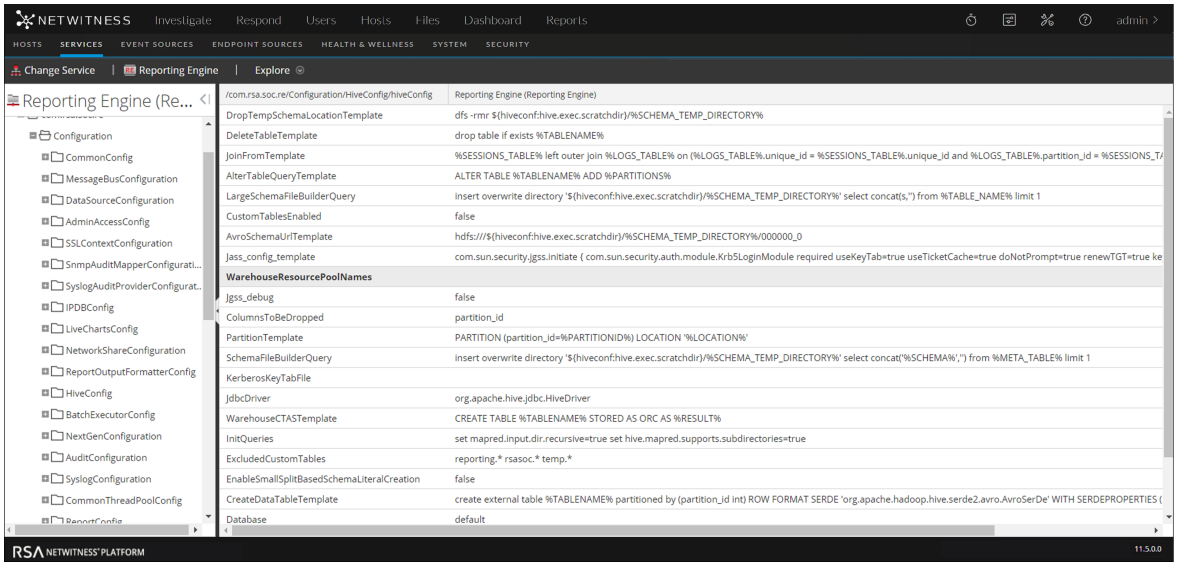
Make sure that you have identified the following:

- Scheduler type and pools or queues you want to use. You can configure only one scheduler for the Reporting Engine. By default the Fair Scheduler is configured.
- Names of the queues or pools, and the resources given to each queue and pool.
- NetWitness does not support multiple queues or pools per cluster. Netwitness recommends that you either provide unique names to queues or pools in all the clusters or use the same queue or pool names in both the clusters. If cluster size is large, there may be more than 3 pools or queues.
- If you are using an unsupported scheduler, the Reporting Engine does not set any property for the jobs that it launches.
- If the name of the pool or queue does not exist in the cluster, then Capacity Scheduler uses the default queue for the report. The Fair Scheduler may not execute the rule or creates a new pool with the lowest share. This is based on the value specified for the Fair Scheduler property `mapred.fairscheduler.allow.undeclared.pools`.
- If you do not specify a pool or queue, the job launched by the test rule is in the `mapr` pool or the default queue. It recommends that you configure a pool mapr with low (around 1/10 of total capacity) share with `maxRunningJobs = 2` so that these rules do not disrupt running reports. Make sure that you do not specify this pool name for any reports.

Specify the Pools and Queues

To specify the pools and queues:

1. Go to **(Admin) > Services**.
2. Select **Reporting Engine** and click > **View** > **Explore**.
3. Select **com.rsa.soc.re** > **Configuration** > **HiveConfig** > **hiveconfig** > **WarehouseResourcePoolNames**.
4. In the **WarehouseResourcePoolNames** field, enter the pool or queue names separated by spaces. For example, to configure four pools or queues with the names pool1, pool2, wrong and default, enter the names separated by a space.



Define Reports, Charts and Alerts

After you configure the Reporting Engine and the required data source based on your requirement, you can generate your Reports, Charts, and Alerts.

Define Reports

After creating the data sources and configuring the user permissions to these data sources, you can now use these data sources to perform the following tasks for Reporting module:

- Define a Rule
- Test a Rule
- Schedule Reports
- Add an Alert
- Add a Chart
- Test a Chart

For more information, see the above topics in the *NetWitness Reporting Guide*.

Define Charts

After creating the data sources and configuring the user permissions to these data sources, you can now use these data sources to perform the following tasks for Reporting module:

- Define a chart and Chart Groups
- Test a Chart
- Investigate Charts
- Manage Charts

For more information, see the above topics in the *NetWitness Reporting Guide*.

Define Alerts

After creating the data sources and configuring the user permissions to these data sources, you can now use these data sources to perform the following tasks for Alerting module:

- Configure Alerts
- Generate Alerts
- Add an Alert
- View an Alert

- View and Alert
- View Alerts Schedule
- Investigate an Alert

For more information, see the above topics in the *NetWitness Reporting Guide*.

Configure Reporting Engine General Settings

After you add and configure the Reporting Engine service, the system settings are defined with default values to achieve optimal results. However, you can modify and customize the Reporting Engine notifications based on your requirement by navigating to the **General** tab in the **Services** > **Config** view for a Reporting Engine.

Access the General Tab

You need to open the General tab to configure the general parameters for Reporting Engine.

To access this view:

1. Go to  **(Admin) > Services**.
2. In the Services list, select a **Reporting Engine** service.
3. Click **View** > **Config**.
4. Select the **General** tab.
5. Click **Apply** after you edit the parameters.

After you navigate to the General tab, you can modify the following parameters.

- System Configuration
- Logging Configuration
- Warehouse Kerberos Configuration

For more information see, General tab for details on the configuration parameters.

Troubleshooting Reporting Engine Configuration

Problem	Solution
When the NetWitness Server IP address is changes, the Reporting Engine datasource configuration does not get updated with the new IP	You must manually configure the new IP. Perform the following manual steps to configure the new IP address: 1. Login in NetWitness Platform. 2. Go to  **(Admin) > Services** > **Reporting Engine** > **View** > **Config**. 3. Click the **Output Actions** tab. 4. Add the new IP address in the **Hostname** field. 5. Click **Apply**. 6. Click the **Sources** tab and add the data sources again.

References

You can customize and make optimum use of the service by modifying the Reporting Engine settings in the **Services** > **Config** view, which has parameters that specifically pertain to the Reporting Engine.

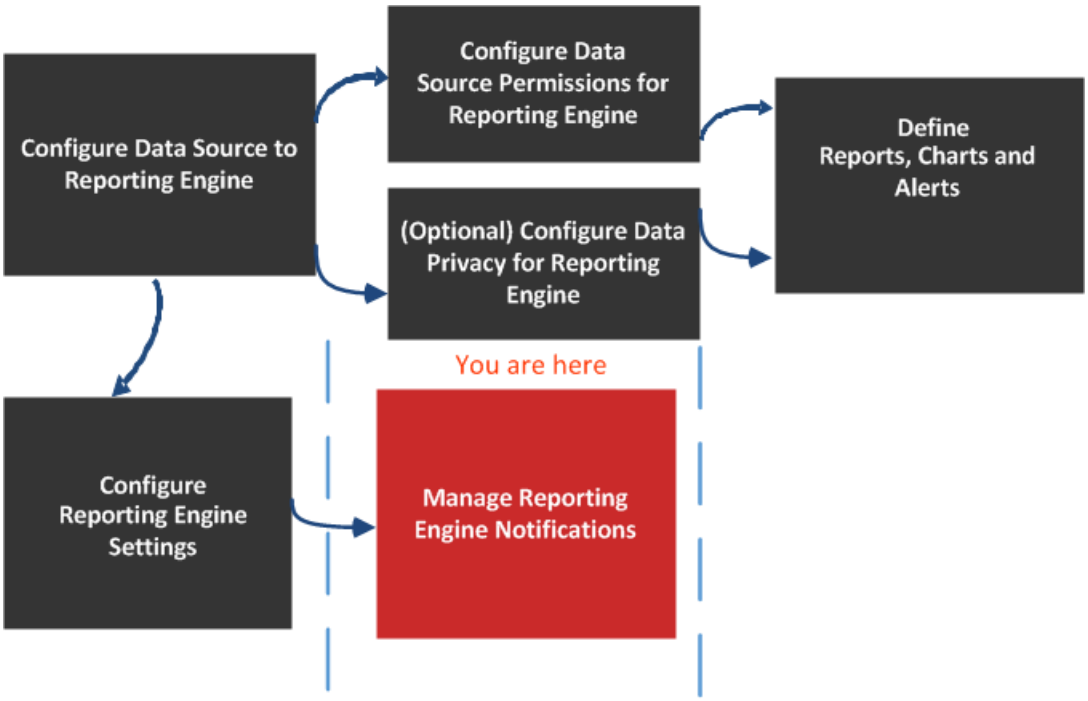
See the following sections for details:

- Reporting Engine General Tab
- Reporting Engine Sources Tab
- Reporting Engine Output Actions Tab
- Reporting Engine Manage Logos Tab

Reporting Engine General Tab

The General tab for the Reporting Engine service controls several settings that can tune the performance of a service and specify the user credentials for the service. Navigate to Services > View > Config > Reporting Engine > General. These settings are used for the Reporting Engine service exclusively.

The required permission to access this view is Manage Services.



What do you want to do?

Role	I want to ...	Refer to...
Administrator	Configure Data Source to Reporting Engine	Configure the Data Sources
Administrator	Configure Data Source Permissions for Reporting Engine	Configure Data Source Permissions
Administrator	Configure Data Privacy for Reporting Engine	Configure Data Privacy for the Reporting Engine
Administrator	Define Reports, Charts, and Alerts	Define Reports, Charts and Alerts
Administrator	Configure Reporting Engine Settings	Configure Reporting Engine Settings
Administrator / SOC Manager	Configure System Settings*	Configure Reporting Engine General Settings

Role	I want to ...	Refer to...
Administrator / SOC Manager	Configure Logging*	Configure Reporting Engine General Settings
Administrator / SOC Manager	Configure Warehouse Kerberos*	Configure Reporting Engine General Settings

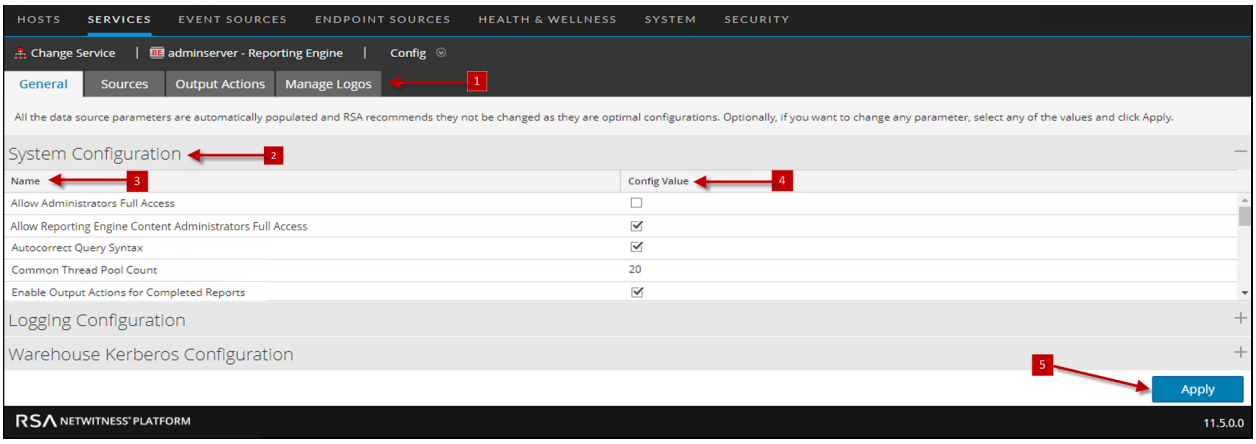
*You can complete these tasks here.

Related Topics

- How Reporting Engine Works

Quick Look

Here is example of the General tab where service configurations are displayed.

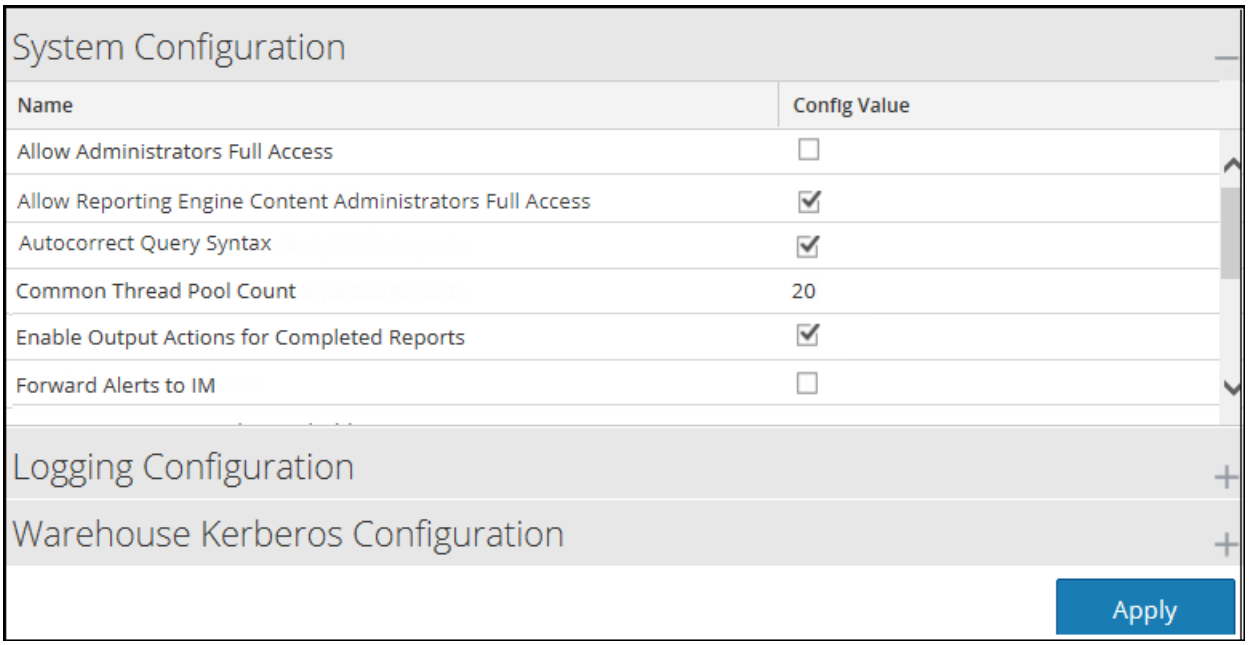


1 Displays all the available configurable tabs.
2 Displays the available configuration parameters for the system.
3 Displays the name of the parameter.
4 Displays the set values for each parameter.
5 Applies the changes.

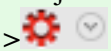
System Configuration

The System Configuration panel parameters for the Reporting Engine manage service configuration for a Reporting Engine service. When you add a Reporting Engine service, default values are in effect. The default values are designed to accommodate most environments and recommends that you do not edit these values because it may adversely affect performance.

The following figure displays the fields that can be configured in the System Configuration panel:



The following table describes the System Configuration panel features.

Name	Config Value
Allow Administrators Full Access	Select the checkbox if you want to access all the Reporting Engine objects (Reports, Rule, Charts, Schedule, and List) created by other users (non-admin). By default, this is not enabled. **Note:** If you enable the checkbox and then disable it, the access on all Reporting Engine objects that were enabled by selecting the checkbox will not be accessible. But, if you have defined the access on specific objects via Permissions window (**Reports > Manage > RE Object >**  **> Permissions**), enabling/disabling this checkbox will not have impact on these objects.
Allow Reporting Engine Content Administrators Full Access	Select the checkbox if you want to access all the Reporting Engine objects (Reports, Rule, Charts, Schedule, and List) created by other users. By default, this is enabled. If disabled, you cannot access the Reporting Engine objects (Reports, Rule, Charts, Schedule, and List) created by other users.
Common thread pool count	The number of thread pools assigned for executing common tasks on the Reporting Engine. A valid value is an integer (**20** default).
Enable Output Actions for Completed Reports	Select the checkbox to process the output actions only for reports with all rule executions successful. By default, this is enabled. If disabled, the output actions are processed for all scenarios (completed, partial, failure).
Forward Alerts to Respond	Select the checkbox to forward all the alerts to Respond. By default, this is not enabled.

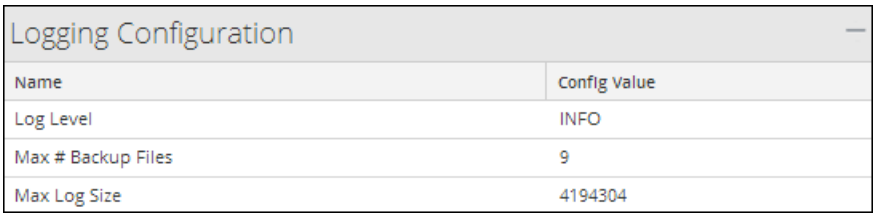
Name	Config Value
Max# of Concurrent Alerts	The maximum number of alerts that can be run simultaneously. This has a direct impact on the RSA service against which the alerts are run, as each alert consumes a query thread on the RSA service. A valid value is an integer (**10** default).
Max # of Concurrent Charts	The maximum number of charts that can be run simultaneously. A valid value is an integer (**10** default).
Max # of Concurrent LookupAndAdd Queries	The maximum number of parallel LookupAndAdd Queries that can be run per NWDB rule. A valid value is an integer (**2** default). When you increase this value, for better performance, you must ensure the NWDB data source is configured to handle the parallel queries.
Max # Concurrent List Value Reports	The maximum number of list value reports per schedule that can be generated in parallel. A valid value is an integer (**1** default).
Max # List Value Reports	The maximum number of list value reports generated, irrespective of the number of values in the list. A valid value is an integer (**10000** default).
Max rows stored per Rule (Billions)	The maximum number of rows that a rule can fetch when queried. A valid value is an integer (**100** default).
Maximum disk space threshold	The maximum disk space threshold allotted (in GB) to execute reports, alerts and charts. The initial value is configured based on the available system space.
Minimum disk space threshold	The minimum disk space threshold allotted (in percentage) required to execute reports, charts, and alerts. By default, this is value is set to 5. **Note:** If the minimum threshold is reached, then the execution of reports, charts and alerts will stop even if the service is running.
NWDB Info Queries Time Out	The info queries time out in seconds for NWDB server. A valid value is an integer (**1800** default).
NWDB Maximum aggregate Rows	The maximum number of rows that is returned when an aggregation is used in the NWDB rule. A valid value is an integer (**1000** default).
NWDB Query Time out	The time out in seconds for NWDB server to time out the rule execution, if it cannot process the result in configured time. A valid value is an integer (**259200000** default).
Process output actions for successful reports only	Select this checkbox to process output actions only for reports whose all rule executions are successful. When you de-select this checkbox, output action will be triggered for partial, completed, and failed reports. **Note:** This is applicable for all output actions except for dynamic list output actions.
Retain Alert history for # days	The maximum number of days to retain the alert history and alert status. A valid value is an integer (**100** default).

Name	Config Value
Retain Chart history for # days	The maximum number of days to retain the chart history and chart status. A valid value is an integer (**30** default).
Retain Report history for # days	The maximum number of days the system retains report history and report status. A valid value is an integer (**100** default).
Schedule Thread pool count	The number of thread pools assigned for scheduled tasks (for example, clearing history) on the Reporting Engine. A valid value is an integer (**5** default).

Logging Configuration

The Logging Configuration panel parameters of the Reporting Engine manages the logging configuration for a Reporting Engine service. When you add a Reporting Engine service, default values are in effect. RSA designed the default values to accommodate most environments and recommends that you do not edit these values because it may adversely affect performance of the Reporting Engine.

The following figure displays the fields that can be configured in the Logging Configuration panel.



The following table describes the Logging Configuration panel features.

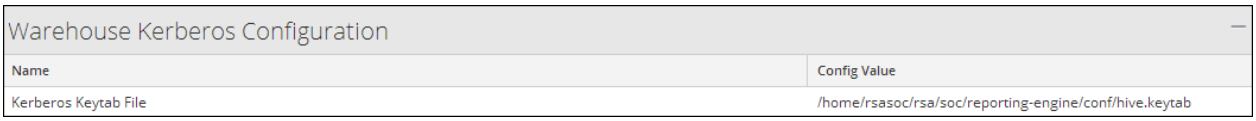
Name	Config Value
Log Level	The logging level that determines the scope of information included in log files. Possible values are: • `ERROR` • `WARN` • `INFO (default)` • `DEBUG` • `ALL`
Maximum # Backup Files	The maximum number of backup log files the system retains. A valid value is an integer (**9** default).
Max Log Size	The maximum size (in bytes) of the primary log file. A valid value is an integer (**4194304** default).

For more information on Reporting Engine logging, see Accessing Reporting Engine Log Files.

Warehouse Kerberos Configuration

The Warehouse Kerberos Configuration panel provides a way to specify the Kerberos Keytab file on this Reporting Engine.

The following figure displays the field that can be configured in the Warehouse Kerberos Configuration panel:



The following table describes the Kerberos Configuration panel features:

Name	Config Value
Kerberos Keytab File	The Kerberos keytab file location. For example: `/var/netwitness/re-server/rsa/soc/reporting-engine/conf/hive.keytab`

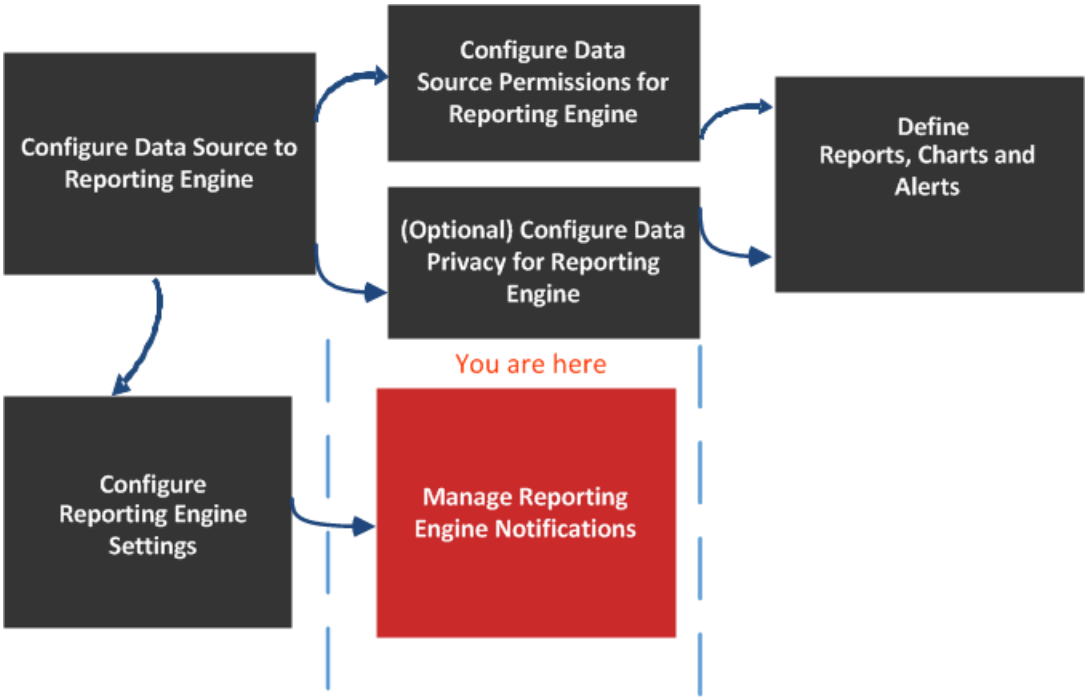
The default Kerberos configuration file is located at, `/etc/kbr5.conf` in the Reporting Engine. You can modify the configuration file to provide details for Kerberos realms and other parameters related to Kerberos.

Added the host name (or FQDN) and IP address of the Horton Works nodes and Warehouse Connector to the DNS server. If the DNS server is not configured, add the host name (or FQDN) and IP address of the Horton Works nodes and Warehouse Connector to the `/etc/hosts` file in the host on which the Warehouse Connector service is installed.

Reporting Engine Sources Tab

The services configuration parameters are available in the Sources tab of the Services Config view for the Reporting Engine. The Sources tab for the Reporting Engine service in the Services Config view controls that data sources associated with a Reporting Engine. The Source tab consists of a single panel with a toolbar and a grid that lists the data sources associated with the Reporting Engine.

Workflow



Role	I want to...	Refer to...
Administrator	Configure Data Source to Reporting Engine	Configure the Data Sources
Administrator	Configure Data Source Permissions for Reporting Engine	Configure Data Source Permissions
Administrator	Configure Data Privacy for Reporting Engine	Configure Data Privacy for the Reporting Engine
Administrator	Define Reports, Charts, and Alerts	Define Reports, Charts and Alerts
Administrator	Configure Reporting Engine Settings	Configure Reporting Engine Settings
Administrator	Add, delete or edit a new or available service*	Configure the Data Sources

Role	I want to...	Refer to...
Administrator	Set a data source as default*	Configure the Data Sources
Administrator	Configure data source permissions*	Configure Data Source Permissions

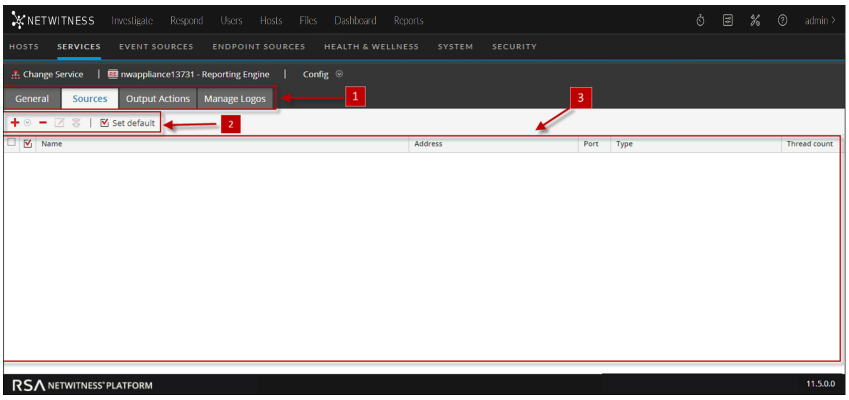
*You can complete these tasks here.

Related Topics

- How Reporting Engine Works

Quick Look

Here is example of the Sources tab where the available services are displayed.



1 Displays all the available configurable tabs.
2 Displays the available configuration parameters for the selected service .
3 Displays the field parameters for the selected service.

The data sources available to the Reporting Engine for which you are defining reports, charts and defining alerts are:

- **NWDB Data Sources** - The NetWitness Database (NWDB) data sources are Decoders, Log Decoders, Brokers, Concentrators, Archiver, and Collection.

  Note: When a data privacy plan has been implemented to limit access to sensitive data on a data source, you must configure different service accounts in Reporting Engine for privileged and non-privileged users. To configure different service accounts for data privacy, you can add more than one NWDB data source. This procedure is available under Configure Reporting Engine Settings.

  IMPORTANT: If you change the admin password on a NetWitness service that is used as a Reporting Engine data source, you must remove and then re-add the service as a data source.

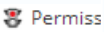
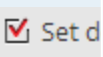
- **Warehouse Data Sources** - The Warehouse data sources are Horton Works and MapR.
- **Respond Data Sources** - Respond is used to generate reports on alerts and incidents. The Respond data sources are Reporting Engine, ESA, Malware, EndPoint, and Web Threat Detection. Respond is used to store the alerts and incidents reports.

If you set a source as the default data source, NetWitness uses that source when you create reports and alerts unless you choose to override it with one of the other sources listed in this tab.

Note: You can manage access control to NWDB and Warehouse Data Sources. For more information, see Configure Reporting Engine Settings.

Features

You can perform the following actions on the Sources tab:

Icon	Actions
	This option adds new services as data sources for Reporting Engine. Add existing services (Archiver, Workbench, and Collection) as data sources for Reporting Engine. For details, see the corresponding topic: • (Optional) Add Archiver as Data Source • (Optional) Add Collection as Data Source to Reporting Engine • (Optional) Add Workbench as Data Source
	This option removes data sources from a Reporting Engine.
	This option configures the Data Source Permissions. This is enabled only for NWDB and Warehouse Data Sources. For more information, see Configure Data Source Permissions.
	This option sets the default data sources for a Reporting Engine. This is the source to which NetWitness defaults in the **Data source** field of the following views: • Rule Definition view. • Create or Modify Alert view.

The NetWitness data sources are listed under the different categories as follows:

- NWDB Data Sources category displays the NetWitness data sources.
- Warehouse Data Sources category displays the Warehouse data sources.

Column	Description
	Clicking the check box selects the data source. After you select it, you can use toolbar to remove the source or set the source as the default.
Name	Displays the name of the data source.
Address	Displays the IP Address of the data source.
Port	Displays the port of the data source.
Type	Displays the service type of the data source.
Thread Count	Displays the thread pool size used for executing rules on the data source.

Reporting Engine Output Actions Tab

You can configure output actions for a Reporting Engine to determine the format you want the data to be presented to you based on your requirements. The service configuration parameters are available in the Output Actions tab of the Services Config view configured for a report or an alert execution. This tab consists of the following panels:

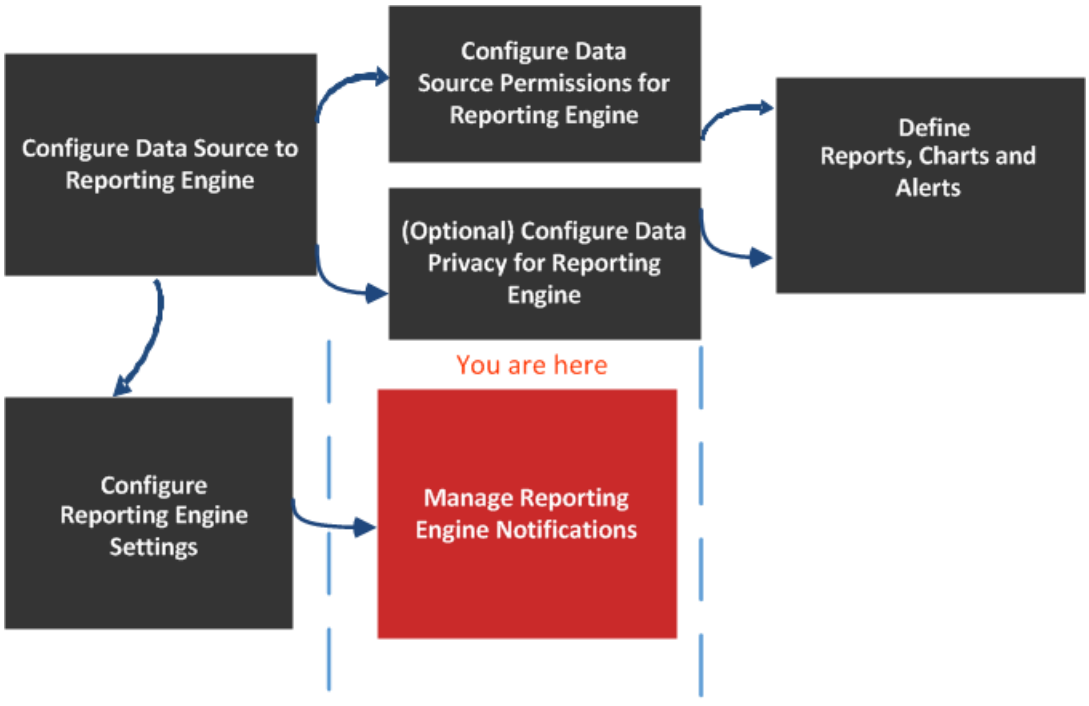
- NetWitness Configuration
- Simple Mail Transfer Protocol (SMTP)
- Simple Network Management Protocol (SNMP)
- Syslog
- Simple File Transfer Protocol (SFTP)
- Uniform Resource Locator (URL)
- Network Share

For instance, Syslog output action is used specifically for Reporting Engine Alerts, whereas, SFTP, URL, and Network Share output action is used specifically for Reporting Engine Reports.

You can configure the required permission to access this view in Manage Services.

You must ensure that the Reporting Engine is up and running and the data source from which you want to generate a report is configured in the NetWitness.

Workflow



What do you want to do?

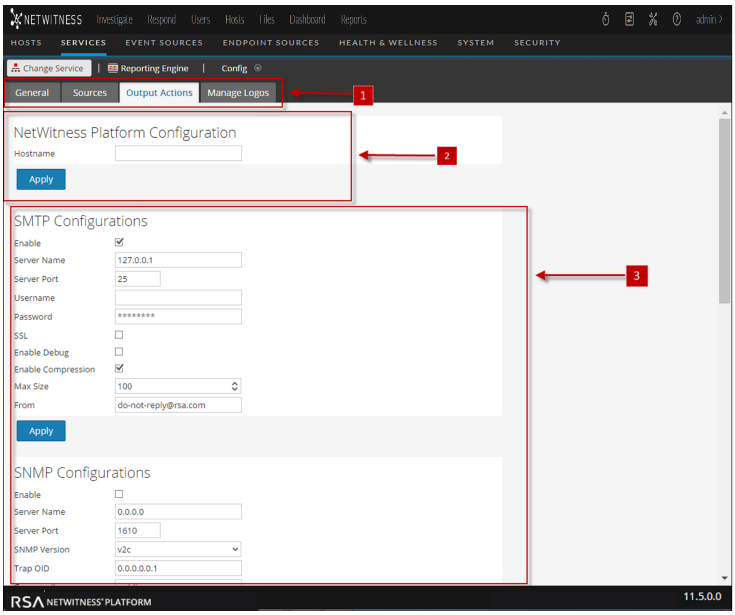
Role	I want to...	Refer to...
Administrator	Configure Data Source to Reporting Engine	Configure the Data Sources
Administrator	Configure Data Source Permissions for Reporting Engine	Configure Data Source Permissions
Administrator	Configure Data Privacy for Reporting Engine	Configure Data Privacy for the Reporting Engine
Administrator	Define Reports, Charts, and Alerts	Define Reports, Charts and Alerts
Administrator	Configure Reporting Engine Settings	Configure Reporting Engine Settings
Administrator	Configure NetWitness Configuration *	Configure Reporting Engine General Settings
Administrator	Configure SMTP Configuration*	Configure Reporting Engine General Settings
Administrator	Configure SNMP Configuration*	Configure Reporting Engine General Settings
Administrator	Configure Syslog Configuration*	Configure Reporting Engine General Settings
Administrator	Configure SFTP Configuration*	Configure Reporting Engine General Settings
Administrator	Configure URL Configuration*	Configure Reporting Engine General Settings
Administrator	Configure Network Share Configuration*	Configure Reporting Engine General Settings

*You can complete these tasks here.

Related Topics

- How Reporting Engine Works

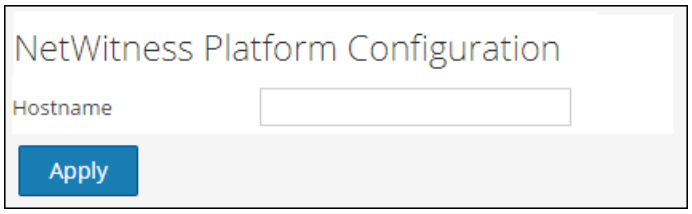
Quick Look



1	Displays all the available configurable tabs.
2	Displays the NetWitness configuration host.
3	Displays all the types of output action that can be configured.

NetWitness Configuration

The following figure shows the NetWitness Configuration on the Output Actions Tab.



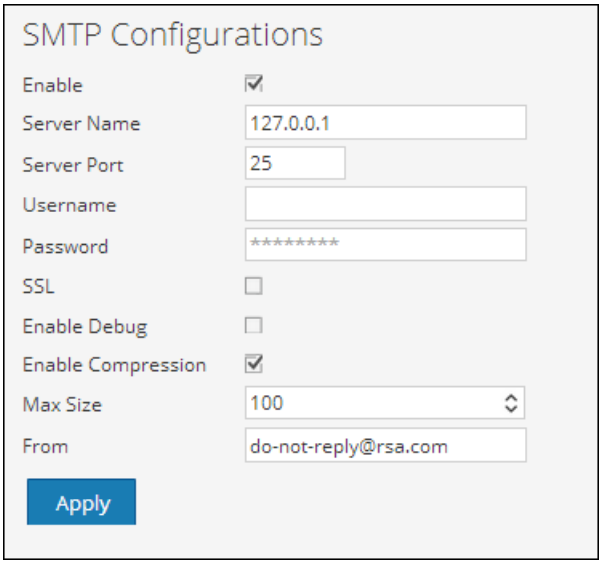
The following parameters identify the NetWitness host that is associated with the Reporting Engine.

Name	Config Value
Host Name	IP Address or Hostname of the NetWitness server. You must specify this parameter for all kind of deployments so that you can refer to this address to create investigation links to NetWitness from Reports, Alerts, and so on. The NetWitness uses this parameter to correctly generate: • SMTP Output Action • SNMP Output Action • Syslog Output Action • SFTP Output Action • URL Output Action • Network Share Output Action • Hyperlinks for meta values in Report PDFs
Apply	Update the configuration.

SMTP

After an execution is completed, an email notification is sent to the user based on the SMTP configuration.

The following figure shows the SMTP Configuration on the Output Actions Tab.



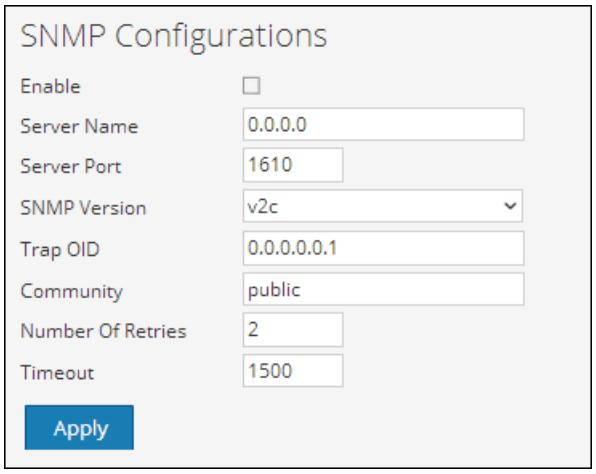
The following parameters manage SMTP (email) output action configuration for a Reporting Engine service. When you add a Reporting Engine service, default values are in effect. You must modify the **Config Values** of these parameters according to the requirements of your enterprise.

Name	Config Value
Enable	Check this box to enable SMTP as an output action for both alert and report from this Reporting Engine. By default, this value is enabled.
Server Name	Specify the hostname or IP Address of the server on which the target SMTP server runs. Default value is 0.0.0.0.
Server Port	Specify the SMTP server port number. Default value is 25.
Username	Specify the username of your SMTP account. Default value is blank. Password Specify
Password	Specify the password of your SMTP account.
SSL	Check this box to use Secure Socket Layer (SSL) to communicate with the SMTP server. Default value is do not use SSL.
Enable Debug	Check this box to enable debugging. Default value is do not enable debug.
Enable Compression	Check this box to enable compression. Default value is enable compression. If this value is enabled, the output files will have `.zip` extension.
Max Size	Specify the maximum size of attachments that can be sent. Default value is 100.
From	Specify the email address from which Security Analytics sends all messages. Default value is `do-not-reply@rsa.com`.
Apply	Update the configuration.

SNMP

After an execution is completed, a trap notification is sent to the user based on the SNMP configuration.

The following figure shows the SNMP Configuration on the Output Actions Tab.



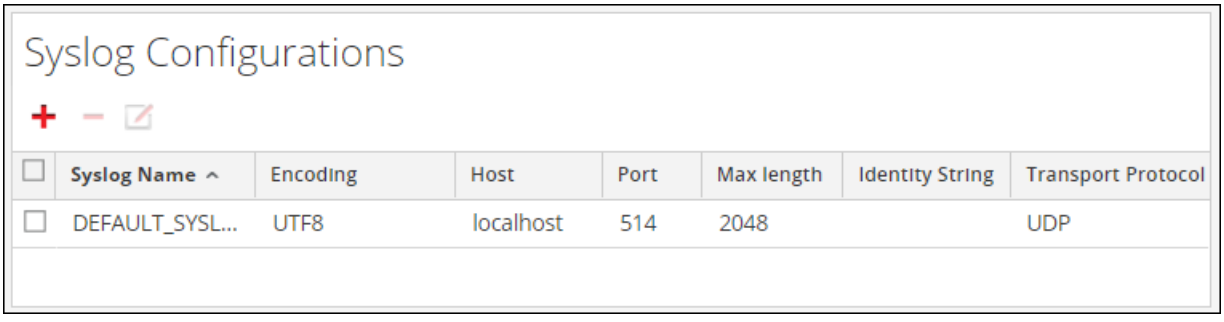
The following parameters manage SNMP (messages to network-attached services) output action configuration for a Reporting Engine service. When you add a Reporting Engine service, default values are in effect. You must modify the **Config Values** of these parameters according to the requirements of your enterprise.

Name	Config Value
Enable	Check this box to enable SNMP output action as an output for alert messages from this Reporting Engine. Default value is Disable.
Server Name	Specify the hostname or IP Address of the server on which the target SNMP server runs. Default value is **0.0.0.0**.
Server Port	Specify the port number of the server on which the target SNMP server listens for faults and exceptions. Default value is **1610**.
SNMP Version	Specify the version number of the SNMP protocol NetWitness uses to send SNMP traps.
Trap OID	Specify the object identification number that identifies the type of trap to send. Default value is **0.0.0.0.0.1**.
Community	Specify the SNMP group to which NetWitness belongs. The default value is **public**.
Number Of Retries	Specify the maximum number of times NetWitness tries to resend the alert message through SNMP. Default value is **2**.
Timeout	Specify the number of seconds after which NetWitness times out (stops trying to send SNMP alerts). Default value is **1500**.
Apply	Update the configuration.

Syslog

After an execution is completed, all notifications are sent via Syslog messages to a particular host based on the Syslog configuration. Multiple Syslog servers can be configured on the Syslog Configuration panel.

The following figure displays the Syslog Configuration on the Output Actions Tab.



The following parameters manage syslog output action configuration for a Reporting Engine service. When you add a Reporting Engine service, you can define values for this output configuration, as no default values are available for this configuration. You must modify the **Config Values** of these parameters according to the requirements of your enterprise.

Name	Config Value
Syslog Name	The name of the Syslog configuration. **Note:** You cannot create a Syslog configuration with a name that already exists in the Reporting Engine Syslog configuration list.
Encoding	Specify the internationalization encoding for Syslog messages. Default value is **UTF8**.
Server Name	Specify the hostname or IP Address of the server on which the target Syslog process runs. Default value is blank.
Server Port	Specify the port number of the server on which the target Syslog server listens for faults and exceptions. Default value is **514**.
Max Length	Specify the maximum size (in bytes) of each Syslog alert message. Default value is **2048**. If **UDP** is the transport type and the Syslog message size is greater than 1024 bytes, you must configure a Syslog server that supports message sizes greater than 1024 bytes.
Identity String	Specify the string NetWitness inserts as a prefix in all Syslog alert messages. Default value is blank.
Include Local Hostname	Check this box to include the local hostname in all Syslog alert messages. Default value is do not include local hostname.
Truncate Message	Check this box to truncate all Syslog alert messages. Default value is do not truncate Syslog messages.
Use Identity	Check this box to use the IDENT protocol. Default value is does not use this protocol.
Include Local Timestamp	Check this box to include the local timestamp in all Syslog alert messages. Default value is do not include local timestamp.
Transport Protocol	Specify the transport type for Syslog message delivery. There are three parts to the Syslog transport type: UDP, TCP, and SECURE_TCP. Default value is **UDP**.
Syslog Message Delimiter	Specify the delimiter for the Syslog message. There are three delimiters: CR, LF, and CRLF. By default the value is **CR**. **Note:** This field populates when you select TCP or SECURE_TCP as the transport protocol.
Trust Store Password	Specify the password for the Trust store. **Note:** This field populates when you select SECURE_TCP as the transport protocol.
Key Store Password	Specify the password for the Key store. **Note:** This field populates when you select SECURE_TCP as the transport protocol.

Name	Config Value
Apply	Save the configuration.

SFTP

After an execution is completed, you can send or transfer files to a remote location based on the SFTP configuration.

The following figure displays the SFTP Configuration on the Output Actions Tab.



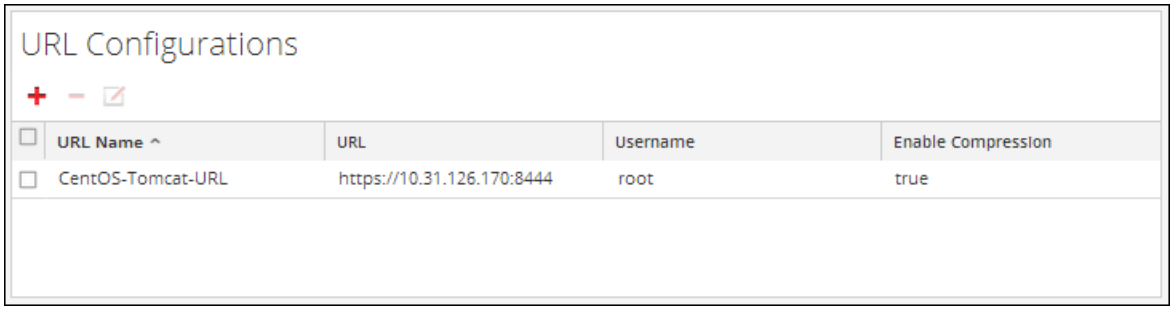
The following parameters manage SFTP (file transfer to a local drive) output action configuration for a Reporting Engine service. When you add a Reporting Engine service, you can define values for this output configuration, as no default values are available for this configuration. You must modify the **Config Values** of these parameters according to the requirements of your enterprise.

Name	Config Value
SFTP Name	The name of the SFTP configuration. **Note:** You cannot create an SFTP configuration with a name that already exists in the Reporting Engine SFTP configuration list.
Host	The IP Address or Hostname of the Reporting Engine server associated with the file transfer.
Port	If you want to use a different port than the default port, enter a port number. Default value is **22**.
Username	Specify the username for the SFTP configuration.
Password	Specify the password for the SFTP configuration.
Custom Folder	Select an SFTP location where you want to transfer the file to. You can use the pre-defined Windows or Linux directory structure in the custom folder path. For example, **/root/Downloaded_Files**. **Note:** If the directory does not exist, RE will create the directory in the custom folder path and copy files to this directory.
Enable Compression	Select this checkbox to enable compression. Default value is enable compression. If this value is enabled, the output files will have ".zip" extension.

URL

After an execution is completed, the output files are published to a URL based on the URL configuration.

The following figure shows the URL Configuration on the Output Actions Tab.



The following parameters manage URL (file transfer to a URL) output action configuration for a Reporting Engine service. When you add an Reporting Engine service, you can define values for this output configuration, as no default values are available for this configuration. You must modify the Config Values of these parameters according to the requirements of your enterprise.

Name	Config Value
URL Name	The name of the URL configuration. **Note:** You cannot create a URL configuration with a name that already exists in the Reporting Engine URL configuration list.
URL	The URL address associated with the file transfer.
Username	Specify the username for the URL configuration.
Password	Specify the password for the URL configuration.
Enable Compression	Select this checkbox to enable compression. Default value is enable compression. If this value is enabled, the output files will have ".zip" extension.

After the URL is configured, the files will be copied under the "URL_OUTPUT_ACTION" directory and the following parameters are sent to the server along with the compressed file.

Name	Config Value
filename	The name of the file.
filesize	The file size in bytes.
filetype	The file type associated with the file.
filechecksum	The number computed from a file that can be used to confirm that this is the one you expect and has been downloaded and stored properly.
hashingalgorithm	The hashing algorithm used to calculate the file checksum.

Name	Config Value
reportname	The name of the downloaded report.
executionid	The execution id associated with the report execution.
reportexecutionstarttime	The start time the report was executed.
status	The report creation status.
status description	The status description.

Network Share

After an execution is completed, you can transfer the output files to a mounted path or shared location based on the Network Share configuration.

The following figure shows the Network Share Configuration on the Output Actions Tab.



The following parameters manage Network Share (file transfer to a shared location on the network) output action configuration for a Reporting Engine service. When you add a Reporting Engine service, you can define values for this output configuration, as no default values are available for this configuration. You must modify the **Config Values** of these parameters according to the requirements of your enterprise.

Name	Config Value
Network Share Name	The name of the Network Share. **Note:** You cannot create a Network Share configuration with a name that already exists in the Reporting Engine Network Share configuration list.
Mounted Path	The path (location) associated with the file transfer. You can use the pre-defined Linux directory structure in the mounted path. For example, **/mnt/win**. **Note:** The 'rsasoc' user must have read-write access to the specified Network Share mounted path.
	Click to view how the mounted path is created. This pop-up notifies that you must manually create the mounted path.

Name	Config Value
Destination Directory	Name of the directory where the transferred file is stored in the shared location.
Enable Compression	Select this checkbox to enable compression. Default value is enable compression. If this value is enabled, the output files will have ".zip" extension.

The following table lists the common operations you can perform in the Syslog, SFTP, URL and Network Share sections.

Operation	Description
	Create a Syslog, SFTP, URL and Network Share configuration.
	Delete a Syslog, SFTP, URL and Network Share configuration.
	Edit a Syslog, SFTP, URL and Network Share configuration.

Reporting Engine Manage Logos Tab

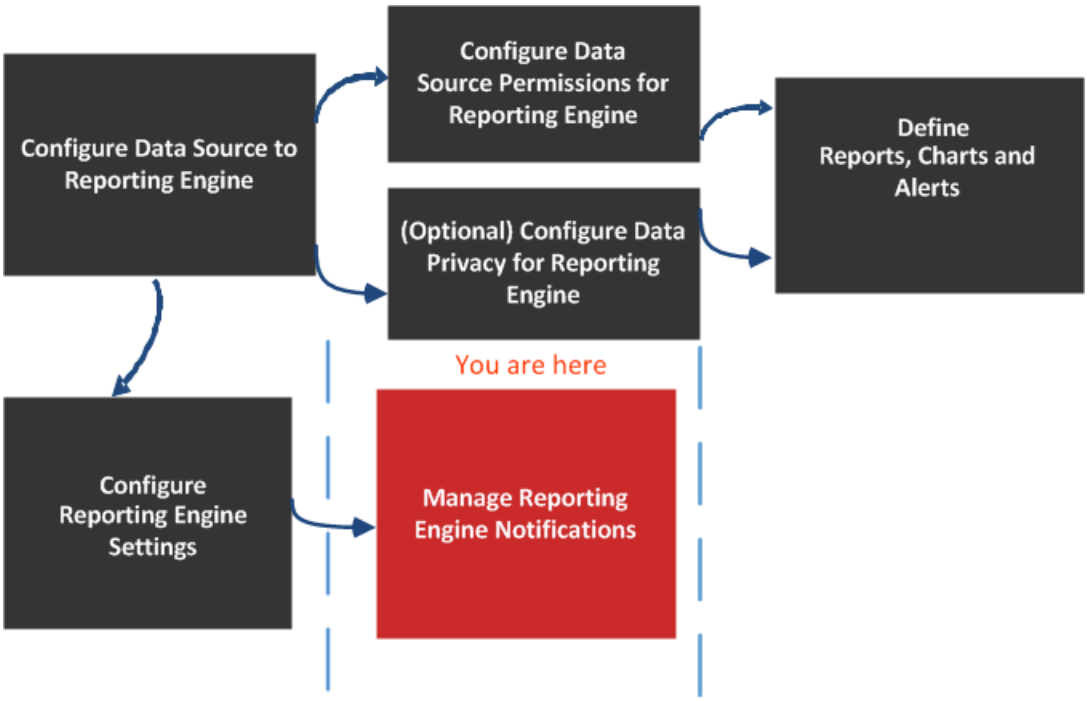
The Manage Logos option available in the **Services Config View** > **Manage Logos** tab, helps you to manage the logos associated with the Reporting Engine. The Manage Logos tab consists of a single panel with a toolbar and a grid that lists the logos.

You can upload the logos that you want to use in your report. After you upload the logo, you can set any logo as a default logo which will be automatically used in all the scheduled reports. You can choose to override the default logo with any other logo listed in this tab when you schedule a report. For more information, see "Select a Logo Dialog" topic in the *Reporting Guide*.

The supported image formats are:

- .jpg
- .png
- .gif

Workflow



What do you want to do?

Role	I want to...	Refer to...
Administrator	Configure Data Source to Reporting Engine	Configure the Data Sources

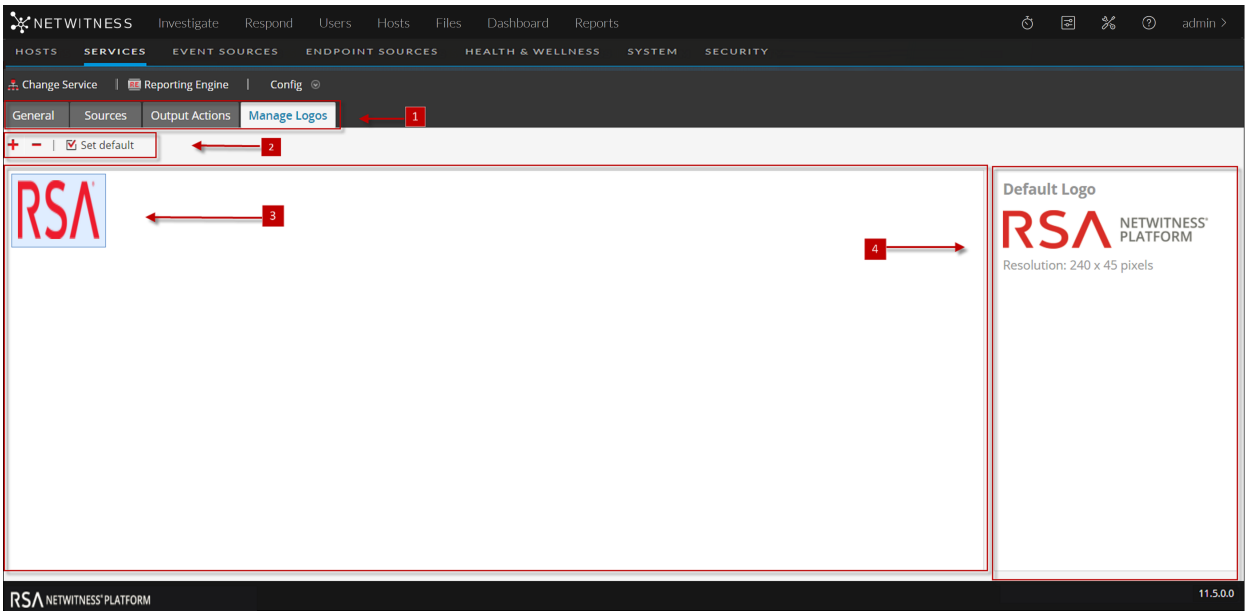
Role	I want to...	Refer to...
Administrator	Configure Data Source Permissions for Reporting Engine	Configure Data Source Permissions
Administrator	Configure Data Privacy for Reporting Engine	Configure Data Privacy for the Reporting Engine
Administrator	Define Reports, Charts, and Alerts	Define Reports, Charts and Alerts
Administrator	Configure Reporting Engine Settings	Configure Reporting Engine Settings
Administrator / SOC Manager	Add, or delete logos*	Configure Reporting Engine General Settings
Administrator / SOC Manager	View the list of logos*	Configure Reporting Engine General Settings
Administrator / SOC Manager	Set a logo as default*	Configure Reporting Engine General Settings

*You can complete these tasks here.

Related Topics

- How Reporting Engine Works

Quick Look



Note: The logo to be uploaded should not exceed 500 KB. The required permission to access this view is Manage Services.

1 Displays all the available configurable tabs.
2 Displays edit actions.
3 Displays all the logos that have been used
4 Displays the default logo used.

You can perform the following actions on the Manage Logos Tab.

Icon	Actions
	Add new logos from the local directory of the system to the Reporting Engine. **Note:** The logo size cannot exceed 500 KB. The logos chosen must be of the following file types: * .jpg * .gif * .png
	Removes logos from the Reporting Engine. **Note:** By performing (Ctrl+click), you can select multiple logos to delete.
	Sets the default logo for a Reporting Engine. This is the logo NetWitness defaults to in the **Log** panel of the Schedule a Report view. **Note:** If no default logo is selected, the logo is displayed.

NetWitness® Platform XDR

Version 12.0

Reporting User Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Reporting Overview **7**
- Reporting Guidelines 10
- Access Control for Reporting 19

Configure and Generate a Report **24**

Configure a Rule **25**
- Create a Rule Group 25
- Create a Rule Using NetWitness Data Source 26
- Create a Rule Using Warehouse Data Source 29
- Create a Rule Using Respond Data Source 34
- Deploy a Rule 38
- Test a Rule 51
- Create a Lists or List Group 53

Create and Schedule a Report **56**
- Create a Report or Report Group 56
- Schedule a Report 57
- Generate a List from the Scheduled Report 62
- Create a Parameterized Report Using Variable 63
- Report with Dynamic Variables 65
- Iterative Report 70
- Create a Report Using a Rule 74

View a Report **75**

Investigate a Report **77**

Manage Lists, Rules or Reports **78**
- Manage a List 78
  - Access Control for a List and List Group 78
  - Edit a List 84
  - Delete a List or List Group 84
  - Duplicate a List 86
  - Export a List or List Group 86
  - Import a List or List Group 88
  - Filter Unused Lists 89
- Manage a Rule 90
  - Access Control for a Rule and Rule Group 90
  - Delete a Rule or Rule Group 98
  - Duplicate a Rule 99

Edit a Rule 100
View Dependents of a Rule 101
Export a Rule or Rule Group 103
Filter Unused Rules 104
Manage a Report 104
Access Control for a Report or Report Group 105
Delete a Report or Report Group 114
Duplicate a Report 115
Edit a Report 116
Refresh a Report Group or Report List 117
Edit a Scheduled Report 117
Delete a Scheduled Report 119
Export a Report 120
Export a Report Group 121
Import a Report or Report Group 121
Enable or Disable a Scheduled Report 122
Start or Stop a Scheduled Report 123
View an Execution History of a Scheduled Report 123
Stop an Individual Execution of a Scheduled Report 124
Manage and Select a Report Logo 125
Search Reporting Details 126
Working with Charts 131
Configure and Generate a Chart 131
Configure a Chart 135
Schedule a Chart 137
View a Chart 137
Test a Chart 138
Investigate a Chart 139
Manage a Chart Group and Chart 140
Working with Alerts 149
Alerting Overview 149
Configure Reporting Engine 153
Configure an Alert 155
Schedule an Alert 157
View an Alert 158
Investigate an Alert 158
Manage an Alert and Alert Template 158
Appendix 166
Rule Syntax 167
NWDB Rule Syntax 167

Sample Supported Queries ... 220
Respond Rule Syntax ... 220
Warehouse DB Simple Rules Syntax ... 225
Warehouse DB Advanced Rules Syntax ... 234
Automated Partition using Custom function ... 247
General syntax ... 247
Task Scheduler for Warehouse Reporting ... 253
Query Aggregates ... 254
Troubleshoot Reporting ... 278
Meta Values in Investigation Link Issue ... 278
Internet Explorer 10 Browser Issue ... 279
Dynamic List Editing Issue ... 279
Deployment Failure Issue ... 279
Respond Server Issue ... 279
Post-Upgrade Issue ... 279
Report Query Timeout Issue ... 280
Reporting References ... 281
Build Chart View ... 282
Build List View ... 285
Build Report View ... 289
Build Rule View ... 295
Chart Permissions Dialog ... 302
Chart View ... 305
Execution History Panel ... 309
Generate List Panel ... 313
Import Chart Dialog ... 316
Import Report Dialog ... 318
Investigate a Chart View ... 320
Lists Permissions Dialog ... 322
List View ... 325
Reports Permissions Dialog ... 328
Report View ... 331
Rule Permissions Dialog ... 335
Rule View ... 338
Select a Logo Dialog ... 342
Schedule a Chart View ... 345
Schedule Report Panel ... 348
Scheduled Reports View ... 355
Test a Chart View ... 360
View a Chart Panel ... 363

View All Charts View 367
View a Report Panel 371
View All Reports View 377
Alerting References 381
Alert List View 382
Alert Permissions Dialog 385
Alert Schedules View 388
Create or Modify Alert Panel 391
Investigate an Alert View 399
Import Alert Dialog 401
Alert Template References 403
Alert Template View 404
Create or Modify Template View 407
View Alerts Schedule View 409
View Alerts View 412

Reporting Overview

Reporting is a collection of data as a result of monitoring the network traffic, which can be used for further analysis. In NetWitness you can run a report against NetWitness Database core services to identify the network activities. For example, if you want to identify the Top Source Countries and Destination Countries, or top Threat and Risk trends that help monitor any changes to the normal categories or monitor the users and services that may potentially have malicious activities etc.

The reporting typically consist of: Reports and Charts. You can report on the log, packet and endpoint data collected, and customize the reports and charts to enhance the visual appearance. You can create real-time reports for historical data. You can create charts and dashlets, that can be added in the real-time chart dashlets as well.

Reporting Engine

Reporting relies on the Reporting Engine to provide data for the reports, alerts and charts. Hence, you must configure the Reporting Engine as a service to NetWitness before you can generate the reports. You must also specify the data source in the Reporting Engine from which the data is extracted.

The data that you can report or alert depends on the configuration of Reporting Engine and the data sources that you specify as part of the rule definition.

Note: Make sure you have access to the components in the Reporting.

Note: Reporting is accessible based on the role based access, defined for the user.

Report

A report is a combination of rules and other formatting objects such as headers and HTML-formatted notes that describe and identify data pertaining to a particular area of interest. Reports are defined and managed in the Build Report page and can be scheduled to run on an adhoc or timely basis. Once a report is run, results are stored centrally and can be automatically sent over email, SFTP, URL, and NFS to users, viewed via the NetWitness web interface, downloaded as PDF and CSV files.

A report consists of the following:

Property	Description	Example
Report Name **Note:** For **Name** field, the icon to extend the column size is not displayed at the end of the column field. You have to hover the mouse a little to the left side to see the icon for extending the column.	Used to identify the report to schedule them at a later time.	Report1
Text	Pre-defined text fields used within a report to make the report more meaningful to the user.	Header1, Comment
Rules	The rules (queries) used to create a report.	select user.dst where ip.src = 10.10.10.1

Note: In the Reporting user interface, the displayed date or time is always according to the user-selected time zone profile.

Rule

A rule is the basic and essential building block in the Reporting. You must create a rule which can be used in a Report, Chart or Alert.

A rule represents a unique query that detects and summarizes the requested information within a collection of network data.

The rule syntax is very similar to that of Standard Query Language (SQL) where you can use the select clause, where clause, sort and group options and limits for the result set. A rule consists of the following:

Property	Description	Example
Name	The name of the rule.	Windows System Account Activity
Select	List of meta types that are returned in the result set. The list of meta types is provided in the Meta Library. Meta Library in the Rule Builder is continually synchronized with the index configuration of the NetWitness host to which NetWitness is connected. The number of meta types that this property can represent depends on how the rule is to be sorted. If the Sort by property is 'None' or Custom, a rule can have more than one select field, for example, for each match, include the ip.src, ip.dst, size, time in the rule result. If a rule is set to be sorted, either by session count, session size, or packet size, then there can only be one field on which to select.	
Where	A clause that is the base query for the rule.	alert='cleartext_ftp_password'
Then (Rule Actions)	A series of functions that manipulate the original result set of a rule in order to make the output in a report more meaningful or add additional functionality other than querying and displaying data.	lookup_and_add ('username','ip.src',10);
Sort By	Determines how the data in the result set is sorted. The various possibilities are: • Total • Value • Column Name	Total
Limit	Designates how large a result set can be for the given rule. Users must note that if a result set is sorted by count or size, the limit represents the top (or bottom) N values to be returned. If the result set is not sorted, the first N values are returned.	20

Note: In the User Interface (UI), the date or time displayed depends on the time zone selected by the user.

Rule Types

There are different rule types in the Reporting. Rule types designate the source of data for the report rule. Following are the rule types:

Rule Type	Description
NetWitness Database (**NetWitness DB**)	The NetWitness database extracts the meta from a Reporting Engine configured to use a Concentrator, Broker and Archiver as the data sources and provides the meta for rules.
Warehouse Database (**Warehouse DB**)	The Warehouse database, also referred to as the NetWitness Warehouse, warehouses large volumes of data. The Warehouse is designed so that you can retrieve large volumes of data easily and efficiently. The Warehouse also extracts the meta from the Reporting Engine.
Respond Database (**Respond DB**)	The Respond database contain alerts and incidents generated from different services and you can create a report on those alerts and incidents.
Unused Rules	The Unused Rules is an option that helps in listing out the rules that are not used by any chart, alert and, report. For more information, see Filter Unused Rules.

Note: In the User Interface (UI), the date or time displayed depends on the time zone selected by the user.

List

A list is a variable that refers to a series of comma-separated values (CSV). You can insert a list into a rule or use it as an argument to a rule action. Lists can act as placeholders for other values, which you can populate and update as needed.

You can create, manage and view lists that can be used to define rules for Reporting and Alerting. You can also filter and delete the unused lists.

Lists cannot be empty or have duplicate or blank values.

Note: If you are defining a report with a rule which has lookup_and_add in the **Then** clause and direct the report output to a list, the list is not populated with the result.
For example, if you create a rule with ip.src in the **Select** clause and lookup_and_add ('ip.dst','ip.src', 10) in the **Then** clause, the report displays the result, but if you have redirected the output to a list, the list will be empty

Chart

Chart is a tabular or grid representation of data. It consists of the following:

Property	Description	Example
Chart Name	Identifies the chart.	Chart1
Rule Basis	Identifies the rule path chosen from the folder hierarchy.	

Any NetWitness DB rule in the Reporting Engine system which is not sorted by none can be used to instantly create a chart. In NetWitness, the chart interval can be adjusted from the chart definition panel itself. Each time a chart runs, it stores its result data locally in the Reporting Engine, so that it can be reviewed in either the Dashboard View or Chart View without any performance considerations. You can also filter the unused charts from Chart View.

Note: In the Reporting user interface, the output for the field where Date and Time are displayed is always according to the user-selected time zone profile.

Note: The Reporting Engine (RE) will automatically check for the available disk space before you execute a Test Rule, Report, Chart and Alert. If the RE disk space (in percentage) is less than the minimum disk space threshold (default value is 5), the RE will stop the current execution and an error message 'Available disk space of Reporting engine home is <5%, please clean up the space to proceed further' is displayed. Additionally, you may also configure the minimum disk space threshold by using the following path: **RE>Config>General>System Configuration>Mini disk space threshold in %.**

Reporting Guidelines

This section lists recommended guidelines to enhance the execution time of your reporting entities such as rules, reports, alerts, charts, and lists. The guidelines are provided for the following:

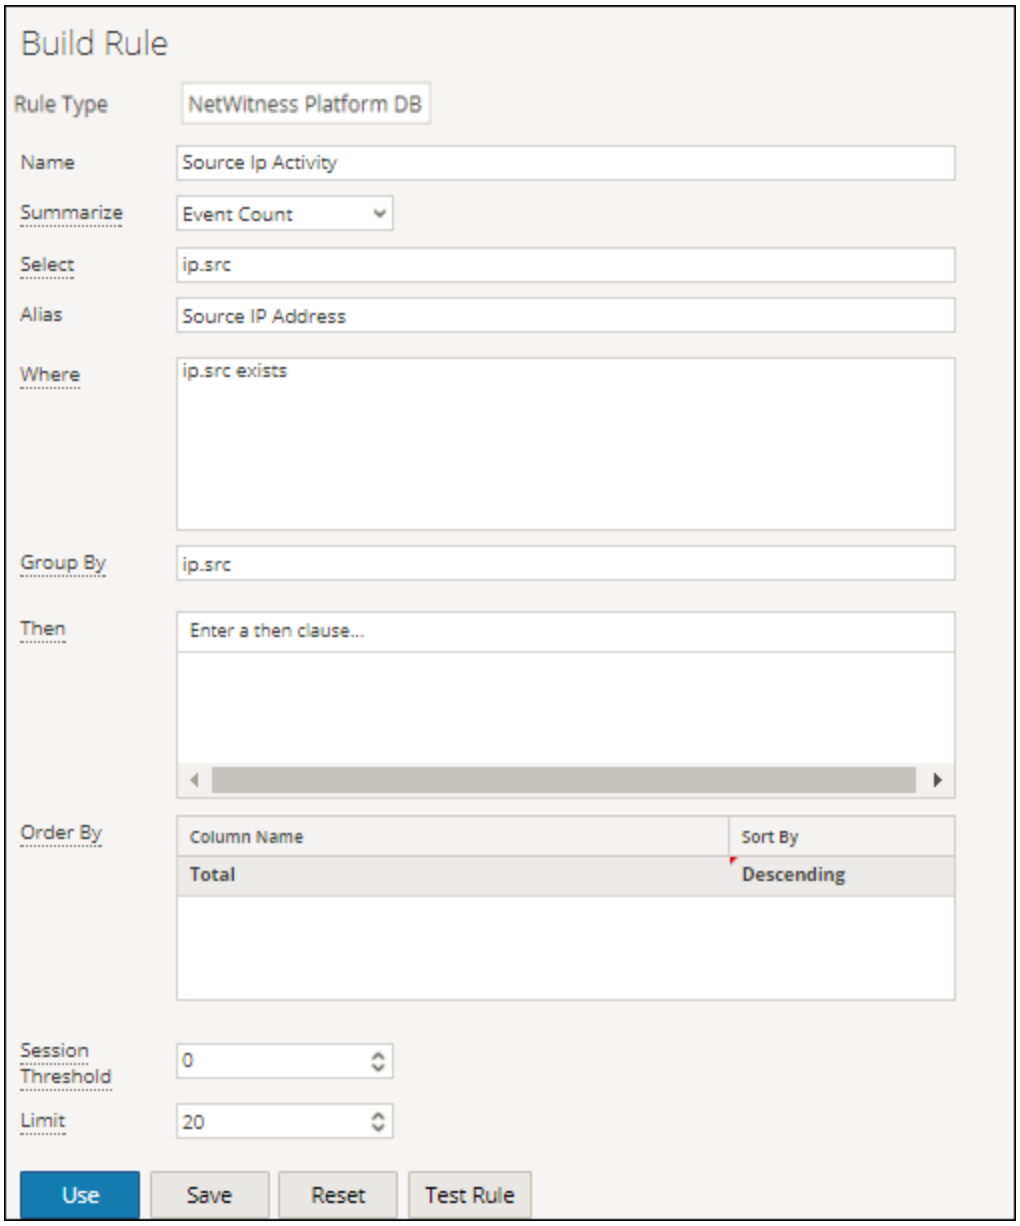
- NWDB Rules
- Timeout Configuration for NWDB Rules
- Lookup and Add rule action
- List value Reports

NWDB Rules

If the reporting entities such as report, alert, or chart contain NWDB rules (in most cases where the query contains Group By) takes a long time to execute, you may do the following:

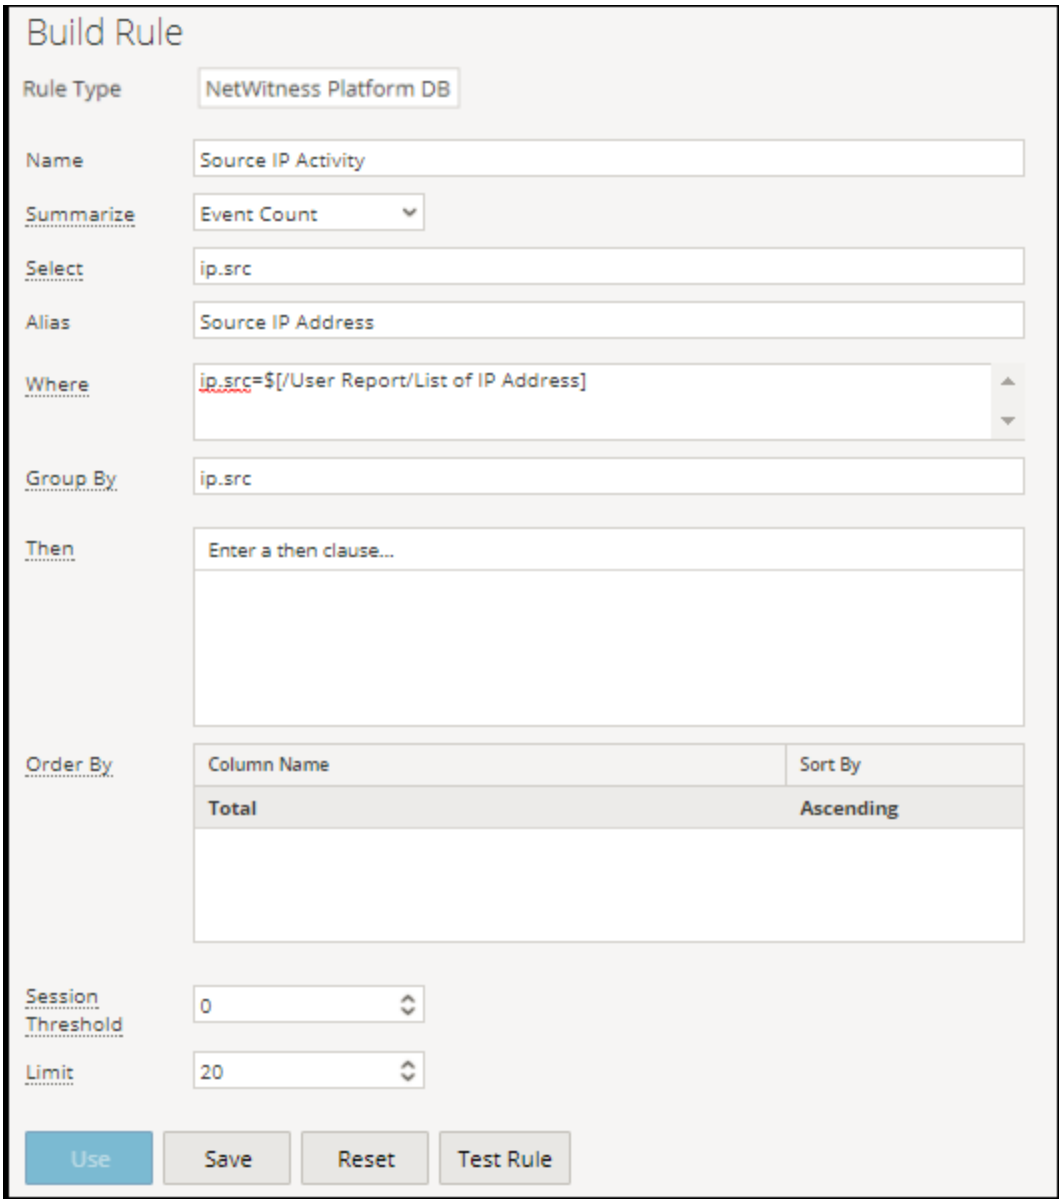
1. Refine the Where clause:

   You may limit the number of sessions scanned by using or refining the Where clause (especially when you use the Group By option). For example, consider the following rule.



If you use a Where clause as mentioned above, the number of sessions aggregated is huge. To avoid this, you can filter only required sessions by specifying the list of IP addresses or creating a List (list of IP Address) that contains relevant IP addresses.

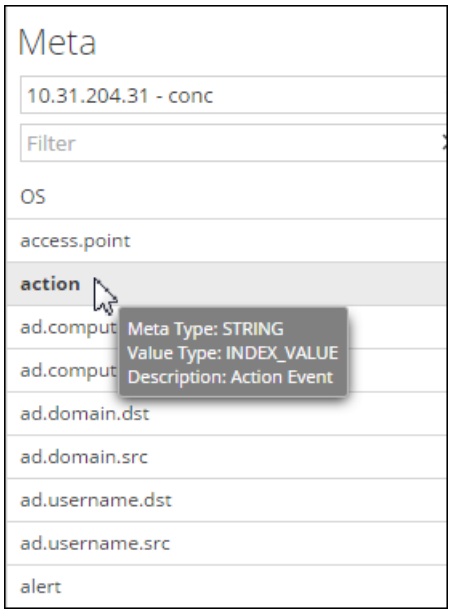
Note: The NWDB rule where clause is appropriately quoted if the syntax has an invalid quote. For example, in case of an invalid meta, or missing separator, the status and the error message is updated appropriately.



2. Using indexed Meta keys in the Where clause:

   To understand if the Meta is indexed or not, mouse hover the Meta List present on the right panel. If the Value Type is INDEX_VALUE, then the Meta is indexed. The Value Type is INDEX_KEY or INDEX_NONE if the Meta is not indexed.

   Below is a snapshot of a Meta key that is indexed.



3. Configure the Timeout option:

   If the query is taking a long time and fails due to timeout issues, you can configure the timeout for the NWDB rule executions. For more information, see below section "Timeout Configuration for NWDB Rules".

4. Schedule the queries to run at different times:

   If multiple query aggregates are concurrently executed and timeout occurs, you may schedule the queries to run at different times without much overlap.

Timeout Configuration for NWDB Rules

Note: It is a good practice to check the statistics of the Reporting Engine and the NWDB data sources before you make any changes to the configuration. For more information, see the "Monitor Service Details" topic for Reporting Engine and "Monitor System Statistics" topic in the *System Maintenance Guide*.

If NWDB rule execution fails due to timeout, you may get the following errors on the View a Report panel:

- Reporting Engine timeout error

  “Data source ‘`10.31.x.x Concentrator`' did not respond within the configured time `30` minutes for the ‘`/sdk/values`' request.”

- NWDB timeout error

  "Error occurred while fetching data from source '`10.31.x.x Concentrator`'. `{Timeout message from NWDB}`”

In such cases, you need do the following:

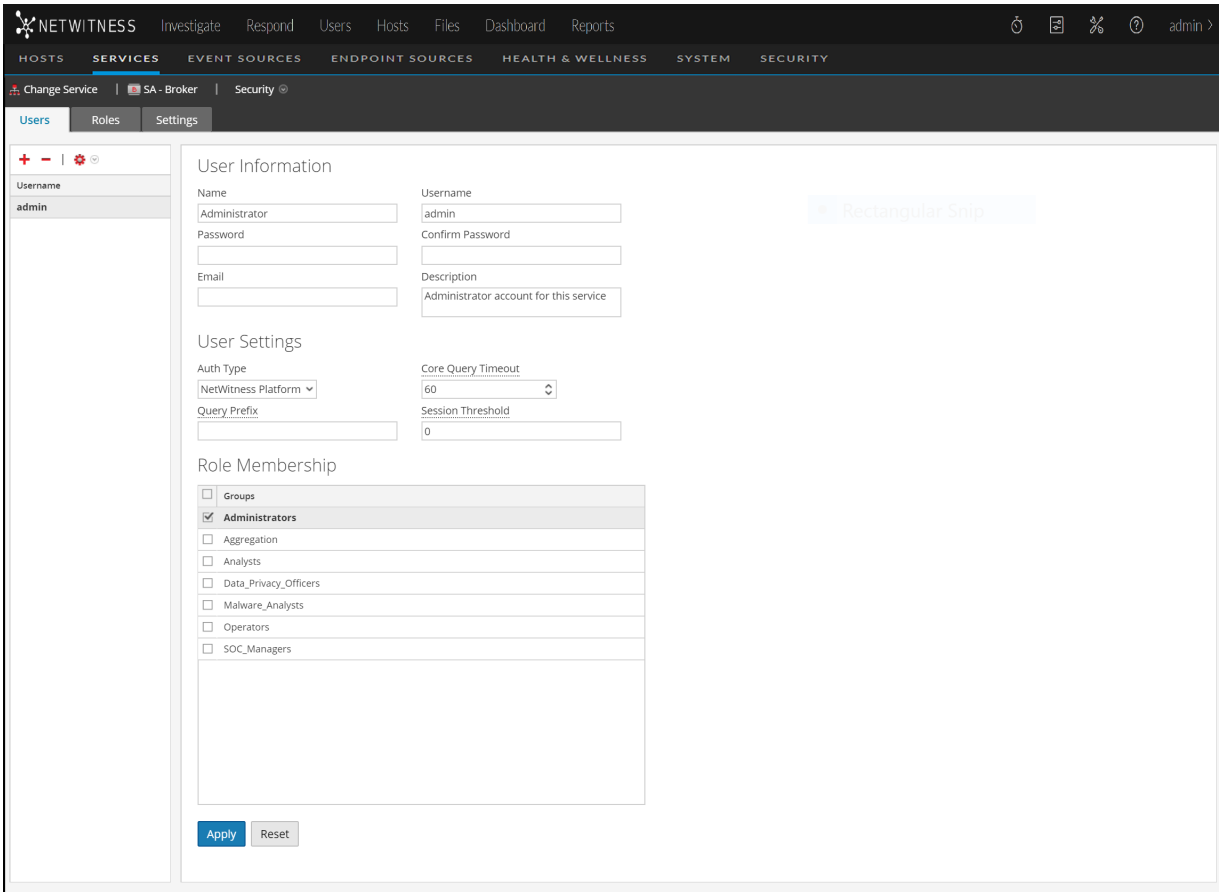
- Reporting Engine timeout

  In case of Reporting Engine timeout, you may set the timeout to a longer duration so the long running queries can be executed. For more information on setting the `NWDB Queries Time Out` and `NWDB Info Queries Time Out` option for the Reporting Engine, see "Step 2. Configure Reporting Engine Settings" topic in the *Reporting Engine Configuration Guide*. RSA recommends you set the `NWDB Query Time Out` to zero minutes (implies no timeout) and `NWDB Info Queries Time Out` to 60 minutes.

- NWDB timeout

  In case of NWDB timeout, you may need to configure the `query.level.timeout` and `max.concurrent.queries` parameters for the NWDB data source based on the recommendations in the *Core Database Tuning Guide* to fine tune the queries.

  The following figure is an example of Explorer view where you can set the parameters for NWDB data source.



- Schedule Reports at different times

  If the NWDB core devices are heavily utilized, you may schedule the reports to run at different times without overlap.

- Split the Report

If you have many rules in a Report, split the report into multiple reports with each report containing logical set of rules. If you have multiple rules, all rules will begin to execute at the same time based on available threads, therefore you may group the rules logically into separate reports.

Lookup_and_Add Rule Action

If a rule that consists of single or multiple `lookup_and_add` rule actions, takes a long time to execute the report, it is because each of the rule action triggers multiple lookup queries on the NWDB data source resulting in longer execution time.

To improve the report execution time, you may do the following:

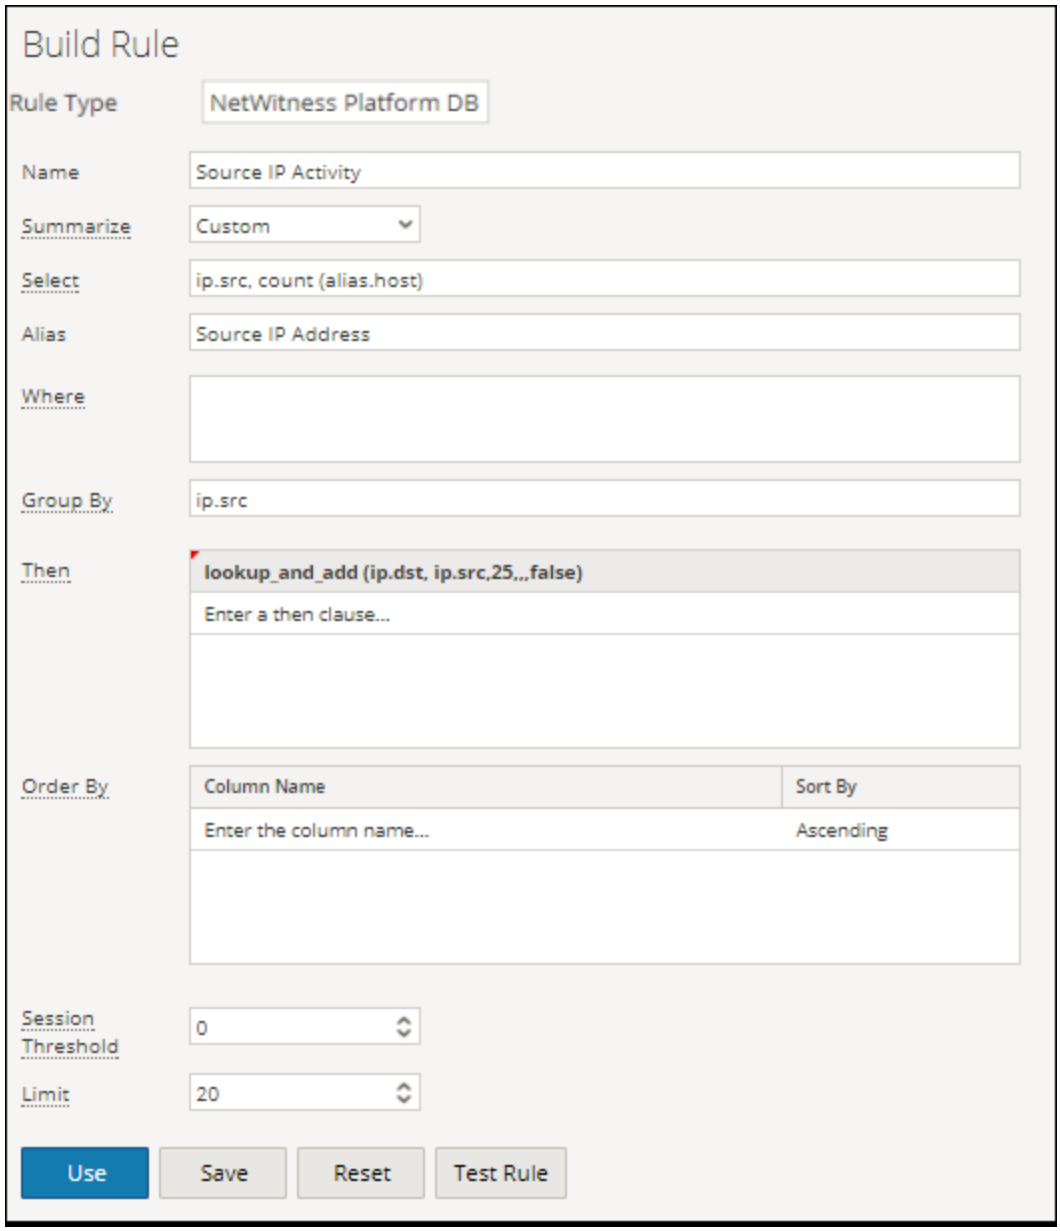
- Refine the Where clause in the following:
  - Rule that contains the `lookup_and_add` rule action
  - `lookup_and_add` rule action
- Set Limits

  You must set appropriate limits for the rule and rule actions. If the limit is high it will result in many queries being triggered and hence the report will take a long time to execute.
- Set the boolean aggregate parameter

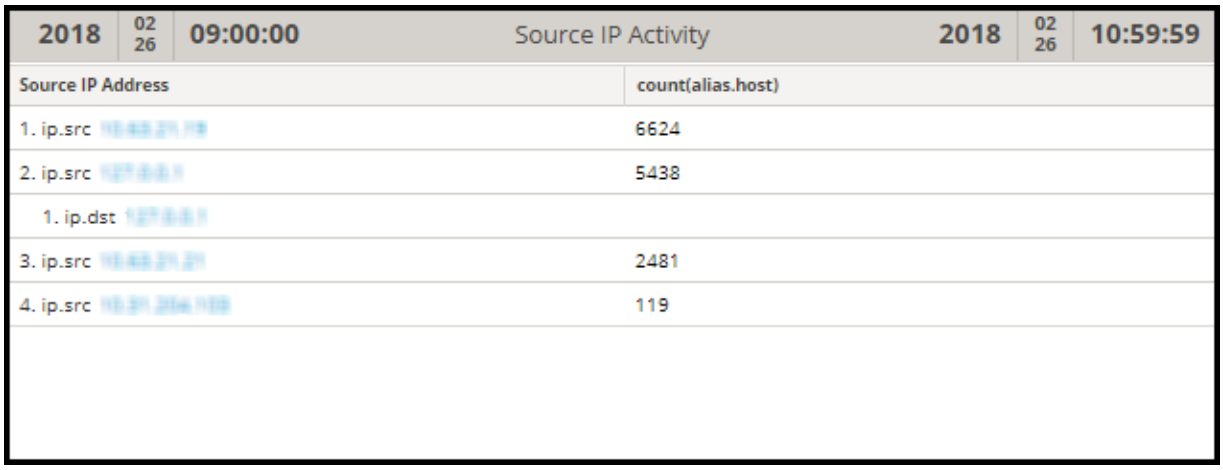
  If you do not want the aggregate value such as `sum(meta), count(meta)` etc. for the lookup values, set the boolean aggregate parameter to false in the `lookup_and_add` rule action. For more information, see the "NWDB Rule Syntax" section in Rule Syntax.

  ```
  lookup_and_add(string select, string field, int limit, boolean inherit,
  string extraWhere, boolean aggregate)
  ```

  Consider the rule with lookup_and_add rule action:



The output is displayed:

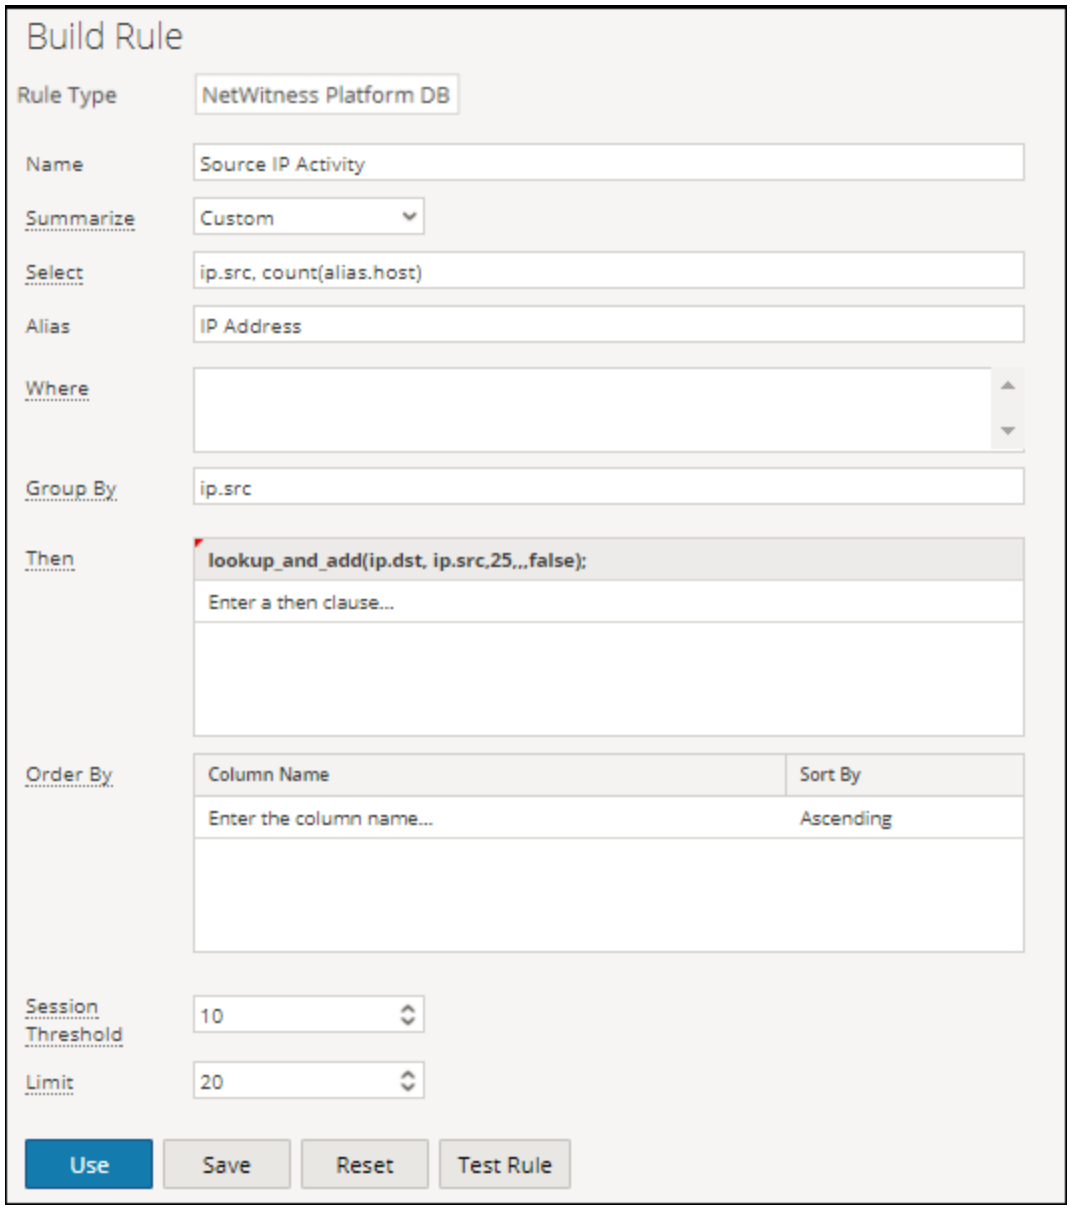


- Each `lookup_and_add` rule action triggers by default two concurrent lookup queries on the data source. NetWitness recommends that you retain the default setting, however if you want to increase the value you may want to ensure the value of `Max # of Concurrent LookupAndAdd Queries` parameter in Reporting Engine is less than the `Max Concurrent Queries` value in the NWDB data source configuration.

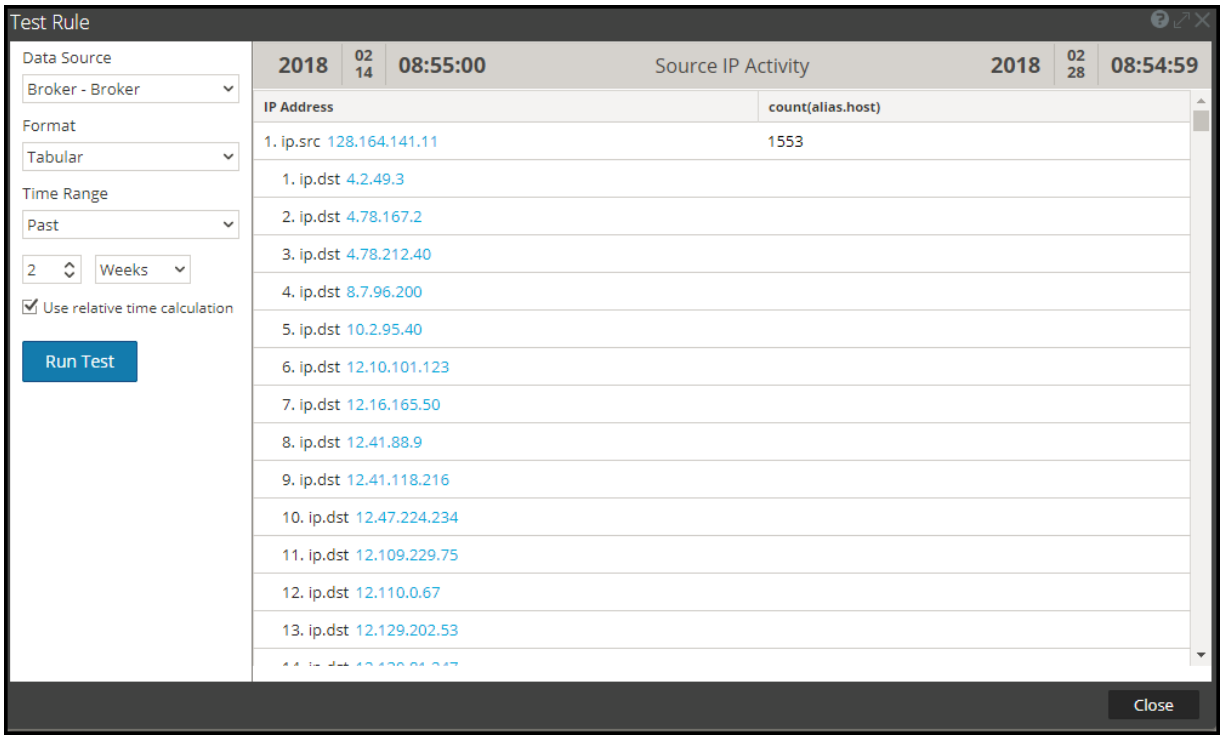
  If the NWDB data source is shared across other services, then you may retain a low value for the `Max # of Concurrent LookupAndAdd Queries` parameter in Reporting Engine as increasing it will impact the queries from other services. For more information, see "Reporting Engine General Tab" topic in *Reporting Engine Configuration Guide*.

- If you are interested only in unique values and not accurate aggregates, then set the `Session Threshold` to a non-zero value for the NWDB rule. For more information, see "Create a Rule Using NetWitness Data Source" section in Configure a Rule. The higher the value, the longer is the rule execution. If the value is set to zero it will take a longer time but will provide accurate aggregates.

  Consider a rule with lookup_and_add rule action and Session Threshold set to 10.



The output is displayed:



List Value Reports

Use a Refined List:

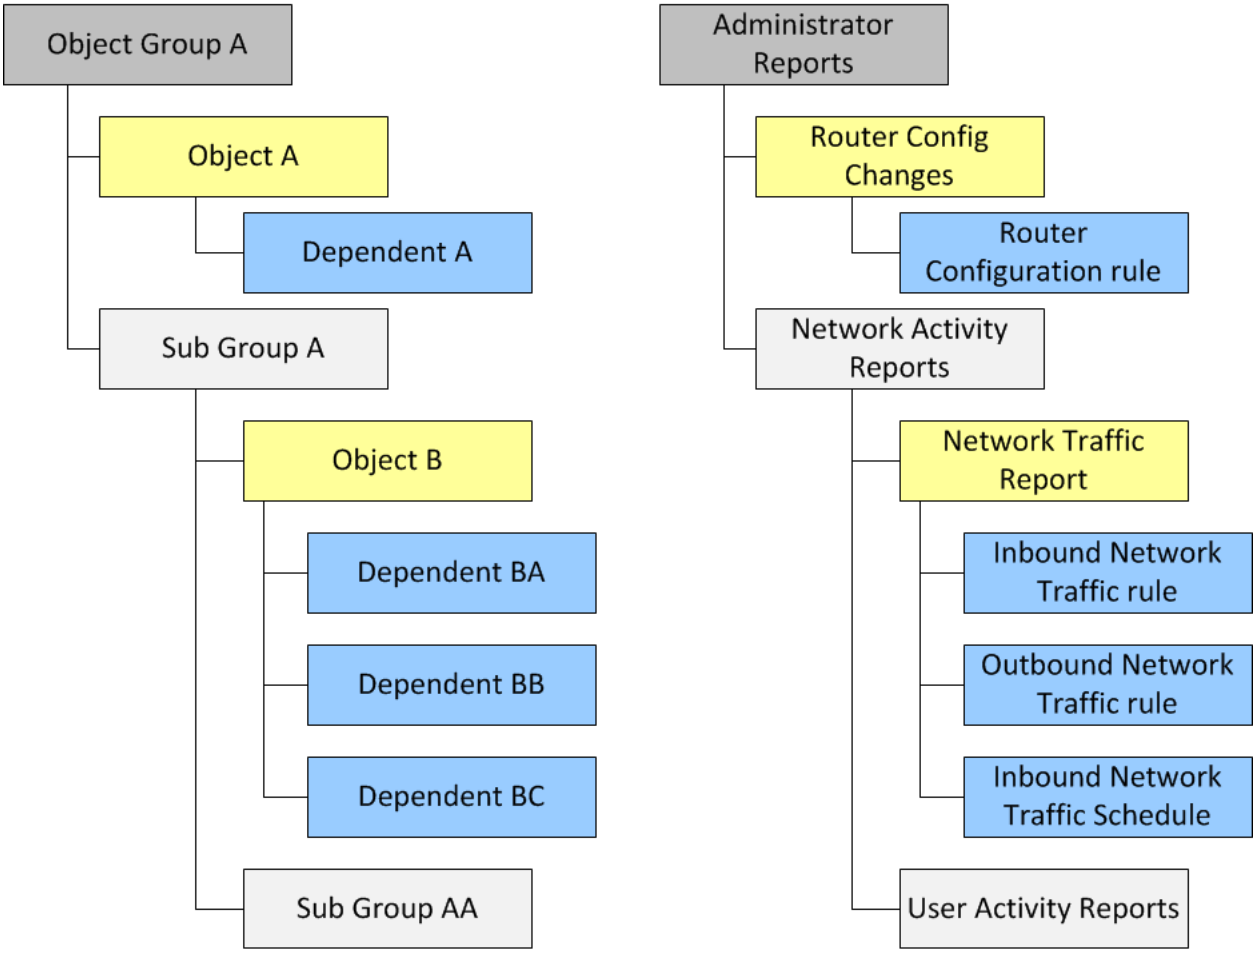
In case of List value reports (for any data source type), individual reports will be generated for each value in the list. Therefore, more the number of values in the list the longer the reports will take to execute. Hence, you must use a refined list to generate such reports.

Access Control for Reporting

Reporting Module provides you the option to set up access control for all the components in the module. In NetWitness, you can define different roles and specify the access control for each of the role from the System Security module. You can define the access control to be provided for the Reporting module for each role. For more information, see "Step 1: Review the Pre-Configured NetWitness Platform Roles" and "Step 2: (Optional) Add a Role and Assign Permission" topics in the *System Security and User Management Guide*.

In the Reports module, you can modify the role permissions or access to the following Reporting objects:

The Following is an example of the hierarchy of the object groups, objects and dependents. This is an illustration of the Report Groups and Reports hierarchy.



Report Groups and Reports Hierarchy

Permission for Object Groups

- You must have the Read & Write permission to set the permissions for the Object Group, Objects, or Dependents. The dependents with “No Access” permission are grayed out and dependents with “Read-Only” permission are indicated with an icon.
- When you set the permission for the Object Group, the Objects and Dependents in the Object Group do not inherit the permission automatically. You must select the "Apply these permissions to sub-groups and <Objects> in this group" option to achieve this. For example, if you do not want Operators roles to access reports in Report Group A, then you must set the permission on Group A to No access for the Operator role and select the "Apply these permissions to sub-groups and Reports in this group" option.
- When you set the permissions for the Object Group and select the "Apply these permissions to sub-groups and <Objects> in this group" option, the dependents such as rules or schedules in the objects do not inherit the permissions automatically. You must use the "Apply Read-only permission to Rules in the <Object>" option to apply the permission to the rules.

- When you set the permissions for the Objects, you must ensure that the Objects in hierarchy should always have a permission that is less than or equal to the one above in the hierarchy for the permission to be applied. For example, if the reports in a Report Group have Read & Write permission and you apply a Read-Only or No Access permission at the Report Group level and select the "Apply these permissions to sub-groups and Reports in this group" option, then the permission on the rules will remain unchanged.
- The permissions are cascaded from top to down in the hierarchy and not vice-versa. For example, if you apply a permission to a rule, it does not change the permission of the Report that contains the rule.

Permission for Objects or Dependents

- You must have the Read & Write permission to set the permissions for the Objects or Dependents.
- You can specify the permission for multiple objects at once instead of setting the permission for each object.
- When you set the permission for the Object, the dependents in the Object do not inherit the permission automatically. You must select the "Apply Read-only permission to Rules in the <Object>" option to achieve this.

When you apply the permission to dependents the permission is applied based on the existing permission for the role. For example, consider an Analyst and a Operator with the following permissions for the different dependents (Report A object has Rule AA, Rule AB, and Rule AC as dependents).

Object or Dependent	Analyst	Operator
Report A	Read & Write	No Access
Rule AA	Read & Write	No Access
Rule AB	Read and Write	Read and Write
Rule AC	Read-only	No Access

When the Analyst applies a Read & Write permission for the Operator role and selects the option "Apply Read-only permission to Rules in the <Object>", the permissions is set for the different dependents as follows:

Modify the Permissions

- **Group Level:** Set the permissions at the Object Group level and for all the object and entities in the Group. For example, if you have 80 reports in the Administrators Reports group and you do not want anyone except the Administrator to add or modify these reports, you can set the permission for all the other roles at the group level to Read-Only and select the option to apply it to all the reports and sub-groups in the report group.
- **Multiple Objects:**Select multiple objects and specify the access for all the selected objects. For example, if you have 10 reports in the Network Traffic sub group with sensitive information that you

do not want anyone to access, select the 10 reports and then set the permission for all the roles as "No Access".

- **Single Object:** Select only the object and specify the permission. For example, select the Network Traffic Report and specify the Read-Write permission for the Security Analyst role or select the Login Failure Alert and specify the Read-Write permission for a Security Analyst role.

Object or Dependent	Operator (Before Permission is applied)	Operator (After Permission is applied)
Report A	No Access	Read & Write
Rule AA	No Access	Read-only
Rule AB	Read and Write	Read & Write
Rule AC	No Access	Read-only

Roles and Permissions for Reporting Module

Although NetWitness has five pre-configured roles, you can add custom roles. For example, in addition to the pre-configured Analysts role, you can add custom roles for AnalystsEurope and AnalystsAsia.

Role	Permission
Administrators	Full system access
Operators	Access to configurations but not to data
Analysts	Access to data but not to configurations
SOC_Managers	Same access as Analysts and an additional permission to handle incidents
Malware_Analysts	Access to malware events only

Depending on the user role, you can set the following access permissions to access the Reporting module components (Rules, Reports, Charts, Alerts, Lists):

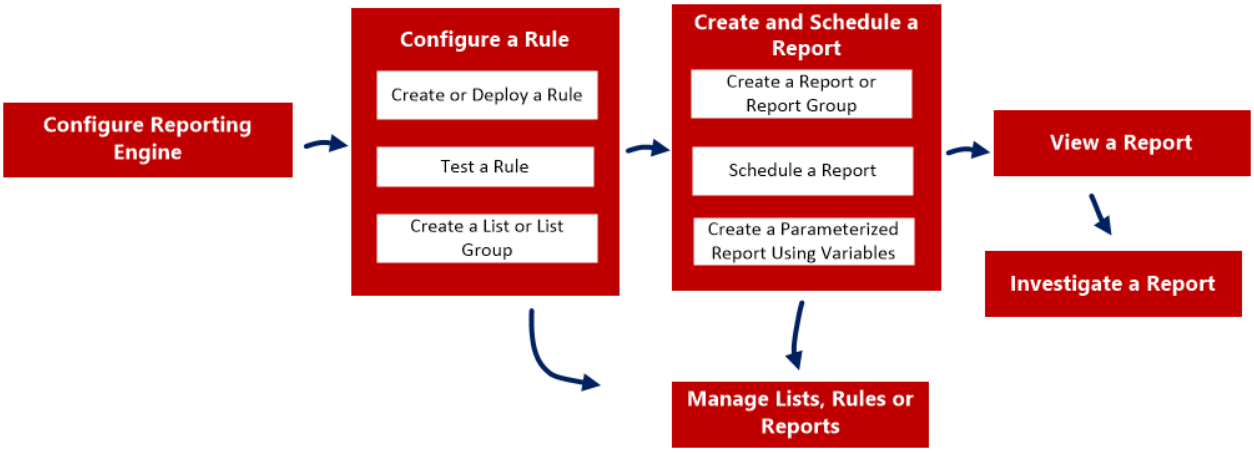
- Create
- Delete
- Export
- Manage
- View

Note: You must enable all these permissions for a user role to be able to define, delete, manage and view each of the Reporting modules. You must also have appropriate permissions for the data source to be listed, while defining the reports, charts, or alerts. For more information, see "Configure Data sources Permissions" topic in the *Reporting Engine Configuration Guide*.

For a detailed list of permissions and how to add a role and assign permissions, see "Role Permissions" and "Step 2. (Optional) Add a Role and Assign Permissions" topics in the *System Security and User Management Guide*.

Configure and Generate a Report

This figure is an overview of the entire process of configuring and generating a report.



To configure and generate a report, perform the following tasks:

1. Configure Reporting Engine - You must configure the Reporting Engine before you can configure and generate a report. You must also specify the data source in the Reporting Engine from which the data is extracted. For more information on how to configure Reporting Engine, see "Configure Reporting Engine" topic in the *Reporting Configuration Guide*.
2. Configure a Rule
3. Create and Schedule a Report
4. View a Report
5. Investigate a Report
6. Manage Lists, Rules or Reports

Configure a Rule

You can create a new rule or deploy an existing rule from the Live Services which can be used in a report. You can use different conditions to refine the data or information in the data sources such as:

- Select clause
- Where clause
- Group By
- Order By and so on

For example, you can write a rule to view the top 20 web addresses that the users visit daily.

You can create different type of rules using different data sources. Based on your requirements you can select any of the following options to create a rule:

- Create a Rule Using NetWitness Data Source
- Create a Rule Using Warehouse Data Source
- Create a Rule Using Respond Data Source

You can also use a list in a rule to refine a search result from the data source. Once a rule is created you can test a rule to see the results returned by the rule.

Create a Rule Group

To create a rule group or rule sub-group, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Do one of the following.
   - To define a rule group:
     a. In the **Rules Groups** Panel, click  .
     The new rule group is added to the Rule Groups panel.
     b. Enter the name for the rule group and press ENTER.
   - To add a rule sub-group:
     a. In the **Rules Groups** panel, select the rule group to which you want to add a sub-group.
     b. Click  .
     The new rule sub-group is added to the rule group.
     c. Enter the name for the rule sub-group and press **ENTER**.

Create a Rule Using NetWitness Data Source

You can create a rule to fetch data or events from a NetWitness data source. The same procedure is used to define a rule to fetch data or events from an Archiver data source.

The Archiver data source can be added in the Services Config View of the Reporting Engine. For more information, see "(Optional) Add Archiver as a Data Source to Reporting Engine" topic in the *Archiver Configuration Guide*.

Prerequisites

Make sure that you understand how custom meta keys are created using custom feeds. For more information, see "Create Custom Meta Keys using Custom Feed" topic in the *Decoder and Log Decoder Configuration Guide*.

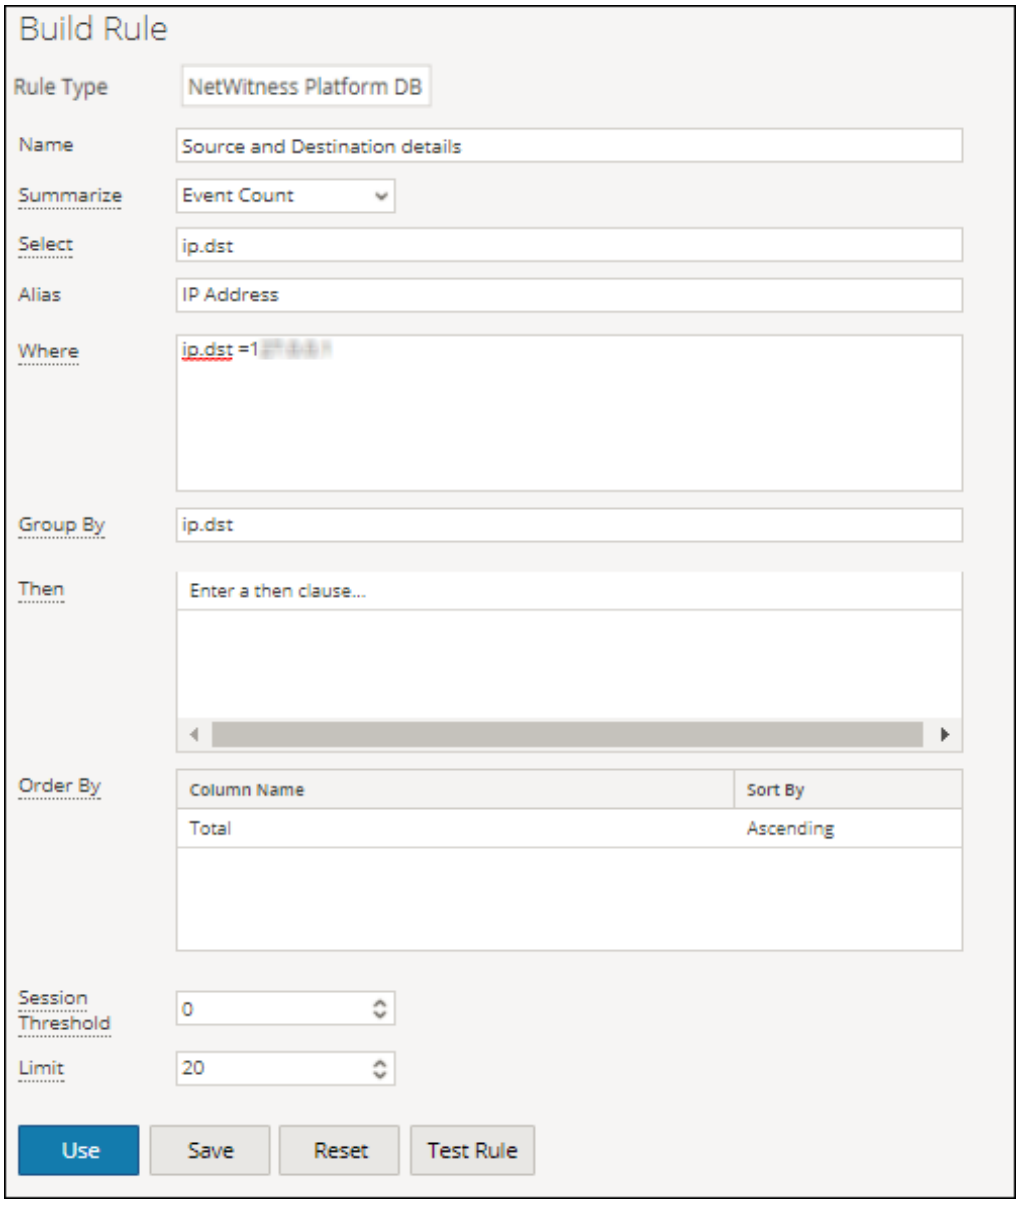
To create a rule to fetch data or events from a NetWitness Data Source, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** toolbar, click + > **NetWitness Platform DB**.

   The Build Rule view tab is displayed.



3. In the **Rule Type** field, **NetWitness Platform DB** is selected by default.
4. In the **Name** field, enter a name that is used to Identify or label the rule in alerts and reports.
5. The **Summarize** field determines the type of summarization or aggregation for the rule. Based on the type of rule to be defined, you must select one of the following:
   - To define a **Non-Aggregate** rule without any grouping, select: **None**
   - To define an **Aggregate** rule with special aggregation like the collection (sessions/events/packets) related aggregates, select one of the following:
     - Event Count
     - Packet Count
     - Session Size

- To define an **Aggregate** rule with meta values and custom aggregates like sum(), count(), and so on, select: **Custom**

  Choosing 'Custom' in the **Summarize** field enables you to define aggregate function of your choice in the *Select* clause. For example, select ip.src, countdistinct(ip.dst), distinct(ip.dst). The supported aggregate functions are:

  - sum (<meta>)
  - count(<meta>)
  - countdistinct(<meta>)
  - min(<meta>)
  - max(<meta>)
  - avg(<meta>)
  - first(<meta>)
  - last(<meta>)
  - len(<meta>)
  - distinct(<meta>)

  For more detailed information about Aggregate and Non-aggregate rule, see "NWDB Rule Syntax section" in Rule Syntax.

6. In the **Select** field, enter a meta or select a meta from the list of available meta types provided in the Meta Library. For more information, see "Meta Panel" in Build Rule View. The meta name to fetch raw log is raw. raw can only be used in the **Select** field. It cannot be used in the **Where** and **Then** fields. Multiple aggregate functions are supported for Custom aggregate rule in the **Select** field. For example, Select: *ip.src, username, service, distinct(country.src), sum(payload).*
7. In the **Alias** field, enter the alias name for columns used in the Select clause.
8. In the **Where** field, enter a meta or select a meta from the list of available meta types and use the operators to construct the Where clause for the base query criteria.
9. The **Group By** field is a read-only field which gets populated with meta that are defined in the Select clause. For a Non-Aggregate function, this field is not visible. A maximum of six meta are supported in the **Group By** field.

   Note: In earlier versions of NetWitness, only one meta was supported for Custom aggregate rule in the **Group By** clause. From now, a maximum of six meta are supported in the **Group By** clause.

10. In the **Then** field, enter the rule actions that manipulate the original result set of a rule in order to make the output in a report more concrete or add additional functionality other than querying data and displaying it, for example, creating a feed from the results. For a complete list of available rule actions, see "NWDB Rule Syntax" in Rule Syntax.

    Note: When a rule is executed for an Archiver data source, it is recommended not to use query intensive rule actions such as lookup_and_add() and show_whats_new().

11. In the **Order By** field, perform the following:

a. In the **Column Name** column, enter the name of the columns by which you want to sort the results. By default, the value is empty. The value gets populated based on the value selected in the **Summarize** field.

- For Summarize 'None', if no **Order By** is selected, then by default it is ordered by session or collection time.
- For other Summarize values, the default sorting is based on the first 'group by' meta selected when no 'order by' is defined. For Event Count, Packet Count, and Session size, the accepted values are Total and Value.

b. In the **Sort by** column, select one of the following ways to sort the results:

- Ascending Order
- Descending Order

12. In the **Session Threshold** field, enter the optimization setting to stop scanning the matching sessions for each possible unique value for the selected meta. The threshold is an integer between 0 (default) and 2147483647.

Note: This is applicable to only NWDB Aggregate rules. If the default value is specified, all the matching sessions will be scanned and the accurate value will be returned. A higher session threshold allows accurate counts for a value. However, this causes longer rule execution time. For example, consider you set the Session Threshold as 1000 for ip.src. If there are 5000 matching sessions then for a particular ip.src value which is present in more than 1000 sessions, NWDB stops the scan after 1000 sessions and returns the extrapolated aggregate value. This optimizes the query execution time. If the value is present in less than 1000 sessions, then the actual value is returned.

13. In the **Limit** field, enter the limit to be put on the query while fetching data from the database. If a result set is sorted by event count, packet count, or session size, the limit represents the top (or bottom) N values to be returned. If the result set is not sorted, the first N values are returned.
14. Click **Save**.

Note: Unlike parsed meta, raw logs are fetched from decoders. When both raw log and parsed meta are queried in a single rule, due to different retention periods, parsed meta might be available and raw logs missing in the same session. So the result will have parsed meta values and empty raw value for those sessions. For example, for the rule Select **ip.src**, **ip.dst**, **service**, **username**, **raw**, the parsed meta might be populated and the **raw** meta remains empty for a few sessions.

Create a Rule Using Warehouse Data Source

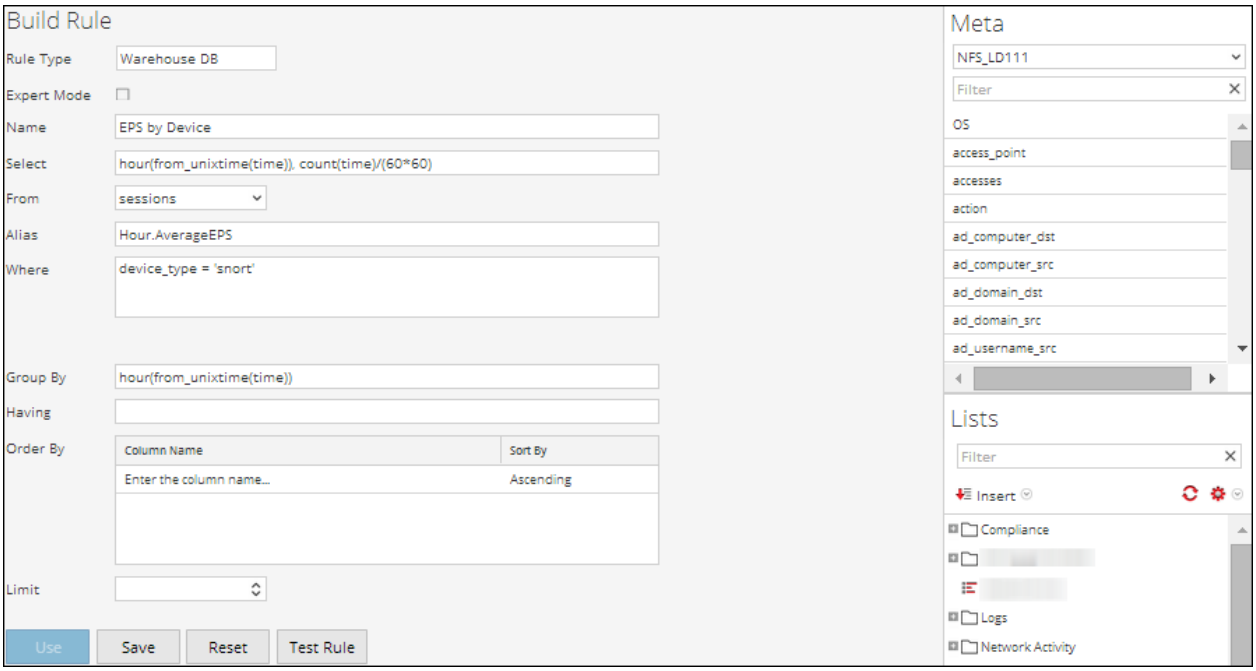
You can create a rule to fetch data or events from a Warehouse event source. You can define the rules in two modes:

- Default Mode
- Expert Mode

Default Mode

In Default Mode, you can create rules containing simple SQL like HIVE queries that contain clauses like Select, Where, Group By, and Having. By default, you can create rules to query sessions or raw logs. For more information on "Simple query syntax and examples", see Warehouse DB Simple Rules Syntax.

The following figure is an example of the **Build Rule view** that displays when you select **Warehouse DB** for **Rule Type** without the Expert Mode selected.

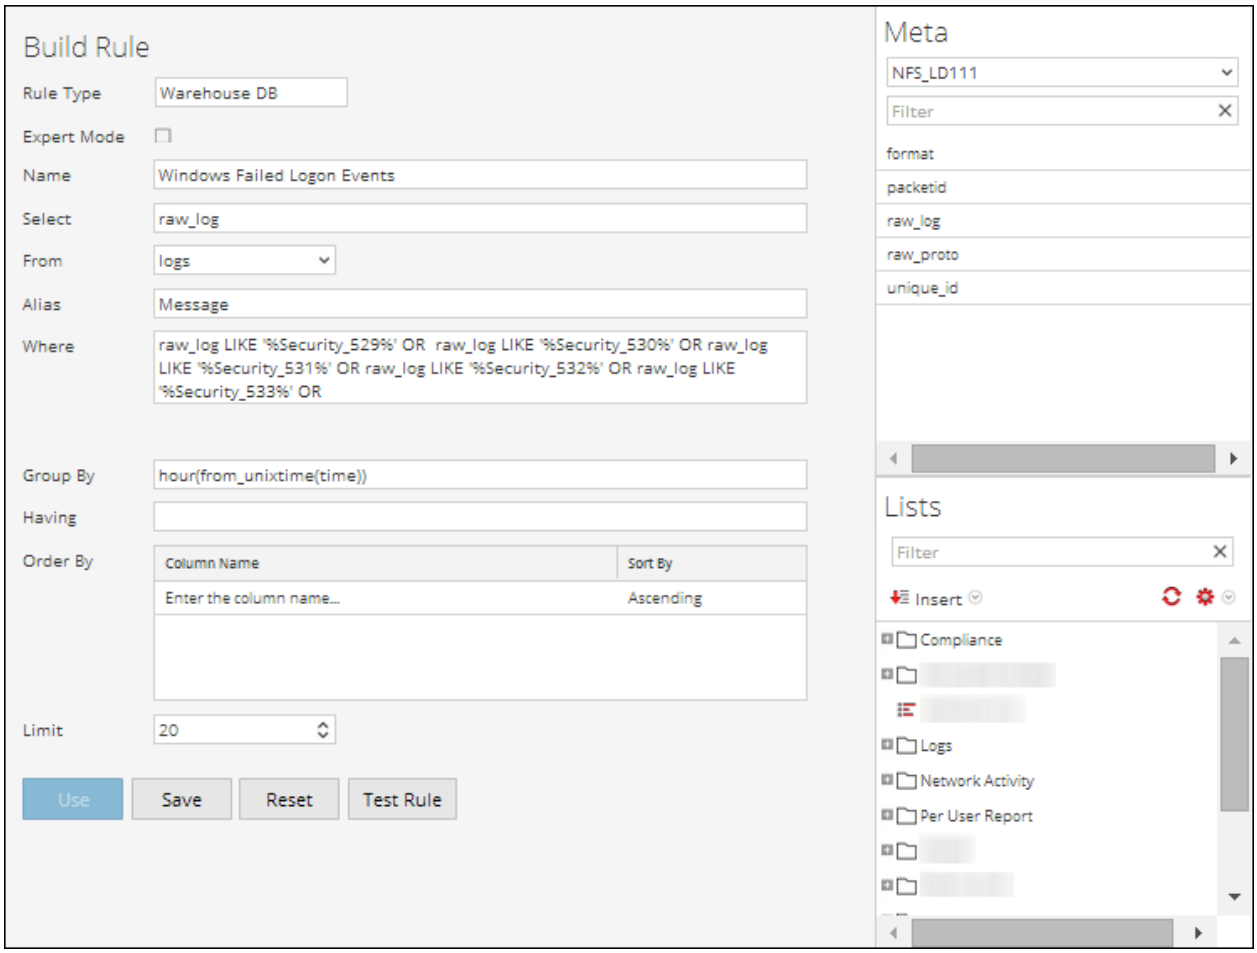


Querying Raw Logs

The raw log format is used in the select or where clause to query for raw logs.

Note: The time range that you can specify in your query is a day (24 hours). If you have specified a time range less than a day in your query, the result set contains data of at least a day (24 hours).

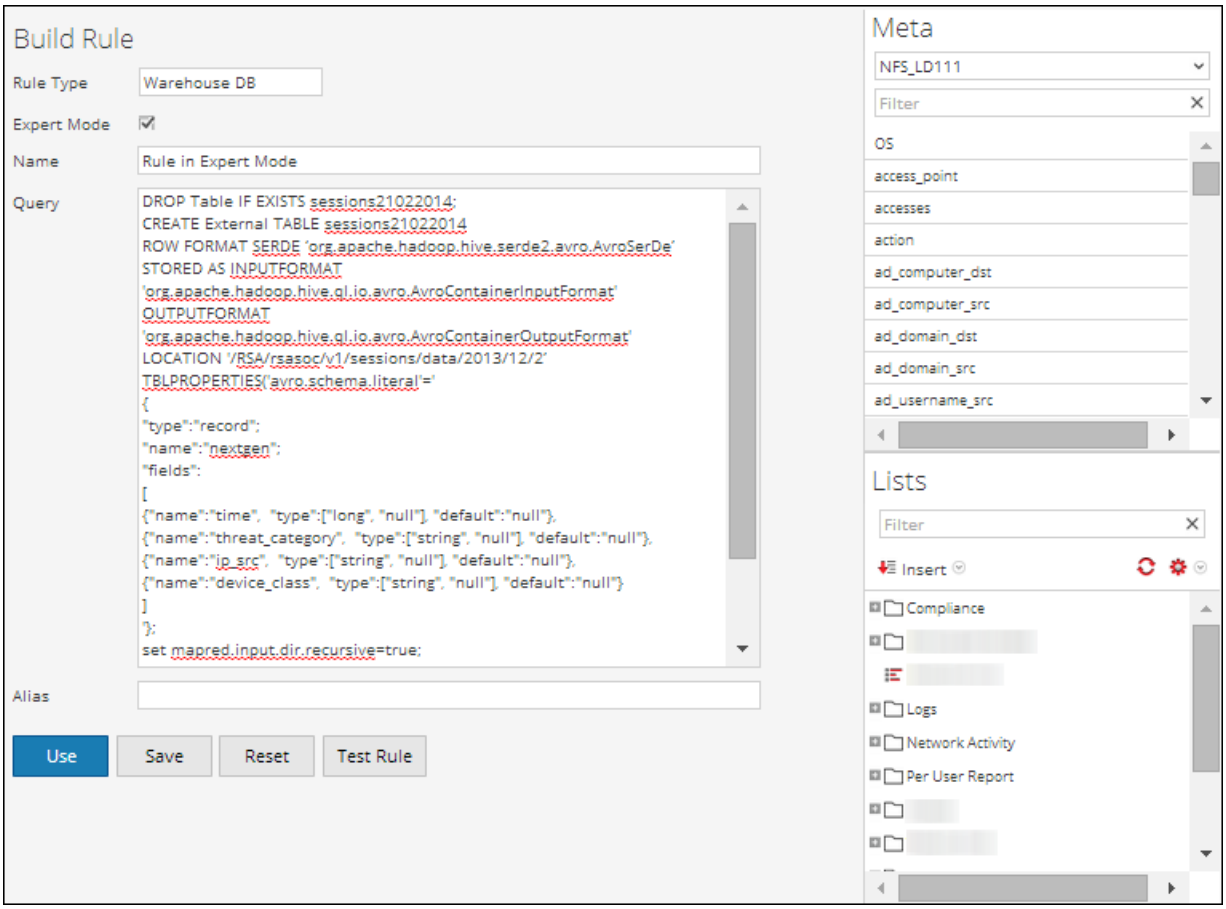
The following figure is an example of the **Build Rule view** that displays when you select **Warehouse DB** for **Rule Type** and create a rule for querying raw logs.



Expert Mode

Advanced rules are defined using complex HIVE queries created using the clauses DROP, CREATE, and so on. Unlike simple rules, we always insert the results into a table. For more information on "Advanced HIVE query language", see *HIVE language manual*.

The following figure is an example of the Build Rule view that is displayed when you select **Warehouse DB** for **Rule Type** with Expert Mode selected.



If you want to generate a report for a specific time range, you need to manually define the time range in the query using the following two variables:

- ${report_starttime} - The starting time of the range in seconds.
- ${report_endtime} - The ending time of the range in seconds.

For example, **`SELECT col1, col2 FROM custom_table WHERE timecol >= ${report_starttime} AND timecol <= ${report_endtime};`**

Note: By default, Reporting Engine treats `${keyword}` as a variable. If you want to specify HIVE variables, you must mention the complete syntax of a variable. For example, `${hiveconf:hive.exec.scratchdir}`.

Prerequisites

Make sure that you understand how custom meta keys are created using custom feeds. For more information, see "Create Custom Meta Keys using Custom Feed" topic in the *Host and Services Configuration Guide*.

To create a rule to fetch data or events from a Warehouse data source, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** toolbar, click  > **Warehouse DB**.

   The Build Rule view is displayed.

3. In the **Rule Type** field, **Warehouse DB** is selected by default.

   If you are defining the rule in Default mode, perform the following:

   a. In the **Name** field, enter a name that is used to Identify or label the rule in alerts and reports.

   b. In the **Select** field, enter a meta or select the meta from the drop-down or select a meta from the list of available meta types provided in the Meta Panel. For more information, see " Meta Panel" in Build Rule View.

   c. In the **From** drop-down menu, select one of the following:

      - Session
      - Logs

   d. In the **Alias** field, enter the alias name for columns used in the Select clause.

   e. In the **Where** field, enter a meta or select a meta from the list of available meta types provided in the Meta Panel. The Where clause provides the base query criteria for the rule.

   f. In the **Group By** field, enter the meta selected in the Select clause, so that the result set is grouped based on the meta.

   g. In the **Having** field, enter the criteria to filter the result set for aggregated queries.

   h. In the **Order By** field, perform the following:

      1. In the **Column Name** column, enter the name of the columns by which you want to group the results.
      2. In the **Sort by** column, select one of the following ways to sort the results:

         - Ascending Order
         - Descending Order

   i. In the **Limit** field, enter the limit to be put on the query while fetching data from the database. If a result set is sorted by session count, packet count, or session size, the limit represents the top (or bottom) N values to be returned. If the result set is not sorted, the first N values are returned.

   j. Click **Save**.

4. If you are defining the rule in Expert mode, select the **Expert Mode** checkbox and perform the following:

   a. In the **Name** field, enter a name that is used to Identify or label the rule in alerts and reports.

   b. In the **Query** field, enter the Hive query statement to query the data source.

   c. In the **Alias** field, enter the alias name for columns used in the Select clause.

   d. Click **Save**.

Create a Rule Using Respond Data Source

You can create a rule to fetch incidents or alerts from a Respond data source.

Prerequisites

Make sure that you:

- Ensure Reporting Engine service is up and running.
- Ensure the Incident Management service is up and running. For more information, see "Configure a Database for the Respond Server Service" topic in the *NetWitness Respond Configuration Guide*.
- (Optional) Ensure the Event Stream Analysis service is up and running. For more information, see "Step 2. Configure Advanced Settings for an ESA Service" topic in the *ESA Configuration Guide*.
- (Optional) Ensure the Malware Analysis service is up and running. For more information, see " (Optional) Configure Auditing on Malware Analysis Host" topic in the *Malware Configuration Guide*.

Note: You need to configure any one of the services (Event Stream Analysis, Reporting Engine, Malware Analysis, or Endpoint) based on your requirement and the type of alerts or incidents you want to generate.

To create a rule to fetch data or events from a Respond Data Source, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** toolbar, click  > **Respond DB**.

   The Build Rule view tab is displayed.

3. In the **Rule Type** field, Respond is selected by default.
4. In the **Name** field, enter a name that is used to Identify or label the rule in alerts and incident reports.
5. The **Summarize** field determines the type of summarization or aggregation for the rule. Based on the type of rule to be defined, you must select one of the following:
   - To define a **Non-Aggregate** rule without any grouping, select **None**
   - To define an **Aggregate** rule with meta values and custom aggregates select **Custom**

     Choosing 'Custom' in the **Summarize** field enables you to define aggregate function of your choice in the *Select* clause based on the report type you have selected.

     For more detailed information about Aggregate and Non-aggregate rule, see Rule Syntax.
6. In the **From** field, based on the type of report output to be displayed, you must select one of the following:
   - Alert
   - Incident

- incidentStats
- incidentUserStats

7. In the **Select**field, enter a meta or select a meta from the list of available meta types provided in the Meta Library. For more information, see "Meta Panel" in Build Rule View. It cannot be used in the **Where** field. Only one aggregate function is supported at a time in a query.

   For example, the supported metas for alert are:

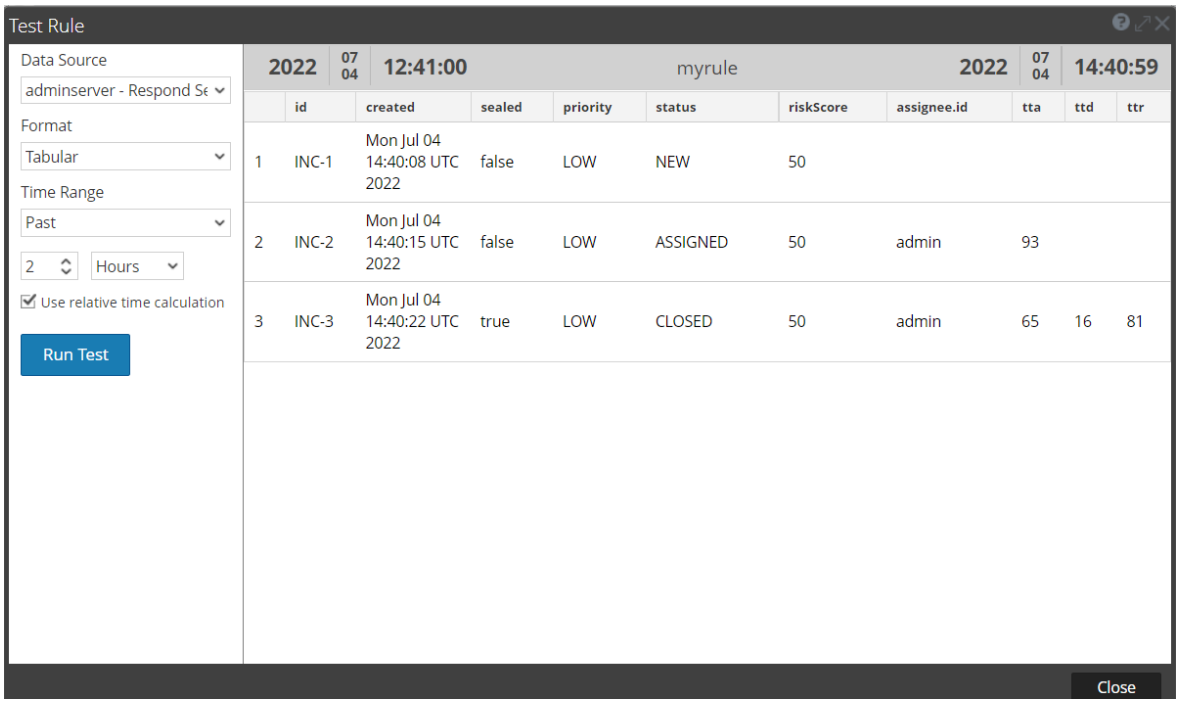
   - alert_host_summary
   - alert.name
   - alert.numEvents
   - alert.severity
   - alert.source
   - alert.timestamp
   - incidentCreated
   - incidentId
   - receivedTime

   For example, the supported metas for incident are:

   - categories
   - created
   - priority
   - riskScore
   - sealed
   - status
   - assignee.id
   - tta (for more information on this meta, see **View Basic Summary Information about the Incident** topic in the *Respond User Guide.*)
   - ttd (for more information on this meta, see **View Basic Summary Information about the Incident** topic in the *Respond User Guide.*)
   - ttr (for more information on this meta, see **View Basic Summary Information about the Incident** topic in the *Respond User Guide.*)

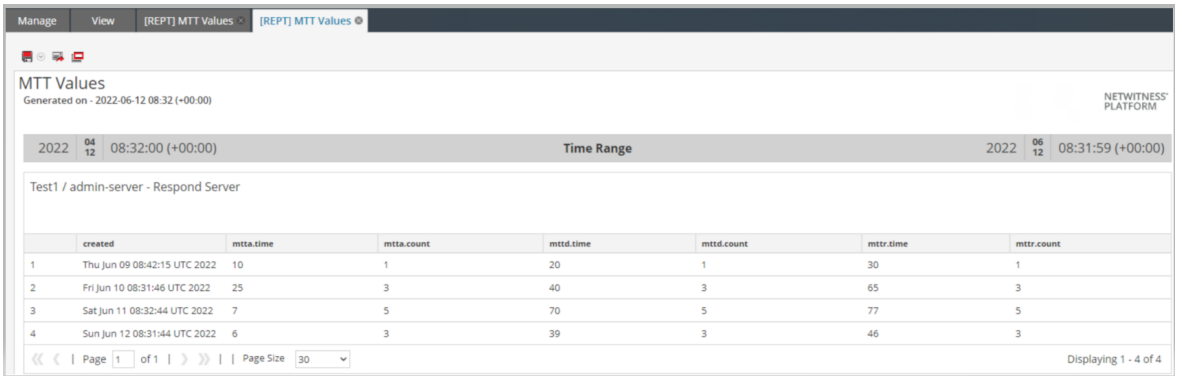
   Note: When an incident is assigned, tta and assignee.id metas are populated. Similarly, when the task assigned is completed and the incident is closed, ttd and ttr metas are populated. Refer to

the following figure.



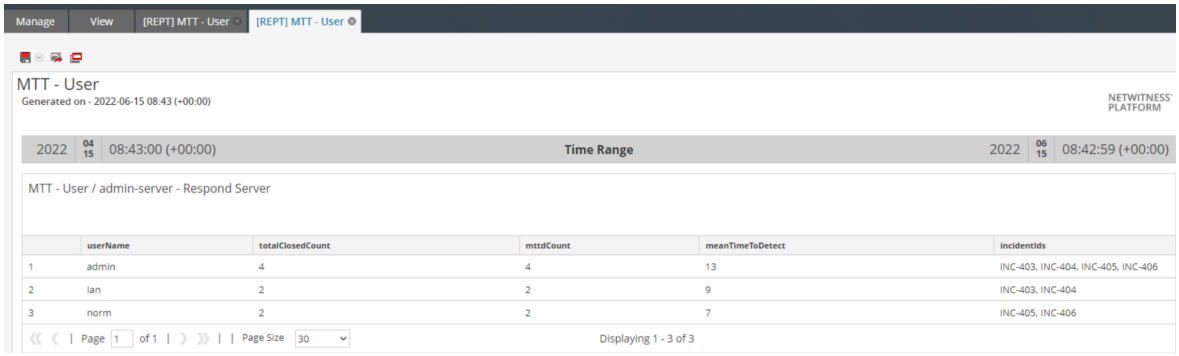
For example, the supported metas for incidentStats are:

- created
- mtta.time - This meta displays the average time taken to acknowledge the incidents in a single day.
- mtta.count - This meta displays the number of incidents acknowledged in a single day.
- mttd.time - This meta displays the average time taken to detect the incidents in a single day.
- mttd.count - This meta displays the number of incidents detected in a single day.
- mttr.time - This meta displays the average time taken to resolve the incidents in a single day.
- mttr.count - This meta displays the number of incidents resolved in a single day.



For example, the supported metas for incidentUserStats are:

- userName - This meta displays the assignee's or the user's ID for the associated user stats.
- totalClosedCount - This meta displays the total number of Incidents closed by the assignee till date.
- meanTimeToDetect - This meta displays the average time taken by the user to detect the incidents in the time range selected.
- mttdCount - This meta displays the count of incidents contributing to the MTTD value computed.
- incidentIds - This meta displays the list of incident IDs closed by the user during the time range selected.



For more detailed information, see "Aggregate and Non-aggregate rule" topic in the Rule Syntax.

8. In the **Alias** field, enter the alias name for columns used in the Select clause.
9. In the **Where** field, enter a meta or select a meta from the list of available meta types and use the operators to construct the Where clause for the base query criteria.
10. The **Group By** field is a read-only field which gets populated with meta that are defined in the Select clause. For a Non-Aggregate function, this field is not visible. A maximum of six meta are supported in the **Group By** field.
11. In the **Order By** field, perform the following:
    a. In the **Column Name** column, enter the name of the columns by which you want to sort the results.

    Note: by default the first meta in the select clause will be dispalyed.

    b. In the **Sort by** column, select one of the following ways to sort the results:
    - Ascending Order
    - Descending Order
12. In the **Limit** field, enter the limit to be put on the query while fetching data from the database. If a result set is sorted by the limit represents the top (or bottom) N values to be returned. If the result set is not sorted, the first N values are returned.
13. Click **Save**.

Deploy a Rule

In NetWitness you can deploy the selected rules on the service (for example, Reporting Engine), using the Deployment Wizard.

Prerequisites

Make sure that:

- The services on which you deploy a rule is up and running.
- The Live Services is configured.

To deploy a rule, perform the following:

1. Go to  **(Configure) > LIVE CONTENT**.
2. In the **Search Criteria** panel, search Live resources (for example, search for the **Application Rule** resource Type).
3. In the **Matching Resources** panel, select **Show Results > Grid**.
4. Select the checkbox to the left or the rules that you want to deploy.
5. In the **Matching Resources** toolbar, click .
6. Click **Next**.
7. Select the service on which you to deploy a rule (For example, Reporting Engine) and click **Next**.
8. Click **Deploy**.
   The rule is deployed successfully.

Use Meta Aliases for Reporting

When you refer to meta data in Reports and Charts, you can only view aliases for the meta names. These aliases makes them more understandable to a broader audience.

You cannot provide alias values for any meta in the WHERE clause because NetWitness uses the WHERE clause to fetch data from the data source (for example, in the Concentrator) and data sources do not support aliases. In other words, you cannot provide the alias value **HTTP** for the HTTP port # 80.

Note: * You cannot create aliases for meta other than the ones that have existing aliases by Reporting Engine. Also, the format of the aliases cannot be changed.
* Aliases are not supported for Alerts and CSV reports.

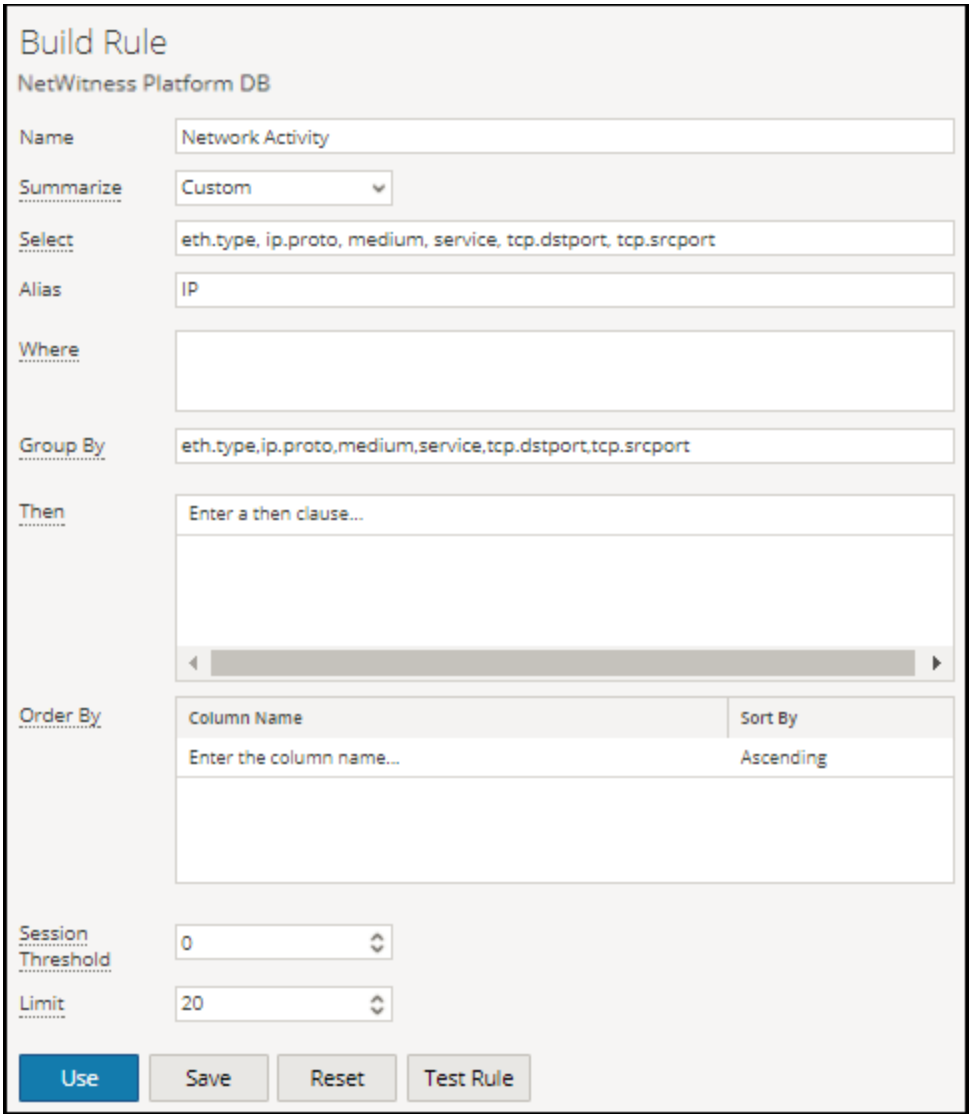
To use alias in a rule, perform the following:

1. Go to **Reports**.
   The Manage tab is displayed.
2. In the **Rules** panel, do one of the following:

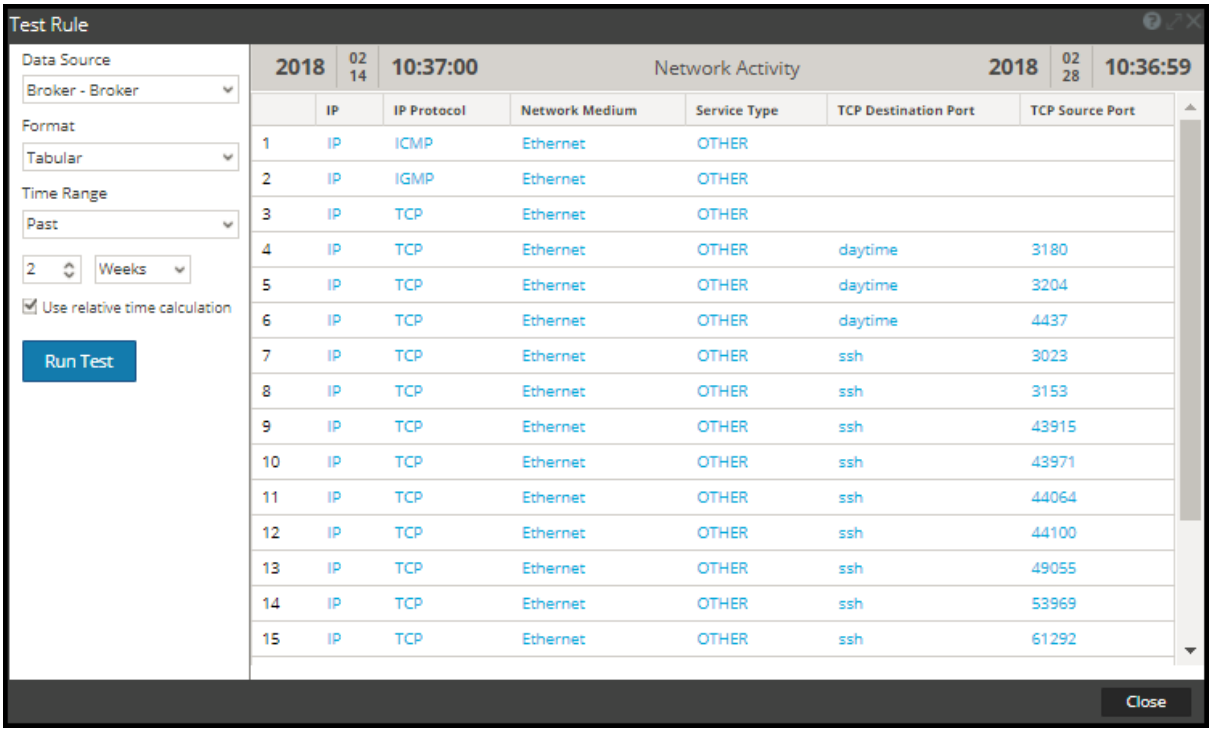
- Select a rule and click in the Rules toolbar.
- Click > **Edit**.

3. Specify the meta key with aliases in the **Select** field.

The following example specifies the **eth.type**, **ip.proto**, **medium**, **service**, **tcp.dstport**, and **tcp.srcport** meta in the Select field.



4. Click **Test Rule**.
   The following example displays the results under the **eth.type**, **ip.proto**, **medium**, **service**, **tcp.dstport**, and **tcp.srcport** alias columns that were specified in the **Select** field of the rule.



	IP	IP Protocol	Network Medium	Service Type	TCP Destination Port	TCP Source Port
1	IP	ICMP	Ethernet	OTHER		
2	IP	IGMP	Ethernet	OTHER		
3	IP	TCP	Ethernet	OTHER		
4	IP	TCP	Ethernet	OTHER	daytime	3180
5	IP	TCP	Ethernet	OTHER	daytime	3204
6	IP	TCP	Ethernet	OTHER	daytime	4437
7	IP	TCP	Ethernet	OTHER	ssh	3023
8	IP	TCP	Ethernet	OTHER	ssh	3153
9	IP	TCP	Ethernet	OTHER	ssh	43915
10	IP	TCP	Ethernet	OTHER	ssh	43971
11	IP	TCP	Ethernet	OTHER	ssh	44064
12	IP	TCP	Ethernet	OTHER	ssh	44100
13	IP	TCP	Ethernet	OTHER	ssh	49055
14	IP	TCP	Ethernet	OTHER	ssh	53969
15	IP	TCP	Ethernet	OTHER	ssh	61292

Alias Definitions

The alias files in this section are examples only and are based on current alias definitions in the Reporting Engine. NetWitness cannot modify these definitions in the Reporting Engine depending on the changes in the concentrator xml file. Since any changes in the Concentrator xml file are not reflected in the Reporting Engine.

The details of different meta are explained in each of the **meta.aliases**.

eth.type

```
ALIAS_FORMAT=$alias
 0=802.3
 257=Experimental
 512=Xerox PUP
 513=Xerox PUP
 1024=Nixdorf
 1536=Xerox NS IDP
 1537=XNS Address Translation (3Mb only)
 2048=IP
 2049=X.75 Internet
 2050=NBS Internet
 2051=ECMA Internet
 2052=CHAOSnet
 2053=X.25 Level 3
 2054=ARP
 2055=XNS Compatibility
 2076=Symbolics Private
 2184=Xyplex
 2304=Ungermann-Bass network debugger
 2560=Xerox IEEE802.3 PUP
```

```
2561=Xerox IEEE802.3 PUP Address Translation
2989=Banyan Systems
2991=Banyon VINES Echo
4096=Berkeley Trailer negotiation
4097=Berkeley Trailer encapsulation for IP
4660=DCA - Multicast
5632=VALID system protocol
6537=Artificial Horizons
6549=Datapoint Corporation (RCL lan protocol)
15360=3Com NBP virtual circuit datagram (like XNS SPP) not registered
15361=3Com NBP System control datagram not registered
15362=3Com NBP Connect request (virtual cct) not registered
15363=3Com NBP Connect repsonse not registered
15364=3Com NBP Connect complete not registered
15365=3Com NBP Close request (virtual cct) not registered
15366=3Com NBP Close response not registered
15367=3Com NBP Datagram (like XNS IDP) not registered
15368=3Com NBP Datagram broadcast not registered
15369=3Com NBP Claim NetBIOS name not registered
15370=3Com NBP Delete Netbios name not registered
15371=3Com NBP Remote adaptor status request not registered
15372=3Com NBP Remote adaptor response not registered
15373=3Com NBP Reset not registered
16972=Information Modes Little Big LAN diagnostic
17185=THD - Diddle
19522=Information Modes Little Big LAN
21000=BBN Simnet Private
24576=DEC unassigned
24577=DEC Maintenance Operation Protocol (MOP) Dump/Load Assistance
24578=DEC Maintenance Operation Protocol (MOP) Remote Console
24579=DECNET Phase IV
24580=DEC Local Area Transport (LAT)
24581=DEC diagnostic protocol (at interface initialization?)
24582=DEC customer protocol
24583=DEC Local Area VAX Cluster (LAVC)
24584=DEC AMBER
24585=DEC MUMPS
24592=3Com Corporation
28672=Ungermann-Bass download
28673=Ungermann-Bass NIUs
28674=Ungermann-Bass diagnostic/loopback
28675=Ungermann-Bass ??? (NMC to/from UB Bridge)
28677=Ungermann-Bass Bridge Spanning Tree
28679=OS/9 Microware
28681=OS/9 Net?
28704=LRT (England) (now Sintrom)
28720=Racal-Interlan
28721=Prime NTS (Network Terminal Service)
28724=Cabletron
32771=Cronus VLN
32772=Cronus Direct
32773=HP Probe protocol
32774=Nestar
32776=AT&amp;amp;T/Stanford Univ.
32784=Excelan
32787=Silicon Graphics diagnostic
32788=Silicon Graphics network games
32789=Silicon Graphics reserved
32790=Silicon Graphics XNS NameServer
32793=Apollo DOMAIN
32814=Tymshare
```

```
32815=Tigan
32821=Reverse Address Resolution Protocol (RARP)
32822=Aeonic Systems
32823=IPX (Novell Netware?)
32824=DEC LanBridge Management
32825=DEC DSM/DDP
32826=DEC Argonaut Console
32827=DEC VAXELN
32828=DEC DNS Naming Service
32829=DEC Ethernet CSMA/CD Encryption Protocol
32830=DEC Distributed Time Service
32831=DEC LAN Traffic Monitor Protocol
32832=DEC PATHWORKS DECnet NETBIOS Emulation
32833=DEC Local Area System Transport
32834=DEC unassigned
32836=Planning Research Corp.
32838=AT&amp;amp;T
32839=AT&amp;amp;T
32840=DEC Availability Manager for Distributed Systems DECamds
32841=ExperData
32859=VMTP
32860=Stanford V Kernel
32861=Evans &amp;amp; Sutherland
32864=Little Machines
32866=Counterpoint Computers
32869=University of Mass. at Amherst
32870=University of Mass. at Amherst
32871=Veeco Integrated Automation
32872=General Dynamics
32873=AT&amp;amp;T
32874=Autophon
32876=ComDesign
32877=Compugraphic Corporation
32878=Landmark Graphics Corporation
32890=Matra
32891=Dansk Data Elektronik
32892=Merit Internodal
32893=Vitalink Communications
32896=Vitalink TransLAN III Management
32897=Counterpoint Computers
32904=Xyplex
32923=EtherTalk - AppleTalk over Ethernet
32924=Datability
32927=Spider Systems Ltd.
32931=Nixdorf Computers
32932=Siemens Gammasonics Inc.
32960=DCA Data Exchange Cluster
32966=Pacer Software
32967=Applitek Corporation
32968=Intergraph Corporation
32973=Harris Corporation
32975=Taylor Instrument
32979=Rosemount Corporation
32981=IBM SNA Services over Ethernet
32989=Varian Associates
32990=TRFS (Integrated Solutions Transparent Remote File System)
32992=Allen-Bradley
32996=Datability
33010=Retix
33011=AppleTalk Address Resolution Protocol (AARP)
33012=Kinetics
```

```
 33015=Apollo Computer
 33023=Wellfleet Communications
 33026=Wellfleet BOFL
 33027=Wellfleet Communications
 33031=Symbolics Private
 33067=Talaris
 33072=Waterloo Microsystems Inc.
 33073=VG Laboratory Systems
 33079=IPX
 33080=Novell Inc
 33081=KTI
 33087=M/MUMPS data sharing
 33093=Vrije Universiteit (NL)
 33094=Vrije Universiteit (NL)
 33095=Vrije Universiteit (NL)
 33100=SNMP
 33103=Technically Elite Concepts
 33169=PowerLAN
 33149=XTP
 33238=Artisoft Lantastic
 33239=Artisoft Lantastic
 33283=QNX Software Systems Ltd.
 33680=Accton Technologies (unregistered)
 34091=Talaris multicast
 34178=Kalpana
 34525=IPv6
 34617=Control Technology Inc.
 34618=Control Technology Inc.
 34619=Control Technology Inc.
 34620=Control Technology Inc.
 34848=Hitachi Cable (Optoelectronic Systems Laboratory)
 34902=Axis Communications AB
 34952=HP LanProbe test?
 36864=Loopback (Configuration Test Protocol)
 36865=3Com XNS Systems Management
 36866=3Com TCP/IP Systems Management
 36867=3Com loopback detection
 43690=DECNET
 64245=Sonix Arpeggio
 65280=BBN VITAL-LanBridge cache wakeups
 34915=PPPoe
 34916=PPPoe
 2056=Frame Relay ARP
 16962=IEEE bridge spanning protocol
 25944=Bridged Ethernet/802.3 packet
 65278=ISO CLNP/ISO ES-IS DSAP/SSAP
```

ip.proto

```
ALIAS_FORMAT=$alias
 0=HOPOPT
 1=ICMP
 2=IGMP
 3=GGP
 4=IP
 5=ST
 6=TCP
 7=CBT
 8=EGP
 9=IGP
 10=BBN-RCC-M
 11=NVP-II
 12=PUP
```

```
13=ARGUS
14=EMCON
15=XNET
16=CHAOS
17=UDP
18=MUX
19=DCN-MEAS
20=HMP
21=PRM
22=XNS-IDP
23=TRUNK-1
24=TRUNK-2
25=LEAF-1
26=LEAF-2
27=RDP
28=IRTP
29=ISO-TP4
30=NETBLT
31=MFE-NSP
32=MERIT-INP
33=SEP
34=3PC
35=IDPR
36=XTP
37=DDP
38=IDPR-CMTP
39=TP++
40=IL
41=IPv6
42=SDRP
43=IPv6-Rout
44=IPv6-Frag
45=IDRP
46=RSVP
47=GRE
48=MHRP
49=BNA
50=ESP
51=AH
52=I-NLSP
53=SWIPE
54=NARP
55=MOBILE
56=TLSP
57=SKIP
58=IPv6-ICMP
59=IPv6-NoNx
60=IPv6-Opts
61=AnyHost
62=CFTP
63=AnyNetwork
64=SAT-EXPAK
65=KRYPTOLAN
66=RVD
67=IPPC
68=AnyFile
69=SAT-MON
70=VISA
71=IPCV
72=CPNX
73=CPHB
```

```
74=WSN
75=PVP
76=BR-SAT-MO
77=SUN-ND
78=WB-MON
79=WB-EXPAK
80=ISO-IP
81=VMTP
82=SECURE-VM
83=VINES
84=TTP
85=NSFNET-IG
86=DGP
87=TCF
88=EIGRP
89=OSPFIGP
90=Sprite-RP
91=LARP
92=MTP
93=AX.25
94=IPIP
95=MICP
96=SCC-SP
97=ETHERIP
98=ENCAP
99=AnyPrivate
100=GMTP
101=IFMP
102=PNNI
103=PIM
104=ARIS
105=SCPS
106=QNX
107=A/N
108=IPComp
109=SNP
110=Compaq-Pe
111=IPX-in-IP
112=VRRP
113=PGM
114=AnyHop
115=L2TP
116=DDX
117=IATP
118=STP
119=SRP
120=UTI
121=SMP
122=SM
123=PTP
124=ISIS
125=FIRE
126=CRTP
127=CRUDP
128=SSCOPMCE
129=IPLT
130=SPS
131=PIPE Pr
132=SCTP St
133=FC Fi
134=RSVP-E2E-
```

```
 255=Reserved
```

medium

```
ALIAS_FORMAT=$alias
 1=Ethernet
 2=Tokenring
 3=FDDI
 4=HDLC
 5=NetWitness
 6=802.11
 7=802.11 Radio
 8=802.11 AVS
 9=802.11 PPI
 10=802.11 PRISM
 11=802.11 Management
 12=802.11 Control
 13=DLT Raw
 32=Logs
```

service

```
ALIAS_FORMAT=$alias
 0=OTHER
 20=FTPD
 21=FTP
 22=SSH
 23=TELNET
 25=SMTP
 53=DNS
 67=DHCP
 69=TFTP
 80=HTTP
 110=POP3
 111=SUNRPC
 119=NNTP
 123=NTP
 135=RPC
 137=NETBIOS
 139=SMB
 143=IMAP
 161=SNMP
 179=BGP
 443=SSL
 502=MODBUS
 520=RIP
 1024=EXCHANGE
 1080=SOCKS
 1122=MSN IM
 1344=ICAP
 1352=NOTES
 1433=TDS
 1521=TNS
 1533=SAMETIME
 1719=H.323
 1720=RTP
 2000=SKINNY
 2040=SOULSEEK
 2049=NFS
 3270=TN3270
 3389=RDP
 3700=DB2
 5050=YAHOO IM
 5060=SIP
```

```
5190=AOL IM
5222=Google Talk
5900=VNC
6346=GNUTELLA
6667=IRC
6801=Net2Phone
6881=BITTORRENT
8000=QQ
8002=YCHAT
8019=WEBMAIL
8082=FIX
20000=DNP3
1000000=KERNEL
1000001=USER
1000003=SYSTEM
1000004=AUTH
1000005=LOGGER
1000006=LPD
1000008=UUCP
1000009=SCHEDULE
1000010=SECURITY
1000013=AUDIT
1000014=ALERT
1000015=CLOCK
```

tcp.dstport

```
ALIAS_FORMAT=$value ($alias)
 7=echo
 9=discard
 13=daytime
 17=qotd
 19=chargen
 20=ftp-data
 21=ftp
 22=ssh
 23=telnet
 25=smtp
 37=time
 42=nameserver
 43=nicname
 53=domain
 70=gopher
 79=finger
 80=http
 88=kerberos
 101=hostname
 102=iso-tsap
 107=rtelnet
 109=pop2
 110=pop3
 111=sunrpc
 113=auth
 117=uucp-path
 119=nntp
 135=epmap
 137=netbios-ns
 139=netbios-ssn
 143=imap
 158=pcmail-srv
 170=print-srv
 179=bgp
 194=irc
 389=ldap
 443=https
 445=cifs
 464=kpasswd
 512=exec
 513=login
 514=cmd
 515=printer
 520=efs
 526=tempo
 530=courier
 531=conference
 532=netnews
 540=uucp
 543=klogin
 544=kshell
 556=remotefs
 636=ldaps
 749=kerberos-adm
 993=imaps
 995=pop3s
 1109=kpop
 1433=ms-sql-s
 1434=ms-sql-m
```

```
 1512=wins
 1524=ingreslock
 1723=pptp
 2053=knetd
 1122=msn im
 1352=notes
 1521=tns
 1533=sametime
 1718=h323
 1720=rtp
 1863=msn im
 2049=nfs
 3389=rdp
 5050=yahoo im
 5060=sip
 5190=aim
 6346=gnuetella
 6667=irc
 9001=tor
 9030=tor
 9535=man
```

tcp.srcport

```
ALIAS_FORMAT=$value ($alias)
 7=echo
 9=discard
 13=daytime
 17=qotd
 19=chargen
 20=ftp-data
 21=ftp
 22=ssh
 23=telnet
 25=smtp
 37=time
 42=nameserver
 43=nicname
 53=domain
 70=gopher
 79=finger
 80=http
 88=kerberos
 101=hostname
 102=iso-tsap
 107=rtelnet
 109=pop2
 110=pop3
 111=sunrpc
 113=auth
 117=uucp-path
 119=nntp
 135=epmap
 137=netbios-ns
 139=netbios-ssn
 143=imap
 158=pcmail-srv
 170=print-srv
 179=bgp
 194=irc
 389=ldap
 443=https
 445=cifs
```

```
 464=kpasswd
 512=exec
 513=login
 514=cmd
 515=printer
 520=efs
 526=tempo
 530=courier
 531=conference
 532=netnews
 540=uucp
 543=klogin
 544=kshell
 556=remotefs
 636=ldaps
 749=kerberos-adm
 993=imaps
 995=pop3s
 1109=kpop
 1433=ms-sql-s
 1434=ms-sql-m
 1512=wins
 1524=ingreslock
 1723=pptp
 2053=knetd
 1122=msn im
 1352=notes
 1521=tns
 1533=sametime
 1718=h323
 1720=rtp
 1863=msn im
 2049=nfs
 3389=rdp
 5050=yahoo im
 5060=sip
 5190=aim
 6346=gnuetella
 6667=irc
 9001=tor
 9030=tor
 9535=man
```

udp.dstport

```
ALIAS_FORMAT=$value ($alias)
 7=echo
 9=discard
 13=daytime
 17=qotd
 19=chargen
 37=time
 39=rlp
 42=nameserver
 53=domain
 67=bootps
 68=bootpc
 69=tftp
 88=kerberos
 111=sunrpc
 123=ntp
 135=epmap
 137=netbios-ns
```

```
138=netbios-dgm
161=snmp
162=snmptrap
213=ipx
443=https
445=cifs
464=kpasswd
500=isakmp
512=biff
513=who
514=syslog
517=talk
518=ntalk
525=timed
533=netwall
550=new-rwho
560=rmonitor
561=monitor
749=kerberos-adm
1167=phone
1433=ms-sql-s
1434=ms-sql-m
1512=wins
1701=l2tp
1812=radiusauth
1813=radacct
2049=nfsd
2504=nlbs
```

Test a Rule

You can test a rule based on the time range and the data source selected.

To test a rule, perform the following:

1. Go to **Reports**.
   The Manage tab is displayed.
2. In the **Rules** panel, do one of the following:
   - Select a rule and click  in the Rules toolbar.
   - Click  > **Edit**.
     The Build Rule view tab is displayed.

3. Click **Test Rule**.
   The Test Rule view is displayed.



> **Note:** When you click **Test Rule**, the rule is not saved. You have to click **Save**in the Build Rule view to save the rule.

4. From the **Data Source** drop-down list, select a data source.
   You must select the appropriate data source for the rule defined.
5. From the **Format** drop-down list, select the format in which you want the result displayed.
6. From the **Time Range** drop-down list, select one of the following.
   - **Past** -To specify number of years, days, weeks, months, days or hours.
   - **Range** - To specify a date range and time period.

   > **Note:** In the User Interface (UI), the date or time displayed depends on the time zone profile selected by the user.

7. **X-Axis** and **Y-Axis** are used to specify the meta to be plotted in charts.
   In **X-Axis**, the Meta for the 'Group by' rule is displayed. In **Y-Axis**, the aggregate functions used in the rule are displayed.

> **Note:** Sum, Count, Countdistinct and Average are the supported aggregate functions for rule. By default, for Custom Rules with multiple 'Group by', you can select only the first meta in **X-Axis**.

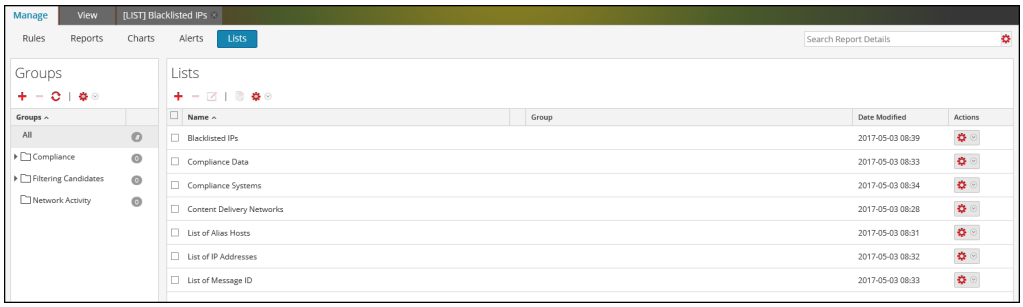
8. Click **Run Test** to execute the rule.
   The rule data (if any) for the selected time range is displayed.

Create a Lists or List Group

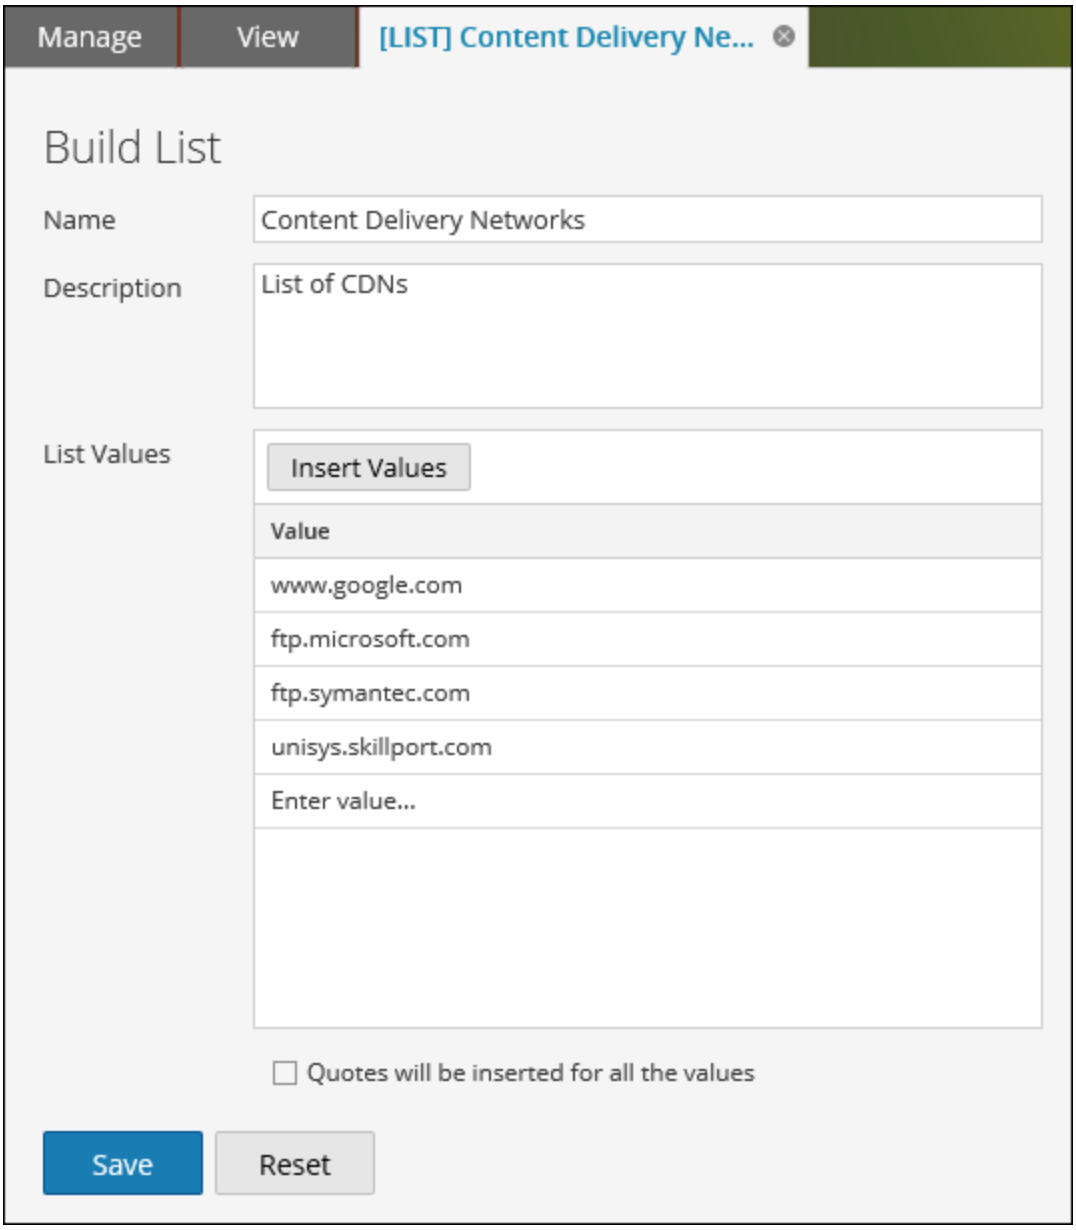
To create a list, perform the following:

Lists can be added within a group or in the root folder.

1. Go to **Reports**, and click **Lists**.
   The List view is displayed.



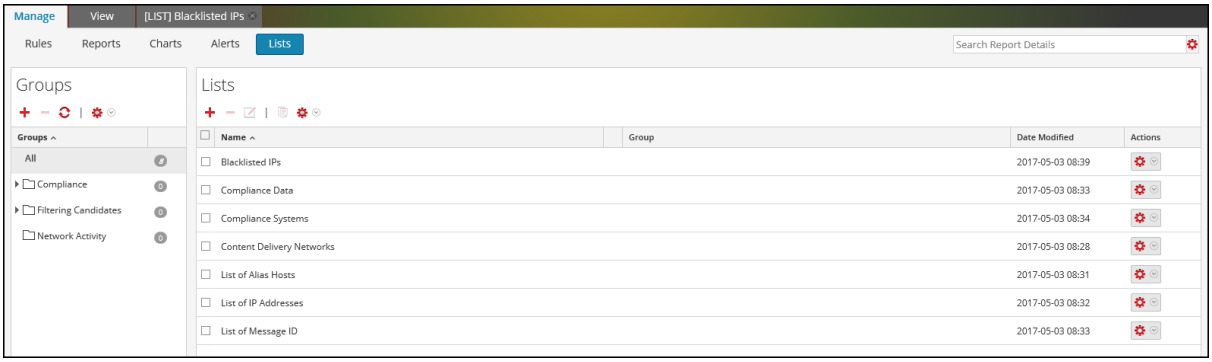
2. In the **Lists** toolbar, click .
   The Build List view tab is displayed.



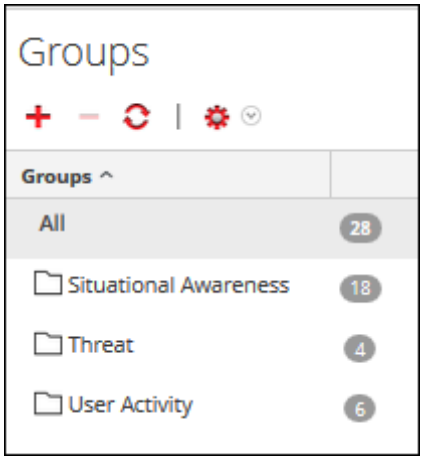
3. In the **Name** field, enter a unique name for the list.
4. In the **Description** field, enter a description for the list.
5. In the **List Values** field, do one of the following:
   - Click **Insert Values** and enter the values separated by commas. You can paste a list of values from a file or other lists.
   - In the **Value** column, enter the values.
6. If you want quotes to be inserted directly for the values at runtime, select **Quotes will be inserted for all the values**.
7. Click **Save**.

To create a list group, perform the following:

1. Go to **Reports**, and click **Lists**.
   The List view is displayed.



2. Do the following:
   - To create a list group:
     a. In the **Lists Groups** panel, click  .
        A new list group is added to the List Groups panel.



     b. Enter the name for the list group and press **ENTER**.
   - To create a list subgroup:
     1. In the **Lists** Groups panel, select the list group to which you want to add a subgroup.
     2. Click  .
        A new list subgroup is added to the list group.
     3. Enter the name for the list subgroup and press **ENTER**.

Create and Schedule a Report

You can create a simple or complex report and configure its execution properties by scheduling a report. A report can include multiple rules and you can schedule different time range to execute the same report. For example, depending on your requirement, you can schedule a report to run daily, weekly or monthly.

When you run a report, the results are stored in Reporting Engine.

After you generate a report, you can perform the following:

- Send the reports by email to other users by configuring the output actions. You can also configure the output actions before generating a report.
- Download the reports as PDF or Comma-Separated Values (CSV) format files.

Note: The cancel operation is not supported for Respond Reports.

Create a Report or Report Group

To create a report to a group or sub-group, perform the following:

1. Go to **Reports**.

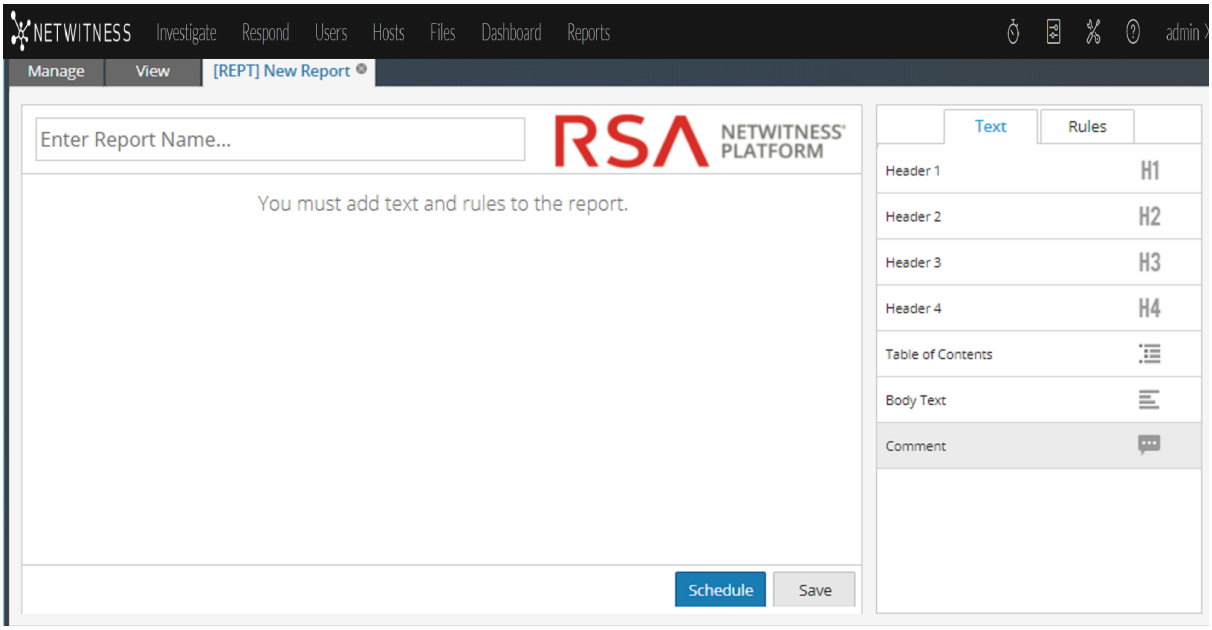
   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** toolbar, click + .

   The Build Report tab is displayed.



4. Enter the name of the report.

5. Drag and drop the text and rules to the report.

   Note: The text entered is optional and you may need this option only when you want to display user-defined headers or content.

6. Click **Save**.

   A confirmation message that the report is saved successfully is displayed.

To create a group to the default folder or add sub-groups under a report group, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports Groups** panel, click  .

   A default group is added in the Report Groups panel.

4. Enter the name of the new group.
5. Press **Enter**.

   The group is added to the Report Groups panel.

Schedule a Report

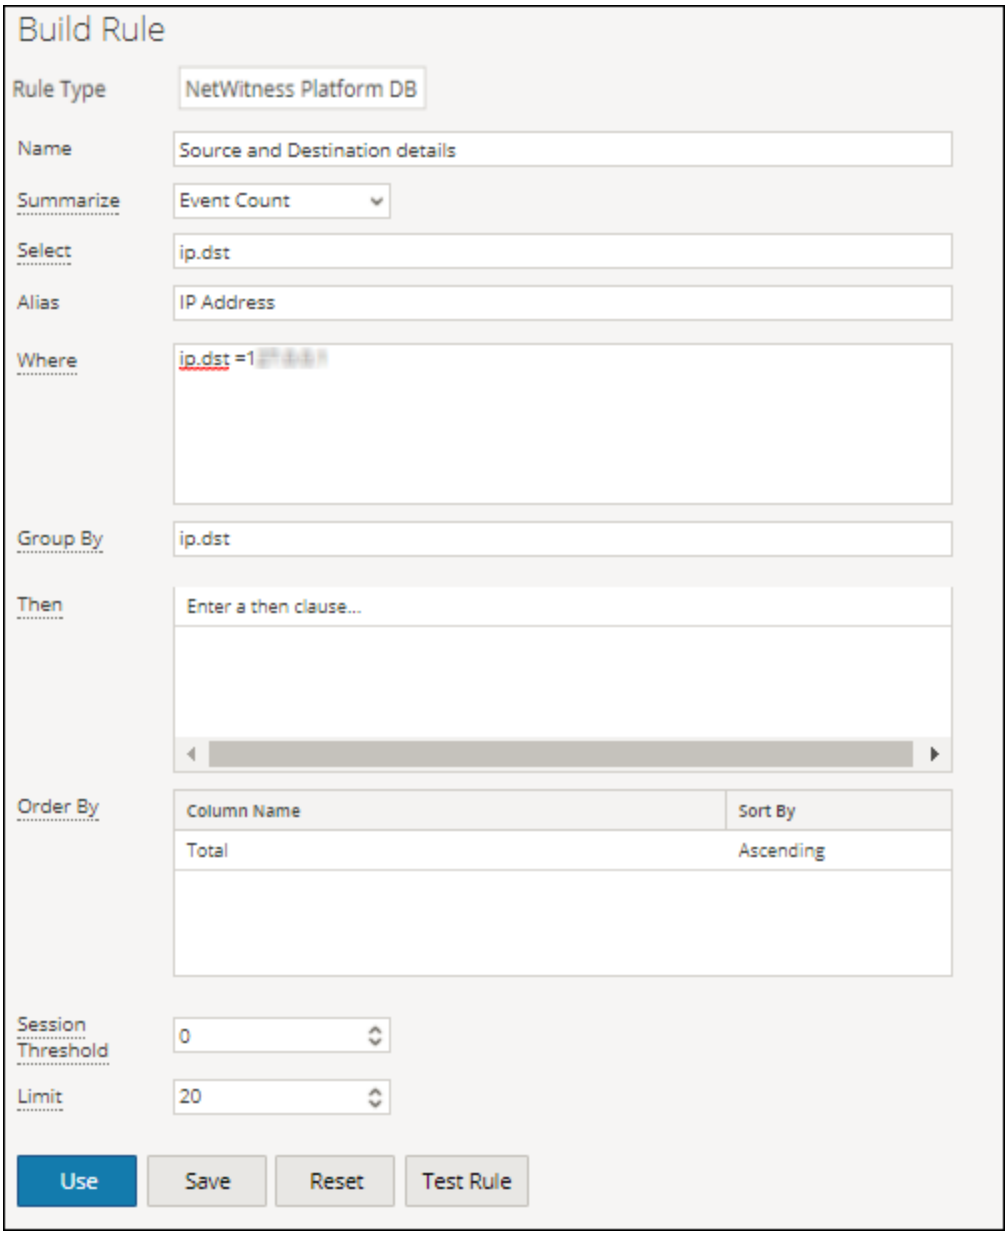
Note: When you schedule a Warehouse report, you can use a supported task scheduler to allocate specific resources in a cluster for the scheduled job. For more information on "supported task schedulers", see Task Scheduler for Warehouse Reporting.

To schedule a report, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** panel, click  to create a rule.
3. Click **Save**.
4. Click **Use**.



5. Go to **Reports**.

   The Manage tab is displayed.

6. Click **Reports**.

   The Report view is displayed.

7. On the **Reports** panel, click + to create a report.
8. Enter the Report Name in the field.
9. Add the rule by drag and drop which has the user defined variable from the Rules tab.
10. Click **Schedule**.

The Schedule Report view is displayed.

If you provide another user with access permissions to a report, you must also provide permissions for the report group, the rules used in the report, and the rule groups otherwise an error message is displayed.

11. To execute the reports as per the schedule, select the **Enable** checkbox.
12. In the **Schedule Name** field, enter a name for the schedule report configuration.
13. From the Data Source field, select the data source.

    Note: If the data source is not listed, then ensure you have **Read**permissions set for the data source. This is applicable for NWDB, Respond and Warehouse data source. For more information, see "Configure Data Source Permissions" topic in *Reporting Engine Configuration Guide*.

14. (Optional) From the **Warehouse Resource Pool** drop-down, select the pools or queues available in the cluster to schedule the report to run on either the pool or queue. This drop-down list is available only if you select a Warehouse DB report.

    Note: All the queues or pools you specified in the Explore page for the Reporting Engine are listed. If no pools or queues are configured in the Explorer page, this drop-down is disabled and the jobs are submitted to the clusters without any a queue or pool name.

    Note: If the pool or queue configured in the report schedule is removed from the Cluster, then in the Capacity Scheduler, the queue name remains undefined. However, in the Fair Scheduler, the specified pool name will be created using the property mapred.fairscheduler.allow.undeclared.pool.

15. From the Time Zone drop-down, select a time zone to display all the time-related data in a report output in the specified format. This setting is configurable from the Reporting Engine Explore view (/com.rsa.soc.re/configuration/reportoutputformatterconfig/reportoutputformatterconfig).
16. From the **Run** field, select the type of run schedule. (For example, Now or Hourly).

    Depending on the type of run schedule, choose one of the following:

    - If you select a **Later** or **Monthly** run schedule, you must provide a value for the day and time in the respective field provided.
    - If you select an **Hourly** run schedule, you must specify the minutes in the **At Minute** field.
    - If you select a **Daily** run schedule, you must enter a value in the **At** field.
    - If you select a **Weekly** run schedule, you must enter a value in the **At** field and also select the week days.

    Note: While scheduling a report, if you select **Past** option or **Range (specific/generic)** option or an end time range very close to the current time, you must ensure that the aggregate data in the data source is returned. If there is an aggregation delay in the data source, the end time you choose must account for the delay, otherwise reports lose non-aggregate data for that time range.

    For information on how to generate a report with variables, see Create a Parameterized Report Using Variable.

17. (Optional) In the **Output Actions** panel, do the following:

a. Enter the email address and subject.

b. Edit the body of the message for the report.

c. Select the format of the attachment.

d. Enter a value for the CSV and Multi-value delimiters.

e. (Optional) In the Other Options field, do the following:

i. Click + and select SFTP, URL, or Network Share output action.

A row gets added with the selected output action.

ii. Select the appropriate options to send the report in PDF or CSV format, or both to the RE configured SFTP, or URL, or Network Share output action.

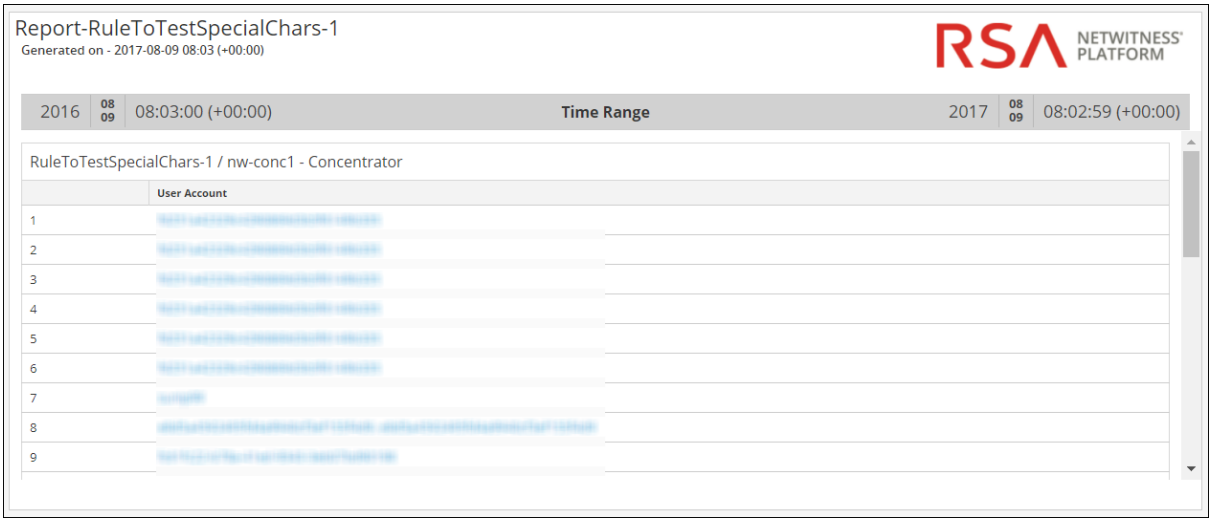
18. (Optional) To add a list in the Dynamic List panel, see Generate a List from the Scheduled Report.

19. (Optional) To choose a logo in the Logo panel, see "Manage and Select a Report Logo" section in Manage Lists, Rules or Reports.

Note: If you do not specify a logo, the default RSA logo will be used.

20. Click **Schedule**.

The scheduled report executes as scheduled and provides the configured outputs.



After you create and Schedule a report, you can perform any of the following tasks:

- You can notify the email recipient when the report execution completes and send reports in PDF and CSV formats as attachments in the email.
- You can generate a list based on the scheduled report and view them in the **Lists** module.
- You can send a scheduled report in PDF or CSV format, or both to the RE configured SFTP location, or URL, or Network Share.
- You can change the default logo and view them in the scheduled report.

- You can modify the NetWitness Reporting Engine config details, by navigating to the Reporting Engine General Tab. See the "Reporting Engine General Tab" topic in the *Reporting Engine Guide*.

Examples

When you schedule reports in the Schedule Report view, by default, the results for the **Past** option are presented based on the user specified time zone. The following examples provide a clear picture on what results to expect when you select **Hours**, **Days**, **Weeks**, **Months**, or **Years** for the **Past** option based on the absolute or relative duration.

Note: By default, the relative duration checkbox is de-selected. This implies that the results for the **Past** option are presented based on the absolute duration.

- **Based on Absolute duration** - Absolute Duration allows a report to be scheduled at an absolute time with respect to the current time, excluding the seconds and considering the time interval as a whole. For example, 12.00pm is the absolute time with respect to the current time (12.45 pm).
  - Hours - Suppose that you select Hours and specify one hour. If the current user specified time is 4.20PM, the report is generated for the time range, 3.00PM to 4.00PM.
  - Days - Suppose that you select Days and specify one day. If the current date is August 27, 2014 and the current user specified time is 10.15AM, the report is generated for the range: August 26, 2014, 12.00AM to August 27, 2014, 12.00AM.
  - Weeks - Suppose that you select Weeks and specify one week. If the current date is August 27, 2014 2.30PM and the day is Wednesday, the report is generated for the range: Saturday, August 16, 2014, 12.00AM to Saturday, August 23, 2014, 12.00AM.
  - Months - Suppose that you select Months and specify one month. If the current date is August 27, 2014 2.30PM, the report is generated for the range:
    July 01, 2014, 12.00AM to July 31, 2014, 12.00AM.
  - Years - Suppose that you select Years and specify one year.If the current date is August 27, 2014 2.30PM, the report is generated for the range:
    January 01, 2013, 12.00AM to December 31, 2013, 12.00AM.
- **Based on Relative duration** - Relative Duration allows a report to be scheduled at a time relative to the current time which might vary based on the current time. For example, 12.45 pm is the relative time with respect to the current time (12.45 pm).
  - Hours - Suppose that you select Hours and specify one hour. If the current user specified time is 4.20PM, the report is generated for the time range, 3.20PM to 4.20PM.
  - Days - Suppose that you select Days and specify one day. If the current date is August 27, 2014 and the current user specified time is 10.15AM, the report is generated for the range: August 26, 2014, 10.15AM to August 27, 2014, 10.15AM.
  - Weeks - Suppose that you select Weeks and specify one week. If the current date is August 27, 2014 12.30PM and the day is Wednesday, the report is generated for the range: Thursday, August 21, 2014 12.30PM to Wednesday, August 27, 2014 12.30PM.
  - Months - Suppose that you select Months and specify one month. If the current date is August 27, 2014, 2.30PM the report is generated for the range:
    July 27, 2014 2.30PM to August 27, 2014 2.30PM.

- Years - Suppose that you select Years and specify one year.If the current date is August 27, 2014 2.30PM, the report is generated for the range: August 27, 2013 2.30PM to August 27, 2014 2.30PM.

Generate a List from the Scheduled Report

You can generate a list from the output of the scheduled report. Make sure that your lists are created in NetWitness prior to generating a list to schedule a report.

To generate a list from the Build Report view, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.
2. Click **Reports**.

   The Report view is displayed.
3. In the **Reports** panel, select a report and in Actions column, click  > **Schedule Report**.

   The Schedule Report view tab is displayed.
4. In the Output Actions section, in **Dynamic List** panel, click .

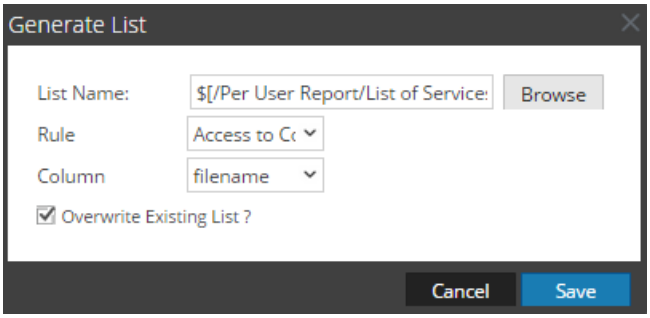
   The Generate List dialog box opens.
5. Click **Browse**.

   The List Selection panel is displayed.
6. Choose a list item and click **Select**.

   The list name gets populated in the List Name field.
7. Select a valid rule to filter the report results further based on the rule definition.
8. Select a value for the **Column** field.

   The column forms the values for the list that gets created.
9. If you want to overwrite the existing list, select the **Overwrite Existing List?** checkbox.
10. Click **Save**.

    The list name gets populated in the Generate List panel.



11. (Optional) Select a list from the Generate List panel and click to delete the selected list.
12. (Optional) Select a list from the Generate List panel and click to edit the list details.

Create a Parameterized Report Using Variable

You use variables for reporting in the NetWitness Reporting module. Parameterized reporting allows you to specify values dynamically at runtime without changing the rule definition so you can view the results based on a particular value. You can achieve parameterize reporting by using variables in the query or rule. For information on adding a rule, see Configure a Rule. At runtime, you can enter the value for the variable or select the value from the list based on which the result set is displayed.

The syntax to specify the variable is as follows:

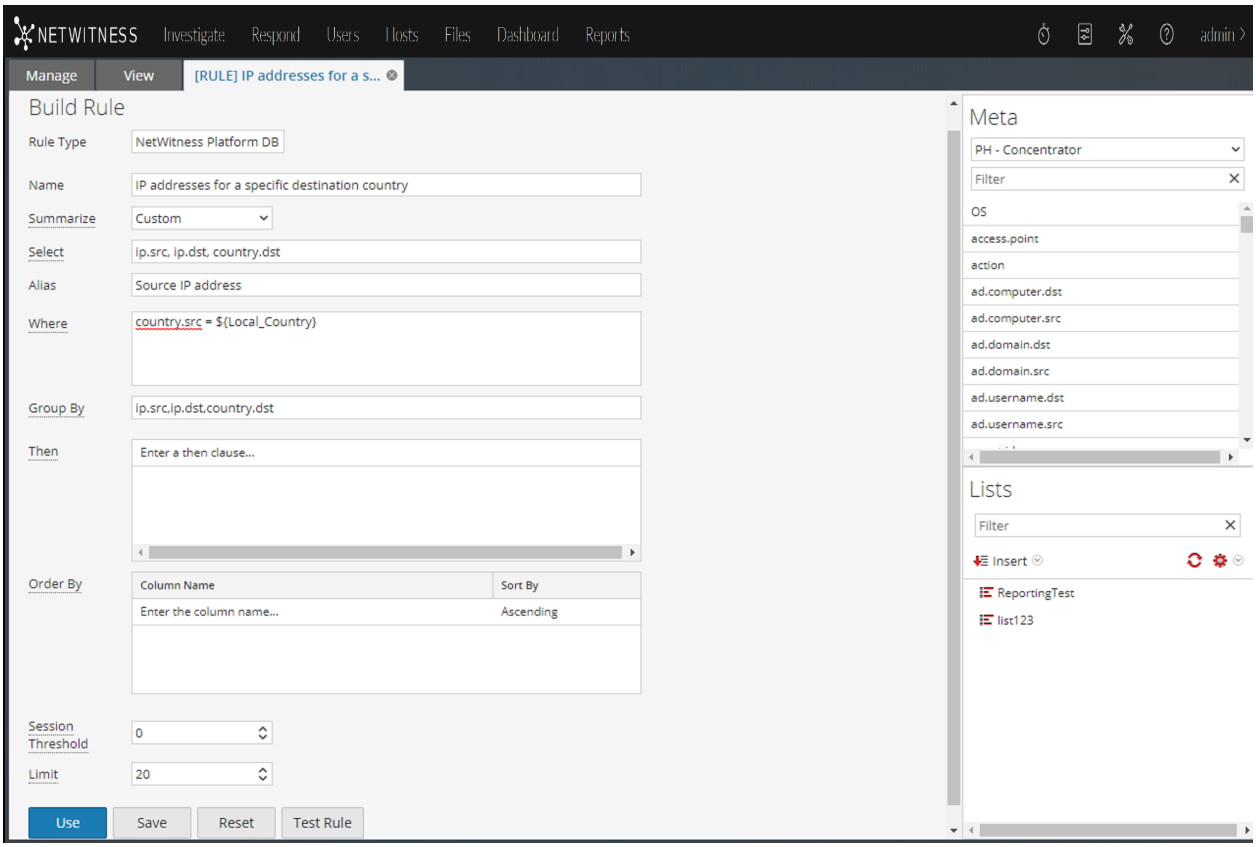
Description	Examples of Supported Syntax
Insert `$` before a variable. Enclose a variable within braces.	`columnname=${<variable>}`

The syntax to define the variable is the same for NetWitness DB and Warehouse DB data sources. When you assign the value of the variable in a Run Configuration, you must enclose the value within single quotes: `'<value>'`.

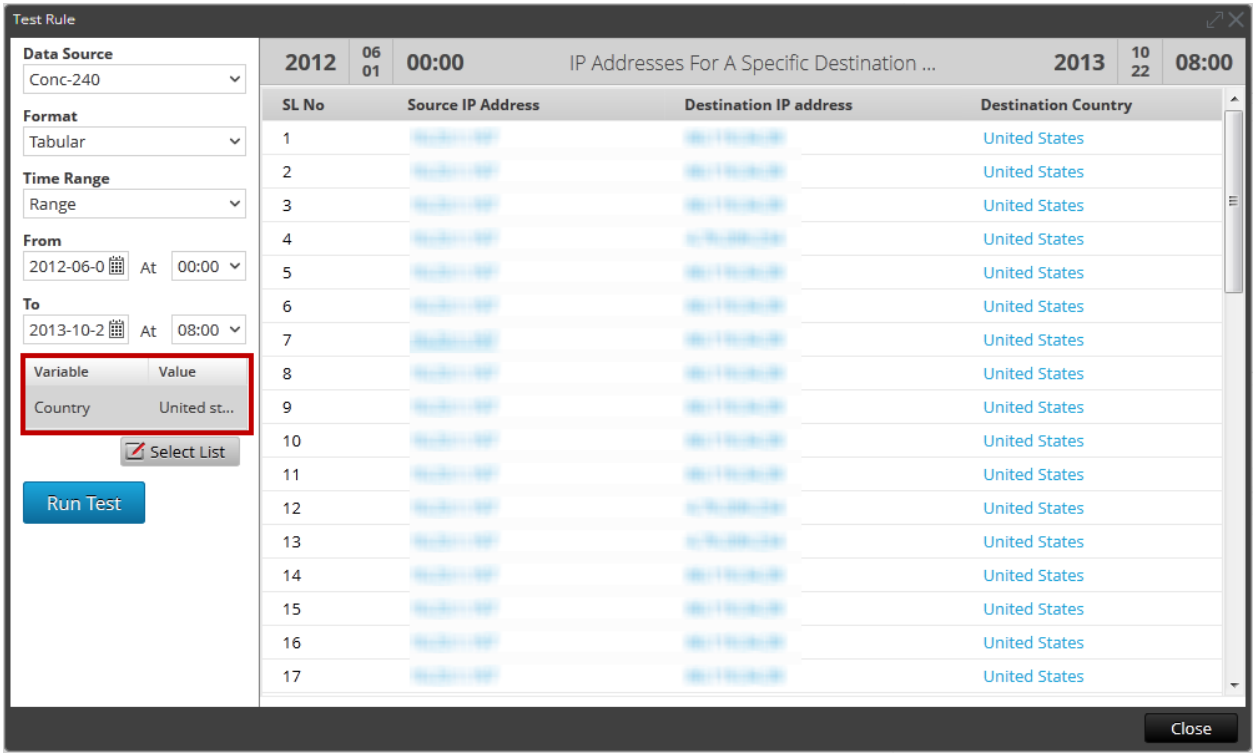
Some examples where a variable can be used are provided in this section.

View Source IP Addresses for a Specific Destination Country

The following is an example of a NetWitness DB rule to view the source and destination ip addresses for a specific destination country. Here the value for source country is defined as a variable `${local_country}`.



At runtime, you are prompted to enter the value for the variable. The figure below shows the `local_Country` variable where you can enter the value. If you enter the value as **United states**, all the source and destination ip addresses with destination country as United states are listed.



You can use the above rule to schedule a report. You can schedule two types of reports:

- Report with Dynamic Variables
- Iterative Report

Report with Dynamic Variables

Dynamic variables allows the user to specify the values for a variable defined in a rule while scheduling a report.

To schedule a report with Dynamic Variable, perform the following:

1. Go to **Reports**.

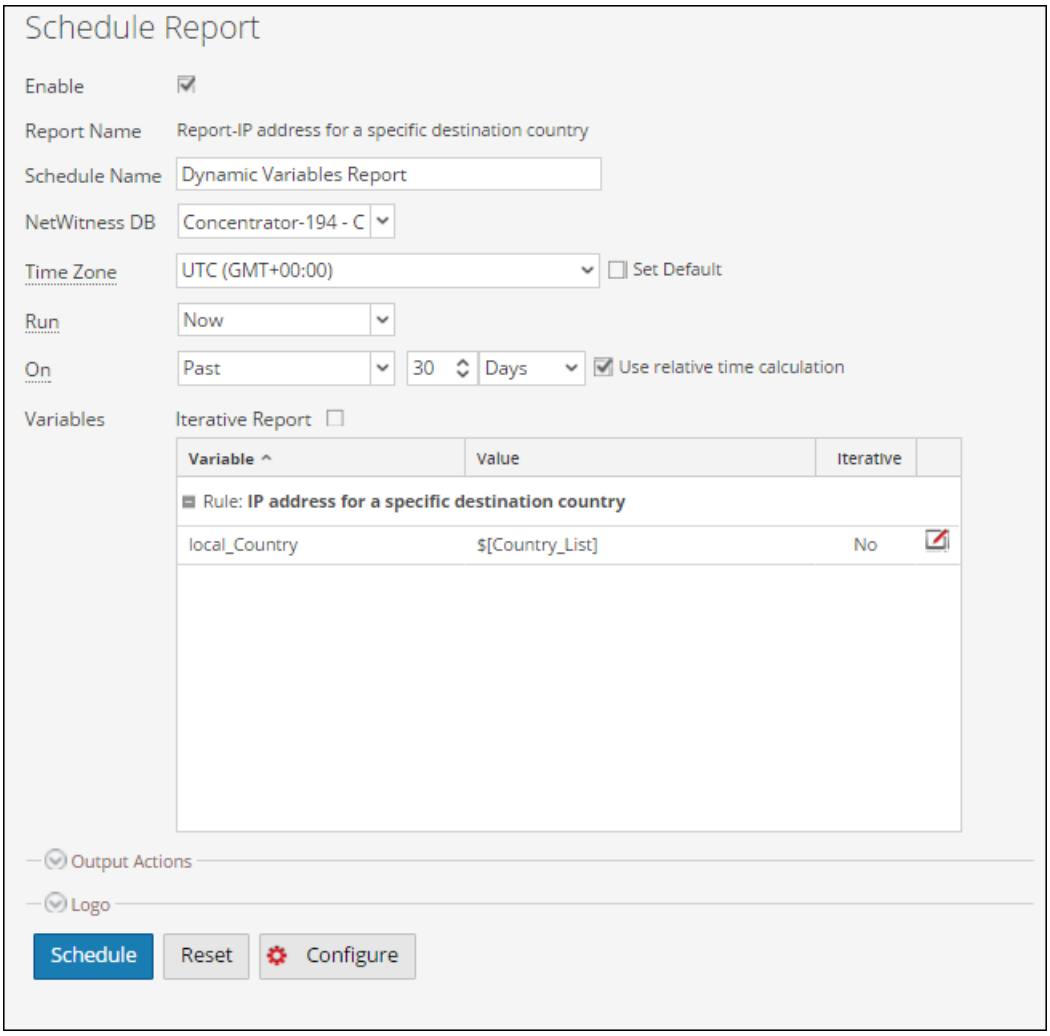
   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. On the **Reports** panel, click + to create a report.
4. Enter the Report Name in the field.
5. Add the rule by drag and drop which has the user defined variable from the Rules tab.
6. Click **Schedule**.

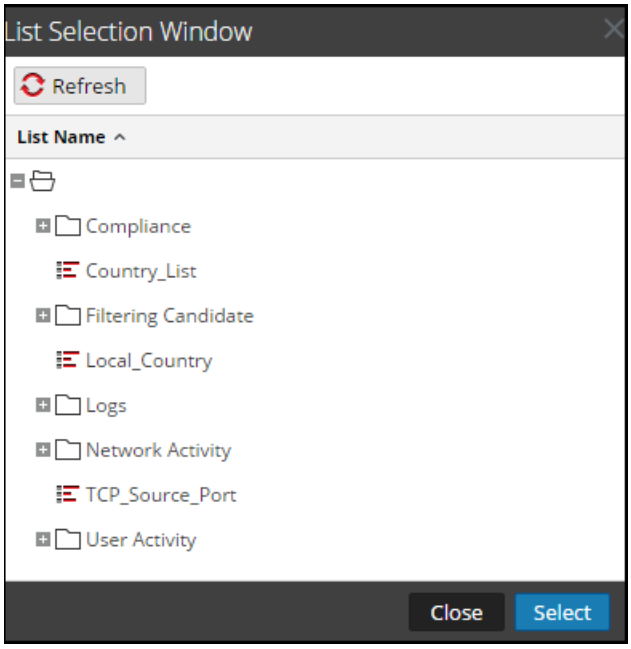
   The Schedule Report view tab is displayed.



7. To execute the reports as per the schedule, select the **Enable** checkbox.
8. In the **Schedule Name** field, enter a name for the schedule report configuration.
9. In the **NetWitness DB** drop-down, select the database.
10. From the Time Zone drop-down, select a time zone to display all the time-related data in a report output in the specified format.This setting is configurable from the Reporting Engine Explore view (/com.rsa.soc.re/configuration/reportoutputformatterconfig/reportoutputformatterconfig).
11. From the **Run** field, select the type of run schedule. (For example, Now or Hourly). Depending on the type of run schedule, do either of the following:
    - If you select a **Later** or **Monthly** run schedule, you must provide a value for the day and time in the respective field provided.
    - If you select an **Hourly** run schedule, you must specify the minutes in the **At Minute** field.
    - If you select a **Daily** run schedule, you must enter a time value in the **At** field.
    - If you select a **Weekly** run schedule, you must enter a value in the **At** field and also select the week days.

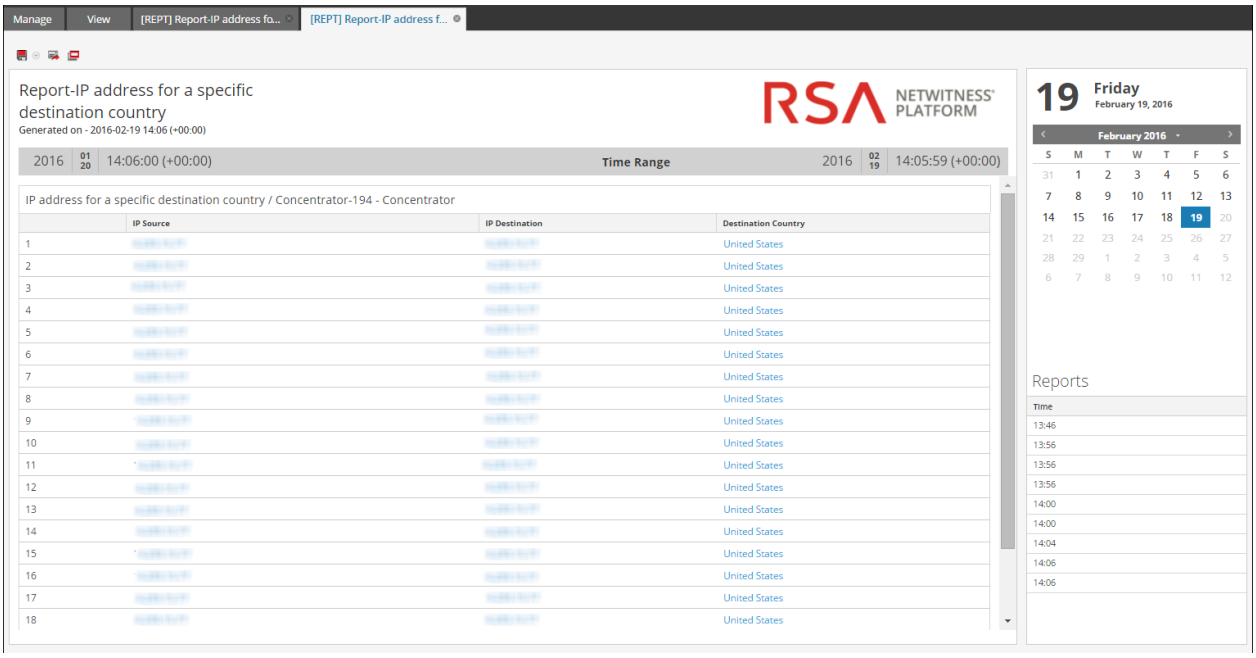
Note: While scheduling a report, if you select **Paste** option or **Range (specific/generic)** option or an end time range very close to the current time, you must ensure that the aggregate data in the data source is returned. If there is an aggregation delay in the data source, the end time you choose must account for the delay, otherwise reports lose non-aggregate data for that time range.

12. In the variables field, click .
13. Do one of the following:
    - Enter the value for the variable, or
    - Choose the list value for the variable.



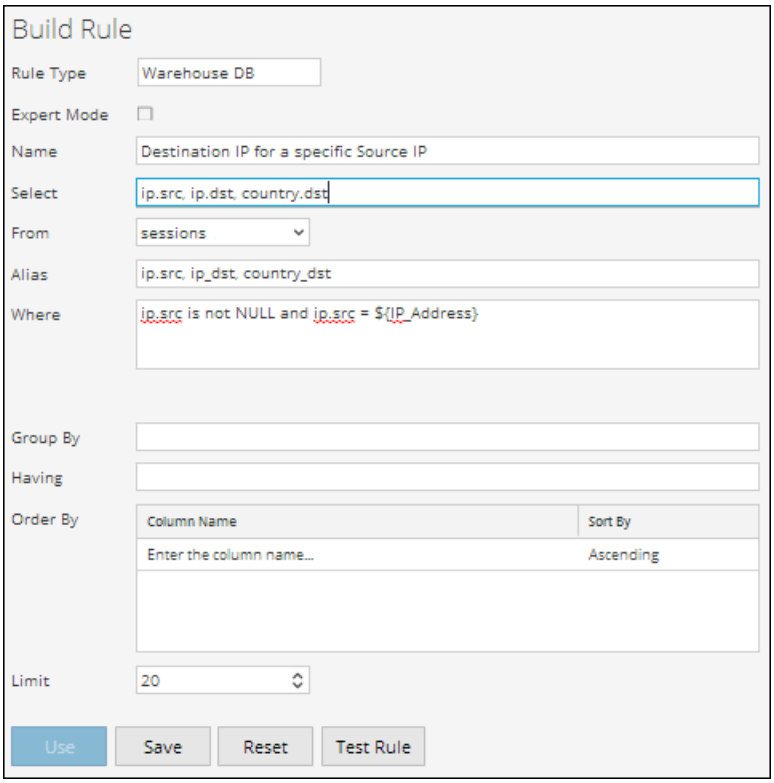
14. Click **Select**.
15. Click **Schedule**.

    The scheduled report executes as scheduled and provides the configured outputs.

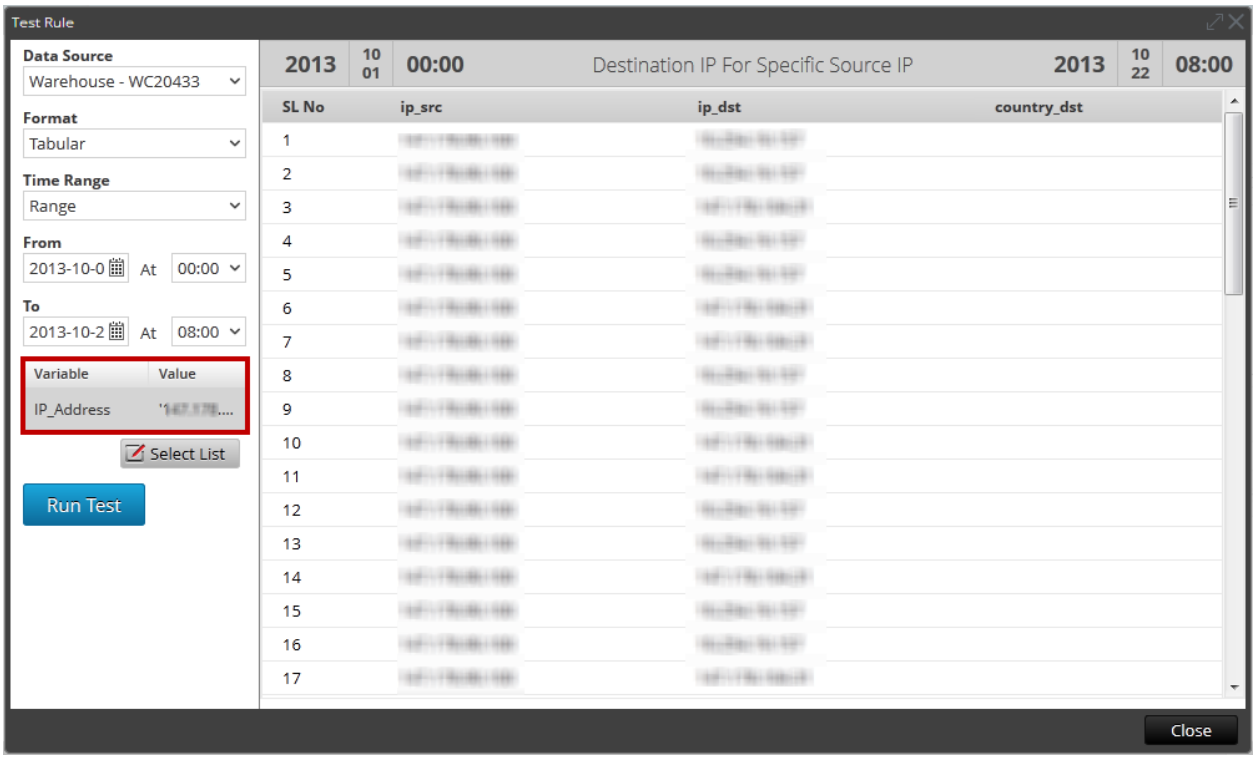


View All Destination IP Addresses for a Source IP Address

The following is an example of a Warehouse rule to view all the destination IP addresses for a specific source IP. The source IP address `ip_src` is defined as a variable `${IP_Address}`.

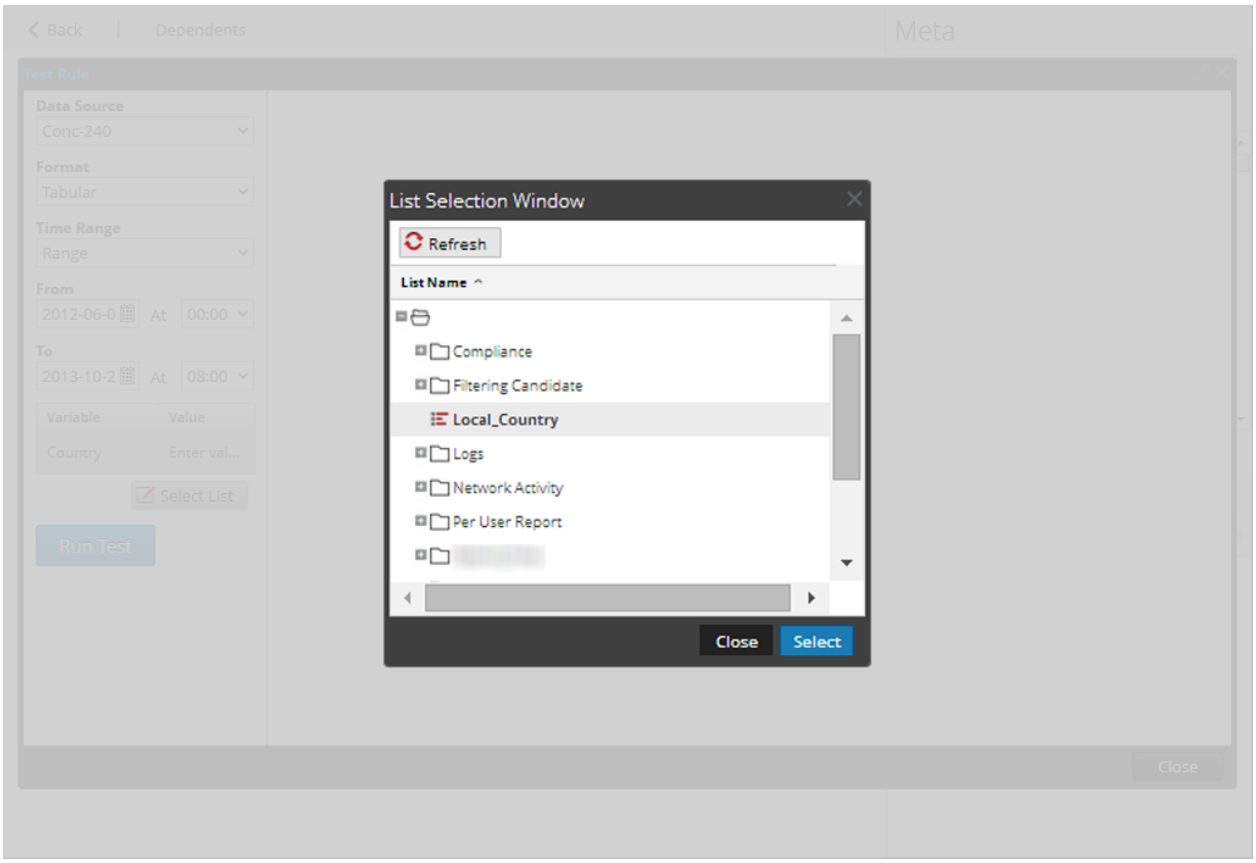


At runtime, you are prompted to enter the source IP address. The figure below shows the `IP_Address` variable, and you can enter a valid source IP address. All the destination IP addresses with the specified source IP are listed.



Associate a Variable to a List of Values

You can associate the variable to a list. For example, you can create a list called `Local_Country` and enter all the country names as values. You can select the list `Local_Country` as the value for the variable `Local_Country`. At Run Configuration, the `Local_Country` list is populated and you can select the country based on which results are displayed.



Iterative Report

An iterative report generates a report for every value in the list.

To schedule an iterative report, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

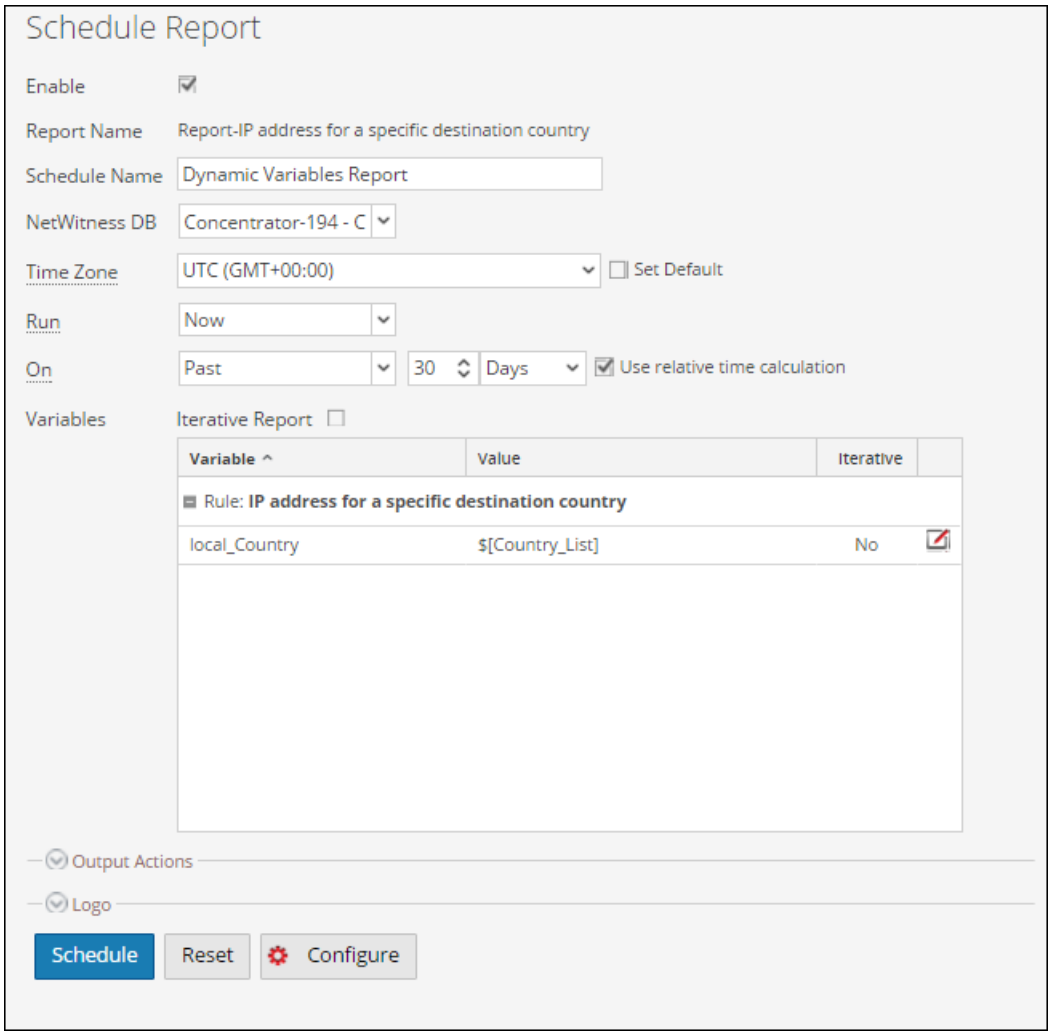
   The Report view is displayed.

3. On the **Reports** panel, click + to create a report.

4. Enter a Report name in the field.

5. Add the rule which has the user defined variable from the Rules tab.

6. Click **Schedule**.
   The Schedule Report view tab is displayed.



7. To execute the reports as per the schedule, select the **Enable** checkbox.
8. In the **Schedule Name** field, enter a name for the schedule report configuration.
9. From the **Data Source** field, select the data source.

> **Note:** If the data source is not listed, then ensure you have **Read** permissions set for the data source. This is applicable for NWDB and Warehouse data source. For more information, see "Configure Data Source Permissions" topic in the *Reporting Engine Configuration Guide*.

10. (Optional) From the **Warehouse Resource Pool** drop-down, select the pools or queues available in the cluster to schedule the report to run on either the pool or queue. This drop-down list is available only if you select a Warehouse DB report.

> **Note:** All the queues or pools you specified in the Explore page for the Reporting Engine are listed. If no pools or queues are configured in the Explorer page, this drop-down is disabled and the jobs are submitted to the clusters without any a queue or pool name.

Note: If the pool or queue configured in the report schedule is removed from the Cluster, then in the Capacity Scheduler, the queue name remains undefined. However, in the Fair Scheduler, the specified pool name will be created using the property `mapred.fairscheduler.allow.undeclared.pool`.

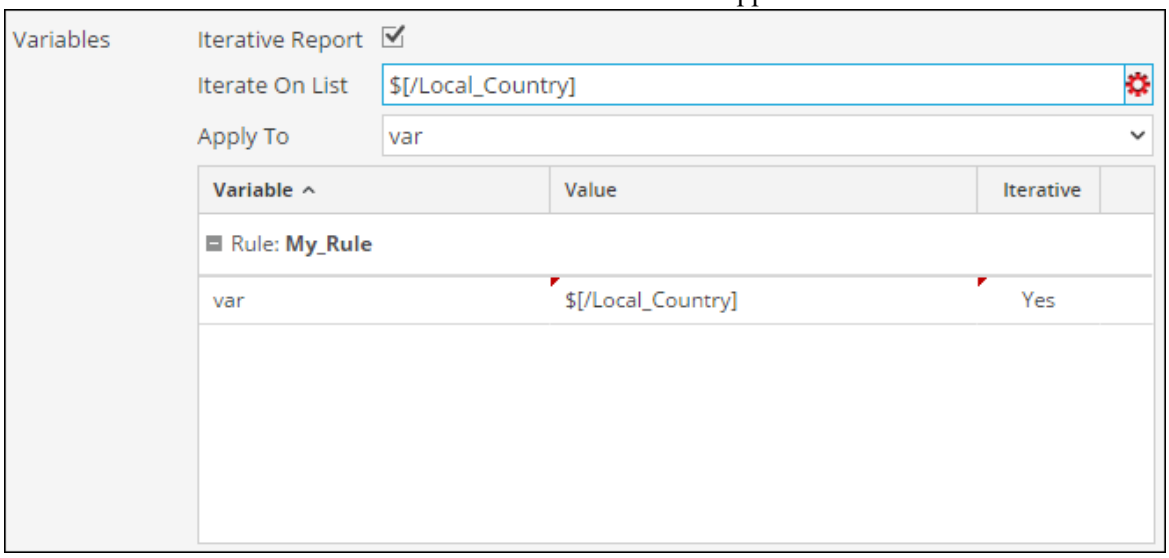
11. From the Time Zone drop-down, select a time zone to display all the time-related data in a report output in the specified format. This setting is configurable from the Reporting Engine Explore view (/com.rsa.soc.re/configuration/reportoutputformatterconfig/reportoutputformatterconfig).
12. From the **Run** field, select the type of run schedule. (For example, Now or Hourly). Depending on the type of run schedule, do either of the following:
    - If you select a **Later** or **Monthly** run schedule, you must provide a value for the day and time in the respective field provided.
    - If you select an **Hourly** run schedule, you must specify the minutes in the **At Minute** field.
    - If you select a **Daily** run schedule, you must enter a time value in the **At** field.
    - If you select a **Weekly** run schedule, you must enter a value in the **At** field and also select the week days.

    Note: While scheduling a report, if you select **Paste** option or **Range (specific/generic)** option or an end time range very close to the current time, you must ensure that the aggregate data in the data source is returned. If there is an aggregation delay in the data source, the end time you choose must account for the delay, otherwise reports lose non-aggregate data for that time range.
13. In the variables field, do the following:
    a. To run iterative reports, select the **Iterative Report** checkbox.
    b. To Iterate on List value, click .

       The List Selection Window opens.
    c. Choose a list and click **Select**.

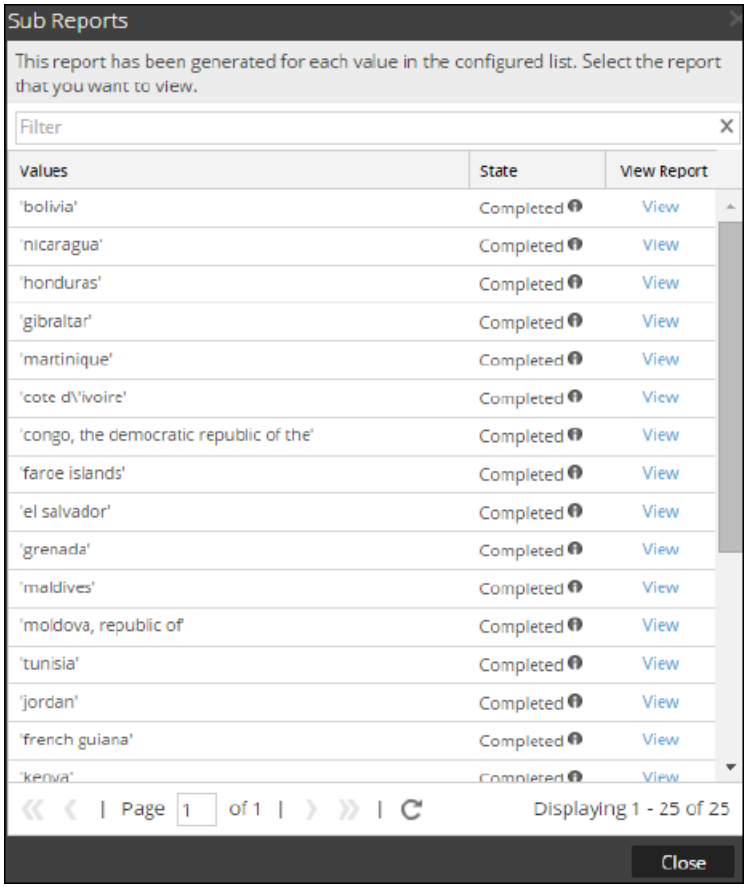
       The list item selected gets added to the **Iterate on List** field.

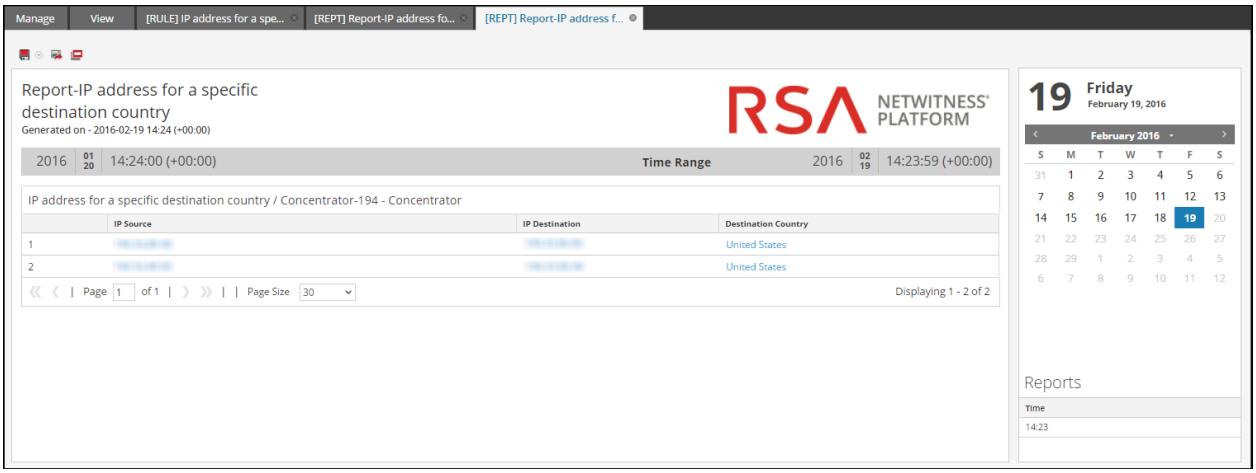
d. Select the variable on which the selected list value has to be applied.



14. Click **Schedule**.
    The scheduled report executes as scheduled and provides the configured outputs.

The following figure shows the Iterative Report view.





Create a Report Using a Rule

You can create a report using a rule. When you create a report using a rule, a default report is created with this single rule. You can further edit the report to add more rules.

To create a report using a rule, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Choose any of the following:

   - Create a report using a rule when you create or edit the rule:

     a. In the **Rules** view, select a rule and click  > **Use**> **Report**.

     The Use Rule dialog is displayed.

   - Select a rule in the Rules panel and click  in the Rule toolbar. From the drop-down menu, select **Use** > **Report**.

   - In the Rules panel click  > **Create Report**.

Note: Custom rules can be used to create a Report and If you select the view for the rule as "Area" or "Pie", a window pops up for **X-Axis** and **Y-Axis** inputs. By default, you can select only the first meta in **X-Axis**.

3. Select **New Report** or **Existing Report** based on your requirement.
4. Click **Select**.

View a Report

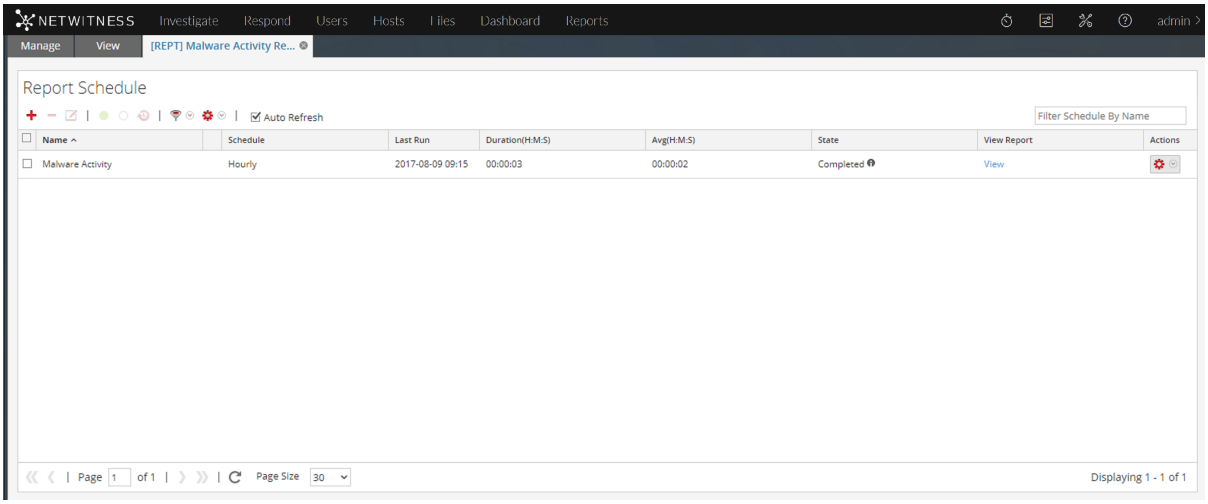
You can view a report or list of all reports. You can also view the scheduled reports to know the state of the scheduled report. If the scheduled report is in a stop or disable state, you can start or enable the scheduled report.

After you view a report, you can perform any of the following tasks:

1. You can print, save, email and view reports on full screen.
2. You can also select a date from the calendar to view a list of successfully run reports for the chosen date.

To view a report, perform the following:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, select a report and in Actions column, click  > **View Scheduled Reports**.
4. Click the **#Schedules** column.
   The Schedule Reports view tab is displayed with the status of each of the scheduled report.



5. Select a scheduled report and click **View**.
   One of the following is displayed:
   - The selected report.
   - The Sub reports panel for a scheduled report having 'Iterative' selected.

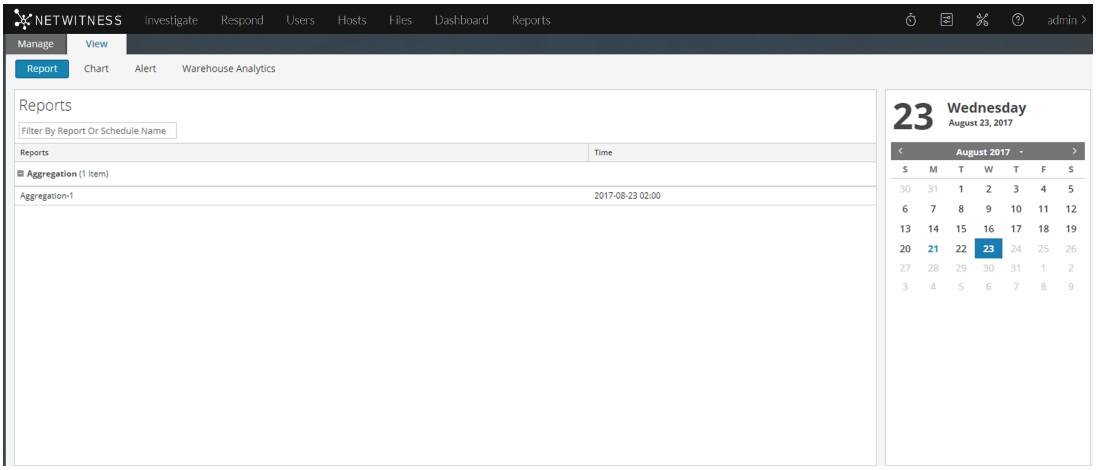
For each value in the configured list a report is displayed.

> **Note:** If the report status is partial or complete, the "last run timestamp" and the "last run (seconds)" are updated. However, the average time taken to run the report is updated only when the report status is complete and not when it is partial.

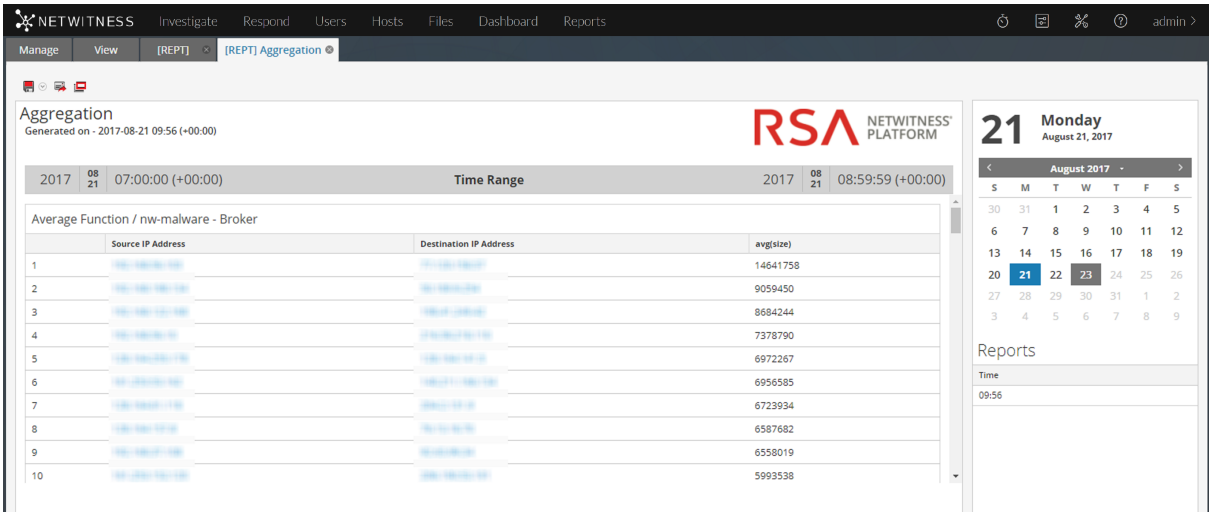
To view a list of all reports, perform the following:

1. Go to **Reports**.
   The **Manage** tab is displayed.
2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, click **View All Reports**.
   A list of reports along with their schedule name and time are displayed on the View tab.

> **Note:** If no list is displayed, select a date from the calendar to view a list of reports for that date.



4. Double-click on a report to view the details of the report.
5. You can select a scheduled report and print, save as PDF/CSV, send email notifications, or view it on full screen.



Investigate a Report

You can investigate a report by directly navigating to the Investigation View from the report. With the Investigate a report option, you can investigate each event mentioned in the report.

To investigate a report, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

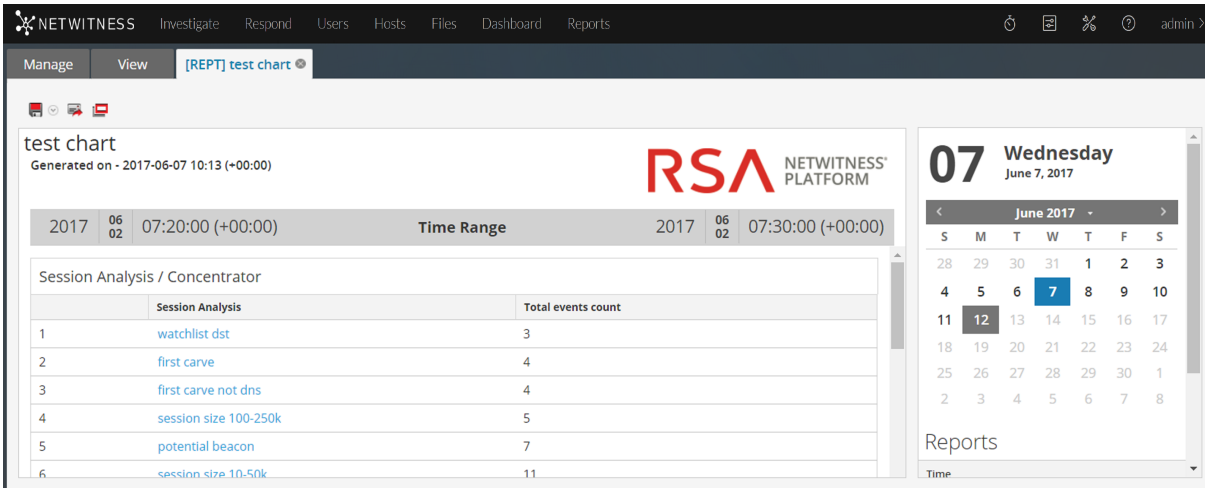
3. In the **Reports** toolbar, click **View All Reports**.

   The View All Reports tab is displayed.

   Note: If no reports are displayed in the View All Reports, select a date for which you want to display the reports.

4. Double-click the report name to view the report details.

   The Report details screen is displayed.



   You can click on the session analysis to investigate on the report.

Note: If you want to manually copy the result data and use it for investigation, make sure that the binary values are prefixed with 'hex:'.

Manage Lists, Rules or Reports

You can set access control, delete, edit, import, or export a list, rule or report.

Manage a List

You can perform the following procedures to manage a list.

- Access Control for a List and List Group
- Edit a List
- Delete a List or List Group
- Duplicate a List
- Export a List or List Group
- Import a List or List Group
- Filter Unused Lists

Access Control for a List and List Group

You can set up the access permissions for the user roles to manage lists or list groups. The Reporting provides access control at the list and list group level. Only a user who has the right set of permissions can perform the tasks in the Reporting. The access control is managed by the administrator from the  (Admin) > Security > Roles tab.

As an administrator, you must ensure that the roles created for specific tasks have access to all the permissions higher in the hierarchy of roles.

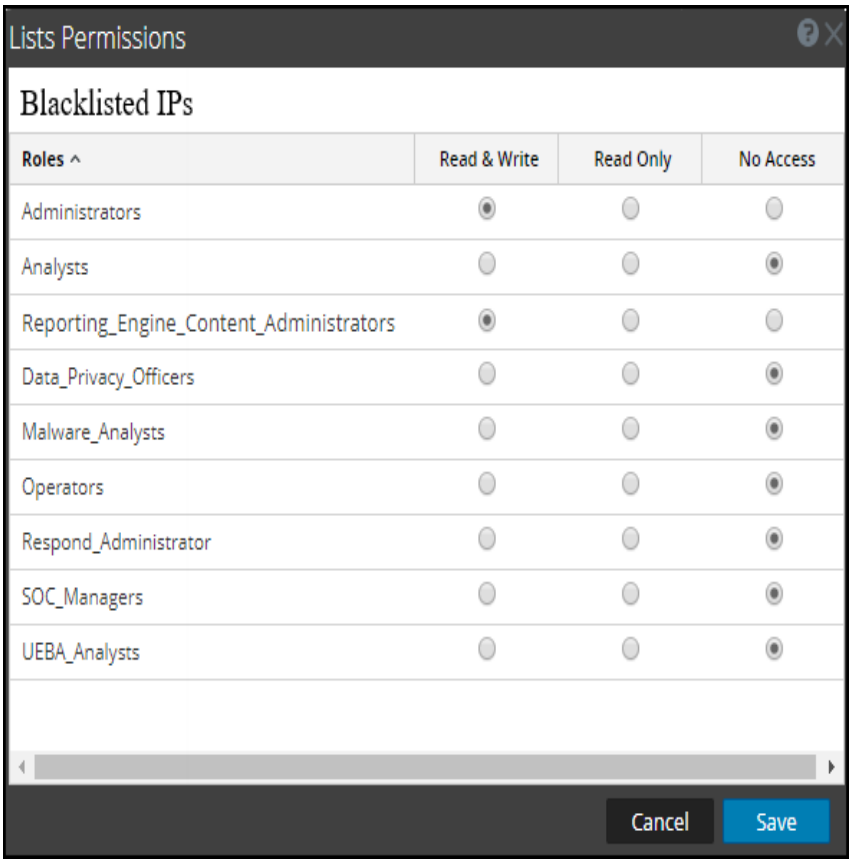
Lists or list groups can be assigned to a specific set of user roles. When users log into NetWitness, they can access only those lists to which they belong. Users who belong to a user role with the **Read & Write** access permission have full access rights on the lists. Further, the access can be strengthened so that lists are accessed only by those who have the **Read Only** access.

Note: You must have **Read Only** permission for a list group to view the lists within that group.

For example, if you want **Security Analysts** to have access to all the lists in a list group, you can set the permission **Read & Write** at the list group level. And, if you do not want the **Operator** role to have access to a specific set of lists in a list group, you can set the permission **No Access** at the list group level.

At the list or list group level, you can set the following access permissions for the user roles in NetWitness. For more information, see List View:

- Read & Write
- Read Only
- No Access



The following table lists the columns in the Lists Permissions panel:

Column	Description
Roles	Describes roles of the users logged into the NetWitness user interface.
Read & Write	Allows users to access, view, edit, delete, import, and export lists on the Lists view. Users can also change the permission on the rule.
Read Only	Allows users to only access and view the list on the lists view.
No Access	Doesn't allow users to access or view the lists.

Access Control for a List

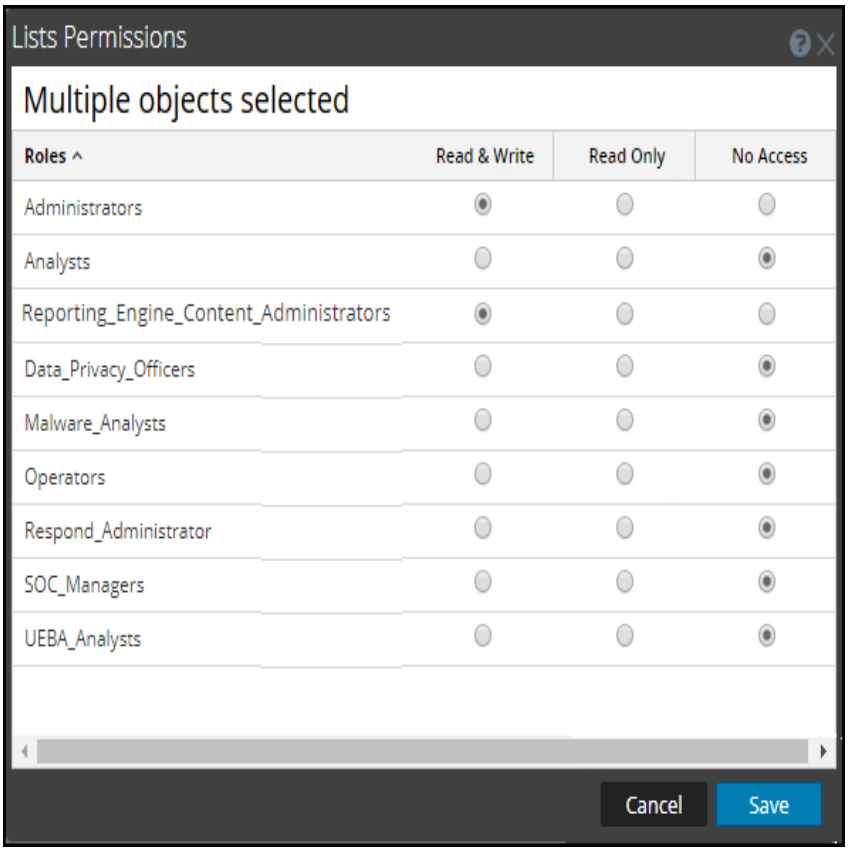
To change the list permissions, you must select a list and set access permissions using the List Permissions panel.

If you want to change the access permission for a specific user role, you must set it at the list level. Except for administrators (with full access in Reporting Engine service Config page) and reporting engine content administrators, the default permission set for all the user roles is **No Access** before applying job permissions.

Access Control Multiple Lists

You can select multiple lists at once and set access permissions using the Lists Permissions Panel. The access permission that you choose is applied to all the selected lists.

Note: The **"*"** beside the role name indicates that other permissions are available for the user role. If you want to change the access permission for the required user role, select the user role and change the access permission.



Note: If a user (other than ADMIN) creates a list, ADMIN cannot access that list.

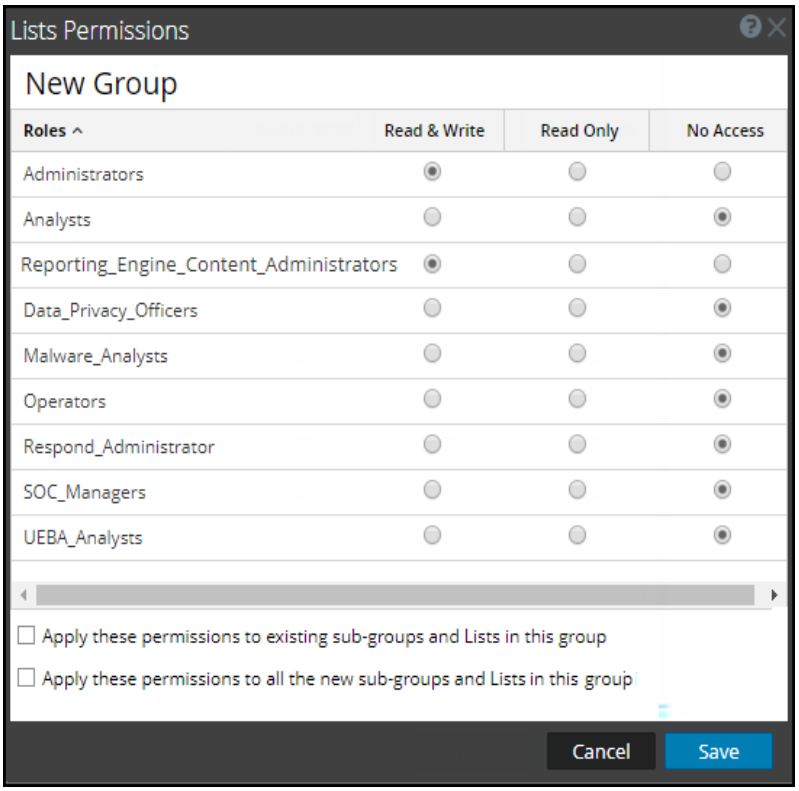
Access Control for a List Group

To change the list group permissions, you must select a list group and set access permissions using the Lists Permissions panel.

If you want to change the access permission for a specific user role, you must set it at the list group level. The default permission set for all the other user roles is **No Access** before applying job permissions.

You can also apply permissions to existing subgroups and lists in the group by selecting the appropriate checkbox.

You can also apply permissions to all the new subgroups and lists in the group by selecting the appropriate checkbox.



The following scenarios describe defining permissions for list groups or subgroups and lists in the groups:

- Scenario 1: Permissions applied to list group or subgroup based on the user role.

  Each of the levels will have a permission set depending on the user role. For example, if a list group is assigned the role of Security Analyst, permissions are set to Read & Write for the list group.

- Scenario 2: Permissions applied to subgroups and lists in the group.

  The access permissions that you set can be applied to subgroups and child objects of this group. Permission at the list group level will be inherited by the subgroups and lists in the group.

Role (Analysts)	Permissions applied to list group or subgroup based on the user role	Permissions applied to subgroup and lists in the group
Group	Read & Write	Read & Write
Subgroup	Read	Read & Write - Inherited
Lists	Read	Read & Write - Inherited

Access permission for a list or list group

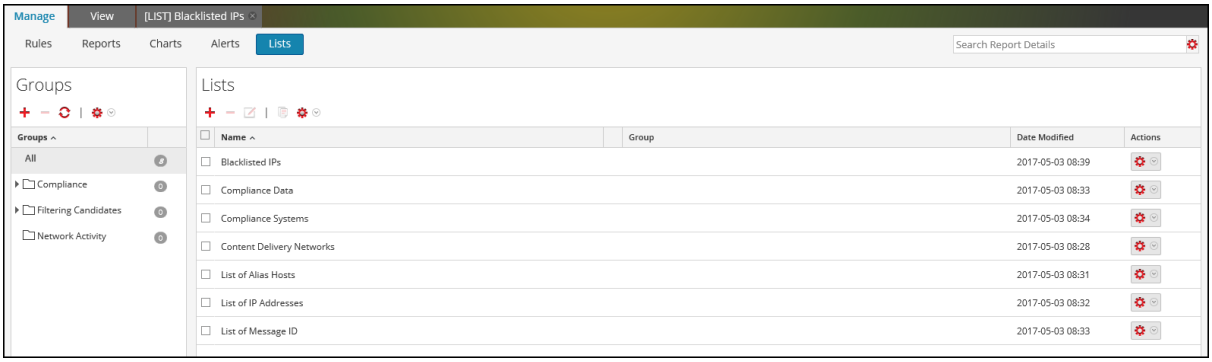
Ensure that you have at least **Read & Write** access permission so that you can set access permissions for lists or list groups.

To set access permission for a list, perform the following:

1. Go to **Reports**.

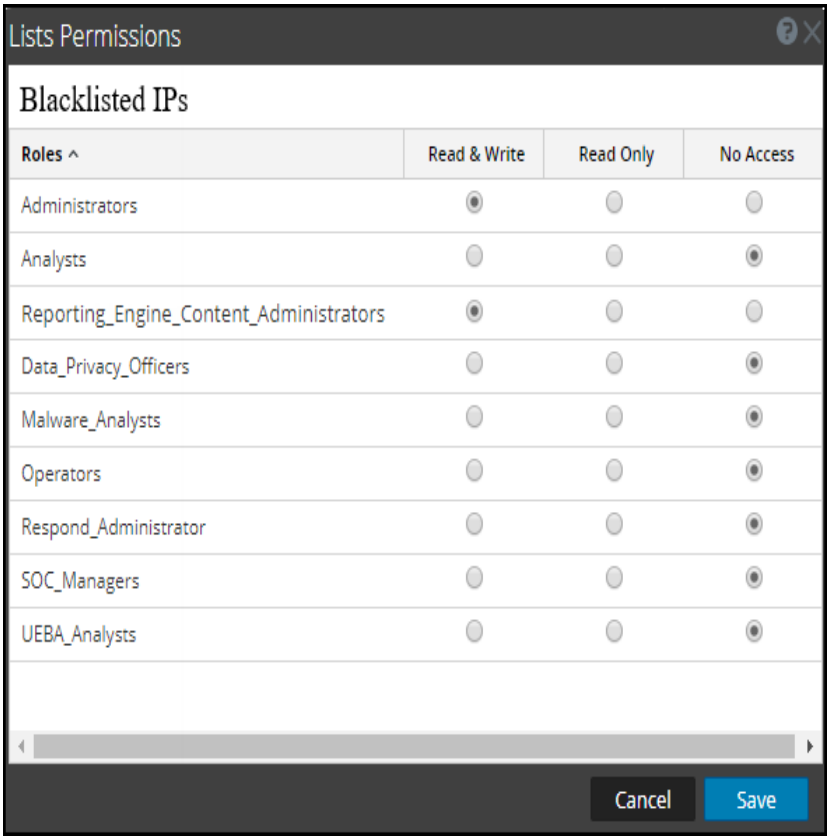
   The Manage tab is displayed.

2. Click **Lists**.
   The List view is displayed.



3. In the **Lists** panel, select a list.

4. Click  > **Permissions** in the List toolbar.
   The List Permissions dialog is displayed.



5. Select the appropriate access permission for each of the user roles and click **Save**.

   A confirmation message that the permission is successfully set for the selected list is displayed.

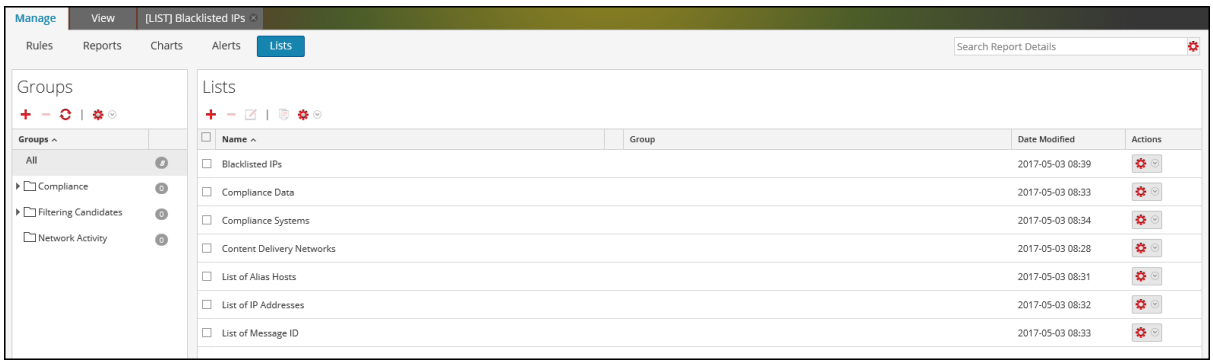
To set access control for a list group, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

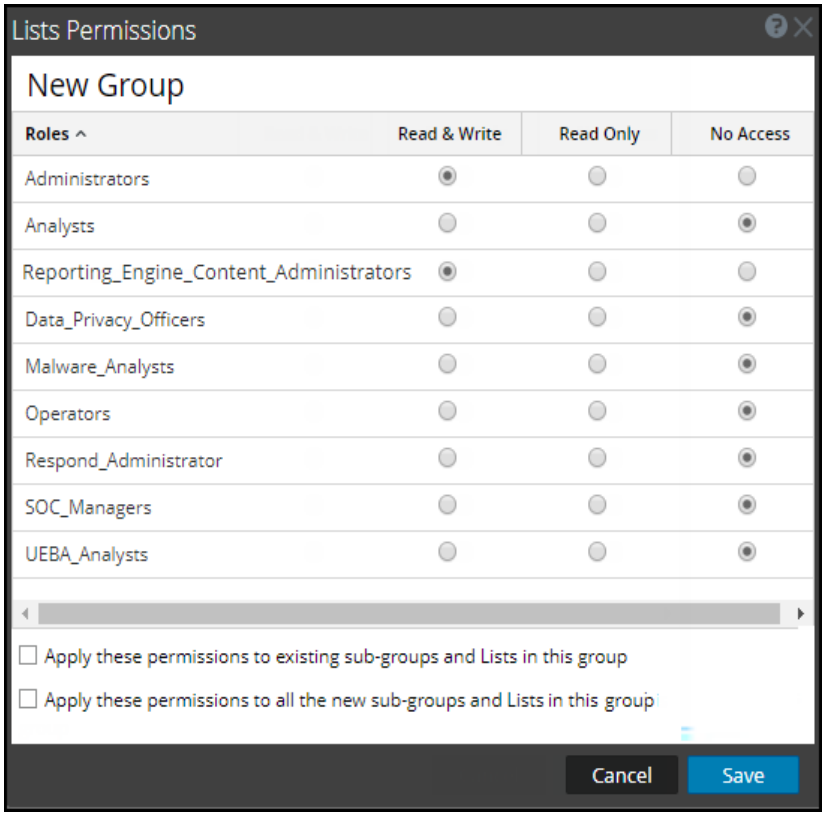
   The List view is displayed.



3. In the **Lists Groups** panel, select a list group.

4. Click  > **Permissions**.

   The List Permissions dialog is displayed.



5. (Optional) Select the appropriate checkbox to apply these permissions to existing subgroups and lists in this group.
6. (Optional) select the appropriate checkbox to apply these permissions to all the new subgroups and lists in this group.
7. Click **Save**.

   A confirmation message that the permission is successfully set for the selected list group is displayed.

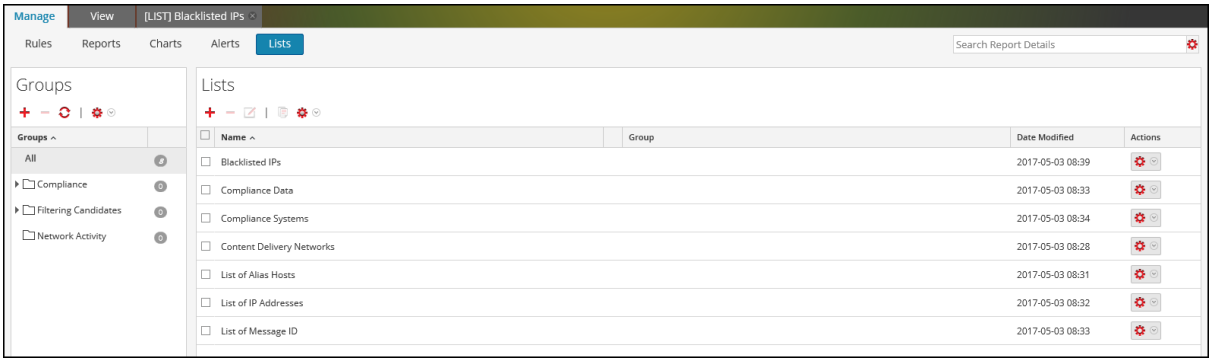
Edit a List

To edit a list, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.
2. Click **Lists**.

   The List view is displayed.



3. In the **Lists** panel, select a list that you want to edit and do one of the following.
   - Click  in the **Lists** toolbar.
   - In the **Lists** panel, click  > **Edit**.

     Note: You can only edit one list at a time.
4. Modify the required fields and add new values to the list.
5. Click **Save**.

   A confirmation message that the list is saved successfully is displayed.

Delete a List or List Group

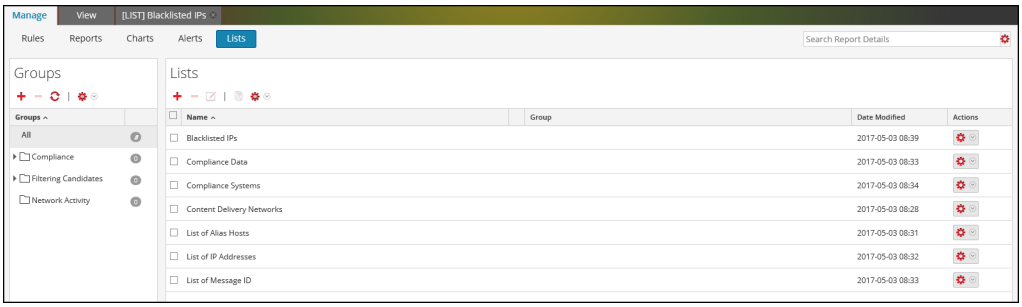
To delete a list, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

   The List view is displayed.



3. In the **Lists** panel, do one of the following:

   - Select a list or multiple lists that you want to delete and click in the **Lists** toolbar.
   - In the **Actions** column, click > **Delete**.

     Note: Before you delete a list, make sure that the list is not associated with any rule.

4. Click **Yes** to delete the list.

   A confirmation message that the list is deleted is displayed and the selected list is deleted from the List View panel.

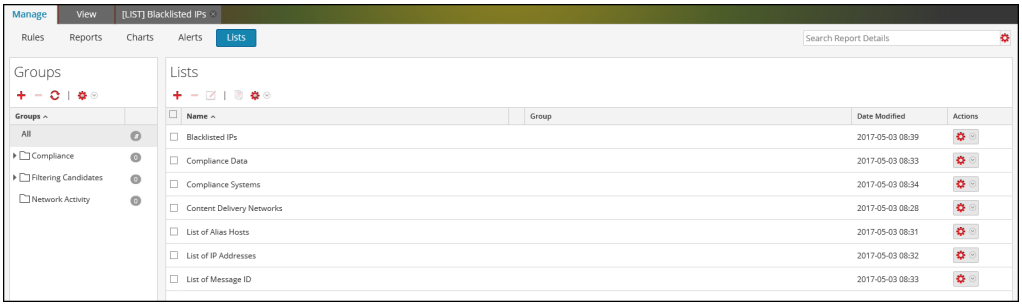
To delete a list group, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

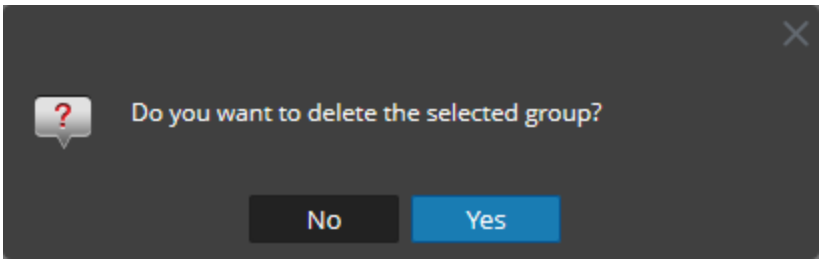
2. Click **Lists**.

   The List view is displayed.



3. In the **Lists Groups** panel, select the group and click .

   A confirmation dialog is displayed.



Caution: If you delete a group, all subgroups and lists in that group are deleted.

4. Click **Yes** to delete the selected group.

Note: If you try to delete a list group that has lists referenced in a rule or an alert, a warning message that **Lists are referenced in a rule** is displayed.

Duplicate a List

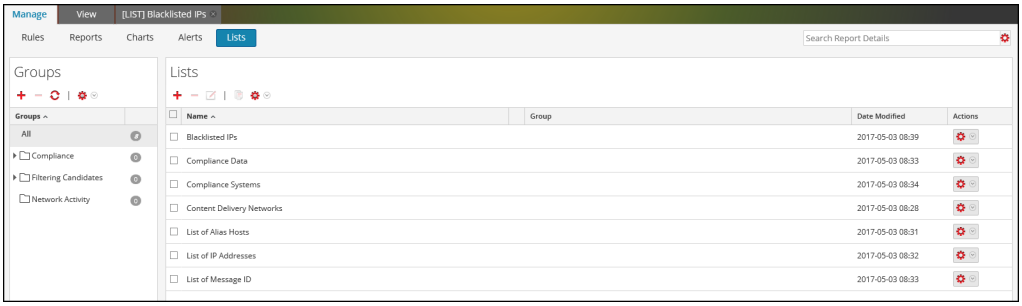
To duplicate a list, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

   The List view is displayed.



3. In the **Lists** panel, select a list that you want to duplicate.

   Note: You can only duplicate one list at a time.

4. In the **Lists** toolbar, click .

Export a List or List Group

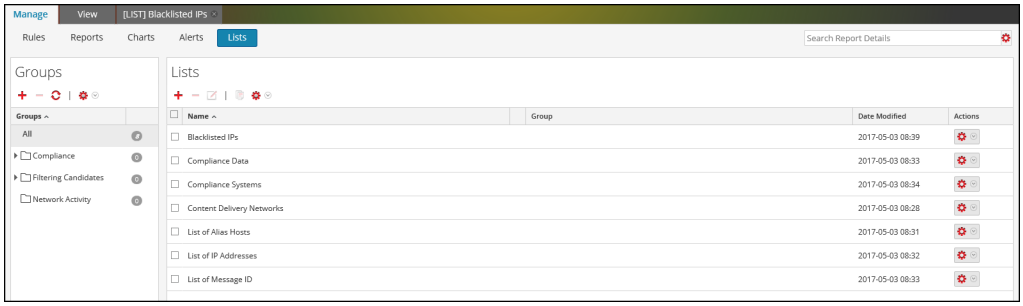
To export a list, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

The List view is displayed.



3. In the **Lists** panel, do one of the following:

   - Select a list and click  **> Export** in the List toolbar.
   - In **Actions** column, click  > **Export**

   You can export multiple lists at a time. To select multiple lists, select the checkbox of the lists to be exported. A browser-specific export dialog may be displayed allowing you to open or save the file.

To export a list group, perform the following:

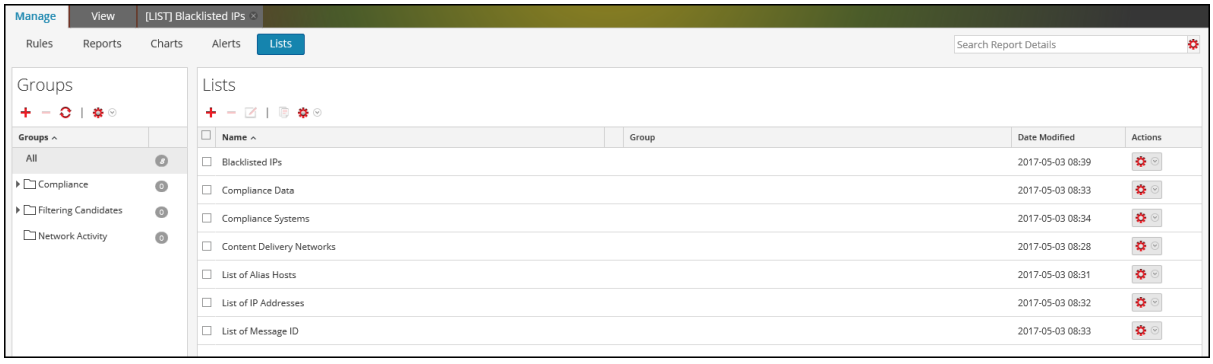
You can export selected list groups to an external file that can be later imported to NetWitness. If nothing is selected in the List Library panel, the entire list tree is exported. When you export, the result is a single export file in binary format.

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

   The List view is displayed.



3. In the **Lists Groups** panel, select the list group containing the lists which you want to export.

4. Click  > **Export**.

   You can export multiple list groups at a time. To select multiple list groups, press and hold the CTRL button and select the list groups to be exported. The exported file is saved to the local drive.

Import a List or List Group

To import a list, perform the following:

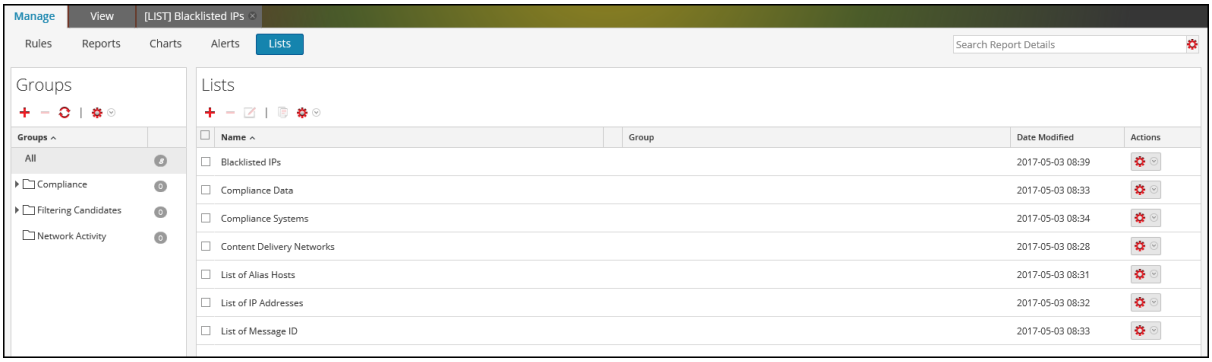
You can import lists from instances of NetWitness into the list tree in the List View panel. Lists must be in a valid binary file exported from a NetWitness instance.

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

   The List view is displayed.

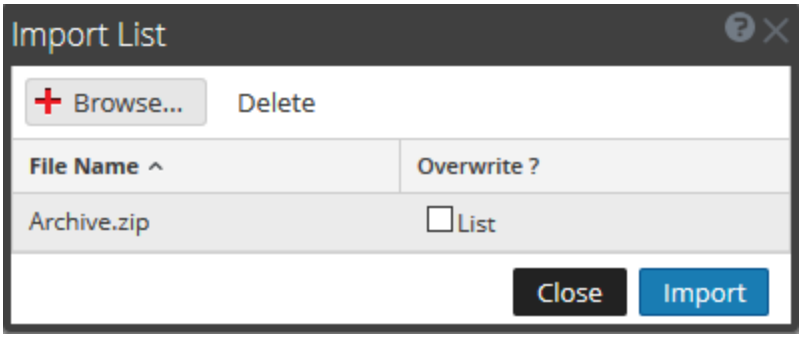


3. In the **Lists** toolbar, click  > **Import**.

   The Import List dialog is displayed.

   Note: You can import multiple lists at a time. To select multiple lists, press and hold the CTRL button and select the lists to be imported.

4. Click **Browse** and select archived file containing the lists.



5. Click **Import**.

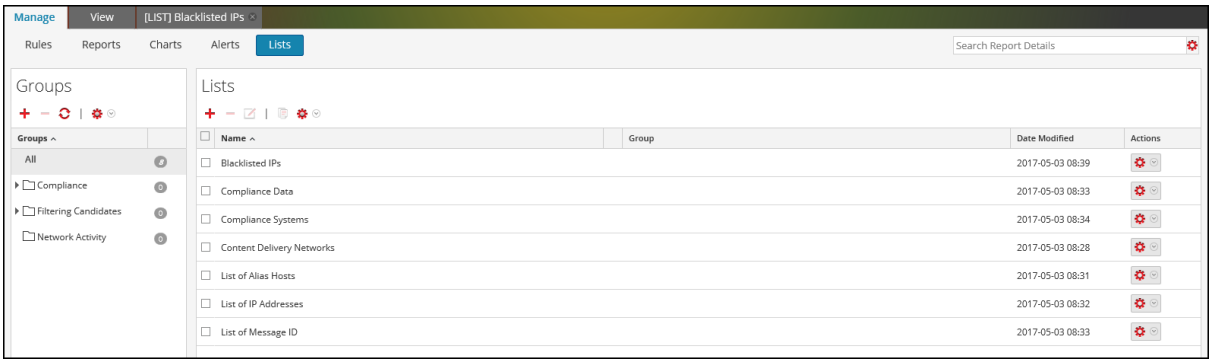
Note: During the import process, if a duplicate list exists and you do not select the overwrite option, the list is imported and no message about duplicate lists is displayed.

To import a list group, perform the following:

You can import list groups from instances of NetWitness into the list tree in the List Groups panel. Lists must be in a valid binary file exported from a NetWitness instance.

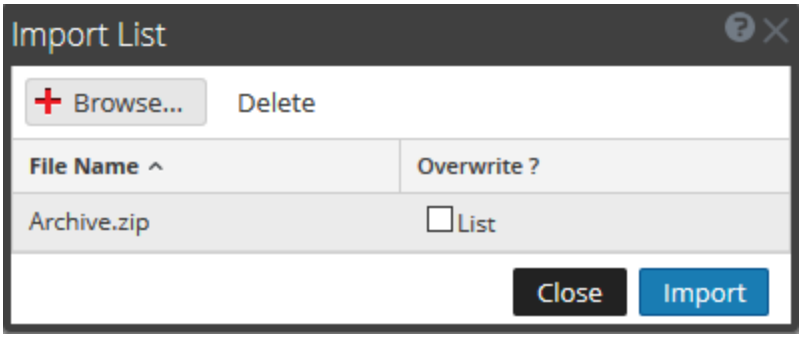
1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Lists**.

   The List view is displayed.



3. In the **Lists Groups** panel, click  > **Import**.

   The Import List dialog box is displayed.
4. Click **Browse** and select archived file containing the list groups.



   You can import multiple list groups at a time. To select multiple list groups, press and hold the CTRL button and select the list groups to be imported.
5. Click **Import**.

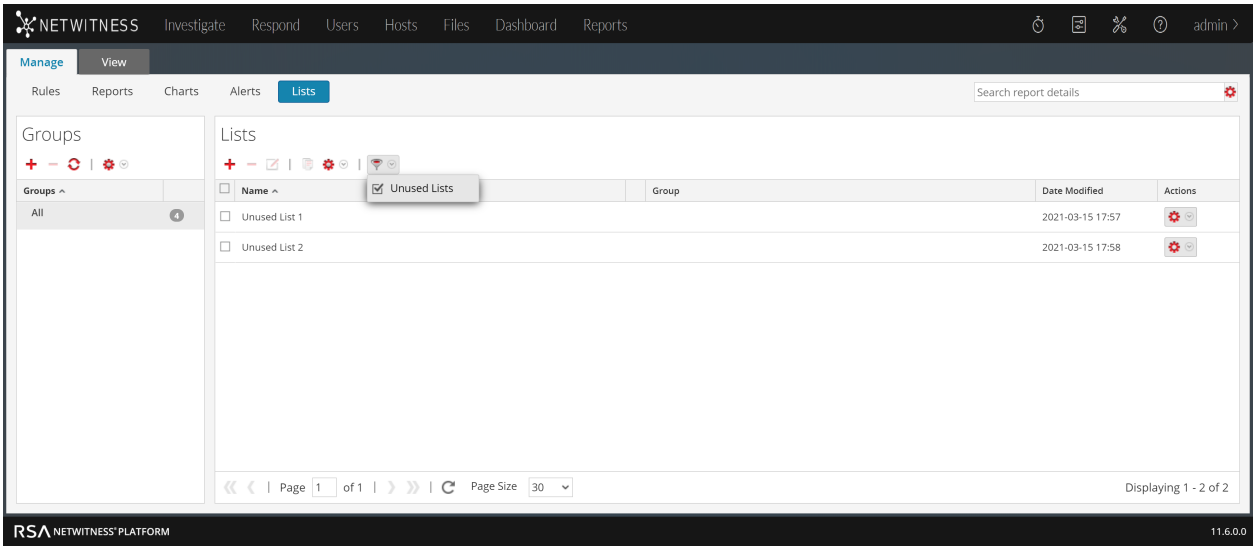
Note: During the import process, if a duplicate list group exists and you do not select the overwrite option, the list group is imported and no message about duplicate list group is displayed.

Filter Unused Lists

You can filter and delete the lists that are not used by any rule. To filter unused lists:

1. Go to **Reports**.
   The **Manage** tab is displayed.
2. In the **Lists** panel, click  and select **Unused Lists**.
   List of unused lists will get displayed.

3. [Optional] Delete the unused lists.
   For more information, see Delete a List or List Group.



Manage a Rule

You can perform the following procedures to manage a rule.

- Access Control for a Rule and Rule Group
- Delete a Rule or Rule Group
- Duplicate a Rule
- Edit a Rule
- View Dependents of a Rule
- Export a Rule or Rule Group
- Filter Unused Rules

Access Control for a Rule and Rule Group

To set access permissions the user will have depending on the user role to manage a rule or rule group. The Reporting provides access control at the rule and rule group level. Only a user who has the right set of permissions can perform the tasks in the Reporting. The access control is managed by the administrator from the **(Admin) > Security > Roles** tab.

When creating users and user roles, the administrator must ensure that the roles created for specific tasks have access to all the permissions higher in the hierarchy of roles.

Rules or Rule Groups can be tied to a specific set of user roles so that when a user logs into NetWitness, the only rules they can access are rules accessible to the group to which the user belongs. Users that belong to a user role with the ‘Read & Write’ access permission have full access rights on the rule. Further, the access can be tightened so that rules are accessed only by those who have the ‘Read Only’ access.

Note: You must at least have ‘Read Only’ permission on a group to view the rules within that group.

At the rule level, you can set the following access permissions for the user roles:

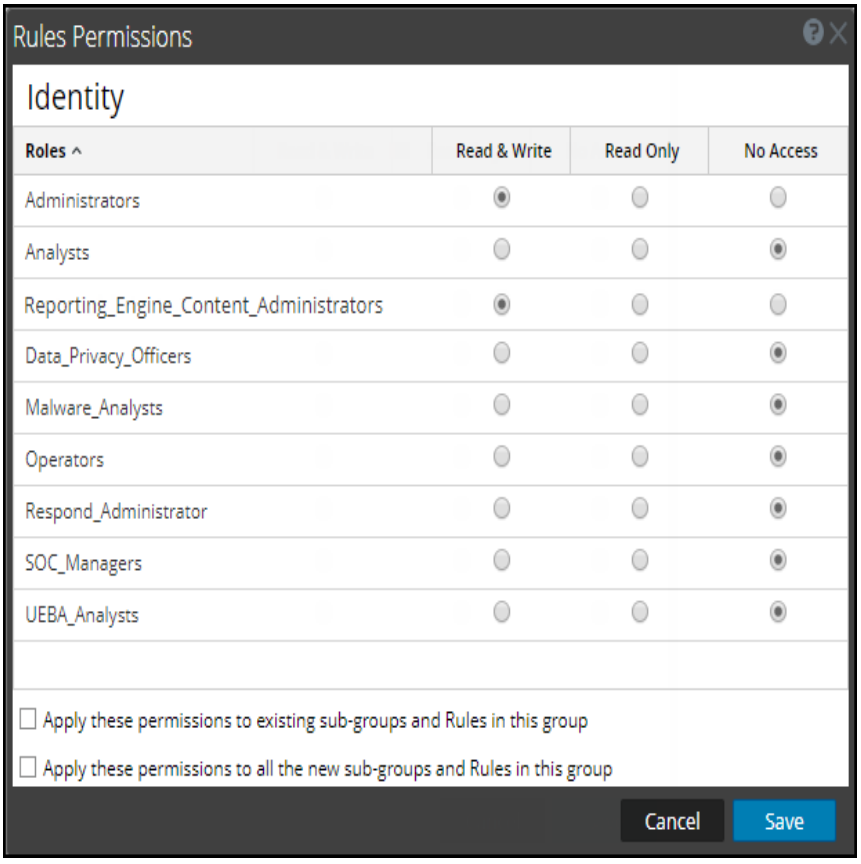
- Read & Write
- Read Only
- No Access

Suppose, you want the **Security Analysts** to have access to all the rules in a Rule Group, you can set the permission '**Read & Write**' at the Rule Group level. And, if you do not want the **Operator** role to have access to a specific set of rules in a rule group, you can set the permission '**No Access**' at the Rule Group level. The permission is set only for the rule group but not the rules or subgroups in the Rule Group.

Access Control for a Rule Group

When you want to change the rule group permissions, you must select a rule group and set access permissions using the Rule Permissions panel.

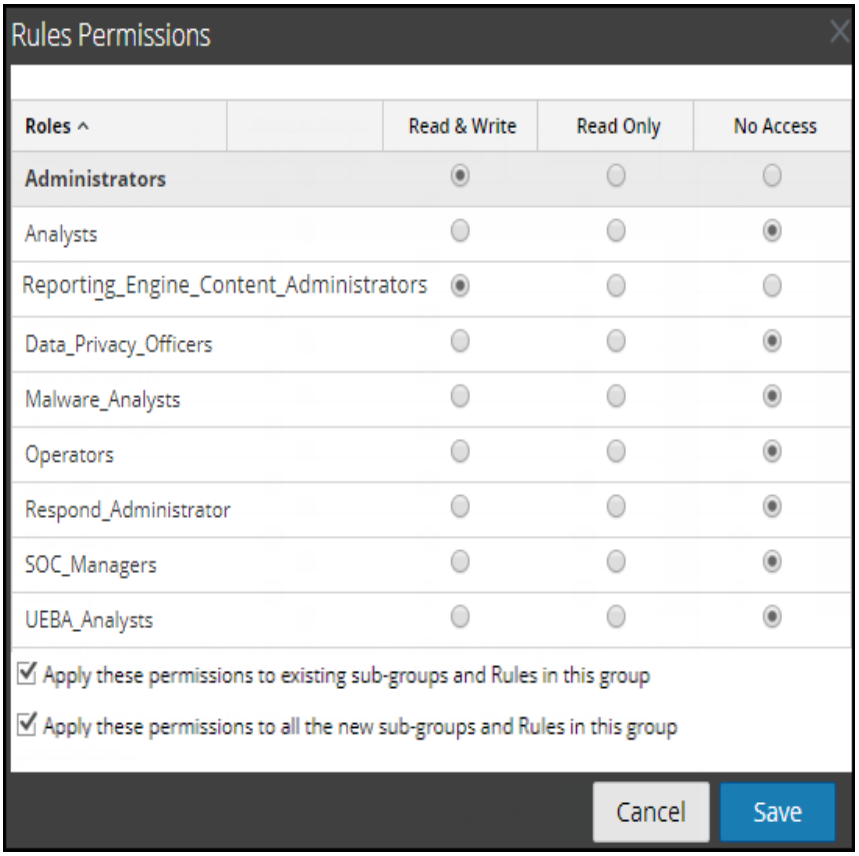
Before applying rule group permissions, the default permission set for all the user roles is 'No Access' permission, and the checkboxes are deselected.



If you want to change the access permission for a specific user role, you must set these at the rule group level, as shown in the figure. Suppose, you want the **Administrators** to have access to all the rules in a Rule Group, you can set the permission '**Read & Write**' in the Rule Group Permissions panel.

You can also apply permissions to existing subgroups and Rules in the group by selecting the appropriate checkbox.

You can also apply permissions to all the new subgroups and Rules in the group by selecting the appropriate checkbox.



You can also apply permissions to subgroups and rules in the group by selecting the checkbox.

The two scenarios are explained in brief:

- Scenario 1: Permissions applied to Rule Group/ Sub Group/ Rules based on the user role.
- Scenario 2: Permissions applied to Sub Group and Rules in the Group.

Role (Analysts)	Permissions applied to Rule Group/ Sub Group/ Rules based on the user role	Permissions applied to Sub group and Rules in the Group
Group	Read & Write	Read & Write
Sub Group	Read	Read & Write - Inherited
Rules	Read	Read & Write - Inherited

The access permissions that you set can be applied to subgroups and child objects of this group.

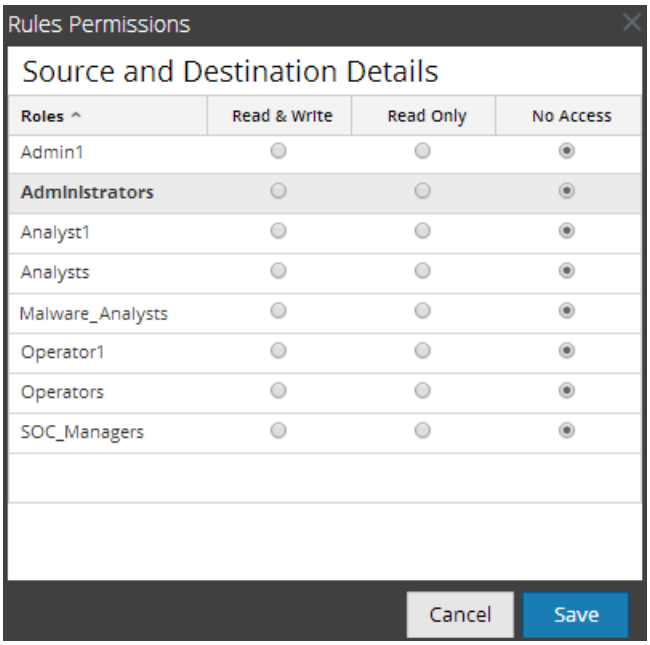
The Rule Group will be assigned the role of a **Security Analyst** and permissions are set to **Read & Write** rule group.

For scenario 1, each of the levels will have a permission set depending on the user role. For scenario 2, the permission at the Rule Group level will be inherited by the Sub Group and Rules in the Group.

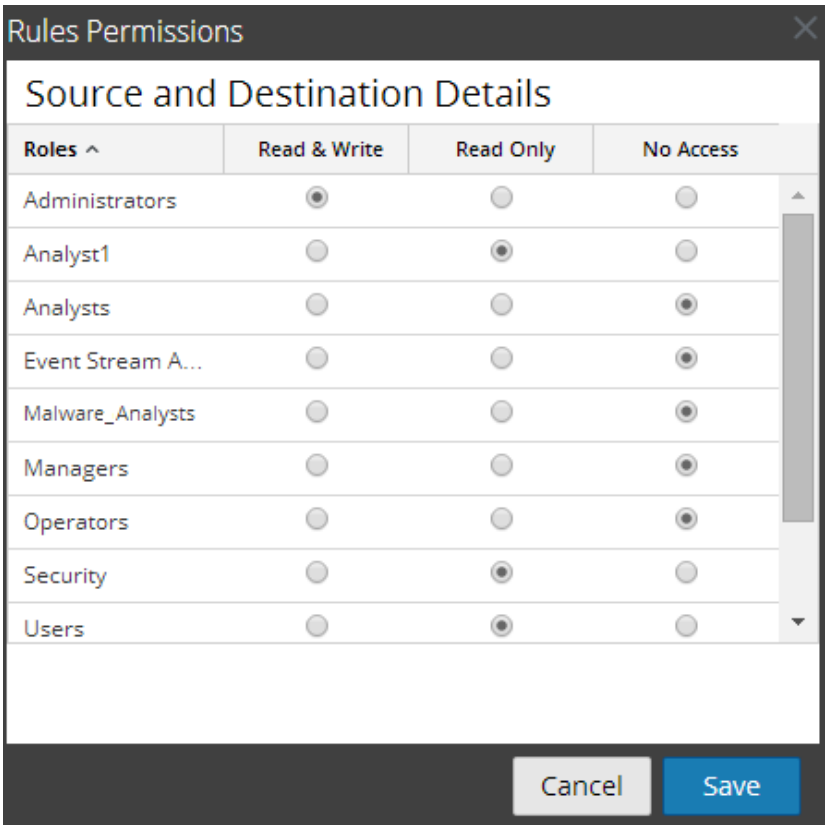
Access Control for a Rule

When you want to change the rule permissions, you must select a rule and set their access permissions using the Rule Permissions panel.

Before applying the Rule permissions, the default permission set for all the user roles is 'No Access' permission and the checkbox is deselected.



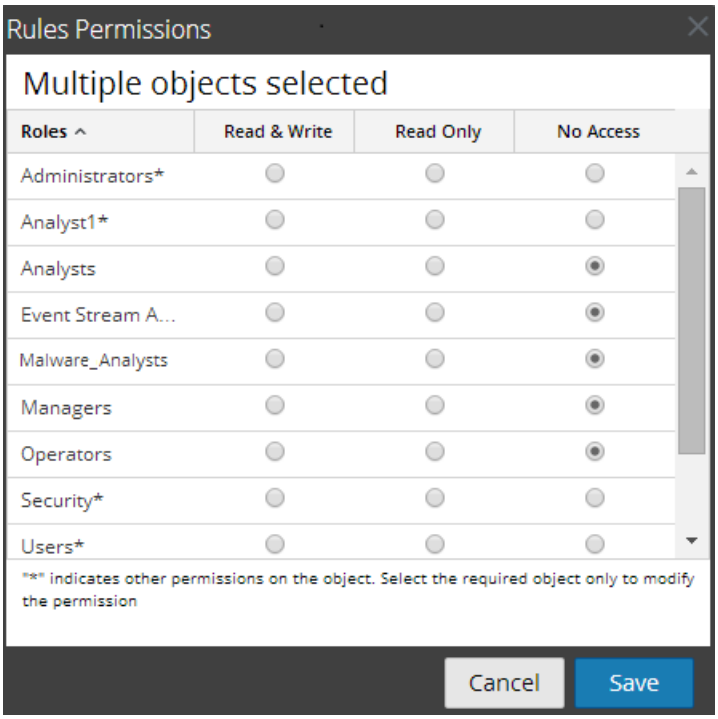
If you want to change the access permission for a specific user role, you must set these at the rule level, as shown in the figure. Suppose, you want the **Administrators** to have access to a specific rule, you can set the permission '**Read & Write**' in the Rule Permissions panel.



Access Control for a Rule When Multiple Rules are Selected

When you want to change permissions of multiple rules, you can select multiple rules at a time and set their access permissions using the Rules Permissions Panel. The access permission that you choose will be applied to all the selected rules.

Note: The '*' besides the role name indicates the other permissions available on the user role. If you want to change the access permission for the required user role, select the user role and change the access permission.



Log in as a specific user and view the access details

When you log in to the NetWitness UI as a user having 'Read access' permission, all the rules will be denoted with the symbol (📖) and when you click on the symbol the 'Read Only' callout is displayed on the Rules panel.

When you log in to the NetWitness

UI as a user not having 'Read & Write' access permission on a Rule, all the rules will be denoted with the symbol ( ) and the rules appear grayed out on the Rules List panel.

The following figure shows the Rules panel when logged in with minimal 'Read & Write' access permission.

Name	Type	Group	Date Modified	Actions
*(raw_log)-RULE	Warehouse	Aggregate Function	2014-07-13 09:46	
[redacted]	Warehouse	Regular	2014-07-16 07:34	
Accounts Created	NetWitness DB	Identity Management	2014-07-14 10:56	
Accounts Created SAW	Warehouse	Compliance_old	2014-07-14 09:40	
Accounts Created SAW	Warehouse	Warehouse	2014-07-25 09:48	
Accounts Created SAW(1)	Warehouse	Warehouse	2014-07-25 09:54	
Accounts Deleted	NetWitness DB	Identity Management	2014-06-26 08:35	

Note: If a user (other than administrator) creates a rule, ADMIN cannot access that rule.

Tabular Listing

The following table lists the columns in the Rules Permissions panel:

Column	Description
Roles	The role of the user logged into the NetWitness user interface.
Read & Write	The user can access, view, edit, delete, import, and export rules on the Rules view. The user can also change the permission on the rule.
Read Only	The user can only access and view the rule on the Rules view
No Access	The user cannot access or view the rule for which this permission is set.

Set Access Control for a Rule

You can set access control for a rule. The Reporting Engine provides access control at the rule level. Only a user who has the right set of permissions can perform tasks on the rule. The administrator when creating users and roles must ensure that the roles created for specific tasks have access to all the permissions higher in the hierarchy of roles.

At the rule level, you can set the following access permissions for the user roles in NetWitness:

- Read & Write – View or edit the rules in the rule group.
- Read Only – View the rules in the rule group.
- No Access – Cannot view or edit the rules in the rule group.

Prerequisites

Make sure that you have a minimal 'Read & Write' access permission to set access permissions for a rule.

To set access control for a rule, perform the following:

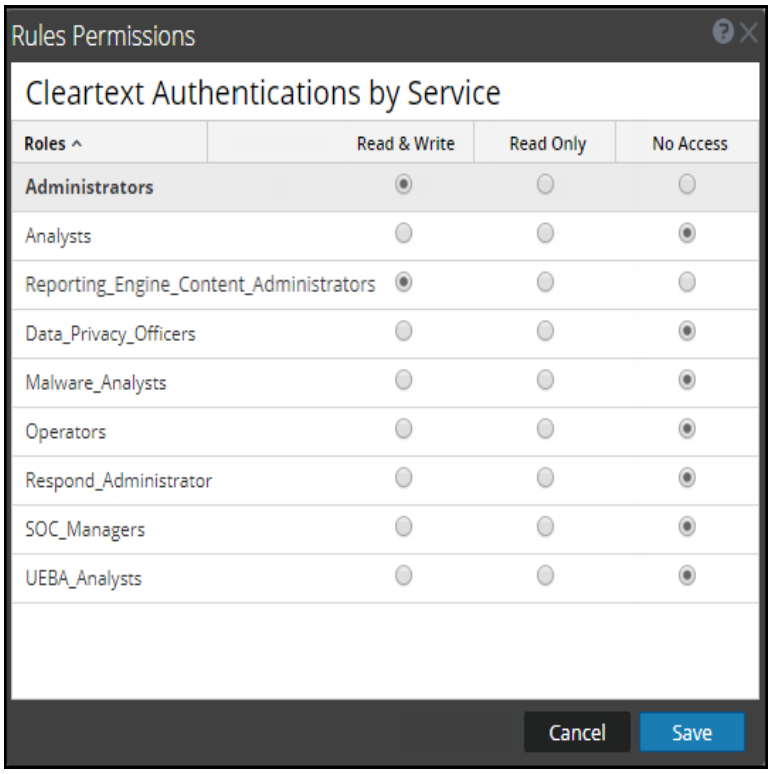
1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** panel, select the rule.

3. Click  > **Permissions** in the Rule toolbar.

   The **Rules Permissions** dialog is displayed.



4. Select the following appropriate access permission for the user role and click **Save**.
   - Read & Write
   - Read Only
   - No Access

Set Access Control for a Rule Group

You can set access control at the rule group level. Only a user who has the right set of permissions can perform the tasks on the rule. The administrator when creating users and roles must ensure that the roles created for specific tasks have access to all the permissions higher in the hierarchy of roles.

At the rule group level, you can set the following access permissions for the user roles in NetWitness:

- Read & Write – View or edit the rules in the rule group.
- Read Only – View the rules in the rule group.
- No Access – Cannot view or edit the rule in the rule groups.

Prerequisites

Make sure that you have a minimal 'Read & Write' access permission to set access permissions for a rule group.

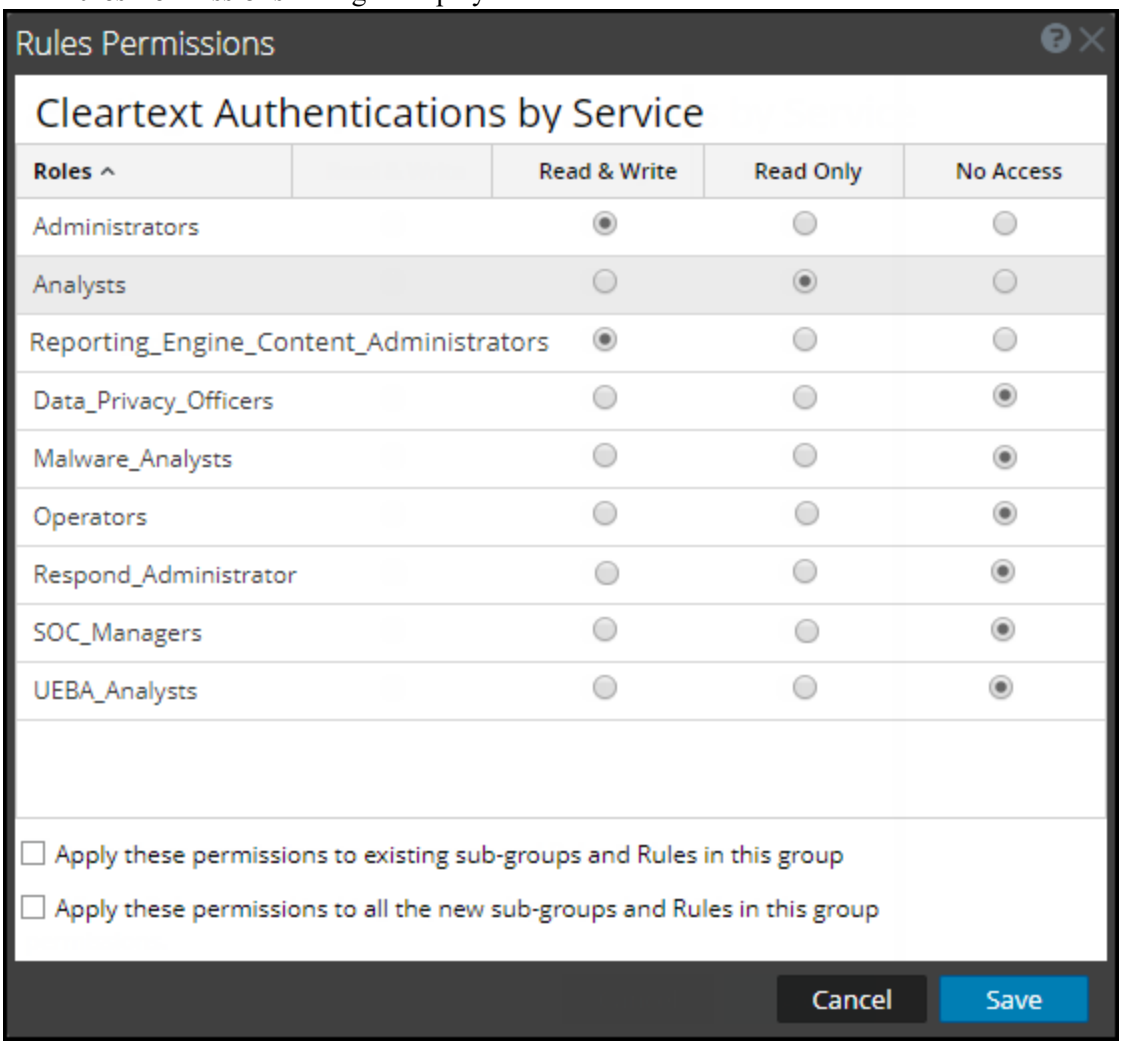
To set access control for a rule group, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules Groups** panel, select the rule group and Click and select **Permissions**.

   The **Rules Permissions** dialog is displayed.



3. (Optional) Select the appropriate checkbox to apply these permissions to existing subgroups and Rules in this group.
4. (Optional) Select the appropriate checkbox to apply these permissions to all the new subgroups and Rules created in this group.
5. Click **Save**.

   A confirmation message that permission is successfully set for the selected rule group is displayed.

Delete a Rule or Rule Group

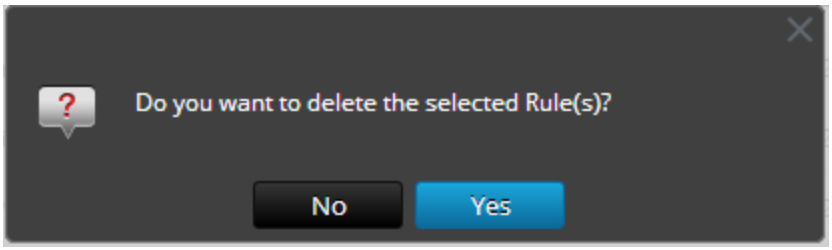
To delete a rule, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** panel, do one of the following.
   - Select a rule and click − in the Rule toolbar.
   - Click ⚙ > **Delete**.

     A confirmation dialog is displayed.



   Note: If a rule is being used in a report, a warning that the rule is in use and cannot be deleted is displayed.

3. Click **Yes** to delete the rule.

   A confirmation message that the rule is deleted successfully is displayed and the selected rule is deleted from the Rules panel.

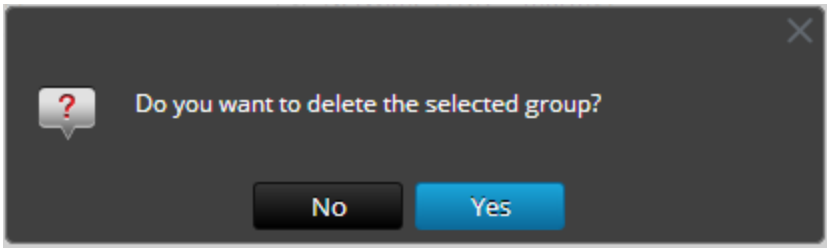
To delete a rule group, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules Groups** panel, select the rule group that you want to delete.
3. Click −.

   A confirmation dialog is displayed.



   Note: If any one of the rules in the group is being used in reports, a warning that the rule is in use and cannot be deleted is displayed.

4. Click **Yes** to delete the group.

   A confirmation message that the group is deleted successfully is displayed and the selected group is deleted from the Rule Groups panel.

Duplicate a Rule

To duplicate a rule, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** panel, select a rule that you want to duplicate.
3. In the Rule toolbar, click .

Edit a Rule

To edit a rule, perform the following:

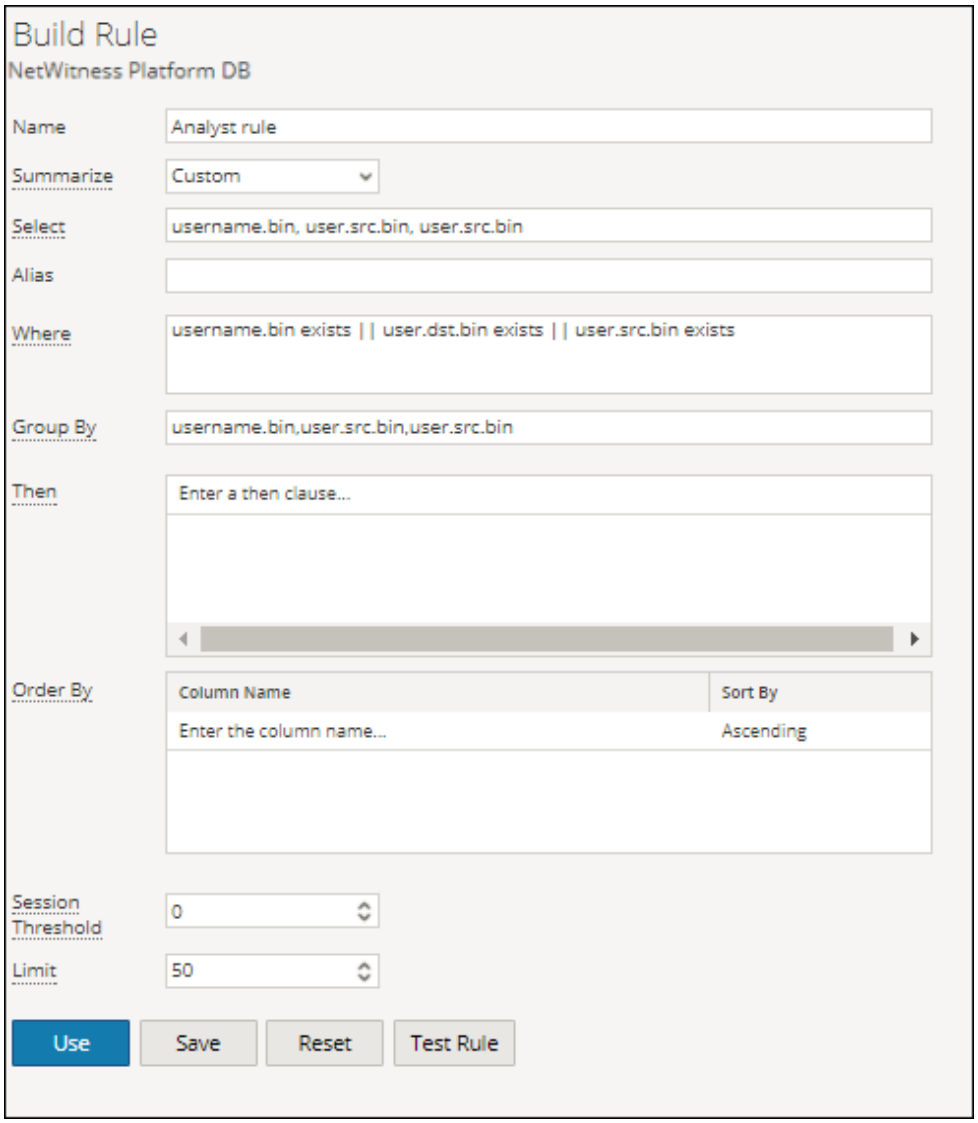
1. Go to **Reports**.

   The Manage tab is displayed.

2. In the **Rules** panel, do one of the following:

   - Select a rule and click  in the Rule toolbar.
   - Click  > **Edit**.

     The Build Rule view tab is displayed.



Note: If a rule is edited, the updated rule definition is applied to the Reports, Charts, and Alerts where the rule is included.

3. Modify the required fields.
4. Click **Save**.

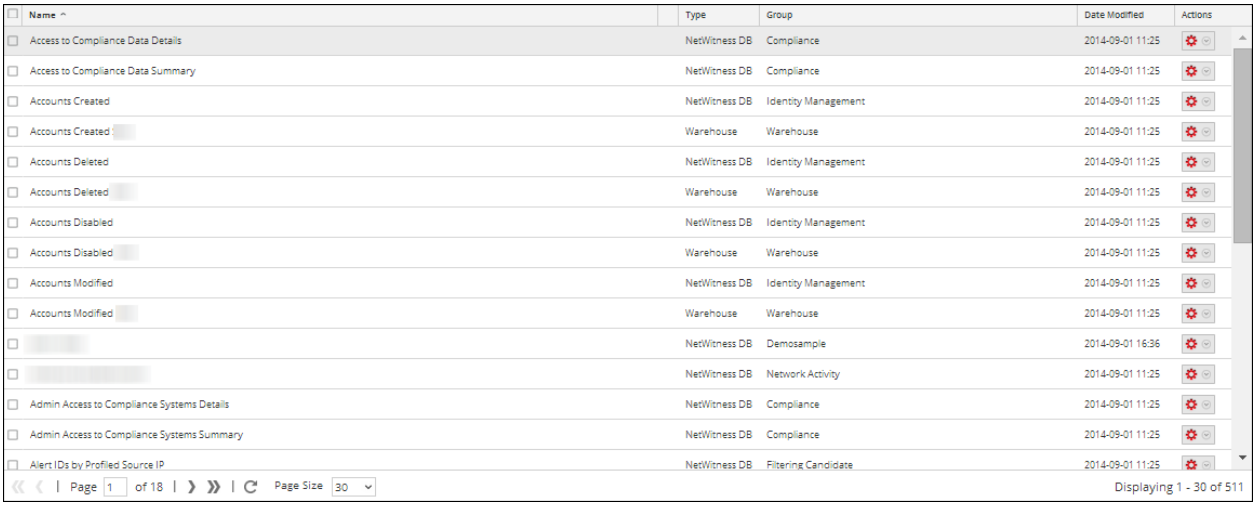
   A confirmation message that the rule is saved successfully is displayed.

When you edit a rule, ensure to re-select the Rule for which you want the Chart to be generated, so that the edited rule is applied. If you do not re-select the Rule and attempt to save or test the rule, the rule is saved and a warning message is displayed.

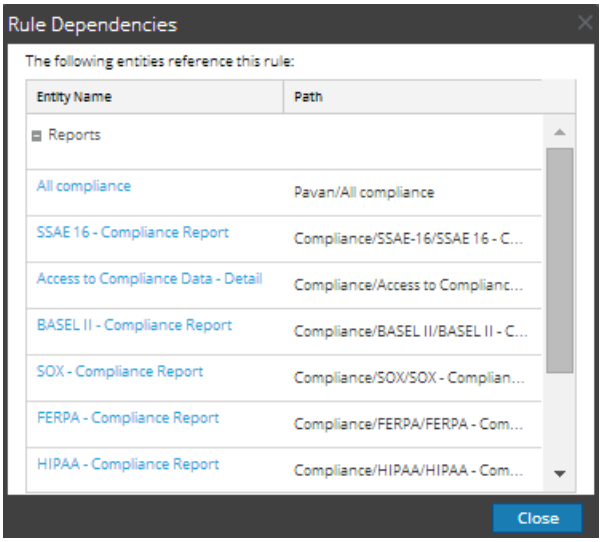
View Dependents of a Rule

You can view dependents of a rule. You must traverse a rule list, select a rule for which you want to identify the dependency over a report, chart, or alert.

The following figure shows the Rule View where you select the rule 'Access to Compliance Data Details'.



The following figure shows the dependency of the rule over alerts and reports.



The following table lists the various columns in the Rule Dependencies dialog and their description.

Column	Description
Entity Name	The name of the entity referencing the rule.
Path	The path where the entity is located in the user interface.

To view dependents of a rule, perform the following:

1. Go to **Reports**.

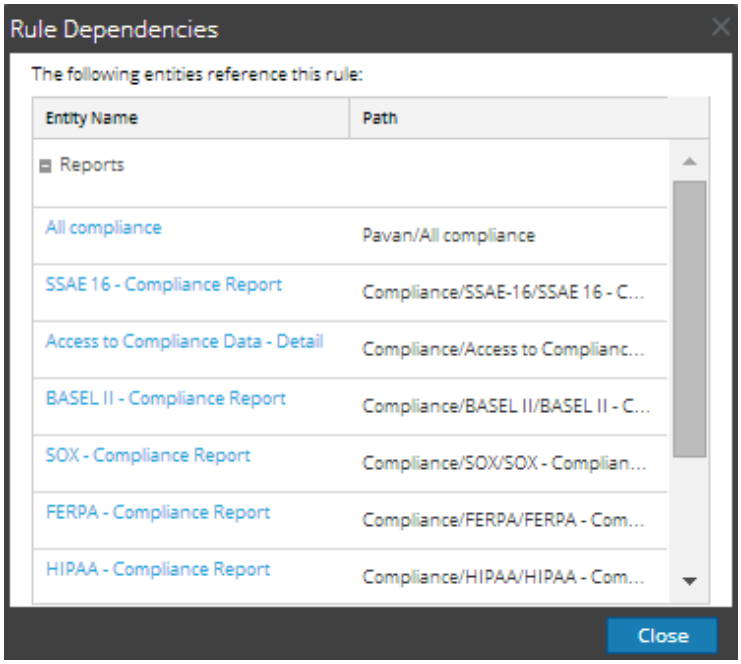
   The Manage tab is displayed.

2. Click **Rules**.

   The Rule view is displayed.

3. In the **Rules** panel, select a rule and in Actions column, click > **Dependents**.

   The Rule Dependencies dialog is displayed.



Export a Rule or Rule Group

Note: Make sure that you have rules in the rule group.

To export a rule, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

   In the **Rules** panel, do one of the following:

   - Select a rule and click > **Export** in the Rule toolbar.
   - Click > **Export**.

2. A browser-specific export dialog may be displayed, allowing you to open or save the file. You can export multiple rules at a time. To select multiple rule, press and hold the CTRL button and select the rules to be exported.

To export a rule group, perform the following :

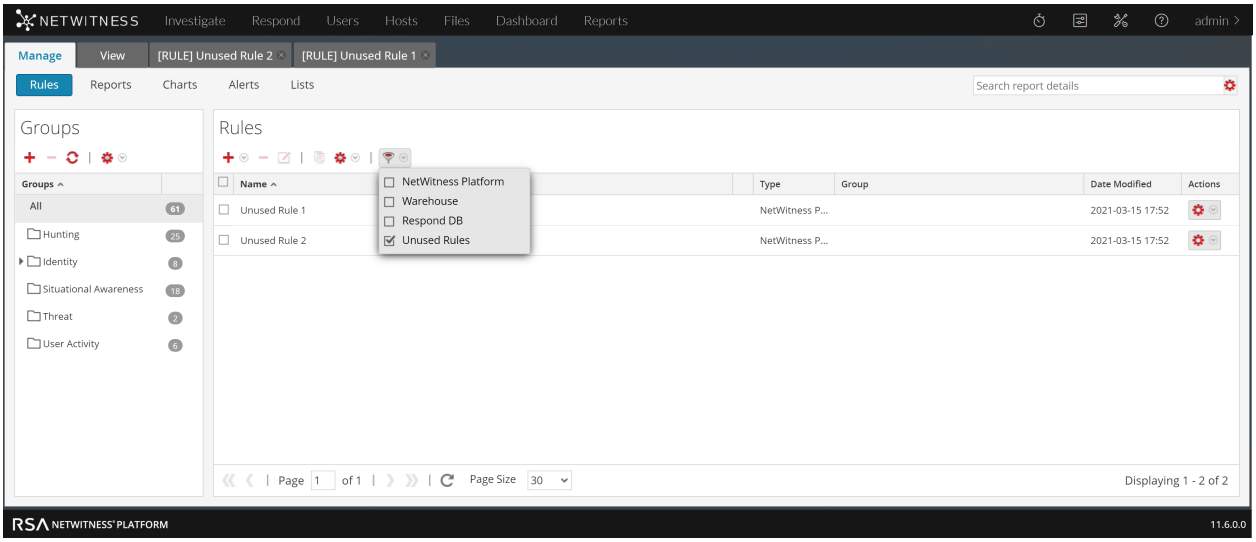
1. Go to **Reports**.
   The Manage tab is displayed.
2. In the **Rules Groups** panel, select the rule group containing the rules which you want to export. You can export multiple rules groups at a time. To select multiple rule groups, press and hold the CTRL button and select the rules groups to be exported.

3. Click  > **Export**.
   A browser-specific export dialog may be displayed allowing you to open or save the file.

Filter Unused Rules

You can filter and delete the rules that are not used by any group in charts, alerts, and reports. To filter unused rules:

1. Go to **Reports**.
   The **Manage** tab is displayed.
2. In the **Rules** panel, click  and select **Unused Rules**.
   List of unused rules is displayed.
3. [Optional] Delete the unused rules.
   For more information, see Delete a Rule or Rule Group.



Manage a Report

You can perform the following procedures to manage a report.

- Access Control for a Report or Report Group
- Delete a Report or Report Group
- Duplicate a Report
- Edit a Report
- Refresh a Report Group or Report List
- Edit a Scheduled Report
- Delete a Scheduled Report

- Export a Report
- Export a Report Group
- Import a Report or Report Group
- Enable or Disable a Scheduled Report
- Start or Stop a Scheduled Report
- View an Execution History of a Scheduled Report
- Manage and Select a Report Logo
- Search Reporting Details

Access Control for a Report or Report Group

This section covers the access permissions the user has depending on the user role to manage a report and report group. The Reporting provides access control at the report and report group level. The user who has the right set of permissions can only perform the tasks in reporting module. The access control is managed by the administrator from the **(Admin) > Security > Roles** tab.

When creating users and user roles, the administrator must ensure that the roles created for specific tasks have access to all the permissions higher in the hierarchy of roles.

Reports and Report Groups can be tied to a specific set of user roles so that when a user logs into NetWitness, the reports with the access rights for the specific user role can be viewed. Users that belong to a user role with the ‘Read & Write’ access permission can define reports. Further, the access can be tightened so that reports are accessed only by those who have the ‘Read Only’ access.

Note: You must have ‘Read Only’ permission for a group to view the reports within that group.

At the report level, you can set the following access permissions for the user roles in NetWitness:

- Read & Write
- Read Only
- No Access

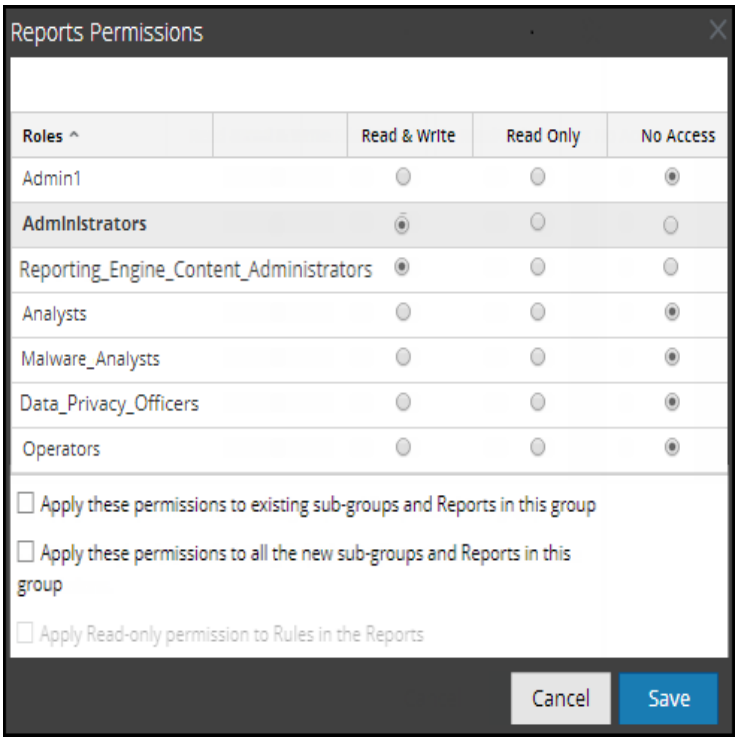
Suppose, you want the NetWitness to have access to all the reports in a Report Group, you can set the permission '**Read & Write**' at the Report Group level. And, if you do not want the **Operator** role to have access to a specific set of reports in a report group, you can set the permission '**No Access**' at the Report Group level.

The permission is set only for the report group but not the reports, rules, or subgroups in the Report Group.

Access Control for a Report Group

When you want to change the report group permissions, you must select a report group and set access permissions using the Reports Permissions panel.

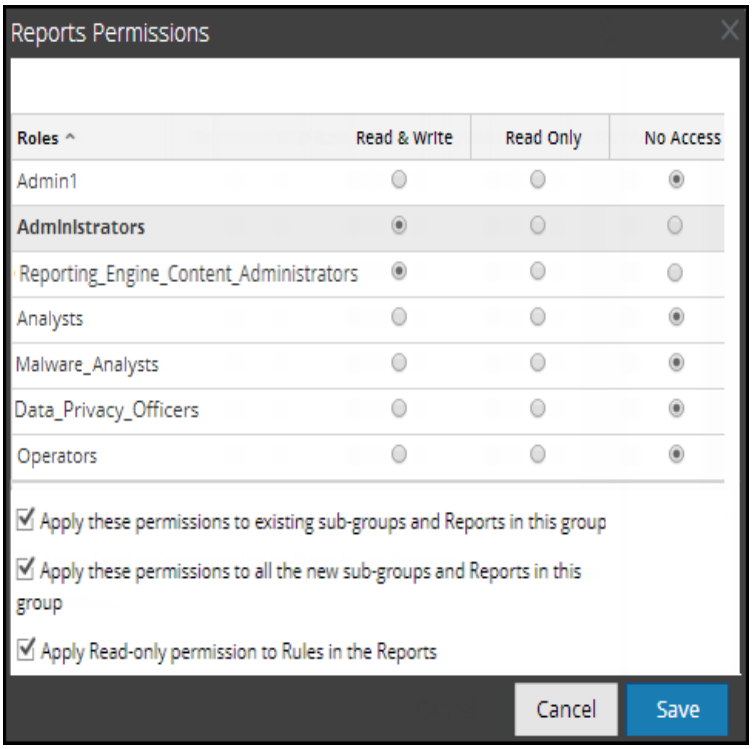
Before applying report group permissions, the default permission set for all the user roles is 'No Access' except for administrators (with full access in Reporting Engine service Config page) and reporting engine content administrators, as shown in the figure.



If you want to change the access permission for a specific user role, you must set these at the report group level, as shown in the figure. <suppose,>Administrators to have access to all the reports in a Report Group, you can set the permission '**Read & Write**' in the Report Group Permissions panel.

You can also apply permissions to existing subgroups and reports in the group, as well as apply read-only permission to rules in the reports by selecting the appropriate checkboxes, as shown in the figure.

You can also apply permissions to all the new subgroups and reports by selecting the appropriate checkbox.



The three scenarios are explained in brief:

- Scenario 1: Permissions applied to Report Group/ Sub Group/ Report based on the user role.
- Scenario 2: Permissions applied to Sub Group and Report in the Group.
- Scenario 3: Read-only permission applied to Rules in the Report.

	Role (Analyst)	Permissions applied to Report Group/ Sub Group/ Report based on the user role	Permissions applied to Sub group and Report in the Group	Permission (Read-only) applied to Rules in the Report
Group	Read & Write	Read & Write	Read & Write	Read & Write
Sub Group	Read	Read	Read & Write - Inherited	Read & Write
Report	Read	Read	Read & Write - Inherited	Read & Write
Rules	Read	Read	Read	Read

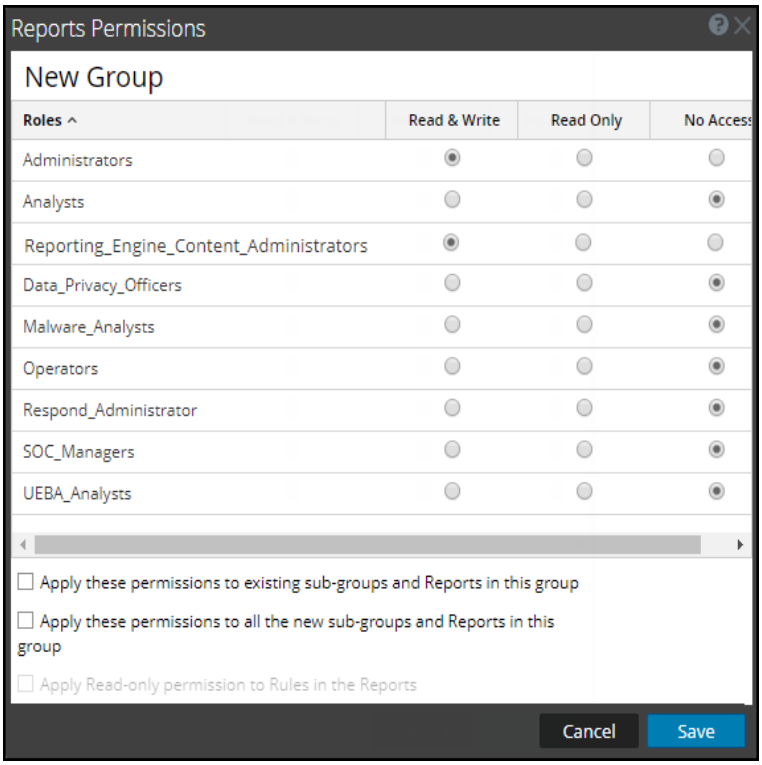
The Report Group will be assigned the role of a **Security Analyst** and permissions are set to **Read & Write** report group.

For scenario 1, each of the levels has a permission set depending on the user role. For scenario 2, the permission at the Report Group level (Read & Write) is inherited by the Sub Group and Reports in the Group. For scenario 3, the Read permission is set for the Rules except that the permission set for the rules cannot be higher than the permissions set for the Report Group.

Access Control for a Report

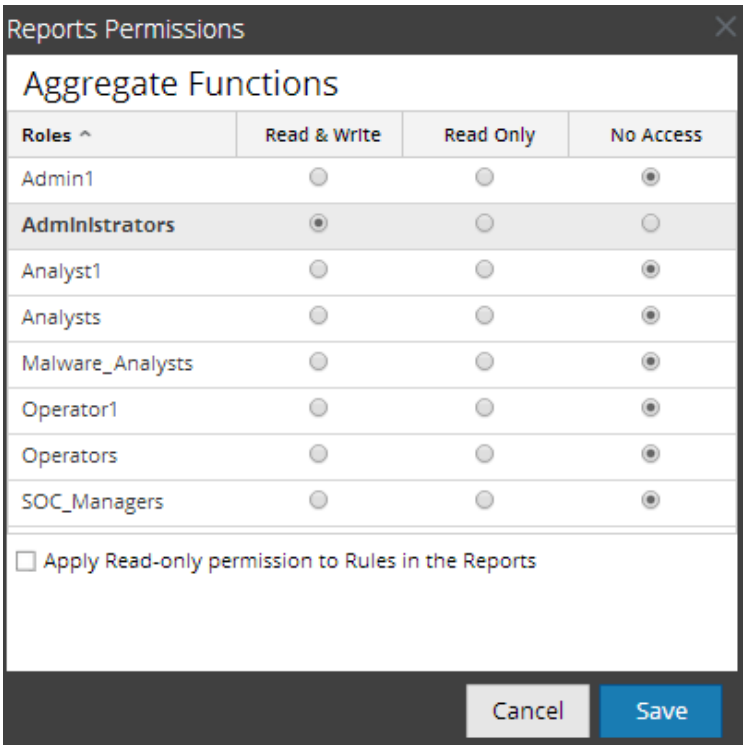
When you want to change the report permissions, you must select a report and set their access permissions using the Report Permissions panel.

Before applying the Report permissions, the default permission set for all the user roles is 'No Access' permission and the checkbox is unchecked, as shown in the figure.



If you want to change the access permission for a specific user role, you must set these at the report level, as shown in the figure. Suppose, you want the **Administrators** to have access to a specific report, you can set the permission '**Read & Write**' in the Report Permissions panel.

You can apply read-only permission to rules in the reports by selecting the checkbox, as shown in the figure.



The two scenarios are explained in brief:

- Scenario 1: Permissions applied to Report Group/ Sub Group/ Report/ Rules.
- Scenario 2: Read-only permission applied to Rules in the Report.

	Role (Analysts)	Permissions applied to Report Group/ Sub Group/ Report/ Rules based on the user role	Permission (Read-only) applied to Rules in the Report
Group	Read & Write	Read & Write	Read & Write
Sub Group	Read	Read	Read & Write
Report	Read	Read	Read & Write
Rules	Read	Read	Read

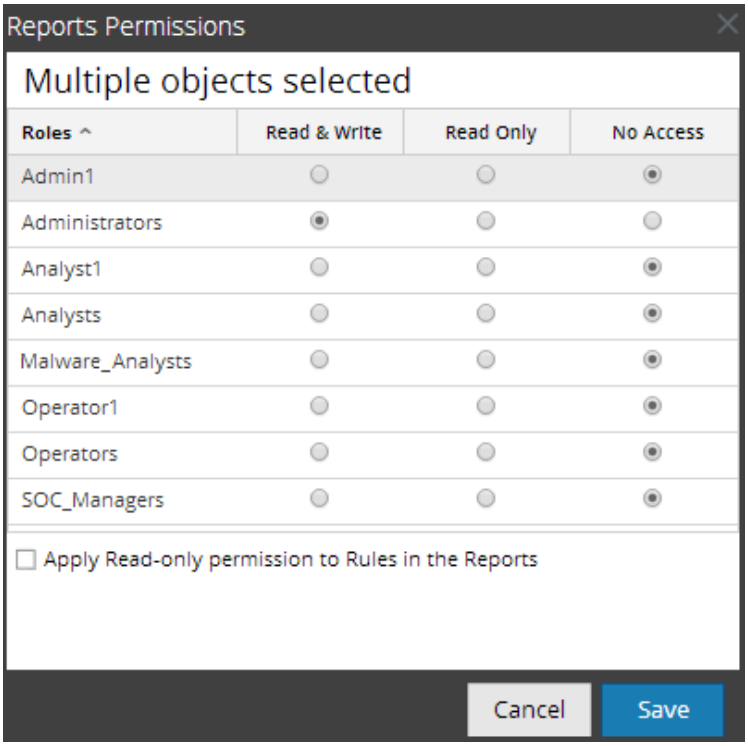
The Report will be assigned the role of a **Security Analyst** and permissions are set to **Read & Write** reports.

For scenario 1, each of the levels has a permission set based on the user role. For scenario 2, the Read permission is set for the Rules except that the permission for the rules cannot be higher than the permission for the Reports.

Note: If the permission for the rules is higher than the permission for the Reports then the permission is be applied only to the reports. For example, if you set the permissions for the Report Group as **No Access** and then specify the option *Apply Read-only permission to Rules in the Reports*, then the read-only permission is not set for the rules.

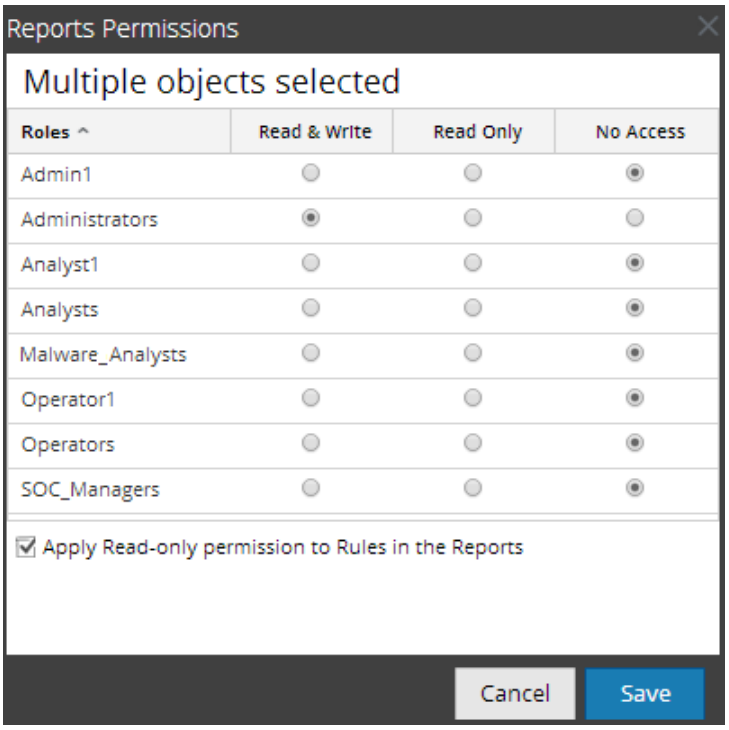
Access Control for a Report When Multiple Reports are Selected

When you want to change permissions of multiple reports, you must select several reports and set their access permissions using the Report Permissions panel. The access permission that you choose is applied to all the selected reports.



Access Control for a Report When Multiple Reports with several rules are Selected

When you want to change permissions when multiple reports with several rules are selected, you must select the checkbox in the Report Permissions panel, as shown in the figure. The read-only access permission is applied to all the rules of the selected reports, provided that the permission of the rules are lower than the permission of the reports.

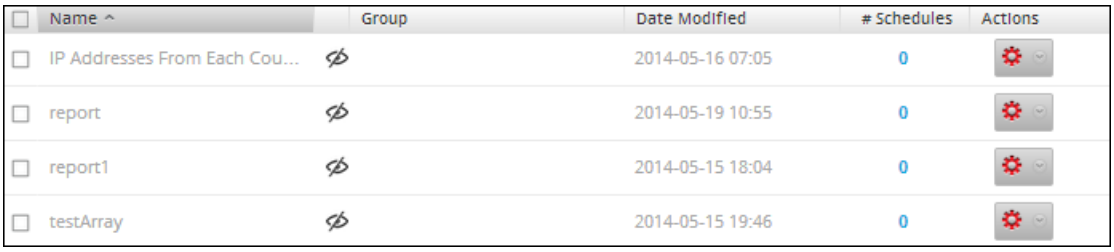


Log in as a Specific User and View the Access Details

When you log in to the NetWitness UI as a user having 'Read access' permission, all the reports is denoted with the symbol (📖) and when you click on the symbol the 'Read Only' callout is displayed on the Reports panel.

When you log in to the NetWitness UI as a user not having 'Read & Write' access permission on a Report, all the reports are denoted with the symbol (⊘) and the reports appear grayed out on the Reports panel.

The following figure shows the Reports panel when logged in with minimal 'Read & Write' access permission.



Note: If a User (other than the super user) creates a report there will be no access to that report for the super user.

Tabular Listing

The following table lists the various columns in the Reports Permissions Panel:

Column	Description
Roles	The role of the user logged into the NetWitness UI.
Read & Write	The user can access, view, edit, import, export, and delete the report on the Reports view. The user can also change the permission on the report.
Read Only	The user can only access and view the report on the Reports view.
No Access	The user cannot access or view the report for which this permission is set.
Apply these permissions to existing subgroups and Reports in this group	Select the checkbox to apply the selected permissions to the existing subgroups in the group and reports in the group. **Note:** This checkbox is populated only when you set access permissions for a Report Group.
Apply these permissions to all new subgroups and Reports in this group	Select the checkbox to apply the selected permissions to the all new subgroups in the group and reports in the group. **Note:** This checkbox is populated only when you set access permissions for a Report Group.
Apply Read-only permission to Rules in the Reports	Select the checkbox to automatically apply permissions to the rules in the reports.

Set Access Control for a Report

Prerequisites

Make sure that you have a minimal 'Read & Write' access permission to set access permissions for a report.

To set access permissions for a report, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

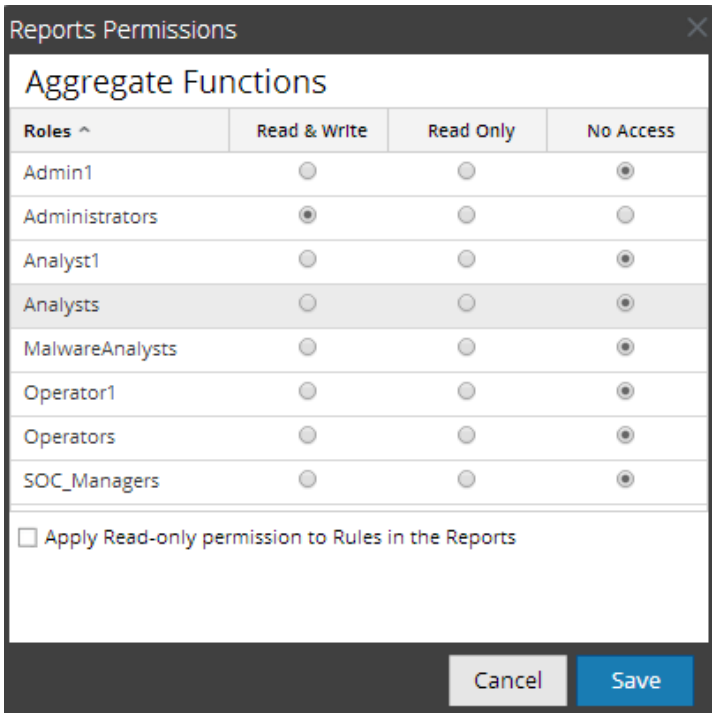
2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** panel, select a report.

4. Click > **Permissions**.

   The Reports Permissions dialog is displayed.



5. Based on the user role, select the appropriate buttons.
6. (Optional) Select the checkbox, if you want to provide read access permission to rules in the reports.

   Note: On selecting the check box, all dependent rules are given READ access permission, provided the permissions for the report is higher compared to the permissions of the rules.

7. Click **Save**.

   A confirmation message that the permission is set for the selected report is displayed.

Set Access Control for a Report Group

Prerequisites

Make sure that you have a minimal 'Read & Write' access permission to set access permissions for a report group.

To set access permissions for a Report Group, perform the following:

1. Go to **Reports**.

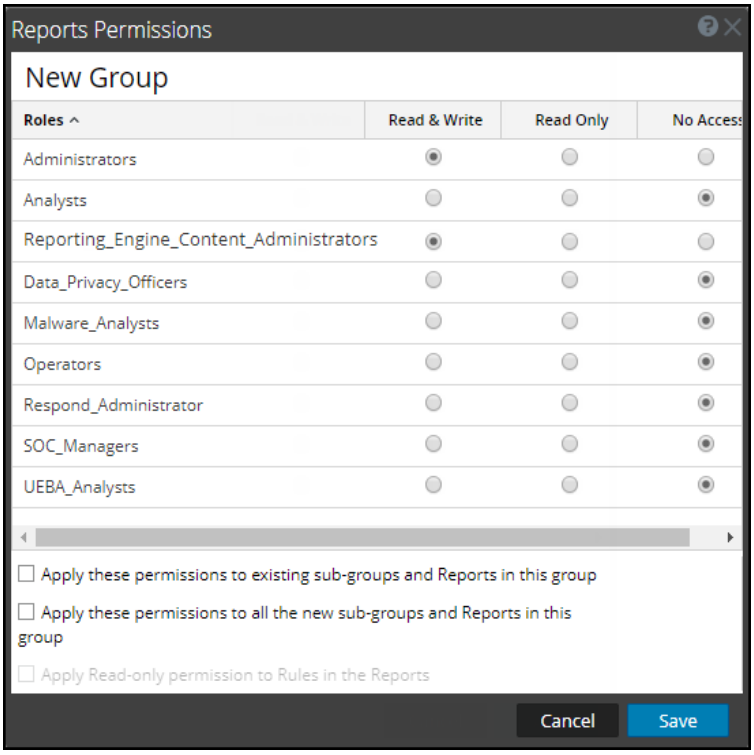
   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports Groups** panel, select or right-click on a report group.
4. Click > **Permissions**.

The Reports Permissions dialog box is displayed.



4. Based on the user role, select the appropriate buttons.
5. (Optional) Select the appropriate checkbox to apply the selected permissions to subgroups and reports in the group.
6. (Optional) Select the appropriate checkbox to apply the selected permissions to all the new subgroups and reports in this group.
7. (Optional) Select the appropriate checkbox to provide read access permission to rules in the reports.

   Note: On selecting the check box, all dependent rules is given READ access permission, provided the permissions for the report is higher compared to the permissions of the rules.

8. Click **Save**.

   A confirmation message that the permission is successfully set for the selected report group is displayed.

Delete a Report or Report Group

To delete reports in a group or subgroup from the Reports panel:

1. Go to **Reports**.

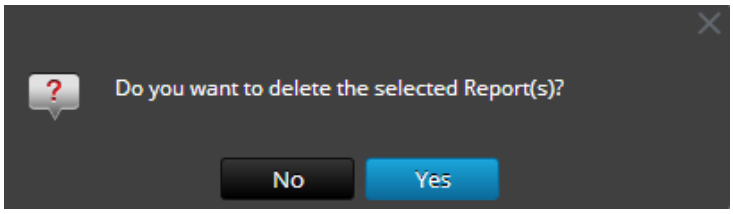
   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** panel, do one of the following:
   - Select the reports and click −.
   - Click ⚙ > **Delete**.

     A confirmation dialog is displayed.



4. Click **Yes** to delete the report.

   A confirmation message that the report is deleted successfully is displayed and the selected report is deleted from the Reports panel.

Delete a Report Group

Prerequisites

Make sure that you have no reports associated with the report group.

To delete report groups in the default folder or subgroups under a report group, perform the following:

1. Go to **Reports**.

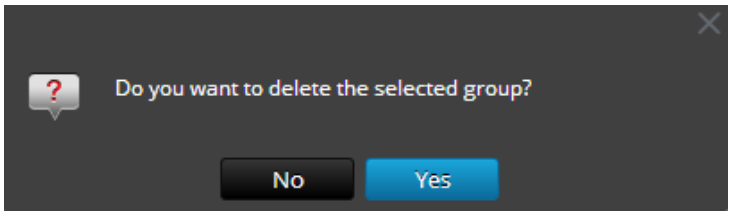
   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports Groups** panel, select the report group and click −.

   A confirmation dialog is displayed.



4. Click **Yes** to delete the group.

   A confirmation message that the group is deleted successfully is displayed and the selected group is deleted from the Report Groups panel.

Duplicate a Report

You can duplicate a report to schedule multiple report for the same report. The duplicated report is displayed in the Reports panel with suffixes. For example, Report (1).

Generally, the duplicate option is used in two scenarios:

- You want to make a copy of the report, to move the same report to another group.
- You want to retain most of the configuration settings for an object but modify few of these settings. For example, when you have a complex query in a rule or several rules in a report, it is very much appropriate to use the duplicate option.

To duplicate an existing report, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** panel, select a report that you want to duplicate and click .

   The report is saved successfully and added to the Report list.

You can move the duplicated report to another group.

Edit a Report

To edit reports in a group or subgroup from the Reports panel, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. Select a report and in Actions column, click  > **Edit**.
   The Build Report view tab is displayed.



4. Modify the text and add more rules to the report (if required).
5. Click **Save**.

   A confirmation message that the report is saved successfully is displayed.

Refresh a Report Group or Report List

You can refresh a report group or reports to view the re-arrangement of groups or reports.

To refresh a report group or reports, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. Do the following to move the group or reports to a new location:

   - In the **Reports Groups** panel, drag and drop the group.
   - In the **Reports** panel, drag and drop the reports to the desired group in the Report Groups panel.
     The report group or reports are moved to the new location.

4. Do the following to refresh a group or report list:

   - In the **Reports Groups** panel, click .

     The report group gets refreshed.

   - In the **Reports** panel, click  .

     The Report list gets refreshed.

Edit a Scheduled Report

To edit a scheduled report from the Scheduled Reports List panel, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

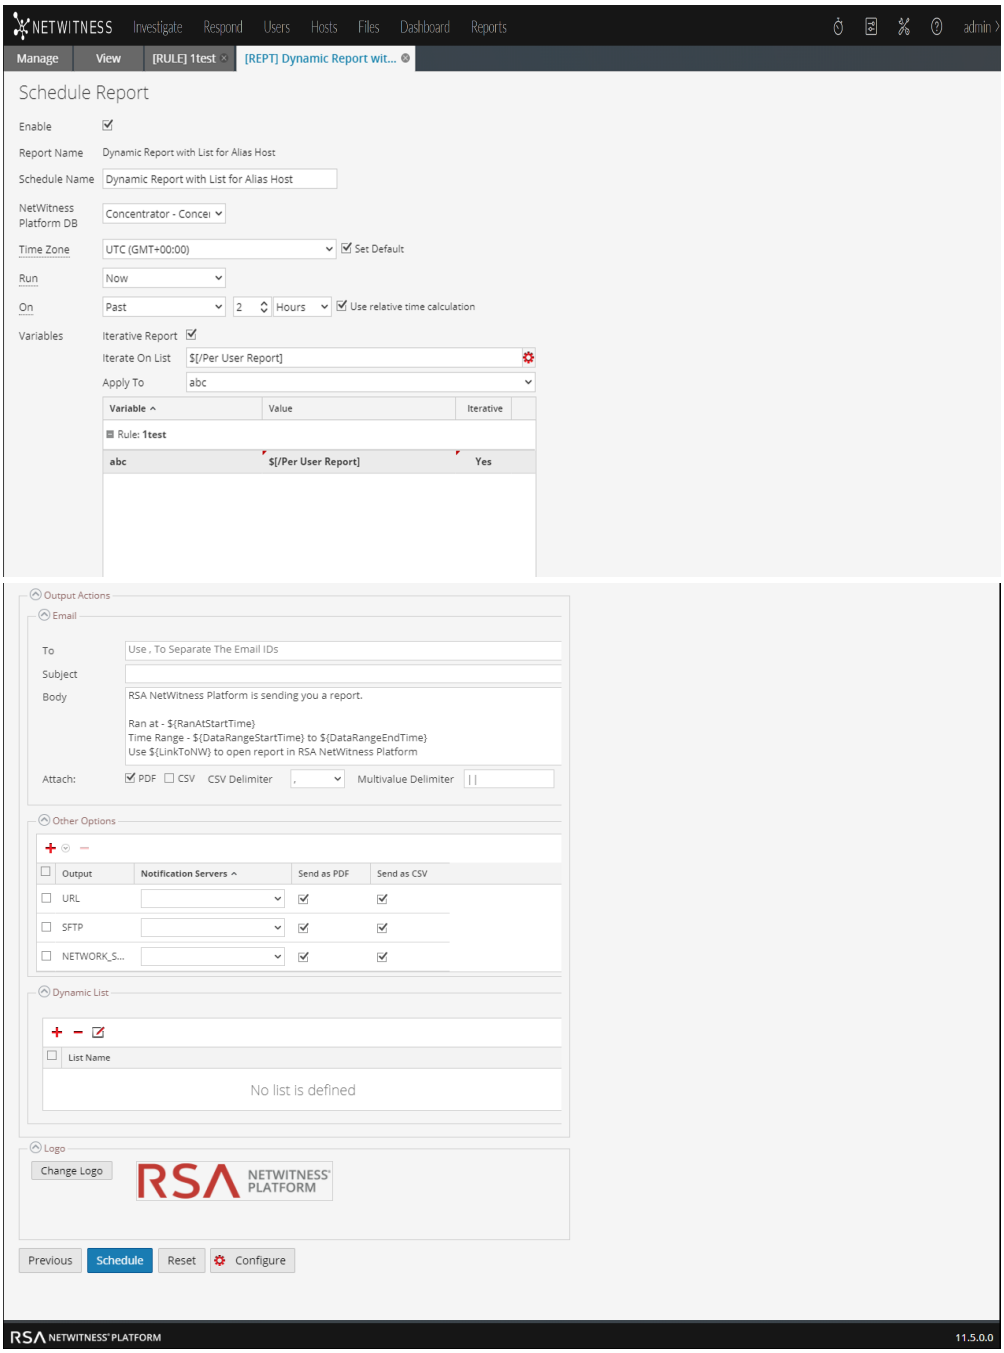
   The Report view is displayed.

3. In the **Reports** panel, select a report and in Actions column, click  > **View Scheduled Reports**.

   The Report Schedule tab is displayed.

4. In the **Report Schedule** panel, do one of the following:

   - Select a report and click .
   - Select a report and click  > **Edit Schedule**.
     The Schedule Report tab is displayed.



5. In the Schedule Report tab, do the following:

   a. In the **Schedule Name** field, modify the name for the schedule report configuration.

   b. To execute the reports as per the schedule, select the **Enable** checkbox.

   c. From the **Data Source** field, select the datasource.

> **Note:** If the data source is not listed, ensure you have **Read** permissions set for the data source. This is applicable for NWDB and Warehouse data source. For more information, see "Configure Data Source Permissions" topic in the *Reporting Engine Configuration Guide*.

6. (Optional) From the **Warehouse Resource Pool** drop-down, select the pool or queue for the report.

   Note: The **Warehouse Resource Pool** drop- down is displayed only if the Warehouse Rule is selected. If no pools or queues are entered for the Reporting Engine, this field is disabled.

7. From the **Run** field, select the type of run schedule. (For example, Now or Hourly).
8. Select the date range to run the query based on absolute duration or select the **Use relative time duration** checkbox to run the query based on relative duration.
9. (Optional) In the Output Actions panel, do the following:
   a. Type the email address and subject.
   b. Edit the body of the message for the report.
   c. Select the format of the attachment.
   d. Type a value for the CSV and Multivalue delimiters.
10. (Optional) In the Other Options field, do the following:
    a. Click  > **SFTP** or **URL** or **Network Share**. Based on the selected option, a row gets added in the Other options field.
    b. Select the appropriate options to send the report in PDF or CSV format to the configured SFTP, URL or Network Share.
11. (Optional) To add a list in the Dynamic List panel, see Generate a List from the Scheduled Report section in Create and Schedule a Report.
12. (Optional) To choose another logo in the Logo panel, see Manage and Select a Report Logo section.

    Note: If you do not specify a logo, the default RSA logo is used.

13. Click **Schedule**.

    The scheduled report executes as scheduled and provides the configured outputs.

Delete a Scheduled Report

To delete a scheduled report from the Scheduled Reports List panel, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

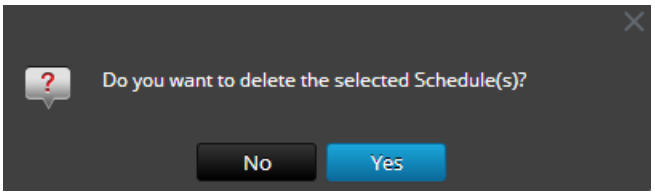
   The Report view is displayed.

3. In the **Reports** toolbar, click **View All Schedules**.

   View Report Schedule is displayed.

4. In the **Report Schedule** panel, select the report.
5. Click  >**Delete Schedule**.

   A confirmation dialog is displayed.



6. Click **Yes** to delete the scheduled report.

   A confirmation message that the scheduled report is deleted successfully is displayed and the selected schedule is deleted from the Scheduled Reports List panel.

Export a Report

You can export the selected reports to an external file that can be later imported to another NetWitness environment.

Prerequisites

Make sure that you have reports in the report group.

To export selected reports in the Report Groups panel to an external file, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

   In the **Reports** panel, do one of the following:

   - Select a report and click  > **Export**.
   - Click  > **Export**.

3. You can export multiple reports at a time. To select multiple reports, check the checkbox of the report to be exported. The exported file is saved to the local drive in an archived format.

Open CSV files with Unicode characters in MS Excel

To open downloaded CSV files containing Unicode characters in MS Excel, follow these steps:

1. Download and save the CSV file.
2. Open Microsoft Excel and navigate to the **Data** tab.
3. Click on **From Text** menu item; find the CSV file that you downloaded and click **Import**.
   The Text Import Wizard is displayed.
4. Select **Delimited** or **Fixed Width** data type from the **Original data type** radio button.
5. Click **File origin** drop down list and select **65001: Unicode (UTF-8)** and click **Next**.
6. Select the delimiter that was used in the file that you imported and click **Next**.
7. Select the data format for each column of data that you want to import and click **Finish**.
   The correct output is displayed in an MS Excel sheet.

Export a Report Group

You can export a selected report groups to an external file that can be later imported to another NetWitness environment.

Prerequisites

Make sure that you have reports in the Report Group.

To export selected report groups in the Report Groups panel to an external file, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports Groups** panel, select a report group and click  and select one of the following:

   - **Export** - This selection exports a report in a .zip file.
   - **Export as Text** - This selection exports all the content from the Reporting Engine in a .zip file which contains the data in text format.

   You can export multiple report groups at a time. To select multiple report groups, press and hold the CTRL button and select the report groups to be exported. The exported file is saved to the local drive.

Import a Report or Report Group

You can import a group containing subgroups and reports from other instances of NetWitness into Report Groups panel. Reports must be in a valid binary file that was exported from another NetWitness instance.
During the import process, you select the binary file and specify whether existing reports with the same name must be overwritten or not by the reports contained in the binary import file.

- If you choose to overwrite, all duplicate rules, lists and reports are overwritten by the contents of the binary import file.
- If you choose not to overwrite, and a duplicate rule, list or report exists in the target folder, the import fails and display a message about duplicate reports.

You cannot import reports to a specific report group. The imported files are stored in the **All**root folder.

Prerequisites

Make sure that you have the reports or report groups exported from other instances of NetWitness.

To import groups containing subgroups and reports from other instances of NetWitness into Report Groups panel, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

The Report view is displayed.

3. In the **Reports Groups** panel, select a folder to import the file.
4. Do one of the following:
   - In the **Reports Groups** toolbar, click > **Import** to import a group.
   - In the **Reports** toolbar, click > **Import** to import a report.
     The Import Report dialog is displayed. You can import multiple reports and report groups at a time. To select multiple reports or report groups, press and hold the CTRL button and select the reports or report groups to be imported.
5. Click **Browse** to select the binary file.

   NetWitnessprovides a file system view of the files.
6. Locate the binary file and click **Open**.

   The file gets added to the Import Report list.
7. (Optional) To overwrite any existing rule in the library with an identically named rule in the binary file when importing, check the **Rule** checkbox. If you do not select the Overwrite option, and an identical rule is encountered in the binary file, the binary file is imported and no error message is displayed.
8. (Optional) To overwrite any existing list in the library with an identically named list in the binary file, check the **List** checkbox. If you do not select the Overwrite option, and an identical list is encountered in the binary file, the binary file is imported and no error message is displayed.
9. (Optional) To overwrite any existing report in the library with an identically named report in the binary file when importing, check the **Report** checkbox. If you do not select the Overwrite option, and an identical report is encountered in the binary file, the binary file is imported and no error message is displayed.
10. Click **Import** to import the binary file.

Enable or Disable a Scheduled Report

To enable or disable a scheduled report from the Scheduled Reports List panel, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.
2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, select a report and click > **View Scheduled Reports**.
   View Report Schedule is displayed.
4. Select a report from the Report Schedule panel.
5. In the Actions column, click > **Enable**.
   The state of the report is changed to ‘Running’, if the report is scheduled to run immediately.

6. In the Actions column, click  > **Disable**.
   The state of the report is changed to 'Inactive'.

Start or Stop a Scheduled Report

To start or stop a scheduled report, perform the following:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, select a report and click  > **View Scheduled Reports**.
   The Report Schedule view is displayed.
4. Select a report from the Scheduled Reports List panel.
5. In the Actions column, click  > **Start**.
   The state of the report is changed to 'Running', if the report is scheduled to run immediately.
6. In the Actions column, click  > **Stop**.
   The state of the report is changed to 'Completed'.

Note: If you have multiple executions running for a scheduled report and the state of the report is in Queued state, the STOP option is disabled till all the previous executions are completed and last execution is in running state. To stop the individual execution of a particular schedule, refer to Stop an Individual Execution of a Scheduled Report.

View an Execution History of a Scheduled Report

You can view the execution history of a scheduled report. You can view the history of a scheduled report that is run. You can view the history based on the following criteria:

- Number of past schedules executed
- Start date and end date for the date range

You can view the details such as how many times the scheduled report was executed, the time of execution (seconds), execution state. You can also view the report generated on a full screen.

To view the execution history of a scheduled report, perform the following:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Report view is displayed.

3. In the **Reports** panel, do one of the following:
   - Click  > **View Scheduled Reports**.
   - Click the **#Schedules** column.
     The Schedule Reports view tab is displayed with the status of each of the scheduled report.
4. Do one of the following:
   - Select a scheduled report and click  > **Execution History**.
   - Select a scheduled report and click  .
     The Execution History view is displayed.

   Note: By default, you can view 10 number of execution history of a scheduled report. The execution history shown depends on the Retain Report History Configuration set on the **General**tab of the  **(Admin) > Services > Reporting Engine Config** view.
   For example, if you set the Retain Report History Configuration to 100 days, the data displayed on the Execution History view. is the past 100 days execution history details considering the current date information.
5. From the **Get history by:** field, select the type of history to be fetched. (For example, Past or Range (Specific))
6. In the **Count** field, enter the number of executions to be displayed.
7. Click **Show History**.
   The execution history of the scheduled report is displayed.

Stop an Individual Execution of a Scheduled Report

If you have multiple executions running for a scheduled report, you can stop the individual execution of the scheduled report in **Execution History** panel.

To stop the individual execution of a scheduled report, perform the following:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, do one of the following:
   - Click  > **View Scheduled Reports**.
   - Click the **#Schedules** column.
     The Schedule Reports view tab is displayed with the status of each of the scheduled report.
4. Do one of the following:
   - Select a scheduled report and click  > **Execution History**.
   - Select a scheduled report and click  .
     The Execution History panel is displayed.

5. Select a schedule, in the **Actions** column, click .
   The state of the execution schedule is changed to 'Cancelled'.

Manage and Select a Report Logo

Prerequisites

Make sure that you have the Reporting Engine service defined prior to managing a logo.

Manage Report Logos

To manage logos, perform the following:

1. Go to **(Admin) > Services**.
   The Services view is displayed.
2. In the **Services** panel, select an Reporting Engine service and click **View** > **Config**.
   The services config view is displayed.
3. Select the **Manage Logos** tab.
   All the available logos are displayed.

Add a Logo

To add a logo, perform the following:

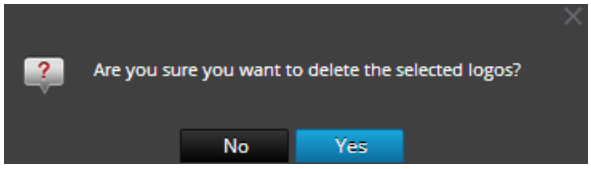
1. In the **Manage Logos** tab, click .
   A file browser opens where you can choose the file from the local drive.
2. Select the logo and click **Open**.
   The selected logo gets added to the Manage Logos section.

Delete a Logo

To delete a logo, perform the following:

1. In the **Manage Logos** tab, do one of the following:
   - Select the logo and click .
   - Perform (Ctrl+click) to select multiple logos and click .

   A confirmation dialog is displayed.



2. If you want to delete the logo, click **Yes**.
   The selected logo is deleted from the Manage Logos section.

Set Default Logo

To set a default logo, perform the following:

In the **Manage Logos** tab, select a logo and click Set default.
The chosen logo is set as the default logo for the RE service.

Select a Logo

To select a logo, perform the following:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** panel, select a report.

4. Click > **View Scheduled Reports**.
   The View scheduled reports view tab is displayed.

5. Select a scheduled report and in Actions column, click > **Edit Schedule**.

   The Schedule Report view tab is displayed.

6. In the Logo panel, click **Change Logo**.

   The Change a Logo dialog box is displayed.

7. Do one of the following:

   - Click **Upload new logo** to upload another logo.
   - Select a logo from the list.

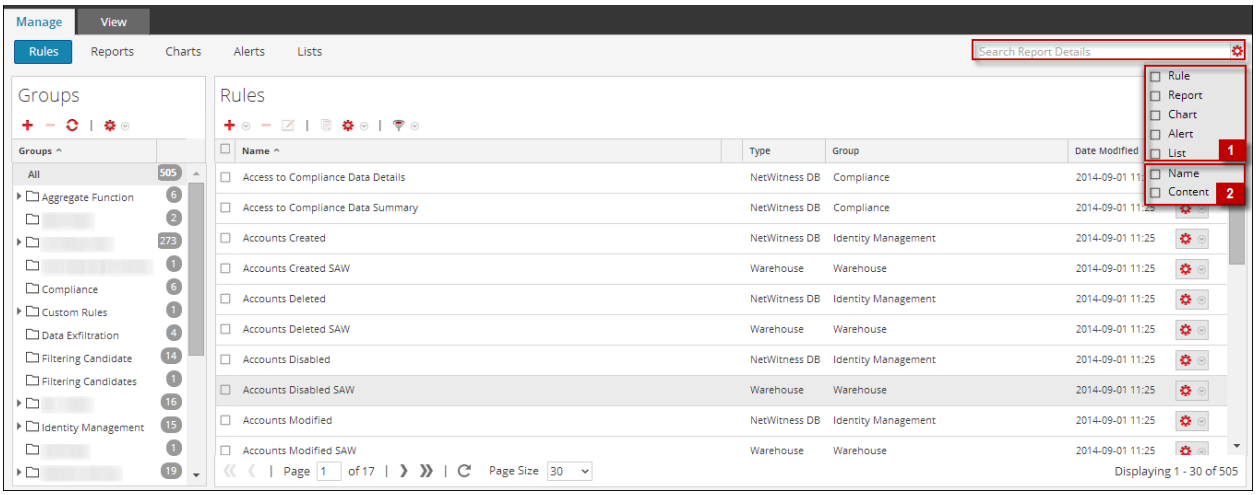
8. Click **Select**.
   The selected logo is available on the Logo panel.

Search Reporting Details

This section provides instructions on how to perform a keyword search on name and content for each of the Reporting components. You can perform a keyword search on name and content for each of the Reporting components (Rule/Report/Chart/Alert/List) on the Reporting UI.

Note: You cannot search based on date and numeric values.

The following figure shows the search parameters available in the Reporting Module:



The following are the search parameters available on the Reporting UI:

1. Search for entities (rule, report, chart, alert, list).
2. Search for the entities based on either the name or content.

Note: Searches are case insensitive. For example, Completed is equivalent to completed.

Prerequisites

In the Reporting Module, you can perform a keyword search based on the name and content (definition). In this context, content implies definition of each of the reporting components. For instance, the value defined in the rule, report, report schedule, chart, and alert panel. You can also prioritize your search by selecting either or all of the components: Rule, Report, Chart, Alert, or List.

Note: You cannot search based on the List values and list path stored in schedule definition panel.

For example, to search for the rule name (ExpertRule), you must select **Rule, Name, and Content** in the **filtering options** drop-down to view all the rule names that matched the search. You can similarly search for a report, chart, alert, or list definition.

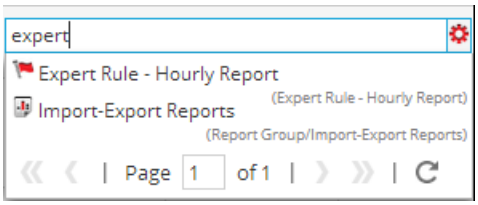
To search for reporting details from the Manage tab, perform the following:

1. Go to **Reports**.

   The **Manage** tab is displayed.

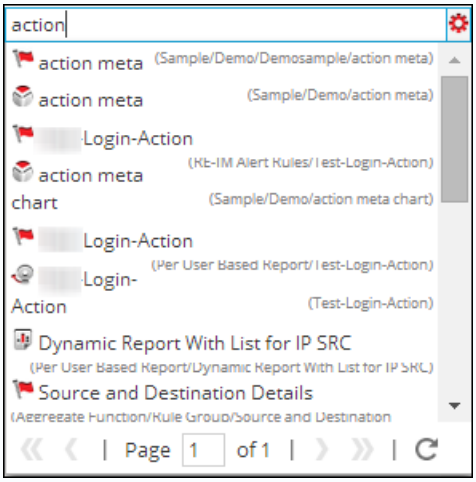
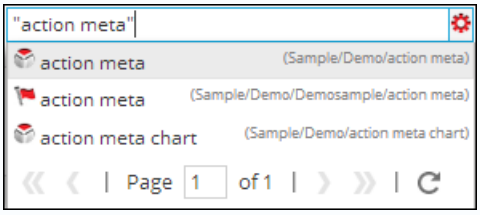
2. Click  and select the appropriate criteria to search.
3. In the **Search** field, enter the text to be searched.

   The search drop-down list is displayed:

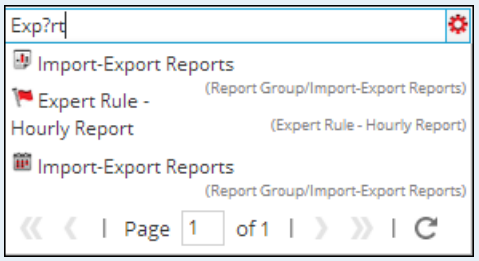
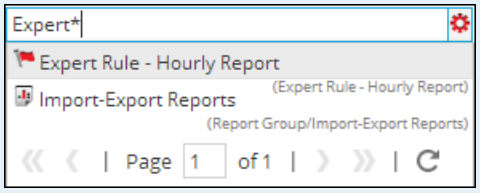
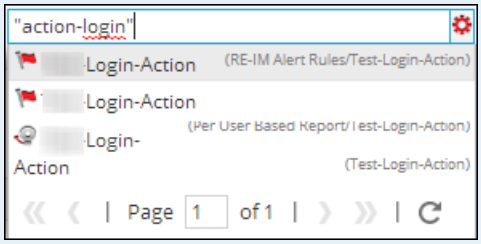


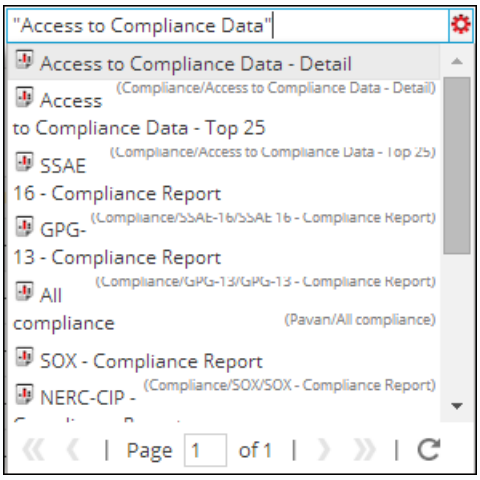
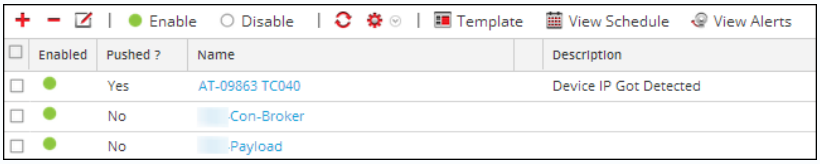
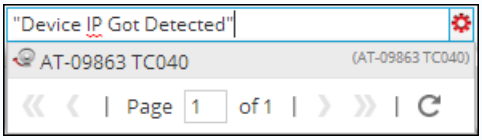
Search Syntax and Different Types of Search

The following table explains the search syntax and the possible searches that can be performed on the Reporting user interface.

Search Types	Description
Word or phrase based search	**Word Based Search:** To search for a word such as "action" or "meta", you must enter the word in the search box. The following figure shows the search results for the text **action**.  **Phrase based search:** A Phrase is a group of words surrounded by double quotes such as "action meta". To search for a phrase, you must enclose phrases in double-quotes in the search box. The following figure shows the search results for the phrase "action meta". 

Search Types	Description

Wildcard Search (Single/ Multiple/ Special Character Search)
 	
The question mark "**?**" symbol is used to perform a single character wild card search and asterisk "*" symbol is used to perform multiple character wildcard search.
	Single character search:
	The single character wildcard search looks for terms that match with the single character replaced. For example, to search for "Expert" or "Export" you can use the search syntax:
	`Exp?rt`
	The following figure shows the search results for the wildcard character **Exp?rt**.
	
	Multiple character search:
	Multiple character wildcard search looks for 0 or more characters. For example, to search for Expert, or Experts, you can use the search syntax:
	`Expert*`
	The following figure shows the search results for the wildcard multiple character **Expert***.
	
	Special character search:
	Certain punctuation and special characters are ignored during search (@#$%^&*(){}"~=+-[]\?|!:,.). For example, a search for action-login will be interpreted during search as "action" "login", that is, if rules exist with name "action-login" and "action@login" and search string is "action-login", the search result will return both the rules.
	

Search Types	Description
Search based on name or content	**Search based on name:** When you want to search based on the name of a report, select **Report**and **Name**box from the filtering options drop-down. For example, to search for the report name "Report With Multiple Rules", you can use the search syntax: "Access to Compliance Data" **Note:** When you search for a report, it implies you can search for the report schedules as well. The search result will return the report containing the specific name.  **Search based on content:** When you want to search for the content within an alert, say alert description, select **Alert** and **Content** box from the filtering options drop-down. For example, to search for the alert description "Device IP Got Detected", you can use the search syntax: "Device IP Got Detected"  The search will return the result having the specific content. 

Working with Charts

The Reporting module user interface provides access to NetWitness charts. The following topics discuss charts:

- Configure and Generate a Chart
- Configure a Chart
- Schedule a Chart
- View a Chart
- Test a Chart
- Investigate a Chart
- Manage a Chart Group and Chart

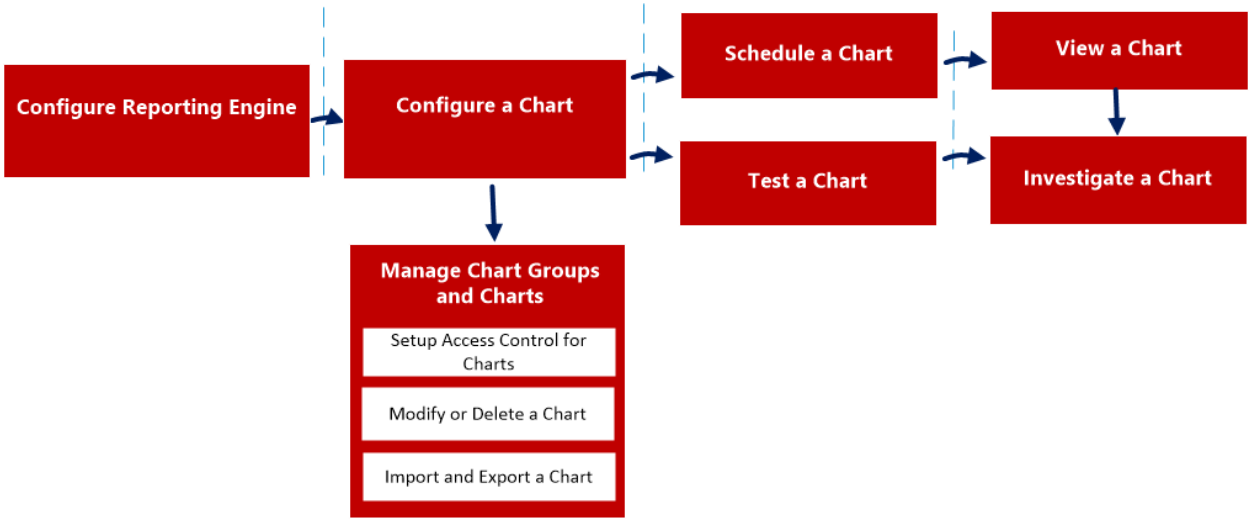
Configure and Generate a Chart

Chart is a graphical visualization of data. You can view different kinds of charts, including multiple types of plot, line, bar, and area charts.

Any NWDB rule in the Reporting Engine system which is not sorted by none can be used to instantly create a chart. For more information on "How to create an NWDB rule", see Configure a Rule.

The chart interval can be adjusted from the chart definition panel itself. Every time a chart is executed, it stores its result data locally in the Reporting Engine, so that it can be reviewed in either the Dashboard View or Chart View without any performance considerations.

The following is an overview of the entire process of configuring and generating a chart.



To configure and generate a chart, perform the following:

1. Configure Reporting Engine
2. Configure an NWDB rule
3. Configure a Chart

4. Schedule a Chart
5. View a Chart
6. Test a Chart
7. Investigate a Chart
8. Manage a Chart Group and Chart

Configure Reporting Engine

You must configure the Reporting Engine before you can configure and generate a chart. You must also specify the data source in the Reporting Engine from which the data is extracted. For more information on how to configure a Reporting Engine, see "Configure Reporting Engine" topic in *Reporting Engine Configuration Guide*.

Configure an NWDB Rule

The NetWitness rule which is not sorted by none is used to create a chart. The NetWitness database extracts the meta from the Reporting Engine and provides the meta for rules. These rules are an essential building block in managing a chart.

Note: If the rule contains the lookup_and_add, sum_count, or sum_values rule actions, the associated chart will not contain data. Do not use lookup_and_add function in a rule to create a chart or a report in a chart format as the output of the lookup_and_add function in a rule will not be displayed on the chart and may result in graphical representation of incorrect information. Also, the PDF for that chart created may show an empty chart or incorrect representation of data.

Configure a Chart

You can configure a chart using the NWDB rules.

Schedule a Chart

After a chart is defined with the required components, you can configure its execution properties by scheduling a chart. Here, you can quickly view, add, and edit the schedule details for a chart.

View a Chart

You can view the scheduled charts in the Chart View.

Test a Chart

You can run the test on a chart and view all the chart details based on the selected time range.

Access Control for a Chart

The Reporting Module provides access control at the chart level. Only a user who has the right set of permissions can perform the tasks in Reporting module. The access control is managed by the administrator from the  **(Admin) > Security > Roles** tab.

When you create users and user roles, ensure that the roles that you create for specific tasks have access to all the necessary permissions. This could require permissions at several levels of the role hierarchy.

Charts can be tied to a specific set of user roles so that when a user logs in NetWitness, the charts with the access rights for the specific user role can be viewed. Users that belong to a user role with the 'Read & Write' access permission can define charts. Further, the access can be tightened so that charts are accessed only by those who have the 'Read Only' access.

At the chart level, you can set the following access permissions for the user roles in NetWitness:

- Read & Write
- Read Only
- No Access

To change the access permission for a specific user role, you must set the permission at the chart level. For example, for **Administrators** to have access to a specific chart, you could set the permission 'Read & Write' in the Charts Permissions dialog.

You can apply read-only permission to rules in the charts by selecting the checkbox.

Two scenarios that describe how to set access control are explained here:

- Scenario 1: Permissions applied to Chart Group/ Subgroup/ Chart/ Rules based on the user role.
- Scenario 2: Read-only permission applied to Rules in the Chart.

	Role (Analyst)	Permissions applied to chart group, subgroup, chart or rules based on the user role	Permissions (Read-only) applied to rules in the chart
Group	Read & Write	Read & Write	Read & Write
Subgroup	Read	Read	Read & Write
Chart	Read	Read	Read & Write
Rules	Read	Read	Read

The chart is assigned the role of a **Security Analyst** and permissions are set to 'Read & Write' charts.

For scenario 1, each of the levels has a permission set based on the user role. For scenario 2, the Read permission is set for the rules except that the permission for the rules cannot be higher than the permission for the charts.

Note: If the permission for the rules is higher than the permission for the chart, the permission is not applied. For example, if you set the permissions for the Report Group as **No Access** and specify the option *Apply Read-only permission to Rules in the Reports*, the read-only permission is not set for the rules.

Access Control for a Chart When Multiple Charts are Selected

To change permissions for multiple charts, you must select several charts and set their access permissions using the Charts Permissions panel. The access permission that you choose is applied to all the selected charts.

Access Control for a Chart When Multiple Charts with Several Rules are Selected

To change access permissions for a specific user role when multiple charts with several rules are selected, select the checkbox in the Charts Permissions panel.

The read-only access permission is applied to all the rules of the selected charts, provided that the permission of the rules are lower than the permission of the charts.

Note: If a user (other than the super user) creates a chart, the super user cannot access that chart.

Access Control for a Chart Group

To change chart group permissions, select a chart group and set the access permissions using the Charts Permissions panel. Before chart group permissions are applied, the default permission set for all the user roles is 'No Access'.

To change the access permission for a specific user role, set the permission at the chart group level. For example, for administrators to have access to all the charts in a Chart Group, set the permission 'Read & Write' in the Charts Group Permissions panel.

You can also apply permissions to subgroups and charts in the group, and apply read-only permission to rules in the charts by selecting the appropriate checkboxes.

Three scenarios that describe how to set access control are explained here:

- Scenario 1: Permissions applied to chart groups, subgroups, or charts based on user roles.
- Scenario 2: Permissions applied to subgroups and charts in the group.
- Scenario 3: Read-only permission applied to rules in the chart.

	Role (Analyst)	Permissions applied to chart groups, subgroups, or charts based on user roles	Permissions applied to subgroups and charts in the group	Permissions (Read-only) applied to rules in the chart
Group	Read & Write	Read & Write	Read & Write	Read & Write
Subgroup	Read	Read	Read & Write - Inherited	Read & Write
Chart	Read	Read	Read & Write - Inherited	Read & Write
Rules	Read	Read	Read	Read

The chart group is assigned the role of a **Security Analyst** and permissions are set to 'Read & Write'.

For scenario 1, each of the levels will have a permission set depending on the user role.

For scenario 2, the permission at the chart group level will be inherited by the subgroup and by charts in the group.

For scenario 3, the Read permission is set for the rules. However, the permission set for the rules cannot be higher than the permissions set for the chart group.

The following table lists the columns in the Charts Permissions panel:

Column	Description
Roles	The role of the user logged into the NetWitness UI.
Read & Write	The user can access, view, edit, import, export, and delete the chart in the Charts view. The user can also change the permission for the chart.
Read Only	The user can only access and view charts on the Charts view.
No Access	The user cannot access or view charts for which this permission is set.
☐ Apply these permissions to existing sub-groups and Charts in this group	Select the checkbox to apply the selected permissions to the existing subgroups in the group and charts in the group. **Note:** This checkbox is populated only when you set access permissions for a Chart Group.
☐ Apply these permissions to all new sub-groups and Charts in this group	Select the checkbox to apply the selected permissions to the new subgroups in the group and charts in the group. **Note:** This checkbox is populated only when you set access permissions for a Chart Group.
☐ Apply Read-only permission to Rules in the Charts	Select the checkbox to automatically apply permissions to the rules in the charts.

Configure a Chart

After a chart is defined with the NetWitness rules with NWDB as the data source, you can configure its execution properties.

Create a Chart Group

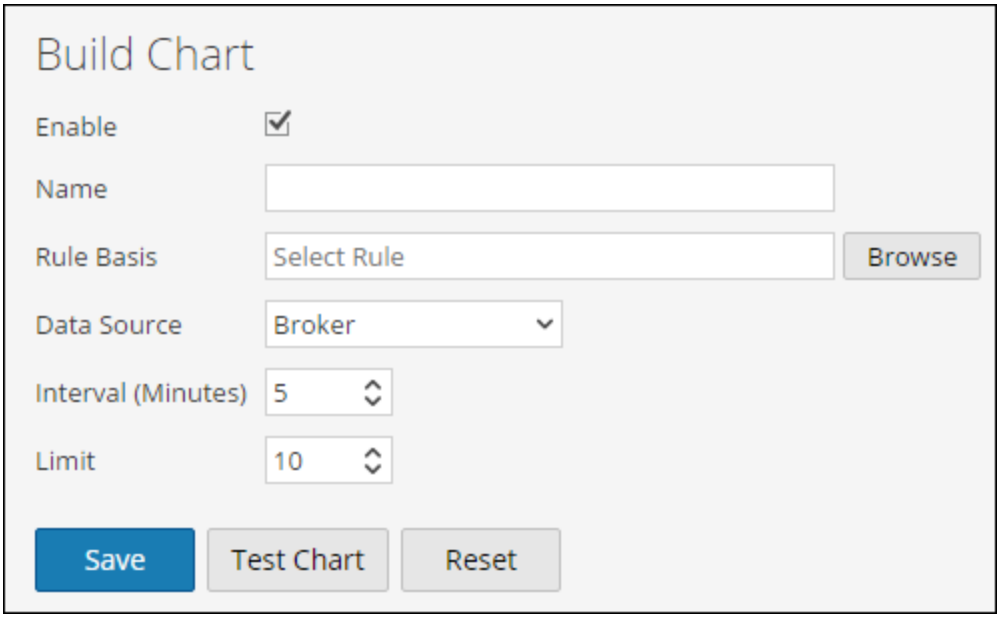
To add groups to the default folder or to add subgroups under a chart group:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Charts**.
   The Chart view is displayed.
3. In the **Charts Groups** panel, click .
   A default group is added in the Chart Groups panel.
4. Enter the name of the new group.
5. Press **Enter**.
   The group is added to the Chart Groups panel.

Create a Chart

To add charts to a group or subgroup:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Charts** to display the Chart view.
3. In the **Charts** toolbar, click .
   The Build Chart tab is displayed.



4. Enter the name of the chart.
5. For the Reporting Engine to collect the data and generate chart results, select the **Enable** checkbox.
6. In the Rule Basis field, do the following:
   a. Click **Browse**. The Add Rule dialog box is displayed.
   b. Navigate the Rule tree and select a rule.
   c. Click **Select**.
7. The Rule appears in the **Rule Basis** field.
8. Select the data source from the **Data Source** drop-down list.

   Note: If the default data source is configured in the Reporting Engine, then the data source is displayed by default on the Build Chart page. If the data source is not displayed, ensure you have Read permissions set for the data source. This is applicable for NWDB and Warehouse data sources. For more information, see the "Configure Data Source Permissions" topic in the *Host and Services Configuration Guide*.

9. (Optional) To modify the Interval value, click the up or down arrow.
   The Interval value is the interval in minutes at which the rule which forms the basis of the chart is

run to collect data.

10. Select the **Limit** value to limit the number of records to be displayed.
11. **X-Axis** and **Y-Axis** are used to specify the meta to be plotted in charts.
    In the **X-Axis**, the meta for the 'Group by' rule is displayed. In the **Y-Axis**, the aggregate functions used in the rule are displayed.

    Note: Sum, Count, Countdistinct and Average are the supported aggregate functions for chart. By default, for Custom Rules with multiple 'Group by', you can select only the first meta in **X-Axis**.
12. Click **Save**.
    A confirmation message that the chart is saved successfully is displayed.

Schedule a Chart

You must schedule a chart to further investigate on the chart details.

By enabling a chart, the chart executes as scheduled and provides the configured output with the state of the chart changed to ‘Scheduled’.

To schedule a chart:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Charts**.
   The Chart view is displayed.
3. In the **Charts** panel, select a chart or several charts that display ○ in the **Enabled** column.
4. Click .
   A confirmation message indicates that the chart(s) state is changed successfully.

View a Chart

After you view a chart, you can perform the tasks:

- You can print, save, email and view charts on full screen.
- You can also select a date from the calendar to view a list of successfully run charts for the chosen date.

To view a chart:

1. Go to **Reports**.

   The Manage tab is displayed.
2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** panel, Select a chart and click  > **View**.

   The View Chart view tab is displayed.

4. In **Chart Options**, do the following:

   a. Select the **Time Range**.

   > **Note:** When you select the Time Range option, you can select a pre-set time range such as last hour, last 3 hours and the Last N Days...or you can customize the selection by choosing Last N Days or Custom. If you select Last N Days option, you can view the historical data for a maximum of 15 days. If you select the Custom option, you can select a start date and end date to view the data for the selected date range.

   b. Select the **Series**, either **Chart Values over Time** or **Chart with Totals**.
   When you select **Chart Values over Time**, the chart displays the change in values for the selected time. When you select **Chart with Totals**, the chart displays a total for each aggregate value for the selected time.

   c. Select **Items To Plot** to define the number of events to view on the chart.

   d. From the **Chart Type** drop-down list, select the chart type.

   e. Click **Reload** to reload the selected chart.
   If there is a delay in retrieving the historical data for the selected time range, a message is displayed.

After the chart is generated, a notification is displayed in the notification tray available in the NetWitness toolbar. For more information on the NetWitness toolbar, see the "Browser Window" topic in the *NetWitness Getting Started Guide*.

View all Charts List

To view a list of all the charts:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** toolbar, click **View All Charts**.

   All the executed charts for the selected date are displayed in a new tab.

   > **Note:**
   > * If no list is displayed, you can select a date from the calendar to view a list of charts.
   > * If you want to view a specific chart, enter the chart name in the search criteria.

4. Click the chart name to view the chart details for that date.

Test a Chart

You can test a chart in the **Test a Chart** view.

To test a chart:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. Do one of the following:

   - In the **Charts** toolbar, click .
   - In the **Charts** panel, double-click a chart or select a chart and click .
   - In the **Charts** panel, select a chart and in Actions column, click  > **Edit**.

   The Build Chart view tab is displayed.

4. Click **Test Chart** to view the chart.

   The View Chart view tab is displayed.

5. In the **Time Range** drop-down, select the hours range.
6. Select the **Series**, either **Chart Values over Time** or **Chart with Totals**.
7. In the **Items To Plot** field, select the number of items to plot.
8. From the **Chart Type** drop-down list, select the chart type.
9. Click **Run Test** to run the test.

The chart data (if any) for the selected time range is displayed.

Investigate a Chart

You can investigate the chart by navigating directly to the Investigation module from the chart.

To investigate a chart:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** toolbar, click **View All Charts**.

   All the executed charts for the selected date from the **Chart Options** panel are displayed on a new tab.

4. Double-click the chart name to view the chart details such as the time at which the chart is executed and the default data source used for the chart execution.
5. Do one of the following:

- Click a data point on the chart to investigate.
- In the toolbar, click **Investigate** to investigate for the entire time range.

Manage a Chart Group and Chart

You can manage chart groups and charts using the following procedures.

Manage a Chart Group

Depending on the access permissions set for the user role, you can modify or delete, import or export, drag and drop a chart, or refresh a chart group.

Modify a Chart Group

To modify a chart group in the default folder or subgroups under a chart group:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts Groups** panel, select the chart group to modify.

The selected chart group is modified and can be viewed on the **Charts Groups** panel.

Delete a Chart Group

To delete a chart group in the default folder or subgroups under a chart group:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts Groups** panel, select the group and click −.

   A confirmation dialog asks for confirmation that you want to delete the selected group.

4. Click **Yes** to delete the group.

The selected group is deleted from the Chart Groups panel.

Import a Chart Group

To import chart groups from other instances of NetWitness:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. From the **Charts Groups** panel, select a folder to import the file.
4. Do one of the following:

   - In the Chart Groups panel, click  > **Import**.

     The **Import Chart** dialog box is displayed.

   - You can import multiple chart groups at the same time. To select multiple chart groups, press and hold the CTRL button and select the chart groups to be imported.

5. Click **Browse** to select the binary file.

   NetWitness provides a file system view of the files.

6. Locate the binary file and click **Open**.

   The file is added to the Import Chart list.

7. (Optional) To overwrite any existing rule in the library with an identically named rule in the binary file when importing, select the **Rule** checkbox. If you do not select the Overwrite option, and an identical rule is encountered in the binary file, the binary file is imported and no error message is displayed.
8. (Optional) To overwrite any existing list in the library with an identically named list in the binary file, select the **List** checkbox. If you do not select the Overwrite option, and an identical list is encountered in the binary file, the binary file is imported and no error message is displayed.
9. (Optional) To overwrite any existing chart in the library with an identically named chart in the binary file when importing, select the **Chart** checkbox. If you do not select the Overwrite option and an identical chart is encountered in the binary file, the binary file is imported and no error message is displayed.
10. Click **Import** to import the binary file.

Export a Chart Group

To export selected chart groups:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts Groups** panel, select a chart group and click  and do one of the following:

   - **Export** - This selection exports a chart in a .zip file.
   - **Export as Text** - This selection exports all the content from the Reporting Engine in a .zip file which contains the data in text format.

   You can export multiple chart groups at the same time. To select multiple chart groups, press and hold the CTRL button and select the chart groups to be exported. The exported file is saved to the local drive.

Drag and Drop Chart to a Group

To drag and drop a chart from the Charts panel to a group in the Charts Groups panel:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. Select a chart from the **Chart** panel and drag and drop the chart to a group in the **Chart Groups** panel.

The chart is copied to the group in the Chart Groups panel.

Refresh a Chart Group

To refresh chart groups:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts Groups** panel, drag and drop the group.

   The chart group is moved to the new location.

4. In the **Charts Groups** panel, Click .

The chart group is refreshed.

Manage a Chart

Depending on the access permissions set for the user role, you can modify or delete, duplicate, import and export, enable or disable charts, search for existing charts, and refresh a chart list.

Access Control for a Chart

To set access permissions for a chart:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** panel, select a chart.

4. Click  > **Permissions**.

   The Charts Permissions dialog box is displayed.

5. Based on the user role, select the appropriate buttons.

6. (Optional) Select the checkbox if you want to provide read access permission to dependent rules.

   Note: On selecting the check box, all dependent rules with No access permission are granted a READ access permission.

7. Click **Save**.

A confirmation message that the permission is successfully set for the selected chart is displayed.

Modify a Chart

To modify a chart in a group or subgroup:

1. Go to **Reports**.

   The Manage tab is displayed.

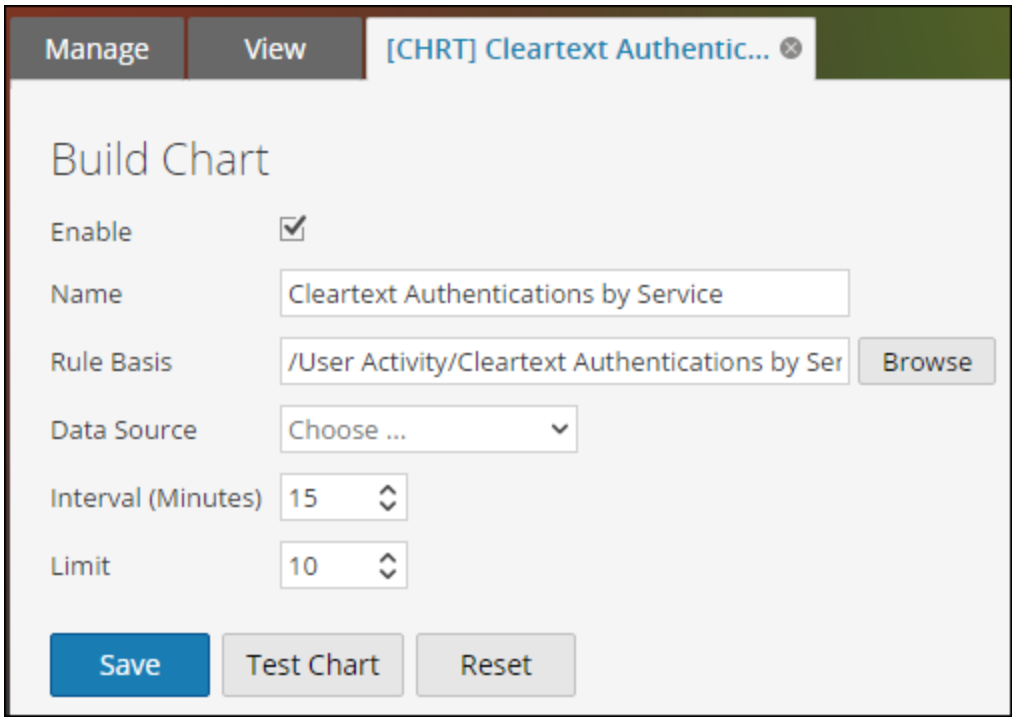
2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** panel, do one of the following:

   - Double-click a chart or select a chart and click .
   - Select a chart and click  > **Edit**.

     The Build Chart view tab is displayed.



4. Modify the name of the chart.
5. For the Reporting Engine to collect the data and generate chart results, select the **Enable** checkbox.
6. (Optional) In the **Rule Basis** field, do the following:

a. Click **Browse**.

The Add Rule dialog is displayed.

b. Navigate the Rule tree and select a rule.

c. Click **Select**.

The Rule appears in the Rule Basis field.

7. Select the data source from the **Data Source** drop-down list.

> **Note:** If the data source is not listed, then ensure you have **Read** permissions set for the data source. This is applicable for NWDB and Warehouse data sources. For more information, see the "Configure Data Source Permissions" topic in the *Host and Services Configuration Guide*.

8. (Optional) To modify the Interval value, click the up or down arrows.
9. Select the limit value to limit the number of records to be displayed.
10. Click **Save**.

A confirmation message that the chart is modified successfully is displayed.

Delete a Chart

To delete a chart in a group or subgroup:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Charts**.
   The Chart view is displayed.
3. In the **Charts** panel, do one of the following:
   - Select the charts and click .
   - Click  > **Delete**.
     A confirmation message asks if you want to delete the selected chart.
4. Click **Yes** to delete the chart.

A confirmation message that the chart is deleted successfully is displayed and the selected chart is deleted from the Charts panel.

Duplicate a Chart

To duplicate an existing chart:

1. Go to **Reports**.

   The Manage tab is displayed.
2. Click **Charts**.

   The Chart view is displayed.

3. From the **Charts** panel, select a chart to be duplicated.
4. In the **Charts** toolbar, click .

The chart is duplicated and gets added to the Charts panel.

Import a Chart

To import charts from other instances of NetWitness:

1. Go to **Reports**.

   The Manage tab is displayed.
2. Click **Charts**.

   The Chart view is displayed.
3. From the **Charts Groups** panel, select a folder from which to import the file.
4. Do one of the following:

   - In the **Chart** toolbar, click  > **Import**.

     The **Import Chart** dialog box is displayed.
   - You can import multiple charts at the same time. To select multiple charts, press and hold the CTRL button and select the charts to be imported.
5. Click **Browse** to select the binary file.

   NetWitness provides a file system view of the files.
6. Locate the binary file and click **Open**.

   The file is added to the Import Chart list.
7. (Optional) To overwrite any existing rule in the library with an identically named rule in the binary file when importing, select the **Rule** checkbox. If you do not select the Overwrite option, and an identical rule is encountered in the binary file, the binary file is imported and no error message is displayed.
8. (Optional) To overwrite any existing list in the library with an identically named list in the binary file, select the **List** checkbox. If you do not select the Overwrite option, and an identical list is encountered in the binary file, the binary file is imported and no error message is displayed.
9. (Optional) To overwrite any existing chart in the library with an identically named chart in the binary file when importing, select the **Chart** checkbox. If you do not select the Overwrite option and an identical chart is encountered in the binary file, the binary file is imported and no error message is displayed.
10. Click **Import** to import the binary file.

Export a Chart

To export selected charts to an external file:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** panel, select a chart and click  and do one of the following:

   - **Export** - This selection exports a chart in a .zip file.
   - **Export as Text** - This selection exports a chart from the Reporting Engine in a .zip file which contains the data in text format.

   You can export multiple charts at the same time. To select multiple charts, select the checkboxes of the charts to be exported. The exported file is saved to the local drive.

Enable a Chart

To enable a chart:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** panel, select a chart or several charts that display ○ in the **Enabled** column.

4. Click ●.

A confirmation message indicates that the chart(s) state is changed successfully.

Disable a Chart

To disable a chart:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** panel, select a chart or several charts that display ● in the **Enabled** column.

4. Click ○.

A confirmation message indicates that the chart(s) status is changed successfully.

Search an Existing Chart

To search for an existing chart:

1. Go to **Reports**.

   The Manage tab is displayed.

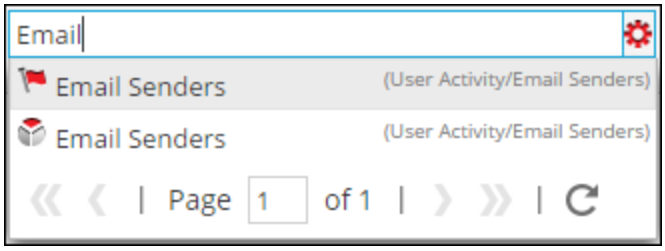
2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** toolbar, enter text in the Search text box.

4. Click > **Chart**.

All charts that match the search string are displayed in the search drop-down list.



Refresh a Chart

To refresh charts:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Charts**.

   The Chart view is displayed.

3. In the **Charts** panel, drag and drop the charts to the desired group in the **Charts Groups** panel.

   The charts are moved to the new location.

4. Do either of the following:

   - In the **Charts** panel, click .
   - In the **Charts** toolbar panel, select **Auto Refresh**.

   The Chart list is refreshed.

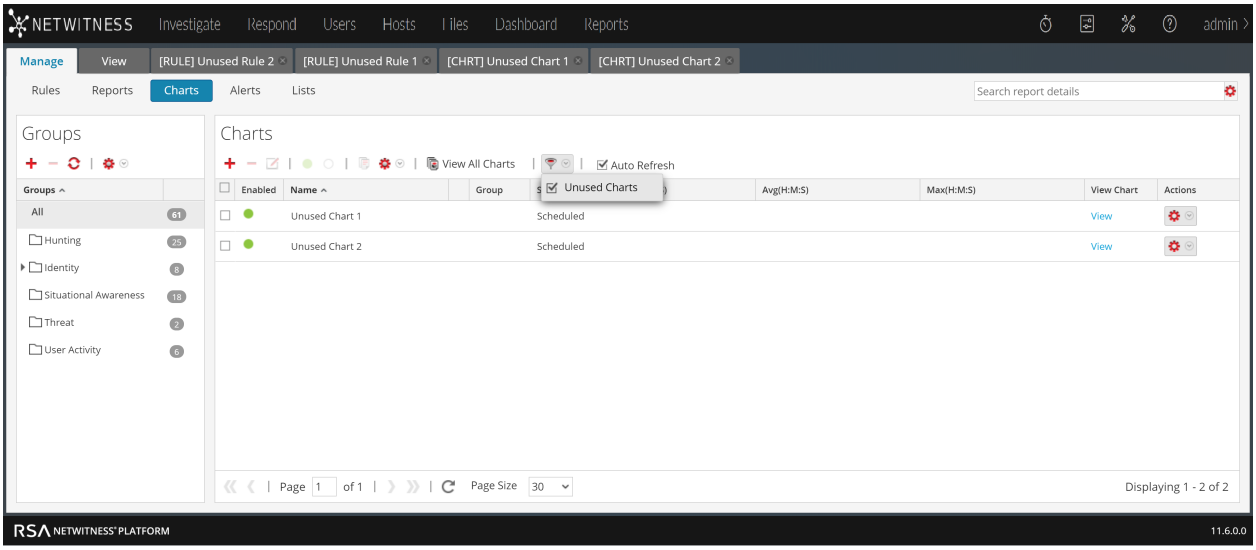
Filter Unused Charts

You can filter and delete the charts that are not used in any dashlet.

To filter unused charts:

1. Go to **Reports**.
   The **Manage** tab is displayed.

2. In the **Charts** panel, click  and select **Unused Charts**.
   List of unused charts will get displayed.
3. Delete the unused charts.
   For more information, see Delete a Chart.



Working with Alerts

The Alerting module user interface provides access to NetWitness alerts. The following topics discuss alerts:

- Alerting Overview
- Configure Reporting Engine
- Configure an Alert
- Schedule an Alert
- View an Alert
- Investigate an Alert
- Manage an Alert and Alert Template

Alerting Overview

Alerts can be used to generate timely insights about current security issues, vulnerabilities, and exploits. For example, when a malicious email is sent from a compromised account, you would need an alert that automatically notifies you when such an event occurs.

The following concepts of alerting will help you understand more about alert rules, conditions, notifications, and templates.

Alert Rules

Alert rules specify the logic for alert generation. Alert rules allow you to set up threshold limits and define how to be notified if these limits are exceeded. For example, you may set up a rule to be alerted if the CPU usage remains abnormally high for 5 minutes or more.

Alert Definitions

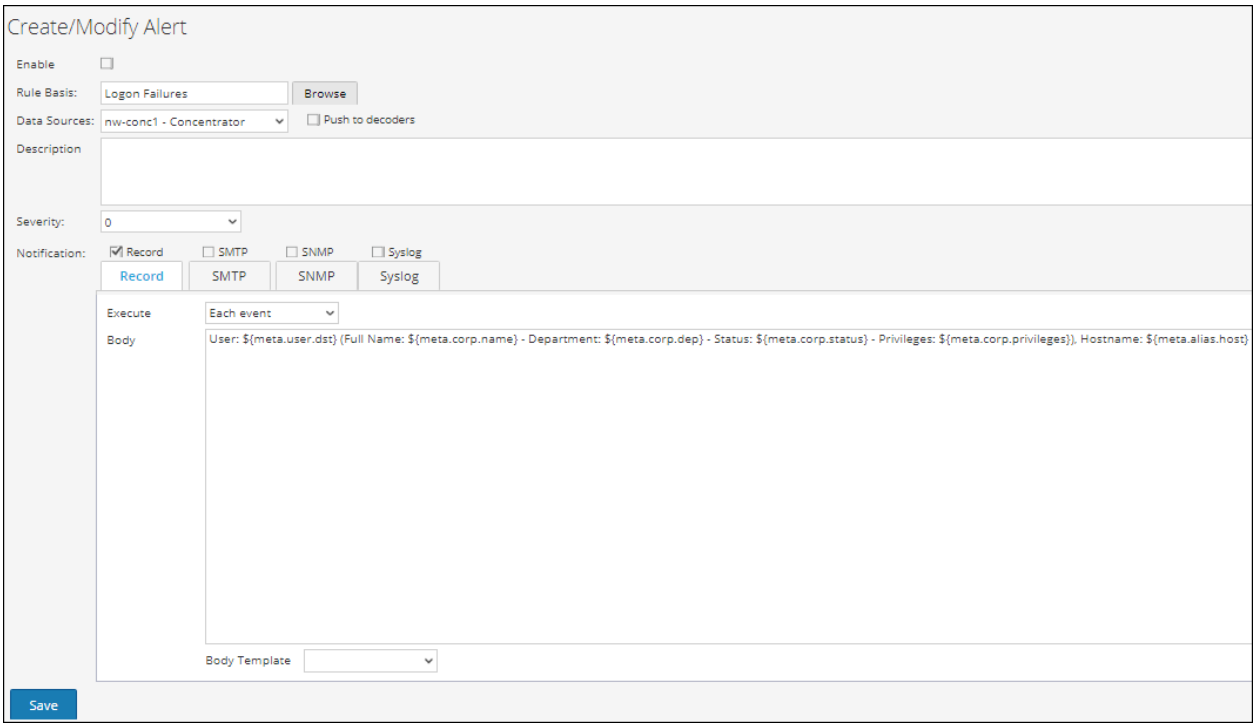
The alert definition is similar to defining rules for reports. These rules must be defined based on your use case. Alert definitions are made by selecting the alert rules you define in the Build Rule view. You select this rule while defining an alert.

Note: You can only alert using rules defined for NetWitness data source.

Once an alert is created, this data is collected from the Reporting Engine and displayed on the user interface.

Once an alert is defined, you can schedule the alert to run every minute (by default), or run at the present time, or run at the near future.

Note: In the NetWitness user interface, wherever Date and Time is displayed, it is always according to the user selected time zone profile.



Alert Notifications

The following are the components required to configure alert notifications:

- Notification server – Notification Server is used to send alert notifications. For example, SMTP mail server. Once you configure a notification server, you can add it to a rule. When the rule triggers an alert, the rule will use that server to send alert notifications.
- Notifications – Alert outputs, which can be email, SMTP, SNMP, and Syslog.
- Templates – The pre-defined format of an alert message.

When ever the rule condition is encountered, alerts get generated based on the severity level and notifies the user depending on the notification method set for that specific alert. The following are the various notification methods:

- Email/ SMTP: Simple Mail Transfer Protocol (SMTP) sends alert emails for system activity. Email alerts can be sent to their intended recipients by selecting SMTP as notification type.
- SNMP: Simple Network Management Protocol (SNMP) sends alerts to multiple computers for SNMP traps. SNMP alerts can be sent to other computers by selecting SNMP as notification type.
- Syslog: Syslog alerts generate notifications from Syslog messages. Syslog alerts can be sent by selecting Syslog as notification type.

Alerts can be configured to notify events that require attention, or as mechanisms to take automated actions based on conditions configured in an alert. An alert is sent when conditions within the entity have met the criteria selected for the alert. The notification criteria determines when and at what frequency the alert is generated.

Alert Templates

Alert templates are pre-defined format for an alert message. You can use these templates to create alerts.

Access Control for Alerting

Depending on the user role, the user is provided with specific set of access permissions in order to manage an alert. The Administrator manages the access rights provided to each user role from the **(Admin) > Security > Roles** tab. You can set access permissions for the user roles to manage an alert. The Reporting module provides access control at the alert level.

Note: Reporting Engine Alert permissions are prefixed with 'RE' to distinguish them from Event Streaming Analysis (ESA).

When you create users and user roles, ensure that the roles that you create for specific tasks have access to all the necessary permissions. This could require permissions at several levels of the role hierarchy.

Alerts can be combined with a specific set of user roles so that when a user logs into NetWitness, the only alerts they can access are alerts accessible by the role to which the user belongs. Users that belong to a user role with the '**Read & Write**' access permission can define alerts. The access can further be tightened so that the alerts are accessed only by those who have the '**Read Only**' access.

At the alert level, you can set the following access permissions for the user roles in NetWitness:

- Read & Write
- Read Only
- No Access

Note: Before applying the Alert permissions, the default permission set for all the user roles is '**No Access**' permission and the checkbox is unchecked.

If you want to change the access permission for a specific user role, you must set it at the alert level. Except for administrators, the default permission set for all the other user roles is '**No Access**' permission.

The two scenarios are explained in brief:

- Scenario 1: Permissions applied to alert/ rules based on the user role.
- Scenario 2: Read-only permission applied to rules in the Alert.

	Role (Analysts)	Permissions applied to Alert/ Rules based on the user role	Permission (Read-only) applied to Rules in the Alert
Alert	**Read & Write**	**Read & Write**	**Read & Write**
Rules	**Read**	**Read**	**Read**

The Alert is assigned the role of a Security Analyst and permissions are set to **Read & Write** alerts.

For scenario 1, each of the levels has a permission set based on the user role. For scenario 2, the **Read** permission is set for the Rules except that the permission for the rules must not be higher than the permission for the Alerts.

If the permission for the rules is higher than the permission for the Alerts, the permission is not applied. For example, if you set the permissions for the Alert as **No Access** and then specify the option *Apply Read-only permission to Rules in the Alerts*, the read-only permission is not set for the rules.

Access Control for an Alert When Multiple Alerts are Selected

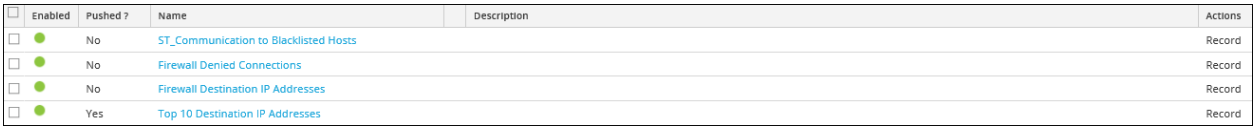
When you want to change permissions of multiple alerts, you must select several alerts and set their access permissions using the Alert Permissions panel. The access permission that you choose is applied to all the selected alerts.

Log in as a specific user and view the access details

When you log in to the NetWitness UI as a user having **Read** access permission, all the alerts will be denoted with the symbol (📖) and when you click on the symbol the 'Read Only' callout is displayed on the Alert panel.

When you log in to the NetWitness UI as a user not having **Read & Write** access permission on an Alert, all the alerts will be denoted with the symbol ( ) and the alerts appear grayed out on the Alert panel.

The following figure shows the Alert panel when logged in with minimal **Read & Write** access permission.

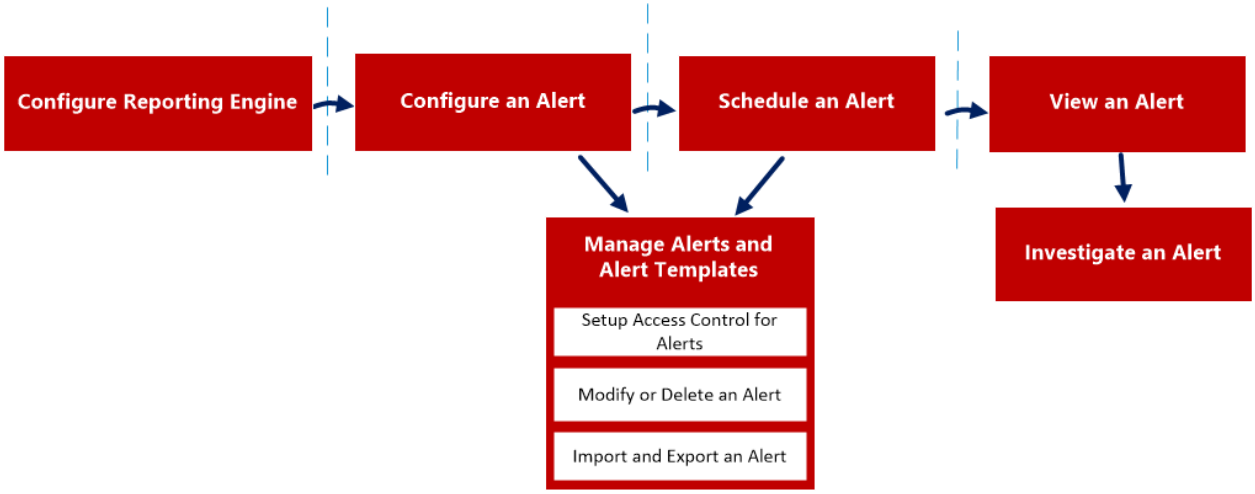


Note: If a user (other than ADMIN) creates an alert, ADMIN cannot access that alert.

The following table lists the various columns in the Alert Permissions panel:

Column	Description
Roles	The role of the user logged into the NetWitness user interface.
Read & Write	The user can access, view, edit, import, export, and delete the alert on the Alerts page. The user can also change the permission on the alert.
Read Only	The user can only access and view the alert on the Alerts page.
No Access	The user cannot access or view the alert for which this permission is set.
Apply Read-only permission to Rules in the Alerts	The user can automatically apply permissions to the rules in the alerts.

The following is an overview of the entire process of alerting:



To configure and generate an alert on Reporting Engine, perform the following:

1. Configure Reporting Engine
2. Configure an Alert
3. Schedule an Alert
4. View an Alert
5. Investigate an Alert
6. Manage an Alert and Alert Template

Configure Reporting Engine

Ensure that:

- You have Decoders that are connected to the Concentrator added to the Reporting Engine for the selected data source, before creating an alert rule.
- You have installed and configured a Syslog server that supports TCP/TLS in your environment. For example, WinSyslog. You can configure the Reporting Engine to send Syslog messages over TCP with Transport Layer Security (TLS) when an alert is triggered.

To configure the Reporting Engine to send Syslog alerts over TCP with Transport Layer Security (TLS):

1. Obtain the required certificates.
2. Append the CA certificate to the ca.pem file on the NetWitness server.
3. Configure the Syslog server to accept messages from client machines.
4. Configure the delivery of alert messages in the NetWitness UI.

Task 1: Obtain the required certificates

To generate certificates for configuring Reporting Engine to send Syslog messages over TCP with TLS:

1. Generate a Certifying Authority (CA) certificate. For more information, see http://www.rsyslog.com/doc/tls_cert_ca.html.

Note: You can ignore this step if you already have a CA running in your environment.

2. Generate a key pair for the Syslog server. For more information, see http://www.rsyslog.com/doc/tls_cert_machine.html.

Note: You can ignore this step if you have already configured security for the Syslog server using the key and certificates generated by the same CA.

Task 2: Append the CA certificate to the ca.pem file on the NetWitness Server

To append an existing CA certificate to the ca.pem file:

1. Manually append the contents of the CA certificate that you generated to the **/etc/pki/CA/certs/ca.pem** file.
2. Run the following command on the NetWitness server to have the certificate populate to the Truststore:
   ```
   keytool -import -file /etc/pki/CA/certs/ca.pem -keystore cacerts
   ```

Task 3: Configure the Syslog Server to accept messages from client machines

To configure the Syslog server to accept messages from client machines that have the same CA certificates:

1. Copy the following files to your secure TCP server target location:
   - ca_cert.pem
   - server_cert.pem
   - server_key.pem

   Where:

   ca_cert.pem - is the CA certificate

   server_cert.pem - is the server certificate

   server_key.pem - is the server key

   For more information, see the documentation specific to your Syslog server. If you are using rsyslog, refer to http://www.rsyslog.com/doc/tls_cert_server.html.

Task 4: Configure the delivery of alert messages in NetWitness

Configure Reporting Engine to send Syslog messages over TCP with Transport Layer Security (TLS) when an alert is triggered by enabling **SECURE_TCP** in the **Output Actions** tab for the Reporting Engine service in the Reporting Engine Services Config View. For more information, see the **"Reporting Engine Output Actions"** topic in the *Host and Services Configuration Guide*.

Configure an Alert

You can configure an alert by setting up alert notifications and adding a notification method to a rule.

Note: Only Administrators can set up these notifications.

To configure an alert:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Alerts**.

   The Alert view is displayed.

3. In the **Alert** toolbar, click .

   The Create/Modify Alert panel is displayed.

4. Click **Enable** to enable the alert.
5. In the **Rule Basis** field:
   a. Click **Browse**.

      The Lookup Rule Basis dialog box is displayed.

   b. Navigate the Rule tree and select a rule.
   c. Click **OK**.

      The Rule name is displayed in the Rule Basis field.

6. From the **Data Sources** drop-down list, select a data source.

   Note: If the data source is not listed, then ensure you have **Read** permissions set for the data source. This is applicable for NWDB and Warehouse Connector data sources. For more information, see **"Configure Data Source Permissions"** topic in the *Host and Services Configuration Guide*.

7. Select the **Push to decoders** checkbox for the Reporting Engine to send the rule to the Decoder.
8. (Optional) Enter an alert description in the **Description** field.
9. From the **Severity** drop-down list, select the severity level.
10. In the **Notification** field:
    a. Select the appropriate notification.
       The selected notification tab is displayed in the Create/Modify Alert dialog box.
    b. (Optional) Deselect the notification to disable the notification tab.
    c. Define an action in one of the **Notification** tabs:

i. In the **Record** tab field:

   a. From the **Execute** drop-down list, select the frequency for recording an alert.

   b. Enter the RECORD message. You can create a new message or select a template in the **Body Template** field and modify the template here.

   c. (Optional) If templates have been defined, select a template for the RECORD message that you can use as is or modify.

ii. In the **SMTP** tab field:

   a. From the **Execute** drop-down list, select a value to identify the number of times to send an email message for the alert.

   b. Enter an email address or comma-separated list of email addresses to send this alert.

   c. Enter the subject of the email message.

   d. Enter the body of the message. You can create a new message or select a template in the **Body Template** field and modify the template here.

iii. In the **SNMP** tab field:

   a. From the **Execute** drop-down list, select a value to identify the number of times that you want to send an SNMP message for the alert.

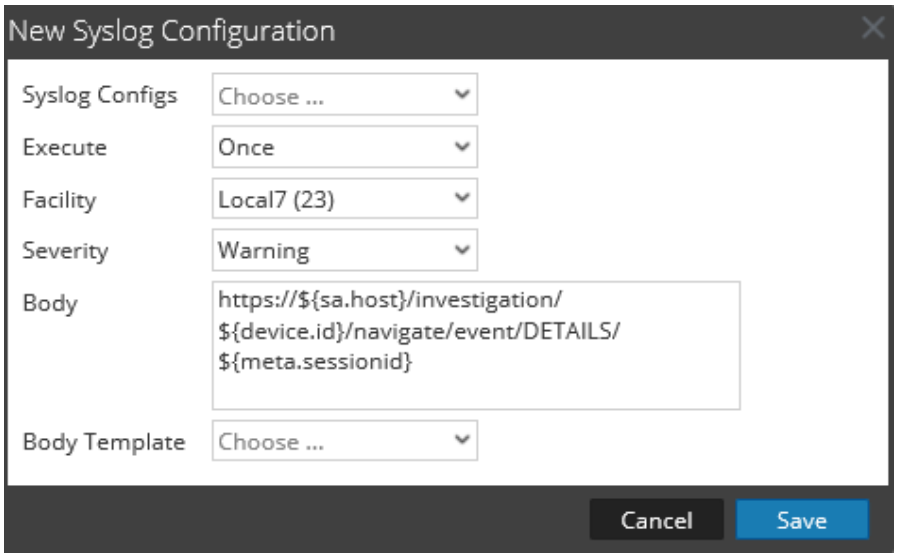
   b. Enter the SNMP message. You can create a new message or select a template in the **Body Template** field and modify the template here.

iv. In the **Syslog** tab field:

> **Note:** You can configure Multiple Syslog servers on the Syslog Configuration panel. For more information, see **"Reporting Engine Output Actions"** topic in the *Host and Services Configuration Guide*.

   a. Click .

   The New Syslog Configuration dialog box is displayed.



b. From the **Syslog Configs** drop-down list, select a value for the syslog configuration.

c. From the **Execute** drop-down list, select a value to identify the number of times to send a Syslog message for the alert.

d. From the **Facility** drop-down list, select the facility.

e. From the **Severity** drop-down list, select the severity level.

f. Enter the Syslog message. You can create a new message or select a template in the **Body Template** field and modify the template here.

> **Note:** If you want to add a metakey, specify the same in the format: ${meta.metakey}. For example, ${meta.ip.dst}.

g. Click **Save**.
The Syslog configuration gets added to the alert.

11. Click **Create**.

NetWitness creates an alert with a confirmation message that the alert is saved successfully. NetWitness generates the alert and executes the output actions every minute.

Schedule an Alert

You must schedule an alert to search for events on a regular schedule.

To schedule an alert:

1. Go to **Reports** to view the Manage tab.
2. Click **Alerts** to open the Alert view.
3. Select an alert to schedule.
4. On the **Alert** toolbar, click **Enable**.
The selected alert is scheduled.

View an Alert

You can view an alert or a list of all alerts.

You can view the alerts triggered and investigate any alert in the Investigation module and customize these views to show alerts for a specific period of time, and set the maximum number of alerts displayed in a single page.

To view an alert:

1. Go to **Reports** to view the Manage tab.
2. Click **Alerts** to open the Alert view.
3. On the **Alert** toolbar, click **View Alerts**.
   The View Alerts view is displayed.

Investigate an Alert

You can investigate every alert that is triggered on the Alert View. For more detailed investigation on a particular alert, you can view the alert on the Investigation module.

To investigate an alert:

1. In the **Alert** section toolbar, click **View Alerts** to navigate to the View Alerts view.
2. Do one of the following:
   - Click the  button against the alert you want to investigate.
     The Investigation module displays the details of the first session that registered the match for the given alert for immediate analysis.
   - Click on the alert name of the alert you want to investigate.
     The Investigation module displays all matches for that particular alert for the hour surrounding the registered alert.

Manage an Alert and Alert Template

You can manage alerts, scheduled alerts, and alert templates using the following procedures.

Manage an Alert

Depending on the access permissions set for the user role, you can modify or delete, import and export, enable or disable alerts, view or refresh an alert list.

Access Control for an Alert When a Single Alert is Selected

To set access permissions for an alert:

1. Go to **Reports**.
   The Manage tab is displayed.

2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** panel, select an alert.
4. Click > **Permissions**.
   The Alert Permissions dialog box is displayed.
5. Based on the user role, select the appropriate options.
6. (Optional) Select the checkbox if you want to automatically provide read access permission to dependent rules.

Note: When the check box is selected, all dependent rules with the No access permission will be given the READ access permission.

7. Click **Save**.
   A confirmation message that the permission is successfully set for the selected alert is displayed.

Access Control for an Alert When Multiple Alerts are Selected

To change permissions of multiple alerts:

1. In the **Alerts** panel, select all the alerts whose permissions must be set.
2. Click > **Permissions**.
   The Alert Permissions dialog box is displayed.
3. Select the permission to set for the respective user role.
4. Click **Save**.
   A confirmation message that the permission is successfully set for all the selected alerts is displayed.

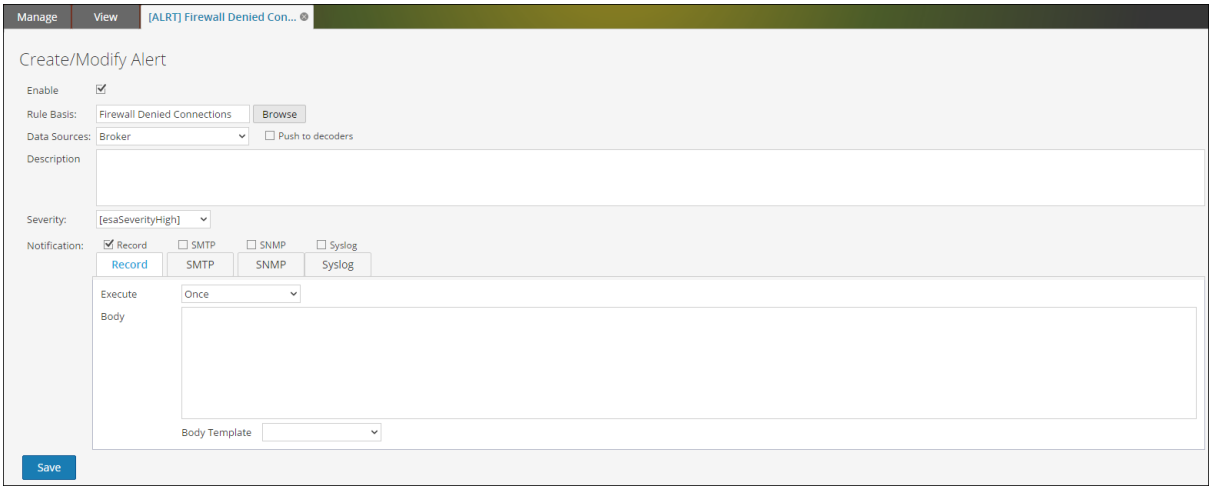
Edit an Alert

For example, if you want to be notified about the alert over an email on a different Email ID, you will have to modify the alert notification section with the new Email ID details to be reverted over an email when an alert is generated. Additionally, you can also modify the alert description and alert notification in the Create or Modify Alert panel.

To edit an alert:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.

3. In the **Alert** panel, select an alert and click .
   The Create or Modify Alert tab is displayed.



4. In the **Rule Basis** field, navigate the rule tree and select another rule.
   The Rule name is displayed in the Rule Basis field.
5. (Optional) Select a data source from the **Data Sources** drop-down list.

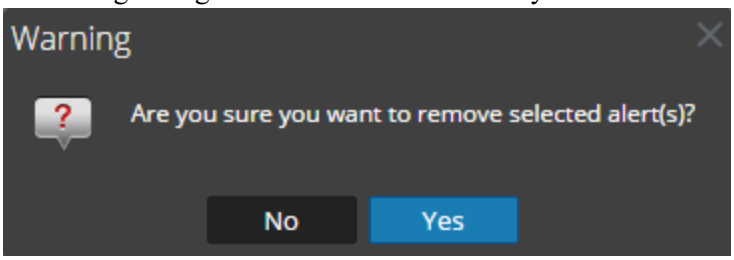
> **Note:** If the data source is not listed, then ensure you have **Read** permissions set for the data source. This is applicable for NWDB and Warehouse data source. For more information, see **"Configure Data Source Permissions"** topic in the *Host and Services Configuration Guide*.

6. (Optional) Modify the alert description in the **Description** field.
7. Modify the appropriate **Notification** tabs – **RECORD**, **SMTP**, **SNMP**, and **Syslog**.
8. Click **Save**.
   A confirmation message that the alert is modified successfully is displayed.

Delete an Alert

To delete an alert:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** panel, select the alert and click .
   A warning dialog asks for confirmation that you want to remove the selected alerts.



4. Click **Yes** to delete the alert.
   A confirmation message that the alert is deleted successfully is displayed and the selected alert is deleted from the Alert panel.

Import an Alert

To import an alert from other instances of NetWitness in the Alerts panel:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** toolbar, click > **Import**.
   The Import Alert dialog box is displayed.
4. Click **Browse** to select the binary file.
   NetWitness provides a file system view of the files. You can import multiple alerts at a time. To select multiple alerts, select the checkbox of the alert to be imported.
5. Locate the binary file, and click **Open**.
   The file is added to the Import Alert list.
6. (Optional) To overwrite any existing alert in the library with an identically named alert in the binary file when importing, select the Alert checkbox. If you do not select the Overwrite option, and an identical alert is encountered in the binary file, the binary file is imported and no error message is displayed.
7. Click **Import** to import the binary file.

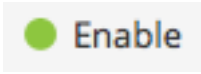
Export an Alert

To export an alert to an external file that can be later imported to NetWitness:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** panel, select an alert and click and do one of the following:
   - **Export** - This selection exports an alert in a .zip file.
   - **Export as Text** - This selection exports all the content from the Reporting Engine in a .zip file which contains the data in text format.
     You can export multiple alerts at a time. To select multiple alerts, check the checkbox of the alert to be exported.
4. Click > **Export**.
   The exported binary file is saved to the local drive.

Enable an Alert

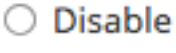
To enable an alert:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** panel, select the alert that displays ○ in the **Enabled** column.
4. Click 

.
   A confirmation message shows that the change to the alert(s) state was successful.

Disable an Alert

To disable an alert:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** panel, select the alert that displays ● in the **Enabled** column.
4. Click 

.
   A confirmation message shows that the alert(s) status is changed successfully.

View an Alert List

To view an alert list:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** toolbar, click **View Alerts**.
   The View Alerts view tab is displayed.
4. Select the last number of days from the drop-down list.
5. Enter a value for the **Max no of alerts**.
   The alerts list is displayed based on the chosen filter value.

Refresh an Alert List

To refresh the list of alerts:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. From the **Alert** toolbar, click  to refresh the alerts list.
   The Alert panel is refreshed.

Manage a Scheduled Alert

You can enable or disable a scheduled alert, and view all scheduled alerts.

Enable a Scheduled Alert

To enable a scheduled alert:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. Click View Schedule.
   The View Alerts Schedule view tab is displayed.
4. In the **Alerts Schedule List** panel, select the scheduled alert (s) to be enabled.
5. Click .
   A confirmation message indicates that the alert(s) status is changed successfully and the alert is now available in the Alert panel.

Disable a Scheduled Alert

To disable a scheduled alert:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. Click View Schedule.
   The View Alerts Schedule view tab is displayed.
4. In the **Alerts Schedule List** panel, select the scheduled alert (s) to be disabled.
5. Click .
   A confirmation message indicates that the alert(s) status is changed successfully and the alert is now available in the Alert panel.

View all Alerts Scheduled

To view all the alerts scheduled:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. In the **Alert** toolbar, click **View Schedule**.
   The View Alerts Schedule view is displayed with a list of all the scheduled alerts.

Manage an Alert Template

You can modify or delete an alert template, and view all alert templates.

Edit an Alert Template

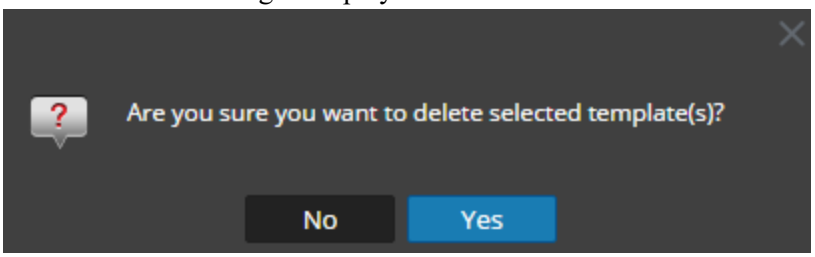
To edit an alert template:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. Click Template.
   The Alert Template view is displayed.
4. In the **Alert Template** panel, select a template and click .
   The Create/Modify Template dialog box is displayed.
5. Click **Save**.
   A confirmation message that the template is modified successfully is displayed.

Delete an Alert Template

To delete an alert template:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.
3. Click Template.
   The Template view tab is displayed.
4. In the **Alert Template** panel, select a template and click .
   A confirmation dialog is displayed.



5. Click **Yes** to delete the template.
   A confirmation message that the template is deleted successfully is displayed.

View all Alert Templates

To view all alert template messages:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Alerts**.
   The Alert view is displayed.

3. In the **Alert** toolbar, click **Template**.
   The Template view tab is displayed with a list of templates.

List of Available Variables

Below are the list of variables available on the Reporting Engine notification for Reporting Engine alerts.

```
Name: ${name}
Severity: ${severity}
alert count: ${count}
Start_session_id = ${sid1}
end_session_id = ${sid2}
data source id = ${device.id}
netwitness host = ${nw.host}
```

See below example for how to use metadata information on the notification template.

To use ip.src and ip.dst meta, use the format `${meta.<meta-name>}`.

```
ip.src meta = ${meta.ip.src}
ip.dst meta = ${meta.ip.dst}
```

Appendix

This section provides detailed information about the supported aggregate functions, rule syntax, advanced rules query syntax in Reporting and task scheduler for Warehouse Reporting.

Rule Syntax

This section describes the different rule syntax supported in the Reporting Engine.

NWDB Rule Syntax

The NWDB rule is one of the rule syntax supported in the Reporting Engine. To enhance the execution time of your reporting entities, see "Reporting Guidelines" section in Reporting Overview.

A Rule is a function that manipulates the result set of a rule in order to make the output in a report more meaningful or add additional functionality to a rule other than querying data and displaying it. Any combination of these rule actions can be used to create unique and interesting representations of the information collected by NetWitness.

The Reporting Engine supports the following categories of NWDB data source rule syntax:

- **select** clause
  - Non-Aggregate Rule
  - Aggregate Rule
- **alias**
- **where** clause
- **where** clause Operators
- **then** clause
- **Limit** field
- Rule Actions
- Rule Operators

Select Clause

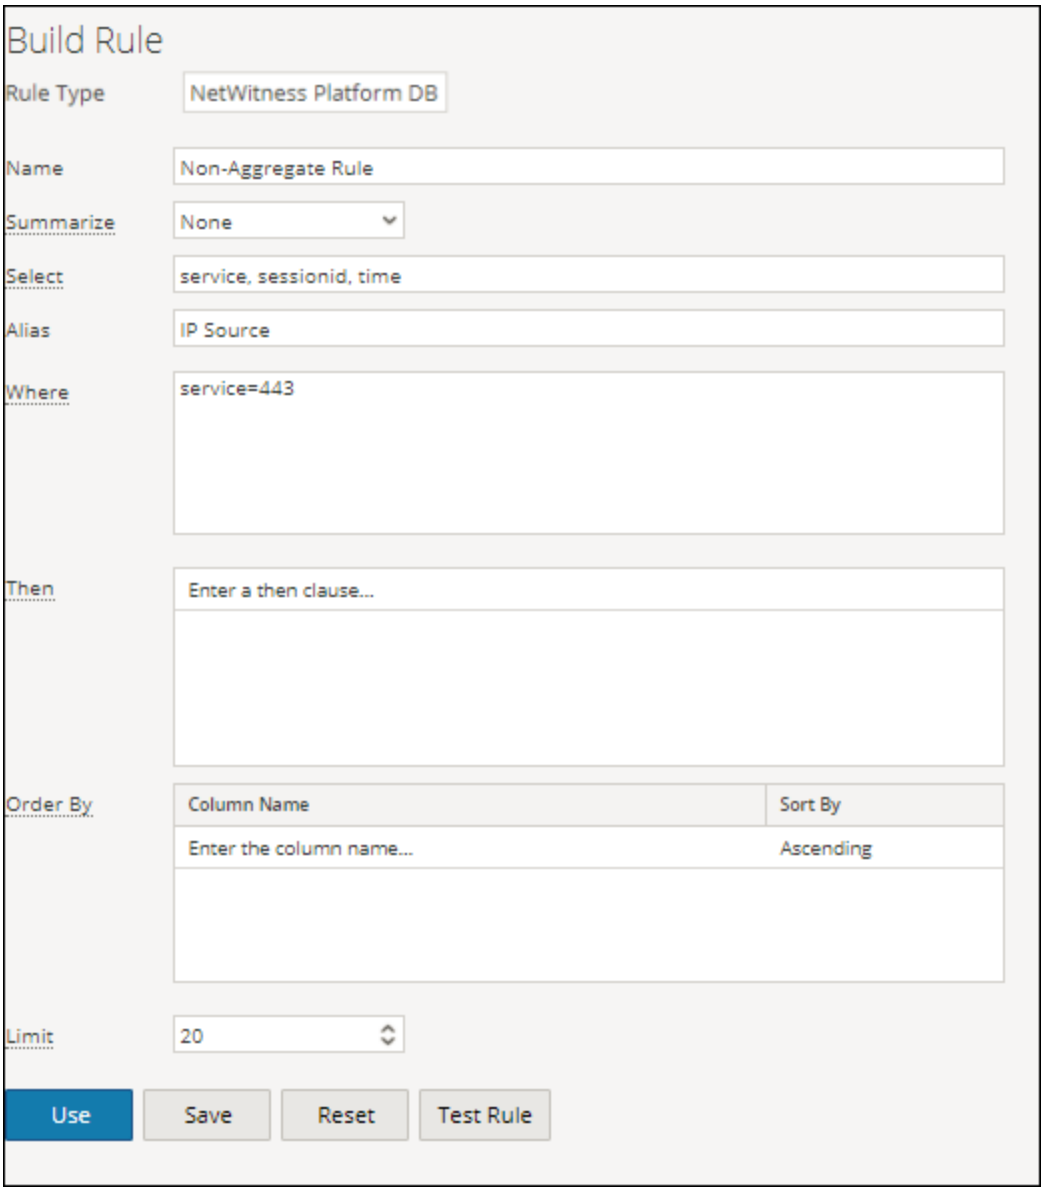
The select clause is a comma separated list of values. For example: select sessionid,time,service.

There are two types of select clause for NWDB Rule:

- Non-aggregate rule
- Aggregate rule

Non-Aggregate Rule

When you want to define a rule without any grouping, choose 'None' in the Summarize field. In a non-aggregate rule, you can select any number of metas in the *Select* clause. For example, select service, sessionid, time.

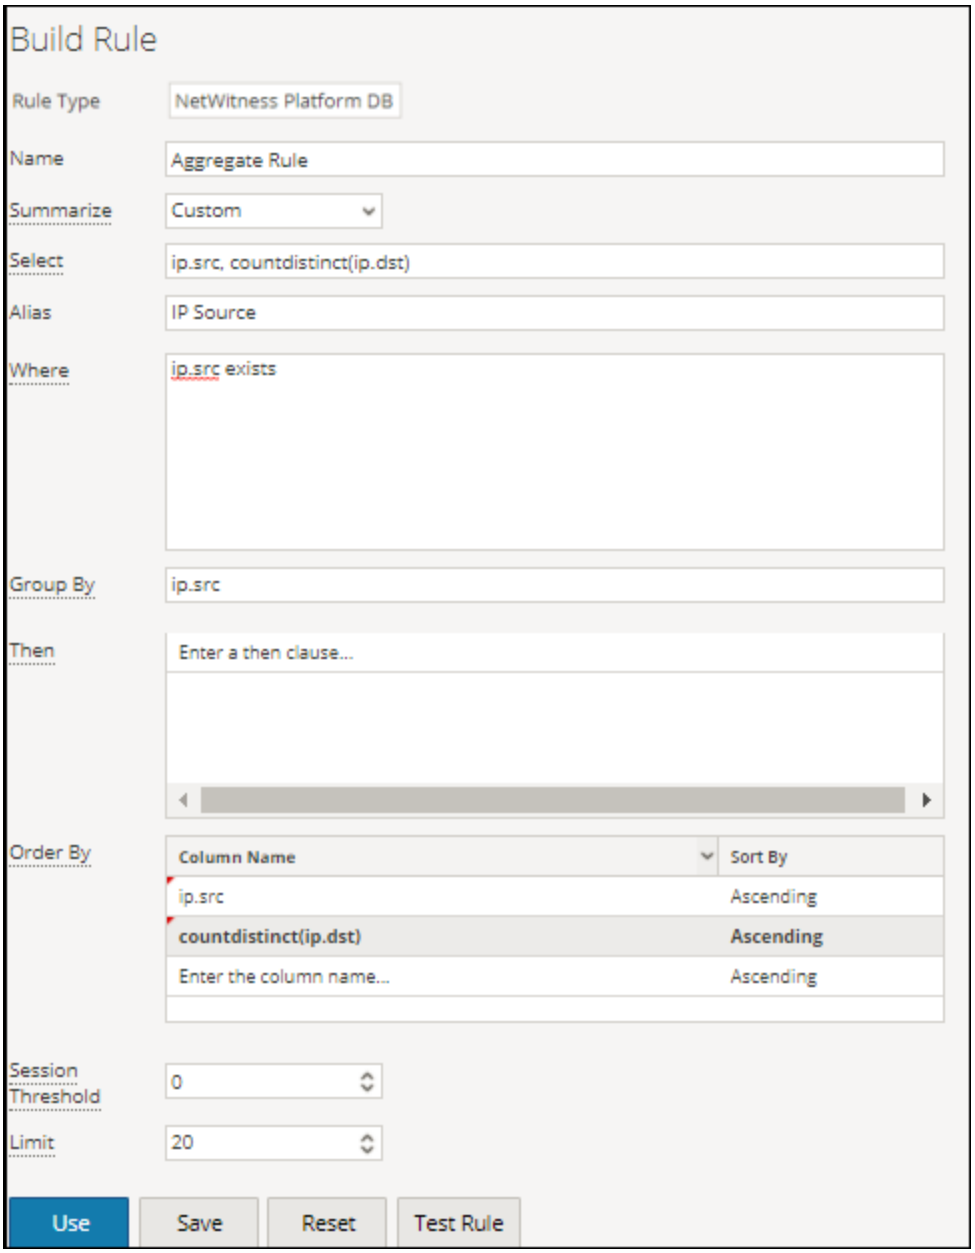


Aggregate Rule

When you want to query for a specific meta and its associated aggregate value then you must use the Aggregate rule. To get an aggregate, you must choose either of the three metas (Event Count, Packet Count, Session Size) or choose 'Custom' in the **Summarize** field to include an aggregate function in the *Select* clause. For example, select ip.src, sum (ip.dst). When Custom aggregate rule is enabled, the following fields are populated in the user interface:

- Group By
- Order By
- Session Threshold

The following figure shows the Build Rule view for Aggregate Rule.



There are two types of aggregate values that can be queried:

- Collection aggregation
- Meta aggregation

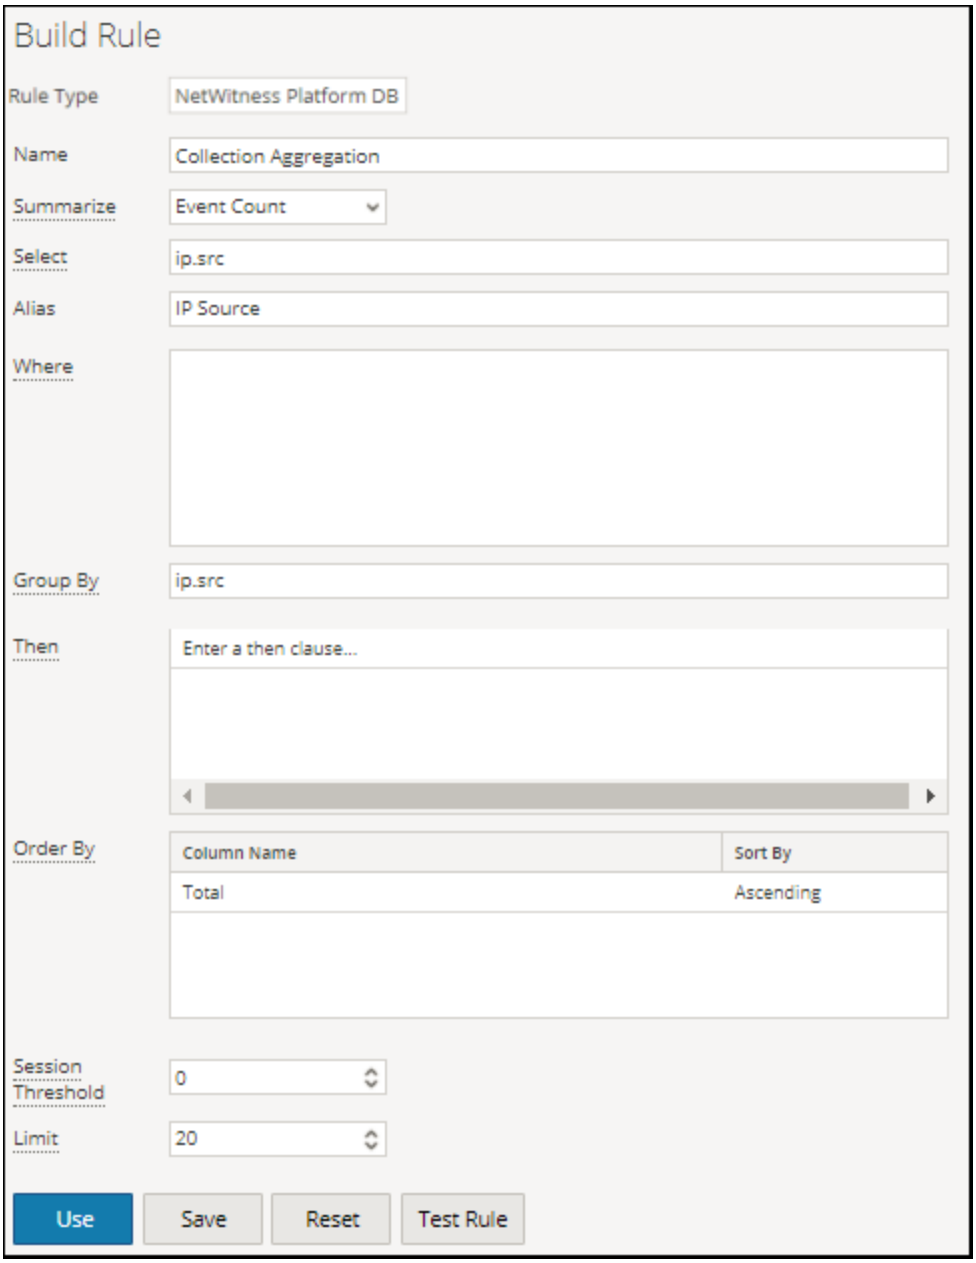
Collection Aggregation

With collection aggregation, you can get aggregates related to Event, Session or Packets. The following values can be queried in a collection aggregation:

- **Event Count**: The total count of events.
- **Packet Count**: The total count of packets.

- **Session Size**: The total session size.

These options are listed in 'Summarize' field and any one of them can be selected in a rule.
For example, choose any of the Collection aggregates (Event Count or Packet Count or Session Size) in the 'Summarize' field and select ip.src.



Meta aggregation

With meta aggregation, you can get aggregates of meta values. The following are the supported meta aggregate functions:

- sum(meta)
- count(meta)
- countdistinct(meta)
- min(meta)
- max(meta)
- avg(meta)
- first(meta)
- last(meta)
- len(meta)
- distinct(meta)

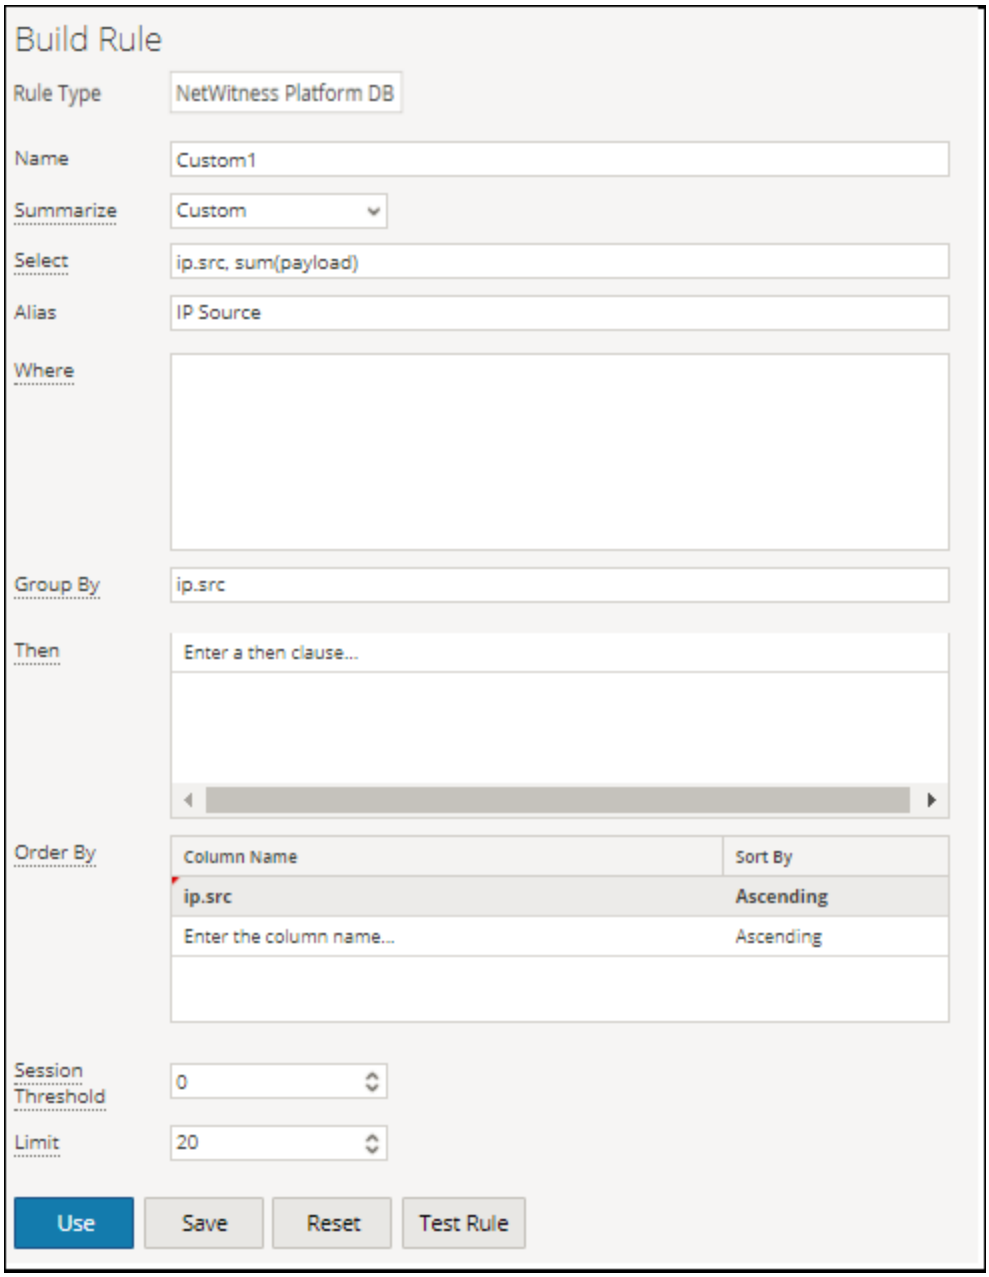
Supported Meta Aggregate Functions

The NWDB service supports the following meta aggregate functions and syntax in this release.

Syntax	Function
sum(<meta>)	The sum of all meta values. For example, if you provide the field sum(payload) in the select clause, the resultset is the sum of payload size. **Note:** The meta field chosen for the sum aggregate function must be of numeric data type.
count(<meta>)	The total number of meta fields that would be returned. For example, if you provide the field count(ip.dst) in the select clause, the resultset is the number of times an ip.dst value is returned.
countdistinct (<meta>)	The total number of distinct meta fields that would be returned. For example, if you provide the field countdistinct(ip.dst) in the select clause, the resultset is the number of times a distinct ip.dst value is returned.
min(<meta>)	The minimum of all meta values. For example, if you provide the field min(payload) in the select clause, the resultset is the min of payload size.
max(<meta>)	The maximum of all meta values. For example, if you provide the field max(payload) in the select clause, the resultset is the max of payload size.
avg(<meta>)	The average of all meta values. For example, if you provide the field avg(payload) in the select clause, the resultset is the avg of payload size. **Note:** The meta field chosen for the avg aggregate function must be of numeric data type.

Syntax	Function
first(<meta>)	The first occurrence of the meta value. For example, if you provide the field first(ip.src) in the select clause, the resultset is the first occurrence of ip.src for that group.
last(<meta>)	The last occurrence of the meta value. For example, if you provide the field last(ip.src) in the select clause, the resultset is the last occurrence of ip.src for that group.
len(<meta>)	Converts all field values to a UInt32 length instead of returning the actual value. This length is the number of bytes to store the actual value, not the length of the structure stored in the meta database. For instance, the meta value "NetWitness" returns a length of 10. All IPv4 fields, like ip.src, returns 4 bytes.
distinct (<meta>)	The distinct values of the meta. For example, if you provide the field distinct(ip.src) in the select clause, the resultset is all the distinct ip.src for that group.

You must select 'Custom' in 'Summarize' field and provide the meta and the meta aggregate functions in the select clause.



Note: Meta aggregate functions cannot be used in a WHERE clause and the rule actions like min_threshold/max_threshold can be used to filter aggregate functions. It is advised to use a more refined WHERE clause to get a better rule performance while using 'group by'.

Aggregate Query for Multiple Meta

To execute aggregate query for multiple Meta, follow these steps:

1. Go to **Reports**.

   The Manage tab is highlighted and the **Rules** view is displayed.

2. In the **Rules** toolbar, click + > **NetWitness Platform DB**.

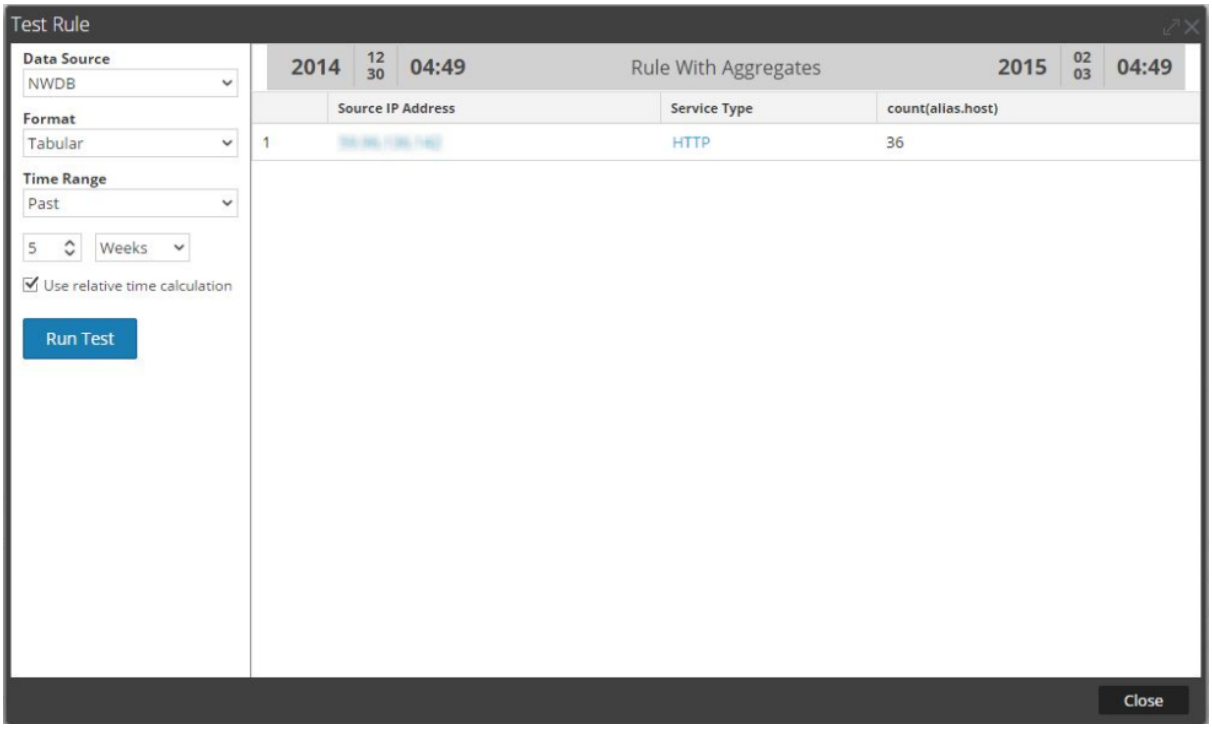
   For example, enter the following meta in the fields highlighted below:

   SELECT: ip.src, service, count(alias.host)
   ALIAS: Source IP Address, Service Type, count(alias.host)
   WHERE: ip.src = 59.96.136.142

   Note: In the alias field you can enter a name for columns used in the select clause. If you do not specify the alias for one of the field in the select clause, then the default description will be used. For example, if the select clause has Field1, Field2, Field3, Field4, and alias has only Field1, Field3, Field4, then for Field2 a default description is used.

3. Click the **Test Rule** button at the bottom of the screen.

   The Test Rule page is displayed.



Summarize

Summarize determines the type of summarization or aggregation for the rule.

Name	Config Value
Summarize	To query metas without any custom grouping, select: • **None**: The data is grouped by session in this case. To get collection (sessions/events/packets) related aggregates, select either of the following: • **EventCount**: The total count of events. • **Packet Count**: The total count of packets. • **Session size**: The total session size. To get meta based aggregates, select: • **Custom**: This indicates that expected meta aggregate function is defined in rule select clause.

Order By

Order By determines how to sort the result set.

Name	Configuration Value
Column Name	The **Column Name** is the name of the columns by which you want to sort the results. By default, the value is empty. When you click on a column, the value gets populated based on the Summarize field. • For 'None' and 'Custom', the value gets populated based on the entries made in the Select field. You can select from this list or add custom name. • For Event Count, Packet Count and Session size, accepted values are Total and Value. • Total - sort by aggregate value • Value - sort by group by meta
Sort By	**Sort By** determines the order in which you want to sort the results. The following are the values: • Ascending Order • Descending Order

Session Threshold

The session threshold is the optimization setting to stop scanning the matching sessions for each possible unique value for the selected meta. The threshold is an integer between 0 (default) and 2147483647. The threshold 0 scans for all matching sessions.

Note: If you provide a non-zero value (a value higher than zero), the aggregate results are inaccurate. This can be used only when you are interested in unique values and not aggregate values.

Supported where Clause

Syntax	Description
where <field1> [<field-operator>] < value1>,<value2>,<value3-value4> <logic-operator> <field2>,and so on	The where clause is a comma separated list of language field values and ranges that is used by NwValues function. In the where clause, string values have to be enclosed within single quotes. For example, where username = 'admin' && service = 22.
where <field1> [<field-operator>] <List1>	You can use a list in the where clause if you have multiple values to report on. For example, where ip.src exists && alias.host exists && alias.host contains $[User Reports/List of Alias Host]. When you use the list you must specify in the format $[<path>/<List name>].

In the where clause, make sure the syntax is correct based on the meta type.
For example,
For all text meta type use quotes for example, username = ‘user1’.
For all IP Addresses, Ethernet Addresses, and Numeric meta types do not use quotes for example, service = 80 && ip.src = 192.168.1.1.
For date and time meta types, if the date and time format is 'YYYY-MM-DD HH:MM:SS‘, use quotes.
If the date and time format is 1448034064 (number of seconds since EPOCH (Jan 1, 1970)), do no use quotes.

Note: If list is used in the rule, make sure that the list values are quoted or unquoted based on the type of the meta used. Checking the **Quotes will be inserted for all the values** checkbox in list definition page (for more information see, "Create Lists or List Groups" section in Configure a Rule) would quote all the list values.

Supported where Clause Operators

Syntax	Description
=	Returns results where the field is equal to any provided value. For example, tcp.dstport = 21-25,110 returns session with TCP destination ports of 21, 22, 23, 24, 25, or 110.
!=	Returns results for fields that do not match the values specified. For example, eth.type !=0x0800 returns sessions outside of hex value (decimal value of 2048) that is all non-IP based protocols.
begins	Checks for a value at the beginning of a text or binary field.
contains	Searches a text or binary value for a partial match.
ends	Checks for a value at the end of a text or binary field.
exists	If the field value exists, regardless of value, the operation evaluates to true.
!exists	If the field value does not exist, the operation evaluates to true.
length	Evaluates the length of the field. For example, username length 20-u returns any username that is 20 or more characters long.

Syntax	Description
regex	Performs a regular expression search against text or binary values.
not	Not operator is used to negate a clause or condition. For example, (not(user.dst ends "$")) will not display values for user destination.

Supported then Clause

Syntax	Description
then <rule action>	The then clause contains a rule action that manipulates the original result set of a rule in order to make the output in a report more concrete or add additional functionality other than querying data and displaying it. For example, dedup (filename).

Limit field

This indicates the limit to be put on the query while fetching data from the database. If a result set is sorted by event
count, packet count, or session size, the limit represents the top (or bottom) N values to be returned. If the result set is not
sorted, the first N values are returned.

Rule Actions

The NWDB data source rule syntax supports the following rule actions:

- dedup
- filter_on
- filter_out
- lookup_and_add
- max_threshold
- min_threshold
- regex
- sum_count
- sum_values
- show_whats_new

dedup (string field)

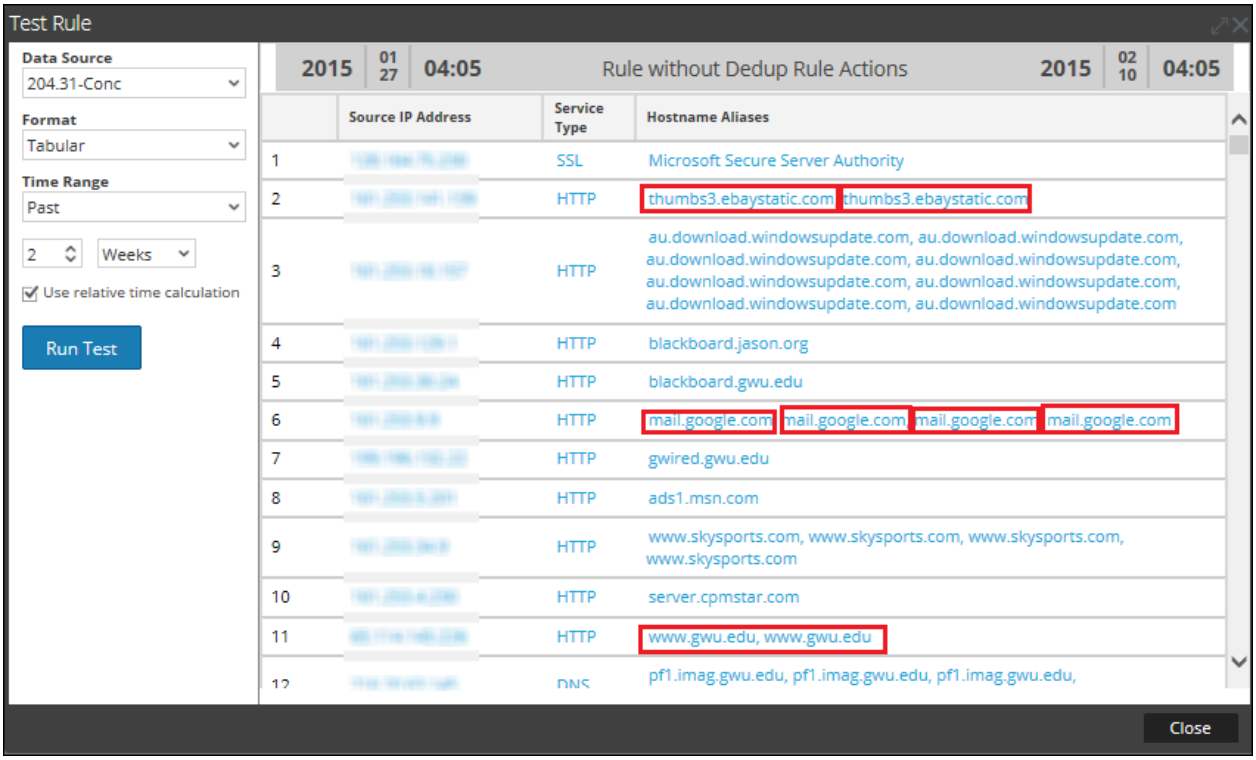
dedup removes the duplicate entries in an unsorted result set and displays only pertinent data.
The dedup rule action removes duplicate entries of a specific field in the report, so that only the first occurrence of that value is listed in the report.

Note: The dedup rule action cannot be used with an aggregate rule.

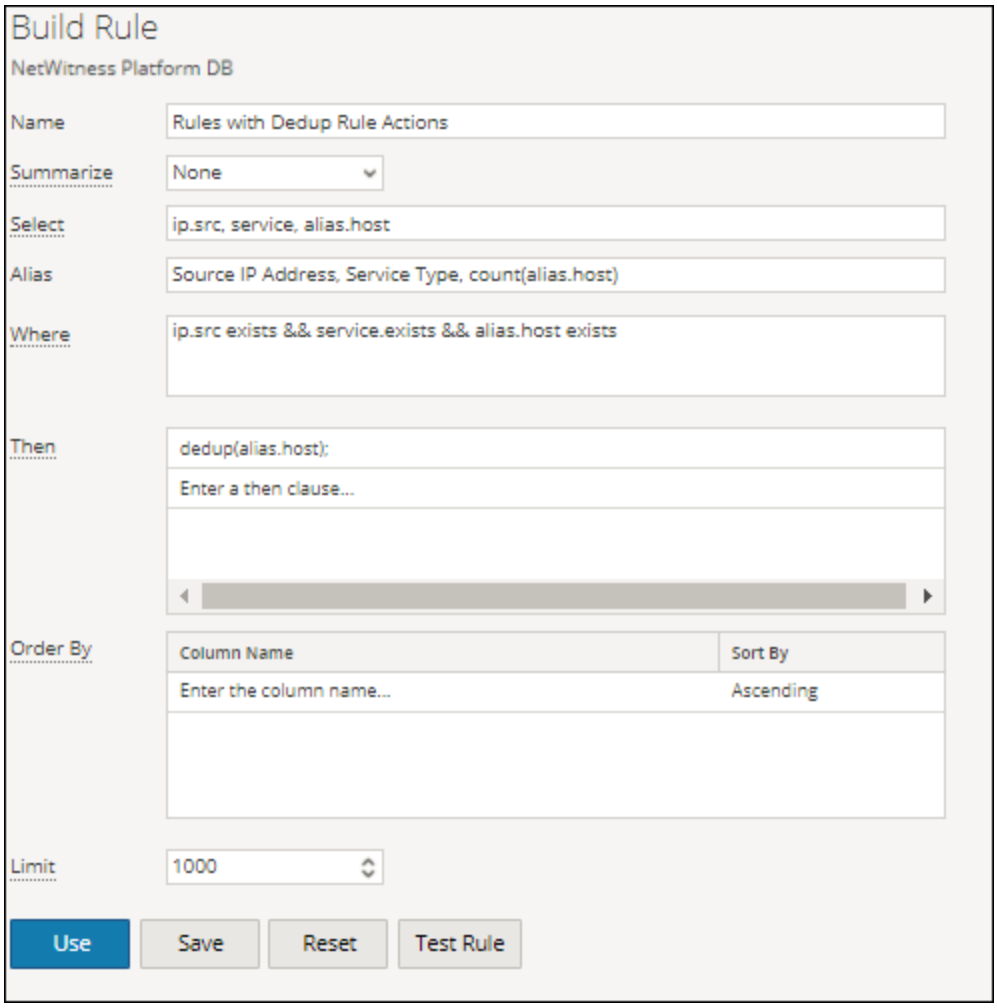
For example, the meta data generated by an individual session is often repetitive, especially when you have sessions with a lot of DNS lookups or web sessions that access the same host multiple times for various resources (such as, javascript, css). To remove the duplicate entries of the host, you can use the dedup rule action.

Example:

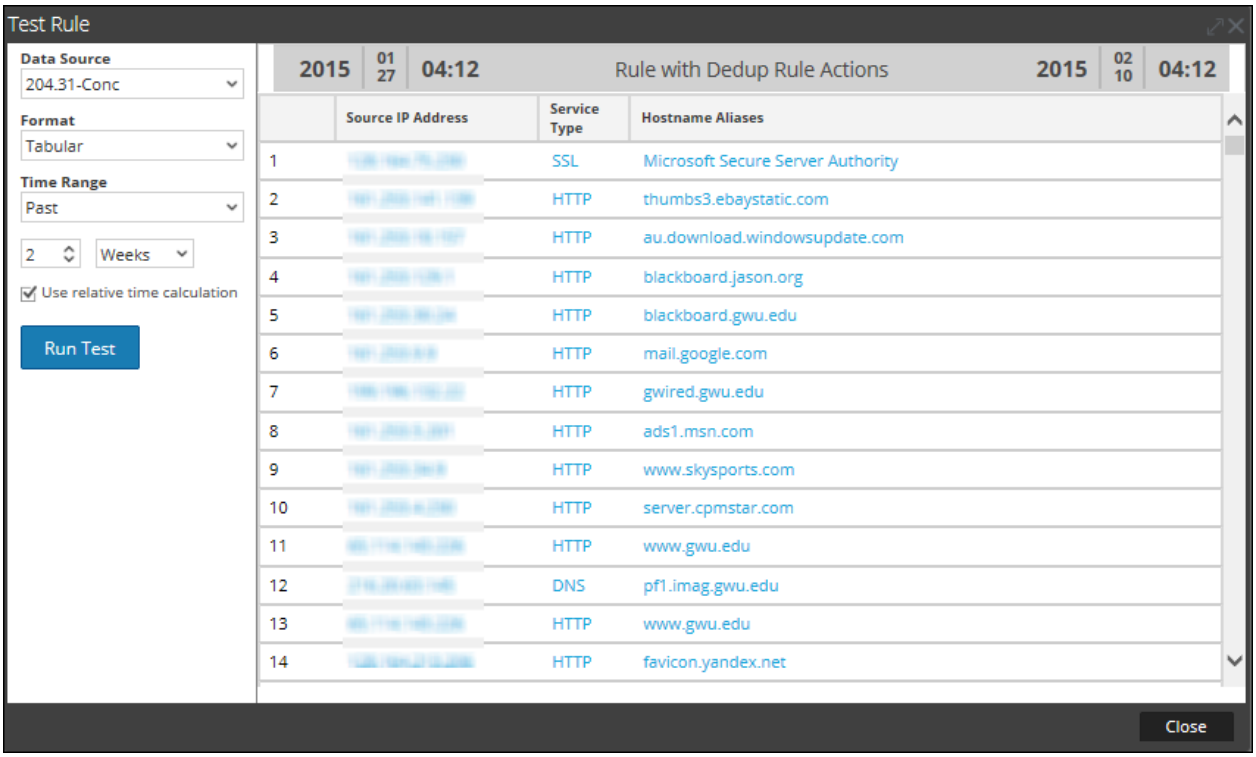
The following example is a lengthy result set that can be trimmed by removing the duplicate values in the same session.



The following figure shows the use of dedup rule action to remove the duplicate entries from the result set.



The duplicate value for each entry in the rule result set is reduced to one value.



filter_on (string filter, string field, bool matchExact)

`filter_on` removes values that do not contain the `filter` criteria from the result set. If the result set contains multiple fields, you must select a specific field to which the filter is applied. To add additional results to a single result set, include function such as lookup_and_add.

The matchExact parameter determines if the match is an exact match or contains a match.

- If matchExact is set to `false`, any value that contains the filter text is considered a match.
- If matchExact is set to `true`, only values that match the provided filter text is included in the result set.

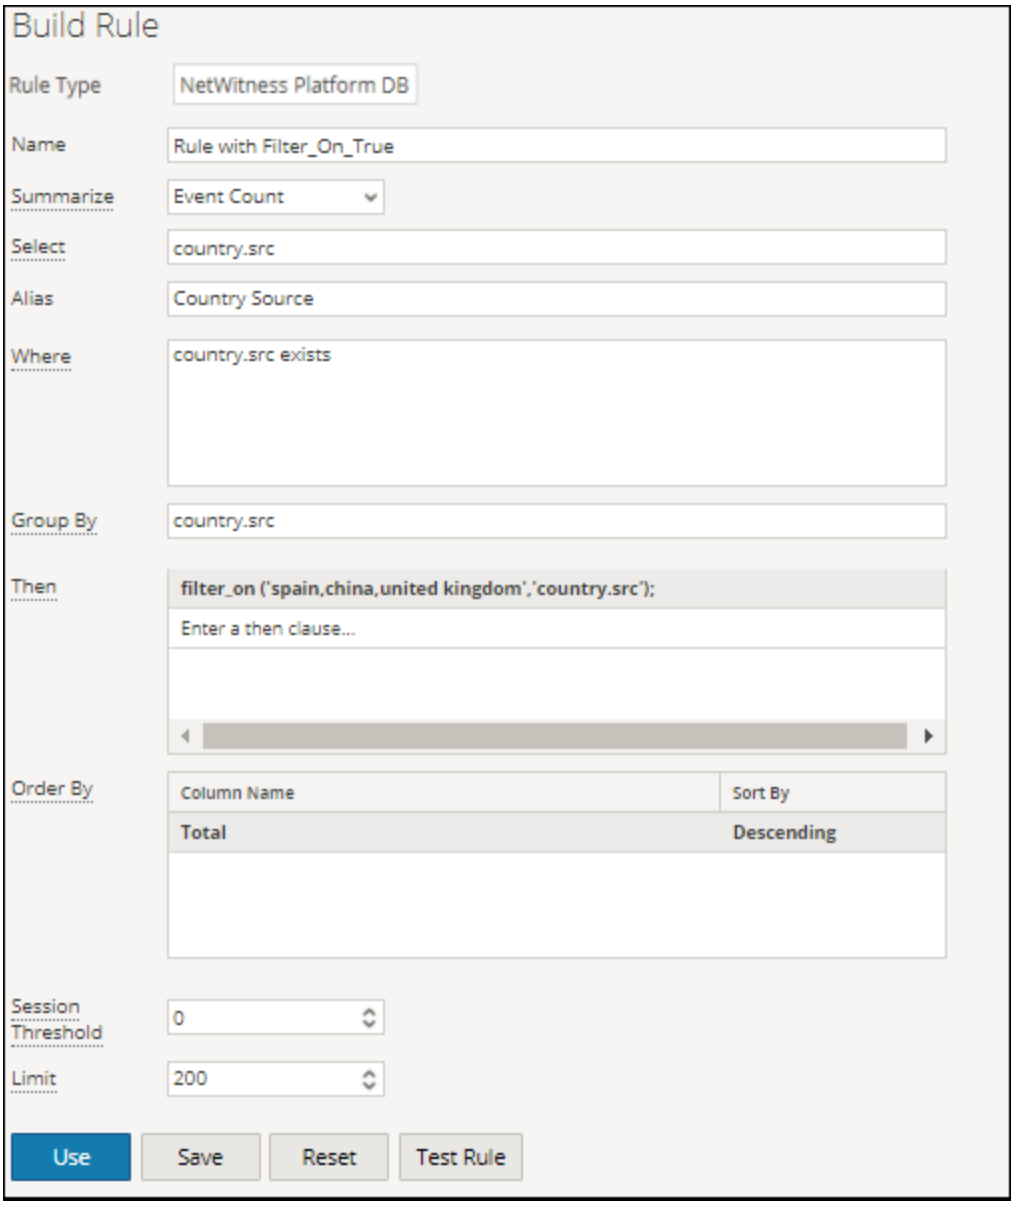
Note: Unless the matchExact parameter is specified, the default behavior of the rule action is to match exactly the text specified in the filter parameter. To specify that results containing the filter text must be kept in the result set, users must set the matchExact parameter to false.

Example:

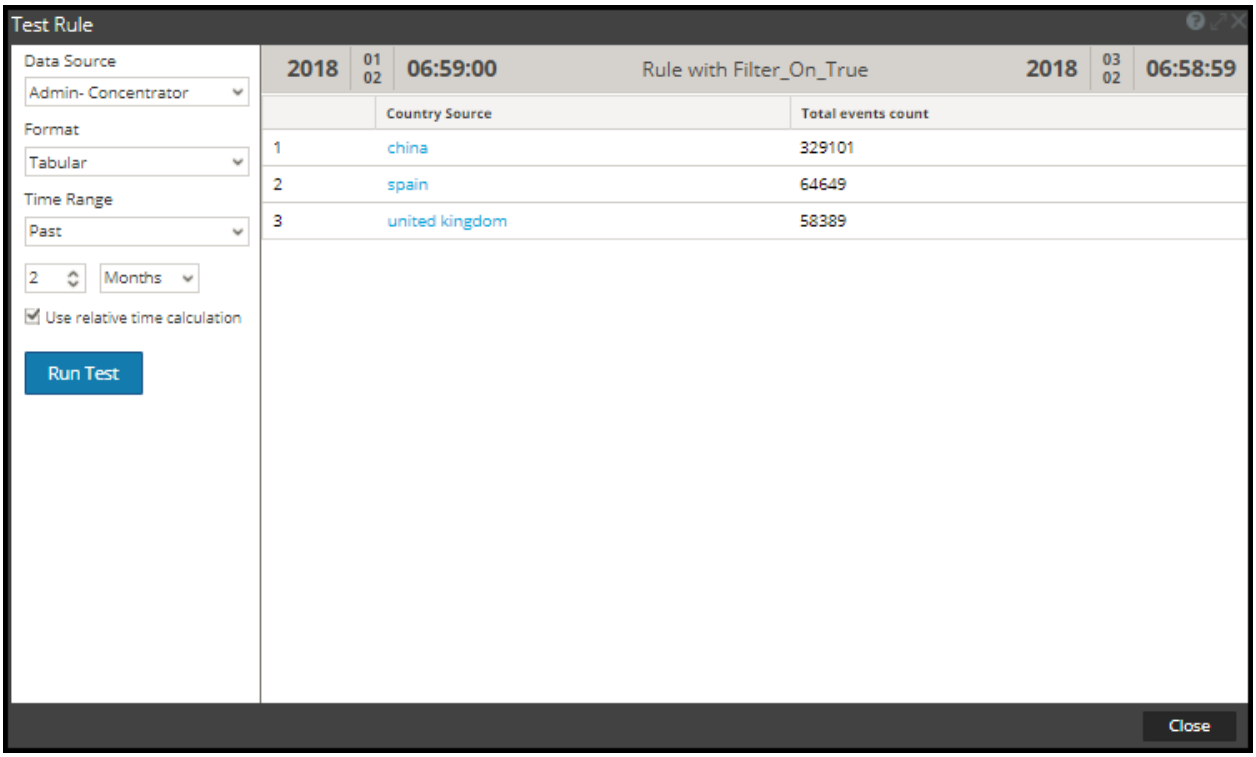
The following figure displays the list of countries and their event count.



The following figure shows a filter_on rule action to filter out countries except Spain, China, United States and United Kingdom from the result set.



The following figure shows the output with the filter_on rule action.



Another way of filtering out the entries from the result set is to create a list of variables which you want to filter out. For example, you can create a list with United Kingdom, France and Germany as values in the list. You can use this list in the rule action to get the same result set. For example, if you create a list called COUNTRY_LIST, you can use the list as follows:

```
filter_on ('$COUNTRY_LIST', 'country.src', 'false');
filter_out (string filter, string field)
filter_out (string filter, string field, bool matchExact)
```

filter_out removes the values that contain the *filter* criteria from the result set. If the result set contains multiple fields, you must select a specific field to which the filter is applied (for example, you can use a lookup_and_add to add results to a single result set).

The matchExact parameter determines if the match is an exact match or contains a match.

- If matchExact is set to false, any value that contains the filter text is considered a match.
- If matchExact is set to true, only values that match the provided filter text is excluded from the result set.

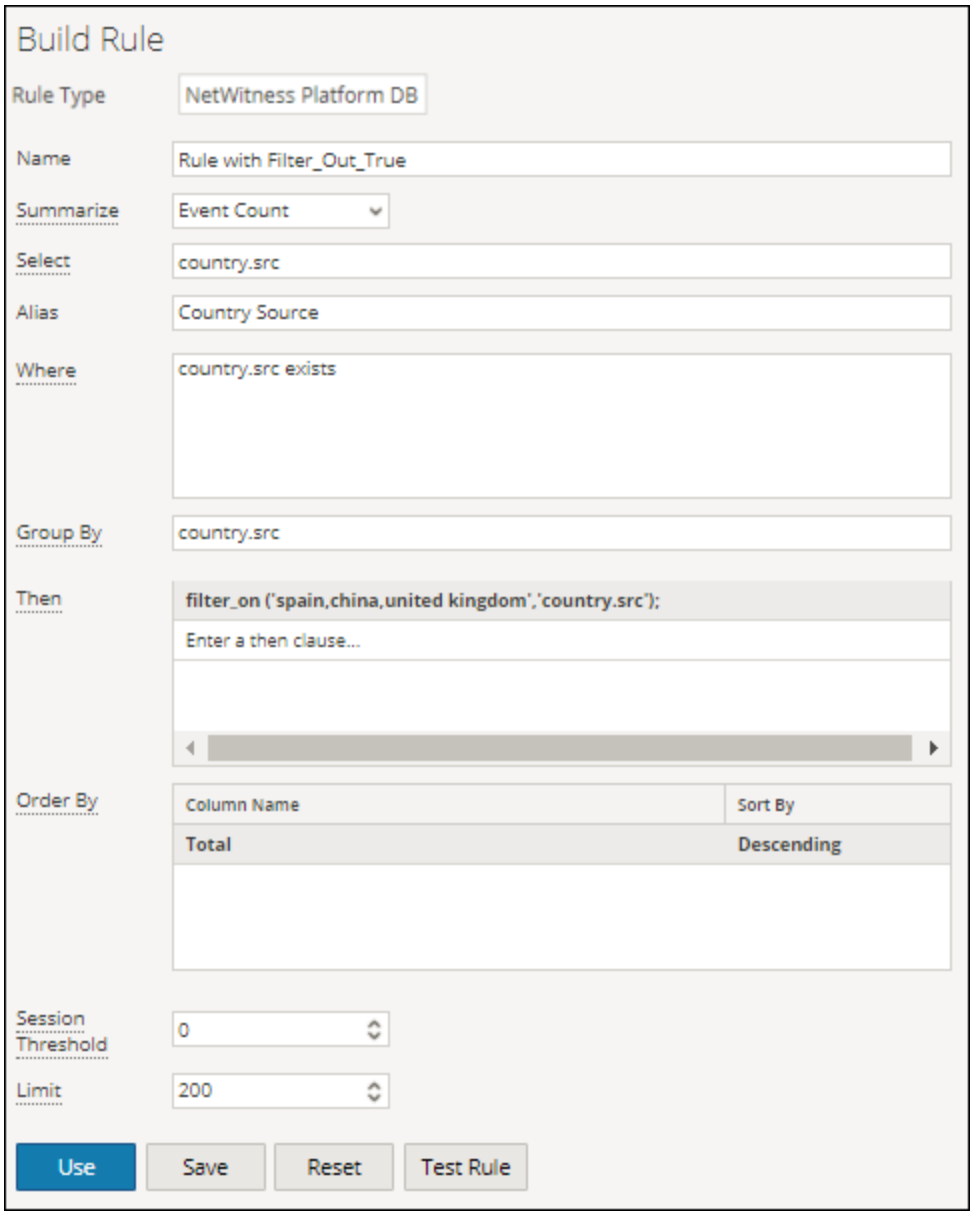
Note: Unless the `matchExact` parameter is specified, the default behavior of the rule action is to match exactly the text specified in the filter parameter. To specify that results containing the filter text must be removed from the result set, users must set the `matchExact` parameter to false.

Example:

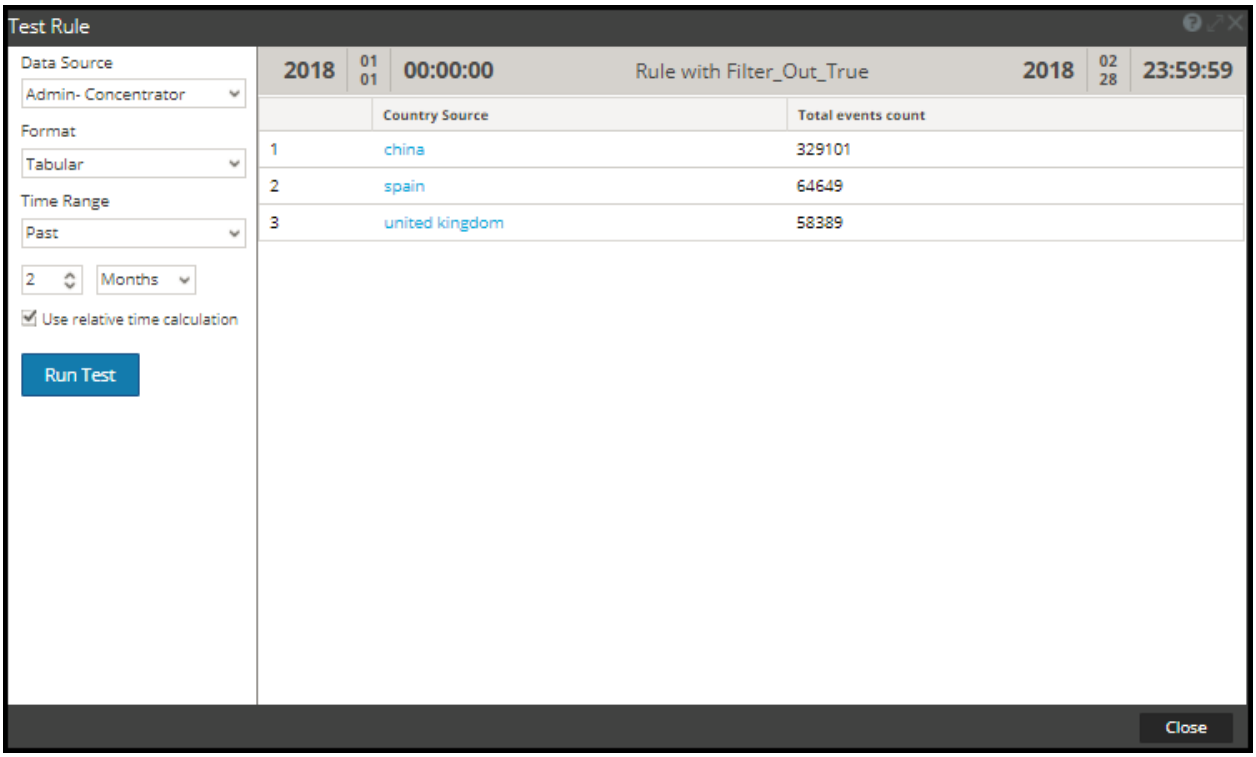
The following figure displays the list of countries and their event count.



The following figure shows the filter_out rule action to remove the event count for Spain, China, United States and United Kingdom from the result set.



The following figure shows the output with the filter_out rule action.




```
lookup_and_add (string select, string field)
lookup_and_add (string select, string field, int limit)
lookup_and_add (string select, string field, int limit, boolean inherit)
lookup_and_add (string select, string field, int limit, boolean inherit, string extraWhere)
lookup_and_add(string select, string field, int limit, boolean inherit, string extraWhere, boolean aggregate)
```

This rule action iterates through a list of values in a result set and lookup additional meta data to further describe the relationships between various elements in a result set.

Note: The lookup_and_add rule action can be used only with an aggregate rule.

The first parameter, select, designates the type of meta data that must be added to elements of the result set. The second parameter, field, specifies where in the result set the append must apply to. Also, a limit can be applied to avoid crowding the result set with a large result set.

By default, subsequent queries to the SDK will inherit the where clause of the parent rule. To use a unique where clause, you can specify a boolean value in the fourth parameter as false and in the fifth parameter you can specify a different where clause.

Note: If you are using a unique where clause in your query, make sure that you use a single quote (') for enclosing arguments and double quotes (") for string values.

Now, with the addition of **Custom** summarization and **Group By** feature, the result can be achieved even without having lookup_and_add rule action. The new rule syntax with groupby displays the result in a flat structure which is better than the earlier rule syntax without groupby Hence it is recommended to manually edit/update rules with lookup_and_add rule action and use groupby clause wherever it is applicable.

Note: Lookup_And_Add rule action is supported only if the SELECT clause has one meta and aggregate function.

For example, see below scenarios: In Example **2a**, lookup_and_add rule action is used. Instead of using lookup_and_add rule action, the same result can be achieved by using **Custom** summarization and **Group By** feature. See Example **2b** below.

But, lookup_and_add rule action is still supported for NWDB rules on the following conditions:

- All versions of NWDB rules with Summarization as Event Count, Packet Count, or Session Size.
- For Custom summarization, the lookup_and_add rule must have only one group by meta with only one aggregate function where the aggregate function must be either sum() or count().

Note: It is not supported for “Summarize-None”.

For example, lookup_and_add rule action can be used for the following rules:

- select ip.src, sum(size) group by ip.src
- select ip.src, count(filename) group by ip.src

It cannot be used for the following rules:

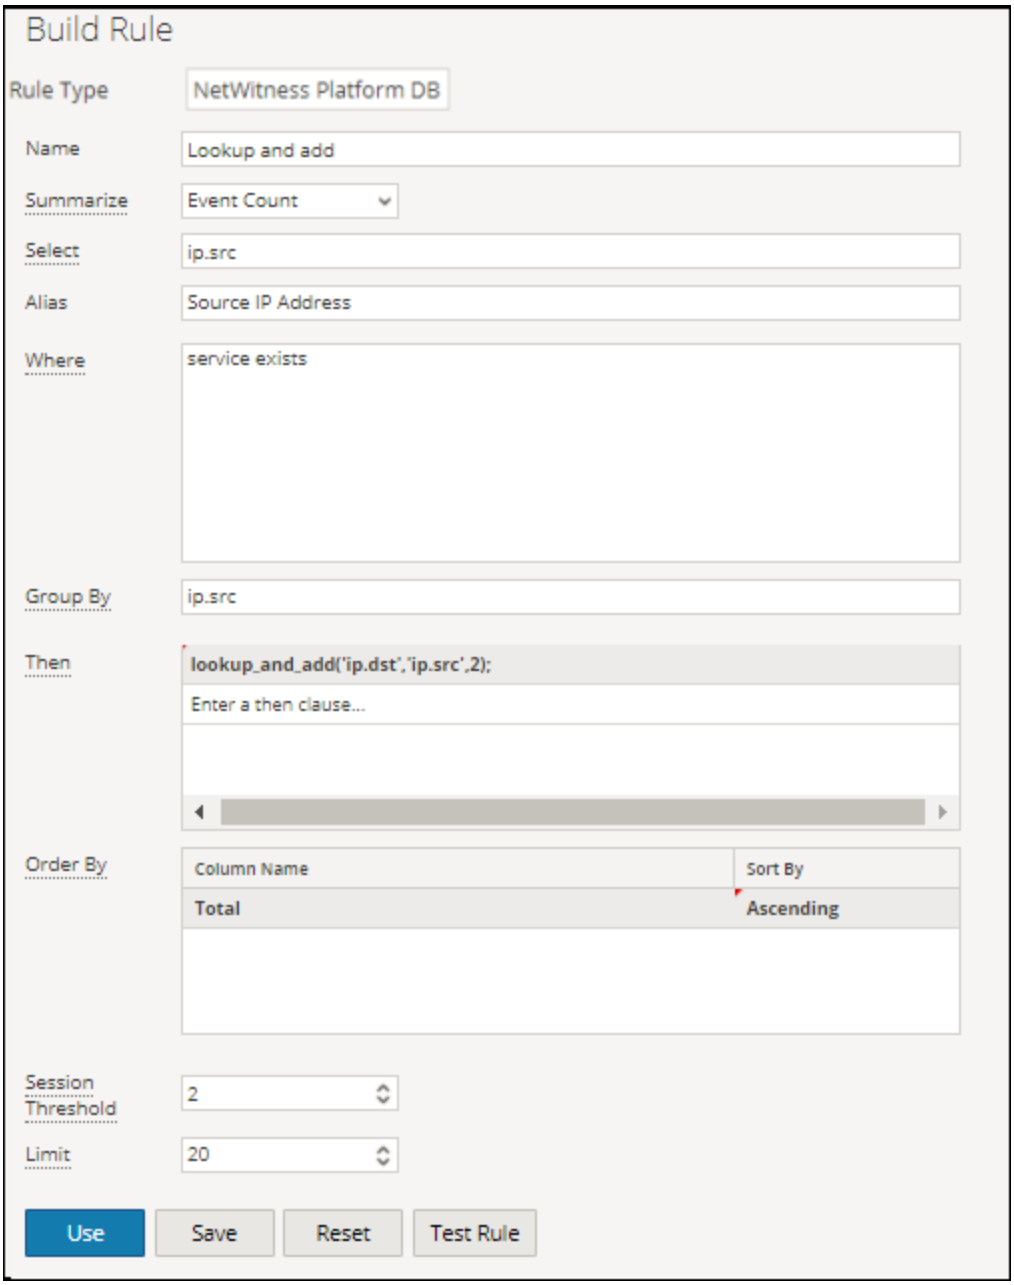
- select ip.src, sum(size),count(filename) group by ip.src
- select ip.src, sum(size),avg(size) group by ip.src
- select ip.src,ip.dst count(filename) group by ip.src,ip.dst

Examples:

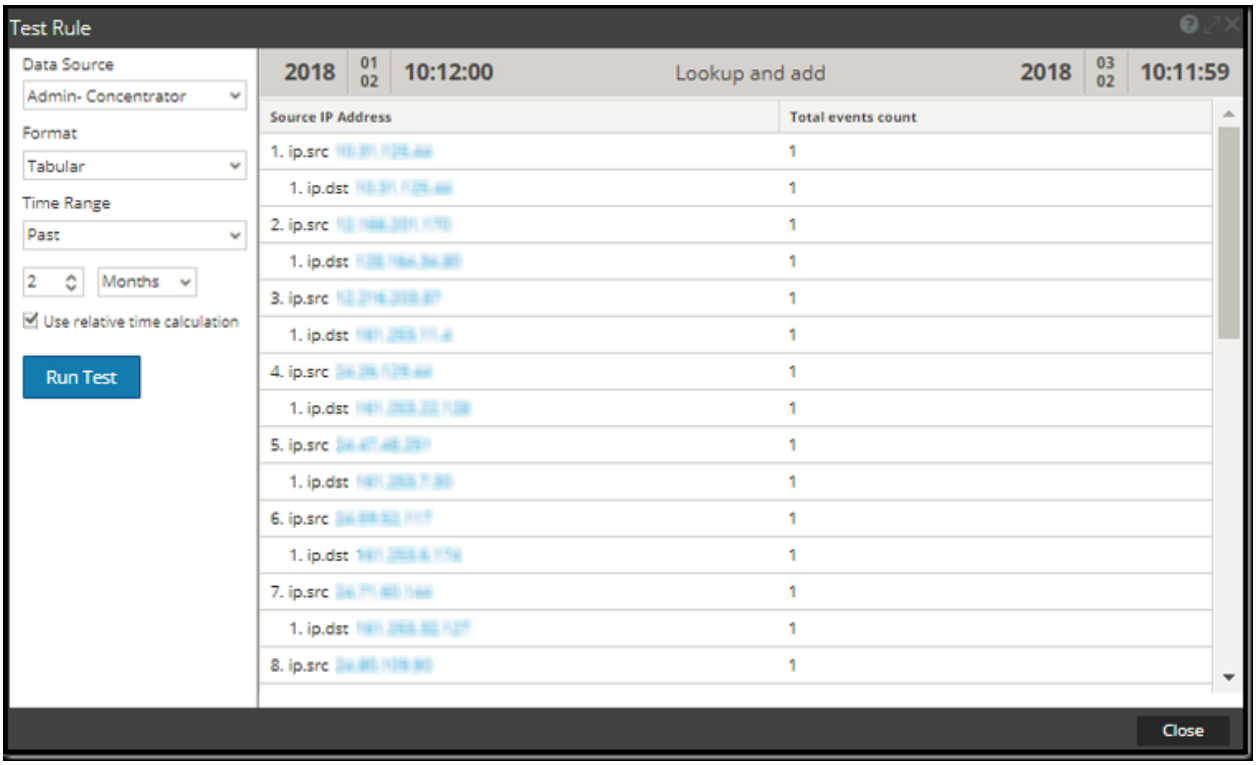
1. lookup_and_add('ip.dst','ip.src', 2);

This rule action would iterate through each ip.src in the initial result set and lookup the top two destination IP addresses with each ip.src.

The following figure shows the rule definition.



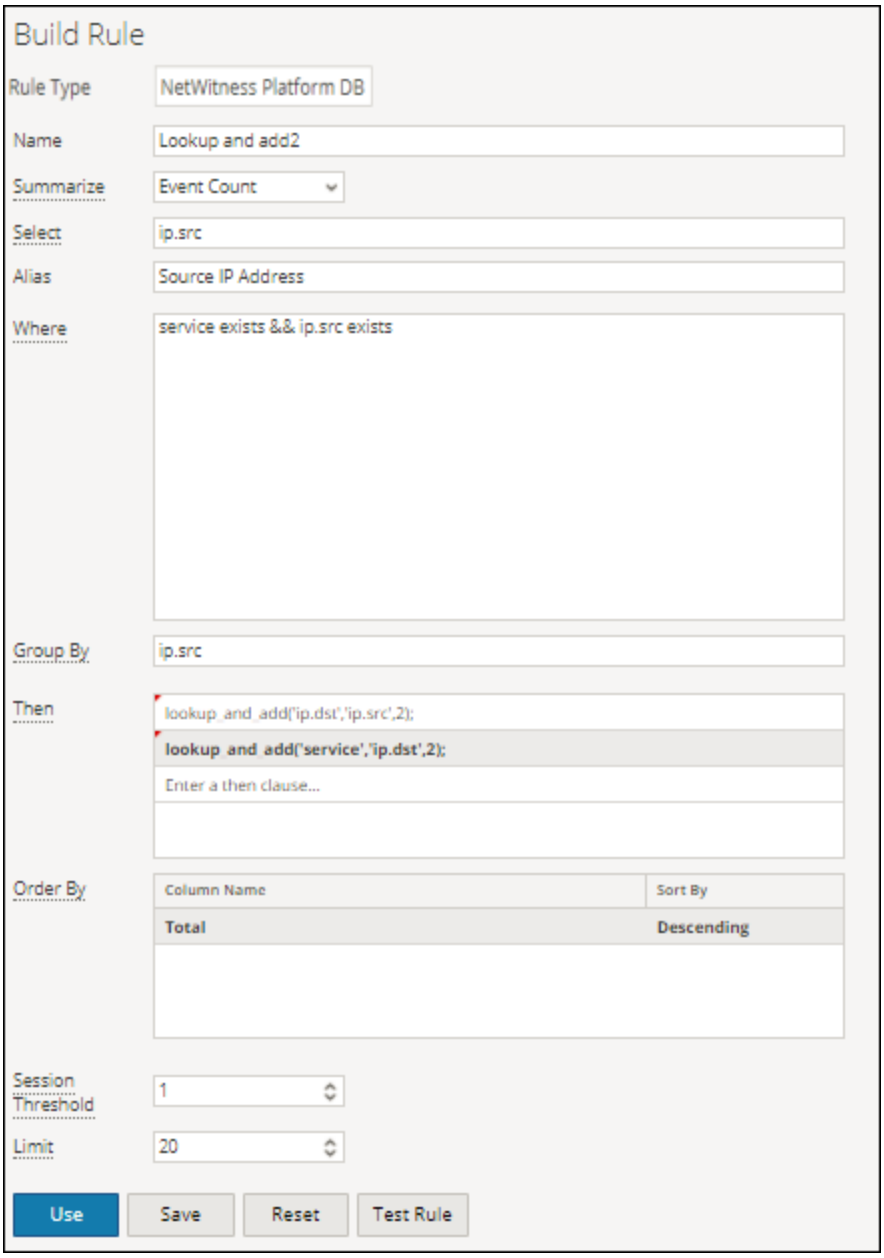
The following figure shows the result set containing the source IP addresses and the top two destination IP addresses with each ip.src.



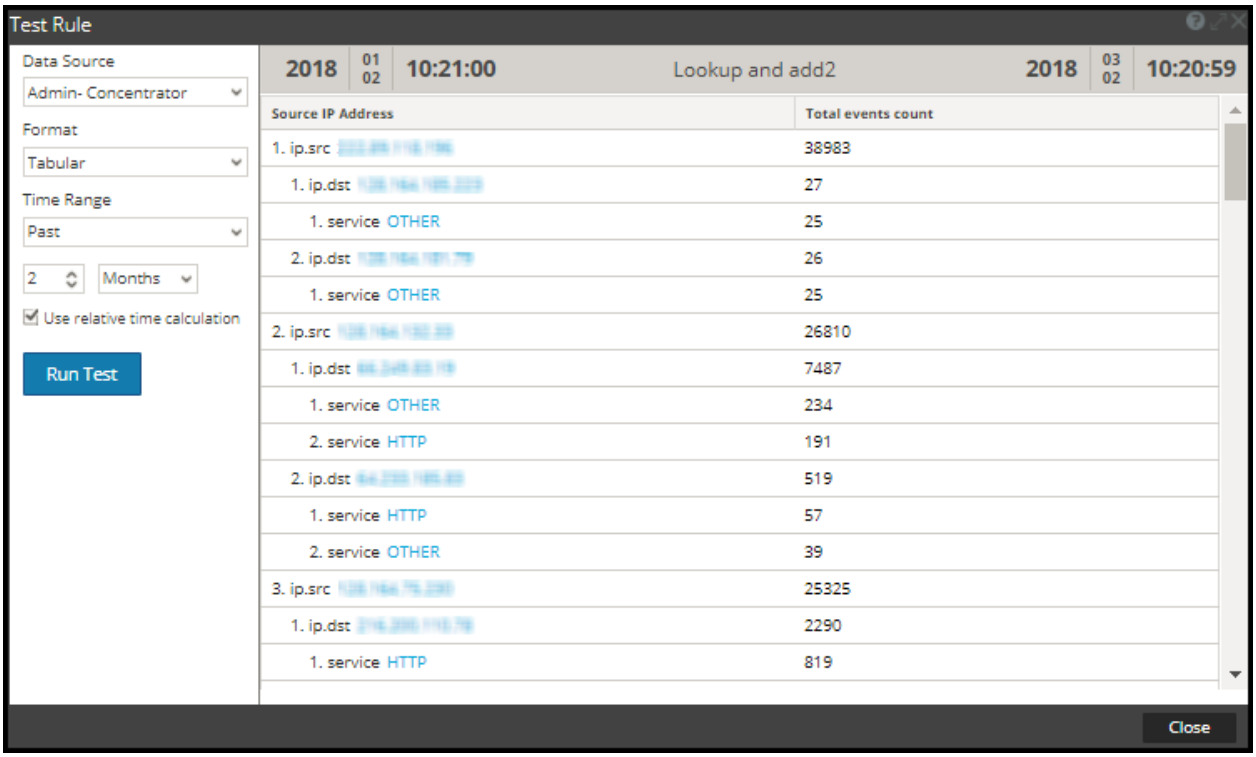
2a. lookup_and_add('ip.dst','ip.src', 2); lookup_and_add('service','ip.src', 3);

This rule action would iterate through each ip.src in the initial result set and lookup the top two destination IP addresses with each ip.src and the top three ports used by each ip.src.

The following figure shows the rule definition.



The following figure shows the result set containing the source IP addresses and the top two destination IP addresses with each ip.src and the top three ports used by each ip.src.

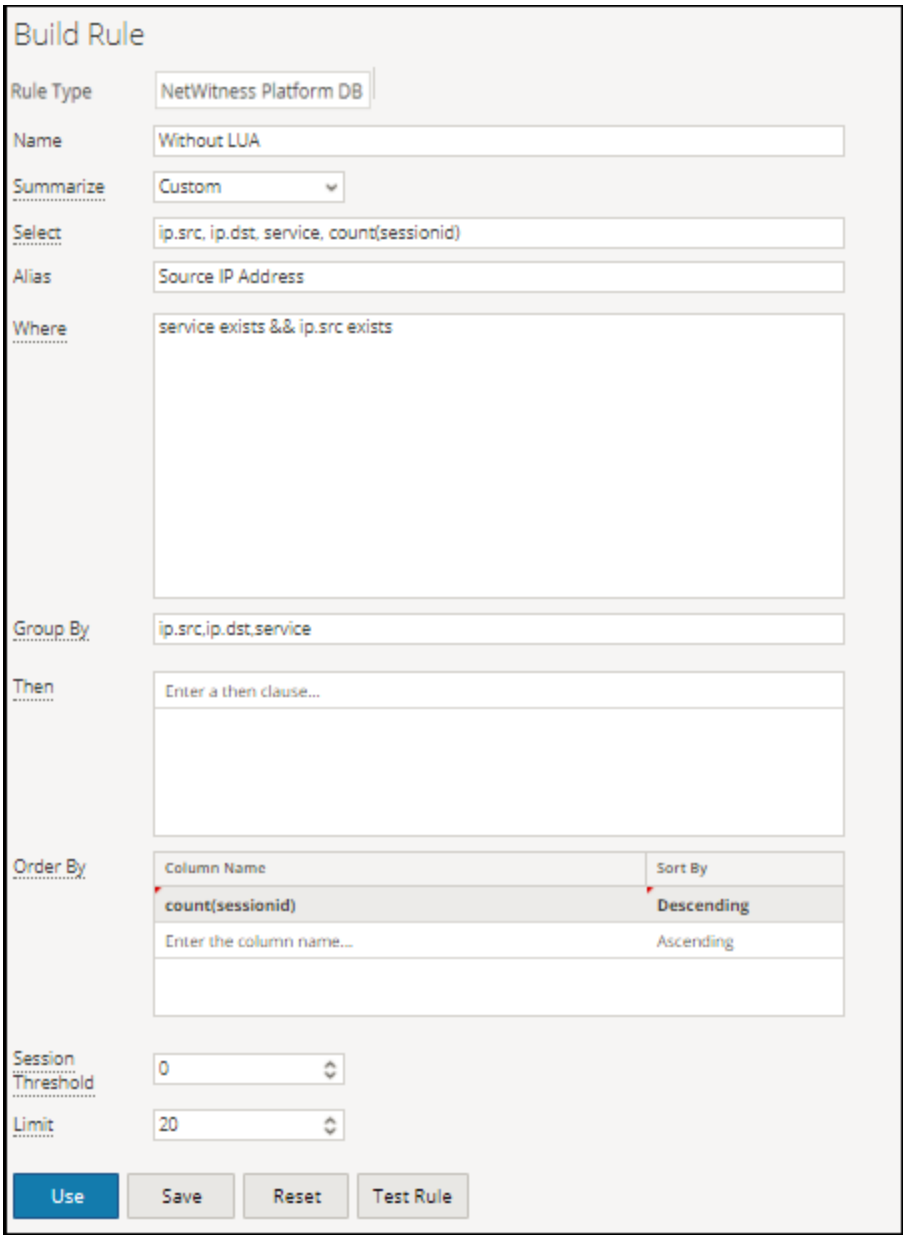


You can make the query as complex as you want by selecting different fields in the result set and by appending to different parts. For example, you may want to know what files each source IP had touched. However, because the parent rule has a WHERE clause of service = 6667 and the default behavior of this rule action is to append to the original WHERE clause, it becomes necessary to override the parent WHERE clause. The easiest way to understand this concept is to look at the previous lookup_ and_add call lookup_and_add('ip.dst','ip.src',2). The actual query that is sent to the server is SELECT ip.dst WHERE service = 6667 &&ip.src = 206.42.199.194. In order to force the WHERE clause to override the service = 6667 portion of the WHERE clause (inherited from the parent rule), the user can specify a 4th parameter of false as shown in example 3.

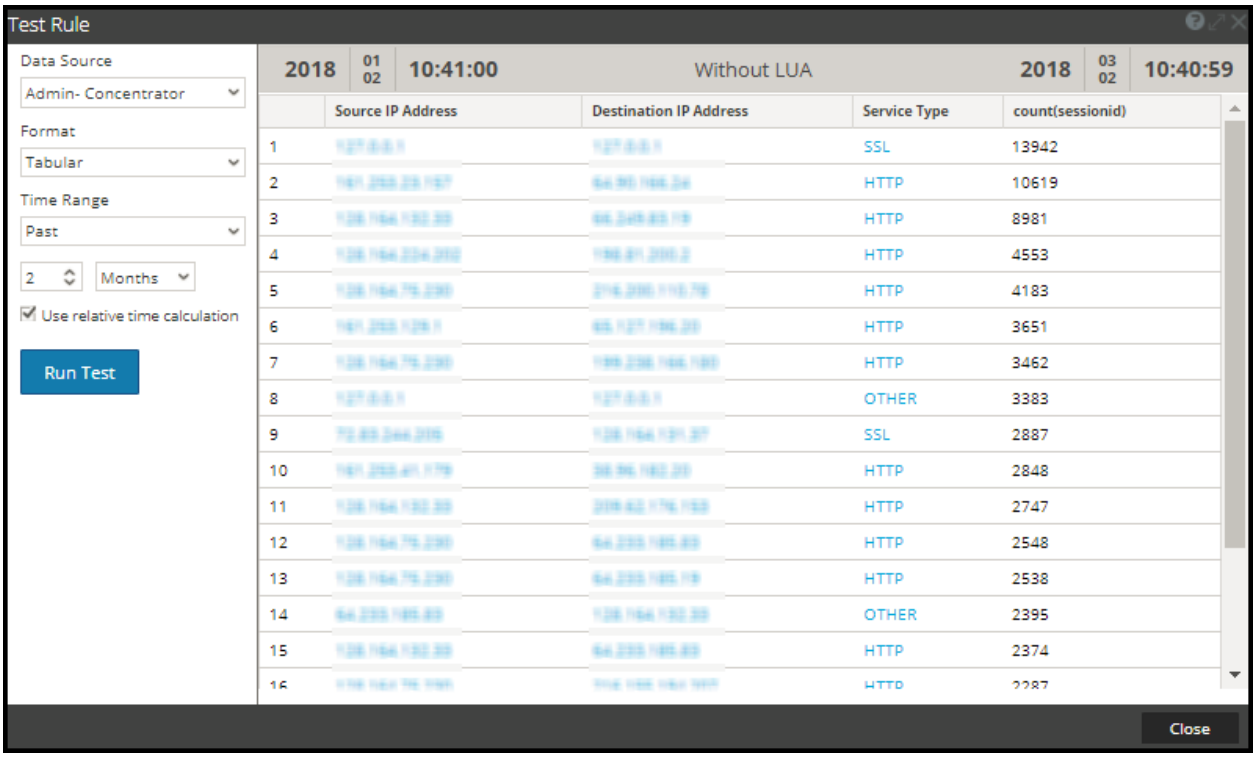
2b. Without Lookup_and_add Rule

This rule uses the Custom summarization and Group By feature to sort the results.

The following figure shows the rule definition.



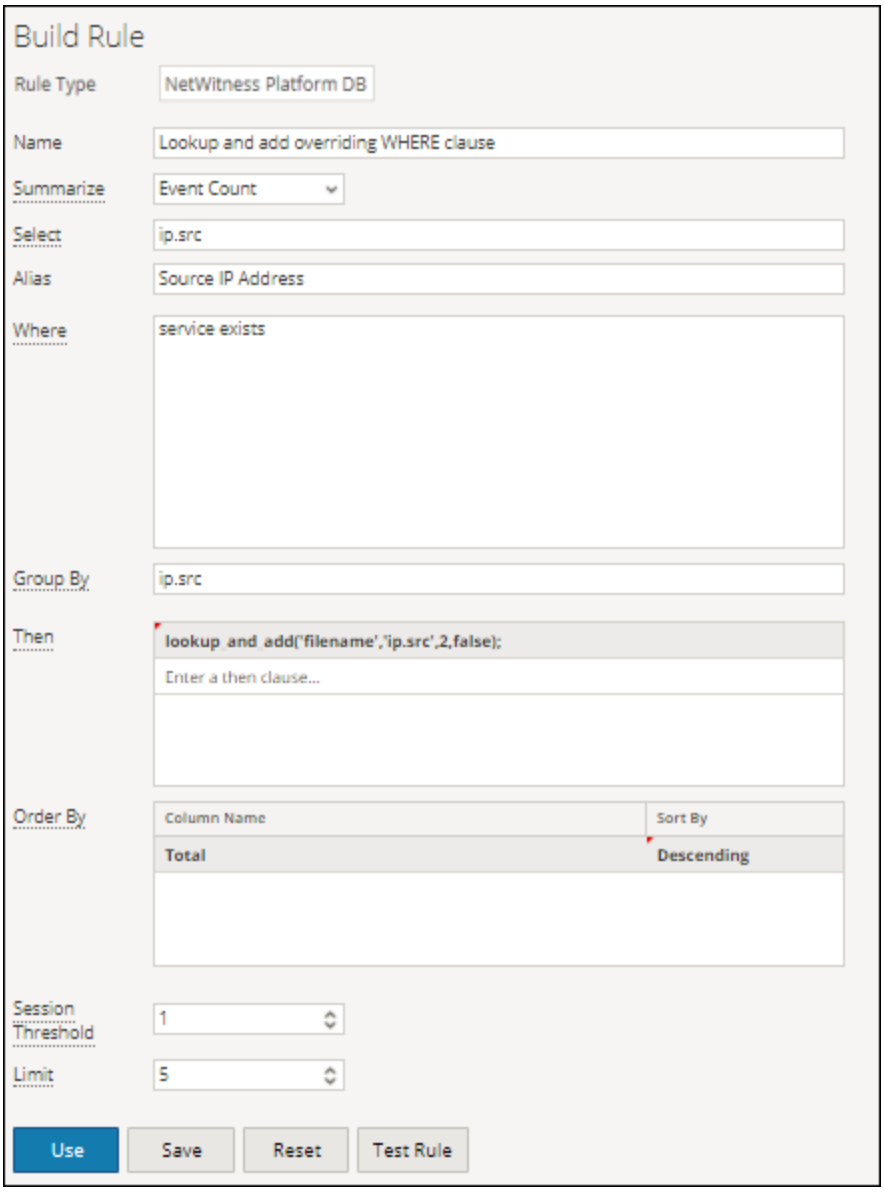
The following figure shows the result set containing the source IP addresses and the top two destination IP addresses with each ip.src and the top three ports used by each ip.src.



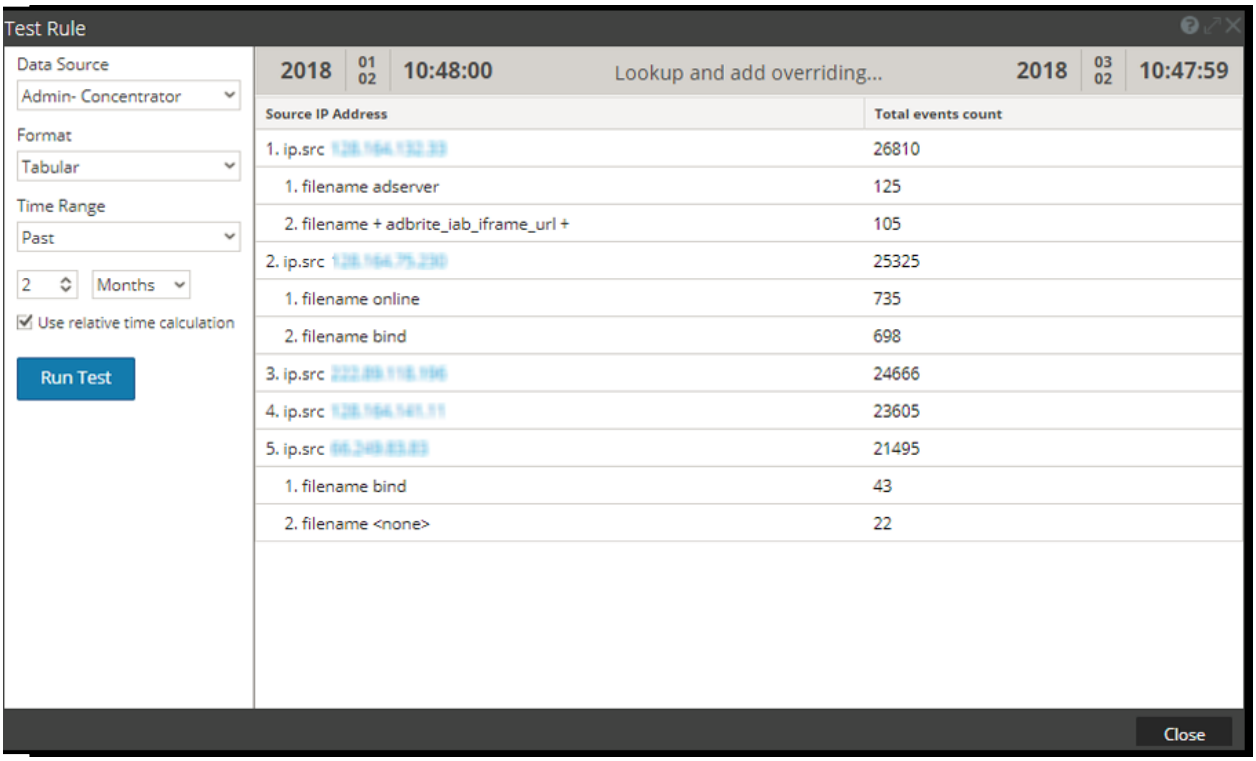
3. lookup_and_add('filename', 'ip.src', 2, false);

This call would issue a query to the server, like `SELECT filename WHERE ip.src = 90.0.0.142` rather than `SELECT filename WHERE service = 6667' && ip.src = 90.0.0.142` because you have specified the rule action to ignore the initial `WHERE` clause of the parent rule.

The following figure shows the rule definition.



The following figure shows the result set.

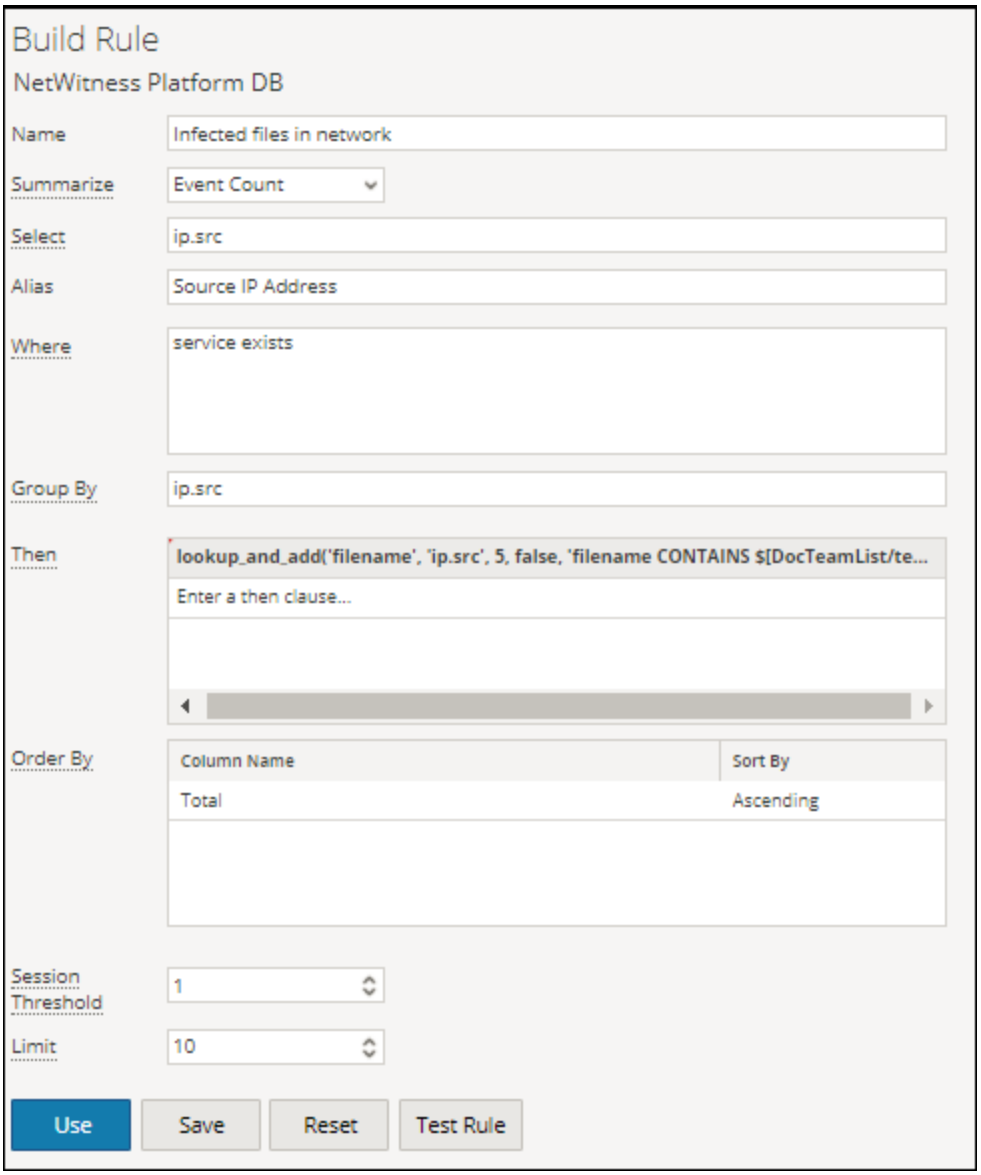


The test list is in a group name netwitness, you can access that list with the following syntax.

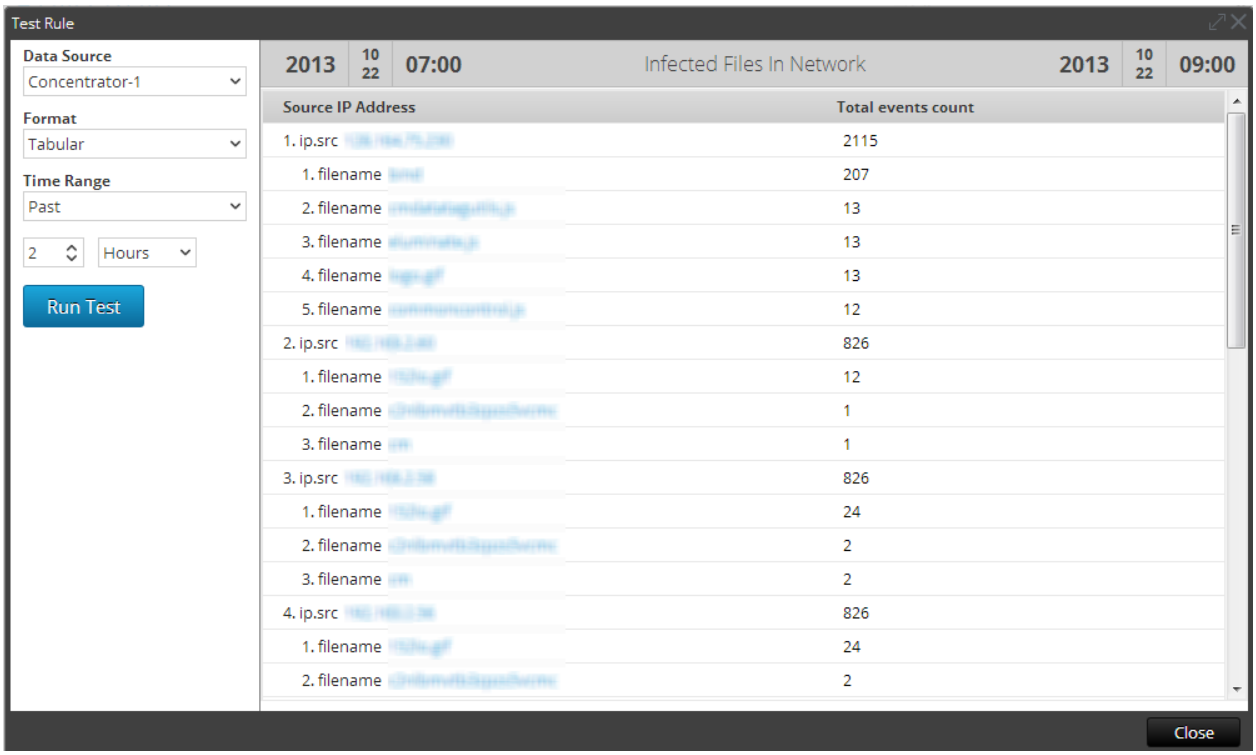
You can even narrow down these appended results even further to only include filenames that have .gif as filename extension by using the fifth parameter in the rule action. The fifth parameter allows you to specify additional WHERE clause criteria. The files with .gif filename extension would be stored in the **test** list within a group named **DocTeamList**. You can access this list with the following syntax: threat.source = $[DocTeamList/test]
This can be referenced in the extra where clause parameter in the following manner:

4. lookup_and_add('filename', 'ip.src', 5, false, 'filename CONTAINS $[DocTeamList/test]');

The following figure shows the rule definition.



The following figure shows the result set.

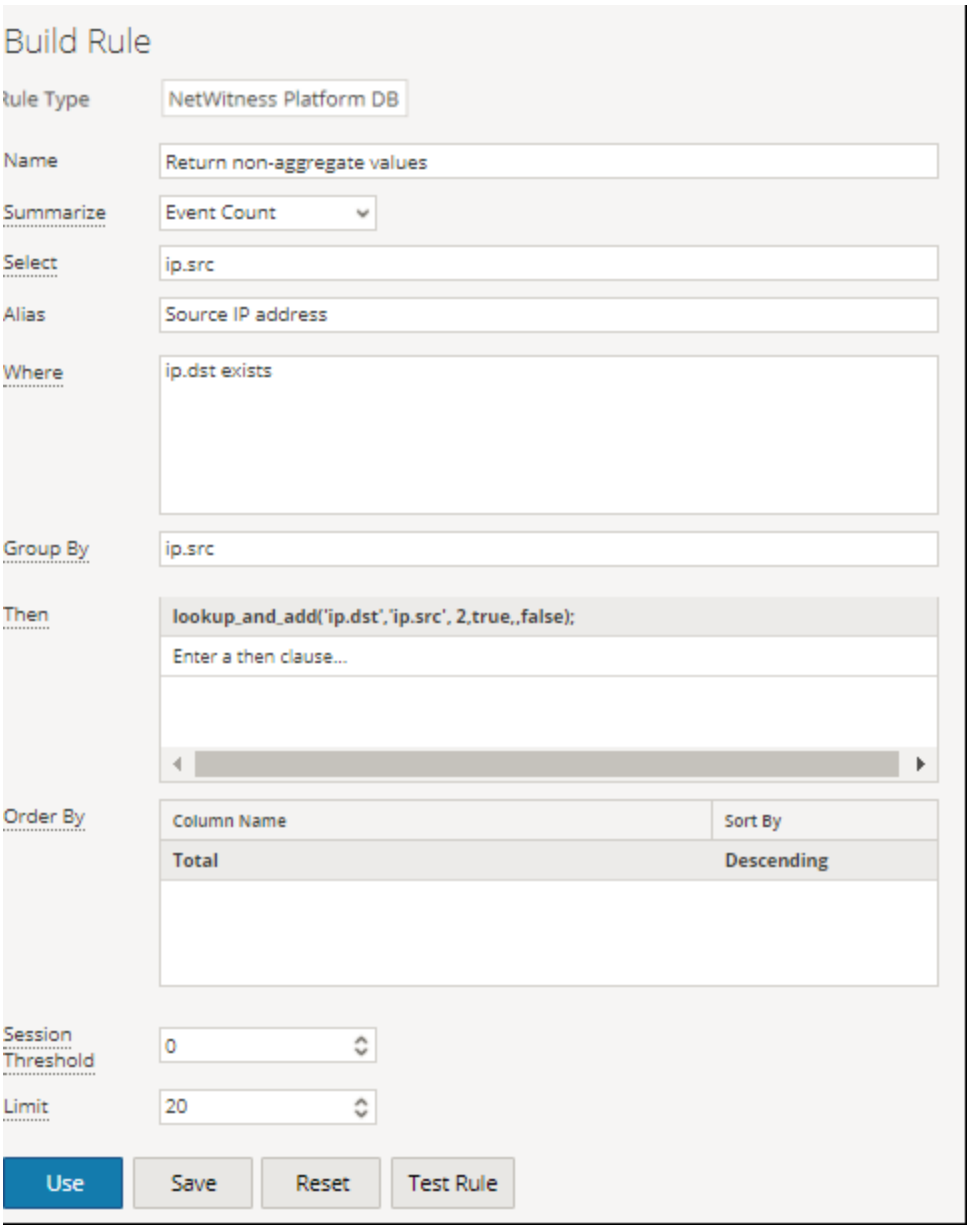


5. lookup_and_add('ip.dst','ip.src', 2,true,,false);

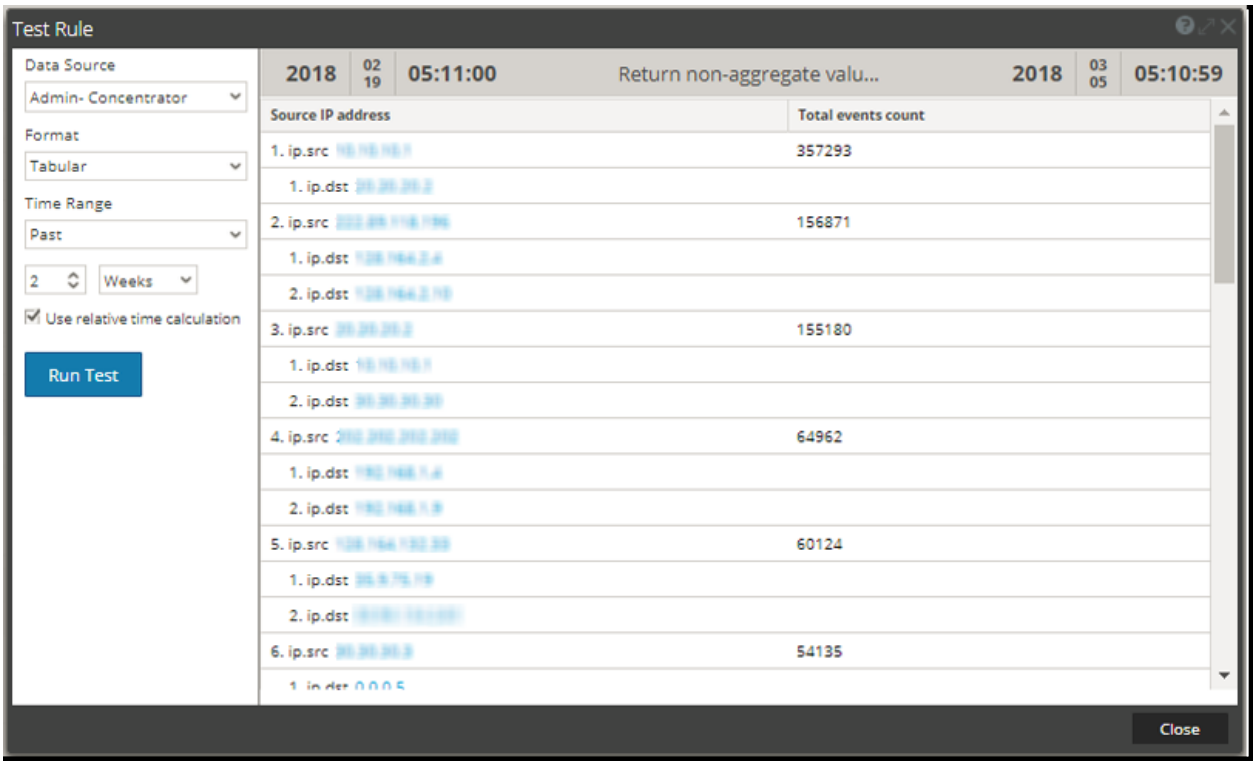
This rule action would iterate through each ip.src in the initial result set and lookup the top two destination IP addresses with each ip.src. The 'aggregate' parameter is set to 'false', this implies that aggregates would be skipped for lookup values and hence the lookup query executions will complete faster.

Note:
The default value for 'aggregate' is 'true'. When 'aggregate' is set to 'false', Reporting Engine passes threshold=1, Sort by='value' and Order=Ascending to NWDB to make lookup queries run faster.
. You must set the 'aggregate' to false, when rule contains aggregate functions or when the rule is run against a wide time range. This helps the rule to complete the execution faster.

The following figure shows the rule definition.



The following figure shows the result set.




```
max_threshold (string quantity)
max_threshold (string quantity, string field)
```

max_threshold removes any results with a quantity that is larger than the maximum threshold quantity from a result set. The quantity can either be in terms of count or size and it is relative to the sorting options of the parent rule. This means that if you sort a rule by size, the rule action expects you to specify the parameter in bytes (you can append KB, MB, GB, TB to the parameter to make size conversion easier).

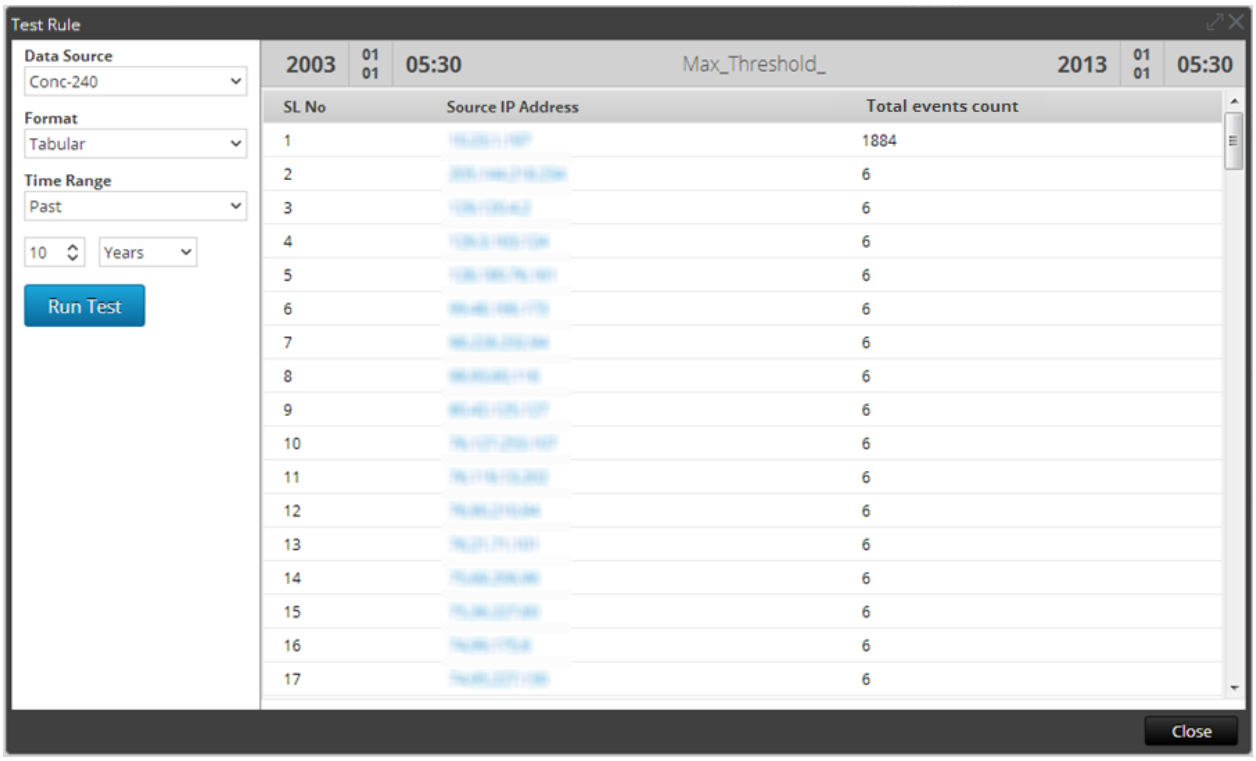
max_threshold rule can also be used to filter values based on the aggregate function values. Use the syntax based on the type of summarization used in the rule as below:

- max_threshold(String quantity): Can be used to filter Event Count, Packet Count, and Session Size.
- max_threshold(String quantity, String field): Can be used to filter values of Custom aggregates or any metas.

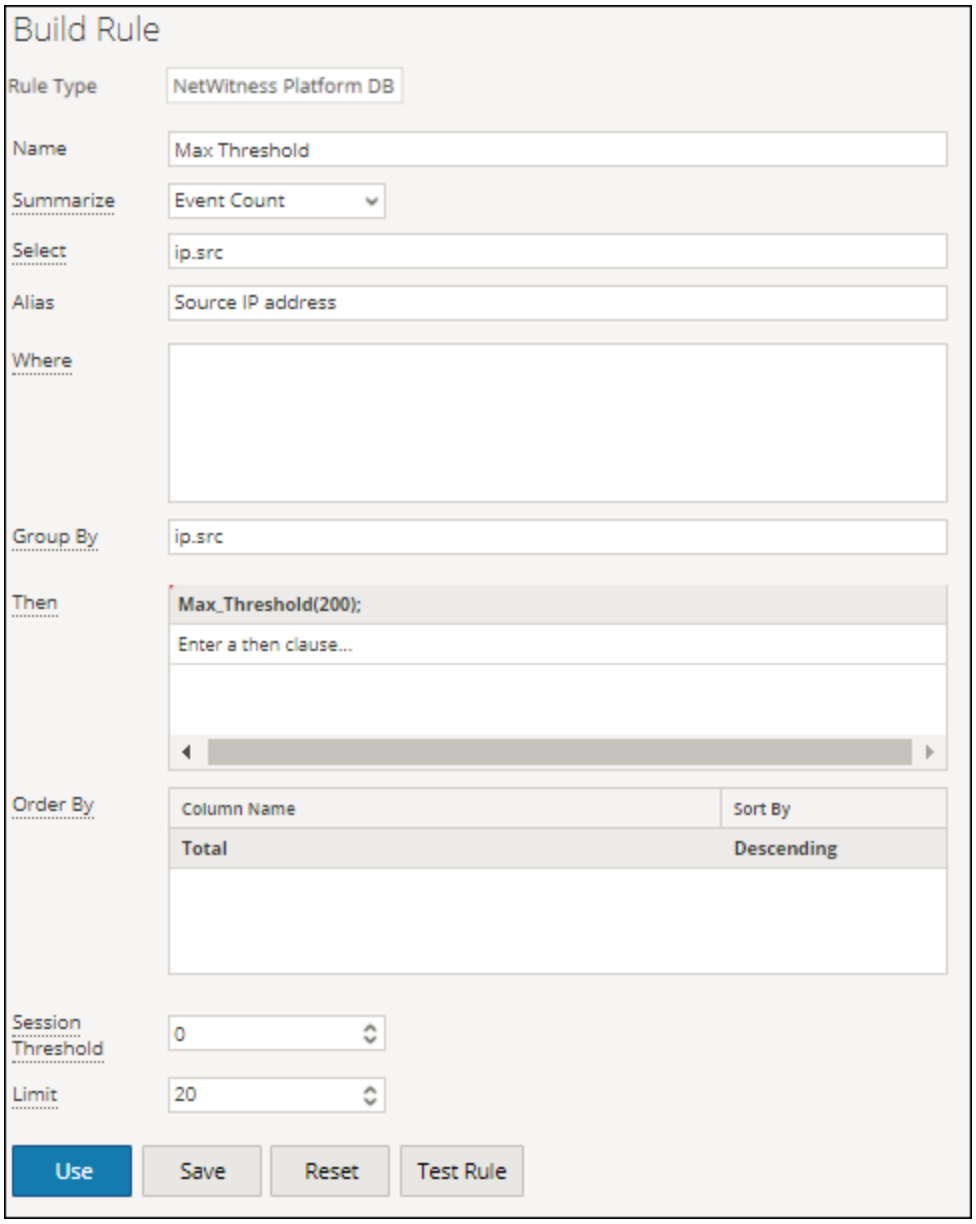
Examples:

1. max_threshold(200);

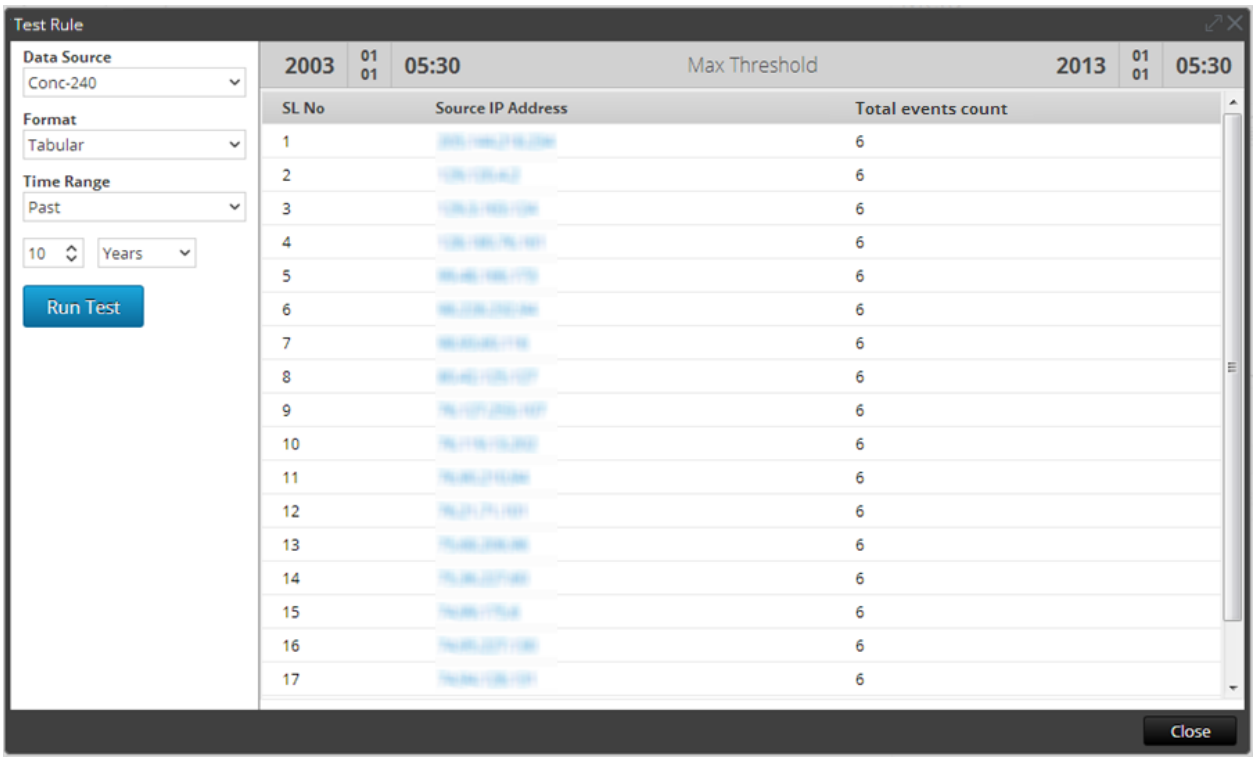
The following figure shows the result without the max_threshold argument. The output results have event counts exceeding 200.



The following figure shows a the max_threshold rule action that puts a limit of 200 bytes on the output. Any output having more than 200 bytes of data are not listed.

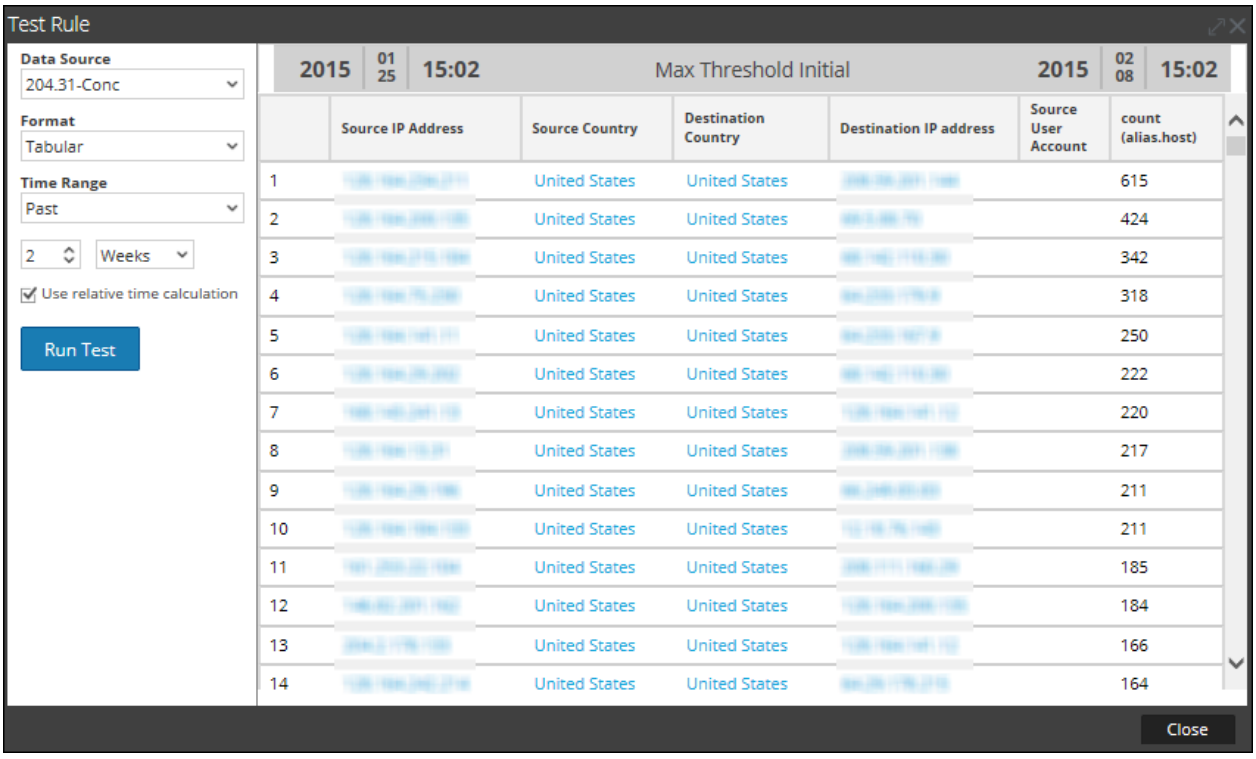


The following figure shows the result when the max_threshold rule action is applied. The result numbered 1 in the above screen capture is removed from the result.

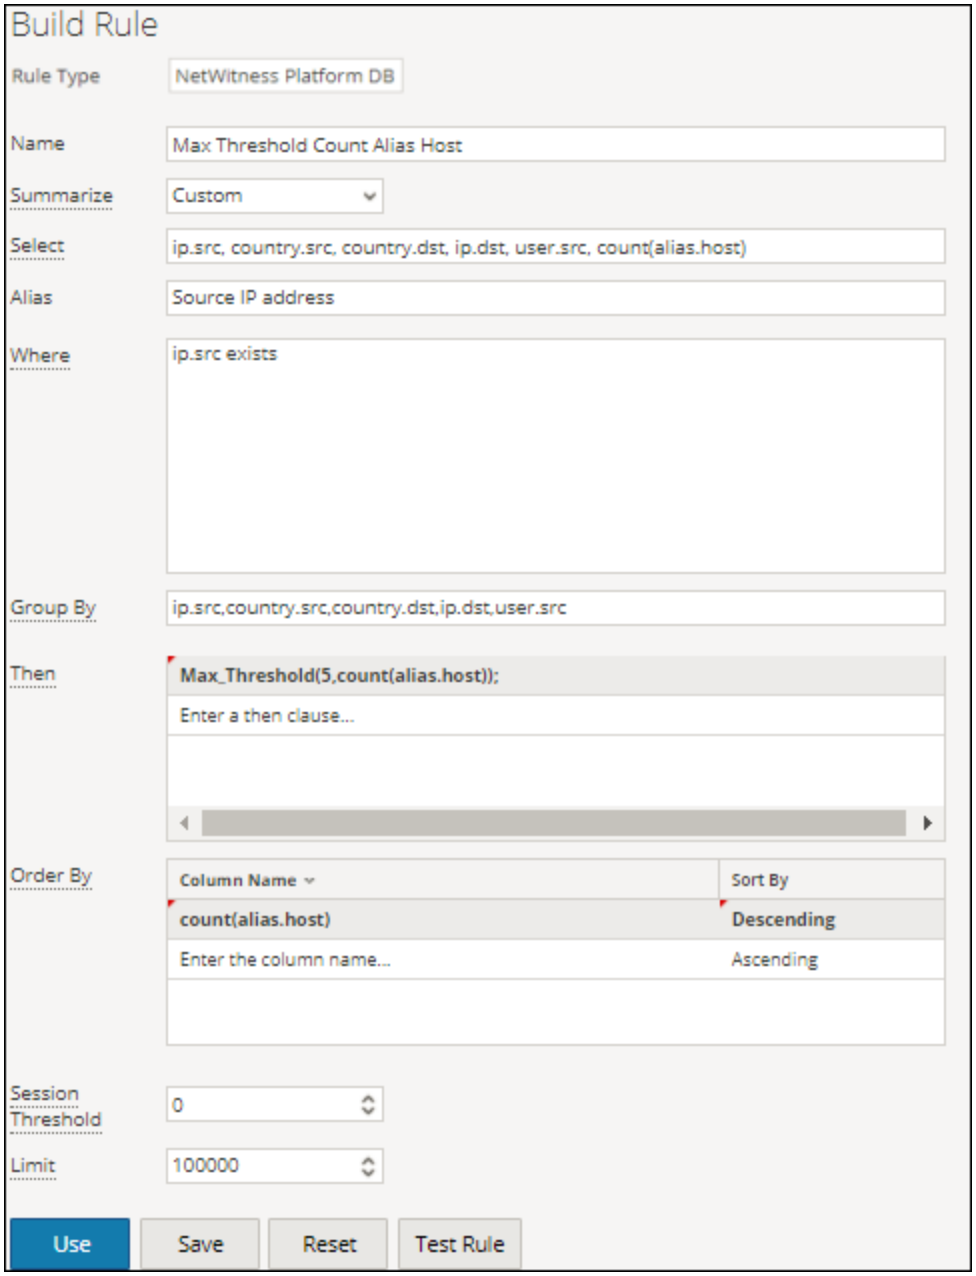


2. max_threshold(5,count(alias.host));

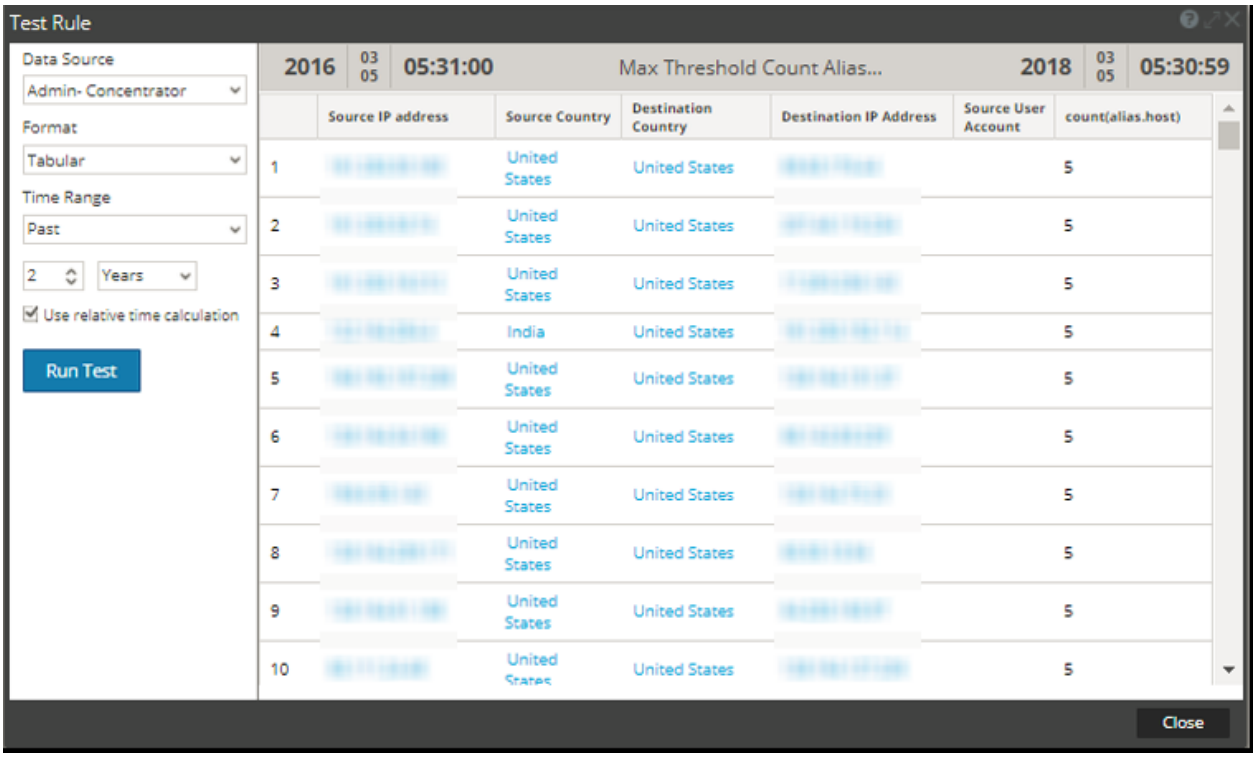
The following figure shows the result without the max_threshold argument. The output results have count of alias.host exceeding 5.



The following figure shows a the max_threshold rule action that puts a limit of 5 on the output. Any output having value more than 5 is not listed.



The following figure shows the result when the max_threshold rule action is applied. Any output having value more than 5 is removed from the result.



min_threshold (string quantity)

min_threshold removes results with a quantity that is smaller than the minimum threshold quantity from a result set. The quantity can either be in terms of count or size and it is relative to the sorting options of the parent rule. This means that if you sort a rule by size, the rule action expects you to specify the parameter in bytes (you can append KB, MB, GB, TB to the parameter to make size conversion easier).

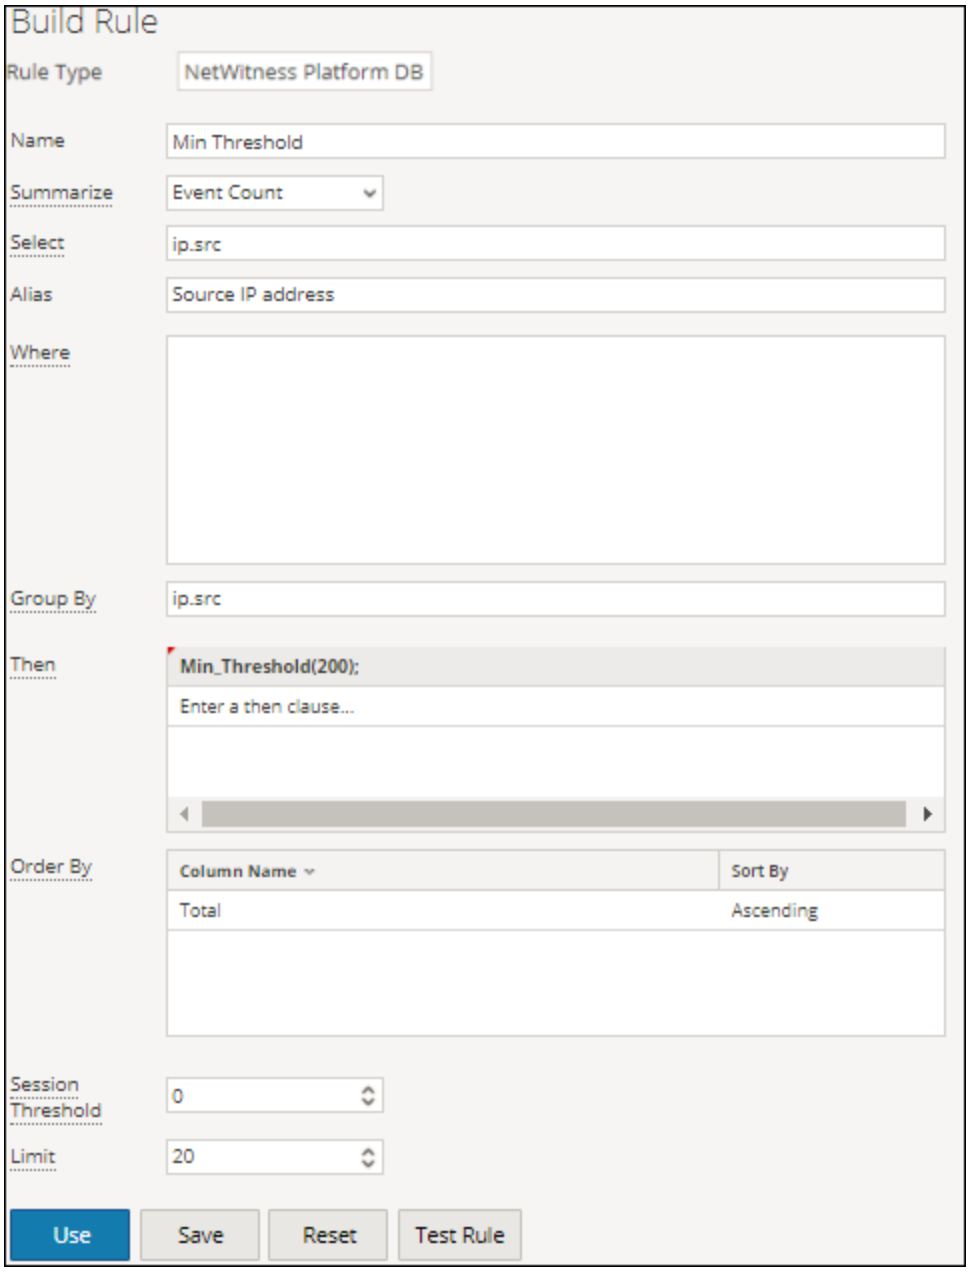
min_threshold rule can also be used to filter values based on the aggregate function values. Use the syntax based on the type of summarization used in the rule as below:

- min_threshold(String quantity): Can be used to filter Event Count, Packet Count, and Session Size.
- min_threshold(String quantity, String field): Can be used to filter values of Custom aggregates or any metas.

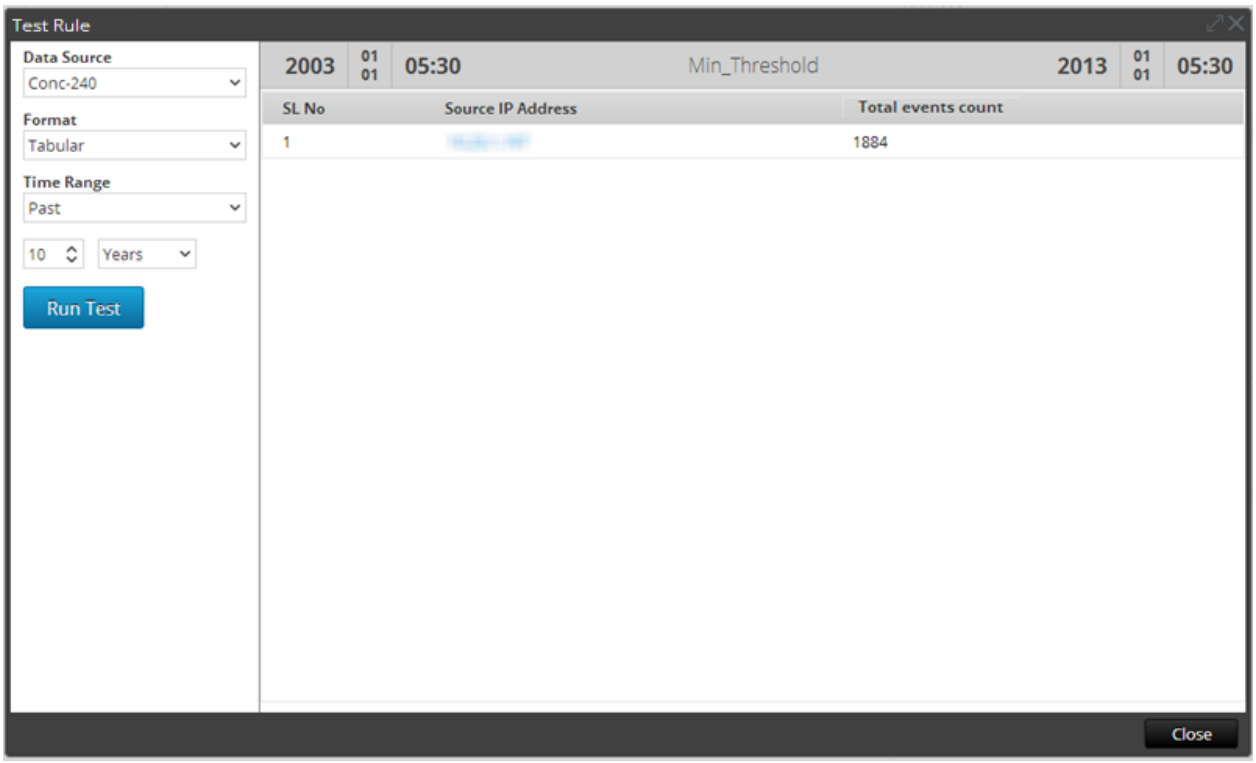
Examples:

1. min_threshold(200);

The following figure shows a sample of the min_threshold query.



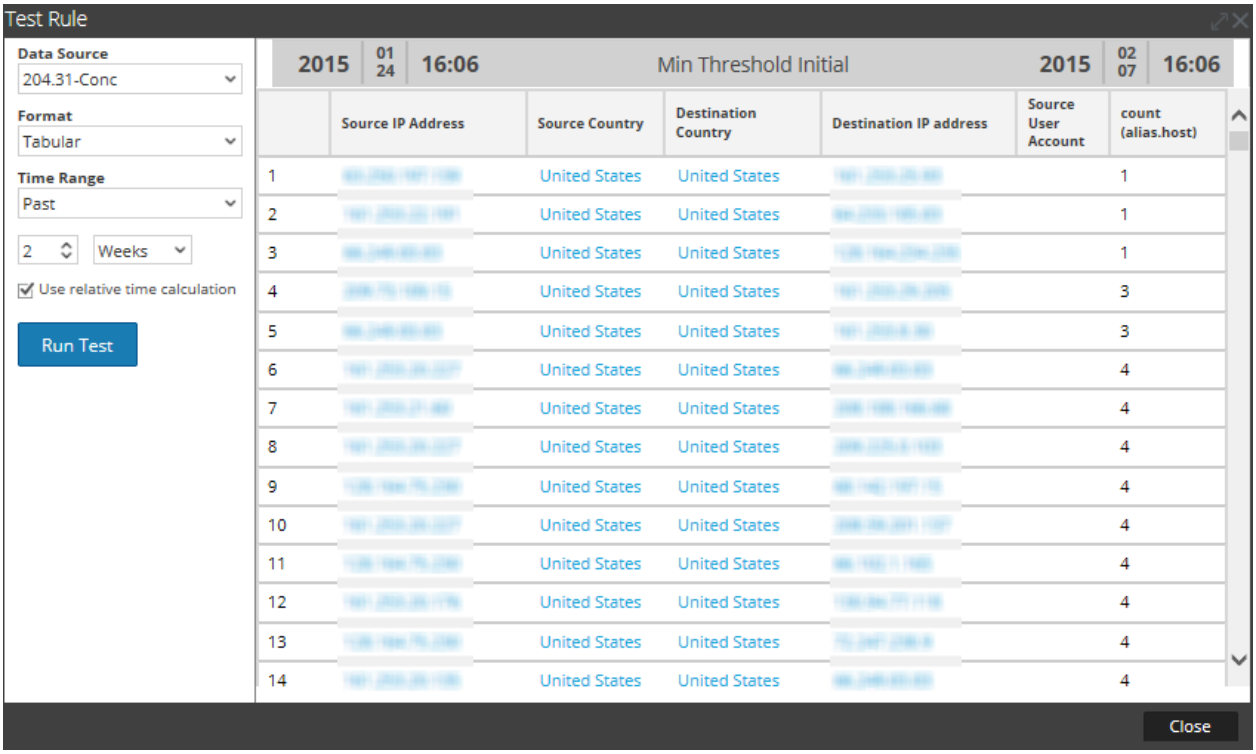
The above figure puts a limit of 200 bytes on the output. Any output having less than 200 bytes of data is not listed. The output with the min_threshold rule action is applied.



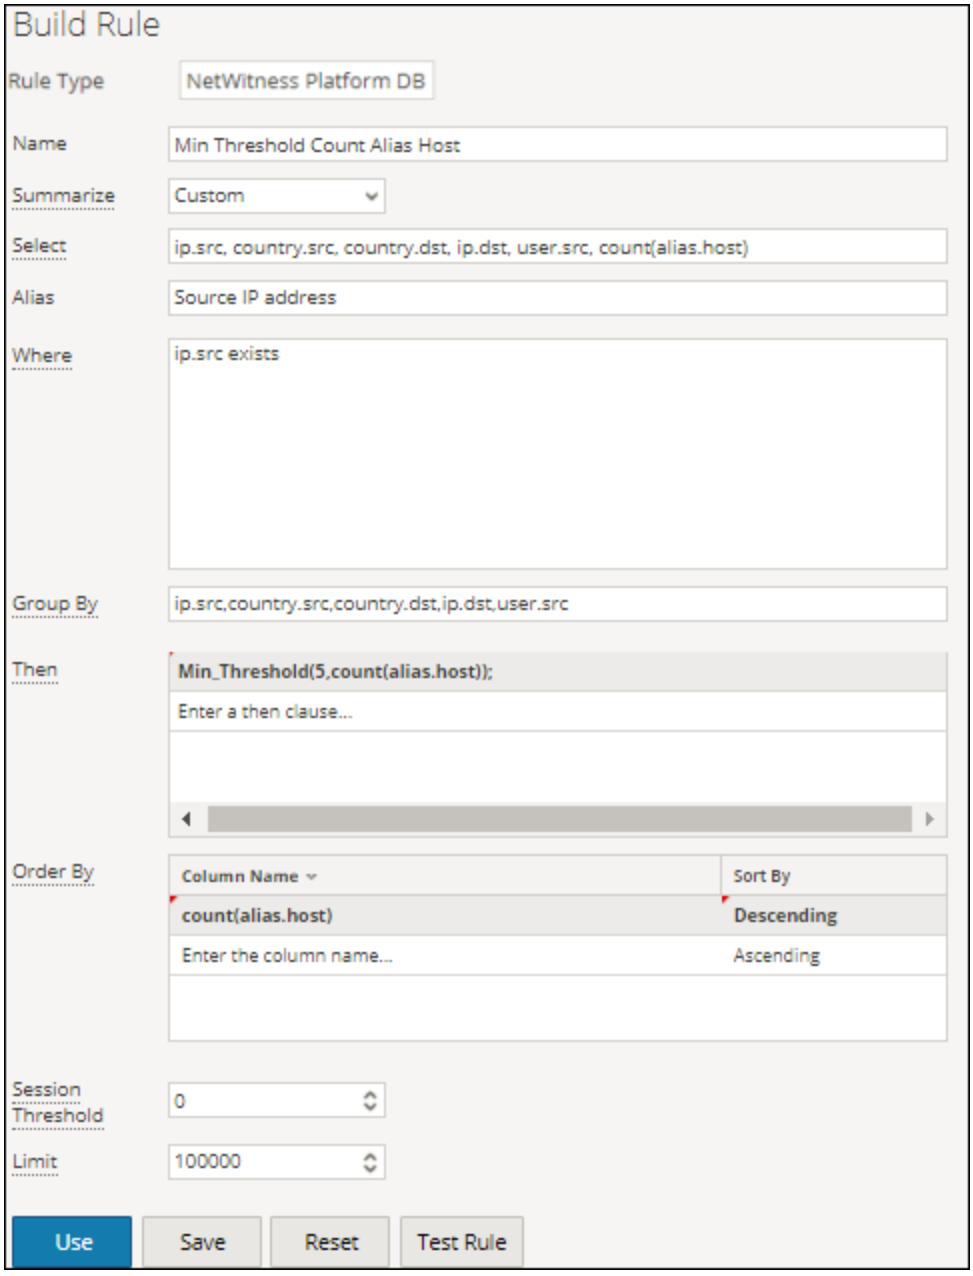
As shown, all the values are greater than 200 bytes.

2. min_threshold(100,count(alias.host));

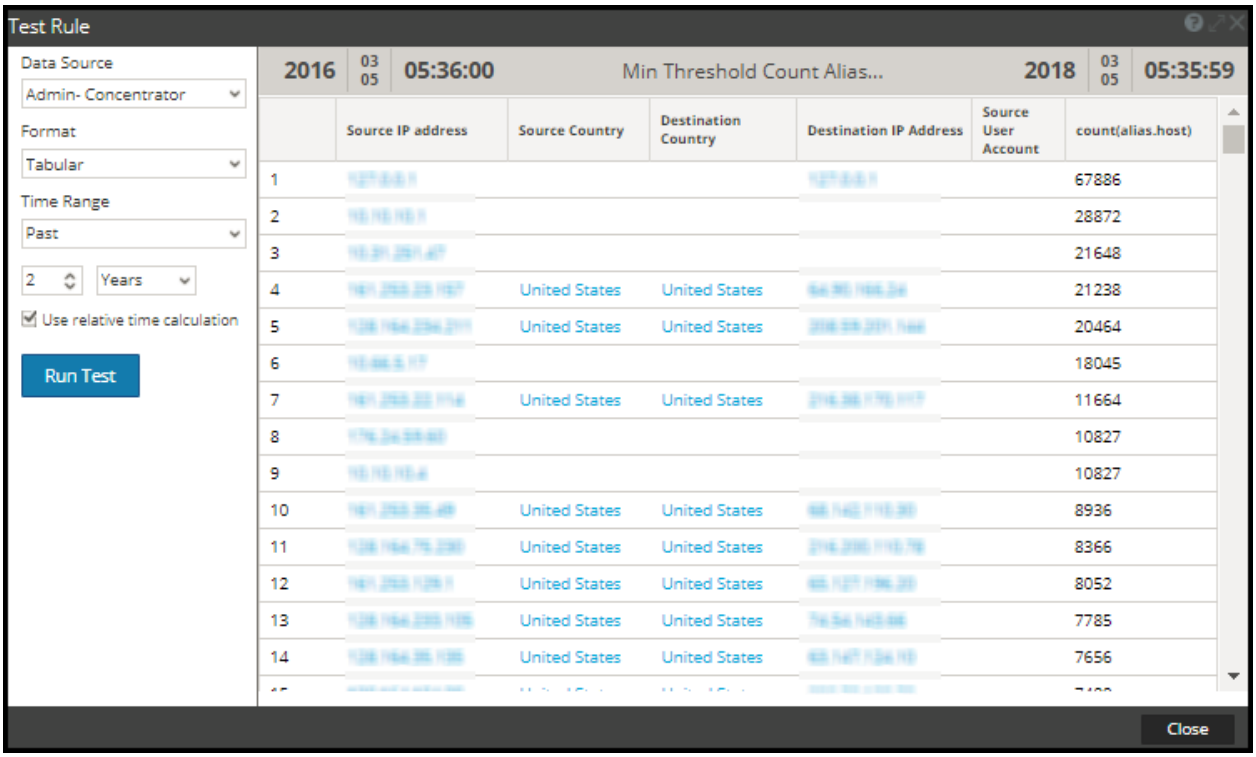
The following figure shows the result without the min_threshold argument. The output results have count of alias.host below 100.



The following figure shows a the min_threshold rule action that sets the minimum limit of 100 on the output. Any output having data less than 100 is not listed.



The following figure shows the result when the min_threshold rule action is applied. Any output having data less than 100 is removed from the result.



regex (string regex, string field)

The regex rule action applies regular expression to the result set. The following is the format of the regex rule action:

regex(regular_expression, meta_name)

Where:

- regular_expression - Regular expression to match the value of the meta.
- meta_name - Meta or field name on which the regex has to be applied.

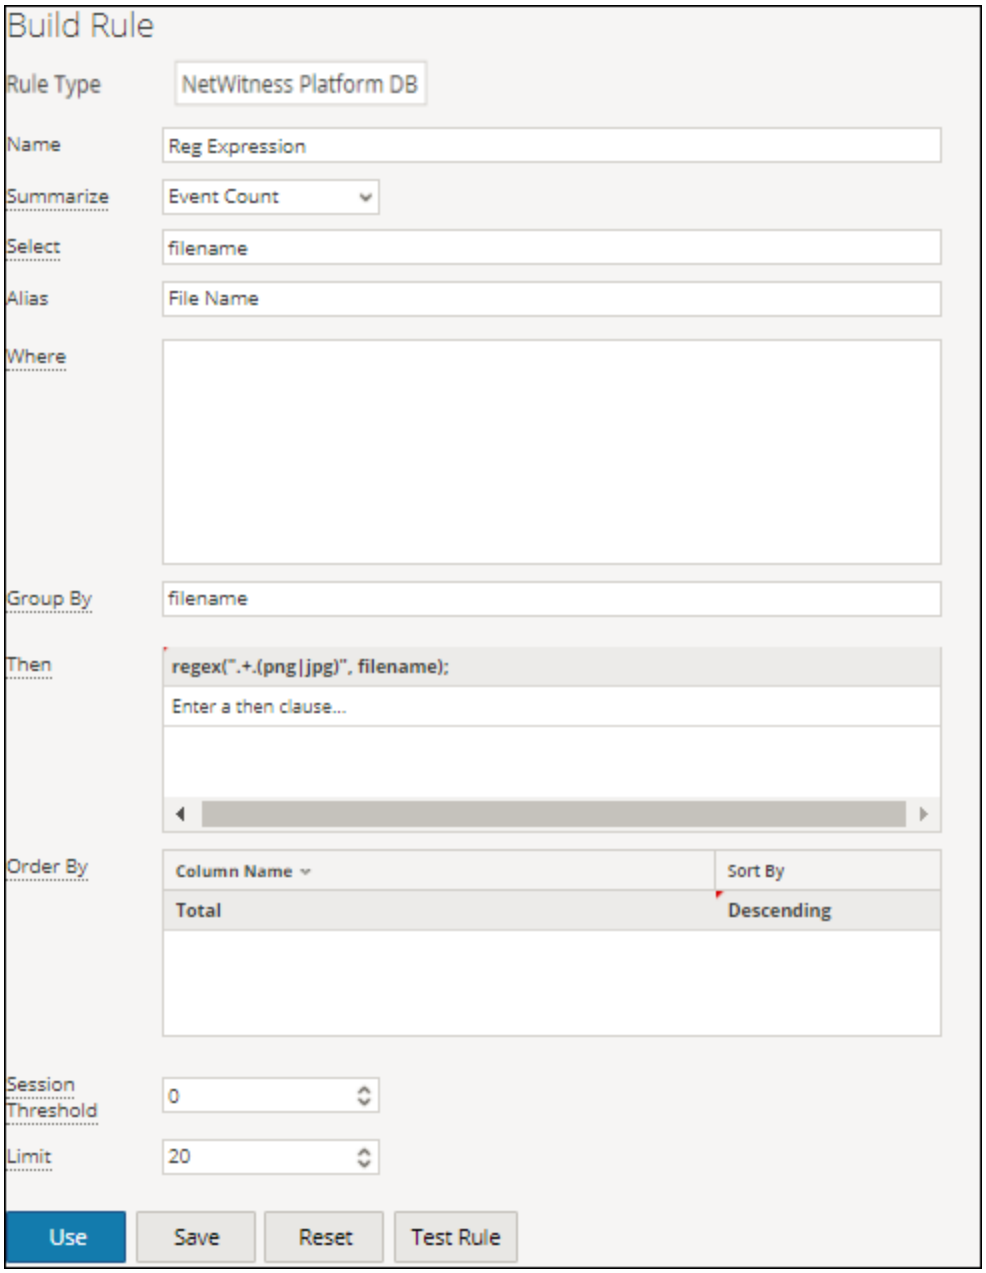
For a comprehensive list of supported regex patterns, refer to http://docs.oracle.com/javase/7/docs/api/java/util/regex/Pattern.html.

Sample regex rule action:

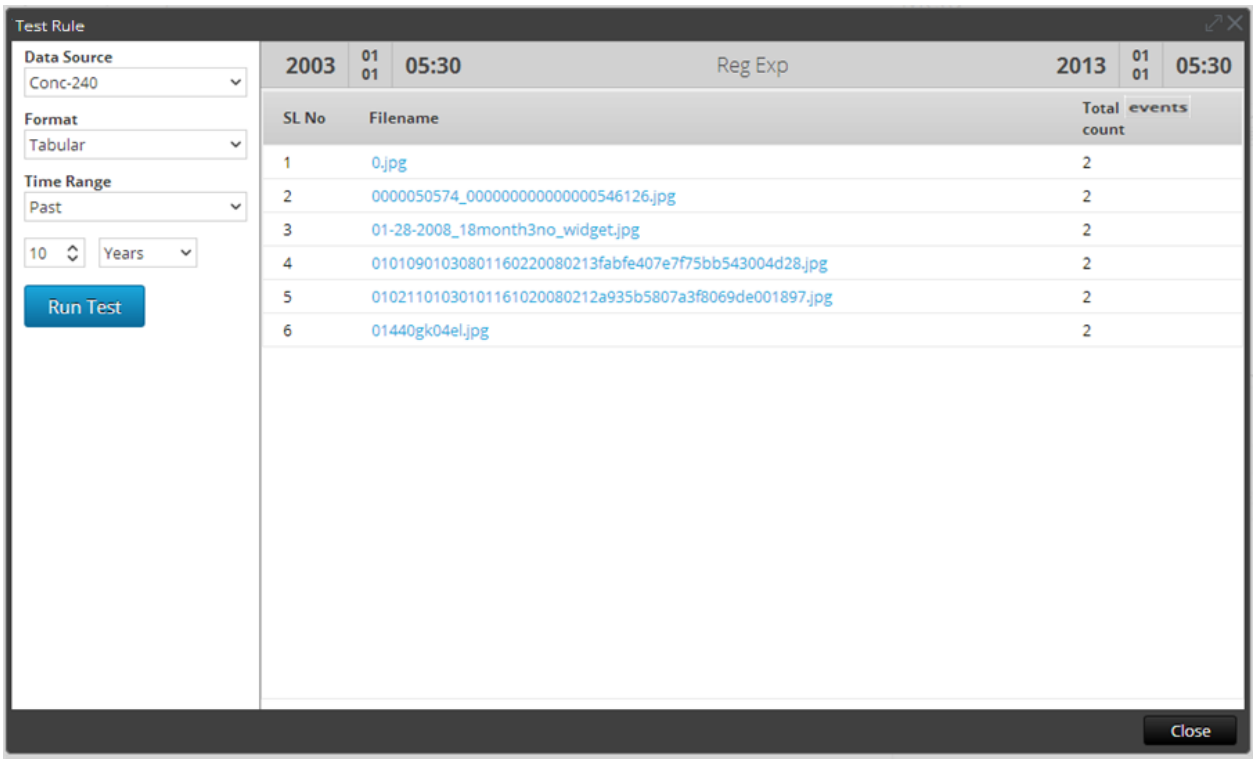
If you want to list filenames of all the PNG and JPEG format files from various sessions, you can write a rule with the following regex rule action:

regex(".+.(png|jpg)", filename);

The following figure shows the rule.



The output with the regex rule action applied is shown in the following figure.

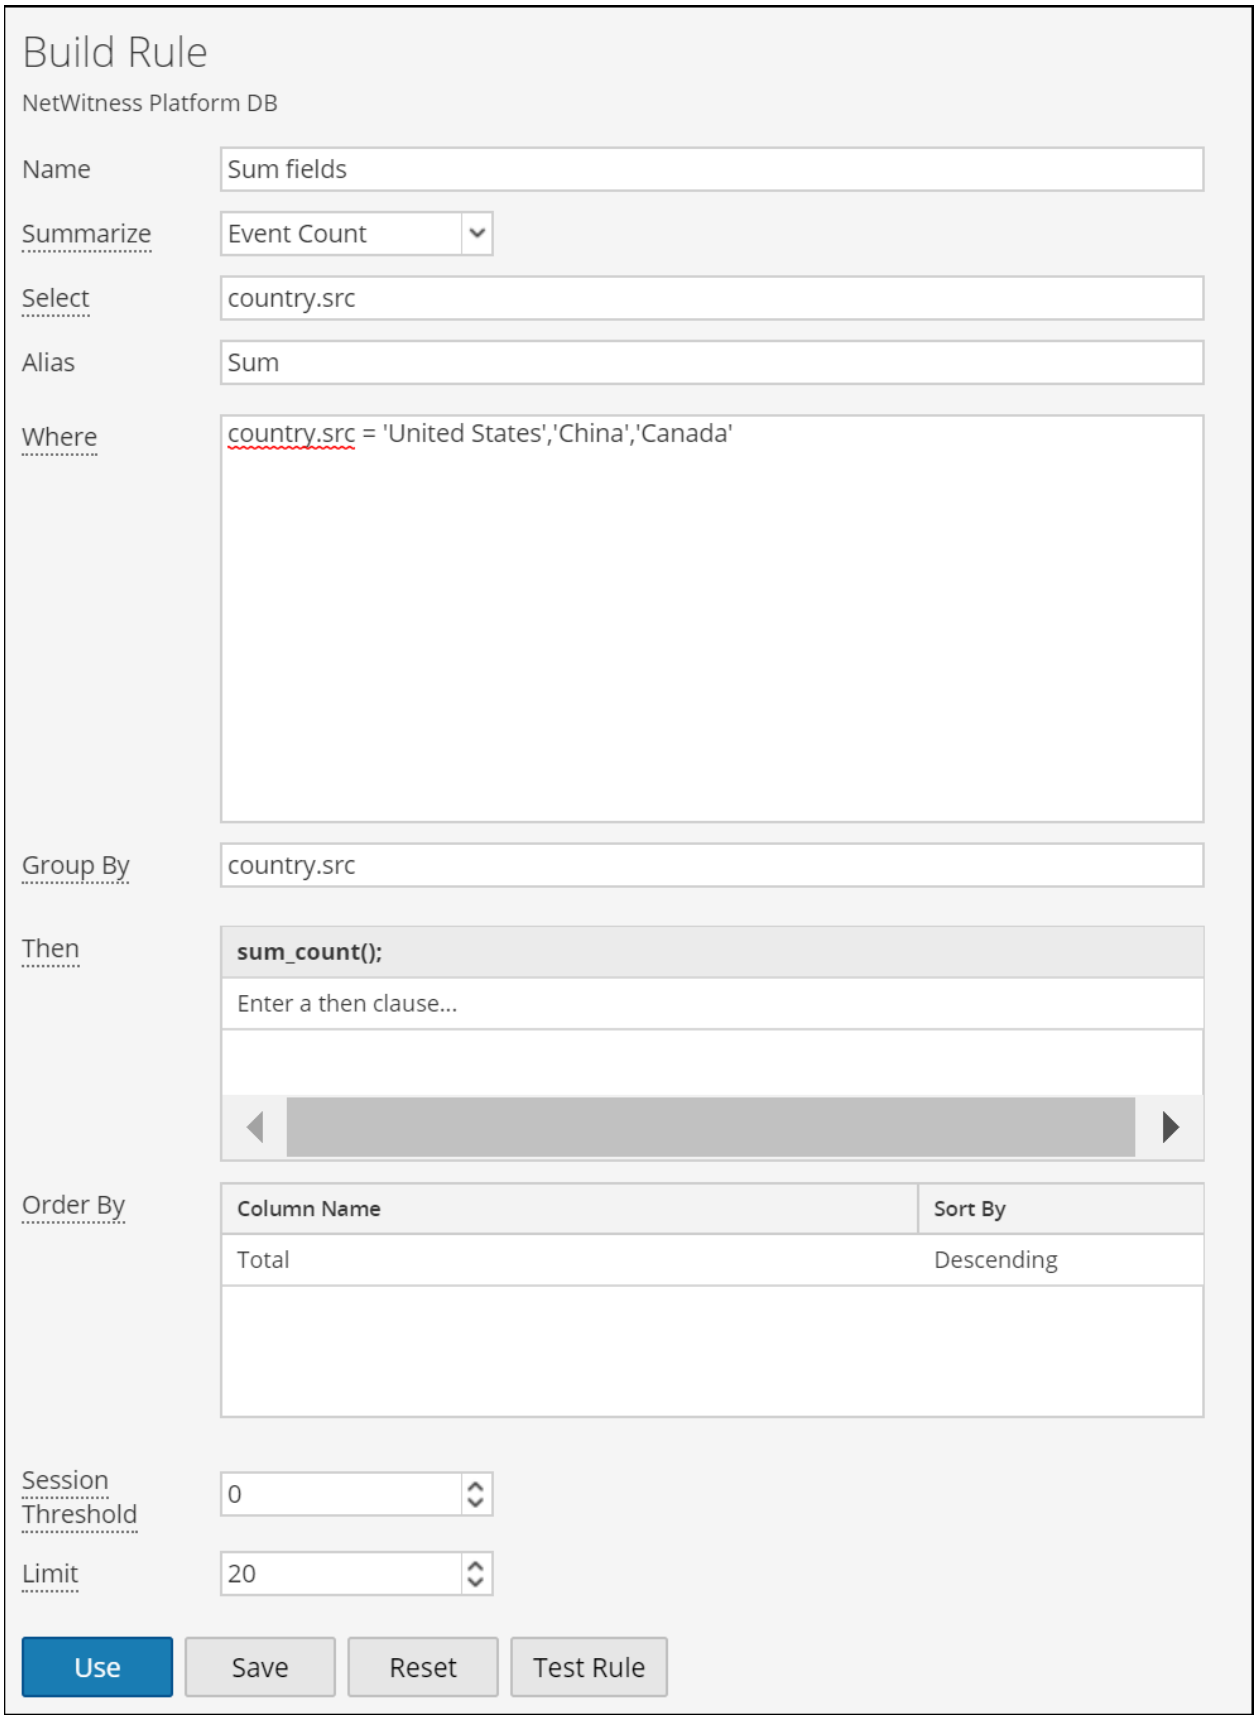


`sum_count()`

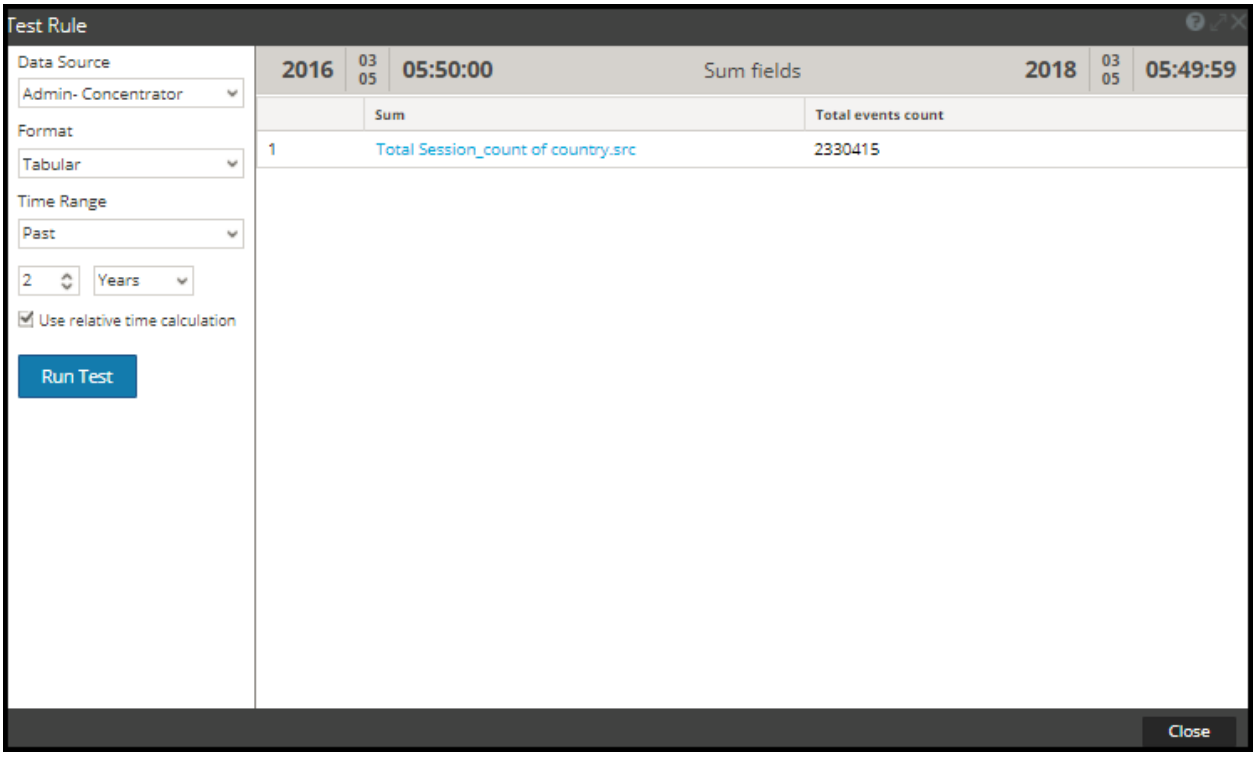
Totals the quantifiers for a given result set. For example, calling a sum_count() for a rule that is sorted by event count totals the size of all values in the result set and displays the total in place of the result set.

Example:

The following figure shows the sum_count() rule action.


Build Rule
NetWitness Platform DB
Name
Sum fields
Summarize
Event Count
Select
country.src
Alias
Sum
Where
country.src = 'United States','China','Canada'
Group By
country.src
Then
sum_count();
Enter a then clause...
Order By
Column Name
Sort By
Total
Descending
Session Threshold
0
Limit
20
Use
Save
Reset
Test Rule

With sum_count() rule action, the output shows the total size of all the event counts.

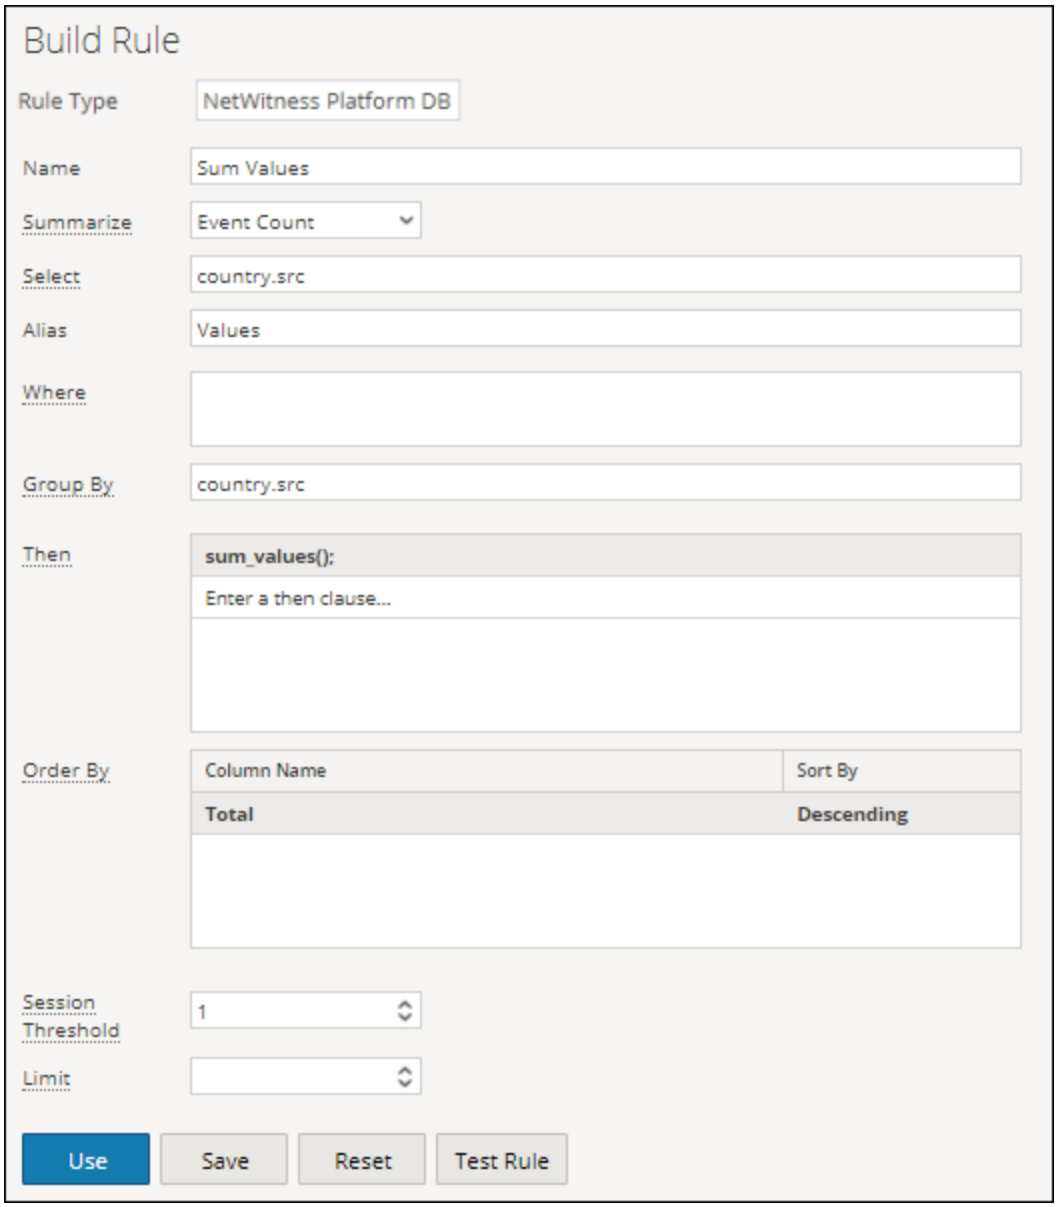


`sum_values()`

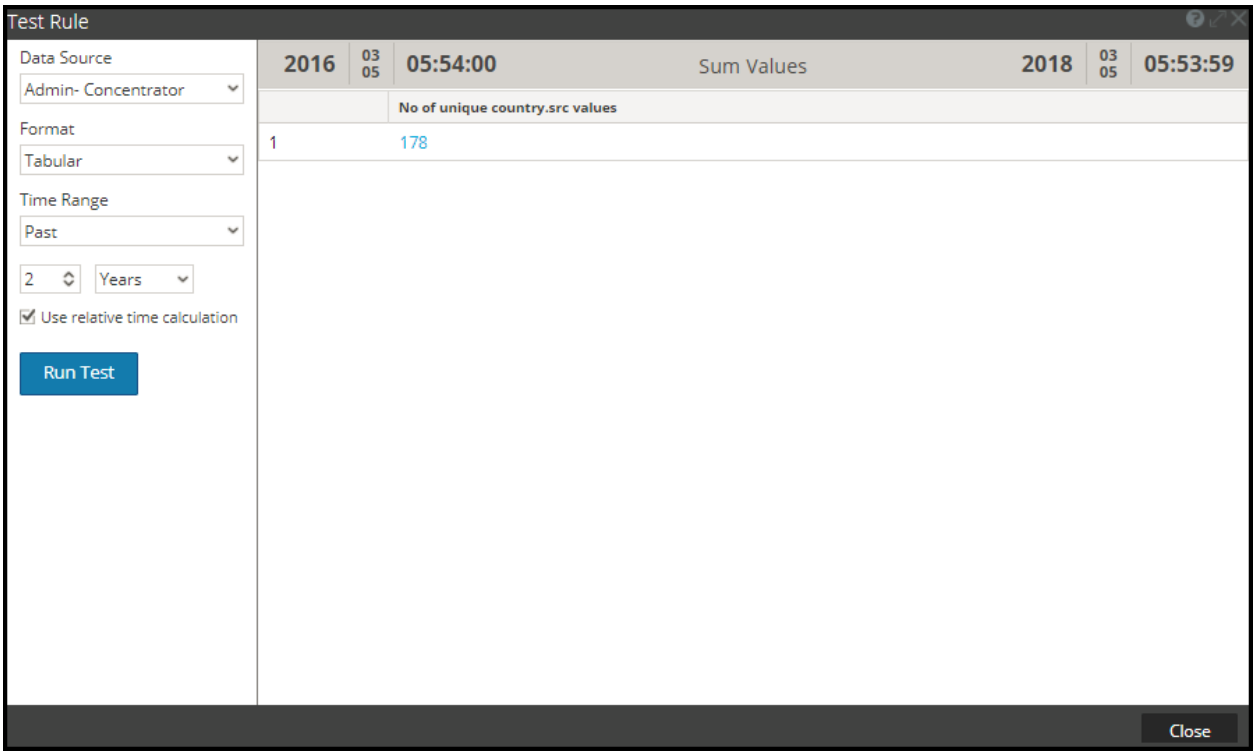
Totals the number of values for a given result set. Use this action to display how many matches exists for a given rule.

Example:

The following figure shows the sum_values() rule action.



The following figure shows the result with sum_values rule action.



show_whats_new()

The show_whats_new() rule action takes any result in a result set and filters out any value that is available in the NetWitness meta database prior to the time frame of the currently running report. When a report runs, NetWitness determines the ID of the first session in the time range of the report. If a value in a result set has a first session id that is greater than the first session id of the report time frame, it did not exist in the NetWitness meta database prior to the report being run and so is new to the NetWitness system relative to the time frame of the report.

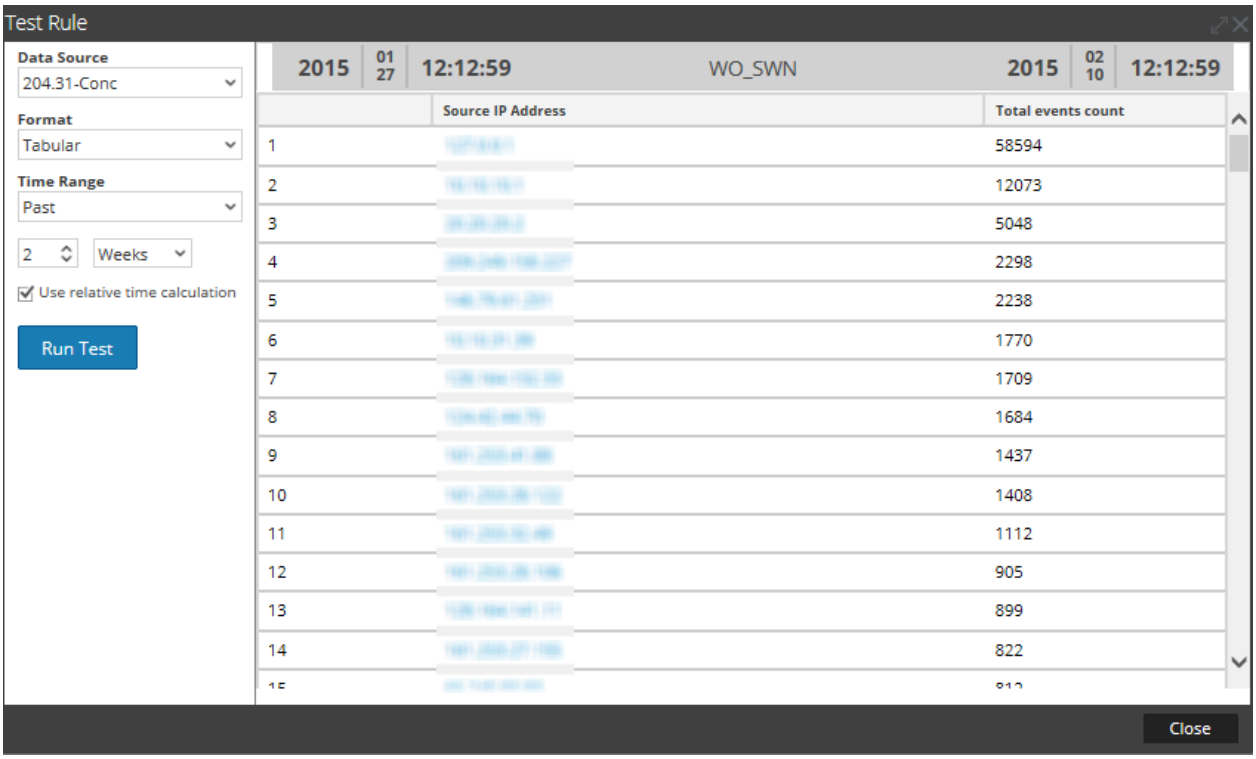
The show_whats_new() rule action is also supported for Custom Aggregate Rule. When multiple meta's are selected in the Custom rule, the first meta is considered for filtering out the old values. See "Example 2" below to understand how this rule action is used for Custom Aggregate Rule.

Note: The show_whats_new() rule action can be used only with an aggregate rule.

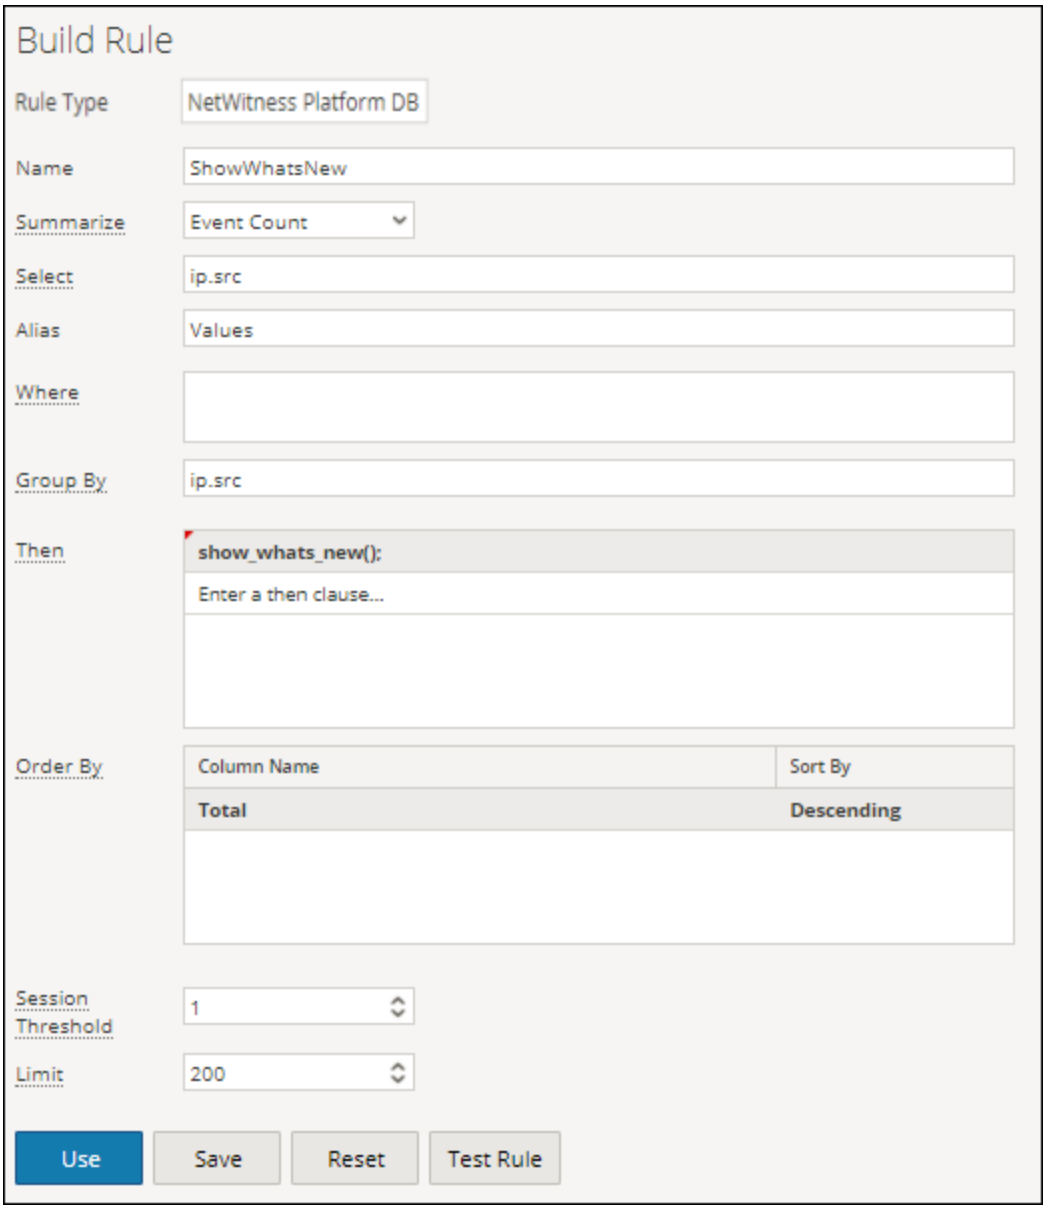
Examples:

1. show_whats_new() for aggregate rule with Event Count

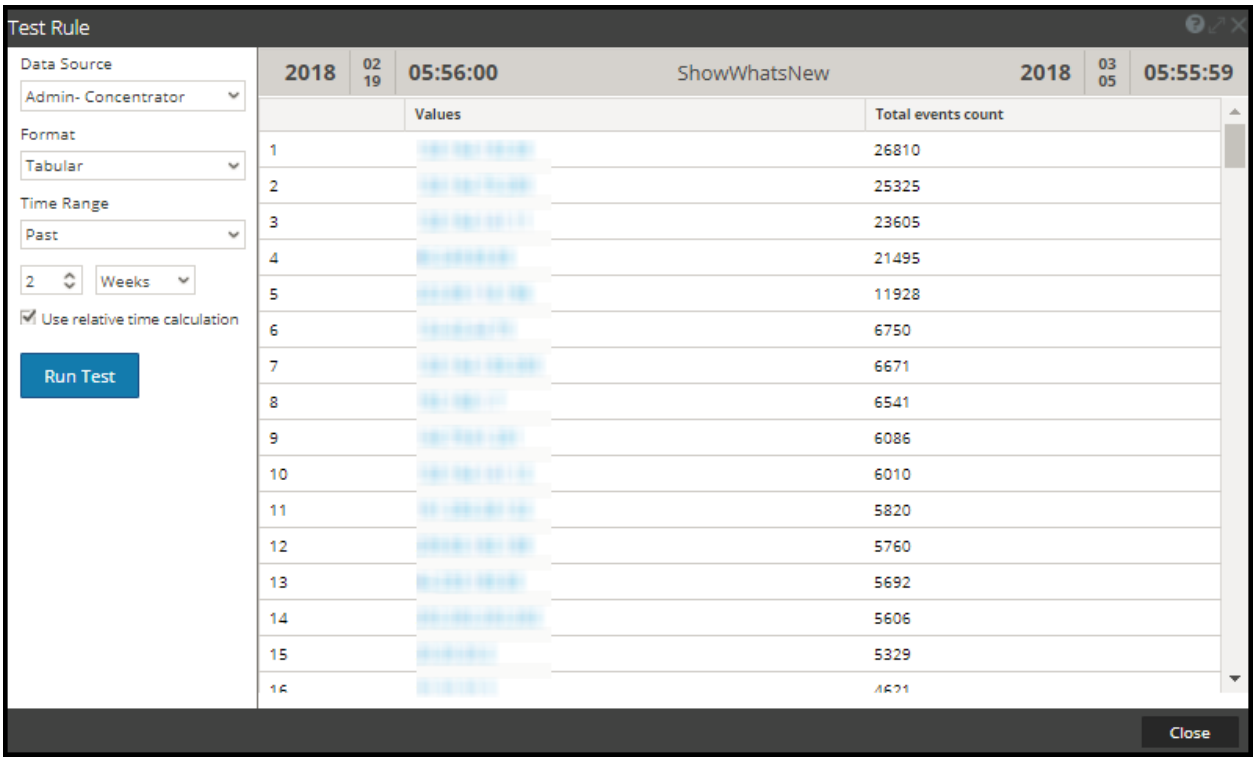
In the following example, all the Source IP Addresses available for the past two weeks are listed.



The following figure shows the use of the show_whats_new rule action to list only the new entries for the past two weeks.

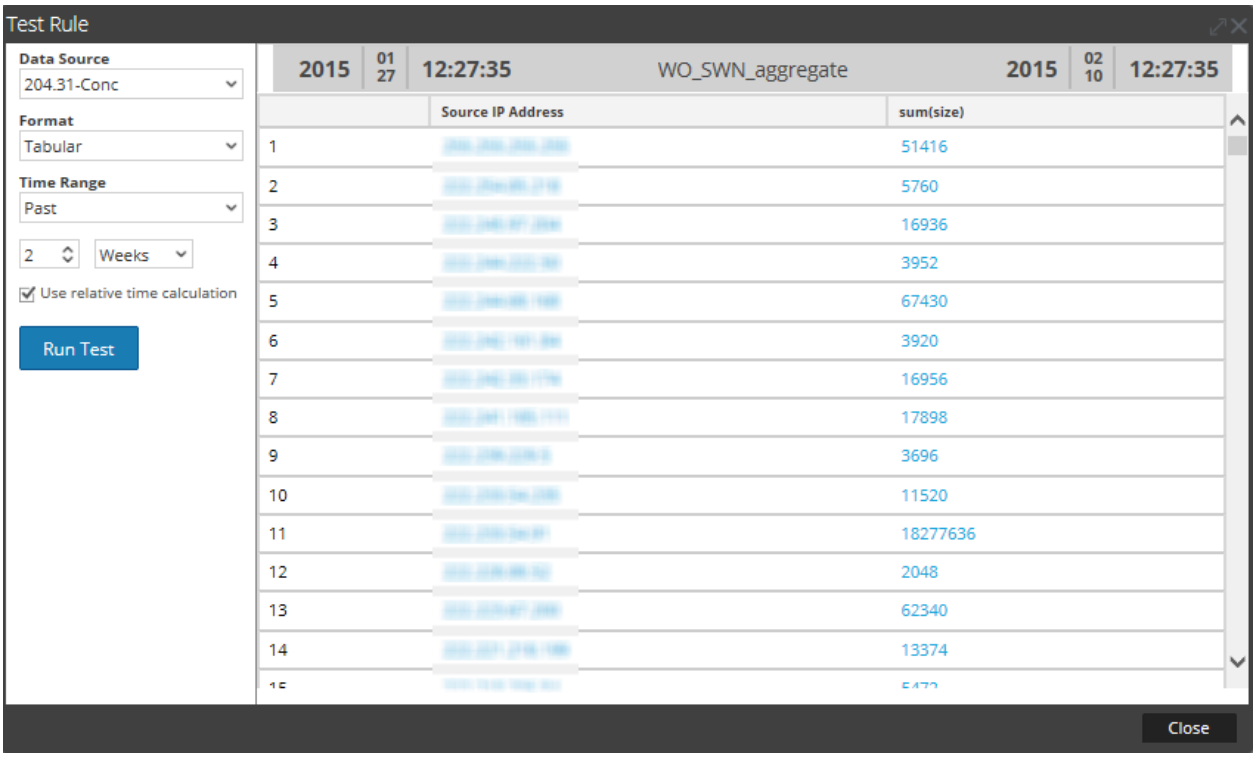


The following figure lists the new entries for the past two weeks.

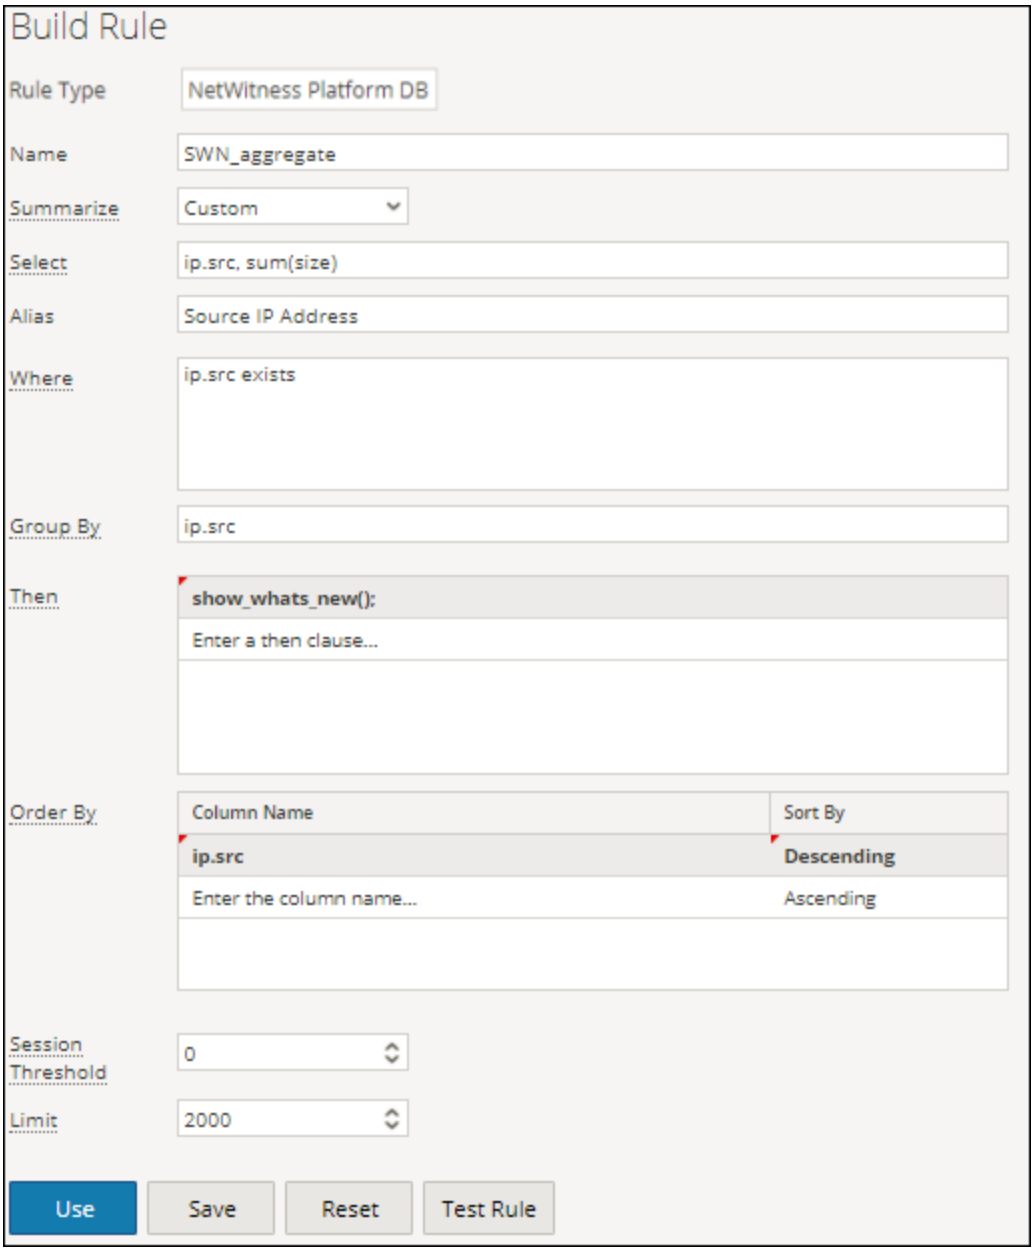


2. show_whats_new() for Custom aggregate rule

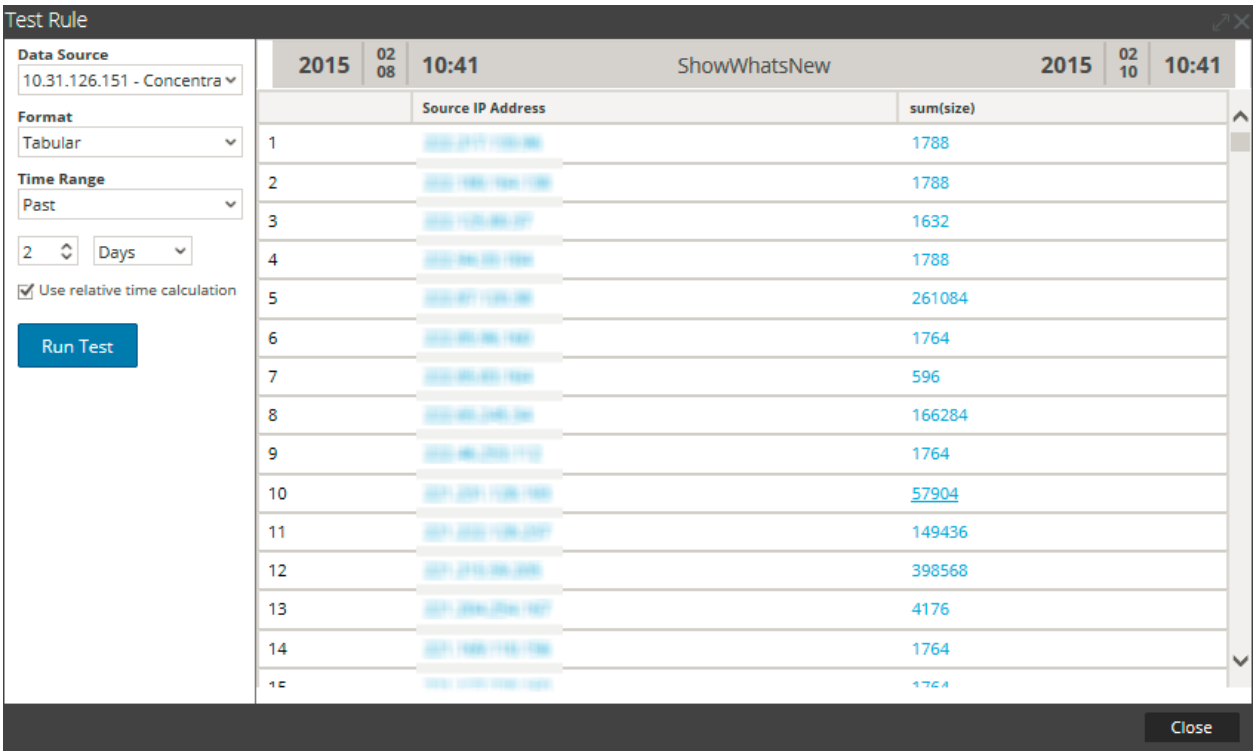
In the following example, all the Source IP Addresses available for the past two weeks are listed.



The following figure shows the use of the show_whats_new rule action to list only the new entries for the past two weeks.



The following figure lists the new entries of Source IP Addresses for the past two weeks.



The power of this feature is that it doesn't matter when the report is run in identifying values that are new to NetWitness. The caveat with this feature is that if a data reset occurs, you will lose your data. However, it is easy to baseline a system and identify changes and new items without a tremendous amount of strain on the system (depending on the size of your result set).

Supported Rule Operators

The NWDB Reporting Engine data source rule syntax supports a subset of rule operators that are supported by NetWitness.

Syntax	Description
*	Use an asterisk (*) as the sole operator in a rule to select all traffic.
=	Equals operator
!=	Does not equal operator
&&	Logical AND operator
||	Logical OR operator
-u	Upper boundary. For example, **tcp.port = 40000-u** selects all TCP ports above 40000.
l-	Lower boundary. For example, **tcp.port = l-40000** selects all TCP ports below 40000.
-	The dash (-) operator only applies to numeric values. Separate the lower and upper boundaries of the range with a dash (-). For example, **tcp.port = 25-443** selects all TCP ports between 25 and 443.

Sample Supported Queries

Respond Rule Syntax

The supported rule syntax for the Respond service through descriptions and examples of supported and unsupported syntax. There is a finite set of syntax that you can use to construct rules for reports using the Respond service in this release.

The Reporting Engine supports the following categories of Respond data source rule syntax:

- **select** clause
  - Non-Aggregate Rule
  - Aggregate Rule
- **alias**
- **where** clause
- **where** clause Operators
- Group By
- Order By
- **Limit** field

Note: List is not supported in Respond Data source rules.

Select Clause

The select clause is a comma separated list of values. For example: select alert.severity, alert.name, count (*).

There are two types of select clause for Respond Rule:

- Non-aggregate rule
- Aggregate rule

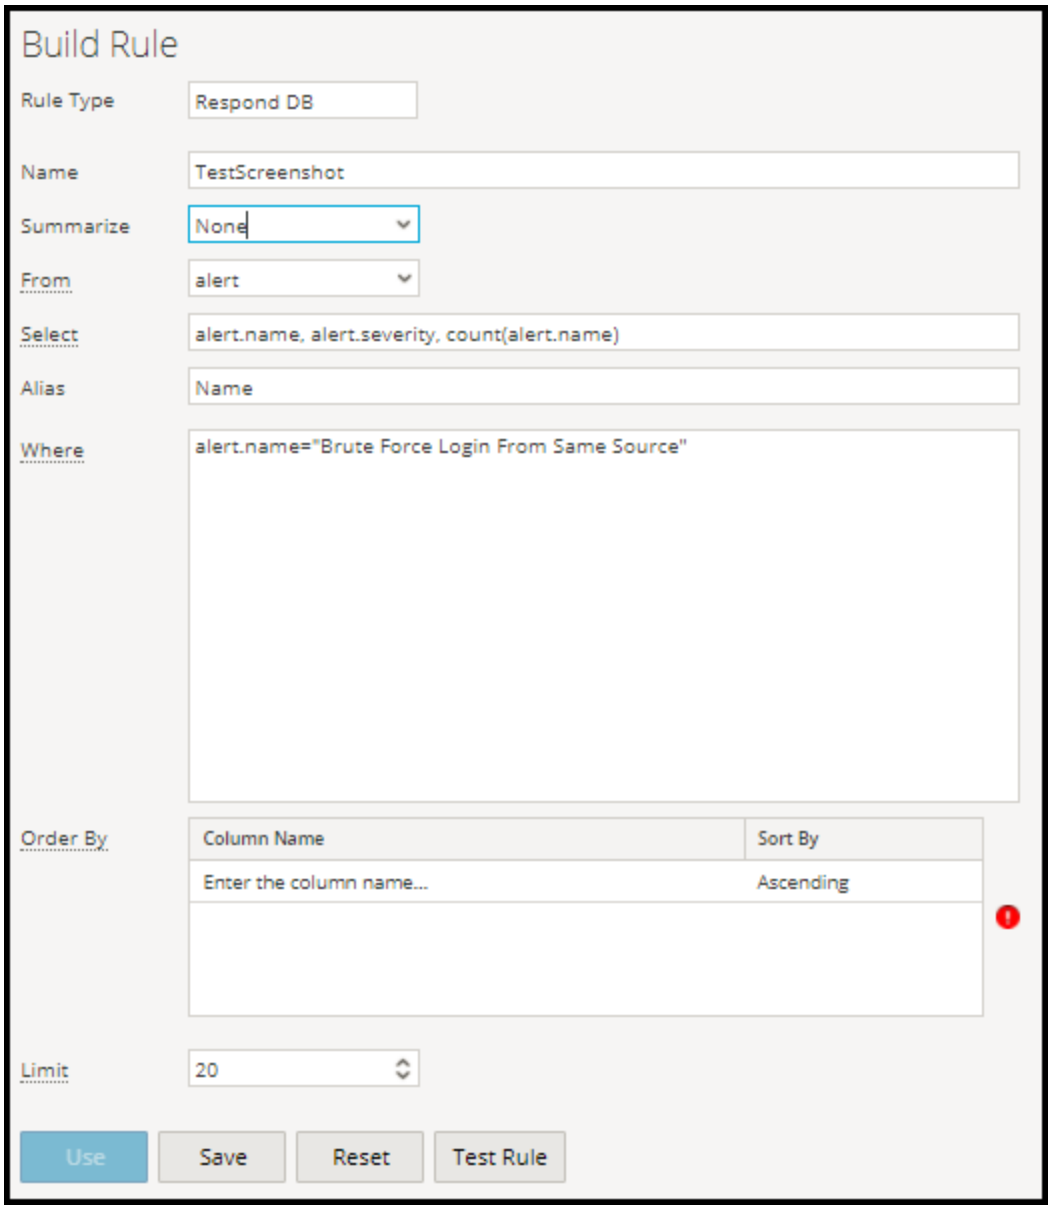
Non-Aggregate Rule

When you want to define a rule without any grouping, choose 'None' in the Summarize field. In a non-aggregate rule, you can select any number of metas in the *Select* clause. For example, select alert.severity, alert.name.

Aggregate Rule

When you want to query for a specific meta and its associated aggregate value then you must use the Aggregate rule. To get an aggregate, you must choose 'Custom' in the **Summarize** field to include an aggregate function in the *Select* clause. For example, select alert.severity, alert.name, count(*).

The following figure shows the Build Rule view for Aggregate Rule.



Supported Aggregate Functions

The rules on Respond service supports the following aggregate functions and syntax.

- count
- max
- min
- sum
- avg

Note: The aggregate functions must be added in the end of a select clause for aggregate query. For example, alert.name, alert.severity, sum(alert.numEvents). By default, a maximum of 10,000 rows results are fetched and this can be configured using the **rsa.response.query.QueryProperties**.

Examples of select Clause Syntax

The following table provides examples of the select Clause Syntax.

Examples	Descriptions
select *column1* ,*column2*,*column3*,...,*columnN*	Select specific metas from an Respond Data Source (You must separate each column with a comma.).

Examples of Supported Select Queries

select alert.name, alert.numEvents, count(alert.numEvents)

select alert.severity, avg(alert.severity)

select alert.timestamp, incidentCreated where alert.timestamp >= 1475658011

Summarize

Summarize determines the type of summarization or aggregation for the rule.

Name	Config Value
Summarize	To query metas without any custom grouping, select: • **None**: To get meta based aggregates, select: • **Custom**: This indicates that expected meta aggregate function is defined in rule select clause.

Alias

Some meta names may not be descriptive, in this case description can be added in the the alias field to make column names more readable. For example, **SELECT**: alert.severity, alert.name, count(*)
ALIAS: Alert Severity, Alert Name

In the alias field you can enter a name for columns used in the select clause. If you do not specify the alias for one of the field in the select clause, then the default description will be used. For example, if the select clause has Field1,Field2,Field3,Field4, and alias has only Field1, ,Field3,Field4, then for Field2 a default description is used.

Where Clause

The where clause is a language field values and ranges that is used by Respond function. In the where clause, string values have to be enclosed within single quotes.

Examples	Descriptions
alert.host summary =' (Primary) Link status "Down" on interface INTNAME.'	For TEXT or string type data, enclose the string or text in single or double quote. If there is any special character such as an apostrophe within the data then you need to add an additional single or double quotes. For example, alert.name = 'top alerts from Cote d"Ivoire'.
alert.timestamp >= 1475658011	For Date and Time (date/timestamp data type columns), use the EPOCH syntax.

Supported Where Clause Operators

Operator	Syntax
= (equals)	*column1* = '*value*'
!= (does not equal)	*column1* != '*value*'
>	*column1* > '*value*'
>=	*column1* >= '*value*'
<	*column1* < '*value*'
<=	*column1* <= '*value*'

Group By

Syntax	Function
group by : alert.severity, alert.timestamp, incidentCreated **Note:** Group by field is enabled for Aggregate queries and are not editable.	Respond picks the metas for Group By field from the selected Select clause automatically.

Order By

Order By determines how to sort the result set and is not case sensitive.

Name	Configuration Value
Column Name	The Column Name is the name of the columns by which you want to sort the results. By default, the value is empty. When you click on a column, the value gets populated based on the Summarize field. • order by alert.name asc • order by incidentCreated desc • order by count(numEvents) • order by status
Sort By	Sort By determines the order in which you want to sort the results such as ascending or descending. **Note:** For all queries, it is mandatory for you to select the order by field.

Limit field

This indicates the limit to be put on the query while fetching data from the database. If a result set is sorted by event count, packet count, or session size, the limit represents the top (or bottom) N values to be returned. If the result set is not sorted, the first N values are returned.

Warehouse DB Simple Rules Syntax

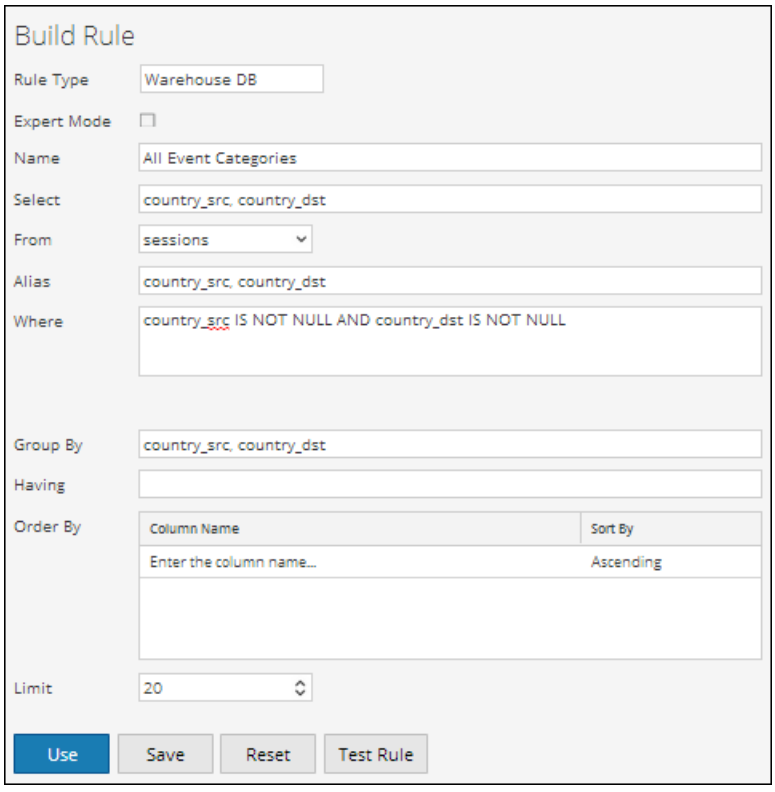
The section explains the simple rules query syntax and examples.

The following examples illustrate simple rules in the default mode:

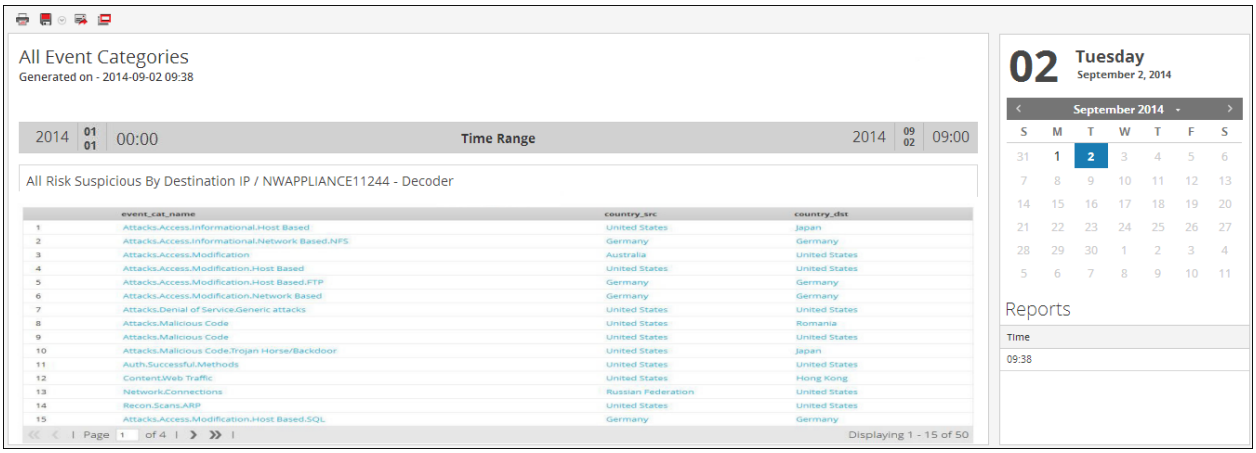
- All Event Categories Report
- Attacks Event Categories Report
- Source: China Event Categories Report
- IP Source and Destination Event Categories Report
- Time Threat Categories Report
- Array Query Report
- Raw Log Query Report

All Event Categories Report

This rule fetches all event categories, source country, and destination country from the **sessions** table by defining alias names (temporary column names) for each of the fields to be fetched from the table, that is, **country_src** for the source country, and **country_dst** for the destination country.

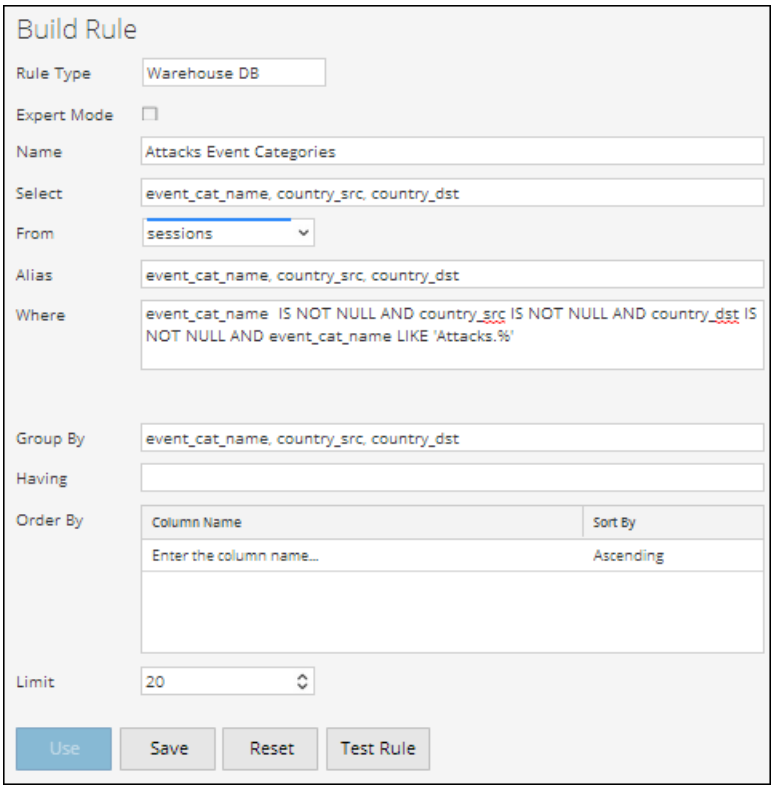


The following figure shows the result set of the All Event Categories rule.

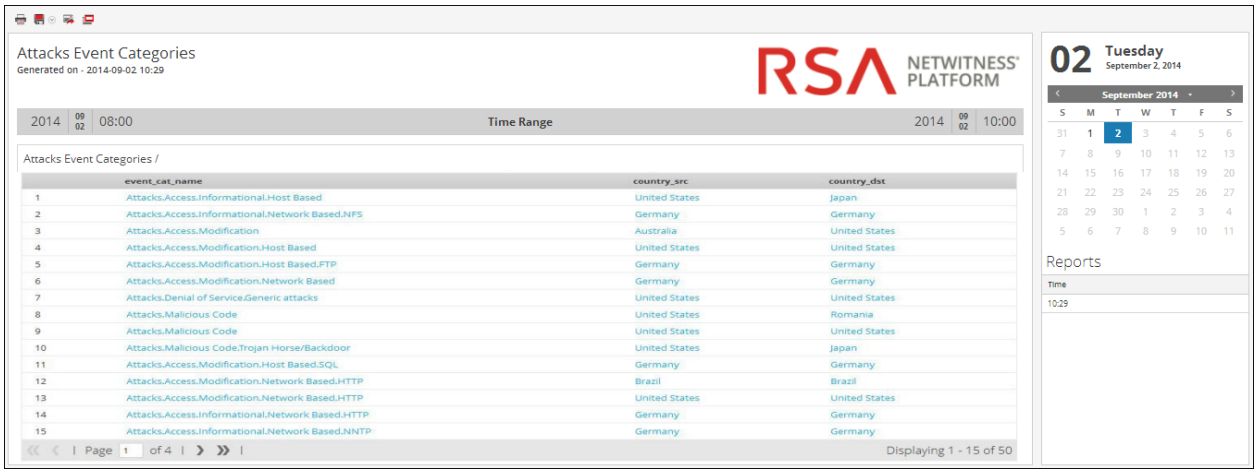


Attacks Event Categories Report

This rule fetches the event categories, source country, and destination country from the **sessions** table by defining alias names (temporary column names) for each of the fields to be fetched from the table and selecting only those columns whose event category name like 'Attacks.%'.

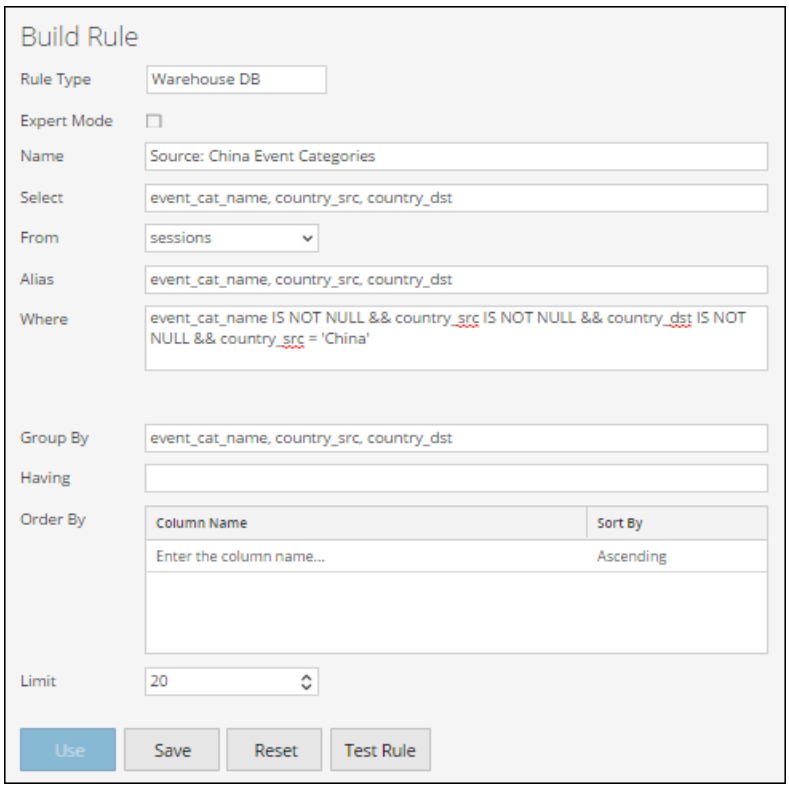


The following figure shows the result set of the Attacks Event Categories rule.

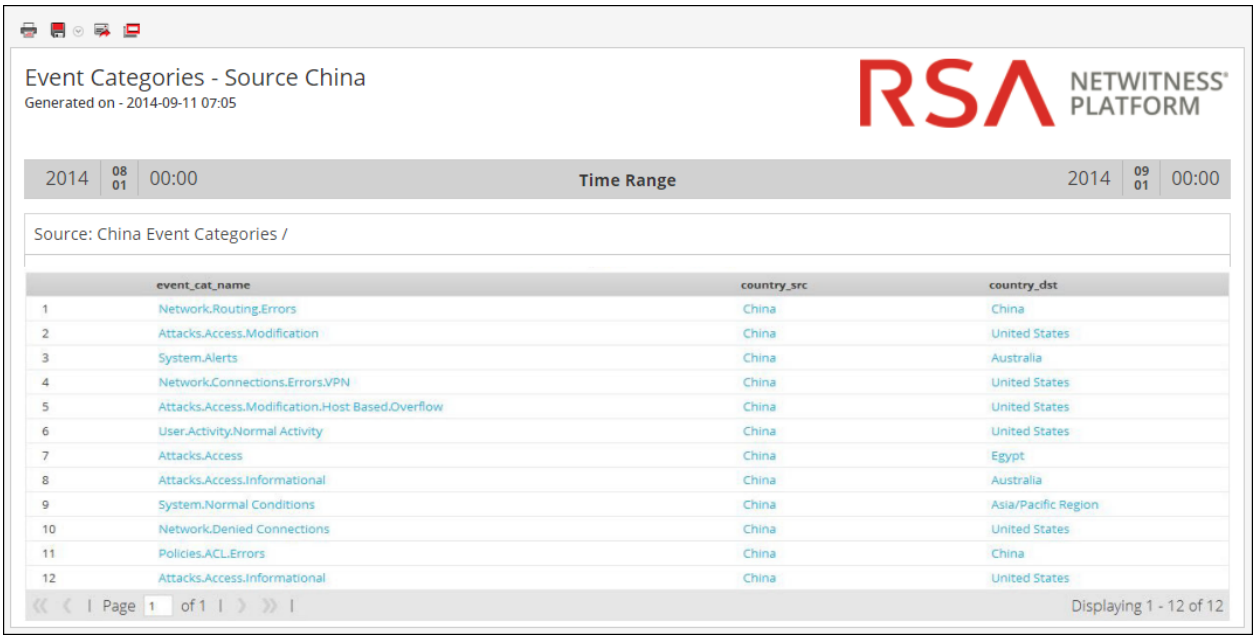


Source: China Event Categories Report

This rule fetches the event categories, source country, and destination country from the **sessions** table by defining alias names (temporary column names) for each of the fields to be fetched from the table and selecting only those columns whose source country is 'China'.



The following figure shows the result set of the Source: China Event Categories rule.



Event Categories - Source China
Generated on - 2014-09-11 07:05

RSA NETWITNESS PLATFORM

2014 08 01 00:00 **Time Range** 2014 09 01 00:00

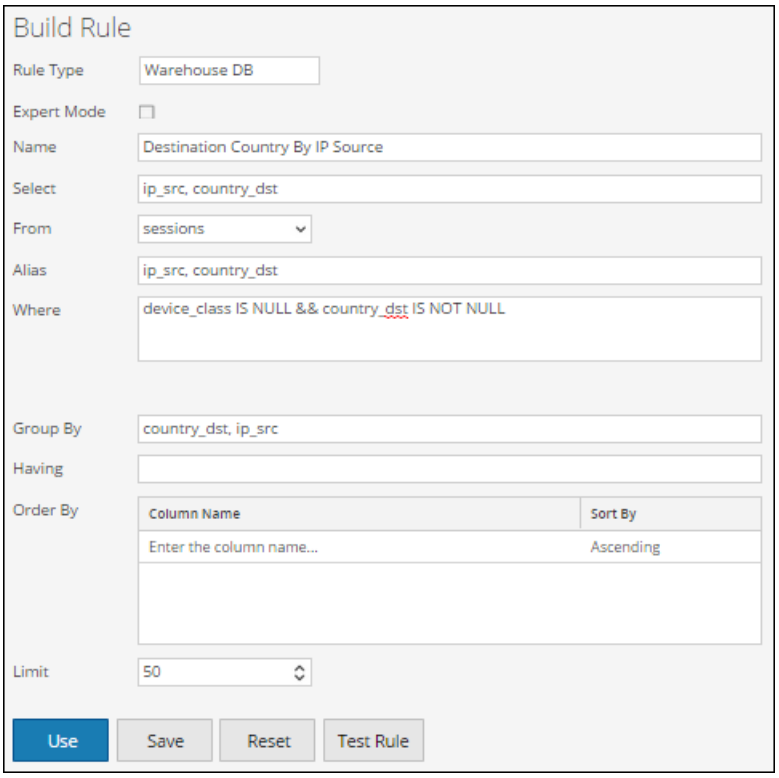
Source: China Event Categories /

	event_cat_name	country_src	country_dst
1	Network.Routing.Errors	China	China
2	Attacks.Access.Modification	China	United States
3	System.Alerts	China	Australia
4	Network.Connections.Errors.VPN	China	United States
5	Attacks.Access.Modification.Host Based.Overflow	China	United States
6	User.Activity.Normal Activity	China	United States
7	Attacks.Access	China	Egypt
8	Attacks.Access.Informational	China	Australia
9	System.Normal Conditions	China	Asia/Pacific Region
10	Network.Denied Connections	China	United States
11	Policies.ACL.Errors	China	China
12	Attacks.Access.Informational	China	United States

Page 1 of 1 Displaying 1 - 12 of 12

IP Source and Destination Event Categories Report

This rule fetches the IP address of source and destination country from the **sessions** table by defining alias names (temporary column names) for each of the fields to be fetched from the table and selecting only those columns whose destination country is NOT NULL.



Build Rule

Rule Type: Warehouse DB

Expert Mode: ☐

Name: Destination Country By IP Source

Select: ip_src, country_dst

From: sessions

Alias: ip_src, country_dst

Where: device_class IS NULL && country_dst IS NOT NULL

Group By: country_dst, ip_src

Having:

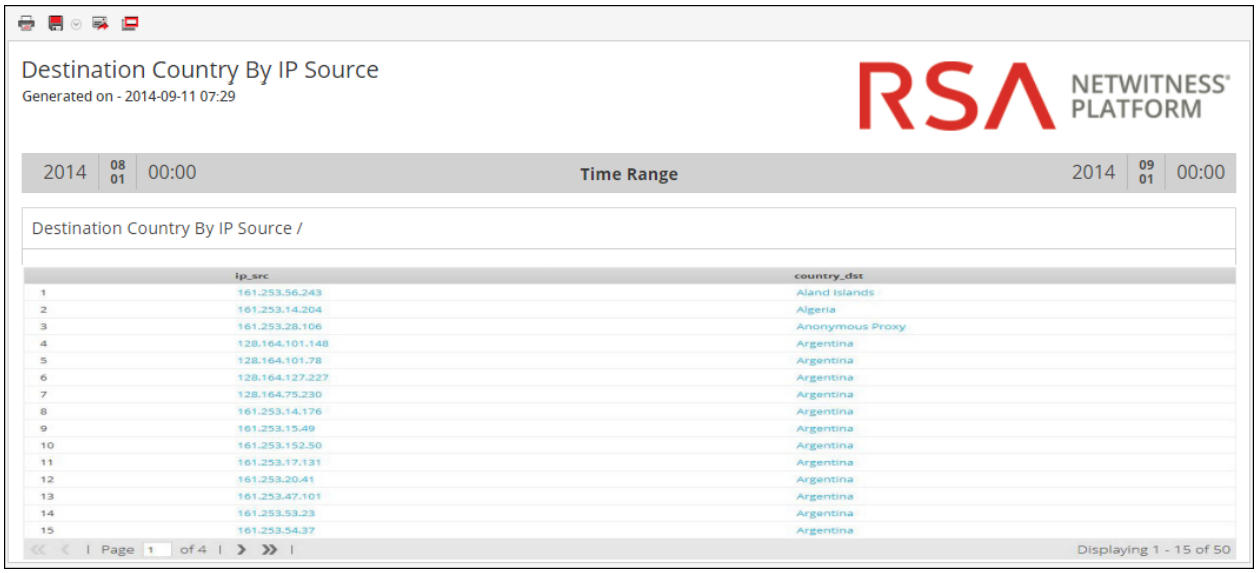
Order By:

Column Name	Sort By
Enter the column name...	Ascending

Limit: 50

Use | Save | Reset | Test Rule

The following figure shows the result set of the IP Source and Destination Event Categories rule.



Destination Country By IP Source

Generated on - 2014-09-11 07:29

RSA NETWITNESS PLATFORM

2014	08 01	00:00	Time Range	2014	09 01	00:00

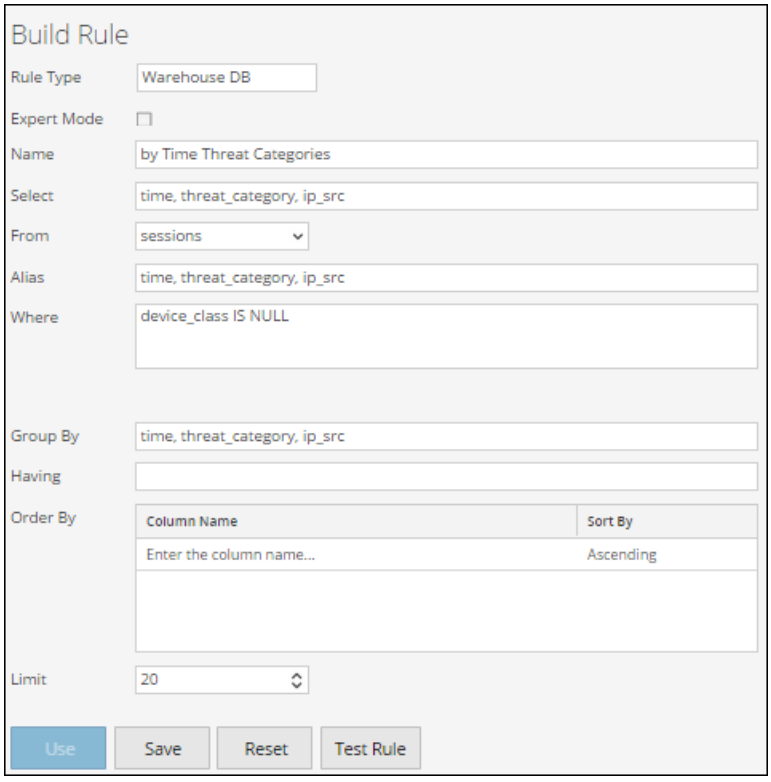
Destination Country By IP Source /

	ip_src	country_dst
1	161.253.56.243	Aland Islands
2	161.253.14.204	Algeria
3	161.253.28.106	Anonymous Proxy
4	128.164.101.148	Argentina
5	128.164.101.78	Argentina
6	128.164.127.227	Argentina
7	128.164.75.230	Argentina
8	161.253.14.176	Argentina
9	161.253.15.49	Argentina
10	161.253.152.50	Argentina
11	161.253.17.131	Argentina
12	161.253.20.41	Argentina
13	161.253.47.101	Argentina
14	161.253.53.23	Argentina
15	161.253.54.37	Argentina

Page 1 of 4 Displaying 1 - 15 of 50

Time Threat Categories Report

This rule fetches the threat category events, the time the log or event was ingested into Log Decoder/Decoder, and the source IP addresses from the **sessions**table by defining alias names (temporary column names) for each of these fields to be fetched from the table.



Build Rule

Rule Type: Warehouse DB

Expert Mode: ☐

Name: by Time Threat Categories

Select: time, threat_category, ip_src

From: sessions

Alias: time, threat_category, ip_src

Where: device_class IS NULL

Group By: time, threat_category, ip_src

Having:

Order By:

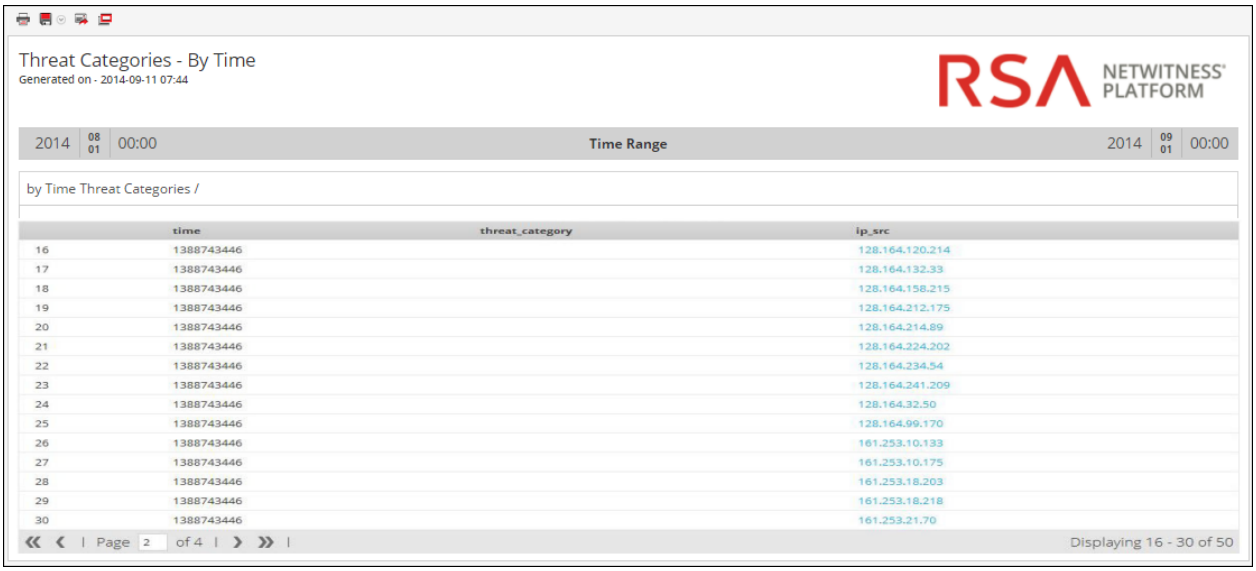
Column Name	Sort By
Enter the column name...	Ascending

Limit: 20

Use | Save | Reset | Test Rule

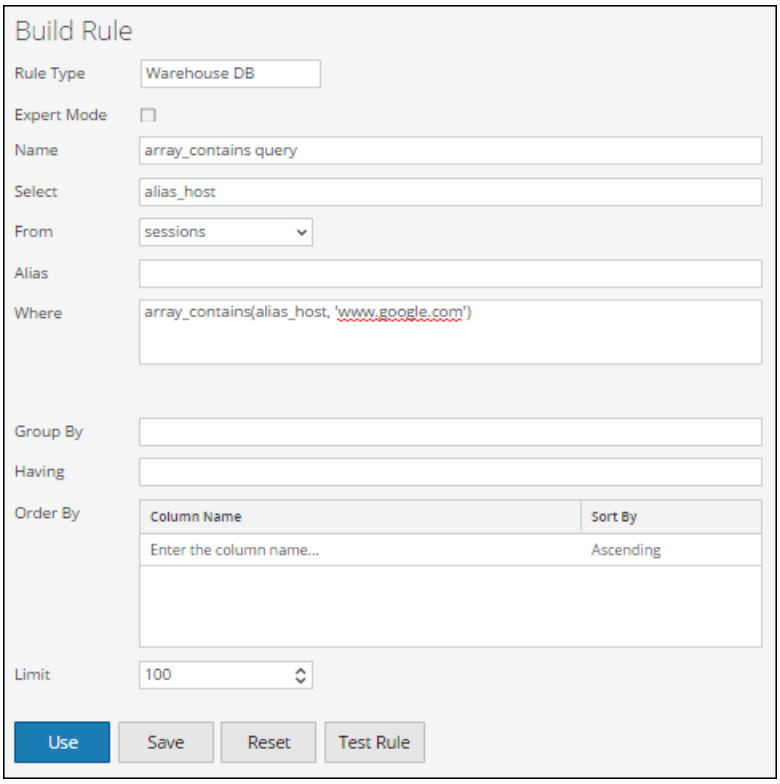
The following figure shows the result set of the by Time Threat Categories rule. The time displayed in the time field is the UNIX time (For example, 1388743446).

Note: In the “Select” clause the syntax would be “UNIX time” to convert to UTC time in report. For example, you can use the Epoch time converter tool to convert UNIX time (1388743446) to UTC (Coordinated Universal Time) (1/3/2014 3:34:06 PM).

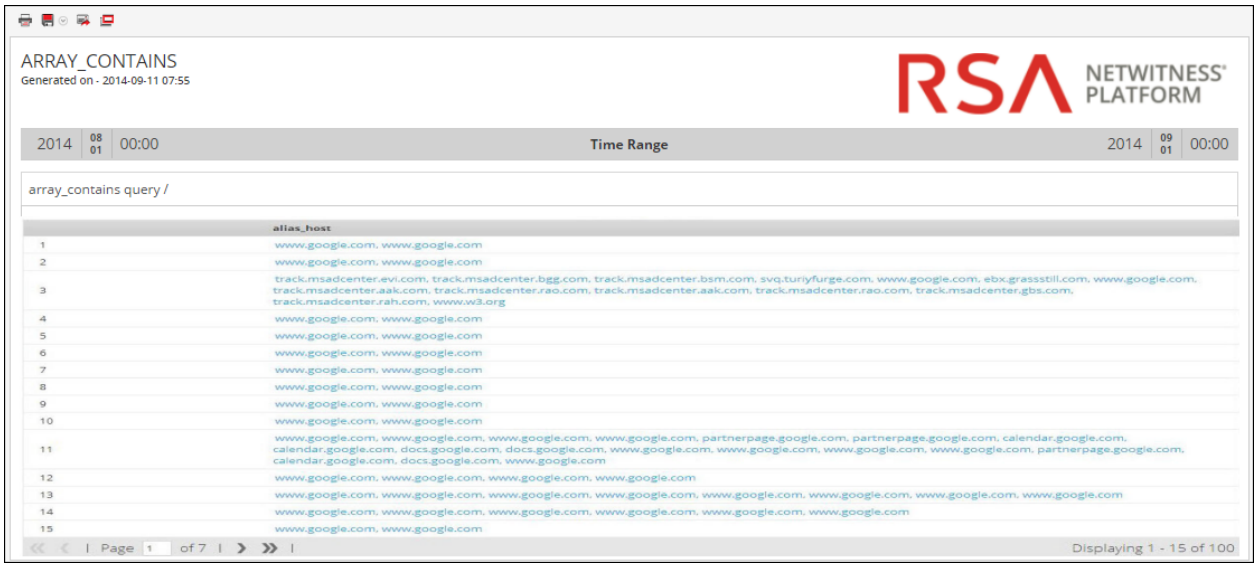


Array Query Report

This rule fetches an array of alias host names from the **sessions** table which contains the value 'www.google.com'.



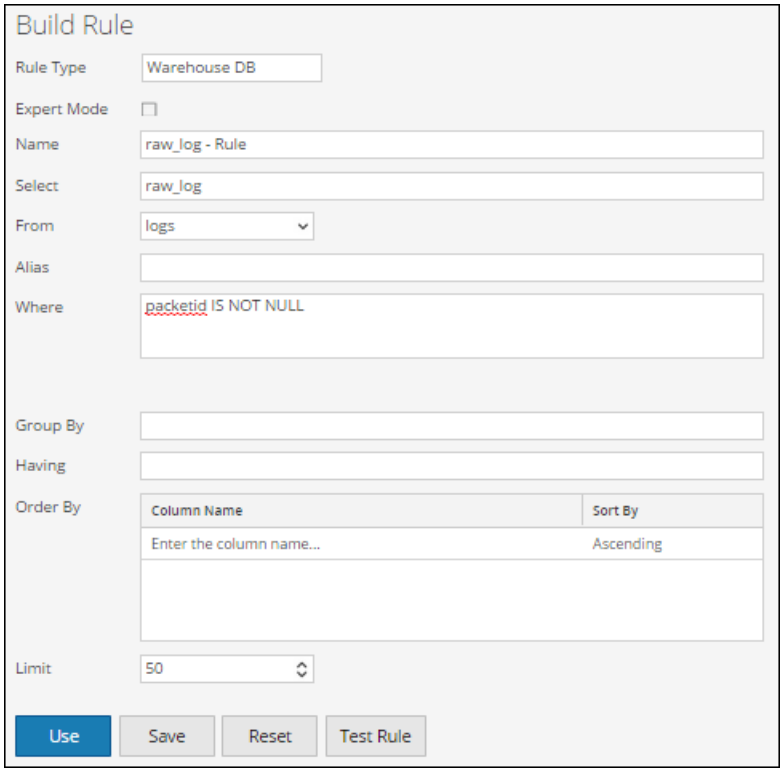
The following figure shows the result set for querying an array from sessions.



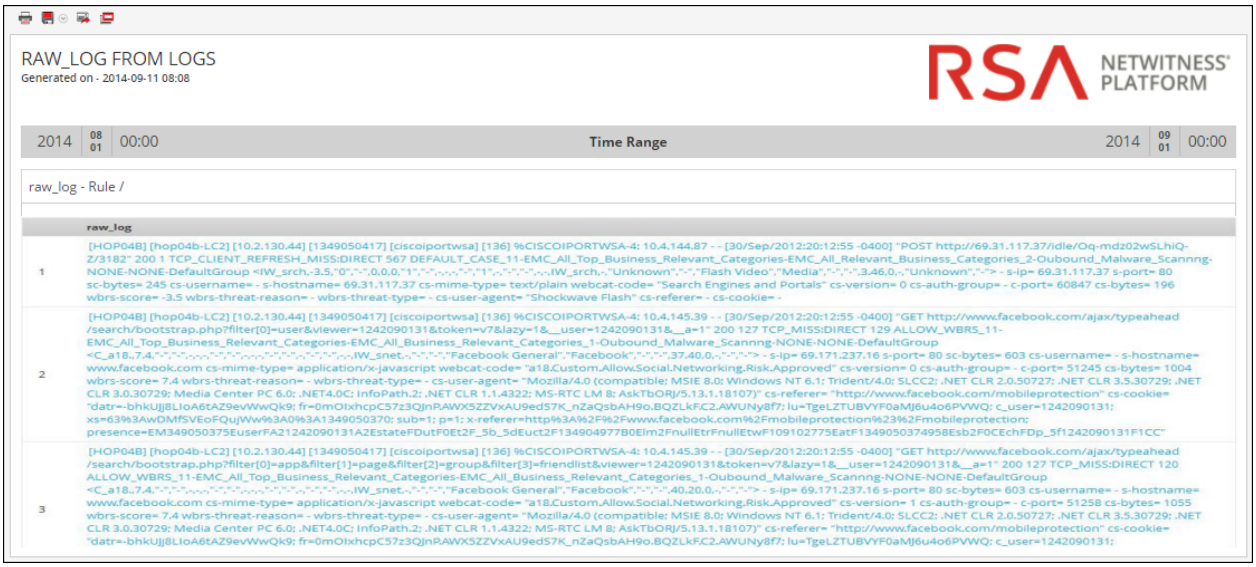
Raw Log Query Report

Raw logs can be queried either from the logs or sessions table.

This rule uses **raw_log** as a meta for querying raw log from logs whose packet ID is NOT NULL.



The following figure shows the result set for querying raw logs from logs.



RAW_LOG FROM LOGS
Generated on - 2014-09-11 08:08

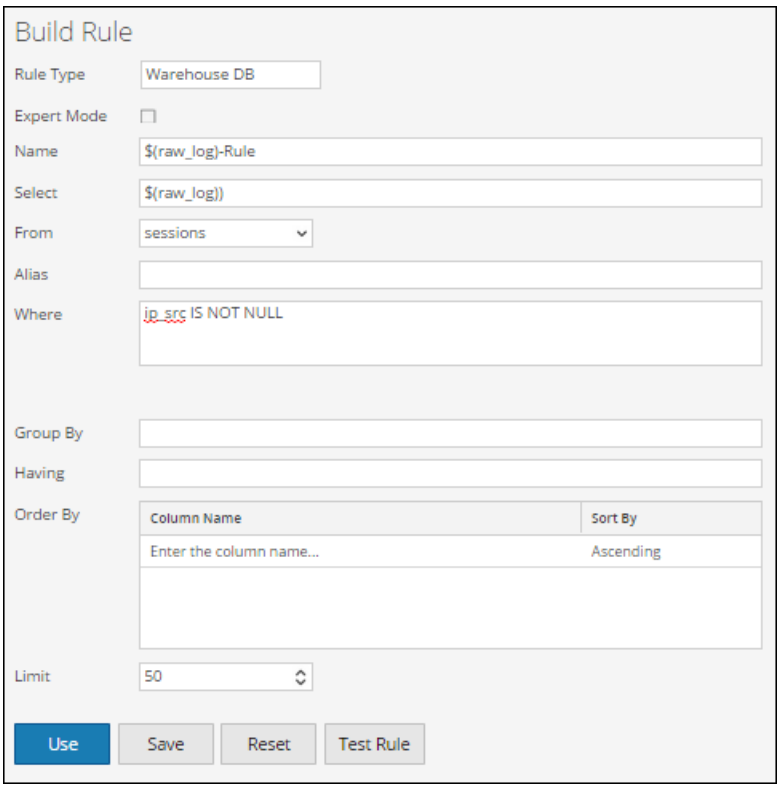
RSA NETWITNESS PLATFORM

2014	08 01	00:00	Time Range	2014	09 01	00:00

raw_log - Rule /

	raw_log
1	[HOP04B] [hop04b-LC2] [10.2.130.44] [1349050417] [ciscoipportwsa] [136] %CISCOIPORTWSA-4: 10.4.144.87 - - [30/Sep/2012:20:12:55 -0400] "POST http://69.31.117.37/idle/Oq-mdz02wSLhiQ-Z/3182" 200 1 TCP_CLIENT_REFRESH_MISS:DIRECT 567 DEFAULT_CASE_11-EMC_All_Top_Business_Relevant_Categories-EMC_All_Relevant_Business_Categories_2-Oubound_Malware_Scannng-NONE-NONE-DefaultGroup <IW_srch,-3.5,"0","-",0,0,0,"1","-",-,-,-,"-","1",-,"-","-","-",-,-,IW_srch,-,"Unknown","-","Flash Video","Media","-","-",3.46.0,-,"Unknown","-"> - s-ip= 69.31.117.37 s-port= 80 sc-bytes= 245 cs-username= - s-hostname= 69.31.117.37 cs-mime-type= text/plain webcat-code= "Search Engines and Portals" cs-version= 0 cs-auth-group= - c-port= 60847 cs-bytes= 196 wbrs-score= -3.5 wbrs-threat-reason= - wbrs-threat-type= - cs-user-agent= "Shockwave Flash" cs-referer= - cs-cookie= -
2	[HOP04B] [hop04b-LC2] [10.2.130.44] [1349050417] [ciscoipportwsa] [136] %CISCOIPORTWSA-4: 10.4.145.39 - - [30/Sep/2012:20:12:55 -0400] "GET http://www.facebook.com/ajax/typeahead/search/bootstrap.php?filter[0]=user&viewer=1242090131&token=v7&lazy=1&__user=1242090131&__a=1" 200 127 TCP_MISS:DIRECT 129 ALLOW_WBRS_11-EMC_All_Top_Business_Relevant_Categories-EMC_All_Business_Relevant_Categories_1-Oubound_Malware_Scannng-NONE-NONE-DefaultGroup <C_a18.,7.4,"-","-",-,-,-,"-","-",-,-,"-","-",-,-,IW_snet,-,"-","-","Facebook General","Facebook","-","-",37.40.0,-,"-","-"> - s-ip= 69.171.237.16 s-port= 80 sc-bytes= 603 cs-username= - s-hostname= www.facebook.com cs-mime-type= application/x-javascript webcat-code= "a18.Custom.Allow.Social.Networking.Risk.Approved" cs-version= 0 cs-auth-group= - c-port= 51245 cs-bytes= 1004 wbrs-score= 7.4 wbrs-threat-reason= - wbrs-threat-type= - cs-user-agent= "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; Trident/4.0; SLCC2; .NET CLR 2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729; Media Center PC 6.0; .NET4.0C; InfoPath.2; .NET CLR 1.1.4322; MS-RTC LM 8; AskTbORJ/5.13.1.18107)" cs-referer= "http://www.facebook.com/mobileprotection" cs-cookie= "datr=-bhkUJj8LIoA6tAZ9evWwQk9; fr=0mOIxhcpC57z3QJnP.AWX5ZZVxAU9ed57K_nZaQsbAH9o.BQZLkF.C2.AWUNy8f7; lu=TgeLZTUBVYF0aMJ6u4o6PVWQ; c_user=1242090131; xs=63%3AwDMfSVEoFQujWw%3A0%3A1349050370; sub=1; p=1; x-referer=http%3A%2F%2Fwww.facebook.com%2Fmobileprotection%23%2Fmobileprotection: presence=EM349050375EuserFA21242090131A2EstateFDutF0Et2F_5b_5dEuct2F134904977B0Elm2FnullEtrFnullEtwF109102775EatF1349050374958Esb2F0CEchFDp_5f1242090131F1CC"
3	[HOP04B] [hop04b-LC2] [10.2.130.44] [1349050417] [ciscoipportwsa] [136] %CISCOIPORTWSA-4: 10.4.145.39 - - [30/Sep/2012:20:12:55 -0400] "GET http://www.facebook.com/ajax/typeahead/search/bootstrap.php?filter[0]=app&filter[1]=page&filter[2]=group&filter[3]=friendlist&viewer=1242090131&token=v7&lazy=1&__user=1242090131&__a=1" 200 127 TCP_MISS:DIRECT 120 ALLOW_WBRS_11-EMC_All_Top_Business_Relevant_Categories-EMC_All_Business_Relevant_Categories_1-Oubound_Malware_Scannng-NONE-NONE-DefaultGroup <C_a18.,7.4,"-","-",-,-,-,"-","-",-,-,"-","-",-,-,IW_snet,-,"-","-","Facebook General","Facebook","-","-",40.20.0,-,"-","-"> - s-ip= 69.171.237.16 s-port= 80 sc-bytes= 603 cs-username= - s-hostname= www.facebook.com cs-mime-type= application/x-javascript webcat-code= "a18.Custom.Allow.Social.Networking.Risk.Approved" cs-version= 1 cs-auth-group= - c-port= 51258 cs-bytes= 1055 wbrs-score= 7.4 wbrs-threat-reason= - wbrs-threat-type= - cs-user-agent= "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; Trident/4.0; SLCC2; .NET CLR 2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729; Media Center PC 6.0; .NET4.0C; InfoPath.2; .NET CLR 1.1.4322; MS-RTC LM 8; AskTbORJ/5.13.1.18107)" cs-referer= "http://www.facebook.com/mobileprotection" cs-cookie= "datr=-bhkUJj8LIoA6tAZ9evWwQk9; fr=0mOIxhcpC57z3QJnP.AWX5ZZVxAU9ed57K_nZaQsbAH9o.BQZLkF.C2.AWUNy8f7; lu=TgeLZTUBVYF0aMJ6u4o6PVWQ; c_user=1242090131;

This rule uses **${raw_log}** as a meta for querying raw log from sessions whose source IP address is NOT NULL.



Build Rule

Rule Type: Warehouse DB

Expert Mode: ☐

Name: $(raw_log)-Rule

Select: $(raw_log))

From: sessions

Alias:

Where: ip_src IS NOT NULL

Group By:

Having:

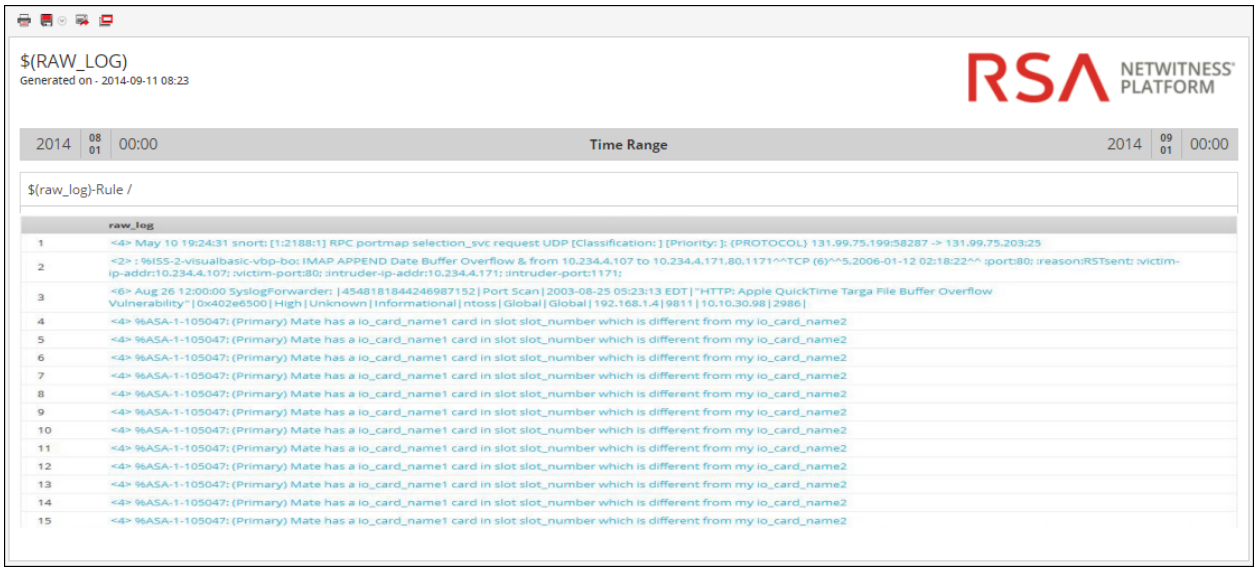
Order By:

Column Name	Sort By
Enter the column name...	Ascending

Limit: 50

Use | Save | Reset | Test Rule

The following figure shows the result set for querying raw logs from sessions.



$(RAW_LOG)
Generated on - 2014-09-11 08:23

RSA NETWITNESS PLATFORM

2014	08 01	00:00	Time Range	2014	09 01	00:00

$(raw_log)-Rule /

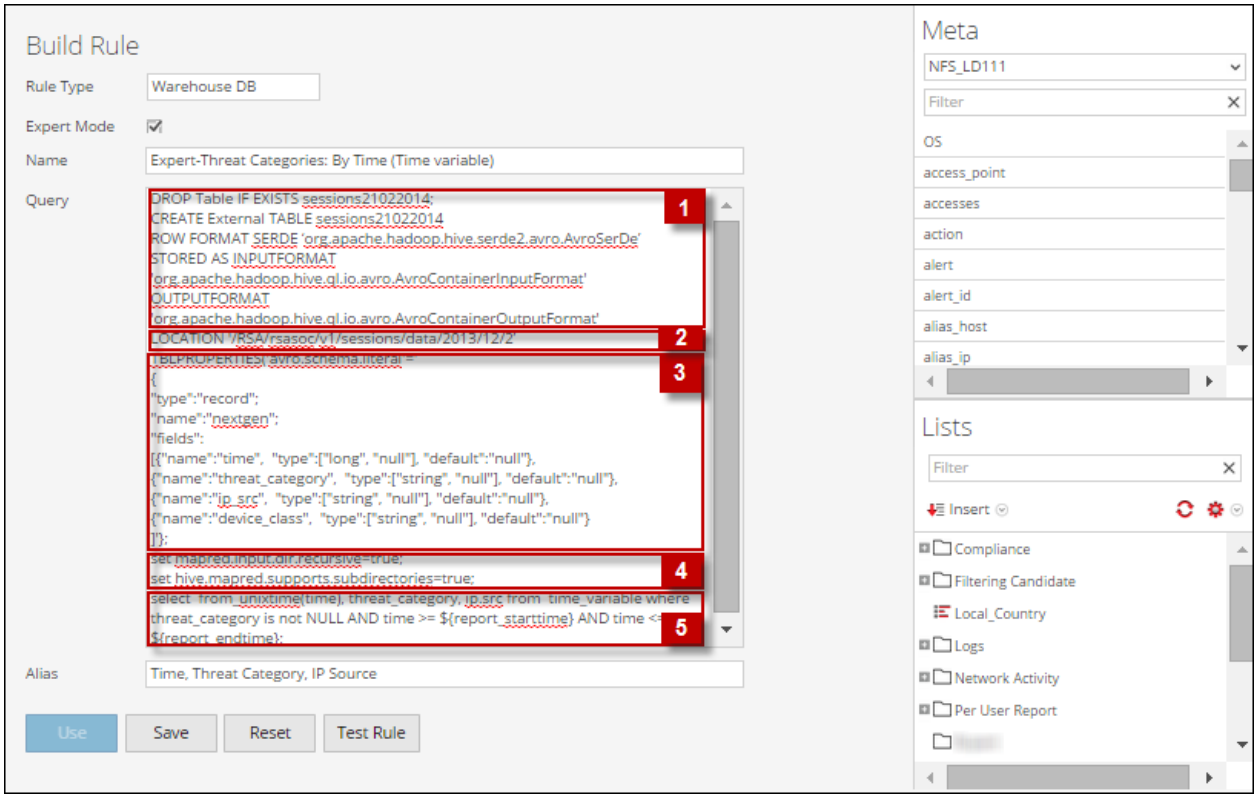
	raw_log
1	<4> May 10 19:24:31 snort: [1:2188:1] RPC portmap selection_svc request UDP [Classification:] [Priority:]: {PROTOCOL} 131.99.75.199:58287 -> 131.99.75.203:25
2	<2> : %IS5-2-visualbasic-vbp-bo: IMAP APPEND Date Buffer Overflow & from 10.234.4.107 to 10.234.4.171.80.1171^^TCP (6)^^5,2006-01-12 02:18:22^^ :port:80; :reason:RSTsent; :victim-ip-addr:10.234.4.107; :victim-port:80; :intruder-ip-addr:10.234.4.171; :intruder-port:1171;
3	<6> Aug 26 12:00:00 SyslogForwarder: |4548181844246987152| Port Scan |2003-08-25 05:23:13 EDT| "HTTP: Apple QuickTime Targa File Buffer Overflow Vulnerability"|0x402e6500|High|Unknown|Informational|ntoss|Global|Global|192.168.1.4|9811|10.10.30.98|2986|
4	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
5	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
6	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
7	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
8	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
9	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
10	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
11	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
12	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
13	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
14	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2
15	<4> %ASA-1-105047: (Primary) Mate has a io_card_name1 card in slot slot_number which is different from my io_card_name2

Warehouse DB Advanced Rules Syntax

The section explains the advanced rules query syntax and examples.

General Syntax of an Advanced Rule

The following figure shows how to define an advanced query.



The following syntax is an example of an advanced query:

```
DROP Table IF EXISTS sessions21022014;
CREATE External TABLE sessions21022014
ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.avro.AvroSerDe'
STORED AS INPUTFORMAT
'org.apache.hadoop.hive.ql.io.avro.AvroContainerInputFormat'
 OUTPUTFORMAT 'org.apache.hadoop.hive.ql.io.avro.AvroContainerOutputFormat'
LOCATION '/RSA/rsasoc/v1/sessions/data/2013/12/2'
TBLPROPERTIES('avro.schema.literal'='
 {
"type":"record";
"name":"nextgen";
"fields":
 [
 {"name":"time",  "type":["long", "null"], "default":"null"},
{"name":"threat_category",  "type":["string", "null"], "default":"null"},
```

```
{"name":"ip_src",  "type":["string", "null"], "default":"null"},
{"name":"device_class",  "type":["string", "null"], "default":"null"}
 ]
 '};
set mapred.input.dir.recursive=true;
set hive.mapred.supports.subdirectories=true;
```

select `from_unixtime(time), threat_category, ip.src` from `time_variable` where `threat_category` is not `NULL` and `time >= ${report_starttime}` and `time <= ${report_endtime};`

Note: Reporting Engine treats a line beginning with <hyphen> <hyphen> as a comment in Expert Warehouse Rule.
For example,

```
set mapred.input.dir.recursive=true;
-- This is an Expert comment
 set hive.mapred.supports.subdirectories=true;
```

The general syntax of an advanced query is as explained below:

1. Drop and create an external table, and then format the row:
   Firstly, we drop the table, if the table already exists and create an external table **sessions21022014**

   ```
   DROP TABLE IF EXISTS sessions21022014
    CREATE EXTERNAL TABLE sessions21022014
   ```

   Note: You must create an external table only if you are using an other table. For example, if you are using an other table apart from **sessions21022014** then you must drop the table and create an external table.

   Then, specify the row format as Avro.SerDe interface to instruct HIVE as to how a record is to be processed. Avro.SerDe allows you to read or write Avro data as HIVE tables and store them as input format and output format.

   ```
   ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.avro.Avro.SerDe'
    STORED AS INPUTFORMAT
   'org.apache.hadoop.hive.ql.io.avro.AvroContainerInputFormat'
    OUTPUTFORMAT 'org.apache.hadoop.hive.ql.io.avro.AvroContainerOutputFormat'
   ```

2. Specify the HDFS location:
   Secondly, you must specify the HDFS location '/RSA/rsasoc/v1/sessions/data/2013/12/2' from where the data is queried before executing the HIVE statements.The location parameter specifies the data to be fetched depending on the date input provided. This is a variable parameter hence you can fetch values depending on the date entered.

3. Define the table schema:
   Thirdly, you define the table schema by defining columns with a specific data type and default value as 'null'.

   ```
   TBLPROPERTIES('avro.schema.literal'='
    {"type":"record";
    "name":"nextgen";
    "fields":
    [
    {"name":"ip_src",  "type":["string", "null"], "default":"null"}
    ]
    '};
   ```

4. Import data from directory which contains sub directories:
   Then, you must enable HIVE to recursively scan all sub-directories and fetch all the data from all sub-directories.
   ```
   set mapred.input.dir.recursive=true;
   set hive.mapred.supports.subdirectories=true;
   ```
5. Fetch data from the HIVE table:
   Once you execute all the above statements, you can query the database with the HIVE query **select** clause to fetch the data from the HIVE table.

The following examples illustrate advanced rules in the expert mode:

- Hourly, daily, weekly, and monthly report
- Table partition based on location report
- Join logs and sessions based on unique_id report
- List report
- Parameterized report
- Partition based table with multiple locations
- Automated partition using custom function (10.5.1 onwards)

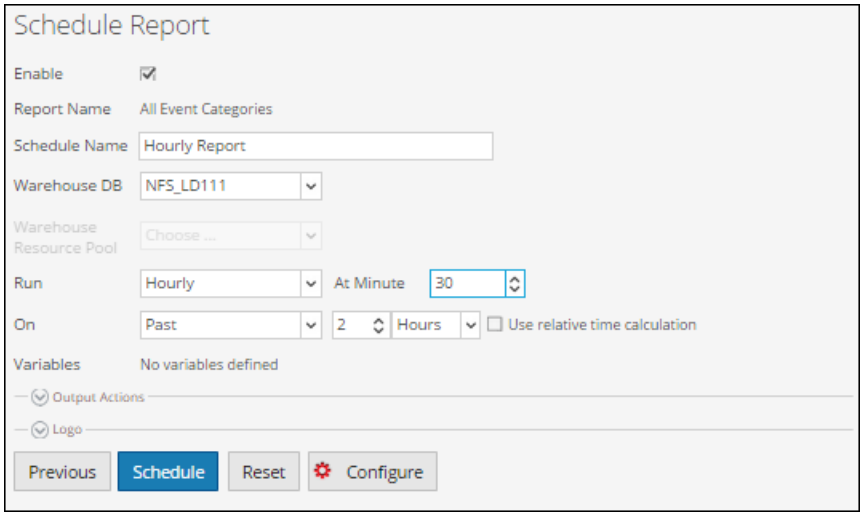
Hourly, Daily, Weekly, and Monthly Report

In these example rules, you can create various reports for December 02, 2013 (as in the below figure). The date variable in the LOCATION statement can be altered, depending on which you can create an hourly, daily, weekly, and monthly report.

Hourly Report

In this example rule, you can create an houly report for December 02, 2013. The LOCATION statement can be altered to generate an hourly report.

LOCATION '/RSA/rsasoc/v1/sessions/data/2013/12/2' - the date input (2013/12/2) indicates year/month/day. The entire data for 02 December, 2013 is retrieved using this location statement.

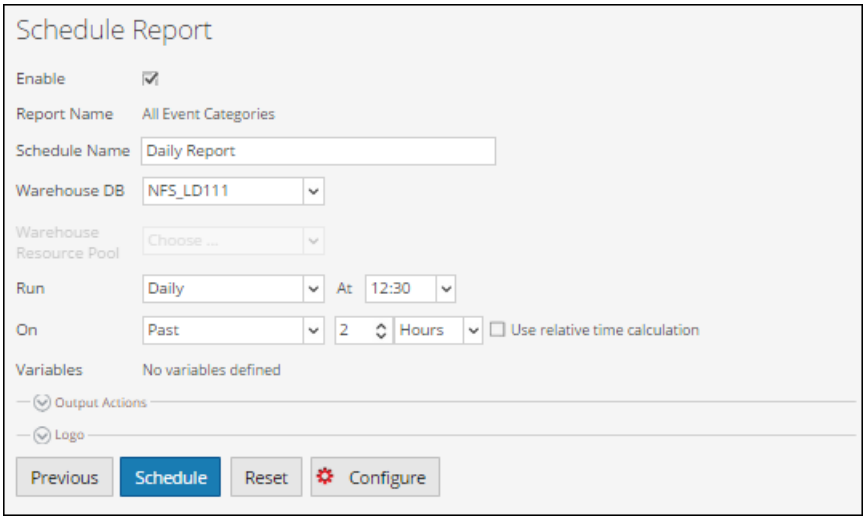


The result set of this query would be an hourly report.

Daily Report

In this example rule, you can create a daily report for December 2013. The LOCATION statement can be altered to generate a daily report.

LOCATION '/RSA/rsasoc/v1/sessions/data/2013/12' - the date input (2013/12) indicates year/month. The entire data for December, 2013 is retrieved using this location statement.

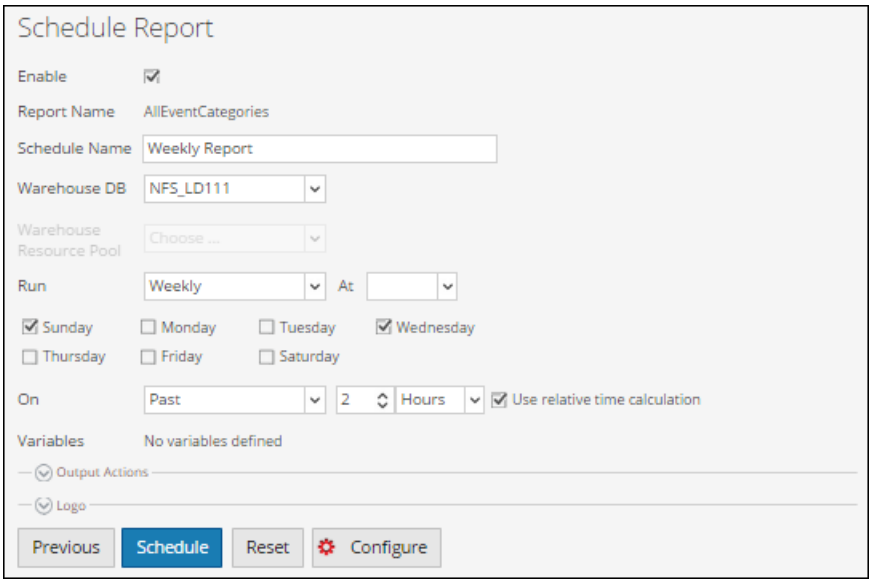


The resultset of this query would be a daily report.

Weekly Report

In this example rule, you can create a weekly report for December 2013. The LOCATION statement can be altered to generate a weekly report.

LOCATION '/RSA/rsasoc/v1/sessions/data/2013/12' - the date input (2013/12) indicates year/month. The entire data for December, 2013 is retrieved using this location statement.

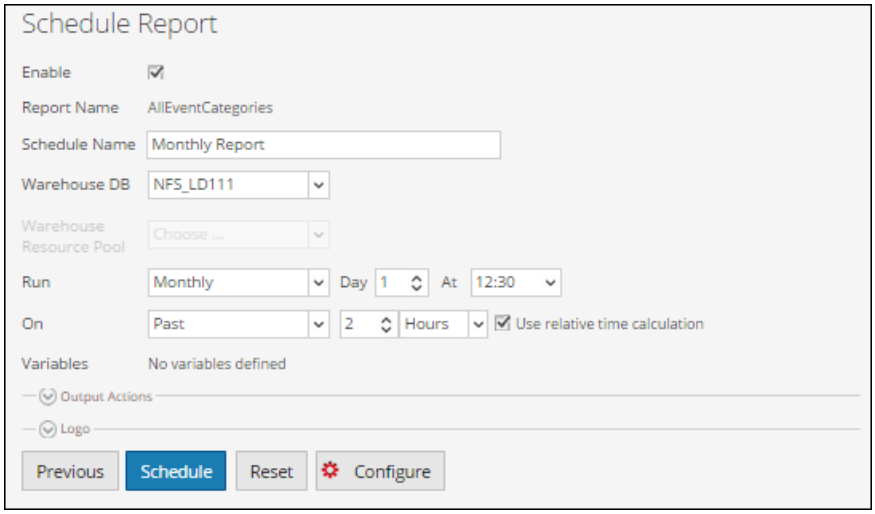


The result set of this query would be a weekly report.

Monthly Report

In this example rule, you can create a monthly report for the year 2013. The LOCATION statement can be altered to generate a monthly report.

LOCATION '/RSA/rsasoc/v1/sessions/data/2013' - the date input (2013) indicates year. The entire data for the year 2013 is retrieved using this location statement.



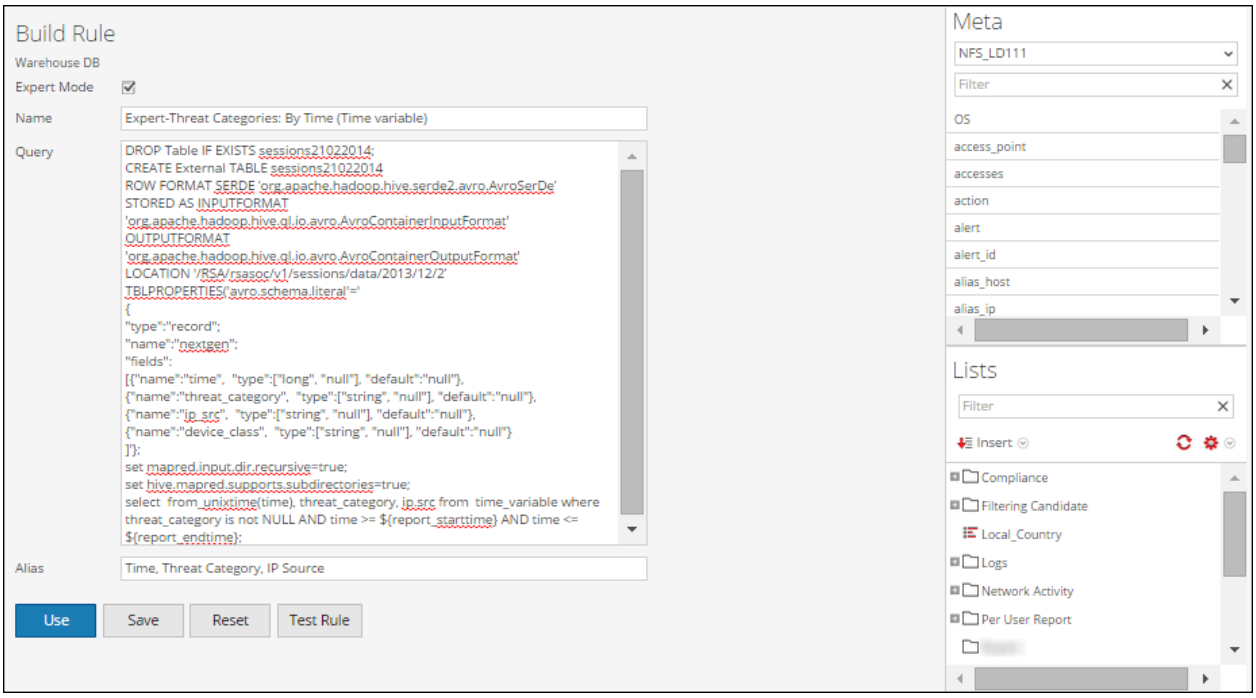
The result set of this query would be a monthly report.

For more information on LOCATION definition, see **Specify the HDFS location** in the **"General Syntax of an Advanced Rule"** section.

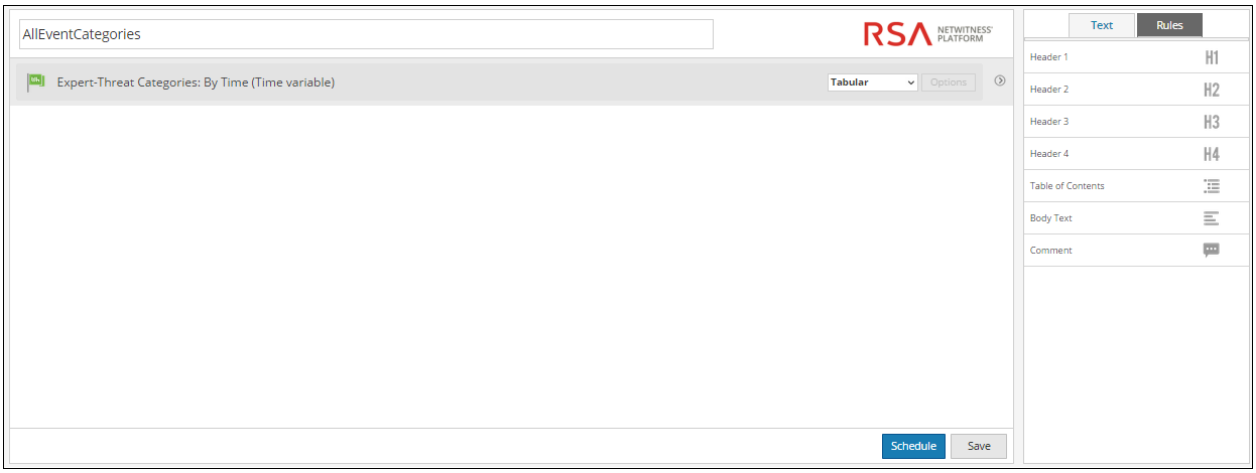
You must perform the following steps in sequence to view the resultset of an advanced rule:

1. Define an Advanced Rule
2. Add an advanced rule to a Report
3. Schedule a Report
4. View a scheduled Report

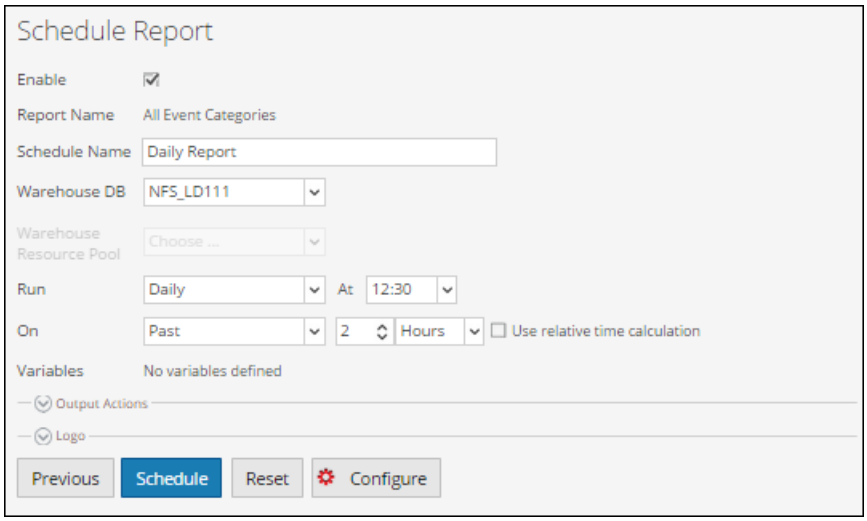
The following figure shows how to define an advanced rule.



The following figure shows how to add an advanced rule to a report (For example, **AllEventCategories**).



The following figure shows how to schedule a daily report.



If you want to generate a report for a specific time range, you need to manually define the time range in the query using the following two variables:

`${report_starttime} - The starting time of the range in seconds.`
`${report_endtime} - The ending time of the range in seconds.`

For example, `SELECT from_unixtime(time), threat_category, ip.src FROM time_variable WHERE threat_category is not NULL AND time >= ${report_starttime} AND time <= ${report_endtime};`

The following figure shows the result set of scheduling a daily report.

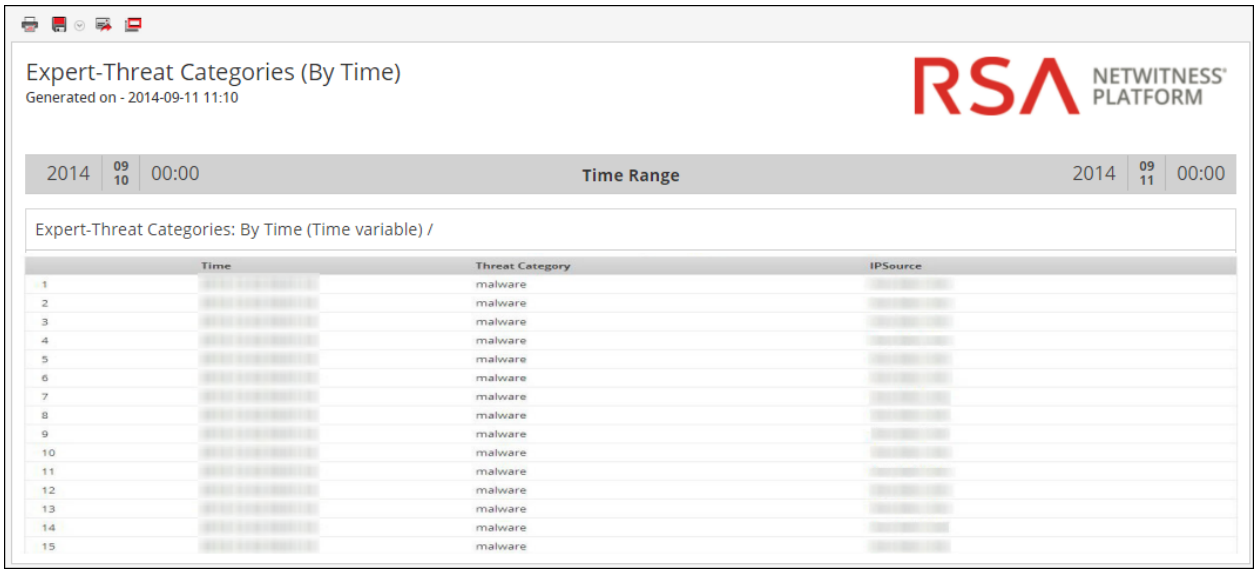
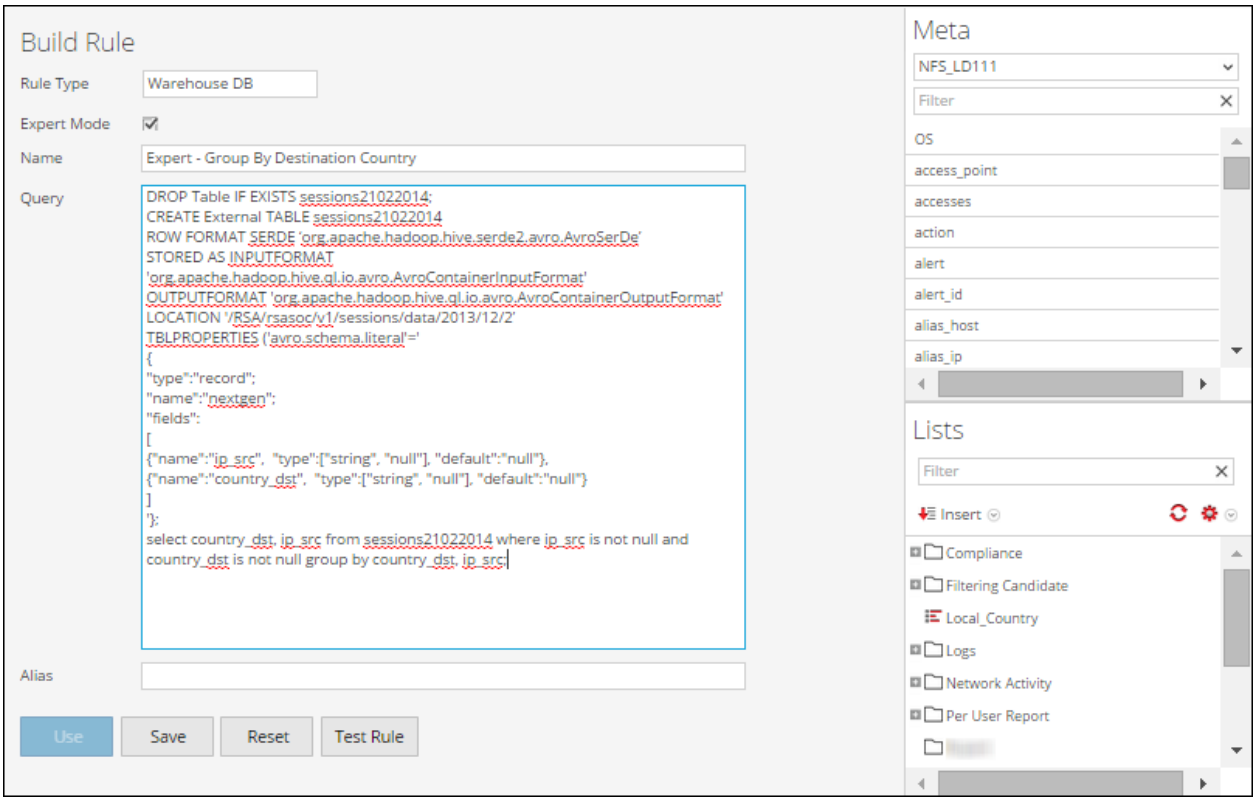


Table Partition Based on Location Report

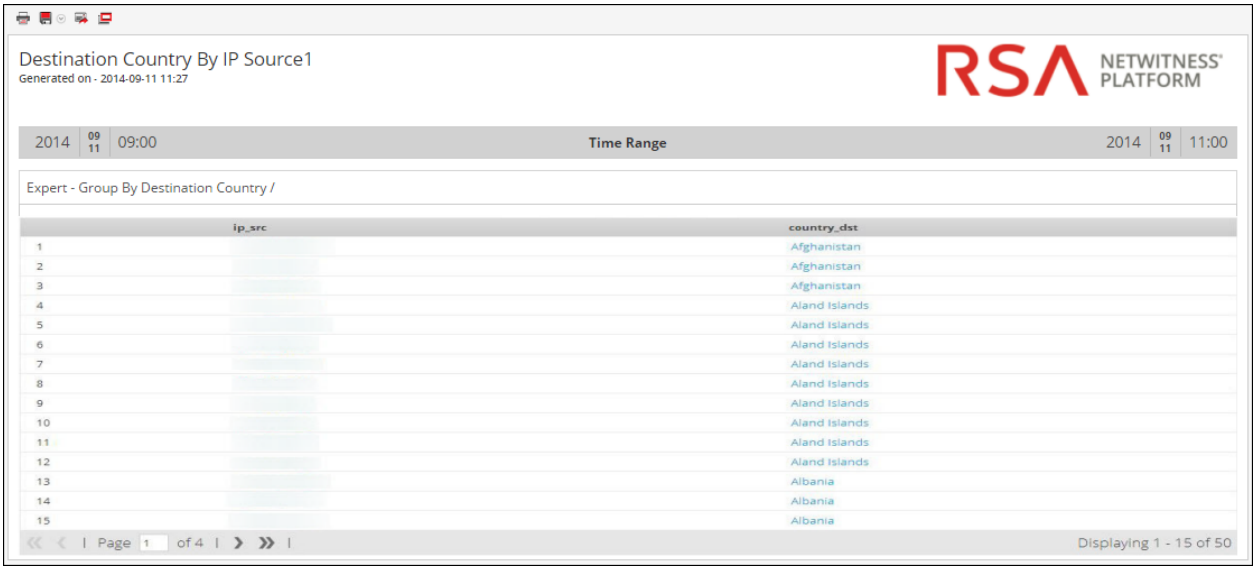
In this example rule, you can create a table partition based on location. Each table can have one or more partition keys which determines how the data is stored. For example, a country_dst of type STRING and an ip_src of type STRING. Each unique value of the partition keys defines a partition of the table.

In the example provided, we execute a HIVE query to fetch destination country and IP address of source from the sessions05032014 table and group the result set by these fields.

This rule provides information about the table created, row formatted, location (directory path) for avro data files in Warehouse, and returns a result set as per the HIVE query to indicate that the query returned a result set. For more information on these statements, see "General Syntax of an Advanced Rule" section.



The following figure shows the result set of creating a table partition based on location report.

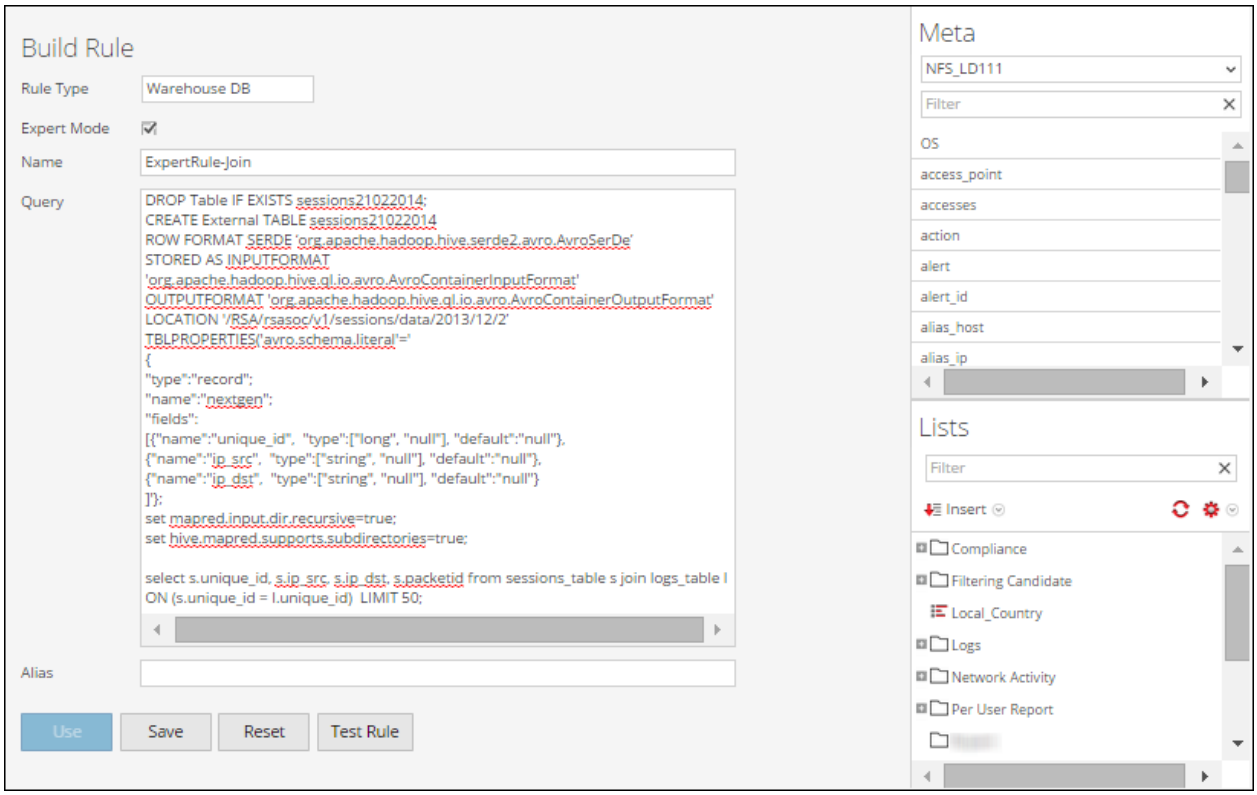


Join Logs and Sessions Based on unique_id Report

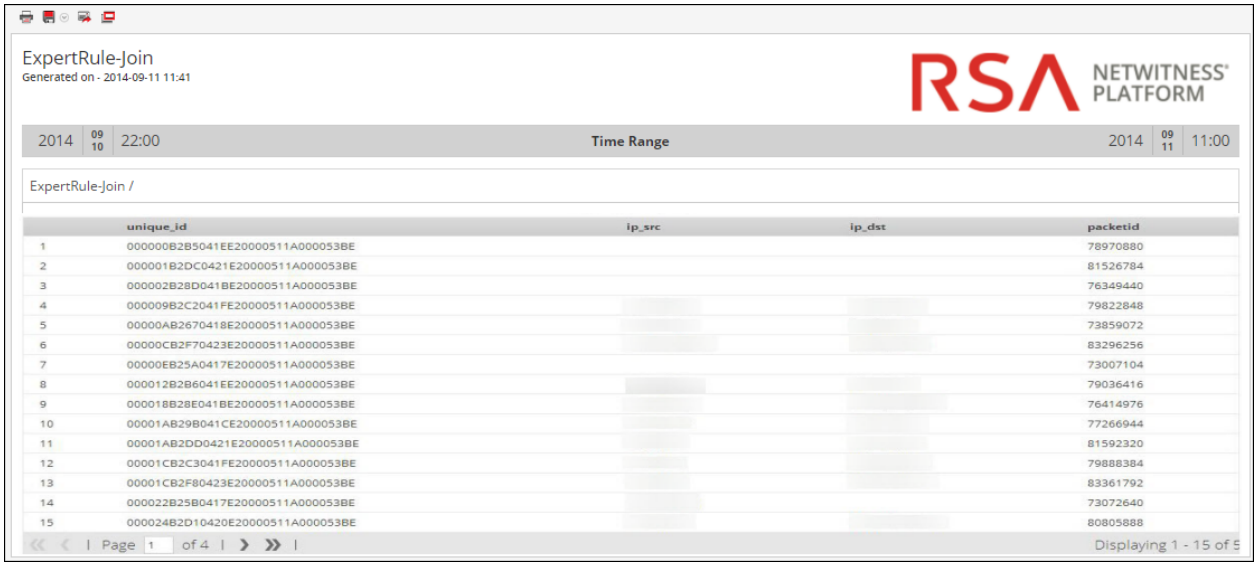
In this example rule, you can create a rule to join logs and sessions table to fetch unique_id, IP address of source and destination, and packet ID based on unique_id.

In the example provided, we execute a HIVE query to fetch certain fields from both the sessions_table and logs_table by performing a join based on the 'unique_id' field.

This rule provides information about the table created, row formatted, location (directory path) for avro data files in Warehouse, and returns a result set as per the HIVE query to indicate that the query returned a result set. For more information on these statements, see the **"General Syntax of an Advanced Rule"** section.



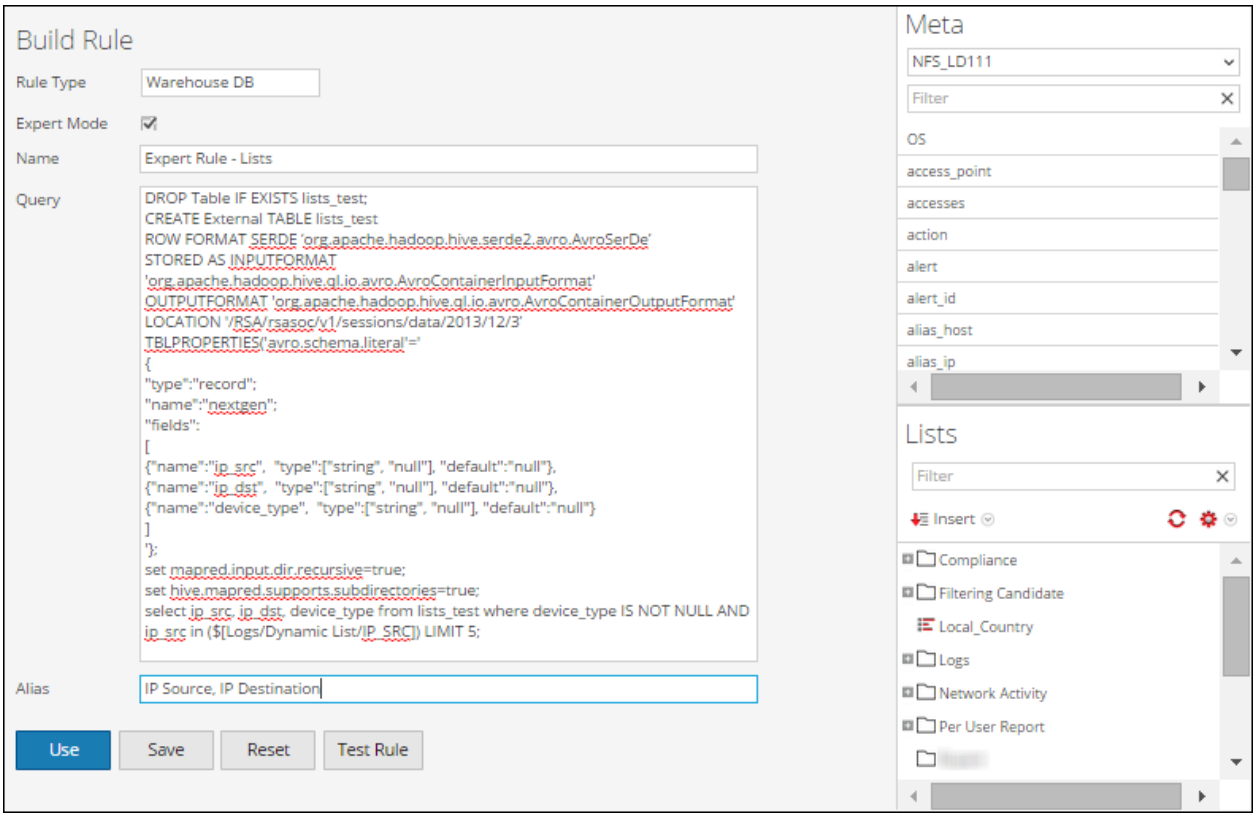
The following figure shows the result set of joining logs and sessions table based on unique_id.



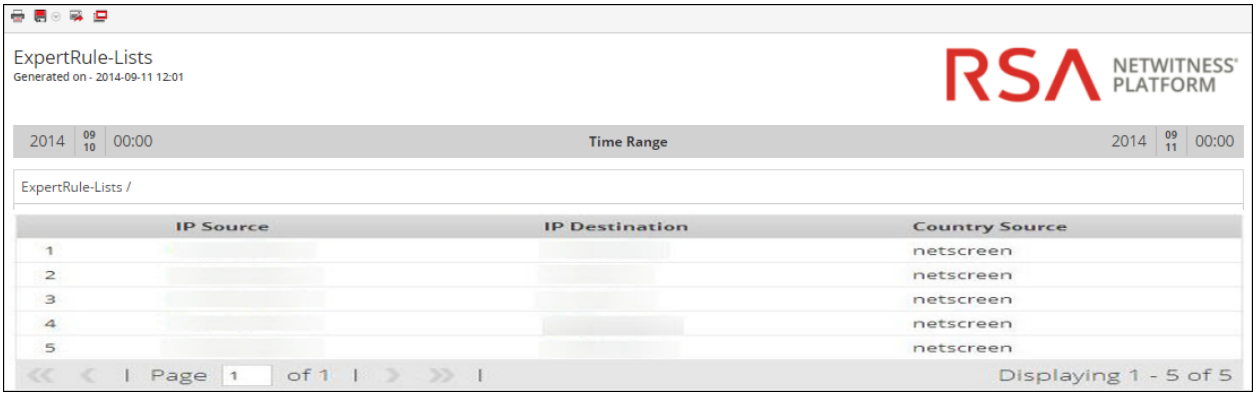
List Report

In this example rule, you can create a List report to fetch IP address of source and destination, and device type from the **lists_test** table where device type is not null and IP address of source is fetched from the appropriate event list.

This rule provides information about the table created, row formatted, location (directory path) for avro data files in Warehouse, and returns a result set as per the HIVE query to indicate that the query returned a result set. For more information on these statements, see the "General Syntax of an Advanced Rule" section.



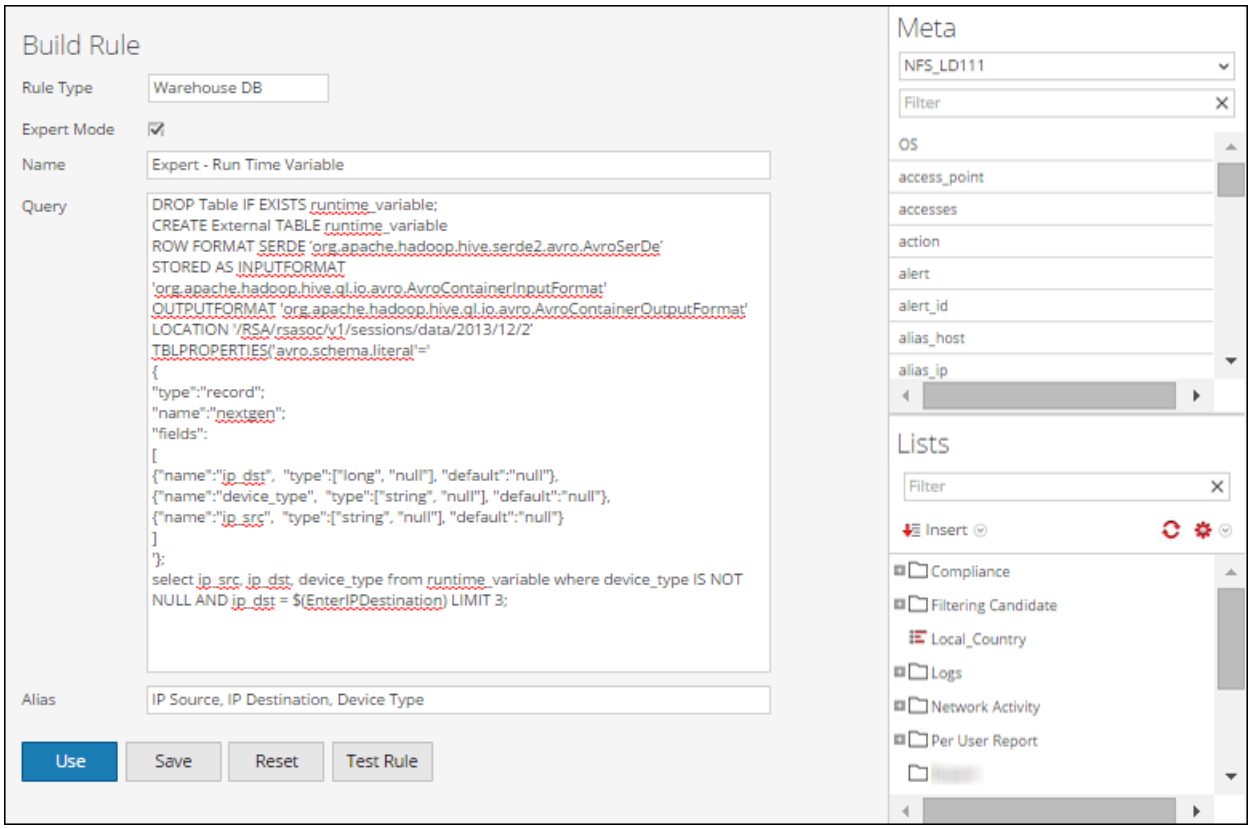
The following figure shows the result set of executing a list report.



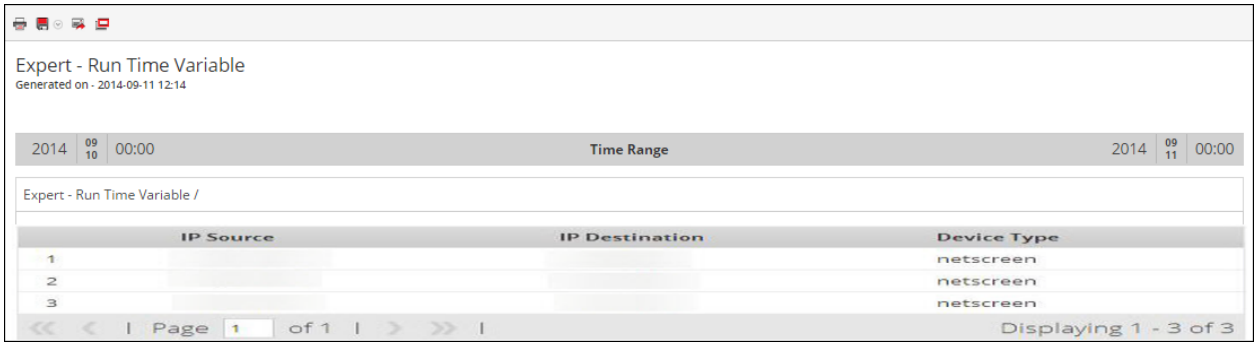
Parameterized Report

In this example rule, you can create a rule to fetch IP addresses of source and destination, and device type from the **runtime_variable** table based on the specified run time variable `${EnterIPDestination}`. At run time, you are prompted to enter a value for the IP address of destination ip_dst. Based on the value entered, the result set is displayed.

This rule provides information about the table created, row formatted, location (directory path) for avro data files in Warehouse, and returns a result set as per the HIVE query to indicate that the query returned a result set. For more information on these statements, see the "General Syntax of an Advanced Rule" section.



The following figure shows the result set of executing a parameterized report.



Partition Based Table with Multiple Locations

The following is an example of partition based table with multiple locations:

```
set mapred.input.dir.recursive=true;
 set hive.mapred.supports.subdirectories=true;
 DROP TABLE IF EXISTS AVRO_COUNT;
 CREATE EXTERNAL TABLE AVRO_COUNT
 PARTITIONED BY (partition_id int)
 ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.avro.AvroSerDe'
 WITH SERDEPROPERTIES (
 'avro.schema.literal'='{
 "name": "my_record", "type": "record",
 "fields": [
 {"name":"sessionid", "type":["null", "long"], "default" : null},
 {"name":"time", "type":["null", "long"], "default" : null}
 ]}'
 )
 STORED AS
 INPUTFORMAT
 'org.apache.hadoop.hive.ql.io.avro.AvroContainerInputFormat'
 OUTPUTFORMAT
 'org.apache.hadoop.hive.ql.io.avro.AvroContainerOutputFormat';
 ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=0) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/8';
 ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=1) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/9';
 ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=2) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/10/';
 ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=3) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/11/';
 ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=4) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/12/';
 SELECT COUNT(*) as TOTAL FROM AVRO_COUNT WHERE time >= ${report_
starttime} AND time
 <= ${report_endtime};
```

The partition based table with multiple location is as explained below:

1. Enable HIVE to recursively scan all sub-directories and read all the data from the sub-directories.
   ```
   set mapred.input.dir.recursive=true;
    set hive.mapred.supports.subdirectories=true;
   ```
2. Drop and create an external table, and then format the rows:
   ```
    DROP TABLE IF EXISTS AVRO_COUNT;
    CREATE EXTERNAL TABLE AVRO_COUNT
    PARTITIONED BY (partition_id int)
    ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.avro.AvroSerDe'
    WITH SERDEPROPERTIES (
    'avro.schema.literal'='{
    "name": "my_record", "type": "record",
    "fields": [
    {"name":"sessionid", "type":["null", "long"], "default" : null},
    {"name":"time", "type":["null", "long"], "default" : null}
    ]}'
    )
   ```

```
STORED AS
INPUTFORMAT
'org.apache.hadoop.hive.ql.io.avro.AvroContainerInputFormat'
OUTPUTFORMAT
'org.apache.hadoop.hive.ql.io.avro.AvroContainerOutputFormat';
```

Note: You must create an external table only if you are using any other table. For example, if you are using any other table apart from **AVRO_COUNT** then you must drop the table and create an external table.

Note: Points to remember when you create a table:
- Dropping a 'non-external' table deletes the data.
- The table is partitioned on a single column called partition_id and this is the standard column for Reporting Engine.
- The default value of any column is null as the AVRO file may not contain the specified column.
- The column names should be in the lowercase as HIVE is case insensitive but AVRO is case sensitive.
- You must specify **avro.schema.literal** in the *SERDEPROPERTIES* .

For more information on the "rule syntax", refer to *Apache HIVE*.

3. Add partitions:
   Once you define a table, you must specify the HDFS locations from where the data needs to be queried before you execute the HIVE statements.The location parameter specifies the data to be fetched depending on the specified date. The data is spread across multiple locations or directories in HDFS. For each location you need to add a partition with unique values assigned to the partition column. The locations can be any directory in the HDFS

```
ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=0) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/8';
ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=1) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/9';
ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=2) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/10/';
ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=3) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/11/';
ALTER TABLE AVRO_COUNT ADD PARTITION(partition_id=4) LOCATION
'/rsasoc/v1/sessions/data/2015/07/22/12/';
```

   Note: HIVE reads each file in these locations as AVRO. Incase if there is a non-AVRO file available in one of these locations then the query may fail.

4. Run the query

```
SELECT COUNT(*) as TOTAL FROM AVRO_COUNT WHERE time >= ${report_
starttime} AND time
<= ${report_endtime};
```

   When a table is created, you can execute specific queries to filter the data. For example, after you create the table you can filter the data as shown in the below examples:
   Sessions with a specific Source IP Address:

```
SELECT * FROM AVRO_COUNT WHERE time >= ${report_starttime} AND time
<= ${report_endtime}  AND ip_src = '127.0.0.1';
```

Group by based on user destination:

```
SELECT * FROM AVRO_COUNT WHERE time >= ${report_starttime} AND time
<= ${report_endtime}  GROUP BY usr_dst;
```

Automated Partition using Custom function

In 10.5.1, you can use the custom function to automate the addition of partitions to a user defined table in the expert mode.

General syntax

```
RE WH CUSTOM ADDPARTITIONS(table, namespace, rollup, [starttime,
endtime])
```

The following table describes the custom function syntax:

S.No	Name	Description
1	table	The table name for which the partition has to be added.
2	namespace	The namespace can be sessions or logs.
3	rollup	This value determines the level of directory path to be included in partitions. The value can be HOUR, DAY, or MINUTE. If Warehouse Connector is configured for Day rollup, setting this value as HOUR produces ZERO results. The number and location of each partition is based on time range used to run the rule and the rollup value.
4	(Optional) starttime, endtime	To generate partitions for a specific time range other than the time range mentioned in the rule, you must specify the starttime and endtime in **Epoch Seconds**. **Note:** Expressions are not supported for the starttime and endtime.

The custom function is invoked when Reporting Engine executes the rule either during test rule or scheduled report. While running a expert rule, whenever Reporting Engine identifies the function declaration, it extracts the required arguments and insert *n* number of ADD PARTITION HiveQL statements and executes them on the Hive Server.

The location and directory structure is determined by the argument passed in the rule and the Hive data source configuration in Reporting Engine. The number of partitions depends on the rollup specified and the time range used while executing the rule. For example, with the rollup as HOUR and the time range as PAST 2 Days results in 48 partitions for 48 Hours while with the rollup as DAY, Reporting Engine creates 2 partitions, one for each day.

The partition query is generated by the Syntax Template as set in Reporting Engine's Hive Configuration attribute AlterTableTemplate.

Note: By default, this function starts adding partitions to a table with partition id from 0 to N-1. Hence this requires that the table must be partitioned by single integer column named partition id.

The following is an example of automated partition using custom function:

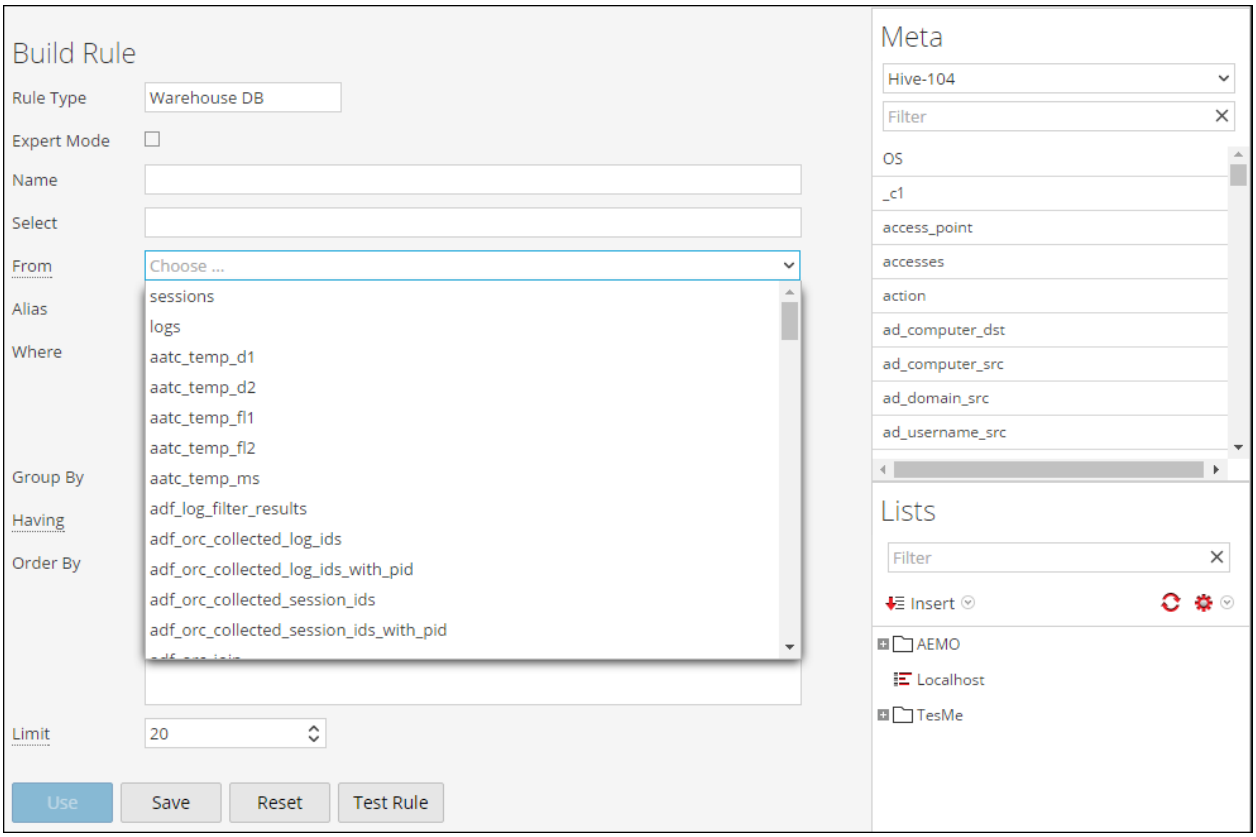
```
set mapred.input.dir.recursive=true;
 set hive.mapred.supports.subdirectories=true;
 DROP TABLE IF EXISTS AVRO_COUNT;

CREATE EXTERNAL TABLE AVRO_COUNT
 PARTITIONED BY (partition_id int)
 ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.avro.AvroSerDe'
 WITH SERDEPROPERTIES (
 'avro.schema.literal'='{
  "name": "my_record", "type": "record",
  "fields": [
  {"name":"sessionid", "type":["null", "long"], "default" : null}
  ,{"name":"time", "type":[ "null" , "long"], "default" : null}
  ,{"name":"unique_id", "type":["null", "string"], "default" : null}
  ]}'
 )
 STORED AS
 INPUTFORMAT
 'org.apache.hadoop.hive.ql.io.avro.AvroContainerInputFormat'
 OUTPUTFORMAT
 'org.apache.hadoop.hive.ql.io.avro.AvroContainerOutputFormat';

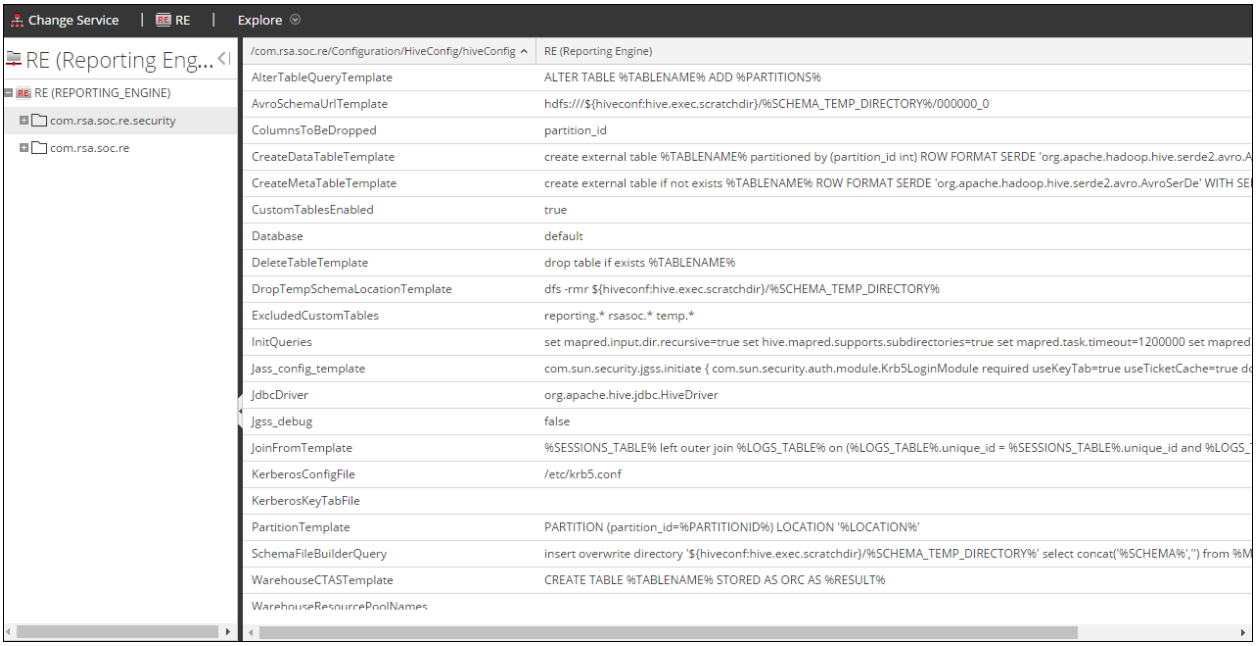
RE_WH_CUSTOM_ADDPARTITIONS(AVRO_COUNT, 'sessions', 'DAY');
 SELECT COUNT(*) as TotalSessions FROM AVRO_COUNT
 WHERE time >= ${report_starttime} AND time <= ${report_endtime};
```

Creating Custom Tables Report

In 10.6.1, you can use and create Custom Tables on the Hive Server. Reporting Engine supports running queries on user defined tables and the ability to create a new table from a Single Rule output. When this feature is enabled in the Warehouse Rule Builder UI, user can see a list of custom tables available in Hive Server.



To enable this feature set **customTablesEnabled** to **TRUE** by navigating to **Reporting Engine** -> **Explore** ->**Hive Config**.



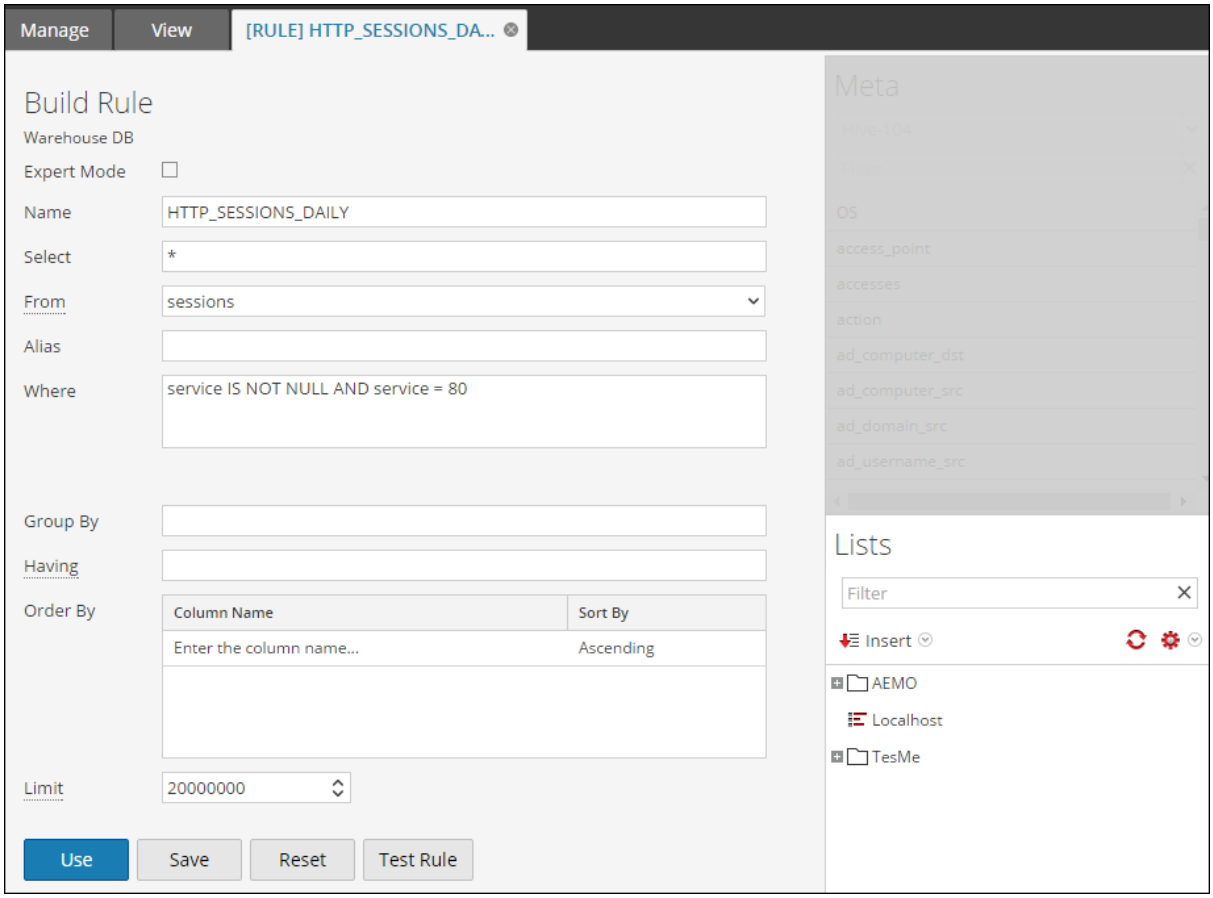
Creating Custom Table from Regular Rules

To schedule a report which contains a single SAW rule, a new text input with a **Warehouse CTAS Name** is added. The user can now specify a Custom Table name that will be created out of the output of the rule in Report.

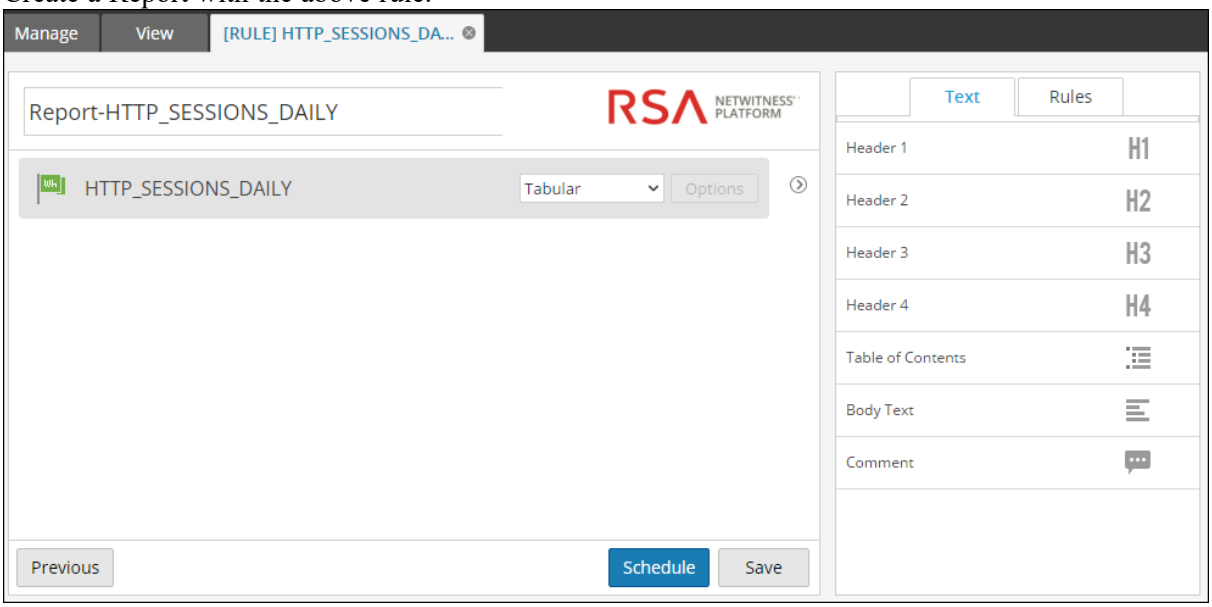
Note: This feature is available only if the Report contains a single SAW rule on the Schedule page. Otherwise, this option is hidden.

The process to use the feature is explained below:

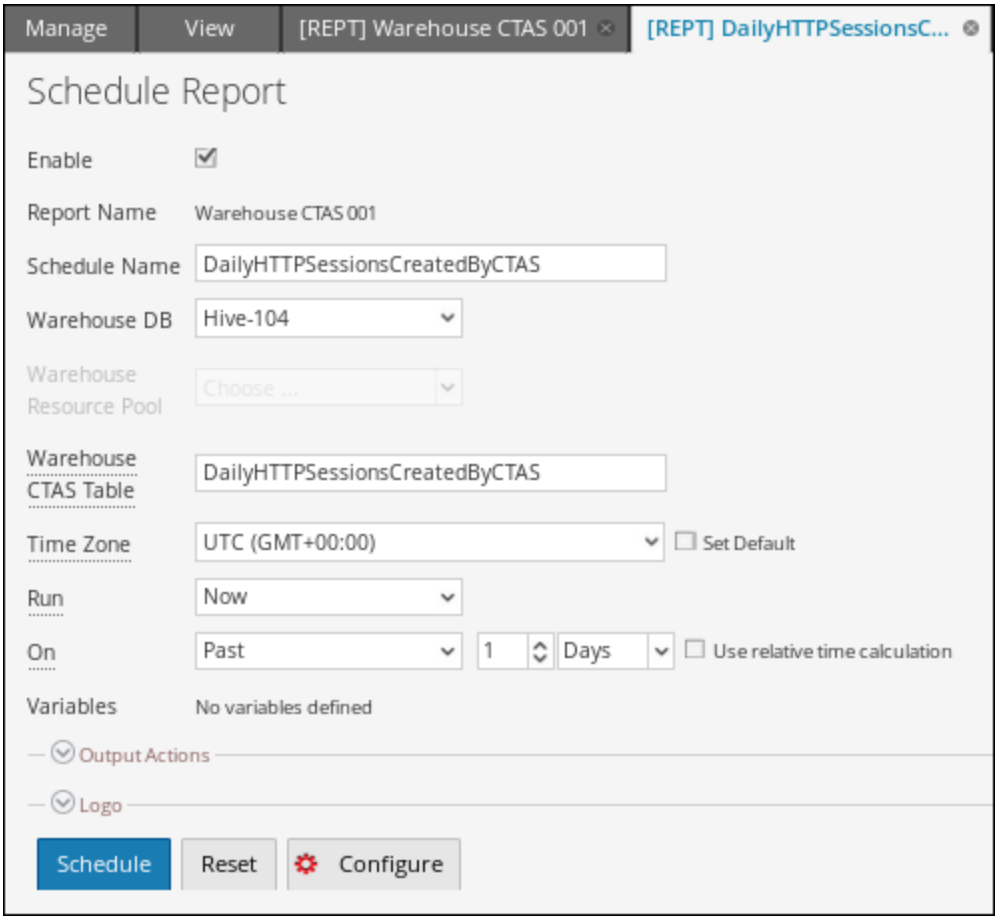
1. Create a rule to filter with data in SAW.



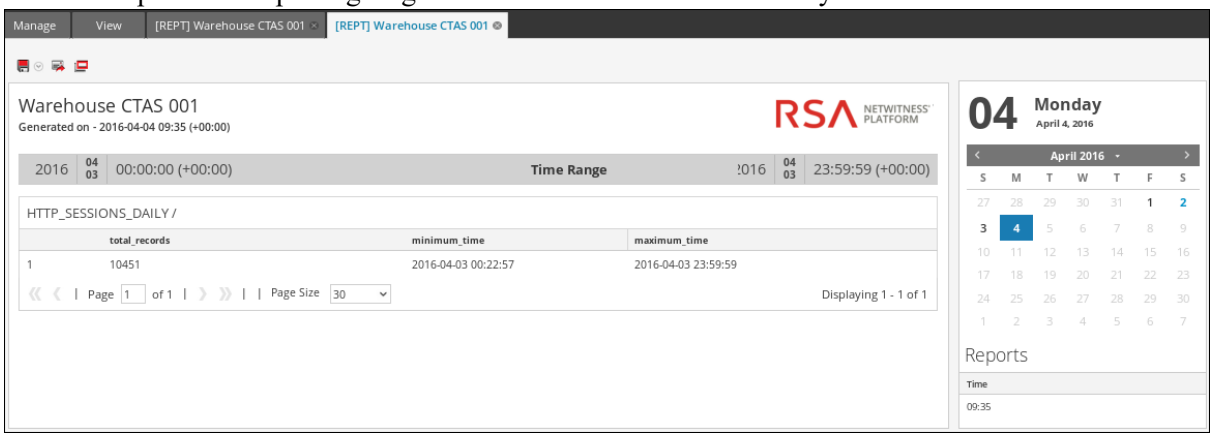
2. Create a Report with the above rule.



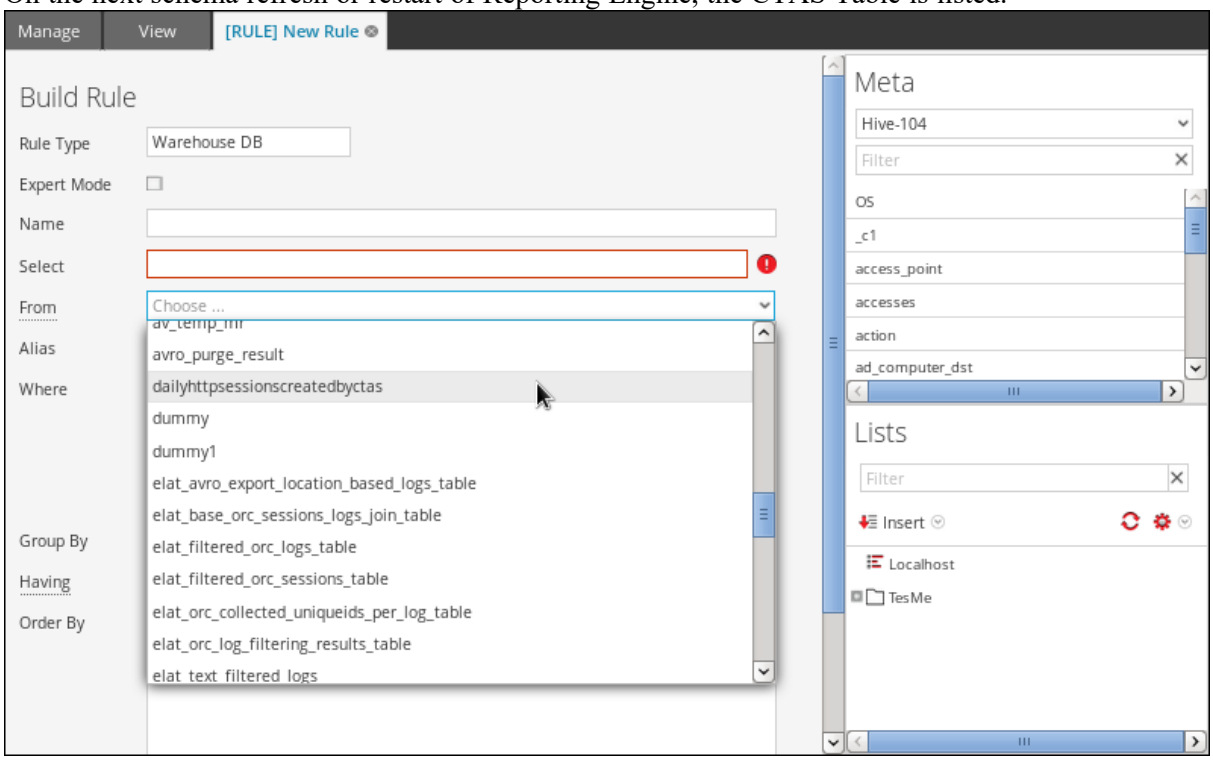
3. Create a Schedule and enter the CTAS Table Name.



4. Run the Report and Reporting Engine will create the Result Summary as below for the Schedule.



5. On the next schema refresh or restart of Reporting Engine, the CTAS Table is listed.



Task Scheduler for Warehouse Reporting

A task scheduler in a Hadoop cluster schedules the jobs consisting of tasks, and allocates specific resources to each job running in a cluster. By default, the task scheduler allocates equal number of resources to all the jobs. For example, if 10 jobs are running they will share resources of the cluster equally. However, you can configure the task scheduler to control the execution of the jobs such that one job runs faster than others by allocating more resources (pools or queues) to the job. This helps you prioritize to run a few reports over others.

Features

NetWitness supports two task schedulers:

- Fair Scheduler (`org.apache.hadoop.mapred.FairScheduler`)
- Capacity Scheduler (`org.apache.hadoop.mapred.CapacityTaskScheduler`)

Fair Scheduler

This scheduler divides the total capacity of the cluster into logical pools. You can submit a job to any one of these pools. All the jobs submitted to a pool share the resources allocated to the pool only. Once a pool has free resources, the freed resources are given to other pools with jobs running. For example, a fair scheduler has 100% resources with two pools namely Pool A and Pool B which share the total resources at 40% and 60% respectively. If Pool A has four jobs running, it allocates 10% resources to each job. When the four jobs are completed, the freed resources are allocated to Pool B.

Note: You can configure a pool to run more than one job in parallel.

Capacity Scheduler

This scheduler divides the total capacity of the cluster into queues. Each queue is allocated a pre-configured share of the total capacity. A job may be submitted to any of these queues. If more than one job is submitted to the same queue, the jobs will be executed sequentially. For example, if a capacity scheduler has 100% resources with three queues namely the Default, Low and High and they share the total resources at 20%, 30% and 50% respectively. If Default has two jobs D1 and D2, Low has three jobs L1, L2 and L3, and High has four jobs H1, H2, H3 and H4, these jobs are executed in their respective queues sequentially. If the jobs in a queue are completed, the freed resources will not be distributed to other queues.

Query Aggregates

This section explains the supported aggregate functions.

Supported Aggregate Functions

The following table lists the supported Aggregate Functions.

Aggregate Function	Description	Input data types	Output data types
count	Returns the count of meta values, which includes duplicate values as well.	Numeric	Numeric
countdistinct	Returns the total number of distinct or unique values.	Numeric	Numeric
distinct	Returns all the unique values.	Any	Any
first	Returns the first occurrence of the meta value.	Any	Same as input
last	Returns the last occurrence of the meta value.	Any	Same as input
sum	Returns a sum of all non-NULL values of metaKey in a group.	Numeric	Numeric
avg (Average)	Returns the average value of all non-NULL values of the metaKey within a group.	Numeric	Numeric
min (Minimum)	Returns the minimum for all values of metaKey in each group. This value is based on order by field.	Any	Any
max (Maximum)	Returns the maximum for all values of metaKey in each group. The maximum value is the value that is returned by order by field.	Any	Any
length	Returns the length of the values of metakey. This is called a "scalar function" in SQL.	Any	Numeric

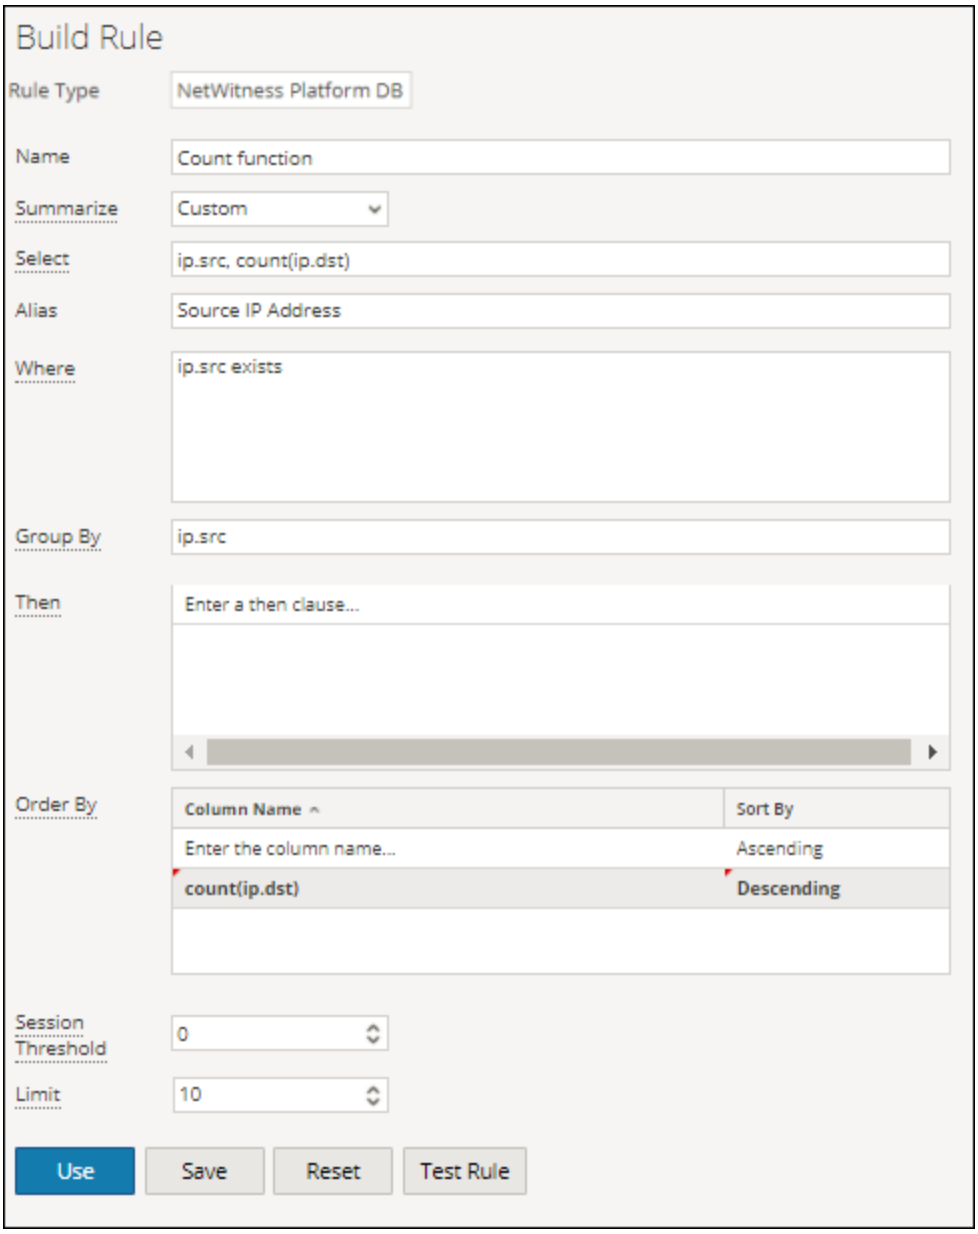
Examples of Queries and Results per Function

Count

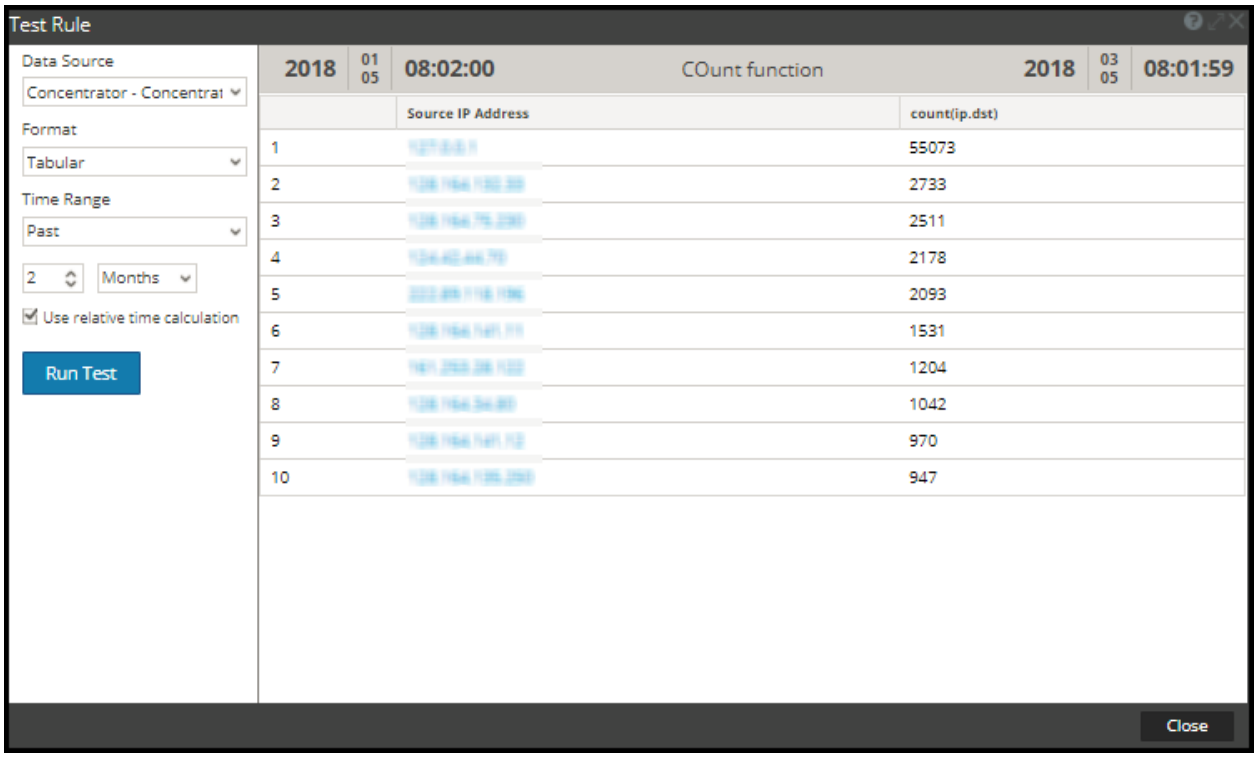
This function returns the number of values for a specified meta key, that exclude null values but include duplicate ones. .

Example

The following figure shows a sample query for count function used for the destination IP and the respective source IP.



The following figure shows the result for the above query.



Here, for each unique ip.src (source IP), the page returns the total number or count of ip.dst (destination IP) values, which include the duplicate values as well.

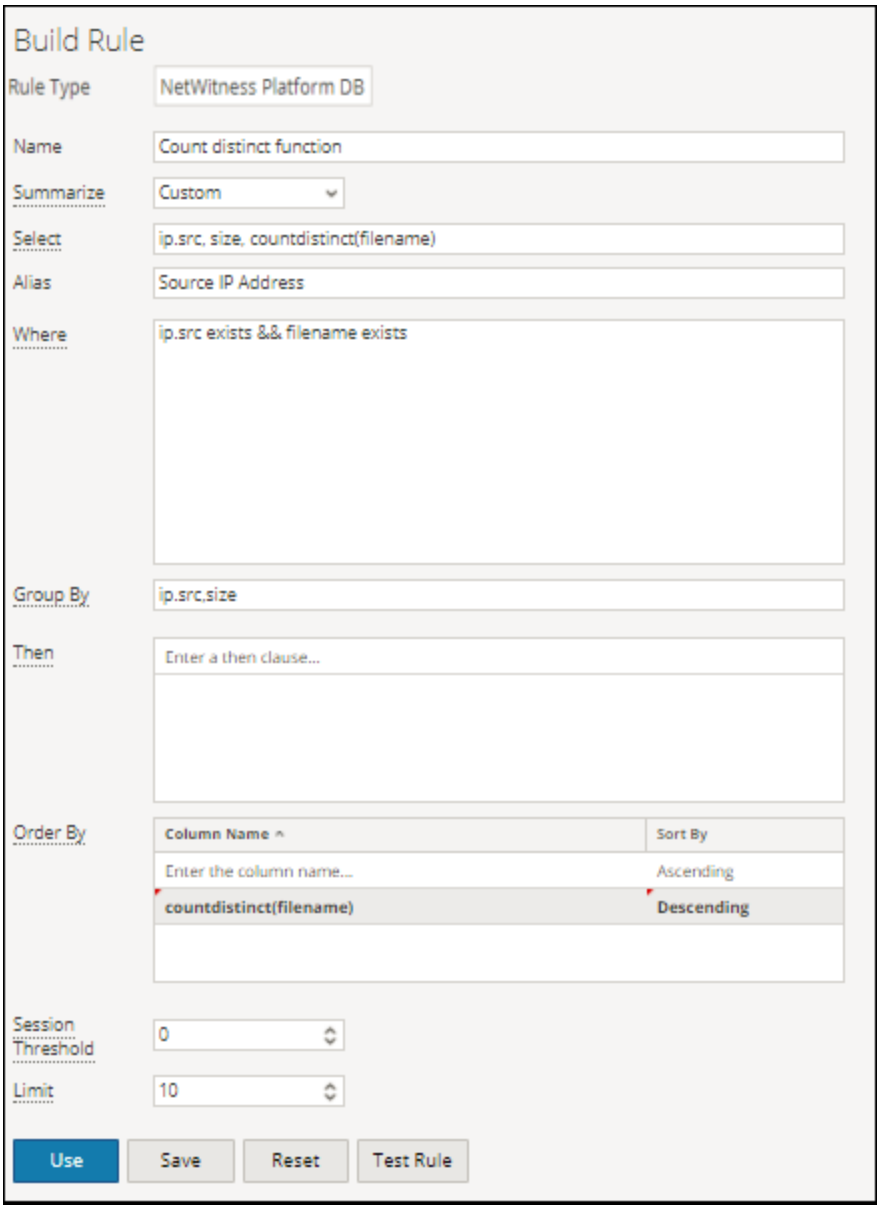
Note: If your NetWitness is currently on 10.5 or newer version and any of the NetWitness Core devices are on 10.3 or 10.4 versions, then some of the aggregate functions may display unexpected errors. However, aggregate functions such as sum() and count() are supported in 10.4 version.

Countdistinct

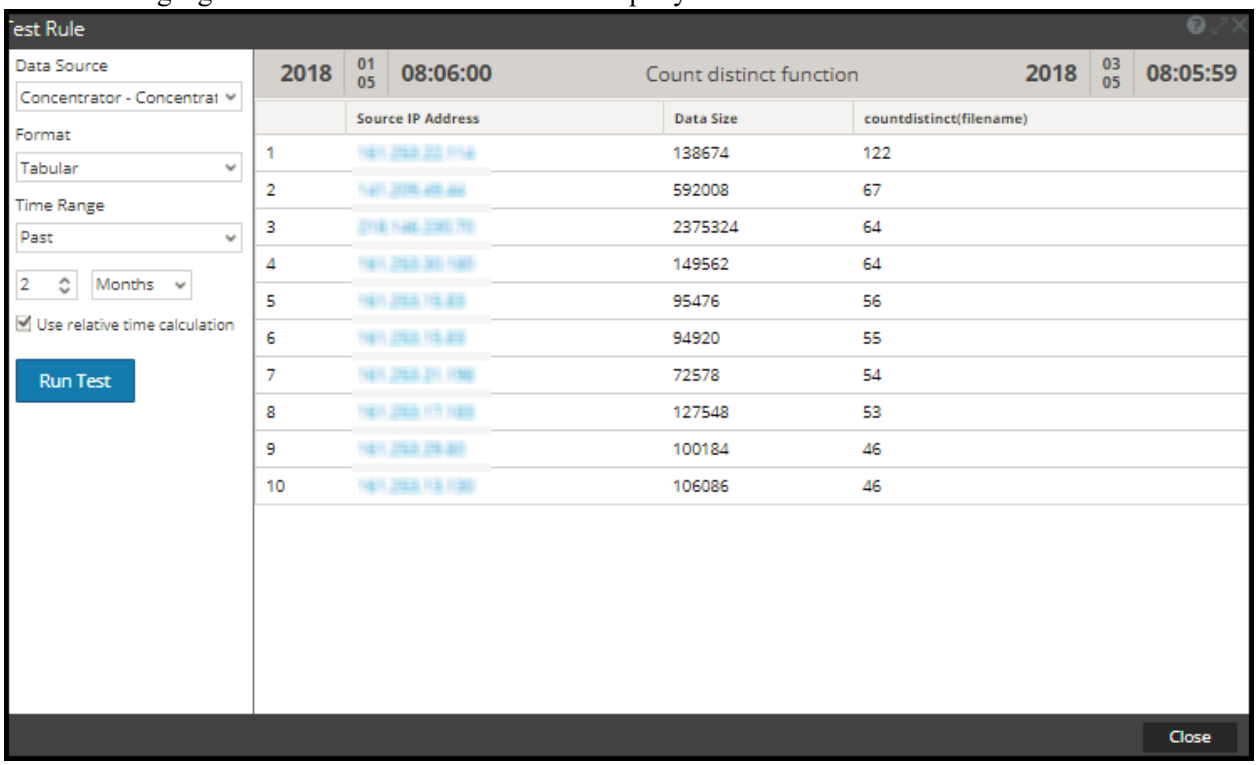
The countdistinct function returns the count of unique or distinct values for the metakey. In other words, countdistinct function can be used to retrieve a number of distinct values for the specified metakey.

The following figure shows a sample query where the countdistinct function is used along with IP source (ip.src) and data size(size).

Example


Build Rule
Rule Type
NetWitness Platform DB
Name
Count distinct function
Summarize
Custom
Select
ip.src, size, countdistinct(filename)
Alias
Source IP Address
Where
ip.src exists && filename exists
Group By
ip.src,size
Then
Enter a then clause...
Order By
Column Name
Sort By
Enter the column name...
Ascending
countdistinct(filename)
Descending
Session Threshold
0
Limit
10
Use
Save
Reset
Test Rule

The following figure shows the result for the above query.



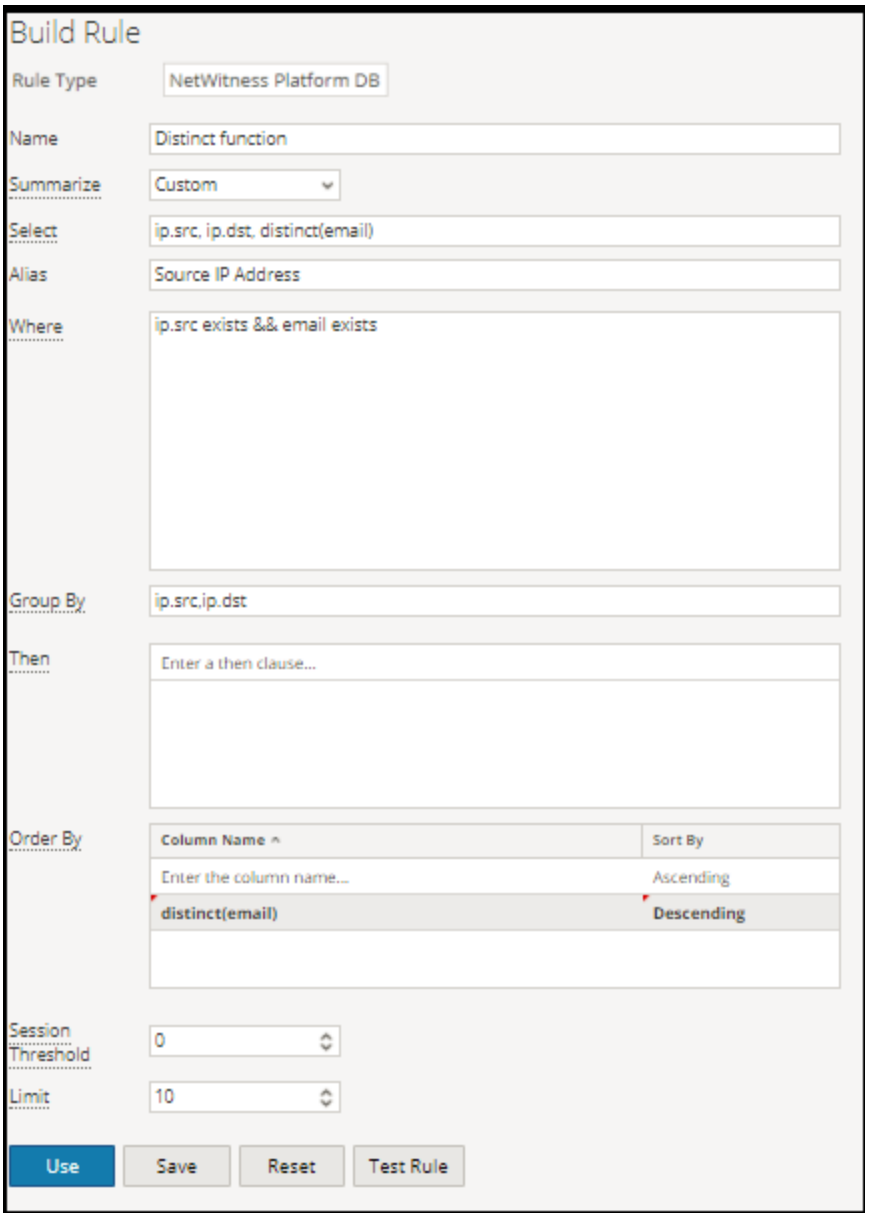
Here, the page displays the data size along with the total number or count of distinct filenames from the respective IP source. Unlike the count function, the countdistinct excludes the duplicate values from the result.

Distinct

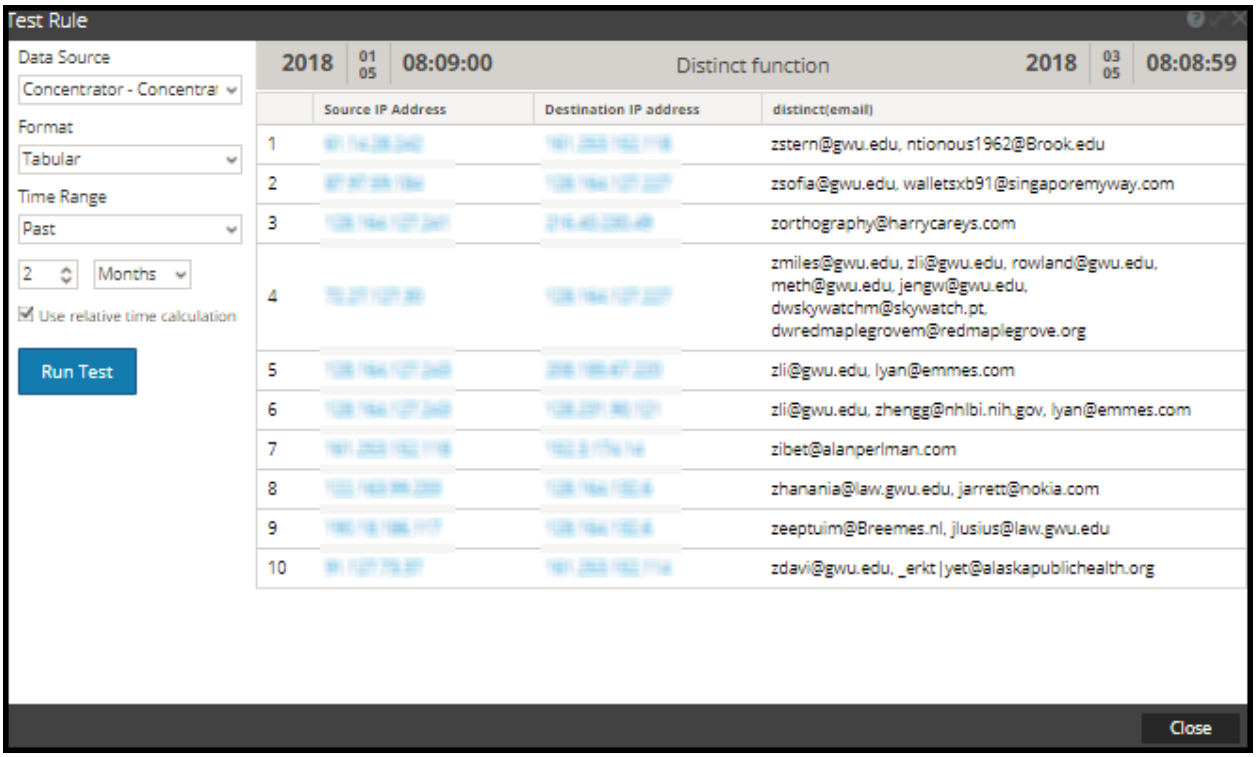
This function returns all the unique or distinct values of the metakey.

Example

The following figure shows a sample query for distinct function used to retrieve e-mails, between various source and destination IP (ip.dst).



The following figure shows the result for the above query.



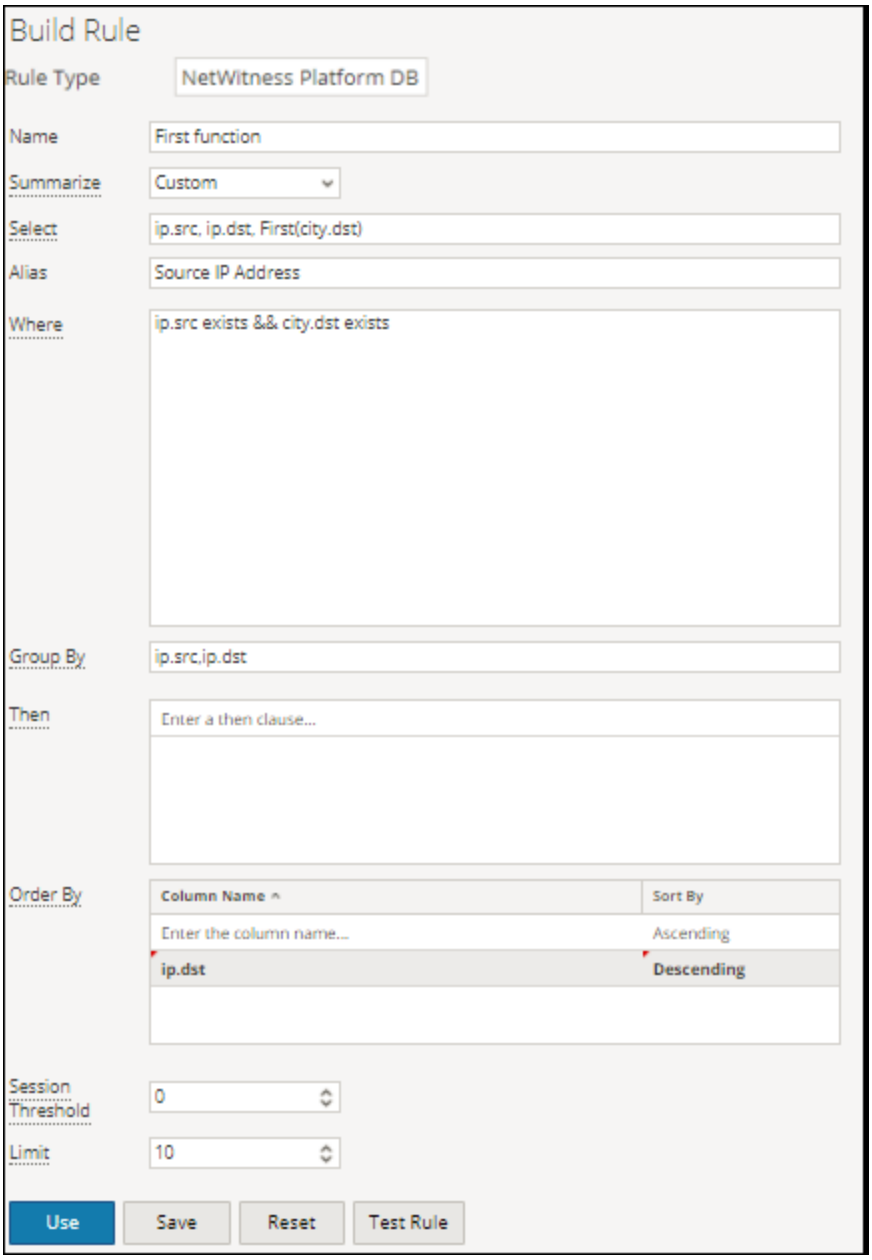
Here, the page displays the list of unique e-mails that were exchanged between the respective IP source and destination.

First

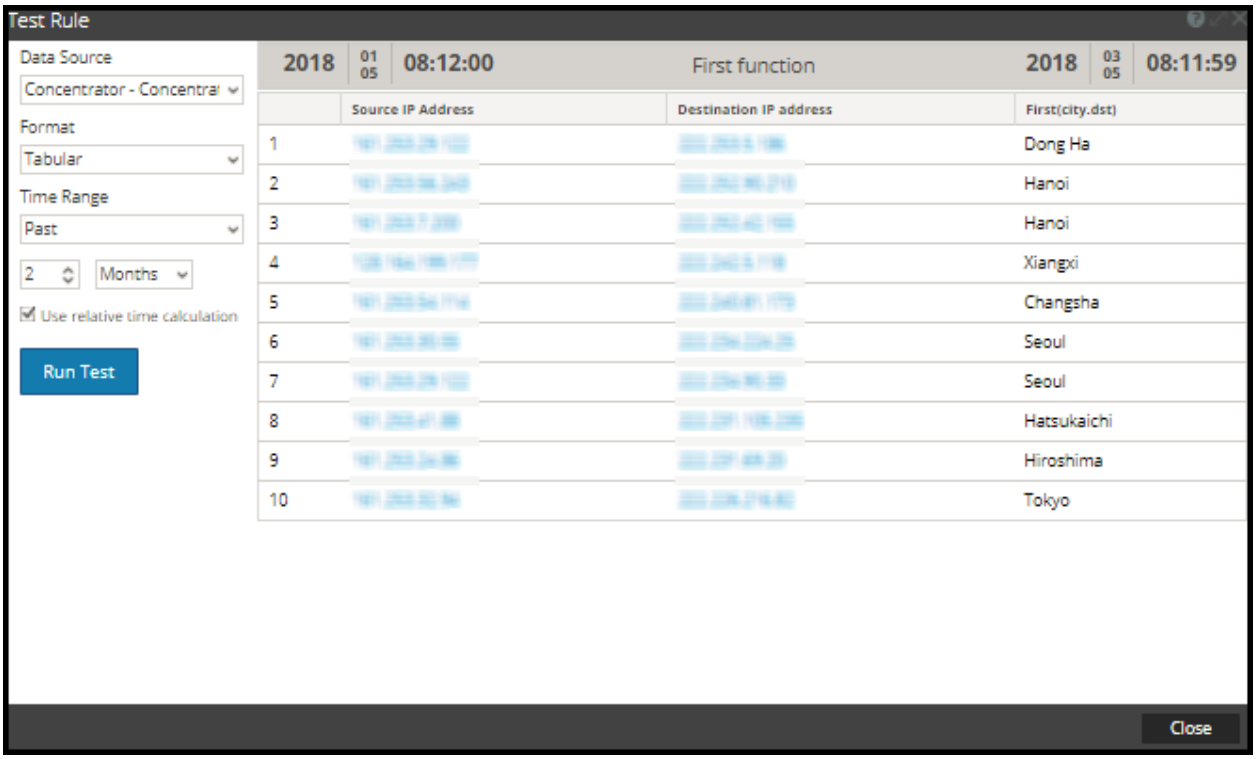
This function is used to retrieve the first value from an ordered sequence of values for a specified metakey.

Example

The following figure shows a sample query for first function used to retrieve the first destination city name.



The following figure shows the result for the above query.



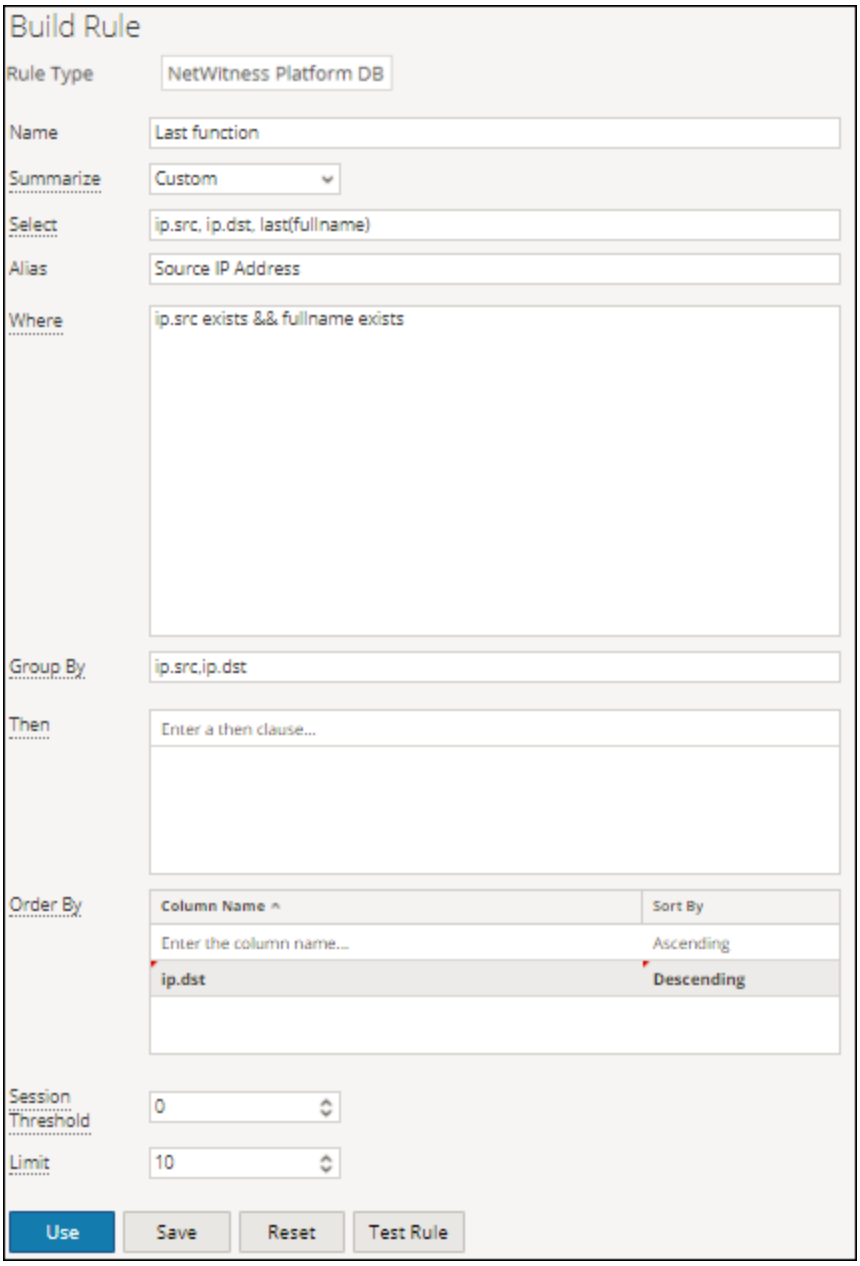
Here, the page displays the the first destination city for the corresponding source and destination IP. You can use the first function to isolate a particular value from a search result.

Last

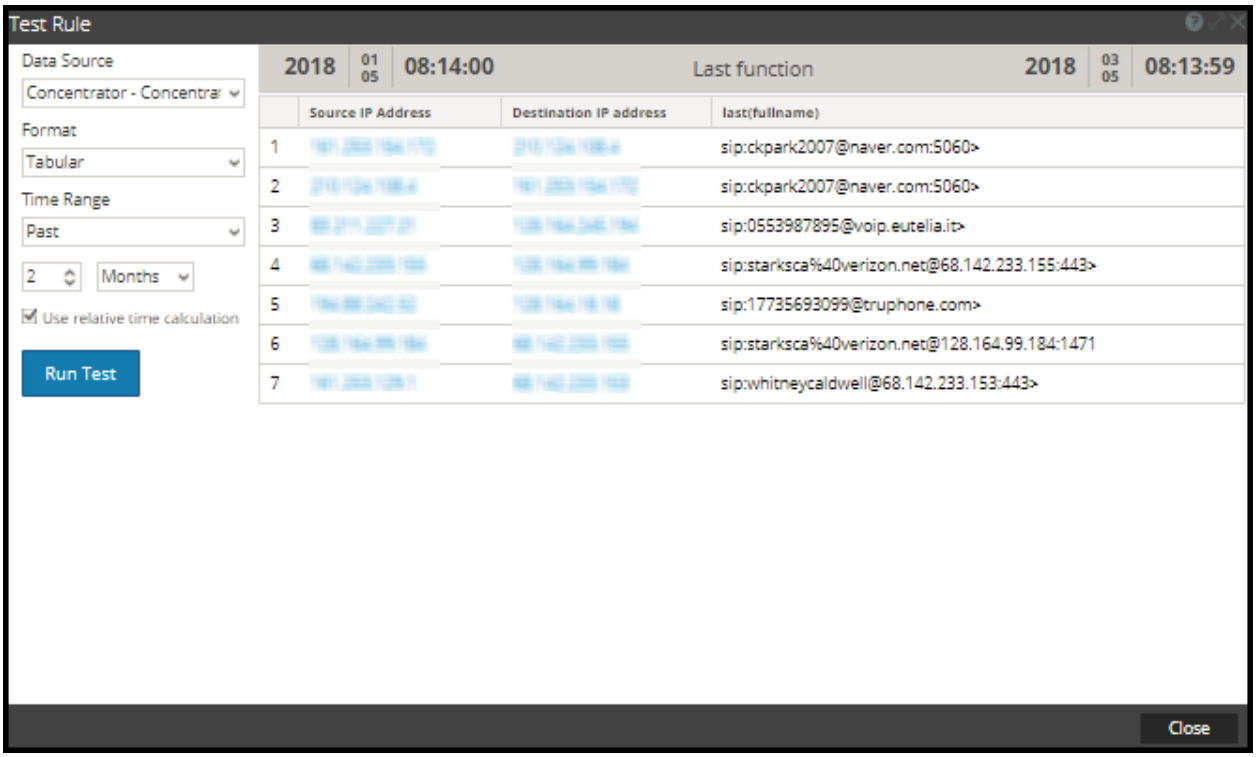
This function is used to retrieve the last value from an ordered sequence of values for a specified metakey.

Example

The following figure shows a sample query for last function used to retrieve the most recent user name.



The following figure shows the result for the above query.



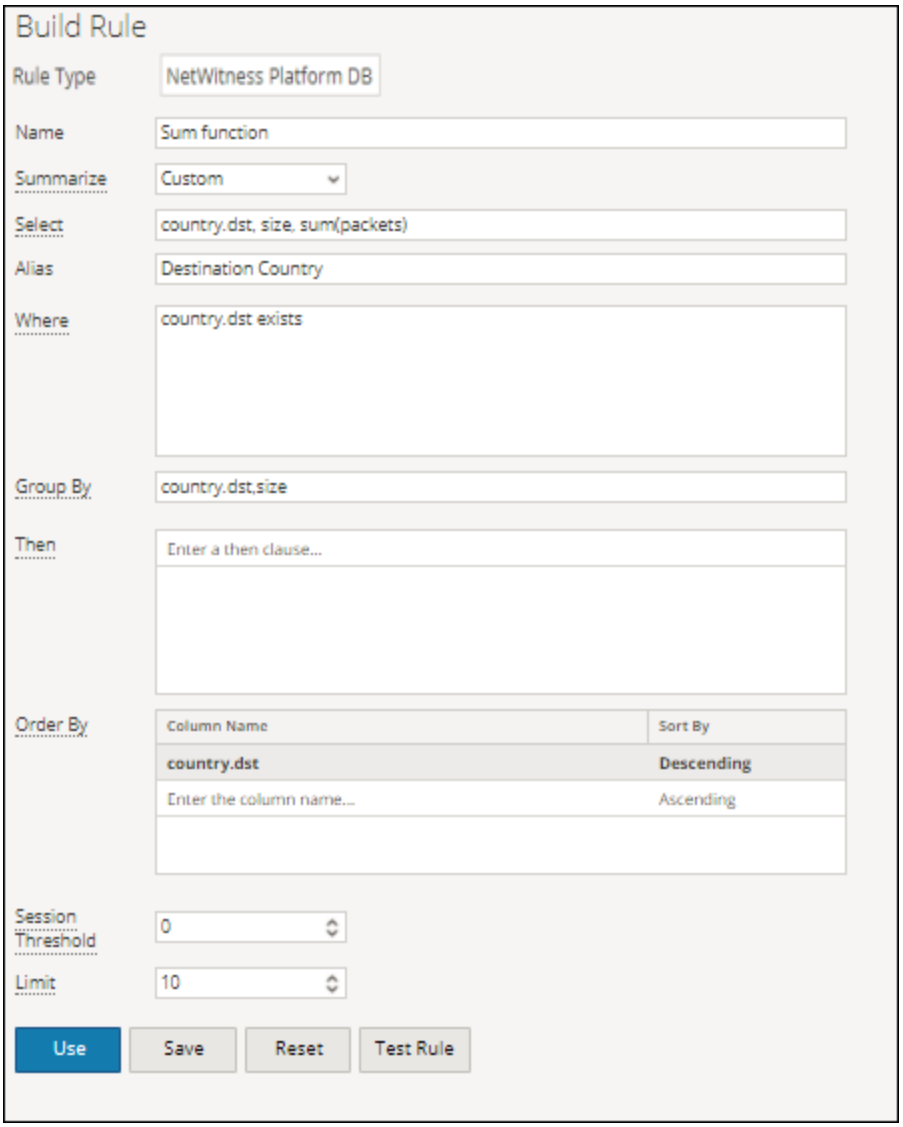
Here, the page displays the list of most recent or last usernames in full, that were exchanged between the source and destination IP.

Sum

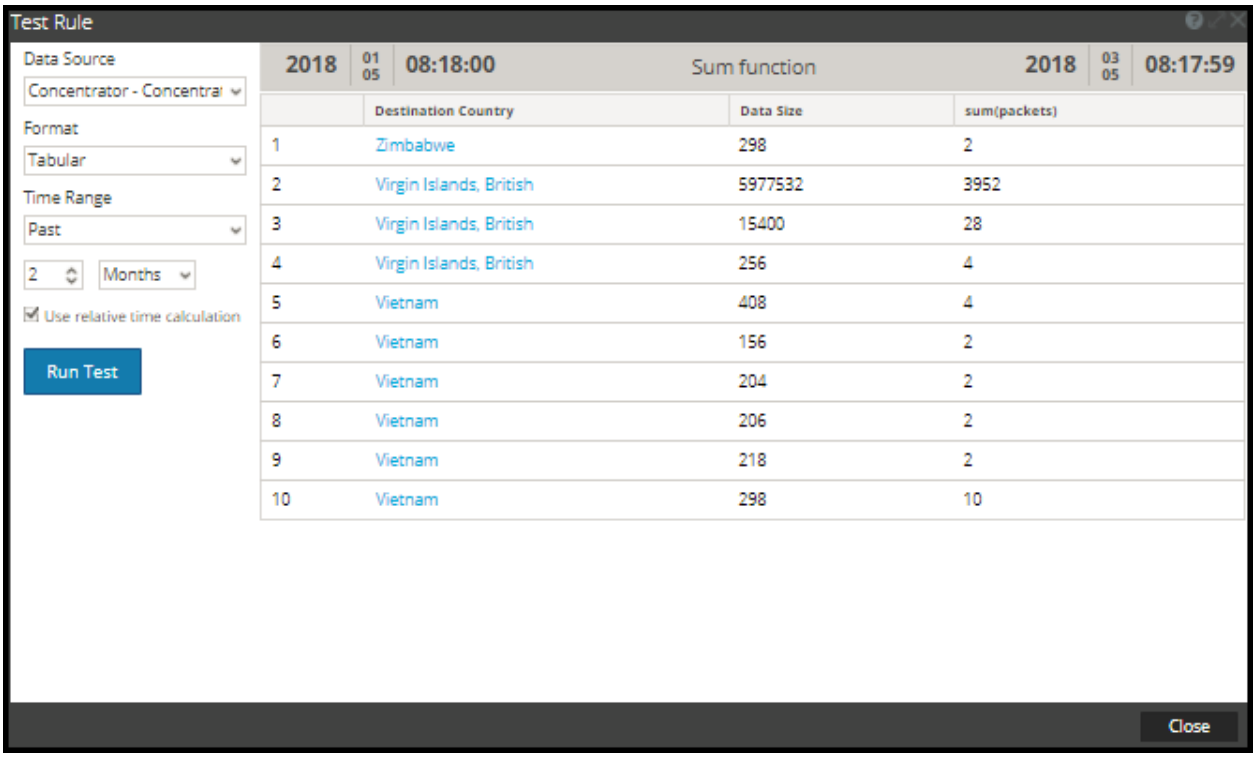
This function returns the total of the non-NULL values of the metaKey within a group.

Example

The following figure shows the query for the Sum function used for packets.



The following figure shows the result of the above query.



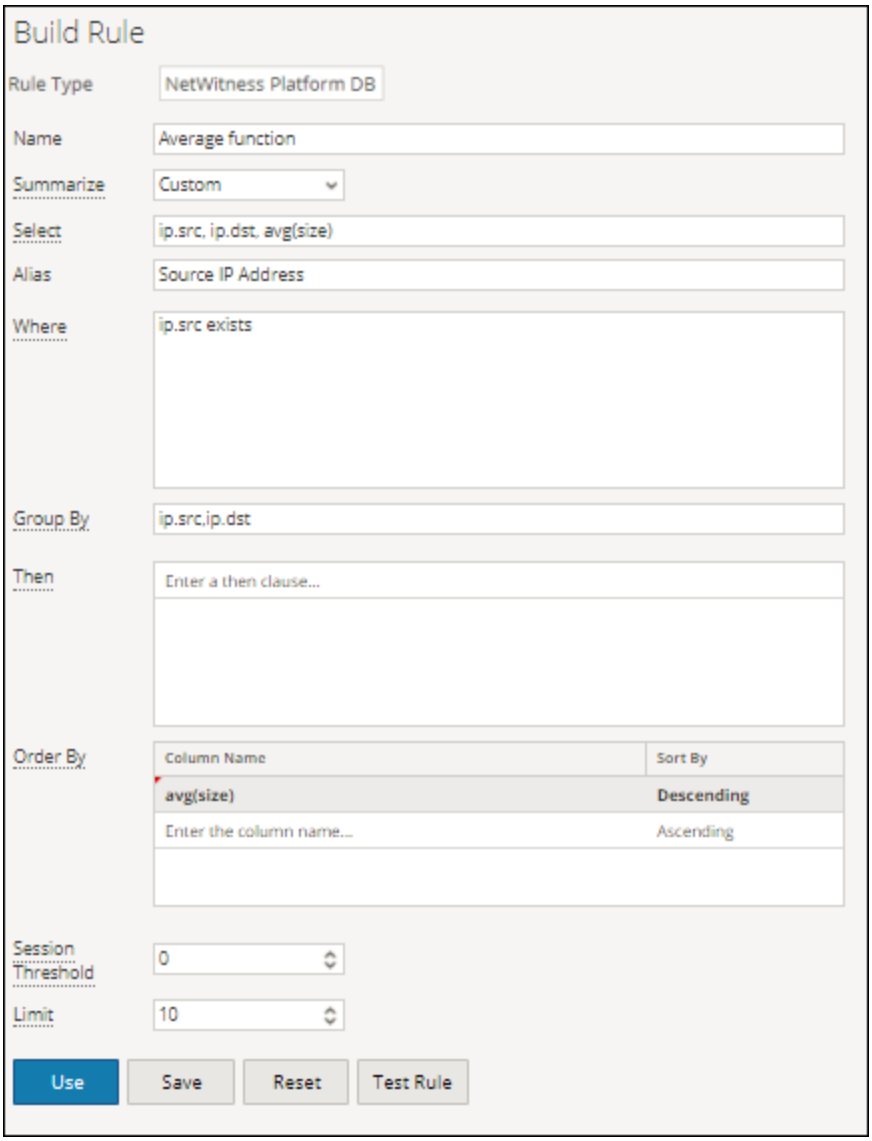
Here the page displays the total or sum of the packets along with the size of the data for the respective destination country.

Avg

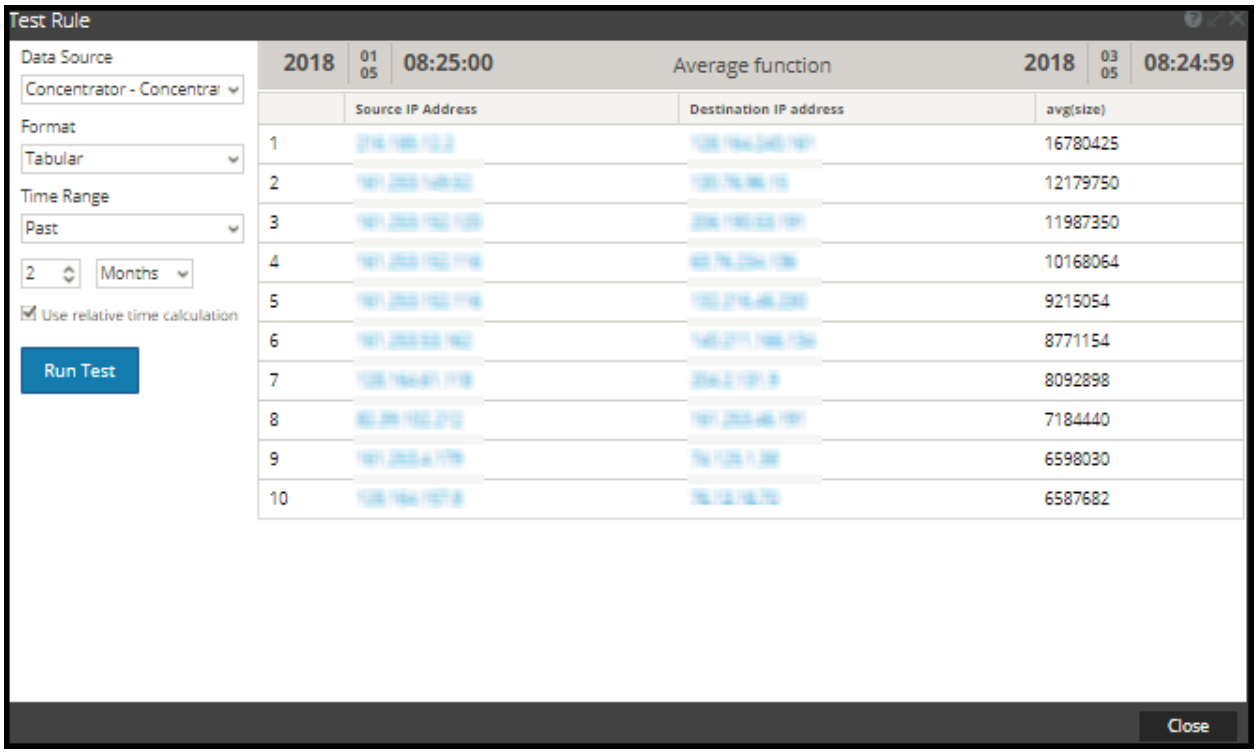
The average function returns the average of non-NULL values of the meta within a group.

Example

The following figure shows a sample query for average data size transmitted between a source and destination IP.



The following figure shows the result for the above query.



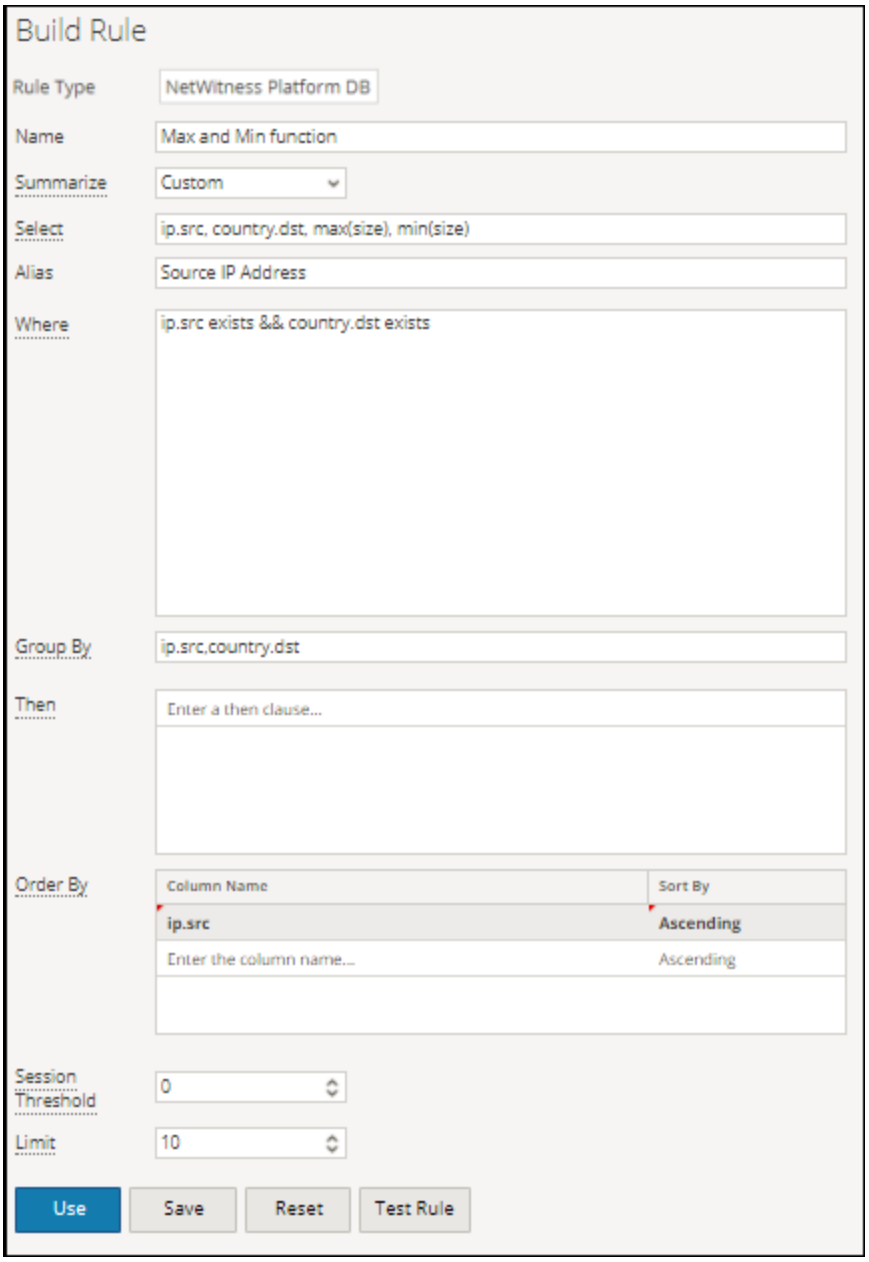
Here, the page displays the average size of data exchanged between source and destination IP:

Max and Min

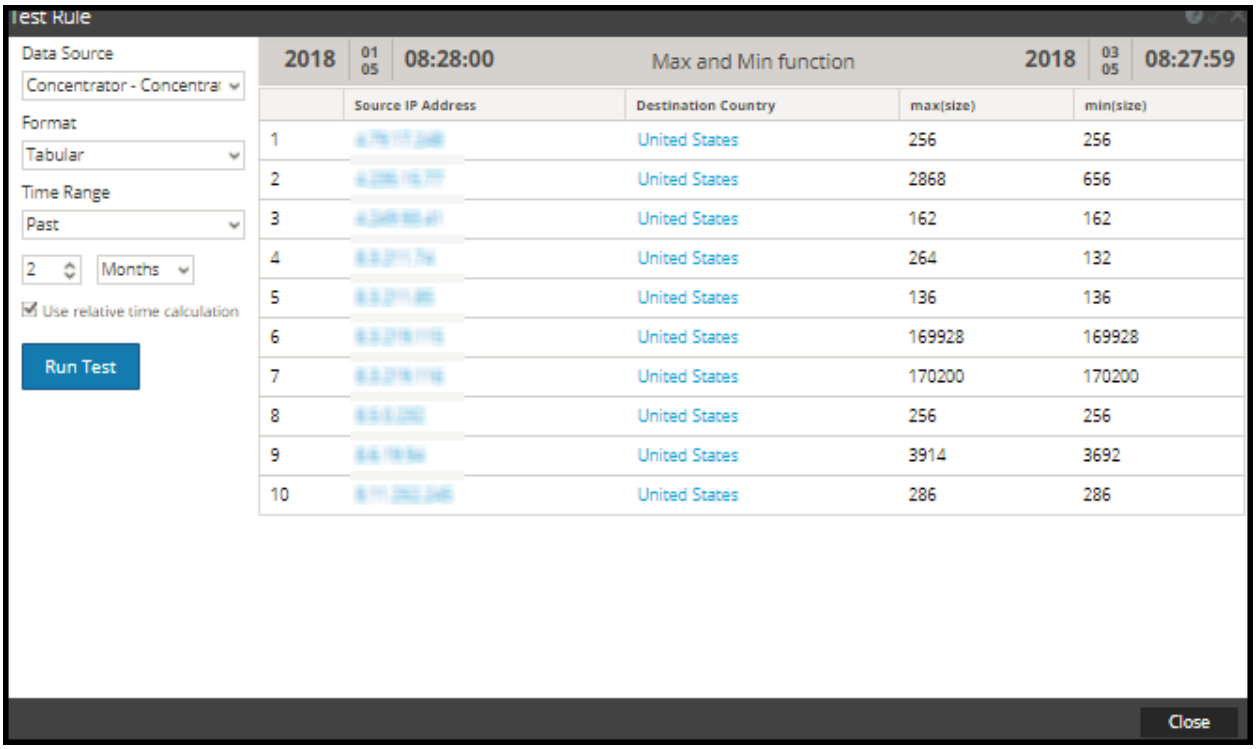
Max and Min functions provide the maximum and minimum for given values of a meta respectively.

The following figure shows a sample query for max and min functions for various data sizes, for source IP and destination country.

Example



The following figure shows the result for the above query.



Here, the page displays the max(size) and min(size) columns, along with the list of source IP and destination country. The max(size) column lists the maximum data sizes exchanged while the min(size) column lists the minimum data sizes that were exchanged.

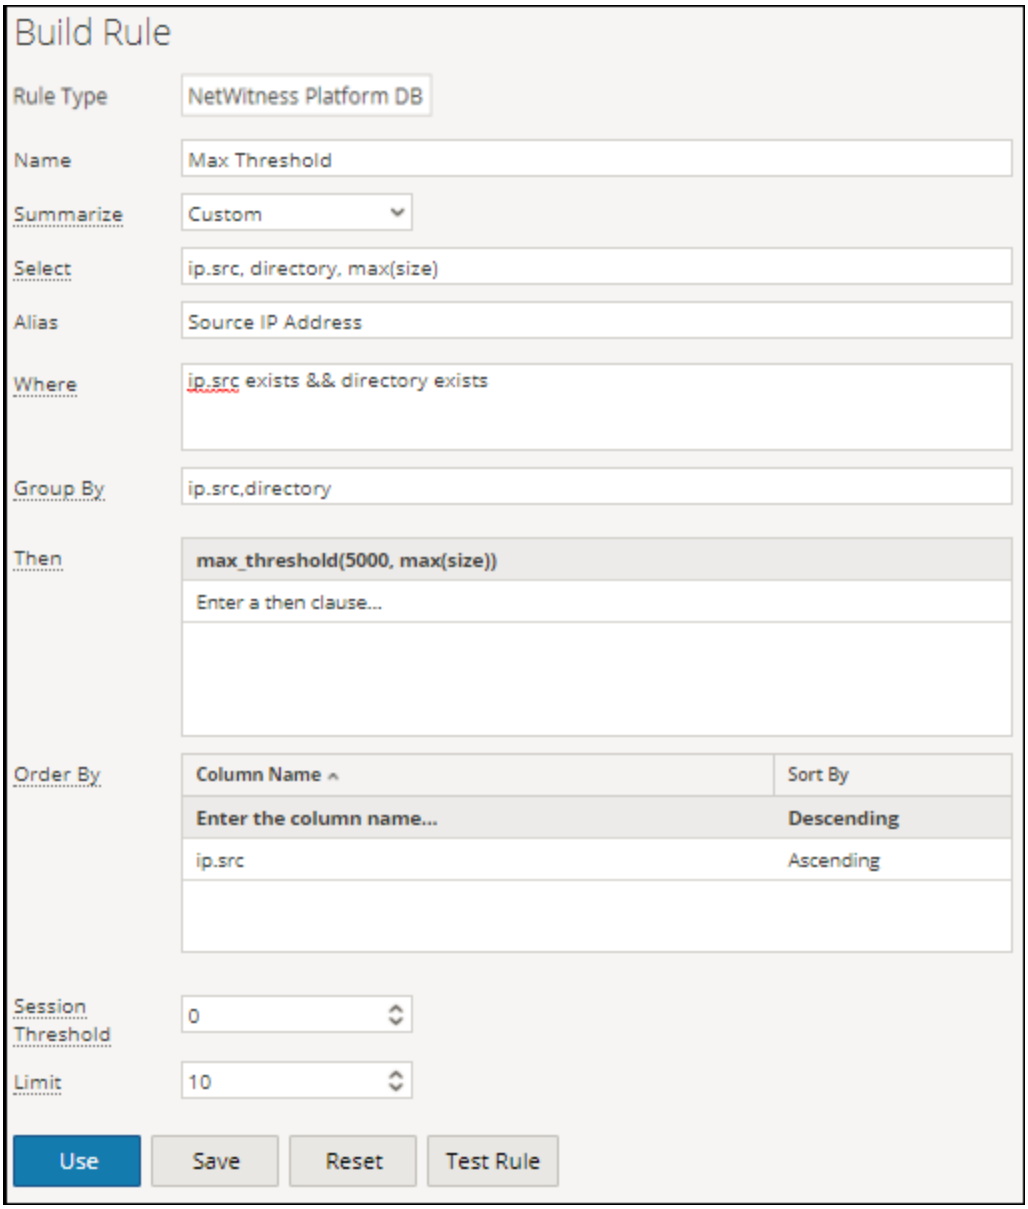
Filter aggregate meta results with Max_threshold

You can further filter the results of any function by using the threshold rule action.

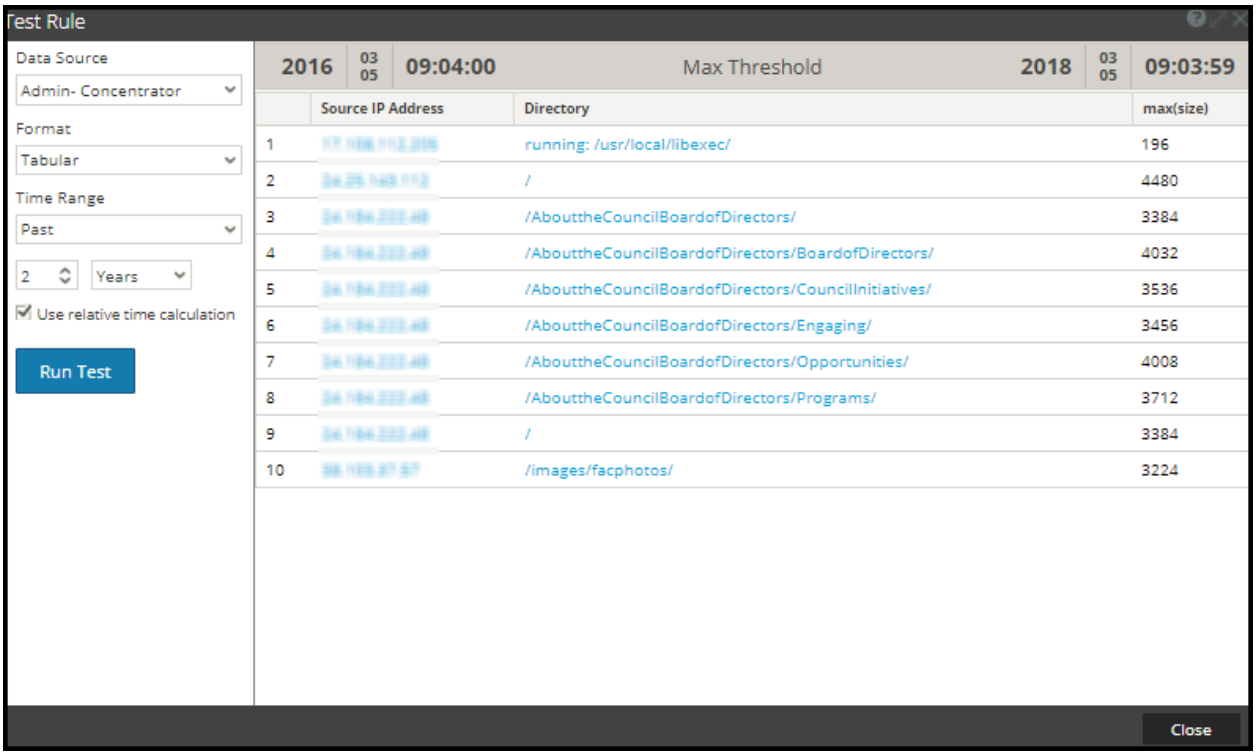
Example

Following is a sample query for max_threshold used along with the Max function in the **Then** field is:
max_threshold(5000,max(size))

The following figure shows the Build Rule screen for the above query.



Here the max_threshold is applied for data size with an upper limit of 5000. The following figure shows the result.



Here, the result page displays the max(size) column, that lists the data sizes lesser than 5000 as this is the maximum threshold in the query, along with the corresponding IP source and the respective directory.

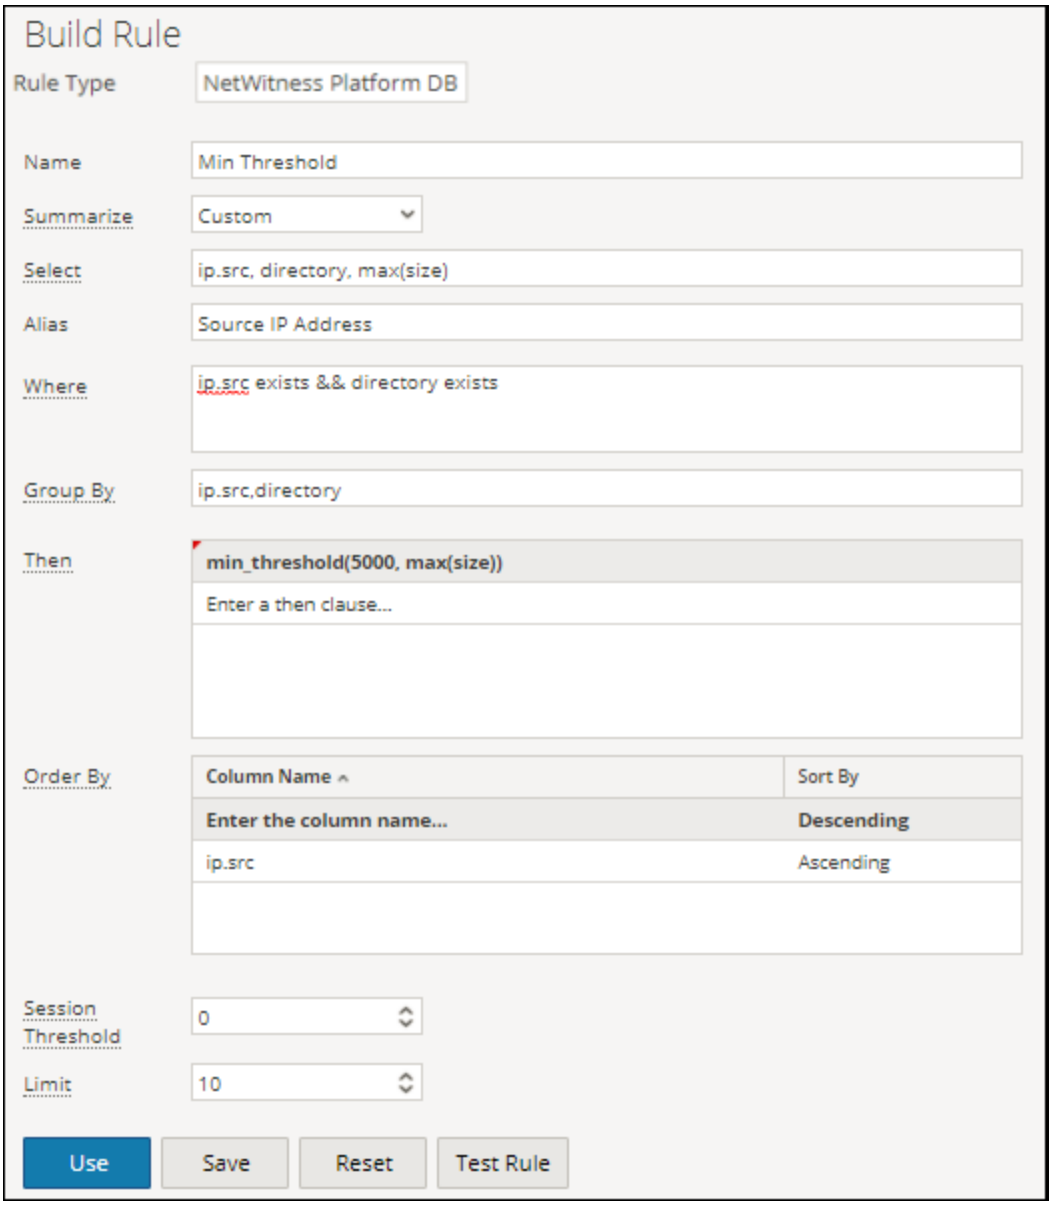
Filter aggregate meta results withMin_threshold

Similarly, min_threshold is used to filter the results for any function. A similar scenario as max_ threshold is considered to explain this.

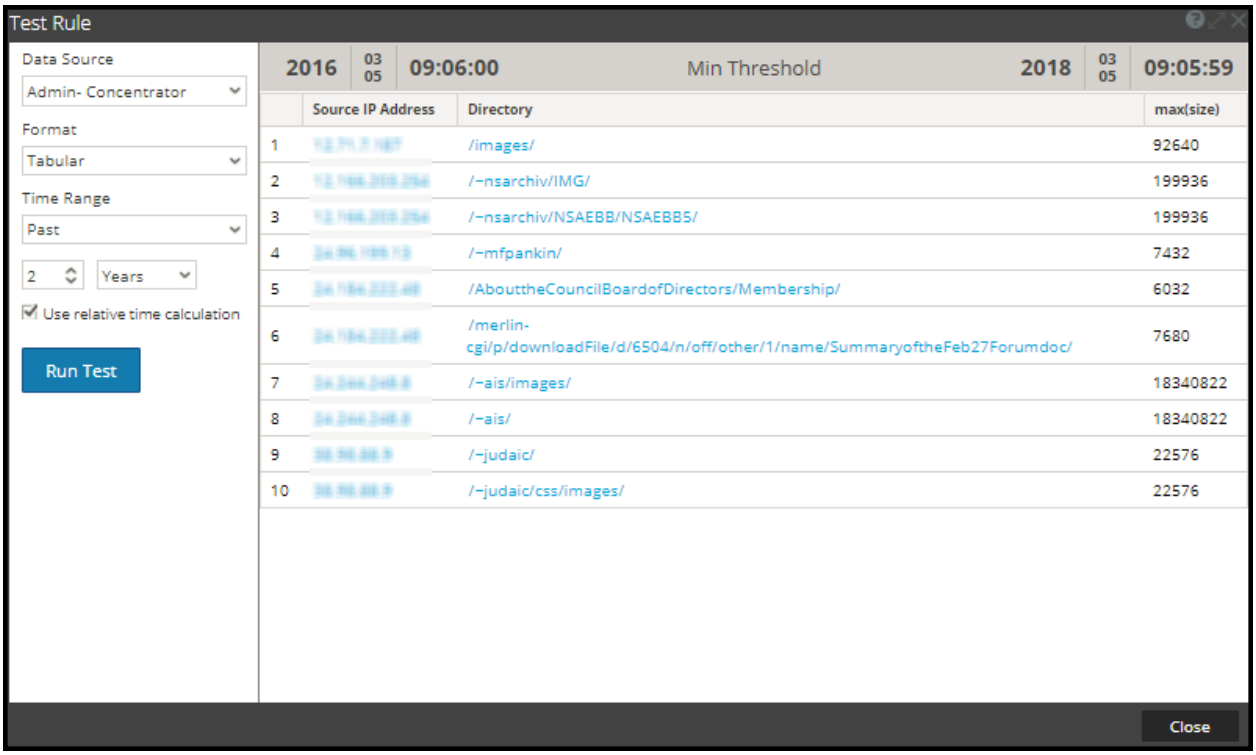
Example

Query for min_threshold used along with the Max function in the **Then** field is:
min_threshold(5000,max(size))

The following figure shows the Build Rule screen for the above query.



Here the min_threshold is applied for data size with a lower limit of 5000. The following figure shows the result.



Here, the result page displays the max(size) column, that lists the data sizes greater than 5000 as this is the minimum threshold in the query, along with the corresponding IP source and the respective directory.

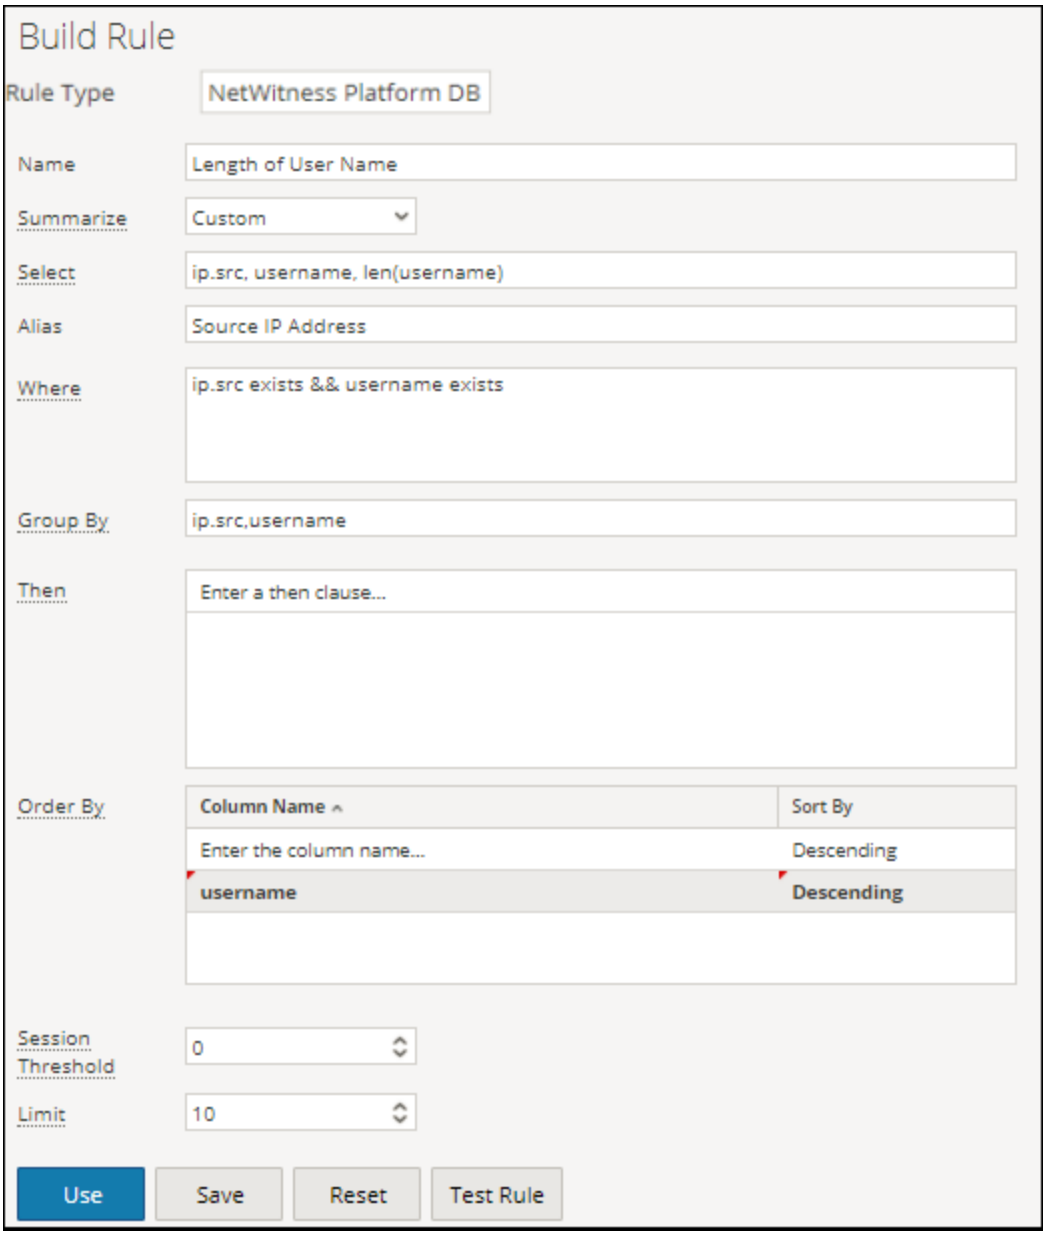
Note: Max_threshold and Min_threshold rule actions are common across all the functions, and can be used along with the other queries in the **Then** field to retrieve the respective output.

Length

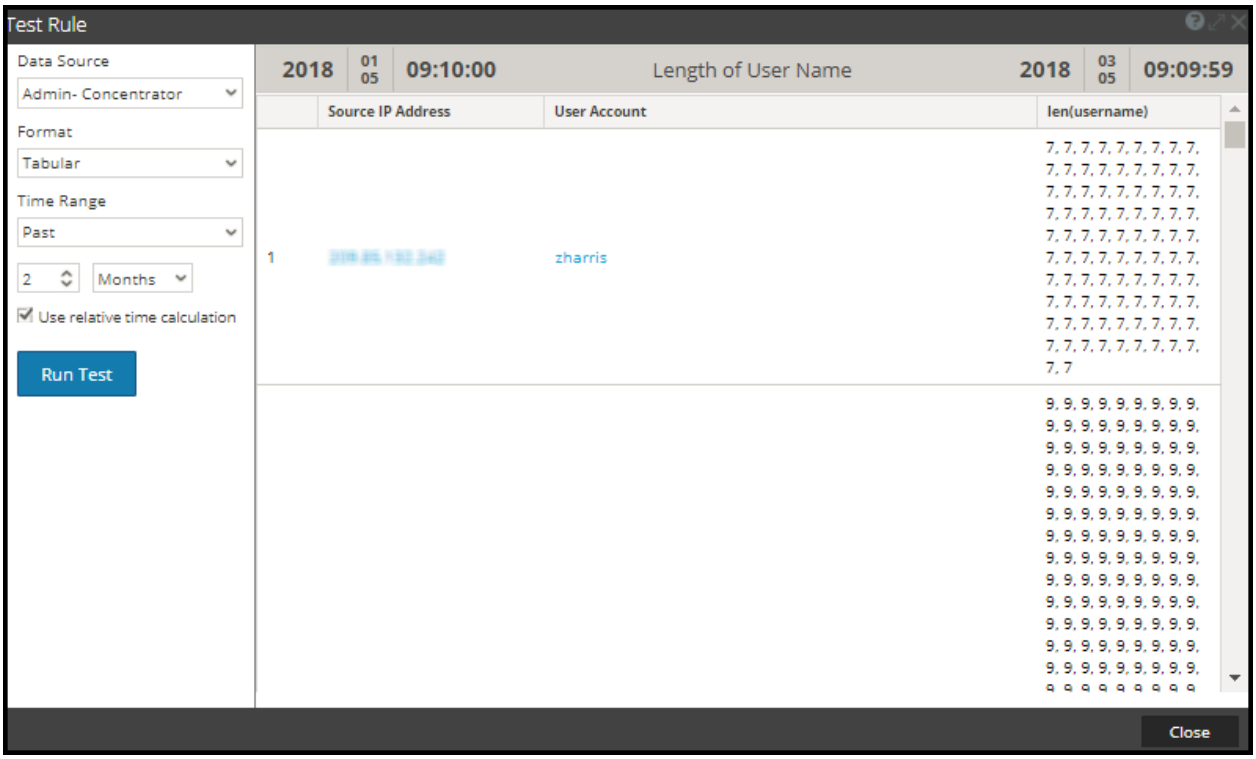
This function returns the length of a meta value. In other words, Length function returns the number of bytes used to store the actual value.
For instance, for the value "Analytics" it returns the length as 9. Similarly, for an IPv4 ip.src, it returns 4 (representing 4 bytes).

Example

The following figure shows a sample query for the length function used for usernames.



The following figure shows the result for the above query.



Here, the page displays the length of the usernames associated with the user account and their respective source IP.

Additional Information

When you query for aggregates (E.g. sum(size)) with **Group By** on a meta which has multiple values in a session, then the session with multiple values is accounted for aggregate calculation for each value of that meta.

Example

When you query for the Count aggregate function with Group By on Alias.host and if the column has multiple values in a session, then the session is counted for each occurrence, including the duplicate values.

Consider the following table.

SessionID	Alias.host	Ip.src	Size
1	host-a, host-b, host-a	a	10
2	host-b, host-c, host-a, host-c	c	20
3	host-b, host-c, host-d	b	30
4	host-c, host-a	a	40

In the above table, alias.host for **host-a** and **host-c** has duplicate values listed for a single session. Let us consider the following query:

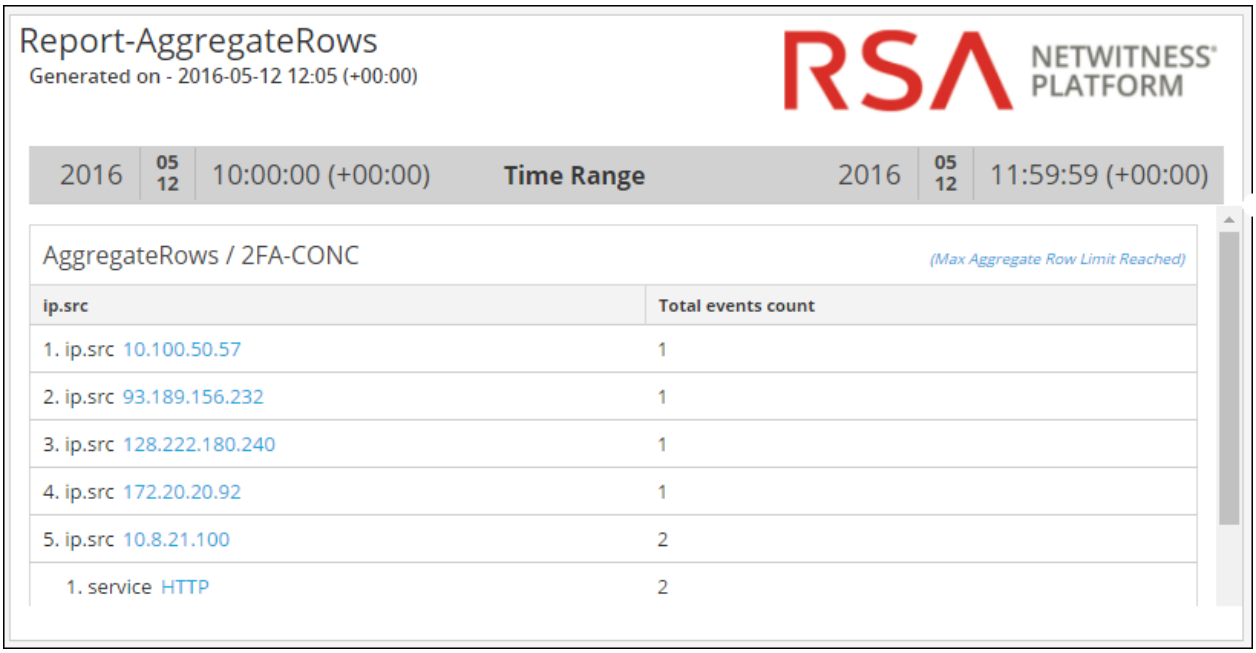
Select : alias.host, count(ip.src), sum(size)
Group By : alias.host

Here, **host-a** and **host-c** are present in 3 sessions and they are duplicated for two different sessions. However, the output is as shown below.

Alias.host	count(lp.src)	Sum (size)
host-a	4	80
host-b	3	60
host-c	4	110
host-d	1	30

Output table shows that the count for **host-a** and **host-c** is 4. This is because for each alias.host value, the entire session is considered. Similarly to calculate sum (size), the same sessions are considered for each alias.host value.

In the report output if the number of rows has reached **NWDB maximum aggregate rows** defined in RE configuration, then a message **Max Aggregate Row Limit Reached** is displayed to indicate that there is more information to be displayed. The default limit is 1000, and you can change this value as per your requirement, in the Reporting Engine Configuration page .



Troubleshoot Reporting

This section provides troubleshooting instructions for issues faced when using the Reporting module in NetWitness.

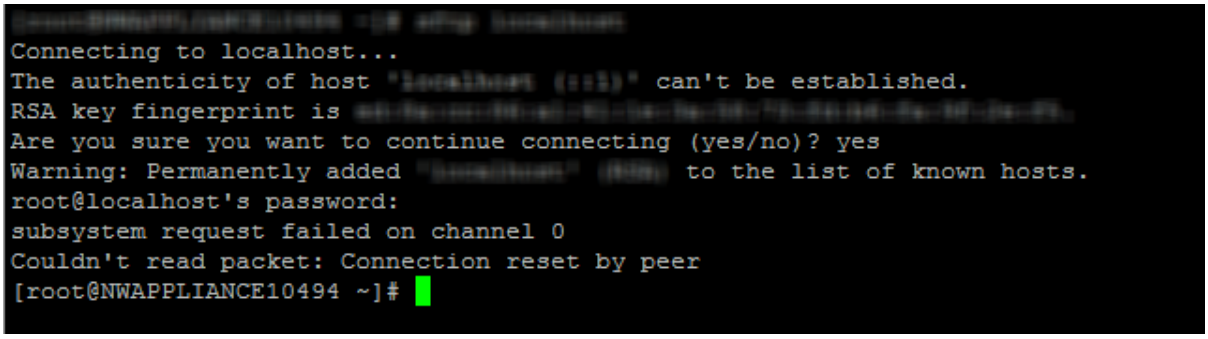
Configuring SFTP Server Issue

Procedure

Try the following steps if you face any issues while configuring the Linux SFTP server:

1. If the Report Output Action for the configured SFTP fails, you must SSH to the SFTP server and try to connect locally to check if SFTP is working fine.

   Connect to SFTP server:

2. If the Local connection fails, open the file **sshd_config> vi /etc/ssh/sshd_config.**
3. Check for the entry in the file:

   ```
   # override default of no subsystems
   Subsystem sftp /usr/libexec/openssh/sftp-server
   ```

4. If this entry does not exist, add the two lines mentioned in Step 3 at the bottom of the file and **Save** it.
5. Restart service from **SSH > service sshd restart**.
6. Retry the SFTP connection now.
7. Make sure SFTP port is not blocked by SA server appliance firewall. Update iptables rules to allow sftp port.

Meta Values in Investigation Link Issue

Issue	When the device information on the datasource is changed, the Investigation link for the meta values of the executed reports is not displayed on the NWDB results page.
Resolution	Remove and re-add the datasource to Reporting Engine. *Note: This workaround is not applicable for reports that are already generated.*

Internet Explorer 10 Browser Issue

Issue	When you click the Test Rule multiple times in quick succession, results with large input data may not displayed in Internet Explorer 10.
Resolution	If this issue occurs, try one of the following steps: • Close the Test Rule window on Internet Explorer 10 and run the test again. • Use other browsers like Chrome or Mozilla Firefox to test the rule execution.

Dynamic List Editing Issue

Issue	A dynamic list cannot be added from the Edit option on the 'View All Schedules' page to an existing schedule.
Resolution	1. Reports > Select the report > 2. Click the #Schedules for the specific report 3. Select the schedule to be modified from the Report Schedule page 4. Edit the schedule

Deployment Failure Issue

Issue	Deployment of reports fail, if the dependencies of certain compliance reports in Live are not deployed prior to the reports.
Resolution	Retry the deployment. If the problem persists, try to deploy the rule or list dependencies first and then deploy the reports.

Respond Server Issue

Issue	When the Forward Alerts to Respond option is enabled and RabbitMQ connections to the Respond Server are blocked, some of the Reporting Engine threads may be blocked.
Resolution	Disable the Forward Alerts to Respond option until the RabbitMQ broker in the NetWitness server at the Respond has begun and accepts the connections.

Post-Upgrade Issue

Issue	Post-upgrade from 10.6.x to 11.2, Categories meta for incident collection is not supported.
Resolution	When using the Categories meta for incident collection, the results rendered are in an incorrect format. Hence this meta is not supported and you cannot use the categories meta in either select clause or where clause. Also, it is not available in the list of metas for selection in the Rule Builder page.

Report Query Timeout Issue

Issue	When scheduling the report, if you choose a high time range, the query may timeout with the error message **`Query on channel 12345 was canceled by the system for exceeding time usage limits. Check timeout values.`**
Resolution	Set the **Summarize** option to **None** in the **Build Rule** view to exclude the **Group-By** clause (that is performance intensive) and then re-schedule the report. If the query still times out after 60 minutes, increase the timeout value from 60 minutes to 90 minutes on all the core services. To increase the timeout value: 1. Go to **Admin** > **Services** > **Security** > **Users** and set the **Core Query Timeout** to 90 minutes. 2. Enter the Service Admin Password and click **Save**. 3. Re-schedule the report.

Reporting References

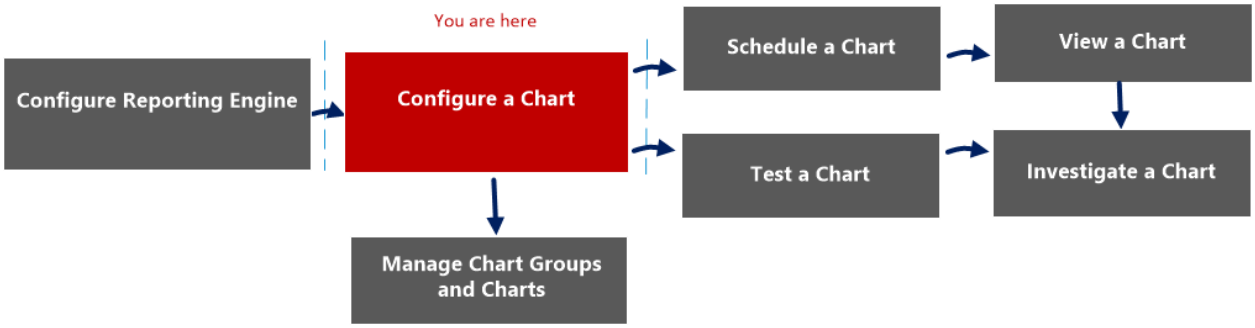
This section provides information about the Reporting user interface. You can look at your place in the workflow for creating and generating a report with theNetWitness, get a quick look at the important features, and follow links to the detailed concepts and procedures.

Build Chart View

In the Build Chart view, you can define and test a chart. You build a chart by assigning a name and then selecting a rule to include.

Note: Only the NetWitness DB rules can be used in charts.

Workflow



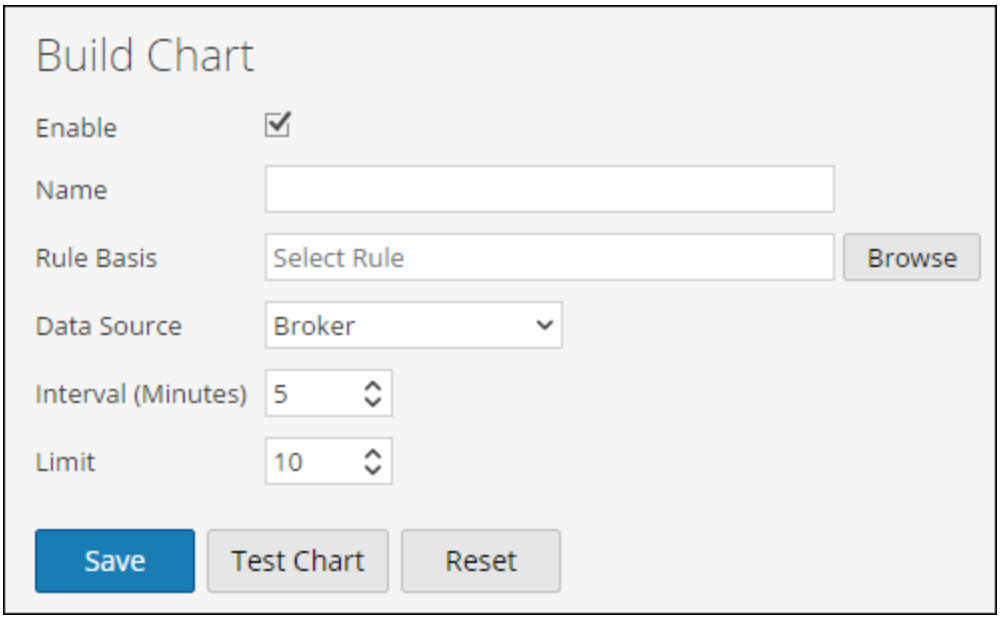
What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart*	Configure a Chart
Administrator/ Analyst	Schedule a chart	Schedule a Chart
Administrator/ Analyst	View a chart	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart	Manage a Chart Group and Chart

*You can complete these tasks here.

Quick View

The following figure is an example of the Build Chart view.



The following table describes the features in the Build Chart view.

Field	Description
Enable	Specifies if the Reporting Engine must collect the data and generate chart results. If the **Enable** checkbox is not selected, the results are not rendered.
Chart Name	Identifies the name of the chart.
Rule Basis	Displays the Add Rules dialog box from which you select a rule that is the basis of a chart. The rule that you select must be a rule which is not sorted by none.
Data Source	If the default data source is configured in the Reporting Engine, the data source is displayed on the Build Chart page. If a chart is configured to run on any other data source, that data source is displayed on the Build Chart page instead of the default data source. The Reporting module works with the following data sources: • Broker • Concentrator • Decoder • Log Decoder • Log Collector
Interval (Minutes)	The chart data refresh interval in minutes.
Limit	The number of records for which a chart is generated.
Save	Saves a chart to the database.
Test Chart	Plots a test chart based on the chart definition.

Field	Description
Reset	Resets the chart details.

Build List View

In the Build List view, you can enter or import values to create a list and save or reset the values. You can use lists when you are writing reporting rules to simplify the process of specifying values in the rule.

Workflow

This workflow shows the procedure to define lists or list groups. You can set access control at the list or list group level so that only users with specific roles can access the lists.

You must ensure that Reporting Engine is configured on the NetWitness.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule*	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

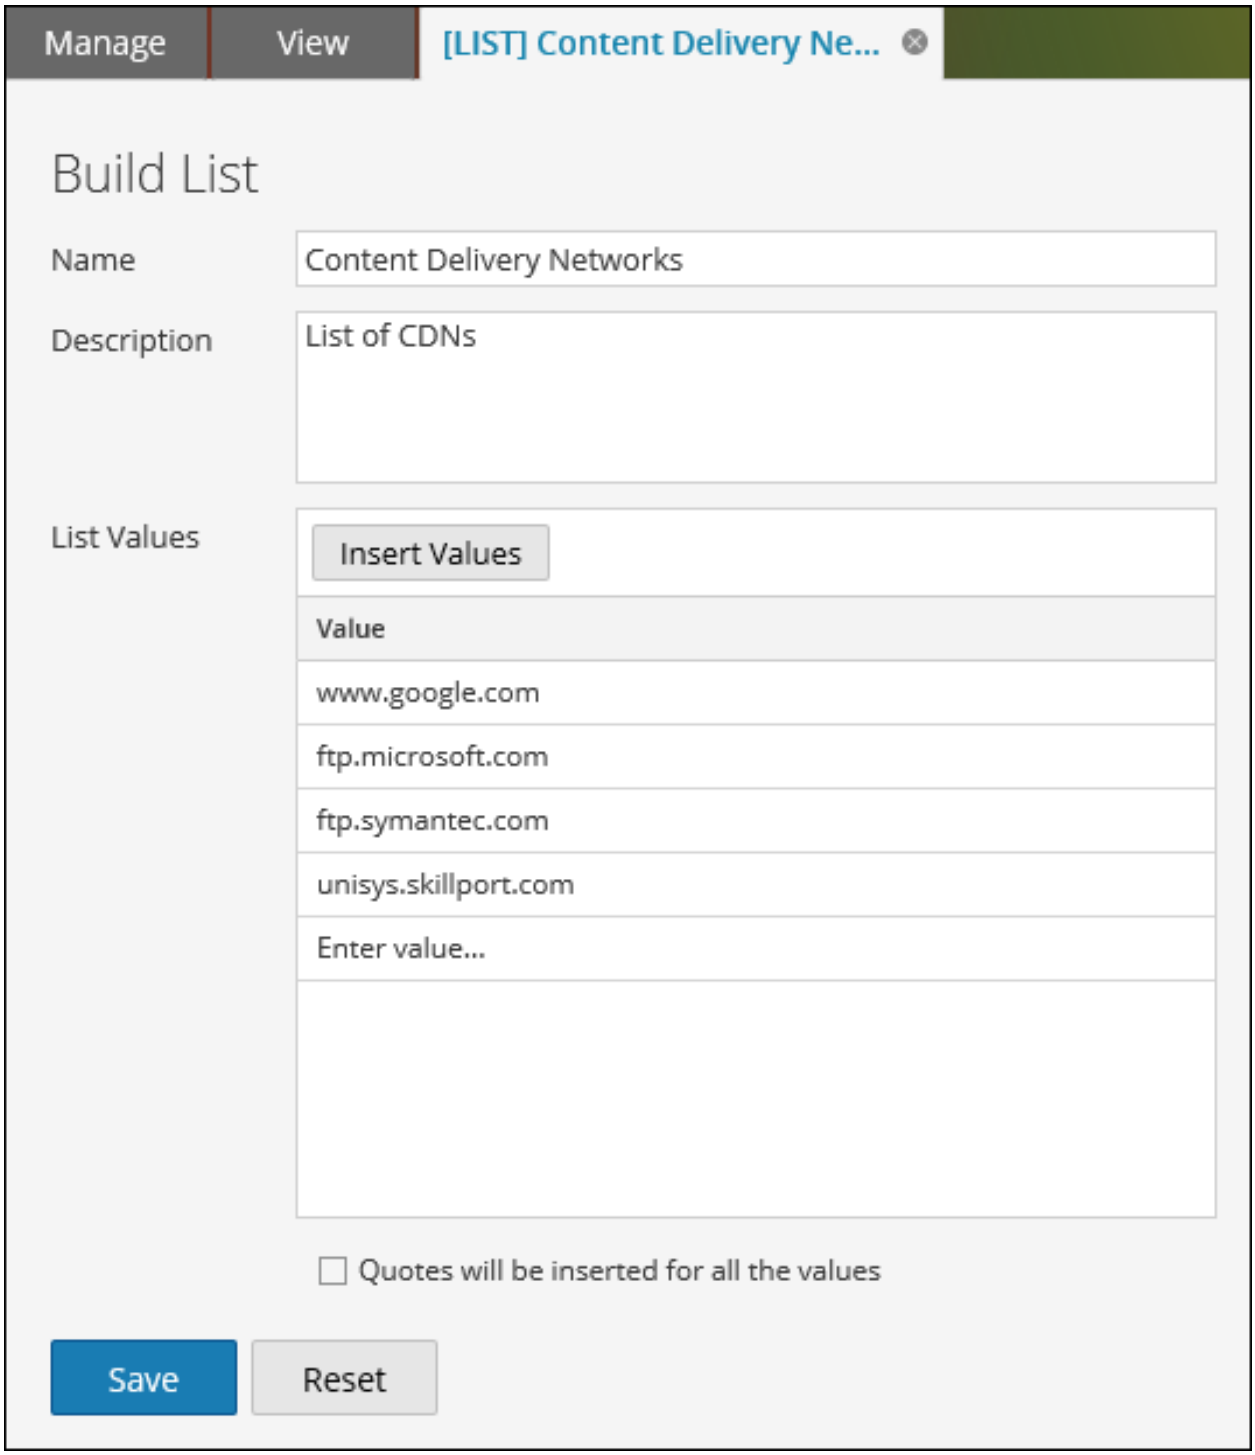
*You can complete these tasks here.

Related Topics

- List View
- Lists Permissions Dialog

Quick View

The following figure shows the Build List View.



To access this view

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

   The Lists view is displayed.

3. In the **Lists** toolbar, click .

   The Build List tab is displayed.

The following table describes the features in the Build List view.

Feature	Description
Name	Identifies and labels the list.
Description	Provides a short description for the list.
List Values	Provides the grid of values associated with selected list from the List Library panel. You can import these values from a file or list. You can also enter values manually.
Quotes will be inserted for all the values	Automatically includes quotes for the values at runtime if checkbox is selected. If the checkbox is not selected and if a value in the list contains a comma, then that value has to be enclosed within single quotes. This syntax does not apply to list values for an NWDB rule.
Save	Saves the rule which can be used to create a report, a chart or an alert.
Reset	Deletes all the information from the fields.

Build Report View

In the Build Report view, you can create a report, add text and rules, and schedule a report.

Workflow

This workflow shows the procedure to create and schedule a report.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report*	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

*You can complete these tasks here.

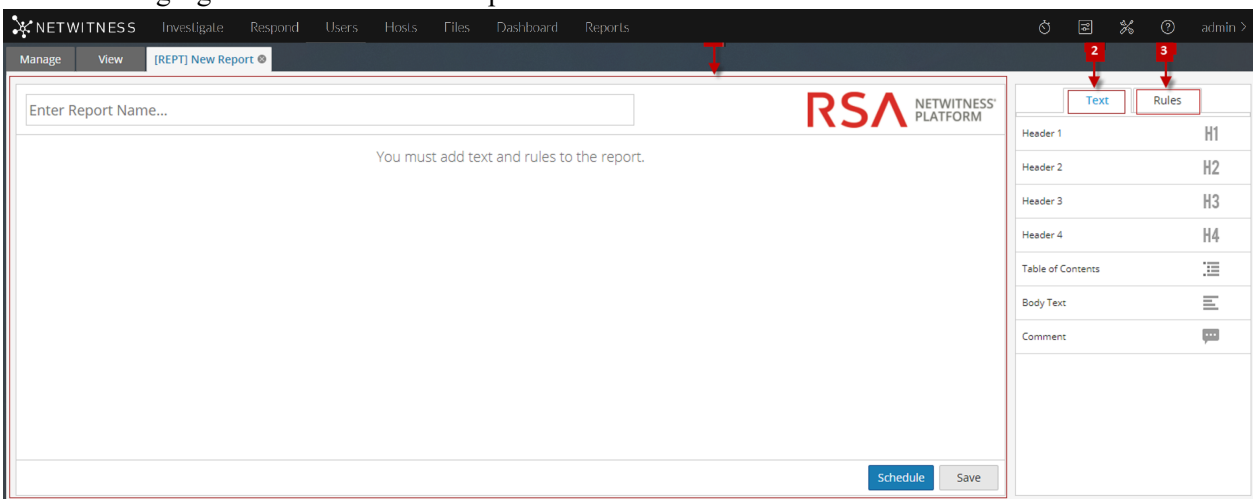
Related Topics

- Configure and Generate a Report
- Report View

- Scheduled Reports View
- Reports Permissions Dialog

Quick View

The following figure shows the Build Report View.



To access this view

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Reports view is displayed.

3. In the **Reports** toolbar, click + .

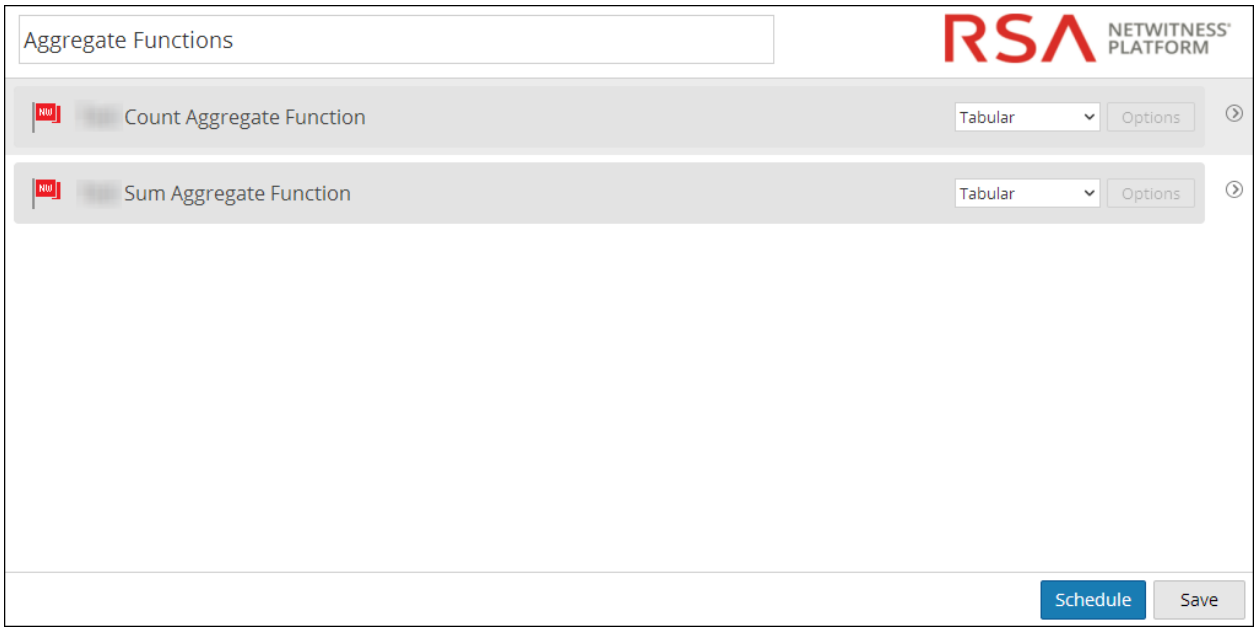
   The Build Report tab is displayed.

The Build Report view consists of the following panels:

1	Report Panel
2	Text Panel
3	Rules Panel

Report Panel

The Report panel allows you to create a report by assigning a name to the report. The content in a report depends on the items selected from the Text and Rules panels.



When you add rules to a report, you can change the output format of these rules either to tabular, area, line or pie by clicking the ⌄ button.

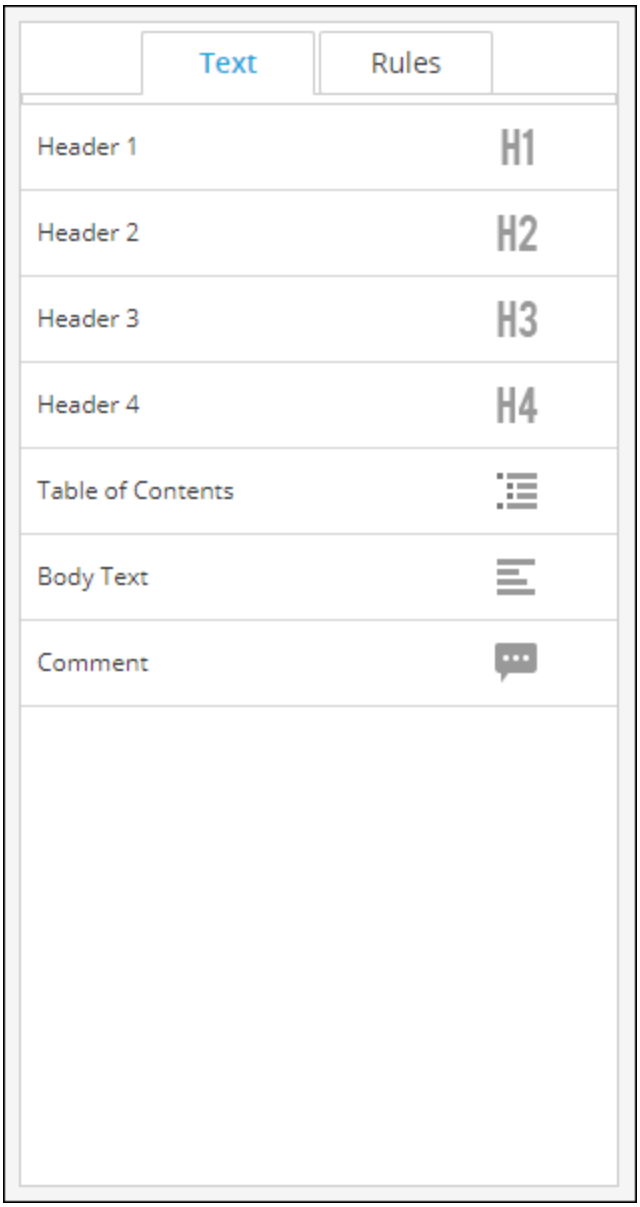
The following table lists the features of the Report Panel and the description.

Feature	Description
Name	This field allows you enter the name of the report.
Options	This field allows you to select the output format of the report such as Tabular, Area, Bar, Bubble, Column, Line, Pie, Step Line, Step Area, Spline Area and Spline.
Schedule	Clicking this option generates the report.
Save	Clicking this option saves the report.

Text Panel

The Text panel consists of a list of text elements that add to the look and feel of the report. You can use these text elements to format the report.

- To add more structure to reports, you can use these headers defined in the Text panel to indent up to four levels. This allows you to identify specific sections in a report that can be included in the Table of Contents for easy navigation in the report result.
- To add headers to the Report panel, drag and drop H1, H2, H3, or H4 onto the Report pane based on the desired level of indentation.



The following table lists the text elements used to format a report:

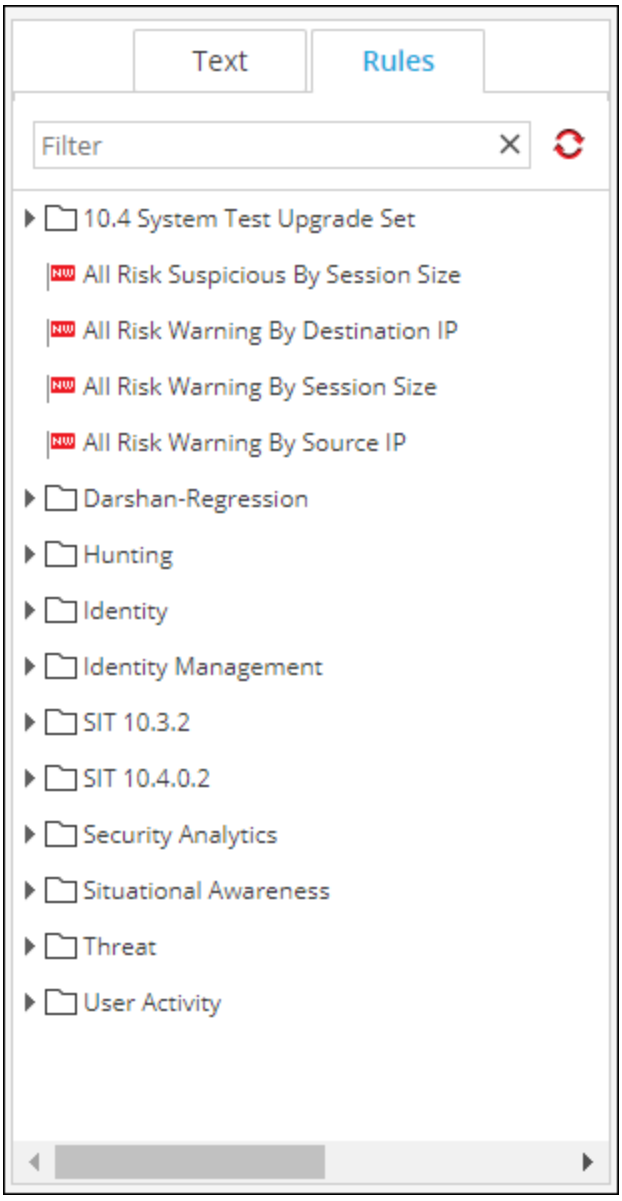
Text Elements	Description
Header 1 H1	The Header 1 element adds a first-level heading to the report definition.
Header 2 H2	The Header 2 element adds a second-level heading to the report definition.
Header 3 H3	The Header 3 element adds a third-level heading to the report definition.

Text Elements	Description
Header 4 	The Header 4 element adds a fourth-level heading to the report definition.
Table of Contents 	The Table of Contents adds table of contents to the report definition.
Body Text 	The Body Text element adds body text to the report definition.
Comment 	The Comment element adds comments to the report definition. **Note:** The Comment element is not displayed when you view all the reports.

Rules Panel

The Rules panel consists of a list of rules that are defined in the Rules. From the rules list, you can drag and drop rules onto the Report panel to associate those rules with the report.

You can search for a specific rule using search text box provided in the Rules panel.


Text
Rules
Filter
10.4 System Test Upgrade Set
All Risk Suspicious By Session Size
All Risk Warning By Destination IP
All Risk Warning By Session Size
All Risk Warning By Source IP
Darshan-Regression
Hunting
Identity
Identity Management
SIT 10.3.2
SIT 10.4.0.2
Security Analytics
Situational Awareness
Threat
User Activity

Build Rule View

The Build Rule view explains the actions and associated procedures that you can perform under Rules.

Workflow

This workflow shows the procedure to create or deploy a rule.



What do you want to do?

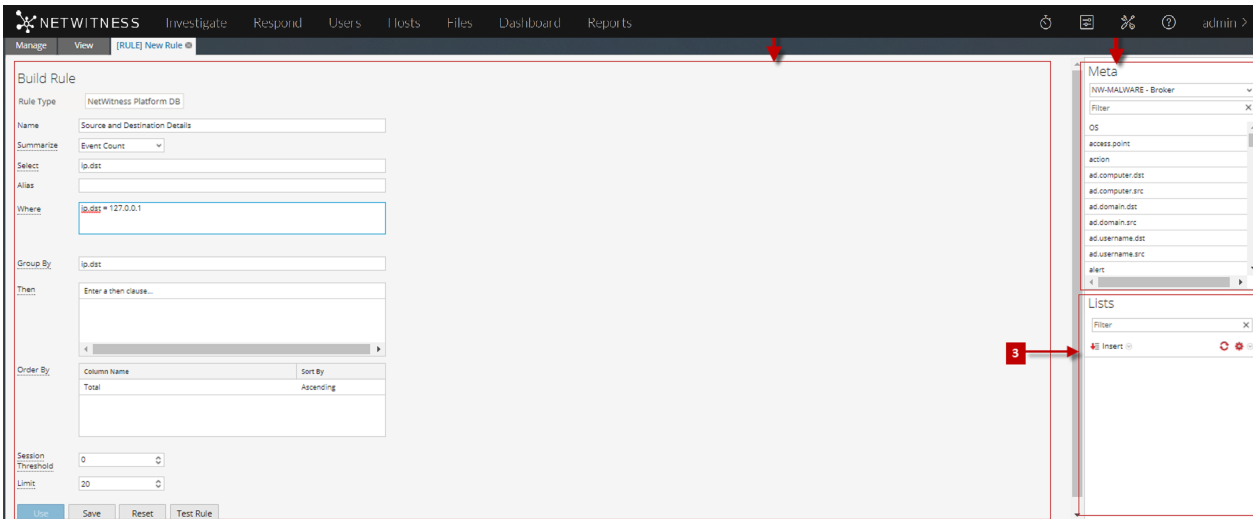
Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule*	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

*You can complete these tasks here.

Related Topics

- Rule Permissions Dialog
- Rule View

Quick View



To access the Build Rule view:

1. Go to **Reports**.
   The Manage tab is displayed.

2. In the **Rules** toolbar, click + > **NetWitness Platform DB**.
   The Build Rule view tab is displayed

Features

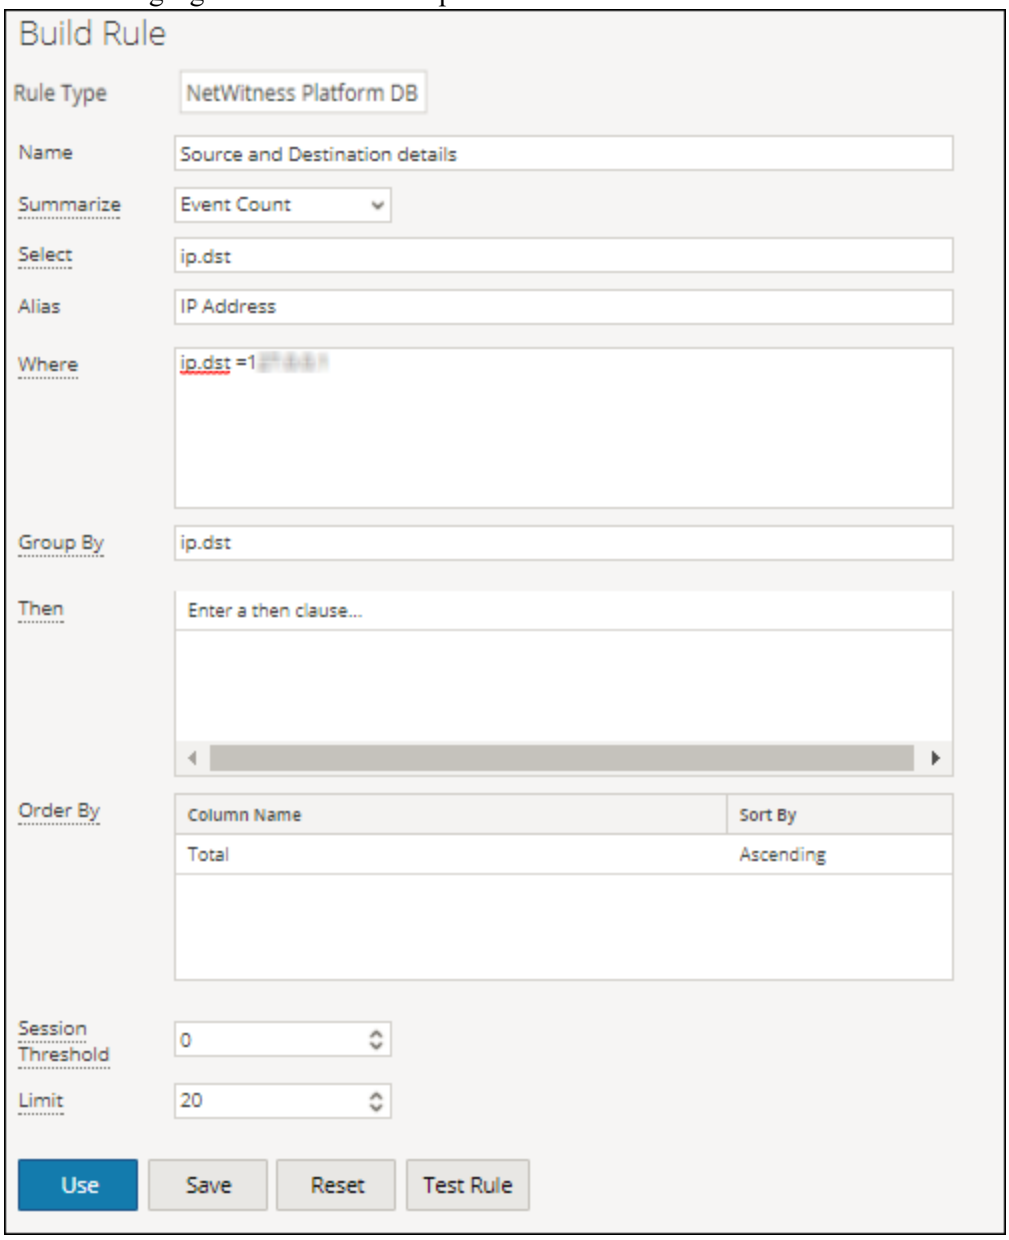
The Build Rule view includes the following panels.

1	Rule panel
2	Meta panel
3	Lists panel

Rule Panel

The Rule panel allows you to create a rule for the selected database type.

The following figure shows the Rule panel.



The following table describes the features in the Rule panel.

Feature	Description
Rule Type	A drop-down list of supported database types for which you can create rules. The options are: NetWitness DB and Warehouse DB.
Name	The name of the rule that you are creating or editing.
Summarize	A drop-down list of summarize options. The options are: None, Event Count, Packet Count, Session Count and Custom.
Select	The meta key for which you need the aggregate values; for example, ip.dest.

Feature	Description
Where	A Where clause that defines the conditions that trigger the rule execution; for example, ip.dest = 127.0.0.1.
Group By	The grouping method for the results. For example, specifying ip.dest produces a report in which like ip.dest values are grouped.
Then	A Then clause that defines the rule actions for additional processing on the output.
Order By	The sequencing method used to show results. For example, specifying Order By the value in the Total column, Ascending, produces a report in which the results are sorted in ascending order based on the value in the Total column.
Session Threshold	A selection list for the session threshold, which specifies maximum number of sessions that should be processed for aggregate functions.
Limit	A selection list for the maximum number of result rows to be fetched.
Use	Clicking Use enables you to use the Rule to generate a Report, Alert of Chart.
Save	Clicking Save saves the rule that you are editing and the Build Rule panel remains open. Before testing a rule, you must save it if you want to keep your changes.
Reset	Clicking Reset clears all the field information .
Test Rule	Clicking test rule opens the Test Rule dialog.

Test Rule Dialog

To access the Test Rule view:

1. Go to **Reports**.

   The Manage tab is displayed.

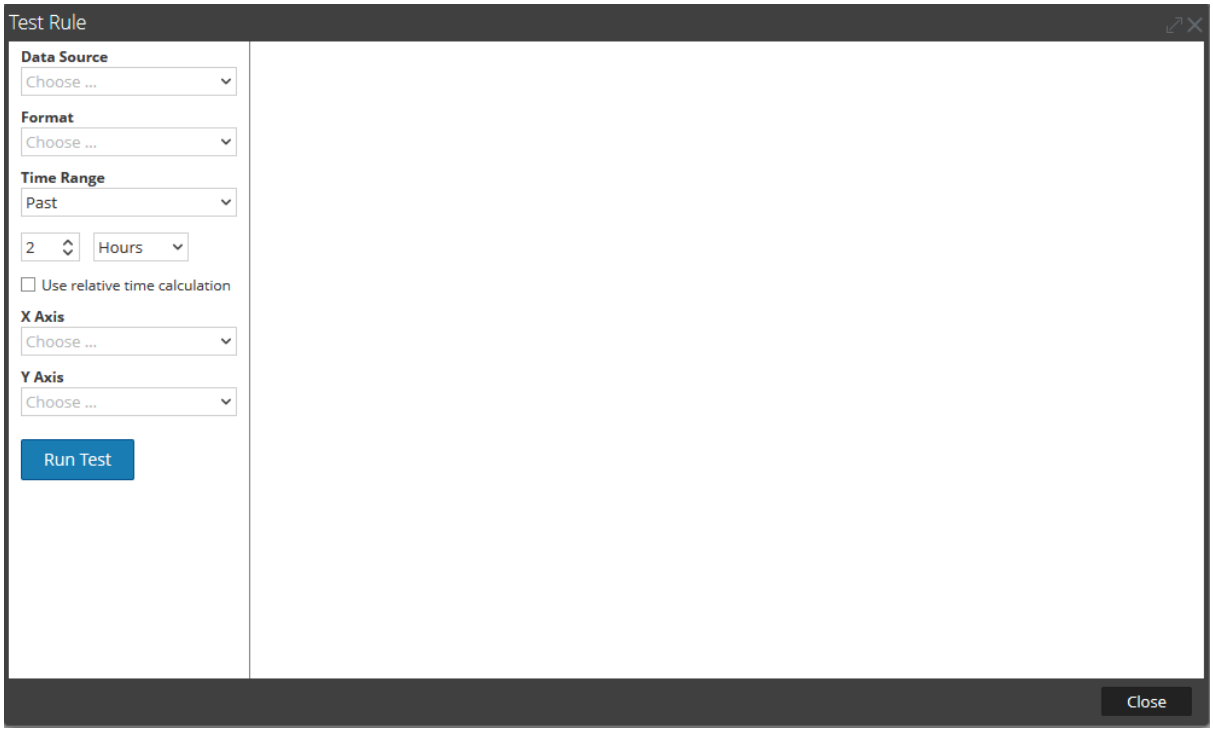
2. In the **Rules** panel, do one of the following:

   - Select a rule and click  in the Rules toolbar.
   - Click  > **Edit**.

     The Build Rule view tab is displayed.

3. Click **Test Rule**.

The Test Rule view is displayed.



The following table describes the features in the Test Rule Dialog.

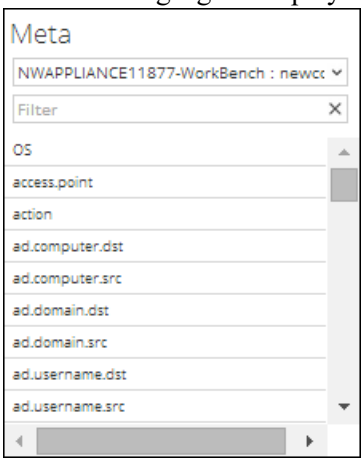
Feature	Description
Data Source	A drop-down list of data sources for the type of rule you are testing. Possible data sources are: Concentrator, Broker, Decoder or Log Decoder.
Format	A drop-down list of the formats for displaying results for the rule. Possible formats are: Tabular, Area, Bar, Bubble, Column, Line, Pie, Step Line, Step Area, Spline Area, and Spline.
Time Range	A drop-down list of time range specification methods. • Selecting Past allows you to specify a number of years, months, days, weeks, or hours. For example, Hours, Days, Weeks, Months, or Years. • Selecting Range allows you to specify a date range and time period. For example, start date to end date. In the user interface, the date or time displayed depends on the time zone profile selected by the user.
Use relative time calculation	Selecting this option calculates the time range relative to the current time.

Feature	Description
X Axis	X-Axis and Y-Axis specify the metadata to be plotted in charts. In the X-Axis drop-down list, the meta types for the `Group by` setting in the rule are listed. You can select multiple meta types when the rule has a single `Group by` setting. For Custom Rules with multiple `Group by` values, you can select only the first meta type for the X-Axis.
Y Axis	In the Y-Axis drop-down list, the aggregate functions used in the rule are listed. Sum, Count, Countdistinct and Average are the supported aggregate functions for rules. You can select one or more aggregate functions.
Run Test	Clicking Run Test executes a test of the rule last saved in the Rule Builder dialog. When the test is complete, the rule data (if any) for the selected time range is displayed.

Meta Panel

The Meta panel provides a list of available meta types that you can use to build the rule. You can use the meta types in the Select, Where, and Then clauses. The Reporting Engine maintains an active list of the available meta names by continuously synchronizing with the data source to which it is connected.

The following figure displays the Meta panel.



The following table describes the features in the Meta panel.

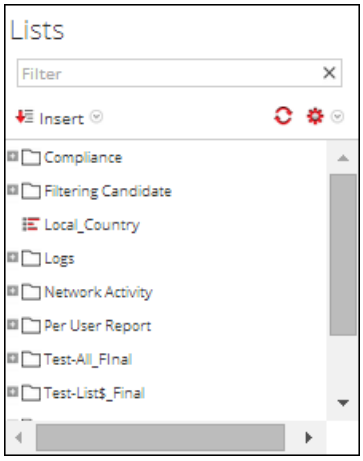
Operation	Description
Choose	Based on the rule type that you have selected, the available data sources are displayed in the drop-down list of the Meta panel. Select the required data source. The available meta types for the data source are displayed. Select a meta.
Filter	Filter the meta for a specific meta value.

Lists Panel

A List is a placeholder for a set of values that you can use in a meta or a variable. For example, you can define a list with all the whitelisted event source IP addresses. Once the List is defined then you can use the List name in the rule. This provides the flexibility of adding, modifying, and deleting the list values.

The Lists panel is a collection of Lists. The Reporting Engine maintains an active list of the available list names by continuously synchronizing with the collection to which it is connected.

The following figure displays the Lists panel.



The following table describes the features in the Lists panel.

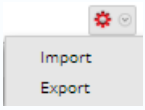
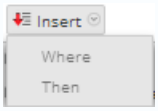
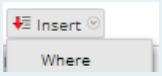
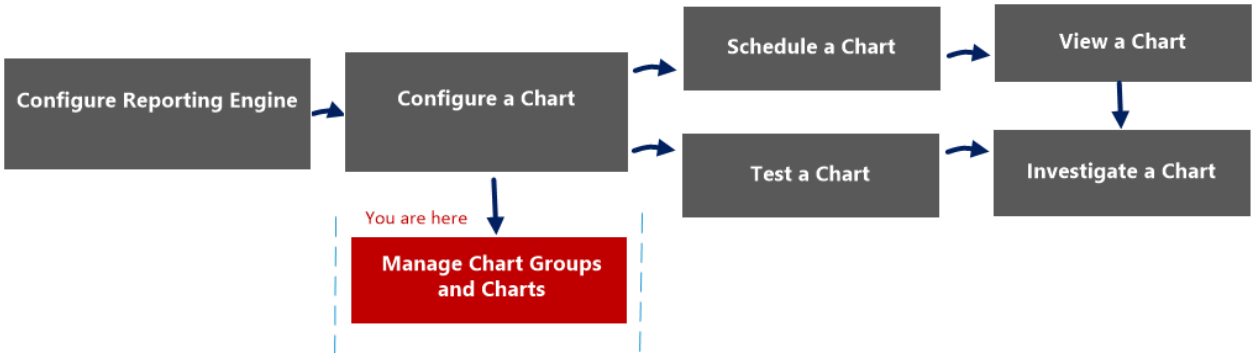
Operation	Description
	Import or Export a list.
	Refresh the Lists.
	If you select the **NetWitness DB** rule type, the options Where and Then are displayed. Insert the list in the Where or Then clause in the rule.
	If you select the **Warehouse DB** rule type, the option Where is displayed. Insert the list in the Where clause in the rule.

Chart Permissions Dialog

In the Chart Permissions dialog, you can manage access permissions for user roles at the chart and chart group level. Only a user with the 'Read & Write' permission can configure the chart in the Reporting module.

Workflow



What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart	Schedule a Chart
Administrator/ Analyst	View a chart	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart*	Manage a Chart Group and Chart

*You can complete these tasks here.

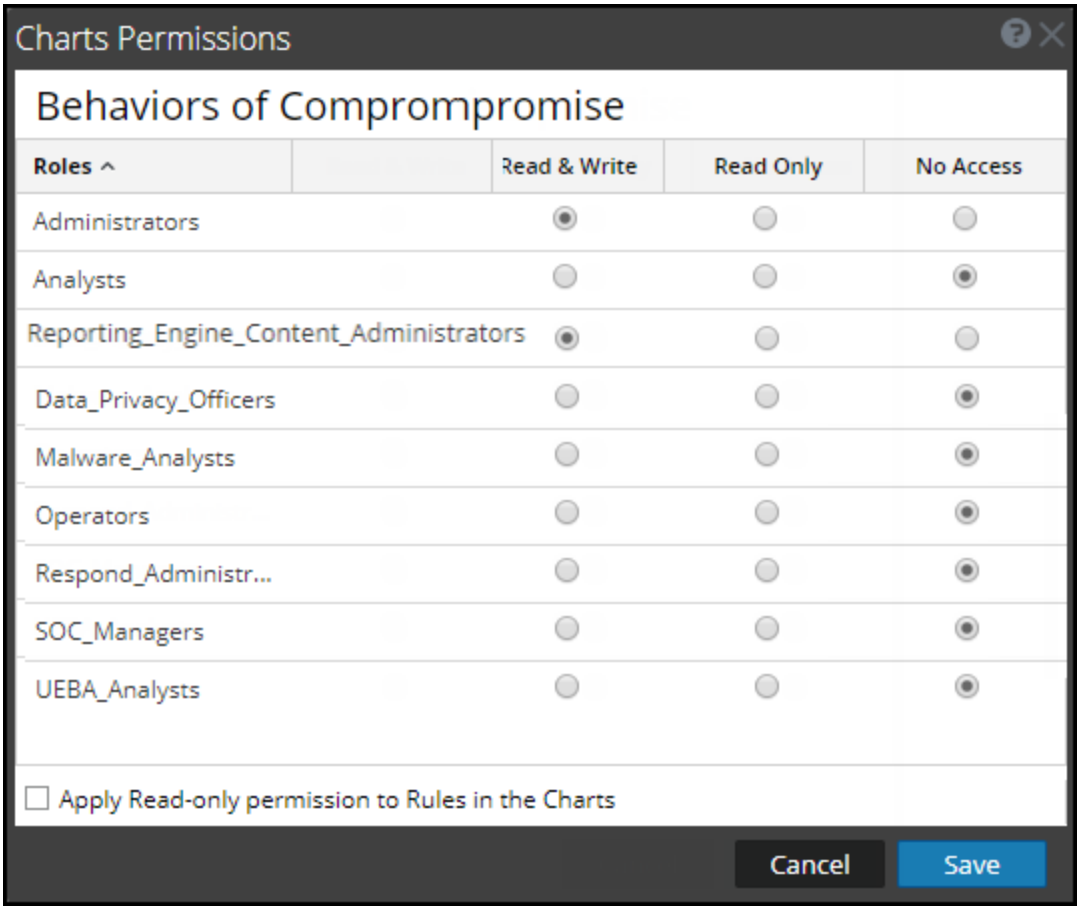
Related Topics

- Configure and Generate a Chart

Quick View

The Chart permissions dialog allows you to set chart permissions depending on the user role.

The following figure is an example with the important features labeled.



1. Click **Reports** to view the Manage tab.
2. Click **Charts** to open the Chart view.
3. In the **Charts** panel, select a report and click > **Permissions**. The Chart Permissions dialog box is displayed.
4. Based on the user role, select the appropriate options.
5. (Optional) Select the checkbox if you want to automatically provide read access permission to dependent rules.
6. Click **Save**.

The following table lists the columns in the Charts Permission dialog.

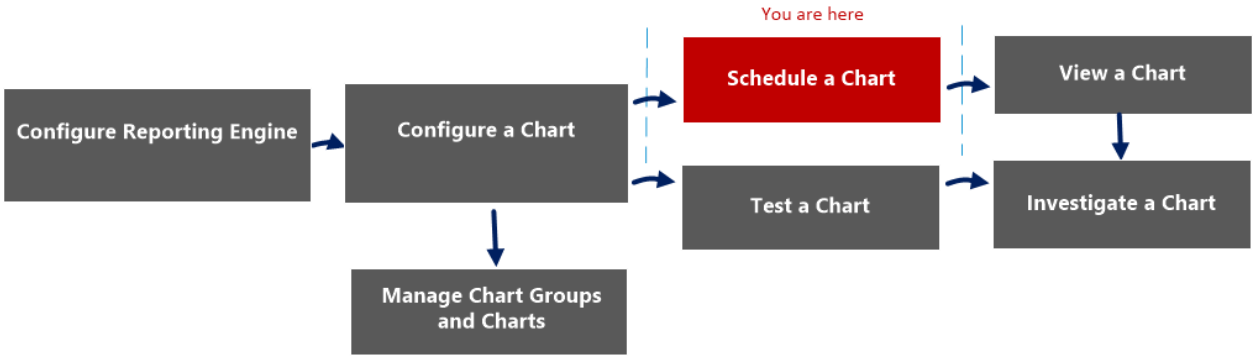
Column	Description
Roles	Displays all the user roles in the NetWitness user interface.
Read & Write	Allows you to apply 'Read&Write' access to the chart.
Read Only	Allows you to apply only 'Read' access to the chart.

Column	Description
No Access	By selecting this permission, you cannot access or view the chart.
☐ Apply these permissions to sub-groups and Charts in this group	Allows you to apply permissions to the chart group, subgroups in the group and charts in the group. **Note:** This checkbox is populated only when you set access permissions for a Chart Group.
☐ Apply Read-only permission to Rules in the Charts	Allows you to automatically apply permissions to the rules in the charts.
Cancel	Cancels all the changes made to the permissions.
Save	Saves the selection and provides access to the role based on the selection.

Chart View

In the Chart View, you can see the available charts and groups in a grid format and also schedule them by enabling the charts.

Workflow



What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart*	Schedule a Chart
Administrator/ Analyst	View a chart	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart	Manage a Chart Group and Chart

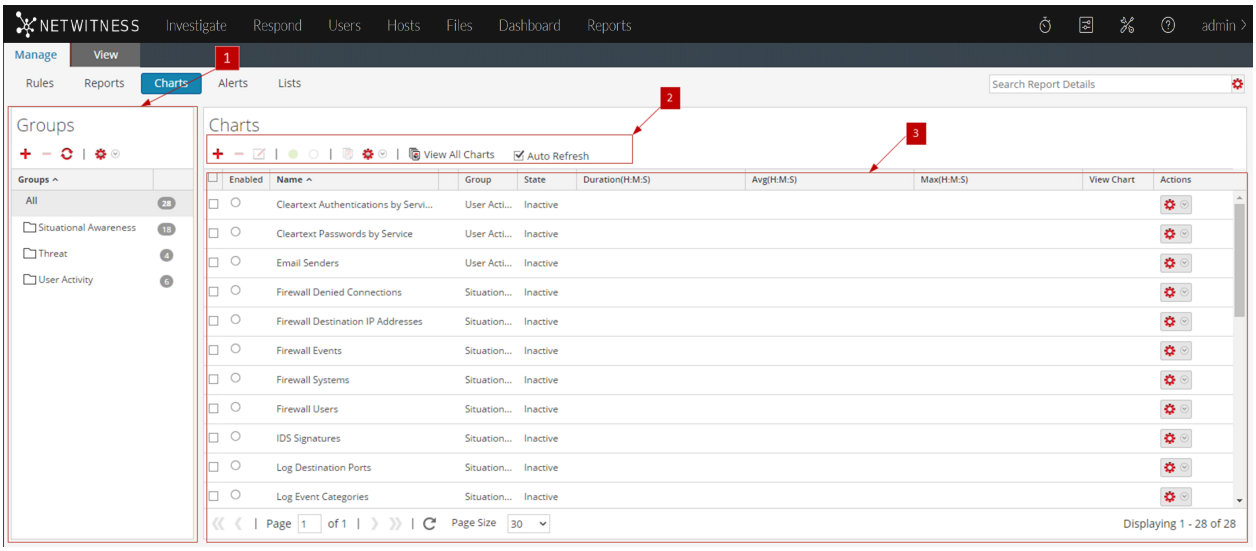
*You can complete these tasks here.

Related Topics

- Configure and Generate a Chart

Quick View

The following figure is an example with the important features labeled.

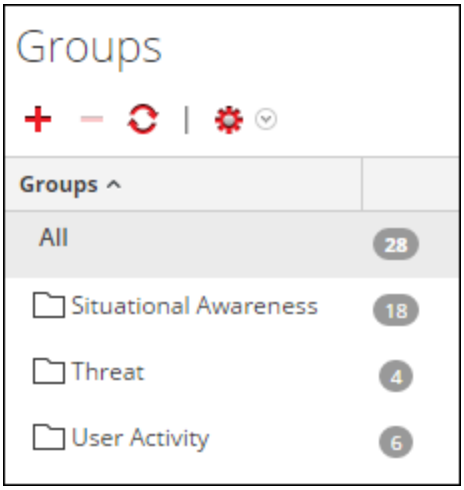


The Chart view includes the following panels:

1 Charts Groups panel
2 Charts toolbar
3 Charts panel

Charts Groups Panel

The Charts Groups panel allows you to organize charts in a group. You can create a group, add charts to the group and move charts among groups. The following figure shows the Charts Groups panel.



The Charts Groups Panel includes the following options:

Feature	Description
+	Adds a new chart to the Reporting module.

Feature	Description
	Deletes one or more selected charts.
	Edits a chart.
	Refreshes the view.
	Provides the following options: Import, Export and Permissions.

Charts Toolbar

The Charts toolbar allows you to add, modify, delete, duplicate, activate, deactivate, import and export a chart. You can also set access permissions for charts in a group.



The Chart toolbar includes the following options:

Feature	Description
	Adds a new chart to the Reporting module.
	Deletes one or more selected charts.
	Edit charts.
	Enables the selected charts.
	Disables the selected charts.
	Creates a duplicate copy of the selected chart.
	Provides the following options: Import, Export, Export as Text and Permissions.
View All Charts	Displays all the executed charts.
Auto Refresh	Automatically refreshes the charts list.
	Allows users to filter unused charts.

Charts Panel

The Charts Panel presents all the charts in a tabular or grid format.



The following table lists the columns in the Chart panel and their description.

Feature	Description
Enabled	● - The chart is enabled. ○ - The chart is disabled.
Name	The name of the chart.
Group	The Chart Group to which the chart belongs.
State	The state of the chart: • Queued • Completed • Failed
Duration (H:M:S)	The time taken to execute the latest chart.
Avg(H:M:S)	The average time taken to run the chart.
Max(H:M:S)	The maximum time taken to run the chart.
View Chart	A hyperlink that redirects to the View a Chart panel.
	The actions menu has the following options: Enable, Disable, View, Delete, Edit, and Export.

Execution History Panel

The Execution History panel allows you to fetch and display history details.

Workflow

This workflow shows the procedure to view report or report groups.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports*	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access control for lists, rules or Reports	Manage Lists, Rules or Reports

*You can complete these tasks here.

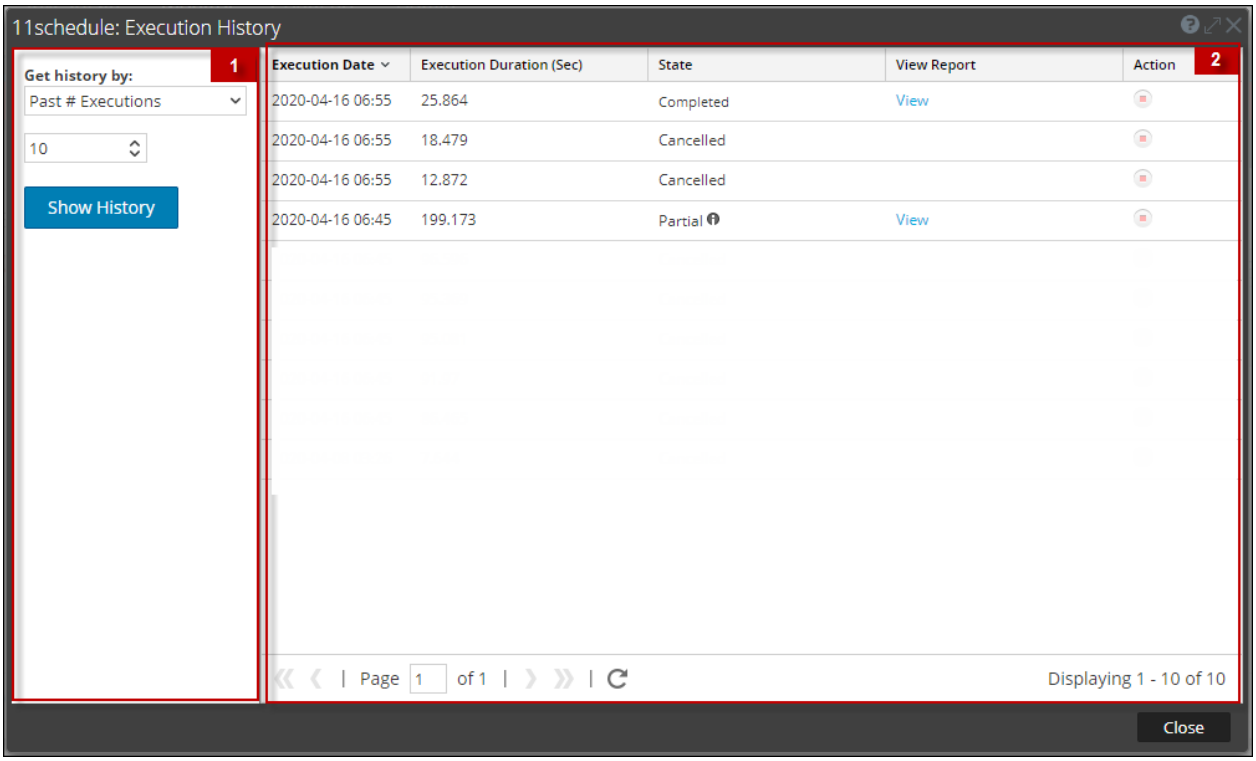
Related Topics

- Configure and Generate a Report
- Generate List Panel

- Scheduled Reports View

Quick View

The following figure is an example of the Execution History view.



Features

The View Execution History has the following panels:

1 Execution History Options panel
2 Execution History Output panel

To access this view:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, do one of the following:

   - Hover the mouse over a report and click > **View Scheduled Reports**.
   - Click **#Schedules** column.
     The Schedule Reports view is displayed with the status of each of the scheduled report.
4. Select a scheduled report and do one of the following:

- Click  > **Execution History**.
- Click  from the **Scheduled Reports** Toolbar Panel.

Execution History Options Panel

The Execution History Options panel allows you to fetch the history details based on either past n number of scheduled reports or a specific date range.

The following table lists the operations in the Execution History Options panel:

Operation	Description
Get history by:	This is the criteria to view the execution history: • **Past # Executions**: The past n number of scheduled reports. By default this option is displayed. • **Range (specific)**: The start date and end date for the date range. **Note:** The **From** and **To** field is populated in the NetWitness UI only when you select 'Range (Specific)' from the **Get history by** list.
From	The start date for the date range.
To	The end date for the date range.
Count	The number of execution history of the scheduled report to be displayed.
Show History	Shows the history details based on the selected criteria.

Execution History Output Panel

The Execution History Output panel displays the history details with the execution date, execution duration (seconds), state of the scheduled report, and a link to view the report.

The following table lists the various columns in the Execution History Output panel:

Column	Description
Execution Date	The date on which the scheduled report was executed. By default, the execution date is in descending order.
Execution Duration (Sec)	The time duration taken to execute the scheduled report.

Column	Description
State	The state of the scheduled report: • Scheduled: If a report is scheduled to run on an hourly, daily, weekly, monthly, or later time, the state of the report is displayed as scheduled, for the first run. • Queued: If a report is still waiting to get executed, the state of the report is displayed as queued. • Running: If the report schedule is in progress, the state of the report is displayed as running. • Partial: If in a report with several rules, a single rule execution failed or an output action failed or creation of PDF/CSV failed, the state of the report is displayed as partial. For example, consider a report with five rules and four rules are executed successfully and one fails, then the state is displayed as Partial. • Failed: If in a report with several rules, all the rule schedule executions failed, the state of the report is displayed as failed. • Completed: If a report schedule is successfully executed, the state of the report is displayed as completed. • Canceled: When cancel request is completed, the state of the report is displayed as canceled. • Inactive: If a report schedule is disabled, the state of the report is displayed as Inactive. • Not available: If the report schedule executed information is not available, the state of the report is displayed as not available.
View Report	The hyperlink to View a Report on full screen.
Action	Displays the icon to stop the execution schedule.
Close	Closes the execution history view.

Generate List Panel

The Generate List dialog allows you to generate and customize a list.

Workflow

This workflow shows the procedure to create and schedule a report.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report*	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports*	Manage Lists, Rules or Reports

*You can complete these tasks here.

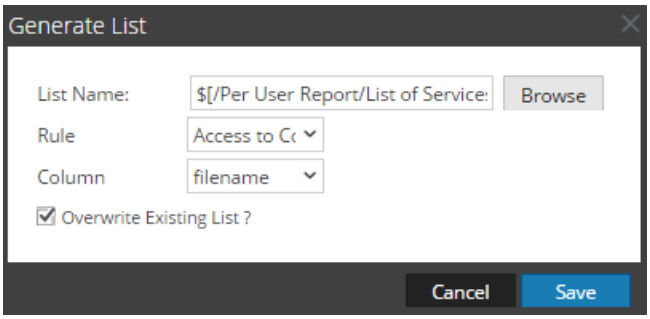
Related Topics

- List View
- Build List View

- Lists Permissions Dialog

Quick View

The following figure is an example of the Generate List dialog.



To access this view:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** panel, select a report and click > **Schedule Report**.

   The Schedule a Report view tab is displayed.

4. In the **Output Actions** section, in **Dynamic List** panel, click .

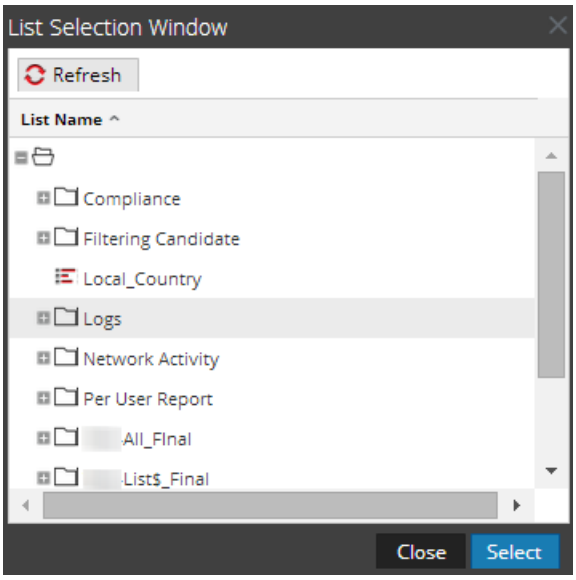
   The Generate List dialog is displayed.

Features

The following table lists the features in the Generate List dialog.

Field	Description
List Name	The name of the list chosen from the List Selection panel.
Browse	Click this button to select a list from the List Selection Window dialog.
Rule	Select a rule to be used to create the list.
Column	Select a value for the column.
Overwrite Existing List?	Overwrites the existing list.
Save	Adds the desired list to the Generate List panel of the Schedule Report view.

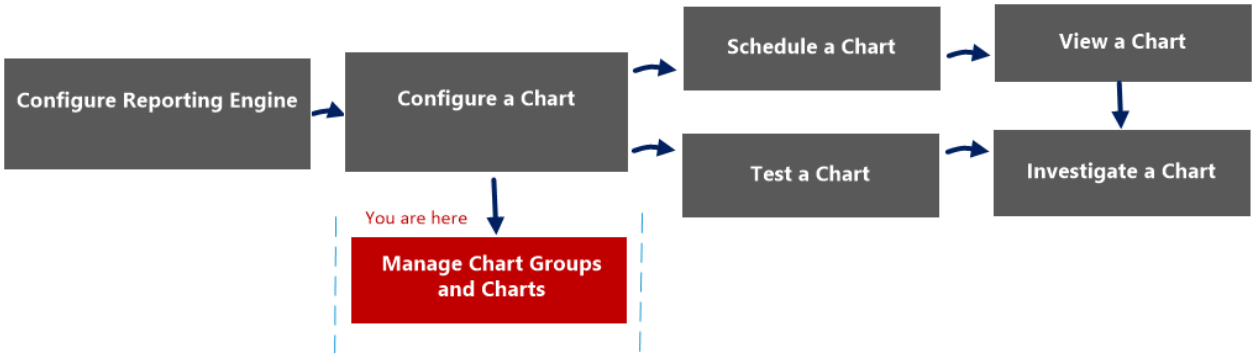
The List Selection Window dialog consists of lists that are defined in the Lists panel. Here, you can select a list to associate it with the report. The following figure shows the dialog.


List Selection Window
Refresh
List Name
Compliance
Filtering Candidate
Local_Country
Logs
Network Activity
Per User Report
All_FInal
List$_Final
Close
Select

Import Chart Dialog

In the Import Chart dialog, you can import charts containing subgroups and charts from other instances of NetWitness into the Chart Groups panel. Charts must be in a valid binary file that was exported from another NetWitness instance.

Workflow



What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart	Schedule a Chart
Administrator/ Analyst	View a chart	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart*	Manage a Chart Group and Chart

*You can complete these tasks here.

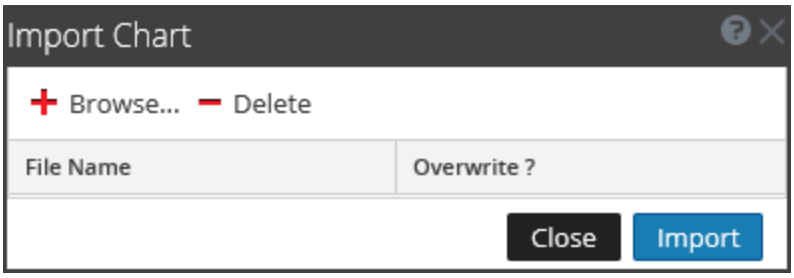
Related Topics

- Configure and Generate a Chart

Quick View

This dialog displays differently when you use it to import groups containing subgroups and charts from other instances of NetWitness into the Chart Groups panel.

The following figure is an example of the Import Chart dialog.



1. Click **Reports** to view the Manage tab.
2. Click **Charts** to open the Chart view.
3. In the **Charts Groups** panel, select a folder to import the file.
4. In the **Charts Groups** panel or **Charts** toolbar, click > **Import** to import the file.

The following table describes the features in the Import Chart dialog.

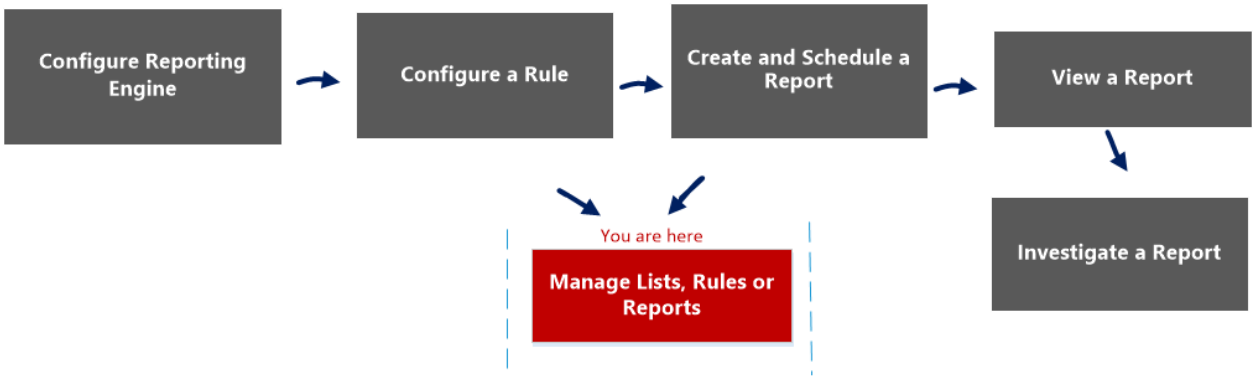
Feature	Description
Browse	Displays a view of the local file system so that you can select the chart to be imported.
Delete	Deletes an imported report from the list of imported charts.
File Name	Displays a list of chart files that will be imported to your Charts module when you click Import.
Overwrite?	Allows you to select the option to overwrite an existing version of the chart you are importing. If you do not select the Overwrite option, a duplicate file is imported and no error message is displayed.
Close	Closes the dialog. If you have charts to select for import, but have not clicked Import. The charts are not imported, and are not saved in this dialog.
Import	Imports the selected charts to your Charts module.

Import Report Dialog

In Import Report dialog, you can import groups containing subgroups and reports from other instances of NetWitness into Report Groups panel. Reports must be in a valid binary file that was exported from another NetWitness instance.

Workflow

This workflow shows the procedure to manage reports or report groups.



What do you want to do?

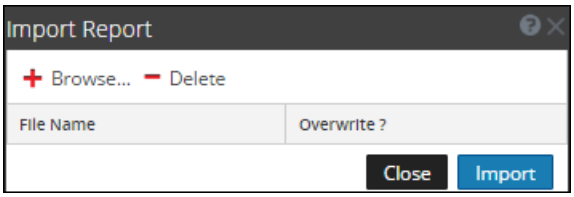
Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports*	Manage Lists, Rules or Reports

*You can complete these tasks here.

Related Topics

- Configure and Generate a Report
- Report View
- Build Report View
- Reports Permissions Dialog

Quick View



To access the Import Report dialog:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports Groups** panel, select a folder to import the file.
4. Do one of the following:
   - In the **Reports Groups** panel, click > **Import** to import a group.
   - In the **Reports** toolbar, click > **Import** to import a report.

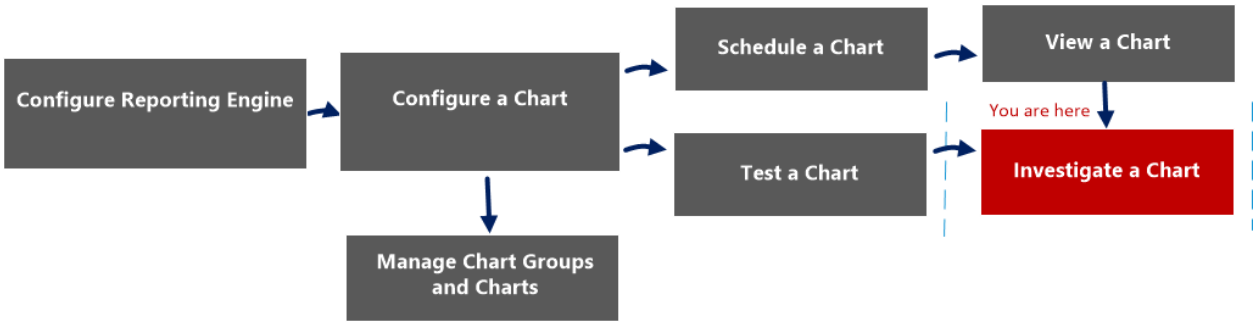
The following table lists the features of the Import Report dialog.

Feature	Description
Browse	This option displays a view of the local file system so that you can select the report to be imported.
Delete	This option deletes an imported report from the list of imported reports.
File Name	Displays a list of report files that will be imported to your Reports module when you click Import.
Overwrite?	Allows you to select the option to overwrite an existing version of the report you are importing. If you do not select the Overwrite option, a duplicate file is imported and no error message is displayed.
Close	This option closes the dialog. If you select a report and not clicked Import. The reports are not imported, and are not saved in this dialog.
Import	This option imports the selected reports to your Reports module.

Investigate a Chart View

In the Investigate a Chart view, you can view and investigate chart details. There are options for filtering and sorting the information in the chart, as well as options for the type of chart, the number of items to chart, and charting values or totals. When viewing a chart, you can open the charted sessions in the Investigation module and save the chart as a PDF.

Workflow



What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engin" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart	Schedule a Chart
Administrator/ Analyst	View a chart	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart*	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart	Manage a Chart Group and Chart

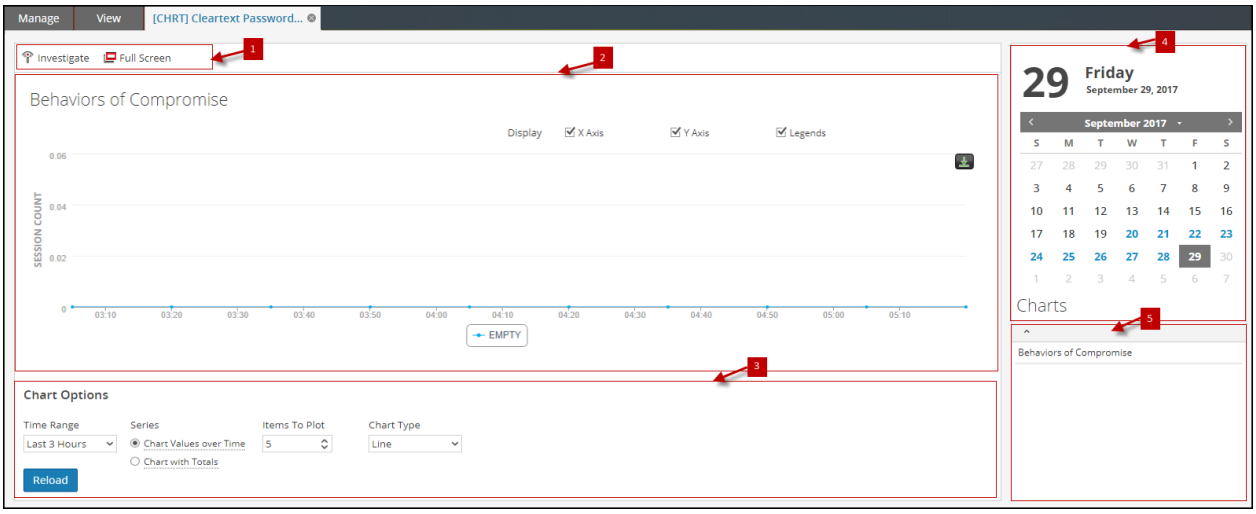
*You can complete these tasks here.

Related Topics

- Configure and Generate a Chart

Quick View

The following figure is an example with the important features labeled.



The View a Chart panel includes the following panels:

1. Chart toolbar
2. Chart Output panel
3. Chart Calendar panel
4. Chart Options panel
5. Chart Executed list

Chart Toolbar

The Chart toolbar has options that allow you to investigate, and view the chart on another screen.



The following table lists the options in the Chart toolbar.

Operation	Description
Investigate	Investigates the chart details.
Full Screen	Displays the chart on a full screen.

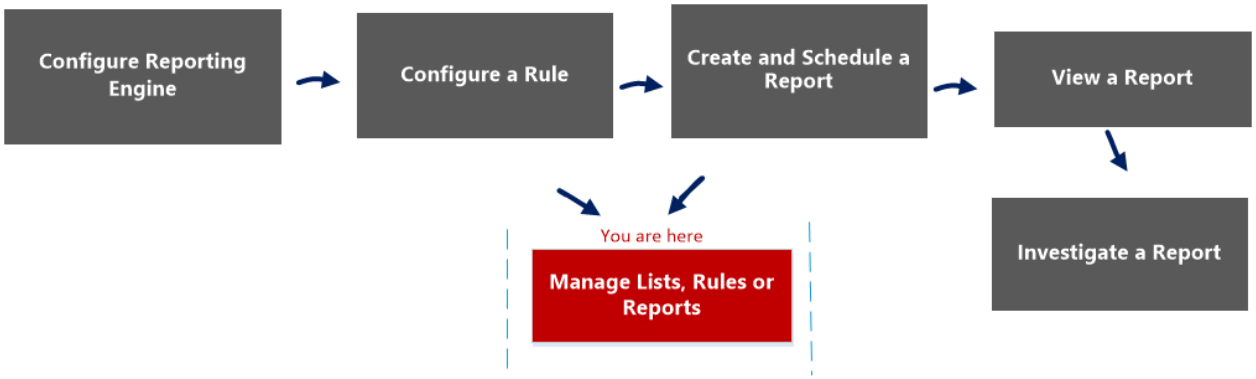
Lists Permissions Dialog

In the Lists Permissions dialog, you can manage access permissions for a user role at the list or list group level. Only a user with **Read and Write** permission can configure the list in the Reporting Module.

Workflow

This workflow shows the procedure to manage lists or list groups. You can set access control at the list or list group level so that only users with specific roles can access the lists. You can use lists to define rules for generating reports, charts and alerts.

You must ensure that Reporting Engine is configured on NetWitness.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports*	Manage Lists, Rules or Reports

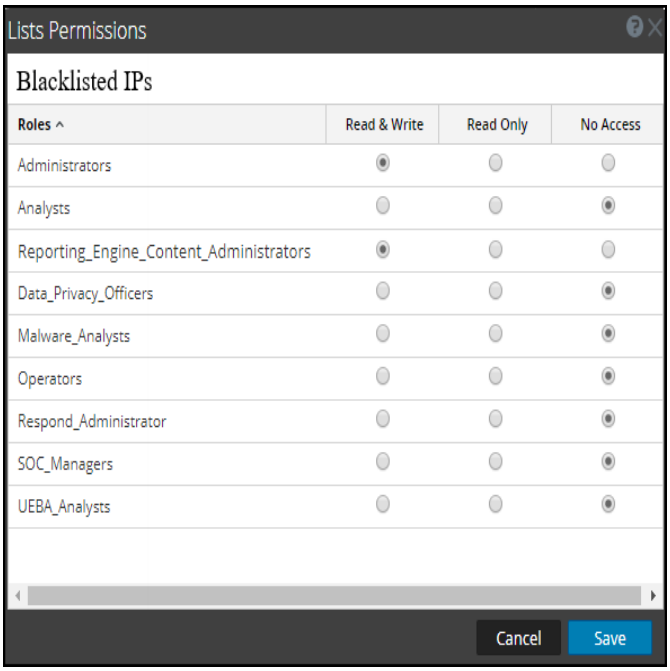
*You can complete these tasks here.

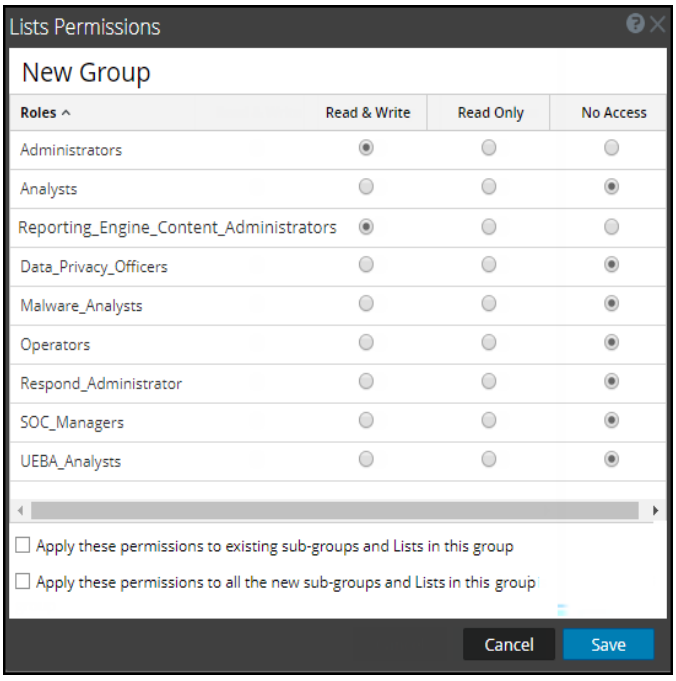
Related Topics

- List View
- List section in the "Role Permissions" topic in the *System Security and User Management Guide*.

Quick View

The following figures are examples of the Lists Permissions dialog and List Group Permission dialog:





To access this view

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Lists**.
   The Lists view is displayed.
3. In the **Lists** view, select a report.
4. In the **Lists** toolbar, click  > **Permissions**.
   The Reports Permissions dialog is displayed.

The following table describes the features in the Lists Permissions dialog:

Feature	Description
Roles	Describes roles of the users logged into the NetWitness user interface.
Read & Write	Allows users to access, view, edit, delete, import, and export lists on the Lists view. Users can also change the permission on the rule.
Read Only	Allows users to only access and view the list on the lists view.
No Access	Doesn't allow users to access or view the lists.
Apply these permissions to subgroups and lists in this groups	Automatically applies permissions to the subgroups and lists in the groups, if checkbox is selected.
Cancel	Cancels all the changes made to the permissions.
Save	Saves the selections and provides access to the roles based on the selections.

List View

In the List view you can see available lists and groups in a grid.

Workflow

This workflow shows the procedure to define lists or list groups. You can set access control at the list or list group level so that only users with specific roles can access the lists. You can use lists to define rules for generating reports, charts and alerts.

You must ensure that Reporting Engine is configured on NetWitness.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule*	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

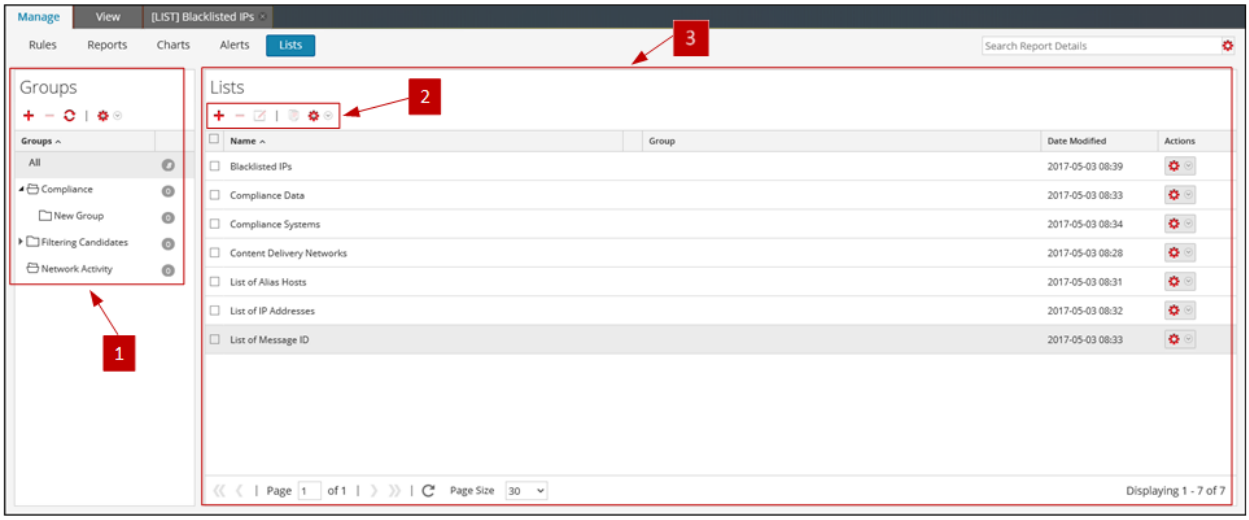
*You can complete these tasks here.

Related Topics

- Lists Permissions Dialog
- Build List View

Quick View

The following figure shows the List view.



To access this view

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Lists**.

   The Lists view is displayed.

The List view includes the following panels:

1 Lists Groups panel
2 Lists toolbar
3 Lists panel

Lists Groups Panel

The Lists Groups panel provides a list of groups used to organize lists and has a toolbar that allows you to create and manage the groups.

Feature	Description
	Allows users to add a new group to the Reporting module.
	Allows users to delete groups.
	Refreshes the view.

Feature	Description
	Allows users to access following options: Import, Export and Permissions.
	Allows users to filter unused lists.

You can perform the following actions using the Lists Groups panel.

- Refresh lists in a group.
- Move lists between different groups. You can move a list from one group to another by dragging and dropping the list in the required group.
- Create list groups.
- Delete list groups.
- Import list groups.
- Export list groups.
- Set access control for list groups.

Lists Toolbar

Feature	Description
	Allows user to add a new list to the Reporting module.
	Allows user to delete one or more selected lists.
	Allows user to edit lists.
	Creates a duplicate copy of the selected list.
	Allows user to access the following options: Import, Export and Permissions.

Lists Panel

The Lists panel displays all the lists defined in a tabular format.

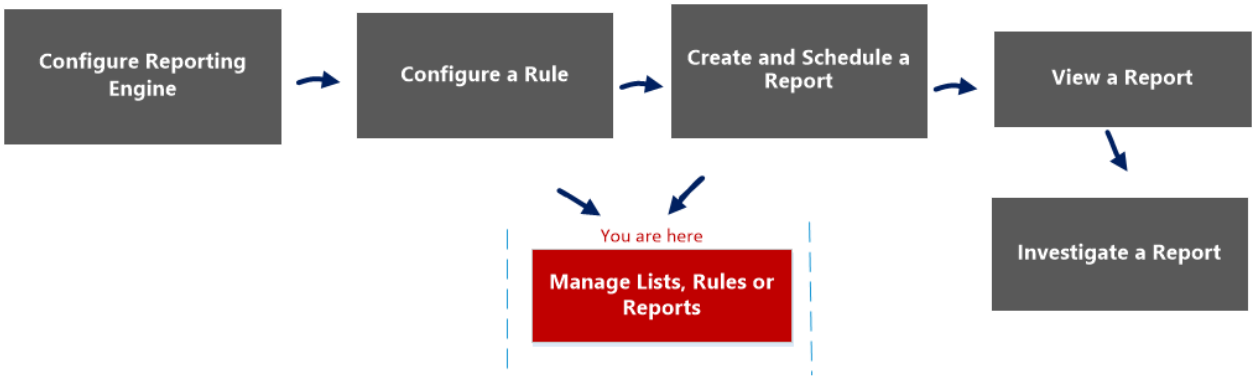
Column	Description
Name	Displays the name of the list. **Note:** For **Name** field, the icon to extend the column size is not displayed at the end of the column field. You have to hover the mouse a little to the left side to see the icon for extending the column.
Group	Displays the list group to which the list belongs.
Date Modified	Displays the date and time when the list was modified.

Reports Permissions Dialog

In the Reports Permissions dialog, the users with 'Read & Write' access permission can configure permissions.

Workflow

This workflow shows the procedure to manage reports or report groups.



What do you want to do?

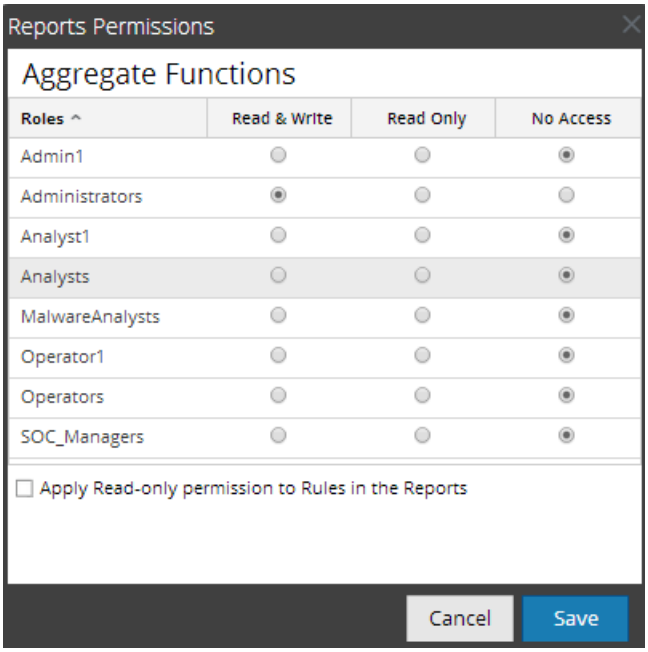
Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports*	Manage Lists, Rules or Reports

*You can complete these tasks here.

Related Topics

- Configure and Generate a Report
- Report View
- Build Report View
- Import Report Dialog

Quick View



To display the Reports Permissions dialog:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, select a report.
4. Click > **Permissions**.
   The Reports Permissions dialog is displayed.

Note: When you select the check box, all dependent rules are given READ access permission, provided the permissions for the report is higher compared to the permissions of the rules.

The following table describes the features in the Reports Permissions dialog.

Feature	Description
Roles	Displays all the roles who can get access to the permissions.

Feature	Description
Read&Write	Allows you to get Read&Write access to the Rules in the Reports.
Read Only	Allows you to get Read Only permissions to the Rules in the Reports.
No Access	If you select this option, you will not get permission to the Rules in the Reports.
Apply Read-only permissions to Rules in the Reports	Allows to set Read Only permissions to the Rules in the Reports for all the roles .
Cancel	This option cancels all the changes made to the permissions.
Save	This option saves the selections and provides access to the roles based on the selections.

Report View

In the Report view, you can create and manage the report or report groups.

Workflow

This workflow shows the procedure to create and schedule a report.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report*	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

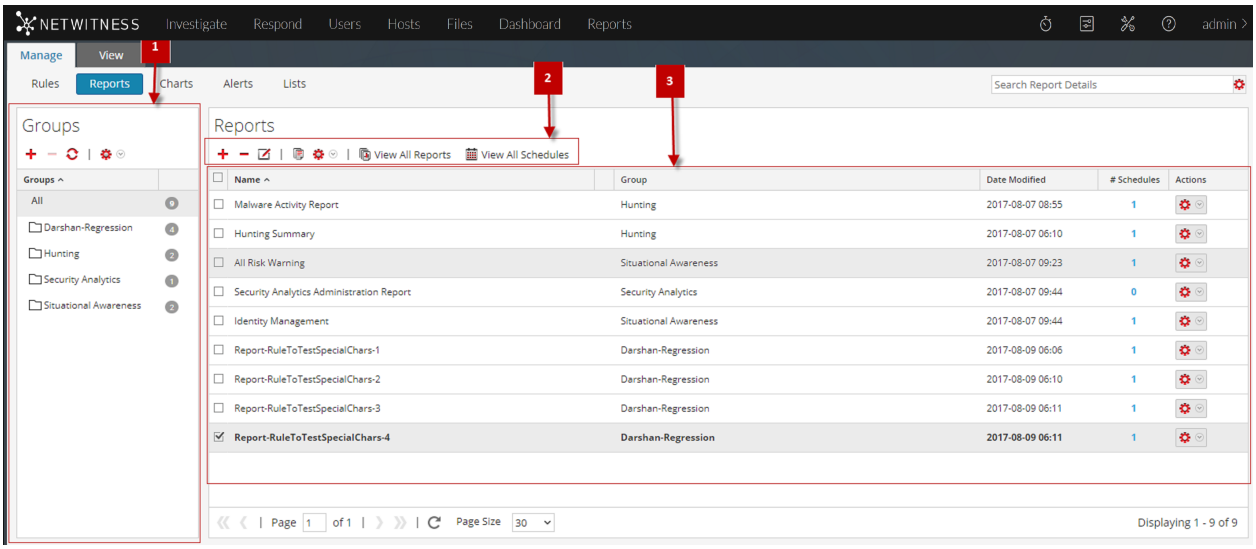
*You can complete these tasks here.

Related Topics

- Build Report View
- Import Report Dialog

- Scheduled Reports View
- Reports Permissions Dialog

Quick View



To access this view:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Reports view is displayed.

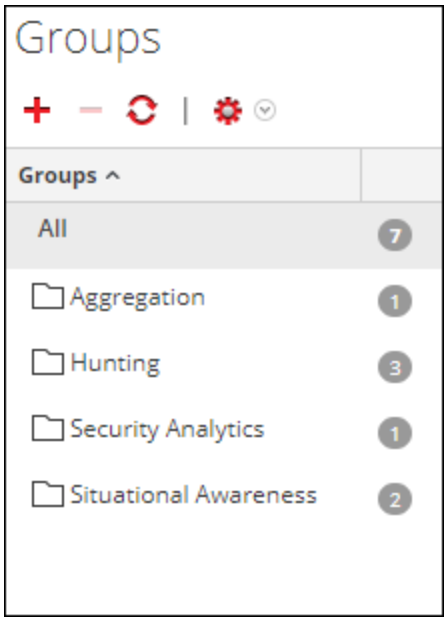
Features

The Report view includes the following sections:

1	Reports Groups panel
2	Reports toolbar
3	Reports panel

Report Groups Panel

The Report Groups panel allows you to organize reports in a group. You can create a report group, add reports to the group, and move reports among groups. You can view all reports by selecting All option under the Groups column.



Feature	Description
+	This option allows you to add a new report to the Reporting module.
−	This option allows you to delete one or more selected report.
↻	This option refreshes the view.
⚙	The actions menu has the following options: Import, Export and Permissions.

Reports Toolbar

The Reports toolbar allows you to add, modify, delete, duplicate, import and export reports. You can also set access permissions for a report in a group.

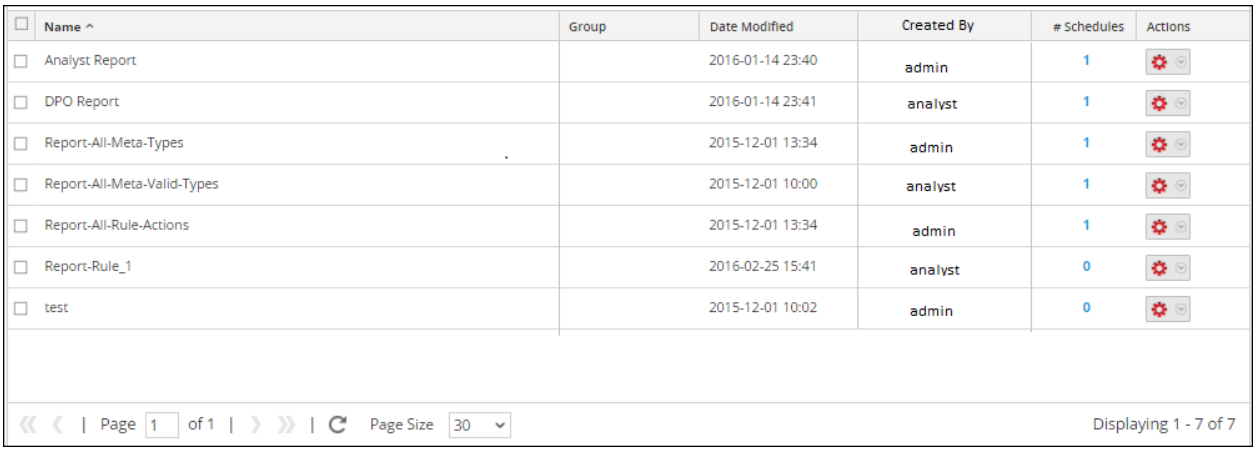


Feature	Description
+	This option allows you to add a new report to the Reporting module.
−	This option allows you to delete one or more selected reports.
✎	This option allows you to edit a chart.
⧉	This option creates a duplicate copy of the selected report.
⚙	The actions menu has the following options: Import, Export , Export as Text and Permissions.

Feature	Description
View All Reports	This option allows you to view a list of reports along with their schedule name and time.
View All	This option allows you to view all the scheduled reports.

Report List Panel

The Report List panel lists all the reports in a tabular format.



The following table describes the columns in the Report List panel.

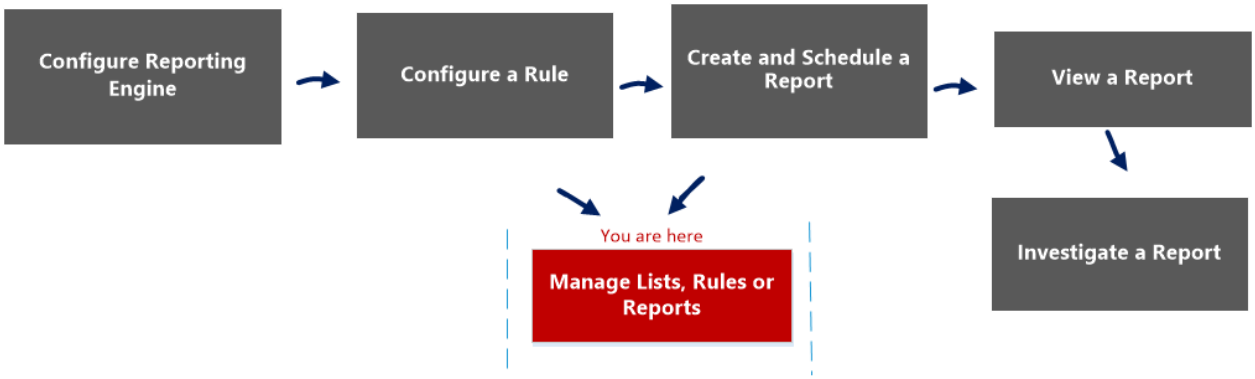
Column	Description
Name	The name of the report.
Group	The Report Group to which the report belongs.
Date Modified	The date and time when the report was modified.
Created By	The creator of the report.
#Schedules	The count indicates the number of schedules created for a report.
Actions	The actions menu has the following options: Schedule Report, View Scheduled Reports, Delete, Edit, and Export.

Rule Permissions Dialog

The Reporting module provides access control at the rule level. Only a user who has the right set of permissions can perform tasks on the rule. When creating user roles, the administrator must ensure that the roles created for specific tasks have access to all the permissions higher in the hierarchy of roles.

Workflow

This workflow shows the procedure to manage rule or rule groups.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports*	Manage Lists, Rules or Reports

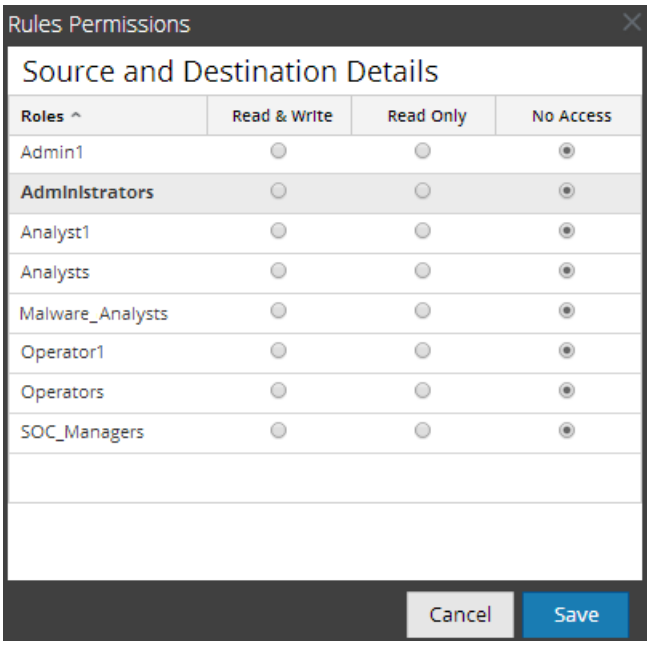
*You can complete these tasks here.

Related Topics

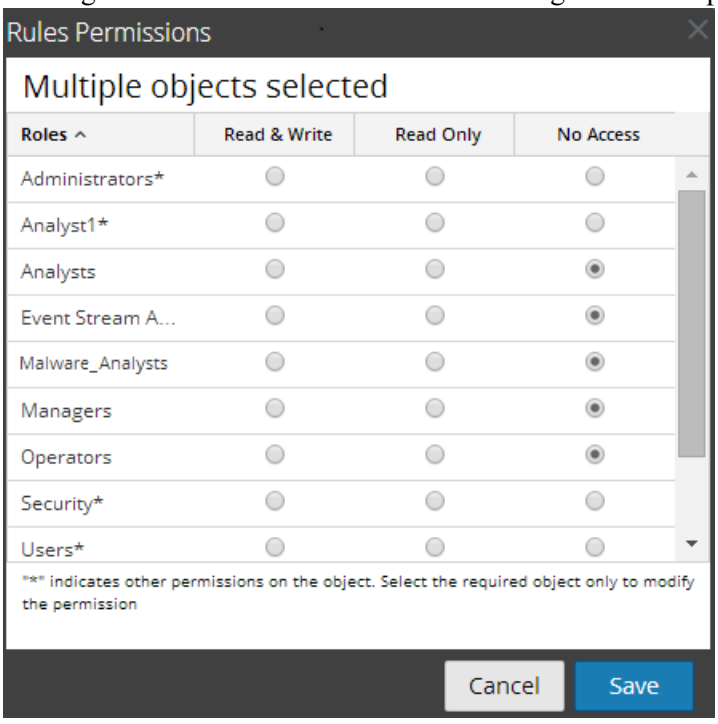
- Rule View

Quick View

This figure shows the Rules Permissions dialog for a single rule.



This figure shows the Rules Permissions dialog when multiple rules are selected.



The dialog has a different appearance for rule groups versus rules. To access the dialog:

1. Go to **Reports**.
   The Manage tab is displayed.
2. In the **Rules** panel, select one or more rules or a rule group.

3. Click  > **Permissions** in the toolbar.

   The Rules Permissions dialog is displayed.

Feature	Description
Roles column	Lists the NetWitness user roles, both built-in and custom roles. Each user who is logged in to NetWitness has user roles assigned. When multiple rules are selected, the asterisk beside the role name, for example, `Security*`, indicates there are other permissions available on that user role. To change the other permissions, you must select the user role and change the access permission.
Read & Write column	When the checkbox in this column is selected, the corresponding user role has permission to view, edit, delete, import, and export rules in the Rules view. The user can also change the permission on the rule.
Read Only column	When the checkbox in this column is selected, the corresponding user role has permission to view the rules in the rule group.
No Access column	When the checkbox in this column is selected, the corresponding user role cannot view or edit the rules in the rule group. Before applying rule permissions, this is the default permission set for all the user roles though the checkbox is unchecked.
Apply these permissions to sub-groups and Rules in this group checkbox	When checked, NetWitness applies permissions to sub-groups and rules in the group.
Cancel option	Clicking Cancel closes the dialog without saving any changes made.
Save option	Clicking Save closes the dialog and updates the rule group permissions for user roles. If specified, the access permissions are applied to subgroups and child objects of this group. When multiple rules are selected, the access permission is applied to all the selected rules.

Rule View

The Rule view is the user interface for managing rules.

Workflow

This workflow shows the procedure to define rule or rule groups.



What do you want to do?

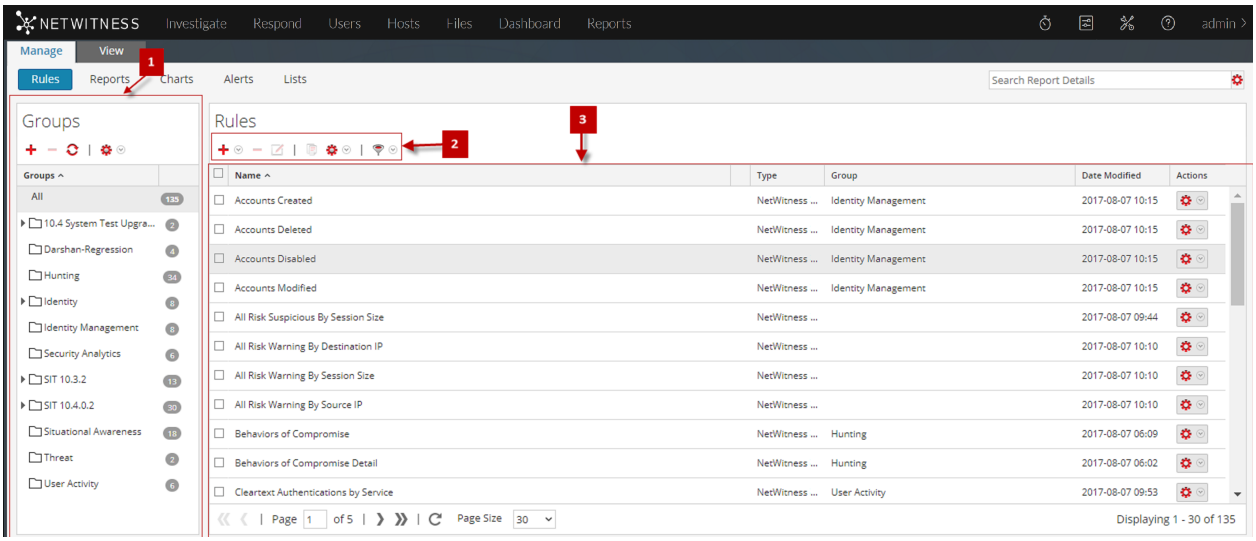
Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule*	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

*You can complete these tasks here.

Related Topics

- Rule Permissions Dialog
- Build Rule View

Quick View



To access the Rules view:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Rules**.

   The Rules view is displayed.

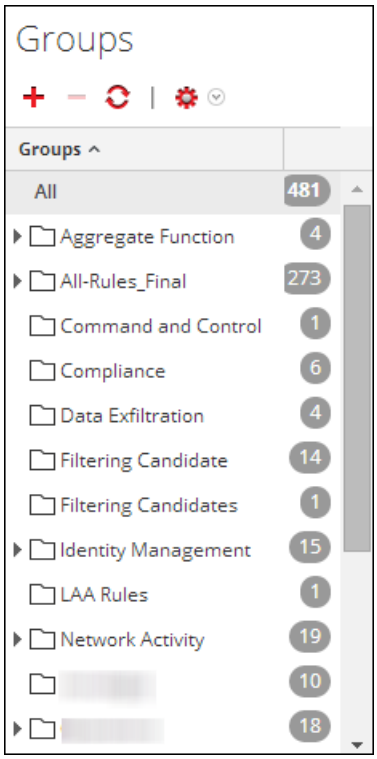
The Rule view includes the following panels.

1	Rules Groups
2	Rules panel
3	Rules Toolbar

Rule Groups Panel

The Rule Groups panel allows you to organize rules into groups using the options in the toolbar. You can create groups and sub-groups and add rules to them. You can also group and move rules between different groups.

The following figure shows the groups in the Rule Groups panel:



The following table describes the features in the Rule Groups Panel.

Feature	Description
+	This option allows you to add a new rule group to the Reporting module.
−	This option allows you to delete one or more rule groups.
↻	This option refreshes the rule group list.
⚙	The actions menu has the following options: Import, Export and Permissions.
All	Displays a list of all the rule groups.

Rule Toolbar

The Rule toolbar allows you to add, delete, edit, and duplicate a rule. The following figure shows the toolbar.



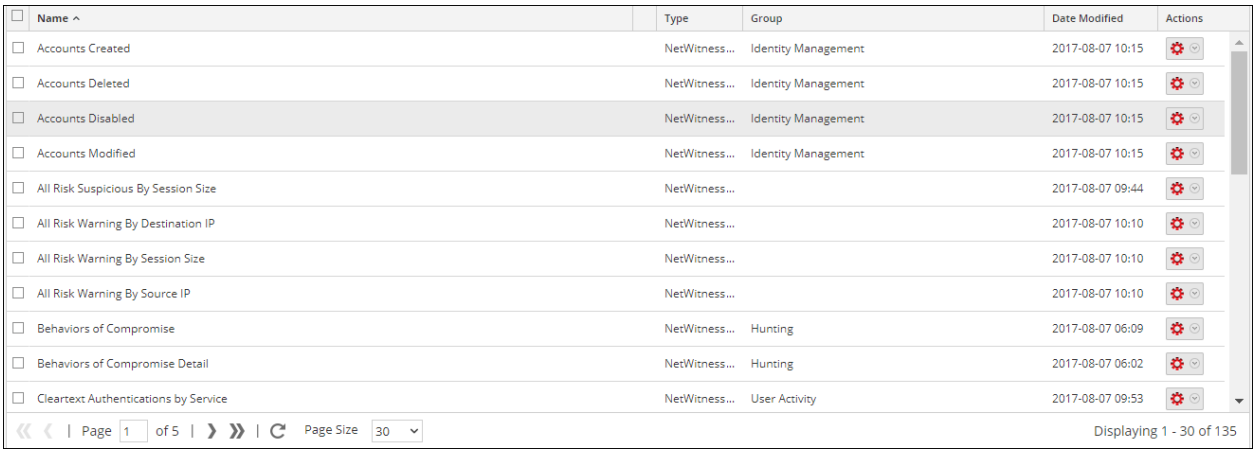
The following table describes the features in the Rule Toolbar

Feature	Description
+	This option allows you to add a new rule to the Reporting module.

Feature	Description
	This option allows you to delete one or more selected rules.
	This option allows you to edit a rule.
	This option allows you to duplicate a rule.
	The actions menu has the following options: Use, Import, Export and Permissions.
	This option allows you to filter the rule type.

Rule List Panel

The following figure shows the list of rules in the Rule List panel.



The following table describes the features in the Rule List Panel.

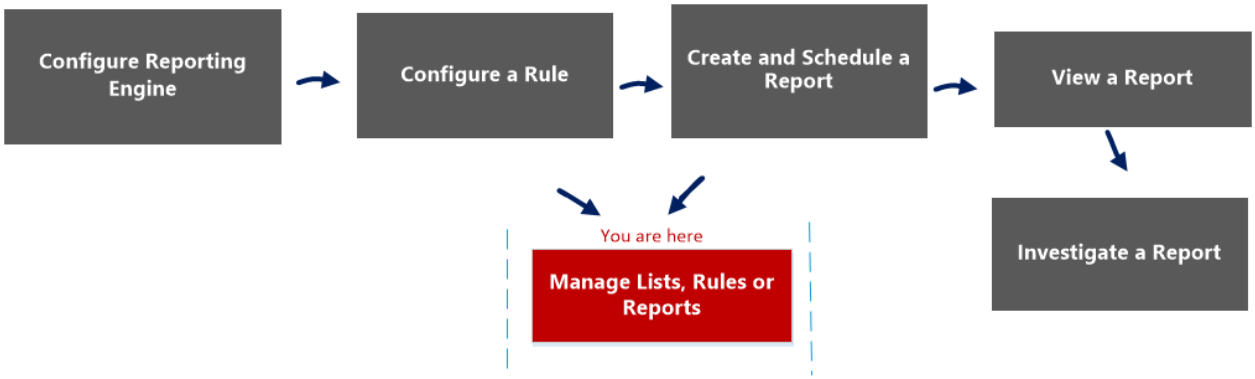
Feature	Description
Name	Displays the name of the rule that you are created or edited. **Note:** For the **Name** field, the icon to extend the column size is not displayed at the end of the column field. You have to hover the mouse a little to the left side to see the icon for extending the column.
Type	Displays the supported database type for the rule you created.
Group	Displays the values which are grouped.
Date Modified	Displays the date when the rule was last modified.
Actions	Displays the actions menu has the following options: Create Alert, Create Chart, Create Report, Delete, Edit, Export, and Dependents.

Select a Logo Dialog

In the Select a Logo dialog, you can upload a new logo that is not available in Reporting Engine Services Config view or choose an existing logo from the Reporting Engine Services Config view.

Workflow

This workflow shows the procedure to manage reports or report groups.



What do you want to do?

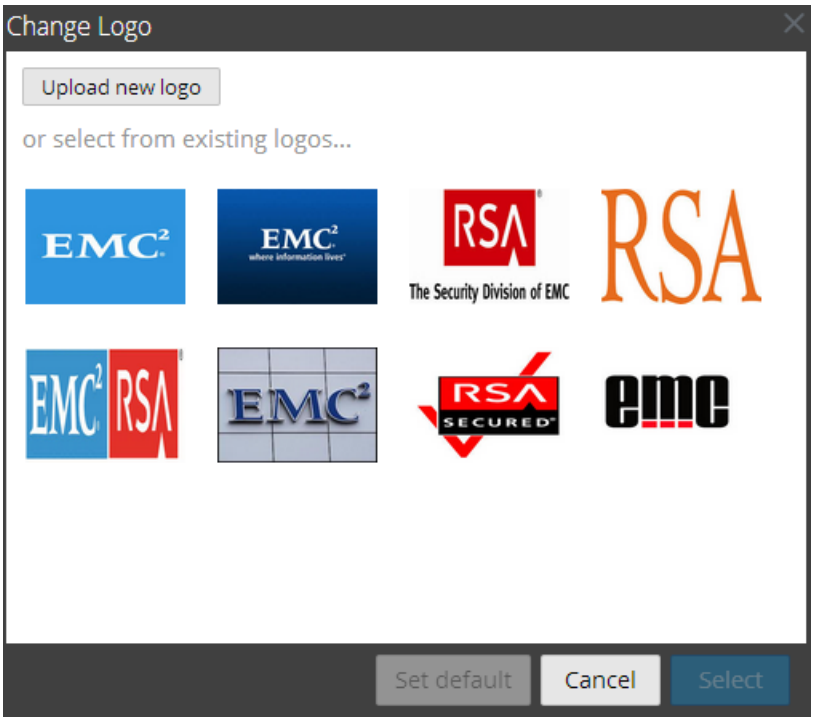
Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports*	Manage Lists, Rules or Reports

*You can complete these tasks here.

Related Topics

- Configure and Generate a Report
- Scheduled Reports View
- Report View

Quick View



To access this dialog:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Reports view is displayed.
3. In the **Reports** panel, select a report.
4. Click > **View Scheduled Reports**.
   The View scheduled reports view tab is displayed.
5. Select a scheduled report and click > **Edit Schedule**.
   The Schedule a Report view tab is displayed.
6. Click the **Logo** panel.
   The Change a Logo dialog box is displayed.

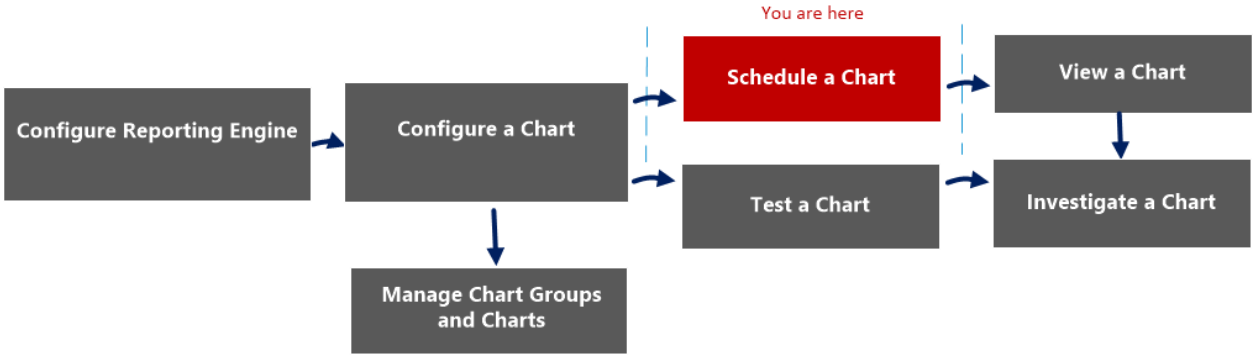
The following table lists the fields in the Select a Logo dialog.

Field	Description
Upload new logo	Click the icon to upload a new logo from the local directory.
Select	Select a logo from the existing list to be used as a logo in the scheduled report.
Cancel	Cancels the logo selection and returns to the Schedule a Report panel.
Set Default	Select a logo to set it as the default logo.

Schedule a Chart View

In the Schedule a Chart View, you can enable or disable a chart.

Workflow



What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart*	Schedule a Chart
Administrator/ Analyst	View a chart	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart	Manage a Chart Group and Chart

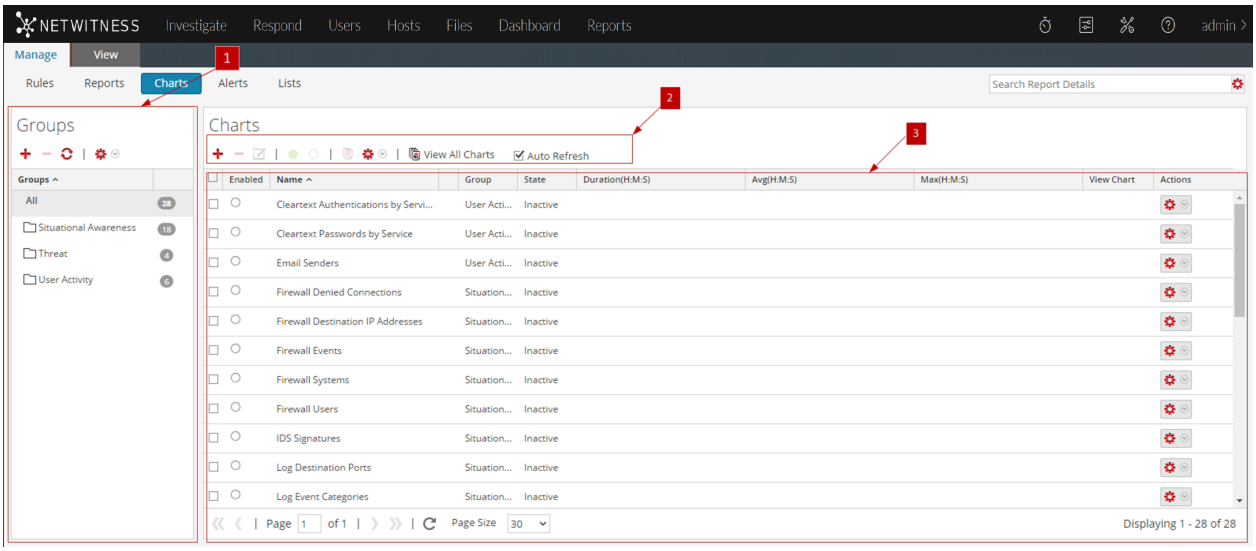
*You can complete these tasks here.

Related Topics

- Configure and Generate a Chart

Quick View

The following figure shows the Schedule a Chart view.



The Schedule a Chart view includes the following panels:

1 Charts Groups panel
2 Charts toolbar
3 Charts panel

Charts Toolbar

The Charts toolbar allows you to add, modify, delete, duplicate, enable, disable, import and export a chart. You can also set access permissions for charts in a group.

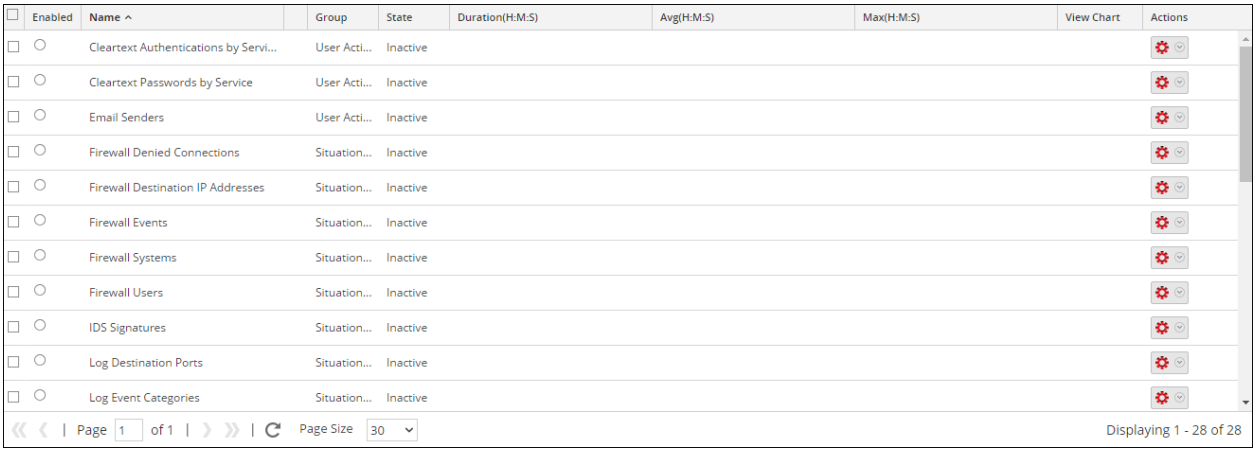


The Charts toolbar includes the following options:

Feature	Description
	Adds a new chart to the Reporting module.
	Deletes one or more selected charts.
	Edit charts.
	Enables the selected charts.
	Disables the selected charts.
	Creates a duplicate copy of the selected chart.
	Provides the following options: Import, Export, Export as Text and Permissions.
View All Charts	Displays all the executed charts.
Auto Refresh	Automatically refreshes the charts list.

Charts Panel

The Charts Panel presents all the charts in a tabular or grid format.



The following table lists the columns in the Charts panel and their description.

Feature	Description
Enabled	- The chart is enabled. - The chart is disabled.
Name	The name of the chart.
Group	The Chart Group to which the chart belongs.
State	The state of the chart: • Queued • Completed • Failed
Duration (H:M:S)	The time taken to execute the latest chart.
Avg(H:M:S)	The average time taken to run the chart.
Max(H:M:S)	The maximum time taken to run the chart.
View Chart	A hyperlink that redirects to the View a Chart panel.
	The actions menu has the following options: Enable, Disable, View, Delete, Edit, and Export.

Schedule Report Panel

The Schedule Report panel allows you to schedule a customized report. Prior to scheduling a report, you can create a dynamic list (with the overwrite option selected) with services added. For more information, see "Generate a List from the Scheduled Report" section in Create and Schedule a Report. Then use the list to generate a report with details in the report like services and host names.

Workflow

This workflow shows the procedure to create and schedule a report.



What do you want to do?

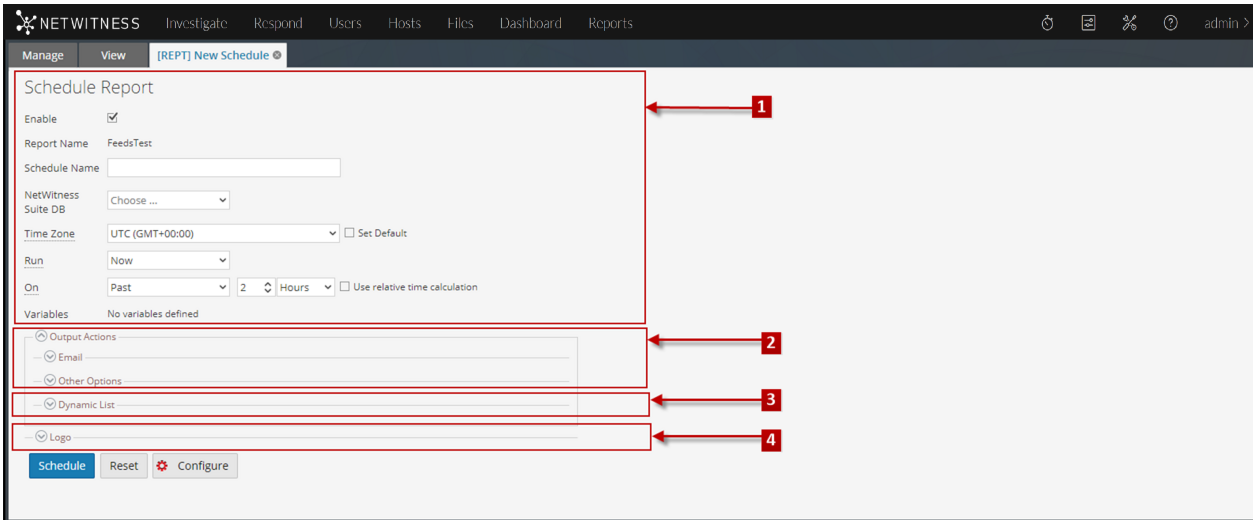
Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report*	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

*You can complete these tasks here.

Related Topics

- Configure and Generate a Report
- Report View
- Build Report View
- Scheduled Reports View

Quick View



To access this view:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Reports view is displayed.
3. In the **Reports** panel, click  > **Schedule Report**.

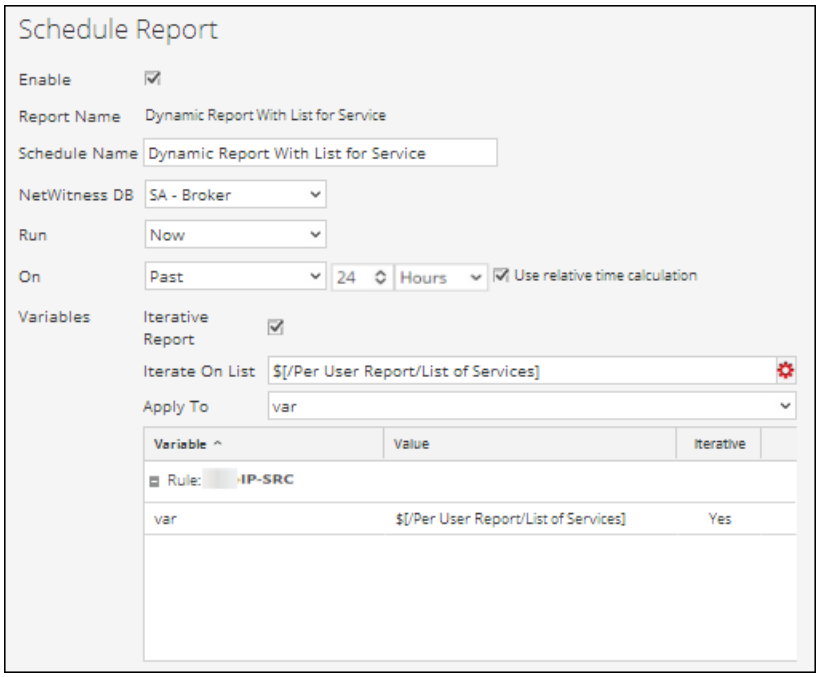
Features

The Schedule Report view consists of the following panels:

1	Schedule Report View
2	Output Actions Panel
3	Dynamic List Panel
4	Logo Panel

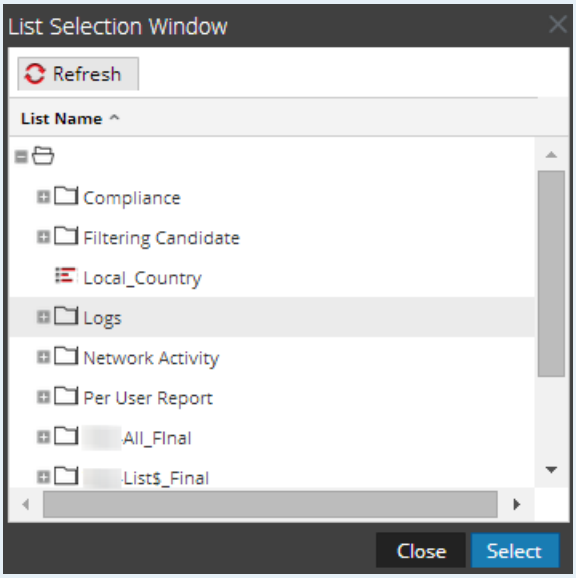
Schedule Report View

The Schedule Report view allows you to schedule reports.



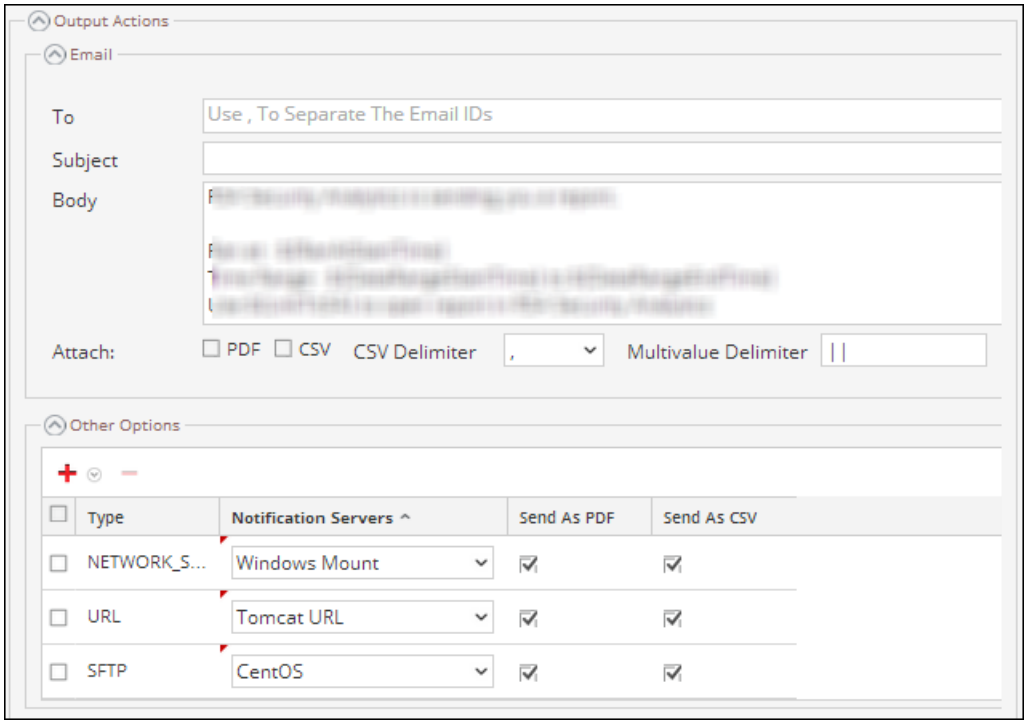
The following table lists the fields in the Schedule Report panel.

Field	Description
Enable	Enables the report schedules and runs the report.
Report Name	The name of the report.
Schedule Name	The name of the scheduled report configuration.
NetWitness DB	The database can be NWDB and Warehouse DB depending on the type of database that you selected in the rule definition. If the report has rules of NWDB and Warehouse DB types, all the database types or rule types are displayed.
Warehouse Resource Pool	If the report has rules of Warehouse DB, the Warehouse Resource Pool drop-down is displayed to select the pools or queues available in the cluster. If no pools or queues are entered for the Reporting engine, this field is disabled. For more information, see "Step 5: Configure Task Scheduler for a Reporting Engine" topic in the *Host and Services Configuration Guide*.
Run	Provides the type of schedule for the run configuration: • Ad-hoc execution • Hourly execution • Daily execution • Weekly execution • Monthly execution
On	The data range on which the query is run.

Field	Description
Use relative time calculation	Uses the relative time duration to schedule a report.
Iterative Report	Select the checkbox to schedule a report for the selected list value.
Iterate on List	Click this button to navigate to the List Selection panel and select a list. The following figure displays this panel:  The List Selection panel is a collection of Lists. The Reporting Engine maintains an active list of the available list names by continuously synchronizing with the collection to which it is connected.
Apply To	Apply list values on the selected variable.
Variables	Displays the rule variables along with their associated values and the iterative properties included in the report. **Note:** Depending on the rule chosen while creating a report, you can view dynamic variables defined for the rule in the **Variables** field of the Schedule Report panel. For example, Test-Country is the rule having the dynamic variable var.
Schedule	Schedules the report.
Reset	Resets the scheduled report.
Configure	Allows you to alter the Reporting Engine configuration details on the "Reporting Engine General Tab" topic in the *Host and Services Configuration Guide*. **Note:** This button is visible on the Schedule Report panel only when you have the 'Manage Device' access permissions on the Reporting module.

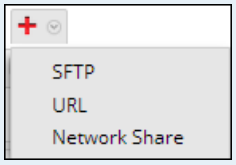
Output Actions Panel

The Output Actions panel specifies output actions to notify the email recipient when the report execution completes and also sends reports in PDF and CSV formats as attachments in the email, based on your selection.



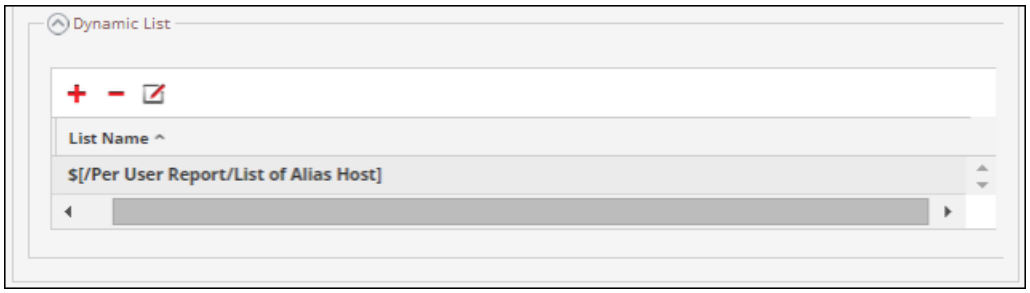
The following table lists the fields in the Output Actions panel.

Field	Description
To	A comma-separated list of email addresses to receive the output.
Subject	The subject entered in the mail.
Body	The body of the email. By default, the body field is populated with pre-defined text that has certain variables that will add meta appropriate to the generated report. In the Reporting Engine, these variables are replaced with actual values. • `${RanAtStartTime}` : The Start time of the report. • `${DataRangeStartTime}` : The Start time of the data time range. • `${DataRangeEndTime}` : The End time of the data time range. • `${LinkToSA}` : The link to the NetWitnessHost from the email which in turn opens the report in NetWitness interface. • `${ReportName}` : The name of the report. • `${DataSource}` : The name of the data source.

Field	Description
Attach:	The output format in which the report is attached to the email, such as PDF or CSV as configured in the Schedule Report dialog.
CSV Delimiter	The default CSV delimiter is comma (,). If the CSV content contains a comma, you must identify a unique separator so the content is stored in its original form. For example, if msg is a column in the report to be saved as CSV and the msg content is as follows: `ASA-SSM-CSC-20 Module in slot 1," application reloading ""CSC SSM""," version ""6.2.1599.0"" CSC SSM scan services are reloading because of a pattern file or configuration update` The above content will be included in three columns due to the commas (,). To avoid this, you must specify a different delimiter such as a pipe line character "|". **Note:** To import the CSV file into Microsoft Excel, use the Data > From Text option in the Excel application. When you import the CSV file you must specify the file type of the file being imported as Delimited and use the same delimiter that you specify to generate the CSV file.
Multivalue Delimiter	The data in multivalued fields are separated by the multivalue delimiter. The default Multivalue delimiter is two pipe line characters (||).
Other Options	You can select an SFTP, URL, or Network Share location configured in ((RE}} and then send the report either in PDF or CSV format based on the requirement.
	Select this option to send the report to the SFTP, URL or Network Share location configured in the Reporting Engine Services Config view.
Type	The type of output action chosen. For example, SFTP, URL or Network Share.
Output Actions	Select the SFTP, URL or Network Share name configured in the Reporting Engine Services Config view.
Send as PDF / Send as CSV	Select these options to send the report either in PDF or CSV format, or both to the configured Notification Server (SFTP, URL or Network Share).

Dynamic List Panel

The Dynamic List panel populates the lists created and you can add, edit or delete the list. The list is generated based on the scheduled report which can be viewed in the Lists view.



The following table lists the operations in the Generate List panel.

Operation	Description
	Adds a new list to the report.
	Deletes all the lists added to the report.
	Displays the Generate List dialog.
List Name	The name of the list chosen from the List Selection panel. For more information on the List Selection panel topic, see Generate List Panel.

Logo Panel

The Logo panel populates the default logo from the Select a Logo panel. For more information on choosing a logo from this panel, see "Manage and Select a Report Logo" section in the Manage Lists, Rules or Reports.

You can set the default logo for a Reporting Engine. This is the logo that is used in the generated reports. For more information on choosing a logo, see Select a Logo Dialog.

Note: If you have not selected any logo then the default RSA logo is used on the report. The option **Save as PDF** for the previously executed reports does not support a new customer logo. It displays the default RSA Logo, if the customer logo must be displayed in the Schedule a Report view.



Scheduled Reports View

The Scheduled Reports view allows you to create, view and manage scheduled reports.

Workflow

This workflow shows the procedure to create and schedule a report.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Configure the Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report*	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports*	Manage Lists, Rules or Reports

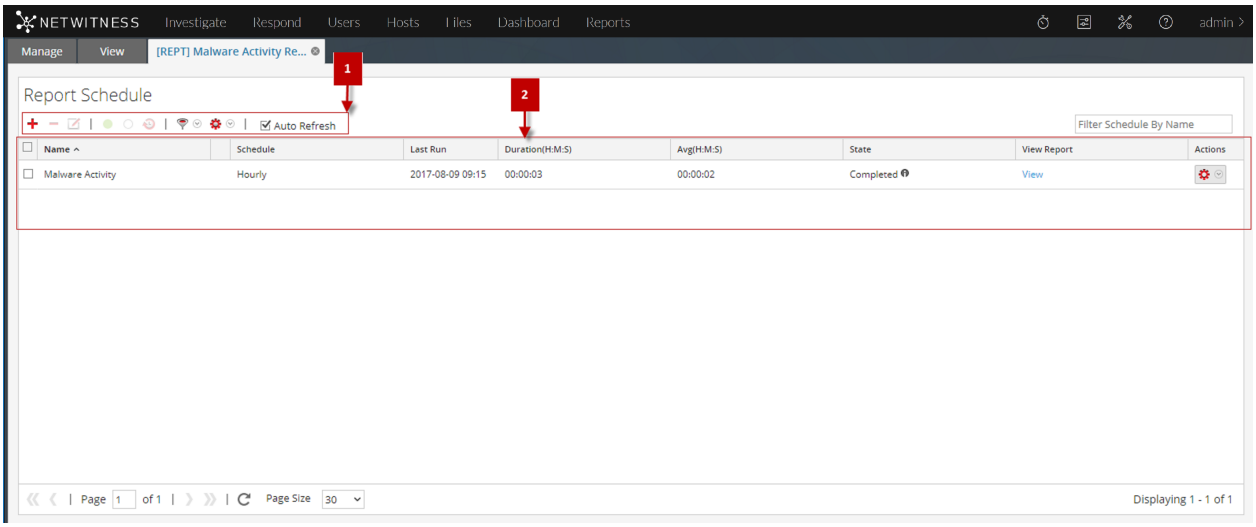
*You can complete these tasks here.

Related Topics

- Build Report View
- Report View

- Schedule Report Panel
- Reports Permissions Dialog

Quick View



To access this view:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** panel, do one of the following:

   - Click  > **View Scheduled Reports**.
   - Click the **Schedules** column.

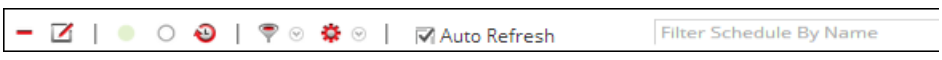
Features

The View Scheduled Reports has the following features:

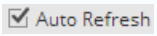
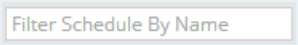
1	Report Schedule Toolbar
2	Report Schedule List panel

Report Schedule Toolbar

The Scheduled Reports has options to add, modify and delete the scheduled report as well as options to enable or disable the selected run configuration.



The following table lists the operations in the Scheduled Reports toolbar.

Operation	Description
	Create a new report schedule.
	Delete the selected report schedule.
	Edit the selected report schedule. **Note:** Double-click on a desired report schedule to edit it.
	Enables the selected report schedule.
	Disables the selected report schedule.
	View the history of the scheduled report.
	Filter schedules based on the type of schedule. (For example, AdHoc)
	Allows you to set permissions for the selected scheduled report.
	Automatically refreshes the scheduled reports list.
	Searches schedules based on the schedule name.

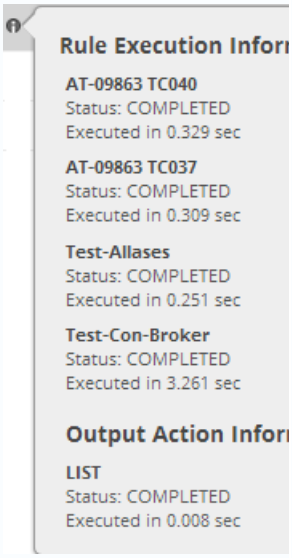
Report Schedule List Panel

The Scheduled Reports List panel lists the scheduled reports in a tabular format.

The following table lists the columns in the Scheduled Reports List panel:

Column	Description
Name	The name of the scheduled report.
Schedule	The type of schedule for the run configuration: • Ad-hoc execution • Hourly execution • Daily execution • Weekly execution • Monthly execution
Last Run	Displays the last time the report was run.
Duration(H:M:S)	Displays the time taken for last execution of the report.

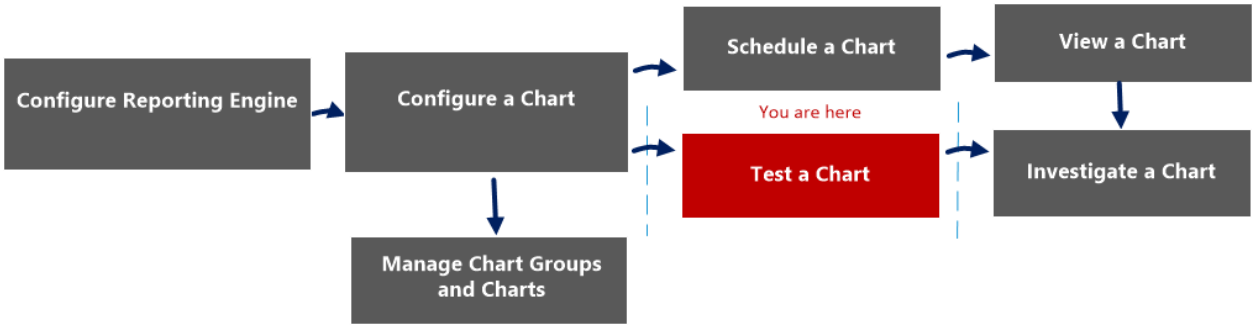
Column	Description
Avg (H:M:S)	Displays the average time taken to run the report.
State	Indicates the state of the scheduled report. • Scheduled: If a report is scheduled to run on an hourly, daily, weekly, monthly, or later time, the state of the report is displayed as scheduled, for the first run. • Queued: If a report is still waiting to get executed, the state of the report is displayed as queued. • Running: If the report schedule is in progress, the state of the report is displayed as running. • Partial: If in a report with several rules, a single rule execution failed or an output action failed or creation of PDF/CSV failed, the state of the report is displayed as partial. For example, consider a report with five rules and four rules are executed successfully and one fails, then the state is displayed as Partial. • Failed: If in a report with several rules, all the rule schedule executions failed, the state of the report is displayed as failed. • Completed: If a report schedule is successfully executed, the state of the report is displayed as completed. • Canceled: When cancel request is completed, the state of the report is displayed as canceled. • Inactive: If a report schedule is disabled, the state of the report is displayed as Inactive. • Not available: If the report schedule executed information is not available, the state of the report is displayed as not available.

Column	Description
	Click to view the rule execution information and output action information. This pop-up notifies the status of multiple rules in a report and the time taken for its execution. **Note:** You can view the rule execution and output action information for a scheduled report having the state **Completed**, **Running**, **Partial** or **Failed**. By default, the Output Actions for Completed Report on Reporting Engine Config page is set to enable, to receive an email when the report status is completed. To receive an email for **Failed** or **Partial** reports, you must disable this option.
View Report	Click to view the rule execution information on the View a Report Panel. You can view the rule execution information for a scheduled report having the state 'running' as well.

Test a Chart View

In the Test a Chart view, you can view and test the charts.

Workflow



What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart	Schedule a Chart
Administrator/ Analyst	View a chart	View a Chart
Administrator/ Analyst	Test a chart*	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart	Manage a Chart Group and Chart

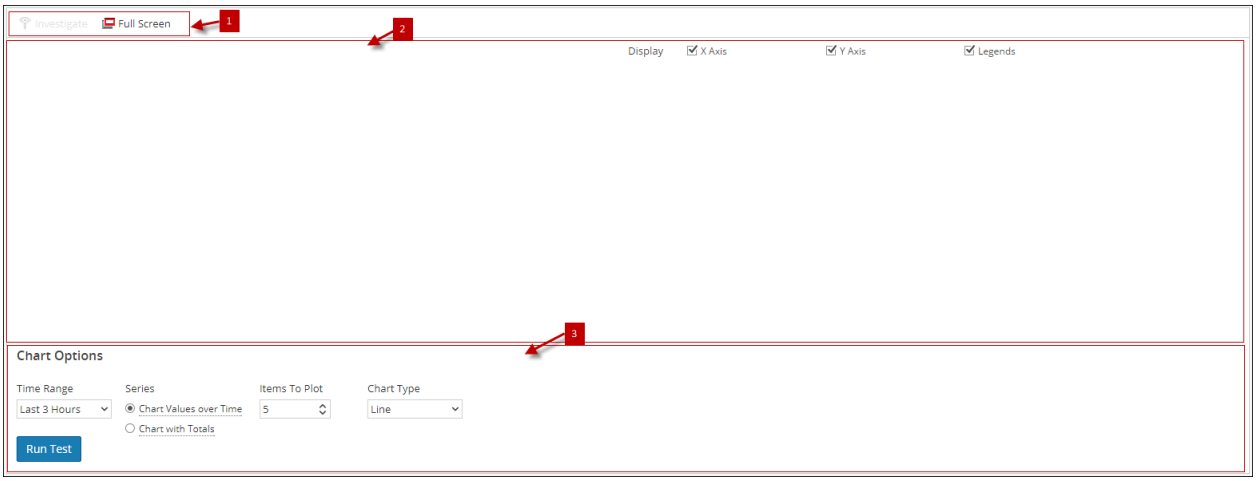
*You can complete these tasks here.

Related Topics

- Configure and Generate a Chart

Quick View

The following figure is an example with the important features labeled.



The Test a Chart view consists of the following panels:

1 Chart toolbar
2 Chart Output panel
3 Chart Options panel

Chart Toolbar

The Charts toolbar allows you to investigate on a particular chart and change the screen to full screen.



Feature	Description
Investigate	Investigates further on the selected chart.
Full Screen	Displays the chart in full screen.

Chart Output Panel

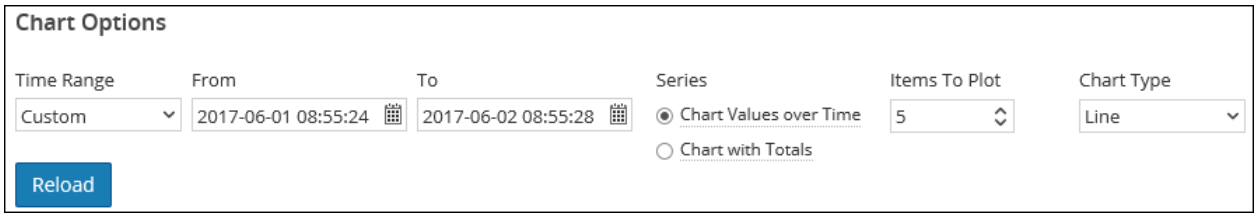
The Chart Output panel displays the information in a chart format for the selected time chart options.

The following table lists the features in the Test a Chart View and their descriptions.

Feature	Description
Display	Allows you select the values that needs to be displayed and have the following options: X Axis, Y Axis and Legends.
X Axis	Displays the session count.
Y Axis	Displays the actual output.
Legends	Displays the list of variables appearing in the chart.

Chart Options Panel

The following figure shows the Chart Options panel, which displays the time range, series, and chart type fields to configure the chart display.



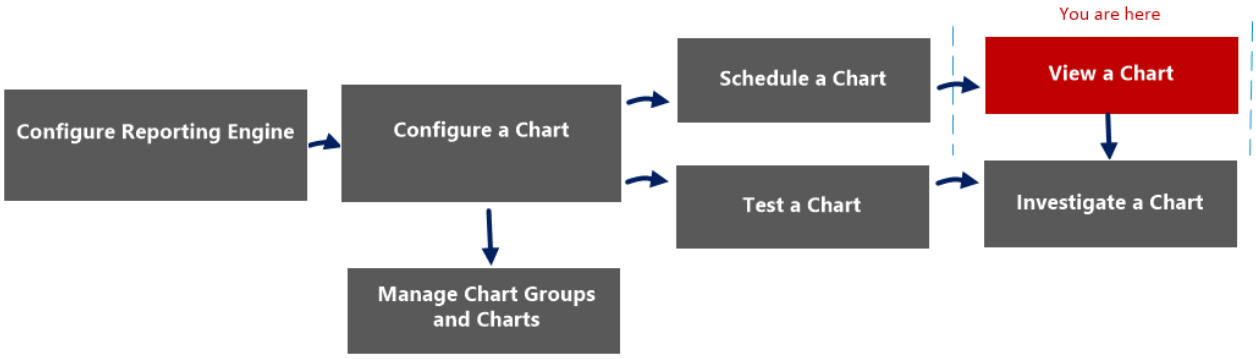
The following table lists the fields in the Charts Options panel and the descriptions.

Feature	Description
Time Range	The default time range is Last 3 Hours. However, you can select a different value from the drop-down list, for example, Last Hour, or Last 6 Hours which are the preset values. Or you can customize by selecting Last N Days or the Custom option.
From	The start date and time. (only for custom options).
To	The end date and time. (only for custom options).
Series	The series field provides you with two options: • Chart Values over Time: Renders the chart for the entire time range selected. • Chart with Totals: Renders the summary of data for the selected date range.
Items to Plot	The maximum number of events the user wants to view on the chart.
Chart Type	The type of chart to be rendered either area, bar, column, line, step line, step area, spline area or spline.

View a Chart Panel

In the View a Chart panel, you can view and manage charts. There are options for filtering and sorting the information in the chart, as well as options for the type of chart, the number of items to chart, and charting values or totals. When viewing a chart, you can open the charted sessions in the Investigation module and save the chart as a PDF.

Workflow



What do you want to do?

Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart	Schedule a Chart
Administrator/ Analyst	View a chart*	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart	Manage a Chart Group and Chart

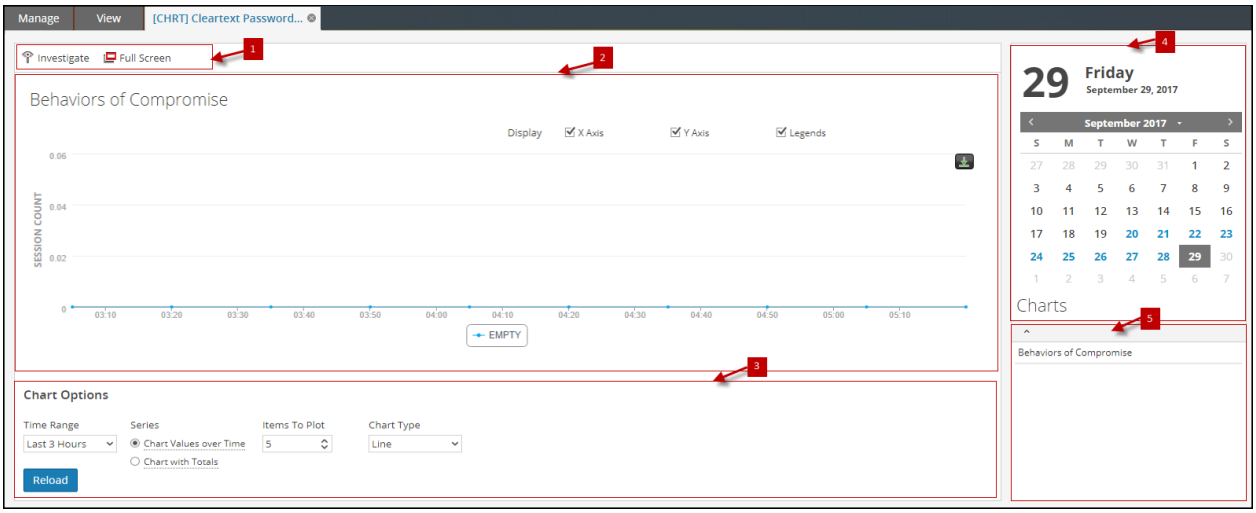
*You can complete these tasks here.

Related Topics

- Configure and Generate a Chart

Quick View

The following figure is an example with the important features labeled.



The View a Chart panel includes the following panels:

1. Chart toolbar
2. Chart Output panel
3. Chart Calendar panel
4. Chart Options panel
5. Chart Executed list

Chart Toolbar

The Chart toolbar has options that allow you to investigate, and view the chart on another screen.



The following table lists the options in the Chart toolbar.

Operation	Description
Investigate	Investigates the chart details.
Full Screen	Displays the chart on a full screen.

Chart Output Panel

The Chart Output panel displays the chart with sortBy on the Y-axis, time on the X-axis and legends.

Note: You can save the chart as PDF using the icon on the Chart Output panel.

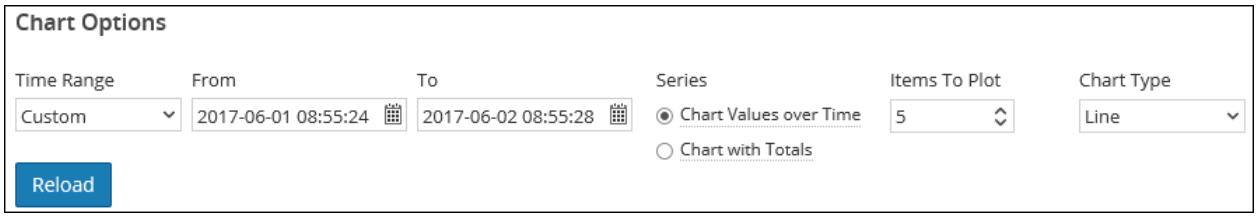
Chart Calendar Panel

The Chart Calendar panel is the default calendar with which you can filter the list of charts depending on the date you select from the Calendar, as shown in the following figure.



Chart Options Panel

The Chart Options panel displays the time range, series, and chart type fields to configure the chart is displayed.



The following table lists the fields in the Chart Options panel.

Field	Description
Time Range	The default time range is Last 3 Hours. However, you can select a different value from the drop-down list, for example, Last Hour, or Last 6 Hours which are the preset values. Or you can customize by selecting Last N Days or the Custom option. **Note:** The time range selected by you for a chart will be saved. When you open the same chart the next time, the time range that is saved will be displayed. This behavior is not applicable for the custom option.
From	The start date and time. (only for custom options)
To	The end date and time. (only for custom options)
Series	The series field provides the user with two options: • Chart Values over Time: Renders the chart for the entire time range selected. • Chart with Totals: Renders the summary of data for the selected date range.

Field	Description
Items to Plot	The maximum number of events the user wants to view on the chart.
Chart Type	The type of chart to be rendered. Either area, bar, column, line, step line, step area, spline area or spline.

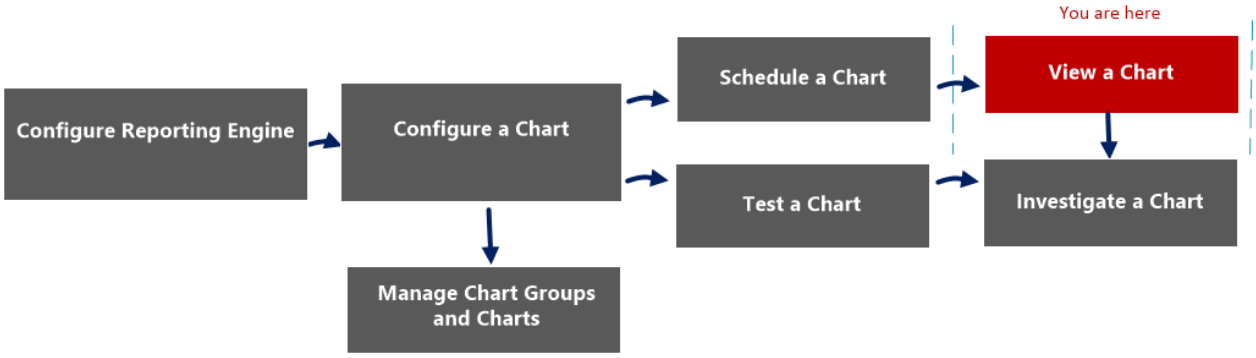
Chart Executed List Panel

The Chart Executed List panel displays all the executions for a particular chart for the selected date. Double-clicking on any chart execution loads the chart on the Chart Output panel. By default, the last executed chart is displayed in the Chart Output panel.

View All Charts View

In the View All Charts view, you can display, print, save and email charts.

Workflow



What do you want to do?

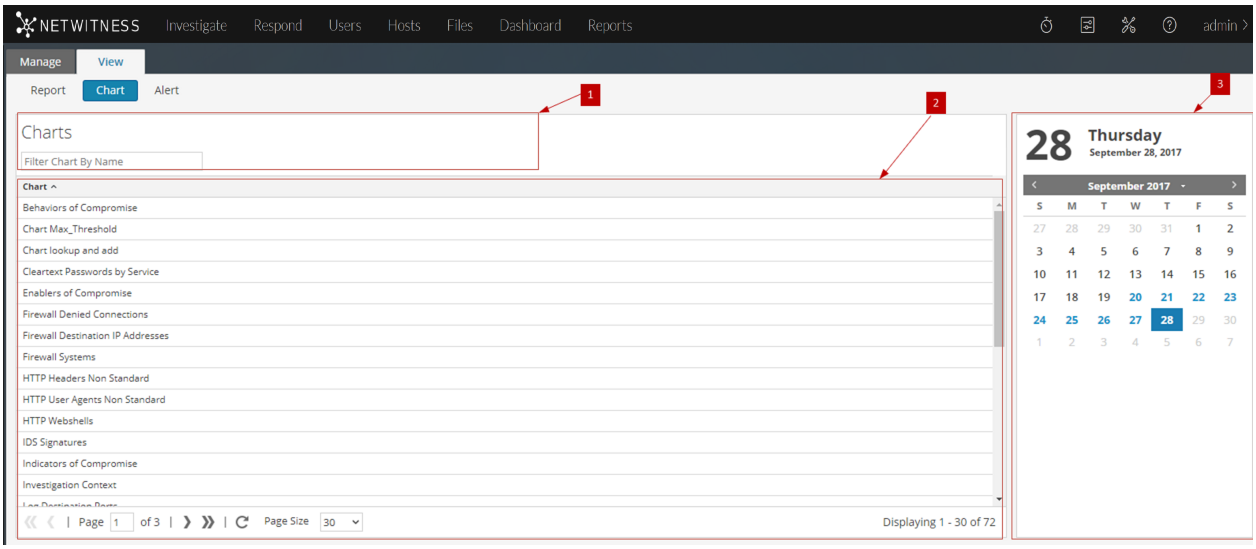
Role	I want to ...	Documentation
Administrator/ Analyst	Configure Reporting Engine	For more information, see "Configure Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator/ Analyst	Configure a chart	Configure a Chart
Administrator/ Analyst	Schedule a chart	Schedule a Chart
Administrator/ Analyst	View a chart*	View a Chart
Administrator/ Analyst	Test a chart	Test a Chart
Administrator/ Analyst	Investigate a chart	Investigate a Chart
Administrator/ Analyst	Manage a chart group and chart	Manage a Chart Group and Chart

*You can complete these tasks here.

Related Topics

- Configure and Generate a Chart

Quick View



The View All Charts panel includes the following panels.

1 Charts Toolbar
2 Charts Output panel
3 Charts Calendar panel

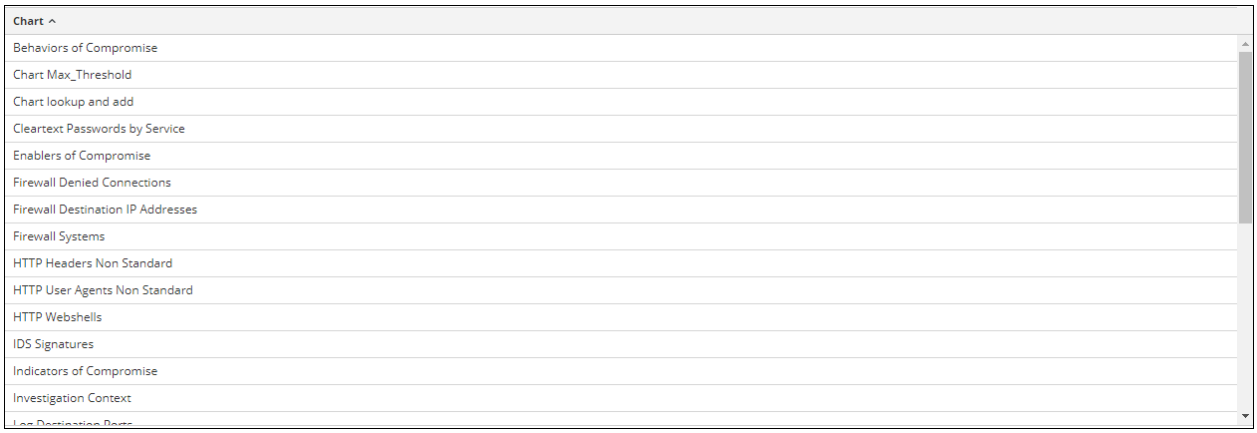
Charts Toolbar

The following table lists the options in the View All Charts toolbar:

Operation	Description
Filter Chart By Name	Searches schedules based on the chart name for a selected calendar day.

Charts Output Panel

The Charts Output panel displays the chart with the chart schedule name.



Feature	Description
Chart	This field displays all the successfully executed charts.

Charts Calendar Panel

The Charts Calendar panel is used to select a date from the Calendar. Based on the date you select, the list of successfully run charts for the date is displayed.


28
Thursday
September 28, 2017
September 2017
S M T W T F S
27 28 29 30 31 1 2
3 4 5 6 7 8 9
10 11 12 13 14 15 16
17 18 19 20 21 22 23
24 25 26 27 28 29 30
1 2 3 4 5 6 7

View a Report Panel

The View a Report panel is used to review the reports.

Workflow

This workflow shows the procedure view a report or list of all reports.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports*	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

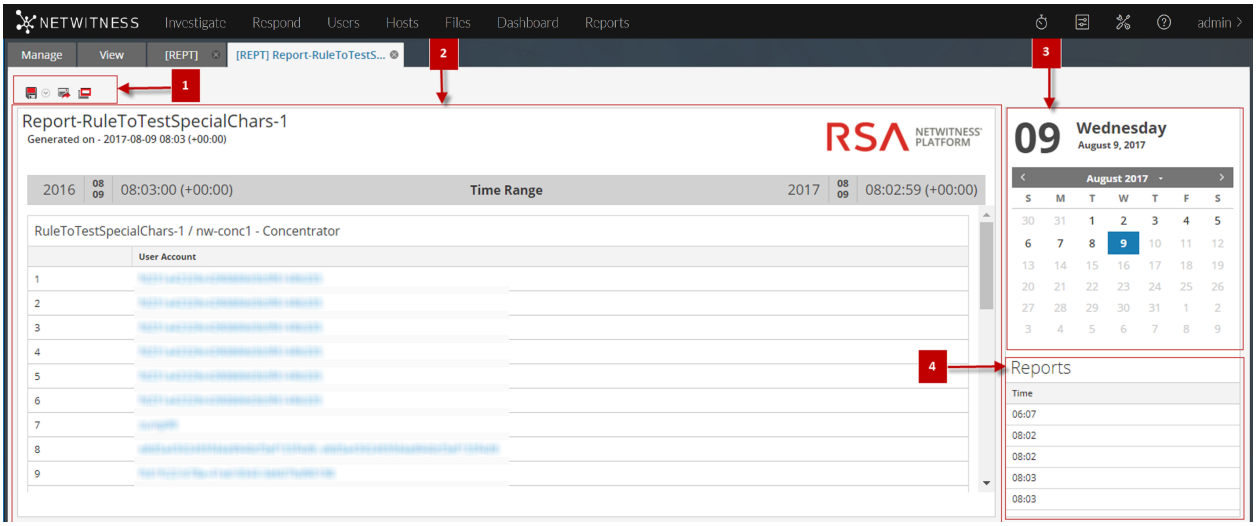
*You can complete these tasks here.

Related Topics

- Configure and Generate a Report
- Build Report View

- Import Report Dialog
- Scheduled Reports View
- Reports Permissions Dialog
- View All Reports View
- Report View

Quick View



To access this view:

1. Go to **Reports**.

   The Manage tab is displayed.

2. Click **Reports**.

   The Report view is displayed.

3. In the **Reports** panel, do one of the following:

   - Click  >**View Scheduled Reports**.
   - Click the **#Schedules** column.

     The Report Schedule view is displayed.

4. Click **View**.

Features

The View a Report panel has the following sections.

1	Reports Toolbar
2	Reports Output panel
3	Reports Calendar panel
4	Reports Time panel

Reports Toolbar

The Reports toolbar allows you to print, save, email, and view reports on full screen.

Note: The Reporting Engine is responsible for generating PDF and CSV output of the reports based on the report definition. The size of the PDF files for a report must not exceed 50,000 cells.

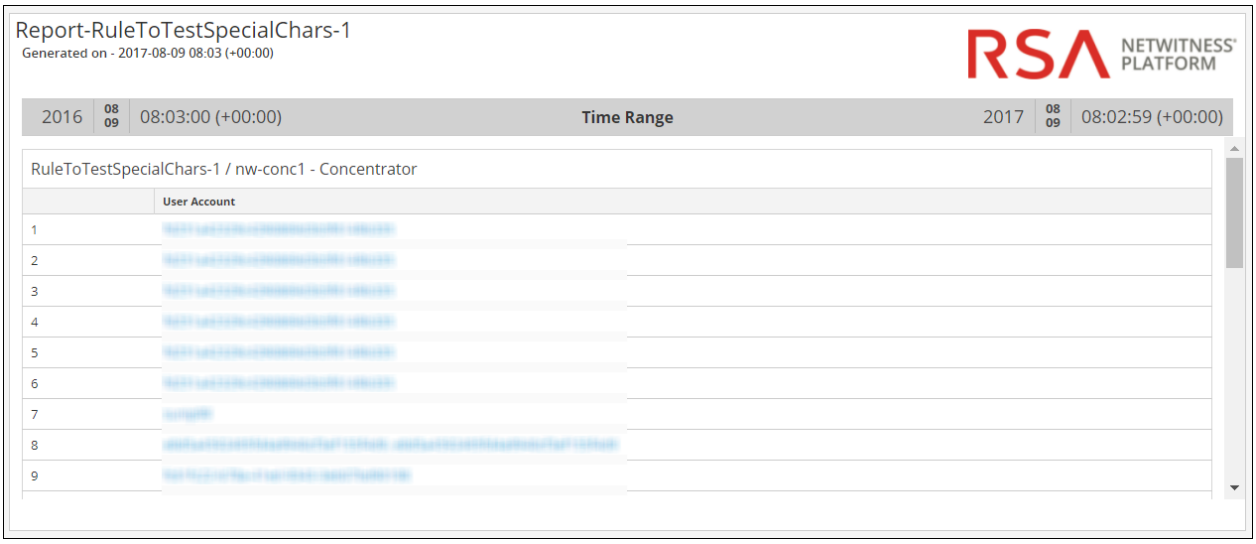


The following table lists the options in the Reports toolbar.

Operation	Description
	Prints the generated report.
	Saves the report as a PDF and a CSV file. **Note:** The **Save As PDF** option is not available for a large report. If you are generating a PDF for a report and it takes a longer time than expected, you get a warning message stating **PDF generation is in progress, please try after some time**. When you click download as a CSV file, the Select Rule to download dialog is displayed. You must select a rule from this dialog to download the rule result in a CSV file. If the file generation takes a while, you can click on the **Notify me** option to be notified once the PDF or CSV is generated. Once the PDF or CSV is generated, you can view the Notifications for the status.
	Emails the report with the PDF or CSV attachment.
	Opens the generated report on a new window.

Reports Output View

The Reports Output panel view the report with the report schedule name, report generated time and the actual report with the selected rule variables.



Feature	Description
Name	This field displays the name of the scheduled report.
Time	This field displays the time when the report is generated.
Report	This field displays the details report with the selected rule variables.

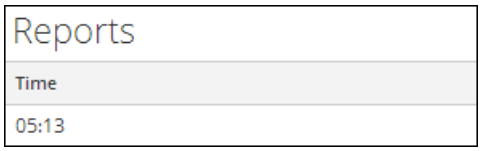
Reports Calendar View

The Reports Calendar view is used to select a date from the Calendar. Based on the date you select, the list of successfully run reports for the date is displayed.

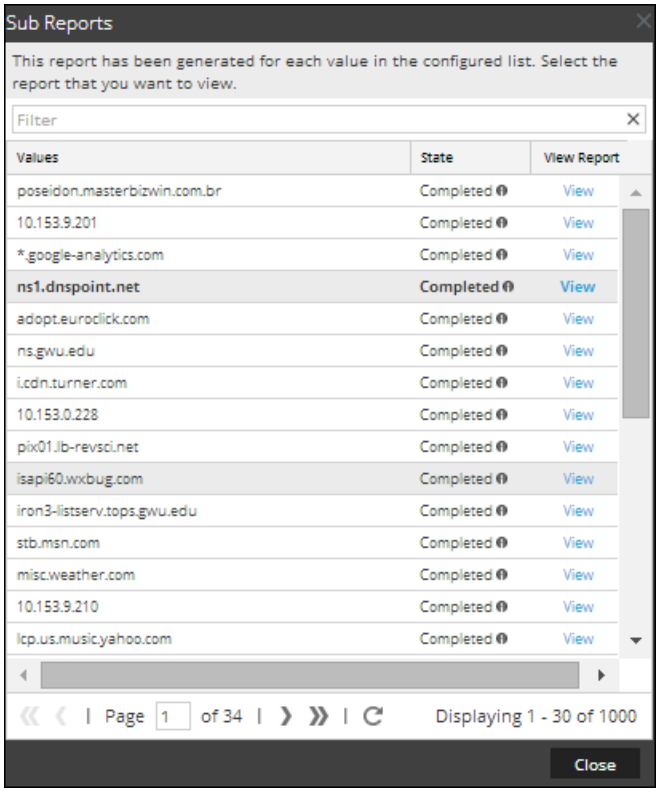


Reports Time View

The Reports Time view displays the time when the report was actually run.



When you click **View** on the scheduled report having **Iterative** selected, the **Sub Reports** panel is displayed. For each value in the configured list a report is generated.



The following table lists the columns in the Sub Reports panel.

Column	Description
Values	The List values chosen for a dynamic variable from the List Selection panel.
State	Indicates the state of the scheduled report for each of the list values. • Partial: If in a report with several rules, a single rule execution failed or an output action failed or creation of PDF/CSV failed, the state of the report is displayed as partial. For example, consider a report with five rules and four rules are executed successfully and one fails, then the state is displayed as Partial. • Failed: If in a report with several rules, all the rule executions failed, the state of the report is displayed as failed. • Completed: If a report is successfully executed, the state of the report is displayed as completed.

Column	Description
View	Clicking on any of the report schedules or sub reports listed and then **View** displays the desired report. **Note:** You can view the completed rules on the **View a Report** page even when the report is 'running'.

View All Reports View

In the View All Reports view, you can display, print, save and email reports.

Workflow

This workflow shows the procedure view a report or list of all reports.



What do you want to do?

Role	I want to ...	Show me how
Administrator / Analyst	Configure Reporting Engine	For more information, see "Step 3: Configure Reporting Engine Data Sources" topic in the *Reporting Engine Configuration Guide*
Administrator / Analyst	Create a List or List Group/Create or Deploy a Rule/Test a Rule	Configure a Rule
Administrator / Analyst	Create and Schedule a Report	Create and Schedule a Report
Administrator / Analyst	View a report or list of all reports*	View a Report
Administrator / Analyst	Investigate a Report	Investigate a Report
Administrator / Analyst	Manage/Access Control for lists, Rules or Reports	Manage Lists, Rules or Reports

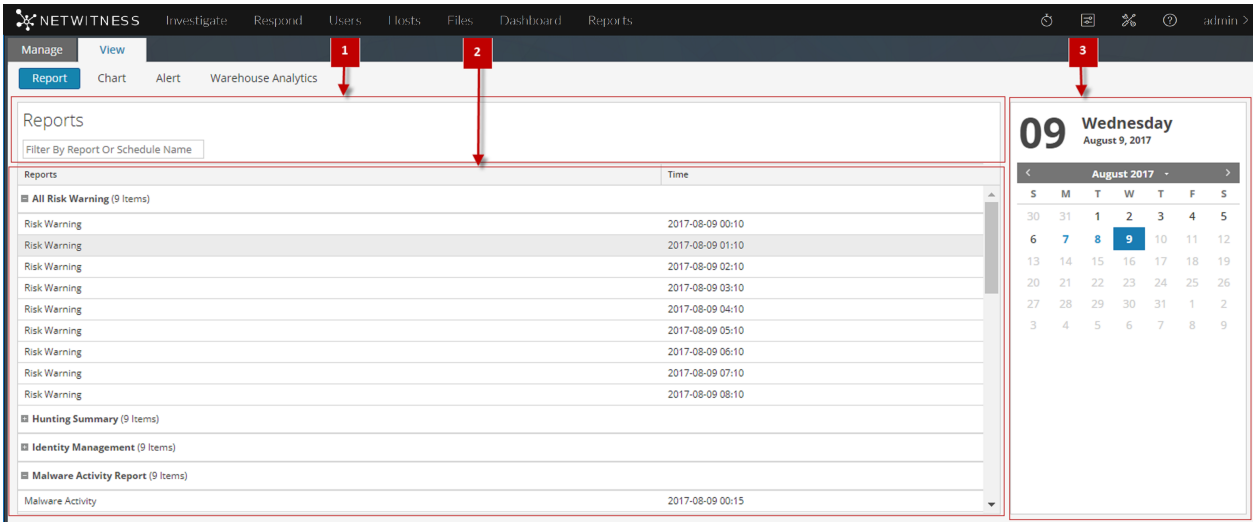
*You can complete these tasks here.

Related Topics

- Configure and Generate a Report
- Build Report View

- Import Report Dialog
- Scheduled Reports View
- Reports Permissions Dialog
- View a Report Panel
- Report View

Quick View



To access this view:

1. Go to **Reports**.
   The Manage tab is displayed.
2. Click **Reports**.
   The Report view is displayed.
3. In the **Reports** panel, click **View All Reports**.
   The Reports panel is displayed, clicking on any of the reports listed allows you to view the report.

Features

The View All Reports panel has the following features.

1	Reports Toolbar
2	Reports Output panel
3	Reports Calendar panel

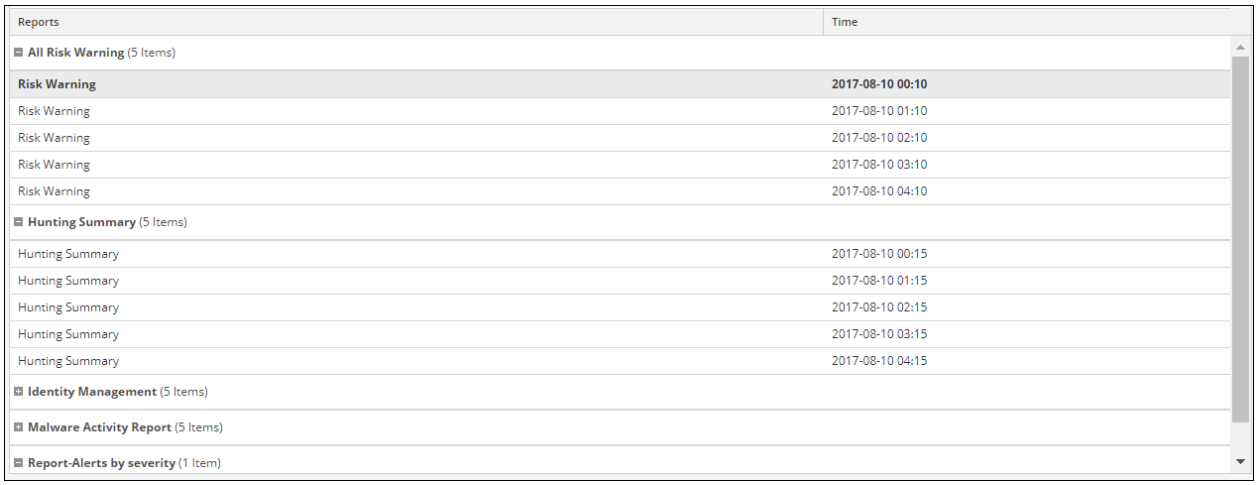
Reports Toolbar

The following table lists the options in the View All Reports toolbar:

Operation	Description
Filter By Report Or Schedule Name	Searches schedules based on the report name or schedule name for a selected calendar day.

Reports Output Panel

The Reports Output panel displays the report with the report schedule name and report generated time.



Feature	Description
Reports	This field displays the detailed report with the selected rule variables.
Time	This field displays the time when the report is generated.

Reports Calendar View

The Reports Calendar view is used to select a date from the Calendar. Based on the date you select, the list of successfully run reports for the date is displayed.


10
Thursday
August 10, 2017
August 2017
S M T W T F S
30 31 1 2 3 4 5
6 7 8 9 10 11 12
13 14 15 16 17 18 19
20 21 22 23 24 25 26
27 28 29 30 31 1 2
3 4 5 6 7 8 9

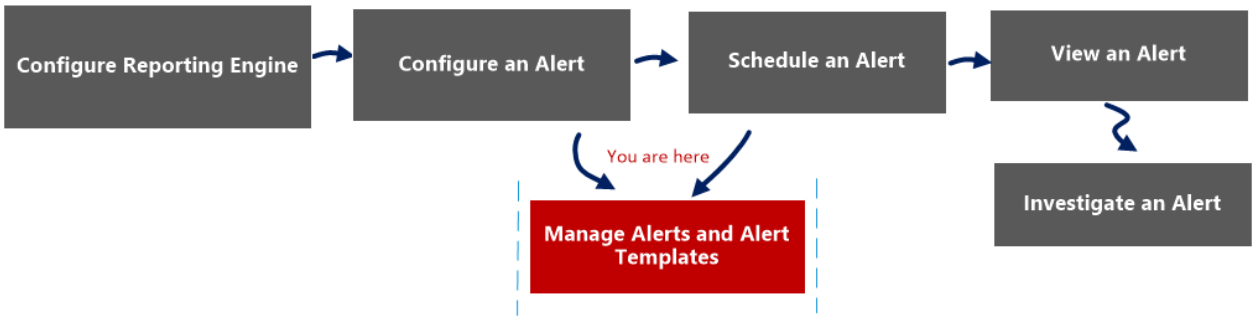
Alerting References

The Reporting module user interface provides access to NetWitness alerts. This topic contains descriptions of the user interface as well as other reference information to help users manage Alerts.

Alert List View

The Alert List view allows you to import, export, manage, and add alerts.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template*	Manage an Alert and Alert Template

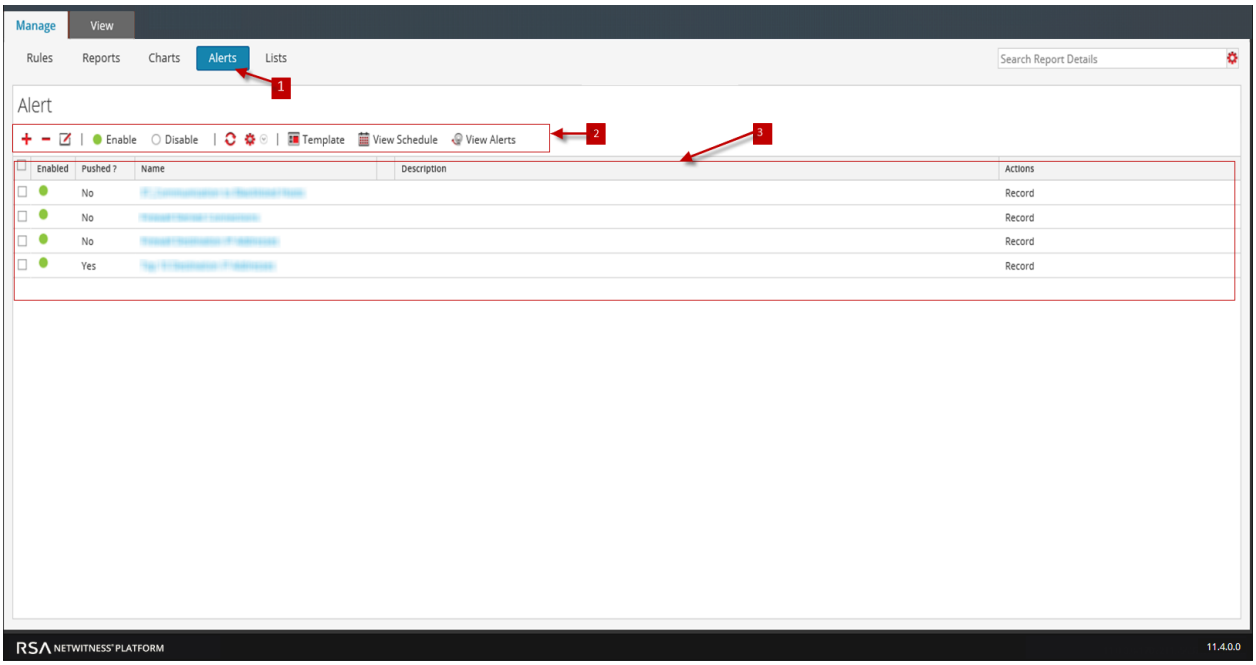
*You can complete these tasks here.

Related Topics

Alerting Overview

Quick View

The following figure is an example with the important features labeled.



1 Click **Alerts** to open the Alert view.

2 The Alert toolbar allows you to add, modify, delete, enable, disable, refresh, import, and export an alert. Using this toolbar, you can also set access permissions for the selected alert.

3 The Alert panel lists all the alerts in a tabular format.

The Alerts List view has the following panels:

- Alert Toolbar
- Alert

Alert Toolbar

The Alert toolbar panel has the following features:

Feature	Description
+	Adds a new alert to the Reporting module.
−	Deletes one or more selected alerts.
	Edits an alert.
Enable	Enables the selected alerts.
Disable	Disables the selected alerts.
	Refreshes the view.
	Enables the following options: Import, Export and Permissions.

Alert

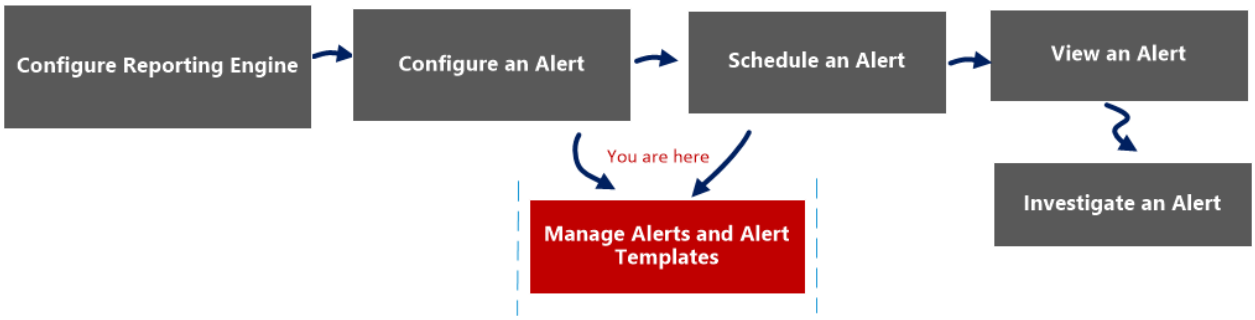
The Alert panel lists all the alerts in a tabular format. The following table lists the columns in the Alert panel and their descriptions.

Feature	Description
Enabled	Displays the state of the alert: • Enabled - the alert is active and fires based on the rule assigned to it. • Disabled - the alert is not active.
Pushed?	Indicates whether the alert is sent to Decoders or Log Decoders: • Yes - Alert is pushed to Decoders or Log Decoders. • No - Alert is not pushed to Decoders or Log Decoders.
Name	Identifies the name of the alert. Clicking the alert name displays the rule on which this alert is based in the Define Rules panel.
Description	Indicates the alert description.
Actions	Indicates the action the system takes when the alert fires. The different available action types are as follows: • Record • SMTP • SNMP • Syslog

Alert Permissions Dialog

In the Alert Permissions dialog, the users with 'Read & Write' access permission can set access permissions for an alert to configure permissions in the Alert Permissions dialog.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template*	Manage an Alert and Alert Template

*You can complete these tasks here.

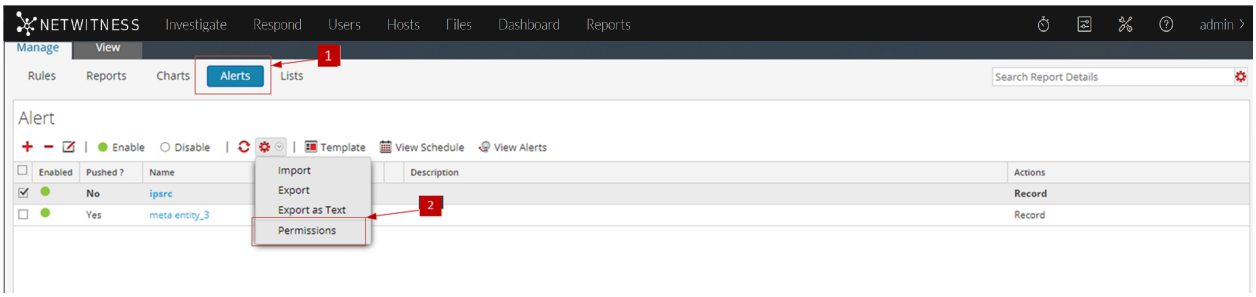
Related Topics

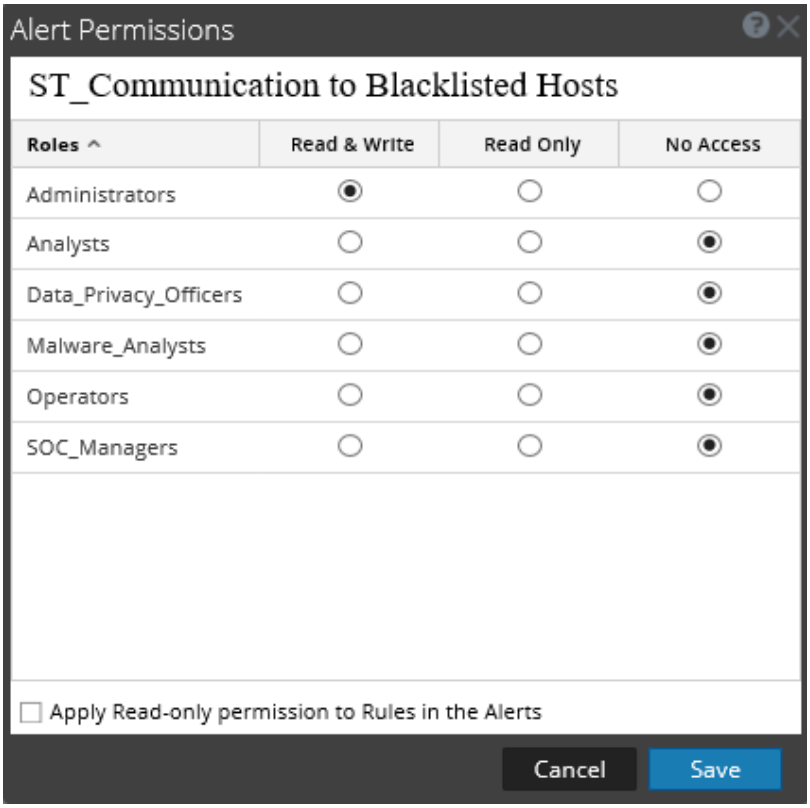
Alerting Overview

Quick View

The Alert permissions dialog allows you to set alert permissions depending on the user role.

The following figure is an example with the important features labeled.





1. Click **Alerts** to open the Alert view.
2. Click ⚙ > **Permissions**. The Alert Permissions dialog box is displayed.
3. Based on the user role, select the appropriate options.
4. (Optional) Select the checkbox if you want to automatically provide read access permission to dependent rules.
5. Click **Save**.

Note: If a User (other than a super user) creates an alert, super users will not be able to access the alert.

The following table lists the columns in the Alert Permissions dialog.

Column	Description
Roles	Displays all the user roles in the NetWitness user interface.

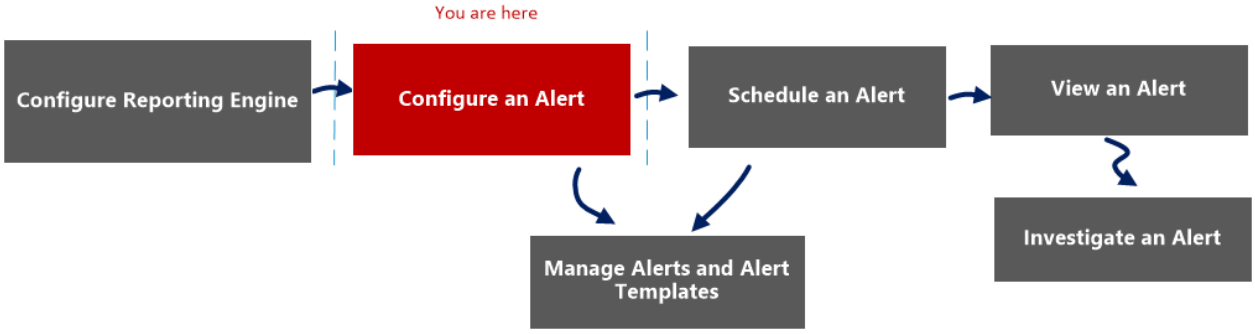
Read & Write	Allows you to apply 'Read&Write' access to the alert.
Read Only	Allows you to apply only 'Read' access to the alert.
No Access	By selecting this permission, you cannot access or view the alert.
Apply Read-only permission to Rules in the Alerts	Allows you to automatically apply permissions to the rules in the alerts.
Cancel	Cancels all the changes made to the permissions.
Save	Saves the selection and provides access to the role based on the selection.

Alert Schedules View

In the Alert Schedules view, you can view all the alerts scheduled. Alternately, you can also disable the scheduled alerts.

Workflow

The following workflow shows the tasks involved in creating or modifying an alert.



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert*	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template	Manage an Alert and Alert Template

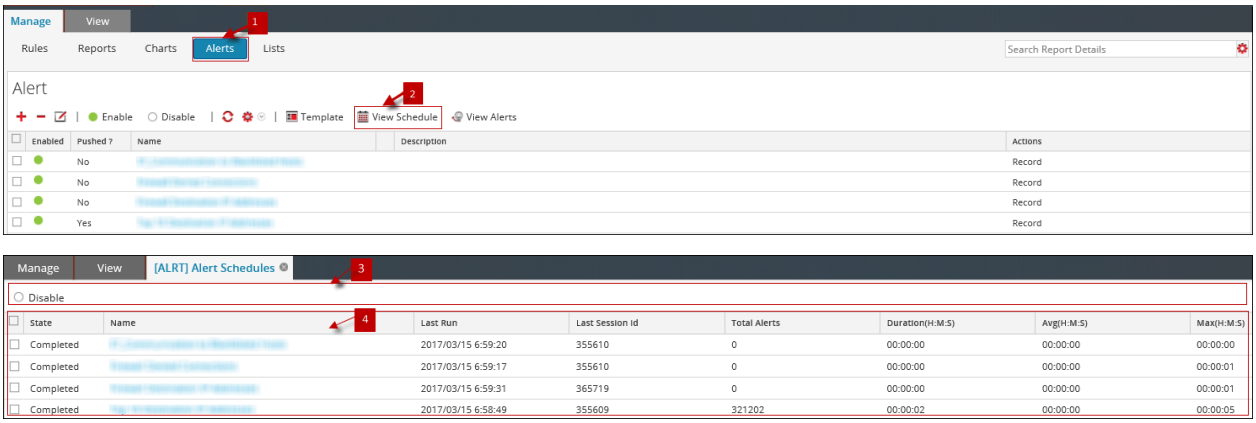
*You can complete these tasks here.

Related Topics

Alerting Overview

Quick View

The following example shows you how to access the Alert Schedules view dialog.



1 Click **Alerts** to open the Alert view.
2 Click **View Schedule** to open the View Alerts Schedule view.
3 The Alerts Schedule toolbar allows you to modify the state of the scheduled alert.
4 The Alerts Schedule List panel lists only the Enabled alerts in a tabular format.

Features

The different panels on the Alert Schedules View dialog are:

- Alerts schedule toolbar panel
- Alerts schedule list panel

Alerts Schedule Toolbar Panel

In the Alerts Schedule Toolbar panel, the Disable icon disables the selected alert. When schedule alerts are no longer needed or are determined to be ineffective, you can disable them so that they are no longer executed. You can select one of more alerts to disable. When an alert is disabled, it is removed from the scheduled alerts list so that you cannot view it here, and it will not execute again unless you manually execute the alert or set up a new schedule for it.

Alerts Schedule List Panel

The following table lists the columns in the Alerts Schedule List panel and their description.

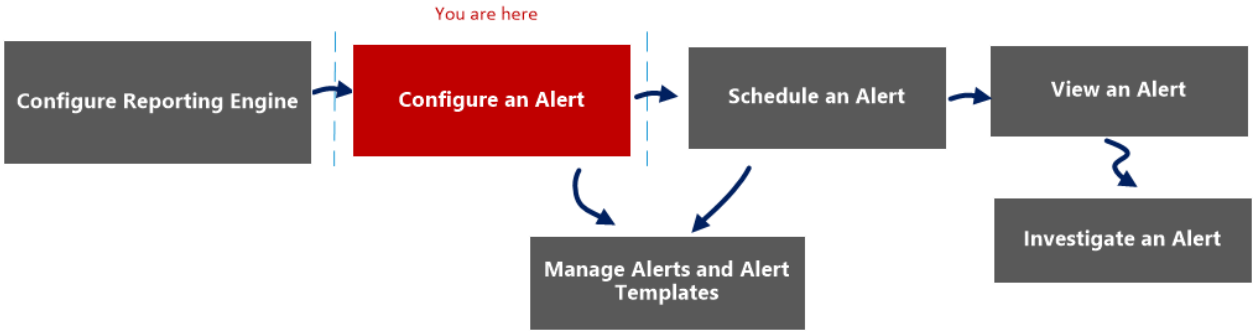
Column	Description
State	The state of the scheduled alert: • Completed • Failed
Name	The name of the scheduled alert.
Last Run {#time}	The last time the scheduled alert was run.
Last Session Id	The Session Id of the last scheduled alert.
Total Alerts	The total number of event occurrences.
Duration	The time taken to run the scheduled alert.

Column	Description
Avg (s)	The average time taken to run the scheduled alert.
Max (s)	The maximum time taken to run the scheduled alert.

Create or Modify Alert Panel

The Create or Modify alert panel is a panel in the Alert List view. This panel allows you to create or modify an alert as per the requirement.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert*	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template	Manage an Alert and Alert Template

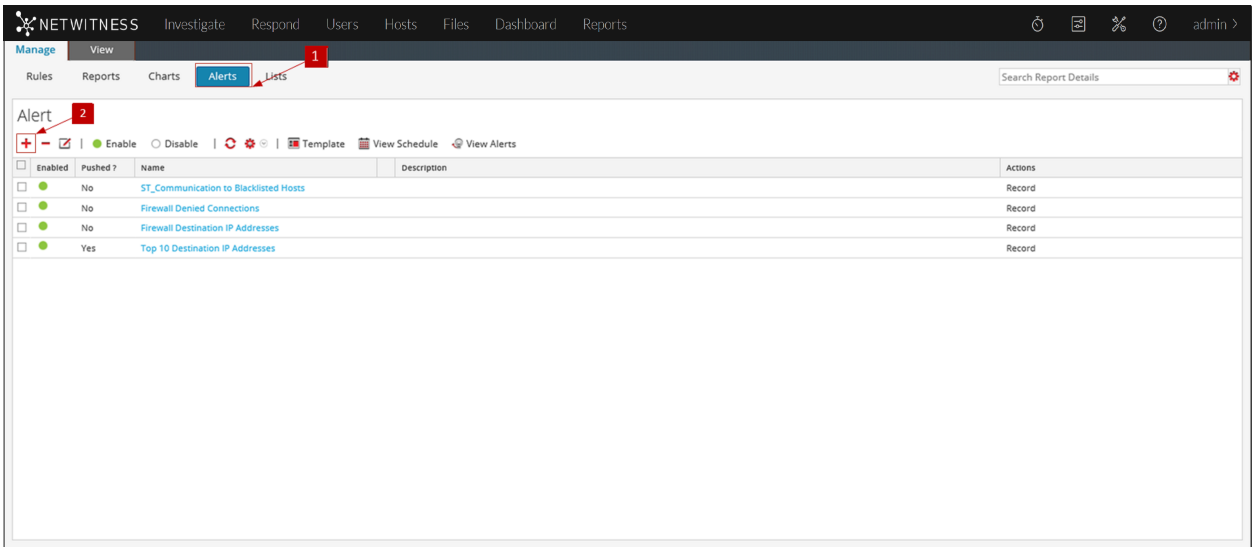
*You can complete these tasks here.

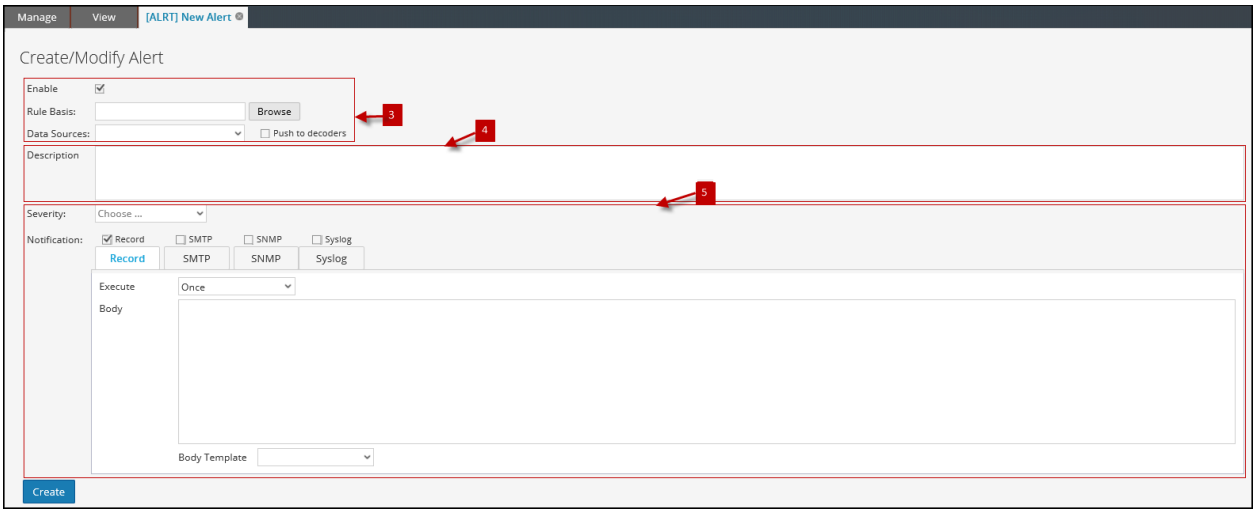
Related Topics

Alerting Overview

Quick View

The following figure is an example with the important features labeled.





1. Click **Alerts** to open the Alert view.
2. Click + to navigate to the Create or Modify Alert panel.
3. Enable the alert, navigate the rule, and select a data source to alert.
4. Enter a brief description of an alert.
5. Define the alert notification methods(RECORD, SMTP, SNMP, Syslog) to alert, when an alert condition is matched.

The Create or Modify Alert panel has the following sections:

- Alert Definition
- Alert Description
- Alert Notification

Alert Definition

The following table describes the fields in the Alert Definition:

Field	Description
Enable	• **Enable** activates the alert. The alert executes and sends output actions every minute (by default) when the alert conditions are met. • **Disable** deactivates the alert. The alert does not execute and does not send any output actions.
Rule Basis	Click **Browse** to display the Rules Library panel from which you select the rule that is the basis of this alert. You must select a rule that has a unique 'where' clause for an alert.
Data Sources	Specifies the data source for the alert.
Push to decoders	Pushes the ‘where’ clause of the alert rule to Decoders connected to the selected NWDB data source. This is the recommended option used to create RE alerts, as the alert conditions are checked on the Decoder itself and the alert queries will be comparatively faster in NWDB. If you deselect this option, the alert rule ‘where’ clause will be queried against the selected NWDB data source. Based on the complexity and metas in the ‘where’ clause of the rule, the alert queries might take more time to process in NWDB. **Note:** RSA NetWitness does not send rules to the Decoder automatically.

Alert Description

The following table describes the fields in the Alert Description:

Field	Description
Description	Describes the alert.
Create	Creates the alert. (This option is displayed when you create an alert.)
Save	Saves the changes made to the alert. (This option is displayed when you modify an alert.)

Alert Notification

The Alert Notification allows you to define the notification action NetWitness takes when an alert is generated, for example, recording or sending the alert using one of the defined output actions. The output actions are Simple Mail Transfer Protocol (SMTP), Simple Network Management Protocol (SNMP), or Syslog message.

The Notification contains the default Record tab, which you use to create an alert. The icon beside the Record tab allows you to select the notification type from the drop-down list for the output to specify for the alert: SMTP, SNMP, or Syslog.

Depending on the selected notification type, the Notification section is populated with predefined text that contains variables that add Meta that is appropriate for the alert. In the Reporting Engine, these variables are replaced with actual values. The following table lists the variables and their descriptions.

Variable	Description
`${meta.<metakey>}`	The meta key value. **Note:** If the <metakey> did not fetch any value, an empty string("") is printed. By default, Reporting Engine displays all the repeated values for a meta key. If you do not want the meta values to repeat in the Alert output, enable the "removeRepeatedMetaValue" option by navigating to **Configuration > Alert Configuration** available for the Reporting Engine in the **Services - Configuration > Explore** view. For example, in an HTTP Session the value for the action is displayed as get, get, put, put, post, get. When this option is enabled, the value is displayed as get, put, post.
`${meta.time}` / `${meta.time:<time_format>}`	`${meta.time}` - The session time is printed in "yyyy-MMM-dd HH:mm:ss" format. `${meta.time:<time_format>}` - The session time is printed in the user-defined custom time format. For example, `${meta.time:dd-MM-yyyy HH:mm:ss}`. For more information on the supported time formats, see http://docs.oracle.com/javase/7/docs/api/java/text/SimpleDateFormat.html **Note:** If the time format provided by the user is invalid, the default time format will be used. The default time format is "yyyy-MMM-dd HH:mm:ss".
`${name}`	The alert name defined in Reporting Engine.
`${count}`	The number of times an alert is detected in a given time frame. (By default, it is one minute)
`${nw.host}`	The NetWitness host name as configured in Reporting Engine.
`${device.id}`	The NetWitness device ID of the data source.

The Alert Notification view has four tabs:

- **Record Tab**
- **SMTP Tab**
- **SNMP Tab**
- **Syslog Tab**

Record Tab

Use the Record tab to define the frequency for recording an alert and the message to generate when an alert is generated.



The following table lists the fields in the Record tab and their description.

Field	Description
Execute	The frequency for recording an alert. • **Once** - Record the alert only once based on the alert interval no matter how often the alert is generated. NetWitness records the number of times the alert has actually generated during that interval in the log file so that analysts know how many times the alert registered a match over a given day. • **Each Event** - Record the alert each time as it generates. If an alert generates unlimited number of times during a day, that alert is often treated as noise and can be ignored, except in case of alerts that require continuous monitoring such as network configuration changes and DDOS attacks. **Note:** Select **Each Event** setting from the **Execute** drop-down list for SNMP and Syslog output actions.
Body	The body of the message.
Body Template	(Optional) If templates have been defined, select a template for the alert message.

SMTP Tab

The SMTP tab allows you to define the SMTP (email) output for this alert.



The following table lists the fields in the SMTP tab and their description.

Field	Description
Execute	The frequency to send an email message for the alert. • **Once** - Sends only one email for an interval, if an alert generates in that interval, irrespective of how many alerts generated. • **Each Event** - Send an email with the alert for every event in which the rule criteria are met.
To	The email addresses to which to send this alert.
Subject	The subject of the email message.
Body	The body of the message.
Body Template	(Optional) If templates have been defined, select a template for the SMTP message that you can use as is or modify.

SNMP Tab

The SNMP tab allows you to define the SNMP output for the alert.

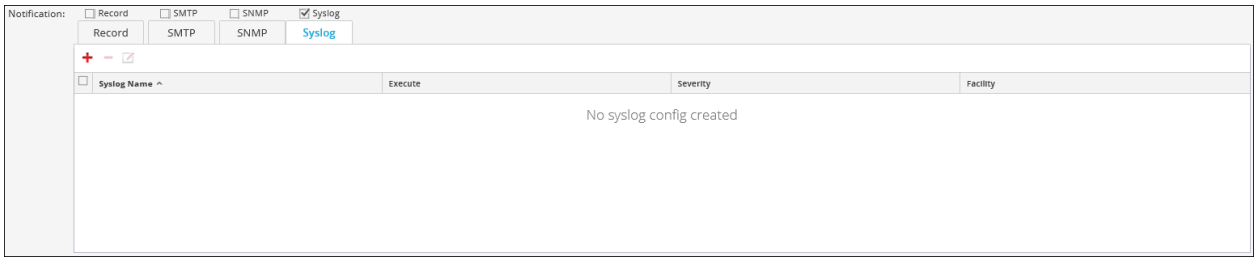


The following table lists the various fields in the SNMP tab and their description.

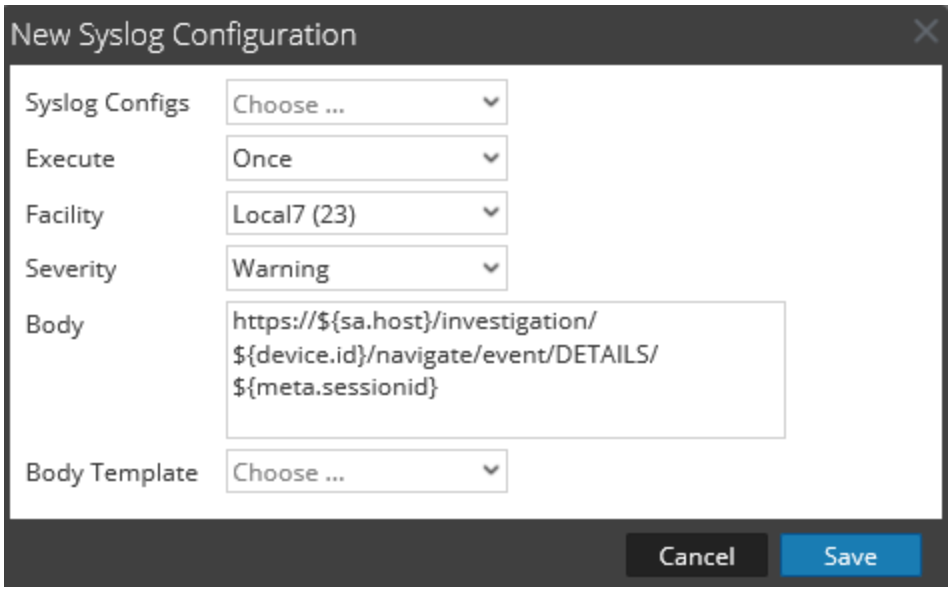
Field	Description
Execute	The frequency to send an SNMP output for an alert. • **Once** - Sends an SNMP message along with an email for an interval, if an alert generates in that interval, irrespective of how many alerts generated. • **Each Event** - Sends an SNMP message with the alert for every event in which the rule criteria are met.
Body	The body of the message.
Body Template	(Optional) If templates have been defined, select a template for the SNMP message to use as is or modify.

Syslog Tab

The Syslog tab allows you to define the Syslog message output for this alert.



Click  to add Syslog configuration to an alert. The New Syslog Configuration dialog box is displayed:



The following table describes the fields in the New Syslog Configuration dialog:

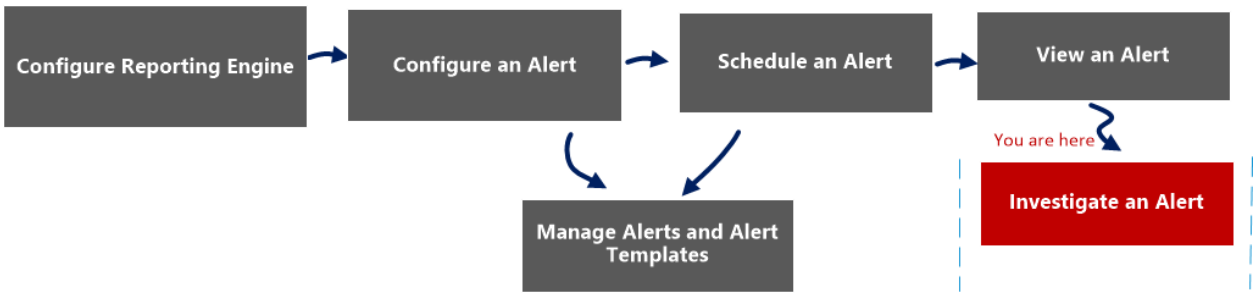
Field	Description
Syslog Configs	The Syslog configuration of the Device Config view located at the Syslog Configuration panel.
Execute	The number of times that you want to send a Syslog output for the alert. • **Once** - Sends a Syslog output along with an email for an interval, an alert generates in that interval, irrespective of how many alerts generated. • **Each Event** - Sends a Syslog output with the alert for every event in which the rule criteria are met.
Facility	The type of program logging the message. Examples for the type of programs are Syslog, Daemon, Mail, and Kernel.

Field	Description
Severity	The severity level of the alert that generated. • Emergency • Alert • Critical • Error • Warning • Notice • Informational • Debug
Body	The body of the message.
Body Template	(Optional) If templates have been defined, select a template for the Syslog message to use as is or modify.

Investigate an Alert View

In the Investigate an Alert view, you can view and investigate alert details. When investigating an alert, you can open the sessions in the Investigation module for further investigation.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert*	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template	Manage an Alert and Alert Template

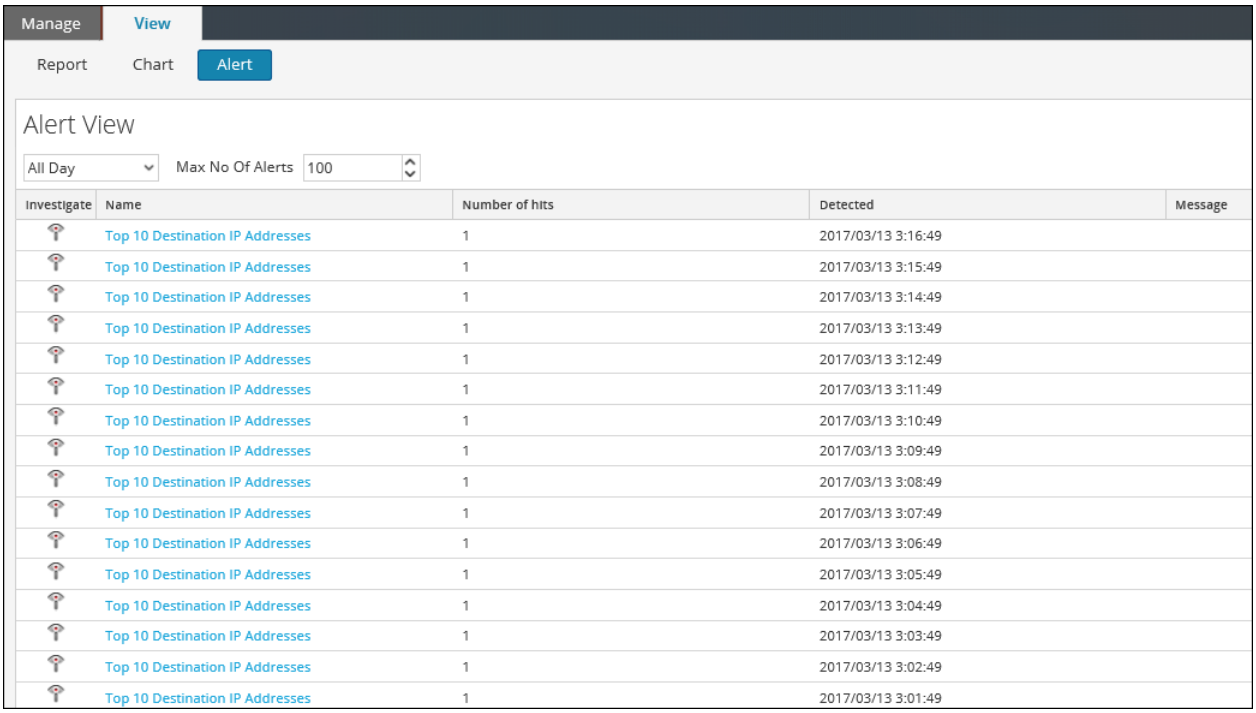
*You can complete these tasks here.

Related Topics

Alerting Overview

Quick View

The following figure is an example with the important features labeled.



The View an Alert view has the following panels:

- View Alerts Toolbar
- View Alerts List

View Alerts List

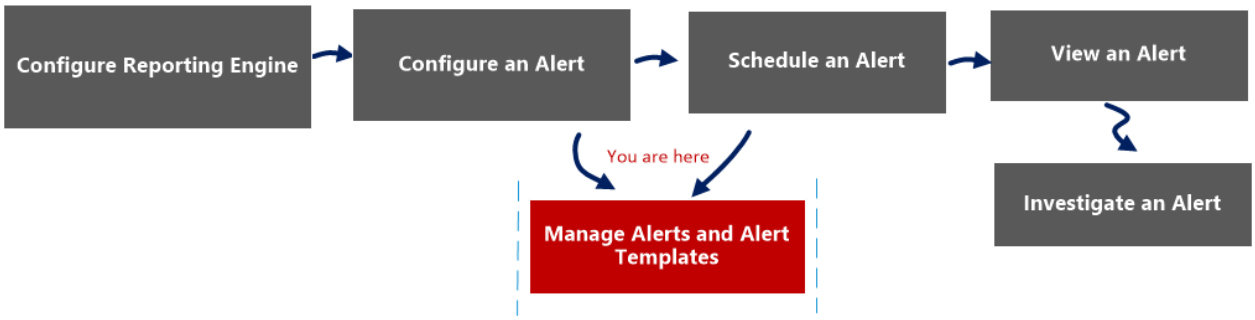
The following table lists the columns in the View Alerts List panel.

Column	Description
	The icon that opens the Investigation module, where the details of the first session that registered the match for the given alert is displayed for immediate analysis. **Note:** You are not redirected to the Investigation module when: -You reconfigure a data source for an existing alert and run an alert on the new data source. -You enter a host name instead of an IP address in the data source field.
Name	The name of the alert that registered the match. The hyperlink on the name opens the Investigation module to view all matches for that particular alert for the hour surrounding the registered alert.
Number of hits	The number of times the alert is generated.
Detected	The date and time at which the alert generates.
Message	The alert message.

Import Alert Dialog

The Import Alert dialog allows you to import an alerts archive and specify whether to overwrite existing rules, lists, and alerts.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template*	Manage an Alert and Alert Template

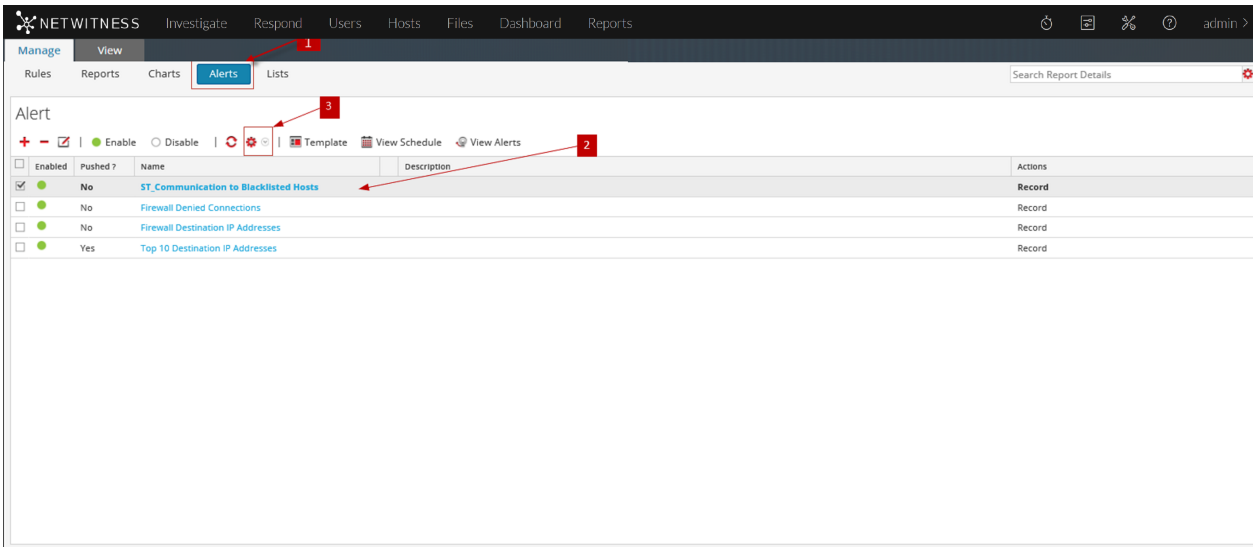
*You can complete these tasks here.

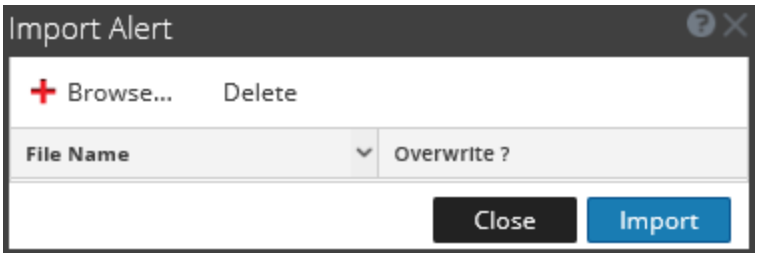
Related Topics

Alerting Overview

Quick View

The following figure is an example with the important features labeled.





1 Click **Alerts** to open the Alert view.
2 In the **Alert** panel, select a folder to import the file.
3 In the **Alert** toolbar, click > **Import** to import an alert.

The following table lists the actions in the Import Alert dialog and their description.

Actions	Description
Browse...	Displays a view of the local zip file system so that you can select the alert to be imported.
−	Deletes the selected alert from the Import Alert dialog.
File Name	Name of the imported binary file.
Overwrite?	Selects the option to overwrite an existing version of the alert you are importing. If you do not select the Overwrite option, a duplicate file is imported and no error message is displayed.
Close	Closes the Import Alert dialog.
Import	Imports the alert with a confirmation message.

Alert Template References

The Reporting module user interface provides access to NetWitness alerts and alert templates as well. This topic contains descriptions of the user interface as well as other reference information to help users manage alert templates.

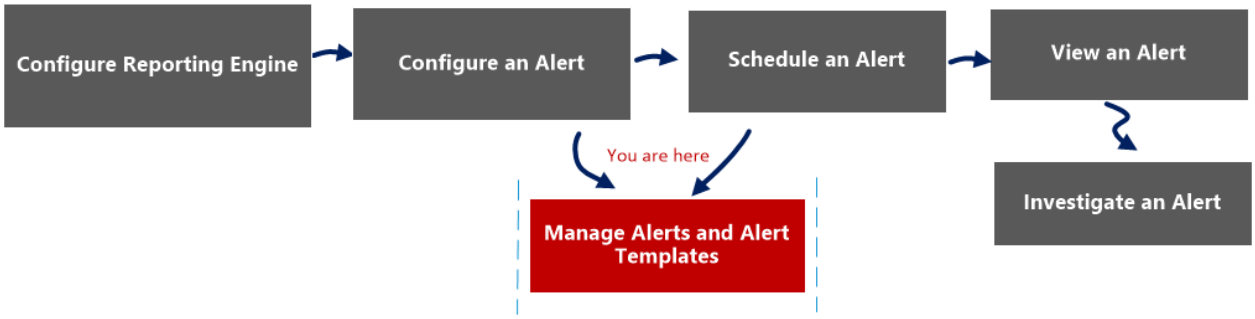
Topics:

- Create or Modify Template View
- Template View

Alert Template View

In the Template view, you can add, modify, view, and delete alert templates.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template*	Manage an Alert and Alert Template

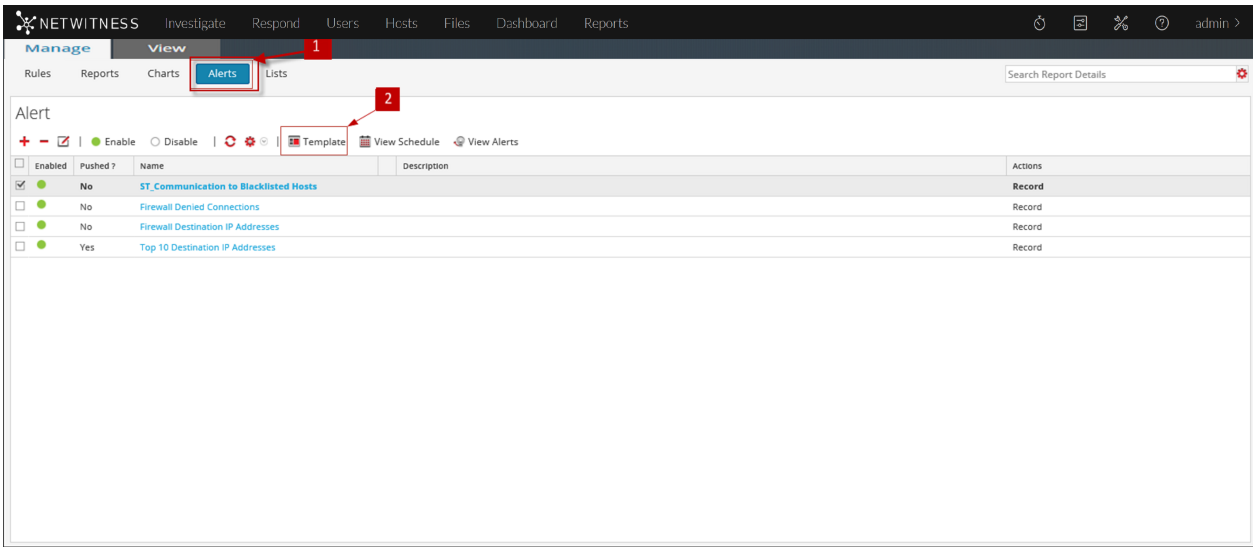
*You can complete these tasks here.

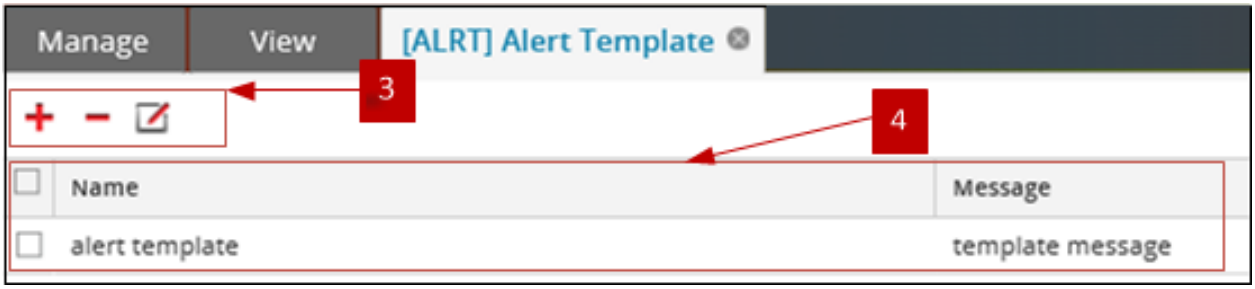
Related Topics

Alerting Overview

Quick View

The following figure is an example with the important features labeled.





1 Click **Alerts** to open the Alert view.
2 Click Template to open the Template view.
3 The Template toolbar allows you to add, modify, and delete alert templates.
4 The Template List panel allows you to view a list of all the templates in a tabular format.

The Alert Template view has the following panels:

- Template Toolbar
- Template List

Template Toolbar

Once the templates are defined, you can select a template to simplify defining and modifying alert messages.

The following table lists the various actions in the Template view and their description.

Actions	Description
+	Creates a new alert template.
−	Deletes the selected alert template.
✎	Edits an existing alert template.

Template List

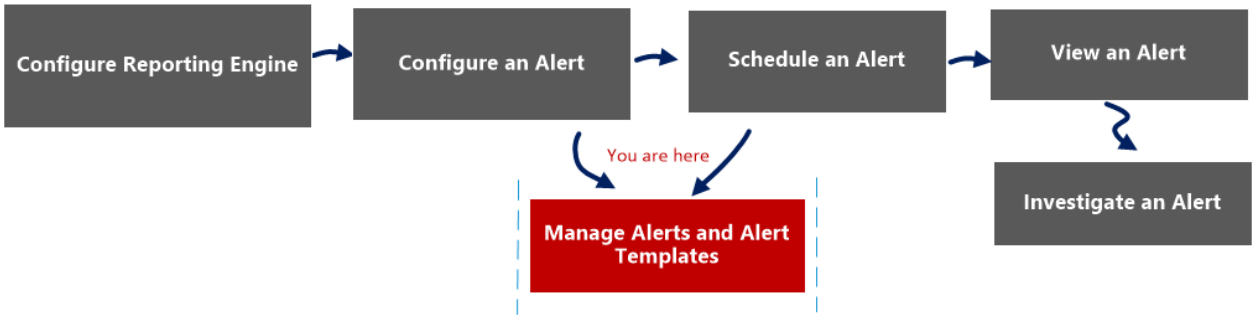
The following table describes the columns in the Templates List panel.

Column	Description
Name	Name of the template.
Message	Alert message defined for the template.

Create or Modify Template View

In the Create/Modify Template view, you can customize alert templates to use when creating alerts.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template*	Manage an Alert and Alert Template

*You can complete these tasks here.

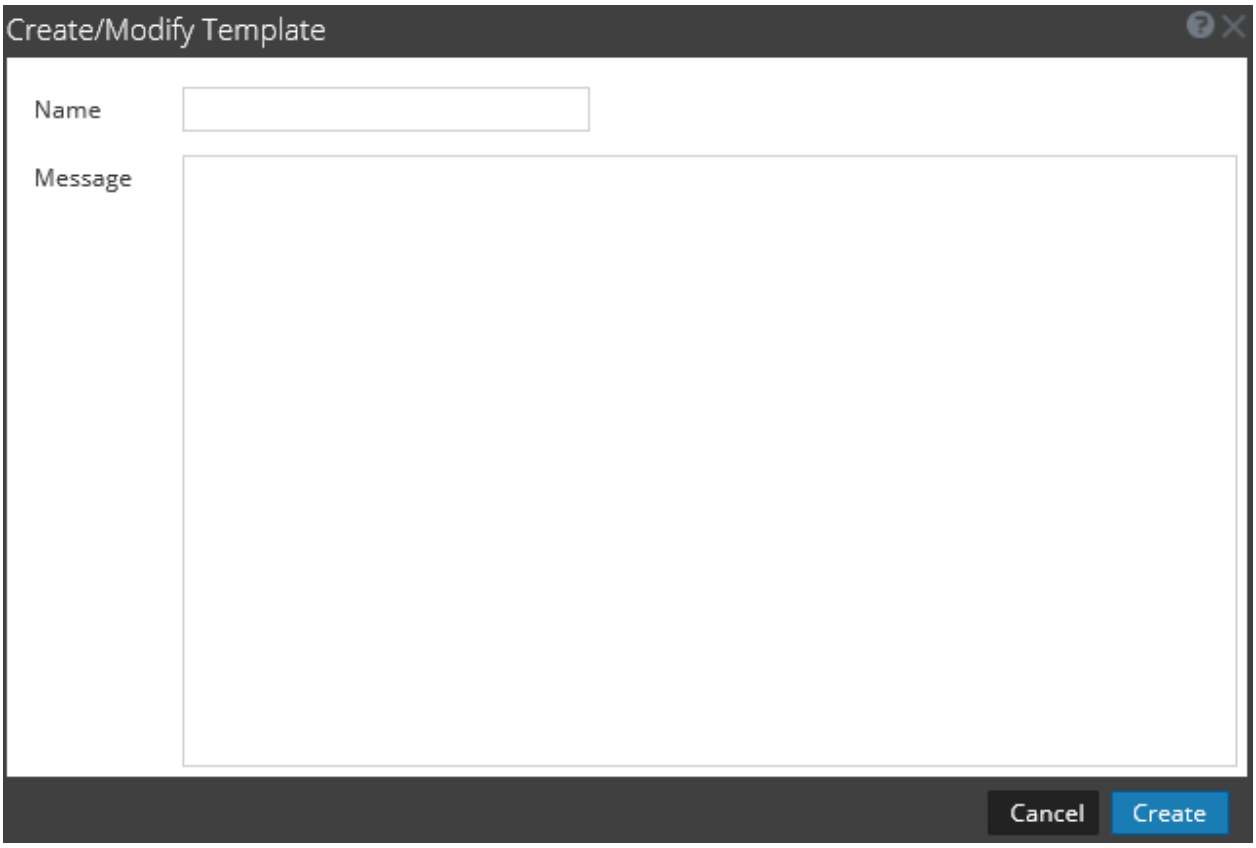
Related Topics

Alerting Overview

Quick View

You can create or modify an alert template name and message on this view.

The following figure is an example of the Create or Modify alert template.



The following table describes the fields in the Create/Modify template.

Feature	Description
Name	Indicates the name of the template for Reporting alerts. For example, source IP.
Message	Specifies the message that will be sent when an alert is triggered.
Create	Creates the template with a confirmation message and becomes available for use in Reporting immediately.
Save	Saves the template with the edited details or when a new template is created. This button is visible only in the edit mode.
Cancel	Closes the dialog without saving the template or any changes made to the template.

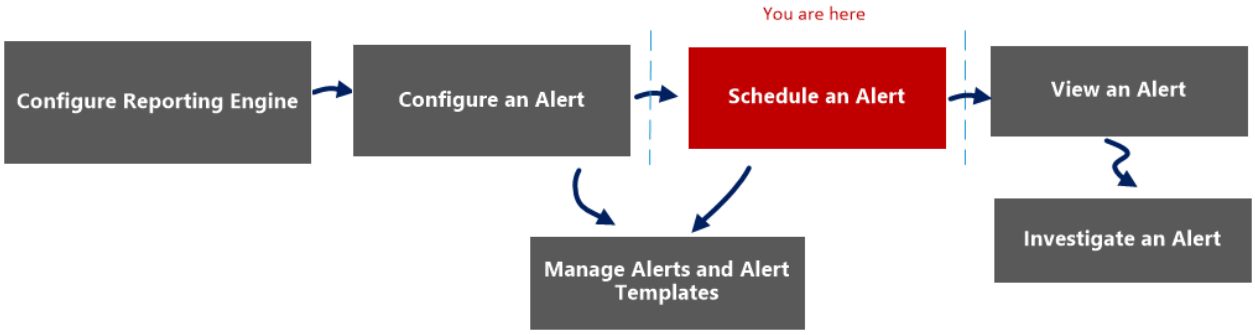
View Alerts Schedule View

In the View Alerts Schedule view, you can view the following information about each of your scheduled alerts.

- Completion status, name, last run time, last session ID, total alerts triggered.
- Statistics about the time taken to run the scheduled alert: duration, average duration, maximum duration.

Note: You can also disable the scheduled alerts.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert*	Schedule an Alert
Administrator/ Analyst	View an alert	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template	Manage an Alert and Alert Template

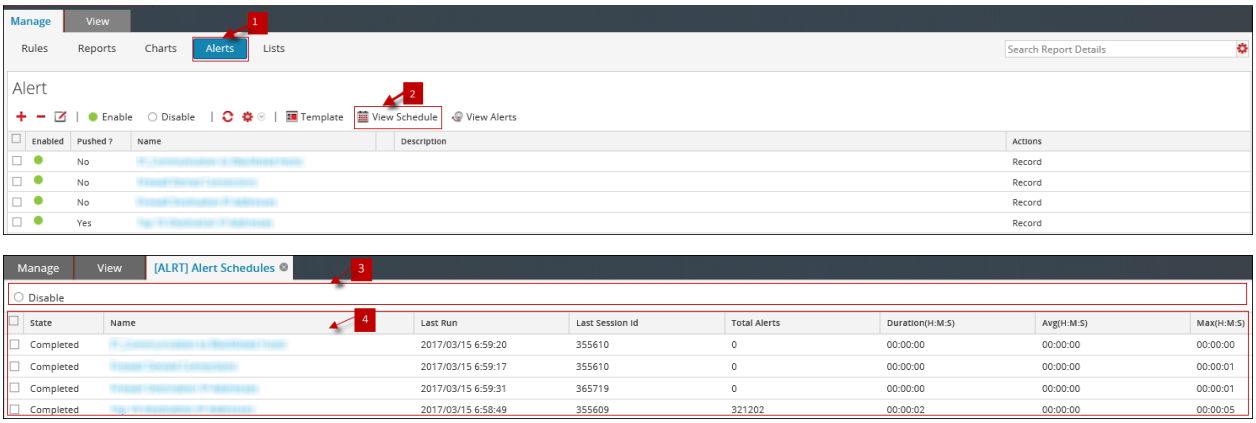
*You can complete these tasks here.

Related Topics

Alerting Overview

Quick View

The following figure is an example with the important features labeled.



1. Click **Alerts** to open the Alert view.
2. Click **View Schedule**to view all the alerts scheduled.
3. The Alerts schedule toolbar allows you to disable the scheduled alert.
4. The Alerts schedule list allows you to view the scheduled alert details.

The View Alerts Schedule view includes the following panels:

1. Alerts Schedule toolbar
2. Alerts Schedule list

Alert Schedule Toolbar

The Alerts Schedule Toolbar panel allows you to modify the state of the scheduled alert.

Feature	Description
Disable	Clicking **Disable** disables the selected alert. When schedule alerts are no longer needed or are determined to be ineffective, you can disable them so that they are no longer executed. You can select one of more alerts to disable. When an alert is disabled, it is removed from the scheduled alerts list so that you can't view it here, and it will not execute again unless you manually execute the alert or set up a new schedule for it.

Alert Schedule List Panel

The Alerts Schedule List panel lists only the Enabled alerts in a tabular format. The following table lists the columns in the Alerts Schedule List panel and their description.

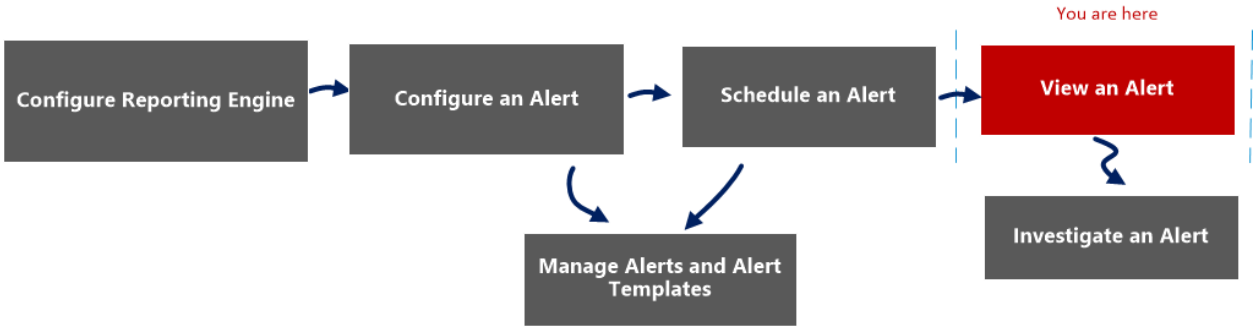
Feature	Description
State	The state of the scheduled alert: • Completed • Failed

Feature	Description
Name	The name of the scheduled alert.
Last Run {#time}	The last time the scheduled alert was run.
Last Session Id	The Session Id of the last scheduled alert.
Total Alerts	The total number of event occurrences.
Duration	The time taken to run the scheduled alert.
Avg (s)	The average time taken to run the scheduled alert.
Max (s)	The maximum time taken to run the scheduled alert.

View Alerts View

In the View Alerts view, you can view all the alerts. Also, you can also customize the view to show alerts for a specific period of time, and set the maximum number of alerts displayed in a single page.

Workflow



What do you want to do?

Role	I want to...	Documentation
Administrator/ Analyst	Configure Reporting Engine	Configure Reporting Engine
Administrator/ Analyst	Configure an alert	Configure an Alert
Administrator/ Analyst	Schedule an alert	Schedule an Alert
Administrator/ Analyst	View an alert*	View an Alert
Administrator/ Analyst	Investigate an alert	Investigate an Alert
Administrator/ Analyst	Manage an alert and alert template	Manage an Alert and Alert Template

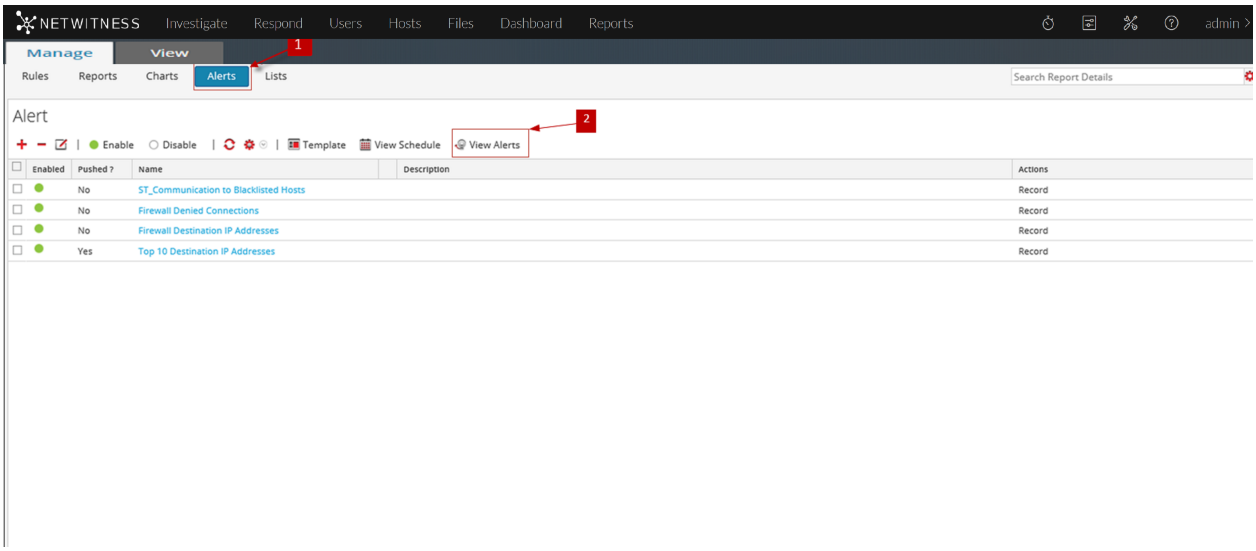
*You can complete these tasks here.

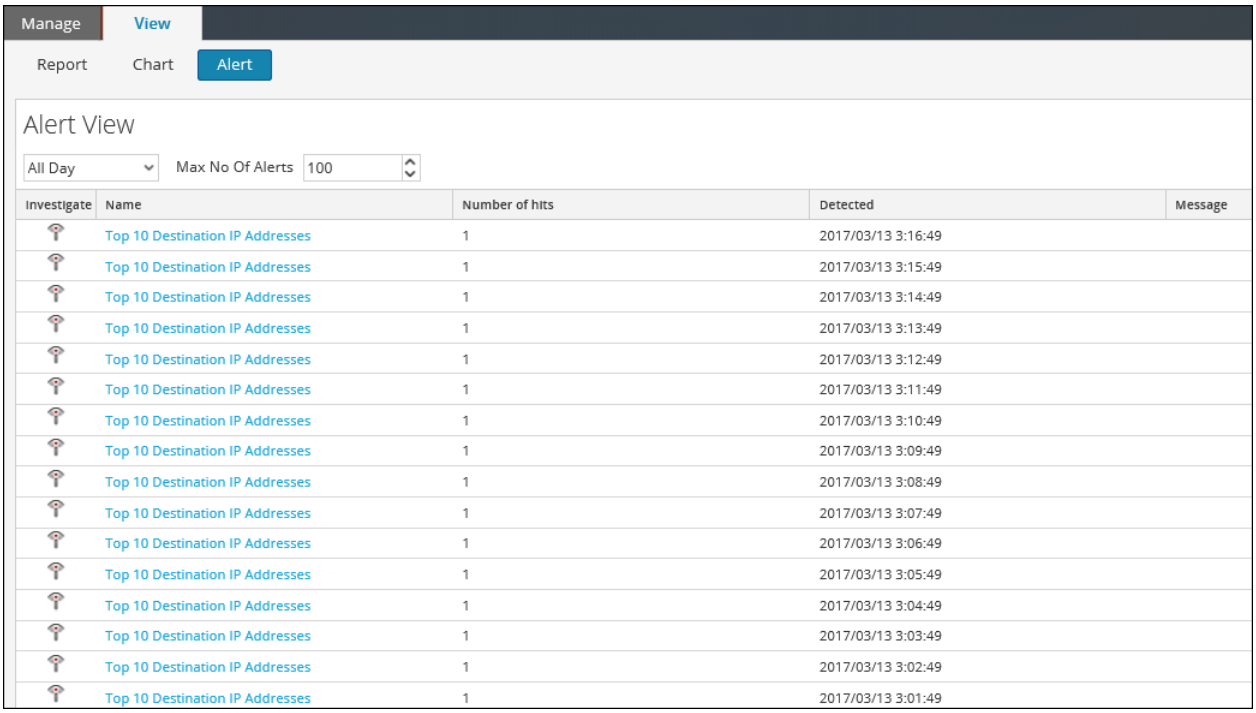
Related Topics

Alerting Overview

Quick View

The following figure is an example with the important features labeled.





1. Click **Alerts** to open the Alert view.
2. Click **View Alerts** to view the different panels on View Alerts.
3. The View Alerts toolbar allows you to filter alerts based on a count, or the start and end date of the alerts.
4. The View Alerts List lists all the filtered alerts in a tabular format.

The View Alerts view has the following panels:

- View Alerts Toolbar
- View Alerts List

View Alerts Toolbar

The following table lists the operations in View Alerts toolbar panel.

Option	Description
Last Hour(s) data	The data fetched from the previous execution.
Max No Of Alerts	The maximum number of alerts that you want to fetch from the Reporting Engine service for a specific time-range.

View Alerts List

The following table lists the columns in the View Alerts List panel.

Column	Description
	The icon that opens the Investigation module, where the details of the first session that registered the match for the given alert is displayed for immediate analysis. **Note:** You are not redirected to the Investigation module when: -You reconfigure a data source for an existing alert and run an alert on the new data source. -You enter a host name instead of an IP address in the data source field.
Name	The name of the alert that registered the match. The hyperlink on the name opens the Investigation module to view all matches for that particular alert for the hour surrounding the registered alert.
Number of hits	The number of times the alert is generated.
Detected	The date and time at which the alert generates.
Message	The alert message.

NetWitness® Platform XDR

Version 12.0

NetWitness Respond Configuration

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

About this Document **6**

NetWitness Respond Configuration Overview **7**

Configuring NetWitness Respond **9**

Step 1. Configure Alert Sources to Display Alerts in the Respond View 10

Prerequisites 10

Configure Reporting Engine to Display Reporting Engine Alerts in the Respond View 10

Configure Malware Analysis to Display Malware Analysis Alerts in the Respond View 11

Configure NetWitness Endpoint to Display NetWitness Endpoint Alerts in the Respond View 11

Step 2. Assign Respond View Permissions 14

Respond-server 15

Incidents 16

Integration-server 16

Investigate-server 16

Incident Email Notification Settings Permissions 17

Respond Event Analysis Permissions 17

Respond Saved Filter Permissions 18

Respond Role Permission Examples 18

Restrict Access to Incidents 19

Step 3. Enable and Create Incident Rules for Alerts 21

Enable Incident Rules 22

Create an Incident Rule 26

Verify the Order of Your Incident Rules 29

Clone an Incident Rule 29

Edit an Incident Rule 29

Export Incident Rules 30

Import Incident Rules 31

Additional Procedures for Respond Configuration **33**

Set Up and Verify Default Incident Rules 34

Set Up the User Behavior Incident Rule 34

Deploy the RSA Live ESA Rules 34

Adjust and Enable the User Behavior Default Rule (or Create It If You Do Not Have It) 35

Set up or Verify a Default Incident Rule 39

Suspected Command & Control Communication By Domain 40

High Risk Alerts: Malware Analysis 41

High Risk Alerts: NetWitness Endpoint 41

High Risk Alerts: Reporting Engine 42
High Risk Alerts: ESA 42
IP Watch List: Activity Detected 43
User Watch List: Activity Detected 44
Suspicious Activity Detected: Windows Worm Propagation 44
Suspicious Activity Detected: Reconnaissance 45
Monitoring Failure: Device Not Reporting 46
Web Threat Detection 46
User Entity Behavior Analytics 47
Detect AI 47
Create a NetWitness Endpoint Incident Rule using File Hash 47
Configure Risk Scoring Settings for Automated Incident Creation 50
Configure Custom Respond Server Alert Normalization 54
Configure Analyst UI for Respond Server Alert Normalization 57
Configure Incident Email Notification Settings 59
Migration Considerations 60
Set a Retention Period for Alerts and Incidents 61
Prerequisites 61
Procedure 61
Result 62
Obfuscate Private Data 63
Mapping File to Obfuscate Meta Keys 63
Prerequisites 63
Procedure 63
Manage Incidents in Archer Cyber Incident & Breach Response 65
Prerequisites 65
Procedure 65
Configure the Option to Send Incidents to Archer 67
Prerequisites 67
Add Archer as a Data Source for Context Hub 67
Configure Threat Aware Authentication 70
Enable Threat Aware Authentication 70
Obtain SecurID API Key 70
Configure RSA SecurID Access API Key 71
Configure Sync Frequency 72
Configure Meta 72
Set a Counter for Matched Alerts and Incidents 74
Edit the Incident Rules Export ZIP File 76
Incident Rule Export Files Editing Guidelines 76
Configure a Database for the Respond Server Service 78

Prerequisites 78
Procedure 78
NetWitness Respond Configuration Reference 80
Configure View 80
Incident Rules View 81
What do you want to do? 81
Related Topics 81
Quick Look 82
Endpoint Risk Scoring Settings 82
Incident Rules 85
Incident Rules Actions 86
Incident Rule Details View 87
What do you want to do? 87
Related Topics 87
Quick Look 87
Group By Meta Key Mappings 91
Incident Email Notification Settings View 93
What do you want to do? 93
Related Topics 93
Quick Look 93
Aggregation Rules Tab (11.0 and earlier) 96
What do you want to do? 96
Related Topics 96
Quick Look 96
Aggregation Rules List 97
Aggregation Rules Toolbar 97
New Rule Tab (11.0 and earlier) 98
What do you want to do? 98
Related Topics 98
Quick Look 98

About this Document

The *NetWitness Respond Configuration Guide for NetWitness® Platform* provides an overview of NetWitness, detailed instructions on how to configure NetWitness Respond in your network, additional procedures that are used at other times, and reference materials that describe the user interface for configuring NetWitness Respond in your network.

Topics

- NetWitness Respond Configuration Overview
- Configuring NetWitness Respond
- Additional Procedures for Respond Configuration
- NetWitness Respond Configuration Reference

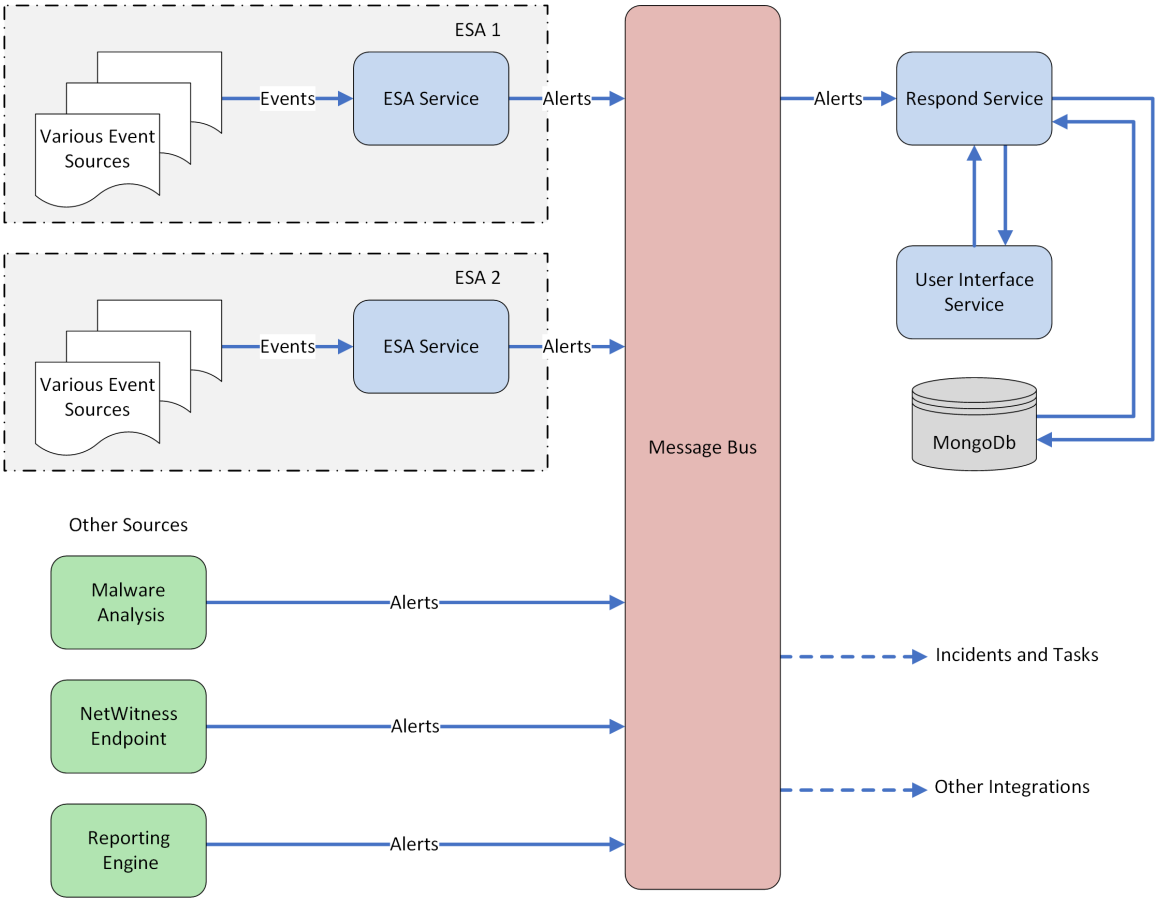
NetWitness Respond Configuration Overview

NetWitness Respond consumes alert data from various sources via the Message Bus and displays these alerts on the NetWitness user interface. The Respond Server service allows you to group the alerts logically and start a NetWitness Respond workflow to investigate and remediate the security issues raised.

The Respond Server service consumes alerts from the message bus and normalizes the data to a common format (while retaining the original data) to enable simpler rule processing. It periodically runs rules to aggregate multiple alerts into an incident and set some attributes of the Incident (for example, severity, category, and so on). The incidents are persisted into MongoDb by the Respond Server service. Incidents are also posted onto the message bus for consumption by other systems (for example, Archer integration).

Note: NetWitness Respond requires an ESA primary server that contains the MongoDb. Alerts, Incidents, and Task records are persisted into this MongoDb by the Respond Server.

The following diagram illustrates the high-level flow of alerts.



You have to configure various sources from which the alerts are collected and aggregated by the Respond Server service.

Configuring NetWitness Respond

This topic provides the high-level tasks required to configure the Respond Server service. The administrator needs to complete the steps in the sequence provided.

Topics

- Step 1. Configure Alert Sources to Display Alerts in the Respond View
- Step 2. Assign Respond View Permissions
- Step 3. Enable and Create Incident Rules for Alerts

Step 1. Configure Alert Sources to Display Alerts in the Respond View

This procedure is required so that alerts from the alert sources are displayed in NetWitness Respond. You have an option to enable or disable the alerts being populated in the Respond view. By default this option is disabled for Detect AI, Reporting Engine, Malware Analysis, and NetWitness Endpoint and enabled only for Event Stream Analysis. So when you install the Respond Server service you need to enable this option in the Reporting Engine, Malware Analysis, and NetWitness Endpoint to populate the corresponding alerts in the Respond view.

Prerequisites

Ensure that:

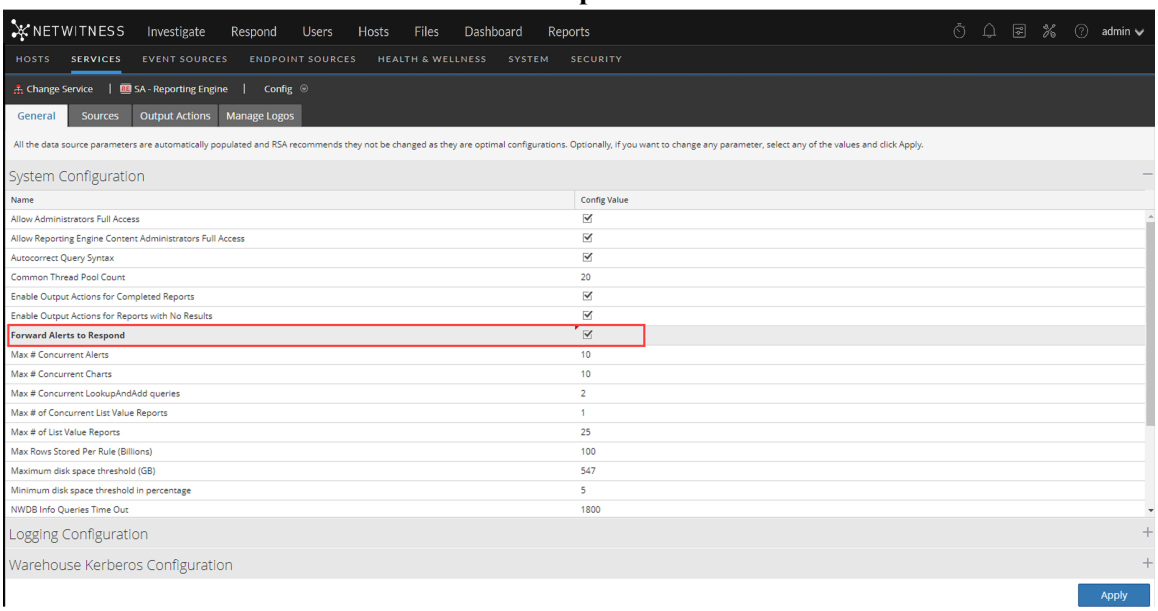
- The Respond Server service is installed and running on NetWitness.
- NetWitness Endpoint is installed and running. This is necessary only if you want to configure NetWitness Endpoint as an alert source in the Respond view.

Configure Reporting Engine to Display Reporting Engine Alerts in the Respond View

The Reporting Engine alerts are by default disabled from being displayed in Respond view. To display and view the Reporting Engine alerts, you have to enable the NetWitness Respond alerts in the Services Config view > General tab for the Reporting Engine.

1. Go to  **(Admin) > Services**, select a Reporting Engine service, and then select  **> View > Config**.
   The Services Config view is displayed with the Reporting Engine General tab open.
2. Select **System Configuration**.

3. Select the checkbox for **Forward Alerts to Respond**.



4. Click **Apply**.
   The Reporting Engine now forwards the alerts to NetWitness Respond.

For details on parameters in the General tab, see the "Reporting Engine General Tab" topic in the *Reporting Engine Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Configure Malware Analysis to Display Malware Analysis Alerts in the Respond View

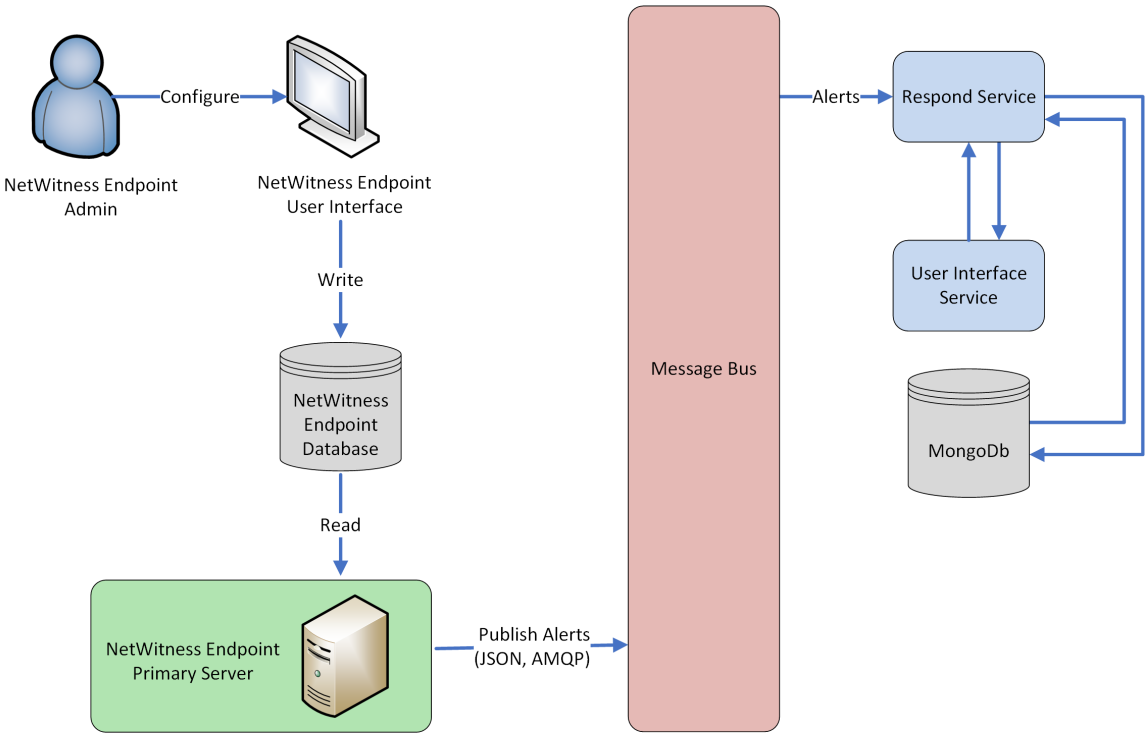
Viewing NetWitness Respond alerts is a function of auditing in Malware Analysis. The procedure for enabling NetWitness Respond alerts is described in the "(Optional) Configure Auditing on Malware Analysis Host" topic in the *Malware Analysis Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Configure NetWitness Endpoint to Display NetWitness Endpoint Alerts in the Respond View

This procedure is required to integrate NetWitness Endpoint with NetWitness so that the NetWitness Endpoint alerts are picked up by the NetWitness Respond component of NetWitness and displayed in the **Respond > Alerts** view.

Note: NetWitness supports NetWitness Endpoint versions 4.3.0.4, 4.3.0.5, 4.4, 4.4.0.2, or later for NetWitness Respond integration. For more detailed information, see "NetWitness Endpoint Integration" in the *NetWitness Endpoint User Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

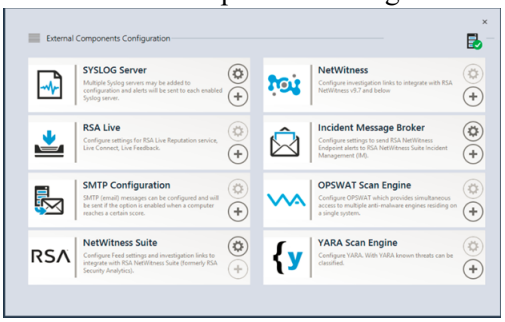
The diagram below represents the flow of NetWitness Endpoint alerts to the NetWitness Respond Server service and its display in the **Respond > Alerts** view.



To configure NetWitness Endpoint to display NetWitness Endpoint alerts in the NetWitness user interface:

1. In the NetWitness Endpoint user interface, click **Configure > Monitoring and External Components**.

   The External Components Configuration dialog is displayed.



2. From the components listed, select **Incident Message Broker** and click + to add a new IM Broker.
3. Enter the following fields:
   a. **Instance Name**: Enter a unique name to identify the IM broker.
   b. **Server Hostname/IP address**: Enter the Host DNS or IP address of the IM Broker (NetWitness Server).
   c. **Port number**: The default port is 5671.
4. Click **Save**.
5. Navigate to the **ConsoleServer.exe.Config** file in **C:\Program Files\RSA\ECAT\Server**.

6. Modify the virtual host configurations in the file as follows:
   `<add key="IMVirtualHost" value="/rsa/system" />`

   Note: In NetWitness 11.0 and later, the virtual host is "/rsa/system". For version 10.6.x and below, the virtual host is "/rsa/sa".

7. Restart the API Server and Console Server.

8. To set up SSL for Respond Alerts, perform the following steps on the NetWitness Endpoint primary console server to set the SSL communications:

   a. Export the NetWitness Endpoint CA certificate to .CER format (Base-64 encoded X.509) from the personal certificate store of the local computer (without selecting the private key).

   b. Generate a client certificate for NetWitness Endpoint using the NetWitness Endpoint CA certificate. (You MUST set the CN name to `ecat`.)

   ```
   makecert -pe -n "CN=ecat" -len 2048 -ss my -sr LocalMachine -a sha1 -sky
   exchange -eku 1.3.6.1.5.5.7.3.2 -in "NWECA" -is MY -ir LocalMachine -sp
   "Microsoft RSA SChannel Cryptographic Provider" -cy end -sy 12
   client.cer
   ```

   Note: In the above code sample, if you upgraded to Endpoint version 4.3 from a previous version and did not generate new certificates, you should substitute `EcatCA` for `NWECA`.

   c. Make a note of the thumbprint of the client certificate generated in step b. Enter the thumbprint value of the client certificate in the IMBrokerClientCertificateThumbprint section of the `ConsoleServer.Exe.Config` file as shown.

   ```
   <add key="IMBrokerClientCertificateThumbprint"
   value="896df0efacf0c976d955d5300ba0073383c83abc"/>
   ```

9. On the NetWitness Server, copy the NetWitness Endpoint CA certificate file in .CER format into the import folder:
   `/etc/pki/nw/trust/import`

10. Issue the following command to initiate the necessary Chef run:
    `orchestration-cli-client --update-admin-node`
    This appends all of those certificates into the truststore.

11. Restart the RabbitMQ server:
    `systemctl restart rabbitmq-server`
    The NetWitness Endpoint account should automatically be available on RabbitMQ.

12. Import the **/etc/pki/nw/ca/nwca-cert.pem** and **/etc/pki/nw/ca/ssca-cert.pem** files from the NetWitness Server and add them to the Trusted Root Certification stores in the Endpoint Server.

Step 2. Assign Respond View Permissions

Add users with the required permissions to investigate incidents and alerts in NetWitness Respond. Users with access to the Respond view need both Incidents and Respond-server permissions. Users with access to configure incident email notification settings need additional Integration-server permissions.

The following pre-configured roles have permissions in the Respond view:

- **Analysts**: The Security Operations Center (SOC) Analysts have access to Alerting, NetWitness Respond, Investigate, and Reporting, but not system configurations.
- **Malware Analysts**: Malware Analysts have access to investigations and malware events.
- **Operators**: Operators have access to configurations, but not Investigate, ESA, Alerting, Reporting and NetWitness Respond.
- **SOC_Managers**: The SOC Managers have the same access as Analysts plus additional permissions to handle incidents and configure NetWitness Respond.
- **Data_Privacy_Officers**: Data Privacy Officers (DPOs) are like Administrators with additional focus on configuration options that manage obfuscation and viewing of sensitive data within the system. See the *Data Privacy Management Guide* for additional information. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
- **Respond_Administrator**: The Respond Administrator has full access to NetWitness Respond.
- **Administrators**: The Administrator has full system access to NetWitness and has all permissions by default.

The NetWitness Respond default permissions are shown in the following tables. You need to assign user permissions from both the **Incidents** and **Respond-server** tabs, which are the Permissions tab names in the  (Admin) > Security view Add or Edit Roles dialogs. You may want to add additional user permissions for Alerting, Context Hub, Investigate, Investigate-server, and Reports.

Caution: It is very important that you assign equivalent user permissions from BOTH the Respond-server tab AND the Incidents tab.

Users who configure incident email notification settings also need permissions in the Integration-server tab.

Respond-server

Permissions	Analysts	SOC Mgrs	DPOs	Respond Admin	Operators	MAs
respond-server.alert.delete			Yes*	Yes*		
respond-server.alert.manage	Yes	Yes	Yes*	Yes*		Yes
respond-server.alert.read	Yes	Yes	Yes*	Yes*		Yes
respond-server.alertrule.manage		Yes	Yes*	Yes*		
respond-server.alertrule.read		Yes	Yes*	Yes*		
respond-server.configuration.manage			Yes*	Yes*		
respond-server.health.read			Yes*	Yes*		
respond-server.incident.delete			Yes*	Yes*		
respond-server.incident.manage	Yes	Yes	Yes*	Yes*		Yes
respond-server.incident.read	Yes	Yes	Yes*	Yes*		Yes
respond-server.journal.manage	Yes	Yes	Yes*	Yes*		Yes
respond-server.journal.read	Yes	Yes	Yes*	Yes*		Yes
respond-server.logs.manage			Yes*	Yes*		
respond-server.metrics.read			Yes*	Yes*		
respond-server.notification.manage (Available in 11.1 and later)		Yes	Yes*	Yes*		
respond-server.notification.read (Available in 11.1 and later)		Yes	Yes*	Yes*		
respond-server.process.manage			Yes*	Yes*		
respond-server.remediation.manage	Yes	Yes	Yes*	Yes*		Yes
respond-server.remediation.read	Yes	Yes	Yes*	Yes*		Yes
respond-server.risk.manage	Yes		Yes*	Yes*		
respond-server.risk.read	Yes		Yes*	Yes*		
respond-server.security.manage			Yes*	Yes*		
respond-server.security.read			Yes*	Yes*		

* Data Privacy Officers and Respond Administrators have the **respond-server.*** permission, which gives them all of the Respond-server permissions.

Incidents

Permissions	Analysts	SOC Mgrs	DPOs	Respond Admin	Operators	MAs
Access Incident Module	Yes	Yes	Yes	Yes		Yes
Configure Incident Management Integration		Yes	Yes	Yes		
Delete Alerts and Incidents			Yes	Yes		
Manage Alert Handling Rules		Yes	Yes	Yes		
View and Manage Incidents	Yes	Yes	Yes	Yes		Yes

The Respond Administrator has all of the Respond-server and Incidents permissions.

Integration-server

Note: The Integration-server permissions are available in NetWitness version 11.1 and later.

Users who configure incident email notification settings also need Integration-server permissions. The following table lists the incident notification permissions in the Integration-server tab assigned to each role.

Permissions	Analysts	SOC Mgrs	DPOs	Respond Admin	Operators	MAs
integration-server.notification.read		Yes	Yes	Yes		
integration-server.notification.manage		Yes	Yes	Yes		

Investigate-server

Users who view Event Analysis in Respond also need Investigate-server permissions. The following table lists the Respond Event Analysis permissions required in the Investigate-server tab and the permissions assigned to each role.

Permissions	Analysts	SOC Mgrs	DPOs	Respond Admin	Operators	MAs
investigate-server.event.read	Yes	Yes	Yes	Yes		Yes
investigate-server.content.reconstruct	Yes	Yes	Yes	Yes		Yes
investigate-server.content.export	Yes	Yes	Yes	Yes		Yes

Incident Email Notification Settings Permissions

Note: Incident email notification setting permissions are available in NetWitness version 11.1 and later.
If you are updating from NetWitness version 11.0 to 11.1 or later, you will need to add additional permissions to your existing built-in NetWitness user roles. For all upgrades to 11.1 or later, you will need to add additional permissions to custom roles.

The following permissions are required for Respond Administrators, Data Privacy Officers, and SOC Managers to access incident email notification settings [ (Configure) > Incident Notifications].

Incidents tab:

- Configure Incident Management Integration

Respond-server tab:

- respond-server.notification.manage
- respond-server.notification.read

Integration-server tab:

- integration-server.notification.read
- integration-server.notification.manage

Respond Event Analysis Permissions

Note: The Event Analysis panel in the Respond view is available in NetWitness version 11.2 and later.

The Events panel in the Respond view, formerly known as the Event Analysis panel, shows the Events view from Investigate for specific indicator events. The following permissions are required to view the Events panel in the Respond view. These permissions are provided by default for users with the Analysts role.

Investigate-server tab:

- investigate-server.event.read
- investigate-server.content.reconstruct
- investigate-server.content.export

Administration tab:

- Access Administration Module

Note: Migrated incidents from NetWitness versions before 11.2 will not show the Events panel in the Respond Incident Details view Indicators panel. Likewise, if you use alerts that were migrated from versions before 11.2 to create incidents in 11.2, you will also not be able to view the Events panel in the Respond view for those incidents.

Respond Saved Filter Permissions

Note: Saved filters for the incidents and alerts lists in Respond are available in NetWitness version 11.5 and later.

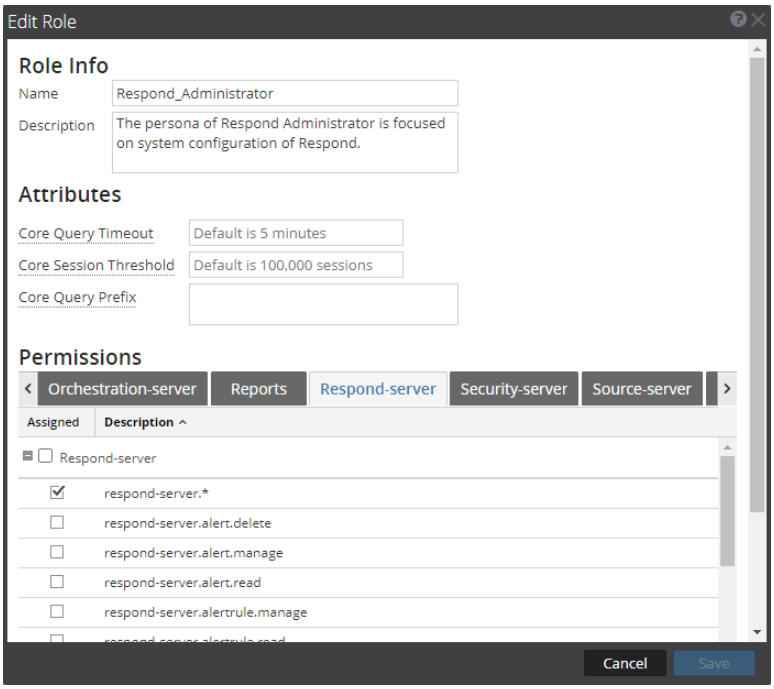
The following permissions are required for the incidents and alerts filters (Respond > Incidents and Respond > Alerts). The Analysts role has the required Respond filter permissions by default.

Respond-server tab:

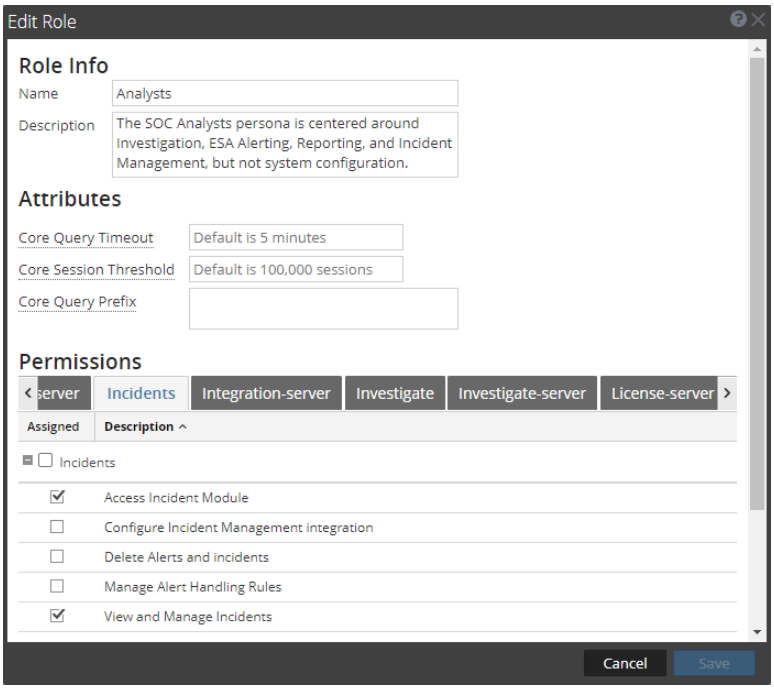
- respond-server.incident.manage
- respond-server.incident.read
- respond-server.alert.manage
- respond-server.alert.read

Respond Role Permission Examples

The following figure shows Respond-server permissions for the default Respond Administrator role. The Respond Administrator role contains all of the NetWitness Respond permissions.



The following figure shows the Incidents permissions for the default Analysts role:



For more information, see "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Restrict Access to Incidents

By default, analysts can view all of the incidents, alerts, and tasks in the Respond view. If you have sensitive or restricted information that should not be shared, you can restrict what analysts and other users can see in the Respond view.

If you restrict access to incidents:

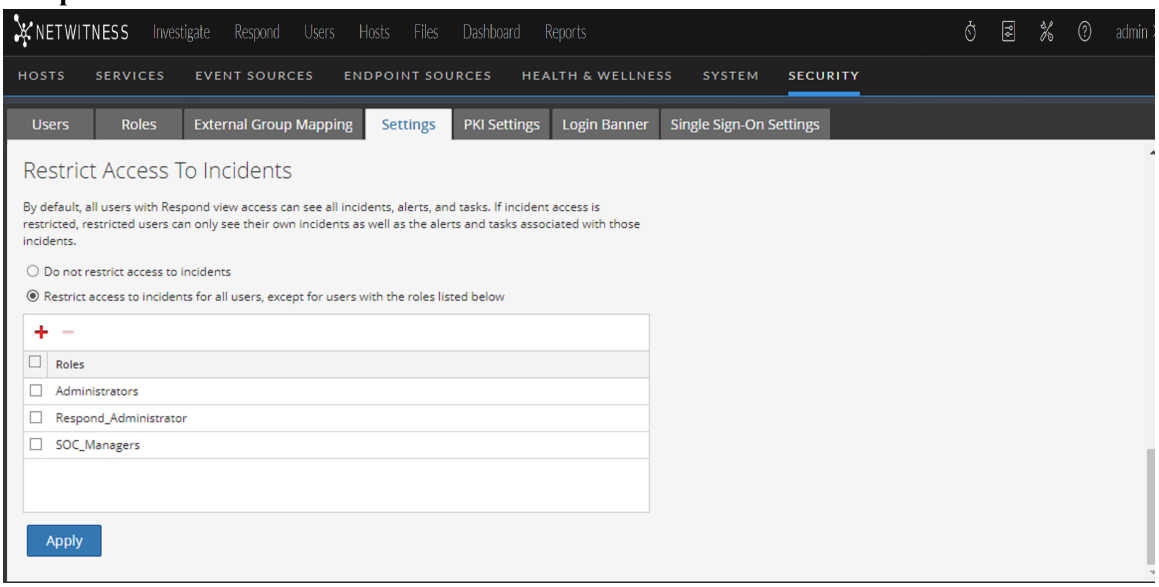
- Analysts can only see incidents assigned to them as well as the alerts and tasks associated with those incidents. Likewise, they can only change the status of and add journal entries (notes) to their own incidents.
- Analysts cannot see the Alerts and Tasks tabs in the Respond view (Respond > Tasks and Respond > Alerts are hidden), so they cannot view all alerts and tasks.
- Analysts cannot see the Assignee button or change the assignee of an incident.
- Analysts cannot see the Related Indicators (alerts) panel (Incident Details view > Find Related tab in the left-side panel).
- When adding events to incidents from the Investigate views, users can only add events to incidents to which they have access. The list of incidents to which users can add events only shows incidents that the user can access.

- When creating incidents from the Investigate views, users must have access to those incidents to view them in the Respond view. For example, when creating incidents from the Investigate view, Analysts must assign the incidents to themselves to view them in the Respond view.

Caution: These restrictions apply to all NetWitness users, except users with the **Administrators**, **Respond_Administrator**, and **SOC_Managers** roles. However, you can adjust the list of user roles whose access to incidents should not be restricted.

To restrict access to incidents:

1. Go to  **(Admin) > Security** and click the **Settings** tab.
2. In the **Restrict Access to Incidents** section, select **Restrict access to incidents for all users, except for users with the roles listed below**.



3. In the list, add the user roles whose access to incidents should not be restricted.
4. Click **Apply**.
   Changes take effect on the next log in to NetWitness.

Step 3. Enable and Create Incident Rules for Alerts

NetWitness Respond incident rules contain criteria to automate the process of creating incidents from alerts. Alerts that meet the rule criteria are grouped together to form an incident. Analysts use these incidents to locate indicators of compromise. Instead of creating an incident for a particular set of alerts and adding the alerts to that incident manually, you can save time by using incident rules to create incidents from alerts for you.

NetWitness provides predefined incident rules that you can use and you can also create your own rules based on your business requirements.

To create incidents automatically, you need to enable at least one incident rule.

When you have two or more incident rules enabled, the order of the rules becomes very important. The highest priority rules are at the top of the Incident Rules list. The highest priority rule has the number 1 in the Order field. The next highest priority rule is number 2 in the Order field, and so on. Alerts can only be part of one incident. If an alert matches more than one rule in the Incident Rule list, it is only evaluated using the highest priority rule that it matches.

In 11.6.1, the Incident rule execution pattern is modified such that the alert aggregation queries refer to the lastRun parameter of the incident rule. So, now multiple incident rules (positioned as per the priority in the order field) in the Incident rules list match the same alert name.

For example,

- INC-rule 1 is at the 4th position in the order field that is associated with the conditions alert.name equal to test and **source** equal to **ESA**.
- INC-rule 2 is at the 24th position in the order field that is associated with the condition **source** equal to **ESA**.

  Before 11.6.1 upgrade, the INC-rule 1 matched the alert name test. On upgrade, the INC-rule 2 at the 24th position in the order field matches the alert name test as the source is ESA.

  To address this scenario, configure each incident rule to have unique conditions. Make sure the conditions in the lower priority incident rules are not duplicate of the conditions in the highest priority rules.

NetWitness has 13 predefined incident rules that you can use. To set up your incident rules, you can do any of the following:

- Enable predefined incident rules
- Add new rules
- Clone rules
- Edit existing rules
- Export and import rules

The Detect AI default incident rule is available in NetWitness 11.6 and later. It captures the network user behavior and uses the deployed RSA Live ESA Rules to create incidents from alerts.

The User Entity Behavior Analytics incident rule is available in 11.3 and later. It captures user entity behavior grouped by Classifier ID to create incidents from alerts. The User Behavior default incident rule is available in NetWitness 11.1 and later. It captures network user behavior and uses deployed RSA Live ESA Rules to create incidents from alerts.

You can select and deploy the RSA Live ESA Rules that you want to monitor. For more information, see Deploy the RSA Live ESA Rules.

To configure the default incident rules or verify your existing default incident rules with the 11.5 default incident rules, see Set Up and Verify Default Incident Rules.

This topic contains the following procedures:

- Enable Incident Rules
- Create an Incident Rule
- Verify the Order of Your Incident Rules
- Clone an Incident Rule
- Edit an Incident Rule
- Export Incident Rules
- Import Incident Rules

Enable Incident Rules

To create incidents automatically, you need to enable at least one incident rule. Predefined (default) incident rules or rules that you create must be enabled before they start creating incidents.

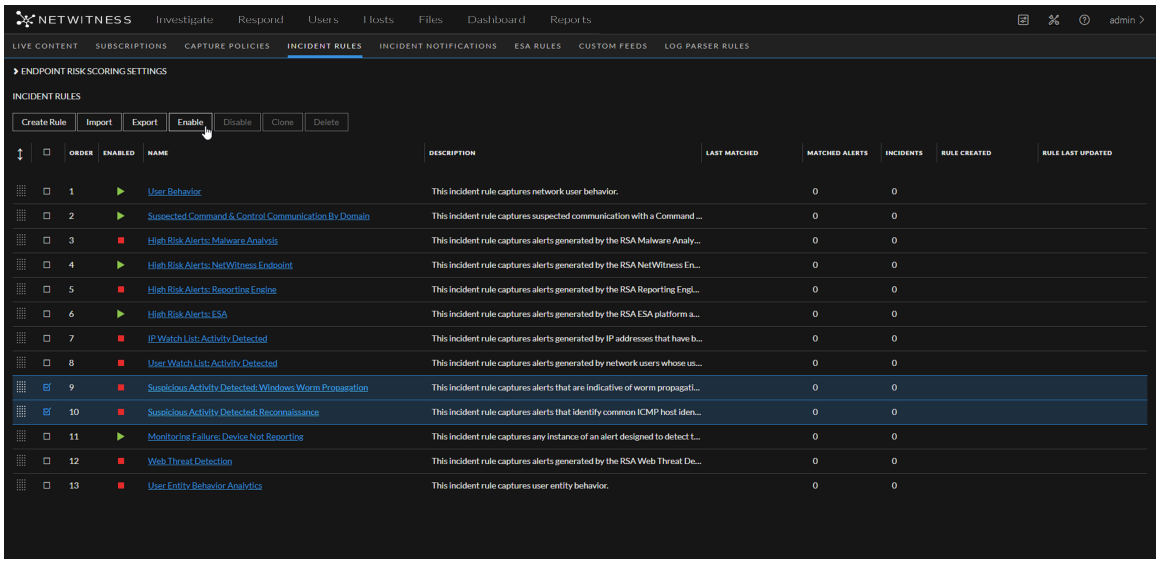
To enable one or more incident rules:

Note: Enabling one or more incident rules from the Incident Rules view is only available in NetWitness version 11.4 and later.

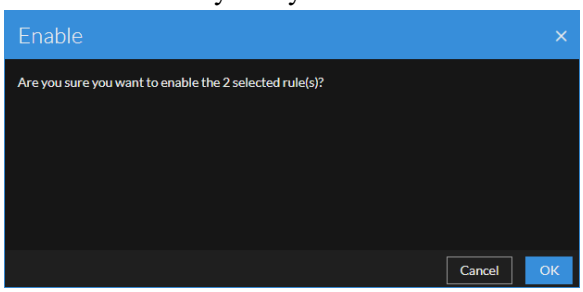
This is the easiest way to enable rules. Use this method after you have made the necessary adjustments to the rules and you just want to quickly enable them.

1. Go to  **(Configure) > Incident Rules**.
   The Incident Rules view is displayed.

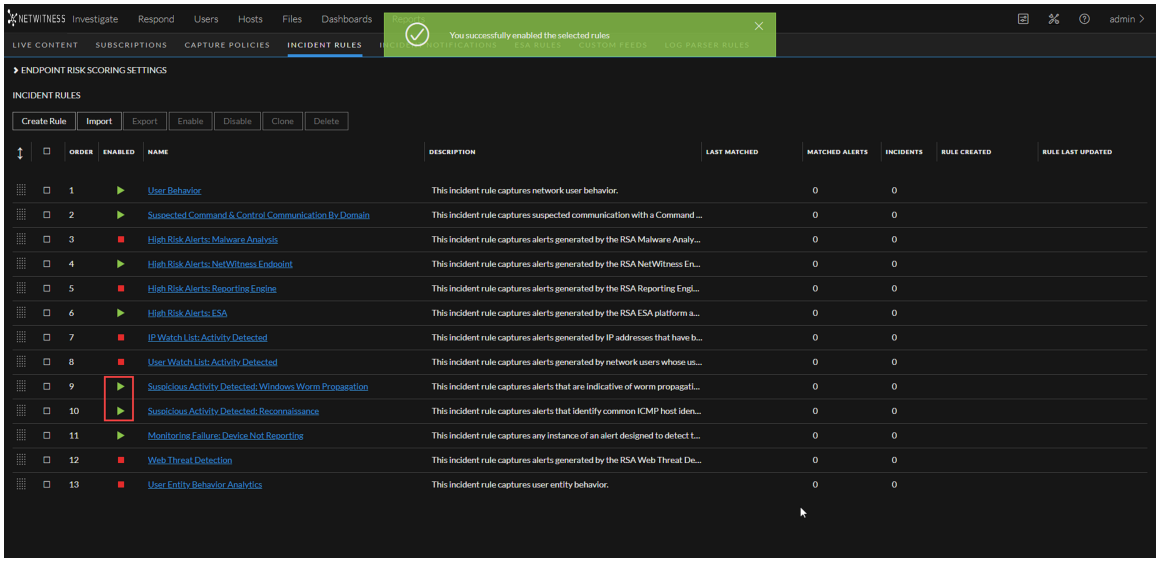
2. Select one or more incident rules and click **Enable**.



3. Click **OK** to verify that you want to enable the selected rules.



In the Incident Rules view, the Enabled column changes from a red square ■ (Disabled) to a green triangle ▶ (Enabled).



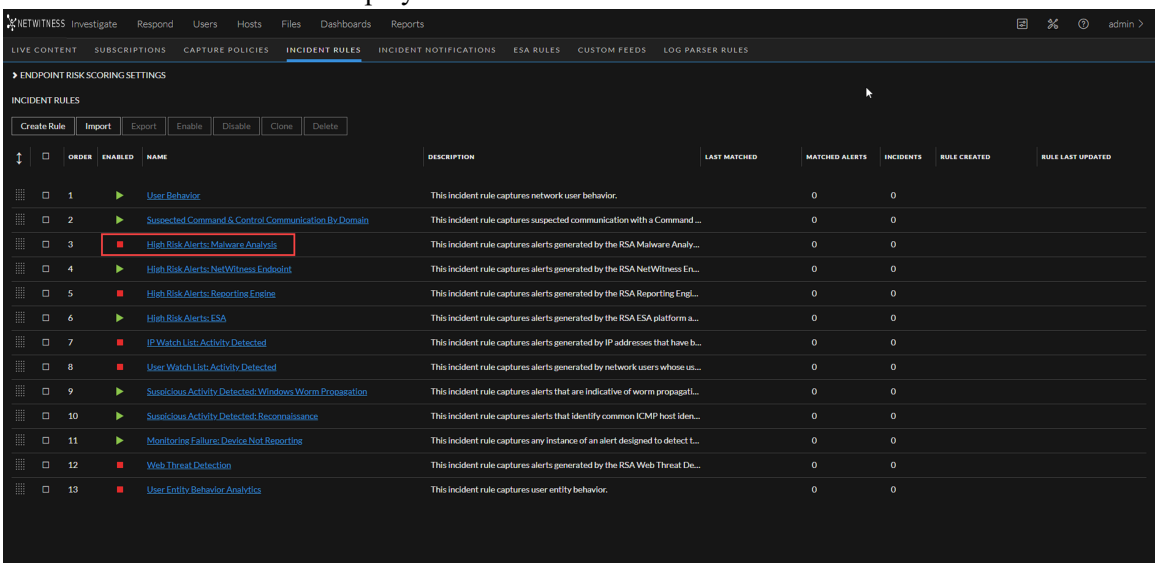
4. Verify the order of your incident rules.

Note: To disable incident rules, follow the above procedure but select the Disable button instead of the Enable button.

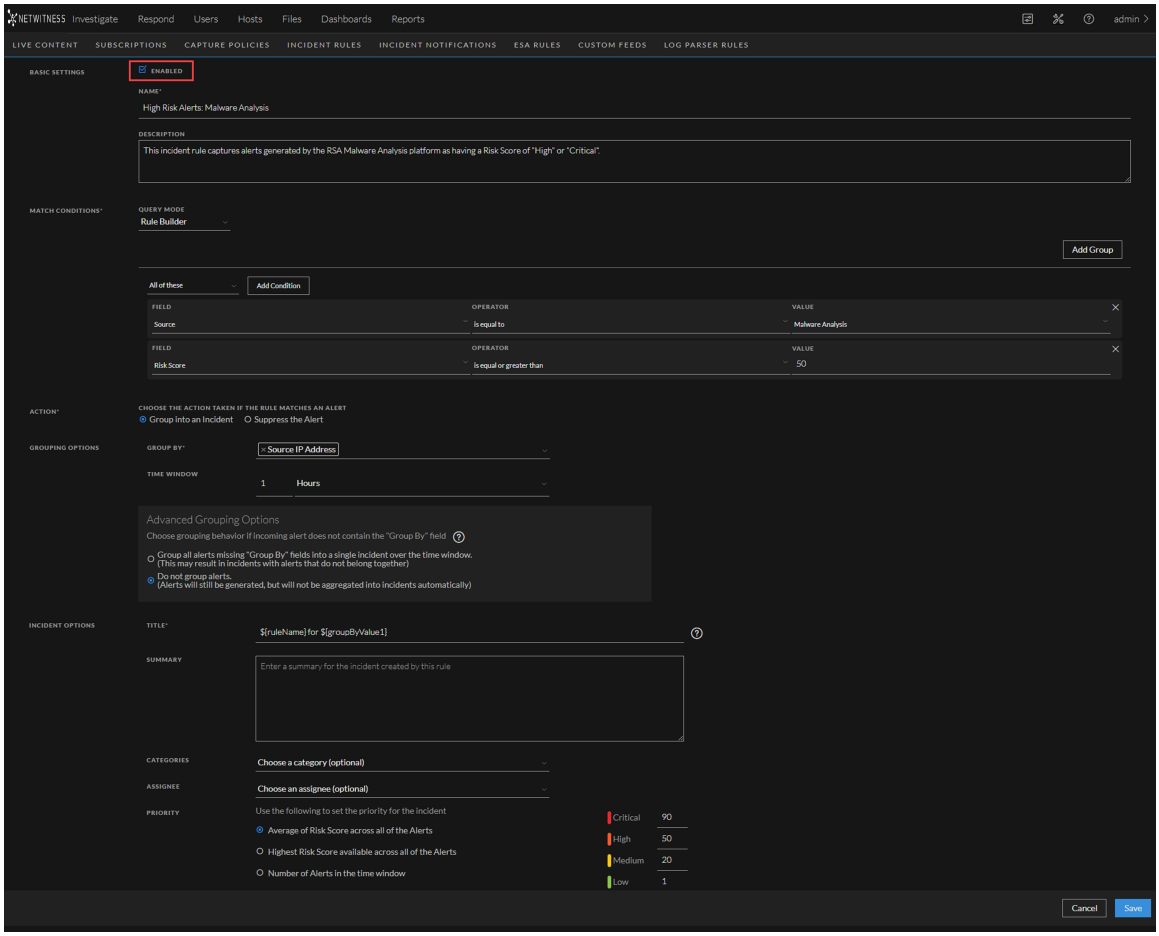
To enable an incident rule from within the incident rule details:

You can enable rules from within the incident rule details when you save your rule adjustments.

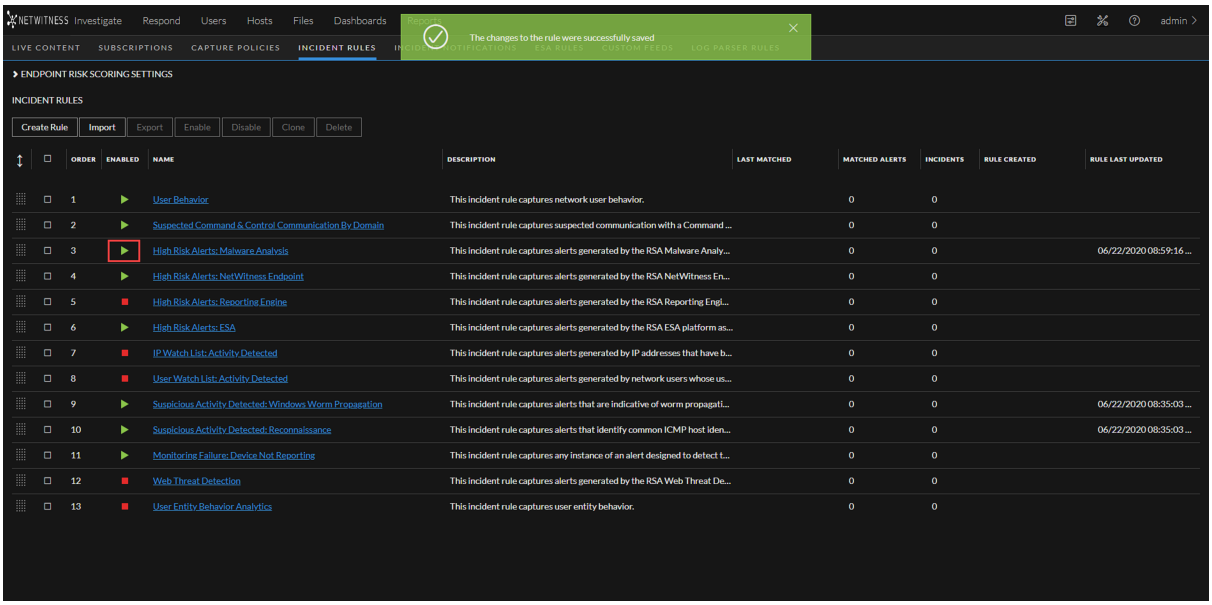
1. Go to **(Configure) > Incident Rules**.
   The Incident Rules view is displayed.



2. Click the link in the **Name** column for the rule that you want to enable.
   The Incident Rule Details view is displayed for the selected rule.



3. Adjust the parameters and conditions of your rule as required. For details about various parameters that can be set as criteria for an incident rule, see Incident Rule Details View. To adjust the default rules, see Set Up and Verify Default Incident Rules.
4. In the Basic Settings section, select the **Enabled** checkbox.
5. Click **Save** to enable the rule.
   Notice that the Enabled column changes from a red square (Disabled) to a green triangle (Enabled).



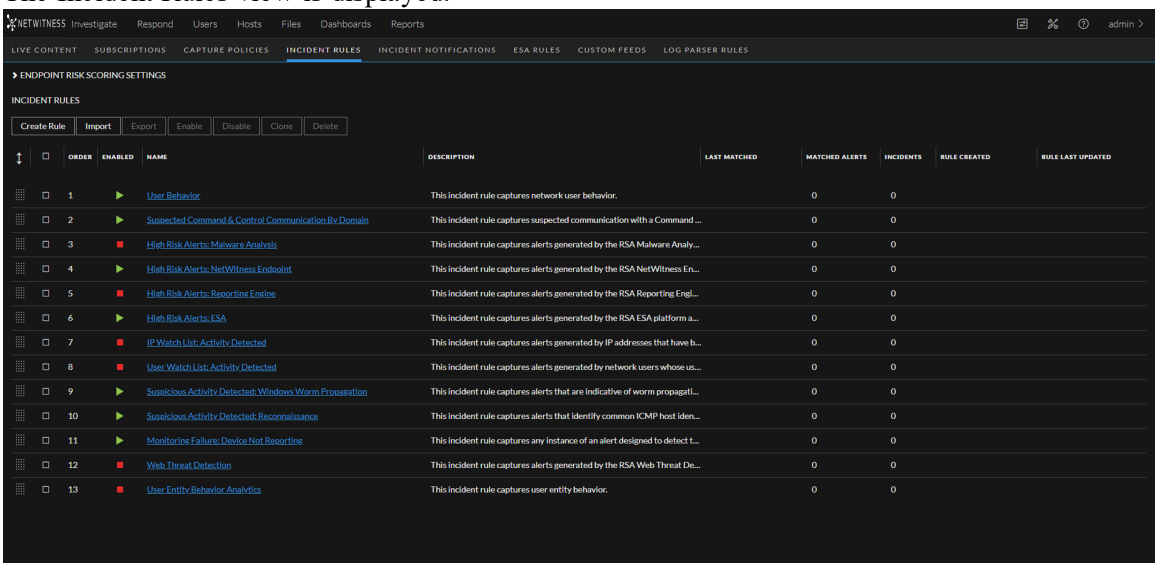
6. Verify the order of your incident rules.

Note: To disable an incident rule in the Incident Rule Details view, follow the above procedure but clear the Enabled checkbox instead of selecting it.

Create an Incident Rule

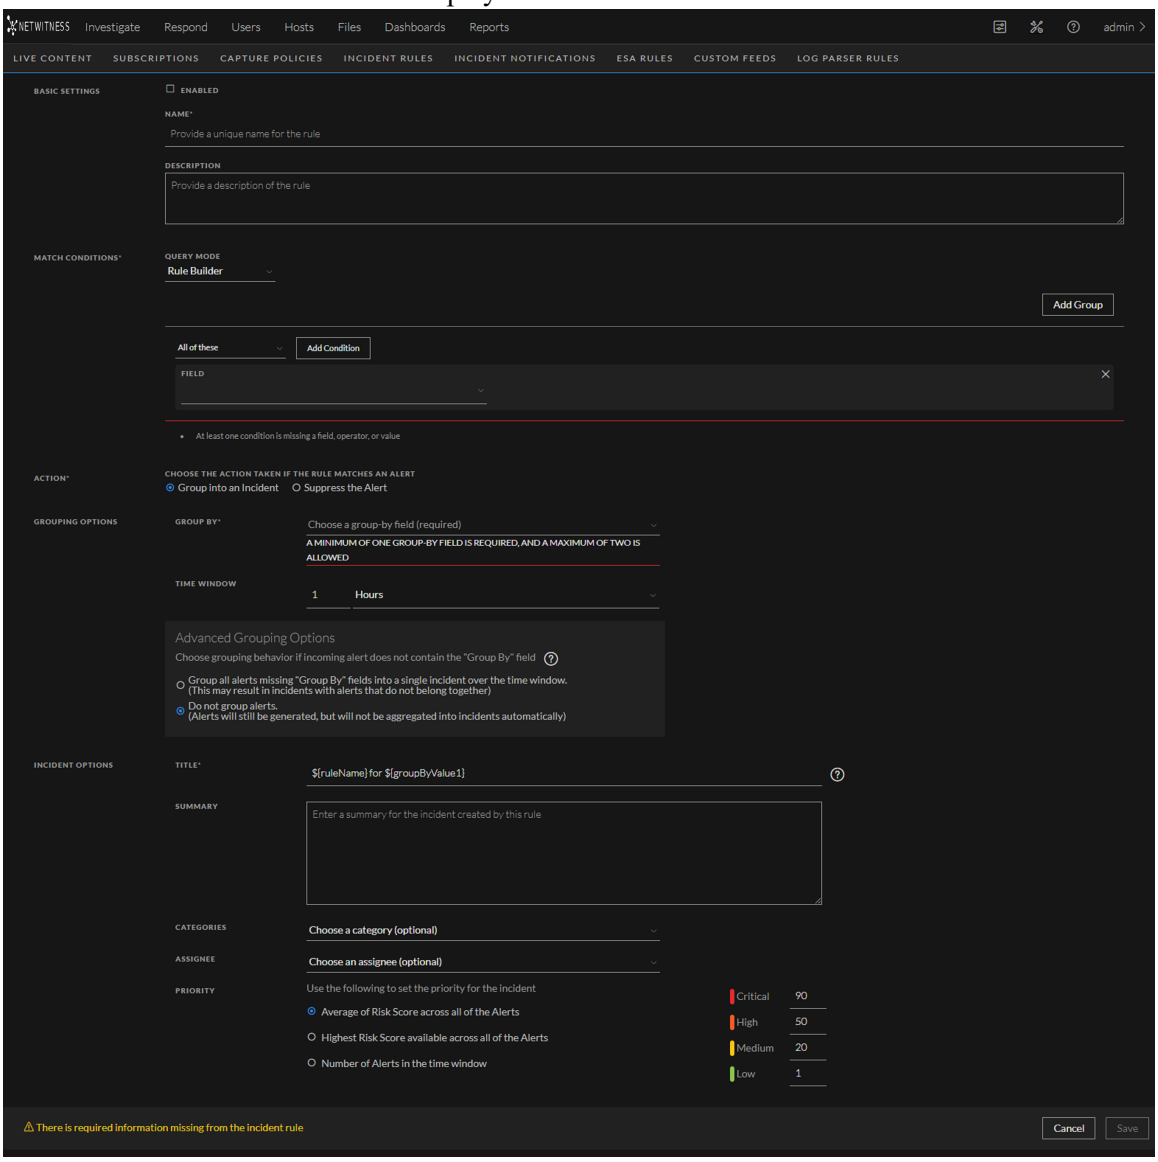
1. Go to **(Configure) > Incident Rules**.

   The Incident Rules view is displayed.



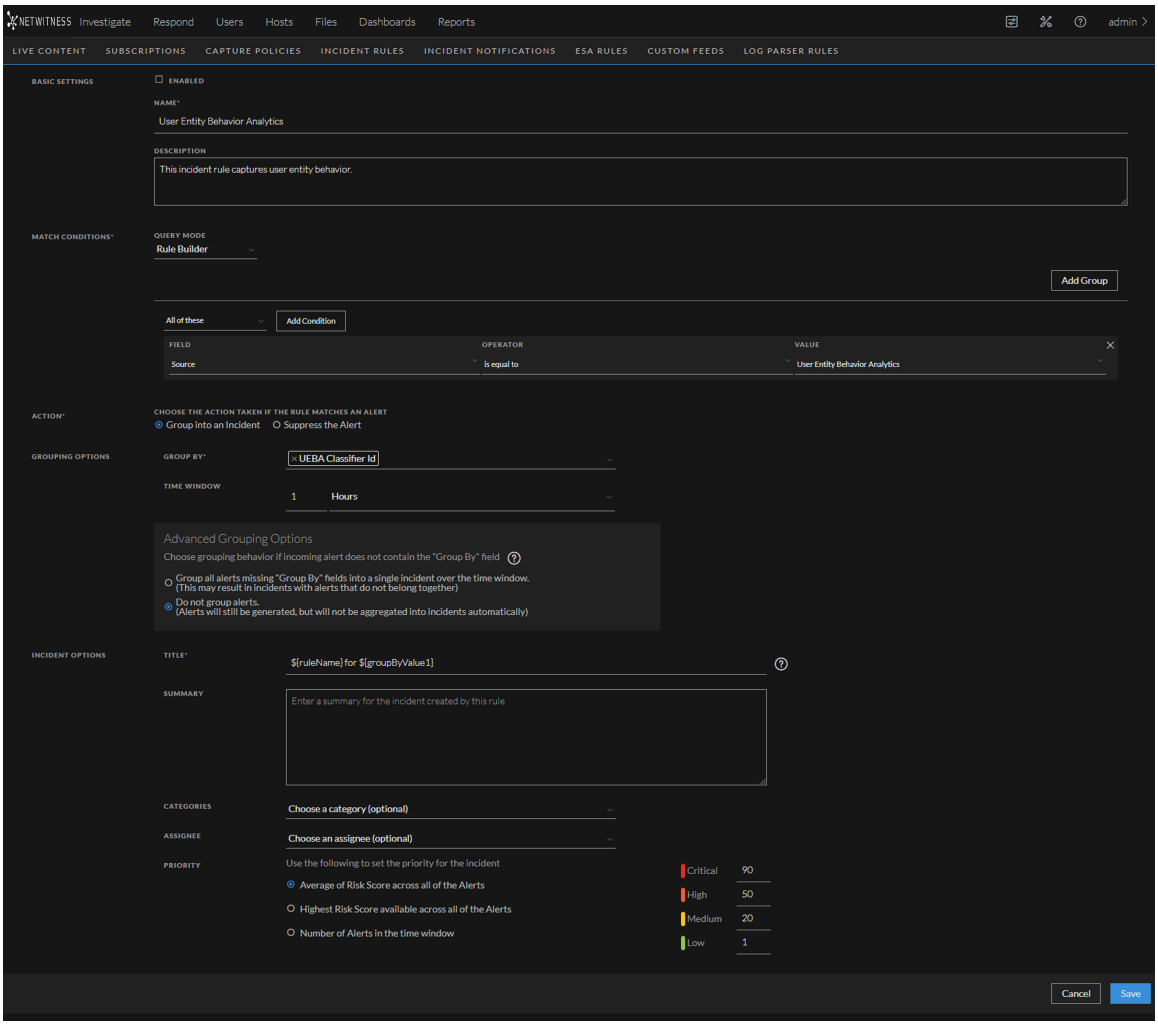
2. To add a new rule, click **Create Rule**.

The Incident Rule Details view is displayed.



3. Enter the parameters and conditions of your rule. All rules need to have at least one condition. For details about parameters that can be set as criteria for an incident rule, see Incident Rule Details View.

   The following figure shows a rule example. It is the default User Entity Behavior Analytics incident rule. It captures user entity behavior grouped by Classifier ID to create incidents from alerts.



4. If you are ready to enable your rule, in the Basic Settings section, select **Enabled**.
5. Click **Save**.

   The rule appears in the Incidents Rules list. If you selected Enabled, the rule is enabled and it starts creating incidents depending on the incoming alerts that match the selected criteria.
6. Verify the order of your incident rules.

Verify the Order of Your Incident Rules

NetWitness Respond evaluates incoming alerts against the incident rules in the order that you define. If alerts match the first rule listed, then that rule creates an incident. If alerts match the second rule listed and those alerts did not match the first rule, then the second rule creates an incident. If alerts match the third rule listed and those alerts did not match the first or second rule listed, then the third rule creates an incident, and so on.

To change the order of the rules, use the drag pads () in front of the rules to move them up and down in the list.
The rule order determines which rule takes effect if the criteria for multiple rules match the same alert. If multiple rules match an alert, only the rule with the highest priority creates an incident.

Clone an Incident Rule

It is often easier to duplicate an existing rule that is similar to a rule that you want to create and adjust it accordingly.

1. Go to  **(Configure) > Incident Rules**.
   The Incident Rules view is displayed.
2. Select the rule that you would like to copy and click **Clone**.
3. Adjust the parameters and conditions of your rule as required. All rules need to have at least one condition.
4. If you are ready to enable your rule, in the Basic Settings section, select **Enabled**.
5. Click **Save** to create the rule.
6. Verify the order of your incident rules.

Edit an Incident Rule

1. Go to  **(Configure) > Incident Rules** and click the link in the **Name** column for the rule that you want to update.
   The Incident Rule Details view is displayed.
2. Adjust the parameters and conditions of your rule as required. All rules need to have at least one condition.
3. If you are ready to enable your rule, in the Basic Settings section, select **Enabled**.
4. Click **Save** to update the rule.
5. Verify the order of your incident rules.

See Also:

- For details about parameters that can be set as criteria for an incident rule, see Incident Rule Details View.

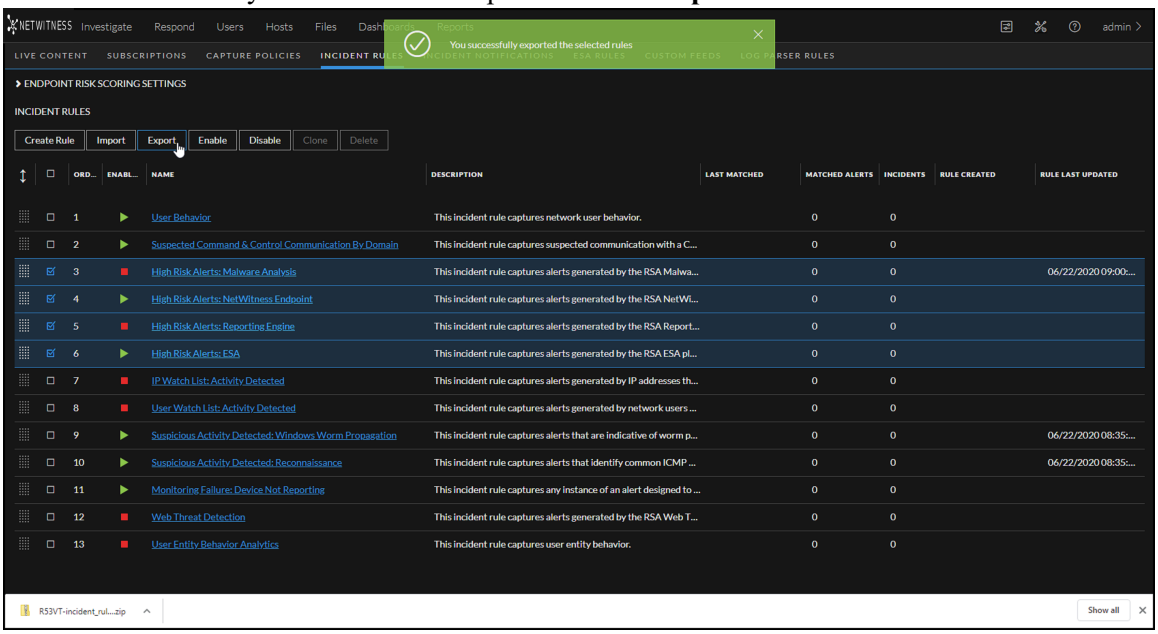
- For details on the parameter and field descriptions in the Incident Rules list, see Incident Rules View.

Export Incident Rules

Note: Exporting and importing incident rules from the Incident Rules view is only available in NetWitness version 11.4 and later.

Exporting incident rules enables you to share incident rules with other NetWitness Servers on the same release version. The exported incident rules file is a ZIP file that contains two JSON files: one file contains the incident rules and the other file contains the incident rule schema. You cannot export Advanced incident rules; the export function only allows incident rules created using Rule Builder.

1. Go to **(Configure) > Incident Rules**.
   The Incident Rules view is displayed.
2. Select the rules that you would like to export and click **Export**.



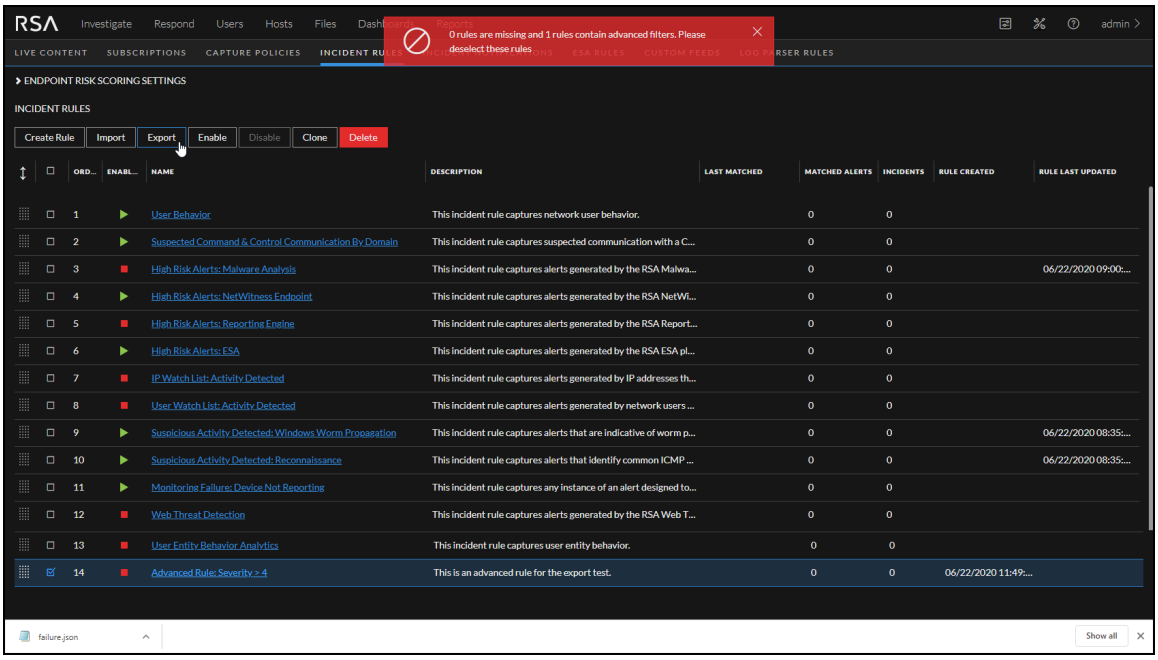
The exported incident rules file is a ZIP file in the format `<random ID>-incident_rules_export.json.zip`, which contains two mandatory JSON files:

- `aggregation_rule_schema.json` contains the incident rule schema.
- `<random ID>-incident_rules_export.json` contains the incident rules.

Note: You cannot export Advanced rules.

You can import this ZIP file on another NetWitness Server on the same release version.

If for some reason the export is not successful, and you receive only a .JSON file, for example, `failure.json`, refresh your browser and try again. This could happen if someone made an adjustment to the incident rules at the same time. You can also receive an error if you attempt to export an Advanced incident rule, which is not allowed.



Import Incident Rules

Note: Exporting and importing incident rules from the Incident Rules view is only available in NetWitness version 11.4 and later.

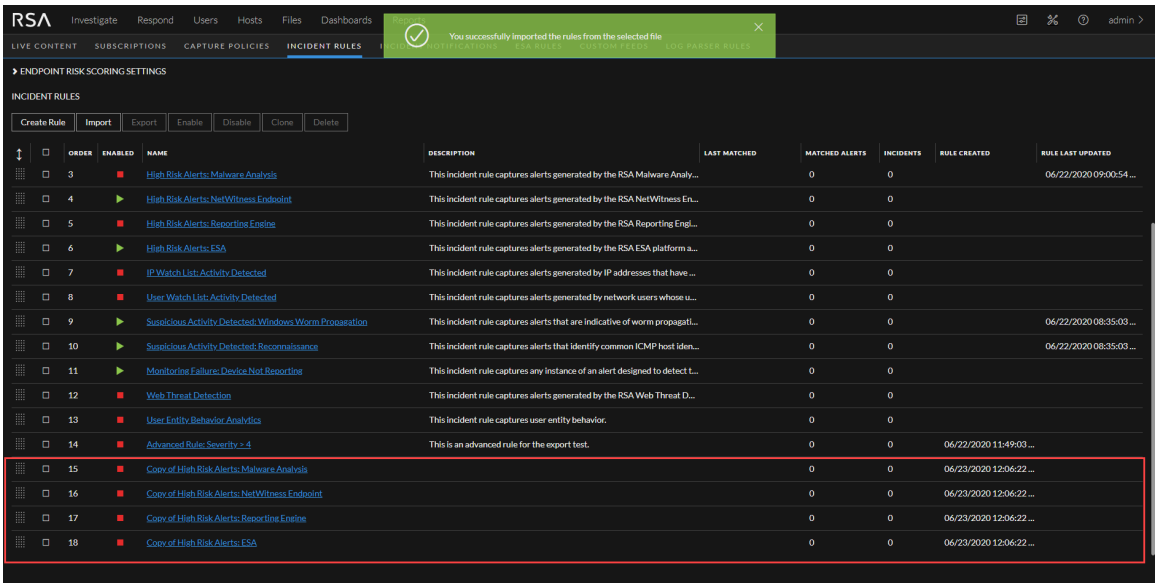
You can import an incident rules ZIP file from NetWitness Servers on the same release version. The incident rules ZIP file must be in the original exported format `<random ID>-incident_rules_export.json.zip` and contain two mandatory JSON files:

- `aggregation_rule_schema.json` contains the incident rule schema.
- `<random ID>-incident_rules_export.json` contains the incident rules.

The import fails if the ZIP file contains additional files or folders. To edit the incident rules ZIP file, see Edit the Incident Rules Export ZIP File.

To import incident rules:

1. Go to **(Configure) > Incident Rules**.
   The Incident Rules view is displayed.
2. Click **Import** and select the incident rules ZIP file to import.
   If the import is successful, a successful import notification is displayed, and the imported incident rules are disabled and shown at the bottom of the incident rules list. The **Rule Created** column shows the date and time of the import.



See Also:

- For details about parameters that can be set as criteria for an incident rule, see Incident Rule Details View.
- For details on the parameter and field descriptions in the Incident Rules list, see Incident Rules View.

Additional Procedures for Respond Configuration

Use this section when you are looking for instructions to perform a specific task after the initial setup of NetWitness Respond.

- Set Up and Verify Default Incident Rules
- Configure Risk Scoring Settings for Automated Incident Creation
- Configure Analyst UI for Respond Server Alert Normalization
- Configure Incident Email Notification Settings
- Set a Retention Period for Alerts and Incidents
- Obfuscate Private Data
- Manage Incidents in Archer Cyber Incident & Breach Response
- Configure the Option to Send Incidents to Archer
- Configure Threat Aware Authentication
- Set a Counter for Matched Alerts and Incidents
- Configure a Database for the Respond Server Service

Set Up and Verify Default Incident Rules

The User Entity Behavior Analytics default incident rule is available in NetWitness 11.3 and later. It captures user entity behavior grouped by Classifier ID to create incidents from alerts.

The Detect AI default incident rule is available in NetWitness 11.6 and later. It captures the anomalies generated by Detect AI.

The User Behavior incident rule, which captures network user behavior, is available in NetWitness 11.1 and later. This rule uses deployed RSA Live ESA Rules to create incidents from alerts. You can select and deploy the RSA Live ESA Rules that you want to monitor.

The following default incident rules changed slightly for 11.1 and later and now have **Source IP Address** as the Group By value:

- High Risk Alerts: Reporting Engine
- High Risk Alerts: Malware Analysis
- High Risk Alerts: ESA

The following default incident rule changed slightly for 11.3 and later and now has the **Host Name** as the Group By value:

- High Risk Alerts: NetWitness Endpoint*

*If you have NetWitness Endpoint, the High Risk Alerts: NetWitness Endpoint default incident rule captures alerts generated by NetWitness Endpoint with a risk score of High or Critical. To aggregate NetWitness Endpoint alerts based on the File Hash instead of Host Name, create another NetWitness Endpoint Rule using the File Hash as the Group By value. See Create a NetWitness Endpoint Incident Rule using File Hash for step-by-step instructions.

To verify your existing default incident rules with the 11.5 default incident rules, look at the default incident rule tables following these procedures. If you are missing a default incident rule, you can create it manually. Review the default incident rules and adjust them to your environment as required.

Set Up the User Behavior Incident Rule

In order to use the default User Behavior incident rule, you need to deploy the RSA Live ESA Rules that you want to monitor from those listed in the User Behavior incident rule conditions. Complete the following procedures to start aggregating alerts for the User Behavior default incident rule:

- Deploy the RSA Live ESA Rules
- Adjust and enable the User Behavior default rule (or create it if you do not have it)

Deploy the RSA Live ESA Rules

1. Go to  **(Configure) > Live Content**.
2. In the **Resource Types** field, select **Event Steam Analysis Rule** and click **Search**.

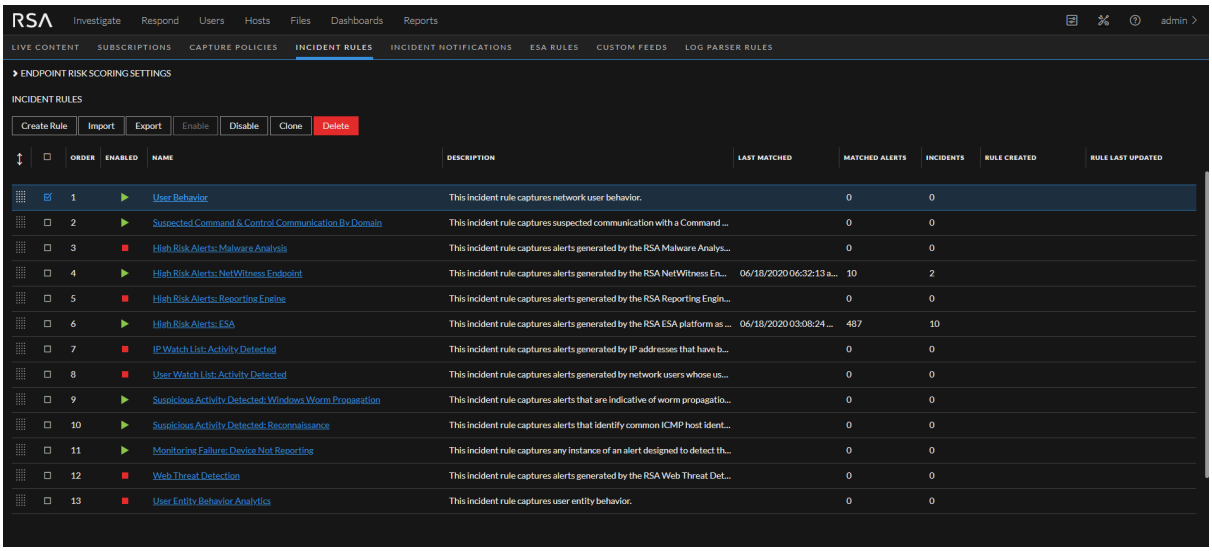
3. In the **Matching Resources** list, select the ESA Rules from the following **User Behavior** table that you are interested in monitoring and deploy them (click **Deploy**).
4. Go to **(Configure) > ESA Rules > Rules** tab, and in the Rule Library **Filter** drop-down list, select **RSA Live ESA Rule**.
5. To add a new ESA rule deployment, in the drop-down list near **Deployments**, click **Add**.
   a. In the ESA Services section, add and then select your ESA service.
   b. In the Data Sources section, click + and add a data source to use for the ESA rule deployment.
   c. In the ESA Rules section, click + and in the Deploy ESA Rules dialog, select the ESA Rules that you selected from the **User Behavior** table, and then click **Save**.
   The selected ESA rules are listed with a status of **Added**.
6. Select the ESA rules that you added from the previous step, and click **Deploy Now**.
   The status of the selected ESA rules changes to **Deployed**.
7. Go to **(Configure) > ESA Rules > Services** tab.
   In the **Deployed Rule Stats** for your ESA service, the rules that you added should have a status of enabled, which is indicated by a green circle in the Enable column.

Adjust and Enable the User Behavior Default Rule (or Create It If You Do Not Have It)

If you have the User Behavior default rule, you can adjust it for your environment and enable it. If you do not have the User Behavior default rule, you can create it manually.

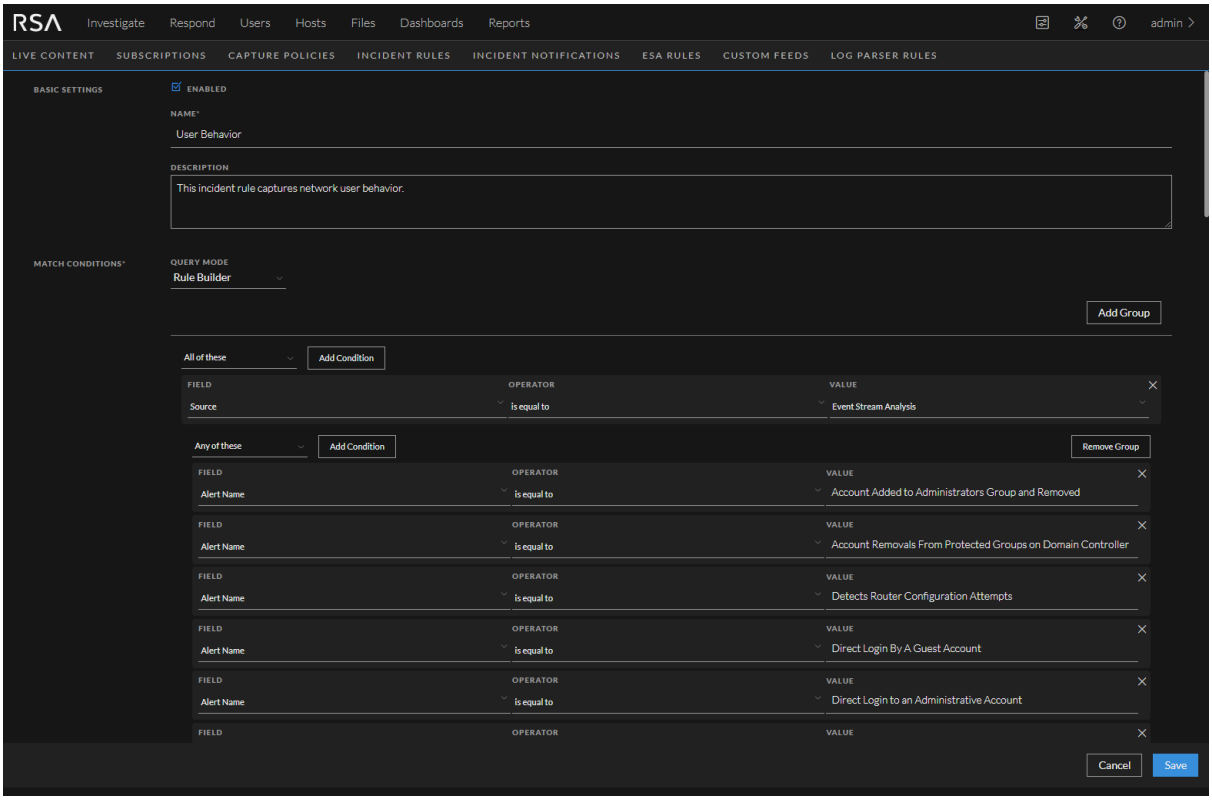
(Optional) To create the User Behavior default rule:

1. Go to **(Configure) > Incident Rules.**
   The Incident Rules view is displayed. (The following figure shows what the User Behavior rule looks like if it was there.)



2. Click **Create Rule** and in the Incident Rule Details view, create the User Behavior default incident rule using the values in the User Behavior table following this procedure. The conditions as well as the values not listed in the table should be set for your business requirements. For details about various parameters that can be set as criteria for an incident rule, see Incident Rule Details View.

   The following figure shows a portion of the User Behavior default rule details. Notice that there are two groups in this rule.



3. If you are ready to enable your rule, in the Basic Settings section, select **Enabled**.

4. Click **Save**.
   The rule appears in the Incidents Rules list. If you selected Enabled, the rule is enabled and it starts creating incidents depending on the incoming alerts that are matched as per the rule criteria.
5. Verify the order of your incident rules. For more information, see Verify the Order of Your Incident Rules.

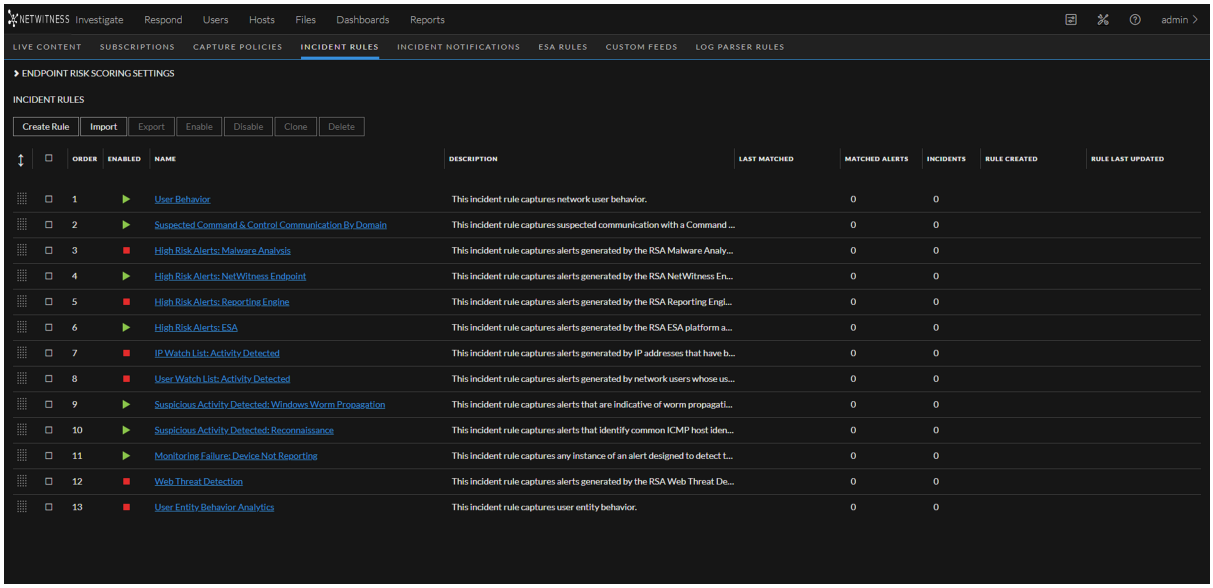
The following table shows the values for the User Behavior default incident rule.

Field	Condition Field	Condition Operator	Value
Name			User Behavior
Description			This incident rule captures network user behavior.
Query Mode:			Rule Builder **Note:** For information about advanced query mode, see Incident Rule Details View
1st Group:			All of these
Condition:	Source	is equal to	Event Stream Analysis
2nd Group:			Any of these
Conditions:	Alert Name	is equal to	Account Added to Administrators Group and Removed
	Alert Name	is equal to	Account Removals From Protected Groups on Domain Controller
	Alert Name	is equal to	Detects Router Configuration Attempts
	Alert Name	is equal to	Direct Login By A Guest Account
	Alert Name	is equal to	Direct Login to an Administrative Account
	Alert Name	is equal to	Failed Logins Followed By Successful Login Password Change
	Alert Name	is equal to	Insider Threat Mass Audit Clearing
	Alert Name	is equal to	Internal Data Posting to 3rd Party Sites
	Alert Name	is equal to	kbrtgt Account Modified on Domain controller
	Alert Name	is equal to	Lateral Movement Suspected Windows
	Alert Name	is equal to	Logins across Multiple Servers
	Alert Name	is equal to	Logins by Same User to Multiple Servers
	Alert Name	is equal to	Malicious Account Creation Followed by Failed Authorization

Field	Condition Field	Condition Operator	Value
	Alert Name	is equal to	Multiple Account Lockouts From Same or Different Users
	Alert Name	is equal to	Multiple Failed Logins Followed By a Successful Login
	Alert Name	is equal to	Multiple Failed Logins from Same User Originating from Different Countries
	Alert Name	is equal to	Multiple Failed Privilege Escalations by Same User
	Alert Name	is equal to	Multiple Intrusion Scan Events from Same User to Unique Destinations
	Alert Name	is equal to	Multiple Login Failures by Administrators to Domain Controller
	Alert Name	is equal to	Multiple Login Failures by Guest to Domain Controller
	Alert Name	is equal to	Multiple Failed Logons from Same Source IP with Unique Usernames
	Alert Name	is equal to	Multiple Successful Logins from Multiple Diff Src to Diff Dest
	Alert Name	is equal to	Multiple Successful Logins from Multiple Diff Src to Same Dest
	Alert Name	is equal to	Privilege Escalation Detected
	Alert Name	is equal to	Privilege Escalation Detected in Unix
	Alert Name	is equal to	Privilege User Account Password Change
	Alert Name	is equal to	Failed Logins Outside Business Hours
	Alert Name	is equal to	DNS Tunneling
	Alert Name	is equal to	User Login Baseline
Group By			Destination User Account
Time Window			1 Hour
Title			${ruleName} for ${groupByValue1}

Set up or Verify a Default Incident Rule

1. Go to  **(Configure) > Incident Rules.**
   The Incident Rules view is displayed.



2. Click the link in the **Name** field of a default incident rule to view the Incident Rule Details view. Set up or verify the default incident rule using the values in the default incident rules tables in this topic. Values not listed in the tables should be set for your business requirements. For details about various parameters that can be set as criteria for an incident rule, see Incident Rule Details View.
3. When you are ready to enable your rule, in the Basic Settings section, select **Enabled**.
4. Click **Save**.
5. Verify the order of your incident rules. For more information, see Verify the Order of Your Incident Rules.

Suspected Command & Control Communication By Domain

The following table shows the values for the Suspected Command & Control Communication By Domain default incident rule.

Field	Condition Field	Condition Operator	Value
Name			Suspected Command & Control Communication By Domain
Description			This incident rule captures suspected communication with a Command & Control server and groups results by domain.
Group:			All of these
Conditions:	Source	is equal to	Event Stream Analysis
	Alert Rule Id	is equal to	Suspected C&C
Group By			Domain for Suspected C& C
Time Window			7 Days
Title			Suspected C&C with ${groupByValue1}
Summary			NetWitness Platform detected communications with ${groupByValue1} that may be command and control malware. 1. Evaluate if the domain is legitimate (online radio, news feed, partner, automated testing, etc.). 2. Review the domain registration for suspect information (Registrant country, registrar, no registration data found, etc). 3. If the domain is suspect, go to the Investigation module to locate other activity to or from it.

High Risk Alerts: Malware Analysis

The following table shows the values for the High Risk Alerts: Malware Analysis default incident rule.

Field	Condition Field	Condition Operator	Value
Name			High Risk Alerts: Malware Analysis
Description			This incident rule captures alerts generated by the NetWitness Malware Analysis platform as having a Risk Score of "High" or "Critical".
Group:			All of these
Conditions:	Source	is equal to	Malware Analysis
	Risk Score	is equal or greater than	50
Group By			Source IP Address
Time Window			1 Hour
Title			${ruleName} for ${groupByValue1}

High Risk Alerts: NetWitness Endpoint

The following table shows the values for the High Risk Alerts: NetWitness Endpoint default incident rule.

Field	Condition Field	Condition Operator	Value
Name			High Risk Alerts: NetWitness Endpoint
Description			This incident rule captures alerts generated by the NetWitness Endpoint platform as having a Risk Score of "High" or "Critical".
Group:			All of these
Conditions:	Source	is equal to	NetWitness Endpoint
	Risk Score	is equal or greater than	50
Group By			Host Name*

Field	Condition Field	Condition Operator	Value
Time Window			1 Hour
Title			${ruleName} for ${groupByValue1}

*To aggregate NetWitness Endpoint alerts based on the File Hash, create another NetWitness Endpoint Rule using the File Hash as the Group By value. See Create a NetWitness Endpoint Incident Rule using File Hash for step-by-step instructions.

High Risk Alerts: Reporting Engine

The following table shows the values for the High Risk Alerts: Reporting Engine default incident rule.

Field	Condition Field	Condition Operator	Value
Name			High Risk Alerts: Reporting Engine
Description			This incident rule captures alerts generated by the NetWitness Reporting Engine as having a Risk Score of "High" or "Critical".
Group:			All of these
Conditions:	Source	is equal to	Reporting Engine
	Risk Score	is equal or greater than	50
Group By			Source IP Address
Time Window			1 Hour
Title			${ruleName} for ${groupByValue1}

High Risk Alerts: ESA

The following table shows the values for the High Risk Alerts: ESA default incident rule.

Field	Condition Field	Condition Operator	Value
Name			High Risk Alerts: ESA
Description			This incident rule captures alerts generated by the NetWitness ESA platform as having a Risk Score of "High" or "Critical".
Group:			All of these

Field	Condition Field	Condition Operator	Value
Conditions:	Source	is equal to	Event Stream Analysis
	Risk Score	is equal or greater than	50
Group By			Source IP Address
Time Window			1 Hour
Title			${ruleName} for ${groupByValue1}

IP Watch List: Activity Detected

The following table shows the values for the IP Watch List: Activity Detected default incident rule.

Field	Condition Field	Condition Operator	Value
Name			IP Watch List: Activity Detected
Description			This incident rule captures alerts generated by IP addresses that have been added as "Source IP Address" *and* "Destination IP Address" conditions of the rule. To add additional IP addresses to the watch list, simply add a new Source and Destination IP Address conditional pair.
Group:			Any of these
Conditions:	Source IP Address	is equal to	1.1.1.1
	Destination IP Address	is equal to	1.1.1.1
	Source IP Address	is equal to	2.2.2.2
	Destination IP Address	is equal to	2.2.2.2
Group By			Source IP Address
Time Window			4 Hours
Title			${ruleName}

User Watch List: Activity Detected

The following table shows the values for the User Watch List: Activity Detected default incident rule.

Field	Condition Field	Condition Operator	Value
Name			User Watch List: Activity Detected
Description			This incident rule captures alerts generated by network users whose user names have been added as a "Source UserName" condition. To add more than one Username to the watch list, simply add an additional Source Username condition.
Group:			Any of these
Conditions:	Source Username	is equal to	jsmith
	Source Username	is equal to	jdoe
Group By			Source Username
Time Window			4 Hours
Title			${ruleName}

Suspicious Activity Detected: Windows Worm Propagation

The following table shows the values for the Suspicious Activity Detected: Windows Worm Propagation default incident rule.

Field	Condition Field	Condition Operator	Value
Name			Suspicious Activity Detected: Windows Worm Propagation
Description			This incident rule captures alerts that are indicative of worm propagation activity on a Microsoft network
1st Group:			All of these
Condition:	Source	is equal to	Event Stream Analysis
2nd Group:			Any of these
Conditions:	Alert Name	is equal to	Windows Worm Activity Detected Logs

Field	Condition Field	Condition Operator	Value
	Alert Name	is equal to	Windows Worm Activity Detected Packets
Group By			Source IP Address
Time Window			1 Hour
Title			${ruleName}

Suspicious Activity Detected: Reconnaissance

The following table shows the values for the Suspicious Activity Detected: Reconnaissance default incident rule.

Field	Condition Field	Condition Operator	Value
Name			Suspicious Activity Detected: Reconnaissance
Description			This incident rule captures alerts that identify common ICMP host identification techniques (i.e. "ping") accompanied by connection attempts to multiple service ports on a host
1st Group:			All of these
Condition:	Source	is equal to	Event Stream Analysis
2nd Group:			Any of these
Conditions:	Alert Name	is equal to	Port Scan Horizontal Packet
	Alert Name	is equal to	Port Scan Vertical Packet
	Alert Name	is equal to	Port Scan Horizontal Log
	Alert Name	is equal to	Port Scan Vertical Log
Group By			Source IP Address
Time Window			4 Hours
Title			${ruleName}

Monitoring Failure: Device Not Reporting

The following table shows the values for the Monitoring Failure: Device Not Reporting default incident rule.

Field	Condition Field	Condition Operator	Value
Name			Monitoring Failure: Device Not Reporting
Description			This incident rule captures any instance of an alert designed to detect the absence of log traffic from a previously reporting device
Group:			All of these
Conditions:	Source	is equal to	Event Stream Analysis
	Alert Name	is equal to	No logs traffic from device in given time frame
Group By			Source IP Address
Time Window			2 Hours
Title			${ruleName}

Web Threat Detection

The following table shows the values for the Web Threat Detection default incident rule.

Field	Condition Field	Condition Operator	Value
Name			Web Threat Detection
Description			This incident rule captures alerts generated by the NetWitness Web Threat Detection platform.
Group:			All of these
Condition:	Source	is equal to	Web Threat Detection
Group By			Alert Rule Id
Time Window			1 Hour
Title			${ruleName} for ${groupByValue1}

User Entity Behavior Analytics

The following table shows the values for the User Entity Behavior Analytics default incident rule.

Field	Condition Field	Condition Operator	Value
Name			User Entity Behavior Analytics
Description			This incident rule captures user entity behavior.
Group:			All of these
Condition:	Source	is equal to	User Entity Behavior Analytics
Group By			UEBA Classifier Id
Time Window			1 Hour
Title			${ruleName} for ${groupByValue1}

Detect AI

The following table shows the values for the Detect AI default incident rule.

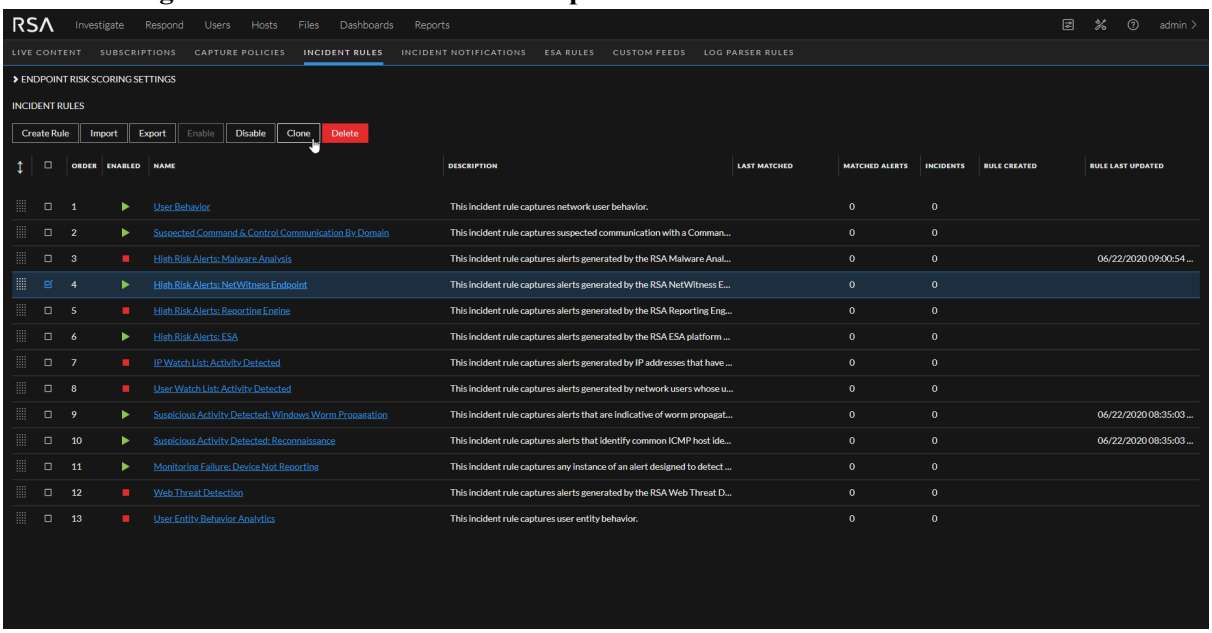
Field	Condition Field	Condition Operator	Value
Name			DetectAI
Description			This incident rule captures anomalies generated by Detect AI
Group:			All of these
Condition:	Source	is equal to	DetectAI
Group By			UEBA Classifier Id, UEBA Entity Name
Time Window			1 Hour
Title			${ruleName} for ${groupByValue2}

Create a NetWitness Endpoint Incident Rule using File Hash

To aggregate NetWitness Endpoint alerts based on the File Hash, create another NetWitness Endpoint Rule using the File Hash as the Group By value. To do this, clone the default NetWitness Endpoint incident rule and change the Group By value.

1. Go to  **(Configure) > Incident Rules**.
   The Incident Rules view is displayed.

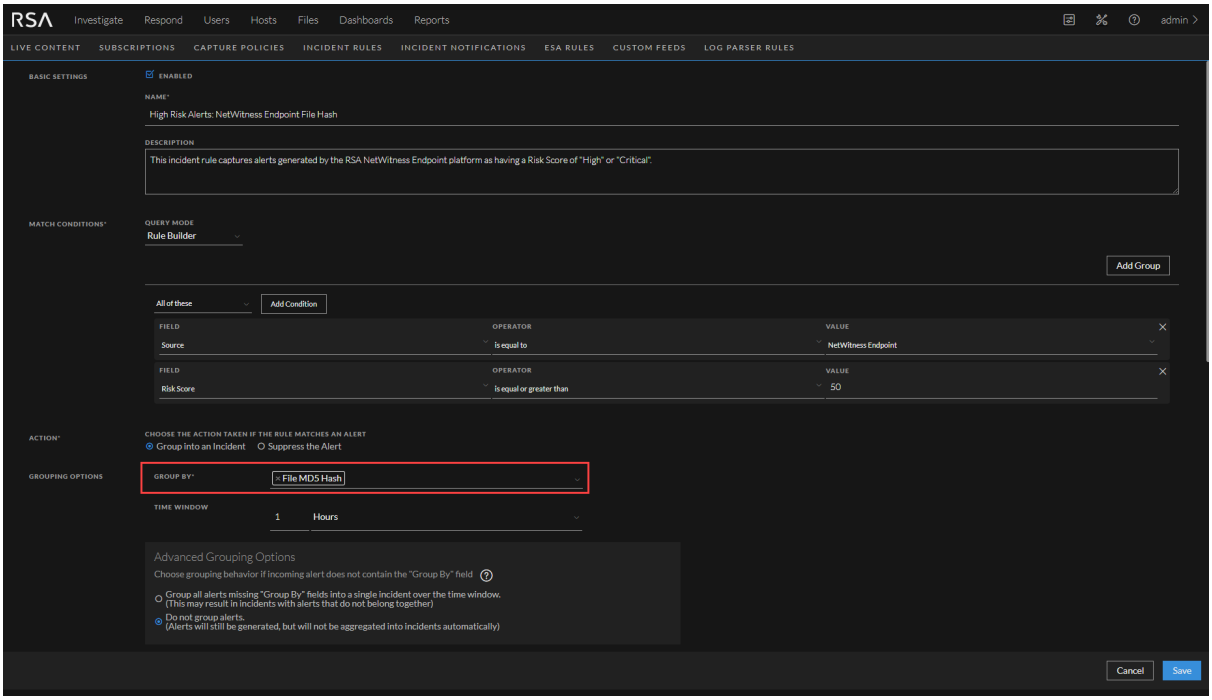
2. Select the **High Risk Alerts: NetWitness Endpoint** default incident rule and click **Clone**.



You will receive a message that you successfully cloned the selected rule.

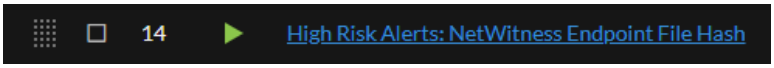
3. Change the **Name** of the rule to an appropriate name, such as High Risk Alerts: NetWitness Endpoint File Hash.

4. In the **Group By** field, remove the previous Group By value and add **File MD5 Hash**.
It is important that File MD5 Hash is the only Group By value listed.



5. If you are ready to enable your rule, in the Basic Settings section, select **Enabled**.

6. Click **Save** to create the rule.
   The Incident Rules view shows your new rule.



7. Verify the order of your incident rules. For more information, see Verify the Order of Your Incident Rules.

Configure Risk Scoring Settings for Automated Incident Creation

Note: The information in this topic applies to NetWitness Version 11.3 and later.

Endpoint Risk Scoring Settings only apply to NetWitness Endpoint.

In addition to automatically creating incidents with predefined rules and rules that you define, NetWitness Respond automatically creates risk scoring incidents for suspicious files and hosts when defined risk score thresholds are crossed. In the background, it monitors the following types of alerts and calculates risk scores for each file and host:

- Critical and High priority alerts from NetWitness Respond
- Medium priority Endpoint alerts from ESA

NetWitness Respond calculates risk score using a combination of the number of distinct alerts and the severity of alerts associated with the file or host. A higher risk score indicates more of these types of alerts. When the calculated risk score exceeds the specified threshold, NetWitness Respond does one of the following during the specified time window, such as 1 day:

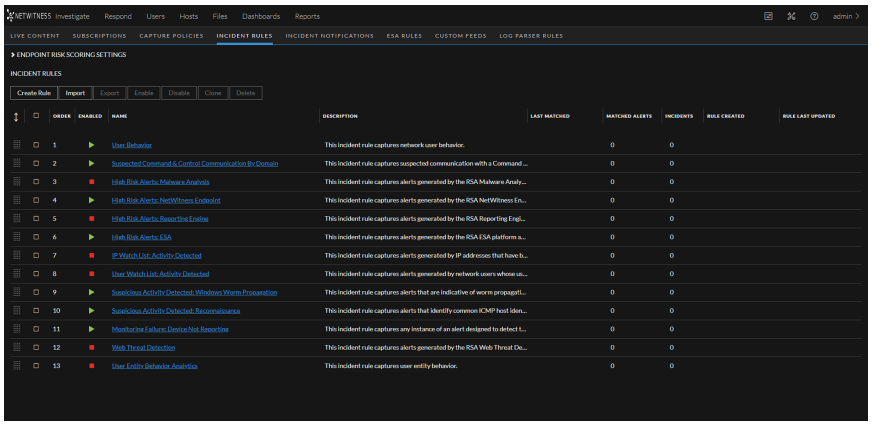
- Creates a risk scoring alert and uses it to create a risk scoring incident
- Adds risk scoring alerts along with associated events to the same incident

For more information on configuring NetWitness Endpoint, see the *NetWitnesss Endpoint Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

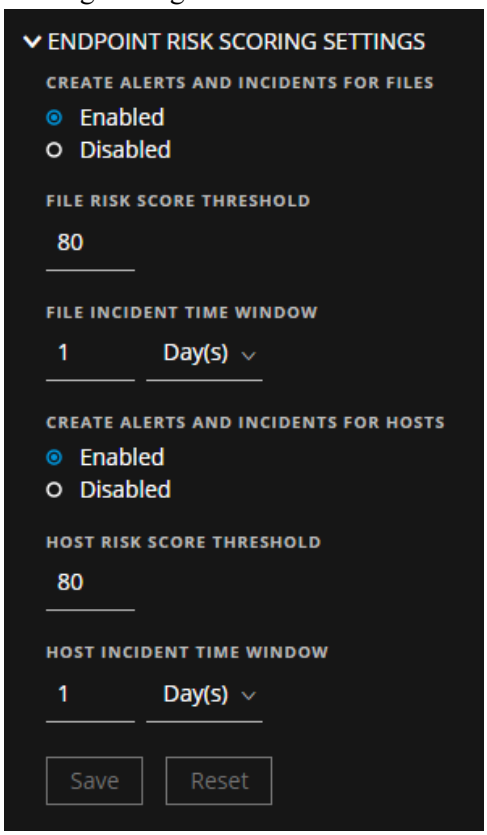
To configure the Endpoint Risk Scoring Settings:

You should leave the Endpoint Risk Scoring Settings at the default values. However, if you are getting too many risk scoring alerts and incidents created, increase the risk score threshold to a higher value. Also, if you are getting too many incidents created for the same hosts or files, increase the time window to add more alerts to the same risk scoring incidents. If you are not seeing many risk scoring incidents, you can either decrease the risk scoring thresholds for hosts and files or decrease the incident time windows.

1. Go to  **(Configure) > Incident Rules**.
   The Incident Rules view is displayed.



2. Click the arrow in front of **ENDPOINT RISK SCORING SETTINGS** to expand the Endpoint Risk Scoring Settings section.



3. In the Endpoint Risk Scoring Settings section, adjust the settings as follows:

   a. **Create Alerts and Incidents for Files**:

   - Select **Enabled** to automatically create risk scoring alerts and incidents for suspicious files. When calculated file risk scores go above the file risk score threshold, it triggers the creation of risk scoring alerts and incidents.
   - Select **Disabled** to stop automatically creating risk scoring alerts and incidents. If you disable it, incidents are not created for suspicious files where risk scores are high.

b. **File Risk Score Threshold**: The File Risk Score Threshold is the risk score level used to trigger alert and incident creation. The File Risk Score Threshold range is from 0-100. For example, if the File Risk Score Threshold is 80 and the calculated risk score of a suspicious Openme.rar file is 81, which is over the Risk Score Threshold of 80, NetWitness Respond creates a risk scoring alert and incident or adds a risk scoring alert to an existing incident depending on the file incident time window.

- If you are seeing too many alerts and incidents, increase the risk score threshold.
- If you are not seeing many alerts and incidents, decrease the risk score threshold.

c. **File Incident Time Window**: The File Incident Time Window is the period of time to wait before creating another incident. The file incident time window range is from 1-24 (hours or days). For example, the suspicious Openme.rar file has a calculated risk score of 81 and a file time window of 1 day. A risk scoring alert and incident is created for the Openme.rar file. During the time window, any similar risk scoring alerts with the same name created for the Openme.rar file get added to the same incident. At the end of the time window (day 1), if the calculated risk score of the file is still over the file risk score threshold and a change occurs with the risk score, another risk scoring alert and incident gets created and any new risk scoring alerts associated with the file get added to the new incident until the next time window (day 3).

- If you are seeing too many alerts and incidents, increase the incident time window.
- If you are not seeing many alerts and incidents, decrease the incident time window.

d. **Create Alerts and Incidents for Hosts**:

- Select **Enabled** to automatically create risk scoring alerts and incidents for suspicious hosts. When calculated host risk scores go above the host risk score threshold, it triggers the creation of risk scoring alerts and incidents.
- Select **Disabled** to stop automatically creating risk scoring alerts and incidents when calculated host risk scores go above the host risk score threshold. If you disable it, incidents are not created for suspicious hosts where risk scores are high.

e. **Host Risk Score Threshold**: The Host Risk Score Threshold is the risk score level used to trigger alert and incident creation. The host risk score threshold range is from 0-100. For example, if the Host Risk Score Threshold is 80 and the calculated risk score of a suspicious host IP address is 81, which is over the Risk Score Threshold of 80, NetWitness Respond creates a risk scoring alert and incident or adds a risk scoring alert to an existing incident depending on the file incident time window.

- If you are seeing too many alerts and incidents, increase the risk score threshold.
- If you are not seeing many alerts and incidents, decrease the risk score threshold.

f. **Host Incident Time Window**: The Host Incident Time Window is the period of time to wait before creating another incident. The host incident time window range is from 1-24 (hours or days). For example, the suspicious host has a calculated risk score of 81 and a Host Time Window of 1 day. During the time window, any similar risk scoring alerts with the same name created for the suspicious host get added to the same incident. At the end of the time window

(day 1), if the calculated risk score of the host is still over the host risk score threshold and a change occurs with the risk score, another risk scoring alert and incident gets created. Any new risk scoring alerts associated with that suspicious host add to that incident until the next time window.

- If you are seeing too many risk scoring alerts and incidents, increase the incident time window.
- If you are not seeing many risk scoring alerts and incidents, decrease the incident time window.

4. Click **Save**.

Configure Custom Respond Server Alert Normalization

Note: This procedure is optional. Administrators can use it to change Respond Server alert normalization.

Analysts who are content experts can create ESA Correlation rules that generate alerts. When a rule is more complex than what can be specified in the ESA Rule Builder, they can write advanced Event Process Language (EPL) rules. After the ESA rules are deployed and the alert criteria is met, ESA Correlation-server forwards the raw alerts to Respond-server.

The schema, meta key selection, and event pattern of the raw alerts are unknown to the Respond-server since they depend on the way the ESA rules are written. In this case, you can customize the logic to parse the raw alert to an acceptable format. The parsing (normalization) logic is written in JavaScript language and NetWitness users who know how to write JavaScript code can customize the Respond normalization script files.

In NetWitness version 11.4 and later, to prevent overwriting future customizations in the Respond normalization scripts, add any custom logic to the **`custom_normalize_<alert type>.js`** files. The custom normalization script files have a `custom_normalize` prefix and are located in the **`/var/lib/netwitness/respond-server/scripts`** directory:

```
data_privacy_map.js
custom_normalize_alerts.js
custom_normalize_core_alerts.js
custom_normalize_ecat_alerts.js
custom_normalize_ma_alerts.js
custom_normalize_ueba_alerts.js
custom_normalize_wtd_alerts.js
utils.js
```

For Example, the `custom_normalize_core_alerts.js` is the normalization script used to add custom logic for ESA alerts. This JavaScript file has a `normalizeAlert` function with the parameters `headers`, `rawAlert`, and `normalizedAlert`. The `normalized` variable is an immutable copy object which has an embedded object with a list of normalized events. So if you have any custom meta keys configured for the events then you have to iterate through the `normalized.events` to populate the appropriate meta keys with values from the `rawAlert.events` object. The following figure shows sample code.

Sample Custom Normalize Core Alerts Code

```
normalizeAlert = function (headers, rawAlert, normalizedAlert) { // normalizedAlert is the immutable copy of ootb normalizer alert, make sure you use // normalized object to update/set the values in your scripts var normalized = Object.assign(normalizedAlert); var custom_events; if(normalized.events !== undefined) { custom_events = normalized.events; } else { custom_events = new Array([]); } for (var i = 0; i < rawAlert.events.length; i++) { custom_events[i].legalentity=Utils.stringValue (rawAlert.events[i].isgs_legalentity); custom_events[i].companycode=Utils.stringValue (rawAlert.events[i].isgs_companycode); } if(normalized.events === undefined){ normalized.events = custom_events; } return normalized; };
```

To Configure Custom Respond Server Alert Normalization:

1. Open the `custom_normalize_core_alerts.js` file.

custom_normalize_core_alerts.js

```
exports.normalizeAlert = function (headers, rawAlert, normalizedAlert) { //
normalizedAlert is the immutable copy of ootb normalizer alert, make sure you use //
normalized object to update/set the values in your scripts var normalized = Object.assign
(normalizedAlert); // Add custom logic below return normalized; };
```

The `normalizedAlert` object has been broadcasted, which already comes through the basic parsing logic flow where some of the meta values have been copied to `normalizedAlert` from `headers` and `rawAlert`.

2. Populate The `normalizedAlert` object with the custom meta values that are not covered in the basic parsing logic:
   a. Add your custom logic below the line: `var normalized = Object.assign (normalizedAlert);`

      The `headers` parameter has very few attributes like:

      `headers.severity`

      `headers.deviceProduct`

      The important parameter is `rawAlert`, which has an embedded `events` object, which is an array. It is basically the list of events associated with the alert. The `events` object has all meta keys governed from the Concentrator. Here are some example meta keys:

      `event.user_dst`
      `event.user_src`
      `event.username`
      `event.domain_dst`
      `event.domain_src`
      `event.host_dst`
      `event.host_src`
      `event.analysis_session`
      `event.analysis_service`
      `event.analysis_file`
      `event.agent_id`
      `event.device_type`
      `event.category`
      `event.action`
      `event.user`
      `event.owner`
      `event.port_dst`
      `event.OS`

`event.process_vid_src`

To view the look-up table for meta keys, go to  **(Configure) > Esa Rules > Settings > Meta Key References**.

The `normalized` object also has an embedded `events` object, which is an array.

b. Iterate through each item of the `normalized.events` and `rawAlert.events` and then copy the custom meta attributes from the `rawAlert.events` to the `normalized.events`.

The following example shows how to add custom meta keys to the `custom_normalize_core_alerts.js` file.

Example of Custom Alert Normalization

```
exports.normalizeAlert = function (headers, rawAlert, normalizedAlert) { // normalizedAlert is the immutable copy of ootb normalizer alert, make sure you use // normalized object to update/set the values in your scripts var normalized = Object.assign (normalizedAlert); var custom_events; if (normalized.events != undefined) { custom_events = normalized.events; } else { custom_events = new Array(); } // iterate through each item for (var i = 0; i < rawAlert.events.length; i++) { custom_events [i].metaKey1=Utils.stringValue(rawAlert.events[i].rawCustomMetaKey1); // Utils is a helper module to stringify the object custom_events[i].metaKey2=Utils.stringValue (rawAlert.events[i].rawCustomMetaKey2); } If (normalized.events == undefined) { normalized.events = custom_events; } return normalized; };
```

The `metaKey1` and `metaKey2` meta keys are now assigned meta keys, which you can view in the NetWitness user interface.

When customizing normalization script files, you can also look at the built-in Respond normalization script files for reference, such as `normalize_alerts.js`.

Configure Analyst UI for Respond Server Alert Normalization

This procedure is optional. Administrators can use it to change Respond Server alert normalization on the Analyst UI.

Note: This option is available in NetWitness version 11.4 and later.

The Analyst UI (Analyst User Interface) enhances the performance of investigations for analysts who work in locations geographically separated from the NetWitness Server host. Respond Server alert normalization is disabled by default on the Analyst UI, but with enough bandwidth you can configure the Respond Server on the Analyst UI to normalize alerts for potential performance gains.

Respond Server alert normalization at the Analyst UI should be very carefully considered. If the Analyst UI is deployed in an environment that is geographically separated from the NetWitness Server (NW Server) and ESA services, depending on available bandwidth, normalizing alerts at the Analyst UI can generate large volumes of traffic, potentially impacting other services on the network. Potential gains from normalizing alerts at the Analyst UI can result in a performance decrease on the NW Server and ESA services.

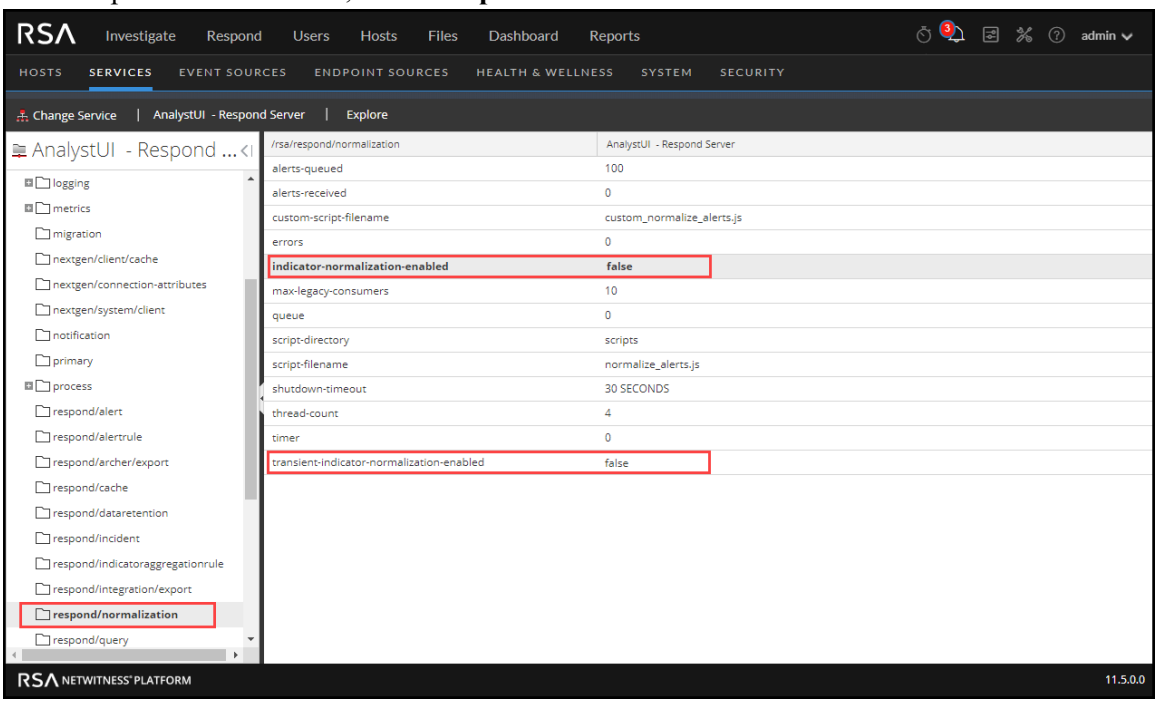
You can configure whether to normalize alerts for any Respond Server (NW Server or Analyst UI) by enabling or disabling alert normalization.

- Normalization is enabled by default for the Respond Server running on the NW Server host.
- Normalization is disabled by default for the Respond Server running on the Analyst UI.

To change the alert normalization settings for the Respond Server running on the Analyst UI:

1. Log in to NetWitness on the NW Server host as administrator.
2. Go to  **(Admin) > Services**, select the Respond Server service running on the Analyst UI, and then select  **> View > Explore**.

3. In the Explore view node list, select **respond/normalization**.



   a. To turn on alert normalization for ESA and other alert generating sources, in the **indicator-normalization-enabled** field, enter **true**. To turn it off, enter **false**.

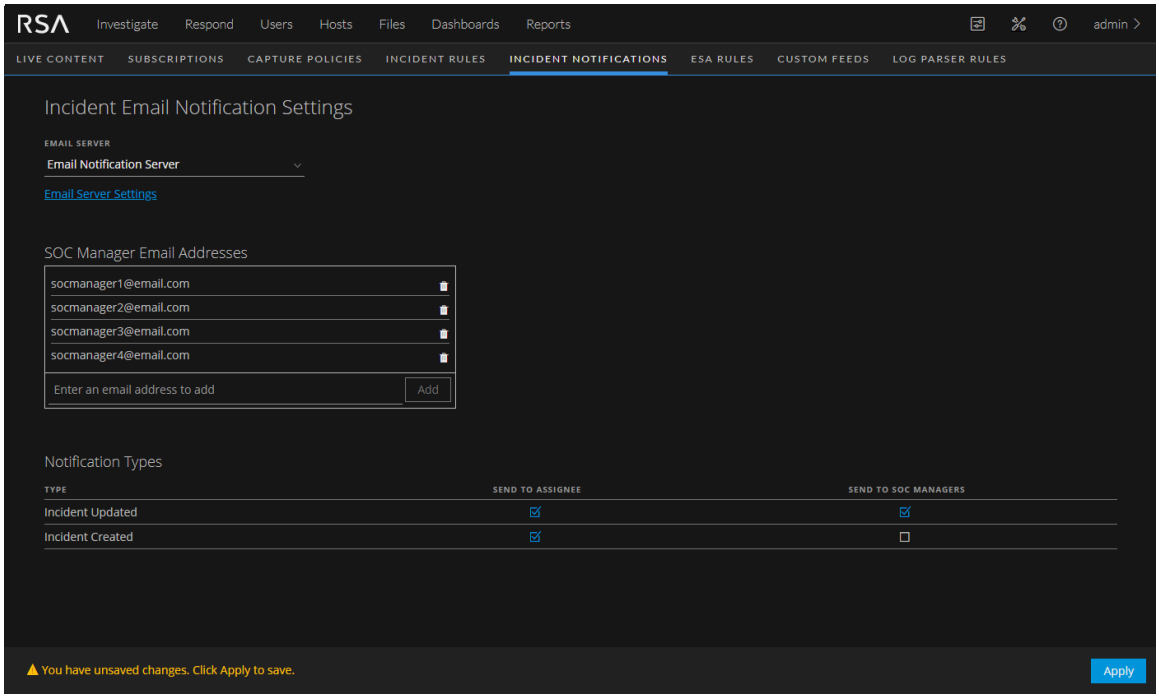
   b. To turn on alert normalization coming from event correlation for risk scoring alerts, in the **transient-indicator-normalization-enabled** field, enter **true**. To turn it off, enter **false**.

4. Restart the Respond Server service on the Analyst UI for the new settings to take effect. To do this, go to  **(Admin) > Services**, select the Respond Server service on the Analyst UI, and then select  **> Restart**.

Configure Incident Email Notification Settings

Incident email notification settings enable email notifications to be sent to SOC Managers and the Analyst assigned to an incident when an incident is created or updated.

1. Go to  **(Configure) > Incident Notifications**.
   The Incident Email Notification Settings view is displayed.



2. In the **Email Server** section, select the email server from the drop-down list that will send out email notifications when the notification settings are enabled.
   If there is no email server configured, you do not see an email server listed in the drop-down list. You have to configure an email server before you can continue with this procedure. To configure an email server, click the **Email Server Settings** link and go to the **Servers** tab. For more information, click the help icon or refer to the *System Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
3. In the **SOC Manager Email Addresses** section, add the email addresses of the SOC Managers that you want to receive email notifications. To add an SOC Manager email address to the list, type it in the field that shows **Enter an email address to add** and click **Add**. To remove an SOC Manager email address from the list, click  next to the email address to be removed.
4. In the **Notification Types** section, select who should receive an email notification when an incident is created and when an incident is updated.

- **Send to Assignee**: An email is sent to the Analyst assigned to the incident.
- **Send to SOC Manager**: An email is sent to all of the addresses listed in the **SOC Manager Email Addresses** list.

5. Click **Apply**. Changes take effect immediately.

Note: If user email address information is updated in the (Admin) > Security > Users tab, it can take up to two minutes for the new email changes to take effect. Any incident creation or incident update email notifications sent during this time go to the old email address.

Migration Considerations

Notification Settings do not migrate from NetWitness version 10.6.x to 11.1 and later. The Incident Management Notification Settings in 10.6.x are different from the Incident Email Notification settings available in 11.1 and later. You will need to manually update the incident notification settings in version 11.1 and later.

Notification Servers from 10.6.x are not displayed in the Email Server drop-down list. The email servers settings must be added to the Global Notification Servers ((Admin) > System > Global Notifications > Server tab).

Custom Incident Management notification templates cannot be migrated to 11.1 and later. No custom templates are supported in 11.1 and later.

Set a Retention Period for Alerts and Incidents

Sometimes data privacy officers want to retain data for a certain period of time and then delete it. A shorter retention period frees up disk space sooner. In some cases, the retention period must be short. For example, laws in Europe state that sensitive data cannot be retained for more than 30 days. After 30 days, the data must be obfuscated or deleted.

Setting a retention period for data is an optional procedure. The time that NetWitness Respond receives alerts and creates an incident determine when retention begins. Retention periods range from 30 to 365 days. If you set a retention period, one day after the period ends data is permanently deleted.

Retention is based on the time that NetWitness Respond receives the alerts and the incident creation time.

Caution: Data deleted after the retention period cannot be recovered.

When the retention period expires, the following data is **permanently deleted**:

- Alerts
- Incidents
- Tasks
- Journal entries

Logs track retention and manual deletion so you can see what has been deleted. You can view Respond Server logs in the following locations:

- **Respond Server Service log**: `/var/log/netwitness/respond-server/respond-server.log`
- **Respond Server Audit log**: `/var/log/netwitness/respond-server/respond-server.audit.log`

The data retention period that you set here does not apply to Archer or other third-party SOC tools. Alerts and incidents from other systems must be deleted separately.

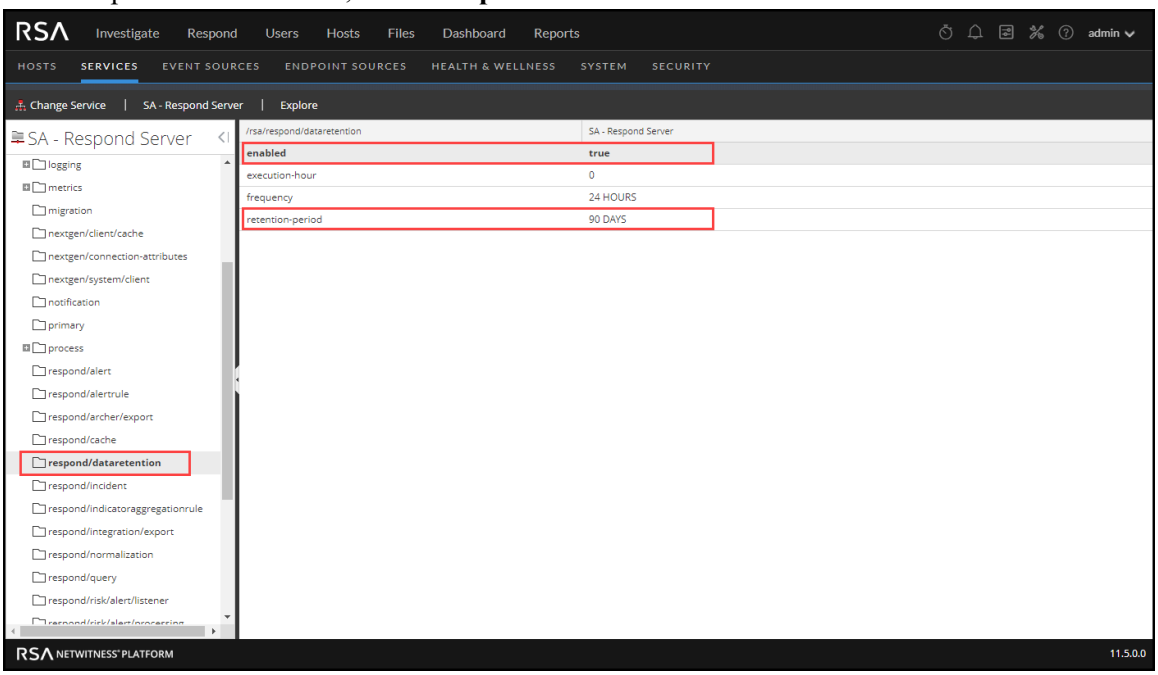
Prerequisites

The Administrator role must be assigned to you.

Procedure

1. Go to  **(Admin) > Services**, select the Respond Server service, and then select  **> View > Explore**.

2. In the Explore view node list, select **respond/dataretention**.



3. In the **enabled** field, select **true** to delete incidents and alerts older than the retention period. The scheduler runs every 24 hours at 23:00.
   You will see a notice that the configuration was successfully updated.
4. In the **retention-period** field, type the number of days to retain incidents and alerts. For example, type 30 DAYS, 60 DAYS, 90 DAYS, 120 DAYS, 365 DAYS, or any number of days.
   A message informs you that the configuration was successfully updated.

Result

Within 24 hours after the retention period ends, the scheduler permanently deletes all alerts and incidents older than the specified period from NetWitness Respond. Journal entries and tasks associated with the deleted incidents are also deleted.

Obfuscate Private Data

The Data Privacy Officer (DPO) role can identify meta keys that contain sensitive data and should display obfuscated data. This topic explains how the administrator maps those meta keys to display a hashed value instead of the actual value.

The following caveats apply to hashed meta values:

- NetWitness supports two storage methods for hashed meta values, HEX (default) and string.
- When a meta key is configured to display a hashed value, all security roles see only the hashed value in the Incidents module.
- You use hashed values the same way you use actual values. For example, when you use a hashed value in rule criteria the results are the same as if you used the actual value.

This topic explains how to obfuscate private data in NetWitness Respond. Refer to the "Data Privacy Management Overview" topic in the *Data Privacy Management Guide* for additional information about data privacy. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Mapping File to Obfuscate Meta Keys

In NetWitness Respond, the mapping file for data obfuscation is data_privacy_map.js. In it you type an obfuscated meta key name and map it to the actual meta key name.

The following example shows the mappings to obfuscate data for two meta keys, ip.src and user.dst:

```
'ip.src.hash' : 'ip.src',
 'user.dst.hash' : 'user.dst'
```

You determine the naming convention for obfuscated meta key names. For example, ip.src.hash could be ip.src.private or ip.src.bin. You must choose one naming convention and use it consistently on all hosts.

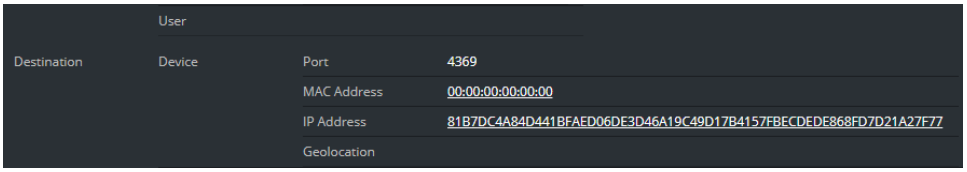
Prerequisites

- DPO role must specify which meta keys require data obfuscation.
- Administrator role must map meta keys for data obfuscation.

Procedure

1. Open the data privacy mapping file:
   /var/lib/netwitness/respond-server/scripts/data_privacy_map.js
2. In the `obfuscated_attribute_map` variable , type the name of a meta key to hold obfuscated data. Then map it to the meta key that does not contain obfuscated data according to this format:
   `'ip.src.hash' : 'ip.src'`
3. Repeat step 2 for every meta key that should display a hashed value.
4. Use the same naming convention as in step 2 and use it consistently on all hosts.

5. Save the file.
   All mapped meta keys will display hashed values instead of actual values.
   In the following figure, a hashed value displays for the destination IP address in the Event Details:



   New alerts will display obfuscated data.

Note: Existing alerts still display sensitive data. This procedure is not retroactive.

Manage Incidents in Archer Cyber Incident & Breach Response

If you want to manage incidents in Archer Cyber Incident & Breach Response instead of NetWitness Respond, you have to configure system integration settings in the Respond Server service Explore view. After you configure the system integration settings, all incidents are managed in Archer Cyber Incident & Breach Response. Incidents created before the integration will not be managed in Archer Cyber Incident & Breach Response.

Caution: If you are managing incidents in Archer Cyber Incident & Breach Response instead of NetWitness Respond, do not use the following in the Respond view: Incidents List view, Incident Details view, and Tasks List view. Do not create incidents from the Respond Alerts List view or from Investigate. In NetWitness 11.4 and later, you can manually create incidents from Respond and Investigate.

For more detailed integration information, see the *NetWitness Archer Integration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Prerequisites

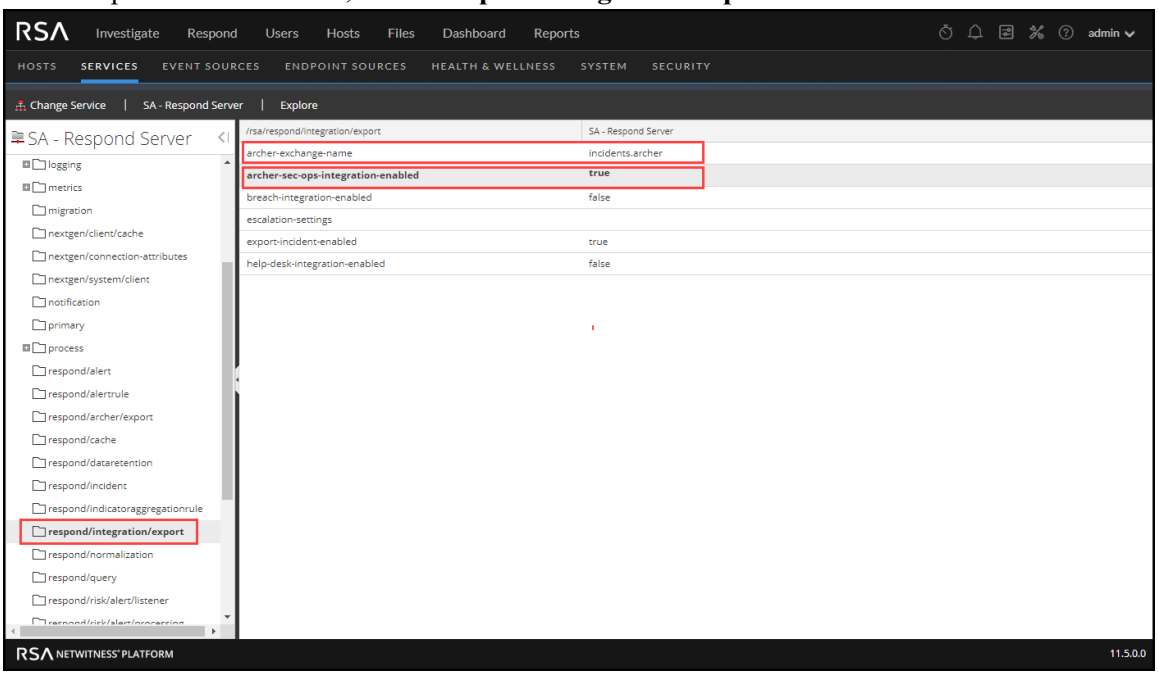
- Archer Cyber Incident & Breach Response 1.3.1.2 (NetWitness 11.0 works only with Archer Cyber Incident & Breach Response 1.3.1.2.)

Procedure

Follow this procedure to configure Respond Server service settings to manage incidents in Archer Cyber Incident & Breach Response.

1. Go to  **(Admin) > Services**, select the Respond Server service, and then select  **> Config > Explore**.

2. In the Explore view node list, select **respond/integration/export**.



3. In the **archer-exchange-name** field, type `incidents.archer`.
   You will see a notice that the configuration was successfully updated.
4. In the **archer-sec-ops-integration-enabled** field, select **true**.
   A message informs you that the configuration was successfully updated.
   Incidents will be managed exclusively in Archer Cyber Incident & Breach Response.

Configure the Option to Send Incidents to Archer

Note: The information in this topic applies to NetWitness Version 11.2 and later.

If you want to manage incidents in NetWitness Respond, you have the option to configure the NetWitness so that you can send incidents to Archer Cyber Incident & Breach Response. If Archer is configured as a data source in Context Hub, you can send incidents to Archer Cyber Incident & Breach Response and you will be able to see a Send to Archer option and a Sent to Archer status in NetWitness Respond. For information on how to use the Send to Archer option and Sent to Archer status, see the *NetWitness Respond User Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Prerequisites

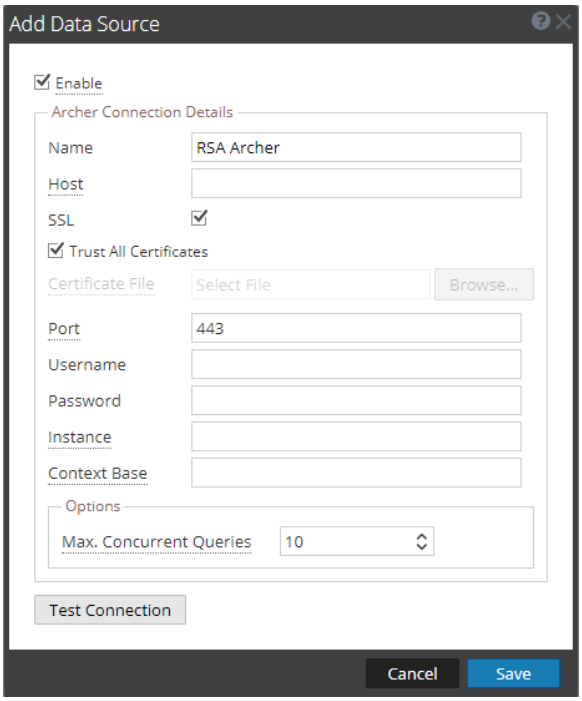
- Archer release 6.6 P4, 6.6 P5, or 6.7 P2 only is required for NetWitness 11.4 and 11.5.
- Archer release 6.4 or later is required for NetWitness 11.2 and 11.3.

Add Archer as a Data Source for Context Hub

To configure sending incidents to Archer Cyber Incident & Breach Response from NetWitness Respond, Archer must be configured as a data source for Context Hub. For more detailed instructions for configuring the Archer data source, see the "Configure Archer as Data Source" topic in the *Context Hub Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

1. Go to  **(Admin) > Services**.
   The Services view is displayed.
2. Select the Context Hub service, and then select  **> View > Config**.
   The Services Config view is displayed.

3. On the **Data Sources** tab, click + > **RSA Archer**.
   The **Add Data Source** dialog is displayed.



4. Provide the following information:

   - By default, the **Enable** checkbox is selected. If this option is unchecked, the save button is disabled, you cannot add the data source, and cannot view the contextual information.
   - Enter the following fields:
     - **Name**: Enter a name for Archer data source.
     - **Host**: Enter the hostname or IP address where Archer server is installed.
     - **SSL**: By default this option is selected and enables SSL communication to Archer .
     - **Trust All Certificates**: Select this checkbox to add the data source without validating the certificate. If you uncheck this option, you need to upload a valid Endpoint server certificate for the connection to be successful.
     - **Port**: The default port is 443.
     - **Username**: Enter the Archer Server username.
     - **Password**: Enter the Archer Server password.
     - **Instance**: Enter the Instance name from which you want to extract data. An Archer instance is a single setup that includes unique content in a database, the connection to the database, the interface, and login. You might have individual instances for each office location or region or for development, test, and production environments. The Instance Database stores the Archer

content for a specific instance.

- **Context Base**: Enter the virtual directory name where the files are stored. For example, rsaarcher located at the Archer web address https://archer.company.com/**rsaarche**r/default.aspx. If the files are stored in the IIS default web address https://archer.company.com/default.aspx, then this field must be empty.
- **Max. Concurrent Queries**: You can configure the maximum number of concurrent queries defined by the Context Hub service to be run against the configured data sources. The default value is 10.

5. Click **Test Connection** to test the connection between Context Hub and the Archer data source.
6. Click **Save**.
   Archer is added as a data source for Context Hub and is displayed in the **Data Sources** tab. A Send to Archer button and Sent to Archer status is visible in NetWitness Respond.

Configure Threat Aware Authentication

Note: The information in this topic applies to NetWitness Version 11.3 and later.

NetWitness Platform creates a list of suspicious users that have an incident created against them and sends it to RSA SecurID Access. The list contains the email IDs of the corresponding suspicious users associated with the incident. RSA SecurID Access maintains this high-risk users list and reduces the access levels or blocks such users using defined policies. When an incident is closed in NetWitness Platform, the associated email IDs are automatically removed from the RSA SecurID high-risk user list.

By default this configuration is disabled in the NetWitness Server. You can enable this feature by editing the yml file located at `/etc/netwitness/respond-server/respond-server.yml`.

Enable Threat Aware Authentication

To enable this configuration:

1. Create a yml file at `/etc/netwitness/respond-server/respond-server.yml`
2. Edit and enter `rsa.respond.securid-integration.enabled: true`

```
[root@adminserver ~]# vi /etc/netwitness/respond-server/respond-server.yml
[root@adminserver ~]# cat /etc/netwitness/respond-server/respond-server.yml
rsa.respond.securid-integration.enabled: true
rsa.security.pki.use-jvm-trust: true
[root@adminserver ~]# service rsa-nw-respond-server restart
Redirecting to /bin/systemctl restart rsa-nw-respond-server.service
[root@adminserver ~]#
```

3. Enter `rsa.security.pki.use-jvm-trust: true` to enable the configuration.
4. Save the yml file and restart the Respond Server service.

Note: Make sure you perform the above configuration if you have enabled a stand-by NW server. In case the primary NW server fails and goes offline, this configuration will allow the standby NW server to connect to RSA SecurID.

Obtain SecurID API Key

A super administrator must generate and download a SecurID API key, and connect to RSA SecurID Access.

To obtain the API key from RSA SecurID Access:

1. Log in to the **RSA SecurID Access Cloud Administration** Console.
2. Click **Platform** > **API Key Management**.
3. Click **ADD**.
   The new key is displayed.

4. Change the **Administrator** role to **Super Administrator**.
5. Click **Save** and **Download** to download and save the API key file.

For more information about generating the API Keys and other related details, see "Manage the Cloud Administration API Keys" at https://community.securid.com/t5/securid-cloud-authentication/manage-the-cloud-administration-api-keys/ta-p/623066 and "Determining Access Requirements for High-Risk Users in the Cloud Authentication Service" topic at https://community.securid.com/t5/securid-cloud-authentication/determining-access-requirements-for-high-risk-users-in-the-cloud/ta-p/623067.

Configure RSA SecurID Access API Key

To configure RSA SecurID Access API key using NetWitness Shell:

1. SSH to the NetWitness Server.
2. Type the command `nw-shell`.
   A console window is displayed.



3. Type `connect --service respond-server.<service-id>` to connect to the Respond Server.
   For example: `connect --service respond-server.36334277-9f93-4402-9523-ed15ad543bfa`.
   You can obtain the `<service_id>` from `cat /etc/netwitness/respond-server/service-id`.
4. Type `login` and enter admin username and password.
5. To set the API key:
   a. Navigate to set-api-key node: `cd /rsa/respond/securid/set-api-key`
   b. type: `invoke --file <path to api key>`

   Note: The path to the API key is the location on the NetWitness Server.

6. Test the connection using the command:
   a. `cd /rsa/respond/securid/test-secur-id-connection`
   b. Type `invoke`.
      A "Connection OK" message is displayed if test connection is successful.
7. To start the process use the command:
   a. `cd /rsa/respond/securid/process-incidents`
   b. `invoke`.

   For more information on how to define policies, see the *RSA SecurID Access Guide* on NetWitness Community.

Configure Sync Frequency

By default, the sync frequency is set to 15 minutes.

To edit the frequency:

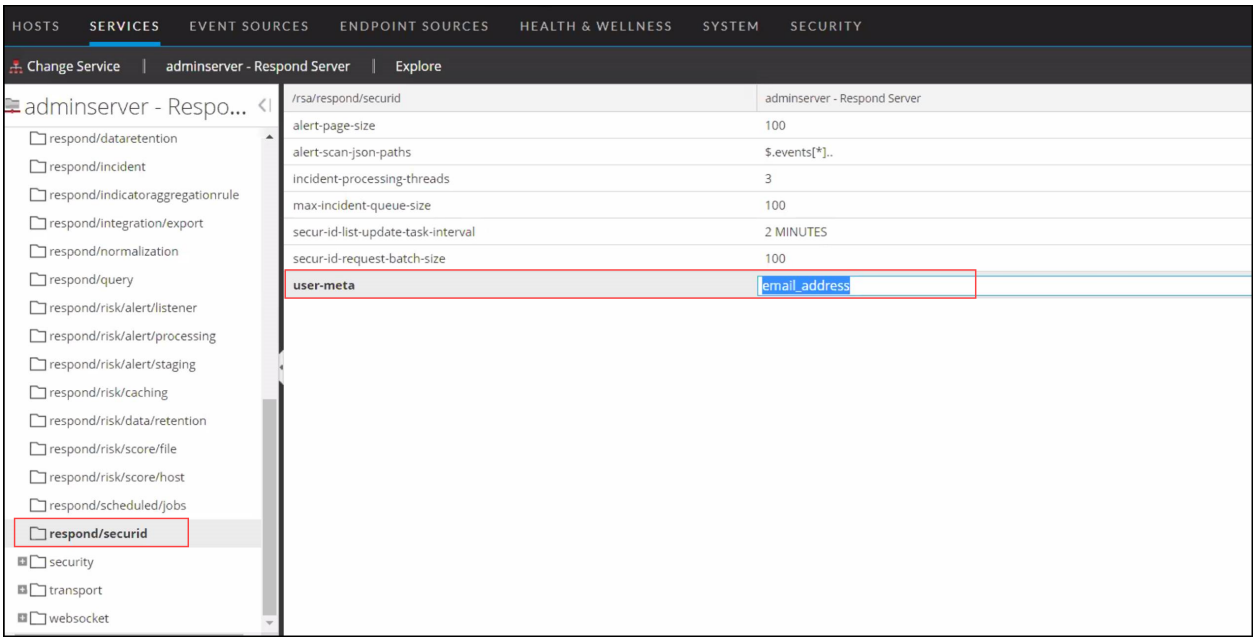
1. Log in to NetWitness.
2. Go to  (Admin) **> Services**, select the Respond Server service, and then select >  **> View > Explore**.
3. Edit the duration at **rsa/respond/securid**.

Configure Meta

You can configure a respond specific meta in an alert to identify a user to be added to SecurID high-risk user's list. By default, the meta is set to `email_address`. Currently, the Respond Server supports metas such as ad_username, and email_address.

To add a Respond Server supported meta:

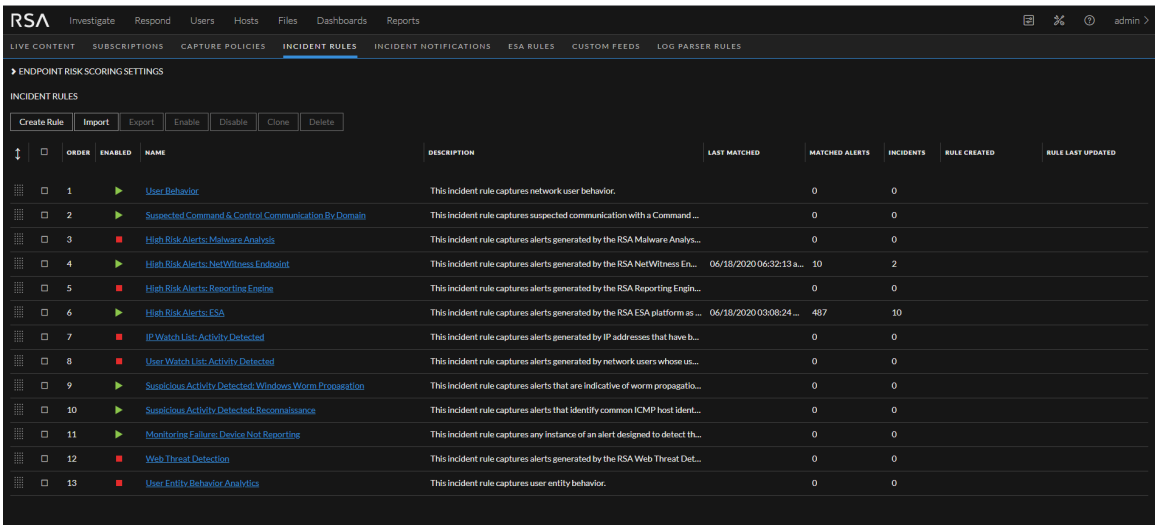
1. Log in to NetWitness Platform.
2. Go to  **(Admin) > Services**, select the Respond Server service, and then select >  **> View > Explore**.
3. In the Explore view node list, select **respond/securid**.
4. Edit and enter a meta in the **user-meta** field.



Note: If at any time you change the meta configuration from email address to ad_username, make sure to run Step 7 from Configure RSA SecurID Access API Key to process the older incidents. In case of a multi-analyst deployment, make sure you configure the same meta on all Respond servers. For example, if you update the meta to ad_username, then the same must be updated on all Respond servers.

Set a Counter for Matched Alerts and Incidents

This procedure is optional. Administrators can use it to change when the count for matched alerts is reset to 0. The Incident Rules view displays these counts in columns on the right.



The following columns provide matched alert information for a rule:

- **Last Matched** column shows the time when the rule last matched alerts.
- **Matched Alerts** column displays the number of matched alerts for the rule.
- **Incidents** column displays the number of incidents created by the rule.

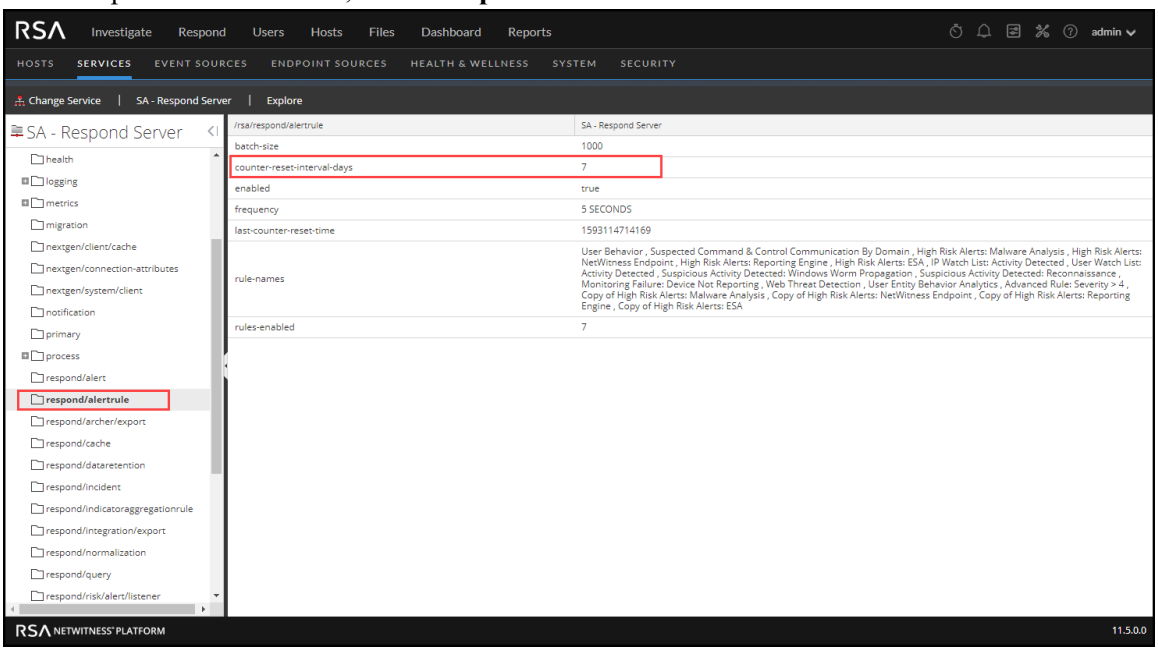
By default, these values reset to zero every 7 days. Depending on how long you want the counts to continue, you can change the default number of days.

Note: When the counter resets to zero, only the numbers in the three columns change to zero. No alerts or incidents are deleted.

To set a counter for matched alerts and incidents:

1. Go to  **(Admin) > Services**, select the Respond Server service, and then select  **> View > Explore**.

2. In the Explore view node list, select **respond/alertrule**.



3. In the right panel, type the number of days in the **counter-reset-interval-days** field.

4. Restart the Respond Server service for the new setting to take effect. To do this, go to  **(Admin) > Services**, select the Respond Server service, and then select  **> Restart**.

Edit the Incident Rules Export ZIP File

This procedure is optional and is for advanced users. When exporting incident rules from the Respond Incident Rules view, the exported incident rules file is a ZIP file in the format `<random ID>-incident_rules_export.json.zip`, which contains two mandatory JSON files:

- `aggregation_rule_schema.json` contains the incident rule schema.
- `<random ID>-incident_rules_export.json` contains the incident rules.

You can import this ZIP file on another NetWitness Server on the same release version.

There may be situations when you need to edit the these files before you import them to another NetWitness Server.

To edit the incident rules export files:

1. Follow the Incident Rule Export Files Editing Guidelines below to edit the export files.
2. Before importing, verify that the ZIP file does not contain additional files or folders. The ZIP file should contain only the mandatory `aggregation_rule_schema.json` and `<random ID>-incident_rules_export.json` files to go through the import. Any files other than these two cause the import to fail.
   For example, when compressing files on a Mac, it adds a temp folder `__MACOSX` that needs to be excluded while zipping the file.

Note: You cannot export Advanced rules.

Incident Rule Export Files Editing Guidelines

Ensure that the following fields have at least one value. Removing a value or having an empty value for the following fields results in abnormal behavior.

Field	Possible Values
name	`A-Z a-z 0-9 " !"#$%&'()*+,-./:;<=>?@[\]^_`{|}~"`
groupByFields	A Minimum of 1 or a Maximum of 2 group_by keys should be present in `aggregation_rule_schema.json`.
timeWindow	A String value in the following accepted formats: Days – Min:`1d` Max: `24d` Hours – Min: `1h` Max: `100h` Minutes – Min: `1m` Max :`100m`

action	Should be one of following values:
	`GROUP_INTO_INCIDENT`
	`SUPPRESS_ALERT`

Field	Possible Values
incidentScoringOptions	Should be one of the following string values: `average`: Average of Risk Score across all of the Alerts `high`: Highest Risk Score available across all of the Alerts `count`: Number of Alerts in the time window
priorityScale	Condition: LOW < MEDIUM < HIGH < CRITICAL

	Sub Fields — **Possible Values**
	`MEDIUM` — `1-100`
	`CRITICAL` — `1-100`
	`HIGH` — `1-100`
	`LOW` — `1-100`

uiFilterConditions	Sample UI Conditions Filter Structure

	Sub Fields — **Possible Values**
	`filterType` — `FILTER`
	`FILTER_GROUP`
	Possible values for `FILTER` are listed below.
	`property` — `value`: fetched from `aggregation_rule_schem.json`
	`operator` — `operators`
	`value` — `type`: dictates the data type.
	Available options:
	`textfield`: String
	`combobox` : from a list of options available in the json
	`datefield`: unix time stamp, for example: 2019-06-12T12:00:00Z
	`numberfield`: Integer

incidentCreationOptions	`ruleSummary`: String `categories`: JSON `array assignee`: JSON

Configure a Database for the Respond Server Service

This procedure is required only if you need to change the database configuration for Respond Server after the deployment of the NetWitness or ESA Primary hosts and their corresponding services. You have to select the ESA Primary server to act as the database host for NetWitness Respond application data, such as alerts, incidents, and tasks. You also have to select the NetWitness Server to act as the database host for NetWitness Respond control data, such as incident rules and categories.

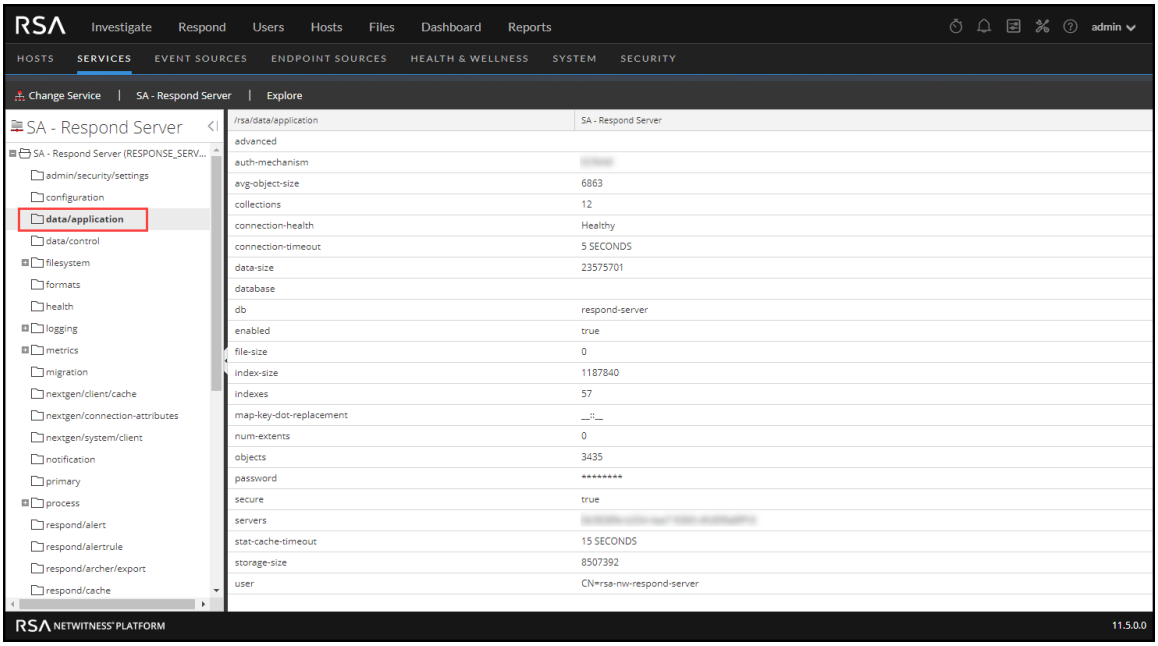
Prerequisites

Ensure that:

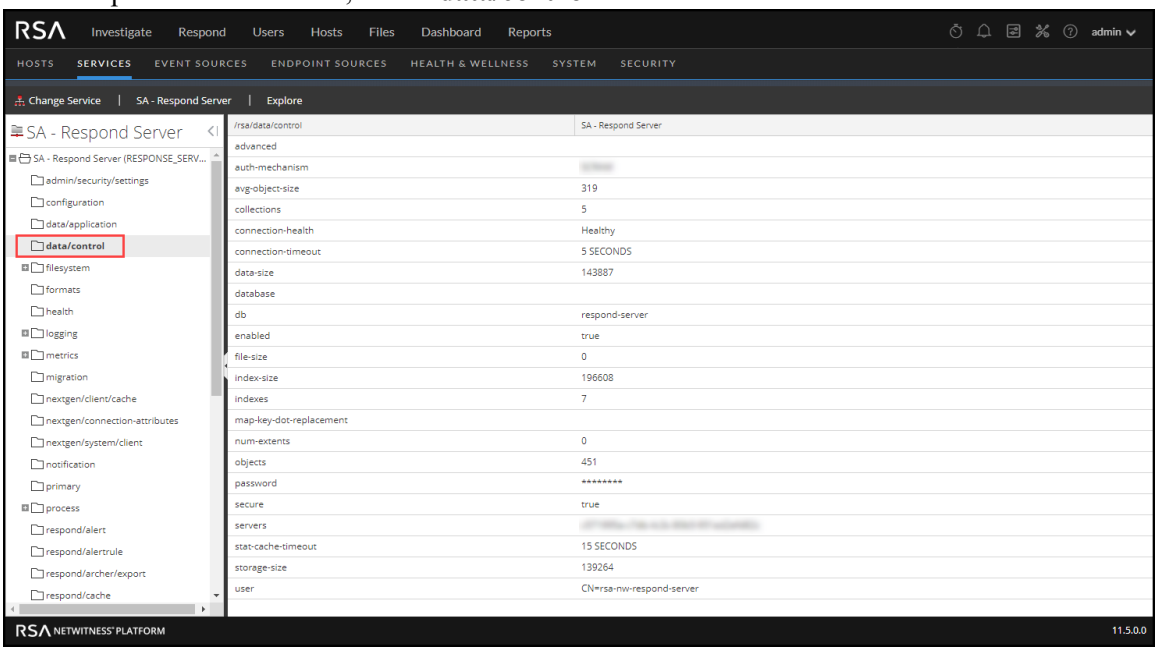
- The Respond Server service is installed and running on NetWitness.
- An ESA host is installed and configured.

Procedure

1. Go to  **(Admin) > Services**.
   The Services view is displayed.
2. In the Services panel, select the Respond Server service and then select  **> View > Explore**.
3. In the Explore view node list, select **data/application**.



4. Provide the following information:
   - **db**: The database name. The default value is respond-server.
   - **password**: The password used for the deployment of the ESA primary server (password for deploy_admin user).
   - **servers**: The hostname or IP address of the **ESA primary server** to act as the database host for NetWitness Respond application data, such as alerts, incidents, and tasks.
   - **user**: Enter **deploy_admin**.
5. In the Explore view node list, select **data/control**.



6. Provide the following information:
   - **db**: The database name. The default value is respond-server.
   - **password**: The password used for the deployment of the NetWitness Server (password for deploy_admin user).
   - **servers**: The hostname or IP address of the **NetWitness Server** to act as the database host for NetWitness Respond control data, such as incident rules and categories.
   - **user**: Enter **deploy_admin**.
7. Restart the Respond Server service. To do this, go to  **(Admin) > Services**, select the Respond Server service, and then select  **> Restart**.

Note: Restarting the Respond Server service is required for the database configuration to be complete.

NetWitness Respond Configuration Reference

This section contains reference information for configuring NetWitness Respond.

Configure View

The Configure view enables you to configure NetWitness Respond functionality.

You can configure incident rules to automate the Respond workflow for automatically creating incidents. You can also configure notification settings to send emails when incidents are created or updated.

Topics

- Incident Rules View
- Incident Rule Details View
- Incident Email Notification Settings View
- Aggregation Rules Tab (11.0 and earlier)
- New Rule Tab (11.0 and earlier)

Incident Rules View

The Incident Rules view enables you to manage the automated incident creation process. NetWitness Respond creates incidents in two ways:

- **Incident Rules**: NetWitness provides preconfigured rules that you can adjust for your environment. You can also create your own rules.
- **Risk Scoring**: (Endpoint Risk Scoring Settings are available in NetWitness version 11.3 and later and only apply to NetWitness Endpoint.) NetWitness Respond uses these settings to automatically create risk scoring incidents for suspicious files and hosts that cross the defined risk score thresholds. If you get too many or too few risk scoring incidents, you can adjust these thresholds.

Note: The information in this topic applies to NetWitness 11.1 and later.

What do you want to do?

Role	I want to ...	Show me how
Analyst, Content Expert, SOC Manager	Create, edit, enable, export, or import an incident rule.	Step 3. Enable and Create Incident Rules for Alerts
Analyst, Content Expert, SOC Manager	Configure the threshold that creates risk scoring alerts and incidents to adjust the amount of alerts and incidents created. Turn off the creation of risk scoring alerts and incidents. Endpoint Risk Scoring Settings only apply to NetWitness Endpoint.	Configure Risk Scoring Settings for Automated Incident Creation
Incident Responders, Analysts, Content Experts, SOC Manager	View the results of my incident rule (View Detected Threats).	See "Responding to Incidents" in the *NetWitness Respond User Guide*.

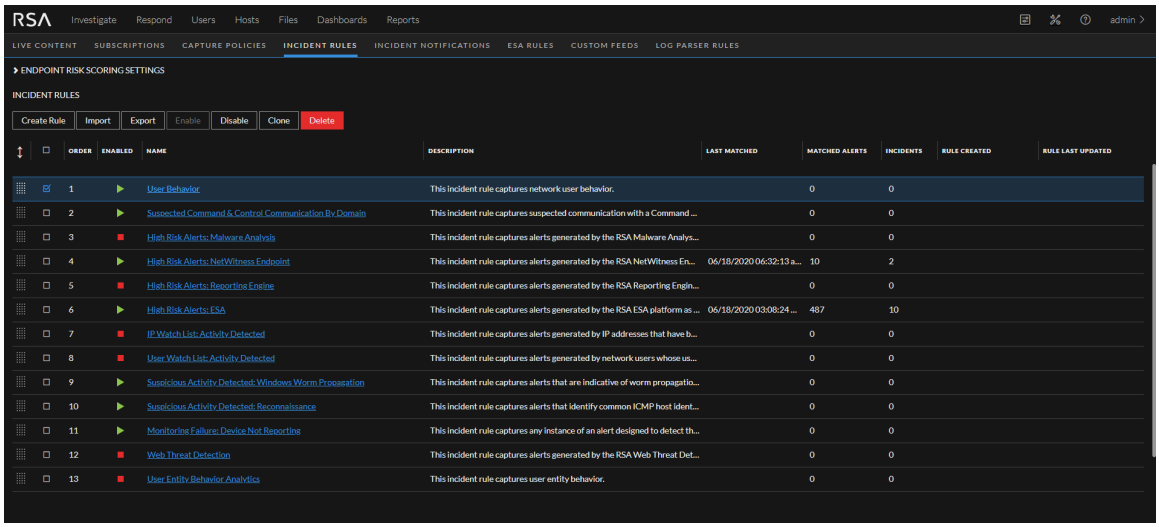
Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Related Topics

- Incident Rule Details View

Quick Look

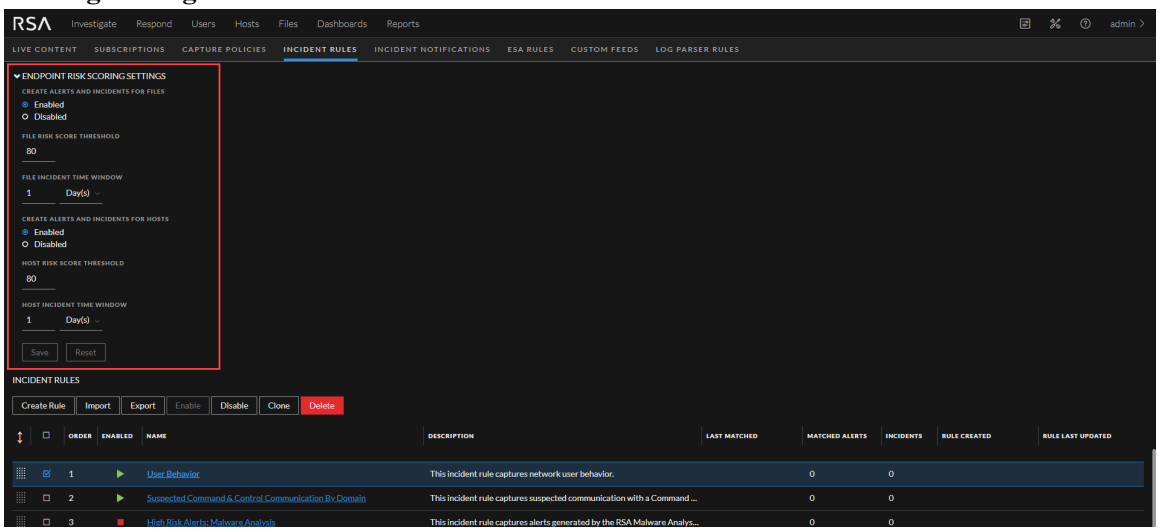
1. To access the Incident Rules view, go to  **(Configure) > Incident Rules**.



   The Incident Rules view has two sections, one for each type of automated incident creation:

   - Endpoint Risk Scoring Settings
   - Incident Rules

2. To view the Endpoint Risk Scoring Settings section, click the arrow in front of **Endpoint Risk Scoring Settings**.

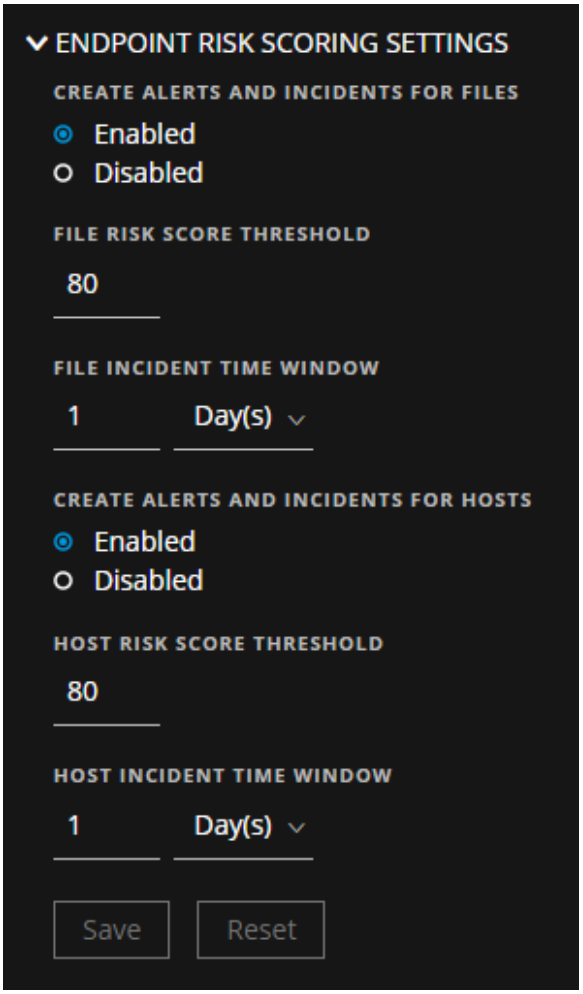


Endpoint Risk Scoring Settings

Note: Endpoint Risk Scoring Settings are available in NetWitness version 11.3 and later and only apply to NetWitness Endpoint. NetWitness Respond uses these settings to automatically create risk scoring incidents for suspicious files and hosts that cross the defined risk score thresholds.

The Endpoint Risk Scoring Settings enable you to configure the thresholds used to automatically create risk scoring alerts and incidents. When calculated risk scores for suspicious files and hosts exceed the specified thresholds, it triggers the creation of risk scoring alerts and incidents. NetWitness recommends that you keep the thresholds at the default values, but you may need to adjust these settings if you get too many or too few alerts and incidents.

For more information on configuring NetWitness Endpoint, see the *NetWitnesss Endpoint Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.



The following table describes the fields in the Endpoint Risk Scoring Settings.

Field / Button	Description
Create Alerts and Incidents for Files	Select **Enabled** to automatically create risk scoring alerts and incidents for suspicious files. When calculated file risk scores go above the file risk score threshold, it triggers the creation of risk scoring alerts and incidents. Select **Disabled** to stop automatically creating risk scoring alerts and incidents. This option is enabled by default.

Field / Button	Description
File Risk Score Threshold	The File Risk Score Threshold is the risk score level used to trigger alert and incident creation. The file risk score threshold range is from 0-100. NetWitness Respond calculates risk score using a combination of the number of distinct alerts and the severity of alerts associated with the file or host. A higher risk score indicates more of these types of alerts. For example, if the file risk score threshold is 80, any calculated file risk score over 80 creates a risk scoring alert and incident or adds a risk scoring alert to an existing incident depending on the file incident time window.
File Incident Time Window	The File Incident Time Window is the period of time to wait before creating another incident. The file incident time window range is from 1-24 (hours or days). For example, an openme.rar file containing suspicious code with enough associated endpoint alerts to get a risk score of 81, which is over the file risk score threshold of 80, automatically creates a risk scoring alert and incident or adds a related risk scoring alert to the same incident within a 1 day time window.
Create Alerts and Incidents for Hosts	Select **Enabled** to automatically create risk scoring alerts and incidents for suspicious hosts. When calculated host risk scores go above the host risk score threshold, it triggers the creation of risk scoring alerts and incidents. Select **Disabled** to stop automatically creating risk scoring alerts and incidents. This option is enabled by default.
Host Risk Score Threshold	The Host Risk Score Threshold is the risk score level used to trigger alert and incident creation. The host risk score threshold range is from 0-100. NetWitness Respond calculates risk score using a combination of the number of distinct alerts and the severity of alerts associated with the file or host. A higher risk score indicates more of these types of alerts. For example, if the host risk score threshold is 80, any calculated host risk score over 80 creates a risk scoring alert and incident or adds a risk scoring alert to an existing incident depending on the host incident time window.
Host Incident Time Window	The Host Incident Time Window is the period of time to wait before creating another incident. The host incident time window range is from 1-24 (hours or days). For example, a suspicious host with enough associated endpoint alerts to get a risk score of 81, which is over the host risk score threshold of 80, automatically creates a risk scoring alert and incident or adds a related risk scoring alert to the same incident within a 1 day time window.

Incident Rules

The Incident Rules section enables you to create and manage incident rules for automating the incident creation process. NetWitness provides preconfigured rules. You can add to and adjust these rules for your own environment.

The Incident Rules section consists of a list and series of buttons. The following table describes the columns in the Incident Rules list.

Column	Description
	Enables you to change the priority order of the rules. Use the drag pad () in front of a rule to move it up and down in the list.
 (Select)	Enables you to select one or more rules in order to take an action, such as Enable or Export. You can select all rules by selecting the checkbox in the column header. You can only select one rule for the Clone or Delete actions.
Order	Shows the order in which the rule is placed. The rule order determines which rule takes effect if the criteria for multiple rules match the same alert. If multiple rules match an alert, only the rule with the highest priority creates an incident. NetWitness Respond evaluates incoming alerts against the incident rules in the order that you define. If alerts match the first rule listed, then that rule creates an incident. If alerts match the second rule listed and those alerts did not match the first rule, then the second rule creates an incident. If alerts match the third rule listed and those alerts did not match the first or second rule listed, then the third rule creates an incident, and so on.
Enabled	Shows whether the rule is enabled or not. The  specifies that the rule is enabled. The  specifies that the rule is not enabled.
Name	Displays the name of the rule with a hyperlink. If you click the link, it opens the Rule Details view, where you can edit the rule.
Description	Displays the description of the rule.
Last Matched	Displays the time when an alert was successfully matched with the rule. This value is reset once a week.
Matched Alerts	Displays the number of matched alerts. This value is reset once a week. To change the setting, see Set a Counter for Matched Alerts and Incidents.
Incidents	Displays the number of incidents created by the rule. This value is reset once a week. To change the setting, see the Set a Counter for Matched Alerts and Incidents.

Column	Description
Rule Created	Shows the date and time that the incident rule was created or imported.
Rule Last Updated	Shows the date and time that the incident rule was last updated.

Incident Rules Actions

The following table shows the operations that can be performed on the Incident Rules list.

Action	Description
Create Rule button	Allows you to add a new incident rule.
Export button	(This option is available in NetWitness 11.4 and later.) Allows you to export one or more incident rules. This enables you to share incident rules with other NetWitness Servers on the same release version. The exported incident rules file is a ZIP file that contains two JSON files: one file contains the incident rules and the other file contains the incident rule schema. You cannot export Advanced incident rules; the export function only allows incident rules created using Rule Builder.
Import button	(This option is available in NetWitness 11.4 and later.) Allows you to import an incident rules ZIP file. This enables you to share incident rules with other NetWitness Servers on the same release version.
Enable button	(This option is available in NetWitness 11.4 and later.) Allows you to enable one or more rules from the incident rules list. You can also enable a rule in the Incident Rule Details view by selecting the Enabled checkbox in the Basic Settings section and then saving the rule.
Disable button	(This option is available in NetWitness 11.4 and later.) Allows you to disable one or more rules from the incident rules list. You can also disable a rule in the Incident Rule Details view by clearing the Enabled checkbox in the Basic Settings section and then saving the rule.
Clone button	Allows you to duplicate one incident rule at a time.
Delete button	Allows you to delete one incident rule at a time.
Name hyperlink	Allows you to edit an incident rule.

Incident Rule Details View

The Incident Rule Details view enables you to create and edit incident rules for creating incidents from alerts. This topic describes the information required when creating or editing a new rule.

Note: The information in this topic applies to NetWitness Version 11.1 and later.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure alert sources for the Respond view.	Step 1. Configure Alert Sources to Display Alerts in the Respond View
Administrator	Assign Respond view permissions.	Step 2. Assign Respond View Permissions
Analyst, Content Expert, SOC Manager	Enable, create, or edit an incident rule.	Step 3. Enable and Create Incident Rules for Alerts
Analyst, Content Expert, SOC Manager	Set up and use the User Behavior default rule. Set up or verify the preconfigured (default) incident rules.	Set Up and Verify Default Incident Rules
Incident Responders, Analysts, Content Experts, SOC Manager	View the results of my incident rule (View Detected Threats).	See "Responding to Incidents" in the *NetWitness Respond User Guide*.

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Related Topics

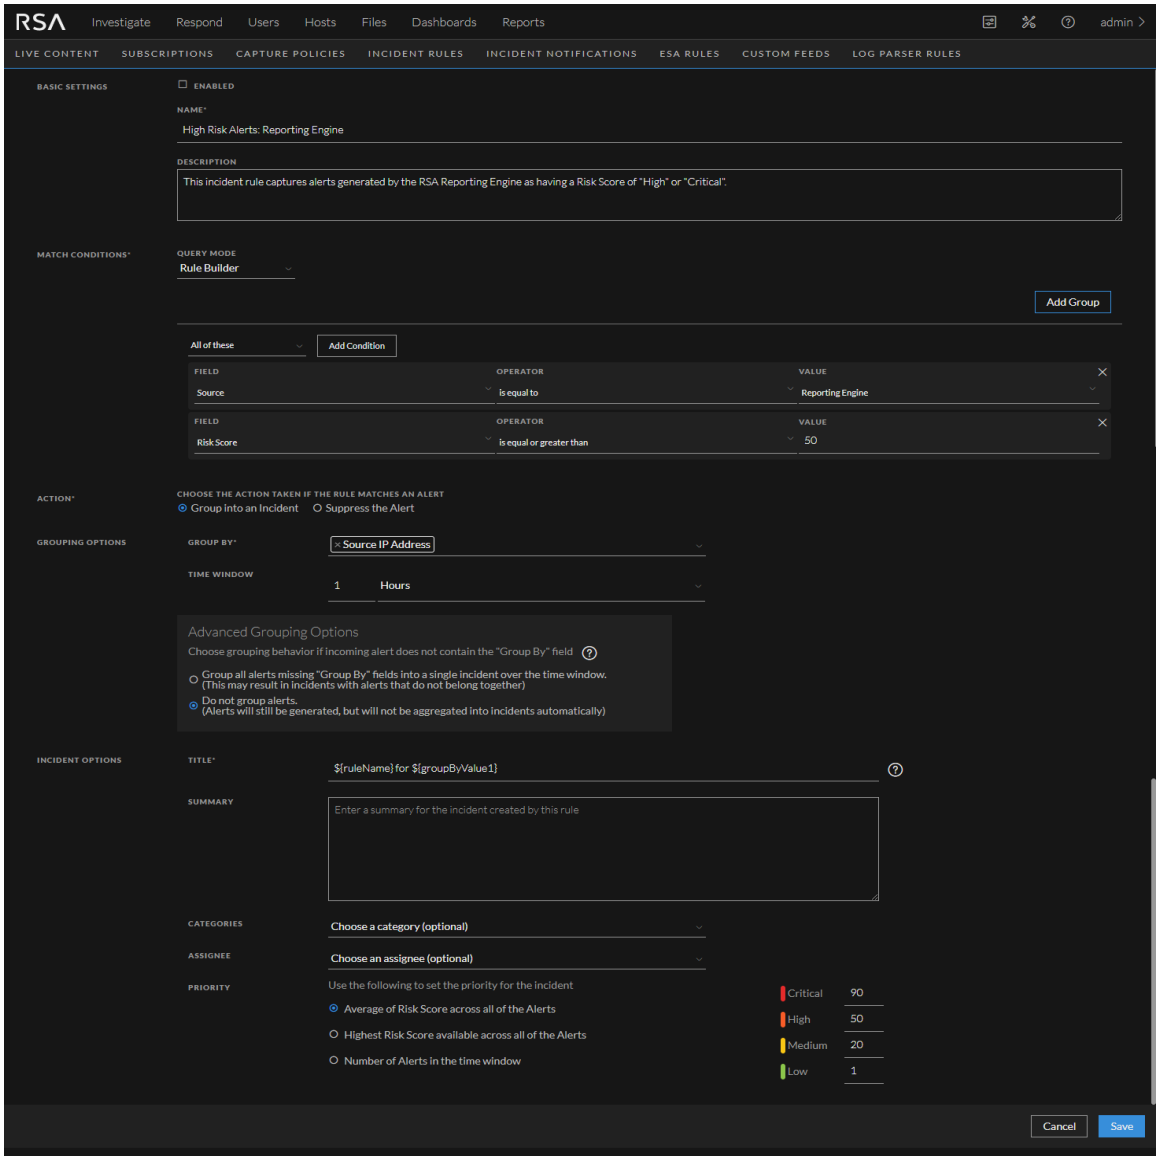
- Incident Rules View

Quick Look

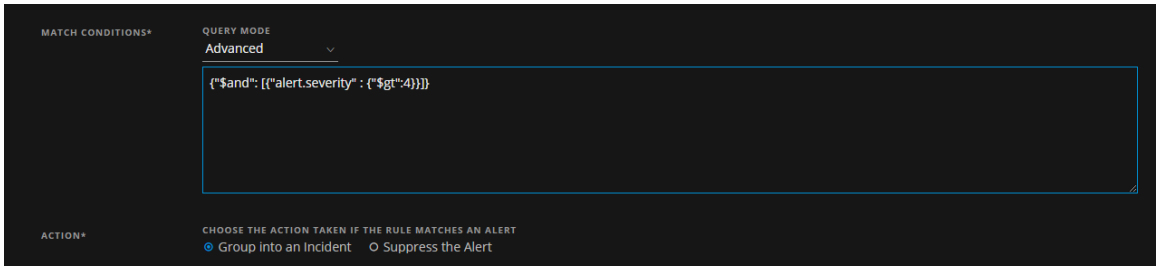
To access the Incident Rule Details view, do one of the following:

- To create a rule, go to  **(Configure) > Incident Rules** and click **Create Rule**.
- To edit a rule, go to  **(Configure) > Incident Rules** and click the link in the **Name** column for the rule that you want to update.

The Incident Rule Details view is displayed. The following figure shows the Incident Rule Details view in Rule Builder query mode.



In the Match Conditions section, if you select Advanced query mode, a field to enter advanced queries is available as shown in the following figure.



The following table describes the options available when creating or editing incident rules.

Section	Field	Description
Basic Settings	Enabled	Select to enable the rule.
	Name*	Name of the rule. *This is a required field.
	Description	A description of the rule to indicate which alerts get aggregated.
Match Conditions*	Query Mode	**Rule Builder**: Select the Rule Builder option if you want to build a query with various conditions that can be grouped. You can also have nested groups of conditions. In the Match Conditions, you can set the value to **All of these**, **Any of these**, or **None of these**. Depending on what you select, the criteria types specified in the Conditions and Group of conditions are matched to group the alerts. For example, if you set the match condition to **All of these**, alerts that match the criteria mentioned in the Conditions and Group Conditions are grouped into one incident. • Add a Condition to be matched by clicking the **Add Condition** button. • Add a Group of Conditions by clicking the **Add Group** button and add conditions by clicking the **Add Condition** button. You can include multiple Conditions and Groups of Conditions that can be matched as per criteria set and group the incoming alerts into incidents. **Advanced**: Select the Advanced query option if you want to use the advanced query builder. You can add a specific condition that needs to be matched as per the matching option selected. For example, you can type the criteria builder format **{"$and": [{"alert.severity" : {"$gt":4}}]}** to group alerts that have severity greater than 4. For advanced syntax, refer to http://docs.mongodb.org/manual/reference/operator/query/ or http://docs.mongodb.org/manual/reference/method/db.collection.find/
Action*	Choose the Action Taken if the Rule Matches the Alert	**Group into an Incident**: If enabled, the alerts that match the criteria set are grouped into an alert. **Suppress the Alert**: If enabled, the alerts that match the criteria are suppressed.

Section	Field	Description
Grouping Options	Group By*	The criteria to group the alerts in accordance with the specified alert fields. You can use a maximum of two fields to group the alerts. You cannot group alerts with fields that do not have values. When alerts are grouped on an alert field, all matching alerts containing the same meta key value for that field are grouped together in the same incident. (See the following **Group By Meta Key Mappings** table.)
	Time Window	The time range for grouping alerts. For example, if the time window is set to 1 hour, all alerts that match the criteria set in the Group By field and that arrive within an hour of each other are grouped into an incident.
	Advanced Grouping Options	Alerts may sometimes come into Respond missing fields that are expected and used for aggregation into incidents (as defined by incident aggregation rules). This can be due to missing parsers, overly generic matching conditions, unexpected data sources, or log sources with unexpected formats. Choose how to handle alerts missing the fields required by the configured incident aggregation rules: • Automatically group them into a single incident over the configured time window. • Do not group them and manually include these orphaned alerts into an incident when required.

Section	Field	Description
Incident Options	Title*	Title of the incident. You can optionally include placeholders in your title. Placeholders enable you to have different titles based on the attributes you grouped. If you do not use placeholders, all incidents created by the rule will have the same title. For example, if you grouped them according to the source, you can name the resulting Incident as Alerts for **${groupByValue1}**, and the incident for all alerts from NetWitness Endpoint would be named **Alerts for NetWitness Endpoint**.
	Summary	(Optional) Summary of the incident created by this rule.
	Categories	(Optional) Category of the incident created. An incident can be classified using more than one category.
	Assignee	(Optional) Name of the user assigned to the incident.
	Priority	**Average of Risk Score across all of the Alerts**: Takes the average of the risk scores across all the alerts to set the priority of the incident created. **Highest Risk Score available across all of the Alerts**: Takes the highest score available across all the alerts to set the priority of the incident created. **Number of Alerts in the time window**: Takes the count of the number of alerts in the time window selected to set the priority of the incident created. **Critical, High, Medium, and Low**: Specify the incident priority threshold of the matched incidents. The defaults are: • Critical: 90 • High: 50 • Medium: 20 • Low: 1 For example, with the Critical priority set to 90, incidents with a risk score of 90 or higher are assigned a Critical priority for this rule.

Group By Meta Key Mappings

When alerts are grouped on an alert field, all matching alerts containing the same meta key value for that field are grouped together in the same incident. For example, if you select the Group By field value **Destination Host**, it uses the mapped meta key `alert.groupby_host_dst`. All alerts with the same meta key value for `alert.groupby_host_dst` are grouped together in the same incident.

The following table shows the mapped meta keys for the Group By field selections.

Group By Field Value	Mapped Meta Key
Alert Name	alert.name
Alert Rule Id	alert.signature_id
Alert Type	alert.groupby_type
Date Created	alert.timestamp
Destination Country	alert.groupby_destination_country
Destination Domain	alert.groupby_domain_dst
Destination Host	alert.groupby_host_dst
Destination IP Address	alert.groupby_destination_ip
Destination Port	alert.groupby_destination_port
Destination User Account	alert.groupby_user_dst
Detector IP Address	alert.groupby_detector_ip
Domain	alert.groupby_domain
Domain for Suspected C&C	alert.groupby_c2domain
File Analysis	alert.groupby_analysis_file
Filename	alert.groupby_filename
File MD5 Hash	alert.groupby_data_hash
Risk Score	alert.risk_score
Service Analysis	alert.groupby_analysis_service
Session Analysis	alert.groupby_analysis_session
Severity	alert.severity
Source	alert.source
Source Country	alert.groupby_source_country
Source Domain	alert.groupby_domain_src
Source Host	alert.groupby_host_src
Source IP Address	alert.groupby_source_ip
Source User Account	alert.groupby_user_src
Source Username	alert.groupby_source_username
User Account	alert.groupby_username

Incident Email Notification Settings View

The Incident Email Notification Settings view enables you to send email notifications when incidents are created or updated to SOC Managers and the Analysts assigned to the incidents.

Note: The information in this topic applies to NetWitness 11.1 and later.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure an email server.	Refer to "Configure the Email Settings as Notification Server" in the *System Configuration Guide*. (To access these settings, click the **Email Server Settings** link and go to the Servers tab or alternatively go to (Admin) > System > Global Notifications > Servers tab.)
Incident Responders, Analysts, Content Experts, SOC Manager	Configure email notifications for when an incident is created or updated.	Configure Incident Email Notification Settings

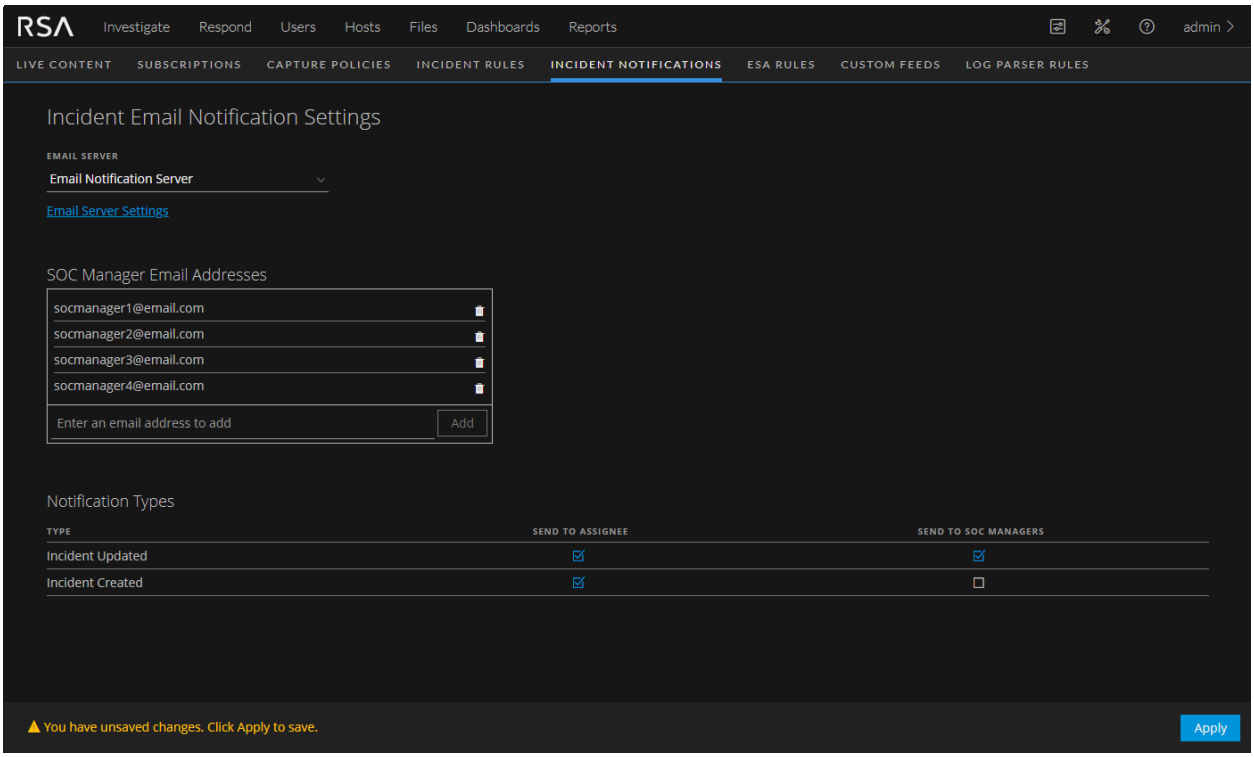
Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Related Topics

- Incident Rules View

Quick Look

To access the incident email notification settings, go to **(Configure) > Incident Notifications**.
The Incident Email Notification Settings view is displayed.



The following table lists the incident email notification settings.

Setting	Description
Email Server	Specifies the Email server that will send the email notifications.
Email Server Settings	Allows you to configure an Email server if the one you want to use for notifications is not listed. Clicking the **Email Server Settings** link goes to (Admin) > System > Global Notifications. Configure the email server on the Servers tab. For instructions, refer to "Configure the Email Settings as Notification Server" in the *System Configuration Guide*.
SOC Manager Email Addresses	Lists the SOC Manager email addresses that receive email notifications when you select **Send to SOC Manager** in the Notification Types section. You can add and remove email addresses as needed.
Notification Types - Incident Created	Specifies who should receive an email notification when an incident is created. • **Send to Assignee**: When an incident is created, an email is sent to the Analyst assigned to the incident. • **Send to SOC Manager**: When an incident is created, an email is sent to all of the addresses listed in the SOC Manager Email Addresses list.

Setting	Description
Notification Types - Incident Updated	Specifies who should receive an email notification when an incident is created. • **Send to Assignee**: When an incident is updated, an email is sent to the Analyst assigned to the incident. • **Send to SOC Manager**: When an incident is updated, an email is sent to all of the addresses listed in the SOC Manager Email Addresses list.
Apply	Applies changes made to the incident notification settings. Changes to these settings take effect immediately.

Note: If user email address information is updated in the  (Admin) > Security > Users tab, it can take up to two minutes for the new email changes to take effect. Any incident creation or incident update email notifications sent during this time go to the old email address.

Aggregation Rules Tab (11.0 and earlier)

The Aggregation Rules tab enables you to create and manage aggregation rules for automating the incident creation process. NetWitness provides 11 preconfigured rules. You can add to and adjust these rules for your own environment.

Note: This topic applies to NetWitness version 11.0 and earlier.

What do you want to do?

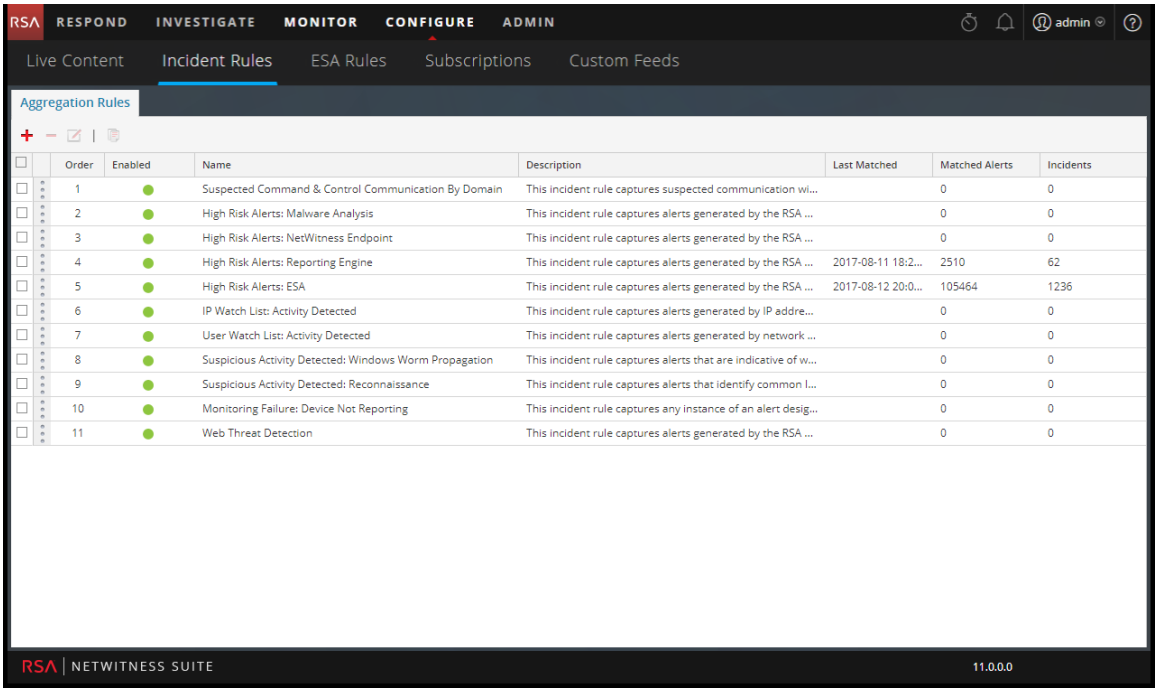
Role	I want to ...	Show me how
Analyst, Content Expert, SOC Manager	Create an aggregation rule.	Step 3. Enable and Create Incident Rules for Alerts
Incident Responders, Analysts, Content Experts, SOC Manager	View the results of my aggregation rule (View Detected Threats).	See "Responding to Incidents" in the *NetWitness Respond User Guide*.

Related Topics

- New Rule Tab (11.0 and earlier)

Quick Look

To access the Aggregation Rules tab, go to **Configure > Incident Rules > Aggregation Rules** tab.



The Aggregation Rules tab consists of a list and toolbar.

Aggregation Rules List

The following table describes the columns in the Aggregation Rules list.

Column	Description
Select	Enables you to select a rule in order to take an action, such as Clone or Delete.
Order	Shows the order in which the rule is placed. The rule order determines which rule takes effect if the criteria for multiple rules match the same alert. If two rules match an alert, only the rule with the highest priority is evaluated.
Name	Displays the name of the rule.
Enabled	Shows whether the rule is enabled or not. The specifies the rule is enabled.
Description	Displays the description of the rule.
Last Matched	Displays the time when an alert was successfully matched with the rule. This value is reset once a week.
Matched Alerts	Displays the number of matched alerts. This value is reset once a week. To change the setting, see Set a Counter for Matched Alerts and Incidents.
Incidents	Displays the number of incidents created by the rule. This value is reset once a week. To change the setting, see the Set a Counter for Matched Alerts and Incidents.

Aggregation Rules Toolbar

The following table shows the operations that can be performed in the Aggregation Rules tab.

Option	Description
+	Allows you to add a new rule.
	Allows you to edit a rule.
−	Allows you to delete a rule.
	Allows you to duplicate a rule.

New Rule Tab (11.0 and earlier)

The New Rules tab enables you to create custom aggregation rules for automating the incident creation process. This topic describes the information required when creating a new rule.

Note: This topic applies to NetWitness version 11.0 and earlier.

What do you want to do?

Role	I want to ...	Show me how
Analyst, Content Expert, SOC Manager	Create an aggregation rule.	Step 3. Enable and Create Incident Rules for Alerts
Incident Responders, Analysts, Content Experts, SOC Manager	View the results of my aggregation rule (View Detected Threats).	See "Responding to Incidents" in the *NetWitness Respond User Guide*.

Related Topics

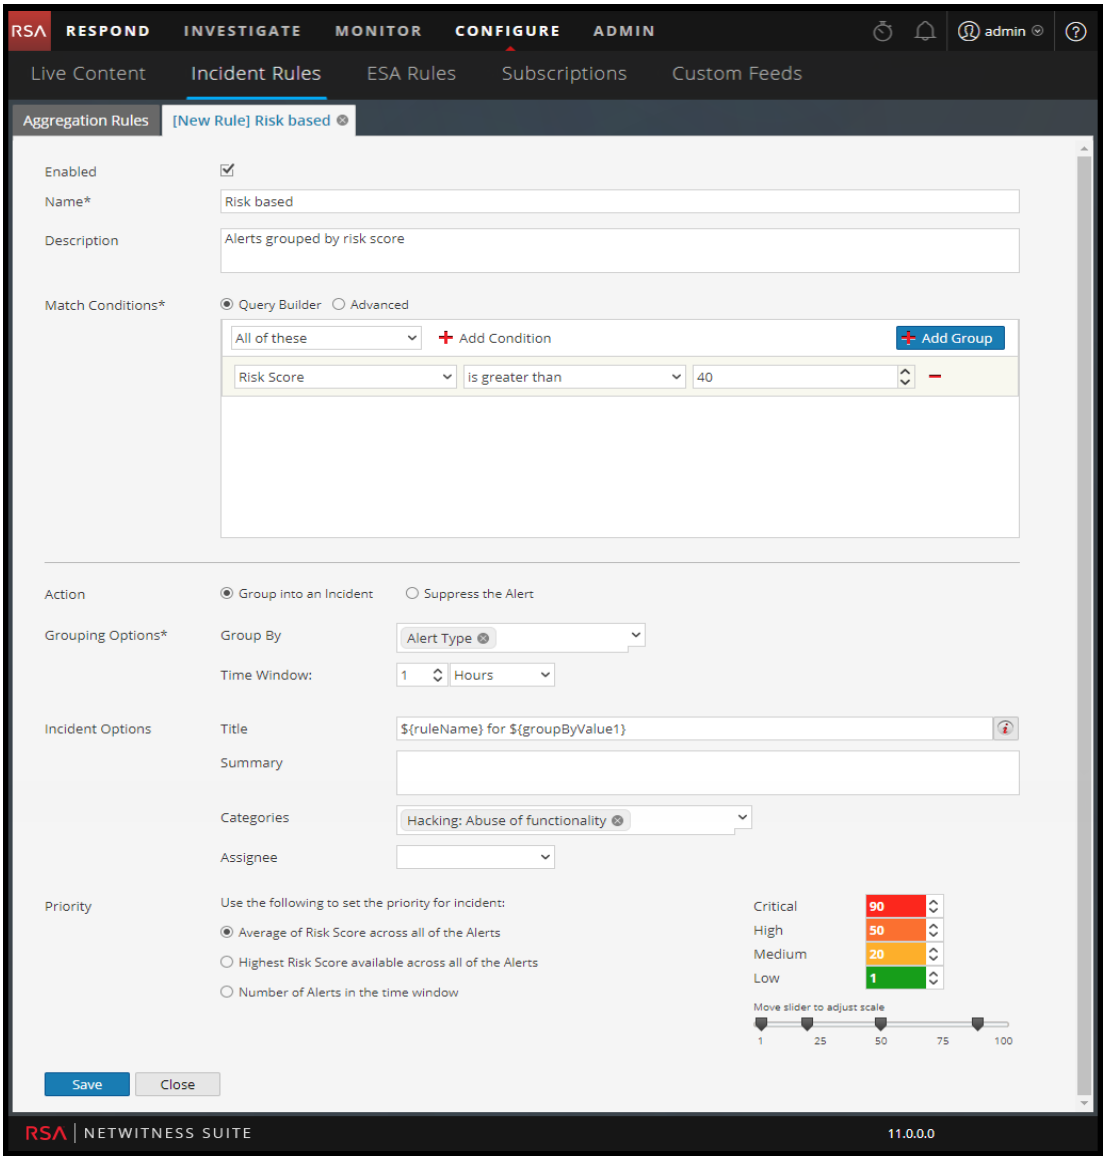
- Aggregation Rules Tab (11.0 and earlier)

Quick Look

To access the New Rule tab view:

1. Go to **Configure > Incident Rules > Aggregation Rules** tab.
2. Click +.

The **New Rule** tab is displayed.



The following table describes the options available when creating customized aggregation rules.

Field	Description
Enabled	Select to enable the rule.
Name*	Name of the rule. *This is a required field.
Description	A description for the rule to give an idea about what alerts get aggregated.

Field	Description
Match Conditions*	**Query Builder** - Select if you want to build a query with various conditions that can be grouped. You can also have nested groups of conditions. Match Conditions - You can set the value to **All of these**, **Any of these**, or **None of these**. Depending on what you select, the criteria types specified in the Conditions and Group of conditions are matched to group the alerts. **For example**, if you set the match condition to **All of these**, alerts that match the criteria mentioned in the Conditions and Group Conditions are grouped into one incident. • Add a Condition to be matched by clicking **Add Condition**. • Add a Group of Conditions by clicking **Add Group** and adding conditions by clicking **Add Condition**. You can include multiple Conditions and Groups of Conditions that can be matched as per criteria set and group the incoming alerts into incidents.
	Advanced - Select if you want to add an advanced query builder. You can add a specific condition that needs to be matched as per the matching option selected. **For example:** you can type the criteria builder format **{"$and": [{"alert.severity" : {"$gt":4}}]}** to group alerts that have severity greater than 4. For advanced syntax, refer to http://docs.mongodb.org/manual/reference/operator/query/ or http://docs.mongodb.org/manual/reference/method/db.collection.find/
Action	**Group into an Incident** - If enabled, the alerts that match the criteria set are grouped into an alert. **Suppress the Alert** - If enabled, the alerts that match the criteria are suppressed.
Grouping Options*	**Group By:** The criteria to group the alerts as per the specified category. You can use a maximum of two attributes to group the alerts. You can group the alerts with one or two attributes. You can no longer group alerts with attributes that do not have values (empty attributes). Grouping on an attribute means that all matching Alerts containing the same value for that attribute are grouped together in the same incident. **Time Window:** The time range specified to group alerts. For example if the time window is set to 1 hour, all alerts that match the criteria set in Group By field and that arrive within an hour of each other are grouped into an incident.

Field	Description
Incident Options	**Title -** (Optional) Title of the incident. You can provide placeholders based on the attributes you grouped. Placeholders are optional. If you do not use placeholders, all Incidents created by the rule will have the same title. For example, if you grouped them according to the source, you can name the resulting Incident as Alerts for **${groupByValue1}**, and the incident for all alerts from NetWitness Endpoint would be named **Alerts for NetWitness Endpoint**.
	Summary - (Optional) Summary of the incident.
	Category - (Optional) Category of the incident created. An incident can be classified using more than one category.
	Assignee - (Optional) Name of the assignee to whom the incident is assigned to.
Priority	**Average of Risk Score across all of the Alerts** - Takes the average of the risk scores across all the alerts to set the priority of the incident created. **Highest Risk Score available across all of the Alerts** - Takes the highest score available across all the alerts to set the priority of the incident created. **Number of Alerts in the time window** - Takes the count of the number of alerts in the time window selected to set the priority of the incident created. **Critical, High, Medium, and Low** - Specify the incident priority threshold of the matched incidents. The defaults are: • Critical: 90 • High: 50 • Medium: 20 • Low: 1 For example, with the Critical priority set to 90, incidents with a risk score of 90 or higher will be assigned a Critical priority for this rule. You can change these defaults by manually changing the priorities or by moving the slider under **Move slider to adjust scale**.

NetWitness® Platform XDR

Version 12.0.0.0

NetWitness Respond User Guide

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

November, 2022

Contents

NetWitness Respond Process **8**

NetWitness Respond Workflow 9

Responding to Incidents **10**

Respond Persist Data 12

Customizing Respond Persist Data 13

Working with Incident Filters 14

Persist and Suspend Persist Events 14

Change Event Retention 15

Responding to Incidents Workflow 17

Review Prioritized Incident List 17

View the Incidents List 18

Filter the Incident List 19

Remove My Filters from the Incidents List View 22

Save the Current Incidents Filter 22

Update a Saved Incidents Filter 23

Delete a Saved Incidents Filter 23

View My Incidents 23

Find an Incident 24

Sort the Incidents List 25

View Unassigned Incidents 26

Assign Incidents to Myself 26

Unassign an Incident 28

Determine which Incidents Require Action **31**

View Incident Details 32

View Basic Summary Information about the Incident 34

View the Indicators and Enrichments 37

View and Study the Events 38

View C2 Enrichment Information for Suspected C&C Incidents 41

View and Study the Entities Involved in the Events on the Nodal Graph 43

Nodal Graph Behaviors and Characteristics 46

Select Node Types to View on the Nodal Graph 49

Filter the Data in the Incident Details View 52

View the Tasks Associated with an Incident 53

View Incident Notes 54

Find Related Indicators 54

Add Related Indicators to the Incident 56
Investigate the Incident 58
View Contextual Information 59
Add an Entity to a Whitelist 61
Create a List 63
View the Reputation Status of a File Hash 64
Pivot to the Investigate > Events View 66
Pivot to the Hosts or Files View 66
Pivot to NetWitness Endpoint Thick Client 67
Pivot to Archer 67
View Event Analysis Details for Indicators 68
Migration Considerations 68
View User Entity Behavior Analytics for Indicators 72
Document Steps Taken Outside of NetWitness 72
View the Journal Entries for an Incident 72
Add a Note 74
Delete a Note 75
Escalate or Remediate the Incident 76
Send an Incident to Archer 76
View All Incidents Sent to Archer 79
Update an Incident 80
Change Incident Status 80
Status Change Workflow 81
Change Events Retention 87
Obtain Retention Usage Details 87
Change Incident Priority 88
Assign Incidents to Other Analysts 91
Rename an Incident 93
View All Incident Tasks 95
Filter the Tasks List 96
Remove My Filters from the Tasks List 98
Create a Task 99
Find a Task 103
Modify a Task 104
Delete a Task 108
Close an Incident 110
Incident Response Use Case Examples 111
Use Case #1: UEBA Anomalous User Activity 111
Use Case #2: Encoded Webshells Detected 114

Reviewing Alerts **119**
View Alerts 119
Filter the Alerts List 121
Remove My Filters from the Alerts List 124
Save the Current Alerts Filter 124
Update a Saved Alerts Filter 125
Delete a Saved Alerts Filter 125
View Alert Summary Information 126
View Event Details for an Alert 127
Investigate Events 131
View Contextual Information 131
Add an Entity to a Whitelist 133
Create a Whitelist 134
Pivot to the Investigate > Navigate View 134
Pivot to the Hosts or Files View 135
Pivot to Endpoint Thick Client 135
Pivot to Archer 135
Create an Incident Manually 136
Add Alerts to an Incident 140
Delete Alerts 142
Review Endpoint Alerts using Process Tree **143**
Process Details Section Values 144
Event Details Section Values 145
NetWitness Respond Reference Information **147**
Incidents List View 148
Workflow 148
What do you want to do? 149
Related Topics 149
Quick Look 149
Incidents List View 150
Incidents List 151
Incident Filters Panel 153
Incident Overview Panel 155
Toolbar Actions 158
Incident Details View 160
Workflow 160
What do you want to do? 161
Related Topics 162
Quick Look 163
Overview Panel 165

Indicators Panel 167
Related Indicators Panel 168
History Panel 169
Events 172
User Entity Behavior Analytics 173
Nodal Graph 174
Nodes 175
Arrows 176
Events List 178
Event Details 181
Journal Panel 183
Tasks Panel 184
Toolbar Actions 185
Alerts List View 187
Workflow 187
What do you want to do? 187
Related Topics 188
Quick Look 188
Alerts List 190
Alert Filters Panel 192
Alert Overview Panel 195
Toolbar Actions 196
Alert Details View 198
Workflow 198
What do you want to do? 198
Related Topics 199
Quick Look 199
Overview Panel 200
Events - Process Tree View 200
Events List 201
Event Details 202
Event Details 202
Event Source or Destination Device Attributes 203
Event Source or Destination User Attributes 204
Toolbar Actions 204
Tasks List View 205
What do you want to do? 205
Related Topics 205
Quick Look 205
Tasks List 206

Task Filters Panel 208
Task Overview Panel 210
Toolbar Actions 211
Add/Remove from List Dialog 212
What do you want to do? 212
Related Topics 212
Quick Look 213
Context Lookup Panel - Respond View 216
What do you want to do? 216
Related Topics 216
Contextual Information Displayed in the Context Lookup Panel 217
Lists Tab 219
Archer Tab 220
Active Directory Tab 221
NetWitness Endpoint Tab 222
Alerts Tab 224
Incidents Tab 225
File Reputation Tab 226
TI Tab 227
REST API Tab 228

NetWitness Respond Process

NetWitness Respond collects alerts from multiple sources and provides the ability to group them logically and start an Incident Respond workflow to investigate and remediate the security issues raised. NetWitness Respond enables you to configure rules that automatically aggregate alerts into incidents. Alerts are normalized by the system to a common format to provide users with a consistent view for the rule criteria regardless of the data source. You can build query criteria based on the alert data with the ability to query on fields that are common as well as specific to data sources.

The rule engine allows you to group similar alerts together into an incident so that the investigation and remediation workflow can be shared across a set of similar alerts. You can create rules that can group alerts into incidents depending on a common value they share for one or two attributes (for example, source hostname) or if they are reported within a limited time window (for example, alerts that are within four hours of each other).

If an alert matches a rule, an incident is created using the criteria. As new alerts are ingested, if an existing incident was already created that matched those criteria, and that incident is not "in progress" yet, the new alerts continue to be added to the same incident. If there is no existing incident for the grouped value (for example, the specific hostname) or the time window, a new incident is created and the alert is added to it.

You can have multiple incident rules. The rules can either group alerts into incidents or suppress alerts from being matched by any rule. The rules are ranked top-to-bottom and only the first rule to match an incoming alert is used to include that alert in an incident. The incidents provide a context for the alerts, provide tools to record the investigation status, and track the progress of associated tasks.

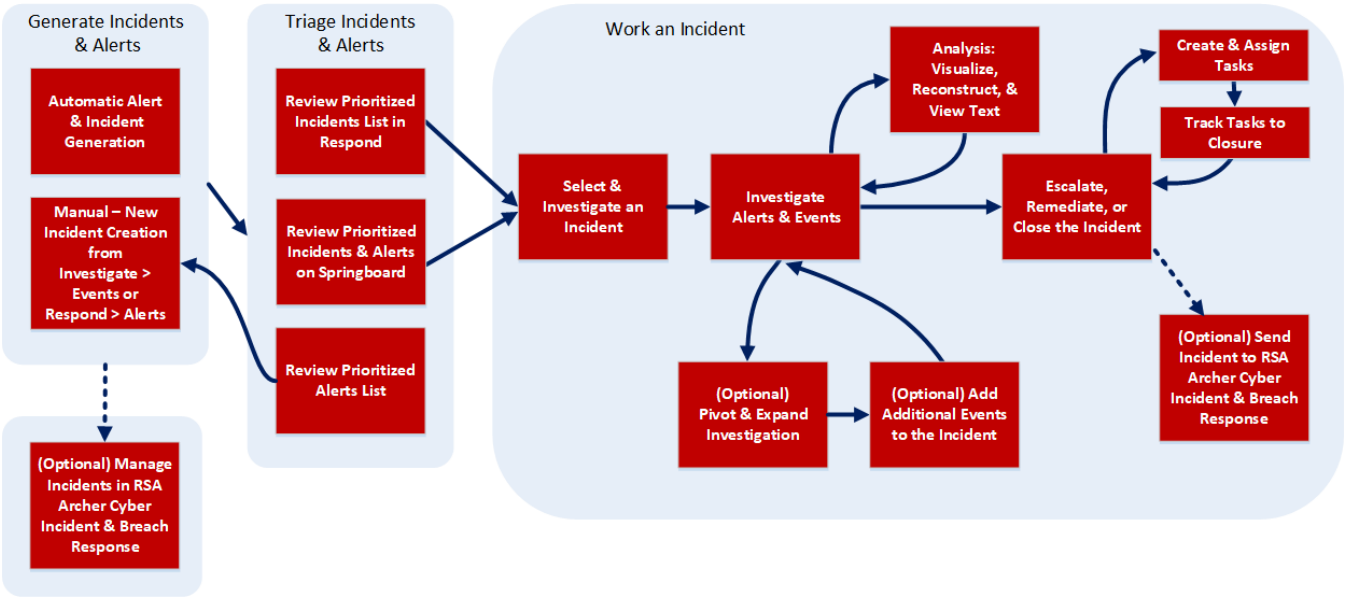
The stages in the NetWitness Respond process are:

- **Generate Incidents & Alerts**
  - Automatic Alert & Incident Generation
  - Manual - New Incident Creation from Investigate > Events or Respond > Alerts
  - (Optional) Manage Incidents in Archer Cyber Incident and Breach Response (If you manage incidents in Archer instead of in NetWitness Respond, the process ends here.)
- **Triage Incidents & Alerts**
  - Review Prioritized Incident List in Respond
  - Review Prioritized Incidents & Alerts on Springboard
  - Review Prioritized Alerts List
- **Work an Incident**
  - Select & Investigate Incident
  - Investigate Alerts & Events
    - Analysis: Visualize, Reconstruct, and View Text
    - (Optional) Pivot & Expand Investigation
    - (Optional) Add Additional Events to Incident

- Escalate, Remediate, or Close the Incident
  - Create & Assign Tasks
  - Track Tasks to Closure
  - (Optional) Send Incidents to Archer Cyber Incident & Breach Response. (In NetWitness version 11.2 and later, if Archer is configured as a data source in Context Hub, you can send incidents to Archer Cyber Incident & Breach Response.)

NetWitness Respond Workflow

The following figure shows the high-level NetWitness Respond workflow process.



Responding to Incidents

An *Incident* is a logically grouped set of alerts created automatically by the Incident Aggregation Engine and grouped by a specific criteria. An incident, available in the Respond view, allows an Analyst to triage, investigate, and remediate these groups of alerts. Incidents can be moved between users, notated, and explored using a nodal graph. Incidents allow users to ensure that they understand the full scope of an attack or event in their NetWitness system and then take action.

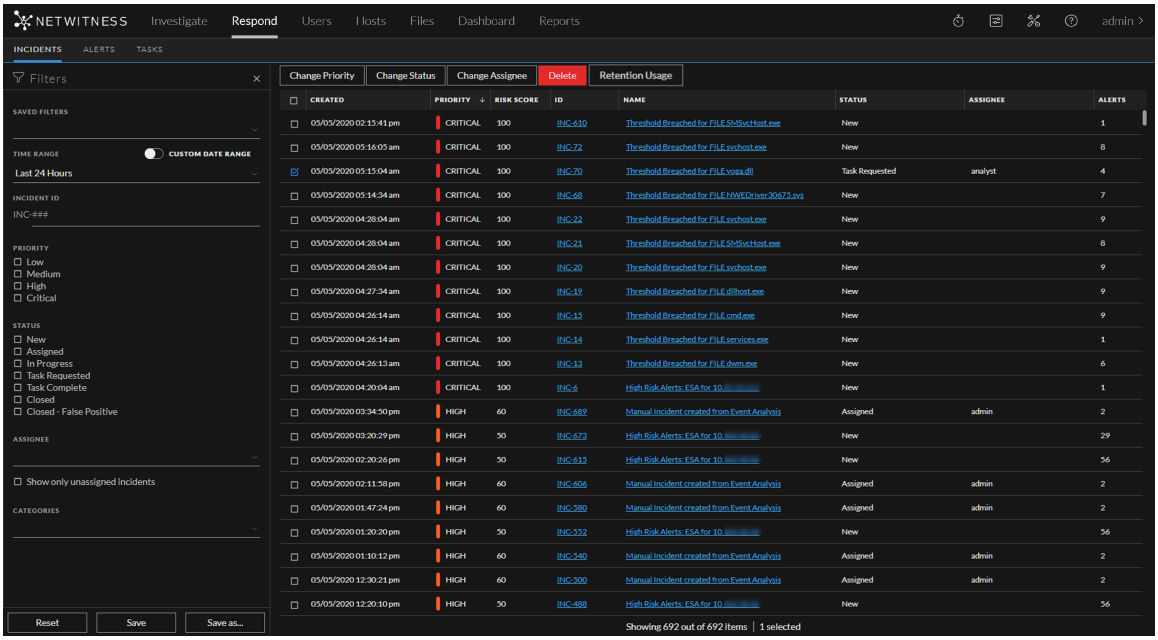
The **Respond** view is designed to help you quickly identify the ongoing issues in your network and work with other Analysts to quickly solve the issues.

The Respond view presents Incident Responders with a queue of incidents in severity order. When you take an incident from the queue, you receive relevant supporting data to help you investigate the incident. This enables you to determine the incident scope so you can escalate or remediate it as appropriate.

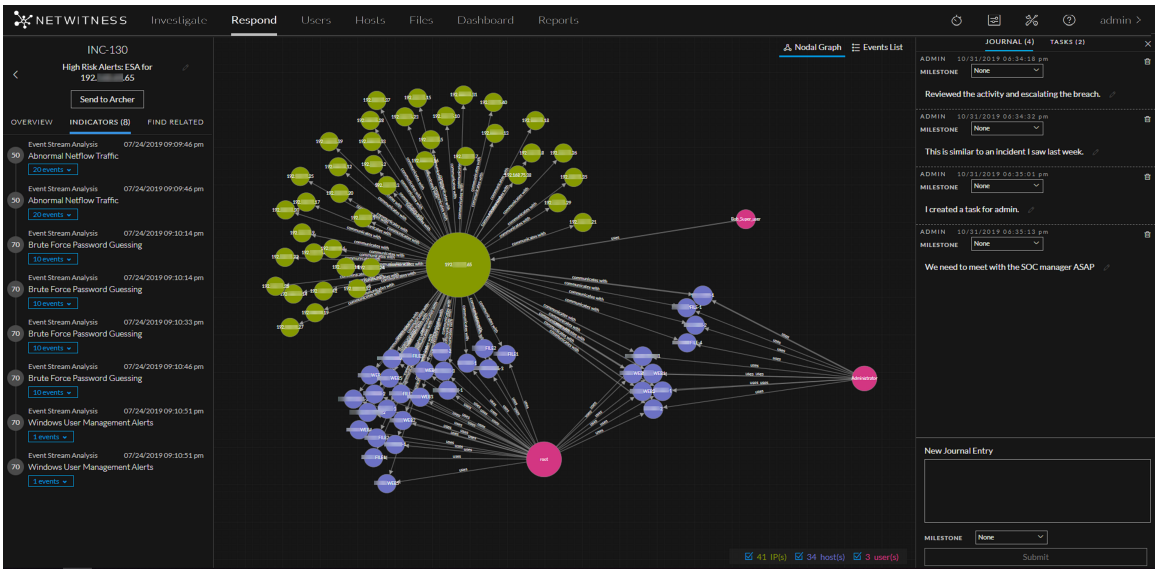
Within the Respond view, you can see Incidents, Alerts, and Tasks:

- **Incidents**: Enables you to respond to and manage incidents from start to finish.
- **Alerts**: Enables you to manage alerts from all sources received by NetWitness and create incidents from selected alerts.
- **Tasks**: Enables you to view and manage the complete list of tasks created for all incidents.

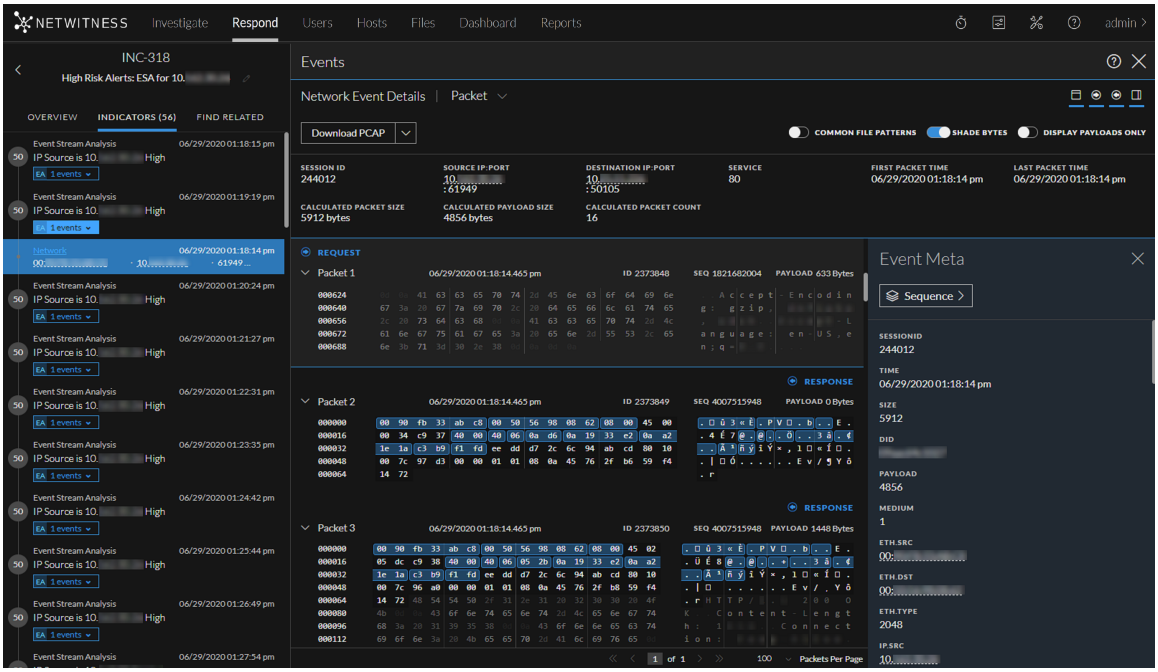
If you navigate to Respond > Incidents, you can see the Incidents List view and from there you can access the Incident Details view for a selected incident. These are the main views that you use to respond to incidents. The following figure shows the list of prioritized incidents in the **Incidents List** view.



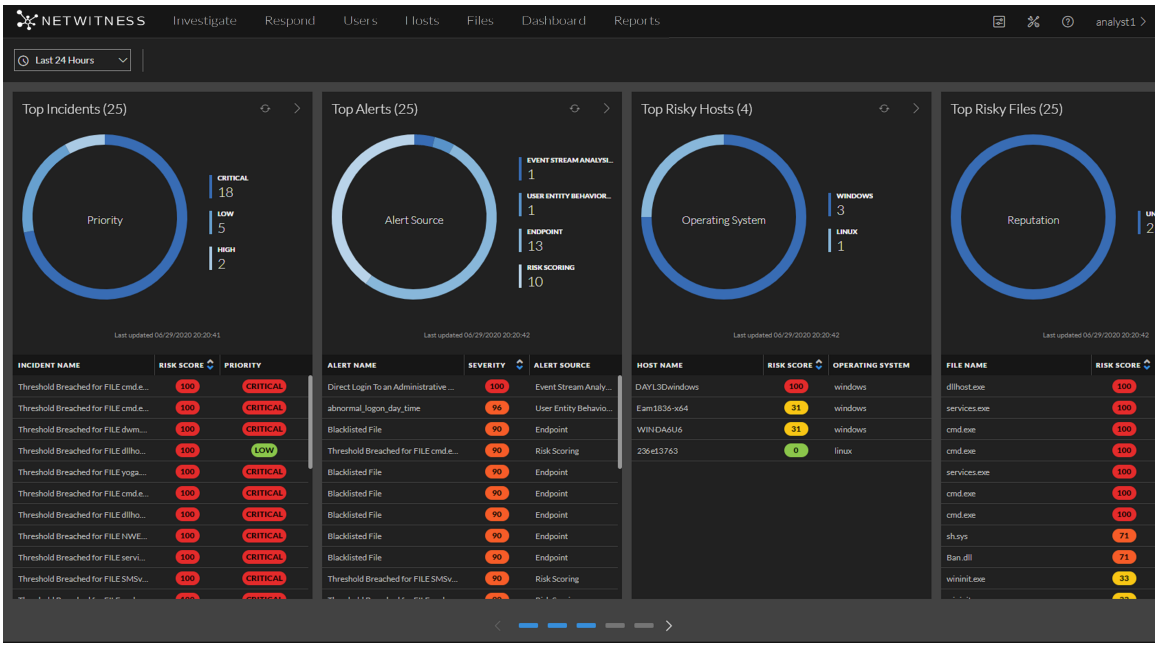
The next figure shows an example of details available in the **Incident Details** view.



The Respond view is designed to make it easy to evaluate incidents, contextualize that data, collaborate with other analysts, and pivot to a deep-dive investigation as needed. The following figure shows an example of an event analysis in the Incident Details view.



In NetWitness Version 11.4 and later, alerts and incidents are also displayed in the Springboard by default. Springboard is a landing page for analysts showing them all risks detected by the platform in a single place. For more information on the Springboard, see "Managing the Springboard" in the *NetWitness Platform Getting Started Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.



Respond Persist Data

The primary objective of this feature is to enable you to investigate events for a longer period of time. When investigating a set of alerts or events, it is important to have the assurance that the underlying data (raw and meta) will be available for both the length of time it takes to run the analysis, and for historical look up. From NetWitness Platform Version 11.6, you can persist events that are associated with particular incidents, thereby enabling you to view the incident in the future, regardless of its age. You can also add a new journal entry in the **JOURNAL** tab for the persisted events for future reference. The event data will always be available for viewing and reconstruction as long as the event is persisted, enabling you to easily refer back to details, even if the original event has rolled over from the NetWitness database.

Once you persist an event, the data is copied from the NetWitness database into a long term storage cache within the data source. You can persist or suspend persist events per incident or per alert within an incident or a specific event in an alert that belongs to an incident. The roll over in the NetWitness database does not impact the events that are already saved in the long term cache.

You can:

- You can persist or suspend persist events per incident or per alert within an incident or a specific event in an alert that belongs to an incident.
- Perform complete event reconstruction for persisted events even after the session is rolled over from the NetWitness database.

- Filter incidents which has persisted events.
- Suspend persist of all events associated with an incident, even if only a few events persisted in it.
- Remove all associated persisted events by deleting an alert or incident.

This topic contains the following basic Respond Persist Data procedures:

- Customizing Respond Persist Data
- Working with Incident Filters
- Persist and Suspend Persist Events
- Change Event Retention

Customizing Respond Persist Data

The persisted events are saved in the directory `/var/netwitness/pin-<servicetype>`, by default. You can manually change the event storage location from the default directory to any other directory, as per the requirement. You can increase the storage space as per your requirement by performing an Network File System (NFS) mount. For more information on how to perform an NFS mount, see Configure the Destination Using NFS.

To customize the persist directory:

1. Go to **Admin** >**Services**.
The **Services** page appears.

2. From the **Filter** pull down menu, select the concentrator and the log decoder services.
The concentrator and log decoder are listed in the **Services** page.

3. Go to **Actions** > **View** > **Explore**.
The **Explore** page appears.

4. Go to **sdk** > **config** .

The **Configuration** page appears.

The following table provides information on the default values for each parameter that are configured in the **Configuration** page.

S.No	Parameter	Default Value
1.	Long Term Cache Behavior (pin.cache.behavior)	fail-on-new
2.	Long Term Cache Directory (pin.cache.dir)	/var/netwitness/pin-<serviceType>***
3.	Long Term Cache Size (pin.cache.size)	10 GB

2. [Optional]*** In the **Long Term Cache Directory (pin.cache.dir)**, enter the path of the custom directory where you want to save the persisted events, if you do not want to use the default location. The default path is `/var/netwitness/pin-<serviceType>`.

Note: When you install NetWitness Platform for the first time or upgrade to the 11.6 version, the directories are pre-configured with default values.

Working with Incident Filters

You can filter incidents based on the persisted events.

To filter incidents:

1. Go to **Respond** > **INCIDENTS**.

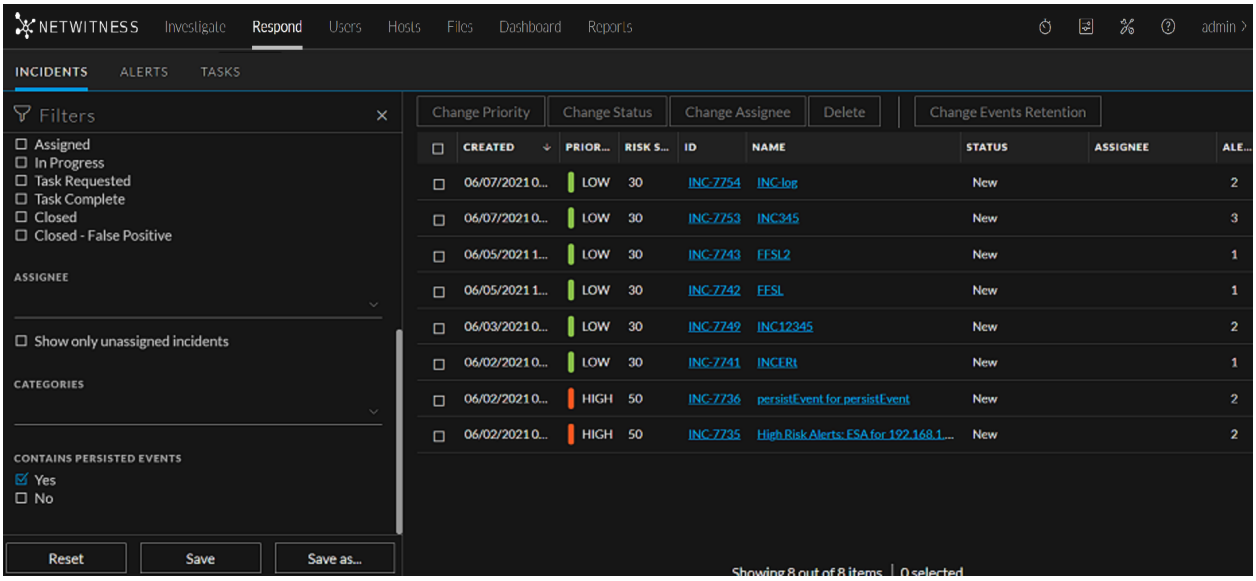
The **INCIDENT** page appears.

2. In the **FILTERS** tab, the **CONTAINS PERSISTED EVENTS** menu provides the filters to view incidents based on the persisted events.

a. Select **Yes** to view the incidents that contain persisted events.

b. Select **No** to view the incidents that do not contain persisted events.

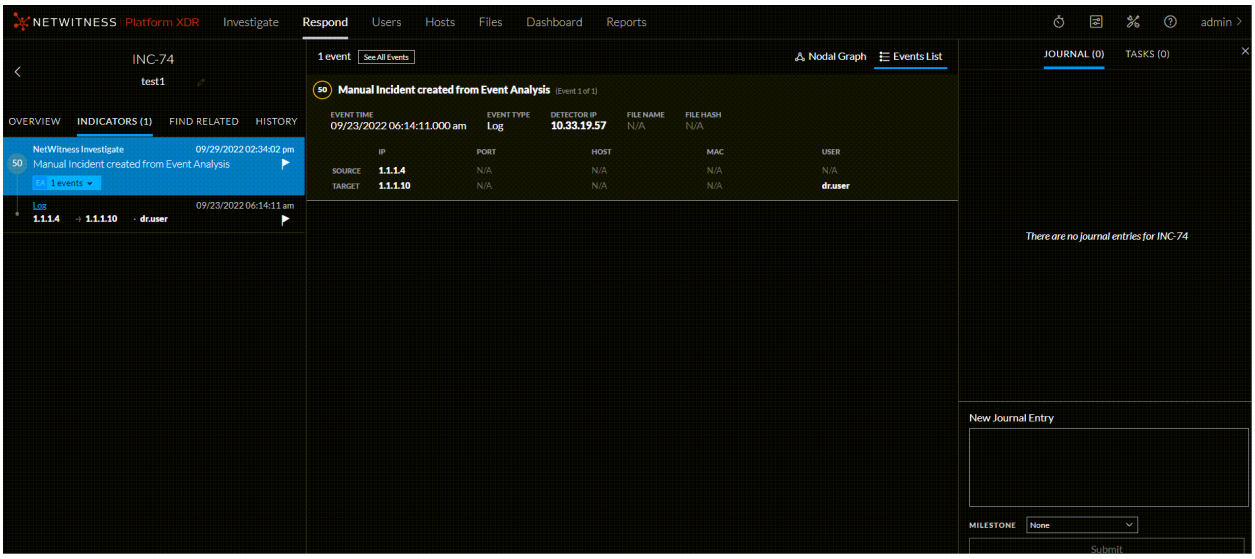
For more information on working with filters, see Filter the Incident List.



Persist and Suspend Persist Events

Persist an event to retain the event and thereby, copy the event data from the regular database into a long-term storage cache within the NetWitness source. Suspend persist the event to delete the event in the long-term storage cache.

- To persist an event, click the flag listed under each of the event. The selected flag is highlighted, and the message Event Persisted successfully appears.
- To suspend persist event, click a highlighted flag (persisted event) once again. The selected flag is no longer highlighted, and the message Suspended Event Persist successfully appears.
- To persist all events in an alert, click the flag at the alert level. The selected flag is highlighted along with all the flags associated with events of the alert.
- To suspend persist all events in an alert, click the flag at the alert level. The selected flag is highlighted along with all the flags associated with events of the alert.



Change Event Retention

Incidents can contain multiple persisted events, for which analysis has been completed. They can be suspended from persisting at a time using this feature.

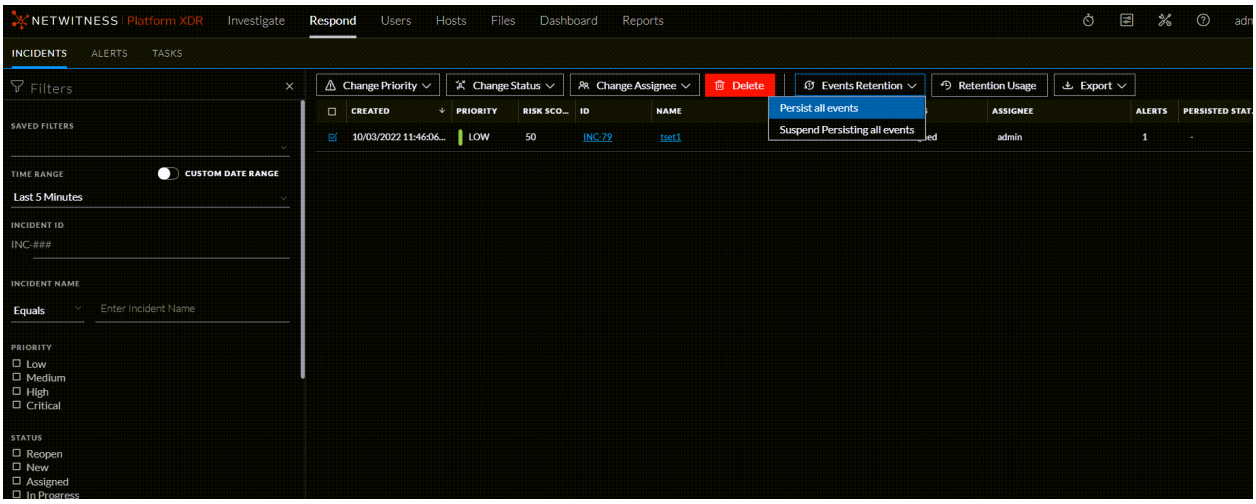
Note: A maximum of 1 incident can be persisted or suspend persisted at a time in the source NetWitness database. If you try to change event retention for more than 1 incidents, an error message will be displayed. This limitation is introduced to optimize and extract the best performance for this feature. This limit cannot be changed.

To change event retention:

1. Go to **Respond** > **INCIDENTS**.
   All the incidents are listed in the **INCIDENTS** page.
2. Click the selection check-box to select the incidents.
3. Click **Change Events Retention** > **Persist all events** or **Suspend Persist all events**.The **Confirm change in Retention** window appears.
4. Click **OK**.
   The events in the selected incident are either persisted or suspended from persisting.

Note: You cannot change the event retention for incidents that are in New or Closed state.

Note: Suspending persist of events in an incident from NetWitness will delete it from the long term cache of the source only. This may not be reversible if the original event data has rolled out in the source database. The events will be deleted permanently.



When you delete incidents or alerts, or suspend persist all the events in an incident, the action will always be successful. However, there are chances for the back-end function to fail due to, but not limited to the following reasons:

- Decoder/Concentrator outage
- Network disruption
- Service Outage

A data retention job is scheduled to run on a daily basis to clean up the events that failed from being deleted. To access the data retention job settings, see To change event retention:

The following table provides information on the parameters that are configured to run the job, which are enabled by default:

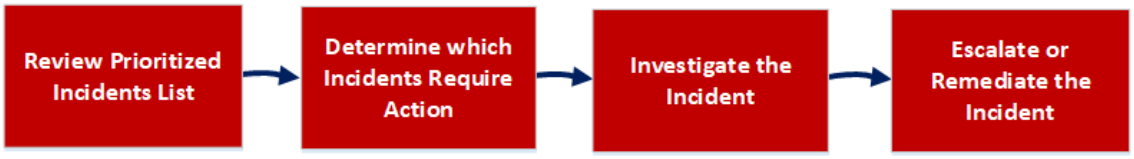
Parameter	Default Setting
`failure-count`	5
`persisted-events-retention-job-enabled`	true - default false.
`persisted-alerts-retention-period`	90 Days

Everyday at a preset time, the data retention job is executed. The default `failure-count` is set to 5 days. If the data retention job is unable to suspend persist an event from the back-end after 5 days, it deletes the event from the repository.

Note: The retention period of risk score context is the same as the retention period of alerts. For example, if the retention period for alerts is set to 90 days, then the retention period of risk score context is also set to 90 days automatically.

Responding to Incidents Workflow

This workflow shows the high-level process that Incident Responders use to respond to incidents in the Respond view.



First, you review the list of prioritized incidents, which shows basic information about each incident, and determine which incidents require action. You can click a link in an incident to get a clearer picture of the incident with supporting details in the Incident Details view. From there, you can further investigate the incident. You can then determine how to respond to the incident, by escalating or remediating it.

These are the basic steps for responding to an incident:

1. Review Prioritized Incident List
2. Determine which Incidents Require Action
3. Investigate the Incident
4. Escalate or Remediate the Incident

Review Prioritized Incident List

In the Respond view, you can view the list of prioritized incidents. The incident list shows both active and closed incidents.

This topic contains the following basic incident list procedures:

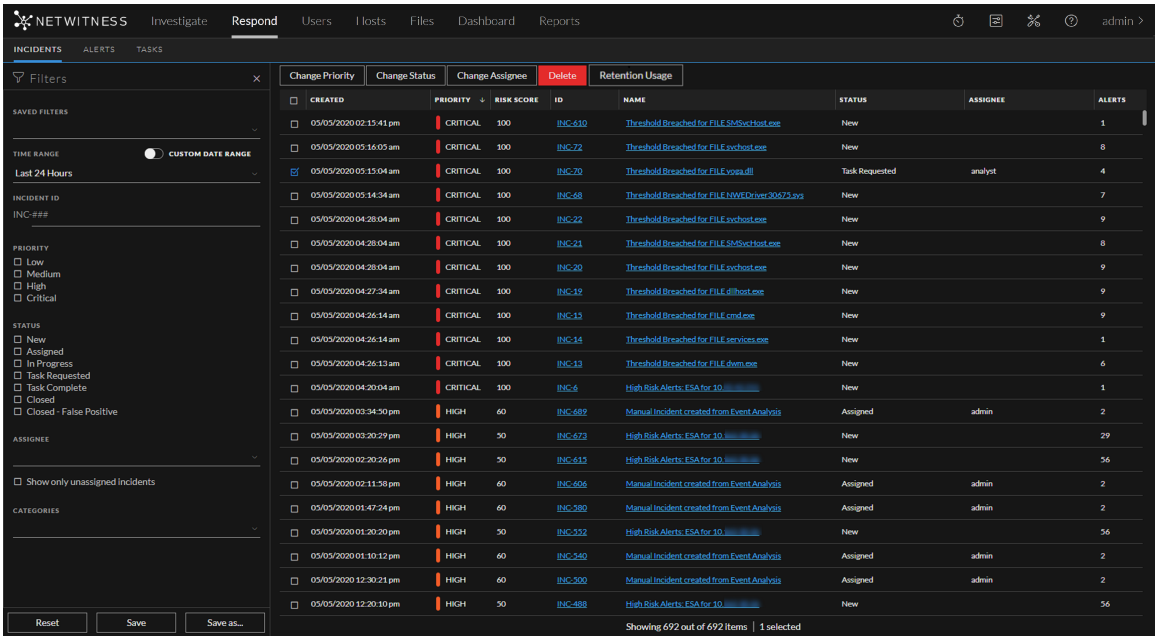
- View the Incidents List
- Filter the Incident List
- Remove My Filters from the Incidents List View
- Save the Current Incidents Filter
- Update a Saved Incidents Filter
- Delete a Saved Incidents Filter
- View My Incidents
- Find an Incident
- Sort the Incidents List
- View Unassigned Incidents

- Assign Incidents to Myself
- Unassign an Incident

View the Incidents List

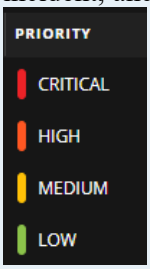
After logging in to NetWitness, most Incident Responders see the Respond view, which is set as the default view. If you have a different initial view, you can navigate to the Respond view.

1. Log in to NetWitness.
   The Respond view shows the list of incidents, also referred to as the Incident List view.



2. If you do not see the incidents list in the Respond view, go to **Respond > Incidents**.
3. Scroll through the incidents list, which shows basic information about each incident as described in the following table.

Column	Description
Created	Shows the creation date of the incident.

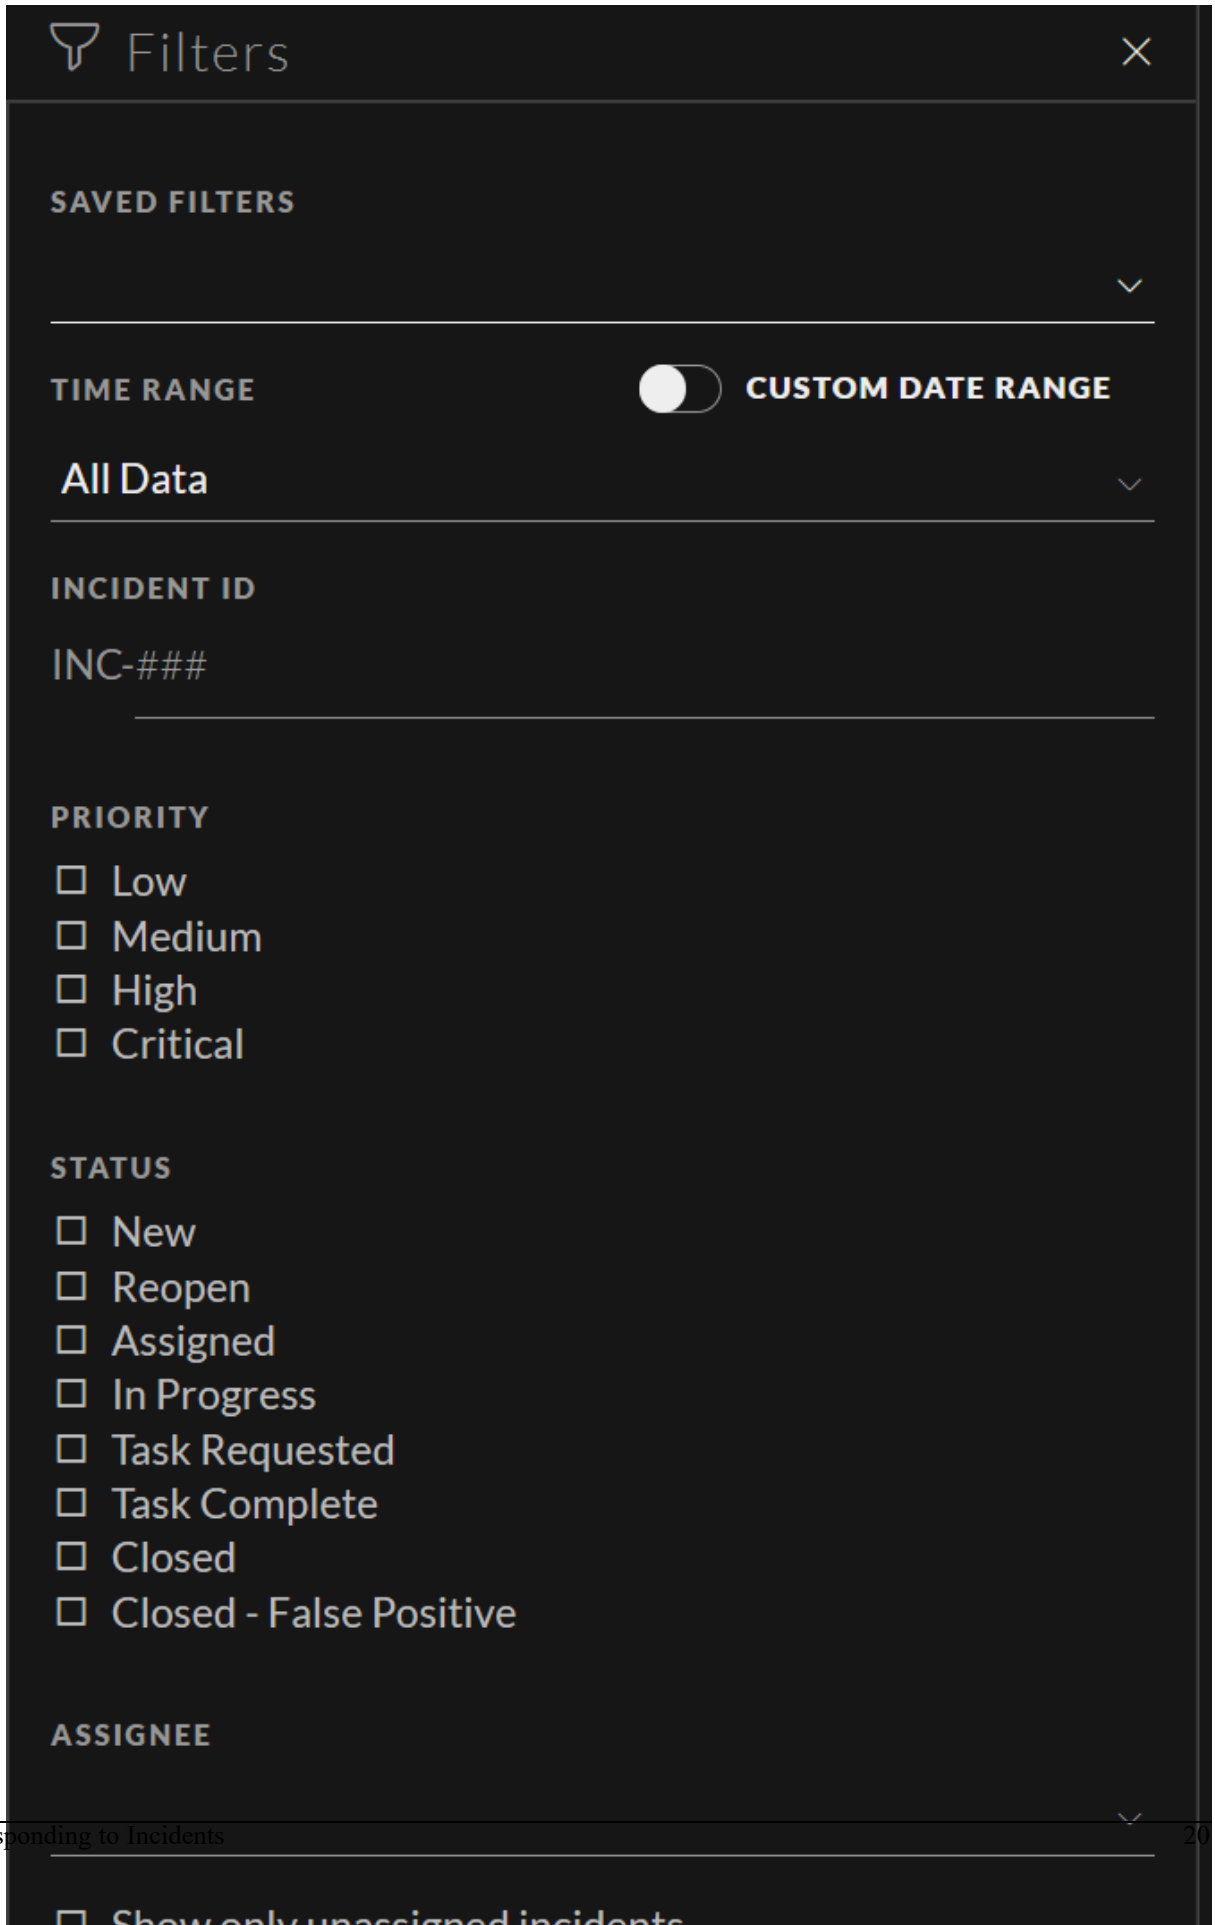
Column	Description
Priority	Shows the incident priority. Priority can be Critical, High, Medium, or Low. The Priority is color coded, where red indicates a **Critical** incident, orange represents a **High** risk incident, yellow indicates a **Medium** risk incident, and green represents a **Low** risk incident. For example: 
Risk Score	Shows the incident risk score. The risk score indicates the risk of the incident as calculated using an algorithm and is between 0-100. 100 is the highest risk score.
ID	Shows the automatically created incident number. Each incident is assigned a unique number that you can use to track the incident.
Name	Shows the incident name. The incident name is derived from the rule used to trigger the incident. Click the link to go to the Incident Details view for the selected incident.
Status	Shows the incident status. The status can be: **Reopen**, **New**, **Assigned**, **In Progress**, **Task Requested**, **Task Complete**, **Closed**, and **Closed - False Positive**.
Assignee	Shows the team member currently assigned to the incident.
Alerts	Shows the number of alerts associated with the incident. An incident may include many alerts. A large number of alerts might mean that you are experiencing a large-scale attack.

At the bottom of the list, you can see the number of incidents on the current page, the total number of incidents, and the number selected. For example: **Showing 1000 out of 1115 items | 3 selected**. The maximum number of incidents that you can view at one time is 1,000.

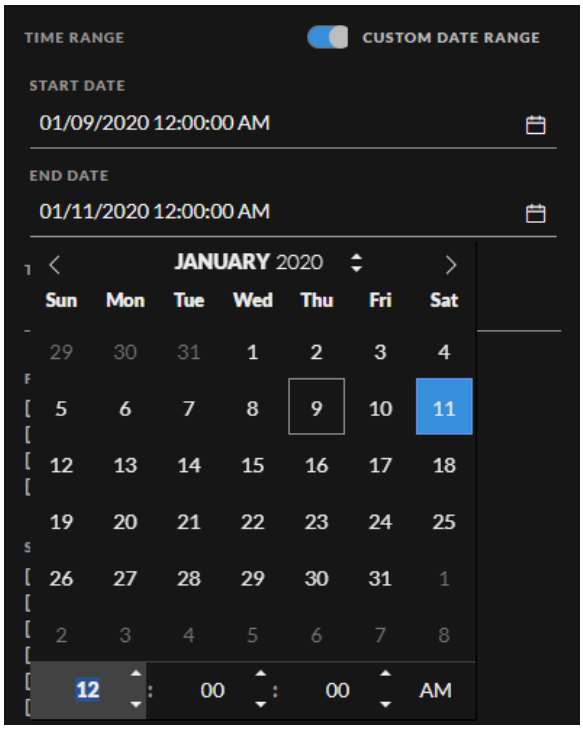
Filter the Incident List

The number of incidents in the Incidents List view can be very large, making it difficult to locate particular incidents. The Filter enables you to specify those incidents that you would like to view. You can also choose the timeframe when those incidents occurred. For example, you may want to view all of the new critical incidents created within the last hour.

1. Verify that the Filters panel appears to the left of the incidents list. If you do not see the Filters panel, in the Incident List view toolbar, click , which opens the Filters panel.


Filters
SAVED FILTERS
TIME RANGE
CUSTOM DATE RANGE
All Data
INCIDENT ID
INC-###
PRIORITY
Low
Medium
High
Critical
STATUS
New
Reopen
Assigned
In Progress
Task Requested
Task Complete
Closed
Closed - False Positive
ASSIGNEE

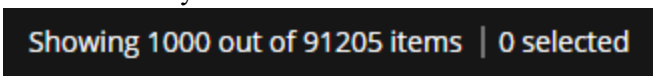
2. In the Filters panel, select one or more options to filter the incidents list:
   - **Time Range**: You can select a specific time period from the Time Range drop-down list. The time range is based on the creation date of the incidents. For example, if you select Last Hour, you can see incidents that were created within the last 60 minutes.
   - **Custom Date Range**: You can specify a specific date range instead of selecting a Time Range option. To do this, click the white circle in front of Custom Date Range to view the Start Date and End Date fields. Select the dates and times from the calendar.



   - **Incident ID**: Type the number of the incident that you would like to locate. For example, for INC-1050, type only the number "1050" to view the incident.
   - **Priority**: Select the priorities that you would like to view.
   - **Status**: Select one or more incident statuses. For example, select Closed - False Positive to view only false positive incidents, which were initially identified as suspicious, but then they were later found to be safe.
   - **Assignee**: Select the assignee or assignees of the incidents that you would like to view. For example, if you only want to view the incidents assigned to Cale or Stanley, select Cale and Stanley from the Assignee drop-down list. If you want to view incidents regardless of the assignee, do not make a selection under Assignee.
     (Available in version 11.1 and later) To view only unassigned incidents, select **Show only unassigned incidents**.
   - **Categories**: Select one or more categories from the drop-down list. For example, if you only want to view incidents classified with the Backdoor or Privilege abuse categories, select Backdoor and Privilege abuse.

- **Sent to Archer**: (In version 11.2 and later, if Archer is configured as a data source in Context Hub, you can send incidents to Archer Cyber Incident & Breach Response and this option will be available in NetWitness Respond.) To view incidents that were sent to Archer, select **Yes**. For incidents that were not sent to Archer, select **No**.
- **CONTAINS PERSISTED EVENTS**: Select a filter to view incidents based on the persisted events.

The incidents list shows a list of incidents that meet your selection criteria. You can see the number of incidents in your filtered list at the bottom of the incident list.



3. If you want to close the Filters panel, click . Your filters remain in place until you remove them.

Remove My Filters from the Incidents List View

NetWitness remembers your filter selections in the Incidents List view. You can remove your filter selections when you no longer need them. For example, if you are not seeing the number of incidents that you expect to see or you want to view all of the incidents in your incident list, you can reset your filters.

1. In the Incident List view toolbar, click .
   The Filters panel appears to the left of the incidents list.
2. At the bottom of the Filters panel, click **Reset**.

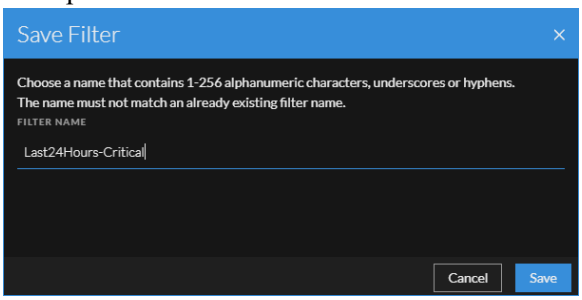
Save the Current Incidents Filter

Note: This option is available in NetWitness Platform Version 11.5 and later.

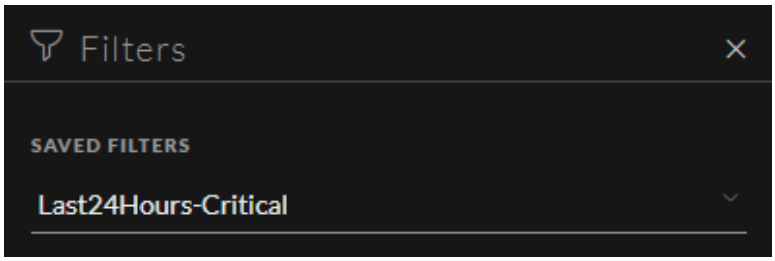
Saved filters provide a way for analysts to save and quickly apply specific filter conditions to the list of incidents. You can also use these filters to customize the Springboard landing page. For example, you may want to create a filter to show only critical incidents over the last 24 hours.

Saved filters are global. You can save a filter for other analysts to use and you can use any saved filter.

1. In the Filters panel, select one or more options to filter the incidents list. For example, in the Time Range field, select Last 24 Hours, and for Priority, select Critical.
2. Click **Save As** and in the **Save Filter** dialog, enter a unique name for the filter and save it, for example Last24Hours-Critical.



The filter is added to the **Saved Filters** list.



Update a Saved Incidents Filter

Note: This option is available in NetWitness Platform Version 11.5 and later.

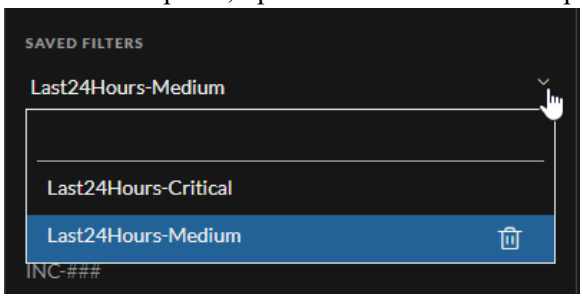
1. In the Filters panel **Saved Filters** drop-down list, select a saved filter.
2. Update your filter selections and click **Save**.

Delete a Saved Incidents Filter

Note: This option is available in NetWitness Platform Version 11.5 and later.

When a saved filter is no longer required, you can remove it from the saved filters list. Filters used in the Springboard cannot be deleted.

1. In the Filters panel, open the **Saved Filters** drop-down list.

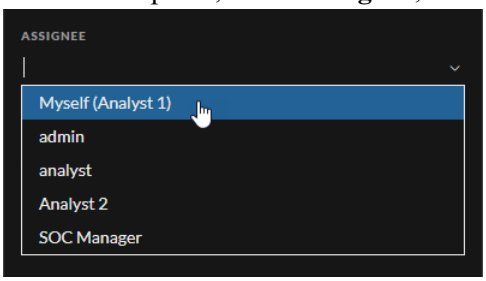


2. Next to the filter name, click  to delete it.

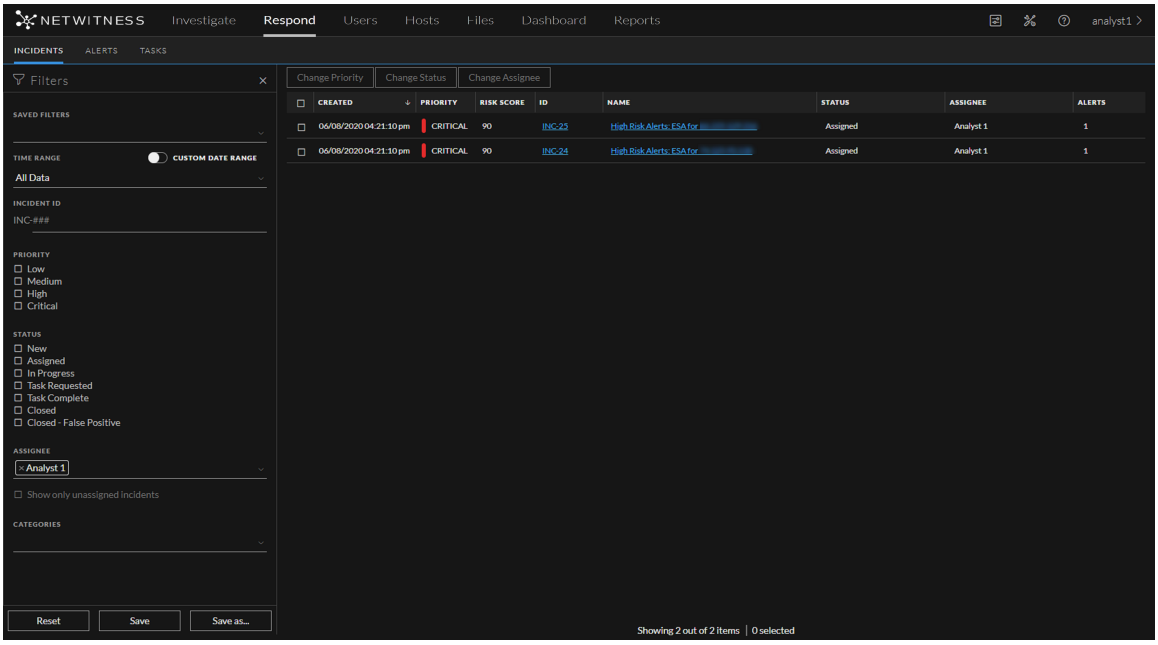
View My Incidents

You can view your incidents by filtering the incidents by your username.

1. If you cannot see the Filter panel, in the Incidents List view toolbar, click .
2. In the Filter panel, under **Assignee**, select **Myself (your full name)** from the drop-down list.



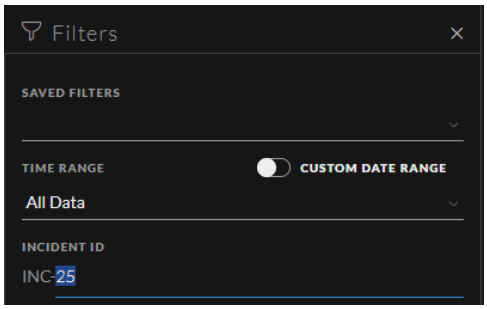
The incidents list shows the incidents that are assigned to you.



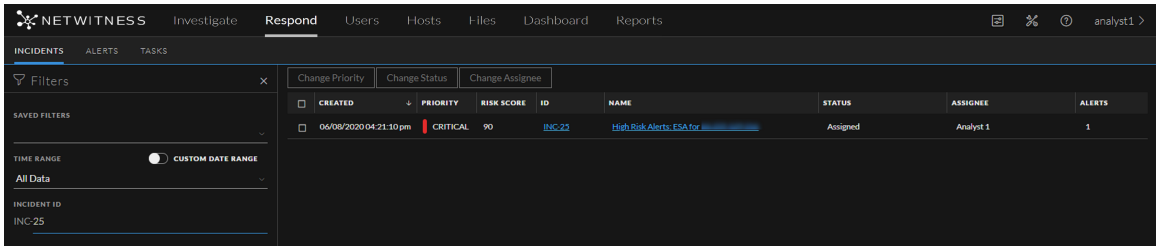
Find an Incident

If you know the Incident ID, you can quickly locate an incident using the Filter. For example, you may want to locate a specific incident out of thousands of incidents.

1. Go to **Respond > Incidents**.
   The Filters panel is located to the left of the incidents list. If you do not see the Filters panel, in the Incident Lists view toolbar, click , which opens the Filters panel.



2. In the **INCIDENT ID** field, type the Incident ID for an incident that you would like to locate, for example, type **25** for INC-25.
   The specified incident appears in your incident list. If you do not see any results, try resetting your filters.



Sort the Incidents List

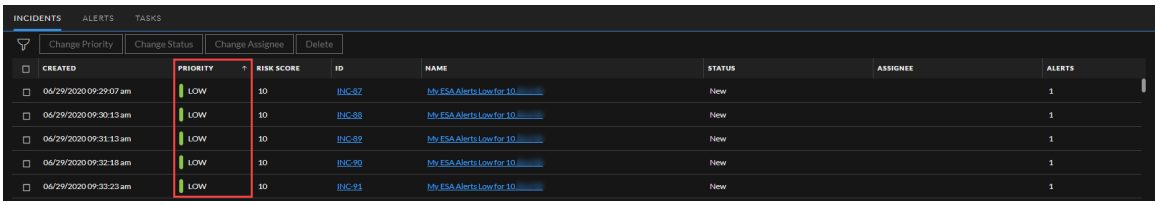
The default sort for the incidents list is by Created date in descending order (newest on the top).



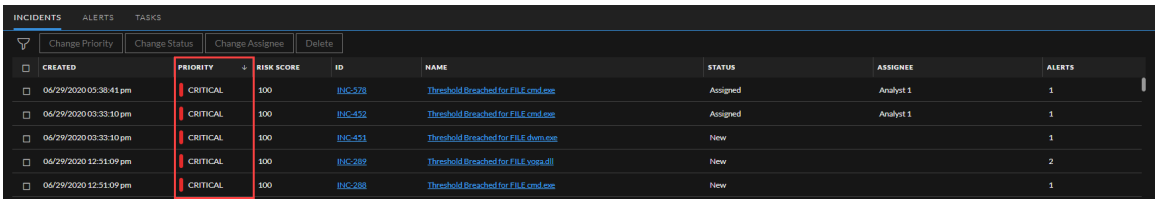
You can change the sort order of the incidents list by clicking a column header in the list.

For example, to prioritize the incidents, you can sort your view by clicking the Priority column header.

The following figure shows the incidents list sorted by Priority in ascending order (lowest priority on top).



To sort by Priority in descending order (highest priority on top), click the Priority column header again. The highest priority incidents are at the top as shown in the following figure.

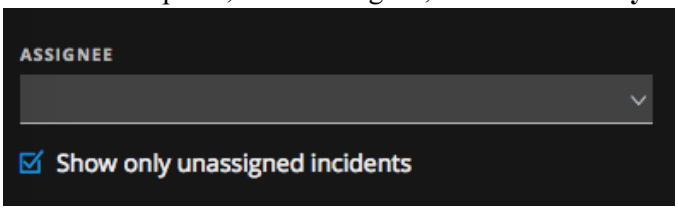


View Unassigned Incidents

Note: This option is available in NetWitness Platform Version 11.1 and later.

You can view unassigned incidents using the Filter.

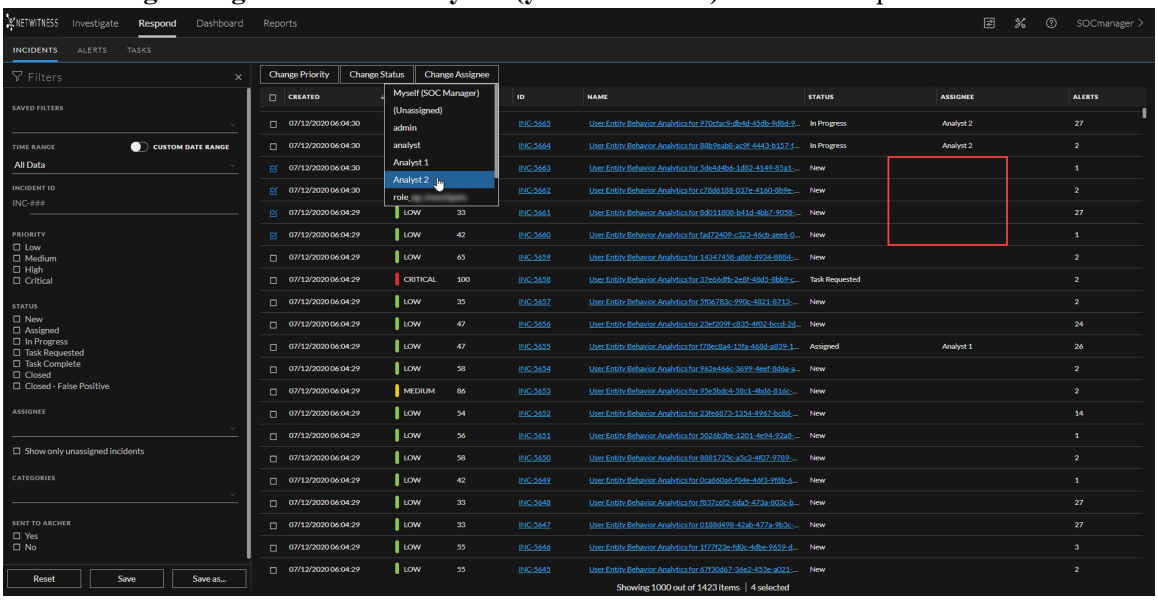
1. If you cannot see the Filter panel, in the Incident List view toolbar, click .
2. In the Filters panel, under Assignee, select **Show only unassigned incidents**.



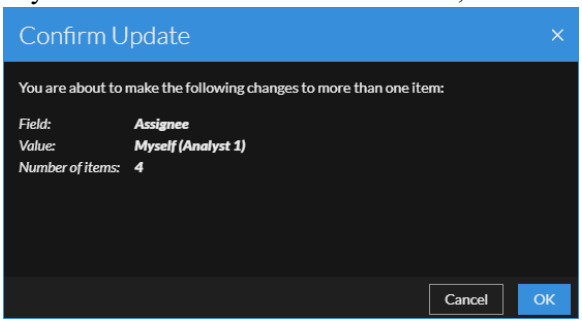
The incidents list is filtered to show unassigned incidents.

Assign Incidents to Myself

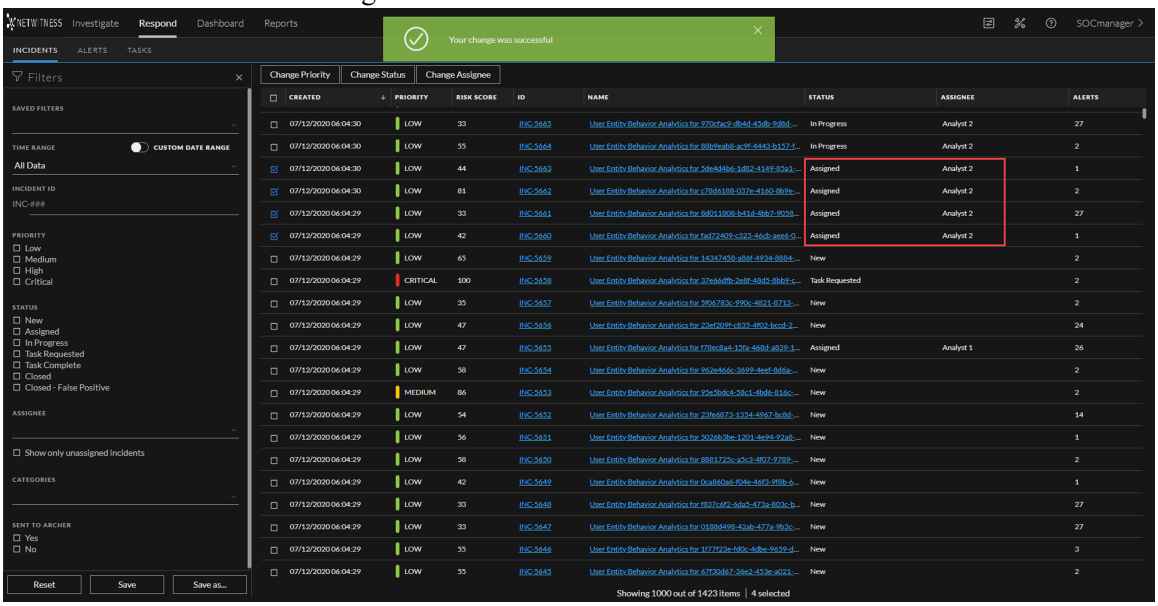
1. In the Incident List view, select one or more incidents that you want to assign to yourself.
2. Click **Change Assignee** and select **Myself (your full name)** from the drop-down list.



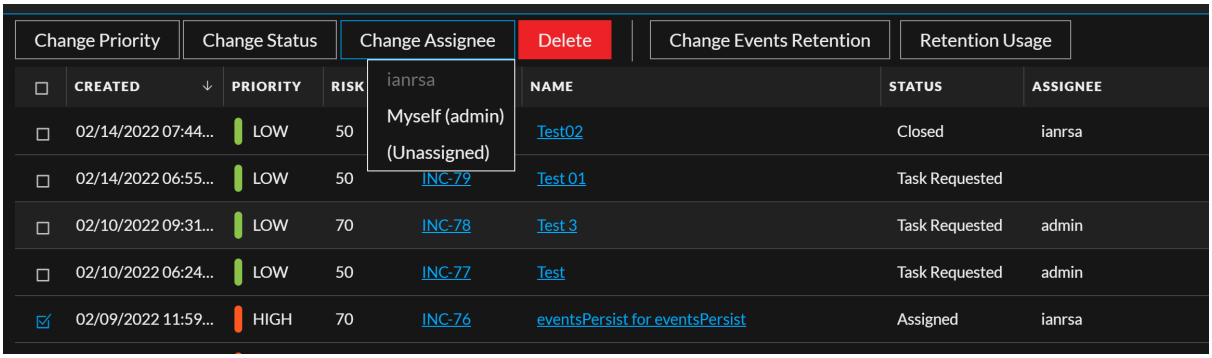
3. If you selected more than one incident, in the Confirm Update dialog, click **OK**.



You can see a successful change notification.

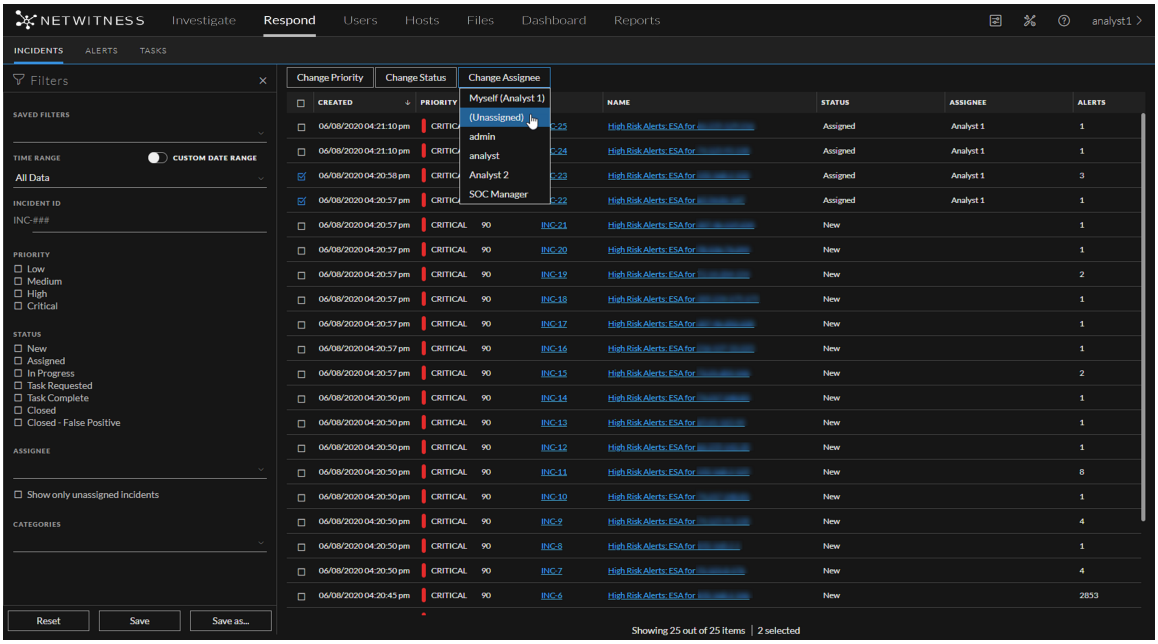


Note: On selecting any particular incident, current assignee name is grayed out under **Change Assignee** drop-down list. This is not applicable in case multiple incidents are selected. Refer the following figure.

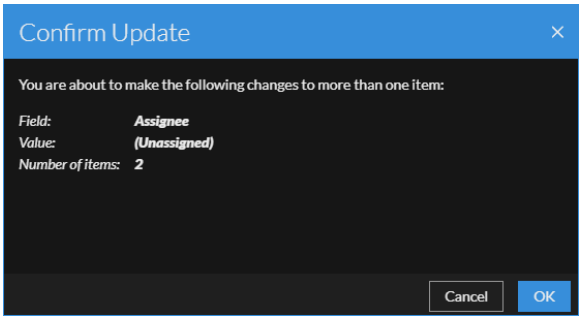


Unassign an Incident

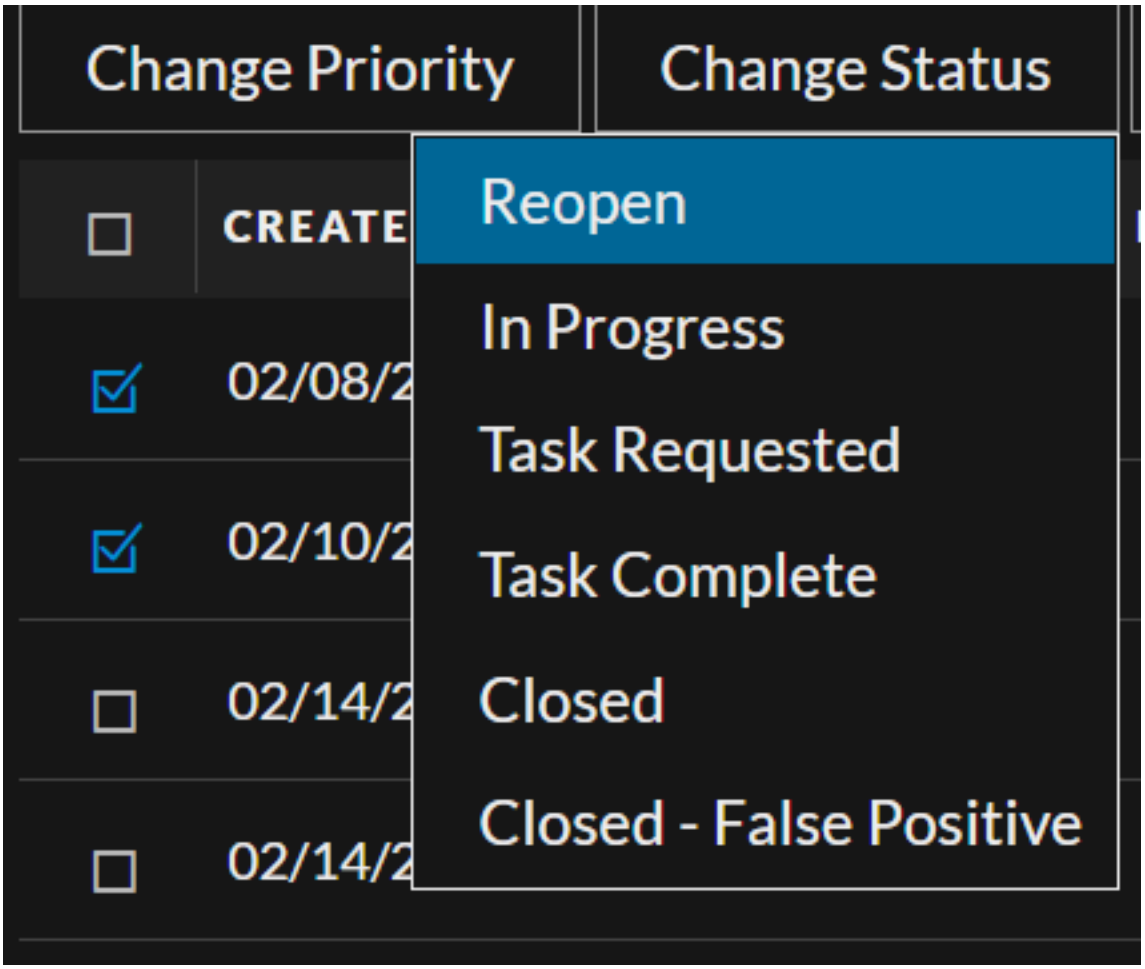
1. In the Incident List view, select one or more incidents that you want to unassign.
2. Click **Change Assignee** and select **(Unassigned)** from the drop-down list.



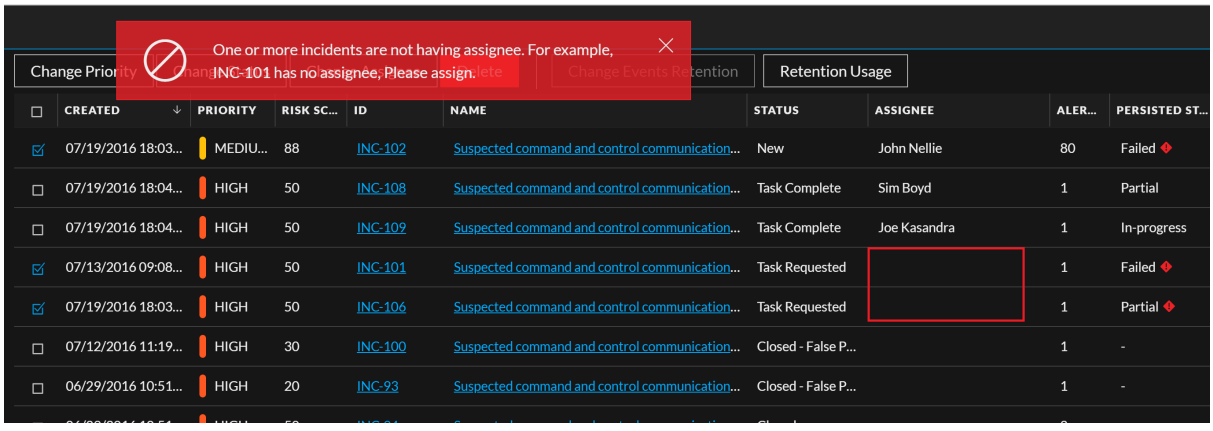
3. If you selected more than one incident, in the Confirm Update dialog, click **OK**.

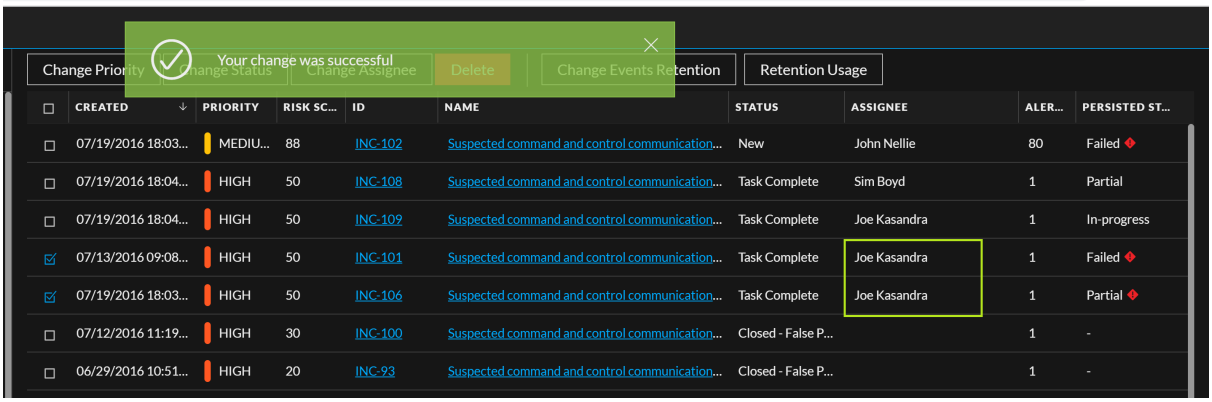


4. Verify that the Status is still correct and make changes as required. To change the status, select one or more incidents, click **Change Status**, and select a new status.



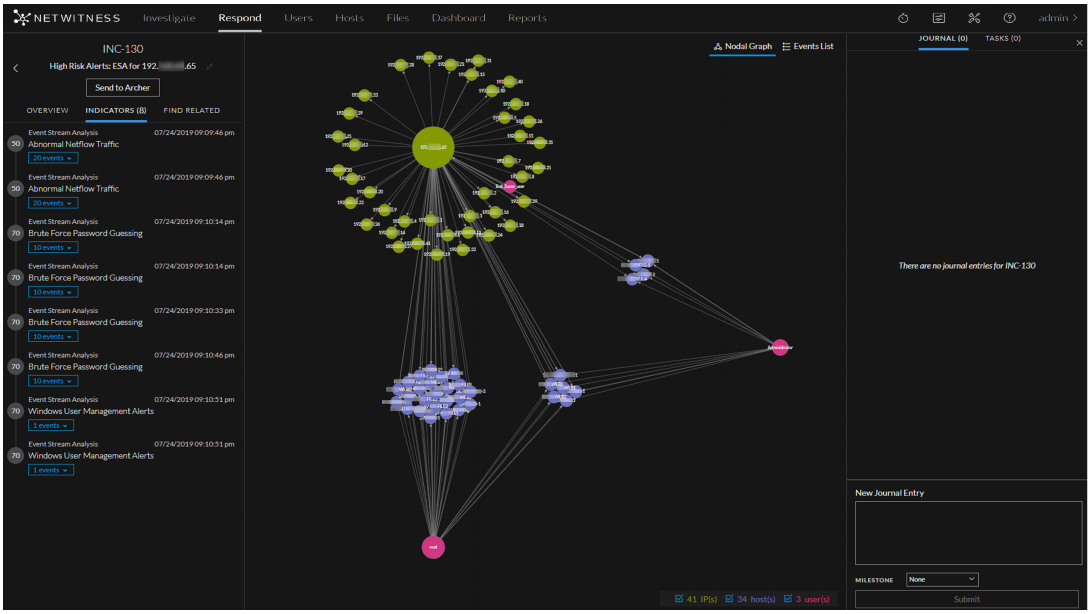
Note: You must assign the incidents to change their status. If you try to change the incident status without assigning the incident, the error message **One or more incidents are not having assignee. For example, INC-x has no assignee, Please assign.** is displayed. Refer the following figures.




Your change was successful
Change Priority
Change Status
Change Assignee
Delete
Change Events Retention
Retention Usage
CREATED
PRIORITY
RISK SC...
ID
NAME
STATUS
ASSIGNEE
ALER...
PERSISTED ST...
07/19/2016 18:03... MEDIU... 88 INC-102 Suspected command and control communication... New John Nellie 80 Failed
07/19/2016 18:04... HIGH 50 INC-108 Suspected command and control communication... Task Complete Sim Boyd 1 Partial
07/19/2016 18:04... HIGH 50 INC-109 Suspected command and control communication... Task Complete Joe Kasandra 1 In-progress
07/13/2016 09:08... HIGH 50 INC-101 Suspected command and control communication... Task Complete Joe Kasandra 1 Failed
07/19/2016 18:03... HIGH 50 INC-106 Suspected command and control communication... Task Complete Joe Kasandra 1 Partial
07/12/2016 11:19... HIGH 30 INC-100 Suspected command and control communication... Closed - False P... 1 -
06/29/2016 10:51... HIGH 20 INC-93 Suspected command and control communication... Closed - False P... 1 -

Determine which Incidents Require Action

Once you get the general information about the incident from the Incident List view, you can go to the Incident Details view for more information to determine the action required.

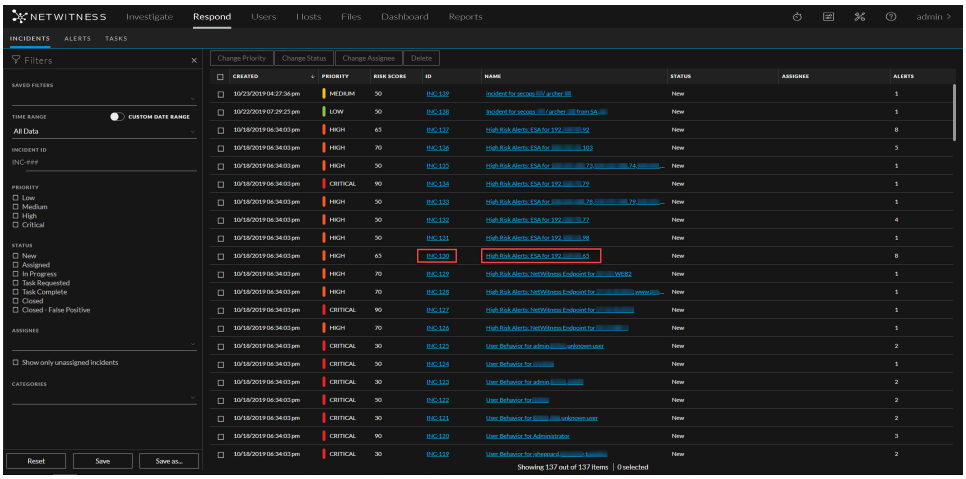


You can perform the following procedures in the Incident Details view to determine the action required on an incident:

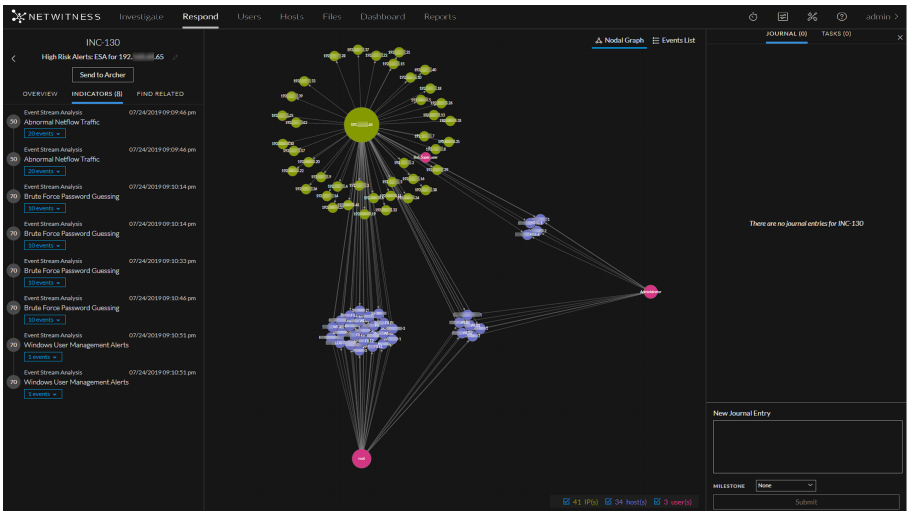
- View Incident Details
- View Basic Summary Information about the Incident
- View the Indicators and Enrichments
- View and Study the Events
- View C2 Enrichment Information for Suspected C&C Incidents
- View and Study the Entities Involved in the Events on the Nodal Graph
- Nodal Graph Behaviors and Characteristics
- Select Node Types to View on the Nodal Graph
- Filter the Data in the Incident Details View
- View the Tasks Associated with an Incident
- View Incident Notes
- Find Related Indicators
- Add Related Indicators to the Incident

View Incident Details

To view details for an incident, in the Incidents List view, choose an incident to view and then click the link in the **ID** or **NAME** column for that incident.



The Incident Details view for the selected incident appears with the Indicators panel, Nodal Graph, and Journal in view.

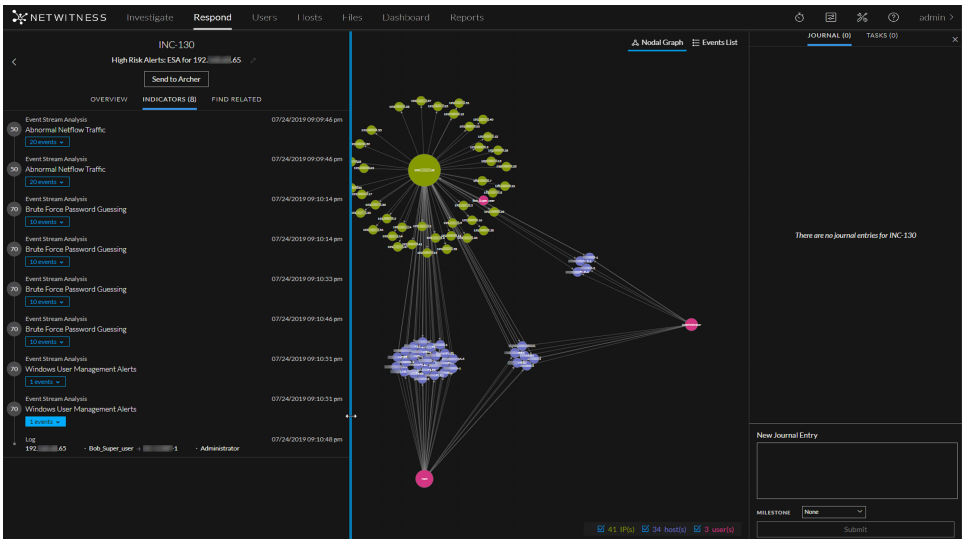


The Incident Details view has the following panels:

- **Overview**: The incident Overview panel contains high-level summary information about the incident, such as the score, priority, alerts, and status. You have the option to send the incident to Archer and change the incident Priority, Status, and Assignee.
- **Indicators**: The Indicators panel contains a chronological listing of indicators. *Indicators* are alerts, such as an ESA alert or a NetWitness Endpoint alert. This listing helps you to connect indicators and notable data. For example, an IP address connected to a command and communication ESA alert might also have triggered a NetWitness Endpoint alert or other suspicious activities.

- **Related Indicators**: The Related Indicators panel enables you to search the NetWitness alerts database to find alerts that are related to this incident. You can also add related alerts that you find to the incident.
- **History**: The History panel allows you to view the different actions performed by the user on an incident. Events such as Incident Assignee change, Incident Status change, Incident Priority change, and Incident creation are recorded in this panel.
- **Nodal Graph**: The nodal graph is an interactive graph that shows the relationship between the entities involved in the incident. An *Entity* is represented by an IP address, MAC address, user, host, domain, file name, or file hash.
- **Events List**: The Events List, also known as the Events table, lists the events associated with the incident. It also shows event source and destination information along with additional information depending on the event type. You can click the top of an event in the list to view the detailed data for that event.
- **Journal**: The Journal panel enables you to access the Journal for the selected incident, which allows you to communicate and collaborate with other analysts. You can post notes to a journal, add Investigation Milestone tags (Reconnaissance, Delivery, Exploitation, Installation, Command and Control, Action on Objective, Containment, Eradication, and Closure), and view the history of activity on your incident.
- **Tasks**: The Tasks panel shows all of the tasks that have been created for the incident. You can also create additional tasks from here.

To view more information in the left-side panel without scrolling, you can hover over the right edge and drag the line to resize the panel as shown in the following figure:

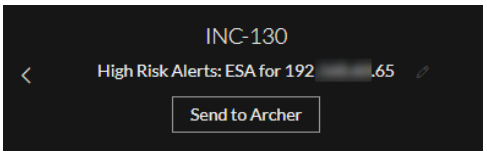


View Basic Summary Information about the Incident

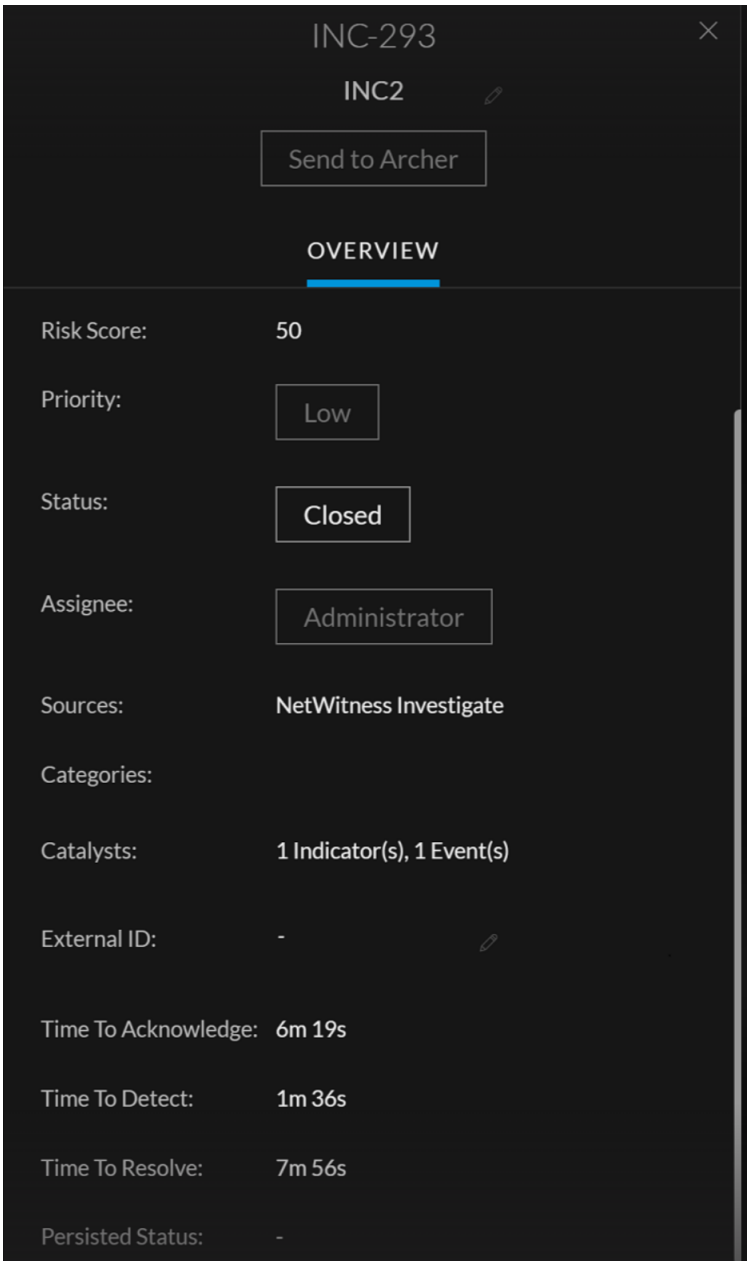
You can view basic summary information about an incident in the Overview panel.

Above the Overview panel, you can see the following information:

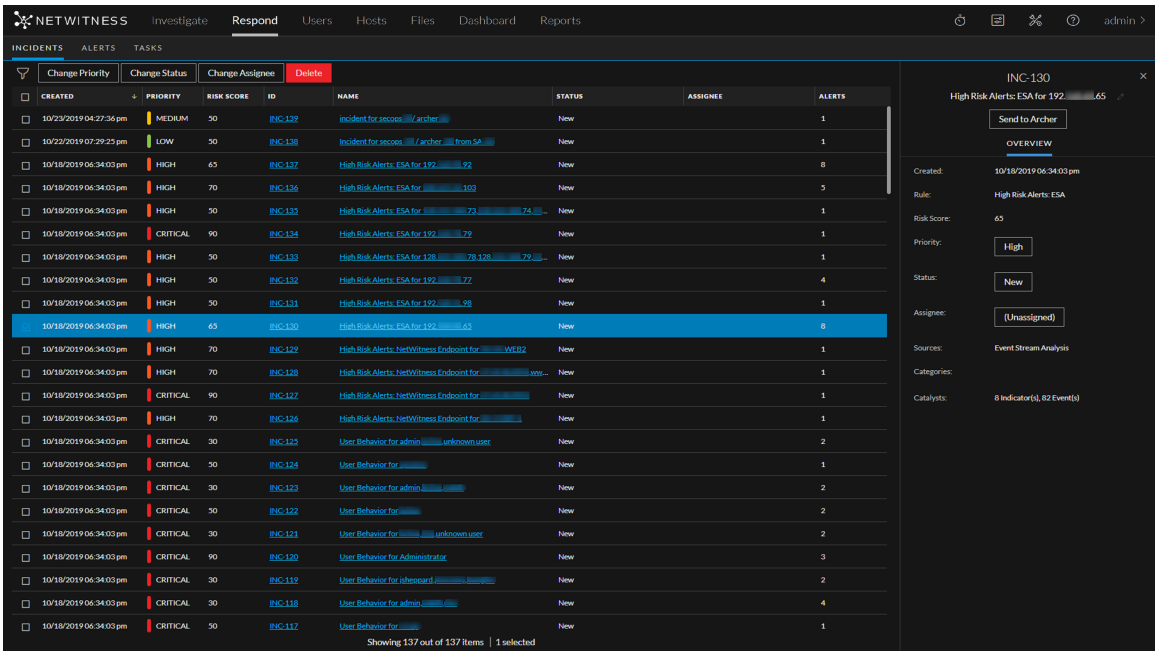
- **Incident ID**: This is an automatically created unique ID assigned to the incident.
- **Name**: The incident name is derived from the rule used to trigger the incident.
- **Send to Archer / Sent to Archer**: (In version 11.2 and later, if Archer is configured as a data source in Context Hub, you can send incidents to Archer Cyber Incident & Breach Response and this option is available in NetWitness Respond.) This shows whether an incident has been sent to Archer Cyber Incident & Breach Response. An incident sent to Archer shows as Sent to Archer. An incident that has not been sent to Archer shows as Send to Archer. You can click the Send to Archer button to send the incident to Archer Cyber Incident & Breach Response.



To view the Overview panel from the Incident Details view, select **Overview** in the left panel.



To view the Overview panel from the Incidents List view, click a row in the incident list. The Overview panel appears on the right.



The Overview panel contains basic summary information about the selected incident:

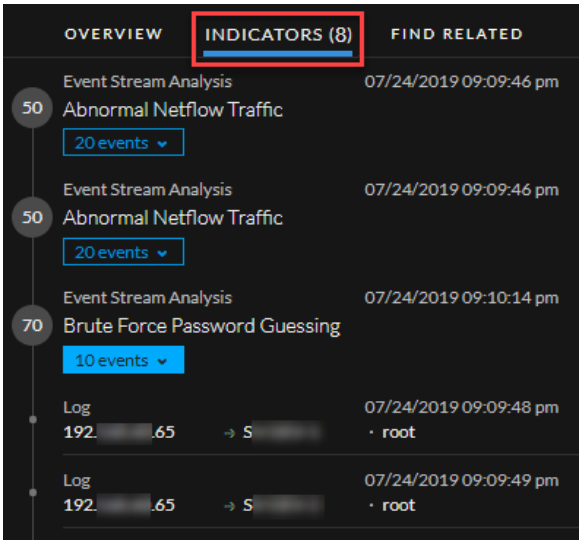
- **Created**: Shows the creation date and time of the incident.
- **Rule / By**: Shows the name of the rule that created the incident or the name of the person who created the incident.
- **Risk Score**: Shows a value between 0 and 100 that indicates the risk of the incident as calculated by an algorithm. 100 is the highest risk score.
- **Priority**: Shows the incident priority. Priority can be Critical, High, Medium or Low.
- **Status**: Shows the incident status. The status can be Reopen, New, Assigned, In Progress, Task Requested, Task Complete, Closed, and Closed - False Positive. After you create a task, the status changes to Task Requested.
- **Assignee**: Shows the team member currently assigned to the incident.
- **Sources**: Indicates the data sources used to locate the suspicious activity.
- **Categories**: Shows the categories of the incident events.
- **Catalysts**: Shows the count of indicators that gave rise to the incident.
- **External ID**: Allows storing the Incident ID referrals from a different platform such as Archer.
- **Time to Acknowledge**: Shows the time taken to assign an Incident after creating it.
- **Time to Detect**: Shows the time taken for completing the task after the Incident is assigned.
- **Time to Resolve**: Shows the time taken for closing the task after the Incident is created.
- **Persisted Status**: Shows the persist status of the Incident. The status can be Complete, Partial, or None (-).

View the Indicators and Enrichments

Note: *Indicators* are alerts, such as an ESA alert or a NetWitness Endpoint alert.

You can find indicators, events, and enrichments on the Indicators panel. The Indicators panel is a chronological listing of indicators that helps you to find enrichments and events related to the triggering indicator. For example, an indicator might be a Command and Control alert, a NetWitness Endpoint alert, a Suspicious Domain (C2) alert, or an alert from an Event Stream Analysis (ESA) rule. The Indicators panel helps you to aggregate and order these indicators (alerts) from different systems so that you can see how they are related and also help you develop a timeline of a given attack.

To view the Indicators panel, in the left panel of the Incident Details view, select **Indicators**.



Indicators are alerts, such as an ESA alert or a NetWitness Endpoint alert. This listing helps you to connect indicators and notable data. For example, indicators can show the data found by your rules. In the Indicators panel, the risk score for an indicator is shown within a solid-colored circle.

Data source information is shown below the names of the indicators. You can also see the creation date and time of the indicator and the number of events in the indicator. When data is available, you can see the number of enrichments. You can click the event and enrichment buttons to view the details.

Note: The maximum number of indicators (alerts) displayed in the Indicators panel is 1,000.

View and Study the Events

You can view and study the events associated with the incident from the Events List. It shows information about the events, such as event time, source IP, destination IP, detector IP, source user, destination user, and file information about the events. The amount of information listed depends on the event type.

There are two types of events:

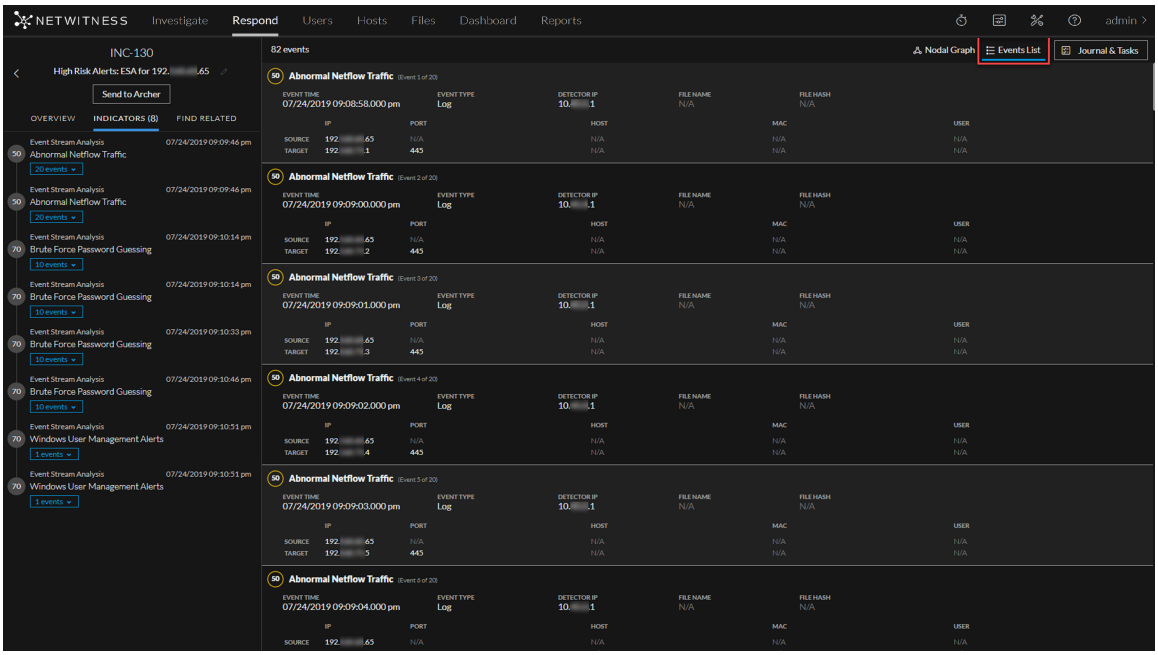
- A transaction between two machines (a Source and a Destination)
- An anomaly detected on a single machine (a Detector)

Some events will only have a Detector. For example, NetWitness Endpoint finds malware on your machine. Other events will have a Source and Destination. For example, packet data shows communication between your machine and a Command and Control (C2) domain.

You can drill further into an event to get detailed data about the event.

To view and study the events:

1. To view the Events List, in the Incident Details view toolbar, click Events List.



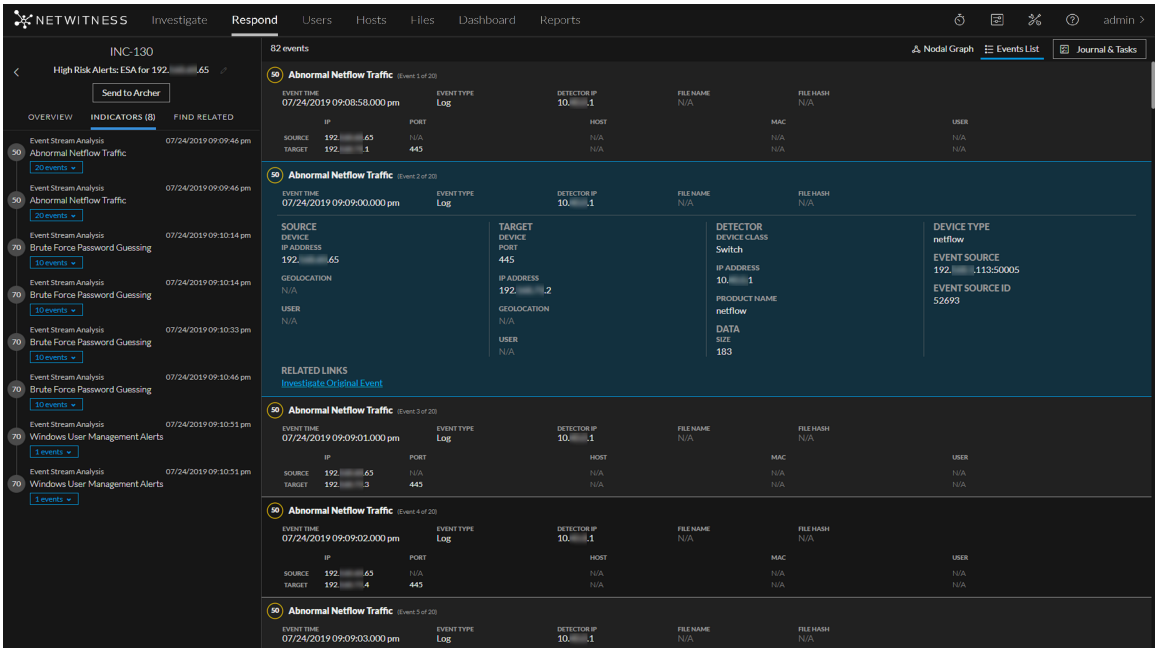
Note: The EVENT TIME displayed on this screen is the same as the COLLECTION TIME from the investigation page.

The Events List shows different information about each event depending on the event type. The maximum number of events displayed in the Events List is 1,000.

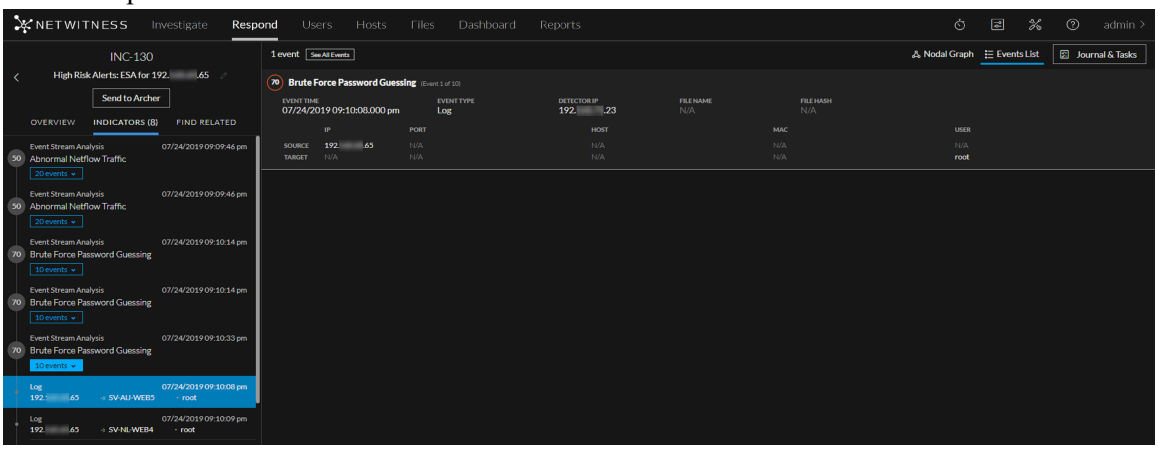
The following table lists typical event information. For details specific to endpoint events, see Events List.

Field	Description
EVENT TIME	Shows the time the event occurred.
EVENT TYPE	Shows the type of alert, such as Log and Network.
DETECTOR IP	Shows the IP address of the machine where an anomaly was detected.
FILE NAME	Shows the file name if a file is involved with the event.
FILE HASH	Shows a hash of the file contents.
SOURCE IP	Shows the source IP address if there was a transaction between two machines.
SOURCE PORT	Shows the source port of the transaction. The source and destination ports can be on the same IP address.
SOURCE HOST	Shows the destination host where the event took place.
SOURCE MAC	Shows the MAC address of the source machine.
SOURCE USER	Shows the user of the source machine.
TARGET IP	Shows the destination IP address if there was a transaction between two machines
TARGET PORT	Shows the destination port of the transaction. The source and destination ports can be on the same IP address.
TARGET HOST	Shows the host name of the destination machine.
TARGET MAC	Shows the MAC address of the destination machine.
TARGET USER	Shows the user of the destination machine.

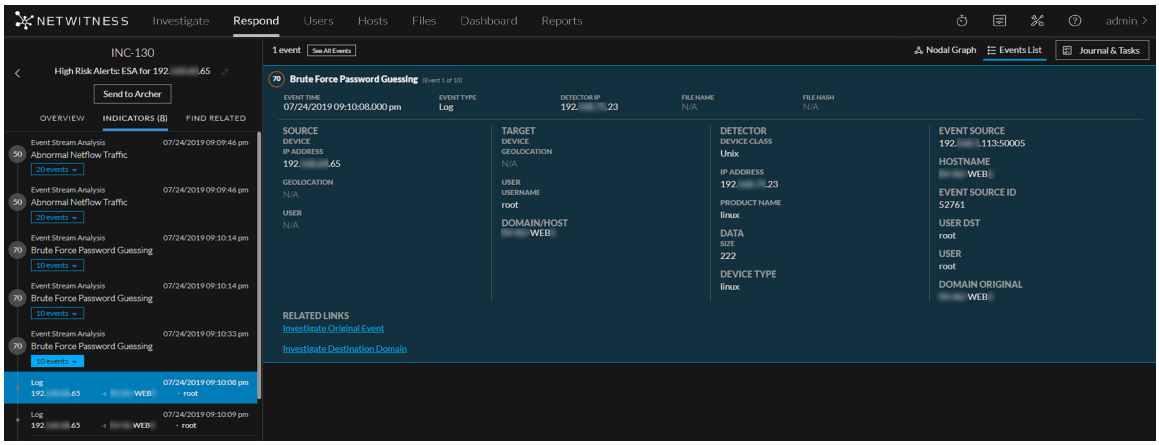
2. Click the top of an event in the Events List to view the event details.
   This example shows the event details for a selected event in the list.



3. To view the events for a specific indicator (alert), go to the Indicators panel on the left and click the indicator to view the events for that indicator in the Events List on the right.
   This example shows one event for a selected indicator.



4. To view event details for a specific indicator event, select an event in the Indicators panel. Click the top of the event to view the details.
   The following example shows information for the selected event.



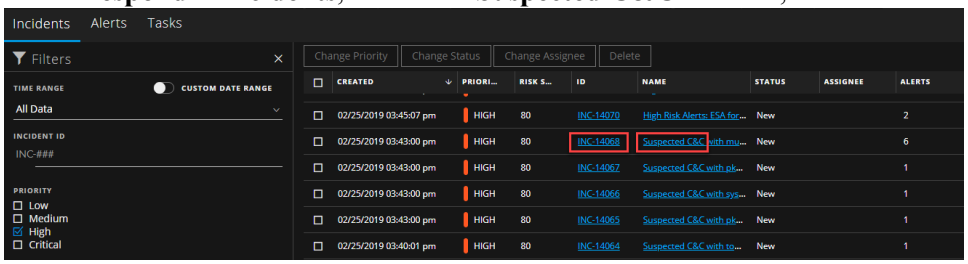
If you have additional Investigate-server permissions, you can also access event analysis details for events. See View Event Analysis Details for Indicators. If you have the UEBA_Analysts role, you can access UEBA details for indicators. See View User Entity Behavior Analytics for Indicators.

View C2 Enrichment Information for Suspected C&C Incidents

Note: This procedure applies only to incidents from ESA Analytics in NetWitness Version 11.3 and 11.4. The Event Stream Analytics Server (ESA Analytics) service, which is used for Automated Threat Detection, is end of life (EOL) and not supported in NetWitness Platform Version 11.5 and later.

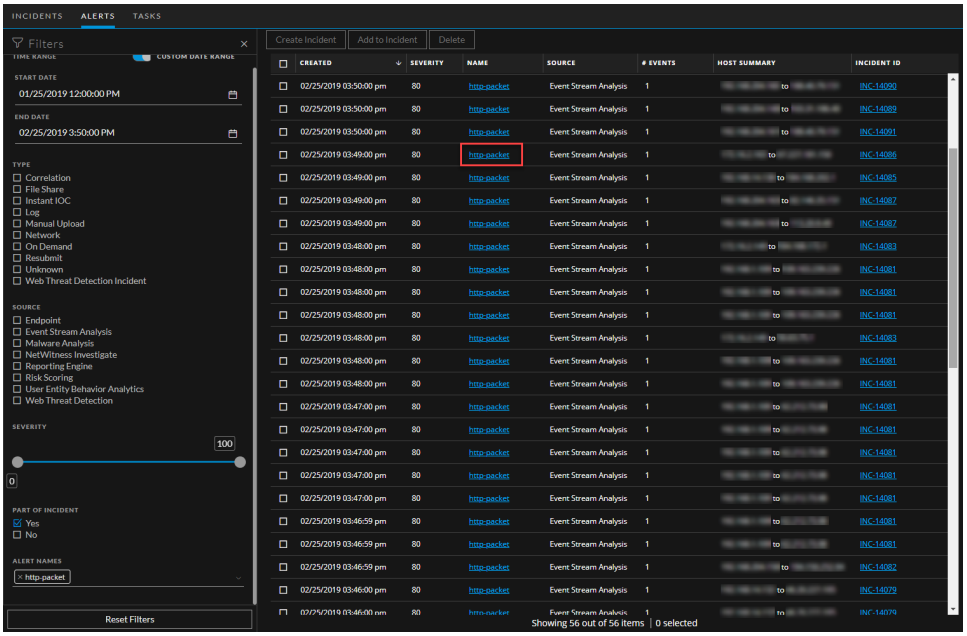
The Events List in version 11.3 and later does not show the Command and Control (C2) enrichment information for HTTP packet alerts in Suspected C&C incidents. However, you can view the C2 enrichment information in the Alert Details view.

1. Go to **Respond > Incidents**, look for a **Suspected C&C** incident, and note the incident ID.

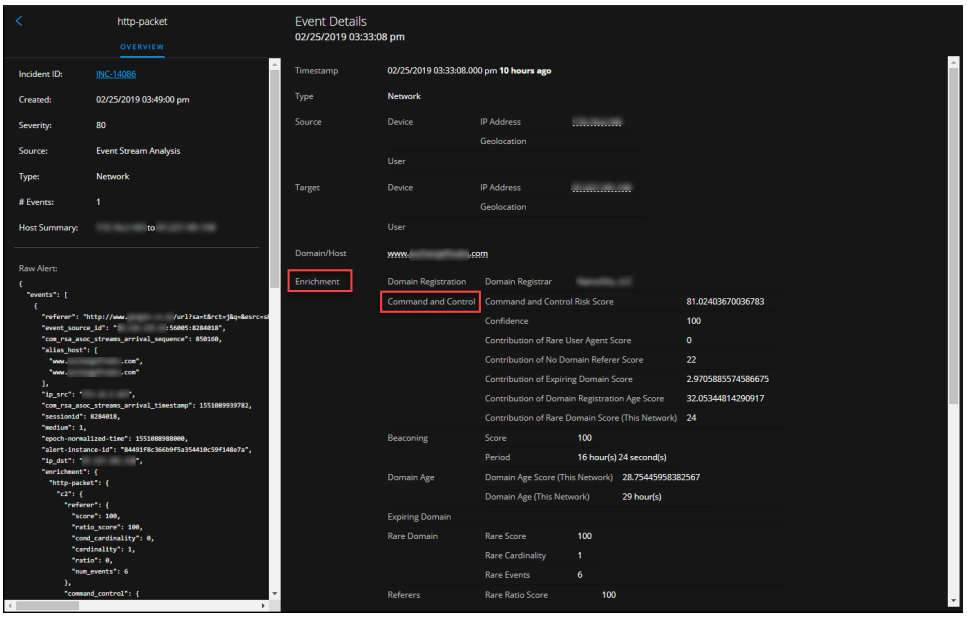


2. Go to **Respond > Alerts** and in the Filters panel, select the following to locate an alert in the Alerts list with the incident ID noted above:
   a. In the **Part of Incident** section, select **Yes**.
   b. In **Alert Names** section, select **http-packet**.

If you are still not able to locate an alert in the Alerts list with the incident ID noted above, try filtering your alerts list more using the time range of the incident.



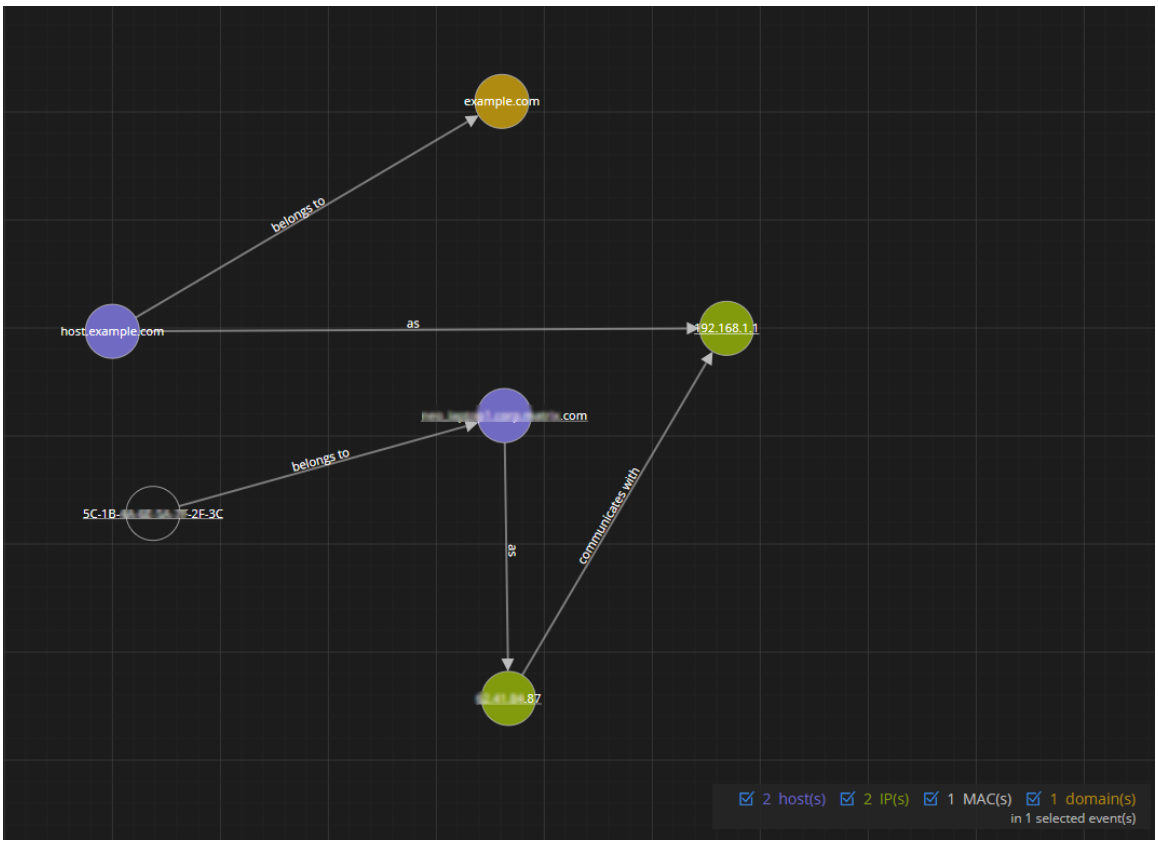
3. In the Alerts list, click the **http-packet** link in the **NAME** field of the alert associated with the incident ID.
The Event Details view shows the C2 enrichment information.



View and Study the Entities Involved in the Events on the Nodal Graph

An *Entity* is either an IP address, MAC address, user, host, domain, file name, or file hash. The nodal graph is an interactive graph that you can move around to get a better understanding of how the entities involved in the events relate to each other. The nodal graphs look different depending on the type of event, the number of machines involved, whether the machines are associated with users, and if there are files associated with the event.

The following figure shows an example nodal graph with six nodes.



If you look closely at the nodal graph, you can see circles that represent nodes. A nodal graph can contain one or more of the following types of nodes:

- **IP address** (If the event is a detected anomaly, you can see a Detector IP. If the event is a transaction, you can see a Destination IP and a Source IP.)
- **MAC address** (You may see a MAC address for each type of IP address.)
- **User** (If the machine is associated with a user, you can see a user node.)
- **Host**
- **Domain**

- **Filename** (If the event involves files, you can see a filename.)
- **File Hash** (If the event involves files, you may see a file hash.)

In NetWitness 11.3 events, nodes for source filename and file hash are supported, but nodes for target filename and file hash are not supported. In NetWitness 11.4 and later events, nodes for both source and target filenames as well as file hashes are supported.

The legend at the bottom of the nodal graph shows the number of nodes of each type and the color coding of the nodes.

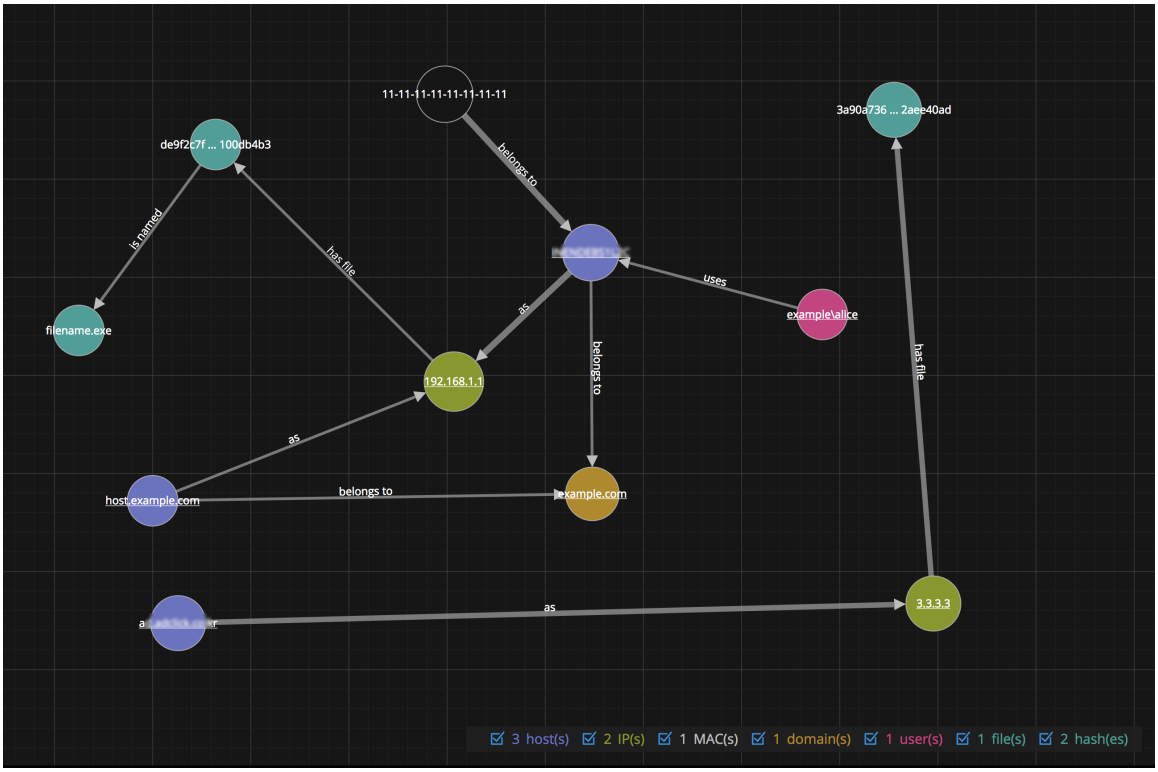
You can click and drag any node to reposition it.

The arrows between the nodes provide additional information about the entity relationships:

- **Communicates with**: An arrow between a Source machine node (IP address or MAC address) and a Destination machine node labeled with "communicates with" shows the direction of the communication.
- **Has file**: An arrow between a machine node (IP address, MAC address, or Host) and a file hash node labeled with "has file" indicates that the IP address has that file.
- **Uses**: An arrow between a User node and a machine node (IP address, MAC address, or Host) labeled with "uses" shows the machine that the user was using during the event.
- **Calls**: (This arrow is available in NetWitness Platform 11.4 and later.) An arrow between two file hash (checksum) nodes labeled with "calls" indicates the direction of the interaction between the associated files. The source file hash "calls" the target (destination) file hash, which indicates that the source file associated with the source file hash is performing an action on the target file associated with the target file hash.
- **As**: (This relationship type represents attributes of the connected node.) An arrow between nodes labeled with "as" provides additional information about the IP address that the arrow points to. In the above example, there is an arrow from the host node circle that points to an IP address node that is labeled with "as". This indicates that the name on the host node circle is the hostname of that IP address and is not a different entity.
- **Is named**: (This relationship type represents attributes of the connected node.) An arrow from a File Hash node to a File Name node labeled with "is named" indicates that the file hash corresponds to a file with that name.
- **Belongs to**: (This relationship type represents attributes of the connected node.) An arrow between two nodes labeled with "belongs to" indicates that they pertain to the same node. For example, an arrow between a MAC address and a Host labeled with "belongs to" indicates that it is the MAC address for the host.

Larger line size arrows indicate more communication between the nodes. Larger nodes (circles) indicate more activity than smaller nodes. The larger nodes are the most common entities mentioned in the events.

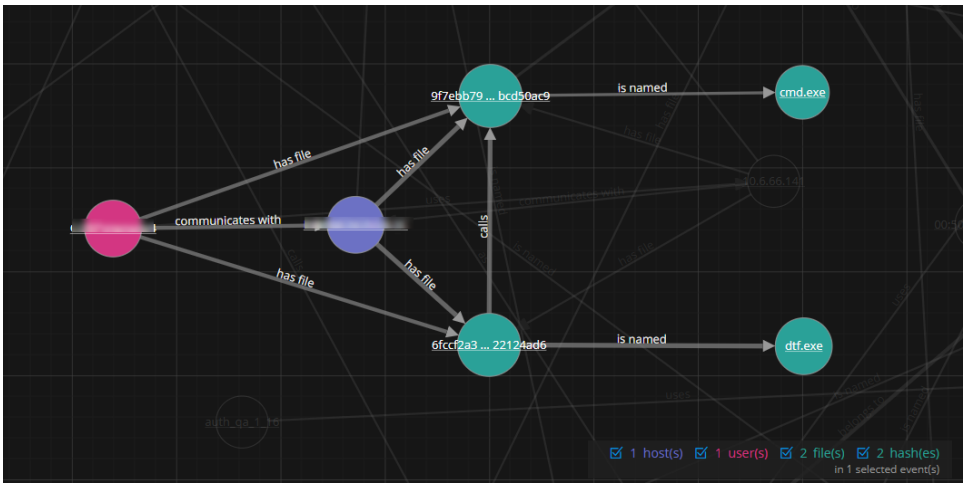
The following nodal graph example has 11 nodes.



In this example, notice that there are two IP nodes. They both have hashed files, but they do not communicate with each other. The IP address at the top (192.168.1.1) represents one machine with two hostnames (host.example.com is one of them) in the example.com domain. The MAC address of the machine is 11-11-11-11-11-11-11-11-11-11 and Alice uses it.

Note: The following example applies to NetWitness Platform 11.4 and later.

In the following nodal graph example, you can see the interaction between source and target (destination) files. There are six nodes in the selected event.



In this example, the user communicates with a host that has `dtf.eve` and `cmd.exe` files. The `dtf.exe` file on the host "calls" (in this case, launches) the `cmd.exe` file, which is suspected to be malicious activity. Notice that the "calls" arrow appears between the source and target file hashes, which are associated with the files.

Nodal Graph Behaviors and Characteristics

Note: These nodal graph behaviors and characteristics are available in NetWitness Platform Version 11.4 and later.

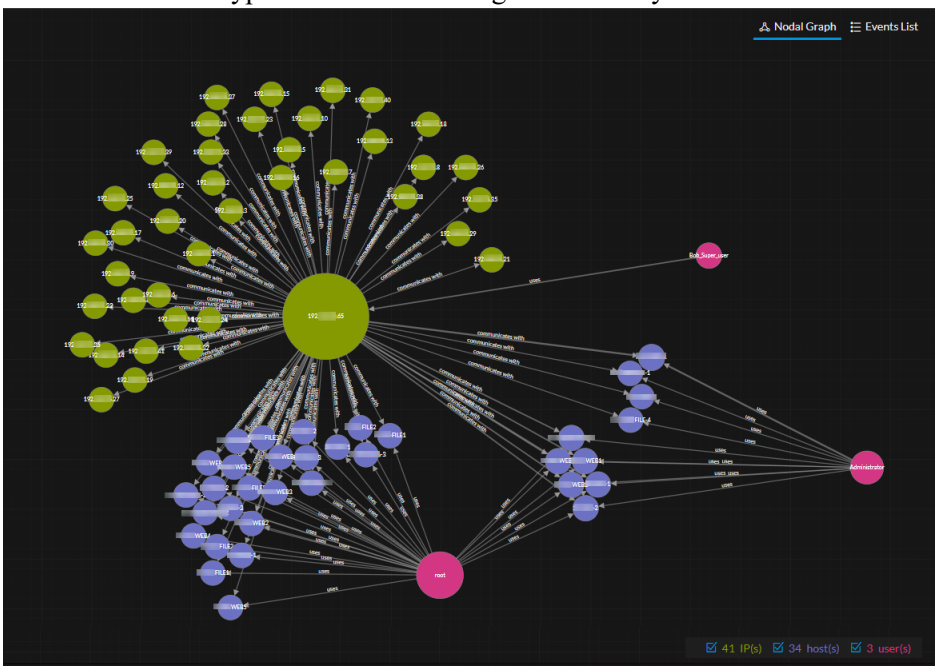
The nodal graph makes it easier for an analyst to get an initial understanding of an incident with minimal effort.

The nodal graph provides the following benefits to an analyst when responding to an incident:

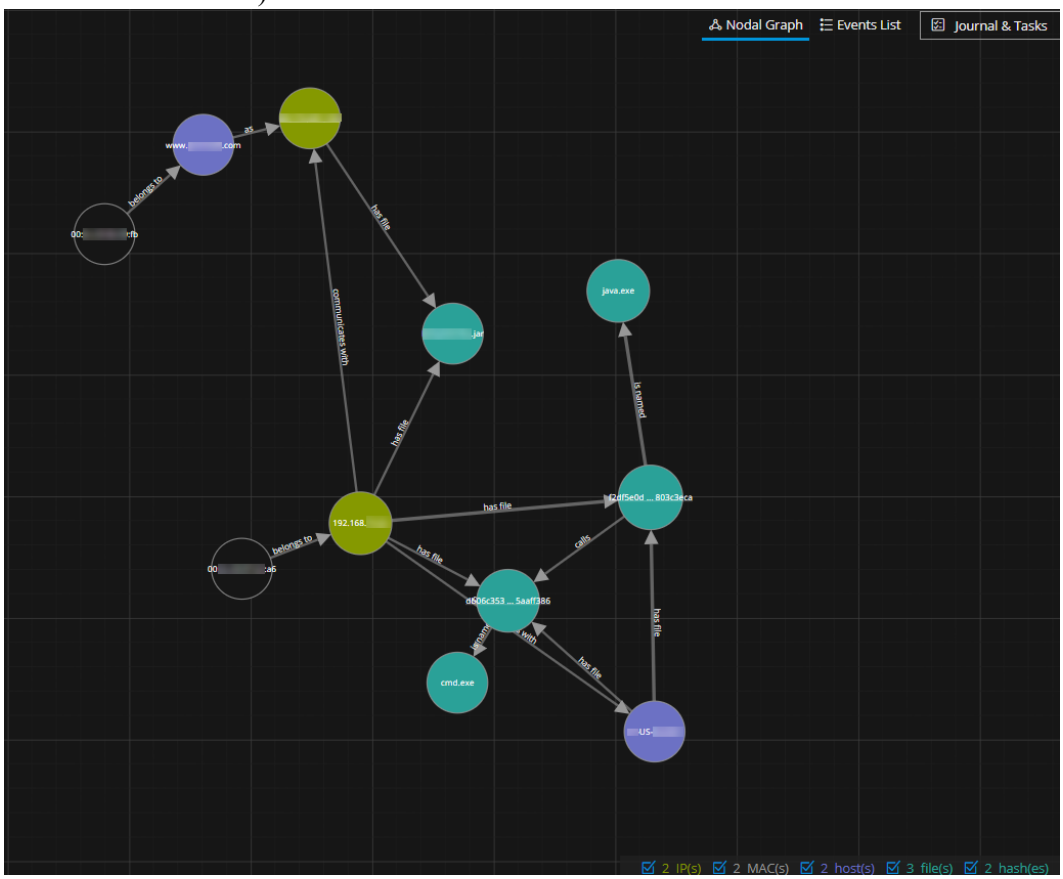
- The nodal graph helps determine scope, commonalities, and outliers in a given dataset, which can be useful context for an analyst.
- In many cases, the initial nodal graph layout presents valuable insight without any interaction from the analyst.
- In cases where the initial layout does not give enough clarity or when an analyst wants to view things differently, a few nodal mouse-drag position adjustments can provide a much faster method of exposing insightful relationships and clusters.

The following behaviors and characteristics are now part of the graph:

- Entities of similar types tend to cluster together visually.

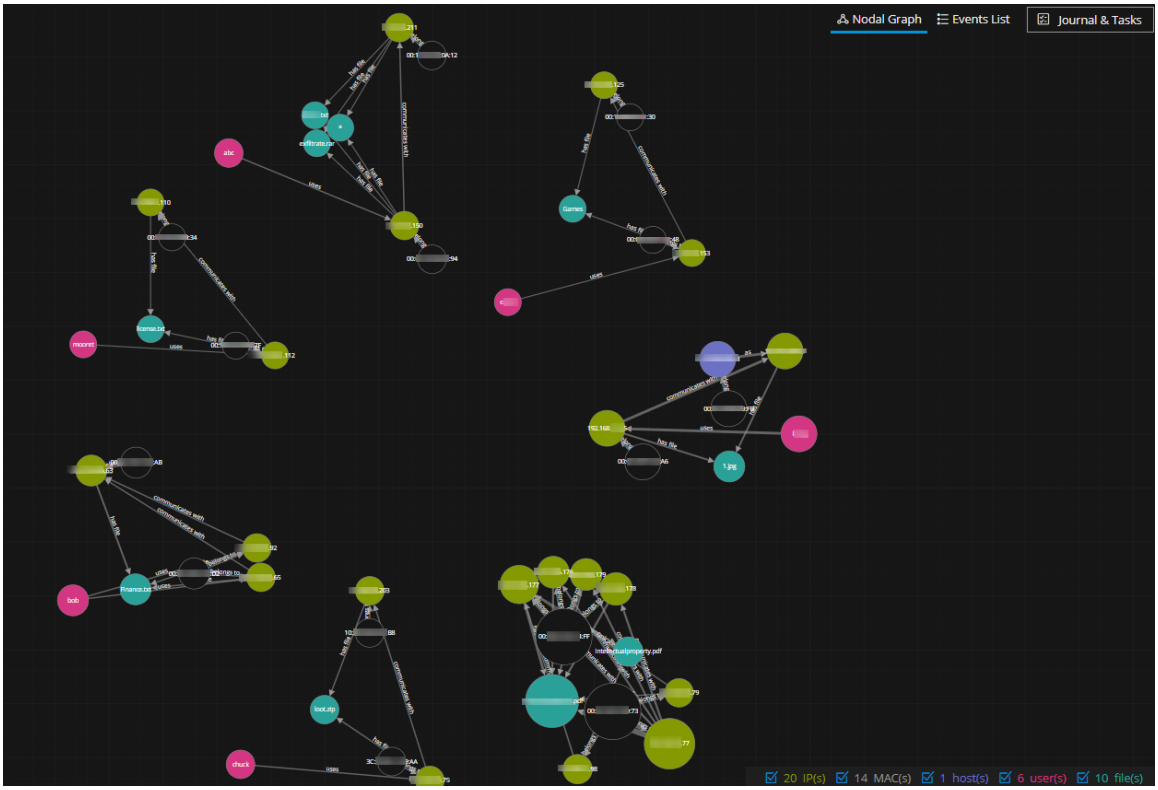


- Attributes and actions are better differentiated. Arrows that represent attributes ("as", "is named", "belongs to", and "has file") tend to be shorter than those representing actions ("call" and "communicates with").



- Leaf nodes, which are nodes that only have a single relationship to a single entity, tend to stay closer together.

- Disjointed graphs, such as clusters of entities and relationships that do not have connections with one another, are forced apart.



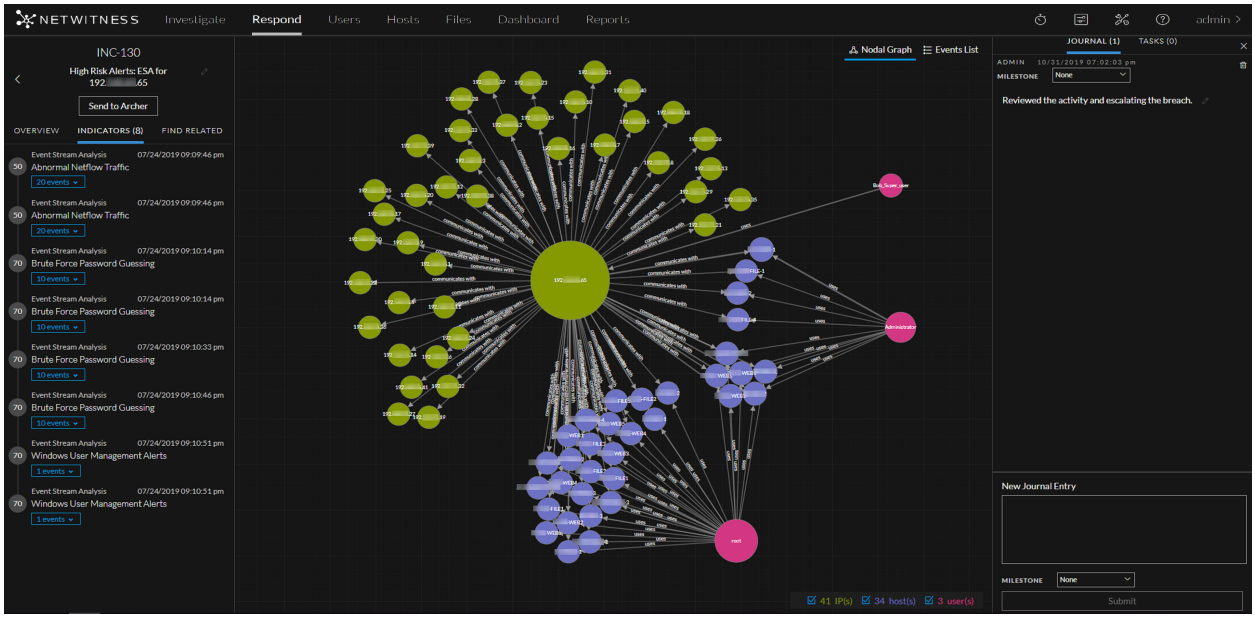
Dragged nodes are pinned in place. Double-click a node to unpin it and allow the forces to apply again to the node.

Select Node Types to View on the Nodal Graph

Note: This option is available in NetWitness Platform Version 11.2 and later.

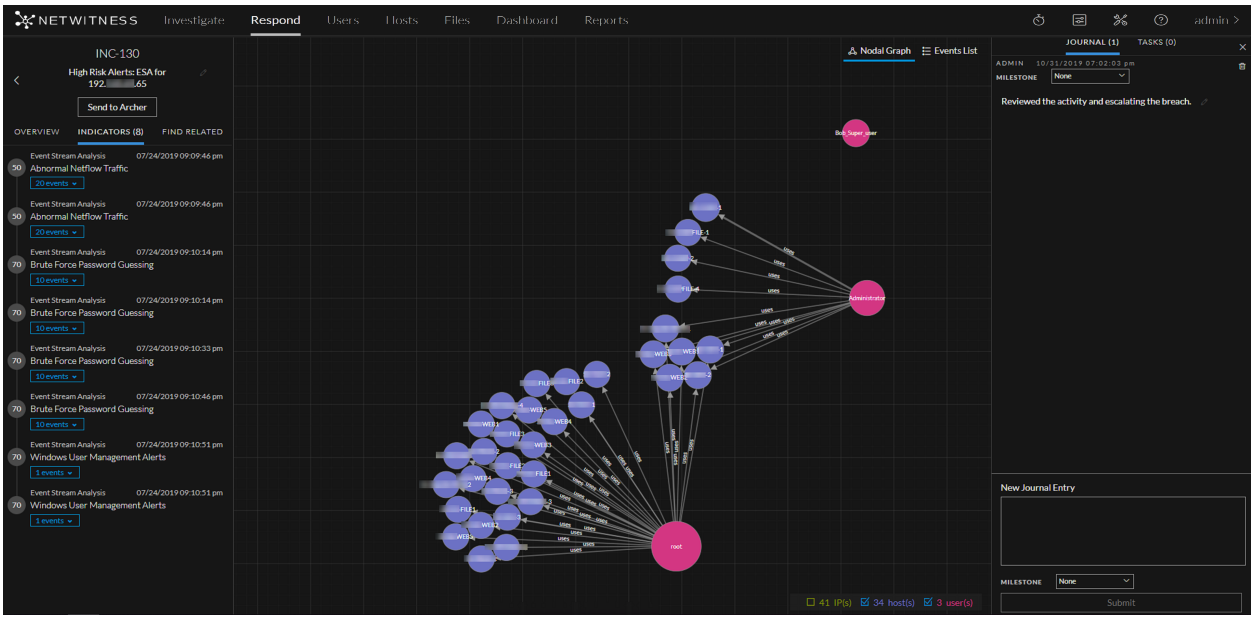
In the Incident Details view nodal graph, you can hide node types to further study the interactions between the entities on the nodal graph.

1. Go to **Respond > Incidents**.
2. In the Incidents List view, choose an incident to view and then click the link in the **ID** or **NAME** column for that incident.
   The Incident Details view for the selected incident appears with the Nodal Graph in view. The legend below the nodal graph has all of the entity node types selected by default.
   If you do not see the nodal graph, click Nodal Graph.

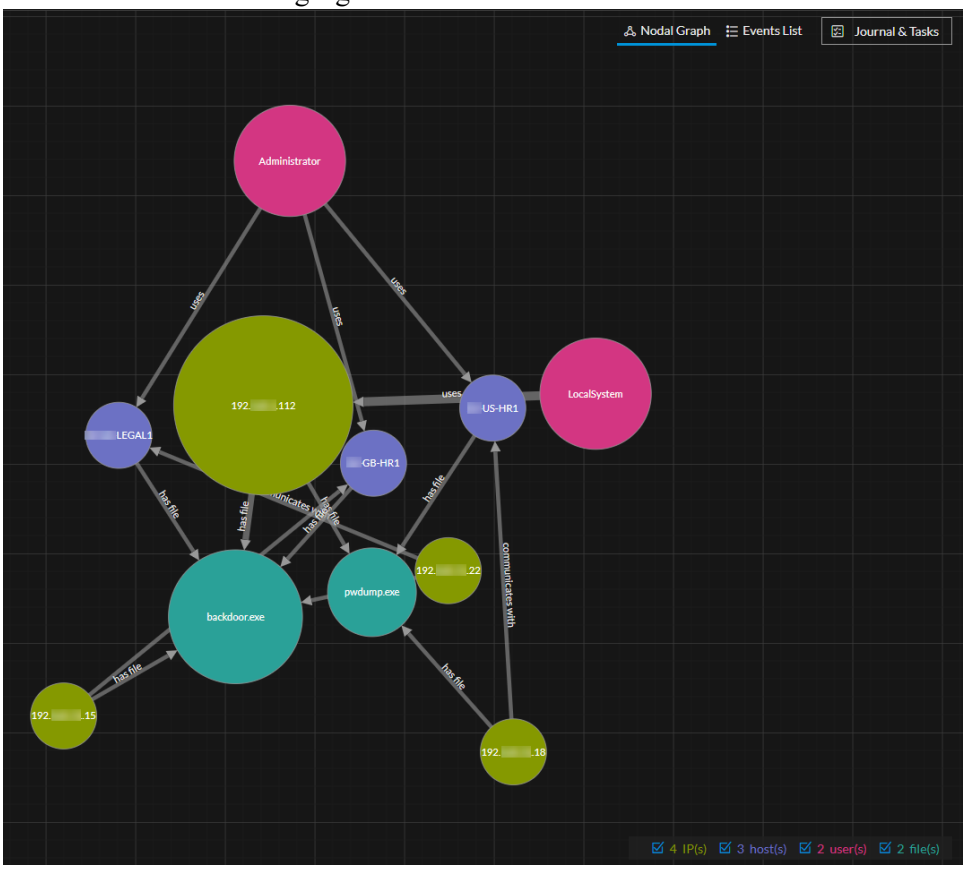


3. To hide node types, in the legend, clear the checkbox for the node types that you would like to hide in the nodal graph.

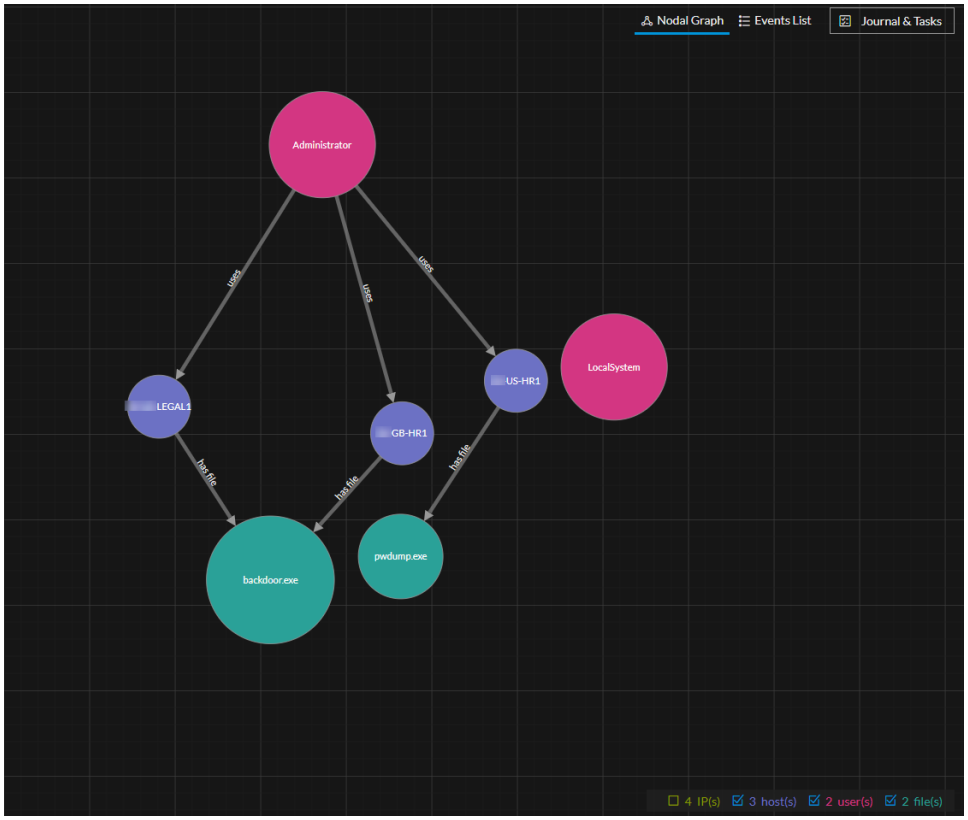
   The following example shows the IP address node type cleared and the IP address nodes are now hidden.



4. To include (unhide) node types, select the checkbox for the node types that you would like to appear in the nodal graph.
   Hiding node types can be especially helpful if the nodal diagram has overlapping entity relationships as shown in the following figure.



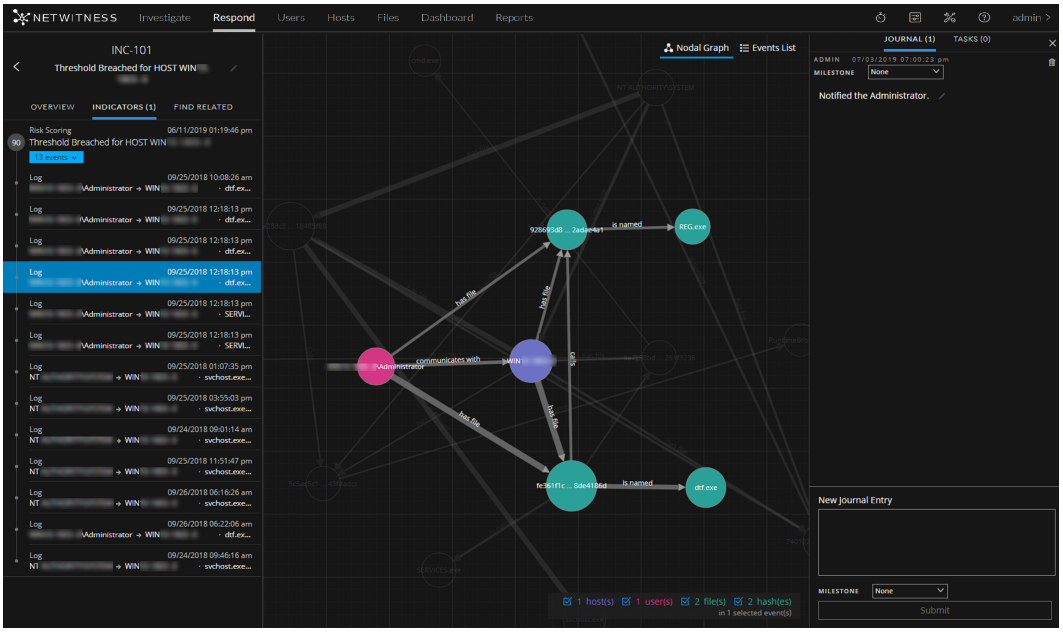
After hiding the IP node types, you can get a better understanding of what is happening with the remaining nodes.



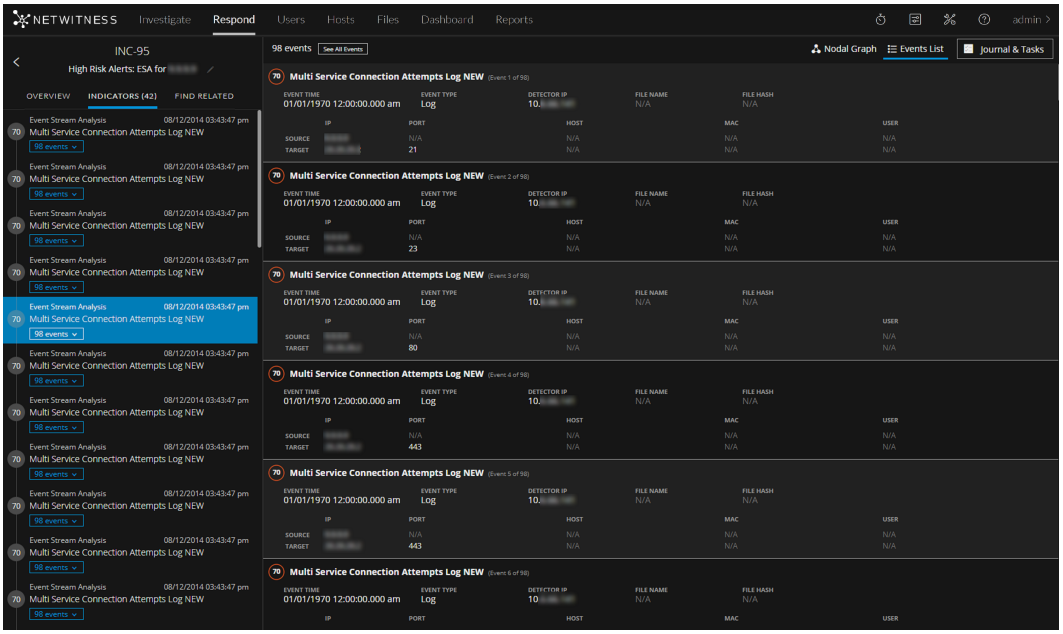
Filter the Data in the Incident Details View

You can click indicators in the Indicators panel to filter what you can see in the Nodal Graph and the Events List.

If you select an indicator to filter the nodal graph, data that is not part of your selection is dimmed, but it is still in view as shown in the following figure.



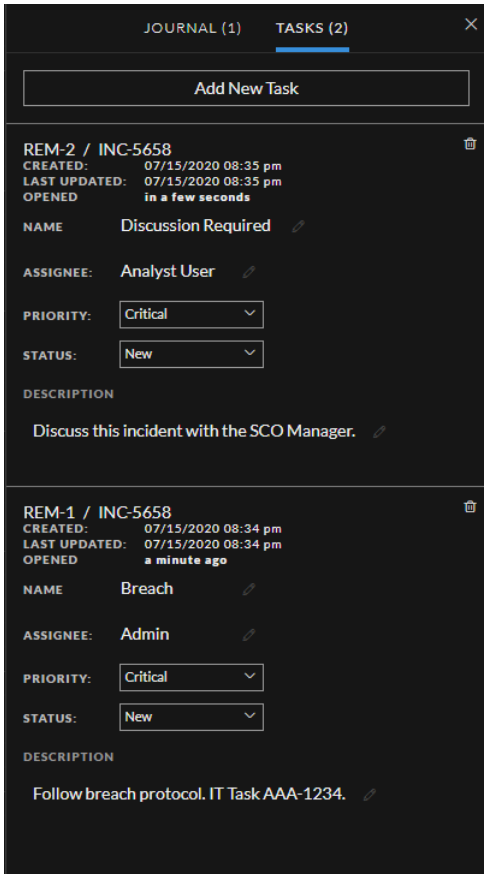
If you select an indicator to filter the Events List, only the events for that indicator are shown in the list. The following figure shows an indicator selected that contains ninety-eight events. The filtered Events List shows those ninety-eight events.



View the Tasks Associated with an Incident

Threat responders and other analysts can create tasks for an incident and track those tasks to completion. This can be very helpful, for example, when you require actions on incidents from teams outside of your security operations. You can view the tasks associated with an incident in the Incident Details view.

1. Go to **Respond > Incidents** and locate the incident that you want to view in the Incidents List.
2. Click the link in the **ID** or **NAME** field of the incident.
3. In the Journal on the right side of the Incident Details view, click the **TASKS** tab.
   If you cannot see the Journal, click **Journal & Tasks** and then click the **TASKS** tab.
   The Tasks panel shows all of the tasks for the incident.

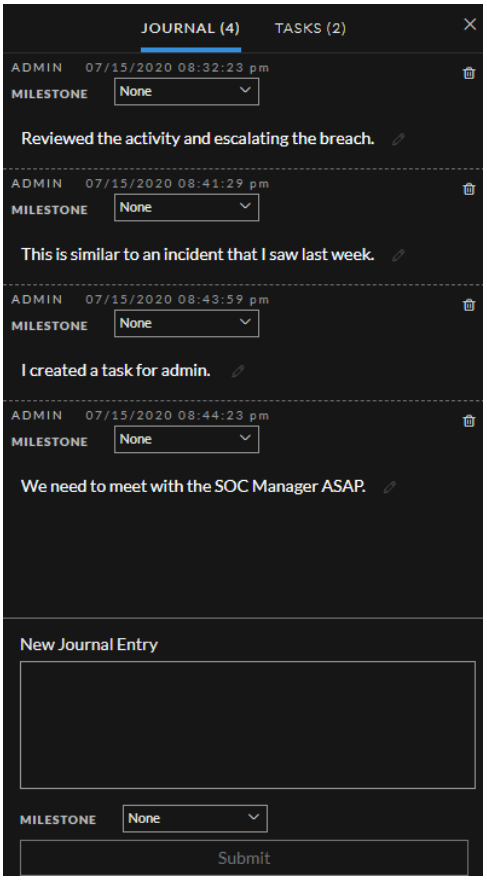


For more information about tasks, see Tasks List View, View All Incident Tasks, and Create a Task.

View Incident Notes

The incident Journal enables you to view the history of activity on your incident. You can view journal entries from other analysts and also communicate and collaborate with them.

1. Go to **Respond > Incidents** and locate the incident that you want to view in the Incidents List.
2. Click the link in the **ID** or **NAME** field of the incident.
   The Journal on the right side of the Incident Details view shows all of the journal entries for the incident.
   If you cannot see the Journal, in the toolbar, click **Journal & Tasks**.

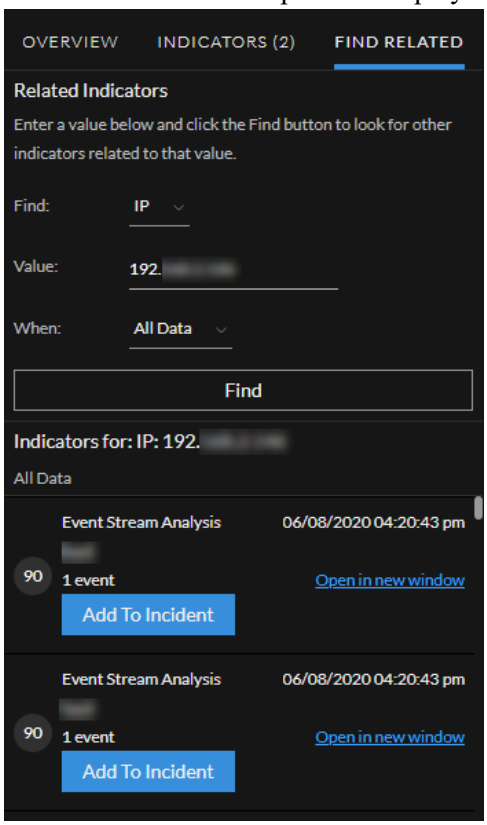


Find Related Indicators

Related Indicators are alerts that were not originally part of the selected incident, but they are related in some way to the incident. The relationship may or may not be obvious. For example, related indicators can involve one or more entities from the incident, but they can also be related due to some intelligence outside of NetWitness.

In the Incident Details view Related Indicators panel, you can search for an entity (such as IP, MAC, Host, Domain, User, Filename, or Hash) in other alerts outside of the current incident.

1. Go to **Respond > Incidents** and locate the incident that you want to view in the Incidents List.
2. Click the link in the **ID** or **NAME** field of the incident.
3. In the left panel of the Incident Details view, click the **FIND RELATED** tab.
   The Related Indicators panel is displayed.

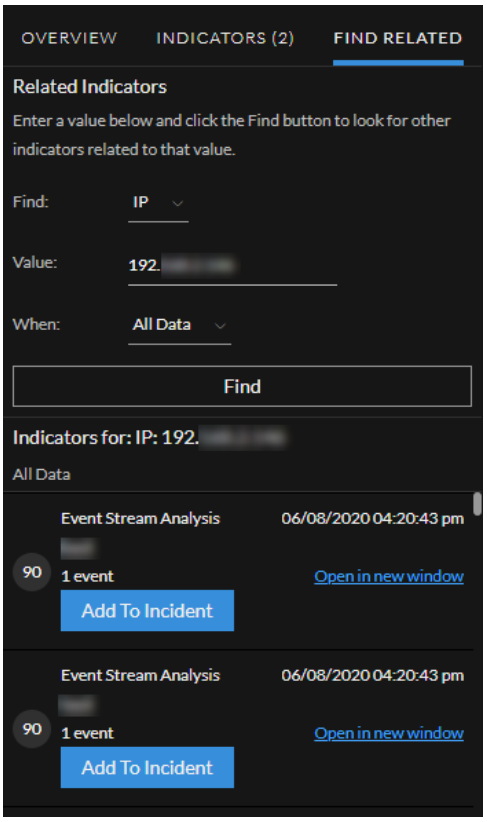


4. In the **Find** field, select the entity type to search, such as IP.
5. In the **Value** field, type a value for the entity, such as a specific IP address.
6. In the **When** field, select the time period to search, such as the Last 24 Hours.
7. Click **Find**.
   A list of related indicators (alerts) appear below the **Find** button in the **Indicators for** section. If an alert is not part of another incident, you can click the **Add to Incident** button to add the related indicator (alert) to the current incident. See Add Related Indicators to the Incident below.

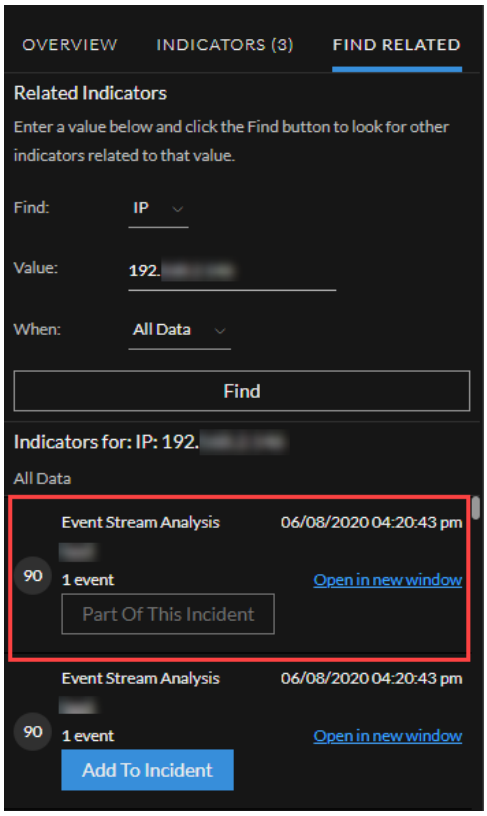
Add Related Indicators to the Incident

You can add related indicators (alerts) to the current incident from Related Indicators panel. An indicator that is already part of an incident cannot be part of another incident. In the search results, if an alert is not already part of an incident, it has an **Add to Incident** button.

1. In the Related Indicators panel, do a search to find related indicators. See Find Related Indicators above.



2. Review the alerts in the search results. The **Indicators for** section (below the Find button) lists the related indicators (alerts).
3. To inspect the details of an alert before adding it as a related indicator to the incident, you can click the **Open in New Window** link to view the alert details for that indicator.
4. For each alert that you want to add to the current incident as a related indicator, click the **Add to Incident** button.
   The button in the Related Indicators panel now shows **Part of This Incident**.



The selected related indicator adds to the Indicators panel. The Indicators tab now shows the additional indicator.



Investigate the Incident

To further investigate an incident within the Incident Details view, you can find links that take you to additional contextual information about the incident when it is available. This additional context can help you understand additional technical context and business context about a specific entity in the incident. It can also provide additional information that you may want to research to ensure that you understand the full scope of the incident.

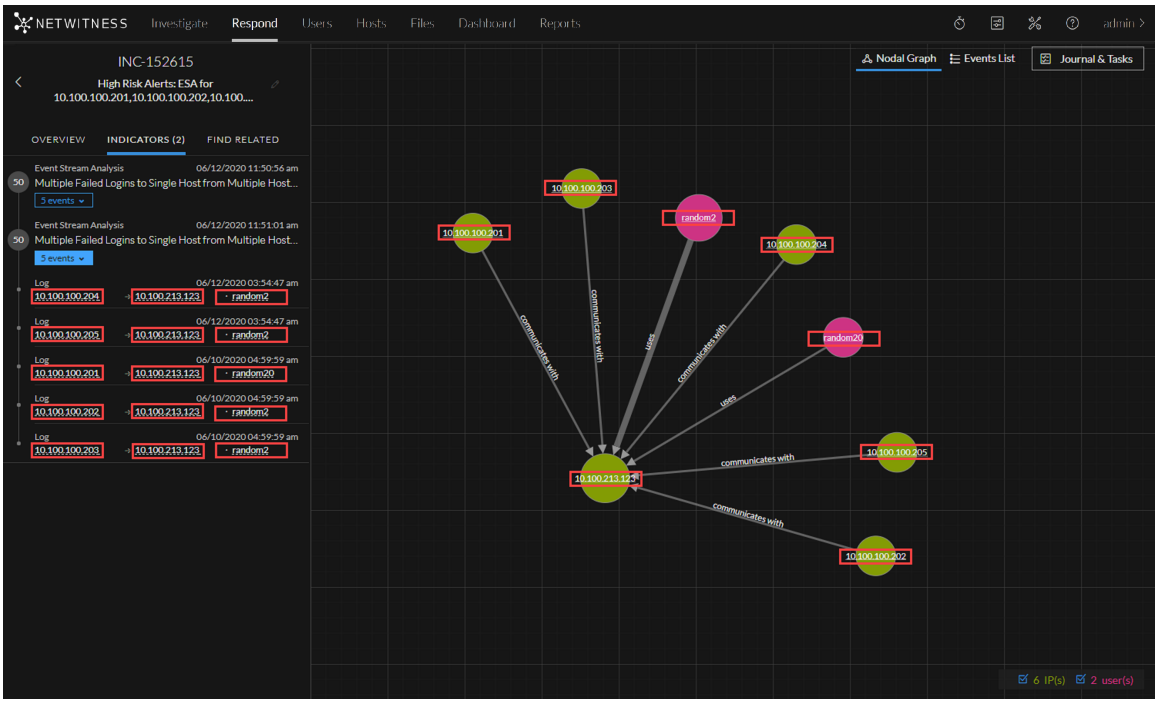
You can perform the following procedures to further investigate an incident:

- View Contextual Information
- Add an Entity to a Whitelist
- Create a List
- View the Reputation Status of a File Hash
- Pivot to the Investigate > Events View
- Pivot to the Hosts or Files View
- Pivot to NetWitness Endpoint Thick Client
- Pivot to Archer
- View Event Analysis Details for Indicators
- View User Entity Behavior Analytics for Indicators
- Document Steps Taken Outside of NetWitness
  - View the Journal Entries for an Incident
  - Add a Note
  - Delete a Note

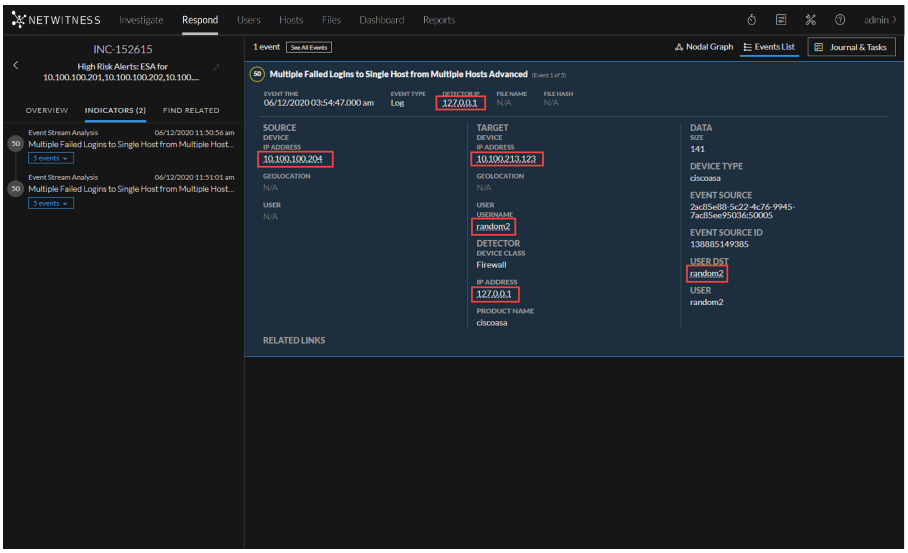
View Contextual Information

In the Indicators panel, Events List, or the Nodal Graph, you can view the underlined entities. If an entity is underlined, NetWitness is populating information about that entity type in the Context Hub. There may be additional information available about that entity type in the Context Hub.

The following figure shows underlined entities in the Indicators panel and the Nodal Graph.



The following figure shows underlined entities in the Events list details.

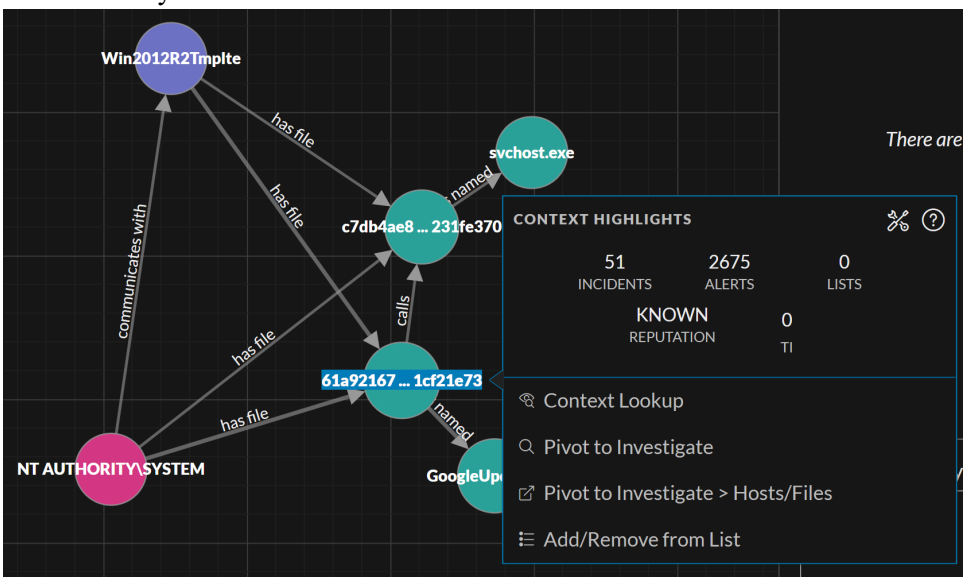


The Context Hub is preconfigured with meta fields mapped to the entities. NetWitness Respond and NetWitness Investigate use these default mappings for context lookup. For information about adding meta keys, see "Configure Settings for a Data Source" in the *Context Hub Configuration Guide*.

Caution: For the Context Lookup to work correctly in the Respond and Investigate views, NetWitness recommends that when mapping meta keys in the **Admin > System > Investigation > Context Lookup** tab, you add only meta keys to the Meta Key Mappings, not fields in the MongoDB. For example, ip.address is a meta key and ip_address is not a meta key (it is a field in the MongoDB).

To view contextual information:

1. In the Indicators panel, Events List, or the Nodal Graph, left or right-click an underlined entity.
   A context tooltip appears with a quick summary of the type of context data that is available for the selected entity.

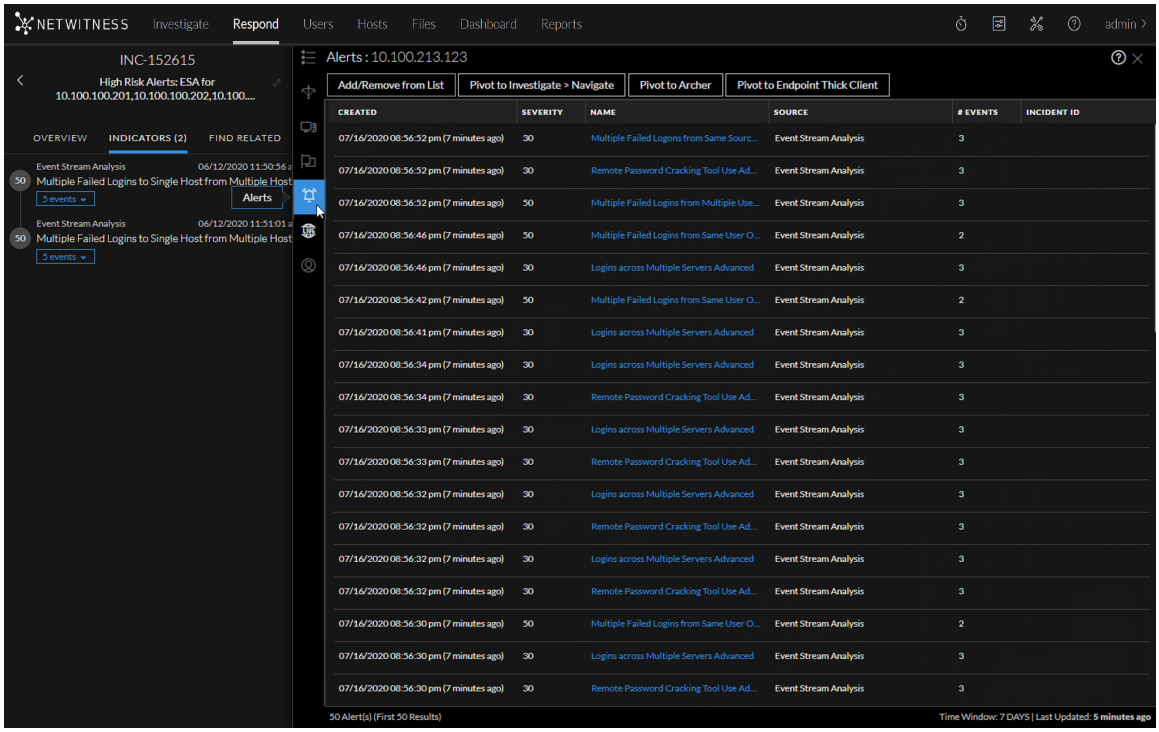


   The information in the **Context Highlights** section helps you to determine the actions that you would like to take. It can show related data for Incidents, Alerts, Lists, Endpoint//, Criticality, Asset Risk, Reputation, and Threat Intelligence (TI). Depending on your data, you may be able to click these items for more information.
   The above example shows 0 related incidents, 6800 alerts, 0 lists for the selected host, 0 incidents for TI, and no information available for Endpoint, Live Connect, Criticality, and Asset Risk.
   TI information comes from the STIX data source configured in Context Hub. For more information, see the *Context Hub Configuration Guide*.
   The **Actions** section lists the available actions. In the above example, the Add/Remove from List, Pivot to Investigate > Navigate, Pivot to Investigate > Hosts/Files and Pivot to Endpoint Thick Client options are available.

Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer data source is not responding. Check that the Archer configuration is enabled and configured properly.

For more information, see Pivot to the Investigate > Events View, Pivot to Archer, Pivot to NetWitness Endpoint Thick Client, Pivot to the Hosts or Files View, and Add an Entity to a Whitelist.

2. To see more details about the selected entity, click the **View Context** button.
   The Context Lookup panel opens and shows all of the information related to the entity.
   The following example shows contextual information for a selected host. It lists all of the incidents that mention that host.

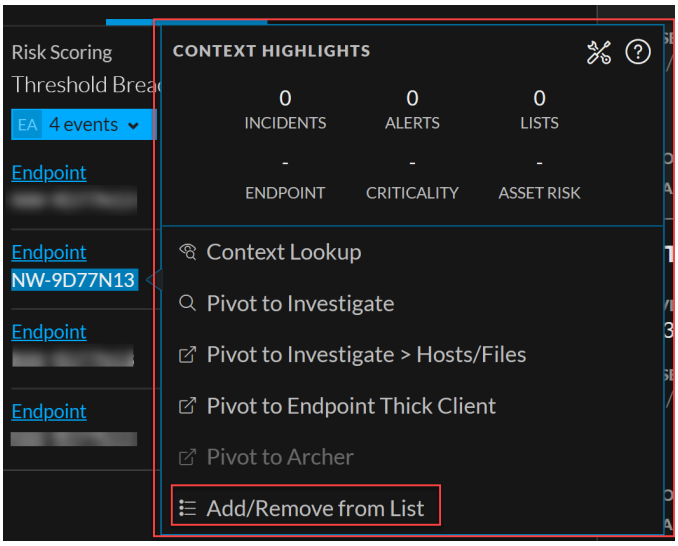


   To understand the different views within the Context Hub Lookup panel, see Context Lookup Panel - Respond View.

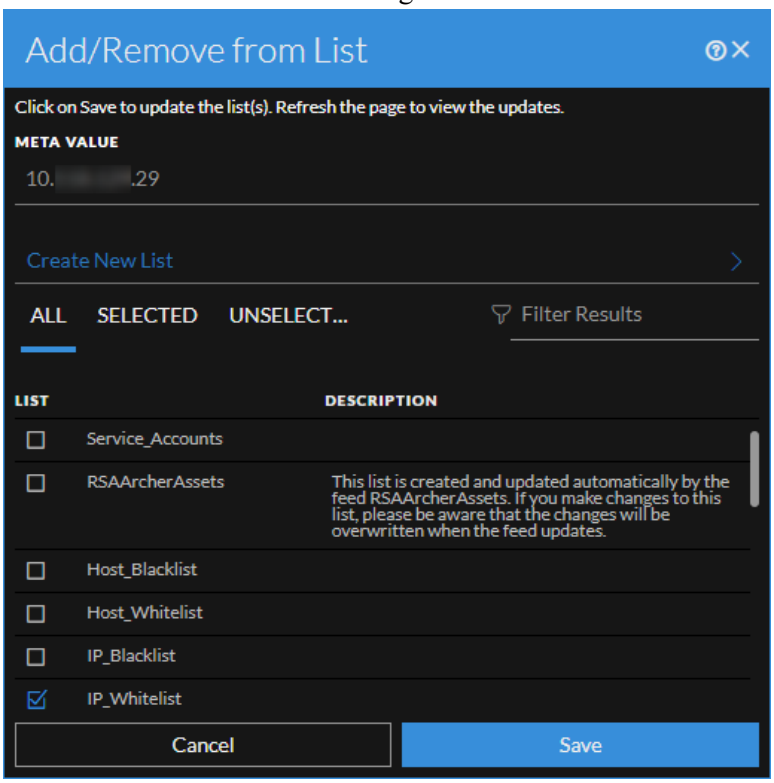
Add an Entity to a Whitelist

You can add any underlined entity to a list, such as a Whitelist or Blacklist, from a context tooltip. For example, to reduce false positives, you may want to whitelist an underlined domain to exclude it from the related entities.

1. In the Indicators panel, Events List, or the Nodal Graph, left or right-click the underlined entity that you would like to add to a Context Hub list.
   A context tooltip appears showing the available actions.



2. In the **ACTIONS** section of the tooltip, click **Add/Remove from List**.
   The Add/Remove from List dialog shows the available lists.



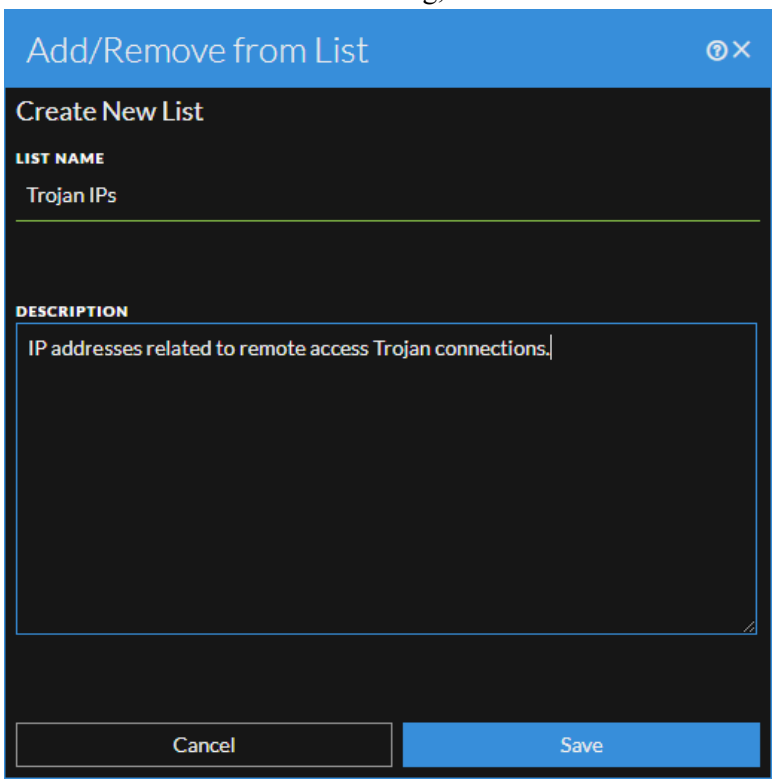
3. Select one or more lists and click **Save**.
   The entity appears on the selected lists.
   Add/Remove from List Dialog provides additional information.

Create a List

You can create lists in Context Hub from the Respond view. In addition to using lists to whitelist and blacklist entities, you can use lists to monitor entities for abnormal behavior. For example, to improve the visibility of a suspicious IP address and Domain under investigation, you may want to include them in two separate lists. One list could be for domains suspected of being related to command and control connections, and another list could be for IP addresses related to remote access Trojan connections. You can then identify indicators of compromise using these lists.

To create a list in Context Hub:

1. In the Indicators panel, Events List, or the Nodal Graph, left or right-click the underlined entity that you would like to add to a Context Hub list.
   A context tooltip opens showing the available actions.
2. In the **Actions** section of the tooltip, click **Add/Remove from List**.
3. In the Add/Remove from List dialog, click **Create New List**.



4. Type a unique **List Name** for the list. The list name is not case sensitive.
5. (Optional) Type a **Description** for the list.
   Analysts with the appropriate permissions can also export lists in CSV format to send to other analysts for further tracking and analysis. The *Context Hub Configuration Guide* provides additional information.

View the Reputation Status of a File Hash

The File Reputation service available on RSA Live checks the reputation of every file hash against an extensive database of known file hashes updated in real-time. The file reputation is displayed in the Investigate and Respond views. In the View Context lookup, if the reputation status changes, Context Hub notifies the change in reputation status to all Endpoint servers. Information about the file hash such as any suspicious or malicious activity on the file is populated from Context Hub. There may be additional information available about that entity in the Context Hub.

The following table describes the file hash reputations.

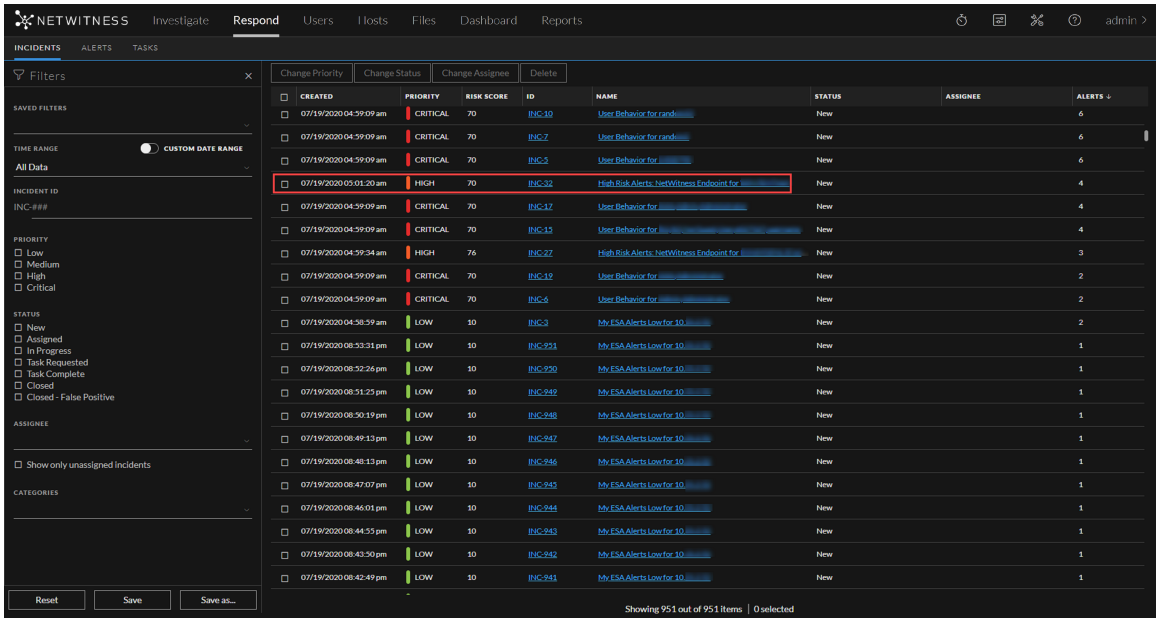
Reputation	Description
Malicious	File hash is labeled as malicious.
Suspicious	File hash is suspected to be malicious.
Unknown	File hash is not known.
Known	File hash information is known to the file reputation service and does not have any previous bad record.
Known Good	File hash information is known good, such as files signed by Microsoft or NetWitness.
Invalid	File hash format is invalid.

Note: A reputation status is visible for a file hash entity only and File Reputation service supports a maximum of 10 million files for a reputation of file hash.

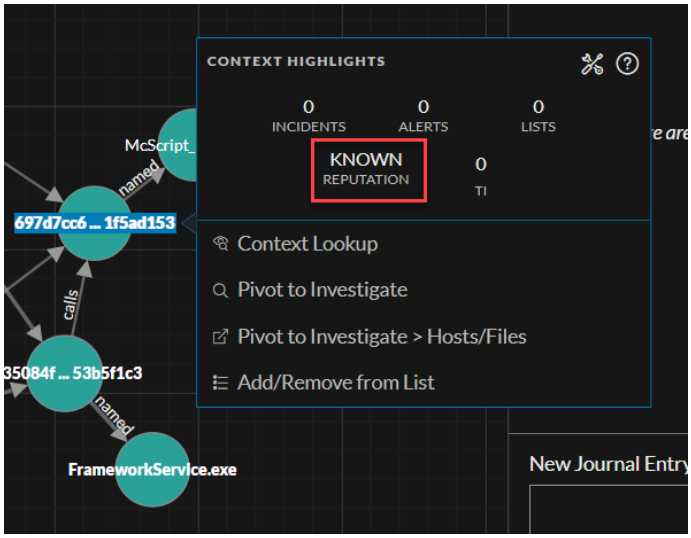
The suspicious or malicious files are available for further analysis in the **Investigate > Navigate** view and **Investigate > Events** view. For more information on the file reputation service, see the *Live Services Management Guide* and the *NetWitness Endpoint User Guide*.

To view the reputation of a file hash:

1. Go to **Respond > Incidents**.
2. In the Incidents List view, choose an incident to view and then click the link in the **ID** or **Name** column for that incident.

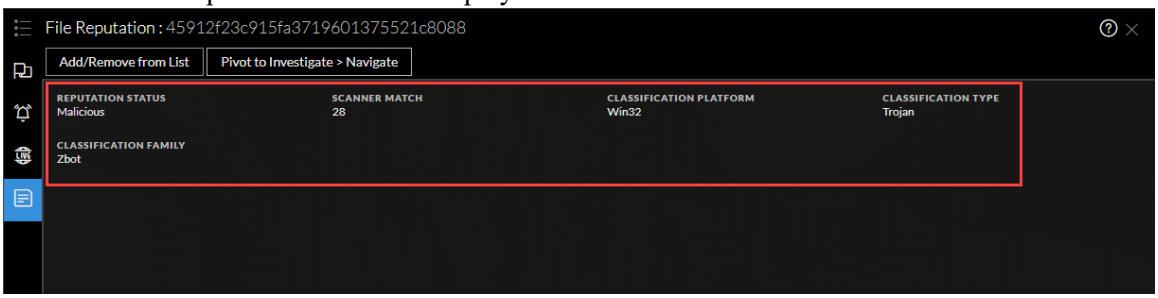


3. In the Incident Details view, left or right click the file hash entity.
   The context tooltip displays the reputation status of the selected file hash entity.



4. Click **Reputation** to view the reputation status information.

5. Click the **File Reputation** icon  to view further details.
   The details for reputation status are displayed.



Pivot to the Investigate > Events View

For a more thorough investigation of the incident, you can access the **Investigate** > **Events**.

1. In the Indicators panel, Events List, or the Nodal Graph, left or right click any underlined entity to access a context tooltip.
2. In the context tooltip panel, select **Pivot to Investigate**.
   The Events view opens, which enables you to perform a deep dive investigation.

For more information, see the *NetWitness Investigate User Guide*. For troubleshooting information with the Investigate > Events link see the *Alerting with ESA Correlation Rules User Guide*.

Pivot to the Hosts or Files View

For a more thorough investigation about specific Hosts and Files, you can access the Hosts and Files views.

1. In the Indicators panel, Events List, or the Nodal Graph, hover over any entity to access a context tooltip.
2. In the context tooltip panel, select **Pivot to Investigate > Hosts/Files**.
   If you hover over a host or IP or MAC address entity and click **Pivot to Investigate > Hosts/Files**, it displays the Hosts view with a specific host listed.
   If you hover over a filename or file hash entity and click **Pivot to Investigate > Hosts/Files** it displays the Files view with a specific file listed.

Note: By default, the search for entities is on the previously selected Endpoint Server. However, you can select a different Endpoint Server to fetch the information or data.

For more information, see the *NetWitness Investigate User Guide*.

Pivot to NetWitness Endpoint Thick Client

If you have the NetWitness Endpoint thick client application installed, you can launch it through the context tooltip. From there, you can further investigate a suspicious IP address, Host, or MAC address.

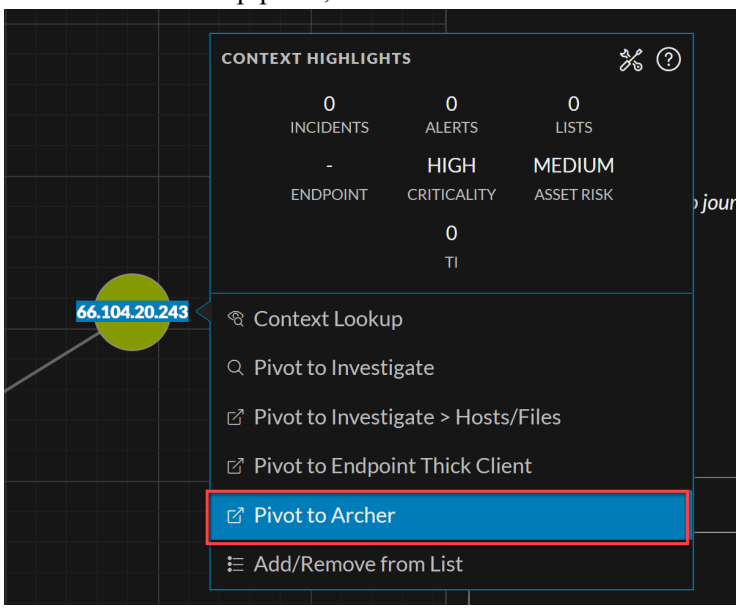
1. In the Indicators panel, Events List, or the Nodal Graph, hover over any underlined entity to access a context tooltip.
2. In the context tooltip panel, select **Pivot to Endpoint Thick Client**.
   The NetWitness Endpoint thick client application opens outside of your web browser.

For more information on the thick client, see the *NetWitness Endpoint User Guide*.

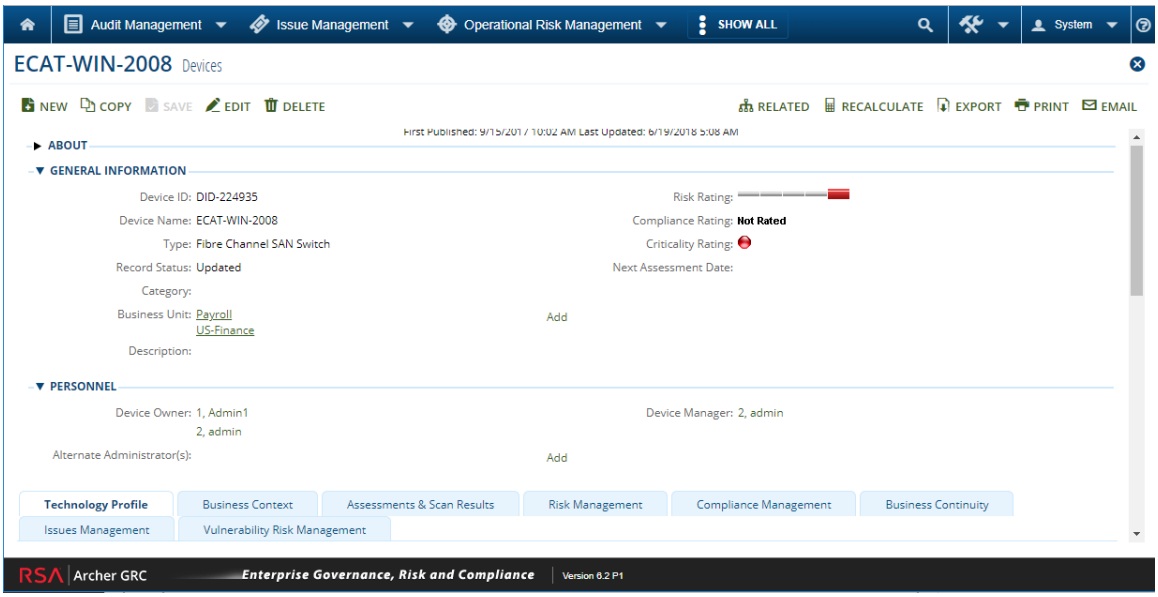
Pivot to Archer

For viewing more details about the device in Archer Cyber Incident & Breach Response, you can pivot to the device details page. This information is displayed only for IP address, host, and Mac address.

1. In the Indicators panel, Events List, or the Nodal Graph, left or right click any underlined entity (IP address, host, and Mac address) to access a context tooltip.
2. In the context tooltip panel, select **Pivot to Archer**.



3. The device details page in **Archer Cyber Incident & Breach Response** opens if you are logged in to the application, otherwise the login screen is displayed.



Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer Datasource is not responding. Check that the Archer configuration is enabled and configured properly.

For more information, see the *NetWitness Archer Integration Guide*.

View Event Analysis Details for Indicators

In the Incident Details view Indicators panel, you can drill deeper into the events associated with the listed indicators to get a better understanding of the events. In the Events panel, you can view raw events and metadata with interactive features that enhance your ability to find meaningful patterns in the data. You can examine network, log, and endpoint events in the Events panel. The Events panel in the Respond view shows the Events view from Investigate for specific indicator events. For detailed information about the Events view, see the *NetWitness Investigate User Guide*.

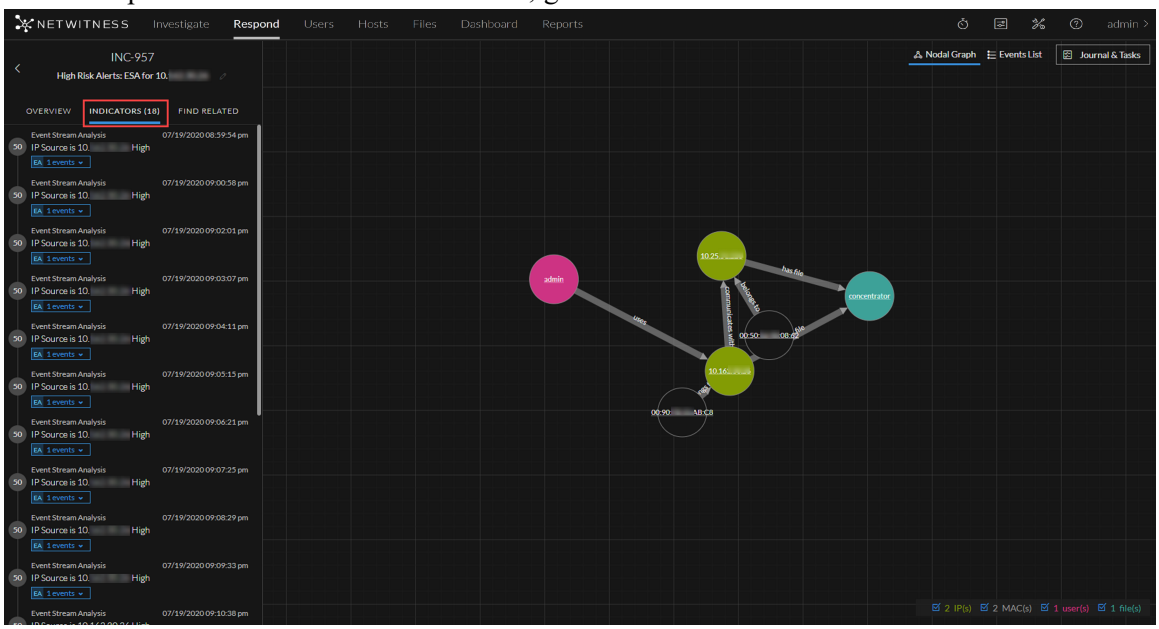
Note: You must have the following Investigate-server permissions to view the Events panel in the Respond view:
`event.read`
`content.reconstruct`
`content.export`
The Events view requires all Core services to be on NetWitness Platform 11.4 or later.

Migration Considerations

Migrated incidents from NetWitness versions before 11.2 will not show the Events panel in the Respond Incident Details view Indicators panel. Likewise, if you use alerts that were migrated from versions before 11.2 to create incidents in 11.5, you will also not be able to view the Events panel in the Respond view for those incidents.

To access event analysis details for an event in the Indicators panel:

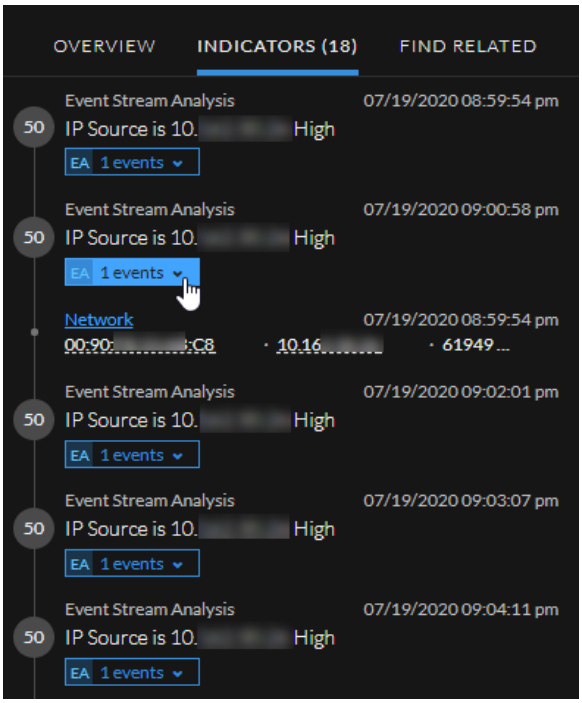
1. Go to **Respond > Incidents**.
2. In the Incidents List view, choose an incident to view and then click the link in the **ID** or **NAME** column for that incident.
   The Incident Details view is displayed.
3. In the left panel of the Incident Details view, go to the **Indicators** tab.



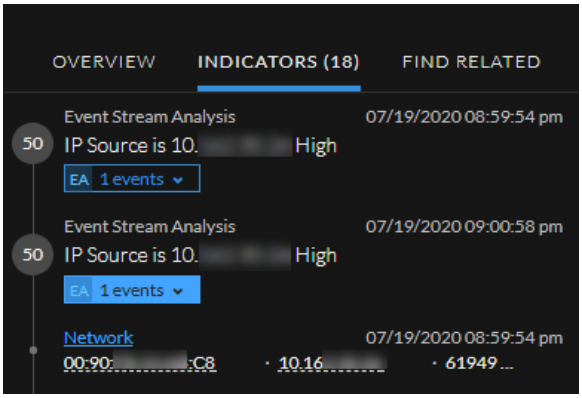
Data source information is shown above the names of the indicators. You can also see the creation date and time as well as the number of events in the indicator. If event analysis (EA) information is available, you can see an **EA** icon in front of the event count as shown in the following figure.



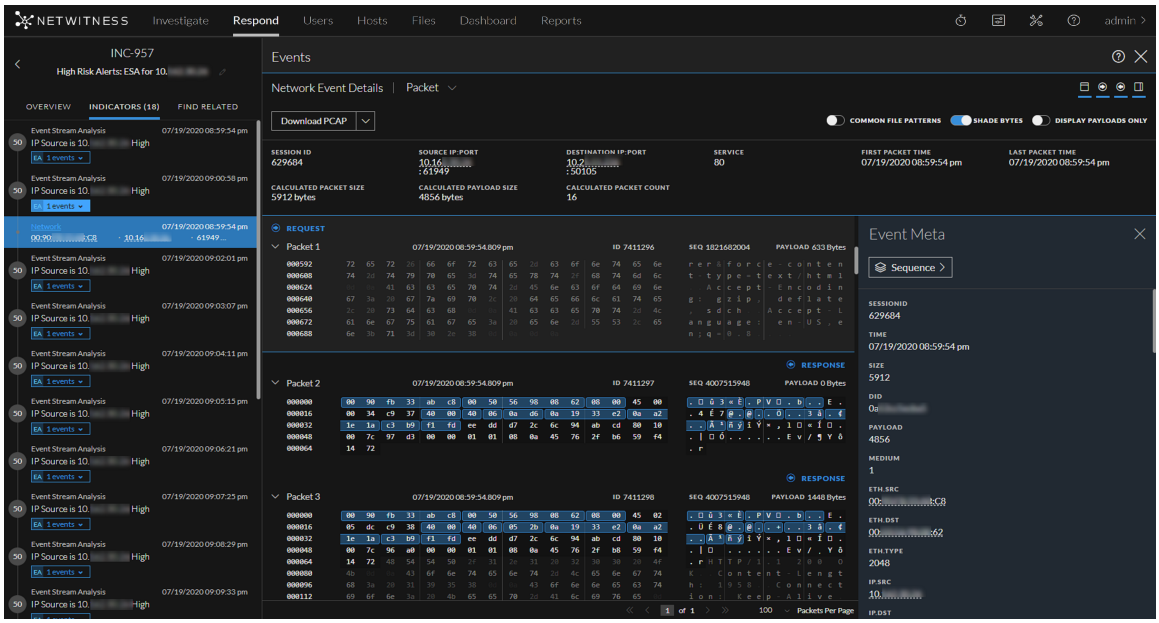
4. Click an event count with an **EA** icon to view additional event information.



5. Click an event type hyperlink within the event to open the Events panel. In the following example, the event type is Network.



The Events panel shows event details for the event, such as packet analysis details. The information available can vary based on the event type.

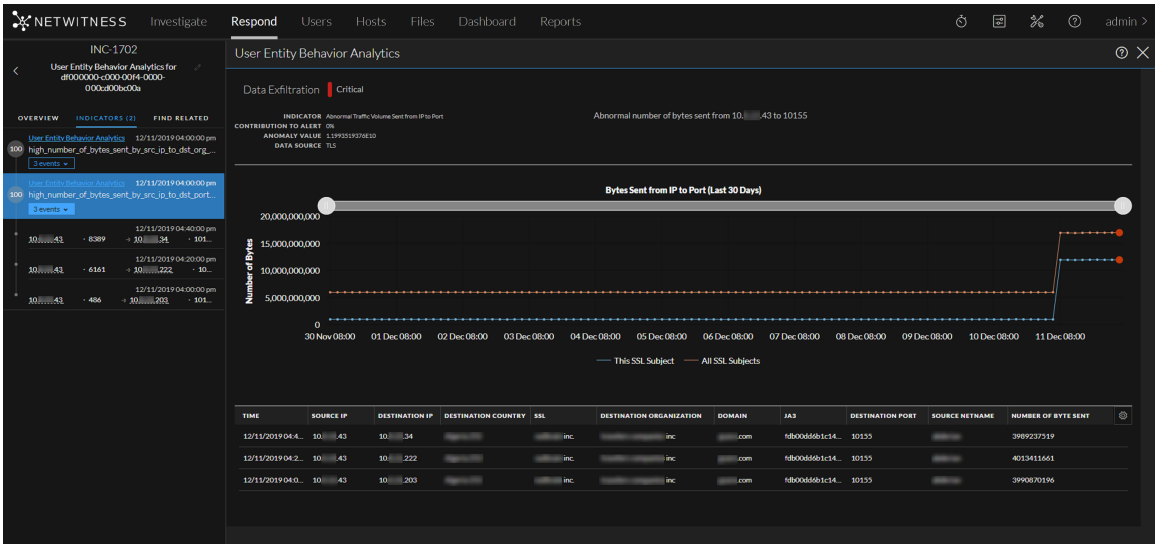


For detailed information about the Events view, see the *NetWitness Investigate User Guide*.

Note: If you want to send the Events URL link to another analyst, you can copy the event type hyperlink, for example Network.

View User Entity Behavior Analytics for Indicators

NetWitness UEBA (User and Entity Behavior Analytics) is an advanced analytics solution for discovering, investigating, and monitoring risky behaviors across all users and entities in your network environment. You can access UEBA from the Respond Incident Details view Indicators panel. Indicators with a **User Entity Behavior Analytics** hyperlink have additional UEBA information available. For detailed information about UEBA, see the *NetWitness UEBA User Guide*.

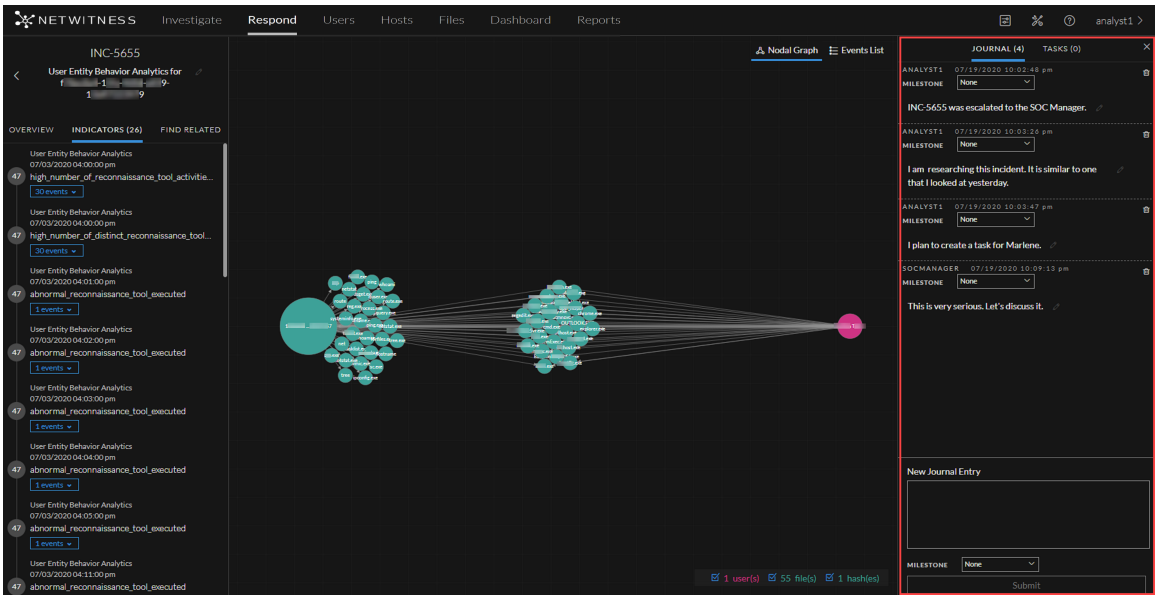


Document Steps Taken Outside of NetWitness

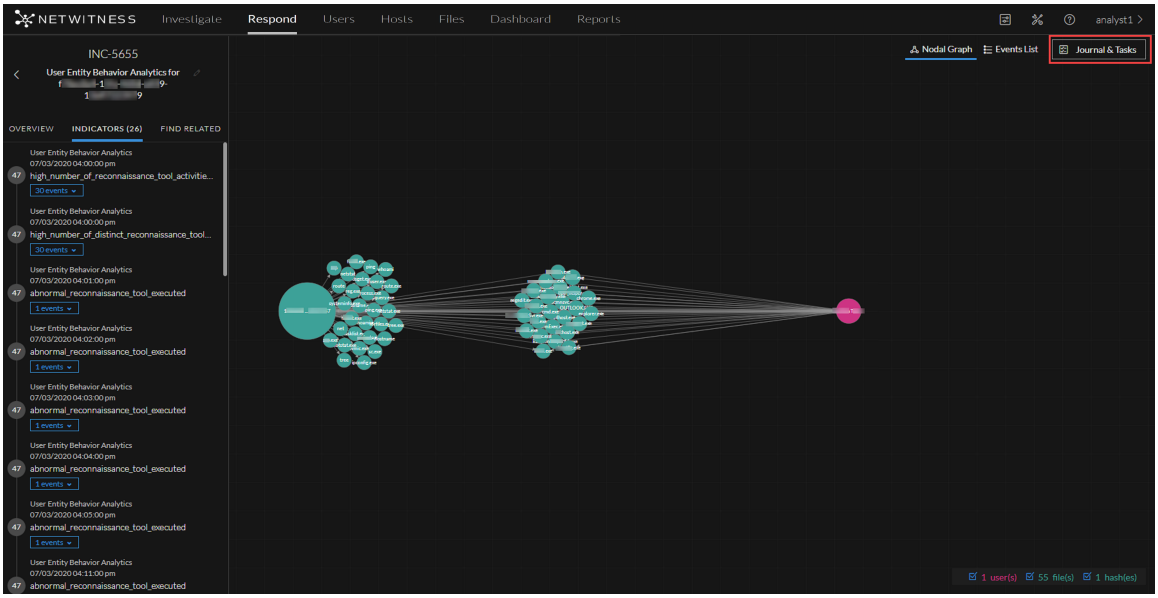
The journal shows notes added by analysts and it enables you to collaborate with your peers. You can post notes to a journal, add Investigation Milestone tags (Reconnaissance, Delivery, Exploitation, Installation, Command and Control, Action on Objective, Containment, Eradication, and Closure), and view the history of activity on your incident.

View the Journal Entries for an Incident

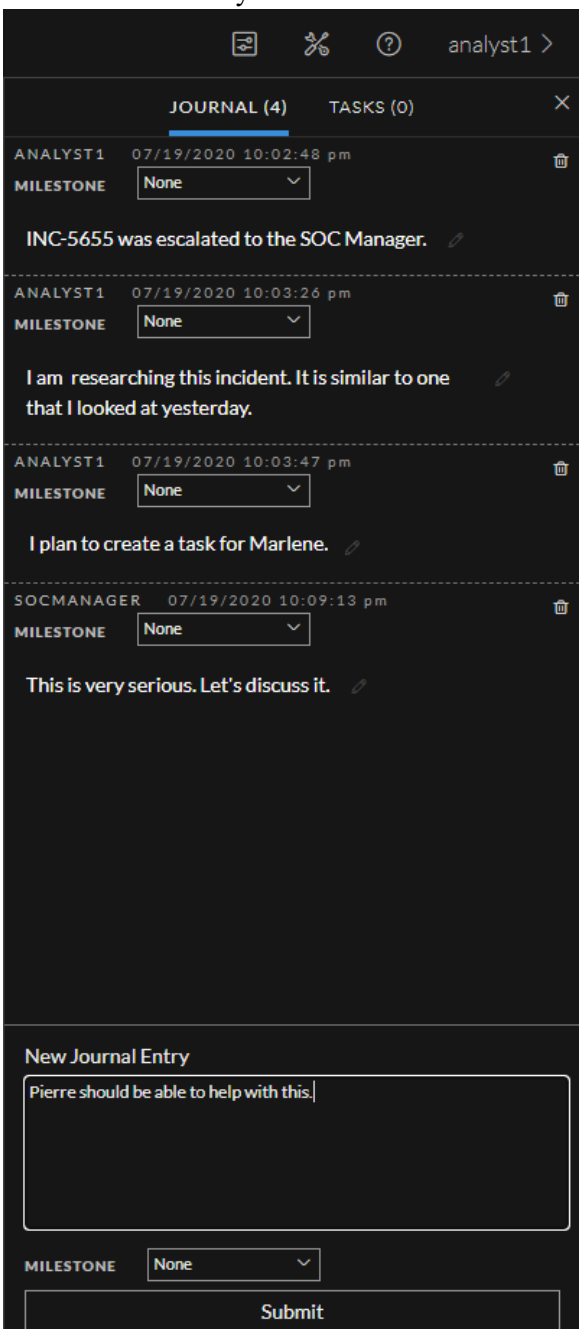
The Journal is on the right side of the Incident Details view.



If you do not see the Journal, in the toolbar, click **Journal & Tasks**.



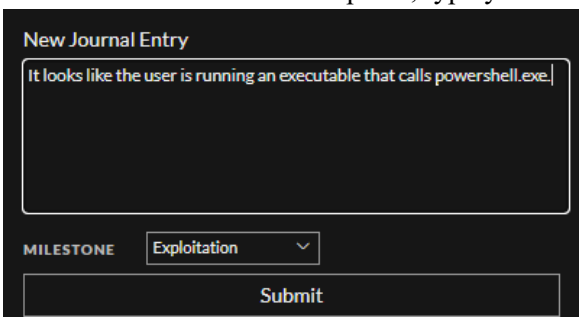
The Journal shows the history of activity on an incident. For each journal entry, you can see the author and time of the entry.



Add a Note

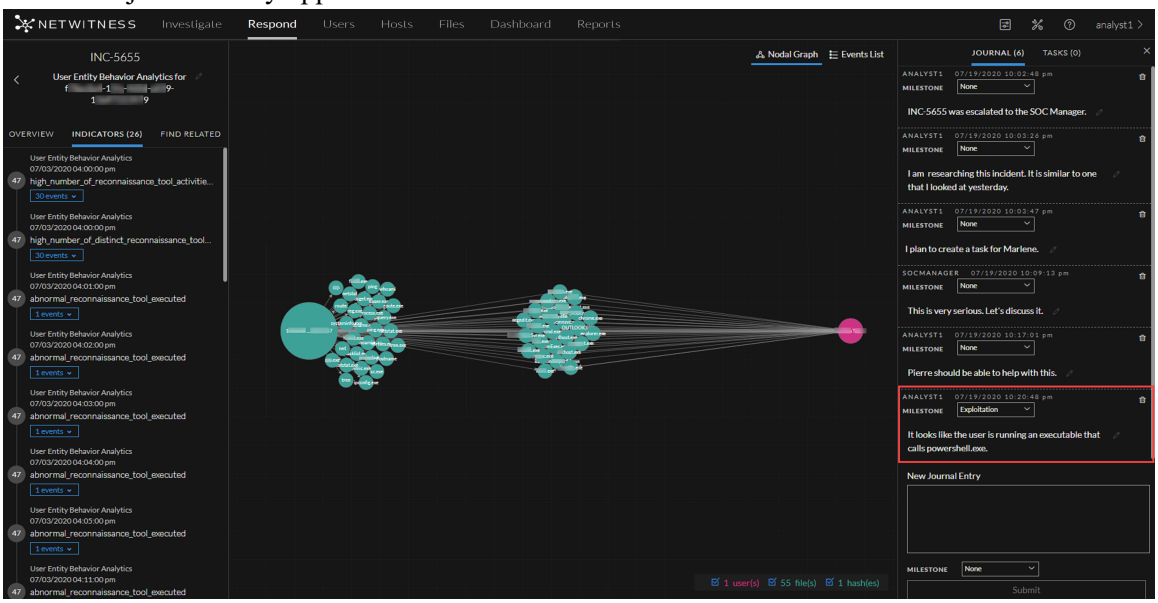
Typically, you will want to add a note to allow another analyst to understand the incident, or add a note for posterity so that your investigative steps are documented.

1. At the bottom of the Journal panel, type your note in the **New Journal Entry** box.



2. (Optional) Select an Investigation Milestone from the drop-down list (Reconnaissance, Delivery, Exploitation, Installation, Command and Control, Action On Objective, Containment, Eradication, and Closure).

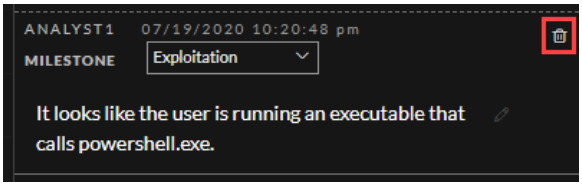
3. After you finish your note, click, **Submit**.
   Your new journal entry appears in the Journal.



Delete a Note

1. In the Journal panel, locate the journal entry that you would like to delete.

2. Click the trash can (delete) icon  next to the journal entry.



3. In the confirmation dialog that appears, click **OK** to confirm that you want to delete the journal entry. This action cannot be reversed.

Escalate or Remediate the Incident

You may want to escalate an incident, assign incidents to another Analyst, or change the status and priority of an incident as you gather more information about it. This is useful if, for example, you upgrade the priority of an incident from high to critical after determining that the incident is a major breach. You may also want to send the incident to Archer Cyber Incident & Breach Response for additional analysis and action.

You can perform the following procedures to escalate or remediate an incident:

- Send an Incident to Archer
- View All Incidents Sent to Archer
- Update an Incident
- Change Incident Status
- Change Events Retention
- Obtain Retention Usage Details
- Change Incident Priority
- Assign Incidents to Other Analysts
- Rename an Incident
- View All Incident Tasks
- Filter the Tasks List
- Remove My Filters from the Tasks List
- Create a Task
- Find a Task
- Modify a Task
- Delete a Task
- Close an Incident

Send an Incident to Archer

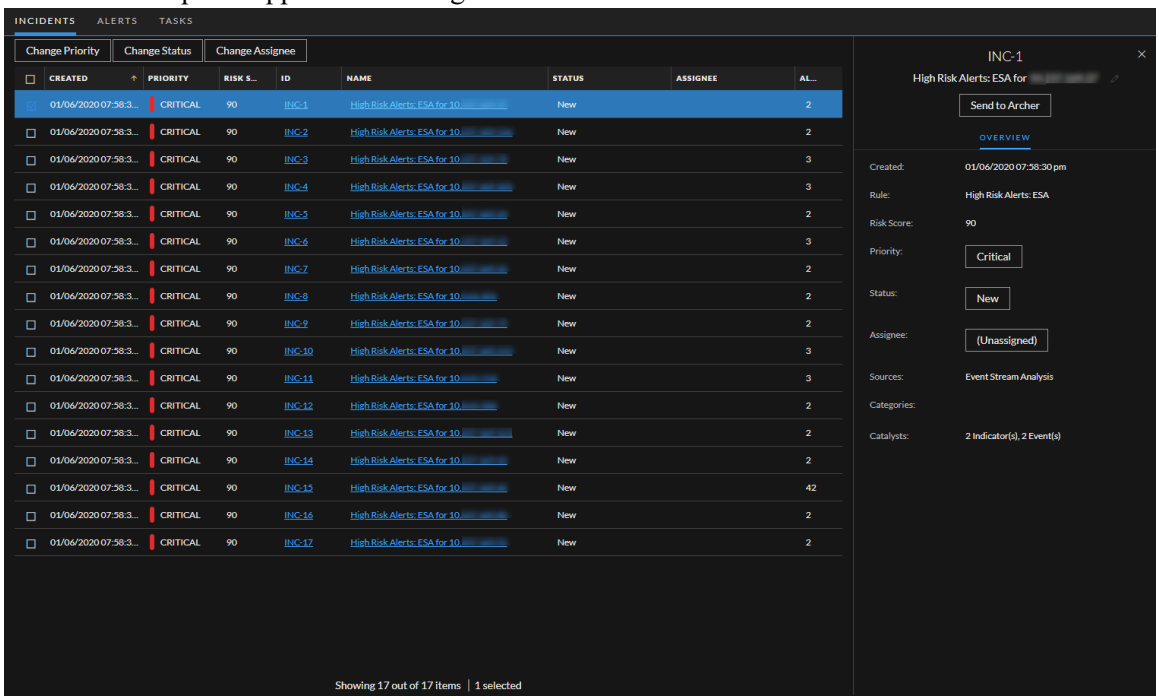
Note: This option is available in NetWitness Version 11.2 and later. If Archer is configured as a data source in Context Hub, you can send incidents to Archer and you can see the Send to Archer option and Sent to Archer Status in NetWitness Respond.

When you send an incident to Archer, a Sent to Archer notification appears within the incident. When configured, the NetWitness Platform can start additional business processes in Archer Cyber Incident & Breach Response. You can view all of the incidents that were sent to Archer Cyber Incident & Breach Response using the filter in the Incident Lists view.

You send an incident to Archer by clicking the Send to Archer button in the Overview panel in the Incident Lists view or the Incident Details view.

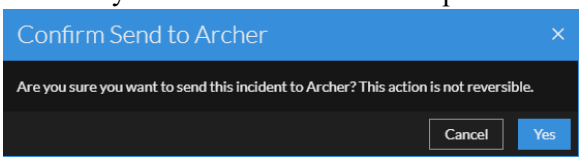
Caution: The **Send to Archer** action is not reversible.

1. Go to **Respond > Incidents**.
2. From the Incidents List view, click the incident that you want to send to Archer Cyber Incident & Breach Response.
   The Overview panel appears on the right.

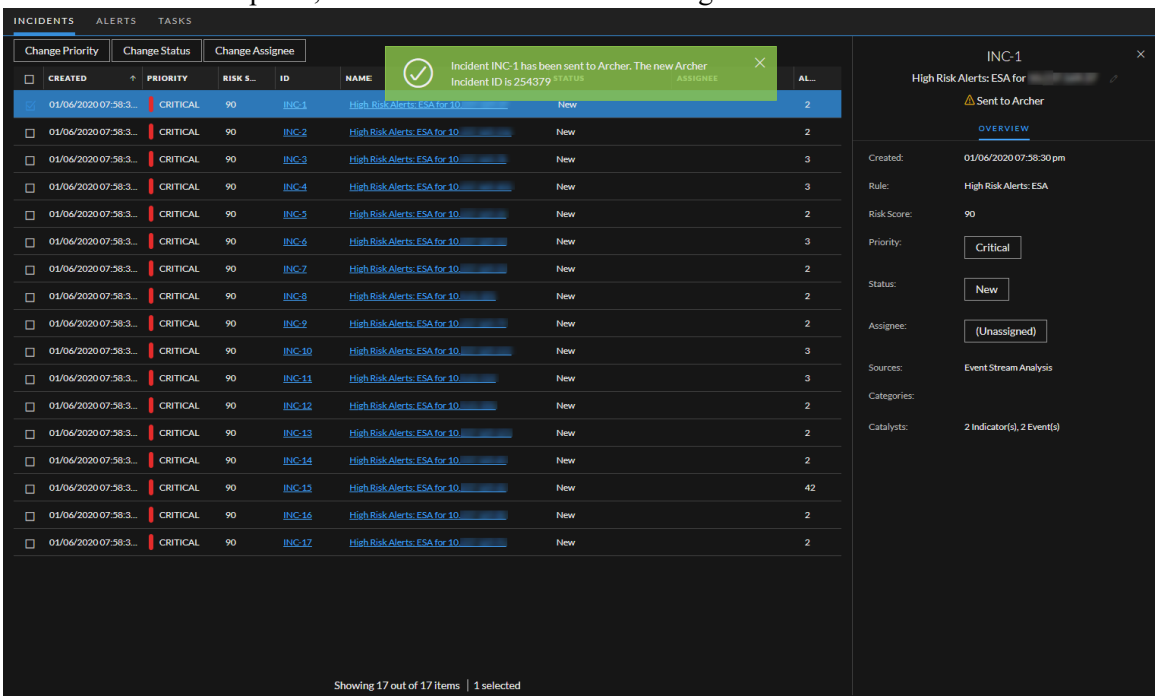


3. In the Overview panel, click **Send to Archer**.

4. Read the **Confirm Send to Archer** dialog and then click **Yes** to confirm sending the incident to Archer Cyber Incident & Breach Response. This action is not reversible.



You will receive a confirmation that the incident was sent to Archer along with an Archer incident ID. In the Overview panel, the Send to Archer button changes to Sent to Archer.



In the Incident Details view (click the link in the ID or NAME field of the incident sent to Archer) you can see the Sent to Archer notification above the Overview and Indicators panels. If you open the Journal, you can see a system journal entry that shows that the incident was sent to Archer and it now has an Archer ID number.



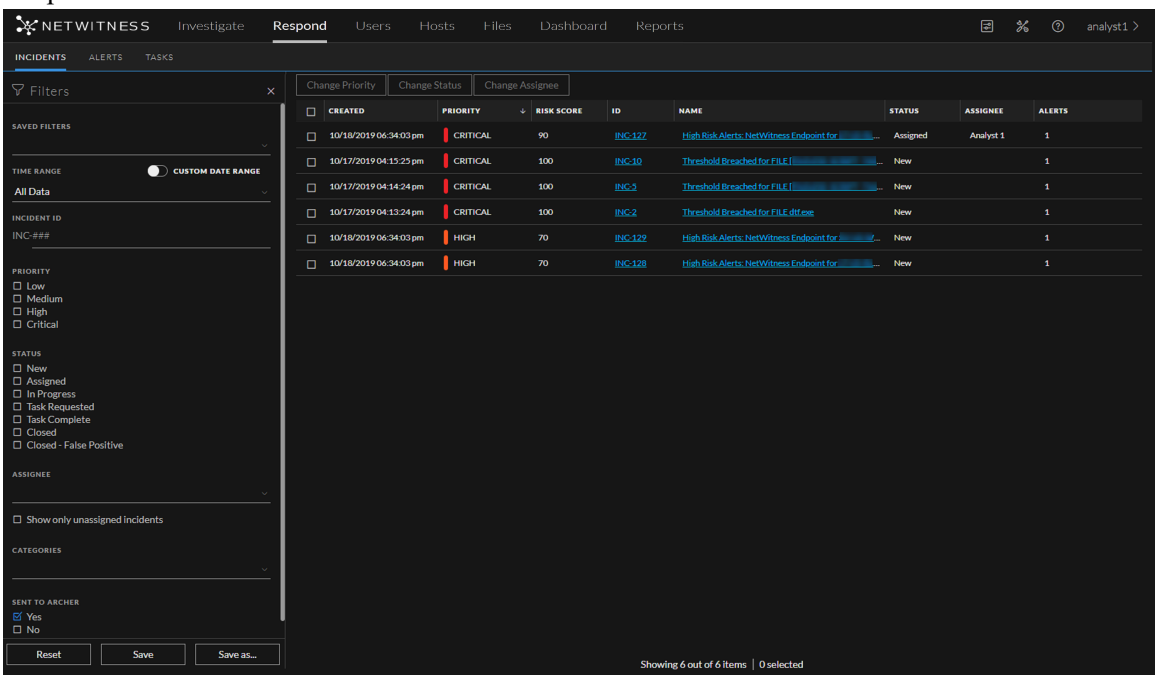
View All Incidents Sent to Archer

Note: This option is available in NetWitness Version 11.2 and later. If Archer is configured as a data source in Context Hub, you can send incidents to Archer and you will be able to see the Sent to Archer option and Sent to Archer Status in NetWitness Respond.

You can view incidents sent to Archer Cyber Incident & Breach Response using the Filter.

1. Go to **Respond > Incidents**. The Incidents List is displayed.
2. If you cannot see the Filters panel, in the Incident List view toolbar, click .
3. In the Filters panel, under **Sent To Archer**, select **Yes**.
   The incidents list will be filtered to show incidents that were sent to Archer Cyber Incident & Breach

Response.



Update an Incident

You can update an incident from several places. You can change the priority, status, or assignee from the Incident List view and the Incident Details view. For example, if you are an Analyst, you may want to assign yourself a case from the Incident List view if you see that it is related to another case you are working on. If you are an SOC Manager or an Administrator, you may want to view unassigned incidents from the Incident List view and assign the incidents as they come in. SOC Managers and Administrators can do bulk updates of the priority, status, or assignee instead of updating them one incident at a time.

From the Details view, you might want to change the status to In Progress once you begin working on an incident, and then update it to Closed or Closed - False Positive after you resolve the issue. Or you might change the priority of the incident to Medium or High as you determine the details of the case.

Change Incident Status

When an incident first appears in the incident list, it has an initial status of New. You can update the status as you complete your work on the incident. The following statuses are available:

- Reopen
- In Progress
- Task Requested
- Task Complete

- Closed
- Closed - False Positive

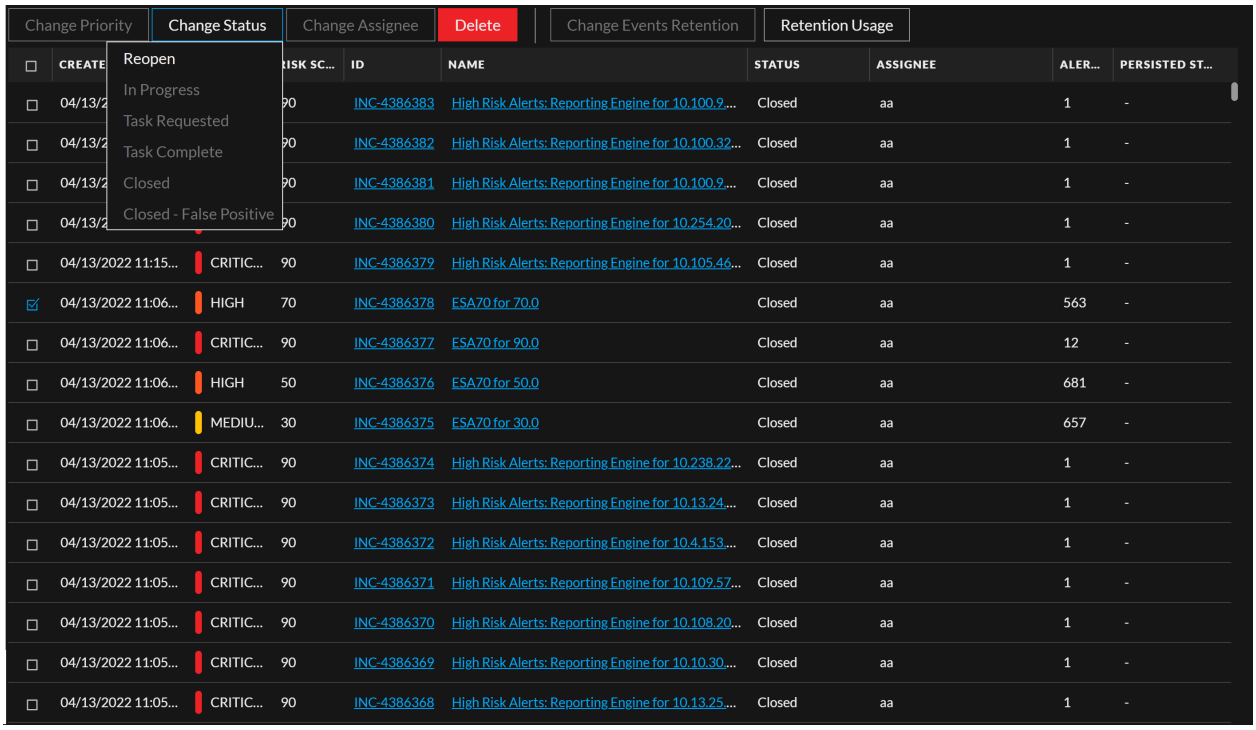
Note: New and Assigned statuses under the Change Status drop-down list are removed in the version 12.0 and later.

Status Change Workflow

The table below lists all the statuses and provides information about specific Status Change Workflow.

Status	New	Reopen	Assigned	In Progress	Task Requested	Task Complete	Closed / Closed - False Positive
New	No	No	Yes	Yes	No	No	Yes
Reopen	No	No	Yes	Yes	No	No	Yes
Assigned	No	No	No	Yes	No	No	Yes
In Progress	No	No	No	No	Yes	Yes	Yes
Task Requested	No	No	No	Yes	No	Yes	Yes
Task Complete	No	No	No	Yes	Yes	No	Yes
Closed / Closed - False Positive	No	Yes	No	No	No	No	No

Note: When you select an incident and click Change Status, all the invalid statuses are grayed out under the Change Status drop-down list. This is not applicable for multi-select of incidents. Refer the following figure.

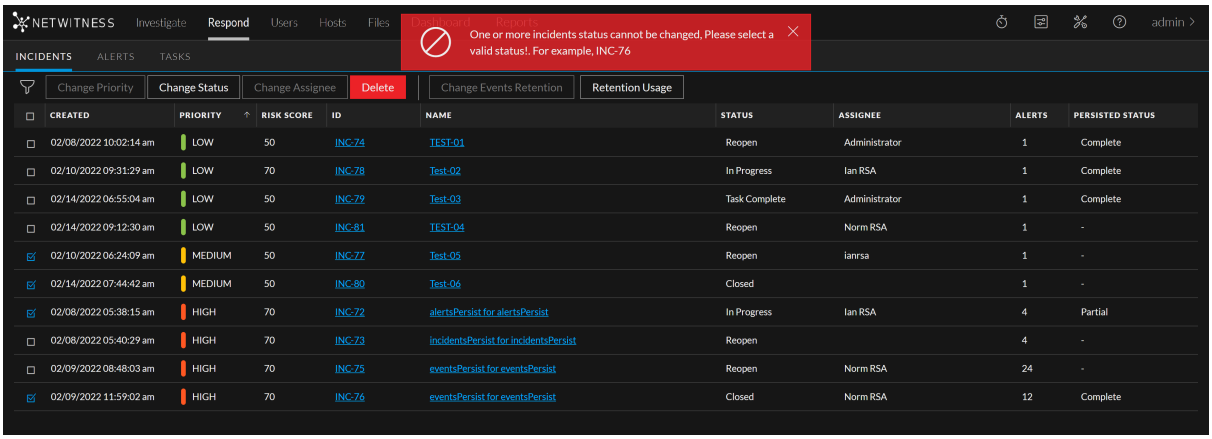


To update the status of multiple incidents:

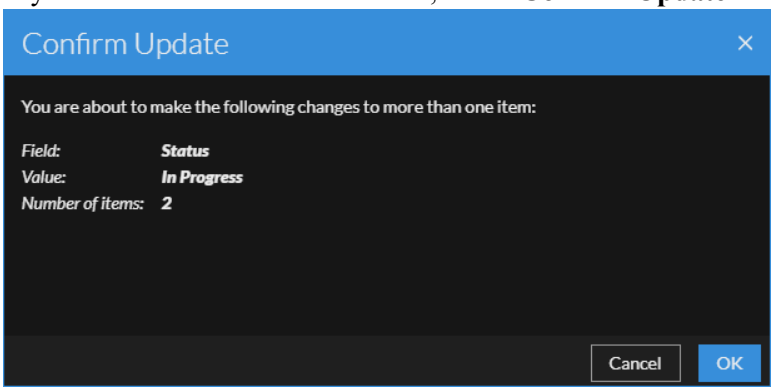
1. In the Incidents List view, select one or more incidents that you would like to change. To select all of the incidents on the page, select the box in the incidents list header row. The number of incidents selected appears in the incidents list footer.
2. Click **Change Status** and select a status from the drop-down list. In this example, the current status is Assigned, but the Assignee would like to change it to In Progress for the selected incidents.



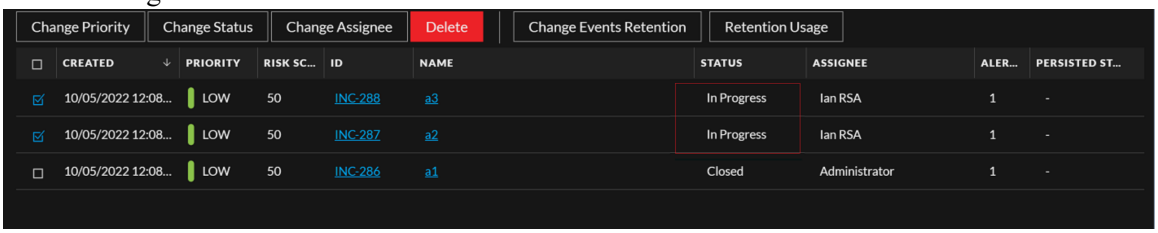
Note: The incident status can be changed to **Reopen** only if the current status of the incident is **Closed** or **Closed - False Positive**. This is also applicable when multiple incidents are selected. Even if one of the multiple incidents selected has the status other than **Closed or Closed - False Positive**, the error message **One or more incidents status cannot be changed, Please select a valid status!. For example, INC-x** is displayed. Refer the following figure.



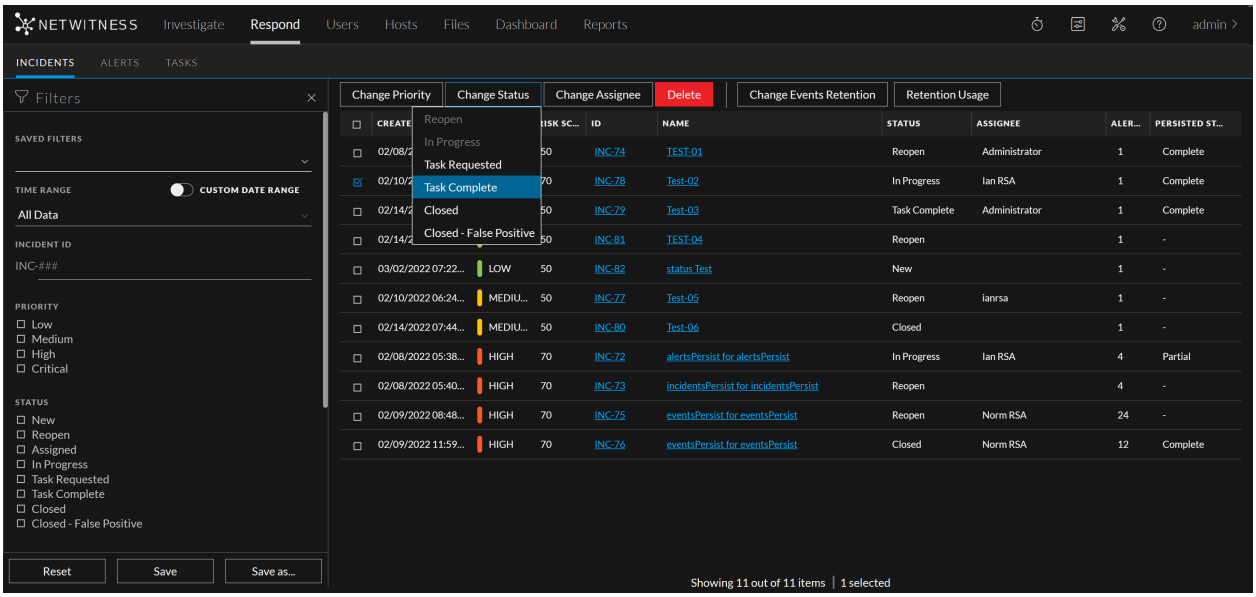
3. If you select more than one incident, in the **Confirm Update** dialog, click **OK**.



You can see a successful change notification. In this example, the status of the updated incidents now show In Progress.

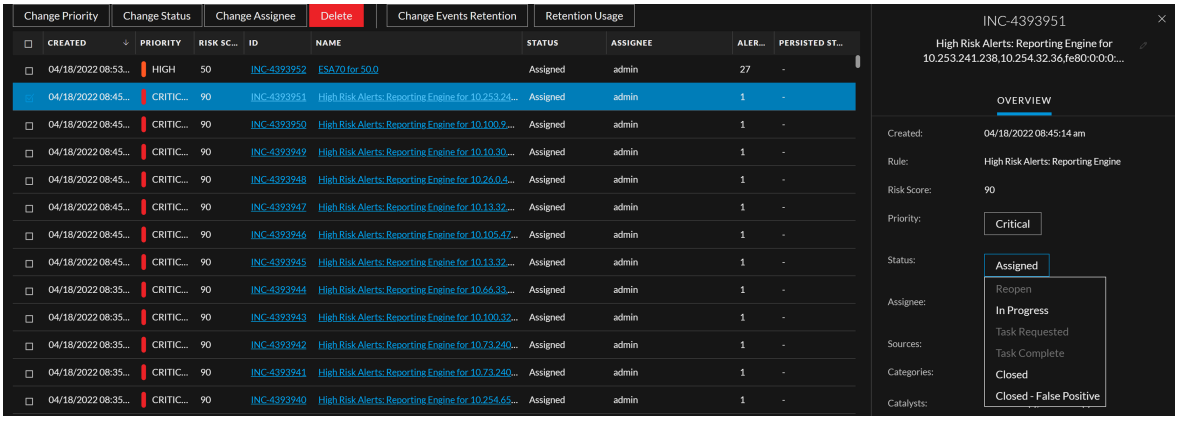


Note: If you select any incident and click **Change Status**, the current status of the incident is grayed out in the drop-down list. This is not applicable if you select multiple incidents. Refer the following figure.

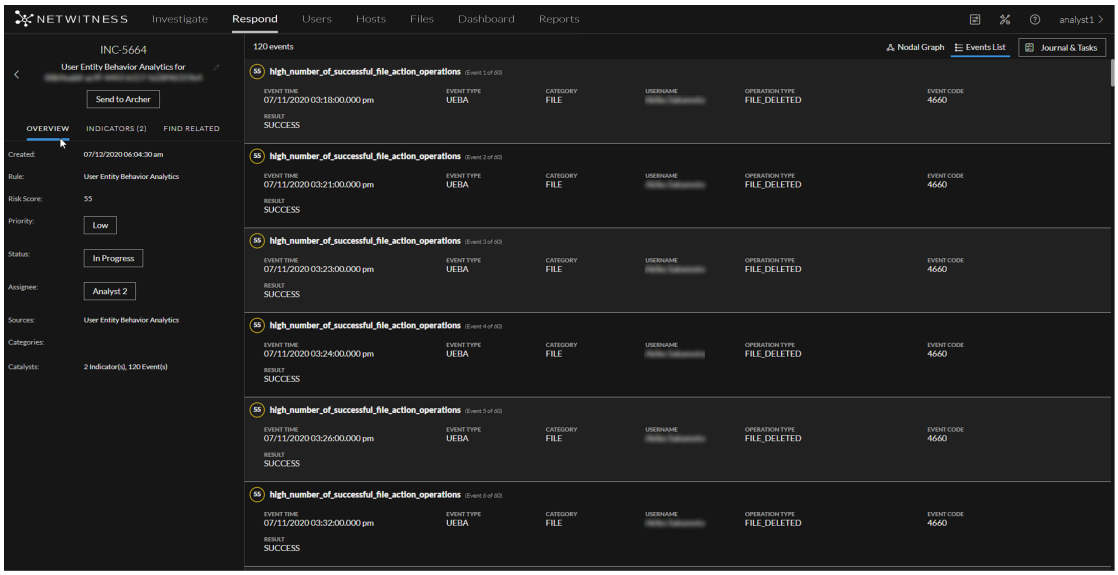


To change the status of a single incident from the Overview panel:

1. To open the Overview panel, do one of the following:
   - From the Incidents List view, click the row of an incident that needs a status update.

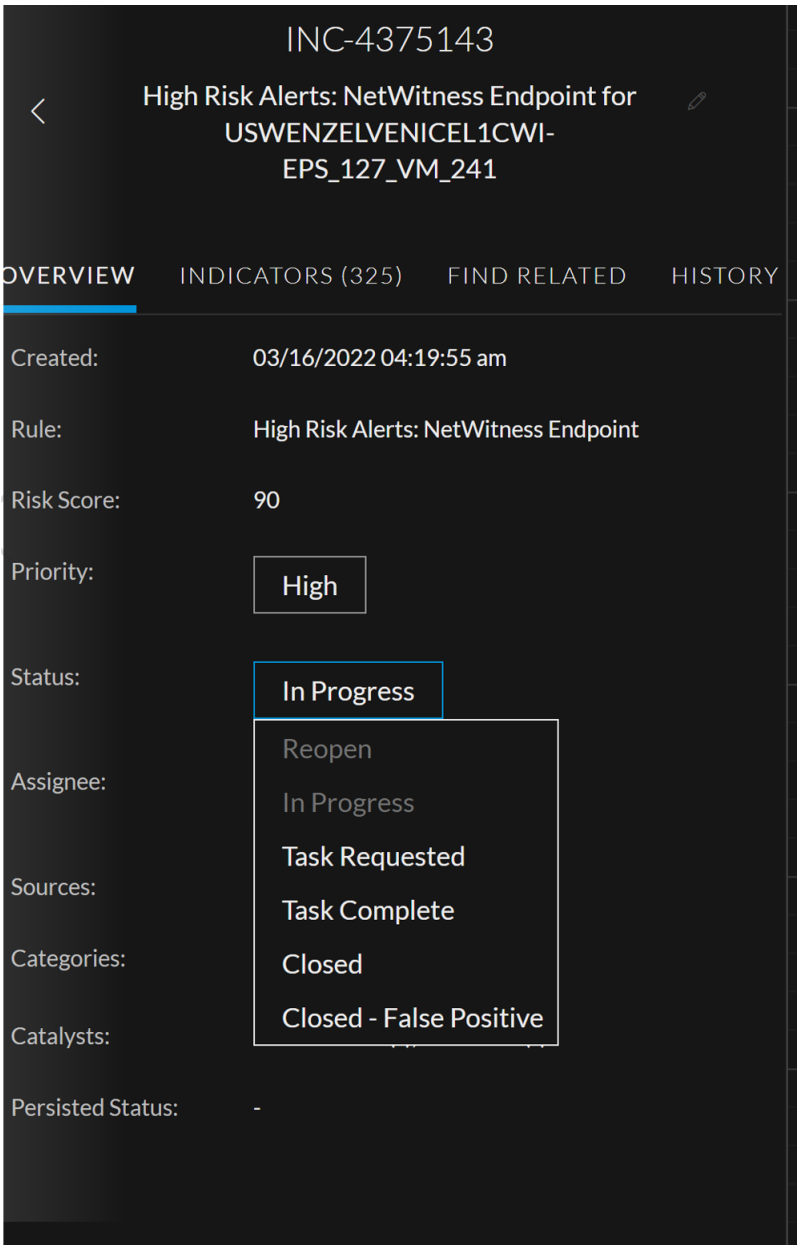


- From the Incident Details view, click the **OVERVIEW** tab.

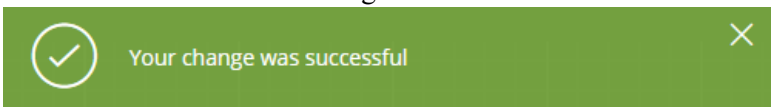


In the Overview panel, the Status button shows the current status of the incident.

2. Click the **Status** button and select a status from the drop-down list.



You can see a successful change notification.



Note: The incident status can be changed to **Reopen** only if the current status of the incident is **Closed** or **Closed - False Positive**.

Change Events Retention

Events retention enables you to persist events that are associated with particular incidents, thereby enabling you to view the incident related events in the future, regardless of its age. The event data will always be available for viewing and reconstruction as long as the event is persisted, enabling you to easily refer back to details, even if the original event has rolled over from the NetWitness database. You can perform the following functions:

- Persist all events
- Suspend persisting all events

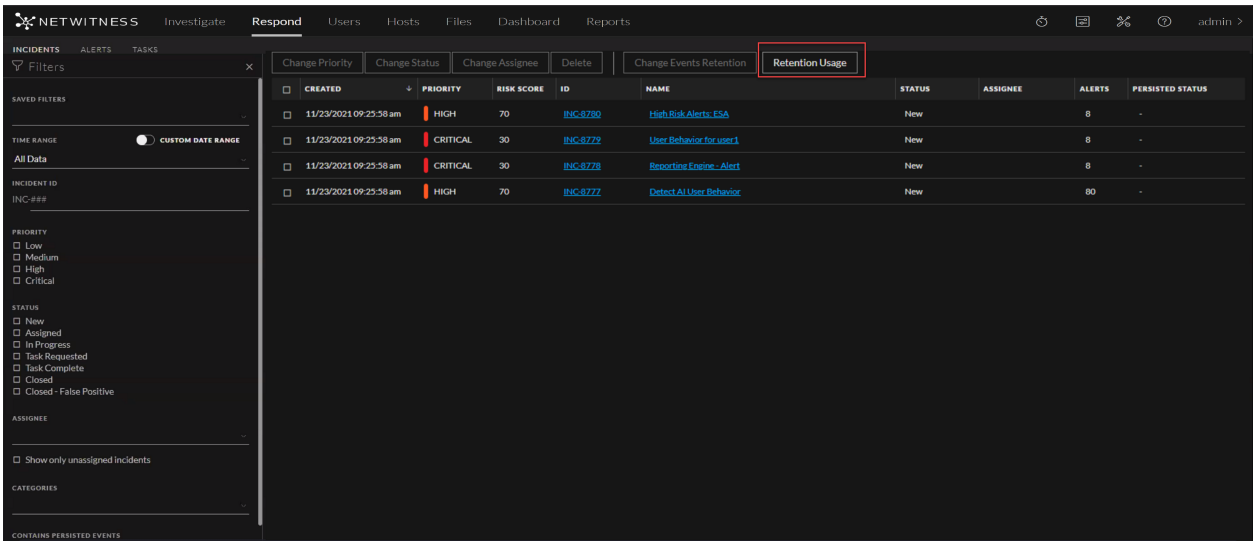
To change event retention:

1. Select the incidents for which you want to change the event retention plan.
2. Select **Persist all events** from the **Change Events Retention** tab to persist all the events that are associated with the selected incidents.
   a. The confirmation message appears. Click **OK** to persist all events.
   Persisting all events in an incident in NetWitness will save the events data in the long term cache of the source.
3. Select **Suspend Persisting all events** from the **Change Events Retention** tab to stop persisting the events that are associated with the selected incidents.
   a. The confirmation message appears. Click OK to suspend persist all events.
   Suspending persist of events in an incident from NetWitness will delete it from the long term cache of the source only. This may not be reversible if the original event data has rolled out in the source database.

Note: You cannot change the event retention for incidents that are in **New** or **Closed** state.

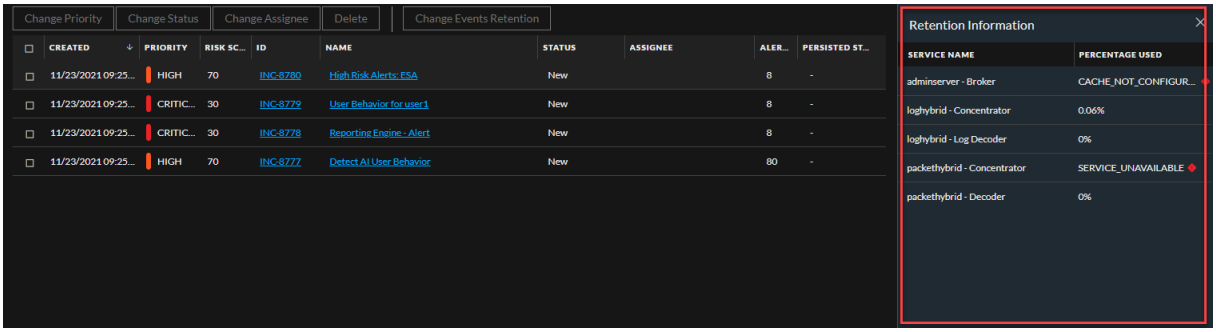
Obtain Retention Usage Details

The Retention Usage tab allows an analyst to fetch all the persisted data disc usage stats of all the configured services and the percentage used by the pinned cache directories. This will enable the analyst to determine if the disk is running out of space and if additional space needs to be added or suspend persist on the existing events in an incident.



After the analyst clicks on the Retention Usage tab, a Retention Information panel is displayed with the following status:

- Percentage of the disk used when data is persisted
- Cache directory is configured or not. In case it is not configured it is explicitly indicated.
- List of all the status of a configured service. In case the service is not available it is explicitly indicated.



Note: In case the disk space exceeds the usage, a warning message displayed. The service threshold can be configured by navigating to **Respond** > **Services** > **respond/core/properties** > **warning-threshold field**.

Change Incident Priority

The incident list is sorted by Priority by default. You can update the priority as you study the details of the case. The following priorities are available:

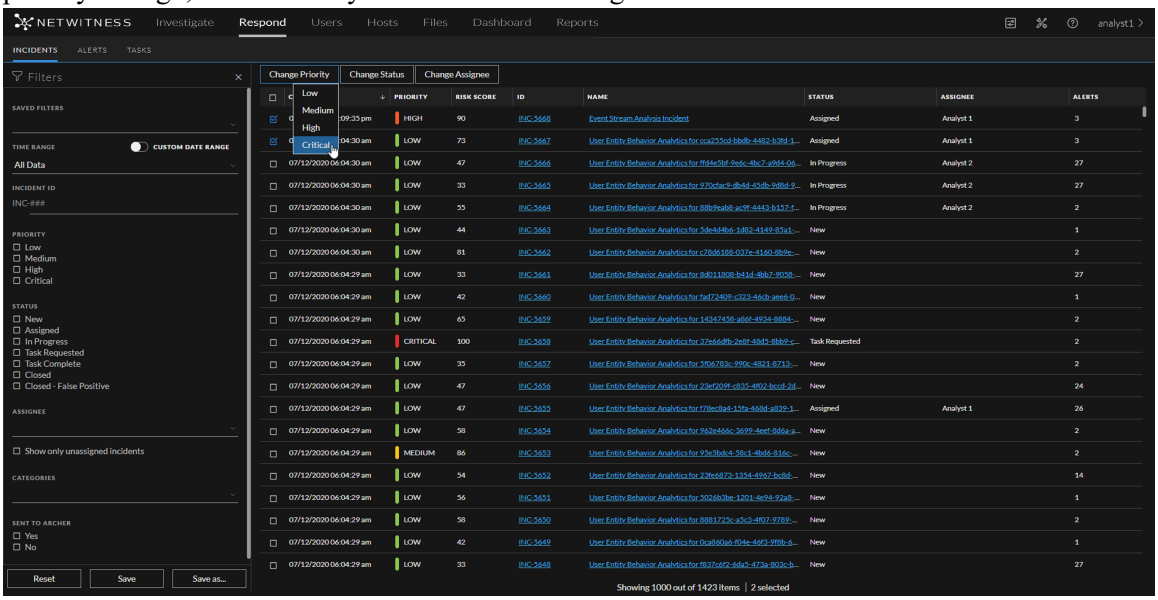
- Critical
- High

- Medium
- Low

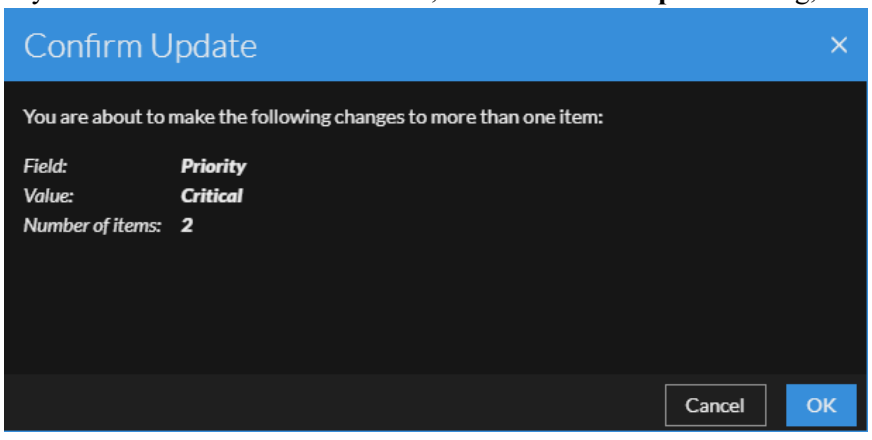
Note: You cannot change the priority of a closed incident.

To update the priority of multiple incidents:

1. In the Incidents List view, select one or more incidents that you would like to change. To select all of the incidents on the page, select the box in the incidents list header row. The number of incidents selected appears in the incidents list footer.
2. Click **Change Priority** and select a priority from the drop-down list. In this example, the current priority is High, but the Analyst would like to change it to Critical for the selected incidents.

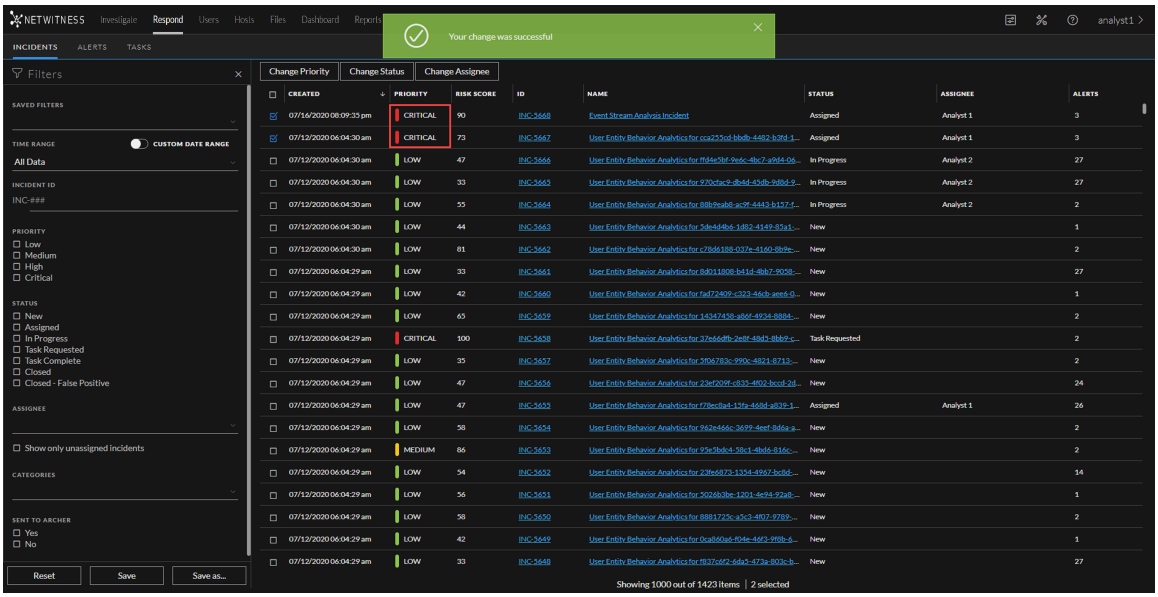


3. If you select more than one incident, in the **Confirm Update** dialog, click **OK**.



You can see a successful change notification. In this example, the status of the updated incidents now

show Critical.

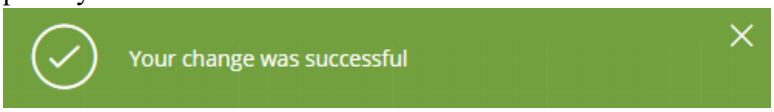


To change the priority of a single incident from the Overview panel

1. To open the Overview panel, do one of the following:
   - From the Incidents List view, click the row of an incident that needs a priority update.
   - From the Incident Details view, click the **Overview** tab in the left panel.
     In the Overview panel, the Priority button shows the current priority of the incident.
2. Click the **Priority** button and select a status from the drop-down list.



You can see a successful change notification. The Priority button changes to show the new incident priority.



Note: Current priority is grayed out under **Priority** drop-down list. You will not be able to select the grayed out priority.

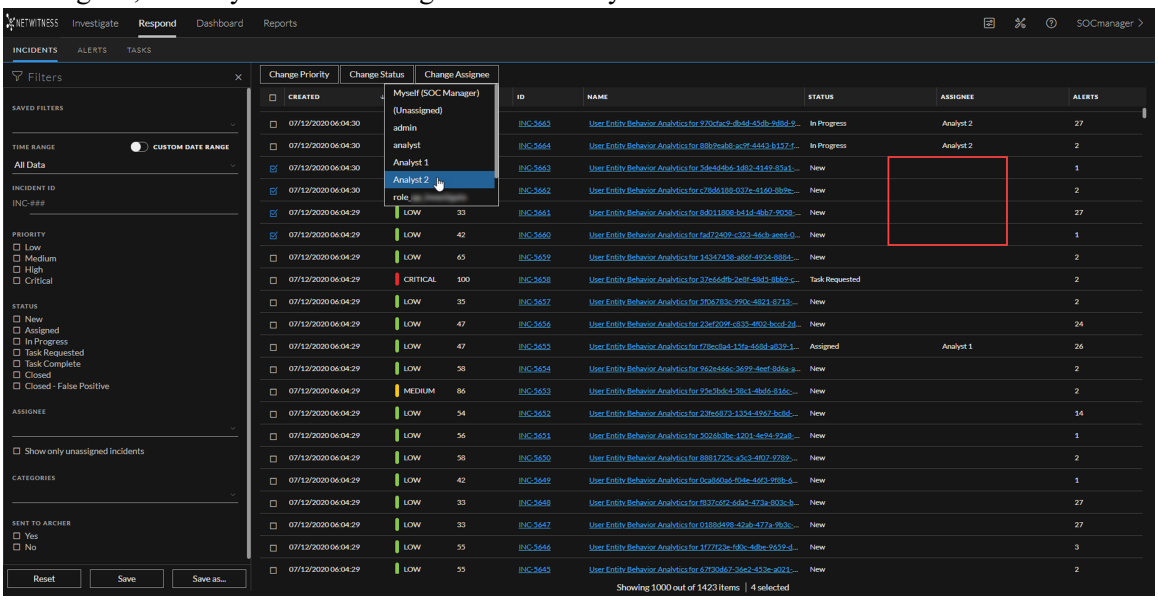
Assign Incidents to Other Analysts

You can assign incidents to other Analysts in the same way as you assign incidents to yourself. SOC Managers and Administrators can assign multiple incidents to a user at the same time.

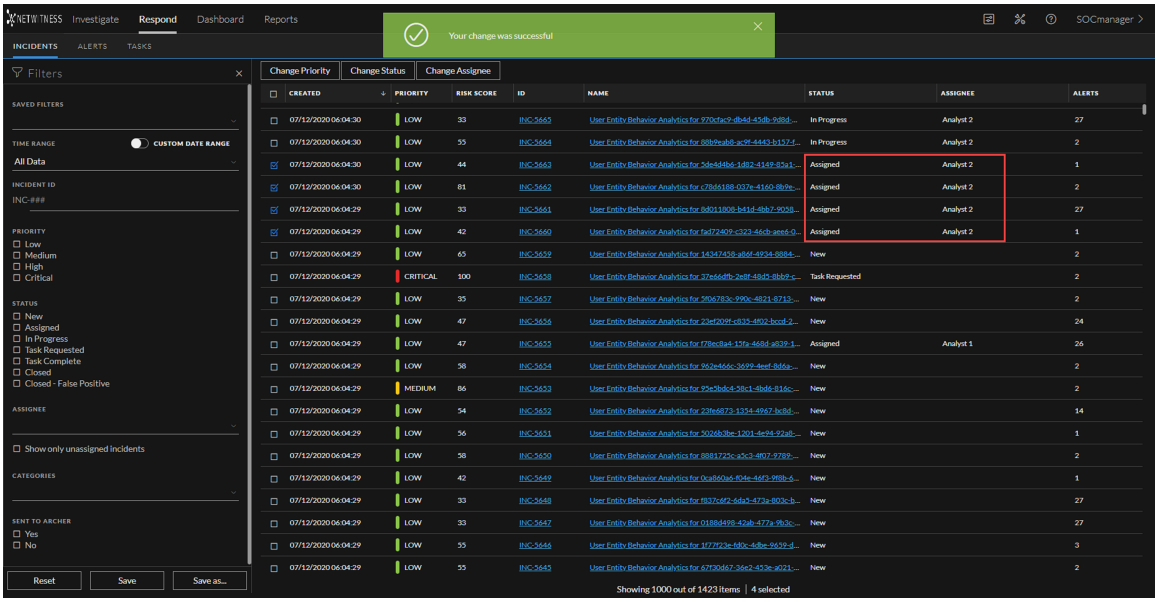
Note: You cannot change the assignee of a closed incident.

To assign multiple incidents to a user:

1. In the Incidents List view, select the incidents that you would like to assign to a user. To select all of the incidents on the page, select the box in the incidents list header row. The number of incidents selected appears in the incidents list footer.
2. Click **Change Assignee** and select a user from the drop-down list. In this example, the incidents are unassigned, but they should be assigned to an Analyst.



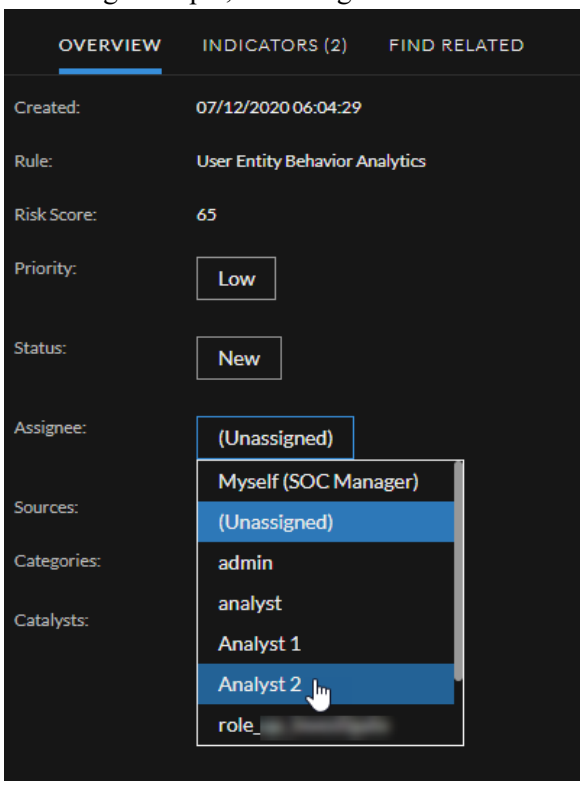
3. If you select more than one incident, in the **Confirm Update** dialog, click **OK**.
   You can see a successful change notification. The assignee changes to the selected user.



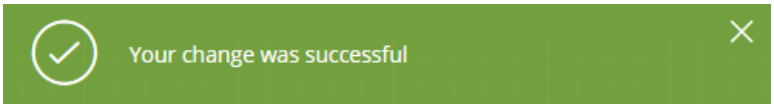
To assign a user to an incident from the Overview panel:

1. To open the Overview panel, do one of the following:
   - From the Incidents List view, click the row of an incident that you would like to assign to a user.
   - From the Incident Details view, click the **Overview** tab in the left panel.
     In the Overview panel, the Assignee button shows the current assignee of the incident. In the

following example, the Assignee button has a current status of Unassigned.



2. Click the **Assignee** button and select a user from the drop-down list.
   You can see a successful change notification. The Assignee button changes to show the assigned user.



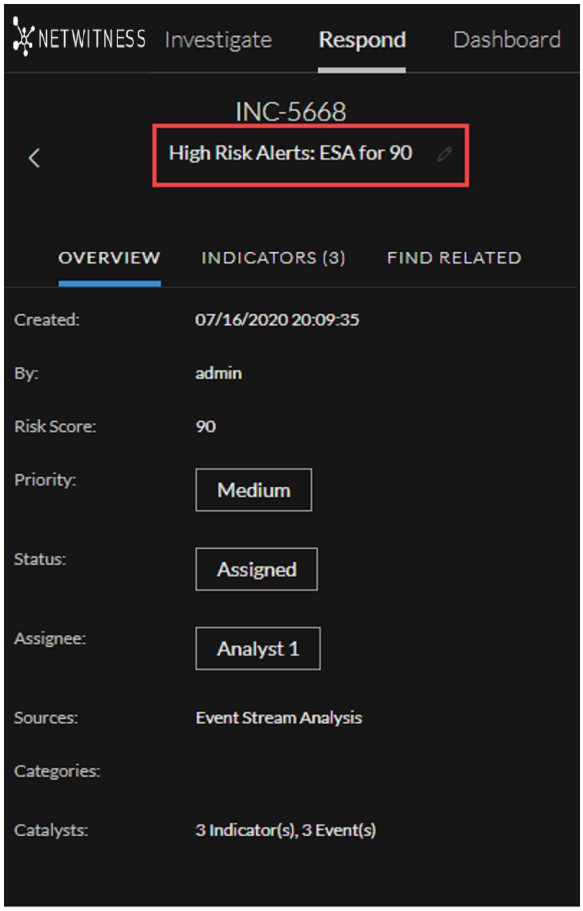
Note: Current assignee name is grayed out under **Assignee** drop-down list. You will not be able to select the grayed out user.

Rename an Incident

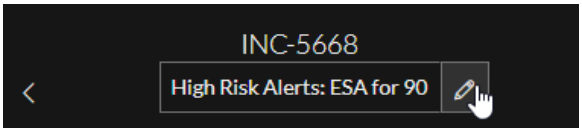
You can rename an incident from the Overview panel in the Incidents List view and the Incident Details view. For example, you may want to rename an incident to provide clarification about the issue, especially if multiple incidents have the same name.

1. Go to **Respond > Incidents**.
2. To open the Overview panel, do one of the following:
   - From the Incidents List view, click the row of an incident that needs a name change.
     The Overview panel opens.

- From the Incident Details view, click the **OVERVIEW** tab in the left panel.
  In the header above the Overview panel, you can see the incident ID and the incident name.



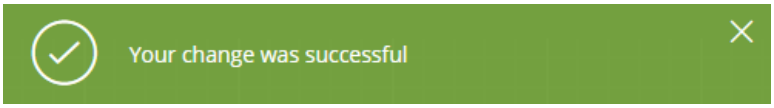
3. Click the incident name in the header to open a text editor.



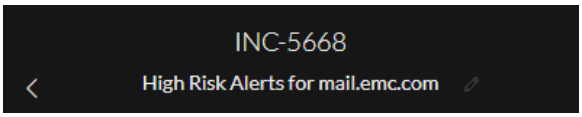
4. Type a new name for the incident in the text editor and click the check mark to confirm the change.



For example, you can change "High Risk Alerts: ESA for 90.0" to "High Risk Alerts for mail.emc.com" for more clarification.
You can see a successful change notification.



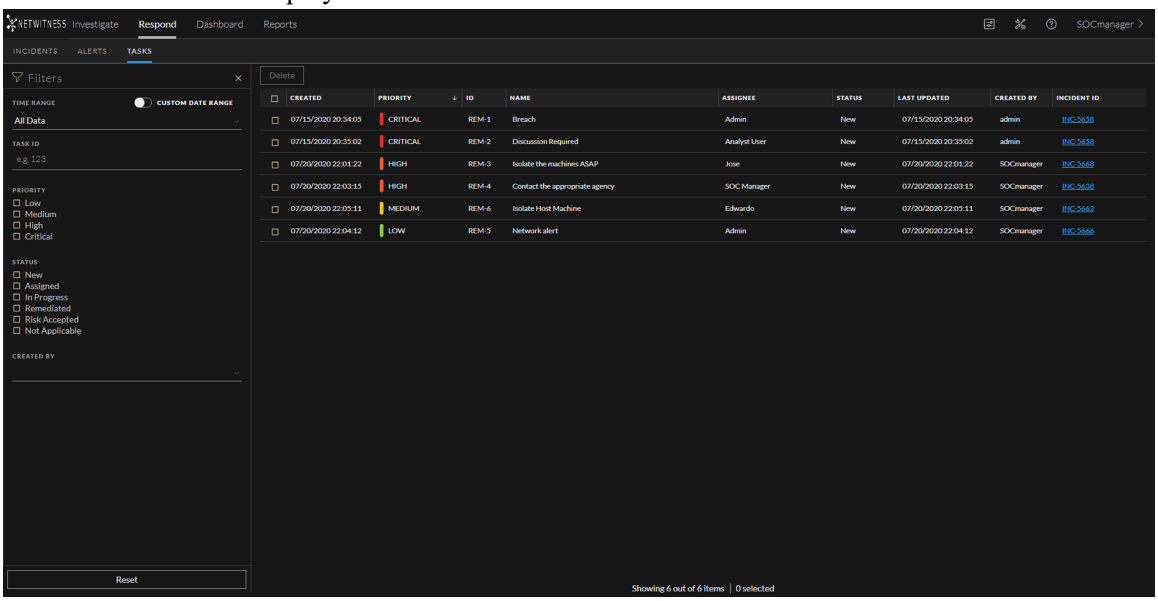
The incident name field shows the new name.



View All Incident Tasks

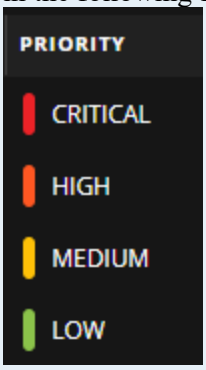
When additional work is required for an incident, you can create tasks for the incident and track the progress on those tasks. This is helpful, for example, when the work being done is outside security operations or you make a request for a computer reimage. In the Tasks List view, you can manage and track the tasks to closure.

1. Go to **Respond > Tasks**.
   The Tasks List view displays a list of all incident tasks.



2. Scroll through the tasks list, which shows basic information about each task as described in the following table.

Column	Description
Created	Displays the date when the task was created.

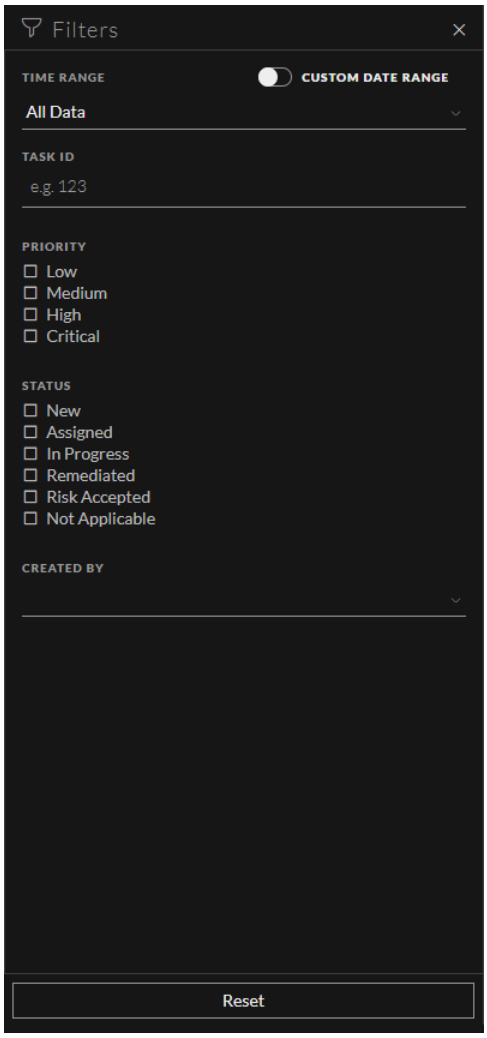
Column	Description
Priority	Displays the priority assigned to the task. The priority can be any of the following: Critical, High, Medium, or Low. The Priority is also color coded, where red indicates **Critical**, orange represents **High** risk, yellow indicates **Medium** risk, and green represents **Low** risk as shown in the following figure: 
ID	Displays the task ID.
Name	Displays the task name.
Assignee	Displays the name of the user assigned to the task.
Status	Displays the status of the task: New, Assigned, In Progress, Remediated, Risk Accepted, and Not Applicable.
Last Updated	Displays the date and time when the task was last updated.
Created By	Displays the user who created the task.
Incident ID	Displays the incident ID for which the task was created. Click the ID to display the details of the incident.

At the bottom of the list, you can see the number of tasks on the current page, the total number of tasks, and the number of tasks selected. For example: **Showing 6 out of 6 items | 2 selected**.

Filter the Tasks List

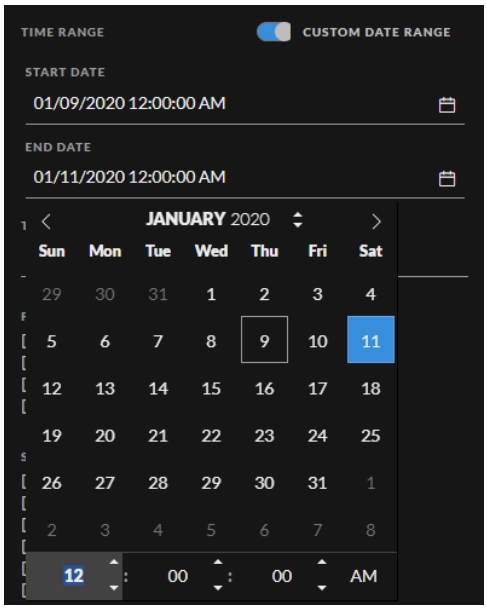
The number of tasks in the Tasks List can be very large, making it difficult to locate particular tasks. The Filter enables you to specify those tasks that you would like to view, such as tasks created within the last 7 days. You can also search for a specific task.

1. Go to **Respond > Tasks**.
   The Filters panel appears to the left of the Tasks list. If you do not see the Filters panel, in the Tasks List view toolbar, click , which opens the Filters panel.



2. In the Filters panel, select one or more options to filter the incidents list:
   - **Time Range**: You can select a specific time period from the Time Range drop-down list. The time range is based on the creation date of the tasks. For example, if you select Last Hour, you can see tasks that were created within the last 60 minutes.
   - **Custom Date Range**: You can specify a specific date range instead of selecting a Time Range option. To do this, click the white circle in front of CUSTOM DATE RANGE to view the Start

Date and End Date fields. Select the dates and times from the calendar.



- **Task ID**: Type the Task ID for a task that you would like to locate, for example REM-123.
- **Priority**: Select the priorities that you would like to view.
- **Status**: Select one or more incident statuses. For example, select Remediated to view completed remediation tasks.
- **Created By**: Select the user who created the tasks that you would like to view. For example, if you only want to view the tasks created by Edwardo, select Edwardo from the CREATED BY drop-down list. If you want to view tasks regardless of the person who created the task, do not make a selection under CREATED BY.

The Tasks List shows a list of tasks that meet your selection criteria. You can see the number of items in your filtered list at the bottom of the tasks list.
For example: **Showing 6 out of 6 items**

3. If you want to close the Filters panel, click **X**. Your filters remain in place until you remove them.

Remove My Filters from the Tasks List

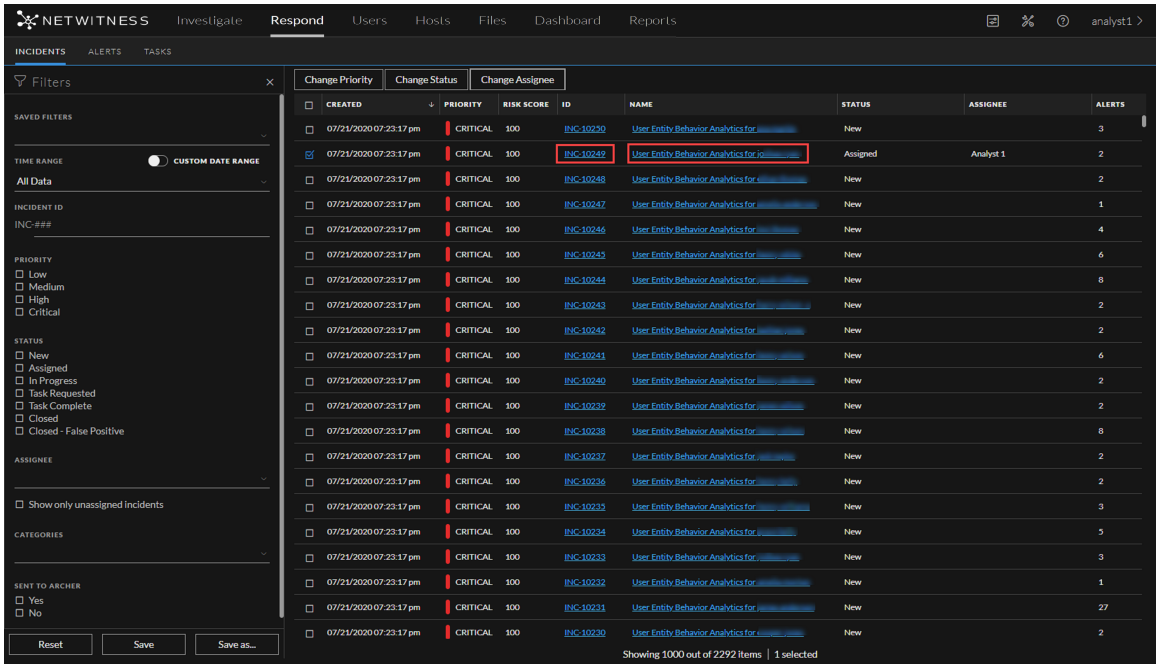
NetWitness remembers your filter selections in the Tasks List view. You can remove your filter selections when you no longer need them. For example, if you are not seeing the number of tasks that you expect to see or you want to view all of the tasks in your tasks list, you can reset your filters.

1. Go to **Respond > Tasks**.
   The Filters panel appears to the left of the tasks list. If you do not see the Filters panel, in the Tasks List view toolbar, click , which opens the Filters panel.
2. At the bottom of the Filters panel, click **Reset Filters**.

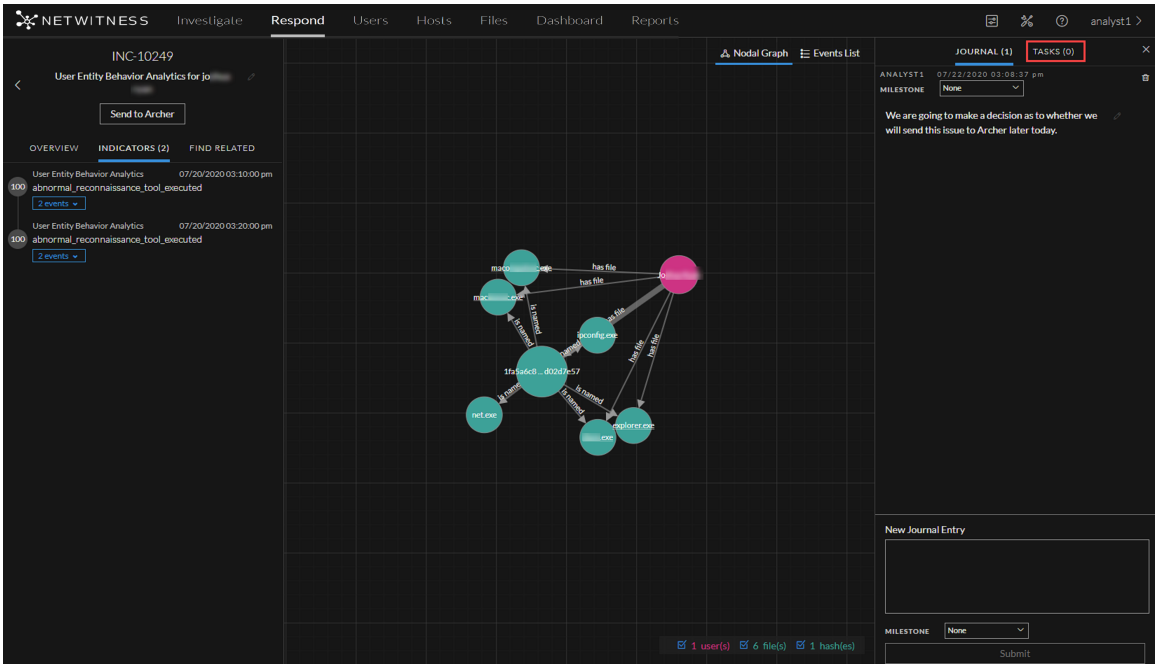
Create a Task

After you investigate an incident and know more about it, you can create a task, assign it to a user, and track it to closure. You create tasks from the Incident Details view.

1. Go to **Respond > Incidents**.
   The Incidents List view displays a list of all of the incidents.
2. Locate the incident that needs a task and click the link in the **ID** or **Name** field.



3. In the Journal panel on the right side of the Incident Details view, click the **Tasks** tab.
   If you do not see the Journal panel, click **Journal & Tasks** and then click the **Tasks** tab.



4. In the Tasks panel, click **Add New Task**.

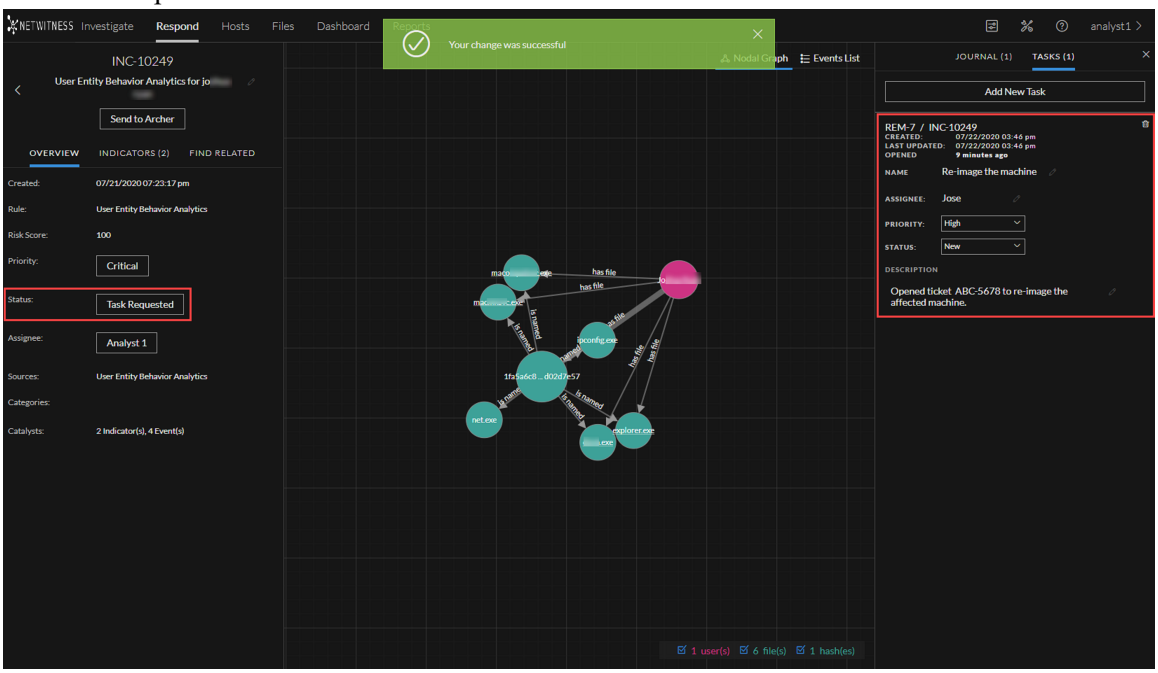


You can see the new task fields.

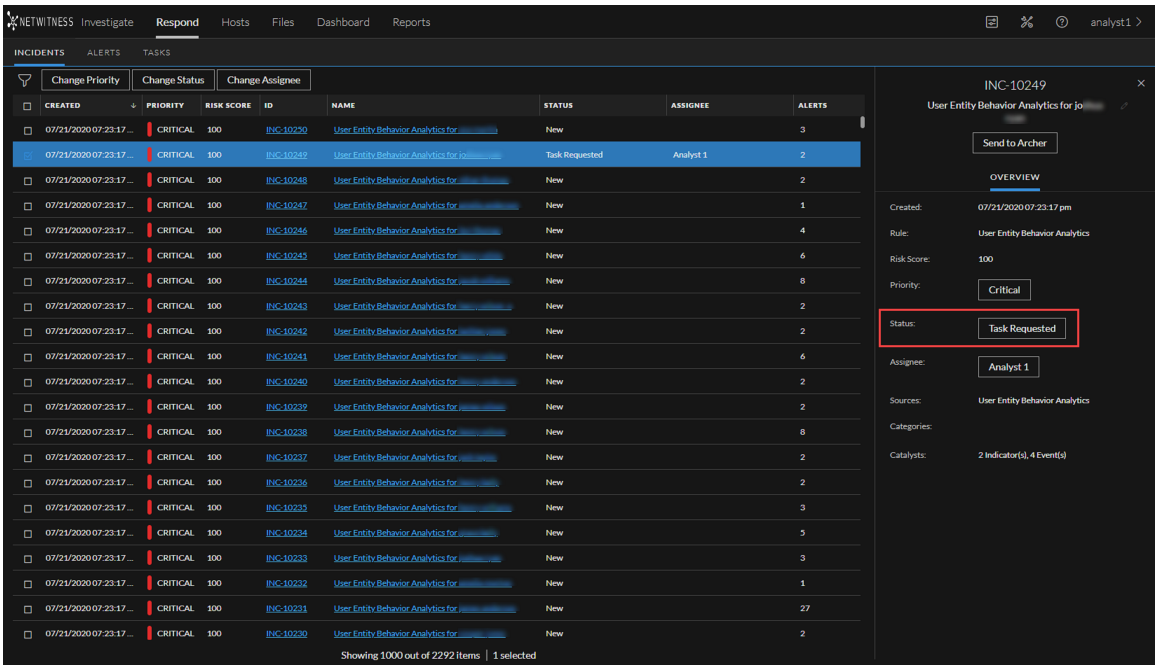


If the incident is in a closed state (Closed or Closed - False Positive), the Add New Task button is disabled.

5. Provide the following information:
   - **Name** - Name of the task. For example: Re-image the machine.
   - **Description** - (Optional) Type information that describes the task. You may want to include any applicable reference numbers.
   - **Assignee** - (Optional) Type the username of the user to whom the task is to be assigned.
   - **Priority** - Click the priority button and select a priority for the tasks from the drop-down list: Low, Medium, High, or Critical.
6. Click **Save**.
   You can see a confirmation that your change was successful. The incident status changes to **Task Requested**. (You may need to refresh the Incident Details view to see the changes.) The task appears in the Tasks panel for this incident.



In the Incidents List view, the incident status also changes to Task Requested.



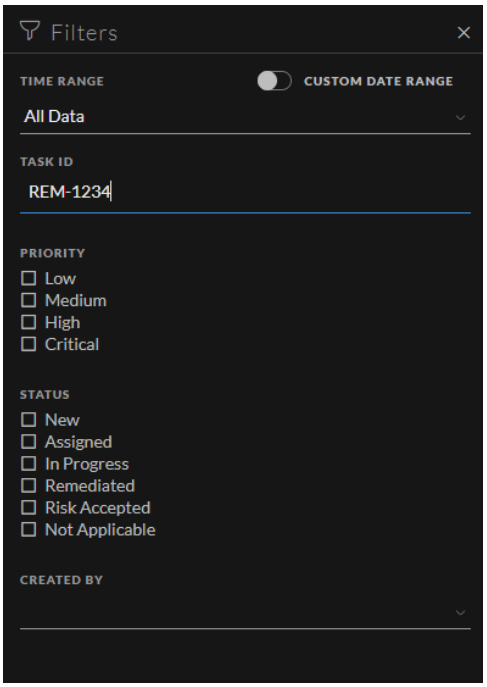
The task also appears in the Tasks list (Respond > Tasks), which shows a list of all incident tasks.

Note: If you do not see the status change, you may need to refresh your internet browser.

Find a Task

If you know the Task ID, you can quickly locate a task using the Filter. For example, you may want to locate a specific task out of thousands of tasks.

1. Go to **Respond > Tasks**.
   The Filters panel appears to the left of the Tasks list. If you do not see the Filters panel, in the Tasks List view toolbar, click , which opens the Filters panel.



2. In the **Task ID** field, type the Task ID for a task that you would like to locate, for example REM-1234.
   The specified task appears in your task list. If you do not see any results, try resetting your filters.

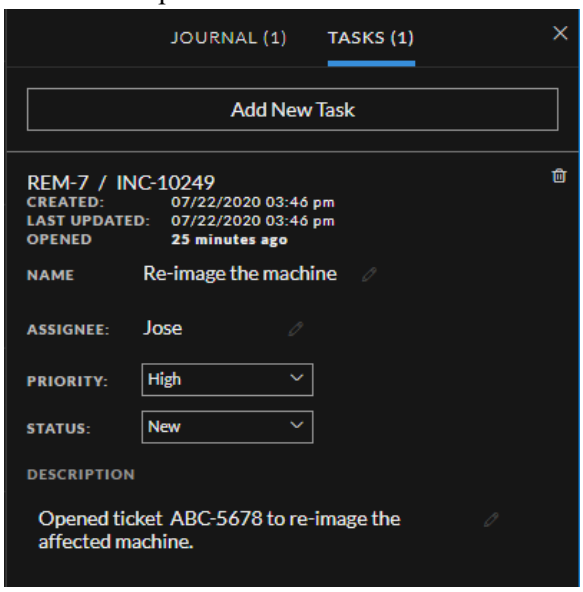
Modify a Task

You can modify a task from within an incident and from the Tasks list. For example, you may want to show the status of the task as In Progress and add some additional information to the task. If the task is in a closed state (Not Applicable, Risk Accepted, or Remediated), you cannot modify the Priority or Assignee.

To modify a Task from within an incident:

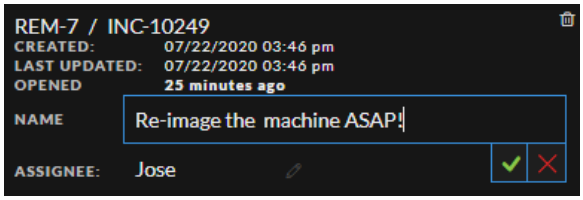
1. Go to **Respond > Incidents**.
   The Incidents List view displays a list of all incidents.
2. Locate the incident that needs a task update and click the link in the **ID** or **Name** field.
3. In the Journal panel on the right side of the Incident Details view, click the **Tasks** tab.
   If you do not see the Journal panel, click **Journal & Tasks** and then click the **Tasks** tab.
   In the Tasks panel, a pencil icon indicates a text field that you can change. A button indicates that

there is a drop-down list to make a selection.



4. You can modify any of the following fields:

- **Name** - Click the current task name to open a text editor.

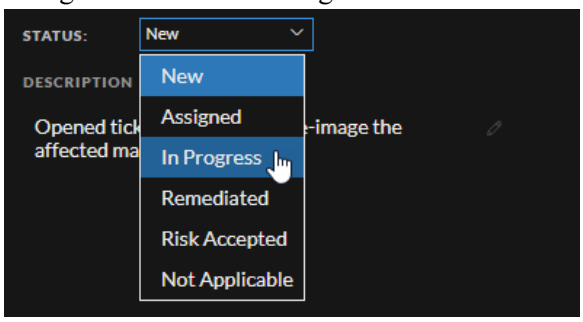


  Click the check mark to confirm the change. For example, you can change "Re-image the machine" to "Re-image the machine ASAP!"

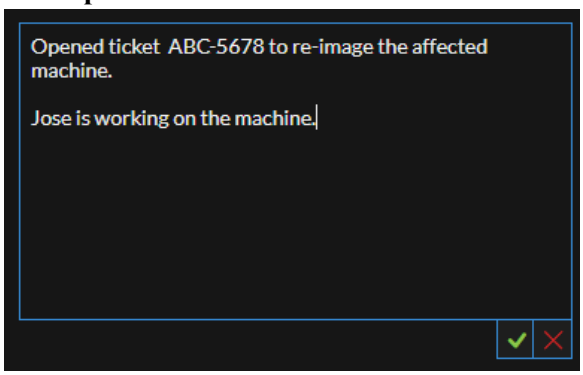
- **Assignee** - Click (Unassigned) or the name of the previous assignee to open a text editor. Type the username of the user to whom the task is to be assigned.
  Click the check mark to confirm the change.

- **Priority** - Click the Priority button and select a priority for the task from the drop-down list: Low, Medium, High, or Critical.

- **Status** - Click the Status button and select a status for the task from the drop-down list: New, Assigned, In Progress, Remediated, Risk Accepted, and Not Applicable. For example, you can change the status to In Progress.



- **Description** - Click the text underneath the description to open a text editor.

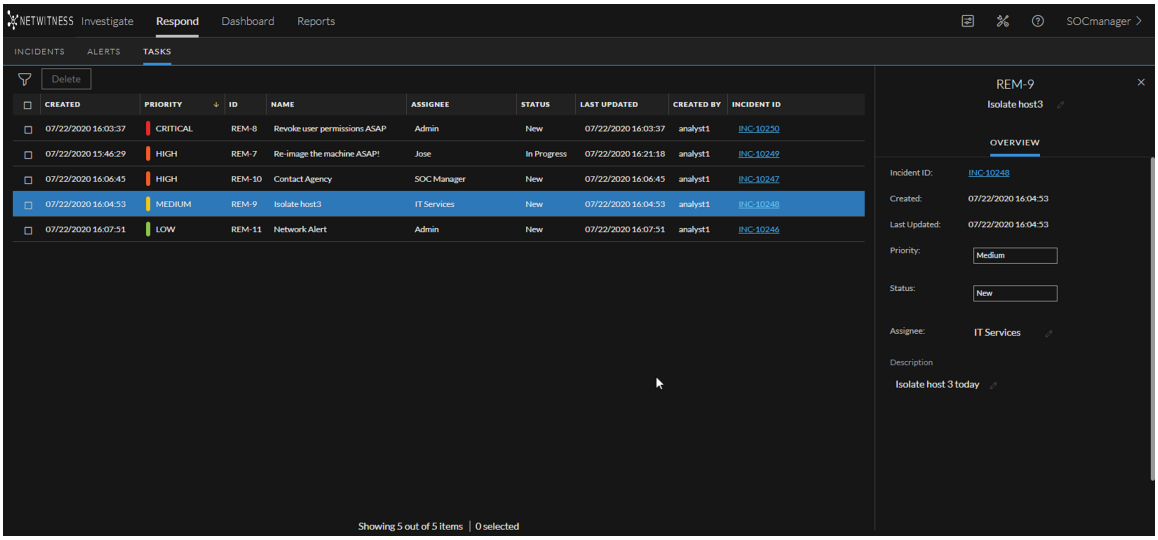


Modify the text and click the check mark to confirm the change.

For each change that you make, you can see a confirmation that your change was successful.

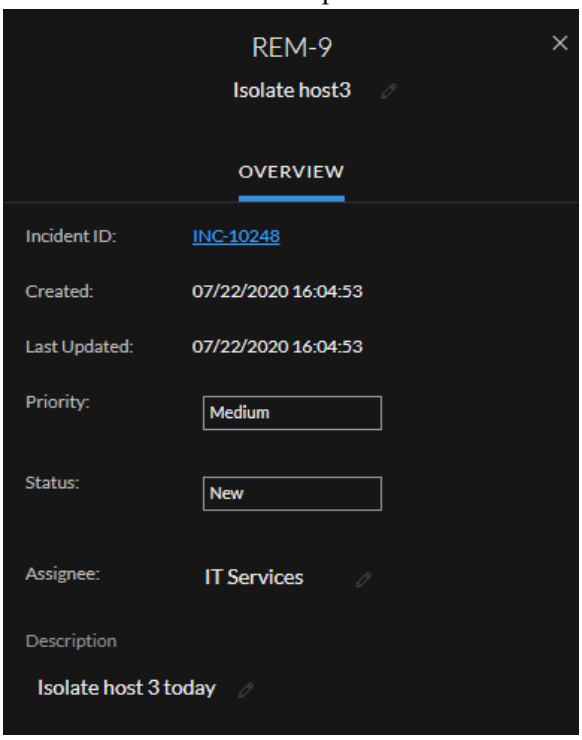
To modify a Task from the Tasks list:

1. Go to **Respond > Tasks**.
   The Tasks List view displays a list of all incident tasks.
2. In the Tasks list, click the task that you want to update.
   The Task Overview panel appears to the right of the tasks list.



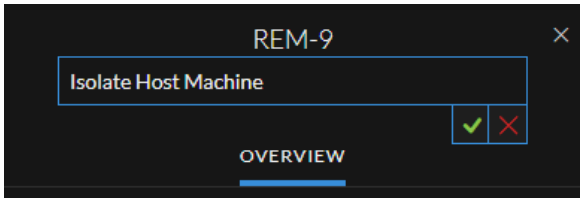
In the Task Overview panel, a pencil icon indicates a text field that you can change. A button

indicates that there is a drop-down list to make a selection.



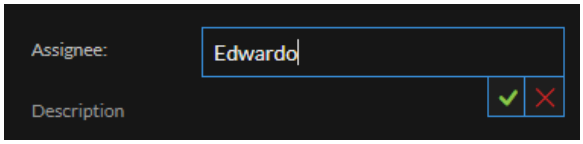
3. You can modify any of the following fields:

- **<Task Name>** - At the top of the Task Overview panel, below the Task ID, click the current task name to open a text editor.



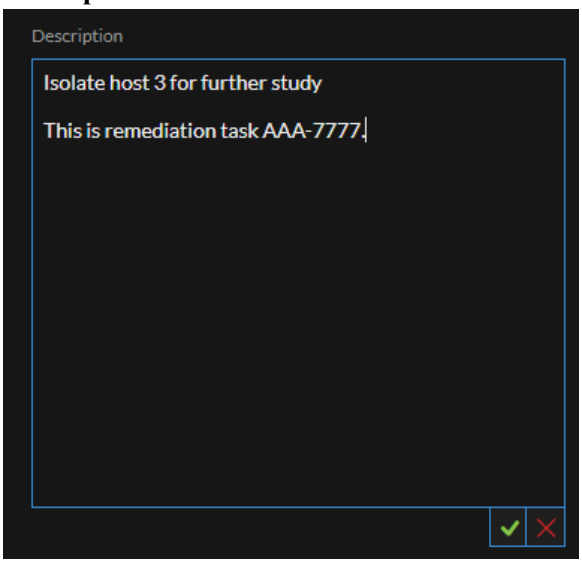
Click the check mark to confirm the change. For example, you can change Isolate Host to Isolate Host Machine.

- **Priority** - Click the Priority button and select a priority for the task from the drop-down list: Low, Medium, High, or Critical.
- **Status** - Click the Status button and select a status for the task from the drop-down list: New, Assigned, In Progress, Remediated, Risk Accepted, and Not Applicable.
- **Assignee** - Click (Unassigned) or the name of the previous assignee to open a text editor. Type the username of the user to whom the task is to be assigned.



Click the check mark to confirm the change.

- **Description** - Click the text underneath the description to open a text editor.



  Modify the text and click the check mark to confirm the change.

For each change that you make, you can see a confirmation that your change was successful.

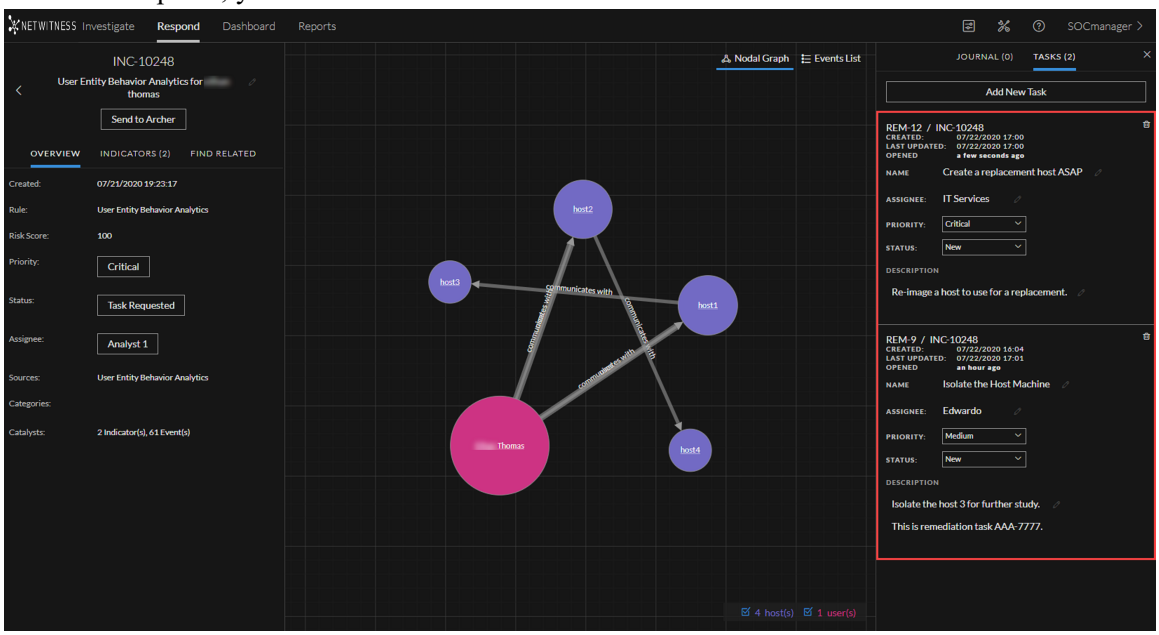
Delete a Task

You can delete a task, if, for example, you created it in error or you find that it is not needed. You can delete a task from within an incident and also from the Tasks List view. In the Tasks List view, you can delete multiple tasks at the same time.

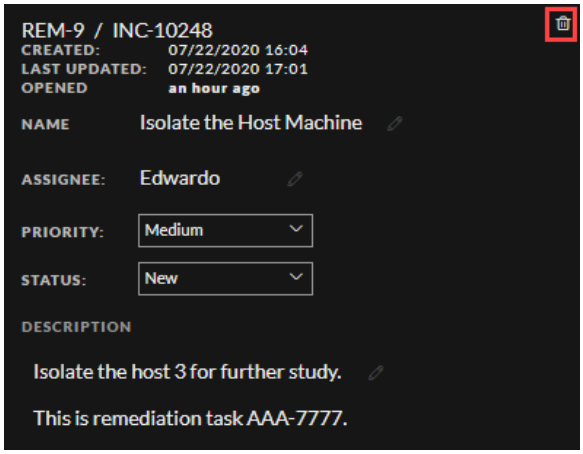
To Delete a Task from within an incident:

1. Go to **Respond > Incidents**.
   The Incidents List view displays a list of all incidents.
2. Locate the incident that needs a task update and click the link in the **ID** or **Name** field.
3. In the Journal panel on the right side of the Incident Details view, click the **Tasks** tab.
   If you do not see the Journal panel, click **Journal & Tasks** and then click the **Tasks** tab.

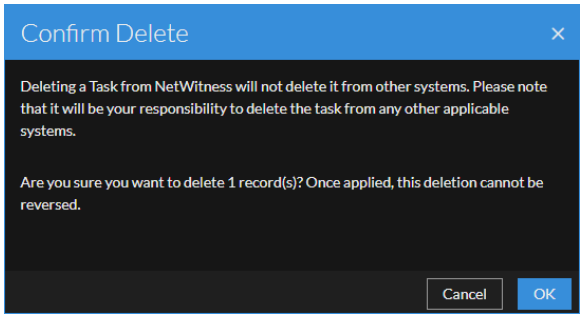
In the Tasks panel, you can see the tasks created for the incident.



4. Click  to the right of the task that you want to delete.



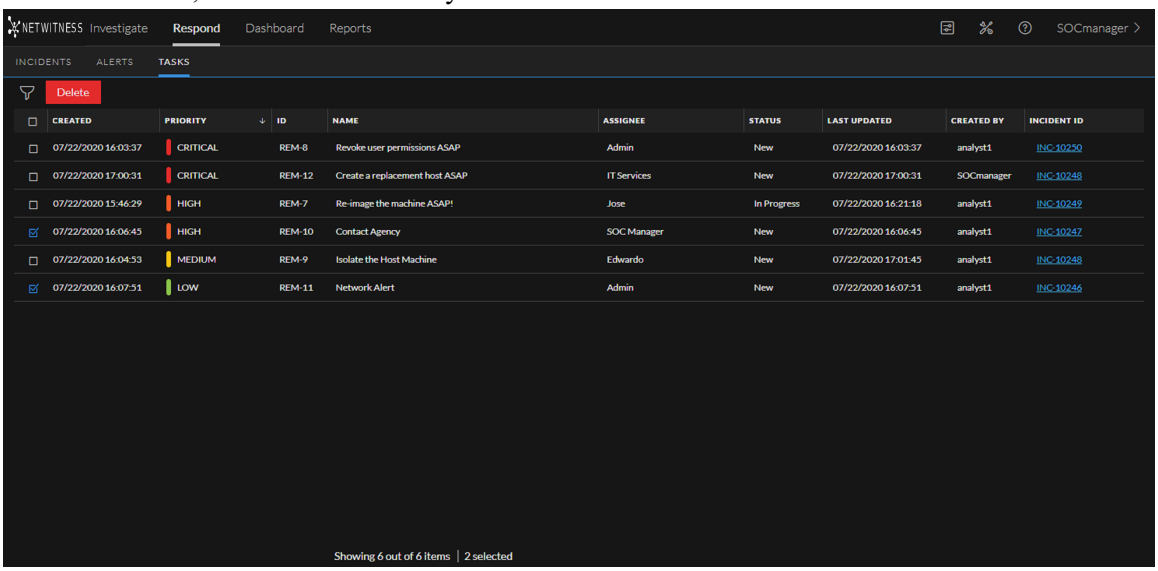
5. Confirm that you want to delete the task and click **OK**.



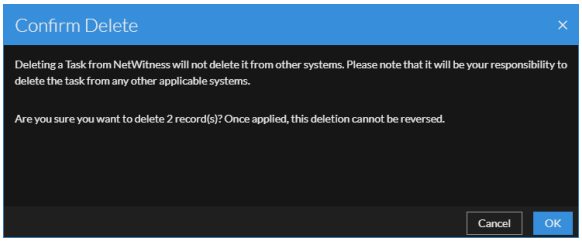
The task is deleted from NetWitness. Deleting tasks from NetWitness does not delete them from other systems.

To Delete Tasks from the Tasks List:

1. Go to **Respond > Tasks**.
   The Tasks List view displays a list of all incident tasks.
2. In the Tasks list, select the tasks that you want to delete and click **Delete**.



3. Confirm that you want to delete the tasks and click **OK**.



   The tasks are deleted from NetWitness. Deleting tasks from NetWitness does not delete them from other systems.

Close an Incident

When you have arrived at a solution after investigating an incident and remediating it, you close the incident.

1. Go to **Respond > Incidents**.
2. In the Incident List view, select the incident that you want to close and click **Change Status**.
3. Select **Closed** from the drop-down list.
   You can see a successful change notification. The incident is now closed. You cannot change the priority or assignee of a closed incident.

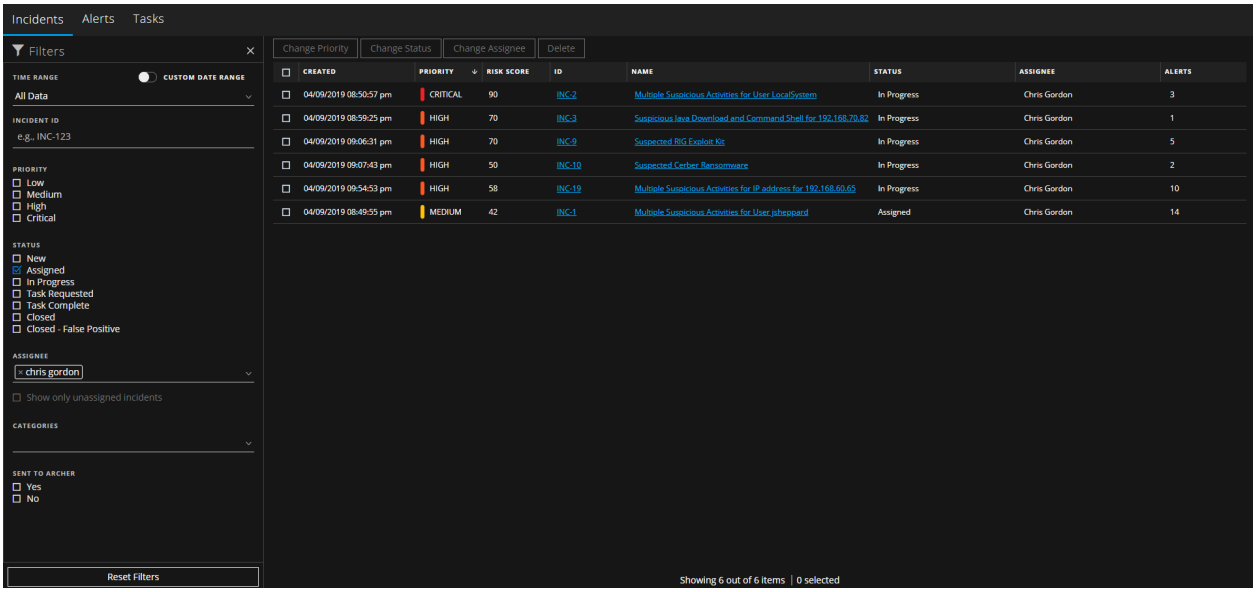
Note: You can also close an incident in the Overview panel. You can close multiple incidents at the same time in the Incident List view. Change Incident Status provides additional details.

Incident Response Use Case Examples

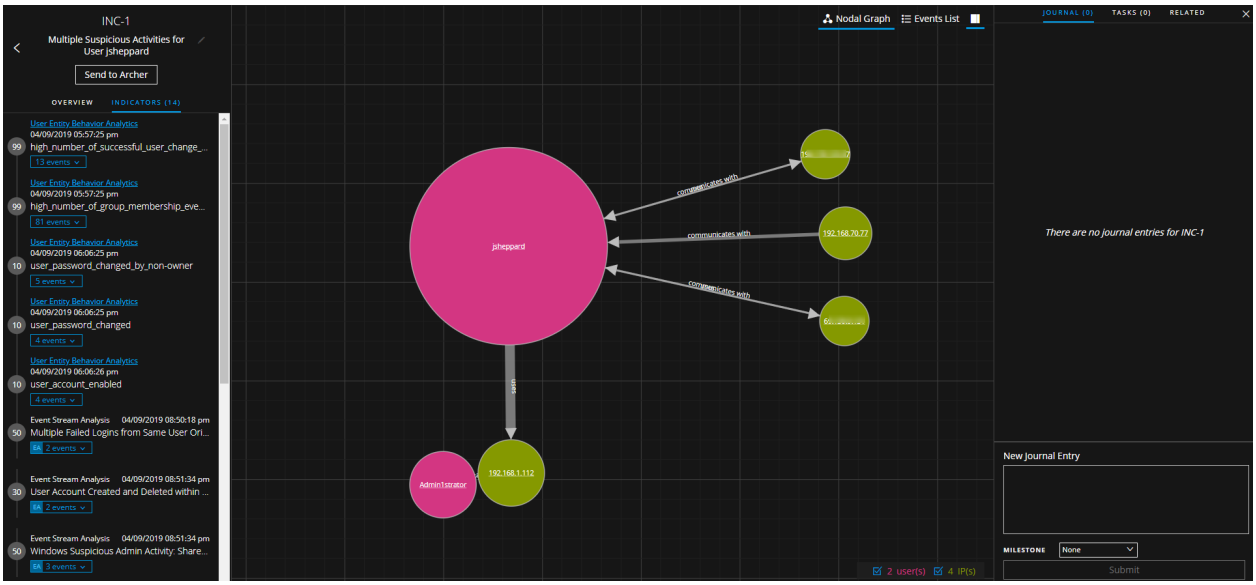
The following use cases provide examples of an analyst using NetWitness to quickly respond to incidents, identify threats, and take action to reduce or eliminate the ability of threat actors to compromise valuable information in the corporate network.

Use Case #1: UEBA Anomalous User Activity

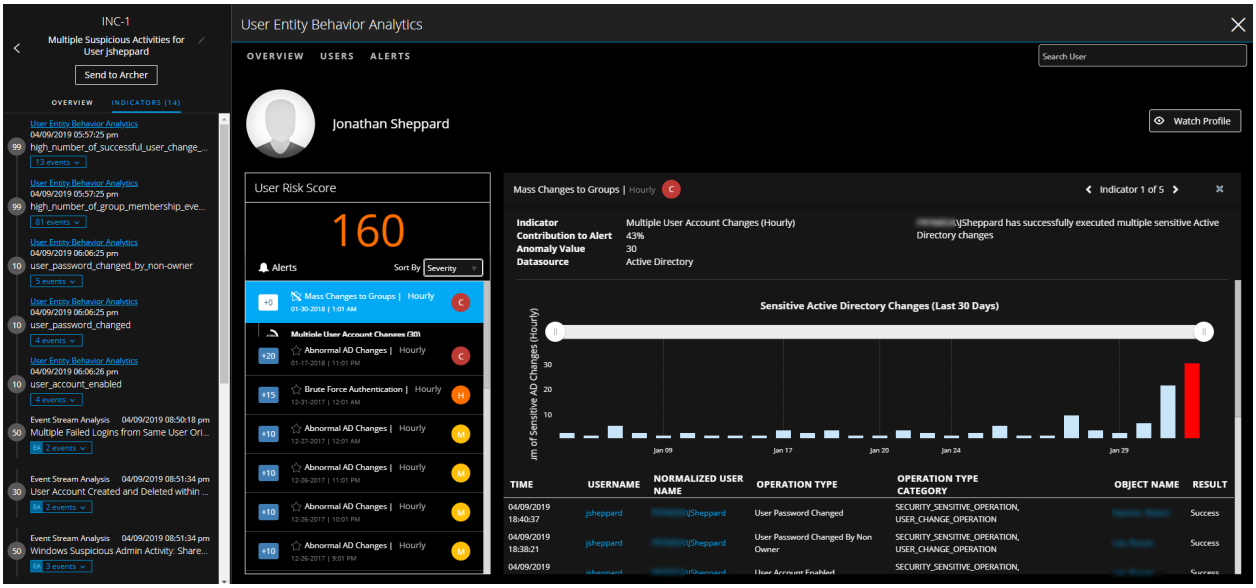
An analyst named Chris logs in, goes to the Incidents List view, and uses the filters on the left-hand side to look at all of the incidents assigned to her. In the list, she notices an incident that she has not yet reviewed (status is Assigned) and opens the incident.



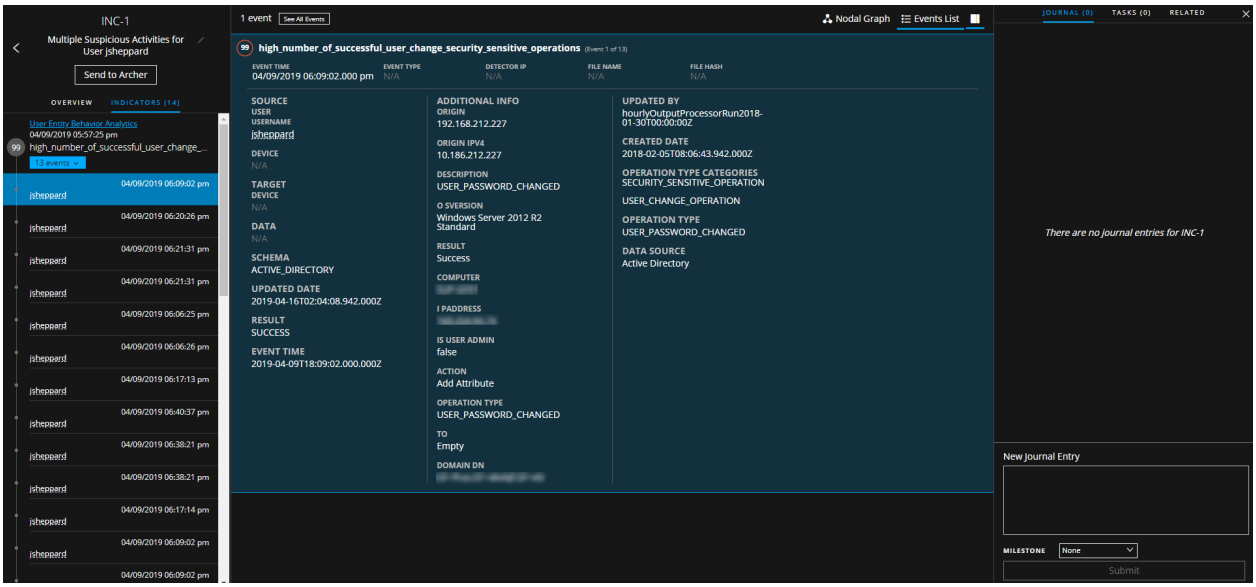
In the Incident Details view, the analyst sees a timeline of contributing events (Indicators panel) on the left, a visualization of the entities involved in the middle, and additional panels on the right where she can keep track of notes and tasks during her review.



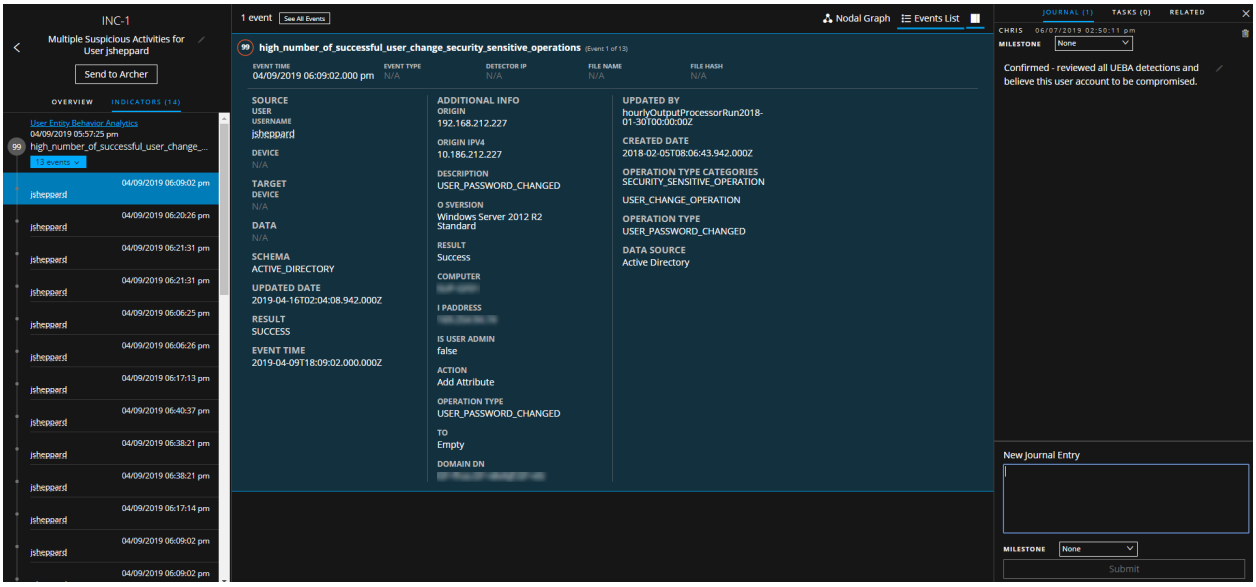
Chris uses the Indicators panel to get finer detail on the events that lead to this incident being created. Clicking the "User Entity Behavior Analytics" link in the Indicators panel exposes the analysis that was done and the anomaly that was detected on the user account "Jonathan Sheppard (jsheppard)." The User Entity Behavior Analytics panel below shows an overall risk score of 160, details the Windows events that contributed to the severity, and shows the actual anomaly in user account changes attributed to this user. She can explore the data in as much depth as required to help validate and understand the alert.



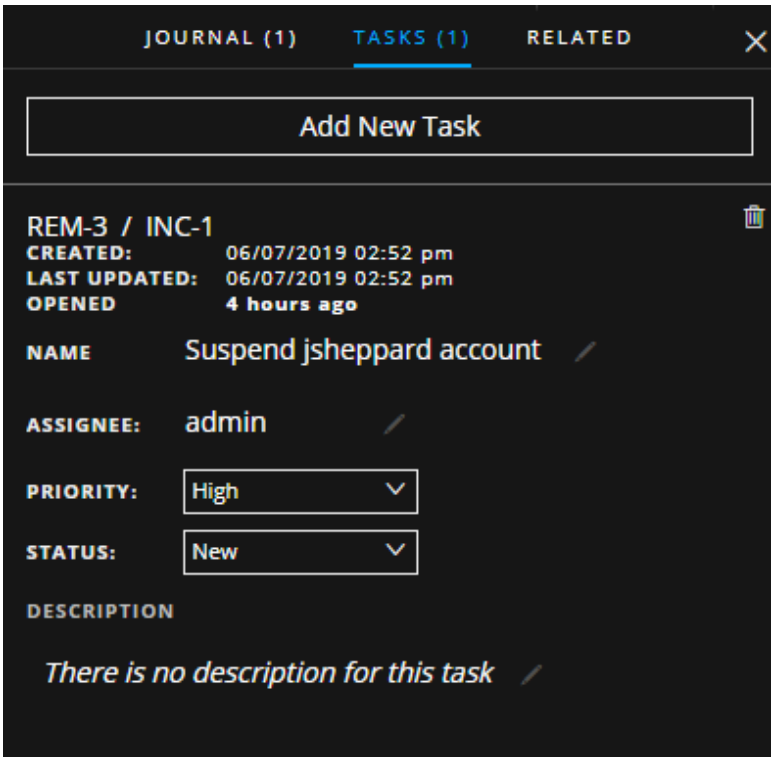
In the Events List, Chris can even inspect the details of the individual log events involved.



During her review and investigation, she can update the incident with her notes, as well as create and assign tasks for herself or other analysts.

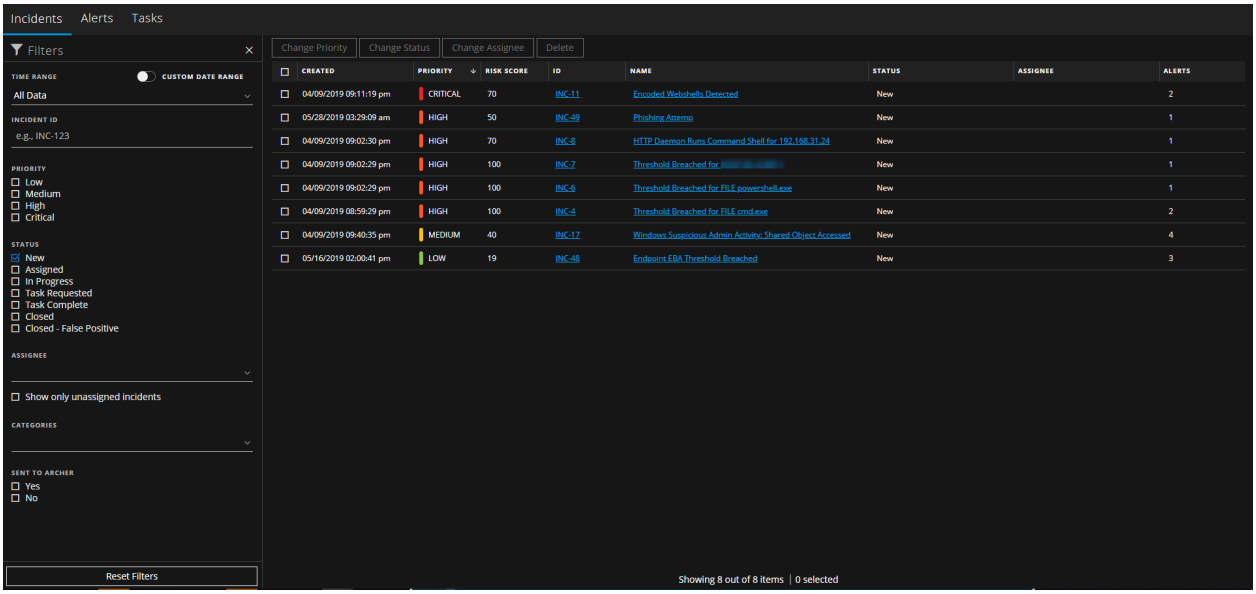


To remediate the incident, Chris opens a task for the administrator (admin) to suspend the jsheppard account.

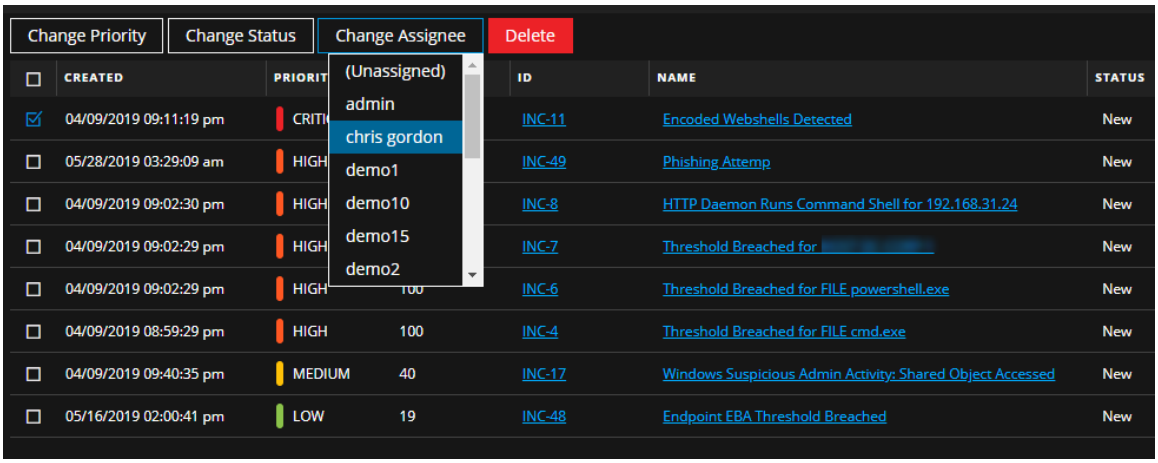


Use Case #2: Encoded Webshells Detected

Analyst Chris logs in, looks at all of the new incidents that have not yet been assigned to anyone, and notices a highly critical incident "Encoded Webshells Detected."



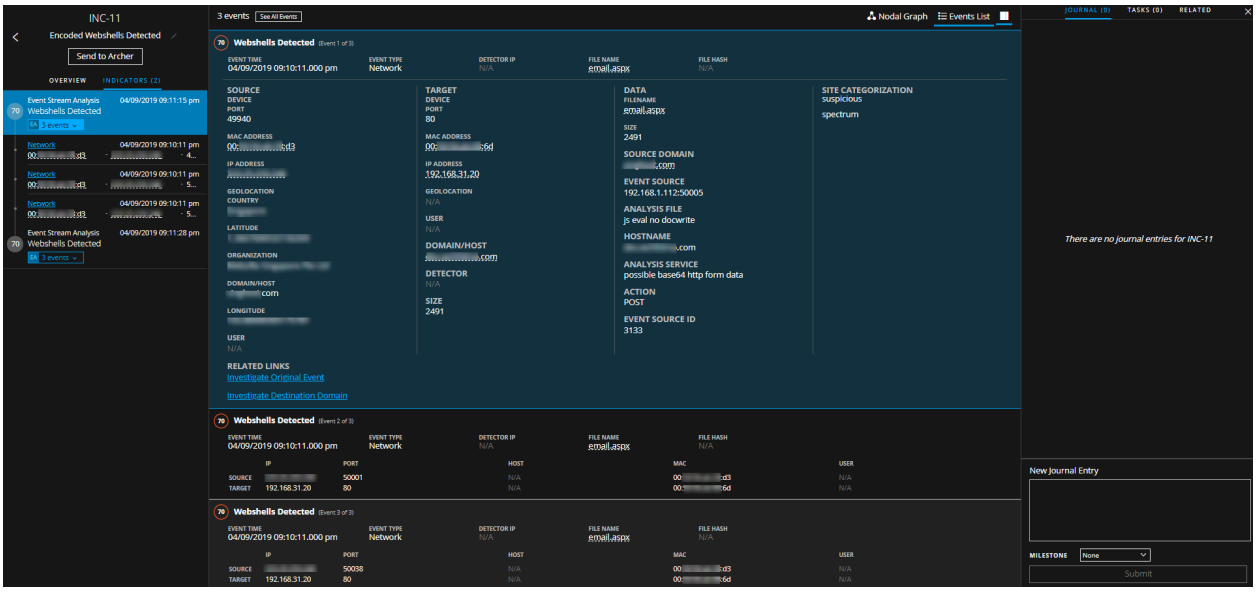
Chris decides to assign this incident to herself and investigate it.



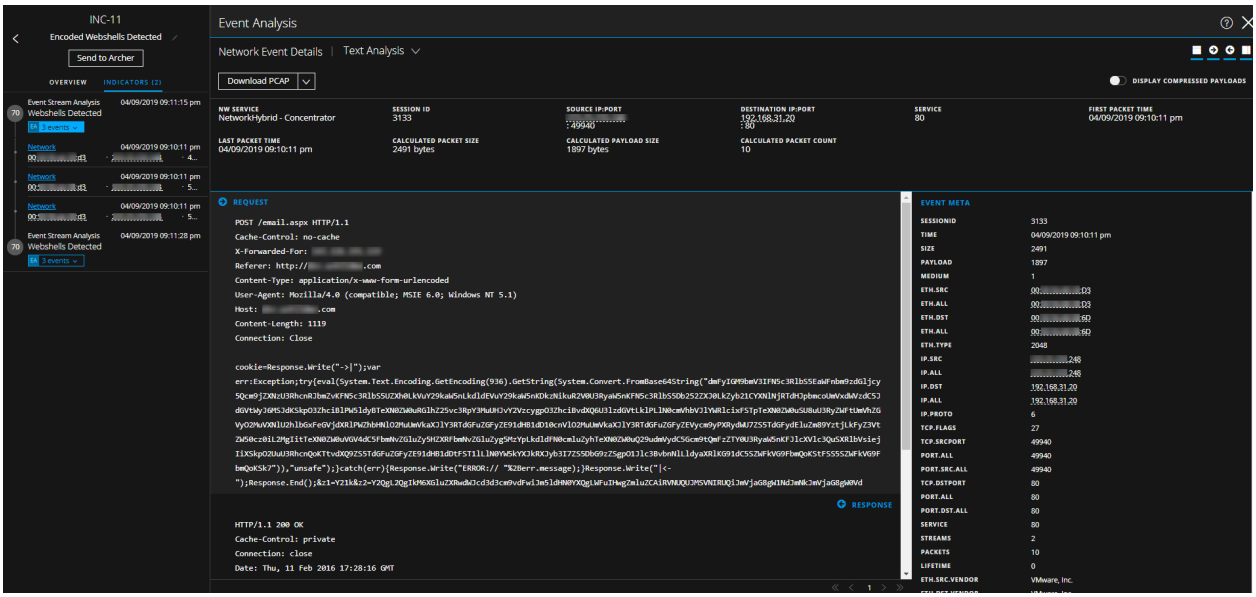
At first glance using the visual summary, Chris can see a couple of specific Event Stream Analysis alerts that kicked off this incident, and the entities associated with the alerts. She sees that a public IP address (xxx.xx.xxx.248) has been detected as interacting with a webshell on the internal web server 192.168.31.20. It looks like the suspicious request was made to the file `email.aspx`. She dives in to investigate.



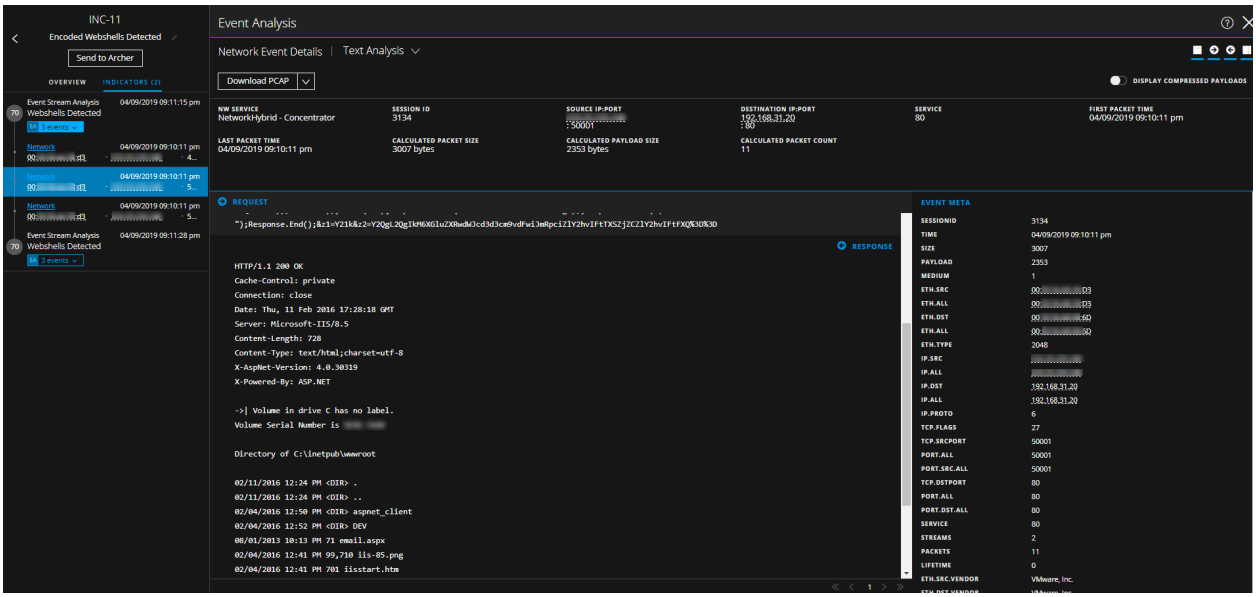
By drilling into the indicator (alert) on the left-hand side, Chris can view the entire list of events associated with the incident and all of the metadata generated by the system, including details about the connections between the external and internal host. She notices that the type of data in this case is "Network," meaning that these events were generated by the full packet capture component of NetWitness.



While the metadata associated with these events does look suspicious to Chris, she wants to investigate even deeper to see what is going on. To do this she can drill into the raw packet reconstruction of the HTTP sessions simply by clicking the "Network" link in the Indicators panel on the left-hand side.

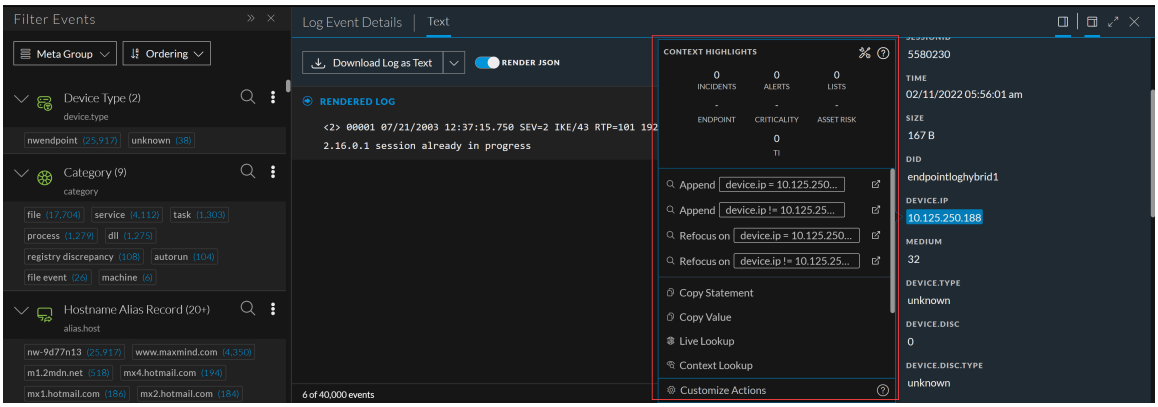


From this deep dive, Chris can examine the raw payload of the suspicious connections and can easily see the strange nature of this request. In the request payload she can see a strange looking request, but most alarmingly, once she scrolls down to the response portion she can see the information that her web server is sending back out to the suspect external IP address.

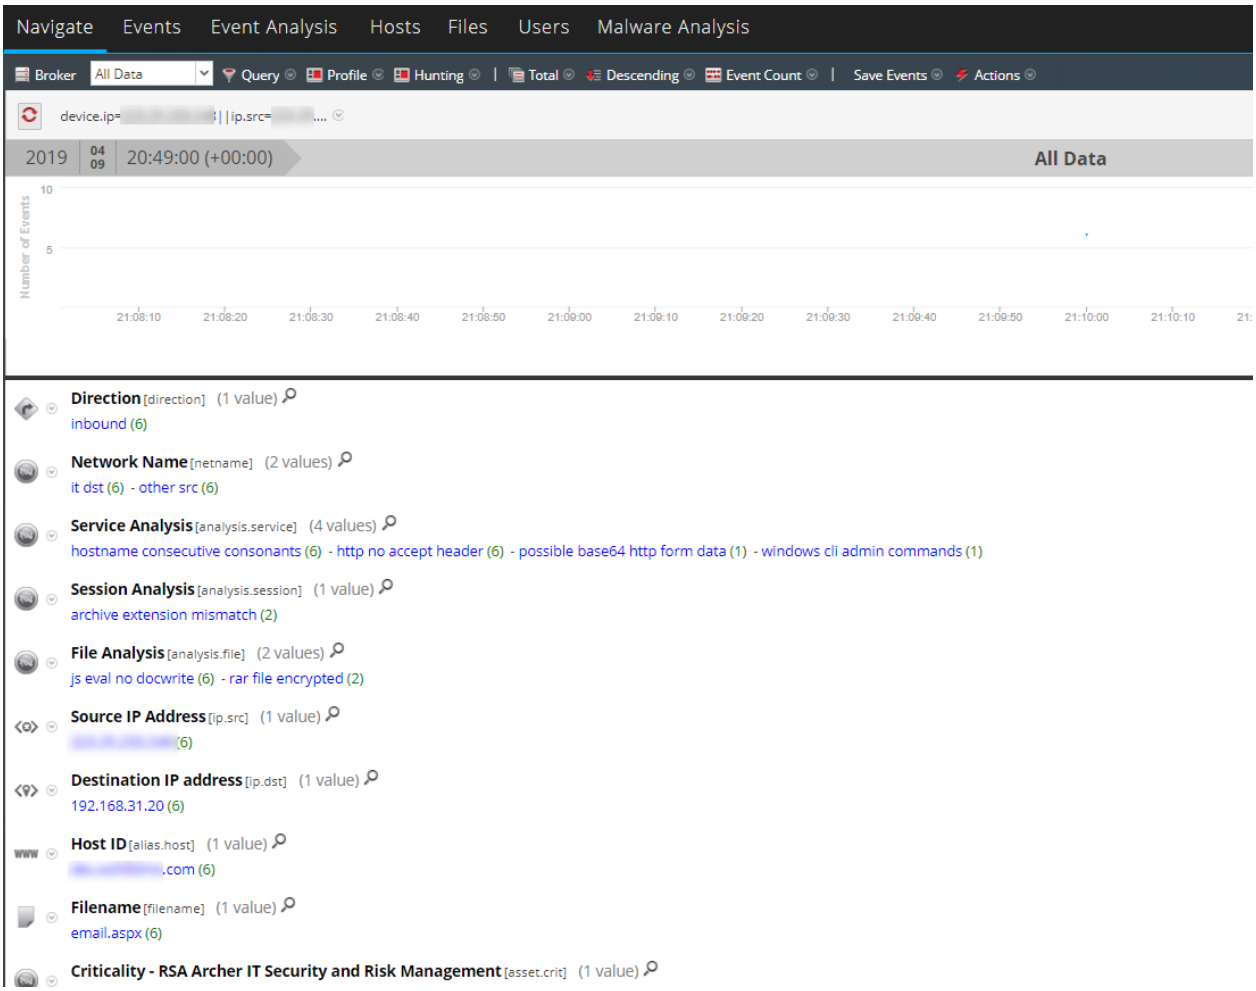


Chris sees what looks to be a directory listing in the packet data, which is not something she would expect from normal web site communications. With this information, she confirms that this is indeed a malicious webshell that has been installed on the internal web server.

From here, Chris can take a number of actions. She can journal her confirmation, assign a task to another user to handle the incident from there, or she can expand her investigation to look for any other activity that has been associated with the malicious external IP address. Chris does this by left or right click the IP address to open a context tooltip and pivoting into Investigate.



This pivot brings Chris into another part of the NetWitness interface where she can perform free-form search and analysis outside of the incident.



In this case she did not uncover anything other than the same network events that were part of the original incident, which is a good step in validating the isolated scope of the incident. If, however, she were to find other interesting events across any log, network, or endpoint data in the system, she could easily add those events into the incident to keep track.

Reviewing Alerts

NetWitness enables you to view a consolidated list of threat alerts generated from multiple sources in one location. You can find these alerts in the Respond > Alerts view. The source of the alerts can be ESA correlation rules, NetWitness Endpoint, Detect AI, Malware Analysis, Reporting Engine, Risk Scoring, as well as many others. You can see the source of the alerts, the alert severity, and additional alert details.

Note: ESA correlation rule alerts can ONLY be found in the Respond > Alerts view.

To better manage a large number of alerts, you have the ability to filter the alerts list based on criteria that you specify, such as severity, time range, and alert source. For example, you may want to filter the alerts to only show those alerts with a severity between 90 and 100 that are not already part of an incident. You can then select a group of alerts to create an incident or add to an existing incident.

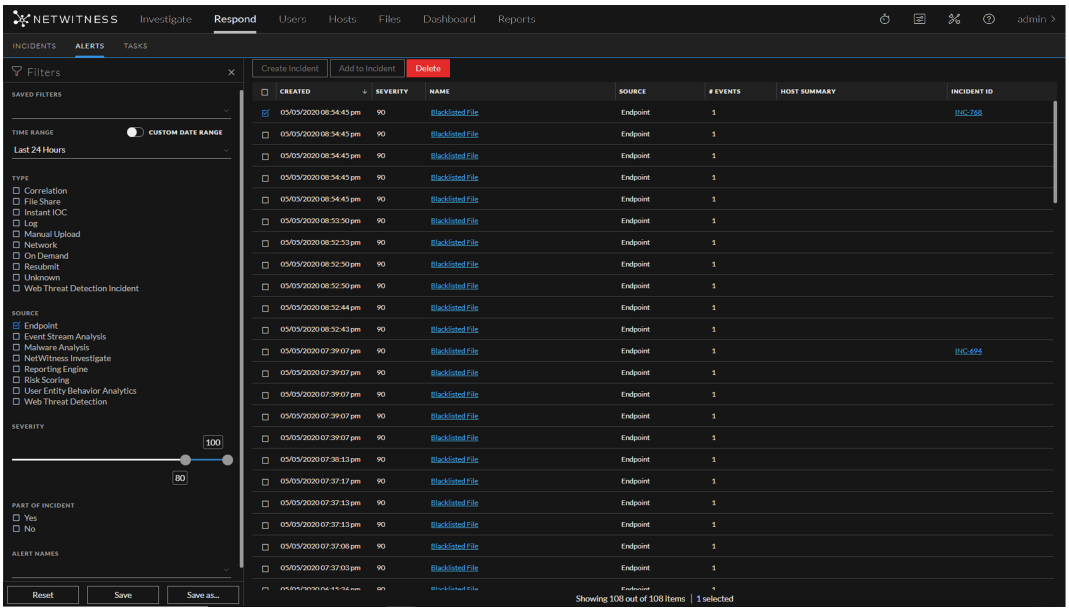
You can perform the following procedures to review and manage alerts:

- View Alerts
- Filter the Alerts List
- Remove My Filters from the Alerts List
- Save the Current Alerts Filter
- Update a Saved Alerts Filter
- Delete a Saved Alerts Filter
- View Alert Summary Information
- View Event Details for an Alert
- Investigate Events
- Create an Incident Manually
- Add Alerts to an Incident
- Delete Alerts

View Alerts

In the Alerts List view, you can browse through various alerts from multiple sources, filter them, and group them to create incidents. This procedure shows you how to access the alerts list.

1. Go to **Respond > Alerts**.
   The Alerts List view displays a list of all NetWitness alerts.



2. Scroll through the alerts list, which shows basic information about each alert as described in the following table.

Column	Description
Created	Displays the date and time when the alert was recorded in the source system.
Severity	Displays the level of severity of the alert. The values are from 1 through 100.
Name	Displays a basic description of the alert.
Source	Displays the original source of the alert. The source of the alerts can be NetWitness Endpoint, Detect AI, Malware Analysis, Event Stream Analysis (ESA Correlation Rules), Reporting Engine, Web Threat Detection, Risk Scoring, and many others. **Note:** In NetWitness Platform 11.3 and later, the **Endpoint** source includes Endpoint alerts from all NetWitness Endpoint versions. If one of the events in an alert has a device_type of `nwendpoint`, the source changes to Endpoint.
# Events	Indicates the number of events contained within an alert. This varies depending on the source of the alert. For example, NetWitness Endpoint and Malware Analysis alerts always have one Event. For certain types of alerts, a high number of events may mean that the alert is more risky.
Host Summary	Displays details of the host like the host name from where the alert was triggered. The details may include information about the source and destination hosts in an Alert. Some alerts may describe events across more than one host .

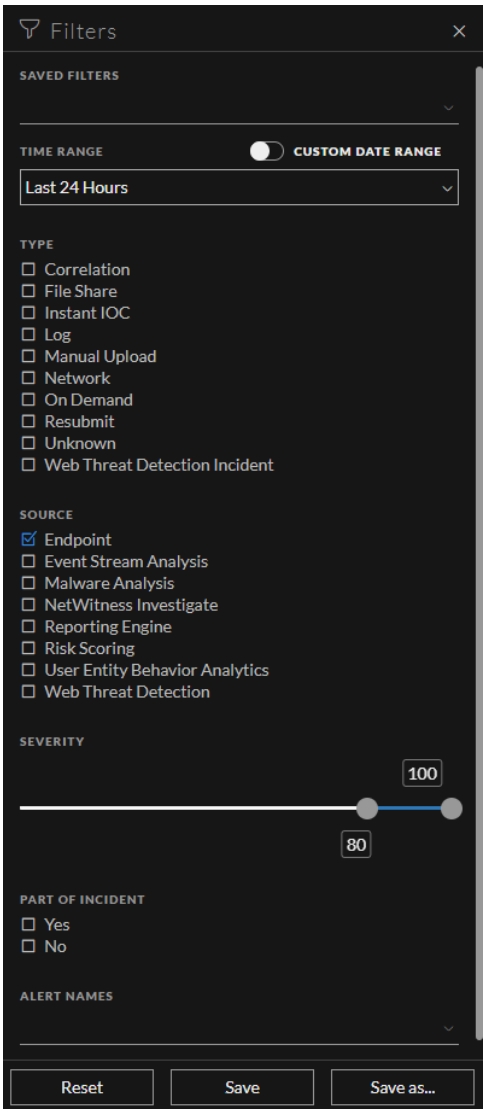
Column	Description
Incident ID	Shows the Incident ID of the alert. If there is no incident ID, the alert does not belong to any incident and you can create an incident to include this alert or the alert can be added to an existing incident.

At the bottom of the list, you can see the number of alerts on the current page and the total number of alerts. For example: **Showing 1000 out of 2069 items**

Filter the Alerts List

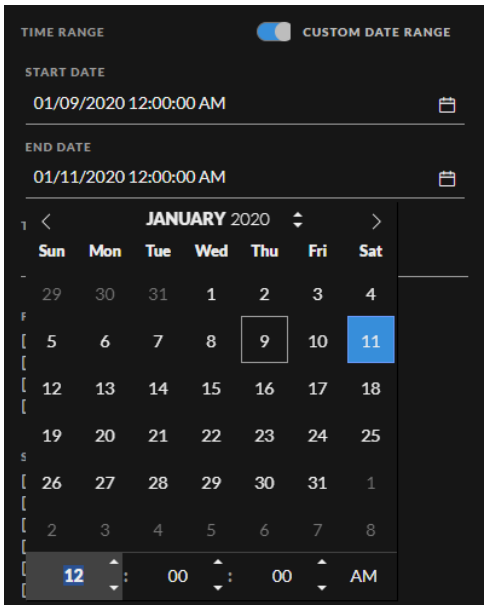
The number of alerts in the Alerts List can be very large, making it difficult to locate particular alerts. The Filter enables you to view the alerts you want to see, for example, alerts from a particular source, alerts of a particular severity, alerts that are not part of an incident, and so on.

1. Go to **Respond > Alerts**.
   The Filters panel appears to the left of the Alerts list. If you do not see the Filters panel, in the Alerts List view toolbar, click , which opens the Filters panel.



2. In the Filters panel, select one or more options to filter the alerts list:
   - **Time Range**: You can select a specific time period from the Time Range drop-down list. The time range is based on the date that the alerts were received. For example, if you select Last Hour, you can see alerts that were received within the last 60 minutes.
   - **Custom Date Range**: You can specify a specific date range instead of selecting a Time Range option. To do this, click the white circle in front of Custom Date Range to view the Start Date and

End Date fields. Select the dates and times from the calendar.



- **Type**: Select the type of events in the alert to view, for example, logs, network sessions, and so on. In NetWitness Platform 11.3 and later, if one of the events in an alert has a device_type of `nwendpoint`, Endpoint is included in the Type field.
- **Source**: Select one or more sources to view alerts triggered by the selected sources. For example, to view NetWitness Endpoint alerts only, select Endpoint as the source. In NetWitness Platform 11.3 and later, the **Endpoint** source includes Endpoint alerts from all NetWitness Endpoint versions. If one of the events in an alert has a device type of `nwendpoint`, the source changes to Endpoint. A **Risk Scoring** source is available in NetWitness Platform 11.3 and later. NetWitness Respond automatically creates incidents from alerts that are over the specified file and host alert thresholds for risk score. For more information, see the *NetWitness Respond Configuration Guide*.
- **Severity**: Select the the level of severity of the alerts to view. The values are from 1 through 100. For example, to concentrate on the highest severity alerts first, you may want to view only those alerts with a severity from 90 to 100.
- **Part of Incident**: To view only alerts that are not part of an incident, select **No**. To view only alerts that are part of an incident, select **Yes**. For example, when you are ready to create an incident from a group of alerts, you can select No to view only those alerts that are not currently part of an incident.
- **Alert Names**: Select the name of the alert to view. You can use this filter to search for all alerts generated by a specific rule, for example, Direct Login to an Administrative Account.

The Alerts List shows a list of alerts that meet your selection criteria. You can see the number of items in your filtered list at the bottom of the alerts list.
For example: **Showing 30 out of 30 items**

3. If you want to close the Filters panel, click **X**. Your filters remain in place until you remove them.

Remove My Filters from the Alerts List

NetWitness remembers your filter selections in the Alerts List view. You can remove your filter selections when you no longer need them. For example, if you are not seeing the number of alerts that you expect to see or you want to view all of the alerts in your alerts list, you can reset your filters.

1. Go to **Respond > Alerts**.
   The Filters panel appears to the left of the alerts list. If you do not see the Filters panel, in the Alerts List view toolbar, click , which opens the Filters panel.
2. At the bottom of the Filters panel, click **Reset Filters**.

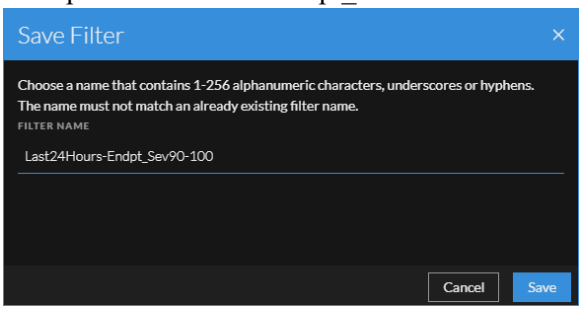
Save the Current Alerts Filter

Note: This option is available in NetWitness Platform Version 11.5 and later.

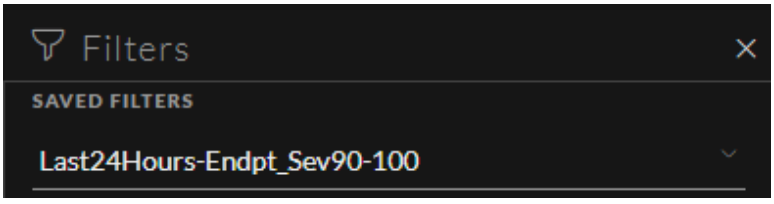
Saved filters provide a way for analysts to save and quickly apply specific filter conditions to the list of alerts. You can also use these filters to customize the Springboard landing page. For example, you may want to create a filter to show only alerts from a specific source and severity level over the last 24 hours. (This option is available in NetWitness Platform 11.5 and later.)

Saved filters are global. You can save a filter for other analysts to use and you can use any saved filter.

1. In the Filters panel, select one or more options to filter the incidents list. For example, in the Time Range field select Last 24 Hours, in the Source field select Endpoint, and for Severity, select the 90 to 100 range.
2. Click **Save As** and in the **Save Filter** dialog, enter a unique name for the filter and save it, for example Last24Hours-Endpt_Sev90-100.



The filter is added to the **Saved Filters** list.



Update a Saved Alerts Filter

Note: This option is available in NetWitness Platform Version 11.5 and later.

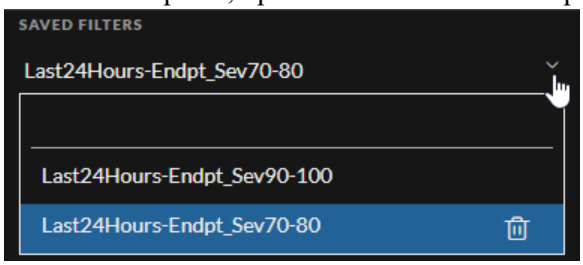
1. In the Filters panel **Saved Filters** drop-down list, select a saved filter.
2. Update your filter selections and click **Save**.

Delete a Saved Alerts Filter

Note: This option is available in NetWitness Platform Version 11.5 and later.

When a saved filter is no longer required, you can remove it from the saved filters list. Filters used in the Springboard cannot be deleted.

1. In the Filters panel, open the **Saved Filters** drop-down list.

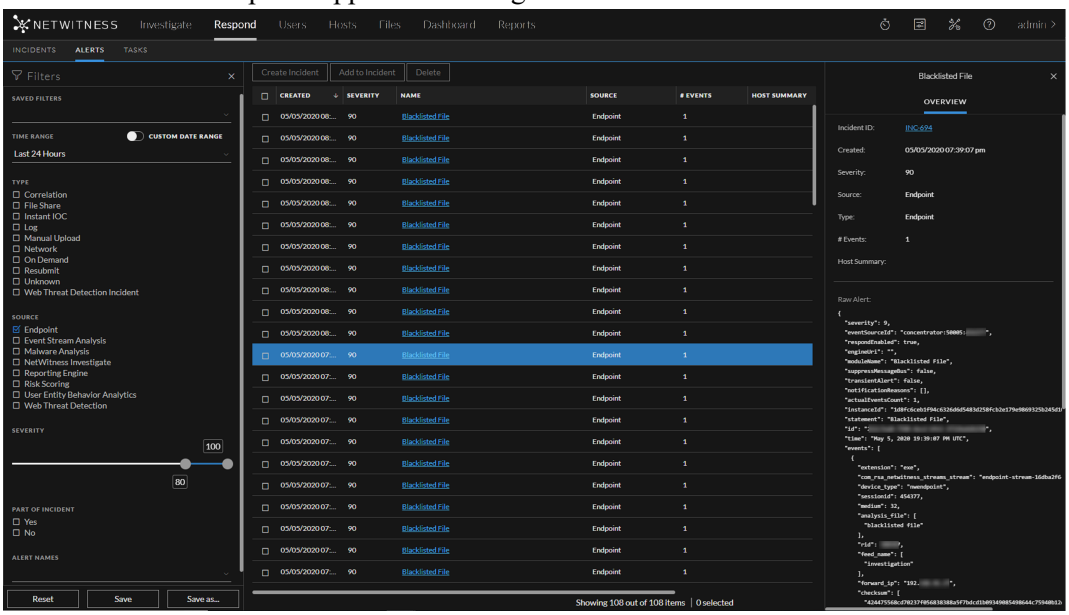


2. Next to the filter name, click to delete it.

View Alert Summary Information

In addition to viewing basic information about an alert, you can also view raw alert metadata in the Overview panel.

1. In the Alerts list, click the alert that you want to view.
   The Alert Overview panel appears to the right of the Alerts list.

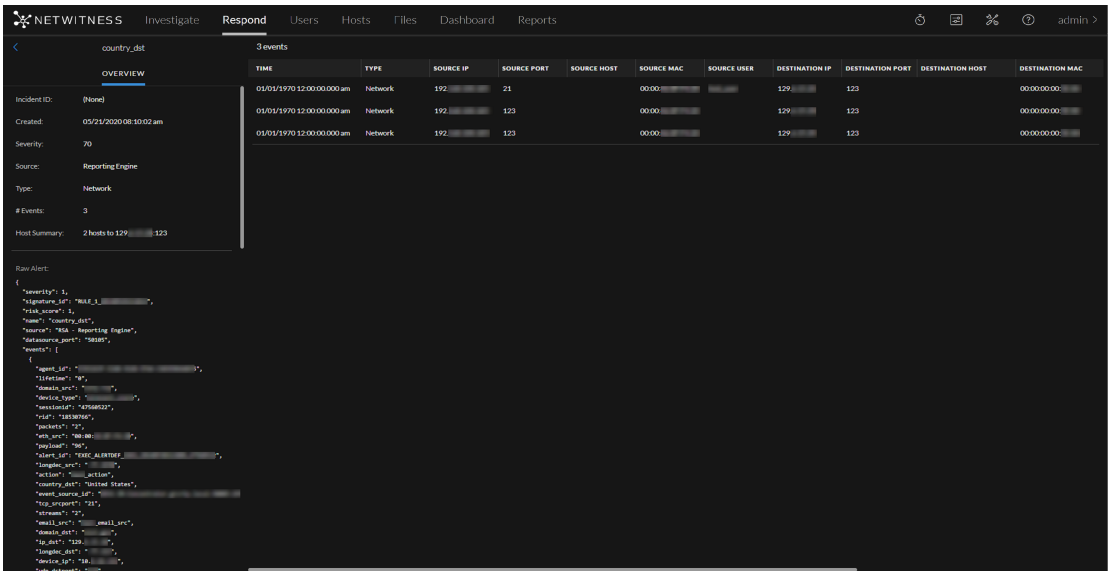


2. In the Overview panel Raw Alert section, you can scroll to view the raw alert metadata.



View Event Details for an Alert

After you review the general information about the alert in the Alerts List view, you can go to the Alert Details view for more detailed information to determine the action required. An alert contains one or more events. In the Alert Details view, you can drill down into an alert to get additional event details and further investigate the alert. The following figure shows an example of the Alert Details view.



The Overview panel on the left has the same information for an alert as the Overview panel in the Alerts List view.

The Events panel on the right shows information about the events in the alert, such as event time, source IP, destination IP, detector IP, source user, destination user, and file information about the events. The amount of information listed depends on the event type.

There are two types of events:

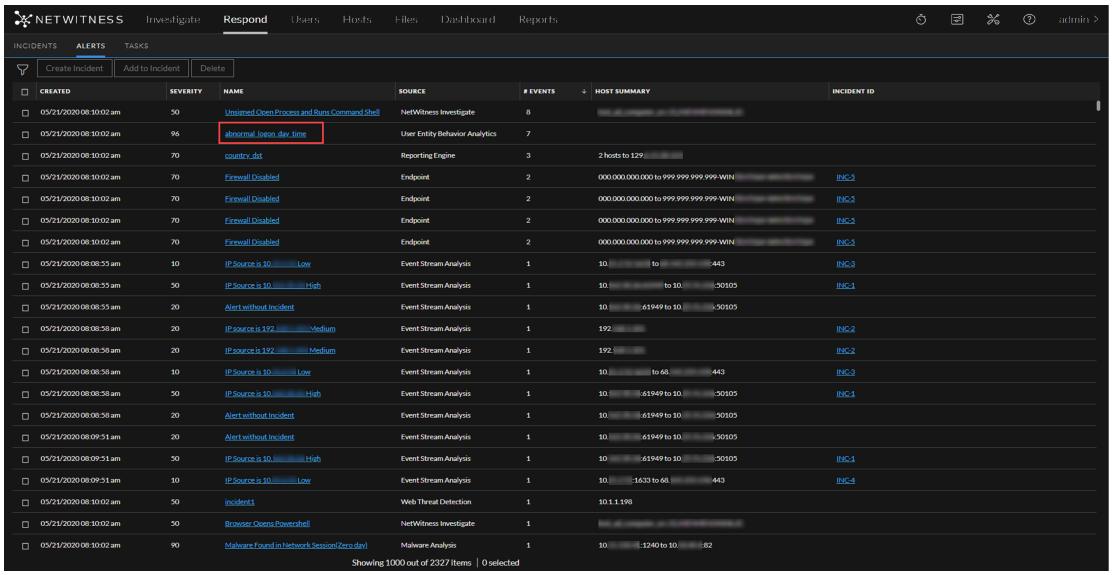
- A transaction between two machines (a Source and a Destination)
- An anomaly detected on a single machine (a Detector)

Some events will only have a Detector. For example, NetWitness Endpoint finds malware on your machine. Other events will have a Source and Destination. For example, packet data shows communication between your machine and a Command and Control (C2) domain.

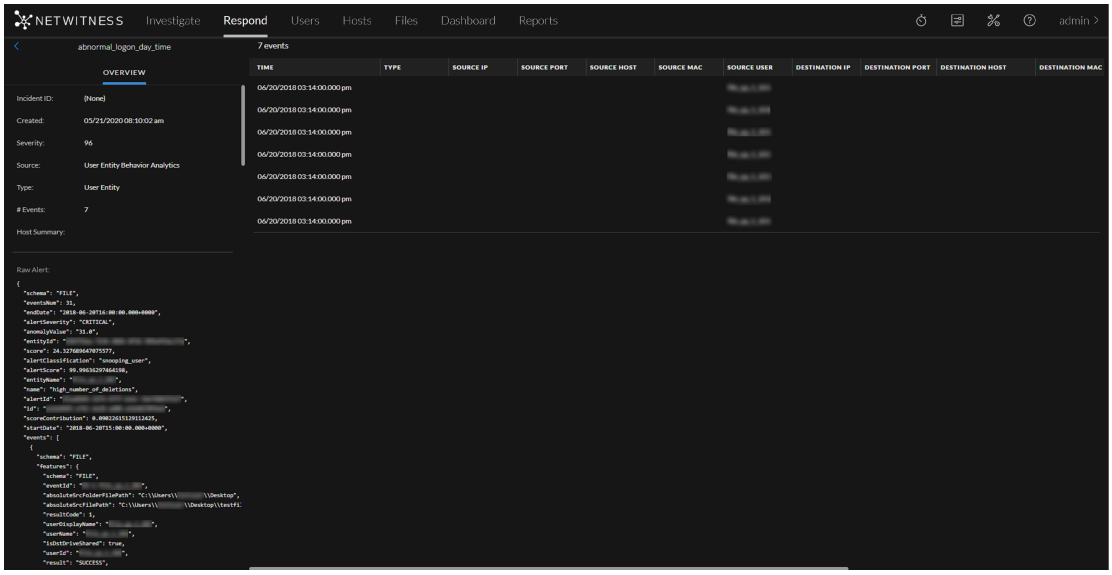
You can drill further into an event to get detailed data about the event.

To View the Event Details for an Alert:

1. To view event details for an alert, in the Alerts List view, choose an alert to view and then click the link in the **Name** column for that alert.



The Alerts Details view shows the Overview panel on the left and the Events panel on the right.



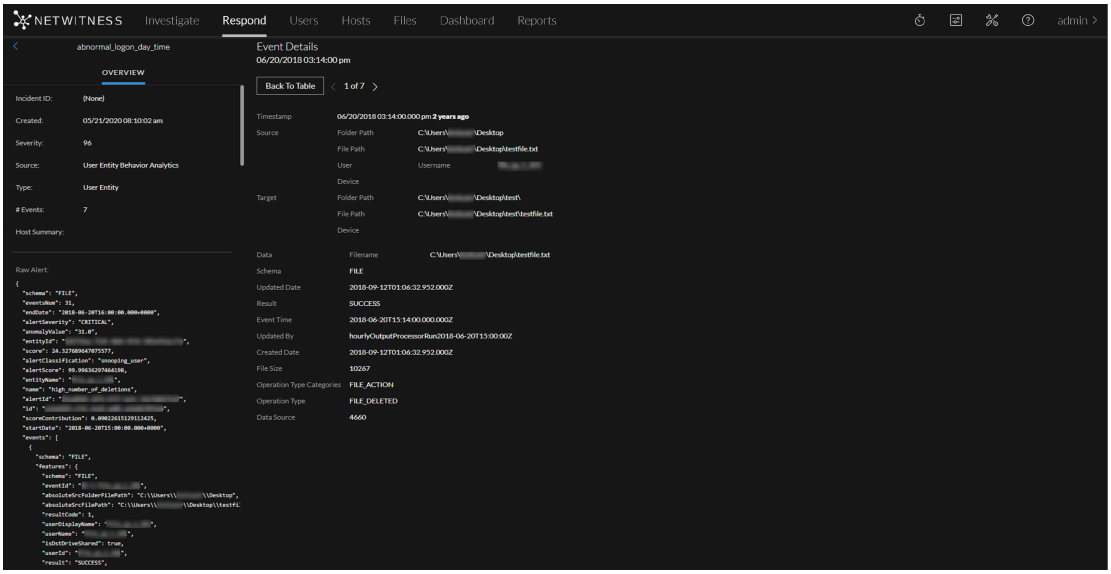
The Events panel shows a list of events with information about each event. The following table shows some of the columns that can appear in the Events List (Events Table).

Column	Description
Time	Shows the time the event occurred.
Type	Shows the type of alert, such as Log and Network.
Source IP	Shows the source IP address if there was a transaction between two machines.
Destination IP	Shows the destination IP address if there was a transaction between two machines
Detector IP	Shows the IP address of the machine where an anomaly was detected.

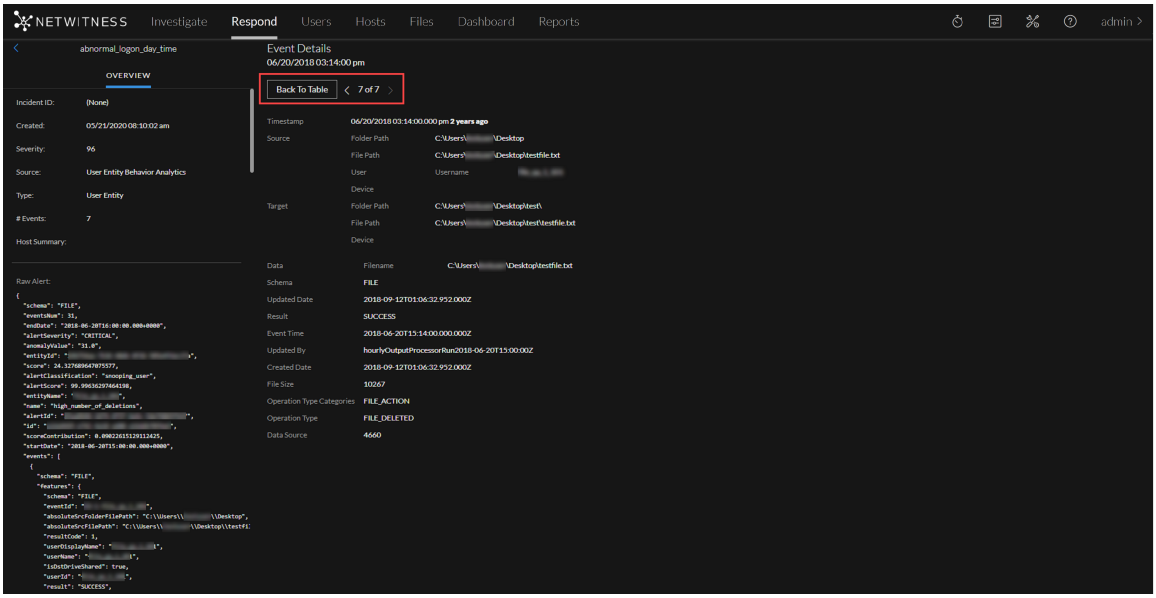
Column	Description
Source User	Shows the user of the source machine.
Destination User	Shows the user of the destination machine.
File Name	Shows the file name if a file is involved with the event.
File Hash	Shows a hash of the file contents.

If there is only one event in the list, you see only the event details for that event instead of a list.

2. Click an event in the Events list to view the Event details.
   This example shows the event details for the first event in the list.



3. Use the page navigation to the right of the Back To Table button to view other events. This example shows the event details for the last event in the list.



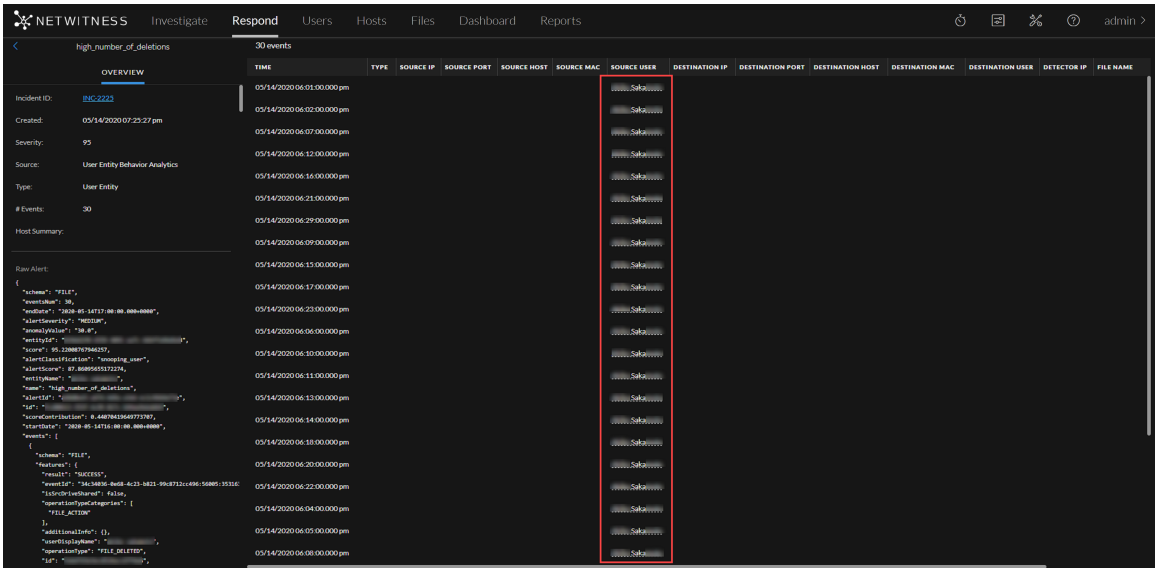
See Alert Details Panel for detailed information about the event data listed in the Alert Details panel.

Investigate Events

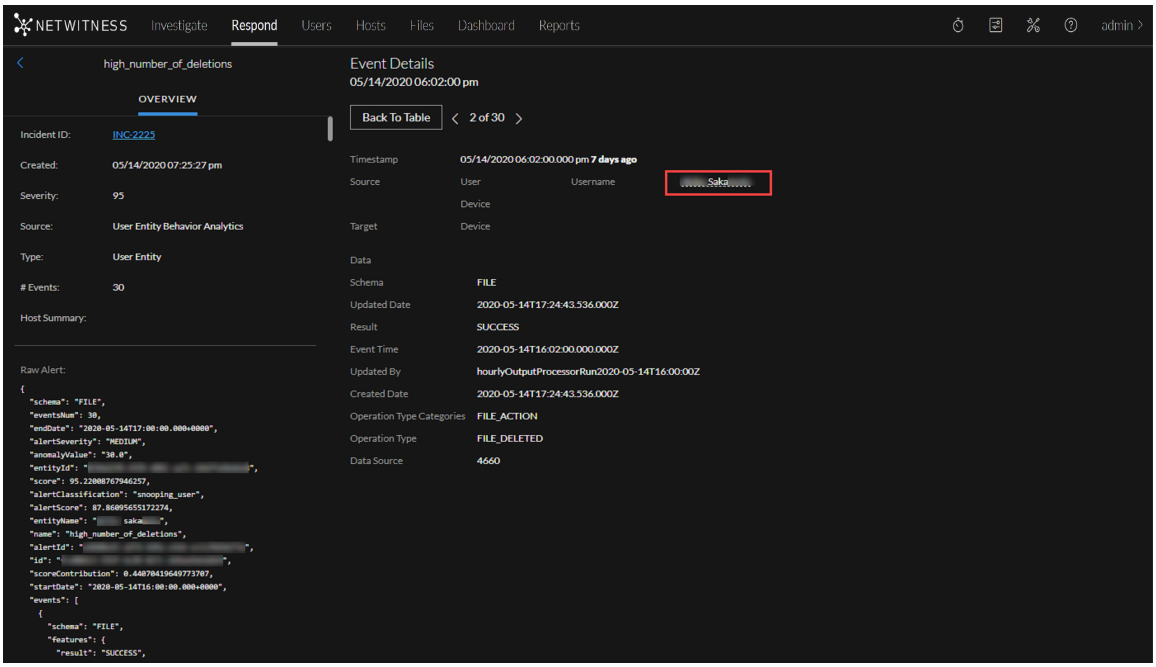
To further investigate the events, you can find links that take you to additional contextual information. From there, you have options available depending on your selection.

View Contextual Information

In the Alert Details view, you can see underlined entities in the Events panel. An underlined entity is considered an entity in the Context Hub and has additional contextual information available. The following figure shows underlined entities in the Events list.



The following figure shows an underlined entity in the Event Details.

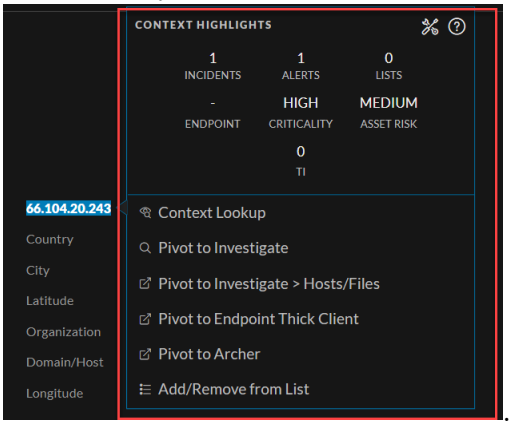


The Context Hub is preconfigured with meta fields mapped to the entities. NetWitness Respond and NetWitness Investigate use these default mappings for context lookup. For information about adding meta keys, see "Configure Settings for a Data Source" in the *Context Hub Configuration Guide*.

Caution: For the Context Lookup to work correctly in the Respond and Investigate views, NetWitness recommends that when mapping meta keys in the **(missing or bad snippet) > System > Investigation > Context Lookup** tab, you add only meta keys to the Meta Key Mappings, not fields in the MongoDB. For example, ip.address is a meta key and ip_address is not a meta key (it is a field in the MongoDB).

To View Contextual Information:

1. In the Alert Details view Events List or Event Details, left or right click an underlined entity.
   A context tooltip appears with a quick summary of the type of context data that is available for the selected entity.



   The information in the **Context Highlights** section helps you to determine the actions that you would like to take. It shows the number of related alerts and incidents. It can show related data for

Incidents, Alerts, Lists, Endpoint, Criticality, Asset Risk, Reputation, and Threat Intelligence (TI). Depending on your data, you may be able to click these numbered items for more information. The above example shows 1 related incidents, 1 related alerts, and one list associated with the selected IP address. There is no information for Endpoint, , Criticality, or Asset Risk. TI information comes from the STIX data source configured in Context Hub. For more information, see the *Context Hub Configuration Guide*.

The other section lists the available actions. In the above example, the Add/Remove From List, Pivot to Investigate, Pivot to Investigate > Hosts/Files, Pivot to Endpoint Thick Client, and and Pivot to Archer options are available.

Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer Datasource is not responding. Check that the Archer configuration is enabled and configured properly.

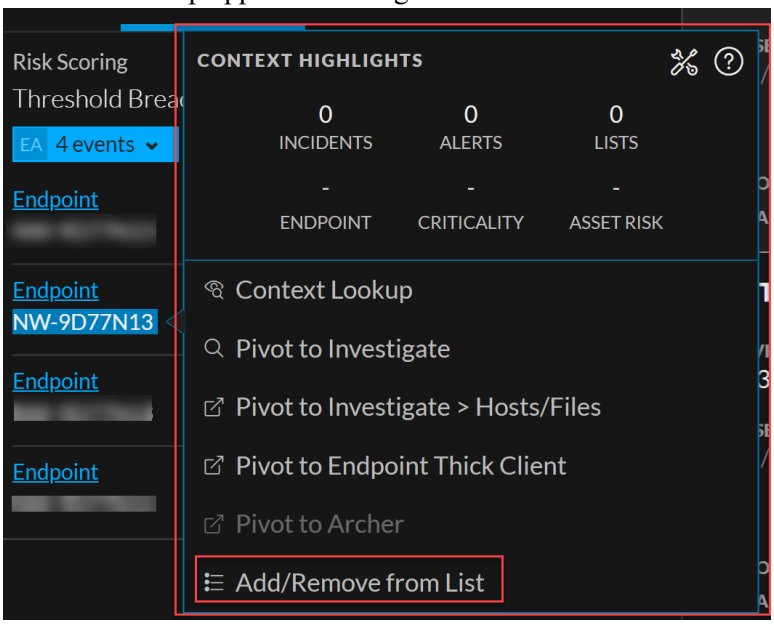
For more information, see Pivot to the Investigate > Navigate View, Pivot to the Hosts or Files View, Pivot to Archer, Pivot to Endpoint Thick Client, and Add an Entity to a Whitelist.

2. To see more details about the selected entity, click the **View Context** button.
   The Context panel opens and shows all of the information related to the entity.
   Context Lookup Panel - Respond View provides additional information.

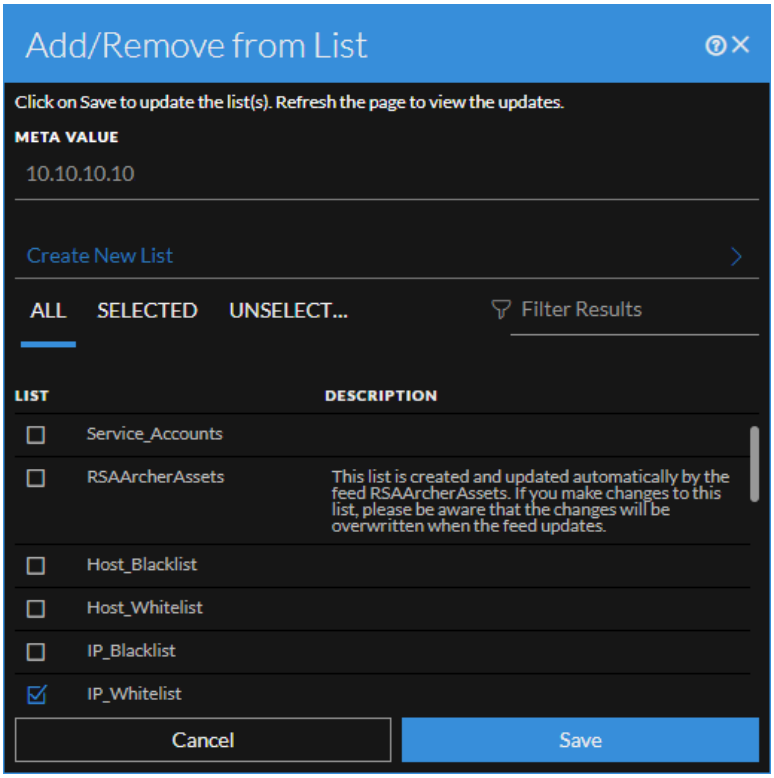
Add an Entity to a Whitelist

You can add any underlined entity to a list, such as a Whitelist or Blacklist, from a context tooltip. For example, to reduce false positives, you may want to whitelist an underlined domain to exclude it from the related entities.

1. In the Alert Details view Events List or Event Details, left or right click the underlined entity that you would like to add to a Context Hub list.
   A context tooltip appears showing the available actions.



2. In the **Actions** section of the tooltip, click **Add/Remove from List**.
   The Add/Remove From List dialog shows the available lists.



3. Select one or more lists and click **Save**.
   The entity appears on the selected lists.
   Add/Remove from List Dialog provides additional information.

Create a Whitelist

You can create a whitelist in the Context Hub in the same way as you would create it in the Incident Details view, see Create a List.

Pivot to the Investigate > Navigate View

For a more thorough investigation of the incident, you can access the Investigate > Navigate view.

1. In the Events List or Event Details in the Alert Details view, hover over any underlined entity to access a context tooltip.
2. In the **Actions** section of the tooltip, select **Pivot to Investigate > Navigate**.
   The Navigate view opens, which enables you to perform a deeper dive investigation.

For more information, see the *NetWitness Investigate User Guide*. For troubleshooting information with the Investigate > Navigate link see the *Alerting with ESA Correlation Rules User Guide*.

Pivot to the Hosts or Files View

For a more thorough investigation about specific Hosts and Files, you can access the Hosts and Files views.

1. In the Events List or Event Details in the Alert Details view, left or right click any entity to access a context tooltip.
2. In the tooltip, select **Pivot to Investigate > Hosts/Files**.
   If you left or right click a host or IP or MAC address entity and click **Pivot to Investigate > Hosts/Files**, it displays the Hosts view with a specific host listed.
   If you left or right click a filename or file hash entity and click **Pivot to Investigate > Hosts/Files** it displays the Files view with a specific file listed.

Note: By default, the search for entities is on the previously selected Endpoint Server. However, you can select a different Endpoint Server to fetch the information or data.

For more information, see the *NetWitness Investigate User Guide*.

Pivot to Endpoint Thick Client

If you have the NetWitness Endpoint thick client application installed, you can launch it through the context tooltip. From there, you can further investigate a suspicious IP address, Host, or MAC address.

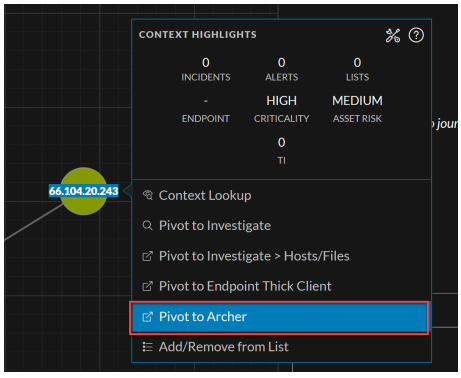
1. In the Events List or Event Details in the Alert Details view, hover over any underlined entity to access a context tooltip.
2. In the **Actions** section of the tooltip, select **Pivot to Endpoint Thick Client**.
   The NetWitness Endpoint thick client application opens outside of your web browser.

For more information on the thick client, see the *NetWitness Endpoint User Guide*.

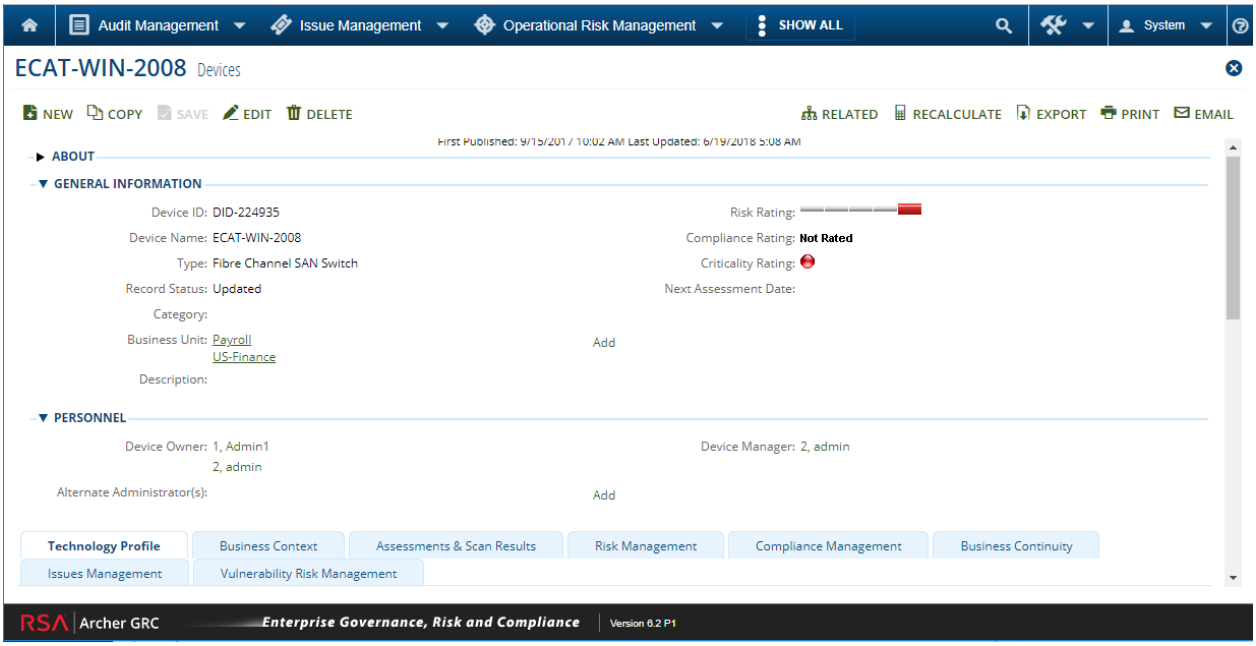
Pivot to Archer

For viewing more details about a device in Archer Cyber Incident & Breach Response, you can pivot to the device details page. This information is displayed only for IP address, host, and Mac address.

1. In the Events List or Event Details in the Alert Details view, left or right click any underlined entity to access a context tooltip.
2. In the **Actions** section, select **Pivot to Archer**.



3. The device details page in **RSA Archer Cyber Incident & Breach Response** opens if you are logged in to the application, otherwise the login screen is displayed.



Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer Datasource is not responding. Check that the Archer configuration is enabled and configured properly.

For more information, see the *NetWitness Archer Integration Guide*.

Create an Incident Manually

You can create incidents manually from alerts in the Alerts List view. The alerts that you select cannot be part of another incident.

In NetWitness Version 11.2 and later, you can change the assignee, category, and priority when you create an incident manually from alerts.

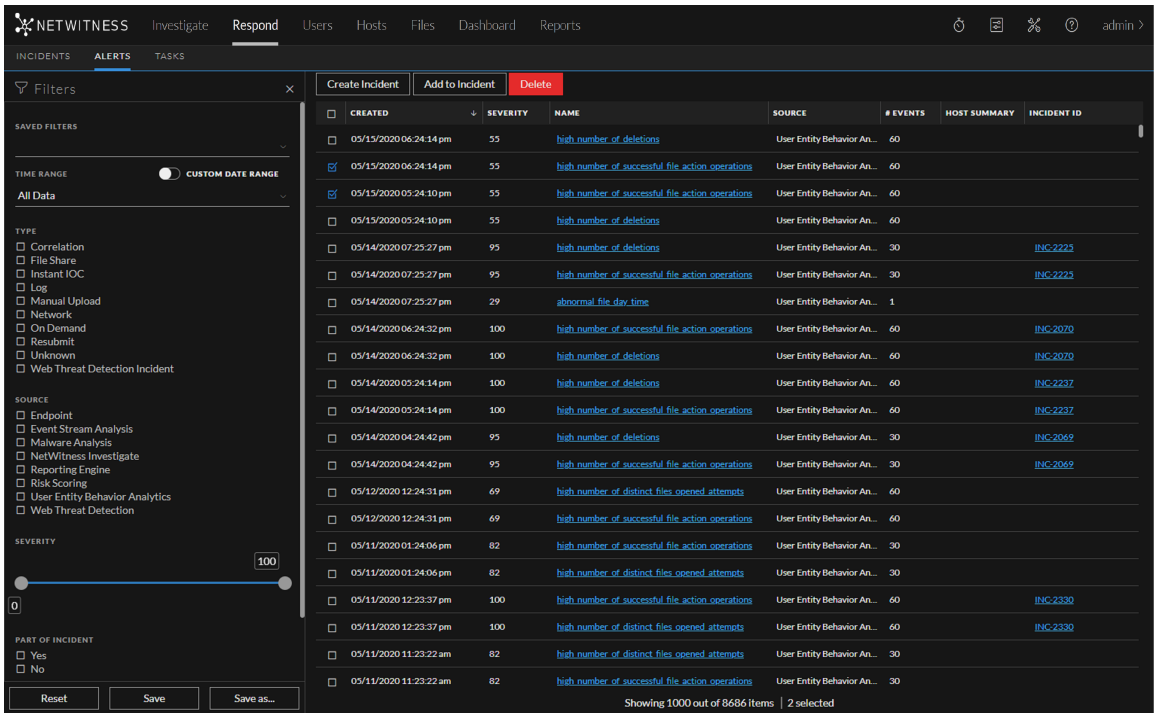
In NetWitness Version 11.1, incidents created manually from alerts default to Low priority, but you can change the priority after you create it. You cannot add categories to manually created incidents in version 11.1.

Note: Incidents can be created manually or automatically. An Alert can only be associated with one Incident. You can create incident rules to analyze the alerts collected and group them into incidents depending on which rules they match. For details, see the "Create an Incident Rule for Alerts" topic in the *NetWitness Respond Configuration Guide*.

To Create an Incident Manually:

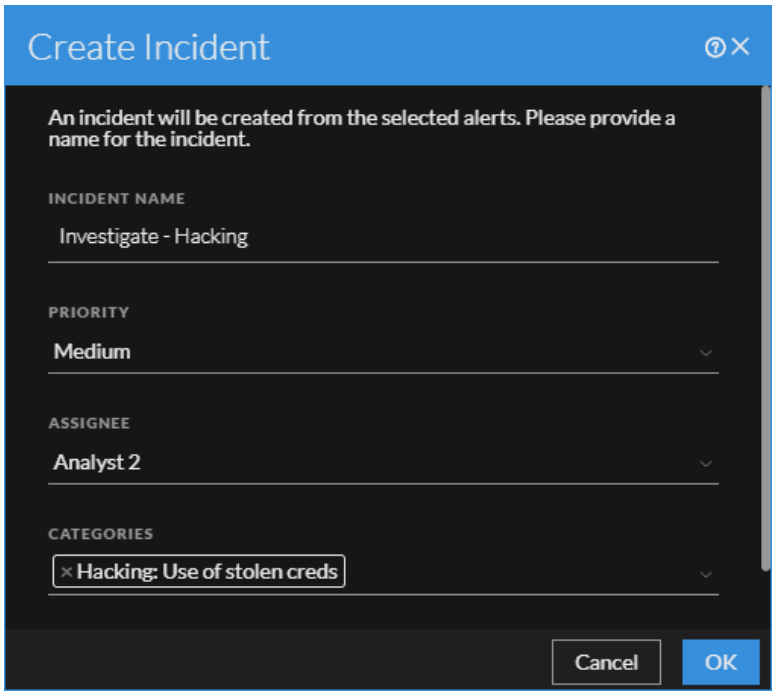
1. Go to **Respond > Alerts**.
2. Select one or more alerts in the Alerts List.

Note: Selecting alerts that do not have incident IDs enable the **Create Incident** button. If the alert is already part of an incident, the button is disabled. You can filter alerts that are not part of an incident by selecting the option **Part of Incidents** as **No** in the Filters panel.



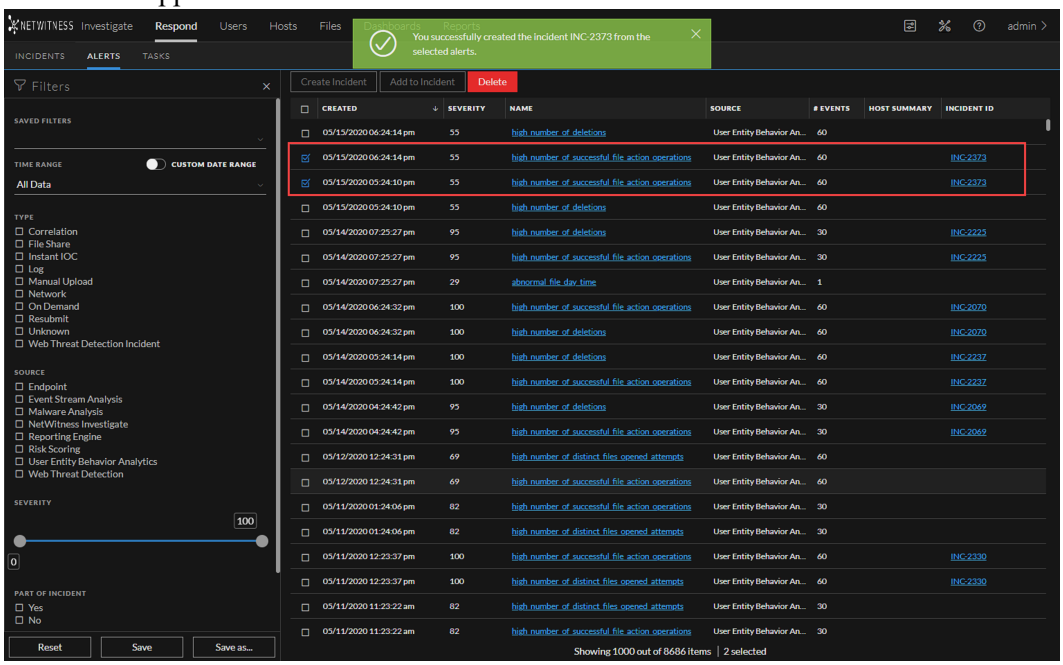
3. Click **Create Incident**.

   The **Create Incident** dialog is displayed.

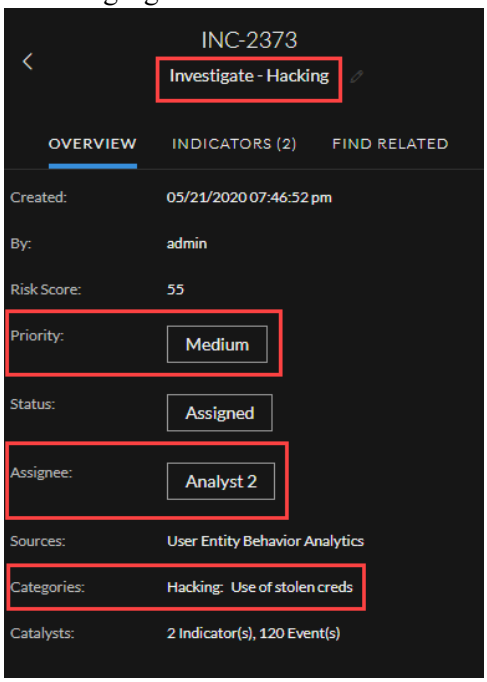


4. In the **Incident Name** field, type a name to identify the incident. For example, Investigate - Hacking.
5. In the **Priority** field, select a priority for the incident. The priority defaults to Low.
6. (Optional) If you are ready to assign the incident, in the **Assignee** field, select a specific user.
7. (Optional) In the **Categories** field, you can select a category to classify the incident, such as Hacking: Use of stolen creds. This is also helpful when trying to locate the incident later using the incidents filter.

8. Click **OK**.
   You can see a confirmation message that an incident was created from the selected alerts. The new incident ID appears as a link in the INCIDENT ID column of the selected alerts.



If you click the link, it takes you to the Incident Details view for that incident, where you can update information, such as changing Priority to high or assigning the incident to another user. The following figure shows the Incident Details view Overview panel for the new incident.



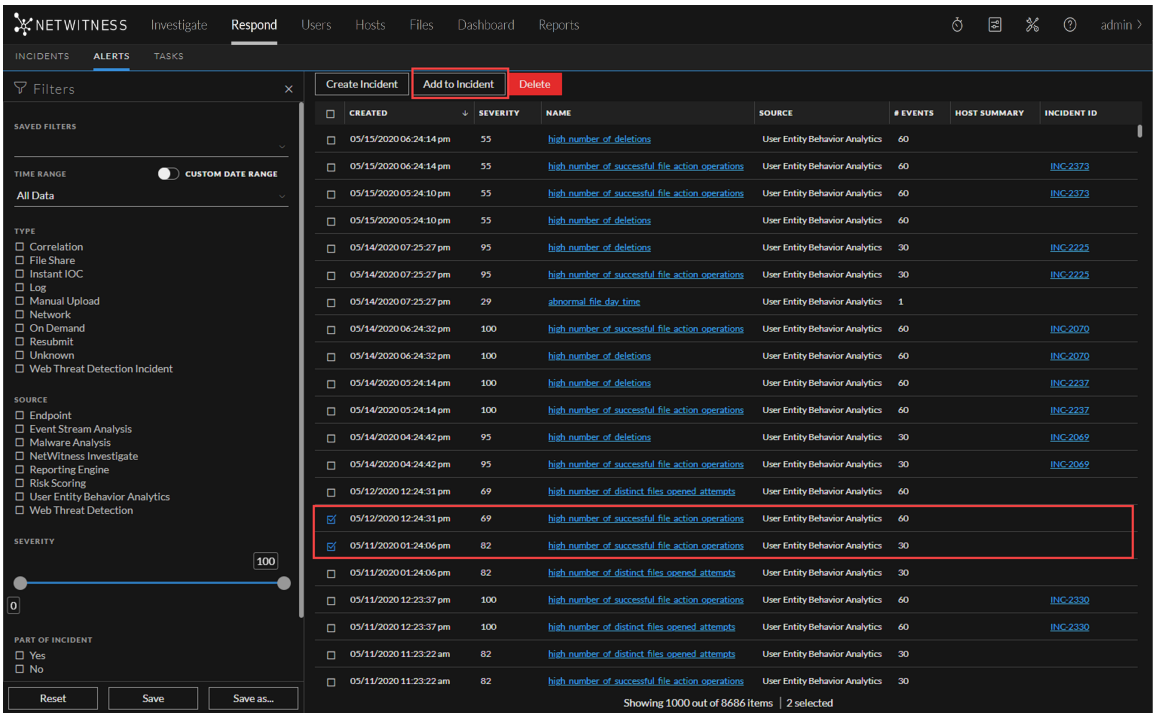
Add Alerts to an Incident

Note: This option is available in NetWitness Version 11.1 and later.

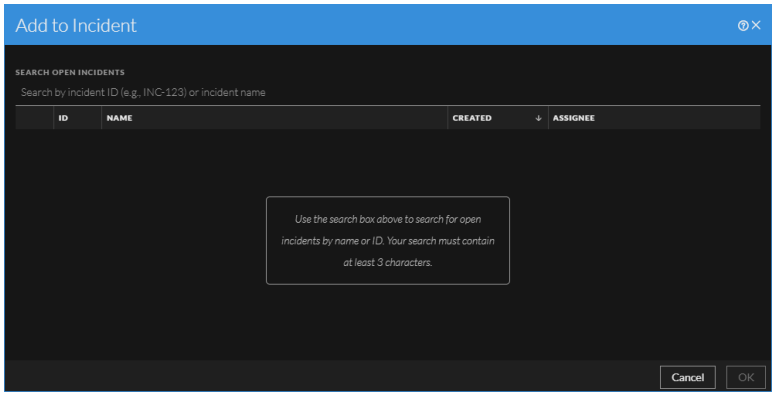
If you have alerts that fit a particular existing incident, you do not have to create a new incident. Instead, you can add alerts to that incident from the Alerts List view. The alerts that you select cannot be part of another incident.

1. Go to **Respond > Alerts**.
2. In the Alerts List, select one or more alerts that you want to add to an incident, and click **Add to Incident**.

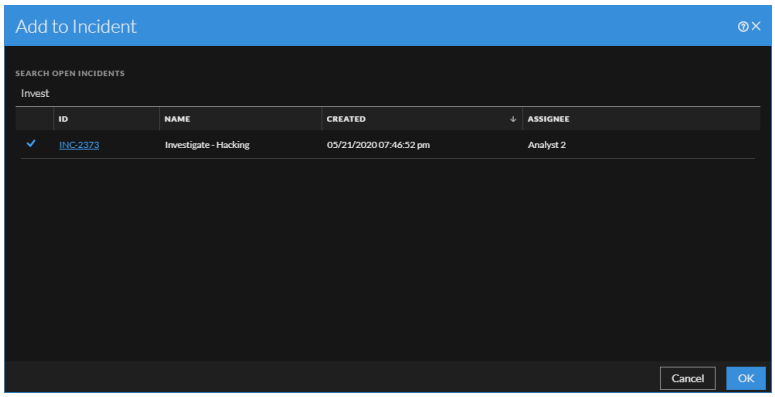
Note: Selecting alerts that do not have incident IDs enables the **Add to Incident** button. If the alert is already part of an incident, the button is disabled. You can filter alerts that are not part of an incident by selecting the option **Part of Incident** as **No** in the Filters panel.



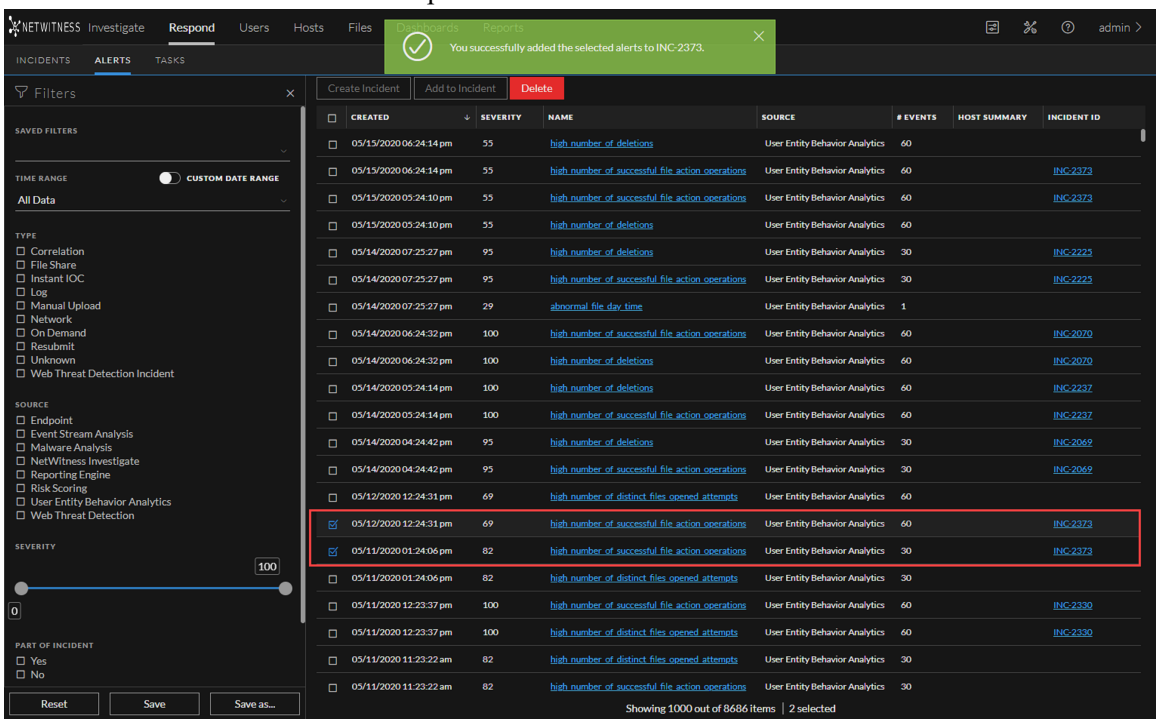
3. In the **Add to Incident** dialog, type at least three characters in the **Search** field to search for the incident by **Name** or **Incident ID**.



4. In the results list, select the incident that will receive the selected alerts and click **OK**.



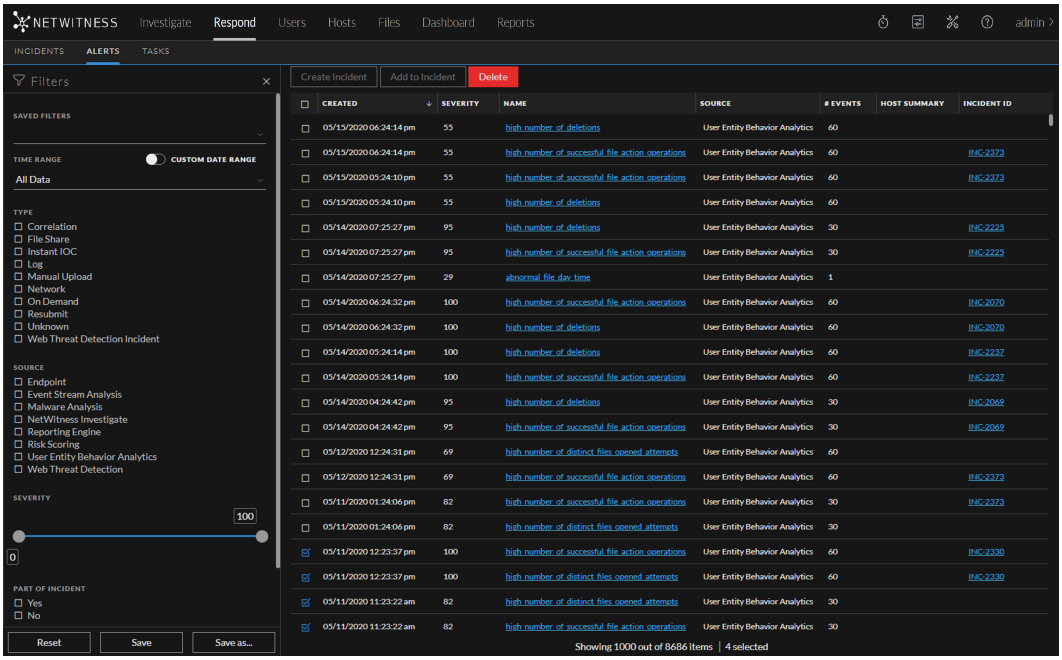
The selected alert or alerts are now part of the selected incident and will have that incident ID.



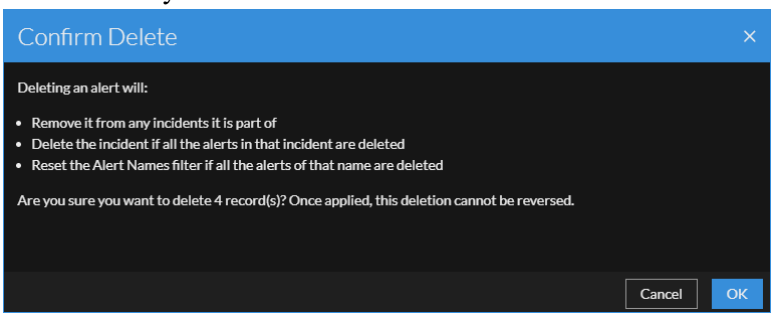
Delete Alerts

Users with the appropriate permissions, such as Administrators and Data Privacy Officers, can delete alerts. This procedure is helpful when you want to remove unnecessary or non-relevant alerts. Deleting these alerts frees up disk space.

1. Go to **Respond > Alerts**.
   The Alerts List view displays a list of all NetWitness alerts.
2. In the Alerts list, select the alerts that you want to delete and click **Delete**.



   If you do not have permission to delete alerts, you will not see the Delete button.
3. Confirm that you want to delete the alerts and click **OK**.

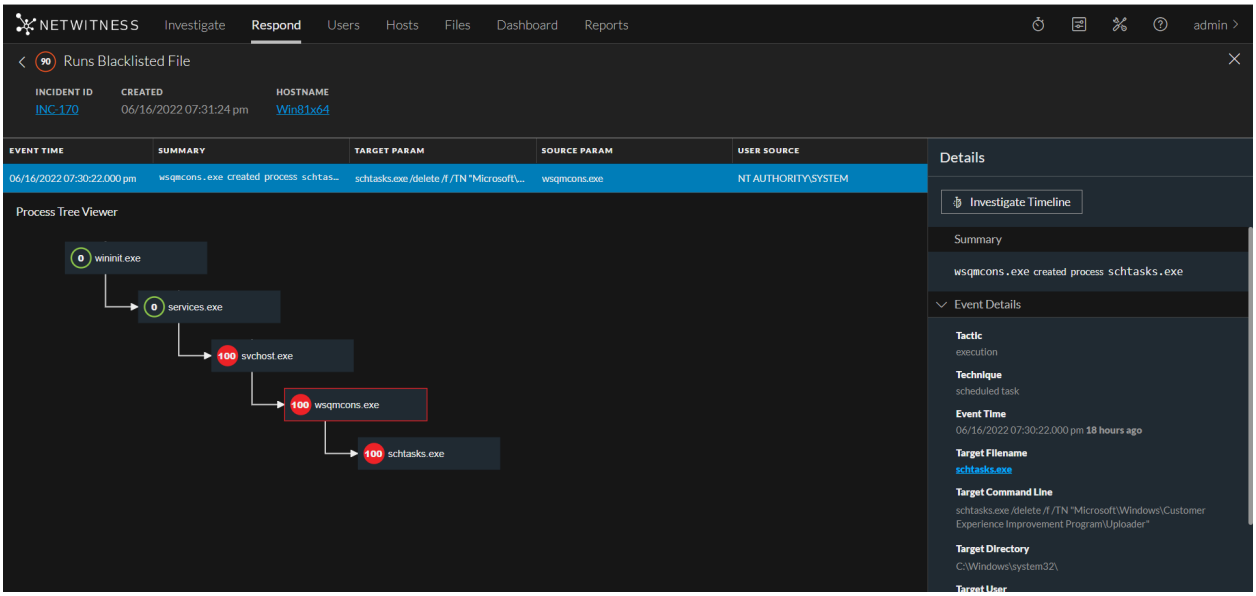


   The alerts are deleted from NetWitness. If a deleted alert is the only alert in an incident, the incident is also deleted. If the deleted alert is not the only alert in an incident, the incident is updated to reflect the deletion.

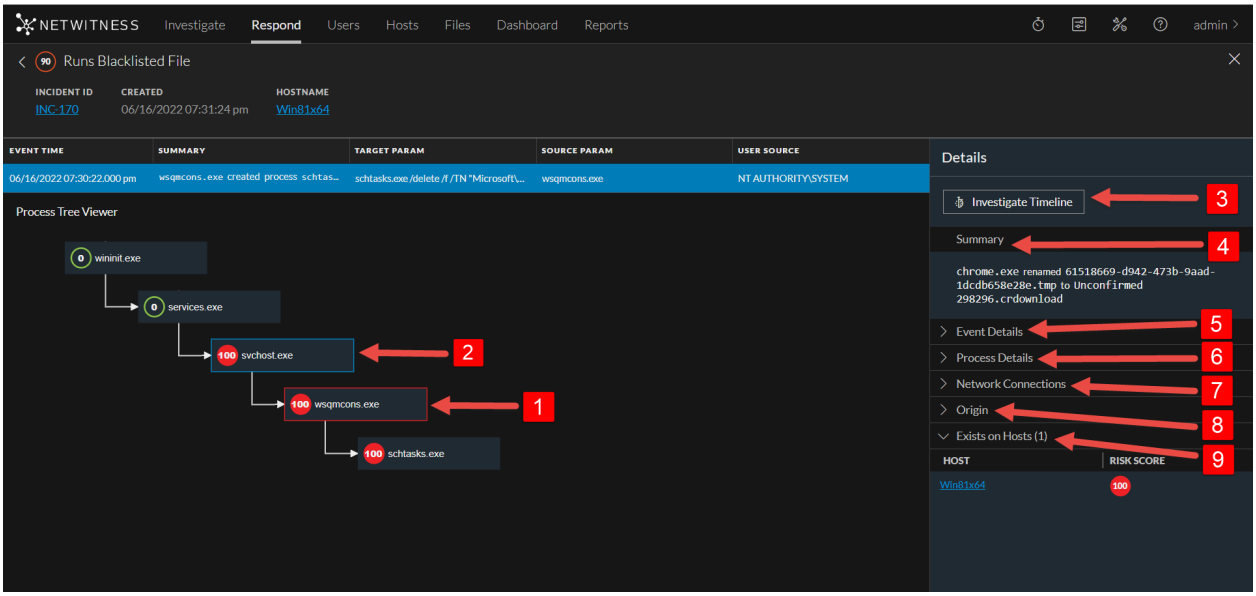
Review Endpoint Alerts using Process Tree

From version 12.0.0.0 and higher, the Alert details page for Endpoint alerts will show a process tree along with the details of Summary, Event details, Process details, etc.

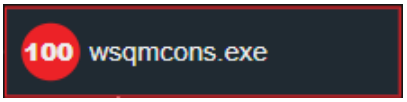
After you filter the Endpoint alerts in the Alerts List view, you can go to the Alert Details view for more detailed information on the Endpoint alerts, to determine the action required. An alert contains one or more events. In the Alert Details view for Endpoint alerts, you can view the alert details in the form of a process tree and additional event details, process details and much more on the right panel. The following figure shows an example of the Alert Details view for Endpoint alerts.



The process tree on the Alert Details view provides a complete picture about where the suspicious/malicious file originated including the path in the form of a process tree.



The **Details** panel on the right has more information for an alert than the Overview panel in the Alerts List view.



- The file that caused the alert is outlined in red.



- Selected file is outlined in blue.

1 - The file that caused the alert, and it is outlined in red. If you click on this file, the red outline will become blue to show it is selected.

2 - The file from which the suspicious/malicious file is originated.

3 - **Investigate Timeline** takes to the Investigate view for the selected alert.

4 - **Summary** shows a short description of the event.

5 - **Event Details** section provided a detailed information about the event that includes the Event Time, Target Filename, Tactic, Technique, Target User etc.

6 - **Process Details** section shows the Directory where the file is stored besides User name, Hash value, Risk score, Signature etc.

7 - **Network Connections** shows any network connection the selected file established since ten minutes before and till ten minutes after the alert triggered time. For example, if the alert was triggered at 16:00 hours, the network connections(if any)established by the selected file from 15:50 hours to 16:10 hours will be shown.

8 - **Origin** section shows how the selected file originated in the host.

9 - **Exists on Hosts** shows the list of hosts(with risk score) the selected file exists.

Process Details Section Values

Name	Description	Example
Tactic	Shows the tactic, as per MITRE ATT&CK frame work, this attempt falls under.	*execution*

Name	Description	Example
Technique	Shows the technique, as per MITRE ATT&CK frame work, this attempt falls under.	*masquerading*
Event Time	Shows the event occurred time.	*06/22/2022 10:14:28.000 am 8 hours ago*
Target Filename	Shows the name of file that is targeted. You can also view it in the process tree, next to the file that caused the alert.	*Unconfirmed 298296.crdownload*
Target Command Line	Shows the command line argument of the target file.	N/A
Target Directory	Shows the targeted directory.	*C:\Users\Administrator\Downloads*
Target User	Shows the user name through which the attempt was made.	*WIxxxxxx\Administrator*
Target Hash	Shows the hash value of the selected file.	*f214c48dc1daxxxx41d327c6bed1b52xxxx492573d85a305d8183eaa0222cc96*

Event Details Section Values

Value	Description	Example
File name	Shows the selected file name with extension	*iexplore.exe*
Command Line	Shows the command line name for the selected file	*IEXPLORE.EXE*
Directory	Shows the location of the selected file	*C:\Program Files\Internet Explorer*
User	Shows the user name	*WIxxxxxx\Administrator*

Hash	Shows the hash value of the selected file	*f214c48dc1daxxxx41d327c6bed1b52xxxx492573d85a305d8183eaa0222cc96*
Risk Score	Risk score of the selected file	*100*
Signature	Shows whether the selected file is signed or not	*microsoft,signed,valid*
Reputation Status	Shows the reputation of a file hash	*Suspicious*
File Status	Shows the file status for the selected file	*Blacklist*

Note: The process tree will be invisible if you drag it to the right end of the screen. Refresh the page to reload the process tree.

NetWitness Respond Reference Information

The Respond view user interface provides access to NetWitness Respond functions. This topic contains descriptions of the user interfaces as well as other reference information to help users understand the functions of NetWitness Respond.

Topics

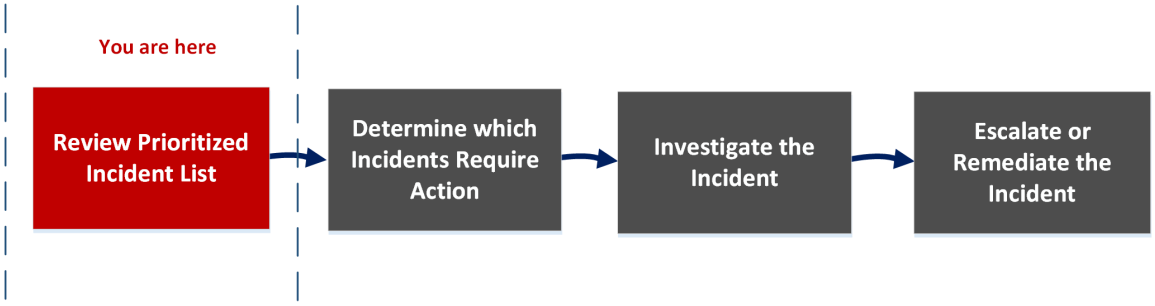
- Incidents List View
- Incident Details View
- Alerts List View
- Alert Details View
- Tasks List View
- Add/Remove from List Dialog
- Context Lookup Panel - Respond View

Incidents List View

The Incidents List view (Respond > Incidents) shows Incident Responders and other Analysts a prioritized results list of incidents created from various sources. For example, your results list could show incidents created from ESA rules or NetWitness Endpoint. From the Incidents List view, you have easy access to the information that you need to quickly triage and manage incidents through completion.

Workflow

This workflow shows the high-level process that Incident Responders use to respond to incidents in NetWitness.



In the Incidents List view, you can review the list of prioritized incidents, which shows basic information about each incident. You can also change the assignee, priority, and status of the incidents. Because the results can be large in the incidents list, you have the option to filter those incidents by time range, incident ID, custom date range, priority, status, assignee, and categories.

What do you want to do?

Role	I want to ...	Show me how
Incident Responders, Analysts, and SOC Manager	View prioritized incidents*	Review Prioritized Incident List
Incident Responders, Analysts, and SOC Manager	Filter and sort the incident list*	Filter the Incident List
Incident Responders, Analysts	View my incidents*	View My Incidents
Incident Responders, Analysts	Assign incidents to myself*	Assign Incidents to Myself
Incident Responders, Analysts, and SOC Manager	Find Incidents*	Find an Incident
Incident Responders, Analysts, and SOC Manager	Send an incident to Archer Cyber Incident & Breach Response or update an incident.*	Escalate or Remediate the Incident
Incident Responders, Analysts	View incident details.	Determine which Incidents Require Action
Incident Responders, Analysts	Further Investigate an incident.	Investigate the Incident
Incident Responders, Analysts, and SOC Manager	Create a task.	Escalate or Remediate the Incident

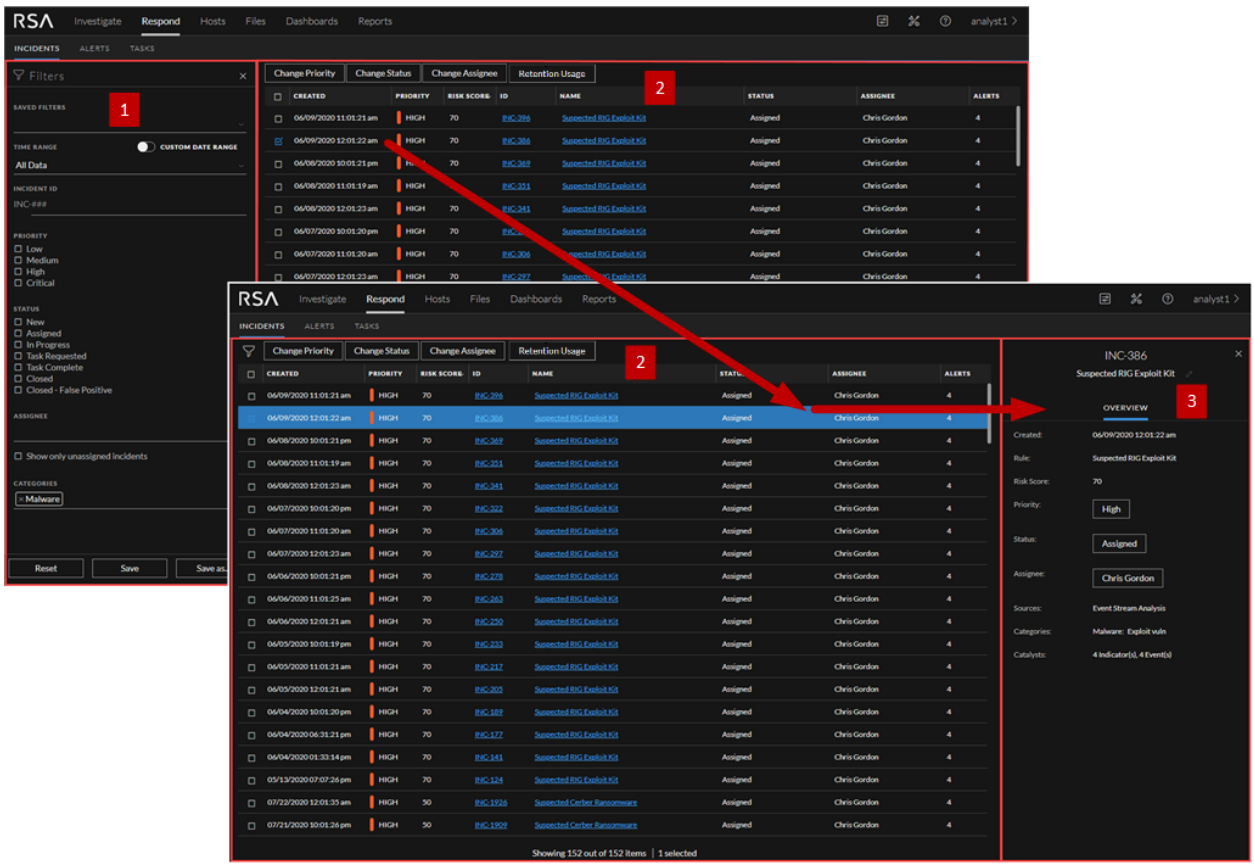
*You can complete these tasks here (that is, in the Incidents List view).

Related Topics

- Incident Details View
- Responding to Incidents

Quick Look

The following example shows the initial Incidents List view with the Filter panel. You can open the Overview panel for an incident by clicking an incident in the Incident List.



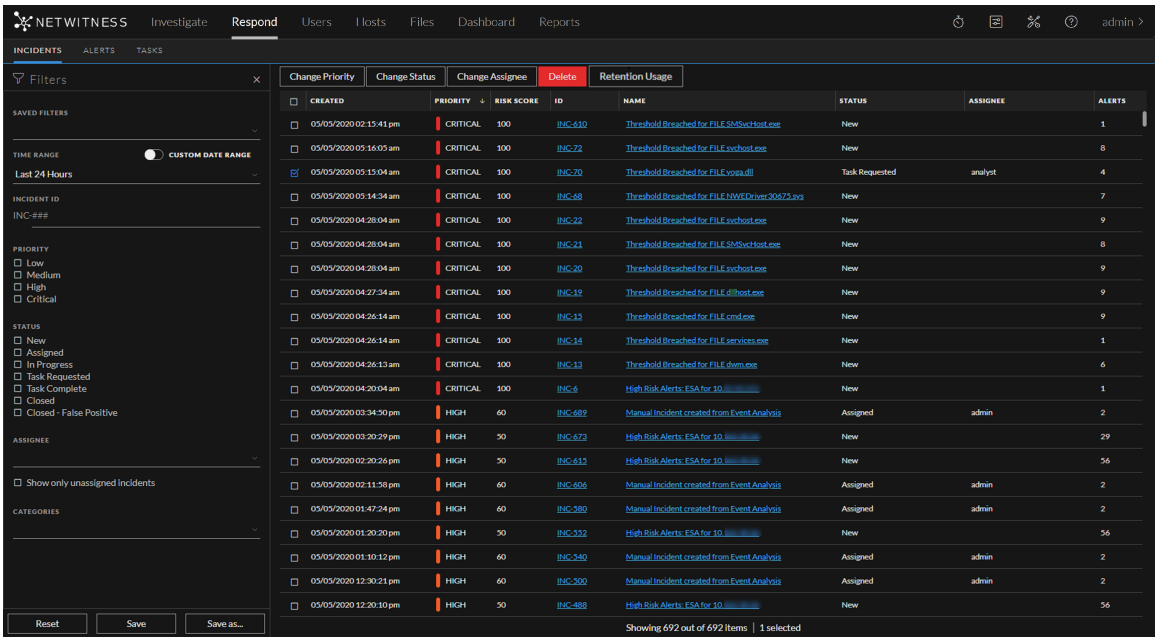
1 Filters Panel
2 Incidents List
3 Overview Panel

You can go directly to the Incident Details view from the Incidents List by clicking the hyperlinked ID or NAME. The Overview panel is also available in the Incident Details view. For more information about the Incidents Details view, see Incident Details View.

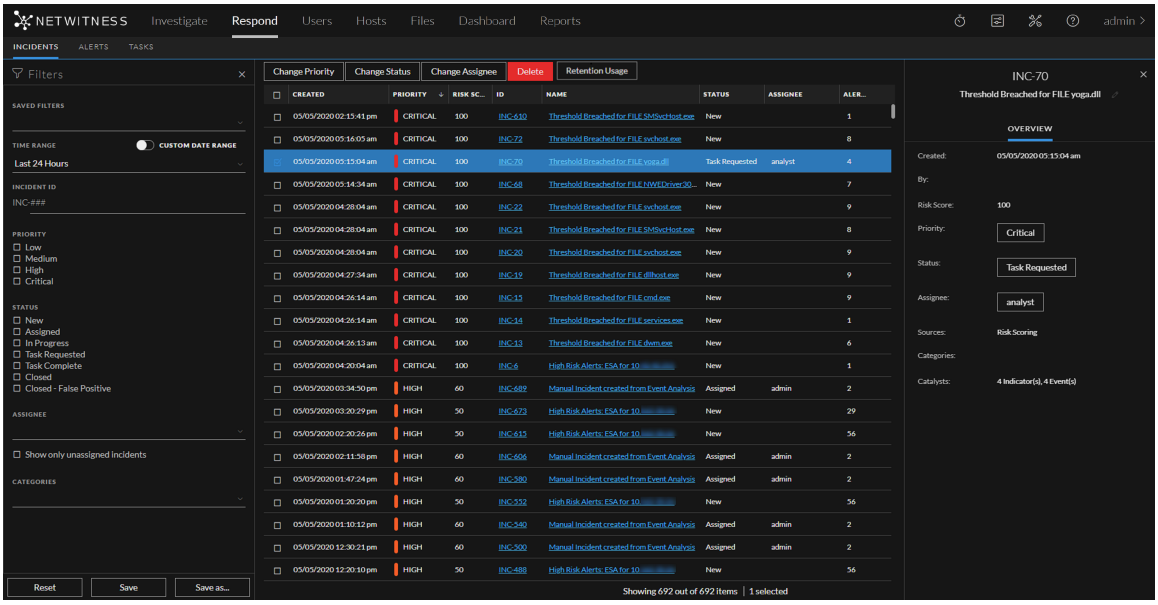
Incidents List View

To access the Incidents List view, go to **Respond > Incidents**. The Incidents List view displays a list of all incidents. The Incidents List view consists of a Filters panel, an Incidents List, and an Incidents Overview panel.

The following figure shows the Filter Panel on the left and the Incidents List on the right.



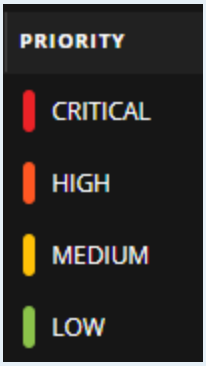
The following figure shows the incident Overview panel on the right.



Incidents List

The Incidents List shows a list of all of the prioritized incidents. You can filter this list to show only incidents of interest.

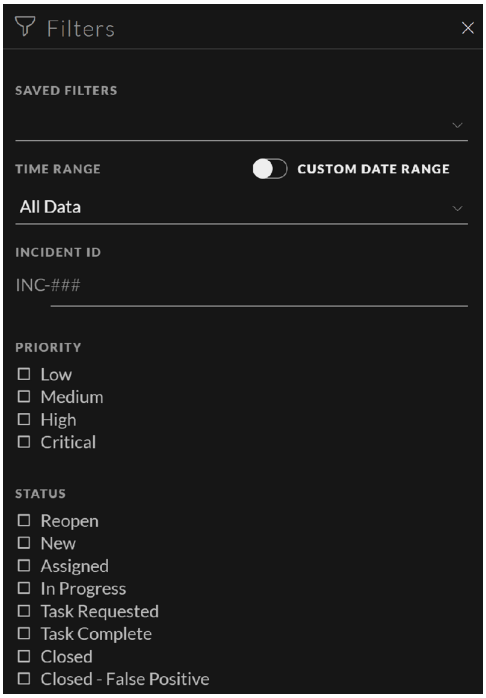
Column	Description
Created	Shows the creation date of the incident.

Column	Description
Priority	Shows the incident priority. Priority can be Critical, High, Medium, or Low. The Priority is color coded, where red indicates a **Critical** incident, orange represents a **High** risk incident, yellow indicates a **Medium** risk incident, and green represents a **Low** risk incident. For example: 
Risk Score	Shows the incident risk score. The risk score indicates the risk of the incident as calculated by an algorithm and is between 0-100. 100 is the highest risk score.
ID	Shows the automatically created incident number. Each incident is assigned a unique number that you can use to track the incident.
Name	Shows the incident name. The incident name is derived from the rule used to trigger the incident. Click the link to go to the Incident Details view for the selected incident.
Status	Shows the incident status. The status can be: Reopen, New, Assigned, In Progress, Task Requested, Task Complete, Closed, and Closed-False Positive.
Assignee	Shows the team member currently assigned to the incident.
Alerts	Shows the number of alerts associated with the incident. An incident may include many alerts. A large number of alerts might mean that you are experiencing a large-scale attack.

At the bottom of the list, you can see the number of incidents on the current page, the total number of incidents, and the number of incidents selected. For example: **Showing 1000 out of 2517 items | 2 selected**. The maximum number of incidents that you can view at one time is 1,000.

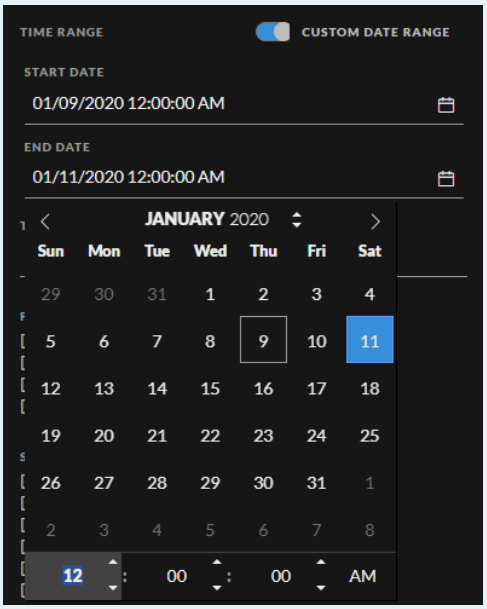
Incident Filters Panel

The following figure shows the filters available in the Filters panel.



The Filters panel, on the left of the Incidents List view, has options that you can use to filter the incidents list. When you navigate away from the Filters panel, the Incidents List view retains your filter selections.

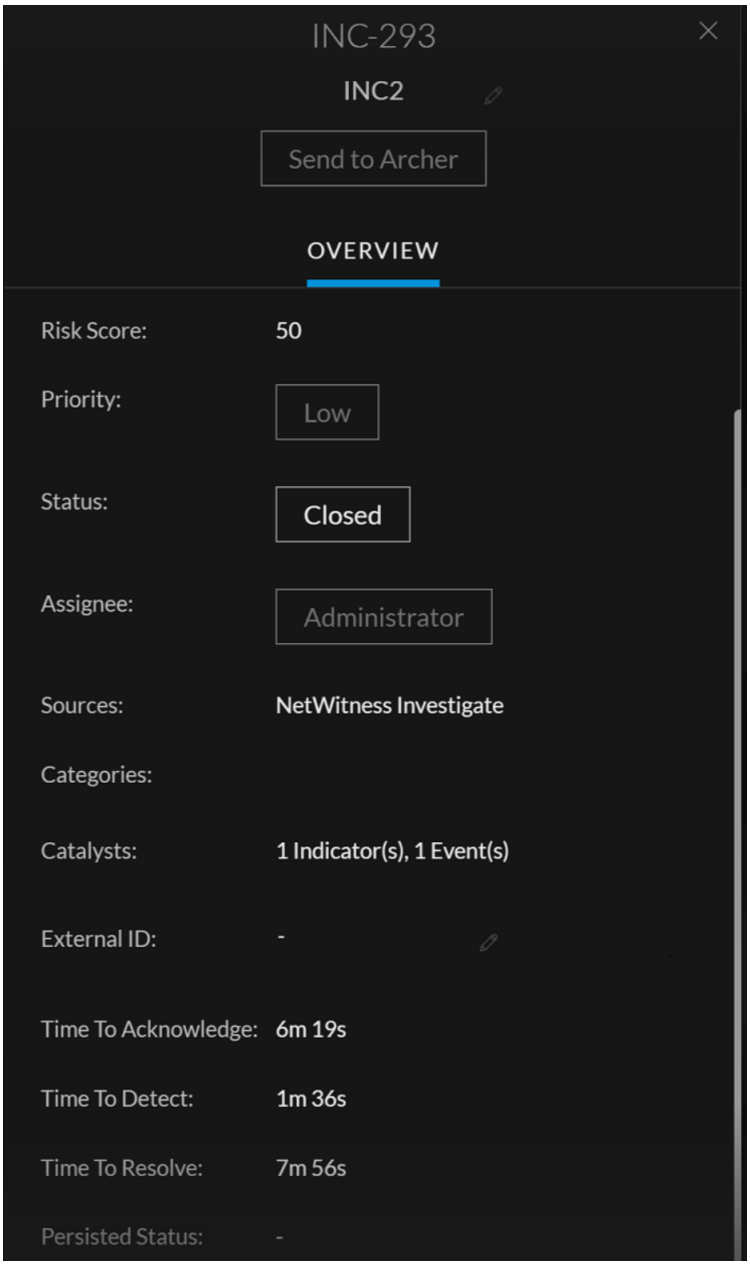
Option	Description
Saved Filters	You can select a saved filter to filter the incident list. Saved filters are global. You can save a filter for other analysts to use and you can use any saved filter. Saved filters are also available for use on the Springboard landing page. Filters used in the Springboard cannot be deleted. (This option is available in NetWitness Platform 11.5 and later.)
Time Range	You can select a specific time period from the Time Range drop-down list. The time range is based on the received date of the alerts. For example, if you select Last Hour, you can see alerts that were received within the last 60 minutes.

Option	Description
Custom Date Range	You can specify a specific date range instead of selecting a Time Range option. To do this, click the white circle in front of Custom Date Range to view the Start Date and End Date fields. Select the dates and times from the calendar. 
Inident ID	Type the number of the incident that you would like to locate. For example, for INC-1050, type only the number "1050" to view the incident.
Priority	Select the priorities that you would like to view.
Status	Select one or more incident statuses. For example, select Closed - False Positive to view only false positive incidents, which were initially identified as suspicious, but then they were later found to be safe.
Assignee	Select the assignee or assignees of the incidents that you would like to view. For example, if you only want to view the incidents assigned to Cale or Stanley, select Cale and Stanley from the Assignee drop-down list. If you want to view incidents regardless of the assignee, do not make a selection under Assignee. (Available in NetWitness Version 11.1 and later) To view only unassigned incidents, select **Show only unassigned incidents**.
Categories	Select one or more categories from the drop-down list. For example, if you only want to view incidents classified with the Backdoor or Privilege abuse categories, select Backdoor and Privilege abuse.
Sent to Archer	(In NetWitness Version 11.2 and later, if Archer is configured as a data source in Context Hub, you can send incidents to Archer Cyber Incident & Breach Response and this option will be available in NetWitness Respond.) To view incidents that were sent to Archer, select **Yes**. For incidents that were not sent to Archer, select **No**.

Option	Description
Reset	Removes your filter selections. If you reset filters on a saved filter, it takes you to the default empty filter.
Save	Saves the currently applied incidents filter or updates a saved filter. For a new filter, choose a unique name that contains 1-256 alphanumeric characters, underscores, or hyphens. (This option is available in NetWitness Platform 11.5 and later.)
Save As	Saves the currently applied incidents filter for future use. Choose a unique name that contains 1-256 alphanumeric characters, underscores, or hyphens. (This option is available in NetWitness Platform 11.5 and later.)

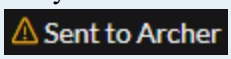
Incident Overview Panel

The Overview panel shows basic summary information about a selected incident. From the Incidents List, you can click an incident to access the Overview panel. The Overview panel in the Incident Details view contains the same information.



The following table lists the fields displayed in the Incident Overview panel.

Field	Description
<Incident ID>	Displays the Incident ID.

Field	Description
Send to Archer / Sent to Archer	(In NetWitness Version 11.2 and later, if Archer is configured as a data source in Context Hub, you can escalate incidents to Archer Cyber Incident & Breach Response and this option will be available in NetWitness Respond.) Shows whether the incident was sent to Archer Cyber Incident & Breach Response: • **Send to Archer**: The incident was not sent to Archer. You can click the **Send to Archer** button to send the incident to Archer Cyber Incident & Breach Response for additional processing. This action is not reversible.  • **Sent to Archer**: The incident was sent to Archer Cyber Incident & Breach Response for additional analysis and action. 
<Incident Name>	Displays the name of the incident. You can click the incident name to change it. For example, rules can create many incidents with the same name. You can change the incident names to be more specific.
Created	Shows the creation date and time of the incident.
Rule / By	Shows the name of the rule that created the incident or the name of the person who created the incident.
RiskScore	Shows a value between 0 and 100 that indicates the risk of the incident as calculated by an algorithm. 100 is the highest risk score.
Priority	Shows the incident priority. Priority can be Critical, High, Medium or Low. To change the priority, you can click the Priority button and select a new priority from the drop-down list.
Status	Shows the incident status. The status can be Reopen, New, Assigned, In Progress, Task Requested, Task Complete, Closed, and Closed - False Positive. To change the status, you can click the Status button and select a new status from the drop-down list.
Assignee	Shows the team member currently assigned to the incident. To change the assignee you can click the Assignee button and select a new assignee from the drop-down list.

Field	Description
Sources	Displays the data sources used to locate the suspicious activity.
Categories	Displays the categories of the incident events.
Catalysts	Displays the count of indicators that gave rise to the incident.
External ID	Allows storing the Incident ID referrals from a different platform. **Note:** Click **Send to Archer** to generate the External ID. The ID generated is automatically stored as External ID.
Time to Acknowledge	Displays the time taken to assign an Incident after creating it.
Time to Detect	Displays the time taken for completing the task after the Incident is assigned.
Time to Resolve	Displays the time taken for closing the task after the Incident is created.
Persisted Status	Displays the persist status of the Incident. The status can be Complete, Partial, or None (-).

Toolbar Actions

This table lists the toolbar actions available in the Incidents List view.

Option	Description
	Enables you to open the Filters panel so that you can specify the incidents that you would like to see in the Incidents List.
	Closes the panel.
Change Priority button	Allows you to change the Priority of one or more selected incidents in the Incidents List.
Change Status button	Allows you to change the Status of one or more selected incidents.
Change Assignee button	Allows you to change the Assignee of one or more selected incidents.
Delete button	Allows you to delete the selected incidents if you have the appropriate permissions, such as an Administrator or Data Privacy Officer.

Option	Description
Retention Usage button	Allows an analyst to fetch all the stats of all the configured services and the percentage used by the pinned cache directories.

Incident Details View

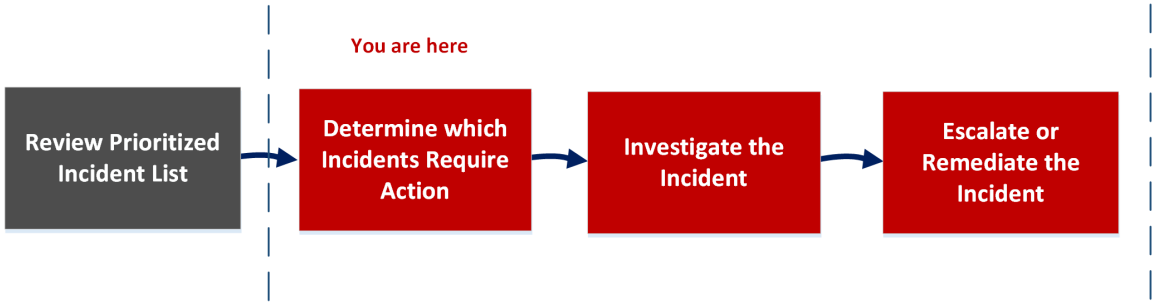
In the Incident Details view (Respond > Incidents > click an ID or NAME hyperlink in the Incidents List), you can view and access extensive incident details. The Incident Details view contains multiple panels that provide the following benefits:

- **Overview**: View an incident summary and update the incident.
- **Indicators**: View the indicators (alerts) involved in the incident, the events within those alerts, and available enrichment information. You can also access Event Analysis details for some events and perform event reconnaissance.
- **Related Indicators**: View indicators (alerts) that are related to the incident and add them to the incident if they are not associated with an incident.
- **History**: View all the actions performed by the user on any incident.
- **Nodal Graph**: Visualize the size and interactions between entities (IP address, MAC address, user, host, domain, file name, or file hash).
- **Events List**: Study the events associated with the incident.
- **Journal**: Add notes and collaborate with other analysts.
- **Tasks**: Create incident tasks and track them to closure.

You can also filter the data in the Incident Details view to study indicators and entities of interest.

Workflow

This workflow shows the high-level process that Incident Responders use to respond to incidents in NetWitness.



In the Incident Details view, you can use the extensive information provided about the incidents to determine which incidents require action. You also have the tools and information to investigate the incident, and then escalate or remediate it.

What do you want to do?

Role	I want to ...	Show me how
Incident Responders, Analysts, and SOC Manager	View prioritized incidents, filter and sort the incident list, find incidents, view my incidents, and assign incidents to myself.	Review Prioritized Incident List
Incident Responders, Analysts	View incident details.*	View Incident Details
Incident Responders, Analysts	View alerts and enrichments.*	View the Indicators and Enrichments
Incident Responders, Analysts	View events.*	View and Study the Events
Incident Responders, Analysts (Additional permissions required)	View event analysis for an event.*	View Event Analysis Details for Indicators
Incident Responders, Analysts	View a graph of the entities involved in the events.*	View and Study the Entities Involved in the Events on the Nodal Graph
Incident Responders, Analysts	Filter the incident data.*	Filter the Data in the Incident Details View
Incident Responders, Analysts	View and add incident notes.*	View Incident Notes and Document Steps Taken Outside of NetWitness
Incident Responders, Analysts	View and create tasks.*	View the Tasks Associated with an Incident and Create a Task
Incident Responders, Analysts	Add related alerts and add them to the incident.*	Find Related Indicators and Add Related Indicators to the Incident
Incident Responders, Analysts	View contextual information about an incident from Context Hub.*	View Contextual Information
Incident Responders, Analysts	Reduce false positives by adding an entity to a whitelist.*	Add an Entity to a Whitelist
Incident Responders, Analysts	Pivot to NetWitness Investigate.*	Pivot to the Investigate > Events View
Incident Responders, Analysts	Pivot to NetWitness Endpoint.*	Pivot to NetWitness Endpoint Thick Client
Incident Responders, Analysts, and SOC Manager	Send an incident to Archer Cyber Incident & Breach Response.*	Send an Incident to Archer

Role	I want to ...	Show me how
Incident Responders, Analysts	Update or close an incident.*	Update an Incident and Close an Incident
Incident Responders, Analysts, and SOC Manager	View all tasks.	Escalate or Remediate the Incident
Incident Responders, Analysts, and SOC Manager	Bulk update incidents and tasks.	Escalate or Remediate the Incident

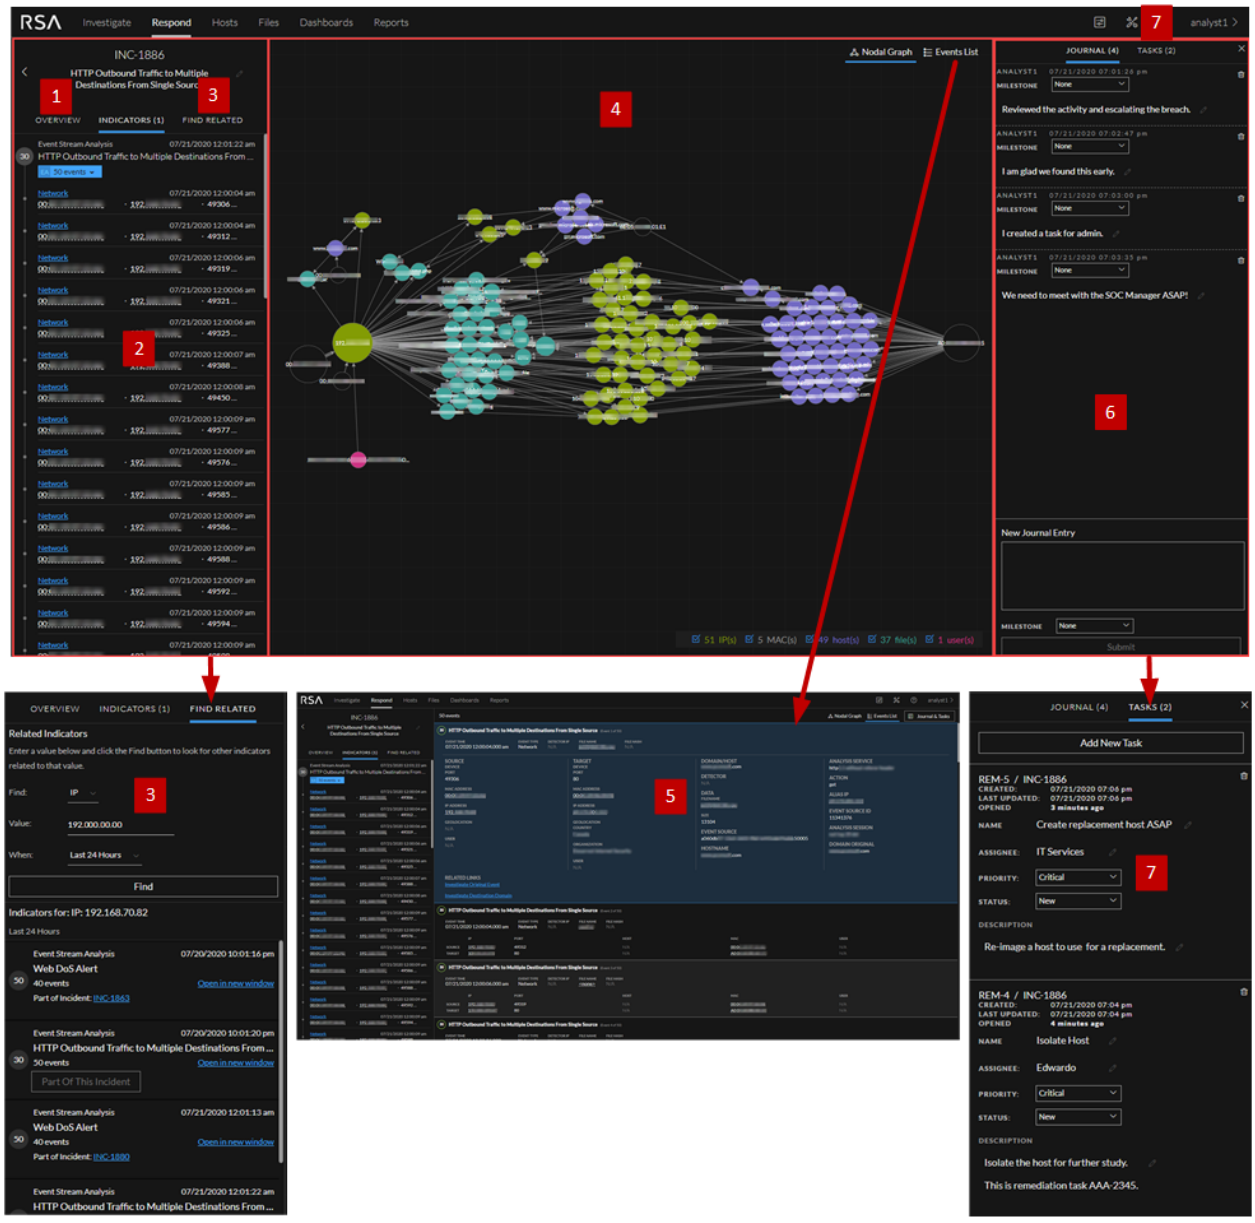
*You can complete these tasks here (that is, in the Incident Details view).

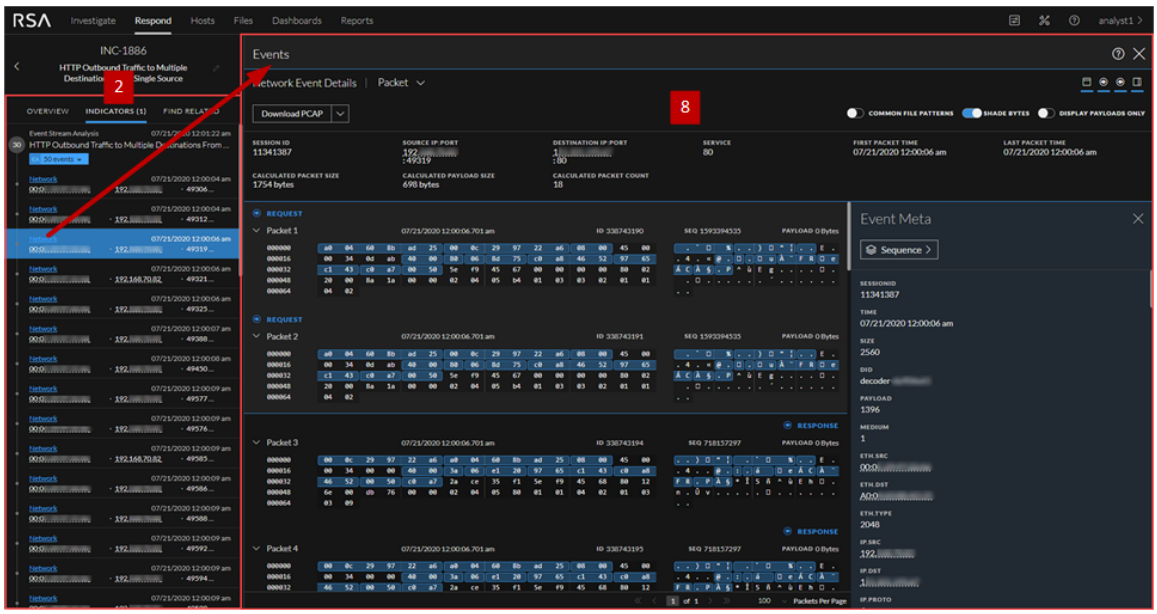
Related Topics

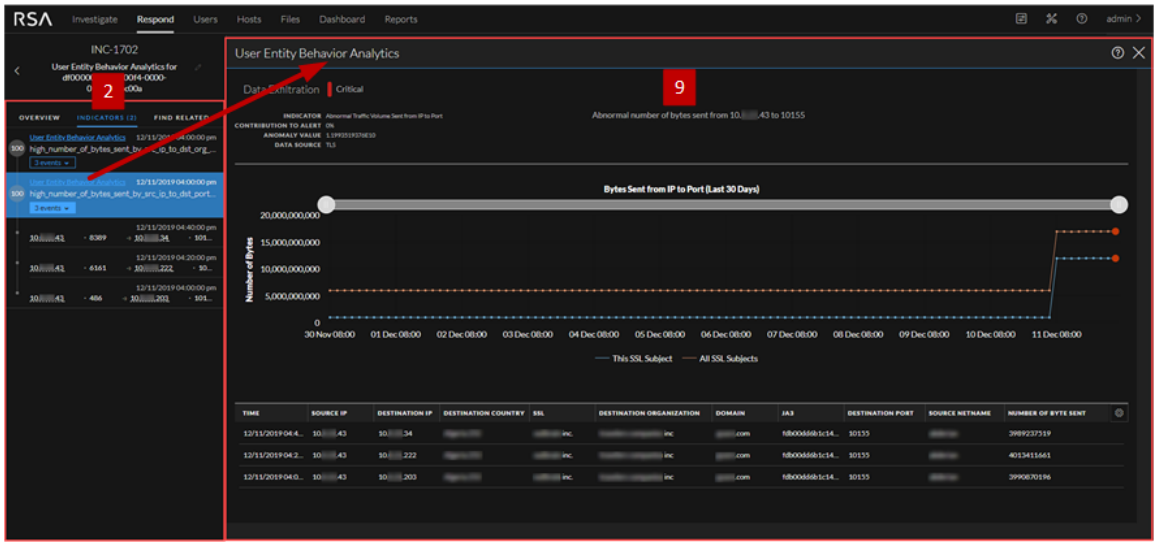
- Incidents List View
- Determine which Incidents Require Action
- Investigate the Incident
- Escalate or Remediate the Incident

Quick Look

The following example shows the locations of the Incident Details view panels.




RSA
Investigate
Respond
Hosts
Files
Dashboards
Reports
INC-1886
HTTP Outbound Traffic to Multiple Destinations from Single Source
2
Events
Network Event Details
Packet
8
Download PCAP
Event Meta
Sequence
Packet 1
Packet 2
Packet 3
Packet 4
REQUEST
RESPONSE


RSA
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
INC-1702
User Entity Behavior Analytics
2
Data Exfiltration
Critical
9
Abnormal number of bytes sent from 10.43 to 10155
Bytes Sent from IP to Port (Last 30 Days)
Number of Bytes
20,000,000,000
15,000,000,000
10,000,000,000
5,000,000,000
0
30 Nov 08:00
01 Dec 08:00
02 Dec 08:00
03 Dec 08:00
04 Dec 08:00
05 Dec 08:00
06 Dec 08:00
07 Dec 08:00
08 Dec 08:00
09 Dec 08:00
10 Dec 08:00
11 Dec 08:00
This SSL Subject
All SSL Subjects



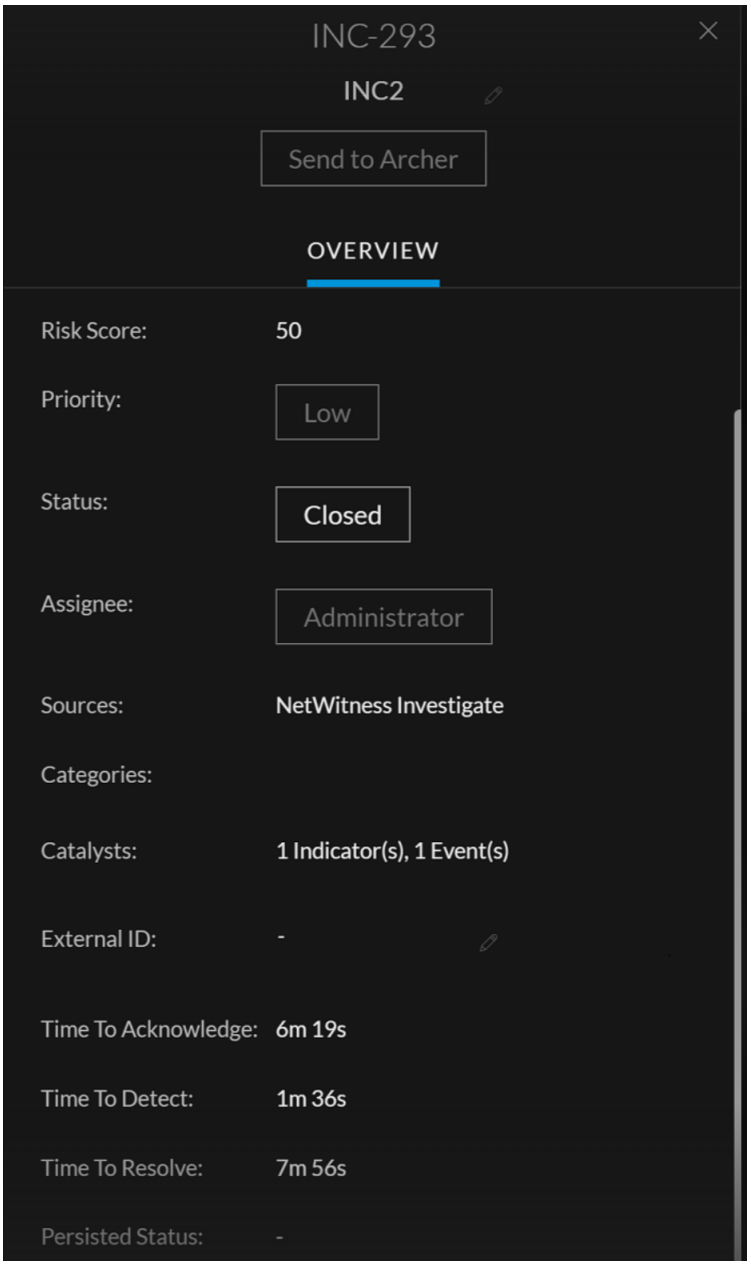
1. Overview (Click the Overview tab to view the Overview panel.)
2. Indicators Panel
3. Related Indicators Panel (Click the Find Related tab to view it.)
4. Nodal Graph
5. Events List (Click the top of an event in the Events List to view event details.)
6. Journal Panel
7. Tasks Panel (Click the Tasks tab to view it.)
8. Events (Click an event type hyperlink in the Indicators panel, such as Network, to view the Events view from Investigate for a specific indicator event.)
9. UEBA (Click a User Entity Behavior Analytics hyperlink in the Indicators panel to view UEBA.)
10. History Panel

Note: Your Incident Details view may not look like these diagrams because the layout changed in NetWitness 11.3.2 and later versions.
The Related tab is renamed as the Find Related tab and is located on the left-side panel.
The journal is open by default on the right-side panel. When the journal is closed, the Journal & Tasks button enables easy access to notes and tasks.

Overview Panel

The Overview panel shows basic summary information about a selected incident. It also allows you to change the incident name and update the incident priority, status, and assignee. The Overview panel in the Incidents List view contains the same information. The Incidents List view Incident Overview Panel topic provides details.

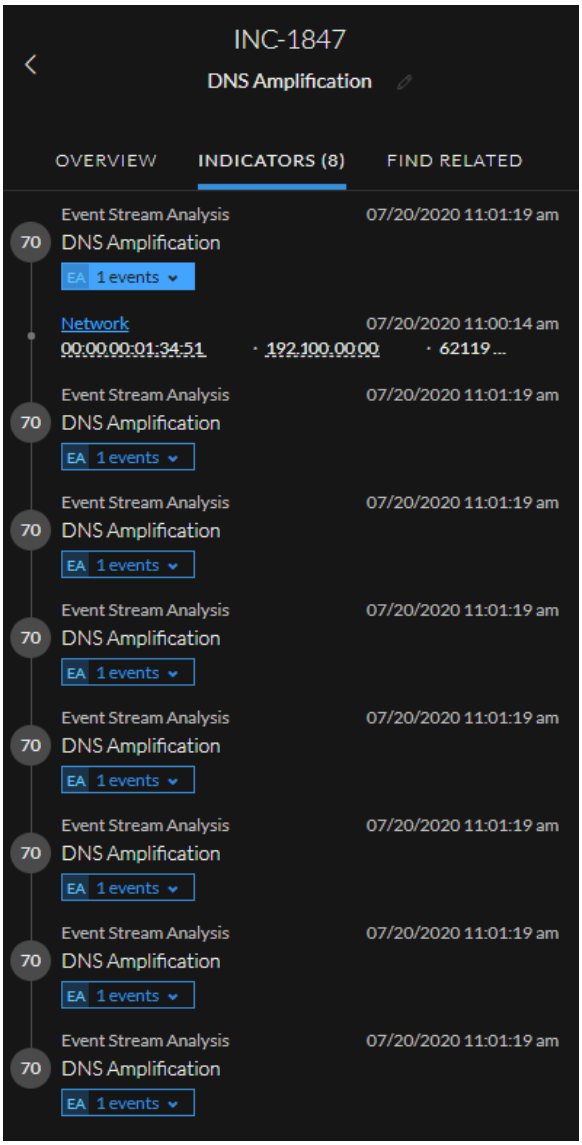
To view the Overview panel in the Incident Details view, click the **Overview** tab in the left panel.


INC-293
INC2
Send to Archer
OVERVIEW
Risk Score: 50
Priority: Low
Status: Closed
Assignee: Administrator
Sources: NetWitness Investigate
Categories:
Catalysts: 1 Indicator(s), 1 Event(s)
External ID: -
Time To Acknowledge: 6m 19s
Time To Detect: 1m 36s
Time To Resolve: 7m 56s
Persisted Status: -

Indicators Panel

The Indicators panel contains a chronological listing of indicators. *Indicators* are alerts, such as an ESA alert or a NetWitness Endpoint alert. (This is different than a timeline, which provides a visual representation of the timing of the events in the incident). This listing helps you to connect indicators and notable data. For example, an IP address connected to a command and communication ESA alert might also have triggered a NetWitness Endpoint alert or other suspicious activities.

To view the Indicators panel, in the left panel of the Incident Details view, click the **Indicators** tab.



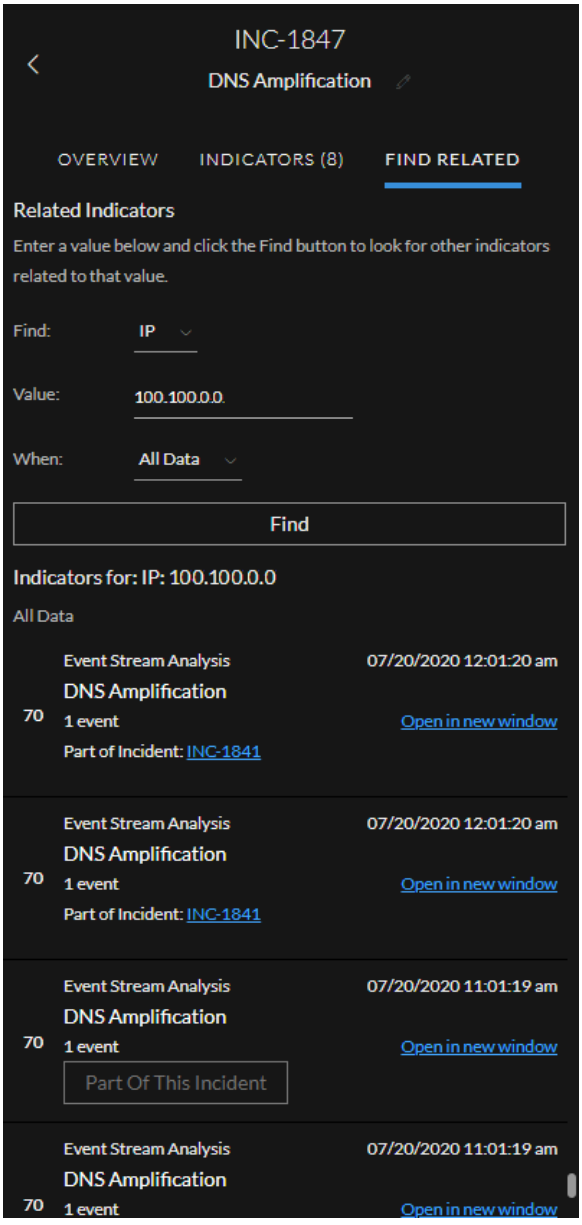
Data source information is shown below the names of the indicators. You can also see the creation date and time of the indicator and the number of events in the indicator. In the Indicators panel, you can drill deeper into the events associated with the listed indicators to get a better understanding of the events.

Note: The maximum number of indicators (alerts) displayed in the Indicators panel is 1,000.

Related Indicators Panel

The Related Indicators panel enables you to search the NetWitness alerts database to find alerts that are related to this incident. You can add alerts that you find to the incident if they are not already associated with an incident.

To view the Related Indicators panel, in the left panel of the Incident Details view, click the **Find Related** tab.



The following table describes the fields in the search section at the top of the panel.

Field	Description
Find	Select the entity that you would like to locate in the alerts. For example, IP.
Value	Type the value of the entity. For example, type the actual IP address of the entity.
When	Select a time range to search for the alerts. For example, Last 24 hours.
Find button	Initiates the search. A list of related indicators appear below the **Find** button in the **Indicators for** section.

The following table describes the options in the **Indicators for** (results) section at the bottom of the panel.

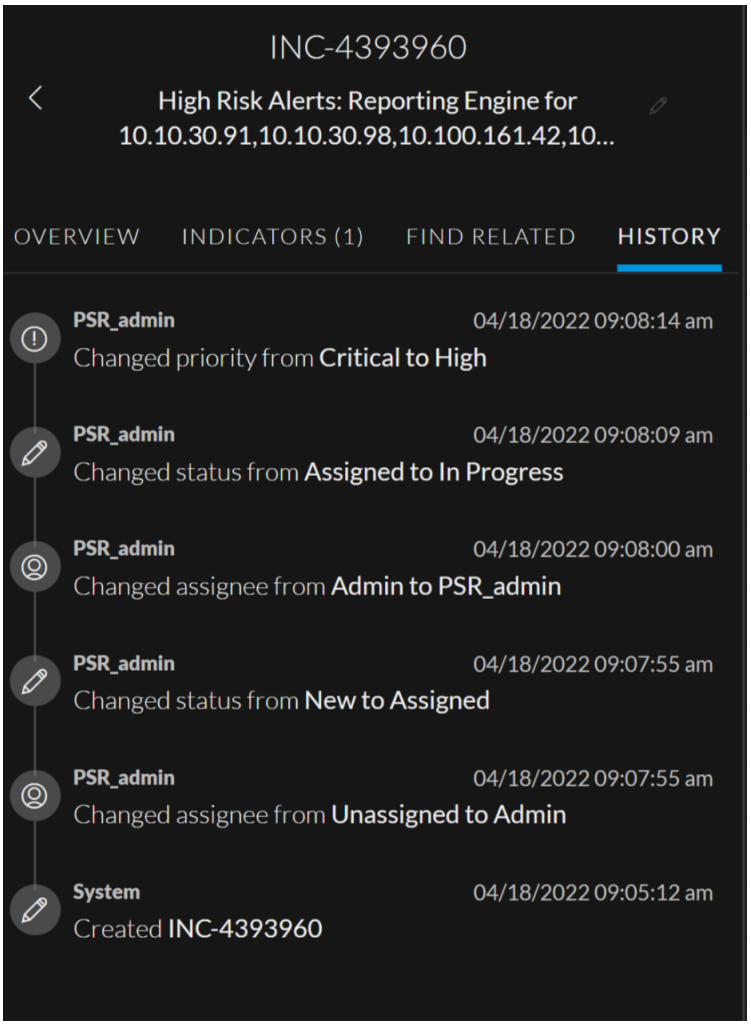
Option	Description
Indicators For:	Shows the search results.
Open in new window link	Shows alert details for the indicator.
Add To Incident button	Adds the related indicator to the incident. The related indicator adds to the Indicators panel.
Part Of This Incident button	Shows that the indicator is already part of the incident.

History Panel

The History panel displays every action performed by the user on an incident. The various actions performed on an incident are as shown below

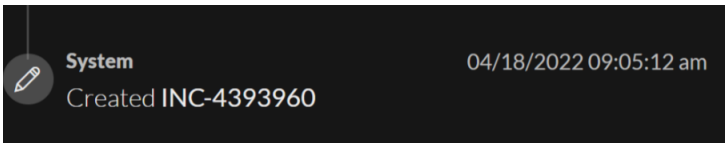
- Incident Assignee Change
- Incident Status Change
- Incident Priority Change
- Incident Creation

Every time a user performs an action on an incident, the date and time also gets recorded and is displayed in the panel. Consider the following example



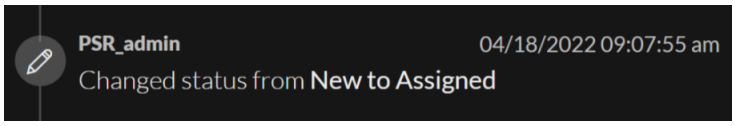
The different actions performed by the user are described below

- In this example, the Incident INC-4393960 was created by the user (System) on 18/04/2022 at 09:05:12 am.

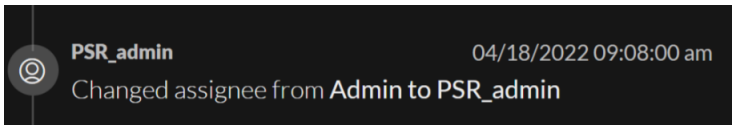


- After the incident creation, PSR_admin assigned the incident to Admin on 18/04/2022 at 09:07:55 am. Hence, the status of the incident is changed from New to Assigned.

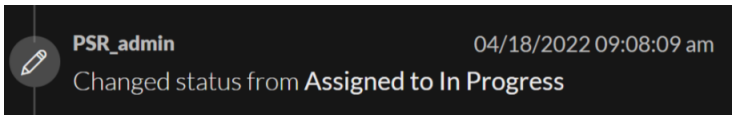




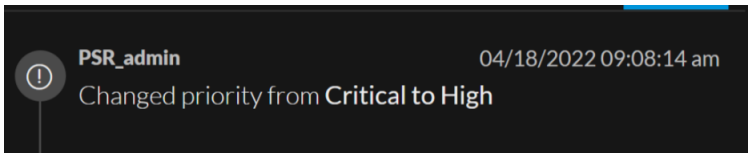
- Later, PSR_admin changed the Incident assignee from Admin to PSR_admin on 18/04/2022 at 09:08:00 am.



- After changing the assignee, PSR_admin changed the Incident status from Assigned to In Progress on 18/04/2022 at 09:08:09 am.



- Later, the Incident priority was changed from Critical to High on 18/04/2022 at 09:08:14 am by PSR_admin.

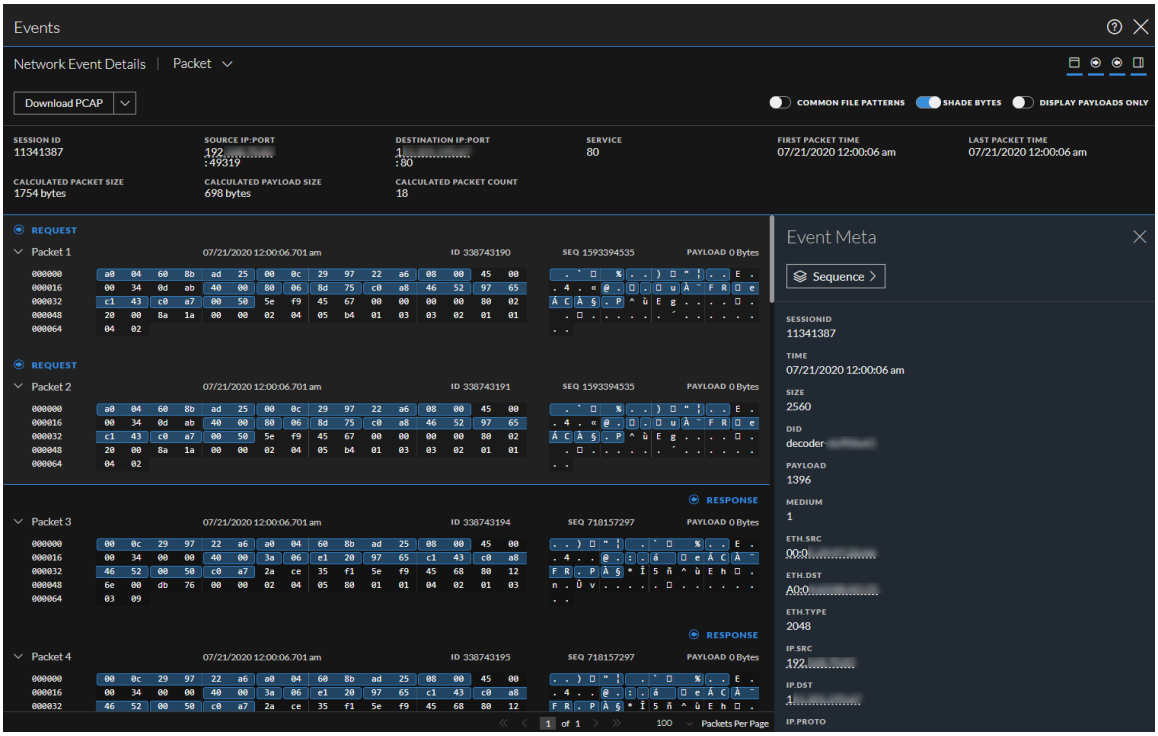


Events

You can perform an event analysis from the Indicators panel. Event counts preceded by an **EA** (event analysis) icon have event reconnaissance information available: 

. You can select an event type hyperlink, such as Network, to access the Events panel for the selected event.

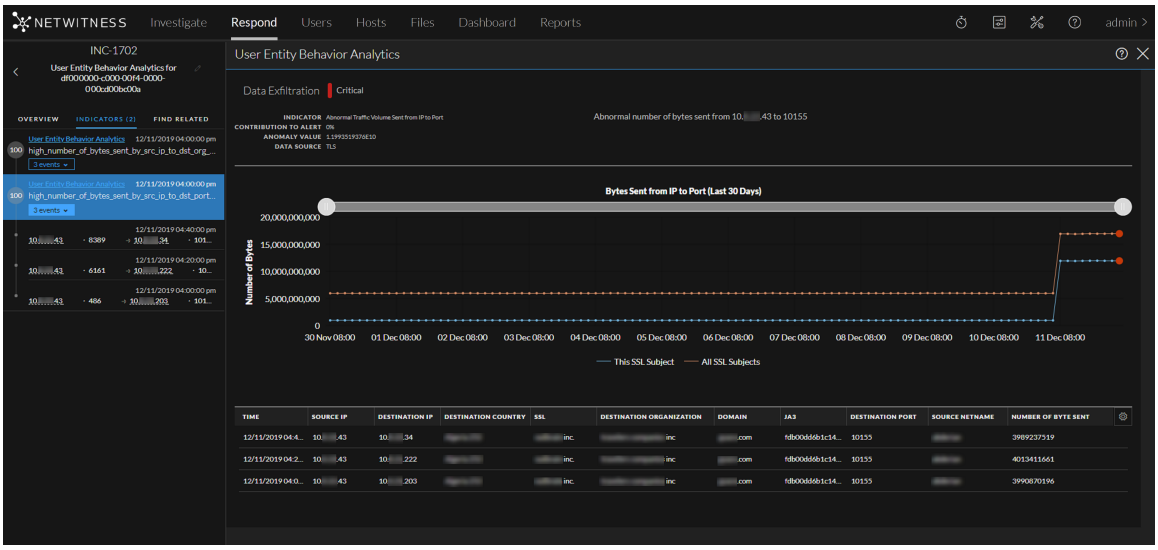
In the Events panel, you can view raw events and metadata with interactive features that enhance your ability to find meaningful patterns in the data. You can examine network, log, and endpoint events. The Events panel in the Respond view shows the Events view from Investigate for specific indicator events. For detailed information about the Events view, see the *NetWitness Investigate User Guide*.



Note: Migrated incidents from NetWitness versions before 11.2 will not show the Events panel in the Respond Incident Details view Indicators panel. Likewise, if you use alerts that were migrated from versions before 11.2 to create incidents in 11.5, you will also not be able to view the Events panel in the Respond view for those incidents.

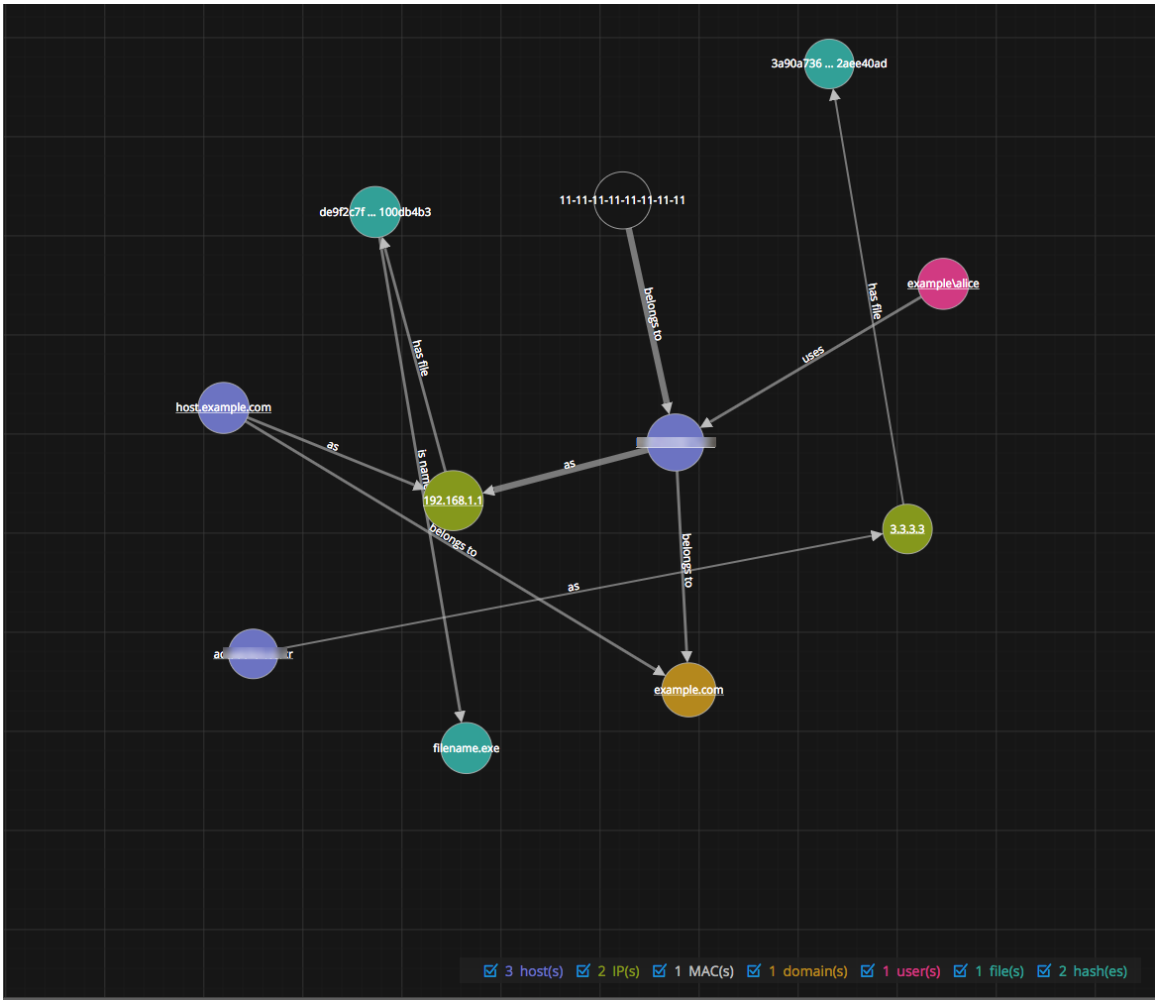
User Entity Behavior Analytics

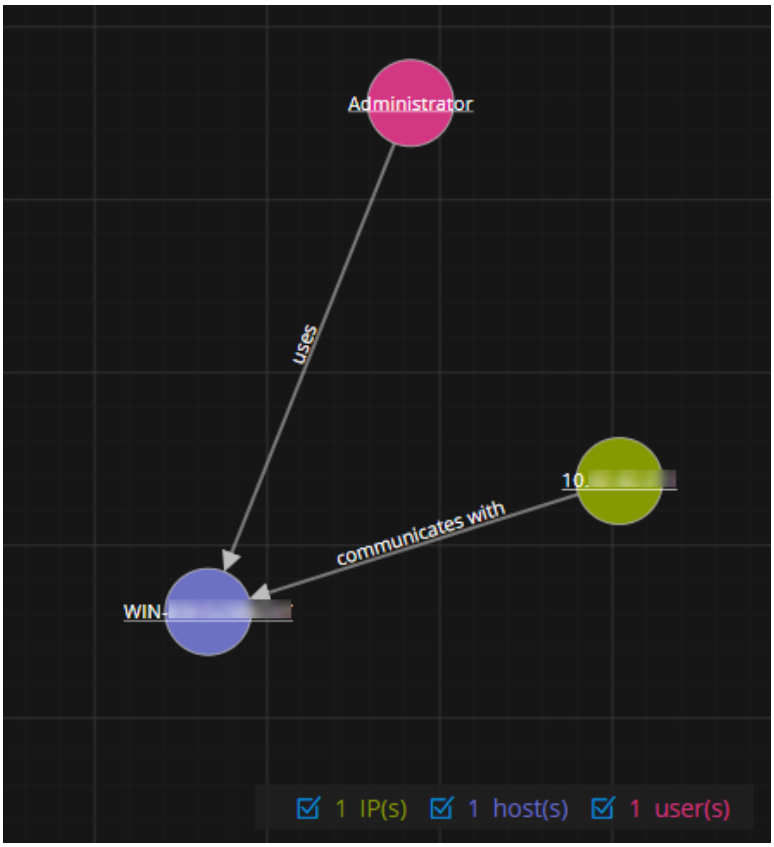
NetWitness UEBA (User and Entity Behavior Analytics) is an advanced analytics solution for discovering, investigating, and monitoring risky behaviors across all users and entities in your network environment. You can access UEBA from the Respond Incident Details view Indicators panel. Indicators with a **User Entity Behavior Analytics** hyperlink have additional UEBA information available. For detailed information about UEBA, see the *NetWitness UEBA User Guide*.



Nodal Graph

The nodal graph is an interactive graph that shows the entities involved in the incident. An *Entity* is represented by an IP address, MAC address, user, host, domain, file name, or file hash.





Nodes

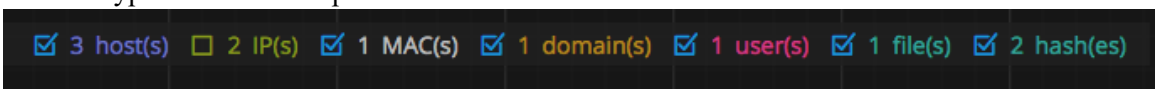
In the nodal graph, circles represent nodes. The following table describes the nodal graph node types.

Node	Description
IP address	If the event is a detected anomaly, you can see a Detector IP. If the event is a transaction, you can see a Destination IP and a Source IP.
MAC address	You may see a MAC address for each type of IP address.
User	If the machine is associated with a user, you can see a user node.
Host	A host can be physical equipment or a virtual machine, designated by a Fully Qualified Domain Name (FQDN) or IP address, on which any service is installed.
Domain	If a host is associated with a domain, you can see a domain node.
Filename	If the event involves files, you can see a filename.
File Hash	If the event involves files, you may see a file hash.

The legend at the bottom of the nodal graph shows the number of nodes of each type and the color coding of the nodes. It also helps you to locate the entities when the values, such as the IP addresses, are hashed.

You can click any node and drag it to reposition it.

In NetWitness Version 11.2 and later, you can select the node types that you want to view by clearing or selecting the checkboxes in the legend. The following figure shows an example nodal graph legend with all node types selected except IP.



Arrows

The arrows between the nodes provide additional information about the entity relationships. The following table describes the nodal graph arrow types.

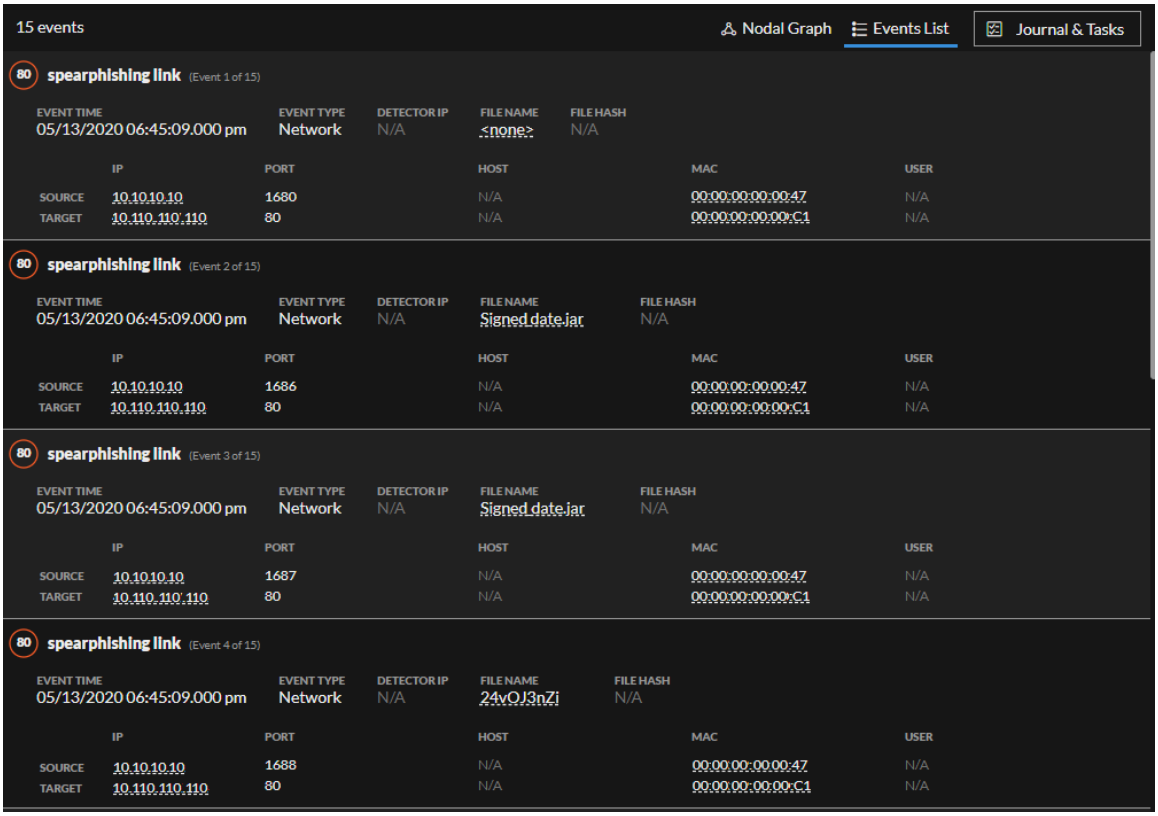
Arrow	Description
Communicates with	An arrow between a Source machine node (IP address or MAC address) and a Destination machine node labeled with "communicates with" shows the direction of the communication.
Has file	An Arrow between a machine node (IP address, MAC address, or Host) and a file hash node labeled with "has file " indicates that the IP address has that file.
Uses	An arrow between a User node and a machine node (IP address, MAC address, or Host) labeled with "uses" shows the machine that the user was using during the event.
Calls	(This arrow is available in NetWitness Platform 11.4 and later.) An arrow between two file hash (checksum) nodes labeled with "calls" indicates the direction of the interaction between the associated files. The source file hash "calls" the target (destination) file hash, which indicates that the source file associated with the source file hash is performing an action on the target file associated with the target file hash.
As	(This relationship type represents attributes of the connected node.) An arrow between nodes labeled with "as" provides additional information about the IP address that the arrow points to. For example, if there is an arrow from the host node circle that points to an IP address node that is labeled with "as", it indicates that the name on the host node circle is the hostname of that IP address and is not a different entity.
Is named	(This relationship type represents attributes of the connected node.) An arrow from a File Hash node to a File Name node labeled with "is named" indicates that the file hash corresponds to a file with that name.
Belongs to	(This relationship type represents attributes of the connected node.) An arrow between two nodes labeled with "belongs to" indicates that they pertain to the same node. For example, an arrow between a MAC address and a Host labeled with "belongs to" indicates that it is the MAC address of the host.

Larger line size arrows indicate more communication between the nodes. Larger nodes (circles) indicate more activity than smaller nodes. The larger nodes are the most common entities mentioned in the events.

Events List

The Events List shows the events associated with the incident. It shows information about the events, such as event time, source IP, destination IP, detector IP, source user, target user, and file information about the events. The amount of information listed depends on the event type. The maximum number of events displayed in the Events List is 1,000.

The following figure shows an Events List for network events.

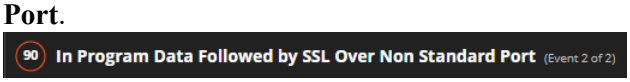


Note: The EVENT TIME displayed on this screen is the same as the COLLECTION TIME from the investigation page.

Each event has a header row with the following information:

- **Risk score**: This is the risk score of the indicator (alert) that contains the event.
- **Title**: This is the name of the event.
- **(Event x of x)**: This indicates the number of the event out of the total number of events in the indicator.

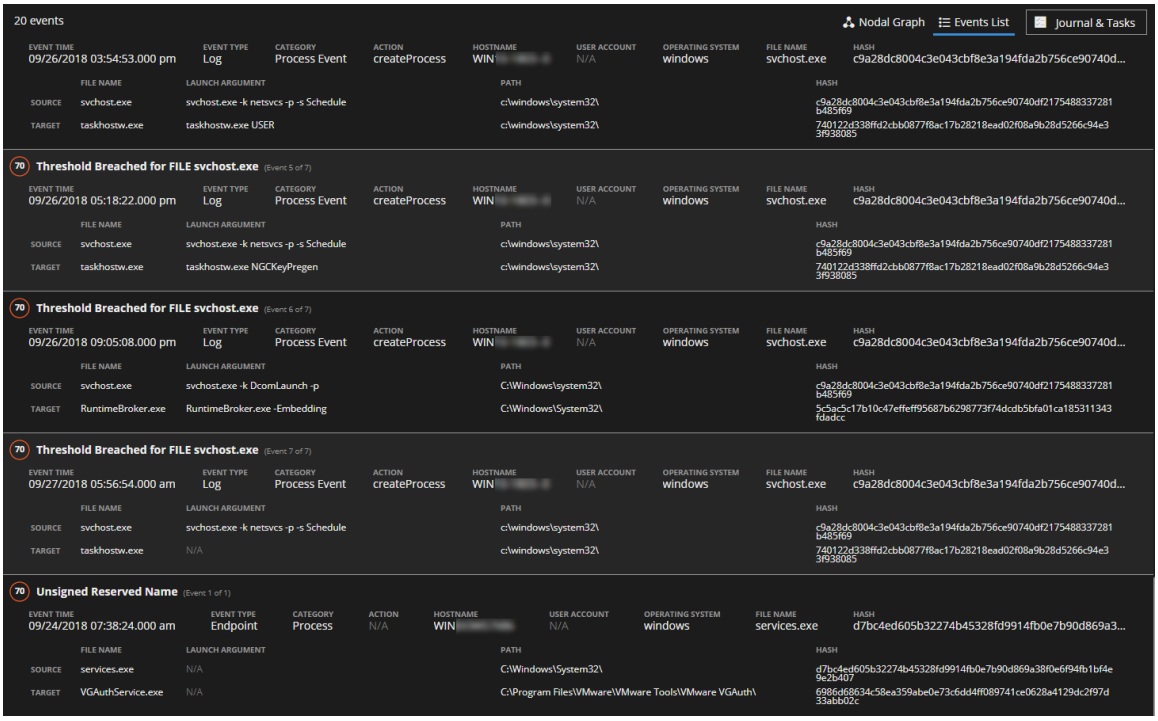
For example, the following event header shows that this event is event 2 of 2 for an indicator (alert) that has a risk score of 90. The event name is **In Program Data Followed by SSL Over Non Standard Port**.



The following table describes the fields in the Events List for network or log events.

Field	Description
EVENT TIME	Shows the time the event occurred.
EVENT TYPE	Shows the type of alert, such as Log and Network.
DETECTOR IP	Shows the IP address of the machine where an anomaly was detected
FILE NAME	Shows the file name if a file is involved with the event.
FILE HASH	Shows a hash of the file contents.
SOURCE IP	Shows the source IP address if there was a transaction between two machines.
SOURCE PORT	Shows the source port of the transaction. The source and destination ports can be on the same IP address.
SOURCE HOST	Shows the destination host where the event took place.
SOURCE MAC	Shows the MAC address of the source machine.
SOURCE USER	Shows the user of the source machine.
TARGET IP	Shows the destination IP address if there was a transaction between two machines
TARGET PORT	Shows the destination port of the transaction. The source and destination ports can be on the same IP address.
TARGET HOST	Shows the HOST name of the destination machine.
TARGET MAC	Shows the MAC address of the destination machine.
TARGET USER	Shows the user of the destination machine.

The following figure shows an Events List for NetWitness Endpoint events.



The following table describes the fields in the Events List for NetWitness Endpoint events. NetWitness Endpoint events have an `Endpoint` Event Type and an `nwendpoint` Device Type. NetWitness Endpoint events from version 4.4.x and earlier can have an Event Type that shows the origin of the event.

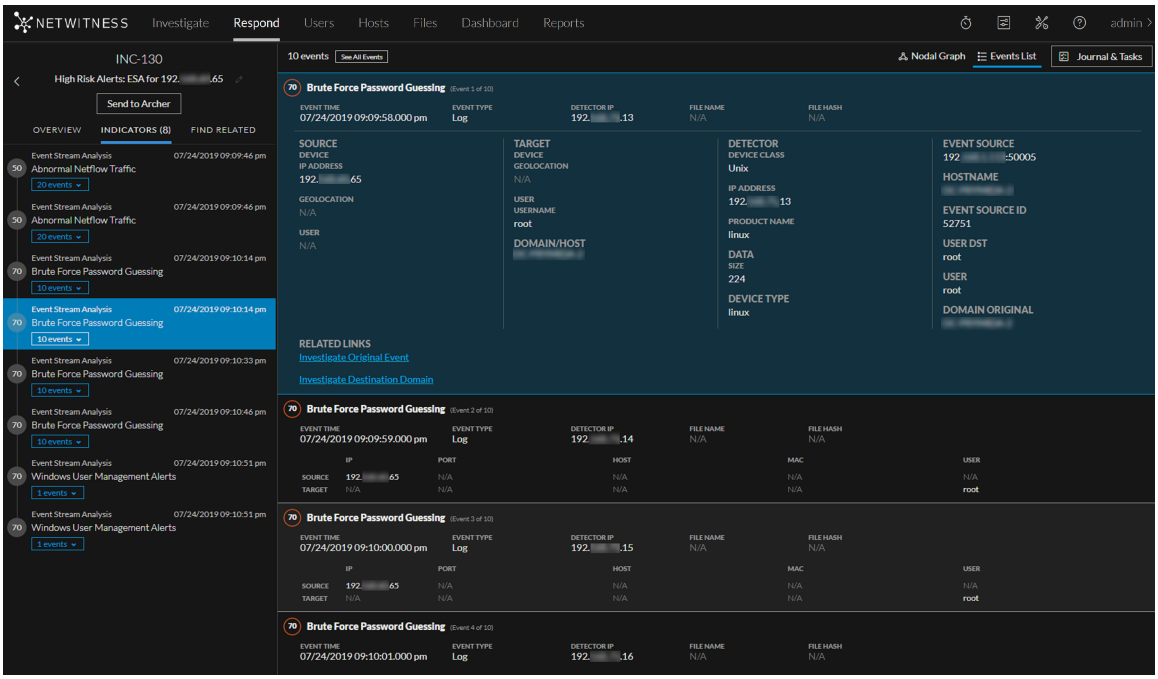
Field	Description
EVENT TIME	Shows the time the event occurred.
EVENT TYPE	Shows the type of alert, such as Endpoint or Log. NetWitness Endpoint events have an `Endpoint` event type. NetWitness Endpoint events from version 4.4.x and earlier can have an Event Type that shows the origin of the event.
CATEGORY	Shows the NetWitness Endpoint category.
ACTION	Shows the action that the file performed.
HOSTNAME	Shows the name of the machine that is running the agent.
USER ACCOUNT	Shows the username of the actively logged in user.
OPERATING SYSTEM	Shows the operating system of the agent.
FILE HASH	Shows the checksum of the file.

Field	Description
SOURCE FILENAME	Shows the name of the source file.
SOURCE LAUNCH ARGUMENT	Shows the command line argument for the running process.
SOURCE PATH	Shows the path of the source file.
SOURCE HASH	Shows the checksum of the source file.
SOURCE IP ADDRESS	Shows the IP address of the agent.
SOURCE PORT	Shows the source port of the connection.
TARGET FILENAME	Shows the name of the target file.
TARGET LAUNCH ARGUMENT	Shows the command line argument for the running process.
TARGET PATH	Shows the path of the target file.
TARGET HASH	Shows the checksum of the target file.
TARGET IP ADDRESS	Shows the destination IP address of this NetWitness activity.
TARGET PORT	Shows the destination port of the connection.
EVENT SOURCE	Shows the hostname or IP address along with the port of the of the Core service that holds the event information.
DEVICE TYPE	Shows the type of the device from which the data is sent or collected. For example, it shows `nwendpoint` for NetWitness Endpoint.

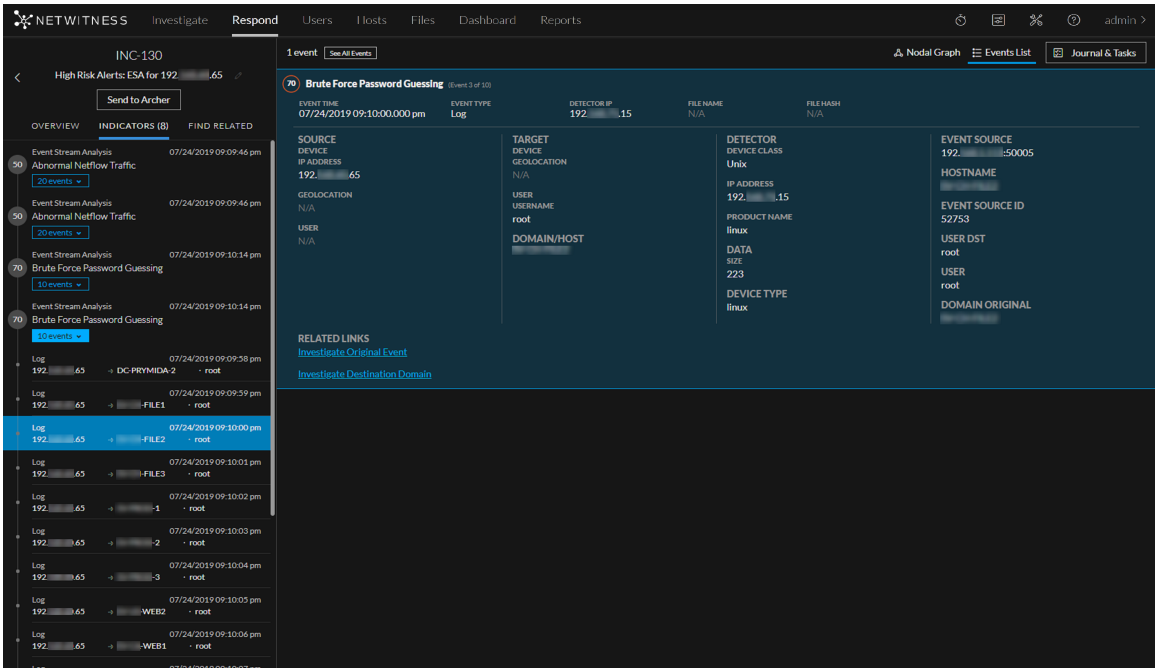
Event Details

To view the event details, you can click the top of an event in the Events List. The details appear below the event. Viewing inline event details enables you to keep the context of the event as it relates to the other events.

The following figure shows an indicator (alert) selected in the Indicators panel. The events for that indicator appear in the Events List on the right.

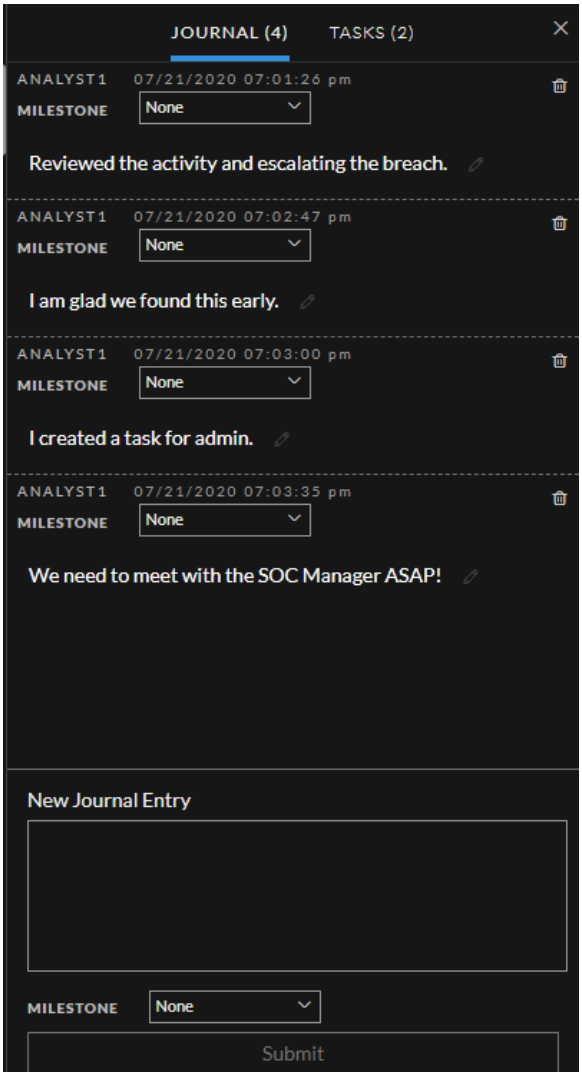


The following figure shows a specific indicator event selected in the Indicators panel. Information about the selected event appears in the Events List on the right.



Journal Panel

The incident Journal shows the history of activity on your incident.

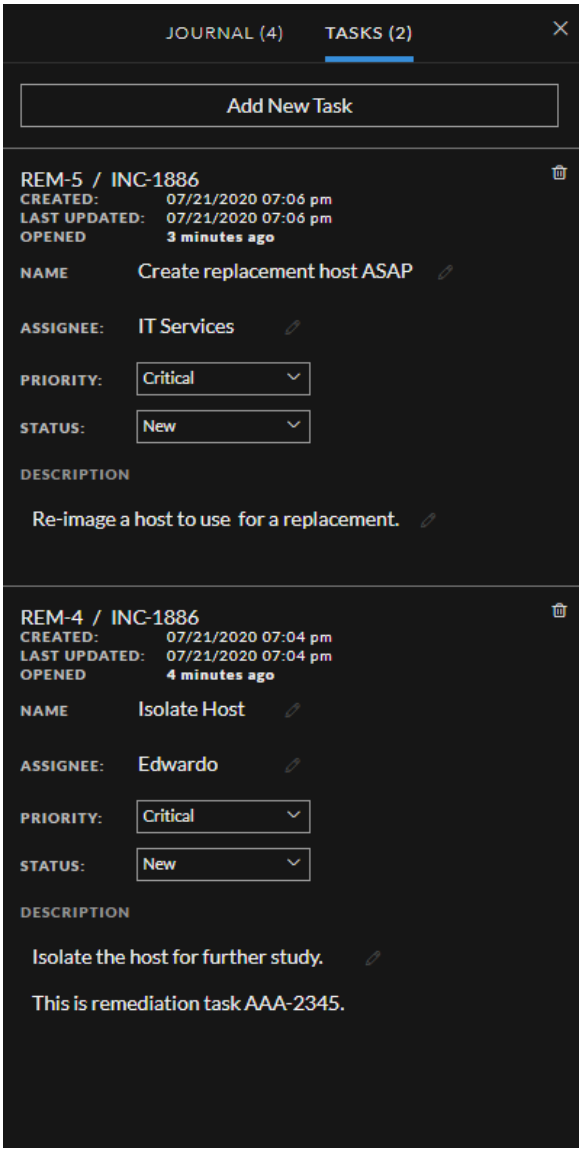


The following table describes the New Journal Entry options.

Field	Description
New Journal Entry	Type your note in the field.
Milestone	(Optional) Select a milestone, if applicable. This field is used to track significant events for the incident.
Submit button	Click submit to add an entry to the journal. You journal entry will be visible to anyone who views the incident.

Tasks Panel

In the Tasks panel, you can manage and track the incident tasks to closure.

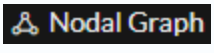


The following table describes the Task fields.

Field	Description
<Task ID / <Incident ID>	The autogenerated Task ID / The incident associated with the task.
Created	The created date of the task.
Last Updated	The date that the task was last modified.

Field	Description
Opened	The time that passed since the task was opened. For example, 3 minutes ago or 2 days ago.
Name	The name of the task. For example: Re-image the machine. You can click this field to edit it.
Assignee	The username of the user assigned to the task. You can click this field to edit it.
Priority	The priority of the task: Low, Medium, High, or Critical. You can click the priority button and select a new priority for the task from the drop-down list.
Status	The status of the task: New, Assigned, In Progress, Remediated, Risk Accepted, and Not Applicable. You can click the status button and select a new status for the task from the drop-down list.
Description	Type information that describes the task. You may want to include any applicable reference numbers. You can click this field to edit it.

Toolbar Actions

Option	Description
	(Back to Incidents) Enables you to navigate back to the Incidents List view.
	Closes the panel.
	Deletes the entry, such as a journal entry or task.
Priority button	(In the Overview panel) Allows you to change the Priority of one or more selected incidents in the Incidents List.
Status button	(In the Overview panel) Allows you to change the Status of one or more selected incidents.
Assignee button	(In the Overview panel) Allows you to change the Assignee of one or more selected incidents.
	Enables you to view the Nodal Graph.
	Enables you to view the incident Events List. Clicking the top of an event enables you to view the event details below it.
	Enables you to view incident notes and tasks.

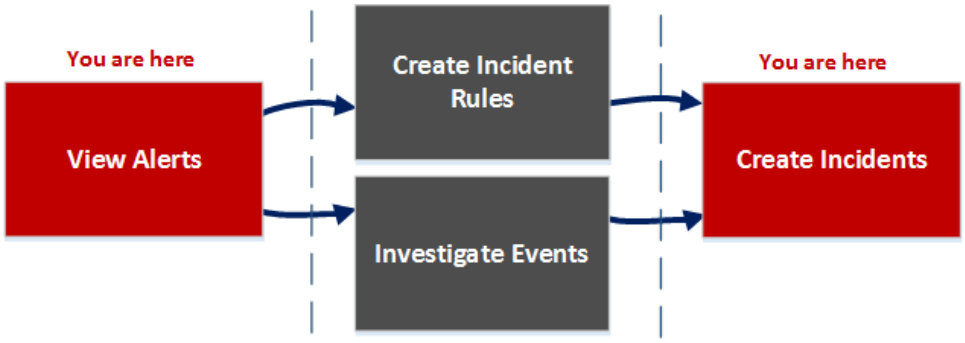
Option	Description
 (Journal, Tasks, and Related)	(This option is available in NetWitness Version 11.3.1 and earlier 11.x versions.) Enables you to view the Journal, Tasks, and Related Indicators panels.
	Enables you to show or hide the event header, request, response, and metadata in the Events panel in the Respond Incident Details view. For more information about event analysis, see "Events View" in the *NetWitness Investigate User Guide*.

Alerts List View

The Alerts List view (Respond > Alerts) enables you to view all of the threat alerts and indicators received by NetWitness in one location. This can include alerts received from ESA Correlation Rules, Malware Analysis, Reporting Engine, NetWitness Endpoint, as well as many others. In the Alerts List view you can browse through various alerts, filter them, and group them to create incidents.

Workflow

This workflow shows the high-level process that Analysts use to review alerts and create incidents.



In the Alerts List view, you can review a list of alerts from all sources received by NetWitness. After that, you can investigate those alerts further and create incidents from the alerts or you can create incident rules to create incidents.

Note: You can use NetWitness Automated Threat Detection to create incidents without manually creating rules.

What do you want to do?

Role	I want to ...	Show me how
Incident Responders, Analysts	View all alerts in NetWitness.*	View Alerts
Incident Responders, Analysts	Filter alerts.*	Filter the Alerts List
Incident Responders, Analysts	View alert overview information and raw alert metadata.*	View Alert Summary Information
Incident Responders, Analysts	Create incidents from alerts.*	Create an Incident Manually

Role	I want to ...	Show me how
Incident Responders, Analysts	(Available in NetWitness Version 11.1 and later) Add alerts to an existing incident.*	Add Alerts to an Incident
Administrators, Data Privacy Officers	Delete alerts.*	Delete Alerts
SOC Managers, Administrators	Create incident rules.	See "Create an Incident Rule for Alerts" in the *NetWitness Respond Configuration Guide*.
Incident Responders, Analysts	Investigate the events in an alert.	View Event Details for an Alert and Investigate Events
Incident Responders, Analysts	Add related alerts to an existing incident.	Add Related Indicators to the Incident

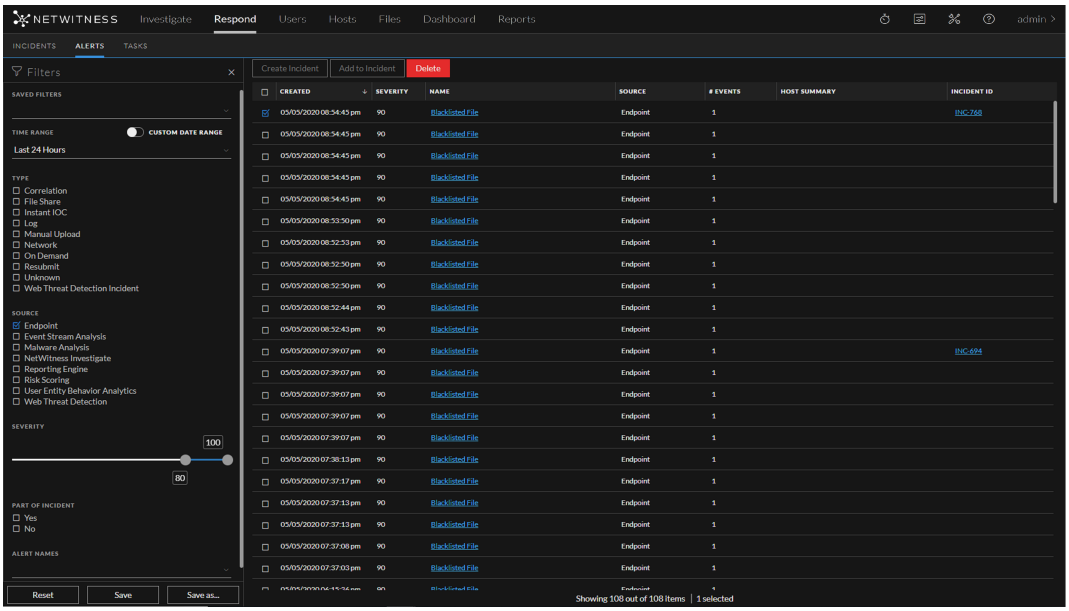
*You can complete these tasks here (that is, in the Alerts List view).

Related Topics

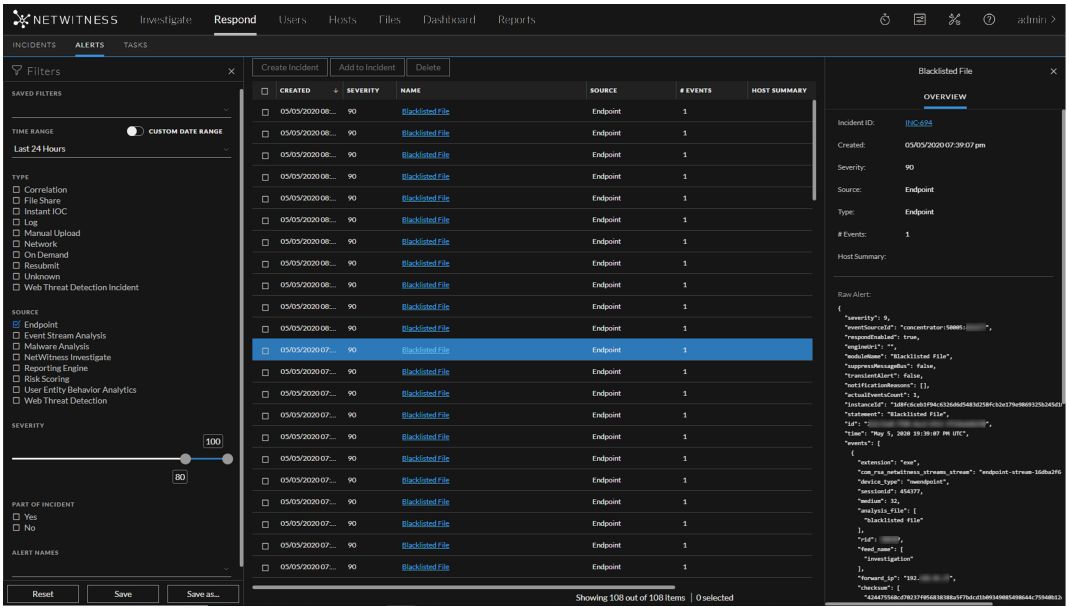
- Alert Details View
- Reviewing Alerts

Quick Look

To access the Alerts List view, go to **Respond > Alerts**. The Alerts List view displays a list of all alerts and indicators received by the Respond Server database in NetWitness. The following figure shows the Filters panel on the left.

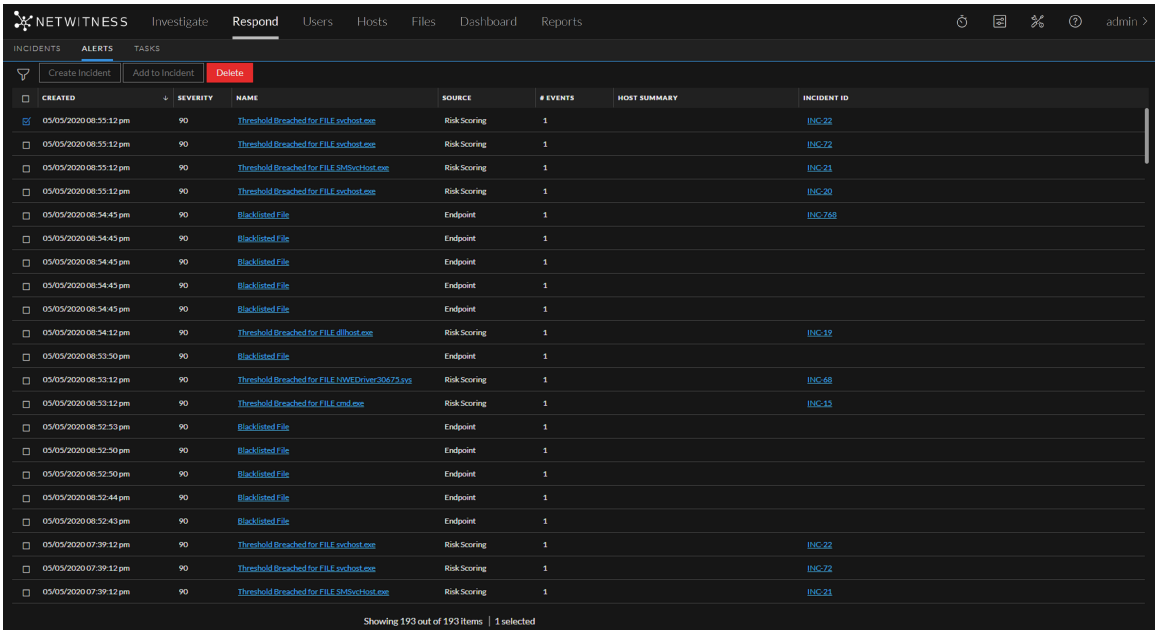


The Alerts List view consists of a Filters panel, an Alerts List, and an Alert Overview panel. You can click an alert in the Alerts list to view the Alert Overview panel on the right.

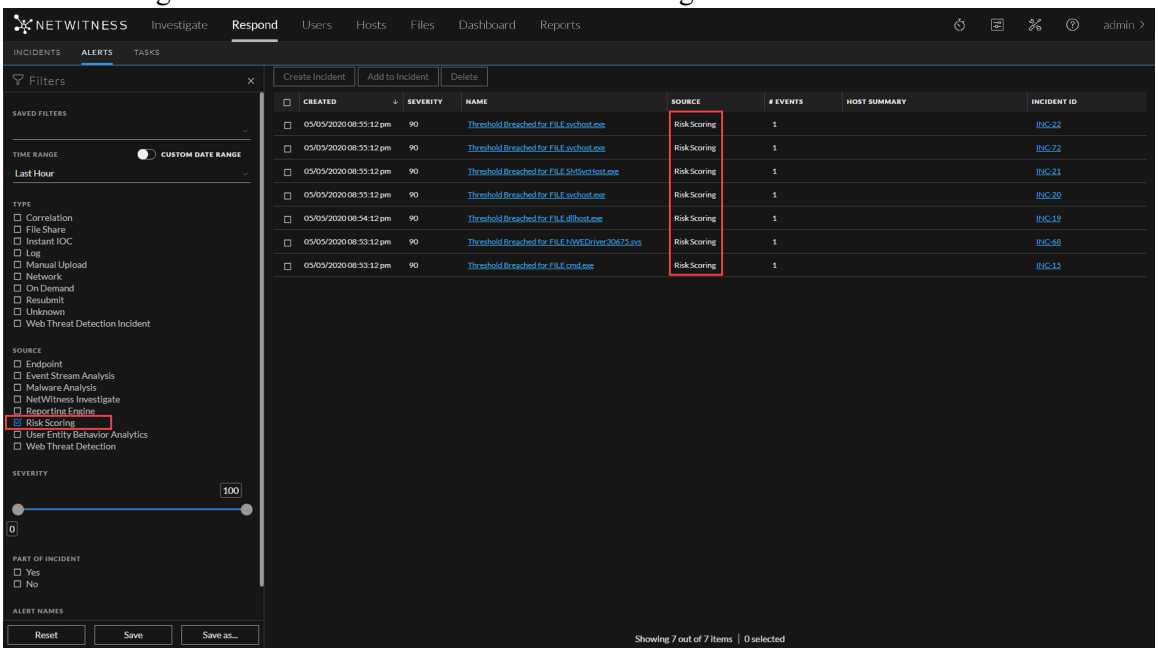


Alerts List

The Alerts List shows all of the alerts in NetWitness. You can filter this list to only show alerts of interest.



The following Alerts List view is filtered for Risk Scoring Alerts.



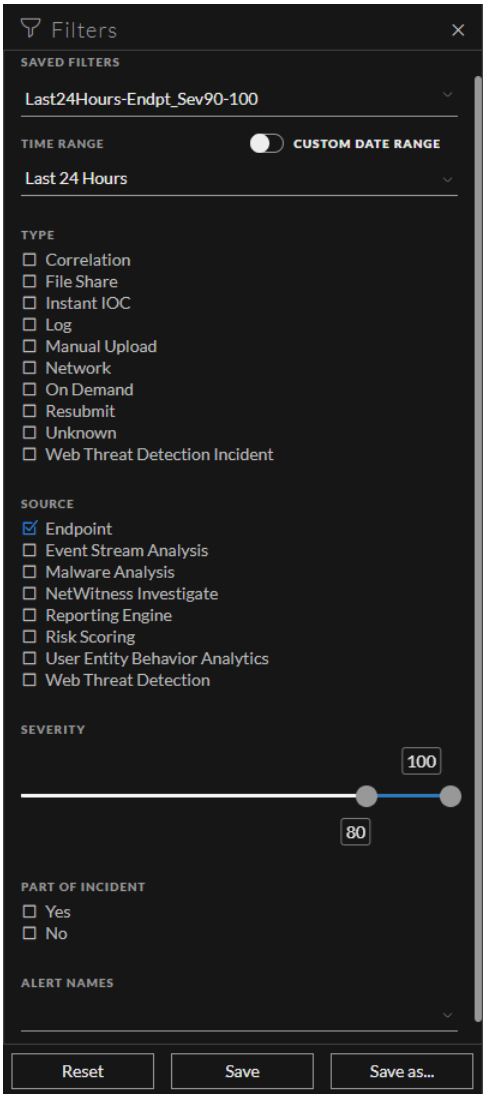
The following table describes the columns in the Alerts List.

Column	Description
	Enables you to select one or more alerts to delete. Users with the appropriate permissions, such as Administrators and Data Privacy Officers, can delete alerts.
Created	Displays the date and time when the alert was recorded in the source system.
Severity	Displays the level of severity of the alert. The values are from 1 through 100.
Name	Displays a basic description of the alert.
Source	Displays the original source of the alert. The source of the alerts can be NetWitness Endpoint, Malware Analysis, ESA correlation rules, Reporting Engine, Risk Scoring, and many others. **Note:** In NetWitness Platform 11.3 and later, the **Endpoint** source includes Endpoint alerts from all NetWitness Endpoint versions. If one of the events in an alert has a device_type of `nwendpoint`, the source changes to Endpoint.
# Events	Indicates the number of events contained within an alert. This varies depending on the source of the alert. For example, NetWitness Endpoint and Malware Analysis alerts always have one Event. For certain types of alerts, a high number of events may mean that the alert is more risky.
Host Summary	Displays details of the host like the host name from where the alert was triggered. The details may include information about the source and destination hosts in an Alert. Some alerts may describe events across more than one host .
Incident ID	Shows the Incident ID of the alert. If there is no incident ID, the alert does not belong to any incident and you can create an incident to include this alert or the alert can be added to an existing incident.

At the bottom of the list, you can see the number of alerts on the current page, the total number of alerts, and the number of alerts selected. For example: **Showing 4 out of 4 items | 1 selected**

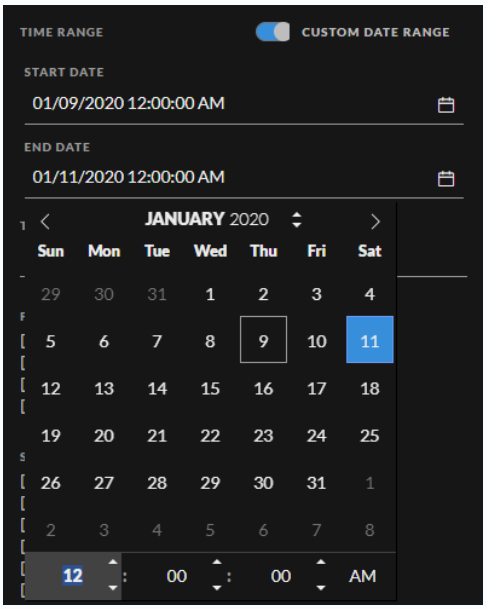
Alert Filters Panel

The following figure shows the filters available in the Filters panel.



The Filters panel, on the left of the Alerts List view, has options that you can use to filter the alerts list. When you navigate away from the Filters panel, the Alerts List view retains your filter selections.

Option	Description
Saved Filters	You can select a saved filter to filter the alerts list. Saved filters are global. You can save a filter for other analysts to use and you can use any saved filter. Saved filters are also available for use on the Springboard landing page. Filters used in the Springboard cannot be deleted. (This option is available in NetWitness Platform 11.5 and later.)

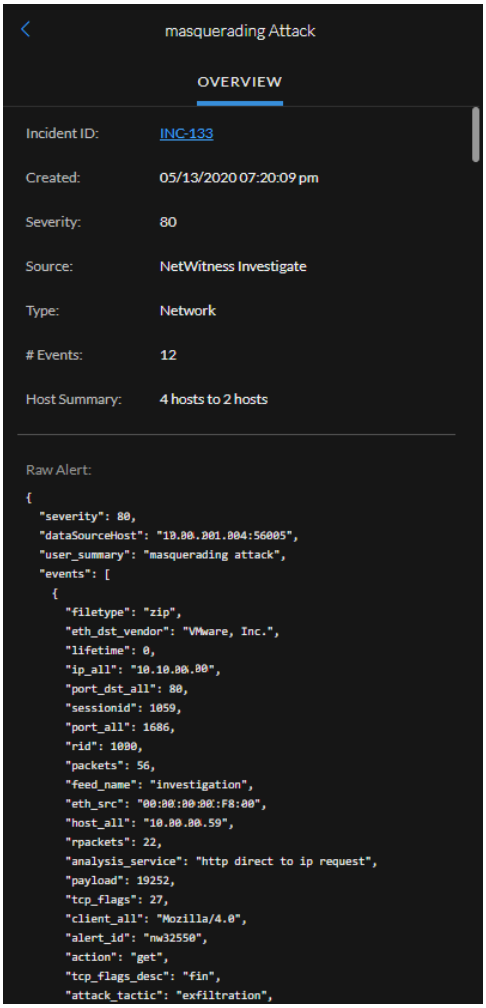
Option	Description
Time Range	You can select a specific time period from the Time Range drop-down list. The time range is based on the received date of the alerts. For example, if you select Last Hour, you can see alerts that were received within the last 60 minutes.
Custom Date Range	You can specify a specific date range instead of selecting a Time Range option. To do this, click the white circle in front of Custom Date Range to view the Start Date and End Date fields. Select the dates and times from the calendar. 
Type	Indicates the type of events in the alert, for example, logs, network sessions, and so on.
Source	Displays the original source of the alert. The source of the alerts can be NetWitness Endpoint, Malware Analysis, Event Stream Analysis (ESA Correlation Rules), Reporting Engine, Web Threat Detection, Risk Scoring, and many others. **Note:** In NetWitness Platform 11.3 and later, the **Endpoint** source includes Endpoint alerts from all NetWitness Endpoint versions. If one of the events in an alert has a device_type of `nwendpoint`, the source changes to Endpoint.
Severity	Displays the level of severity of the alert. The values are from 1 through 100.
Part of Incident	Categorizes alerts on whether or not they are associated with an incident. Select **Yes** to view alerts that are part of an incident. Select **No** to view alerts that are not part of an incident. For example, before you create incidents from alerts, you may want to select No to view only those alerts that are not already part of an incident.

Option	Description
Alert Names	Shows the names of the alerts being filtered. You can use this filter to search for all alerts generated by a specific rule, for example, Direct Login to an Administrative Account.
Reset	Removes your filter selections. If you reset filters on a saved filter, it takes you to the default empty filter.
Save	Saves the currently applied alerts filter or updates a saved filter. For a new filter, choose a unique name that contains 1-256 alphanumeric characters, underscores, or hyphens. (This option is available in NetWitness Platform 11.5 and later.)
Save As	Saves the currently applied alerts filter for future use. Choose a unique name that contains 1-256 alphanumeric characters, underscores, or hyphens. (This option is available in NetWitness Platform 11.5 and later.)

The Alerts List shows a list of alerts that meet your selection criteria. You can see the number of items in your filtered list at the bottom of the alerts list. For example: **Showing 4 out of 4 items**

Alert Overview Panel

The Overview panel shows basic summary information about a selected alert and raw alert metadata. The Overview panel in the Alert Details view contains the same information, but in the Alerts Details view, you can expand the panel to view more information.



The following table lists the fields displayed in the Alert Overview panel.

Field	Description
<Alert Name>	Displays the name of the alert.
Incident ID	Displays the Incident ID associated with the alert. You can click the incident ID link to go to the Incident Details view of the associated incident. If there is no incident ID, the alert does not belong to an incident. You can create an incident for this alert or you can add it to an incident.
Created	Displays the date and time when the alert was created.

Field	Description
Severity	Displays the level of severity of the alert. The values are from 1 through 100.
Source	Displays the original source of the alert. The source of the alerts can be NetWitness Endpoint, Malware Analysis, ESA correlation rules, Reporting Engine, Risk Scoring, and many others. **Note:** In NetWitness Platform 11.3 and later, the **Endpoint** source includes Endpoint alerts from all NetWitness Endpoint versions. If one of the events in an alert has a device_type of `nwendpoint`, the source changes to Endpoint.
Type	Indicates the type of events in the alert, for example, logs, network sessions, and so on. There can be multiple types listed. **Note:** In NetWitness Platform 11.3 and later, the **Endpoint** source includes Endpoint alerts from all NetWitness Endpoint versions. If one of the events in an alert has a device_type of `nwendpoint`, the source changes to Endpoint.
# Events	Indicates the number of events contained within an alert. This varies depending on the source of the alert. For example, NetWitness Endpoint and Malware Analysis alerts always have one Event. For certain types of alerts, a high number of events may mean that the alert is more risky.
Raw Alert	Shows the raw alert metadata.

Toolbar Actions

This table lists the toolbar actions available in the Alerts List view.

Option	Description
	Enables you to open the Filters panel so that you can specify the alerts that you would like to see in the Alerts List.
	Closes the panel.
Create Incident button	Enables you to create incidents from alerts. The alerts cannot be part of an incident. To get a list of alerts without incidents, you can filter the Alerts List. In thePart of Incident section, select No.

Option	Description
Add to Incident button	(This option is available in NetWitness Version 11.1 and later.) Enables you to add selected alerts to an incident. The alerts cannot be part of an incident. To get a list of alerts without incidents, you can filter the Alerts List. In the Part of Incident section, select No.
Delete button	Allows you to delete alerts.

Alert Details View

In the Alert Details view (Respond > Alerts > click on a row in the Alerts List), you can view the overview of an alert, such as the source of the alert, the number of events within the alert, Incident ID, if it is part of an incident. You can also view the raw alert that contains detailed information about the events.

Workflow

This workflow shows the high-level process that Analysts use to review alerts and create incidents.



After reviewing the alerts list, you can investigate those alerts further and create incidents from the alerts, in the Alert Details view. In the Configure > Incident Rules view, you can create incident rules to create incidents.

Note: You can also use NetWitness Automated Threat Detection to create incidents without manually creating rules.

What do you want to do?

Role	I want to ...	Show me how
Incident Responders, Analysts	View all alerts in NetWitness.	View Alerts
SOC Managers, Administrators	Create incident rules.	See "Create an Incident Rule for Alerts" in the *NetWitness Respond Configuration Guide*.
Incident Responders, Analysts	View a list of events in the alert.*	View Event Details for an Alert
Incident Responders, Analysts	View event metadata for each event in the alert.*	View Event Details for an Alert

Role	I want to ...	Show me how
Incident Responders, Analysts	Further investigate the events in the alert.*	Investigate Events
Incident Responders, Analysts	Add alerts to an existing incident.	Add Alerts to an Incident Add Related Indicators to the Incident
Incident Responders, Analysts	Create incidents from alerts.	Create an Incident Manually
Data Privacy Officers, Administrators	Delete alerts.	Delete Alerts

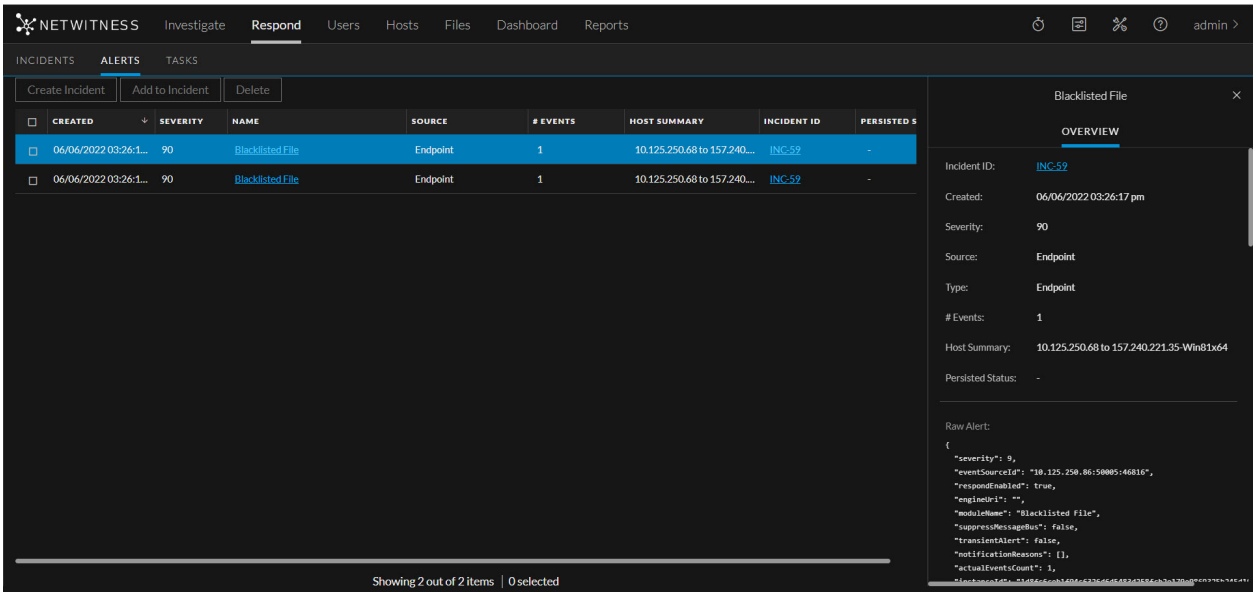
*You can complete these tasks here (that is, in the Alerts Details view).

Related Topics

- Alerts List View
- Reviewing Alerts

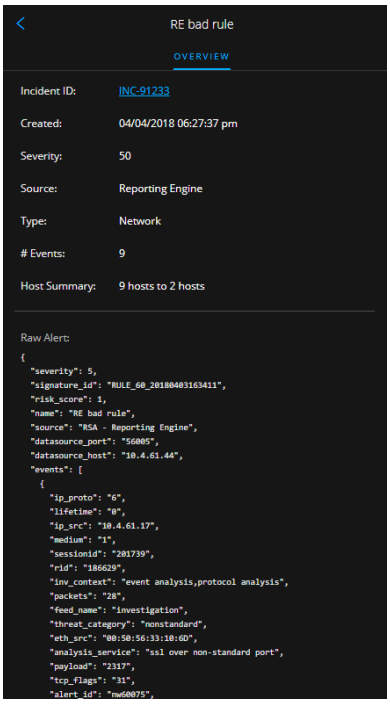
Quick Look

1. To access the Alert Details view, go to **Respond > Alerts**.
2. In the Alerts list, choose an alert to view and then click the link in the Name column for that alert. The Alert Details view has an Overview panel on the right. You can resize the panels to show more information as shown in the following figure.



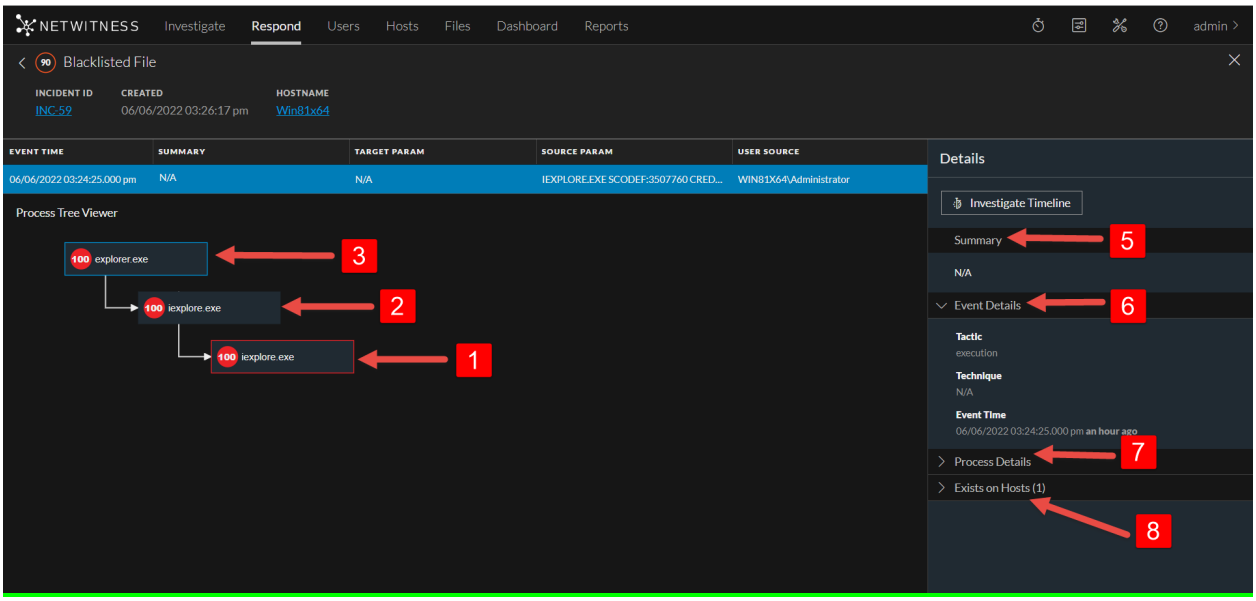
Overview Panel

The Overview panel shows basic summary information about a selected alert. The Overview panel on the Alerts List view contains the same information. The Alerts List view Alert Overview Panel topic provides details.



Events - Process Tree View

Click on an event name link to view the event details. The Process Tree Viewer opens and displays the process that caused the alerts and the processes it originated from.



1-The process that caused the alert is highlighted with a red-colored outline.

2 & 3 - The processes from which the highlighted process originated.

4 - Summary of the alert.

5 - Event Details section shows the tactics, techniques, and event time stamp.

6 - Process Details section provides detailed insights about the selected process.

7 - Shows the details of Network Connections established by the process; You can view the network connections that took place up to ten minutes before and after the alert triggered time. Network connections details are available only for the process that caused the alert.

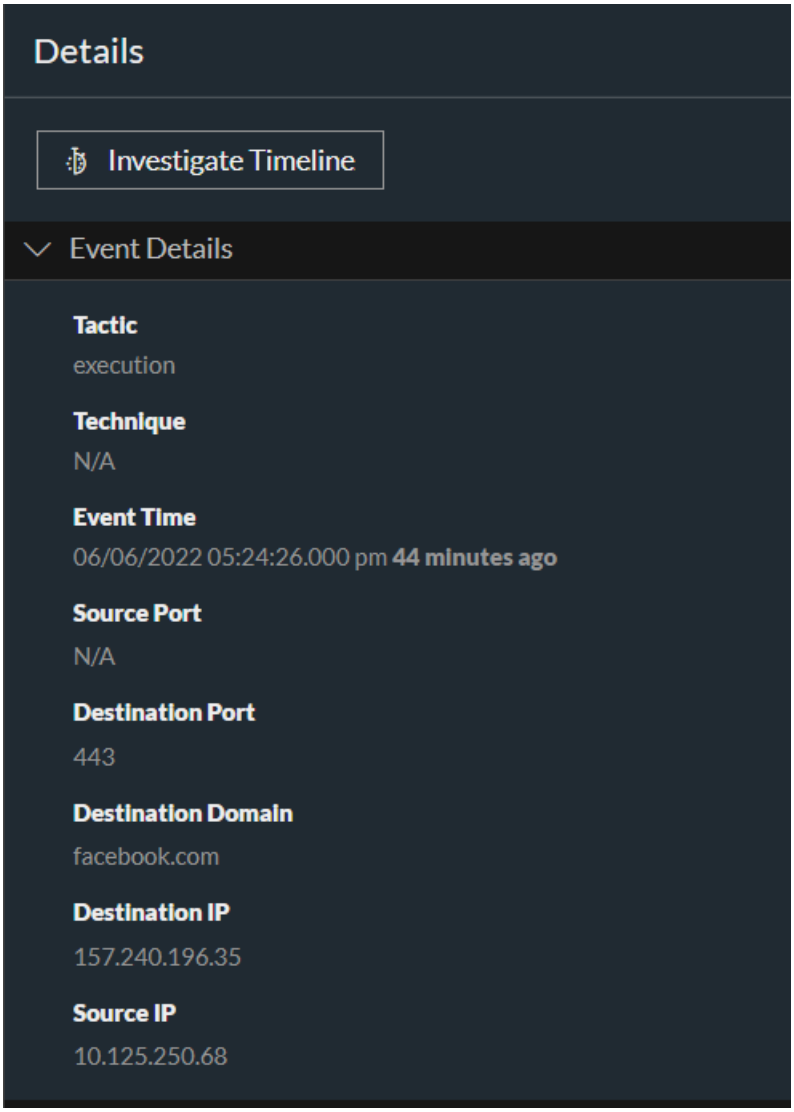
8 - Shows the name and a link to the host where the process exists.

Events List

The Events List for a selected alert shows all of the events contained in that alert.

Event Details

The Event Details in the Events panel shows the event metadata for each event in the alert.



Event Details

The following table lists some event details section and subsections shown in the Event Details. This is not an extensive list.

Section	Subsection	Description
Summary		Shows a summary of the event.
Event		Shows the destination device and user.

Section	Subsection	Description
	Device	Shows information about the destination device. See Event Source or Destination Device Attributes below.
	User	Shows information about the user or users of the destinationSee Event Source or Destination User Attributes below.
Detector		Shows the host or software product that detected the issue. This is most relevant for malware scanners and logs.
	Device Class	Shows the device class of the product that detected the alert.
	IP Address	Shows the IP address of the product that detected the alert.
	Product Name	Shows the name of the product that detected the alert.
Domain		Shows the domain associated with the event.
Enrichment		Shows available enrichment information.
Related Links		If available, it shows a link back to the user interface (UI) of the source product.
	Type	Shows the type of event, such as investigate_original_event.
	URL	Shows the URL link back to the UI of the source product.
Size		Shows the size of the transmission or file involved.
Source		Shows the source device and user.
	Device	Shows information about the source machine. See Event Source or Destination Device Attributes below.
	User	Shows information about the user or users of the source machine. See Event Source or Destination User Attributes below.
Timestamp		Shows the time that the event occurred.
Type		Shows the type of the alert, such as log, network, correlation, Resubmit, Manual Upload, On Demand, File Share, or Instant IOC.

Event Source or Destination Device Attributes

The following table lists attributes for an event source or destination device that can be shown in the Events Details.

Name	Description
Asset Type	Displays the type of device, for example, desktop, laptop, server, network equipment, tablet, and so on.
BusinessUnit	Shows the business unit associated with the device.

Name	Description
Compliance Rating	Shows the compliance rating of the device. It can be Low, Medium, or High.
Criticality	Shows how critical the device is to the business (business criticality).
Facility	Shows the location of the device.
Geolocation	Shows the geographic location for the host. It can contain the following attibutes: city, country, latitude, longitude, organization, and domain.
IP Address	Shows the IP address of the device.
MAC Address	Shows the MAC address of the device.
Netbios Name	Shows the NetBIOS name for the device.
Port	Displays the TCP port, UDP port, or the IP Src port (the first one available) used to connect to and from the host.

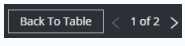
Event Source or Destination User Attributes

The following table lists attributes for an an event source or destination user that can be shown in the Events Details.

Attribute Name	Description
AD Domain	Shows the Active Directory domain.
AD Username	Shows the Active Directory username.
Email Address	Shows the email address of the user.
Username	Shows a general name if you do not know the source of the username, such as UNIX or a username in a particular system.

Toolbar Actions

This table lists the toolbar actions available in the Alert Details view.

Option	Description
	(Back to Alerts) Enables you to navigate back to the Alerts List view.
	Click the arrows to navigate through the event meta details for each event in the alert. The numbers, such as "1 of 2" show the number of the event that you are currently viewing. Click **Back to Table** to go back to the Events List view, which is also known as the Events Table.

Tasks List View

After investigating incidents, in the Tasks List view (Respond > Tasks), you can create and track incident tasks. For example, you can create remediation tasks when you require actions on incidents from teams outside of your security operations. You can reference external ticket numbers within the tasks and then track those tasks to completion. You can also modify and delete tasks as required, depending on your user permissions.

What do you want to do?

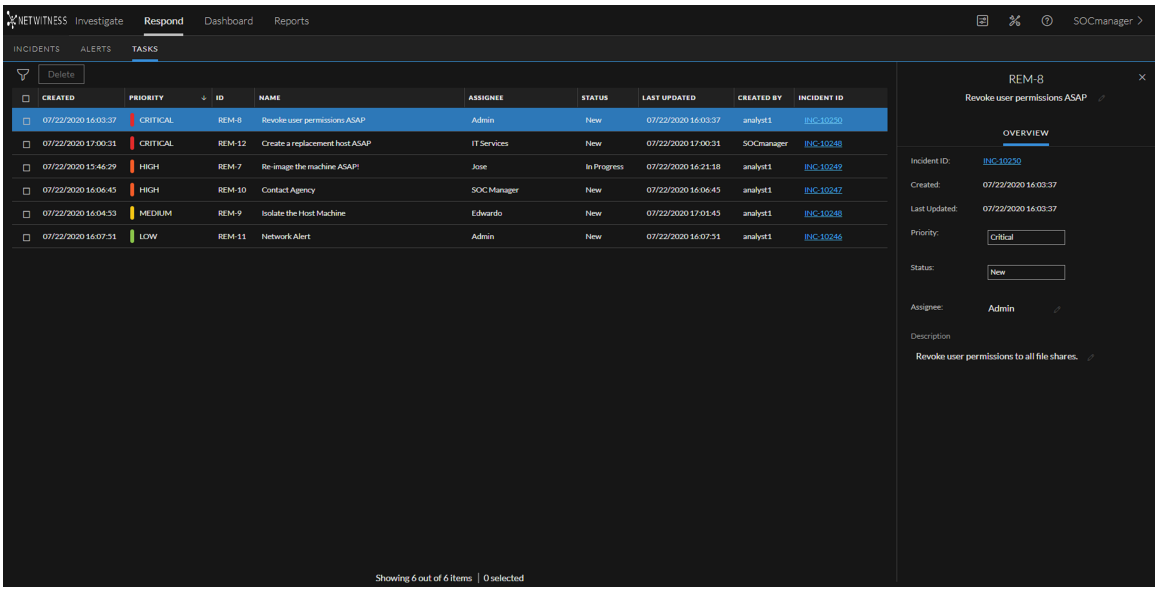
Role	I want to ...	Show me how
Incident Responders, Analysts	View tasks	View All Incident Tasks and View the Tasks Associated with an Incident
Incident Responders, Analysts	Filter tasks.	Filter the Tasks List
Incident Responders, Analysts	Create a task.	Create a Task
Incident Responders, Analysts	Find and modify tasks.	Find a Task and Modify a Task
Incident Responders, Analysts	Close a task (Change the Status to Remediated, Risk Accepted, or Not Applicable).	Modify a Task
Incident Responders, Analysts, SOC Managers	Delete a task.	Delete a Task

Related Topics

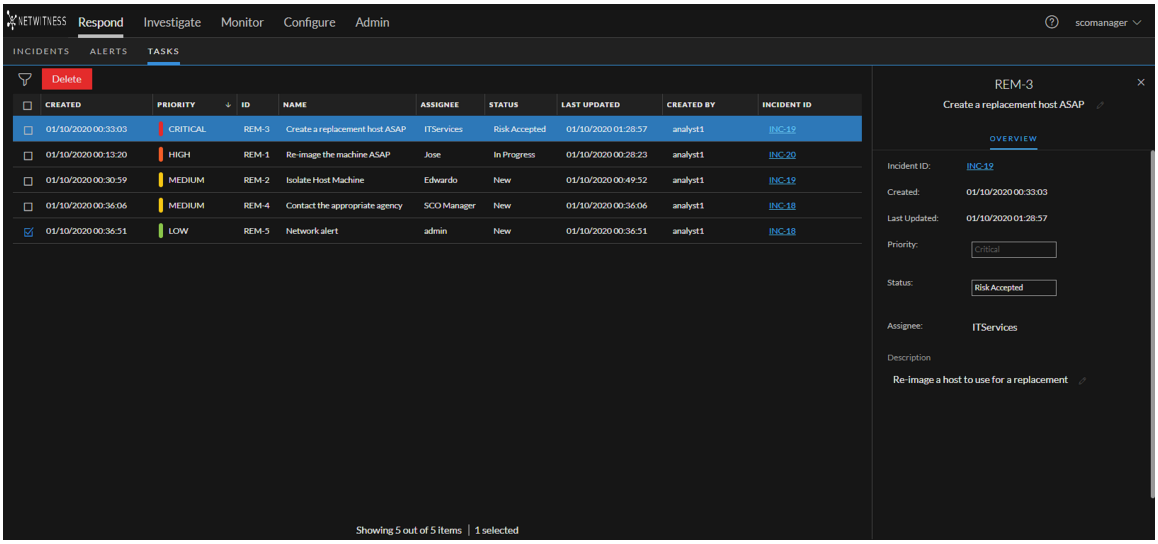
- Incident Details View
- Escalate or Remediate the Incident

Quick Look

To access the Tasks List view, go to **Respond > Tasks**. The Tasks List view displays a list of all incident tasks.



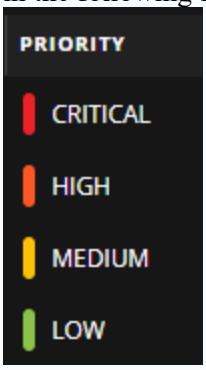
The Tasks List view consists of a Filters panel, a Tasks List, and a Task Overview panel. The following figure shows the Tasks List and the Overview panel.



Tasks List

The Tasks List shows all of the incident tasks. You can filter this list to show only tasks of interest.

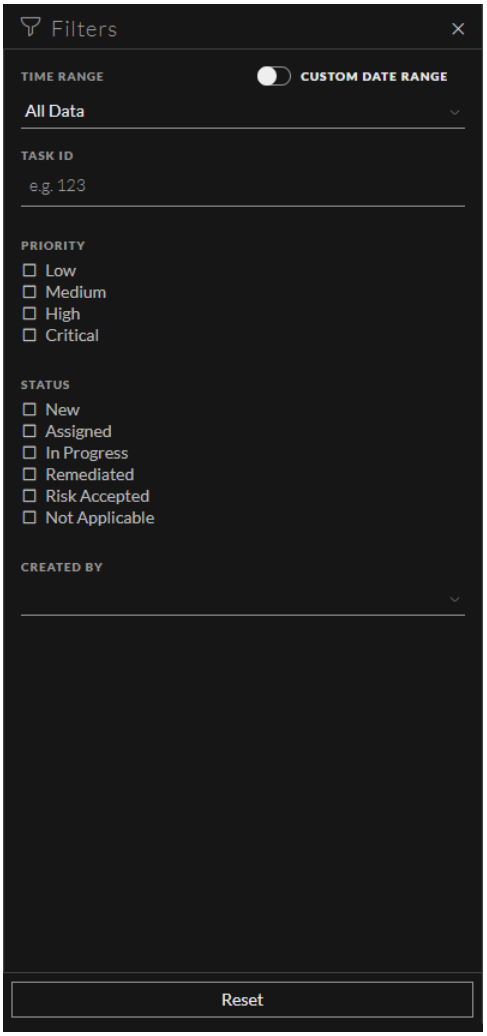
Column	Description
	Enables you to select one or more tasks to modify or delete. Users with the appropriate permissions can make bulk updates and delete tasks, such as SOC Managers. For example, an SOC Manager may want to assign multiple tasks to a user at the same time.
Created	Displays the date when the task was created.

Column	Description
Priority	Displays the priority assigned to the task. The priority can be any of the following: Critical, High, Medium, or Low. The Priority is also color coded, where red indicates **Critical**, orange represents **High** risk, yellow indicates **Medium** risk, and green represents **Low** risk as shown in the following figure: 
ID	Displays the task ID.
Name	Displays the task name.
Assignee	Displays the name of the user assigned to the task.
Status	Displays the status of the task: New, Assigned, In Progress, Remediated, Risk Accepted, and Not Applicable.
Last Updated	Displays the date and time when the task was last updated.
Created By	Displays the user who created the task.
Incident ID	Displays the incident ID for which the task was created. Click the ID to display the details of the incident.

At the bottom of the list, you can see the number of tasks on the current page and the total number of tasks. For example: **Showing 23 out of 23 items**

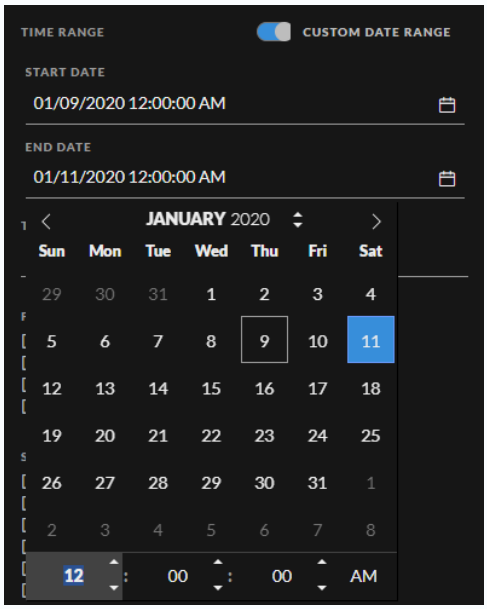
Task Filters Panel

The following figure shows the filters available in the Filters panel.



The Filters panel, on the left of the Tasks List view, has options that you can use to filter the incident tasks.

Option	Description
Time Range	You can select a specific time period from the Time Range drop-down list. The time range is based on the creation date of the tasks. For example, if you select Last Hour, you can see tasks that were created within the last 60 minutes.

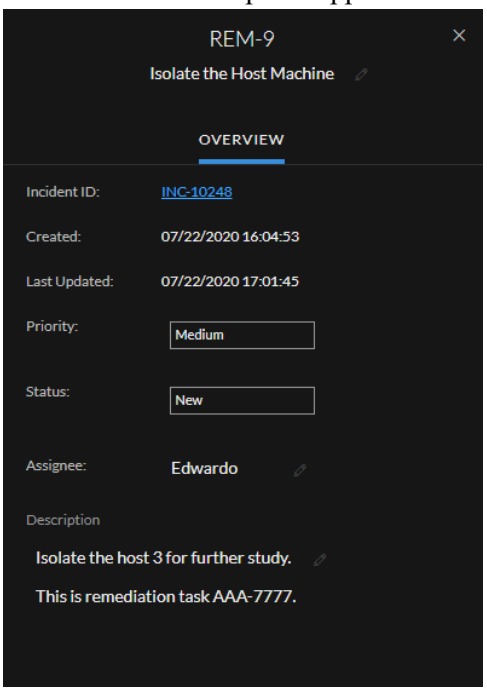
Option	Description
Custom Date Range	You can specify a specific date range instead of selecting a Time Range option. To do this, click the white circle in front of Custom Date Range to view the Start Date and End Date fields. Select the dates and times from the calendar. 
Task ID	You can type the Task ID for a task that you would like to locate, for example REM-123.
Priority	You can select the priorities that you would like to view. If you make one or more selections, the Tasks list shows only those tasks with the selected priorities. For example: If you select Critical, the Tasks list shows only the tasks with a priority set to Critical.
Status	You can select the statuses that you would like to view. If you make one or more selections, the Tasks list shows only those tasks with the selected statuses. For example: If you select Assigned, the Tasks panel shows only the tasks that are assigned to users.
Created By	You can select the user who created the tasks that you would like to view. For example, if you only want to view the tasks created by Edwardo, select Edwardo from the CREATED BY drop-down list. If you want to view tasks regardless of the person who created the task, do not make a selection under CREATED BY.
Reset Filters	Removes your filter selections.

The Tasks List shows a list of tasks that meet your selection criteria. You can see the number of items in your filtered list at the bottom of the tasks list. For example: **Showing 18 out of 18 items**

Task Overview Panel

To access the Task Overview panel:

1. Go to **Respond > Tasks**.
2. In the Task list, click the task that you want to view.
   The Task Overview panel appears to the right of the Tasks list.



The following table lists the fields displayed in the Task Overview panel.

Field	Description
<Task ID>	Displays the automatically assigned task ID.
<Task Name>	Displays the task name. This is an editable field. To change the task name, you can click the current task name to open a text editor. For example, you can change a task name from "Reimage a Laptop" to "Reimage a Server".
Incident ID	Displays the Incident ID for which the task was created. Click the ID to display the details of the Incident.
Created	Displays details about the date and time when the task was created.
Last Updated	Displays the date and time when the task was last updated.
Priority	Displays the priority of the task: Low, Medium, High, or Critical. To change the priority, you can click the priority button and select a priority for the task from the drop-down list.

Field	Description
Status	Displays the status of the task: New, Assigned, In Progress, Remediated, Risk Accepted, and Not Applicable. To change the status, you can click the status button and select a status for the task from the drop-down list.
Assignee	Displays the user assigned to the task. To change the user assigned to the task, you can click (Unassigned) or the name of the previous assignee to open a text editor.
Description	Shows task details. To modify the description, you can click the text underneath the description to open a text editor.

Toolbar Actions

This table lists the toolbar actions available in the Tasks List view.

Option	Description
	Enables you to open the Filters panel so that you can specify the tasks that you would like to see in the Tasks List.
	Closes the panel.
Delete button	Allows you to delete the selected tasks.

Add/Remove from List Dialog

The Add/Remove from List dialog allows you to add or remove an entity or meta value to an existing list or create a new list. For example, when you look up an IP address and you find it suspicious or interesting, you can add it to a relevant list, which has been added a data source. This improves the visibility of the suspicious IP addresses. You can also add entities or meta values to different lists. For example, you can add them to one list for suspected domains related to command and control connections and to another list for Trojan connections IP addresses related to remote access. If a list is not available, you can create a list. You can also remove the entity or meta value from a list.

Note: From the Add/Remove from List dialog, you can only add or remove entities or meta values from single column lists added as a datasource, not multi-column lists. And when you edit a list or a value in a list from the nodal view or the context lookup view, ensure to refresh the web page to view the updated data.

What do you want to do?

Role	I want to ...	Show me how
Incident Responders, Analysts	Add an entity to a list.	From the Incident Details view, see Add an Entity to a Whitelist. From the Alert Details view, Add an Entity to a Whitelist.
Incident Responders, Analysts	Create a whitelist, blacklist, or other list.	Create a List
Administrators	Add a Context Hub list as a data source.	See "Configure Lists as a Data Source" in the *Context Hub Configuration Guide*.
Administrators	Import or export a list for Context Hub.	See "Import or Export Lists for Context Hub" in the *Context Hub Configuration Guide*.

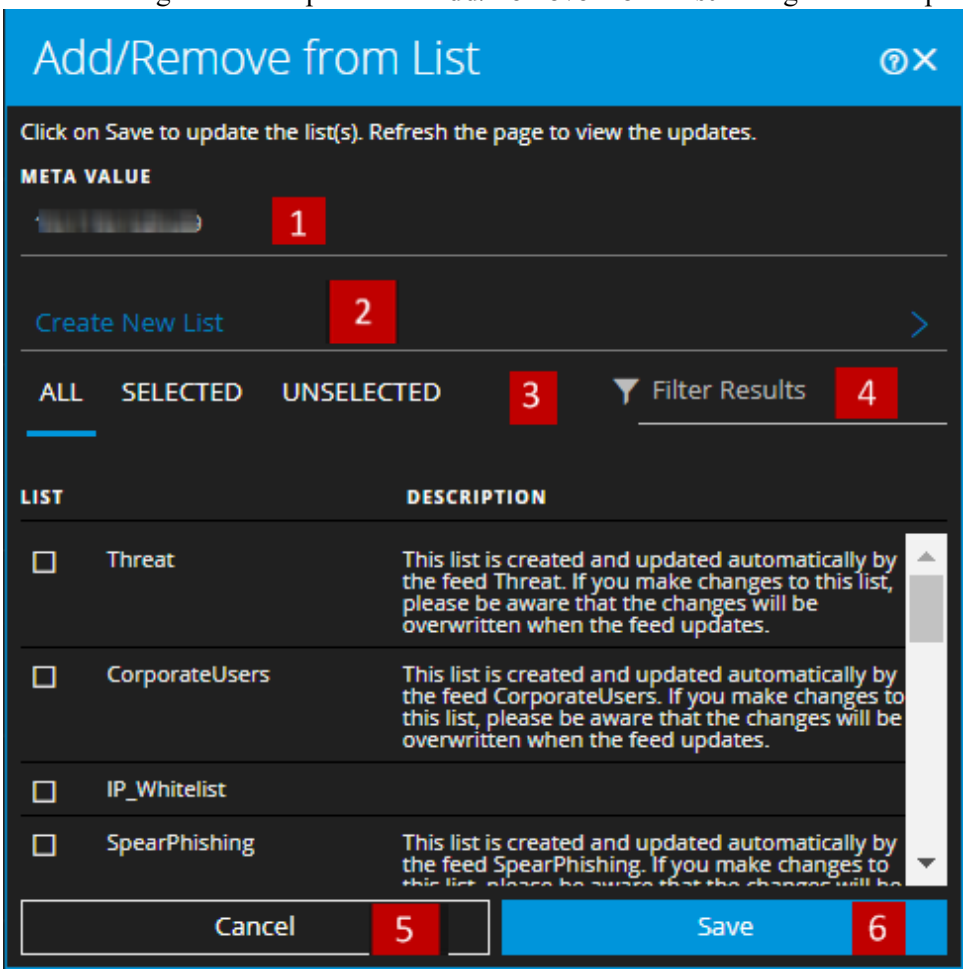
Related Topics

- Investigate the Incident
- Reviewing Alerts
- View Contextual Information (Incident Details view)
- View Contextual Information (Alert Details view)

Note: You cannot delete a list, but you can delete values within a list.

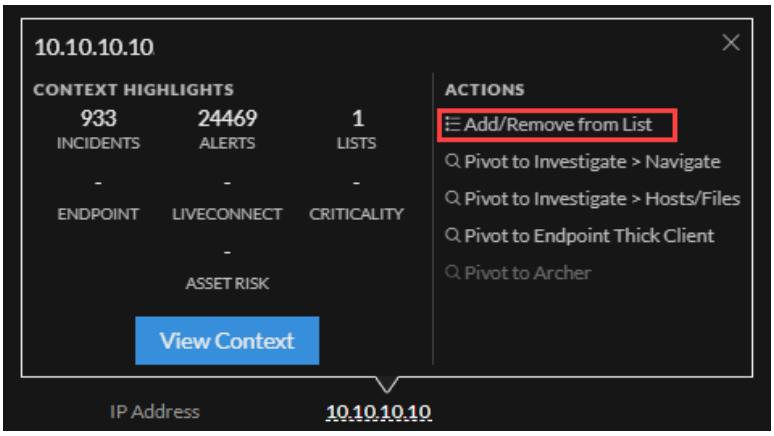
Quick Look

The following is an example of the **Add/Remove from List** dialog in the Respond view.

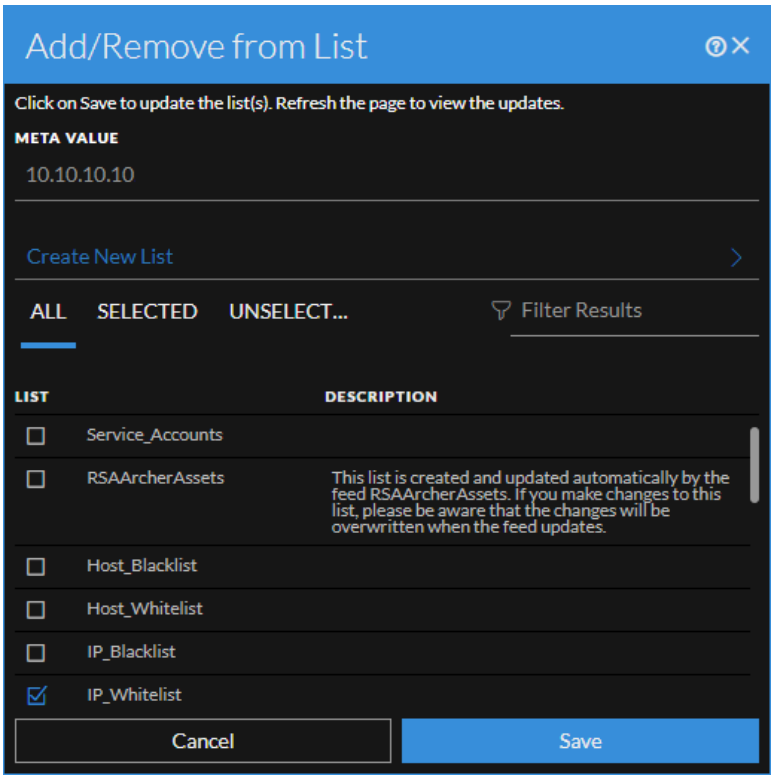


1 Entities or meta values to be added or removed.
2 Create a new list using the selected meta values.
3 Select any of the tabs: All, Selected, or Unselected.
4 Search using the list name or description.
5 Cancel the action.
6 Save to update lists or create a new list.

To access the Add/Remove from List dialog, in the Incident Details view or the Alert Details view, hover over the underlined entity that you would like to add or remove from a Context Hub list. A context tooltip appears showing the available actions.



In the Actions section of the tooltip, click Add/Remove from List. The Add/Remove From List dialog shows the available lists.



The following table shows the options in the Add/Remove from List dialog.

Option	Description
Meta Value	Displays the selected entity or meta value that needs to be added to or removed from one or more lists. You can also a create a new list using the selected value.
Create New List	When clicked, it displays a dialog to create a new list using the selected meta value.

Option	Description
All	Shows all of the available Context Hub lists. The lists that contain the selected entity or meta value are selected. Select a checkbox to add an entity or meta value to a list. Clear a checkbox to remove it from the list.
Selected	Shows only the lists that contain the selected entity or meta value. (All lists are selected.)
Unselected	Shows only the lists that do not contain the selected entity or meta value. (All lists are unselected.)
Filter Results	Enter the name or description of a specific list to search from multiple lists.
List	Displays the name of all the lists.
Description	Displays information about the selected list. The description that you provide when creating a list appears in this dialog. For example: This list contains all of the blacklisted IP addresses.
Cancel	Cancels the operation.
Save	Saves the changes.

Context Lookup Panel - Respond View

The Context Hub service brings together contextual information from several data sources into the Respond view so that analysts can make better decisions during their analysis and take appropriate action. Seeing the entities, meta values, and contextual information in a single interface helps analysts to prioritize and identify areas of interest. For example, recently created incidents and alerts from the Respond view involving a given entity or meta value will be displayed when the analyst queries for additional information for that entity or meta value. The Context Lookup panel displays contextual information for the selected entities or meta values such as IP address, User, Host, Domain, File Name, or File Hash. The data available depends on the configured sources in the Context Hub.

The Context Lookup panel displays the contextual information based on the data available on the configured sources in the Context Hub.

What do you want to do?

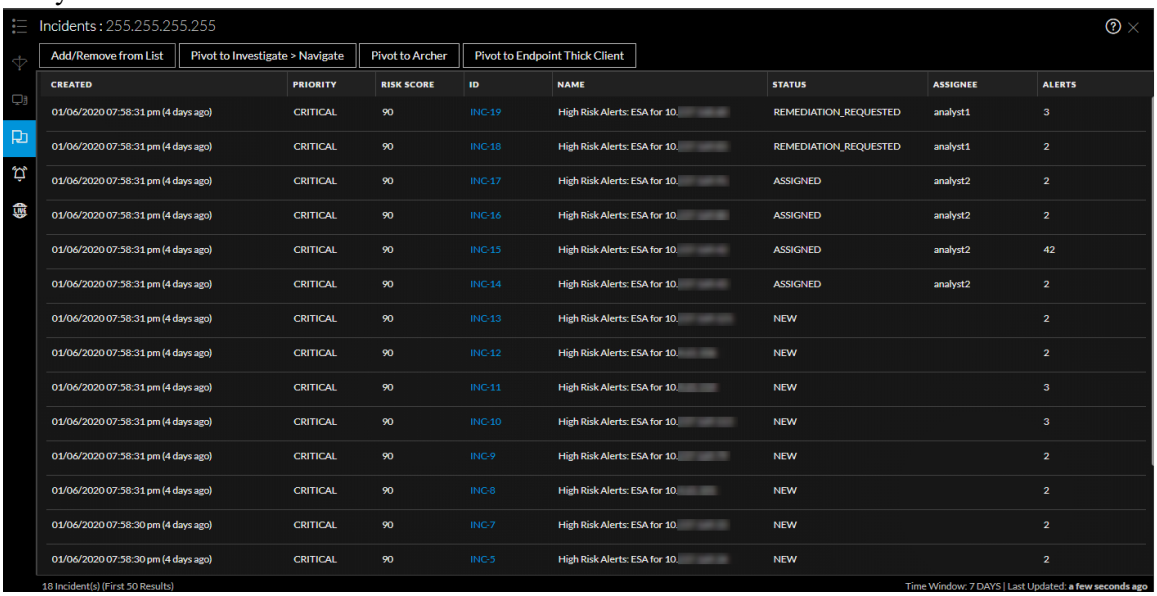
Role	I want to ...	Show me how
Incident Responders, Analysts, Threat Hunters	Navigate to the Context Lookup panel.	From the Incident Details view, see View Contextual Information. From the Alert Details view, see View Contextual Information.
Incident Responders, Analysts, Threat Hunters	Understand the information in the Context Lookup panel for a selected entity.	See the information in this topic.
Administrator	Configure Data Sources for Context Hub.	See "Configure Data Sources for Context Hub" in the *Context Hub Configuration Guide*.
Administrator	Configure Context Hub settings.	See "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.

Related Topics

- Investigate the Incident
- Reviewing Alerts

Contextual Information Displayed in the Context Lookup Panel

The contextual information or query results displayed in the Context Lookup panel depends on the selected entity and the associated data sources. The Context Lookup panel has separate tabs for each of the data sources. The tabs are: List data source, Archer, Active Directory, Endpoint, Incidents, Alerts, Live Connect, and REST API. The following figure shows the Context Lookup panel for a selected entity in the Incident Details view.



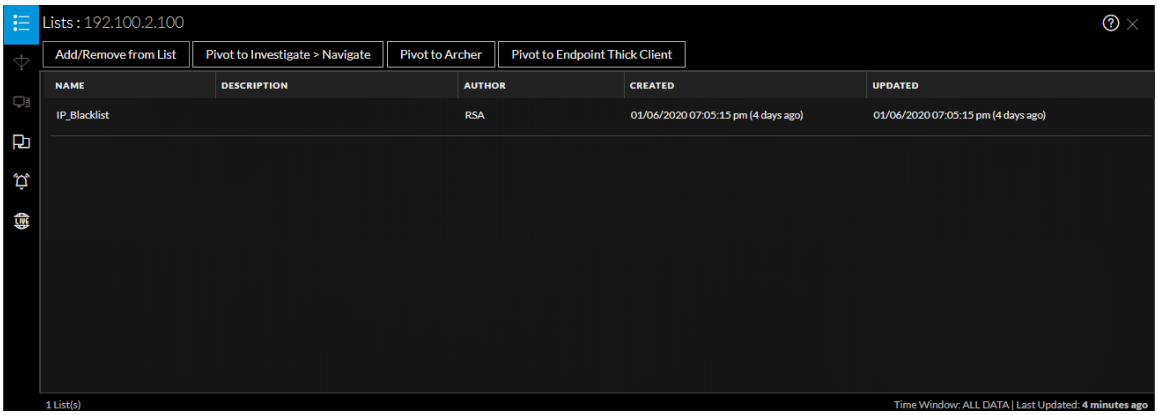
The following table describes the data available on each tab and the supported entities.

Tab	Description	Supported Entities
 (Lists)	Displays all of the list data associated with the selected entity or meta value. The result is sorted by the last updated list.	All entities
 (Archer)	Displays asset information along with criticality ratings using the Archer data source.	IP, Host, and Mac
 (Active Directory)	Displays all user information for the selected user.	User

Tab	Description	Supported Entities
 (NetWitness Endpoint)	Displays the NetWitness Endpoint data source information for the selected entity or meta value, which includes the Machines, Modules, and IIOC levels. Modules are by highest IOC score to lowest IIOC score and IIOC levels are sorted by highest IOC levels to lowest IOC levels.	IP, MAC address, and Host
 (Incidents)	Displays the list of incidents associated with the selected entity or meta value. The result is sorted by newest incidents to oldest incidents.	All entities
 (Alerts)	Displays the list of alerts associated with the selected entity or meta value. The result is sorted by newest alerts to oldest alerts.	All entities
 (Live Connect)	Displays information related to Live Connect.	IP, Domain, and Filehash
 (File Reputation)	Displays file reputation status for Filehash entities.	Filehash entities
 TI	Displays information for STIX data sources.	IP address, email address, domain, filename, URL's, and file hash. **Note:** The context lookup for email address and URL will be displayed only if these metas are mapped. Navigate to  (Admin) > System > Investigation > Context Lookup.
 REST API	Displays the list of REST APIs (enabled in Context Hub) associated with selected the entity.	All entities

Lists Tab

The Context Lookup panel for Lists shows one or more lists associated with the selected entity or meta value. The following figure is an example of the Context Panel for Lists, and the table describes the fields.



Field	Description
Name	The name of the list (defined while creating the list).
Description	The description of the list (defined while creating the list).
Author	The owner who created the list.
Created	The date when the list was created.
Updated	The date when the list was last updated or modifed.
Count	The number of lists in which the selected entity or meta value is available.
Time Window	The time window based on the value set for the "Query Last" field in the Configure Responses dialog. By default, all Lists data is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

Archer Tab

The Context Lookup panel for Archer displays asset information along with criticality ratings using the Archer data source for IP, Host, and Mac entities. The following figure is an example of the Context Lookup panel for Archer, and the table describes each field.



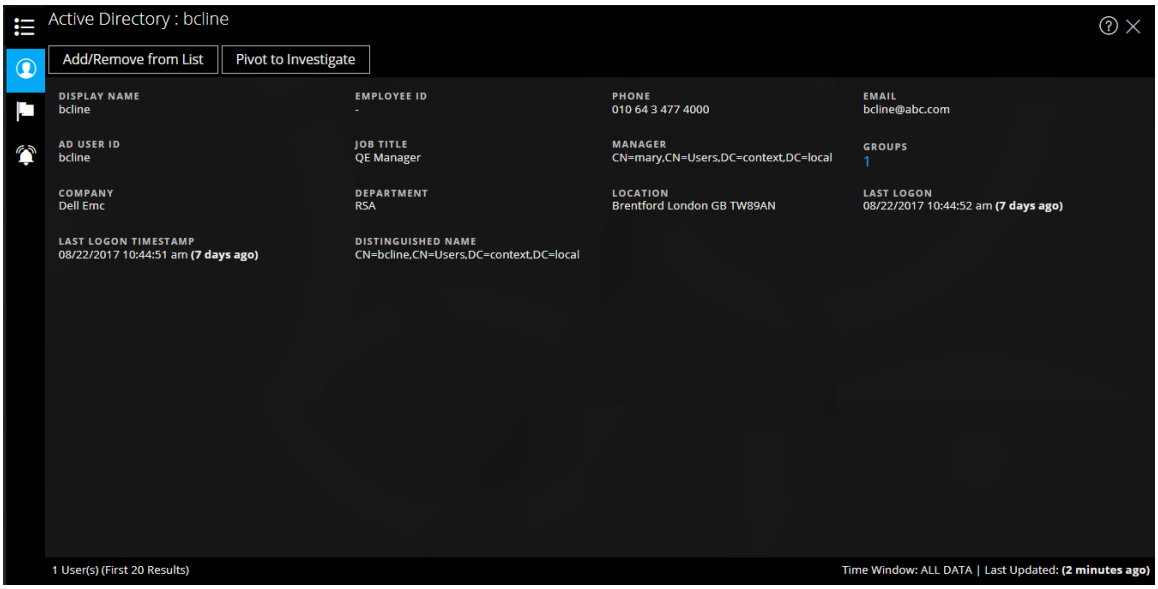
Field	Description
Criticality Rating	The device operational criticality based on the applications it supports. The criticality ratings can be set as Not Rated, Low, Medium-Low, Medium, Medium-High, or High.
Risk Rating	The calculated risk rating for the device based on the most recent assessment and the average risk rating of facilities using the device. The risk rating can be set as Severe, High, Medium, Low, or Minimal.
Device Name	The unique name of the device.
Host Name	The host name of the device.
IP Address	The primary internal IP address of the device.
Device ID	The automatically populated value that uniquely identifies the record across all applications within the system.
Type	The device type, for example, server, laptop, desktop, and others.
Facilities	Links to records in the Facilities application that are related to this device.
Business Unit	Links to records in the Business Unit application that are related to this device. For more than three business unit values, you can hover over the field to view the values.
Device Owner	The person who is responsible for the device and receives read and update rights of the record.

Field	Description
Count	The number of assets available.
Time Window	The time window based on the value that is set for the "Query Last" field in the Configure Responses dialog. By default, all data for Archer is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

Note: In the localized versions, only these twelve fields are displayed: Criticality Rating, Risk Rating, Device Owner, Business Unit, Host Name, MAC Address, Facilities, IP Address, Type, Device ID, Device Name, and Business Processes.

Active Directory Tab

The following figure is an example of a Context Lookup panel for Active Directory.



The Context Lookup panel for Active Directory displays all the related information, incidents, and alerts for a user. You can perform a look up using the following formats:

- userPrincipalName
- Domain\UserName
- sAMAccountName

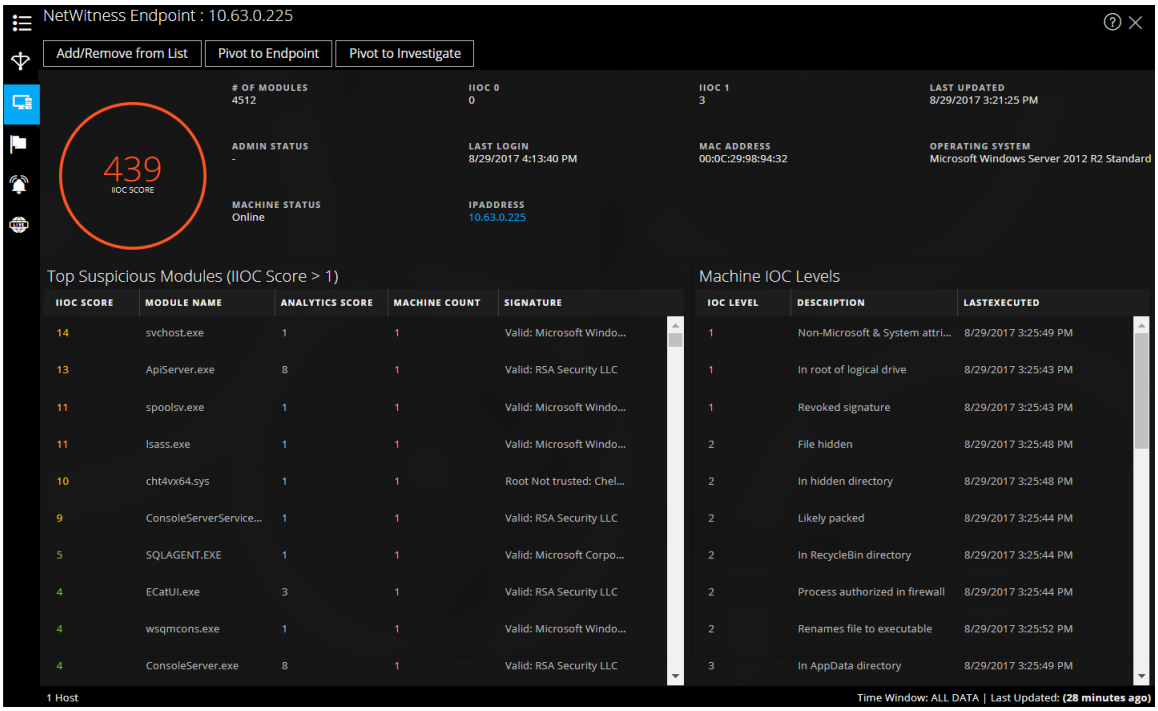
The following information is displayed for Active Directory.

Field	Description
Display Name	The name of the user.
Employee ID	The employee ID of the user.

Field	Description
Phone	The phone number of the user.
Email	The email ID of the user.
AD User ID	The unique identification of the user within an organization.
Job Title	The designation of the user.
Manager	The name of the user's manager.
Groups	The list of groups of which the user is a member.
Company	The name of the user's company.
Department	The department name to which the user belongs within the organization.
Location	The location of the user.
Last Logon	The time when the user logged into the system, only if the Global Catalogue is defined.
Last Logon TimeStamp	The time when the user logged into the system.
Distinguished Name	The unique name assigned to the user.
Count	The number of users.
Time Window	The time window based on the value that is set for the "Query Last" field in the Configure Data Source Settings dialog. By default, all data for Active Directory is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

NetWitness Endpoint Tab

The following figure is an example of the Context Lookup panel for NetWitness Endpoint.



The following information displayed for IIOCs.

Field	Description
# Of Modules	The number modules that are looked up.
Admin Status	The admin status (if any).
Last Updated	The time when the data was last refreshed.
Last Login	The time when the user last logged in.
MAC Address	The Machine MAC Address.
Operating System	The Version of the Operating System used by the NetWitness Endpoint machine.
Machine Status	The state of the module being viewed: Online, Offline, Active, or Inactive.
IP Address	The IP address of the specific module.

The following information is displayed for modules.

Field	Description
IIOC Score	A machine IIOC score is an aggregated score based on the module scores. This is based on the value set for Minimum IIOC Score field in the Context Hub Data Source Settings dialog. The default value for Minimum IIOC Score is 500. See "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.

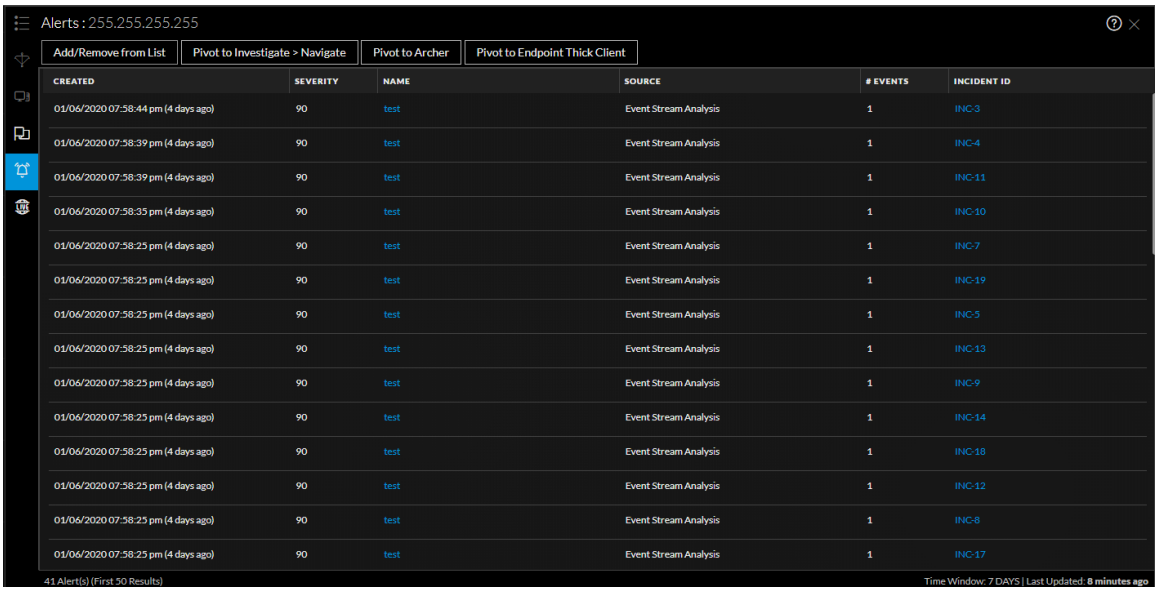
Field	Description
Module Name	The name of the module that is being looked up.
Analytic Score	The number of active files for the selected machine.
Machine Count	The number of machines on which that particular IOC got triggered.
Signature	Indicator of whether the file is signed or unsigned, valid or invalid, and signatory information. For example, Google, Apple, and so on.

The following information is displayed for machines.

Field	Description
IOC Levels	The IOC levels.
Description	The description for he IOC level if available.
Last executed	The time when the action was executed.
Count	The number of hosts that are being looked up.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, all data for NetWitness Endpoint is fetched.
Last Updated	The time when scan results were last updated in NetWitness Endpoint database.

Alerts Tab

The following figure is an example of Context Panel for Alerts that is displayed based on time first (Newest to Oldest) and then severity.

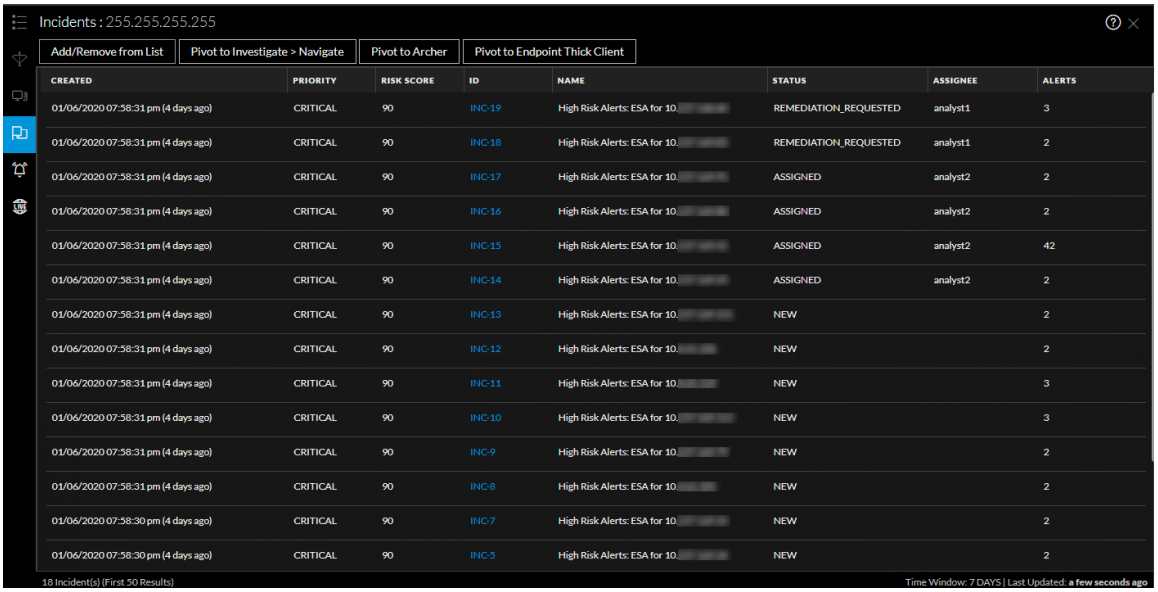


The following information is displayed in the Context Lookup panel for Alerts.

Field	Description
Created	The date and time when the alert was created.
Severity	The severity value of the alerts.
Name	The name of the alert. You can click the name to view the details of a specific alert.
Source	The alert source name from which the alert is triggered.
#Events	The number of events associated with the alert.
Incident ID	The ID of the incident (if any) with which the alert is associated. You can click the ID to view the details of a specific alert.
Count	The number of alerts. By default only the first 100 alerts are displayed. For more information on how to configure the settings, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, the alert data for last 7 days is fetched.
Last Updated	The time when contextual data was last fetched from data source.

Incidents Tab

The following figure is an example of the Context Panel for Incidents, which is based on time first (Newest to Oldest) and then priority status.

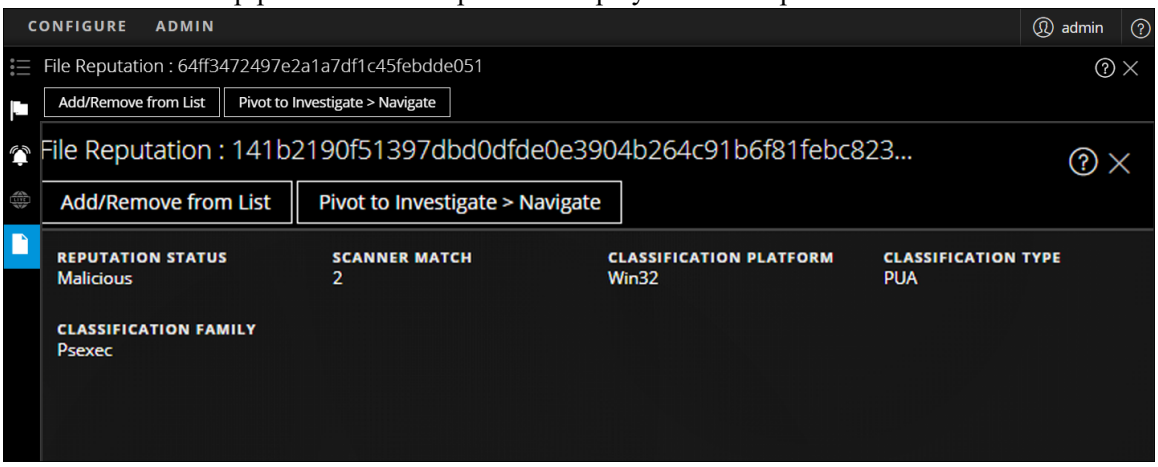


The following information is displayed in the Context Lookup panel for Incidents.

Field	Description
Created	The date when the incident was created.
Priority	The priority status of the incidents.
Risk Score	The risk score of the incidents.
ID	The Incident ID of the incident. You can click on the ID to display further details about the incident.
Name	The incident name.
Status	The status of the incident
Assignee	The current owner of the incident.
Alerts	The number of alerts associated with the incident.
Count	The number of incidents. By default only the first 100 incidents are displayed. For more information on how configure the settings, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, the alert data for last 7 days is fetched.
Last Updated	The time when contextual data was last fetched from data source.

File Reputation Tab

The Context Lookup panel for File Reputation displays the file reputation status of a file.

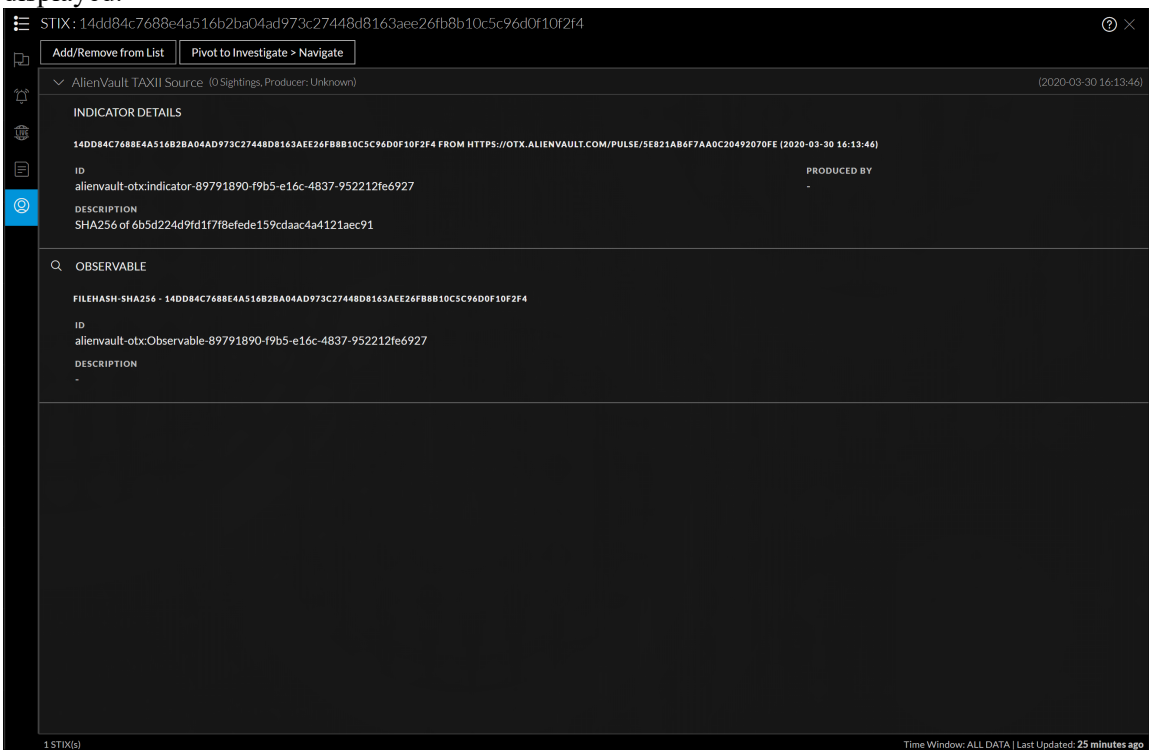


Field	Description
Reputation Status	Reputation Status of filehash. For more information about reputation status, see "View Reputation of files" in the *UEBA User Guide*.

Field	Description
Scanner Match	Number of scanners that detected malware or suspicious activity in the last scan.
Classification Platform	Classification for the queried filehash based on the platform. For example, the platform can be Win 32.
Classification Type	Classification for the queried filehash based on the type.
Classification Family	Classification for the queried filehash based on the malware family name.

TI Tab

The following figure is an example of a Context Panel for TI, and the table describes the information displayed.



Field	Description
Data Source name	Displays the STIX data source name from where the data is retrieved.
Timestamp	The time when the event was created.

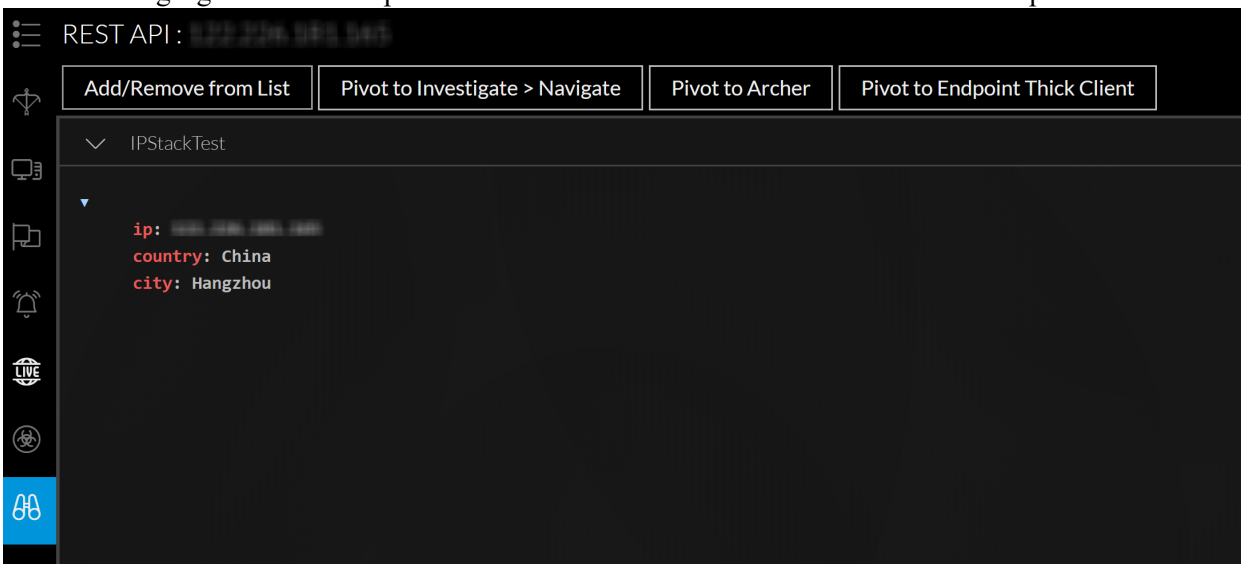
Field	Description
Indicator Details	Indicator Title: Displays the details that contains a pattern that can be used to detect suspicious or malicious cyber activity. ID: Displays the ID of the selected indicator. Produced by: Displays the user role who requested for the STIX data. Description: Displays details about the selected IP address which are being watch listed.
Observable	Observable Title: Displays and conveys information about cyber security related entities such as files, systems, and networks using the STIX Cyber-observable Objects (SCOs). ID: Displays the ID of the selected observable.
(Optional) SightingsREST	Sightings Title: Displays the name of the sighting source. Confidence: Displays the criticality of the sighting. Reference: Displays the reference URL of the sighting source.

REST API Tab

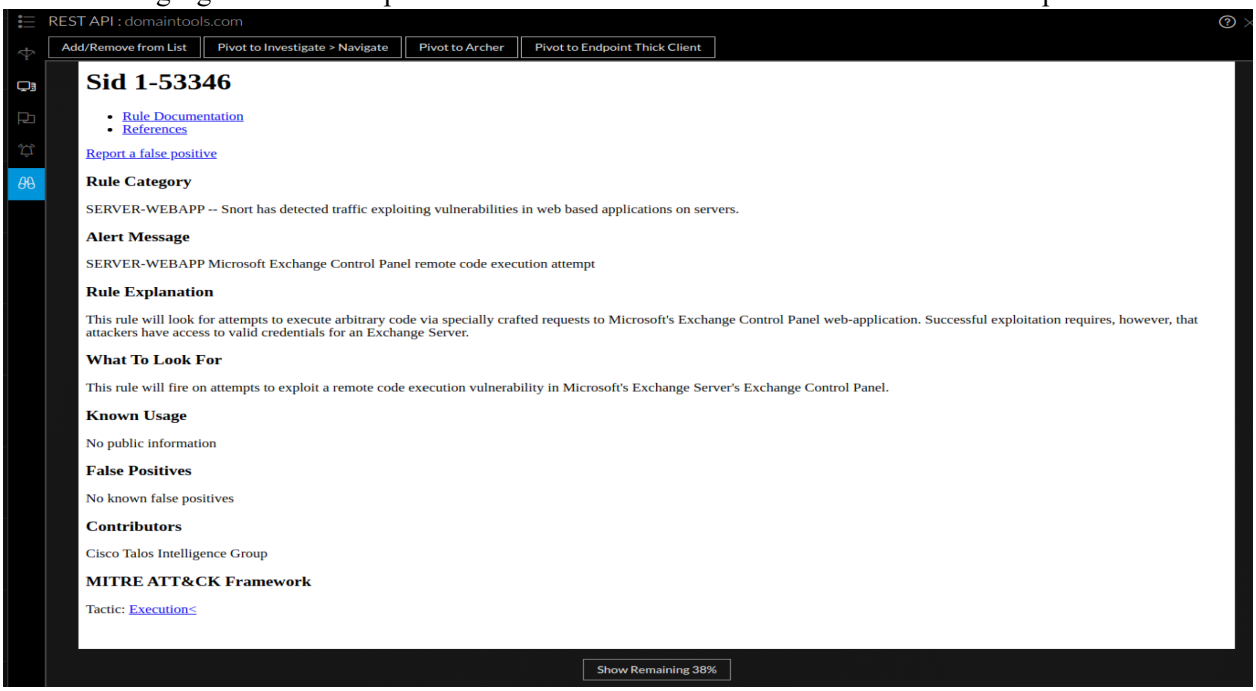
The Context Lookup panel for REST API shows HTML or JSON response (based on the response type configured) associated with the selected entity or meta value.

Note: For JSON response type, the fields that are mapped with friendly names (during REST API configuration) are only displayed for context Lookup. If you have not mapped any fields, all fields are displayed for context lookup.

The following figure is an example of the Context Panel for REST API with JSON response:



The following figure is an example of the Context Panel for REST API with HTML response:



NetWitness® Platform XDR

Version 12.0

RESTful API User Guide for 12.0

NETWITNESS

Platform XDR

NetWitness Core Services

NetWitness Core Services 1
Contact Information 30
Trademarks 30
License Agreement 30
Third-Party Licenses 30
Note on Encryption Technologies 30
Distribution 30
Overview 31
Current Version 31
Schema 31
HTTP Usage 32
HTTP Verbs 32
HTTP Status Codes 32
Service API Ports 33
Case Sensitive 33
Error Response 33
Authentication and Authorization 34
Authorization 34
Decoder service 35
Introduction 35
API Usage 35
Sample Request 36
Sample Response 36
HTTP Headers 36
JSON Response 36
/connections node 37
API Messages 37
closeAll 37
count 37
help 38

info 38
ls 38
netSpeed 39
/database node 39
API Messages 40
count 40
dbState 40
dump 40
hashInfo 41
help 42
info 42
ls 43
manifest 43
optimize 44
reconfig 44
resetMax 45
sizeRoll 46
stagger 47
timeRoll 48
wipe 49
/decoder node 49
API Messages 49
count 49
dpdk 49
help 50
info 50
ls 51
reconfig 51
reset 52
resetMax 53
select 53
sslKeys 53
start 54

stop 54
whoAgg 55
/decoder/parsers node 55
API Messages 55
attrib 55
content 56
count 56
delete 56
devices 57
disable 57
enable 57
export 58
feed 58
help 59
info 59
ls 59
reload 60
schema 60
snrtStat 61
upload 61
/index node 61
API Messages 61
cacheClr 61
count 62
help 62
info 62
inspect 62
language 63
ls 64
profile 64
reconfig 65
save 65
schema 65

sizeRoll 65
values 67
/logs node 67
API Messages 67
count 67
download 67
help 68
info 69
ls 69
pull 70
timeRoll 70
/rest node 71
API Messages 71
count 71
help 71
info 72
ls 72
/sdk node 73
API Messages 73
aliases 73
cancel 73
content 74
count 75
delCache 75
help 75
hierarch 76
info 76
keyrefs 76
language 76
ls 77
msearch 78
packets 79
pin 80

precache 80
query 81
reconfig 82
search 82
session 83
summary 83
timeline 84
validate 85
values 86
xforms 87
/sys node 87
API Messages 87
caCert 87
count 88
fileEdit 88
help 88
info 89
ls 89
peerCert 90
save 90
servCert 91
shutdown 91
statHist 91
telemtry 92
/users node 93
API Messages 93
addOrMod 93
auths 94
count 95
delete 95
help 95
info 96
ls 96

unlock 96
whoAmI 97
LogDecoder service 97
Introduction 97
API Usage 97
Sample Request 98
Sample Response 98
HTTP Headers 98
JSON Response 98
/collections node 99
API Messages 99
count 99
create 99
help 100
info 100
ls 100
mounts 101
/connections node 101
API Messages 101
closeAll 101
count 102
help 102
info 103
ls 103
netSpeed 103
/database node 104
API Messages 104
count 104
dbState 104
dump 105
hashInfo 106
help 106
info 107

ls 107
manifest 108
optimize 108
reconfig 109
resetMax 110
sizeRoll 110
stagger 112
timeRoll 112
wipe 113
/decoder node 114
API Messages 114
count 114
help 114
info 115
logStats 115
ls 116
reconfig 117
reparse 117
reset 118
resetMax 119
select 119
start 120
stop 120
whoAgg 120
/decoder/parsers node 121
API Messages 121
attrib 121
content 121
count 122
delete 122
devices 122
disable 123
dyndvmap 123

enable 124
export 124
feed 124
headers 125
help 125
info 126
ipdevice 126
iptmzone 127
ls 127
patterns 128
reload 129
schema 129
tzinfo 129
upload 131
/index node 131
API Messages 131
cacheClr 132
count 132
help 132
info 132
inspect 133
language 133
ls 134
profile 134
reconfig 135
save 135
schema 135
sizeRoll 136
values 137
/logs node 137
API Messages 137
count 137
download 137

help ... 139
info ... 139
ls ... 139
pull ... 140
timeRoll ... 141
/rest node ... 141
API Messages ... 141
count ... 142
help ... 142
info ... 142
ls ... 142
/sdk node ... 143
API Messages ... 143
aliases ... 143
cancel ... 143
content ... 144
count ... 145
delCache ... 145
help ... 145
hierarch ... 146
info ... 146
keyrefs ... 146
language ... 146
ls ... 147
msearch ... 148
packets ... 149
pin ... 150
precache ... 151
query ... 151
reconfig ... 152
search ... 153
session ... 153
summary ... 154

timeline 154
validate 155
values 156
xforms 158
/sys node 158
API Messages 158
caCert 158
count 158
fileEdit 159
help 159
info 160
ls 160
peerCert 160
save 161
servCert 161
shutdown 161
statHist 162
telemtry 163
/users node 164
API Messages 164
addOrMod 164
auths 165
count 165
delete 165
help 166
info 166
ls 167
unlock 167
whoAmI 167
Concentrator service 168
Introduction 168
API Usage 168
Sample Request 169

Using URL encoded request 169
Using special content type 169
Sample Response 170
HTTP Headers 170
JSON Response 170
/concentrator node 171
API Messages 171
add 171
count 172
delete 172
edit 172
help 173
info 174
ls 174
reconfig 174
reset 175
resetMax 176
start 176
status 176
stop 177
whoAgg 177
/connections node 177
API Messages 177
closeAll 177
count 178
help 178
info 179
ls 179
netSpeed 180
/database node 180
API Messages 180
count 180
dbState 180

dump 181
hashInfo 182
help 182
info 183
ls 183
manifest 184
optimize 184
reconfig 185
resetMax 186
sizeRoll 186
stagger 188
timeRoll 188
wipe 189
/index node 190
API Messages 190
cacheClr 190
count 190
help 190
info 191
inspect 191
language 192
ls 192
profile 193
reconfig 193
save 194
schema 194
sizeRoll 194
values 195
/logs node 196
API Messages 196
count 196
download 196
help 197

info 198
ls 198
pull 199
timeRoll 199
/rest node 200
API Messages 200
count 200
help 200
info 201
ls 201
/sdk node 202
API Messages 202
aliases 202
cancel 202
content 203
count 204
delCache 204
deviceId 204
help 205
hierarch 205
info 205
keyrefs 205
language 206
ls 207
msearch 207
packets 208
pin 209
precache 210
query 211
reconfig 212
search 212
session 213
summary 213

timeline 214
validate 215
values 215
xforms 217
/sys node 217
API Messages 217
caCert 217
count 218
fileEdit 218
help 218
info 219
ls 219
peerCert 220
save 220
servCert 220
shutdown 221
statHist 221
telemtry 222
/users node 223
API Messages 223
addOrMod 223
auths 224
count 224
delete 224
help 225
info 225
ls 226
unlock 226
whoAmI 226
Broker service 227
Introduction 227
API Usage 227
Sample Request 228

Using URL encoded request 228
Using special content type 228
Sample Response 228
HTTP Headers 229
JSON Response 229
Sample Request 230
Using special content type 230
Sample Response 230
HTTP Headers 230
JSON Response 230
/broker node 232
API Messages 232
add 232
count 233
delete 233
edit 233
help 234
info 235
ls 235
repair 236
reset 236
resetMax 237
start 237
status 237
stop 238
whoAgg 238
/connections node 238
API Messages 238
closeAll 238
count 239
help 239
info 240
ls 240

netSpeed 241
/index node 241
API Messages 241
count 241
gaps 241
help 242
info 242
ls 243
ranges 243
save 244
/logs node 244
API Messages 244
count 244
download 244
help 245
info 246
ls 246
pull 247
timeRoll 247
/rest node 248
API Messages 248
count 248
help 248
info 249
ls 249
/sdk node 250
API Messages 250
aliases 250
cancel 250
content 251
count 252
delCache 252
deviceId 252

help 253
hierarch 253
info 253
keyrefs 254
language 254
ls 255
msearch 255
packets 257
pin 258
precache 259
query 259
reconfig 261
search 261
session 262
summary 262
timeline 263
validate 264
values 265
xforms 267
/sys node 267
API Messages 267
caCert 267
count 267
fileEdit 268
help 268
info 269
ls 269
peerCert 269
save 270
servCert 270
shutdown 270
statHist 271
telemtry 272

/users node 273
API Messages 273
addOrMod 273
auths 274
count 274
delete 274
help 275
info 275
ls 276
unlock 276
whoAmI 276
Archiver service 276
Introduction 277
API Usage 277
Sample Request 278
Using URL encoded request 278
Using special content type 278
HTTP Headers 279
JSON Response 279
/archiver node 280
API Messages 280
add 280
count 281
delete 281
edit 282
help 283
info 283
ls 284
reconfig 284
repair 285
reset 285
resetMax 286
start 286

status 286
stop 287
whoAgg 287
/connections node 288
API Messages 288
closeAll 288
count 288
help 288
info 289
ls 289
netSpeed 290
/index node 290
API Messages 290
count 290
gaps 291
help 291
info 291
ls 292
ranges 292
save 293
/logs node 293
API Messages 293
count 293
download 293
help 294
info 295
ls 295
pull 296
timeRoll 296
/rest node 297
API Messages 297
count 297
help 297

info ... 298
ls ... 298
/sdk node ... 299
API Messages ... 299
aliases ... 299
cancel ... 299
content ... 300
count ... 301
delCache ... 301
deviceId ... 301
help ... 302
hierarch ... 302
info ... 302
keyrefs ... 303
language ... 303
ls ... 304
msearch ... 305
packets ... 306
pin ... 307
precache ... 308
query ... 309
reconfig ... 310
search ... 310
session ... 311
summary ... 311
timeline ... 312
validate ... 313
values ... 314
xforms ... 316
/sys node ... 316
API Messages ... 316
caCert ... 316
count ... 317

fileEdit 317
help 317
info 318
ls 318
peerCert 319
save 319
servCert 320
shutdown 320
statHist 320
telemtry 321
/users node 322
API Messages 322
addOrMod 322
auths 325
count 325
delete 325
help 325
info 326
ls 326
unlock 327
whoAmI 327
Appliance service 327
Introduction 327
API Usage 327
Sample Request 328
Sample Response 328
HTTP Headers 328
JSON Response 328
/appliance node 329
API Messages 329
addFSMon 329
clock 329
cntrlcnt 330

count 330
delFSMon 330
devlist 330
dpToNic 330
ethtool 331
file 331
fwdLogs 331
help 332
info 332
ls 332
nicToDp 333
partClr 333
partNew 334
powerOff 334
raidList 335
reboot 335
reset 335
serial 336
setNTP 336
setSNMP 337
srvList 337
start 337
stop 338
upgrade 338
vgs 339
vmstat 339
/connections node 340
API Messages 340
closeAll 340
count 341
help 341
info 341
ls 341

netSpeed 342
/logs node 342
API Messages 343
count 343
download 343
help 344
info 344
ls 345
pull 345
timeRoll 346
/rest node 347
API Messages 347
count 347
help 347
info 347
ls 348
/sys node 348
API Messages 348
caCert 348
count 349
fileEdit 349
help 349
info 350
ls 350
peerCert 351
save 351
servCert 351
shutdown 352
statHist 352
telemtry 353
/users node 354
API Messages 354
addOrMod 354

auths 357
count 357
delete 357
help 357
info 358
ls 358
unlock 359
whoAmI 359
Cloud service 359
Introduction 359
Device Registration 359
API Usage 360
/cloud node 360
API Messages 360
activate 360
add 360
count 361
delete 361
edit 362
help 363
info 364
ls 364
reconfig 364
refresh 365
resetMax 365
start 365
status 365
stop 366
/connections node 366
API Messages 366
closeAll 366
count 367
help 367

info 368
ls 368
netSpeed 369
/logs node 369
API Messages 369
count 369
download 369
help 370
info 371
ls 371
pull 372
timeRoll 372
/rest node 373
API Messages 373
count 373
help 373
info 374
ls 374
/sys node 375
API Messages 375
caCert 375
count 375
fileEdit 375
help 376
info 376
ls 377
peerCert 377
save 377
servCert 378
shutdown 378
statHist 378
telemtry 379
/users node 380

API Messages 380
addOrMod 380
auths 383
count 383
delete 383
help 384
info 384
ls 384
unlock 385
whoAmI 385
Manuals 385
/database 386
/database dbState 386
/database hashInfo 387
/database manifest 387
/database reconfig 388
/database sizeRoll 388
/database stagger 390
/database timeRoll 391
/decoder 392
/decoder reset 393
/decoder select 393
/decoder sslKeys 393
Premaster 394
TLS 1.3 Keys 395
Private Keys or PEM files 396
Key Management 397
Return Values 398
Viewing Unencrypted Traffic 398
/decoder start 399
/decoder stop 399
/decoder whoAgg 399
/sdk 399

/sdk content ... 400
Render Types ... 401
NetWitness Content Render Flags ... 403
RESTful API ... 404
/sdk hierarch ... 405
/sdk keyrefs ... 405
/sdk language ... 405
C API ... 406
Language Flags ... 407
NetWitness Language Result Flags ... 407
/sdk msearch ... 409
`msearch` Flags ... 410
`msearch` Index Search Mode ... 411
Text Search Syntax ... 411
Search Syntax And Index Modes ... 411
`msearch` Tips ... 412
/sdk packets ... 412
RESTful API ... 412
C API ... 414
/sdk pin ... 415
Miscellaneous Operations ... 415
pinId ... 416
/sdk query ... 416
`where` Clauses ... 419
Query Operators ... 420
Text Values ... 421
IP Addresses ... 421
MAC Addresses ... 421
Numeric Values ... 421
Bucketed Numeric Indexes ... 421
Numeric Value Aliases ... 422
Date and Time Expressions ... 422
Relative Time Points ... 422

Special Range Values 423
`group by` Clause (since 10.5) 424
`order by` Clause (since 10.5) 425
`search` parameter 427
/sdk search 428
/sdk session 428
/sdk summary 428
/sdk timeline 429
/sdk validate 431
/sdk values 431
Parameters 432
`values` Flags 434
`values` Call Example 435
Values call and bucketing mode 435
Suggestion Mode 435
`search` parameter 436
/sys 436
/sys caCert 437
/sys fileEdit 437
/sys peerCert 438
/sys save 439
/sys servCert 439
/sys shutdown 439
/sys statHist 439
Configuration of the Stats Database 444
/decoder/parsers 445
Configuration 445
Receiving Protobuf Encoded Logs 447
Messages 447
/decoder/parsers dyndvmap 447
Parameters and Results 448
Display Mapped Sources 448
Purge Mappings 448

Force Persistence 448
/decoder/parsers tzinfo 449
Parameters and Results 449
Version 449
Timezone Names 449
Timezone Abbreviations 450
Duplicates Abbreviations 450
Overridden Abbreviations 450
Add/Edit an Abbreviation Override 451
Remove an Abbreviation Override 451
/sdk deviceId 451
/appliance 452
Summary of Array Configuration Commands 452
Commands that deal with direct-attach RAID volumes 453
Commands that deal with allocating block devices as storage 453
Commands that deal with allocating storage to services 453
/appliance partClr 453
Uninitialize a block storage device 453
/appliance partNew 454
Initialize a block storage device 454
Usage 454
Example: Assign devices `sdd` and `sde` to Decoder. 454
Example: Assign device `sdc` to Archiver 455
Netwitness Service Volume Reference 455
/appliance vgs 458
List Volume Groups 458

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Overview

Netwitness Core Services API can be accessed through the host and port of each core service.The core services comprises Decoder, Concentrator, LogDecoder, Broker, Archiver and Appliance.
The API can be accessed via REST through their respective ports.

Current Version

All requests to the API will use latest version of API available with the service.

```
Netwitness-Version: 12.0.0
```

Schema

The core services API are structured in tree form which includes nodes and node messages, it has predictable service and node oriented URLs. The format of the URL consists of the following components:

```
https://<hostname or
IP>:<port>/[node1][/node2][...]?msg=<message
name>[&param1=value1][&param2=value2][...]
```

The API accepts input requests by using URL parameters and by POSTing application/json.

- The special content type, application/x-netwitness-string-params passes parameters as:
    - plain text in the format: param1=value1 param2="value \"2\"
    - `Note`: Quotes, as part of the value, must be preceded by the backslash \ character. Any character can be escaped in this manner, including the backslash itself `\\` .

The following are acceptable output formats:

- application/json
- text/plain
- text/xml
- text/html
- application/octet-stream for binary content

The content type that is returned can be controlled through the HTTP Accept header. It is possible to set the parameter force-content-type to one of the previous values.

Any fields containing human readable timestamps or dates would be in the format

```
YYYY-MM-DD HH:MM:SS
```

HTTP Usage

API tries to adhere as closely as possible to standard HTTP and REST conventions in its use of HTTP verbs and status codes.

> Code in the 2xx range indicates success.
>
> Code in the 4xx range indicates an error that failed for the given information supplied (e.g a required parameter was omitted, URL path doesn't exist etc..).
>
> Code in the 5xx range indicates an error with server (these are rare).

HTTP Verbs

Verb	Usage
`GET`	Used to retrieve resource. Along with message options on the URL `GET` can also be used to create, update and delete a resource.
`POST`	Used to create or upload a new resource.
`DELETE`	Used to delete an existing resource.

HTTP Status Codes

Stauts Code	Usage
`200 OK`	Request completed successfully.
`400 Bad Request`	The request was malformed. The response body will include an error providing further information.
`401 Unauthorized`	Similar to `403 Forbidden`, but specifically for use when authentication is required and has failed or has not yet been provided.

Stauts Code	Usage
`403 Forbidden`	The request was valid, but the server is refusing the action. The user might not have the necessary permissions for a resource.
`404 Not Found`	The requested resource does not exist.
`500 Internal Server Error`	An unexpected error has occurred. The response body will include a message providing further information.

Service API Ports

The following are default ports that each core service listen for REST requests.

Service	Port
Decoder	50104
LogDecoder	50102
Concentrator	50105
Broker	50103
Archiver	50108
Appliance	50106

Case Sensitive

All URLs, request parameters and JSON fields are case sensitive.

Error Response

Error response along with HTTP status code would be returned in the JSON format

Authentication and Authorization

All requests must include the username and password in the `Authentication` parameters or as `Authorization` header (base64 encoded).

An example `curl` Authentication would be one of the following:

- Using `username and password`

  ```
  curl -u username:password
  ```

- Using `Basic Authorization Header`

  ```
  curl -H 'Authorization: Basic <Base64 encoded username:password>'
  ```

Where:

- `username`: The username of the service account.
- `password`: The password of that account.

The API supports both HTTP & HTTPS protocols and for secure communication prefer to use HTTPS and certificates.

> This document refers `curl` commands to access REST API and to use ssl certificates for curl commands refer https://curl.se/docs/sslcerts.html
>
> (Optional) To configure custom certificates for the REST interface refer https://community.rsa.com document `RESTful API User Guide` section `(Optional) Configure Custom SSLCertificate for the REST Interface`

API access over HTTPS is secured using TLS, all credentials will be encrypted in transit.

Authorization

In order to make requests through the API, users must belong to groups associated with roles that have the access permissions, as well as any underlying permissions required to fulfill the request. For details on service user roles and permissions refer https://community.rsa.com section `Services Security View - Service User Roles and Permissions`

`Note:` Through this document the API message details also include required security.roles permissions for accessing the API message.

Decoder service

Introduction

Decoder captures packets, sessionizes them, parses sessions, generates meta and index, and then writes the results to disk. The results are written to a session database, meta database, packet database and index (query engine). Typically the index on a decoder is just time and no other fields are indexed, though they certainly can be.

The metadata generated from raw packet data capture is enriched with security context through session parsing.

API Usage

Decoder service provides API interface through its tree form as nodes, node messages and child nodes.

The following are few examples of the several API nodes and messages supported by Decoder:

- /decoder node messages:
    - `select` capture adapters
    - `start` and `stop` packet capture
    - `reconfig` decoder for optimal settings
    - `sslkeys` to push ssl crypto information for ssl decryption before session parsing
    - `/decoder/stats` node for decoder capture stats information
- /database node messages:
    - `reconfig` database for optimal database settings
    - `dbState` to get comprehensive information about current databases
    - `/database/stats` node for database stats information
- /sdk node messages:
    - `reconfig` sdk for optimal settings
    - `summary` for retrieving summary information from the databases

    - `query` for performing query against meta database
    - `/sdk/stats` node for queries active, pending and other sdk operations
- /sys node messagges:
    - `statHist` to retrieve historical stats from stats database
    - `/sys/stats` node for service and system stats information

Sample Request

The following is a sample request to get list of capture adapters on the decoder using select message:

Request path is `/decoder` and message is `msg=select`

```
$ curl 'https://decoder-host:50104/decoder?msg=select'  \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-www-form-urlencoded;charset=ISO-8859-1' \
   -u 'username:password'
```

Sample Response

HTTP Headers

```
HTTP/1.1 200 OK
Content-Length: 887
Connection: Keep-Alive
Pragma: no-cache
Expires: -1
Cache-Control: no-cache, no-store, must-revalidate
Content-Type: application/json
```

JSON Response

```
{
  "flags": 1073938433,
  "params": {
     "1": "packet_mmap_,em0 (bpf)",
     "2": "packet_mmap_,em1 (bpf)",
     "3": "packet_mmap_,lo (bpf)",
     "4": "packet_mmap_,ALL (bpf)",
     "5": "PFRINGZC,em0 (bpf)",
     "6": "PFRINGZC,em1 (bpf)",
     "7": "PFRINGZC,lo (bpf)",
     ...
```

```
  }
}
```

`Note`: The API Messages section for the path nodes provide more details on API messages available and any associated manual pages.

/connections node

API Messages

closeAll

Description: Closes all connections

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: native|rest|amqp}	Indicates the type of connection to close, either native, rest or amqp. Default is all types.
limit	uint32, optional	If non-zero, will stop closing connections after limit is reached. NOTE: It starts closing the most idle connections first.
idleLongerThan	string, optional	A time duration in HH:MM:SS or ##{d,h,m,s} format. Any connection that has been idle longer than the passed in duration is eligible to be closed. Connections that have seen activity within the duration window will not be closed. Examples: "01:05:10" (1 hour, 5 mins, 10 secs) or "65m" (65 minutes)
exceptMe	bool, optional	If true, will not close the connection that initiated the command.

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyz ABCDEFGHIJKLMNOPQRSTU VWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

netSpeed

Description: Network Speed Test

Security.roles: connections.manage,aggregate

Parameters:

Parameter	Type, Options	Description
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 1024 and 536870912, inclusive}	The size of each message sent, default is 64k
count	uint32, optional, {unsigned:The value must be between 1 and 536870912, inclusive}	The number of messages to send, default is 16k

/database node

Manual for /database

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

dbState

Description: Returns comprehensive information about the current database state or persists the current state to disk for fast reload

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: filelist|summary|save}	The operation to perform (filelist, summary or save)
type	string, {enum-any:The value must be one or more of the following: session|meta|packet}	The database(s) to operate on
options	string, optional, {enum-one:The value must be one of the following: bytes|pretty-print}	Output display options like 'bytes' or 'pretty-print'. Pretty print (the default) will convert all sizes from bytes to human readable quantities like MB or GB.

Manual for dbState

dump

Description: Dumps information out of the database in nwd formatted files

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to dump
type	string, optional, {enum-one:The value must be one of the following: db|nwd}	The dump type
source	string, optional, {enum-any:The value must be one or more of the following: s|m|p}	The types of data to dump, default is all
verbose	bool, optional	If true (default is false), dumps more information
file	string, optional	Optional filename to use for NWD type, otherwise filename is assumed to be _sessionid_.nwd
limit	uint32, optional	If assigned, will only output the number of metas/packets passed in. Otherwise will print everything.

hashInfo

Description: Retrieves hash information for database files that containing session/meta/packet objects for a set of sessions or date range.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma delimited list of sessions and session ranges to retrieve file hashes for.
beginDate	string, optional	The beginning of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.

Parameter	Type, Options	Description
endDate	string, optional	The end of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.
directories	string, optional	A semi-colon delimited list of additional directories to search for hash files. This may be used to account for changes to the hash.dir configuration.

Manual for hashInfo

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

manifest

Description: If a manifest directory is defined, it will allow operations on the manifest files (such as a time based query) for database files in cold storage.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: query|compress}	The operation to perform (defaults to query)
time1	date-time, optional	The beginning time (UTC) for matching offline database files

Parameter	Type, Options	Description
time2	date-time, optional	The ending time (UTC) for matching offline database files
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format that is returned (posix, simple), default is posix

Manual for manifest

optimize

Description: Runs a series of tests to determine the optimimal *.write.block.size based on configured drives and using data from existing databases.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: session|meta|packet}	The database types to optimize, default is all
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 104857600 and 107374182400, inclusive}	The amount of data to write for each test
compression	string, optional, {enum-one:The value must be one of the following: none|gzip|bzip2|lzma}	Determines if compression block sizes should be tested, default is none
compressionLevel	uint32, optional, {unsigned:The value must be between 0 and 9, inclusive}	Determines the compression level, 0-9, where 1 is fastest, 9 is the best compression and zero is a predetermined balance between speed and compression

reconfig

Description: Calculates new drive sizes and free space for the session, meta and/or packet directories. No directories are removed and the assumption is each directory is mounted on a separate filesystem and will only be used for storage of that database.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective *.dir and *.free.space.min settings, otherwise it will just output the calculations
type	string, optional, {enum-any:The value must be one or more of the following: session|meta|packet}	The database types to reconfigure, default is all
percent	uint32, optional, {unsigned:The value must be between 1 and 99, inclusive}	The drive percentage to use for the calculated size per directory, default is 95
op	string, optional, {enum-one:The value must be one of the following: normal|10g}	The operating mode, 'normal' or '10g'. Default is normal.

Manual for reconfig

resetMax

Description: Resets all max database stats or just the ones listed.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
stats	string, optional, {enum-any:The value must be one or more of the following: session.rate.max|meta.rate.max|packet.rate.max}	The stats to reset, default is all

sizeRoll

Description: Delete database files based on the total size of all databases (passed with 'type' parameter) or space remaining on shared volume(s). This command should not be used on databases that don't share storage.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-any:The value must be one or more of the following: session|meta|packet}	The databases to consider for removing the oldest data based on total size or space remaining
log	bool, optional	If true, will log all actions taken, otherwise all activity will happen silently. Default is true.
maxSize	size, optional	The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Hot tier.
maxSizeWarm	size, optional	The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Warm tier.
maxPercent	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Hot tier.
maxPercentWarm	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When

Parameter	Type, Options	Description
		exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Warm tier.
minFree	size, optional	The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once again exceeds target. This only applies to the Hot tier
minFreeWarm	size, optional	The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once again exceeds target. This only applies to the Warm tier

Manual for sizeRoll

stagger

Description: Staggers database files to optimize read/write performance across multiple volumes. Can be used after adding an empty mount point

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: session|meta|packet}	The databases files to stagger
dryRun	bool, optional	If true (the default), will only return a description of the operations that would be performed. If false, then the files

Parameter	Type, Options	Description
		will actually be moved to an optimal read/write pattern. Please turn off capture or aggregation before actually performing the stagger operation.

Manual for stagger

timeRoll

Description: Delete database files that exceed a given age

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-any:The value must be one or more of the following: session|meta|packet}	The database files to remove
timeCalc	string, optional, {enum-one:The value must be one of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of days
date	string, optional	Remove database files older than the given UTC date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

Manual for timeRoll

wipe

Description: Overwrites all packets and/or meta for a session with a pattern (for eliminating sensitive information). Meta keys sessionid, time and size always remain untouched.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id whose packets will be wiped
payloadOnly	bool, optional	If true (default), will only overwrite the packet payload
pattern	string, optional	The pattern to use, by default it uses all zeros
metaList	string, optional	Comma separated list of meta to wipe, default (empty) is all meta
source	string, optional, {enum-any:The value must be one or more of the following: m|p}	The types of data to wipe, meta and/or packets, default is just packets

/decoder node

Manual for /decoder

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

dpdk

Description: perform DPDK administration

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
migrate	string, {not-empty:Value cannot be empty}	perform migration of existing capture configuration to DPDK:
commit	bool, optional	commit changes

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

reconfig

Description: Calculates optimal settings for decoder pools and buffers based on the installed hardware.

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false,on,off}	If true (default is false), will automatically update the respective decoder settings, otherwise it will just output the calculations
op	string, optional, {enum-one:The value must be one of the following: normal|10g|ndr}	The operating mode, 'normal', '10g', or 'ndr'. Default is normal.

Parameter	Type, Options	Description
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

Manual for reconfig

reset

Description: Reset data, index, manifests, stats, configuration, or logs for this service. Data automatically deletes index and stats, unless filesCreatedAfter is specified. Service is automatically restarted. Example arguments: data=1 config=1 log=1 This example will reset data, index, logs, and configuration index=1 This example will reset the index only manifest=1 This example will delete everything in the manifest directory filesCreatedAfter="2015-12-01 14:00:00" This example will delete all session, meta and packet files created on or after Dec 1st, 2015 2pm (UTC) from the service. All other files will remain. The index will not be touched, but upon restart will be truncated to match the last session in the session database.

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
data	bool, optional	Reset data, automatically resets index (may not be applicable to all services)
index	bool, optional	Reset index, only valid if data is not reset (may not be applicable to all services)
manifest	bool, optional	Reset all manifests stored in the long term manifest directory (may not be applicable to all services)
config	bool, optional	Reset configuration settings to default
stats	bool, optional	Reset the stats database
log	bool, optional	Reset the log database

Parameter	Type, Options	Description
filesCreatedAfter	date-time, optional	Delete all database files (session, meta, packets) created after a certain date. The index will automatically get rolled back on restart. Not valid with option index.
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart
force	bool, optional	Force reset without confirmation

Manual for reset

resetMax

Description: Resets all max stats to zero.

Security.roles: decoder.manage

select

Description: Selected a new capture device

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
adapter	string, optional	The new adapter. If ommitted will display a list of available adpaters. You may specify a comma-separated list of adapter indexes to enable multi-interface capture.

Manual for select

sslKeys

Description: Push SSL crypto information to enable SSL decryption of a session's packets prior to parsing

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
clear	bool, optional	Clears all existing keys from storage. Cannot be used with any other parameters.
maxKeys	uint32, optional	Sets the total number of keys that can be held in memory before aging out begins. Cannot be used with any other parameters.
counts	bool, optional	Returns key counts for sslKeys. Cannot be used with any other parameters.
random	string, optional	Adds the random that identifies the session key exchange.
premaster	string, optional	Adds the encryption key for the random.
pemFilename	string, optional	Gives the name of a PEM file that is the binary part of the request. This should be a private key that can be used for decrypting SSL traffic.
listPems	bool, optional	Returns the list of all installed PEM files used for decryption.
deletePem	string, optional	Given the PEM name, deletes it from the service. Can be repeated.

Manual for sslKeys

start

Description: Starts aggregation

Security.roles: decoder.manage

Manual for start

stop

Description: Stops aggregation

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
flush	bool, optional	If true (default), process all sessions and packets currently in memory. Otherwise, stop immediately, discarding unprocessed sessions and packets.

Manual for stop

whoAgg

Description: Returns information on who is aggregating from this service

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
clean	bool, optional	If true, will delete any aggregation trackers that no longer have a valid channel

Manual for whoAgg

/decoder/parsers node

API Messages

attrib

Description: Manage attributes for content files.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|update}	The operation to perform. If list (default), return a list of attributes for specified type of contents. If update, set the attributes for specified content file.

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices).
file	string, optional	The content file name.
values	string, optional	List of attributes to set for the specified content file in name value pairs.

content

Description: Get parsing content file information.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
file	string, optional	The content filename (type must be specified)
type	string, optional, {enum-one:The value must be one of the following: feeds|parsers}	The content type
uuid	string, optional	The content uuid

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete parsing content files.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
file	string, optional	The content filename (type must be specified)
type	string, optional, {enum-one:The value must be one of the following: feeds|parsers}	The content type (feeds|parsers)
uuid	string, optional	The content uuid

devices

Description: Returns the log parsers

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
uuid	string, optional	The content uuid

disable

Description: Disable the specified content.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices)
uuid	string	The content uuid

enable

Description: Enable the specified content.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices)
uuid	string	The content uuid

export

Description: Returns the content as a package.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices)
uuid	string, optional	The content uuid
file	string, optional	The content file name

feed

Description: Manages the feed parsers

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: notify|reload|remove|delete}	The feed operation to perform
file	string, optional	The feed filename

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

reload

Description: Reloads all meta parsers

Security.roles: parsers.manage

schema

Description: Returns the parsers meta schema

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
onlyShowEnabled	bool, optional	If true, only return information on parsers that are currently enabled
showRegisteredMeta	bool, optional	If true (the default), show the registered meta for each parser
showOptions	bool, optional	If true, shows the options supported by each parser. Passing true changes the XML structure returned so each parser (potentially) has

Parameter	Type, Options	Description
prettyPrint	bool, optional	If true, returns the XML formatted to be easier to read. Default is false.

snrtStat

Description: Get Snort rule statistics

Security.roles: parsers.manage

upload

Description: Upload feed/parser/geoip2 files to this service

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: start|attributes|cancel|finished}	The operation to perform (start, cancel, finished)
size	uint64, optional	The size of the incoming file
filename	string, optional	The name of the incoming file
reload	uint32, optional	Reload all parsers when upload is complete for the appropriate file types (default). A setting of 0 when sending a finished op will disable the reload.

/index node

API Messages

cacheClr

Description: Clear index internal caches

Security.roles: index.manage

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

inspect

Description: Provides overall information about the index or can inspect specific keys and values

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string, optional	The key to inspect
value	string, optional	The value to inspect
format	string, optional, {enum-one:The value must be one of the following: Text|IPv4|Int8|UInt8|Int16|UInt16|Int32|UInt32|Int64|UInt64|UInt128|Float32|Float64|TimeT|DayOfWeek|HourOfDay|Binary|IPv6|MAC}	Only include keys that have this format
valueCount	uint64, optional	The number of values to return from each key
allSlices	bool, optional	Return information on all slices, instead of just most recently created slice
summarizeAllValues	bool, optional	Calculate total summary and page usage from all values in each slice

language

Description: Returns language information

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
values	bool, optional	If set to 1, will describe the value overrides defined in the language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

profile

Description: Returns information about the index used by the Index Profile tool

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string, optional	The key to profile
value	string, optional	The value to profile
max	uint32, optional	Max number of pages to return
flag	uint32, optional	Bitwise option when key is specified. 1: Break report to values. 2: List page sizes. 4: Categorize page counts in different

Parameter	Type, Options	Description
		range. 8: Sessions count Per page with each algorithm. 16: Verify session page

reconfig

Description: Calculates new sizes for the summary and page directories. NOTE: Any size changes require a reset of the index before they take effect.

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective config node settings, otherwise it will just output the calculations
percent	uint32, optional, {unsigned:The value must be between 1 and 99, inclusive}	The drive percentage to use for the index database, default is 90

save

Description: Saves the index to disk

Security.roles: index.manage

schema

Description: Introspects all open Meta DB files for meta keys and merges those keys with the defined language to return a fairly comprehensive list of all meta keys in existence.

Security.roles: index.manage

sizeRoll

Description: Delete index and other DB files (see 'type' param) based on the total size or space remaining on shared volumes. This command is not intended to be used on databases that do not share storage.

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
log	bool, optional	If true, will log all actions taken, otherwise all activity will happen silently. Default is true.
type	string, optional, {enum-any:The value must be one or more of the following: session|meta|packet}	The databases to consider for removing the oldest data based on total size or space remaining
maxSize	size, optional	The maximum size of the index hot storage. When exceeded, oldest data is deleted first until total size is less than maxSize.
maxSizeWarm	size, optional	The maximum size of the index warm storage. When exceeded, oldest data is delete first until total size is less that maxSizeWarm
maxPercent	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This parameter applies to hot storage only.
maxPercentWarm	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This parameter applies to warm storage only.
minFree	size, optional	The minimum allowed free space on the volumes before oldest data is deleted. This

Parameter	Type, Options	Description
		parameter applies to hot storage only.
minFreeWarm	size, optional	The minimum allowed free space on the volumes before oldest data is deleted. This parameter applies to warm storage only.

values

Description: Returns the top n values for the given key

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string	The index key for values
top	uint32, {unsigned:The value must be between 0 and 10000, inclusive}	The number of top entries to return

/logs node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

download

Description: Downloads log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting log number to download
id2	uint64, optional	The ending log number to download
time1	date-time, optional	The starting log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
time2	date-time, optional	The ending log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
op	string, {enum-one:The value must be one of the following: start|next|cancel}	The operation to perform (start, next, cancel)
logTypes	string, optional	The type(s) of log messages to retrieve (debug,info,audit,warning,failure)
match	string, optional	Case insensitive string to match in each log message
regex	string, optional	Regular expression to match in each log message
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)
batchSize	uint32, optional, {unsigned:The value must be between 2000 and 10000, inclusive}	The number of logs to retrieve per batch, default is 2000

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is

whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number

Parameter	Type, Options	Description
		(bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

pull

Description: Downloads N log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The first log id number to retrieve, this is mutually exclusive with id2
id2	uint64, optional	The last log id number that will be sent, defaults to most recent log message when id1 or id2 is not sent
count	uint32, optional, {unsigned:The value must be between 1 and 10000, inclusive}	The number of logs to pull
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)

timeRoll

Description: Delete log entries that exceed a given age

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
timeCalc	string, optional, {enum-any:The value must be one or more of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 60, inclusive}	Remove log entries older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 24, inclusive}	Remove log entries older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 365, inclusive}	Remove log entries older than the given number of days
date	string, optional	Remove log entries older than the given date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

/rest node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

/sdk node

Manual for /sdk

API Messages

aliases

Description: Retrieves aliases for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string	The language key to retrieve aliases for
value	string, optional	An optional value to return a specific alias for

cancel

Description: Cancel a running query, regardless of the user

Security.roles: sdk.manage,sdk.meta

Parameters:

Parameter	Type, Options	Description
handle	uint32	The channel handle running the query
uuid	string, optional	The device UUID if taking a device offline instead of canceling the

Parameter	Type, Options	Description
		entire query (Broker only). Passing uuid will propagate the request upstream to find the first matching device to cancel.

content

Description: Returns the packet content for a session

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session id to retrieve packet content for
packet	uint64, optional	The packet id of the first packet of the session, in case the session has rolled out
pinId	string, optional	The Pin ID of the session to retrieve from long term storage. If provided, the packet and session parameters must not be provided.
maxSize	size, optional	The max number of bytes to return, zero means no limit
failOnCacheMiss	bool, optional	If true, the content call fails if the NWD is not in the cache. This is used by optimized caching code only.
renderType	string, optional	The render type (see NwContentTypes in NwSDK.h), pass zero to pick best option. Not specifying this parameter always results in an NWD and renderFlags and renderOptions are not considered. The following values are also supported: 'pcap', 'file-list-json', 'session-meta-file-list-json', 'nwd', 'log', 'files'
renderFlags	uint32, optional	Bitwise mask to control options, (see NwContentFlags in NwSDK.h)

Parameter	Type, Options	Description
renderOptions	string, optional	An encoded string params containing additional options for controlling the operation.

Manual for content

count

Description: Returns the number of child nodes

Security.roles: everyone

delCache

Description: Removes all cached NWD files from either the normal NWD cache or the pin cache

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-one:The value must be one of the following: nwd|pin|all}	Specifies which cache to delete all NWD files from. The default is 'nwd'. Passing 'all' will delete from all caches.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

hierarch

Description: Report the set of devices attached to this service

Security.roles: sdk.meta

Manual for hierarch

info

Description: Returns detailed information about the node

Security.roles: everyone

keyrefs

Description: Retrieves the set of entity keys

Security.roles: sdk.meta

Manual for keyrefs

language

Description: Retrieves the language definition

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting language id
id2	uint64, optional	The ending language id
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
options	string, optional	Extra options, currently unused
flags	string, optional	Optional flags to configure how the results are returned. Can be a number (bitwise mask) or comma separated values like default, no-count, quick-count or full-count.
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1

Parameter	Type, Options	Description
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

msearch

Description: Search for pattern matches in many sessions or packets

Security.roles: sdk.content & sdk.meta

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	The session ID ranges to search
packets	string, optional	The packet ID ranges to search
search	string	The search string to use
where	string, optional	The where clause used to identify sessions to consider for the search
limit	uint64, optional	The maximum number of sessions to traverse for this search
size	uint64, optional	The maximum number of results to return for this search
flags	string, {enum-any:The value must be one or more of the following: regex|sp|sm|ci|pre|post|ds|precache|si}	Flags to use for search. This is a comma separated list of one or more of the flag values: regex,sp,sm,ci,pre,post,ds,precache,si

Parameter	Type, Options	Description
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for msearch

packets

Description: Stream packets back based on the input parameters provided

Security.roles: sdk.content & sdk.packets

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: setup|start|cancel|processed}	The operation to perform (start, cancel, processed)
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) whose packets will be streamed back.
packets	string, optional	A comma separated list of packet ids or session/packet ids (#&#) that will be streamed back.
where	string, optional	Where clause which will be evaluated to determine which sessions to stream back
time1	date-time, optional	A starting time range (UTC) where matching packets will be streamed back ("2010-Apr-20 09:00:00")
time2	date-time, optional	An ending time range (UTC) where matching packets will be streamed back ("2010-Apr-20 10:00:00")
flags	uint32, optional	Additional flags as defined by the NwPackets SDK function

Parameter	Type, Options	Description
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for packets

pin

Description: Pin a session in the long term cache so it can be retrieved even after the meta or packets have rolled out. To retrieve, you must make a content call with the Pin ID that is returned from this command. The op={unpin,validate} commands all take one or more comma separated Pin IDs in the pinId parameter and returns the status for the Pin IDs that were recognized by a service.

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session ID to pin. This cannot be used with the 'sessions' parameter.
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) to pin. This cannot be used with the 'session' parameter.
pinId	string, optional	One or more Pin IDs, separated by commas
op	string, optional, {enum-one:The value must be one of the following: ls|unpin|validate}	Various operations that can be performed. Will be submitted to upstream devices if not configured for long term cache.

Manual for pin

precache

Description: Efficiently caches NWD files for future retrieval via the content message

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) that will be saved in the NWD cache directory
packets	string, optional	A comma separated list of session and packet ids (#p#) that will be saved in the NWD cache directory
waitOnComplete	bool, optional	If true, will wait for precache to complete before returning a response.

query

Description: Performs a query against the meta database

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id (to run the query from most recent to oldest meta, make id1 larger than id2)
id2	uint64, optional	The ending meta id
size	uint32, optional	The max number of entries to return, or just stream back all results if zero
query	string, optional	The query string to use
flags	string, optional	The flags to use for query. Can be a number (bitwise mask) or comma separated values like query-log.
threshold	uint64, optional	Query optimization to stop processing results after the threshold is reached (useful with select aggregate functions). Zero means no threshold (the default).
partition	uint32, optional	Used to partition the result set across multiple clients. Must be greater than

Parameter	Type, Options	Description
		zero and less than or equal to totalPartitions
totalPartitions	uint32, optional	The total number of clients that will be submitting the same query for partitioning.
search	string, optional	A text search clause to apply to the where clause of this query
streamLimit	uint32, optional	Limit the number of sessions returned by a streaming query
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for query

reconfig

Description: Calculates default values for some of the config nodes.

Security.roles: sdk.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the config nodes with the defaults
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

search

Description: Searches for matches in session/packet content

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to search
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
search	string	The search string to use

Manual for search

session

Description: Retrieves the meta id range for the session range

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64	The starting session id (inclusive). Pass zero to scope to lowest valid id.
id2	uint64	The ending session id (inclusive). Pass zero to scope to highest valid id.

Manual for session

summary

Description: Retrieves summary information from the databases

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
flags	string, optional	Optional SDK flags. Can be a number (bitwise mask) or comma separated values like default or ignore-cache.

Manual for summary

timeline

Description: Returns the count of sessions/size/packets in discrete time intervals

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS), if greater than time2, then the whole time range is returned
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
timezone	int32, optional	The timezone of the local computer to receive the data, hour offset from GMT [-13 to 13]
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for the timeline. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, order-ascending, order-descending, ignore-cache or clear-cache.
where	string, optional	Optional where clause for filtering the data
bookendHours	uint32, optional	The number of hours outside the given time1, time2 which will become part of the cache. Default is 24 hours.

Parameter	Type, Options	Description
fieldName	string, optional	The time field to retrieve values for, defaults to 'time'
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query

Manual for timeline

validate

Description: Validate a query or values call without execution

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: query|values}	The operation to validate
validateLanguage	bool, optional	If true (default is false), language tokens and syntax will be validated, otherwise just syntax will be validated
fieldName	string, optional	Required if validating a 'values' operation. The field to retrieve values for.
where	string, optional	Required if validating a 'values' operation. The filter criteria to be validated.
query	string, optional	Required if validating a 'query' operation. The query string to be validated.
json	bool, optional	Structure response in JSON syntax

Manual for validate

values

Description: Performs a value count query and returns the matching values for a report

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id
id2	uint64, optional	The ending meta id
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for values. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, aggregate, sort-total, sort-value, order-ascending, order-descending, ignore-cache, clear-cache, or database-scan.
threshold	uint64, optional	Query optimization to stop processing large session counts
fieldName	string	The field to retrieve values for
where	string, optional	The filter criteria for the values
aggregateFunction	string, optional, {enum-one:The value must be one of the following: count|sum}	The aggregate operation to be applied where the aggregate flag is set
aggregateFieldName	string, optional	The meta field to aggregate in the aggregateFunction

Parameter	Type, Options	Description
min	string, optional	The lower limit of values to return. Only values greater than this value will be returned.
max	string, optional	The upper limit of values to return. Only values less than this value will be returned.
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query

Manual for values

xforms

Description: Retrieves transforms for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string, optional	An optional language key to retrieve transform for, all transforms returned if not provided

/sys node

Manual for /sys

API Messages

caCert

Description: Display or delete trusted CA certs

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: list|delete|add}	CA action to perform (list,delete,add)
id	string, optional	ID of CA certificate to remove when op=delete

Manual for caCert

count

Description: Returns the number of child nodes

Security.roles: everyone

fileEdit

Description: View and edit system configuration files

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: dir|get|getBackup|put|create|delete}	The operation to perform (dir, get, getBackup, put, create, delete)
type	string, optional, {enum-one:The value must be one of the following: parser|filter}	The type of file, if not provided will check file extension
filename	string, optional, {length:The value length must be between 1 and 64, inclusive}	File to create, edit, get or delete

Manual for fileEdit

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1

Parameter	Type, Options	Description
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

peerCert

Description: Display or modify trusted peer certificates

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|delete|add}	Operation to perform on the list of trusted peers (list,delete,add)
id	string, optional	ID of peer certificate to remove when op=delete

Manual for peerCert

save

Description: Forces a save of system settings

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
force	uint32, optional	If 1, forces the configuration file to saved, regardless if changes are detected, default is 0

Manual for save

servCert

Description: Display the current server certificate in PEM format.

Security.roles: sys.manage

Manual for servCert

shutdown

Description: Stop the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
reason	string, optional, {length:The value length must be between 0 and 512, inclusive}	The reason for shutting down the service
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart

Manual for shutdown

statHist

Description: Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time for retrieving stats

Parameter	Type, Options	Description
time2	date-time, optional	The ending time for retrieving stats
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format for each stat snapshot, default is posix (seconds since 1970)
include	string, optional	Comma separated list of stats to include (wildcards allowed)
exclude	string, optional	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)
onChange	string, optional	Comma separated list of stats that reduce the result set to only when at least one of the stats change from their previous value
reduce	bool, optional	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result
showAll	bool, optional	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Manual for statHist

telemtry

Description: Returns telemetry information about the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
level	uint32, optional, {signed:The value	The level of information to be sent by the telemetry request. A value of 1 (default) will send all stat nodes, a

Parameter	Type, Options	Description
	must be between 1 and 3, inclusive}	value of 2 will send stat and config nodes, and a value of 3 will send stat nodes, config nodes and usage metrics for rules, parsers and devices (where applicable).
include	string, optional	Comma separated list of nodes to include (wildcard usage is ok).
exclude	string, optional	Comma separated list of nodes to exclude (wildcard usage is ok). If "include" is also specified, exclusions are applied to the set of nodes specified by "include".
url	string, optional	Endpoint URL to send the JSON response.
op	string, optional, {enum-one:The value must be one of the following: get|configure}	The telemetry operation to execute. Default is 'get'.

/users node

API Messages

addOrMod

Description: Add a new user or update an existing user in the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to add or update, must be alphanumeric or @!#$%&'+-=?^_`{|}~.[]\
password	string, optional	The user's password

Parameter	Type, Options	Description
pwdIsHashed	bool, optional	Set to true if the password has already been hashed
groups	string, optional	The groups the user belongs to
authType	string, optional, {enum-one:The value must be one of the following: netwitness|pam}	The authentication system to use, "netwitness" is the default
queryTimeout	uint32, optional, {unsigned:The value must be between 0 and 43200, inclusive}	The maximum number of minutes that a query is allowed to execute for this user. Zero means unlimited.
displayName	string, optional	Display name for this user account
email	string, optional	Email address for this account
description	string, optional	Description of this user account
queryPrefix	string, optional	Query filter applied to every query performed by this user account
sessionThreshold	uint32, optional, {unsigned:The value must be between 0 and 4294967295, inclusive}	Query optimization which will extrapolate the remaining session counts when they exceed this value

auths

Description: Get supported authentication types

Security.roles: everyone

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete a user from the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to delete

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following:	The specific help operation to perform (e.g., op=manual would

Parameter	Type, Options	Description
	messages|parameters|description|values|roles|extra|manual}	return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

unlock

Description: Unlock a locked out account

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string	The username or account to unlock

whoAmI

Description: Returns information about the authenticated user

Security.roles: everyone

LogDecoder service

Introduction

LogDecoder captures logs, parses the logs into meta, indexes them and writes the results to disk. It also has a session, meta and packet database (which is really the log message) and index. Similar to a Decoder the LogDecoder also has minimal index with fewer fields along with time.

The metadata generated from raw log capture is enriched with security context through session parsing.

API Usage

LogDecoder service provides API interface through its tree form as nodes, node messages and child nodes.

The following are few examples of the several API nodes and messages supported by LogDecoder:

- /decoder node messages:
    - `start` and `stop` logs capture
    - `reconfig` logdecoder for optimal settings
    - `/decoder/stats` node for logdecoder capture stats information
- /decoder/parsers node messages:
    - `content` to get parsing file contents
    - `devices` list of log parser devices
    - `ipdevice` map IP to Device type in log parsing

    - `schema` log parsers meta schema
- /database node messages:
    - `reconfig` database for optimal database settings
    - `dbState` to get comprehensive information about current databases
    - `/database/stats` node for database stats information
- /sdk node messages:
    - `reconfig` sdk for optimal settings
    - `summary` for retrieving summary information from the databases
    - `query` for performing query against meta database
    - `/sdk/stats` node for queries active, pending and other sdk operations
- /sys node messagges:
    - `statHist` to retrieve historical stats from stats database.
    - `/sys/stats` node for service and system stats information

Sample Request

The following is sample request to describe ip device mappings on the logdecoder.

Request path is `/decoder/parsers`, message is `ipdevice` and parameters are `op=describe` .

```
$ curl 'https://logdecoder-host:50102/decoder/parsers?msg=ipdevice&op=describe'  \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-www-form-urlencoded;charset=ISO-8859-1' \
   -u 'username:password'
```

Sample Response

HTTP Headers

```
HTTP/1.1 200 OK
Content-Length: 278
Connection: Keep-Alive
Pragma: no-cache
Expires: -1
Cache-Control: no-cache, no-store, must-revalidate
Content-Type: application/json
```

JSON Response

The sample response here is JSON encoded ip-device mapping schema in xml format. It is possible for response values to contain xml content for certain requests.


```
{
 "flags": 1073872897,
 "string": "<addressmap-schema>
  <DeviceEntry ipv4=\"X.X.X.X\" device=\"ciscomeraki\" lastUpdated=\"1586804116\"
soft=\"false\"/>
  <DeviceEntry ipv6=\"::1\" device=\"ciscomeraki\" lastUpdated=\"1596696544\"
soft=\"true\"/>
  </addressmap-schema>"
}
```


`Note`: The API Messages section for the path nodes provide more details on API messages available and any associated manual pages.

/collections node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

create

Description: Create an empty collection

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
name	string, {collection-name:Collection names must be valid path names}	The name of the collection to be created.
path	string, optional, {not-empty:Value cannot be empty}	The path of the collection. All database and index folders will be created under this path.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

mounts

Description: Returns the common mount points for all collections with disk statistics

Security.roles: everyone

/connections node

API Messages

closeAll

Description: Closes all connections

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: native|rest|amqp}	Indicates the type of connection to close, either native, rest or amqp. Default is all types.
limit	uint32, optional	If non-zero, will stop closing connections after limit is reached.

Parameter	Type, Options	Description
		NOTE: It starts closing the most idle connections first.
idleLongerThan	string, optional	A time duration in HH:MM:SS or ##{d,h,m,s} format. Any connection that has been idle longer than the passed in duration is eligible to be closed. Connections that have seen activity within the duration window will not be closed. Examples: "01:05:10" (1 hour, 5 mins, 10 secs) or "65m" (65 minutes)
exceptMe	bool, optional	If true, will not close the connection that initiated the command.

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

netSpeed

Description: Network Speed Test

Security.roles: connections.manage,aggregate

Parameters:

Parameter	Type, Options	Description
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 1024 and 536870912, inclusive}	The size of each message sent, default is 64k
count	uint32, optional, {unsigned:The value must be between 1 and 536870912, inclusive}	The number of messages to send, default is 16k

/database node

Manual for /database

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

dbState

Description: Returns comprehensive information about the current database state or persists the current state to disk for fast reload

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: filelist|summary|save}	The operation to perform (filelist, summary or save)

Parameter	Type, Options	Description
type	string, {enum-any:The value must be one or more of the following: session|meta|packet}	The database(s) to operate on
options	string, optional, {enum-one:The value must be one of the following: bytes|pretty-print}	Output display options like 'bytes' or 'pretty-print'. Pretty print (the default) will convert all sizes from bytes to human readable quantities like MB or GB.

Manual for dbState

dump

Description: Dumps information out of the database in nwd formatted files

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to dump
type	string, optional, {enum-one:The value must be one of the following: db|nwd}	The dump type
source	string, optional, {enum-any:The value must be one or more of the following: s|m|p}	The types of data to dump, default is all
verbose	bool, optional	If true (default is false), dumps more information
file	string, optional	Optional filename to use for NWD type, otherwise filename is assumed to be _sessionid_.nwd
limit	uint32, optional	If assigned, will only output the number of metas/packets passed in. Otherwise will print everything.

hashInfo

Description: Retrieves hash information for database files that containing session/meta/packet objects for a set of sessions or date range.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma delimited list of sessions and session ranges to retrieve file hashes for.
beginDate	string, optional	The beginning of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.
endDate	string, optional	The end of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.
directories	string, optional	A semi-colon delimited list of additional directories to search for hash files. This may be used to account for changes to the hash.dir configuration.

Manual for hashInfo

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

manifest

Description: If a manifest directory is defined, it will allow operations on the manifest files (such as a time based query) for database files in cold storage.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: query|compress}	The operation to perform (defaults to query)
time1	date-time, optional	The beginning time (UTC) for matching offline database files
time2	date-time, optional	The ending time (UTC) for matching offline database files
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format that is returned (posix, simple), default is posix

Manual for manifest

optimize

Description: Runs a series of tests to determine the optimimal *.write.block.size based on configured drives and using data from existing databases.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more	The database types to optimize, default is all

Parameter	Type, Options	Description
	of the following: session|meta|packet}	
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 104857600 and 107374182400, inclusive}	The amount of data to write for each test
compression	string, optional, {enum-one:The value must be one of the following: none|gzip|bzip2|lzma}	Determines if compression block sizes should be tested, default is none
compressionLevel	uint32, optional, {unsigned:The value must be between 0 and 9, inclusive}	Determines the compression level, 0-9, where 1 is fastest, 9 is the best compression and zero is a predetermined balance between speed and compression

reconfig

Description: Calculates new drive sizes and free space for the session, meta and/or packet directories. No directories are removed and the assumption is each directory is mounted on a separate filesystem and will only be used for storage of that database.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective *.dir and *.free.space.min settings, otherwise it will just output the calculations
type	string, optional, {enum-any:The value must be one or more	The database types to reconfigure, default is all

Parameter	Type, Options	Description
	of the following: session|meta|packet}	
percent	uint32, optional, {unsigned:The value must be between 1 and 99, inclusive}	The drive percentage to use for the calculated size per directory, default is 95
op	string, optional, {enum-one:The value must be one of the following: normal|10g}	The operating mode, 'normal' or '10g'. Default is normal.

Manual for reconfig

resetMax

Description: Resets all max database stats or just the ones listed.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
stats	string, optional, {enum-any:The value must be one or more of the following: session.rate.max|meta.rate.max|packet.rate.max}	The stats to reset, default is all

sizeRoll

Description: Delete database files based on the total size of all databases (passed with 'type' parameter) or space remaining on shared volume(s). This command should not be used on databases that don't share storage.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-any:The value must be one or more of the following: session|meta|packet}	The databases to consider for removing the oldest data based on total size or space remaining
log	bool, optional	If true, will log all actions taken, otherwise all activity will happen silently. Default is true.
maxSize	size, optional	The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Hot tier.
maxSizeWarm	size, optional	The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Warm tier.
maxPercent	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Hot tier.
maxPercentWarm	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Warm tier.
minFree	size, optional	The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once

Parameter	Type, Options	Description
		again exceeds target. This only applies to the Hot tier
minFreeWarm	size, optional	The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once again exceeds target. This only applies to the Warm tier

Manual for sizeRoll

stagger

Description: Staggers database files to optimize read/write performance across multiple volumes. Can be used after adding an empty mount point

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: session|meta|packet}	The databases files to stagger
dryRun	bool, optional	If true (the default), will only return a description of the operations that would be performed. If false, then the files will actually be moved to an optimal read/write pattern. Please turn off capture or aggregation before actually performing the stagger operation.

Manual for stagger

timeRoll

Description: Delete database files that exceed a given age

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-any:The value must be one or more of the following: session|meta|packet}	The database files to remove
timeCalc	string, optional, {enum-one:The value must be one of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of days
date	string, optional	Remove database files older than the given UTC date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

Manual for timeRoll

wipe

Description: Overwrites all packets and/or meta for a session with a pattern (for eliminating sensitive information). Meta keys sessionid, time and size always remain untouched.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id whose packets will be wiped
payloadOnly	bool, optional	If true (default), will only overwrite the packet payload
pattern	string, optional	The pattern to use, by default it uses all zeros
metaList	string, optional	Comma separated list of meta to wipe, default (empty) is all meta
source	string, optional, {enum-any:The value must be one or more of the following: m|p}	The types of data to wipe, meta and/or packets, default is just packets

/decoder node

Manual for /decoder

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

logStats

Description: Retrieve count and last receive times for log devices, sources and forwarders.

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: clear|remove|save|buckets}	The log stats operation. If clear, all log stats state will be reset and no results will be returned. If save, the current stats are saved for the next time the log decoder is restarted.

Parameter	Type, Options	Description
deviceFilter	string, optional	A regex that will only return matching devices.
forwarderFilter	string, optional	A regex that will only return matching forwarders.
sourceFilter	string, optional	A regex that will only return matching sources.
discoveredFilter	string, optional	A regex that will only return matching discovered device types.
timeFilter	string, optional	This parameter forms a window from the current time which will return any device that falls outside that window. So if you pass 300, it will return all devices that have not sent a log in over 5 minutes. You can pass a hard datetime (YYYY-MM-DD HH:MM:SS) or a time duration (HH:MM:SS or just number of seconds). Not compatible with the lastReceived parameter.
lastReceived	string, optional	Compare the last received log time from each device with the passed in operator (
lastUpdated	string, optional	Compare the last updated log time from each device with the passed in operator (
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format that is returned (posix, simple), default is simple (YYYY-MM-DD HH:MM:SS)

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

reconfig

Description: Calculates optimal settings for decoder pools and buffers based on the installed hardware.

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false,on,off}	If true (default is false), will automatically update the respective decoder settings, otherwise it will just output the calculations
op	string, optional, {enum-one:The value must be one of the following: normal|10g|ndr}	The operating mode, 'normal', '10g', or 'ndr'. Default is normal.
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

Manual for reconfig

reparse

Description: Reparse a specified set of packet IDs

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
packets	string	A comma separated list of packet ids or packet id ranges to reparse
collection	string, optional	The target collection into which to write the newly created sessions. If this is left blank then the new sessions will be written to the collections determined by the collections rules.
source	string, optional	The source collection from which to read packets. If left blank then packets will be read from the default collection

reset

Description: Reset data, index, manifests, stats, configuration, or logs for this service. Data automatically deletes index and stats, unless filesCreatedAfter is specified. Service is automatically restarted. Example arguments: data=1 config=1 log=1 This example will reset data, index, logs, and configuration index=1 This example will reset the index only manifest=1 This example will delete everything in the manifest directory filesCreatedAfter="2015-12-01 14:00:00" This example will delete all session, meta and packet files created on or after Dec 1st, 2015 2pm (UTC) from the service. All other files will remain. The index will not be touched, but upon restart will be truncated to match the last session in the session database.

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
data	bool, optional	Reset data, automatically resets index (may not be applicable to all services)
index	bool, optional	Reset index, only valid if data is not reset (may not be applicable to all services)

Parameter	Type, Options	Description
manifest	bool, optional	Reset all manifests stored in the long term manifest directory (may not be applicable to all services)
config	bool, optional	Reset configuration settings to default
stats	bool, optional	Reset the stats database
log	bool, optional	Reset the log database
filesCreatedAfter	date-time, optional	Delete all database files (session, meta, packets) created after a certain date. The index will automatically get rolled back on restart. Not valid with option index.
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart
force	bool, optional	Force reset without confirmation

Manual for reset

resetMax

Description: Resets all max stats to zero.

Security.roles: decoder.manage

select

Description: Selected a new capture device

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
adapter	string, optional	The new adapter. If ommitted will display a list of available adpaters. You may specify a comma-separated list of adapter indexes to enable multi-interface capture.

Manual for select

start

Description: Starts aggregation

Security.roles: decoder.manage

Manual for start

stop

Description: Stops aggregation

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
flush	bool, optional	If true (default), process all sessions and packets currently in memory. Otherwise, stop immediately, discarding unprocessed sessions and packets.

Manual for stop

whoAgg

Description: Returns information on who is aggregating from this service

Security.roles: decoder.manage

Parameters:

Parameter	Type, Options	Description
clean	bool, optional	If true, will delete any aggregation trackers that no longer have a valid channel

Manual for whoAgg

/decoder/parsers node

Manual for /decoder/parsers

API Messages

attrib

Description: Manage attributes for content files.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|update}	The operation to perform. If list (default), return a list of attributes for specified type of contents. If update, set the attributes for specified content file.
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices).
file	string, optional	The content file name.
values	string, optional	List of attributes to set for the specified content file in name value pairs.

content

Description: Get parsing content file information.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
file	string, optional	The content filename (type must be specified)
type	string, optional, {enum-one:The value must be one of the following: feeds|parsers}	The content type
uuid	string, optional	The content uuid

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete parsing content files.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
file	string, optional	The content filename (type must be specified)
uuid	string, optional	The content uuid
type	string, optional, {enum-one:The value must be one of the following: feeds|parsers|device}	The content type (feeds|parsers|device)

devices

Description: Returns the log parsers

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
uuid	string, optional	The content uuid

disable

Description: Disable the specified content.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices)
uuid	string	The content uuid

dyndvmap

Description: Retrieve dynamic mappings of sources to device types in log parsing.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: describe|clear|save}	The operation to perform. If describe (default), return a listing of mappings. If clear, the all mappings will be removed. If save, the mappings will be persisted.
deviceFilter	string, optional	A regex that will only return matching devices.
sourceFilter	string, optional	A regex that will only return matching sources.

Parameter	Type, Options	Description
forwarderFilter	string, optional	A regex that will only return matching forwarders.

Manual for dyndvmap

enable

Description: Enable the specified content.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices)
uuid	string	The content uuid

export

Description: Returns the content as a package.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: feeds|parsers|devices}	The content type (feeds|parsers|devices)
uuid	string, optional	The content uuid
file	string, optional	The content file name

feed

Description: Manages the feed parsers

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: notify|reload|remove|delete}	The feed operation to perform
file	string, optional	The feed filename

headers

Description: Return header information about particular devices.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
devices	string, optional	A CSV of devices for which to get header information. If omitted, all header information is returned.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ipdevice

Description: Map IP to Device type in log parsing. Multiple device types mapped to the same ip/host are prioritized in the order in which they are listed. Takes effect immediately.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: add|edit|remove|describe}	The operation to perform.
entries	string, optional	The IP entries. StringParam in format of 'ip=device'. The edit op accepts a preceding '+' or '-' for each param (e.g. '+ip=device').The add op or edit op with + means adding or editing a map entry, the remove op or edit op with
time	string, optional	Used with 'describe' to return the mapped devices that have been

Parameter	Type, Options	Description
		updated AFTER the time specified (eg. "2010-Apr-20 10:00:00" will return all elements updated after April 20, 2010 10:00).
sources	string, optional	Used with 'describe' to return the mapped devices of the specified CSV of sources. (eg. sources="192.168.1.114,192.168.112" will return all mappings from those devices only).
soft	bool, optional	Used with 'op=add'. Indicates that the mappings to be added are soft mappings.

iptmzone

Description: Map IP to time zone in log parsing. Takes effect after parser reload.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: add|edit|remove|describe}	The operation to perform (edit|describe).
entries	string, optional	The IP entries. StringParam in format of '+/-ip=
mdformat	string, optional, {enum-one:The value must be one of the following: none|mdy|dmy}	Override the parsing order of Month/Day or Day/Month (none|mdy|dmy). In Month/Day/Year or Day/Month/Year sequences, the parsing order of Month and Day can reversed to match the order observed in the messages. A value of 'dmy' will override parsing to be Day/Month/Year. A value of 'mdy' will perform the reverse.

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

patterns

Description: API for interfacing UI managed content nodes on the Log Decoder

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
op	string, optional	Use 'parse' to get detailed information about how parse rules are parsed, otherwise leave blank for default behavior.
devices	string, optional	A CSV of device names or '*'. If empty, a list of all devices that have at least one content node defined is returned. If '*', a list of content nodes for all devices. Otherwise, the content nodes of the requested devices.
stems	string, optional, {enum-any:The value must be one or more of the following: default|custom|tokens}	A CSV of source content filename stems to list. Default is all stems.

Parameter	Type, Options	Description
showDisabled	bool, optional	Include disabled rules. Default is false.
type	string, optional, {enum-one:The value must be one of the following: RULE|DataType}	If not specified then returns all content types, otherwise the name of the content type requested.

reload

Description: Reloads all meta parsers

Security.roles: parsers.manage

schema

Description: Returns the parsers meta schema

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
onlyShowEnabled	bool, optional	If true, only return information on parsers that are currently enabled
showRegisteredMeta	bool, optional	If true (the default), show the registered meta for each parser
showOptions	bool, optional	If true, shows the options supported by each parser. Passing true changes the XML structure returned so each parser (potentially) has
prettyPrint	bool, optional	If true, returns the XML formatted to be easier to read. Default is false.

tzinfo

Description: Timezones available for normalizing timestamps parsed from log events.

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: version|tznames|abbreviations|duplicates|overrides|add|remove}	The operation to perform.
abbr	string, optional	A timezone abbreviation to add or remove to the overrides list. This parameter is required for the add and remove operations.
gmtoffset	string, optional	The offset from GMT of a timezone abbreviation to add to the override list. Format is [-+]HH[MM].
tzname	string, optional	A timezone name to override the default timezone for a duplicate abbreviation. The named

Parameter	Type, Options	Description
		timezone must exist and the abbreviation must be valid for that timezone (see tznames).

Manual for tzinfo

upload

Description: Upload feed/parser/geoip2 files to this service

Security.roles: parsers.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: start|attributes|cancel|finished}	The operation to perform (start, cancel, finished)
size	uint64, optional	The size of the incoming file
filename	string, optional	The name of the incoming file
reload	uint32, optional	Reload all parsers when upload is complete for the appropriate file types (default). A setting of 0 when sending a finished op will disable the reload.

/index node

API Messages

cacheClr

Description: Clear index internal caches

Security.roles: index.manage

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

inspect

Description: Provides overall information about the index or can inspect specific keys and values

Security.roles: index.manage

Parameters:

Parameter	Type, option	Description
key	string, optional	The key to inspect
value	string, optional	The value to inspect
format	string, optional, {enum-one:The value must be one of the following: Text|IPv4|Int8|UInt8|Int16|UInt16|Int32|UInt32|Int64|UInt64|UInt128| Float32|Float64|TimeT|DayOfWeek|HourOfDay|Binary|IPv6|MAC}	Only include keys that have this format
valueCount	uint64, optional	The number of values to return from each key
allSlices	bool, optional	Return information on all slices, instead of just most recently created slice
summarizeAllValues	bool, optional	Calculate total summary and page usage from all values in each slice

language

Description: Returns language information

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
values	bool, optional	If set to 1, will describe the value overrides defined in the language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

profile

Description: Returns information about the index used by the Index Profile tool

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string, optional	The key to profile

Parameter	Type, Options	Description
value	string, optional	The value to profile
max	uint32, optional	Max number of pages to return
flag	uint32, optional	Bitwise option when key is specified. 1: Break report to values. 2: List page sizes. 4: Categorize page counts in different range. 8: Sessions count Per page with each algorithm. 16: Verify session page

reconfig

Description: Calculates new sizes for the summary and page directories. NOTE: Any size changes require a reset of the index before they take effect.

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective config node settings, otherwise it will just output the calculations
percent	uint32, optional, {unsigned:The value must be between 1 and 99, inclusive}	The drive percentage to use for the index database, default is 90

save

Description: Saves the index to disk

Security.roles: index.manage

schema

Description: Introspects all open Meta DB files for meta keys and merges those keys with the defined language to return a fairly comprehensive list of all meta keys in existence.

Security.roles: index.manage

sizeRoll

Description: Delete index and other DB files (see 'type' param) based on the total size or space remaining on shared volumes. This command is not intended to be used on databases that do not share storage.

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
log	bool, optional	If true, will log all actions taken, otherwise all activity will happen silently. Default is true.
type	string, optional, {enum-any:The value must be one or more of the following: session|meta|packet}	The databases to consider for removing the oldest data based on total size or space remaining
maxSize	size, optional	The maximum size of the index hot storage. When exceeded, oldest data is deleted first until total size is less than maxSize.
maxSizeWarm	size, optional	The maximum size of the index warm storage. When exceeded, oldest data is delete first until total size is less that maxSizeWarm
maxPercent	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This parameter applies to hot storage only.
maxPercentWarm	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This parameter applies to warm storage only.

Parameter	Type, Options	Description
minFree	size, optional	The minimum allowed free space on the volumes before oldest data is deleted. This parameter applies to hot storage only.
minFreeWarm	size, optional	The minimum allowed free space on the volumes before oldest data is deleted. This parameter applies to warm storage only.

values

Description: Returns the top n values for the given key

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string	The index key for values
top	uint32, {unsigned:The value must be between 0 and 10000, inclusive}	The number of top entries to return

/logs node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

download

Description: Downloads log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting log number to download
id2	uint64, optional	The ending log number to download
time1	date-time, optional	The starting log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
time2	date-time, optional	The ending log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
op	string, {enum-one:The value must be one of the following: start|next|cancel}	The operation to perform (start, next, cancel)
logTypes	string, optional	The type(s) of log messages to retrieve (debug,info,audit,warning,failure)
match	string, optional	Case insensitive string to match in each log message
regex	string, optional	Regular expression to match in each log message
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)
batchSize	uint32, optional, {unsigned:The value must be between 2000	The number of logs to retrieve per batch, default is 2000

Parameter	Type, Options	Description
	and 10000, inclusive}	

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

pull

Description: Downloads N log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The first log id number to retrieve, this is mutually exclusive with id2
id2	uint64, optional	The last log id number that will be sent, defaults to most recent log message when id1 or id2 is not sent
count	uint32, optional, {unsigned:The value must be between 1 and 10000, inclusive}	The number of logs to pull

Parameter	Type, Options	Description
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)

timeRoll

Description: Delete log entries that exceed a given age

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
timeCalc	string, optional, {enum-any:The value must be one or more of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 60, inclusive}	Remove log entries older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 24, inclusive}	Remove log entries older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 365, inclusive}	Remove log entries older than the given number of days
date	string, optional	Remove log entries older than the given date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

/rest node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

/sdk node

Manual for /sdk

API Messages

aliases

Description: Retrieves aliases for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string	The language key to retrieve aliases for
value	string, optional	An optional value to return a specific alias for

cancel

Description: Cancel a running query, regardless of the user

Security.roles: sdk.manage,sdk.meta

Parameters:

Parameter	Type, Options	Description
handle	uint32	The channel handle running the query
uuid	string, optional	The device UUID if taking a device offline instead of canceling the entire query (Broker only). Passing uuid will propagate the request upstream to find the first matching device to cancel.

content

Description: Returns the packet content for a session

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session id to retrieve packet content for
packet	uint64, optional	The packet id of the first packet of the session, in case the session has rolled out
pinId	string, optional	The Pin ID of the session to retrieve from long term storage. If provided, the packet and session parameters must not be provided.
maxSize	size, optional	The max number of bytes to return, zero means no limit
failOnCacheMiss	bool, optional	If true, the content call fails if the NWD is not in the cache. This is used by optimized caching code only.
renderType	string, optional	The render type (see NwContentTypes in NwSDK.h), pass zero to pick best option. Not specifying this parameter always results in an NWD and renderFlags and renderOptions are not considered. The following values are also supported: 'pcap', 'file-list-json', 'session-meta-file-list-json', 'nwd', 'log', 'files'

Parameter	Type, Options	Description
renderFlags	uint32, optional	Bitwise mask to control options, (see NwContentFlags in NwSDK.h)
renderOptions	string, optional	An encoded string params containing additional options for controlling the operation.

Manual for content

count

Description: Returns the number of child nodes

Security.roles: everyone

delCache

Description: Removes all cached NWD files from either the normal NWD cache or the pin cache

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-one:The value must be one of the following: nwd|pin|all}	Specifies which cache to delete all NWD files from. The default is 'nwd'. Passing 'all' will delete from all caches.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

hierarch

Description: Report the set of devices attached to this service

Security.roles: sdk.meta

Manual for hierarch

info

Description: Returns detailed information about the node

Security.roles: everyone

keyrefs

Description: Retrieves the set of entity keys

Security.roles: sdk.meta

Manual for keyrefs

language

Description: Retrieves the language definition

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting language id
id2	uint64, optional	The ending language id
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
options	string, optional	Extra options, currently unused
flags	string, optional	Optional flags to configure how the results are returned. Can be a number (bitwise mask) or comma separated values like default, no-count, quick-count or full-count.
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

msearch

Description: Search for pattern matches in many sessions or packets

Security.roles: sdk.content & sdk.meta

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	The session ID ranges to search
packets	string, optional	The packet ID ranges to search
search	string	The search string to use
where	string, optional	The where clause used to identify sessions to consider for the search
limit	uint64, optional	The maximum number of sessions to traverse for this search
size	uint64, optional	The maximum number of results to return for this search
flags	string, {enum-any:The value must be one or more of the following:	Flags to use for search. This is a comma separated list of one or more of the flag values:

Parameter	Type, Options	Description
	regex|sp|sm|ci|pre|post|ds|precache|si}	regex,sp,sm,ci,pre,post,ds,precache,si
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for msearch

packets

Description: Stream packets back based on the input parameters provided

Security.roles: sdk.content & sdk.packets

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: setup|start|cancel|processed}	The operation to perform (start, cancel, processed)
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) whose packets will be streamed back.
packets	string, optional	A comma separated list of packet ids or session/packet ids (#&#) that will be streamed back.
where	string, optional	Where clause which will be evaluated to determine which sessions to stream back
time1	date-time, optional	A starting time range (UTC) where matching packets will be streamed back ("2010-Apr-20 09:00:00")
time2	date-time, optional	An ending time range (UTC) where matching packets will be

Parameter	Type, Options	Description
		streamed back ("2010-Apr-20 10:00:00")
flags	uint32, optional	Additional flags as defined by the NwPackets SDK function
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for packets

pin

Description: Pin a session in the long term cache so it can be retrieved even after the meta or packets have rolled out. To retrieve, you must make a content call with the Pin ID that is returned from this command. The op={unpin,validate} commands all take one or more comma separated Pin IDs in the pinId parameter and returns the status for the Pin IDs that were recognized by a service.

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session ID to pin. This cannot be used with the 'sessions' parameter.
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) to pin. This cannot be used with the 'session' parameter.
pinId	string, optional	One or more Pin IDs, separated by commas
op	string, optional, {enum-one:The value must be one of the following: ls|unpin|validate}	Various operations that can be performed. Will be submitted to upstream devices if not configured for long term cache.

Manual for pin

precache

Description: Efficiently caches NWD files for future retrieval via the content message

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) that will be saved in the NWD cache directory
packets	string, optional	A comma separated list of session and packet ids (#p#) that will be saved in the NWD cache directory
waitOnComplete	bool, optional	If true, will wait for precache to complete before returning a response.

query

Description: Performs a query against the meta database

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id (to run the query from most recent to oldest meta, make id1 larger than id2)
id2	uint64, optional	The ending meta id
size	uint32, optional	The max number of entries to return, or just stream back all results if zero
query	string, optional	The query string to use

Parameter	Type, Options	Description
flags	string, optional	The flags to use for query. Can be a number (bitwise mask) or comma separated values like query-log.
threshold	uint64, optional	Query optimization to stop processing results after the threshold is reached (useful with select aggregate functions). Zero means no threshold (the default).
partition	uint32, optional	Used to partition the result set across multiple clients. Must be greater than zero and less than or equal to totalPartitions
totalPartitions	uint32, optional	The total number of clients that will be submitting the same query for partitioning.
search	string, optional	A text search clause to apply to the where clause of this query
streamLimit	uint32, optional	Limit the number of sessions returned by a streaming query
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for query

reconfig

Description: Calculates default values for some of the config nodes.

Security.roles: sdk.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the config nodes with the defaults
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

search

Description: Searches for matches in session/packet content

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to search
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
search	string	The search string to use

Manual for search

session

Description: Retrieves the meta id range for the session range

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64	The starting session id (inclusive). Pass zero to scope to lowest valid id.
id2	uint64	The ending session id (inclusive). Pass zero to scope to highest valid id.

Manual for session

summary

Description: Retrieves summary information from the databases

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
flags	string, optional	Optional SDK flags. Can be a number (bitwise mask) or comma separated values like default or ignore-cache.

Manual for summary

timeline

Description: Returns the count of sessions/size/packets in discrete time intervals

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS), if greater than time2, then the whole time range is returned
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)

Parameter	Type, Options	Description
timezone	int32, optional	The timezone of the local computer to receive the data, hour offset from GMT [-13 to 13]
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for the timeline. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, order-ascending, order-descending, ignore-cache or clear-cache.
where	string, optional	Optional where clause for filtering the data
bookendHours	uint32, optional	The number of hours outside the given time1, time2 which will become part of the cache. Default is 24 hours.
fieldName	string, optional	The time field to retrieve values for, defaults to 'time'
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query

Manual for timeline

validate

Description: Validate a query or values call without execution

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: query|values}	The operation to validate
validateLanguage	bool, optional	If true (default is false), language tokens and syntax will be validated, otherwise just syntax will be validated
fieldName	string, optional	Required if validating a 'values' operation. The field to retrieve values for.
where	string, optional	Required if validating a 'values' operation. The filter criteria to be validated.
query	string, optional	Required if validating a 'query' operation. The query string to be validated.
json	bool, optional	Structure response in JSON syntax

Manual for validate

values

Description: Performs a value count query and returns the matching values for a report

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id
id2	uint64, optional	The ending meta id
size	uint32, {unsigned:The value must be between 1 and	The max number of entries to return

Parameter	Type, Options	Description
	1677721, inclusive}	
flags	string, optional	The flags to use for values. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, aggregate, sort-total, sort-value, order-ascending, order-descending, ignore-cache, clear-cache, or database-scan.
threshold	uint64, optional	Query optimization to stop processing large session counts
fieldName	string	The field to retrieve values for
where	string, optional	The filter criteria for the values
aggregateFunction	string, optional, {enum-one:The value must be one of the following: count|sum}	The aggregate operation to be applied where the aggregate flag is set
aggregateFieldName	string, optional	The meta field to aggregate in the aggregateFunction
min	string, optional	The lower limit of values to return. Only values greater than this value will be returned.
max	string, optional	The upper limit of values to return. Only values less than this value will be returned.
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query

Manual for values

xforms

Description: Retrieves transforms for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string, optional	An optional language key to retrieve transform for, all transforms returned if not provided

/sys node

Manual for /sys

API Messages

caCert

Description: Display or delete trusted CA certs

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: list|delete|add}	CA action to perform (list,delete,add)
id	string, optional	ID of CA certificate to remove when op=delete

Manual for caCert

count

Description: Returns the number of child nodes

Security.roles: everyone

fileEdit

Description: View and edit system configuration files

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: dir|get|getBackup|put|create|delete}	The operation to perform (dir, get, getBackup, put, create, delete)
type	string, optional, {enum-one:The value must be one of the following: parser|filter}	The type of file, if not provided will check file extension
filename	string, optional, {length:The value length must be between 1 and 64, inclusive}	File to create, edit, get or delete

Manual for fileEdit

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

peerCert

Description: Display or modify trusted peer certificates

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|delete|add}	Operation to perform on the list of trusted peers (list,delete,add)
id	string, optional	ID of peer certificate to remove when op=delete

Manual for peerCert

save

Description: Forces a save of system settings

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
force	uint32, optional	If 1, forces the configuration file to saved, regardless if changes are detected, default is 0

Manual for save

servCert

Description: Display the current server certificate in PEM format.

Security.roles: sys.manage

Manual for servCert

shutdown

Description: Stop the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
reason	string, optional, {length:The value length must be between 0 and 512, inclusive}	The reason for shutting down the service
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart

Manual for shutdown

statHist

Description: Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time for retrieving stats
time2	date-time, optional	The ending time for retrieving stats
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format for each stat snapshot, default is posix (seconds since 1970)

Parameter	Type, Options	Description
include	string, optional	Comma separated list of stats to include (wildcards allowed)
exclude	string, optional	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)
onChange	string, optional	Comma separated list of stats that reduce the result set to only when at least one of the stats change from their previous value
reduce	bool, optional	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result
showAll	bool, optional	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Manual for statHist

telemtry

Description: Returns telemetry information about the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
level	uint32, optional, {signed:The value must be between 1 and 3, inclusive}	The level of information to be sent by the telemetry request. A value of 1 (default) will send all stat nodes, a value of 2 will send stat and config nodes, and a value of 3 will send stat nodes, config nodes and usage metrics for rules, parsers and devices (where applicable).
include	string, optional	Comma separated list of nodes to include (wildcard usage is ok).

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes to exclude (wildcard usage is ok). If "include" is also specified, exclusions are applied to the set of nodes specified by "include".
url	string, optional	Endpoint URL to send the JSON response.
op	string, optional, {enum-one:The value must be one of the following: get|configure}	The telemetry operation to execute. Default is 'get'.

/users node

API Messages

addOrMod

Description: Add a new user or update an existing user in the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to add or update, must be alphanumeric or @!#$%&'+-=?^_`{|}~.[]\
password	string, optional	The user's password
pwdIsHashed	bool, optional	Set to true if the password has already been hashed
groups	string, optional	The groups the user belongs to

Parameter	Type, Options	Description
authType	string, optional, {enum-one:The value must be one of the following: netwitness|pam}	The authentication system to use, "netwitness" is the default
queryTimeout	uint32, optional, {unsigned:The value must be between 0 and 43200, inclusive}	The maximum number of minutes that a query is allowed to execute for this user. Zero means unlimited.
displayName	string, optional	Display name for this user account
email	string, optional	Email address for this account
description	string, optional	Description of this user account
queryPrefix	string, optional	Query filter applied to every query performed by this user account
sessionThreshold	uint32, optional, {unsigned:The value must be between 0 and 4294967295, inclusive}	Query optimization which will extrapolate the remaining session counts when they exceed this value

auths

Description: Get supported authentication types

Security.roles: everyone

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete a user from the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	ption
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to delete

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

unlock

Description: Unlock a locked out account

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string	The username or account to unlock

whoAmI

Description: Returns information about the authenticated user

Security.roles: everyone

Concentrator service

Introduction

Concentrator stores and indexes metadata for fast queries and retrieval of raw data capture. It does not capture data, it aggregates the sessions and meta from a Decoder/LogDecoder and indexes them.
Typically the index on a concentrator is much larger and this meant to do the heavy lifting for queries, leaving the Decoders focused on capture and parsing. When the raw data (packets or logs) are requested from a Concentrator, the request is forwarded to the Decoder with the data. The response (with the raw data) is then forwarded back to the client. In this respect, raw data access is seamless from the client's point of view, even though the Concentrator never stores the data. It simply knows where the data is kept, requests the data on the client's behalf and streams it back seamlessly.

API Usage

Concentrator service provides API interface through its tree form as nodes, node messages and child nodes.

The following are few examples of the several API nodes and messages supported by Concentrator:

- /concentrator node messages:
    - `start` and `stop` aggregation
    - `reconfig` concentrator for optimal settings based on memory and hardware
    - `add` a device to concentrator
    - `/concentrator/stats` node for concentrator aggregation stats information
- /sdk node messages:
    - `reconfig` sdk for optimal settings
    - `summary` to retrieve summary information from the databases
    - `values` to perform a value count query and return the matching values for a report
    - `query` to perform query against meta database
    - `language` of meta keys for key info, format, description etc...

    - `/sdk/stats` node for queries active, pending and other sdk operations
- /index node messages:
    - `language` of meta keys index level, valueMax, format, description etc...
    - `inspect` for overall information about the index
    - `save` for saving index to disk
    - `/index/stats` node for index stats , slices, memory usage, reindex status etc...
- /database node messages:
    - `reconfig` database for optimal database settings
    - `dbState` to get comprehensive information about current databases
    - `/database/stats` node for database stats information
- /sys node messages:
    - `statHist` to retrieve historical stats from stats database
    - `/sys/stats` node for service and system stats information

Sample Request

The following is a sample request to get session count in sdk query on the concentrator.

Request path is `/sdk` and URL request parameters for message query are `msg=query query="select count(sessionid) where did='decoder-pdec-0'" size=5`.

For queries that are expected to take longer time on execution an additional URL parameter `expiry=0` can be added to disable default REST request timeout and continue with user profile timeout.

Using URL encoded request

```
$ curl "https://concentrator-host:50105/sdk?msg=query&query=select%20count(sessionid)%20where%20did%3D%27decoder-pdec-0%27&size%3D5&expiry=0" \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1' \
   -u 'username:password'
```

Using special content type

```
Content Type: application/x-netwitness-string-params

$ curl 'https://concentrator-host:50105/sdk?msg=query&expiry=0' \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-netwitness-string-params' \
   -d "query=\"select count(sessionid) where did='decoder-pdec-0'\" size=5"
   -u 'username:password'
```

Sample Response

HTTP Headers

```
HTTP/1.1 200 OK
Content-Length: 186
Connection: Keep-Alive
Pragma: no-cache
Expires: -1
Cache-Control: no-cache, no-store, must-revalidate
Content-Type: application/json
```

JSON Response


```
{
   "flags": 1074200577,
   "results": {
      "id1": 3216363401111,
      "id2": 3216363401110,
      "fields": [
         {
            "id1": 0,
            "id2": 0,
            "count": 6864785,
            "format": 8,
            "type": "sessionid",
            "flags": 2,
            "group": 0,
            "value": "6864785"
         }
      ]
   }
}
```


`Note`: The API Messages section for the path nodes provide more details on API messages available and any associated manual pages.

/concentrator node

API Messages

add

Description: Add a new device for aggregation

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, optional, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to add (ip:port or [ipv6]:port), either this or ip must be set
ip	string, optional	The IP address of the device, either this or device must be set
port	uint32, optional, {unsigned:The value must be between 0 and 65535, inclusive}	The port number the device is listening on
username	string	The username to run as
password	string, optional	The account password. Assumes service trust aggregation if no password is provided.
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup
reconnect	string, optional	If true, will automatically reconnect when not connected
name	string, optional	Device collection name if device is a collection

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete an existing aggregation device

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to delete
name	string, optional	Device collection name if device is a collection

edit

Description: Edit an existing aggregation device

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to edit
username	string, optional	The username to run as

Parameter	Type, Options	Description
password	string, optional	The account password. For service trust aggregation you must not provide this parameter
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup
reconnect	string, optional	If true, will automatically reconnect when not connected
name	string, optional	Device collection name if device is a collection
switchToTrustedAuth	bool, optional	If true, removes any password associated with this account and will use trusted authentication henceforth. When switching to trusted, only the 'device' parameter is allowed.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

reconfig

Description: Calculates optimal settings for concentrator pools and buffers based on the installed hardware.

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective concentrator settings, otherwise it will just output the calculations
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

reset

Description: Reset data, index, manifests, stats, configuration, or logs for this service. Data automatically deletes index and stats, unless filesCreatedAfter is specified. Service is automatically restarted. Example arguments: data=1 config=1 log=1 This example will reset data, index, logs, and configuration index=1 This example will reset the index only manifest=1 This example will delete everything in the manifest directory filesCreatedAfter="2015-12-01 14:00:00" This example will delete all session, meta and packet files created on or after Dec 1st, 2015 2pm (UTC) from the service. All other files will remain. The index will not be touched, but upon restart will be truncated to match the last session in the session database.

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
data	bool, optional	Reset data, automatically resets index (may not be applicable to all services)
index	bool, optional	Reset index, only valid if data is not reset (may not be applicable to all services)

Parameter	Type, Options	Description
manifest	bool, optional	Reset all manifests stored in the long term manifest directory (may not be applicable to all services)
config	bool, optional	Reset configuration settings to default
stats	bool, optional	Reset the stats database
log	bool, optional	Reset the log database
filesCreatedAfter	date-time, optional	Delete all database files (session, meta, packets) created after a certain date. The index will automatically get rolled back on restart. Not valid with option index.
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart
force	bool, optional	Force reset without confirmation

resetMax

Description: Resets all max stats, including device max stats back to zero.

Security.roles: concentrator.manage

start

Description: Starts aggregation

Security.roles: concentrator.manage

status

Description: Change the online/offline status of an aggregation device

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to change status
status	string, optional, {enum-one:The value must be one of the following: online|offline}	Whether to take device offline or put online
name	string, optional	Device collection name if device is a collection

stop

Description: Stops aggregation

Security.roles: concentrator.manage

whoAgg

Description: Returns information on who is aggregating from this service

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
clean	bool, optional	If true, will delete any aggregation trackers that no longer have a valid channel

/connections node

API Messages

closeAll

Description: Closes all connections

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: native|rest|amqp}	Indicates the type of connection to close, either native, rest or amqp. Default is all types.
limit	uint32, optional	If non-zero, will stop closing connections after limit is reached. NOTE: It starts closing the most idle connections first.
idleLongerThan	string, optional	A time duration in HH:MM:SS or ##{d,h,m,s} format. Any connection that has been idle longer than the passed in duration is eligible to be closed. Connections that have seen activity within the duration window will not be closed. Examples: "01:05:10" (1 hour, 5 mins, 10 secs) or "65m" (65 minutes)
exceptMe	bool, optional	If true, will not close the connection that initiated the command.

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

netSpeed

Description: Network Speed Test

Security.roles: connections.manage,aggregate

Parameters:

Parameter	Type, Options	Description
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 1024 and 536870912, inclusive}	The size of each message sent, default is 64k
count	uint32, optional, {unsigned:The value must be between 1 and 536870912, inclusive}	The number of messages to send, default is 16k

/database node

Manual for /database

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

dbState

Description: Returns comprehensive information about the current database state or persists the current state to disk for fast reload

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: filelist|summary|save}	The operation to perform (filelist, summary or save)
type	string, {enum-any:The value must be one or more of the following: session|meta|packet}	The database(s) to operate on
options	string, optional, {enum-one:The value must be one of the following: bytes|pretty-print}	Output display options like 'bytes' or 'pretty-print'. Pretty print (the default) will convert all sizes from bytes to human readable quantities like MB or GB.

Manual for dbState

dump

Description: Dumps information out of the database in nwd formatted files

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to dump
type	string, optional, {enum-one:The value must be one of the following: db|nwd}	The dump type
source	string, optional, {enum-any:The value must be one or more of the following: s|m|p}	The types of data to dump, default is all
verbose	bool, optional	If true (default is false), dumps more information
file	string, optional	Optional filename to use for NWD type, otherwise filename is assumed to be _sessionid_.nwd

Parameter	Type, Options	Description
limit	uint32, optional	If assigned, will only output the number of metas/packets passed in. Otherwise will print everything.

hashInfo

Description: Retrieves hash information for database files that containing session/meta/packet objects for a set of sessions or date range.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma delimited list of sessions and session ranges to retrieve file hashes for.
beginDate	string, optional	The beginning of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.
endDate	string, optional	The end of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.
directories	string, optional	A semi-colon delimited list of additional directories to search for hash files. This may be used to account for changes to the hash.dir configuration.

Manual for hashInfo

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated

Parameter	Type, Options	Description
		values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

manifest

Description: If a manifest directory is defined, it will allow operations on the manifest files (such as a time based query) for database files in cold storage.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: query|compress}	The operation to perform (defaults to query)
time1	date-time, optional	The beginning time (UTC) for matching offline database files
time2	date-time, optional	The ending time (UTC) for matching offline database files
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format that is returned (posix, simple), default is posix

Manual for manifest

optimize

Description: Runs a series of tests to determine the optimimal *.write.block.size based on configured drives and using data from existing databases.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: session|meta}	The database types to optimize, default is all
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 104857600 and 107374182400, inclusive}	The amount of data to write for each test
compression	string, optional, {enum-one:The value must be one of the following: none|gzip|bzip2|lzma}	Determines if compression block sizes should be tested, default is none
compressionLevel	uint32, optional, {unsigned:The value must be between 0 and 9, inclusive}	Determines the compression level, 0-9, where 1 is fastest, 9 is the best compression and zero is a predetermined balance between speed and compression

reconfig

Description: Calculates new drive sizes and free space for the session, meta and/or packet directories. No directories are removed and the assumption is each directory is mounted on a separate filesystem and will only be used for storage of that database.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective *.dir and *.free.space.min settings, otherwise it will just output the calculations
type	string, optional, {enum-any:The value must be one or more of the following: session|meta}	The database types to reconfigure, default is all
percent	uint32, optional, {unsigned:The value must be between 1 and 99, inclusive}	The drive percentage to use for the calculated size per directory, default is 95
op	string, optional, {enum-one:The value must be one of the following: normal|10g}	The operating mode, 'normal' or '10g'. Default is normal.

Manual for reconfig

resetMax

Description: Resets all max database stats or just the ones listed.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
stats	string, optional, {enum-any:The value must be one or more of the following: session.rate.max|meta.rate.max}	The stats to reset, default is all

sizeRoll

Description: Delete database files based on the total size of all databases (passed with 'type' parameter) or space remaining on shared volume(s). This command should not be used on databases that don't share storage.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-any:The value must be one or more of the following: session|meta}	The databases to consider for removing the oldest data based on total size or space remaining
log	bool, optional	If true, will log all actions taken, otherwise all activity will happen silently. Default is true.
maxSize	size, optional	The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Hot tier.
maxSizeWarm	size, optional	The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Warm tier.
maxPercent	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Hot tier.
maxPercentWarm	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Warm tier.
minFree	size, optional	The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once

Parameter	Type, Options	Description
		again exceeds target. This only applies to the Hot tier
minFreeWarm	size, optional	The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once again exceeds target. This only applies to the Warm tier

Manual for sizeRoll

stagger

Description: Staggers database files to optimize read/write performance across multiple volumes. Can be used after adding an empty mount point

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-one:The value must be one of the following: session|meta}	The databases files to stagger
dryRun	bool, optional	If true (the default), will only return a description of the operations that would be performed. If false, then the files will actually be moved to an optimal read/write pattern. Please turn off capture or aggregation before actually performing the stagger operation.

Manual for stagger

timeRoll

Description: Delete database files that exceed a given age

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
type	string, {enum-any:The value must be one or more of the following: session|meta}	The database files to remove
timeCalc	string, optional, {enum-one:The value must be one of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 10000, inclusive}	Remove database files older than the given number of days
date	string, optional	Remove database files older than the given UTC date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

Manual for timeRoll

wipe

Description: Overwrites all packets and/or meta for a session with a pattern (for eliminating sensitive information). Meta keys sessionid, time and size always remain untouched.

Security.roles: database.manage

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id whose packets will be wiped
payloadOnly	bool, optional	If true (default), will only overwrite the packet payload
pattern	string, optional	The pattern to use, by default it uses all zeros
metaList	string, optional	Comma separated list of meta to wipe, default (empty) is all meta
source	string, optional, {enum-any:The value must be one or more of the following: m|p}	The types of data to wipe, meta and/or packets, default is just packets

/index node

API Messages

cacheClr

Description: Clear index internal caches

Security.roles: index.manage

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is

whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

inspect

Description: Provides overall information about the index or can inspect specific keys and values

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string, optional	The key to inspect
value	string, optional	The value to inspect
format	string, optional, {enum-one:The value must be one of the following: Text|IPv4|Int8|UInt8|Int16|UInt16|Int32|UInt32|Int64|UInt64|UInt128|Float32|Float64|TimeT|DayOfWeek|HourOfDay|Binary|IPv6|MAC}	Only include keys that have this format
valueCount	uint64, optional	The number of values to return from each key
allSlices	bool, optional	Return information on all slices, instead of just most recently created slice
summarizeAllValues	bool, optional	Calculate total summary and page usage from all values in each slice

language

Description: Returns language information

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
values	bool, optional	If set to 1, will describe the value overrides defined in the language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

profile

Description: Returns information about the index used by the Index Profile tool

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string, optional	The key to profile
value	string, optional	The value to profile
max	uint32, optional	Max number of pages to return
flag	uint32, optional	Bitwise option when key is specified. 1: Break report to values. 2: List page sizes. 4: Categorize page counts in different range. 8: Sessions count Per page with each algorithm. 16: Verify session page

reconfig

Description: Calculates new sizes for the summary and page directories. NOTE: Any size changes require a reset of the index before they take effect.

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective config node settings, otherwise it will just output the calculations
percent	uint32, optional, {unsigned:The value must be between 1 and 99, inclusive}	The drive percentage to use for the index database, default is 90

save

Description: Saves the index to disk

Security.roles: index.manage

schema

Description: Introspects all open Meta DB files for meta keys and merges those keys with the defined language to return a fairly comprehensive list of all meta keys in existence.

Security.roles: index.manage

sizeRoll

Description: Delete index and other DB files (see 'type' param) based on the total size or space remaining on shared volumes. This command is not intended to be used on databases that do not share storage.

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
log	bool, optional	If true, will log all actions taken, otherwise all activity will happen silently. Default is true.

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: session|meta}	The databases to consider for removing the oldest data based on total size or space remaining
maxSize	size, optional	The maximum size of the index hot storage. When exceeded, oldest data is deleted first until total size is less than maxSize.
maxSizeWarm	size, optional	The maximum size of the index warm storage. When exceeded, oldest data is delete first until total size is less that maxSizeWarm
maxPercent	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This parameter applies to hot storage only.
maxPercentWarm	string, optional	The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This parameter applies to warm storage only.
minFree	size, optional	The minimum allowed free space on the volumes before oldest data is deleted. This parameter applies to hot storage only.
minFreeWarm	size, optional	The minimum allowed free space on the volumes before oldest data is deleted. This parameter applies to warm storage only.

values

Description: Returns the top n values for the given key

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
key	string	The index key for values
top	uint32, {unsigned:The value must be between 0 and 10000, inclusive}	The number of top entries to return

/logs node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

download

Description: Downloads log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting log number to download
id2	uint64, optional	The ending log number to download
time1	date-time, optional	The starting log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
time2	date-time, optional	The ending log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: start|next|cancel}	The operation to perform (start, next, cancel)
logTypes	string, optional	The type(s) of log messages to retrieve (debug,info,audit,warning,failure)
match	string, optional	Case insensitive string to match in each log message
regex	string, optional	Regular expression to match in each log message
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)
batchSize	uint32, optional, {unsigned:The value must be between 2000 and 10000, inclusive}	The number of logs to retrieve per batch, default is 2000

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

pull

Description: Downloads N log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The first log id number to retrieve, this is mutually exclusive with id2
id2	uint64, optional	The last log id number that will be sent, defaults to most recent log message when id1 or id2 is not sent
count	uint32, optional, {unsigned:The value must be between 1 and 10000, inclusive}	The number of logs to pull
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)

timeRoll

Description: Delete log entries that exceed a given age

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
timeCalc	string, optional, {enum-any:The value must be one or more of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must	Remove log entries older than the given number of minutes

Parameter	Type, Options	Description
	be between 0 and 60, inclusive}	
hours	uint32, optional, {unsigned:The value must be between 0 and 24, inclusive}	Remove log entries older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 365, inclusive}	Remove log entries older than the given number of days
date	string, optional	Remove log entries older than the given date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

/rest node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

/sdk node

Manual for /sdk

API Messages

aliases

Description: Retrieves aliases for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string	The language key to retrieve aliases for
value	string, optional	An optional value to return a specific alias for

cancel

Description: Cancel a running query, regardless of the user

Security.roles: sdk.manage,sdk.meta

Parameters:

Parameter	Type, Options	Description
handle	uint32	The channel handle running the query

Parameter	Type, Options	Description
uuid	string, optional	The device UUID if taking a device offline instead of canceling the entire query (Broker only). Passing uuid will propagate the request upstream to find the first matching device to cancel.

content

Description: Returns the packet content for a session

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session id to retrieve packet content for
packet	uint64, optional	The packet id of the first packet of the session, in case the session has rolled out
pinId	string, optional	The Pin ID of the session to retrieve from long term storage. If provided, the packet and session parameters must not be provided.
maxSize	size, optional	The max number of bytes to return, zero means no limit
failOnCacheMiss	bool, optional	If true, the content call fails if the NWD is not in the cache. This is used by optimized caching code only.
renderType	string, optional	The render type (see NwContentTypes in NwSDK.h), pass zero to pick best option. Not specifying this parameter always results in an NWD and renderFlags and renderOptions are not considered. The following values are also supported: 'pcap', 'file-list-json', 'session-meta-file-list-json', 'nwd', 'log', 'files'
renderFlags	uint32, optional	Bitwise mask to control options, (see NwContentFlags in NwSDK.h)

Parameter	Type, Options	Description
renderOptions	string, optional	An encoded string params containing additional options for controlling the operation.

Manual for content

count

Description: Returns the number of child nodes

Security.roles: everyone

delCache

Description: Removes all cached NWD files from either the normal NWD cache or the pin cache

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-one:The value must be one of the following: nwd|pin|all}	Specifies which cache to delete all NWD files from. The default is 'nwd'. Passing 'all' will delete from all caches.

deviceId

Description: Given a session id, this message retrieves the session id and the device it maps to

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id that will be translated into the device name and session id.

Manual for deviceId

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

hierarch

Description: Report the set of devices attached to this service

Security.roles: sdk.meta

Manual for hierarch

info

Description: Returns detailed information about the node

Security.roles: everyone

keyrefs

Description: Retrieves the set of entity keys

Security.roles: sdk.meta

Manual for keyrefs

language

Description: Retrieves the language definition

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting language id
id2	uint64, optional	The ending language id
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
options	string, optional	Extra options, currently unused
flags	string, optional	Optional flags to configure how the results are returned. Can be a number (bitwise mask) or comma separated values like default, no-count, quick-count or full-count.
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
queryPriority	int32, optional, {signed:The value must be zero OR	The query priority for this query

Parameter	Type, Options	Description
	between -20 and 20, inclusive}	

Manual for language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

msearch

Description: Search for pattern matches in many sessions or packets

Security.roles: sdk.content & sdk.meta

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	The session ID ranges to search
packets	string, optional	The packet ID ranges to search

Parameter	Type, Options	Description
search	string	The search string to use
where	string, optional	The where clause used to identify sessions to consider for the search
limit	uint64, optional	The maximum number of sessions to traverse for this search
size	uint64, optional	The maximum number of results to return for this search
flags	string, {enum-any:The value must be one or more of the following: regex|sp|sm|ci|pre|post|ds|precache|si}	Flags to use for search. This is a comma separated list of one or more of the flag values: regex,sp,sm,ci,pre,post,ds,precache,si
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for msearch

packets

Description: Stream packets back based on the input parameters provided

Security.roles: sdk.content & sdk.packets

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: setup|start|cancel|processed}	The operation to perform (start, cancel, processed)
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) whose packets will be streamed back.

Parameter	Type, Options	Description
packets	string, optional	A comma separated list of packet ids or session/packet ids (#&#) that will be streamed back.
where	string, optional	Where clause which will be evaluated to determine which sessions to stream back
time1	date-time, optional	A starting time range (UTC) where matching packets will be streamed back ("2010-Apr-20 09:00:00")
time2	date-time, optional	An ending time range (UTC) where matching packets will be streamed back ("2010-Apr-20 10:00:00")
flags	uint32, optional	Additional flags as defined by the NwPackets SDK function
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
devices	string, optional	Optional list of comma separated device names that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /concentrator/devices
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for packets

pin

Description: Pin a session in the long term cache so it can be retrieved even after the meta or packets have rolled out. To retrieve, you must make a content call with the Pin ID

that is returned from this command. The op={unpin,validate} commands all take one or more comma separated Pin IDs in the pinId parameter and returns the status for the Pin IDs that were recognized by a service.

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session ID to pin. This cannot be used with the 'sessions' parameter.
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) to pin. This cannot be used with the 'session' parameter.
pinId	string, optional	One or more Pin IDs, separated by commas
op	string, optional, {enum-one:The value must be one of the following: ls|unpin|validate}	Various operations that can be performed. Will be submitted to upstream devices if not configured for long term cache.

Manual for pin

precache

Description: Efficiently caches NWD files for future retrieval via the content message

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) that will be saved in the NWD cache directory
packets	string, optional	A comma separated list of session and packet ids (#p#) that will be saved in the NWD cache directory

Parameter	Type, Options	Description
waitOnComplete	bool, optional	If true, will wait for precache to complete before returning a response.

query

Description: Performs a query against the meta database

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id (to run the query from most recent to oldest meta, make id1 larger than id2)
id2	uint64, optional	The ending meta id
size	uint32, optional	The max number of entries to return, or just stream back all results if zero
query	string, optional	The query string to use
flags	string, optional	The flags to use for query. Can be a number (bitwise mask) or comma separated values like query-log.
threshold	uint64, optional	Query optimization to stop processing results after the threshold is reached (useful with select aggregate functions). Zero means no threshold (the default).
partition	uint32, optional	Used to partition the result set across multiple clients. Must be greater than zero and less than or equal to totalPartitions
totalPartitions	uint32, optional	The total number of clients that will be submitting the same query for partitioning.
search	string, optional	A text search clause to apply to the where clause of this query

Parameter	Type, Options	Description
streamLimit	uint32, optional	Limit the number of sessions returned by a streaming query
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query

Manual for query

reconfig

Description: Calculates default values for some of the config nodes.

Security.roles: sdk.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the config nodes with the defaults
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

search

Description: Searches for matches in session/packet content

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to search
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
search	string	The search string to use

Manual for search

session

Description: Retrieves the meta id range for the session range

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64	The starting session id (inclusive). Pass zero to scope to lowest valid id.
id2	uint64	The ending session id (inclusive). Pass zero to scope to highest valid id.

Manual for session

summary

Description: Retrieves summary information from the databases

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
flags	string, optional	Optional SDK flags. Can be a number (bitwise mask) or comma separated values like default or ignore-cache.

Manual for summary

timeline

Description: Returns the count of sessions/size/packets in discrete time intervals

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS), if greater than time2, then the whole time range is returned
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
timezone	int32, optional	The timezone of the local computer to receive the data, hour offset from GMT [-13 to 13]
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for the timeline. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, order-ascending, order-descending, ignore-cache or clear-cache.
where	string, optional	Optional where clause for filtering the data
bookendHours	uint32, optional	The number of hours outside the given time1, time2 which will become part of the cache. Default is 24 hours.
fieldName	string, optional	The time field to retrieve values for, defaults to 'time'
queryPriority	int32, optional, {signed:The value must be zero OR	The query priority for this query

Parameter	Type, Options	Description
	between -20 and 20, inclusive}	
search	string, optional	A text search clause to apply to the where clause of this query

Manual for timeline

validate

Description: Validate a query or values call without execution

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: query|values}	The operation to validate
validateLanguage	bool, optional	If true (default is false), language tokens and syntax will be validated, otherwise just syntax will be validated
fieldName	string, optional	Required if validating a 'values' operation. The field to retrieve values for.
where	string, optional	Required if validating a 'values' operation. The filter criteria to be validated.
query	string, optional	Required if validating a 'query' operation. The query string to be validated.
json	bool, optional	Structure response in JSON syntax

Manual for validate

values

Description: Performs a value count query and returns the matching values for a report

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id
id2	uint64, optional	The ending meta id
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for values. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, aggregate, sort-total, sort-value, order-ascending, order-descending, ignore-cache, clear-cache, or database-scan.
threshold	uint64, optional	Query optimization to stop processing large session counts
fieldName	string	The field to retrieve values for
where	string, optional	The filter criteria for the values
aggregateFunction	string, optional, {enum-one:The value must be one of the following: count|sum}	The aggregate operation to be applied where the aggregate flag is set
aggregateFieldName	string, optional	The meta field to aggregate in the aggregateFunction
min	string, optional	The lower limit of values to return. Only values greater than this value will be returned.
max	string, optional	The upper limit of values to return. Only values less than this value will be returned.

Parameter	Type, Options	Description
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query

Manual for values

xforms

Description: Retrieves transforms for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string, optional	An optional language key to retrieve transform for, all transforms returned if not provided

/sys node

Manual for /sys

API Messages

caCert

Description: Display or delete trusted CA certs

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: list|delete|add}	CA action to perform (list,delete,add)
id	string, optional	ID of CA certificate to remove when op=delete

Manual for caCert

count

Description: Returns the number of child nodes

Security.roles: everyone

fileEdit

Description: View and edit system configuration files

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: dir|get|getBackup|put|create|delete}	The operation to perform (dir, get, getBackup, put, create, delete)
type	string, optional, {enum-one:The value must be one of the following: parser|filter}	The type of file, if not provided will check file extension
filename	string, optional, {length:The value length must be between 1 and 64, inclusive}	File to create, edit, get or delete

Manual for fileEdit

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is

whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder,

Parameter	Type, Options	Description
		session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

peerCert

Description: Display or modify trusted peer certificates

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|delete|add}	Operation to perform on the list of trusted peers (list,delete,add)
id	string, optional	ID of peer certificate to remove when op=delete

Manual for peerCert

save

Description: Forces a save of system settings

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
force	uint32, optional	If 1, forces the configuration file to saved, regardless if changes are detected, default is 0

Manual for save

servCert

Description: Display the current server certificate in PEM format.

Security.roles: sys.manage

Manual for servCert

shutdown

Description: Stop the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
reason	string, optional, {length:The value length must be between 0 and 512, inclusive}	The reason for shutting down the service
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart

Manual for shutdown

statHist

Description: Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time for retrieving stats
time2	date-time, optional	The ending time for retrieving stats
timeFormat	string, optional, {enum-one:The value must be one of the	Specify the time format for each stat snapshot, default is posix (seconds since 1970)

Parameter	Type, Options	Description
	following: posix|simple}	
include	string, optional	Comma separated list of stats to include (wildcards allowed)
exclude	string, optional	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)
onChange	string, optional	Comma separated list of stats that reduce the result set to only when at least one of the stats change from their previous value
reduce	bool, optional	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result
showAll	bool, optional	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Manual for statHist

telemtry

Description: Returns telemetry information about the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
level	uint32, optional, {signed:The value must be between 1 and 3, inclusive}	The level of information to be sent by the telemetry request. A value of 1 (default) will send all stat nodes, a value of 2 will send stat and config nodes, and a value of 3 will send stat nodes, config nodes and usage metrics for rules, parsers and devices (where applicable).

Parameter	Type, Options	Description
include	string, optional	Comma separated list of nodes to include (wildcard usage is ok).
exclude	string, optional	Comma separated list of nodes to exclude (wildcard usage is ok). If "include" is also specified, exclusions are applied to the set of nodes specified by "include".
url	string, optional	Endpoint URL to send the JSON response.
op	string, optional, {enum-one:The value must be one of the following: get|configure}	The telemetry operation to execute. Default is 'get'.

/users node

API Messages

addOrMod

Description: Add a new user or update an existing user in the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to add or update, must be alphanumeric or @!#$%&'+-=?^_`{|}~.[]\
password	string, optional	The user's password
pwdIsHashed	bool, optional	Set to true if the password has already been hashed
groups	string, optional	The groups the user belongs to

Parameter	Type, Options	Description
authType	string, optional, {enum-one:The value must be one of the following: netwitness|pam}	The authentication system to use, "netwitness" is the default
queryTimeout	uint32, optional, {unsigned:The value must be between 0 and 43200, inclusive}	The maximum number of minutes that a query is allowed to execute for this user. Zero means unlimited.
displayName	string, optional	Display name for this user account
email	string, optional	Email address for this account
description	string, optional	Description of this user account
queryPrefix	string, optional	Query filter applied to every query performed by this user account
sessionThreshold	uint32, optional, {unsigned:The value must be between 0 and 4294967295, inclusive}	Query optimization which will extrapolate the remaining session counts when they exceed this value

auths

Description: Get supported authentication types

Security.roles: everyone

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete a user from the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to delete

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

unlock

Description: Unlock a locked out account

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string	The username or account to unlock

whoAmI

Description: Returns information about the authenticated user

Security.roles: everyone

Broker service

Introduction

Broker aggregates and store database ranges from multiple core devices and provides a Concentrator like SDK API to treat multiple devices as a single query-able entity. In common deployments Brokers aggregate from multiple Concentrators and sub Brokers.

Broker does not aggregate data like session and meta.

API Usage

Broker service provides API interface through its tree form as nodes, node messages and child nodes.

The following are few examples of the several API nodes and messages supported by Broker:

- /broker node messages:
    - `start` and `stop` aggregation
    - `repair` to reset the global summary based on the device ranges
    - `add` a new device for aggregation
    - `/broker/stats` node for broker aggregation stats information
- /sdk node messages:
    - `reconfig` sdk for optimal settings
    - `summary` to retrieve summary information from the database ranges
    - `values` to perform a value count query and return the matching values for a report
    - `query` to perform query against meta database of upstream devices
    - `language` of meta keys for key info, format, description etc...
    - `/sdk/stats` node for queries active, pending and other sdk operations
- /index node messages:
    - `ranges` to show sessionid ranges for all devices

    - `gaps` to show the gaps and session virtual range mappings per device
    - `save` to save broker device mappings to disk
- /sys node messages:
    - `statHist` to retrieve historical stats from stats database
    - `/sys/stats` node for service system stats information

Sample Request

The following is a sample request to get top 3 values of ip.src through sdk values on the Broker.
Request path is `/sdk`, message is `msg=values` and parameters are `fieldName=ip.src where=\"time=\'2022-07-06 16:30:00\' - \'2022-07-06 16:34:59\'\" size=3 flags=sessions,sort-total,order-descending threshold=100000`

For queries that are expected to take longer time on execution an additional URL parameter `expiry=0` can be added to disable default REST request timeout and continue with user profile timeout.

Using URL encoded request

```
$ curl 'https://broker-host:50103/sdk?msg=values&fieldName=ip.src%20where=\"time=\'2022-07-06%2016:30:00\'%20-%20\'2022-07-06%2016:34:59\'\"&size=3&flags=sessions,sort-total,order-descending&threshold=100000&expiry=0' \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1' \
   -u 'username:password'
```

Using special content type

```
Content Type: application/x-netwitness-string-params

$ curl 'https://broker-host:50103/sdk?msg=values&expiry=0' \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-netwitness-string-params' \
   -d "fieldName=ip.src where=\"time=\'2022-07-06 16:30:00\' - \'2022-07-06 16:34:59\'\" size=3 flags=sessions,sort-total,order-descending threshold=100000" \
   -u 'username:password'
```

Sample Response

HTTP Headers

```
HTTP/1.1 200 OK
Content-Length: 769
Connection: Keep-Alive
Pragma: no-cache
Expires: -1
Cache-Control: no-cache, no-store, must-revalidate
Content-Type: application/json
```

JSON Response


```
{
 "flags": 1074200577,
 "results": {
  "id1": 0,
  "id2": 0,
  "fields": [
   {
    "id1": 79390446,
    "id2": 85993604,
    "count": 104240,
    "format": 128,
    "type": "ip.src",
    "flags": 1,
    "group": 100000,
    "value": "10.XX.XX.81"
   },
   {
    "id1": 79115795,
    "id2": 85854695,
    "count": 101206,
    "format": 128,
    "type": "ip.src",
    "flags": 1,
    "group": 100000,
    "value": "10.XX.XX.83"
   },
   {
    "id1": 79114441,
    "id2": 85709910,
    "count": 100845,
    "format": 128,
    "type": "ip.src",
    "flags": 1,
    "group": 100000,
    "value": "10.XX.XX.152"
   }
```

```
  ]
 },
 "resultsInfo": {
  "devices": "XX.XX.XX.85:50005=on XX.XX.XX.84:56005=on",
  "offlineDeviceCount": "0",
  "onlineDeviceCount": "2",
  "deviceElapsedTime": "XX.XX.XX.85:50005=00:00:00 XX.XX.XX.84:56005=00:00:00",
  "uuidMapping": "XX.XX.XX.85:50005=7bae6dd3-c277-4db8-ba06-edc77298601c
XX.XX.XX.84:56005=58794d2a-e3ff-4a23-a931-c9e0fbc2ebc8"
 }
}
```


Sample Request

The following is a sample request to select 3 values of ip.src in ssl traffic originating from a country using sdk query on the Broker.
Request path is `/sdk`, message is `msg=query` and parameters are `query="select ip.src where time=\'2022-07-06 18:50:00\' - \'2022-07-06 18:54:59\' && tcp.dstport=443 && country.src='United States'" size=3`

Using special content type

```
$ curl 'https://broker-host:50103/sdk?msg=values&expiry=0' \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-netwitness-string-params' \
   -d "query=\"select ip.src where time=\'2022-07-06 18:50:00\' - \'2022-07-06
18:54:59\' && tcp.dstport=443 && country.src='United States'\" size=3" \
   -u 'username:password'
```

Sample Response

HTTP Headers

```
HTTP/1.1 200 OK
Content-Length: 786
Connection: Keep-Alive
Pragma: no-cache
Expires: -1
Cache-Control: no-cache, no-store, must-revalidate
Content-Type: application/json
```

JSON Response

```
{
 "flags": 1074200577,
 "results": {
  "id1": 2009111894,
  "id2": 2370947558,
  "fields": [
   {
    "id1": 2009111743,
    "id2": 2009111743,
    "count": 0,
    "format": 128,
    "type": "ip.src",
    "flags": 0,
    "group": 79178639,
    "value": "11.XX.XX.152"
   },
   {
    "id1": 2009111773,
    "id2": 2009111773,
    "count": 0,
    "format": 128,
    "type": "ip.src",
    "flags": 0,
    "group": 79178640,
    "value": "11.XX.XX.152"
   },
   {
    "id1": 2009111893,
    "id2": 2009111893,
    "count": 0,
    "format": 128,
    "type": "ip.src",
    "flags": 0,
    "group": 79178644,
    "value": "11.XX.XX.198"
   }
  ]
 },
 "resultsInfo": {
  "devices": "10.XX.XX.85:50005=on 10.XX.XX84:56005=on",
  "offlineDeviceCount": "0",
  "onlineDeviceCount": "2",
  "deviceElapsedTime": "10.XX.XX.85:50005=00:00:00 10.XX.XX.84:56005=00:00:00",
  "uuidMapping": "10.XX.XX.85:50005=7bae6dd3-c277-4db8-ba06-edc77298601c
10.XX.XX.84:56005=58794d2a-e3ff-4a23-a931-c9e0fbc2ebc8"
 }
}
```

`Note`: The API Messages section for the path nodes provide more details on API messages available and any associated manual pages.

/broker node

API Messages

add

Description: Add a new device for aggregation

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, optional, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to add (ip:port or [ipv6]:port), either this or ip must be set
ip	string, optional	The IP address of the device, either this or device must be set
port	uint32, optional, {unsigned:The value must be between 0 and 65535, inclusive}	The port number the device is listening on
username	string	The username to run as
password	string, optional	The account password. Assumes service trust aggregation if no password is provided.
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup

Parameter	Type, Options	Description
reconnect	string, optional	If true, will automatically reconnect when not connected
name	string, optional	Device collection name if device is a collection

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete an existing aggregation device

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to delete
name	string, optional	Device collection name if device is a collection

edit

Description: Edit an existing aggregation device

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to edit
username	string, optional	The username to run as
password	string, optional	The account password. For service trust aggregation you must not provide this parameter
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup
reconnect	string, optional	If true, will automatically reconnect when not connected
name	string, optional	Device collection name if device is a collection
switchToTrustedAuth	bool, optional	If true, removes any password associated with this account and will use trusted authentication henceforth. When switching to trusted, only the 'device' parameter is allowed.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

repair

Description: Reset the global summary based on the device ranges

Security.roles: index.manage

reset

Description: Reset data, index, manifests, stats, configuration, or logs for this service. Data automatically deletes index and stats, unless filesCreatedAfter is specified. Service is automatically restarted. Example arguments: data=1 config=1 log=1 This example will reset data, index, logs, and configuration index=1 This example will reset the index only manifest=1 This example will delete everything in the manifest directory filesCreatedAfter="2015-12-01 14:00:00" This example will delete all session, meta and packet files created on or after Dec 1st, 2015 2pm (UTC) from the service. All other files will remain. The index will not be touched, but upon restart will be truncated to match the last session in the session database.

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
data	bool, optional	Reset data, automatically resets index (may not be applicable to all services)
index	bool, optional	Reset index, only valid if data is not reset (may not be applicable to all services)
manifest	bool, optional	Reset all manifests stored in the long term manifest directory (may not be applicable to all services)
config	bool, optional	Reset configuration settings to default
stats	bool, optional	Reset the stats database

Parameter	Type, Options	Description
log	bool, optional	Reset the log database
filesCreatedAfter	date-time, optional	Delete all database files (session, meta, packets) created after a certain date. The index will automatically get rolled back on restart. Not valid with option index.
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart
force	bool, optional	Force reset without confirmation

resetMax

Description: Resets all max stats, including device max stats back to zero.

Security.roles: concentrator.manage

start

Description: Starts aggregation

Security.roles: concentrator.manage

status

Description: Change the online/offline status of an aggregation device

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to change status

Parameter	Type, Options	Description
status	string, optional, {enum-one:The value must be one of the following: online|offline}	Whether to take device offline or put online
name	string, optional	Device collection name if device is a collection

stop

Description: Stops aggregation

Security.roles: concentrator.manage

whoAgg

Description: Returns information on who is aggregating from this service

Security.roles: concentrator.manage

Parameters:

Parameter	Type, Options	Description
clean	bool, optional	If true, will delete any aggregation trackers that no longer have a valid channel

/connections node

API Messages

closeAll

Description: Closes all connections

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: native|rest|amqp}	Indicates the type of connection to close, either native, rest or amqp. Default is all types.
limit	uint32, optional	If non-zero, will stop closing connections after limit is reached. NOTE: It starts closing the most idle connections first.
idleLongerThan	string, optional	A time duration in HH:MM:SS or ##{d,h,m,s} format. Any connection that has been idle longer than the passed in duration is eligible to be closed. Connections that have seen activity within the duration window will not be closed. Examples: "01:05:10" (1 hour, 5 mins, 10 secs) or "65m" (65 minutes)
exceptMe	bool, optional	If true, will not close the connection that initiated the command.

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

netSpeed

Description: Network Speed Test

Security.roles: connections.manage,aggregate

Parameters:

Parameter	Type, Options	Description
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 1024 and 536870912, inclusive}	The size of each message sent, default is 64k
count	uint32, optional, {unsigned:The value must be between 1 and 536870912, inclusive}	The number of messages to send, default is 16k

/index node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

gaps

Description: Shows the gaps and other error checking in session virtual range mappings per device

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting session id
id2	uint64, optional	The ending session id

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

ranges

Description: Shows the session id ranges for all devices or a single device

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
device	string, optional	The device to report ranges for, defaults to 'all' devices
name	string, optional	The device collection name if the device is a collection
id1	uint64, optional	The starting session id
id2	uint64, optional	The ending session id
format	string, optional, {enum-one:The value must be one	The format of the returned report

Parameter	Type, Options	Description
	of the following: tab|comma}	

save

Description: Saves the broker device mappings to disk

Security.roles: index.manage

/logs node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

download

Description: Downloads log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting log number to download
id2	uint64, optional	The ending log number to download
time1	date-time, optional	The starting log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
time2	date-time, optional	The ending log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: start|next|cancel}	The operation to perform (start, next, cancel)
logTypes	string, optional	The type(s) of log messages to retrieve (debug,info,audit,warning,failure)
match	string, optional	Case insensitive string to match in each log message
regex	string, optional	Regular expression to match in each log message
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)
batchSize	uint32, optional, {unsigned:The value must be between 2000 and 10000, inclusive}	The number of logs to retrieve per batch, default is 2000

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters:	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')

Parameter	Type, Options	Description
	abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

pull

Description: Downloads N log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The first log id number to retrieve, this is mutually exclusive with id2
id2	uint64, optional	The last log id number that will be sent, defaults to most recent log message when id1 or id2 is not sent
count	uint32, optional, {unsigned:The value must be between 1 and 10000, inclusive}	The number of logs to pull
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)

timeRoll

Description: Delete log entries that exceed a given age

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
timeCalc	string, optional, {enum-any:The value must be one or more of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 60, inclusive}	Remove log entries older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 24, inclusive}	Remove log entries older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 365, inclusive}	Remove log entries older than the given number of days
date	string, optional	Remove log entries older than the given date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

/rest node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

/sdk node

Manual for /sdk

API Messages

aliases

Description: Retrieves aliases for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string	The language key to retrieve aliases for
value	string, optional	An optional value to return a specific alias for

cancel

Description: Cancel a running query, regardless of the user

Security.roles: sdk.manage,sdk.meta

Parameters:

Parameter	Type, Options	Description
handle	uint32	The channel handle running the query

Parameter	Type, Options	Description
uuid	string, optional	The device UUID if taking a device offline instead of canceling the entire query (Broker only). Passing uuid will propagate the request upstream to find the first matching device to cancel.

content

Description: Returns the packet content for a session

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session id to retrieve packet content for
packet	uint64, optional	The packet id of the first packet of the session, in case the session has rolled out
pinId	string, optional	The Pin ID of the session to retrieve from long term storage. If provided, the packet and session parameters must not be provided.
maxSize	size, optional	The max number of bytes to return, zero means no limit
failOnCacheMiss	bool, optional	If true, the content call fails if the NWD is not in the cache. This is used by optimized caching code only.
renderType	string, optional	The render type (see NwContentTypes in NwSDK.h), pass zero to pick best option. Not specifying this parameter always results in an NWD and renderFlags and renderOptions are not considered. The following values are also supported: 'pcap', 'file-list-json', 'session-meta-file-list-json', 'nwd', 'log', 'files'
renderFlags	uint32, optional	Bitwise mask to control options, (see NwContentFlags in NwSDK.h)

Parameter	Type, Options	Description
renderOptions	string, optional	An encoded string params containing additional options for controlling the operation.

Manual for content

count

Description: Returns the number of child nodes

Security.roles: everyone

delCache

Description: Removes all cached NWD files from either the normal NWD cache or the pin cache

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-one:The value must be one of the following: nwd|pin|all}	Specifies which cache to delete all NWD files from. The default is 'nwd'. Passing 'all' will delete from all caches.

deviceId

Description: Given a session id, this message retrieves the session id and the device it maps to

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id that will be translated into the device name and session id.

Manual for deviceId

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyz ABCDEFGHIJKLMNOPQRSTU VWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

hierarch

Description: Report the set of devices attached to this service

Security.roles: sdk.meta

Manual for hierarch

info

Description: Returns detailed information about the node

Security.roles: everyone

keyrefs

Description: Retrieves the set of entity keys

Security.roles: sdk.meta

Manual for keyrefs

language

Description: Retrieves the language definition

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting language id
id2	uint64, optional	The ending language id
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
options	string, optional	Extra options, currently unused
flags	string, optional	Optional flags to configure how the results are returned. Can be a number (bitwise mask) or comma separated values like default, no-count, quick-count or full-count.
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
queryPriority	int32, optional, {signed:The value must be zero OR	The query priority for this query

Parameter	Type, Options	Description
	between -20 and 20, inclusive}	
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /broker/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

msearch

Description: Search for pattern matches in many sessions or packets

Security.roles: sdk.content & sdk.meta

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	The session ID ranges to search
packets	string, optional	The packet ID ranges to search
search	string	The search string to use
where	string, optional	The where clause used to identify sessions to consider for the search
limit	uint64, optional	The maximum number of sessions to traverse for this search
size	uint64, optional	The maximum number of results to return for this search
flags	string, {enum-any:The value must be one or more of the following: regex|sp|sm|ci|pre|post|ds|precache|si}	Flags to use for search. This is a comma separated list of one or more of the flag values: regex,sp,sm,ci,pre,post,ds,precache,si
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /broker/devices or match the /sys/stats/uuid of each device to be queried.

Parameter	Type, Options	Description
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for msearch

packets

Description: Stream packets back based on the input parameters provided

Security.roles: sdk.content & sdk.packets

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: setup|start|cancel|processed}	The operation to perform (start, cancel, processed)
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) whose packets will be streamed back.
packets	string, optional	A comma separated list of packet ids or session/packet ids (#&#) that will be streamed back.
where	string, optional	Where clause which will be evaluated to determine which sessions to stream back
time1	date-time, optional	A starting time range (UTC) where matching packets will be streamed back ("2010-Apr-20 09:00:00")
time2	date-time, optional	An ending time range (UTC) where matching packets will be streamed back ("2010-Apr-20 10:00:00")

Parameter	Type, Options	Description
flags	uint32, optional	Additional flags as defined by the NwPackets SDK function
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
devices	string, optional	Optional list of comma separated device names that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /broker/devices
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for packets

pin

Description: Pin a session in the long term cache so it can be retrieved even after the meta or packets have rolled out. To retrieve, you must make a content call with the Pin ID that is returned from this command. The op={unpin,validate} commands all take one or more comma separated Pin IDs in the pinId parameter and returns the status for the Pin IDs that were recognized by a service.

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session ID to pin. This cannot be used with the 'sessions' parameter.

Parameter	Type, Options	Description
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) to pin. This cannot be used with the 'session' parameter.
pinId	string, optional	One or more Pin IDs, separated by commas
op	string, optional, {enum-one:The value must be one of the following: ls|unpin|validate}	Various operations that can be performed. Will be submitted to upstream devices if not configured for long term cache.

Manual for pin

precache

Description: Efficiently caches NWD files for future retrieval via the content message

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) that will be saved in the NWD cache directory
packets	string, optional	A comma separated list of session and packet ids (#p#) that will be saved in the NWD cache directory
waitOnComplete	bool, optional	If true, will wait for precache to complete before returning a response.

query

Description: Performs a query against the meta database

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id (to run the query from most recent to oldest meta, make id1 larger than id2)
id2	uint64, optional	The ending meta id
size	uint32, optional	The max number of entries to return, or just stream back all results if zero
query	string, optional	The query string to use
flags	string, optional	The flags to use for query. Can be a number (bitwise mask) or comma separated values like query-log.
threshold	uint64, optional	Query optimization to stop processing results after the threshold is reached (useful with select aggregate functions). Zero means no threshold (the default).
partition	uint32, optional	Used to partition the result set across multiple clients. Must be greater than zero and less than or equal to totalPartitions
totalPartitions	uint32, optional	The total number of clients that will be submitting the same query for partitioning.
search	string, optional	A text search clause to apply to the where clause of this query
streamLimit	uint32, optional	Limit the number of sessions returned by a streaming query
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then

Parameter	Type, Options	Description
		all devices are queried. The device names must match the list defined under /broker/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for query

reconfig

Description: Calculates default values for some of the config nodes.

Security.roles: sdk.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the config nodes with the defaults
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

search

Description: Searches for matches in session/packet content

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to search
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
search	string	The search string to use

Manual for search

session

Description: Retrieves the meta id range for the session range

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64	The starting session id (inclusive). Pass zero to scope to lowest valid id.
id2	uint64	The ending session id (inclusive). Pass zero to scope to highest valid id.

Manual for session

summary

Description: Retrieves summary information from the databases

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
flags	string, optional	Optional SDK flags. Can be a number (bitwise mask) or comma separated values like default or ignore-cache.

Manual for summary

timeline

Description: Returns the count of sessions/size/packets in discrete time intervals

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS), if greater than time2, then the whole time range is returned
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
timezone	int32, optional	The timezone of the local computer to receive the data, hour offset from GMT [-13 to 13]
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for the timeline. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, order-ascending, order-descending, ignore-cache or clear-cache.
where	string, optional	Optional where clause for filtering the data
bookendHours	uint32, optional	The number of hours outside the given time1, time2 which will become part of the cache. Default is 24 hours.
fieldName	string, optional	The time field to retrieve values for, defaults to 'time'
queryPriority	int32, optional, {signed:The value must be	The query priority for this query

Parameter	Type, Options	Description
	zero OR between -20 and 20, inclusive}	
search	string, optional	A text search clause to apply to the where clause of this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /broker/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for timeline

validate

Description: Validate a query or values call without execution

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: query|values}	The operation to validate
validateLanguage	bool, optional	If true (default is false), language tokens and syntax will be validated, otherwise just syntax will be validated

Parameter	Type, Options	Description
fieldName	string, optional	Required if validating a 'values' operation. The field to retrieve values for.
where	string, optional	Required if validating a 'values' operation. The filter criteria to be validated.
query	string, optional	Required if validating a 'query' operation. The query string to be validated.
json	bool, optional	Structure response in JSON syntax

Manual for validate

values

Description: Performs a value count query and returns the matching values for a report

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id
id2	uint64, optional	The ending meta id
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for values. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, aggregate, sort-total, sort-value, order-ascending, order-descending, ignore-cache, clear-cache, or database-scan.
threshold	uint64, optional	Query optimization to stop processing large session counts

Parameter	Type, Options	Description
fieldName	string	The field to retrieve values for
where	string, optional	The filter criteria for the values
aggregateFunction	string, optional, {enum-one:The value must be one of the following: count|sum}	The aggregate operation to be applied where the aggregate flag is set
aggregateFieldName	string, optional	The meta field to aggregate in the aggregateFunction
min	string, optional	The lower limit of values to return. Only values greater than this value will be returned.
max	string, optional	The upper limit of values to return. Only values less than this value will be returned.
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /broker/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for values

xforms

Description: Retrieves transforms for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string, optional	An optional language key to retrieve transform for, all transforms returned if not provided

/sys node

Manual for /sys

API Messages

caCert

Description: Display or delete trusted CA certs

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: list|delete|add}	CA action to perform (list,delete,add)
id	string, optional	ID of CA certificate to remove when op=delete

Manual for caCert

count

Description: Returns the number of child nodes

Security.roles: everyone

fileEdit

Description: View and edit system configuration files

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: dir|get|getBackup|put|create|delete}	The operation to perform (dir, get, getBackup, put, create, delete)
type	string, optional, {enum-one:The value must be one of the following: parser|filter}	The type of file, if not provided will check file extension
filename	string, optional, {length:The value length must be between 1 and 64, inclusive}	File to create, edit, get or delete

Manual for fileEdit

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

peerCert

Description: Display or modify trusted peer certificates

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|delete|add}	Operation to perform on the list of trusted peers (list,delete,add)
id	string, optional	ID of peer certificate to remove when op=delete

Manual for peerCert

save

Description: Forces a save of system settings

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
force	uint32, optional	If 1, forces the configuration file to saved, regardless if changes are detected, default is 0

Manual for save

servCert

Description: Display the current server certificate in PEM format.

Security.roles: sys.manage

Manual for servCert

shutdown

Description: Stop the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
reason	string, optional, {length:The value length must be between 0 and 512, inclusive}	The reason for shutting down the service
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart

Manual for shutdown

statHist

Description: Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time for retrieving stats
time2	date-time, optional	The ending time for retrieving stats
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format for each stat snapshot, default is posix (seconds since 1970)
include	string, optional	Comma separated list of stats to include (wildcards allowed)
exclude	string, optional	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)

Parameter	Type, Options	Description
onChange	string, optional	Comma separated list of stats that reduce the result set to only when at least one of the stats change from their previous value
reduce	bool, optional	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result
showAll	bool, optional	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Manual for statHist

telemtry

Description: Returns telemetry information about the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
level	uint32, optional, {signed:The value must be between 1 and 3, inclusive}	The level of information to be sent by the telemetry request. A value of 1 (default) will send all stat nodes, a value of 2 will send stat and config nodes, and a value of 3 will send stat nodes, config nodes and usage metrics for rules, parsers and devices (where applicable).
include	string, optional	Comma separated list of nodes to include (wildcard usage is ok).
exclude	string, optional	Comma separated list of nodes to exclude (wildcard usage is ok). If "include" is also specified, exclusions are applied to the set of nodes specified by "include".
url	string, optional	Endpoint URL to send the JSON response.

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: get|configure}	The telemetry operation to execute. Default is 'get'.

/users node

API Messages

addOrMod

Description: Add a new user or update an existing user in the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to add or update, must be alphanumeric or @!#$%&'+-=?^_`{|}~.[]\
password	string, optional	The user's password
pwdIsHashed	bool, optional	Set to true if the password has already been hashed
groups	string, optional	The groups the user belongs to
authType	string, optional, {enum-one:The value must be one of the following: netwitness|pam}	The authentication system to use, "netwitness" is the default
queryTimeout	uint32, optional, {unsigned:The value must be between 0 and 43200, inclusive}	The maximum number of minutes that a query is allowed to execute for this user. Zero means unlimited.

Parameter	Type, Options	Description
displayName	string, optional	Display name for this user account
email	string, optional	Email address for this account
description	string, optional	Description of this user account
queryPrefix	string, optional	Query filter applied to every query performed by this user account
sessionThreshold	uint32, optional, {unsigned:The value must be between 0 and 4294967295, inclusive}	Query optimization which will extrapolate the remaining session counts when they exceed this value

auths

Description: Get supported authentication types

Security.roles: everyone

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete a user from the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters:	The username to delete

Parameter	Type, Options	Description
	abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

unlock

Description: Unlock a locked out account

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string	The username or account to unlock

whoAmI

Description: Returns information about the authenticated user

Security.roles: everyone

Archiver service

Introduction

Archiver is similar to a Concentrator in all but one respect. In addition to aggregating sessions and meta, it would also aggregate the raw logs. It typically has smaller index meant for compliance reports and not investigations. Because it aggregates logs, it is almost always paired with one or more LogDecoders, not Decoders.
The aggregated data is stored in collections which can be configured through rules for example:
syslog collection for all logs with `device.type=syslog`. By default all the aggregated data would be part of `default` collection.

Requests for the raw data (logs) are not forwarded to LogDecoder since the data is already stored on the Archiver.

API Usage

Archiver service provides API interface through its tree form as nodes, node messages and child nodes.

The following are few examples of the several API nodes and messages supported by Archiver:

- /archiver node messages:
    - `start` and `stop` aggregation
    - `repair` to reset the global summary based on the collection ranges
    - `add` a new device for aggregation
    - `/archiver/stats` node for archiver aggregation stats information
- /sdk node messages:
    - `reconfig` sdk for optimal settings
    - `summary` to retrieve summary information from the database ranges
    - `values` to perform a value count query and return the matching values for a report
    - `query` to perform query against meta database
    - `language` of meta keys for key info, format, description etc...
    - `/sdk/stats` node for queries active, pending and other sdk operations
- /index node messages:
    - `ranges` to show sessionid ranges for all collections

    - `gaps` to show the gaps and session virtual range mappings per collection
    - `save` to save arciver collection mappings to disk
- /sys node messages:
    - `statHist` to retrieve historical stats from stats database
    - `/sys/stats` node for service and system stats information

Archiver also support per collection API node paths and messages. In order to access API node and messags for a specific collection the request path needs to be `/archiver/collections/{collection-name}` e.g for default collection it would be `/archiver/collections/default`.

`Note`: Queries issued on root `/sdk` path would be run against all available collections and queries issued on collection sdk path would be run against that collection.

Sample Request

The following is a sample request query on Archiver to get 4 values of distinct ip.dst, grouped by ip.src and ordered by distinct ip.dst in descending order.

Request path is `/sdk`, message is `msg=query` and parameters are `query=\"select distinct(ip.dst) where time=\'2022-07-06 16:50:00\' - \'2022-07-06 16:54:59\' GROUP BY ip.src ORDER BY distinct(ip.dst) desc \" size=4"`

For queries that are expected to take longer time on execution an additional URL parameter `expiry=0` can be added to disable default REST request timeout and continue with user profile timeout.

Using URL encoded request

```
$ curl 'https://archiver-host:50108/sdk?msg=query&expiry=0&query=select%20distinct(ip.dst)%20where%20time='2022-07-06%2016:50:00'%20-%20'2022-07-06%2016:54:59'%20GROUP%20BY%20ip.src%20ORDER%20BY%20distinct(ip.dst)%20desc&size=4'\
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1' \
   -u 'username:password'
```

Using special content type

```
Content Type: application/x-netwitness-string-params
```

```
$ curl 'https://archiver-host:50108/sdk?msg=query&expiry=0' \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-netwitness-string-params' \
   -d "query=\"select distinct(ip.dst) where time=\'2022-07-06 16:50:00\' - \'2022-07-06
16:54:59\' GROUP BY ip.src ORDER BY distinct(ip.dst) desc  \" size=4" \
   -u 'username:password'
```

HTTP Headers

```
HTTP/1.1 200 OK
Content-Length: 561
Connection: Keep-Alive
Pragma: no-cache
Expires: -1
Cache-Control: no-cache, no-store, must-revalidate
Content-Type: application/json
```

JSON Response


```
{
 "flags": 1074200577,
 "results": {
  "id1": 1231176344,
  "id2": 1267146902,
  "fields": [
   {
    "id1": 1126481600,
    "id2": 1212151333,
    "count": 142,
    "format": 128,
    "type": "ip.dst",
    "flags": 6,
    "group": 0,
    "value": "192.XX.XX.246"
   },
   {
    "id1": 1126676313,
    "id2": 1231176343,
    "count": 342,
    "format": 128,
    "type": "ip.dst",
    "flags": 6,
    "group": 0,
    "value": "10.XX.XX.22"
   },
   {
    "id1": 1126779641,
    "id2": 1225198229,
```

```
      "count": 320,
      "format": 128,
      "type": "ip.dst",
      "flags": 6,
      "group": 0,
      "value": "10.XX.XX.174"
    },
    {
      "id1": 1126936749,
      "id2": 1224923290,
      "count": 660,
      "format": 128,
      "type": "ip.dst",
      "flags": 6,
      "group": 0,
      "value": "10.XX.XX.235"
    }
  ]
 }
}
```


`Note`: The API Messages section for the path nodes provide more details on API messages available and any associated manual pages.

/archiver node

API Messages

add

Description: Add a new device for aggregation

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
device	string, optional, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to add (ip:port or [ipv6]:port), either this or ip must be set

Parameter	Type, Options	Description
ip	string, optional	The IP address of the device, either this or device must be set
port	uint32, optional, {unsigned:The value must be between 0 and 65535, inclusive}	The port number the device is listening on
username	string	The username to run as
password	string, optional	The account password. Assumes service trust aggregation if no password is provided.
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup
reconnect	string, optional	If true, will automatically reconnect when not connected
name	string, optional	Device collection name if device is a collection

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete an existing aggregation device

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to delete
name	string, optional	Device collection name if device is a collection

edit

Description: Edit an existing aggregation device

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to edit
username	string, optional	The username to run as
password	string, optional	The account password. For service trust aggregation you must not provide this parameter
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup
reconnect	string, optional	If true, will automatically reconnect when not connected

Parameter	Type, Options	Description
name	string, optional	Device collection name if device is a collection
switchToTrustedAuth	bool, optional	If true, removes any password associated with this account and will use trusted authentication henceforth. When switching to trusted, only the 'device' parameter is allowed.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

reconfig

Description: Calculates optimal settings for archiver pools and buffers based on the installed hardware.

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective archiver settings, otherwise it will just output the calculations
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids

Parameter	Type, Options	Description
		and/or all-in-one machines where the service will need to share memory.

repair

Description: Reset the global summary based on the collection ranges -or- delete a specified collection's mapping ranges and regenerate them anew.

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
name	string, optional	The name of the collection if regenerating ranges

reset

Description: Reset data, index, manifests, stats, configuration, or logs for this service. Data automatically deletes index and stats, unless filesCreatedAfter is specified. Service is automatically restarted. Example arguments: data=1 config=1 log=1 This example will reset data, index, logs, and configuration index=1 This example will reset the index only manifest=1 This example will delete everything in the manifest directory filesCreatedAfter="2015-12-01 14:00:00" This example will delete all session, meta and packet files created on or after Dec 1st, 2015 2pm (UTC) from the service. All other files will remain. The index will not be touched, but upon restart will be truncated to match the last session in the session database.

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
data	bool, optional	Reset data, automatically resets index (may not be applicable to all services)

Parameter	Type, Options	Description
index	bool, optional	Reset index, only valid if data is not reset (may not be applicable to all services)
manifest	bool, optional	Reset all manifests stored in the long term manifest directory (may not be applicable to all services)
config	bool, optional	Reset configuration settings to default
stats	bool, optional	Reset the stats database
log	bool, optional	Reset the log database
filesCreatedAfter	date-time, optional	Delete all database files (session, meta, packets) created after a certain date. The index will automatically get rolled back on restart. Not valid with option index.
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart
force	bool, optional	Force reset without confirmation

resetMax

Description: Resets all max stats, including device max stats back to zero.

Security.roles: archiver.manage

start

Description: Starts aggregation

Security.roles: archiver.manage

status

Description: Change the online/offline status of an aggregation device

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to change status
status	string, optional, {enum-one:The value must be one of the following: online|offline}	Whether to take device offline or put online
name	string, optional	Device collection name if device is a collection

stop

Description: Stops aggregation

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
flush	bool, optional	If true (default), process all sessions and packets currently in memory. Otherwise, stop immediately, discarding unprocessed sessions and packets.

whoAgg

Description: Returns information on who is aggregating from this service

Security.roles: archiver.manage

Parameters:

Parameter	Type, Options	Description
clean	bool, optional	If true, will delete any aggregation trackers that no longer have a valid channel

/connections node

API Messages

closeAll

Description: Closes all connections

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: native|rest|amqp}	Indicates the type of connection to close, either native, rest or amqp. Default is all types.
limit	uint32, optional	If non-zero, will stop closing connections after limit is reached. NOTE: It starts closing the most idle connections first.
idleLongerThan	string, optional	A time duration in HH:MM:SS or ##{d,h,m,s} format. Any connection that has been idle longer than the passed in duration is eligible to be closed. Connections that have seen activity within the duration window will not be closed. Examples: "01:05:10" (1 hour, 5 mins, 10 secs) or "65m" (65 minutes)
exceptMe	bool, optional	If true, will not close the connection that initiated the command.

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

netSpeed

Description: Network Speed Test

Security.roles: connections.manage,aggregate

Parameters:

Parameter	Type, Options	Description
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 1024 and 536870912, inclusive}	The size of each message sent, default is 64k
count	uint32, optional, {unsigned:The value must be between 1 and 536870912, inclusive}	The number of messages to send, default is 16k

/index node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

gaps

Description: Shows the gaps and other error checking in session virtual range mappings per device

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting session id
id2	uint64, optional	The ending session id

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

ranges

Description: Shows the session id ranges for all devices or a single device

Security.roles: index.manage

Parameters:

Parameter	Type, Options	Description
device	string, optional	The device to report ranges for, defaults to 'all' devices
name	string, optional	The device collection name if the device is a collection
id1	uint64, optional	The starting session id

Parameter	Type, Options	Description
id2	uint64, optional	The ending session id
format	string, optional, {enum-one:The value must be one of the following: tab|comma}	The format of the returned report

save

Description: Saves the broker device mappings to disk

Security.roles: index.manage

/logs node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

download

Description: Downloads log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting log number to download
id2	uint64, optional	The ending log number to download
time1	date-time, optional	The starting log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS

Parameter	Type, Options	Description
time2	date-time, optional	The ending log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
op	string, {enum-one:The value must be one of the following: start|next|cancel}	The operation to perform (start, next, cancel)
logTypes	string, optional	The type(s) of log messages to retrieve (debug,info,audit,warning,failure)
match	string, optional	Case insensitive string to match in each log message
regex	string, optional	Regular expression to match in each log message
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)
batchSize	uint32, optional, {unsigned:The value must be between 2000 and 10000, inclusive}	The number of logs to retrieve per batch, default is 2000

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values

Parameter	Type, Options	Description
		like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

pull

Description: Downloads N log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The first log id number to retrieve, this is mutually exclusive with id2
id2	uint64, optional	The last log id number that will be sent, defaults to most recent log message when id1 or id2 is not sent
count	uint32, optional, {unsigned:The value must be between 1 and 10000, inclusive}	The number of logs to pull
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)

timeRoll

Description: Delete log entries that exceed a given age

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
timeCalc	string, optional, {enum-any:The value must be one or more of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 60, inclusive}	Remove log entries older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 24, inclusive}	Remove log entries older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 365, inclusive}	Remove log entries older than the given number of days
date	string, optional	Remove log entries older than the given date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

/rest node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder,

Parameter	Type, Options	Description
		session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

/sdk node

Manual for /sdk

API Messages

aliases

Description: Retrieves aliases for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string	The language key to retrieve aliases for
value	string, optional	An optional value to return a specific alias for

cancel

Description: Cancel a running query, regardless of the user

Security.roles: sdk.manage,sdk.meta

Parameters:

Parameter	Type, Options	Description
handle	uint32	The channel handle running the query

Parameter	Type, Options	Description
uuid	string, optional	The device UUID if taking a device offline instead of canceling the entire query (Broker only). Passing uuid will propagate the request upstream to find the first matching device to cancel.

content

Description: Returns the packet content for a session

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session id to retrieve packet content for
packet	uint64, optional	The packet id of the first packet of the session, in case the session has rolled out
pinId	string, optional	The Pin ID of the session to retrieve from long term storage. If provided, the packet and session parameters must not be provided.
maxSize	size, optional	The max number of bytes to return, zero means no limit
failOnCacheMiss	bool, optional	If true, the content call fails if the NWD is not in the cache. This is used by optimized caching code only.
renderType	string, optional	The render type (see NwContentTypes in NwSDK.h), pass zero to pick best option. Not specifying this parameter always results in an NWD and renderFlags and renderOptions are not considered. The following values are also supported: 'pcap', 'file-list-json', 'session-meta-file-list-json', 'nwd', 'log', 'files'
renderFlags	uint32, optional	Bitwise mask to control options, (see NwContentFlags in NwSDK.h)

Parameter	Type, Options	Description
renderOptions	string, optional	An encoded string params containing additional options for controlling the operation.

Manual for content

count

Description: Returns the number of child nodes

Security.roles: everyone

delCache

Description: Removes all cached NWD files from either the normal NWD cache or the pin cache

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-one:The value must be one of the following: nwd|pin|all}	Specifies which cache to delete all NWD files from. The default is 'nwd'. Passing 'all' will delete from all caches.

deviceId

Description: Given a session id, this message retrieves the session id and the device it maps to

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id that will be translated into the device name and session id.

Manual for deviceId

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

hierarch

Description: Report the set of devices attached to this service

Security.roles: sdk.meta

Manual for hierarch

info

Description: Returns detailed information about the node

Security.roles: everyone

keyrefs

Description: Retrieves the set of entity keys

Security.roles: sdk.meta

Manual for keyrefs

language

Description: Retrieves the language definition

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting language id
id2	uint64, optional	The ending language id
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
options	string, optional	Extra options, currently unused
flags	string, optional	Optional flags to configure how the results are returned. Can be a number (bitwise mask) or comma separated values like default, no-count, quick-count or full-count.
size	uint32, {unsigned:The value must be between 1	The max number of entries to return

Parameter	Type, Options	Description
	and 1677721, inclusive}	
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /archiver/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for language

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1

Parameter	Type, Options	Description
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

msearch

Description: Search for pattern matches in many sessions or packets

Security.roles: sdk.content & sdk.meta

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	The session ID ranges to search
packets	string, optional	The packet ID ranges to search
search	string	The search string to use
where	string, optional	The where clause used to identify sessions to consider for the search
limit	uint64, optional	The maximum number of sessions to traverse for this search
size	uint64, optional	The maximum number of results to return for this search
flags	string, {enum-any:The value must be one or more of the following: regex|sp|sm|ci|pre|post|ds|precache|si}	Flags to use for search. This is a comma separated list of one or more of the flag values: regex,sp,sm,ci,pre,post,ds,precache,si

Parameter	Type, Options	Description
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /archiver/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for msearch

packets

Description: Stream packets back based on the input parameters provided

Security.roles: sdk.content & sdk.packets

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: setup|start|cancel|processed}	The operation to perform (start, cancel, processed)
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) whose packets will be streamed back.

Parameter	Type, Options	Description
packets	string, optional	A comma separated list of packet ids or session/packet ids (#&#) that will be streamed back.
where	string, optional	Where clause which will be evaluated to determine which sessions to stream back
time1	date-time, optional	A starting time range (UTC) where matching packets will be streamed back ("2010-Apr-20 09:00:00")
time2	date-time, optional	An ending time range (UTC) where matching packets will be streamed back ("2010-Apr-20 10:00:00")
flags	uint32, optional	Additional flags as defined by the NwPackets SDK function
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
devices	string, optional	Optional list of comma separated device names that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /archiver/devices
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for packets

pin

Description: Pin a session in the long term cache so it can be retrieved even after the meta or packets have rolled out. To retrieve, you must make a content call with the Pin ID

that is returned from this command. The op={unpin,validate} commands all take one or more comma separated Pin IDs in the pinId parameter and returns the status for the Pin IDs that were recognized by a service.

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64, optional	The session ID to pin. This cannot be used with the 'sessions' parameter.
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) to pin. This cannot be used with the 'session' parameter.
pinId	string, optional	One or more Pin IDs, separated by commas
op	string, optional, {enum-one:The value must be one of the following: ls|unpin|validate}	Various operations that can be performed. Will be submitted to upstream devices if not configured for long term cache.

Manual for pin

precache

Description: Efficiently caches NWD files for future retrieval via the content message

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
sessions	string, optional	A comma separated list of session ids or session id ranges (#-#) that will be saved in the NWD cache directory
packets	string, optional	A comma separated list of session and packet ids (#p#) that will be saved in the NWD cache directory
waitOnComplete	bool, optional	If true, will wait for precache to complete before returning a response.

query

Description: Performs a query against the meta database

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id (to run the query from most recent to oldest meta, make id1 larger than id2)
id2	uint64, optional	The ending meta id
size	uint32, optional	The max number of entries to return, or just stream back all results if zero
query	string, optional	The query string to use
flags	string, optional	The flags to use for query. Can be a number (bitwise mask) or comma separated values like query-log.
threshold	uint64, optional	Query optimization to stop processing results after the threshold is reached (useful with select aggregate functions). Zero means no threshold (the default).
partition	uint32, optional	Used to partition the result set across multiple clients. Must be greater than zero and less than or equal to totalPartitions
totalPartitions	uint32, optional	The total number of clients that will be submitting the same query for partitioning.
search	string, optional	A text search clause to apply to the where clause of this query
streamLimit	uint32, optional	Limit the number of sessions returned by a streaming query
queryPriority	int32, optional, {signed:The	The query priority for this query

Parameter	Type, Options	Description
	value must be zero OR between -20 and 20, inclusive}	
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /archiver/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for query

reconfig

Description: Calculates default values for some of the config nodes.

Security.roles: sdk.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the config nodes with the defaults
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

search

Description: Searches for matches in session/packet content

Security.roles: sdk.content

Parameters:

Parameter	Type, Options	Description
session	uint64	The session id to search
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
search	string	The search string to use

Manual for search

session

Description: Retrieves the meta id range for the session range

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64	The starting session id (inclusive). Pass zero to scope to lowest valid id.
id2	uint64	The ending session id (inclusive). Pass zero to scope to highest valid id.

Manual for session

summary

Description: Retrieves summary information from the databases

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
flags	string, optional	Optional SDK flags. Can be a number (bitwise mask) or comma separated values like default or ignore-cache.

Manual for summary

timeline

Description: Returns the count of sessions/size/packets in discrete time intervals

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS), if greater than time2, then the whole time range is returned
time2	date-time, optional	The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS)
timezone	int32, optional	The timezone of the local computer to receive the data, hour offset from GMT [-13 to 13]
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for the timeline. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, order-ascending, order-descending, ignore-cache or clear-cache.
where	string, optional	Optional where clause for filtering the data
bookendHours	uint32, optional	The number of hours outside the given time1, time2 which will become part of the cache. Default is 24 hours.

Parameter	Type, Options	Description
fieldName	string, optional	The time field to retrieve values for, defaults to 'time'
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under /archiver/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for timeline

validate

Description: Validate a query or values call without execution

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of	The operation to validate

Parameter	Type, Options	Description
	the following: query|values}	
validateLanguage	bool, optional	If true (default is false), language tokens and syntax will be validated, otherwise just syntax will be validated
fieldName	string, optional	Required if validating a 'values' operation. The field to retrieve values for.
where	string, optional	Required if validating a 'values' operation. The filter criteria to be validated.
query	string, optional	Required if validating a 'query' operation. The query string to be validated.
json	bool, optional	Structure response in JSON syntax

Manual for validate

values

Description: Performs a value count query and returns the matching values for a report

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting meta id
id2	uint64, optional	The ending meta id
size	uint32, {unsigned:The value must be between 1 and 1677721, inclusive}	The max number of entries to return
flags	string, optional	The flags to use for values. Can be a number (bitwise mask) or comma separated values like sessions, size,

Parameter	Type, Options	Description
		packets, aggregate, sort-total, sort-value, order-ascending, order-descending, ignore-cache, clear-cache, or database-scan.
threshold	uint64, optional	Query optimization to stop processing large session counts
fieldName	string	The field to retrieve values for
where	string, optional	The filter criteria for the values
aggregateFunction	string, optional, {enum-one:The value must be one of the following: count|sum}	The aggregate operation to be applied where the aggregate flag is set
aggregateFieldName	string, optional	The meta field to aggregate in the aggregateFunction
min	string, optional	The lower limit of values to return. Only values greater than this value will be returned.
max	string, optional	The upper limit of values to return. Only values less than this value will be returned.
queryPriority	int32, optional, {signed:The value must be zero OR between -20 and 20, inclusive}	The query priority for this query
search	string, optional	A text search clause to apply to the where clause of this query
devices	string, optional	Optional list of comma separated device names or UUIDs that will be queried. If this parameter is not present, then all devices are queried. The device names must match the list defined under

Parameter	Type, Options	Description
		/archiver/devices or match the /sys/stats/uuid of each device to be queried.
excludeDevices	string, optional	Optional list of UUIDs that will be excluded from being queried. If this parameter is not present, then all devices are queried. The device UUIDs must match the UUID defined under /sys/stats/uuid of each device to be queried.

Manual for values

xforms

Description: Retrieves transforms for the specified key

Security.roles: sdk.meta

Parameters:

Parameter	Type, Options	Description
key	string, optional	An optional language key to retrieve transform for, all transforms returned if not provided

/sys node

Manual for /sys

API Messages

caCert

Description: Display or delete trusted CA certs

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: list|delete|add}	CA action to perform (list,delete,add)
id	string, optional	ID of CA certificate to remove when op=delete

Manual for caCert

count

Description: Returns the number of child nodes

Security.roles: everyone

fileEdit

Description: View and edit system configuration files

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: dir|get|getBackup|put|create|delete}	The operation to perform (dir, get, getBackup, put, create, delete)
type	string, optional, {enum-one:The value must be one of the following: parser|filter}	The type of file, if not provided will check file extension
filename	string, optional, {length:The value length must be between 1 and 64, inclusive}	File to create, edit, get or delete

Manual for fileEdit

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is

whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1

Parameter	Type, Options	Description
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

peerCert

Description: Display or modify trusted peer certificates

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|delete|add}	Operation to perform on the list of trusted peers (list,delete,add)
id	string, optional	ID of peer certificate to remove when op=delete

Manual for peerCert

save

Description: Forces a save of system settings

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
force	uint32, optional	If 1, forces the configuration file to saved, regardless if changes are detected, default is 0

Manual for save

servCert

Description: Display the current server certificate in PEM format.

Security.roles: sys.manage

Manual for servCert

shutdown

Description: Stop the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
reason	string, optional, {length:The value length must be between 0 and 512, inclusive}	The reason for shutting down the service
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart

Manual for shutdown

statHist

Description: Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time for retrieving stats
time2	date-time, optional	The ending time for retrieving stats

Parameter	Type, Options	Description
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format for each stat snapshot, default is posix (seconds since 1970)
include	string, optional	Comma separated list of stats to include (wildcards allowed)
exclude	string, optional	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)
onChange	string, optional	Comma separated list of stats that reduce the result set to only when at least one of the stats change from their previous value
reduce	bool, optional	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result
showAll	bool, optional	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Manual for statHist

telemtry

Description: Returns telemetry information about the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
level	uint32, optional, {signed:The value must be between 1 and 3, inclusive}	The level of information to be sent by the telemetry request. A value of 1 (default) will send all stat nodes, a value of 2 will send stat and config nodes, and a value of 3 will send stat nodes, config nodes and usage metrics for rules, parsers and devices (where applicable).

Parameter	Type, Options	Description
include	string, optional	Comma separated list of nodes to include (wildcard usage is ok).
exclude	string, optional	Comma separated list of nodes to exclude (wildcard usage is ok). If "include" is also specified, exclusions are applied to the set of nodes specified by "include".
url	string, optional	Endpoint URL to send the JSON response.
op	string, optional, {enum-one:The value must be one of the following: get|configure}	The telemetry operation to execute. Default is 'get'.

/users node

API Messages

addOrMod

Description: Add a new user or update an existing user in the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to add or update, must be alphanumeric or @!#$%&'+-

Parameter	Type, Options	Description
		=?^_`{|}~.[]\
password	string, optional	The user's password
pwdIsHashed	bool, optional	Set to true if the password has already been hashed
groups	string, optional	The groups the user belongs to
authType	string, optional, {enum-one:The value must be one of the following: netwitness|pam}	The authentication system to use, "netwitness" is the default
queryTimeout	uint32, optional, {unsigned:The value must be between 0 and 43200, inclusive}	The maximum number of minutes that a query is allowed to execute

Parameter	Type, Options	Description
		for this user. Zero means unlimited.
displayName	string, optional	Display name for this user account
email	string, optional	Email address for this account
description	string, optional	Description of this user account
queryPrefix	string, optional	Query filter applied to every query performed by this user account
sessionThreshold	uint32, optional, {unsigned:The value must be between 0 and 4294967295, inclusive}	Query optimization which will extrapolate the remaining

Parameter	Type, Options	Description
		session counts when they exceed this value

auths

Description: Get supported authentication types

Security.roles: everyone

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete a user from the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to delete

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1

Parameter	Type, Options	Description
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

unlock

Description: Unlock a locked out account

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string	The username or account to unlock

whoAmI

Description: Returns information about the authenticated user

Security.roles: everyone

Appliance service

Introduction

Appliance is a service that co-exist along with other core services and provide options to collect appliance stats like disk, cpu, temperature etc..., also manage core services and configure storage attached to the machine.

API Usage

Appliance service provides API interface through its tree form as nodes, node messages and child nodes.

The following are few examples of the several API nodes and messages supported by Appliance:

- /appliance node messages:
    - `start` and `stop` core service
    - `srvList` to list core services on the appliance
    - `devlist` list storage devices
    - `dfAlloc` create a default configuration for storage used by a core service
    - `partNew` to make partition on block device
    - `/appliance/stats` node for appliance stats information
- /sys node messages:
    - `statHist` to retrieve historical stats from stats database
    - `/sys/stats` node for service system stats information

Sample Request

The following is a sample request to get list of core services on the appliance.

Request path is `/appliance` and message is `msg=srvList`

```
$ curl 'https://appliance-host:50106/appliance?msg=srvList'  \
   -H 'Accept: application/json;charset=UTF-8' \
   -H 'Content-Type: application/x-www-form-urlencoded; charset=ISO-8859-1' \
   -u 'username:password'
```

Sample Response

HTTP Headers

```
HTTP/1.1 200 OK
Content-Length: 657
Connection: Keep-Alive
Pragma: no-cache
Expires: -1
Cache-Control: no-cache, no-store, must-revalidate
Content-Type: application/json
```

JSON Response

It is possible for response values to contain text/plain key-value pairs in content for certain requests.

```
{
 "flags": 1073938433,
 "params": {
  "localhost:56004": "type=decoder
  packet.dir=\"/var/netwitness/decoder/packetdb=137.67 GB\"
  meta.dir=\"/var/netwitness/decoder/metadb=137.67 GB\"
  session.dir=\"/var/netwitness/decoder/sessiondb=137.67 GB\"
  index.dir=\"/var/netwitness/decoder/index=130.43 GB\"
  ... "
 }
}
```

`Note`: The API Messages section for the path nodes provide more details on API messages available and any associated manual pages.

/appliance node

Manual for /appliance

API Messages

addFSMon

Description: Add a filesystem monitor Example arguments: path=/var/lib/netwitness/decoder/packetdb

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
path	string, {file:The file or directory must exist or be creatable on the system}	Filesystem path to monitor

clock

Description: Set the appliance local clock

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
set	date-time, optional	Set a new time. Specify the time with format "YYYY-MM-DD HH:MM:SS".

cntrlcnt

Description: Return the number of controllers

Security.roles: appliance.manage

count

Description: Returns the number of child nodes

Security.roles: everyone

delFSMon

Description: Delete a filesystem monitor Example arguments: path=/var/lib/netwitness/decoder/packetdb

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
path	string, {file-system-monitor:Validates the file system exists}	Monitored filesystem path to remove

devlist

Description: list storage devices

Security.roles: appliance.manage

dpToNic

Description: Remove an Ethernet device from DPDK

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
eth	string, optional, {enum-one:The value must be one of the following: }	Network interface to remove from DPDK

ethtool

Description: Show interface information

Security.roles: appliance.manage

file

Description: List, retrieve and delete files from approved directories on this device

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: ls|get|delete}	The operation to perform (ls, get, delete)
location	string, {enum-one:The value must be one of the following: crashreporter|config|update|feeds|parsers}	The location on the device where the operation will be performed
filename	string, optional	The name of the file to retrieve or delete

fwdLogs

Description: Enable or disable syslog forwarding Example arguments: host=loghost.localdomain

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
host	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789=:/.-}	Remote syslog server

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

nicToDp

Description: Attach a NIC to DPDK

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
nic	string, optional, {enum-one:The value must be one of the following: eth0|lo}	Network interface to send to DPDK

partClr

Description: remove partitions on a block device

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
name	string, {enum-one:The value must be one of the following: sda|sdb|sdc}	block device name
commit	bool, optional	commit changes

Manual for partClr

partNew

Description: make partitions on a block device

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
name	string, {enum-one:The value must be one of the following: }	block device name
service	string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}	service that will use storage
volume	string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}	volume to create
commit	bool, optional	commit changes

Manual for partNew

powerOff

Description: Shut down this appliance No arguments required

Security.roles: appliance.manage

raidList

Description: list drive shelves attached to this appliance

Security.roles: appliance.manage

reboot

Description: Reboot this appliance No arguments required

Security.roles: appliance.manage

reset

Description: Reset data, index, manifests, stats, configuration, or logs for this service. Data automatically deletes index and stats, unless filesCreatedAfter is specified. Service is automatically restarted. Example arguments: data=1 config=1 log=1 This example will reset data, index, logs, and configuration index=1 This example will reset the index only manifest=1 This example will delete everything in the manifest directory filesCreatedAfter="2015-12-01 14:00:00" This example will delete all session, meta and packet files created on or after Dec 1st, 2015 2pm (UTC) from the service. All other files will remain. The index will not be touched, but upon restart will be truncated to match the last session in the session database.

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
data	bool, optional	Reset data, automatically resets index (may not be applicable to all services)
index	bool, optional	Reset index, only valid if data is not reset (may not be applicable to all services)
manifest	bool, optional	Reset all manifests stored in the long term manifest directory (may not be applicable to all services)
config	bool, optional	Reset configuration settings to default
stats	bool, optional	Reset the stats database

Parameter	Type, Options	Description
log	bool, optional	Reset the log database
filesCreatedAfter	date-time, optional	Delete all database files (session, meta, packets) created after a certain date. The index will automatically get rolled back on restart. Not valid with option index.
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart
force	bool, optional	Force reset without confirmation

serial

Description: Retrieve the appliance serial number

Security.roles: everyone

setNTP

Description: Set the clock source for this appliance Example arguments: source=tictoc.localdomain

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
source	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789=:/.-}	The hostname or address of an NTP server that will serve as the hostnam

Parameter	Type, Options	Description
		e for this appliance. Specify "local" to use the built-in clock as the time source.

setSNMP

Description: Enabled or Disable the SNMP service Example arguments: enable=1

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
enable	uint32, {unsigned:The value must be between 0 and 1, inclusive}	Enable the SNMP service. Set to 1 to enable the service, and 0 to disable the service

srvList

Description: list of core services on this appliance

Security.roles: appliance.manage

start

Description: Start a Netwitness service on this appliance Example arguments: service=decoder

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
service	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789=:/.-}	The name of the service to start (decoder, concentrator, etc.)

stop

Description: Stop a Netwitness service on this appliance Example arguments: service=decoder

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
service	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789=:/.-}	The name of the service to stop (decoder, concentrator, etc.)

upgrade

Description: Upgrade the server process

Security.roles: appliance.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: start|check|cancel|finished}	The operation to perform (start, cancel, finished)
size	uint64, optional	The size of the incoming upgrade package
filename	string, optional	The name of the incoming upgrade package
saveFile	bool, optional	If one, the upgrade file will be saved, otherwise it is deleted
crcFile	bool, optional	If true, will perform a CRC of the file and compare to the client's CRC
force	bool, optional	If true, forces the upgrade to take place even if server is up-to-date
rpmforce	bool, optional	If true, forces the rpm to use the force option
shutdown	bool, optional	If true, all RSA NetWitness NextGen services will be shutdown and restarted after upgrade

vgs

Description: list volume groups

Security.roles: appliance.manage

Manual for vgs

vmstat

Description: Get CPU and VM statistics

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
count	uint32, optional, {unsigned:The value must be between 1 and 30, inclusive}	The number of times to gather statistics and print the report. Default is 5

/connections node

API Messages

closeAll

Description: Closes all connections

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: native|rest|amqp}	Indicates the type of connection to close, either native, rest or amqp. Default is all types.
limit	uint32, optional	If non-zero, will stop closing connections after limit is reached. NOTE: It starts closing the most idle connections first.
idleLongerThan	string, optional	A time duration in HH:MM:SS or ##{d,h,m,s} format. Any connection that has been idle longer than the passed in duration is eligible to be closed. Connections that have seen activity within the duration window will not be closed. Examples: "01:05:10" (1 hour, 5 mins, 10 secs) or "65m" (65 minutes)
exceptMe	bool, optional	If true, will not close the connection that initiated the command.

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

netSpeed

Description: Network Speed Test

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 1024 and 536870912, inclusive}	The size of each message sent, default is 64k
count	uint32, optional, {unsigned:The value must be between 1 and 536870912, inclusive}	The number of messages to send, default is 16k

/logs node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

download

Description: Downloads log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting log number to download
id2	uint64, optional	The ending log number to download
time1	date-time, optional	The starting log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
time2	date-time, optional	The ending log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
op	string, {enum-one:The value must be one of the following: start|next|cancel}	The operation to perform (start, next, cancel)
logTypes	string, optional	The type(s) of log messages to retrieve (debug,info,audit,warning,failure)
match	string, optional	Case insensitive string to match in each log message
regex	string, optional	Regular expression to match in each log message
timeFormat	string, optional, {enum-one:The value must be	The time format used in each log message, default is posix time (seconds since 1970)

Parameter	Type, Options	Description
	one of the following: posix|simple}	
batchSize	uint32, optional, {unsigned:The value must be between 2000 and 10000, inclusive}	The number of logs to retrieve per batch, default is 2000

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

pull

Description: Downloads N log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The first log id number to retrieve, this is mutually exclusive with id2
id2	uint64, optional	The last log id number that will be sent, defaults to most recent log message when id1 or id2 is not sent
count	uint32, optional, {unsigned:The value must be between 1	The number of logs to pull

Parameter	Type, Options	Description
	and 10000, inclusive}	
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)

timeRoll

Description: Delete log entries that exceed a given age

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
timeCalc	string, optional, {enum-any:The value must be one or more of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 60, inclusive}	Remove log entries older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 24, inclusive}	Remove log entries older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 365, inclusive}	Remove log entries older than the given number of days
date	string, optional	Remove log entries older than the given date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

/rest node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

/sys node

Manual for /sys

API Messages

caCert

Description: Display or delete trusted CA certs

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: list|delete|add}	CA action to perform (list,delete,add)

Parameter	Type, Options	Description
id	string, optional	ID of CA certificate to remove when op=delete

Manual for caCert

count

Description: Returns the number of child nodes

Security.roles: everyone

fileEdit

Description: View and edit system configuration files

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: dir|get|getBackup|put|create|delete}	The operation to perform (dir, get, getBackup, put, create, delete)
type	string, optional, {enum-one:The value must be one of the following: parser|filter}	The type of file, if not provided will check file extension
filename	string, optional, {length:The value length must be between 1 and 64, inclusive}	File to create, edit, get or delete

Manual for fileEdit

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

peerCert

Description: Display or modify trusted peer certificates

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|delete|add}	Operation to perform on the list of trusted peers (list,delete,add)
id	string, optional	ID of peer certificate to remove when op=delete

Manual for peerCert

save

Description: Forces a save of system settings

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
force	uint32, optional	If 1, forces the configuration file to saved, regardless if changes are detected, default is 0

Manual for save

servCert

Description: Display the current server certificate in PEM format.

Security.roles: sys.manage

Manual for servCert

shutdown

Description: Stop the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
reason	string, optional, {length:The value length must be between 0 and 512, inclusive}	The reason for shutting down the service
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart

Manual for shutdown

statHist

Description: Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time for retrieving stats
time2	date-time, optional	The ending time for retrieving stats
timeFormat	string, optional, {enum-one:The value must be one	Specify the time format for each stat snapshot, default is posix (seconds since 1970)

Parameter	Type, Options	Description
	of the following: posix|simple}	
include	string, optional	Comma separated list of stats to include (wildcards allowed)
exclude	string, optional	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)
onChange	string, optional	Comma separated list of stats that reduce the result set to only when at least one of the stats change from their previous value
reduce	bool, optional	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result
showAll	bool, optional	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Manual for statHist

telemtry

Description: Returns telemetry information about the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
level	uint32, optional, {signed:The value must be between 1 and 3, inclusive}	The level of information to be sent by the telemetry request. A value of 1 (default) will send all stat nodes, a value of 2 will send stat and config nodes, and a value of 3 will send stat nodes, config nodes and usage metrics for rules, parsers and devices (where applicable).

Parameter	Type, Options	Description
include	string, optional	Comma separated list of nodes to include (wildcard usage is ok).
exclude	string, optional	Comma separated list of nodes to exclude (wildcard usage is ok). If "include" is also specified, exclusions are applied to the set of nodes specified by "include".
url	string, optional	Endpoint URL to send the JSON response.
op	string, optional, {enum-one:The value must be one of the following: get|configure}	The telemetry operation to execute. Default is 'get'.

/users node

API Messages

addOrMod

Description: Add a new user or update an existing user in the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to add or update, must be alphanumeric or @!#$%&'+-

Parameter	Type, Options	Description
		=?^_`{|}~.[]\
password	string, optional	The user's password
pwdIsHashed	bool, optional	Set to true if the password has already been hashed
groups	string, optional	The groups the user belongs to
authType	string, optional, {enum-one:The value must be one of the following: netwitness|pam}	The authentication system to use, "netwitness" is the default
queryTimeout	uint32, optional, {unsigned:The value must be between 0 and 43200, inclusive}	The maximum number of minutes that a query is allowed to execute

Parameter	Type, Options	Description
		for this user. Zero means unlimited.
displayName	string, optional	Display name for this user account
email	string, optional	Email address for this account
description	string, optional	Description of this user account
queryPrefix	string, optional	Query filter applied to every query performed by this user account
sessionThreshold	uint32, optional, {unsigned:The value must be between 0 and 4294967295, inclusive}	Query optimization which will extrapolate the remaining session counts when they

Parameter	Type, Options	Description
		exceed this value

auths

Description: Get supported authentication types

Security.roles: everyone

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete a user from the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to delete

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.

Parameter	Type, Options	Description
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

unlock

Description: Unlock a locked out account

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string	The username or account to unlock

whoAmI

Description: Returns information about the authenticated user

Security.roles: everyone

Cloud service

Introduction

The NwCloud core service is part of the XDR Sensor package that is part of our SaaS offering. NwCloud runs alongside NwDecoder and acts as a "go between" between NwDecoder and AWS. In other words, it acts as an aggregator that downloads session/meta information from NwDecoder, converts that information to JSON (for now), and then pushes that data to Kinesis.

Device Registration

In order to connect to AWS, the NwCloud service must be registered using a Device Activation Package (or DAP). The DAP is a file named `device-actvation-package.json`, and can be downloaded from the tenant's portal. The file can then be

either uploaded to the NwCloud service using the activate message, or it can manually be placed in NwCloud's configuration folder.

The activate command can be used through the REST service at the endpoint `/cloud?msg=activate` This requires a program like Postman that will allow you include the DAP's JSON in the body of the request. In this case the Content-Type of the request must be `application/octet-stream`. Once the request is sent, the DAP file will simply be written to NwCloud's configuration folder. Alternatively, users can SCP or copy the DAP into NwCloud's configuration folder manually.

API Usage

Cloud service provides API interface through its tree form as nodes, node messages and child nodes.

- /cloud node messages:
    - `activate` the NwCloud sensor
    - `refresh` the AWS Credentials

`Note`: The API Messages section for the path nodes provide more details on API messages available and any associated manual pages.

/cloud node

API Messages

activate

Description: Activate the NwCloud sensor

Security.roles: cloud.manage

add

Description: Add a new device for aggregation

Security.roles: cloud.manage

Parameters:

Parameter	Type, Options	Description
device	string, optional, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to add (ip:port or [ipv6]:port), either this or ip must be set
ip	string, optional	The IP address of the device, either this or device must be set
port	uint32, optional, {unsigned:The value must be between 0 and 65535, inclusive}	The port number the device is listening on
username	string	The username to run as
password	string, optional	The account password. Assumes service trust aggregation if no password is provided.
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup
reconnect	string, optional	If true, will automatically reconnect when not connected
name	string, optional	Device collection name if device is a collection

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete an existing aggregation device

Security.roles: cloud.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to delete
name	string, optional	Device collection name if device is a collection

edit

Description: Edit an existing aggregation device

Security.roles: cloud.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to edit
username	string, optional	The username to run as
password	string, optional	The account password. For service trust aggregation you must not provide this parameter
options	string, optional	The device options
ssl	string, optional	If true, the device requires SSL to connect
backup	string, optional	If true, the device will be used as a backup
reconnect	string, optional	If true, will automatically reconnect when not connected

Parameter	Type, Options	Description
name	string, optional	Device collection name if device is a collection
switchToTrustedAuth	bool, optional	If true, removes any password associated with this account and will use trusted authentication henceforth. When switching to trusted, only the 'device' parameter is allowed.

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

reconfig

Description: Calculates optimal settings for cloud pools and buffers based on the installed hardware.

Security.roles: cloud.manage

Parameters:

Parameter	Type, Options	Description
update	bool, optional	If true (default is false), will automatically update the respective cloud settings, otherwise it will just output the calculations

Parameter	Type, Options	Description
memory	size, optional	If non-zero, will use that setting as installed physical RAM (e.g., memory="24 GiB"). This is useful for hybrids and/or all-in-one machines where the service will need to share memory.

refresh

Description: Refresh the AWS Credentials

Security.roles: cloud.manage

Parameters:

Parameter	Type, Options	Description
details	bool, optional	If not specified (default) then AWS credentials will be refreshed. Otherwise show the time of the next expiration.

resetMax

Description: Resets all max stats, including device max stats back to zero.

Security.roles: cloud.manage

start

Description: Starts aggregation

Security.roles: cloud.manage

status

Description: Change the online/offline status of an aggregation device

Security.roles: cloud.manage

Parameters:

Parameter	Type, Options	Description
device	string, {ip-address:The value must be an IP address and port in the following format: [ip-address]:[0-65535]}	The device to change status
status	string, optional, {enum-one:The value must be one of the following: online|offline}	Whether to take device offline or put online
name	string, optional	Device collection name if device is a collection

stop

Description: Stops aggregation

Security.roles: cloud.manage

Parameters:

Parameter	Type, Options	Description
flush	bool, optional	If true (default), process all sessions and packets currently in memory. Otherwise, stop immediately, discarding unprocessed sessions and packets.

/connections node

API Messages

closeAll

Description: Closes all connections

Security.roles: connections.manage

Parameters:

Parameter	Type, Options	Description
type	string, optional, {enum-any:The value must be one or more of the following: native|rest|amqp}	Indicates the type of connection to close, either native, rest or amqp. Default is all types.
limit	uint32, optional	If non-zero, will stop closing connections after limit is reached. NOTE: It starts closing the most idle connections first.
idleLongerThan	string, optional	A time duration in HH:MM:SS or ##{d,h,m,s} format. Any connection that has been idle longer than the passed in duration is eligible to be closed. Connections that have seen activity within the duration window will not be closed. Examples: "01:05:10" (1 hour, 5 mins, 10 secs) or "65m" (65 minutes)
exceptMe	bool, optional	If true, will not close the connection that initiated the command.

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

netSpeed

Description: Network Speed Test

Security.roles: connections.manage,aggregate

Parameters:

Parameter	Type, Options	Description
size	size, optional, {byte-range:The value (in bytes, size modifiers like MB are accepted) must be between 1024 and 536870912, inclusive}	The size of each message sent, default is 64k
count	uint32, optional, {unsigned:The value must be between 1 and 536870912, inclusive}	The number of messages to send, default is 16k

/logs node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

download

Description: Downloads log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The starting log number to download
id2	uint64, optional	The ending log number to download

Parameter	Type, Options	Description
time1	date-time, optional	The starting log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
time2	date-time, optional	The ending log date to download, can be posix time (seconds since 1970) or DD-MM-YY HH:MM:SS
op	string, {enum-one:The value must be one of the following: start|next|cancel}	The operation to perform (start, next, cancel)
logTypes	string, optional	The type(s) of log messages to retrieve (debug,info,audit,warning,failure)
match	string, optional	Case insensitive string to match in each log message
regex	string, optional	Regular expression to match in each log message
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)
batchSize	uint32, optional, {unsigned:The value must be between 2000 and 10000, inclusive}	The number of logs to retrieve per batch, default is 2000

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

pull

Description: Downloads N log entries

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
id1	uint64, optional	The first log id number to retrieve, this is mutually exclusive with id2
id2	uint64, optional	The last log id number that will be sent, defaults to most recent log message when id1 or id2 is not sent
count	uint32, optional, {unsigned:The value must be between 1 and 10000, inclusive}	The number of logs to pull
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	The time format used in each log message, default is posix time (seconds since 1970)

timeRoll

Description: Delete log entries that exceed a given age

Security.roles: logs.manage

Parameters:

Parameter	Type, Options	Description
timeCalc	string, optional, {enum-any:The value must be one or more of the following: current|last-write}	The time calculation to use (current date or last write date), default is current
minutes	uint32, optional, {unsigned:The value must be between 0 and 60, inclusive}	Remove log entries older than the given number of minutes
hours	uint32, optional, {unsigned:The value must be between 0 and 24, inclusive}	Remove log entries older than the given number of hours
days	uint32, optional, {unsigned:The value must be between 0 and 365, inclusive}	Remove log entries older than the given number of days
date	string, optional	Remove log entries older than the given date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters

/rest node

API Messages

count

Description: Returns the number of child nodes

Security.roles: everyone

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

/sys node

Manual for /sys

API Messages

caCert

Description: Display or delete trusted CA certs

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: list|delete|add}	CA action to perform (list,delete,add)
id	string, optional	ID of CA certificate to remove when op=delete

Manual for caCert

count

Description: Returns the number of child nodes

Security.roles: everyone

fileEdit

Description: View and edit system configuration files

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, {enum-one:The value must be one of the following: dir|get|getBackup|put|create|delete}	The operation to perform (dir, get, getBackup, put, create, delete)

Parameter	Type, Options	Description
type	string, optional, {enum-one:The value must be one of the following: parser|filter}	The type of file, if not provided will check file extension
filename	string, optional, {length:The value length must be between 1 and 64, inclusive}	File to create, edit, get or delete

Manual for fileEdit

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

peerCert

Description: Display or modify trusted peer certificates

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
op	string, optional, {enum-one:The value must be one of the following: list|delete|add}	Operation to perform on the list of trusted peers (list,delete,add)
id	string, optional	ID of peer certificate to remove when op=delete

Manual for peerCert

save

Description: Forces a save of system settings

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
force	uint32, optional	If 1, forces the configuration file to saved, regardless if changes are detected, default is 0

Manual for save

servCert

Description: Display the current server certificate in PEM format.

Security.roles: sys.manage

Manual for servCert

shutdown

Description: Stop the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
reason	string, optional, {length:The value length must be between 0 and 512, inclusive}	The reason for shutting down the service
cl	string, optional	A (cl="name=value name2=value2") set of options for setting the command line on the next restart

Manual for shutdown

statHist

Description: Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char, * to

match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
time1	date-time, optional	The starting time for retrieving stats
time2	date-time, optional	The ending time for retrieving stats
timeFormat	string, optional, {enum-one:The value must be one of the following: posix|simple}	Specify the time format for each stat snapshot, default is posix (seconds since 1970)
include	string, optional	Comma separated list of stats to include (wildcards allowed)
exclude	string, optional	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)
onChange	string, optional	Comma separated list of stats that reduce the result set to only when at least one of the stats change from their previous value
reduce	bool, optional	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result
showAll	bool, optional	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Manual for statHist

telemtry

Description: Returns telemetry information about the service

Security.roles: sys.manage

Parameters:

Parameter	Type, Options	Description
level	uint32, optional, {signed:The value must be between 1 and 3, inclusive}	The level of information to be sent by the telemetry request. A value of 1 (default) will send all stat nodes, a value of 2 will send stat and config nodes, and a value of 3 will send stat nodes, config nodes and usage metrics for rules, parsers and devices (where applicable).
include	string, optional	Comma separated list of nodes to include (wildcard usage is ok).
exclude	string, optional	Comma separated list of nodes to exclude (wildcard usage is ok). If "include" is also specified, exclusions are applied to the set of nodes specified by "include".
url	string, optional	Endpoint URL to send the JSON response.
op	string, optional, {enum-one:The value must be one of the following: get|configure}	The telemetry operation to execute. Default is 'get'.
arn	string, optional	SNS ARN to send the JSON response.
decoderInclude	string, optional	Nodes to include on Decoder

/users node

API Messages

addOrMod

Description: Add a new user or update an existing user in the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to add or update, must be alphanumeric or @!#$%&'+-=?^_`{|}~.[]\
password	string, optional	The user's password
pwdIsHashed	bool, optional	Set to true if the password has already been hashed
groups	string, optional	The groups the user belongs to
authType	string, optional, {enum-one:The value must be one of the following: netwitness|pam}	The authentication system to use,

Parameter	Type, Options	Description
		"netwitness" is the default
queryTimeout	uint32, optional, {unsigned:The value must be between 0 and 43200, inclusive}	The maximum number of minutes that a query is allowed to execute for this user. Zero means unlimited.
displayName	string, optional	Display name for this user account
email	string, optional	Email address for this account
description	string, optional	Description of this user account
queryPrefix	string, optional	Query filter applied to every query perform

Parameter	Type, Options	Description
		ed by this user account
sessionThreshold	uint32, optional, {unsigned:The value must be between 0 and 4294967295, inclusive}	Query optimization which will extrapolate the remaining session counts when they exceed this value

auths

Description: Get supported authentication types

Security.roles: everyone

count

Description: Returns the number of child nodes

Security.roles: everyone

delete

Description: Delete a user from the system

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789@!#$%&'+-=?^_`{|}~.[]\}	The username to delete

help

Description: Describes this node and its supported messages. NOTE: Command parameters are passed as ="" ="" etc., where must be in double quotes if there is whitespace. To pass a quote in the value, you must escape it by preceding it with a backslash \.

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
msg	string, optional, {chars:The value can only contain the characters: abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ?}	The name of the message to retrieve detailed help about (aliases are 'm' or 'message')
op	string, optional, {enum-one:The value must be one of the following: messages|parameters|description|values|roles|extra|manual}	The specific help operation to perform (e.g., op=manual would return a man page on this node or the specified message)
format	string, optional, {enum-one:The value must be one of the following: default|xml|html}	The format of the response, default returns in a human friendly format

info

Description: Returns detailed information about the node

Security.roles: everyone

ls

Description: Returns the list of child nodes

Security.roles: everyone

Parameters:

Parameter	Type, Options	Description
depth	uint32, optional	How many levels deep to return node info, default is 1
options	string, optional	What types of nodes to return information about, default is all nodes. Can be a number (bitwise mask) or comma separated values like config, stat, folder, session, connection, channel, restart-needed or pretty-print.
exclude	string, optional	Comma separated list of nodes or node pathnames to exclude (wildcards allowed: * and ?)

unlock

Description: Unlock a locked out account

Security.roles: users.manage

Parameters:

Parameter	Type, Options	Description
name	string	The username or account to unlock

whoAmI

Description: Returns information about the authenticated user

Security.roles: everyone

Manuals

/database

The Database tree holds the main functionality related to the session, meta and/or packet/log databases.

The following is a brief overview of the Database commands and what they do:

- **hashInfo** - Retrieves hash information for database files that containing session/meta/packet objects for a set of sessions or date range.
- **dbState** - Returns comprehensive information about the current database state or persists the current state to disk for fast reload.
- **dump** - Provides comprehensive information about a single session in the databases.
- **manifest** - If a manifest directory is defined, it will allow operations on the manifest files (such as a time based query) for database files in cold storage.
- **optimize** - Runs a series of tests to determine the optimal **.write.block.size** based on configured drives and using data from existing databases.
- **reconfig** - Reconfigure database settings based on hardware changes or passed in memory sizes.
- **resetMax** - Resets all max database stats or just the ones listed.
- **sizeRoll** - Delete database files based on the total size of all databases (passed with 'type' parameter) or space remaining on shared volume(s). This command should not be used on databases that don't share storage.
- **stagger** - Staggers database files to optimize read/write performance across multiple volumes. Typically used after adding an empty mount point.
- **timeRoll** - Delete database files that exceed a given age.
- **wipe** - Overwrites all packets and/or meta for a session with a pattern (for eliminating sensitive information).

/database dbState

This command is used internally to persist the state of the databases so they can be reloaded in seconds upon restart. There is an automatically scheduled command to write out the state every so often and it will always be written out on a normal shutdown. There are also options to view key information about a specific database that can be useful when troubleshooting.

Supported parameters:

- **op** - The operation to perform. *save* will persist the database state. *summary* returns a short summary of key information about a specific database (see *type* parameter). *filelist* gives detailed information about a database.
- **type** - The database(s) to perform the operation on. One or more of: *session*, *meta* or *packet*. Separate multiple databases with a comma.
- **options** - Two options that affect returned file sizes: *bytes* or *pretty-print*. *pretty-print* returns sizes in human readable formats like GBs.

/database hashInfo

This command can be used to retrieve the hash information for database files. This is only useful if hashing has been enabled (see configuration for *hash.algorithm*, *hash.databases* and *hash.dir*).

- **sessions** - A comma delimited list of sessions and session ranges to retrieve file hashes for. Use this or *beginDate*, *endDate* but not both.
- **beginDate** - The beginning of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.
- **endDate** - The end of a date range (YYYY-MM-DD HH:MM:SS) of database file creation dates to retrieve file hashes for.
- **directories** - A semi-colon delimited list of additional directories to search for hash files. This may be used to account for changes to the hash.dir configuration.

/database manifest

Manifest files are created with every session, meta, and packet (log) DB file and index slice directory. A manifest file is a file that describes several key pieces of information about the data to which it refers. Manifest files are written as a JSON record. Manifest files travel with the data they represent from tier to tier. If the data they represent is deleted, the manifest file is also deleted, except in the following special case. If the service has `/database/config/manifest.dir` configured to a valid directory, at the point when the manifest data is deleted, a copy of the manifest file is placed into the directory pointed at by `manifest.dir` (the directory is created if it does not exist). This enables a NetWitness Platform feature called historical manifest searching.

The intention of this process is to keep historical manifest files for years, in one location for offline querying. As you might imagine from a service running for many years, this can potentially generate hundreds of thousands of files. This should not be a concern

however, as the service automatically compresses files into a single archive in order to save space when they grow too numerous. Manifest files are very small and compress well.

- **op** - The operation to perform (defaults to query). Other option is 'compress', which will compress all stored manifest files.
- **time1** - The beginning time (UTC) for matching offline database files (YYYY-MM-DD HH:MM:SS)
- **time2** - The ending time (UTC) for matching offline database files (YYYY-MM-DD HH:MM:SS)
- **timeFormat** - Specify the time format that is returned (posix, simple), default is posix

Parameters:

Example: manifest time1="2014-04-20 11:00:00" time2="2014-04-11 11:20:00" timeFormat=simple

More information on manifests can be found in the Core Database Tuning Guide.

/database reconfig

Calculates new drive sizes and free space for the session, meta and/or packet directories. No directories are removed and the assumption is each directory is mounted on a separate filesystem and will only be used for storage of that database.

Parameters:

- **type** - The database types to reconfigure (session, meta or packet), default is all.
- **update** - If true (default is false), will automatically update the respective settings, otherwise it will just output the calculations for viewing.
- **percent** - The drive percentage to use for the calculated size per directory, default is 95.
- **op** - (PACKET DECODER ONLY) The operating mode, 'normal' or '10g'. Default is normal.

If you want to apply the recommended settings, then make sure to pass "update=true".

/database sizeRoll

The sizeRoll command is used to keep two or more databases synchronized wrt data retention. Under normal conditions, each database (session, meta, packet or index) is responsible for managing it's own size and when data should be aged out. However, while this is good for preventing disks from filling up, it's not ideal when one database has too much retention and the oldest data could be orphaned because the other database is aging out faster. This is where the *sizeRoll* command comes into play.

Parameters:

- **type** - The databases (session, meta or packet) to consider for removing the oldest data based on total size or space remaining.
- **log** - If true, will log all actions taken, otherwise all activity will happen silently. Default is true.
- **maxSize** - The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Hot tier.
- **maxSizeWarm** - The maximum size of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxSize. This applies only to the Warm tier.
- **maxPercent** - The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Hot tier.
- **maxPercentWarm** - The maximum percentage of all the volumes of all databases passed in 'type' parameter combined. When exceeded, oldest data is deleted first until total size is less than maxPercent of total volumes. This applies only to the Warm tier.
- **minFree** - The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once again exceeds target. This only applies to the Hot tier.
- **minFreeWarm** - The minimum allowed free space on all volumes for all databases in 'type' parameter before oldest data is deleted. When minimum free space drops below target, the oldest data is deleted first until free space once again exceeds target. This only applies to the Warm tier.

sizeRoll is used to keep the specified databases aging out at roughly the same time period. In order for this to work, the databases being managed by *sizeRoll* must all shared the same storage space(s). Since they share the same storage, there is a tradeoff in I/O performance. Typically, Archiver uses the *sizeRoll* command to manage it's space, as maximum retention is more important that I/O performance.

There are two *sizeRoll* commands, one on the *database* node and one on the *index* node. The only difference is the *index* node command implicitly includes managing the index size. If you use the one on the *database* node, it will not manage the index size.

The *sizeRoll* command is per use. The typical use case is to set it up as a scheduled command so it runs periodically, aging out files based on your criteria. For more information on running scheduled commands, see the Core Database Tuning Guide.

/database stagger

The stagger command is typically only useful for a 10G Decoder and usually just for the packet database. Maximum performance is achieved for storing and retrieving packets when multiple volumes are present. In this scenario, the Decoder always fills the volume with the most free space. When the volumes are roughly the same size, this results in a staggered write pattern, which allows maximum throughput for reading and writing across all volumes. However, this only naturally occurs when multiple packet storage volumes are present at the time the Decoder is first deployed.

A typical use case is adding more storage to an existing Decoder to increase retention. However, when adding storage to an deployment that has already filled the existing volumes with stored packets, the Decoder will naturally fill the new storage with packets before rolling out any packets on the existing storage. This results in a suboptimal read/write pattern because most reads will occur on the same volume that is currently being written to. In a 10G deployment, reads are blocked from the volume when writes are occurring. This doesn't stop ALL reads on that volume, because the file is buffered in memory before being written, but it does result in suboptimal read performance.

With the **stagger** command, you can add more storage and then have the service naturally stagger the files across ALL volumes (existing and new) so that read performance is optimized. This command should only be performed AFTER the storage is mounted and the Decoder configured to use it (e.g., after adding the mount point(s) to *packet.dir*). The downside to this command is it can take some time to stagger and the Decoder should **not** be capturing during the stagger operation.

Recommended workflow:

1. Add all storage and configure mount points
2. Add new storage mount points to packet.dir (or session.dir/meta.dir) and restart service (very important!)
3. Ensure capture is stopped
4. Run stagger operation but make sure the connection that initiated the stagger operation is never terminated until the operation is complete! If the connection is terminated, then the stagger operation will be canceled. If the operation is canceled, the files that were already staggered will remain in

place. The operation can be resumed by rerunning the same command (the work already done will not need to be done again). If running **stagger** from NwConsole, run the "timeout 0" command before sending the stagger command. This will prevent the normal 30 second command timeout.

5. Start capture after stagger command finishes.

The following are the parameters for the command:

- **type** - The database that will be staggered (session, meta or packet). Typically only the packet database is useful for staggering, but it is possible to do the session or meta database when multiple volumes are present for those databases. Since the session and meta databases write far less data than the packet db, typically staggering those databases results in less noticeable performance gains.
- **dryRun** - If true (the default), will only return a description of the operations that would be performed. If false, then the files will actually be moved to an optimal read/write pattern. You MUST pass false to actually stagger the files!

Example usage from NwConsole:

```
login <decoder>:50004 <username> <password>
timeout 0
send /database stagger type=packet dryRun=false
```

If you run this command via the RESTful API, please pass the additional parameter **expiry=0** to prevent a timeout from the service. You will also need to ensure the HTTP client does not disconnect before the operation completes.

/database timeRoll

The *timeRoll* command is used to limit the amount of data based on the age of the data when it was written. For instance, if you have a requirement to only store packet data for 7 days, but no longer than that, then you can use the timeRoll command to "age out" packets that were written more than 7 days ago.

Parameters:

- **type** - The database(s) to age out. Can be one or more of session, meta or packet.

- **timeCalc** - The time calculation to use ("current" date or "last-write" date). Default is current. Unless capture or aggregation is stopped, both settings usually amount to pretty much the same thing.
- **minutes** - Remove database files older than the given number of minutes. Minutes are additive if *hours* or *days* is also supplied.
- **hours** - Remove database files older than the given number of hours. Hours are additive if *minutes* or *days* is also supplied.
- **days** - Remove database files older than the given number of days. Days are additive if *minutes* or *hours* is also supplied.
- **date** - Remove database files older than the given UTC date (YYYY-MM-DD HH:MM:SS), not compatible with minutes, hours, or days parameters.

Example:

```
timeRoll type=session,meta,packet timeCalc=current hours=12
days=3
```

The above command will age out all session, meta and packet files that are older than 3.5 days from the current time. If you stop capture, then it will continue to age out the files until only the most recent file is left. It is not permitted to remove the most recent file being written, so it's not possible to completely remove all data using the *timeRoll* command. To do that, you would need to run the "reset" command on the /decoder (or /<service> node for other services). If you wanted to always keep 3.5 days worth of data, even if capture is stopped, then change the timeCalc parameter to be "last-write".

The *timeRoll* command is per use. The typical use case is to set it up as a scheduled command so it runs periodically, aging out files based on your criteria. For more information on running scheduled commands, see the Core Database Tuning Guide.

/decoder

The Decoder tree holds the main functionality related to capture, parsing and configuration of those processes.

The following is a brief overview of the Decoder commands and what they do:

- **sslKeys** - Provides an interface to upload premaster keys, TLS 1.3 keys or private keys so that the captured encrypted packets that match the keys can be decrypted before the parsing step.
- **start** - Starts capture of the selected device

- **stop** - Stops capture and flushes any remaining data in the pipeline. This can take several minutes.
- **reset** - Used to wipe the device of all data, rebuild the index, or start from a fresh configuration. Use with caution!
- **whoAgg** - Returns information on all the services that are currently aggregating from this Decoder.
- **select** - Used to select the capture device to use for capture.
- **import** - Import raw data directly thru this interface. Cannot be performed while normal capture is running.
- **reconfig** - Used to perform automatic ("easy") configuration of the service based on the input parameters.
- **resetMax** - Returns all statistics to zero.

/decoder reset

Used to wipe the device of all data, rebuild the index, or start from a fresh configuration. Use with caution!

/decoder select

Used to select the capture device to use for capture.

/decoder sslKeys

This command is all about decrypting incoming packets, before the parsing step, so that the enabled parsers will see the unencrypted packet payload and create meta accordingly. Otherwise, most parsers will only see encrypted garbage and will fail to create meaningful metadata. Decrypting packets in real-time requires a non-trivial amount of extra work in the parsing stage. You should plan accordingly by making sure the incoming traffic bandwidth does not overwhelm the available compute power. In other words, you may need more Decoders to decrypt traffic than you would if not decrypting.

Packets captured on a Decoder normally have a timeout of around 60 seconds (in the assembly stage) before they are sent to the parse step. If the Decoder is under memory pressure (very high bandwidth), then the lifetime of the packets in Assembler can be shortened. This timeout can be adjusted through configuration and by increasing the amount of memory available to hold packets in Assembly. Regardless, in order to perform decryption of the packets, the decryption key must be received by the Decoder before the parsing stage.

```
NOTE: Currently with default builds TLS 1.2 and earlier
protocols can be decrypted and you must have the native HTTPS
(not Lua) parser enabled. This parser is port agnostic and
will attempt to decrypt any SSL/TLS traffic when it has the
corresponding decryption key. If FIPS is enabled, this will
restrict the list of ciphers for decryption to only those that
are FIPS approved.

      TLS 1.3 protocol decryption is ONLY supported in nonfips
builds.
```

The following types of encryption keys can be used:

- **premaster** - This is the ephemeral symmetric key actually used in the TLS payload stream for encryption and decryption.
- **private** - The asymmetric private key used during the TLS handshake that encrypts the premaster.
- **tls 1.3 keys** - The TLS 1.3 ephemeral keys which are used in TLS handshake and payload stream encryption and decryption. In the decoder TLS 1.3 decryption is only supported with nonfips builds.

Premaster

The *premaster* key is generated randomly and is ephemeral for the life of one specific TLS session. Normally, there is not an easy way to get *premaster* keys to a Decoder in time for the parsing step. However, both Chrome and Firefox can write the premaster keys they generate to a file. This is useful for testing purposes. To configure your browser to do this, all you have to do is create an environment variable called SSLKEYLOGFILE and assign it the pathname of a text file to write the keys to. Decoder will accept the file exactly as it is written and will use all the decryption keys in the file for any encrypted traffic it captures. The following is a sample NwConsole script that uploads the file to a Decoder:

```
login <decoder>:50004 <username> <password>
send /decoder sslKeys --file-data=SSLKeys.txt
```

or you could use the following curl command (with the RESTful port):

```
curl -u "<username>:<password>" -H "Content-Type:
application/octet-stream" --data-binary @"/path/SSLKeys.txt" -
X POST "http://<hostname>:50104/decoder?msg=sslKeys"
```

Once the symmetric keys are uploaded, they will immediately be used for any necessary decryption. Symmetric keys are stored in memory and there is a limit to how many can be stored at any point in time. As more are added, the earliest keys will be aged out. You can

also add premaster keys by just passing the *random* and *premaster* parameters to **sslKeys**.

TLS 1.3 Keys

TLS 1.3 protocol uses a set of ephemeral symmetric keys also called as key secrets for the life of one specific TLS 1.3 session. The TLS 1.3 session also encrypts handshake, all the TLS records after Server Hello are encrypted , and at each stage a corresponding key is used for the encryption e.g: client handshake, server handshake, client application and server application etc.. . So in order for decrypting a TLS 1.3 session all the required session keys are needed to be available.

```
NOTE: The asymmetric private key decryption is not supported
by TLS 1.3 protocol and only perfect forward secrecy through
ephemeral keys is supported.
```

Similar to *premaster* key, capturing TLS 1.3 keys and uploading to Decoder in time for parsing is not an easy way. However, well known browsers like Chrome , Firefox and IE can write the TLS 1.3 keys they generate to a file and this is useful for testing purposes. To configure your browser to do this, create an environment variable called SSLKEYLOGFILE and set with a pathname to text file to write the keys. Decoder will accept the file exactly as it is written and will use all the decryption keys in the file for any encrypted traffic it captures.

The following are keys used in TLS 1.3 session:

- **CETS** : The CLIENT_EARLY_TRAFFIC_SECRET is used by TLS client to encrypt and send early payload there by notifying server that client would like to use TLS 1.3 and also saves a round trip. Hence it is also termed as 0-RTT.
- **CHTS** : The CLIENT_HANDSHAKE_TRAFFIC_SECRET is used by TLS client to encrypt handshake records.
- **SHTS** : The SERVER_HANDSHAKE_TRAFFIC_SECRET is used by TLS server to encrypt subsequent handshake.
- **CTS0** : The CLIENT_TRAFFIC_SECRET_0 is the first application traffic key used by TLS client to encrypt application payload of the session. Subsequent application keys are recomputed based on previous application keys on Key Update request.
- **STS0** : The SERVER_TRAFFIC_SECRET_0 is the first application traffic key used by TLS server to encrypt application payload of the session. Subsequent application keys are recomputed based on previous application keys on Key Update request.

The following is a sample NwConsole script that uploads the file to a Decoder:

```
login <decoder>:50004 <username> <password>
send /decoder sslKeys --file-data=SSLKeys.txt

Optionally the parameters can be passed in as list in the
upload file,  ex:
send /decoder sslKeys --file-data=SSLKeys.txt --file-
format=params-list
```

or you could use the following curl command (with the RESTful port):

```
curl -u "<username>:<password>" -H "Content-Type:
application/octet-stream" --data-binary @"/path/SSLKeys.txt" -
X POST "http://<hostname>:50104/decoder?msg=sslKeys"
```

Once the symmetric keys are uploaded, they will immediately be used for any necessary decryption. Symmetric keys are stored in memory and there is a limit to how many can be stored at any point in time. As more are added, the earliest keys will be aged out. You can also add available TLS 1.3 keys by just passing
the *random*, *CETS* , *CHTS* , *SHTS* , *CTS0* and *STS0* parameters to **sslKeys**.

```
NOTE: For most of the sessions the CLIENT_EARLY_TRAFFIC_SECRET
key may not be generated or used by TLS clients and it is
normal. However, if handshake and application keys are not
generated and not available then TLS 1.3 decryption wouldn't
be successful.
```

Private Keys or PEM files

Private keys are normally stored in PEM files and are the asymmetric keys generated by services that accept TLS traffic. These keys are used during the TLS handshake to encrypt the premaster symmetric key that will be used for the rest of the payload encryption. For example, if you have a web server whose traffic you want visibility into, then you would upload the private key it uses to encrypt traffic. You only need to do this once, as it is stored permanently (or until a delete command). Once installed, all TLS handshakes that use that private key will be able to be decrypted by the Decoder. After upload, a parser reload command needs to be issued so the newly installed key becomes visible to the HTTPS parser. The following are some sample commands that will upload a PEM file for decryption.

```
NOTE: Not all ciphers suites compute the symmetric key using
the RSA private key (e.g., Ephemeral Diffie Hellman, which is
the E in DHE and ECDHE, has forward secrecy).  Encrypted
traffic with those ciphers cannot be decrypted unless the
premaster key is uploaded to Decoder before the session is
parsed.  For those ciphers, the RSA key is only used to sign
```

```
the ephemeral keys to protect the key exchange against MITM
attacks.
```

For NwConsole:

```
send /decoder sslKeys pemFilename=MyKey.pem --file-
data=/path/MyKey.pem
```

Using the RESTful interface (you must provide the pemFilename parameter in the URL):

```
curl -u "<username>:<password>" -H "Content-Type:
application/octet-stream" --data-binary @"/path/MyKey.pem" -X
POST
"http://<hostname>:50104/decoder?msg=sslKeys&pemFilename=MyKey
.pem"
```

Private keys are automatically encrypted before storing to protect them.

Key Management

The following is the full list of parameters that can be used with the **sslKeys** command:

- **clear** - Removes all premaster keys from memory. Will not delete any PEM files installed on the system.
- **maxKeys** - Changes the maximum number of premaster keys that will be stored in memory.
- **listPems** - Returns a list of all installed private key PEM files.
- **deletePem** - Deletes the named PEM file from the file system. You can pass this parameter more than once to remove multiple files.
- **random** - The random hash used to identify the premaster key or TLS 1.3 keys.
- **premaster** - The premaster key that will be installed for the previous random parameter. They must show up in pairs and random must be first.
- **CETS** - The TLS 1.3 client early traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
- **CHTS** - The TLS 1.3 client handshake traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.
- **SHTS** - The TLS 1.3 server handshake traffic secret key that will be installed for the previous random parameter. It must show up with random parameter.

- **CTS0** - The TLS 1.3 client traffic secret 0 key that will be installed for the previous random parameter. It must show up with random parameter.
- **STS0** - The TLS 1.3 server traffic secret 0 key that will be installed for the previous random parameter. It must show up with random parameter.

Example: The random and TLS 1.3 keys can be passed as space separated parameters to sskKeys, ex:

random =<value> *CETS* =<value> *CHTS* =<value> *SHTS* =<value> *CTS0* =<value> *STS0* =<value>

Return Values

Most commands return name/value pairs of stats about the symmetric keys in memory. The following explains what they mean:

- **added** - The number of premaster and TLS 1.3 keys just added during this command
- **total** - The total number of premaster and TLS 1.3 keys loaded in memory
- **agedOut** - The total number of premaster and TLS 1.3 keys that were removed during this command (this is not a lifetime stat)
- **maxKeys** - The current maximum allowed premaster and TLS 1.3 keys
- **premaster** - The total number of premaster keys loaded in memory.
- **tls1.3** - The total number of TLS 1.3 keys loaded in memory.

Viewing Unencrypted Traffic

Even though the packets are decrypted during the parse stage, only the encrypted packets are written to disk. However, the matching premaster key used for decrypting is written to the *tls.premaster* meta key. This key can then be used to subsequently decrypt the packets on the fly.

Incase of TLS 1.3 session the matching keys are written to the following meta and these keys can then be used to subsequently decrypt the packets on the fly.

- *tls.client.early* : Hex encoded client early traffic key
- *tls.client.hdshk* : Hex encoded client handshake traffic key
- *tls.server.hdshk* : Hex encoded server handshake traffic key
- *tls.client.app* : Hex encoded client application traffic key
- *tls.server.app* : Hex encoded server application traffic key

One Decoder API that can be used to see this is the /sdk/content RESTful service. All you need to know is the Session ID of the encrypted packets and the *flags* parameter masked

to the value 128 (or 0x80 in hex). Point your browser to the Decoder's RESTful interface and type in something like this:

```
http://<decoder>:50104/sdk/content?session=<id>&flags=128&renderer=text
```

You should get back a simple web page showing the packets after they are decrypted. Remove the *flags* parameter to see what it looks like encrypted (or at least remove the 0x80 mask). For more information on the /sdk/content service, see the manual page for /sdk content.

/decoder start

Starts capture on the selected interface.

/decoder stop

Stops capture and flushes all data currently in the pipeline.

/decoder whoAgg

Returns information on all the services that are currently aggregating from this Decoder.

/sdk

The SDK is the primary means to access parsed metadata and the raw data that generated it. There are three main mechanisms for performing queries in the database, the *query*, *values*, and *msearch* commands. Most SDK commands over the RESTful interface override the default expiry of 30 seconds to be unlimited. Why? Because there are already mechanisms in place to cancel long running queries based on configured settings and having a small expiry value only causes confusion.

The following is a brief overview of the commands and what they do:

- **query** - Selects meta from the meta database based on the query that is passed in, possibly using the index for fast retrieval.

- **values** - Returns groups of unique meta values sorted by some criteria. It is optimized to return a subset of the unique values sorted by an aggregate function such as count.
- **msearch** - Takes text search terms as it's input, and returns matching sessions that match the search terms. It can search within indexes, meta, raw packets, or raw logs.
- **packets** - Returns raw packets or log data based on a time range, session ID list or where clause.
- **summary** - Returns name=value pairs of information regarding the active databases of the service.
- **timeline** - Returns session counts for a time period. Usually used for charting sessions over time.
- **deviceId** - Converts a session ID to a equivalent session on a remote service (e.g., which device and session ID was this aggregated from).
- **content** - Convert raw packets data to some other consumable format. Typically used to extract files out of well known protocols like HTTP or SMTP/POP.
- **aliases** - Returns the textual representations of values that are normally integer based (e.g., service 80 is HTTP).
- **language** - Returns the language schema of this service's index.
- **keyrefs** - Returns list of meta entities and renamed keys.
- **validate** - Validate a query or values call without execution.
- **pin** - Pins one or more sessions for long term storage outside of the normal session/meta/packet rollover mechanics.

/sdk content

The content call is used to retrieve a Session stored on disk and return it as one of many different representations. Or it may just return the Session as the raw data, which would include all the metadata and the raw packets or log. This command supports a wide range of representations (hex, packets, web, VOIP, email, meta, text, etc.), but many of those representations are only useful for the Investigator thick client. Nowadays, most third party API access revolves around extracting one or more files out of a specific protocol like HTTP.

Parameters:

- **session** - The session ID that will be retrieved

- **packet** - The packet ID of the first packet of the session. This should never be used by the caller as it's for internal API use only.
- **pinId** - The pinId of the session in long term cache. If provided, session cannot be used.
- **maxSize** - The max number of bytes to return, zero means no limit. This parameter is used to control the maximum bytes that a large network session should return and is mainly meant to prevent an extraordinary large network session from consuming a large number of resources during the transfer. Be careful setting this parameter to zero.
- **failOnCacheMiss** - If true, the content call fails if the NWD is not in the cache. This is used by optimized caching code only and is not normally used by the caller.
- **renderType** - The render type (see below). Not specifying this parameter always results in an NWD and renderFlags and renderOptions are not considered.
- **renderFlags** - Bitwise mask to control option (see below).
- **renderOptions** - An encoded string params containing additional options. For internal use only.

Render Types

C Enum	Value	String Equivalent	Description
NW_CONTENT_TYPE_AUTO	0	apptype	Auto Select HTML Content View
NW_CONTENT_TYPE_DETAILS	1	meta	HTML Meta Details View
NW_CONTENT_TYPE_TEXT	2	text	HTML Text View
NW_CONTENT_TYPE_HEX	3	hex	HTML Hex View
NW_CONTENT_TYPE_PACKETS	4	packets	HTML Packet View

C Enum	Value	String Equivalent	Description
NW_CONTENT_TYPE_MAIL	5	mail	HTML Mail View
NW_CONTENT_TYPE_WEB	6	web	HTML Web Page View
NW_CONTENT_TYPE_VOIP	7	voip	HTML VOIP View
NW_CONTENT_TYPE_IM	8	im	HTML IM View
NW_CONTENT_TYPE_FILES	9	filelist	HTML Listing of all files found in session
NW_CONTENT_TYPE_PCAP	100	pcap	Pcap Packet File
NW_CONTENT_TYPE_RAW	102	raw	Raw Content
NW_CONTENT_TYPE_XML	103		Meta XML File
NW_CONTENT_TYPE_CSV	104		Meta Comma Separated File
NW_CONTENT_TYPE_TXT	105		Meta Tab Separated File
NW_CONTENT_TYPE_NWD	106	nwd	NetWitness Data File
NW_CONTENT_TYPE_FILE_EXTRACTOR	107	files	Extract files from common protocols
NW_CONTENT_TYPE_LOGS	108	logs	Log extract (captured

C Enum	Value	String Equivalent	Description
			logs, LF delimited)
NW_CONTENT_TYPE_PROTOBUF	109		Content of the NWD as a Google Protocol Buffer object

NetWitness Content Render Flags

C Enum	Value	Comment
NW_CONTENT_FLAG_STREAM1	0x00001	Return only request stream content.
NW_CONTENT_FLAG_STREAM2	0x00002	Return only response stream content.
NW_CONTENT_FLAG_SINGLE_COLUMN	0x00004	Format generated web page as a single column with requests and responses interleaved.
NW_CONTENT_FLAG_PACKET_PAYLOAD	0x00008	Include only session payload.
NW_CONTENT_FLAG_DECODEAS_EBCDIC	0x00010	Convert session payload from EBCDIC to ASCII.
NW_CONTENT_FLAG_DO_NOT_EMBED	0x00020	Do not embed application/audio/video traffic into the generated web page.
NW_CONTENT_FLAG_UNCOMPRESS_TEXT	0x00040	Unzip web content in text view.
NW_CONTENT_FLAG_DECODE_SSL	0x00080	Attempt to decrypt SSL sessions if the

C Enum	Value	Comment
		encryption key is provided.
NW_CONTENT_FLAG_STRIP_STYLE_TAGS	0x00100	Removes all html style tags from the original html document.
NW_CONTENT_FLAG_IGNORE_CACHE	0x01000	Ignore any content in cache and requery, affects only request.
NW_CONTENT_FLAG_NO_EMBEDDED_EXE	0x02000	Do not look for or extract hidden/embedded PE files when performing file extraction.
Protobuf only flags		
NW_CONTENT_FLAG_INCLUDE_DUPS	0x04000	Include packets otherwise removed by assembly
NW_CONTENT_FLAG_INCLUDE_HEADERS	0x08000	Include packet header meta information
NW_CONTENT_FLAG_CAPTURE_ORDER	0x10000	Do not assemble packets, return in capture order

RESTful API

The content command is exposed over the RESTful interface by appending /content to the path of the /sdk node (normally just /sdk/content, but would be different for Workbench collections). The RESTful API converts the native content messages that are normally sent over the wire into one of several supported output formats that will be returned over HTTP. The primary use case of /sdk/content involves extracting files from well known protocols. For instance, it is possible to extract all files transfered over HTTP, SMTP, POP3 or SMB, assuming the full session was captured. When requesting files over the RESTful interface, the files will be returned as multipart/mixed.

Example usage for extracting files from a session with a supported protocol:

`/sdk/content?session=12345&maxSize=0&render=files&expiry=0`
The *expiry* parameter is the server side RESTful timeout (in seconds). Setting it to zero means do not timeout.

/sdk hierarch

The `hierarch` call returns information about the hierarchy of devices attached to the collection represented in this database.

A hierarchy consists of this device, plus any devices that this device is connected to. For each device, the contents of the `/sys/stats` folder is returned. This information includes the device name, it's UUID, and it's version information.

The hierarchy command returns it's information as a MessagePack object, which may be translated into different representations depending on what API you are using to access the Core service. For example, using the REST API it is translated to a JSON object.

For devices that connect to upstream devices, such as a Broker or Concentrator, the hierarchy message will contain a `devices` member. The `devices` member is an array that holds the contents of the `hierarch` message as executed on each upstream device. In this way, the `hierarch` message forms a hierarchical directory of all services that the device connects to, both directly and indirectly.

/sdk keyrefs

Returns list of meta entities and renamed keys. This messages does not accept any parameters.

1. Meta Entities. These are defined as entity elements in the index configuration.
2. Renamed Keys. These are defined as keys that contain a rename element.

The returned list has two types of entries:

Renamed keys can be distinguished from entities in the result set because renamed keys also include the key itself in the list of referenced keys.

/sdk language

Returns the current meta database schema as well as the level of indexing for each key as defined by the current configuration on the service. Each result returned describes a single meta key.

Parameters:

- **id1** - The starting language id. This is optional and is usually left off or set to zero. The only time this parameter should be used is if the full language is not retrieved in a single call and the next batch must be retrieved. In this case, you would set *id1* to the next valid ID to be retrieved based on the last valid ID retrieved from the previous call.
- **id2** - The ending language id. Normally not used, but see remarks for *id1*.
- **time1** - An optional starting date and time for retrieving language counts. Only valid when *flags* is set to *quick-count* or *full-count*.
- **time2** - An optional ending date/time for retrieving language counts.
- **options** - Extra optional arguments, currently unused.
- **flags** - Optional flags to configure how the results are returned. Can be a number (bitwise mask as defined in the C API) or comma separated values like *default*, *no-count*, *quick-count* or *full-count*.
- **size** - The max number of entries to return. In most cases, you will want to set this number high (e.g., 1000) to retrieve the full language with one call.

C API

Each NwField result will describe a single language element as follows:

- NwField::type will contain the name (accessible through the @p name member)
- NwField::variant will contain an nwText value with a description
- NwField::flags will specify the index level and other information as defined by @ref NwLanguageResultFlags. Please use the provided masks to extract the required information from the uint32.

```
After calling this function, you can optionally call
NwResultsInfo to check for additional information about the
query results.
```

Parameters:

- **collectionHandle** - Handle of the collection.
- **lang1** - The first language id, usually pass zero.
- **lang2** - The second language id, usually pass zero.

- **time1** - The starting time in seconds since 1970.
- **time2** - The ending time in seconds since 1970.
- **reserved** - Reserved for future use, pass NULL.
- **flags** - Additional NwLanguageFlags to control the data returned.
- **result** - NwField array, large enough to hold size results.
- **resultSize** - The size of the results array, returns size of results.

Language Flags

C Enum Flag	Value	Comment
NW_LANGUAGE_FLAG_DEFAULT	0	No value counting is done. Time1 and time2 parameters are ignored.
NW_LANGUAGE_FLAG_NO_COUNT	1	Same as default.
NW_LANGUAGE_FLAG_QUICK_COUNT	2	Returns the count of values from the most recent non-empty time slice only. Slices of index between time1 and time2 are considered.
NW_LANGUAGE_FLAG_FULL_COUNT	4	Scan the time slices between time1 and time2 and return the unique value counts between those times in the results.

NetWitness Language Result Flags

```
These flags and masks are used to decipher the results
returned from the NwLanguage
function.  Specifically, the nwuint32 flags member in NwField.
```

C Enum Flag	Value	Comment
NW_LANGUAGE_KEY_INDEX_MASK	0x000F	Mask for the index level.
NW_LANGUAGE_KEY_INDEX_FILTER	0	Indicates the key is filtered, so data related to the key should be ignored.
NW_LANGUAGE_KEY_INDEX_NONE	1	Indicates the data is kept but not indexed
NW_LANGUAGE_KEY_INDEX_KEYS	2	Indicates the data is indexed only at the key level.

C Enum Flag	Value	Comment
NW_LANGUAGE_KEY_INDEX_VALUES	3	Indicates the data is indexed at key and value levels.
NW_LANGUAGE_KEY_SPECIAL_MASK	0x00F0	Special flags about each language key
NW_LANGUAGE_KEY_SPECIAL_SINGLETON	0x0010	Token appears only once per session
NW_LANGUAGE_KEY_SPECIAL_VALUE_ALIASES	0x0040	Indicates this key has value aliases
NW_LANGUAGE_KEY_SPECIAL_TRANSFORM	0x0080	Indicates this key has a transform
NW_LANGUAGE_KEY_ACTION_MASK	0x0F00	Mask indicating default behavior in Investigator
NW_LANGUAGE_KEY_ACTION_HIDDEN	0x0100	Indicates the default is to hide the key, user configurable from index-*.xml file.
NW_LANGUAGE_KEY_ACTION_OPEN	0x0200	Indicates the default is to open the key, user configurable from index-*.xml file.
NW_LANGUAGE_KEY_ACTION_CLOSE	0x0300	Indicates the default is to close the key, user configurable from index-*.xml file.
NW_LANGUAGE_KEY_ACTION_AUTO	0x0400	Indicates the default is to auto open the key, user configurable from index-*.xml file.
NW_LANGUAGE_KEY_PROTECTED_VALUE	0x10000	Values of this type should be treated as protected.
NW_LANGUAGE_KEY_PROTECTED_TOKEN	0x20000	Indicates that the key is intended to contain tokenized values.

C Enum Flag	Value	Comment
NW_LANGUAGE_KEY_TRANSIENT	0x800000	Indicates that values for this key are not persisted and only exist during parsing on a decoder.
NW_LANGUAGE_KEY_RESULT	0x80000000	Indicates this result is a response to a language call

Returns:

- 1 on success, 0 on failure

/sdk msearch

The index provides a low-level `msearch` function to perform text searches against all meta types. This type of search does not require users to define their queries in terms of known meta types. Instead, it searches all parts of the database for matches. `msearch` is used by the Events view text search. See the "Filter and Search Results in the Events View" topic in the *Investigation and Malware Analysis Guide* for detail on the accepted search forms and examples.

`msearch` parameters:

```
msearch-params  = search-param, {space, where-param}, {space,
limit-param}, {space, size-param}, {space, flags-param};
search-param    = "search=", ? free-form search string ? ;
where-param     = "where=", ? optional where clause ? ;
limit-param     = "limit=", ? optional session scan limit ? ;
size-param     = "size=", ? optional result count limit ? ;
flags           = "flags=", {msearch-flag, {"," msearch-flag}
};
msearch-flag    = "sp" | "sm" | "si" | "ci" | "regex" ;
```

The `msearch` algorithm works as follows:

1. A set of sessions is identified from the index by finding the intersection of three sets:
    - (Set 1) All sessions in the database
    - (Set 2) Sessions that match the `where` clause parameter
    - (Set 3) If the `si` flag is specified, sessions that indexed values that match the search string parameter.

2. If the search specifies the `sm` parameter, all meta items from the set of sessions identified in step 1 are read and scanned to see if they match the search string parameter. The meta items will be read from the service nearest to the point where the search was executed. For example, if the search is performed on a Broker, the meta items may be read from the Concentrator nearest to the broker, but if the search is performed on an Archiver the meta items will be read from the Archiver itself.
3. If the search specifies the `sp` parameter, all raw packet or log entries from the set of sessions identified in step 1 are read and scanned to see if they match the search string parameter. The packets will be read from the service nearest to the point where the search was executed. For example, if the search is performed on a Concentrator, the packet data will be read from the Decoder, but if the search is performed on an Archiver, the packet data will be read from the Archiver itself.
4. Matches from step 2 and step 3 are returned as they are found, up to the point where the `limit` parameter is reached or the `size` count is reached, whichever occurs first. The `limit` parameter specifies the maximum number of sessions for which meta and packet data will be scanned. If `limit` is not specified, the entire set of sessions determined in step 1 is scanned. The `size` parameter specifies the maximum number of results that will be returned. In practice, the `size` parameter acts more as a suggestion. It is possible that slightly more results than specified will be returned, but fewer results will never be returned. If the `size` parameter is not specified, all results matching the search will be returned.

`msearch` Flags

Flag	Description
`sp`	Scans raw packet data
`sm`	Scans all meta data
`si`	Does index lookups for all search parameters before scanning meta
`ci`	Performs a case insensitive search. Returned results are case-preserving.
`regex`	Treats the search parameter as a regular expression. Only a single regular expression can be specified, but the regular expression may be arbitrarily complex.

`msearch` Index Search Mode

Using the index search mode, specified by using the `si` flag, causes results to be returned significantly faster than any other mode. The main limitation of this mode is that it only returns matches on text terms that match value-indexed meta values.

- The `si` parameter must be combined with the `sm` flag. The `si` parameter implies the search only matches indexed meta.
- The `si` parameter can be used with regex searches, however only text indexed values will match. IP addresses and numbers will not match the regex.

Text Search Syntax

The search parameter given to `msearch` is composed of 1 or more words, separated by whitespace. For example, searching for `foo` returns sessions that contain the word foo.

If multiple terms are provided for the search, they are implicitly considered to be an AND operation. For example, searching for `foo bar` returns sessions that contain both foo AND bar. Sessions that contain only foo or only bar are filtered out. If you want to search for sessions containing any of two or more terms, you must explicitly separate the terms with the word OR. For example, searching for `foo OR bar` returns sessions that contain either foo or bar.

Search Syntax And Index Modes

The searches given to the `msearch` command are interpreted according to the index level on all the indexes. `msearch` works on the value-indexed keys in the index. Search terms provided to `msearch` will find values that are an exact match to values that were indexed.

As of version 11.1, there are new index modes available that allow `msearch` to locate text that is not an exact match to the search input. `msearch` supports wildcard searches on the word meta index, if the word meta index has the ngram option enabled. For details on the ngram option, see the topic Index Customization.

The wildcard search allows the use of the `*` and `?` characters as wildcards in search terms. The `*` can stand for 0 or more characters, while the `?` may stand for any single character. To search for those characters in an N-gram enabled index, you may escape them with a backslash character.

If the word index has the 'edge' N-gram option enabled, then it can be used to locate searches that end in a wildcard. This means it is only useful for finding text that begins with a known prefix.

If the word index has the 'all' N-gram option enabled, then wildcards may appear anywhere in the search term.

This table summarizes the relationship between word index level, and the types of searches that `msearch` will locate.

Search input	Non-indexed	IndexValues	IndexValues with Edge N-grams	IndexValues with All N-grams
"foo"	no match	"foo"	words starting with "foo"	words containing "foo"
"foo*"	no match	literal "foo*"	words starting with "foo"	words starting with "foo"
"*foo"	no match	literal "*foo"	no match	words ending with "foo"
"*foo*"	no match	literal "*foo*"	words staring with foo	words containing "foo"
"foo*"	no match	literal "foo*"	literal "foo*"	literal "foo*"

`msearch` Tips

- Always use the `where` clause to specify a time range for the search.
- To search for IP address ranges, specify them in the `where` clause.
- Use the `limit` parameter when not using the index search mode. Without it, there will be an extremely large amount of data read by the meta and packet databases.

/sdk packets

The packets call returns raw packets or logs based on the user's selection criteria. This command corresponds to the C SDK command NwPackets() and the RESTful URL at /sdk/packets. Both APIs operate by calling the "/sdk packets" command on the server to do all the heavy lifting.

NwConsole also supports saving pcap or log files via the internal packets command. For more information, type "help packets" while running NwConsole. Detailed help on the command can be found by looking at the manual page ("man packets").

RESTful API

The packets command is exposed over the RESTful interface by appending /packets to the path of the /sdk node (normally just /sdk/packets, but would be different for Workbench collections). The RESTful API converts the native packet messages that are normally sent over the wire into one of several supported output formats that will be returned over HTTP.

The following are the arguments used to select the session's packets/logs to be returned:

- **sessions** - A CSV list of Session IDs or Session ID ranges. Example: sessions=1,3,5-10,12-15,17
- **where** - An optional where clause that will be used to match which session's packets or logs will be returned
- **time1** - The starting time (either POSIX time or "YYYY-MM-DD HH:MM:SS" format). Cannot be used with the parameters *sessions* and *where*.
- **time2** - The ending time (either POSIX time or "YYYY-MM-DD HH:MM:SS" format). Cannot be used with the parameters *sessions* and *where*.

The *time1* and *time2* parameters are extremely fast but they do not filter any raw data within the specified time range.

The following are the supported output formats, which are selected by passing the *render* parameter with one of the following values:

- **pcap** - Returns the packets in a pcap file (log sessions are ignored). Content-Type: application/octet-stream
- **logs** - Returns raw logs in a plain text file (network sessions are ignored). Content-Type: text/plain
- **text/csv** - Return a line delimited CSV file containing only logs (network sessions are ignored). Each entry/line consists of 5 fields: timestamp, source, forwarder, lccid, log. The last field (log) is the actual log text and is not escaped, any comma's that occur after the 4th delimiter are part of the log text. Content-Type: text/plain
- **text/xml** - Return an XML file containing only logs (network sessions are ignored). Each log entry will include timestamp and optionally source, forwarder and lccid if known. The log text will be contained in a CDATA, and split across consecutive CDATA if the logs contains a CDATA terminator (]]>). Content-Type: text/xml
- **application/json** - Return a JSON file containing only logs (network sessions are ignored). Each log entry will include timestamp, source, forwarder, lccid and log text. Content-Type: application/json

Example:

```
Grab 5 minutes worth of packets: time1="2015-09-12 15:00:00"
time2="2015-09-12 15:05:00"
/sdk/packets?expiry=0&time1=2015-09-
15%2016%3A00%3A00&time2=2015-09-15%2016%3A05%3A00
```

The *expiry* parameter is the server side RESTful timeout (in seconds). Setting it to zero means do not timeout.

C API

```
nwuint32 NwPackets(
      nwuint32 collectionHandle,
      const char* parameters,
      nwuint32 flags,
      NwPacketsHandler packetsHandler,
      void* userData);
```

Streams packets from a local or remote collection based on a list of session IDs or a time range. The packets are retrieved from the passed in NwPacketsHandler (which cannot be NULL). The handler is called once for every packet. userData can be anything, usually a pointer to a structure, and is passed verbatim to the packets handler.

Examples of the parameters string:

```
Session Range w/Where Clause
      sessions=1,5,8,11-15,21-30,45-23221 where="service=80 &&
time='2015-03-19 00:00:00'-'2015-03-20 00:00:00'"

Time Range
      time1="2010-2-13 13:00:00" time2="2010-2-13 13:05:00"
```

This call can be canceled using NwCancel() from another thread or from the packets handler.

NOTE: The time range is only guaranteed to work correctly if the collection has not imported any packets out of order.

Parameters:

- **collectionHandle** - Handle of the collection.
- **parameters** - A string containing either a comma separated list of session IDs and session ranges, a where clause or a beginning and ending time range. The time range (*time1*/*time2*) is mutually exclusive with *sessions* and/or *where*.

- **flags** - Additional flag specifiers. Requesting session ids is only possible when passing the "sessions" parameter.
- **packetsHandler** - The callback function, called once for each packet returned.
- **userData** - User defined pointer returned in the progress handler.

Returns 1 on success, 0 on failure

/sdk pin

The idea behind the `pin` command is to store sessions (with meta and packets) for long term retention, outside of the normal rollover mechanics. For instance, if several sessions indicate there might be a attack on a customer's network, those sessions could be pinned (perhaps by NetWitness Respond) so they can be viewed at a later time without worry that the meta or packets might be rolled out before the threat is accessed.

There are two parameters that control how a session is pinned:

- **session** - The session ID that should be pinned
- **sessions** - One or more comma separated session IDs or session ranges that should be pinned. You can use either parameter interchangeably, but they cannot both be submitted at the same time. Example: "sessions=1,5,8-11", this will pin session IDs 1, 5, 8, 9, 10 and 11.

The response from pinning one or more sessions is called the Pin ID. For every session pinned, you will receive back a response that contains the original session ID, the Pin ID and the total size (bytes) of the pinned session in long term storage. When pinning multiple sessions, you will get back multiple responses, each response will be for a single session.

The Pin ID should be stored for later retrieval. For instance, Respond will store the Pin IDs for each incident in which the customer requested the sessions to be pinned. Later on, you can view those pinned sessions via a content call. See the /sdk content call for more details. Pinned sessions can only be viewed via a content call. They will not show up in a query or other SDK command once the sessions have rolled out.

Miscellaneous Operations

There is one additional parameter `op` than can be used for some additional operations around pinned sessions. The following are the valid values for the `op` parameter:

- **ls** - Returns a listing of all pinned sessions for all online services. The returned listing will contain the Pin ID, the creation date of the pinned session and the storage size of the pinned session (in bytes).
- **validate** - Validates that one or more Pin IDs are still valid and in Pin storage. The response to each Pin ID (see the `pinId` parameter below) passed in will either indicate that the Pin ID is valid, not-found or if the no Core service recognizes the Pin ID, then no response will be given for those Pin IDs. If no response is received for a specific Pin ID, you can only assume that the service that stores those Pin IDs is offline, is no longer operational or that the original service for the Pin ID is no longer configured to store pinned sessions.
- **unpin** - This parameter is used to delete previously pinned sessions. Like `validate` above, it requires one or more Pin IDs to be passed in. The responses received are also nearly identical to `validate`, with the exception that instead of indicating the Pin ID is valid, it indicates the Pin ID was deleted.

pinId

This parameter accepts one or more comma separated Pin IDs. It must be provided when `op=validate` or `op=unpin` is used.

/sdk query

The `query` message has the following syntax:

```
query-params     = size-param, space, query-param, {space,
start-meta-param}, {space, end-meta-param}, {space, search-
param}, {space, streamlimit-param};
size-param       = "size=", ? integer between 0 and 1,677,721
? ;
query-param      = "query=", query-string ;
start-meta-param = "id1=", metaid ;
end-meta-param   = "id2=", metaid ;
search-param     = "search=", search-string ;
streamlimit-param = "streamLimit=", streamSize ;
metaid           = ? any meta ID from the meta database ? ;
```

The `id1`, `id2`, and `size` parameters form a paging mechanism for returning a large number of results from the database. Their usage mostly benefits developers who are writing applications directly against the NetWitness Core database. Normally, results are

returned in the order of oldest to newest data (higher meta IDs are always more recent). In order to return results from most recent to oldest, reverse the IDs such that `id1` is larger than `id2`. This has a slight performance penalty, because the where clause must be completely evaluated before processing in reverse order can begin.

When size is left off or set to zero, the system streams back all results without paging. For the RESTful interface, this results in the full response to be returned with chunked-encoding. The native protocol returns the results over multiple messages.

The `streamLimit` parameter to the query specifies how many groups of results will be returned when the query is operating in streaming mode. For example, the query may be in streaming mode by virtue of not having a 'size' parameter, but the streamLimit can be used to limit how many results are actually returned. Specifying the `streamLimit` implies that streaming mode is requested. If the client requests a 'streamLimit' but the query cannot be processed in streaming mode due, then an error is generated.

The `query` parameter is a `query` command string with its own NetWitness-specific syntax:

```
query-string      = select-clause {, where-clause} {, group-
by-clause {, order-by-clause } } ;
select-clause     = "select ",  ( "*" | meta-or-aggregate {,
meta-or-aggregate} ) ;
where-clause      = " where ",  { where-criteria } ;
meta-or-entity    = (meta_key | entity) ;
meta-or-aggregate = meta-or-entity | aggregate_func, "(",
meta-or-entity, ")" ;
aggregate_func    = "sum" | "count" | "min" | "max" | "avg" |
"distinct" | "first" | "last" | "len" | "avglen" |
"countdistinct" ;
group-by-clause   = " group by ", meta-key-list
meta-key-list     = meta-or-entity {, meta-key-list}
order-by-clause   = " order by ", order-by-column
order-by-column   = meta-or-aggregate { "asc" | "desc" } {,
order-by-column}
```

The `select` clause allows you to specify either `*` to return all the meta in all the sessions that match the where clause, or a set of meta field names and aggregate functions to select a subset of the meta with each session.

The `select` clause may contain entity names in the place of meta key names. If an entity name is in the select clause, meta items returned by the query will have their key name set to the entity name, rather than their actual meta key name stored in the session. Thus, the names of the meta items returned in the query will match the names of the metas in the select clause. For example, if there is an entity `ip` that consists of `ip.dst` and `ip.src`, then a query containing `select ip` will only return `ip` fields, with nothing to distiguish `ip.dst` meta items from `ip.src` meta items in the result set.

The `select` clause may contain renamed meta key names. Any fields appearing in the result set as a result of a renamed key in the `select` clause will be returned with the meta key name matching the name used in the `select` clause. For example, if the key `port_src` is used to rename `tcp.srcport`, then a query containing `select port_src` will only return `port_src` fields, even if the underlying meta had type `tcp.srcport`.

Note: Usage of renamed meta key pairs in the `select` clause cannot be combined with fixed-size result paging for a query. Doing so causes discrepancies in the results returned to Brokers. The reason for the discrepancies is that Concentrators cannot return only one of the key values of a renamed meta key pair and still preserve the correctness of the result set for the requested size. Hence, the Concentrator omits renamed meta key pair results to preserve the correctness of the result set, which causes the Broker to pull the result from the next Concentrator and advance the IDs that are returned.

Example: `select ip.proto,ipv6.proto` cannot be combined with `size=10` (a paging query) `size=10 flags=0 threshold=0 query="select time,ip.src,ip.dst, ip.proto,ipv6.proto,eth.type,size,payload,lifetime,client,did`

The aggregate functions have the following effect on the query result set.

Function	Result
`sum`	Add all meta values together; only works on numbers
`count`	The total number of meta fields that would have been returned
`min`	The minimum value seen
`max`	The maximum value seen
`avg`	The average value for the number
`distinct`	Returns a list of all unique values seen
`countdistinct`	Returns the number of unique values seen. `countdistinct` is equivalent to the number of metas that would have been returned by the distinct function.
`first`	Returns the first value seen
`last`	Returns the last value seen
`len`	Converts all field values to a UInt32 length instead of returning the actual value. This length is the number of bytes to store the actual value, not the length of the structure stored in the meta database.

Function	Result
	For example, the word "NetWitness" returns a length of 10. All IPv4 fields, like `ip.src`, return 4 bytes.
`avglen`	Returns a single value which is the average value returned from the `len` function. The result is always a float64 value.

`where` Clauses

The `where` clause is a filter specification that allows you to select sessions out of the collection by using the index.

Syntax:

```
where-criteria    = criteria-or-group, { space, logical-op,
space, criteria-or-group } ;
criteria-or-group = criteria | group ;
criteria          = (meta-key | entity), ( unary-op | binary-
op meta-value-ranges ) ;
group             = ["~"], "(" where-clause ")" ;
logical-op        = "&&" | "||" ;
unary-op          = "exists" | "!exists" ;
binary-op         = "=" | "!=" | "<" | ">" | ">=" | "<=" |
"begins" | "contains" | "ends" | "regex" ;
meta-value-ranges = meta-value-range, { ",", meta-value-range
} ;
meta-value-range  = (meta-value | "l" ), [ "-", ( meta-value |
"u" ) ] ;
meta-value        = number | quoted-value | ip-address | mac-
address | relative-time ;
number            = ? any numeric value ? | ( '"' text '"' )
quoted-value      = ( '"' text '"' ) | ( '"' date-time '"' ) ;
relative-time     = "rtp(" , time-boundary , "," ,  positive-
integer , time-unit, ")" ;
time-boundary     = "earliest" | "latest" | "now" ;
positive-integer  = ? any non-negative integral number ?
time-unit         = "s" | "m" | "h" ;
```

When specifying rule criteria, the `meta-value` part of the clause is expected to match the type of the meta specified by the `meta-key`. For example, if the key is `ip.src` the `meta-value` should be an IPv4 address. Entity names are allowed in any location where a meta-key name is required.

Queries using a `meta-key` name will match meta items corresponding both to the `meta-key` name as well as to the names of any "renames" specified for the key. See "Key Renaming" under the Index Customization topic for details on key renaming.

Query Operators

The following table describes the function of each operator.

Operator	Function
`=`	Match sessions containing the meta value exactly. If a range of values is specified, any of the values is considered a match.
`!=`	Matches all sessions that would not match the same clause as if it were written with the `=` operator.
`<`	For numeric values, matches sessions containing meta with the numeric value less than the right side. If the right side is a range, the first value in the range is considered. If multiple ranges are specified, the behavior is undefined. For text metas, a lexicographical comparison is performed.
`<=`	Same behavior as `<`, but sessions containing meta that equals the value exactly are also considered matches.
`>`	Similar to the `<` operator, but matches sessions where the numeric value is greater than the right side. If the right side is a range, the last value in the range is considered for the comparison.
`>=`	Same behavior as `>`, but sessions containing meta that equals the value exactly are also considered matches.
`begins`	Matches sessions that contain text meta value that starts with the same characters as the right side.
`ends`	Matches sessions that contain text meta that ends with the same characters as the right side.
`contains`	Matches sessions that contain text meta that contains the substring given on the right side.
`regex`	Matches sessions that contain text meta that matches the regex given on the right side. The regex parsing is handled by boost::regex.
`exists`	Matches sessions that contain any meta value with the given meta key.
`!exists`	Matches sessions that do not contain any meta value with the given meta key.
`length`	Matches sessions that contain text meta values of a certain length. The expression on the right side must be a non-negative number.

Text Values

The system expects quoted text values. Unless it can be parsed as a time (see below), a quoted value is interpreted as text.

It is also important to quote any text value that may contain `-` so that it is not interpreted as a range.

For text values, the backslash character `\` is used as an escape value. This character is used when you need to search for a value containing quote characters. If you need to search for a backslash character, then the backslash itself must be escaped, as `\\` . Note that if you are wrapping the query parameters within another language, such as the parameter fields of the REST interface, you may need to add additional escape levels as required by whatever API or interface you are using to interact with the core service.

IP Addresses

IP addresses can be expressed using standard text representations for IPv4 and IPv6 addresses. In addition, the query can use CIDR notation to express a range of addresses. If CIDR notation is used, it is expanded to the equivalent value range.

MAC Addresses

A MAC address can be specified using standard MAC address notation: `aa:bb:cc:dd:ee:ff`

Numeric Values

In a where clause, you can specify numeric search values. Numbers should not be surrounded by quotes.

Bucketed Numeric Indexes

Meta keys indexed with bucketing can be used like any other numeric search value. Under most situations such searches will return sessions that have a meta value that exactly matches the requested search criteria.

Special behavior is invoked for queries that select only `sessionid`, for example a query of the form `select sessionid where size = 2048`.
Selecting `sessionid` explicitly bypasses all meta database read operations, and only returns index information. If selecting `sessionid` only, and if the numeric value specified is exactly equal to one of the bucket values, then the system will return all sessions that match somewhere in the bucket, rather than an exact match. For example, the search term `size = 2048` will match all sessions in the 2 KB bucket, which is the range from 2048 to 3171 bytes. However, if the search values does not match a bucket

values, then the system will return only matches for the exact byte value. For example, the search term `size = 2049` will only match sessions with a size meta value exactly 2049. In this mode of operation, specifying a non-bucket value in a where clause is slower than searching within a bucket value. The 'where' clause parameter to the `values` API also invokes this optimization.

Using bucketed values in other forms of `query` does not invoke special behavior. The same is true for the `msearch` API. For those APIs, the use of a bucketed index in the where clause is evaluated accurately, without special meaning applied to bucket values. To search within an entire bucket using these APIs, specify the bucket range explicitly. For example `size=2048-3171`.

More information on how to tell if an index is bucketing is in the topic Index Customization.

Numeric Value Aliases

For numeric values, aliases specified in the index can be used in a query as a quoted string in place of where a literal numeric value would be used; e.g.,

```
select * where service = "NFS"
```

Numeric value aliases can be used anywhere a numeric literal might be used: as a single value, as the beginning or end of a range, or in a comma-delimited list of values (and/or ranges).

Refer to the topic Index Customization for details of how value aliases can be specified in the index.

Date and Time Expressions

In NetWitness Platform, dates are represented using Unix epoch time, which is the number of seconds since Jan 1, 1970 UTC. In queries, you can express the time as this number of seconds, or you can use the string representation. The string representation for the date and time is `"YYYY-mmm-DD HH:MM:SS"`. A three-letter abbreviation represents the month. You can also express the Month as a two-digit number, 01-12.

Time values must be quoted.

All times specified in queries are expected to be in UTC.

Relative Time Points

Relative time points allow a where clause to reference a value at some fixed offset relative to the earliest or latest time metas seen in the collection. It can also be used to reference a point in time relative the the current time.

A relative time point expression has the syntax `rtp(boundary, duration)`.

The boundary is either `earliest`, `latest`, or `now`.

The duration is an expression of hours, minutes, or seconds. For example, `24h`, `60m`, or `60s`. When the boundary is `earliest`, the duration represents the amount of time **after** the earliest time present in the collection. If the boundary is `latest`, the duration represents the amount of time **before** the earliest time present in the collection. If the boundary is `now`, the duration represents the amount **before** the current time.

When the boundary is `now`, the system clock of the Core service host is used to determine what time it is.

Boundary can be represented as 0 seconds if you wish to specify the relative time point with no duration offset. This is most useful in the case of the `now` boundary, since it is possible that the highest, latest, time observed in the collection may be much earlier than the current time.

Relative time points can only be used in SDK operations, where there is a collection from which to get the boundaries for earliest and latest time metas.

Relative time points only work on indexed meta types. The default indexed meta types are `time` and `event.time`.

Examples:

```
Last 90m of collection time:
time = rtp(latest, 90m) - u

First 2 days of event time:
event.time = l - rtp(earliest, 48h)

Events added in the last hour:
time = rtp(now, 60m) - rtp(now,0s)
```

Special Range Values

Ranges are normally expressed with the syntax *`smallest`* – *`largest`*, but there are some special placeholder values you can use in range expressions. You can use the letter `l` to represent the lower-bound of the all meta values as the start of the range, and `u` to represent the upper bound. The bounds are determined by looking at the smallest or largest meta value found in the index out of all the meta values that have already entered the index.

If you use the `l` or `u` tag, it should be unquoted.

For example, the expression `time = "2014-may-20 11:57:00" - u` would match all time from that 2014-may-20 11:57:00 to the most recent time found in the collection.

Notice that it is easy to confuse a range expression with a text string. Make sure that text values that contain – are quoted, and that hyphens within range expressions are not within quoted text.

`group by` Clause (since 10.5)

The query API has the ability to generate aggregate groups from the results of a query call. This is done using a `group by` clause on the query. When `group by` is specified, the result set for the query is subdivided into groups. Each group of results is uniquely identified by the meta values indicated in the group by clause.

For example, consider the query `select count(ip.dst)`. This query returns a count of all ip.dst metas in the database. However, if you add a `group by` clause, like this: `select count(ip.dst) group by ip.src`, the query returns a count of the ip.dst metas found for each unique ip.src.

As of version 10.5, you can utilize up to 6 meta fields in a `group by` clause.

The `group by` clause shares some of the same functionality as the `values` call, but it offers significantly more advanced groups at the expense of longer query times. Producing the results of a grouped query involves reading the meta from the meta database for all sessions that match the `where` clause, while a values call can produce its aggregates by reading the index only.

The contents of each group returned by the query are defined by the `select` clause. The `select` clause can contain any of the aggregate functions or meta fields selected. If multiple aggregates are selected, the result of the aggregate function is defined for each group. If nonaggregate fields are selected, the meta fields are returned in batches for each group.

The result set of a `group by` query is encoded with the following rules:

1. All meta items associated with a group are delivered with the same group number.
2. The first meta items returned to the group identify the group key. For example, if the `group by` clause specifies `group by ip.src`, then the first meta item of each group will be an `ip.src`.
3. The normal, nonaggregate meta items are returned after the `group key`, but they all will have the same group number as the group key metas.
4. The aggregate result meta fields for each group are returned next.
5. All fields within a group are returned together. Different group results will not be interleaved.

If one of the `group by` meta items is missing from one of the sessions matched by the `where` clause, that meta field is treated as a NULL for the purposes of that group. When the results for that group are returned, the NULL-valued parts of the group key will be omitted from the group's results, since the database has no concept of NULL.

The semantics of a `group by` query differ from a SQL-like database in terms of what meta fields are returned. SQL databases require you to select the `group by` columns explicitly in the `select` clause if you want them to be returned in the result set. The NetWitness Core database always implicitly returns the group columns first.

A query with a `group by` clause honors the result set `size` parameter if one is provided. However, due to the nature of the grouping, it puts an additional burden on the caller to page and reform groups if a fixed-size result set is requested. For this reason, you should not specify an explicit result size when making a `group by` call. By not specifying an explicit size, the entire result set will be delivered as partial results.

`group by` clauses allow results to be grouped by an entity definition.

The following table describes the database honors configuration parameters that limit I/O or memory impact of a group by query.

Parameter	Function
`/sdk/config/max.query.groups`	This is the limit on how many groups can be held in memory to calculate aggregates. This parameter allows you to limit the overall memory usage of the query.
`/sdk/config/max.where.clause.sessions`	This is the limit on how many sessions from the where clause can be processed in a query. This parameter allows you to set a limit on the number of sessions that have to be read from the meta and session databases to resolve a query.

`order by` Clause (since 10.5)

An `order by` clause can be added to a query that contains a `group by` clause. The `order by` clause causes the set of grouped results to be returned in sorted order.

An `order by` consists of a set of items to sort by in ascending or descending order. Sorting can be performed on any data field that will be returned in the result set. This includes meta specified by the `select` clause, aggregate function results specified by the `select` clause, or `group by` meta fields.

The `order by` clause can sort over many columns. There is no limit on the number of `order by` columns allowed in the query; but a practical limit exists in that each of the `order by` columns must refer to something returned by the `select` clause or `group by` clause. The multiple column sort is imposed lexicographically, meaning that if two groups have equal values for the first column, then they are sorted by the second columns. If they are equal in the second column, they are sorted by the third column, and so on for however many `order by` columns are provided. Groups that do not contain any of the metas referenced by the `order by` clause are sorted first in the result set in the case of an ascending sort, and last in the case of a descending sort.

The NetWitness Core database is unique in that the groups of results returned by a query may each have many values for a selection. For example, it is possible to select all meta items that match a meta type and organize them into groups, and it is possible to use the `distinct()` function to return groups of distinct meta values.

If an `order by` clause references one of the fields in the group that has multiple values, the behavior of the returned result set depends on whether the query is processed as a **streaming** query or a **non-streaming query**.

Streaming queries are those queries where the results are returned to the client as soon as they are found in the database. This is the common use case for queries where the order by cause refers to an indexed meta key. The index for the meta key is used to determine which sessions to stream back to the client first. Streaming queries allow for fast server-side sorting of results. They can be used to allow `order by` clauses on much larger data sets than non-streaming queries, and they return results much faster than non-streaming queries. When a streaming query is sorted, the behavior is as follows:

1. The set of meta values matching the order by clause is fetched from the index, and sorted according to the `order by` direction (`ascending` or `descending`).
2. For each sorted meta value, the sessions containing that value are fetched from the database and returned to the client one value at a time. As those sessions groups are found, they are sent to the client and then discarded to free memory.

Note that the effect this has on sessions that might get fetched for more than one meta distinct value in step 2. When that happens the session might appear more than once in the result set stream. This is logically valid because the session might have more than one 'correct' position in the sorted result set.

Non-streaming queries are those queries that cannot be streamed back to the client and require the Core database to hold the entire result set in memory before determining the final sort order. The most common case for this is if the ordering criteria is an aggregate function. The value generated by the aggregate function has to be calculated before the sort order can be determined, so the results cannot be streamed back to the client as they are found.

1. Within each group, the fields with multiple matching values are ordered by the ordering clause
2. All the groups are sorted by comparing the first occurrence of the ordered field found within each group

The `order by` clause is only available in queries that have a `group by` clause, since groups are required to organize the meta fields into distinct records. If you wish to sort an arbitrary query as if there were no grouping applied, use `group by sessionid`. This ensures that results are returned in groups of distinct sessions or events.

`group by` clauses are naturally returned in ascending group key order; but, an `order by` clause can be used to return groups in a different order.

If an `order by` column does not specify `asc` or `desc`, the default ordering is ascending.

Examples:

```
select countdistinct(ip.dst) GROUP BY ip.src ORDER BY
countdistinct(ip.dst)
select countdistinct(ip.dst) GROUP BY ip.src ORDER BY
countdistinct(ip.dst) desc
select countdistinct(ip.dst),sum(size) GROUP BY ip.src ORDER
BY sum(size) desc, countdistinct(ip.dst)
select sum(size) GROUP BY ip.src, ip.dst ORDER BY ip.dst desc
select user.dst,time GROUP BY sessionid ORDER BY user.dst
select * GROUP BY sessionid ORDER BY time
```

`search` parameter

The `query` API supports a `search` parameter to perform free text searching. The syntax of the search parameter is identical to the search parameter utilized by the `msearch` API call. Refer to the `msearch` documentation for a description of the search field syntax.

The `search` parameter acts as an extension of the `where` clause in the `query` parameter. This means that the `query` and `search` parameters work together. Use the `query` parameter to specify the `select` clause, the `group by` clause, or the `order by` clause. Any `where` clause criteria specified in the `query` parameter are combined with the search filter as if they were joined with an `AND` operation.

Searches through the `query` API are always done against indexed meta, in case-insensitive mode. It has the same behavior as specifying flags `si,sm,ci` to the `msearch` API.

/sdk search

The *search* command is deprecated in favor of the more powerful *msearch* command.

/sdk session

The *session* command maps a session range on the device to it's corresponding meta range.

Parameters:

- **id1** - The starting session ID. If zero, will scope up to the first valid Session ID on device.
- **id2** - The ending session ID. If zero, will scope up to the last valid Session ID on device.

Returns:

- **id1** - The starting session ID.
- **id2** - The ending session ID.
- **field1** - The starting meta ID that corresponds to *id1*.
- **field2** - The ending meta ID that corresponds to *id2*.

/sdk summary

The summary call returns a lot of encoded information that provides global statistics on the databases. The values are encoded as <name>="<value>" parameters separated by whitespace. The quotes are optional if no whitespace is in the value.

Definitions:

- **mid1** = The starting valid meta id
- **mid2** = The ending valid meta id
- **msize** = The current size of the meta database (bytes)
- **mmax** = The maximum size of the meta database (bytes)
- **pid1** = The starting valid packet id
- **pid2** = The ending valid packet id
- **psize** = The current size of the packet database (bytes)

- **pmax** = The maximum size of the packet database (bytes)
- **ptime1** = The time of the first packet in the packet database (POSIX time)
- **ptime2** = The time of the last packet in the packet database (POSIX time)
- **time1** = The first valid time from the index (POSIX time, seconds since 1970)
- **time2** = The last valid time from the index (POSIX time, seconds since 1970)
- **sid1** = The first valid session id
- **sid2** = The last valid session id
- **ssize** = The current size of the session database (bytes)
- **smax** = The maximum size of the session database (bytes)
- **stotalsize** = The total size of the summary index database (bytes)
- **isize** = The total size of index in memory (bytes)

/sdk timeline

Retrieve counts for session/size/packets across a time range, results are given in discrete time intervals of minutes, hours or days. This command only works with meta fields that are of type *TimeT*.

Internally, the *timeline* call performs a *values* call on the passed in fieldName (usually "time"). The *values* call returns the counts for the where clause in minute intervals (you can literally get the same data just using *values*). However, what the *timeline* does do is perform what's known as a "roll up". For request intervals less than or equal to 24 hours, it just returns the counts for the minutes within the time range. For intervals more than a day but less than a month, it will roll up the counts into "hours" and return the aggregate counts for each hour within the time range (taking into account the passed in timezone). If the time range is larger than a month, then it will perform the roll up as "days", again taking into account the passed in timezone of the client.

```
Roll Up Reasoning: The main function of this API is for
graphing session counts over a time period.  The "roll up"
helps provide a readable graph over large time periods.  Most
monitors can't display a graph at the minute level over more
than a couple of days, so the roll up keeps the data transfer
reasonable and provides an easy to read graph.  Should the
user zoom in further (less than a day), the graph will then
provide greater detail at the minute level.  If for some
reason you don't want the "roll up", it's always possible to
make a _values_ call for the same data at the minute level
only.

Another feature of the _timeline_ call is an in memory cache
of time periods greater than a day.  The cache really improves
performance over large time ranges by taking the load off the
index for often requested data that changes very little.  The
```

```
cache stores results in hour increments and ages it out after
an hour.  It is possible for the cache to grow stale for data
received within the last hour, so it's possible to pass the
"ignore-cache" flag so it hits the index for the latest data.
Just keep in mind that this will result in slower queries and
is definitely *not recommended* when the time range is very
large (weeks or months).
```

Parameters:

- **time1** - The starting time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS), if greater than time2, then the whole time range is returned, which is the default behavior is neither **time1** or **time2** is supplied.
- **time2** - The ending time in seconds since 1970 or a time string (YYYY-MMM-DD HH:MM:SS). If not provided, the default value is zero (EPOCH).
- **timezone** - The timezone of the local computer to receive the data, hour offset from GMT [-13 to 13]. If zero, results are returned in UTC.
- **size** - The max number of entries to return
- **flags** - The flags to use for the timeline. Can be a number (bitwise mask) or comma separated values like sessions, size, packets, order-ascending, order-descending, ignore-cache or clear-cache.
- **where** - Optional where clause for filtering the data. Only counts of sessions that match the where clause will be returned.
- **bookendHours** - The number of hours outside the given time1, time2 which will become part of the cache. Default is 24 hours. A larger number of hours can make the current call slower, but can make subsequent calls much faster that fall within this timerange + bookend hours.
- **fieldName** - The meta field to return time values on. The default value is *time*. NOTE: *time* is highly optimized in the index and will typically perform much better than other *TimeT* fields like *event.time*.

The time intervals returned are in minutes, hours or days. The interval cannot be specified by the caller, but is controlled by the amount of data requested. Any time1/time2 interval of 24 hours or less will always be returned in minute intervals. There isn't an option to return the values in seconds. If you need exact values, then you must use the *values* or *query* command. Keep in mind that *time* is always rounded down to the minute in the index (for maximum performance), whereas *event.time* is maintained at the second level. However, if you need to retrieve exact *time* values, you can use the *query* command to retrieve them from the Meta DB, which stores them accurately.

Each result returned is made up of the following fields:

- **id1** - The minimum session ID for this time period

- **id2** - The maximum session ID for this time period
- **count** - The count of sessions/size/packets as requested in the *flags* parameter for this time period
- **format** - Always 32 which is date-time
- **value** - The time period for this result. Will always be set to the start of the minute, hour or day as defined by *type*.
- **type** - Indicates whether the result is a *minute*, *hour* or *day* interval.
- **flags** - Always zero
- **group** - Always zero

/sdk validate

Validate a query or values call without execution.

This API is used to do syntax validation on the `query` and `values` call. The two modes of operation are selected using the `op` paramter.

If validating a `values` call, provide the where clause and fieldName using their respective parameters.

If validating a `query` call, provide the `query` syntax using the query parameter.

/sdk values

The index provides a low-level `values` function to access the unique meta values that have been stored in the index. This function allows developers to perform more advanced operations on groups of unique meta values.

The `values` call parameter syntax:

```
values-params          = field-name-param, space, where-param,
space, size-param, {space, flags-param} {space, start-meta-
param}, {space, end-meta-param}, {space, threshold-param},
{space, aggregate-func-param}, {space, aggregate-field-param},
{space, min-param}, {space, max-param}, {space, search-param}
;
field-name-param       = "fieldName=", (meta-key | entity) ;
where-param            = "where=", where-clause ;
size-param             = "size=", ? integer between 1 and
1,677,721 ? ;
start-meta-param       = ? same as query message ?
```

```
end-meta-param         = ? same as query message ?
flags-param            = "flags=", {values-flag, {"," values-
flag} } ;
values-flag            = "sessions" | "size" | "packets" |
"sort-total" | "sort-value" | "order-ascending" | "order-
descending" ;
threshold-flag         = "threshold=", ? non-negative integer ?
;
aggregate-func-param   = "aggregateFunction=", { aggregate-
func-flag } ;
aggregate-func-flag    = "count" | "sum" ;
aggregate-field-param = "aggregateFieldName=", ( meta-key |
entity ) ;
min-param              = "min=", meta-value ;
max-param              = "max=", meta-value ;
search-param           = "search=", search-string ;
```

The `values` call provides the function of returning a set of unique meta values for a given meta key. For each unique value, the `values` call can provide an aggregate total count. The function used to generate the total is controlled by the flags parameter.

Parameters

The following table describes the function of each parameter.

Parameter	Function
`fieldName`	This is the meta key name for which you retrieve unique values. For example, if `fieldName` is `ip.src`, this function returns the unique source IP values in the collection. Entities can be used for the field name, in which case the result is defined as the combined set of field values for all the referenced meta keys. If the `fieldName` refers to a key with rename references, the result is defined as the combined set of field values for the given meta key name plus all of the references' meta keys.
`where`	This is a `where` clause which filters the set of sessions for which the unique values are returned. For example, if the `fieldName` is `ip.src`, and the `where` clause is `ip.src = 192.168.0.0/16`, only values in the range of `192.168.0.0` to `192.168.255.255` are returned. For information on the `where` clause syntax, see *Where Clauses*.

Parameter	Function
`size`	The size of the set of unique values to return. This function is optimized to return a small subset of the possible unique values in the database.
`id1`, `id2`	These optional parameters limit the scope of the search for unique values to a specific region of the meta database and the index. Setting the `id1` and `id2` parameters to a limited range of the meta database is very important to running searches quickly on large collections.
`flags`	Flags control how the values are sorted and totaled. Flags are described in the following Values Flags section.
`threshold`	Setting the `threshold` parameter allows the `values` call to short-cut collection of the total associated with each value once the threshold is reached. By providing a threshold, the caller can reduce the amount of index and meta items that must be retrieved from the database. If the `threshold` parameter is omitted or set to 0, this optimization is not used.
`aggregateFunction`	Optional parameter used to change the default behavior from counting sessions, packets, or size to counting or summing the numeric field defined by `aggregateFieldName`. Both parameters must be specified when either is defined. Pass either `sum` or `count` to specify which behavior to perform.
`aggregateFieldName`	The meta field on which to perform the `aggregateFunction`. Both `aggregateFunction` and `aggregateFieldName` parameters must be specified when the `aggregate` flag is set. Performing a `values` call using one of the aggregate functions can be significantly slower than a `values` call that collects totals of sessions, packets, or size. The reason for this is that each session that matches the `where` clause must be retrieved from the meta database. This scan causes a large portion of the query to be I/O bound on the meta DB volumes. The time taken to run an aggregate `values` call is linearly proportional to the number of sessions that match the `where` clause.
`min`, `max`	The minimum and maximum value that should be returned from the call. These parameters are used to iterate (or page) over an extremely large number of values, typically more values than could be returned from a single call.

Parameter	Function
	Primarily used in conjuction with the flags `sort-value,sort-ascending` such that the highest value returned would be used in a subsequent call as the `min` parameter value. The values are exclusive. If `min="rsa"` was specified and `rsa` was a valid value, `rsa` would not be returned; instead, the next highest value would be returned.
`search`	Text search pattern to be used to further refine the `where` parameter

`values` Flags

The `flags` parameter controls how the values call operates. There are three groups of flags that correspond to the different modes of operation as shown in the following table.

Flag	Description
`sessions`, `size`, `packets`	The `values` call allows you to specify one of these flags to determine how the total for each value is calculated. If the flag is `sessions`, the `values` call returns a count of sessions that contain each value. If the flag is `size`, the `values` call totals the size of all sessions that contain each unique value, and reports the total size for each unique value. If the flag is `packets`, the values call totals the number of packets in all sessions that contain each unique value, and then reports that total for each unique value.
`sort-total`, `sort-value`	These flags control how results are sorted. If the flag is `sort-total`, the result set is sorted in order of the totals collected. If the flag is `sort-value`, the results are returned in order of the sorting order of the values.
`order-ascending`, `order-descending`	These flags control the sort order of the result set. For example, if sorting by total in descending order, the values with the greatest total are returned first.
`suggest`	Enables suggestion mode for the values API. All other flags are ignored if this flag is set
`database-scan`	Make the values call bypass the index and instead collect unique values as if where traversing the meta database.

Flag	Description
	This mode is slow on most cases, but it can be fast if the where clause matches very few sessions.
`ignore-cache`, `clear-cache`	These flags control result set caching on the set of values returned by this call. Normally these should not be used.

`values` Call Example

The `values` call is used extensively by the Navigation view in NetWitness Platform. The default view generates calls that look like this:

```
/sdk/values id1=198564099173 id2=1542925695937 size=20
flags=sessions,sort-total,order-descending threshold=100000
fieldName=ip.src where="time=\"2014-May-20 13:12:00\"-\"2014-
May-21 13:11:59\""
```

In this example, the Navigation view requests unique values for `ip.src`. It requests unique values of `ip.src` in the time range given. It asks for the count of sessions that match each `ip.src`, and the results are the top 20 `ip.src` values when sorted by the number total count of sessions in descending order. In addition, the Navigation view has a meta ID range in order to provide an optimization hint to the query engine.

Values call and bucketing mode

When a values call is executed with a `fieldName` parameter that specifies a bucketed indexed meta, the system will only return the bucket values present within the rest of the criteria. This has the side effect of producing counts and totals that represent all sessions within each returned bucket. This is useful because it summarizes size meta into groups that represent human-readable ranges like 1 MB, 2 MB, and so one.

When the number of sessions scanned by the values call drops below 1000 sessions, the values call operates in meta-scanning mode, and at that point it returns exact values for numeric value indexes, regardless of the bucketing setting on the index.

Suggestion Mode

The `values` call has an additional execution mode that is used to provide suggested search values. In this mode of operation, the `values` call only identifies unique values stored with the given meta key name. It provides these results within milliseconds. To achieve this it does not provide any session counts, it will not utilize any other sort flags. The suggestion mode does utilize the `where` parameter to refine suggestions, but it only

utilizes the time range clause if provided. Other portions of the `where` clause are not utilized to refine suggestion.

Suggestion mode is enabled by setting the `suggest` flag in the `flags` parameter.

Suggestion mode gives special meaning to the `min` parameter. The `min` parameter can contain the starting point for the suggested values. The return values of `suggest` mode will only include values that start with the text provided in the `min` parameter.

`search` parameter

The `values` API supports a `search` parameter to perform free text searching. The syntax of the search parameter is identical to the search parameter utilized by the `msearch` API call. Refer to the `msearch` documentation for a description of the search field syntax.

The `search` parameter acts as an extension of the `where` parameter. This means that the `where` and `search` parameters work together. Any `where` parameter specified is combined with the search filter as if they were joined with an `AND` operation.

Searches through the `values` API are always done against indexed meta, in case-insensitive mode.

The `values` API only operates on index entries for most requests, in order to provide fast totals over a large numbers of events. When the search parameter operates over in-exact indexes, such as N-Gram indexes, it may include sessions with near matches instead of narrowing the search to exact matches.

/sys

The system tree holds the main functionality related to overall service configuration, scheduler and access to the Stats Database.

The following is a brief overview of the system commands and what they do:

- **save** - Saves the current configuration to disk. Any configuration change will be written to disk within a couple of minutes or shutdown, so this is typically redundant.
- **caCert** - List, add or delete trusted certs on the system.
- **shutdown** - Saves all data and then performs a clean shutdown of the service. Typically the service will restart based on systemd settings.
- **fileEdit** - Used to view and edit configuration files, index files, parsers and even the scheduler jobs.

- **peerCert** - List, add or delete trusted peer certs on the system. This enables clients to connect using a trusted SSL connection.
- **servCert** - Displays the current server certificate in PEM format
- **statHist** - Query historical stats stored in the Stats Database.

/sys caCert

Parameter	Description
op	Operation to perform on the CA certificate list (list, delete, add)
id	ID of CA certificate to remove (only valid with op=delete)

Manipulate the certificate authority (CA) list used by this service. The Certificate Authority list is used to verify the authenticity of peer certificates.

You must specify what operation to perform on the certificate authority list: list, delete, or add

The **op=list** command will output a list of the known certificate authority certificates.

The **op=delete** command will remove a certificate from the certificate authority list. The delete command must be accompanied by an id parameter. The id refers to the id of the certificate authority certificate to remove from the certificate authority list. Use the list command to see a list of certificate authority id's.

The **op=add** command will upload an additional certificate to the service, and add it to the certficiate authority list. This command can only be invoked programmatically. The REST interface provides a simple HTML interface to upload a certificate authority certificate using the add operation. This is available at /sys/cauploa on the REST interface.

/sys fileEdit

Parameter	Description
op	The operation to perform (dir, get, getBackup, put, create, delete)
type	The type of file (parser, filter). If not provided, it will use the file extension of the filename parameter.
filename	The filename to create, edit, retrieve or delete.

The purpose of this command is to provide an API to retrieve the contents of files that can be edited by the user. This includes Lua parsers, custom index files, scheduler jobs and a few other miscellaneous files. The workflow around this API is thus:

1. Retrieve the list of files that can be edited by sending the parameter **op=dir**. This will return a file list pair containing the filename and it's size in bytes.
2. Using the **op=get filename=<filename>** command, retrieve the contents of the file to edit.
3. After editing, you can then save the file back on the service by using **op=put filename=<filename>**.

When using the put operation, the actual contents of the file must come after the parameters using the message type called PARAMS_BINARY. This message type can be sent by NwConsole using the following example command:

```
send /sys fileEdit op=put filename=index-concentrator-
custom.xml --file-data="/Users/user/Documents/index-
concentrator-custom.xml"
```

The **op=create type=parser filename=MyLuaParser.lua** command can be used to add new Lua parsers to the system. This will not automatically reload the parser.

To delete files, you can use **op=delete filename=<filename to delete>**.

/sys peerCert

Parameter	Description
op	Operation to perform on the peer certificate list (list, delete, add)
id	ID of CA certificate to remove (only valid with op=delete)

Manipulate the peer certificate list used by this service. The peer certificate list is used to grant super-user, or "trusted" access to a connection. The if the client provides a certificate that matches one of the certificates in the peer certificate list, the client is trusted. The peer certificate is *not* trusted based on certificate chain trust. Trust is only granted if the client's certificate is explicitly added to the peer list.

You must specify what operation to perform on the peer certificiate list: list, delete, or add

The **op=list** command will output a list of the known peer certificates.

The **op=delete** command will remove a certificate from the peer certificate list. The delete command must be accompanied by an id parameter. The id refers to the id of the peer certificate to remove from the peer list. Use the list command to see a list of peer id's.

The **op=add** command will upload an additional certificate to the service, and add it to the peer certificate list. This command can only be invoked programmatically. The REST interface provides a simple HTML interface to upload a peer certificate using the add operation. This is available at /sys/trustpeer on the REST interface.

/sys save

Parameter	Description
force	Forces the configuration to be written to disk, regardless if there are any changes that require a save

Saves the current configuration to disk. This command is redundant because the system automatically saves changes within a couple of minutes and always on shutdown.

/sys servCert

Display the public certificate for this service, in PEM format.

/sys shutdown

Parameter	Description
reason	The reason for the shutdown. This is saved to the log.
cl	Short for command line. This is an advanced parameter used to set options on the command line of the service upon restart.

This will perform a normal shutdown of the service, saving all files and then exit cleanly. Typically, the service will restart based on the systemd configuration.

/sys statHist

Parameter	Description
time1	The starting time for retrieving stats.
time2	The ending time for retrieving stats.
timeFormat	Specify the time format for each stat snapshot (posix, simple), default is posix (seconds since 1970).
include	Comma separated list of stats to include (wildcards allowed)
exclude	Comma separated list of stats to exclude (wildcards allowed, has precedence over include)
reduce	If true, reduces data transmission by replacing stat pathnames with a shorthand and provides a lookup table as first result.
onChange	Comma separated list of stats for tracking changes. Only results where at least one stat in the list has changed will be returned.
showAll	If true, changes each result to show all stats previously seen, even if the value hasn't changed. It has a moderate performance impact and is best used with include, exclude or onChange parameters to reduce the returned result set size.

Retrieve historical stats from the stats db. Don't send time1/time2 to get bounding times about stats db. Supported wildcards are ? to match any single char and * to match zero or more characters, not including slash /, ** to match zero or more characters including slash /.

With no params, you get information about the how many records are in the database and the bounding time range:

```
/sys statHist
[
       total: 1188
       time1: 1336497862
       time2: 1336499092
]
```

By default, it returns it in posix time, which is the number of seconds since 1970. You can ask for a readable date:

```
/sys statHist timeFormat=simple
[
       total: 1260
       time1: 2012-May-08 17:24:22
```

```
        time2: 2012-May-08 17:46:08
]
```

You can craft a query to return stats within a time range like so (some results snipped for readability):

```
/sys statHist time1="2017-01-19 15:00:00" time2="2017-01-19
15:00:02" timeFormat=simple
[
        /logs/stats/first.id=202028111
        /logs/stats/last.failure.id=205165789
        /logs/stats/last.id=207753128
        /logs/stats/last.warning.id=205849692
        /logs/stats/total=5725018
        /rest/stats/service.status=Ready
        /sdk/stats/queries.active=0
        /sdk/stats/queries.pending=0
        /sdk/stats/values.calls=10369
        /sdk/stats/values.calls.cached=1
        /storedproc/stats/storedproc.active=0
        /storedproc/stats/storedproc.pending=0
        /sys/stats/config.filename=NwConcentrator.cfg
        /sys/stats/cpu=0%
        /sys/stats/hostname=LOKI
        /sys/stats/memory.process=31838142464
        /sys/stats/memory.process.max=135460544512
        /sys/stats/memory.system=116689436672
        /sys/stats/memory.total=135460544512
        /sys/stats/module=concentrator
        /sys/stats/release=0.0.0
        /sys/stats/version=11.0.0.0
        time="2017-Jan-19 15:00:00"
]
[
        /logs/stats/last.id=207753130
        /logs/stats/total=5725020
        time="2017-Jan-19 15:00:01"
]
[
        /sys/stats/memory.system=116689952768
        time="2017-Jan-19 15:00:02"
]
```

Notice how the first second has a bunch of stats, but the next 2 do not? What happens is the stat db only stores values that have changed since the last update. However, that would be a problem for stats that rarely change if you wanted to plot them in a chart, so every 5 mins on the 5 min mark, *every* stat is stored again. When the service first starts, it also stores every stat in the db, then follows with only stat changes until the next 5 min mark. Bottom line: when charting, it's best to make sure time1 is on a 5 minute mark.

You can also limit the stats you are interested in by using include or exclude:

```
/sys statHist time1="2017-01-19 15:00:00" time2="2017-01-19
15:00:05" timeFormat=simple include=/sys/**,/logs/stats/*
```

This command would show the results above, but only for stats under the whole /sys tree and the immediate stats under /logs/stats.

If you would like to see every stat value with every time period, you can use the parameter **showAll=true**. This parameter will carry forward every stat even if the value hasn't changed. However, performance is noticeably impacted when doing this, especially when querying large time ranges. It will also greatly expand the size of the result set and is best used with **include**, **exclude** or **onChange** to reduce the returned size.

For example, the sample above would look like this with **showAll=true** :

```
/sys statHist time1="2017-01-19 15:00:00" time2="2017-01-19
15:00:02" timeFormat=simple showAll=true
[
       /logs/stats/first.id=202028111
       /logs/stats/last.failure.id=205165789
       /logs/stats/last.id=207753128
       /logs/stats/last.warning.id=205849692
       /logs/stats/total=5725018
       /rest/stats/service.status=Ready
       /sdk/stats/queries.active=0
       /sdk/stats/queries.pending=0
       /sdk/stats/values.calls=10369
       /sdk/stats/values.calls.cached=1
       /storedproc/stats/storedproc.active=0
       /storedproc/stats/storedproc.pending=0
       /sys/stats/config.filename=NwConcentrator.cfg
       /sys/stats/cpu=0%
       /sys/stats/hostname=LOKI
       /sys/stats/memory.process=31838142464
       /sys/stats/memory.process.max=135460544512
       /sys/stats/memory.system=116689436672
       /sys/stats/memory.total=135460544512
       /sys/stats/module=concentrator
       /sys/stats/release=0.0.0
       /sys/stats/version=11.0.0.0
       time="2017-Jan-19 15:00:00"
]
[
       /logs/stats/first.id=202028111
       /logs/stats/last.failure.id=205165789
       /logs/stats/last.id=207753130
       /logs/stats/last.warning.id=205849692
```

```
        /logs/stats/total=5725020
        /rest/stats/service.status=Ready
        /sdk/stats/queries.active=0
        /sdk/stats/queries.pending=0
        /sdk/stats/values.calls=10369
        /sdk/stats/values.calls.cached=1
        /storedproc/stats/storedproc.active=0
        /storedproc/stats/storedproc.pending=0$
        /sys/stats/config.filename=NwConcentrator.cfg
        /sys/stats/cpu=0%
        /sys/stats/hostname=LOKI
        /sys/stats/memory.process=31838142464
        /sys/stats/memory.process.max=135460544512
        /sys/stats/memory.system=116689436672
        /sys/stats/memory.total=135460544512
        /sys/stats/module=concentrator
        /sys/stats/release=0.0.0
        /sys/stats/version=11.0.0.0
        time="2017-Jan-19 15:00:01"
]
[
        /logs/stats/first.id=202028111
        /logs/stats/last.failure.id=205165789
        /logs/stats/last.id=207753130
        /logs/stats/last.warning.id=205849692
        /logs/stats/total=5725020
        /rest/stats/service.status=Ready
        /sdk/stats/queries.active=0
        /sdk/stats/queries.pending=0
        /sdk/stats/values.calls=10369
        /sdk/stats/values.calls.cached=1
        /storedproc/stats/storedproc.active=0
        /storedproc/stats/storedproc.pending=0
        /sys/stats/config.filename=NwConcentrator.cfg
        /sys/stats/cpu=0%
        /sys/stats/hostname=LOKI
        /sys/stats/memory.process=31838142464
        /sys/stats/memory.process.max=135460544512
        /sys/stats/memory.system=116689952768
        /sys/stats/memory.total=135460544512
        /sys/stats/module=concentrator
        /sys/stats/release=0.0.0
        /sys/stats/version=11.0.0.0
        time="2017-Jan-19 15:00:02"
]
```

As you can see, the result size grew and even the stat values that didn't change are included.

A very useful parameter is **onChange=<stat pathname>[,<stat pathname>]**. This will only return results when the listed stats change. Only one stat has to change, not all of them, to be returned.

For example, suppose you want to see when Decoder had packet drops and the capture rate at that time:

```
time1="2017-01-19 15:00:00" time2="2017-01-19 15:10:00"
timeFormat=simple
include=/decoder/stats/capture.rate,/decoder/stats/capture.dro
pped,time onChange=/decoder/stats/capture.dropped
[ /decoder/stats/capture.rate=3876
/decoder/stats/capture.dropped=102 time="2017-Jan-19 15:05:57"
]
[ /decoder/stats/capture.dropped=210 time="2017-Jan-19
15:08:02" ]
```

So, in the result set above, the stat **capture.dropped** changed twice in the 10 minute period requested. But for the second result, we don't know what the capture rate was because in that time period, it didn't change from a previous time period and wasn't recorded. This is where **showAll=true** comes in handy:

```
time1="2017-01-19 15:00:00" time2="2017-01-19 15:10:00"
timeFormat=simple
include=/decoder/stats/capture.rate,/decoder/stats/capture.dro
pped,time onChange=/decoder/stats/capture.dropped showAll=true
[ /decoder/stats/capture.rate=3876
/decoder/stats/capture.dropped=102 time="2017-Jan-19 15:05:57"
]
[ /decoder/stats/capture.rate=4205
/decoder/stats/capture.dropped=210 time="2017-Jan-19 15:08:02"
]
```

Now we know what the capture rate was at the time of the second drop. The query kept track of previous values for **capture.rate** and placed the value in the second result. This can be very useful when doing adhoc queries, but keep in mind it will be somewhat slower to return results.

Configuration of the Stats Database

There are 3 configuration nodes that affect how the stat database works. They are all under /sys/config.

- **stat.compression** - Determines whether or not the records in the stat db are compressed or not. The values can be none or gzip. If you enable gzip

compression, the retention of the historical stats will be greatly increased. However, there will be a moderate performance penalty when doing a query, as it will have to uncompress each record on the fly to return results.

- **stat.dir** - The directory on the device where the stats db will be stored. The default location is /var/lib/netwitness/<service>/statdb. More importantly, this value also indicates how large the stat db can grow before it goes into rollover. Each directory can be assigned a maximum size (e.g., /var/lib/netwitness/<service>/statdb=1 GB). If you would like to store more stats, you can increase this size, as long as you have enough disk space.
- **stat.exclude** - This is a comma separated list of stat paths that will be excluded from getting stored in the stats database. The more you exclude, the longer the retention for the stats that are included. A sample value is something like this: /users/roles/*,/connections/**,/services/**,/sdk/stats/queries/**,/decoder/devices/*. As you can see, wildcards are supported, so you can exclude whole subtrees. A single asterisk means only exclude those stats that match at that subtree level, but it is not recursive. Two asterisks means recursively exclude stats in subtrees as well.

/decoder/parsers

Configuration

Dynamic Header Thresholds

The Dynamic Header Thresholds allow headers to identify devices without a message match.

Across devices (parsers), a priority is assigned to each header reflecting the complexity of the `content` attribute (pattern) of the header. The more generic a header is the lower the match priority. The more distinct and complex a header is the higher the match priority.

As example, the following header has a very high match priority (**4503**):

```
%CACHEFLOWELFF_syslog:<hfld1>,cs-
method="<messageid>",<!payload:hfld1>
```

An example of a header with a very low match priority (**8**):

```
<msgIdPart1> <msgIdPart2>, <!payload:msgIdPart1>
```

The explicit match priority value for each header can be viewed in the logs when debug logging is enabled.

As the header matches are tracked, high and low thresholds are dynamically applied proportionately against the header with this highest observed match priority. Header matches above the high threshold can be considered even without a known message group ID. Matches below the low threshold which have a hardcoded message group ID or fail to match a message are removed from consideration. This logic helps increases the likelihood of a single device match for strong headers.

- Threshold Percent (**header.threshold.percent**) - Percentage used to establish the high and low match priority thresholds. Setting of '0' disables this feature. Changes take effect on parser reload.
- Minimum Score (**header.threshold.min.score**) - Minimum observed match priority required to engage the high threshold. Setting of '0' disables the high threshold. Changes take effect on parser reload.

Header Prioritization

Message matches resulting from a Message Group ID extracted from message payload can be prioritized over matches using a hardcoded ID. Though principally content driven, default treatment is configurable to minimize the number of headers which will require edits.

- Shortcut Score (**header.prioritize.shortcut.score**) - Priority score, at and below which, matches resulting from a Message Group ID extracted from message payload are prioritized over those with a hardcoded ID. Content specified treatment takes precedence over configuration. Changes to this value could require changes to content to compensate for changes in header treatment.

Dynamic Mapping

Message matches when flagged by the content elements `<HEADER>` or `<MESSAGE>` with the attribute `discoverable="map"` trigger a temporary event source mapping which directs future parsing to the matched device type. A mapped source is valid for a configured amount of time. Each time a subsequent match is flagged by content, the expiration time is extended for the source.

If directed parsing to the mapped device does not provide a match, parsing will fall back to discovery. This allows the source to be, if required, dynamically mapped to multiple device types. It also allows for handling repurposed source addresses. Each mapped device type maintains its own expiration time.

- Mapping Minutes (**dynamic.mapping.minutes**) - Delta time in minutes for which a content defined mapping should be retained.

Receiving Protobuf Encoded Logs

Logdecoder can receive 4 byte length
prefixed `com.rsa.netwitness.carlos.nextgen.Event` protobuf objects on port 50202 (default). The connection must be initiated with byte string:

```
0x4E 0x47 0x43 0x50 0x01 0x00
```

- `raw_message` - [required] if not present or present with a length of 0, the log will not be re-written. Otherwise, this value, along with the log header will be the serialized packet.
- `collection_meta`
- `lc.msgtype` - [required] only protobuf events that have this field and its value is Unstructured will be re-written as raw logs
- `lc.srcid` - [optional] this will be added to the log header in the source field and typed appropriately (text/ipv4/ipv6), text values will be parsed to determine if they are ipv4/ipv6 addresses.
- `lc.cid` - [optional] this is expected as text and will be added to the log header if present
- `content_meta`
- `syslog.pri` - [optional] Will be added to the log header if present (0 otherwise), expects either a text or uint32 value (text must be parseable as a uint32).

LD checks the following fields in the `Event` instances:

All fields of collection and content meta will be checked for entries in the table map to determine if meta should be created.

Messages

- **tzinfo** - Retrieve and update timezone settings used for log event processing.
- **dyndvmap** - Display, purge and force persistence of dynamically mapped event sources.

/decoder/parsers dyndvmap

The `dyndvmap` command provides an interface to display, purge and force persistence of dynamically mapped event sources.

Parameters and Results

- **op** - Operation to perform: `describe`, `clear` or `save`. Defaults to `describe`
- **deviceFilter** - A regex that will only return matching devices.
- **sourceFilter** - A regex that will only return matching sources.
- **forwarderFilter** - A regex that will only return matching forwarders.

Display Mapped Sources

Expired mappings are only purged on synchronization between parse threads. The reported mappings from this command will only include those which have **not** expired. To view the size of the mappings container query the `parsers` stat `mappings.dynamic.source.count`.

View all **valid** dynamically mapped event sources.

```
dyndvmap op=describe
```

View only mappings for a particular device type.

```
dyndvmap op=describe deviceFilter=rhlinux
```

View only mappings for a particular source.

```
dyndvmap op=describe sourceFilter=1.1.1.1
```

Purge Mappings

Flush **all** mapped sources.

```
dyndvmap op=clear
```

Force Persistence

Force the mappings to be saved to disk.

```
dyndvmap op=save
```

/decoder/parsers tzinfo

The **tzinfo** command provides an interface to view timezone information used for converting dates parsed from a captured log to normalized UTC timestamps. This information is derived from the IANA Time Zone Database installed on the service host. As timezone abbreviations do not necessarily represent a single GMT offset, **tzinfo** also provides a means to override the definition of an abbreviation as well as define entire new abbreviations.

Parameters and Results

- **op** - The operation to perform: `version`, `names`, `abbreviations`, `duplicates`, `overrides`, `add` or `remove`
- **abbr** - a timezone abbreviation (e.g. `EST`)
- **gmtoffset** - an offset from GMT in the format [+/-]HH[MM]
- **tzname** - A timezone name (e.g. `America/New_York`) which must be a known timezone (see **Timezone Names**)

Version

View the version of the IANA Time Zone Database currently in use.

```
tzinfo op=version
```

Timezone Names

View the name, abbreviations and GMT offsets for all known timezones.

```
tzinfo op=tznames
```

For each known timezone, the following will be returned:

- **tzname** - The timezone name
- **stdAbbr** - The abbreviation the timezone currently uses for standard time
- **stdOffset** - The offset from GMT the timezone currently uses for standard time
- **dstAbbr** - The abbreviation the timezone currently uses for daylight savings time

- **dstOffset** - The offset from GMT the timezone currently uses for daylight savings time

An empty value for **dstAbbr** indicates that a timezone does not currently observe daylight savings time. Also note that the offsets returned here reflect the entries in the underlying Time Zone Database and are not impacted by **Overridden Abbreviations**.

Timezone Abbreviations

View all current timezone abbreviations, the default GMT offset and the **Abbreviation Override** offset if one exists.

```
tzinfo op=abbrs
```

For each abbreviation definition, the following will be returned:

- **abbr** - The abbreviation
- **gmtoffset** - The offset from GMT as defined by the timezone database
- **override** - A specified GMT offset used (see **Overridden Abbreviations**)

Duplicates Abbreviations

View all abbreviations that have multiple values for GMT offset in distinct timezones.

```
tzinfo op=duplicates
```

For each abbreviation definition that has multiple offset definitions in the timezone database, the following will be returned:

- **abbr** - The abbreviation
- **gmtoffset** - The offset from GMT
- **tzname** - The timezone that defines this abbreviation and GMT offset

Abbreviations that are currently overridden and have a GMT offset or timezone explicitly set will not be displayed as a duplicate.

Overridden Abbreviations

View all abbreviations that have had GMT offset or timezone explicitly set.

```
tzinfo op=overrides
```

For each overridden abbreviation, the following will be returned:

- **abbr** - The abbreviation
- **gmtoffset** - If the override is a specific offset, offset from GMT
- **tzname** - If the override specifies a timezone
- **valid** - Indicates wether or not the setting was successfully loaded and is currently active

Add/Edit an Abbreviation Override

Add a new abbreviation or override the timezone database definition of an existing abbreviation. Adding an existing abbreviation override will update that abbreviation.

An abbreviation can be defined as a specific GMT offset. The following creates the abbreviation `EQT` with a GMT offset of -5 hours:

```
tzinfo op=add abbr=EQT gmtoffset=-0500
```

A duplicate abbreviation can be overridden from the default definition. The following changes the default offset of `CST` from -6 hours to -5 hours as defined by the timezone `Cuba`:

```
tzinfo op=add abbr=CST tzname=Cuba
```

The timezone must be known (see **Timezone Names**) and define the provided abbreviations as standard or daylight.

Remove an Abbreviation Override

Remove an abbreviation that was added to **Overridden Abbreviations**.

```
tzinfo op=remove abbr=EQT
```

/sdk deviceId

The *deviceId* command takes a session ID on the current device and returns the originating device name and remote session ID on the originating device. This command is only useful for sessions that have been aggregated from another service.

```
NOTE: This command does not track a session down to the
originating device (e.g., Broker -> Concentrator -> Decoder),
only one step removed.  So if you run this on a Broker, it
will give you the session on the Concentrator, not the
Decoder.
```

Parameters:

- **session** - The session ID to look up

Returns:

- **device** - The originating device name that *session* was aggregated from
- **session** - The session ID on the originating device

/appliance

Storage Array Configuration Utilities

The steps for attaching storage to your NetWitness Platform are:

1. Physically attach the storage and make it available to this host. For direct-attach storage, the RAID manipulation commands will perform the task of constructing hardware RAID volumes and making these volumes appear as drives. For SAN storage, you must allocate the storage through the SAN's management tools and present them to this host. For either type of storage, the result will be that the storage appears as block devices, such as /dev/sdc in the case of direct-attach drives, or /dev/emcpowera in the case of an EMC SAN drive.
2. Allocate the block devices as usable storage. You can use the `partNew` command to take block storage and organize it into filesystems to be used by storage. The result of this is that the block device is put into a volume group, such as `decoder`, `concentrator`, or `index`
3. Allocate the storage to a service. You can use the `srvAlloc` command to assign the filesystems that are created to services that are running on this host.

Summary of Array Configuration Commands

Each of the commands listed below has built-in help that describes their function and usage. If you are using the REST interface, select the command from the drop-down menu to see the help text.

Commands that deal with direct-attach RAID volumes

- **raidList** - List the RAID controllers and direct-attach enclosures that are present on this host
- **raidNew** - Allocate direct-attach enclosures into block devices
- **raidClr** - De-allocate direct-attach enclosures used as block devices

Commands that deal with allocating block devices as storage

- **devlist** - List the block devices that are present on this host
- **partNew** - Allocate partitions on a block device and create volume groups
- **partClr** - Remove volume groups from a block device
- **vgs** - Summarize how block devices are organized into volume groups

Commands that deal with allocating storage to services

- **srvList** - List services on the host and their allocated storage paths
- **srvAlloc** - Allocate a volume group to a service
- **srvFree** - Remove a volume gorup from a service

/appliance partClr

Uninitialize a block storage device

This command clears data from a block storage device. It performs these tasks:

1. Unmount any filesystems within volume groups on the device.
2. Remove any volume groups the device is a part of.
3. Wipe the physical volume labels from the device.
4. Create a new, empty partition table on the device.

Notes:

- You must remove the block device from all service configurations.
- You must stop the service if you are removing the last block device from that service.
- All `decoder` volumes must be removed before `decodersmall` volumes
- All `logdecoder` volumes must be removed before `logdecodersmall` volumes
- All `index` volumes must be removed before `concentrator` volumes

/appliance partNew

Initialize a block storage device

This command prepares a storage device to use in NetWitness. It performs these tasks:

1. Creates the partition table on the block device.
2. Creates the Linux Volume Manager physical device partition.
3. Creates a volume group containing the physical device.
4. Creates logical volumes in the volume group.
5. Creates XFS filesystems on each logical volume.
6. Creates /etc/fstab entries for each logical volume.
7. Mounts each logical volume.

Usage

You must provide the block device name, for example, "sdc" or "/dev/sdc". Use the devlist command to locate unused block devices. You must provide a name for the service that will be used with the storage, for example, "decoder" or "concentrator". You may optionally provide the volume type that will be created. The default volume type has the same name as the service.

By default, the partNew command does **not** make changes. Instead it prints a report of the actions that will be taken. To actually make the changes to the system, add the commit=true parameter to the command.

Example: Assign devices `sdd` and `sde` to Decoder.

```
send /appliance partNew name=sdd service=decoder
volume=decodersmall
send /appliance partNew name=sde service=decoder
volume=decoder
```

Note that the decodersmall volume must be created before the decoder volume, because it holds the small filesystem mounted at /var/netwitness/decoder.

Example: Assign device `sdc` to Archiver

```
send /appliance partNew name=sdc service=archiver
```

Note that we do not need to specify a volume type. It is assumed that the volume type is "archiver".

Netwitness Service Volume Reference

Service volume names

Service	Volume Name	Filesystems Created
decoder	decoder	packetdb
decoder	decodersmall	decoder root, index, sessiondb, metadb
decoder	hybrid-decoder-meta	decoder root
decoder	packet	packetdb
logdecoder	logdecoder	packetdb
logdecoder	logdecodersmall	logdecoder root, index, sessiondb, metadb
logdecoder	hybrid-logdecoder-meta	logdecoder root
logdecoder	logpacket	packetdb
logdecoder	endpoint-log-hybrid	mongo, packetdb
logdecoder	log-indexed-decoder	logdecoder root
logdecoder	logindex	logdecoder index
concentrator	concentrator	concentrator root, metadb, sessiondb

Service	Volume Name	Filesystems Created
concentrator	hybrid-concentrator	concentrator root
concentrator	index	index
archiver	archiver	database

Volume sizing

Volume	Filesystem	Mount Point	Size
decodersmall	decoroot	/var/netwitness/decoder	10 GB
decodersmall	index	/var/netwitness/decoder/index	30 GB
decodersmall	sessiondb	/var/netwitness/decoder/sessiondb	600 GB
decodersmall	metadb	/var/netwitness/decoder/metadb	all remaining free space on decodersmall volume
decoder	packetdb	/var/netwitness/decoder/packetdb	all space on decoder volume
packet	packetdb	/var/netwitness/decoder/packetdb	all space on packet volume
hybrid-decoder-meta	decoroot	/var/netwitness/decoder	all space on hybrid-decoder-meta volume
logdecodersmall	decoroot	/var/netwitness/logdecoder	10 GB
logdecodersmall	index	/var/netwitness/logdecoder/index	30 GB
logdecodersmall	sessiondb	/var/netwitness/logdecoder/sessiondb	600 GB
logdecodersmall	metadb	/var/netwitness/logdecoder/metadb	all remaining free space on

Volume	Filesystem	Mount Point	Size
			logdecodersmall volume
hybrid-logdecoder-meta	decoroot	/var/netwitness/logdecoder	all space on hybrid-logdecoder-meta volume
log-indexed-decoder	decoroot	/var/netwitness/logdecoder	all space on log-indexed-decoder volume
logdecoder	packetdb	/var/netwitness/logdecoder/packetdb	all space on logdecoder volume
logpacket	packetdb	/var/netwitness/logdecoder/packetdb	all space on logpacket volume
logindex	index	/var/netwitness/logdecoder/index	all space on logindex volume
concentrator	root	/var/netwitness/concentrator	30 GB
concentrator	sessiondb	/var/netwitness/concentrator/sessiondb	10% of free space on concentrator volume
concentrator	metadb	/var/netwitness/concentrator/metadb	all remaining free space on concentrator volume
hybrid-concentrator	root	/var/netwitness/concentrator	all space on hybrid-concentrator volume
index	index	/var/netwitness/concentrator/index	all space on index volume

Volume	Filesystem	Mount Point	Size
archiver	database	/var/netwitness/archiver/database	all space on archiver volume
endpoint-log-hybrid	mongo	/var/netwitness/mongo	50% of free space on endpoint-log-hybrid volume
endpoint-log-hybrid	packetdb	/var/netwitness/logdecoder/packetdb	all remaining free space on endpoint-log-hybrid volume

/appliance vgs

List Volume Groups

This command enumerates all the volume groups on this host. In addition, it displays the physical volumes that the volume group consists of, and the logical volumes within the volume group.

NetWitness® Platform XDR

Version 12.0

Security Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

Security Configuration Guide Overview 7
Security Configuration Settings 8
Access Control Settings 9
User Authentication 10
- Multi-Factor Authentication 10
- NetWitness Platform Core Trusted Connection 10
- User Accounts 10
- Configuring New Accounts 11
- Authentication Configuration 11
- User Passwords 12
  - NetWitness Users 12
  - NetWitness Core Service Users 12
- Security Parameter Settings 12
- Kibana Health and Wellness Interface 13

User Authorization 14
- Access Roles 14

Component Authentication 15
- Host Configuration and Service Authentication 15
- Changing Credentials for Default Configuration Service Accounts 15
- Configuring Live Account Authentication 15
- Configuring Lockbox Authentication 16
- Change Root Account Password On All NW Hosts 16
- Display Logon Banner for Remote SSH Connections 16

Secure Boot Loader 16
Disable Interactive Startup 17
- Create a Customized Logon Banner for NetWitness Platform 17

Log Settings 18
- Log Description 18
- Log Management and Retrieval 18

Communication Security Settings 20
- Port Usage 20

NetWitness Network Architecture Diagram 21
NetWitness Network (Packets) Architecture Diagram with Ports 22
NetWitness Logs Architecture Diagram with Ports 23
Event Stream Analysis Network (Packets) Architecture Diagram with Ports 24

Event Stream Analysis (Logs) Architecture Diagram with Ports ... 25
NetWitness Firewall Requirements Summary ... 26
Comprehensive List of NetWitness Host, Service, and iDRAC Ports ... 30
NW Server Host (Primary and Warm Standby NW Server Host) ... 31
Analyst UI Host ... 32
Archiver Host ... 33
Broker Host ... 34
Concentrator Host ... 35
Endpoint Log Hybrid ... 36
Endpoint Relay Server ... 37
Event Stream Analysis (ESA) Host ... 38
New Health and Wellness ... 39
New Health and Wellness on Different Subnet ... 39
iDRAC Ports ... 40
Log Collector Host ... 41
Log Decoder Host ... 42
Log Hybrid Host ... 43
Log Hybrid - Retention Host ... 44
Malware Host ... 45
Network Decoder Host ... 46
Network Hybrid Host ... 47
UEBA Host ... 48
NetWitness Endpoint Architecture ... 49
NetWitness Endpoint 4.4 Integration with NetWitness Platform ... 50
NetWitness Endpoint Architecture with Ports ... 50
How to Change UDP Port for Endpoint Log Hybrid ... 51
Task 1 - Tell All Agents to Use a New UDP Port ... 51
Task 2 - Update the Port on All Endpoint Log Hybrid Hosts in Your Environment ... 51
Network Encryption ... 53
NetWitness Platform Web Server Communications ... 53
Reporting Engine, ESA and Warehouse Connector : External Communication ... 53
Log Collector Service ... 53
Enabling HTTPS on REST Interfaces for Core Services ... 55
Data Security Settings ... 56
Securing Data ... 56
Data Privacy ... 56
Default Storage Passwords ... 57
Alert System Settings ... 58
FIPS Compliance ... 59
NetWitness Platform Components working in FIPS mode ... 59

Common Criteria Compliance **60**
Disabling Unencrypted Ports For NetWitness Core Services 60
STIG Compliance **62**
Overview 62
STIG Limits Account Access 62
Other Security Considerations **63**
Changing the RabbitMQ Management Password for Windows Legacy Collectors 63
Hardening the NetWitness Platform Core service 64
Example: 64
NFS Access Controls 65
Secure Deployment and Usage Settings **67**
Security Controls Map **68**
Secure Enclave 68
Secure Deployment Guidelines 69
Firewall Rules **71**
DMZ to Corporate Network 71
Corporate Network to Site 71
Site to Site 72
Live CMS to DMZ 73
RSA Download Central to DMZ 73
External Email Server to DMZ 73
Syslog Server to Site 73
SNMP Server to Site 73
Secure Deployment Settings **75**
Secure Maintenance **76**
Security Patch Upgrade 76
Security Patch Management 76
Virus Scanning 76
Ongoing Monitoring and Auditing 77
Hardware Replacement 77
Physical Security Controls Recommendations **78**
Supporting Users **79**
Preventing Social Engineering Attacks 79
Confirming User Identities 79
Advice for Your Users 80
Appendix A: Customer Provided Certificates **81**
Appendix B: Reissue Certificates **83**
Introduction 83

CA Certificate Reissue 83
Service Certificate Reissue 83
Reissuing Service Certificate 84
When to Use the --host-key Argument 85
cert-reissue Arguments and Options for All Hosts 85
When to Use the Individual Host Arguments (--host-id <id>, --host-name <display-name>, --host-addr <ip/hostname>) 86
cert-reissue Arguments and Options for a Single Host 87
Reissuing Certificates for All Hosts Except Windows Legacy Collection (WLC) host 88
Running the Cert-Reissue Command for All Hosts 88
Running the Cert-Reissue Command for an Individual Host 88
Reissuing Certificates for a WLC Host 88
Successful Reissue Summary Report 89
Unsuccessful Reissue Summary Reports 89
Reissue Failed for Host and Aborted Command 89
Reissue Certificate Partially Executed 90
Appendix C. Troubleshooting Cert-Reissue Command 91
Argument Options Used for Troubleshooting 91
Problems and How to Troubleshoot Them 92

Security Configuration Guide Overview

This guide provides information about the security configuration settings and security best practices for NetWitness.

This guide applies to NetWitness version 11.0 or later. There will be periodic updates made to the content.

Anyone using this guide should possess experience as a network engineer, equivalent to at least that of a journeyman, and also have a strong understanding of network concepts and TCP/IP communications.

Security Configuration Settings

This topic describes information about various security configuration settings that are designed to help you securely operate NetWitness.

You can adjust the following security configuration settings:

- Access Control Settings
- Log Settings
- Communication Security Settings
- Data Security Settings
- Alert System Settings
- Other Security Considerations

Access Control Settings

Access control settings are designed to enable the protection of resources against unauthorized access or by external components.

- User Authentication
- User Authorization
- Component Authentication

User Authentication

User authentication settings are designed to control the process of verifying an identity claimed by a user for accessing NetWitness.

Multi-Factor Authentication

You can set up Multi-Factor Authentication (MFA) for NetWitness using one of the following methods:

- ADFS Log in to NetWitness with SecurID Passcode
- PAM SecurID Log in to NetWitness for AD Users

For more information, see "Set Up Multi-Factor Authentication" and "Set Up Single Sign-On Authentication" in the *System Security and User Management Guide*.

NetWitness Platform Core Trusted Connection

NetWitness Core has the ability to connect and authenticate over SSL without having to provide user account information on the service itself. This feature is only available over the native port and not the REST interface. For more information on trusted connection, see *Host and Services Getting Started Guide*.Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

User Accounts

The following table identifies the default NetWitness user roles including the Administrator (admin) account and several service accounts. When deploying, you must enter a password for the System Administrator account and all the service accounts. For more information on passwords and password strength, see "Settings Tab" in the *System Security and User Management Guide*.

Note: Custom roles can be added as required. For instructions, see "Add a Service User Role" in the *Host and Services Getting Started Guide*.

User Roles	Description
Administrators	Full system access
Operators	Access to configurations but not to metadata and session content.
Analysts	Access to metadata and session content but not to configurations.
SOC_Managers	Same access as Analysts plus additional permission to handle incidents.
Malware_ Analysts	Access to malware events and to metadata and session content.
Data_Privacy_ Officers	Access to metadata and session content as well as configuration options that manage obfuscation and viewing of sensitive data within the system (see Data Privacy Management).

User Roles	Description
Respond_ Administrator	Access to all Respond server and Incidents permissions.
UEBA_Analysts	Access to the User and Entity Behavior Analytics (UEBA) service in the **Investigate** > **Users** view. UEBA is an advanced analytics solution for discovering, investigating, and monitoring risky behaviors across all entities in your network environment. **Note:** You do not need to set up specific permissions for this role. You only need to assign this role to a user, and that user will have access to UEBA.

Configuring New Accounts

Each NetWitness user must have an account to log on to the UI. For more information on how to add new user accounts, see "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

Caution: NetWitness recommends that you ensure that users are approved by the company for logging on to the system before creating an account for them. Even if users are approved, NetWitness recommends that you only assign the minimum set of access permissions that enable the users to perform their jobs.

Authentication Configuration

User authentication settings are designed to control the process of verifying an identity claimed by a user for accessing NetWitness. For more information, see "Set Up System Security" in the *System Security and User Management Guide*.

Below are recommendations for some of the configurations:

- **Default System Administrator Account**: NetWitness recommends that you instruct NetWitness administrators on your corporate IT policy and security best practices to generate and manage passwords for the default System Administrator account. NetWitness recommends that you change the default System Administrator password and the admin passwords for the service accounts per your company's password policy. For more information on password strength settings, see "Password Strength" in the *System Security and User Management Guide*. You should change the System Administrator password using the Admin user preferences. For instructions on how to change the password, see "Change the Default admin Passwords" in the *System Security and User Management Guide*.
- **External Authentication**: NetWitness supports external authentication. For more information, see "Configure External Authentication" in the *System Security and User Management Guide*.

User Passwords

NetWitness Users

Administrators can set the appropriate level of password strength for the user and can force users to change their passwords when password strength policy changes. Administrators can specify the global default user password expiration period and the notification period for the password expiry. For more information, see "Configure System-Level Security Settings" in the *System Security and User Management Guide*.

The following table shows the default security parameters settings for passwords.

Caution: NetWitness recommends that you change these settings in accordance with your corporate policy. Users must ensure the idle period and session time-out is specified.

Parameter	Default Setting
Global Default User Password Expiration Period	0
Notify User <n> Days Prior to Password Expiry	5

NetWitness Core Service Users

Administrators can change the password of a service user and replicate the new password to all the NetWitness Core services with the defined user account. For more information, see "Change a Service User Password" in the *Host and Services Getting Started Guide*.

You can also change your password from the Preferences panel in the Profile view. For more information, see "User Preferences" in the *NetWitness Getting Started Guide*.

Security Parameter Settings

The following table shows the default security parameters settings.

Caution: NetWitness recommends that you change these settings in accordance with your corporate policy.

Parameter	Default Setting
Lockout Period	20 minutes
Idle Period	10 minutes
Session Timeout	480 minutes
Max Login Failures	5

For more information on security parameter settings, see "Configure System-Level Security Settings" in the *System Security and User Management Guide*.

Kibana Health and Wellness Interface

You can change the default password of the Kibana UI. For more information, see the "Changing Kibana Password" in the *System Maintenance Guide*.

User Authorization

User authorization settings are designed to control rights or permissions that are granted to a user for accessing a resource managed by NetWitness.

Access Roles

NetWitness allows you to create access roles that you can assign to users. Each access role is mapped to a list of user authorization settings.

For more information, see the following NetWitness topics:

- "Role Permissions" in the *Alerting Using ESA Guide*.
- "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

Note: NetWitness recommends that you review users' task permissions on a routine basis to ensure that each user is granted the correct task permissions.

NetWitness allows access roles to be assigned to users through external group membership or directly to user accounts. NetWitness recommends that you assign permissions through group membership and not assign permissions directly to user accounts. For more information, see "Map User Roles to External Groups" in the *System Security and User Management Guide*.

The user roles assigned also control permissions that are granted to accounts that need access to a specific component of NetWitness.

Component Authentication

This topic describes how component authentication settings control the process of verifying an identity claimed by an external or internal system or component.

Host Configuration and Service Authentication

When you install or upgrade to NetWitness 11.x, trusted connections are established by default with two settings:

1. SSL is enabled.
2. NetWitness is connected to core services using the encrypted SSL port.

NetWitness allows secure authentication services for the following hosts as SSL is enabled by default:

- NetWitness Server
- Decoder
- Log Decoder
- Concentrator
- Broker
- Log Collector
- Archiver
- ESA
- Malware Analysis
- Endpoint
- UEBA

Note: By default all the services on the hosts have SSL enabled.

For more information, see the *Host and Service Getting Started Guide*.

Changing Credentials for Default Configuration Service Accounts

For information on how to reset the password for an admin of the host service accounts, see "Users Tab" in the *Host and Services Configuration Guide*.

Note: The default user name of the host service accounts (admin) cannot be modified.

Configuring Live Account Authentication

NetWitness supports secure authentication for the Live account connection to the Content Management System (CMS) as the SSL is enabled by default. The default communications port on the CMS is 443. For information on how to configure this setting, see "Configure Live Settings" in the *System Configuration Guide*.

Configuring Lockbox Authentication

Lockbox provides an encrypted file that Warehouse Connector or Log Collector uses to store and protect sensitive data. You need to create the lockbox by providing a lockbox password while configuring the Warehouse Connector or Log Collector for the first time. For more information on lockbox setup, see the following topics:

- "Log Collector - Set Up a Lockbox" in the *Log Collection Guide*.
- "Warehouse Connector - Create Lockbox" in the *Host and Services Getting Started Guide*.

Change Root Account Password On All NW Hosts

You must change the default SSH root account password to a strong password in all the hosts in the NetWitness Deployment.

Display Logon Banner for Remote SSH Connections

NetWitness allows you to customize the logon banner to display standard government or corporate warning signs for SSH remote connections to the hosts.

For example:

"This system is private. Use or misuse may be logged and invalid access pursued."

1. Log on to the appliance using root credentials.
2. Type **cd /etc/** to switch to the **/etc/** directory.
3. Edit the **/etc/issue.net** file with the required banner text.
4. Save the changes and exit.
5. Type **cd /etc/ssh** to switch to the **/etc/ssh** directory.
6. Edit the **/etc/ssh/sshd_config** file to remove the comment for the banner and provide the location of the banner text file (For example, **/etc/issue.net**).
   The following file is an example of an **sshd.config** file before being modified:

   ```
   # no default banner path
   #Banner none
   ```

   The following file is an example of an **sshd.config** file after being modified:

   ```
   # no default banner
   #Banner /etc/issue.net
   ```
7. Save the changes and exit.
8. Type **service sshd restart** in order to restart the sshd service.

Secure Boot Loader

A boot loader password is set to prevent unauthorized modification of boot menu entries. Change the default password to a strong password.

To change the boot loader password:

1. Run the grub2-setpassword command as root.

   ```
   # grub2-setpassword
   ```

2. Enter and confirm the password:

   ```
   Enter password:
   Confirm password:
   ```

Disable Interactive Startup

To prevent users from starting up the system interactively, as root, disable the **PROMPT** parameter in the file:

```
"/etc/sysconfig/init"
PROMPT= no
```

Create a Customized Logon Banner for NetWitness Platform

NetWitness allows you to customize and create a login banner that is displayed when users log on to NetWitness. For more information, see "Create a Customized Login Banner" in the *System Security and User Management Guide*.

Log Settings

A log is a chronological record of system activities that enables the reconstruction and examination of the sequence of environments and activities surrounding or leading to an operation, procedure, or event in a security-relevant transaction.

Global Audit Logging provides NetWitness Auditors with consolidated visibility into user activities within NetWitness in real-time from one centralized location. NetWitness audit logs are collected in a centralized system that converts them into the required format and forwards them to an external syslog system. The external syslog system can be a third-party syslog server or a Log Decoder. For more information, see "Global Audit Logging Overview" in the *System Configuration Guide*.

Log Description

The following table shows the security-relevant logs provided by NetWitness.

Component	Reference
Appliance and Service Logs	See the "Services Explore View" and "Services Logs View" in the *Host and Services Getting Started Guide* and "Configure Log File Settings" in the *System Configuration Guide*.
Audit Logs	See "Configure Global Audit Logging" in *System Configuration Guide*.
Syslogs	See "Configure Syslog and SNMP Settings" in the *System Configuration Guide*.

Log Management and Retrieval

For more information on:

- Log settings, see "Configure Log File Settings" in the *System Configuration Guide*.

  Note: NetWitness recommends that you set the maximum log file size in accordance to your corporate policy.

- Log forwarding, see "Set Syslog Forwarding" in the *Host and Services Getting Started Guide*.
- Setting log overrides:
  You may override the default logging levels if you want to include messages generated by specific modules.

  `Syntax: <module>=<level>`

  `SDK-Language=none`

  Where level is one or more of "`none|debug|info|warning|failure|audit|all`", all options must be separated by a pipe |

  `none`

  and

  `all`

are mutually exclusive with each other and all other options.

Overrides are useful for query auditing (that is, those modules that begin with SDK-) or for debugging by module (that is, Index). The following are the type of logs available :

- Data
- Engine
- Index
- Network
- Packet
- Parse
- Decoder
- Rules
- Concentrator
- Appliance
- SDK
- SDK-Query
- SDK-Values
- SDK-Language
- SDK-Info
- SDK-Session
- SDK-Timeline
- SDK-Content
- SDK-Search

Note: NetWitness recommends that you restrict permissions to the log files folder to the appropriate user.

Communication Security Settings

Communication security settings are designed to enable the establishment of secure communication channels between NetWitness components, as well as between NetWitness components and external systems or components.

Port Usage

To help ensure security, NetWitness recommends that you configure your firewall rules and access control lists to expose only the ports and protocols necessary for the operation of NetWitness. The services, such as Reporting Engine, Respond Service, Malware, Log Collector, Live account, Broker, Concentrator, Decoder, and Log Decoder, use specific TCP ports to communicate with each other and the following:

- Web user client interfaces
- Live CMS
- LDAP synchronization
- Third-party email server
- NetWitness console

All communication from NetWitness uses the native NetWitness Core ports. The additional native NetWitness Core port per host allows an administrator to enable secure (SSL) network communications while still being able to use non-secure (HTTP and NetWitness Core native) connectivity methods for communication between services that are present on the same host system. Administrators can toggle the ports on and off to support only SSL, only non-SSL, or both.

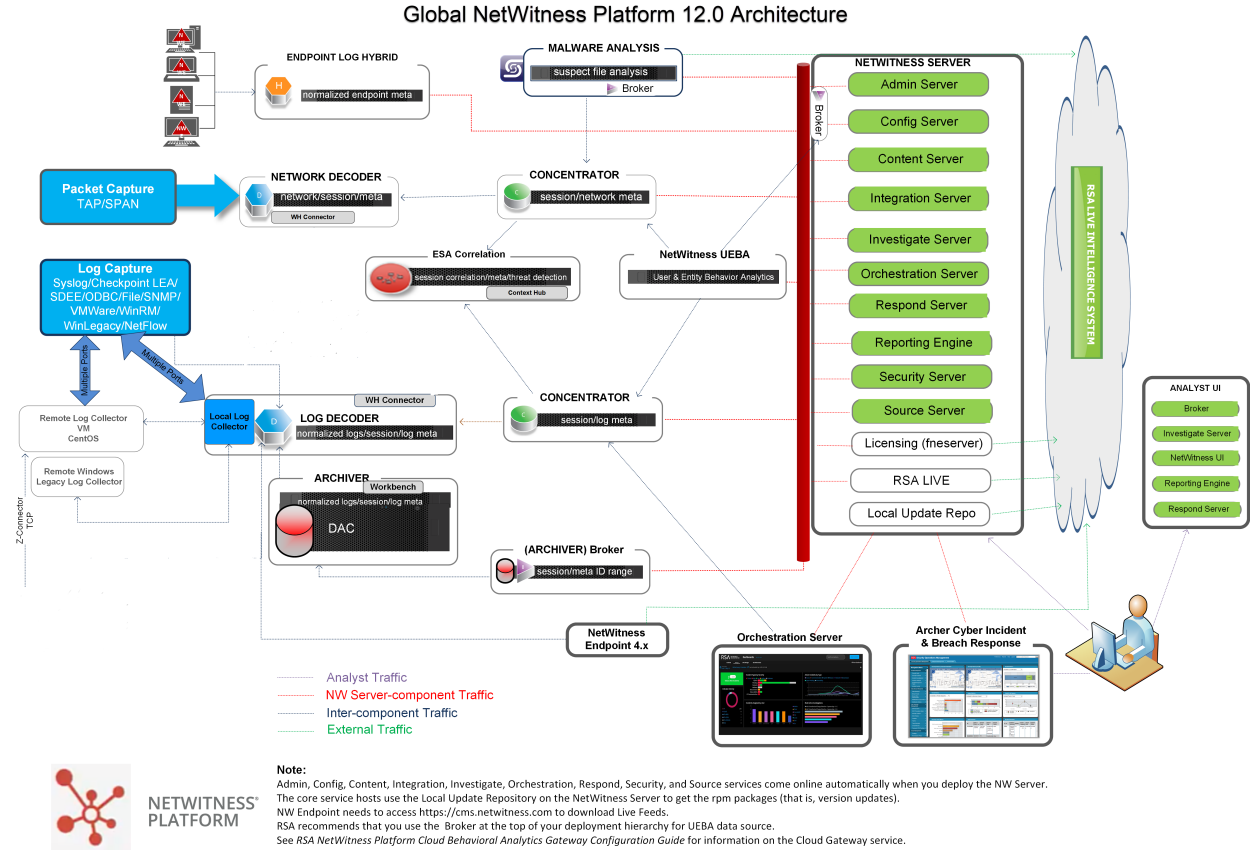
Refer to the following diagram and port table to ensure that all the relevant ports are opened for components in your NetWitness deployment to communicate with each other.

For more information on individual Endpoint Architectural diagrams, see NetWitness Endpoint Architecture at the end of this topic.

Note: enVision Local Collector is not more supported in versions 11.x. Warehouse is supported from version 11.5 and earlier versions.

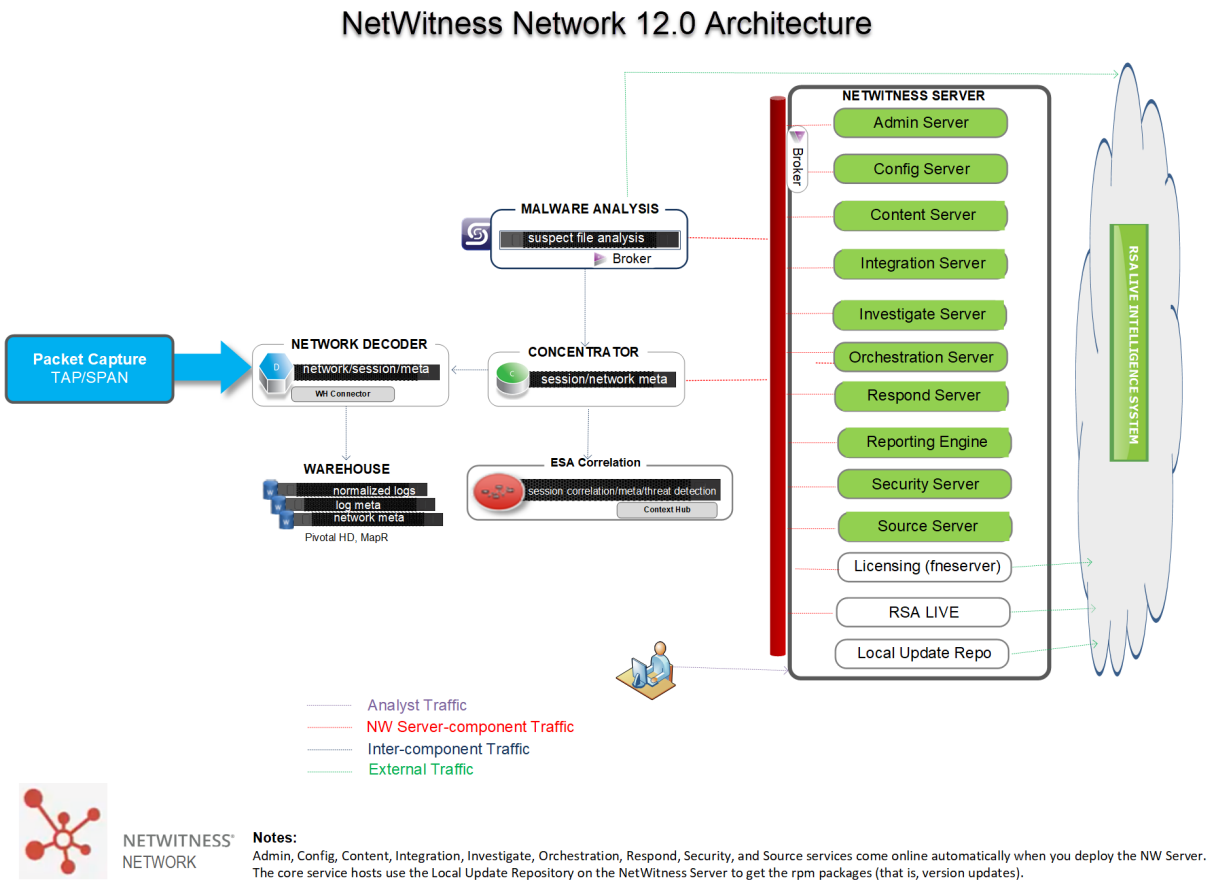
NetWitness Network Architecture Diagram

The following diagram illustrates the NetWitness network architecture including all of its component products.

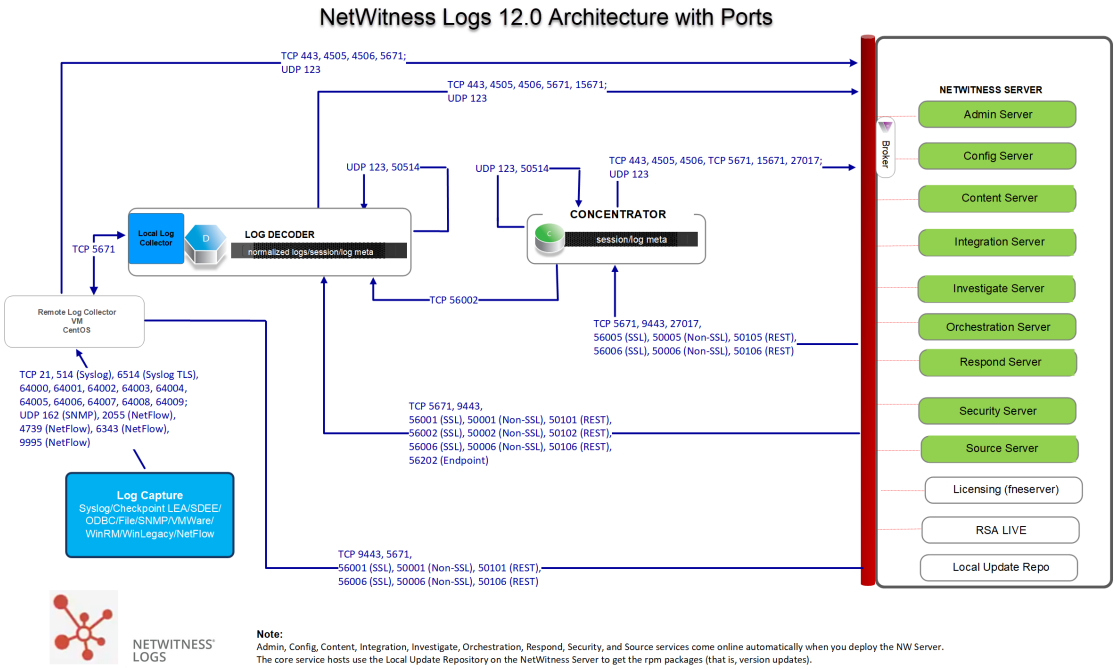


Note: NetWitness core hosts must be able to communicate with the Netwitness Server through UDP port 123 for Network Time Protocol (NTP) time synchronization.

NetWitness Network (Packets) Architecture Diagram with Ports

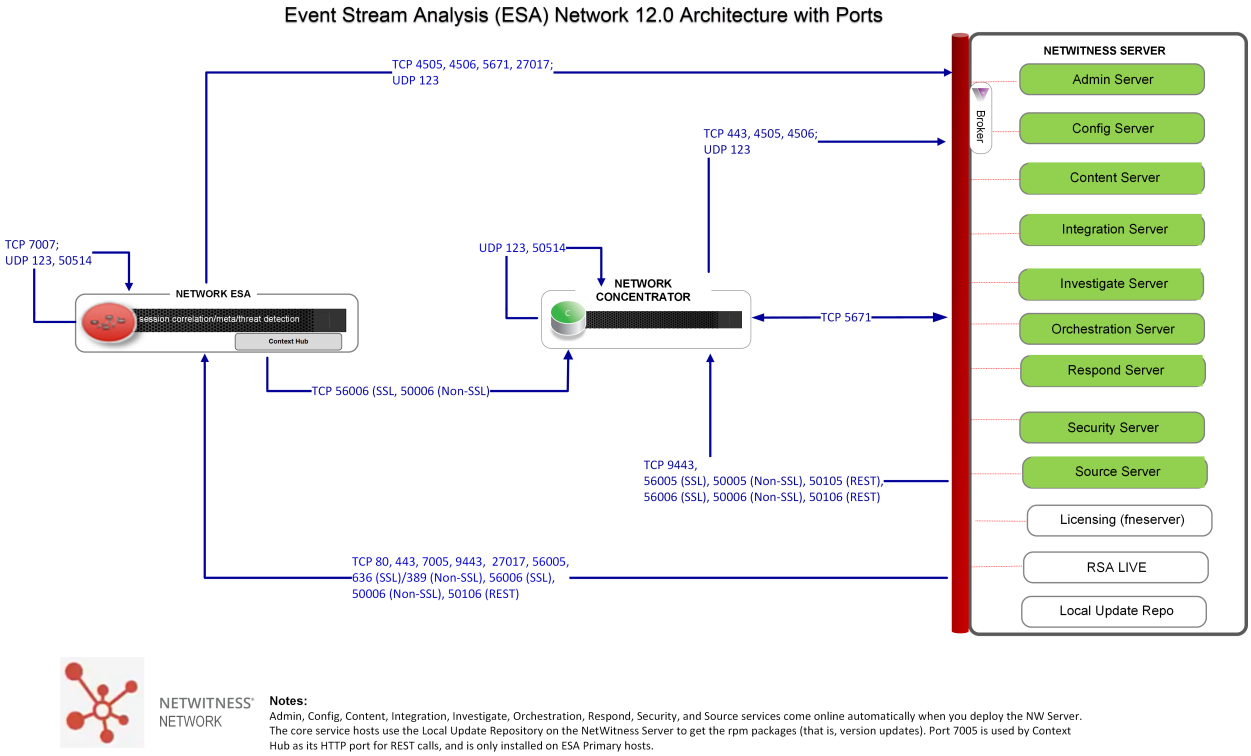


NetWitness Logs Architecture Diagram with Ports



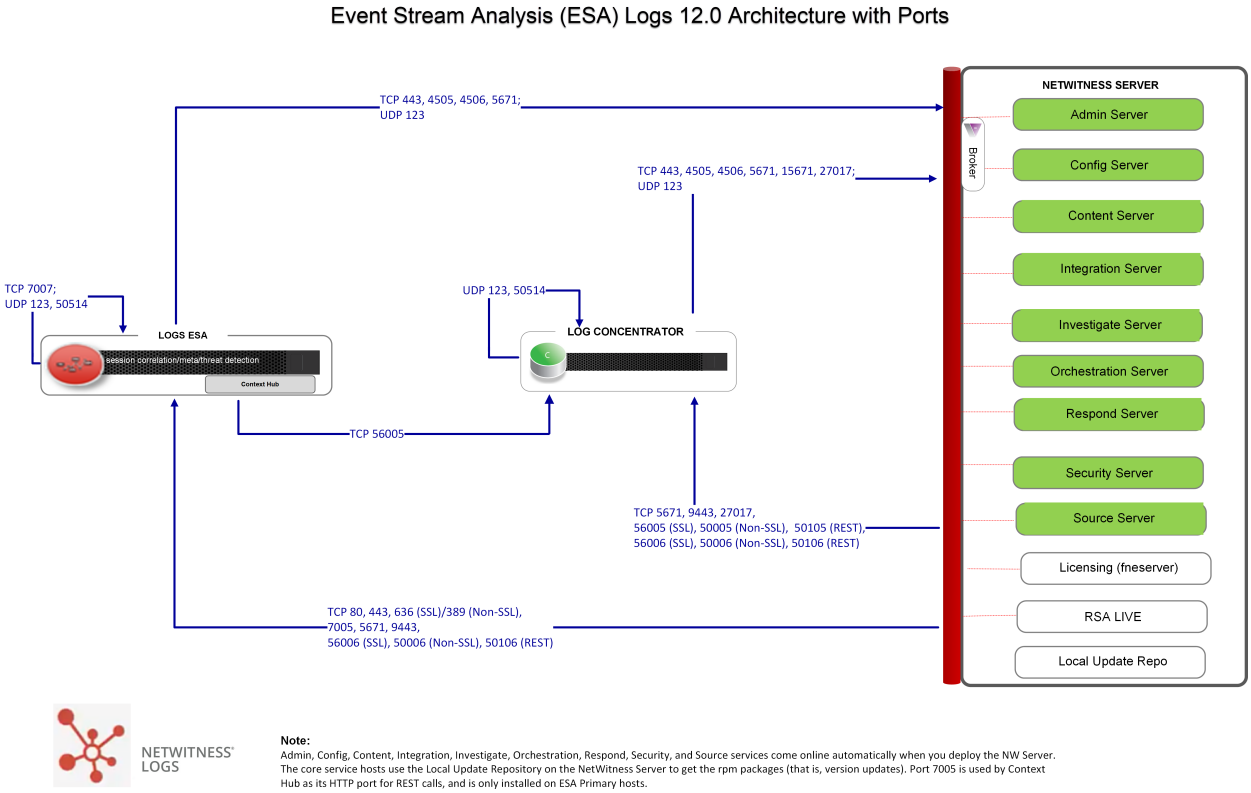
Event Stream Analysis Network (Packets) Architecture Diagram with Ports

The following diagram illustrates the Event Stream Analysis network architecture with packet capture.



Event Stream Analysis (Logs) Architecture Diagram with Ports

The following diagram illustrates the Event Stream Analysis network architecture with log collection.



NetWitness Firewall Requirements Summary

The following table lists all the ports that need to be open in your firewall by host.

Note: The "NW Server" host ports apply to both the Primary and Warm Standby NW Server. Synchronization between the Primary and Warm Stanby is done through TCP Port 22.

Source Host	Destination Host	Ports
NW Server	ESA Primary	**TCP:** 22, 80, 443, 5671, 7005, 50030 (SSL), 50035 (SSL), 50036 (SSL) **UDP:**123
NW Server	ESA	**TCP:** 22, 80, 443, 5671, 50035 (SSL), 50036 (SSL) **UDP**: 123
NW Server	Network Decoder	**TCP:** 22, 5671, 50004 (Non-SSL), 50006 (Non-SSL), 50104 (REST), 50106 (REST), 56004 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Broker	**TCP:** 5671, 50003 (Non-SSL), 50006 (Non-SSL), 50103 (REST), 50106 (REST) 56003 (SSL), 56006 (SSL) **UDP:**123
NW Server	Concentrator (Network & Logs)	**TCP:** 22, 5671, 50005 (Non-SSL), 60006 (Non-SSL), 50105 (REST), 50106 (REST), 56005 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Network Hybrid	**TCP:** 22, 5671, 50004 (Non-SSL), 50005 (Non-SSL), 50006 (Non-SSL), 50104 (REST), 50105 (REST), 50106 (REST), 56004 (SSL), 56005 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Log Decoder	**TCP:** 22, 5671, 50001 (Non-SSL), 50002 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Log Hybrid	**TCP:** 22, 5671, 50001 (Non-SSL), 50002 (Non-SSL), 50005 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50105 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56005 (SSL), 56006 (SSL) **UDP:** 123

Source Host	Destination Host	Ports
NW Server	Log Hybrid - Retention	**TCP:** 22, 5671, 50001 (Non-SSL), 50002 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50105 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Endpoint Log Hybrid	**TCP:** 22, 5671, 7050, 7054, 50001 (Non-SSL), 50002 (Non-SSL), 50005 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50102 (REST), 50105 (REST), 50106 (REST), 56001 (SSL), 56002 (SSL), 56005 (SSL), 56006 (SSL), 56202 (Endpoint) **UDP:** 123
NW Server	VLC	**TCP:** 22, 5671, 50001 (Non-SSL), 50006 (Non-SSL), 50101 (REST), 50106 (REST), 56001 (SSL), 56006 (SSL) **UDP:** 123
NW Server	Archiver	**TCP:** 22, 514, 5671, 6514, 50006(Non-SSL), 50007 (Non-SSL), 50008 (Non-SSL), 50106 (REST), 50107 (REST), 50108 (REST), 56006 (SSL), 56007 (SSL), 56008 (SSL) **UDP:** 123, 514
NW Server	Malware	**TCP:** 22, 5671, 5432, 50003 (Non-SSL), 50006 (Non-SSL), 50103 (REST), 50106 (REST), 56003 (SSL), 56006 (SSL), 60007 **UDP:** 123
NW Server	UEBA	**TCP:** 22 **UDP:** 123
ESA	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017 **UDP:** 123, 53
ESA	Active Directory	**TCP:** 389 (Non-SSL), 636 (SSL)
ESA	Archer	**TCP:** 80 (Non-SSL), 443 (SSL),
ESA Secondary	ESA Primary	**TCP:** 27017
ESA Primary or Secondary	Concentrator	**TCP:** 50005 (Non-SSL), 56005 (SSL)
Network Decoder	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017, **UDP:** 53, 123

Source Host	Destination Host	Ports
Concentrator (Network & Logs)	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017 **UDP:** 53, 123
Network Hybrid	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017) **UDP:** 53, 123
Log Decoder	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017 **UDP:** 53, 123
Log Hybrid	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671,15671, 27017 **UDP:** 53, 123
Log Hybrid - Retention	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671,15671, 27017 **UDP:** 53, 123
VLC	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671,15671, 27017 **UDP:** 53, 123
VLC	Log Collector	**TCP:** 5671
Log Collector	VLC	**TCP:** 5671
Endpoint Log Hybrid	NW Server	**TCP:** 53, 80, 443, 5671, 4505, 4506, 15671, 27017 **UDP:** 53, 123
Endpoint Log Hybrid	Log Decoder	**TCP:** 50202 (Non-SSL), 50102 (REST), 56202 (SSL) **UDP:** 514
Endpoint Agent	Log Decoder	**TCP:** 514, 6514 **UDP:** 514
Endpoint Agent	Endpoint Log Hybrid	**TCP:** 443 **UDP:** 444
UEBA	NW Server	**TCP:** 53, 80, 443, 4505, 4506, 5671, 15671, 27017, 50003 (Broker-Non-SSL), 50103 (Broker/REST), 56003 (Broker/SSL) **UDP:** 53, 123
UEBA	Concentrator	**TCP:** 50005 (Non-SSL), 50105 (REST), 56005 (SSL)

www connections

Source Host	Destination Host	Ports
NW Server	cloud.netwitness.com cms.netwitness.com download.rsasecurity.com panacea.threatgrid.com quantum.subscribenet.com rsasecurity.subscribenet.com smcupdate.emc.com	**TCP:** 80, 443
ESA (Primary & Secondary)	cloud.netwitness.com cms.netwitness.com download.rsasecurity.com panacea.threatgrid.com quantum.subscribenet.com rsasecurity.subscribenet.com smcupdate.emc.com	**TCP:** 80, 443
Malware	panacea.threatgrid.com cloud.netwitness.com	**TCP:** 443

Comprehensive List of NetWitness Host, Service, and iDRAC Ports

Note: For ports used in event collection through the NetWitness Logs, see the "The Basics" in the *NetWitness Platform Log Collection Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

This section contains the port specifications for the following hosts.

NW Server Host	iDRAC Ports
Analyst UI Host	Log Collector Host
Archiver Host	Log Decoder Host
Broker Host	Log Hybrid Host
Concentrator Host	Log Hybrid - Retention
Endpoint Log Hybrid Host	Malware Host
Endpoint Relay Server	Network Decoder Host
Event Stream Analysis Host	Network Hybrid Host
New Health & Wellness	UEBA Host

NW Server Host (Primary and Warm Standby NW Server Host)

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	NW Server	TCP 443, 80	nginx - NetWitness UI
Admin Workstation	NW Server	TCP 15671	RabbitMQ Management UI
Admin Workstation	NW Server	TCP 22	SSH Primary to Standby NW Server sy nchronization port.
NW Hosts	NW Server	TCP 53 UDP 53	DNS
NW Hosts	NW Server	TCP 15671	RabbitMQ Management UI
NW Hosts	NW Server	TCP 4505, 4506	Salt Master Ports
NW Hosts	NW Server	TCP 443	RSA Update Repository
NW Hosts	NW Server	TCP 5671	RabbitMQ-amqp
NW Hosts	NW Server	UDP 123	NTP
NW Hosts	NW Server	TCP 27017	MongoDB
NW Server	cloud.netwitness.com	TCP 443	Live
NW Server	cms.netwitness.com	TCP 443	Live
NW Server	smcupdate.emc.com	TCP 443	Live
NW Server	NFS Server	TCP 111, 2049, UDP 111, 2049	iDRAC Installations
NW Server	NW Hosts	UDP 123	NTP
NW Server	NW Endpoint	TCP 443, 9443	For NW Endpoint 4.x integrations

Analyst UI Host

Source Host	Destination Host	Destination Ports	Comments
Analyst UI	NW Server	TCP 7006	The Content Server is listening on this port.
Analyst UI	NW Server	TCP 7009	The Admin Server is listening on this port.
Analyst UI	NW Server	TCP 7012	The Integration Server is listening on this port.
Analyst UI	NW Server	TCP 7015	The Source Server is listening on this port.
Analyst UI	NW Server	TCP 7016	The License Server is listening on this port.
NW Hosts	Analyst UI	TCP 5671	RabbitMQ-amqp
Analyst UI	NW Server	UDP 123	NTP

Note: Make sure you restrict DNS access on the Netwitness Server. You must allow access to DNS service on NetWitness Server only from NeWitness Hosts by adding the firewall rules in NetWitness Server.

Archiver Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Archiver	TCP 15671	RabbitMQ Management UI
Archiver	NW Server	TCP 15671	RabbitMQ Management UI
Archiver	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Archiver	TCP 22	SSH
NW Server	Archiver	TCP 50008 (Non-SSL), 56008 (SSL), 50108 (REST)	Archiver Application Ports
NW Server	Archiver	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Archiver	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
NW Server	Archiver	TCP 50007 (Non-SSL), 56007 (SSL), 50107 (REST)	Workbench Application Ports
Archiver	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations

Broker Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Broker	TCP 15671	RabbitMQ Management UI
Broker	Concentrator	TCP 50005 (Non-SSL), 56005	Concentrator Application Port
Broker	NW Server	TCP 15671	RabbitMQ Management UI
Broker	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Broker	TCP 22	SSH
NW Server	Broker	TCP 50003 (Non-SSL), 56003 (SSL), 50103 (REST)	Broker Application Ports
NW Server	Broker	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Broker	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Broker	NW Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Endpoint Broker	NW Server	TCP 443	RSA Update Repository

Concentrator Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Concentrator	TCP 15671	RabbitMQ Management UI
Concentrator	Log Decoder	TCP 50002 (Non-SSL), 56002 (SSL)	Log Decoder Application Port
Concentrator	Network Decoder	TCP 56004, 50004 (Non-SSL)	Network Application Port
Concentrator	NW Server	TCP 15671	RabbitMQ Management UI
Concentrator	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Concentrator	TCP 22	SSH
NW Server	Concentrator	TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports
Malware	Concentrator	TCP TCP 50005 (Non-SSL), 56005 (SSL)	Malware
NW Server	Concentrator	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Concentrator	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Concentrator	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations

Endpoint Log Hybrid

Source Host	Destination Host	Destination Ports	Comments
Endpoint Agent	Endpoint Log Hybrid	TCP 443 UDP 444	NGINX HTTPS NGINX UDP. If UDP port 444 is not acceptable in your environment, see How to Change UDP Port for Endpoint Log Hybrid.
Endpoint Agent	Log Decoder or Virtual Log Collector	TCP 514 (Syslog) UDP 514 (Syslog) TLS 6514	Windows Log Collection
Endpoint Log Hybrid	Log Decoder (External)	TCP 50102 (REST) 56202 (Protobuf SSL) 50202 (Protobuf)	To forward meta to an external Log Decoder
Endpoint Log Hybrid	NW Server	TCP 443	RSA Update Repository
NW Server	Endpoint Log Hybrid	TCP 7050	UI web traffic
Endpoint Log Hybrid	NW Server	TCP 5671	Message Bus
Endpoint Log Hybrid	NW Server	TCP 27017	MongoDB
NW Server	Endpoint Log Hybrid	TCP 7054	UI web traffic
NW Server	NFS Server	TCP 111, 2049 UDP 111, 2049	iDRAC Installations
NW Server	Endpoint Log Hybrid	TCP 50001 (Non-SSL), 56001 (SSL), 50101(REST)	Log Collector application ports
NW Server	Endpoint Log Hybrid	TCP 50002 (Non-SSL), 56002 (SSL), 56202 (Endpoint), 50102 (REST)	Log Decoder application ports
Admin Workstation	Endpoint Log Hybrid	TCP 15671	RabbitMQ Management UI
Endpoint Log Hybrid	NW Server	TCP 15671	RabbitMQ Management UI

Endpoint Relay Server

Source Host	Destination Host	Destination Ports	Comments
Endpoint Agent	Relay Server	TCP 443	To forward host data to the Relay Server
Endpoint Log Hybrid	Relay Server	TCP 443	Pull host data from the Relay Server

Event Stream Analysis (ESA) Host

Note: The ports in this table are for the ESA Primary an ESA Secondary hosts. The Content Hub, Correlation and ESA Analytics services are co-located on the ESA Primary host. The Correlation and ESA Analytics services are co-located on the ESA Secondary host.

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	ESA	TCP 15671	RabbitMQ Management UI
ESA Primary and Secondary	NW Server	TCP 15671	RabbitMQ Management UI
ESA Primary and Secondary	NW Server	TCP 443	RSA Update Repository
Admin Workstation	ESA	TCP 22	SSH
NW Server, ESA Secondary	ESA Primary	TCP 27017	MongoDB
NW Server	ESA Primary	TCP 7005	Context Hub Launch Port - (ESA Primary)
NW Server	ESA	TCP 50030 (SSL)	ESA Application Port
NW Server	ESA	TCP 50035 (SSL)	ESA Application Port
NW Server	ESA	TCP 50036 (SSL)	ESA Application Port
NW Server	ESA	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
ESA Primary and Secondary	cms.netwitness.com	TCP 443	Live
ESA Primary and Secondary	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
ESA Primary and Secondary	Active Directory	636 (SSL)/389 (Non-SSL)	
NW Server	ESA	80 (HTTP)/ 443 (HTTPS)(REST)	
ESA Primary	Archer	443 (SSL)/80 (Non-SSL)	
ESA Primary	ESA Primary	TCP 7007	Launch Port

New Health and Wellness

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	New Health and Wellness	TCP 22	SSH
Admin Workstation	New Health and Wellness	TCP 5601	Kibana UI
NW Hosts	New Health and Wellness	TCP 9200	Elasticsearch REST API Port
NW Server	New Health and Wellness	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
NW Server	New Health and Wellness	TCP 15671	RabbitMQ Management UI
NW Server	New Health and Wellness	TCP 7018	Metrics Server Launch Port
NW Server	New Health and Wellness	TCP 7020	Node Infra Server Launch Port

New Health and Wellness on Different Subnet

If the New Health and Wellness is on a different subnet, you must open the respective NetWitness Platform services port.

Example:
New Health and Wellness is on subnet A: 10.10.1.0/24 and NetWitness Platform host LogHybrid is on subnet B: 10.10.2.0/24. In this case, you must open ports for Log Decoder, Log Collector, Concentrator on Metrics Server (New Health and Wellness host) to allow ports in the firewall for communication.

Source Host	Destination Host	Destination Ports	Comments
New Health and Wellness	Log Decoder	50002(Non-SSL),56002 (SSL)	Log Decoder Application Ports
New Health and Wellness	Log Collector	50001(Non-SSL),56001 (SSL)	Log Collector Application Ports
New Health and Wellness	Concentrator	50005(Non-SSL)/56005 (SSL)	Concentrator Application Ports

iDRAC Ports

Port	Function	Comments
22*	SSH	Default, configurable port through which iDRAC listens for connections
443*	HTTP	Default, configurable port through which iDRAC listens for connections
5900*	Virtual Console keyboard and mouse redirection, Virtual Media, Virtual Folders, and Remote File Share.	Default, configurable port through which iDRAC listens for connections
111, 2049	TCP	NetWitness Platform hosts to NFS Server
111, 2049	UDP	NetWitness Platform hosts to NFS Server

Log Collector Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Collector	TCP 15671	RabbitMQ Management UI
Log Collector	NW Server	TCP 15671	RabbitMQ Management UI
Log Collector	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Collector	TCP 22	SSH
Log Collector	Log Event Sources	See *Log Collection Configuration Guide*. (missing or bad snippet)	
Log Event Sources	Log Collector	TCP 514 (Syslog) UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)"	Log Collection Ports
Log Event Sources	Log Collector	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008,64009	Log Collection FTP/S Ports
NW Server	Log Collector	TCP 50001 (Non-SSL), 56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Collector	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Collector	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Collector	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC installations
Log Collector	Virtual Log Collector	TCP 5671	In Pull Mode
Virtual Log Collector	Log Collector	TCP 5671	In Push Mode

Log Decoder Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Decoder	TCP 15671	RabbitMQ Management UI
Log Decoder	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Decoder	TCP 22	SSH
Log Decoder	Log Event Sources	See *Log Collection Configuration Guide*. (missing or bad snippet)	
Log Event Sources	Log Decoder	TCP 514 (Syslog), UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)	Log Collection Ports
Log Event Sources	Log Decoder	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008, 64009	Log Collection FTP/S Ports
NW Server	Log Decoder	TCP 50001 (Non-SSL),56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Decoder	TCP 50002 (Non-SSL), 56002 (SSL),56202 (Endpoint), 50102 (REST)	Log Decoder Application Ports
NW Server	Log Decoder	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Decoder	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Decoder	Log Collector	TCP 6514	
Log Decoder	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations
Log Decoder	NW Server	TCP 15671	RabbitMQ Management UI

Log Hybrid Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Hybrid	TCP 15671	RabbitMQ Management UI
Log Hybrid	NW Server	TCP 15671	RabbitMQ Management UI
Log Hybrid	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Hybrid	TCP 22	SSH
Log Collector	Log Event Sources	See *Log Collection Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.	
Log Event Sources	Log Hybrid	TCP 514 (Syslog), UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)	Log Collection Ports
Log Event Sources	Log Hybrid	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008, 64009	Log Collection FTP/S Ports
NW Server	Log Hybrid	TCP 50001 (Non-SSL), 56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Hybrid	TCP 50002 (Non-SSL), 56002 (SSL), 56202 (Endpoint), 50102 (REST)	Log Decoder Application Ports
NW Server	Log Hybrid	TCP TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports
NW Server	Log Hybrid	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Hybrid	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Hybrid	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations

Log Hybrid - Retention Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Log Hybrid - Retention	TCP 15671	RabbitMQ Management UI
Log Hybrid - Retention	NW Server	TCP 15671	RabbitMQ Management UI
Log Hybrid - Retention	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Log Hybrid - Retention	TCP 22	SSH
Log Collector	Log Event Sources	See *Log Collection Configuration Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.	
Log Event Sources	Log Hybrid - Retention	TCP 514 (Syslog), UDP 162 (SNMP), 514 (Syslog), 2055 (NetFlow), 4739 (NetFlow), 6343 (NetFlow), 9995 (NetFlow)	Log Collection Ports
Log Event Sources	Log Hybrid - Retention	TCP 21, 64000, 64001, 64002, 64003, 64004, 64005, 64006, 64007, 64008, 64009	Log Collection FTP/S Ports
NW Server	Log Hybrid - Retention	TCP 50001 (Non-SSL), 56001 (SSL), 50101 (REST)	Log Collector Application Ports
NW Server	Log Hybrid - Retention	TCP 50002 (Non-SSL), 56002 (SSL), 56202 (Endpoint), 50102 (REST)	Log Decoder Application Ports
NW Server	Log Hybrid - Retention	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Log Hybrid - Retention	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Log Hybrid - Retention	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations

Malware Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Malware	TCP 15671	RabbitMQ Management UI
Malware	NW Server	TCP 15671	RabbitMQ Management UI
Malware	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Malware	TCP 22	SSH
NW Server	Malware	TCP 60007	Malware Application Ports
NW Server	Malware	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Malware	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
NW Server	Malware	TCP 5432	Postgresql
NW Server	Malware	TCP 56003 (SSL), 50103 (REST)	Broker Application Ports
Malware	panacea.threatgrid.com	TCP 443	Threatgrid
Malware	cloud.netwitness.com	TCP 443	Community evaluation / Opswat
Malware	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations

Network Decoder Host

Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Network Decoder	TCP 15671	RabbitMQ Management UI
Network Decoder	NW Server	TCP 15671	RabbitMQ Management UI
Network Decoder	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Network Decoder	TCP 22	SSH
NW Server	Network Decoder	TCP 56004 (SSL), 50104 (REST), 50004 (Non-SSL)	Network Decoder Application Ports
NW Server	Network Decoder	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Network Decoder	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Network Decoder	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations

Network Hybrid Host

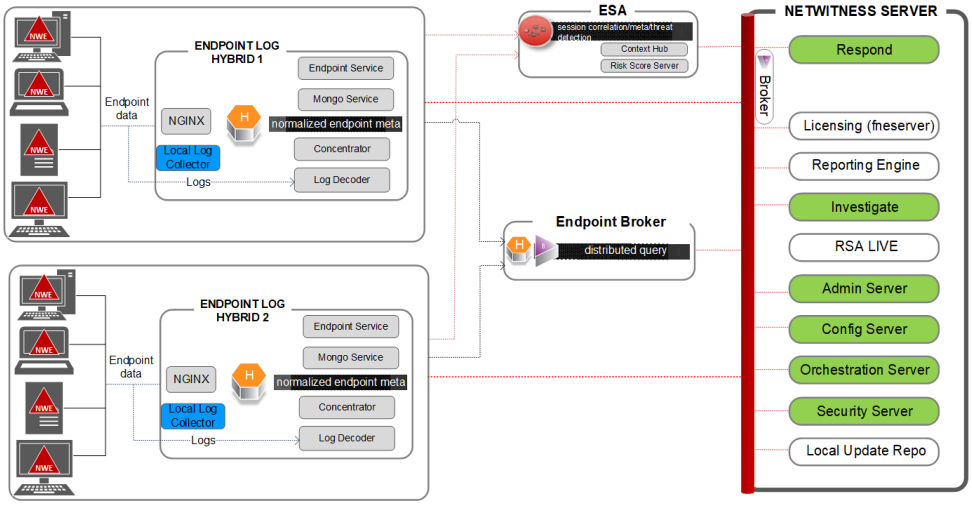
Source Host	Destination Host	Destination Ports	Comments
Admin Workstation	Network Hybrid	TCP 15671	RabbitMQ Management UI
Network Hybrid	NW Server	TCP 15671	RabbitMQ Management UI
Network Hybrid	NW Server	TCP 443	RSA Update Repository
Admin Workstation	Network Hybrid	TCP 22	SSH
NW Server	Network Hybrid	TCP 56004 (SSL), 50104 (REST), 50004 (Non-SSL)	Network Decoder Application Ports
NW Server	Network Hybrid	TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports
NW Server	Network Hybrid	TCP 56006 (SSL), 50106 (REST)	NetWitness Appliance Ports
NW Server	Network Hybrid	TCP 5671	RabbitMQ (AMQPS) message bus for all NW hosts.
Network Hybrid	NFS Server	TCP 111 2049 UDP 111 2049	iDRAC Installations

UEBA Host

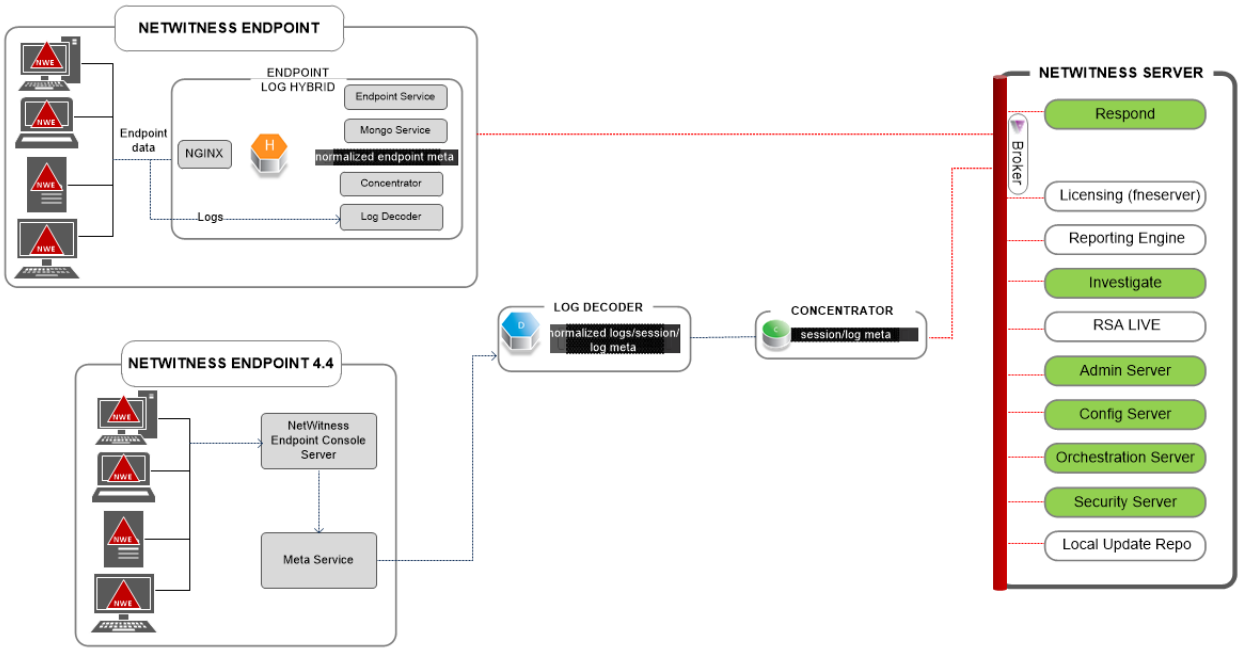
Source Host	Destination Host	Destination Ports	Comments
UEBA Server	NW Server	TCP 443	RSA Update Repository
UEBA Server	Broker	TCP 56003 (SSL), 50103 (REST)	Broker Application Ports
UEBA Server	Concentrator	TCP TCP 50005 (Non-SSL), 56005 (SSL), 50105 (REST)	Concentrator Application Ports
Admin Workstation	UEBA Server	443	UEBA Monitoring
Admin Workstation	UEBA Server	22	SSH
UEBA Server	NW Server	15671	UEBA Alerts forwarding to Respond
NW Server	NFS Server	TCP 111, 2049 UDP 111, 2049	iDRAC Installations

NetWitness Endpoint Architecture

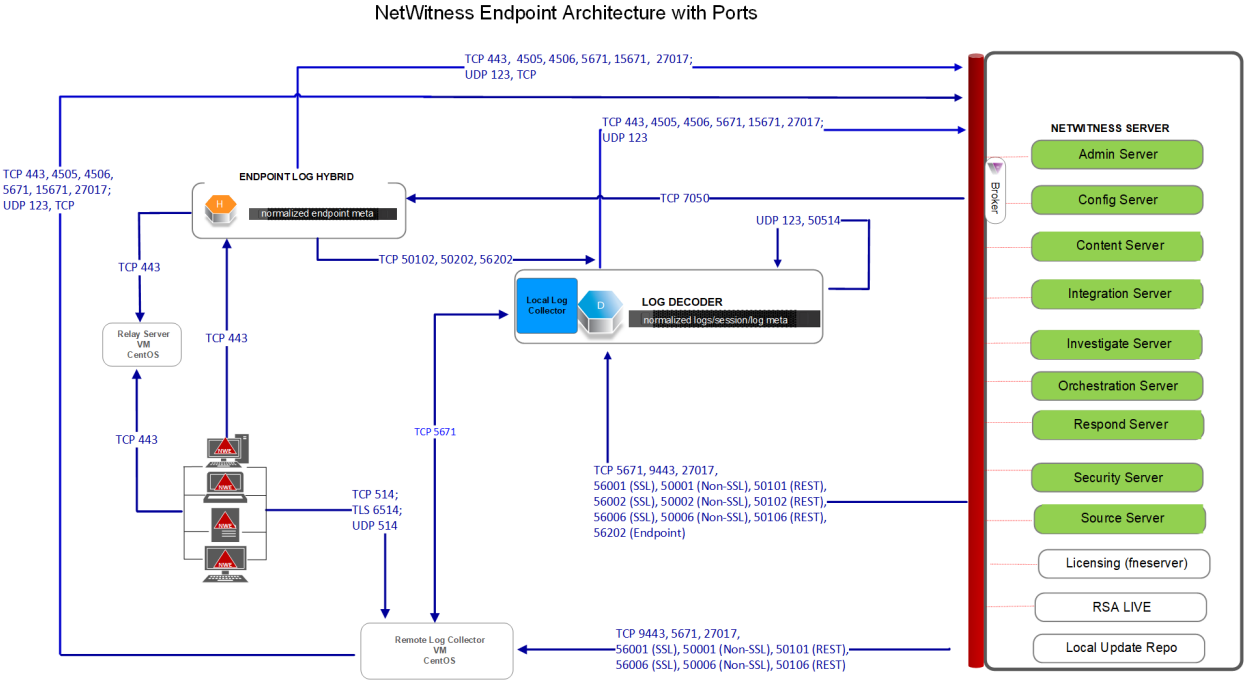
NetWitness Endpoint Architecture



NetWitness Endpoint 4.4 Integration with NetWitness Platform



NetWitness Endpoint Architecture with Ports



For more information on the services running on Endpoint Log Hybrid, see *NetWitness Endpoint Configuration Guide*.

How to Change UDP Port for Endpoint Log Hybrid

The following steps tell you how to change the Endpoint Log Hybrid default UDP port 444 if it is not acceptable in your environment. 555 is the example this procedures uses as a replacement for 444 UDP port.
There are two tasks you need to o change the Endpoint Log Hybrid default UDP port 444:

Task 1 - Tell All Agents to Use a New UDP Port

Task 2 - Update the Port on All Endpoint Log Hybrid Hosts in Your Environment

Note: If you did not select the custom firewall rules option when you ran the nwsetup-tui, NetWitness platform overwrites the firewall rules are overwritten after a period of time. Refer to the following Knowledge Base Article 00036446 (https://community.netwitness.com/t5/netwitness-knowledge-base/how-to-add-custom-firewall-rules-after-nwsetup-tui-has-completed/ta-p/5900) if this is the case.

Task 1 - Tell All Agents to Use a New UDP Port

Complete the following steps to update the UDP port in the default Enterprise Data Replication (EDR) policy, and all other policies you have, to tell all agents to use a new UDP port.

1. In the **NetWitness** menu, select **(missing or bad snippet)> Endpoint Sources** > **Policies**.
   The **Policies** view is displayed.
2. Select the **Default EDR Policy** and click **Edit** from the toolbar.
3. roll down to find the **UDP PORT** and change the value (for example, change from **444** to **555**).
4. Click **Publish Policy** at the bottom of the view.

Task 2 - Update the Port on All Endpoint Log Hybrid Hosts in Your Environment

SSH to each Endpoint Log Hybrid host in your environment with `admin` credentials and make the following updates.

1. Update the `iptables` rules to allow `555` in place of `444`.
   a. Replace `444` with `555` in the following file.
      ```
      vi /etc/sysconfig/iptables
      ```
   b. Restart `iptables` with the following command string.
      ```
      systemctl restart iptables
      ```
   c. Verify the change with the following command string.
      ```
      iptables -L -n
      ```
      The following is an example of what is displayed for a correct change.
      ```
      ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp multiport dports 555 /*
      EndpointNginxPort */ ctstate NEW
      ```

2. Update the SELinux policy. 555 is a privileged port, so you must update SELinux policy to allow this port.
   a. Run the following command string.
      `semanage port -a -t http_port_t -p udp 555`
      If you received any python errors or warnings, ignored them.
   b. Verify the change with the following command string.
      `semanage port -l | grep http_port_t`
      The following is an example of what is displayed for a correct change.
      `http_port_t udp 555, 444`
   c. (Optional) Remove `444`.
3. Update nginx config.
   a. Edit the following file.
      `vi /etc/nginx/nginx.conf`
   b. Search for the following string.
      `listen 444 udp;`
   c. Replace `444` with `555`.
   d. Restart nginx with the following command string.
      `systemctl restart nginx`
4. Verify that agents are communicating over the new port.
   a. Run the following command string.
      `tcpdump -i eth0 port 555`
   b. Wait for 30 seconds because the port sends out a beacon every 30 seconds. If t everything is working correctly, information similar to the following will be displayed.
      `09:20:12.571316 IP 10.40.15.103.60807 > EPS1.rsa.lab.emc.com.dsf: UDP, length 20`
      `09:20:12.572433 IP EPS1.rsa.lab.emc.com.dsf > 10.40.15.103.60807: UDP, length 1`
      Both lines must be returned. One is the size request (`20` bytes) and the other is the response size (`1` byte).

Note: For the latest architecture and port usage information, see "Network Architecture and Ports" topic in the *NetWitness Deployment Guide*.

Network Encryption

You can configure NetWitness Platform to send or receive data from external data sources.

Note: NetWitness recommends that whenever you have the option to choose between unsecured and secured versions of a communication protocol, you choose the secured version.

NetWitness Platform Web Server Communications

The NetWitness Platform UI or web server communicates with the Live Service (CMS) over port 443 using the HTTPS protocol.

Note: During installation, the system is engineered to set the default communication protocol to HTTPS over port 443.

Reporting Engine, ESA and Warehouse Connector : External Communication

NetWitness recommends that you use the secure TCP protocol and enable an SSL connection while configuring Reporting Engine, ESA, Warehouse Connector, Licensing, and Malware.

For more information on Reporting Engine, see "Configure Output Actions" in the *Host and Services Getting Started Guide*.

For more information on Malware external communication, see " Configure Malware Analysis Operating Environment" in the *Host and Services Getting Started Guide*.

For more information on ESA, see " Notification Methods" in the *Alerting Using ESA Guide*.

For more information on the Warehouse Connector, see "Configure Warehouse Connector" in the *Host and Services Getting Started Guide*.

For more information on Licensing, see "Configure NetWitness Platform Notifications" in the *Licensing Management Guide*.

Log Collector Service

To help secure communication between the Log Collector service running on the Log Decoder and the event sources, NetWitness recommends the following protocols.

Event Source	Protocol	Resources
File	SFTP, SCP, FTPS	For more information, see " File Collection Protocol Configuration" in the *Log Collection Guide*.

Event Source	Protocol	Resources
ODBC	ODBC	For more information on configuring an ODBC event source, see "ODBC Collection Configuration" in the *Log Collection Guide*. **Note:** Depending on the event source, administrators can configure additional progress driver parameter for secure connections. For more information, see the Progress document at https://www.progress.com/odbc/resources/documentation/books-and-readme-file. For more information on using a certificate, see the certificate creation kit at http://openssl.org/. For more information on securing communication with SQL Server, Oracle, and ODBC, see the URLs: http://technet.microsoft.com/en-us/l...QL.105%29.as http://technet.microsoft.com/en-us/l.../cc754431.aspx http://www.oracle.com/technetwork/database/options/advanced-security/overview/index.html http://www.psdn.progress.com/progres...92/odr/odr.pdf
Windows	HTTPS	For more information on configuring a Windows event source to use certificates and enable HTTPS, see the NetWitness Platformhelp topics in the *Windows Collection Configuration Guide*.
Check Point	OPSEC LEA	For more information on configuring a Check Point event source to use certificates, see the NetWitness Platform help topics in the *Check Point Collection Configuration Guide*.
Netflow	Netflow	For more information on configuring a Netflow event source to use certificates, see the NetWitness Platform help topics in the *Netflow Collection Configuration Guide*.
SDEE	SDEE	For more information on configuring a SDEE event source to use certificates, see the NetWitness Platform help topics in the *SDEE Collection Configuration Guide*.
SNMP	SNMP	For more information on configuring a SNMP event source to use certificates, see the NetWitness Platform help topics in the *SNMP Collection Configuration Guide*.
VMware		For more information on configuring a VMware event source to use certificates, see the NetWitness Platform help topics in the *VMware Collection Configuration Guide*.
Legacy Windows and NetApp		For more information on configuring a Legacy Windows event source to use certificates, see the NetWitness Platform help topics in the *Legacy Windows and NetApp Collection Configuration Guide*

Event Source	Protocol	Resources
Amazon Web Services (AWS) Cloud Trail	HTTPS	For more information on configuring an AWS Cloud Trail event source to use certificates, see the NetWitness Platform help topics in the *AWS (CloudTrail) Collection Configuration Guide*.

Note: For more information on enabling SSL for component communications, see Component Authentication.

Enabling HTTPS on REST Interfaces for Core Services

To enable HTTPS on REST interfaces:

1. Log in to REST interface.
2. Go to the **rest** > **config** node.
3. Set **SSL** config to **on**.
4. Restart the service.

Data Security Settings

Data security settings are designed to enable the definition of controls to prevent data permanently stored by NetWitness from being disclosed in an unauthorized manner.

Securing Data

To help protect online data, such as current database, log, and configuration files, NetWitness recommends that you restrict access to the files, database and configure permissions so that only trusted administrators are allowed to access them.

NetWitness recommends that you back up your sensitive data, encrypt it, and keep it in a secure physical location in accordance with your corporate disaster recovery and business continuity policies.

The backup can be done in the following ways:

- Regular backup of Configuration and Data files – You can back up and restore data and configuration files for the core host and services and all the modules of NetWitness. For more information, see the *NetWitness Recovery Tool User Guide*.
- Regular backup of critical configuration – You can export configurations using the Export option available on the UI. For example, you can take a backup of critical rules, reports, alerts, ESA rules, dashboards, investigation profiles, metadata groups, event sources, global notifications, and so on. For more information, see topics:
  - "Export a Rule, Export an Alert and Export a Report" in the *Reporting Guide*.
  - "Rule Library View and Dashboard" in the *Alerting using ESA Guide*.
  - "Manage Profiles Dialog and Export a Meta Group" in the *Investigation and Malware Analysis Guide*.
  - "Events View and Export Event Sources" in the *Event Source Management Guide*.
  - "Global Notifications Panel Toolbar" in the *System Configuration Guide*.

Data Privacy

Data Privacy is very integral and helps you manage privacy-sensitive data. You can achieve data privacy using the Data Privacy Officer (DPO) role. The DPO can configure NetWitness to limit the exposure of metadata and raw content (packets and logs) using a combination of techniques. The methods available to protect data in NetWitness include:

- Data Obfuscation
- Data Retention Enforcement
- Auditing Logging

For more information, see topics in the *Data Privacy Management Guide*.

Default Storage Passwords

The default storage passwords for database accounts that store alerts in ESA, and Respond Service can be changed. For more information, see "Change Default Storage Passwords" in the *Host and Services Getting Started Guide*.

Alert System Settings

For instructions on configuring NetWitness to send alerts or notifications, see the following topics in the *System Configuration Guide*:

- "Email Configuration Panel"
- "Global Audit Logging Configurations Panel"
- "Global Notifications Panel"

FIPS Compliance

This topic provides information on the Federal Information Processing Standards (FIPS) compliant mode for NetWitness. The FIPS publications are guidelines that set best practices for software and hardware security products for the protection of valuable and sensitive information.

When the FIPS compliant mode is used, products that support one or more FIPS standards can be set into a mode where the product uses FIPS approved algorithms and methods only.

NetWitness supports both FIPS approved and non-FIPS approved functions. In FIPS mode, the product is incapable of using any non-FIPS approved methods.

NetWitness Platform Components working in FIPS mode

The table below lists the NetWitness components that work in FIPS mode. The method you use to activate or deactivate FIPS depends on the type of security library used by your NetWitness services. Your NetWitness services use the security libraries, as mentioned in the following table.

Services	Security Library
Event Stream Analysis (ESA), Malware Analysis, Reporting Engine, NetWitness Host, Respond Service, Context Hub and Endpoint	BSAFE
Broker, Concentrator, Decoder, Log Decoder, Warehouse Connector, Archiver, and Workbench	BSAFE

FIPS 140-2 Certified Cryptographic Modules are enabled for all services that perform cryptographic operations. For the following services, although the FIPS Cryptographic Module is leveraged, the use of FIPS cipher suites is not being enforced:

- NTP: UPD Port 123
- TCP: SSH Port 22
- TCP: Salt API Loopback Port 8000
- CollectD
- Log Collector
- Log Decoder

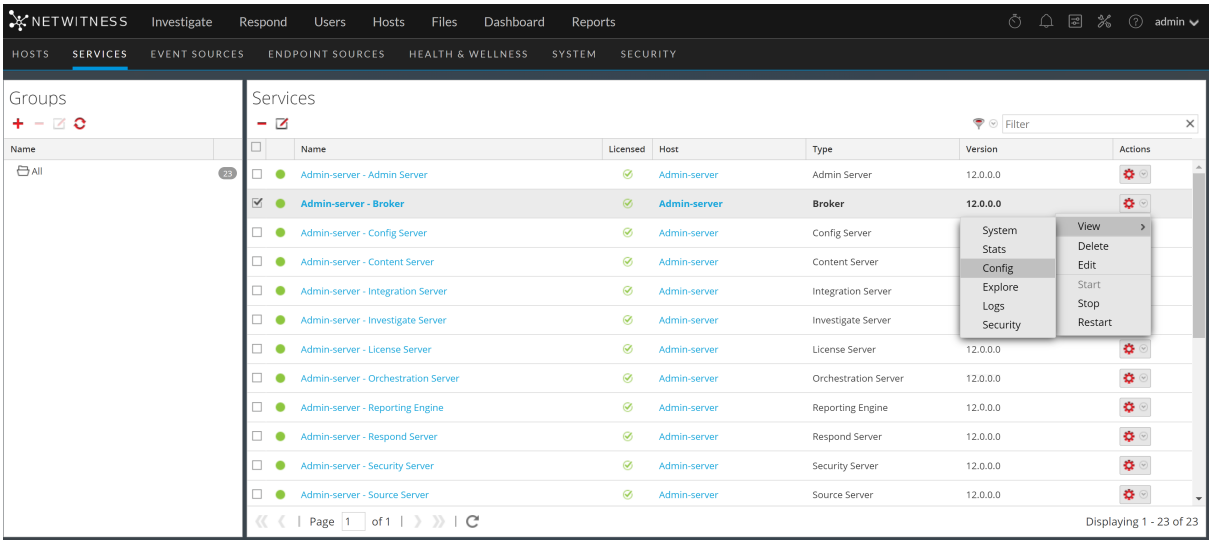
Common Criteria Compliance

This topic provides information on Common Criteria Compliance for NetWitness NetWitness.To support this requirement, you must ensure that only a secure communication is configured for the core services. To achieve this you must disable the unencrypted ports for the NetWitness core services.

Disabling Unencrypted Ports For NetWitness Core Services

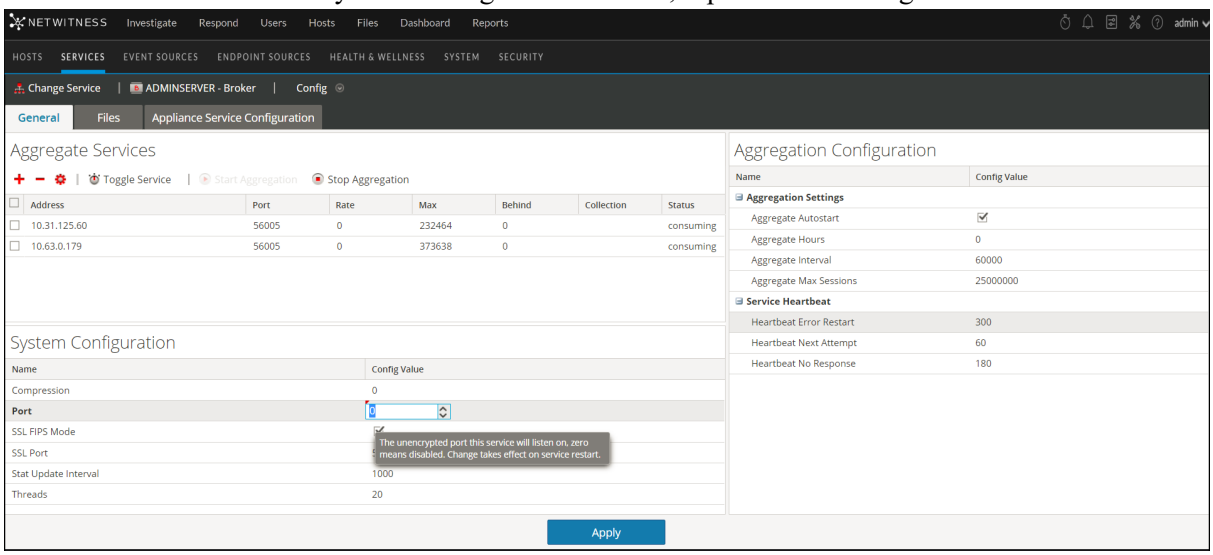
To disable an unencrypted port for a NetWitness core service:

1. Log in to NetWitness Platform.
2. Go to  **(Admin) > Services**.
   The Services page is displayed.
   Select a core service to configure.
3. Click  and select **View** > **Config**.



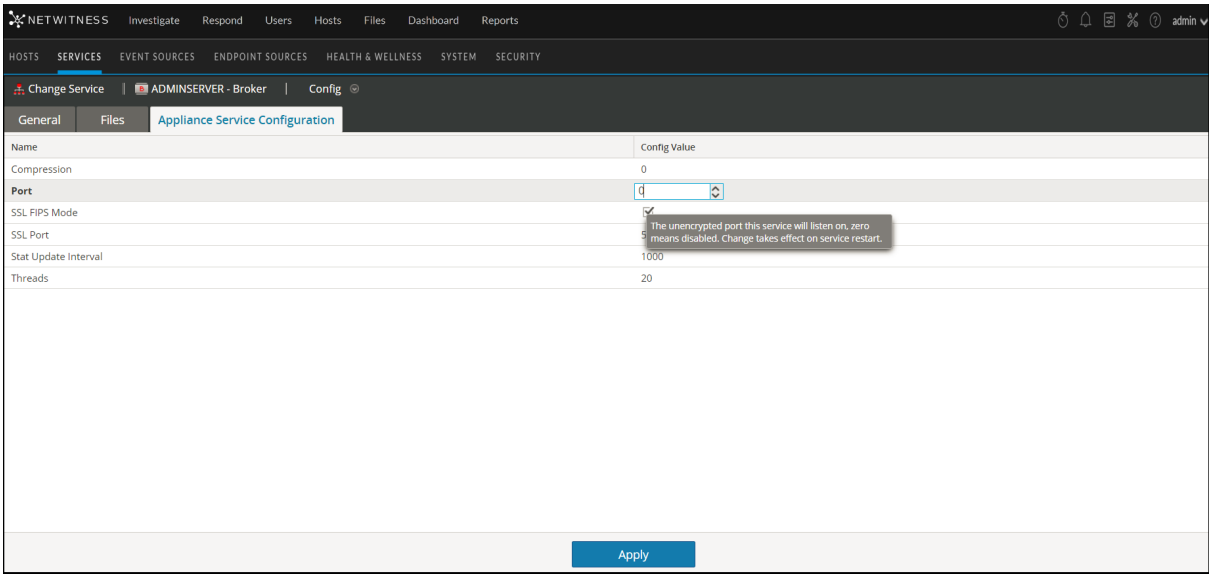
4. Select the General tab.

5. In the Port field under the System Configuration section, replace the existing value with **0**.



6. Now click the Appliance Service Configuration tab.
7. In the Port field, replace the existing value with **0**

.



8. Click **Apply** and restart the service, if prompted.

Note: After you apply the changes only the SSL port is configured for the service and no unencrypted interfaces are available to interact with the service.

STIG Compliance

Overview

Note: Security Technical Implementation Guide (STIG) Compliance is available in NetWitness.

This topic provides information on Security Technical Implementation Guide (STIG) Compliance.

STIG Limits Account Access

NetWitness 11.3.1 and later, is STIG compliant which helps to lock down information, systems, and software, which might otherwise be vulnerable to a malicious computer attack by limiting account access to a system.

For example, STIG:

- Ensures that the account password has a length, complexity, expiration period, and lockout period that are in accordance with DISA's best practices.
- Applies auditing and logging of user actions on the appliance.

To be STIG compliant, your organization must implement policies that ensure strong passwords. For more information about STIG and STIG exceptions, see DISA STIG in the *System Maintenance Guide*.

Other Security Considerations

This topic describes various other security configuration settings that are not covered in previous sections.

Changing the RabbitMQ Management Password for Windows Legacy Collectors

For Windows Legacy Log Collectors (WLCs), a default password is used for the "logcollector" username to access the RabbitMQ broker on that machine. NetWitness recommends that you change the password for WLCs, per the procedure outlined, which involves changing the RabbitMQ password for the Log Collector and for the RabbitMQ broker.

Note: For CentOS, changing the RabbitMQ password is not supported.

If you are using a Log Collector, you may have to initialize the "lockbox". For instructions, see "Set Up a Lockbox" in the *Log Collection Guide*.

To change the RabbitMQ password:

1. Change the RabbitMQ password in Log Collector:
   a. Go to the Explore view for the Log Collector service.
   b. Right-click the **event-broker** node and select **config**.
   c. Type the new password in the **amqp_password** field. The password is encrypted by a key that is managed through the lockbox of this Log Collector. This only changes the password on the Log Collector side.

   Note: Most of the settings should not be changed. Ensure you do NOT change the Message Queue User Name "amqp_username" because it is referred to in some certificate checks.
2. Change the RabbitMQ password for the RabbitMQ broker:
   a. Go to the Explore view for the Log Collector service.
   b. Right-click the **event-broker** node and select **properties**.
   c. Select **passwd** in the drop-down list.
   d. In the **Parameters** field, type the old and new password.
   Ensure you remember your old password. If it was never changed, it should be "netwitness" by default.
   Example: `Parameters: oldpw=<netwitness> newpw=<YourNewPasswordHere>`
   e. Click **Send**.

Hardening the NetWitness Platform Core service

By default, all NetWitnessCore services ship with a default username and password and with SSL turned off. To harden the service, you have to run it with the command line option `-s harden=true`

Using a Decoder, here's an example command line:

```
NwDecoder -s harden=true -s defaultUsername=<username> -s
defaultPassword=<password>
```

The above command does the following:

1. Removes the default admin account (with caveats, see below).
2. Creates a new account <username> with a password of <password> (thus meeting the password requirements below).
3. Enables SSL on both the native and REST ports.
4. Strengthens default password requirements:
   - /users/config/account.lockout.time = 60
   - /users/config/password.alpha.lowercase.min = 1
   - /users/config/password.alpha.uppercase.min = 1
   - /users/config/password.length.min = 8
   - /users/config/password.numeric.min = 1
   - /users/config/password.symbol.min = 1
5. Sets /rest/config/user.agent.whitelist = Apache-HttpClient/\d\.\d\.\d

   Note: This setting prevents the browsers to connect to the REST port.

The caveat for changing the default user account is that there cannot be an already existing configuration file. This is always true the first time the service is run or before the service is licensed. To harden a service, you must run it before a configuration is written or delete whatever configuration file exists and then harden.

To alter the command line for a service that writes its own upstart script without actually SSHing into the box and modifying the script, there is a new parameter that you can pass to either the /sys shutdown or /decoder reset command (substitute decoder for the actual service name) and this parameter is called "cl" for command line. What you do is pass name=value pairs to the "cl" parameter and those parameters will take affect on the next restart of the service.

Example:

```
/sys shutdown reason="Restart because license was applied"
cl="harden=true default
Username=<username> defaultPassword=<password>"
```

The above command shuts down the service (which should be restarted by Linux upstart) and the command line parameters will be applied on the restart. This command line exactly matches the command line given above for the decoder service. If you want to do a configuration reset, you can use the following:

```
 /broker reset config=true cl="harden=true defaultUsername=<username>
defaultPassword=<password>"
```

This will delete the broker configuration file and create a new default configuration that is automatically hardened with the given default account and credentials. The admin account will not exist when the broker restarts, only the <username> account exists.

NFS Access Controls

By default, the NFS mounts are wide open. To lock them down to a specific address, you must edit the exports file and specify the IP addresses that are allowed to interact with the SAW.

The SAW NFS service is managed from the command line using *mapr-nfsserver*.

```
[root@saw-node1 ~]# service mapr-nfsserver
Usage: /etc/init.d/mapr-nfsserver {start|stop|status|restart|}
[root@saw-node1 ~]# service mapr-nfsserver status
nfsserver (pid 5692 5691) is running...
[root@saw-node1 ~]#
```

If nfs-utils is installed on the node, you can execute a *showmount* on the localhost to see the exposed exports.

```
[root@saw-node1 ~]# showmount -e localhost
Export list for localhost:
/mapr *
/mapr/saw *
[root@saw-node1 ~]#
```

Exports are controlled using the exports file in the */opt/mapr/conf* directory.

```
[root@saw-node1 ~]# cat /opt/mapr/conf/exports
# Sample Exports file
# for non /mapr exports
# <Path> <comma separated cldb addresses=host:port> <exports_
control>
# for /mapr exports
# <Path> <exports_control>
#access_control -> order is specific to default
# list the hosts before specifying a default for all
# a.b.c.d,1.2.3.4(ro) d.e.f.g(ro) (rw)
# enforces ro for a.b.c.d & 1.2.3.4 and everybody else is rw
# special path to export clusters in mapr-clusters.conf. To disable
exporting,
```

```
# comment it out. to restrict access use the exports_control
#
/mapr (rw)
#to export only certain clusters, comment out the /mapr & uncomment.
# Note: this will cause /mapr to be unexported
#/mapr/clustername (rw)
#to export /mapr only to certain hosts (using exports_control)
#/mapr a.b.c.d(rw),e.f.g.h(ro)
# export /mapr/cluster1 rw to a.b.c.d & ro to e.f.g.h (denied for
others)
#/mapr/cluster1 a.b.c.d(rw),e.f.g.h(ro)
# export /mapr/cluster2 only to e.f.g.h (denied for others)
#/mapr/cluster2 e.f.g.h(rw)
# export /mapr/cluster3 rw to e.f.g.h & ro to others
#/mapr/cluster2 e.f.g.h(rw) (ro)
[root@saw-node1 ~]#
```

To restrict the SAW exports to a certain IP address or group of IPs, you must first edit the exports file and then restart the *mapr-nfsserver service*.

```
[root@saw-node1 ~]# vi /opt/mapr/conf/exports
[root@saw-node1 ~]# cat /opt/mapr/conf/exports | grep ^/mapr
/mapr 10.42.1.87(rw)
[root@saw-node1 ~]# showmount -e localhost
Export list for localhost:
/mapr *
/mapr/saw *
[root@saw-node1 ~]# service mapr-nfsserver restart
[root@saw-node1 ~]# showmount -e localhost
Export list for localhost:
/mapr 10.42.1.87
/mapr/saw 10.42.1.87
[root@saw-node1 ~]# mount -t nfs -o nolock,tcp localhost:/mapr/saw
/saw
mount.nfs: access denied by server while mounting
localhost:/mapr/saw
[root@saw-node1 ~]#
```

Note: Trying to mount the export on the localhost will fail as only a specific host IP is now allowed to use the NFS mount.

Secure Deployment and Usage Settings

This topic describes the settings for secure deployment and usage. It is very important to protect all physical, local, and remote access to the NetWitness appliances. It is also important to restrict all access methods to the absolute minimum required to maintain NetWitness.

Note: NetWitness recommends that you do not set up the test environments to be exact copies of the full production environment. If the test environment is identical to the production environment, you should take the same precautions to protect the test environment as you do in the production environment.

Security Controls Map

This topic describes the security controls map. An NetWitness deployment can consist of the following components:

- Decoder
- Log Decoder
- Concentrator
- Broker
- Log Collector
- Context Hub
- Malware
- ESA
- Archiver
- NetWitness Server
- Endpoint
- UEBA Server

NetWitness supports integration with products such as NetWitness 4.x and Archer.

NetWitness recommends that you access the host on secure client machines within the network. If you must access the host through remote access, NetWitness recommends that you connect to the network through a secure VPN connection. Only allow remote access to NetWitness hosts for secure maintenance using the Remote Desktop Protocol (RDP) through a secure VPN connection.

Caution: NetWitness recommends that you deploy the hosts in a secure location, where physical access to the hosts are restricted only to the personnel who manage the hosts.

Secure Enclave

To help protect NetWitness against unauthorized authentication and access by end users or machines, NetWitness recommends that you deploy NetWitness hosts such as Broker, Concentrator, and Decoder.

You can help create a secure enclave by separating the low security corporate network from the high security network with firewalls. To help create a secure enclave, NetWitness recommends that you:

- Implement basic physical security elements, policies, procedures, and processes for the low security network.
- Provide access to the hosts within the secure enclave through a secure virtual private network (VPN) tunnel only, such as IPSec tunnel, to establish encryption and authentication of all network traffic to and from the hosts.

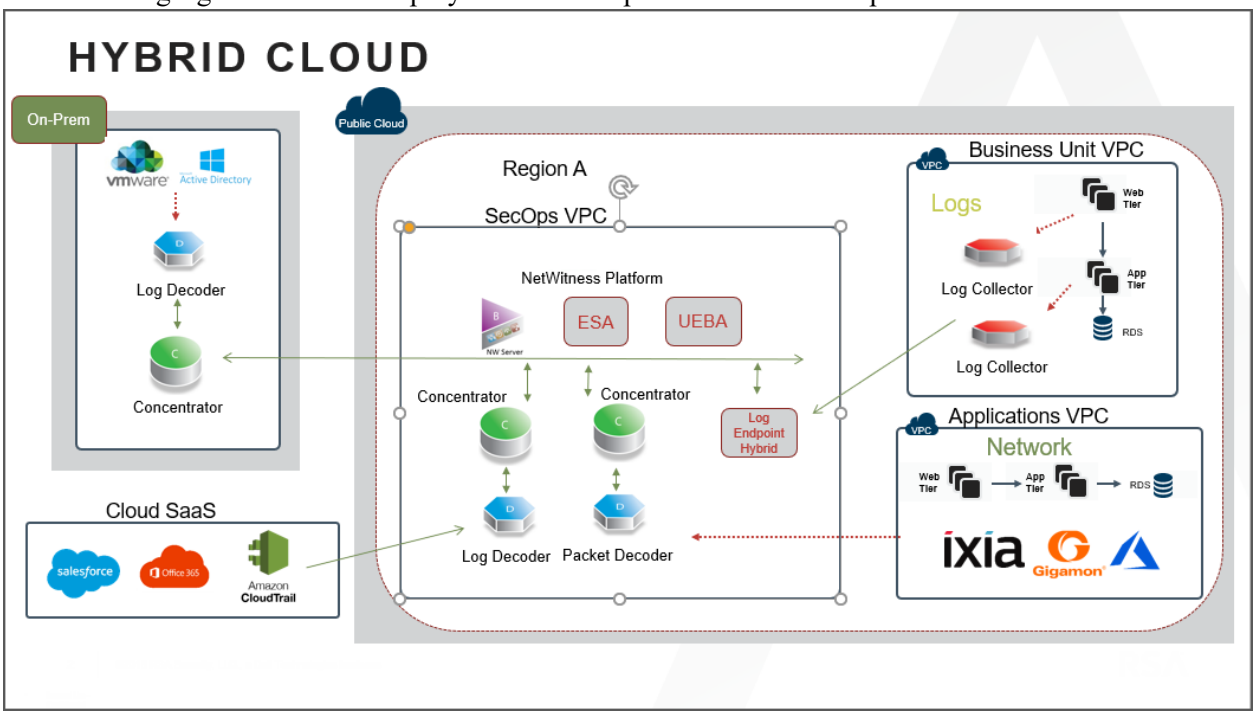
Note: The client machines through which you access NetWitness can be present outside of the Secure Enclave.

Secure Deployment Guidelines

To help ensure a secure deployment, NetWitness recommends that you:

- Deploy multiple hosts in the corporate network. The multiple hosts in the example are in two geographic locations and include the following components:
  - NetWitness
  - Broker
  - Packet Decoder
  - Log Decoder
  - Concentrator
  - ESA
  - Archiver
  - NetWitness Warehouse
  - Malware Analysis
  - Endpoint
  - UEBA Server
- Ensure that all the components are connected to the same subnetwork.
- Deploy firewalls at each site to ensure the secure transfer of data from an instance of NetWitness at one site to another instance of NetWitness located at a different site.
- Configure firewall rules to control all communication between different sites and other components in the network as depicted in the previous figure.
- Implement data transfer between sites using a secure tunnel IPSec.

The following figures show the deployment of multiple sites within a corporate network:



Firewall Rules

It is important that you use a firewall to restrict network traffic between NetWitness and external systems. NetWitness recommends that you configure firewall rules to help ensure secure communication for the following connections:

- Demilitarized zone (DMZ) to corporate network
- Corporate network to site sub network
- Site to site
- Live CMS to DMZ
- External email server to DMZ

Note: NetWitness recommends that you restrict access from client hosts to only known IP addresses. For example, if you set up the NetWitness Client UI on IP address 192.168.0.1, configure your firewall to allow only the IP address 192.168.0.1 to connect to the NetWitness host.

Note: The firewall rules should be configured on an external firewall and not on any of the NetWitness host.

NetWitness recommends that you configure firewall rules as described in the sections below. These recommendations are based on the following assumptions:

- You have a stateful firewall, indicating that only the establishment of TCP ports is considered.
- You specify the direction of communication for the UDP ports because the connections are sessionless.
- You deploy NetWitness as shown in the Security Controls Map.
- The firewall processes the rules top to bottom, finishing with a generic drop of all the packets.

DMZ to Corporate Network

NetWitness recommends that you:

- Configure whitelist communication from the VPN server in the DMZ to the client machines on which you run NetWitness applications such as NetWitness Web UI.
- Create firewall rules for all the machines from which you intended to remotely access the corporate network through Remote Desktop Protocol (RDP).

Corporate Network to Site

NetWitness recommends that the firewall at each NetWitness site allow access only from designated client machines through a whitelisted IP address and port.

NetWitness recommends that you secure the following ports to ensure secure communication between the client machine that is set as the NetWitness Web UI and the NetWitness site:

- TCP 443

For more information, see Communication Security Settings. To help ensure secure communication between the client machines that access the NetWitness UI and a site, you must set up the firewall rules as shown below:

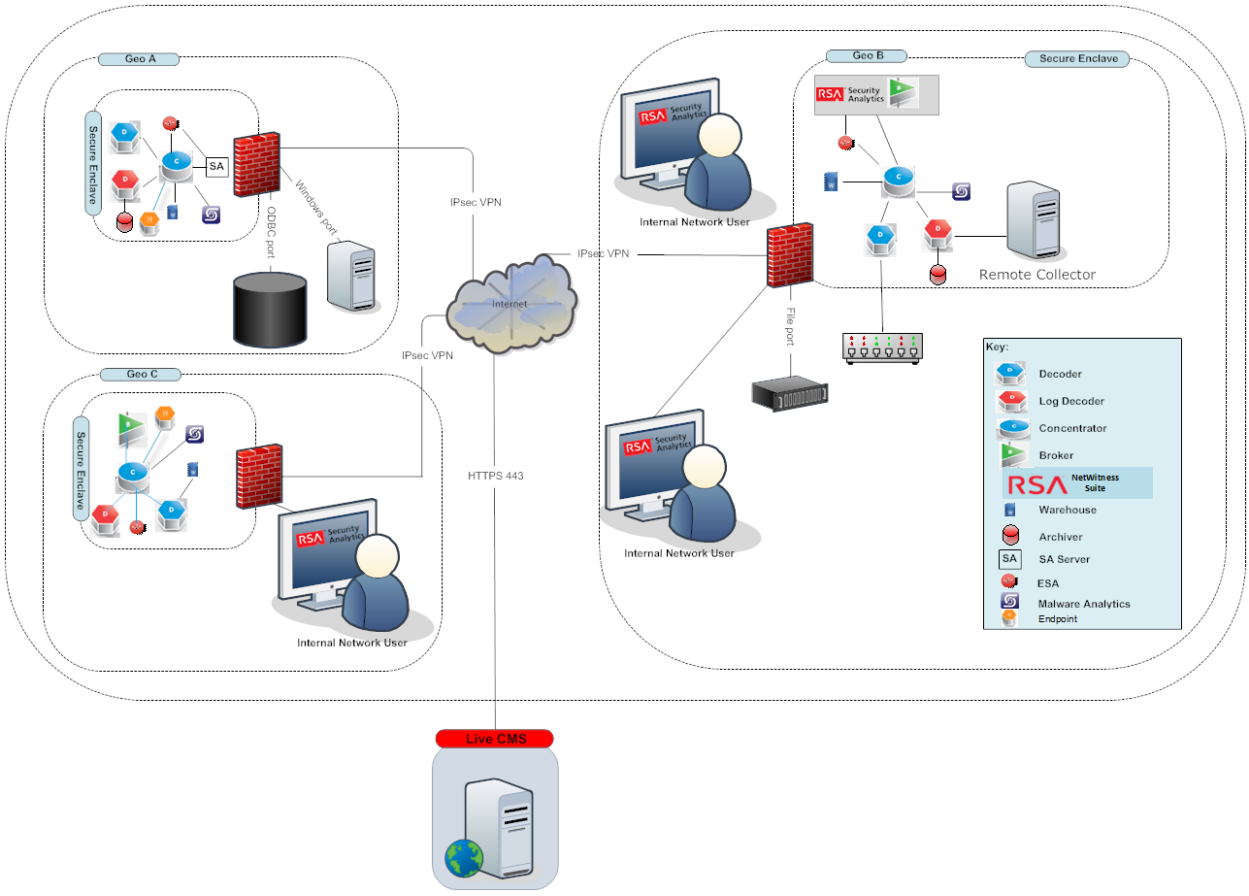
ALLOW $nwclient_IP to $nwsite_IP on port 443/tcp

DROP all from * to *

where **nwclient_IP** is the IP address assigned to the client machine that is set as the NetWitness web UI and **nwsite_IP** is the IP address assigned to the Broker host within which the NetWitness web server is running.

Site to Site

NetWitness may run in multiple sub-networks within your corporate network, called sites. You can configure NetWitness to allow the hosts located in one site to communicate with the hosts in another site.



For this scenario, NetWitness recommends that you do the following:

- Ensure that the firewall at each NetWitness site allows communication between two sites only through a whitelisted IP address and port. For a graphical depiction of the site-to-site security control map showing the site firewalls, see the above figure.
- NetWitness system update uses port 80. That means NetWitness site to another site (where brokers, decoders exist), port 80 should be open.

Live CMS to DMZ

To ensure secure communication between the NetWitness site and Live CMS, set up the firewall rules as shown below:

ALLOW $nw_site_IP to $Live_IP on port 443/tcp

DROP all from * to *

where **nw_site_IP** and **Live_IP** are the IP addresses assigned to the NetWitness site and the Live CMS respectively.

Note: If you are using proxy server with self-signed certificate, you must add exception in proxy server rule to allow traffic between Live CMS server (cms.netwitness.com, port 443) and NetWitness.

RSA Download Central to DMZ

To ensure secure communication between the NetWitness site and RSA Download Central, set up the firewall rules as shown below:

ALLOW $nw_site_IP to $rsa_DLC_IP on port 80/tcp

DROP all from * to *

where **nw_site_IP** and **rsa_DLC_IP** are the IP addresses assigned to the NetWitness site and RSA Download Central respectively.

External Email Server to DMZ

To ensure secure communication between the NetWitness site and the External Email Server, set up the firewall rules as shown below:

ALLOW $nw_site_IP to $Email_IP on port 443/tcp

DROP all from * to *

where **nw_site_IP** and **Email_IP** are the IP addresses assigned to the NetWitness site and the external email server respectively.

Syslog Server to Site

If you have enabled the syslog port, to ensure secure communication between the NetWitness site and the Syslog Server, set up the firewall rules as shown below:

ALLOW $nw_site_IP to $Syslog_IP on port 514/udp

DROP all from * to *

where **nw_site_IP** and **Syslog_IP** are the IP addresses assigned to the NetWitness site and the syslog server respectively.

SNMP Server to Site

If you have enabled the SNMP port, to ensure secure communication between the NetWitness site and the SNMP Server, set up the firewall rules as shown below:

ALLOW $nw_site_IP to $SNMP_IP on port 1610/SNMP

DROP all from * to *

where **nw_site_IP** and **SNMP_IP** are the IP addresses assigned to the NetWitness site and the SNMP server respectively.

Secure Deployment Settings

The following table shows the security controls that NetWitness recommends putting in place to help secure the deployment.

Default Deployment Setting	Secure Deployment Settings	Pros of Secure Deployment Settings	Cons of Secure Deployment Settings	Instructions on how to configure secure deployment settings
HTTPS is enabled by default between the NetWitness client and the server	For the best possible security between the client and the server, use certificates from CA.	Provides a high level of protection for the communication between client and server by avoiding tampering, spoofing, and man-in-the middle attacks.	May have impact on performance	For instructions on installing external certificates, see SSL Certificate Guidance for NetWitness

Secure Maintenance

This topic describes some common solutions to help ensure secure maintenance.

Security Patch Upgrade

To apply security patches to any of the required release, see the *Upgrade Guide* for the specific upgrade version.

Security Patch Management

All security patches for NetWitness originate at NetWitness and are available for you via the NetWitness User Interface. For more information, see "Manage NetWitness Updates" in the *System Maintenance Guide* .

The following table lists the third-party components for which patches are needed.

Third-party Component for which patch is needed	Frequency of Patch	EMC Responsibility (Y/N)	Customer Responsibility (Y/N)	Reference to instructions for Applying Patch
NetWitness Hosts	Monthly and Quarterly	Y	Y	Based on NetWitness recommendations

Virus Scanning

NetWitness recommends that you:

- Deploy anti-virus client software on the deployed servers in accordance with your enterprise requirements.
- Run anti-virus and anti-malware tools with the most current definition files on the deployed servers.
- Scan all files/drivers before uploading on the deployed server.
- Follow best practices for patch management and regularly review available patches for all anti-virus and anti-malware software.

Ongoing Monitoring and Auditing

As with any critical infrastructure component, NetWitness recommends that you constantly monitor your system and perform periodic and random audits (for example, configuration, permissions, and security logs). You should ensure that the configurations and user access settings match your company policies and needs. For more information, see "Global Audit Logging Configurations Panel" in the *System Configuration Guide*.

Hardware Replacement

If NetWitness hardware fails or is faulty, order a replacement by contacting NetWitness Customer Support. While awaiting a replacement, the Redundant Array of Independent Disks (RAID) configuration is designed to ensure that there is no data loss due to a hardware failure.

The RAID configuration on NetWitness:

- Hosts are RAID 1.
- Direct Attach Capacity (DAC) disk shelves is RAID 5.

Physical Security Controls Recommendations

Physical security controls help to enable the protection of resources against unauthorized physical access and physical tampering. NetWitness recommends that the physical devices and servers for NetWitness are deployed in a secure data center leveraging the organization's best practices for physically securing a data center, server rack, and/or server.

Supporting Users

This topic describes well-defined policies around help desk procedures for your NetWitness installation.

It is important to have well-defined policies around help desk procedures for your NetWitness installation. NetWitness recommends that your help desk administrators understand the importance of password strength and the sensitivity of data, such as user logon names and passwords. Creating an environment where an end user is frequently asked for this kind of sensitive data increases the opportunity for social engineering attacks. Train end users to provide, and help desk administrators to request, the least amount of information needed in each situation.

Preventing Social Engineering Attacks

Fraudsters frequently use social engineering attacks to trick unsuspecting employees or individuals into divulging sensitive data that can be used to gain access to protected systems. NetWitness recommends that you use the following guidelines to help reduce the likelihood of a successful social engineering attack:

- Help desk administrators should only ask for User IDs over the phone when receiving help desk calls. Help desk administrators should never ask for user passwords.
- The help desk telephone number should be well known to all users.
- Help desk administrators should perform an action to authenticate the user's identity before performing any administrative action on a user's behalf. For example, ask users one or more questions to which only they know the answer.
- If help desk administrators need to initiate contact with a user, they should not request any user information. Instead, users should be instructed to call the help desk back at a well-known help desk telephone number to ensure that the original request is legitimate.

Confirming User Identities

It is critical that your help desk administrators verify end users' identities before performing any help desk operations on their behalf. NetWitness recommends that you verify user identity using the following methods:

- Call the end user back on a phone owned by the organization and on a number that is already stored in the system.

Caution: Be wary of using mobile phones for identity confirmation, even if they are owned by the company because mobile phone numbers are often stored in locations that are vulnerable to tampering or social engineering.

- Send the user an email to a company email address. If possible, use encrypted email.
- Work with the employee's manager to verify the user's identity.
- Verify the identity in person.

- Use multiple open-ended questions from employee records. For example, "Name one person in your group" or "What is your badge number?" Avoid yes or no questions.

Advice for Your Users

NetWitness recommends that you instruct your users to do the following:

- Never give passwords to anyone.
- Change passwords at regular intervals.
- Inform your users of what information requests to expect from help desk administrators.
- Always log off from the web interface when finished.
- Always lock their desktops when stepping away from their computers.
- Regularly close their browser and clear their cache of data.

Note: Consider regular training to communicate this guidance to users.

Appendix A: Customer Provided Certificates

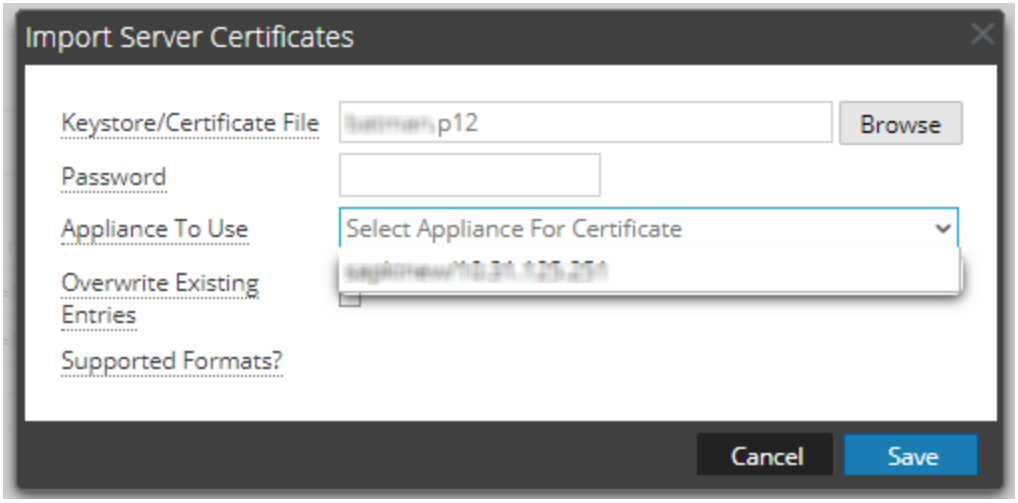
The procedure tells you how to replace the internally generated NetWitness web server certificate (NGINX front-door) with a customer issued certificate. This enables client browsers to establish a trusted SSL connection.

Caution: The cert files and key files must be `.pem` format. All the files must have the same name and permissions as the original files generated by NetWitness.

1. Rename your certificate files and save them in for NGINX.
   - Rename the customer provided `cert.pem` certificate pem file to `web-server-cert.pem`.
   - Rename the customer provided `key.pem` key pem file to `web-server-key.pem`.
   - Rename customer provided `cert.chain` certificate chain file to `web-server-cert.chain`.
   - Rename `cert.p7b` certificate p7b file to `web-server-cert.p7b`.
2. SSH to the NW Server.
3. Replace the existing NetWitness Platform generated `/etc/pki/nw/web/web-server-cert.pem`, `/etc/pki/nw/web/web-server-key.pem`, `/etc/pki/nw/web/web-server-cert.chain` and `/etc/pki/nw/web/web-server-cert.p7b` files with the files you renamed in step 1.
4. Restart NGINX service.
   ```
   service nginx restart
   ```

To add Custom Server Certificate in PFX, P12 and JKS format.

1. Go to  **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **PKI Settings** tab.
3. In the **Server Certificates** section, click .
   The Import Server Certificates dialog is displayed.



4. In the **Keystore/Certificate File** field, click **Browse** and select the keystore.
5. In the **Password** field, enter the keystore password.
6. In the **Appliance To Use** field, select the appliance for which you want to use this certificate.
7. (Optional) Select the **Overwrite Existing Entries** checkbox to overwrite the entries of the certificate that is already added.
8. Click **Save.**
   The NetWitness Server certificate with its private key is successfully added to NetWitness.

Note: When the certificate is being applied on the selected appliance, no other operation on PKI can be performed until the process is completed.
Double-click on the added entries to view the details of the certificate.

9. To apply the server certificate on a server, select a certificate and click .

Note: Uploading a keystore will add the server certificate and its private key locally. To apply a server certificate on a server, you need to select a server certificate and click the synchronization button .
All server certificates are also synchronized on the appliances when PKI is enabled.

Appendix B: Reissue Certificates

Introduction

For a secure deployment, NetWitness has installed internal NetWitness issued certificates such as CA Certificate and Service certificates.

The validity for NetWitness certificates are as follows:

- CA root certificate for 11.x deployment is valid for 10 years
- CA root certificate for 10.6.x deployment is valid for 5 years
- Service certificates are valid for 1000 days

When these certificates are about to expire or have expired, you must renew and reissue the certificates as soon as possible to avoid any issues with your NetWitness deployment.

Note: You can view the expiration details, by executing the `ca-expire-test-sh` script on the NetWitness Server. For more information, see Reissue root CA security certificates on NetWitness Platform 11.x and download the script.

CA Certificate Reissue

To renew the CA certificates, do the following:

- Before you upgrade from 10.6.x to 11.x, check the expiry and reissue those certificates. For more information, see the Reissue root CA security certificates on NetWitness Platform 11.x.
- If you are on 10.6.x , check the expiry and reissue all the certificates. For more information, see the Reissuing security certificates on NetWitness Platform 10.6.x.

Note: If you have Windows Legacy Collectors (WLC) in your deployment, renew the CA certificate of the WLC after renewing the CA certificate of the NetWitness Admin Server.

Service Certificate Reissue

To renew the Service certificates, do the following:

- In case of a fresh installation of 11.3 or later, you must use the `cert-reissue` script. For more information, see the Reissuing Service Certificate .
- In case of a fresh installation of 11.1 or 11.2 or if you upgrade from 10.6.x to 11.x, you must use the `nw-root-ca-update` script. For more information, see the Reissue root CA security certificates on NetWitness Platform 11.x.

Note: If you have a host that is decommissioned or plan to remove, do not renew the certificate for that host.

Reissuing Service Certificate

You can reissue service certificates in the following two ways.

- All at once
  Reboot NW Server host after the `cert-reissue --host-key` command completes.
- One at a time
  Reissue the NW Server host certificates first, restart the host,and then reissue each component host. Ensure that you restart the component host as well.

IMPORTANT: If you are reissuing certificates for each host individually (one at a time), you must reissue the certificate for the NW Server host before you can reissue certificates for any other host.

When to Use the `--host-key` Argument

Use the `cert-reissue --host-key` command string if you have a large number of hosts. Make sure that:

- All your hosts are running 11.4.0.0 or later.
- All your hosts are online.
- The NW Server host run time services are running.

`cert-reissue` Arguments and Options for All Hosts

The following tables lists the argument you can use to reissue certificates for all hosts at one time. See Appendix C. Troubleshooting Cert-Reissue Command for additional options you can use with Customer Support to troubleshoot errors.

Arguments	Description
`--host-key`	Reissues certificates for all hosts at one time applying system health checks and restarts services. **Note:** If even one host is not online, this command fails. If you have numerous hosts in your deployment, make sure that all hosts are up and running.

Caution: Make sure you do not run this argument on a node or host that you plan to remove or decommission.

When to Use the Individual Host Arguments (`--host-id <id>`, `--host-name <display-name>`, `--host-addr <ip/hostname>`)

The `cert-reissue --host-id <id>`, `cert-reissue --host-name <display-name>`, or `cert-reissue --host-addr <ip/hostname>` reissues a certificate for an individual host. You may want to reissue certificates for an individual host if you have a small number of hosts.

Make sure that:

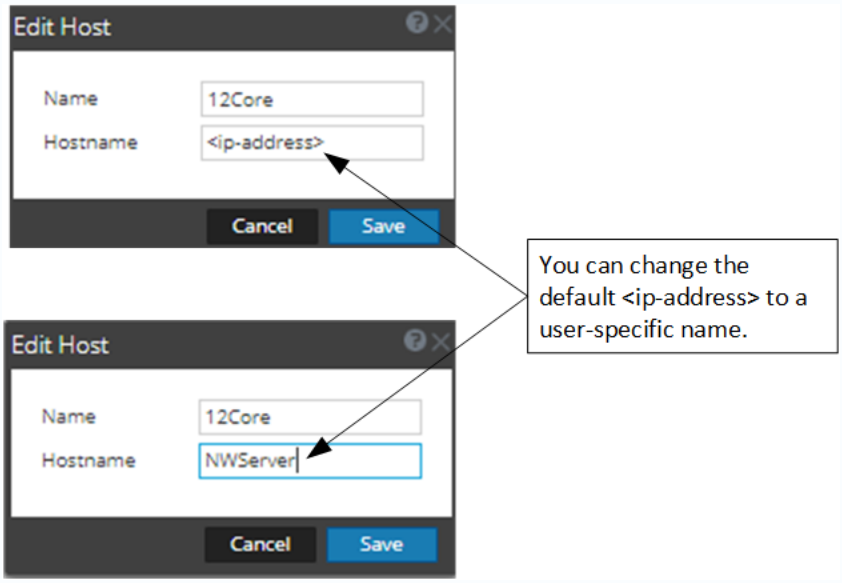
- Each host is running 11.4.0.0 or later.
- Each host is online.
- The NW Server host run time services are running
- You reissue certificates for the NW Server host first.

`cert-reissue` Arguments and Options for a Single Host

The following tables lists the arguments and options you can use to reissue certificates for a single host (one host at a time). For more information, see the Appendix C. Troubleshooting Cert-Reissue Command section on the additional options you can use with NetWitness Customer Support to troubleshoot errors.

Note: You must run the command for the NW Server host first and reboot that host before you run the command for each component host.

Arguments	Description
`--host-id <id>`	Reissues certificate for the host identified by `<id>` (host identification code).
`--host-name <display-name>`	Reissues certificate for host identified by `<display-name>`. `display-name` is the value shown under **Name** in the **(missing or bad snippet) > Hosts** View in the NetWitness Platform Interface.

`--host-addr <Hostname-in UI>`
or
`--host-addr <hostname>`
	Reissues certificate for the host identified by the value shown under **Hostname** in the**(missing or bad snippet) > Hosts > Edit** dialog in the NetWitness Platform Interface. This value can be an ip-addres (default) or a user-specified name.
	

Reissuing Certificates for All Hosts Except Windows Legacy Collection (WLC) host

Use the `cert-reissue` command to reissue certificates for all hosts except the WLC host with the following procedures.

Running the Cert-Reissue Command for All Hosts

1. SSH to the NW Server host.
2. Submit the appropriate command string.
   `cert-reissue --host-key`

Running the Cert-Reissue Command for an Individual Host

1. SSH to the NW Server host.
2. Submit the appropriate command string (that is `cert-reissue --host-id` or `--host-name` or `--host-addr`). Each of the following command strings is an example of how you reissue certificates for a specific host.
   - `--host-id <host-identification-code>`
   - `--host-name <named-displayed-under-Name-in-Hosts-view>`
   - `--host-addr <ip-address-default-hostname-or-user-specified-hostname>`

Reissuing Certificates for a WLC Host

You must use the `wlc-cli-client` utility to reissue certificates for a WLC host (you cannot use the `cert-reissue` command). You also need to specify a number of WLC identification parameters with this utility.

Note: The certificates for a Windows Legacy Server host are stored in the following directories on the host.
`C:\ProgramData\netwitness\ng\logcollector_cert.pem`
`C:\ProgramData\netwitness\ng\logcollector_dh2048.pem`
Th validity period of WLC certificates can range from 2 to 20 years. If you rename or remove the files and restart **NwLogCollector** Service, NetWitness regenerates them.
`/ssl/truststore.pem` - is no longer used in 11.x
Every reissue of a certificate on the Windows Legacy server creates a new private key.

To reissue certificates on a WLC host.

1. SSH to the NW Server host.
2. Submit the following command string.
   ```
   wlc-cli-client --cert-renew --host <wlc-host-ip-address> --port 50101 --use-ssl false --username <wlc-username> --password <wlc-password> --ss-username <deploy-admin-username> --ss-password <deploy-admin-password>
   ```

Successful Reissue Summary Report

When you run `cert-reissue --host-key`, the following summary report will be displayed if all hosts are online, all run time services are running, and all hosts on version 11.4.0.0 or later.

```
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
+---------+--------------------+--------+-----------------------------------------------------+
```

Unsuccessful Reissue Summary Reports

You must contact NetWitness Customer Support to troubleshoot problems. You know there is a problem if any `<host-id>` does not return a `Success`**Status**. `Success` indicates that certificates were reissued for a host. The following examples illustrate unsuccessful reissues.

Reissue Failed for Host and Aborted Command

The following three examples illustrate the failure of certificate reissuing for any hosts.

```
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Failed! |failed to connect, is host online?                   |
|<host-id>| <IP-address>       |Failed! |service(s) down                                      |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
+---------+--------------------+--------+-----------------------------------------------------+
```

```
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Failed! |version <version-earlier-than-11.3.0.0> not supported|
|<host-id>| <IP-address>       |Failed! |version <version-earlier-than-11.3.0.0> not supported|
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
+---------+--------------------+--------+-----------------------------------------------------+
```

Reissue Certificate Partially Executed

The NW Server Host certificates were reissued but failed to properly distribute the reissued certificates to one or more component hosts.

```
+---------+--------------------+
|         | Host               |
+---------+--------------------+
|<host-id>| <IP-address>       |
|<host-id>| <IP-address>       |
+---------+--------------------+
...
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Partial |Reissue completed, triggers failed                   |
|<host-id>| <IP-address>       |N/A     |[ Skipped… ]                                         |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
+---------+--------------------+--------+-----------------------------------------------------+
```

Appendix C. Troubleshooting Cert-Reissue Command

You must contact NetWitness Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897) to troubleshoot problems. You know there is a problem if any `<host-id>` does not return a `Success`**Status**. `Success` indicates that certificates were reissued for a host.

Argument Options Used for Troubleshooting

You use the following argument options with `cert-reissue --host-all` to troubleshoot problems.

You can run `cert-reissue --host--all<arguments>` multiple times without an adverse effect.

Note: Use the following Argument Options with caution. They force the `cert-reissue` command to execute for all the hosts.

Argument Option	Description
`--skip-health-checks`	Reissues certificates for all hosts at one time without applying system health checks (force Reissue). This means that the command does not: • verify that all hosts are online line. • verify that all services are running. **Use case:** You have numerous hosts and you know that a small minority of them will fail. This updates all the hosts that conform to the checking rules and you can reissue certificates for the others subsequently with the help of Customer Support.
`--skip-version-checks`	Do not verify that hosts are running version 11.3.0.0 or later. **Use Case:** You have numerous hosts and your know that some of them are not updated to 11.3 or later. This reissues certificates for all the hosts that are at 11.3 or later and you can reissue certificates for the others subsequently with the help of Customer Support.
`--ignore-trigger-errors`	Ignore any errors that trigger failures. This option forces the cert reissue process to continue disregarding the errors instead of aborting or failing the cert reissue command quickly. When a cert reissue for a host succeeds, the reissued certificates on that host are not provisioned to other dependent hosts (referred to as trusts). In this case, the: • host with reissued certificates is reported as “Partial.” • the hosts with trusts that failed to update are listed separately in the summary table to tell you that these hosts may require a refresh using the new `--refresh-trusts-only` option.
`--refresh-trusts-only`	Refreshes trusts exclusively for host identified by `<id>` (does not reissue certificates for that host).

Problems and How to Troubleshoot Them

This section describes solutions to problems that you may encounter when running the cert-reissue command to reissue certificates with suggested causes and solutions.

Status	Failed!
Error Message	... 2019-02-06 13:34:39.646 INFO 8540 --- [main] c.r.n.i.o.client.OrchestrationClient : Checking host connections... ... 2019-02-06 13:34:57.861 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Host '192.168.200.99' (nw-platform-esa-primary) verification failed! ... 2019-02-06 13:34:57.862 INFO 8540 --- [main] c.r.n.i.o.client.OrchestrationClient : Checking status of services... 2019-02-06 13:35:57.931 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Service 'nw-platform-node-zero - Investigate Server' not available! ... +---------+--------------------+--------+---+ | | Host | Status | Message | +---------+--------------------+--------+---+ |<host-id>| <IP-address> |Failed! |failed to connect, is host online? | |<host-id>| <IP-address> |Failed! |service(s) down | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | +---------+--------------------+--------+---+
Cause	`cert-reissue --host-all` failed because one or more hosts are offline or one or more run time services are unreachable. You can force this command to run in spite of this error by specifying the `--skip-health-checks` option, that is: `cert-reissue --host-all--skip-health-checks`
Solution	1. Bring appropriate hosts back online or make sure the NW Server hosts run time services are running. 2. Run `cert-reissue` for the hosts affected.

Status	Failed!
Error Message	... 2019-02-06 13:34:39.643 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Host '192.168.200.102' (nw-platform-decoder) version '11.2.0.0' not supported, minimum required version: 11.3.0.0 2019-02-06 13:34:39.644 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Host '192.168.200.101' (nw-platform-concentrator) version '11.2.0.0' not supported, minimum required version: 11.3.0.0 ... +---------+--------------------+--------+--+ | | Host | Status | Message | +---------+--------------------+--------+--+ |<host-id>| <IP-address> |Failed! |version <version-earlier-than-11.3.0.0> not supported| |<host-id>| <IP-address> |Failed! |version <version-earlier-than-11.3.0.0> not supported| |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | +---------+--------------------+--------+--+
Cause	`cert-reissue -host-all` command string failed because one or more hosts are running a version earlier than 11.3.0.0 **Note:** You can force the reissue of certificates for the remaining hosts using the `-skip-version-checks` argument.
Solution	Update the host to 11.3 or later and run `cert-reissue` for that host again.

Status	Partial
Error Message	... 2019-02-06 02:27:09.078 ERROR 20647 --- [main] c.r.n.i.o.client.OrchestrationClient : Trigger failed for host '<IP-address>' (nw-platform-decoder) 2019-02-06 02:27:09.079 ERROR 20647 --- [main] c.r.n.i.o.client.OrchestrationClient : Trigger failed for host '<IP-address>' (nw-platform-concentrator) ... 2019-02-06 02:27:09.118 WARN 20647 --- [main] c.r.n.i.o.client.OrchestrationClient : One or more host(s) may require manual refresh due to failed triggers: +---------+--------------------+ | | Host | +---------+--------------------+ |<host-id>| <IP-address> | |<host-id>| <IP-address> | +---------+--------------------+ ... +---------+--------------------+--------+---+ | | Host | Status | Message | +---------+--------------------+--------+---+ |<host-id>| <IP-address> |Partial |Reissue completed, triggers failed | |<host-id>| <IP-address> |N/A |[Skipped…] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | +---------+--------------------+--------+---+
Cause	`cert-reissue` command completed on NW Server host however one or more triggers failed. This aborted the `cert-reissue` command for other hosts.
Solution	Address all the errors and run the `cert-reissue --host--all<arguments>` command string again.

Status	Partial
Error Message	... 2019-02-06 14:18:03.208 ERROR 17800 --- [main] c.r.n.i.o.client.OrchestrationClient : Trigger failed for host '192.168.200.82' (nw-platform-node-x) ... ... 2019-02-06 14:29:05.200 WARN 17800 --- [main] c.r.n.i.o.client.OrchestrationClient : One or more host(s) may require manual refresh due to failed triggers: +---------+--------------------+ | | Host | +---------+--------------------+ |<host-id>| <IP-address> | +---------+--------------------+ ... +---------+--------------------+--------+---+ | | Host | Status | Message | +---------+--------------------+--------+---+ |<host-id>| <IP-address> |Failed! |Cert reissue failed! | |<host-id>| <IP-address> |Partial |Reissue completed, triggers failed | |<host-id>| <IP-address> |Success | Cert reissue successful | |<host-id>| <IP-address> |Success | Cert reissue successful | |<host-id>| <IP-address> |Success | Cert reissue successful | +---------+--------------------+--------+---+
Cause	One or more hosts did not pass system health checks. In addition, one or more of the unhealthy hosts are running core services, which will result in the NW Server host `cert-reissue` to fail (because of failed triggers explained above). By disabling health checks and trigger errors, you can continue the process and reissue certificates for the remaining hosts. The NW Server host **Status** is reported as `Partial` because the `cert-reissue` command completed for the NW Server but downstream triggers failed for other hosts.
Solution	Manually refresh the failed core hosts (to synchronize trust peers). Submit the following command string to reissue certificates for healthy hosts. `cert-reissue --host-all --skip-health-checks --ignore-trigger-errors`

Status	Failed!

Error Message	When cert reissue fails for Respond server only when ESA is not present in the stack, the following messages are displayed in the "`/var/log/netwitness/config-management/chef-solo.log`":
	`1. [2022-05-31T02:11:20+00:00] ERROR: Running exception handlers`
	`2. [2022-05-31T02:11:20+00:00] ERROR: Exception handlers complete`
	`3. 2-05-31T02:11:20+00:00] FATAL: Stacktrace dumped to /var/lib/netwitness/config-management/cache/chef-stacktrace.out`
	`4. [2022-05-31T02:11:20+00:00] FATAL: Please provide the contents of the stacktrace.out file if you file a bug report`
	`5. [202[2022-05-31T02:11:20+00:00] FATAL: Mixlib::ShellOut::ShellCommandFailed: nw_pki_bootstrap_launch[reissue certs for rsa-nw-respond-server] (rsa-response::certreissue line 13) had an error: Mixlib::ShellOut::ShellCommandFailed: execute[respond-server-get-operational-csr] (/var/lib/netwitness/config-management/cache/cookbooks/nw-pki/resources/bootstrap_launch.rb line 242) had an error: Mixlib::ShellOut::ShellCommandFailed: Command execution failed. STDOUT/STDERR suppressed for sensitive resource`

Solution	Ignore the cert renewal for the respond server service and go ahead with other cert renewal process of other service present in the stack. **Note:** You might see other devices are offline status on UI with temporary cert-mismatch after reboot Admin server in this case if you reissue the service certificate "One at a time" way. However, you can also proceed to cert renewal with other present node-x services.

NetWitness® Platform XDR

Version 12.0

Service Configuration Properties Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Introduction **8**

Admin-server Configuration **9**

ContextConfigProperties 9

UsageTrackingConfigurationProperties 9

SamlProperties 9

WebSocketMessagingProperties 11

Analysis-server Configuration **12**

AnalysisProperties 12

NwDbProperties 12

AnalysisProperties$Temporal 13

Config-server Configuration **14**

ConfigServerProperties 14

Content-server Configuration **14**

DeployProperties 14

LogDeviceParserProperties 14

Contexthub-server Configuration **16**

LiveConfig 16

LiveConnectPathConfig 16

ServiceProperties 16

ServiceDataProperties 17

GlobalQueryResponseCachePolicy 17

AsyncServiceProperites 18

ReputationProperties 19

StixProperties 20

TaxiiServiceProperties 20

HttpProxyProperties 21

rryCorrelation-server Configuration **22**

AlertProperties 22

ContextHubProperties 23

DataPrivacyProperties 23

DebugProperties 24

EndpointProperties 24

EngineProperties 24

EsperProperties 25

FileMapProperties 26

GeoIpProperties 26
HealthProperties 26
MetricProperties 27
ServiceProperties 27
RuleProperties 27
StatsProperties 28
StreamProperties 28
MigrationProperties 32
RecordStreamMetrics 32
Endpoint-broker-server Configuration 33
AggregationProperties 33
ExportProperties 33
FileContextProperties 34
Endpoint-server Configuration 35
CertificateStatusProperties 35
AgentCommandProperties 35
DataRetentionProperties 36
DownloadedDataRetentionProperties 37
InactiveMachineRetentionProperties 37
DataStoreHealthProperties 37
DataStoreProperties 38
RepositoryProperties 38
FileDownloadProperties 38
ExecutionRetryProperties 40
ExportProperties 40
FileDownloadDiskHealthProperties 41
FileCacheProperties 41
FileReputationStatusProperties 41
RiskScoreProperties 42
FileContextProperties 42
FileStatusProperties 42
GroupPolicyProperties 43
MachineFileProperties 43
MachineServiceProperties 44
MachineFileScoreConfigurationProperties 45
MetaForwardProperties 45
PackagerProperties 46
MachineDataHandlerProperties 46
QueueFileSystemPersistenceProperties 47
RelayCommunicationProperties 47

RelayInstallerProperties 48
RelayMetricsProperties 49
SslContextProperties 49
ThrottlingConfigurationProperties 49
UdpProperties 50
Enrichment-server Configuration 51
EngineProperties 51
StreamProperties 51
Integration-server Configuration 52
DeploymentNotificationProperties 52
Investigate-server Configuration 54
AliasesProperties 54
ColumnGroupProperties 54
EventAnalysisProperties 54
IncidentProperties 55
KeyrefsProperties 55
MetaKeyCacheProperties 55
ReconstructionProperties 55
ResponseProperties 59
EventsStreamProperties 59
Launch-framework Configuration 60
ConfigurationModuleProperties 60
ContentProperties 60
DataProperties 60
FileSystemProperties 62
HealthCheckProperties 62
LoggingAuditProperties 62
LogForwarderProperties 63
LoggingProperties 63
LoggingOperationalProperties 63
MetricsAggregationProperties 64
MetricsElasticProperties 64
MetricsHistoricalProperties 64
MetricsJmxProperties 65
MetricsProperties 65
NotificationProperties 65
ProcessJvmMemoryProperties 66
ProcessProperties 66
AuthenticationProperties 67
AuthorizationProperties 67

CertificateAuthorityProperties 68
PkiProperties 68
ServiceAccountProperties 69
TransportBusProperties 70
TransportBusSubscriptionProperties 71
TransportProperties 71
TransportHttpProperties 71
License-server Configuration 73
LicenseProperties 73
FneProperties 73
Metrics-server Configuration 76
MetricsProperties 76
AlertProcessingProperties 76
ElasticRetentionProperties 77
ElasticProperties 77
ElasticClientCertAuthenticatorProperties 78
ElasticJwtAuthenticatorProperties 79
ElasticServiceProperties 80
KibanaProperties 80
MetricsRetryProperties 80
TelemetryProperties 81
MongoDataRetentionProperties 81
Node-infra-server Configuration 82
AdminNodeTrackerProperties 82
ChefRunProperties 82
No-op-server Configuration 84
TestProperties 84
Orchestration-server Configuration 85
HostProperties 85
SaltClientProperties 85
ChefConfigurationProperties 86
ProvisionHostProperties 87
DeploymentProperties 87
TaskExecutionProperties 87
Relay-server Configuration 89
NchanProperties 89
RelayPkiProperties 89
Respond-server Configuration 90
MigrationProperties 90

RespondPrimaryProperties 90
AlertRuleProperties 90
ArcherIntegrationProperties 91
RespondCacheProperties 91
DataRetentionConfiguration 91
IndicatorAggregationJobConfig 92
IntegrationExportProperties 92
NormalizationProperties 93
QueryProperties 94
RiskProcessingProperties 94
RiskCachingProperties 95
RiskRetentionProperties 96
RespondScheduledJobsProperties 96
SecurIdIntegrationProperties 96
Security-server Configuration 98
MigrationProperties 98
PasswordPolicyProperties 98
PkiAuthenticationProperties 99
AuthenticationPolicyProperties 99
DeploymentProperties 99
SamlUserAccountProperties 100
Source-server Configuration 101
FeatureProperties 101
UsmProperties 101

Introduction

This guide provides descriptions and consolidated information for each NetWitness Platform management service configuration property in one place, for example, for services like security-server, investigate-server, and correlation-server. All of these configuration properties have valid default values wherever applicable and do not need to be changed. Please use caution if you need to modify any of these values.

Note: Some of the service configuration values are set during the deployment of NetWitness Platform and are not reflected in this guide. If any of these values need to be overridden, the user must have Administrator privileges.

You can use the NetWitness nw-shell utility or the NetWitness Platform user interface to modify these values. For information about how to use the nw-shell utility, see the *Shell User Guide for NetWitness Platform*. For information about how to use the user interface to modify these values, see the *Hosts and Services Getting Started Guide for NetWitness Platform* and the *Deployment Guide for NetWitness Platform*.

Admin-server Configuration

ContextConfigProperties

Name	Default value	Type	Description
rsa.admin-server.contexthub.enabled	false	boolean	Context Hub integration in Admin Server
rsa.admin-server.contexthub.host		string	
rsa.admin-server.contexthub.port	0	integer	
rsa.admin-server.contexthub.query-timeout		seconds	timeout time for async context response query

UsageTrackingConfigurationProperties

Name	Default value	Type	Description
rsa.admin.ceip-viewed	false	boolean	CEIP(customer enhancement improvement program) tracking value, disabled by default
rsa.admin.usage-tracking-enabled	true	boolean	Usage tracking configuration value, enabled by default

SamlProperties

Name	Default value	Type	Description
rsa.security.authentication.web.saml.auto-lookup-idp-metadata	false	boolean	Controls remote lookup of IDP metadata
rsa.security.authentication.web.saml.default-idp		identityprovider metadata$identityprovider	SAML IDP to be used by default

Name	Default value	Type	Description
rsa.security.authentication.web.saml.entity-id		string	A globally unique identifier used to identify this deployment of Netwitness as a client entity in the identity provider. Every SAML message contains the entity ID.
rsa.security.authentication.web.saml.global-logout-enabled	false	boolean	Flag to determine if global logout request has to be sent to IDP
rsa.security.authentication.web.saml.idp-metadata-url		string	URL to fetch IDP metadata
rsa.security.authentication.web.saml.keystore-alias	nw-saml	string	Alias for storing signing/encryption keys in the SAML keystore
rsa.security.authentication.web.saml.metadata-reload-interval	15 minutes	seconds	Time interval between reloading IDP metadata. Defaults to 15 minutes
rsa.security.authentication.web.saml.saml-response-skew	60 minutes	seconds	Sets the maximum allowed difference between the clocks of the IDP and SP systems. Defaults to 60 seconds.
rsa.security.authentication.web.saml.sp-metadata-filename	nw_saml_metadata.xml	string	Default filename to be set when exporting the service provider metadata

Name	Default value	Type	Description
rsa.security.authentication.web.saml.sso-enabled	false	boolean	Flag to enable or disable SAML based SSO authentication
rsa.security.authentication.web.saml.trust-all-certs-for-idp-metadata	false	boolean	Flag to ignore certificate verification while downloading IDP metadata from the given URL
rsa.security.authentication.web.saml.use-proxy	false	boolean	Determines if requests to IDP has to be routed through a proxy (if configured)

WebSocketMessagingProperties

Name	Default value	Type	Description
rsa.websocket.messaging.batch-size	10	long	The number of entries to send, per message.

Analysis-server Configuration

AnalysisProperties

Name	Default value	Type	Description
rsa.analysis.auto-analyze	true	boolean	{@code true} to automatically start analyze
rsa.analysis.definition-files		set	{@link Set} of definition files to be analyzed
rsa.analysis.nws	false	boolean	Enable NWS device support.
rsa.analysis.periodic-task-wait-time	1	seconds	Time period to wait between tasks.
rsa.analysis.time-allowance	15	seconds	Time to wait before issuing the query. Ex: If the time frame is 9-10pm and time allowance is 15 minutes, query at 10:15pm.
rsa.analysis.time-frame-interval	1	seconds	Time Interval to use for each analysis time frame as we advance forward

NwDbProperties

Name	Default value	Type	Description
rsa.analysis.nwdb.host	localhost	string	NetWitness Core Query Broker host address
rsa.analysis.nwdb.password		string	NetWitness Core connection password
rsa.analysis.nwdb.port	50005	integer	NetWitness Core Query Broker port
rsa.analysis.nwdb.ssl	false	boolean	NwtWitness Core SSL flag
rsa.analysis.nwdb.start-time		string	The start time to analyze with the format of "yyyy-MM-dd HH:mm" UTC.
rsa.analysis.nwdb.user	admin	string	NetWitness Core connection user

AnalysisProperties$Temporal

Name	Default value	Type	Description
rsa.analysis.temporal.enabled	true	boolean	Use temporal services instead of Repositories
rsa.analysis.temporal.max-entries	1000	integer	Maximum number of {@link java.util.Map} entries.
rsa.analysis.temporal.retention-period	30	seconds	The period of time to retain the data.

Config-server Configuration

ConfigServerProperties

Name	Default value	Type	Description
rsa.configuration.notify-delay	5	seconds	The length of time to delay notifications to wait for more changes for an identity

Content-server Configuration

DeployProperties

Name	Default value	Type	Description
rsa.content-server.deploy.deploy-timeout	2	seconds	Max time allotted to deploy a single content to an endpoint Default is 2 minutes Minimum of 30 seconds Maximum of 1 hour
rsa.content-server.deploy.max-content-deployer-threads	4	integer	The size of the content deployer thread pool
rsa.content-server.deploy.max-policy-deployer-threads	2	integer	The size of the policy deployer thread pool

LogDeviceParserProperties

Name	Default value	Type	Description
rsa.content.parser.cache-duration	24	seconds	Time it takes for the cache that stores data to expire
rsa.content.parser.cache-duration-for-service-status	5	seconds	Time it takes for the cache that stores service status data to expire

Name	Default value	Type	Description
rsa.content.parser.database-initialization-enabled	true	boolean	Enable database initialization task
rsa.content.parser.json-feature	true	boolean	Feature toggle for whether JSON data types are supported.
rsa.content.parser.log-decoder-sync-interval	12 hours	seconds	The interval between log decoder sync sweeps.
rsa.content.parser.max-try-counter-for-ld-sync-failure	3	integer	Switch to a new Log decoder from environment after trying below configured times to previously/user configured LD
rsa.content.parser.preferred-log-decoder-name-for-sync		string	Display name of the preferred log decoder for sync job. Sync job will use this only if it is online and contains parser details
rsa.content.parser.previously-synced-log-decoder-id		string	ID of the previously synced log decoder
rsa.content.parser.remove-previous-sync-parsers-for-new-log-decoder	false	boolean	Flag to indicate that remove previously synced log device parser information in the case of new log decoder found.
rsa.content.parser.retrieval-timeout	30	seconds	Timeout to wait for nextgen response
rsa.content.parser.sleep-interval-after-notification	60	seconds	Sleep interval before triggering a task after receiving notification
rsa.content.parser.sync-from-log-decoder-enabled	true	boolean	Enable automatic sync of parsers task

Contexthub-server Configuration

LiveConfig

Name	Default value	Type	Description
rsa.cms.client.host		string	
rsa.cms.client.password		string	
rsa.cms.client.port	0	integer	
rsa.cms.client.use-ssl	false	boolean	
rsa.cms.client.username		string	

LiveConnectPathConfig

Name	Default value	Type	Description
rsa.cms.client.lc-auth-path	/authlive/authenticate/LIVECONNECT	string	
rsa.cms.client.lc-feedback-path	/liveconnect/v2/feedback/meta	string	

ServiceProperties

Name	Default value	Type	Description
rsa.contexthub.backup-data-path		string	Migration data backup location
rsa.contexthub.config-dir-path		string	Config directory which contains all the configuration files
rsa.contexthub.data-dir-path		string	Data Directory
rsa.contexthub.file-system-service		filesystemservice	
rsa.contexthub.jobs-dir-path		string	Jobs directory which contains all the job configurations

Name	Default value	Type	Description
rsa.contexthub.max-entries-for-list	100000	integer	List datasource max limit
rsa.contexthub.ootb-list-version		string	OOTB List Version
rsa.contexthub.prefetch-pool-size	3	integer	Prefetch job pool size
rsa.contexthub.replace-config	false	boolean	Replace the batch config files on service boot
rsa.contexthub.templates-dir-path		string	Template directory which contains all the templates
rsa.contexthub.tried-adding-respond-server	false	boolean	Now when the service boots-up and if there is Respond-Server already in the deployment, we should try to add that as a CH data-source, but this should be done already once, so in case user deletes the source we shouldn't add it again.

ServiceDataProperties

Name	Default value	Type	Description
rsa.contexthub.data.disk-size	120	bytes	Max database disk space allocated for the Contexthub service.
rsa.contexthub.data.used-disk-upper-threshold	95	double	

GlobalQueryResponseCachePolicy

Name	Default value	Type	Description
rsa.contexthub.query-response-cache.available-memory	0	long	
rsa.contexthub.query-response-cache.cache-name		string	

Name	Default value	Type	Description
rsa.contexthub.query-response-cache.cache-store-bulk-insert-size	20	integer	
rsa.contexthub.query-response-cache.enabled	true	boolean	
rsa.contexthub.query-response-cache.max-seconds-in-cache	1800	seconds	
rsa.contexthub.query-response-cache.modification-queue	20	integer	
rsa.contexthub.query-response-cache.percentage-of-heap-as-cache	50	double	
rsa.contexthub.query-response-cache.preload	true	boolean	
rsa.contexthub.query-response-cache.thread-pool	2	integer	
rsa.contexthub.query-response-cache.used-cache-upper-threshold	100	double	

AsyncServiceProperites

Name	Default value	Type	Description
rsa.contexthub.query-threads.core-pool-size	20	integer	
rsa.contexthub.query-threads.max-pool-size	250	integer	
rsa.contexthub.query-threads.max-seconds-before-results-expire	0	long	
rsa.contexthub.query-threads.queue-capacity	1000	integer	

ReputationProperties

Name	Default value	Type	Description
rsa.contexthub.reputation.batch-size	1000	integer	Size of the batch to be send to Reputation Server
rsa.contexthub.reputation.max-hashes-to-be-queried	595000	integer	Maximum number of hashes to be considered for refreshing in 1 day
rsa.contexthub.reputation.max-query-supported-by-reputation-server	600000	integer	Maximum number of queries supported by the live reputation service.
rsa.contexthub.reputation.max-staged-count-for-refresh	1	long	Maximum staged entries in staging store if present, prefetch will be retried in refreshCheckInterval seconds. This is done to give priority to new reputation queries. Eg: Prefetch starts at 9 PM on a day. CH checks whether the "Staged" entries in Staging store are less than maxStagedCountForRefresh. If it is less prefetch starts else prefetch is skipped for this time and retried in refreshCheckInterval seconds.
rsa.contexthub.reputation.preferred-hashing-algorithm	md5	string	Algorithms that should be used while interaction with Reputation Server
rsa.contexthub.reputation.refresh-batch-interval	15	seconds	Time Interval between 2 batches sent to reputation server for refresh
rsa.contexthub.reputation.refresh-check-interval	15	seconds	Configuration to handle "Staging Store has entries" or Any other exception/error in case of refresh (prefetch). Staged entries should get priority over refresh job. And in case of any failures CH should retry refresh after this interval. Eg: At t1, CH started refresh job but finds that there are some entries in staging store with status - "Staged" CH will check if the no. of entries are > minStagedCountForRefresh, and if thats the case it will retry refresh after seconds configured here until that day’s 11:55 PM UTC

Name	Default value	Type	Description
rsa.contexthub.reputation.reputation-query-batch-interval	2	seconds	If CH does not gets any batch from staging store for RS, next time it queries staging store for a batch is after the seconds configured here. Eg: At t1, CH got a batch from staging store and RS was queried At t2, CH did not get any batch from staging store Now at t3 (t2 + reputationQueryBatchInterval) , CH again queried staging store for a batch

StixProperties

Name	Default value	Type	Description
rsa.contexthub.stix.data-store-read-page-size	100	integer	When reading the entire data store of a STIX source, this property determines the page size

TaxiiServiceProperties

Name	Default value	Type	Description
rsa.enrichment.stix.config.disabled-xml-features		list	
rsa.enrichment.stix.config.max-taxii-poll-window	7	seconds	Maximum time range to query TAXII server in one cycle. Defaults to 7 days. E.g., If the total range to query the TAXII server is 30 days, the range will be divided into shorter time windows of 7 days each.
rsa.enrichment.stix.config.taxii-service-max-attempts	2	integer	Max number of attempts to query TAXII Service
rsa.enrichment.stix.config.taxii-service-retry-wait-in-sec	10	integer	Max number of attempts to query TAXII Service

HttpProxyProperties

Name	Default value	Type	Description
rsa.transport.http.proxy.enabled	false	boolean	
rsa.transport.http.proxy.host		string	
rsa.transport.http.proxy.ntlm-domain		string	
rsa.transport.http.proxy.password		string	
rsa.transport.http.proxy.port		integer	
rsa.transport.http.proxy.reinitialize-proxy	false	boolean	
rsa.transport.http.proxy.use-ntlm-auth	false	boolean	
rsa.transport.http.proxy.use-ssl	false	boolean	Flag indicating whether we should use HTTP or HTTPS
rsa.transport.http.proxy.user		string	

rryCorrelation-server Configuration

AlertProperties

Name	Default value	Type	Description
rsa.correlation.alert.keep-alive-time	0	long	The keepAlive time for threads
rsa.correlation.alert.max-alerts-queue-size	10000	integer	The max rabbitmq alerter queue size
rsa.correlation.alert.num-threads	3	integer	No. of threads to process
rsa.correlation.alert.respond-enabled	true	boolean	The respond is enabled globally
rsa.correlation.alert.respond-endpoint-severities		list	The list of severities which can be consumed by respond and are related to app-rules
rsa.correlation.alert.retry-delay	1	seconds	retry time for each interval
rsa.correlation.alert.risk-score-severities		list	The list in severities which can be consumed by risk score and are related to app-rules
rsa.correlation.alert.sleep-time	1000	long	The max time to sleep in thread
rsa.correlation.alert.statement-name-max-length	128	integer	The maximum length of the entire statement @Name
rsa.correlation.alert.statement-name-place-holder-max-length	64	integer	The maximum length for each place holder value in the statement @Name
rsa.correlation.alert.timeout-retry-policy	3650	seconds	Retry time in seconds for total timeout
rsa.correlation.alert.total-threads	10	integer	The total number of threads in the pool
rsa.correlation.alert.transient-enabled	true	boolean	The transient is enabled globally. Currently used only for key-value rule and not in basic rule

ContextHubProperties

Name	Default value	Type	Description
rsa.correlation.contexthub.data-expired-in-seconds	5	integer	The duration of time before the ContextHub content is too old and need to be re-retrieved.
rsa.correlation.contexthub.fail-on-retrieve-retry-count	3	integer	Number of times to retry when failed to retrieve data from ContextHub.
rsa.correlation.contexthub.fail-on-retrieve-wait-between-retries	5	seconds	Wait duration between retries when failed to retrieve data from ContextHub.
rsa.correlation.contexthub.fail-on-set-entries-wait-between-retries	5	seconds	Wait duration between retries when failed to add/delete entries to/from ContextHub.
rsa.correlation.contexthub.file-backed-dir		string	Location on local disk where to store the paged files.
rsa.correlation.contexthub.mapped-memory-size	0	integer	Total number of bytes of data that are kept in memory.
rsa.correlation.contexthub.notification-handler-thread-pool-size	8	integer	Number of concurrent notification handler threads.
rsa.correlation.contexthub.page-file-size	4096	integer	The size of each paged file stored on local disk.
rsa.correlation.contexthub.set-entries-thread-pool-size	128	integer	Number of RSAContext set entries concurrent Threads pool size.

DataPrivacyProperties

Name	Default value	Type	Description
rsa.correlation.data-privacy.global-private-fields		list	List of fields that are always removed from the output for data privacy, regardless of source

DebugProperties

Name	Default value	Type	Description
rsa.correlation.debug.actions		string	
rsa.correlation.debug.enabled	false	boolean	
rsa.correlation.debug.resource-ids		string	

EndpointProperties

Name	Default value	Type	Description
rsa.correlation.endpoint.app-rules-paths		list	{@link List} of Endpoint App Rules candidate paths of the resource file.
rsa.correlation.endpoint.enabled	true	boolean	{@code true} if Endpoint Rules processing is enabled.

EngineProperties

Name	Default value	Type	Description
rsa.correlation.engine.auto-start	true	boolean	Determines if all {@link Engine} should start on service deployed.
rsa.correlation.engine.concurrent-deployment	10	integer	Number of asynchronous Engine deployment Tasks.
rsa.correlation.engine.send-event-heart-beat-frequency	1	seconds	Log send Event heartbeat frequency.
rsa.correlation.engine.startup-error-retry-interval	10	seconds	Retry interval if error occurs during startup.

EsperProperties

Name	Default value	Type	Description
rsa.correlation.esper.background-metrics-enabled	true	boolean	Set to {@code false} to get Esper metrics on demand.
rsa.correlation.esper.background-metrics-frequency	5	seconds	How often should the background Esper metrics process should be performed.
rsa.correlation.esper.config-resource	classpath:esper/esper-config.xml	string	Esper Configuration xml Resource.
rsa.correlation.esper.enable-statement-metric	false	boolean	Set true if esper Metrics needs to be enabled. By default it is set to false by Esper. Making it true will allow to capture additional esper-metrics but note that activating Esper metrics may cause performance impacts
rsa.correlation.esper.metrics-memory-back-off	1	seconds	How long to back off for after reaching a metrics timeout error.
rsa.correlation.esper.metrics-num-threads	16	integer	The number of threads to use for calculating metrics, per engine. Each thread gets metrics for a single rule.
rsa.correlation.esper.metrics-timeout	15	seconds	How long we should allow for retrieval of metrics for a single rule. Counting memory for rules that are using a lot of memory takes a lot of time and cpu that blocks processing of new events. <p> In the case of a timeout, we will capture the error for reporting purposes.
rsa.correlation.esper.snapshot-dir		string	RSAPersist snapshot directory.
rsa.correlation.esper.snapshot-frequency	5	seconds	Taking snapshot periodic duration.
rsa.correlation.esper.use-external-clock	true	boolean	{@code true} for Esper to process CurrentTimeEvent.

FileMapProperties

Name	Default value	Type	Description
rsa.correlation.filemap.file-backed-dir		string	Location on local disk where to store the paged files.
rsa.correlation.filemap.page-file-size	4096	integer	The size of each paged file stored on local disk.
rsa.correlation.filemap.total-memory-size	0	integer	Total number of bytes of data that are kept in memory.

GeoIpProperties

Name	Default value	Type	Description
rsa.correlation.geoip.city-resource		string	The City database Resource.
rsa.correlation.geoip.local-dir		string	Local store folder where to store the database files.
rsa.correlation.geoip.org-resource		string	The Organization database Resource.

HealthProperties

Name	Default value	Type	Description
rsa.correlation.health.check-every	15	seconds	The
rsa.correlation.health.fatal-percentage	90	integer	The percentage of memory consumption at which it is considered to be in fatal state
rsa.correlation.health.health-check-id	memory-check	string	The name which is required to set the HealthCheck
rsa.correlation.health.warning-percentage	80	integer	The percentage of memory consumption at which it is considered to be Warning in Warning state

MetricProperties

Name	Default value	Type	Description
rsa.correlation.metric.collectd-max-value-length	64	integer	CollectD field value maximum length.

ServiceProperties

Name	Default value	Type	Description
rsa.correlation.re-deployment-cycle	0	integer	The current re-deployment cycle.
rsa.correlation.re-deployment-required	0	integer	The number of re-deployment required.
rsa.correlation.send-re-deployment-notification	true	boolean	{@code true} to notify SA to re-deploy all active {@code Engine}s.
rsa.correlation.version		string	Project version.
rsa.correlation.wait-before-checking-for-success-re-deployment	1	seconds	Wait duration before checking to see if SA response re-deployment is successful.

RuleProperties

Name	Default value	Type	Description
rsa.correlation.rule.fired-rules-heart-beat		integer	Number of permits for a duration.
rsa.correlation.rule.fired-rules-heart-beat-every	1	seconds	A length of time to apply the permits. Minimum of 1 second and max at 1 day.
rsa.correlation.rule.log-fired-rules	false	boolean	Should we log the rules as soon as it fired with the relevant events.
rsa.correlation.rule.max-constituent-events	0	integer	Maximum number of Events in the List sent to AlertManager.

StatsProperties

Name	Default value	Type	Description
rsa.correlation.stats.days-to-keep-stats-file	3	integer	

StreamProperties

Name	Default value	Type	Description
rsa.correlation.stream.aggregation-queue-size	10	integer	Size of the queue that holds aggregation Events staging them before sending them to Rule Engine.
rsa.correlation.stream.batch-size	0	integer	Controls how many records do we ask for at a time.
rsa.correlation.stream.big-integer-to-long	true	boolean	Choose if we want to convert {@code BigInteger} {@code Meta} value to {@code Long} like sessionid.
rsa.correlation.stream.buffer-size	0	integer	Controls the number of records the stream can keep outstanding.
rsa.correlation.stream.check-supply	false	boolean	Should this source check for supply
rsa.correlation.stream.collection-duration-in-minutes	0	integer	For query based aggregation this parameter determines if it should operate on continuous mode or finite mode. By Default it is 0 which means continuous mode. CollectionDuration should be specified in minutes.
rsa.correlation.stream.compression	0	integer	The number of bytes in each message before it will be compressed. Zero is no compression at all. range:0 to 131071
rsa.correlation.stream.compression-level	0	integer	The level of compression. 1 is fastest and 9 is the best compression. A value of zero means pick the best balance between speed and compression. range:0 to 9
rsa.correlation.stream.connection-time-out	0	integer	Override connection timeout in sources. Only if greater than 0.

Name	Default value	Type	Description
rsa.correlation.stream.default-multi-valued		list	New multi-valued fields for this version. These fields should all be migrated to multi-valued with Rule changes. A warning message will be logged if multi-valued does NOT contain all of these fields.
rsa.correlation.stream.default-single-valued		list	New single-valued fields for this version. These fields should all be migrated to single-valued with Rule changes. A warning message will be logged if single-valued does NOT contain all of these fields.
rsa.correlation.stream.dots-to-underscores	true	boolean	Choose if we want to translate "user.dst" to "user_dst".
rsa.correlation.stream.event-batch-size	1000	integer	Number of Events in a batch store in the queue.
rsa.correlation.stream.event-enrichment-queue-size	10	integer	Size of the queue to be used to enrich the {@code Event} before offer to {@code Rule} {@code Engine}.
rsa.correlation.stream.event-enrichment-thread-pool-size	8	integer	Concurrent Event enrichment Thread pool size.
rsa.correlation.stream.event-polling-timeout-in-milli-seconds	1000	long	Event polling from queue timeout in milliseconds.
rsa.correlation.stream.event-source-id	false	boolean	Controls whether we need to add the event source identifier (ESA compatibility)
rsa.correlation.stream.filter		string	Filter to be sent across to the source
rsa.correlation.stream.idle-retry-interval	0	integer	Controls how long to wait (in milli-seconds) before retrying an idle source.
rsa.correlation.stream.lag-time	15	seconds	Lag time is the expected time an event takes to pass through the different levels of capture/parse etc and become available to query in the concentrator.
rsa.correlation.stream.lowercase		list	Choose if the fields to translate to lower case

Name	Default value	Type	Description
rsa.correlation.stream.max-sessions	0	integer	Controls the number of sessions in a batch. The more you filter out ESA data source traffic, the lower you should set this value.
rsa.correlation.stream.mechanism		string	NextGen core devices send and receive type 'AGGREGATION' or 'QUERY'.
rsa.correlation.stream.minutes-back	5	integer	Controls how far back in time should we go for a fresh start.
rsa.correlation.stream.multi-valued		list	Choose the fields considered as multi-valued.
rsa.correlation.stream.multi-valued-as-array	false	boolean	{@code true} to convert multi-valued Collection to Array.
rsa.correlation.stream.no-system-meta	false	boolean	Controls the addition of system meta to records.
rsa.correlation.stream.pre-fetch	0	integer	Controls how many batches to pull and keep ready in anticipation of demand
rsa.correlation.stream.query		string	Query Based RecordStream select clause for all sources.
rsa.correlation.stream.reader-buffer-size	1048576	integer	
rsa.correlation.stream.retrieve-record-stream-stats-every	2	seconds	How often should the {@code RecordStream} status be retrieved.
rsa.correlation.stream.retrieve-schema-every	5	seconds	A length of time to apply the permits. Minimum of 1 second and max at 1 day.
rsa.correlation.stream.retrieve-schema-frequency	1	integer	Number of permits for a duration.
rsa.correlation.stream.retry-timeout	0	integer	Controls how long to wait (in milli-seconds) before retrying a failed source.
rsa.correlation.stream.save-position-every	1	seconds	A length of time to apply the permits. Minimum of 1 second and max at 1 day.
rsa.correlation.stream.save-position-frequency	1	integer	Number of permits for a duration.

Name	Default value	Type	Description
rsa.correlation.stream.single-valued		list	Uses by Rules deployment process to ensure that these fields are not be treated as multi-valued.
rsa.correlation.stream.socket-timeout	0	integer	Override socket timeout in sources. Only if greater than 0.
rsa.correlation.stream.source-poll-interval	0	integer	Controls the parameters passed to {@code RecordSourceSubscription}.
rsa.correlation.stream.start-session-id	0	long	Override StartSession Id in sources for debug purposes. Only if greater than 0.
rsa.correlation.stream.tcp-no-delay	false	boolean	
rsa.correlation.stream.thread-pool-size	0	integer	Controls the size of the thread pool used the stream executor. Default to 100.
rsa.correlation.stream.time-batch-in-seconds	0	integer	Determines the batch size for the query based aggregation in seconds. By default it will be a 60 second window. This for now will not be configurable for user. This is because concentrator operates most efficiently when the time window is a minute.
rsa.correlation.stream.time-measured-in-seconds	true	boolean	{@code true} if time meta is measured in seconds in the event.
rsa.correlation.stream.time-meta-field	time	string	Decides what field should be used for time.
rsa.correlation.stream.time-order-by-field		string	Controls the name of the field that we consider the timestamp. This must be a long value.
rsa.correlation.stream.time-order-hold-interval	0	integer	To order records from multiple sources, we need to allow some "hold" time for sessions within a time window to arrive from all sources. This parameter specifies the hold interval (in milli-seconds)
rsa.correlation.stream.time-order-no-inflow-give-up-interval	0	integer	Controls the interval (in milliseconds) after which we take a "quiet" source out of the equation to allow progress on a time ordered stream. The default value is 0, which implies that we wait forever for events to arrive.

Name	Default value	Type	Description
rsa.correlation.stream.time-order-offline-give-up-interval	0	integer	Controls the interval (in milli-seconds) after which we take an offline source out of the equation to allow progress on a time ordered stream. The default value is 0, which implies that we wait forever. This parameter does not affect the re-connection retries; those which are performed in all cases.
rsa.correlation.stream.time-ordered	false	boolean	Enables source time synchronization and ordering.
rsa.correlation.stream.use-direct-buffer	false	boolean	
rsa.correlation.stream.use-event-time-for-esper	false	boolean	{@code true} to use the timeMetaField in the Event for Esper CurrentTimeEvent.

MigrationProperties

Name	Default value	Type	Description
rsa.migration.home-data-path	/var/netwitness/esa	string	The location of ESA home directory

RecordStreamMetrics

Name	Default value	Type	Description
rsa.records.stream.version		string	

Endpoint-broker-server Configuration

AggregationProperties

Name	Default value	Type	Description
rsa.endpoint.broker.cache-enabled	true	boolean	Cache enabled
rsa.endpoint.broker.cache-entry-limit	1000	integer	Cache entries limit
rsa.endpoint.broker.fetched-machines-limit	100	integer	The number of machine infos fetched for a given checksum. This is used to fetch the top 'n' risky machine-infos for a given file.
rsa.endpoint.broker.last-accessed-interval	120	seconds	Time interval when server would fetch the accessible endpoint servers in active session
rsa.endpoint.broker.last-updated-interval	600	seconds	Time interval when server would fetch the accessible endpoint servers when no request is sent for a long time
rsa.endpoint.broker.max-response-pages	10	integer	Max pages that an be viewed in UI host and file list views.
rsa.endpoint.broker.ping-timeout	2	seconds	Max time server will wait for the ping query to complete
rsa.endpoint.broker.query-timeout	10	seconds	Max timeout for individual endpoint server query to complete
rsa.endpoint.broker.threads	30	integer	Max number of request handler threads

ExportProperties

Name	Default value	Type	Description
rsa.endpoint.broker.export.directory-context	ExportDirectory	string	Represents the directory context (reference name) for the files to be exported
rsa.endpoint.broker.export.file-cleanup-interval	1800	seconds	Schedule interval for cleanup of files/directories
rsa.endpoint.broker.export.file-expiration-time	3600	seconds	Expiration time for the file(s) created

Name	Default value	Type	Description
rsa.endpoint.broker.export.path-prefix	temp/export	string	Represents the path prefix for files to be exported

FileContextProperties

Name	Default value	Type	Description
rsa.endpoint.broker.file-search.timeout	30	seconds	File context keyword search operation time out in seconds

Endpoint-server Configuration

CertificateStatusProperties

Name	Default value	Type	Description
rsa.endpoint.certificate.status.ignored-notifications-retry-interval	60	seconds	Notifications are ignored once posting file status fails. These ignored notifications are queried periodically. The property defines the interval.
rsa.endpoint.certificate.status.new-files-query-for-automatic-status-interval	300	seconds	Time (in seconds) between subsequent querying of new files for automatic assignment of file status to be send to Contexthub server
rsa.endpoint.certificate.status.query-batch-size	3000	integer	* Max number of thumbprints those should be fetched from repository in a single query
rsa.endpoint.certificate.status.request-batch-size	500	integer	* Max number of thumbprints those should be part of the request sent to Contexthub-Server
rsa.endpoint.certificate.status.request-interval	300	seconds	Time (in seconds) between querying for any new Certificates seen in endpoint server Defaulting to 5 minutes

AgentCommandProperties

Name	Default value	Type	Description
rsa.endpoint.command.cache-reload-delay	5	seconds	Interval delay to reload pending commands cache
rsa.endpoint.command.cancel-interval	24	seconds	Interval to cancel expired commands
rsa.endpoint.command.expiration-count	5	integer	Indicates the maximum number of times command would be resent to agent(s)

Name	Default value	Type	Description
rsa.endpoint.command.expiration-time	20	seconds	Indicates the duration until when command will not be resent to agent(s)

DataRetentionProperties

Name	Default value	Type	Description
rsa.endpoint.config.data-retention.enabled	true	boolean	Indicates if all machine data older than configured threshold {@code #thresholdInDays}, is to be deleted. This is enabled by default.
rsa.endpoint.config.data-retention.initial-rollover-delay	1	seconds	Time to delay before the first execution of the storage size based retention job
rsa.endpoint.config.data-retention.recurrence-interval	0 0 0 * * *	string	Indicates the time and frequency to run the deletion task. Configured to run everyday at 12:00:00 AM, by default.
rsa.endpoint.config.data-retention.rollover-after	80	double	The threshold (in %) indicating the storage size used, after which data should be cleaned up from the database
rsa.endpoint.config.data-retention.rollover-chunk-size	10	double	The chuck of data that should be cleanup up from the database. For example, 10 indicates 10% of the data should be cleaned up. Used for storage size based data retention job.
rsa.endpoint.config.data-retention.rollover-delay	10	seconds	Delay between invocations of the storage size based retention job
rsa.endpoint.config.data-retention.size-based-rollover-enabled	true	boolean	Indicates if storage size based retention job is enabled. This involves clearing up the disk, if it reaches a certain threshold {@see #rolloverAfter}. This is enabled by default.
rsa.endpoint.config.data-retention.threshold-in-days	30	integer	The retention threshold specified (in days)

DownloadedDataRetentionProperties

Name	Default value	Type	Description
rsa.endpoint.config.downloaded-data-retention.enabled	true	boolean	Indicates retention active status. This is enabled by default.
rsa.endpoint.config.downloaded-data-retention.recurrence-interval	0 0 0 * * *	string	Indicates the time and frequency to run the deletion task. Configured to run everyday at 00:00:00 AM, by default.
rsa.endpoint.config.downloaded-data-retention.threshold-in-days	90	integer	The retention threshold specified (in days)

InactiveMachineRetentionProperties

Name	Default value	Type	Description
rsa.endpoint.config.inactive-machine-retention.enabled	true	boolean	Indicates if all machines inactive for more than the configured threshold {@code #thresholdInDays}, is to be deleted. This is enabled by default.
rsa.endpoint.config.inactive-machine-retention.recurrence-interval	0 0 1 * * *	string	Indicates the time and frequency to run the deletion task. Configured to run everyday at 01:00:00 AM, by default.
rsa.endpoint.config.inactive-machine-retention.threshold-in-days	90	integer	The retention threshold specified (in days)

DataStoreHealthProperties

Name	Default value	Type	Description
rsa.endpoint.data-store-thresholds.fatal-percent	95	integer	
rsa.endpoint.data-store-thresholds.warning-percent	85	integer	

DataStoreProperties

Name	Default value	Type	Description
rsa.endpoint.data.application.compression-factor	2.5	double	Indicates the compression ratio used by mongo while writing to the filesystem
rsa.endpoint.data.application.db-path		string	Specify the path/directory allocated for the database files. Assumed to be /var/netwitness/mongo by default

RepositoryProperties

Name	Default value	Type	Description
rsa.endpoint.datastore.index-creation-enabled	true	boolean	Determines whether the indexes should be be created on the service startup

FileDownloadProperties

Name	Default value	Type	Description
rsa.endpoint.download.agent-beacon-threshold	5	seconds	Indicates the agent beacon time considered to (re)attempt file download
rsa.endpoint.download.base-path		string	Path in endpoint server where downloaded files are stored. Assumed to be /var/netwitness/endpoint-server by default
rsa.endpoint.download.batch-size	1000	integer	Number of entries to fetch & process from {@link CollectionConstants#GLOBAL_FILE_DOWNLOAD_REQUEST_COLLECTION} collection
rsa.endpoint.download.command-expiration-time	20	seconds	Indicates the expiration time for automatic file download commands, after which command would be cancelled

Name	Default value	Type	Description
rsa.endpoint.download.disk-check-interval	5	seconds	Indicates the interval to check the health of disk to which files will be downloaded
rsa.endpoint.download.download-threads	10	integer	Max number of auto file download handler threads
rsa.endpoint.download.downloaded-files-cache-size	2000000	integer	Max number of entries to store as part of downloaded files cache
rsa.endpoint.download.file-processor-batch-size	100	integer	Maximum number of concurrent processing requests that should be handled by server
rsa.endpoint.download.max-attempts	50	integer	Maximum number of agents that will be tried against in order to get the file downloaded to server, following which the next server takes over (if any)
rsa.endpoint.download.max-pending-commands	50	integer	Defines the maximum cap of unprocessed file download commands that can exist for a given agent, i.e. although {@link AgentCommandRequestType#Manual} commands can still be created, it is used to restrict addition of {@link AgentCommandRequestType#Automatic} file download commands
rsa.endpoint.download.periodic-cleanup-delay	2	seconds	Interval between successive lookups and attempts made by the server to delete pending requests which are no longer required
rsa.endpoint.download.periodic-hash-cleanup-delay	1	seconds	Interval between successive lookups and attempts made by the server to delete requests for files identified to be downloaded
rsa.endpoint.download.periodic-marking-delay	5	seconds	Interval between marking requests to be considered for processing by server(s)
rsa.endpoint.download.periodic-processing-delay	1	seconds	Interval between successive lookups and attempts made by server to process pending file download requests, i.e. to create/issue file download commands
rsa.endpoint.download.periodic-retry-processing-delay	5	seconds	Interval between successive lookups and attempts made by server to retry processing of older pending file download requests

Name	Default value	Type	Description
rsa.endpoint.download.request-cache-size	2000000	integer	Max number of entries to store as part of downloaded files request cache
rsa.endpoint.download.threads	2	integer	Max number of request handler threads
rsa.endpoint.download.update-interval	5	seconds	Interval in which downloaded status of newly added files is updated

ExecutionRetryProperties

Name	Default value	Type	Description
rsa.endpoint.execution.retry.file-persistence-delay	50	seconds	Indicates the wait time for retrying file data persistence
rsa.endpoint.execution.retry.max-delay	2	seconds	Indicates the maximum delay to be used between retries
rsa.endpoint.execution.retry.min-delay	30	seconds	Indicates the minimum delay to be used between retries

ExportProperties

Name	Default value	Type	Description
rsa.endpoint.export.directory-context	ExportDirectory	string	Represents the directory context (reference name) for the files to be exported
rsa.endpoint.export.file-cleanup-interval	1800	seconds	Schedule interval for cleanup of files/directories
rsa.endpoint.export.file-expiration-time	3600	seconds	Expiration time for the file(s) created
rsa.endpoint.export.max-exportable-entries	100000	integer	Maximum entries that can be exported into csv from the database, for files
rsa.endpoint.export.path-prefix	temp/export	string	Represents the path prefix for files to be exported

FileDownloadDiskHealthProperties

Name	Default value	Type	Description
rsa.endpoint.file-download-disk-thresholds.fatal-percent	70	integer	
rsa.endpoint.file-download-disk-thresholds.warning-percent	60	integer	

FileCacheProperties

Name	Default value	Type	Description
rsa.endpoint.file.cache.expiration-time	1800	seconds	Expiration threshold, since last access of item(s)
rsa.endpoint.file.cache.size	100000	long	Maximum items in the cache

FileReputationStatusProperties

Name	Default value	Type	Description
rsa.endpoint.file.reputation.ignored-notifications-query-interval	300	seconds	Time (in seconds) between subsequent check for ignored notifications querying
rsa.endpoint.file.reputation.known-signed-providers	microsoft,apple	string	List of signature providers for which we don’t need to compute the reputation. This is only accounted when filterOutKnowFiles = true/
rsa.endpoint.file.reputation.query-batch-size	2000	integer	* Max number of hashes those should be fetched from repository in a single query
rsa.endpoint.file.reputation.request-batch-size	500	integer	* Max number of hashes those should be part of the request sent to Contexthub-Server

Name	Default value	Type	Description
rsa.endpoint.file.reputation.request-interval	10	seconds	Time (in seconds) between subsequent requests to be send to Reputation-Server
rsa.endpoint.file.reputation.skip-known-good-files	true	boolean	Should reputation be computed for files from know sources ? This can be files that are signed by known CA's or maybe what the customer might have configured to be white-listed files

RiskScoreProperties

Name	Default value	Type	Description
rsa.endpoint.file.score.query-batch-size	2000	integer	Max number of file/machines to be fetched from repository in a single query
rsa.endpoint.file.score.request-interval	20	seconds	Time (in seconds) between subsequent requests to be sent

FileContextProperties

Name	Default value	Type	Description
rsa.endpoint.file.search.timeout	30	seconds	File context keyword search operation time out in seconds
rsa.endpoint.file.search.total-count	100	integer	Max number of results that will be returned for a/any snapshot response

FileStatusProperties

Name	Default value	Type	Description
rsa.endpoint.file.status.ignored-notifications-query-interval	300	seconds	Time (in seconds) between subsequent check for ignored notifications querying

Name	Default value	Type	Description
rsa.endpoint.file.status.query-batch-size	3000	integer	* Max number of hashes those should be fetched from repository in a single query
rsa.endpoint.file.status.request-batch-size	500	integer	* Max number of hashes those should be part of the request sent to Contexthub-Server
rsa.endpoint.file.status.request-interval	10	seconds	Time (in seconds) between subsequent requests to be send to Reputation-Server

GroupPolicyProperties

Name	Default value	Type	Description
rsa.endpoint.group-policy.bulk-write-count	1000	integer	Number of items to be written as part of a batch/bulk write operation performed, to assign/update group-policy to machines present in the deployment
rsa.endpoint.group-policy.initial-sync-delay	20	seconds	Time to wait for the initial group-policy details to be synced
rsa.endpoint.group-policy.periodic-evaluation-delay	30	seconds	Interval between successive evaluations performed (if required), to assign/update group-policy to machines present in the deployment

MachineFileProperties

Name	Default value	Type	Description
rsa.endpoint.machine-file.delete-task-delay	5	seconds	Initial delay to clean-up {@link CollectionConstants#MACHINE_FILE_ COLLECTION} collection for un-managed agents and decrement host count
rsa.endpoint.machine-file.fetch-limit	50	integer	Number of documents to be fetched from {@link CollectionConstants#MACHINE_ FILE_STAGE_COLLECTION} collection amd merge to the {@link CollectionConstants#MACHINE_FILE_ COLLECTION} collection

Name	Default value	Type	Description
rsa.endpoint.machine-file.periodic-bookmark-update-time	60	seconds	Interval between successive merging of {@link CollectionConstants#MACHINE_FILE_STAGE_COLLECTION} collection to {@link CollectionConstants#MACHINE_FILE_COLLECTION} collection
rsa.endpoint.machine-file.periodic-merge-delay	30	seconds	Interval between successive merging of {@link CollectionConstants#MACHINE_FILE_STAGE_COLLECTION} collection to {@link CollectionConstants#MACHINE_FILE_COLLECTION} collection
rsa.endpoint.machine-file.refresh-time	86400	seconds	Time interval to refresh the files present in a machine. The min value is set to 8h and max value is 48h.
rsa.endpoint.machine-file.refresh-time-delay	900	seconds	This is the time interval to check if agent files needs to be refreshed and create command for the agent if so.
rsa.endpoint.machine-file.retry-count	500	integer	Indicates the number of times it must be retried
rsa.endpoint.machine-file.retry-wait-time	10	seconds	Indicates the wait time for retrying to save machineFileHistory
rsa.endpoint.machine-file.staged-machine-file-deletion-delay	10	seconds	Delay between cleaning up of machine file data from {@link CollectionConstants#MACHINE_FILE_STAGE_COLLECTION} collection
rsa.endpoint.machine-file.threads	20	integer	Max number of merge machine file handler threads
rsa.endpoint.machine-file.update-history-limit	1000	integer	Number of documents to be updated into {@link CollectionConstants#MACHINE_FILE_HISTORY_COLLECTION} collection

MachineServiceProperties

Name	Default value	Type	Description
rsa.endpoint.machine.fetched-machines-limit	100	integer	The number of machine infos fetched for a given checksum. This is used to fetch the top 'n' risky machine-infos for a given file.

Name	Default value	Type	Description
rsa.endpoint.machine.search-query-timeout	10	seconds	Max timeout for machine detail to query to complete in milliseconds
rsa.endpoint.machine.status-persistence-interval	30	seconds	Interval in seconds in which machine/agent status will be persisted to db. Since it is a costly operation higher value is preferred and more higher the value is more inaccuracy will be in status related db queries

MachineFileScoreConfigurationProperties

Name	Default value	Type	Description
rsa.endpoint.machine.file.score.limit-of-checksums-in-batch	500	integer	
rsa.endpoint.machine.file.score.min-delay-for-refresh-seconds	120	seconds	

MetaForwardProperties

Name	Default value	Type	Description
rsa.endpoint.meta.enabled	false	boolean	Enable/Disable Meta integration
rsa.endpoint.meta.ld-buffer-check-enabled	true	boolean	Configuration option to disable the throttling on Log decoder buffer availability.
rsa.endpoint.meta.ld-buffer-limit-percentage	75	integer	Pool.packet.capture / pool.packet.page percentage at which we need to throttle.
rsa.endpoint.meta.logdecoder-host-id		string	The unique identifier of the host in which the Log decoder resides.
rsa.endpoint.meta.logdecoder-port	0	integer	Log decoder Port to which metas are to be posted
rsa.endpoint.meta.logdecoder-rest-password		string	Password to access the logdecoder rest port

Name	Default value	Type	Description
rsa.endpoint.meta.logdecoder-rest-port	0	integer	Log decoder REST Port to which metas are to be posted. This port number is used to query the available buffer before sending the meta.
rsa.endpoint.meta.logdecoder-rest-username		string	Username to access the logdecoder rest port
rsa.endpoint.meta.protobuf-ssl-enabled	false	boolean	SSL or Non SSL communication
rsa.endpoint.meta.rest-ssl-enabled	false	boolean	REST SSL or Non REST SSL communication
rsa.endpoint.meta.logdecoder-host		string	Log decoder Ip or hostname to which metas are to be posted @deprecated (since 6.0.0, To handle DHCP scenarios as well as manual IP change / load balancing scenarios, .use of ` logdecoderHost ` for LD communication is deprecated.The ` logdecoderHostId ` will be used instead of ` logdecoderHost ` for all log decoder communicationin future versions.

PackagerProperties

Name	Default value	Type	Description
rsa.endpoint.packager.agent-cert-name	client.p12	string	
rsa.endpoint.packager.beacon-interval	600	seconds	
rsa.endpoint.packager.packager-dir	/usr/lib/netwitness/endpoint-agents	string	

MachineDataHandlerProperties

Name	Default value	Type	Description
rsa.endpoint.queue.file-properties-drain-at-close	false	boolean	Optionally drain the queued files data to disk when the service is shutdown normally

Name	Default value	Type	Description
rsa.endpoint.queue.file-size	100	integer	Max number of concurrent data requests that should be handled by server for processing file data
rsa.endpoint.queue.file-threads	20	integer	Max number of file persistence threads
rsa.endpoint.queue.size	100	integer	Max number of concurrent data requests that should be handled by server
rsa.endpoint.queue.threads	10	integer	Max number of request handler threads

QueueFileSystemPersistenceProperties

Name	Default value	Type	Description
rsa.endpoint.queue.file.directory-context	dataDirectory	string	Represents the directory context (reference name) for the files to be persisted from file queues
rsa.endpoint.queue.file.path-prefix	temp/queue/files	string	Represents the path prefix for files to be persisted from Files queues

RelayCommunicationProperties

Name	Default value	Type	Description
rsa.endpoint.relay.communication.connect-timeout	30	seconds	Common connect timeout for all connections.
rsa.endpoint.relay.communication.initial-delay	30	seconds	Time to wait before attempting to connect to relay server
rsa.endpoint.relay.communication.max-connections	100	integer	Maximum number of connections allowed to nchan from relay server

Name	Default value	Type	Description
rsa.endpoint.relay.communication.nchan-base-url	https://localhost:7056	string	
rsa.endpoint.relay.communication.publish-channel	/agent/publish	string	
rsa.endpoint.relay.communication.request-timeout	30	seconds	Common request timeout for all connections.
rsa.endpoint.relay.communication.retry-interval	10	seconds	Delay between connection attempts
rsa.endpoint.relay.communication.subscribe-channel	/endpoint_ server/subscribe	string	
rsa.endpoint.relay.communication.subscribe-request-timeout	5	seconds	0s is infinite time.
rsa.endpoint.relay.communication.thread-pool-size	100	integer	

RelayInstallerProperties

Name	Default value	Type	Description
rsa.endpoint.relay.installer.cert-name	relay-server-cert.p12	string	Relay-server certificate file name
rsa.endpoint.relay.installer.dependency-dir	/var/netwitness/ endpoint-server/relay	string	Directory where relay-server dependencies will be downloaded. Non root user must have read, write access.
rsa.endpoint.relay.installer.download-on-restart	true	boolean	Flag to decide whether to delete local copy of relay-server dependencies and download from configured yum repo on every endpoint server restart. It might take sometime for the downloading to complete, during which user will not be able to download relay-server installer.

Name	Default value	Type	Description
rsa.endpoint.relay.installer.init-delay	20	seconds	Delay for Background task which will download relay-server dependencies.

RelayMetricsProperties

Name	Default value	Type	Description
rsa.endpoint.relay.metrics.periodic-evaluation-delay	300	seconds	Time interval to evaluate if any relay-server config was modified and update the metrics if required
rsa.endpoint.relay.metrics.refresh-time	300	seconds	Time interval to refresh the metrics from relay-server server

SslContextProperties

Name	Default value	Type	Description
rsa.endpoint.ssl.ssl-session-cache-size	0	integer	Max number of sessions to be kept in ssl session cache
rsa.endpoint.ssl.ssl-session-timeout	0	seconds	Max time an SSL session can be reused

ThrottlingConfigurationProperties

Name	Default value	Type	Description
rsa.endpoint.throttling.enabled	true	boolean	
rsa.endpoint.throttling.max	70	integer	

UdpProperties

Name	Default value	Type	Description
rsa.transport.udp.enabled	true	boolean	Boolean to indicate if server can consume Udp packet
rsa.transport.udp.port	0	integer	UDP port
rsa.transport.udp.size	5000	integer	Max number of concurrent data requests that should be handled by server
rsa.transport.udp.threads	20	integer	Max number of request handler threads

Enrichment-server Configuration

EngineProperties

Name	Default value	Type	Description
rsa.enrichment.engine.auto-start	true	boolean	Determines if all {@link Engine} should start on service deployed.
rsa.enrichment.engine.startup-error-retry-interval	10	seconds	Retry interval if error occurs during startup.

StreamProperties

Name	Default value	Type	Description
rsa.enrichment.stream.buffer-size	40000	integer	Controls the number of records the stream can keep outstanding.
rsa.enrichment.stream.connection-time-out	0	long	Override connection timeout in sources. Only if greater than 0.
rsa.enrichment.stream.dots-to-underscores	true	boolean	Choose if we want to translate "user.dst" to "user_dst".
rsa.enrichment.stream.event-source-id	true	boolean	Controls whether we need to add the event source identifier (ESA compatibility)
rsa.enrichment.stream.lag-time	15	seconds	Lag time is the expected time an event takes to pass through the different levels of capture/parse etc and become available to query in the concentrator.
rsa.enrichment.stream.mechanism		string	{@link StreamSettings.Mechanism}.
rsa.enrichment.stream.minutes-back	5	integer	Controls how far back in time should we go for a fresh start.
rsa.enrichment.stream.multi-valued		list	Choose the fields considered as multi-valued.
rsa.enrichment.stream.reader-buffer-size	1048576	integer	

Name	Default value	Type	Description
rsa.enrichment.stream.save-position-every	1	seconds	A length of time to apply the permits. Minimum of 1 second and max at 1 day.
rsa.enrichment.stream.save-position-frequency	1	integer	Number of permits for a duration.
rsa.enrichment.stream.socket-timeout	0	long	Override socket timeout in sources. Only if greater than 0.
rsa.enrichment.stream.source-poll-interval	1000	integer	Controls the parameters passed to RecordSource#poll(int, TimeUnit).
rsa.enrichment.stream.start-session-id	0	long	Override StartSession Id in sources for debug purposes. Only if greater than 0.
rsa.enrichment.stream.tcp-no-delay	true	boolean	
rsa.enrichment.stream.time-batch-in-seconds	60	integer	Determines the batch size for the query based aggregation in seconds. By default it will be a 60 second window. This for now will not be configurable for user. This is because concentrator operates most efficiently when the time window is a minute.
rsa.enrichment.stream.time-measured-in-seconds	true	boolean	{@code true} if time meta is measured in seconds in the event.
rsa.enrichment.stream.time-meta-field	time	string	Decides what field should be used for time.
rsa.enrichment.stream.time-ordered	false	boolean	Decides if the stream should be time ordered.
rsa.enrichment.stream.use-event-time-for-esper	false	boolean	{@code true} to use the timeMetaField in the Event for Esper CurrentTimeEvent.

Integration-server Configuration

DeploymentNotificationProperties

Name	Default value	Type	Description
rsa.notification.startup-retry-interval	15s	seconds	The service startup failure retry interval.

Investigate-server Configuration

AliasesProperties

Name	Default value	Type	Description
rsa.investigate.aliases.cache-duration	24	seconds	Time it takes for the cache that stores aliases to expire
rsa.investigate.aliases.retrieval-timeout	30	seconds	Timeout to wait for aliases sdk response

ColumnGroupProperties

Name	Default value	Type	Description
rsa.investigate.column.group.number-of-visible-columns	15	integer	

EventAnalysisProperties

Name	Default value	Type	Description
rsa.investigate.eventanalysis.legacy-events-enabled	false	boolean	Flag to determine if legacy events tab and related links have to be enabled
rsa.investigate.eventanalysis.limit	5000	integer	The default event limit
rsa.investigate.eventanalysis.role-event-limit		map	The per-role event limit

IncidentProperties

Name	Default value	Type	Description
rsa.investigate.incident.max-events-per-alert	60	long	Max. number of events that should be added to a single alert when creating incidents from events

KeyrefsProperties

Name	Default value	Type	Description
rsa.investigate.keyrefs.cache-duration	2	seconds	Time it takes for the cache that stores aliases to expire
rsa.investigate.keyrefs.retrieval-timeout	30	seconds	Timeout to wait for aliases sdk response

MetaKeyCacheProperties

Name	Default value	Type	Description
rsa.investigate.metakey.cache.cache-duration	7	seconds	Number of seconds a metakey should live in the cache. Default: 1 WEEK

ReconstructionProperties

Name	Default value	Type	Description
rsa.investigate.reconstruction.clear-cache-older-than	24	seconds	Cache files which are older than this time interval would be cleared
rsa.investigate.reconstruction.compressed-file-password	netwitness	string	Default zip file password for email recon downloads

Name	Default value	Type	Description
rsa.investigate.reconstruction.content-type-file-extractor-max-size	4	bytes	From NetWitness Core documentation <p> The max number of bytes to return, zero means no limit. This parameter is used to control the maximum bytes that a large network session should return and is mainly meant to prevent an extraordinary large network session from consuming a large number of resources during the transfer. Be careful setting this parameter to zero.
rsa.investigate.reconstruction.email-attachment-hash-provider		reconstructionproperties $emailattachmenthashprovider	The calculated hash type for any email attachments
rsa.investigate.reconstruction.email-full-render	true	boolean	Flag to enable/disable full rendering of email messages. <p> When set to true email bodies will be fully reconstructed which will benefit email's with HTML body content. Styling will be preserved as best as possible—external styles and references must be removed—and inline content (images), if included in the session, will be displayed. Placeholders will be shown for content that is not available or cannot be rendered and any inline script should be made inactive but displayed to the user for informational purposes. <p> If set to false, standard rendering is used which will render the email body as best as possible and return it as text in the `bodyContent` field of the {@link Email} object. <p> This setting is dependent on the Reconstruction Object Cache being enabled. (see {@link ReconstructionProperties#objectCacheEnabled}) It is ignored otherwise.

Name	Default value	Type	Description
rsa.investigate.reconstruction.endpoint-enrichment-time-window	30	seconds	Endpoint Enrichment Query time window in seconds. Network events will be correlated with endpoint events triggered within this time window of the network event's time. If network event time is x, endpoint events will be queried from 'x - @endpointEnrichmentTimeWindow' time
rsa.investigate.reconstruction.endpoint-enrichment-time-window-buffer	5	seconds	Additional buffer time range used to query for endpoint enrichment data. If network event time is x, endpoint events will be queried till 'x + @endpointEnrichmentTimeWindow Buffer' time
rsa.investigate.reconstruction.endpoint-events-query-time-out	5	seconds	Max time allowed in seconds for all endpoint core queries to complete
rsa.investigate.reconstruction.enrichment-instance-init-delay	1	seconds	Initial delay to fetch the investigate service details from all orchestrated endpoint services
rsa.investigate.reconstruction.image-placeholder-url		uri	Url used in Email recon for web email when original images cannot be loaded
rsa.investigate.reconstruction.object-cache-enabled	true	boolean	Flag to enable/disable reconstruction object cache. In addition to caching the content (protobuf files) that are downloaded from core devices, the investigate service will attempt to cache any objects and files that are created while reconstructing sessions. For this release (11.4—the first release with the object cache) this only pertains to email reconstruction.
rsa.investigate.reconstruction.reactive-message-size	256	bytes	Used in reactive streaming to configure the maximum buffer size for holding reconstructed data.

Name	Default value	Type	Description
rsa.investigate.reconstruction.reactive-text-streaming	true	boolean	Flag to turn on reactive streaming for text reconstruction. Reactive streaming prevents web socket overload by sending as many reconstructed text blocks that fit into a known buffer size and stopping until the caller tells the service to proceed.
rsa.investigate.reconstruction.session-enrichment-time-out	10	seconds	Max time allowed in seconds for all enrichment queries to complete including core, endpoint and other enrichment queries
rsa.investigate.reconstruction.support-script-urls		uri[]	If html is generated in reconstruction, that is served to the UI via an IFRAME (as to not interfere with the functionality/styling of the main application) this setting stores an array of strings (url's) to javascript files that will be injected into the html. The javascript is injected via <script /> elements at the time of HTML file creation and therefore will be saved to the object-cache. Any updates to this array would require clearing of the object-cache and/or a service restart.
rsa.investigate.reconstruction.sync-core-timeout	600	seconds	Max time to wait for operations for caching core content to complete to prevent deadlocks. Internal setting. Not recommended for customer use.
rsa.investigate.reconstruction.wire-size-provider		reconstructionproperties $wiresizeprovider	The method used to determine object size when transmitting objects via websocket

ResponseProperties

Name	Default value	Type	Description
rsa.investigate.response.events-batch-size	5000	long	Number of data size to send per message. If client send request with stream batch size and it is smaller than this, the client batch size will be used instead.

EventsStreamProperties

Name	Default value	Type	Description
rsa.investigate.stream.events.factor-of-multiple-meta-values-with-same-key	5	integer	Like the above property. This property is used to calculate a safety threshold if not specified. It's a factor to allow for multiple meta values existing in the same key and should be something reasonably high.
rsa.investigate.stream.events.safe-num-of-column-selected	50	integer	Used to calculate a safety value for "threshold" in the query to avoid the query going unbounded if threshold is not specified. The value of threshold is calculate by the formula below: threshold = (num of sessions desired) * (num of column selected) * (factor of multiple meta values with same key) If the above (num of column selected) can't be inferred from "select" field, this default value would be used.

Launch-framework Configuration

ConfigurationModuleProperties

Name	Default value	Type	Description
rsa.configuration.backoff-duration	1	seconds	Amount of time to wait until a retry is attempted if the config-server is unavailable
rsa.configuration.connection-timeout	30	seconds	A timeout how long to wait if the config-server is unavailable
rsa.configuration.remote-enabled	false	boolean	If the configuration server is even attempted
rsa.configuration.schema-synchronization-retry-interval	1	seconds	This property controls how long to wait before retrying a failed schema synchronization attempt.

ContentProperties

Name	Default value	Type	Description
rsa.content.disk-path		path	The path where the content resides on disk

DataProperties

Name	Default value	Type	Description
rsa.data.application.advanced		map	A set of advanced properties specific to the data provider
rsa.data.application.auth-mechanism	SCRAM	string	Default username/password authentication "SCRAM". Alternative: "PLAIN"
rsa.data.application.connection-timeout	5	seconds	How long to wait before giving up on a connection attempt
rsa.data.application.database		string	The database name

Name	Default value	Type	Description
rsa.data.application.enabled	false	boolean	If true will enable database support
rsa.data.application.map-key-dot-replacement		string	Mongo disallows "." in map keys, if a value is provided, dots in map keys are replaced by it.
rsa.data.application.password		string	The connection password
rsa.data.application.secure	false	boolean	Use an SSL/TLS connection to the database
rsa.data.application.servers	[localhost]	string[]	A comma separated list of database servers
rsa.data.application.stat-cache-timeout	15	seconds	How long to wait before refreshing database statistics?
rsa.data.application.user		string	The connection user
rsa.data.control.advanced		map	A set of advanced properties specific to the data provider
rsa.data.control.auth-mechanism	SCRAM	string	Default username/password authentication "SCRAM". Alternative: "PLAIN"
rsa.data.control.connection-timeout	5	seconds	How long to wait before giving up on a connection attempt
rsa.data.control.database		string	The database name
rsa.data.control.enabled	false	boolean	If true will enable database support
rsa.data.control.map-key-dot-replacement		string	Mongo disallows "." in map keys, if a value is provided, dots in map keys are replaced by it.
rsa.data.control.password		string	The connection password
rsa.data.control.secure	false	boolean	Use an SSL/TLS connection to the database
rsa.data.control.servers	[localhost]	string[]	A comma separated list of database servers
rsa.data.control.stat-cache-timeout	15	seconds	How long to wait before refreshing database statistics?
rsa.data.control.user		string	The connection user

FileSystemProperties

Name	Default value	Type	Description
rsa.filesystem.conf-path	/etc/netwitness	string	The path to directory where all service configuration resides. Ignored if prefix is specified.
rsa.filesystem.data-path	/var/lib/netwitness	string	The path to directory where all service data resides. Ignored if prefix is specified.
rsa.filesystem.logs-path	/var/log/netwitness	string	The path to directory where all service logs reside. Ignored if prefix is specified.
rsa.filesystem.prefix		string	If not empty the prefix specifies the root for all service file system state. When empty, the individual values are used.

HealthCheckProperties

Name	Default value	Type	Description
rsa.health.check-every	15	seconds	Rate at which health checks are scheduled to run.
rsa.health.concurrency	5	integer	Number of concurrent threads that runs health checks.
rsa.health.timeout	30	seconds	Time out for a {@link com.rsa.asoc.launch.api.health.HealthCheck} when service health checks are run. If a component is unable to respond with health status with in this period, it is marked as {@link com.rsa.asoc.launch.api.health.Health.Status#Unhealthy}

LoggingAuditProperties

Name	Default value	Type	Description
rsa.logging.audit.max-file-count	10	integer	The maximum number of archive files to retain.
rsa.logging.audit.max-file-size	10	bytes	The maximum size a log file is allowed to grow

LogForwarderProperties

Name	Default value	Type	Description
rsa.logging.forward.categories		string[]	The log categories to choose for forwarding
rsa.logging.forward.destination		logforwarderproperties$destination	The forwarding destination
rsa.logging.forward.enabled	true	boolean	Is forwarding enabled?
rsa.logging.forward.host	localhost	string	The destination host address
rsa.logging.forward.port	50514	integer	The destination port
rsa.logging.forward.secure	false	boolean	Use TLS for forwarding (only supported with LOGSTASH_TCP)

LoggingProperties

Name	Default value	Type	Description
rsa.logging.levels		string	Service log levels specified as a comma separated sequence of "logger:level". Note logger names are case sensitive.

LoggingOperationalProperties

Name	Default value	Type	Description
rsa.logging.operations.max-file-count	10	integer	The maximum number of archive files to retain.

Name	Default value	Type	Description
rsa.logging.operations.max-file-size	10	bytes	Maximum file size of each file allowed to grow

MetricsAggregationProperties

Name	Default value	Type	Description
rsa.metrics.aggregation.enabled	true	boolean	Is the reporter enabled?
rsa.metrics.aggregation.filter-prefixes		list	What to report? The default behavior is to report everything, if a selection of metrics must be reported add their prefixes to this list.
rsa.metrics.aggregation.host		string	The host name of the aggregator.
rsa.metrics.aggregation.interval	1	seconds	How often to report?
rsa.metrics.aggregation.port	0	integer	The port number.

MetricsElasticProperties

Name	Default value	Type	Description
rsa.metrics.elastic.enabled	true	boolean	Is the reporter enabled?

MetricsHistoricalProperties

Name	Default value	Type	Description
rsa.metrics.historical.enabled	true	boolean	Is the reporter enabled?
rsa.metrics.historical.filter-prefixes		list	What to report? The default behavior is to report everything, if a selection of metrics must be reported add their prefixes to this list.
rsa.metrics.historical.interval	1	seconds	How often to report?

Name	Default value	Type	Description
rsa.metrics.historical.max-file-count	10	integer	The maximum number of archive files to retain.
rsa.metrics.historical.max-file-size	10	bytes	Maximum file size of each file allowed to grow

MetricsJmxProperties

Name	Default value	Type	Description
rsa.metrics.jmx.enabled	true	boolean	Is the reporter enabled?
rsa.metrics.jmx.filter-prefixes		list	What to report? The default behavior is to report everything, if a selection of metrics must be reported add their prefixes to this list.
rsa.metrics.jmx.interval	1	seconds	How often to report?

MetricsProperties

Name	Default value	Type	Description
rsa.metrics.profile-api-invocation	true	boolean	Profiles timing of all {@link com.rsa.asoc.launch.api.annotation.LaunchApi} methods.

NotificationProperties

Name	Default value	Type	Description
rsa.notification.drain-at-shutdown	true	boolean	The flag to control if we drain the notification queue before shutdown. If there are a lot of pending notifications this may cause noticeable delays in shutdown time, particularly if the deployment integration server is unavailable and each forward goes through the retry attempts.

Name	Default value	Type	Description
rsa.notification.max-pending	1000	integer	The maximum number of notifications left pending.
rsa.notification.max-threads	1	integer	The size of the thread pool.
rsa.notification.retry-at-shutdown	true	boolean	The flag to control if we should retry failed notifications when the service is going down. This is true by default but can lead to delayed shutdowns if notifications cannot be forwarded.
rsa.notification.retry-attempts		integer	The number of times we retry if a notification cannot be forwarded to the centralized notification service. The default setting is to never give up but can be changed to smaller value (e.g. 10) if it is OK to drop some notifications.
rsa.notification.retry-delay	10s	seconds	The delay between successive retry attempts.

ProcessJvmMemoryProperties

Name	Default value	Type	Description
rsa.process.jvm.memory-thresholds.fatal-percent	95	integer	Percent of heap memory usage, above which JVM health is marked Fatal
rsa.process.jvm.memory-thresholds.warning-percent	80	integer	Percent of heap memory usage, above which JVM health is marked Unhealthy

ProcessProperties

Name	Default value	Type	Description
rsa.process.shutdown-delay	5	seconds	The delay between a request to shutdown and the eventual shutdown trigger.

AuthenticationProperties

Name	Default value	Type	Description
rsa.security.authentication.prefetch-before	5	seconds	If prefetch service token enabled, fetch next when there is this much seconds left for expiry of current token
rsa.security.authentication.prefetch-service-token	true	boolean	Prefetch service tokens before they expire
rsa.security.authentication.refresh-token-lifetime	30	seconds	The time-to-live on a refresh token.
rsa.security.authentication.remote-enabled	false	boolean	Support remote authentication.
rsa.security.authentication.remote-timeout	30	seconds	The time to wait for a response before failing a remote authentication.
rsa.security.authentication.token-lifetime	8	seconds	The time-to-live on a token.
rsa.security.authentication.trusted-channel-enabled	true	boolean	Support trusted channel authentication.

AuthorizationProperties

Name	Default value	Type	Description
rsa.security.authorization.permission-cache-expiry	15	seconds	This property controls cache expiry interval for the role to permissions mapping.
rsa.security.authorization.permission-cache-size	100	integer	This property controls number of role definitions cached in the service.
rsa.security.authorization.permission-synchronization-retry-interval	1	seconds	This property controls how long to wait before retrying a failed permission synchronization attempt.
rsa.security.authorization.remote-enabled	true	boolean	This property controls if the service must synchronize its permissions to the deployment Security Server.

CertificateAuthorityProperties

Name	Default value	Type	Description
rsa.security.ca.alias	Service CA	string	The alias for the CA keypair.
rsa.security.ca.auto-sign-operational-certificate	false	boolean	Should the service operational certificate be automatically signed by the embedded CA?
rsa.security.ca.certificate-lifetime	1000	seconds	The certificate validity lifetime
rsa.security.ca.issued-time-allowance	10	seconds	The certificate issued time can allow some clock drift.
rsa.security.ca.store-certificates	false	boolean	Should the service store certificates it signs

PkiProperties

Name	Default value	Type	Description
rsa.security.pki.audit-tls-hand-shakes	true	boolean	Enables auditing of TLS handshakes
rsa.security.pki.ciphers	[TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256]	string[]	This property controls the list of SSL cipher suites used by the service.
rsa.security.pki.client-auth		ssl$clientauth	This property controls the SSL client authentication preference.
rsa.security.pki.tls-protocols		string[]	This property controls the TLS protocol versions supported by the applications.

Name	Default value	Type	Description
rsa.security.pki.trust-synchronization-retry-interval	1	seconds	This property controls how long to wait before retrying a deployment trust synchronization attempt.
rsa.security.pki.use-deployment-trust	true	boolean	This property controls if the deployment security-server must be trusted.
rsa.security.pki.use-jvm-trust	false	boolean	This property controls if the JVM trust store should be used to validate peer certificates.
rsa.security.pki.verify-certificates	true	boolean	This property controls whether we must verify server certificates.

ServiceAccountProperties

Name	Default value	Type	Description
rsa.security.serviceaccounts.auth-request-validity	1	seconds	Validity of the auth request
rsa.security.serviceaccounts.max-request-cache-size	10000	integer	Cache hashes of previously authenticated requests to prevent renegotiation attacks
rsa.security.serviceaccounts.min-interval-between-authentications	30	seconds	Minimum interval between authentications to prevent brut-force attacks
rsa.security.serviceaccounts.token-lifetime	1	seconds	The time-to-live for a service account token.

TransportBusProperties

Name	Default value	Type	Description
rsa.transport.bus.advanced		map	A map that holds any other RabbitProperties configuration
rsa.transport.bus.enabled	true	boolean	Defines if Bus is enabled
rsa.transport.bus.host	localhost	string	The connection host
rsa.transport.bus.max-consumers	16	integer	Maximum number of consumers reading the queues
rsa.transport.bus.message-prefetch	1	integer	Number of messages to be handled in a single request
rsa.transport.bus.min-consumers	8	integer	Minimum number of consumers reading the queues
rsa.transport.bus.password		string	The connection password
rsa.transport.bus.port	5672	integer	The connection port
rsa.transport.bus.recovery-interval	15	seconds	The time to wait between attempts to recover a broken RabbitMQ broker connection.
rsa.transport.bus.reply-timeout	15	seconds	The time to wait for replies to arrive before giving up. AMQP is not connection-oriented so the absence of a service on the bus can only be determined by the absence of a reply. This timeout determines how long the framework waits before giving up.
rsa.transport.bus.secure	false	boolean	Use an SSL/TLS connection to the broker
rsa.transport.bus.shutdown-timeout	5	seconds	The time to wait for workers after the container is stopped, and before the connection is forced closed. If any workers are active when the shutdown signal comes they will be allowed to finish processing as long as they can finish within this timeout. Otherwise the connection is closed and messages remain unacked (if the channel is transactional).
rsa.transport.bus.username		string	The connection user
rsa.transport.bus.virtual-host	/rsa/system	string	The connection virtual host

TransportBusSubscriptionProperties

Name	Default value	Type	Description
rsa.transport.bus.subscription.subscribe-retry-interval	5	seconds	The interval to retry declaring the bindings for subscriptions if the exchange is unavailable or missing.

TransportProperties

Name	Default value	Type	Description
rsa.transport.detailed-errors-enabled	false	boolean	A boolean indicating whether the server should return detailed errors that may contain additional implementation details.

TransportHttpProperties

Name	Default value	Type	Description
rsa.transport.http.access-token-headers		string[]	Defines the HTTP headers to check for an access token
rsa.transport.http.basic-auth-enabled	false	boolean	Defines if web security basic authentication should be enabled.
rsa.transport.http.enabled	true	boolean	Defines if HTTP is enabled
rsa.transport.http.keep-alive-timeout	60s	seconds	The number of seconds this Connector will wait for another HTTP request before closing the connection. Setting the value to 0 will indicate no (i.e. infinite) timeout.
rsa.transport.http.max-keep-alive-requests	100	integer	The maximum number of HTTP requests which can be pipelined until the connection is closed by the server. Setting this to -1 will allow an unlimited amount of pipelined or keep-alive HTTP requests.

Name	Default value	Type	Description
rsa.transport.http.port	8080	short	Defines the HTTP port
rsa.transport.http.secure	true	boolean	Defines if HTTPS must be used
rsa.transport.http.session-id-random-algorithm	HMACDRBG	string	This property controls the algorithm to use for the SecureRandom used to generate HTTP session IDs.

License-server Configuration

LicenseProperties

Name	Default value	Type	Description
rsa.license.auto-cleanup-enabled	true	boolean	Whether to enable data retention job
rsa.license.auto-cleanup-interval	24 hours	seconds	Data retention job interval. The job cleans License usage and history data older than `dataRetention` property value
rsa.license.auto-refresh-interval	60 minutes	seconds	Periodic entitlement refresh interval
rsa.license.compliance-cache-expiry	60 minutes	seconds	Compliance will be evaluated and cached for this duration before evaluating again. Set to 0 to disable cache. Maximum cache expiry duration is 6 hours
rsa.license.data-retention-period	365 days	seconds	Days to retain license usage and history data in repository

FneProperties

Name	Default value	Type	Description
rsa.license.fne.base-url	/api/1.0/	string	Base URL path of REST API, all other URLs are appended to baseUrl
rsa.license.fne.capability-request-url	/capability_request/offline	string	URL to generate a new capability request
rsa.license.fne.capability-response-url	/capability_response	string	URL to post a response of capability request
rsa.license.fne.delay-before-proxy-update	2	seconds	Wait time to hold the proxy update job before receiving all the proxy config change notifications.

Name	Default value	Type	Description
rsa.license.fne.feature-page-size	25	integer	Requests to local flexnet server is paginated. This property controls the number of features that FlexnetLS returns in a single request. Defaults to 25.
rsa.license.fne.features-url	/features	string	URL to fetch features, %s is replaced with FNE License Server instance id
rsa.license.fne.fne-restart-timeout	5 minutes	seconds	Maximum time to wait for the FNE server to restart before timing out.
rsa.license.fne.host	localhost	string	Host where FNE License Server is running, defaults to localhost
rsa.license.fne.host-id-url	/hostids/selected	string	URL to fetch Unique identifier of the environment that maps licenses
rsa.license.fne.host-ids-url	/hostids	string	URL to fetch all host identifiers available in the environment
rsa.license.fne.local-config-file	/opt/netwitness/flexnetls/local-configuration.yaml	string	Local Configuration file
rsa.license.fne.max-pages	50	integer	Maximum number of pages to be requested from FlexnetLS features response. This is to prevent requesting infinite number of pages from FlexnetLS because of any reason which may cause end-of-response to not be detected.
rsa.license.fne.port	3333	integer	Port number where FNE License Server accepts REST requests, defaults to 3333
rsa.license.fne.proxy-config-file	/etc/netwitness/flexnetls/proxy.conf	string	Proxy configuration file
rsa.license.fne.sync-url	/capability_request	string	URL to force an online sync from FNE License Server to FNO remote service to fetch Features

Name	Default value	Type	Description
rsa.license.fne.trust-store-path	/var/netwitness/flexnetls	string	Path to the directory which contains FNE server trust store files
rsa.license.fne.use-ssl	false	boolean	Is the REST API secured

Metrics-server Configuration

MetricsProperties

Name	Default value	Type	Description
rsa.metrics.content.content-sync-interval	2 minutes	seconds	Time interval at which content available should be tried to upload to Elastic-Kibana
rsa.metrics.content.elastic-health-check-interval	1 minute	seconds	Time interval at which Elastic and Kibana health should be checked
rsa.metrics.content.metrics-health-check-interval	30 minutes	seconds	Time interval at which metrics collection health for Netwitness services should be checked
rsa.metrics.content.orchestration-sync-interval	30 Seconds	seconds	Interval at which orchestration-server should be queried to get deployment details
rsa.metrics.content.overwrite-index-template	false	boolean	Controls if NW index template need to be overridden
rsa.metrics.content.saved-objects-created	false	boolean	Whether to create saved objects

AlertProcessingProperties

Name	Default value	Type	Description
rsa.metrics.elastic.alert.metrics-interval	900	seconds	Duration (in seconds) after which Alarm metrics is uploaded to elastic
rsa.metrics.elastic.alert.retry-delay	1	seconds	Interval between subsequent retries
rsa.metrics.elastic.alert.wait-time	1	seconds	Wait time post which, first retry attempt on Alert Notification will be done
rsa.metrics.elastic.alert.work-interval	60	seconds	Frequency (in seconds) how often Alerts need to collected from Elastic

ElasticRetentionProperties

Name	Default value	Type	Description
rsa.metrics.elastic.data.retention.alert-retention-threshold	60	seconds	Time duration (in days) for which alerts need to be retained
rsa.metrics.elastic.data.retention.allocated-size	100	bytes	Disk space permissible to be used for Netwitness indexes in Elastic
rsa.metrics.elastic.data.retention.time-retention-frequency	4	seconds	Interval between invocations of the time based retention job
rsa.metrics.elastic.data.retention.time-threshold	30	seconds	Time (in days) for which Additional Telemetry data needs to be retained

ElasticProperties

Name	Default value	Type	Description
rsa.metrics.elastic.host		string	Host Elastic node/cluster is running on
rsa.metrics.elastic.password		string	Password for Elastic
rsa.metrics.elastic.port	9200	integer	Port Elastic node/cluster is running on
rsa.metrics.elastic.secure	true	boolean	Secure to decide protocol for Elastic
rsa.metrics.elastic.username		string	Username to connect Elastic
rsa.metrics.elastic.verify-hostname	false	boolean	Whether to verify Elastic hostname on SSL validation

ElasticClientCertAuthenticatorProperties

Name	Default value	Type	Description
rsa.metrics.elastic.security.authenticator.clientcert.http-enabled	false	boolean	Whether HTTP authentication requests are enabled
rsa.metrics.elastic.security.authenticator.clientcert.order	0	integer	Numerical order of this authentication domain
rsa.metrics.elastic.security.authenticator.clientcert.roles-attribute		string	DN attribute in X509 certificate that identifies role of the subject principal. In NW product suite, username is mapped to a role name in elasticsearch
rsa.metrics.elastic.security.authenticator.clientcert.transport-enabled	false	boolean	Whether authentication enabled at Transport level
rsa.metrics.elastic.security.authenticator.clientcert.username-attribute		string	DN attribute in X509 Certificate that identifies subject principal (username), defaults to 'cn' (Common name)

ElasticJwtAuthenticatorProperties

Name	Default value	Type	Description
rsa.metrics.elastic.security.authenticator.jwt.http-enabled	false	boolean	Whether HTTP authentication requests are enabled
rsa.metrics.elastic.security.authenticator.jwt.jwt-header		string	Authentication backend defines how to validate username or credentials, this is specific to authenticator
rsa.metrics.elastic.security.authenticator.jwt.order	0	integer	Numerical order of this authentication domain
rsa.metrics.elastic.security.authenticator.jwt.roles-key		string	Key name in the JWT claim that contains to user roles
rsa.metrics.elastic.security.authenticator.jwt.signing-key		string	Public key from the issuer certificate that issued JWT token in Base64 format without headers (i.e.without `----BEGIN/END Publi Key----`)
rsa.metrics.elastic.security.authenticator.jwt.subject-key		string	Name of the key in JWT claim that identifies subject principal (user)
rsa.metrics.elastic.security.authenticator.jwt.transport-enabled	false	boolean	Whether authentication enabled at Transport level

ElasticServiceProperties

Name	Default value	Type	Description
rsa.metrics.elastic.service.password		string	Password for nw-service user to be used for metrics collection

KibanaProperties

Name	Default value	Type	Description
rsa.metrics.kibana.base-path		string	
rsa.metrics.kibana.connection-time-out	30	seconds	Time to changes connection timeout
rsa.metrics.kibana.host		string	Host Kibana node/cluster is running on
rsa.metrics.kibana.password		string	Password for kibana
rsa.metrics.kibana.port	0	integer	Port Kibana node/cluster is running on
rsa.metrics.kibana.secure	false	boolean	Secure to decide protocol for Kibana
rsa.metrics.kibana.username		string	Username to connect kibana
rsa.metrics.kibana.verify-hostname	false	boolean	Whether to verify kibana hostname on SSL validation

MetricsRetryProperties

Name	Default value	Type	Description
rsa.metrics.retry.retry-attempts	5	integer	The number of times we retry if a service can be configured with the latest configuration.
rsa.metrics.retry.retry-delay	4min	seconds	The delay between successive retry attempts.

TelemetryProperties

Name	Default value	Type	Description
rsa.telemetry.collect-click-stream-metrics	false	boolean	Controls if the click stream stats are to be collected
rsa.telemetry.config-load-retry-interval	2	seconds	Controls the interval to load the telemetry config

MongoDataRetentionProperties

Name	Default value	Type	Description
rsa.telemetry.mongo.data.retention.additional-data-threshold	7	seconds	Time (in days) for which Additional Telemetry data needs to be retained
rsa.telemetry.mongo.data.retention.disk-allocated-size	50	integer	Size of the data (in GB) that is allowed to be consumed by Telemetry data
rsa.telemetry.mongo.data.retention.mandatory-data-threshold	365	seconds	Time (in days) for which Mandatory Telemetry data needs to be retained
rsa.telemetry.mongo.data.retention.size-retention-frequency	1	seconds	Interval between invocations of the storage size based retention job

Node-infra-server Configuration

AdminNodeTrackerProperties

Name	Default value	Type	Description
rsa.platform.node.admin-node-ip-list-save-path	/etc/netwitness/node-infra-server/admin-node-ip-list.json	string	Path to persist ip list of possible admin nodes
rsa.platform.node.call-home-interval	15	seconds	Call home interval Reaches out to orchestration server at this given interval
rsa.platform.node.chef-status-check-interval	1	seconds	Override wait time between checks for chef status checks
rsa.platform.node.enable-call-home	true	boolean	Enable call home feature
rsa.platform.node.enable-track-admin-hosts	true	boolean	Enable tracking of admin hosts feature
rsa.platform.node.path-for-admin-node-primary-check	/nodeMode	string	Part of the URL path to find information about the admin node being primary or not
rsa.platform.node.primary-admin-host-ip-path	/etc/netwitness/node-infra-server/primary-admin-host-ip	string	Path to save the primary admin host's ip before running chef. Chef / Ohai plugin reads this file.
rsa.platform.node.restart-after-admin-node-ip-change	true	boolean	Restart service after Admin node Ip change
rsa.platform.node.switch-master-run-list	nw-dns-client	string	Chef run list to switch to new master ip
rsa.platform.node.track-admin-hosts-interval	5	seconds	Node mode check interval Reaches out to admin node to verify if it is still the primary

ChefRunProperties

Name	Default value	Type	Description
rsa.platform.node.chefrun.chef-client-command	/usr/bin/chef-client	string	Path to chef client

Name	Default value	Type	Description
rsa.platform.node.chefrun.chef-client-log-level		chefrunproperties$ chefclientloglevel	The log level of chef client
rsa.platform.node.chefrun.chef-config-location	/var/lib/netwitness/config-management/client.rb	string	The location of the Chef solo configuration
rsa.platform.node.chefrun.chef-log-location	/var/log/netwitness/config-management/chef-client.log	string	The location to write the log file for a Chef client execution
rsa.platform.node.chefrun.chef-run-timeout	1	seconds	Timeout for chef run before marking task as failed
rsa.platform.node.chefrun.sudo	true	boolean	Execute command with sudo

No-op-server Configuration

TestProperties

Name	Default value	Type	Description
rsa.noop.test.array-string		string[]	
rsa.noop.test.map-uri		map	
rsa.noop.test.pojo-array		testproppojo[]	
rsa.noop.test.pojo-list		list	
rsa.noop.test.pojo-map		map	
rsa.noop.test.simple-long		long	
rsa.noop.test.simple-pojo		testproppojo	
rsa.noop.test.simple-string		string	
rsa.noop.test.some-interval		seconds	
rsa.noop.test.some-size		bytes	

Orchestration-server Configuration

HostProperties

Name	Default value	Type	Description
rsa.orchestration.engine.host-infra.back-off-interval	15	seconds	Backoff interval in seconds
rsa.orchestration.engine.host-infra.max-backoff-factor	5	integer	Maximum call-home delay factor. Used to exponentially back off for subsequent host refresh failures.

SaltClientProperties

Name	Default value	Type	Description
rsa.orchestration.engine.salt.client.api-uri	https://localhost:8000/run	uri	The URI for the SaltStack API
rsa.orchestration.engine.salt.client.external-authentication	pam	string	The Salt API authentication mechanism
rsa.orchestration.engine.salt.client.host-verification-enabled	true	boolean	The certificate host validation when communicating with the Salt API
rsa.orchestration.engine.salt.client.password		string	The Salt API password
rsa.orchestration.engine.salt.client.retry-interval	5	seconds	The amount of time in between retries
rsa.orchestration.engine.salt.client.retry-timeout	2	seconds	The maximum amount of time to wait during retries of Salt API commands {@link #retryInterval}
rsa.orchestration.engine.salt.client.username		string	The Salt API username

ChefConfigurationProperties

Name	Default value	Type	Description
rsa.orchestration.engine.salt.component.chef-cache-run-list-location	/var/lib/netwitness/config-management/nodes	string	Directory where chef stores the json from the previous chef runs
rsa.orchestration.engine.salt.component.chef-config-location	/var/lib/netwitness/config-management/client.rb	string	The location of the Chef solo configuration
rsa.orchestration.engine.salt.component.chef-log-location	/var/log/netwitness/config-management/chef-solo.log	string	The location to write the log file for a Chef solo execution
rsa.orchestration.engine.salt.component.chef-run-list-location	/etc/netwitness/config-management/node.json	string	The location to write the Chef run list
rsa.orchestration.engine.salt.component.descriptor-location	/etc/netwitness/component-descriptor/descriptor	string	The location of the component descriptor JSON file
rsa.orchestration.engine.salt.component.use-stable-package-versions	true	boolean	When enabled, the RPM version in the component descriptor file will be used. If disabled, the latest RPM version available in the yum repository will be used instead.

ProvisionHostProperties

Name	Default value	Type	Description
rsa.orchestration.engine.salt.task.provision-host.interval	5	seconds	The amount of time in between the availability checks.
rsa.orchestration.engine.salt.task.provision-host.ping-timeout	30	seconds	Timeout for ping
rsa.orchestration.engine.salt.task.provision-host.timeout	1	seconds	The amount of time to wait for the host/minion to connect to the Salt master. The checks will happen at the configured {@link #interval}

DeploymentProperties

Name	Default value	Type	Description
rsa.orchestration.platform.node-mode	active	string	

TaskExecutionProperties

Name	Default value	Type	Description
rsa.orchestration.task.async-keep-alive-time	10	long	maximum time that excess idle threads will wait for new tasks before terminating for async tasks
rsa.orchestration.task.async-keep-alive-time-unit		timeunit	Time unit for above
rsa.orchestration.task.async-pool-size	1	integer	the maximum number of threads to allow for asynchronous executor service
rsa.orchestration.task.cleanup-interval	7	seconds	The interval at which to delete existing Tasks
rsa.orchestration.task.freeze-detection-threshold	15	seconds	Time after which a stuck monitor thread is considered frozen

Name	Default value	Type	Description
rsa.orchestration.task.number-of-threads	5	integer	Number of threads in the pool for the sync execution service
rsa.orchestration.task.retain-task-duration	30	seconds	The amount of time to keep the jobs since its created. After the configured time elapses since the task created, task will be deleted.
rsa.orchestration.task.salt-job-schedule-wait-time	5	seconds	Grace period to wait for a salt job to be scheduled
rsa.orchestration.task.shutdown-timeout	30	seconds	The amount of time to wait, after receiving a shutdown request, for an executing task to finish before interrupting the execution.
rsa.orchestration.task.task-monitor-interval	30	seconds	Interval to check for new tasks and submit them
rsa.orchestration.task.update-interval	15	seconds	Amount of time between status check of currently running asynchronous jobs

Relay-server Configuration

NchanProperties

Name	Default value	Type	Description
rsa.netwitness.relay.agent-subscribe-channel	/agent/subscribe	string	Subscription nchan URL to get the responses published by endpoint-server
rsa.netwitness.relay.connection-timeout	60	seconds	Max time to wait for the connection to be created
rsa.netwitness.relay.max-connections	100	integer	Maximum number of connections allowed to nchan from relay server
rsa.netwitness.relay.nchan-base-url	https://localhost	string	URL of the Nchan service
rsa.netwitness.relay.read-timeout	120	seconds	Max time to wait for the response before throwing error
rsa.netwitness.relay.server-metrics-url	/nchan_stub_ status	string	URL to get the nchan metrics
rsa.netwitness.relay.server-publish-channel	/endpoint_ server/publish	string	Publish nchan URL to forward the requests and endpoint-server will read from there

RelayPkiProperties

Name	Default value	Type	Description
rsa.netwitness.relay.pki.certificate-path	/etc/pki/nw/service/import	string	location of the certificate and private key
rsa.netwitness.relay.pki.initial-delay	5	seconds	Time to wait before reloading the keystore if needed
rsa.netwitness.relay.pki.relay-cert-name	relay-server-cert.chain	string	the file name of the relay-server certificate chain
rsa.netwitness.relay.pki.relay-key-name	relay-server-key.pem	string	the file name of the relay-server private key

Respond-server Configuration

MigrationProperties

Name	Default value	Type	Description
rsa.migration.im-data-path	/opt/rsa/im	string	The location of the 10.x IM service
rsa.migration.max-retries	200	integer	Number of time respond attempts to run the migration in case unable to connect mongo or mongo is down.
rsa.migration.time-to-wait-between-retries	60	seconds	Frequency (in seconds) how often respond try to connect mongo

RespondPrimaryProperties

Name	Default value	Type	Description
rsa.primary.host	true	boolean	Determine whether the current respond service is running on the primary
rsa.primary.mode		respondprimaryproperties$scheduledjobsmode	Mode of current respond server

AlertRuleProperties

Name	Default value	Type	Description
rsa.respond.alertrule.batch-size	1000	long	The number of alerts to be processed by rule in a batch
rsa.respond.alertrule.counter-reset-interval-days	7	integer	How often should rule counters be reset
rsa.respond.alertrule.enabled	true	boolean	Alert rules enabled

Name	Default value	Type	Description
rsa.respond.alertrule.frequency	5	seconds	The frequency of the alert rule job
rsa.respond.alertrule.last-counter-reset-time	0	long	Timestamp for when the rule counters were reset

ArcherIntegrationProperties

Name	Default value	Type	Description
rsa.respond.archer.export.user-domain		string	Archer UserDomain, to be set only when LDAP is enabled on Archer

RespondCacheProperties

Name	Default value	Type	Description
rsa.respond.cache.user-cache-expiry	2	seconds	How often to query security server for the latest user information like their email
rsa.respond.cache.user-cache-size	1000	integer	Total size of the user cache

DataRetentionConfiguration

Name	Default value	Type	Description
rsa.respond.dataretention.enabled	false	boolean	Is the data retention job enabled
rsa.respond.dataretention.execution-hour	0	integer	Hour at which to run the job
rsa.respond.dataretention.frequency	24	seconds	How often should the job to delete old alerts/incidents run
rsa.respond.dataretention.retention-period	90	seconds	How long should alerts/incidents be stored

IndicatorAggregationJobConfig

Name	Default value	Type	Description
rsa.respond.indicatoraggregationrule.schedule-delay	0	long	Delay and frequency of indicator aggregation jobs
rsa.respond.indicatoraggregationrule.schedule-rate	5000	long	
rsa.respond.indicatoraggregationrule.seek-ahead-days	0	integer	How many days ahead should indicator aggregation go from incident window close time.
rsa.respond.indicatoraggregationrule.seek-back-days	1	integer	How many days back should indicator aggregation go from first alert received time when aggregating indicators

IntegrationExportProperties

Name	Default value	Type	Description
rsa.respond.integration.export.archer-exchange-name	incidents.archer	string	
rsa.respond.integration.export.archer-sec-ops-integration-enabled	false	boolean	
rsa.respond.integration.export.breach-integration-enabled	false	boolean	
rsa.respond.integration.export.escalation-settings		map	
rsa.respond.integration.export.export-incident-enabled	true	boolean	
rsa.respond.integration.export.help-desk-integration-enabled	false	boolean	

NormalizationProperties

Name	Default value	Type	Description
rsa.respond.normalization.alerts-queued	100	integer	The number of alerts to queue from rabbit before waiting to consume further The more you increase it, the higher chance of losing alerts if respond goes down during normalization
rsa.respond.normalization.custom-script-filename	custom_ normalize_ alerts.js	string	The name of the main custom JavaScript file used to normalize alerts.
rsa.respond.normalization.indicator-normalization-enabled	true	boolean	Determines whether the legacy and indicator bindings should be created or not
rsa.respond.normalization.max-legacy-consumers	10	integer	The maximum number of consumers that can consume from the legacy alerting exchange.
rsa.respond.normalization.script-directory	scripts	string	The name of the directory, relative to the service home directory, that contains the normalization JavaScript files.
rsa.respond.normalization.script-filename	normalize_ alerts.js	string	The name of the main JavaScript file used to normalize alerts.
rsa.respond.normalization.shutdown-timeout	30	seconds	The maximum amount of time to wait to finish processing alerts that have been received before shutting down the service.
rsa.respond.normalization.thread-count	4	integer	The number of threads to use to normalize and persist alerts.
rsa.respond.normalization.transient-indicator-normalization-enabled	true	boolean	Determines whether the low priority transient alerts binding should be created or not

QueryProperties

Name	Default value	Type	Description
rsa.respond.query.default-batch-size	100	long	Default chunk/batch size to send a stream of items to the client (client may override)
rsa.respond.query.default-query-limit	1000	long	Default number of items to send to the client in response to a single request (client may override)
rsa.respond.query.max-query-limit	5000	long	Max number of items to send to the client in response to a single request

RiskProcessingProperties

Name	Default value	Type	Description
rsa.respond.risk.alert.processing.concurrent-processors	4	integer	Concurrent number of staging that should be done.
rsa.respond.risk.alert.processing.context-limit	1000	integer	Maximum number of alert contexts per rule in a category
rsa.respond.risk.alert.processing.default-files	cmd.exe, powershell.exe, wscript.exe, cscript.exe, rundll32.exe	string	Name of files those are considered to be default OS provided files
rsa.respond.risk.alert.processing.page-size	100	integer	Page size for query while querying for persisted alerts
rsa.respond.risk.alert.processing.persisted-collection-interval	30	seconds	Interval at which alert collection should be queried for persisted alerts
rsa.respond.risk.alert.processing.staging-cleanup-interval	5	seconds	Cleanup interval for processed AlertRule from staging collection

Name	Default value	Type	Description
rsa.respond.risk.alert.processing.staging-fetch-size	5000	integer	Number of AlertRule to be fetched from staging in a single request
rsa.respond.risk.alert.processing.staging-work-interval	10	seconds	Frequency (in seconds) how often staged entries need to fetched for processing
rsa.respond.risk.alert.processing.track-file-name-change	false	boolean	Over time file-name might change for a hash, should that changes be tracked and latest name should be saved
rsa.respond.risk.alert.processing.track-host-name-change	true	boolean	Over time host-name might change for a host, should that changes be tracked and latest name should be saved

RiskCachingProperties

Name	Default value	Type	Description
rsa.respond.risk.caching.expiration-time	60	seconds	Time (in minutes) since last access of entry post which it will expire from cache.
rsa.respond.risk.caching.grouped-cache-expiration-time	5	seconds	Time (in minutes) since last access of entry post which it will expire from grouped cache
rsa.respond.risk.caching.grouped-cache-size	10000	integer	Max number of entries to be stored in the grouped cache
rsa.respond.risk.caching.size	500000	integer	Size of entries to be stored in cache

RiskRetentionProperties

Name	Default value	Type	Description
rsa.respond.risk.data.retention.frequency	1	seconds	Frequency to run the retention job
rsa.respond.risk.data.retention.retention-period	30	seconds	The retention threshold specified (in days)
rsa.respond.risk.data.retention.roll-up-to-day	false	boolean	Controls if the rollup-time needs to be calculate from start of the day when the task is executed.

RespondScheduledJobsProperties

Name	Default value	Type	Description
rsa.respond.scheduled.jobs.aggregation-job-enabled	true	boolean	Determine whether the aggregation job enabled/disabled
rsa.respond.scheduled.jobs.data-retention-job-enabled	true	boolean	Determine whether the data retention job enabled/disabled
rsa.respond.scheduled.jobs.risk-scoring-enabled	true	boolean	Determine whether the risk scoring functionality enabled/disabled

SecurIdIntegrationProperties

Name	Default value	Type	Description
rsa.respond.securid.alert-page-size	100	integer	Alerts are fetched from incidents pagewise. This property controls the maximum number of alerts to be fetched per page
rsa.respond.securid.alert-scan-json-paths	$.events[*]..	list	List of JSONPaths to scan the given userMetas in an alert. By default, it has just one JSONPath enough to read all direct occurrences of the given userMeta values from source and destination metas in all events in an alert.

Name	Default value	Type	Description
rsa.respond.securid.incident-processing-threads	3	integer	Number of threads to process the incident update events
rsa.respond.securid.max-incident-queue-size	100	integer	Maximum size of the queue used to hold the incident change events for processing.
rsa.respond.securid.secur-id-list-update-task-interval	15 minutes	seconds	Interval of the periodic task which updates the high-risk users' list in the SecurId cloud
rsa.respond.securid.secur-id-request-batch-size	100	integer	Maximum number of users to be sent in a single request to SecurId cloud.
rsa.respond.securid.user-meta	email_ address	string	The "respond specific" meta in an alert that identifies the user to be added to SecurID high-risk users' list Defaults to email_ address

Security-server Configuration

MigrationProperties

Name	Default value	Type	Description
rsa.migration.enable	true	boolean	Flag to be used in case of unit and integration tests for disabling migration check on first boot.
rsa.migration.ui-data-path	/backup/var/lib/netwitness/uax	string	The location of 10.x SA UI data

PasswordPolicyProperties

Name	Default value	Type	Description
rsa.security.account.password-policy.cannot-include-id	false	boolean	Can the password include the account identifier in it?
rsa.security.account.password-policy.min-chars	8	integer	The minimum number of characters the password must have.
rsa.security.account.password-policy.min-lower-chars	0	integer	The minimum number of lower-case characters the password must have.
rsa.security.account.password-policy.min-non-latin-chars	0	integer	The minimum number of non-latin characters a password must have
rsa.security.account.password-policy.min-numeric-chars	0	integer	The minimum number of numeric characters the password must have.
rsa.security.account.password-policy.min-special-chars	0	integer	The minimum number of special characters the password must have.
rsa.security.account.password-policy.min-upper-chars	0	integer	The minimum number of upper-case characters the password must have.
rsa.security.account.password-policy.passwords-expire-after	30	seconds	The maximum time an account password is valid before it must be changed.
rsa.security.account.password-policy.special-chars	~!@#$%^&*_-+=`$|(){}[]:;'"<>,.?/	string	The characters that are considered "special"

PkiAuthenticationProperties

Name	Default value	Type	Description
rsa.security.authentication.pki.enabled	false	boolean	This property controls if the PKI authentication feature is enabled
rsa.security.authentication.pki.retry-interval	1	seconds	This property controls how long to wait before retrying an initialization task.

AuthenticationPolicyProperties

Name	Default value	Type	Description
rsa.security.authentication.policy.account-lockout	20	seconds	Account lockout interval
rsa.security.authentication.policy.auto-create-external-users	true	boolean	Automatically create user profiles for external accounts
rsa.security.authentication.policy.case-sensitive	true	boolean	Case sensitivity of account identifiers
rsa.security.authentication.policy.max-successive-failures-before-lockout	5	integer	Accounts are locked after successive failures
rsa.security.authentication.policy.must-have-a-role	true	boolean	Require role mapping for external users

DeploymentProperties

Name	Default value	Type	Description
rsa.security.deployment.bootstrap-retry-interval	30	seconds	The amount of time to wait before retrying a failed bootstrap attempt.

SamlUserAccountProperties

Name	Default value	Type	Description
rsa.security.saml.account.user-account-type		accounttype	Account type of the users authenticated by SAML

Source-server Configuration

FeatureProperties

Name	Default value	Type	Description
rsa.features.file-policy-feature	true	boolean	rsa.usm feature toggle to control if creating File policies are supported
rsa.features.view-sources-feature	false	boolean	rsa.usm feature toggle to control if creating File policies are supported

UsmProperties

Name	Default value	Type	Description
rsa.usm.cache-duration	24	seconds	Time it takes for the cache that stores data to expire
rsa.usm.content-initialization-enabled	true	boolean	Enable content initialization task
rsa.usm.content-sync-interval	1	seconds	The interval between content sync tasks.
rsa.usm.database-initialization-enabled	true	boolean	Enable database initialization task
rsa.usm.page-size	1000	integer	Default page size for paged response
rsa.usm.policy-ip-to-uuid-migration-enabled	true	boolean	Enable policy ip to uuid migration
rsa.usm.source-count-sync-interval	10	seconds	The interval between source count sync tasks.

NetWitness® Platform XDR

Version 12.0

Storage Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Storage Overview **6**

Storage Requirements **7**

Drive Specifications 7

Required NetWitness Platform Storage Volumes 7

Performance Recommendations 9

Input/Output Operations Per Second 9

General Description of How NetWitness Platform Hosts Store Data 9

Prepare Physical Storage **10**

Configure Block devices for Drive pack 10

Benefits of Series 6/6E Drive Pack 10

Decoder Meta Use Cases 10

Concentrator Index Use Cases 11

Configure Block Devices for PowerVaults 11

Configure Block Device for Decoder / Log Decoder 12

Configure Block Device for Concentrator 14

Configure Storage 14

Configure Storage for Decoder / LogDecoder 14

Configure Storage for Concentrator 15

Enable Security on SED Capable Drives 15

Prepare Virtual or Cloud Storage **16**

Decoder, Log Decoder, Concentrator, Archiver 16

NW Server, ESA Primary, ESA Secondary and Malware Analysis 16

Log Collector 17

Endpoint Log Hybrid 17

Additional Endpoint Log Hybrid Partitions 21

UEBA 22

Configure Storage Using the REST API **23**

REST API Storage Configuration Commands 23

Storage Configuration Tasks 24

Task 1 - Attach Storage to the Host and Access the REST API Storage Commands 24

Task 2 - (Conditional) RAID Configuration for PowerVault and DACs 26

Task 3 - Allocate Block Devices to Partitions, Volume Groups, and Logical Volumes 29

Task 4 - Allocate Volume Groups to NetWitness Services - srvAlloc 31

Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture 31

Prepare Unity Storage **34**
Task 1 - Access Unisphere User Interface (UI) 35
Task 2 - Create Pools 36
Task 3 - Create LUNS 39
Task 4 - Register Hosts 41
Task 5 - Assign LUNS to Hosts 43
Task 6 - Install PowerPath 45
Migrate Data to Another Storage Type **47**
Migrate Data Using the Warm and Hot Tier Option 47
Stop the Service 47
Set Up PowerVault 47
Configure The Mount Points 48
Set up Warm and Hot Tiers 49
Decommision the DAC 51
Move Data From DAC to PowerVault 52
Data on PowerVault After Move from DAC 55
Appendix A. How NetWitness Platform Hosts Store Data **56**
Decoder Hosts 56
Concentrator Host 56
Archiver Host 57
Hybrid Hosts 57
Options for SAN Configurations 57
Performance Recommendations 57
Enable Security on SED Capable Drive groups on Host with a mix of SED and NON SED Drives 57
Appendix B. Encrypt a Series 6E Core or Hybrid Host (encryptSedVd.py) **61**
Enable SED on configured Drive Groups 63
Enable Virtual Drives / Drive Groups - PERC H740 (Mini) Adaptors (Internal storage) 66
Enable SED on configured Virtual Drives/ Drive Groups on Power Vault (PERC 840) 68
Enable Virtual Drives / Drive Groups - PERC H840 Adaptors 68
Appendix C. Troubleshooting **76**
Reconfigure Pre-Configured DAC Attached to Decoder Using REST API 76
Appendix D. Sample Storage Configuration Scenarios for 15-Drive DACs **77**
Configure Storage for Archiver 77
Configure Storage for Network (Packet) Decoder 80
Configure Storage for Network Concentrator 90
Configure Storage for Log Decoder Hybrid 96
Appendix E. Sample Storage Configuration Scenarios for 8 or 12-Drive PowerVault **101**
Configure Storage for Archiver using NW-PV-A/NW-PV-A-N 101

Configure Storage for Decoder using NW-PV-B/NW-PV-B-N 103
Configure Storage for Concentrator using NW-PV-C/NW-PV-C-N 109
Configure Storage for Concentrator using NW-PV-D/NW-PV-D-N 112
Appendix F. Sample Storage Configuration Scenarios for Meta Disk Kits 116
Meta-Only (One kit – 3 SED drives configured as RAID 5) 116
Maximize PowerVault Storage Capacity (One Kit – 3 SED Drives Configured as RAID 5) 121
Appendix G. Sample Storage Configuration for Concentrator Index with One Meta Disk Kit 130
Revision History 137

Storage Overview

This guide provides you with storage requirements and the instructions on how to allocate storage for physical (DACs, PowerVaults, Unity) and virtual storage devices for NetWitness Platform. It also includes the following topics.

- Detect Encryption on Existing PowerVault
- Migrate Data to Another Device

Refer to the following Hardware Setup Guides for information on how to connect these device to NetWitness Platform Core and Hybrid physical hosts:

- PowerVault (MD 1400) Setup Guide (see the "Enclosure Options" section of "Hardware Description") - NetWitness Community.
- 60-Drive DAC Setup Guide - NetWitness Community.
- 15-Drive DAC Setup Guide - NetWitness Community.

Storage Requirements

This section contains all the storage requirements needed to successfully attach storage to your NetWitness Platform deployment host systems. It contains the required drive types, appropriate volumes, and performance IOPS that are needed.

Drive Specifications

General specifications for core NetWitness Platform Hosts are:

- IO size 490/Dec
- Response/Latency < 20ms
- Decoder 10/90 read/write (low random I/O)
- Concentrator 50/50 read/write (high random I/O)

RAID Group	Suitable Volumes
NL-SAS or 10K SAS	All Packet Decoder volumes All Log Decoder volumes All Archiver volumes Concentrator meta volume
SSD	Concentrator index volume

Required NetWitness Platform Storage Volumes

Service Volume Names

Service	Volume Name	File Systems Created
Network Decoder	decoder	packetdb
Network Decoder	decodersmall	decoder root, index, sessiondb, metadb
Log Decoder	logdecoder	packetdb
Log Decoder	logdecodersmall	logdecoder root, index, sessiondb, metadb
Concentrator	concentrator	concentrator root, metadb, sessiondb
Concentrator	index	index
Archiver	archiver	database

Volume Sizing

The volume sizes below are automatically created when using the NetWitness Platform storage tool, described in Configure Storage Using the REST API.

Volume	Filesystem	Mount Point	Size
decodersmall	decoroot	/var/netwitness/decoder	10 GB
decodersmall	index	/var/netwitness/decoder/index	30 GB
decodersmall	sessiondb	/var/netwitness/decoder/sessiondb	600 GB
decodersmall	metadb	/var/netwitness/decoder/metadb	100% of free space on decodersmall volume
decoder	packetdb	/var/netwitness/decoder/packetdb	100% of free space on decoder volume
logdecodersmall	decoroot	/var/netwitness/logdecoder	10 GB
logdecodersmall	index	/var/netwitness/logdecoder/index	30 GB
logdecodersmall	sessiondb	/var/netwitness/logdecoder/sessiond	600 GB
logdecodersmall	metadb	/var/netwitness/logdecoder/metadb	100% of free space on logdecodersmall volume
logdecoder	packetdb	/var/netwitness/logdecoder/packetdb	100% of free space on logdecoder volume
concentrator	root	/var/netwitness/concentrator	30 GB
concentrator	sessiondb	/var/netwitness/concentrator/sessiondb	10% of free space on concentrator volume
concentrator	metadb	/var/netwitness/concentrator/metadb	100% of free space on concentrator volume
index	index	/var/netwitness/concentrator/index	100% of free space on index volume
archiver	database	/var/netwitness/archiver/database	100% of free space on archiver volume

Performance Recommendations

NetWitness recommends that Packet and Log Decoders receive two LUNs or Block Devices, one for Packet data, the other for all other databases. This allows you to segregate the high-bandwidth Packet Database from the other databases so they do not compete for I/O bandwidth with other activity.

Concentrators require a separate SSD-based index volume for best performance. You must house this index volume on a different RAID group than the Concentrator Meta database volume, which you can stored on NL-SAS. Archivers can use a single large NL-SAS storage volume per appliance.

Input/Output Operations Per Second

The following table lists the IOPS requirements for the Decoder and Concentrator hosts.

Logs	Log Decoder	Concentrator
10K EPS	400	8,000
20K EPS	550	10,300
25K EPS	1,200	10,800

Packets	Network Decoder	Concentrator
1Gbps	600	6,050
2 Gbps	950	8,300
4 Gbps	1,650	12,800
6 Gbps	2,400	17,300
8 Gbps	3,200	21,800

General Description of How NetWitness Platform Hosts Store Data

For information about how NetWitness Platform hosts store data, see Appendix A. How NetWitness Platform Hosts Store Data.

Prepare Physical Storage

IMPORTANT: NetWitness recommends you to create a block device for NetWitness Storage.

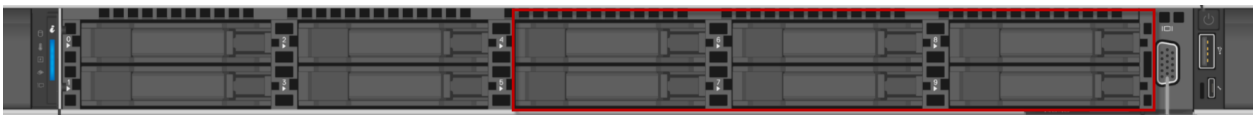
This section provides two options to configure block device:

- Configure Block devices for Drive pack
- Configure Block Devices for PowerVaults

Note: Block device is also referred to as Virtual Drive or Drive Group

Configure Block devices for Drive pack

You can add additional drives to the Series 6 or 6E appliances to accommodate various use cases. These drives provide the capability for the decoder meta or concentrator index volumes to reside on the appliance. A minimum of 2 drives and a maximum of 6 drives are possible. The number of drives will depend on how much meta cache or index is needed.



Benefits of Series 6/6E Drive Pack

- **Maximize PowerVault Storage Capacity** - Traditionally, PowerVault storage allocates a volume for the Decoder metadata. This reduces the usable storage on the PowerVault. Drive Packs reduce this issue by providing 20TB of extra usable PV storage.
- **Reduces Cost for Meta Only Use Case** - For metadata-only deployments, drive pack fits for a customer who want to purchase hardware from NetWitness. This provides more cost-effective solution, because a drive pack can substitute a PowerVault.
- **Enable existing deployments to utilize compression options**.
- **Provides capability for expanding meta keys and associated indexing**.

Decoder Meta Use Cases

- Meta-Only
- Maximize Power Vault Storage

Two or more 2.4TB 10K SAS SED drives can be added to a Decoder for the decodersmall or logdecodersmall volumes. These volumes are used to store the meta cache on the Decoders.

Both the Log Decoders and Network Decoders parse out meta data from the raw captured traffic. The meta data is then aggregated to a Concentrator for indexing.

The host requires storage to store a cache for the meta extracted during the data capture for Concentrator aggregation. The meta cache on a Decoder is generally fixed in size, but you can expand it to support additional cache to avoid possible connectivity loss between the Decoder and the corresponding Concentrator.

Typically, the decodersmall or logdecodersmall volumes are stored on the first three drives of the first and second (10G config only) PowerVault enclosures. By utilizing the drive pack option, these three drives can instead be used for the packetdb (maximizing Power Vault storage).



For meta-only scenarios, the decodersmall volume would be stored on the drive pack, therefore eliminating the need for a Power Vault.

Concentrator Index Use Cases

- Support Additional Meta-Key Indexing
- Capability to Enable compression for Existing Deployments

Two or more 3.84 TB SSD SED drives can be added to a Concentrator to increase the index volume. The index storage needs are scaled based on the NetWitness Platform deployment retention requirements. If additional meta keys are enabled and indexed, it may impact index retention.

For existing deployments, an SSD index drive pack is required if you need to enable compression. When compressing the packetdb and metadb, additional index is needed to support compression of those databases.

See Appendix F. Sample Storage Configuration Scenarios for Meta Disk Kits for sample storage configuration scenarios for Meta Disk Kits.

Configure Block Devices for PowerVaults

The Physical, Virtual, or Cloud NetWitness hosts for Decoders, Log Decoders, Concentrators, and Archivers need block storage attached. Make sure that the allocated storage meets all of the storage requirements. Specifically, make sure that the required storage volumes are created (for more information, see Required NetWitness Platform Storage Volumes in Storage Requirements), and:

- At least two block devices are created for Decoders (Meta, Session and Packet volumes)

Note: The larger block devices hold the packet volume, and the smaller block devices hold the meta and session.

- At least two block devices are created for Concentrators (Index and Meta volumes).
- Ensure that the block device meets the minimum IOPS for expected ingestion rates.

Configure Block Device for Decoder / Log Decoder

While creating the block device RAID configuration, the best practice is to configure a RAID 6 for the larger NL-SAS drives and RAID 5 or 1 for any 10k SAS or SSD type drives.

1. Identify the controller ID (Ctl) for '**PERC H840P Adaptor**'.

   `/opt/MegaRAID/perccli/perccli64 show`

   In the below figure the controller ID is '**1**' corresponds to '**PERCH840PAdaptor**'.

```
[root@116S6Core1 ~]# /opt/MegaRAID/perccli/perccli64 show
Status Code = 0
Status = Success
Description = None

Number of Controllers = 2
Host Name = 116S6Core1
Operating System  = Linux 3.10.0-1160.21.1.el7.x86_64
StoreLib IT Version = 07.0400.0200.0400
StoreLib IR3 Version = 15.50-0

System Overview :
===============

-----------------------------------------------------------------------------
Ctl Model           Ports PDs DGs DNOpt VDs VNOpt BBU sPR DS EHS ASOs Hlth
-----------------------------------------------------------------------------
  0 PERCH740PMini       8  10   2     0   2     0 Opt On  -  N      0 Opt
  1 PERCH840Adapter     8  12   0     0   0     0 Opt On  -  N      0 Opt
-----------------------------------------------------------------------------

Ctl=Controller Index|DGs=Drive groups|VDs=Virtual drives|Fld=Failed
PDs=Physical drives|DNOpt=DG NotOptimal|VNOpt=VD NotOptimal|Opt=Optimal
Msng=Missing|Dgd=Degraded|NdAtn=Need Attention|Unkwn=Unknown
sPR=Scheduled Patrol Read|DS=DimmerSwitch|EHS=Emergency Hot Spare
Y=Yes|N=No|ASOs=Advanced Software Options|BBU=Battery backup unit
Hlth=Health|Safe=Safe-mode boot
```

2. Identify the Enclosure ID (**EID**) for controller '**1**'. In this case the **EID** is '**247**'.

   `/opt/MegaRAID/perccli/perccli64 /c1 /eall show`

```
[root@116S6Core1 ~]# /opt/MegaRAID/perccli/perccli64 /c1 /eall show
Controller = 1
Status = Success
Description = None


Properties :
==========

-------------------------------------------------------------------
EID State Slots PD PS Fans TSs Alms SIM Port# ProdID VendorSpecific
-------------------------------------------------------------------
247 OK       12 12  2    4   7    0   2 01 x4 MD1400
-------------------------------------------------------------------

EID-Enclosure Device ID |PD-Physical drive count |PS-Power Supply count|
TSs-Temperature sensor count |Alms-Alarm count |SIM-SIM Count
```

3. Identify the existing block devices on the host. The block device name is identified under **NAME** column. The block devices names shown below are **sda**,**sdb** and **sdc**. Use '**lsblk**' to list the block devices.

```
[root@116S6Core1 ~]# lsblk
NAME                          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                             8:0    0  1.1T  0 disk
├─sda1                          8:1    0    1G  0 part /boot
└─sda2                          8:2    0  1.1T  0 part
  ├─netwitness_vg00-root      253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0   10G  0 lvm  /home
sdb                             8:16   0  2.2T  0 disk
└─sdb1                          8:17   0  2.2T  0 part
  └─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
sdc                             8:32   0  8.7T  0 disk
[root@116S6Core1 ~]#
```

4. Create the Virtual Drive or Drive Group (**DG**) on PERCH840PAdaptor using disks in slot 0 through 9 (for example, all the drives) using below command.

> **Warning:** Every decoder needs a logdecodersmall or decodersmall volume for meta. This example assumes the meta volume already exists on another PowerVault or Drive Pack. If this enclosure will account for the meta volume, the first two or 3 drives would need to be allocated for the meta volume block device. Another block device would need to be created with the remaining drives for the packetdb volume.

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd type=raid6 drives=247:0-11
strip=128 force
```

```
[root@116S6Core1 ~]# /opt/MegaRAID/perccli/perccli64 /c1 add vd type=raid6 drives=247:0-11 strip=128 force
Controller = 1
Status = Success
Description = Add VD Succeeded


[root@116S6Core1 ~]#
```

5. To view the Virtual Drive created in the above step:

```
/opt/MegaRAID/perccli/perccli64 /c1 /vall show
```

```
[root@116S6Core1 ~]# /opt/MegaRAID/perccli/perccli64 /c1 /vall show
Controller = 1
Status = Success
Description = None


Virtual Drives :
==============

-------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache Cac sCC       Size Name
-------------------------------------------------------------
0/0   RAID6 Optl  RW     No      RWBD  -   OFF 106.918 TB
-------------------------------------------------------------

Cac=CacheCade|Rec=Recovery|OfLn=OffLine|Pdgd=Partially Degraded|Dgrd=Degraded
Optl=Optimal|RO=Read Only|RW=Read Write|HD=Hidden|TRANS=TransportReady|B=Blocked|
Consist=Consistent|R=Read Ahead Always|NR=No Read Ahead|WB=WriteBack|
FWB=Force WriteBack|WT=WriteThrough|C=Cached IO|D=Direct IO|sCC=Scheduled
Check Consistency
```

6. Identify the new block device on the host. The block device name is identified under **NAME** column. The new block device corresponding to the above virtual drive is **sdd**. This block device name is required when configuring storage. Use '**lsblk**' to list the block devices.

   `lsblk`

```
[root@116Decoder perccli]# lsblk
NAME                         MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                            8:0    0   931G  0 disk
├─sda1                         8:1    0     1G  0 part /boot
└─sda2                         8:2    0   930G  0 part
  ├─netwitness_vg00-root     253:0    0    30G  0 lvm  /
  ├─netwitness_vg00-swap     253:1    0     4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome   253:2    0   2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog   253:3    0    10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome  253:4    0    10G  0 lvm  /home
sdb                            8:16   0   1.8T  0 disk
└─sdb1                         8:17   0   1.8T  0 part
  └─netwitness_vg00-nwhome   253:2    0   2.7T  0 lvm  /var/netwitness
sdc                            8:32   0   8.7T  0 disk
sdd                            8:48   0 106.9T  0 disk
[root@116Decoder perccli]#
```

7. You must Configure Storage for Decoder / LogDecoder and Concentrator Configure Storage to complete the configuration.

Configure Block Device for Concentrator

Block Devcies must be configured on PowerVaults before configuring the PowerVaults as storage to Concentrator. The Block Devices can be configured using the steps similar to Configure Block Device for Decoder / Log Decoder using percli64 utility. Use the SSD drives for index and the remaining drives for the Meta or Session DB.

Configure Storage

Configure Storage for Decoder / LogDecoder

Use REST API tool to configure the above block devices or virtual drives as Storage for Decoder / Log Decoder or Concentrator. For more information, see *Storage Configuration Tasks* (Task 3 and Task 4) for Decoder or Log Decoder and Task 1 to Task 5 for Concentrator in Configure Storage Using the REST API topic.

Service	Controller	Volume	Block Device
Decoder / Log Decoder	PERC H740 Mini Adaptor	decodersmall	Refer to step 7 in Prepare Physical Storage (Decoder / Log Decoder) in this example the block device is '**sdc**'.
Decoder / Log Decoder	PERC H840 Adaptor	decoder	Refer to Step 6 in Configure Block Device for Decoder / Log Decoder in this example the block device is '**sdd**'.

Configure Storage for Concentrator

Use REST API tool to configure the block devices created on Drive Pack and/or Power Vaults. The block devices created on SSD is allocated to Index database and the one created on HDD to Meta/Session Database. See Storage Configuration Tasks (Task 3 and Task 4) for Concentrator in Configure Storage Using the REST API topic.

Enable Security on SED Capable Drives

To enable Security on the SED Capable Drive Group on PERC H740 Mini and PERC H840 Adaptors, see Appendix B. Encrypt a Series 6E Core or Hybrid Host (encryptSedVd.py) .

Prepare Virtual or Cloud Storage

This section describes how to set up virtual or cloud storage for the following types of component hosts:

- Decoder, Log Decoder, Concentrator, Archiver
- NW Server, ESA Primary, ESA Secondary and Malware Analysis
- Log Collector
- Endpoint Log Hybrid
- Additional Endpoint Log Hybrid Partitions
- UEBA

Decoder, Log Decoder, Concentrator, Archiver

Virtual or Cloud NetWitness hosts for Decoders, Log Decoders, Concentrators, and Archivers need block storage attached. Make sure that the allocated storage meets all of the storage requirements. Specifically, make sure that the required storage volumes are created (see "Required NetWitness Platform Storage Volumes" in Storage Requirements), and:

- At least two Block Devices are created for Decoders (meta /session and packet volumes)
- At least two block devices are created for Concentrators (index and meta volumes)
- Ensure that block devices can meet the minimum IOPS for expected ingestion rates

Attach the allocated storage to the NetWitness host by following the hosting platforms native procedure.

- VmWare – Vsphere Console (add disk to VM)
- Hyper-V – Manager Console (add disk to VM)
- Azure – Add Managed Disks to virtual instance
- AWS – Add EBS Storage to virtual instance
- Google Cloud Platform (GCP) - Add storage to virtual instance

Once the disk is attached to the virtual host, configure block devices for services. For more information, see Virtual Hosts Additional Storage Playbook.

NW Server, ESA Primary, ESA Secondary and Malware Analysis

For an extension of `/var/netwitness/` partition, attach an external volume.

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `/dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`

3. `lvresize --resizefs --extents +100%FREE /dev/netwitness_vg00/nwhome`
4. `xfs_growfs /dev/netwitness_vg00/nwhome`

NetWitness recommends the following partition definitions. However, you can change these values based on the retention days.

LVM	Folder	Block Storage
/dev/netwitness_vg00/nwhome	/var/netwitness/	Refer to the Cloud Provider Block Storage setup (storage) tables.

Log Collector

For an extension of `/var/netwitness/` partition, attach an external volume

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvresize --resizefs --extents +100%FREE /dev/netwitness_vg00/nwhome`
4. `xfs_growfs /dev/netwitness_vg00/nwhome`

NetWitness recommends the following partition definitions. However, you can change these values based on the retention days.

LVM	Folder	Block Storage
/dev/netwitness_ vg00/nwhome	/var/netwitness/	Refer to the Cloud Provider Block Storage setup (storage) tables.

Endpoint Log Hybrid

The total disk size required depends on the data retention period. You can use the below per day disk usage indicative values to calculate the required disk size for your deployment. For example, to retain 30 days of data, multiply the below per day disk usage values with 30.

The following table provides disk usage for one full scan. The full scan disk usage values are based on the below event count:

- Files count -1100
- Processes count -100
- Dlls count - 500
- Drivers count -150
- Services count - 500
- Tasks count -100

Endpoint Log Hybrid(50K Advance Agents - Disk usage per full scan)					
	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	220 GB	12 GB	5 GB	NA	237 GB
Concentrator	230 GB	NA	5 GB	6 GB	241 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 30 GB (Subsequent per scan increase)

The following tables provide per day disk usage for tracking data. The total tracking events per agent per day is 29000.

Endpoint Log Hybrid (50K Advance Agents - Tracking data without Expanded Network Visibility)					
	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	1500 GB	140 GB	46 GB	NA	1,686 GB
Concentrator	1600 GB	NA	46 GB	30 GB	1,676 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 1.5 GB (Tracking data per day increase)

The following tables provide per day disk usage for tracking data. Total tracking events per agent per day is 33000

Endpoint Log Hybrid (50K Advance Agents - Tracking data with Expanded Network Visibility)					
	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	1800 GB	152 GB	55 GB	NA	2007 GB
Concentrator	1900 GB	NA	55 GB	36 GB	1991 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 1.5 GB (Tracking data per day increase)

The following table provides per day disk usage for insight agents. The total tracking data per agent per day is 10800 plus 1 full scan daily.

Endpoint Log Hybrid (50K Insights Agents with Expanded Network Visibility)					
	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	500 GB	52 GB	18 GB	NA	570 GB

Endpoint Log Hybrid (50K Insights Agents with Expanded Network Visibility)					
Concentrator	600 GB	NA	18 GB	13 GB	631 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 30 GB (Subsequent per scan increase)

The following table provides Endpoint Agents sizing based on the feature.

Feature	Description	Agent or Endpoint Server
Endpoint Only	Only scan and tracking data	Maximum 50K Endpoint Agents only
Windows Logs Only	Only Windows Logs from agents. Assuming 20K events per second supported by Hybrid.	Maximum 20K Agents: • Generates 20K log events per second
File Collection Only	Only File Collection from agents. Assuming 20K events per second supported by Hybrid	Maximum 20K Agents : • Generates 20K log events per second
Endpoint and Windows Logs	Event per second per agent • (For Windows Logs) 1 event sent by 1 agent every second • (For Tracking Events) 0.4 event sent by 1 agent every second • 20K events per second supported by Hybrid **Note:** Total agents should be calculated as below: Hybrid events per second/ (Windows Logs Endpoint Server of 1 agent + Tracking Event Endpoint Server for 1 agent) For example, 20000 / (1.0 + 0.4)	Maximum 15K (approximately) Agents: • Generates 15K (approximately) Windows log events Plus • Generates 15K (approximately) Agents EDR data

Feature	Description	Agent or Endpoint Server
Endpoint, Windows Logs and File Collection	Event per second per agent: • (For Windows Logs) 1 event sent by 1 agent every second • (For Tracking Events) 0.4 event sent by 1 agent every second • (For File Collection) 1 event sent by 1 agent every second • 20, 000 events per second supported by Hybrid **Note:** Total agents should be calculated as below: Hybrid events per second/ (Windows Logs Endpoint Server of 1 agent + Tracking Event Endpoint Server for 1 agent + File Collection) For example, 20000 / (1.0 + 1.0 + 0.4)	Maximum 10K (approximately) Agents: • Generates 10K (approximately) Windows log events Plus • Generates 10K (approximately) Endpoint Agents data Plus • Generates 10K (approximately) Agents File Collection data

Extending File Systems

For Endpoint Server, attach external disk for extension of `/var/netwitness/` partition, create an external disk with suffix as `nwhome`.

Follow these steps:

1. Ensure you have added a new disk. For more information, see "Task 1. Add New Disk" in the *Virtual Hosts Installation Guide for NetWitness Platform*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
2. Execute `lsblk` and get the physical volume name
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvresize --resizefs --extents +100%FREE /dev/netwitness_vg00/nwhome`
6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

NetWitness recommended partition for Endpoint Server (can be changed based on the retention days).

LVM	Folder	Size	Disk Type
`/dev/netwitness_vg00/nwhome`	`/var/netwitness/`	6TB	HDD

For Mongo DB, attach external disk for extension of `/var/netwitness/mongo` partition, create an external disk with suffix as `nwhome`.

Follow these steps:

1. Ensure you have added a new disk. For more information, see "Task 1. Add New Disk" in the *Virtual Hosts Installation Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.
2. Execute `lsblk` and get the physical volume name
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc1`
4. `vgextend hybrid /dev/sdc1`
5. `lvresize --resizefs --extents +100%FREE /dev/hybrid-vlmng`
6. `xfs_growfs /dev/mapper/hybrid-vlmng`

NetWitness recommended partition for Mongo DB (Can be changed based on the retention days). Minimum recommended size for var/netwitness is 500 GB.

LVM	Folder	Size	Disk Type
`/dev/hybrid-vlmng`	`/var/netwitness/mongo`	6TB	HDD

Additional Endpoint Log Hybrid Partitions

The following partition should be on the volume group endpoint and should be in a single RAID 0 array.

Folder	LVM	Volume Group
/var/netwitness/mongo	hybrid-mongo	endpoint
/var/netwitness/concentrator	concentrator-concroot	endpoint
/var/netwitness/concentrator/index	hybrid-concinde	endpoint
/var/netwitness/logdecoder	hybrid-ldecroot	endpoint

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md0`
2. `vgcreate -s 32 endpoint /dev/md0`
3. `lvcreate -L <disk_size> -n <lvm_name> endpoint`
4. `mkfs.xfs /dev/ endpoint /<lvm_name>`
5. Repeat the above steps for all the LVMs mentioned.

NetWitness recommends the following partitions. However, you can change these values based on the retention days.

LVM	Folder	Block Storage
/dev/netwitness_vg00/nwhome	/var/netwitness/	Refer to the Cloud Provider Block Storage setup (storage) tables.

LVM	Folder	Block Storage
/dev/endpoint/hybridmongo	/var/netwitness/mongo	Refer to the Cloud Provider Block Storage setup (storage) tables.
/dev/endpoint/concentratorconcroot	/var/netwitness/concentrator	Refer to the Cloud Provider Block Storage setup (storage) tables.
/dev/endpoint/hybridconcinde	/var/netwitness/concentrator/index	Refer to the Cloud Provider Block Storage setup (storage) tables.
/dev/endpoint/hybridldecroot	/var/netwitness/logdecoder	Refer to the Cloud Provider Block Storage setup (storage) tables.

UEBA

The following procedure attaches an external disk and extends the `/var/netwitness/` partition. You must use `nwhome` as the eternal disk suffix. This procedure illustrates how to add a 2TB disk.

Note: `/var/netwitness` is the only partition that can reside on this volume.

1. List the physical volume name.
   `lsblk` (for example, `dev/mapper/sdc`)
2. Extend the `/var/netwitness/` partition.

```
pvcreate <pv_name>where pv_name is dev/mapper/sdc
vgextend netwitness_vg00 /dev/mapper/sdc
lvresize --resizefs --extents +100%FREE /dev/mapper/netwitness_vg00/nwhome
xfs_growfs /dev/mapper/netwitness_vg00-nwhome
```

This partition is the NetWitness recommended partition for UEBA. You can change it based on retention days.

Configure Storage Using the REST API

In NetWitness Platform 11.3 and later releases, you use the REST API for all storage configuration operations. For information about how to use the REST API, see the *RESTful API User Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

REST API Storage Configuration Commands

Each of the commands listed below has built-in help that describes their function and usage. If you are using the REST interface, select the command from the drop-down menu to see the help text. For examples of REST API storage configuration commands, see Appendix D. Sample Storage Configuration Scenarios for 15-Drive DACs.

Commands for Direct-Attached RAID Volumes

- `raidList` : List the RAID controllers and direct-attach enclosures that are present on this host.
- `raidNew` : Allocate direct-attached enclosures to block devices.

Commands for Allocating Block Devices as Storage

- `devlist` : List available block devices on the host.
- `partNew` : Allocate partitions on a block device and create volume groups.
- `vgs` : Summarize how block devices are organized into volume groups.

Commands for Allocating Storage to Services

- `srvList` :List services on the host and their allocated storage paths.
- `srvAlloc` : Allocate a volume group to a service.
- `srvFree` : Remove a volume group from a service.
- `multipath-II` : To verify if SAN devices are attached.

Command to Reconfigure Services to Detect and Use All of the New Storage

- `reconfig` - After configuring new storage, detect and use new storage on the associated service and database.

Storage Configuration Tasks

Task 1 - Attach storage to the host and access the REST API storage configuration commands.
Task 2 - (Conditional) Configure RAID if necessary.
Task 3 - Allocate block devices to partitions, volume groups, and logical volumes.
Task 4 - Allocate volume groups to NetWitness services.
Task 5 - Reconfigure services and databases to detect and appropriately use new storage.

Task 1 - Attach Storage to the Host and Access the REST API Storage Commands

IMPORTANT: Task 1 is not applicable for NetWitness version 11.5.0.0 and 11.5.0.1.

Complete the following steps to attach an external storage device to a host and access the storage configuration commands available through the REST API.

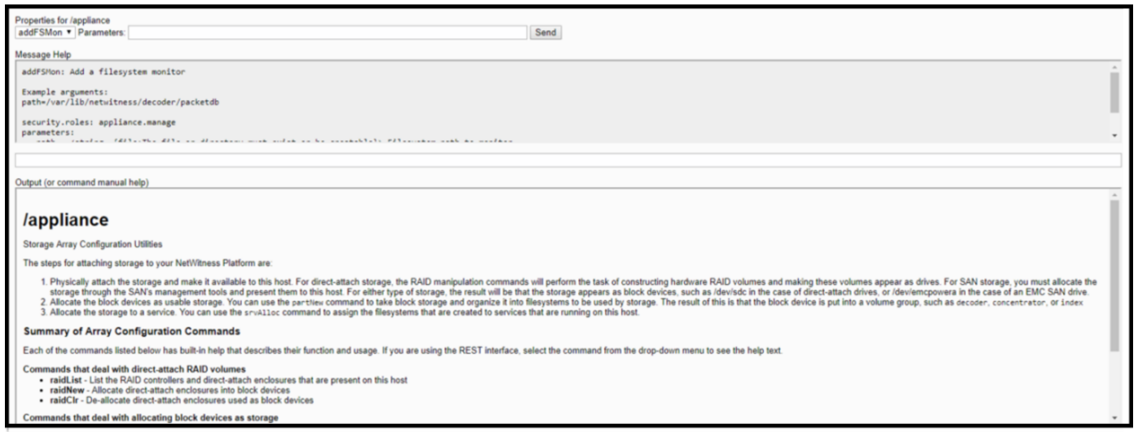
1. Attach the storage and make it available to this host.
   - To attach PV storage, refer to the *PowerVault (Dell MD 1400) Setup Guide*.
   - For third-party storage, create the RAID groups to match the volumes listed in Storage Requirements
2. There are two ways that you can access the REST API storage commands: from a Browser, or from the **Services** > **Explore** view from the User Interface.

   Note: Once you have accessed the REST API, the steps that you perform are the same, no matter which method you used to access it.

- From a Browser.
  a. Open a Browser and specify the ip-address of the host with port **50106**.
  The following example is the Decoder, but you need to use port 50106 for any host hardware for which you are configuring storage using the REST API.
  `https://<decoder-ip-address>:50106`
  b. Log in with the `admin` account credentials.
  The following REST API menu is displayed.



  c. Click on the `(*)` next to **appliance** to access the REST command set.
  The **Properties for /appliance** dialog is displayed under the initial REST menu. The **Output (or command manual help)** section describes the commands that the REST API can send to the device, their usage, and their parameters.



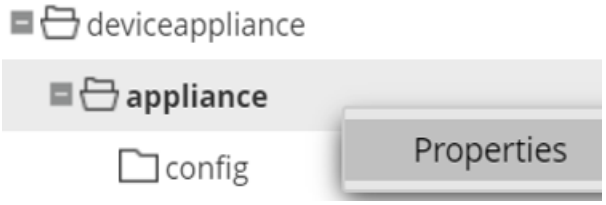
- From the User Interface.

  a. In the **NetWitness** menu, go to  (Admin) > **SERVICES**.

  b. Select the service (for example, a Concentrator).

  c. Under  (actions), select **View** > **Explore**.

  d. Navigate to **deviceappliance/appliance**, right click, and click **Properties**.



Note: If you are on NetWitness version 11.5.0.0 or 11.5.0.1, you must navigate to **System** > **Host Tasks** >**Task**.

You can now access the storage commands from the **Properties** dialog.

3. Proceed to:

- Task 2 if you need to configure RAID for PowerVault or DACs.
- Task 3 if you do not need to configure RAID and already have a block device available.

Task 2 - (Conditional) RAID Configuration for PowerVault and DACs

IMPORTANT: Task 2 is mandatory if you are on NetWitness version 11.5.0.0 or 11.5.0.1.

NetWitness Platform hardware uses direct-attached SAS drives for storage. These drives are housed in a SAS enclosure. SAS enclosures are shelves of drives attached to the NetWitness node by a cable connected to the SAS host bus adapter.

SAS enclosures are also known as other names, such as "DAC" (Direct-Attached Capacity), or "JBOD" (Jumbo Box of Disks), or "Dell PowerVault".

NetWitness Platform utilizes Dell PERC SAS host bus adapters. NetWitness Platform devices typically include two SAS host bus adapters. One is used for controller drives that are internal to the NetWitness Node, and another is used for controlling drives attached to the SAS enclosures. The internal controller and drives are configured when the node is built, but the external SAS enclosures are not. You execute the `raidList` and `raidNew` commands to identify and configure the external SAS enclosures.

These commands work with the following SAS enclosure types:

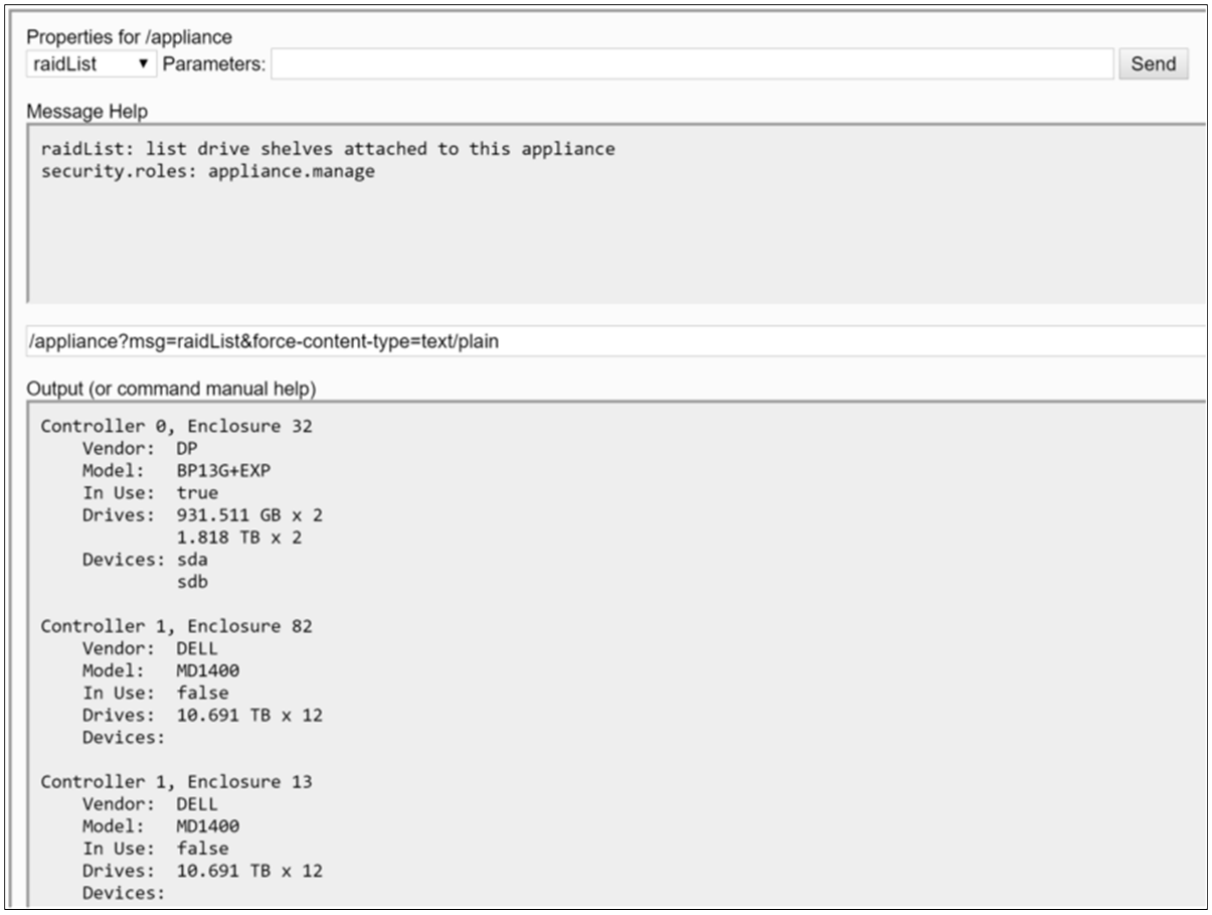
- EMC ESAS 15-drive enclosures
- EMC ESAS 60-drive enclosures
- Dell PowerVault 12-drive enclosures
- Dell PowerVault 8-drive enclosures

Note: EMC 60-drive enclosures are logically organized as four separate 15-drive sub-enclosures. They behave as if there are four 15-drive enclosures, each of which can be configured independently.

The `raidList` and `raidNew` commands operate on entire enclosures. Execute `raidList` to identify the enclosues. execute `raidNew` to configure an enclosure to perform one of the pre-determined roles within a NetWitness Platform node.

After you attach storage to the host and access the REST API storage commands, complete the following steps to create RAID if required.

1. Execute the `raidList` command to identify the controllers and enclosures that are attached to the system.
   In the following example, Controller 1 does not display any block devices. This indicates the array is not configured.



2. Select a RAID layout scheme for the Enclosure.

The following table lists the PV to Supported Hosts Mapping.

Type	SKU	Specification	Supported Hosts
High Performance	NW-PV-C	Dell Storage MD 1400 6 x 12TB NL-SAS SED, 2 x 3.8TB SSD SED	Concentrator

Type	SKU	Specification	Supported Hosts
High Performance	NW-PV-D	Dell Storage MD 1400 9 x 12TB NL-SAS SED, 3 x 3.8TB SSD SED	Concentrator
High Density	NW-PV-A	Dell Storage MD1400 12 x 12 TB NL-SAS SED	Decoder, LogDecoder, Archiver, hybrid (log & packet)
High Density	NW-PV-B	Dell Storage MD 1400 8 x 12TB NL-SAS SED	Decoder, LogDecoder, Archiver, hybrid (log & packet)

The following tables show you the supported allocation schemes.

Note: For RAID configuration, when the Decoder is configured for 10G capture, use the `decoder` scheme for the **first two enclosures** and the `archiver` scheme for subsequent enclosures.
When you are not configuring for 10G capture, use the `decoder` scheme for the **first enclosure** and the `archiver` scheme for subsequent enclosures.
These configurations will maximize storage capacity and performance.

Scheme	Drives Required	Allocation
decoder or logdecoder	8 or 12 or 15 HDDs	3x drives in RAID 5 for decodersmall or logdecodersmall, all remaining drives in RAID 5
archiver	8 or 12 or 15 HDDs	All drives in RAID 6 for archiver or decoder database volume
networkhybrid	8 or 12 or 15 HDDs	3x drives in RAID 5 for meta expansion, all remaining drives in RAID 5 for packet expansion
loghybrid	8 or 12 or 15 HDDs	Half of the drives in RAID 5 for meta expansion, half the drives in RAID 5 for packet expansion
concentrator	2 or more SSDs, 4 or more HDDs	All SSDs in RAID 1 or RAID 5 for index, all HDDs in RAID 6 for meta

3. After the controller, enclosure, and scheme are identified, execute the `raidNew` command to create RAID Volumes. For example:
   ```
   send /appliance raidNew controller=1 enclosure=82 scheme=decoder preferSecure=false
   ```
   Add the `commit=1` parameter to actually execute this operation. Execute the `raidList` command to list the created block devices.
4. (Optional) Configure SEDs (Self-Encrypting Drives). If the `raidNew` command detects self-encrypting drives and a security key has been set on the controller, the `raidNew` command will attempt to create a secure array. To set a security key on the controller, execute the `raidKey` command. For example:
   ```
   send /appliance raidKey controller=1 key=myPasssphrase keyId=1
   ```
   - To create a secured (that is, encrypted) array on physical devices attached to a controller with a security key set, specify `preferSecure=true` when using `raidNew`

- To create an unsecured (that is, unencrypted) array on physical devices attached to a controller with a security key set, specify `preferSecure=false` when using `raidNew`.

5. Go to Task 3 - Allocate Block Devices to Partitions, Volume Groups, and Logical Volumes, after you create RAID volumes.

Task 3 - Allocate Block Devices to Partitions, Volume Groups, and Logical Volumes

The `partNew` command prepares a storage device to use in NetWitness Platform. It performs the following tasks.

- Creates the partition table on the block device.
- Creates the Linux Volume Manager physical device partition.
- Creates a volume group containing the physical device.
- Creates logical volumes in the volume group.
- Creates XFS filesystems on each logical volume.
- Creates `/etc/fstab` entries for each logical volume.
- Mounts each logical volume.

Complete the following steps to allocate block devices to partitions, volume groups, and logical volumes.

1. Run the `devlist` command to locate unused block devices. The following example shows the `devlist` command output.

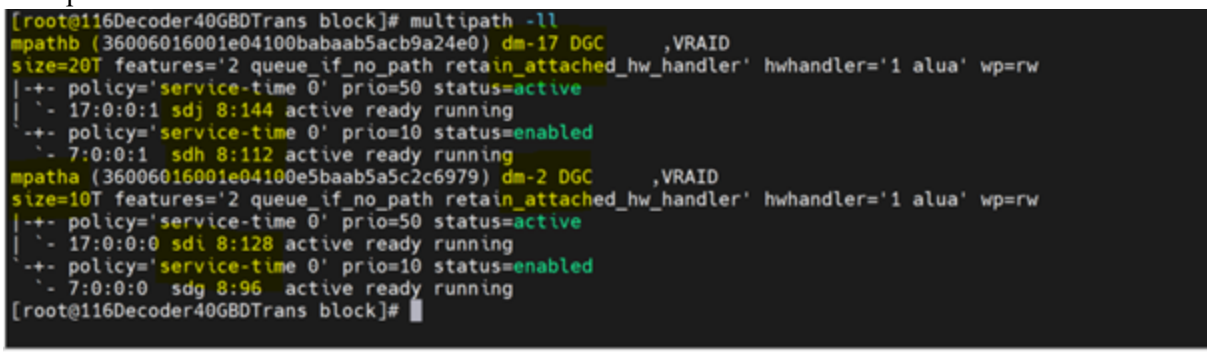
Output (or command manual help)

```
sda: vendor=DELL model="PERC H730P Mini" size="931 GB" used=1
sdb: vendor=DELL model="PERC H730P Mini" size="1.81 TB" used=1
sdc: vendor=DELL model="PERC H830 Adp" size="21.38 TB" used=1
sdd: vendor=DELL model="PERC H830 Adp" size="85.53 TB" used=1
```

Also,you must provide a name for the service that will be used with the storage, for example, **decoder** for the Network Decoder service, or **concentrator** for the Concentrator service. You have the option of providing the volume type. The default volume type has the same name as the service.

Note: Run the `devlist` command to see if the multipath user-friendly names are listed correctly.

2. Run the `multipath_II` command to make sure that SAN devices are attached. The following is an example when SAN devices are attached.



Note: Block devices should be configured with a user-friendly name such as mpatha, mpathb etc.

3. Execute the `partNew` command to allocate block devices to partitions, volume groups, and logical volumes.

   By default, the `partNew` command does not make changes. It displays the actions that will be taken if you commit the command string. To actually make the changes to the system, add the `commit=true` parameter to the command.

   For example, to assign devices `sdd` and `sde` to Decoder:

```
send /appliance partNew name=sdc service=decoder volume=decodersmall
commit=true
send /appliance partNew name=sdd service=decoder volume=decoder commit=true
```

Caution: For the **decoder** and **concentrator** services, you must create storage volumes in a specific order.
- The **decoder** has the **decodersmall** and **decoder** volumes. Create the **decodersmall** volume before the **decoder** volume because **decodersmall** contains the small filesystem mounted at `/var/netwitness/decoder`.
- The **concentrator** has the **concentrator** and **index** volumes. Create the **concentrator** volume before **index** volume or it will fail and you receive the following message.

```
Failed to process message partNew for /appliance
com.rsa.netwitness.carlos.transport.TransportException: Volumes for index
require mount point /var/netwitness/concentrator to be created and
mounted first.
```

4. Execute the `vgs` command to validate that the `partNew` command created the correct Logical Volumes.

   The output of this command:

   - Enumerates all the volume groups on this host.
   - Displays the physical volumes that the volume group consists of, and the logical volumes within the volume group.

5. Go to Task 4 - Allocate Volume Groups to NetWitness Services- `srvAlloc`.

Task 4 - Allocate Volume Groups to NetWitness Services - `srvAlloc`

The `srvAlloc` command configures services on a host to use storage in a volume group. You must provide the name of the service to configure and the volume group to assign to the service (the service you provide must be installed on the host). For information about NetWitness Platform service volumes, see "NetWitness Platform Service Volume Reference" in Storage Requirements.

Allocate services in the following order:

- For the Decoder, allocate `decodersmall` first then the `decoder`
- For a Concentrator, allocate `concentrator` first then `index`.

Note: By default, the `srvAlloc` command does not make changes. You must append the `commit=1` parameter to the command string to actually make the changes to the system and restart the specified service after making changes.

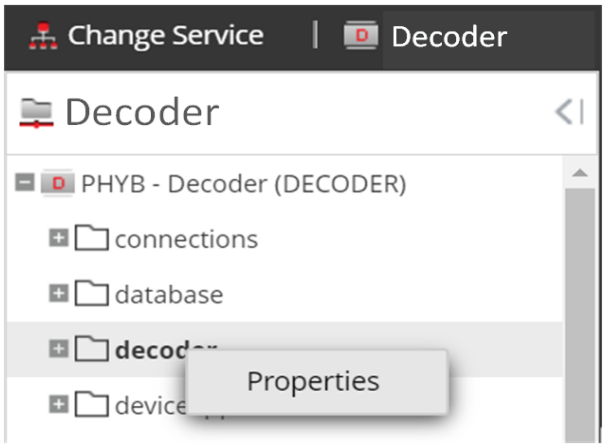
1. Execute the `srvList` command to see a list of services installed on this host.
   The `srvList` command communicates with the service through the SSL port. You install a Category on a host. A Category can be a single service, or multiple related services, located on the same host.
2. Execute the `srvAlloc` command to configure a service on a host to use storage in a volume group.
   For example:
   ```
   service=concentrator volume=concentrator commit=1
   service=concentrator volume=index commit=1
   ```
3. Go to Task 5 - Reconfigure Services and Databases to Detect and Appropriately Use New Storage.

Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture

You need to reconfigure the Decoder service and databases for 10G capture. Complete the following steps so that the Network Decoder service and its database detect and use new free space.

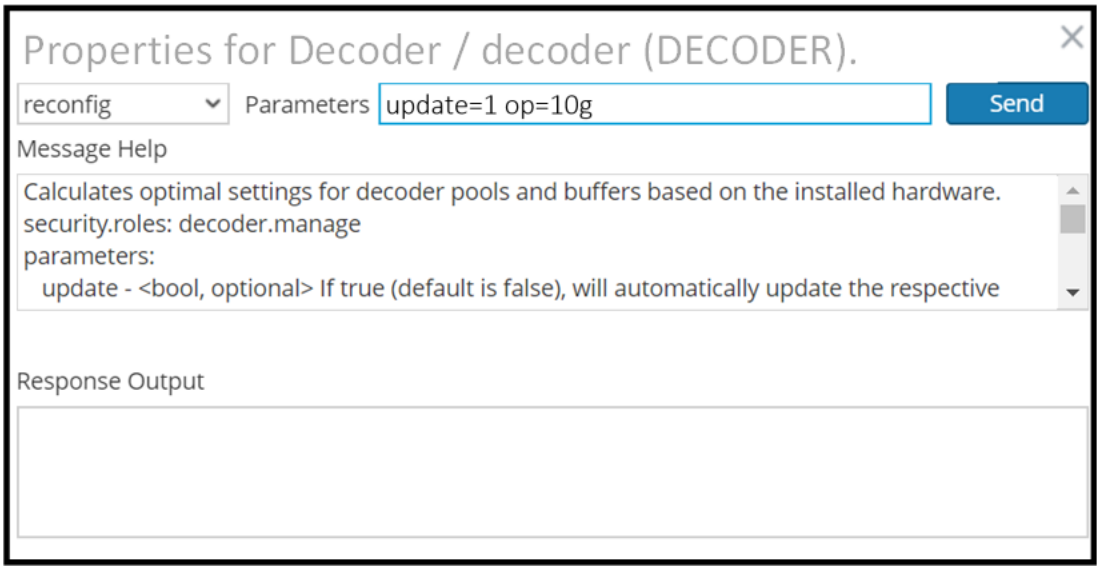
1. In the **NetWitness** menu, go to  (Admin) > **SERVICES**.
   The **SERVICES** view is displayed.
2. Select the **decoder**.
3. Under  (actions), select **View** > **Explore**.
   The **Explore** tree for the service is displayed.

4. Reconfigure space on the **decoder** service.
   a. Navigate to the **decoder**, right click, and click **Properties**.

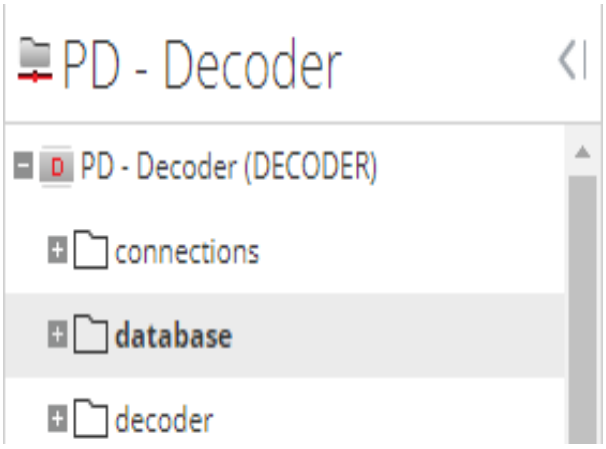


   The **Properties** dialog is displayed.

   b. Execute the `reconfig` command by selecting it from the drop-down list, specifiy `update=1 op=10g` in **Parameters**, and click **Send**.

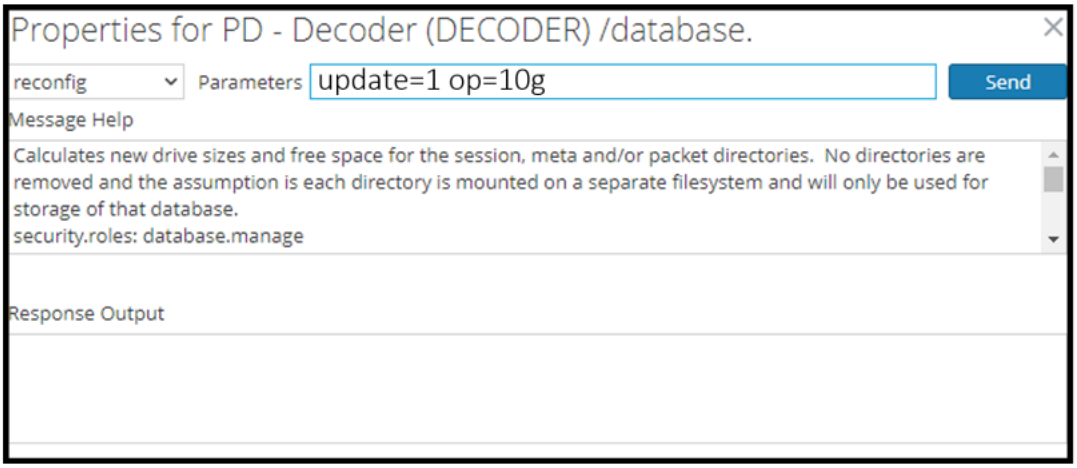


5. Reconfigure space on the database.
   a. Navigate to **database** in the service **Explore** tree, right click, and click **Properties**.



   The **Properties** dialog is displayed.

   b. Execute the `reconfig` command by selecting it from the drop-down list, specifiy `update=1 op=10g` in **Parameters**, and click **Send**.



Prepare Unity Storage

You must work with your Dell EMC Storage Engineer to allocate storage within your Unity environment for the NetWitness Platform and ensure the allocated storage meets all of the NetWitness Platform Storage Requirements. Specifically, make sure that:

- You have at least two LUNS created for Decoders (meta /session and packet volumes).
- You have at least two LUNS created for Concentrators (index and meta volumes).
- Ensure block devices can meet the minimum IOPS for expected ingestion rates.

You must add every NetWitness host that uses the Unity storage as a host within the Unity interface. After you create hosts and LUNs, you must assign the LUNs to the hosts. Assigning the LUNs to hosts makes the storage visible to the hosts so they can locate the storage through the host-based Dell EMC PowerPath software.

Note: A Dell EMC engineer will configure the following Unity Array.

You need to perform the following tasks to prepare Unity Storage.

Task 1 - Access Unisphere User Interface (UI)

Task 2 - Create Pools

Task 3 - Create LUNS

Task 4 - Register Hosts

Task 5 - Assign LUNS to Hosts

Task 6 - Install PowerPath

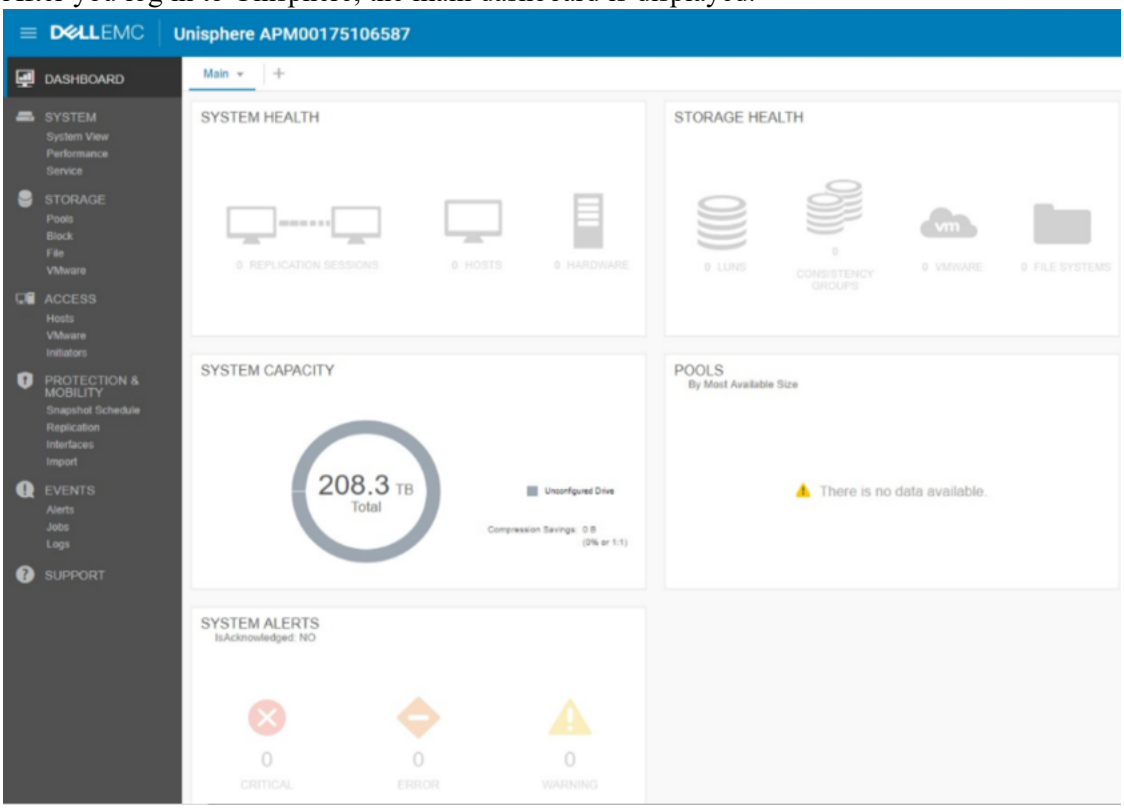
Task 1 - Access Unisphere User Interface (UI)

1. Connect your workstation on the same subnet as the UNITY.
2. Open a browser and go to **http://<unisphereIP>** to connect to the Unisphere UI.
3. Log in with the credentials provided by the DellEMC CE. The default credentials are **admin/Password123#**.

Note: Unisphere will ask you to change password the first time log in. It also asks you to install the license before you can configure array (DellEMC CE may do this for you. You must get the new admin password from them).



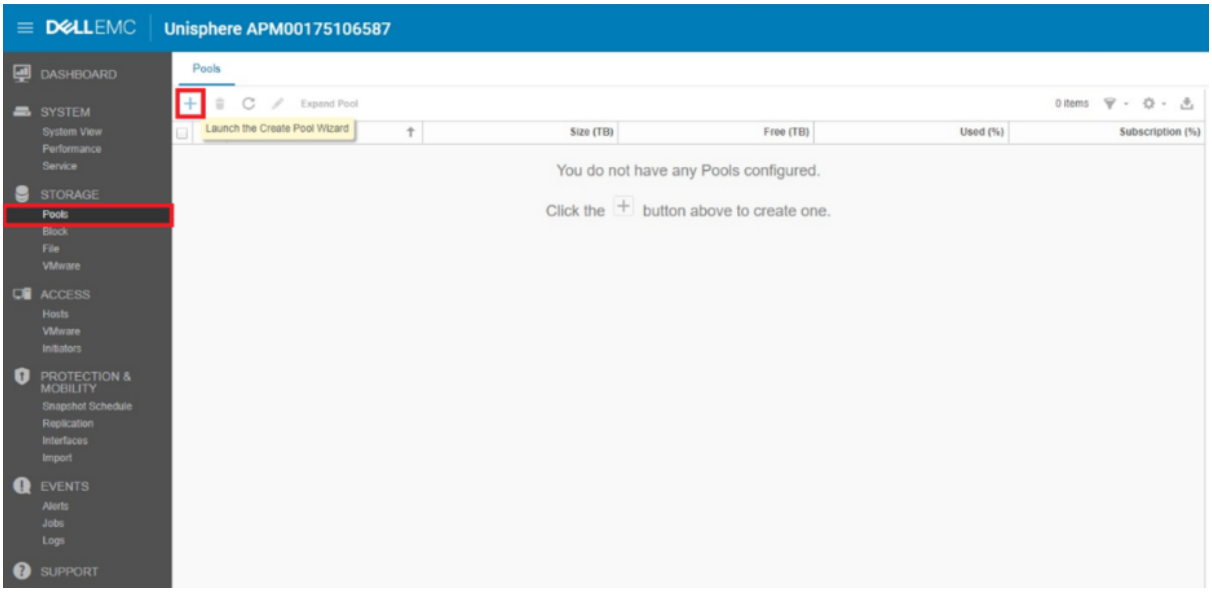
After you log in to Unisphere, the main dashboard is displayed.



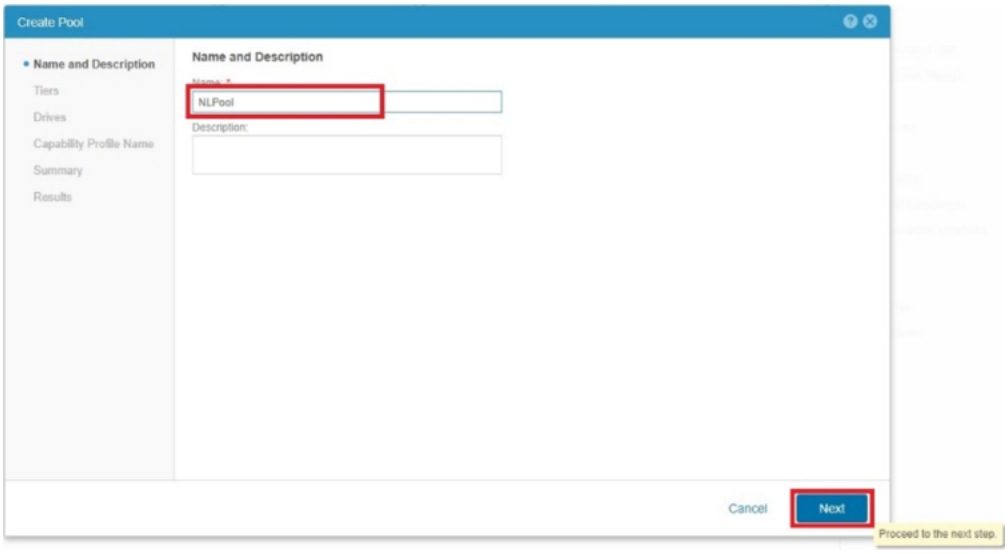
Task 2 - Create Pools

The NetWitness configuration consists of two different pools. One pool is dedicated to the NL-SAS drives and the other pool is dedicated to the SSDs.

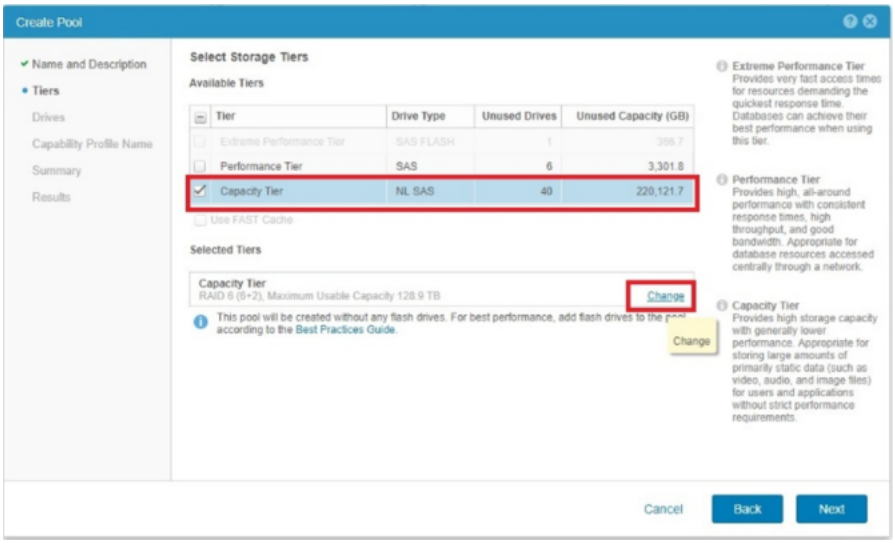
1. From **Storage Section**,click > **Pools** > (Add) to launch the Create Pool Wizard.



2. Enter in a name for the pool (for example, **NLPool**) and click **Next**. Optionally, you can also enter a description for the pool.

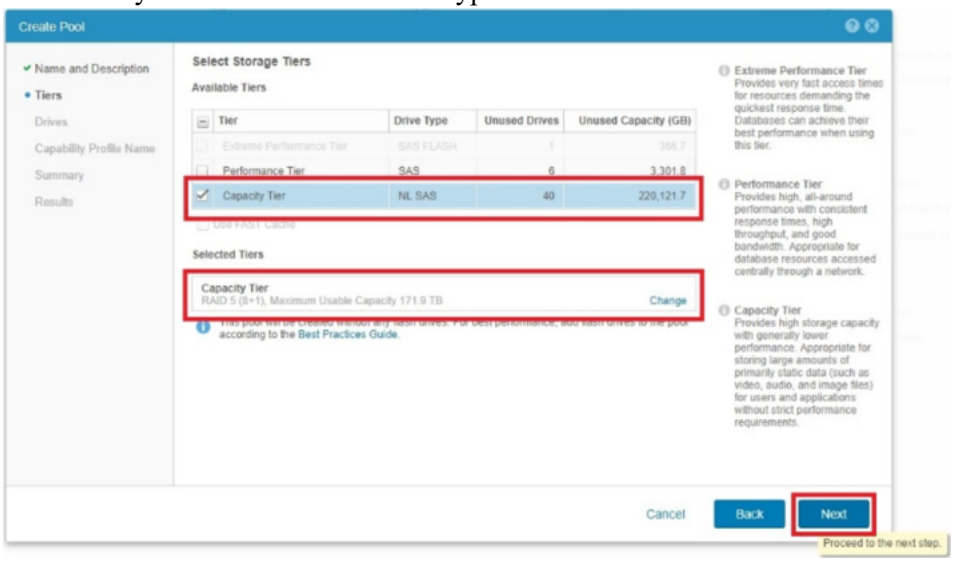


3. Select **Capacity Tier** under **Tier** for the tier type (drive type) and click **Change**.

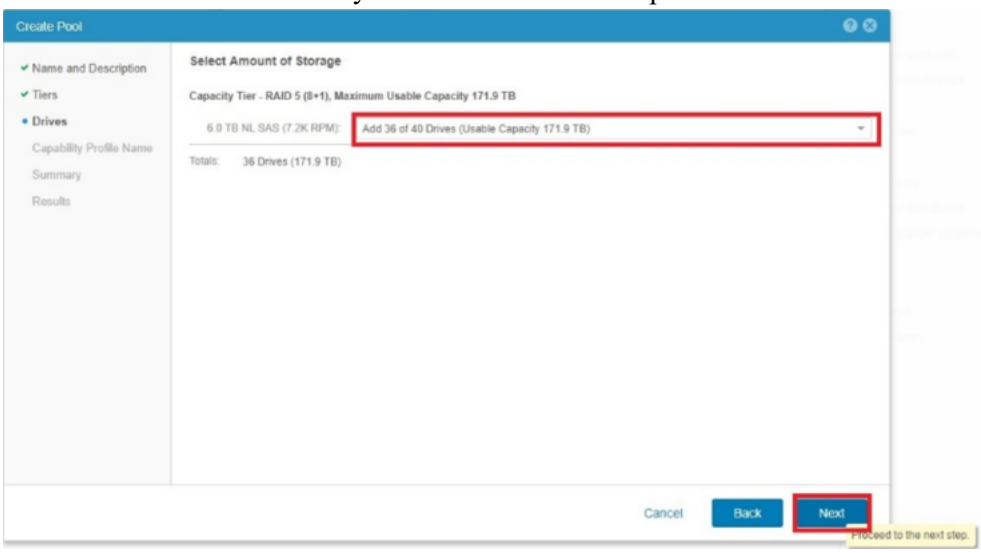


4. Choose the RAID type and from the drop down and select the RAID size.
   The RAID type and size are a customer preference. The only requirement is to make sure you have enough IOPS within the pool to accommodate the log or packet capture and queries. In the following example, a **RAID 5 (8+1)** configuration is selected, however some customers may prefer a **RAID 6 (10+2 or 12 +2)**.

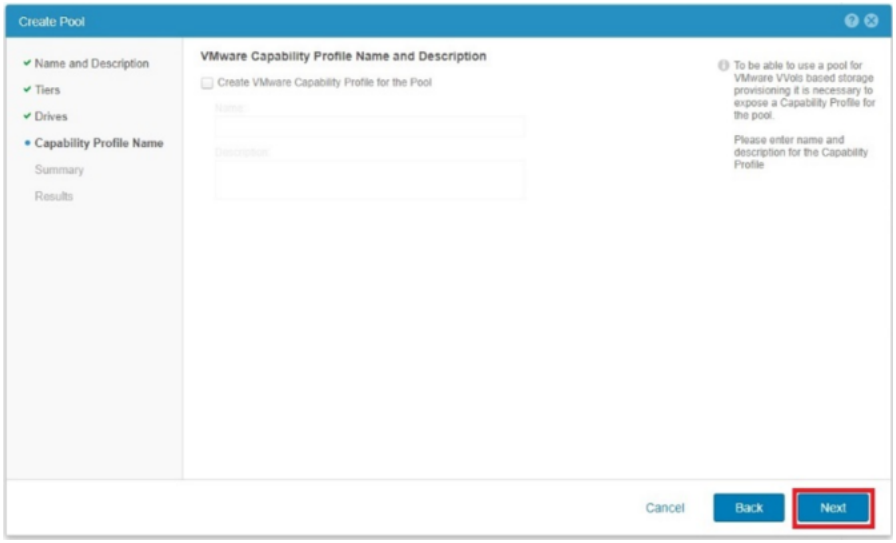
5. Make sure you have the correct Raid type and size selected.



6. Choose the number of drives you want to add into the pool and click **Next**.

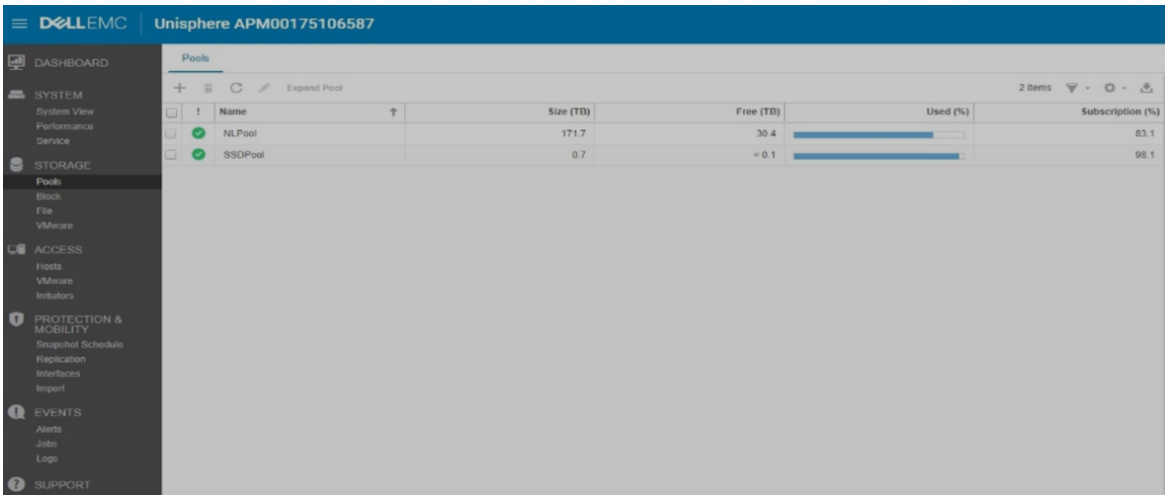


7. Skip the **VMware Capability** section and click **Next**.



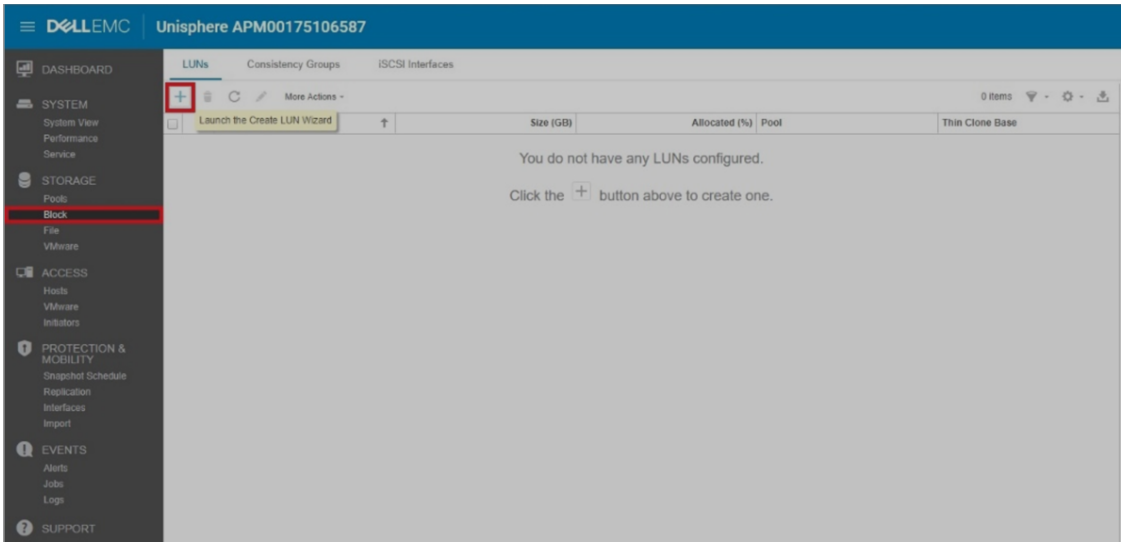
8. Make sure that everything is correct on the Summary tab, and click **Finish**.
9. Create another pool for the SSDs using steps 2 – 8.
   a. Enter in a name for theother pool (for example, **SDDPool**) and click **Next**. Optionally, you can also enter a description for the pool.
   b. Select **Extreme Performance Tier** under **Tier** for the tier type (drive type) and click **Change**.
   c. Choose the RAID type and from the drop down, select the RAID size, and click **OK**.

> **Note:** Raid 5 (4+1) RAID Configuration is different then Capacity Tier.



Task 3 - Create LUNS

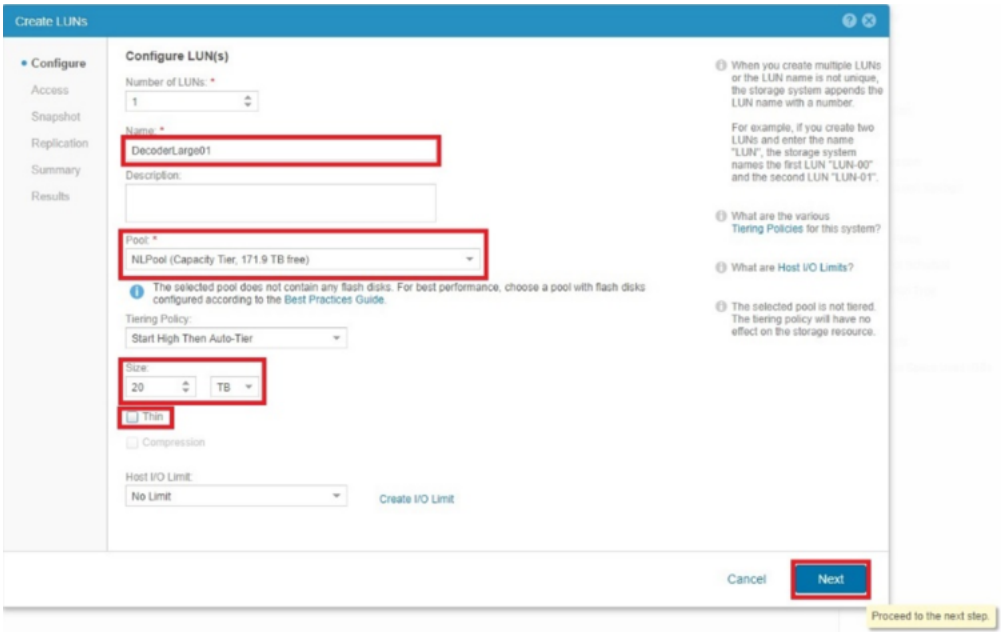
1. From **Storage** section, click **Block** > (Add) to launch the **Create LUN Wizard**.



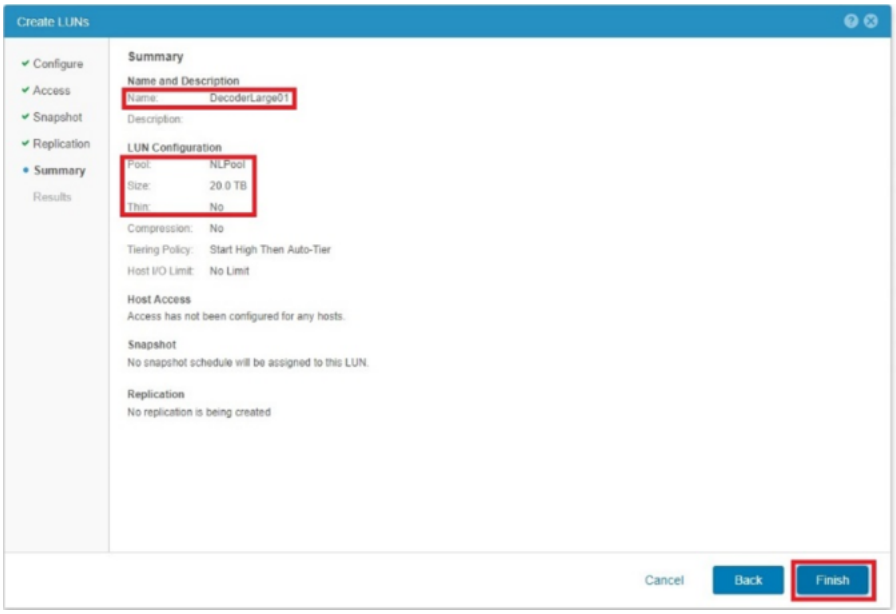
The table below list all of the possible LUNS you may need to create. The ConIndex is the only LUN you need to assign to the SSD Pool. Make sure that the LUN sizes do not exceed what is listed below.

DecoderLarge01	75 TB orLess	NL Pool	No
DecoderSmall01	20 TB or Less	NL Pool	No
Concentrator01	15 TB or Less	NL Pool	No
Archiver01	75 TB or Less	NL Pool	No
ConIndex01	3 TB or Less	**SSD Pool**	No

2. Enter the LUN Name from the list. Optionally, you can enter a description of LUN.
3. Select the appropriate pool from the list on the drop-down menu.
4. Deselect the **Thin** checkbox (These will be fully provisioned LUNs).
5. Select **Next** to proceed to the next menu.



6. Click **Next** until you get to the summary section.
7. Verify that the **Name**, **Pool**, **Size** and **Thin** selections are all correct.
8. Click **Finish** to complete LUN creation.

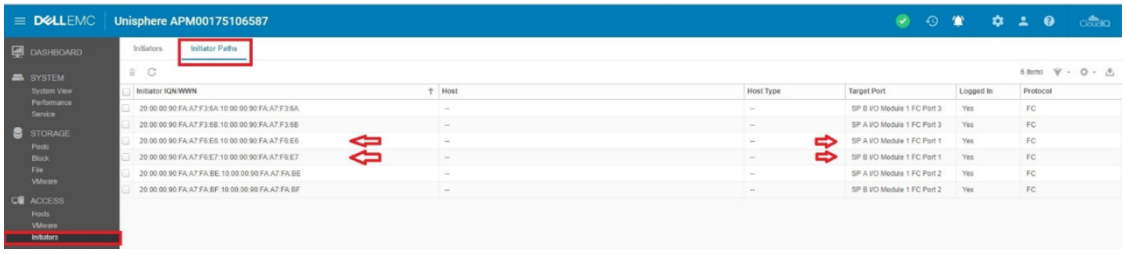


9. Repeat steps 2- 8 for the remaining LUN creations.

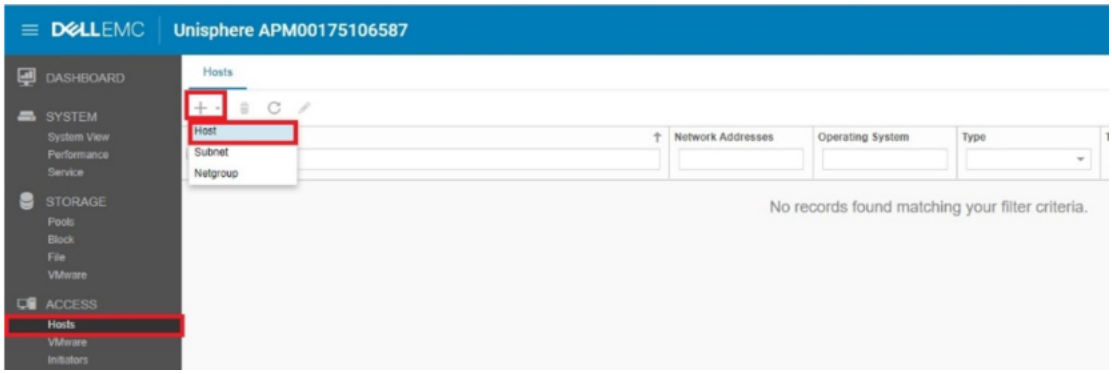
Task 4 - Register Hosts

Before proceeding, record the hostname and IP address of the Head Unit and make sure that the HBAs in the head unit are properly cabled to the UNITY.

1. From the Access section, click **Initiators**.
2. Under the **Initiator Paths** tab, make sure that the correct HBAs are selected that you will use to register the Head Unit.
   You should see two initiators per Head Unit. This represents the fiber connection from port 1 to SPA and port 1 to SPB. If you have multiple head units, the easiest method is to power each down and then power them up and register one by one.

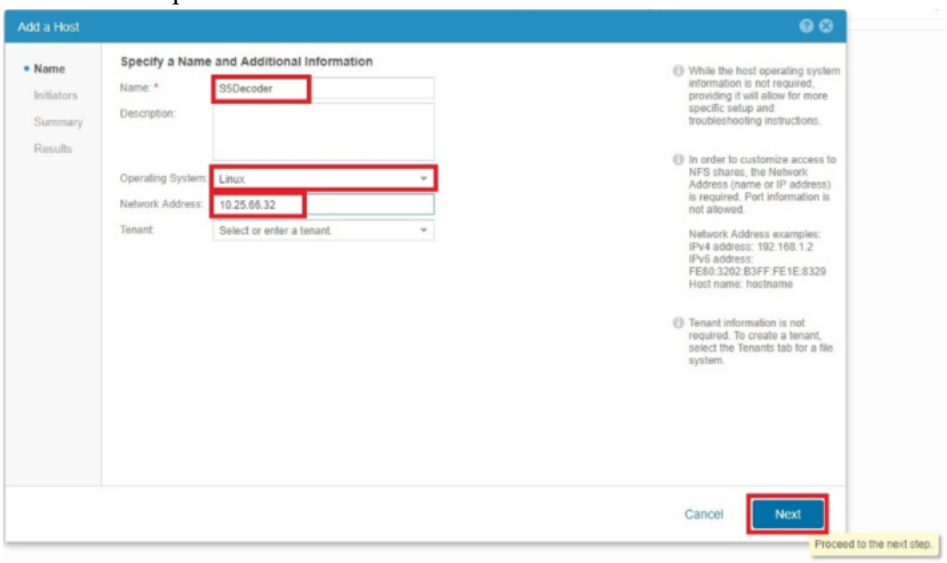


3. From the **Access** section, click **Hosts** > (Add) > **Host** to add a host configuration.

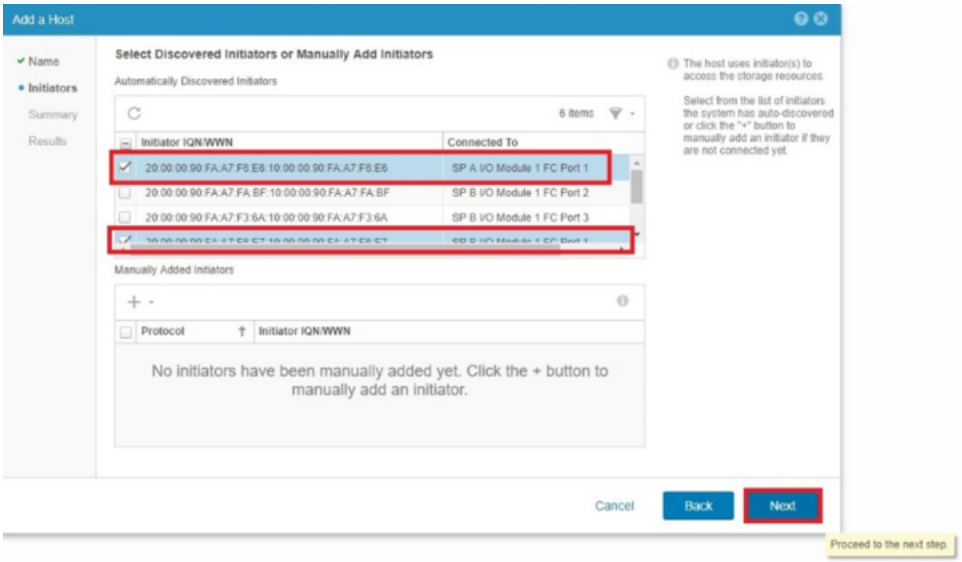


4. Enter the Hostname of the Head Unit.
5. Under **Operating System**, select **Linix** from the-drop down menu.
6. Enter the IP address of the Head Unit.

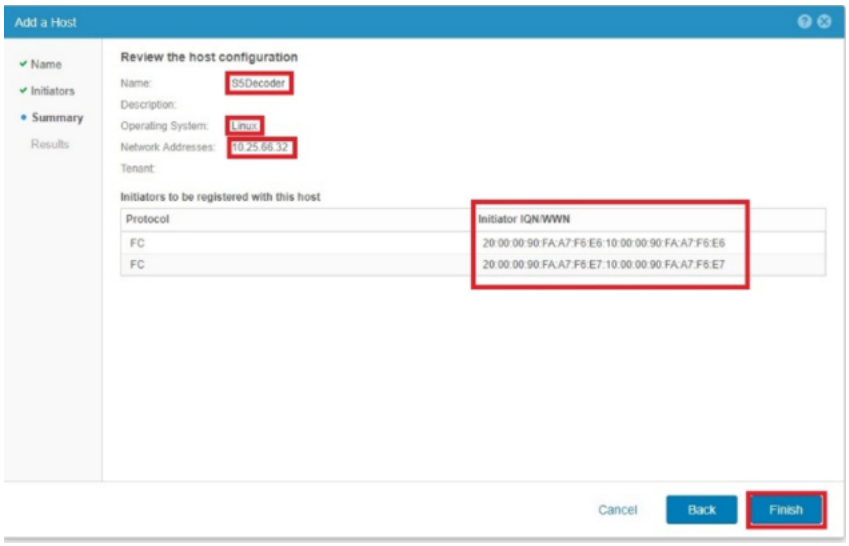
7. Click **Next** to proceed to the next section.



8. In the Initiators section, select the two initiators that correspond to the correct port associated with the Head Unit and click **Next** to proceed.



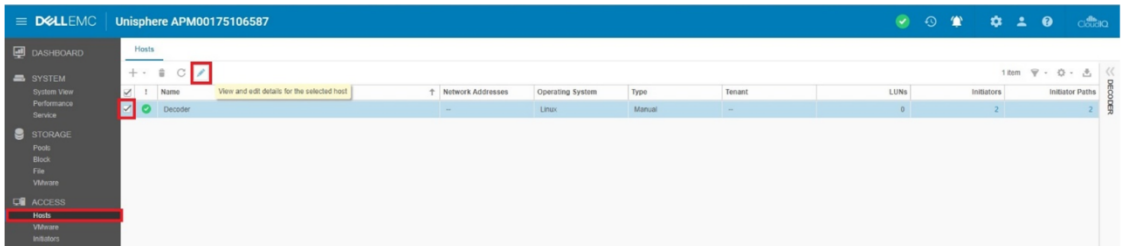
9. Make sure that the **Name**, **OS**, **IP** and **WWNs** are correct and click **Finish**.



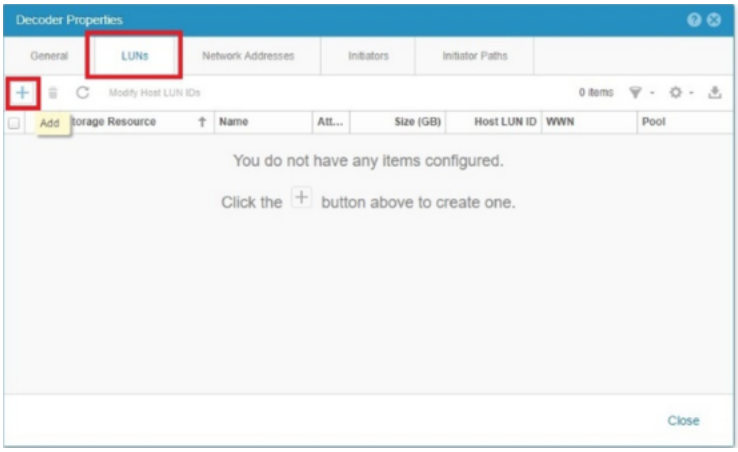
10. Repeat steps 2-9 for all Head Units.
11. In the Initiators section, select the two initiators that correspond to the correct port associated with the Head Unit. Then click “Next” to proceed.

Task 5 - Assign LUNS to Hosts

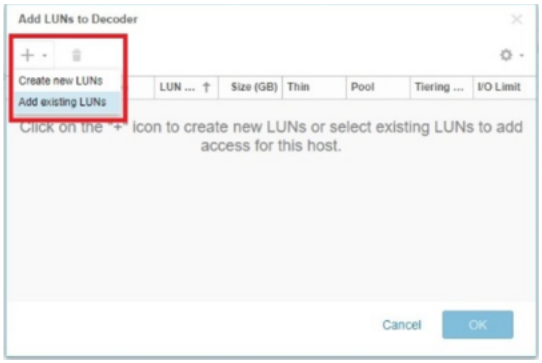
1. From the **Access** section, click **Hosts**, select the head unit (for example, **Decoder**) and click  (edit) to view and edit details for the selected host.



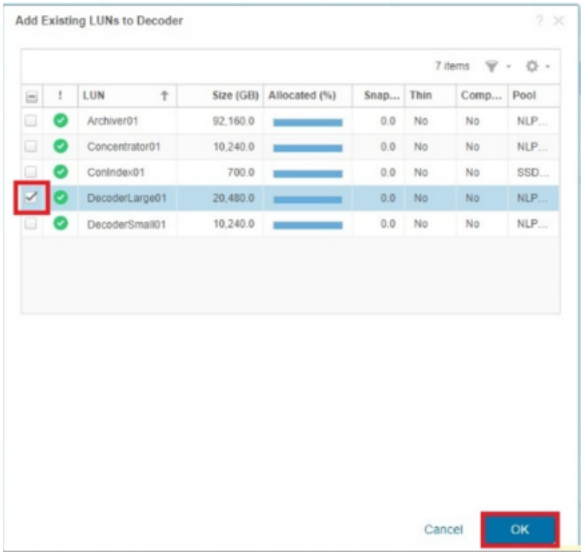
2. In the **Properties** section, select the **LUNS** tab and click + (Add icon).



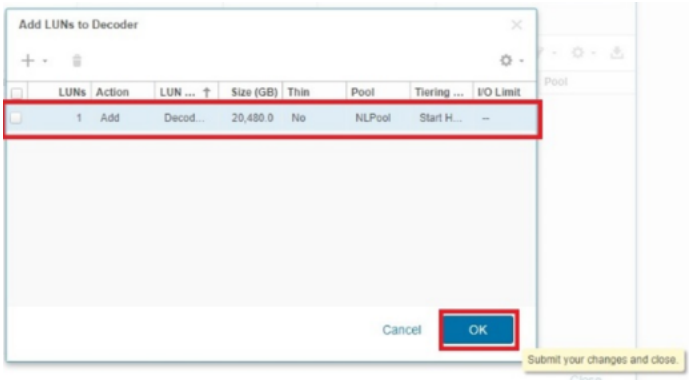
3. From the **Add LUNs to <Host>** popup, click + > **Add existing LUNs**.



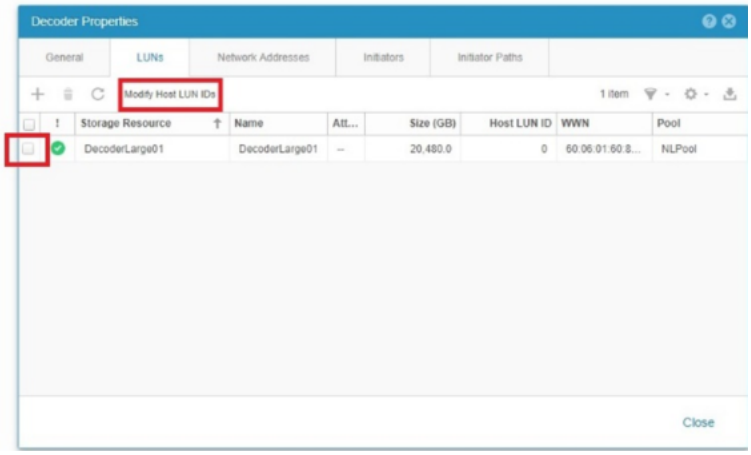
4. Select the LUN to add to the Head Unit and **OK**.



5. Make sure that the correct LUN was added to the host and click **OK**.



6. (OPTIONAL) If you need to modify the HLU (Host LUNN Unique ID):
   a. Select the LUN you want to change.
   b. Click **Modify Host LUN IDs**.



7. Click  (edit), change the HLU to the number you want, and click **OK**.

Task 6 - Install PowerPath

1. Make sure that the Emulex ports on the Decoder host are attached to the Unity.
2. Log in to `root` on the Decoder attached to the Unity with the `admin` credentials.
3. Install PowerPath and register the Dell EMC PowerPath licenses for Unity hardware.
   `yum install DellEMCPower.LINUX-6.4.0.00.00-95.RHEL7.x86_64.rpm`

> **Note:** When you purchase an NetWitness Provided Unity, PowerPath licenses are sent to you. You can download PowerPath at support.dell.com.

Note: It is possible that the RPM downloaded from Dell EMC is not signed with a cert that the NetWitness device has available, which can cause the installation to fail with the `package not signed` error. Run the yum install with the `--ngpgcheck` option to enable the software to install.

4. Make sure that all the PowerPath connections are correct.
   `powermt display dev=all`
   The following output is an example of valid PowerPath connections.

```
============================================================================
--------------- Host ---------------   - Stor -   -- I/O Path --   -- Stats ---
###  HW Path                I/O Paths    Interf.   Mode     State   Q-IOs Errors
============================================================================
  15 lpfc                    sde        SP A6     active   alive       0      0
  18 lpfc                    sdg        SP B6     active   alive       0      0

Pseudo name=emcpowerb
Unity ID=APM00174407815 [Host_62]
Logical device ID=600601609D90460069967 45A46B60AB6 [DecoderSmall01]
state=alive; policy=CLAROpt; queued-IOs=0
Owner: default=SP A, current=SP A       Array failover mode: 4
============================================================================
--------------- Host ---------------   - Stor -   -- I/O Path --   -- Stats ---
###  HW Path                I/O Paths    Interf.   Mode     State   Q-IOs Errors
============================================================================
  15 lpfc                    sdd        SP A6     active   alive       0      0
  18 lpfc                    sdf        SP B6     active   alive       0      0
```

5. Verify that the PowerPath license is installed using the `emcpreg` command.

```
[root@NWAPPLIANCE24932 ~]# emcpreg -list
Key BQPO-DB4M-VFC2-Q24R-ML9Z-EQTU
 Product: PowerPath
 Capabilities: Al
```

6. Add the following string to the `/etc/lvm/lvm.conf` file to filter the LVM (Logical Volume Manager) so that it ignores duplicate volumes.

```
filter = [ "a|^/dev/sda2$|", "a|^/dev/sdb1$|",
"a|^/dev/emcpower.*|", "r|.*/|" ]
```

7. Run the following commands in this order:
   a. `systemctl enable PowerPath.service`
   b. `systemctl start PowerPath.service`
8. Reboot the Decoder.
9. Complete the instructions in Configure Storage Using the REST API to complete storage configuration.

Migrate Data to Another Storage Type

This section provides two options for moving data from DACs to PowerVaults:

Migrate Data Using the Warm and Hot Tier Option

Move Data From DAC to PowerVault

Refer to the Hardware Setup Guides on NetWitness Community for detailed instructions for setting up NetWitness Platform host and storage hardware.

Migrate Data Using the Warm and Hot Tier Option

In this procedure, you configure a warm tier for the DAC's, so that they do not write any new data. The warm tier continues to be available for analyst operations. You configure the PowerVaults as a hot tier, where new data can be written and available for analysts. When the required data retention is available on the hot tier, the warm tier can be decommissioned.

To set up the warm and hot tiers, perform the following tasks:

- Stop the Service
- Set Up PowerVault
- Configure The Mount Points
- Set up Warm and Hot Tiers
- Decommision the DAC

Stop the Service

1. Log in to the NetWitness Platform user interface.
2. Go to  (Admin) > **SERVICES** and select the service (for example, Log Decoder).
3. Click  > **View** > **Config**, and under Log Decoder Configuration, clear the **Capture Autostart** checkbox, and then click **Apply**.
4. In the menu bar, click the down arrow next to **Config**, select **System**, and at the top of the panel, click **Stop Capture**.
5. From the command line interface in NwConsole, stop the service by running the following command:
   `systemctl stop nwlogdecoder`

Set Up PowerVault

1. Go to the REST API for the service by entering the IP address of the service, in this example, the Log Decoder. For example, `172.16.0.1:50106`.
2. Click the asterisk (*) next to the service. for example, **decoder (*)**.

3. Under **Properties for /decoder**, click the down arrow, select **RaidNew** and enter the following parameters, entering the name of the service for `scheme`. In this example, we use `logdecoder`.
   ```
   controller=1 enclosure=75 scheme=logdecoder commit=1
   ```
4. Click **Send**.
5. To configure the partitions, click the down arrow again, select **PartNew**, and enter the following parameters,
   ```
   name=sde service=logdecoder volume=logdecodersmall commit=1
   ```
6. Click **Send**.
7. With **PartNew** still selected, enter the following parameters:
   ```
   name=sdf service=logdecoder volume=logdecoder commit=1
   ```

   Note: To validate the partition definitions before committing them, you can enter these parameters without `commit=1`, and click **Send**. After you validate the parameters, add `#commit=1` and then click **Send** to commit the parameter settings.

Configure The Mount Points

1. On the NwConsole at the root level of the service (for example, the Log Decoder), run `df -h`.
   A list of mounted partitions is displayed.
2. Unmount all of the old storage points of the DAC and copy all the data to the Log Decoder. At the root level, run the `umount` command and the path name of each partition. You can concatonate the path names, for example:
   ```
   umount /var/netwitness/logdecoder/index
   /var/netwitness/logdecoder/sessiondb /var/netwitness/logdecoder/metadb
   /var/netwitness/logdecoder/packetdb /var/netwitness/logdecoder/index0
   /var/netwitness/logdecoder/sessiondb0 /var/netwitness/logdecoder/metadb0
   /var/netwitness/logdecoder/packetdb0
   ```
3. Temporarily mount the petitions in the `decoroot` folder in the `/mnt` directory in order to access the files. For example:
   ```
   mount /dev/mapper/logdecodersmall-decoroot /mnt/decoroot/
   ```
4. Copy the contents of decoroot from /mnt to /var/netwitness/logdecoder, answering Y (yes) to the prompts:
   ```
   cp -R statdb /var/netwitness/logdecoder/
   ```
5. Unmount `/mnt/decoroot`.
   ```
   umount /mnt/decoroot
   ```
6. Comment out `decoroot` from `/etc/fstab`, as this was on the DAC and the DAC will be decomissioned.
   ```
   #/dev/logdecodersmall/decoroot
   /var/netwitness/logdecoder/xfs/noatime,nosuid 1 2
   ```
7. Mount all the remaining file systems.
   ```
   mount -a
   ```
8. Start the nwlogdecoder service (with capture still disabled).
   ```
   systemctl start nwlogdecoder
   ```

Set up Warm and Hot Tiers

Caution: Before you set up warm and hot tiers, be sure that you know the right warm and hot tier entries for each collection so that you can set them up accurately.

1. Go to (Admin) > **SERVICES** and select the service (for example, Log Decoder).
2. For the Log Decoder service, click > **View** > **Explore**, and go to **database** > **config**.
   a. Copy the contents of `meta.dir` and paste them to `meta.dir.warm` as shown in the following example:



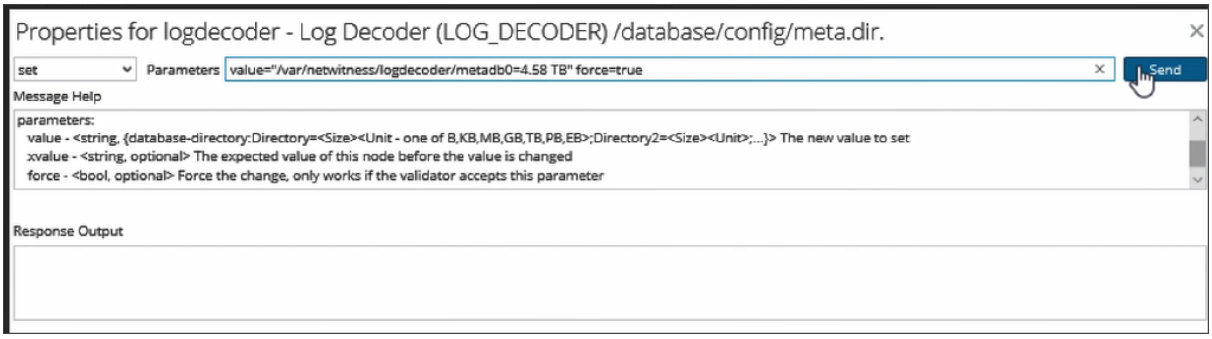
   b. In the same way, copy the packet database in `packet.dir` to `packet.dir.warm`.
   c. Copy the session database in `session.dir` to `session.dir.warm`.
3. Go to **index** > **config** and copy `index.dir` to `index.dir.warm`.

Note that the new volumes end in 0, so PowerVault will write to the directories ending in 0, for example:

```
[root@logdecoder ~]# df -h
Filesystem                              Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root         30G  3.3G   27G  11% /
devtmpfs                                 63G     0   63G   0% /dev
tmpfs                                    63G   12K   63G   1% /dev/shm
tmpfs                                    63G   34M   63G   1% /run
tmpfs                                    63G     0   63G   0% /sys/fs/cgroup
/dev/sda1                              1019M   96M  924M  10% /boot
/dev/mapper/netwitness_vg00-nwhome       3.3T  1.2G  3.3T   1% /var/netwitness
/dev/mapper/netwitness_vg00-usrhome       10G   33M   10G   1% /home
/dev/mapper/netwitness_vg00-varlog        10G  1.5G  8.6G  15% /var/log
tmpfs                                     13G     0   13G   0% /run/user/0
/dev/mapper/logdecodersmall-index         30G   54M   30G   1% /var/netwitness/logdecoder/index
/dev/mapper/logdecodersmall-sessiondb    600G  733M  599G   1% /var/netwitness/logdecoder/sessiondb
/dev/mapper/logdecodersmall-metadb       4.9T   11G  4.9T   1% /var/netwitness/logdecoder/metadb
/dev/mapper/logdecoder-packetdb           31T   12G   31T   1% /var/netwitness/logdecoder/packetdb
/dev/mapper/logdecodersmall0-index        30G   33M   30G   1% /var/netwitness/logdecoder/index0
/dev/mapper/logdecodersmall0-sessiondb   600G   34M  600G   1% /var/netwitness/logdecoder/sessiondb0
/dev/mapper/logdecodersmall0-metadb       21T   34M   21T   1% /var/netwitness/logdecoder/metadb0
/dev/mapper/logdecoder0-packetdb          86T   35M   86T   1% /var/netwitness/logdecoder/packetdb0
[root@logdecoder ~]#
```

Update the Decoder configuration with the path to the PowerVault mount by adding a 0 to the path.

1. In the /database/config column, right-click **meta.dir** and click **Properties**.
2. In **Properties for logdecoder**, select **set**, and in **Parameters**, enter `value="var/netwitness/logdecoder/metadb0=4.58 TB' and add force=true`, as shown in this example, and then click **Send**.



3. Repeat step 2 for **session.dir**, **packet.dir**, and **index.dir**. Do not be concerned if the size is the same as the DAC in "=xx GB". This will be updated in the next step.

   Note: We are only putting the PowerVault paths into the *.dir values.

4. Update the sizes for the live PowerVault volumes.
   a. In the Log Decoder Explore view, in the left panel, right-click **database** and click **Properties**.
   b. Select **reconfig** and in **Parameters**, enter `update=1` and click **Send**.
   c. Repeat steps a and b for **index**.
5. Restart the service.
   `systemctl restart nwlogdecoder`

6. Go to  (Admin) > **SERVICES**, select the Log Decoder service, and click  > **View** > **System**.
7. Click **Start Capture**.
8. Go to the **Config** view, select **Capture Autostart**, and click **Apply**.
9. Reboot the host.

Decommision the DAC

When the DAC data has aged, you should go back into the Explore view and remove all of the *.dir.warm configurations for session, meta, packet and index. You can determine when the DAC data has aged by going to the Log Decoder  > **View** Explore view. Since we have a hot and warm tier, there are two sets of configuration stats that you need to be aware of. For example, for a packet Decoder, when you look at the packet oldest time in `packet.oldest.file.time`, look at the packet.oldest.file.time.hot value and if you see that your DAC had storage up until 30 days ago you can take your DAC offline and decommission it.

These are the basic steps for decommissioning a DAC. NetWitness recommends that you work with your Customer Support representative when you decommission your DACs.

1. Go to  (Admin) > **SERVICES** and select the service (for example, Log Decoder).
2. Click  > **View** > **Config**, and under Log Decoder Configuration, clear the **Capture Autostart** checkbox, and then click **Apply**.
3. In the menu bar, click the down arrow next to **Config**, select **System**, and at the top of the panel, click **Stop Capture**.
4. From the commandline interface in NwConsole, stop the service by running the following command:
   `systemctl stop nwlogdecoder`
5. Unmount the warm tier. At the root level, run the `umount` command and the path name of each partition. You can concatonate the path names, for example:
   ```
   umount /var/netwitness/logdecoder/index
   /var/netwitness/logdecoder/sessiondb /var/netwitness/logdecoder/metadb
   /var/netwitness/logdecoder/packetdb /var/netwitness/logdecoder/index0
   /var/netwitness/logdecoder/sessiondb0 /var/netwitness/logdecoder/metadb0
   /var/netwitness/logdecoder/packetdb0
   ```
6. Comment out all the old DAC dbs from `/etc/fstab`, so that only the PowerVault dbs remain.
7. Start the service.
   `systemctl start nwlogdecoder`
8. In the user interface, go to  (Admin) > **SERVICES** and select the Log Decoder service.
9. Click  > **View** > **Explore** and remove the warm tier configurations:
   a. In **database** > **config**, delete the content for `meta.dir.warm`, `packet.dir.warm`, `session.dir.warm`.

b. In **index** > **config**, delete the content for `index.dir.warm`.

c. Go to the **Config** view, select **Capture Autostart**, and click **Apply**

d. Go to the **System** view and click **Start Capture**.

10. Restart the service.
    `systemctl restart nwlogdecoder`

The DAC is now unmounted, and is no longer configured in the Decoder for warm storage and is ready to be wiped clean.

1. Remove the logical volume. Run `lvscan` to get a list of the logical volumes.
2. Run `lvremove` on the old logical volumes, for example:
   `/dev/logdecodersmall/decoroot /devlvremove /dev/logdecodersmall/index /dev/logdecodersmall/sessiondb /dev/logdecodersmall/metadb /dev/logdecodersmall/packetdb`
3. Remove the volume groups. Run `vgscan` to get a list of volume groups.
4. Run `vgremove` on the old volume groups (be careful not to remove any volume groups that end in 0, as they are PowerVault).
5. Run `pvscan` to view block devices that are freed up.
6. When the DAC has been successfully removed, reboot the host.

Move Data From DAC to PowerVault

The following procedure describes how to move data from DAC to PowerVault. Before you move data from 2 DACs to 2 PowerVaults, a table, similar to the following table, is displayed if you run the `pvs` (Physical Volume Size) command from the Decoder Linux console (or SSH to the Decoder) with 2 DACs attached and configured to the Decoder. The column headings are Physical Volume (PV), Volume Group(VG), Linux Format (Fmt), Linux Attribute (Attr), Physical Volume Size (PSize), and Physical Volume Free Space(PFree).

PV	VG	Fmt	Attr	PSize	PFree
/dev/sda2	netwitness_vg00	lvm2	a--	<930.00g	0
/dev/sdb1	netwitness_vg00	lvm2	a--	<1.82t	0
/dev/sdc	decodersmall	lvm2	a--	<5.46t	0
/dev/sdd	decoder	lvm2	a--	<27.29t	0
/dev/sde	decodersmall0	lvm2	a--	<5.46t	0
/dev/sdf	decoder0	lvm2	a--	<27.29t	0

Complete the following steps to move data from a DAC to a PowerVault.

1. Attach two PowerVaults to a separate PERC controller on the Decoder.
2. Create the devices.
   a. Open a Browser and specify the ip-address of the Network Decoder and port **50106** to access the REST tool.
   b. Log in with the `admin` account credentials.



   c. Click on the `(*)` next to **appliance** to access the REST command set.
   d. Run `raidList` to display the Controller/Enclosure combination with the new PowerVault enclosures.
   In the following example, the output shows **dev/sdg** and **/dev/sdh** on **Controller 2, Enclosure 246**.

```
Controller 2, Enclosure 246
    Vendor:  DELL
    Model:   MD1400
    In Use:  true
    Drives:  10.691 TB x 12
    Devices: sdg
             sdh
```

   e. Under **Properties for /appliance**, select `raidNew`, specify `controller=<PowerVault-controller-id> enclosure=<PowerVault-enclosure-id> scheme=decoder preferSecure=false`, and click **Send**.

> **Note:** You specify **preferSecure=false** if the PowerVault drives are not SED drives. If PowerVault drives are SED drives and you do not want to encrypt them you specify **preferSecure=false**. You must specify **preferSecure=true** if PowerVault drives are SED drives and you want to encrypt them.

3. Go to the Decoder Linux console or SSH to the Decoder and run the following commands.

```
parted -s /dev/sdg mklabel gpt
parted -s -a optimal /dev/sdg mkpart LVM 0% 100%
pvcreate -f /dev/sdg
parted -s /dev/sdh mklabel gpt
parted -s -a optimal /dev/sdh mkpart LVM 0% 100%
```

```
pvcreate -f /dev/sdh
```

If the volume is created successfully, the following message is displayed.

```
Physical volume "/dev/sdg" successfully created
```

Note: Repeat this step for every block device. The block device names may be different depending on how many enclosures per perc card slot.

4. Run the following command strings to extend the DAC volume group (**decoder**, **decodersmall**) to the Powervault Physical volume.

```
vgextend decoder /dev/sdg
vgextend decodersmall /dev/sdh
```

5. Run the following command strings to move the data from the DAC to the PowerVault. In this following command string, the DAC is **/dev/sdc** and the PowerVault is **/dev/sdg**.

```
pvmove /dev/sdc /dev/sdg
pvmove /dev/sdd /dev/sdh
```

Note: 1.) The `pvmove` command synchronizes data across volumes so that NetWitness can continue ingesting or aggregating data while the migration is executing. You can run the `pvmove` command multiple times if it fails. 2.) Depending on the amount of data on the drives, the move can take a long time complete depending on the amount of data. For example, in a test, it took four hours to move one TB of data.

6. After the move is complete, run the following commands to reduce and remove the DAC drive.

```
vgreduce decoder /dev/sdc
pvremove /dev/sdc
vgreduce decodersmall /dev/sdd
pvremove /dev/sdd
```

7. Detach the physical connections from the DACs to the host.

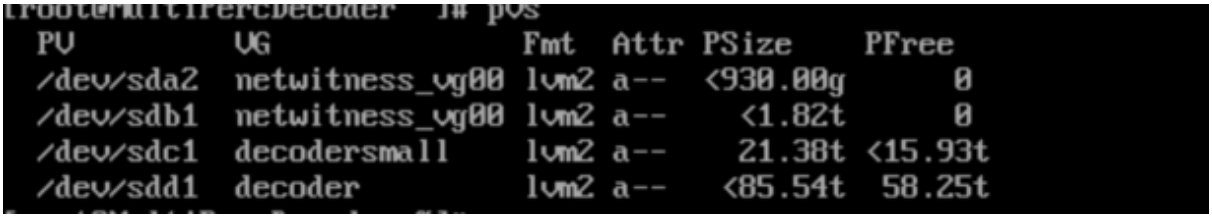
8. Verify that the Physical volumes are moved from the DACs to the PowerVaults.

   a. Reboot the host.

   ```
   reboot
   ```

   b. Verify that the **/etc/fstab** file is correct.

   c. Run the `pvs` command and make sure that the **PSize** and **PFree** values are correct on the PowerVault.



Data on PowerVault After Move from DAC

After you move data from 2 DACs to 2 PowerVaults, a table, similar to the following table, is displayed if you run the `pvs` (Physical Volume Size) command from the Decoder Linux console (or SSH to the Decoder) with 2 PowerVaults attached and configured to the Decoder. The column headings are Physical Volume (PV), Volume Group(VG), Linux Format (Fmt), Linux Attribute (Attr), Physical Volume Size (PSize), and Physical Volume Free Space(PFree).

PV	VG	Fmt	Attr	PSize	PFree
/dev/sda2	netwitness_vg00	lvm2	a--	<930.00g	0
/dev/sdb1	netwitness_vg00	lvm2	a--	<1.82t	0
/dev/sdc1	decodersmall	lvm2	a--	21.38t	<15.93t
/dev/sdd1	decoder	lvm2	a--	<85.54t	58.25t

Appendix A. How NetWitness Platform Hosts Store Data

In most deployments, NetWitness Platform Decoders, Log Decoders, Concentrators, Archivers, and Hybrid hosts require external storage to house their data. Each host uses the external storage in different ways and with different expectations on throughput and performance of the external storage. Some hosts have a higher occurrence of sequential writes and some hosts have a higher occurrence of random reads and writes.

Decoder Hosts

Log Decoders and Network Decoders capture data and parse meta. The difference between these two hosts is in the type of data they capture:

- Log Decoder captures logs.
- Network Decoder captures packets.

Both Log Decoders and Network Decoders parse out meta data from the raw captured traffic. The meta data is then aggregated to a Concentrator for indexing. The host requires storage to house the raw payload data (raw packets or raw logs) and a cache for the meta extracted during data capture for Concentrator aggregation.

Your retention requirements is a key factor in determining the amount of storage you need for the raw packets or raw logs. In most deployments, you add storage over time based on increased retention requirements and increased capture rates. The storage for the raw data must support a high amount of sequential writes with random reads. Especially in the case of higher speed Network Decoder environments, it is recommended to have a minimum of two partitions exposed to the host to support the throttling between partitions for reads and writes.

The meta cache on a Decoder is generally fixed in size but you can expand it to support additional cache the possible loss of connectivity between the Decoder and a corresponding Concentrator. The meta cache must support a random IOPS rate for sustained writes from the Decoder of meta extracted and the corresponding reads from the Concentrator as meta is aggregated to a Concentrator.

Concentrator Host

A Concentrator aggregates and indexes the meta data from a Decoder. Both the meta and index storage needs are scaled based on your NetWitness Platform deployment retention requirements. Similar to raw data stored on the Decoders, you may need to increase the storage for both meta data and index data over time to meet your retention requirements.

The meta storage houses all meta data extracted from either a Network Decoder or Log Decoder. Although the ratio of how much meta is extracted may change, the expectations for performance against meta storage is the same for both packet capture and log capture environments. The meta storage must support a sustained amount of sequential writes with random reads of meta data.

The index storage houses the live index generated from the meta data aggregated to a Concentrator. The size of the index is directly related to the size of the meta store. In addition to supporting IOPS for sustained writes, the index also needs to support a much higher rate IOPS for reads than meta based on interactive queries run through analyst interaction and reports and alerts.

Archiver Host

The Archiver host requires a single partition for both meta and raw log storage. The storage pool deals primarily with sequential writes for long term data written from a Log Decoder or Network Decoder and random reads for reports and analysis.

Hybrid Hosts

A Hybrid hosts two or more services on a single host. For example:

- A Network Hybrid hosts both the Decoder and Concentrator services handling packets exclusively. It captures packet data and indexes this data to the Concentrator service. Expectations for storage performance match what is outlined for a dedicated Network Decoder host and dedicated Concentrator host.
- A Log Hybrid hosts both the Log Decoder and Concentrator services handling logs exclusively. It captures log data and indexes the data to a Concentrator service. Expectations for performance match what is outlined for a dedicated Log Decoder and dedicated Concentrator.
- An Endpoint Log Hybrid hosts the Endpoint Server, Log Decoder, Concentrator, Log Collector, and Endpoint Broker services. It collects and manages endpoint (host) data from Windows, Mac, and Linux hosts, collects log files and Windows logs from Windows hosts, and generates metadata to correlate endpoint data with sessions from other events sources, such as logs and packets.

Options for SAN Configurations

If you want to use a Storage Area Network (SAN) , use the same basic drive groups and partition organization that you use for the other NetWitness storage devices. Depending on the SAN configuration and overhead, SAN configurations may require more enclosures and drives to operate with the same performance as on PowerVault or DAC. When deciding whether to use SAN, PowerVault, or DAC, any additional overhead on the SAN will be important to determine the minimum storage required.

Performance Recommendations

NetWitness recommends that Packet and Log Decoders receive two LUNs or Block Devices, one for Packet data, the other for all other databases. This allows you to segregate the high-bandwidth Packet Database from the other databases so they do not compete for I/O bandwidth with other activity.

Concentrators require a separate SSD-based index volume for best performance. You must house this index volume on a different RAID group than the Concentrator Meta database volume, which you can stored on NL-SAS. Archivers can use a single large NL-SAS storage volume per appliance.

Enable Security on SED Capable Drive groups on Host with a mix of SED and NON SED Drives

The encryptSedVd.py may fail to identify the SED Capable Virtual Drives when there is mix of both SED and NON-SED drives on the appliance. The below steps are applicable when both SED and NON-SED capable virtual drives exist on the host.

1. SSH to the appliance and enable security on the PERC H740 (mini) Adaptor. The controller number for this adaptor is **0**. The PERC H840 Adaptor is shown as **1**.
   To list all the controllers on the appliance:
   `/opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'`
   The first column (**Ctl**) lists out the controller index on the appliance. In this case, the controller '**0**' corresponds to '**PERC H740 Mini**' and controller '**1**' corresponds to '**PERC H840 Adaptor**' . The columns '**DGs**' and '**VDs**' displays the virtual drives and drive groups on the controller.

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'
Ctl Model          Ports PDs DGs DNOpt VDs VNOpt BBU sPR DS EHS ASOs Hlth
-------------------------------------------------------------------------
  0 PERCH740PMini      8  10   3     0   3     0 Opt On  -  N      0 Opt
  1 PERCH840Adapter    8  12   1     0   1     0 Opt On  -  N      0 Opt
[root@116Decoder perccli]#
```

2. To enable the security on the 'PERC H740 (mini) Adaptor', for example, Controller '**0**', execute the following command:
   `/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='<SOME_STRING_VALUE>'!' keyid='< SOME_STRING_VALUE >'`
   Example:
   `/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!' keyid=1`
   `'Netwitness1' is the securityKey and '1' is ID. Preserve both the Key and keyID securely.`

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!' keyid='netwitness'
Controller = 0
Status = Success
Description = None


Controller Properties :
=====================

--------------------
Ctrl Method  Result
--------------------
   0 set Key Success
--------------------
```

3. Identify the correct Drive group (DG) / Virtual Drive (VD) corresponding to the SED Capable drives that you are trying to enable security.
   `/opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'`
   Refer to first two and last column to identify the correct Drive Group (DG) / Virtual Drive (VD) correspond to the 6 SED enabled drives. On Series 6 appliances, there is only one DG/VD with **RAID6**. '**NAME**' column can be used to identify the VD/DG. In this case, the DG/VD is '**2**'. Using a combination of '**Type**', '**Name**' and '**Size**' columns (these are defined by the user when the VDs are created above).

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'
DG/VD TYPE  State Access Consist Cache Cac sCC     Size Name
-------------------------------------------------------------------------
0/0   RAID1 Optl  RW     Yes     RWBD  -   OFF 931.0 GB
1/1   RAID1 Optl  RW     Yes     RWBD  -   OFF 1.818 TB
2/2   RAID6 Optl  RW     Yes     RWBD  -   OFF 8.730 TB Virtual Disk 2
-------------------------------------------------------------------------
[root@116Decoder perccli]#
```

4. To turn on Security on the disk group (created out of the 6 SED Capable drives), execute the below command:

```
/opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
```

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
Controller = 0
Status = Success
Description = Success
```

5. Get the Enclosure ID (**EID**) using on the controller '**0**'. In this case, it is '64'

```
/opt/MegaRAID/perccli/perccli64 /c0 /eall show
```

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /eall show
Controller = 0
Status = Success
Description = None


Properties :
==========

-------------------------------------------------------------------------
EID State Slots PD PS Fans TSs Alms SIM Port#      ProdID    VendorSpecific
-------------------------------------------------------------------------
 64 OK       10 10  0    0   0    0   1 00 & 00 x8 BP14G+EXP  +
-------------------------------------------------------------------------

EID-Enclosure Device ID |PD-Physical drive count |PS-Power Supply count|
TSs-Temperature sensor count |Alms-Alarm count |SIM-SIM Count


[root@116Decoder perccli]#
```

6. To confirm that the drives / Drive Groups (DG) are **SED Enabled** and **Secured**, run the below command and verify the **SED Capable**, **Secured**, **SED Enabled** flags are set as '**Yes**' for drives in slots 4 (**s4**) through 9 (**s9**).

```
/opt/MegaRAID/perccli/perccli64 /c0 /e64/sall show all | egrep -i '
(Policies/Settings |SED Capable|Secured|SED Enabled)'
```

```
Drive /c0/e64/s0 Policies/Settings :
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s1 Policies/Settings :
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s2 Policies/Settings :
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s3 Policies/Settings :
SED Capable = No
SED Enabled = No
```

```
Secured = No
Drive /c0/e64/s4 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s5 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
```

Appendix B. Encrypt a Series 6E Core or Hybrid Host (encryptSedVd.py)

NetWitness Series 6E Core and Hybrid hosts have Self-Encrypting Drives (SED). The `encryptSedVd.py` script:

- Validates that the Series 6E host has the correct setup for encryption.
- Encrypts unencrypted drives.

Note: For external storage devices such as PowerVault, refer to "Configure Storage Using the REST API" under "Using the REST API to Configure Storage" for instructions on how to encrypt their SED drives.

The following scenarios are examples of why you would use the `encryptSedVd.py` script.

- You want to know if a physical host has encryption. In this case, if the script determines that the device does not have encryption, it gives you the opportunity to encrypt it.
- You set up a device without encryption and you want to encrypt it.

You will find this script in the `rsa-sa-tools` directory for releases 11.4.0.0 and later. The following directory is for 11.4.0.0.

`rsa-sa-tools-11.4.0.0-xxxx.noarch.rpm`

The following procedure illustrates how to use the script.

1. Log in as `root`.
2. Change the directory to the rsa-sa-tools RPM base directory:

   ```
   cd /opt/rsa/saTools/supportScript/
   ```

3. Execute the following command:

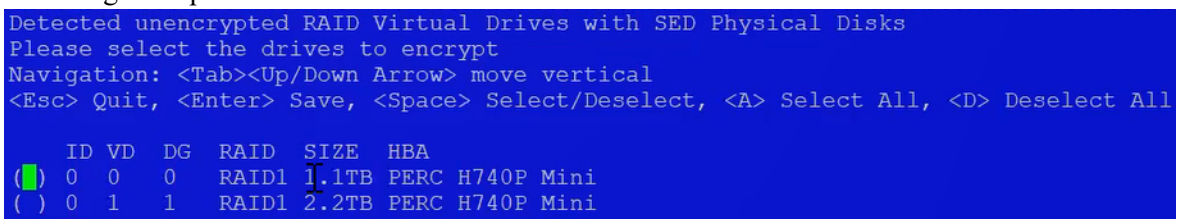
   ```
   OWB_ALLOW_NON_FIPS=1 ./encryptSedVd.py
   ```

   The script tells you if the disks are encrypted or not encrypted.

   - If the drives are encrypted, the script displays the following message.

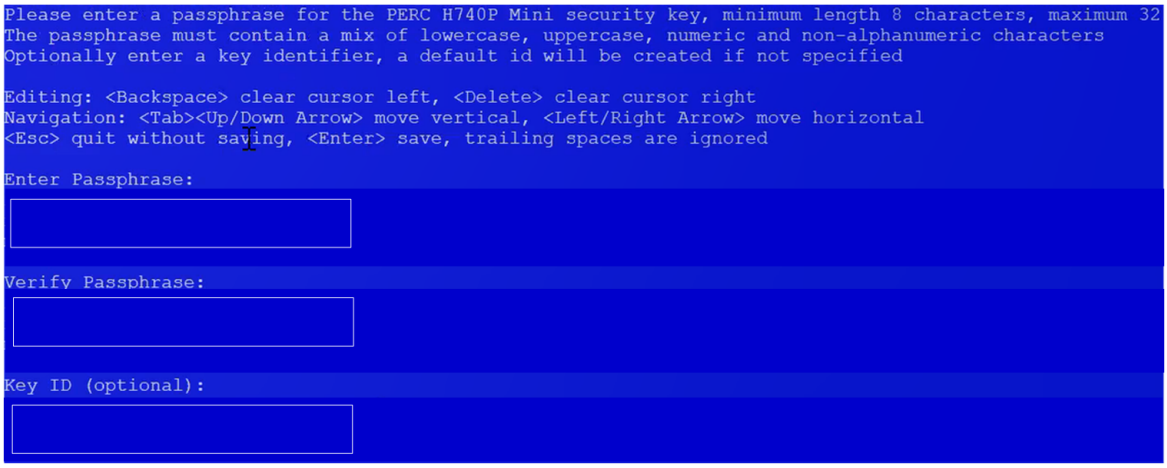
     ```
     No unencrypted RAID virtual drives with SED physical drives found.
     ```

   - If the drives are not encrypted, the script identifies the unencrypted drives as shown in the following example.



4. If the drives are not encrypted and you want to encrypt them:

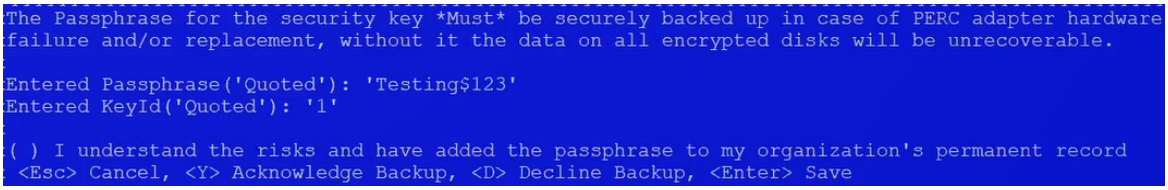
   a. Select the drives you want to encrypt with the space bar and press **Enter**.
   The following prompt is displayed.



   b. In the **Enter Passphrase** text box, type the `<passphrase>`, for example `nFreDaW$792`, and press **Tab**.

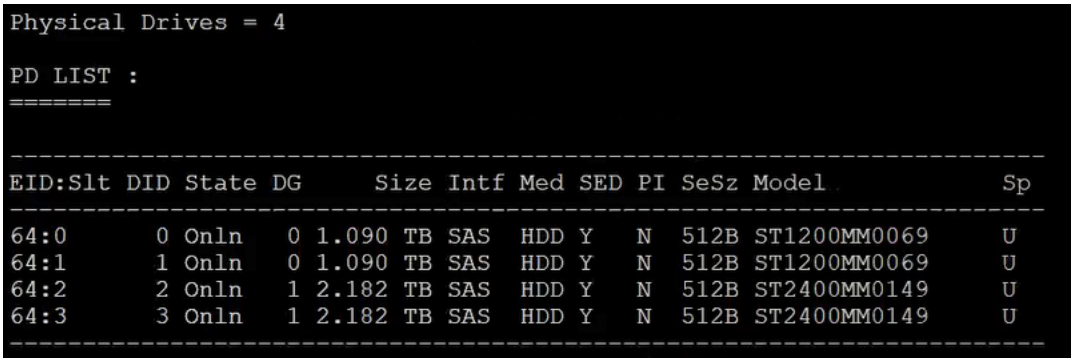
   c. In the **Verify Passphrase** text box, re-enter passphrase again for validation.

   d. In the **Key ID (optional)** text box, enter an optional ID string for the security key less than 256 characters or press Enter for none.
   The following prompt is displayed.



   e. Select `<Y>` and press **Enter** to confirm that you added the Passphrase.

   f. Submit the following command string to verify that the SED drives are encrypted.
   `/opt/MegaRAID/perccli/perccli64 /c0 show more`
   The following information is displayed. You can see that all four SED drives are encrypted (that is, Y is displayed for each drive in the SED column).



Note: The SED Enabled and Secured label values are set to Yes, if the drives are SED enabled and secured.
To check the drives on controller 0 and enclosure 247 use the below command:
`/opt/MegaRAID/perccli/perccli64 /c1 /e247/sall show all | egrep -i '(Policies/Settings|SED Capable|Secured|SED Enabled)'`

You will find detailed information on `perccli` commands in the Dell PowerEdge RAID Controller CLI Reference Guide (http://l4u-00.jinr.ru/pub/misc/h-w/LSI/dell-sas-hba-12gbps_reference-guide_en-us.pdf).

Enable SED on configured Drive Groups

Virtual Drives configured are SED Capable but are NOT SED Enabled.

To enable virtual drives or drive groups using PERC H840 Adaptors (External storage):

1. SSH to the appliance and run the below script to encrypt the virtual drive (on external storage).

Note: The encryptSedVd.py script turn on the SED feature only on Virtual Drives or Drive Groups on the PERC H840 Adaptors (external storage) and NOT on PERC H740 mini. Refer to Enable Virtual Drives / Drive Groups - PERC H740 (Mini) Adaptors (Internal storage) to enable SED on PERC H740 Mini .

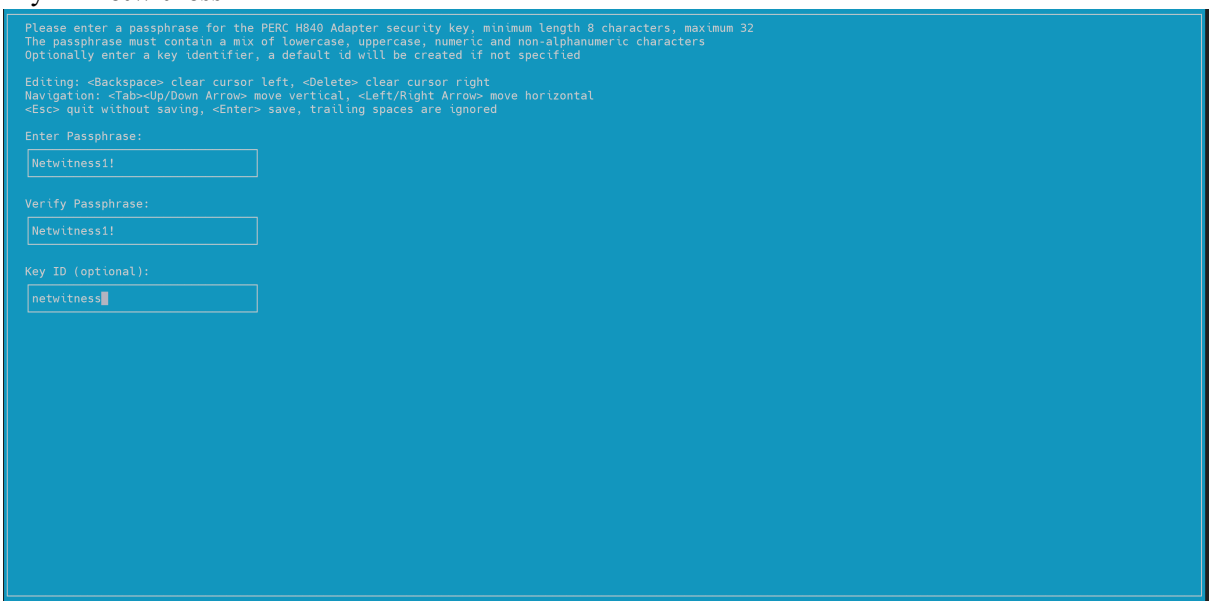
```
OWB_ALLOW_NON_FIPS=true /opt/rsa/saTools/supportScript/encryptSedVd.py
```

2. Select the Virtual Drive and press **Enter**.
   Passphrase screen is displayed.

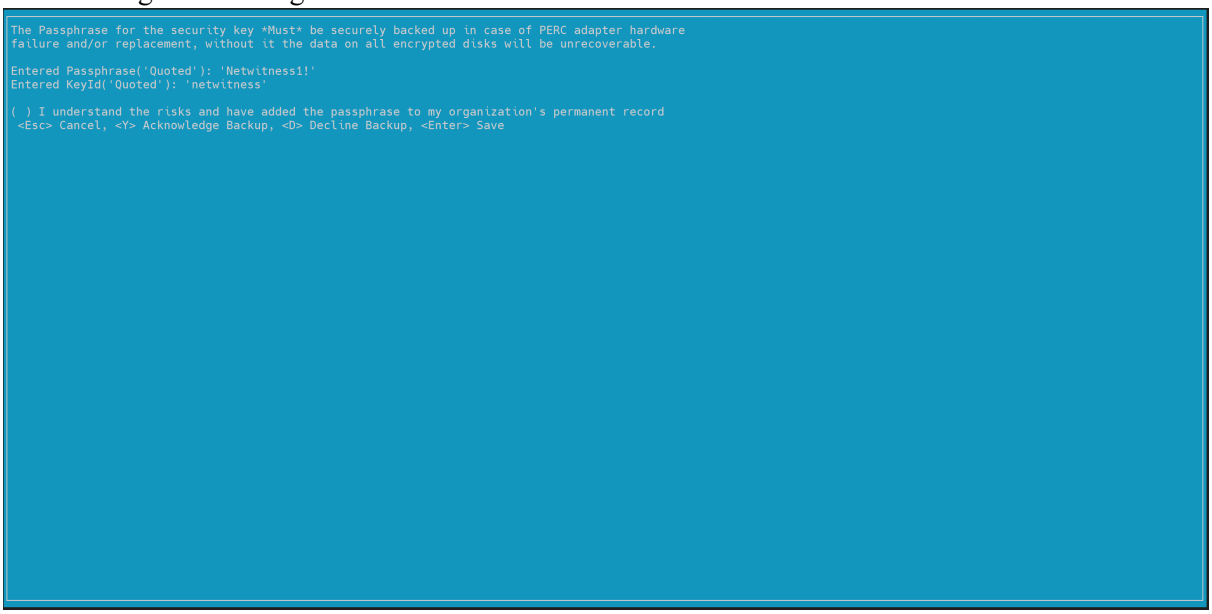
3. Enter the Passphrase and press **Enter**.
   For Example,

   Passphrase : **Netwitness1!**

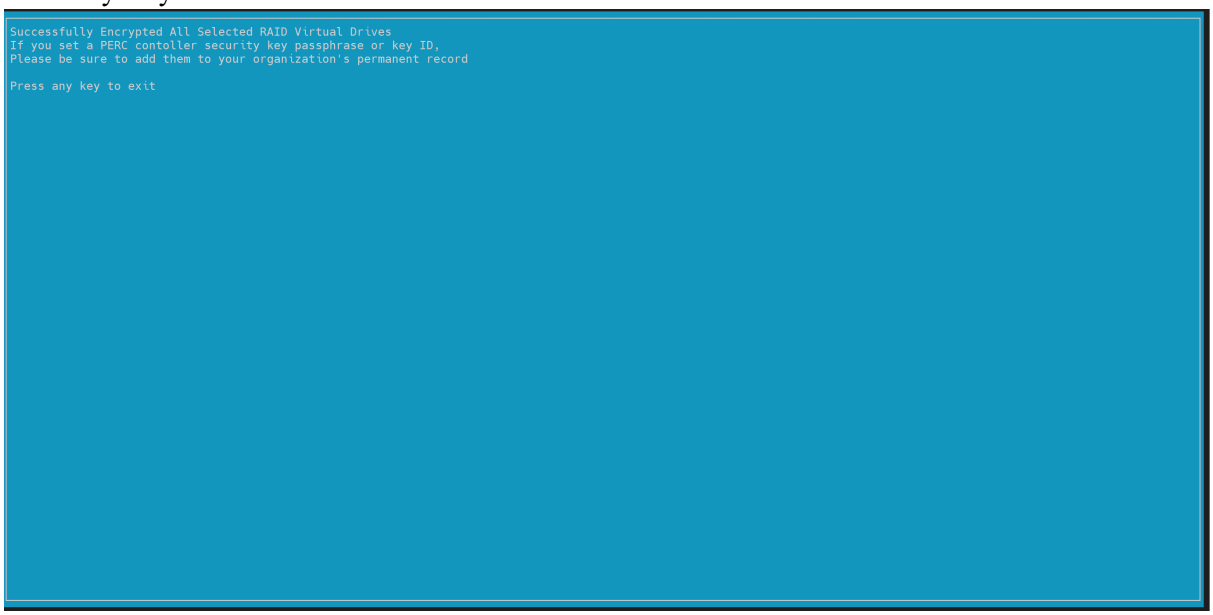
keyID: **netwitness**



4. Acknowledge the message and Press **Enter** to Save.



5. Press any Key to Exit.



6. To confirm that the drives are SED Enabled and secured, run the following command and verify the SED Enabled and Secured returns Yes.

```
/opt/MegaRAID/perccli/perccli64 /c1 /e247/sall show all | egrep -i '
(Policies/Settings|SED Capable|Secured|SED Enabled)'
Drive /c1/e247/s0 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s1 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s2 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s3 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s4 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s5 Policies/Settings :
SED Capable = Yes
```

```
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s10 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s11 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
```

Enable Virtual Drives / Drive Groups - PERC H740 (Mini) Adaptors (Internal storage)

You can enable the SED capability on the Virtual Drive or Drive Groups created out of on-board SED capable drives (in slots 4 through 9 – total of 6 drives) using the percli64 utility. You cannot use `/opt/rsa/saTools/supportScript/encryptSedVd.py` to turn on Security on the Virtual drives on the PERC H740 (mini) Adaptor.

1. SSH to the appliance and enable security on the PERC H740 (mini) Adaptor. The controller number for this adaptor is **0**. The PERC H840 Adaptor is shown as **1**.

To list all the controllers on the appliance run the following command:

```
/opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'
```

The first column (**Ctl**) lists out the controller index on the appliance. In this case, the controller **0** corresponds to **PERC H740 Mini** and controller **1** corresponds to **PERC H840 Adaptor**. The columns **DGs** and **VDs** dislpays the virtual drives and drive groups on the controller.

2. To enable the security on the **PERC H740 (mini) Adaptor**, for example, Controller **0**, run the following command:
   ```
   /opt/MegaRAID/perccli/perccli64 /c0 set securitykey=’<String>’!'
   keyid=’<String>’
   ```
   Example:
   ```
   /opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!'
   keyid='netwitness'
   ```
   ```
   'Netwitness1’ is the securityKey and ‘netwintess’ is ID.
   ```
   Make a note of both the Key and keyID securely.
3. Identify the correct Drive group (DG) or Virtual Drive (VD) corresponding to the SED Capable drives that you want to enable security.
   ```
   /opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'
   ```
   Check the first two and last column to identify the correct Drive Group / Virtual Drive correspond to the 6 SED enabled drives that are SED Capable. On Series 6 appliances, there is only one DG or VD with **RAID6** type. Name column can be used to identify the VD or DG. In this case, the DG or VD is **2**. Using a combination of **Type**, **Name** and **Size** columns (these are defined when you created VDs above).
4. To turn on Security on the disk group (created out of the 6 SED Capable drives) for **decodersmall** volume group, run the below command:
   ```
   /opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
   ```
5. Get the Enclosure ID (**EID**) using on the controller **0**. In this case, it is **64**
   ```
   /opt/MegaRAID/perccli/perccli64 /c0 /eall show
   ```
6. To confirm that the drives or drive groups are **SED Enabled** and **Secured**, run the below command and verify the **SED Capable**, **Secured**, **SED Enabled** flags are set as **Yes** for drives in slots 4 (s4) through 9 (s9).
   ```
   /opt/MegaRAID/perccli/perccli64 /c0 /e64/sall show all | egrep -i '
   (Policies/Settings |SED Capable|Secured|SED Enabled)'
   ```

Drive /c0/e64/s0 Policies/Settings :

```
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s1 Policies/Settings :
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s2 Policies/Settings :
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s3 Policies/Settings :
```

```
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s4 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s5 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
```

Enable SED on configured Virtual Drives/ Drive Groups on Power Vault (PERC 840)

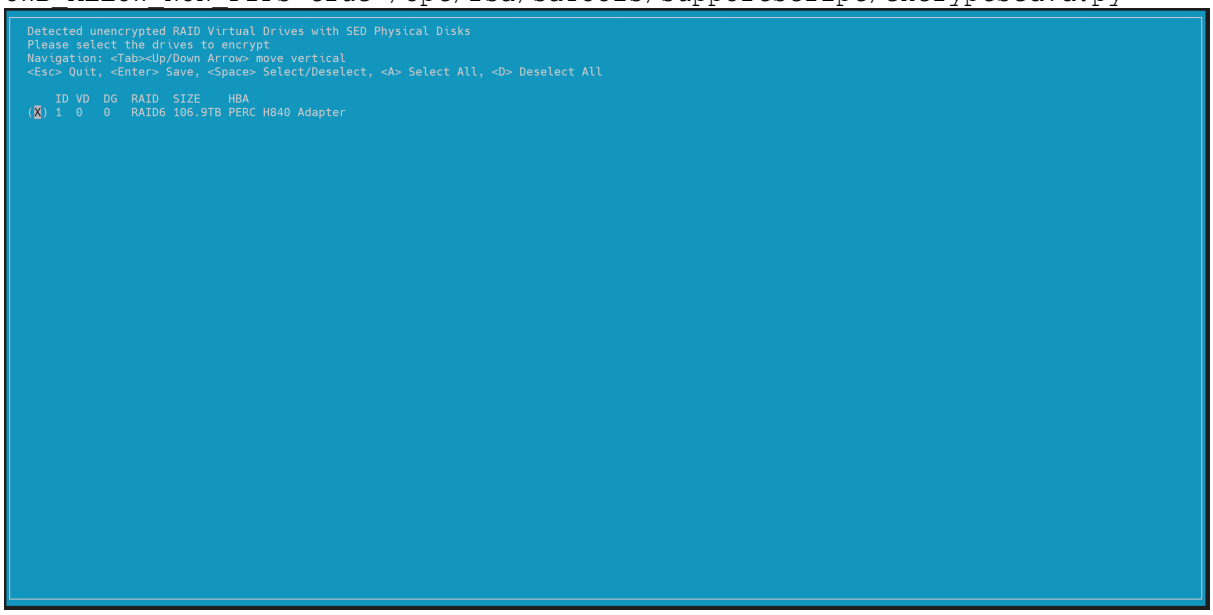
Enable Virtual Drives / Drive Groups - PERC H840 Adaptors

Note: The virtual disk created in *Configure Block Devices for PowerVaults* section in Prepare Physical Storage is SED capable but NOT SED Enabled.

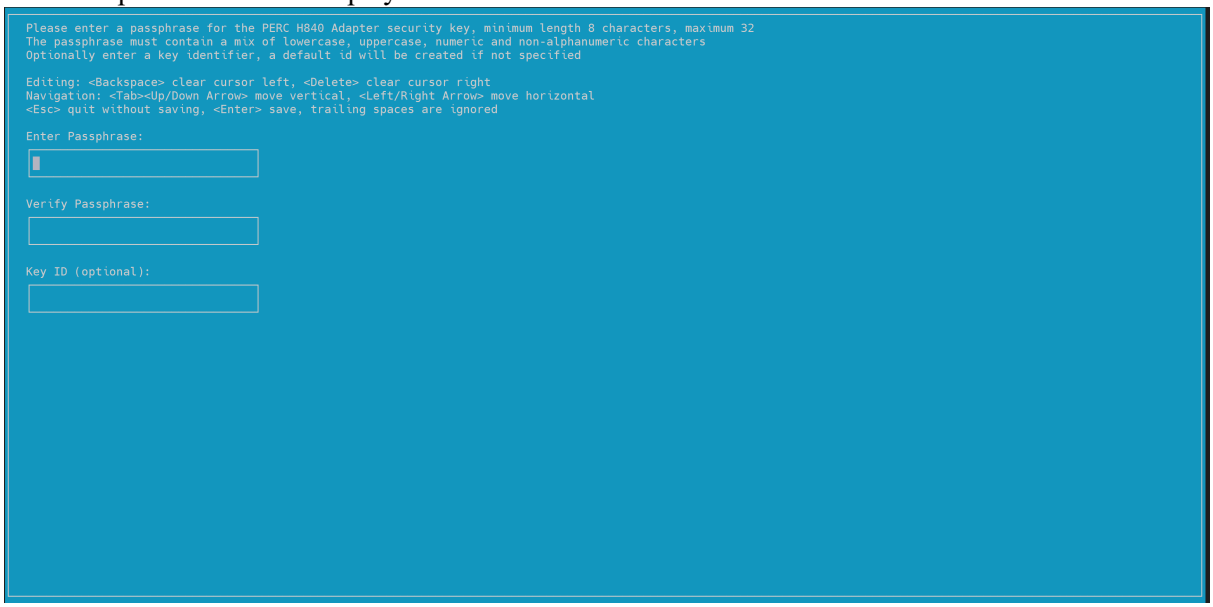
1. To enable, SSH into the appliance and run the below script to encrypt the VD (on external storage).
   `OWB_ALLOW_NON_FIPS=true /opt/rsa/saTools/supportScript/encryptSedVd.py`

   Note: The `encryptSedVd.py` script turn on the SED feature only on Virtual Drives or Drive Groups on the PERC H840 Adaptors (external storage) and NOT on PERC H740 mini. Refer to **Enable Virtual Drives / Drive Groups - PERC H740 (Mini) Adaptors (Internal storage)** to enable SED on PERC H740 Mini

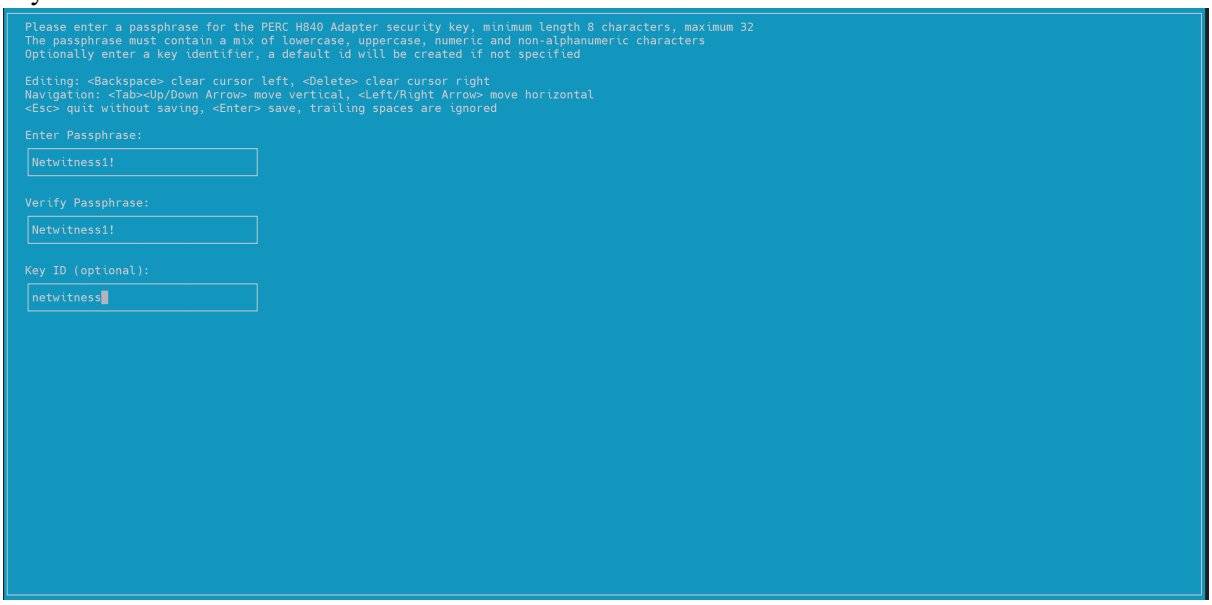
   `OWB_ALLOW_NON_FIPS=true /opt/rsa/saTools/supportScript/encryptSedVd.py`



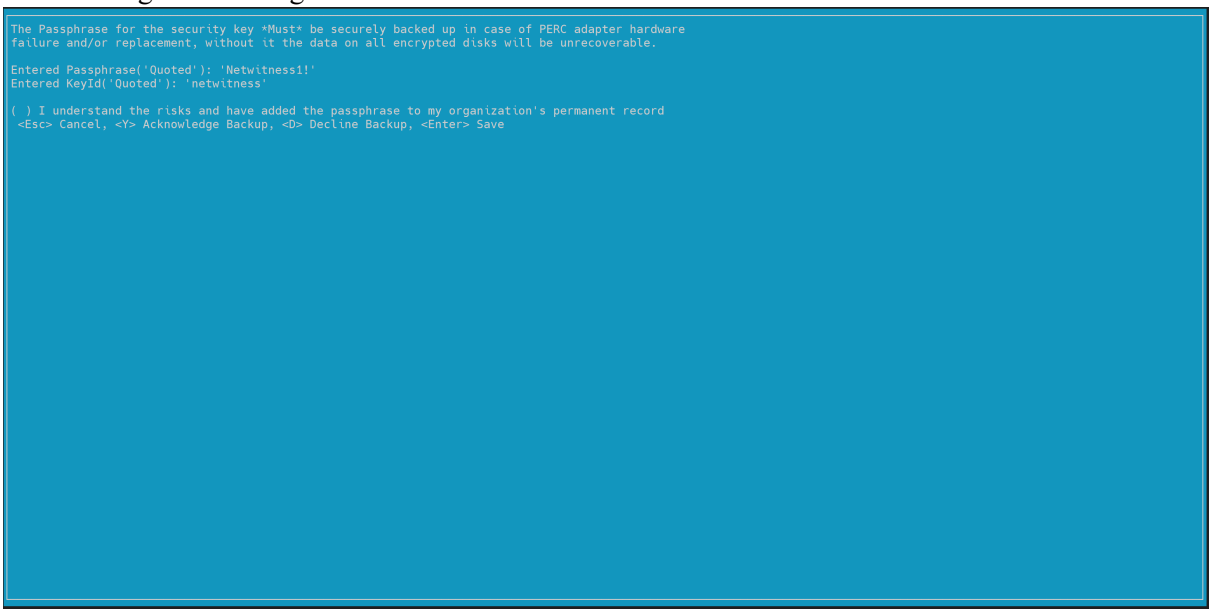
2. Select both the **Virtula Disks** and press **Enter**.
   The Passphrase screen is displayed.



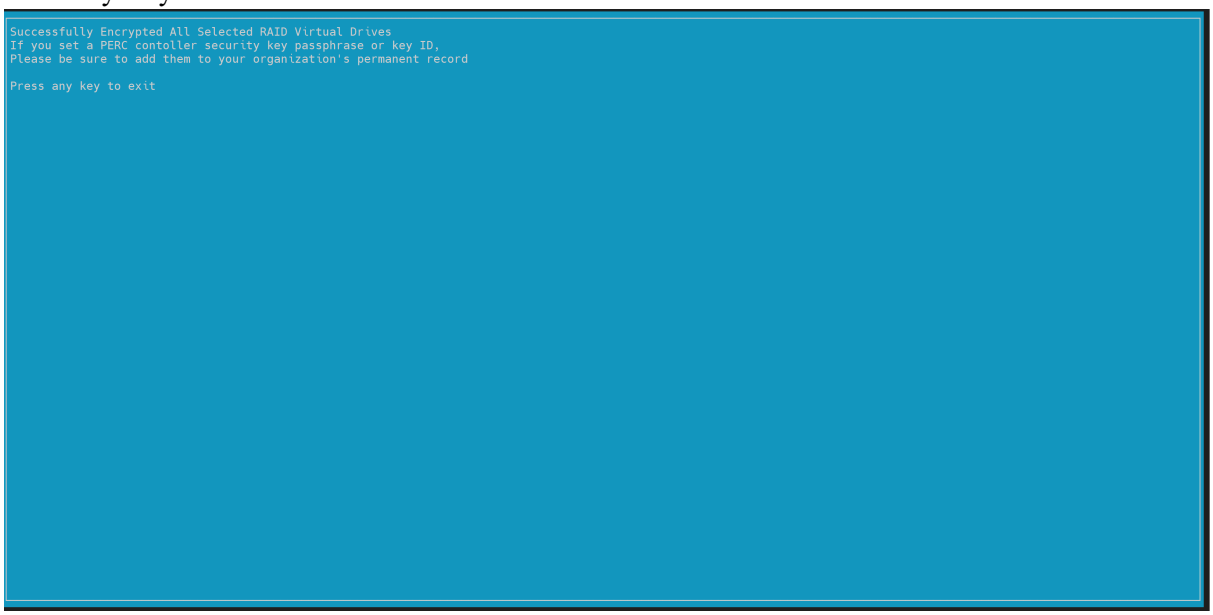
3. Enter the Passphrase and press **Enter**.
   For example,
   Passphrase : **Netwitness1!**
   keyID: **netwitness**



4. Acknowledge the message and Press **Enter** to Save.



5. Press any Key to Exit.



6. To confirm that the drives are SED Enabled and secured, run the below command and verify the SED Enabled and Secured returns Yes.
   `/opt/MegaRAID/perccli/perccli64 /c1 /e247/sall show all | egrep -i '(Policies/Settings|SED Capable|Secured|SED Enabled)'`
   Drive /c1/e247/s0 Policies/Settings :

```
SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s1 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s2 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s3 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s4 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes
```

```
Secured = Yes
Drive /c1/e247/s5 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s10 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s11 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
```

Enable Security on SED Capable Drive groups on Host with a mix of SED and NON SED Drives

The encryptSedVd.py may fail to identify the SED Capable Virtual Drives when there is mix of both SED and NON-SED drives on the appliance. The below steps are applicable when both SED and NON-SED capable virtual drives exist on the host.

1. SSH to the appliance and enable security on the PERC H740 (mini) Adaptor. The controller number for this adaptor is **0**. The PERC H840 Adaptor is shown as **1**.
   To list all the controllers on the appliance:
   `/opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'`
   The first column **(Ctl)** lists out the controller index on the appliance. In this case, the controller '**0**' corresponds to '**PERC H740 Mini**' and controller '**1**' corresponds to '**PERC H840 Adaptor**' . The columns '**DGs**' and '**VDs**' displays the drive groups and virtual drives on the controller.

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'
Ctl Model          Ports PDs DGs DNOpt VDs VNOpt BBU sPR DS EHS ASOs Hlth
--------------------------------------------------------------------------
  0 PERCH740PMini      8  10   3     0   3     0 Opt On  -  N      0 Opt
  1 PERCH840Adapter    8  12   1     0   1     0 Opt On  -  N      0 Opt
[root@116Decoder perccli]#
```

2. To enable the security on the 'PERC H740 (mini) Adaptor' i.e Controller '**0**', execute the following command:
   `/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='<SOME_STRING_VALUE>'!' keyid='< SOME_STRING_VALUE >'`

   Example:
   `/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!' keyid=1`

   'Netwitness1' is the securityKey and '1' is ID. Preserve both the Key and keyID securely.

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!' keyid='netwitness'
Controller = 0
Status = Success
Description = None


Controller Properties :
=====================

---------------------
Ctrl Method  Result
---------------------
   0 set Key Success
---------------------
```

3. Identify the correct Drive group (DG) / Virtual Drive (VD) corresponding to the SED Capable drives that we are trying to enable security.
   `/opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'`

   Refer to first two and last column to identify the correct Drive Group (DG) / Virtual Drive (VD) correspond to the 6 SED enabled drives. On Series 6 appliances, there is only one DG/VD with **RAID6**. '**Name**' column can be used to identify the VD/DG. In this case, the DG/VD is '**2**'. Using a combination of '**Type**', '**Name**' and '**Size**' columns (these were defined by the user when the VDs are created above).

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'
DG/VD TYPE  State Access Consist Cache Cac sCC     Size Name
------------------------------------------------------------------------
0/0   RAID1 Optl  RW     Yes     RWBD  -   OFF 931.0 GB
1/1   RAID1 Optl  RW     Yes     RWBD  -   OFF 1.818 TB
2/2   RAID6 Optl  RW     Yes     RWBD  -   OFF 8.730 TB Virtual Disk 2
------------------------------------------------------------------------
[root@116Decoder perccli]#
```

4. To turn on Security on the disk group (created out of the 6 SED Capable drives), execute the below command:

```
/opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
```

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
Controller = 0
Status = Success
Description = Success
```

5. Get the Enclosure ID (**EID**) using on the controller '**0**'. In this case, it is '64'

```
/opt/MegaRAID/perccli/perccli64 /c0 /eall show
```

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /eall show
Controller = 0
Status = Success
Description = None


Properties :
==========

------------------------------------------------------------------------------
EID State Slots PD PS Fans TSs Alms SIM Port#      ProdID    VendorSpecific
------------------------------------------------------------------------------
 64 OK       10 10  0    0    0    0    1 00 & 00 x8 BP14G+EXP  +
------------------------------------------------------------------------------

EID-Enclosure Device ID |PD-Physical drive count |PS-Power Supply count|
TSs-Temperature sensor count |Alms-Alarm count |SIM-SIM Count


[root@116Decoder perccli]#
```

6. To confirm that the drives / Drive Groups (DG) are **SED Enabled** and **Secured**, run the below command and verify the **SED Capable**, **Secured**, **SED Enabled** flags are set as '**Yes**' for drives in slots 4 (**s4**) through 9 (**s9**).

```
/opt/MegaRAID/perccli/perccli64 /c0 /e64/sall show all | egrep -i '
(Policies/Settings |SED Capable|Secured|SED Enabled)'

Drive /c0/e64/s0 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No

Drive /c0/e64/s1 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No

Drive /c0/e64/s2 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No
```

```
Drive /c0/e64/s3 Policies/Settings :
SED Capable = No
SED Enabled = No
Secured = No
Drive /c0/e64/s4 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s5 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
```

Appendix C. Troubleshooting

This section contains instructions on how to resolve various storage tasks using the REST API.

Reconfigure Pre-Configured DAC Attached to Decoder Using REST API

Th scenario covers how to reconfigure a DAC using the REST API that was configured using another tool and clear any pre-existing data (if no longer need or backed up to another storage device).

The following information describes the state of the host and storage hardware prior to the attempt to reconfigure the storage devices using the REST API.

When the DAC was added, it had old data and was configured (but not using the REST API). This prevented the REST API from executing the `raidNew` command and returned the `Physical disk does not have appropriate attributes` error message.

The following steps describe the scenario and with its resolution.

1. From the Decoder Linux console (or SSH to Decoder), submitted the following command string.
   `/opt/MegaRAID/perccli/perccli64 /c2/fall del`
   You will find detailed information on perccli commands in the **Dell PowerEdge RAID Controller CLI Reference Guide** (https://topics-cdn.dell.com/pdf/dell-sas-hba-12gbps_reference-guide_en-us.pdf).
   This deleted all foreign configuration from controller 2 and cleared all data from the DAC.
2. Tried to partition the DAC, but the `partNew` command failed because that information was already defined on the DAC. `partNew` displayed that you must use one an available device, but `devList` showed it in use.
3. Assuming that the partitions were defined, tried to allocate the storage devices, but this did not work because the DAC was not mounted.
4. Tried to mount the DAC from the command line, but received `mount failed: structure needs to be cleaned` error message.
5. There was no data that needed to be preserved on the DAC, so submitted the following command strings to clean the structure.
   ```
   mkfs.xfs -f /dev/decoder0/packetdb
   mkfs.xfs -f /dev/decoder1/packetdb
   ```
6. Mounted devices to their appropriate locations in `/var/netwitness/decoder`.
7. Completed the remainder of the applicable steps as described in Configure Storage Using the REST API to reconfigure the DAC

Appendix D. Sample Storage Configuration Scenarios for 15-Drive DACs

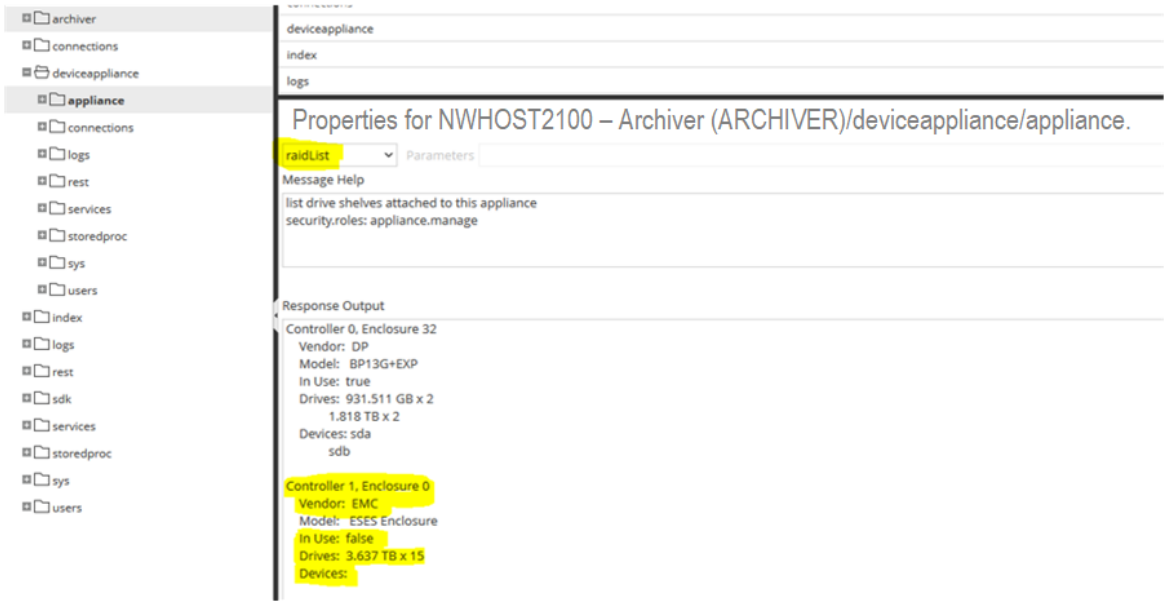
This appendix illustrates the following example of how to configure storage on two non-encrypted 15-drive DAC external storage devices.

- Configure Storage for Archiver
- Configure Storage for Network (Packet) Decoder
- Configure Storage for Network Concentrator
- Configure Storage for Log Decoder Hybrid

Configure Storage for Archiver

The following scenario configures storage on one, non-encrypted, 15-Drive DAC for an Archiver physical host.

1. Execute the `raidList` command.
   a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
   You should see the following information.
   ```
   In Use: FALSE
   Devices: <empty>
   ```
   b. Verify the Drive Count, Size, and Vendor.
   The following example illustrates what you should see before you create a RAID array.

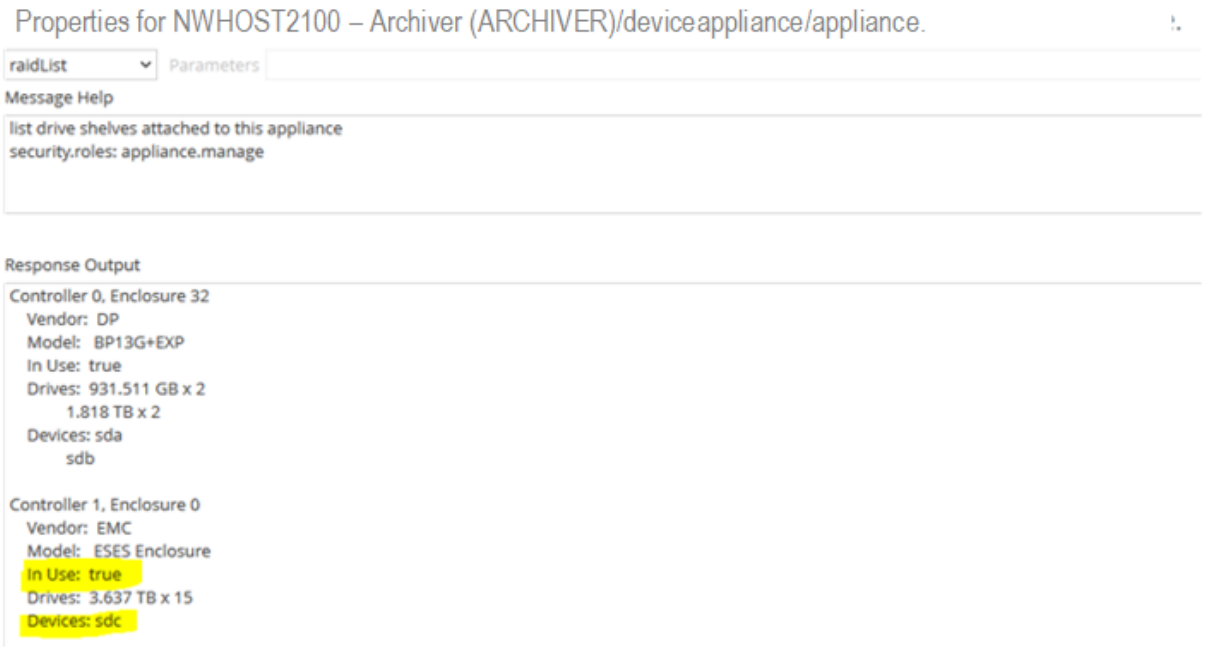


2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.
   ```
   controller=1 enclosure=0 scheme=archiver commit=1
   ```

The following example illustrates what you should see after you create a RAID array.

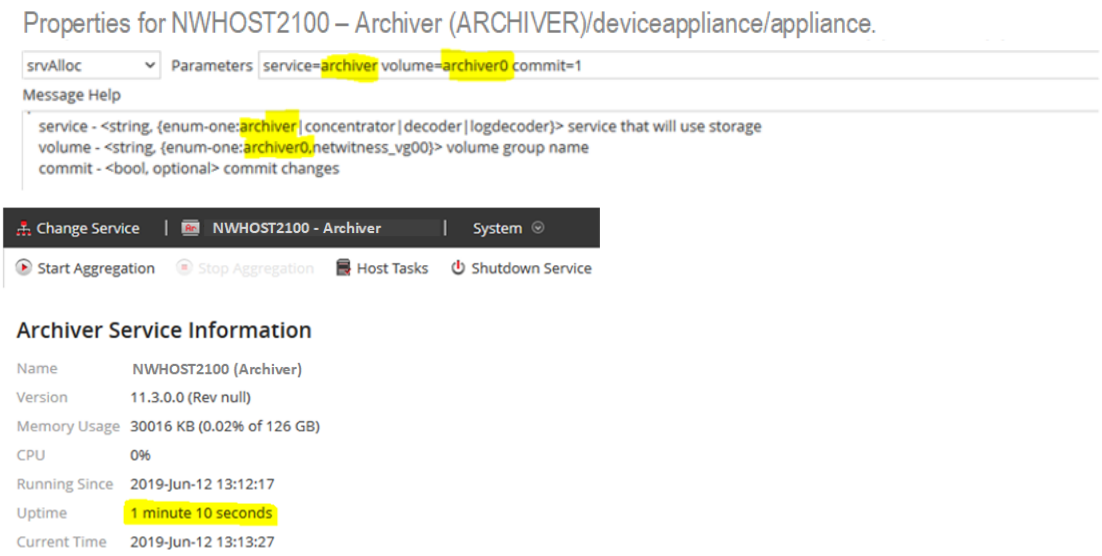


3. Execute the `raidList` command to verify the new RAID array.
   You should now see the following information.
   `In Use: TRUE`
   `Devices: <device>` (for example, `sdc`)

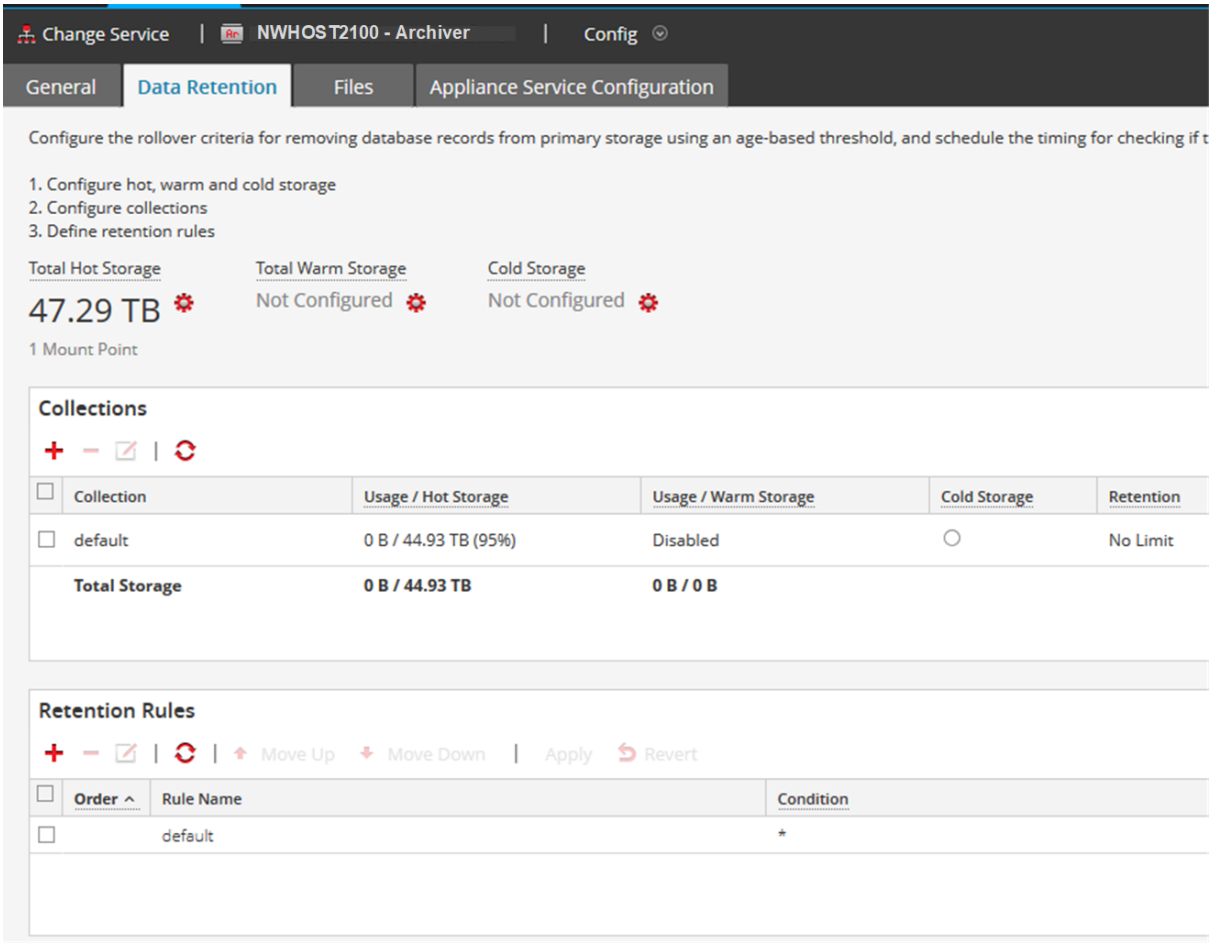


4. Execute the `partNew` command with the following parameters to create partitions and mount points in the `etc/fstab` file.
   `name=<device> (for example, sdc) service=archiver volume=archiver commit=1`

5. Execute the `srvAlloc` command with the following parameters to allocate the space to the archiver service. This adds storage to the archiver service configuration and restarts the service every time it is executed.
   `service=archiver volume=archiver0 commit=1`



6. Confirm the "Hot Storage" in "Data Retention".



7. Reconfigure the following Archiver service to detect and take advantage of all of the free space as described in Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture.

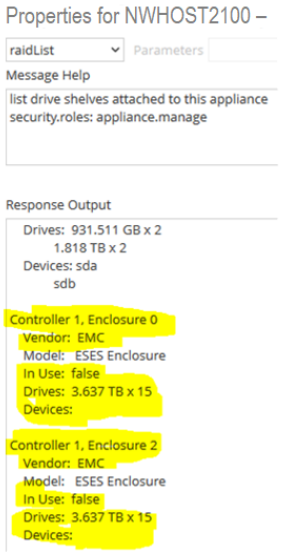
Configure Storage for Network (Packet) Decoder

The following scenario configures storage on two, non-encrypted, 15-Drive DACs for a Network Decoder for 10G Capture physical host.

1. Execute the `raidList` command.

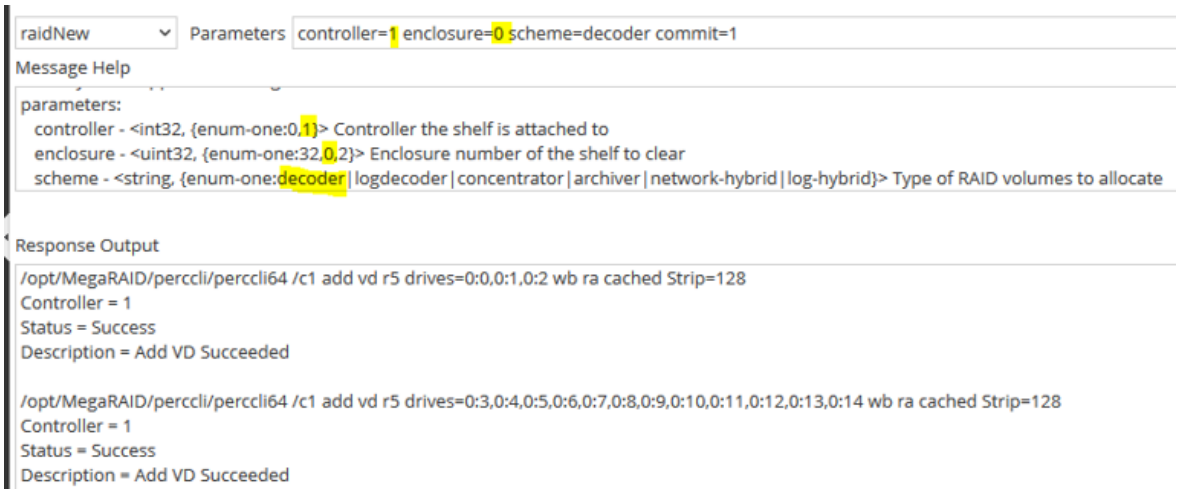
   a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
   You should see the following information.
   ```
   In Use: FALSE
   Devices: <empty>
   ```

   b. Verify the Drive Count, Size, and Vendor.
   The following example illustrates what you should see before you create a RAID array.

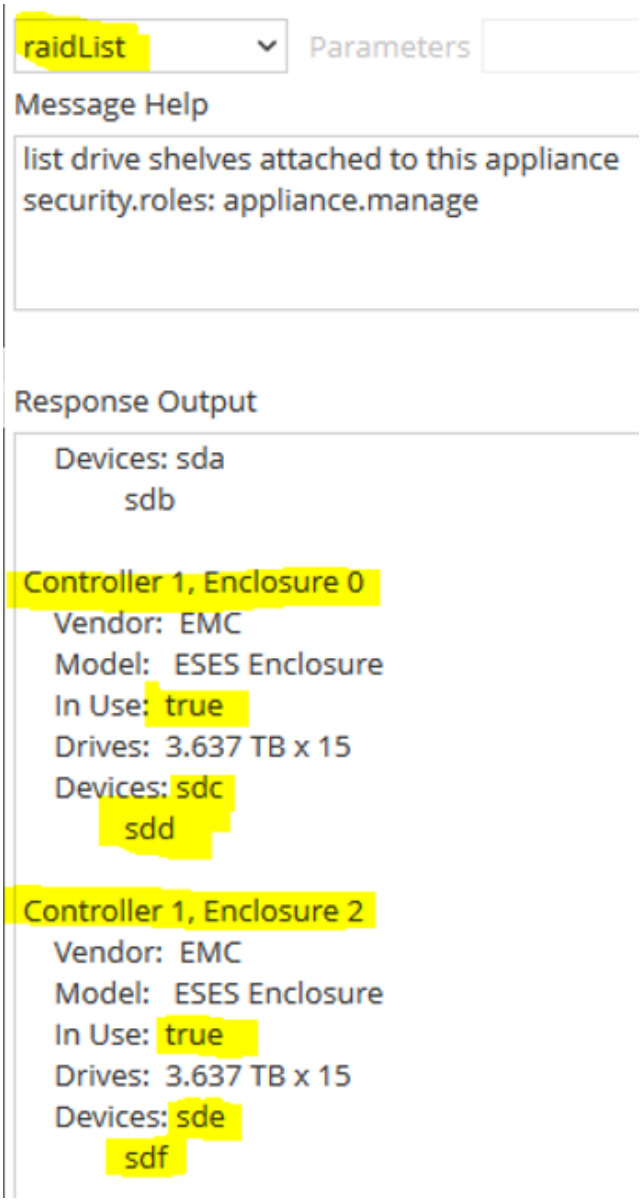


2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

   - Parameters for the first enclosure:
   ```
   controller=1 enclosure=0 scheme=decoder commit=1
   ```

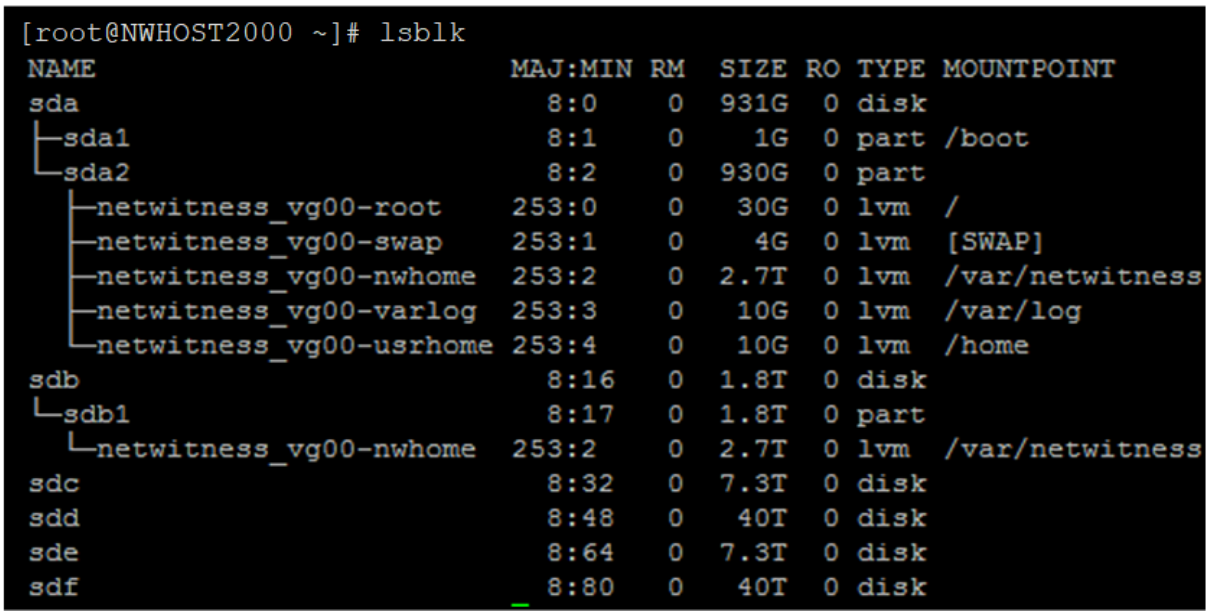

raidNew
Parameters
controller=1 enclosure=0 scheme=decoder commit=1
Message Help
parameters:
controller - <int32, {enum-one:0,1}> Controller the shelf is attached to
enclosure - <uint32, {enum-one:32,0,2}> Enclosure number of the shelf to clear
scheme - <string, {enum-one:decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid}> Type of RAID volumes to allocate
Response Output
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=0:0,0:1,0:2 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=0:3,0:4,0:5,0:6,0:7,0:8,0:9,0:10,0:11,0:12,0:13,0:14 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded

- Parameters for the second enclosure:
  `controller=1 enclosure=2 scheme=decoder commit=1`



3. Use the `raidList` command to display block devices for enclosures so you can verify `In Use: TRUE`.

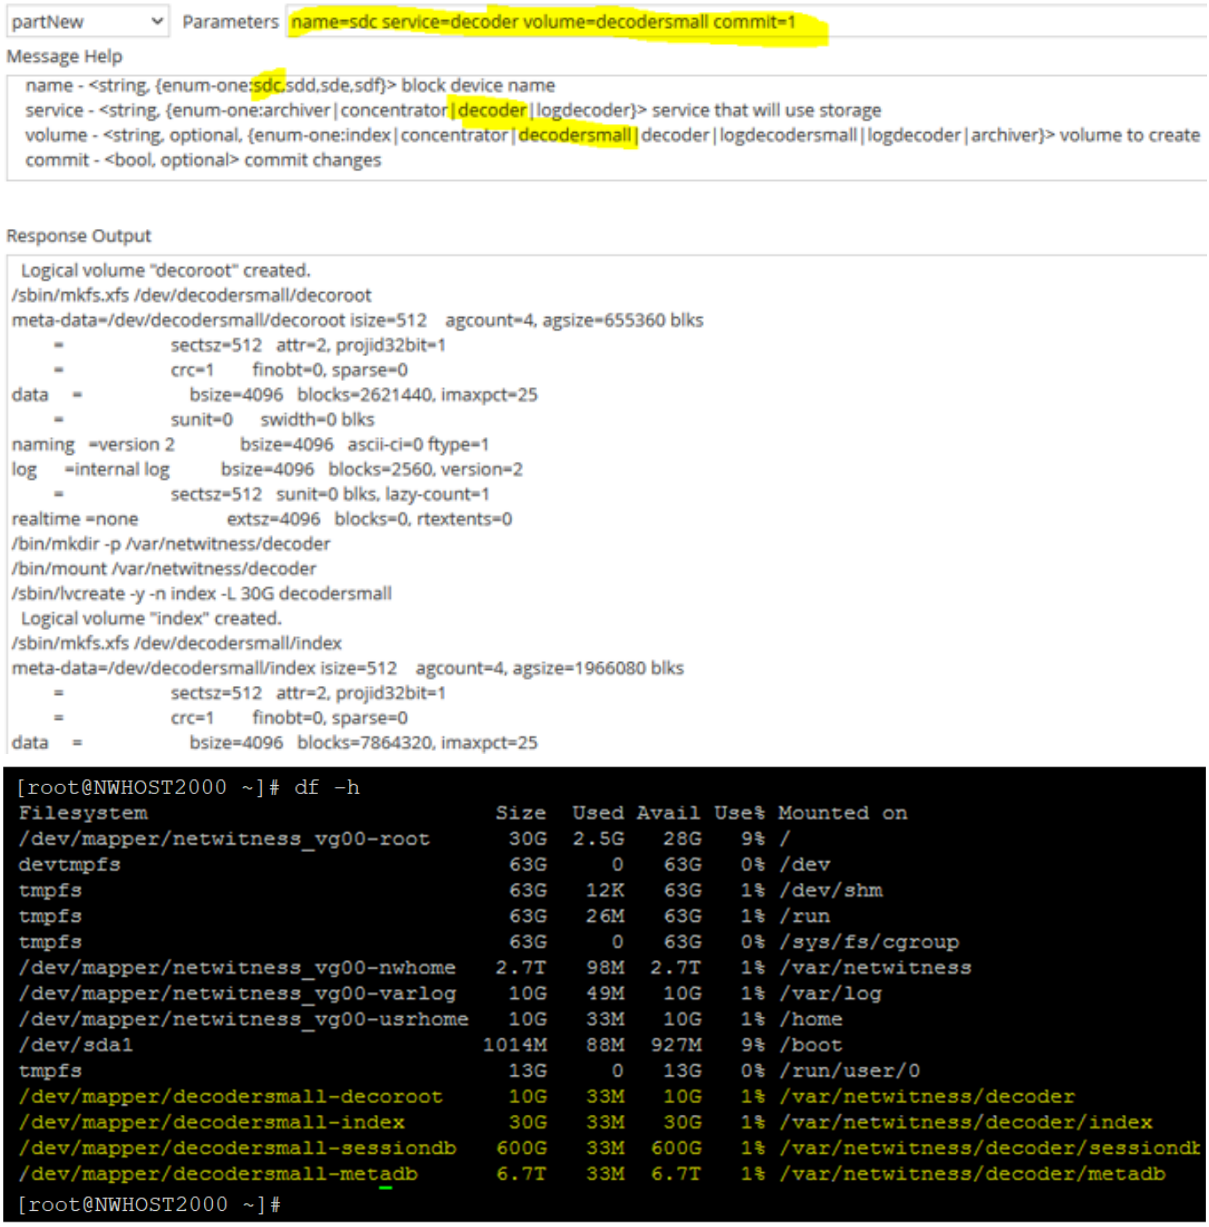
4. SSH to the Network Decoder and use the `lsblk` command to confirm sizes for **decodersmall**.



Note: For RAID configuration, when you use the decoder for 10G Capture you use **decoder** for both enclosures for performance reasons. When you do not use the **decoder** for 10G Capture, you use the decoder and archiver for the enclosures to maximize storage for becuase the second enclosure is a single RAID under the **archiver** configuration.

5. Execute the `partNew` command to create the **decodersmall** partition first (decoder dir, index, metadb, sessiondb) (First Enclosure, SDC, SDD) with the following parameters.

```
name=sdc service=decoder volume=decodersmall commit=1
```



partNew Parameters name=sdc service=decoder volume=decodersmall commit=1

Message Help

```
name - <string, {enum-one:sdc,sdd,sde,sdf}> block device name
service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create
commit - <bool, optional> commit changes
```

Response Output

```
 Logical volume "decoroot" created.
/sbin/mkfs.xfs /dev/decodersmall/decoroot
meta-data=/dev/decodersmall/decoroot isize=512    agcount=4, agsize=655360 blks
      =                   sectsz=512   attr=2, projid32bit=1
      =                   crc=1        finobt=0, sparse=0
data     =                  bsize=4096   blocks=2621440, imaxpct=25
      =                   sunit=0      swidth=0 blks
naming   =version 2          bsize=4096   ascii-ci=0 ftype=1
log      =internal log       bsize=4096   blocks=2560, version=2
      =                   sectsz=512   sunit=0 blks, lazy-count=1
realtime =none               extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder
/bin/mount /var/netwitness/decoder
/sbin/lvcreate -y -n index -L 30G decodersmall
 Logical volume "index" created.
/sbin/mkfs.xfs /dev/decodersmall/index
meta-data=/dev/decodersmall/index isize=512    agcount=4, agsize=1966080 blks
      =                   sectsz=512   attr=2, projid32bit=1
      =                   crc=1        finobt=0, sparse=0
data     =                  bsize=4096   blocks=7864320, imaxpct=25
```

```
[root@NWHOST2000 ~]# df -h
Filesystem                              Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root         30G  2.5G   28G   9% /
devtmpfs                                 63G     0   63G   0% /dev
tmpfs                                    63G   12K   63G   1% /dev/shm
tmpfs                                    63G   26M   63G   1% /run
tmpfs                                    63G     0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-nwhome      2.7T   98M  2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog       10G   49M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome      10G   33M   10G   1% /home
/dev/sda1                              1014M   88M  927M   9% /boot
tmpfs                                    13G     0   13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot        10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index           30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb      600G   33M  600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb         6.7T   33M  6.7T   1% /var/netwitness/decoder/metadb
[root@NWHOST2000 ~]#
```

6. Execute the `partNew` command to create the decoder volume (packetdb) (First Enclosure, SDC, SDD) with the following parameters.

   `name==sdd service=decoder volume=decoder commit=1`

partNew Parameters name=sdd service=decoder volume=decoder commit=1

Message Help

```
name - <string, {enum-one:sdc,sdd,sde,sdf}> block device name
service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create
commit - <bool, optional> commit changes
```

Response Output

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
  Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f decoder /dev/sdd1
  Volume group "decoder" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder
  Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder/packetdb
meta-data=/dev/decoder/packetdb  isize=512    agcount=41, agsize=268435455 blks
      =                  sectsz=512   attr=2, projid32bit=1
      =                  crc=1        finobt=0, sparse=0
data     =                 bsize=4096   blocks=10742791168, imaxpct=5
      =                  sunit=0      swidth=0 blks
naming   =version 2         bsize=4096   ascii-ci=0 ftype=1
log      =internal log      bsize=4096   blocks=521728, version=2
      =                  sectsz=512   sunit=0 blks, lazy-count=1
realtime =none              extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/packetdb
/bin/mount /var/netwitness/decoder/packetdb
```

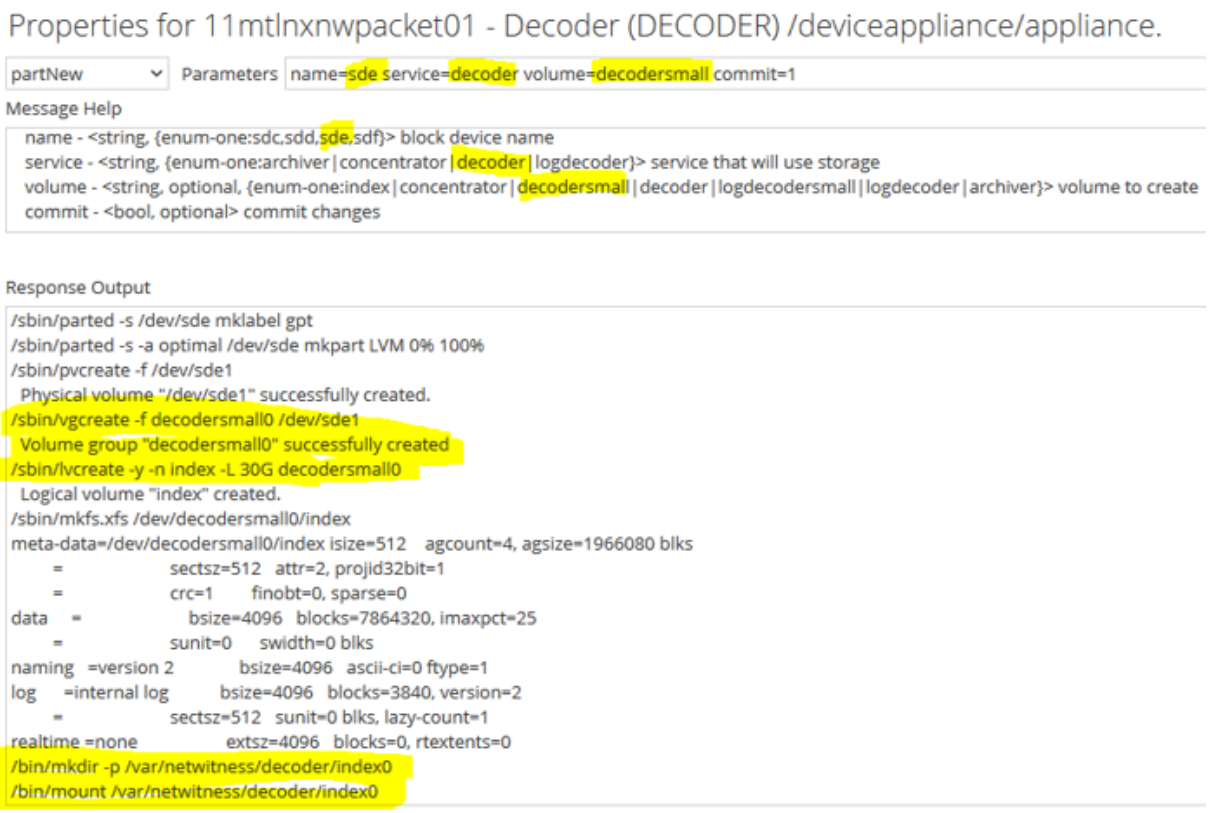
```
[root@NWHOST2000 ~]# df -h
Filesystem                          Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root     30G  2.5G   28G   9% /
devtmpfs                             63G     0   63G   0% /dev
tmpfs                                63G   12K   63G   1% /dev/shm
tmpfs                                63G   26M   63G   1% /run
tmpfs                                63G     0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-nwhome  2.7T   98M  2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog   10G   50M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome  10G   33M   10G   1% /home
/dev/sda1                          1014M   88M  927M   9% /boot
tmpfs                                13G     0   13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot    10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index       30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb  600G   33M  600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb     6.7T   33M  6.7T   1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb         41T   34M   41T   1% /var/netwitness/decoder/packetdb
```

In the following example, the following partions are created for SDC, SDD (Enclosure 0).

```
[root@NWHOST2000 ~]# lsblk
NAME                          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                             8:0    0  931G  0 disk
├─sda1                          8:1    0    1G  0 part /boot
└─sda2                          8:2    0  930G  0 part
  ├─netwitness_vg00-root      253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0  2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0   10G  0 lvm  /home
sdb                             8:16   0  1.8T  0 disk
└─sdb1                          8:17   0  1.8T  0 part
  └─netwitness_vg00-nwhome    253:2    0  2.7T  0 lvm  /var/netwitness
sdc                             8:32   0  7.3T  0 disk
└─sdc1                          8:33   0  7.3T  0 part
  ├─decodersmall-decoroot     253:5    0   10G  0 lvm  /var/netwitness/decoder
  ├─decodersmall-index        253:6    0   30G  0 lvm  /var/netwitness/decoder/index
  ├─decodersmall-sessiondb    253:7    0  600G  0 lvm  /var/netwitness/decoder/sessiondb
  └─decodersmall-metadb       253:8    0  6.7T  0 lvm  /var/netwitness/decoder/metadb
sdd                             8:48   0   40T  0 disk
└─sdd1                          8:49   0   40T  0 part
  └─decoder-packetdb          253:9    0   40T  0 lvm  /var/netwitness/decoder/packetdb
sde                             8:64   0  7.3T  0 disk
sdf                             8:80   0   40T  0 disk
```

At this point, you add the second DAC enclosure.

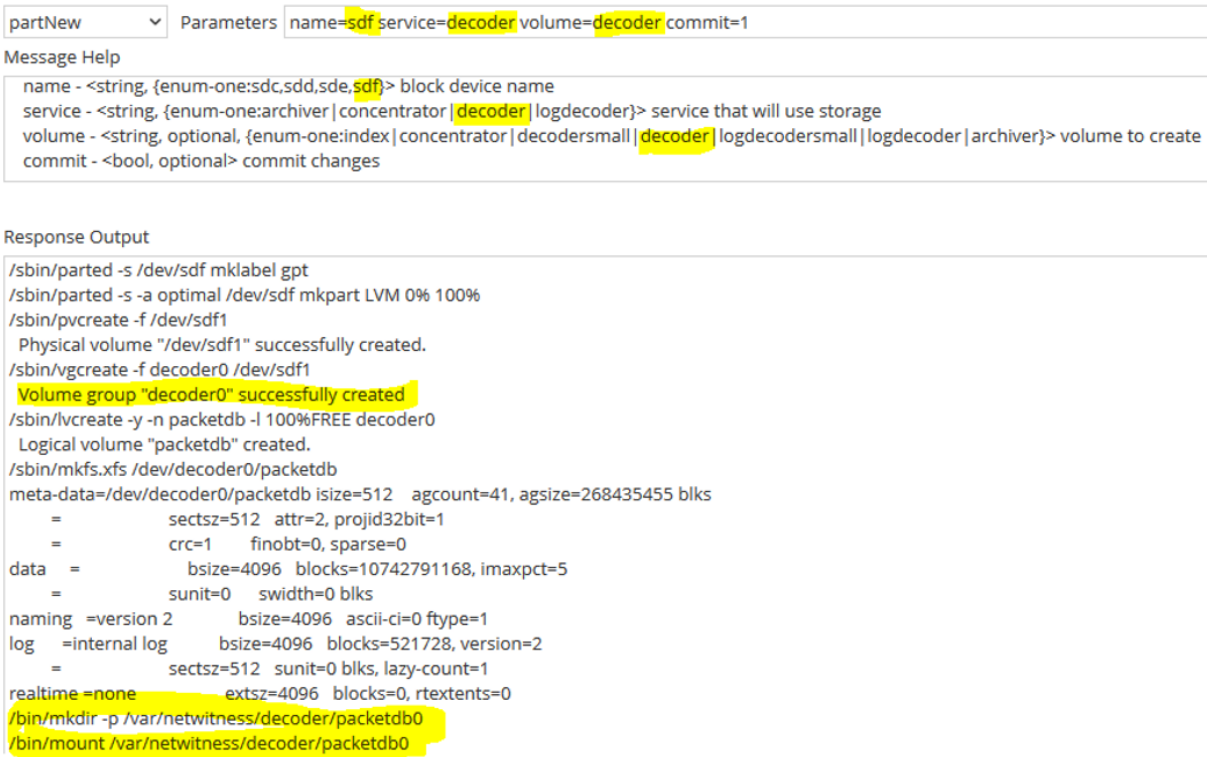
7. Execute the `partNew` command to create the decodersmall partition first (Second Enclosure, SDE, SDF) with the following parameters.
   `name=sde service=decoder volume=decodersmall commit=1`



```
[root@NWHOST2000 ~]# df -h
Filesystem                              Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root         30G  2.5G   28G   9% /
devtmpfs                                 63G     0   63G   0% /dev
tmpfs                                    63G   12K   63G   1% /dev/shm
tmpfs                                    63G   26M   63G   1% /run
tmpfs                                    63G     0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-nwhome      2.7T   98M  2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog       10G   50M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome      10G   33M   10G   1% /home
/dev/sda1                              1014M   88M  927M   9% /boot
tmpfs                                    13G     0   13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot        10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index           30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb      600G   33M  600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb         6.7T   33M  6.7T   1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb_            41T   34M   41T   1% /var/netwitness/decoder/packetdb
```

8. Execute the `partNew` command to create the `packetdb` decoder volume (Second Enclosure SDE, SDF) with the following parameters.

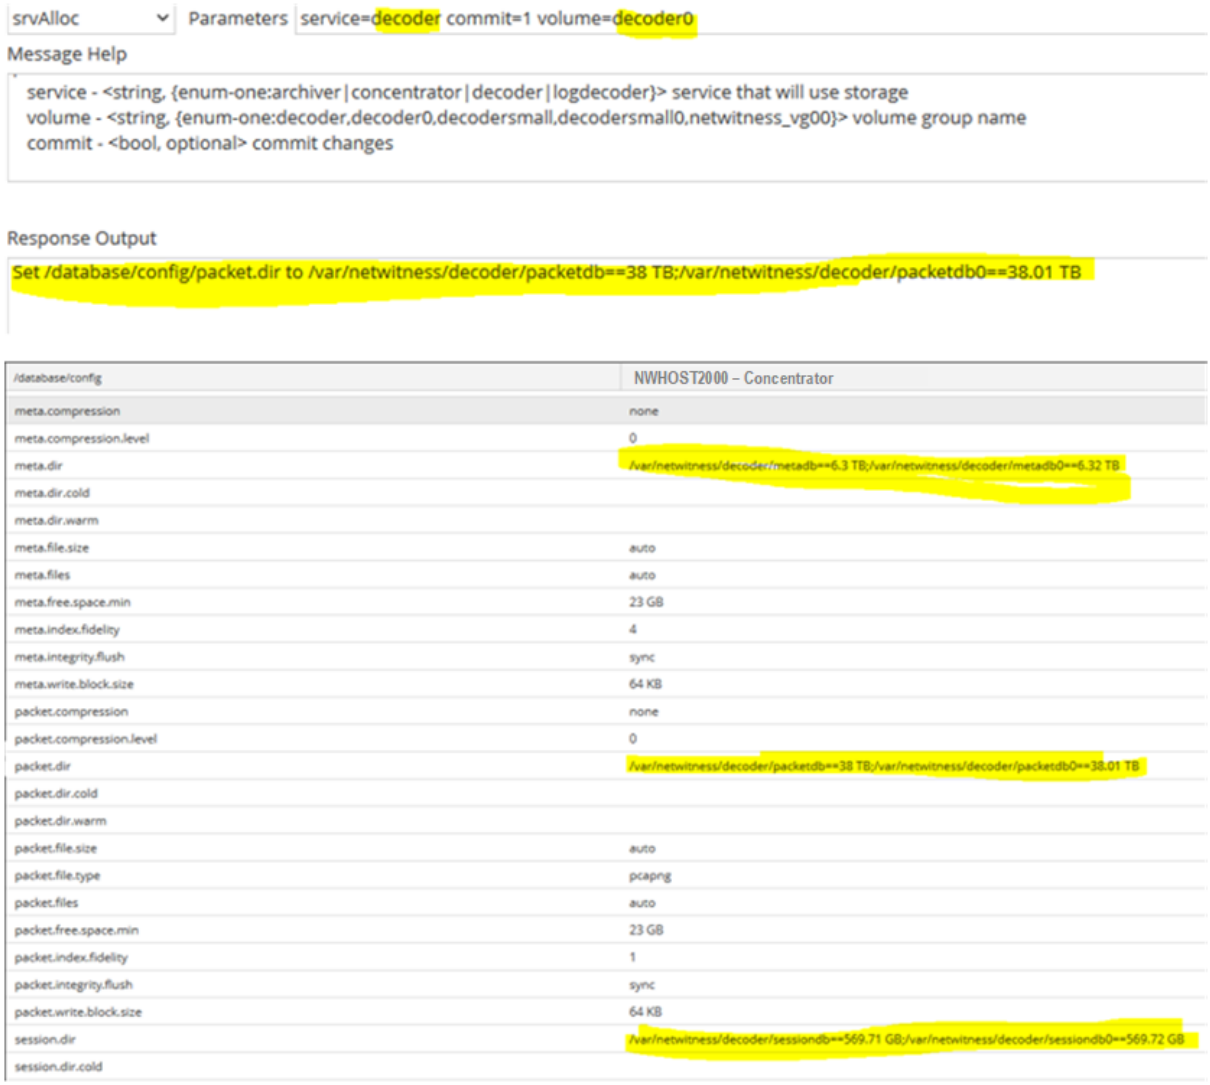
   `name=sdf service=decoder volume=decoder commit=1`



```
[root@NWHOST2000 ~]# df -h
Filesystem                          Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root     30G  2.5G   28G   9% /
devtmpfs                             63G     0   63G   0% /dev
tmpfs                                63G   12K   63G   1% /dev/shm
tmpfs                                63G   27M   63G   1% /run
tmpfs                                63G     0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-nwhome  2.7T   98M  2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog   10G   50M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome  10G   33M   10G   1% /home
/dev/sda1                          1014M   88M  927M   9% /boot
tmpfs                                13G     0   13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot    10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index       30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb  600G   33M  600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb     6.7T   33M  6.7T   1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb         41T   34M   41T   1% /var/netwitness/decoder/packetdb
/dev/mapper/decodersmall0-index      30G   33M   30G   1% /var/netwitness/decoder/index0
/dev/mapper/decodersmall0-sessiondb 600G   33M  600G   1% /var/netwitness/decoder/sessiondb0
/dev/mapper/decodersmall0-metadb    6.7T   33M  6.7T   1% /var/netwitness/decoder/metadb0
/dev/mapper/decoder0-packetdb        41T   34M   41T   1% /var/netwitness/decoder/packetdb0
```

```
[root@NWHOST2000 ~]# lsblk
NAME                        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                           8:0    0  931G  0 disk
├─sda1                        8:1    0    1G  0 part /boot
└─sda2                        8:2    0  930G  0 part
  ├─netwitness_vg00-root    253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap    253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome  253:2    0  2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog  253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome 253:4    0   10G  0 lvm  /home
sdb                           8:16   0  1.8T  0 disk
└─sdb1                        8:17   0  1.8T  0 part
  └─netwitness_vg00-nwhome  253:2    0  2.7T  0 lvm  /var/netwitness
sdc                           8:32   0  7.3T  0 disk
└─sdc1                        8:33   0  7.3T  0 part
  ├─decodersmall-decoroot   253:5    0   10G  0 lvm  /var/netwitness/decoder
  ├─decodersmall-index      253:6    0   30G  0 lvm  /var/netwitness/decoder/index
  ├─decodersmall-sessiondb  253:7    0  600G  0 lvm  /var/netwitness/decoder/sessiondb
  └─decodersmall-metadb     253:8    0  6.7T  0 lvm  /var/netwitness/decoder/metadb
sdd                           8:48   0   40T  0 disk
└─sdd1                        8:49   0   40T  0 part
  └─decoder-packetdb        253:9    0   40T  0 lvm  /var/netwitness/decoder/packetdb
sde                           8:64   0  7.3T  0 disk
└─sde1                        8:65   0  7.3T  0 part
  ├─decodersmall0-index     253:10   0   30G  0 lvm  /var/netwitness/decoder/index0
  ├─decodersmall0-sessiondb 253:11   0  600G  0 lvm  /var/netwitness/decoder/sessiondb0
  └─decodersmall0-metadb    253:12   0  6.7T  0 lvm  /var/netwitness/decoder/metadb0
sdf                           8:80   0   40T  0 disk
└─sdf1                        8:81   0   40T  0 part
  └─decoder0-packetdb       253:13   0   40T  0 lvm  /var/netwitness/decoder/packetdb0
```

9. Execute the `srvAlloc` command with the following parameters to add the storage information into the Service Configuration settings.
   - `service=decoder volume=decodersmall commit=1`
   - `service=decoder volume=decodersmall0 commit=1`
   - `service=decoder volume=decoder commit=1`
   - `service=decoder volume=decoder0 commit=1`



srvAlloc | Parameters | service=decoder commit=1 volume=decoder0

Message Help

service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
volume - <string, {enum-one:decoder,decoder0,decodersmall,decodersmall0,netwitness_vg00}> volume group name
commit - <bool, optional> commit changes

Response Output

Set /database/config/packet.dir to /var/netwitness/decoder/packetdb==38 TB;/var/netwitness/decoder/packetdb0==38.01 TB

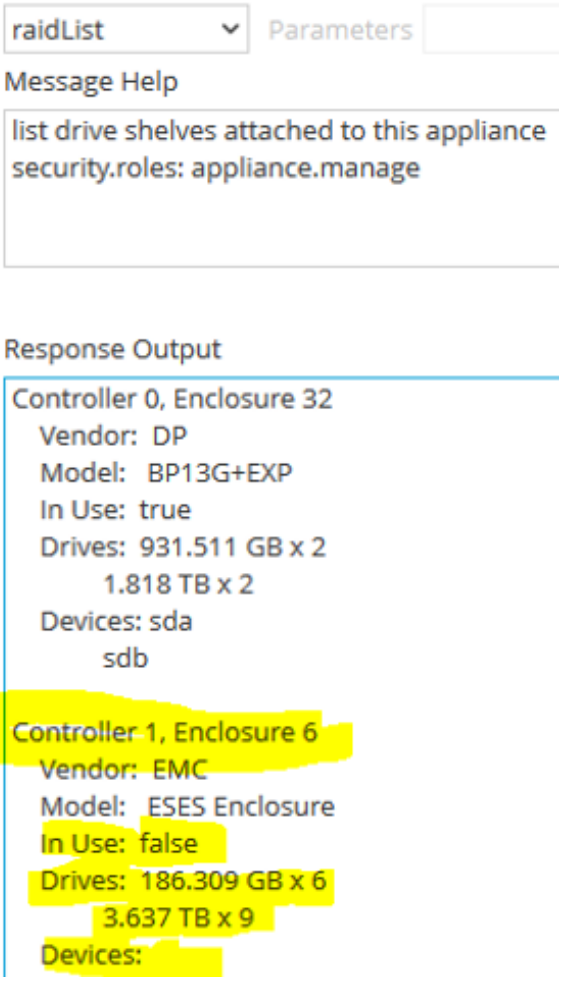
/database/config	NWHOST2000 – Concentrator
meta.compression	none
meta.compression.level	0
meta.dir	/var/netwitness/decoder/metadb==6.3 TB;/var/netwitness/decoder/metadb0==6.32 TB
meta.dir.cold	
meta.dir.warm	
meta.file.size	auto
meta.files	auto
meta.free.space.min	23 GB
meta.index.fidelity	4
meta.integrity.flush	sync
meta.write.block.size	64 KB
packet.compression	none
packet.compression.level	0
packet.dir	/var/netwitness/decoder/packetdb==38 TB;/var/netwitness/decoder/packetdb0==38.01 TB
packet.dir.cold	
packet.dir.warm	
packet.file.size	auto
packet.file.type	pcapng
packet.files	auto
packet.free.space.min	23 GB
packet.index.fidelity	1
packet.integrity.flush	sync
packet.write.block.size	64 KB
session.dir	/var/netwitness/decoder/sessiondb==569.71 GB;/var/netwitness/decoder/sessiondb0==569.72 GB
session.dir.cold	

10. Reconfigure the following Network Decoder service and its database to detect and take advantage of all of the free space as described in Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture.

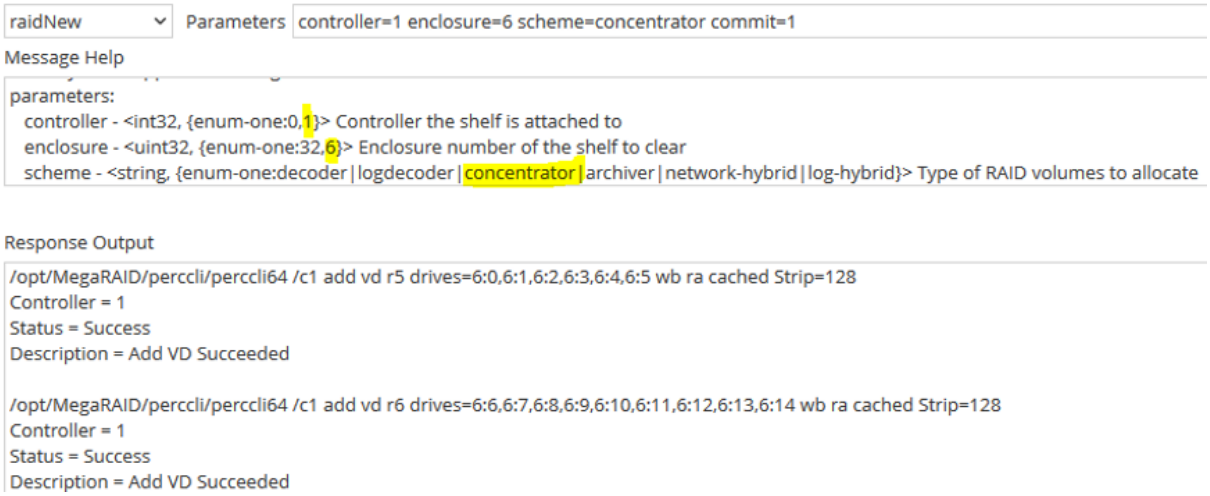
Configure Storage for Network Concentrator

The following scenario configures storage on one, non-encrypted, 15-Drive DAC for a Network Concentrator physical host.

1. Execute the `raidList` command.



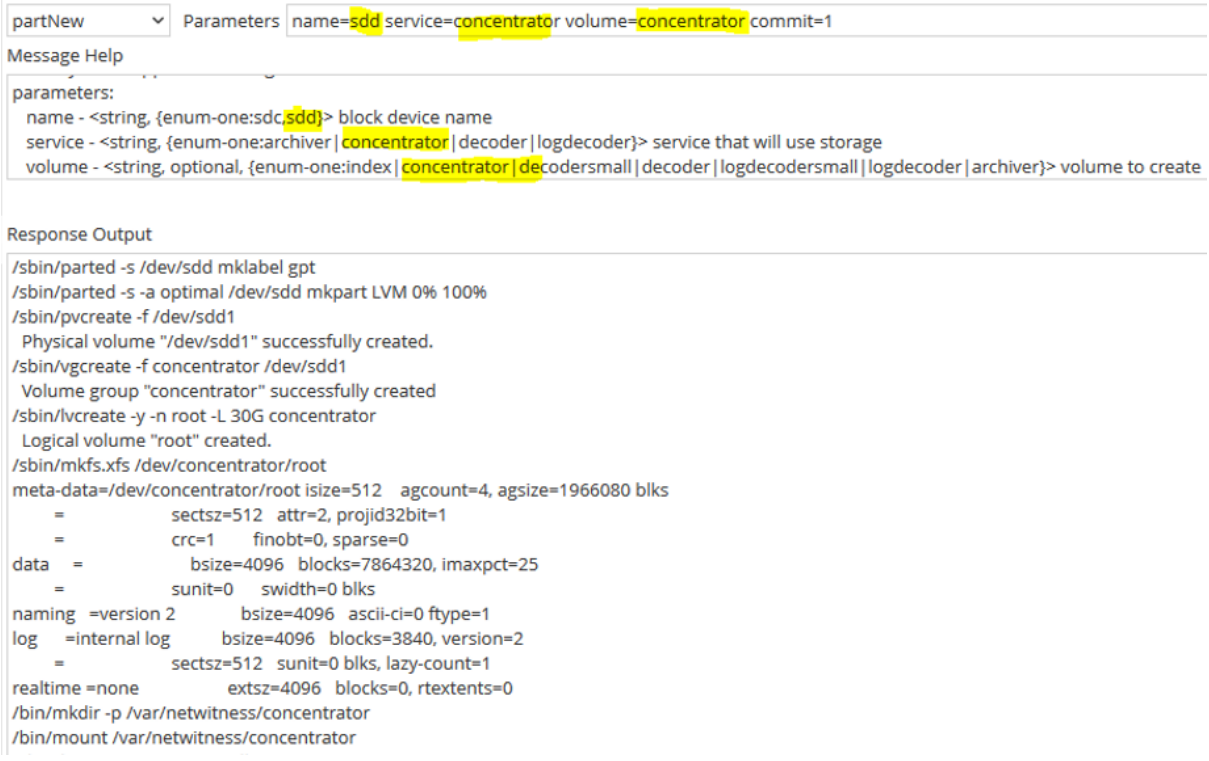
2. Execute the `raidNew` command with the following parameters.
   `controller=1 enclosure=6 scheme=concentrator`




```
[root@NWHOST1500 ~]# lsblk
NAME                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                             8:0    0   931G  0 disk
├─sda1                          8:1    0     1G  0 part /boot
└─sda2                          8:2    0   930G  0 part
  ├─netwitness_vg00-root      253:0    0    30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0     4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0   2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0    10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0    10G  0 lvm  /home
sdb                             8:16   0   1.8T  0 disk
└─sdb1                          8:17   0   1.8T  0 part
  └─netwitness_vg00-nwhome    253:2    0   2.7T  0 lvm  /var/netwitness
sdc                             8:32   0 928.8G  0 disk
sdd                             8:48   0  25.5T  0 disk
[root@NWHOST1500 ~]#
```

3. Execute the `partNew` command to create the **concentrator** partition first with the following parameters. You must create the **concentrator** volume before **index** volume or it will fail.

```
name=sdd service=concentrator volume=concentrator commit=1
```




```
[root@NWHOST1500 ~]# lsblk
NAME                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                             8:0    0   931G  0 disk
├─sda1                          8:1    0     1G  0 part /boot
└─sda2                          8:2    0   930G  0 part
  ├─netwitness_vg00-root      253:0    0    30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0     4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0   2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0    10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0    10G  0 lvm  /home
sdb                             8:16   0   1.8T  0 disk
└─sdb1                          8:17   0   1.8T  0 part
  └─netwitness_vg00-nwhome    253:2    0   2.7T  0 lvm  /var/netwitness
sdc                             8:32   0 928.8G  0 disk
sdd                             8:48   0  25.5T  0 disk
└─sdd1                          8:49   0  25.5T  0 part
  ├─concentrator-root         253:5    0    30G  0 lvm  /var/netwitness/concentrator
  ├─concentrator-sessiondb    253:6    0   600G  0 lvm  /var/netwitness/concentrator/sessiondb
  └─concentrator-metadb       253:7    0  24.9T  0 lvm  /var/netwitness/concentrator/metadb
```

4. Execute the partNew command with the following parameters with the following parameters to create an index on SSDs.

   ```
   name=sdc service=concentrator volume=index commit=1
   ```

partNew Parameters name=sdc service=concentrator volume=index commit=1

Message Help

parameters:
name - <string, {enum-one:sdc,sdd}> block device name
service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create

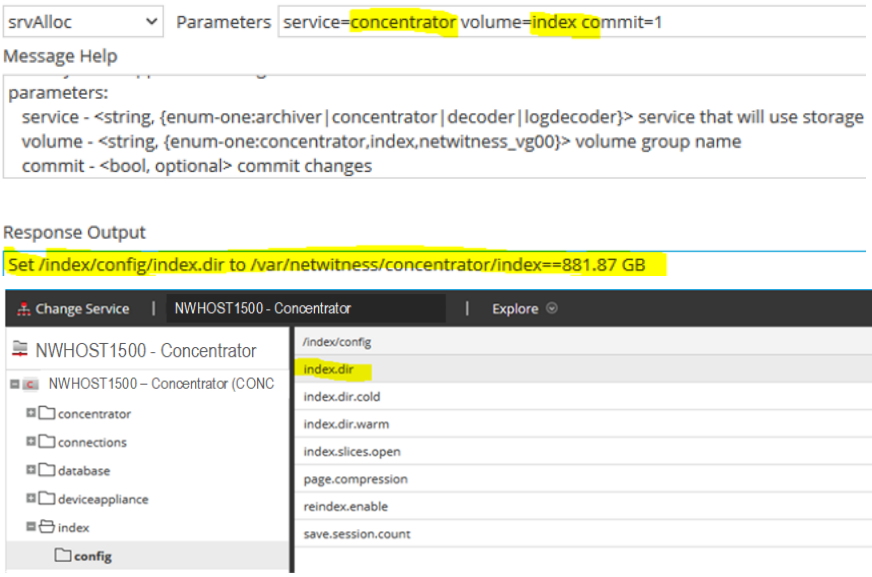
Response Output

```
/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
  Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f index /dev/sdc1
  Volume group "index" successfully created
/sbin/lvcreate -y -n index -l 100%FREE index
  Wiping xfs signature on /dev/index/index.
  Logical volume "index" created.
/sbin/mkfs.xfs /dev/index/index
meta-data=/dev/index/index      isize=512    agcount=4, agsize=60866304 blks
         =                      sectsz=4096  attr=2, projid32bit=1
         =                      crc=1        finobt=0, sparse=0
data     =                      bsize=4096   blocks=243465216, imaxpct=25
         =                      sunit=0      swidth=0 blks
naming   =version 2             bsize=4096   ascii-ci=0 ftype=1
log      =internal log          bsize=4096   blocks=118879, version=2
         =                      sectsz=4096  sunit=1 blks, lazy-count=1
realtime =none                  extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/index
```

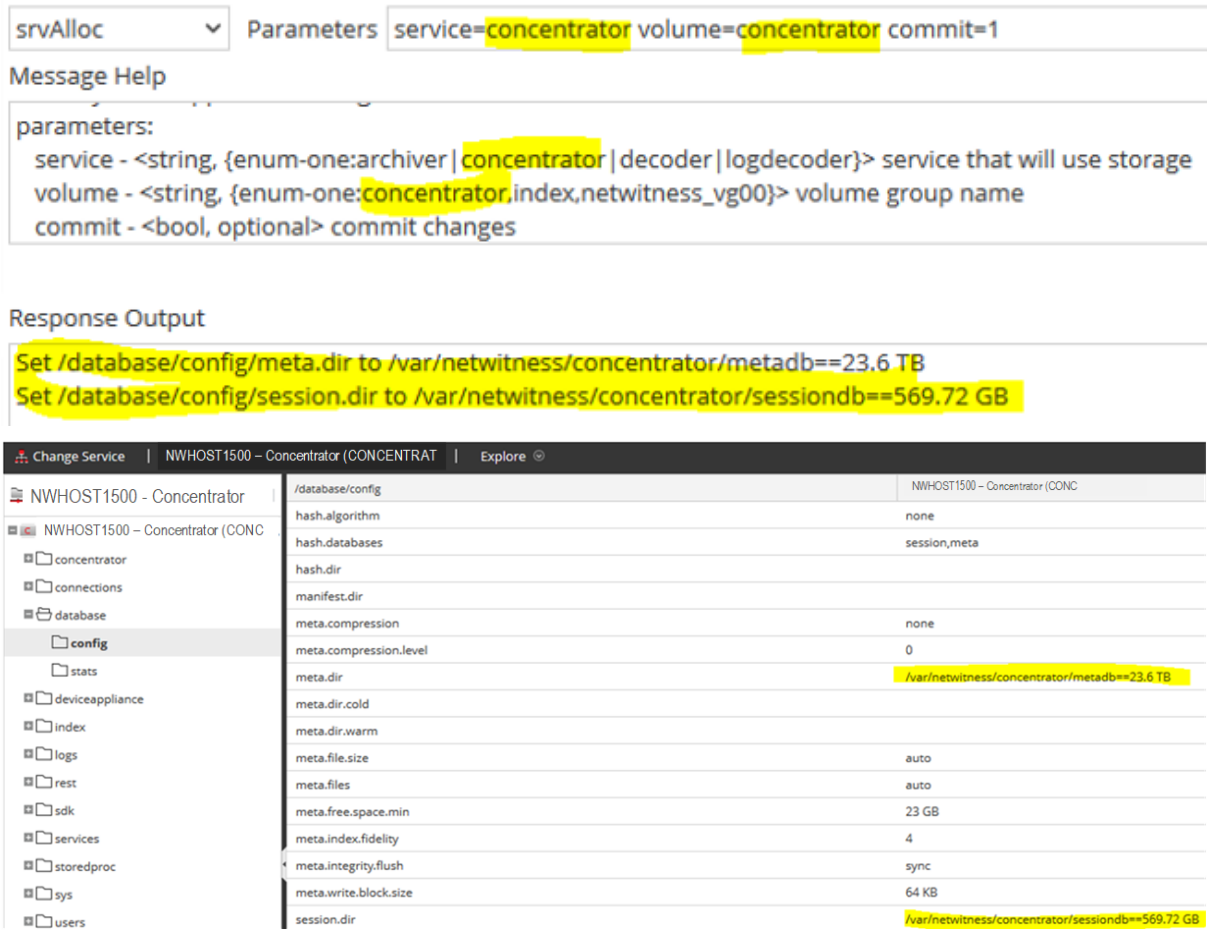
```
[root@NWHOST1500 ~]# lsblk
NAME                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                             8:0    0   931G  0 disk
├─sda1                          8:1    0     1G  0 part /boot
└─sda2                          8:2    0   930G  0 part
  ├─netwitness_vg00-root      253:0    0    30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0     4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0   2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0    10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0    10G  0 lvm  /home
sdb                             8:16   0   1.8T  0 disk
└─sdb1                          8:17   0   1.8T  0 part
  └─netwitness_vg00-nwhome    253:2    0   2.7T  0 lvm  /var/netwitness
sdc                             8:32   0 928.8G  0 disk
└─sdc1                          8:33   0 928.8G  0 part
  └─index-index               253:8    0 928.8G  0 lvm  /var/netwitness/concentrator/index
sdd                             8:48   0  25.5T  0 disk
└─sdd1                          8:49   0  25.5T  0 part
  ├─concentrator-root         253:5    0    30G  0 lvm  /var/netwitness/concentrator
  ├─concentrator-sessiondb    253:6    0   600G  0 lvm  /var/netwitness/concentrator/sessiondb
  └─concentrator-metadb       253:7    0  24.9T  0 lvm  /var/netwitness/concentrator/metadb
```

```
[root@NWHOST1500 ~]# df -h
Filesystem                              Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root         30G  2.1G   28G   7% /
devtmpfs                                 63G     0   63G   0% /dev
tmpfs                                    63G   12K   63G   1% /dev/shm
tmpfs                                    63G   10M   63G   1% /run
tmpfs                                    63G     0   63G   0% /sys/fs/cgroup
/dev/sda1                              1014M   91M  924M   9% /boot
/dev/mapper/netwitness_vg00-varlog       10G   52M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome      10G   33M   10G   1% /home
/dev/mapper/netwitness_vg00-nwhome      2.7T   98M  2.7T   1% /var/netwitness
tmpfs                                    13G     0   13G   0% /run/user/0
/dev/mapper/concentrator-root            30G   33M   30G   1% /var/netwitness/concentrator
/dev/mapper/concentrator-sessiondb      600G   33M  600G   1% /var/netwitness/concentrator/sessiondb
/dev/mapper/concentrator-metadb          25T   33M   25T   1% /var/netwitness/concentrator/metadb
/dev/mapper/index-index                 929G   33M  929G   1% /var/netwitness/concentrator/index
```

5. Execute the `srvAlloc` command with the following parameters.
   `service=concentrator volume=index commit=1`



6. Execute the `srvAlloc` command with the following parameters.
   `service=concentrator volume=concentrator commit=1`

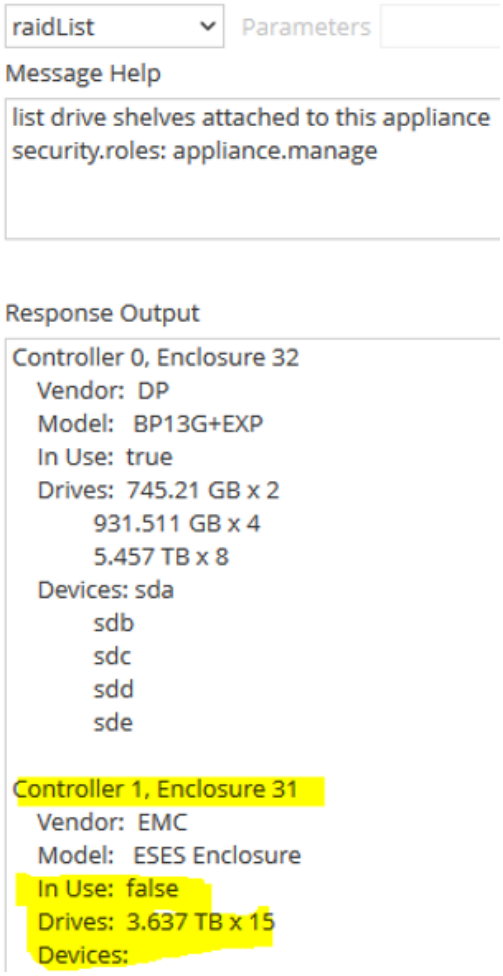


7. Reconfigure the following Network Concentrator service and its database to detect and take advantage of all of the free space as described in Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture.

Configure Storage for Log Decoder Hybrid

The following scenario configures storage on one, non-encrypted, 15-Drive DAC for a Log Decoder Hybrid physical host.

1. Execute the `raidList` command.



2. Execute the `raidNew` command with the following parameters.

```
controller=1 enclosure=31 scheme=log-hybrid commit=1
```

raidNew Parameters controller=1 enclosure=31 scheme=log-hybrid commit=1

Message Help

```
controller - <int32, {enum-one:0,1}> Controller the shelf is attached to
enclosure - <uint32, {enum-one:32,31}> Enclosure number of the shelf to clear
scheme - <string, {enum-one:decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid}> Type of RAID volumes to allocate
preferSecure - <bool, optional, {bool:0,1,yes,no,true,false,on,off}> Prefer creation of a secure array given compatible physical drives and a controller with a security key set
```

Response Output

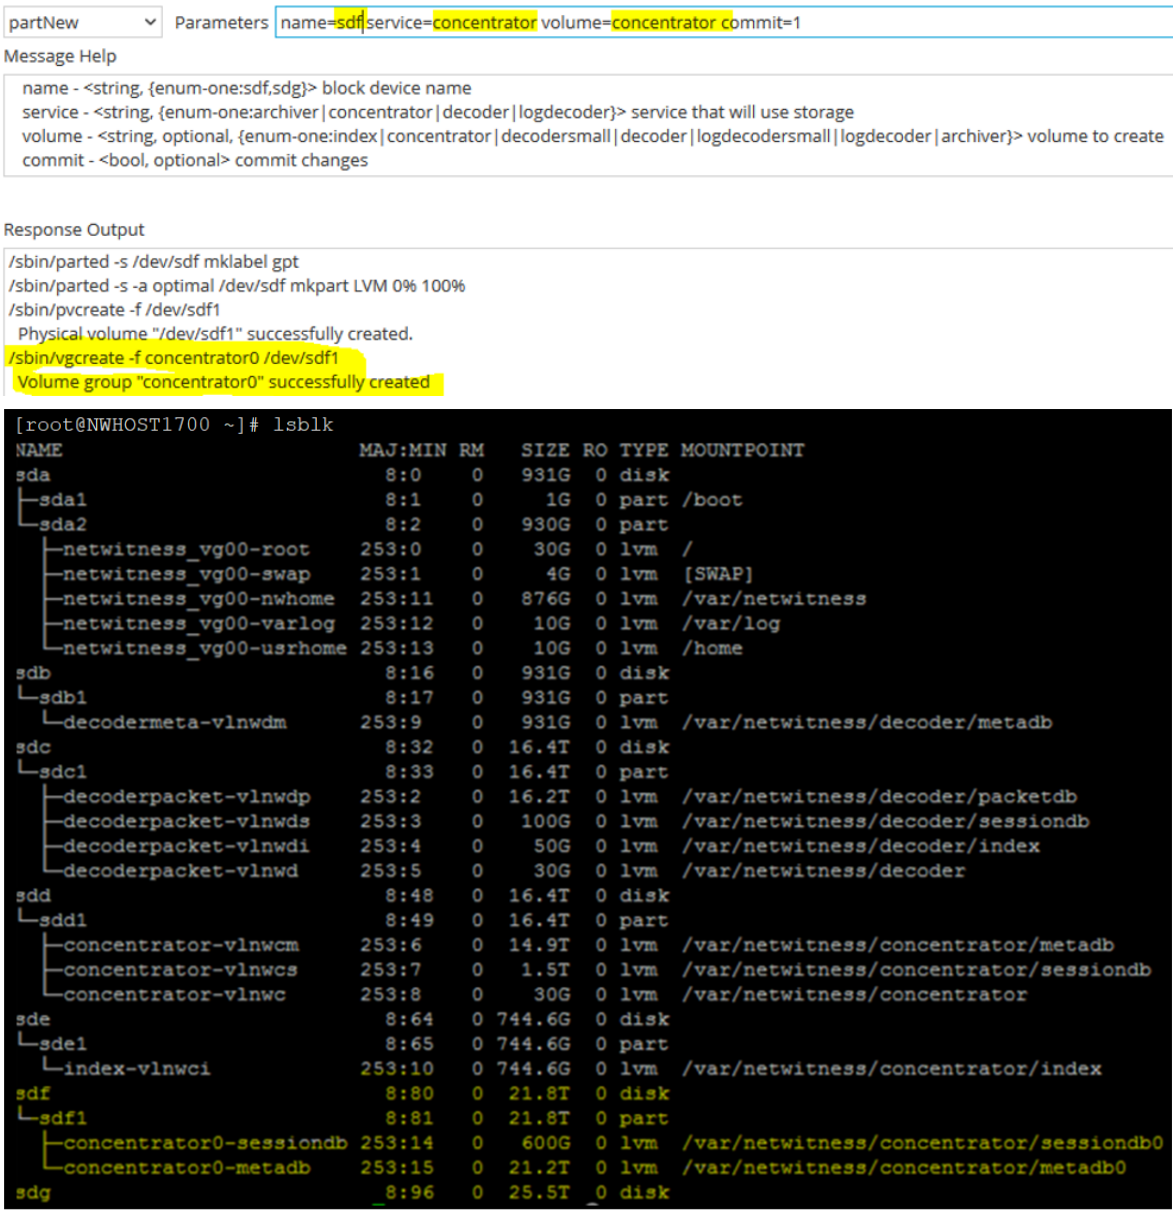
```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=31:0,31:1,31:2,31:3,31:4,31:5,31:6 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded

/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=31:7,31:8,31:9,31:10,31:11,31:12,31:13,31:14 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded
```

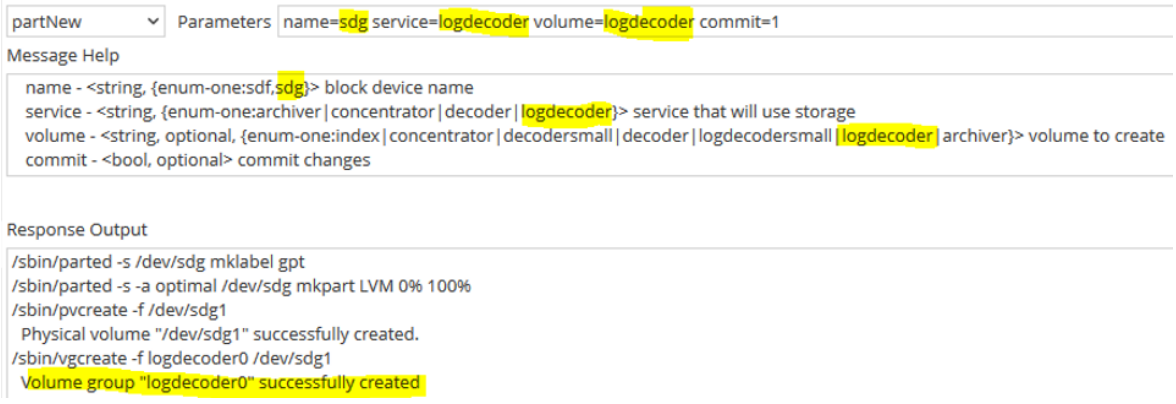
```
[root@NWHOST1700 ~]# lsblk
NAME                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                             8:0    0   931G  0 disk
├─sda1                          8:1    0     1G  0 part /boot
└─sda2                          8:2    0   930G  0 part
  ├─netwitness_vg00-root      253:0    0    30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0     4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:11   0   876G  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:12   0    10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:13   0    10G  0 lvm  /home
sdb                             8:16   0   931G  0 disk
└─sdb1                          8:17   0   931G  0 part
  └─decodermeta-vlnwdm        253:9    0   931G  0 lvm  /var/netwitness/decoder/metadb
sdc                             8:32   0  16.4T  0 disk
└─sdc1                          8:33   0  16.4T  0 part
  ├─decoderpacket-vlnwdp      253:2    0  16.2T  0 lvm  /var/netwitness/decoder/packetdb
  ├─decoderpacket-vlnwds      253:3    0   100G  0 lvm  /var/netwitness/decoder/sessiondb
  ├─decoderpacket-vlnwdi      253:4    0    50G  0 lvm  /var/netwitness/decoder/index
  └─decoderpacket-vlnwd       253:5    0    30G  0 lvm  /var/netwitness/decoder
sdd                             8:48   0  16.4T  0 disk
└─sdd1                          8:49   0  16.4T  0 part
  ├─concentrator-vlnwcm       253:6    0  14.9T  0 lvm  /var/netwitness/concentrator/metadb
  ├─concentrator-vlnwcs       253:7    0   1.5T  0 lvm  /var/netwitness/concentrator/sessiondb
  └─concentrator-vlnwc        253:8    0    30G  0 lvm  /var/netwitness/concentrator
sde                             8:64   0 744.6G  0 disk
└─sde1                          8:65   0 744.6G  0 part
  └─index-vlnwci              253:10   0 744.6G  0 lvm  /var/netwitness/concentrator/index
sdf                             8:80   0  21.8T  0 disk
sdg                             8:96   0  25.5T  0 disk
```

3. Execute the `partNew` command with the following parameters with the following parameters.

   - `name=sdf service=concentrator volume=concentrator commit=1`



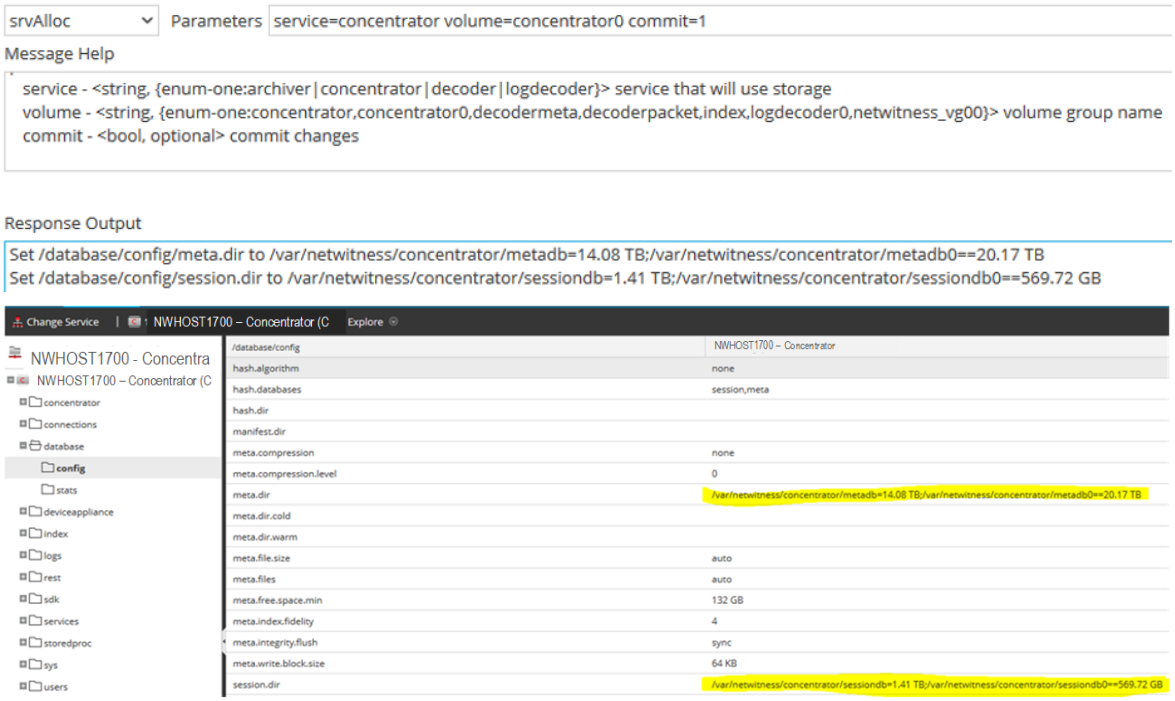
- `name=sdg service=logdecoder volume=logdecoder commit=1`




```
[root@NWHOST1700 ~]# lsblk
NAME                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                             8:0    0   931G  0 disk
├─sda1                          8:1    0     1G  0 part /boot
└─sda2                          8:2    0   930G  0 part
  ├─netwitness_vg00-root      253:0    0    30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0     4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:11   0   876G  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:12   0    10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:13   0    10G  0 lvm  /home
sdb                             8:16   0   931G  0 disk
└─sdb1                          8:17   0   931G  0 part
  └─decodermeta-vlnwdm        253:9    0   931G  0 lvm  /var/netwitness/decoder/metadb
sdc                             8:32   0  16.4T  0 disk
└─sdc1                          8:33   0  16.4T  0 part
  ├─decoderpacket-vlnwdp      253:2    0  16.2T  0 lvm  /var/netwitness/decoder/packetdb
  ├─decoderpacket-vlnwds      253:3    0   100G  0 lvm  /var/netwitness/decoder/sessiondb
  ├─decoderpacket-vlnwdi      253:4    0    50G  0 lvm  /var/netwitness/decoder/index
  └─decoderpacket-vlnwd       253:5    0    30G  0 lvm  /var/netwitness/decoder
sdd                             8:48   0  16.4T  0 disk
└─sdd1                          8:49   0  16.4T  0 part
  ├─concentrator-vlnwcm       253:6    0  14.9T  0 lvm  /var/netwitness/concentrator/metadb
  ├─concentrator-vlnwcs       253:7    0   1.5T  0 lvm  /var/netwitness/concentrator/sessiondb
  └─concentrator-vlnwc        253:8    0    30G  0 lvm  /var/netwitness/concentrator
sde                             8:64   0 744.6G  0 disk
└─sde1                          8:65   0 744.6G  0 part
  └─index-vlnwci              253:10   0 744.6G  0 lvm  /var/netwitness/concentrator/index
sdf                             8:80   0  21.8T  0 disk
└─sdf1                          8:81   0  21.8T  0 part
  ├─concentrator0-sessiondb   253:14   0   600G  0 lvm  /var/netwitness/concentrator/sessiondb0
  └─concentrator0-metadb      253:15   0  21.2T  0 lvm  /var/netwitness/concentrator/metadb0
sdg                             8:96   0  25.5T  0 disk
└─sdg1                          8:97   0  25.5T  0 part
  └─logdecoder0-packetdb      253:16   0  25.5T  0 lvm  /var/netwitness/decoder/packetdb0
```

4. Execute the `srvAlloc` command with the following parameters.

   - `service=concentrator volume=concentrator0 commit=1`



   - `service=logdecoder volume=logdecoder0 commit=1`



5. Reconfigure the following Log Decoder service and its database to detect and take advantage of all of the free space as described in Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture.

Appendix E. Sample Storage Configuration Scenarios for 8 or 12-Drive PowerVault

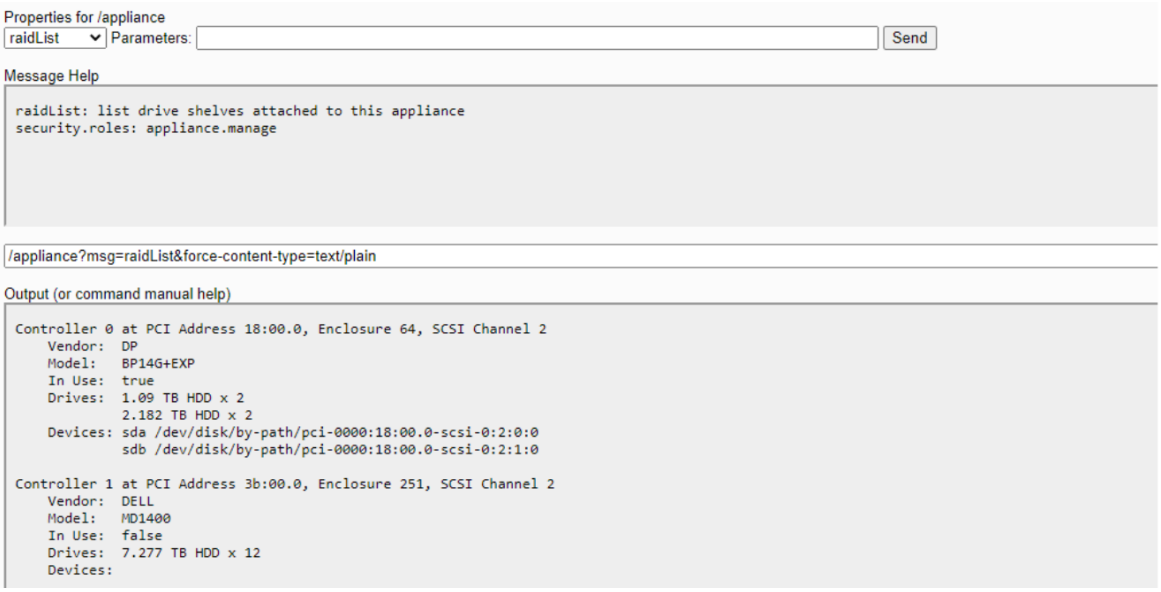
This appendix illustrates the following example of how to configure storage on one non-encrypted 8 or 12-drive PowerVault external storage devices.

- Configure Storage for Archiver using NW-PV-A/NW-PV-A-N
- Configure Storage for Decoder using NW-PV-B/NW-PV-B-N
- Configure Storage for Concentrator using NW-PV-C/NW-PV-C-N
- Configure Storage for Concentrator using NW-PV-D/NW-PV-D-N

Configure Storage for Archiver using NW-PV-A/NW-PV-A-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for an Archiver physical host.

1. Execute the `raidList` command.
   a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
   You should see the following information.
   ```
   In Use: FALSE
   Devices: <empty>
   ```
   b. Verify the Drive Count, Size, and Vendor.
   The following example illustrates what you should see before you create a RAID array.

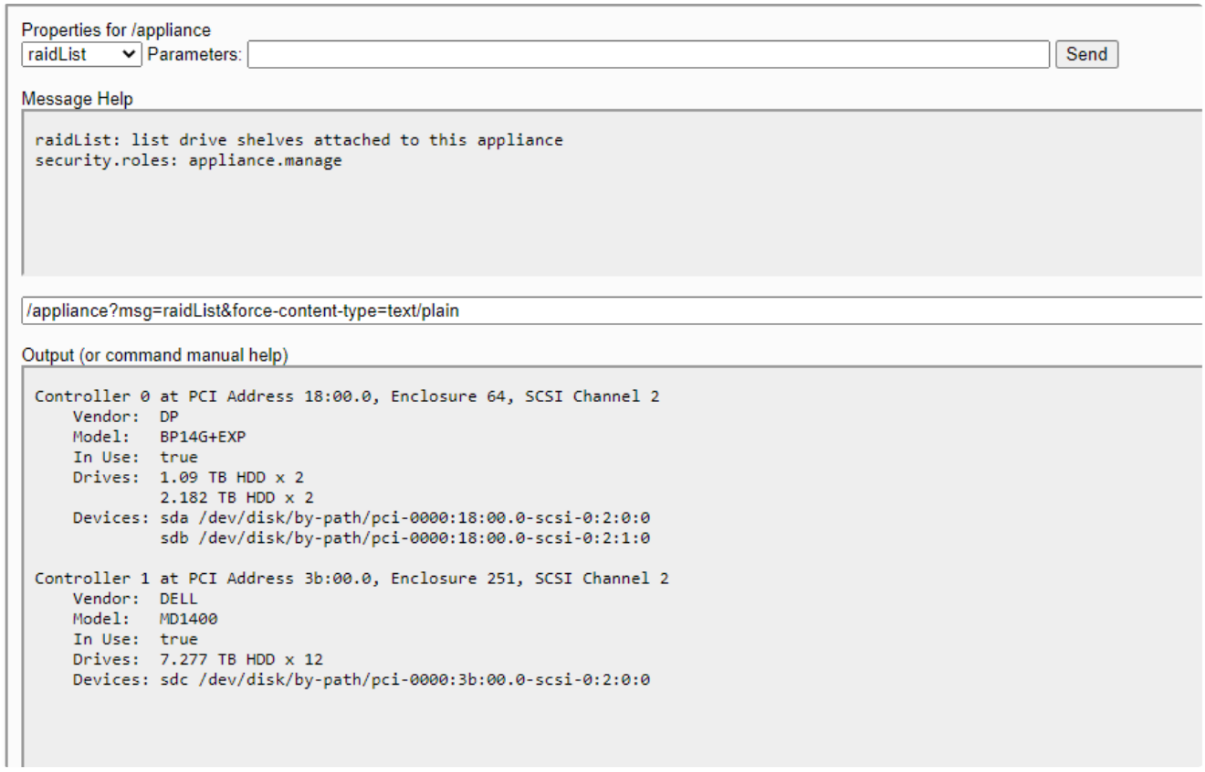


2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.
   ```
   controller=1 enclosure=251 scheme=archiver commit=1
   ```

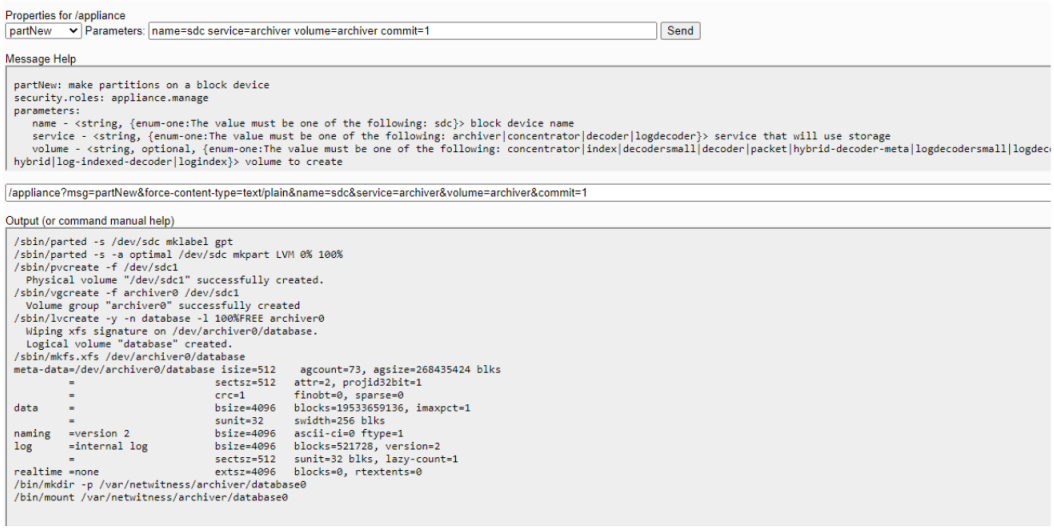
The following example illustrates what you should see after you create a RAID array.



3. Execute the `raidList` command to verify the new RAID array.
   You should now see the following information.
   `In Use: TRUE`
   `Devices: <device>` (for example, `sdc`)



4. Execute the `partNew` command with the following parameters to create partitions and mount points in the `etc/fstab` file.
   `name=sdc service=archiver volume=archiver commit=1`



5. Execute the `srvAlloc` command with the following parameters to allocate the space to the archiver service. This adds storage to the archiver service configuration and restarts the service every time it is executed.

   ```
   service=archiver volume=archiver0 commit=1
   ```

Configure Storage for Decoder using NW-PV-B/NW-PV-B-N

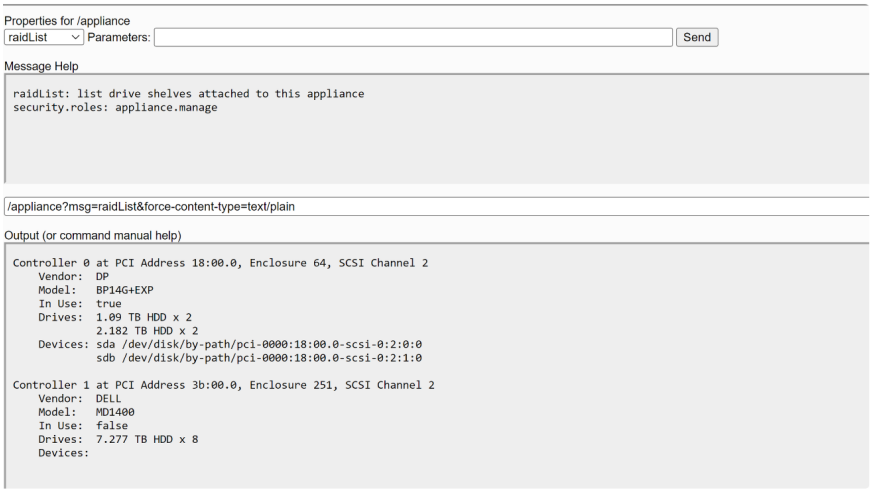
The following scenario configures storage on one, non-encrypted, 8-Drive PowerVault for a Network Decoder physical host.

1. Execute the `raidList` command.

   a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
   You should see the following information.

   ```
   In Use: FALSE
   Devices: <empty>
   ```

b. Verify the Drive Count, Size, and Vendor.
The following example illustrates what you should see before you create a RAID array.



2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded. Parameters for the first enclosure:
`controller=1 enclosure=0 scheme=decoder commit=1`

Scheme	Enclosure Type	Drives Required	Allocation
decoder	External	8 or 12 or 15 HDDs	3x drives in RAID 5 for decodersmall, all remaining drives in RAID 5 for decoder
log decoder	External	8 or 12 or 15 HDDs	Same as decoder configuration
decoder hot spare	External	8 or 12 or 15 HDDs	2x drives in RAID 1 for decodersmall, 1 drive as hot spare, all remaining drives in RAID 5 for decoder
log decoder hot spare	External	8 or 12 or 15 HDDs	Same as decoder hot spare configuration
archiver	External	8 or 12 or 15 HDDs	All drives in RAID 6 for archiver database volume
packet expansion	External	8 or 12 or 15 HDDs	All drives in RAID 6 for decoder volume. No drives allocated for decodersmall.
network hybrid	External	8 or 12 or 15 HDDs	3x drives in RAID 5 for meta expansion, all remaining drives in RAID 5 for packet expansion
network hybrid	Internal	S5 / S6 hybrid drive set	2x small HDD RAID 1 for decoder meta, 5x large HDD decoder, 3x large HDD concentrator, 2x SSD index
log hybrid	External	8 or 12 or 15 HDDs	Half of the drives in RAID 5 for meta expansion, half the drives in RAID 5 for packet expansion
log hybrid	Internal	S5 / S6 hybrid drive set	2x small HDD RAID 1 for decoder meta, 4x large HDD decoder, 4x large HDD concentrator, 2x SSD index
endpoint hybrid	Internal	S5 / S6 hybrid drive set	2x small HDD RAID 1 for decoder meta, 4x large HDD RAID 10 for log decoder and endpoint, 4x large HDD RAID 5 for concentrator, 2x SSD index
log indexed decoder	Internal	S6E hybrid drive set	10x HDD RAID 6 for log decoder meta and packet, 2x SSD index

Properties for /appliance

raidNew Parameters: controller=1 enclosure=251 scheme=decoder commit=1 Send

Message Help

```
raidNew: allocate RAID devices in a drive shelf
security.roles: appliance.manage
parameters:
   controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to
   enclosure - <uint32, optional, {enum-one:The value must be one of the following: 64|251}> Enclosure number of the shelf to clear.  Required i
   scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-ir
expansion|decoder-hotspare|logdecoder-hotspare}> Type of RAID volumes to allocate
```

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=251&scheme=decoder&commit=1

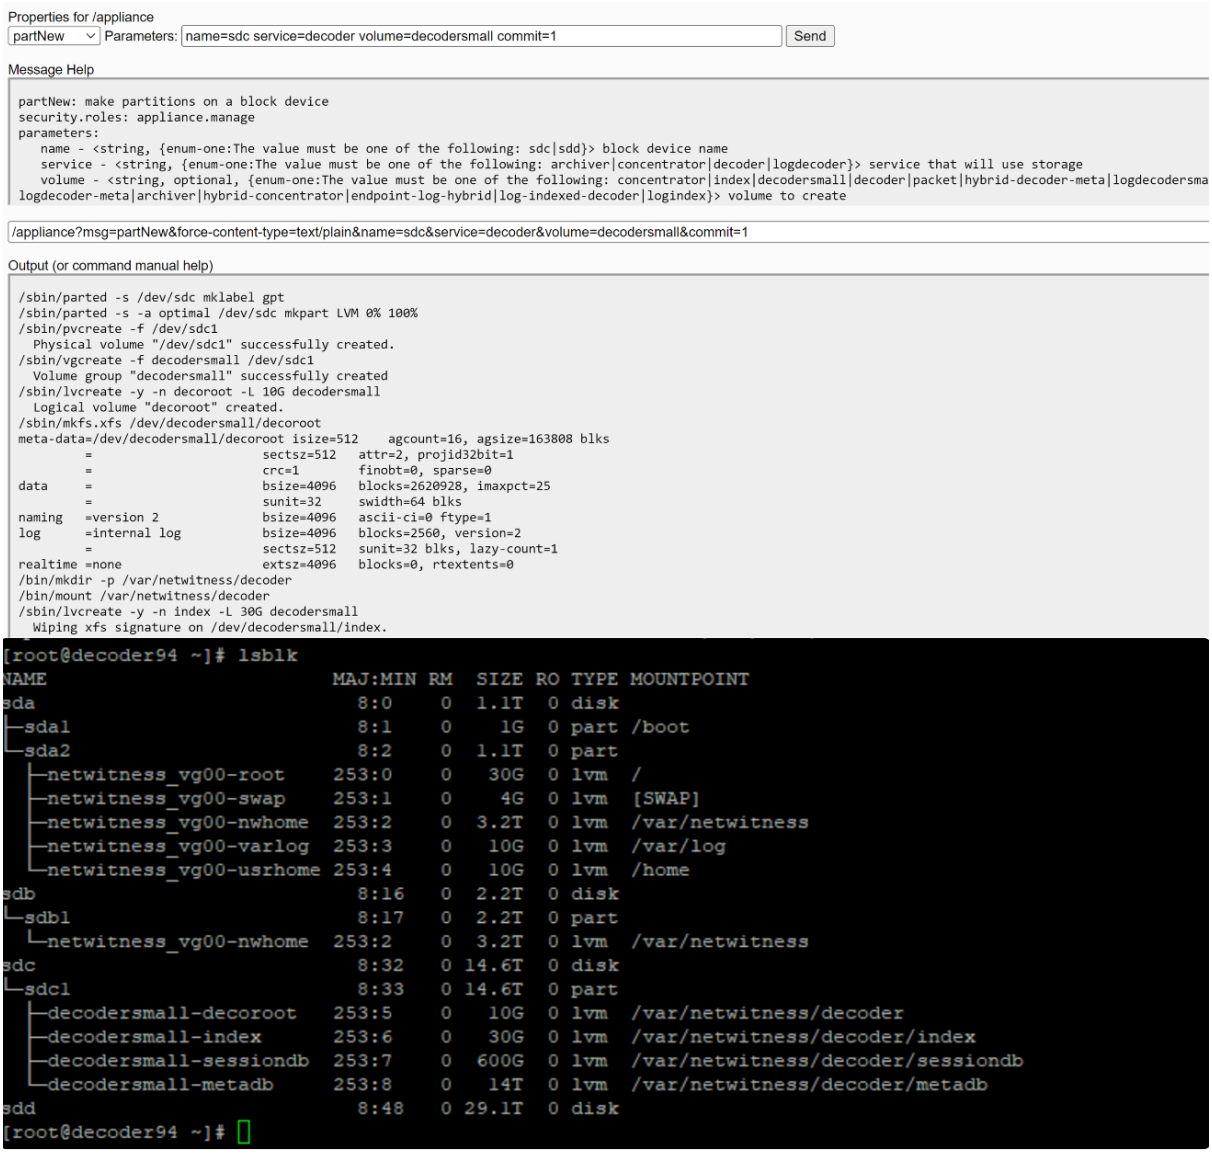
Output (or command manual help)

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=251:0,251:1,251:2 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.66.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.


/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=251:3,251:4,251:5,251:6,251:7 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.66.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.
```

```
[root@decoder94 ~]# lsblk
NAME                          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                             8:0    0  1.1T  0 disk
├─sda1                          8:1    0    1G  0 part /boot
└─sda2                          8:2    0  1.1T  0 part
  ├─netwitness_vg00-root      253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0   10G  0 lvm  /home
sdb                             8:16   0  2.2T  0 disk
└─sdb1                          8:17   0  2.2T  0 part
  └─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
sdc                             8:32   0 14.6T  0 disk
sdd                             8:48   0 29.1T  0 disk
[root@decoder94 ~]#
```

3. Execute the `partNew` command to create the **decodersmall** partition first (decoder dir, index, metadb, sessiondb) (First Enclosure, SDC, SDD) with the following parameters.
   `name=sdc service=decoder volume=decodersmall commit=1`



4. Execute the `partNew` command to create the decoder volume (packetdb) (First Enclosure, SDC, SDD) with the following parameters.
   `name==sdd service=decoder volume=decoder commit=1`

Properties for /appliance

partNew Parameters: name=sdd service=decoder volume=decoder commit=1 Send

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
   name - <string, {enum-one:The value must be one of the following: sdc|sdd}> block device name
   service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
   volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-m
logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create
```

/appliance?msg=partNew&force-content-type=text/plain&name=sdd&service=decoder&volume=decoder&commit=1

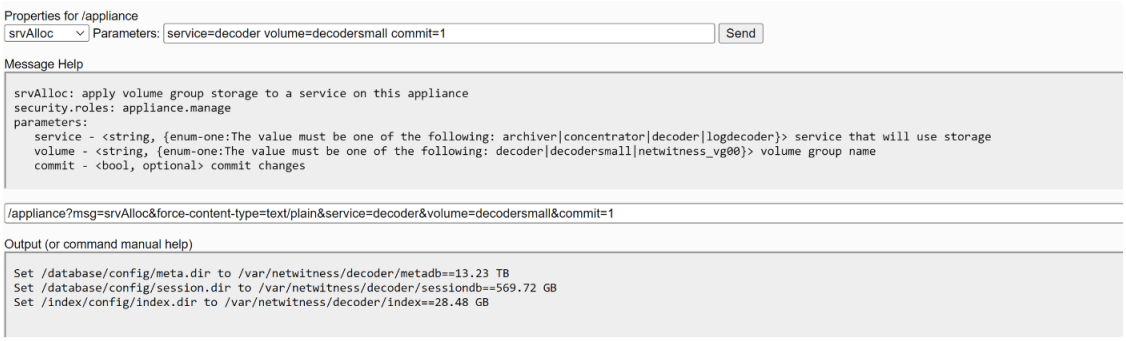
Output (or command manual help)

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
  Wiping xfs signature on /dev/sdd1.
  Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f decoder /dev/sdd1
  Volume group "decoder" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder
  Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder/packetdb
meta-data=/dev/decoder/packetdb  isize=512    agcount=32, agsize=244170720 blks
         =                       sectsz=512   attr=2, projid32bit=1
         =                       crc=1        finobt=0, sparse=0
data     =                       bsize=4096   blocks=7813463040, imaxpct=5
         =                       sunit=32     swidth=128 blks
naming   =version 2              bsize=4096   ascii-ci=0 ftype=1
log      =internal log           bsize=4096   blocks=521728, version=2
         =                       sectsz=512   sunit=32 blks, lazy-count=1
realtime =none                   extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/packetdb
/bin/mount /var/netwitness/decoder/packetdb
```

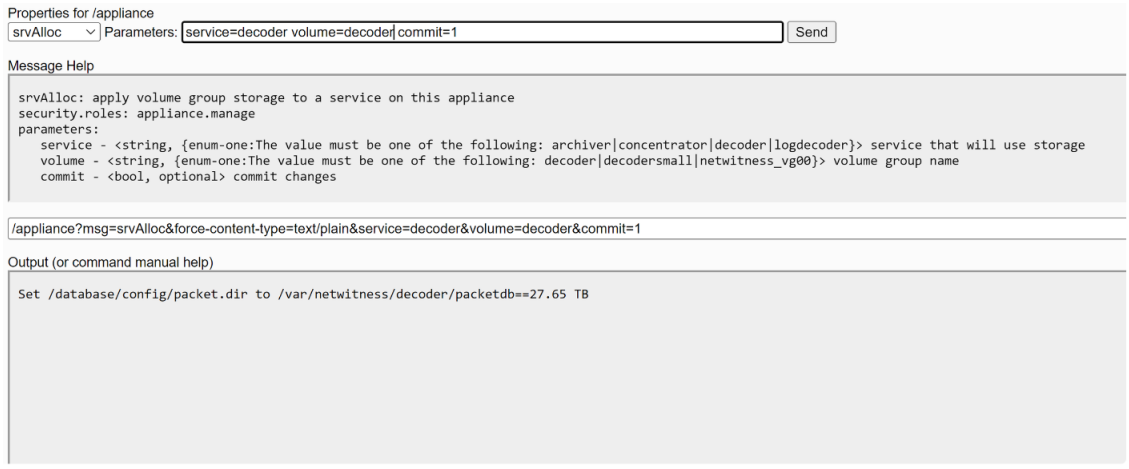
```
[root@decoder94 ~]# lsblk
NAME                          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                             8:0    0  1.1T  0 disk
├─sda1                          8:1    0    1G  0 part /boot
└─sda2                          8:2    0  1.1T  0 part
  ├─netwitness_vg00-root      253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0   10G  0 lvm  /home
sdb                             8:16   0  2.2T  0 disk
└─sdb1                          8:17   0  2.2T  0 part
  └─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
sdc                             8:32   0 14.6T  0 disk
└─sdc1                          8:33   0 14.6T  0 part
  ├─decodersmall-decoroot     253:5    0   10G  0 lvm  /var/netwitness/decoder
  ├─decodersmall-index        253:6    0   30G  0 lvm  /var/netwitness/decoder/index
  ├─decodersmall-sessiondb    253:7    0  600G  0 lvm  /var/netwitness/decoder/sessiondb
  └─decodersmall-metadb       253:8    0   14T  0 lvm  /var/netwitness/decoder/metadb
sdd                             8:48   0 29.1T  0 disk
└─sdd1                          8:49   0 29.1T  0 part
  └─decoder-packetdb          253:9    0 29.1T  0 lvm  /var/netwitness/decoder/packetdb
[root@decoder94 ~]#
```

```
[root@decoder94 ~]# df -kh
Filesystem                            Size  Used Avail Use% Mounted on
devtmpfs                               63G     0   63G   0% /dev
tmpfs                                  63G   40K   63G   1% /dev/shm
tmpfs                                  63G   11M   63G   1% /run
tmpfs                                  63G     0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root       30G  3.7G   27G  13% /
/dev/sda1                            1014M  122M  893M  12% /boot
/dev/mapper/netwitness_vg00-nwhome    3.3T  543M  3.3T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog     10G  138M  9.9G   2% /var/log
/dev/mapper/netwitness_vg00-usrhome    10G   33M   10G   1% /home
tmpfs                                  13G     0   13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot      10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index         30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb    600G   34M  600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb        14T   34M   14T   1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb           30T   34M   30T   1% /var/netwitness/decoder/packetdb
[root@decoder94 ~]#
```

5. Execute the `srvAlloc` command with the following parameters to add the storage information into the Service Configuration settings.
   - `service=decoder volume=decodersmall commit=1`



   - `service=decoder volume=decoder commit=1`



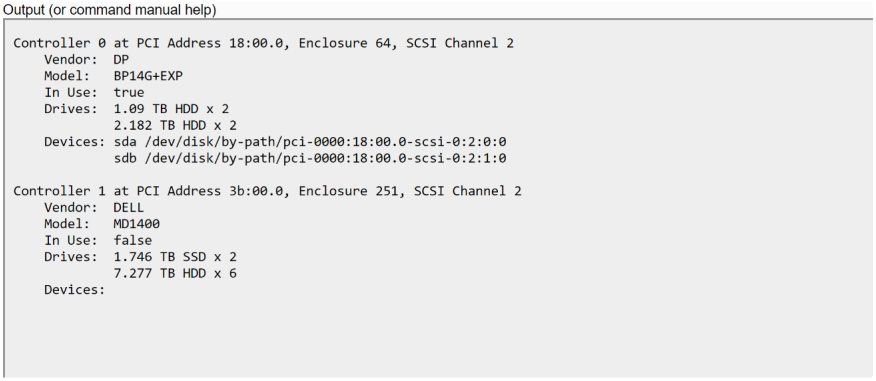
Configure Storage for Concentrator using NW-PV-C/NW-PV-C-N

The following scenario configures storage on one, non-encrypted, 8-Drive PowerVault for a Network Concentrator physical host.

1. Execute the `raidList` command.

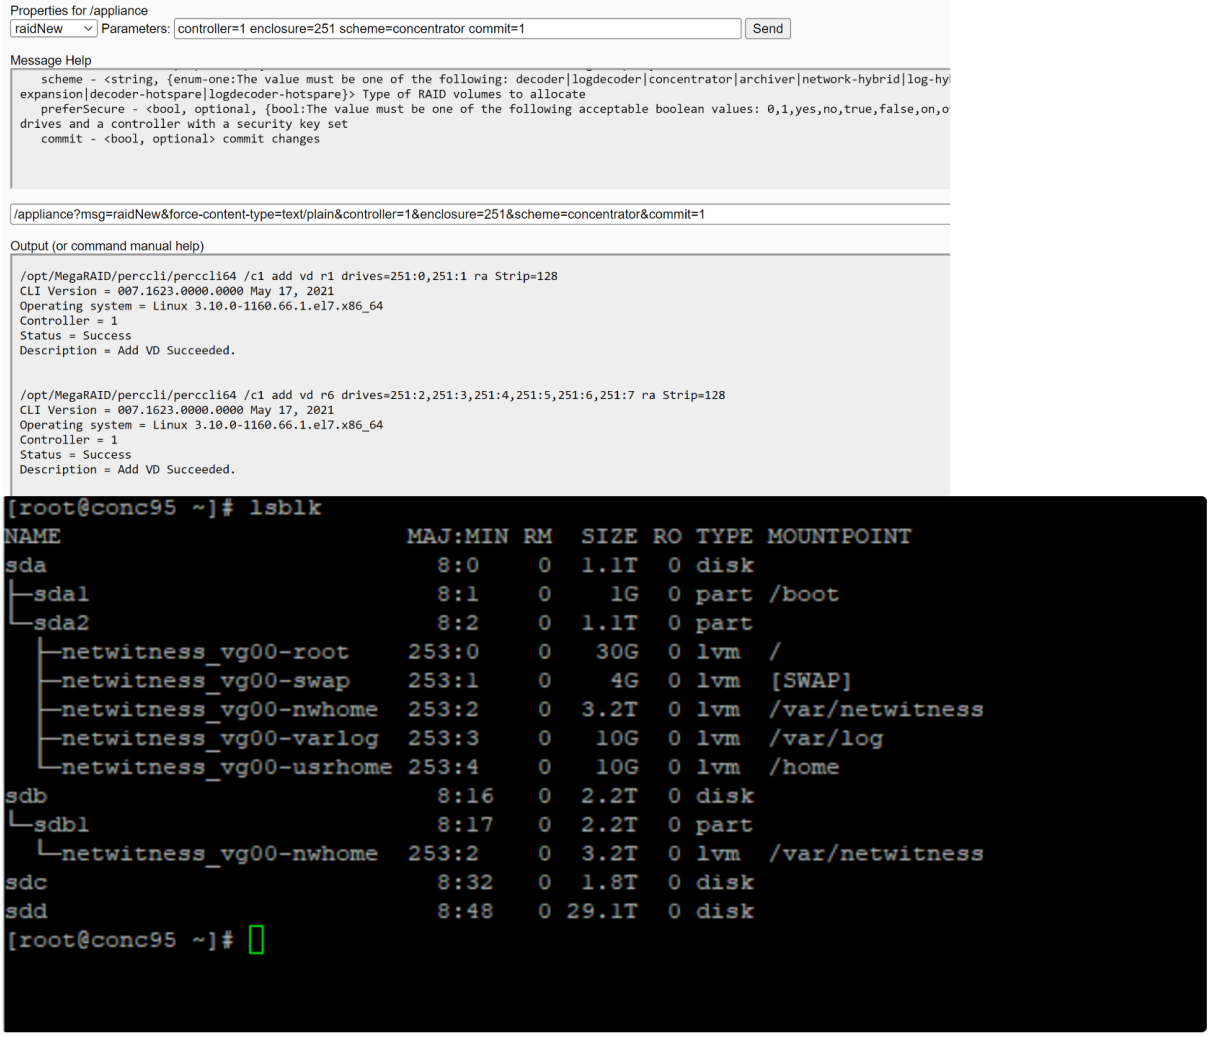
   a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
   You should see the following information.
   `In Use: FALSE`
   `Devices: <empty>`

   b. Verify the Drive Count, Size, and Vendor.
   The following example illustrates what you should see before you create a RAID array.



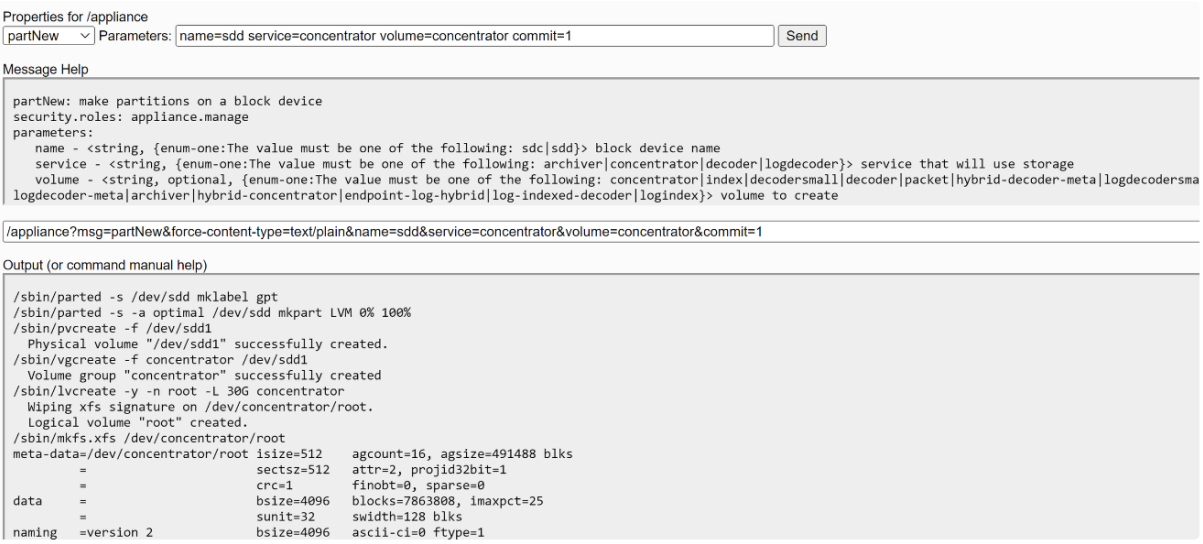
2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.
   `controller=1 enclosure=251 scheme=concentrator commit=1`

The following example illustrates what you should see after you create a RAID array.



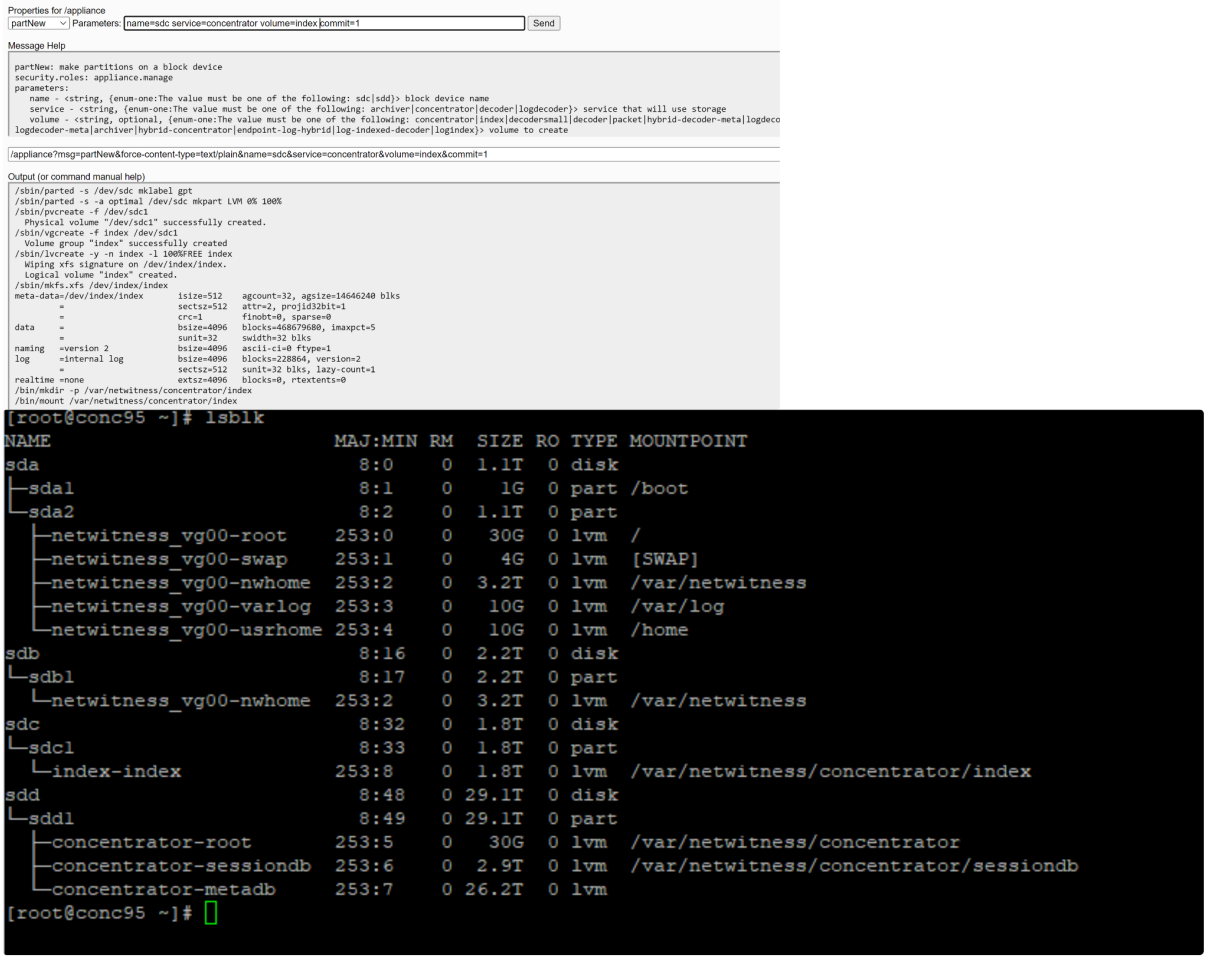
3. Execute the `partNew` command to create the **concentrator** partition first with the following parameters. You must create the **concentrator** volume before **index** volume or it will fail.

```
name=sdd service=concentrator volume=concentrator commit=1
```



4. Execute the partNew command with the following parameters to create an **index** on SSDs.

```
name=sdc service=concentrator volume=index commit=1
```




```
[root@conc95 ~]# lsblk
NAME                          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                             8:0    0  1.1T  0 disk
├─sda1                          8:1    0    1G  0 part /boot
└─sda2                          8:2    0  1.1T  0 part
  ├─netwitness_vg00-root      253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:4    0   10G  0 lvm  /home
sdb                             8:16   0  2.2T  0 disk
└─sdb1                          8:17   0  2.2T  0 part
  └─netwitness_vg00-nwhome    253:2    0  3.2T  0 lvm  /var/netwitness
sdc                             8:32   0  1.8T  0 disk
└─sdc1                          8:33   0  1.8T  0 part
  └─index-index               253:8    0  1.8T  0 lvm  /var/netwitness/concentrator/index
sdd                             8:48   0 29.1T  0 disk
└─sdd1                          8:49   0 29.1T  0 part
  ├─concentrator-root         253:5    0   30G  0 lvm  /var/netwitness/concentrator
  ├─concentrator-sessiondb    253:6    0  2.9T  0 lvm  /var/netwitness/concentrator/sessiondb
  └─concentrator-metadb       253:7    0 26.2T  0 lvm
[root@conc95 ~]#
```

5. Execute the `srvAlloc` command with the following parameters.

```
service=concentrator volume=concentrator commit=1
```

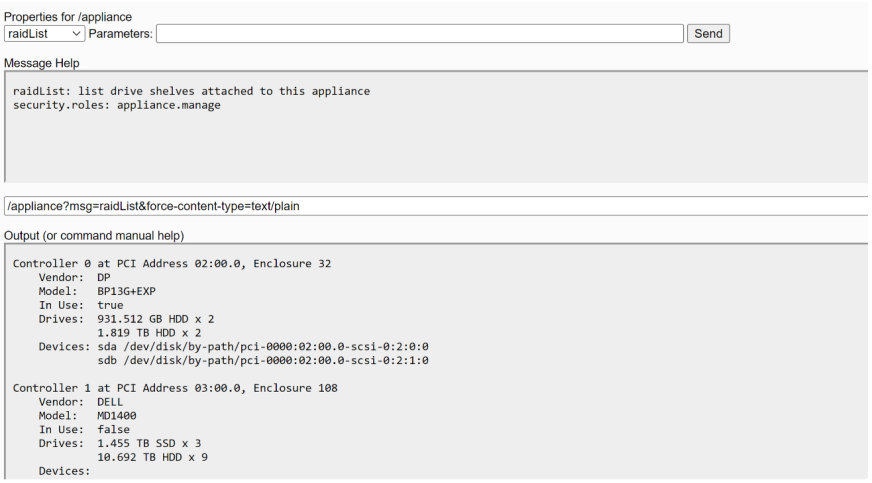
Configure Storage for Concentrator using NW-PV-D/NW-PV-D-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for a Log Decoder Hybrid physical host.

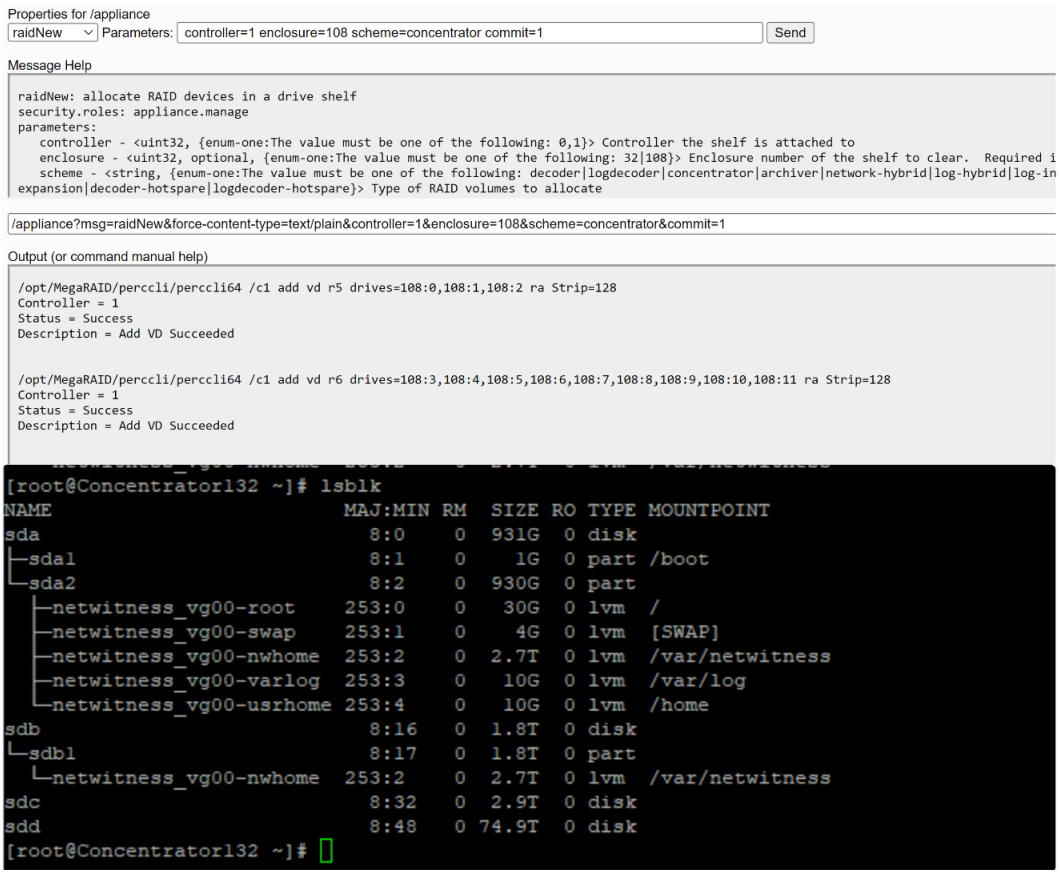
1. Execute the `raidList` command.

   a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
   You should see the following information.
   ```
   In Use: FALSE
   Devices: <empty>
   ```

   b. Verify the Drive Count, Size, and Vendor.
   The following example illustrates what you should see before you create a RAID array.

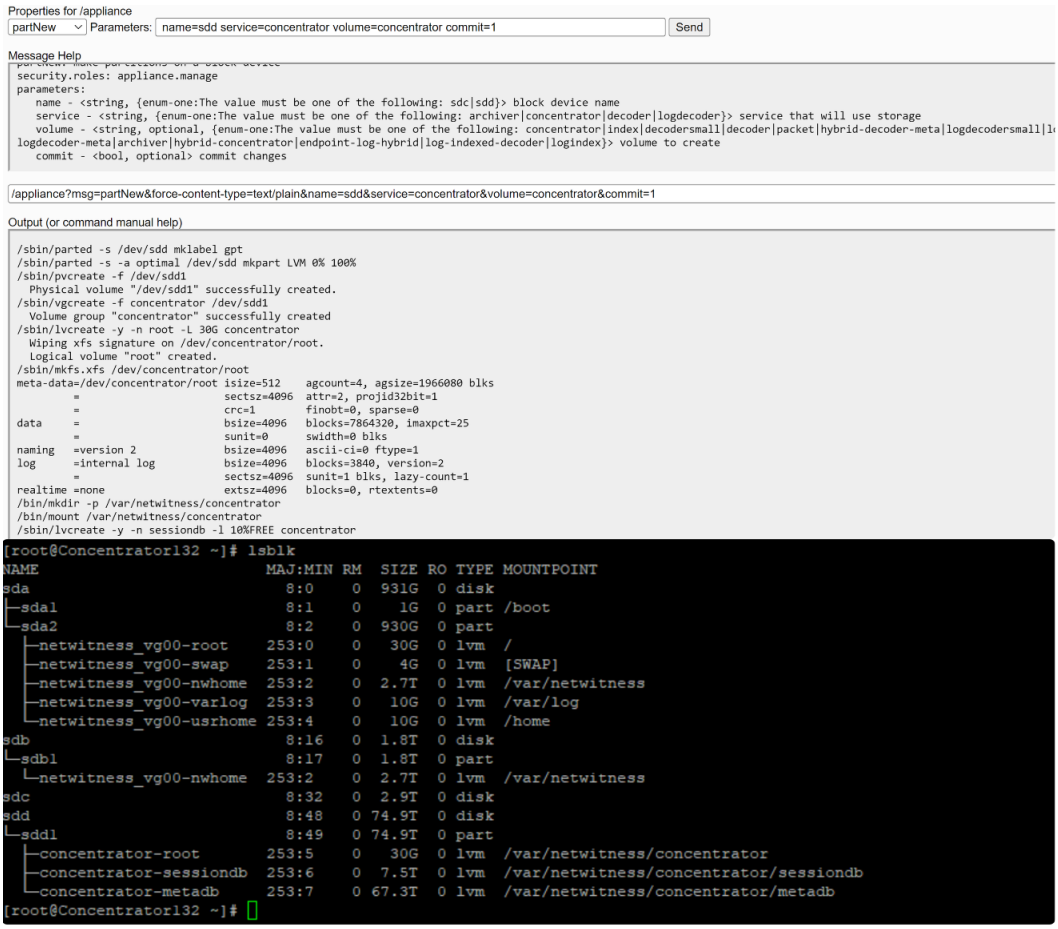


2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.
   `controller=1 enclosure=108 scheme=concentrator commit=1`
   The following example illustrates what you should see after you create a RAID array.



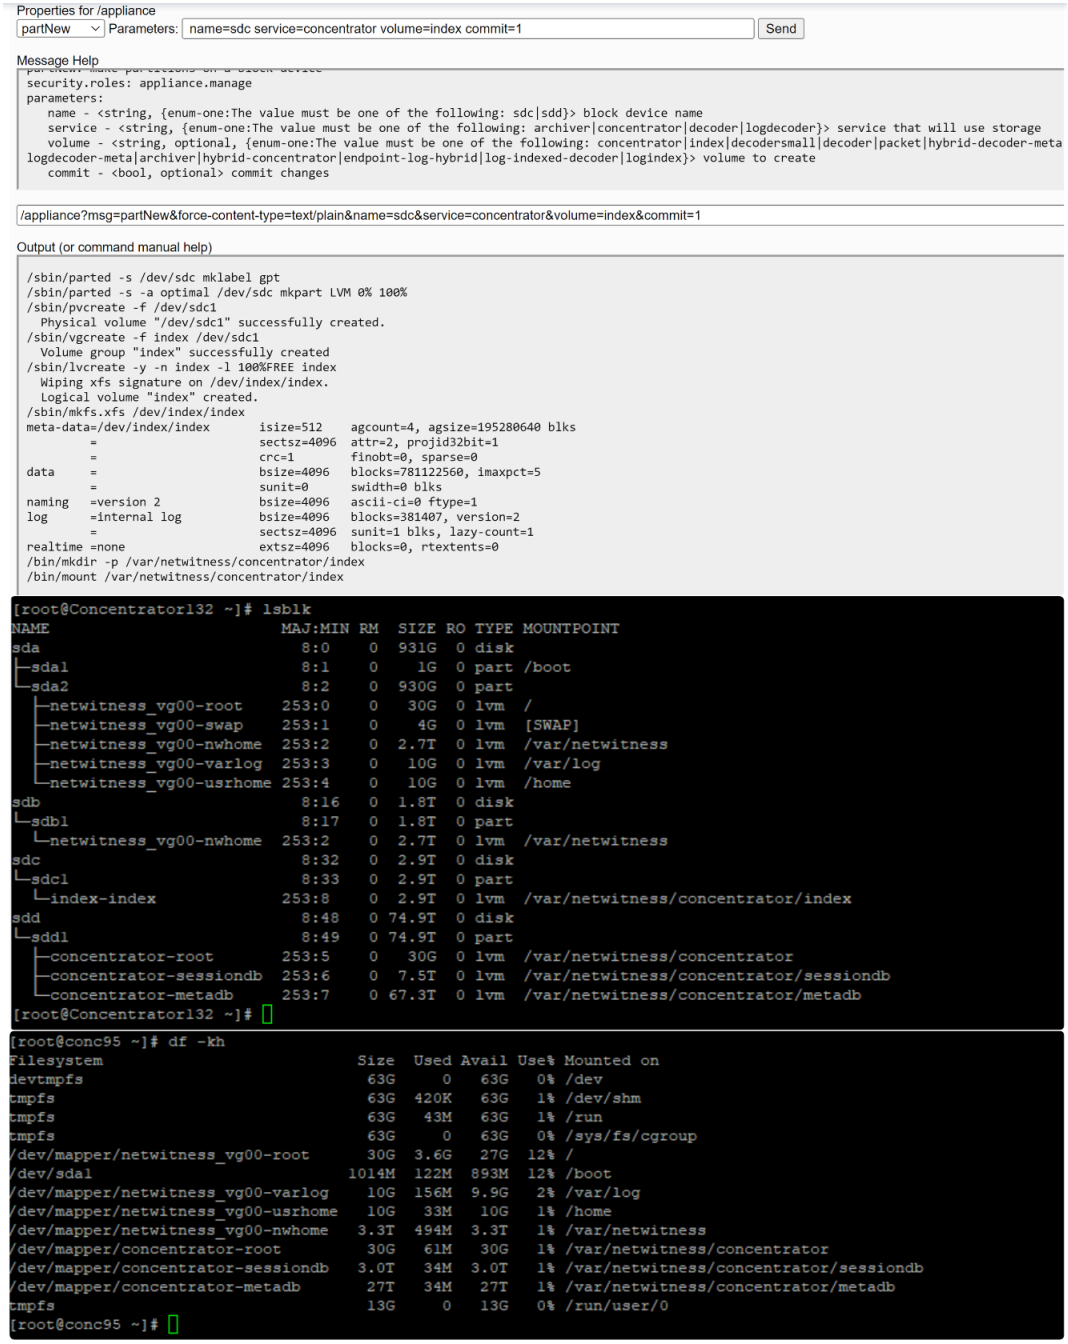
3. Execute the `partNew` command to create the **concentrator** partition first with the following parameters. You must create the concentrator volume before index volume or it will fail.

   `name=sdd service=concentrator volume=concentrator commit=1`



4. Execute the `partNew` command with the following parameters to create an **index** on SSDs.

   `name=sdc service=concentrator volume=index commit=1`



5. Execute the `srvAlloc` command with the following parameters.

   `service=concentrator volume=concentrator commit=1`

Appendix F. Sample Storage Configuration Scenarios for Meta Disk Kits

This appendix illustrates the examples of how to configure a Meta Disk Kit described in Configure Block devices for Drive pack section for Network Decoder in the below scenarios.

- Meta-Only: No Externally Attached Storage
- Maximize PowerVault Storage capacity for Network Decoder

Meta-Only (One kit – 3 SED drives configured as RAID 5)

Configure one Meta Disk kit as decodersmall volume for Network Decoder:

Note:
1. The configuration for Log Decoder is the similar as that of a Network Decoder. Substitute the service and volume names that correspond to Log Decoder.
2. A single Meta Disk pack is configured as RAID5 (3 drives) and two Meta Disk packs (6 drives) are configured as RAID6.
3. When configuring two Meta Disk Packs or adding the second Meta Disk pack, use appropriate disk slot numbers when creating the virtual drive (step 5 below).

On the Series 6 (Dell R640) appliance, the Meta Drive Pack disks are installed in slots 4, 5 and 6. The virtual drive configuration requires identifying the controller ID and Enclosure ID (EID). On Series 6 appliance, the controller ID and Enclosure IDs are 0 and 64. However, the nwraidtool.py script that is installed on every server can help to confirm these ID numbers, or the PercCLI commands as shown in the below steps.

Identify controller/enclosures:

1. Identify the controller ID (Ctl) for the internal controller (PERC H740P Mini/ PERC H750). In the below figure the controller ID is 0 and highlighted in yellow. The drive count is displayed under PDs.

   ```
   /opt/MegaRAID/perccli/perccli64 show
   ```

```
[root@117Decoder ~]# /opt/MegaRAID/perccli/perccli64 show
Status Code = 0
Status = Success
Description = None

Number of Controllers = 2
Host Name = 117Decoder
Operating System  = Linux 3.10.0-1160.36.2.el7.x86_64
StoreLib IT Version = 07.0400.0200.0400
StoreLib IR3 Version = 15.50-0

System Overview :
===============

-----------------------------------------------------------------------------
Ctl Model          Ports PDs DGs DNOpt VDs VNOpt BBU sPR DS EHS ASOs Hlth
-----------------------------------------------------------------------------
  0 PERCH740PMini      8 [redacted]     0 [redacted]     0 Opt On  -  N      0 Opt
  1 PERCH840Adapter    8  24   1     0   1     0 Opt On  -  N      0 Opt
-----------------------------------------------------------------------------

Ctl=Controller Index|DGs=Drive groups|VDs=Virtual drives|Fld=Failed
PDs=Physical drives|DNOpt=DG NotOptimal|VNOpt=VD NotOptimal|Opt=Optimal
Msng=Missing|Dgd=Degraded|NdAtn=Need Attention|Unkwn=Unknown
sPR=Scheduled Patrol Read|DS=DimmerSwitch|EHS=Emergency Hot Spare
Y=Yes|N=No|ASOs=Advanced Software Options|BBU=Battery backup unit
Hlth=Health|Safe=Safe-mode boot
```

2. Identify the Enclosure ID (EID) for controller ‘0’. In this case the EID is ‘64’ and highlighted in yellow.

   `/opt/MegaRAID/perccli/perccli64 /c0 /eall show`

```
[root@117Decoder ~]# /opt/MegaRAID/perccli/perccli64 /c0 /eall show
Controller = 0
Status = Success
Description = None


Properties :
==========

-----------------------------------------------------------------------------
EID State Slots PD PS Fans TSs Alms SIM Port#      ProdID    VendorSpecific
-----------------------------------------------------------------------------
 64 OK       10 [redacted]  0    0    0    0    0   1 00 & 00 x8 BP14G+EXP  +
-----------------------------------------------------------------------------

EID-Enclosure Device ID |PD-Physical drive count |PS-Power Supply count|
TSs-Temperature sensor count |Alms-Alarm count |SIM-SIM Count
```

3. Identify the new Meta Disk kit disk slot numbers (Slots 4 through 9) on the controller PERC H740P Mini. These drives do not belong to any Drive Group (DG). The DG column for these drives display a dash ‘-‘, the State column shows ‘UGood’, and the SED value as ‘Y’. See the entries highlighted in yellow.

```
[root@116S6Core1 ~]# /opt/MegaRAID/perccli/perccli64 /c0 /eall /sall show
Controller = 0
Status = Success
Description = Show Drive Information Succeeded.


Drive Information :
==================

-----------------------------------------------------------------------
EID:Slt DID State DG     Size Intf Med SED PI SeSz Model            Sp
-----------------------------------------------------------------------
64:0      0 Onln   0 1.090 TB SAS  HDD Y   N  512B ST1200MM0069     U
64:1      1 Onln   0 1.090 TB SAS  HDD Y   N  512B ST1200MM0069     U
64:2      2 Onln   1 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:3      3 Onln   1 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:4      4 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:5      5 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:6      6 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:7      7 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:8      8 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:9      9 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
-----------------------------------------------------------------------

EID-Enclosure Device ID|Slt-Slot No.|DID-Device ID|DG-DriveGroup
DHS-Dedicated Hot Spare|UGood-Unconfigured Good|GHS-Global Hotspare
UBad-Unconfigured Bad|Onln-Online|Offln-Offline|Intf-Interface
Med-Media Type|SED-Self Encryptive Drive|PI-Protection Info
SeSz-Sector Size|Sp-Spun|U-Up|D-Down/PowerSave|T-Transition|F-Foreign
UGUnsp-Unsupported|UGShld-UnConfigured shielded|HSPShld-Hotspare shielded
CFShld-Configured shielded|Cpybck-CopyBack|CBShld-Copyback Shielded
```

4. Identify the existing block devices (sda and sdb) on the host used by the OS and NetWitness. Use the command ‘lsblk’ to list out the block devices.

   ```
   lsblk
   ```

```
[root@116S6Core1 ~]# lsblk
NAME                        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                           8:0    0  1.1T  0 disk
├─sda1                        8:1    0    1G  0 part /boot
└─sda2                        8:2    0  1.1T  0 part
  ├─netwitness_vg00-root    253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap    253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome  253:2    0  3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog  253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome 253:4    0   10G  0 lvm  /home
sdb                           8:16   0  2.2T  0 disk
└─sdb1                        8:17   0  2.2T  0 part
  └─netwitness_vg00-nwhome  253:2    0  3.2T  0 lvm  /var/netwitness
```

Configure the virtual drive(s) on the installed Meta Disk kit (slots 4,5, and 6):

5. Create the Virtual Drive/Drive Group (DG) on the internal controller using the disks in slot 4 through 6 using the below command.

   ```
   /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6
   strip=128
   ```

```
[root@117Decoder ~]# /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6 strip=128
Controller = 0
Status = Success
Description = Add VD Succeeded
```

6. The new virtual drive shows up as ‘2/2’ under DG/VD column. See the entry highlighted in yellow.

   `/opt/MegaRAID/perccli/perccli64 /c0 /vall show`

```
[root@117Decoder perccli]#  /opt/MegaRAID/perccli/perccli64 /c0 /vall show
Controller = 0
Status = Success
Description = None


Virtual Drives :
==============

-------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache Cac sCC     Size Name
-------------------------------------------------------------
0/0   RAID1 Optl  RW     Yes     RWBD  -   OFF 1.090 TB
1/1   RAID1 Optl  RW     Yes     RWBD  -   OFF 2.182 TB
2/2   RAID5 Optl  RW     Yes     RWBD  -   OFF 4.365 TB
-------------------------------------------------------------

Cac=CacheCade|Rec=Recovery|OfLn=OffLine|Pdgd=Partially Degraded|Dgrd=Degraded
Optl=Optimal|RO=Read Only|RW=Read Write|HD=Hidden|TRANS=TransportReady|B=Blocked|
Consist=Consistent|R=Read Ahead Always|NR=No Read Ahead|WB=WriteBack|
FWB=Force WriteBack|WT=WriteThrough|C=Cached IO|D=Direct IO|sCC=Scheduled
Check Consistency
```

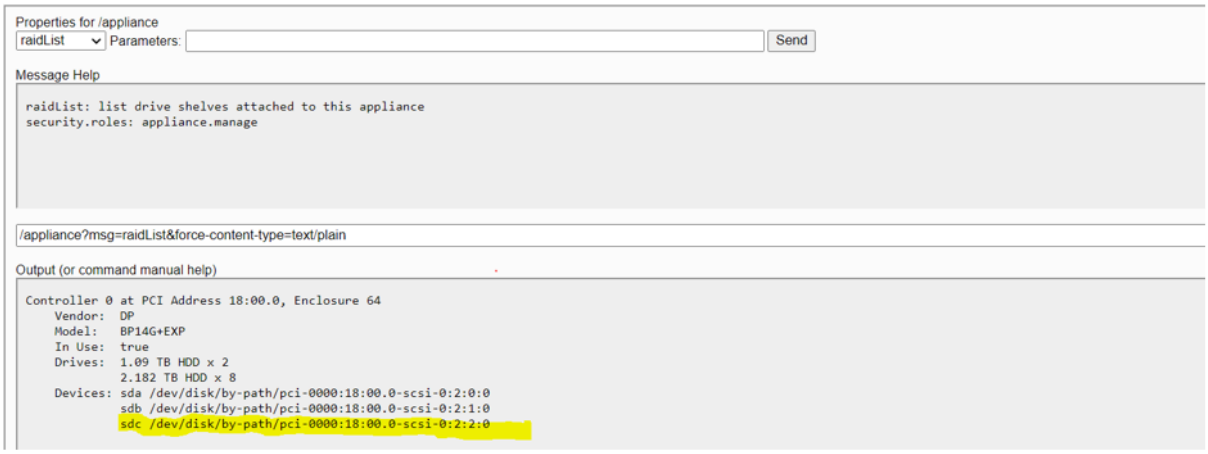
7. Identify the new block device on the host. The block device name is identified under the NAME column. In this example the new block device is 'sdc'. This block device name is required when configuring storage. Use ‘lsblk’ to list the block devices. See entry highlighted in yellow.

   `lsblk`

```
[root@117Decoder perccli]# lsblk
NAME                      MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                         8:0    0   1.1T  0 disk
├─sda1                      8:1    0     1G  0 part /boot
├─sda2                      8:2    0   256M  0 part /boot/efi
└─sda3                      8:3    0   1.1T  0 part
  ├─netwitness_vg00-root  253:0    0    30G  0 lvm  /
  ├─netwitness_vg00-swap  253:1    0     4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome 253:2   0   3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog 253:3   0    10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome 253:4  0    10G  0 lvm  /home
sdb                         8:16   0   2.2T  0 disk
└─sdb1                      8:17   0   2.2T  0 part
  └─netwitness_vg00-nwhome 253:2   0   3.2T  0 lvm  /var/netwitness
sdc                         8:32   0   4.4T  0 disk
```

8. Now, you must Configure the above block device as decodersmall for Network Decoder. Use the REST API (Ex: **https://<HostIP>:50106**, HostIP is the ip address of the decoder host) to retrieve the existing raid list. Note the block device name corresponding to the Meta Disk kit. In this case it is ‘sdc’.

   Note: If a second Meta Disk kit is being configured, the block device would be ‘sdd’.

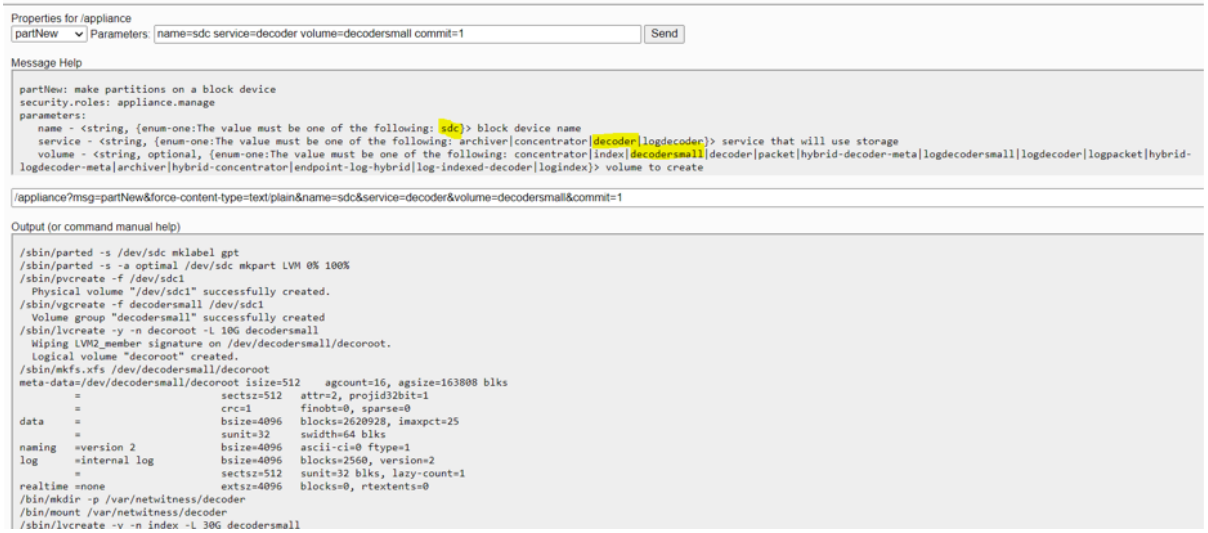


9. Make partitions on the block device (‘sdc’) for decoder service with the following parameters:

   `name=sdc service=decoder volume=decodersmall commit=1`

   `For logdecoder use: name=sdc service=logdecoder volume=logdecodersmall commit=1`

   Note: If a second Meta Disk kit is being configured, the volume for decoder would be ‘decodersmall0’. For logdecoder, it is ‘logdecodersmall0’.

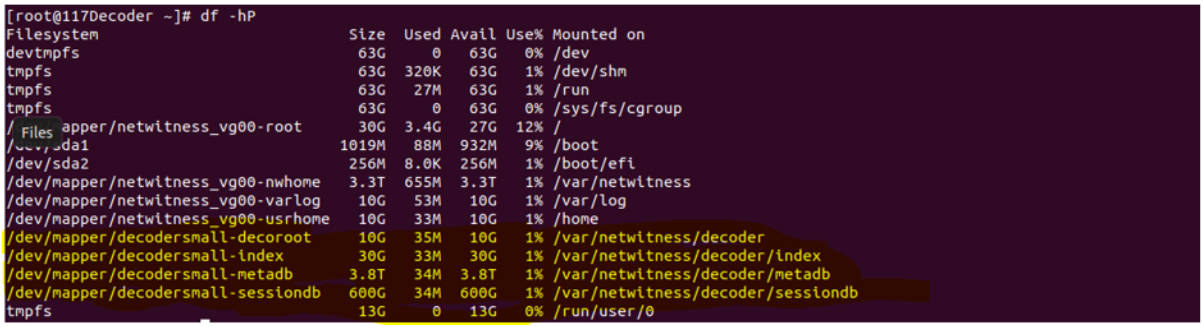


10. Allocate the decodersmall to Decoder service using ‘srvAlloc’.

    `service=decoder volume=decodersmall commit=1`

    Note: For logdecoder: service=logdecoder volume=logdecodersmall commit=1





Maximize PowerVault Storage Capacity (One Kit – 3 SED Drives Configured as RAID 5)

Configure one Meta Disk kit as decodersmall volume for Network Decoder maximizing the PowerVault Storage.

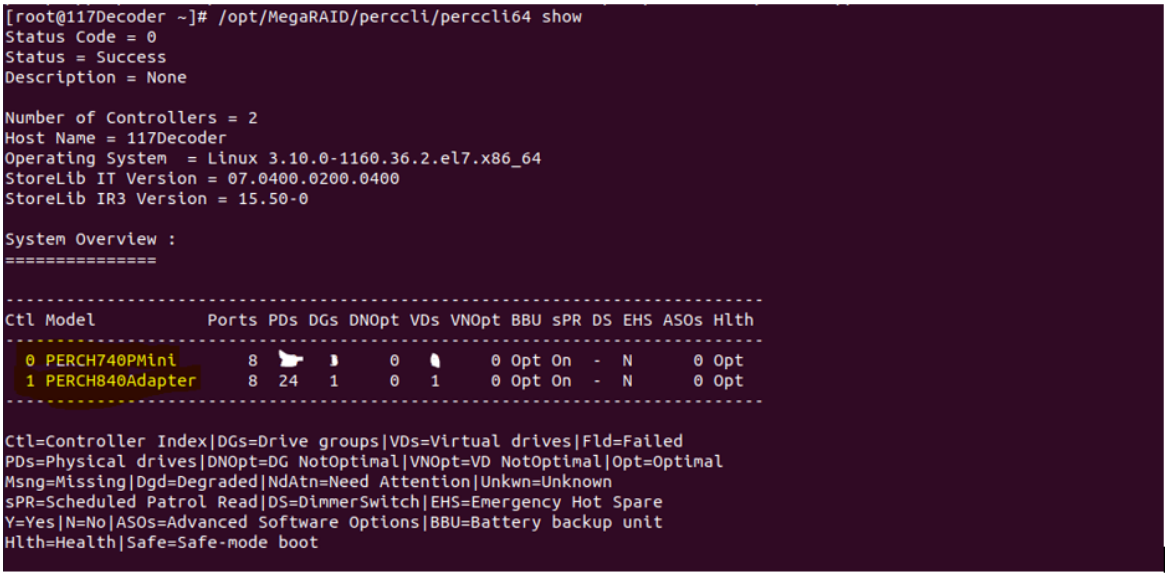
Note:
1. Best practice – when 1-4 PowerVaults are configured, a single Meta Disk kit is recommended and when 5-8 PowerVaults are configured two (2) Meta Disk kits are recommended (optional).
2. The configuration for Log Decoder is the similar as that of a Network Decoder. Substitute the service and volume names that correspond to Log Decoder.
3. A singe Meta Disk kit is configured as RAID5 (3 drives) and two Meta Disk kits (6 drives) are configured as RAID6.
4. When configuring two Meta Disk kits or adding thea second Meta Disk kit, use appropriate disk slot numbers when creating the virtual drive (step 5 below).

On the Series 6 (Dell R640) appliance, the Meta Drive Pack disks are installed in slots 4, 5 and 6. The virtual drive configuration requires identifying the controller ID and Enclosure ID (EID). On Series 6 appliance, the controller ID and Enclosure IDs are 0 and 64. However, the nwraidtool.py script that is installed on every server can help to confirm these ID numbers, or the PercCLI commands as shown in the below steps.

Identify controller/enclosures:

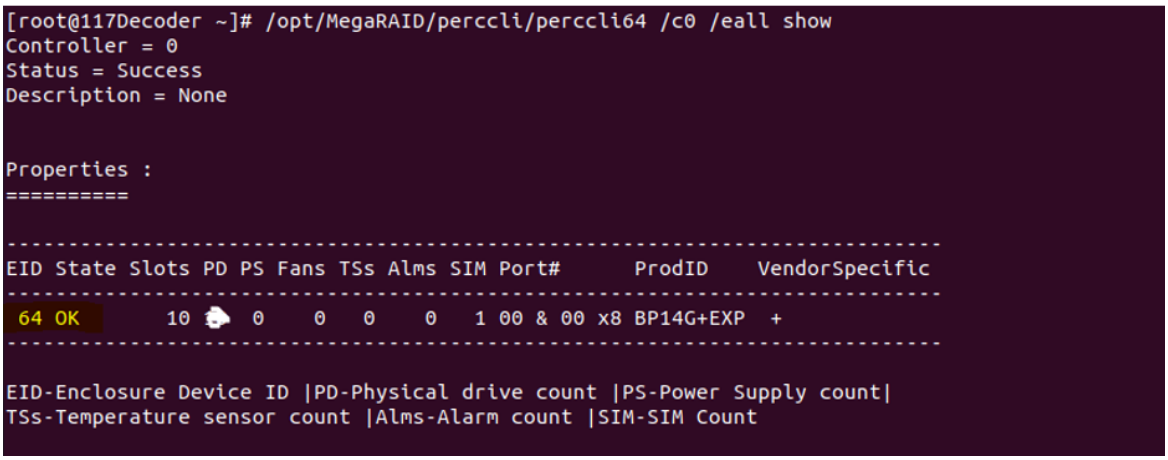
1. Identify the controller ID (Ctl) for the internal controller (PERC H740P Mini/ PERC H750). In the below figure the controller ID is 0 and highlighted in yellow. The drive count is displayed under PDs.

   `/opt/MegaRAID/perccli/perccli64 show`



2. Identify the Enclosure ID (EID) for controller ‘0’. In this case the EID is ‘64’ (highlighted in yellow).

   `/opt/MegaRAID/perccli/perccli64 /c0 /eall show`



3. Identify the new Meta Disk kit disk slot numbers (Slots 4 through 9) on the controller PERC H740P Mini. These drives do not belong to any Drive Group (DG). The DG column for these drives display a dash ‘-‘, the State column shows ‘UGood’, and the SED value as ‘Y’. See the entries highlighted in yellow.

```
[root@116S6Core1 ~]# /opt/MegaRAID/perccli/perccli64 /c0 /eall /sall show
Controller = 0
Status = Success
Description = Show Drive Information Succeeded.


Drive Information :
==================

-----------------------------------------------------------------------
EID:Slt DID State DG     Size Intf Med SED PI SeSz Model            Sp
-----------------------------------------------------------------------
64:0      0 Onln   0 1.090 TB SAS  HDD Y   N  512B ST1200MM0069     U
64:1      1 Onln   0 1.090 TB SAS  HDD Y   N  512B ST1200MM0069     U
64:2      2 Onln   1 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:3      3 Onln   1 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:4      4 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:5      5 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:6      6 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:7      7 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:8      8 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
64:9      9 UGood  - 2.182 TB SAS  HDD Y   N  512B ST2400MM0149     U
-----------------------------------------------------------------------

EID-Enclosure Device ID|Slt-Slot No.|DID-Device ID|DG-DriveGroup
DHS-Dedicated Hot Spare|UGood-Unconfigured Good|GHS-Global Hotspare
UBad-Unconfigured Bad|Onln-Online|Offln-Offline|Intf-Interface
Med-Media Type|SED-Self Encryptive Drive|PI-Protection Info
SeSz-Sector Size|Sp-Spun|U-Up|D-Down/PowerSave|T-Transition|F-Foreign
UGUnsp-Unsupported|UGShld-UnConfigured shielded|HSPShld-Hotspare shielded
CFShld-Configured shielded|Cpybck-CopyBack|CBShld-Copyback Shielded
```

4. Identify the existing block devices (sda and sdb) on the host used by the OS and NetWitness. Use the command 'lsblk' to list out the block devices.

   ```
   lsblk
   ```

```
[root@116S6Core1 ~]# lsblk
NAME                        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                           8:0    0  1.1T  0 disk
├─sda1                        8:1    0    1G  0 part /boot
└─sda2                        8:2    0  1.1T  0 part
  ├─netwitness_vg00-root    253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap    253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome  253:2    0  3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog  253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome 253:4    0   10G  0 lvm  /home
sdb                           8:16   0  2.2T  0 disk
└─sdb1                        8:17   0  2.2T  0 part
  └─netwitness_vg00-nwhome  253:2    0  3.2T  0 lvm  /var/netwitness
```

Configure virtual drive(s) on installed meta pack disks (slots 4,5 and 6):

5. Create the Virtual Drive/Drive Group (DG) on the internal controller using the disks in slot 4 through 6 using the below command.

   ```
   /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6
   strip=128
   ```

```
[root@117Decoder ~]# /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6 strip=128
Controller = 0
Status = Success
Description = Add VD Succeeded
```

6. The new virtual drive shows up as '2/2' (highlighted in yellow) under DG/VD column.

   `/opt/MegaRAID/perccli/perccli64 /c0 /vall show`

```
[root@117Decoder ~]#  /opt/MegaRAID/perccli/perccli64 /c0 /vall show
Controller = 0
Status = Success
Description = None


Virtual Drives :
==============

-------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache Cac sCC     Size Name
-------------------------------------------------------------
0/0   RAID1 Optl  RW     Yes     RWBD  -   OFF 1.090 TB
1/1   RAID1 Optl  RW     Yes     RWBD  -   OFF 2.182 TB
2/2   RAID5 Optl  RW     Yes     RWBD  -   OFF 4.365 TB
-------------------------------------------------------------

Cac=CacheCade|Rec=Recovery|OfLn=OffLine|Pdgd=Partially Degraded|Dgrd=Degraded
Optl=Optimal|RO=Read Only|RW=Read Write|HD=Hidden|TRANS=TransportReady|B=Blocked|
Consist=Consistent|R=Read Ahead Always|NR=No Read Ahead|WB=WriteBack|
FWB=Force WriteBack|WT=WriteThrough|C=Cached IO|D=Direct IO|sCC=Scheduled
Check Consistency
```

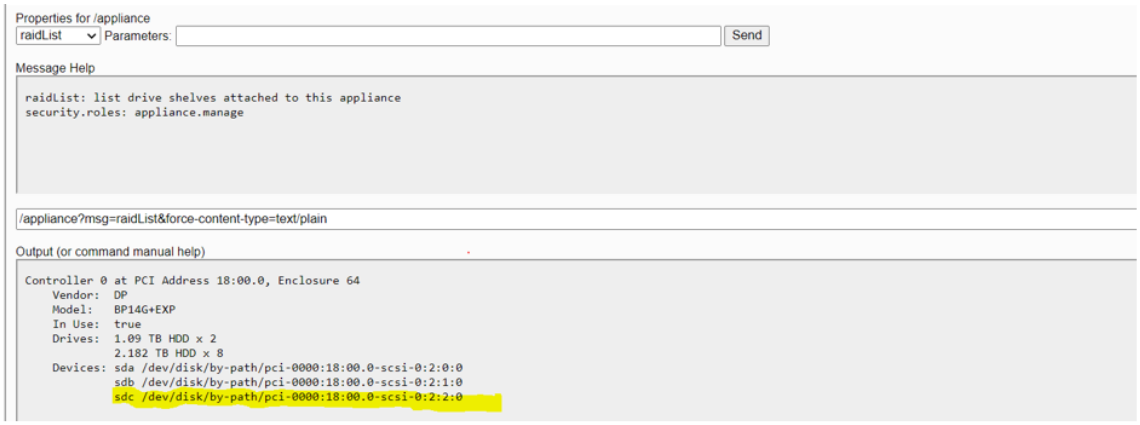
7. Identify the new block device on the host. The block device name is identified under the NAME column. In this example the new block device is 'sdc'. This block device name is required when configuring storage. Use 'lsblk' to list the block devices.

   `lsblk`

```
[root@117Decoder ~]# lsblk
NAME                        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                           8:0    0  1.1T  0 disk
├─sda1                        8:1    0    1G  0 part /boot
├─sda2                        8:2    0  256M  0 part /boot/efi
└─sda3                        8:3    0  1.1T  0 part
  ├─netwitness_vg00-root    253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap    253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome  253:2    0  3.2T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog  253:3    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome 253:4    0   10G  0 lvm  /home
sdb                           8:16   0  2.2T  0 disk
└─sdb1                        8:17   0  2.2T  0 part
  └─netwitness_vg00-nwhome  253:2    0  3.2T  0 lvm  /var/netwitness
sdc                           8:32   0  4.4T  0 disk
```

8. Configure the above block device as decodersmall volume for Decoder (volume name for Log Decoder is logdecodersmall). Use REST API (Ex: https://<HostIP>:50106, HostIP is the ip address of the decoder host) to retrieve the existing raid list. Note the block device name corresponding to the Meta Disk kit. In this case, it is 'sdc' (highlighted in yellow)

   Note: If a second Meta Disk kit is being configured, the block device would be 'sdd'.

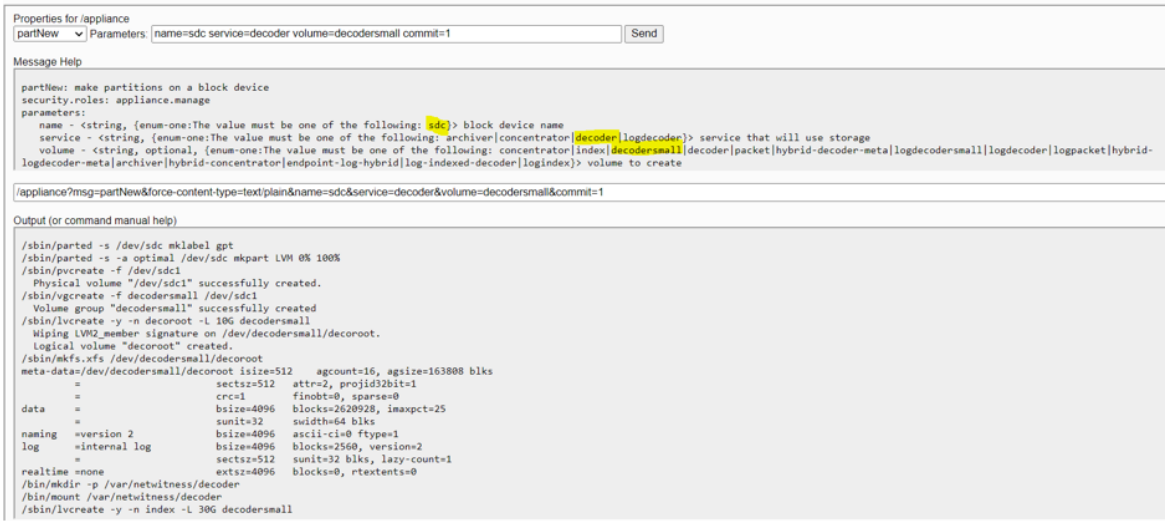


9. Make partitions on the block device (‘sdc’) with the following parameters:

   `name=sdc service=decoder volume=decodersmall commit=1`

   `For logdecoder service, use: name=sdc service=logdecoder volume=logdecodersmall commit=1`

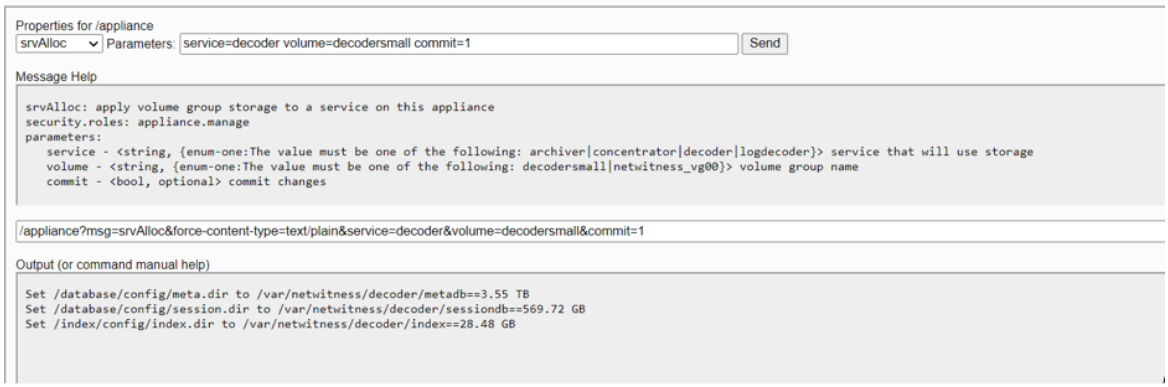
> **Note:** If a second Meta Disk kit is being configured, the volume for decoder service is ‘decodersmall0’ and ‘logdecodersmall0’ for logdecoder service.

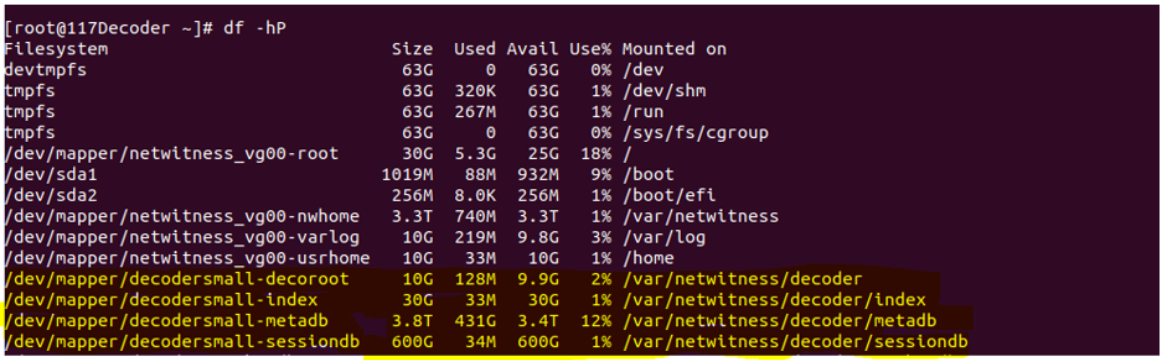


10. Allocate the decodersmall to Decoder service

    `service=decoder volume=decodersmall commit=1For logdecoder use: service=logdecoder volume=logdecodersmall commit=1`

> **Note:** If a second Meta Pack is being configured, the volume for decoder is ‘decodersmall0’ and ‘logdecodersmall0’ for logdecoder service.





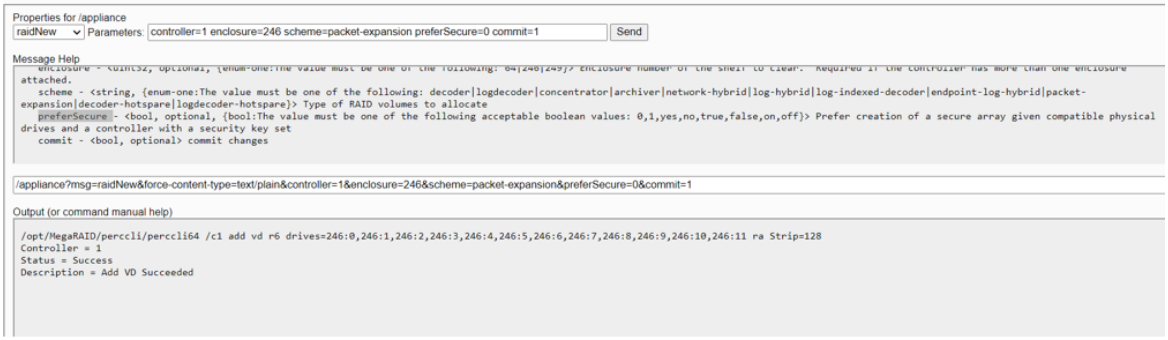
Configure external storage (PowerVault) as packetDB for Decoder:

11. Use REST API (Ex: https://<HostIP>:50106, HostIP is the ip address of the decoder host) to configure raid on the attached PowerVault. Identify the controller and enclosure information using raidList. In this example, two PowerVaults are connected to the Decoder appliance (controller 1 with enclosure 246 and 249). REST API is used in the next steps to configure the PowerVault (Enclosure 246) connected to Controller 1.

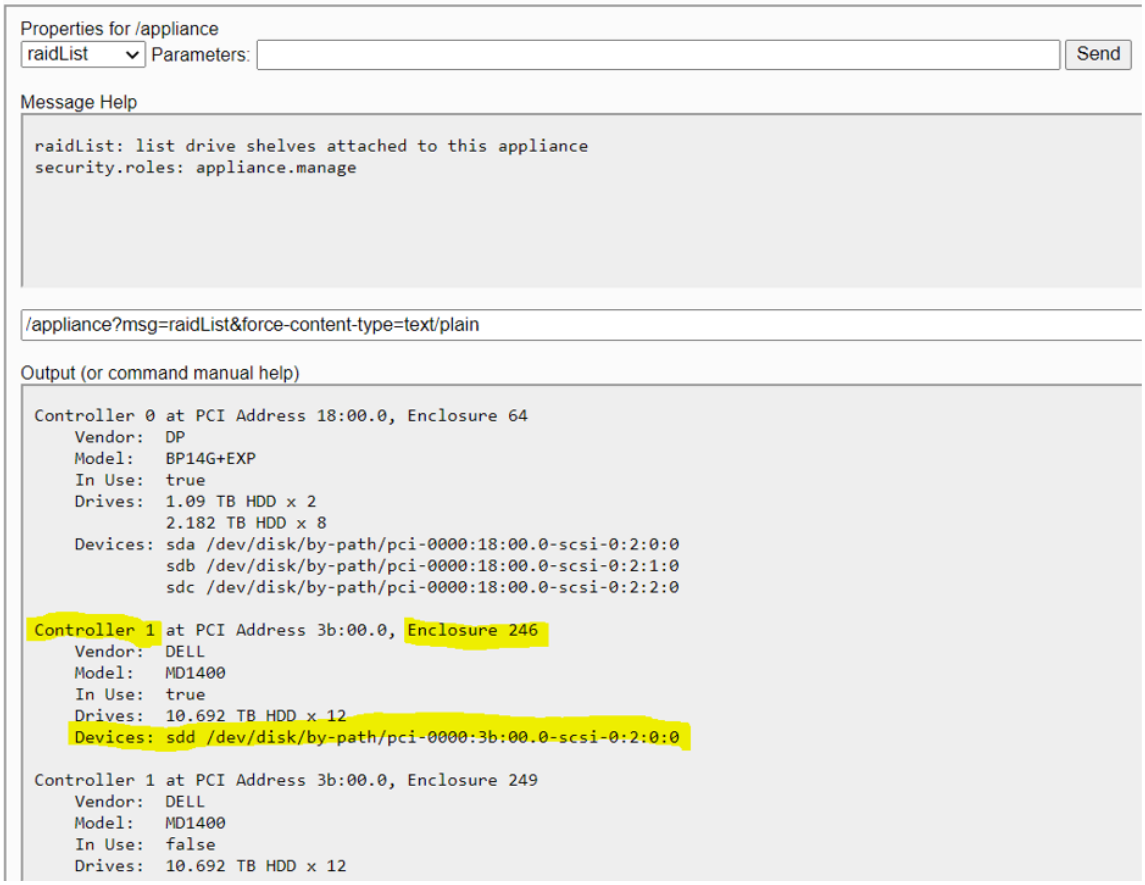


12. Create the raid on the attached PowerVault. Identify the controller and enclosure from ‘raidList’. Use appropriate value for preferSecure to enable encryption. To turn on encryption on SED capable drives, use preferSecure=1 after setting the security key on the controller. Encryption can also be turned on after configuring storage using the steps listed in Appendix B. Encrypt a Series 6E Core or Hybrid Host of the storage guide.

```
controller=1 enclosure=246 scheme=packet-expansion preferSecure=0
commit=1
```

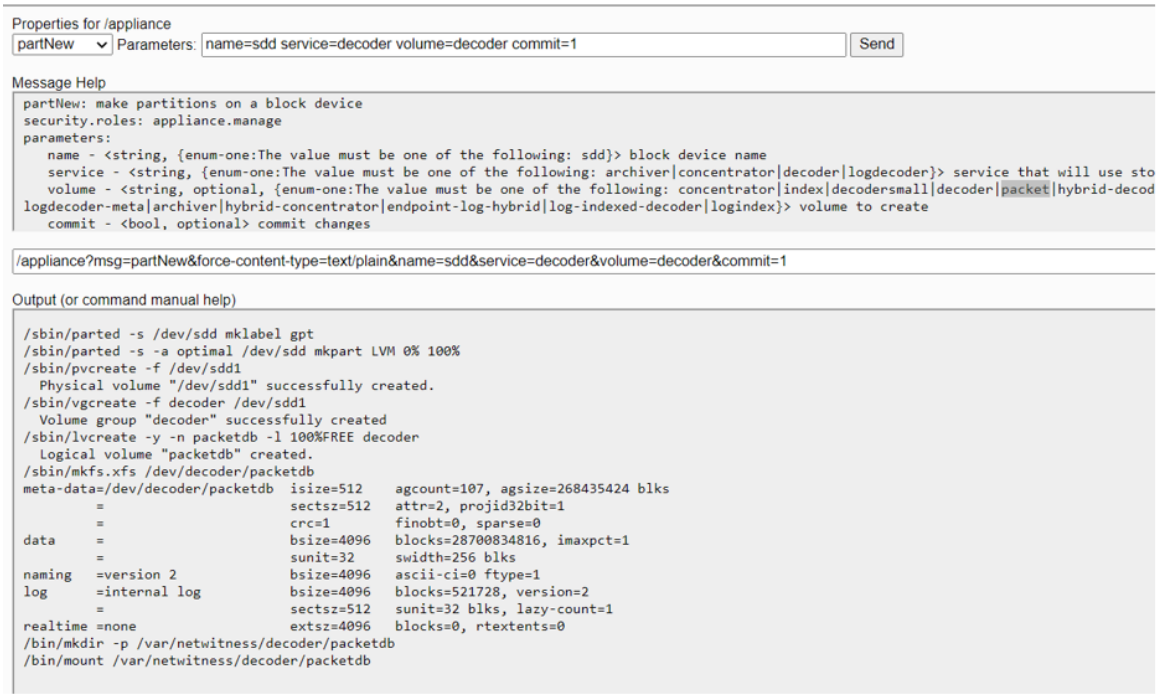


13. Identify the block device created in the earlier step. Use ‘raidList’ to retrieve all the block devices. In this case, it is ‘sdd’.



14. Make partitions on the block device (‘sdd’) using ‘partNew’.

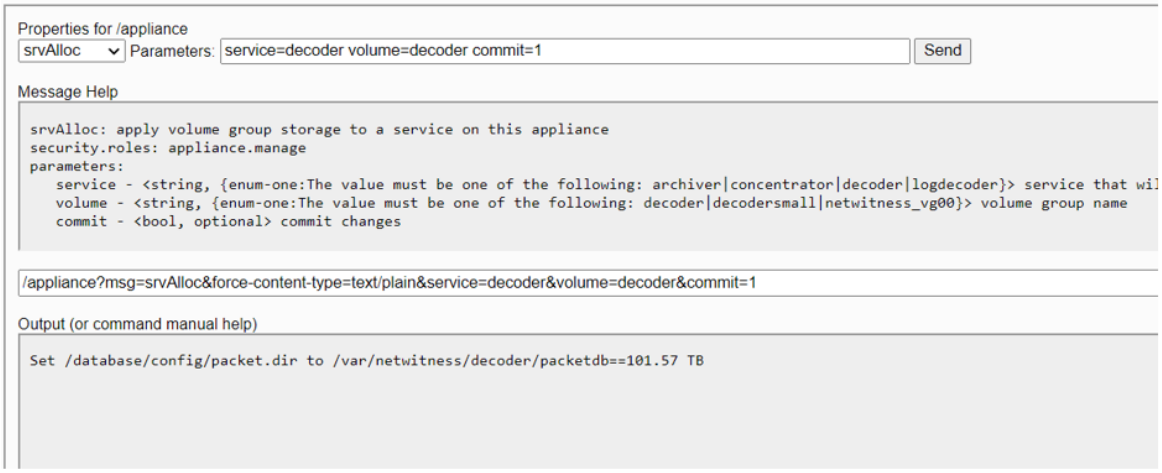
```
name=sdd service=decoder volume=decoder commit=1
```




```
[root@117Decoder ~]# df -hP
Filesystem                              Size  Used Avail Use% Mounted on
devtmpfs                                 63G     0   63G   0% /dev
tmpfs                                    63G  160K   63G   1% /dev/shm
tmpfs                                    63G   27M   63G   1% /run
tmpfs                                    63G     0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root         30G  3.4G   27G  12% /
/dev/sda1                              1019M   88M  932M   9% /boot
/dev/sda2                               256M  8.0K  256M   1% /boot/efi
/dev/mapper/netwitness_vg00-nwhome      3.3T  656M  3.3T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog       10G   53M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome      10G   33M   10G   1% /home
/dev/mapper/decodersmall-decoroot        10G   36M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index           30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-metadb         3.8T   34M  3.8T   1% /var/netwitness/decoder/metadb
/dev/mapper/decodersmall-sessiondb      600G   34M  600G   1% /var/netwitness/decoder/sessiondb
tmpfs                                    13G     0   13G   0% /run/user/0
/dev/mapper/decoder-packetdb            107T   38M  107T   1% /var/netwitness/decoder/packetdb
[root@117Decoder ~]#
```

15. Allocate the decoder volume to Decoder service using ‘srvAlloc’.

```
service=decoder volume=decoder commit=1
```



Note: Additional PowerVaults shall be configured using REST API as outlined in the Storage Guide.

Appendix G. Sample Storage Configuration for Concentrator Index with One Meta Disk Kit

This appendix illustrates the examples of how to configure concentrator index with one Meta Disk Kit.

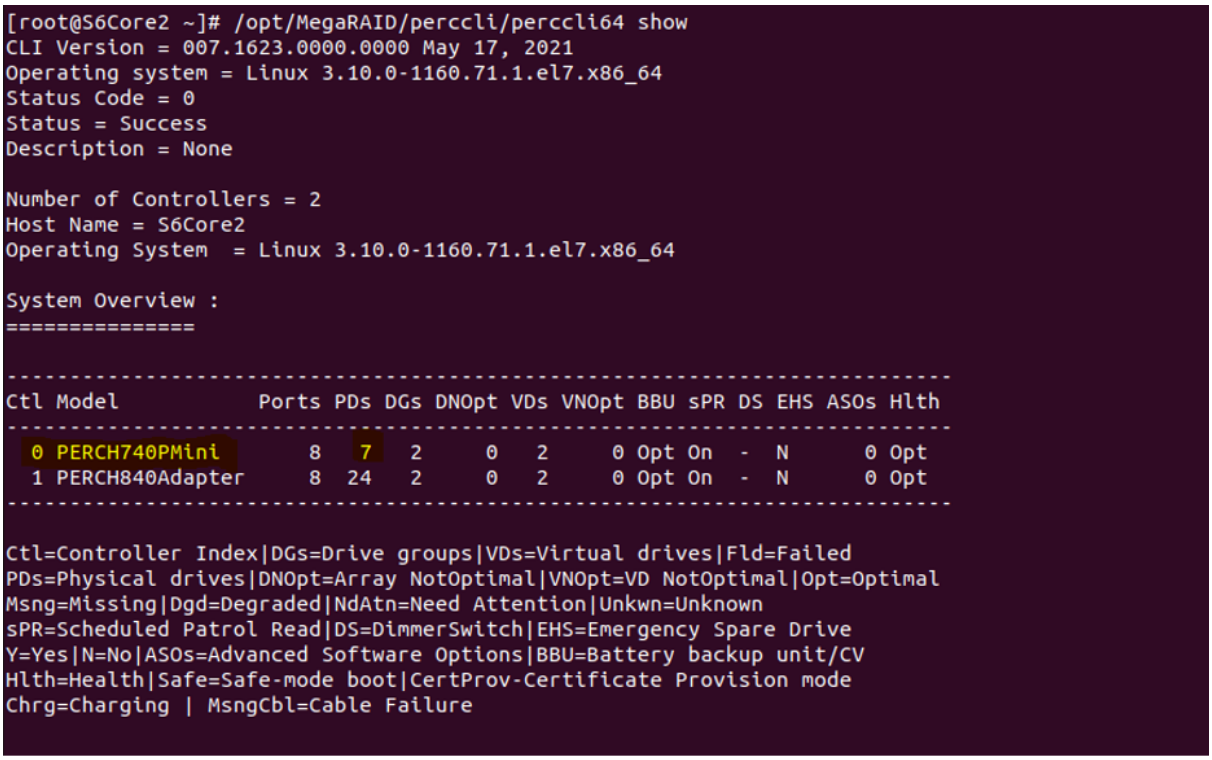
Note:
- Concentrator index meta kit supports a minimum of two to a maximum of 6 drives (SSD).
- A single Meta Disk pack is configured as RAID5 (3 drives) or two Meta Disk packs (6 drives) are configured as RAID6. Two drive configuration is also supported. The two drives are added in slots 4,5 and configured as RAID1.
- When configuring two Meta Disk Packs or adding the second Meta Disk pack, use appropriate disk slot numbers when creating the virtual drive (step 5 below).

On the Series 6 (Dell R640) appliance, the Meta Drive Pack disks are installed in slots 4, 5 and 6. The virtual drive configuration requires identifying the controller ID and Enclosure ID (EID). On Series 6 appliance, the controller ID and Enclosure IDs are 0 and 64. However, the nwraidtool.py script that is installed on every server can help to confirm these ID numbers, or the PercCLI commands as shown in the below steps.

Identify Controller/Enclosures

1. Identify the controller ID (Ctl) for the internal controller (PERC H740P Mini/ PERC H750). In the below figure the controller ID is 0 and highlighted in yellow. The drive count is displayed under PDs.

   ```
   /opt/MegaRAID/perccli/perccli64 show
   ```



2. Identify the Enclosure ID (EID) for controller ‘0’. In this case the EID is ‘64’ and highlighted in yellow.

   `/opt/MegaRAID/perccli/perccli64 /c0 /eall show`

```
[root@S6Core2 ~]# /opt/MegaRAID/perccli/perccli64 /c0 /eall show
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.71.1.el7.x86_64
Controller = 0
Status = Success
Description = None


Properties :
==========

------------------------------------------------------------------------
EID State Slots PD PS Fans TSs Alms SIM Port# ProdID    VendorSpecific
------------------------------------------------------------------------
 64 OK       10  7  0    0   0    0   1 -     BP14G+EXP  +
------------------------------------------------------------------------

EID=Enclosure Device ID | PD=Physical drive count | PS=Power Supply count
TSs=Temperature sensor count | Alms=Alarm count | SIM=SIM Count | ProdID=Product ID
```

3. Identify the new Meta Disk kit disk slot numbers (Slots 4 through 9) on the controller PERC H740P Mini. These drives do not belong to any Drive Group (DG). The DG column for these drives display a dash ‘-‘, the State column shows ‘UGood’, and the SED value as ‘Y’. See the entries highlighted in yellow.

   `/opt/MegaRAID/perccli/perccli64 /c0 /eall /sall show`

```
[root@S6Core2 ~]# /opt/MegaRAID/perccli/perccli64 /c0 /eall /sall show
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.71.1.el7.x86_64
Controller = 0
Status = Success
Description = Show Drive Information Succeeded.


Drive Information :
=================

-------------------------------------------------------------------------------
EID:Slt DID State DG       Size Intf Med SED PI SeSz Model            Sp Type
-------------------------------------------------------------------------------
64:0      0 Onln   0 931.000 GB SAS  HDD N   N  512B ST1000NX0473     U  -
64:1      1 Onln   0 931.000 GB SAS  HDD N   N  512B ST1000NX0473     U  -
64:2      2 Onln   1   1.818 TB SAS  HDD N   N  512B ST2000NX0463     U  -
64:3      3 Onln   1   1.818 TB SAS  HDD N   N  512B ST2000NX0463     U  -
64:4      4 UGood  -   3.492 TB SAS  SSD Y   N  512B KPM5WVUG3T84     U  -
64:5      5 UGood  -   3.492 TB SAS  SSD Y   N  512B KPM5WVUG3T84     U  -
64:6      6 UGood  -   3.492 TB SAS  SSD Y   N  512B KPM6WVUG3T84     U  -
-------------------------------------------------------------------------------

EID=Enclosure Device ID|Slt=Slot No|DID=Device ID|DG=DriveGroup
DHS=Dedicated Hot Spare|UGood=Unconfigured Good|GHS=Global Hotspare
UBad=Unconfigured Bad|Sntze=Sanitize|Onln=Online|Offln=Offline|Intf=Interface
Med=Media Type|SED=Self Encryptive Drive|PI=Protection Info
SeSz=Sector Size|Sp=Spun|U=Up|D=Down|T=Transition|F=Foreign
UGUnsp=UGood Unsupported|UGShld=UGood shielded|HSPShld=Hotspare shielded
CFShld=Configured shielded|Cpybck=CopyBack|CBShld=Copyback Shielded
UBUnsp=UBad Unsupported|Rbld=Rebuild
```

4. Identify the existing block devices (sda and sdb) on the host used by the OS and NetWitness. ‘sdc’ and ‘sdd’ are block devices corresponding to configured PowerVault. Use the command `lsblk` to list out the block devices.

```
[root@S6Core2 ~]# lsblk
NAME                          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                             8:0    0  931G  0 disk
├─sda1                          8:1    0    1G  0 part /boot
└─sda2                          8:2    0  930G  0 part
  ├─netwitness_vg00-root      253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:6    0  2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:7    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:8    0   10G  0 lvm  /home
sdb                             8:16   0  1.8T  0 disk
└─sdb1                          8:17   0  1.8T  0 part
  └─netwitness_vg00-nwhome    253:6    0  2.7T  0 lvm  /var/netwitness
sdc                             8:32   0  2.9T  0 disk
└─sdc1                          8:33   0  2.9T  0 part
  └─index-index               253:5    0  2.9T  0 lvm  /var/netwitness/concentrator/index
sdd                             8:48   0 74.9T  0 disk
└─sdd1                          8:49   0 74.9T  0 part
  ├─concentrator-root         253:2    0   30G  0 lvm  /var/netwitness/concentrator
  ├─concentrator-sessiondb    253:3    0  7.5T  0 lvm  /var/netwitness/concentrator/sessiondb
  └─concentrator-metadb       253:4    0 67.3T  0 lvm  /var/netwitness/concentrator/metadb
```

Configure Virtual Drive(s) on the Installed Meta Disk kit (Slots 4,5, and 6):

5. Create the Virtual Drive/Drive Group (DG) on the internal controller using the disks in slot 4 through 6 using the below command.

```
/opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6
strip=128
```

```
[root@S6Core2 ~]# /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6 strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.71.1.el7.x86_64
Controller = 0
Status = Success
Description = Add VD Succeeded.
```

6. The new virtual drive shows up as ‘2/2’ under DG/VD column.

```
/opt/MegaRAID/perccli/perccli64 /c0 /vall show
```

```
[root@S6Core2 ~]#  /opt/MegaRAID/perccli/perccli64 /c0 /vall show
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.71.1.el7.x86_64
Controller = 0
Status = Success
Description = None


Virtual Drives :
==============

--------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache Cac sCC       Size Name
--------------------------------------------------------------
0/0   RAID1 Optl  RW     Yes     RWBD  -   OFF 931.000 GB
1/1   RAID1 Optl  RW     Yes     RWBD  -   OFF   1.818 TB
2/2   RAID5 Optl  RW     No      RWBD  -   OFF   6.984 TB
--------------------------------------------------------------

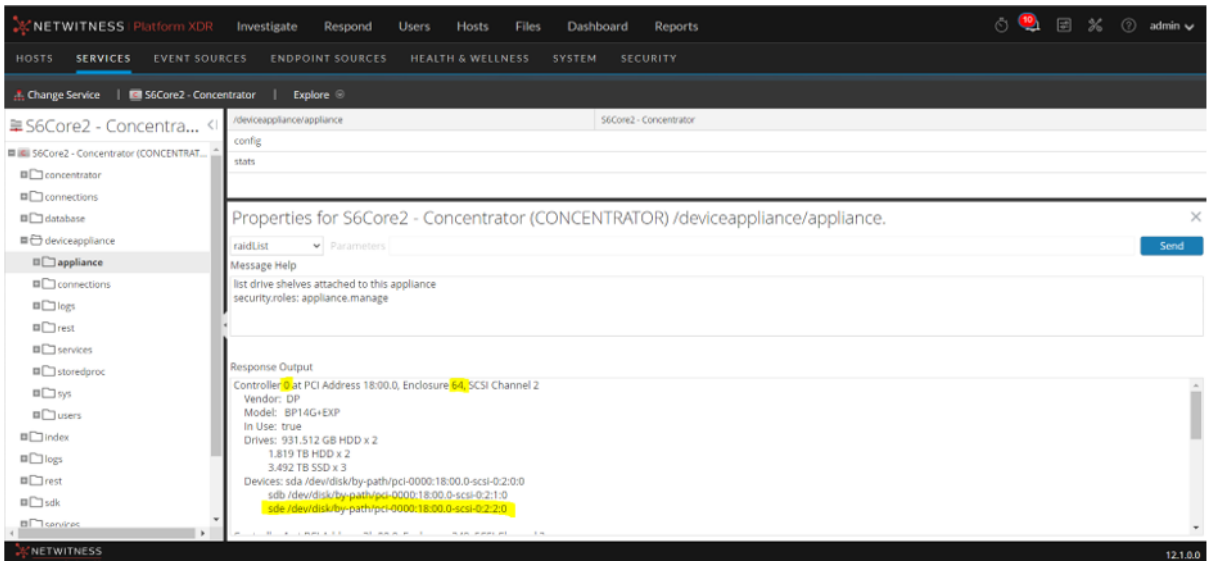
VD=Virtual Drive| DG=Drive Group|Rec=Recovery
Cac=CacheCade|Rec=Recovery|OfLn=OffLine|Pdgd=Partially Degraded|Dgrd=Degraded
Optl=Optimal|dflt=Default|RO=Read Only|RW=Read Write|HD=Hidden|TRANS=TransportReady
B=Blocked|Consist=Consistent|R=Read Ahead Always|NR=No Read Ahead|WB=WriteBack
FWB=Force WriteBack|WT=WriteThrough|C=Cached IO|D=Direct IO|sCC=Scheduled
Check Consistency
```

7. Identify the new block device on the host. The block device name is identified under the NAME column. In this example the new block device is 'sdc'. This block device name is required when configuring storage. Use `lsblk` to list the block devices.

```
[root@S6Core2 ~]# lsblk
NAME                          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                             8:0    0  931G  0 disk
├─sda1                          8:1    0    1G  0 part /boot
└─sda2                          8:2    0  930G  0 part
  ├─netwitness_vg00-root      253:0    0   30G  0 lvm  /
  ├─netwitness_vg00-swap      253:1    0    4G  0 lvm  [SWAP]
  ├─netwitness_vg00-nwhome    253:6    0  2.7T  0 lvm  /var/netwitness
  ├─netwitness_vg00-varlog    253:7    0   10G  0 lvm  /var/log
  └─netwitness_vg00-usrhome   253:8    0   10G  0 lvm  /home
sdb                             8:16   0  1.8T  0 disk
└─sdb1                          8:17   0  1.8T  0 part
  └─netwitness_vg00-nwhome    253:6    0  2.7T  0 lvm  /var/netwitness
sdc                             8:32   0  2.9T  0 disk
└─sdc1                          8:33   0  2.9T  0 part
  └─index-index               253:5    0  2.9T  0 lvm  /var/netwitness/concentrator/index
sdd                             8:48   0 74.9T  0 disk
└─sdd1                          8:49   0 74.9T  0 part
  ├─concentrator-root         253:2    0   30G  0 lvm  /var/netwitness/concentrator
  ├─concentrator-sessiondb    253:3    0  7.5T  0 lvm  /var/netwitness/concentrator/sessiondb
  └─concentrator-metadb       253:4    0 67.3T  0 lvm  /var/netwitness/concentrator/metadb
sde                             8:64   0    7T  0 disk
└─sde1                          8:65   0    7T  0 part
```

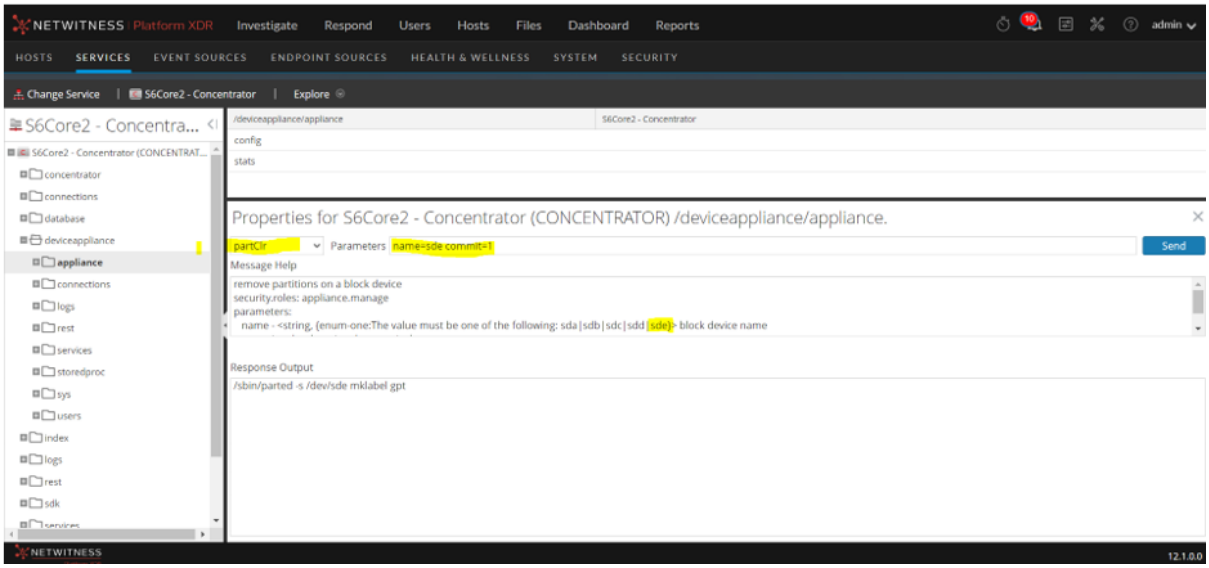
8. Now, you must Configure the above block device as index for Concentrator. Use the Explorer view (Login into UI, select Hosts->Service->Action->View->Explorer->deviceAppliance->appliance->properties) to complete configuration. Use ‘raidList’ to retrieve the existing raid list. Note the block device name corresponding to the Meta Disk kit. In this case it is ‘sde’.

Note: If a second Meta Disk kit is being configured, the block device would be ‘sdf’.



9. Clear partitions on the above block device (‘sde’) with the following parameters:

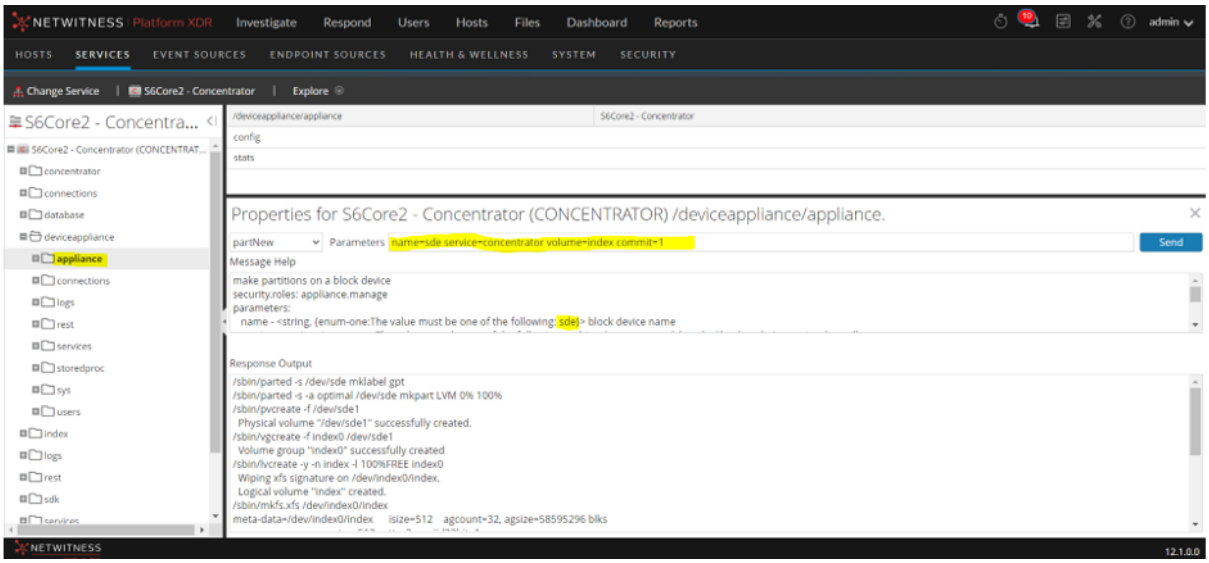
   name=sde commit=1



10. Make partitions on the block device (‘sde’) with the following parameters:

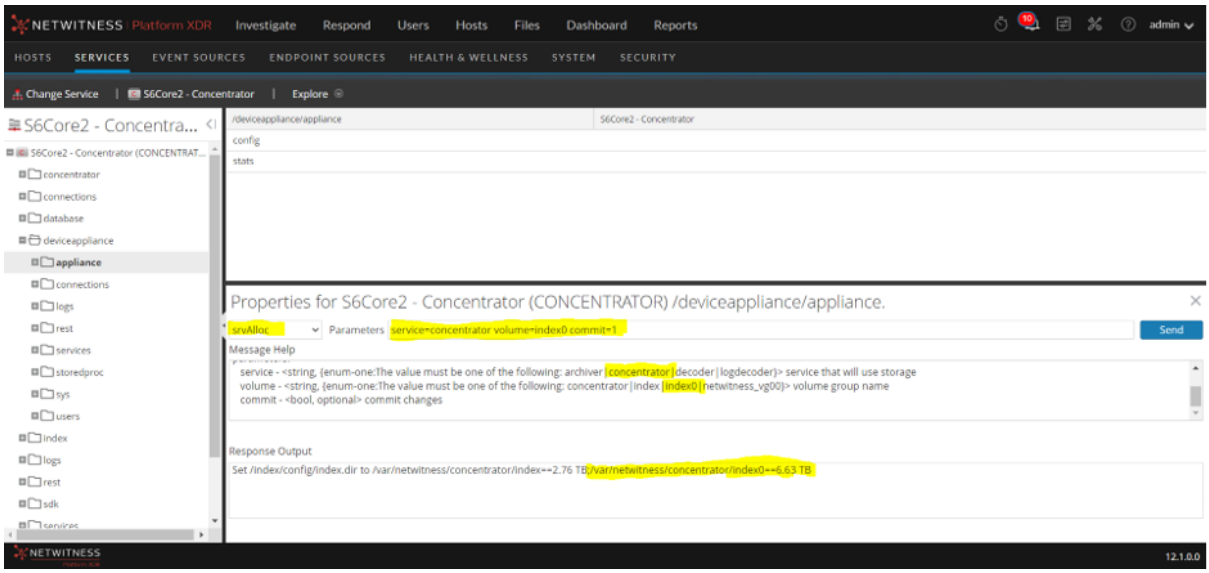
    name=sde service=concentrator volume=index commit=1

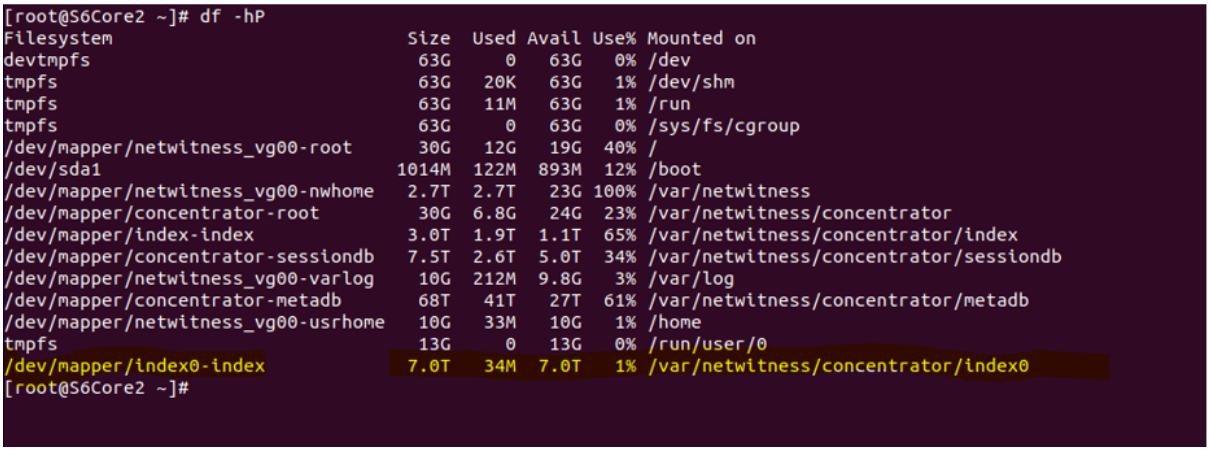
> **Note:** If a second Meta Disk kit is being configured, the volume would be ‘index1’.



11. Allocate the index to Concentrator service using ‘srvAlloc’

    service=concentrator volume=index commit=1




[root@S6Core2 ~]# df -hP
Filesystem Size Used Avail Use% Mounted on
devtmpfs 63G 0 63G 0% /dev
tmpfs 63G 20K 63G 1% /dev/shm
tmpfs 63G 11M 63G 1% /run
tmpfs 63G 0 63G 0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root 30G 12G 19G 40% /
/dev/sda1 1014M 122M 893M 12% /boot
/dev/mapper/netwitness_vg00-nwhome 2.7T 2.7T 23G 100% /var/netwitness
/dev/mapper/concentrator-root 30G 6.8G 24G 23% /var/netwitness/concentrator
/dev/mapper/index-index 3.0T 1.9T 1.1T 65% /var/netwitness/concentrator/index
/dev/mapper/concentrator-sessiondb 7.5T 2.6T 5.0T 34% /var/netwitness/concentrator/sessiondb
/dev/mapper/netwitness_vg00-varlog 10G 212M 9.8G 3% /var/log
/dev/mapper/concentrator-metadb 68T 41T 27T 61% /var/netwitness/concentrator/metadb
/dev/mapper/netwitness_vg00-usrhome 10G 33M 10G 1% /home
tmpfs 13G 0 13G 0% /run/user/0
/dev/mapper/index0-index 7.0T 34M 7.0T 1% /var/netwitness/concentrator/index0
[root@S6Core2 ~]#

Revision History

Revision	Date	Description
1.0	August, 2022	12.0 Release Updates

NetWitness® Platform XDR

Version 12.0

System Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

August, 2022

Contents

System Configuration Overview **9**

Standard Procedures **10**

Access System Settings 11

Configure Notification Servers 12

Notification Servers Overview 12

Configure the Email Settings as Notification Server 13

Configure Script as a Notification Server 14

Configure the SNMP Settings as Notification Server 15

Configure a Syslog Notification Server 16

Configure Notification Outputs 18

Notification Outputs Overview 18

Configure Email as a Notification 19

Configure Script as a Notification 20

Configure SNMP as a Notification 21

Configure Syslog as a Notification 22

Configure Templates for Notifications 24

Configure Global Notifications Templates 24

Add a Template 25

Duplicate a Template 26

Edit a Template 27

Delete a Template 27

Define a Template for ESA Alert Notifications 27

ESA Data Model 27

Import and Export a Global Notifications Template 29

Import a Template 29

Export a Template 30

Configure Email Servers and Notification Accounts 30

Configure Global Audit Logging 32

Global Audit Logging - High-Level Procedure 33

Configure a Destination to Receive Global Audit Logs 34

Configure a Syslog Notification Server for a Third-Party Syslog Server 35

Configure a Syslog Notification Server for a Log Decoder 36

Next Steps 38

Define a Template for Global Audit Logging 38

Define a Global Audit Logging Template for a Log Decoder 38

Define a Custom Global Audit Logging Template 40

Next Step ... 42
Define a Global Audit Logging Configuration ... 43
Prerequisites ... 43
Add a Global Audit Logging Configuration ... 43
Edit a Global Audit Logging Configuration ... 44
Delete a Global Audit Logging Configuration ... 45
Verify Global Audit Logs ... 45
Example CEF Output ... 47
Configure Centralized Audit Logging ... 47
Filtering the Aggregated Logs ... 48
Log Retention Policy ... 48
Disable Centralized Audit Logging ... 49
Configure Investigation Settings ... 49
Map Context Hub Meta Types ... 49
Configure Common Settings ... 49
Configure Navigate and Legacy Events View Settings ... 50
Clear Reconstruction Cache for Services ... 52
Configure Events View Settings ... 52
Configure the Sync Core Timeout to Remedy Deadlocks in Events View Reconstructions ... 54
Configure Live Services Settings ... 55
Prerequisite ... 55
About Live Feedback Participation ... 55
Access the Live Services Configuration Panel ... 56
Configure Live Account ... 57
Configure the Live Content Synchronization Interval and Notification ... 58
Force Immediate Synchronization ... 59
Live Feedback Overview ... 60
About Live Feedback Participation ... 60
JSON File ... 60
Upload Data to RSA for Live Feedback ... 67
Download Live Feedback Historical Data ... 68
Share Data with RSA ... 68
Configure Log File Settings ... 69
Configure System Log File Size and Backup Count ... 69
Set the Log Level for an Individual Package ... 70
Configure Syslog and SNMP Settings ... 71
Configure and Enable Syslog Settings ... 71
Configure and Enable SNMP Settings ... 72
Disable Syslog or SNMP Settings ... 73

Additional Procedures **74**
Add Custom Context Menu Actions 74
View Context Menu Actions in NetWitness 75
Add a Context Menu Action 75
Edit a Context Action 77
Delete a Context Action 78
Export Context Menu Actions 78
Import Context Menu Actions 79
Example Procedure: Context Menu Action to Investigate ip.dst from alias.ip 81
Configure NTP Servers 83
Add an NTP Server 83
Modify an NTP Server 84
Configure Proxy for NetWitness Platform 86
Import Certificates for HTTPS Service 87
Troubleshoot System Configuration **88**
Troubleshoot Global Audit Logging 88
Basic Troubleshooting 88
Advanced Troubleshooting 88
Verify the Packages and Services on the Hosts 89
Possible Issues 89
Possible Solutions 90
Solution Examples 94
Troubleshoot Issues identified in the NTP Settings Panel or Log Files Messages 97
Troubleshoot Global Notifications 98
References **99**
Global Audit Logging Configurations Panel 100
Add New Configuration Dialog 103
Features 103
User Actions Logged 104
Supported CEF Meta Keys 112
Supported Common Event Format (CEF) Meta Keys 112
Supported Global Audit Logging Meta Key Variables 117
Supported Global Audit Logging Meta Key Variables 117
Global Audit Logging Operation Reference 119
CARLOS 119
ESA 119
Investigation 120
Reporting Engine 122
Warehouse Connector 123
Health & Wellness 123

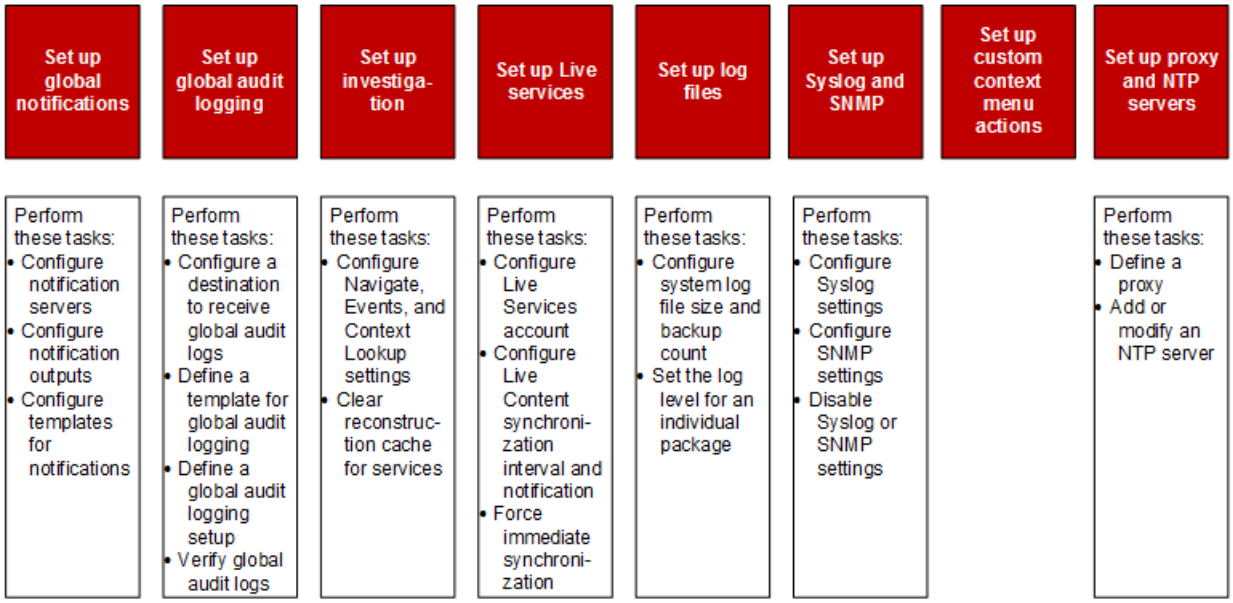
NetWitness Core Services 124
Malware Analysis 127
NetWitness User Interface 130
Respond 134
Investigate Server 134
Security Server 136
Admin Server 136
Config Server 137
Context Hub Server 137
Local Audit Log Locations 138
Global Notifications Panel 140
WorkFlow 140
What do you want to do? 140
Related Topics 140
Quick Look 141
Toolbar and Features 141
Global Notifications Panel Toolbar 142
Define Notification Server Dialogs 145
Email 145
SNMP 147
Syslog 150
Script 152
Define Notification Output Dialogs 153
Features 153
Email 153
SNMP 154
Syslog 155
Script 157
Define Notification Template Dialog 159
Output Tab 163
Workflow 163
What do you want to do? 163
Related Topics 164
Quick Look 164
Servers Tab 166
Workflow 166
What do you want to do? 166
Related Topics 166
Quick Look 167
Templates Tab 169
Workflow 169

What do you want to do? 169
Related Topics 169
Quick look 169
HTTP Proxy Settings Panel 171
Related topics 171
Quick Look 171
Email Configuration Panel 173
Workflow 173
What do you want to do? 173
Related Topics 173
Quick Look 174
Email Server Settings 174
Email Statistics 175
Investigation Configuration Panel 176
Workflow 176
What do you want to do? 176
Related Topics 176
Quick Look 177
Common Settings Tab 177
Investigate Tab 178
Render Threads Setting 178
Parallel Coordinates Settings 178
Legacy Events Tab 179
Enable Legacy Events 180
Event Search Settings 180
Reconstruction Settings 180
Web View Reconstruction Settings 182
Reconstruction Cache Settings 184
Context Lookup Tab 185
Events Tab 186
Live Services Configuration Panel 187
Workflow 187
What do you want to do? 187
Related Topics 188
Live Services Quick Look 188
Live Account Section 189
Live Content Section 190
Force Immediate Synchronization 191
Additional Live Services 192
About Live Feedback Participation 193
NTP Settings Panel 194

Workflow 194
What you need to do? 194
Related Topics 194
Quick Look 194
Context Menu Actions Panel 196
Workflow 196
What do you want to do? 196
Quick Look 196
Legacy Notifications Configuration Panel 199
Workflow 199
What do you want to do? 199
Related Topics 199
Quick Look 200
Syslog Settings 200
SNMP Settings 201

System Configuration Overview

In the Administration System view, administrators can configure system settings to receive optimal performance from NetWitness. This diagram show the available configuration options.



In this guide, the standard procedures provide instructions for administrators who want to customize settings that apply across the system in NetWitness. Although some of these settings have default values, the administrator needs to view and evaluate all default values.

Additional procedures are not essential for the set up of NetWitness, they include certain customization options that are beyond the usual setup; for example, adding custom context menus or setting up a proxy.

In addition, reference topics and troubleshooting topics supply detailed information about the user interface and suggestions for resolving possible issues.

The following sections describe system configuration:

- Standard Procedures provide instructions for administrators who want to customize settings that apply across the system in NetWitness.
- Additional Procedures provide instructions for setting up customization options that are beyond the usual system configuration.

Standard Procedures

The topics in this section provide instructions for administrators who want to customize settings that apply across the system in NetWitness. Although some of these settings have default values, the administrator needs to view and evaluate all default values. The procedures can be performed in any sequence and are listed alphabetically.

Access System Settings

Configure Notification Servers

Configure Notification Outputs

Configure Templates for Notifications

Configure the Email Settings as Notification Server

Configure Email Servers and Notification Accounts

Configure Global Audit Logging

Configure Investigation Settings

Configure Live Services Settings

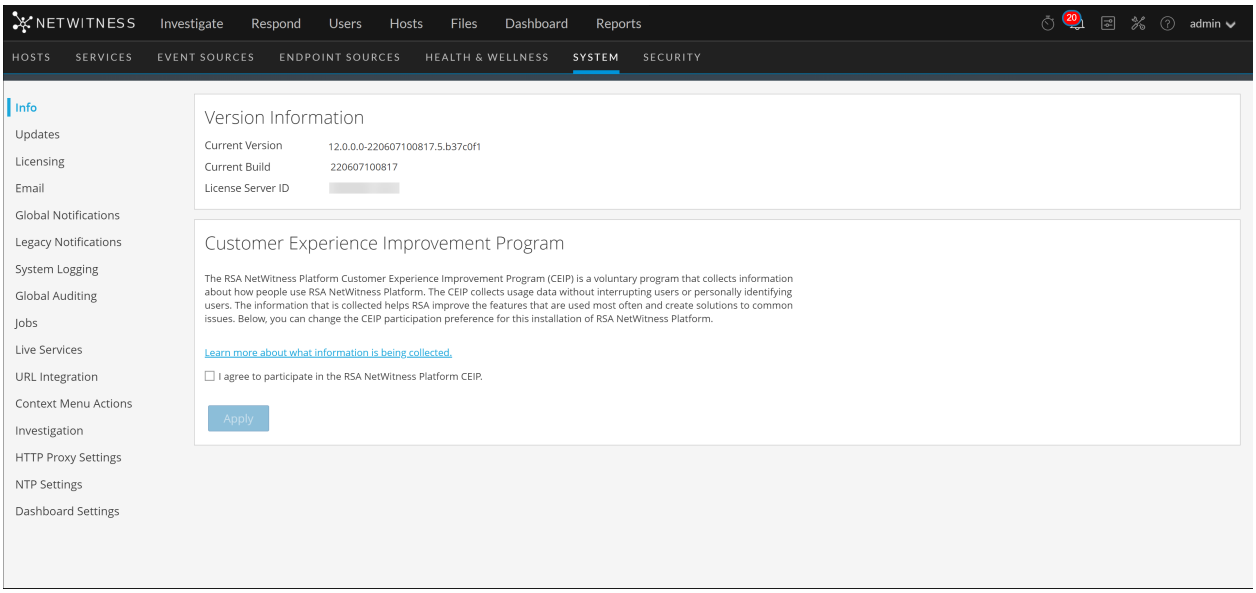
Configure Log File Settings

Access System Settings

This topic introduces the system configuration capabilities of NetWitness in the Administration System view. Administrators can configure notifications, email notifications, global audit logging, logging settings, connection to Live Services, and URL integration in NetWitness.

To access the system settings:

Go to **(Admin) > System**.
The Administration System view is displayed.



On the left panel of the Administration System view is an options panel listing all system nodes available for configuration. When you select a node, the associated content is displayed in the right panel.

Configure Notification Servers

This topic provides instructions on how to configure notification servers. For ESA, notification servers are required to define an ESA rule. A notification server is also required to configure global audit logging.

Global Notifications configurations define notifications settings for Event Source Management (ESM), Health and Wellness, New Health and Wellness, Global Audit Logging, Event Stream Analysis (ESA), and Respond. Notification Servers define the servers from which you want to receive notifications from the system. For Global Audit Logging, define Log Decoders as Syslog Notification Servers.

You can define, delete, edit, import, and export a notification server in NetWitness. Individual topics describe the relevant procedures. For more information on ESA alert configuration, see "Notification Methods" in the *Alerting with ESA Correlation Rules User Guide*. You delete, edit, import, and export notification outputs and notification servers in the same way as templates. You cannot disable or delete notification servers associated with global audit logging configurations.

Notification Servers Overview

This topic provides an overview of notification servers. You configure notification servers in the Administration System view ( (Admin) > System > Global Notifications > Servers tab).

Global Notifications are used by a variety of components in NetWitness, such as Event Stream Analysis (ESA), Respond, Health and Wellness, New Health and Wellness, Event Source Management (ESM), and Global Audit Logging. Notification settings are called **Notification Servers**.

Event Stream Analysis sends notifications to users through email, SNMP, or Syslog about various system events. In ESA, these alert notification settings are called Notification Servers. You can configure multiple notification servers and use them while defining an ESA rule, for example, you can configure multiple mail servers or Syslog servers and use the settings while defining an ESA rule.

Note: New Health and Wellness supports only Email and Syslog notifications.

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

You can configure the following notification servers:

- Email
- SNMP
- Syslog
- Script

Email notification servers enable you to configure email server settings to send alert notifications. SNMP notification servers enable you to configure SNMP trap host settings as a notification server to send alert notifications.

Syslog notification servers enable you to configure Syslog settings as a notification server to send notifications. When enabled, Syslog provides auditing through the use of the RFC 5424 Syslog protocol. Syslog has proven to be an effective format to consolidate logs, as there are many open source and proprietary tools for reporting and analysis. For Global Audit Logging, you can only use Syslog Notification Servers.

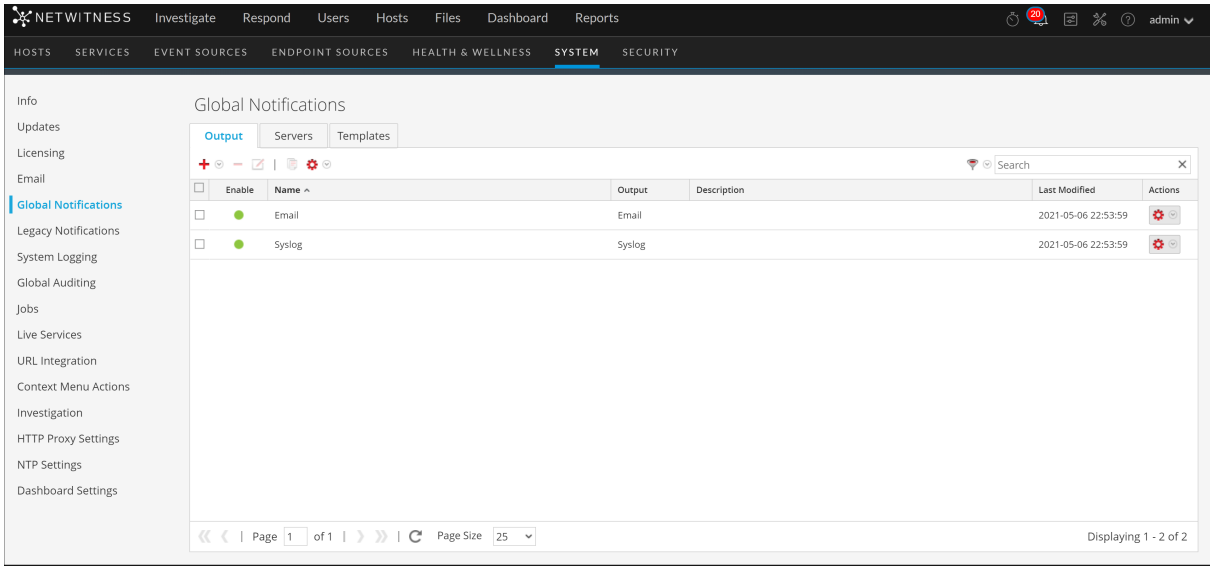
Script notification servers enable you to configure Script as a notification server.

For detailed information on the different notification server configurations, including parameters and descriptions, see Define Notification Server Dialogs.

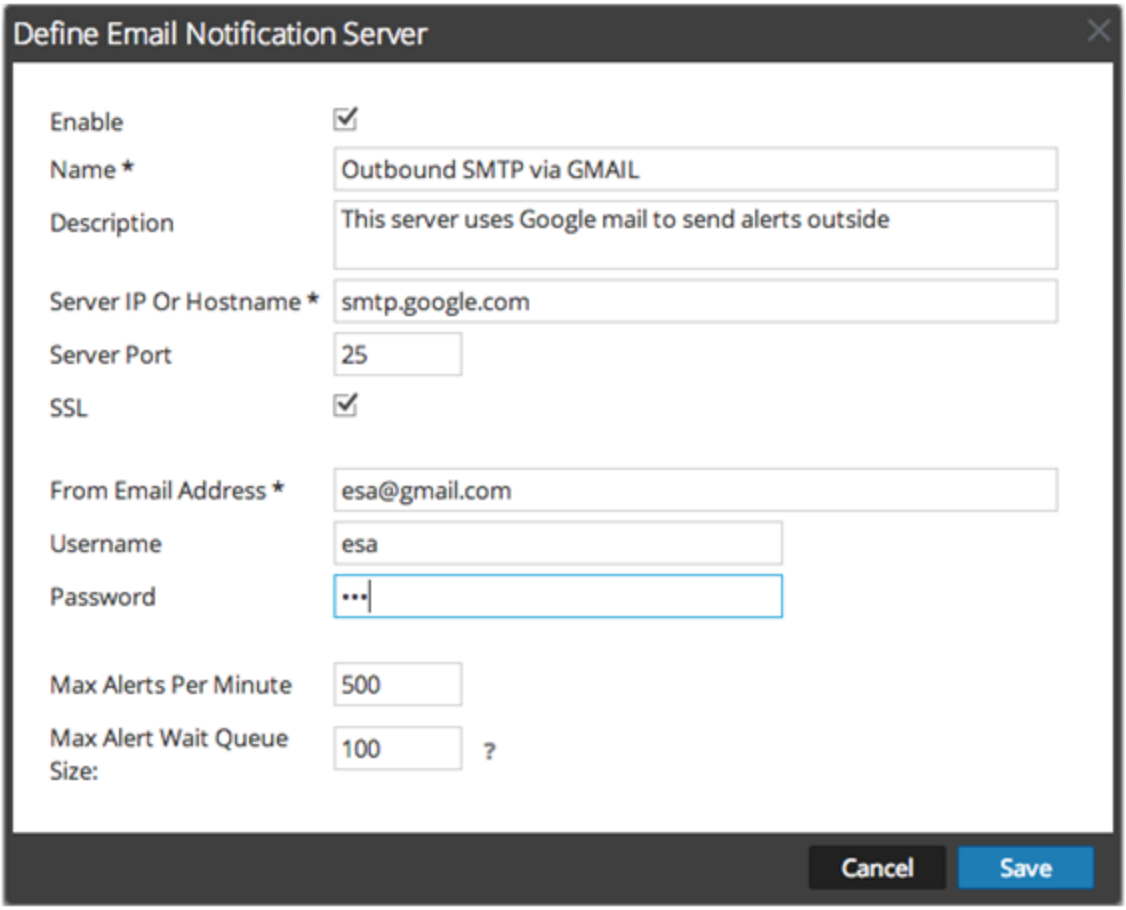
Configure the Email Settings as Notification Server

To configure email server settings as a notification server to send alert notifications:

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
   The **Notifications** configuration panel is displayed with the **Output** tab open.
3. Click the **Servers** tab.



4. From the  drop-down menu, select **Email**.



5. In the **Define Email Notification Server** dialog, provide the required information and click **Save**.

Note: For ESM/SMS and ESA notifications, you must specify only the hostname/FQDN in the Server IP or Hostname field.

For details of the parameters and descriptions, see Define Notification Server Dialogs

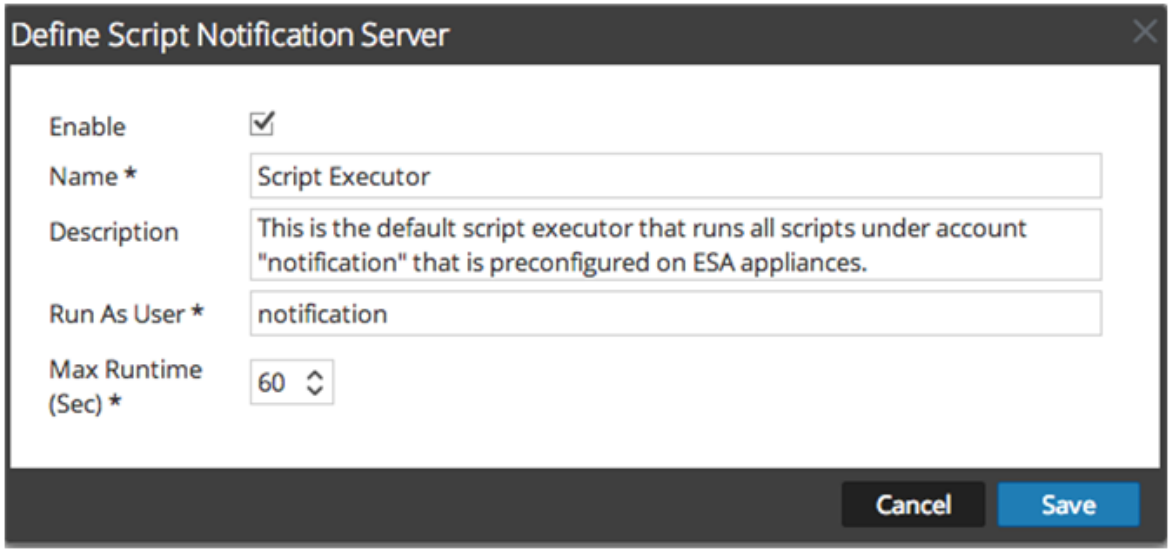
Configure Script as a Notification Server

ESA allows you to run scripts in response to ESA alerts. However, you must first configure the user identity and other details that are required to run the scripts.

To configure Script as a notification server:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Servers** tab.

4. From the  drop-down menu, select **Script**.



5. In the **Define Script Notification Server** dialog, provide the required information and click **Save**.

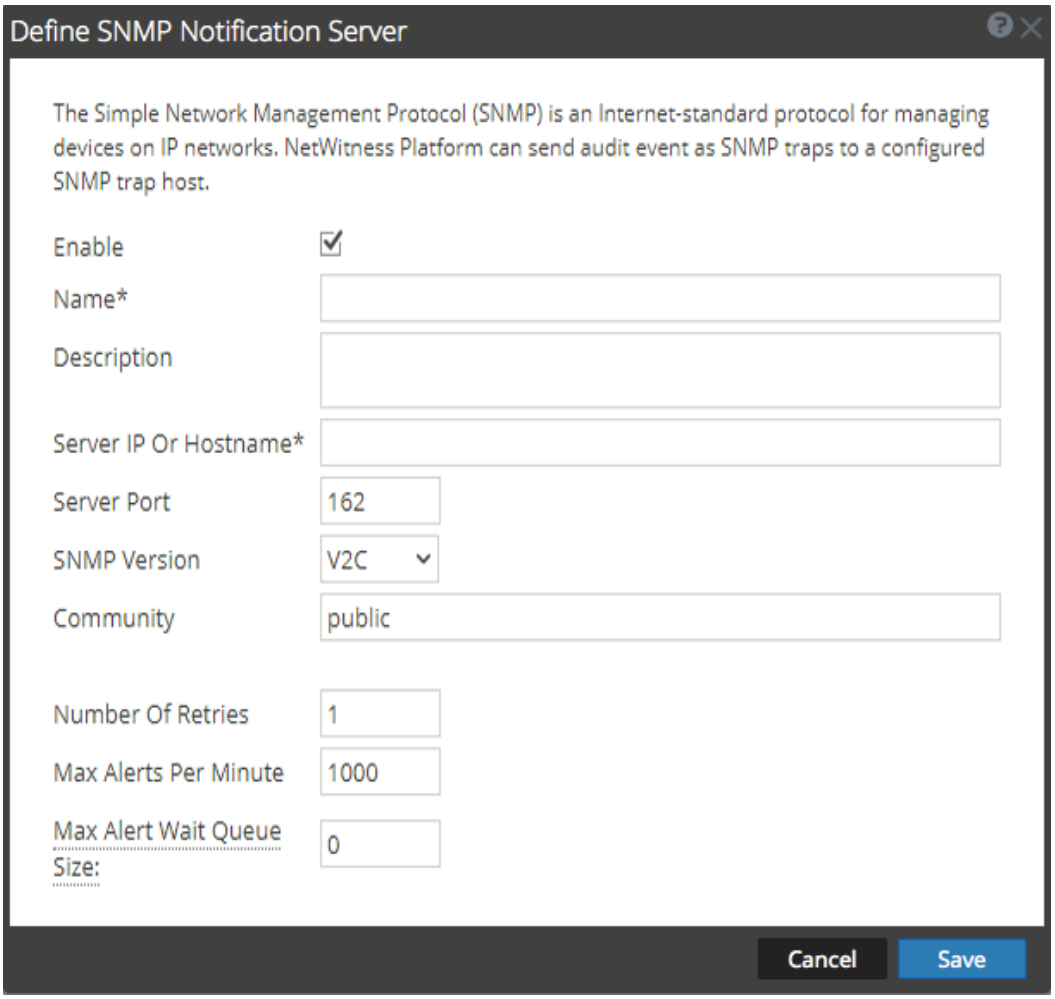
For details of the parameters and descriptions, see Define Notification Server Dialogs.

Configure the SNMP Settings as Notification Server

To configure the SNMP trap host settings as a notification server to send alert notifications:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Servers** tab.

4. From the  drop-down menu, select **SNMP**.



5. In the **Define SNMP Notification Server** dialog, provide the required information and click **Save**.

For details of the parameters and descriptions, see Define Notification Server Dialogs.

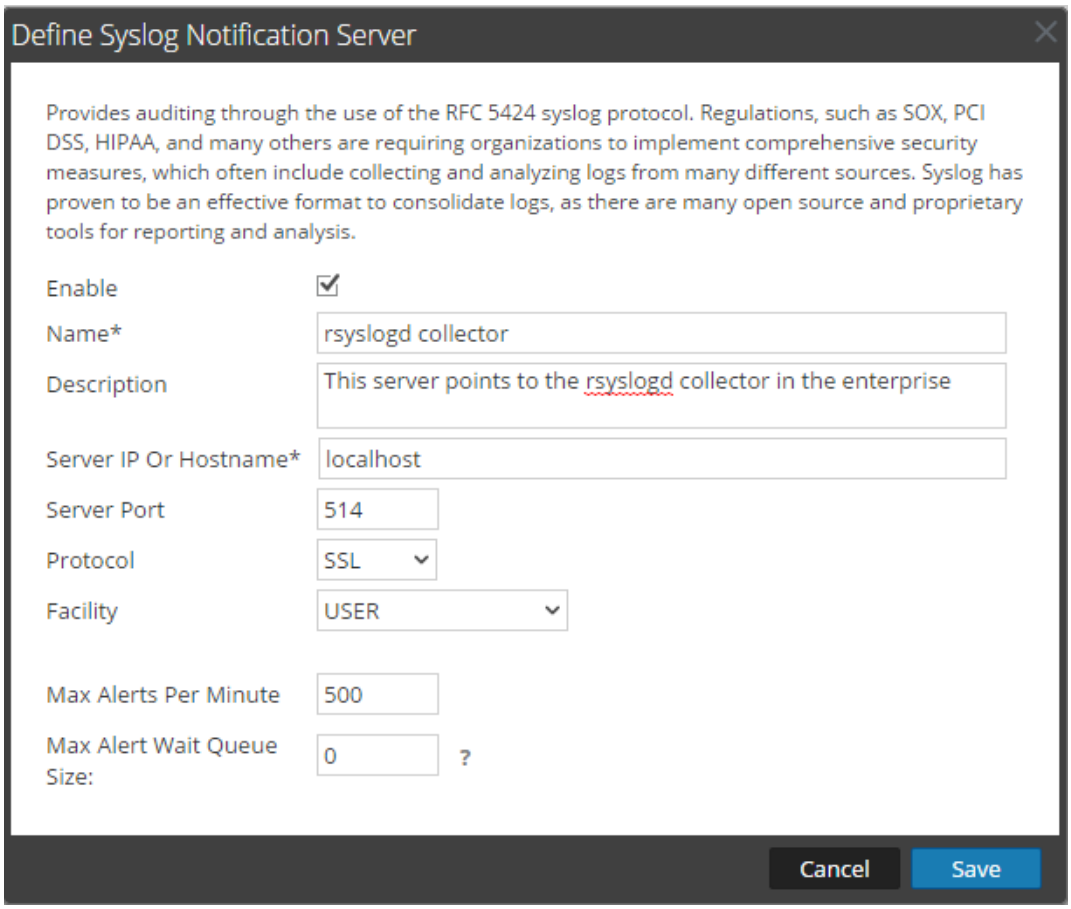
Configure a Syslog Notification Server

This topic provides instructions on how to configure a Syslog notification server. When enabled, Syslog provides auditing through the use of the RFC 5424 Syslog protocol. Syslog has proven to be an effective format to consolidate logs, as there are many open source and proprietary tools for reporting and analysis.

To configure Syslog as a notification server:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Servers** tab.

4. From the drop-down menu, select **Syslog**.



5. In the **Define Syslog Notification Server** dialog, provide the required information and click **Save**.

For details of the parameters and descriptions, see Define Notification Server Dialogs.

Configure Notification Outputs

This topic provides instructions on how configure notification outputs. These notification outputs are required to define an ESA rule.

Global Notifications configurations define notifications settings for Event Source Management (ESM), Health and Wellness, New Health and Wellness, Global Audit Logging, Event Stream Analysis (ESA), and Respond.

Note: You do not need to configure the Output tab for Global Audit Logging.

Notification Output configurations define email addresses and subject lines, SNMP trap OID settings, syslog output settings, and script code.

You can define, delete, edit, import, and export notification outputs in NetWitness. Individual topics describe the relevant procedures. For more information on ESA alert configuration, see "Notification Methods." in the *Alerting with ESA Correlation Rules User Guide*. You delete, edit, import, and export notification outputs and notification servers in the same way as templates. If you attempt to delete a notification output being used by alerts, you will receive a warning confirmation message that the alerts using the notification will not function properly. The message shows the number of alerts in use.

Notification Outputs Overview

This topic provides an overview of notification outputs. notification outputs are required when defining an ESA rule. You can configure notification outputs in the Administration System view ((Admin) > System > Global Notifications > Outputs tab).

Global Notifications configurations define notifications settings for Event Source Management (ESM), Health and Wellness, New Health and Wellness, Global Audit Logging, Event Stream Analysis (ESA), and Respond.

Note: You do not need to configure notification outputs (the Output tab) for Global Audit Logging.

Notification outputs are the destinations used for sending notifications. For ESA, notification outputs enable you to define how you want to receive the ESA alerts. The following are the different notification outputs supported by NetWitness:

- Email
- SNMP
- Syslog
- Script

Note: New Health and Wellness supports only Email and Syslog notification outputs.

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

Email notification settings define the destination email address to which you can send the alerts. You can also add a custom description in the subject of the email and define multiple destination email addresses.

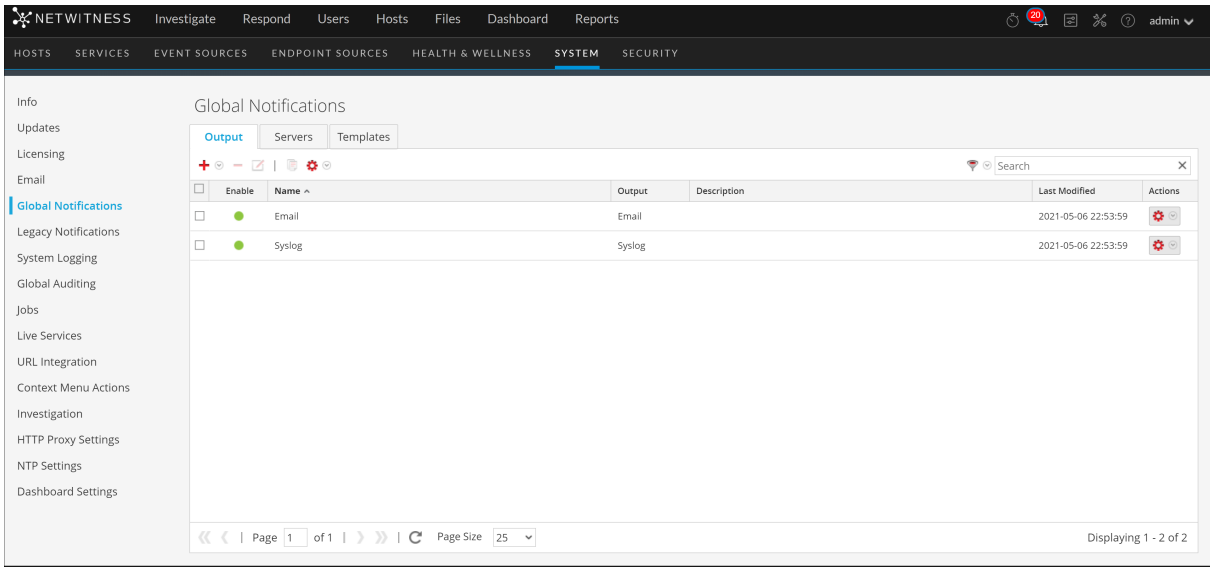
SNMP notification settings enable you to define the SNMP settings to send alert notifications. Syslog notifications enable you to define the Syslog settings used to send alert notifications. Script notifications enable you to define the Script that executes in response to the alert.

For detailed information on the notification configurations, including parameters and descriptions, see Define Notification Server Dialogs.

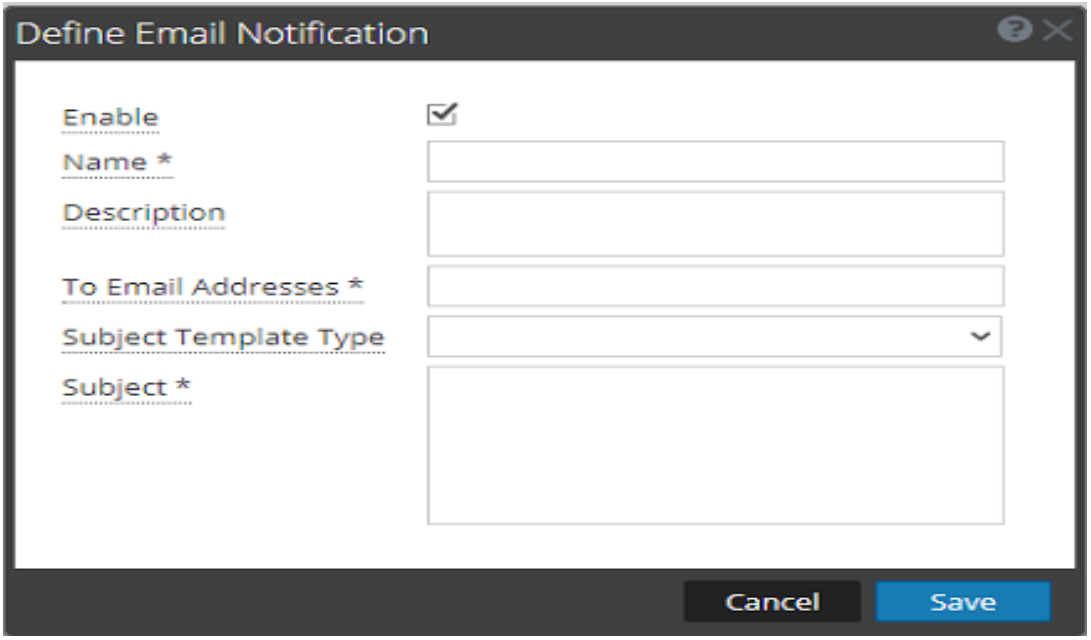
Configure Email as a Notification

To configure Email as a notification:

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.



3. On the **Output** tab, from the + drop-down menu, select **Email**.



4. In the **Define Email Notification** dialog, provide the required information and click **Save**.

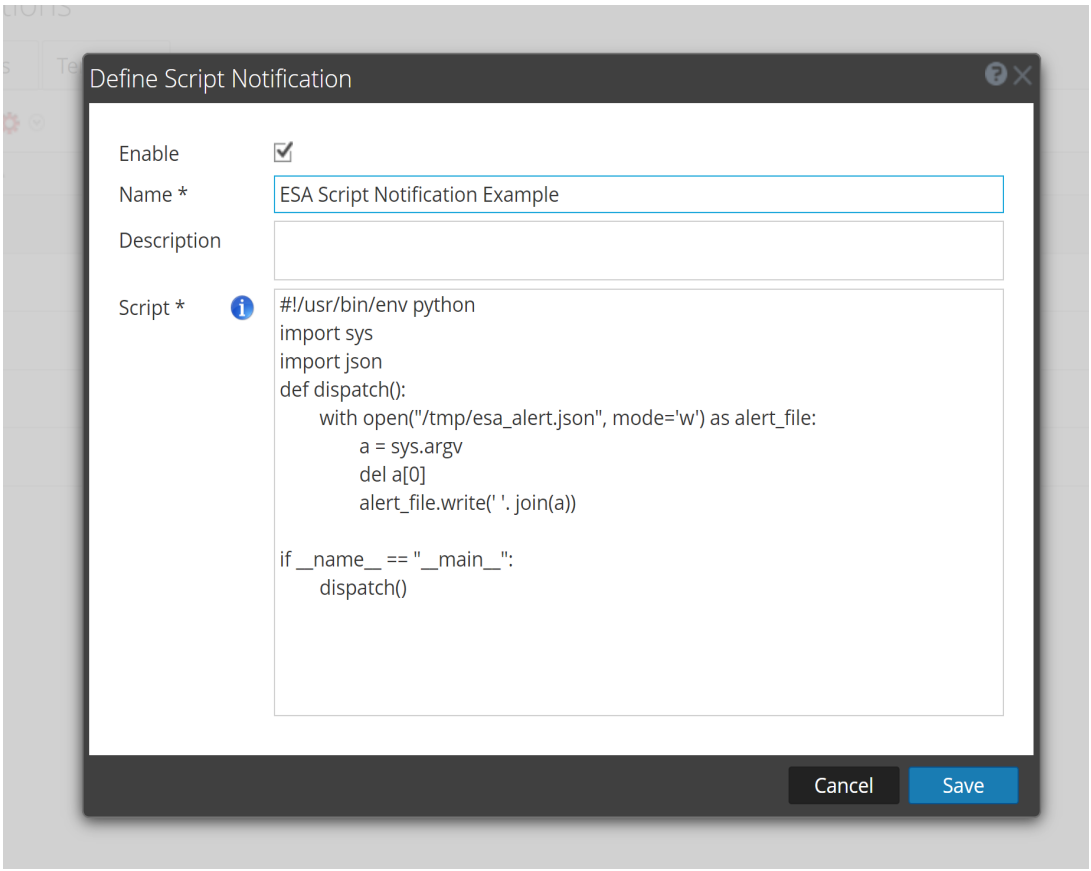
For details of the parameters and descriptions, see Define Notification Server Dialogs.

Configure Script as a Notification

This topic provides instructions to define and configure a Script as a notification output. ESA allows you to run scripts in response to ESA alerts. You need to define the script using the (Admin) > System > Notifications > Output tab. You can use any script for ESA notifications.

To configure the script as a notification:

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. On the Output tab, from the + drop-down menu, select **Script**.



4. In the **Define Script Notification** dialog, provide the required information and click **Save**.

Note: To retrieve alerts information in the scripts, use command line arguments based on the scripting language. For Example:
- If you are using Python as the scripting language, use sys.argv (command line arguments) to retrieve alerts information.
- If you are using Bash as the scripting language, use $*, $1, $2, and $@ (command line arguments) to retrieve alerts information.

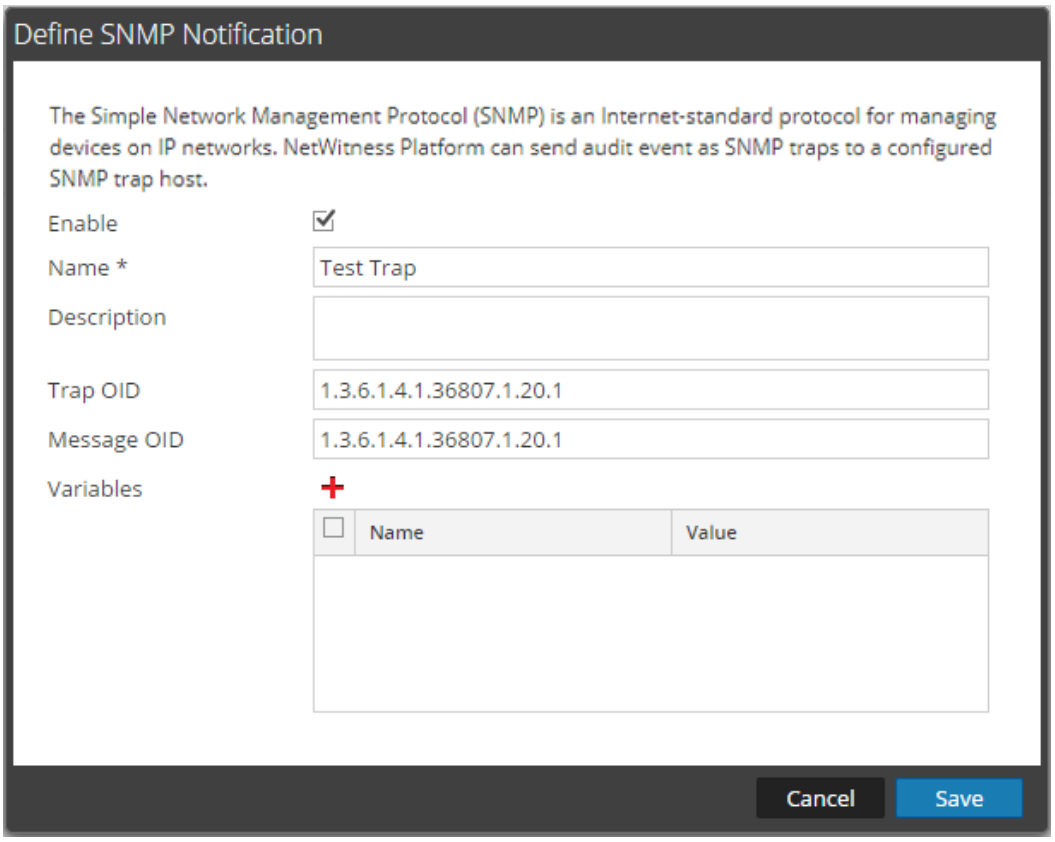
Note: Use **Temp** folder to create files or folders as a part of the script.

For details of the parameters and descriptions, see Define Notification Server Dialogs.

Configure SNMP as a Notification

To configure SNMP as a notification output to send alert notifications:

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. On the Output tab, from the drop-down menu, select **SNMP**.



4. In the SNMP Notification dialog, provide the required information and click **Save**.

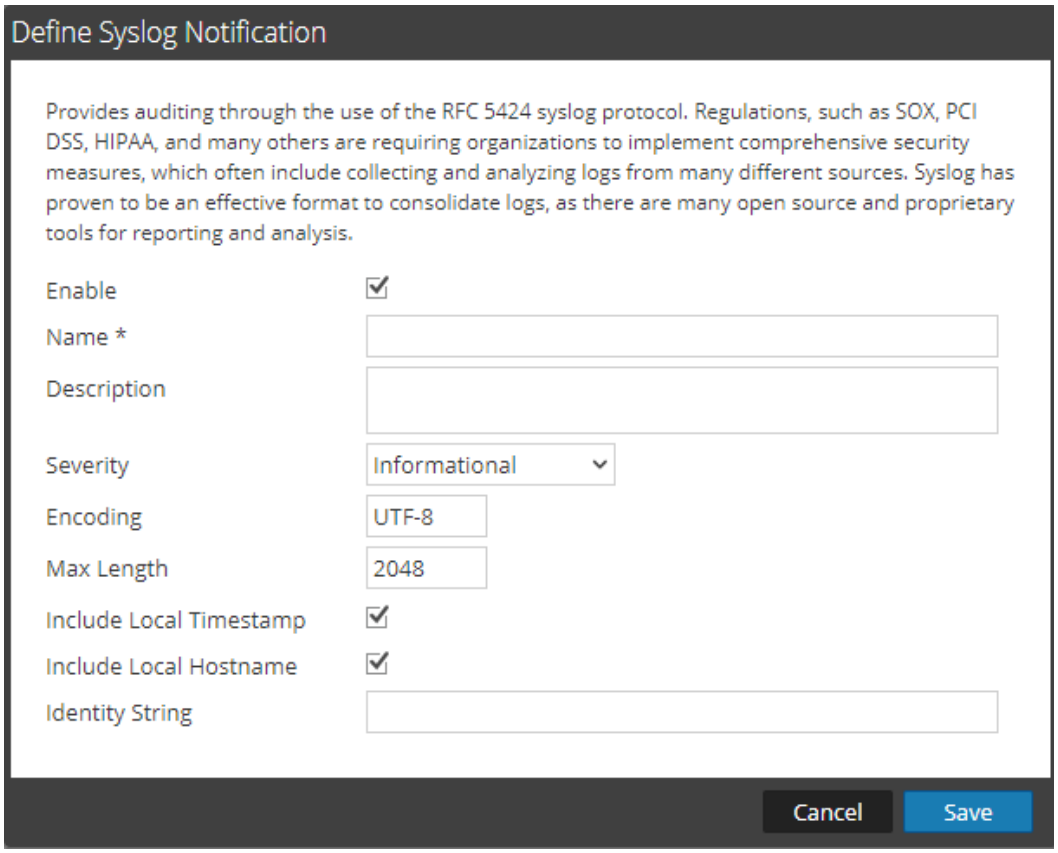
For details of the parameters and descriptions, see Define Notification Server Dialogs.

Configure Syslog as a Notification

To configure Syslog as a notification output when sending alert notifications:

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.

3. On the Output tab, from the drop-down menu, select **Syslog**.



4. In the **Define Syslog Notification** dialog, provide the required information and click **Save**.

For details of the parameters and descriptions, see Define Notification Server Dialogs.

Configure Templates for Notifications

You configure notification templates in the Administration System view ((Admin) > System > Global Notifications > Templates tab). A notification template defines the format and message fields of the notifications. There are different template types for the notifications that you can configure:

- Audit Logging
- Event Stream Analysis
- Event Source Monitoring
- Health Alarms
- New Health & Wellness Alarms(Version 11.5 or Later).

You can use the available default templates or you can configure your own templates for Email (SMTP), SNMP, Syslog, and Script, depending on the template type. To learn how to modify or configure your own template, see Define Notification Template Dialog.

Global audit logging sends audit logs in the format specified in the Audit Logging template. You can use the default audit logging templates or you can define your own audit logging template. For more information on how to define an Audit Logging template, see Define a Template for Global Audit Logging.

Event Stream Analysis (ESA) sends notifications in the format specified in the Event Stream Analysis templates. The default Event Stream Analysis templates for email, SNMP, Syslog, and Script are available on installation. You can customize these templates as well as create new templates which you can use for the notifications. For more information on how to define ESA templates, see Define a Template for ESA Alert Notifications.

For more information on ESA alert configuration, see "Notification Methods" in the *Alerting with ESA Correlation Rules User Guide*. You cannot delete templates associated with global audit log configurations.

To learn how to define, delete, edit, duplicate, import, and export a notification template in NetWitness, see:

Configure Global Notifications Templates

Define a Template for ESA Alert Notifications

Import and Export a Global Notifications Template

Configure Global Notifications Templates

This topic provides instructions for adding, editing, duplicating, and deleting global notifications templates.

Note: New Health and Wellness supports only Email and Syslog notification outputs.

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

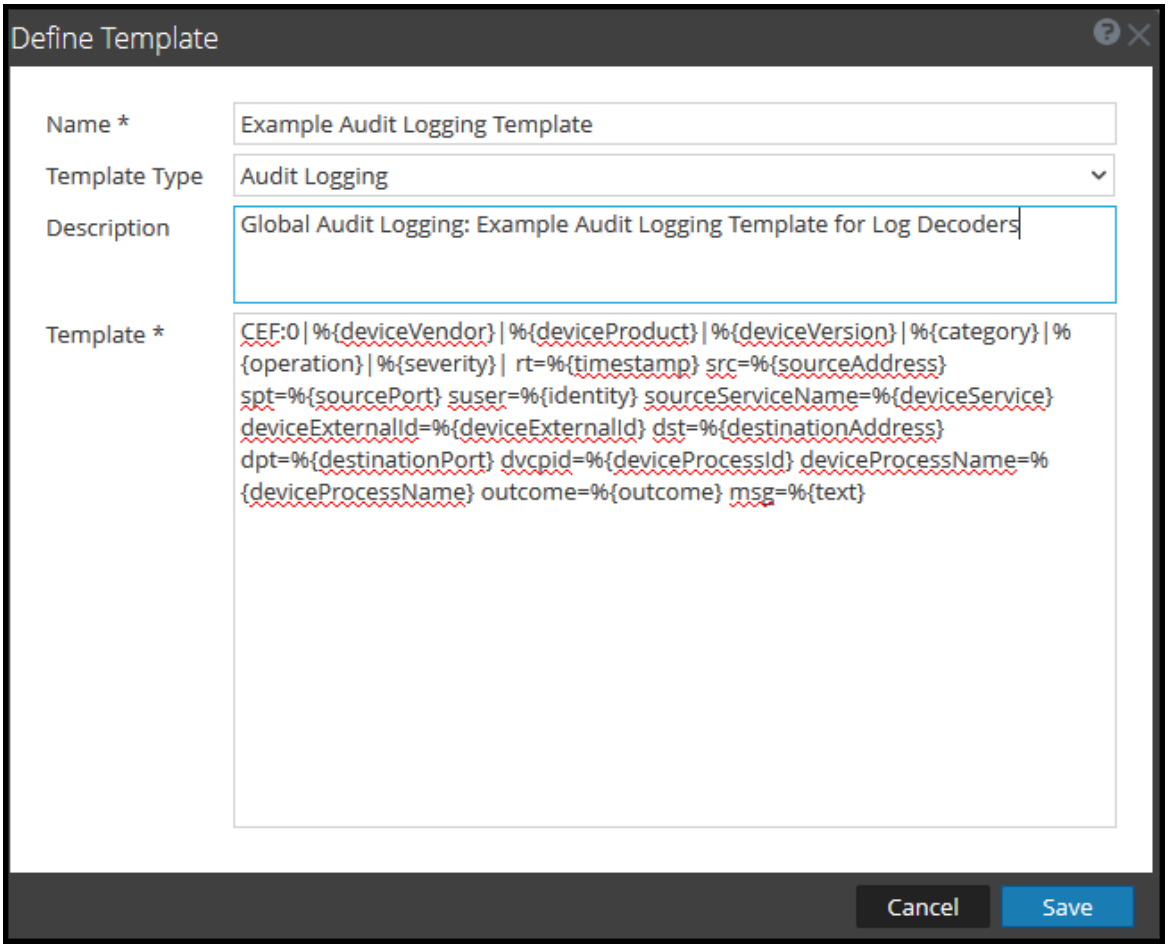
When upgrading from version 10.4, all existing notification templates migrate to the Event Stream Analysis template type.

Add a Template

You can use the default templates provided or you can configure your own templates. To configure your own template:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Templates** tab.
4. Click  to configure a template.
5. In the **Define Template** dialog, provide the following information:
   a. In the **Name** field, type the name for the template.
   b. In the **Template Type** field, select the type of template you want to create. For example, if you are creating a template for global audit logging, select the Audit Logging template type. For New Health and Wellness, select New Health & Wellness Alarms template type.
   c. In the **Description** field, type a brief description for the template.
   d. In the **Template** field, specify the format for the template, see Define Notification Template Dialog.

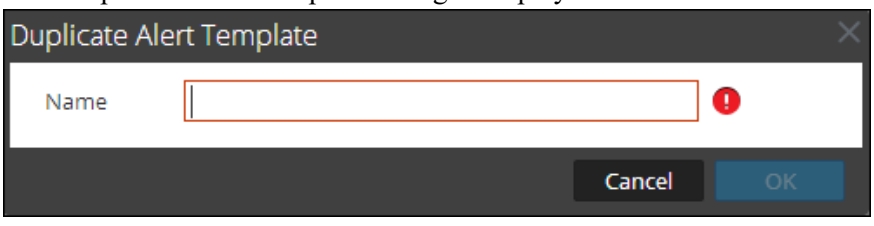
e. Click **Save** to save the template.
The following is an example of a template for Audit Logging.



Duplicate a Template

You can make a copy of an existing default or user-defined template. To duplicate a template:

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Templates** tab.
4. Select the template that you want to duplicate and click .
The Duplicate Alert Template dialog is displayed.



5. Type the name for the duplicate template.
6. Click **OK**.

You can modify a default or user-defined template. When you edit a template, the changes are reflected only when the alert is triggered.

Edit a Template

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Templates** tab.
4. Select a template and click .
5. In the **Define Template** dialog, modify the **Name**, **Template Type**, **Description**, and **Template** fields as required.
6. Click **Save** to save the template.

Delete a Template

You can delete a user-defined template. When you delete a template that is used in an ESA rule, the Event Stream Analysis default template is used for alerts. You cannot delete templates associated with global audit logging configurations.

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Templates** tab.
4. Select one or more templates and click .
   A confirmation dialog is displayed.
5. Click **Yes**.
   The selected template is deleted.

Define a Template for ESA Alert Notifications

This topic describes how you can define a template for alert notifications. Event Stream Analysis (ESA) allows you to define useful templates for alerts. You need to have a good understanding of FreeMarker and the ESA data model to define a template. For more information on FreeMarker, see FreeMarker Template Author's Guide.

ESA Data Model

Consider an ESA alert rule as shown below:

```
@Name('module_144d43f5_f0b4_4cd0_8c6c_5ce65c37e624_Alert')
@Description('Brute Force Login To Same Destination')
@RSAAlert(oneInSeconds=0, identifiers={"ip_dst"})
SELECT* FROMEvent (ec_activity = 'Logon',ec_theme = 'Authentication',ec
outcome = 'Failure',ip_dst IS NOT NULL)
.std:groupwin(ip_dst)
.win:time_length_batch(60 seconds, 2)
GROUPBYip_dst HAVING COUNT(*) = 2;
```

When a rule like the above is fired, the alert generated has two constituent events, each resembling a NextGen session with multiple meta values. The alert data-object passed to the FreeMarker template evaluator are as follows:

```
(root)
  |
  +- id = "4e67012f-9c53-4f0b-ac44-753e2c982b79"                    // Unique identifier for each alert
  |
  +- severity = 1                                                    // The severity of the alert
  +- time = 2018-12-31T11:02Z                                        // The alert time (needs a ?datetime for
proper rendering)
  |  +- moduleType = "ootb"                                          // The module type
  |
  +- moduleName = "Brute Force Login To Same Destination"            // A description of the module
  |
  +- statement = "module_144d43f5_f0b4_4cd0_8c6c_5ce65c37e624_Alert" // The name of the EPL statement
  |  +- events                                                       // The constituent events - as a sequence
of event maps
       |   +- [0]                                                    // offset 0 (i.e. the first constituent
event)
       |   |       |   |
       |   +- event_cat_name = "User.Activity.Failed Logins"
       |   +- device_class = "Firewall"                             // event meta (accessible as ${events
[0].device_class}$)
       |   |       |   +- event_source_id = "uttam:50002:1703395"    // Investigation URI to the individual
session (used by SA)
       |   |
       |   +- ...                                                   // Other meta
       |   |
       |   +- sessionid = 1703395                                   // NextGen sessionid
       |   |
       |   +- time = 1388487764                                     // event/session time at NextGen source
(as a long Unix timestamp)
       |   |
       |   +- user_dst = "user5"
       |
       +- [1]                                                       // offset 1 (i.e. the second consituent
event)
           |
           +- device_class = "Firewall"
           |
           +- event_cat_name = "User.Activity.Failed Logins"
           |
           +- event_source_id = "uttam:50002:1703405"
           |
           +- ...
           |
           +- sessionid = 1703405
           |
           +- time = 1388487766
           |
           +- user_dst = "user5"
```

There are two types of template variables available in the data model:

- **Alert Meta Data:** These hold alert level details like statement name, module name, alert id, alert time, severity, and others. In FreeMarker terminology, these are top level variables associated with the alert instance itself and can be referenced simply by their names like `${moduleName}`. The `time`

meta is special because it is of type `Date` and it needs to be suffixed with a `?datetime` to be properly rendered.

- **Constituent Event Meta Data**: These include the session meta fields from individual events that constitute the alert. An alert can have multiple constituent events, so there can be more than one such map in the same alert. These show up as a sequence of hashes to the FreeMarker template evaluator and must be referenced. For instance, the alert has two constituent events the event_source_id for the first is available as `${events[0].event_source_id}` and the same for the second is accessible as ${`events[1].event_source_id`.} You also need to be aware of which meta fields are multi-valued because those need be treated as sequences, for example ${`events[0].alias_host`} will not work because it is a sequence.

Note: The metadata available in the constituent events for a given alert is determined by the EPL SELECT clause. For example, alerts from `SELECT sessionid, time FROM ...` have only two meta values available (sessionid, time). Constituent events in `SELECT * FROM Event ...` will carry all meta fields from the `Event` type with **non-null** values.

If your template uses meta keys that are not present in all alert output, you should consider using the FreeMarker provisions for default values.

For example, if a template with text Id=${id},ec_outcome=${ec_outcome} is evaluated for an alert which does not include the meta key `ec_outcome` then the template evaluation fails. In such cases, you can use the missing value placeholder `${ec_outcome!"default"}`.

Import and Export a Global Notifications Template

This topic provides instructions on how to import and export a template for notifications.

- You can export default or user-defined templates.
- You can import a template that has been exported from the NetWitness instance. If you import a template with the same name as an existing template, then the existing template will be overwritten.

Import a Template

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Templates** tab.
4. In the toolbar, select  **> Import**.
   The **Import** dialog is displayed.
5. In the **Select File**field, type the filename or click **Browse** and select the file to be imported.
6. Click **Import**.

Export a Template

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Templates** tab.
4. Select the template you want to export.

   Note: You can export all the templates using the > **Export All** option.

5. In the **Actions** column, select > **Export**.
   The **Export** dialog is displayed.
6. In the **Enter File Name** field, type the filename.
7. Click **Save**.

Configure Email Servers and Notification Accounts

This topic provides instructions for configuring email so that users can receive notifications in NetWitness. NetWitness Platform can send notifications to users through email about various system events. To be able to configure these email notifications, you must first configure the SMTP email server. The Email Configuration panel provides a way to:

- Configure the email server.
- Set up an email account to receive notifications.
- View statistics on email operations.

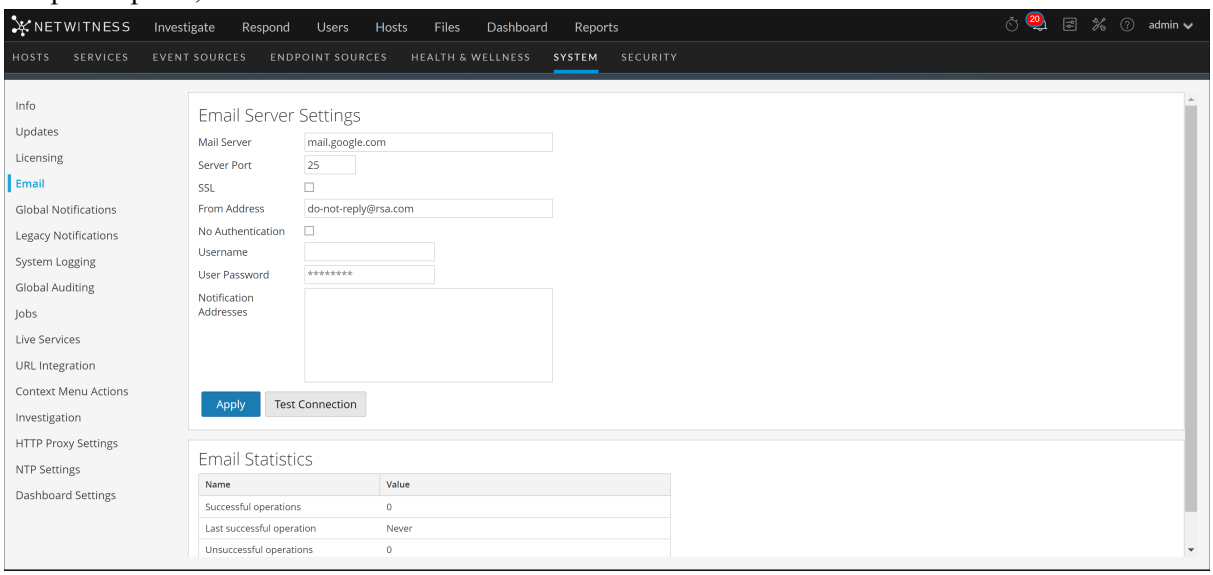
NetWitness requires access to an SMTP mail server in order to send reports to users. Each user account can be configured to receive emailed reports. These reports can be generated manually, through the user interface, or automatically, through the auditing system. The following guidelines apply:

- Any SMTP mail host can be used to deliver emails, and each host requires a different configuration. The SMTP provider provides the settings for configuration.
- Some SMTP servers require user authentication in order to relay emails successfully. Typically, this is the login and password for the email account.
- Best practice is to create a new, dedicated email account on the SMTP email server for NetWitness reports.

To configure NetWitness email notifications:

1. Go to **(Admin) > System**.
   The Administration System view is displayed.

2. In options panel, select **Email**.



3. If you want to change the default mail server, specify the **Mail server** name and **Server port**.
4. If the email server communicates with NetWitness using SSL, set the box next to **Use SSL**.
5. In the **From address** field, type the name of the email account sending NetWitness email notifications.
6. If the SMTP server requires user authentication to relay emails successfully, type the **Username** and **User Password** for logging in to the email account.
7. To activate the settings, click **Apply**.
   You can now configure NetWitness modules to receive various notifications by email.

Configure Global Audit Logging

Global Audit Logging provides NetWitness Auditors with consolidated visibility into user activities within NetWitness in real-time from one centralized location. This visibility includes audit logs gathered from the NetWitness system and the different services throughout the NetWitness infrastructure.

NetWitness audit logs collect in a centralized system that converts them into the required format and forwards them to an external syslog system. The external syslog system can be a third-party syslog server or a Log Decoder.

You configure global audit logging in the Global Audit Logging Configurations panel. An audit logging template defines the format and message fields of the audit log entries. A Syslog Notification Server configuration defines the destination to send the audit logs. If you want to forward audit logs to a Log Decoder, configure a Syslog type of Notification Server for the Log Decoder.

The following are some of the user actions logged from NetWitness:

- User logouts
- All UI pages accessed
- Committed configuration changes
- Queries performed by the user
- Data export operations

Note: For examples of some of the user actions logged, see Add New Configuration Dialog

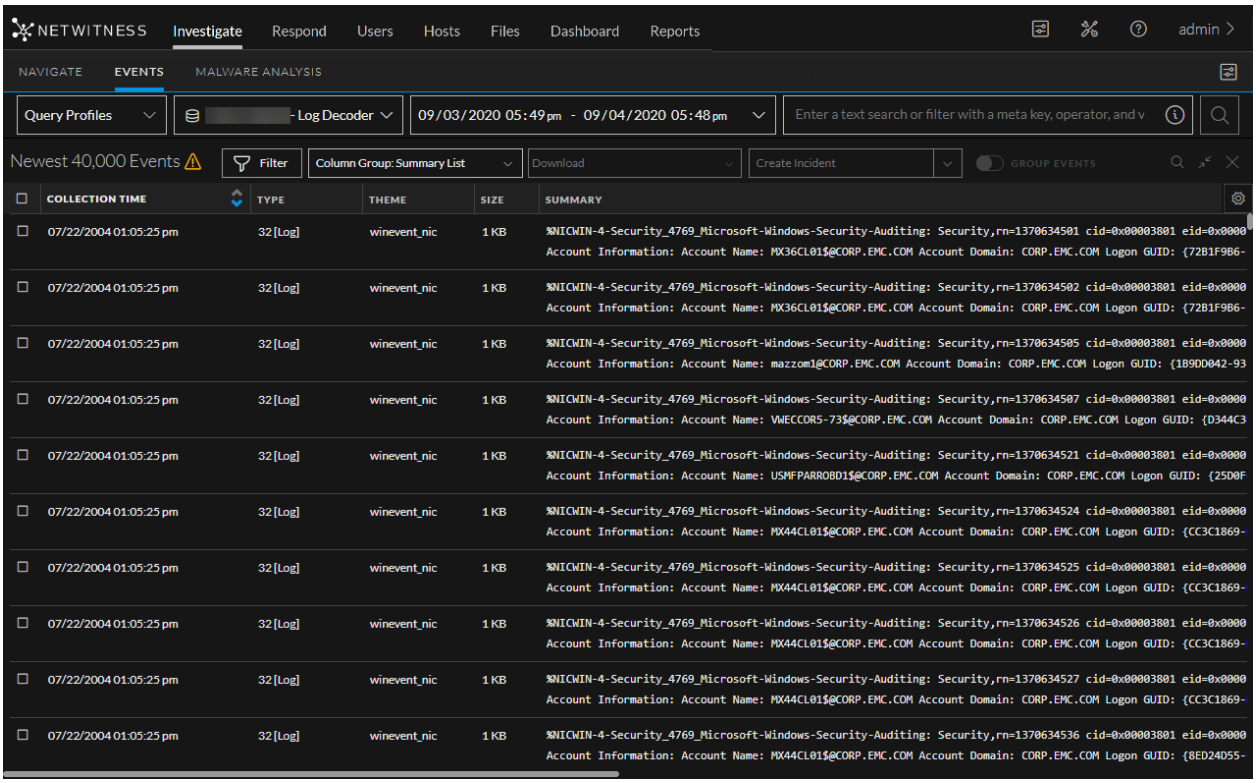
After you create a global audit logging configuration, audit logs containing these user actions automatically go to the external syslog system in the format specified in the selected Audit Logging template. You can create multiple global audit logging configurations for different destinations that use different templates. For example, you can create a global audit logging configuration for an external Syslog server with a template that contains all of the available meta keys and another configuration for a Log Decoder with a template that contains selected meta keys.

For Log Decoders, you use the Default Audit CEF Template. You can add or remove fields from the Common Event Format (CEF) template if you have specific requirements. Define a Template for Global Audit Logging provides instructions and Supported CEF Meta Keys describes the CEF meta keys available to use in the audit logging templates.

For third-party syslog servers, you can use a default audit logging template or define your own format (CEF or non-CEF). Define a Template for Global Audit Logging provides instructions and Supported Global Audit Logging Meta Key Variables describes the available variables.

Auditors can view the audit logs on the selected Log Decoder or third-party syslog server. If using a Log Decoder, auditors can view the audit logs using NetWitness Investigations or Reports.

The following figure shows global audit logs in Investigation (INVESTIGATE > Events).



For examples of some of the user actions logged, see Add New Configuration Dialog. For a list of message types being logged by the various NetWitness components, see Global Audit Logging Operation Reference.

Global Audit Logging - High-Level Procedure

Global Audit Logging is configured in the Global Audit Logging Configurations panel, which is accessed from (Admin) > System view > Global Auditing. Before you can configure Global Audit Logging, you need to configure a Syslog Notification Server and an Audit Logging template. A Syslog Notification Server defines the destination to send the audit logs. An Audit Logging template defines the format and message fields of the audit log entry.

The Global Audit Logging Configuration panel provides a **view settings** link that takes you to the Global Notifications panel ((Admin) > System view > Global Notifications) where you can configure the Syslog Notification Server and Audit Logging template.

Perform the following procedures in the order shown to configure Global Audit Logging.

Procedures	Reference / Instructions
1. Configure a Syslog Notification Server.	Configure a Syslog Notification Server to use for Global Audit Logging. You can define a third-party syslog server or Log Decoder as a destination to receive the audit logs. Configure a Destination to Receive Global Audit Logs. Global Audit Logging configurations use the Syslog notification server type. If you want to forward audit logs to a Log Decoder, create a Notification Server of the Syslog type.
2. Select or configure an Audit Logging template to use.	Select an Audit Logging template for the Syslog notification server. You can use a default Audit Logging template or define your own audit logging template. Global Audit Logging configurations use the Audit Logging template type and a Syslog notification server. Configure Templates for Notifications provides additional information. For Log Decoders, use the **Default Audit CEF Template**. You can add or remove fields from the Common Event Format (CEF) template if you have specific requirements. Define a Template for Global Audit Logging provides instructions. For third-party syslog servers, you can use a default audit logging template or define your own format (CEF or non-CEF). Define a Template for Global Audit Logging provides instructions and Supported Global Audit Logging Meta Key Variables describes the available variables.
3. (Optional - Only if consuming with a Log Decoder) Deploy the Common Event Format parser to your Log Decoder from Live.	Ensure that you have deployed and enabled the latest Common Event Format parser from Live. Find and Deploy Live Resources and Enable and Disable Log Parsers provide instructions.
4. Define a global audit logging configuration, which defines how the global audit logs are forwarded to external Syslog systems.	Define a Global Audit Logging Configuration provides instructions. After you add a Global Audit Logging configuration, audit logs are forwarded to the selected Notification Server in the configuration.
5. Verify that the global audit logs show the audit events.	Test your audit logs to ensure that they show the audit events as defined in your audit logging template. Verify Global Audit Logs provides instructions.

Configure a Destination to Receive Global Audit Logs

In Global Audit Logging, Syslog Notification Servers are the configurations that define the destinations to receive global audit logs. You need to configure a Syslog Notification Server to use Global Audit Logging. You can define a third-party syslog server or a Log Decoder as the destination to receive the audit logs.

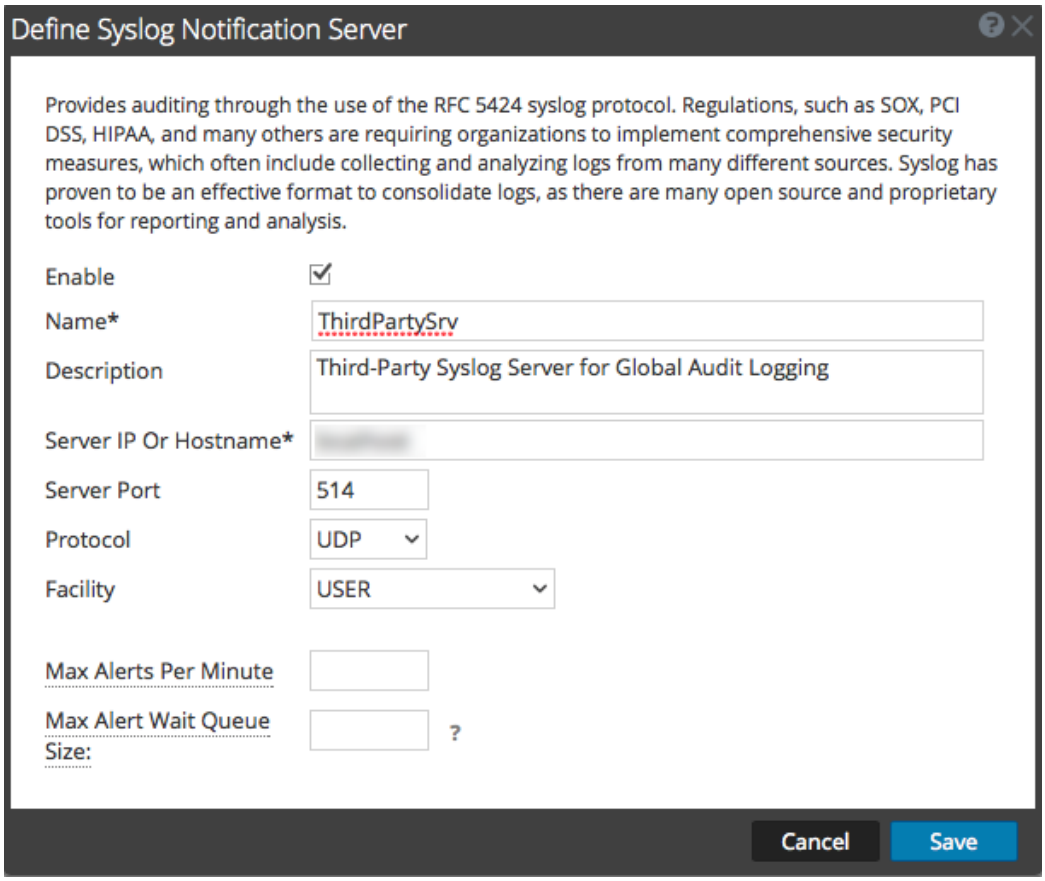
Configure a Syslog Notification Server for a Third-Party Syslog Server

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Servers** tab.

   Note: You do not need to configure the Output tab for Global Audit Logging.

4. From the drop-down menu, select **Syslog**.

   The **Define Syslog Notification Server** dialog is displayed.



5. Configure the Syslog notification server as described in the following table.

Field	Description
Enable	Select to enable the notification server.
Name	A name to identify or label the third-party syslog server.
Description	(Optional) A brief description of the notification server.

Field	Description
Server IP or Hostname	The third-party syslog server hostname or IP address.
Server Port	The port number where the target syslog process is listening.
Protocol	The protocol to be used for transferring formatted audit logs to the third-party syslog server.
Facility	The syslog facility to be used for writing formatted audit logs to the third-party syslog server.

The **Max Alerts Per Minute** and **Max Alert Wait Queue Size** fields are not used for Global Audit Logging.

6. Click **Save**.

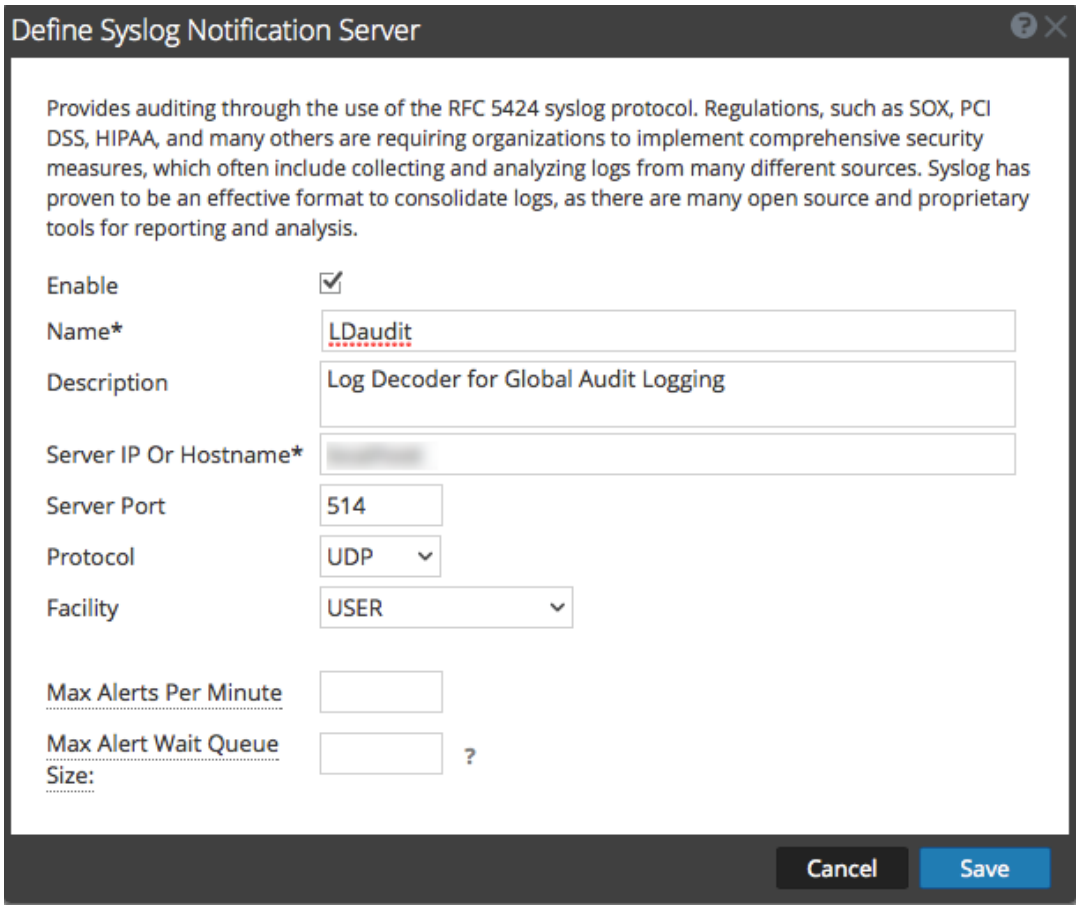
Configure a Syslog Notification Server for a Log Decoder

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Servers** tab.

Note: You do not need to configure the Output tab for Global Audit Logging.

4. From the  drop-down menu, select **Syslog**.

The **Define Syslog Notification Server** dialog is displayed.



5. Configure the Syslog notification server as described in the following table.

Field	Description
Enable	Select to enable the notification server.
Name	A name to identify or label the Log Decoder syslog notification server.
Description	(Optional) A brief description of the notification server.
Server IP or Hostname	The Log Decoder hostname or IP address.
Server Port	The port number where the target syslog process is listening.
Protocol	The protocol to be used for transferring formatted audit logs to the Log Decoder.
Facility	The Syslog facility to be used for writing formatted audit logs to the Log Decoder.

The **Max Alerts Per Minute** and **Max Alert Wait Queue Size** fields are not used for Global Audit Logging.

6. Click **Save**.

Next Steps

Select a default Audit Logging template to use for Global Audit Logging. If necessary, you can define your own custom template. Define a Template for Global Audit Logging provides additional information.

Define a Template for Global Audit Logging

This topic provides instructions on how to define an audit logging template to use for Global Audit Logging. Before you configure Global Audit Logging, configure a Syslog notification server and select an Audit Logging template. You can choose to use a default audit logging template or you can define your own template.

NetWitness includes two default audit logging templates:

- **Default Audit CEF Template**: You can use this template for Log Decoders and third-party syslog servers.
- **Default Audit Human-Readable Format**: You can use this template only for third-party syslog servers. Do not forward messages from this template to a Log Decoder.

The first procedure provides instructions on how to define an audit logging template for a Log Decoder. The audit logging template defines the format and message fields of the audit logs sent to the Log Decoder or third-party syslog server.

Global audit logging templates that you define for a Log Decoder use Common Event Format (CEF) and must meet the following specific standard requirements:

- Include the CEF headers in the template.
- Use only the extensions (Key=Value) listed in the Supported CEF Meta Keys table.
- Ensure that the extensions are in the `key=%{string}<space>key=%{string}` format.

The second procedure provides instructions on how to define a custom global audit logging template in human-readable format for a third-party syslog server. For third-party syslog servers, you can define your own format (CEF or non-CEF).

Define a Global Audit Logging Template for a Log Decoder

You can use the **Default Audit CEF Template** to send global audit logs to a Log Decoder. To define your own template:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Click the **Templates** tab.
4. Click  to configure a template.

5. In the **Define Template** dialog, provide the following information:
   a. In the **Name** field, type the name for the template.
   b. In the **Template Type** field, select the **Audit Logging** template type.
   c. In the **Description** field, type a brief description for the template.
   d. In the **Template** field, enter the format for the audit logging template.
   The following format is a customized template provided as an example. It differs from the default CEF template.

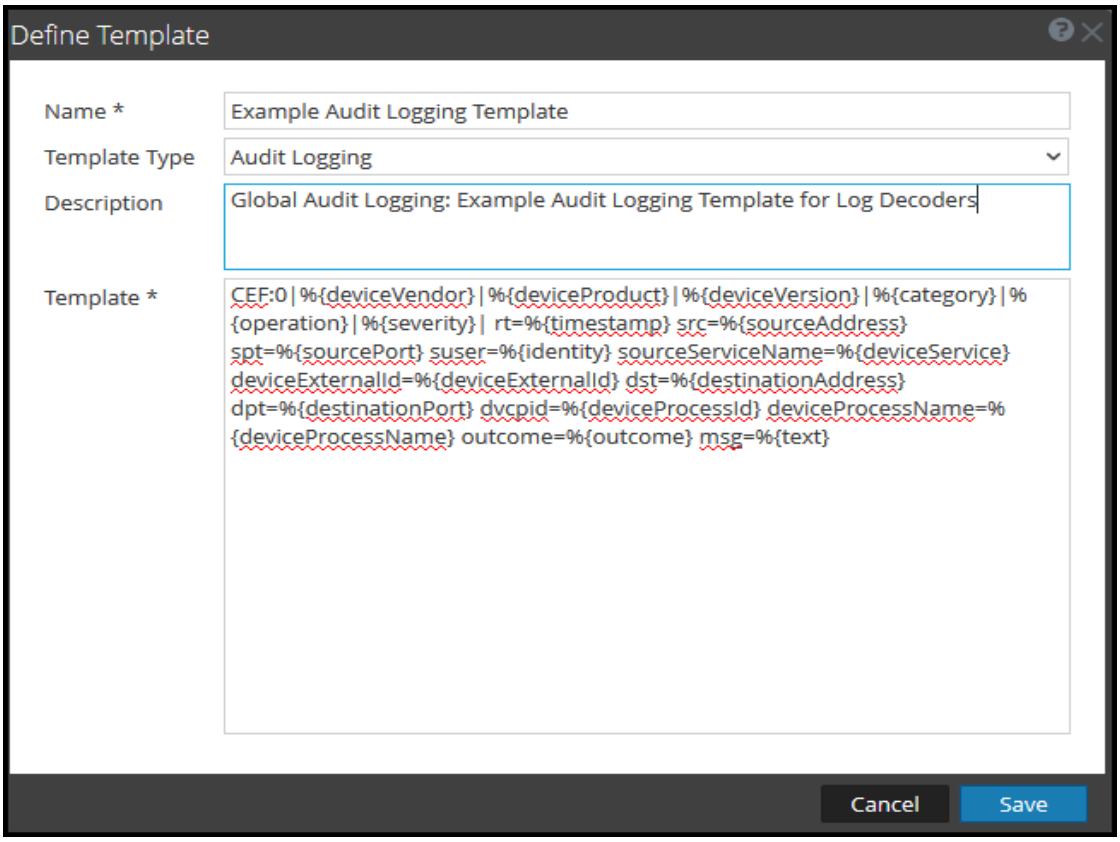
```
CEF:0|%{deviceVendor}|%{deviceProduct}|%{deviceVersion}|%{category}|
%{operation}|%{severity}|
rt=%{timestamp} src=%{sourceAddress} spt=%{sourcePort} dpt=
%{destinationPort} dst=%{destinationAddress} dvcpid=%{deviceProcessId}
tpt=%{transportProtocol} sessionId=%{sessionId} scope=%{scope} suser=
%{identity} sourceServiceName=%{deviceService} deviceExternalId=
%{deviceExternalId} deviceProcessName=%{deviceProcessName} device
Facility=%{deviceFacility} outcome=%{outcome} msg=%{text} remoteAddress
=%{remoteAddress} reasonForFailure=%{reasonForFailure} reason=%{reason}
arguments=%{Arguments} user=%{User} referrerURL=%{referrer} role=%{Role}
id=%{id} account=%{Account} deviceIDs=%{deviceIDs} file=%{file} account
Provider=%{AccountProvider} uri=%{uri} addRole=%{Add.Role} addPermission
=%{Add.Permission} userAgent=%{userAgent} userGroup=%{userGroup}
userRole=%{userRole} key=%{key} value=%{value} paramKey=%{Key}
paramValue=%{Value} alert=%{alert} incident=%{incident} action=%{action}
notification Binding=%{NotificationBinding} name=%{name} enabled=
%{enabled} disabled=%{disabled} params=%{parameters}
```

The highlighted CEF syslog header is required to conform to the CEF standard and is a requirement for the CEF parser in the Log Decoder. The other keys are optional and you can configure them. See all the supported meta keys that are supported by the CEF parser in the Log Decoder in the Supported CEF Meta Keys table.

Note: Use all of the extensions in the following format:
`deviceProcessName=%{deviceProcessName} outcome=%{outcome}`
Include a `<space>` between each key=%{string} pair in the extension keys section.

Note: After you upgrade to 11.x from earlier versions, then '$' is replaced with '%' automatically

6. Click **Save**.



After you define the CEF audit logging template, ensure that you have deployed and enabled the latest Common Event Format (CEF) parser from Live. "Find and Deploy Live Resources" in the *Live Services Management Guide* provides instructions.

Note: If you need to use a specific meta key for Investigations and Reporting, ensure that the meta keys that you select are indexed in the **table-map-custom.xml** file on the Log Decoder. If they are not indexed, follow the instructions in the "Maintain the Table Map Files" topic in the *Host and Services Configuration Guide* procedure to update the table mappings. Ensure that the meta keys are also indexed in the **index-concentrator-custom.xml** on the Concentrator. See the "Edit a Service Index File" topic in the *Host and Services Configuration Guide* for additional information.

Define a Custom Global Audit Logging Template

For third-party syslog servers, you can define your own template format (CEF or non-CEF). You can use the **Default Audit Human-Readable Format** template to send global audit logs to a third-party syslog server in a format that is easier to read than the CEF format. If you want to define your own template in human-readable format, follow this procedure.

For Log Decoders, you must use a CEF template with some specific requirements. The *Define an Audit Logging Template for a Log Decoder* procedure above provides instructions for creating a template in CEF format.

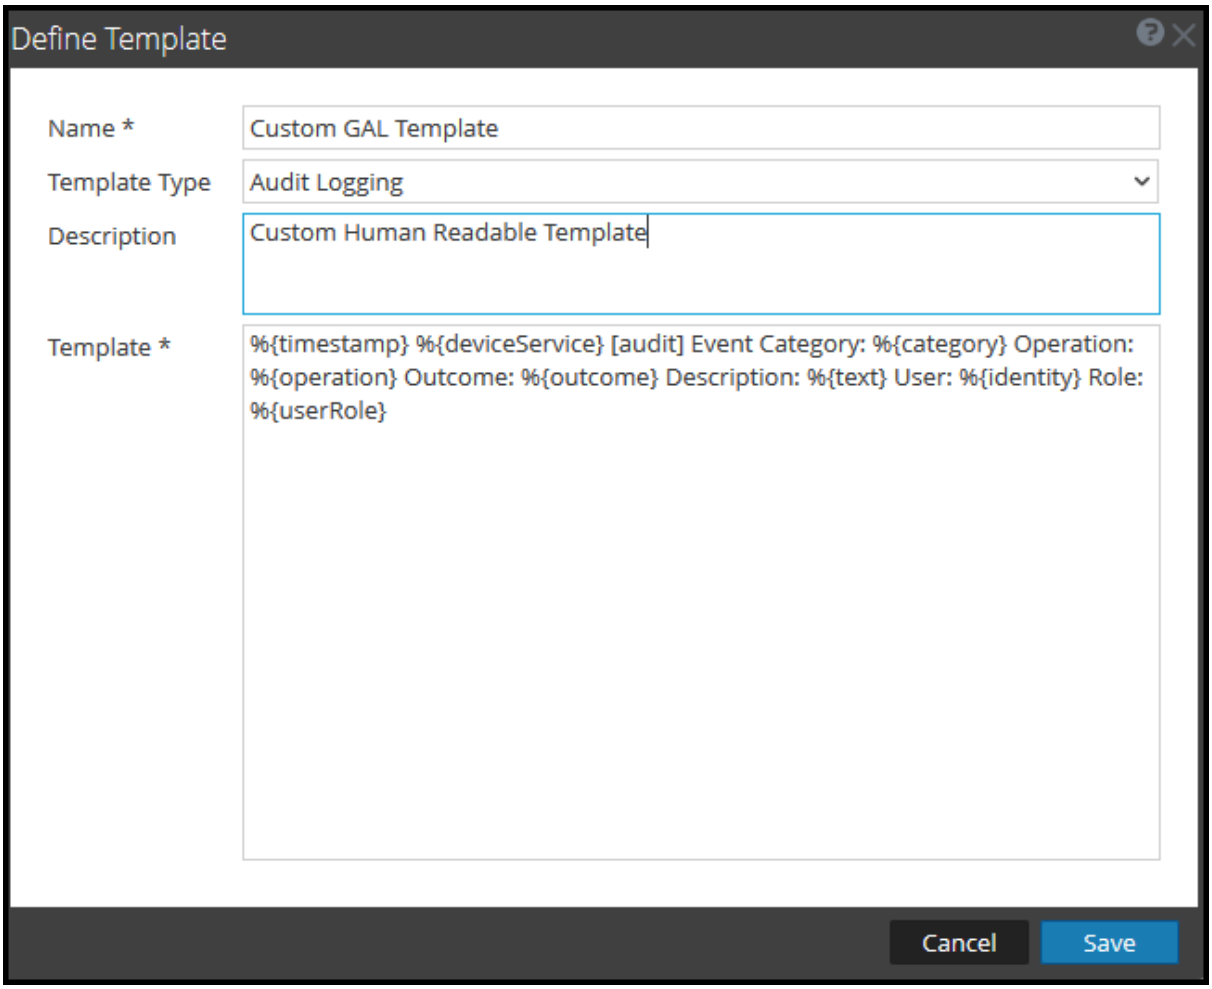
To define a custom global audit logging template in human-readable format:

1. Go to  **(Admin) > System**.
2. In the left navigation panel, select **Notifications**.
3. Click the **Templates** tab.
4. Click  to configure a template.
5. In the **Define Template** dialog, provide the following information:
   a. In the **Name** field, type the name for the template.
   b. In the **Template Type** field, select the **Audit Logging** template type.
   c. In the **Description** field, type a brief description for the template.
   d. In the **Template** field, enter the format for the audit logging template. The following example is in human-readable format with selected meta key variables.

```
%{timestamp} %{deviceService} [audit] Event Category: %{category}
Operation: %{operation} Outcome: %{outcome} Description: %{text}
User: %{identity} Role: %{userRole} Parameters: %{parameters}
```

   You can use any of the meta key variables that are supported by global audit logging shown in the Supported Global Audit Logging Meta Key Variables table.

6. Click **Save**.



The following example shows global audit logs in human-readable format for this template:

```
Jun 11 2019 04:53:54 UpdateStackConcentrator Jun 11 2019 04:53:54 CONCENTRATOR
[audit] Event Category: DATA_ACCESS Operation: sdk.info Outcome: pending
Description: has requested the SDK summary info User: admin Role: null
params=flags\=1
```

```
Jun 11 2019 04:53:55 updatestackadminserver Jun 11 2019 04:53:55 source-server
[audit] Event Category: API Operation: /rsa/process/ready Outcome: success
Description: null User: Netwitness Web(nw-web) Role: null params=
{"Arguments":"[]"}
```

```
Jun 11 2019 05:15:46 UpdateStackeplh Jun 11 2019 05:15:46 LOG_DECODER [audit]
Event Category: MANAGEMENT Operation: upload Outcome: pending Description: has
started uploading file User: escalateduser Role: null params=file\=esmFeed.zip
```

Next Step

Define a Global Audit Logging Configuration provides instructions for defining a global audit logging configuration for NetWitness.

Define a Global Audit Logging Configuration

This topic tells administrators how to define a global audit logging configuration. This procedure is required only if you choose to set up centralized audit logging in your environment. These global audit logging configurations define how the global audit logs are forwarded to external syslog systems or Log Decoders. Audit logs are forwarded to the selected Notification Servers.

Prerequisites

Before starting this procedure, configure the following to use for global audit logging:

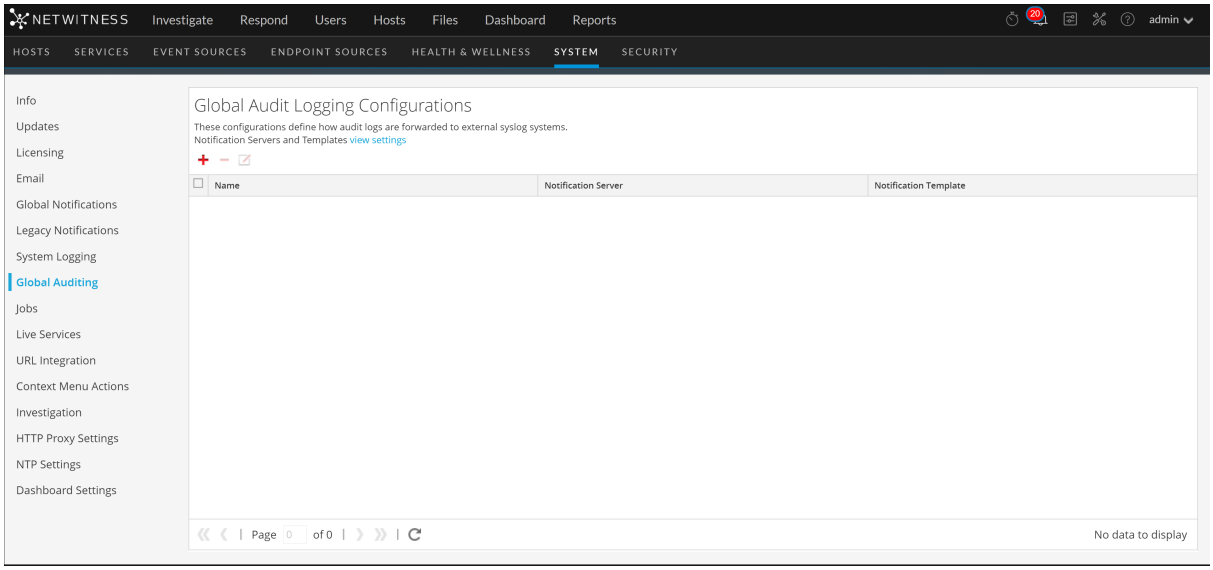
- Syslog Notification Server
- Audit Logging Template

You configure the notification server and template on the Global Notifications panel. You can access the Global Notifications panel by going to (Admin) > System and select Global Notifications. You can only define a Syslog type of Notification Server for global audit logging. For Log Decoders, use a Syslog type of Notification Server and a Common Event Format (CEF) audit logging template. You can use a default audit logging template or define your own template. You can create multiple audit logging templates and Syslog Notification Servers to use for your global audit logging configurations.

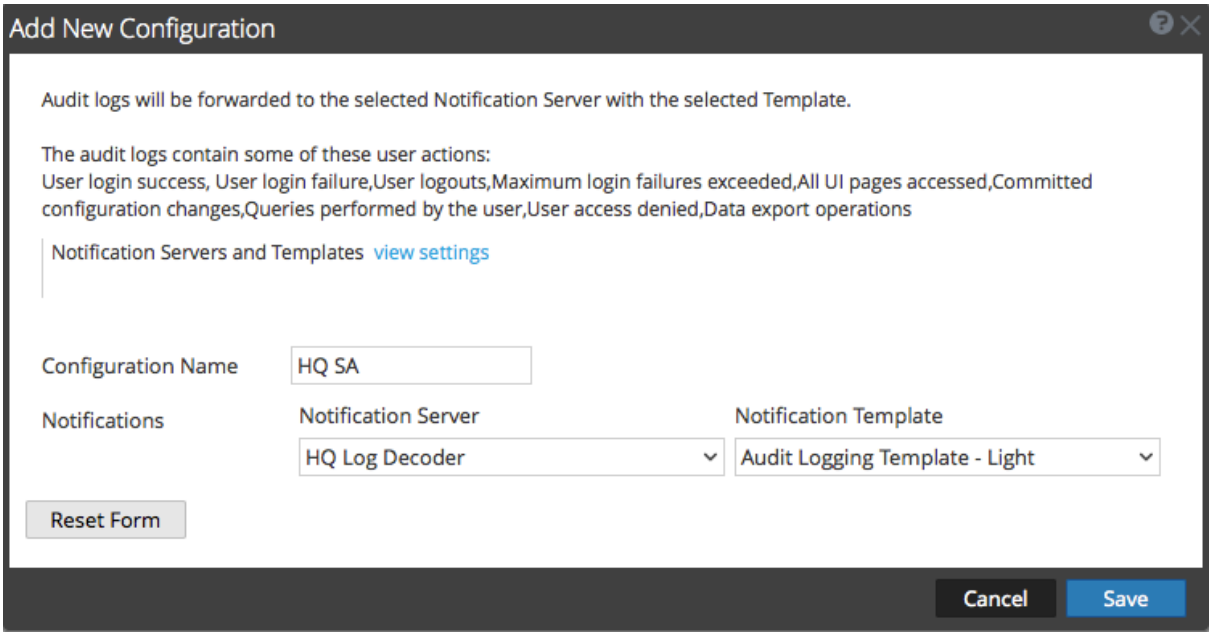
If you are forwarding global audit logs to a Log Decoder, deploy the Common Event Format parser to your Log Decoder from Live.

Add a Global Audit Logging Configuration

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Auditing**.
   The **Global Audit Logging Configurations** panel is displayed.



3. Click + to add a global audit logging configuration.
   The **Add New Configuration** dialog is displayed.



4. In the **Configuration Name** field, type a unique name for the global audit logging configuration. For example, you can create a configuration for a specific type of global audit logging configuration, such as HQ NW for a NetWitness headquarters configuration.
5. In the **Notifications** section, select the syslog **Notification Server** to use for this configuration. The notification server is the destination to send the global audit logs.
6. Select the audit logging **Notification Template** to use for this configuration. The Audit Logging template defines the format and audit log message fields to be sent.
7. Click **Save**.

Add New Configuration Dialog provides additional information and examples of the user actions logged. For a list of message types being logged by the various NetWitness components, see Global Audit Logging Operation Reference.

Edit a Global Audit Logging Configuration

This topic provides instructions on how to edit a global audit logging configuration. You can edit a global audit logging configuration to change the destination of the global audit logs for your user audits by selecting a different Notification Server. You can also change the format and message fields of the global audit log entries by selecting a different Notification Template. You make changes to the Notification Server or Template on the Global Notifications panel. You can access the Global Notifications panel by clicking the **view settings** link on the Global Audit Logging Configurations panel.

You cannot change which NetWitness user actions are logged and sent in the global audit logs.

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Auditing**.

3. In the **Global Audit Logging Configurations** panel, select a configuration to edit and click .
4. In the **Add New Configuration** dialog, modify the global audit logging configuration as required. You can modify the **Configuration Name** and select a different **NotificationServer** or **Template**.
5. Click **Save.**

Delete a Global Audit Logging Configuration

Deleting a global audit configuration does not delete the associated notification server and template. After you delete a global audit logging configuration, the forwarding of global audit logs specified in that configuration is discontinued.

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Auditing**.
3. In the **Global Audit Logging Configurations** panel, select a configuration to delete and click .
   A confirmation dialog is displayed.
4. Click **Yes**.
   The selected configuration is deleted.

Verify Global Audit Logs

This topic provides instructions on how to verify global audit logs. After you have configured global audit logging, you need to test your global audit logs to ensure that they show the audit events as defined in your global audit logging template.

In version 11.5 and later, audit logging provides information about the aggregation account and the actual user who submitted the query. For example, the information is displayed as follows in the audit log:

```
User aggAccount (session 478, [::1]:1133, on behalf of <username of submitter>) has requested the SDK transforms.
```

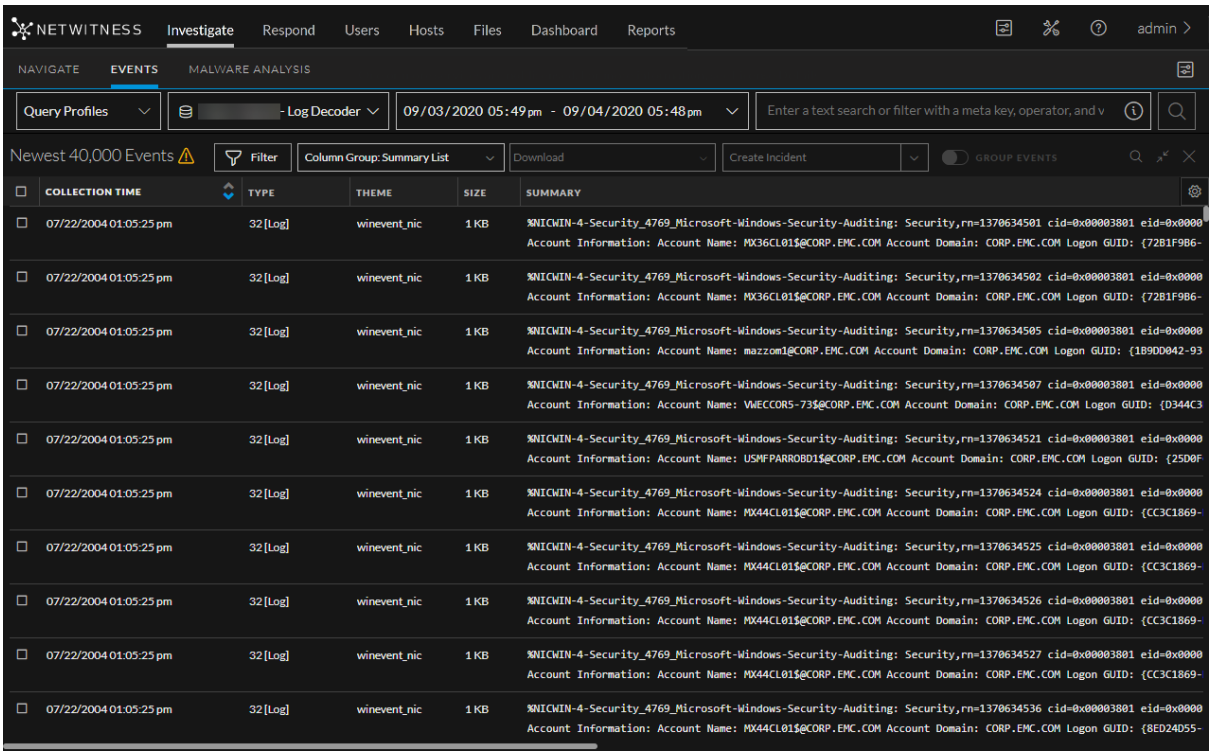
This information is available through multiple levels of Brokers and Concentrators.

Note: If you are running a mixed-version environment, any version earlier than 11.5 will not provide the real user information.

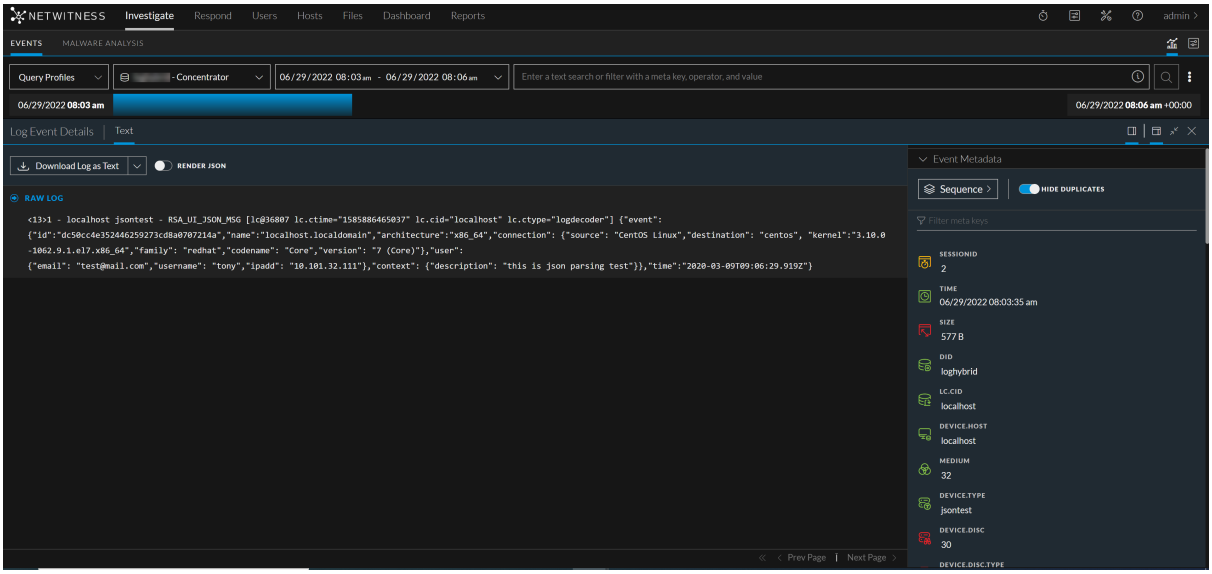
Before starting this task, complete the steps detailed in Configure Global Audit Logging.

To view and verify the global audit logs if you are using a Log Decoder:

1. Go to **Investigate > Events**, select the Log Decoder service and click the submit query icon ( ) to the right of the query bar.



2. Compare the fields in the global audit logs with the fields defined in the global audit logging template that you used in your global audit logging configuration.
3. Double-click a log to open the reconstruction and click  to open the Event Meta panel.
4. Verify that the meta that you want to audit is correct.



Example CEF Output

The following example shows global audit logs for an audit logging Common Event Format (CEF) template.

Template:

```
CEF:0|%{deviceVendor}|%{deviceProduct}|%{deviceVersion}|%{category}|%
{operation}|%{severity}|
t=%{timestamp} src=%{sourceAddress} spt=%{sourcePort} tpt=%{transport
Protocol} scope=%{scope} suser=%{identity} sourceServiceName=%{device Service}
deviceExternalId=%{deviceExternalId} deviceProcessName=%{device ProcessName}
outcome=%{outcome} msg=%{text} remoteAddress=%{remoteAddress}
reasonForFailure=%{reasonForFailure} reason=%{reason} arguments= %{Arguments}
user=%{User} referrerURL=%{referrer} role=%{Role} id=%{id} account=%{Account}
deviceIDs=%{deviceIDs} file=%{file} accountProvider= %{AccountProvider} uri=%
{uri} addRole=%{Add.Role} addPermission= %{Add.Permission} userAgent=%
{userAgent} userGroup=%{userGroup} userRole= %{userRole} key=%{Key} value=%
{Value} alert=%{alert} incident=%{incident} action=%{action}
notificationBinding=%{NotificationBinding} name=%{name} enabled=%{enabled}
disabled=%{disabled} params=%{parameters}
```

Example logs:

```
Jun 07 2019 09:06:05 UpdateStackConcentrator CEF:0|RSA|NetWitness Audit|
11.3.1.0|AUTHENTICATION|logoff|6|rt=Jun 07 2019 09:06:05 src=101.101.101.
101 spt=55060 scope=scope suser=admin sourceServiceName=CONCENTRATOR
deviceExternalId=3ebf91d9-e879-4727-a473-72d309e1741d deviceProcessName=
NwConcentrator outcome=success \r\n
```

```
Jun 07 2019 09:06:11 UpdateStackConcentrator CEF:0|RSA|NetWitness Audit|
11.3.1.0|AUTHENTICATION|login|6|rt=Jun 07 2019 09:06:11 src=101.101.101.101
spt=55060 scope=scope suser=admin sourceServiceName=CONCENTRATOR device
ExternalId=3ebf91d9-e879-4727-a473-72d309e1741d deviceProcessName=
NwConcentrator outcome=success userGroup=Administrators userRole=admin.owner,
aggregate,concentrator.manage,connections.manage,database.manage,everyone,
index.manage,logs.manage,rules.manage,sdk.content,sdk.manage,sdk.meta,
sdk.packets,services.manage,storedproc.execute,storedproc.manage,sys.manage,
users.manage \r\n
```

Configure Centralized Audit Logging

NetWitness Platform collects audit logs from all the NetWitness services and aggregates it into a single file in a centralized location on the NetWitness Admin Server. This aggregated log file provides the advantage for faster access and easy analysis of the audit logs.

The aggregated logs from all services are sent to the following centralized location:

- `/var/netwitness/logstash/logs/rsa-netwitness-audit.log` (JSON format)
- Syslog running on the local host (human-readable format)

Centralized audit logging is enabled by default. To forward the aggregated logs to the external syslog system (a third-party Syslog server or Log Decoder), you must configure the Global Audit logging in  **(Admin) > System > Global Auditing**. The aggregated logs are sent in the format specified in the selected Audit Logging template. A Syslog Notification Server configuration defines the destination to send the audit logs. To forward the audit logs to a Log Decoder, configure a Syslog type of Notification Server for the Log Decoder.

- For instructions on how to define a template, see Define a Template for Global Audit Logging
- For instructions on how to configure a syslog notification server, see Configure a Destination to Receive Global Audit Logs
- For instructions on how to configure global Audit logging, see Define a Global Audit Logging Configuration

Filtering the Aggregated Logs

Before the logs are aggregated, standard filters are applied to the logs to reduce redundancy and filter out logs that are not useful. The filters contain entries that control the content written to the aggregated log file. The following default filters are available in `/etc/logstash/` location.

- `json-filter-action.yaml`
- `json-filter-category.yaml`

`json-filter-action.yaml` - This filter blacklists the log messages based on the `operation` meta key and stops the log message being written to aggregated log file. For example, if `"/rsa/process/ready":"true"` is entered in `json-filter-action.yaml`, any raw log that contains `"/rsa/process/ready"` in the `operation` meta key is blacklisted and not written to aggregated log file.

Note: If you do not want to apply filters, then delete all the default entries and replace with `{}` character. Note that this increases the log size and the logs may be redundant.

`json-filter-category.yaml` - This filter whitelists the log messages based on the `category` meta key and writes the log message to the aggregated log file. For example, if `'\b(?i)SECURITY\b':"true"` is entered in `json-filter-category.yaml`, any raw log that contains 'SECURITY' in the `category` meta key is whitelisted and written to the aggregated log file.

Log Retention Policy

The aggregated log file is retained as per the following default settings:

- If the file size reaches 250 MB, the file is compressed as a single zip file.
- If the number of zip files exceed 90, the oldest zip file in the directory is automatically deleted.

You can modify the log retention policy in the file `logstash` available in `/etc/logrotate.d/` location.

Note: The size of the aggregated log file depends on the filters applied, so make sure to set up filters correctly for optimal directory space.

Disable Centralized Audit Logging

If you do not want the logs to be aggregated, in the `json-filter-category.yaml` filter, delete all the default entries and replace with the `{}` character.

Configure Investigation Settings

This topic provides instructions for administrators who are configuring the settings that apply to all investigations on the NetWitness instance being configured. The settings for configuring and tuning behavior of NetWitness Investigate are available in the  (Admin) > System > Investigation panel. These settings apply to all investigations and reconstructions on the current instance of NetWitness.

Map Context Hub Meta Types

The Context Hub is preconfigured with meta fields mapped to entities. NetWitness Respond and Investigate use these default mappings for context lookup. For information about adding meta keys, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.

Caution: For the Context Lookup to work correctly in the Respond and Investigate views, when mapping meta keys in the  (Admin) > System > Investigation > Context Lookup tab, it is best practice to add only meta keys to the Meta Key Mappings. Do not add fields in the MongoDB to the Meta Key Mappings. Here is a sample meta key and Mongo DB field; `ip.address` is a meta key and `ip_address` is a field in the MongoDB.

In the **Context Lookup** tab, you can manage mapping of Context Hub meta types with meta keys in Investigate. You can add or remove meta keys in the list of meta types supported in Investigate by Context Hub. Procedures associated with this tab are provided in "Manage Context Hub Lists and List Values in the Navigate and Events Views" in the *NetWitness Investigate User Guide*.

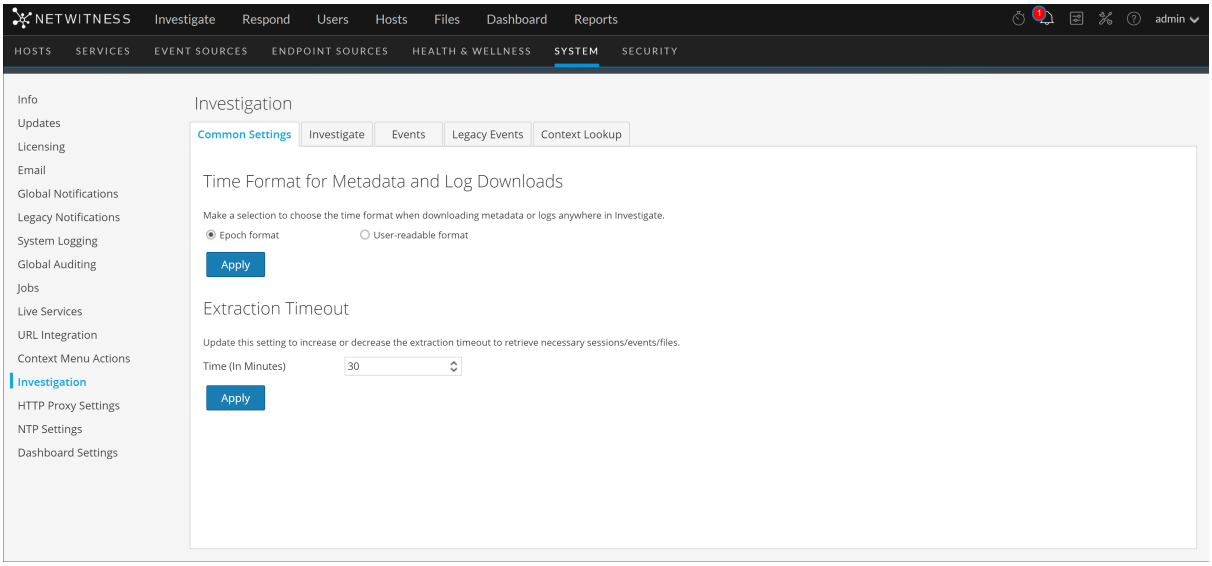
Configure Common Settings

In version 11.5 and later, the Common Settings tab allows you to configure settings that apply to the Navigate view, the Events view, and the Legacy Events view. Initially, the only setting that you can set for all views is the time format used when downloading metadata and logs, and other settings may be added in future versions.

By default, the time format for downloads is Epoch format, which shows the time as a numerical value representing the number of seconds from the Unix epoch, January 1, 1970. The resulting number requires a conversion to be understood. You can change the setting to get a more understandable format that combines the user preference time zone, date format, and time format into an easily understood representation, which follows the industry standard ISO 8601 representation when possible.

This setting applies to all 11.7 Investigate views.

1. Go to  **(Admin) > System**, and in the options panel, select **Investigation**.
   The Investigation Configuration panel is displayed.



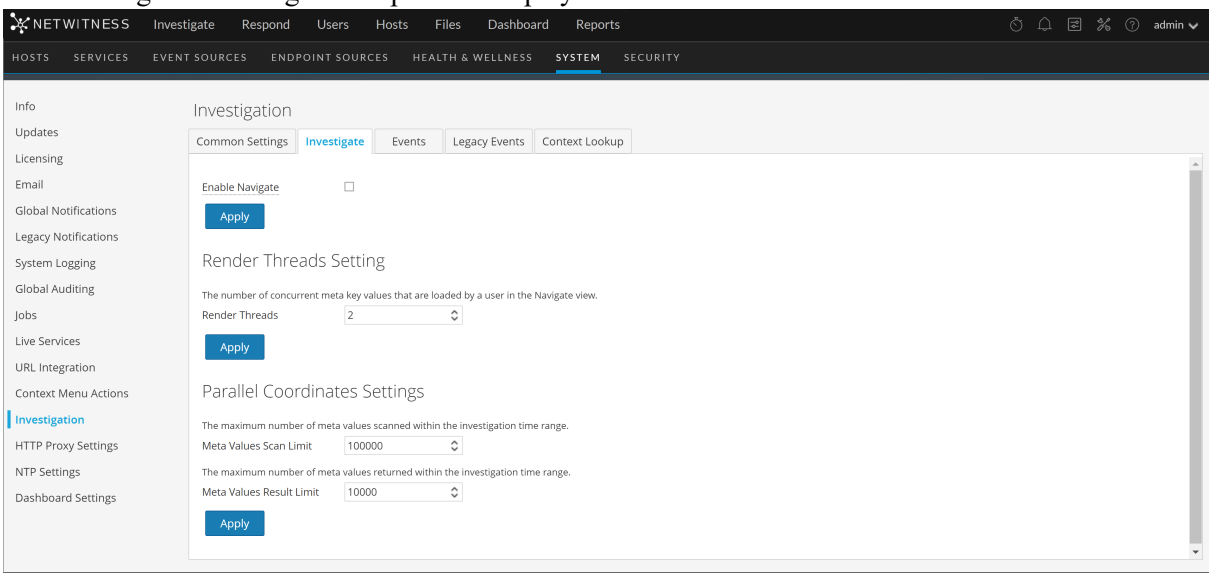
2. In the **Common Settings** tab, do the following:
   a. Select the time format to be used in metadata and log downloads in Investigate.
   b. Edit the time in minutes to set the extraction timeout when the logs are downloaded before the session expires.
      By default, it is set to 30 minutes and can be set to a maximum of 60 minutes.
3. Click **Apply**.
   The setting goes into effect immediately.

Configure Navigate and Legacy Events View Settings

The name of the Version 11.3 and earlier Events tab was changed to Legacy Events tab in Version 11.4.

1. Go to  **(Admin) > System.**
2. In the options panel, select **Investigation.**

The Investigation Configuration panel is displayed.



3. In the **Investigate** tab, in the **Render Threads Settings** field, select the maximum number of concurrent meta key values that are loaded by a single user in the Navigate view. Click **Apply**.
4. In the **Investigate** tab, in the **Parallel Coordinates Settings** section, set the maximum limits for meta values scanned and meta value results that can be included in a parallel coordinates visualization. For better performance, these are the recommended settings: Meta Values Scan Limit - 100000 and Meta Values Result Limit to 1,000-10,000
Click **Apply**.
5. In the **Legacy Events** tab, in the **Enable Legacy Events** section, select the check box to view the legacy sub menus and legacy events in classic view. Click **Apply**.
6. In the **Legacy Events** tab, in the **Event Search Settings** section, set the maximum numbers of events scanned and event results displayed when an analyst is conducting an event search in the Legacy Events view. The actual number of events scanned and displayed may be slightly greater than the limit set here. Click **Apply**.
7. In the **Legacy Events** tab, in the **Reconstruction Settings** section, set the limits for the amount of data processed in the reconstruction of a single event. The default values are 500 maximum packets and 2097152 bytes. If analysts are seeing slow performance when reconstructing sessions in Investigate, the reconstruction settings may need adjustment. Click **Apply**.

Caution: Setting a higher value affects the performance of the NetWitness Server by increasing the time and memory taken to create a reconstruction of an event. Setting the value to zero disables any limit and may lead to a NetWitness Server crash.

7. (Optional) In the **Legacy Events** tab, in the **Web View Reconstruction Settings** section, enable the use of supporting files in a web view reconstruction, and configure the additional settings to calibrate web view reconstructions. These include the time range (in seconds) to scan for related events, the maximum number of related events to scan, and overrides to Reconstruction Settings for use with web view reconstructions. Click **Apply**.

Clear Reconstruction Cache for Services

Under Reconstruction Cache Settings, administrators can clear the cache for one or more services. For example, the administrator can clear the cache for a single Broker, a Broker and Decoder, or all connected services. These are a few examples of causes for stale cache being used in a reconstruction.

- The downstream services may have their sessions invalidated or data reset. As an example, if Investigate is browsing a Broker and a downstream Concentrator or Decoder has a data reset, the metadata and session data for the investigating service (Broker) does not match the content if the downstream service has reset and repopulated. The reconstruction in Investigate shows content from cache, which does not match the real content. Even if the Decoder is offline, content is still displayed in the Broker reconstruction. Clearing cache on the Broker causes the NetWitness to reach out to the Decoder and an error is returned because the Decoder is offline.
- Another case where cache may be stale is when a service ID for a downstream service changes. This can happen when exporting, importing, deleting, and adding services to NetWitness because NetWitness can reuse service IDs. In this case, clearing the cache on the Broker causes NetWitness to request data from the services.

To clear reconstruction cache, do one of the following:

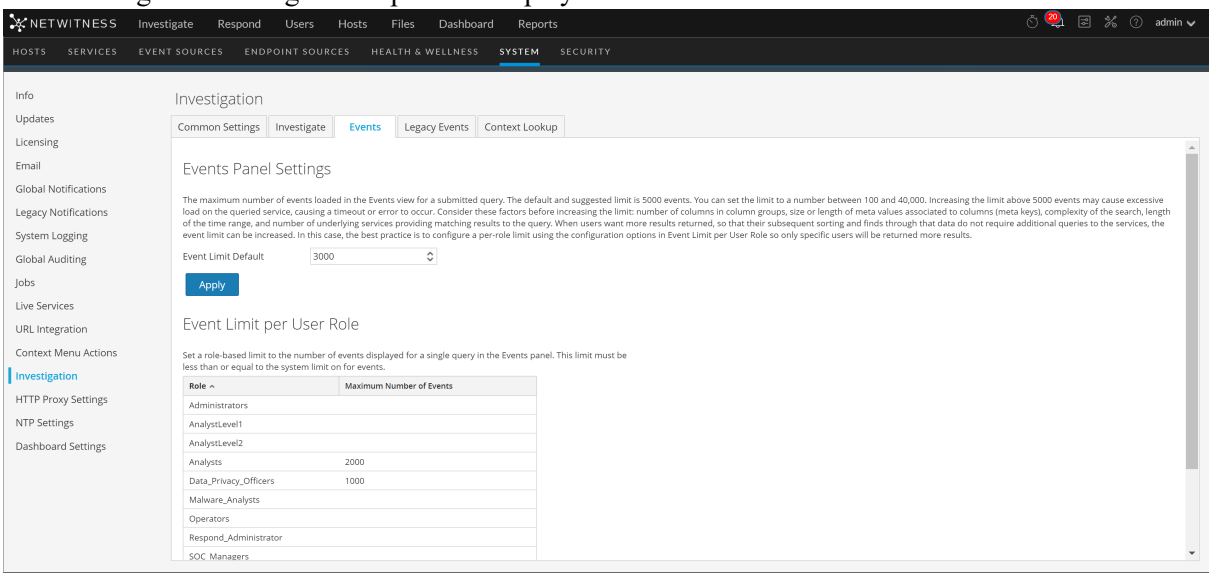
1. To clear cache for one or more services, select the services and click **Clear Cache for the Selected Services**.
2. To clear the cache for all listed services, click **Clear Cache for All Services.**
   The reconstruction cache for the selected services is cleared. NetWitness sends a request for data to the services.

Configure Events View Settings

These settings apply to the 11.3 and earlier Event Analysis view and the 11.4 and later Events view.

1. Go to  **(Admin) > System**, and in the options panel, select **Investigation**.

   The Investigation Configuration panel is displayed.



2. In the Events tab, in the **Event Limit Default** field under **Events Panel Settings**, select the maximum number of events loaded in the Events panel when a query is submitted.
   The default and suggested value is 10,000 events, and you can select a value between 100 and 40,000 events. Increasing the limit above 10,000 events may cause excessive load on the queried service, causing a timeout or error to occur. Consider these factors before increasing the limit: number of columns in column groups, size or length of meta values associated to columns (meta keys), complexity of the search, length of the time range, and number of underlying services providing matching results to the query.
   When users want more results returned, so that their subsequent sorting and finds through that data do not require additional queries to the services, the event limit can be increased. In this case, the best practice is to configure a per-role limit using the configuration options in Event Limit per User Role so only specific users are returned more results. For example, set the global Event Limit Default to 5,000, and then create different Analyst roles that can be set to higher limits, up to the maximum 40,000 events.
3. If a query returns more events than the configured Event Limit Default, the Events panel title shows the analyst that more results are available but are not listed due to the limit. Increasing the limit may place additional load on the queried service; the ideal limit is determined by your environment.
4. Click **Apply**.

   The change becomes effective immediately, and applies to any new queries submitted by analysts.
5. Under **Event Limit Per User Role**, select the maximum number of events loaded for a single query for individual user roles. This limit must be less than or equal to the system events limit of 40,000 events; it can be larger than the default or configured limit set under Event Limit Default.
6. Click **Apply**.

The change becomes effective immediately, and applies to any new queries submitted by users assigned to the user role.

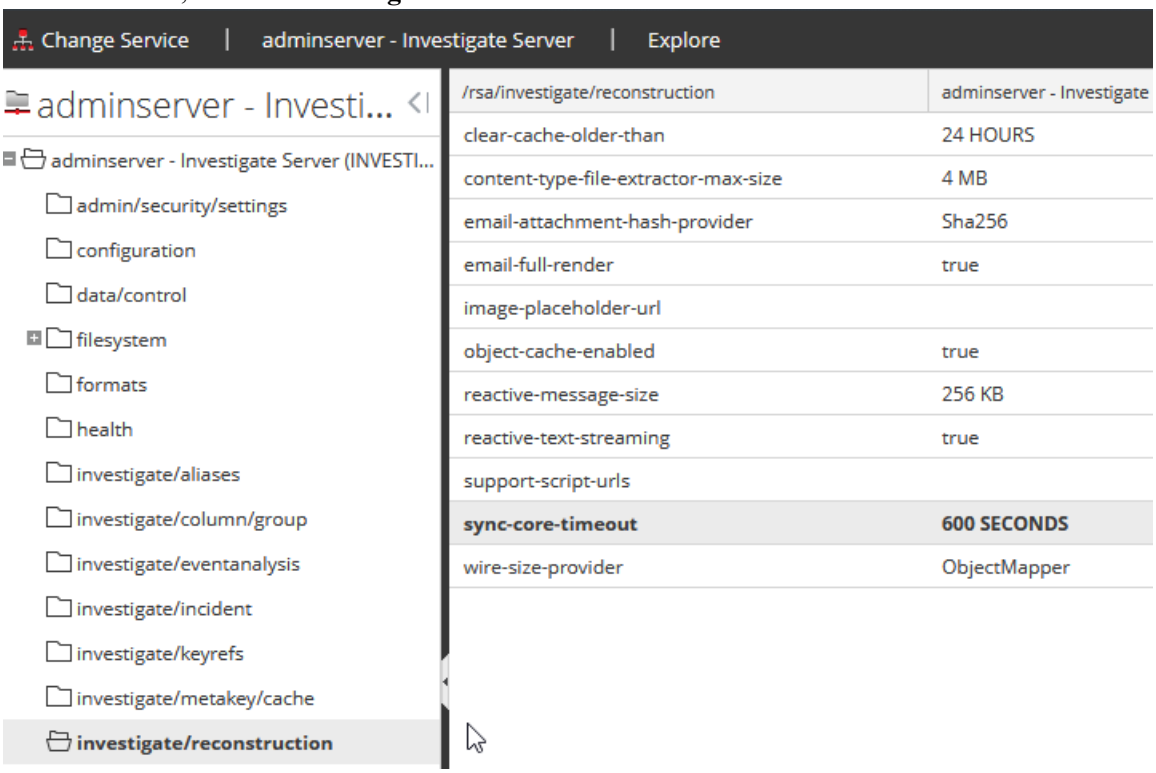
Configure the Sync Core Timeout to Remedy Deadlocks in Events View Reconstructions

The `sync-core-timeout` is a setting in the `/investigate/reconstruction` node that determines the maximum time to wait for operations for caching core content to complete to prevent deadlocks. The default value is 600 seconds (10 minutes) and needs no adjustment under most circumstances. If analysts are seeing a spinner for a very long time (>10 minutes) when loading a reconstruction in the Events view, for example from events on a 10G Decoder, increasing the length of this timeout may improve the ability to reconstruct events.

Caution: Changing the timeout setting to more than 600 seconds may lead to stability issues.

To adjust the sync-core-timeout:

1. Go to  **(Admin) > Services > Investigate-server** and **View > Explorer**.
2. In the node list, click the **investigate-reconstruction** node.



3. In the **sync-core-timeout** field, type a new value for the number of seconds before timeout and press **RETURN**.
   The setting is applied and goes into effect immediately.

Configure Live Services Settings

Options for configuring Live Services are in the System view > Live Services Configuration panel. The Live Configuration panel allows you to configure:

- The Live account.
- The Live Content update schedule and preferences for notification of updates.
- Participation in Live Services Feedback (Version 11.4.0 and earlier).
- Sharing Live Content Usage
- RSA Live Connect (Version 11.5.0 and earlier).

Prerequisite

To activate your Live account for NetWitness, please contact NetWitness Customer Support. When you have a confirmation that your Live account has been set up, you can configure and test the CMS server connection.

For information on Analyst Behaviors and Data Sharing, see "NetWitness Feedback and Data Sharing" topic in the *Live Services Management Guide*.

About Live Feedback Participation

Once you sign up for a Live account, Live Feedback automatically collects relevant information for further improvement and anonymously sends it to RSA. The shared data is protected in accordance with the applicable license agreement. For information on Live Feedback, see Live Feedback Overview. For information, see Configure Live Services Settings

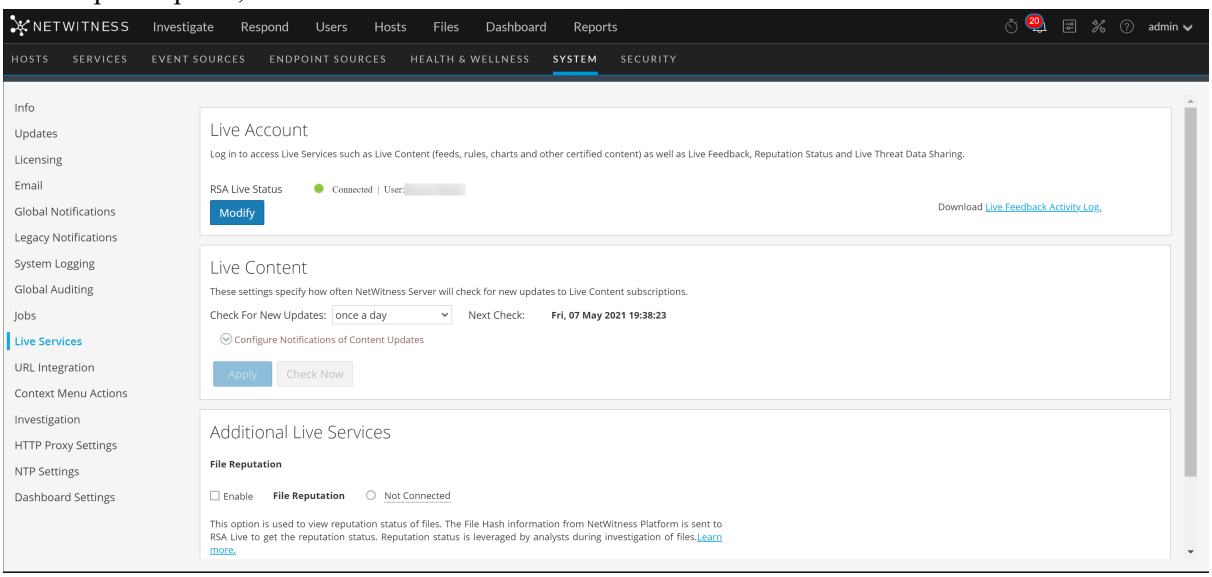
If needed, you can manually download historical usage data and share it with RSA. For information on how to download historical usage data and share it with RSA, see Upload Data to RSA for Live Feedback.

This topic contains the following procedures:

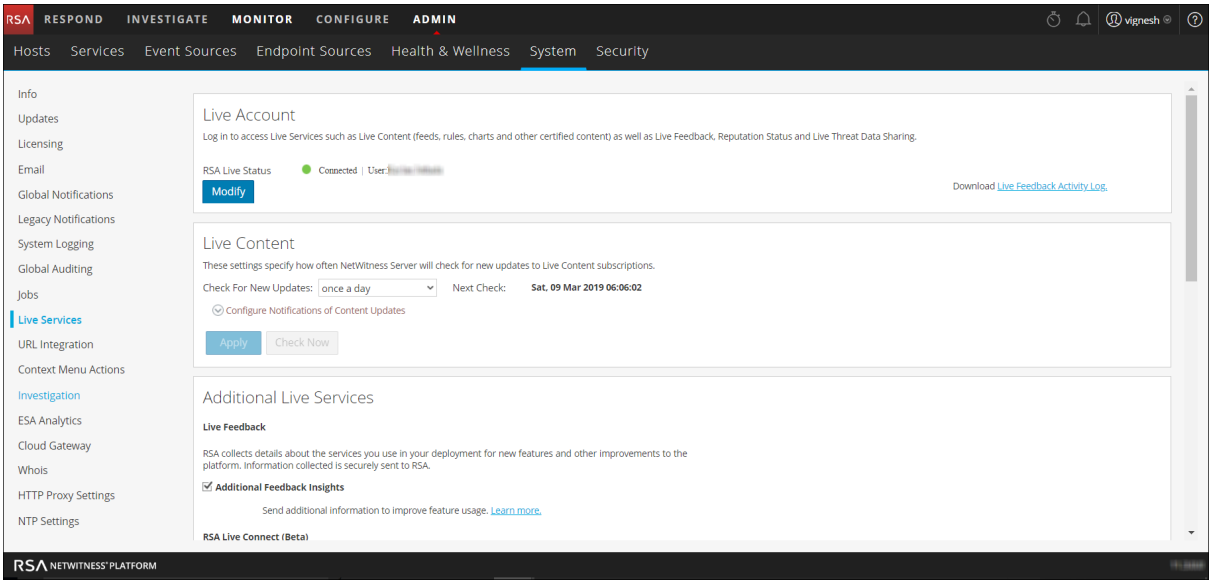
- Access the Live Services Configuration Panel
- Configure Live Account
- Configure the Live Content Synchronization Interval and Notification
- Force Immediate Synchronization

Access the Live Services Configuration Panel

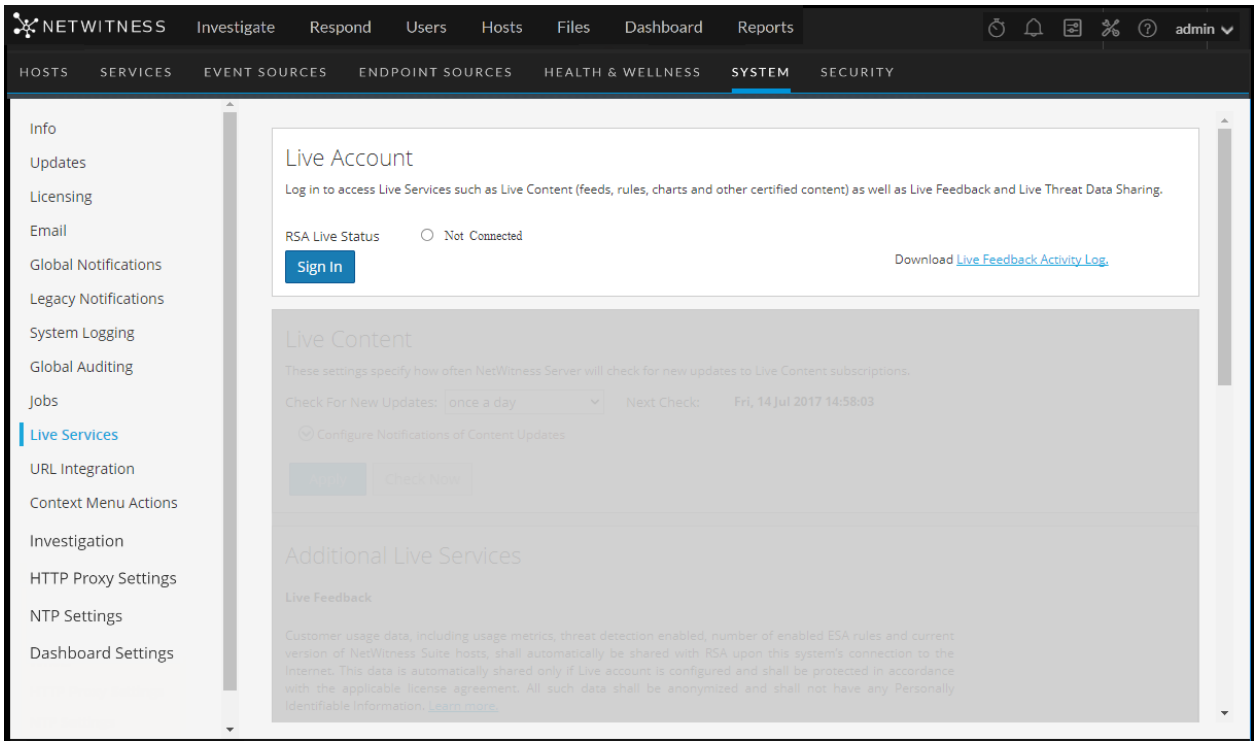
1. Go to  **(Admin) > System**.
2. In the options panel, select **Live Services**.



If Version 11.4.0.1 or earlier is installed, this is the panel.



Note: If you are not signed in with your Live Account credentials, a masked screen is displayed.



Configure Live Account

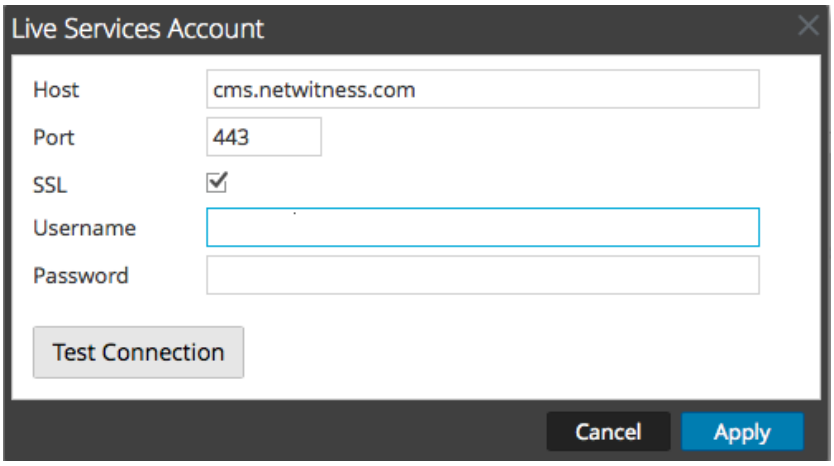
In the **Live Account** section, you must set up the user's Live account. The information needed to set up the user's Live account consists of the Username, Password, and Live URL for the Content Management System. This information is provided by Customer Care.

To configure a Live account:

1. In the **Live Account** section, click **Sign In**.

   Note: The **Modify** button shows that the live account is configured. Click **Modify** to change the user that is accessing Live Services.

2. In the Live Services Account dialog box, enter the Host (typically **cms.netwitness.com**) and type your username and password.

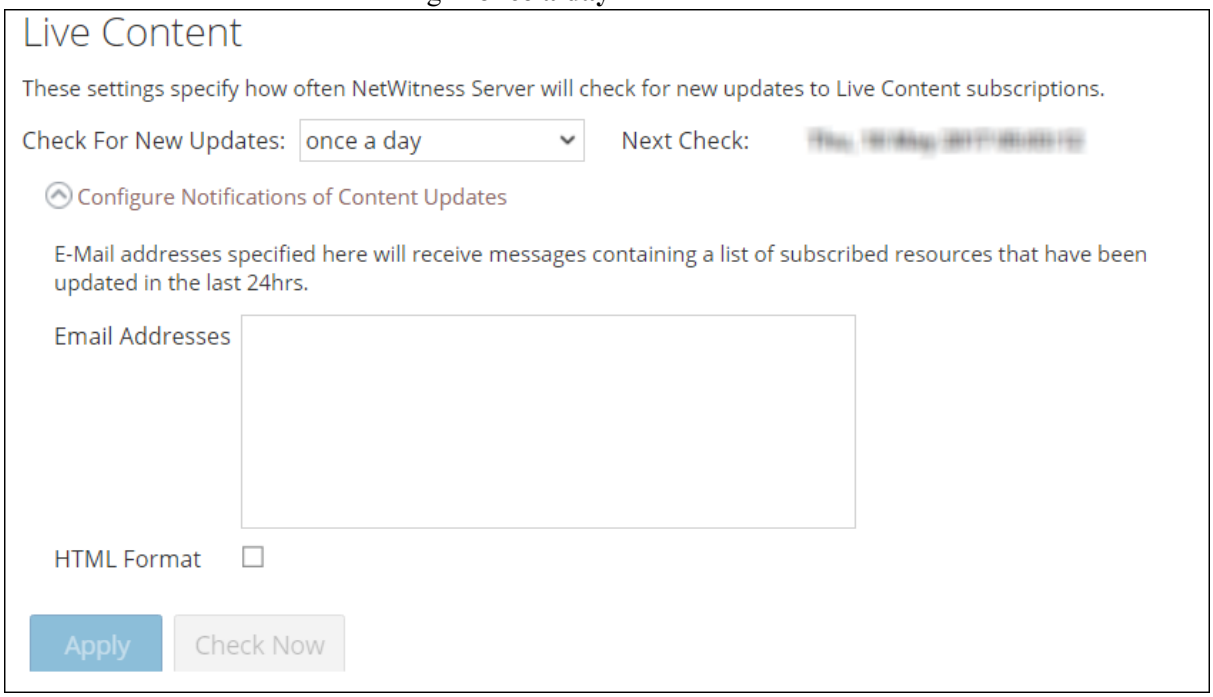


3. (Optional) If you are using a different CMS, type the host URL for the Content Management System. The default points to the CMS at **cms.netwitness.com**.
4. (Optional) If you are using a different CMS, type the communications port for Live to send requests to the Content Management System. The default for this field is **443**, which is the communications port on the Content Management System.
5. (Optional) If you do not want to use SSL, uncheck the **SSL** option. (SSL is enabled by default.)
6. Click **Test connection** to test the connection to CMS.
7. To save and apply the configuration, click **Apply**.

Configure the Live Content Synchronization Interval and Notification

You can change the interval at which NetWitness checks for new updates to Live Content:

1. Use the **Check for New Updates** field to change the interval. Select an interval from the drop-down list. The default value for this setting is **once a day**.



2. To configure Live Services to send update reports to one or more people, in the **Email Addresses** field, type the email addresses as a comma-separated list, for example, **john@company.com,ted@company.com,brian@company.com**
3. (Optional) To receive messages in HTML format rather than plain text, select **HTML Format**.
4. To save and apply, click **Apply**.

   The time and date of the next scheduled Live synchronization based on the configured interval for checking is displayed.

Force Immediate Synchronization

Instead of waiting for the next scheduled resource cycle, this option forces Live to begin immediate synchronization of the subscribed resources in this instance of NetWitness. One use for this is to see the immediate impact of a configuration change. For example, a new service has been added, or new resources have been toggled for automatic deployment. The scheduled synchronization could take place hours later if Live Services is set to synchronize a few times a day.

Caution: Synchronization can cause a parser reload if a FlexParser is deployed in the update cycle. This is acceptable once or twice a day, but a number of back-to-back parser reloads can cause packet loss at the Decoder. If this is the initial setup and you haven't configured Live resource subscriptions, do not Synchronize Now. Wait until you have configured subscriptions.

To force immediate synchronization, click **Check Now**. NetWitness checks for updates in subscribed resources.

File Reputation

File Reputation provides analysts the opportunity to view reputation status of files.

By default, **File Reputation** is enabled in **Additional Live Services** section.

Live Feedback Overview

This topic provides an introduction to Live Feedback. Live Feedback collects relevant information such as the Licensing usage data for Network Decoder, Log Decoder and Malware Analysis, Threat Detection Enabled or Disabled status, Number of enabled ESA rules,and version number details of all the services of NetWitness. For more information about the licensing usage data for Packer Decoder, Log Decoder and Malware Analysis, see the **License Details** tab topic in the *Licensing Guide*. The information is collected to improve future releases of NetWitness. You will automatically be signed on to live feedback and you cannot disable this option.

In addition to this, information on the Live Content Usage can also be shared with NetWitness. Live Content usage metrics for resource types from  **(Configure) > Live Content > Search Criteria** such as total count of NetWitness Application Rule, NetWitness Correlation Rule etc. can be shared with NetWitness. The information collected is used to improve the use of Live Content. For more information about sharing live content configuration, see Live Services Configuration Panel.

About Live Feedback Participation

Once you sign up for a Live account, Live Feedback automatically collects relevant information for further improvement and anonymously sends it to RSA. The shared data is protected in accordance with the applicable license agreement. For information on Live Feedback, see Live Feedback Overview. For information, see Configure Live Services Settings

If needed, you can manually download historical usage data and share it with RSA. For information on how to download historical usage data and share it with RSA, see Upload Data to RSA for Live Feedback.

Note: Live Feedback is activated only if you have configured your Live account.

The Live Feedback data is in JSON format as mentioned below. When you sign up with your Live Account credentials, a single encrypted JSON file is automatically uploaded to the NetWitness servers everyday.

JSON File

The JSON file consists of usage data information for a component or a set of components. In case of a set of components with the same license id, the usage data for all the components is aggregated and represented as a component called Entitlement. However, even if there is a single component such as a log decoder or decoder, an Entitlement component will be generated and will display the usage data for a single component. This aggregation is for components namely log decoders, decoders or malware analysis.

Note: The version of Entitlement is always null as it is the aggregate for a license data.

For example, if there are three Decoders with the same license id "xxx" with the following usage data:
Decoder1 = 150 MB
Decoder2 = 250 MB
Decoder3 = 100 MB
The aggregated usage data of 500 MB is displayed.

This JSON file is described in the following sections:

- Components
- Metrics
- Other Product Details
- Sample

Components

Details of each service in your NetWitness deployment. This is represented as Component. For each component the following details are displayed.

Component	Description
Version	Version number of the component in the NetWitness deployment. For example, 11.1.0.0.x.x.x.x.
ID	This is the unique Component ID that represents the host and is used to link to the metrics generated.
Properties	• **Name** - This is the name of the property for that component. For example, malware analysis, ESA, log decoder, etc. • **Value** - This is the unique value to identify the component.

Metrics

Metrics of the components (hosts) such as Log Decoder, Decoder and Malware Analysis. The license usage data for each host is shared. For Live Content usage metrics, resource types from  **(Configure) > Live Content > Search Criteria** such as total count of NetWitness Application Rule, NetWitness Correlation Rule and so on are shared.

Component	Description
Usage	• **Value** - This is the value generated for the specific component ID for each component. • **Name** - This is the name of the statistics for which the metrics is collected. For example, Capture Total Bytes.
StartTimeUTC	This is the time from when the metrics is collected. (in EPOCH format).
EndTimeUTC	This is the time when the metrics collection is complete (in EPOCH format).
Component ID	This is the ID of the component for which the value is recorded.

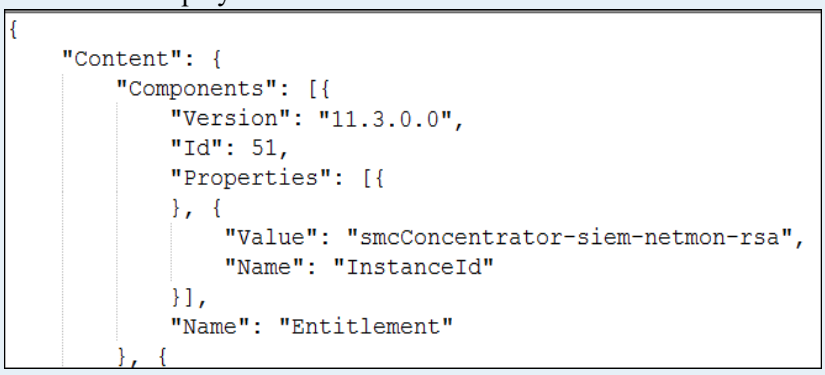
Other Product Details

- **End Time** - This is the time when the metrics collection is complete (in EPOCH format).
- **Product Name** - This is the name of the product. In this example, the Product Type is **NetWitness**.
- **Version** - This is the version of the JSON file which tracks the changes made to the file format.

- **Start Time** - This is the time from when the metrics is collected. (in EPOCH format).
- **Product Type** - This is the name of the product. In this example, the Product Type is **NetWitness**.
- **Product Version** - This is the version of the product from which the metrics is collected. In this example, the Product Version is **11.3.0.0-SNAPSHOT**.
- **Product Instance** - This is the License Server ID.
- **Checksum** - This is the information which is used for integrity checks.

The following table describes details of the JSON file with examples.

Metrics	Description
Content	Displays the content that contains all the Components, Metrics, Product Type and Product Instance data except Checksum.

Metrics	Description
Components	The details of all the services in NetWitness are represented as a Component. The details of the component such as the version number of the component, the name, and the value is displayed.  **Version**: Displays the version of NetWitness service. For example, 11.3.0.0. **ID**: Displays an unique id which is generated for the NetWitness service and is used to link to the metrics for that particular component. In this example, the ID for Malware Analysis is 5 and the metrics is displayed for ComponentId 5 in bytes:  **Properties**: Displays the properties for the component such as name and value as shown in the above figure. **Value**: Displays the value of the property which is an internal UUID for a component as shown in the above figure This is generated by NetWitness. For example, For malware analysis the value displayed as "55f7a0b30e502231c42d063f" **Name: "InstanceId"**: Displays the name of the property as shown in the above figure.

Metrics	Description
	Name": "malwareanalysis": Displays the name of component which is a service name such as LogDecoder, Decoder, or MalwareAnalysis.

Metrics	Displays the list of the metrics with the usage data for components namely MacHosts, LinuHosts and WinHosts.
	In this example, the metrics is displayed for ComponentId 1 in bytes.
	
	}],
	"Metrics": [{
	"Usage": [{
	"Value": "0.0",
	"Name": "MacHosts"
	}, {
	"Value": "0.0",
	"Name": "LinuxHosts"
	}, {
	"Value": "0.0",
	"Name": "WinHosts"
	}, {
	"Value": "0.0",
	"Name": "TotalHosts"
	}],
	"StartTimeUTC": 1539043200000,
	"EndTimeUTC": 1539129599000,
	"ComponentId": 1
	}, {

	StartTimeUTC: Displays the time when the metrics is collected, in the EPOCH format.
	Usage: Displays the usage value and usage type statistics of the component.
	Value: Displays the value of the statistics. For example, "Value": "1582940012678" as shown in the above figure.
	Name: Displays the name of the statistics. For example, Capture Total Bytes or Total File bytes.
	EndTimeUTC: Displays the time when the metrics collection is complete, in the EPOCH format.
	ComponentId: Displays the component id for which the metric values are collected. This is the same as the "ID" in the Components section.
Content	Displays the content that contains all the Components, Metrics, Product Type and Product Instance data except Checksum.
ProductType	Displays the product type that generates the file. For example, `"ProductType": "NetWitness Platform"`
ProductInstance	Displays the License server Id and is unique per NetWitness. For example, `"ProductInstance": "00-0C-29-6C-66-E3"`

Metrics	Description
Checksum	Displays the Checksum for the "Content" section in the file. Used by NetWitness for integrity check. For example, `"Checksum": "883DACF97E4BCD9F590A1461A4DD0A312B5883A6CF82E0518E77AAB6A6DDB654"`

Example

Here is a sample JSON file.


```
{
    "Content": {
        "Components": [{
            "Version": "11.3.0.0",
            "Id": 7,
            "Properties": [{
                "Value": "57470c96e4b0cf62c7bfbd53",
                "Name": "InstanceId"
            }],
            "Name": "esa"
        },
        {
            "Version": "11.3.0.0",
            "Id": 4,
            "Properties": [{
                "Value": "5714c78be4b0ea5bd2b96e69",
                "Name": "InstanceId"
            }],
            "Name": "incidentmanagement"
        },
        {
            "Version": "11.3.0.0",
            "Id": 2,
            "Properties": [{
                "Value": "5714c78be4b0ea5bd2b96e65",
                "Name": "InstanceId"
            }],
            "Name": "sa"
        },
        {
            "Version": "11.3.0.0",
            "Id": 1,
            "Properties": [{
                "Value": "5714c78be4b0ea5bd2b96e63",
                "Name": "InstanceId"
            }],
            "Name": "malwareanalysis"
        },
        {
            "Version": "11.3.0.0",
            "Id": 3,
            "Properties": [{
                "Value": "5714c78be4b0ea5bd2b96e67",
                "Name": "InstanceId"
            }],
            "Name": "reportingengine"
        }],
        "Metrics": [{
            "StartTimeUTC": 1464480000000,
            "Stats": [{
                "Value": "Disabled",
                "Name": "Threat Detection"
            },
            {
                "Value": "3.0",
                "Name": "Number Of Enabled ESA Rules"
            }],
            "EndTimeUTC": 1464566399000,
            "ComponentId": 7
        }],
        "EndTime": 1464566399000,
        "Version": "1.0",
        "StartTime": 1464479999000,
        "ProductType": "Security Analytics",
        "ProductInstance": "00-0C-29-A2-57-B4"
    },
    "Checksum": "6445C704D3F9E67D24DBA8F11EB6C003CBCC0E199576342E6E6D2545524F583F"
}
```


The JSON file includes details of all the licenses currently available on the appliance. Here is a sample of the Entitlement information within the JSON file for a appliance license for Broker.

```
    }, {
        "Version": "2015.0506",
        "Id": 14,
        "Properties": [{
            "Value": "M133206102",
            "Name": "SerialNumber"
        }, {
            "Value": "Broker",
            "Name": "DeviceType"
        }, {
            "Value": "PERPETUAL",
            "Name": "FeatureType"
        }, {
            "Value": "-1",
            "Name": "Threshold"
        }, {
            "Value": "1000654868",
            "Name": "AccountId"
        }, {
            "Value": "B02E-03A1-08A6-EC3B",
            "Name": "ActivationId"
        }, {
            "Value": "2015-05-05 00:00:00",
            "Name": "LicenseStartDate"
        }, {
            "Value": "permanent",
            "Name": "LicenseEndDate"
        }, {
            "Value": "20t-52osb7",
            "Name": "FeatureId"
        }, {
            "Value": "smcBroker",
            "Name": "Name"
        }, {
            "Value": "20t-52osb7",
            "Name": "InstanceId"
        }],
        "Name": "Entitlement"
    }, {
```


Upload Data to RSA for Live Feedback

This topic provides instructions for a NetWitness administrator to export the metrics in NetWitness for Live Feedback.

If the Live Account is not configured, you can manually upload the usage data to RSA. For more information, see Live Services Configuration Panel.

The Live Account section has a Live Feedback Activity Log which enables you to download the usage data required for Live Feedback. This is active regardless of the Live Account configuration.

You can first download the Live Feedback historical data, and then upload it to share with RSA.

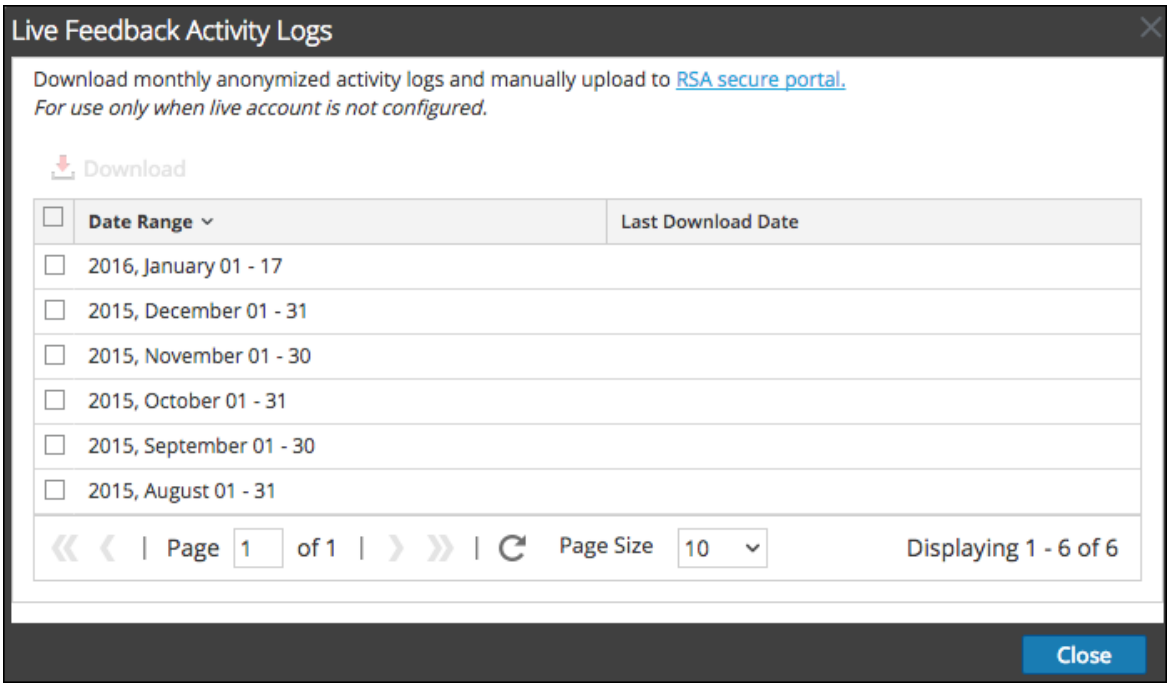
Download Live Feedback Historical Data

To download the Live Feedback historical data:

1. Go to **(Admin) > System**.
2. In the options panel, select **Live Services**.

   The **Live Account** screen is displayed which consists of the **RSA Live Status** and Download **Live Feedback Activity Log**.
3. Click the **Live Feedback Activity Log**.

   The **Live Feedback Activity Log** window opens which allows the NetWitness user to download the required Live Feedback historical data.



4. Select one or multiple entries by setting the checkboxes and click **Download**.

   Note: If you select multiple entries in the history table, the downloaded zip file consists of an individual JSON file for each month.

The downloaded Live Feedback data is in JSON format, and is bundled as a .zip file. For more information, see Live Feedback Overview.

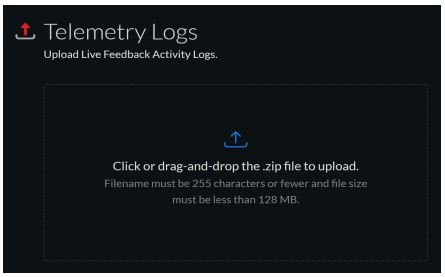
Share Data with RSA

After you download the Live Feedback data, you can then upload it using the following procedure.

Note:
- To download the Live Feedback data, see topic Download Live Feedback Historical Data.
- You can share data using the new live registration portal. For more information, Create Live Account.

To share the data to RSA:

1. Log in to the NetWitness Platform Live using your credentials.
2. Click  on the left panel.
   The **Telemetry Logs** dialog is displayed.



Note:
- You can upload only .zip files.
- Filename must be 255 characters or less and file size must be less than 128 MB.

3. Click or drag-and-drop a file onto this area to upload.

Configure Log File Settings

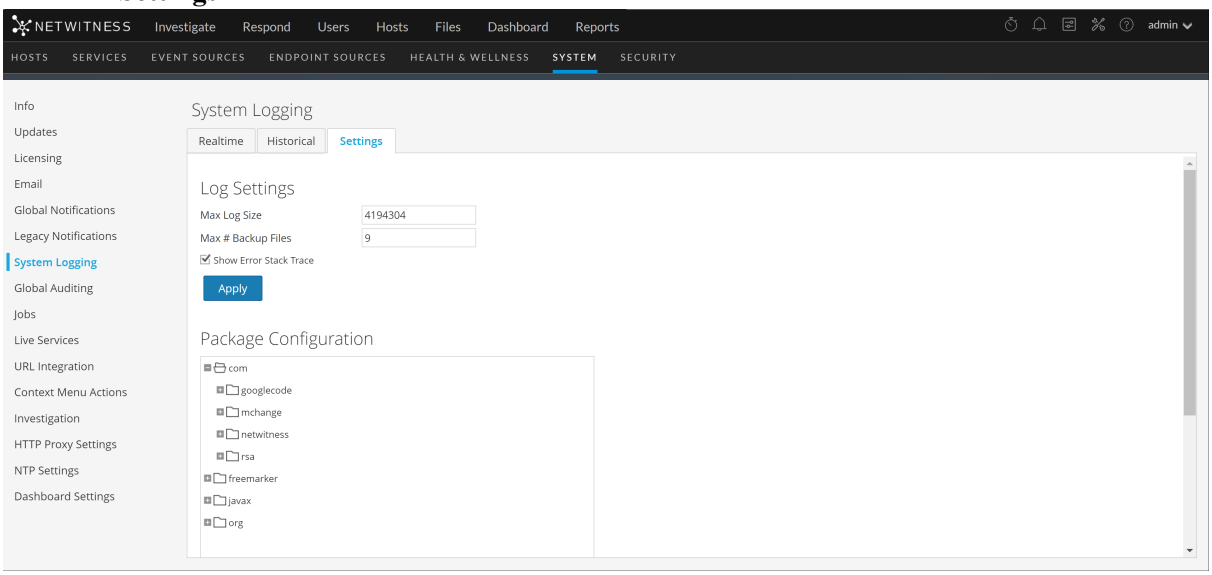
In NetWitness Platform, you can configure the size of the log files, the number of backup log files maintained, as well as the default logging levels for the packages within NetWitness.

Configure System Log File Size and Backup Count

The log file size and backup count are configured with default values. If you want to change the default values for the log file size and number of backups:

1. Go to  **(Admin) > System.**
2. In options panel, select **System Logging**.
   The System Logging Configuration panel opens to the Realtime tab by default.

3. Click the **Settings** tab.

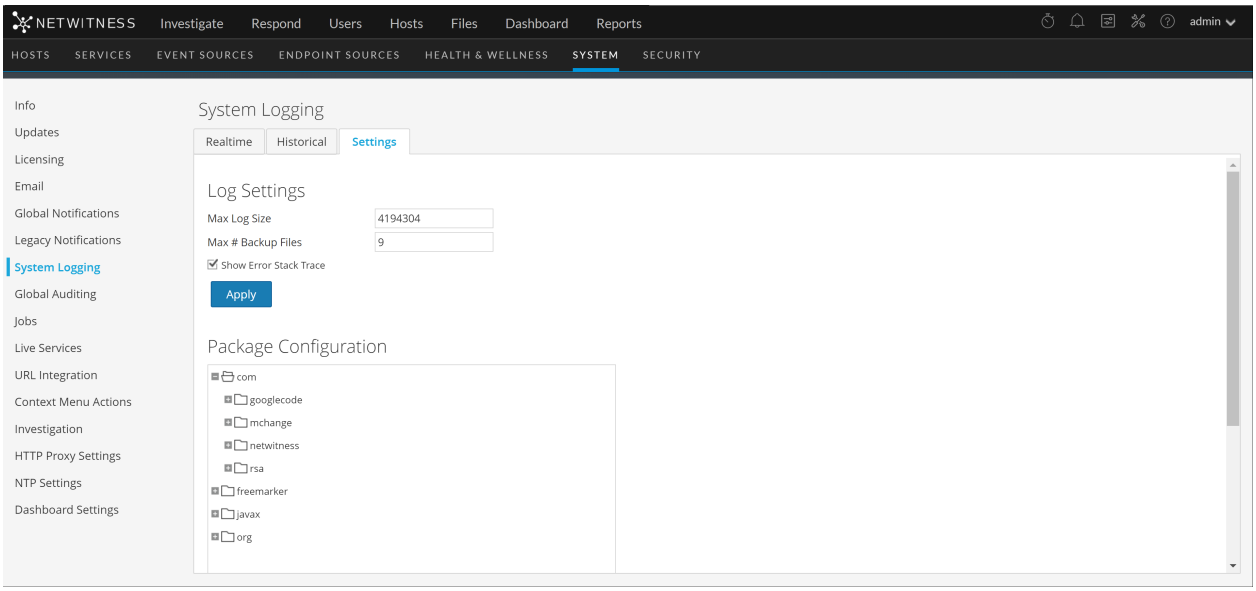


4. In the **Max Log Size** field, type the maximum size in bytes. The minimum value for this setting is **4096**.
5. In the **Max # Backup Files** field, type the maximum number of backup logs to maintain. The minimum value for this setting is **0**. When the maximum number of log files is attained, and a new backup file is made, the oldest backup is discarded.
6. Click **Apply**.
   The changes go into effect immediately.

Set the Log Level for an Individual Package

The Package Configuration section shows the NetWitness Network in a tree structure. The tree contains all the packages used within NetWitness. You can drill down into the tree to view the log levels of each package. The log level for all packages that are not explicitly set is the same as the **root** log level. To set the log level for a package:

1. Select the package in the **Package** tree.
   The name of the package is displayed in the **Package** field. If a log level is already set for the package, that level is shown.



2. Select the **Log Level** in the drop-down list.
3. Click **Apply**.
   The new log level becomes effective immediately.
4. (Optional) If you want to revert to the default log level specified for **root**, click **Reset**.

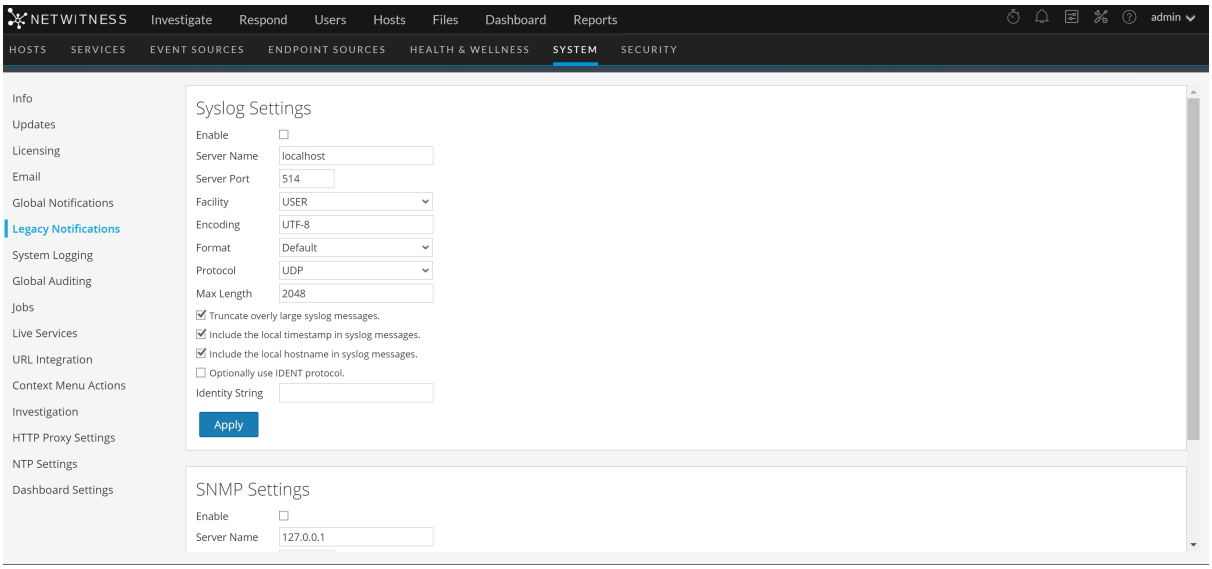
Configure Syslog and SNMP Settings

On the Legacy Notifications panel, you can configure syslog and SNMP notification settings. These configurations are used for Entitlement, legacy Event Source Management (ESM), Warehouse Connector monitoring, and Archiver monitoring.

Configure and Enable Syslog Settings

1. Go to  **(Admin) > System**.
2. In the options panel, select **Legacy Notifications**.

   The Legacy Notifications Configuration panel is displayed.



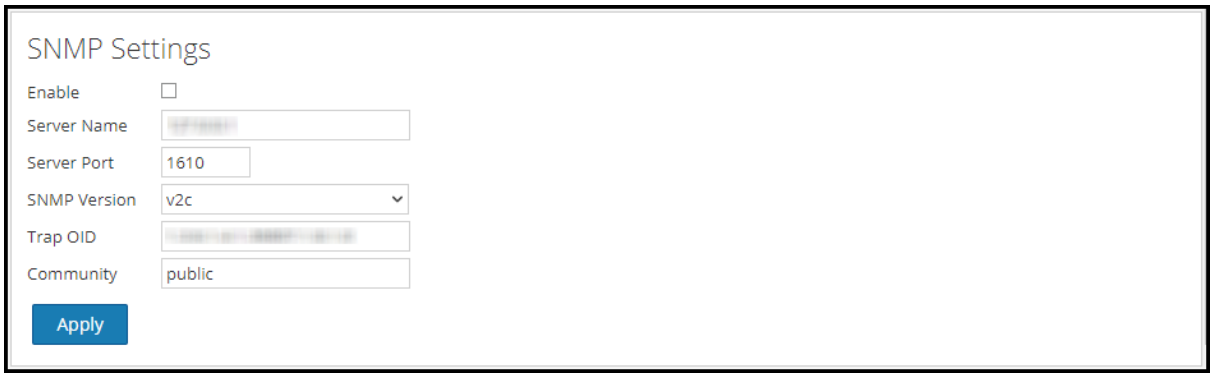
3. In the **Server Name** and **Server Port** fields under **Syslog Settings**, type the host name where the target syslog process is running and the port where the target syslog process is listening.
4. In the **Facility**, **Encoding**, **Format**, and **Max length** fields, specify the syslog facility, message text encoding, message format, and maximum message length.
5. In the **Protocol** field, select either UDP or TCP.
6. (Optional) Select the options for what to include in messages: **Truncate overly large syslog messages**, **Include the local timestamp in syslog messages**, and **Include the local hostname in syslog messages**.
7. (Optional) Configure syslog to prepend an Identity String before each syslog alert.
8. Set the **Enable** checkbox.
9. Click **Apply**.

Syslog notifications are immediately enabled. Legacy Notifications Configuration Panel provides detailed information about these settings.

Configure and Enable SNMP Settings

1. Go to  **(Admin) > System**.
2. In the options panel, select **Legacy Notifications**.

   The Legacy Notifications Configuration panel is displayed, with SNMP Settings at the bottom of the panel.



3. In the **Server Name** and **Server Port** fields under **SNMP Settings**, type the host name and listening port of the SNMP trap host.
4. Select the **SNMP version** in the drop-down menu, **v1** or **v2c**.
5. In the **Trap OID** field. specify the object ID for the SNMP trap on the trap host that receives the audit event. The default value is **0.0.0.0.0.1**.
6. In the **Community** field, specify the community string used to authenticate on the SNMP trap host, the default value is **public**.
7. Set the **Enable** checkbox.
8. Click **Apply**.

SNMP notifications are immediately enabled. Legacy Notifications Configuration Panel provides detailed information about these settings.

Disable Syslog or SNMP Settings

To disable syslog or SNMP settings on this NetWitness instance:

1. Clear the appropriate **Enable** checkbox.
2. Click **Apply**.
   The selected settings are immediately disabled.

Additional Procedures

Additional procedures are not essential for the set up of NetWitness, they include certain customization options that are beyond the usual setup; for example, adding custom context menus or setting up a proxy.

Add Custom Context Menu Actions

Configure NTP Servers

Configure Proxy for NetWitness Platform

Add New Configuration Dialog

Supported CEF Meta Keys

Supported Global Audit Logging Meta Key Variables

Global Audit Logging Operation Reference

Local Audit Log Locations

Add Custom Context Menu Actions

In the Context Menu Actions panel, Data Privacy Officer, Administrator, Analyst, and SOC Manager can view, add, edit, delete, import, and export context menu actions for the current instance of NetWitness. Each context menu action applies to a specific context in the NetWitness user interface, and appears as an option when you right-click a specific location in the user interface.

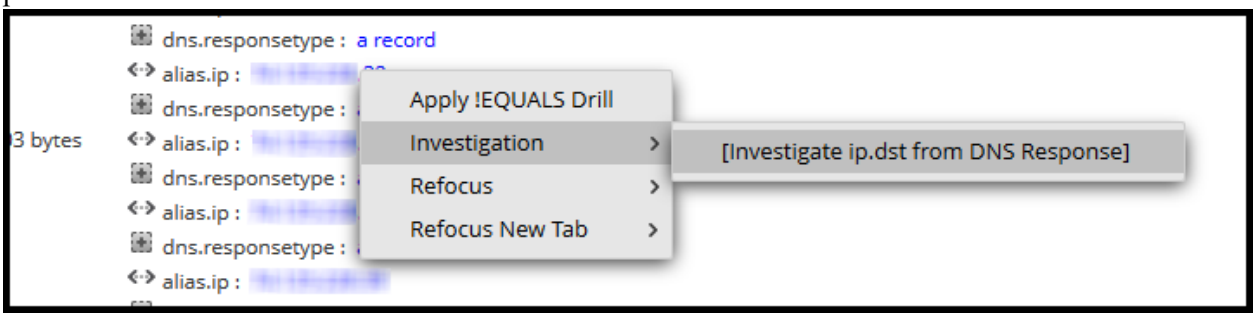
Note: All the context menu actions available in Events in NetWitness Suite 11.1 are also now available under Events in NetWitness Platform 12.0.

If you want to create a custom variation of a built-in context menu action, you can copy the configuration to a new context menu action and modify the custom context menu action. To copy, switch to the Advanced view, open the action and copy the JSON configuration file, create a new action/edit an existing action and paste. A context menu action is defined by:

- Action: The title of the action in the context menu.
- Component: The NetWitness module in which the context menu is available.
- Meta key: The content to which the action applies.
- Definition: The definition of the action.

Note: All context menu actions created before you upgrade to 11.3, functions as configured.

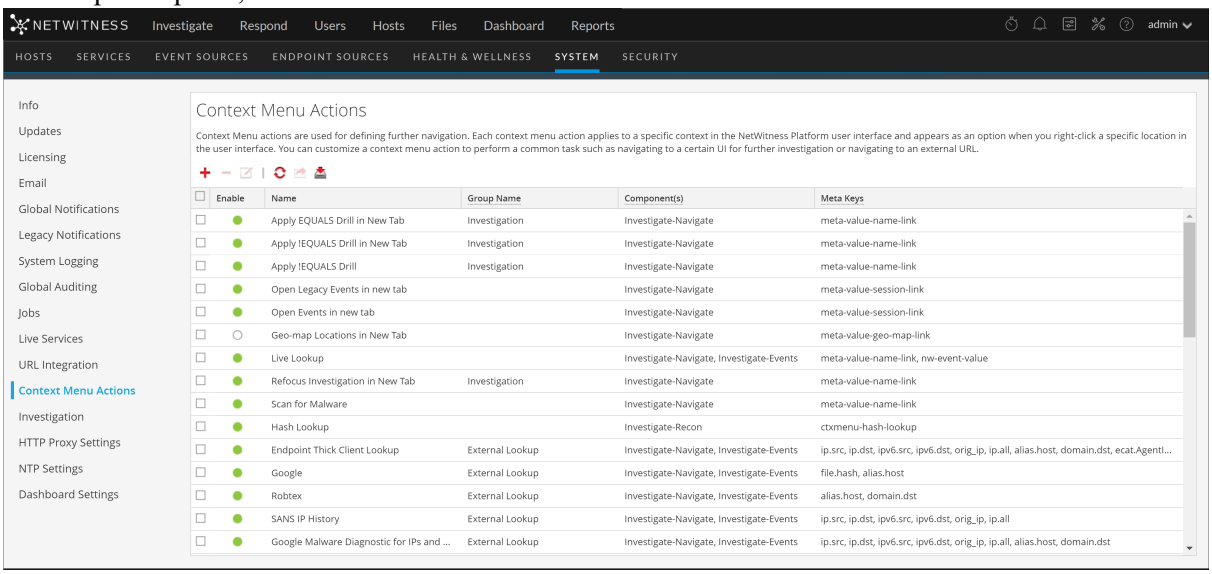
This is an example of a custom context menu action; the steps to create this example are provided as a procedure below.



View Context Menu Actions in NetWitness

To view existing context actions in NetWitness both default and custom:

1. Go to **(Admin) > System.**
2. In the options panel, select **Context Menu Actions**.

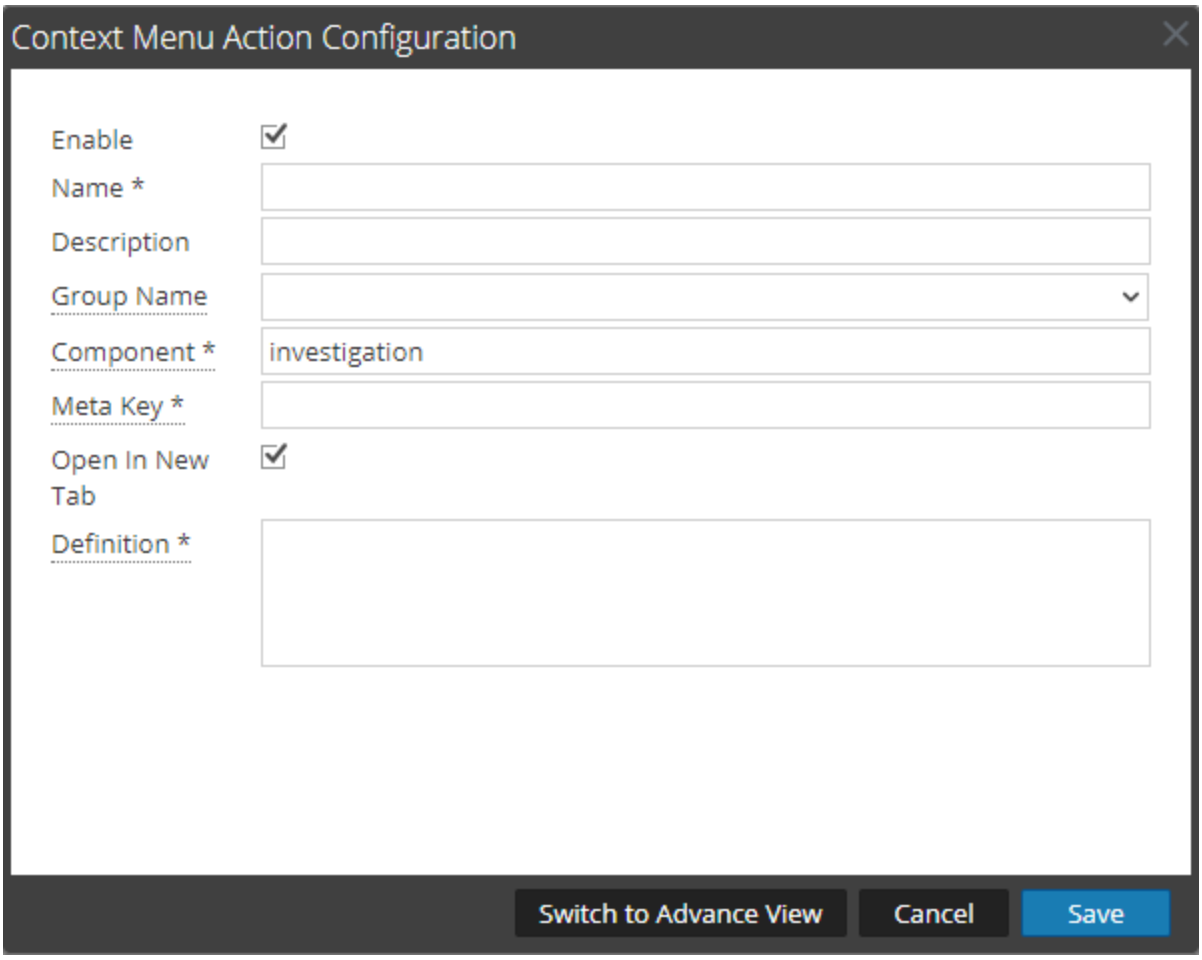


All the new actions which were available in NetWitness Suite 11.1 in the Investigate > Events tab can now be configured using the context menu actions. Details of the information in the Context Menu Action panel are provided in Context Menu Actions Panel.

Add a Context Menu Action

To add a context menu action in NetWitness:

1. In the toolbar, click .
   The Context Menu Action Configuration dialog is displayed.



Fill the required fields:

a. Enable: Select Enable to enable this context menu action.

b. Name: Enter the name of the context menu action.

c. Description: Enter a description of the context menu action.

d. Group Name: Select the group name from the drop-down menu. The action appears under this group in Context menu.

e. Component: The name of the component under which action will appear in the user interface. For example. under Investigate, the Context menu action can appear under Investigate-Navigate, Investigate-Legacy Events, Investigate-Event Recon and Investigate-Events.

Note: The Investigate-Legacy Events option and related data is displayed only if the Enable Legacy Events checkbox is enabled under **(Admin)** > **System** > **Investigation** > **Legacy Events**.

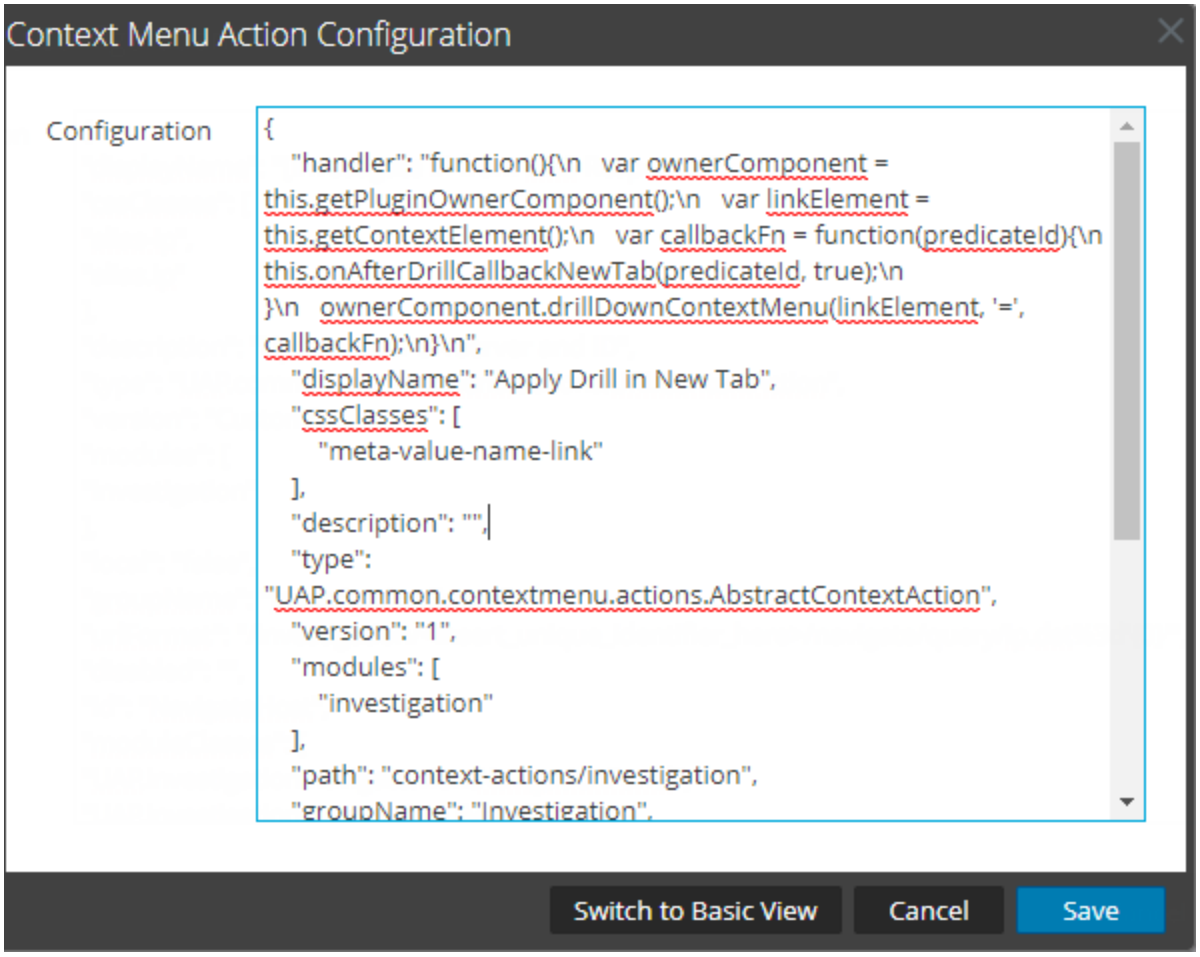
f. Meta Key: Enter the meta key separated by commas to further narrow-down scope for the context menu action. The action will appear on these meta key. Context menu actions have to be defined specifically for each meta key, the key references in a meta key do not inherit a context menu actions. For example, a context menu action created for ip.all are not created for ip.src as well. A

separate action has to be created for the sub-category or key reference of a meta.

g. Open in New Tab: Select this option to open the context menu action in a new tab.

h. Definition: Enter further action performed for this context menu action. For example, open a certain user interface or navigate to an external URL.

2. You can also type the CSS code to define the context menu action. The example procedure at the end of this topic provides step-by-step instructions that you can use to create a useful context menu action. Click **Switch to Advance View** to add the context menu action.



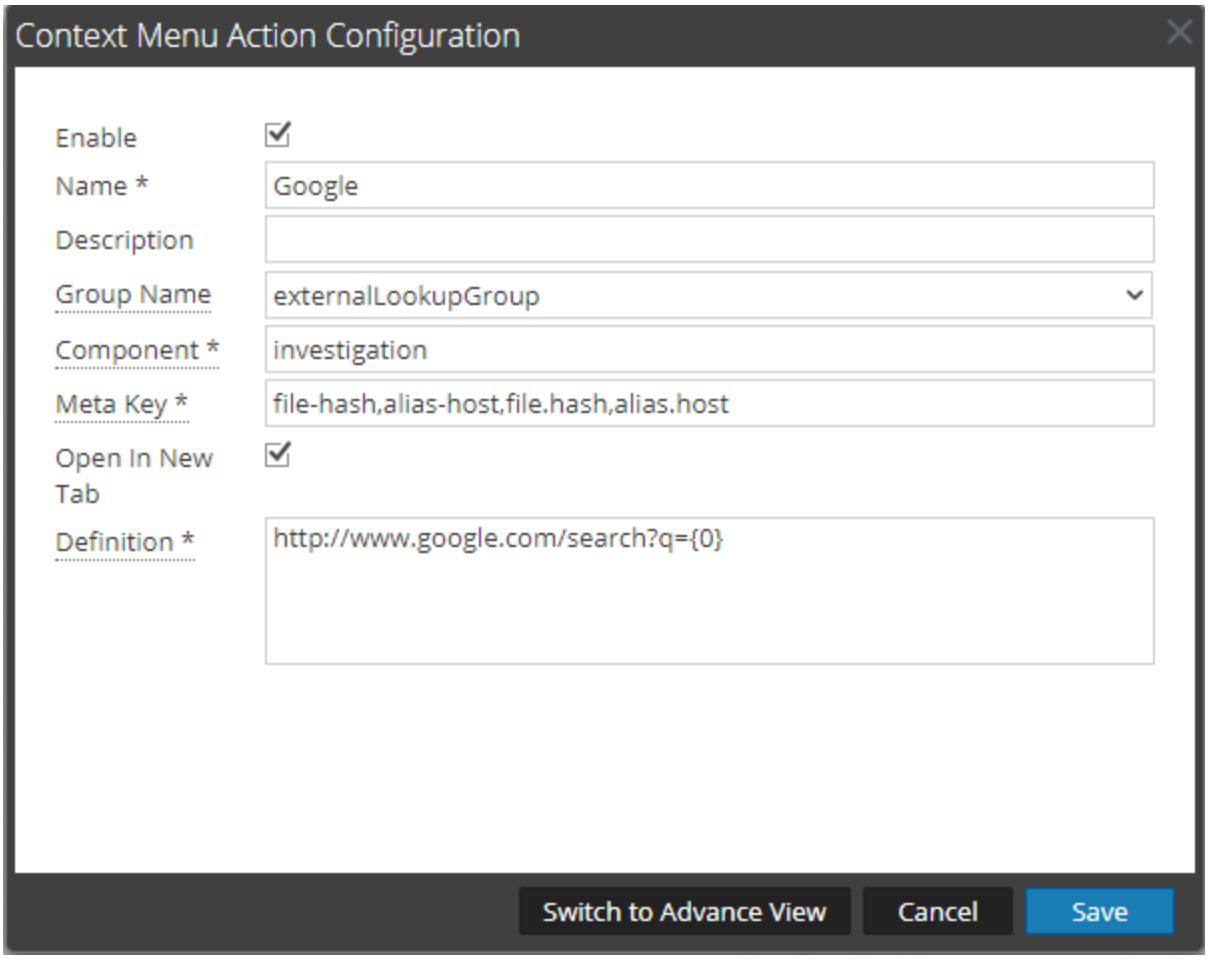
3. Click **OK.**
The new context menu action is created and added at the end of the list of context menu actions.

4. The context menu action becomes available in the configured location.

Edit a Context Action

To edit a context action:

1. Select the row in the grid and either **double-click** the row or click .
The **Context Menu Action Configuration Dialog** is displayed.



2. Edit the **Configuration**.
3. To save the changes, click **OK**.

Delete a Context Action

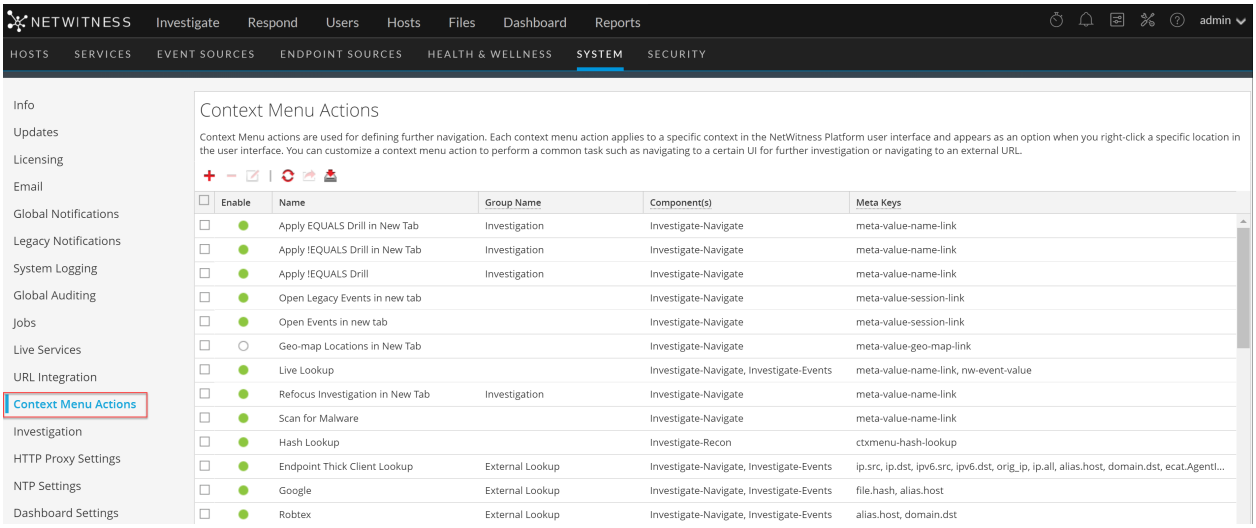
To remove a context menu action from NetWitness entirely:

1. Select the action.
2. Click .
   A dialog requests confirmation that you want to delete the context menu action.
3. Click **Yes**.
   The option is removed from the Context Menu Actions panel.

Export Context Menu Actions

You can export context menu action to a zip file. The zip file contains the JSON files with each each JSON file mapping to a context menu action. To export the context menu action, follow these steps:

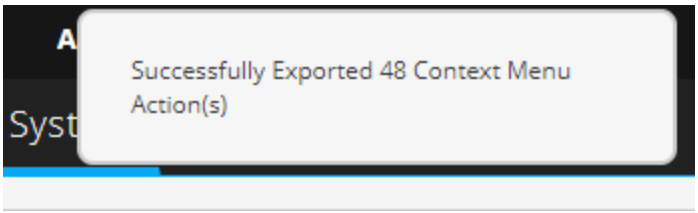
1. Go to **(Admin) > System**.
2. Click **Context Menu Actions**.



3. Click to select a context menu action to import. Click the header to select ALL the context menu actions.



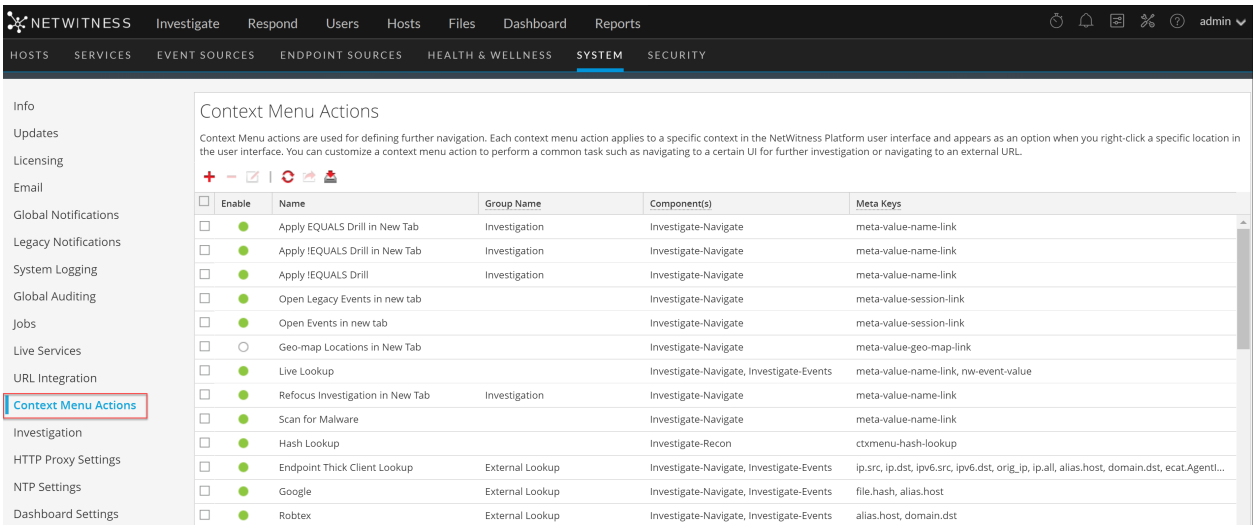
4. Click Export Action(s) under Context Menu Actions.
5. The success message confirming the actions uploaded successfully is displayed.



Import Context Menu Actions

You can import ccontext actions in Context Menu Actions tab. These actions can then be edited or used as is for investigating context where applicable. Follow these steps to import a context menu action(s):

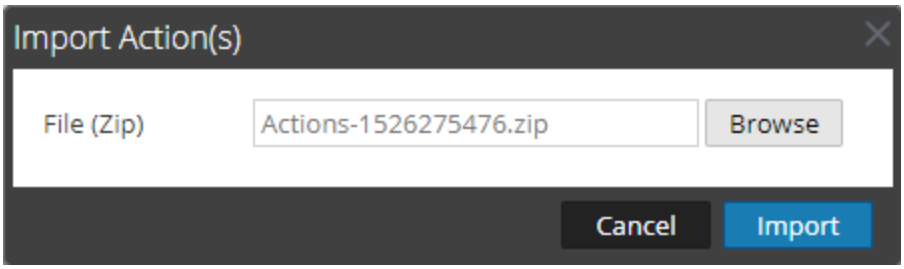
1. Go to **(Admin) > System**.
2. Click **Context Menu Actions**.



3. Click Import Action under Context Menu Actions.
4. In Import Action click **Browse** to locate and select the file. The zip file typically contains the json files containing context menu actions exported previously.

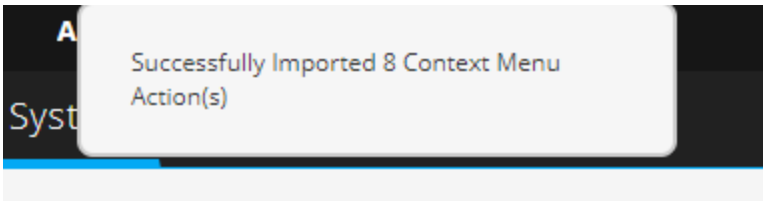


5. Select the Zip file and click **Open**.
6. Click **Import**



Note: There is no validation for an action for Events with a Javascript function.

7. The success message confirming the actions uploaded successfully are displayed.



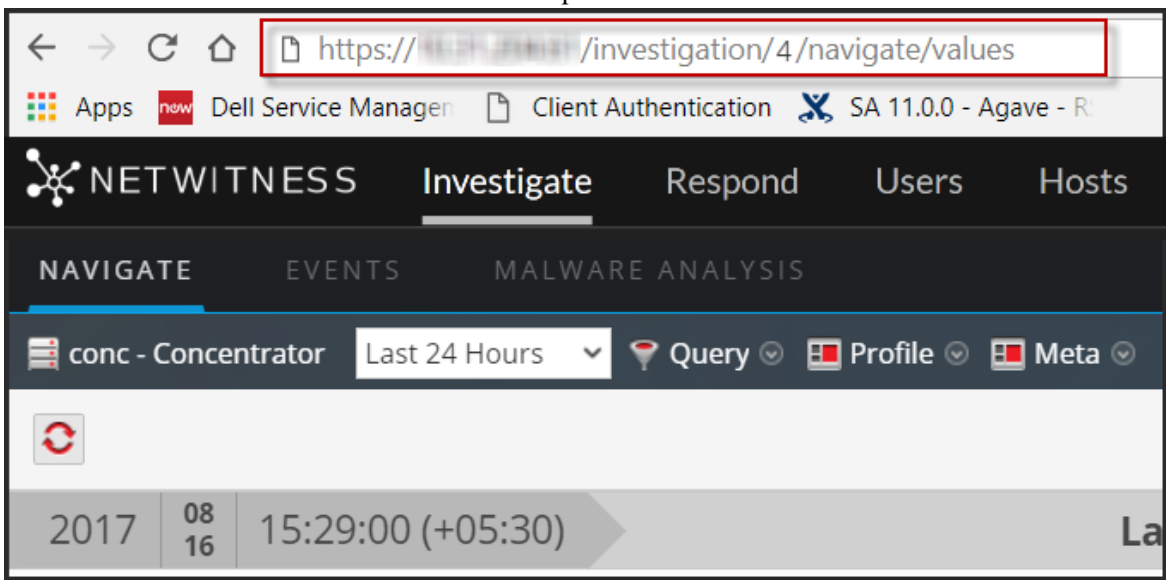
Note: If an error message is displayed, check the log files and try importing the context menu actions file again.

Example Procedure: Context Menu Action to Investigate ip.dst from alias.ip

This example adds a context menu action that allows analysts to pivot from the `alias.ip` values (the IP addresses returned from a DNS request) to the `ip.dst` meta key. It helps analysts to locate any detected traffic to the IP address that was returned for a DNS query.

To implement the context menu action:

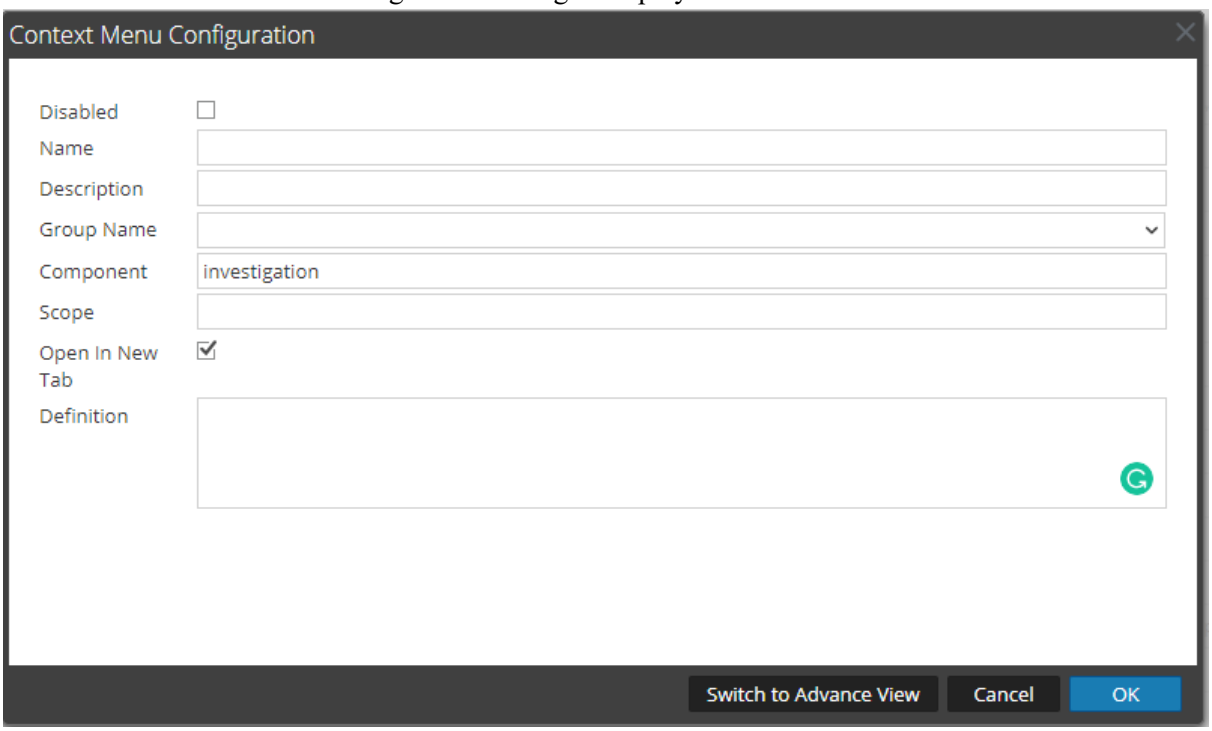
1. Determine the unique identifier for your NetWitness Server as follows:
   a. Log onto NetWitness , go to **Investigate > Navigate**, choose a service (for example, a Concentrator) to investigate, and wait for the values to load.
   b. Look for the URL and locate the number after `investigation`. In this example, the unique identifier for the action is 4. You need this unique identifier to add to the context menu action.



2. Go to **(Admin) > System > Context Menu Actions**

   In the toolbar, click .

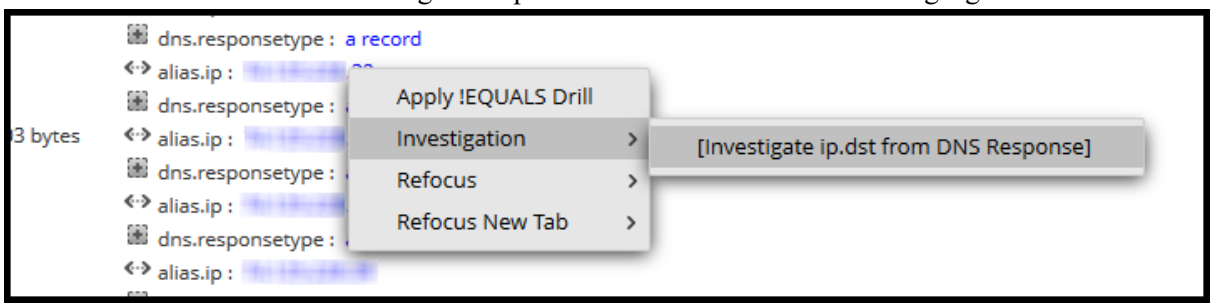
The Context Menu Action Configuration dialog is displayed.



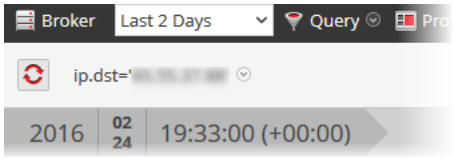
3. Copy the entire sample code block below and paste it in the window.

```
{
    "displayName": "[Investigate IP from DNS Response]",
    "cssClasses": [
        "alias-ip",
        "alias.ip"
    ],
    "description": "Update your NW server and ID",
    "type": "UAP.common.contextmenu.actions.URLContextAction",
    "version": "Custom",
    "modules": [
        "investigation"
    ],
    "local": "false",
    "groupName": "investigationGroup",
    "urlFormat": "/investigation/<insert_unique_identifier_
here>/navigate/query/ip.dst%3d'{0}'",
    "disabled": "",
    "id": "NavigateHost",
    "moduleClasses": [
        "UAP.investigation.navigate.view.NavigationPanel",
        "UAP.investigation.events.view.EventGrid"
    ],
    "openInNewTab": "true"
}
```

4. In the **urlFormat** line replace **<insert-unique_identifier_here>** with your unique identifier. The URL should look like this:
   `"/investigation/4/navigate/query/ip.dst%3d'{0}'"`
5. Click **OK**, and restart your browser.
6. To test the action, open an investigation in the Navigate view and right-click on the meta key `alias.ip`.
   The context menu with the Investigation option should look like the following figure.



7. Should produce a pivot like this.



8. If you are using this example for DNS traffic investigation, you may want to consider creating a meta group specific to DNS traffic as described in "Manage User-Defined Meta Groups" in the *NetWitness Investigate Guide*.

Configure NTP Servers

This topic provides instructions on how to configure Network Time Protocol (NTP) servers. NTP is a protocol designed to synchronize host machine clocks over a network. For more information on NTP go to their home page (http://www.ntp.org/).

Note: NetWitness Server Core hosts must be able to communicate with the NetWitness Server host with UDP port 123 for NTP time synchronization.

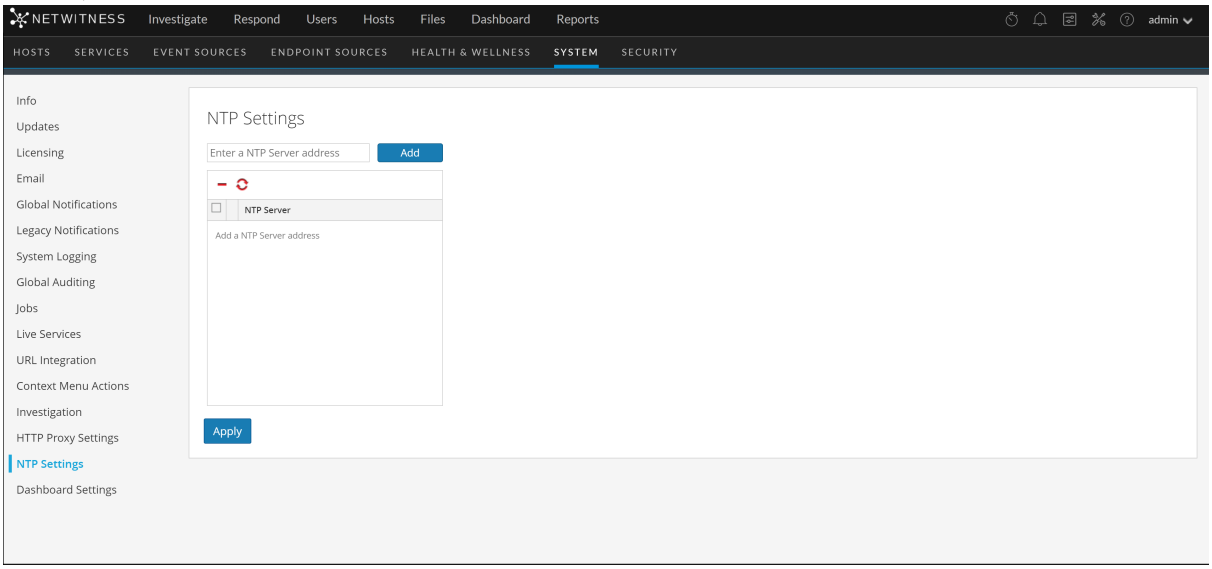
You use the **(Admin)** > **System** > **NTP Settings** view to configure one or more NTP servers. After you configure an NTP server, NetWitness uses NTP to synchronize the host machine clocks. You configure multiple NTP servers for Fail Over purposes. This topic contains the following procedures:

- Add an NTP Server
- Modify an NTP Server

Add an NTP Server

To add an NTP server:

1. Go to  **(Admin) > System**.
2. In the options panel, select **NTP Settings**.
   The NTP Settings panel is displayed prompting you to enter the hostname (that is, the IP Address or FQDN) of an NTP server.



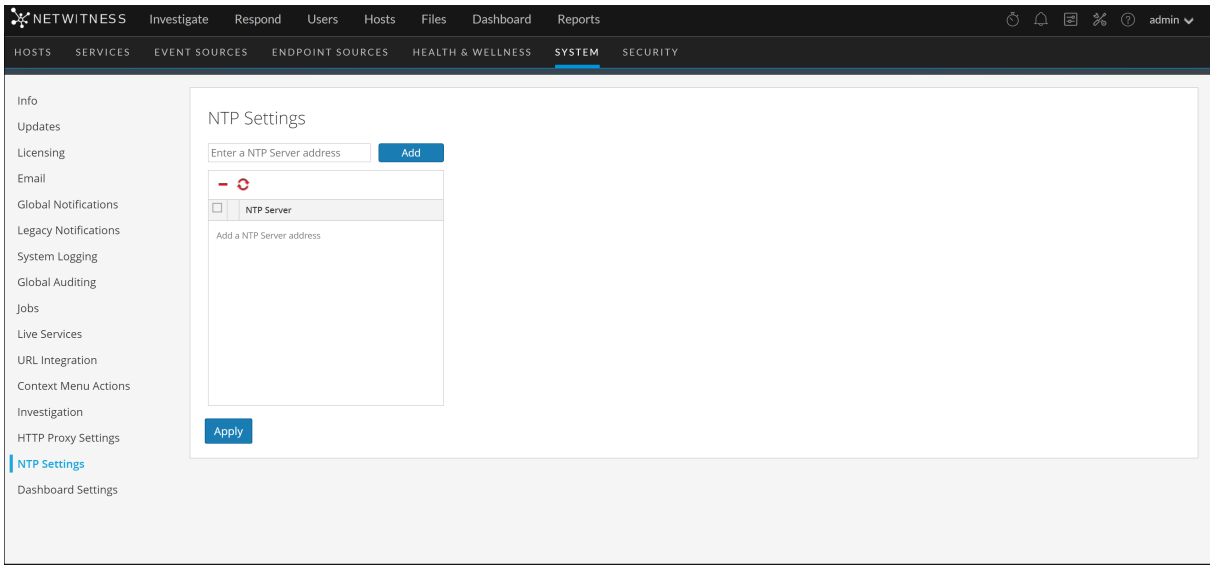
3. Enter the IP address or FQDN for an NTP server.
   If the hostname syntax is invalid, NetWitness disables the **Add** and **Apply** buttons and displays **Entered an invalid hostname**.
4. Click **Add**.
   - If the hostname syntax is valid and NetWitness can reach the server, it displays **Validating**.
   - If the hostname syntax is valid and NetWitness cannot reach a server, the following is displayed, where ***hostname*** is the hostname that you attempted to add: **The NTP server *hostname* is unreachable. Please verify the address or check your firewall settings.**
5. Click **Apply**.
   A dialog displays notification that the settings have been saved and requests confirmation that you want to apply the settings now.
6. Click **Yes**.
   The NTP server specified now ensures that your host machine clocks are synchronized. If you decide to configure multiple NTP servers and a server is down, NetWitness will fail over to next server configured.

For details of the parameters and descriptions, see NTP Settings Panel.

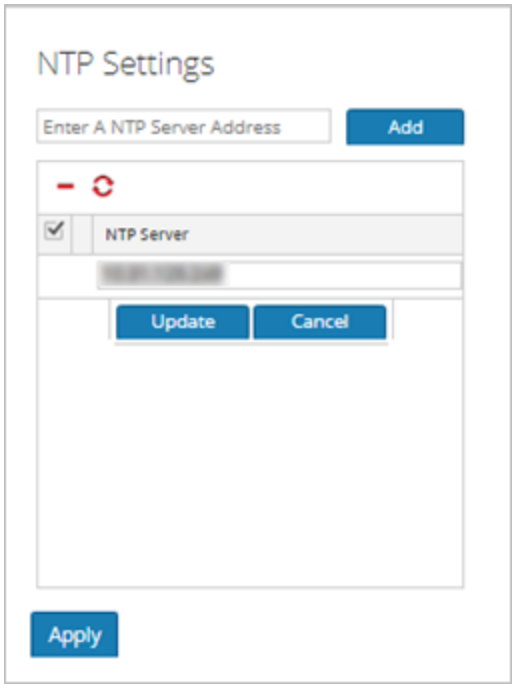
Modify an NTP Server

To modify an existing NTP server:

1. Go to  **(Admin) > System**.
2. In the options panel, select **NTP Settings**.
   The NTP Setting panel is displayed.



3. Double-click the **NTP Server** hostname that you want to modify.
   The NTP Server textbox becomes editable and the Update and Cancel buttons are displayed.



4. Edit the hostname, click **Update**, and click **Apply.** (click **Cancel** before you click **Apply** to cancel the edit.)
   NetWitness changes the hostname according to your edits.

Configure Proxy for NetWitness Platform

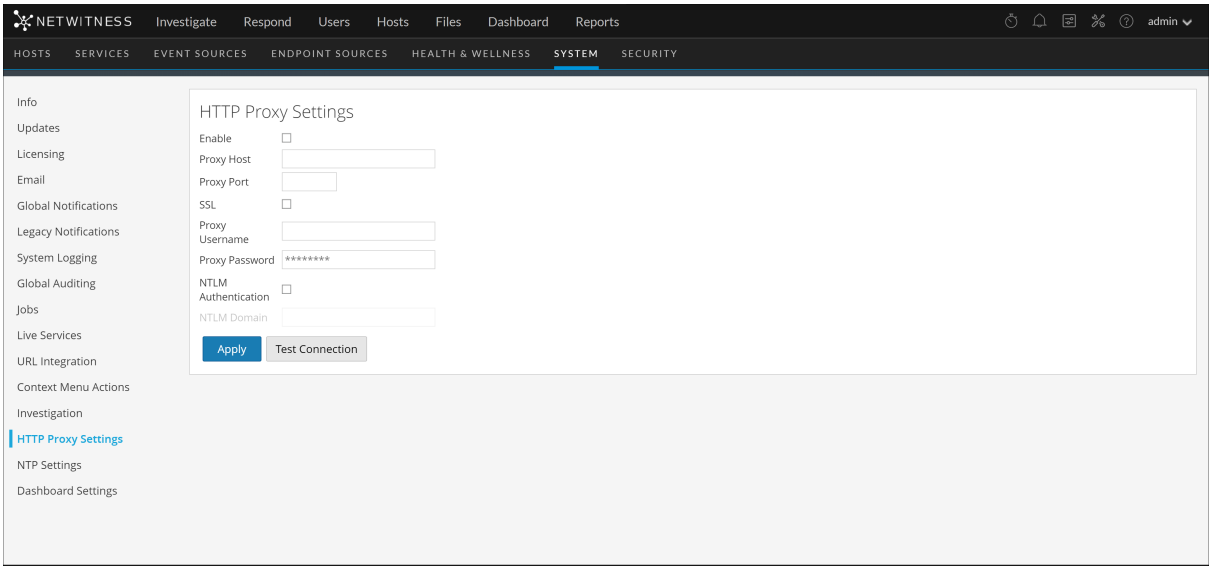
This topic provides a procedure for setting up a proxy that is used across NetWitness modules and services.

Note: Proxy support is only for HTTP and HTTPS proxies and not SOCKS5.

You can configure a proxy that is used across NetWitness modules and services in the System View > Advance Configuration panel. The Proxy Settings in the Advanced Configuration panel set up a proxy to be used wherever a proxy is needed in NetWitness. These settings override any proxy settings configured for an individual service or module, such as Malware Analysis or Live.

To configure a proxy for use across NetWitness modules:

1. Go to  **(Admin) > System**.
2. In the options panel, select **HTTP Proxy Settings**. The HTTP Proxy Settings panel is displayed.



3. Click the **Enable** checkbox.

   The fields where you configure the proxy settings are activated.

4. Type the hostname for the proxy server and the port used for communications on the proxy server.
5. (Optional) Type the username and password that serve as credentials to access the proxy server if authentication is required.
6. (Optional) Enable **Use NTLM Authentication** and type the NTLM domain name.
7. (Optional) Enable **Use SSL** if communications use Secure Socket Layer. If you enable **Use SSL**, ensure that you import the required certificates for the services to retrieve information.

   You need to import certificates and add the certificate in the head node for all the specific services that communicate externally over a proxy. If there are any other service that communicates with external resources over the internet, ensure that you add the certificates. For more on how to add the certificates, see Import Certificates for HTTPS Service

8. The proxy is immediately available for use throughout NetWitness modules and services, for example, Live and Malware Analysis.
9. To save and apply the configuration, click **Apply**.

Import Certificates for HTTPS Service

Import certificates to communicate with the HTTPS services:

1. SSH to the NW node and copy the CA certificate located in the following directory:
   /etc/pki/ca-trust/source/
2. Execute the following command to update the certificates:
   `update-ca-trust`
3. Execute the following command to add the certificate to the java keystore:
   `keytool -list -keystore /etc/pki/java/cacerts -storepass changeit |& head`
4. Restart the service on the NW node.

Note: Perform the procedure for all the HTTPS servers.
Example: HTTPS proxy server and HTTPS feed server.

Troubleshoot System Configuration

The topics in this section provide troubleshooting information for administrators who are configuring settings that apply across the system in NetWitness.

- Troubleshoot Global Audit Logging
- Troubleshoot Issues identified in the NTP Settings Panel or Log Files Messages
- Troubleshoot Global Notifications

Troubleshoot Global Audit Logging

This topic provides information about possible issues that NetWitness users may encounter when implementing Global Audit Logging in NetWitness. Look for explanations and solutions in this topic.

After you configure Global Audit Logging, you should test your audit logs to ensure that they show the audit events as defined in your audit logging template. If you cannot view the audit logs on your third-party syslog server or Log Decoder, or the audit logs do not appear as expected, look at the basic troubleshooting suggestions below. If you are still having issues, you can look at the advanced troubleshooting suggestions.

Basic Troubleshooting

If you cannot view audit logs on a third-party syslog server or Log Decoder:

- Verify that RabbitMQ is up and running.
- Verify the syslog notification server configuration and make sure it is enabled.

  (This configuration is located at  (Admin) > System > Global Notifications. Do not select Legacy Notifications.)
- Check the Global Audit Logging configuration.

Configure Global Audit Logging and Verify Global Audit Logs provide instructions. If you are sending audit logs to a Log Decoder:

- Ensure that the Log Decoder is aggregating on the Concentrator on the same host:

   (Admin) > Services > (Select Concentrator) >  > View > Config.
- Verify that the latest CEF parser is deployed and enabled.
- Check the audit logging notification template. You must use a CEF template and all logs feeding into the Log Decoder must use a CEF template.

If you are sending audit logs to a third-party syslog server, Ensure that the destination port configured for the third-party syslog server is not blocked by a firewall.

Advanced Troubleshooting

In order to use Global Audit Logging on your network, RabbitMQ must be functioning.

For centralized audit logging, each of the NetWitness services writes audit logs to rsyslog listening on port 50514 using UDP on the local host. The rsyslog plugin provided in the audit logging package adds additional information and uploads these logs to RabbitMQ. Logstash running on the NetWitness Server host aggregates audit logs from all of the NetWitness services, coverts them to the required format, and sends them to a third-party syslog server or Log Decoder for investigation. You configure the format of the global audit logs and the destination used by Logstash through the NetWitness user interface.

Define a Global Audit Logging Configuration provides instructions.

Verify the Packages and Services on the Hosts

NetWitness Host

The following packages or services must be present on the NetWitness Server host:

- rsyslog-8.4.1
- rsa-audit-rt
- logstash-5.6.4
- rsa-audit-plugins
- rabbitmq server

Services on a Host other than the NetWitness Host

The following packages or services must be present on each of the NetWitness hosts other than the NetWitness Server host:

- rsyslog-8.4.1
- rsa-audit-rt
- rabbitmq server

Log Decoder

If you forward global audit logs to a Log Decoder, the following parser should be present and enabled:

- CEF

Possible Issues

What if I perform an action on a service but audit logs do not reach the configured third-party syslog server or Log Decoder?

The possible causes could be one or all of the following:

- A service is not logging to the local syslog server.
- Audit logs are not getting uploaded to RabbitMQ from the local syslog.
- Audit logs are not aggregated on the NetWitness Server host.
- Aggregated logs on the NetWitness Server host are not being forwarded to the configured third-party syslog server or Log Decoder.

- The Log Decoder is not configured to receive global audit logs in CEF format:
  - Log Decoder capture is not turned on
  - CEF Parser is not present
  - CEF Parser is not enabled

Possible Solutions

The following table provides possible solutions for the issues.

Issue	Possible Solutions

A service is not logging to the local syslog server.	• Ensure that rsyslog is up and running.
	You could use the following command:
	`service rsyslog status`
	• Ensure that rsyslog is listening on port 50514 using UDP.
	You could use the following command:
	`netstat -tulnp|grep rsyslog`
	• Ensure the application or component is sending audit logs to port 50514. Run the tcpdump utility on the local interface for port 50514.
	You could use the following command:
	`sudo tcpdump -i lo -A udp and port 50514`
	See "Solution Examples" below to view the command outputs.

Audit logs are not getting uploaded to RabbitMQ from the local syslog.	• Ensure that the rsyslog plugin is up and running.
	You could use the following command:
	`ps -ef|grep rsa_audit_onramp`
	• Ensure the RabbitMQ server is up and running.
	You could use the following command:
	`service rabbitmq-server status`
	See "Solution Examples" to view the command outputs.

Audit logs are not aggregated on the NetWitness Server host.	• Ensure Logstash is up and running.
	You could use the following commands:
	`ps -ef|grep logstash`
	`service logstash status`
	• Ensure the RabbitMQ server is up and running.
	You could use the following command:
	`service rabbitmq-server status`
	• Ensure the RabbitMQ server is listening on port 5672.

Issue	Possible Solutions

	You could use the following command:
	`netstat -tulnp|grep 5672`
	• Check for any errors generated at the Logstash level.
	You could use the following command for the location of the log files:
	`ls -l /var/log/logstash/logstash.*`
	See "Solution Examples" to view the command outputs.

Aggregated logs on the NetWitness Server host are not being forwarded to the configured third-party syslog server or Log Decoder.	• Ensure Logstash is up and running.
	You could use the following commands:
	`ps -ef|grep logstash`
	`service logstash status`
	• Check for any errors generated at the Logstash level.
	You could type the following command for the location of the log files:
	`ls -l /var/log/logstash/logstash*`
	See "Solution Examples" below to view the command outputs.
	• Ensure that the destination service is up and running.
	• Ensure that the destination service is listening on the correct port using the correct protocol.
	• Ensure that the configured port on the destination host is not blocked.

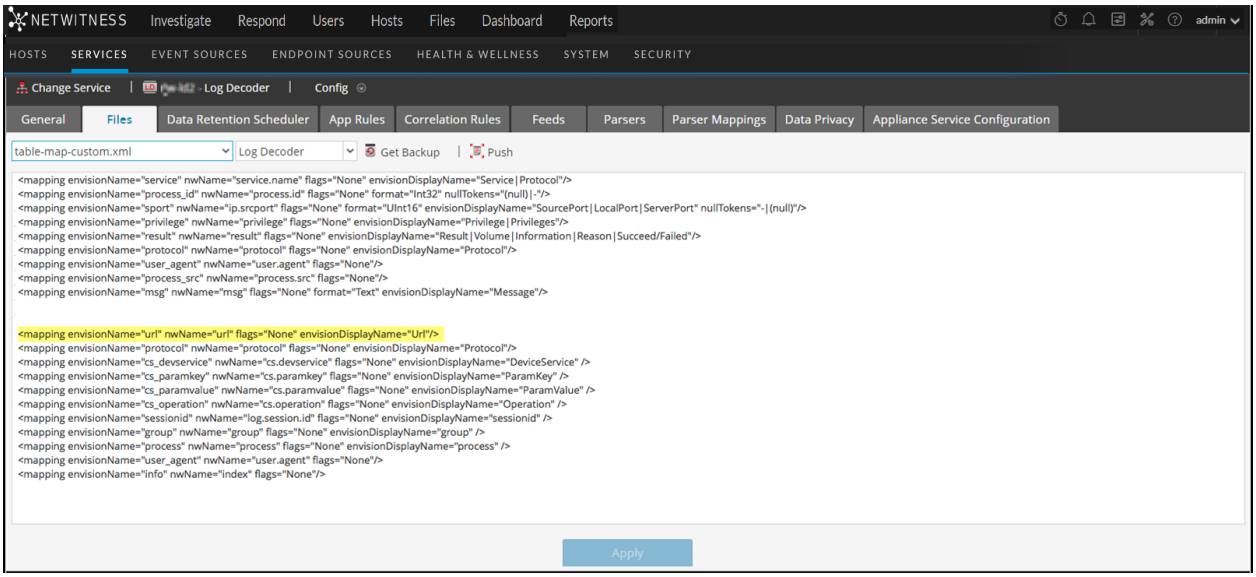
Audit logs forwarded from the Logstash lead to parse failure at the Log Decoder.	• Ensure that you are using an appropriate notification template. Audit Logs parsed by a Log Decoder must be in CEF format. The destination from which audit logs directly or indirectly make their way to the Log Decoder must also use a CEF Template. • The Notification Template must follow the CEF standard. Follow the steps in this guide to either use the default CEF template or create a custom CEF template following strict guidelines. Define a Template for Global Audit Logging provides additional information. • Verify the Logstash configuration.

Why can't we see the custom metadata in Investigation?

Usually, if a meta key is not visible in Investigation, it is not being indexed. If you need to use custom meta keys for Investigations and Reporting, ensure that the meta keys that you select are indexed in the **table-map-custom.xml** file on the Log Decoder. Follow the "Maintain the Table Map Files" procedure to modify the **table-map-custom.xml** file on the Log Decoder.

Ensure that the custom meta keys are also indexed in the **index-concentrator-custom.xml** on the Concentrator. "Edit a Service Index File" provides additional information.

The following figure shows an example **table-map-custom.xml** file in NetWitness Server ((Admin) > Services > (select the Log Decoder) > >View > Config) with a custom meta `url` example highlighted.

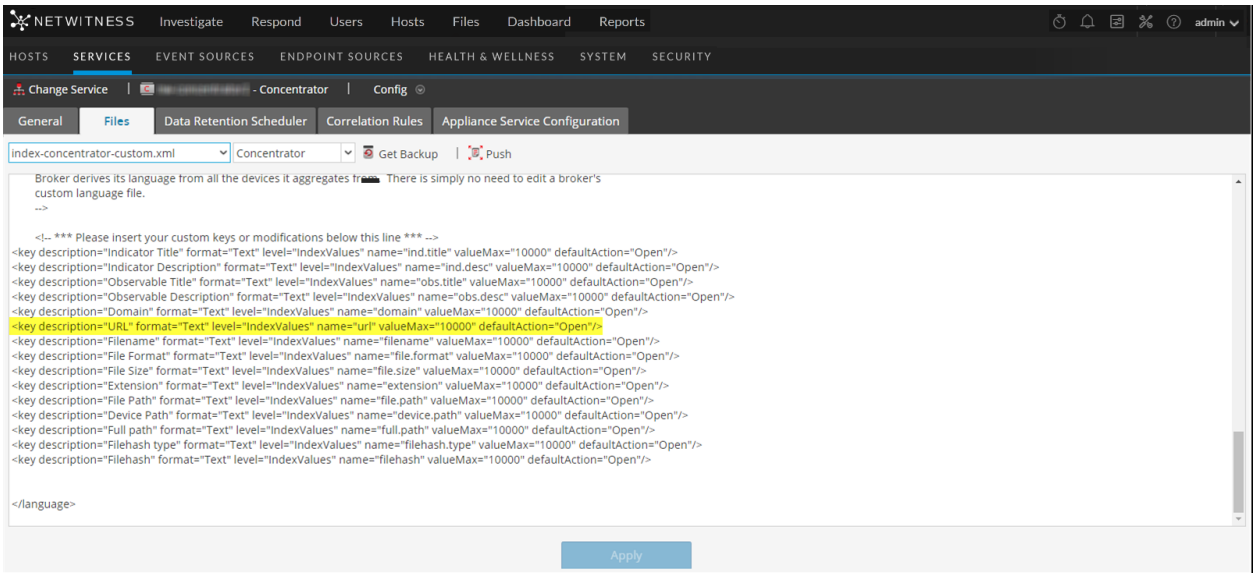


The `url` custom meta example is highlighted in the following code sample from the **table-map-custom.xml** file above:

```
<mapping envisionName="url" nwName="url" flags="None"
envisionDisplayName="Url"/>
<mapping envisionName="protocol" nwName="protocol" flags="None"
envisionDisplayName="Protocol"/><mapping envisionName="cs_devservice"
nwName="cs.devservice" flags="None" envisionDisplayName="DeviceService"
/><mapping envisionName="cs_paramkey" nwName="cs.paramkey" flags="None"
envisionDisplayName="ParamKey" /><mapping envisionName="cs_paramvalue"
nwName="cs.paramvalue" flags="None" envisionDisplayName="ParamValue"
/><mapping envisionName="cs_operation" nwName="cs.operation" flags="None"
envisionDisplayName="Operation" /><mapping envisionName="sessionid"
nwName="log.session.id" flags="None" envisionDisplayName="sessionid"
/><mapping envisionName="group" nwName="group" flags="None"
```

```
envisionDisplayName="group" /><mapping envisionName="process" nwName="process"
flags="None" envisionDisplayName="process" /><mapping envisionName="user_
agent" nwName="user.agent" flags="None"/><mapping envisionName="info"
nwName="index" flags="None"/>
```

The following figure shows an example **index-concentrator-custom.xml** file in NetWitness Server ((Admin) > Services > (select the Concentrator) > > View > Config) with a custom meta `url` example highlighted.



The `url` custom meta example is highlighted in the following code sample from the **index-concentrator-custom.xml** file above:

```
<key description="Severity" level="IndexValues" name="severity"
valueMax="10000" format="Text"/><key description="Result" level="IndexValues"
name="result" format="Text"/><key level="IndexValues"  name="ip.srcport"
format="UInt16" description="SourcePort"/><key description="Process"
level="IndexValues" name="process" format="Text"/><key description="Process
ID" level="IndexValues" name="process_id" format="Text"/><key
description="Protocol" level="IndexValues" name="protocol" format="Text"/><key
description="UserAgent" level="IndexValues" name="user_agent"
format="Text"/><key description="DestinationAddress" level="IndexValues"
name="ip.dst" format="IPv4"/><key description="SourceProcessName"
level="IndexValues" name="process.src" format="Text"/><key
description="Username" level="IndexValues" name="username"
format="Text"/><key description="Info" level="IndexValues" name="index"
```

```
format="Text"/><key description="customdevservice" level="IndexValues"
name="cs.devservice"  format="Text"/>
<key description="url" level="IndexValues" name="url"  format="Text"/>
<key description="Custom Key" level="IndexValues" name="cs.paramkey"
format="Text"/><key description="Custom Value" level="IndexValues"
name="cs.paramvalue"  format="Text"/><key description="Operation"
level="IndexValues" name="cs.operation"  format="Text"/><key description="CS
Device Service" level="IndexValues" name="cs.device" format="Text"
valueMax="10000" defaultAction="Closed"/>
```

Solution Examples

The following possible solution examples show the outputs of the example commands. See the above table for the complete listing of possible solutions.

Ensure that rsyslog is up and running

You can use the following command:

```
service rsyslog status
```

```
[root@NWAPPLIANCE22574 ~]# service rsyslog status
rsyslogd (pid  1293) is running...
[root@NWAPPLIANCE22574 ~]#
```

Ensure that rsyslog is listening on port 50514 using UDP

You can use the following command:

```
netstat -tulnp|grep rsyslog
```

```
[root@NWAPPLIANCE22574 ~]# netstat -tulnp|grep rsyslog
udp        0      0 127.0.0.1:50514           0.0.0.0:*                           1293/rsyslogd
[root@NWAPPLIANCE22574 ~]#
```

Ensure that the application or component is sending audit logs to port 50514

The following figure shows the output of running the tcpdump utility on the local interface for port 50514.

You can use the following command:

```
sudo tcpdump -i lo -A udp and port 50514
```

```
[root@NWAPPLIANCE22574 ~]# sudo tcpdump -i lo -A udp and port 50514
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on lo, link-type EN10MB (Ethernet), capture size 65535 bytes
08:54:46.536420 IP NWAPPLIANCE22574.34822 > NWAPPLIANCE22574.50514: UDP, length 593
E..m..@.@.:~..........R.Y.m<38>2015-04-24T08:54:46Z NWAPPLIANCE22574 SA_SERVER {"category":"DATA_ACCESS","deviceProduct":"Security Analytics","deviceService":"SA_SERVER","devic
eVendor":"RSA","deviceVersion":"10.5.0.0","identity":"Unknown identity","operation":"/poll/cda459a3-4e9d-ce1f-20f2-8cb1e31ef198","outcome":"Success","parameters":"{referrer=http
s://10.31.252.196/unified/dashboard/1, method=DELETE, userAgent=Mozilla/5.0 (Windows NT 6.1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/42.0.2311.90 Safari/537.36, queryStrin
g=ctoken=b33b67c5-6ae9-47b4-b435-560ecd38b760, remoteAddress=10.30.97.119}","severity":6}

08:54:46.615748 IP NWAPPLIANCE22574.34822 > NWAPPLIANCE22574.50514: UDP, length 365
E.....@.@.;b..........R.u..<38>2015-04-24T08:54:46Z NWAPPLIANCE22574 SA_SERVER {"category":"DATA_ACCESS","deviceProduct":"Security Analytics","deviceService":"SA_SERVER","devic
eVendor":"RSA","deviceVersion":"10.5.0.0","identity":"admin","key":"user.general.contextmenu","operation":"Users.preferences.","severity":6,"userRole":"Administrators+Administra
tors+PRIVILEGED_CONNECTION_AUTHORITY"}

08:54:46.618691 IP NWAPPLIANCE22574.34822 > NWAPPLIANCE22574.50514: UDP, length 367
E.....@.@.;`..........R.w..<38>2015-04-24T08:54:46Z NWAPPLIANCE22574 SA_SERVER {"category":"DATA_ACCESS","deviceProduct":"Security Analytics","deviceService":"SA_SERVER","devic
eVendor":"RSA","deviceVersion":"10.5.0.0","identity":"admin","key":"user.notifications.enabled","operation":"Users.preferences.","severity":6,"userRole":"Administrators+Administ
rators+PRIVILEGED_CONNECTION_AUTHORITY"}

08:54:46.623411 IP NWAPPLIANCE22574.34822 > NWAPPLIANCE22574.50514: UDP, length 369
E.....@.@.;^..........R.y..<38>2015-04-24T08:54:46Z NWAPPLIANCE22574 SA_SERVER {"category":"DATA_ACCESS","deviceProduct":"Security Analytics","deviceService":"SA_SERVER","devic
eVendor":"RSA","deviceVersion":"10.5.0.0","identity":"admin","key":"user.browser_timezone_zoneId","operation":"Users.preferences.","severity":6,"userRole":"Administrators+Admini
strators+PRIVILEGED_CONNECTION_AUTHORITY"}

08:54:46.626311 IP NWAPPLIANCE22574.34822 > NWAPPLIANCE22574.50514: UDP, length 369
E.....@.@.;^..........R.y..<38>2015-04-24T08:54:46Z NWAPPLIANCE22574 SA_SERVER {"category":"DATA_ACCESS","deviceProduct":"Security Analytics","deviceService":"SA_SERVER","devic
eVendor":"RSA","deviceVersion":"10.5.0.0","identity":"admin","key":"user.browser_timezone_zoneId","operation":"Users.preferences.","severity":6,"userRole":"Administrators+Admini
strators+PRIVILEGED_CONNECTION_AUTHORITY"}
```

Ensure that the rsyslog plugin is up and running

You can use the following command:

```
ps -ef|grep rsa_audit_onramp
```

```
[root@NWAPPLIANCE22574 ~]# ps -ef|grep rsa_audit_onramp
root      1636  1293  0 06:05 ?        00:00:03 /usr/sbin/rsa_audit_onramp --node_id=96b08193-a9d0-4a79-b362-87b56851f411
root     22248  6921  0 09:09 pts/0    00:00:00 grep rsa_audit_onramp
[root@NWAPPLIANCE22574 ~]#
```

Ensure the RabbitMQ server is up and running

You can use the following command:

```
service rabbitmq-server status
```

```
[root@NWAPPLIANCE22574 ~]# service rabbitmq-server status
Status of node sa@localhost ...
[{pid,1862},
 {running_applications,
     [{rabbitmq_federation_management,"RabbitMQ Federation Management",
         "3.4.2"},
      {rabbitmq_management,"RabbitMQ Management Console","3.4.2"},
      {rabbitmq_web_dispatch,"RabbitMQ Web Dispatcher","3.4.2"},
      {webmachine,"webmachine","1.10.3-rmq3.4.2-gite9359c7"},
      {mochiweb,"MochiMedia Web Server","2.7.0-rmq3.4.2-git680dba8"},
      {rabbitmq_federation,"RabbitMQ Federation","3.4.2"},
      {rabbitmq_stomp,"Embedded Rabbit Stomp Adapter","3.4.2"},
      {rabbitmq_management_agent,"RabbitMQ Management Agent","3.4.2"},
      {rabbit,"RabbitMQ","3.4.2"},
      {ssl,"Erlang/OTP SSL application","5.3.2"},
      {public_key,"Public key infrastructure","0.21"},
      {crypto,"CRYPTO version 2","3.2"},
      {asn1,"The Erlang ASN1 compiler version 2.0.4","2.0.4"},
      {os_mon,"CPO  CXC 138 46","2.2.14"},
      {inets,"INETS  CXC 138 49","5.9.7"},
      {mnesia,"MNESIA  CXC 138 12","4.11"},
      {amqp_client,"RabbitMQ AMQP Client","3.4.2"},
      {rabbitmq_auth_mechanism_ssl,
          "RabbitMQ SSL authentication (SASL EXTERNAL)","3.4.2"},
      {xmerl,"XML parser","1.3.5"},
      {sasl,"SASL  CXC 138 11","2.3.4"},
      {stdlib,"ERTS  CXC 138 10","1.19.4"},
      {kernel,"ERTS  CXC 138 10","2.16.4"}]},
 {os,{unix,linux}},
 {erlang_version,
     "Erlang R16B03 (erts-5.10.4) [source] [64-bit] [smp:2:2] [async-threads:30] [kernel-poll:true]\n"},
 {memory,
```

Ensure logstash is up and running

You can use the following commands:

```
ps -ef | grep logstash
service logstash status
```

```
[root@NWAPPLIANCE22574 ~]# ps -ef|grep logstash
logstash  1583     1  0 06:05 ?        00:01:09 /usr/bin/java -Djava.io.tmpdir=/var/lib/logstash -Xmx500m -XX:+UseParNewGC -XX:+UseConcMarkSweepGC -Djava.awt.headless=true -XX:C
MSInitiatingOccupancyFraction=75 -XX:+UseCMSInitiatingOccupancyOnly -jar /opt/logstash/vendor/jar/jruby-complete-1.7.11.jar -I/opt/logstash/lib /opt/logstash/lib/logstash/runner
.rb agent --pluginpath /opt/logstash -f /etc/logstash/conf.d -l /var/log/logstash/logstash.log
root      8509  6921  0 09:31 pts/0    00:00:00 grep logstash
[root@NWAPPLIANCE22574 ~]# service logstash status
logstash is running
[root@NWAPPLIANCE22574 ~]#
```

Ensure the RabbitMQ server is listening on port 5672

For example, type the following command:

```
netstat -tulnp | grep 5672
```

```
[root@NWAPPLIANCE22574 ~]# netstat -tulnp|grep 5672
tcp        0      0 127.0.0.1:5672              0.0.0.0:*                   LISTEN      1862/beam.smp
tcp        0      0 0.0.0.0:25672               0.0.0.0:*                   LISTEN      1862/beam.smp
[root@NWAPPLIANCE22574 ~]#
```

Check for any errors generated at the Logstash level

You can type the following command for the location of the log files:

```
ls -l /var/log/logstash/logstash.*
```

```
[root@NWAPPLIANCE22574 ~]# ls -l /var/log/logstash/logstash.*
-rw-r--r--. 1 root     root        0 Apr 24 06:05 /var/log/logstash/logstash.err
-rw-r--r--. 1 logstash logstash 1043 Apr 24 06:04 /var/log/logstash/logstash.log
-rw-r--r--. 1 root     root       57 Apr 24 06:12 /var/log/logstash/logstash.stdout
[root@NWAPPLIANCE22574 ~]#
```

See the Possible Solutions table above for the complete listing of issues and possible solutions.

Troubleshoot Issues identified in the NTP Settings Panel or Log Files Messages

This section provides troubleshooting information for issues identified by messages NetWitness displays in the NTP Settings panel and log files.

Issue	Possible Solutions
Message	User Interface: **Unexpected error occurred. First check the logs then contact Customer Care to resolve error.** **System Log:** **Timestamp Level Message** *yyyy-dd-mmThh:mm:ss:ms* ERROR com.rsa.smc.sa.adm.exception.MCOAgent Exception: No request sent, we did not discover any nodes
Possible Cause	Low level NetWitness configuration is in error or supporting service is not running.
Solution	Contact Customer Care.
Message	User Interface: **Specified an invalid Hostname syntax**.
Possible Cause	Tried to enter NTP server hostname that does not confirm to IP address or FQDN syntax.
Solution	Reenter hostname in using correct syntax.
Message	User Interface: **Specified NTP server that already exists.**
Possible Cause	Tried to enter NTP server hostname that is already defined in NetWitness.
Solution	Enter hostname for an NTP server not configured in NetWitness.
Message	User Interface: **Cannot reach NTP server *hostname***. Please verify the server address and your firewall settings.
Possible	The server address or firewall settings may be in error.

Issue	Possible Solutions
Cause	
Solution	Verify the server address and your firewall settings and correct them if required.

Troubleshoot Global Notifications

This topic provides information about possible issues that NetWitness users may encounter when implementing Global Notifications in NetWitness.

Issue	Possible Solution
We are not receiving notifications that were configured for a service, but the service log file does not show any errors.	For any notification-related troubleshooting, check the `integration-server log` file in addition to the log file of the service creating the notification. For example, when troubleshooting ESA rule notifications, check both the ESA Correlation service log files (`/var/log/netwitness/correlation-server/correlation-server.log`) AND the Integration-Server log files on the NetWitness Server (`/var/log/netwitness/integration-server/integration-server.log`). If the `integration-server.log` file shows a failure when the Integration-Server attempts to send a notification to the notification server, you should check the notification server configuration in the Global Notifications settings (**(Admin) > System > Global Notifications > Servers** tab).

References

This topic provides reference materials that describe the user interface for configuring system settings in NetWitness and define parameters. Administrators use options in the Administration System view to configure system settings. Each panel is described in a separate topic.

- Global Audit Logging Configurations Panel
- Global Notifications Panel
  - Define Notification Server Dialogs
  - Define Notification Output Dialogs
  - Define Notification Template Dialog
  - Output Tab
  - Servers Tab
  - Templates Tab
- HTTP Proxy Settings Panel
- Email Configuration Panel
- Investigation Configuration Panel
- Live Services Configuration Panel
- NTP Settings Panel
- Context Menu Actions Panel
- Legacy Notifications Configuration Panel

Global Audit Logging Configurations Panel

In the **Global Audit Logging Configurations** panel ((Admin) > System > Global Auditing), you configure global audit logging by adding configurations that define how global audit logs are forwarded to external syslog systems. Global audit logs are forwarded to the selected Notification Server in your global audit logging configuration using the selected Notification Template.

Global Audit Logging provides auditors with consolidated visibility into user activities within NetWitness in real-time from one centralized location.

Workflow

This workflow shows the necessary procedures to configure and verify Global Audit Logging.



Before you can define a Global Audit Logging configuration, you need to create a Syslog Notification Server on the Global Notifications > Server tab. The Syslog Notification Server is the destination that receives the global audit logs. Next, you need to select or define an Audit Logging template on the Global Notifications > Templates tab. The Audit Logging template defines the format and message fields of the audit logs sent to the Log Decoder or third-party syslog server. If you are consuming with a Log Decoder, deploy the Common Event Format parser to your Log Decoder from Live.

Note: You do not need to configure the Global Notifications > Output tab for Global Audit Logging.

After you add a Global Audit Logging configuration here, audit logs are forwarded to the selected Notification Server in the configuration. Verify your audit logs to ensure that they show the audit events as defined in your audit logging template.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Create a Syslog Notification Server.	Configure a Destination to Receive Global Audit Logs
Administrator	Choose an Audit Logging template.	Define a Template for Global Audit Logging

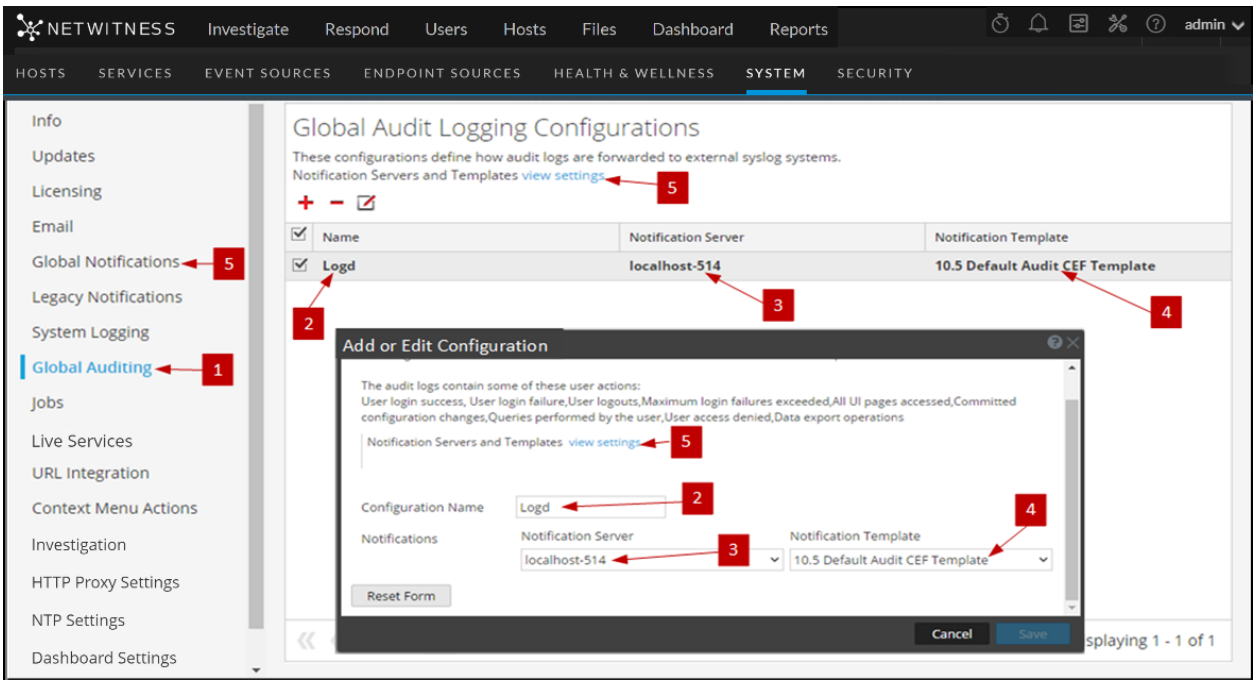
Role	I want to ...	Show me how
Administrator	Configure Global Audit Logging	Define a Global Audit Logging Configuration For the complete procedure, see "Global Audit Logging - High-Level Procedure" in Configure Global Audit Logging.
Administrator	Verify Global Audit logs	Verify Global Audit Logs

Related Topics

- Troubleshoot Global Audit Logging
- Add New Configuration Dialog
- Supported CEF Meta Keys
- Supported Global Audit Logging Meta Key Variables
- Global Audit Logging Operation Reference
- Local Audit Log Locations

Quick Look

The following example illustrates a Global Audit Logging configuration. The configuration defines how NetWitness forwards global audit logs to external syslog systems.



1 Displays the Global Audit Logging Configurations panel.
2 Name that identifies the Global Audit Logging configuration.
3 Notification Server assigned to the Global Audit Logging configuration.
4 Notification Template assigned to the Global Audit Logging configuration.

5 Displays the Global Notifications panel where you set up Servers and Templates required to configure a Global Audit Logging configuration.

Toolbar

The following table describes the toolbar actions

Icon	Description
	Adds a global audit logging configuration.
	Deletes a global audit logging configuration. Deleting a global audit configuration does not delete the associated notification server and template. After you delete a global audit logging configuration, the forwarding of global audit logs specified in that configuration is discontinued.
	Edits a global audit logging configuration. You can change the destination of the global audit logs for your user audits by selecting a different Notification Server. You can also change the format and message fields of the global audit log entries by selecting a different Notification Template. You cannot change which NetWitness user actions are logged and sent in the global audit logs.

Configurations

The following table describes the listed configurations.

Title	Description
	To select an individual configuration, select the checkbox next to the configuration. To select all configurations, select the checkbox in the title bar of the table.
Name	Displays the name of the global auditing configuration. For example, you can name the configurations based on the destination of the global audit logs, such as HQ SA and My Syslog Server.
Notification Server	Displays the Syslog Notification Server selected as the destination for the global audit logs. If you want to forward global audit logs to a Log Decoder, create a Syslog type of Notification Server. Configure a Destination to Receive Global Audit Logs provides instructions on how to create a Syslog Notification Server for global audit logging.
Notification Template	Displays the Audit Logging Notification Template selected for the configuration. It defines the format and message fields of the audit log entries. For Log Decoders, use the **Default Audit CEF Template**. You can add or remove fields from the Common Event Format (CEF) template if you have specific requirements. Define a Template for Global Audit Logging provides instructions and Supported CEF Meta Keys describes the available CEF meta keys. For, third-party syslog servers, you can use a default audit logging template or define your own format (CEF or non-CEF). Define a Template for Global Audit Logging provides instructions and Supported Global Audit Logging Meta Key Variables describes the available meta key variables.

Add New Configuration Dialog

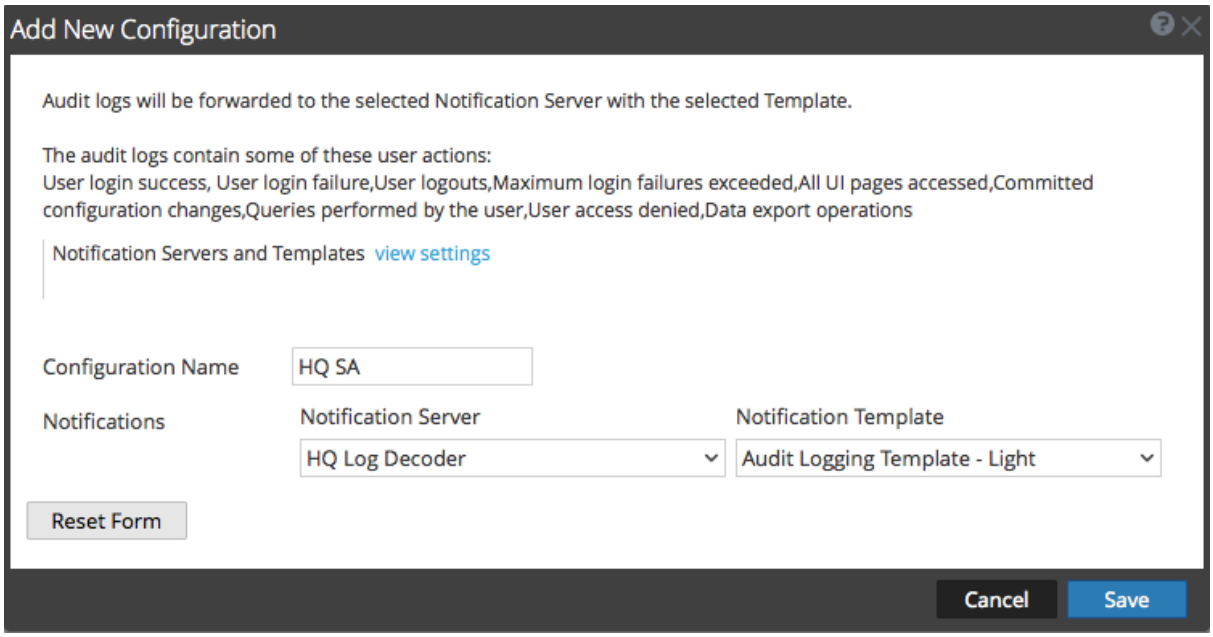
In the NetWitness Platform Administration System view Global Audit Logging Configurations panel, you can create multiple global audit logging configurations. These configurations are used to forward global audit logs to a central location to perform user audits.

Procedures related to global audit logging are described in Configure Global Audit Logging.

To access the **Add New Configuration** dialog:

1. Go to select **(Admin) > System**.
2. In the options panel, select **Global Auditing**.
3. In the **Global Audit Logging Configurations** panel, click .

   The **Add New Configuration** dialog is displayed.



   The Notifications section enables you to select a syslog notification server for the global audit logging configuration and a template to use for the global audit logs. The template defines the details of the global audit log entries.

Features

The following table describes the features in the Add New Configuration and Edit Configuration dialogs.

Feature	Description
Notifications Servers and Templates **view settings** link	Takes you to the Global Notifications panel where you can view or configure the notification server and template settings. A syslog notification server and an audit logging template are required before you can create a global audit configuration.
Configuration Name	Specifies the unique name used to identify the global audit logging configuration.
Notification Server	Specifies the syslog notification server to send the selected audit log information. Configure a Destination to Receive Global Audit Logs provides instructions on how to create a Syslog Notification Server for global audit logging.
Notification Template	Specifies the template to use for the global audit logging configuration. The template should be an Audit Logging template. For Log Decoders, use the **Default Audit CEF Template**. You can add or remove fields from the Common Event Format (CEF) template if you have specific requirements. Define a Template for Global Audit Logging provides instructions. For third-party syslog servers, you can use a default audit logging template or define your own format (CEF or non-CEF). Define a Template for Global Audit Logging provides instructions and Supported Global Audit Logging Meta Key Variables describes the available variables.
Reset Form button	Clears the configuration settings in the dialog.

User Actions Logged

The following table provides examples of some of the user actions logged from NetWitness. These actions are the minimum user actions logged when applicable.

User Action	Example
User login success	A user logs on with valid credentials.
User login failure	A user tries to log on using invalid credentials.
User logouts	A user logs out from NetWitness (Administration > Sign Out) or a user logs out due to a session timeout.
Max login failures exceeded	A user tries to log on using invalid credentials five times. Five (5) is the number of Max Login Failures defined in Administration Security view > Settings tab (Administration > Security > Settings tab).
All UI pages accessed	When a user accesses the Reporting module (Administration > Reports), it logs as `[REP] Reports`. When a user accesses the Administration System view (Administration > System), it logs as `[ADM] System`.

User Action	Example
Committed configuration changes	A user changes his or her password and or any security setting (Administration > Security > Settings tab).
Queries performed by the user	A user performs an investigation query.
User access denied	A user tries to access a module and does not have permissions to access it.
Data export operations	A user exports data from the Events view (Investigation > Events > Actions > Export).

The following table shows examples of internal audit logs logged from NetWitness

User Actions	Audit Log Examples
User Login success	{"type":"fileclone","hostname":"UpdateStackAdminServer", "timegenerated":"2019-05-23T13:55:42.764124+00:00","syslogtag": "ADMIN-SERVER","@version":"1","fromhost-ip":"110.10.10.1”, "deviceVendor":"RSA","deviceService":"admin-server”, "deviceVersion":"11.3.1.0” “uri":"/oauth/token", "referrer":"https://10.111.201.10/login","success":"true", "identity":"AdminNorm","action":"Logon-Web","deviceServiceId": "247cedcb-cXXX-4XXX-8XXX-5XXXXa","deviceProduct":"NetWitness", "category":"Security","operation":"Logon-Web",”outcome": "success","remoteAddress":"101.181.15.10","message":null, "logTime":"2019-05-23T13:55:42.769Z","@timestamp":"2019-05 -23T13:55:42.769Z","timereported":"2019-05-23T13:55:42+00:00", "node_id":"e0XXX8-4XXX-4XXX-8XXXX-6d4b8XXXX09"}
User Login Failure	{"type":"fileclone","hostname":"UpdateStackAdminServer", "timegenerated":"2019-05-23T13:42:38.485701+00:00","syslogtag": "ADMIN-SERVER","@version":"1","fromhost-ip":"111.1.10.11", "deviceVendor":"RSA","deviceService":"adminserver","device Version":"11.3.1.0","uri":"/oauth/token","referrer":"https: //10.111.201.10/login","success":"false","identity":"AdminNor m", "reasonForFailure":"Bad Credentials","action":"Logon-Web", "deviceServiceId":"2XXXX-cXXX-4XXX-8XXX-5feXXXX2a","device Product":"NetWitness","category":"Security","operation": "Logon-Web","outcome":"failed","remoteAddress":"101.181.15.10", "message":null,"logTime":"2019-05- 23T13:42:38.494Z","@timestamp": "2019-05-23T13:42:38.494Z","timereported":"2019-05-23T13: 42:38+00:00","node_id":"e0XXXX-4XXX-4XXX-8XXX-6dXXXXX809"}

User Actions	Audit Log Examples
User Logouts	{"type":"fileclone","hostname":"UpdateStackAdminServer", "timegenerated":"2019-06-06T13:43:57.112760+00:00","syslogtag": "SOURCE-SERVER","@version":"1","fromhost-ip":"107.0.110.1,"device Vendor":"RSA","deviceService":"source-server","deviceVersion": "11.3.1.0","size":"0","success":"true","identity":"system", "action":"sourceCountUpdate","deviceServiceId":"c872d520-b06b-46cb-b5c1-8e240b105020","deviceProduct":"NetWitness","category": "SystemOperation","operation":"sourceCountUpdate","parameters": " {\"size\":\"0\"}","outcome":"success"},"message":null,"logTime" :"2019-06-06T13:43:57.117Z","@timestamp":"2019-06-06T13:43:57 .117Z","timereported":"2019-06-06T13:43:57+00:00","node_id": "e07b16f8-4xxx-4xx1-895b-6xxxxx09"}
All UI pages accessed	{"type":"fileclone","hostname":"UpdateStackAdminServer", "timegenerated":"2019-05-23T14:03:16.094611+00:00","syslogtag": "SA_SERVER","@version":"1","fromhost-ip":"117.10.10.11", "json": {"severity":"6","deviceVendor":"RSA","identity":"AdminNorm", "deviceService":"SA_ SERVER","deviceProduct":"NetWitness","device Version":"11.3.1.0","category":"DATA_ACCESS","userRole": "Administrators","operation":"HttpRequest",outcome":"Success"}, "message":null,"logTime":"2019-05-23T14:03:16.115Z","@timestamp": "2019-05-23T14:03:16.115Z","timereported":"2019-05-23T14:03:16Z", "node_id":"e0XXXX-4XXX-4XXX-8XXX-6d5XXXX09"} {"type":"fileclone","hostname":"UpdateStackAdminServer", "timegenerated":"2019-05-23T14:04:17.305585+00:00","syslogtag": "SA_SERVER","@version":"1","fromhost-ip":"117.10.10.1","json": {"severity":"6","deviceVendor":"RSA","identity":"AdminNorm", "deviceService":"SA_ SERVER","deviceProduct":"NetWitness","device Version":"11.3.1.0","category":"SYSTEM","userRole": "Administrators","operation":"Page Accessed","key":"[ADM] Hosts", "outcome":"Success"},"message":null,"logTime":"2019-05-23T14: 04:17.309Z","@timestamp":"2019-05-23T14:04:17.309Z","time reported":"2019-05-23T14:04:17Z","node_id":"e07XXXX-4XXX-4XXX-8XXX-6d55XXXXX09"}

User Actions	Audit Log Examples
Committed configuration changes	{"type":"fileclone","hostname":"UpdateStackAdminServer", "timegenerated":"2019-05-23T14:09:09.741982+00:00","syslogtag": "SA_SERVER","@version":"1","fromhost-ip":"117.101.0.11","json": {"severity":"6","deviceVendor":"RSA","deviceService":"SA_ SERVER" ,"deviceVersion":"11.3.1.0","identity":"AdminNorm","device Product":"NetWitness","category":"CONFIGURATION","userRole": "Administrators","operation":"Modified","parameters":"save", "value":"[10.10.201.10]","key":"ntp-servers","outcome":"Success"}, "message":null,"logTime":"2019-05-23T14:09:09.748Z","@timestamp": "2019-05-23T14:09:09.748Z","timereported":"2019-05-23T14:09:09Z", "node_id":"e07XXXX-4XXX-4XXX-8XXX-6dXXXXX9"}
Queries performed by the user	{"type":"fileclone","hostname":"UpdateStackAdminServer", "timegenerated":"2019-05-23T14:12:02.909062+00:00","syslogtag": "INVESTIGATE-SERVER","@version":"1","fromhost-ip":"117.10.10.11", "json":{"deviceVendor":"RSA","deviceService":"investigate-server","deviceVersion":"11.3.1.0*","success":"true","identity": "AdminNorm","action":"update","deviceServiceId":"f8XXXX5-bXXX-4XXX-bXXX-fXXXXX6","deviceProduct":"NetWitness","category": "Predicate","operation":"update","updated":"UserPredicateEntity (id=5cXXXXXXXXX9dd, userId=AdminNorm, predicateEntity=Predicate Entity(id=ff53, legacyId=null, query=user.all='solay', display Name=user.all='solay'), lastUsed=2019-05-23T14:12:02.897Z", "outcome":"success"},"message":null,"logTime":"2019-05-23T14:12:02.920Z","@timestamp":"2019-05-23T14:12:02.920Z","time reported":"2019-05-23T14:12:02+00:00","node_id": "e0XXXX-4XXX-4XXX-8XXX-6d5XXXXX09"}
Data export operations	2019-02-11 11:20:30,188 deviceVersion: "11.3.0.0" deviceService: "SA_SERVER" category: DATA_ACCESS operation: "submitExtractPcap" parameters: "deviceId=6 collectionName= predicateHandle=c6cf sessionIds=[9285468, 9286362, 9628535, 9629308, 10013047, 10017581, 10428756, 10439924, 10819088, 10820894, 11164416] startDate=2019-02-11T08:20:00.000Z endDate=2019-02-11T11:19:59.000Z id1=1 id2=287399592" outcome: "Success" identity: "admin" userRole: "Administrators"

The following table shows examples of Global Audit Logs using the default Common Event Format (CEF) template. After you create a Global Audit Logging configuration, audit logs automatically go to the external syslog system in the format specified in the selected Audit Logging template.

User Actions	CEF Examples
User Login Success	May 23 2019 13:52:39 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|Security|Logon-Web|6|rt=May 23 2019 13:52:39 scope=scope suser=AdminNorm sourceServiceName=admin-server deviceExternalId=eXXXX-4XXX-4XXX-8XXX-6dXXXXX09 deviceProcessName=ADMIN-SERVER outcome=success remoteAddress=110.10.10.1 uri=/oauth/token referrerURL=https://10.111.201.10/login
User Login Failure	May 23 2019 13:42:38 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|Security|Logon-Web|6|rt=May 23 2019 13:42:38 scope=scope suser=AdminNorm sourceServiceName=admin-server deviceExternalId=eXXXX-4XXX-4XXX-8XXX-6XXXXXX09 deviceProcessName=ADMIN-SERVER outcome=failed remoteAddress=110.10.10.1 reasonForFailure=Bad credentials uri=/oauth/token referrerURL=https://10.111.201.10/login
User Logouts	Jun 06 2019 13:01:25 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|Security|Logoff|6|rt=Jun 06 2019 13:01:25 scope=scope suser=admin sourceServiceName=admin-server deviceExternalId=e07b16f8-4xxx-4xx1-895b-6dxxxxx809 deviceProcessName=ADMIN-SERVER outcome=success remoteAddress=101.101.007.101 reason=User Triggered referrerURL=https://10.111.117.115/respond/incidents uri=/oauth/logout action=Logoff,"uri":"/oauth/logout”
All UI pages accessed	May 23 2019 14:01:13 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|DATA_ACCESS|HttpRequest|6|rt=May 23 2019 14:01:13 scope=scope suser=AdminNorm userRole=Administrators sourceServiceName=SA_SERVER deviceExternalId=e0XXX8-4XXX-4XXX-8XXX-6XXXXXX09 deviceProcessName=SA_SERVER outcome=Success remoteAddress=110.11.10.1 uri=/admin/appliances referrerURL=https://10.111.201.10/admin/services May 23 2019 14:01:13 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|SYSTEM|Page Accessed|6|rt=May 23 2019 14:01:13 scope=scope key=[ADM] Hosts suser=AdminNorm userRole=Administrators sourceServiceName=SA_SERVER deviceExternalId=e0XXXX-4XXX-4XXX-8XXX-6d5XXXXX09 deviceProcessName=SA_SERVER outcome=Success
Committed configuration changes	May 23 2019 14:08:03 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|CONFIGURATION|Modified|6|rt=May 23 2019 14:08:03 scope=scope key=ntp-servers value={10.10.20.10\=true} suser=AdminNorm userRole=Administrators sourceServiceName=SA_SERVER deviceExternalId=e07XXX-4XXX1-4XXX-8XXX-6d5XXXXX809 deviceProcessName=SA_SERVER params=validate

User Actions	CEF Examples
Queries performed by the user	May 23 2019 14:12:32 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|Predicate|update|6|rt=May 23 2019 14:12:32 scope=scope suser=AdminNorm sourceServiceName=investigate-server deviceExternalId=e0XXXX-4XXX-4XXX-8XXX-6d5XXXXX09 deviceProcessName=INVESTIGATE-SERVER outcome=success “updated":"UserPredicateEntity(id\=5cXXXXXXdd, userId\=AdminNorm, predicateEntity\=PredicateEntity (id\=ff53, legacyId\=null, query\=user.all\='solay', displayName\=user.all\='solay'), lastUsed\=2019-05-23T14:12:32.406Z)"}
Data export operations	May 23 2019 14:17:05 updatestackadminserver CEF:0|RSA|NetWitness Audit|11.3.1.0|DATA_ ACCESS|submitExtractPcap|6|rt=May 23 2019 14:17:05 scope=scope suser=AdminNorm userRole=Administrators sourceServiceName=SA_SERVER deviceExternalId=e0XXXX8-4XXX-4XXX-8XXX-6d5XXXXX9 deviceProcessName=SA_SERVER outcome=Success params=deviceId\=17 collectionName\= predicateHandle\=8629 sessionIds\=null startDate\=2019-05-23T10:59:00.000Z endDate\=2019-05-23T13:58:59.999Z id1\=1 id2\=393378

The following table shows examples of global audit logs using the default human-readable format template on a third-party syslog server.

User Actions	Human-Readable Format Output
User Login Success	Jun 11 2019 05:02:07 UpdateStackAdminServer Jun 11 2019 05:02:07 BROKER [audit] Event Category: AUTHENTICATION Operation: login Outcome: success Description: null User: admin Role: admin.owner, aggregate,concentrator.manage,connections.manage,everyone,index. manage,logs.manage,sdk.content,sdk.manage,sdk.meta,sdk.packets, services.manage,storedproc.execute,storedproc.manage,sys.manage, users.manage params=null
User Login Failure	Jun 11 2019 05:22:11 updatestackadminserver Jun 11 2019 05:22:11 admin-server [audit] Event Category: Security Operation: Logon-Web Outcome: failed Description: null User: admin Role: null params= {"referrer":"https://10.101.101.101/login","method":"POST", "reasonForFailure":"Bad credentials","userAgent":"Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/75.0.3770.80 Safari/537.36","uri":"/oauth/token", "remoteAddress": "10.101.102.103"}

User Actions	Human-Readable Format Output
User Logouts	Jun 11 2019 02:06:24 updatestackadminserver Jun 11 2019 02:06:24 admin-server [audit] Event Category: Security Operation: Logoff Outcome: success Description: null User: admin Role: null params ={"reason":"User Triggered","referrer":"https://10.101.101.101/respond/incidents","method":"POST","userAgent":"Mozilla/5.0 (Windows NT 10.0; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/74.0.3729.169 Safari/537.36","uri":"/oauth/logout", "remoteAddress":"10.101.102.103"}
All UI pages accessed	Jun 11 2019 02:06:25 updatestackadminserver Jun 11 2019 02:06:25 SA_SERVER [audit] Event Category: DATA_ACCESS Operation: Http Request Outcome: Success Description: null User: Unknown identityRole: null params= {referrer\=https://10.101.101.101/login, method\=GET, X-Forwarded-For\=10.201.111.111, userAgent\=Mozilla/ 5.0 (Windows NT 10.0; WOW64)AppleWebKit/537.36 (KHTML, like Gecko) Chrome/74.0.3729.169 Safari/537.36,queryString\=, uri\=/display/ security/securitybanner/get, remoteAddress\=10.101.102.103}
Committed configuration changes	Jun 11 2019 02:08:13 updatestackadminserver Jun 11 2019 02:08:13 security-server [audit] Event Category: Authorization Operation: update Outcome: success Description: null User: admin Role: null params={"Role":"analystl1","Add.Permission":"[admin-server.process.manage, admin-server.configuration.manage, admin-server. health.read, admin-server.security.manage, admin-server.metrics. read, admin-server.security.read, admin-server.logs.manage]"}

User Actions	Human-Readable Format Output
Queries performed by the user	Jun 11 2019 02:12:57 UpdateStackConcentrator Jun 11 2019 02:12:57 CONCENTRATOR[audit] Event Category: DATA_ACCESS Operation: query Outcome: success Description:has finished query (channel 2085, queued 00:00:00, execute 00:00:00) User: adminRole: null params= queryPriority\=20 id1\=1 id2\=324751797 size\=0 flags\=0 threshold\=0 query\="select event.time , sessionid , alias.host, reference.id ,host.src , user.dst , event.type , result.code , event.source.id , host.dst ,service.name , logon.type , device .type , event.cat.name , ec.activity , ec.outcome,analysis. service , event.desc , action , user.src , result where ((reference.id \='4624','4625','4769','4648') || (device.type \= 'rsaacesrv' && ec.activity \= 'Logon') || ((action \= '/usr/sbin/sshd' || action\='/usr/bin/login') && device.type \= 'rhlinux')) and event.time \= 1539444840-1539444899 "
Data export operations	Jun 11 2019 02:14:04 updatestackadminserver Jun 11 2019 02:14:04 SA_SERVER [audit] Event Category: DATA_ACCESS Operation: submit ExtractLogs Outcome: Success Description: null User: admin Role: Administrators params=deviceId\=17 collectionName\=predicate Handle\= sessionIds\=null exportFormat\=RAWLOGS startDate\=2019- 06-10T23:14:00.000Z endDate\=2019-06-11T02:13:59.999Z id1\=1 id2\ =1694664

For lists of message type being logged by the various NetWitness components, see Global Audit Logging Operation Reference.

Supported CEF Meta Keys

This topic describes the Common Event Format (CEF) meta keys that NetWitness global audit logging supports.

Global audit logging templates that you define for a Log Decoder use Common Event Format (CEF) and must meet the following specific standard requirements:

- Include the CEF headers in the template.
- Use only the extensions and custom extensions in a (Key=Value) format from the meta key table below.
- Ensure that the extensions and custom extensions are in the `key=%{string}<space>key=%{string}` format.

For third-party syslog servers, you can define your own format (CEF or non-CEF).

Procedures related to this table are described in Define a Template for Global Audit Logging and Configure Global Audit Logging.

Supported Common Event Format (CEF) Meta Keys

The following table describes the CEF Syslog meta keys that NetWitness global audit logging supports. The Datetime and Hostname fields in the Syslog Prefix are not configurable and not included in the template, but they are prepended to every log message by default. The CEF Header is required to conform to the CEF standard and for any CEF parser. The Extensions and Custom Extensions are optional. The Default Audit CEF Template contains many of the fields in this table. You can add any of the Extensions and Custom Extensions listed to the global audit logging template that you define.

CEF Field	String	Description	NW Meta Keys	Index in Log Decoder
Syslog Prefix				
Datetime	Not Configurable	Syslog Header date time	event.time.str	Transient
Hostname	Not Configurable	Syslog Header hostname	alias.host	None
CEF Header		The CEF Header fields are required to conform to the CEF standard and for any CEF parser.		
CEF:Version	CEF:0	CEF Header	--STATIC--	N/A
DeviceVendor	%{deviceVendor}	The product vendor, NetWitness	-	N/A
DeviceProduct	%{deviceProduct}	The product family. This is always NetWitness Audit.	product	Transient
DeviceVersion	%{deviceVersion}	Host/Service version	version	Transient

CEF Field	String	Description	NW Meta Keys	Index in Log Decoder
Signature ID	%{category}	Identifier of the audit event. It specifies the the category of the audit event.	event.type	None
Name	%{operation}	Description of the event	event.desc	None
Severity	%{severity}	Severity of the audit event	severity	Transient
Extensions				
deviceExternalId	%{deviceExternalId}	Unique ID of the host or service generating the audit event	hardware.id	Transient
deviceFacility	%{deviceFacility}	Syslog facility used when writing the event to syslog daemon. For example, authpriv.	cs.devfacility	Custom
deviceProcessName	%{deviceProcessName}	Name of the executable corresponding to dvcpid	process	None
dpt	%{destinationPort}	Destination Port	ip.dstport	None
dst	%{destinationAddress}	Destination IP Address	ip.dst	None
dvcpid	%{deviceProcessId}	ID of the process generating the event, which is the process ID of the NetWitness service	process.id	Transient
msg	%{text}	Free text, extra information, or actual description for the event	msg	Transient
outcome	%{outcome}	Outcome of the operation performed corresponding to the audit event	result	Transient
tpt	%{transportProtocol}	Network protocol used	protocol	Transient
userAgent	%{userAgent}	Browser detail of the user accessing the page	user.agent	Transient
rt	%{timestamp}	Time at which the event is reported	event.time	None
sourceServiceName	%{deviceService}	The service that is responsible for generating this event	service.name	Transient

CEF Field	String	Description	NW Meta Keys	Index in Log Decoder
spt	%{sourcePort}	Source Port	ip.srcport	Transient
userRole	%{userRole}	User role permissions assignment. For example: admin.owner, appliance.manage, connections.manage, everyone, logs.manage, services.manage, storedproc.execute, storedproc.manage, sys.manage, users.manage	user.role	Transient
src	%{sourceAddress}	Source IP Address	ip.src	None
suser	%{identity}	Identity of the logged on user responsible for generating the audit event	user.dst	None
Custom Extensions				
params	%{parameters}	API and Operation parameters, which capture specific parameters about a query	index	Transient
paramKey	%{key}	A configuration item key. It is the config param for which the audit event is captured. For example: /sys/config/stat.interval	obj.name	None
paramValue	%{value}	A configuration value. It is the value captured during the update.	no meta key	Custom
userGroup	%{userGroup}	Role assignment. For example: Administrators, Analysts, MalwareAnalysts, Malware_Analysts, Operators, PRIVILEGED_CONNECTION_AUTHORITY, SOC_Managers	group	None

CEF Field	String	Description	NW Meta Keys	Index in Log Decoder
referrerURL	%{referrer}	The parent URL that refers to the current URL	referer	None
sessionId	%{sessionId}	Session or connection identifier	log.session.id	Transient
remoteAddress	%{remoteAddress}	Ip address of the destination	ip.src	None
reasonForFailure	%{reasonForFailure}	reason for failure for the certain action performed	result	None
reason	%{reason}	Reason for certain action performed	result	None
addRole	%{Add.Role}	User role Assignment	user.role	Transient
id	%{id}	Incident id or host id	no meta key	Transient
arguments	%{arguments}	Value passes between programs or functions	index	Transient
uri	%{uri}	Directory	directory	None
user	%{User}	Name of the user from the source or destination	user.dst	None
accountProvider	%{AccountProvider}	Authentication account for the user. For example, PAM, and PKI.	index	Transient
file	%{file}	Name of the content file used for deployment	filename	File
deviceIDs	%{deviceIDs}	Device id for the particular service	hardware.id	Transient
role	%{Role}	User role assignment	user.role	Transient
account	%{Account}	user account	user.dst	None
addPermission	%{Add.Permission}	User role permission assignment	permissions	Transient
key	%{Key}	Name of a configuration/rule	obj.name	None
value	%{Value}	Value of a configuration change. For example, "Value":"HR12". In this example, hours format is changed to 12 hours.	no meta key	Custom

CEF Field	String	Description	NW Meta Keys	Index in Log Decoder
alert	%(alert}	Id of the alert, For example, id:5ce457afec6c0f02ffb85ace	alert	Transient
moduleSettings	%{ModuleSettings}	Message or name of a setting	index	Transient
incident	%{incident}	Id of the incident. For example, INC-313	context	None
action	%{action}	Action performed by the user. For example, service.stop	action	None
notificationBinding	%{NotificationBinding}	Type of notification. For example, incident created, alert, incident removed	index	Transient
name	%{name}	name of a configuration or rule	alert	Transient
enabled	%{enabled}	Enable the rule	no meta key	Custom
disabled	%{disabled}	Disable the rule	no meta key	Custom

Note: Use all of the extensions in the following format:
`deviceProcessName=%{deviceProcessName} outcome=%{outcome}`
Include a `<space>` between a value and a tagname.

By default, all meta keys are not indexed. In the above table, the **Index in Log Decoder** column shows the state of the `flags` keyword (Transient, None, and Custom). If a key is set to `Transient`, it is parsed but not stored in the database. If it is set to `None`, it is indexed and stored in the database. A key listed as "Custom" does not exist in the table-map.xml file and, therefore, it is not stored or parsed at all.

For more information, see the following documentation:

- The "Maintain the Table Map Files" section in the "Hosts and Services Procedures" topic in the *Hosts and Services Getting Started Guide* provides instructions for verifying and updating the table mappings.
- The "Edit a Service Index File" section in the "Hosts and Services Procedures" topic in the *Hosts and Services Getting Started Guide* provides information on updating the custom index file on the Concentrator.

Supported Global Audit Logging Meta Key Variables

This topic describes the meta key variables that NetWitness global audit logging supports.

NetWitness provides predefined global audit logging templates that you can use for your global audit logging configurations. For third-party syslog servers, you can define your own template format (CEF or non-CEF) using supported meta key variables.

Procedures related to this table are described in Define a Template for Global Audit Logging and Configure Global Audit Logging.

Supported Global Audit Logging Meta Key Variables

The following table describes the meta key variables that NetWitness global audit logging supports. Use these values to create a custom audit logging template for a third-party syslog server.

Variable	Description
%{category}	Identifier of the audit event. It specifies the the category of the audit event.
%{destinationAddress}	Destination IP Address
%{destinationPort}	Destination Port
%{deviceExternalId}	Unique ID of the service generating the audit event
%{deviceFacility}	Syslog facility used when writing the event to syslog daemon. For example, authpriv.
%{deviceProcessId}	ID of the process generating the event, which is the process ID of the NetWitness service
%{deviceProcessName}	Name of the executable corresponding to dvcpid
%{deviceProduct}	The product family. This is always NetWitness Audit.
%{deviceService}	Service responsible for generating the event
%{deviceVendor}	The product vendor, RSA
%{deviceVersion}	Host/Service version
%{identity}	Identity of the logged on user responsible for generating the audit event
%{key}	A configuration item key. It is the config param for which the audit event is captured.
%{operation}	Description of the event

Variable	Description
%{outcome}	Outcome of the operation performed corresponding to the audit event
%{parameters}	API and Operation parameters, which capture specific parameters about a query
%{referrerUrl}	The parent URL that refers to the current URL
%{sessionId}	Session or connection identifier
%{severity}	Severity of the audit event
%{sourceAddress}	Source IP Address
%{sourcePort}	Source Port
%{sourceService}	The service that is responsible for generating this event
%{text}	Free text, extra information, or actual description for the event
%{timestamp}	Time at which the event is reported
%{transportProtocol}	Network protocol used
%{userAgent}	Browser detail of the user accessing the page
%{userGroup}	Role assignment
%{userRole}	User role permissions assignment
%{value}	A configuration value. It is the value captured during the update

Global Audit Logging Operation Reference

This topic lists message types being logged by the various NetWitness components. Most messages plainly state the operation being logged; when necessary the meaning of the message is explained.

After you create a global audit logging configuration, audit logs automatically go to the external syslog system in the format specified in the selected audit logging template. The message types being logged by the various NetWitness components are shown in the following tables.

CARLOS

The following table lists the operations logged by CARLOS.

Serial #	Operation Name	Meaning
1	SetProviderConfiguration	A new notification server (for example, SMTP server) was added or updated
2	SetInstanceConfiguration	A new notification type (for example, email destination) was added or updated
3	SetTemplateDefinition	A new template was added or updated
4	RemoveProviderConfiguration	A notification server was removed
5	RemoveInstanceConfiguration	A notification type was removed
6	RemoveTemplateDefinition	A template definition was removed
7	Commit	A configuration bean change was committed
8	Set	A JMX property value was set via NetWitness Explore view

ESA

The following table lists the operations logged by the Event Stream Analysis (ESA).

Serial #	Operation Name	Meaning
9	SetSourceRequest	A concentrator was added or updated to ESA as source
10	RemoveSourceRequest	A concentrator was removed from ESA as source
11	SetEplModule	An EPL module was deployed or updated to ESA
12	RemoveEplModule	An EPL module was removed from ESA

Serial #	Operation Name	Meaning
13	SetEnrichmentSourceRequest	An ESA enrichment source was added/updated
14	RemoveEnrichmentSourceRequest	An ESA enrichment source was removed
15	SetDatabaseReference	An enrichment database reference was made to ESA
16	UpdateEnrichmentData	Data rows added to an ESA enrichment source
17	SetEnrichmentConnection	A connection was made between an EPL module and an enrichment source
18	RemoveEnrichmentConnection	A connection between an EPL module and an enrichment source was removed
19	DisableTrialModule	ESA Trial rules were disabled

Investigation

The following table lists the operations logged by Investigations.

Serial #	Operation Name	Meaning
1	VisualizePreferences	Operations related to Informer Visualization Request.
2	ParallelCoordinates	Operations related to Loading of Co-Ordinate View Navigation.
3	TimeLine	Operations related to Loading of Timeline View Navigation.
4	ExteralQuery	Operation when a Direct Query is fired via URL.
5	PrintView	Operations to open Investigation in Print View.
6	submitExtractFiles	Operation to submit a Request to Extract files from Sessions.
7	submitExtractLogs	Operation to submit a Request to Extract Logs from Sessions.
8	submitExtractPcap	Operation to submit a Request to Extract Sessions from Sessions.
9	DataScienceDrill	Operation to investigate from Data Science Report.
10	breadCrumbs	Operation to access the Query Breadcumbs.
11	Create	Operation when a new Investigation Query is being saved as a predicate to be used for URL Integration.
12	userPredicates	Operation to access Recent Queries of a user.

Serial #	Operation Name	Meaning
13	chartDefaultMetas	Operation to access last used Meta for generating Coordinate Chart.
14	defaultDevice	Operation to access the Default Investigation Device.
15	deleteDefaultDevice	Operation to delete the Default Investigation Device.
16	chartPreferences	Operation to edit an Investigation Navigation Chart Parameters such as Height.
17	devicePreferences	Operation to save the preferences about the Investigation Device such asTime Range, Profile, Meta Groups etc.
18	topValues	Operation to get the Top Values for Metas. Normally called from Top Values Dashlet.
19	MetaLanguages	Operation to read the Meta Languages from a Device.
20	MetaGroups	Operations related to Investigation Meta Groups.
21	DefaultMetaKeys	Operations related to Investigation Default Meta Keys.
22	UpdateDefaultMetaKeys	Operations to update Investigation Default Meta Keys.
23	UpdateMetaGroup	Operations to update Investigation Meta Groups.
24	ApplyMetaGroup	Operations to use Investigation Meta Groups.
25	DeactivateMetaGroup	Operations to reset Investigation Meta Groups in UI.
26	DeleteMetaGroup	Operations to remove Investigation Meta Group.
27	DeleteMetaGroups	Operations to remove multiple Investigation Meta Groups.
28	ImportMetaGroups	Operations to import Investigation Meta Groups.
29	ExportMetaGroup	Operations to export multiple Investigation Meta Groups.
30	GeoMap	Operation to access the Geo Map View of Investigation.
31	deleteEndpointCache	Operation to clear Reconstruction Cache of a Device.
32	delete	Operation to delete Alert Templates.
33	CustomColumnGroup	Operation to apply or read Custom Column Group.
34	Import	Operations related to Import of Column Group or Profiles.
35	Export	Operations related to Export of Column Group or Profiles.
36	SaveProfile	Operation to save an Investigation Profile.
37	ApplyProfile	Operation to apply an Investigation Profile.

Serial #	Operation Name	Meaning
38	DeactivateProfile	Operation to deactivate an Investigation Profile.
39	DeleteProfile	Operation to delete an Investigation Profile.
40	DeleteProfiles	Operation to delete multiple Investigation Profiles.

Reporting Engine

The following table lists the operations logged by the Reporting Engine.

Serial #	Operation Name	Meaning
1	TEMPLATE	For all operations related to template
2	CHART	For all operations related to chart
3	REPORT	For all operations related to report
4	RULE	For all operations related to rule
5	IMAGE	For all operations related to Logo Images used in Reports.
6	LIST	For all operations related to list
7	ALERT	For all operations related to alert
8	CONFIG	For all operations related to configuration change
9	SCHEDULE	For all operations related to schedule
10	ROLE	For all operations related to role/authorization
11	BATCH_JOB	For all operations related to batch jobs
12	SCHEDULER	For all operations related to scheduler
13	QUERYPROCESSOR	For all operations related to queryprocessor
14	FORMATTER	For all operations related to formatter
15	OUTPUTACTION	For all operations related to outputaction
16	STATUSMANAGER	For all operations related to statusmanager
17	BATCH_RUNDEF	For all operations related to batch rundef
18	CHARTGROUP	For all operations related to chart group
19	REPORTGROUP	For all operations related to report group
20	RULEGROUP	For all operations related to rule group

Serial #	Operation Name	Meaning
21	LISTGROUP	For all operations related to list group
22	DISKSPACE	For all operations related to disk space

Warehouse Connector

The following table lists the operations logged by the Warehouse Connector.

Serial #	Operation Name	Meaning
1	LockBox Password Create	Operation to create LockBox Password.
2	LockBox Password Update	Operation to update LockBox Password.
3	LockBox Password Refresh	Operation to refresh LockBox Password.
4	Adding Stream	Operation to add a Stream.
5	Adding Source	Operation to add a Source.
6	Adding Destination	Operation to add a Destination.
7	Removing	Operation to remove a Source, Stream, or Destination.
8	Changing Password	Operation to change the Password.
9	Updating Source	Operation to update a Source.
10	Adding Source to Stream	Operation to add a Source to a Stream.
11	Deleting Source from Stream	Operation to delete a Source from a Stream.
12	Setting Destination to Stream	Operation to set a Destination to a Stream.
13	Finalizing Stream	Operation to finalize a Stream and initiate the aggregation.
14	Stopping Stream	Operation to stop a Stream.
15	Starting Stream	Operation to start a Stream.
16	Reloading Stream	Operation to reload a Stream.

Health & Wellness

The following table lists the operations logged by Health & Wellness.

Serial #	Operation Name	Meaning
1	SavePolicyRequest	Operation while adding or modifying a Policy.
2	RemovePolicyRequest	Operation while removing a Policy.

NetWitness Core Services

The following table lists the operations logged by NetWitness Core Services.

Serial #	Operation Name	Meaning
1	FILE-Command	Operation to list, retrieve and delete files from approved directories on this device.
2	SERVICE-Start	Service started
3	SERVICE-Stop	Service stopped
4	REDIRECT-Syslog	Operation for syslog forwarding.
5	ADD-Monitor	Issuing a filesystem monitor operation
6	DELETE-Monitor	Issuing a filesystem monitor deletion operation
7	SHUTDOWN-Service/shutdown.service	Shutting down appliance service
8	REBOOT-Service	Restarting appliance service
9	CONFIGURE-Network	Issuing Network Configuration change
10	SET-NTP	Issuing NTP set operation
11	STOP-NTP	Issuing NTP stop operation
12	NTP-Timesync	Issuing NTP time sync operation
13	SET-SNMP	Issuing SNMP set
14	UPGRADE/upgrade	Issuing upgrade operation
15	create.collection	Operation to create an empty collection.
16	restore	Issuing restore
17	session.aggregation	Issuing aggregation start/stop
18	add.device	Adding a device for aggregation
19	edit.device	Editing a device used for aggregation
20	delete.device	Deleting a device used for aggregation
21	capture.start	Starting capture operation
22	capture.stop	Stopping capture operation
23	select.interface	Selecting capture interface
24	export	Operation to export packets or sessions.

Serial #	Operation Name	Meaning
25	reload	Issuing a parser reload
26	schema	Issuing a schema request for loaded parsers
27	upload/file.upload	Issuing file upload
28	notify	Issuing feed notify
29	delete	Issuing file deletion
30	edit.config	Configuration change operation
31	parsers.transforms	Perform a language key transformation
32	data.reset	Data reset operation
33	timeout	REST request timeout
34	cancel	Cancel a running query
35	timeroll	Operation to delete the database files that exceed a given limit.
36	dump	Operation to dump information out of the database in nwd formatted files.
37	session.wipe	Issuing a session wipe operation
38	REPLACE-Rule	Issuing a rule replace operation
39	MERGE-Rule	Issuing a rule merge operation
40	ERASE-Rule	Issuing deletion of a set of all rules
41	ADD-Rule	Issuing a rule addition operation
42	DELETE-Rule	Issuing deletion of a set of rules
43	sdk.info	Issuing SDK summary info.
44	sdk.session	Issuing SDK session info.
45	sdk.language	Issuing SDK language
46	sdk.aliases	Issuing SDK alias request
47	sdk.transform	Issuing SDK transformation request
48	sdk.search	Issuing session content search request
49	sdk.cache	Operation related to session content cache
50	sdk.content	Issuing session content request

Serial #	Operation Name	Meaning
51	check.authorization	Operation to check user roles for permissions to execute an operation.
52	close.connection	Issuing a connection close operation
53	handshake	Issuing an SSL handshake
54	logon/login	Operation to login from NW to the other services, mostly to privileged users.
55	STOREDPROCOP	Issuing file upload cancel/start
56	ADD-Task	Added scheduled task
57	DELETE-Task	Deleted scheduled task
58	logoff	Issuing logout operation
59	list.cacerts	Issuing list trusted CA certificate operation
60	delete.cacerts	Issuing delete trusted CA certificate operation
61	add.cacerts	Issuing addition of trusted CA certificate operation
62	restart.command	Issuing restart command line option
63	delete.file/file.delete	Operation to delete system configuration files.
64	update.file/file.update	Operation to update system configuration file.
65	create.file	Issuing file creation operation
66	query	Issue a database query
67	unlock	Issuing unlock user account operation
68	user.add	Operation to create user accounts on individual devices.
69	user.delete	Operation to delete a user on individual devices.
70	group.create	Operation to add a new group to the system.
71	user.remove	Remove a user account from a group
72	group.delete	Delete a group from the /users/groups tree
73	add.user	Issuing add user command to collection
74	delete.user	Issuing delete user command to collection
75	remove.user	Removing an user from collection
76	collection.open	Issuing an open command for a collection

Serial #	Operation Name	Meaning
77	collection.close	Issuing a close command for a collection
78	collection.delete	Issuing collection deletion command
79	reingest.start	Operation to start reingesting of packet data in collection.
80	feed.notify	Issuing a feed notify command
81	collect	Issuing a collect command
82	collect.start	Issuing a data collection start
83	collection.global	Issuing import parser command
84	parser.reload	Issuing parser reload command
85	reingest	Operation to reingest packet data in collection.
86	collection.create	Issuing a create collection command
87	collection.restore	Issuing a restore collection command
88	collection.clone	Issuing a clone collection command
89	parser.reload	Issuing parser reload command
90	sdk.query	Performs a query against the meta database
91	sdk.msearch	Search for pattern matches in many sessions or packets
92	sdk.values	Performs a value count query and returns the matching values for a report
93	sdk.timeline	Returns the count of sessions/size/packets in discrete time intervals

Malware Analysis

The following table lists the operations logged by the Malware Analysis (MA) component.

Serial #	Operation Name	Meaning
1	GetDashBoardSummaryRequest	Get dashboard analysis statistics
2	GetFileScoreSummaryRequest	Get aggregated file scores by score type and risk level
3	CountEventsAndFilesRequest	Get count of events and files over a time frame
4	GetAvVendorDetectionRequest	Get AV vendor analysis results

Serial #	Operation Name	Meaning
5	GetAVVendorsRequest	Get list of AV Vendors supported
6	SetInstalledAVVendors	Request Update list of installed AV Vendors in config
7	CountEventByCriteriaRequest	Count events by criteria
8	FindEventByIdRequest	Get event by id
9	FindEventByCriteriaRequest	Get event by criteria
10	DeleteEventRequest	Delete event
11	CommentOnEventRequest	Add comment to event
12	ReSubmitEventRequest	Resubmit event for analysis
13	FindEventScoreByIdRequest	Get event score by event id
14	FindEventScoreByCriteriaRequest	Get event score by criteria
15	FindMetaByIdRequest	Get meta by id
16	FindMetaByCriteriaRequest	Get meta by criteria
17	FindMetaValueByCriteriaRequest	Get meta value by criteria
18	CountByDistinctMetaValueRequest	Count distinct meta values
19	CountByMetaNameAndValueWithDate RangeIntervalRequest	Count meta and values with interval for charting
20	CountByValueAndAverageOverallScore Request	Count meta and map to overall scores for events
21	CountByValueAndAverageGroupScore Request	Count meta and map to group scores for events
22	CountFileEntryByCriteriaRequest	Count files by criteria
23	FindFileEntryByIdRequest	Get file by id
24	FindFileEntryByCriteriaRequest	Get file by criteria
25	ReSubmitFileEntryRequest	Resubmit file for analysis
26	FileDownloadRequest	Download file from repository
27	FileUploadRequest	Upload file for analysis
28	FindFileScoreByIdRequest	Get file score by id
29	FindFileScoreByCriteriaRequest	Get file score by criteria

Serial #	Operation Name	Meaning
30	FindHashValueByIdRequest	Get whitelist/blacklist Hash value by id
31	FindHashValueByCriteriaRequest	Get whitelist/blacklist Hash value by criteria
32	AddHashValueRequest	Add whitelist/blacklist Hash value
33	UpdateHashValueRequest	Update whitelist/blacklist Hash value
34	DeleteHashValueRequest	Delete whitelist/blacklist Hash value
35	FindHashValueByMd5Request	Find whitelist/blacklist Hash value by md5
36	AddHashValueInFileRequest	Add File to repository as well as hash value
37	GetDefaultRulesRequest	Get default IOC Rules configuration
38	ResetToDefaultRulesRequest	Reset IOC Rules configuration to default
39	GetAllOverrideRulesRequest	Get IOC Rules user created override configuration
40	FindOverrideRuleByIdRequest	Find IOC override rule by id
41	AddOverrideRuleRequest	Add IOC override rule
42	UpdateOverrideRuleRequest	Update IOC override rule
43	DeleteOverrideRuleRequest	Delete IOC override rule
44	SubmitOnDemandNextGenRequest	Submit new ondemand nextgen scan
45	FindOnDemandJobEntryByIdRequest	Get ondemand job entity by id
46	FindOnDemandJobEntryByCriteria Request	Get ondemand job entity by criteria
47	GetOnDemandJobInfoRequest	Get ondemand job reference entity by id
48	GetOnDemandDefaultConfiguration	Request Get ondemand default configuration
49	CancelOnDemandJobRequest	Cancel ondemand job in progress
50	DeleteOnDemandJobRequest	Delete ondemand job
51	ReSubmitOnDemandJobRequest	Resubmit ondemand job
52	SubscriptionRequest	Subscribe to MA Cloud communication
53	UnSubscribeRequest	Unsubscribe from MA Cloud communication
54	GetTopEventInfluencesRequest	Get Top N event influences

Serial #	Operation Name	Meaning
55	GetServerInfoRequest	Get server info, such as server time
56	DataResetRequest	Reset database
57	OnDemandJobStatusNotification	Report ondemandjob progress to subscribers
58	LicenseStatusNotification	Report license status - num samples analyzed
59	DataResetNotification	Report that data was reset
60	GetIocSummaryRequest	Get IOC rules aggregated by event/file scores
61	FindAlertTemplatesByCriteriaRequest	Get rabbitmq alert templates by criteria
62	SaveAlertTemplateRequest	Update alert template
63	DeleteAlertTemplateRequest	Delete alert template
64	GetJobStatusRequest	Get in progress job analysis thread status
65	GetEventTypeCountSummaryRequest	Get event analysis counts by date chart
66	Logon	Logon to the MA Service
67	Modified	Modifying config changes
68	GetNextGenSummaryRequest	Get nextgen dashboard summary statistics

NetWitness User Interface

The following table lists the operations logged by the NetWitness User Interface component.

Serial #	Operation Name	Meaning
1	uploadTrialLicense	Upload Trial License
2	LicenseEntitle	Entitle License
3	LicenseDeactivation	Deactivate License
4	ExpiredLicense	License Expired
5	LicenseOutOfComplianceAcknowledgement	EULA Acknowledgement
6	resetLicense	Reset License
7	usageDateExport	License data usage - csv/pdf

Serial #	Operation Name	Meaning
8	refreshLicense	Refresh LLS license
9	LicenseOutOfCompliance	Out of Compliance
10	OOTBEntitlementOutOfCompliance	OOTB Trial license Out of Compliance
11	OOTBEntitlementFirstLoginTimeModified	OOTB time modified
12	OOTBEntitlementFileDeleted	OOTB File deleted
13	OOTBEntitlementDataTampering	OOTB data tampering
14	uploadOfflineResponse	Upload offline response
15	offlineDownloadCapRequest	Download offline request
16	movePerpetualToThroughput	Move Appliance license to Throughput
17	moveThroughputToPerpetual	Mover Throughput to Appliance license
18	mapApplianceLicense	Map Service to Real license
19	delete	Operation to delete Alert Templates.
20	HttpRequest	Operation for Audit Logging of the accessed URL.
21	Page Accessed	Operation for Audit Logging of the accessed page.
22	Navigate	Operation to navigate to the accessed page.
23	Events	Operation to view the accessed event page.
24	Recon	Operation for Event Reconstruction requested.
25	Services	Operation while reading the list of available devices for investigation.
26	Service	Operation for a List of devices requested to be investigated.
27	Collections	Operation to view the list of collections requested.
28	Profiles	Operation to apply a Profile.
29	ColumnGroups	Operation to apply or read Column Group.

Serial #	Operation Name	Meaning
30	ParallelCoordinates	Operations related to Loading of co-ordinate view navigation.
31	Timeline	Operations related to loading of timeline view navigation.
32	PrintView	Operations to open investigation in print view.
33	Preferences	Operations related to Informer Request.
34	import	Operations related to Import of Column Group or Profiles.
35	export	Operations related to Export of Column Group or Profiles.
36	Predicate	Operations related to Queries (Predicates) used for Investigation.
37	Languages	Operation for Language requested from a Device.
38	CancelLanguageLoad	Operation for Language Load Canceled from Navigate Page.
39	summary	Operation for a summary requested from a Device.
40	languages	Operation for a language requested from a device.
41	aliases	Operation for meta aliases requested from a device.
42	query	Operation for SDK Query requested from a device.
43	msearch	Operation for a meta search requested from a device.
44	nodeListing	Node Listing for a node requested from a Device.
45	content	SDK Content call requested from a Device for downloading a PCAP or Log.
46	Export Files	File Listing Requested for a Session in File View or Extraction jobs.
47	packets	Packets requested for sessions in Packet View or Extraction Jobs.

Serial #	Operation Name	Meaning
48	deleteEndpointCache	Operation to clear reconstruction cache of a device.
49	Logon	Operation for user to sign in to NetWitness User Interface.
50	Logoff	Operation for user to sign out of NetWitness User Interface.
51	defaultDevice	Operation to access the Default SA UI Device.
52	deleteDefaultDevice	Operation to delete the Default investigation device.
53	submitExtractFiles	Operation to submit a request to Extract files from Sessions.
54	submitExtractLogs	Operation to submit a Request to Extract Logs from Sessions.
55	submitExtractPcap	Operation to submit a Request to Extract Sessions from Sessions.
56	MetaGroup	Operations related to SA UI Meta Groups.
57	ExternalQuery	Operation when a Direct Query is fired via URL.
58	GeoMap	Operation to access the Geo Map View of Investigation.
59	SaveProfile	Operation to save an Investigation Profile.
60	ApplyProfile	Operation to apply an Investigation Profile.
61	DeleteProfile	Operation to apply an Investigation Profile.
62	DeactivateProfile	Operation to apply an Investigation Profile.
63	VisualizePreferences	Operations related to Informer Visualization Request.
64	ExportMetaGroup	Operations to export multiple SA UI Meta Groups.
65	userPredicates	Operations to export multiple SA UI Meta Groups.

Serial #	Operation Name	Meaning
66	FileView	Operation for reconstruction request for File View.
67	resource.update	Operation when Live Subscription State changes.

Respond

The following table lists the operations logged by the Respond component.

Serial #	Operation Name	Meaning
1	update	Update notification setting
2	update	Update integration settings configuration
3	delete	Delete Alerts
4	create	Create new incident
5	update	Update incident details
6	read	Read incident details
7	delete	Delete incidents
8	read	Read remediation tasks
9	delete	Delete Remediation tasks
10	update	Update remediation tasks
11	create	Create new rule
12	update	Update existing alert rule
13	reorder	Reorder priority of alert rules

Investigate Server

Serial #	Operation name	Meaning
1	Aliases	Fetch aliases
2	BackgroundJob	Category for all background job
3	ColumnGroup	Category for all column group operation

Serial #	Operation name	Meaning
4	Count-query	Default Investigate
5	Countdistict-query	Default Investigate
6	Content	Default Session Content
7	Create	User entity
8	Create	User preferences
9	Create	Predicate
10	Delete	User preferences
11	Delete	Predicate
12	Extract-content	Extract Content from Session
13	Folders	Category for all folder operation
14	Files	SDK content request
15	InvestigateExport	Category for extraction invocation
16	Key-refs	Fetch meta keys
17	Languages	SDK language call
18	MetaGroup	Category for all meta group operation
19	Metakey	Category for all metakey operation
20	MailReconstruction	Category for all mail reconstruction operation
21	ParsingRequest	Category for all request parsing operation
22	PacketReconstruction	Category for all packet reconstruction operation
23	Predicate InvestigateIncidents	Category for all predicate operation
24	Query	SDK query
25	Reconstruction	Category for common shared reconstruction operation
26	ReconstructionCache	Category for reconstruction caching operation
27	ReconstructionStreaming	Category for reconstruction streaming operation
28	ReconstructDataSecurity	Category for reconstruction security operations (data-scrubbing)
29	Session-Meta	SDK query to get session meta

Serial #	Operation name	Meaning
30	Summary	SDK summary
31	Search-meta-value	Search meta-value based on field name
32	TextReconstruction	Category for all text reconstruction operation
33	Timeline	Timeline request
34	UserPreferences	Category for all user preferences operation
35	Update	User entity
36	Update	User preferences
37	Update	Profile group
38	Update	Predicate
39	Values	SDK values
40	Validate-query	Validate SDK query

Security Server

The following table lists the events logged by the Security Server.

Log Category	Description
Create:	Add record for a role with new ID.
Create:	Add user record with new ID.
Update:	Update the user account with a new ID.
Authentication:	Logs events pertaining to user logins and logouts.
Authorization:	Logs events pertaining to user access checks and RBAC management.
UserAccount	Logs events pertaining to NetWitness domain account management.
ExternalProvider:	Tracks events pertaining with external account providers (for example, Active Directory).

The following example shows an event logged by the Security Server:

```
2018-03-13 16:25:02,938 UserAccount{action=ExpirePassword, success=true,
identity=admin, parameters={id=Justin}}
```

Admin Server

The following table lists the events logged by the Admin Server.

Log Category	Description
Restore:	System operation to restore a data springboard.

Config Server

The following table lists the events logged by the Admin Server.

Log Category	Description
newregistration:	Record the registration to config server that manages the collection storage.

Context Hub Server

The following table lists the events logged by the Admin Server.

Log Category	Description
verifyconnection:	System operation to check if the connection is live.
addconnection	Create new connection to access data.

Local Audit Log Locations

NetWitness has global audit logging capabilities. When you configure global audit logging, audit logs from all NetWitness components collect in a centralized system, which converts them into the required format and forwards them to a third-party syslog server or a Log Decoder.

To view audit logs from the individual services, you can look at the local audit log locations.The following table shows the local directory paths of the audit logs for the NetWitness user interface and the various NetWitness services.

Service/Module	Audit Log Location
NetWitness User Interface (NetWitness Web Server)	The NetWitness user interface sends audit logs to the following locations: • `/var/lib/netwitness/uax/logs/audit/audit.log` (human-readable format) • Syslog running on the local host (JSON format) The NetWitness user interface uses the AUTH facility of syslog to write audit logs to syslog. You can only see audit logs in the first location (`/var/lib/netwitness/uax/logs/audit/audit.log`).
Core Services (Decoder, Log Decoder, Concentrator, Broker, and Archiver), Log Collector, Warehouse Connector, and Workbench	The Core services and similar services send audit logs to Syslog running on the local host. Path: `/var/log/secure` (JSON format) The Core services use the AUTHPRIV facility of syslog to write audit logs to syslog.

Service/Module	Audit Log Location

Reporting Engine,
Malware Analysis,
Respond,
ESA Correlation (11.3 and later),
and Event Stream Analysis
(11.2 and earlier)
	These services send audit logs to the following locations:
	• `<application-home-directory>/logs/audit/audit.log` (human-readable format)
	• Syslog running on the local host (JSON format)
	
	The following are the audit log locations of these services:
	Reporting Engine:
	`/var/netwitness/re-server/rsa/soc/reporting-engine/logs/audit/audit.log`
	Respond Server:
	`/var/log/netwitness/respond-server/respond-server.audit.log`
	Malware Analysis:
	`/var/lib/netwitness/malware-analytics-server/spectrum/logs/audit/audit.log`
	ESA Correlation (11.3 and later):
	`/var/log/netwitness/correlation-server/correlation-server.audit.log`
	Event Stream Analysis (11.2 and earlier):
	`/opt/rsa/esa/logs/audit/audit.log`
	These services use the AUTH facility of syslog to write audit logs to syslog. You can only see audit logs in the first location (`<application-home-directory>/logs/audit/audit.log`).

Health & Wellness, Event Source Management (ESM), and Appliance and Service Grouping (ASG)	These Services send audit logs to the following locations: • `/opt/rsa/sms/logs/audit/audit.log` (human-readable format) • Syslog running on the local host (JSON format) These services use the AUTH facility of syslog to write audit logs to syslog. You can only see audit logs in the first location (`/opt/rsa/sms/logs/audit/audit.log`).
Aggregated Audit Logs	The aggregated audit logs from all the services are sent to the following locations: • `/var/netwitness/logstash/logs/rsa-netwitness-audit.log` (JSON format) • Syslog running on the local host (human-readable format)

Global Notifications Panel

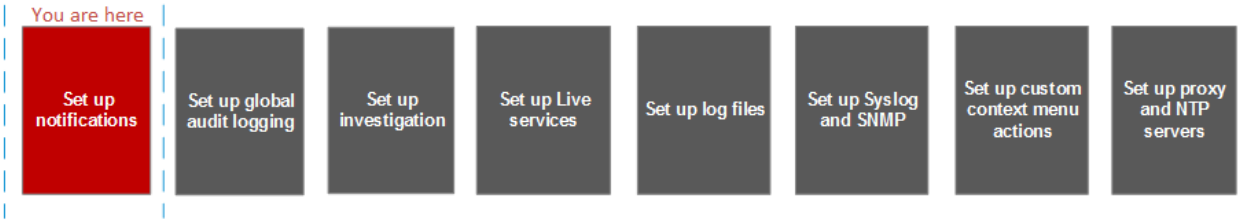
Global Notifications panel introduces the features for configuring notification settings. Global Notifications configurations define notifications settings for Event Source Management (ESM), Health and Wellness, Global Audit Logging, Event Stream Analysis (ESA), and Respond.

In the Global Notifications panel, you can configure the following global notification settings:

- Notification Outputs
- Notification Servers
- Templates

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

WorkFlow



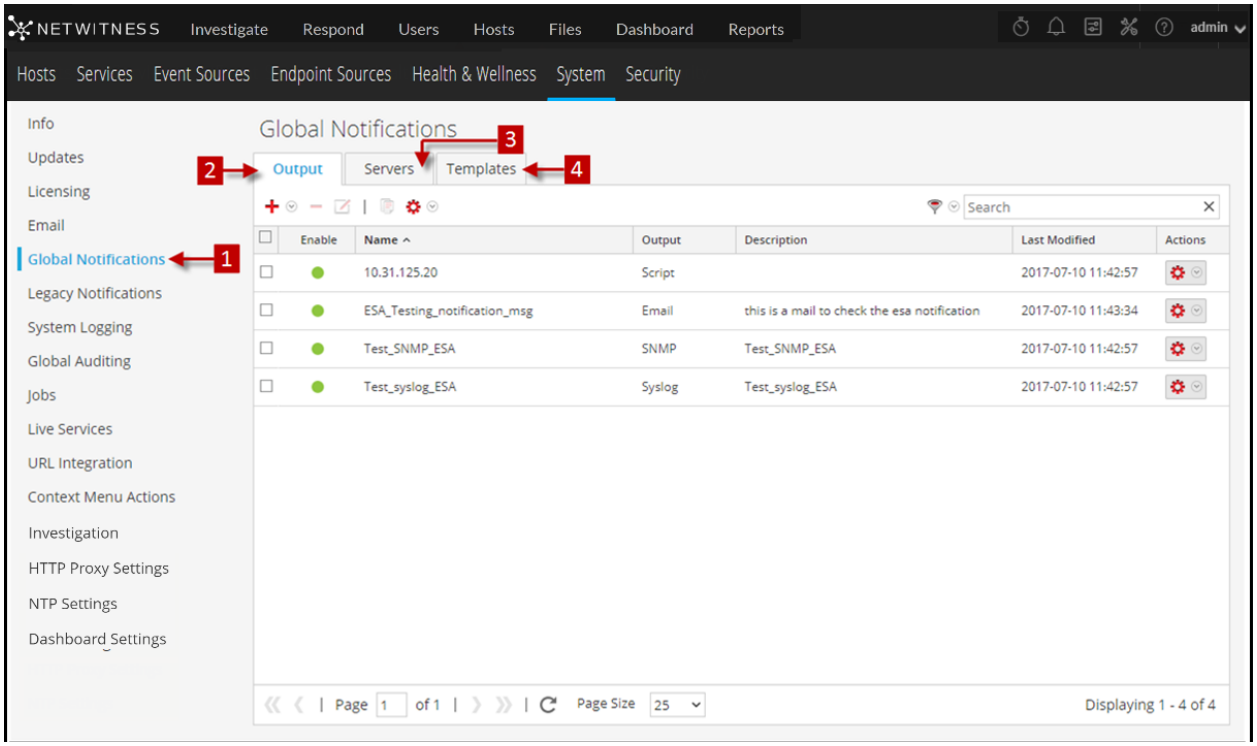
What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure Notification Servers	Servers Tab
Administrator	Configure Notification Outputs	Output Tab
Administrator	Configure Notification Templates	Templates Tab

Related Topics

- Configure a Syslog Notification Server
- Configure Script as a Notification Server

Quick Look



1 Displays the Global Notification Panel.
2 Displays the Output Tab
3 Displays the Servers Tab
4 Displays the Templates Tab

Toolbar and Features

The Global Notifications panel has three tabs: Output, Servers, and Templates.

Feature	Description
Output tab	This tab enables you to configure notification outputs. See Output Tab for more information.
Servers tab	This tab enables you to configure notification servers. See Servers Tab for more information.
Templates tab	This tab enables you to configure notification templates. See Templates Tab for more information.

This table describes the columns in the grid for Notification Outputs and Notification Servers.

Column	Description
	Selects a row for an action in the toolbar. Clicking the checkbox in the column title selects or deselects all rows in the grid.
Enable	Indicates whether the configuration is enabled. A solid colored green circle indicates that a configuration is enabled. A blank white circle indicates that a configuration is not enabled.
Name	A name that identifies or labels the configuration.
Output	The configuration output. The outputs are Email, SNMP, Syslog, and Script.
Description	A brief description about the configuration.
Last Modified	Shows the date and time of the last configuration change.
Actions	Provides an Actions menu for the selected configuration with actions that can be taken on the configuration. The Actions menu enables you to delete, edit, duplicate, and export the configuration.

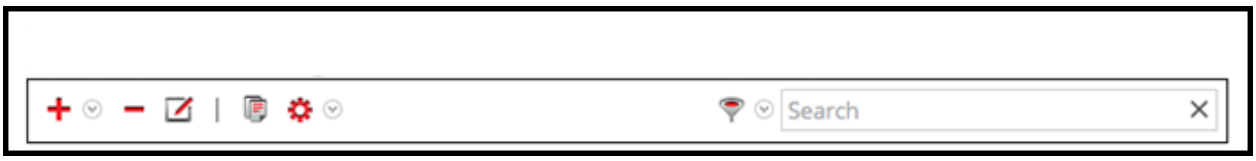
This table describes the columns in the grid for Notification Templates.

Column	Description
	Selects a row for an action in the toolbar. Clicking the checkbox in the column title selects or deselects all rows in the grid.
Name	A name that identifies or labels the template.
Template Type	The type of template. The types are Audit Logging, Event Stream Analysis, Event Source Monitoring, and Health Alarms.
Description	A brief description about the template.
Actions	Provides an Actions menu for the selected configuration with actions that can be taken on the template. The Actions menu enables you to delete, edit, duplicate, and export the template.

Global Notifications Panel Toolbar

The Global Notifications panel toolbar is at the top of the Output, Servers, and Templates tabs.

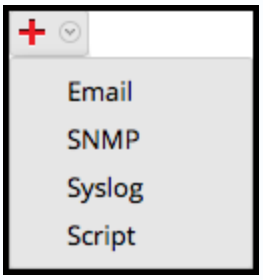
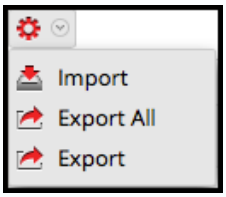
The following figure shows the toolbar on the Output and Servers tabs.

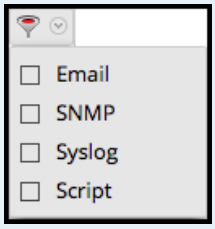


The following figure shows the toolbar on the Templates tab.



The following table describes the features of the Global Notifications panel toolbar.

Feature	Description
 	Adds a notification server on the Servers tab, adds a notification output (notification) on the Output tab, and adds a notification template on the Templates tab. On the Servers and Output tabs, you can select to configure Email, SNMP, Syslog, and Script notification settings.
	Removes a selected notification configuration. You cannot delete notification servers and notification types that are associated with global audit log configurations. If you attempt to delete a notification output (notification) being used by alerts, you will receive a warning confirmation message that the alerts using the notification will not function properly. The message shows the number of alerts in use. You can also delete a configuration by selecting a configuration and then in the Actions column, selecting  > Delete.
	Edits a selected notification configuration. You can also edit a configuration by selecting a configuration and then in the Actions column, selecting  > Edit.
	Duplicates a selected notification configuration. You can also duplicate a configuration by selecting a configuration and then in the Actions column, selecting  > Duplicate.
	Displays the following options: • **Import**: Imports a notification server, type, or template. For example, on the Servers tab, you can import a notification server configuration. • **Export All**: Exports all of the configurations. For example, if you are on the Servers tab, you can export all of the notification server configurations. • **Export**: Exports a selected configuration. You can also export a configuration by selecting a configuration and then in the Actions column, selecting  > Export.

Feature	Description
	Filters by Email, SNMP, Syslog, or Script.
	Searches configurations in the grid.

Define Notification Server Dialogs

This topic describes the Define Notification Server dialogs used to configure the settings of the various types of notification servers. You configure notification servers in the (Admin) > System > Global Notifications > Servers tab.

Notifications are used by a variety of components in NetWitness, such as Event Stream Analysis (ESA), Respond, and Global Audit Logging. Notification settings are called Notification Servers. In the Servers tab of the Administration System view Notifications panel, you can create multiple Notification Server configurations.

You can configure the following types of notification server settings in NetWitness:

- Email
- SNMP
- Syslog
- Script

For Global Audit Logging, you can only use Syslog Notification Servers.

Procedures related to notification servers are described in Configure Notification Servers.

To access the Define Notification Server dialogs:

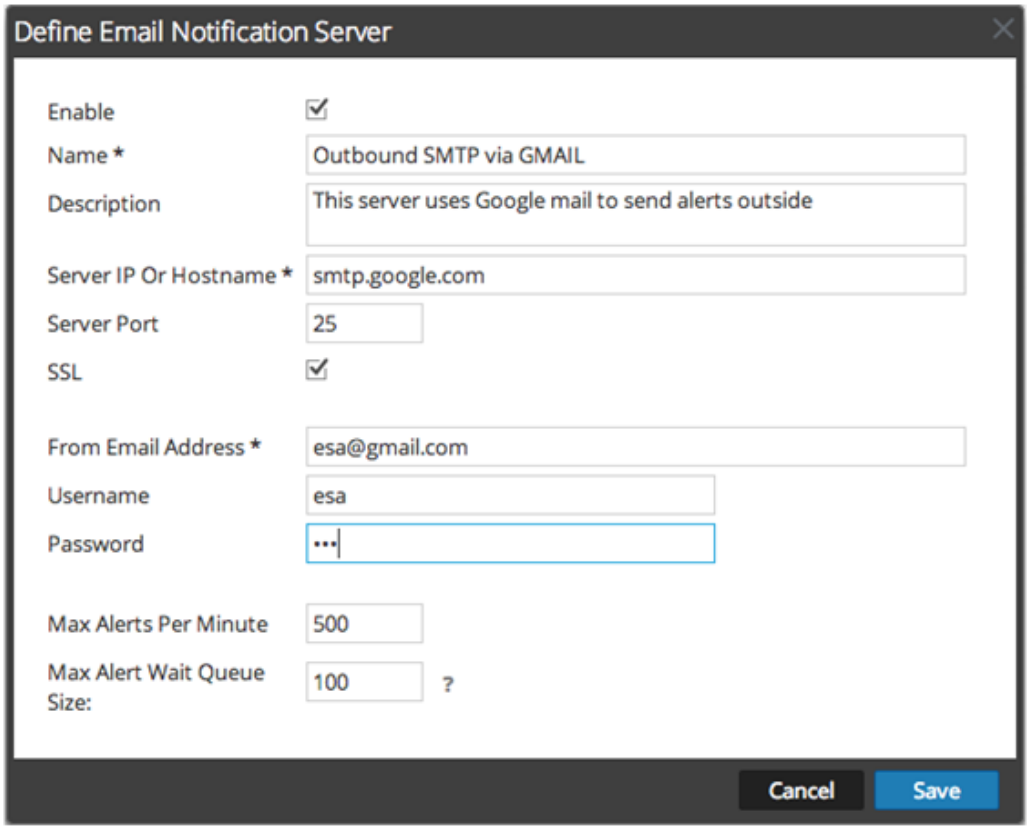
1. Go to **(Admin) > System**.
2. In the left navigation panel, select **Global Notifications**.
3. In the **Notifications Servers** panel, click and then select a type of notification server (Email, SNMP, Syslog, or Script)
   The Define Notification Server dialog is displayed for your selection.

There are four notification server dialogs, which allow you to configure notification servers.

Email

Email notification servers enable you to configure email server settings to send alert notifications.

The following figure shows the Define Email Notification Server dialog.



The following table lists the various parameters that you need to define for the email notification servers.

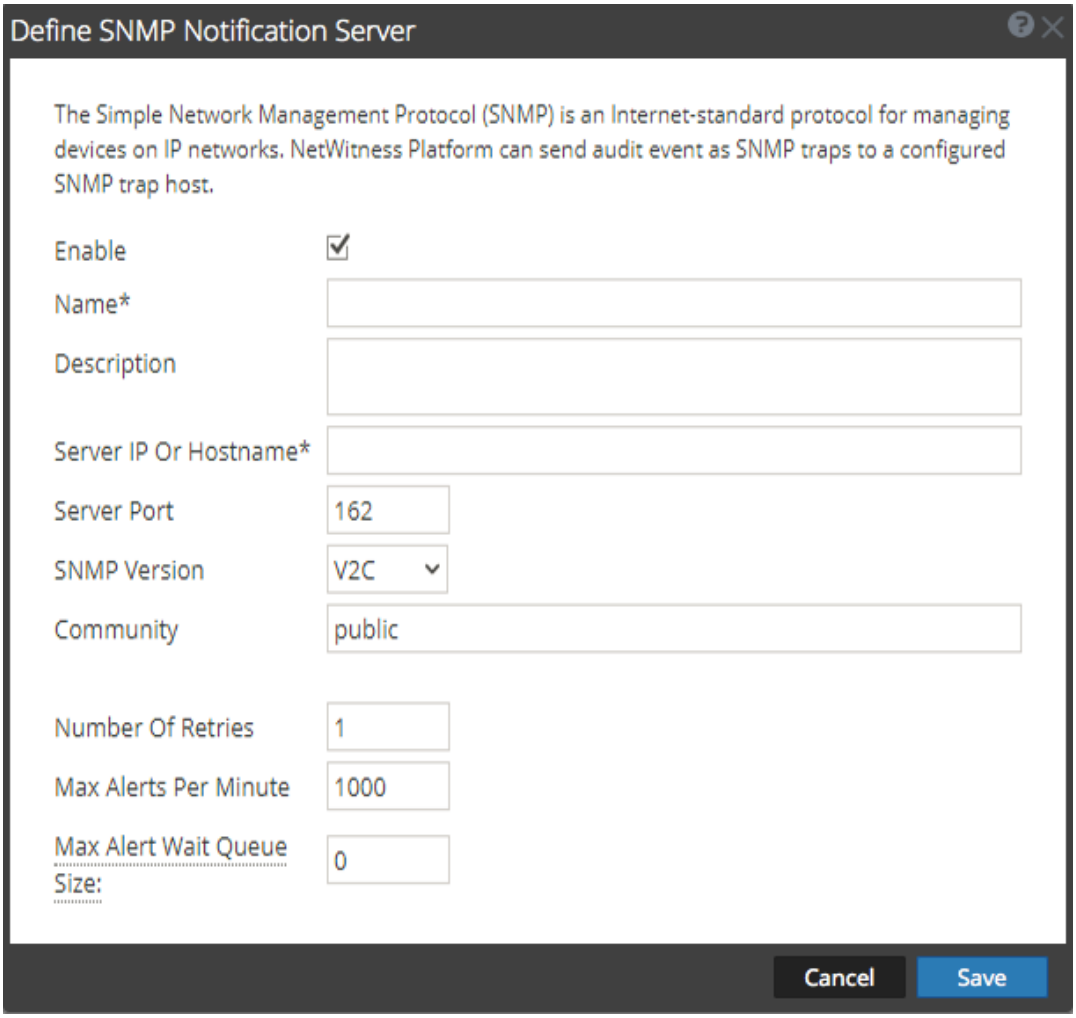
Parameters	Description
Enable	Select to enable the notification server.
Name	A name to identify or label the notification server.
Description	A brief description about the notification server.
Server IP Or Hostname	Hostname of the email server. For ESM/SMS and ESA notifications, you must specify only the hostname/FQDN.
Server Port	The server port.
SSL	Select the option if you want the communication to happen through SSL.
From EMail Address	Email account from which you want to send email notifications.
Username	Username for logging into the email account if the SMTP server requires user authentication to relay emails successfully.

Parameters	Description
Password	User password for logging into the email account if the SMTP server requires user authentication to relay emails successfully.
Max Alerts Per Minute	Describes the maximum number of alerts per minute.
Max Alert Wait Queue Size	Describes the maximum number of alerts to be queued before they are dropped.

SNMP

SNMP notification servers enable you to configure SNMP trap host settings as a notification server to send alert notifications.

The following figure shows the Define SNMP Notification Server dialog.



The following table lists the various parameters that you need to define for the SNMP notification servers.

Parameters	Description
Enable	Select to enable the notification server.
Name	A name to identify or label the notification server.
Description	A brief description about the notification server.
Server IP Or Hostname	SNMP trap host IP address or hostname.
Server Port	Listening port number on the SNMP trap host.

<table>
<tr><th>Parameters</th><th>Description</th></tr>
<tr><td>SNMP Version</td><td>SNMP version. The following are the options:
V1
V2C
V3
If you select SNMP Version 3 (v3), the following parameters are displayed:
<table>
<tr><th>Parameters</th><th>Description</th></tr>
<tr><td>Notification Type</td><td>Based on the notification type a SNMP messages are sent each time an alert is generated.
The following notification types are supported:
Inform - Inform is acknowledged trap. The sender gets an acknowledgement from the receiver.
Trap - Trap is unacknowledged notification</td></tr>
<tr><td>Authoritative Engine ID (This optioin is availabe only for notification type TRAP)</td><td>An identifier which is used to identify the agents. Authoritative engine ID along with the username is used to uniquely identify the agent.</td></tr>
<tr><td>Security Level</td><td>Define the security level. The following are the options:
Unauthenticated and Unencrypted
Authenticated and Unencrypted
Authenticated and Encrypted</td></tr>
<tr><td>Auth Protocol (This option is available only for security level Authenticated and Unencrypted and Authenticated and Encrypted)</td><td>Authentication protocol which is used to validate a user before providing an access to the server. The options are:
SHA
MD5</td></tr>
<tr><td>Auth Key (This option is available only for security level Authenticated and</td><td>A password that you want to use for authentication.</td></tr>
</table></td></tr>
</table>

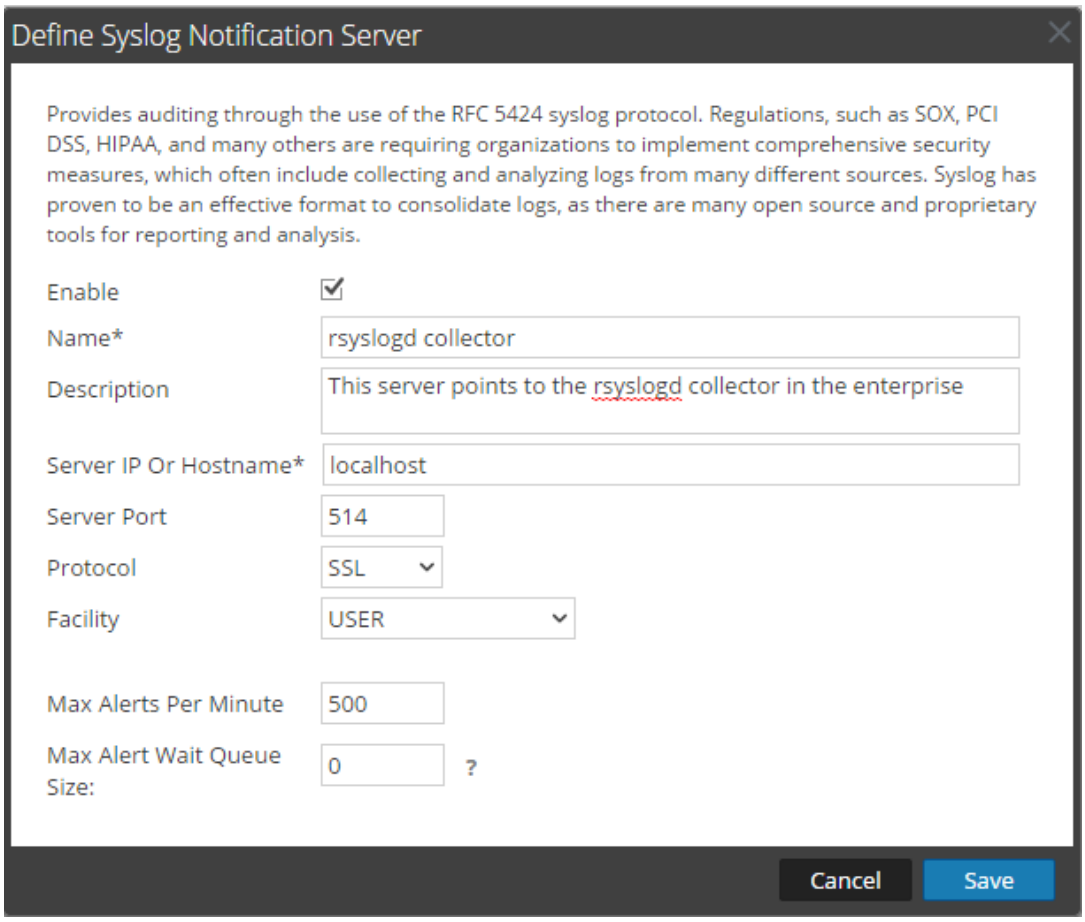
Parameters	Description
	Unencrypted and Authenticated and Encrypted) Privacy Protocol (This option is available only for security level Authenticated and Encrypted) — Privacy protocol is an encryption technique for data communication. Private Key (This option is avaliable only for security level Authenticated and Encrypted) — A password that you want to use for encryption.
Community	Community string used to authenticate on the SNMP trap host. The default value is **public**.
Number of Retries	Number of retries for the trap.
Max Alerts Per Minute	Maximum number of alerts per minute.
Max Alert Wait Queue Size	Maximum number of alerts to be queued before they are dropped.

Syslog

Syslog notification servers allow you to configure Syslog settings as a notification server to send notifications. When enabled, Syslog provides auditing through the use of the RFC 5424 Syslog protocol. Syslog has proven to be an effective format to consolidate logs, as there are many open source and proprietary tools for reporting and analysis.

You cannot disable notification servers associated with global audit logging configurations.

The following figure shows the Define Syslog Notification Server dialog.



The following table lists the various parameters that you need to define for the Syslog notification servers.

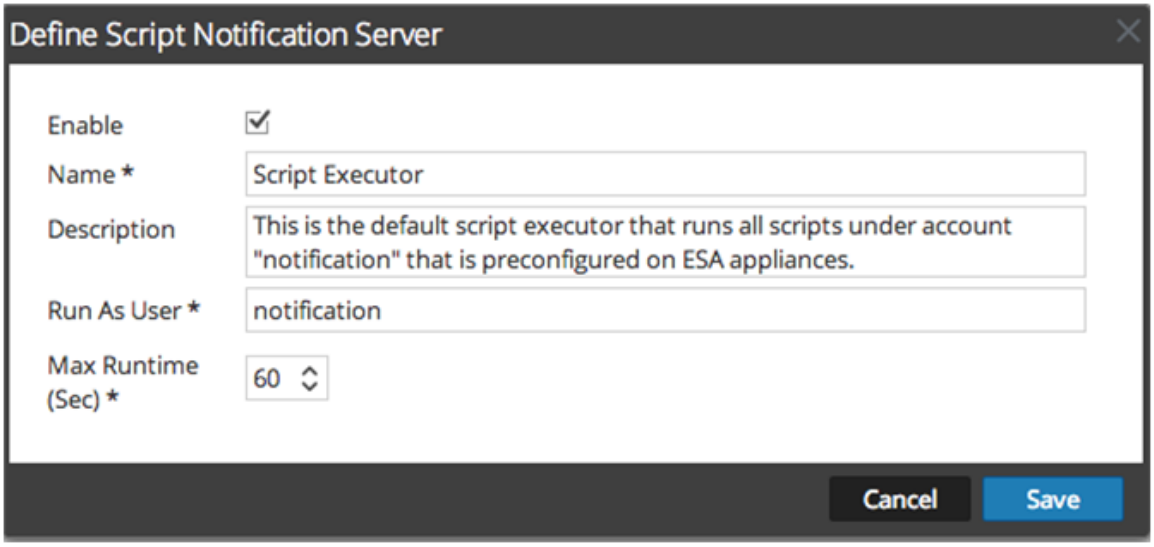
Parameters	Description
Enable	Select to enable the notification server.
Name	A name to identify or label the notification server.
Description	A brief description about the notification server.
Server IP Or Hostname	The hostname of the host where the target Syslog process is running.
Server Port	The port number where the target Syslog process is listening.
Protocol	The protocol to be used to transfer the Syslog files.
Facility	The designated Syslog facility to use for all outgoing messages. It is used to specify what type of program is logging the message. Some possible values are KERN, USER, MAIL, and DAEMON. This lets the configuration file specify that messages from different facilities will be handled differently.

Parameters	Description
Max Alerts Per Minute	Maximum number of alerts per minute. This field is not used for Global Audit Logging.
Max Alert Wait Queue Size	Maximum number of alerts to be queued before they are dropped. This field is not used for Global Audit Logging.

Script

Script notification servers enable you to configure Script as a Notification Server.

The following figure shows the Define Script Notification Server dialog.



The following table lists the various parameters that you need to define for the Script notification servers.

Parameters	Description
Enable	Select to enable the notification server.
Name	A name to identify or label the notification server.
Description	A brief description about the notification server.
Run As User	Name of the user identity under which the script is executed. The default user identity is **notification**. For ESA, you cannot set this to anything else unless you have created the account on the ESA host.
Max Runtime (Sec)	The maximum time (in seconds) the script is allowed to run.

Define Notification Output Dialogs

This topic provides descriptions of the various notification output dialogs. You configure notification outputs in the (Admin) > System > Global Notifications > Output tab. Notifications are basically the destinations used for sending notifications. For ESA, notifications enable you to define how you want to receive the ESA alerts. The following are the different notifications supported by NetWitness:

- Email
- SNMP
- Syslog
- Script

Procedures related to notifications are described in Configure Notification Outputs.

To access the Define Notification dialogs:

1. Go to **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. On the **Output** tab, click and then select a notification output (Email, SNMP, Syslog, or Script) The Define Notification dialog is displayed for your selection.

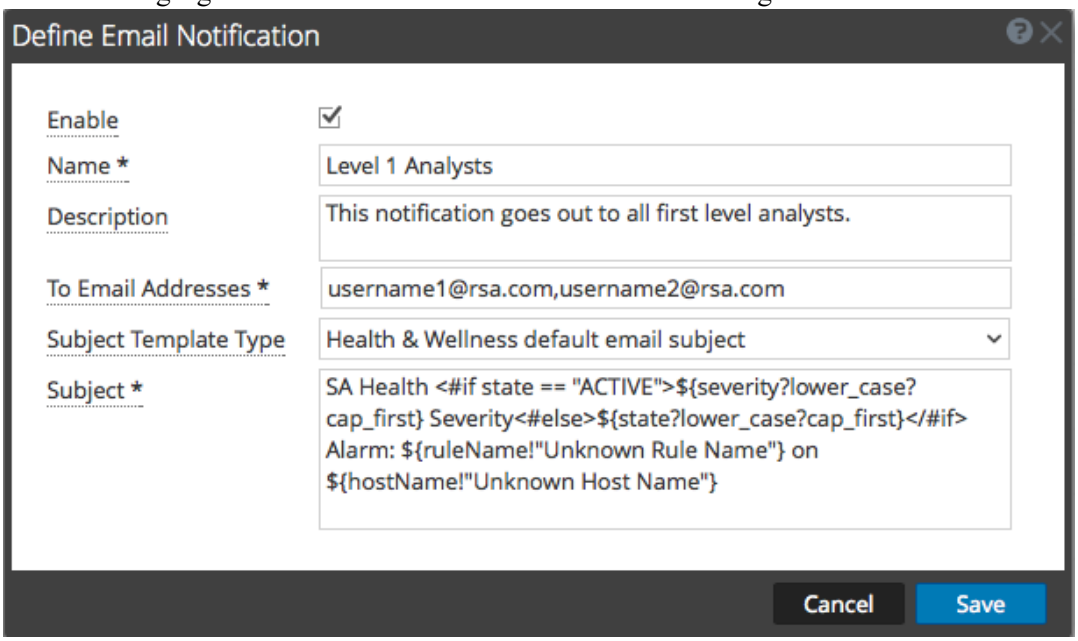
Features

There are four notification dialogs, which allow you to configure notification outputs.

Email

Email notifications enable you to define the destination email address to which you can send the alerts. It also enables you to add a custom description in the subject of the email and also to define multiple destination email addresses.

The following figure shows the Define Email Notification dialog.



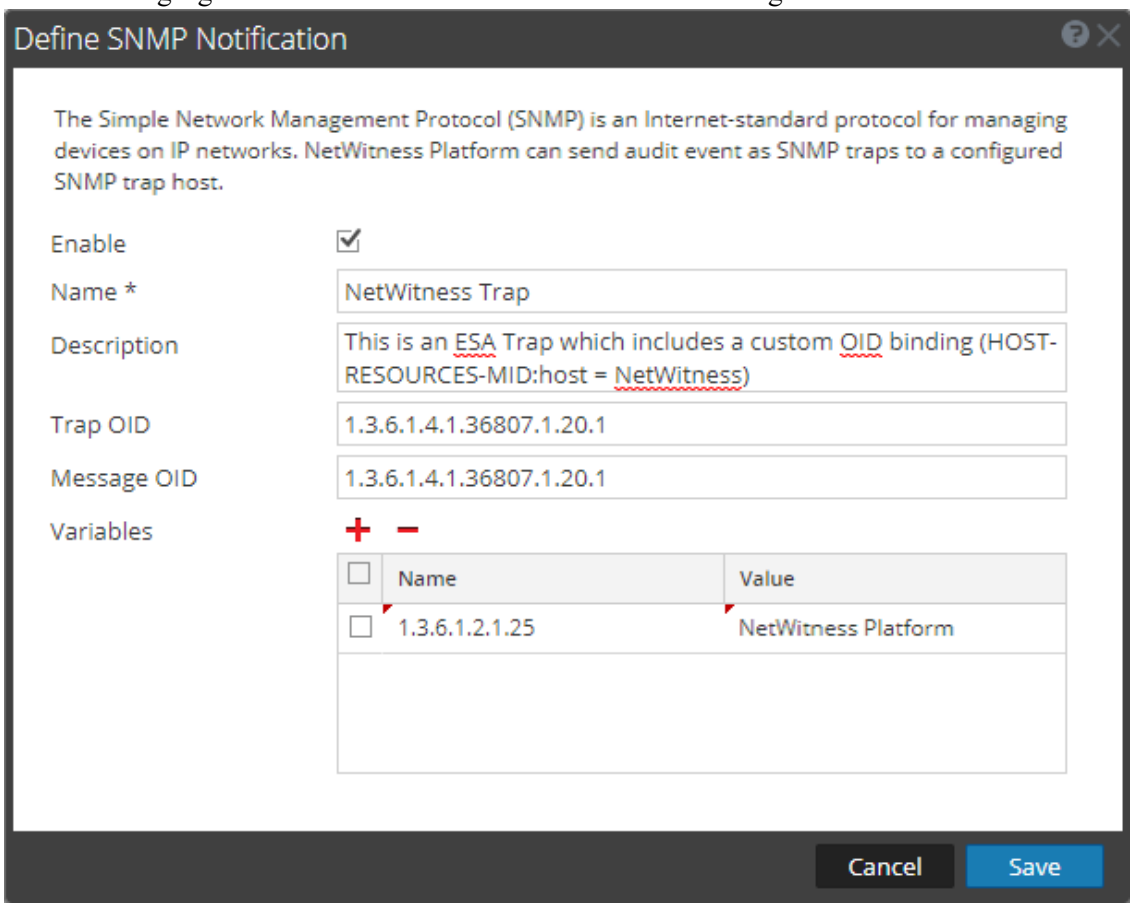
The following table lists the various parameters that you need to define for the email notifications.

Parameter	Description
Enable	Select to enable the notification.
Name	A name to identify or label the notification.
Description	A brief description about the notification.
To Email Addresses	Describes the destination email address to which the alert needs to be sent. **Note:** You can define multiple email addresses.
Subject Template Type	Lists available templates for creating a subject. When you choose a template, the Subject field is automatically filled in with the code for your chosen template. Example, for New Health and Wellness, you must select **New Health & Wellness default email subject**.
Subject	Custom description about the triggered alert. This information is automatically filled in if you choose one of the predefined templates from the Subject Template Type drop-down menu. **Note:** To provide a custom subject, please refer to "Include the Default Email Subject Line" topic in the *System Maintenance Guide*.

SNMP

SNMP notifications enable you to define the SNMP settings to send alert notifications.

The following figure shows the Define SNMP Notification dialog.



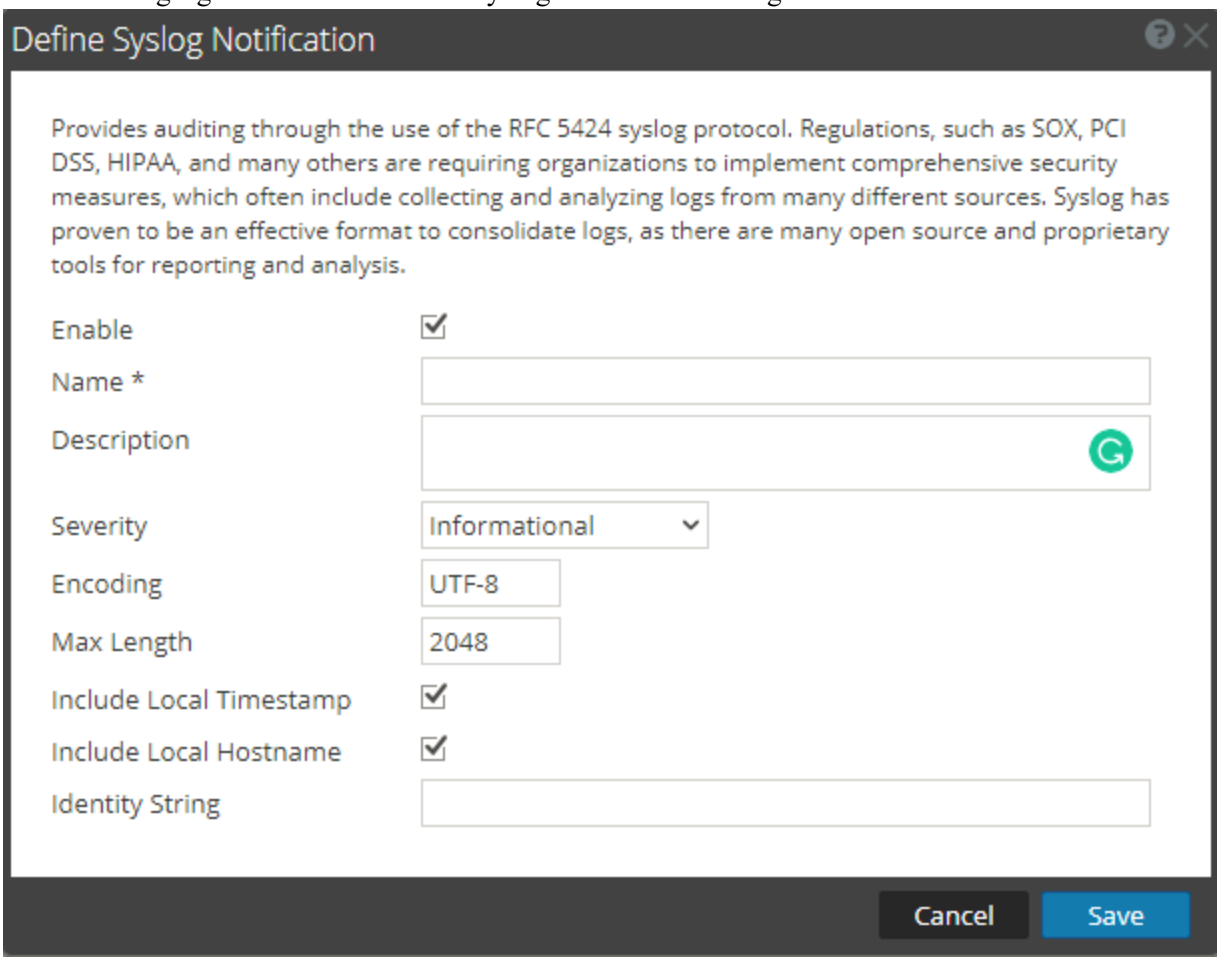
The following table lists the various parameters that you need to define for the SNMP notifications.

Parameter	Description
Enable	Select to enable the notification.
Name	A name to identify or label the notification.
Description	A brief description about the notification.
Trap OID	The object ID for the SNMP trap on the trap host that receives the event. The default value is **1.3.6.1.4.1.36807.1.20.1**. This value is a hierarchical name that represents the system that generates the trap. 1.3.6.1.4.1 is the common prefix for all enterprises and 36807.1.20.1 identifies NetWitness.
Message OID	The message object identifier for the SNMP trap.
Variables	Additional information that should be included within the trap. It is a variable that is a name value pair.

Syslog

Syslog notifications enable you to define the Syslog settings to send alert notifications.

The following figure shows the Define Syslog Notification dialog.



The following table lists the various parameters that you need to define for the Syslog notifications.

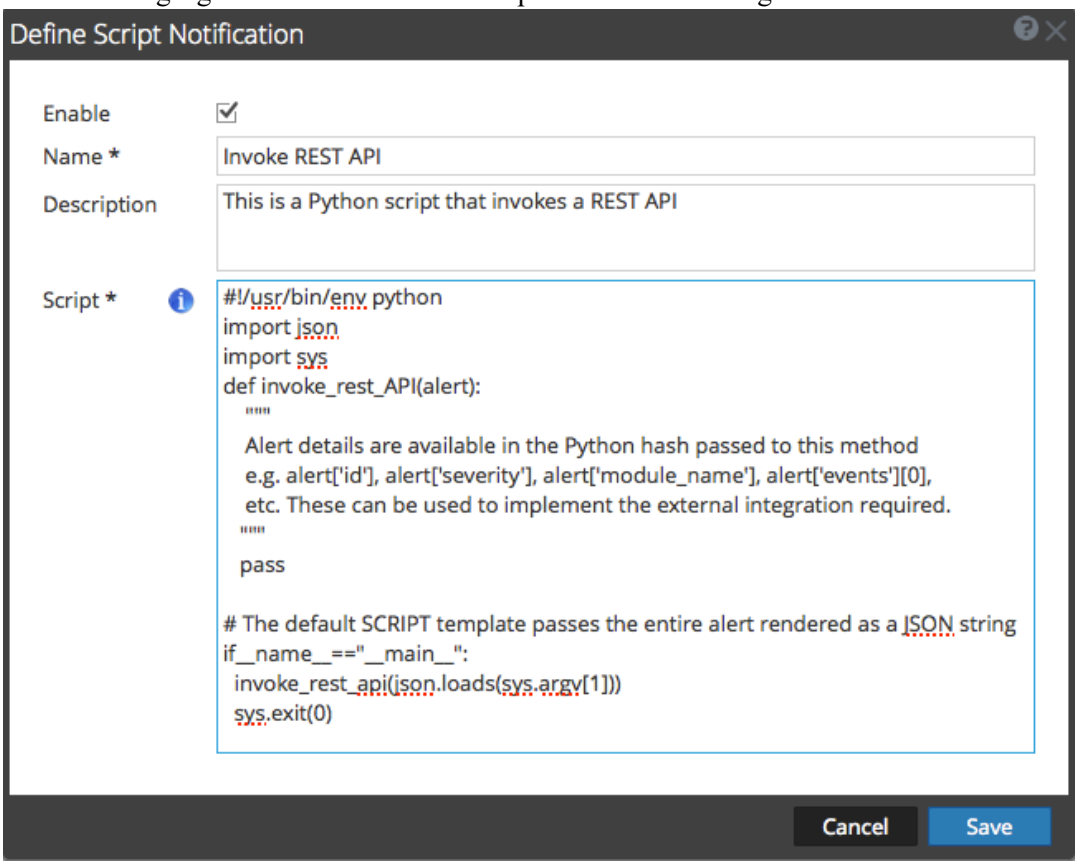
Parameter	Description
Enable	Select to enable the notification.
Name	A name to identify or label the notification.
Description	A brief description about the notification.
Severity	Defines the severity of the alert.
Encoding	Defines the encoding format. In some environments where no regular character sets are used (for example, Japanese characters), this field will help selecting the right encoding of the characters.
Max Length	The maximum length of a Syslog message in bytes. The default value is **2048**. Messages that exceed the maximum length are truncated when the **Truncate overly large syslog messages** checkbox is selected, which is found in Administration > System > Legacy Notifications. Legacy Notifications Configuration Panel provides additional information.

Parameter	Description
Include Local Timestamp	Select to include the local timestamp in messages.
Include Local Hostname	Select to include the local hostname in Syslog messages.
Identity String	An identity string to be prefixed to each Syslog alert. If the string is blank, no identity string is prefixed to the outgoing Syslog alerts. You can use this to identify the alerts from ESA.

Script

Script notifications enable you to define the Script that executes in response to the alert. You can use any script for ESA notifications.

The following figure shows the Define Script Notification dialog.



The following table lists the various parameters that you need to define for the Script notifications.

Parameter	Description
Enable	Select to enable the notification.
Name	A name to identify or label the notification.

Parameter	Description
Description	A brief description about the notification.
Script	Defines the script.

Define Notification Template Dialog

In the Global Notifications panel, you can configure global notification settings for Notification Servers, Notification Outputs, and Notification Templates. On the Templates tab, you configure the templates for various notifications. The notification template defines the format and message fields of the notifications. You can select a default template or you can use the Define Template dialog to configure and edit templates.

You can select a default template and use it or modify a default template based on your requirement. You can also use one of the following template types and create a template:

- Audit Logging
- Event Stream Analysis
- Event Source Monitoring
- Health Alarms
- New Health & Wellness Alarms (Version 11.5 and later).

These notification templates are created in an HTML, FreeMarker (FTL) format, a combination of both HTML and FTL, or Common Event Format (CEF) format:

- Email (SMTP) notification output type is created in both HTML and FTL formats.
- SNMP and Script notification output types are created in FTL format.
- Syslog notification output type is created in CEF format.

You need to have a good understanding of HTML, FTL, and CEF formats to successfully configure your own notification template. Click on the respective links to understand the specific formats.

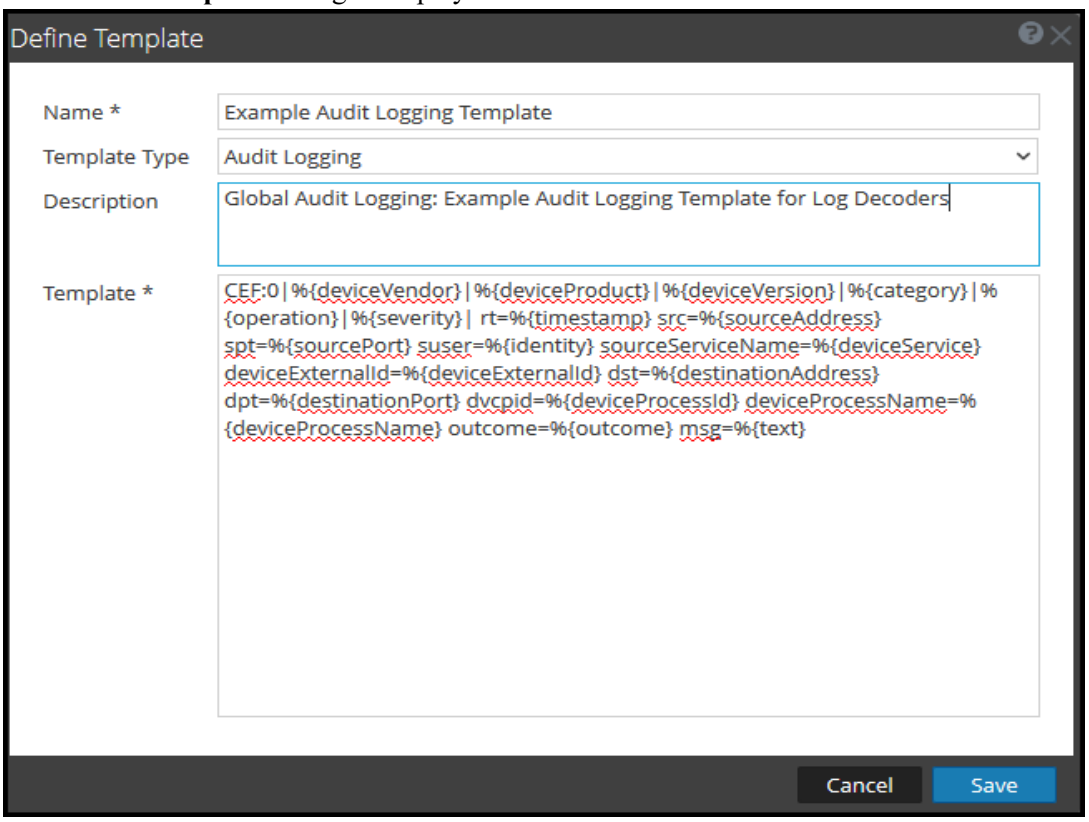
- HTML format, Introduction to HTML.
- FTL format, Overall structure - Apache FreeMarker Manual.
- CEF format, Freemarker Tips & Tricks in NetWitness.

Procedures related to notification templates are described in Configure Templates for Notifications.

To access the Define Template dialog:

1. Go to  **(Admin) > System**.
2. In the left navigation panel, select **Global Notifications > Template Tab**.
3. In the **Notifications Configurations** panel, click , or select a configuration and click  to modify.

The **Define Template** dialog is displayed.

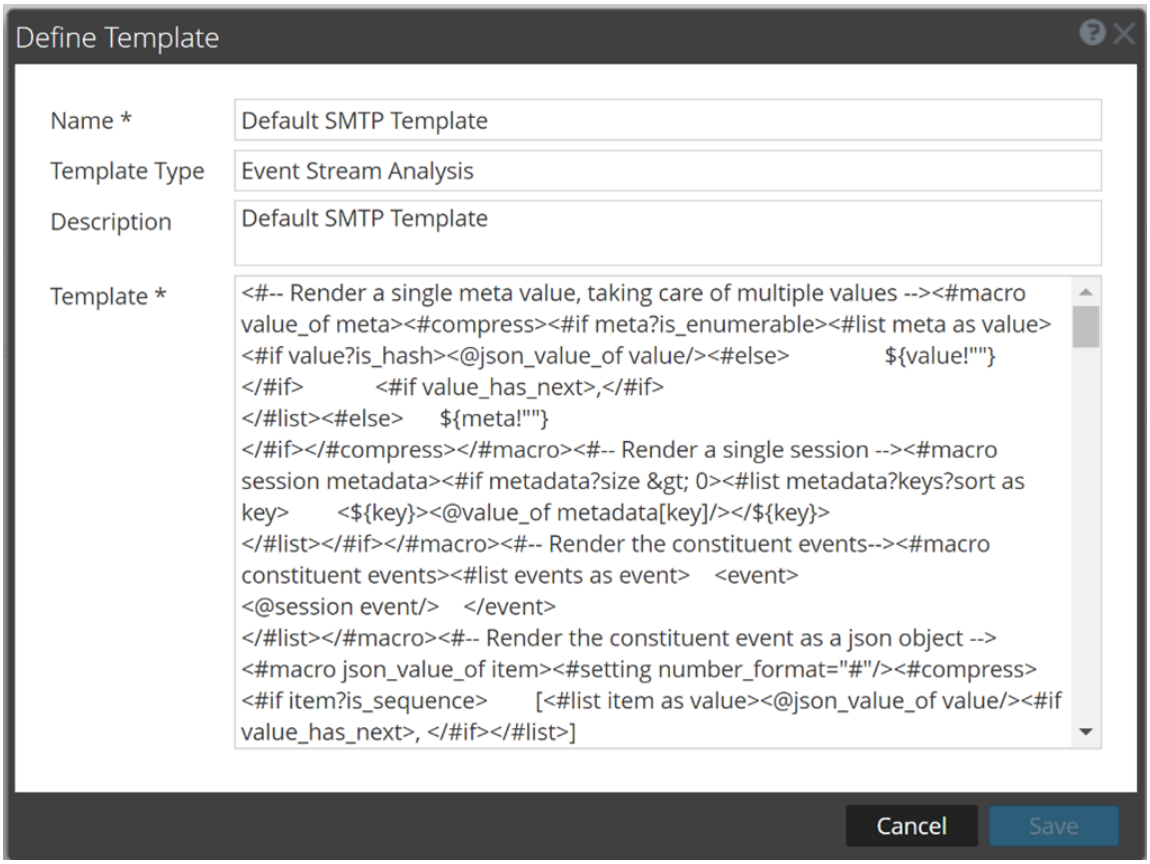


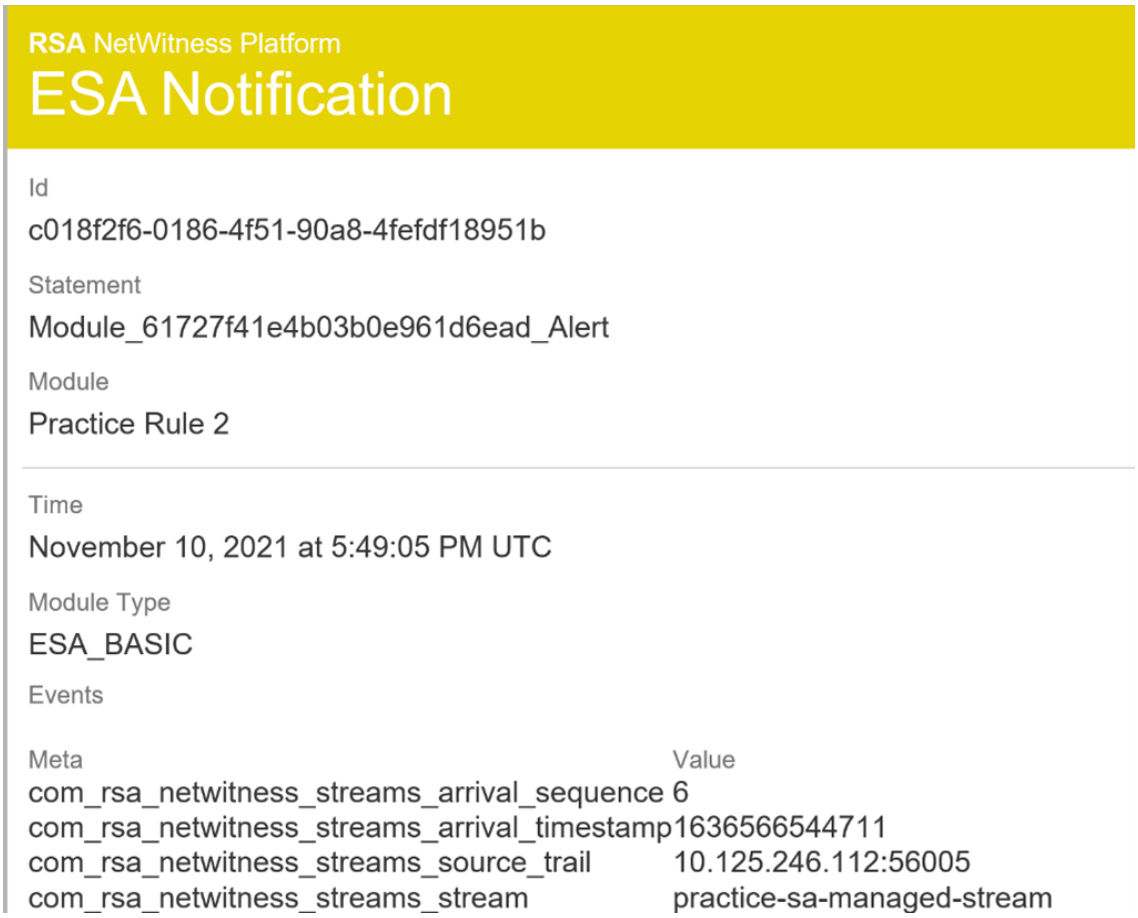
The following table describes the features in the Define Template dialog.

Field	Description
Name	Type a unique name for the notification template.
Template Type	Select the type of template that you want to create: • **Audit Logging**: Use this template for Global Audit Logging. • **Event Stream Analysis**: Use this template type for ESA alert notifications. • **Event Source Monitoring**: Use this template type for ESM notifications. • **Health Alarms**: Use this template type for Health and Wellness notifications. • **New Health and Wellness Alarms**: Use this template type for New Health and Wellness notifications.
Description	Add a description for the template. For example, if you create a notification template for Log Decoders to use for Global Audit Logging, you could mention that information in the description.

Field	Description
Template	Specify mandatory **CEF:** prefix when you create a template in CEF format. Define a Template for Global Audit Logging provides instructions on how to define an audit logging template to use for Global Audit Logging. To define a template for Event Stream Analysis (ESA), see Define a Template for ESA Alert Notifications. **Note:** Use Key references available in default notification templates.

Below is an example of a defined SMTP (email) template and output.





RSA NetWitness Platform

ESA Notification

Id

c018f2f6-0186-4f51-90a8-4fefdf18951b

Statement

Module_61727f41e4b03b0e961d6ead_Alert

Module

Practice Rule 2

Time

November 10, 2021 at 5:49:05 PM UTC

Module Type

ESA_BASIC

Events

Meta	Value
com_rsa_netwitness_streams_arrival_sequence	6
com_rsa_netwitness_streams_arrival_timestamp	1636566544711
com_rsa_netwitness_streams_source_trail	10.125.246.112:56005
com_rsa_netwitness_streams_stream	practice-sa-managed-stream

Output Tab

In the **Global Notifications** panel, in the **Output** tab ((Admin) > System > Global Notifications > Output), you configure notification outputs. Global Notifications configurations define notifications settings for Event Source Management (ESM), Health and Wellness, New Health and Wellness, Global Audit Logging, Event Stream Analysis (ESA), and Respond.

Notification Output configurations define email addresses and subject lines, SNMP trap OID settings, syslog output settings, and script code.

Notifications are the destinations configured for the alert notifications that are sent by ESA service. You can configure the following as destinations using the Output tab:

- Email
- SNMP
- Syslog
- Script

Note: You do not need to configure the Output tab for Global Audit Logging. For detailed steps, see Configure Global Audit Logging.

Workflow

This workflow shows the necessary procedures to configure and verify the output for Global Notifications. You can perform the following:

- Configure the Email settings as notification.
- Configure SNMP settings as notification.
- Configure Syslog settings as notification.
- Configure a Script as notification.



What do you want to do?

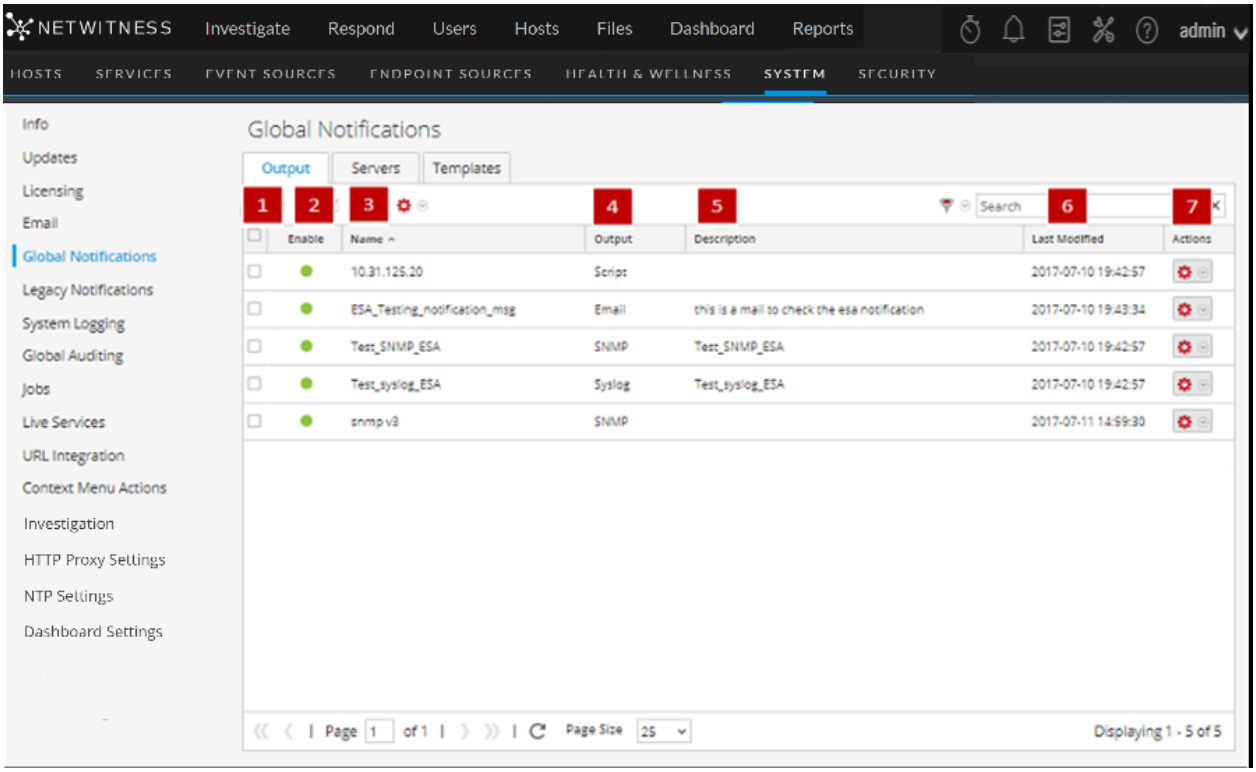
Role	I want to ...	Show me how
Administrator	Define notification outputs.	Configure Notification Outputs

Related Topics

- Notification Outputs Overview
- Configure Email as a Notification
- Configure Script as a Notification
- Configure SNMP as a Notification
- Configure Syslog as a Notification

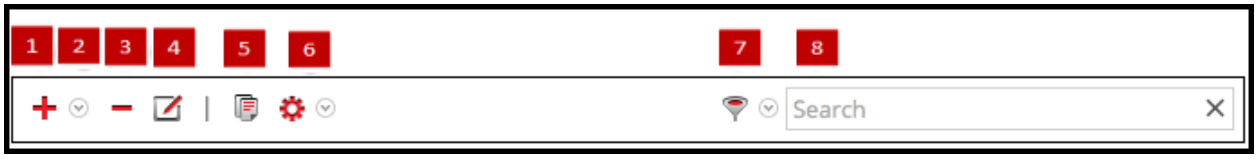
Quick Look

The following example illustrates Global Notification Outputs configuration.



1	Selects a row for an action in the toolbar. Selecting the check box in the column title selects or deselects all rows in the grid.
2	Indicates whether the configuration is enabled. A solid colored green circle indicates that a configuration is enabled. A blank white circle indicates that a configuration is not enabled.
3	Identifies or labels the configuration.
4	Identifies the configuration output. The outputs are Email, SNMP, Syslog, and Script.
5	Describes the configuration.
6	Shows the date and time of the last configuration change.
7	Provides an Actions menu for the selected configuration with actions that can be taken on the configuration. The Actions menu enables you to delete, edit, duplicate, and export the configuration.

The Global Notifications panel toolbar is at the top of the Output tag and provides the following options:



1. Adds a notification output
2. Configures Email, SNMP, Syslog, and Script notification settings.
3. Removes a selected notification configuration. You cannot delete notification servers and notification types that are associated with global audit log configurations. If you attempt to delete a notification output (notification) being used by alerts, you will receive a warning confirmation message that the alerts using the notification will not function properly. The message shows the number of alerts in use. You can also delete a configuration by selecting a configuration and then in the Actions column, selecting > Delete.
4. Edits a selected notification configuration. You can also edit a configuration by selecting a configuration and then in the Actions column, selecting > Edit.
5. Duplicates a selected notification configuration. You can also duplicate a configuration by selecting a configuration and then in the Actions column, selecting > Duplicate.
6. Displays the following options:
   - **Import**: Imports a notification server, type, or template. For example, on the Servers tab, you can import a notification server configuration.
   - **Export All**: Exports all of the configurations. For example, if you are on the Servers tab, you can export all of the notification server configurations.
   - **Export**: Exports a selected configuration. You can also export a configuration by selecting a configuration and then in the Actions column, selecting > Export.
7. Filters by Email, SNMP, Syslog, or Script.
8. Searches configurations in the grid.

Servers Tab

Servers Tab describes the components of the Global Notifications > Servers tab. This tab enables you to configure notification servers. Global Notifications configurations define notifications settings for Event Source Management (ESM), Health and Wellness, Global Audit Logging, Event Stream Analysis (ESA), and Respond.

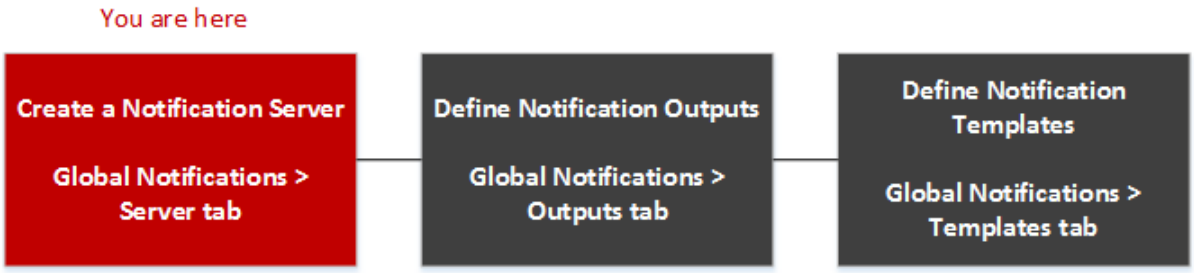
Configure **Notification Servers** in the Servers tab. On the Servers tab, add the servers from which you want to receive notifications from the system. For Global Audit Logging, define Log Decoders as Syslog Notification Servers.

Event Stream Analysis can send notifications to users through email, SNMP, or Syslog when an alert is triggered on the ESA service. These alert notification senders are called Notification Servers. You can configure multiple notification settings and use them while defining an ESA rule, for example, you can configure multiple mail servers or Syslog servers and use the settings while defining an ESA rule.

Workflow

The workflow shows the necessary procedures to configure and verify the Servers for Global Notifications. You can perform the following:

- Configure the Email settings as a notification server.
- Configure SNMP settings as a notification server.
- Configure Syslog settings as a notification server.
- Configure a Script as a notification server.



What do you want to do?

Role	I want to ...	Show me how
Administrator	Define notification Servers	Configure Notification Servers

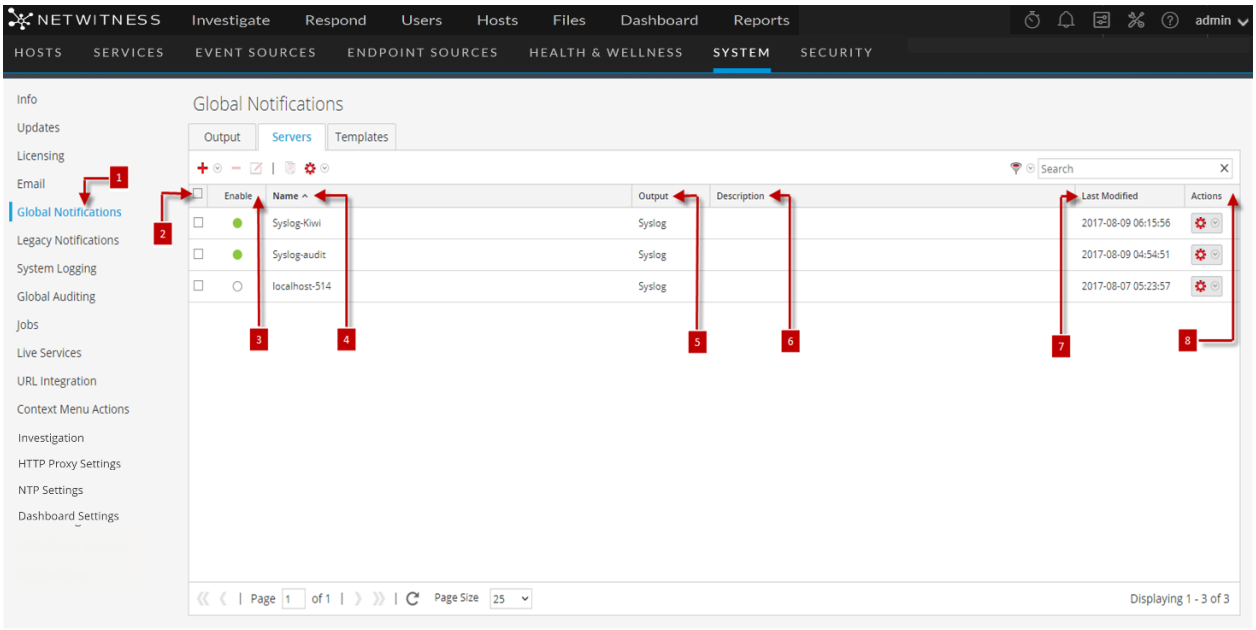
Related Topics

- Notification Servers Overview
- Configure the Email Settings as Notification Server

- Configure Script as a Notification Server
- Configure the SNMP Settings as Notification Server
- Configure a Syslog Notification Server

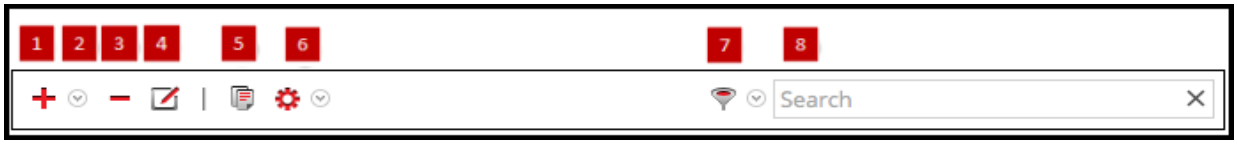
Quick Look

The following example illustrates Global Notification Servers configuration.



1	Displays the Server Tab Panel.
2	Selects a row for an action in the toolbar. Selecting the checkbox in the column title selects or deselects all rows in the grid.
3	Indicates whether the configuration is enabled. A solid colored green circle indicates that a configuration is enabled. A blank white circle indicates that a configuration is not enabled.
4	Identifies or labels the configuration.
5	Identifies the configuration output. The outputs are Email, SNMP, Syslog, and Script.
6	Describes the configuration.
7	Shows the date and time of the last configuration change.
8	Provides an Actions menu for the selected configuration with actions that can be taken on the configuration. The Actions menu enables you to delete, edit, duplicate, and export the configuration.

The Global Notifications panel toolbar is at the top of the Output tag and provides the following options:



1	Adds a notification output
2	Configures Email, SNMP, Syslog, and Script notification settings.
3	Removes a selected notification configuration. You cannot delete notification servers and notification types that are associated with global audit log configurations. If you attempt to delete a notification output (notification) being used by alerts, you will receive a warning confirmation message that the alerts using the notification will not function properly. The message shows the number of alerts in use. You can also delete a configuration by selecting a configuration and then in the Actions column, selecting  > Delete.
4	Edits a selected notification configuration. You can also edit a configuration by selecting a configuration and then in the Actions column, selecting  > Edit.
5	Duplicates a selected notification configuration. You can also duplicate a configuration by selecting a configuration and then in the Actions column, selecting  > Duplicate.
6	Displays the following options: • **Import**: Imports a notification server, type, or template. For example, on the Servers tab, you can import a notification server configuration. • **Export All**: Exports all of the configurations. For example, if you are on the Servers tab, you can export all of the notification server configurations. • **Export**: Exports a selected configuration. You can also export a configuration by selecting a configuration and then in the Actions column, selecting  > Export.
7	Filters by Email, SNMP, Syslog, or Script.
8	Searches configurations in the grid.

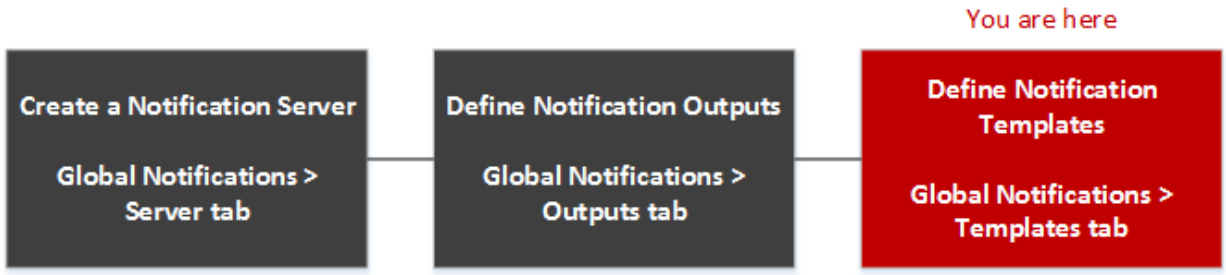
Templates Tab

The Notification Templates tab enables to configure notification templates. Global Notifications configurations define notifications settings for Event Source Management (ESM), Health and Wellness, New Health and Wellness, Global Audit Logging, Event Stream Analysis (ESA), and Respond. Notification templates define the format and message fields of the notifications.

Select a default template or configure templates for Email, SNMP, Syslog, and Script, depending on the template type. For Event Stream Analysis (ESA) templates, configure Email, SNMP, Syslog, and Script. For Audit Logging templates, configure Syslog.

Event Stream Analysis templates are not specific to any type of alert notifications, that is, the same template can be used for all types of notifications.

Workflow



What do you want to do?

Role	I want to ...	Show me how
Administrator	Define notification Templates	Configure Templates for Notifications

Related Topics

Configure Global Notifications Templates

Add a Template

Define a Template for ESA Alert Notifications

Delete a Template

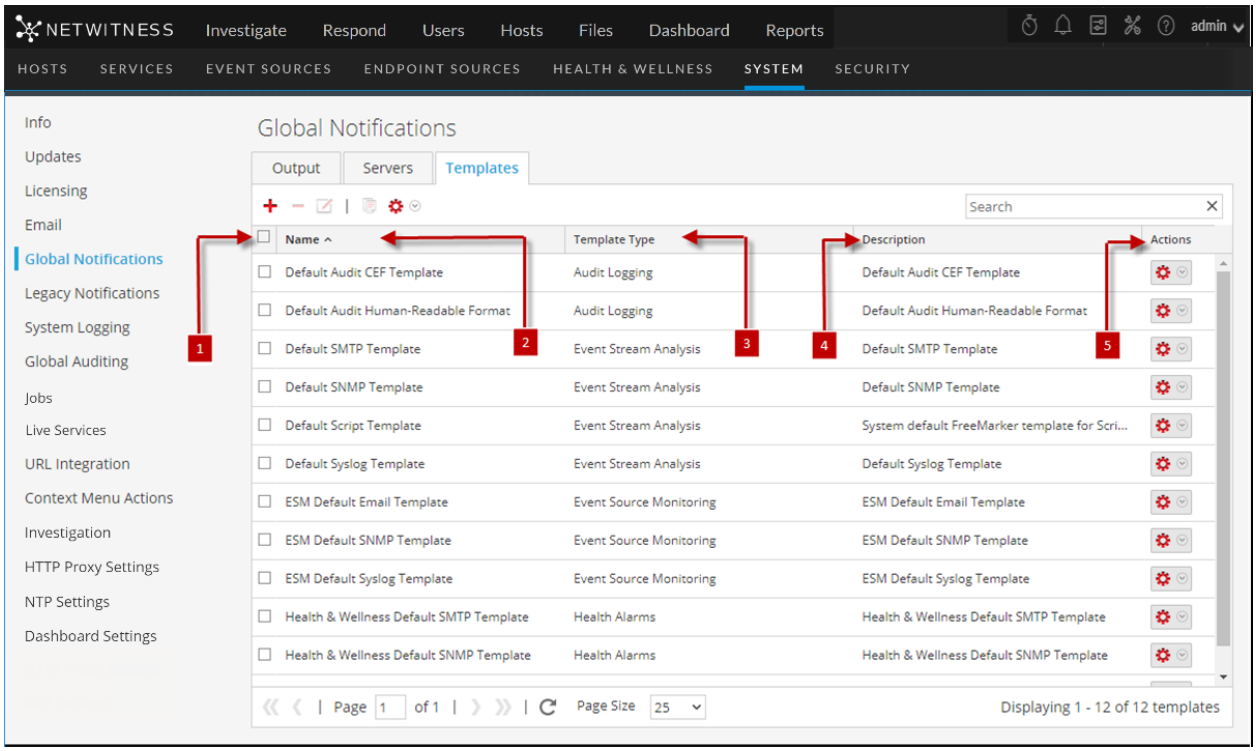
Duplicate a Template

Edit a Template

Import and Export a Global Notifications Template

Quick look

The following example illustrates Global Notification Templates Tab.



1. Selects a row for an action in the toolbar. Selecting the check box in the column title selects or deselects all rows in the grid.
2. Identifies or labels the templates
3. Choose a Template Type
4. Describes the templates
5. Provides an Actions menu  for the selected templates with actions that can be taken on the Templates. The Actions menu enables you to delete, edit, duplicate, and export the configuration.

HTTP Proxy Settings Panel

HTTP Proxy Settings Panel introduces the proxy support features of the (Admin) > **System** > **HTTP Proxy Settings** panel.

Note: Proxy support is only for HTTP and HTTPS proxies and not SOCKS5.

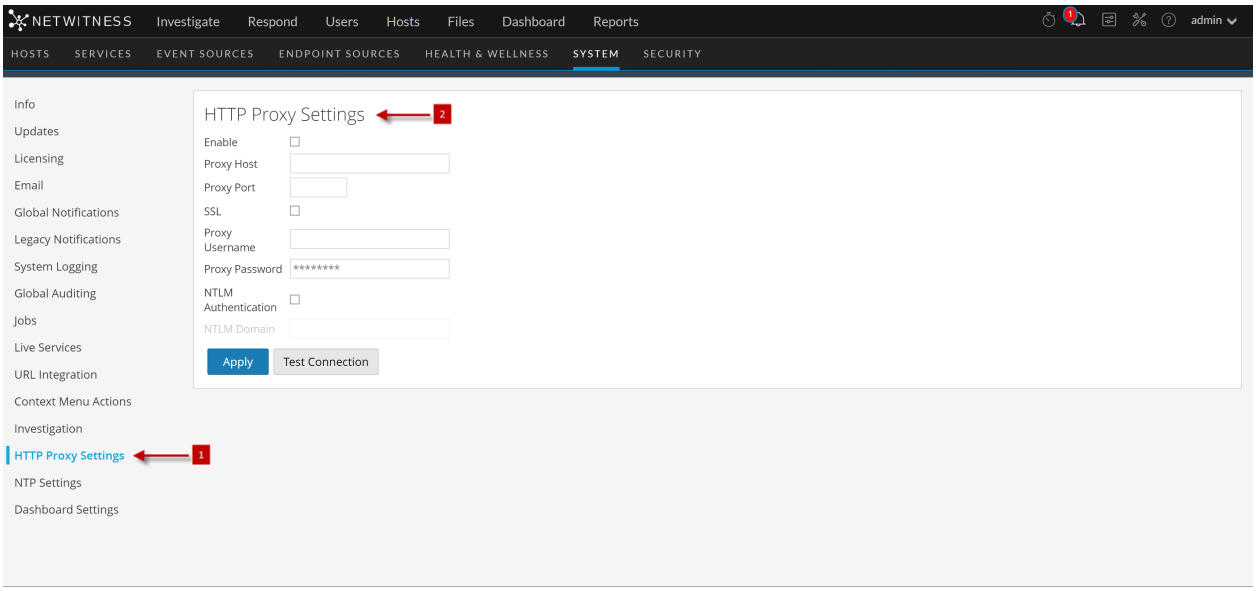
The HTTP Proxy Settings panel provides a user interface for configuring a proxy for use across NetWitness modules and services. The Proxy Settings set up a proxy to be used wherever a proxy is needed in NetWitness. The settings in this panel override any proxy settings configured for an individual service such as Malware Analysis or Live.

Related topics

"Configure Proxy for NetWitness Platform" in Additional Procedures

Quick Look

The following example illustrates an HTTP Proxy Settings Panel.



1 Displays the HTTP Proxy Settings Panel.
2 Allows the user to configure HTTP Proxy Sittings.

This table describes the features in the HTTP Proxy Settings section.

Feature	Description
Enable	Enable the system proxy configuration for use in NetWitness.
Proxy Host	The hostname for the proxy host.

Feature	Description
Proxy Port	The port used for communication on the proxy host.
SSL	(Optional) Enable communication using SSL.
Proxy Username	(Optional) The user name used to log on to the proxy host if the proxy requires authentication.
Proxy Password	(Optional) The user password used to log on to the proxy host if the proxy requires authentication.
NTLM Authentication	Use NT LAN Manager authentication and session security protocols.
NTLM Domain	The name of NTLM domain.
Apply	Applies any changes made, and they become effective immediately.

Email Configuration Panel

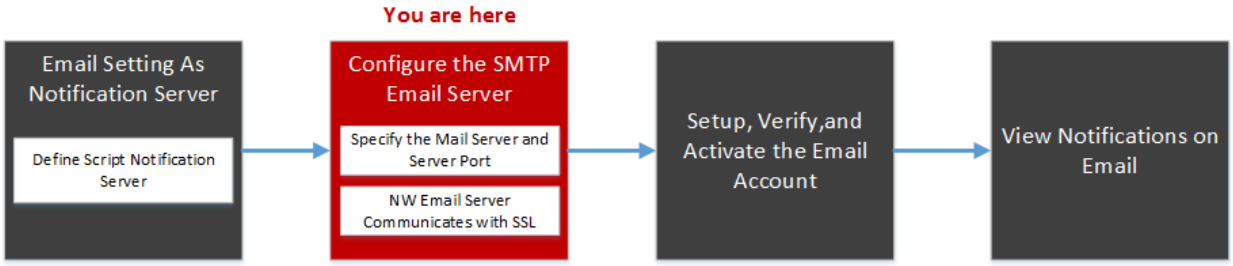
The Email Configuration Panel provides information about email configuration settings in the (Admin) > System > Email Configuration panel. NetWitness Platform sends notifications to users vwith email about various system events. To be able to configure these email notifications, first configure the SMTP email server (See Configure Email Servers and Notification Accounts).

The Email Configuration panel provides a way to:

- Configure the email server.
- Set up an email account to receive notifications.
- View statistics on email operations.

Workflow

This workflow shows the necessary procedures to configure and verify Email Panel.



What do you want to do?

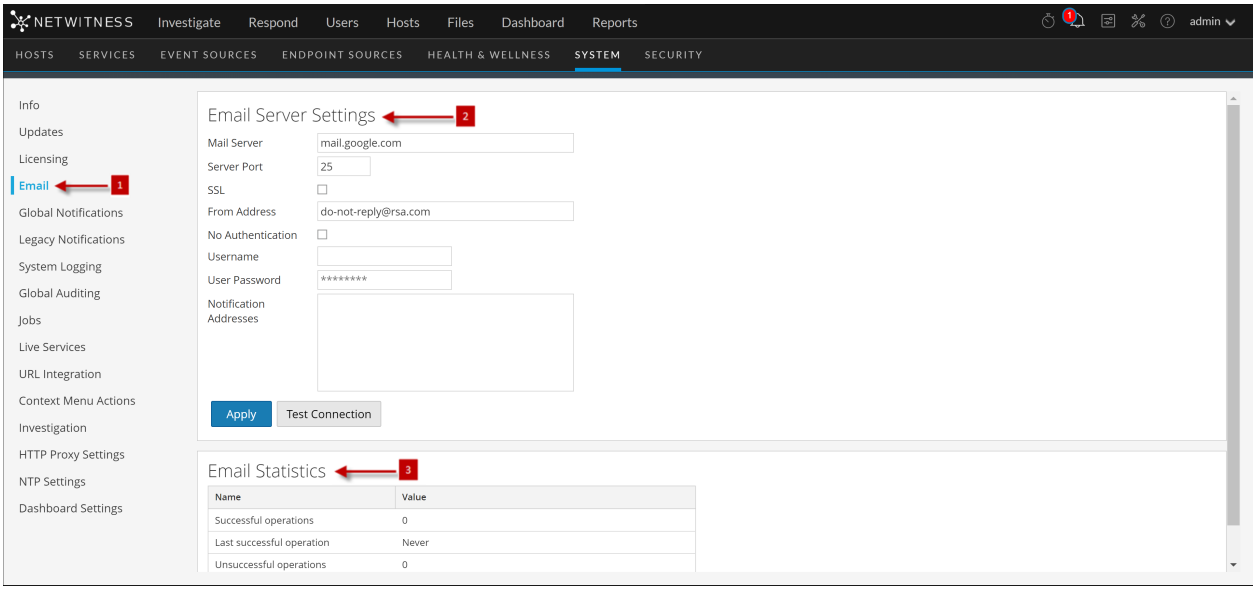
Role	I want to ...	Show me how
Administrator	Configure the SMTP Email Server	Configure Email Servers and Notification Accounts
Administrator	Email Setting as Notification Server	Configure the Email Settings as Notification Server
Administrator	Setup, Verify and Activate the Email Account	Receive Notification on Email

Related Topics

- Configure the Email Settings as Notification Server
- Configure Email as a Notification
- Configure Email Servers and Notification Accounts

Quick Look

The following example illustrates an Email configuration. The configuration defines how events are notified on Email.



1 Displays the Email Configuration Panel.
2 Allows the user to configure Email Server sittings.
3 Provides feedback on Email operations.

The **Email Configuration** panel has two sections: **Email Server Settings** and **Email Statistics**.

Email Server Settings

In the **Email Server Settings** section, you configure the following parameters.

Feature	Description
Mail server	The email server name. The default value is **mail.google.com**.
Server port	The server port used to send and receive emails. The default value is **25**.
Use SSL	The preference for SSL use in communications between the email server and NetWitness. The default value is to not use SSL (unchecked).
From address	The address that appears in all emails from NetWitness. The default from address for emails is **do-not-reply@rsa.com**.
Username	The username to access the email server. The default value is **blank**.
User password	The user password to access the email server. The default value is **blank**.

Feature	Description
Test connection	Tests the connection to the email server.
Apply	Applies the email configuration to this instance of NetWitness.

Email Statistics

The Email Statistics section provides feedback on the number of successful and failed email operations as well as the time of the last successful and unsuccessful email operation. For each statistic the name of the statistic and the value is displayed.

Investigation Configuration Panel

The (Admin) > System > Investigation Configuration panel provides the user interface for administrators to configure the system-wide settings that NetWitness Investigate uses when analyzing data and reconstructing an event.

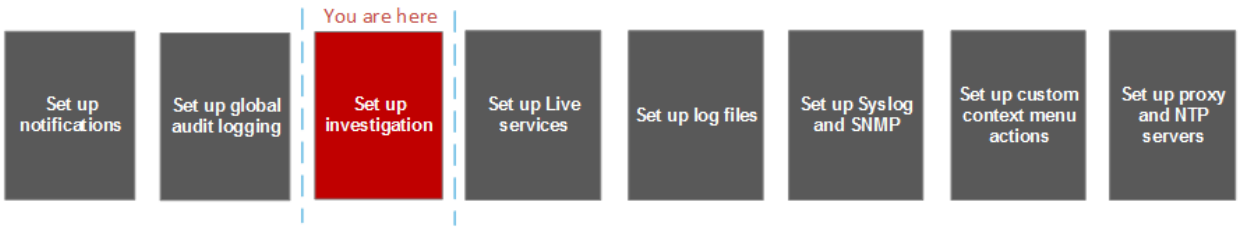
The settings allow an administrator to manage application performance for Investigate. As analysts analyze and reconstruct sessions that they are investigating, performance can be affected by operations that involve loading, searching, visualizing, and reconstructing large amounts of data.

Note: Analysts can also set individual preferences for Investigate in the Profiles view and in the Navigate, Legacy Events, and Events views.

To access the Investigation Configuration panel:

1. Go to **(Admin) > System.**
2. In the options panel, select **Investigation.**

Workflow



What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure Navigate, Legacy Events, and Events view settings	Configure Investigation Settings
Administrator	Map Context Hub Meta Types	Configure Investigation Settings
Administrator	Clear reconstruction cache for services	Configure Investigation Settings

Related Topics

- Standard Procedures

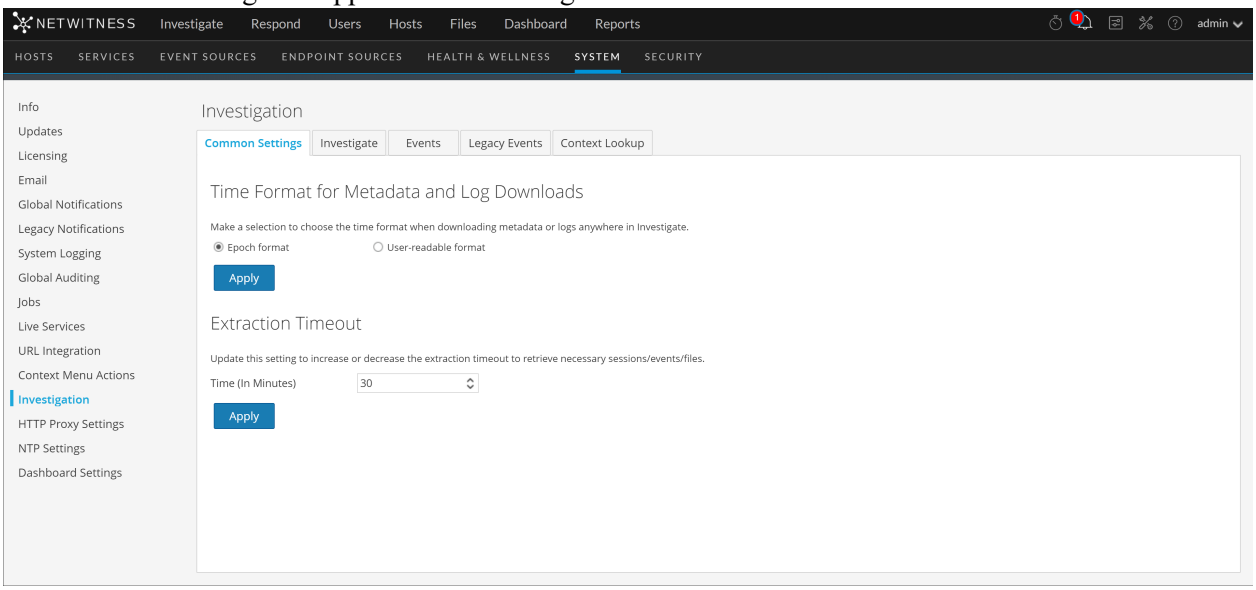
Quick Look

The Investigation Configuration panel has four tabs: Common Settings (Version 11.5 and later), Navigate, Events, Legacy Events, and Context Lookup.

Though most fields in the tabs have a selection list with specific increments through the range of possible values, you can enter a value within the allowed range manually. An invalid entry is signaled by the field highlighted in red. When valid values are selected, clicking Apply in a given section puts the changes into effect immediately.

Common Settings Tab

The Common Settings tab applies to all Investigate views.



The following table describes the options in this tab.

Parameter	Description
Time Format for Metadata and Log Downloads	
Epoch format	Use the Epoch format in downloads from Investigate (default value). These are Epoch representations of `2020-04-13 09:34AM`: • 12-hour representation = `61547519856000` • 24-hour representation = `61547519976000`

Parameter	Description

User-readable format	Use a user-readable time in downloads from Investigate. The downloads use a more understandable format that combines the user preference time zone, date format, and time format into an easily understood representation, which follows the industry standard ISO 8601 representation when possible.
	Examples of 12-Hour Format
	`04-13-2020T09:17:36AM-07:00`
	`13-04-2020T09:17:36AM-07:00`
	`2020-13-13T09:17:36AM-07:00`
	Examples of 24-Hour Format
	`04-13-2020T09:19:14-07:00`
	`13-04-2020T09:19:14-07:00`
	`2020-04-13T09:19:14-07:00`

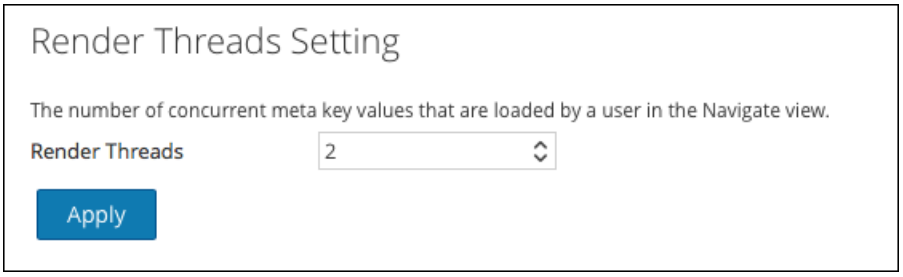
Extraction Timeout	
Time (In Minutes)	Use this to set the time limit before the session expires when downloading the logs from the service.

Investigate Tab

The Investigate tab has two sections: Render Threads Setting and Parallel Coordinates Settings. The following figure shows the Navigate tab.

Render Threads Setting

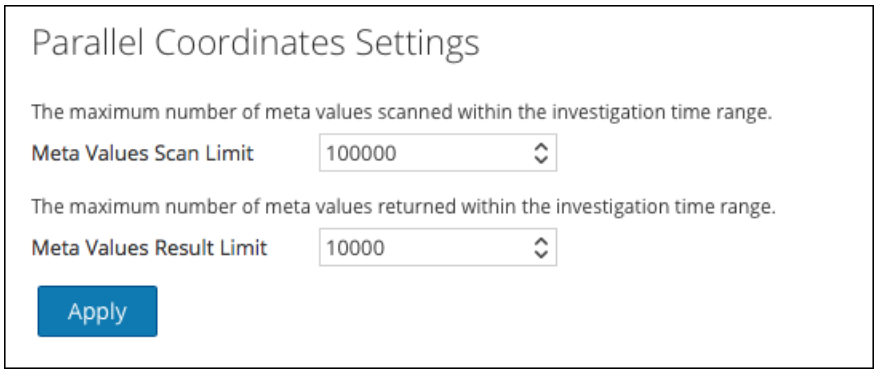
The Render Threads Setting is a selectable value between 1 and 20, which defines the number of concurrent (Values) loads in the Navigate view. The default value is 1.



Parallel Coordinates Settings

The Parallel Coordinates Settings apply to the Parallel Coordinates visualization in the Navigate view. There is a fixed limit on the amount of data that can be rendered as a parallel coordinates chart. In NetWitness the administrator can configure parallel coordinates limits here.

Note: For better performance, recommended settings are **Meta Values Scan Limit: 100000** and **Meta Values Result Limit: 1000-10000**.

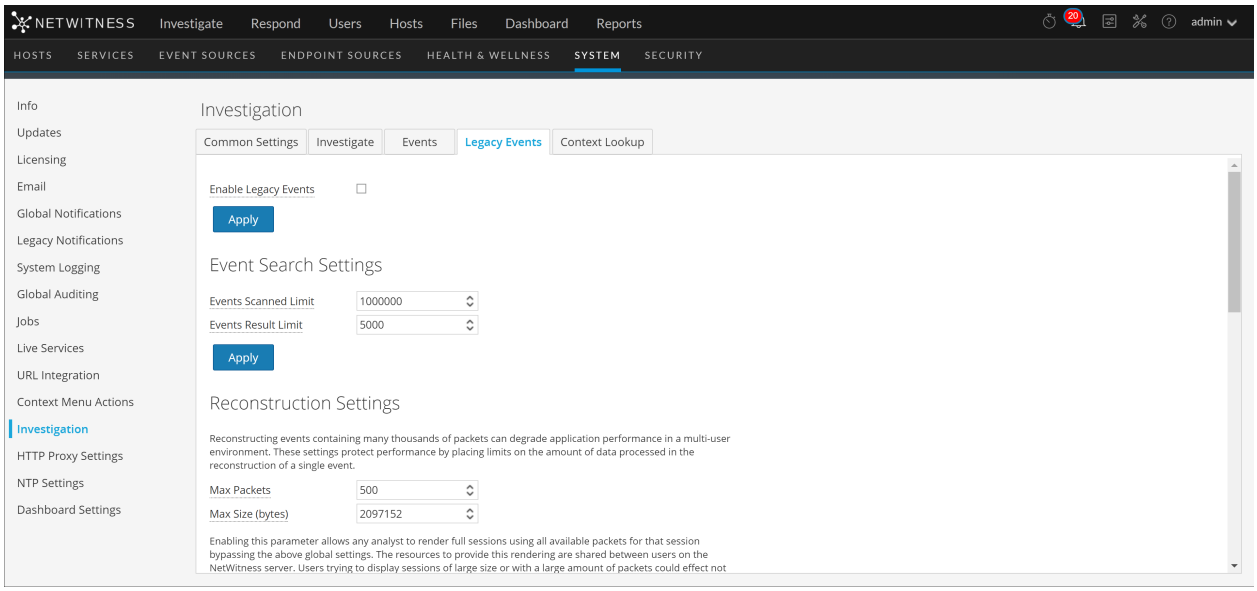


The following table describes the Parallel Coordinates Settings.

Parameter	Description
Meta Values Scan Limit	The maximum number of meta values scanned within the time range the analyst has selected in the Navigate view. Possible values are in the range of 1,000 to 10,000,000. The default value is 100,000.
Meta Values Result Limit	The maximum number of meta values returned within the time range the analyst has selected in the Navigate view. Possible values are in the range of 100 to 1,000,000,000. The default value is 10,000.

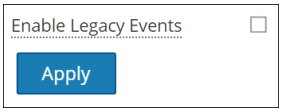
Legacy Events Tab

The Events tab provides configurable settings that affect the investigation of events. This tab has five sections: Enable Legacy Events, Event Search Settings, Reconstruction Settings, Web View Reconstruction Settings, and Reconstruction Cache Settings. The following figure shows the Events tab.



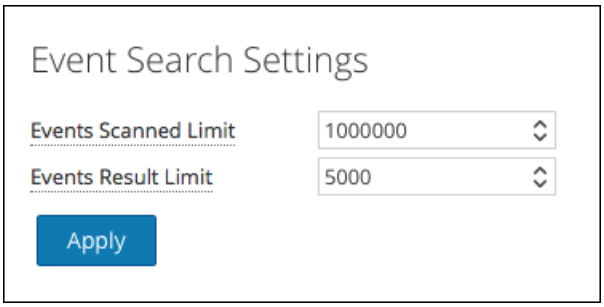
Enable Legacy Events

The Enable Legacy Events checkbox helps to enable and view the legacy events tab and view the classic events page on the UI. By default, this option is disabled.



Event Search Settings

The Event Search Settings help to limit the number of events scanned when searching in the Events view.



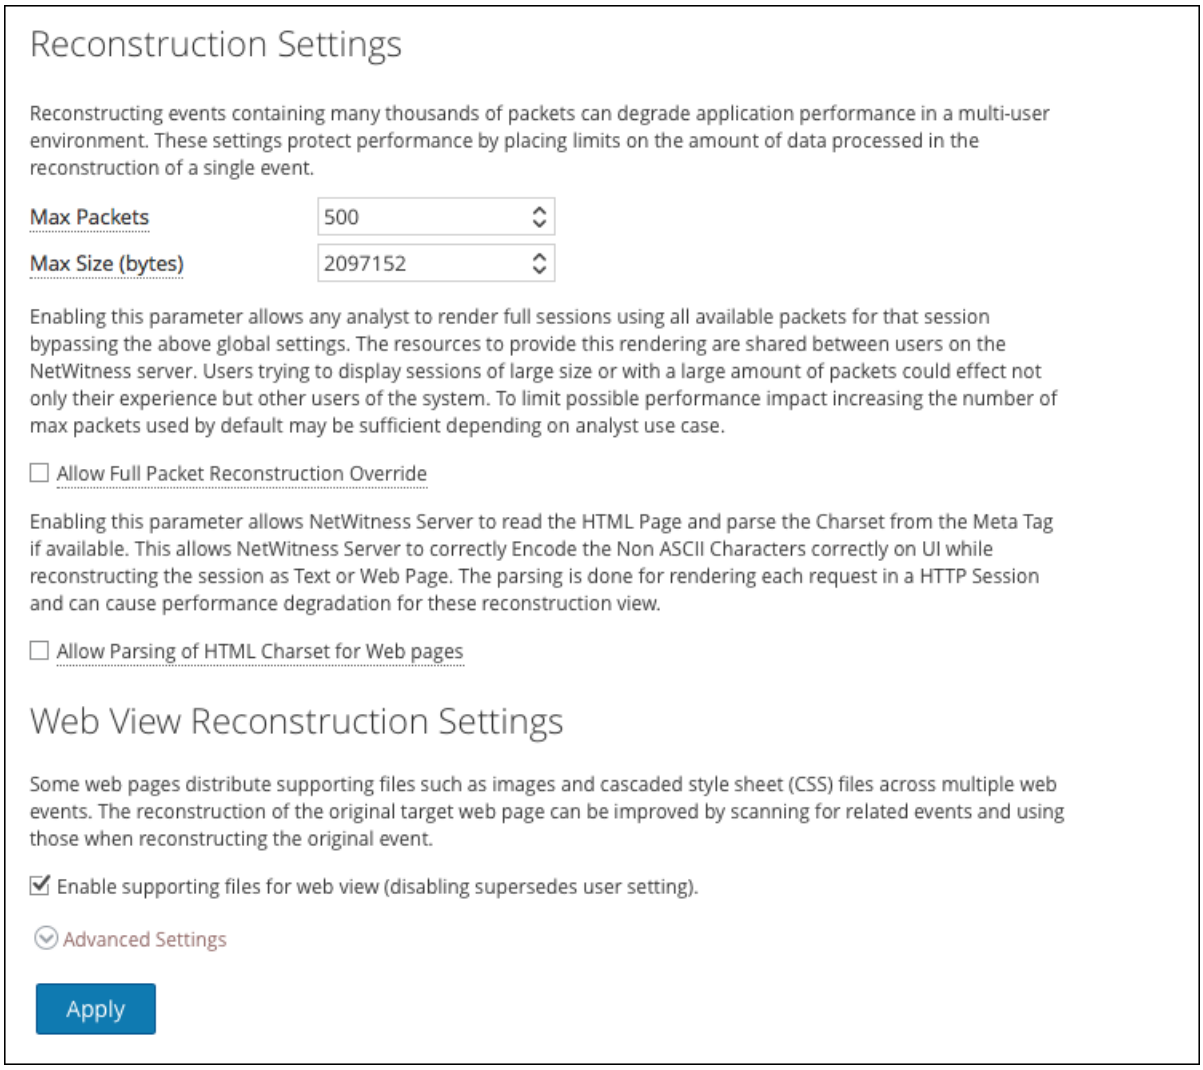
The following table describes the Event Search Settings.

Parameter	Description
Events Scanned Limit	The maximum number of events to scan when searching in the Events view. The actual number of events scanned may be slightly greater than the limit set here.
Events Result Limit	The maximum number of results to return when searching in the Events view. The actual number of results returned may be slightly greater than the limit here.

Reconstruction Settings

As analysts reconstruct sessions that they are investigating, some events can be very large and contain many thousands of source packets. Reconstructing these sessions, especially in a multi-user environment, can degrade application performance. The Reconstruction Settings allow an administrator to limit the number of packets and the size of a single event during reconstruction.

Note: An override to the Reconstruction Settings section is configurable for web views (under Web View Reconstruction Settings).



The following table describes the Reconstruction Settings features.

Parameter	Description
Maximum number of packets for a single event	This setting protects performance by placing a limit on the number of packets processed for a single event reconstruction. Possible values are in the range from 100 to 10,000 packets, using manual entry or increments of 100 from the selection list. The default value is 100 packets.
Maximum size, in bytes of a single event	This setting protects performance by placing a limit on the maximum size, in bytes, of a single event reconstruction. Possible values are in the range from 102,400 to 104,857,600 bytes, using manual entry or increments of 10,240 from the selection list. The default value is 2,097,152 bytes.

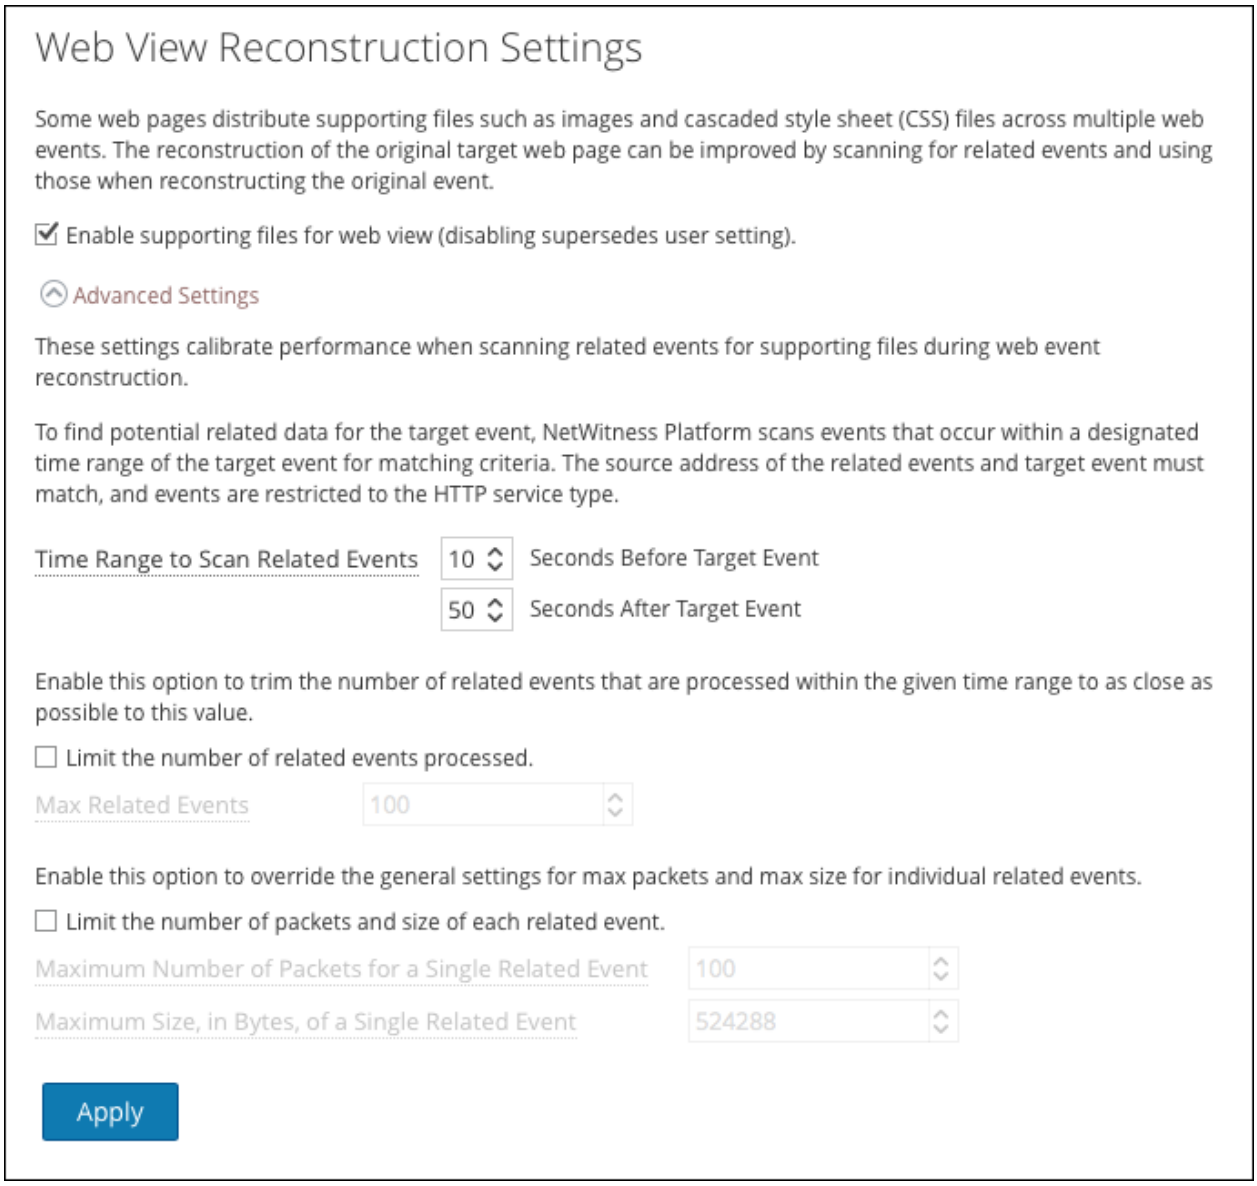
Parameter	Description
Allow Full Packet Reconstruction Override	When this checkbox is selected, the analysts is provided with a Use More Packets button in the Reconstruction Panel. This enables the NW Server to regenerate events using all the packets available in the Event.
Allow Parsing of HTML Charset for Web pages	This option allows the NetWitness Server to identify the web page encoding defined in the HTML meta tag instead of the HTTP header. The default setting is disabled.

Web View Reconstruction Settings

The Web View Reconstruction Settings allow an administrator to configure settings that improve the reconstruction of a web view by scanning and reconstructing related events that contain the same supporting files. When NetWitness is reconstructing a web view that spans multiple events, it is possible to improve the reconstruction of the target event by scanning and reconstructing related events that contain the same supporting files, such as images and cascaded style sheet (CSS) files.

- The only related events scanned are HTTP service type events with the same source address as the target event, and a time stamp within a specified time range before and after the target event.
- The maximum number of related events to scan is configurable.

Clicking on the Advanced Settings option displays all configurable settings in this section.



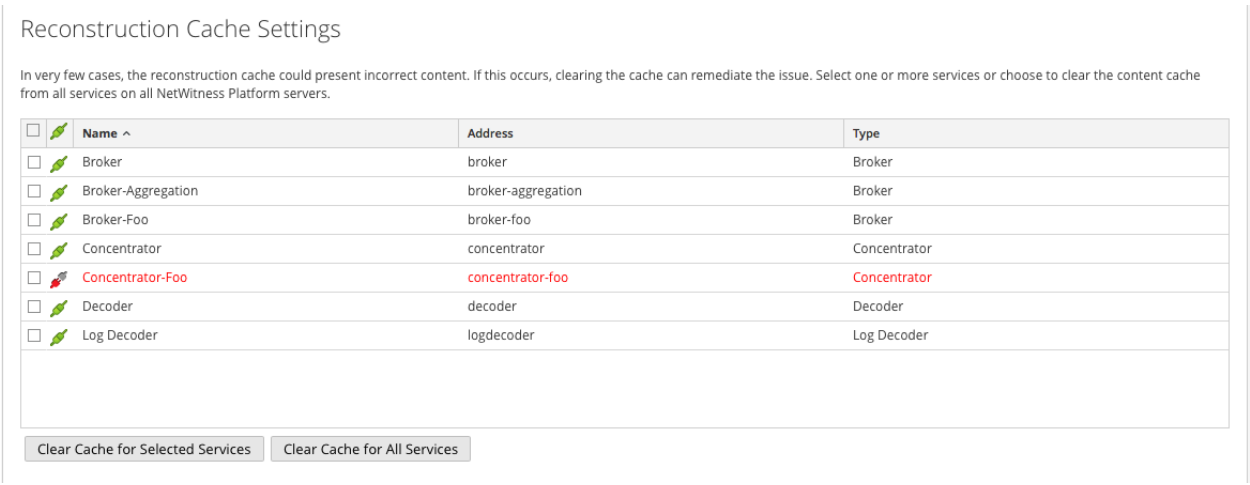
The following table describes the Web View Reconstruction Settings.

Parameter	Description
Enable supporting files for web view	This option determines how web views that have related data in other sessions are reconstructed. The default setting is enabled. When enabled, supporting files from related events can be used in the reconstruction of web views. Additional settings for calibrating the performance are enabled in this section, and Analysts have the option to enable CSS use in reconstructions. When disabled, supporting files from related events are not used and the setting for analysts to enable CSS use in reconstructions is disabled.

Parameter	Description
Time Range to Scan Related Events	Available when **Enable supporting files for web view** is checked. Configures the time range within which NetWitness scans related events that are of the service type HTTP and have the same source address as the target event. This is a value between 0 and 60. • Seconds Before Target Event • Seconds After Target Event
Limit the number of related events processed	Allows configuration of the maximum number of related events that NetWitness scans within the specified time range to discover supporting files for the target event. By default, this is disabled. When enabled, the Maximum Related Events field becomes active.
Max Related Events	When **Limit the number of events processed** is enabled, this field specifies the maximum number of related events that NetWitness scans within the specified time range to discover supporting files for the target event. This is a selectable value between 10 and 1,000, using an increment of 100. The default value is 100.
Limit the number of packets and size of each related event	Overrides the general settings for the maximum number of packets and maximum size (in bytes) for individual related events.
Maximum Number of Packets for a Single Related Event	Possible values are in the range from 100 to 10,000 packets, using increments of 100 from the selection list. The default value is 100 packets.
Maximum Size, in Bytes, of a Single Related Event	Possible values are in the range from 102,400 to 104,857,600 bytes, using increments of 10,240 from the selection list. The default value is 524,288 bytes.

Reconstruction Cache Settings

In some cases, the reconstruction cache can present incorrect content; for this reason NetWitness removes reconstructions that are older than a day from the cache. The cache is cleaned every day at midnight. Between the daily cache cleanings, certain actions may result in stale cache being used for a reconstruction, and if the need arises, administrators can manually clear cache for one or more services that are connected to the current NetWitness Server.

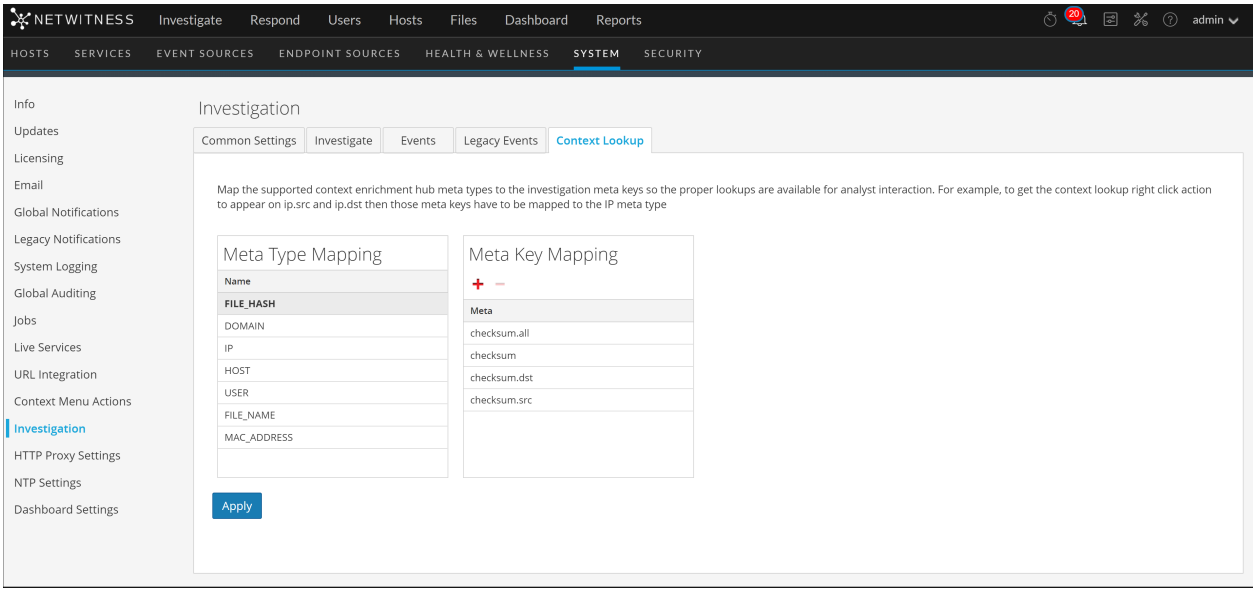


The following table describes the Reconstruction Cache Settings features.

Feature	Description
Selection box	Selection box in individual rows and in the title bar allow selection of one or more, or all services that need to have cache cleared manually.
Clear Cache for Selected Services	Clears the reconstruction cache for each selected service.
Clear Cache for All Services	Clears the reconstruction cache for all services.

Context Lookup Tab

Procedures associated with this panel are provided in "Manage Meta Type and Meta Key Mapping" in the *Context Hub Configuration Guide*. The following figure shows the Context Lookup tab.



The Context Lookup tab enables the administrator to configure the Investigate meta keys and meta type mapping. The administrator can add or remove meta keys found in Investigate to the list of meta types supported by Context Hub service. NetWitness Respond and Investigate use these default mappings for context lookup. For information about adding meta keys, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.

Caution: For the Context Lookup to work correctly in the Respond and Investigate views, this is the best practice: When mapping meta keys in the  (Admin) > SYSTEM > Investigation > Context Lookup tab, add only meta keys to the Meta Key Mappings. Do not add fields in the MongoDB. For example, `ip.address` is a meta key and `ip_address` is not a meta key (it is a field in the MongoDB).

The following table describes the features of the Context Lookup tab.

Feature	Description
	Adds a meta key to the selected meta type supported by Context Hub.
	Deletes the meta key from the selected meta type.
Apply	Saves the changes made to the Context Lookup tab.

Events Tab

The following figure shows the Event tab.

The Events tab provides configurable settings that affect the number of events displayed in the Events panel. This tab has two sections: Events Panel Settings and Event Limit per User Role.

Feature	Description
Event Limit Default	Specifies the maximum number of events loaded in the Events panel when a query is submitted. Possible values are integers between 100 and 40,000, and the default value is 5,000 events. If a query returns more events than the configured Event Limit Default, the Events panel title shows the analyst that more results are available but are not listed due to the limit. Increasing the limit may place additional load on the queried service; the ideal limit is determined by your environment.
Event Limit Per User Role	Specifies the maximum number of events loaded for a single query for individual user roles. This limit must be less than or equal to the system events limit of 40,000, but it can be greater than the Event Limit Default.
Apply	Each setting has an Apply button, which saves the change. The change becomes effective immediately, and applies to any new queries submitted by users.

Live Services Configuration Panel

Live Services Configuration Panel introduces the features for setting up your Live account and the CMS server connection.

Live Account consists of two sections, namely RSA Live Status and Download Live Feedback Activity Log. **Sign In** by entering your Live Account credentials to access the Live Services. To activate your Live account for NetWitness, please contact NetWitness Customer Support. When you have confirmation that your Live account has been set up, you can configure the CMS server connection as described in Configure Live Services Settings

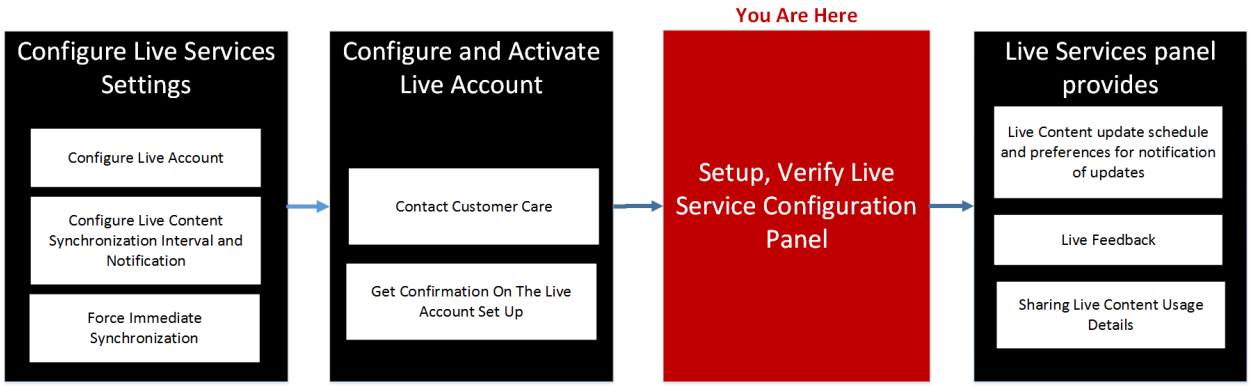
The Live Services panel provides the user interface for:

- The Live account
- The Live Content update schedule and preferences for notification of updates
- Participation in Live Feedback
- Sharing Live Content Usage Details

For information on Live Feedback, see Live Feedback Overview

For information on Analyst Behaviors and Data Sharing, see the "NetWitness Platform Feedback and Data Sharing topic in the *Live Services Management Guide*.

Workflow



What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure Live Account, CMS Server Connection	Configure the Email Settings as Notification Server
Administrator	Upload Data to NetWitness for Live Feedback	Upload Data to RSA for Live Feedback

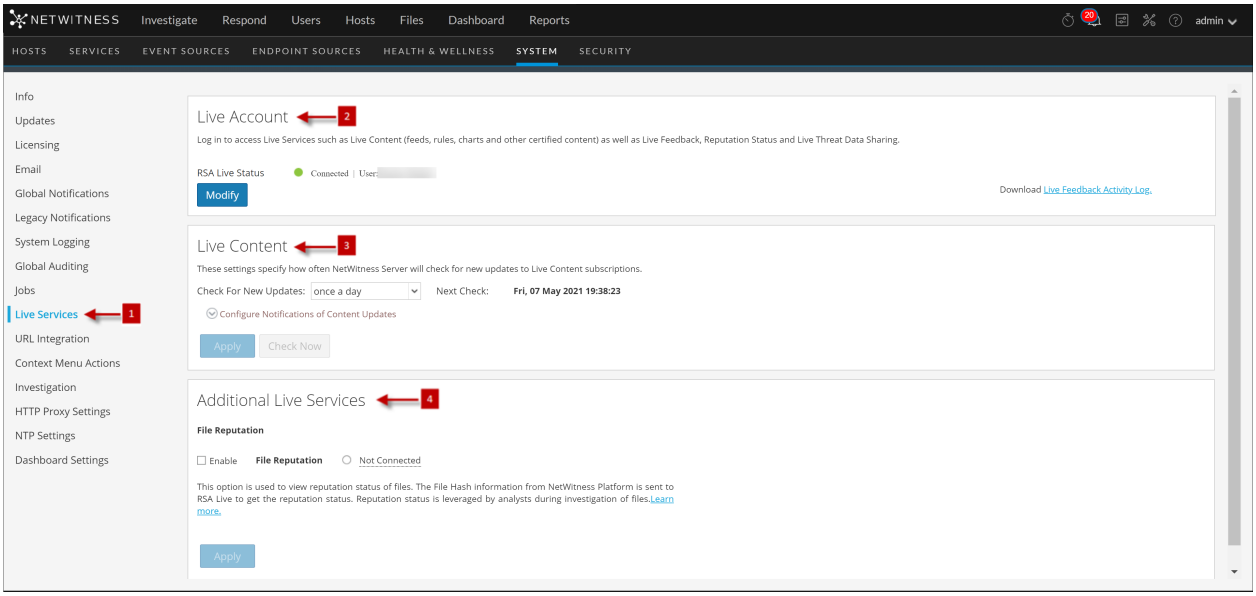
Role	I want to ...	Show me how
Administrator	Setup, Verify Live Service Configuration Panel	Live Services Configuration Panel
Administrator	Overview On Live Feedback	Live Feedback Overview

Related Topics

- Live Feedback Overview
- Configure Live Services Settings
- Upload Data to RSA for Live Feedback
- *Live Services Management Guide*

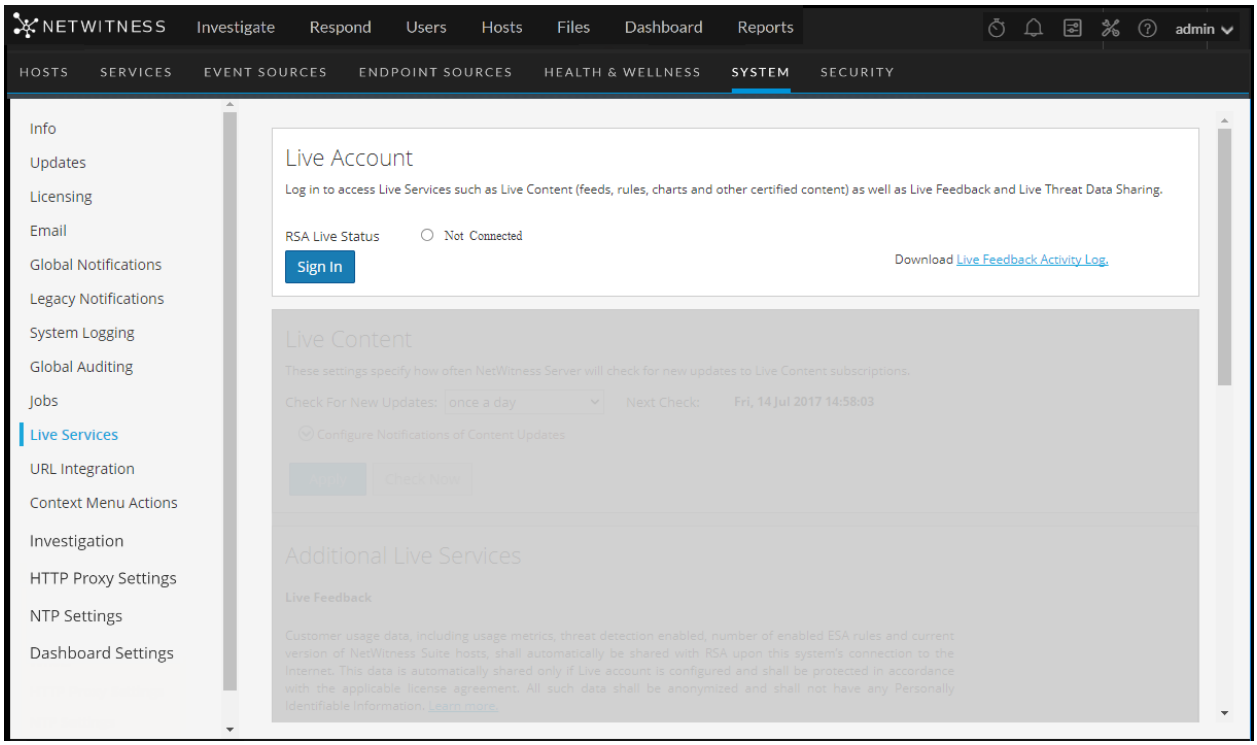
Live Services Quick Look

You access this view in the **(Admin) > System > Live Services**.



1 Displays the Live Services Configuration Panel.
2 Enter Live Account Credentials with the help of Customer Care.
3 Provides updates on Live Content.
4 Additional Live Services provide Live feedback.

Note: If you are not signed in with your Live Account credentials, a masked screen is displayed as shown here.



The Live Configuration panel has three sections: Live Account, Live Content, and Additional Live Services.

Live Account Section

In the **Live Account** section, you must enter the Live credentials. The information needed to set up the user's Live account consists of the Username, Password, and Live URL for the NetWitness Content Management System. This information is provided by Customer Care.

The following table describes the Live Account section features.

Feature	Description
Host	The Live URL for the Content Management System. The default value points to the NetWitness CMS at **cms.netwitness.com**.
Port	The communications port for Live to send requests to the Content Management System. The default value for this field is **443**, which is the communications port on the Content Management System.
SSL	Allows the user to communicate via SSL.
Username	The Live account user name as provided by NetWitness Customer Support.
Password	The Live account user password as provided by NetWitness Customer Support.
Test connection	Tests if the connection is successful or not.
Apply	Saves and applies the configuration.

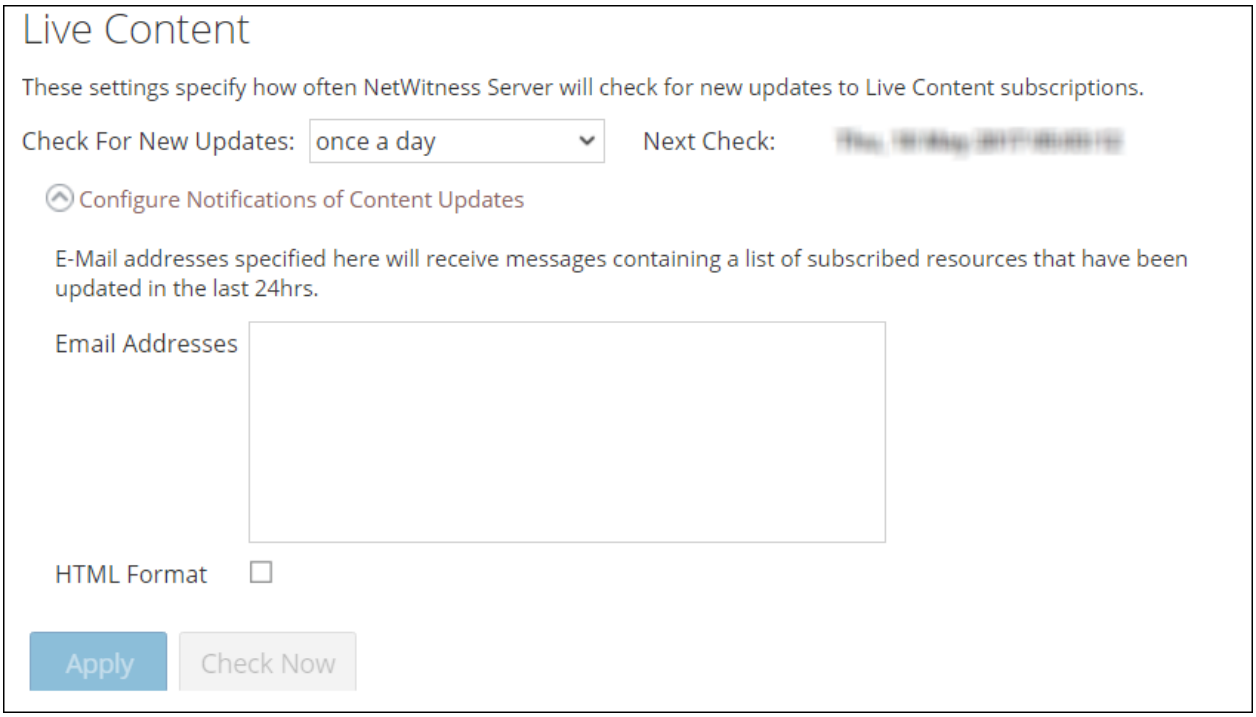
The Live Account section provides an option to download and share the Live Feedback historical data by clicking Live Feedback Activity Log.

For more information about how to download historical data, see Upload Data to RSA for Live Feedback

Live Content Section

You can configure the Live Content Synchronization interval and notification at which NetWitness checks for new updates to Live Content:

Use the **Check for New Updates** field to change the interval. Select an interval from the drop-down list. The default value for this setting is **once a day**.



The following table describes the Live Content features.

Feature	Description
Check for new updates	This setting dictates how often NetWitness checks for new updates to Live Subscriptions and synchronizes subscribed resources and tags: • once a day • twice a day • four times a day • every hour • every other hour • every half hour The default value for this setting is once a day.
Next Check	Displays the time and date of the next scheduled Live synchronization based on the configured interval for checking.
Email Addresses	Email addresses specified here receive messages containing a list of subscribed resources that have been updated in the last 24 hours.
HTML format	Specifies the format of email messages. • Set = HTML • Cleared = text
Check Now	Instead of waiting for the next scheduled resource cycle, this option forces Live to begin immediate synchronization of the subscribed resources in this instance of NetWitness. **Caution:** Use this feature with caution because synchronization can cause a parser reload if a Lua Parser or Flex Parser is deployed in the update cycle. This is acceptable once or twice a day, but a number of back-to-back parser reloads can cause packet loss at the Decoder. If this is the initial setup and you haven't configured Live resource subscriptions, do not Synchronize Now. Wait until you have configured subscriptions.
Apply	Applies the changed configuration to the subscription synchronization behavior. The changes become effective immediately. The **Next Live synchronization is scheduled for** field is updated if the time changed.

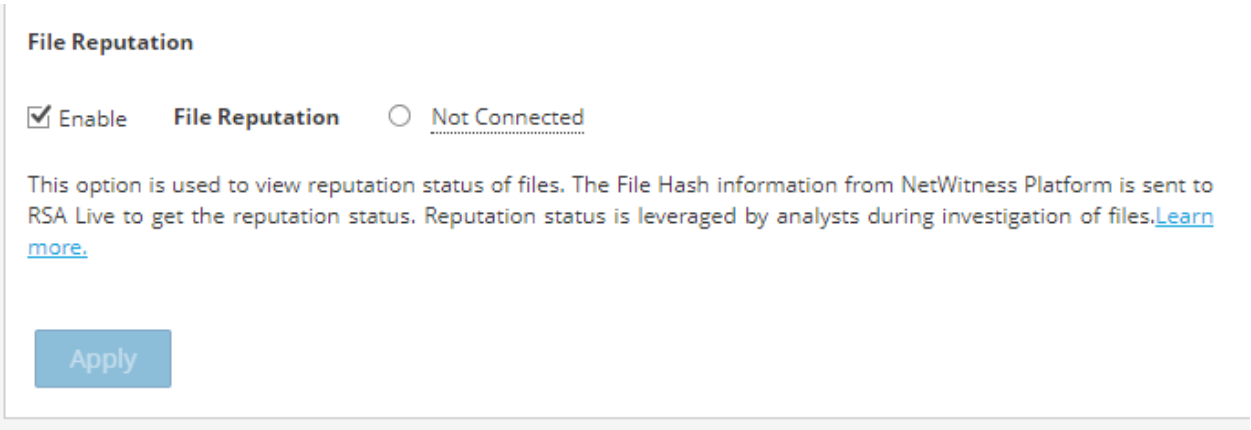
Force Immediate Synchronization

To force immediate synchronization, click **Check Now**. NetWitness checks for updates in subscribed resources.

Instead of waiting for the next scheduled resource cycle, this option forces Live to begin immediate synchronization of the subscribed resources in this instance of NetWitness. One use for this is to see the immediate impact of a configuration change. For example, a new service has been added, or new resources have been toggled for automatic deployment. The scheduled synchronization could take place hours later if Live Services is set to synchronize a few times a day.

Caution: Synchronization can cause a parser reload if a Flex Parser is deployed in the update cycle. This is acceptable once or twice a day, but a number of back-to-back parser reloads can cause packet loss at the Decoder. If this is the initial setup and you haven't configured Live resource subscriptions, do not Synchronize Now. Wait until you have configured subscriptions.

Additional Live Services



Note: Click on Learn more to know more about the data NetWitness is collecting. For more information, see Live Feedback Overview

The following tables describes the Additional Live Services features.

Feature	Description

Live Feedback	Lists the types of data NetWitness is collecting:
	• Product Name
	• Product Version
	• Product Instance
	• Activation Key
	• Details of each Component such as:
	  • ID
	  • Name
	  • Version
	  • Instance ID
	• Metrics for each component

Additional Feedback Insights	Enables NetWitness to send anonymous, technical data about the content usage metrics to NetWitness. This option is enabled by default.

About Live Feedback Participation

Once you sign up for a Live account, Live Feedback automatically collects relevant information for further improvement and anonymously sends it to RSA. The shared data is protected in accordance with the applicable license agreement. For information on Live Feedback, see Live Feedback Overview. For information, see Configure Live Services Settings

If needed, you can manually download historical usage data and share it with RSA. For information on how to download historical usage data and share it with RSA, see Upload Data to RSA for Live Feedback.

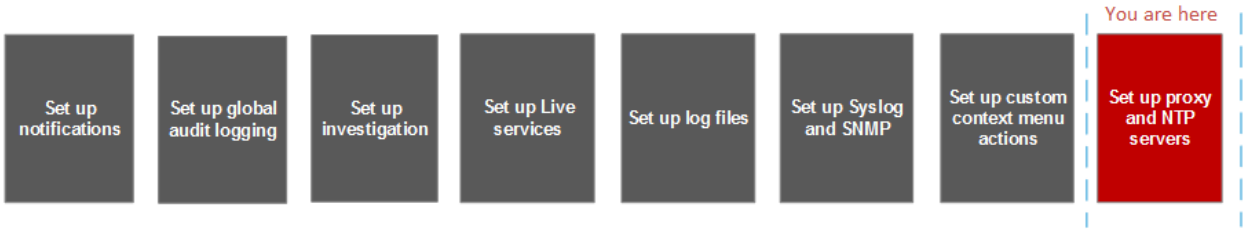
NTP Settings Panel

NTP setting panel is a protocol designed to synchronize the host machine clocks over a network. For more information on NTP see the home page (http://www.ntp.org/).

Note: NetWitness core hosts must be able to communicate with the NetWitness Server host with UDP port 123 for NTP time synchronization.

You use the **(Admin)** > **System** > **NTP Settings** view to configure one or more NTP servers. After you configure an NTP server, NetWitness uses NTP to synchronize the host machine clocks. You configure multiple NTP servers for Fail Over purposes.

Workflow



What you need to do?

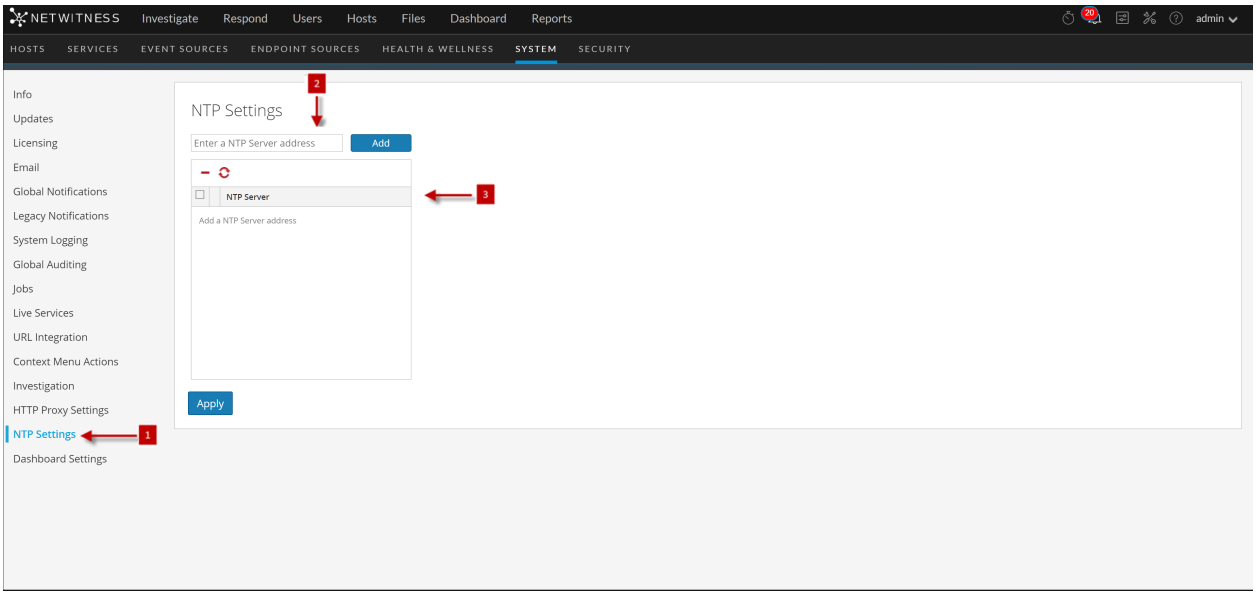
Role	I want to ...	Show me how
Administrator	Add or Modify an NTP Server	Configure NTP Servers

Related Topics

- Configure NTP Servers
- Troubleshoot Issues identified in the NTP Settings Panel or Log Files Messages

Quick Look

The following example illustrates an NTP setting panel. The panel defines how to add NTP server to NTP setting panel.



1 Displays the NTP setting panel.
2 Enter the NTP Server IP Address or hostname.
3 Click on an existing hostname.

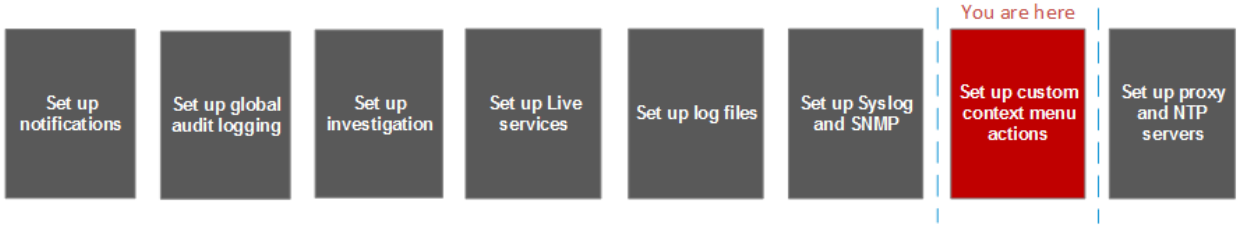
The following table describes the settings in the NTP Settings panel.

Setting	Description
Add	Enter the NTP Server IP Address or hostname to add the NTP server to NetWitness.
	Delete the selected NTP server.
	Synchronizes the selected NTP server.
	Selects the NTP server that you want to delete or synchronize.
NTP Server	NTP Server IP Address or hostname. If you click on an existing hostname, NetWitness makes the hostname editable and displays the following command buttons: • **Update** - Applies your edits. • **Cancel** - Cancels your edits.
Apply	Applies the NTP server settings and synchronizes host machine clocks to NTP.

Context Menu Actions Panel

In the Context Menu Actions panel, Administrators can view built-in context menu actions, and add, edit, or delete custom context menu actions that appear as options in a context menu.

Workflow

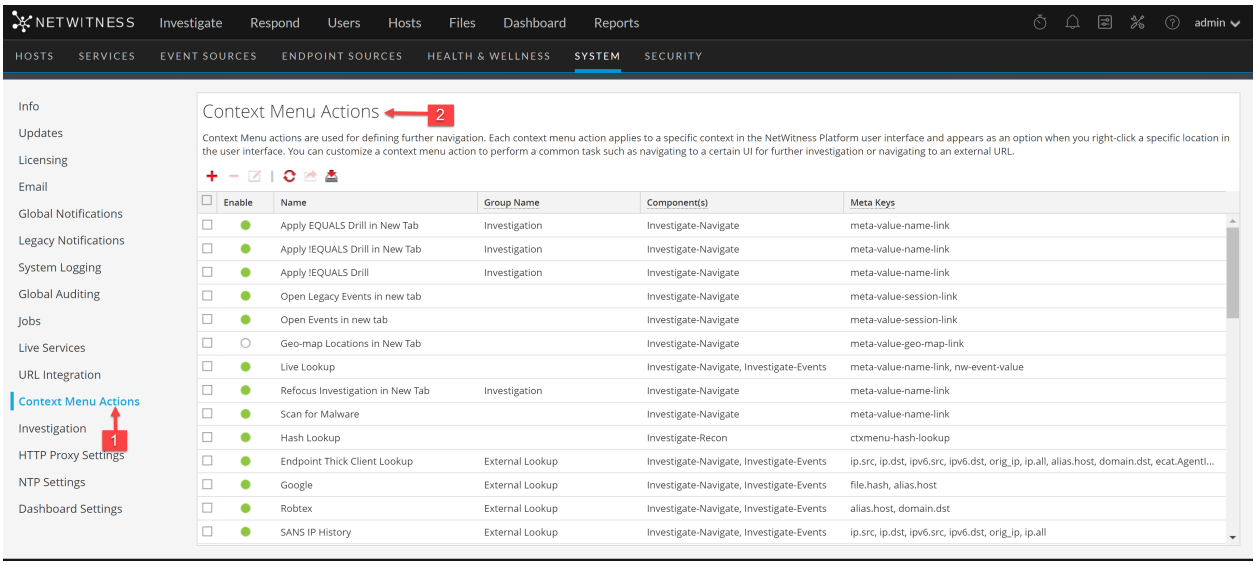


What do you want to do?

Role	I want to ...	Show me how
Administrator	Custom Context Menu Actions panel	Add Custom Context Menu Actions.

Quick Look

The following figure is an example of the Context Menu Actions panel.



1 Displays the Context Menu Actions Panel.
2 Toolbar allows you to Add, Edit, Delete Context Menu Actions.

The Context Menu Actions panel has a list and a toolbar. The following table describes the toolbar options and grid features.

Features	Description
	Displays the Context Menu Configuration dialog, in which you can create a new context action.
	Refreshes the list.
	Deletes the selected context actions. NetWitness does not request confirmation that you want to delete the action. The selected actions are immediately deleted with no opportunity to cancel.
	Displays the Edit Context Action dialog, in which you can edit an existing context action.
Visibility	Displays whether the context menu action is enabled or disabled.
Action Name	The name of the context menu action as it appears on the meta when a user right-clicks to initiate action.
Action Group	The action group under which this context menu action is grouped.
Component	The UI component to which the Action Name and Action Group belong.

Meta Keys	The names of the modules in which the context action is available. Currently all built-in context menu actions are for the Investigation module.
	When creating a context menu action, the parameter is `modules`.
	Here is a line of sample code:
	`"modules": [`
	`"investigation"`
	`],`

CSS Classes and Examples

CSS classes can be meta keys and non-meta keys.

Meta Key CSS Classes

One type of CSS class that you can add is meta keys. For meta keys that have a period, change the period to a dash when defining a CSS class. For example, the meta key `alias.host` becomes the CSS class `alias-host`. The meta key `ip.src` becomes the CSS class `ip-src`.

Non-Meta Key CSS Classes

Built-in non-meta key CSS Classes are also available. The classes in the following table define actions and the part of the user interface where the action is available.

CSS Class	Type	Description
meta-value-session-link	Action	Open on meta session count number
meta-value-name-link	Action	Open on meta value name
nw-event-value	Action	Use for reconstruction context actions on meta value
UAP.investigation.navigate.view. NavigationPanel	User interface	Applies to Navigate view
UAP.investigation.events.view. EventGrid	User interface	Applies to Event View
UAP.investigation.reconstruction.view. content.ReconstructedEventDataGrid	User interface	Applies to Event Reconstruction View

Example

This is a commented example of a context menu action to validate the user agent from the Client Application (client) meta key. The comments are removed automatically once applied in the Administration System view. The new menu item is displayed after restarting the browser.

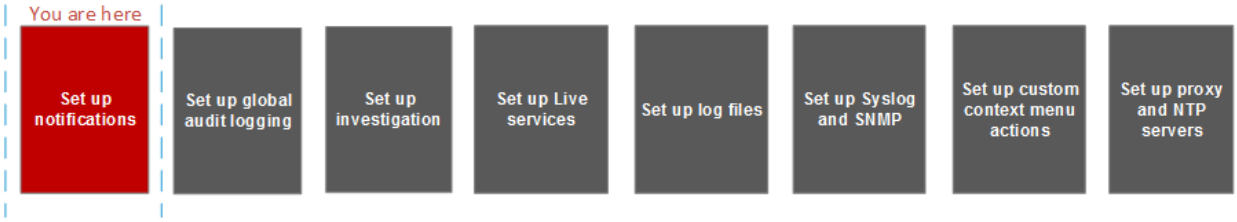
```
 {
      "displayName": "User Agent String Lookup”, <!-- What name shows up in NW
UI -->
      "cssClasses": [
          “client”  <!-- What meta key to launch from -->
      ],
      "description": "",
      "type": "UAP.common.contextmenu.actions.URLContextAction",
      "version": "1",
      "modules": [
          "investigation"
      ],
      "local": "false",
      "groupName": "externalLookupGroup”, <!-- What group to show link in.
Remove line to show in main list -->
      "urlFormat": "http://www.useragentstring.com/?uas={0}&getText=all”, <!-- The {0}
gets replaced with whatever was right clicked on -->
      "disabled": "",
      "id": "UserAgentStringAction",
     "moduleClasses": [
          "UAP.investigation.navigate.view.NavigationPanel”, <-- Enabled in
Navigate pane-->
         “UAP.investigation.events.view.EventGrid” <-- Enabled in Event View
pane -->
      ],
      "openInNewTab": "true",
      "order": "15"
  }
```

Legacy Notifications Configuration Panel

The Legacy Notifications Configuration panel provides the ability to configure syslog and SNMP notification settings. These configurations are used for Entitlement, legacy Event Source Management (ESM), Warehouse Connector monitoring, and Archiver monitoring.

Procedures related to these settings are described in Configure Syslog and SNMP Settings.

Workflow



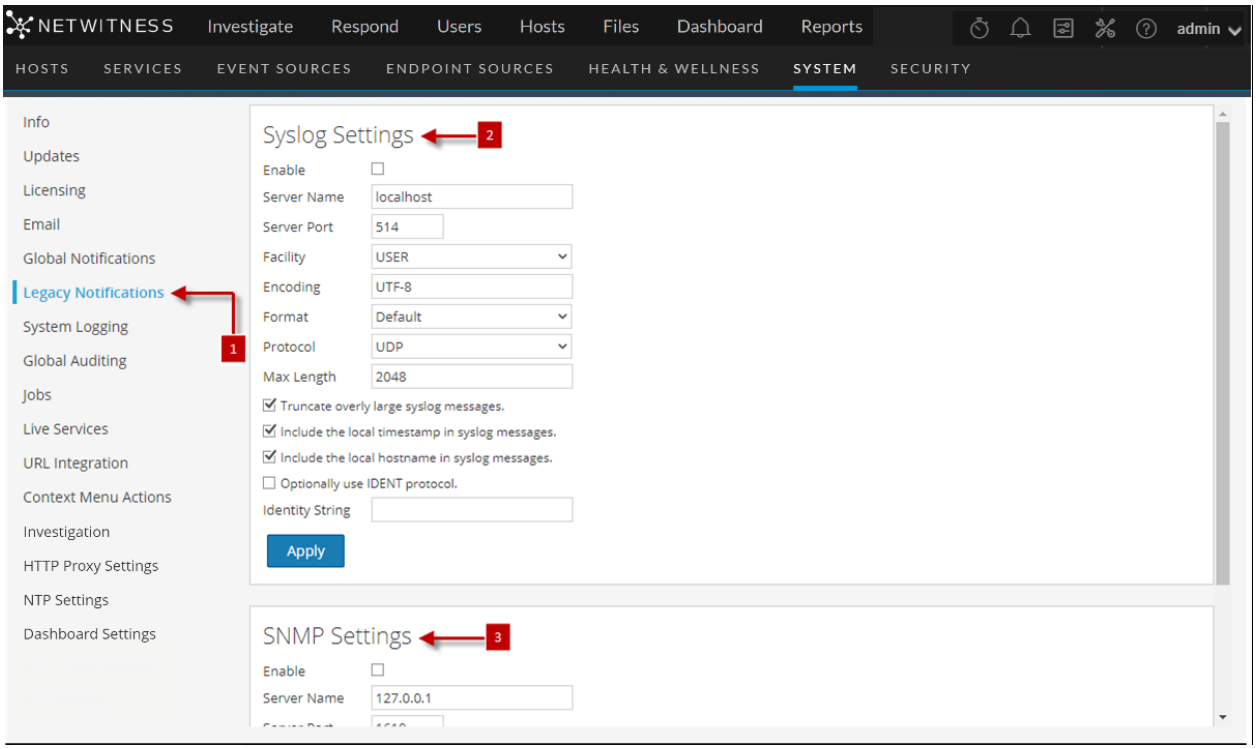
What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure Syslog Settings	Configure Syslog and SNMP Settings
Administrator	Configure SNMP Settings	Configure Syslog and SNMP Settings

Related Topics

- Configure Syslog and SNMP Settings

Quick Look



1	Displays the Legacy Notification Configuration Panel.
2	Allows the user to configure syslog notifications for Entitlement, legacy Event Source Management (ESM), Warehouse Connector monitoring, and Archiver monitoring.
3	Allows the user to configure SNMP notifications for Entitlement, legacy Event Source Management (ESM), Warehouse Connector monitoring, and Archiver monitoring.

The Legacy Notifications Configuration Panel consists of two sections: Syslog Settings and SNMP Settings.

Syslog Settings

The following table describes the available options for configuring syslog notifications for Entitlement, legacy Event Source Management (ESM), Warehouse Connector monitoring, and Archiver monitoring.

Feature	Description
Enable	Enables the syslog settings configured here.
Server Name	Specifies the host where the target syslog process is running.
Server port	Specifies the port where the target syslog process is listening.

Feature	Description
Facility	Specifies the designated syslog facility to use for all outgoing messages. Possible values are KERN, USER, MAIL, DAEMON, AUTH, SYSLOG, LPR, NEWS, UUCP, CRON, AUTHPRIV, FTP, LOCAL1 through LOCAL7.
Encoding	Specifies the encoding to use for text in syslog messages, for example, UTF-8.
Format	Specifies the message format. Possible values are: Default, PCI DSS, or SEC.
Protocol	Specifies the communications protocol used when sending syslogs: UDP or TCP. By default, the UDP protocol is selected.
Max length	Specifies the maximum length in bytes of any syslog message. The default value is **2048**. Messages that exceed the maximum length are truncated when the **Truncate overly large syslog messages** checkbox is selected.
Truncate overly large syslog messages	When checked, any messages exceeding the maximum length are truncated.
Include the local timestamp in syslog messages	When checked, NetWitness includes the local timestamp in messages.
Include the local hostname in syslog messages	When checked, NetWitness includes the local hostname in syslog messages.
Optionally use IDENT protocol	When checked, NetWitness prepends the identity string to outgoing syslog alerts.
Identity string	This is an identity string to be prepended to each syslog alert. If the string is blank, no identity string is prepended to the outgoing syslog alerts. You can use this to identify the source of the alert. Users conventionally set it to the name of the program that sends the syslog message.
Apply	Applies the syslog configuration settings.

SNMP Settings

The following table describes the available options for configuring SNMP notifications for Entitlement, legacy Event Source Management (ESM), Warehouse Connector monitoring, and Archiver monitoring.

Feature	Description
Enable	Enables the SNMP settings configured here.
Server Name	Specifies the SNMP trap host.
Server port	Specifies the listening port on the SNMP trap host
SNMP version	Specifies the SNMP version, **v1** or **v2c**.

Feature	Description
Trap OID	Specifies the object ID for the SNMP trap on the trap host that receives the audit event. The default value is **0.0.0.0.0.1**.
Community	Specifies the community string used to authenticate on the SNMP trap host, the default value is **public**.
Enable	Enables SNMP notifications as configured here.
Apply	Applies the SNMP configuration settings.

NetWitness® Platform XDR

Version 12.0

System Maintenance Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

October, 2022

Contents

System Maintenance **13**

Review Best Practices **14**

Safeguarding Assets with NetWitness Supplied Policies 14

Safeguarding Assets with Policies Based on Your Environment 14

Creating Rules and Notifications Judiciously 14

Troubleshooting Issues 14

Monitor NetWitness Platform Health **15**

Monitor Health and Wellness using NetWitness Platform UI 16

Manage Policies 17

Add a Policy 17

Add Policy Example 20

Edit a Policy 22

Duplicate a Policy 23

Assign Services or Groups 24

Remove Services or Groups 26

Add or Edit a Rule 26

Hide or Show Rule Conditions Columns 27

Delete a Rule 28

Suppress a Rule 28

Suppress a Policy 29

Add an Email Notification 29

Delete an Email Notification 29

Include the Default Email Subject Line 30

Monitor System Statistics 33

Filter System Statistics 34

View Historical Graphs of System Statistics 37

Monitor Service Statistics 38

Add Statistics to a Gauge or Chart 39

Create a Gauge for a Statistic 39

Create a Timeline Chart for a Statistic 39

Search for a Statistic in the Chart Stats Tray 40

Edit Properties of Statistics Gauges 41

Edit Properties of a Gauge 41

Add Stats to the Gauges Section 41

Edit Properties of Timeline Charts 43

Edit Properties of a Timeline 43

Edit Properties of a Historical Timeline 43
Add Stats to Timeline Charts 44
Monitor Hosts and Services 45
Filter Hosts and Services in the Monitoring View 46
Monitor Host Details 48
Monitor Service Details 49
Monitor Event Sources 52
Monitor Alarms 53
Monitor Health and Wellness Using SNMP Alerts 55
SNMP Configuration 55
Thresholds 55
Configure SNMPv3 for a Host 55
Set the Threshold for a Service 56
SNMP Traps for System Status 57
Troubleshooting Health & Wellness 58
Issues Common to All Hosts and Services 58
Issues Identified by Messages in the Interface or Log Files 58
Issues Not Identified by the User Interface or Logs 63
Monitor New Health and Wellness 65
Dashboard 65
Visualization 65
Monitors 66
Notifications 66
Installing New Health and Wellness 66
Accessing New Health and Wellness Dashboards 66
Configuring Notifications 67
Adding Alert Notifications 68
Suppressing Notifications 69
Monitoring through Dashboards 70
Creating a Custom Dashboard 71
Monitoring through Alerts 72
Creating Custom Monitors 74
Adding Custom Trigger to an Existing Monitor 75
Managing Dashboards and Alerts 77
Modify a Dashboard 77
Delete a Dashboard 77
Delete a Visualization 78
Modify an Existing Trigger 78
Managing Notifications 79
Modify a Notification 79
Modify a Notification Suppression Policy 79

Advanced Configurations 80
Restore Default Content 80
Enable Services 80
Disable Services 81
Update an Interval 82
Update the Default Configuration 83
Configure the Data Retention Policy 85
Backup and Restore New Health and Wellness 89
Troubleshooting New Health and Wellness 90
Appendices 92
New Health and Wellness Dashboards 93
Deployment Health Overview Dashboard 93
Hosts Dashboard 95
Logs Dashboard 97
Packet Overview Dashboard 99
Analysis Dashboard 101
Endpoint Dashboard 103
ESA Correlation Overview Dashboard 105
Logstash Input Plugin Dashboard 106
License Usage Dashboard 108
New Health and Wellness Monitors 113
Uninstall New Health and Wellness 116
Manage NetWitness Platform Updates 120
Reissue Certificates 121
Introduction 121
CA Certificate Reissue 121
Service Certificate Reissue 121
Reissuing Service Certificate 122
When to Use the --host-all Argument 123
cert-reissue Arguments and Options for All Hosts 123
When to Use the Individual Host Argument (--host-key <ID, IP, hostname or display name of host>) 124
Reissuing Certificates for All Hosts Except Windows Legacy Collection (WLC) host 125
Running the Cert-Reissue Command for All Hosts 125
Running the Cert-Reissue Command for an Individual Host 125
Reissuing Certificates for a WLC Host 125
Successful Reissue Summary Report 126
Unsuccessful Reissue Summary Reports 126
Reissue Failed for Host and Aborted Command 126
Reissue Certificate Partially Executed 127

Display System and Service Logs **128**
View System Logs 128
Display Service Logs 128
Filter Log Entries 129
Show Details of a Log Entry 129
Access Reporting Engine Log File 131
All Log Files 131
Upstart Logs 131
Search and Export Historical Logs 132
Display the Historical System Log 132
Display a Historical Service Log 132
Search Log Entries 133
Show Details of a Log Entry 133
Page Through Log Entries 134
Export a Log File 134
Maintain Queries Using URL Integration **135**
Edit a Query 135
Delete a Query 136
Clear All Queries 136
Use a Query in a URI 136
Manage the deploy_admin Account **138**
Change the deploy_admin Account Password 138
Change the deploy_admin Account Password in a Mixed Version Environment 138
Change the deploy_admin Account Password for a Component Host that is Unavailable 139
NW Server Host Secondary IP Configuration Management **140**
Change Host Network Configuration **141**
Change Host Network Configuration 141
SSO 144
Reporting Engine 145
UCF 146
PAM 146
ECAT 146
RSA NetWitness Orchestrator (By Demisto) 149
Audit Logging 150
Health and Wellness 150
Malware Analysis 150
Windows Legacy Collection 151
Change Network Configuration for Warm Standby (Secondary) Server 152
Reconnecting Component Hosts with NW Server Hosts 152

Manage Custom Host Entries **154**
Manage Custom Host Entries in /etc/hosts 154
Manage Public or NAT IPv4 Addresses for Hosts 154
Manage Custom Jetty Configuration 154
Configure FIPS Support **156**
FIPS support for Log Collectors 157
FIPS support for Log Decoders and Decoders 157
DISA STIG **159**
How STIG Limits Account Access 159
NetWitness Passwords 159
Generate the OpenSCAP Report 159
Disable Rules in OpenSCAP Report that Hang the Report 160
Install OpenSCAP 160
Sample Report 160
Report Fields 161
Create the OpenSCAP Report 162
Create Report in HTML Only 162
Create Report in XML Only 162
Create Report in Both XML and HTML 163
Manage STIG Controls Script (manage-stig-controls) 163
Commands 163
Control Groups 164
Other Arguments 165
Rules List 166
Exceptions to STIG Compliance 181
Key to Elements in Exception Descriptions 181
CCE Number 181
Control Group ID 181
Check 181
Comments 182
Customer Responsibility Exceptions 182
CCE-26952-2 Configure Periodic Execution of AIDE (Control Group = audit) 182
CCE-27096-7 Install AIDE (Control Group = n/a) 182
CCE-27218-7 Remove the X Windows Package Group 182
CCE-27295-5 Use Only FIPS 140-2 Validated Ciphers (Control Group = n/a) 183
CCE-27334-2 Ensure SELinux State is Enforcing 183
CCE-27445-6 Disable SSH Root Login (Control Group = ssh-prevent-root) 183
CCE-80127-4 Install McAfee Virus Scanning Software (Control Group = n/a) 183
CCE-80129-0 Virus Scanning Software Definitions Are Updated (Control Group = n/a) 183
CCE-80207-4 Enable Smart Card Login (Control Group = n/a) 184

CCE-80359-3 Enable FIPS Mode in GRUB2 (Control Group = fips-kernel) 184
CCE-80374-2 Configure Notification of Post-AIDE Scan Details (Control Group = n/a) 184
CCE-80375-9 Configure AIDE to Verify Access Control Lists (Control Group = n/a) 185
CCE-80376-7 Configure AIDE to Verify Extended Attributes (Control Group = n/a) 185
CCE-80377-5 Configure AIDE to Use FIPS 140-2 for Validating Hashes (Control Group = n/a) 185
CCE-80519-2 Install Smart Card Packages For Multi-Factor Authentication (Control Group = n/a) 185
Exceptions That Are Not a Finding 186
CCE-26404-4 Ensure /var Located On Separate Partition (Control Group = n/a) 186
CCE-26828-4 Set GNOME Login Inactivity timeout (Control Group = n/a) 186
CCE-26884-7 Set Lockout Time For Failed Password Attempts (Control Group = auth) 186
CCE-26971-2 Ensure /var/log/audit Located On Separate Partition (Control Group = audit) 187
CCE-27127-0 Enable Randomized Layout of Virtual Address Space (Control Group = n/a) 187
CCE-27157-7 Verify File Hashes with RPM (Control Group = n/a) 187
CCE-27339-1 Record Events that Modify the System's Discretionary Access Controls - chmod 188
CCE-27209-6 Verify and Correct File Permissions with RPM (Control Group = n/a) 188
CCE-27303-7 (Control ID = 2) Modify the System Login Banner (Control Group = ssh) 189
CCE-27311-0 Very Permissions on SHH Server *.pub Key Files (Control Group = na) 189
CCE-27314-4 Enable SSH Warning Banner (Control Group = na) 189
CCE-27349-0 Set Default firewalld Zone for Incoming Packets (Control Group = n/a) 190
CCE-27386-2 Ensure Default SNMP Password Is Not Used (Control Group = n/a) 190
CCE-27455-5 Use Only FIPS 140-2 Validated MACs (Control Group = na) 190
CCE-27471-2 Disable SSH Access via Empty Passwords (Control Group = n/a) 190
CCE-27485-2 Very Permissions on SHH Server Private *.key Key Files (Control Group = na) 190
CCE-80156-3 Disable Kernel Parameter for Sending ICMP Redirects for All Interfaces (Control Group = n/a) 191
CCE-80157-1 Disable Kernel Parameter for IP Forwarding (Control Group = n/a) 191
CCE-80158-9 Configure Kernel Parameter for Accepting ICMP Redirects for All Interfaces (Control Group = n/a) 191
CCE-80163-9 Configure Kernel Parameter for Accepting ICMP Redirects By Default (Control Group = n/a) 191
CCE-80165-4 Configure Kernel Parameter to Ignore ICMP Broadcast Echo Requests (Control Group = n/a) 192
CCE-80225-6 Print Last Log (Control Group = n/a) 192
CCE-80226-4 Enable Encrypted X11 Forwarding (Control Group = n/a) 192
CCE-80348-6 Ensure gpgcheck Enabled for Repository Metadata (Control Group = n/a) 192
CCE-80383-3 Record Attempts to ALter Logon Events - failock (Control Group = na) 193
CCE-80399-9 Ensure auditd Collects Information on the Use of Privileged Commands - userhelper (Control Group = na) 193
CCE-80437-7 Configure PAM in SSSD Services (Control Group = n/a) 193
CCE-80438-5 Configure Multiple DNS Servers in /etc/resolv.conf (Control Group = n/a) 193

CCE-80439-3 Configure Time Service Maxpoll Interval (Control Group = na) 193
CCE-80447-6 Configure the Firewalld Ports (Control Group = n/a) 194
CCE-80515-0 Configure SSSD LDAP Backend Client CA Certificate Location (Control Group = n/a) 194
CCE-80545-7 Verify and Correct Ownership with RPM (Control Group = n/a) 194
CCE-80546-5 Configure SSSD LDAP Backend to Use TLS For All Transactions (Control Group = n/a) 195
CCE-80998-8 Verify firewall Enabled 195
CCE-82035-7 Ensure /var/log/audit Located On Separate Partition 196
CCE-82053-0 Ensure /tmp Located On Separate Partition 196
CCE-82353-4 Ensure /var Located On Separate Partition 196
Rules Supported in a Future Release 196
CCE-27277-3 Disable Modprobe Loading of USB Storage Driver (Control Group = services) 197
CCE-27309-4 Set Boot Loader Password in grub2 (Control Group = fips-kernel) 197
CCE-80179-5 Configure Kernel Parameter for Accepting IPv6 Source-Routed Packets for All Interfaces 197
CCE-80660-4 Record Any Attempts to Run setfiles (Control Group = audit) 198
CCE-80661-2 Ensure auditd Collects Information on Kernel Module Loading - create_module (Control Group = audit) 198

Troubleshoot NetWitness Platform 199

Debugging Information 200
NetWitness Log Files 200
Files of Interest 200
Error Notification 202
Miscellaneous Tips 203
Audit Log Messages 203
NwConsole for Health & Wellness 203
Thick Client Error: remote content device entry not found 203
View Example Parsers 203
Configure WinRM Event Sources 203
Troubleshoot Feeds 204
Overview 204
Details 204
How it Works 204
Feed File 204
Troubleshooting 205
Feed File Existence 205
Group Meta Populated on LD 205
Device Group Meta on Concentrator 206
SMS Log File 206
Verify Logstats Data is Getting Read and Published by ESMReader and ESMAggregator 207

Configure JMX Feed Generator Job Interval 209
Troubleshooting Cert-Reissue Command 211
Argument Options Used for Troubleshooting 211
Problems and How to Troubleshoot Them 212
References 216
Health and Wellness View 217
Health and Wellness View - Alarms View 218
What do you want to do? 218
Related Topics 218
Quick Look 218
Alarm Details Panel 220
Event Source Monitoring View 221
Health and Wellness Historical Graphs 222
Historical Graph View for Events Collected from an Event Source 223
Historical Graph for System Stats 224
Health and Wellness Settings View - Archiver 226
What do you want to do? 226
Quick Look 226
Features 226
Health and Wellness Settings View - Event Sources 228
Health and Wellness Settings View - Warehouse Connector 229
Access the Warehouse Connector Monitoring view 229
What do you want to do? 229
Related topics 229
Quick Look 229
Warehouse Connector Monitoring parameters 230
Monitoring View 231
What do you want to do? 231
Quick Look 231
Groups Panel 231
Hosts Panel 232
Archiver Details View 234
Broker Details View 236
Concentrator Details View 237
Decoder Details View 238
ESA Correlation Details View 240
Health Stats Tab 240
JVM Tab 241
ESA Analytics Details View 242
Host Details View 243

Log Collector Details View 245
Collection Tab 245
Event Processing Tab 245
Log Decoder Details View 247
Malware Details View 249
Warehouse Connector Details View 250
Policies View 251
What do you want to do? 251
Quick Look 251
Policies Panel 251
Policy Detail Panel 252
Groups dialog 254
Rules Dialog 254
Threshold Operators 256
Health & Wellness Email Templates 257
Health & Wellness Default SMTP Template 257
Alarms Template 258
NetWitness Platform Out-of-the-Box Policies 259
System Stats Browser View 265
What do you want to do? 265
Related Topics 265
Quick Look 265
Filters 266
Commands 266
System Stats View Display 266
New Health & Wellness Tab 268
What do you want to do? 268
Related Topics 268
Quick Look 268
System View - System Info Panel 270
System Updates Panel - Settings Tab 272
What do you want to do? 272
Related Topics 272
Quick Look 272
Features 272
System Logging - Settings View 274
What do you want to do? 274
Related Topics 274
Quick Look 274
Features 275

Log Settings 275
Package Configuration 275
System Logging - Realtime Tab 277
What do you want to do? 277
Related Topics 277
Quick Look 277
Features 278
Toolbar 278
Log Grid Columns 279
System Logging - Historical Tab 280
What do you want to do? 280
Related Topics 280
Quick Look 280
Features 281
Search Log Entries 282
Show Details of a Log Entry 282
Page Through the Entries 283
Export 283

System Maintenance

This guide tells administrators how to manage hosts and services in the network, maintain and monitor the network, run jobs, and tune performance after initial network setup.

The following diagram shows the different system maintenance tasks available to you.



The following topics describe these tasks:

- Review Best Practices
- Monitor Health and Wellness using NetWitness Platform UI
- Manage NetWitness Platform Updates
- Display System and Service Logs
- Maintain Queries Using URL Integration
- Configure FIPS Support
- Change Host Network Configuration
- DISA STIG
- Troubleshoot NetWitness Platform

Review Best Practices

Review the following best practices to maintain your NetWitness Platform deployment.

Safeguarding Assets with NetWitness Supplied Policies

The purpose of the NetWitness core policies delivered with NetWitness are for safeguarding your NetWitness domain assets immediately (before you configure rules specific to your environment and your security policy).

NetWitness recommends that you set up email notifications to the appropriate asset owners for these policies as soon as possible. This will notify them when performance and capacity thresholds are crossed so they can take action immediately.

NetWitness also recommends that you evaluate the core policies and disable a policy or change its service and group assignments according to your specific monitoring requirements.

Safeguarding Assets with Policies Based on Your Environment

NetWitness core policies are generic and may not provide sufficient monitoring coverage for your environment. NetWitness recommends that you gather issues over a period of time, that are not identified by the NetWitness core policies, and configure rules to help you prevent these issues.

Creating Rules and Notifications Judiciously

NetWitness recommends that you make sure that each rule and policy is necessary before you implement it, if possible. NetWitness also recommends that you review implemented policies on a regular basis for their validity. Invalid alarms and email notifications can adversely affect the focus of the asset owners.

Troubleshooting Issues

NetWitness recommends that you review Troubleshooting Health & Wellness and Troubleshoot NetWitness Platform when you receive error messages in the user interface and log files from hosts and services.

Monitor NetWitness Platform Health

Monitor the Health & Wellness of the NetWitness using any of the following:

- Monitor Health and Wellness using NetWitness Platform UI
- Monitor New Health and Wellness

Monitor Health and Wellness using NetWitness Platform UI

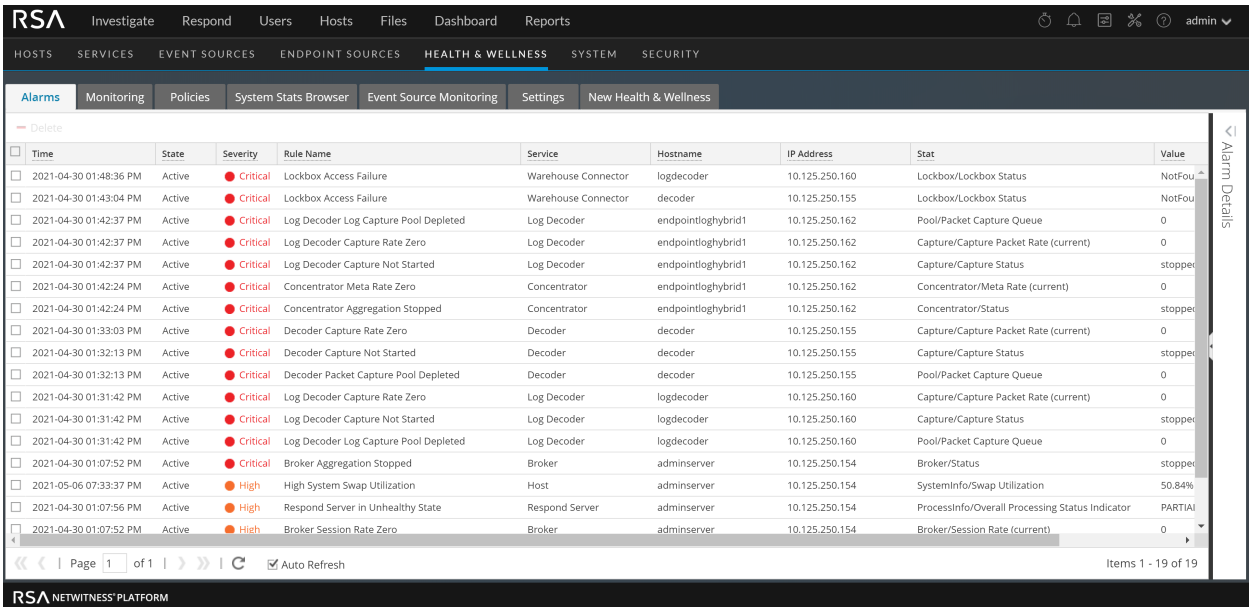
The Health & Wellness module of NetWitness enables you to:

- View the current health of all the hosts, all services running on the hosts, and various aspects of the health of your hosts.
- Monitor the hosts and services in your network environment.
- View details of various event sources configured with NetWitness.
- View system stats for the selected hosts by filtering the views as required.

You can also configure Archiver and Warehouse Connector monitoring, monitor host statistics, and work with system logs to monitor NetWitness.

Note: All users have permission to view the entire Health and Wellness interface by default. The Administrator and the Operator roles are the only roles that can manage the Policies view by default. Refer to the "Role Permissions" topic in the *Security User Management Guide* for a complete list of the default permissions for the NetWitness Interface.

The following figure displays the Health & Wellness module of the NetWitness user interface.



Manage Policies

Policies are either user-defined or supplied by NetWitness. A policy defines:

- Services and hosts to which the policy applies.
- Rules that specify statistical thresholds that govern alarms.
- When to suppress the policy.
- Who to notify when an alarm triggers and when to notify them.

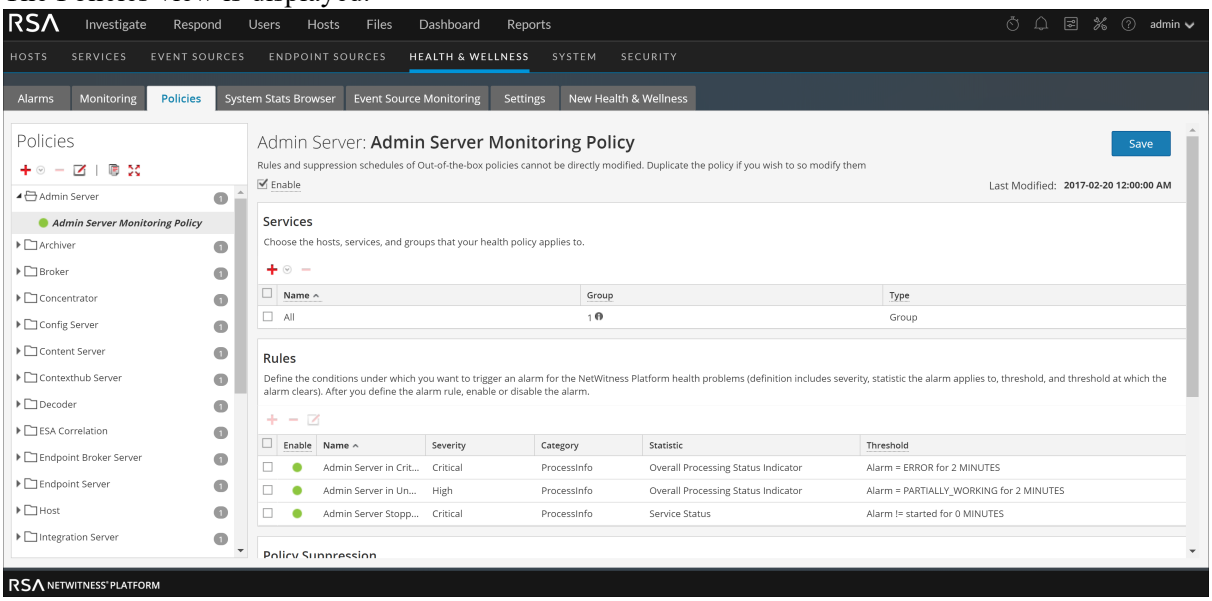
For related reference topics, see NetWitness Platform Out-of-the-Box Policies

Note: You can now configure a policy to notify Public Key Infrastructure (PKI) certificate expiration status.

Add a Policy

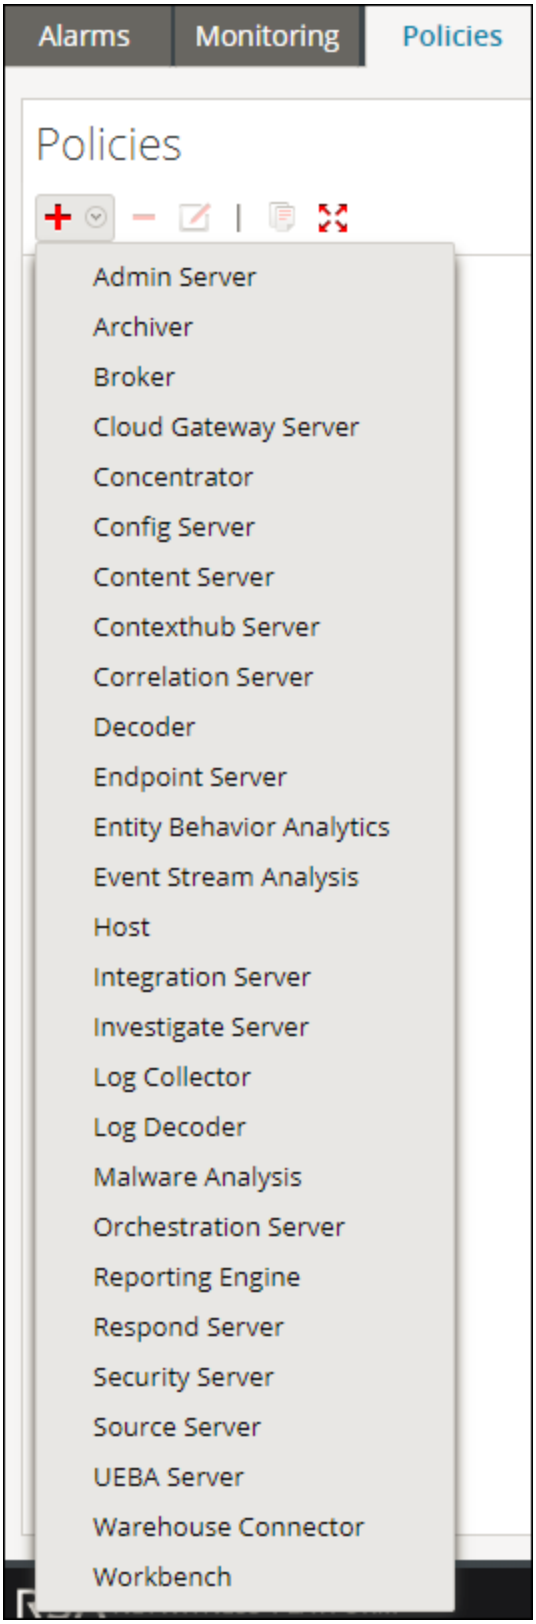
1. Go to  **(Admin) > Health & Wellness.**
2. Click the **Policies** tab.

   The Policies view is displayed.

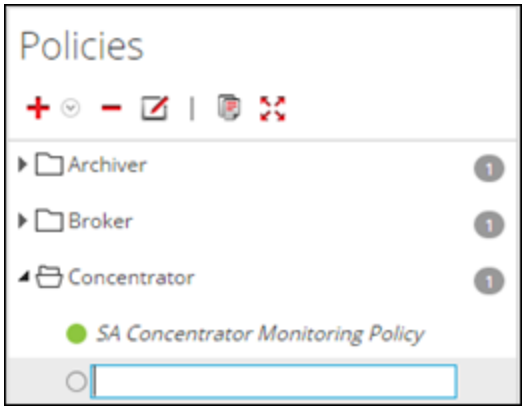


3. Click  in the **Policies** panel.

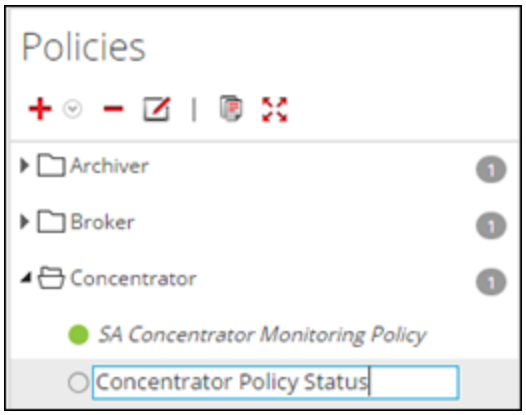
A list of your hosts and services displays for which you can create health policies.



4. Select a host or service (for example, **Concentrator**).
   For a PKI policy, you must select a host (for example, Host).
   The host or service is displayed in the Policies panel with a blank Policy Detail panel.



5. Enter a name for the Policy (for example, **Concentrator Policy Status**) in the **Policies** panel.



   The name (for example, **Concentrator Policy Status**) is now displayed as the policy name in Policy Detail panel.

6. Create a Policy in the Policy Detail panel:

   a. Select the **Enable** checkbox.

   b. Add relevant services (in this example, any relevant Concentrator services) that you want to monitor for health statistics.
      For a PKI policy, you must select the LOCALHOST to monitor for health statistics.

   c. Add rule conditions to configure the policy.

   d. Suppress enforcement of the policy for the time periods you want.

   e. Add any email notifications you want for the policy.

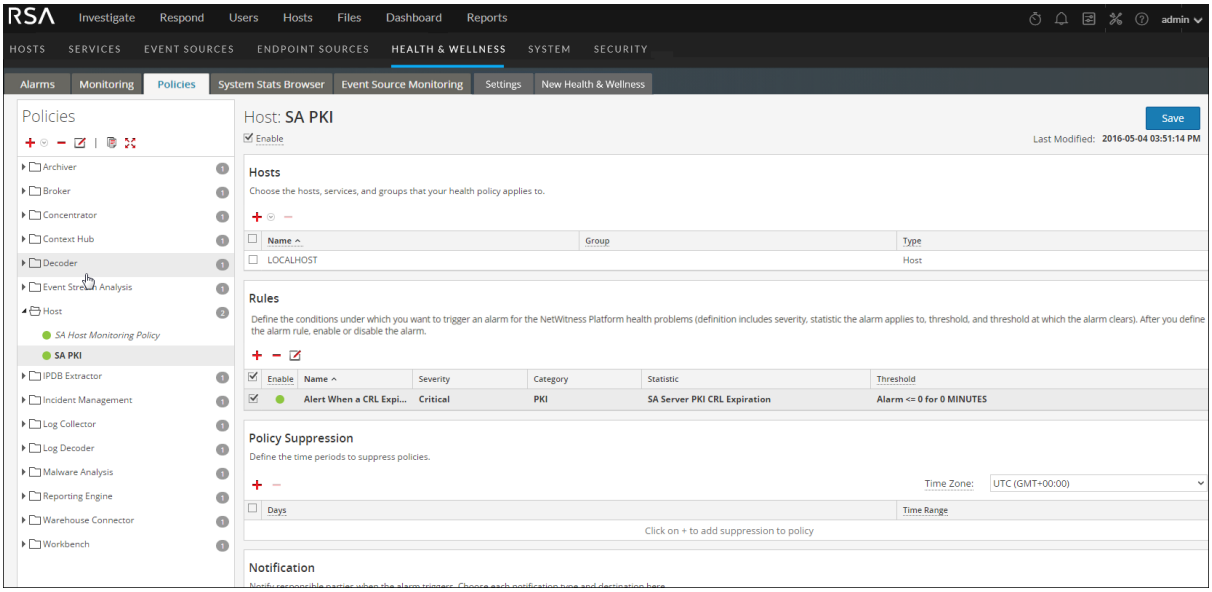
   f. Click **Save** in the Policy Detail panel.

      The Policy is added.

Add Policy Example

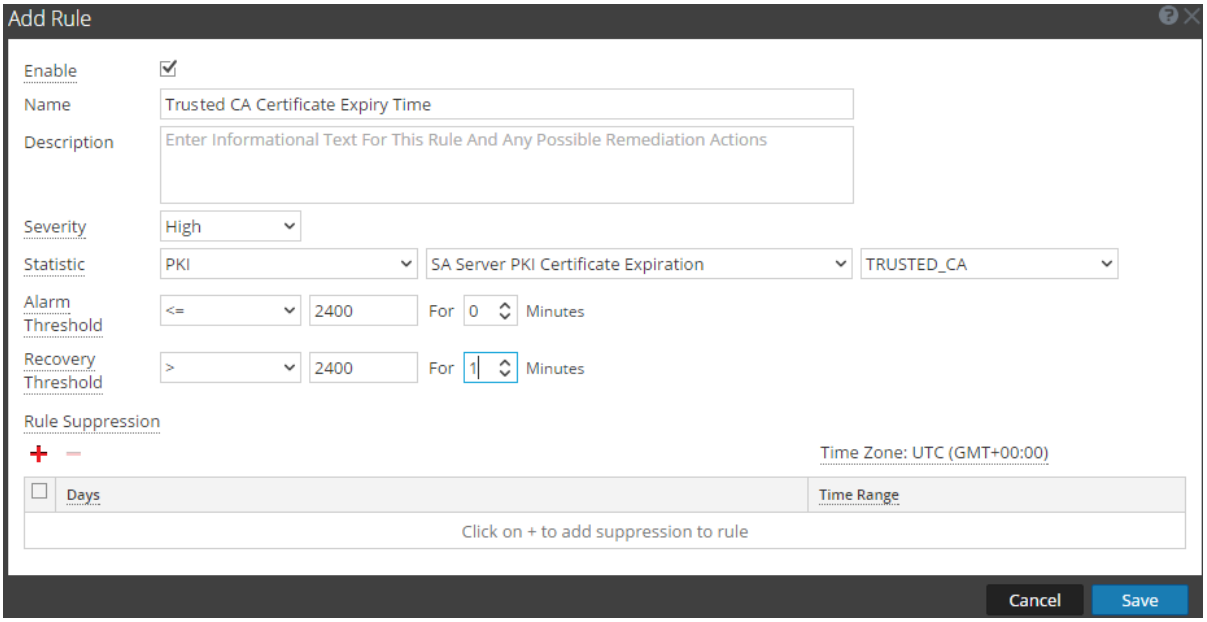
Below is a high-level example of configuring a PKI policy:

1. Add a new PKI policy.

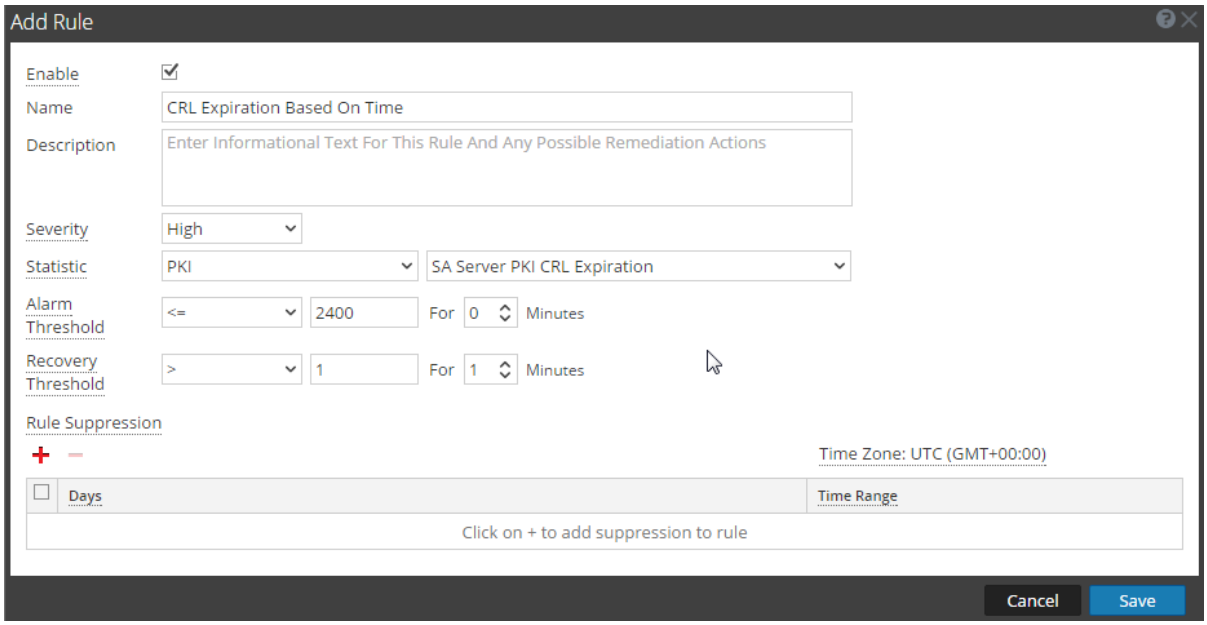


2. Add a Rule with Statistics:

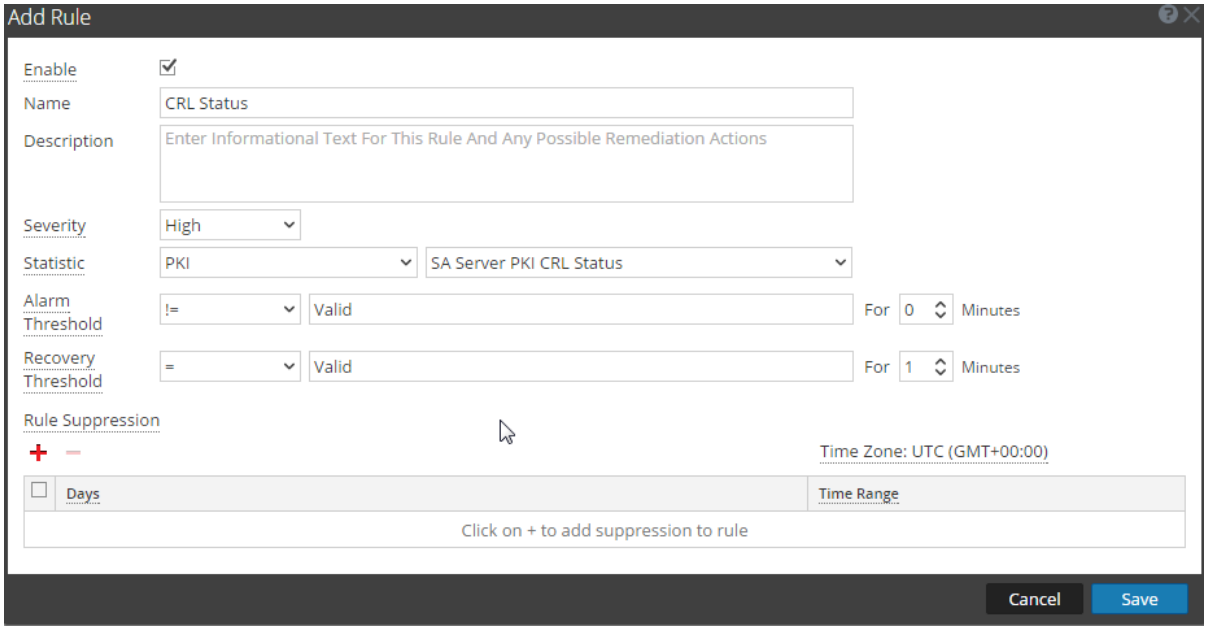
   - For CA Expiration



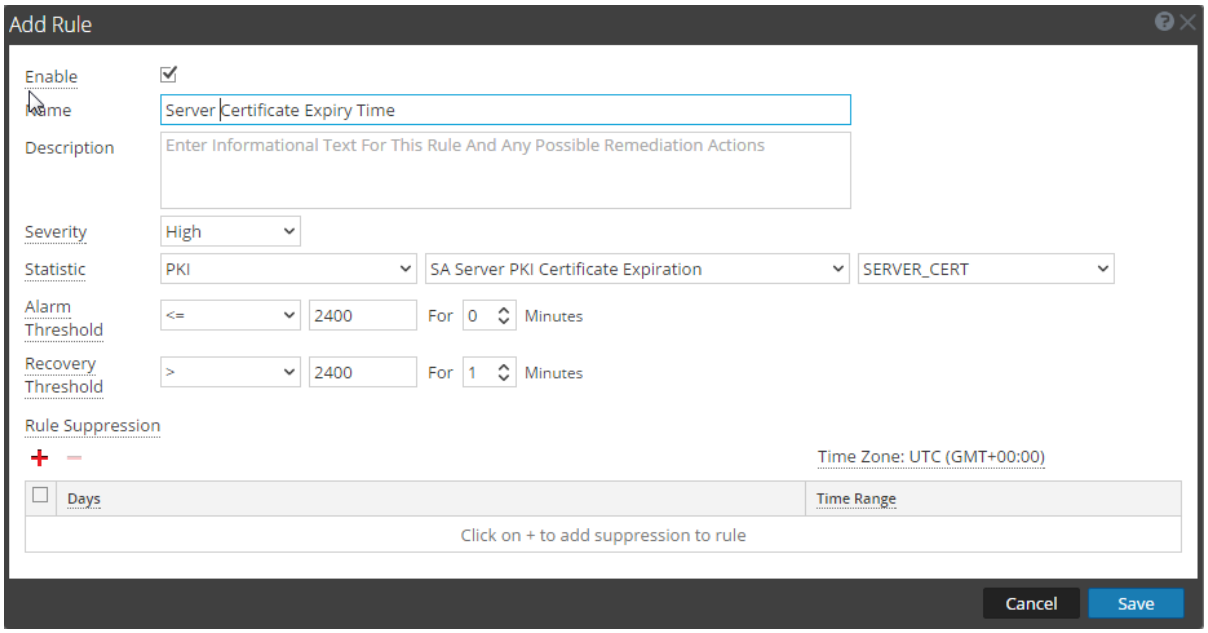
   - For CRL Expiration



- For CRL Status



- For Server Certificate Expiration



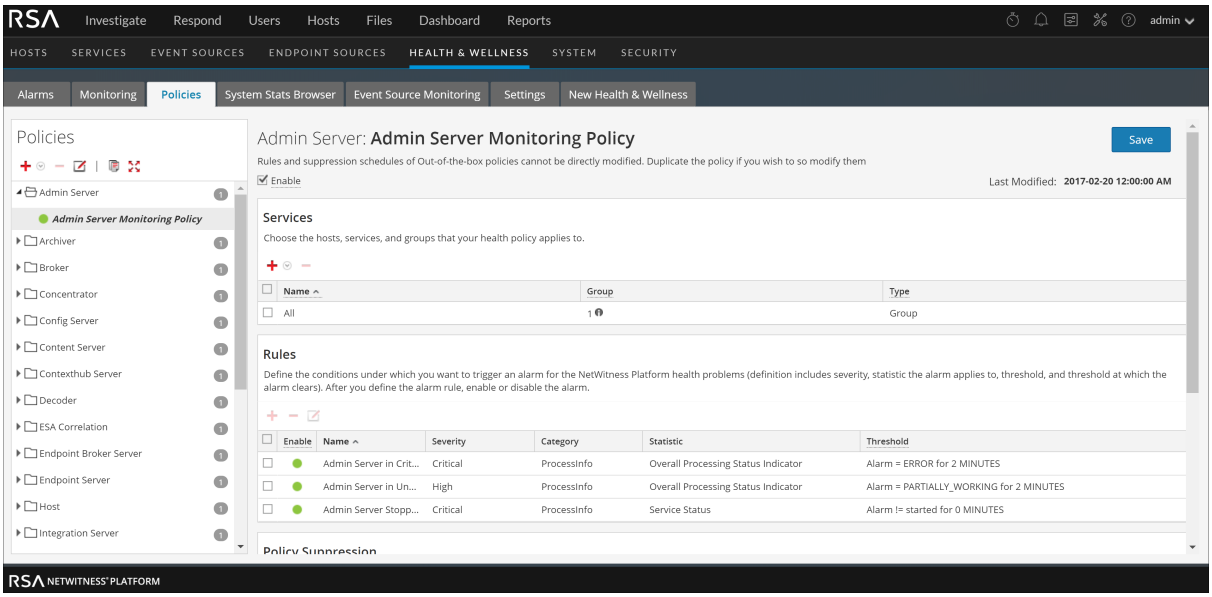
Edit a Policy

1. Go to **(Admin) > Health & Wellness.**
2. Click the **Policies** tab.

   The Policies view is displayed.
3. Select a policy (for example, **Concentrator Policy Status**) under a host or service.

   The Policy Detail is displayed.
4. Click .

   The policy name (for example, **Admin Server Monitoring Policy**) and policy detail panel become editable.

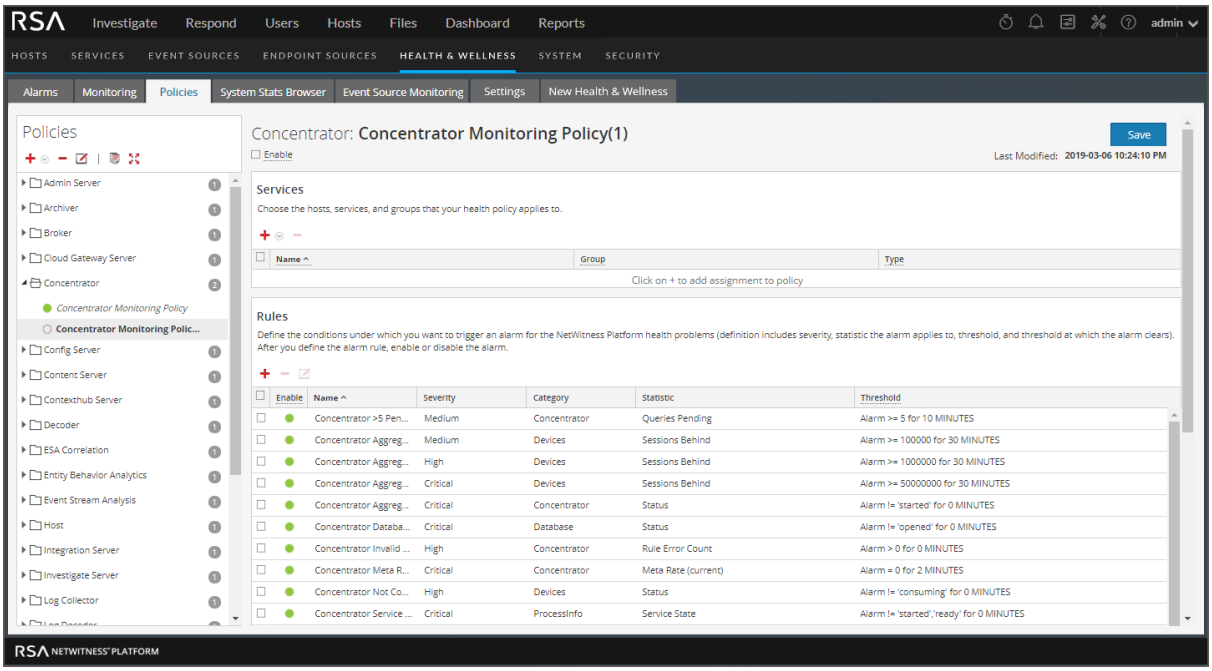


5. Make the required changes and click **Save** in the Policy Detail panel. You can:
   - Edit the policy name.
   - Enable or disable the policy.
   - Add or delete hosts and services in the policy.
   - Add, delete or modify rules in the policy.
   - Add, edit, or delete suppressions in the policy.
   - Add, edit, or delete notifications in the policy.

Note: Save applies the policy rules based on the selection of enable or disable. It also resets the rule condition timers for changed rules, and the entire policy.

Duplicate a Policy

1. Go to  **(Admin) > Health & Wellness**.
2. Click the **Policies** tab.
3. Select a policy (for example, **Concentrator Monitoring Policy**) under a host or service.
4. Click .NetWitness copies the policy and lists it with **(1)** appended to the original policy name.



5. Click  and rename the Policy [for example, rename **Concentrator Monitoring Policy(1)**] to **New Concentrator Policy**.

> **Note:** A duplicated policy is disabled by default and the host and service assignments are not duplicated. Assign any relevant hosts and services to the duplicated policy before you use it to monitor health and wellness of the NetWitness infrastructure.

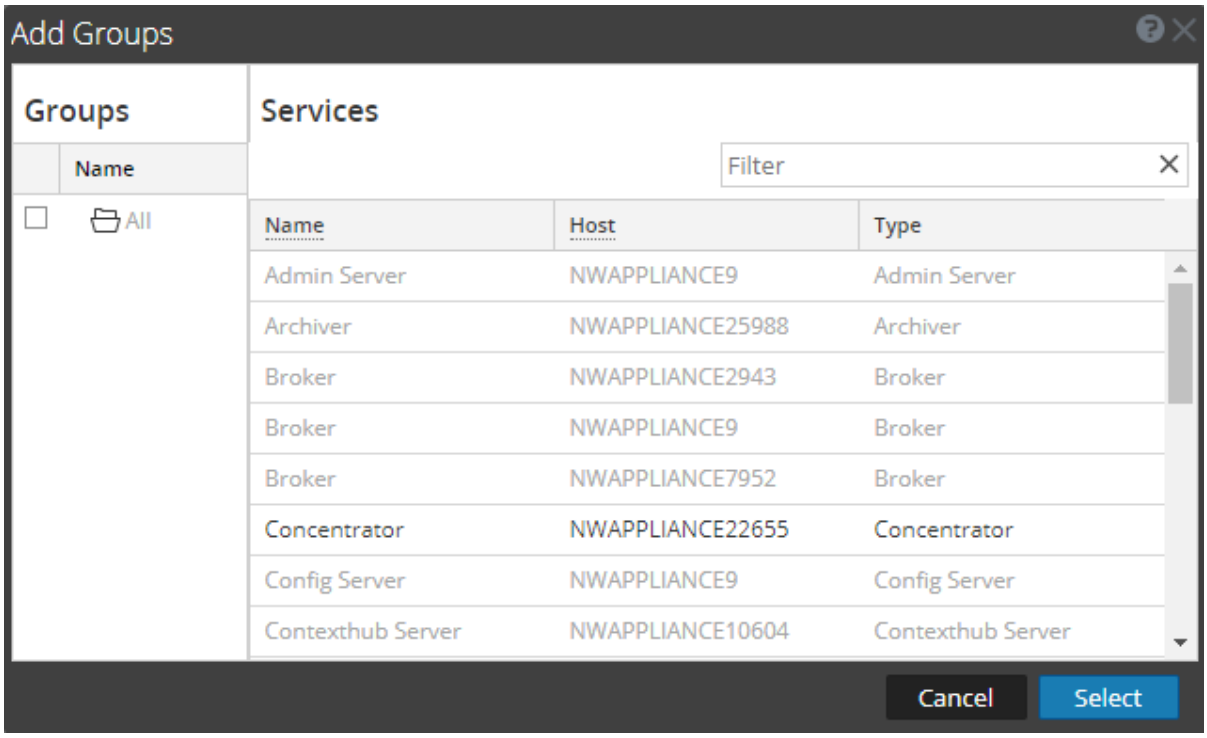
Assign Services or Groups

To assign hosts or services to a policy:

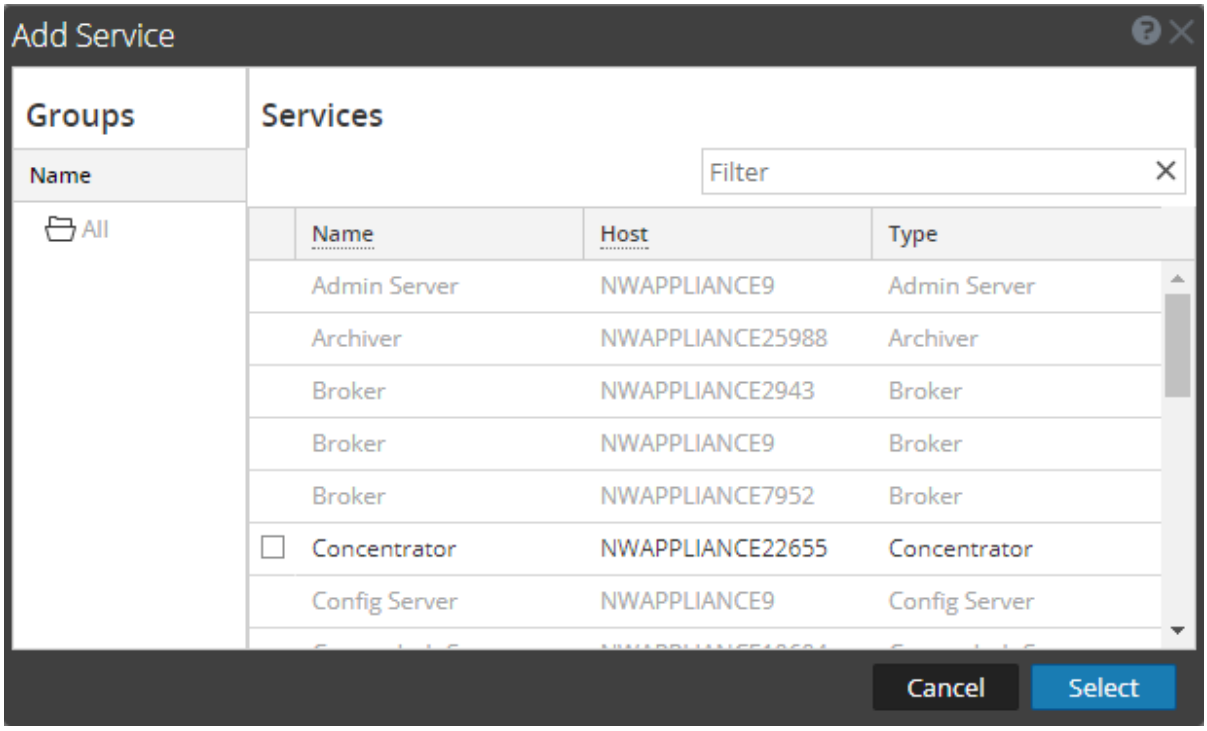
1. Go to  **(Admin) > Health & Wellness.**
2. Click the **Policies** tab.

   The Policies view is displayed.
3. Select a policy (for example, **First Policy**) under a host or service.

   The Policy Detail view is displayed.
4. Click  in the Services and Groups list toolbar.
5. Choose one of the following actions:
   - For hosts, select **Groups** or **Hosts** from the selection menu.
   - For services, select **Groups** or **Services** from the selection menu.
6. Depending on whether you are assigning services or groups, perform one of the following actions:
   - **Groups**, the **Groups** dialog is displayed from which you can select predefined groups of hosts or services.



- **Services**, the **Services** dialog is displayed from which you can select individual services.



7. Select the checkbox next to the groups or services you want to assign to the policy, click **Select** in the dialog, and click **Save** in the Policy Detail panel.

Note: Services are filtered for selection based on the type of policies. For example, you can only select Concentrator services for a Concentrator type of policy.

Remove Services or Groups

To remove a host or service from a policy:

1. Go to **(Admin) > Health & Wellness**.
2. Click the **Policies** tab.

   The Policies view is displayed.
3. Select a policy under a service.

   The Policy Detail view is displayed.
4. Select a host or service.
5. Click .

   The host or service is removed from the policy.

Add or Edit a Rule

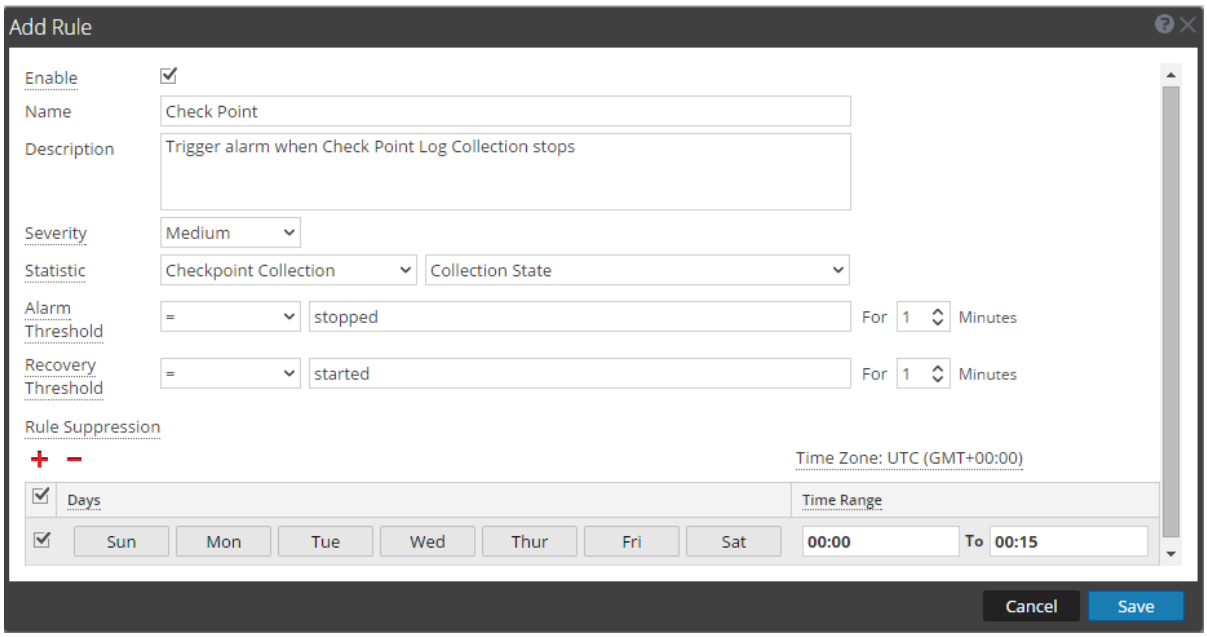
To add a rule to a policy:

1. Go to **(Admin) > Health & Wellness.**
2. Click the **Policies** tab.

   The Policies view is displayed.
3. Select a policy (for example, **Checkpoint**) under a host or service.

   The Policy Detail view is displayed.
4. Depending on whether you are adding or editing rule, do the following:

   - To add a rule, click in the Rules list toolbar.
   - To edit a rule, select a rule from the Rules list and click .
5. Complete the dialog to define or update the rule.
6. Add a description as shown in the following example.



7. Click **OK**.

   The rule is added (or updated) to the policy.

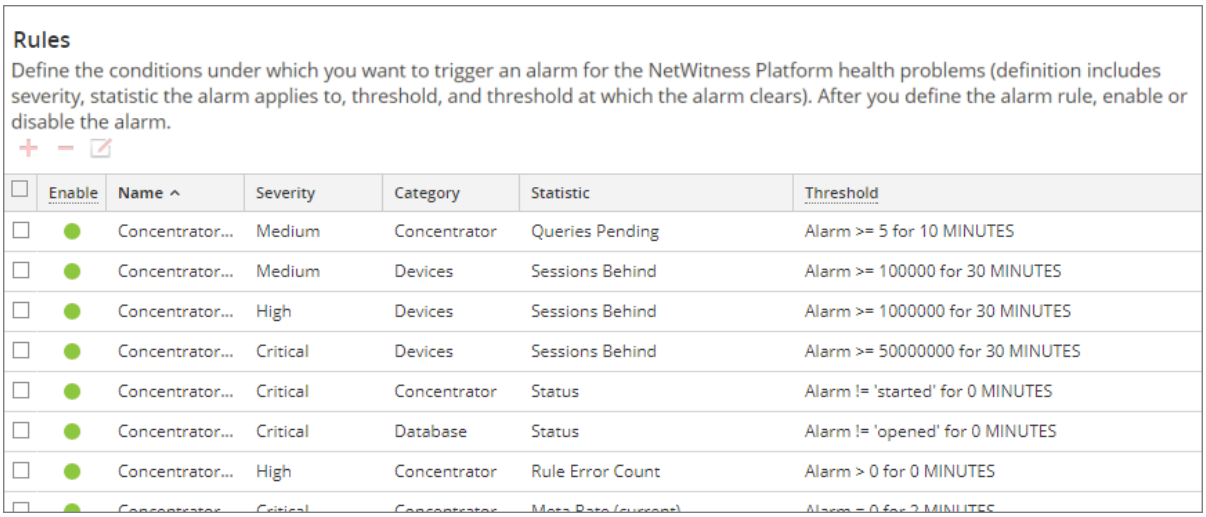
Hide or Show Rule Conditions Columns

To hide or show rule conditions columns in the Rules panel:

1. Go to **(Admin) > Health & Wellness.**
2. Click **Policies** tab.

   The Policies view is displayed.
3. Select a policy under a service.

   The Policy Detail view is displayed.
4. Go to the **Rules** panel.



5. Click **v** to the right of **Category** , set **Columns**, and clear the **Static** and **Threshold** rule conditions.

   You can set or clear any Rules column to show or hide it.
   The **Rules** panel displays without the rule conditions.

Delete a Rule

To remove a host or service from a policy:

1. Go to **(Admin) > Health & Wellness.**
2. Click the **Policies** tab.
   The Policies view is displayed.
3. Select a policy under a service.
   The Policy Detail view is displayed.
4. Select a rule from the **Rules** list (for example, **Checkpoint**).
5. Click **-**.
   The rule is removed from the policy.

Suppress a Rule

1. Click the **Policies** tab.
   The Policies view is displayed.
2. Select a policy under a service.
   The Policy Detail view is displayed. You can specify rule suppressions time ranges when you initially add it or you can edit the rule and specify suppression time ranges.
3. Add or edit a rule.
4. In the **Rules Suppression** panel of the **Add** or **Edit Rule** dialog, specify the days and time ranges during which you want the rule suppressed.

Suppress a Policy

1. Add or edit a policy.
   The Policies view is displayed.
2. In the **Policy Suppression** panel:
   a. Select a time zone from the **Time Zone** drop-down list.
      This time zone applies to the entire policy (both policy suppression and rule suppression).
   b. Click  in the toolbar.
   c. Specify the days and time ranges during which you want the policy suppressed.

Add an Email Notification

To add an email notification to a policy:

1. Add or edit a policy.
   The Policies view is displayed.
2. In the **Notification** panel:
   a. Click  in the toolbar.
      A blank EMAIL notification row is displayed.
   b. Select the email:
      - Notification types in the Recipient column (see "Configure Notification Outputs" in the *NetWitness System Configuration Guide* for the source of the values in this drop-down list).
      - Notification server in the Notification Server column (see 'Configure Notification Servers" in the *NetWitness System Configuration Guide* for the source of the values in this drop-down list).
      - Template server in the Template column (see "Configure Notification Templates" in the *NetWitness System Configuration Guide*for the source of the values in this drop-down list).

      Note: Refer to Include the Default Email Subject Line if you want to include the default Email subject line from the Health & Wellness template in your Health & Wellness Email notifications for specified recipients.

Delete an Email Notification

To add an email notification to a policy:

1. Add or edit a policy.
   The Policies view is displayed.
2. In the **Notification** panel:
   a. Select an email notification.
   b. Click .
      The notification is removed.

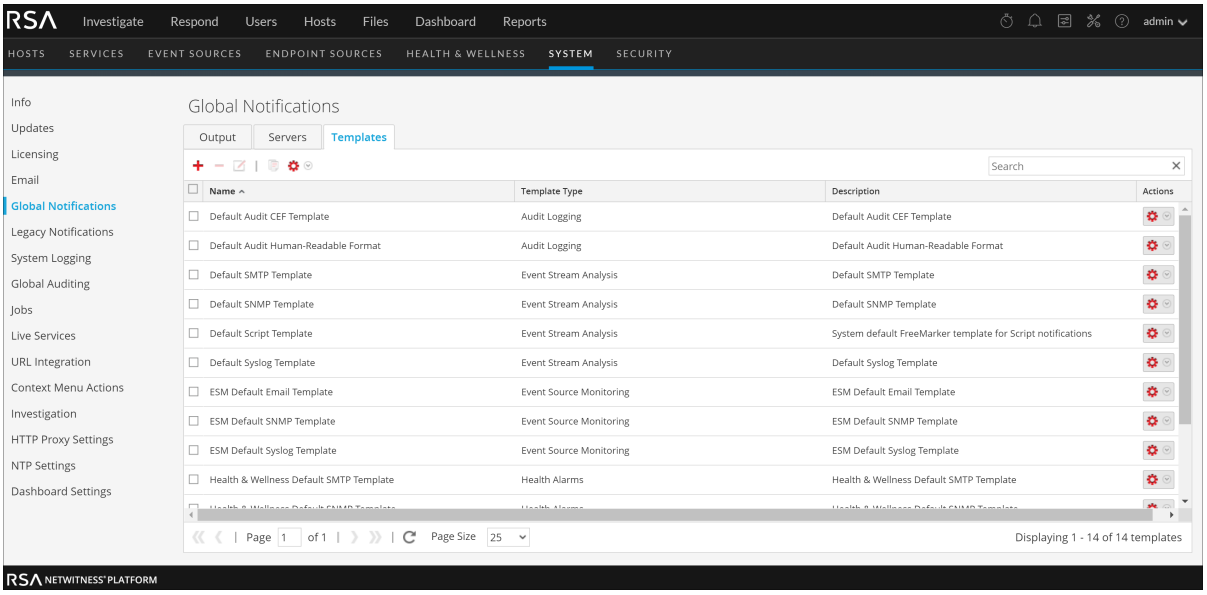
Include the Default Email Subject Line

The emails generated by the notifications you set up for policies do not include the subject line from the Health & Wellness Default Email Notification templates. You need to specify the subject line in the do not include subject lines. This procedure shows you how to insert a subject line into the templates.

For related reference topics, see Policies View and NetWitness Platform Out-of-the-Box Policies.

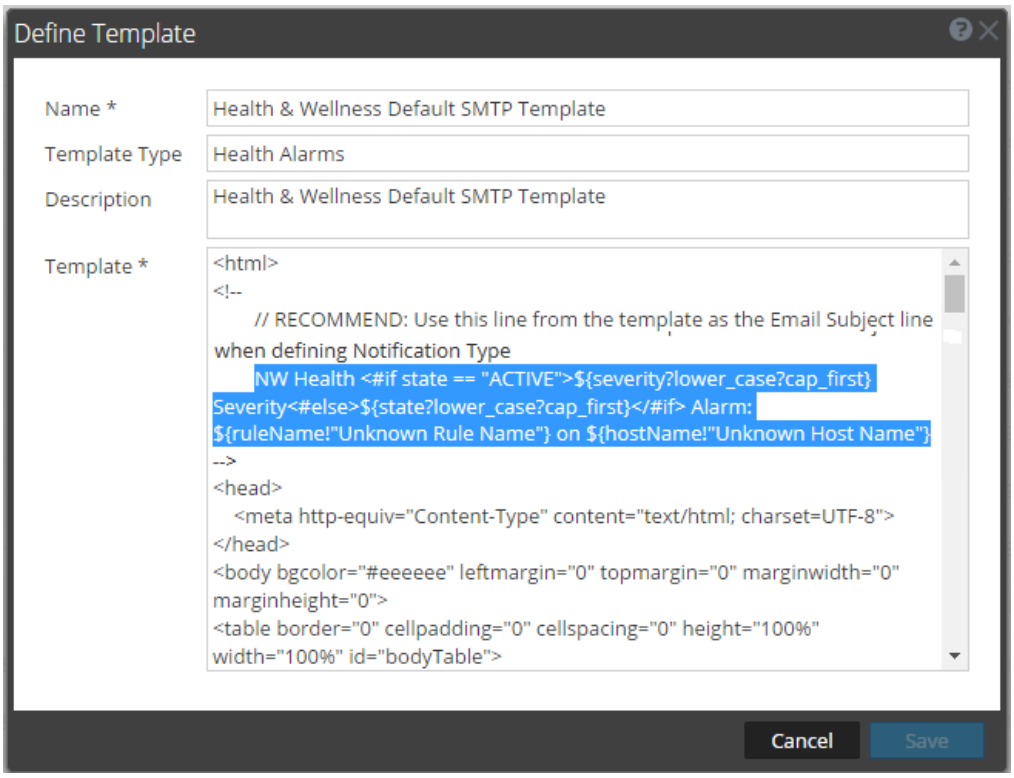
To include the subject line from a Health & Wellness email template in your email notification:

1. Go to  **(Admin) > System**.
2. In the options panel, select **Global Notifications**.
3. Select a Health & Wellness Email Template (for example, **Health & Wellness Default SMTP Template**).



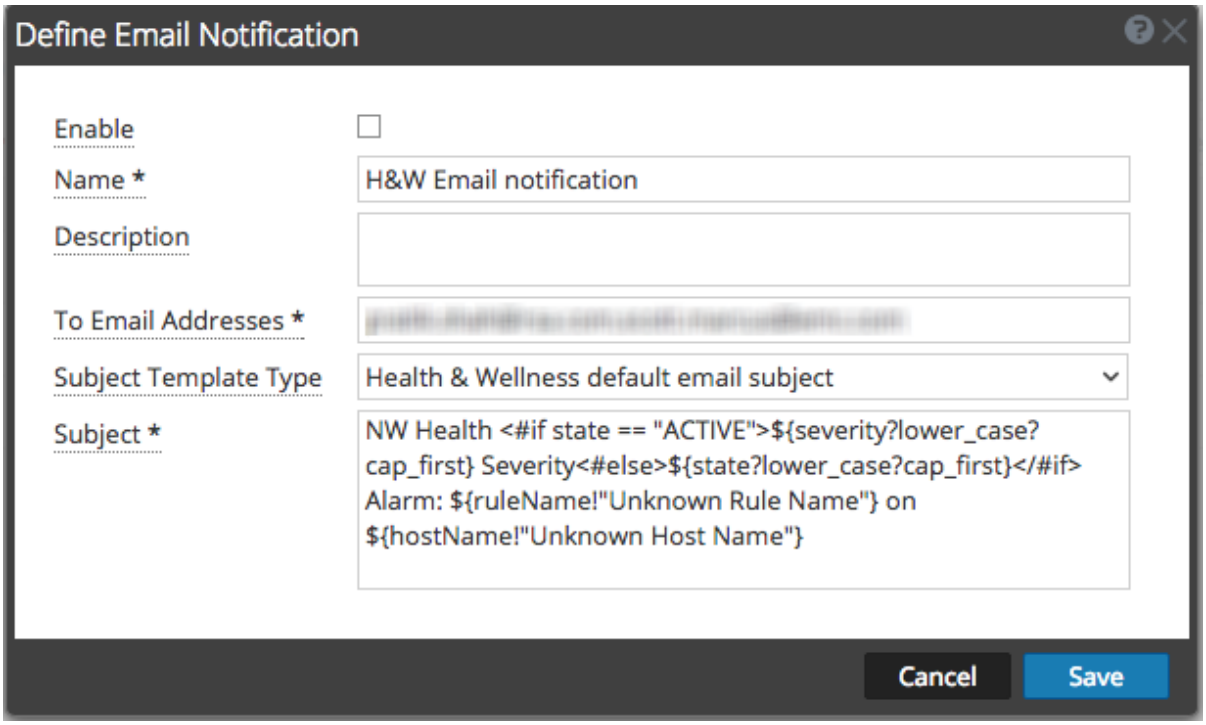
The Define Template dialog is displayed.

4. Click , then in the **Template** field, copy the Subject Line (Highlight the subject line and press Ctrl-C) into the buffer.



5. Click **Cancel** to close the Template.
6. Click the **Output** tab and select a notification (for example **Health & Wellness**).
7. Click .

   The **Define Email Notification** dialog is displayed.
8. Replace the value in **Subject** field text box with the subject line that you have in the buffer (highlight the existing text and press Ctl-V).


Define Email Notification
Enable
Name *
H&W Email notification
Description
To Email Addresses *
Subject Template Type
Health & Wellness default email subject
Subject *
NW Health <#if state == "ACTIVE">${severity?lower_case?cap_first} Severity<#else>${state?lower_case?cap_first}</#if> Alarm: ${ruleName!"Unknown Rule Name"} on ${hostName!"Unknown Host Name"}
Cancel
Save

9. Click **Save**.

Monitor System Statistics

The System Stats Browser filters statistics by the selected host, component running on the host, statistical category, individual statistic, or any combination of host, component, category, and statistic. You can also choose the order in which to display this information.

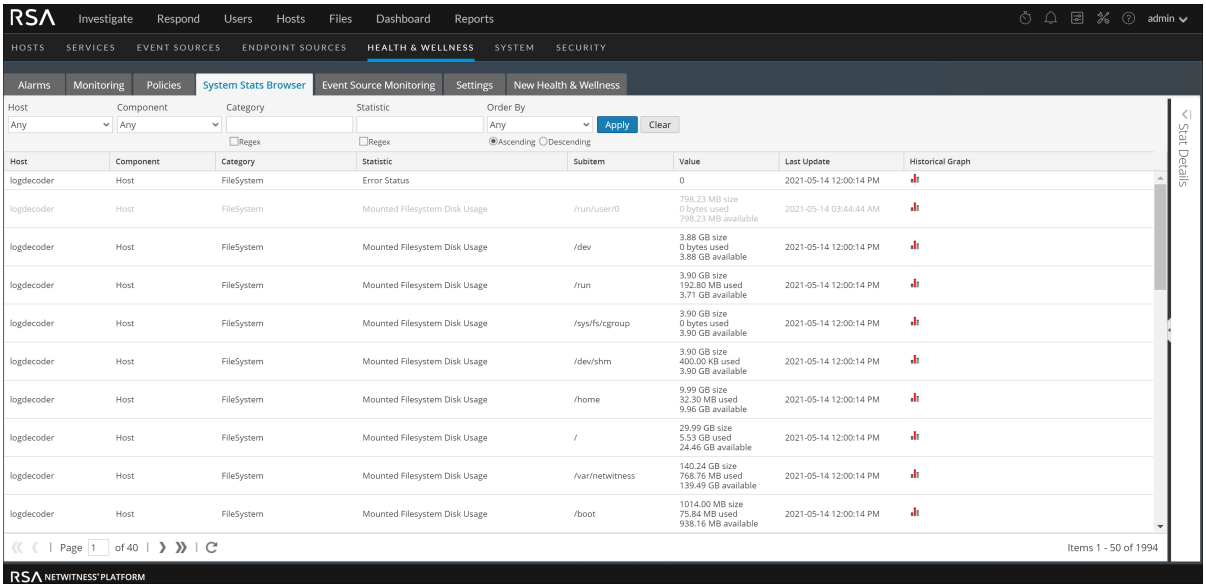
To access the System Stats browser:

1. Go to **(Admin)** > **Health & Wellness**.

   The Health & Wellness view is displayed with the Alarms tab open.

2. Click the **System Stats Browser** tab.

   The System Stats Browser tab is displayed.



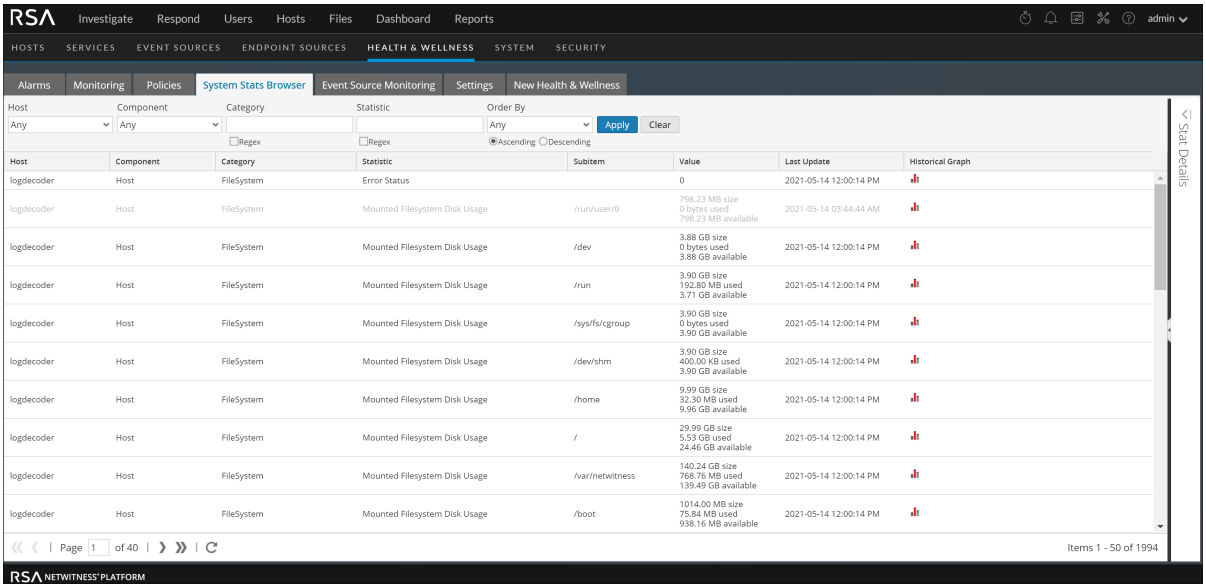
Filter System Statistics

You can filter system statistics to monitor:

- Statistics collected for a particular host
- Statistics collected for a particular component
- Statistics collected of a particular type or that belong to a certain category
- Statistics listed in an ordered way as per the selection chosen

To filter the list of system statistics:

1. Go to  **(Admin) > Health & Wellness**.
   The Health & Wellness view is displayed with the Alarms tab open.
2. Click **System Stats Browser**.
   The System Stats Browser tab is displayed.

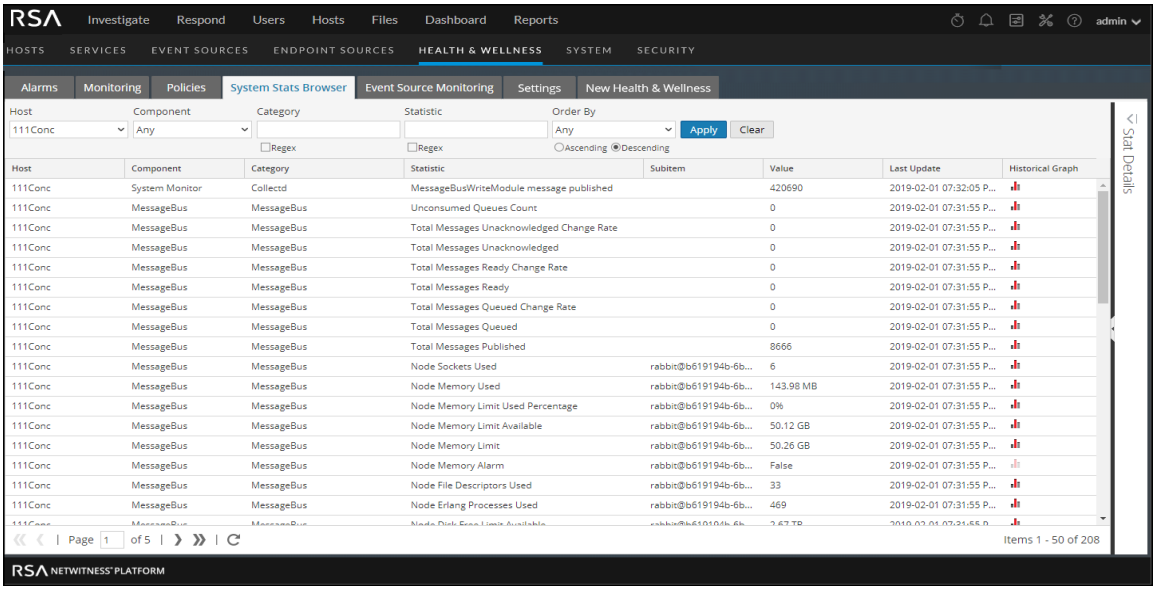


Filter the list of system statistics in one of the following ways:

- To view system statistics for a particular host, select the host in the **Host** drop-down list.
  The system statistics for the selected host is displayed.
- To view system statistics for a particular component, select the component in the **Component** drop-down list.
  The system statistics for the selected component are displayed.
- To view system statistics for a particular category, type the category name in the **Category** field. Select **Regex** to enable the Regex filter. It performs a regular expression search against text and lists the specified category. If Regex is not selected, it supports globbing pattern matching.
  The System Stats for the selected category is displayed.
- To list statistics in a preferred order, you can set the order in the **OrderBy** column.

- To view a particular statistic across hosts, type the statistic name in the **Statistic** field. Select **Regex** to enable the Regex filter. It performs a regular expression search against text and lists the specified category. If Regex is not selected it supports globbing pattern matching. The information for the selected statistics is displayed.

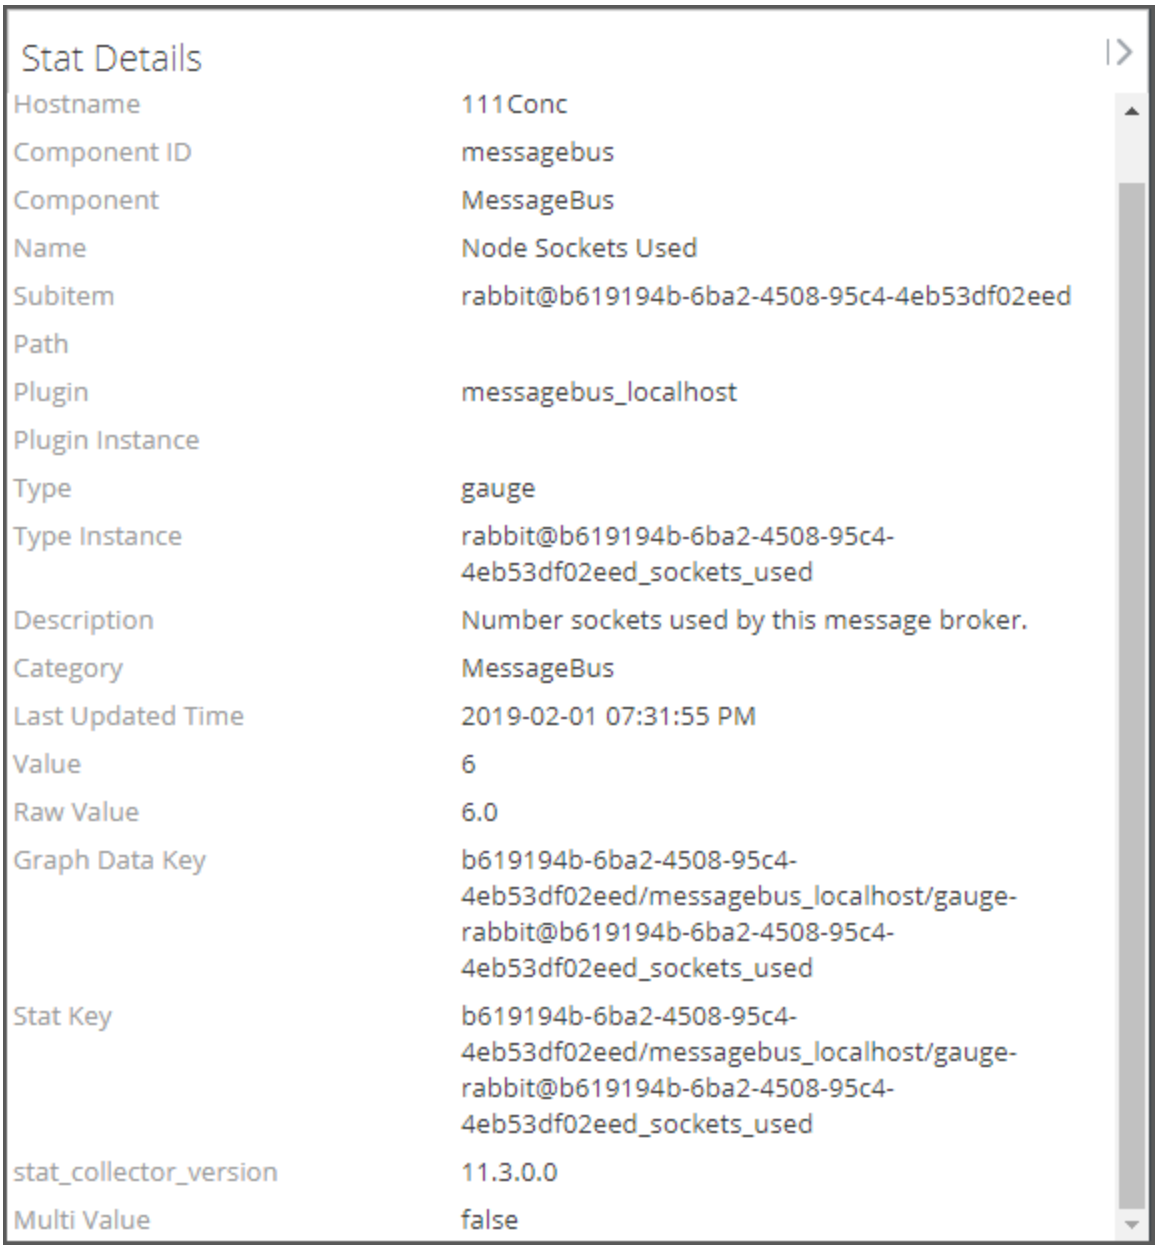
  The following figure shows the System Stats Browser filtered by the 111Conc host listed in descending statistical category order.



3. To view the details for an individual statistic:

   a. Select a row to select a statistic.

   b. Click .
   The Stat Details pane is displayed.



For details on various parameters in the **(Admin)** > **Health & Wellness** > **System Stats Browser** view, see System Stats Browser View

View Historical Graphs of System Statistics

The historical graph of the collected system stats gives you information about the variation of the stats over a selected time frame.

To view a historical graph:

1. Go to **(Admin) > Health & Wellness**.

   The Health & Wellness view is displayed with the Alarms tab open.

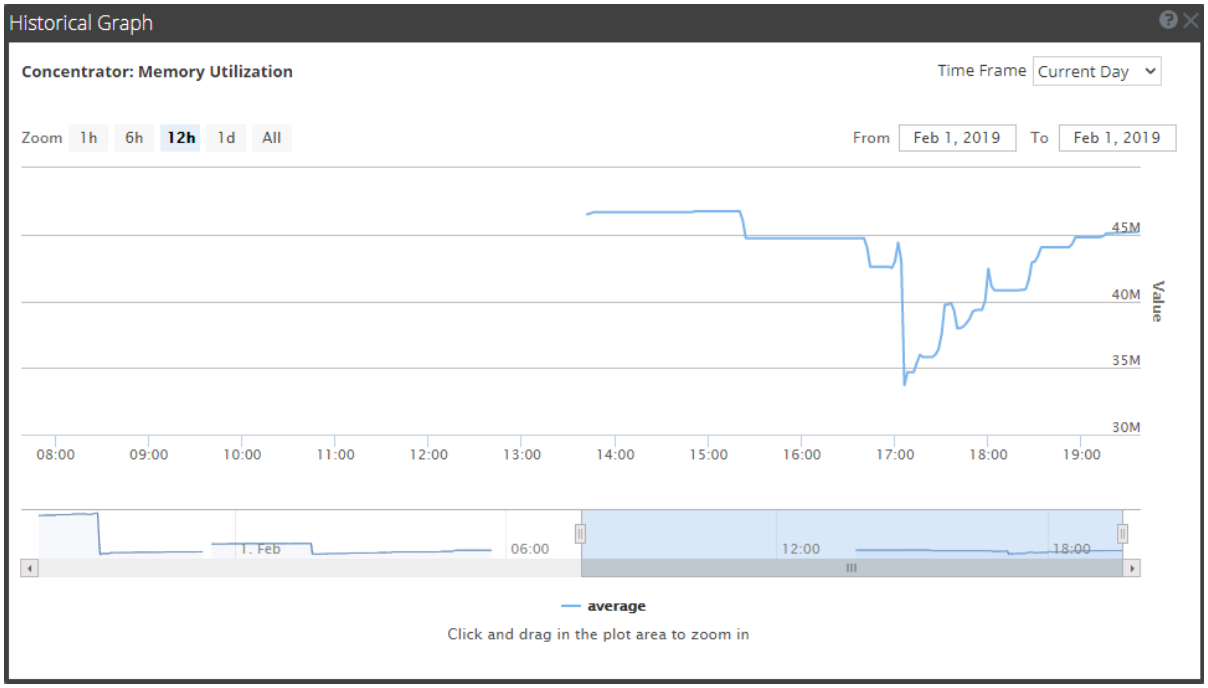
2. Click the **System Stats Browser** tab.

3. In the System Stats Browser tab, specify the filter criteria to display the statistics you want.

4. In the **Historical Graph** column, select .

   The Historical graph for the selected statistic is displayed.

   The figure below gives an example of the historical graph for the Memory Utilization statistic for a host.



   The graphical view is customized to display the statistics collected for the current day and the values are zoomed in for an interval of an hour (10.15 - 11.15 hrs). Hover over the graph to view the details at a particular instant. For example, in the figure it displays the memory utilization at 12.00 hrs.

Note: You can customize the graph view by selecting the Time Frame and Date range. You can zoom in using the zoom in value, time window, or by just clicking and dragging in the plot area. For details on the parameters to customize and zoom in functions, see Historical Graph for System Stats. Any break or gap in the chart line indicates that the service or host was down during that time.

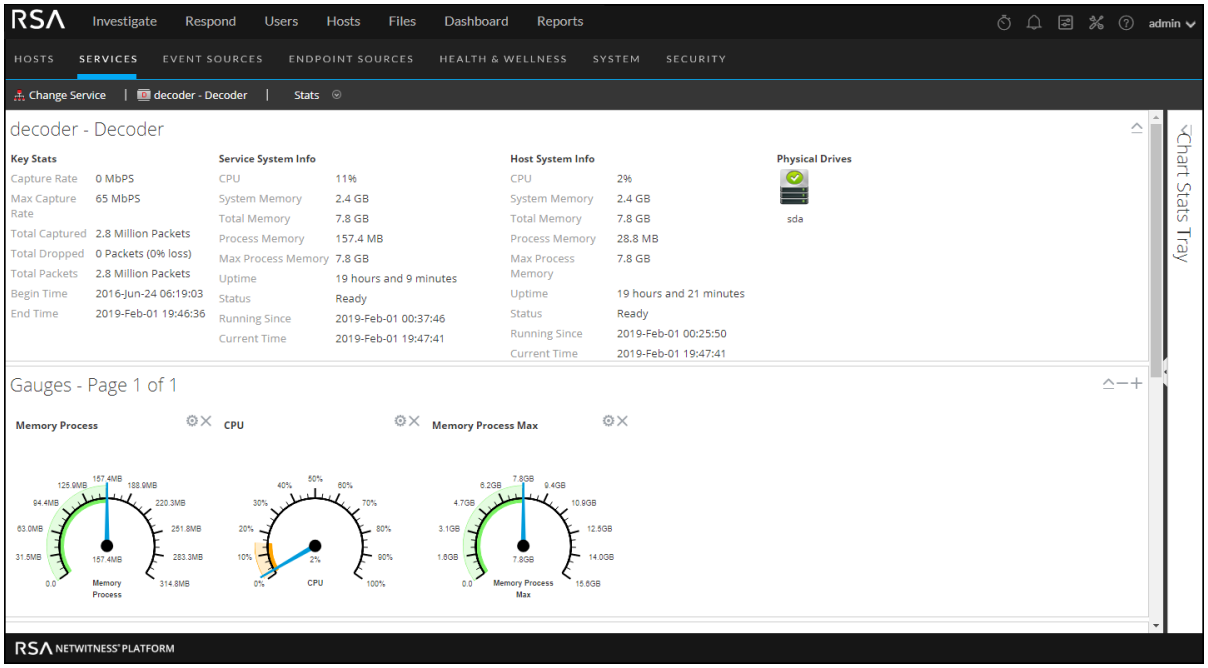
Monitor Service Statistics

NetWitness provides a way to monitor the status and operations of a service. The Service Stats view displays key statistics, service system information, and host system information for a device. More than 80 statistics are available for viewing as gauges and in timeline charts. Only statistics for session size, sessions, and packets are viewable in historical timeline charts.

Although different statistics are available for different types of services, certain elements are common for any Core device.

To monitor service statistics in NetWitness:

1. Go to  **(Admin) > Services**.
   The Services view is displayed.
2. Select a service, and select **View > Stats** in the Actions column.



3. To customize the view, collapse or expand charts. For example, expand the Chart Stats Tray to see available charts, and then drag a section up or down to change the sequence. Or, drag the Gauges section to the top so that it is above the Summary Stats section.

Add Statistics to a Gauge or Chart

In the Services Stats view, you can customize the monitored statistics for individual services. The Chart Stats Tray lists all available statistics for the service. The number of statistics varies according to the type of service being monitored. Any statistic in the Chart Stats Tray can be displayed in a gauge or a timeline chart. Only statistics for session size, sessions, and packets are viewable in historical timeline charts.

Create a Gauge for a Statistic

To create a gauge for a statistic in the Services Stats view:

1. Go to **(Admin) > Services**.

   The Admin Services View is displayed.

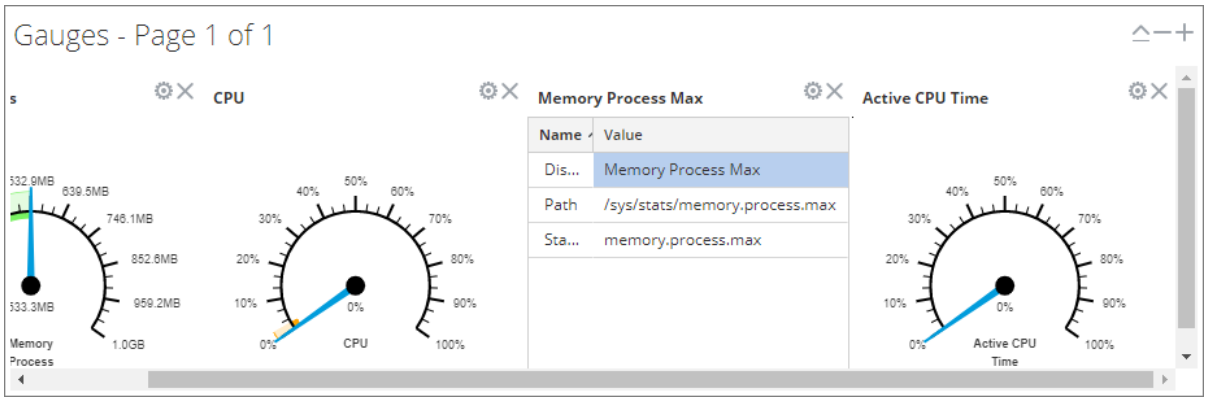
2. Select a service and select **View > Stats** in the Actions column.

   The Chart Stats Tray is displayed on the right side.

3. If the tray is collapsed, click to view the list of available statistics.

4. From the **Chart Stats Tray**, click on any statistic and drag it into the **Gauges** section.

   A gauge is created for the statistic. If there is no space for the gauge, a new page is created in the Gauges section and the gauge is added to the new page. In the example, the Active CPU Time chart was added to the Gauges section by dragging it from the Chart Stats Tray.

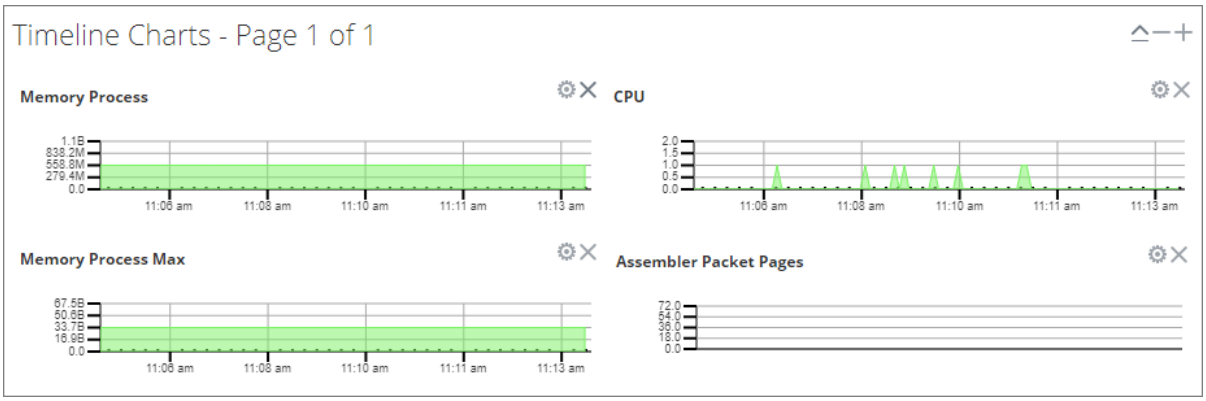


Create a Timeline Chart for a Statistic

To create a timeline for a statistic:

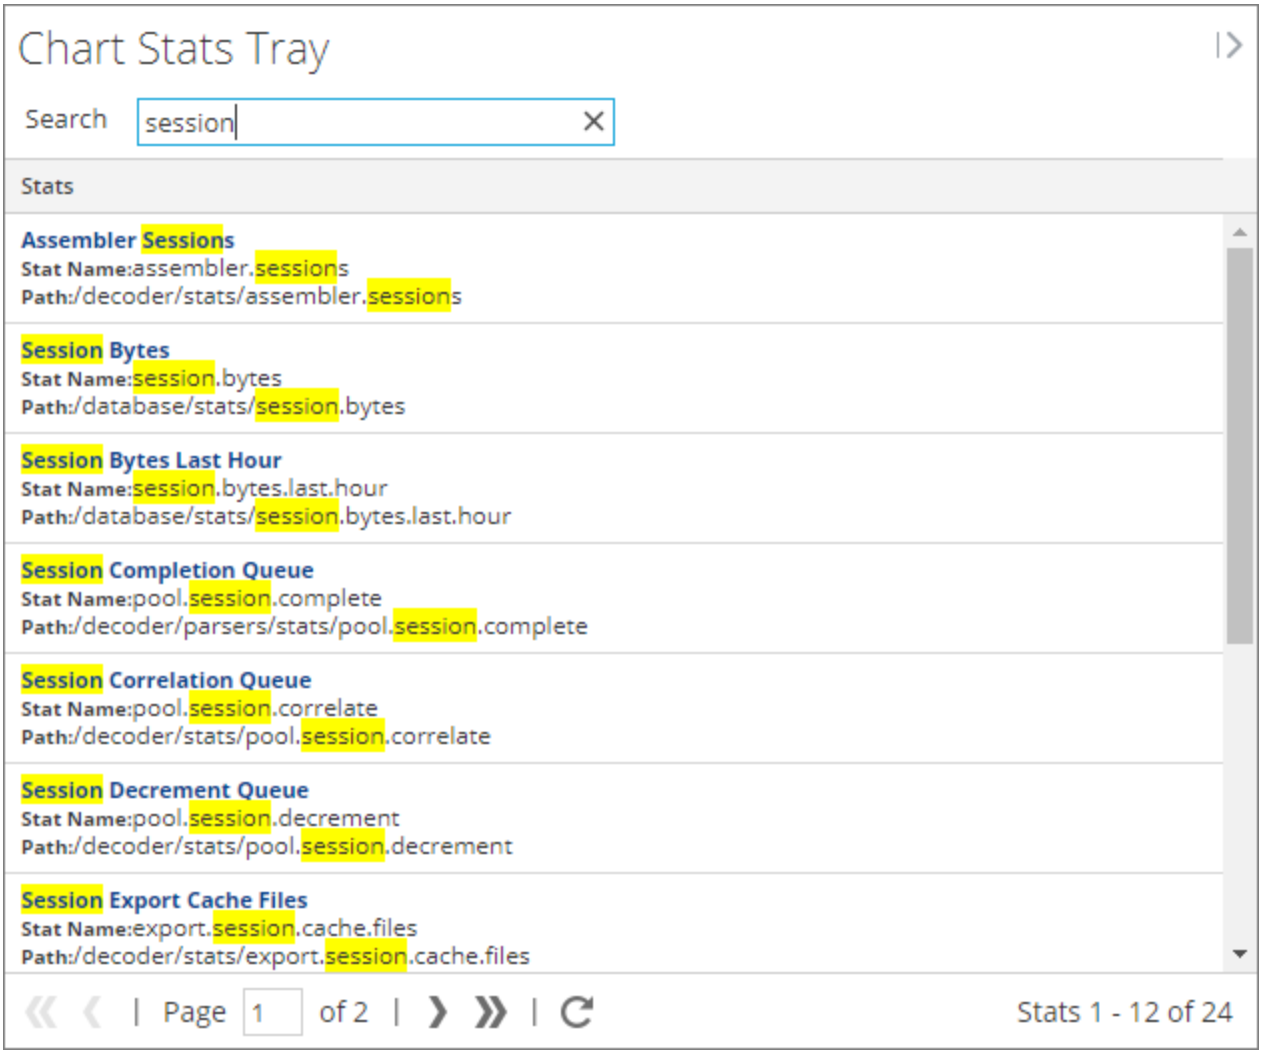
From the **Chart Stats Tray**, click on a statistic and drag it into the **Timeline Charts** or the **Historical Timeline Charts** section.

A timeline chart is created for the statistic. If there is no space for the chart, a new page is created in the Timeline Chart section and the chart is added to the new page. In the example, the Assembler Packet Pages chart was added to the Timeline Charts section by dragging it from the Chart Stats Tray.



Search for a Statistic in the Chart Stats Tray

To search for a statistic, type a search term; for example, **session**, in the Search field and press **Enter**. Statistics that match are displayed with the matching word highlighted.

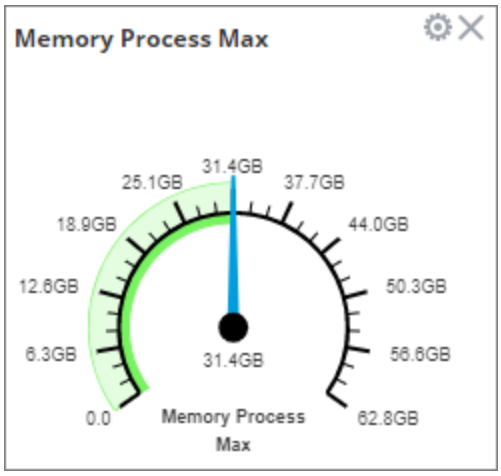


Edit Properties of Statistics Gauges

The Gauges section of the Service Stats view presents statistics in the form of an analog gauge. The properties of each individual gauge are editable; all gauges have an editable title and some have additional editable properties.

Edit Properties of a Gauge

1. Go to **(Admin) > Services**
   The Admin Services view is displayed.
2. Select a service and select **View > Stats** in the Actions column.
   The Service Stats view includes the Gauges section.
3. Go to the gauge for which you want to edit properties (for example, **Memory Process**).



4. Click the Properties icon ( ) to display the parameter names and values.
5. To highlight the value of the **Display Name** field, double-click on the value; for example, **Memory Process**.

   Note: Clicking the other two values does nothing because the properties are not editable in the gauge.

6. Type a new value for the Display Name and click the **Properties** icon ( ).
   The new title replaces **Memory Process**.

Add Stats to the Gauges Section

You can add more gauges by dragging a statistic from the **Chart Stats Tray** into the **Gauges** section.

1. To expand the Chart Stats Tray, click .
2. Scroll down and select a statistic, for example, **Session Rate (maximum)**.

3. Drag the statistic to the **Gauges** section.
   The new gauge is displayed in the Gauges section.

Edit Properties of Timeline Charts

Timeline charts display statistics in a running timeline. The Service Stats view includes two types of timelines; current time and historical. You can drag any statistic available in the Chart Stats Tray to the Timeline Charts section. Only statistics for session size, sessions, and packets are viewable in historical timeline charts. The properties of an individual timeline chart are editable; all timeline charts have an editable title and some have additional editable properties.

To access the charts:

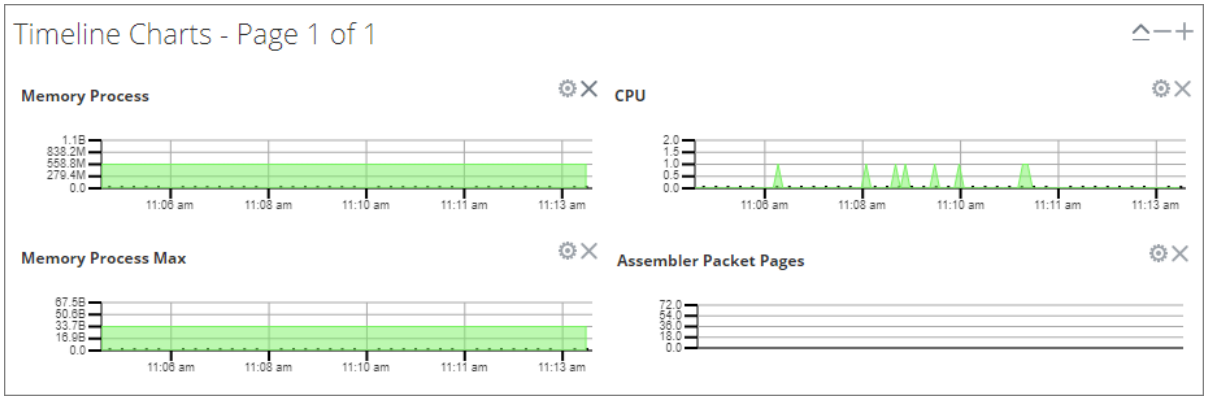
1. Go to **(Admin) > Services**.
2. Select a service and click **Stats**.

   The Services Stats view is displayed. The charts are in this view.

Edit Properties of a Timeline

To edit properties of a timeline chart:

1. Go to the timeline chart for which you want to edit properties (for example, **Memory Process**).



2. Click the **Properties** icon ( ) to display the parameter names and values.
3. Double-click on a value (for example, the **Display Name** field) to make the value editable.

   Note: Clicking the other two values does nothing because the properties are not editable in the chart.

4. Type a new value and click the **Properties** icon.

   The timeline chart is displayed with new values.

Edit Properties of a Historical Timeline

To edit properties of a historical timeline chart:

1. Go to Historical Timeline Charts.
2. Click the **Properties** icon ( ) to display the parameter names and values.

3. Click on a value (for example, **01/27/2019** for the **Begin Date** field) to make the value editable.
4. Type a new value.
5. Edit the **End Date** and **Display Name** if required.
6. Click the **Properties** icon ( ).
   The historical timeline is displayed with new values.

Note: To return the properties of the historical timeline chart back to the default so that the values dynamically update, remove the Begin Date and the End Date, place your cursor in the Begin Date field, and refresh your browser.

Add Stats to Timeline Charts

You can add timeline charts by dragging a statistic from the Chart Stats Tray into the Timelines section.

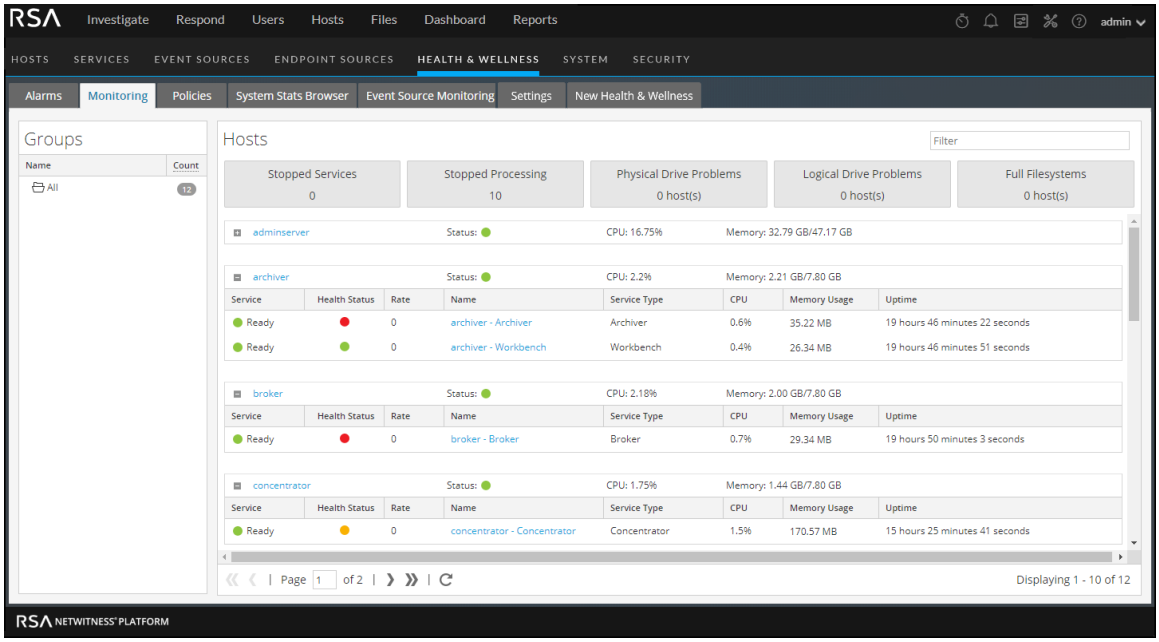
1. To expand the Chart Stats Tray, click .
2. Scroll down and select a statistic; for example, **Session Rate (maximum)**.
3. Drag the statistic to the **Timelines Section**.
   The new timeline is displayed in the Timelines section.

Monitor Hosts and Services

NetWitness provides a way to monitor the status of the hosts and services installed in your environment. You can view the current health of all the hosts, services running on the hosts, their CPU usage and memory consumption, and the host and service details.

To monitor hosts and services in NetWitness:

1. Go to **(Admin) > Health & Wellness**.
   The Health & Wellness view is displayed with the Alarms tab open.
2. Select the **Monitoring** tab.
   A list of all hosts and their associated services that belong to the group **All** is displayed by default.
   The operational status, CPU usage, and memory usage for each host is displayed.



A list of services installed on the host is shown below the host. If you cannot see the services, click ⊞ to the left of a host to display the services.
The name, operating status, CPU usage, memory usage, and the time operating for each service is displayed.

Filter Hosts and Services in the Monitoring View

You can filter hosts and services in the monitoring view in one of the following ways:

- Hosts belonging to a particular group
- A specific host and its associated services
- Hosts whose services are stopped
- Hosts whose services have stopped processing or processing has been turned off
- Hosts that have physical drive problems
- Hosts that have logical drive problems
- Hosts that have full file systems

For the related reference topic, see Monitoring View.

To filter hosts and services:

1. Go to  **(Admin) > Health & Wellness**.

   The Health & Wellness view is displayed with the Alarms tab open by default.
2. Select the **Monitoring** tab.
3. Filter the hosts and services in one of the following ways:
   - To view a list of hosts and their associated services belonging to a particular group, select the group in the Groups panel.

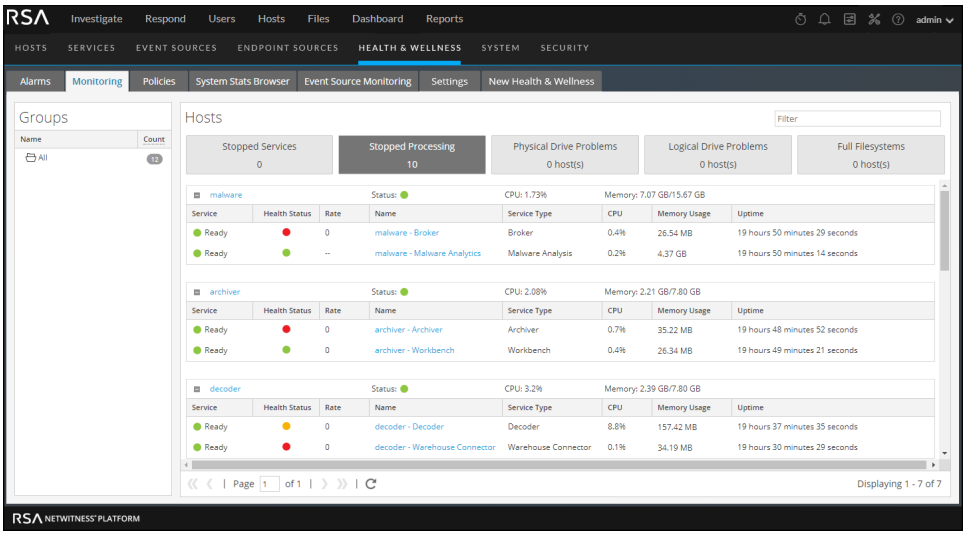
     All hosts and their associated services belonging to the specified group are displayed in the Hosts panel.

     Note: The grouping of hosts is derived from the groups created in the Admin Hosts view. All groups created in the Admin Hosts view are displayed here.

     For example, if you select the group **LC_Group** in the Groups panel, a list of all hosts that are part of the group are displayed.
   - To view a list of all services that have stopped processing, click **Stopped Processing** in the Hosts panel.

     A list of all the hosts that have at least one service with the status as stopped processing is displayed.

     Note: The buttons on the top display the system statistics for all of the hosts configured in NetWitness and does not change with the application of filters on the groups.



- In a similar way, you can filter the list of hosts and the associated services by choosing the appropriate filter:
  - Click **Stopped Services** to display a list of all stopped services.
  - Click **Physical Drive Problems** to display a list of host with physical drive problems.
  - Type the host name in the Filter box to display only the required host and the services running on the host.

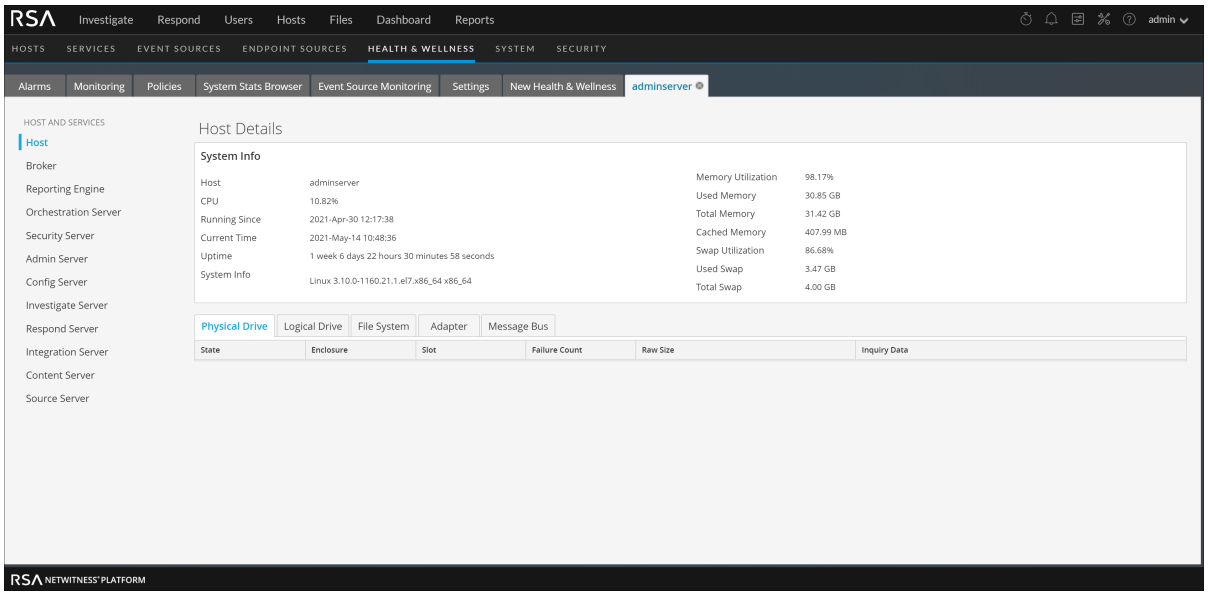
Monitor Host Details

You can view the details of a host, its memory and CPU usage, system information, physical drive, logical drive, and file system details to investigate potential problems with the host.

To view host details:

1. Go to  **(Admin) > Health & Wellness > Monitoring** tab.
2. Click a host in the **Hosts** panel.

   The Host Details view shows important system information about the selected host, such as memory utilization and file system usage.



Monitor Service Details

You can view the details of a service, its memory and CPU usage, system information, and various details depending on the service selected.

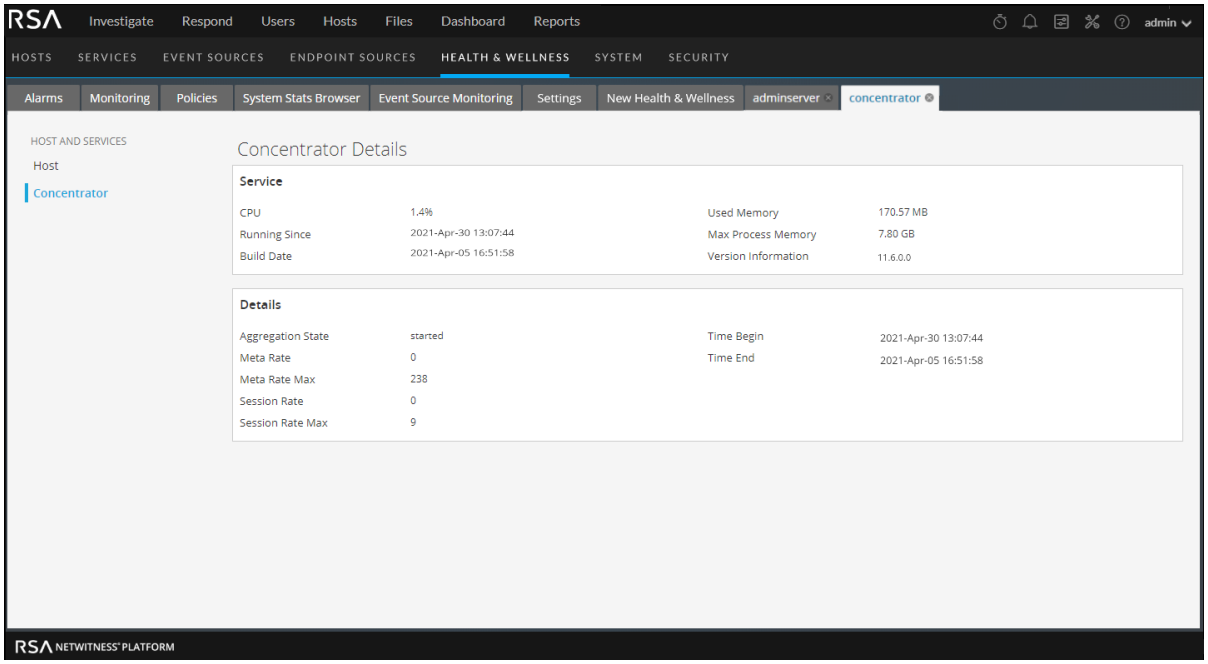
To view service details:

1. Go to **(Admin) > Health & Wellness > Monitoring** tab.

   The Hosts panel shows the services running on each host.

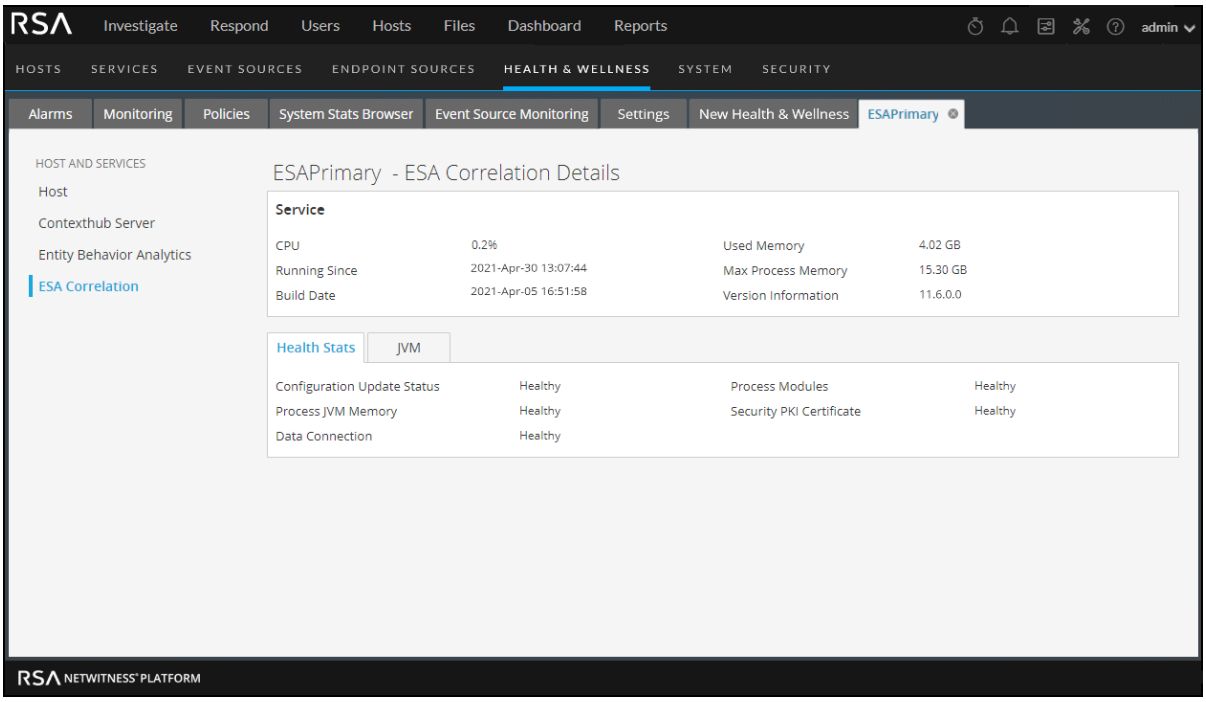
2. In the Hosts panel, click a service name link to get more information.

   The service details view shows the health status of the selected service. The Decoder service details include capture statistics and the Concentrator and Broker details include aggregation statistics.

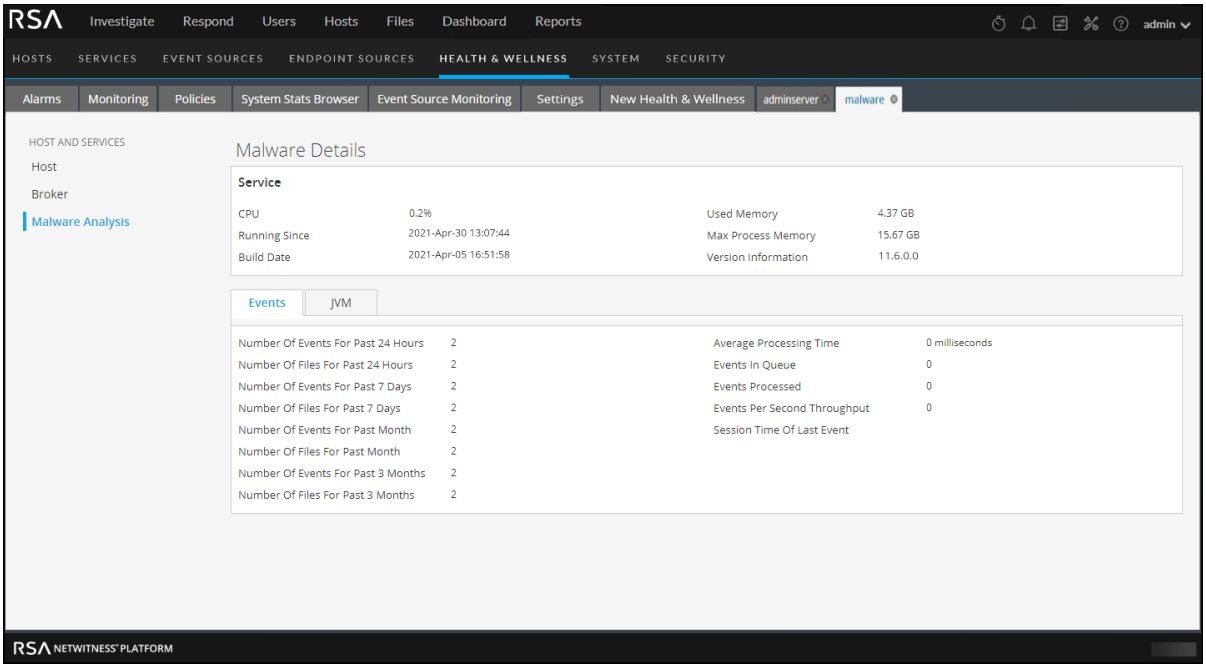


Many services, including the ESA Correlation service, have a **Health Stats** tab that provides information about the health status of the service. The **JVM tab** shows the total memory used by the selected service and the total memory capacity of the host. For more information, see Health Stats Tab and JVM Tab.

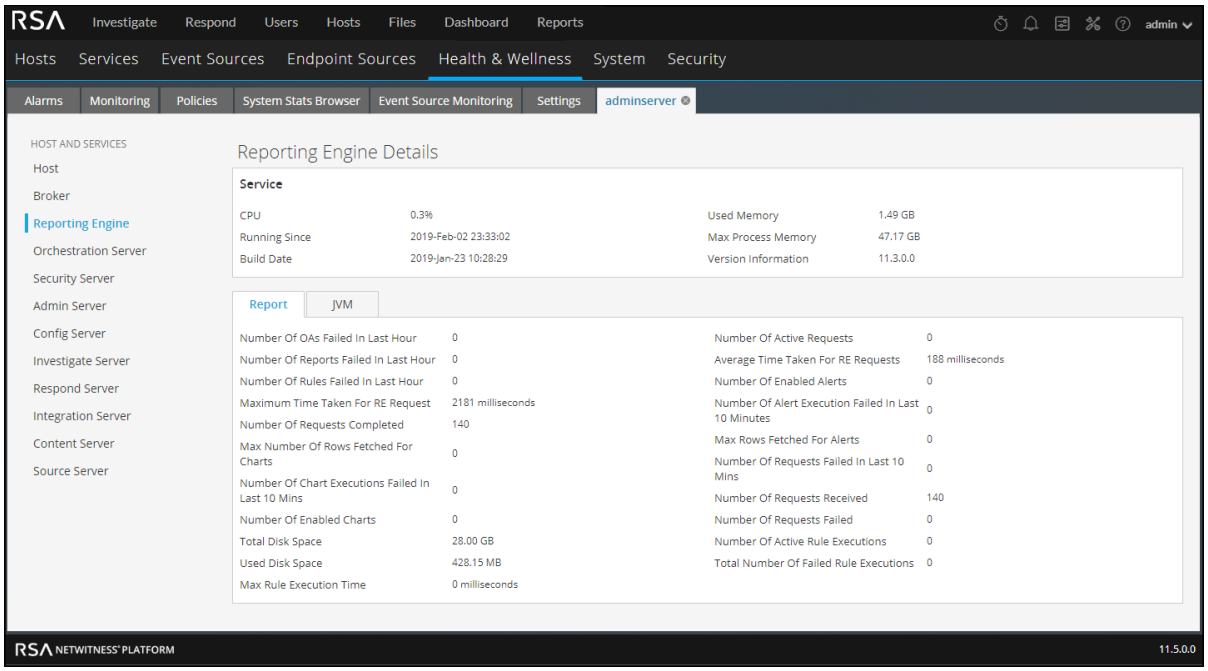
For more information on the ESA Correlation service and ESA Rule memory usage, see the *Alerting with ESA Correlation Rules User Guide*.



The Malware Analysis service details view has **Service** information plus the **Event,** and **JVM** tabs that show additional statistics. The **Events** tab shows event processing statistics.



The Reporting Engine service details view has **Service** information plus the **Report** and **JVM** tabs that show additional statistics.



You can also view the details of other services by clicking the services listed in the options panel on the left.

Refer to Monitoring View for a detailed description of the Details view for each service.

Monitor Event Sources

Note: For NetWitness 11.4.1, this view has been deprecated. To manage Event Sources, use the  (Admin) > Event Sources view. For details, see "About Event Source Management" in the*NetWitness Event Source Management Guide*.

Monitor Alarms

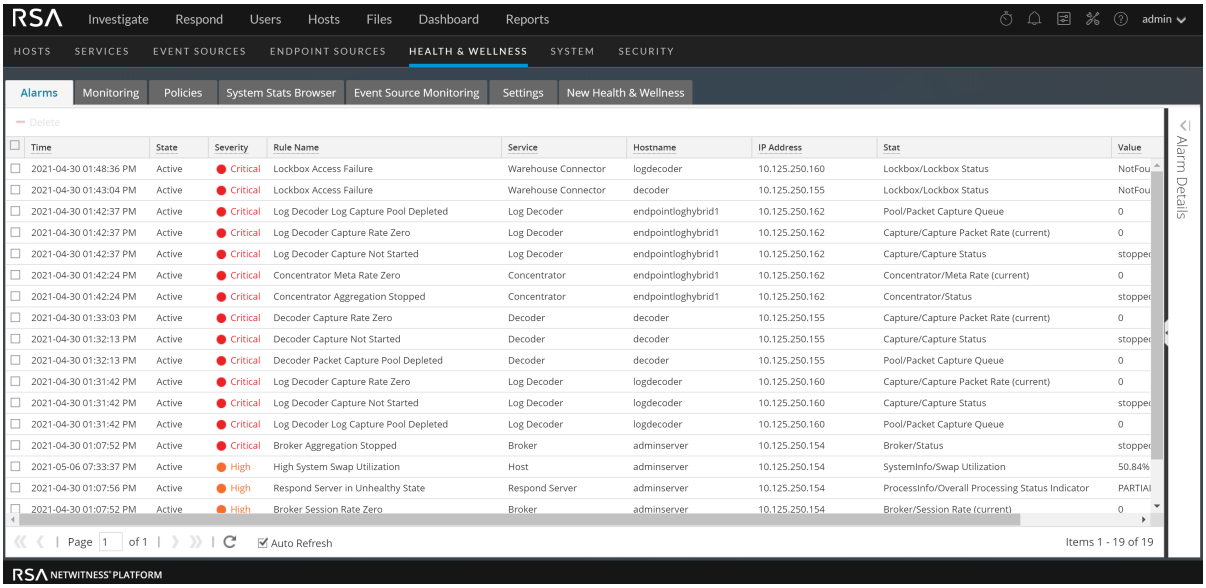
You can set up alarms and monitor them in the Health and Wellness interface for the hosts and services in your NetWitness domain. Alarms display in the view as **Active** when the statistical thresholds for hosts and services have been crossed. Alarms are grayed out and change to the **Cleared** status when the clearing threshold has been crossed.

You set up the parameters for alarms in Manage Policies. For the related reference topic, see Health and Wellness View - Alarms View.

To monitor alarms:

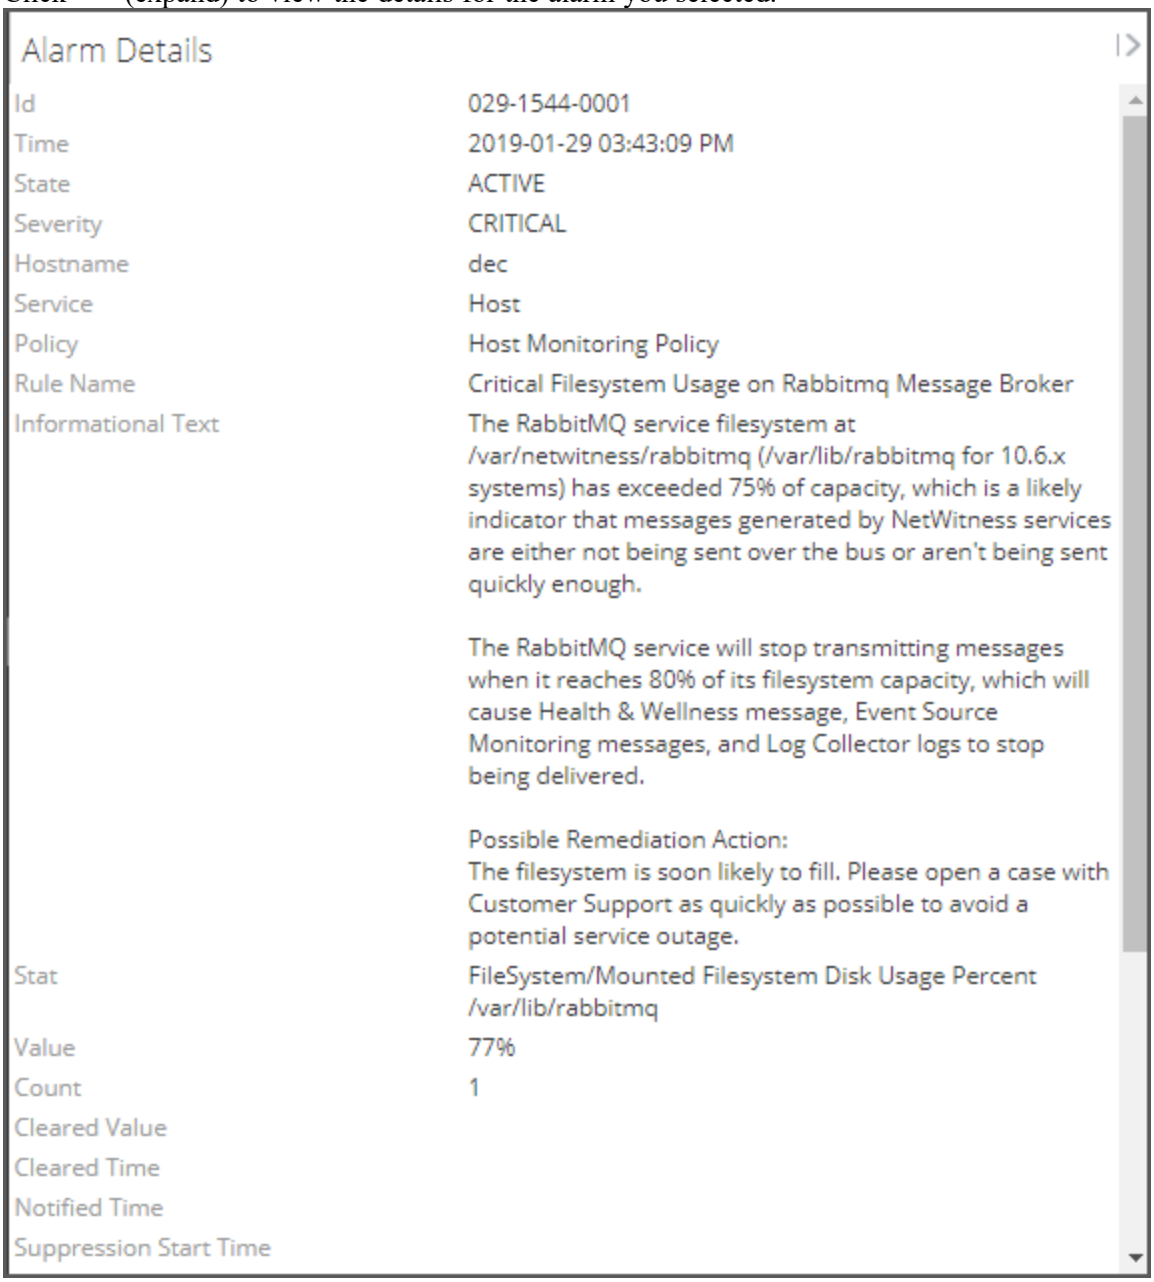
1. Go to  **(Admin) > Health & Wellness**.

   The Health & Wellness view is displayed with the Alarms tab open by default.



2. Click on the alarm for which you want to display details in the Details Panel.

3. Click (expand) to view the details for the alarm you selected.



Monitor Health and Wellness Using SNMP Alerts

You can monitor a NetWitness component to proactively send alerts, using Simple Network Management Protocol (SNMP) that is based on thresholds or system failures.

You can monitor the following for NetWitness components:

- CPU utilization that reaches a defined threshold
- Memory utilization that reaches a defined threshold
- Disk utilization that reaches a defined threshold

SNMP Configuration

NetWitness Servers can be configured to send out SNMPv3 threshold traps and monitor traps. Threshold traps are sent in conjunction with node thresholds that are configured by the NetWitness Core applications. Monitor traps are sent by the SNMP daemon for the items indicated in the SNMP configuration file. You must set up the SNMP daemon on another service to receive SNMP traps from NetWitness. You can set up SNMP on NetWitness in the configuration setting for the NetWitness Server. For more information, see "Service Configuration Settings" in the *NetWitness Host and Services Getting Started Guide* for a specific type of host.

Thresholds

Thresholds can be set on any service statistics that can accept the `setLimit` message. You can retrieve current thresholds using the `getLimit` message. To set a limit, you can pass a low and high threshold value.

When the value of a statistic crosses either the low or high threshold, an SNMP trap is triggered, indicating that the threshold has been crossed. The trap is not triggered if the value is below the low and above the high value, but another trap is triggered if it crosses back into the normal range (above the low and below the high).

You must set the threshold for the service using the Service Explorer view or the REST API.

This example shows a sample threshold for monitoring CPU usage (below 10% or above 90%):

```
/sys/stats/cpu setLimit low=10 high=90
```

This example shows how the threshold is set using REST API:

```
http://<log decoder>:50102/sys/stats/cpu?msg=setLimit&low=10&high=90
```

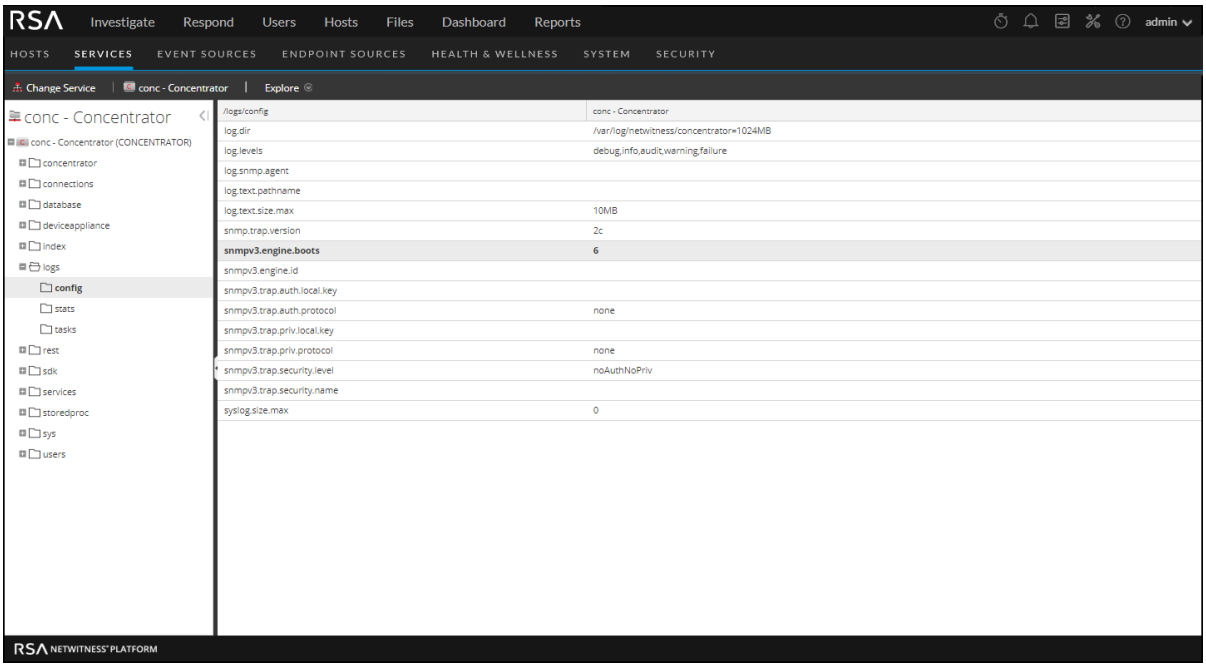
If the CPU usage spikes to 90% or higher, an SNMP trap is generated:

```
23435333 2018-Dec-16 11:08:35 Threshold warning path=/sys/stats/cpu old=77%
new=91
```

Configure SNMPv3 for a Host

1. Go to  **(Admin) > Services**.
   The Services view is displayed.
2. Select the service.
3. In the Actions column, select **View > Explore**.

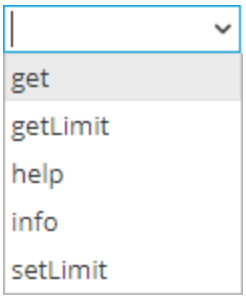
4. In the nodes list, expand the list and select a configuration folder. For example, **logs** > **config**
5. Set the SNMPv3 configuration.



Set the Threshold for a Service

1. Go to  **(Admin) > Services**.
   The Services view is displayed.
2. Select the service.
3. In the Actions column, select **View > Explore**.
4. In the nodes list, expand the list and select a stat folder.
5. Select a stat, for example, CPU, and right-click.
6. From the drop-down menu, select **Properties.**

   The Properties panel is displayed. The Properties panel has a drop-down list of available messages for the parameter.



7. Select **setLimit**.
8. Specify the low and high values.

SNMP Traps for System Status

The threshold mechanism can also be used to monitor string-valued stats generated by Core services. There are two ways to monitor string-valued stats:

1. Generate a trap whenever the status value is NOT an expected value. For example, if you want monitor the stat `/broker/stats/status` and generate a trap whenever the value is not `started`, set the `high` limit on the stat to the expected value. You would use the `setLimit` message on `/broker/stats/status` as follows:
   ```
   setLimit high=started
   ```
2. Generate a trap whenever the status value matches an expected value. This is accomplished by using the `low` limit on the stat. For example, if you wanted generate a trap when the stat `/sys/stats/service.status` has the value `"Initialization Failure"`, you would use the `setLimit` message on `/sys/stats/service.status` as follows:
   ```
   setLimit low="Initialization Failure"
   ```

In both of these scenarios, it is possible to check for multiple values by using a comma-separated list of values to check for.

Troubleshooting Health & Wellness

Issues Common to All Hosts and Services

You may see the wrong statistics in the Health & Wellness interface if:

- Some or all the hosts and services are not provisioned and enabled correctly.
- You have a mixed-version deployment (that is, hosts updated to different NetWitness versions).
- Supporting services are not running.

Issues Identified by Messages in the Interface or Log Files

This section provides troubleshooting information for issues identified by messages NetWitness displayed in the Health & Wellness Interface or included in the Health & Wellness log files.

Message	User Interface: **Cannot connect to System Management Service** System Management Service (SMS) logs: `Caught an exception during connection recovery! java.io.IOException at com.rabbitmq.client.impl.AMQChannel.wrap (AMQChannel.java:106) at com.rabbitmq.client.impl.AMQChannel.wrap (AMQChannel.java:102) at com.rabbitmq.client.impl.AMQConnection.start(AMQConnection.java:346) at com.rabbitmq.client.impl.recovery. RecoveryAwareAMQConnectionFactory. newConnection (RecoveryAwareAMQConnectionFactory.java:36) at com.rabbitmq.client.impl.recovery. AutorecoveringConnection. recoverConnection(AutorecoveringConnection.java:388) at com.rabbitmq.client.impl.recovery. AutorecoveringConnection.beginAutomaticRecovery (AutorecoveringConnection.java:360) at com.rabbitmq.client.impl.recovery.AutorecoveringConnection. access$000(AutorecoveringConnection.java:48) at com.rabbitmq.client.impl.recovery. AutorecoveringConnection$1.shutdownCompleted (AutorecoveringConnection.java:345) at com.rabbitmq.client.impl.ShutdownNotifierComponent. notifyListeners(ShutdownNotifierComponent.java:75) at com.rabbitmq.client.impl.AMQConnection$MainLoop.run (AMQConnection.java:572) at java.lang.Thread.run(Thread.java:745) Caused by: com.rabbitmq.client.ShutdownSignalException: connection error at com.rabbitmq.utility.ValueOrException.getValue`

	(ValueOrException.java:67) at com.rabbitmq.utility.BlockingValueOrException. uninterruptibleGetValueBlockingValueOrException.java:33) at com.rabbitmq.client.impl.AMQChannel$BlockingRpcContinuation. getReply (AMQChannel.java:343) at com.rabbitmq.client.impl.AMQConnection.start (AMQConnection.java:292) ... 8 more Caused by: java.net.SocketException: Connection reset at java.net.SocketInputStream.read (SocketInputStream.java:189) at java.net.SocketInputStream.read (SocketInputStream.java:121) at java.io.BufferedInputStream.fill (BufferedInputStream.java:246) at java.io.BufferedInputStream.read (BufferedInputStream.java:265) at java.io.DataInputStream.readUnsignedByte (DataInputStream.java:288) at com.rabbitmq.client.impl.Frame.readFrom(Frame.java:95) at com.rabbitmq.client.impl.SocketFrameHandler.readFrame (SocketFrameHandler.java:139) at com.rabbitmq.client.impl.AMQConnection$MainLoop.run (AMQConnection.java:532)
Possible Cause	RabbitMQ service not running on the NetWitness Server.
Solution	Restart the RabbitMQ, SMS, and NetWitness services using the following commands. systemctl restart rabbitmq-server systemctl restart rsa-sms systemctl restart jetty

Message/ Problem	User Interface: **Cannot connect to System Management Service**
Cause	The System Management Service, RabbitMQ, or Mongo service is not running.
Solution	Run the following commands on NetWitness Server to make sure all these services are running. [root@nwserver ~]# systemctl status rsa-sms RSA NetWitness SMS :: Server is not running. [root@nwserver ~]# systemctl start rsa-sms Starting RSA NetWitness SMS :: Server... [root@nwserver ~]# systemctl status rsa-sms RSA NetWitness SMS :: Server is running (5687). [root@nwserver ~]# systemctl status mongod mongod (pid 2779) is running... systemctl status rabbitmq-server Status of node nw@localhost ... [{pid,2501}, {running_applications,

	[{rabbitmq_federation_management,"RabbitMQ Federation Management", "3.3.4"},

Message/ Problem	User Interface: **Cannot connect to System Management Service**
Possible Cause	`/var/lib/rabbitmq` partition usage is 70% or greater.
Solution	Contact Customer Care.

Message/ Problem	User Interface: **Host migration failed.**
Possible Cause	One or more NetWitness services may be in a **stopped** state.
Solution	Make sure that the following services are running then restart the NetWitness Server: Archiver, Broker, Concentrator, Decoder, Event Stream Analysis, Response Server, IPDB Extractor, Log Collector, Log Decoder, Malware Analysis, Reporting Engine, Warehouse Connector, Workbench.

Message/ Problem	User Interface: **Server Unavailable.**
Possible Cause	One or more NetWitness services may be in a **stopped** state.
Solution	Make sure that the following services are running then restart the NetWitness Server: Archiver, Broker, Concentrator, Decoder, Event Stream Analysis, Response Server, IPDB Extractor, Log Collector, Log Decoder, Malware Analysis, Reporting Engine, Warehouse Connector, Workbench.

Message/ Problem	User Interface: **Server Unavailable**

Possible Cause	System Management Service (SMS), RabbitMQ, or Mongo service is not running.

Solution 1	Run the following commands on NetWitness Server to make sure all these services are running.
	`[root@nwserver ~]# systemctl status rsa-sms`
	`RSA NetWitness SMS :: Server is not running.`
	`[root@nwserver ~]# systemctl start rsa-sms`
	`Starting RSA NetWitness SMS :: Server...`
	`[root@nwserver ~]# systemctl status rsa-sms`
	`RSA NetWitness SMS :: Server is running (5687).`
	`[root@nwserver ~]# systemctl status mongod`
	`mongod (pid 2779) is running...`
	`systemctl status rabbitmq-server`
	`Status of node nw@localhost ...`
	`[{pid,2501},`
	`{running_applications,`
	`[{rabbitmq_federation_management,"RabbitMQ Federation Management",`
	`"3.3.4"},`

Solution 2	Make sure `/var/lib/rabbitmq` partition is less than 75% full
Solution 3	Check NetWitness Server log files (`var/lib/netwitness/uax/logs/nw.log`) for any errors.

Message/ Problem	ContextHub stops and does not allow you to add or edit data sources and lists.
Possible Cause	The storage is full by 95% or above.
Solution 1	Increase the storage by updating the YML file, located at /etc/netwitness/contexthub-server/ contexthub-server.yml. For example, to increase storage from 120 to 150 GB, enter a value (in bytes) by editing the relevant parameter: `rsa.contexthub.data.disk-size: 161061273600`
Solution 2	Delete unwanted or unused large list.
Solution 3	Configure the TTL index for the list to automatically delete STIX and TAXI data and to clean up storage space.

Message/ Problem	Context Hub runs on a fixed memory and 50% is reserved for cache. When cache is 100% full, the cache response stops. For all new lookups the response will be slow.

Possible Cause	The cache is full by 50% or above.
Solution 1	By default, Context Hub cleans the cache every 30 minutes. Reduce the cache expiration time of data sources.
Solution 2	Disable cache for data sources.
Solution 3	Increase the RAM of the CH Java process by editing the `-Xmx` option available in the /etc/netwitness/contexthub-server/contexthub-server.conf file. In `JAVA_OPTS`, search for the `-Xmx` option. For example, edit the entry as follows: `-Xmx8G` where `8G` represents 8GB space. Then restart the ContextHub service. **Note:** The memory is less than the available system memory. Be aware that there are many other services running on the host.

Message/ Problem	List Data Source displays an unhealthy stats or status.
Possible Cause 1	Unable to: • access the data source • parse or read a CSV file • schema mismatched CSV
Possible Cause 2	Unable to authenticate when accessing the data source.
Solution 1	Make sure to save the csv file at correct location i.e/var/lib/netwitness/contexthub-server/data/ and verify the required read permissions.
Solution 2	Make sure the csv file schema specified while configuring the data source matches. If not, then either create a new data source with the new schema or edit the csv file to match the schema. For example, if you configure a List Data Source with a schema with column1, column2, and column3. And next time you update the csv file where the number of column increase or decrease or the order of the columns are changed. In this case there is a schema mismatch and the configured list data source will show “Unhealthy” in Health and Wellness stats.
Solution 3	Make sure the password is correct. To confirm edit the data source, enter the password and click test connection. For more information related the above solutions, see "Configure Lists as a Data Source" topic in the *Context Hub Configuration Guide*.

Issues Not Identified by the User Interface or Logs

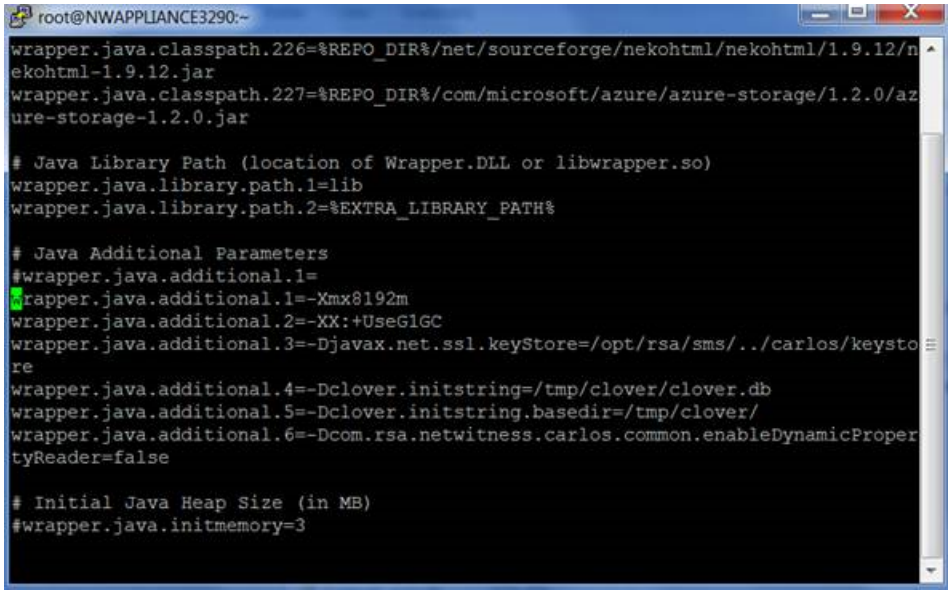
This section provides troubleshooting information for issues that are not identified by messages NetWitness displays in the Health & Wellness Interface or includes in the Health & Wellness log files. For example, you may see incorrect statistical information in the Interface.

Problem	Incorrect statistics displayed in Health and Wellness interface.
Possible Cause	SMS service is not running. SMS service must be running on the NetWitness Server.
Solution	Restart SMS service.

Problem	NetWitness does not show the version to which you upgraded until you restart jettysrv (jeTTy server).
Possible Cause	When NetWitness checks a connection, it polls a service every 30 seconds to see if it is active. During that 30 seconds, if the service comes back up, it will not get the new version.
Solution	1. Manually stop the service. 2. Wait until you see that it is it offline. 3. Restart the service. NetWitness displays the correct version.

Problem	NetWitness Server does not display the **Service Unavailable** page.
Possible Cause	After you upgrade to NetWitness version 10.5, JDK 1.8 is not default version and this causes the jettysrv (jeTTy server) to fail to start. Without the jeTTy server, the NetWitness server cannot display the **Service Unavailable** page.
Solution	Restart jettysrv.

Problem	The SMS service is stopped and the following error is displayed in the log file: `java.lang.OutOfMemoryError: Java heap space`
Solution	You can use the following solution to increase the memory according to your needs. 1. Open /opt/rsa/sms/conf/wrapper.conf



2. Replace `wrapper.java.additional.1=-Xmx16g` with:
   `wrapper.java.additional.1=-Xmx20g`
3. Restart the SMS service:
   `systemctl start rsa-sms`

Monitor New Health and Wellness

NetWitness New Health and Wellness is an advanced monitoring and alerting system that provides insights on the operational state of the host and services in your deployment, and helps identify potential issues. NetWitness Platform is prepackaged with a third-party tool that renders interactive dashboards and visualizations.

New Health and Wellness provides:

- Dashboards with interactive visualization.
- Easy-to-create customized content.
- Alerts on your data and customize alert conditions.
- Ability to add alert notifications (for example, Email and Syslog notifications).
- Ability to suppress alert notifications for a time period as required.

New Health and Wellness provides default content, such as dashboards, visualizations, and monitors, to set up monitoring and alerting.

Please direct any New Health and Wellness feedback to nw.health.wellness.feedback@rsa.com.

Dashboard

The dashboard is a collection of intuitive visualizations for the administrator to monitor the health of the host and services, identify trends, track performance, and drill down to specific details.

After installation of the New Health and Wellness service, the following default dashboards are available to begin monitoring.

- Deployment Health Overview dashboard
- Hosts dashboard
- Logs dashboard
- Packets Overview dashboard
- Analysis dashboard
- Endpoint dashboard
- ESA Correlation Overview dashboard

For more information on the dashboards, see New Health and Wellness Dashboards .

Visualization

A visualization is a graphical representation of data in your deployment. You can create new visualizations or use the existing visualizations to build dashboards. Depending on the visualization you select the data is displayed in the dashboard.

Monitors

A monitor is a job that runs on a defined schedule, which queries Elasticsearch to evaluate the system health. You can define one or more triggers for a monitor and assign a severity level based on the threshold. When one or more trigger conditions are met, New Health and Wellness generates an alerts. You can create new monitors or customize the existing monitors based on your requirement.

Notifications

Notifications can be sent when health alerts are generated, for example, email and syslog notifications.

If you do not want to receive notifications, you can suppress them for a time period. For example, if you want to suppress low-severity alerts during weekends, you can do so by specifying a suppression policy.

After the suppression time period, the notifications will be sent for the alerts triggered during the suppression time.

Installing New Health and Wellness

You can deploy the New Health and Wellness feature on any one of the following, listed in the order of preferred deployment method with most preferred first:

- Standalone virtual host (Most preferred recommendation to ensure no performance impact on any other functionality of deployed nodes)
- Physical machine:
  - Broker
  - Admin Server
  - ESA

Installing New Health and Wellness enables all hosts and services in your deployment to start sending metrics for monitoring. For more information on installing New Health and Wellness, see the "Deployment Optional Setup Procedures" topic in the *Deployment Guide*.

Accessing New Health and Wellness Dashboards

After you deploy New Health and Wellness, you can access New Health and Wellness dashboards. By default, only Administrator can access New Health and Wellness dashboards. Other roles do not have permission to view New Health and Wellness dashboards.

To access New Health and Wellness dashboards:

1. Log in to NetWitness Platform.
2. Go to  **(Admin) > Health & Wellness.**
3. Click **New Health & Wellness**.
4. Click **Pivot to Dashboard**.

Note: To view dashboards, your browser must be configured to allow popups and redirects.

Configuring Notifications

You can choose to receive notifications of alert, when an alert (monitor) is triggered. You can configure the following types of notifications for the alerts.

- Email
- Syslog

To set up notifications, you must configure the following :

- **Notification Server**: This is the source of the notifications and must be configured to specify the Email or Syslog server settings.
- **Notifications Output**: This is the notification type, namely Email or Syslog. When you set up a notification, you must specify the notification output for an alert.
- **Templates**: The message format of an alert notification is defined using template. Default template is available for New Health and Wellness notifications or you can create new templates.

Adding Alert Notifications

To add notifications such as email or syslog, you must specify the notification details as follows:

- Output Type
- Recipient
- Notification Server
- Template

Note: You must ensure that the notification server and template is configured before you set up the notifications. You can use an existing server if required. For detailed procedures, see the *System Configuration Guide*.

To add a notification:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness**.
3. Click **View Notifications Settings**.
4. Specify the following:
   - **Output Type -** Notification type such as Email or Syslog.
   - **Recipient** - Select the recipient based on the output type selected.
   - **Notification Server** - Select the server that will send the notification. If you want to add a new notification server, see "Configure Notification Servers" in the *System Configuration Guide*.
   - **Template** - Notification template such as Email or Syslog.
5. If you want to add another notification, click **Add Condition** and repeat step 4.

Note: You can specify a maximum of four conditions in the notification settings.

6. Click **Save**.

Suppressing Notifications

You can suppress notifications for a time period by specifying a suppression policy. For example, you do not want to receive notifications of low-severity alerts during peak hours or weekends. This ensures that notifications are not sent during the selected time period. The notifications triggered during this time period will be sent after the suppression period.

A suppression policy consists of condition that defines the day and time range for suppressing a notification.

Note: You can specify a maximum of seven conditions in the policy.

To suppress alert notification:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness**.
3. In the **New Health & Wellness** tab, do anyone of the following:
   - Select the alert or alerts for which you want to apply a suppression policy.
   - Select all to apply a suppression policy to all alerts.
   - Filter alerts on **suppression applied**, **monitor name**, **trigger name**, and **severity** and apply the suppression policy.
4. Click **View Suppression Policy**.
   The policy suppression dialog is displayed.
5. Specify the days and time range during which you want to suppress the notification.

Note: Time range is based on the Time Zone configured in the User Preferences panel as described in "Setting User Preferences" in the *Getting Started Guide*.

6. To add an additional suppression policy, click **Add condition** and repeat step 5.

Note: The maximum number of suppression policy supported for an alert is seven, after which the **Add Condition** option is disabled.

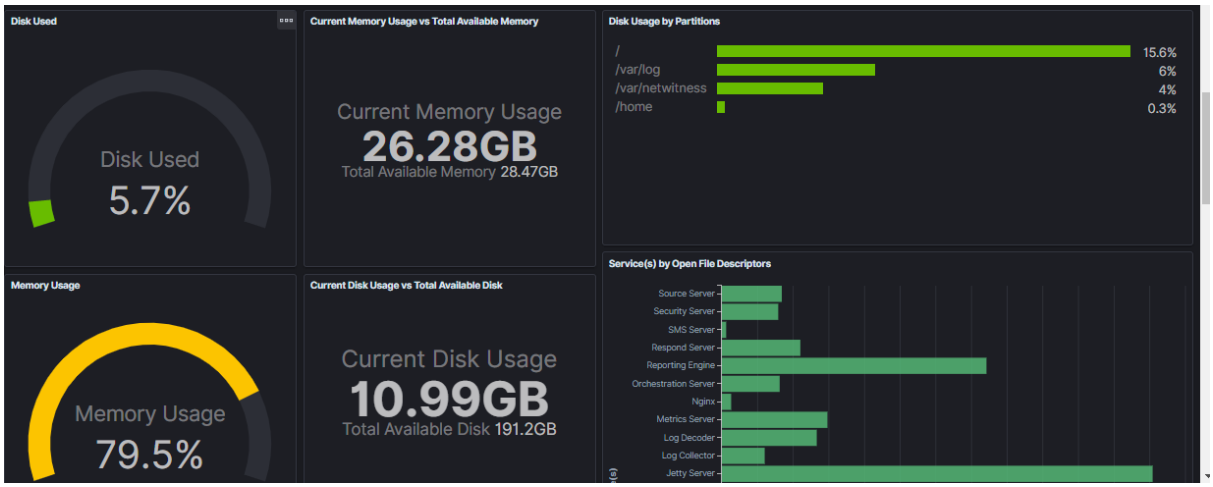
7. Click **Save**.

Monitoring through Dashboards

For a thorough analysis of the health of NetWitness Platform hosts and services using the different dashboards, you can pivot to dashboard without any additional authentication. By default, the last 6 hours of data is displayed in the dashboard. For more information on the default dashboards, see New Health and Wellness Dashboards.

To monitor through the dashboard:

1. Go to **(Admin) > Health & Wellness**.
2. Click **New Health & Wellness**.
3. Click **Pivot to Dashboard**.
   The Deployment Health Overview dashboard is displayed.
4. Go to > **Dashboard** to view all the available dashboards.
5. Select the dashboard you want to view.
6. Click the dashboard link. For example, Hosts.
   Once the dashboard view is displayed you can look at the visualizations (charts, tables, and so on) to view current disk usage of hosts, incoming and outgoing traffic of the host, active queries on a service, and so on.



7. You can adjust the time range on the top right corner and apply filters to view the statistics.

Creating a Custom Dashboard

You can create a new dashboard by adding one or more existing visualizations or a new visualization.

To create a new dashboard:

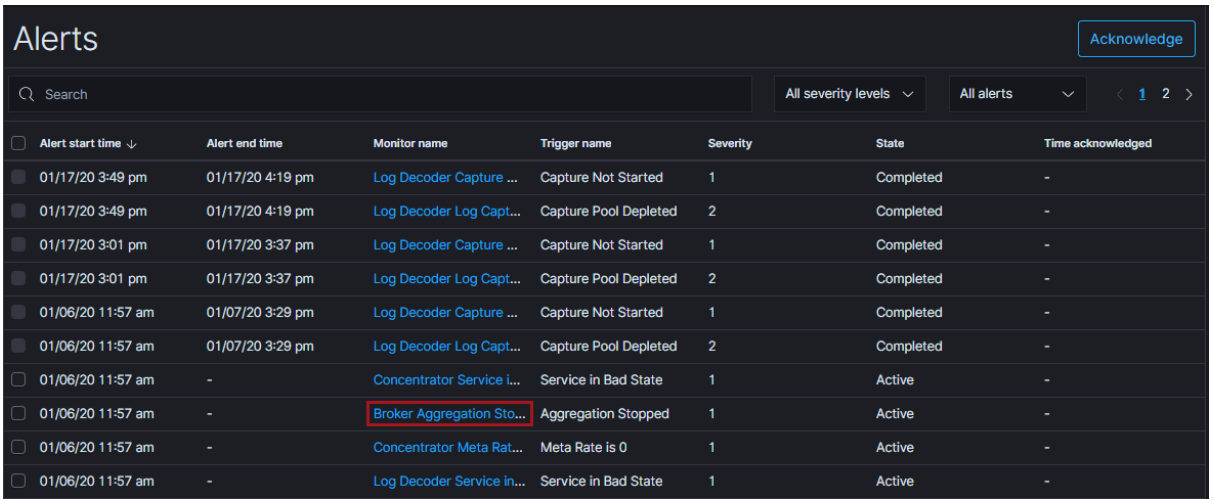
1. Log in to NetWitness Platform.
2. Go to  **(Admin) > Health & Wellness.**
3. Click **New Health & Wellness**.
4. Click **Pivot to Dashboard**.
5. Go to  > **Kibana** > **Dashboard**.
6. In the **Dashboards** panel, click **Create dashboard**.
7. Click **Create new**.
8. Select a visualization that you want to add to the dashboard. For more information on Visualization, see "Visualize" topic in the Kibana 7.10.0 guide.
9. Click **Save**.

Monitoring through Alerts

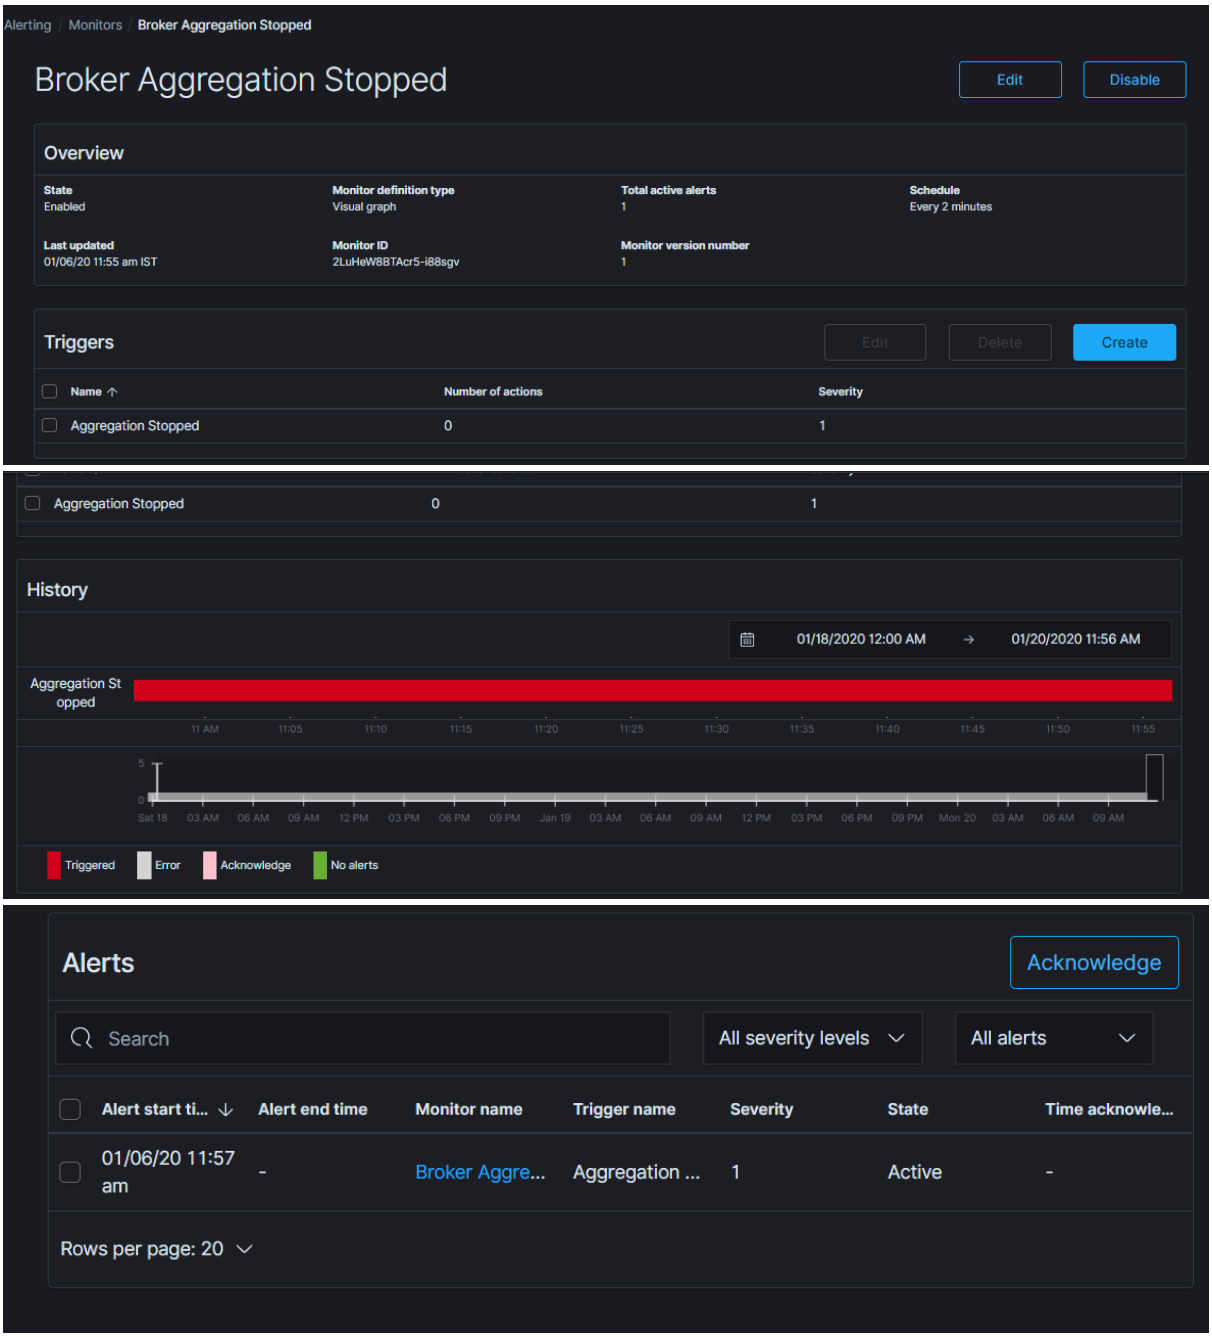
You can monitor the health of NetWitness Platform hosts and services using the alerts.

To monitor using alerts:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness**.
3. Go to  >**Open Distro for Elasticsearch** > **Alerting**.
   The alerts view summarizes the alerts generated over the period of time along with the trigger, severity (Critical, High, Medium or Low) and state of the alert.
4. To view the monitors and triggers associated with the alert, click the **Monitor name** link.



For example, if an alert is generated by the monitor name Broker Aggregation Stopped, you can view more details by clicking on the Broker Aggregation Stopped monitor link.


Alerting / Monitors / Broker Aggregation Stopped
Broker Aggregation Stopped
Edit
Disable
Overview
State
Enabled
Monitor definition type
Visual graph
Total active alerts
1
Schedule
Every 2 minutes
Last updated
01/06/20 11:55 am IST
Monitor ID
2LuHeW8BTAcr5-i88sgv
Monitor version number
1
Triggers
Edit
Delete
Create
Name
Number of actions
Severity
Aggregation Stopped
0
1
Aggregation Stopped
0
1
History
01/18/2020 12:00 AM
01/20/2020 11:56 AM
Aggregation Stopped
11 AM
11:05
11:10
11:15
11:20
11:25
11:30
11:35
11:40
11:45
11:50
11:55
5
0
Sat 18
03 AM
06 AM
09 AM
12 PM
03 PM
06 PM
09 PM
Jan 19
03 AM
06 AM
09 AM
12 PM
03 PM
06 PM
09 PM
Mon 20
03 AM
06 AM
09 AM
Triggered
Error
Acknowledge
No alerts
Alerts
Acknowledge
Search
All severity levels
All alerts
Alert start ti...
Alert end time
Monitor name
Trigger name
Severity
State
Time acknowle...
01/06/20 11:57 am
-
Broker Aggre...
Aggregation ...
1
Active
-
Rows per page: 20

Creating Custom Monitors

You can create a new monitor for the host and services and define a trigger.

To create monitors:

1. Log in to NetWitness Platform.
2. Go to  **(Admin) > Health & Wellness.**
3. Click **New Health & Wellness**.
4. Click **Pivot to Dashboard**.
5. Go to  > **Open Distro for Elasticsearch** > **Alerting**.
6. In the **Monitors** tab, click **Create monitors**.
7. In the **Create Monitors** section, specify the required details.
8. Click **Create**.

   After a monitor is created, you can add a trigger to this monitor.
9. In the **Create Trigger** view, provide the required details:
   a. **Trigger name** - Specify the name of the trigger.
   b. **Severity level** - Set the severity level from range 1–5. 1 is the highest severity and 5 is the lowest severity.
   c. **Trigger condition** - Set the trigger condition with the value. The options are IS ABOVE, IS BELOW, IS EQUAL. For example, IS ABOVE 200.
10. Click **Create** to save the trigger.

Note: After the specified duration, defined as **FOR THE LAST** in the **Define monitor** view, the alert state will change if the trigger condition is met. For more information, refer https://nw-corp.atlassian.net/browse/ASOCKB-65.

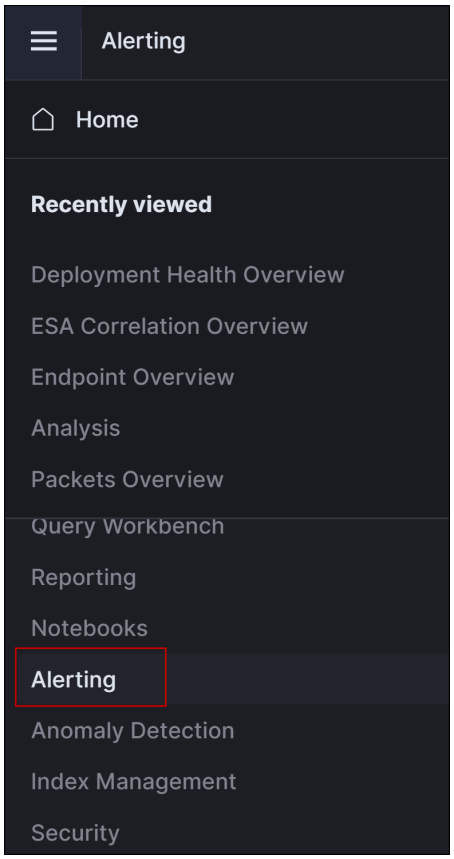
For more information on creating monitors, see "Alerting" in the Open Distro for Elasticsearch 1.12.0 guide.

Adding Custom Trigger to an Existing Monitor

You can add one or more triggers to an existing monitor and assign severity level.

To add trigger to an existing monitor:

1. Log in to NetWitness Platform.
2. Go to  **(Admin) > Health & Wellness.**
3. Click **New Health & Wellness**.
4. Click **Pivot to Dashboard**.
5. Go to  > **Open Distro for Elasticsearch** > **Alerting**.



6. In the **Monitor** section, click the **monitor** for which new trigger need be added.
7. In the **Triggers** section, select **Create**.
8. In the **Create Trigger** view, provide the required details:
   a. **Trigger name** – specify the name of the trigger.
   b. **Severity level** - Set the severity level from range 1-5. 1 is the highest severity and 5 is the lowest severity.

c. **Trigger condition** – Set the trigger condition with the value. The options are IS ABOVE, IS BELOW, IS EQUAL. For example, IS ABOVE 200.

9. Click **Create** to save the trigger.

Note: After the specified duration, defined as **FOR THE LAST** in the **Define monitor** view, the alert state will change if the trigger condition is met. For more information, refer https://nw-corp.atlassian.net/browse/ASOCKB-65.

Managing Dashboards and Alerts

You can modify the dashboards and alerts to monitor details of interest.

Modify a Dashboard

You can modify the dashboard to edit the visualization, delete visualization, customize the panel title, or change the positions of visualization. You can organize the visualization in the dashboard to display data of greatest interest on the top.

Note: Any changes to the visualization in a dashboard modifies the visualization content.

To modify a dashboard:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness**.
3. Click **Pivot to Dashboard**.
4. Go to  > **Kibana** > **Dashboard**.
5. Click the dashboard link you want to modify. For example, Hosts overview.
6. Click **Edit** and make the necessary changes to the dashboard. For example, you can edit or delete visualization, customize panel.
7. Click **Save and return**.

Delete a Dashboard

Once the dashboard is deleted, you cannot monitor the details specific to the dashboard.

To delete a dashboard:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness**.
3. Click **Pivot to Dashboard**.
4. Go to  > **Kibana** > **Dashboard**.
5. Select the check box of the dashboard you want to delete and click **Delete**.

You can delete one or more dashboards at a time.

Delete a Visualization

To delete a visualization:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness**.
3. Click **Pivot to Dashboard**.
4. Go to  > **Kibana** > **Visualize**.
5. In the **Visualizations** view, select the visualizations you want to delete.
6. Click **Delete visualization**.

   You can delete one or more visualization at a time.

Modify an Existing Trigger

To modify an existing trigger:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness** tab.
3. Click **Pivot to Dashboard**.
4. Go to  > **Open Distro for Elasticsearch** > **Alerting**.
5. In the **Dashboard** tab, select the monitor whose trigger is to be modified.
6. In the **Triggers** section, select the trigger you want to modify from list of Triggers and select **Edit**.
7. In the **Edit Trigger** view, make the necessary changes. You can change the Trigger name, severity level, Trigger condition.
8. Click **Update** to save the changes.

Managing Notifications

You can manage alert (monitor) notifications based on your requirement.

Modify a Notification

To modify a notification:

1. Log in to NetWitness Platform.
2. Go to  **(Admin) > Health & Wellness.**
3. Click **New Health & Wellness**.
4. Click **View Notifications Settings**.
5. Make the necessary changes.
6. Click **Save**.

Modify a Notification Suppression Policy

When you modify a suppression policy, you can change the day or time ranges. If you select multiple alerts to modify the policy, any changes would overwrite all the existing policies.

To modify suppression policy:

1. Go to  **(Admin) > Health & Wellness.**
2. Click **New Health & Wellness**.
3. Select one or multiple rows for which you to modify suppression policy.
4. Click **View Suppression Policy**.

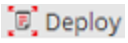
Note: If you have selected multiple alerts to modify the suppression policy,
- Only policies that are common are displayed. Policy suppression settings would be empty if there is no common polices.
- If you edit the common policies, the changes would overwrite all the selected alert policies.

5. Make the necessary changes.
6. Click **Save**.

Advanced Configurations

Restore Default Content

This allows you to bring back all the default content such as dashboards, visualizations, monitors to its original or default state. This overwrites any changes made to the default content. For example, if you have deleted any dashboard or visualization and want to bring back the default content.

1. Log in to NetWitness Platform UI.
2. Click  (Configure) > **LIVE CONTENT**.
3. In the **Search Criteria** panel, select the **Resource Types** as:
   - Health and Wellness Dashboards
   - Health and Wellness Monitors
4. Click **Search**.
5. In the **Matching Resources** view, select the checkbox to the left of the resources that you want to deploy.
6. In the **Matching Resources** toolbar, click 

.
7. In the **Deployment Wizard** > **Resources** tab, click **Next**.
8. In the **Services** tab, select the Metrics Server service.
9. Click **Next**.
10. Click **Deploy**.
    The **Deploy** page is displayed. The Progress bar turns green when you have successfully deployed the resources to the selected services.
11. Click **Close**.

Enable Services

This is used to enable all the services to start sending metrics for monitoring. For example, if you have disabled few services from sending metrics for monitoring and want to enable all those disabled services to start sending again.

1. SSH to the Admin Server.
2. Enter the following command:
   `nw-shell`
   The console window is displayed.



3. Connect to metrics-server using the following command:

   ```
   connect --service metrics-server
   ```

4. Enter the login command:
   ```
   login
   ```
5. Enter the admin username and password.
6. Navigate to the enable option using the following command:

   ```
   cd /rsa/metrics/elastic/enable-all
   ```

7. Execute the following command to enable all services:

   ```
   invoke
   ```

Disable Services

This is used to disable all the services to send metrics for monitoring. Once disabled, none of the services sends alerts to the Elasticsearch and the dashboards are not updated, and alerts will not be triggered.

1. SSH to the Admin Server.
2. Enter the following command:
   ```
   nw-shell
   ```
   The console window is displayed.



3. Connect to metrics-server using the following command:

   ```
   connect --service metrics-server
   ```

4. Enter the login command:
   ```
   login
   ```
5. Enter the admin username and password.
6. Navigate to Elasticsearch using the following command:

   ```
   cd /rsa/metrics/elastic/disable-all
   ```

7. Execute the following command to disable all services to stop writing to Elasticsearch:

   ```
   invoke
   ```

Note: This disables all services to send metrics to Elasticsearch but does not stop metric beat to send system level metrics to Elasticsearch. You need to manually stop metric beat on all hosts if you wish to stop using Health and Wellness.

Update an Interval

You can update a common interval for all the services to send data for monitoring. For example, if all the services are set to different intervals and you want to configure all the services to send data to elastic search on the same interval.

The interval can be set in seconds, minutes and hours.

1. SSH to the Admin Server.
2. Enter the following command:
   ```
   nw-shell
   ```
   The console window is displayed.



3. Connect to metrics-server using the following command:

   `connect --service metrics-server`

4. Enter the login command:
   `login`
5. Enter the admin username and password.
6. Navigate to the `update-interval` directory using the following command:

   `cd /rsa/metrics/elastic/update-interval`

7. Execute the following command to set a common interval for all the services:

   `invoke <interval>`

   For example, `invoke 30seconds`

Update the Default Configuration

By default, New Health and Wellness configurations are applied after the New Health and Wellness is enabled successfully. To change the configuration of a service, you need to update the existing configuration. After the configuration is updated, the service is notified of the changes.

To update the configuration, perform the following:

1. SSH to the Admin Server.
2. Enter the following command:
   `nw-shell`
   The console window is displayed.



3. Connect to the metrics-server using the following command:

   `connect --service metrics-server`

4. Enter the login command:

   `login`

5. Enter the admin username and password.
6. To get configuration of a service, execute following commands:
   a. `cd /rsa/metrics/elastic/get-config`
   b. `invoke <service-id>`

Note:
To get the service id for core services:
1) Go to **(Admin)** > Core service.
2) Click > **View** > **Explore**.
3) Expand the **sys/stats** node list.
4) In the **UUID** filed, copy the value.
To get the service id for launch services:
1) Go to **(Admin)** > Launch service.
2) Click > **View** > **Explore**.
3) Click the process node.
4) In the **service-id** field, copy the value.
To get the service id for Carlos services:
1) SSH to host in which the Carlos service is deployed.
2) Execute the following command:
For Reporting Engine:
`cat /var/netwitness/re-server/rsa/soc/reporting-engine/service-id`
For Legacy Web Server:
`cat /var/netwitness/uax/service-id`

Note: The core services are Archiver, Broker, Concentrator, Decoder, Log Decoder and; Carlos services are Reporting Engine, Legacy Web Server. All the other services that are not included in Core and Carlos services are part of launch services.

7. Copy the configuration and save it in a file.
   a. Copy the configuration from step 6 and exit from `nw-shell` using command:
      `exit`
   b. Create a file under `/root` in admin server, copy the configurations to the file and save it.

   For example, For the Reporting Engine service, create a file `reporting-engine.json` under `/root/` and copy the configurations obtained from step 6 and save.
8. To set configurations for a service:
   a. `cd /rsa/metrics/elastic/set-config`
   b. `invoke --file <absolute path of the path>`
      For example, `invoke --file /root/reporting-engine.json`

Configure the Data Retention Policy

You can configure the retention policy for monitors (alerts triggered) and metrics based on age and size.

By default, 90 days of data with 100 GB of size for monitors (alerts triggered) and 30 days of data with 100 GB of size for metrics are retained.

To change the configure for monitor (alerts triggered) retention:

1. SSH to the Admin Server.
2. Enter the following command:
   `nw-shell`
   The console window is displayed.



3. Connect to metrics-server using the following command:
   `connect --service metrics-server`

4. Enter the login command:
   `login`
5. Enter the admin username and password.
6. Go to alert-retention-threshold using command:
   `cd /rsa/metrics/elastic/data/retention/alert-retention-threshold`
7. Set the value between 1day to 90days using command:
   `set <number of days>`
   For example, `set 50days`
8. Exit from nw-shell using the command:
   `exit`
9. SSH to the host on which New Health and Wellness is installed.
10. Restart the metrics server on which New Health and Wellness is installed using the command:
    `service rsa-nw-metrics-server restart`

To change the configuration for metrics time threshold:

1. SSH to the Admin Server.
2. Enter the following command:
   `nw-shell`
   The console window is displayed.



3. Connect to metrics-server using the following command:
   `connect --service metrics-server`
4. Enter the login command:
   `login`
5. Enter the admin username and password.

6. Go to time-threshold using command:
   `cd /rsa/metrics/elastic/data/retention/time-threshold`
7. Set the value from **1day** to **90days** using command:
   `set <number of days>`
   For example, `set 40days`
8. Exit from `nw-shell` using the command:
   `exit`
9. SSH to the host on which New Health and Wellness is installed.
10. Restart the metrics server on which the New Health and Wellness is installed using the command:
    `service rsa-nw-metrics-server restart`

To change the size configuration:

1. SSH to the Admin Server.
2. Enter the following command:
   `nw-shell`
   The console window is displayed.



3. Connect to metrics-server using the following command:
   `connect --service metrics-server`
4. Enter the login command:
   `login`
5. Enter the admin username and password.
6. Go to allocated-size using command:
   `cd /rsa/metrics/elastic/data/retention/allocated-size`
7. Set the value using command:
   `set <size to be allocated>`

For example, `set 200GB`

8. Exit from nw-shell using the command:

   `exit`

9. SSH to the host on which New Health and Wellness is installed.
10. Restart the metrics server on which the New Health and Wellness is installed using the command:

    `service rsa-nw-metrics-server restart`

Note: Make sure the `/var/netwitness` partition on standalone New Health and Wellness has enough disk space. After you review your datastore configuration, you may determine that you need to add a new volume. For more information on adding a new volume, see "Add New Volume and Extend Existing File Systems" topic in the *Virtual Host Installation Guide*.

Backup and Restore New Health and Wellness

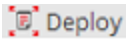
Perform the following steps in the order given:

1. Back up the Admin Server. For more information on the instructions, see "Disaster Recovery (Back Up and Restore)" topic in *Recovery Tool User Guide*.
2. At the root level, type the following command on the host on which New Health and Wellness is installed :

   ```
   nw-recovery-tool --export --dump-dir /var/netwitness/backup --category Search
   ```

   Note: Make sure that the Admin Server is restored, up and running successfully.
3. Restore the New Health and Wellness using the following command:

   ```
   nw-recovery-tool --import --dump-dir /var/netwitness/backup --category Search
   ```

   Note: Restore the search category (New Health and Wellness) on the same host on which New Health and Wellness is installed.
4. Reboot the Host on which search category (New Health and Wellness) is restored.
5. Restore the default content using RSA Live:
   a. Log in to NetWitness Platform UI.
   b. Click  (Configure) > **LIVE CONTENT**.
   c. In the **Search Criteria** panel, select the **Resource Types** as:
      - Health and Wellness Dashboards
      - Health and Wellness Monitors
   d. Click **Search**.
   e. In the **Matching Resources** view, select the checkbox to the left of the resources that you want to deploy.
   f. In the **Matching Resources** toolbar, click 

   .
   g. In the **Deployment Wizard** > **Resources** tab, click **Next**.
   h. In the **Services** tab, select the Metrics Server service.
   i. Click **Next**.
   j. Click **Deploy**.
   The **Deploy** page is displayed. The Progress bar turns green when you have successfully deployed the resources to the selected services.
   k. Click **Close**.

Troubleshooting New Health and Wellness

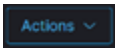
This topic describes how to troubleshoot New Health and Wellness issues.

Issue	An error ‘n of m shards failed’ or ‘unknown field in the index’ in the New Health and Wellness dashboards.
Resolution	Refresh the index patterns, perform the following: 1. Log in to NetWitness Platform. 2. Go to  **(Admin) > Health & Wellness.** 3. Click **New Health & Wellness**. 4. Click **Pivot to Dashboard**. 5. Go to  > **Stack Management** > **Index Patterns**. 6. Click **nw* index pattern**. 7. Click **Refresh** to refresh the index pattern on top right corner. **Note:** If the issue still persists, refresh other index patterns such as **nw-metricbeat*** or **nw-concentrator*** and so on.

Issue	Unable to send data to elastic search once disk usage reaches 85%.
Explanation	If the Elasticsearch disk usage reaches 85%, the saved objects (index patterns, dashboards, visualizations etc) becomes read-only mode. And, services does not write new metrics to Elasticsearch or allow to edit any saved objects.
Resolution	To change the indexes to write mode, execute the following command on the host in which Elasticsearch is installed: `curl -k --cert /etc/pki/nw/elastic/elasticsearch-cert.pem --key /etc/pki/nw/elastic/elasticsearch-key.pem -X PUT -H "Content-Type: application/json" -d '{"index.blocks.read_only_allow_delete": null }' https://localhost:9200/_all/_settings` **Note:** This command is supported only with certificates.

Issue	**New Health and Wellness** > **Pivot to Dashboard** fails due to time synchronization issue.
Explanation	Pivot to dashboard fails if there is no time synchronization between the NetWitness Server and the host on which New Health and Wellness is

Resolution	installed. You must synchronize the time and pivot to dashboard. To synchronize the time do one of the following: • Configure the NTP Server. For more information, see "Configure NTP Servers" in the *System Configuration Guide*. • Run the following commands on the host on which New Health and Wellness is installed 1. SSH to NetWitness host. 2. Run the following commands. ◦ `systemctl stop ntpd` ◦ `ntpdate nw-node-zero` ◦ `systemctl start ntpd`

Issue	NW Host High Swap Utilization monitor generates many false alerts.
Explanation	By default the set threshold is > 50 %, which might generate many false alerts on NW Host High Swap Utilization monitor for Linux hosts, which is considered normal. **Note:** On upgrade to 11.5.3.0 version or later, the monitor will be automatically deleted on the New Health & Wellness monitors view.
Resolution	On version 11.5.2.0 or lower, you can perform the following steps to delete the monitor: 1. Log in to the NetWitness Platform. 2. Go to **(missing or bad snippet)> Health & Wellness**. 3. Click **New Health & Wellness**. 4. Click **Pivot to Dashboard**. The **Deployment Health Overview** dashboard is displayed. 5. Go to  > **Open Distro for Elasticsearch** > **Alerting**. The Dashboards tab is displayed by default. 6. Click **Monitors** tab. 7. Select the NW Host High Swap Utilization monitor and click  > **Delete**. The monitor is deleted.

Appendices

This section is a collection of appendices for new health and wellness:

- New Health and Wellness Dashboards
- New Health and Wellness Monitors

New Health and Wellness Dashboards

This topic provides the list of default New Health and Wellness dashboards and associated visualizations and metrics.

Deployment Health Overview Dashboard

This dashboard provides the overall health of the NetWitness Platform hosts and services. The following table provides the information on default visualizations available on this dashboard.

Note: The parameters and metrics listed below are the default values. You can customize the parameters and metrics of any visualization based on your requirement. For example, you can customize a visualization to view the CPU utilization for all the core services or any particular service.

Visualization	Parameters and Metrics	Objective	Description
Alarms Summary	• Count of active alert • Alert severity	Provides the summary of active health alarms based on the severity.	Displays the active alarms grouped by severity (Critical, High, Medium, Low).
Offline Services	• Service name • Status Time • Refresh time 15 minutes	Identifies the list of unavailable services.	Displays the list of offline services.
Stopped Archiver Aggregation	• Count of archivers where aggregation is stopped • Refresh time 15 minutes	Identifies the number of Archivers where aggregation is stopped.	Displays the number of Archivers where aggregation is stopped. For more information, see Notifications.
Stopped Broker Aggregation	• Count of Brokers where aggregation is stopped • Refresh time 15 minutes	Identifies the number of Brokers where aggregation is stopped.	Displays the number of Brokers where aggregation is stopped. For more information, see Notifications.
Stopped Concentrator Aggregation	• Count of Concentrators where aggregation is stopped • Refresh time 15 minutes	Identifies the number of Concentrators where aggregation is stopped.	Displays the number of Concentrators where aggregation is stopped. For more information, see Notifications.

Visualization	Parameters and Metrics	Objective	Description
Stopped Decoder/Log Decoder Capture	• Count of Decoders or Log Decoders where capture is stopped • Refresh time 15 minutes	Identifies the number of Decoders or Log Decoders where capture is stopped.	Displays the number of Decoder or Log Decoder where capture is stopped. For more information, see Notifications.
Total vs Offline Services	• Total number of services • Count of offline services • Refresh time 15 minutes	Identifies the number of offline services versus total number of services.	Displays the total number of services and the number of services that are offline.
Stopped State Aggregation & Capture	• Services name • Host name • Service version	Provides the list of services where aggregation and capture are stopped.	Displays the list of services where aggregation and capture are stopped.
NetWitness Services Version Status	• Service version	Provides the status of NetWitness Platform service versions.	Displays the status of NetWitness Platform service versions.
NetWitness Services – Uptime Summary	• Service name • Host name • Running since	Provides an overview on the uptime of the services in the deployment.	Displays the list of services and their uptime.
Memory Utilization Trend	• Service name • Memory usage	Provides the memory utilization trend to detect any high utilizations and take necessary action.	Displays the memory utilization trend of the hosts.
Current CPU Usage	• Services name • CPU usage	Provides the CPU usage trend of the hosts to identify any high utilizations and take necessary action.	Displays the current CPU usage of the services.
Current Disk Usage	• Services name • Disk usage	Provides the disk utilization in the real time to identify any high utilizations and take necessary action.	Displays the current disk usage of the hosts.
Capture Rate for Log Decoders	• Service name • Capture rate	Provides the capture rate trend to identify any high values and take necessary action.	Displays the trend of Log Decoders capture rate.

Visualization	Parameters and Metrics	Objective	Description
Capture Rate for Network Decoders	• Service name • Capture rate	Provides the capture rate trend to identify any high values and take necessary action.	Displays the trend of Network Decoders capture rate.
Session Aggregation Rate and Trend for Concentrators	• Service name • Session aggregation rate	Provides an overview on the session rate of the Concentrators to identify any high values and take necessary action.	Displays the session aggregation rate and trend of Concentrator.
Retention Summary	• Service id • Service name • Running on host • Oldest meta file time • Oldest packet file time • Oldest session file time	Provides a quick view on the current retention of the Decoders, Concentrators and Archivers to check if the retention is lower than the configured retention.	Displays the oldest date for meta, session, packet present in decoders, logdecoders and concentrators
Total CPU Usage Trend for Services	• CPU usage • Service name	Provides the CPU usage trend of the services to detect the high utilization and take necessary action.	Displays the top 20 services where CPU usage is high.
Total Memory Usage Summary for Services	• Service name • Memory usage	Provides the memory usage summary of NetWitness Platform services to detect any high usage and take necessary actions.	Displays the top services that are utilizing the resident memories.

Hosts Dashboard

This dashboard provides the resource utilization and health of NetWitness hosts in your deployment. The following table provides information on default **Visualizations** available on this dashboard.

Visualization	Metrics	Objective	Description
Disk Used	• Disk usage	Provides the current disk usage of the hosts to detect the high utilization and take immediate action.	Displays the current disk usage of the host.

Visualization	Metrics	Objective	Description
Current Memory Usage vs Total Available	• Current memory usage • Total available memory	Provides the current memory usage versus total available memory to identify high usage and take necessary action.	Displays the current memory usage and total available memory of the host.
Current Disk Usage vs Total Available Disk	• Current disk usage • Total available disk	Provides the current disk usage versus total available disk to identify high usage and take necessary action.	Displays the current disk usage versus total available disk.
Disk Usage by Partitions	• Disk partition • Disk usage	Provides the disk usage by different partitions to identify high usage and take necessary action.	List of partitions and associated disk percentage.
Resident Memory Usage by Services	• Service name • Resident memory usage	Provides the resident memory usage per service to identify high usage and take necessary action.	Displays the resident memory usage of the service.
Memory Usage	Memory usage	Provides the current memory usage percentage of the hosts to identify high memory usage and take necessary action.	Displays the memory usage of the host.
CPU Usage	CPU usage	Provides the CPU usage percentage to identify high usage and take necessary action.	Displays the CPU usage of the host.
CPU Usage by Services	• Service name • CPU usage	Provides the CPU Percentage per service to detect high usage and take necessary action.	Displays the CPU usage of the service.
Interfaces by Incoming Traffic	Incoming traffic on interfaces	Provides the trend on interfaces incoming traffic to detect any deviation on time.	Display the incoming traffic interfaces.
Interfaces by Outgoing Traffic	Outgoing traffic on interfaces	Provides the trend on interfaces outgoing traffic to detect any deviation on time.	Display the interfaces outgoing traffic.
Services by Open File Descriptors	• Services • Open file descriptor	Provides the list of open file descriptor associate with a service.	Displays the list of open file descriptor associated with a service.

Visualization	Metrics	Objective	Description
TOP APPLIANCES BY DISK IO READ (Line) Vs WRITE (Bar)	• Service name • Disk IO Read • Disk IO Write	Provides the list of top appliances by disk IO read and write to detect any high usage and take necessary action.	Displays top appliances based on disk IO read and write usage.
Total Inbound Traffic for All Interfaces	• Count of inbound traffic on Interfaces • Total transferred traffic	Provides the total inbound traffic to detect any deviation on time.	Displays the current inbound traffic and total transferred traffic.
Total Outbound Traffic for All Interfaces	• Count of outbound traffic on Interfaces • Total transferred traffic	Provides the total outbound traffic to detect any deviation on time.	Display the current outbound traffic and total transferred traffic.

Logs Dashboard

This dashboard provides information on various NetWitness Platform logs. The following table provides information on default **Visualizations** available on this dashboard.

Visualization	Metrics	Objective	Description
Log Decoders by Capture Rate	• Service name • Capture Rate	Provides the capture rate of Log Decoders to detect high capture rate on time and take necessary action.	Displays the Log Decoders by capture rate.

Visualization	Metrics	Objective	Description
Log Decoders by Capture Packet Rate	• Service name • Capture Packet Rate	Provides the capture packet rate of Log Decoder to detect high capture packet rate on time and take necessary action.	Displays the Log Decoders by capture packet rate.
Log Decoders by CPU Percentage	• Service name • CPU usage	Identifies the Log Decoders by CPU usage to detect high usage and take necessary action.	Display the Log Decoders by CPU usage..
Log Decoders by Resident Memory Usage	• Service name • Resident Memory Usage	Identifies the Log Decoders by resident memory usage to detect high usage and take necessary action.	Display Log decoder by resident memory usage.
SDK Active Queries on Concentrators	• Service name • Count of active queries	Identifies the concentrators by SDK active queries.	Display concentrators by SDK active queries.
Concentrators Status	• Service running on host • Service type • Service version • Aggregation status • Average session rate • Max session rate • Active queries	Provides the concentrator status.	Display the list of concentrators and its status.

Visualization	Metrics	Objective	Description
Concentrator Session Aggregation Rate [Trend]	• Service name • Session rate	Provides the trend of Concentrator session aggregation rates to detect high session rates and take necessary action.	Displays Concentrator session aggregation rate.
SDK Active Queries on Brokers	• Service name • Count of Active Queries	Identifies the Brokers by SDK active queries.	Lists Brokers by SDK active queries.
Brokers Status	• Service running on host • Service type • Service version • Aggregation status • Average session rate • Max session rate • Active queries	Provides the Broker status.	Displays the list of Brokers and their status.

Packet Overview Dashboard

This dashboard provides information on NetWitness Platform network data. The following table provides information on default **Visualizations** available on this dashboard.

Visualization	Metrics	Objective	Description
Network Decoders by Capture Rate	• Service name • Capture rate	Identifies the capture rate of Network Decoder to detect high value and take necessary action.	Displays Network Decoders by capture rate.

Visualization	Metrics	Objective	Description
Network Decoders by Capture Drop	• Service name • Capture drop percentage	Identifies the capture drop rate of Network Decoders to detect drop rate and take necessary action.	Displays Network Decoders by capture drop.
Network Decoders by CPU Percentage	• Service name • CPU usage	Identifies the Network Decoders by CPU usage to detect high usage and take necessary action.	Displays Network Decoder by CPU used.
Network Decoders by Resident Memory Usage	• Service name • Resident memory usage	Identifies the Network Decoders by resident memory usage to detect high usage and take necessary action.	Displays Network Decoder by resident memory usage.
SDK Active Queries on Concentrators	• Service name • Count of active queries	Identifies the concentrators by SDK active queries.	Displays Concentrators by SDK active queries.
Concentrators Status	• Service running on host • Service type • Service version • Aggregation status • Average session rate • Max session rate • Active queries	Provides the Concentrator status.	Displays the list of Concentrators and their status.

Visualization	Metrics	Objective	Description
Concentrator Session Aggregation Rate [Trend]	• Service name • Session rate	Provides the trend of Concentrator session aggregation rate to detect high value and take necessary action.	Displays the trend of concentrator session aggregation rate.
SDK Active Queries on Brokers	• Service name • Count of active queries	Identifies the Brokers by SDK active queries.	Display the Broker by SDK active queries.
Brokers Status	• Service running on host • Service type • Service version • Aggregation status • Average session rate • Max session rate • Active queries	Provides the Broker status.	Displays the list of brokers and its status.

Analysis Dashboard

This dashboard provides details about Reporting Engines on Primary UI or Analyst UI. The following table provides the information on default **Visualizations** available on this dashboard.

Visualization	Metrics	Objective	Description
Reporting Engine Rule Query Executions	• Hostname • Failed rule executions • Cancelled rule execution • Active rule execution • Total rule execution	Provides the status of the queries executed by Reporting Engine to detect any deviations on time.	Displays the queries executed by Reporting Engine.

Visualization	Metrics	Objective	Description
Reporting Engine Reports Executions	• Hostname • Failed in last hour • Running more than one hour • Cancelled in last hour • Output actions failed in last hour	Provides the status of the reports executed by Reporting Engine to detect any deviations on time.	Displays the Reporting Engine reports.
Reporting Engine Alerts Execution	• Enabled alerts • Execution failed • Execution skipped in las 10 minutes • Running alerts • Output actions failed in last 10 minutes	Provides the status of the alerts generated by Reporting Engine to detect any deviations on time.	Displays the Reporting Engine alerts.
Reporting Engine Charts Executions	• Hostname • Enabled charts • Execution failed • Execution cancelled in last 10 minutes	Provides the status of the charts executed by Reporting Engine to detect deviations on time.	Displays Reporting Engine charts.
Reporting Engine Disk Usage	• Disk Used • Total disk space	Provides the disk usage by Reporting Engine to detect any deviations high usage and take necessary action.	Displays the disk used by Reporting Engine.
Unassigned Open Incidents	Count of unassigned open incidents	Identifies unassigned incidents to assist Administrator to take necessary action.	Displays the unassigned incidents.

Visualization	Metrics	Objective	Description
Incidents Sent to Archer	Count of incidents sent to archer	Provides statistics on the incidents sent to Archer to assist Administrator to take necessary action.	Displays the incidents sent Archers.

Endpoint Dashboard

This dashboard provides information on NetWitness Endpoints and agents installed on Endpoints. The following table provides information on default Visualizations available on this dashboard.

Visualization	Metrics	Objective	Description
Endpoint Server to Agent Communication Queued	• Service name • Count of queued request to Agent	Provides an overview of the queued agent communication to the Endpoint Server to identify any issues around the queued communication.	Displays the queued request to agent.
Endpoint Server to Agent Communication Rejected Count	• Service name • Count of rejected request to agent	Provides an overview of rejected agent communication to the Endpoint Server to identify any issues related to the rejected count.	Displays the rejected request to agent.
Endpoint Agent Overview	• Hostname • Total active agents • Active advanced • Active insights agents • Active advanced windows agents • Active advanced linux agents • Active advanced mac agents	Provides an overview of Endpoint Agents.	Displays list of agents and its details.

Visualization	Metrics	Objective	Description
Relay Servers Overview	• Hosts • Total relay servers • Agents communicated via relay server • Agents communicated in last two days via relay server	Provides an overview of the Relay Servers.	Displays the Relay server details.
Files Count by File Status	• Count of blacklisted files • Count of graylisted files • Count of netural file • Count of whitelisted files	Provides an overview of file status by count to assist an Administrator on the overall statistics of Endpoint actions on files.	Displays the file count of file statuses.
Files Count by Certificate Status	• Count of blacklisted certificates • Count of gray listed certificates • Count of neutral certificates • Count of whitelisted certificates	Provides an overview on certificate status to assist an Administrator to take necessary action.	Displays the count of certificate statuses.

Visualization	Metrics	Objective	Description
File Count by Reputation Status	• Count of unknown status • Count of suspicious status • Count of malicious status • Count of known good status • Count of known status • Count of invalid status	Provides an overview on the reputation status to assist an Administrator to take necessary action.	Displays the count of files reputation status.
Endpoint Hosts with Risk Score Greater than 90	Count of hosts with risk score greater than 90	Identifies the number of hosts with risk score higher than 90 for immediate attention.	Displays the count of hosts with risk score greater than 90.
Endpoint Files with Risk Score Greater than 90	Count of files with risk score greater than 90	Identifies the number of files with risk score higher than 90 for immediate attention.	Displays the count of files with risk score greater than 90.

ESA Correlation Overview Dashboard

This dashboard provides health statistics and trends on the ESA deployment. The following table provides the information on default **Visualizations** available on this dashboard.

You can choose the ESA host and Deployment name for the Dashboard view source using the filter.

Visualization	Metrics	Objective	Description
Sessions Behind by Sources	Count of sessions behind by sources.	Provides the session behind trend for the sources to take necessary actions when the session behind goes higher.	Displays the count of sessions behind by sources.
Sessions Rate by Sources	Count of sessions rate by sources.	Provides the session rate trend for the sources to take necessary actions when the session rate goes higher.	Displays the count of sessions rate by sources.
Top Rules by Memory	Memory used by rules.	Provides the memory usage per rule to identify the rule with high memory usage and take necessary action.	Displays the top rules based on memory usage.
Top Rules by CPU	CPU used by rules.	Provides the CPU usage per rule to identify the rule with high CPU usage and take necessary action.	Displays the top rules based on CPU usage.

Visualization	Metrics	Objective	Description
ESA Correlation Resident Memory Usage	Resident memory usage.	Provides resident memory usage trend to be able to detect high usage and take necessary action.	Displays the trend of ESA correlation resident memory usage.
ESA Correlation CPU Usage	CPU usage.	Provides CPU usage trend to detect high usage and take necessary action.	Displays the trend of ESA correlation CPU usage.
ESA CR - Event Rate by Deployments	Event rate of each ESA correlation deployment.	Identify the event rate by each deployment under ESA Correlation to detect high usage and take necessary action.	Displays the trend of ESA correlation event rate of each deployment.

Logstash Input Plugin Dashboard

The Logstash Input Plugin dashboard provides insight on Logstash event source and the NetWitness Input Plugin.

Prerequisites

- You must install the New Health and Wellness. For more information, see New Health and Wellness
- You must ensure to download the Logstash Input Plugin Dashboard from RSA Live. For more information, see Advanced Configurations

The following table provides the information on default **Visualizations** available on this dashboard.

Visualization	Metrics	Objective	Description
Logstash Inactive Pipelines	• Inactive pipeline • Total number of Pipeline	Provides the summary of inactive pipeline.	Displays the total number of inactive pipeline.
Logstash Inactive Sources	• Inactive sources • Total number of sources	Provides the summary of inactive sources.	Displays the total number of inactive sources.
Logstash Incoming Rate	• Incoming rate of Logstash event sources.	Provides the trend of incoming rate for Logstash event source.	Displays the trend of Logstash incoming rate.
Logstash Outgoing Rate	Outgoing rate of Logstash event sources.	Provides the trend of outgoing rate for Logstash event sources.	Displays the trend of Logstash outgoing rate.

Visualization	Metrics	Objective	Description
Logstash Data Persistence	Count of persisted Logstash event source.	Provides the trend of persisted Logstash event source to take necessary actions when the persisted data count goes higher.	Displays the trend of Logstash persisted data.
Logstash Pipelines status	• Inactive pipelines • Active pieplines • Total number of pipelines	Provides the status of Logstash pipeline.	Displays the status of Logstash pipeline.
Sources by Input event rate	• Input rate of sources	Provides the trend of source input rate to identify any high values and take necessary action.	Displays the trend of source input event rate.
Sources by Sessions behind	Logstash session behind	Provides the trend of Logstash session behind to identify any high values and take necessary action..	Displays the trend of Logstash session behind.
Events to consumer (Output Plugin)	Output event rate	Provides the trend of output event rate to identify any high values and take necessary action.	Displays the trend of Logstash output event rate.

Visualization	Metrics	Objective	Description
Logstash Service Uptime	• Service start time • Host name • Service name	Provides the time when the service is active.	Displays the time when the service is running.
Source Status	• Stream status • Source host • Service type • Service version • Logstash host • Average input event rate • Average source session rate • Average session behind	Provides overall status of the services. For example, stream status of the Decoder connected with the Logstash.	Displays the list of all services and its status.
Logstash Services by CPU Usage	• CPU usage • Service name	Provides the CPU usage trend of the Logstash services to identify any high utilizations and take necessary action.	Displays the CPU utilization trend of the Logstash services.
Logstash Service by Resident Memory Usage	• Resident memory usage • Service name	Provides the resident memory usage trend of the Logstash services to identify any high utilizations and take necessary action.	Displays the resident memory utilization trend of the Logstash services.

License Usage Dashboard

The License Usage Dashboard provides License usage statistics and trends on the Log Decoder, Packet Decoder, Endpoint, UEBA, and Malware services and at aggregated levels under throughput license. It provides an overview of the usage of all types of Throughput licenses in your deployment.

It helps you:

- Track daily license usage for individual hosts
- Track daily usage of Throughput licenses for all the hosts in your deployment
- Download license usage reports

Prerequisites

- You must install the New Health and Wellness. For more information, see New Health and Wellness
- You must ensure to download the License Usage Dashboard from RSA Live. For more information, see Advanced Configurations

IMPORTANT: The date shown in the user interface is set by the browser's time zone. The user's browser may be different from the time zone of the Admin server host. To make a date shown in the user interface to match the Admin server, change the time zone setting under Stack Management > Advanced Settings. If you update the time zone under Advanced settings, it affects other DateTime displays throughout the user interface.

The following table provides the information on default **Visualizations** available on this dashboard. You can choose any host and Deployment name for the Dashboard view source using the filter.

Visualization	Metrics	Objective	Description
Packets Analyzed	• Usage in bytes • Per day	Provides the packet data analyzed in bytes for Packet Decoder. Helps to track daily usage.	Displays the packet data analyzed daily in bytes.
Packets on Disk	• Usage in bytes • Per day	Provides the amount of data stored on the disk daily in bytes for the Packet Decoder. Helps to track daily usage.	Displays the amount of packet data usage on the disk daily in bytes.
Logs Processed	• Usage in bytes • Per day	Provides the Log data processed in bytes for Log Decoders. Helps to track daily usage.	Displays the Log data processed daily in bytes.
Users Analyzed	• Users Analyzed • Per day	Provides the number of users analyzed. Helps to track daily usage.	Displays the number of users analyzed daily.

Visualization	Metrics	Objective	Description
Files Analyzed	• Usage in bytes • Per day	Provides the files analyzed in bytes. Helps to track daily usage.	Displays the files analyzed daily in bytes.
Hosts Analyzed	• Hosts Analyzed • Per day	Provides the number of hosts analyzed. Helps to track daily usage.	Displays the number of hosts analyzed daily.
Aggregate Usage - Logs Processed	• Usage in bytes • Per day	Provides an aggregated log of data from all the log services under the throughput license. Helps to track daily usage, detect high value, and take necessary action.	Displays an aggregated log of data from all the log services daily.
Aggregate Usage - Packet Analyzed	• Usage in bytes • Per day	Provides an aggregated packet of data from all the packet services under the throughput license. Helps to track daily usage, detect high value, and take necessary action.	Displays an aggregated packet of data from all the packet services daily.

Visualization	Metrics	Objective	Description
Aggregate Usage - Packets on Disk	• Usage in bytes • Per day	Provides an aggregated packet data stored on the disk from all the packet services under the throughput license. Helps to track daily usage, detect high value, and take necessary action.	Displays an aggregated packet data stored on the disk from all the packet services daily.
Aggregate Usage - File Analyzed	• Usage in bytes • Per day	Provides an aggregated file usage in bytes from all the Malware services under throughput license. Helps to track daily usage, detect high value, and take necessary action.	Displays an aggregated file usage in bytes from all the Malware services daily.
Aggregate Usage - Host Analyzed	• Hosts analyzed • Per day	Provides an aggregated list of hosts from all the endpoint servers under the throughput license. Helps to track daily usage, detect high value, and take necessary action.	Displays an aggregated list of hosts from all the hosts daily.

Visualization	Metrics	Objective	Description
Aggregate Usage - User Analyzed	• Users analyzed • Per day	Provides an aggregated users from all the UEBA servers under the throughput license. Helps to track daily usage, detect high value, and take necessary action.	Displays an aggregated users from all the UEBA servers daily.

New Health and Wellness Monitors

This topic lists the default New Health and Wellness monitors.

Number	Monitor
1.	Respond Server Risk Scoring Unprocessed Alerts High Count
2.	Respond Server Risk Scoring Unprocessed Alert Older Than 24 Hours
3.	Respond Server Risk Scoring Transient Alerts Ignored High Count
4.	Reporting Engine Shared Task Critical Utilization
5.	Reporting Engine Schedule Task Pool Critical Utilization
6.	Reporting Engine Rule(s) Execution Failed
7.	Reporting Engine Report(s) Running > 1 hour
8.	Reporting Engine Report(s) Executions Failed
9.	Reporting Engine Chart(s) Execution Failed
10.	Reporting Engine Available Disk < 5%
11.	Reporting Engine Available Disk < 20%
12.	Reporting Engine Available Disk < 2%
13.	Reporting Engine Available Disk < 10%
14.	NW Offline Service
15.	NW Host High Swap Utilization
16.	NW Host Filesystem Disk Full
17.	NW Host Critical Memory Usage
18.	NW Host Critical Disk Usage
19.	NW Host Critical CPU Usage
20.	Log Decoder Service in Bad State
21.	Log Decoder Log Capture Pool Depleted
22.	Log Decoder Invalid Rules Detected
23.	Log Decoder Dropping > 5% of Logs
24.	Log Decoder Dropping > 10% of Logs
25.	Log Decoder Dropping > 1% of Logs
26.	Log Decoder Database(s) Not Open

Number	Monitor
27.	Log Decoder Capture Rate Zero
28.	Log Decoder Capture Not Started
29.	Endpoint Server to Agent - Incoming UDP Packets Requested
30.	Endpoint Server to Agent - Incoming UDP Packets Rejected
31.	Endpoint Server to Agent - Incoming UDP Packets Queued
32.	Endpoint Server to Agent - Incoming UDP Packets Dropped
33.	Endpoint Server to Agent - Incoming UDP Packets Delayed
34.	Endpoint Server - Machine Persistence Failed
35.	Endpoint Server - Inactive Machine Retention Failed
36.	ESA Correlation - Sessions Behind on Datasources
37.	ESA Correlation - ESA Rule High Memory Usage
38.	ESA Correlation - ESA Rule High CPU Usage
39.	ESA Correlation - ESA Rule Critical Memory Usage
40.	ESA Correlation - ESA Rule Critical CPU Usage
41.	Decoder Service in Bad State
42.	Decoder Packet Capture Pool Depleted
43.	Decoder Invalid Rules Detected
44.	Decoder Dropping > 5% of Packets
45.	Decoder Dropping > 10% of Packets
46.	Decoder Dropping > 1% of Packets
47.	Decoder Database(s) Not Open
48.	Decoder Capture Rate Zero
49.	Decoder Capture Not Started
50.	Contexthub Server Query Response Cache Usage > 80%
51.	Contexthub Server High Query Response Cache Usage
52.	Contexthub Server Database High Disk Usage
53.	Contexthub Server Database Critical Disk Usage
54.	Contexthub Server Critical Query Response Cache Usage
55.	Concentrator Service in Bad State

Number	Monitor
56.	Concentrator Meta Rate Zero
57.	Concentrator Individual Rule(s) Detected
58.	Concentrator Database(s) Not Open
59.	Concentrator Aggregation Stopped
60.	Concentrator > 5 Pending Queries
61.	Broker Session Rate Zero
62.	Broker Service in Bad State
63.	Broker Aggregation Stopped
64.	Broker > 5 Pending Queries
65.	Archiver Service in Bad State
66.	Archiver Aggregation Stopped
67.	Logstash offline
68.	Logstash Persisting data

Uninstall New Health and Wellness

To uninstall New Health and Wellness, perform the following:

1. Take a backup of NetWitness Server host. For more information, see “Disaster Recovery (Back Up and Restore)” topic in the *NetWitness Recovery Tool User Guide*.

   ```
   nw-recovery-tool --export --dump-dir /some/folder --category AdminServer --
   category Search
   ```

   Note: If New Health and Wellness is not installed on NetWitness Server, you must take a backup of the host on which New Health and Wellness is installed.

2. Make sure that the installation or upgrades are not in progress and stop the orchestration server on NetWitness Server host:

   ```
   systemctl stop rsa-nw-orchestration-server
   ```

3. Remove the New Health and Wellness service category (“Search") from the host:

   a. SSH to Admin server

   b. Fetch host details where New Health and Wellness is installed using the following command:
   ```
   mongo localhost/orchestration-server -u deploy_admin -p <deploy_admin-
   password> --authenticationDatabase admin --eval 'db.host.find({
   "installedServices": /.*Search.*/i })'
   ```

   Sample output

   ```
   { "_id" : "56f2a90b-1f03-d09a-fb71-42c2a93958a8", "hostname" :
   "10.10.10.11", "ipv4" : "10.10.10.11", "ipv4Public" : "", "displayName"
   : "adminserver", "version" : { "major" : 11, "minor" : 5, "servicePack"
   : 0, "patch" : 0, "snapshot" : false, "rawVersion" : "11.5.2.0" },
   "lastFailedRefreshAttempt" : NumberLong(0), "refreshAttemptDelayFactor"
   : 0, "thirdParty" : false, "installedServices" : [ "Search",
   "AdminServer" ], "meta" : { "node-zero" : true }, "_class" :
   "com.rsa.asoc.orchestration.host.HostEntity" }
   ```

   c. Remove the "Search" from the `installedServices`.

   IMPORTANT: Do not remove any other category names.

   d. Replace <LIST-OF-CATEGORIES-EXCEPT-SEARCH> with a comma-delimited AND double-quoted list of all the existing installed services found earlier EXCEPT "Search":
   ```
   mongo localhost/orchestration-server -u deploy_admin -p <deploy_admin-
   password> --authenticationDatabase admin --eval 'db.host.update({ "_id"
   : "<hw-node-uuid>" },{$set: {"installedServices" : [ <LIST-OF-
   CATEGORIES-EXCEPT-SEARCH> ]}})'
   ```

   Example
   ```
   mongo localhost/orchestration-server -u deploy_admin -p netwitness --
   authenticationDatabase admin --eval 'db.host.update({ "_id" : "56f2a90b-
   1f03-d09a-fb71-42c2a93958a8" },{$set: {"installedServices" : [
   "AdminServer" ]}})'
   ```

   Sample output

   ```
   MongoDB shell version v4.0.19
   ```

```
connecting to: mongodb://localhost:27017/orchestration-
server?authSource=admin&gssapiServiceName=mongodb

Implicit session: session { "id" : UUID("04e32380-347e-4b7d-a63e-
a094536d7242") }

MongoDB server version: 4.0.19

WriteResult({ "nMatched" : 1, "nUpserted" : 0, "nModified" : 1 })
```

e. Make sure that the "Search" category is removed in the updated host record in the `installedServices` :

```
mongo localhost/orchestration-server -u deploy_admin -p <deploy_admin-
password> --authenticationDatabase admin --eval 'db.host.find({ "_id" :
"<hw-node-uuid>" })'
```

Example

```
mongo localhost/orchestration-server -u deploy_admin -p netwitness --
authenticationDatabase admin --eval 'db.host.find({ "_id" : "56f2a90b-
1f03-d09a-fb71-42c2a93958a8" })'
```

Note: Any inconsistencies can result in unrecoverable errors.

Sample output

```
{ "_id" : "56f2a90b-1f03-d09a-fb71-42c2a93958a8", "hostname" :
"10.10.10.11", "ipv4" : "10.10.10.11", "ipv4Public" : "", "displayName" :
"adminserver", "version" : { "major" : 11, "minor" : 5, "servicePack" : 0,
"patch" : 0, "snapshot" : false, "rawVersion" : "11.5.2.0" },
"lastFailedRefreshAttempt" : NumberLong(0), "refreshAttemptDelayFactor" :
0, "thirdParty" : false, "installedServices" : [ "AdminServer" ], "meta" :
{ "node-zero" : true }, "_class" :
"com.rsa.asoc.orchestration.host.HostEntity" }
```

4. Stop the New Health and Wellness services:

```
systemctl stop rsa-nw-metrics-server elasticsearch opendistro-performance-
analyzer kibana
```

5. Disable the New Health and Wellness services:

```
systemctl disable rsa-nw-metrics-server elasticsearch opendistro-performance-
analyzer kibana
```

6. Uninstall the New Health and Wellness packages using the command:

```
yum erase -y rsa-nw-metrics-server opendistroforelasticsearch
opendistroforelasticsearch-kibana
```

Note: `rsa-nw-shell` (installed with metrics server) is a shared package and should not be removed.

7. Remove the configuration folders or files:

- `/etc/netwitness/metrics-server`
- `/etc/netwitness/platform/elasticsearch`
- `/etc/netwitness/platform/nodeinfo/metrics-server`
- `/etc/netwitness/platform/nodeinfo/elasticsearch-open-distro`

- /etc/netwitness/platform/nodeinfo/kibana-open-distro
- /etc/systemd/system/rsa-nw-metrics-server.service.d
- /etc/systemd/system/elasticsearch.service.d
- /etc/pki/nw/service/bootstrap/metrics-server.completed
- /etc/pki/nw/service/rsa-nw-metrics-server-cert.pem
- /etc/pki/nw/service/rsa-nw-metrics-server.chain
- /etc/pki/nw/elastic
- /etc/pki/nw/kibana
- /var/log/netwitness/metrics-server
- /var/log/kibana
- /etc/collectd.d/rsa-metrics-server.conf
- /etc/logrotate.d/kibana
- /etc/elasticsearch
- /etc/kibana
- /var/lib/elasticsearch
- /var/lib/kibana
- /var/netwitness/elasticsearch

8. Start the orchestration Server on NetWitness Server:
   `systemctl start rsa-nw-orchestration-server`
9. Unregister the New Health and Wellness from the `installedService`:
   a. Find the service IDs for metrics-server, elasticsearch-open-distro, and kibana-open-distro

   Note: Make sure you look for service IDs for the correct host; do not unregister elastic or kibana on an UEBA host.

   ```
   orchestration-cli-client --list-services | grep <hw-node-IP-address>
   ```

   Sample output

   ```
   ID=50082d04-320c-4ce2-8379-00f38ae2d1df, NAME=metrics-server,
   HOST=192.168.1.2:7018, TLS=true

   ID=530ff46a-8793-4e8e-be9c-742193d1705a, NAME=elasticsearch-open-distro,
   HOST=192.168.1.2:9200, TLS=true

   ID=4bad6ea8-e3a4-46ab-a342-34356bea65bb, NAME=kibana-open-distro,
   HOST=192.168.1.2:5601, TLS=true

   ... (other services) ...
   ```

   b. Remove the service IDs returned above for metrics-server, elasticsearch-open-distro, and kibana-open-distro (associated with New Health new Wellness host):

   ```
   orchestration-cli-client --remove-service --id <metrics-server-service-id>
   ```

```
orchestration-cli-client --remove-service --id <elasticsearch-open-distro-service-id>

orchestration-cli-client --remove-service --id <kibana-open-distro-service-id>
```

c. Verify if the services are removed:

```
orchestration-cli-client --list-services | grep <hw-node-IP-address>
```

10. On all hosts, except for UEBA, stop and disable metricbeat:

```
systemctl stop metricbeat

systemctl disable metricbeat
```

Note: For NetWitness Platform without UEBA, you can stop and disable metricbeat on all hosts through salt:
`salt '*' cmd.run 'systemctl stop metricbeat && systemctl disable metricbeat'`

11. (Optional) - If you are not reinstalling New Health and Wellness (on same or other hosts), you can also remove metricbeat package and configuration:

a. Package to uninstall:
`metricbeat`

b. Service configurations to uninstall:

- `/etc/metricbeat`
- `/var/log/metricbeat`
- `mongo account`
- `systemd configuration`

12. Refresh the New Health and Wellness host:
`nw-manage --refresh-host --host-key <node-ip>`
Make sure that the New Health and Wellness service is not installed or running and metricbeat service is not active on the New Health and Wellness host.

13. If you are not reinstalling New Health and Wellness on another host, you must refresh UI hosts (NetWitness Server host and Analyst UI) to update NGNIX:

```
nw-manage --refresh-host --host-key <node-ip>
```

Note: After uninstalling New Health and Wellness, if you want to install New Health and Wellness again, see "New Health and Wellness" in the *Deployment Guide*.

Manage NetWitness Platform Updates

NetWitness issues NetWitness software version updates on a regular basis as it strives to continually improve the product. A software version update consists of a release, service pack, or patch (including security patch) and ancillary software on which the release, service pack, or patch depends. User guides are provided for each software version update release, which include detailed steps for installing the update. It is important that you download the update guide for the release from NetWitness Community (https://community.netwitness.com/) and follow the steps described there. Additional information is available in the "Apply Version Updates to a Host" topic in the *Hosts and Services Getting Started Guide* and in System Updates Panel - Settings Tab.

Reissue Certificates

Introduction

For a secure deployment, NetWitness has installed internal NetWitness-issued certificates such as CA Certificate and Service certificates .

The validity for NetWitness certificates are as follows:

- CA root certificate for 11.x deployment is valid for 10 years
- CA root certificate for 10.6.x deployment is valid for 5 years
- Service certificates are valid for 1000 days

> **Note:** The certificate expiration warning is triggered 30 days prior to expiration.

When these certificates are about to expire or have expired, you must renew and reissue the certificates as soon as possible to avoid any issues with your NetWitness deployment.

> **Note:** You can view the expiration details, by executing the `ca-expire-test-sh` script on the NetWitness Server. For more information, see Reissue root CA security certificates on NetWitness Platform 11.x and download the script.

CA Certificate Reissue

To renew the CA certificates, do the following:

- Before you upgrade from 10.6.x to 11.x, check the expiry and reissue those certificates. For more information, see the Reissue root CA security certificates on NetWitness Platform 11.x.
- If you are on 10.6.x , check the expiry and reissue all the certificates. For more information, see the Reissuing security certificates on NetWitness Platform 10.6.x.

> **Note:** If you have Windows Legacy Collectors (WLC) in your deployment, renew the CA certificate of the WLC after renewing the CA certificate of the NetWitness Admin Server.

Service Certificate Reissue

To renew the Service certificates, do the following:

- If your hosts are on NetWitness Platform 11.3 or later, you must use the `cert-reissue` script. For more information, see the Reissuing Service Certificate .
- If your hosts are on 11.1.x or 11.2.x, you must upgrade the NetWitness Platform to 11.3 or later and run the `cert-reissue` script.

Note: If you have a host that is decommissioned or plan to remove, do not renew the certificate for that host.

Reissuing Service Certificate

You can reissue service certificates in the following two ways.

- All at once
  Reboot NW Server host after the `cert-reissue --host-all` command completes.
- One at a time
  Reissue the NW Server host certificates first, restart the host, then reissue each component host.

IMPORTANT: If you are reissuing certificates for each host individually (one at a time), you must reissue the certificate for the NW Server host before you can reissue certificates for any other host.

When to Use the `--host-all` Argument

Use the `cert-reissue --host-all` command string if you have a large number of hosts. Make sure that:

- All your hosts are running 11.3.0.0 or later.
- All your hosts are online.
- The NW Server host run time services are running.

`cert-reissue` Arguments and Options for All Hosts

The following tables lists the argument you can use to reissue certificates for all hosts at one time. See Troubleshooting Cert-Reissue Command for additional options you can use with Customer Support to troubleshoot errors.

Arguments	Description
`--host-all`	Reissues certificates for all hosts at one time applying system health checks and restarts services. **Note:** If even one host is not online, this command fails. If you have numerous hosts in your deployment, make sure that all hosts are up and running.

Caution: Make sure you do not run this argument on a node or host that you plan to remove or decommission.

When to Use the Individual Host Argument (`--host-key <ID, IP, hostname or display name of host>`)

The `cert-reissue --host-key <ID, IP, hostname or display name of host>` command reissues a certificate for an individual host. You may want to reissue certificates for an individual host if you have a small number of hosts.

Make sure that:

- Each host is running 11.3.0.0 or later.
- Each host is online.
- The NW Server host run time services are running
- You reissue certificates for the NW Server host first.

Note: You must run the command for the NW Server host first and reboot that host before you run the command for each component host.

Reissuing Certificates for All Hosts Except Windows Legacy Collection (WLC) host

Use the `cert-reissue` command to reissue certificates for all hosts except the WLC host with the following procedures.

Running the Cert-Reissue Command for All Hosts

1. SSH to the NW Server host.
2. Submit the appropriate command string.
   `cert-reissue --host-all`

Running the Cert-Reissue Command for an Individual Host

1. SSH to the NW Server host.
2. Submit the appropriate command string:
   `cert-reissue --host-key <ID, IP, hostname or display name of host>`

Reissuing Certificates for a WLC Host

You must use the `wlc-cli-client` utility to reissue certificates for a WLC host (you cannot use the `cert-reissue` command). You also need to specify a number of WLC identification parameters with this utility.

Note: The certificates for a Windows Legacy Server host are stored in the following directories on the host.
`C:\ProgramData\netwitness\ng\logcollector_cert.pem`
`C:\ProgramData\netwitness\ng\logcollector_dh2048.pem`
Th validity period of WLC certificates can range from 2 to 20 years. If you rename or remove the files and restart **NwLogCollector** Service, NetWitness regenerates them.
`/ssl/truststore.pem` - is no longer used in 11.x
Every reissue of a certificate on the Windows Legacy server creates a new private key.

To reissue certificates on a WLC host.

1. SSH to the NW Server host.
2. Submit the following command string.
   ```
   wlc-cli-client --cert-renew --host 10.129.43.13 --port 50101 --use-ssl
   false --username <nwadmin service account> --password <'nwadmin service
   account password'> --ss-username <deploy_admin> --ss-password <'deploy_
   admin password'>
   ```

   Note: nwadmin service account is the WLC rest UI User and **'nwadmin service account password'** is the WLC rest UI password.

Successful Reissue Summary Report

When you run `cert-reissue --host-all` , the following summary report will be displayed if all hosts are online, all run time services are running, and all hosts on version 11.4.0.0 or higher.

```
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
|<host-id>| <IP-address>       |Success |Cert reissue successful                              |
+---------+--------------------+--------+-----------------------------------------------------+
```

Unsuccessful Reissue Summary Reports

You must contact Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897) to troubleshoot problems. You know there is a problem if any `<host-id>` does not return a `Success`**Status**. `Success` indicates that certificates were reissued for a host. The following examples illustrate unsuccessful reissues.

Reissue Failed for Host and Aborted Command

The following three examples illustrate the failure of certificate reissuing for any hosts.

```
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Failed! |failed to connect, is host online?                   |
|<host-id>| <IP-address>       |Failed! |service(s) down                                      |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
+---------+--------------------+--------+-----------------------------------------------------+
```

```
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Failed! |version <version-earlier-than-11.3.0.0> not supported|
|<host-id>| <IP-address>       |Failed! |version <version-earlier-than-11.3.0.0> not supported|
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
+---------+--------------------+--------+-----------------------------------------------------+
```

Reissue Certificate Partially Executed

The NW Server Host certificates were reissued but failed to properly distribute the reissued certificates to one or more component hosts.

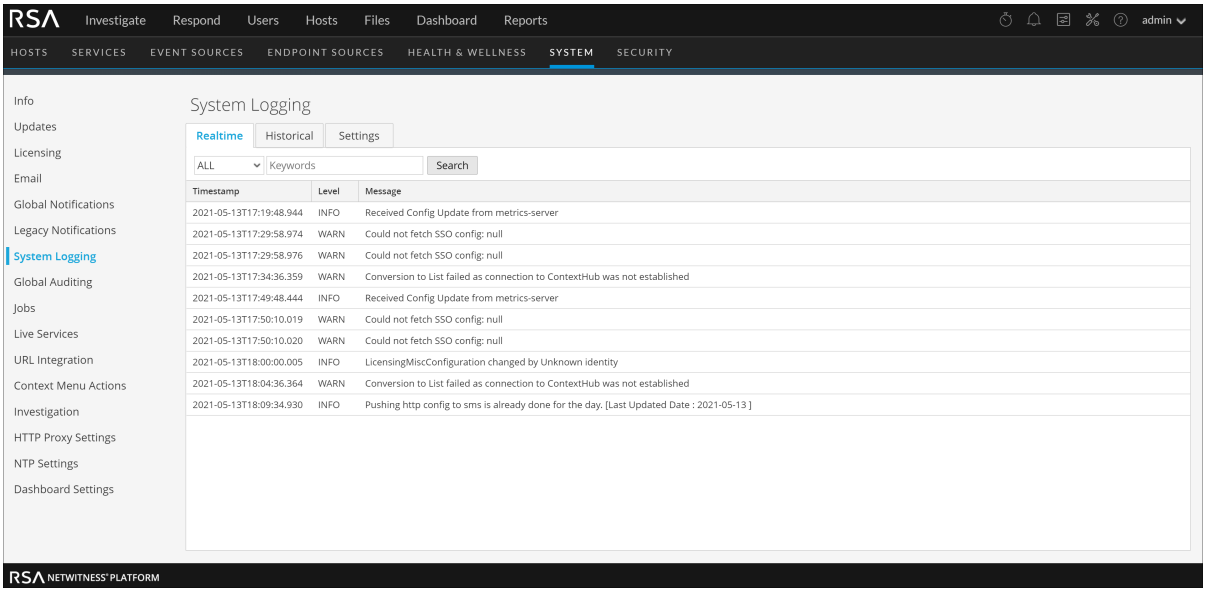
```
+---------+--------------------+
|         | Host               |
+---------+--------------------+
|<host-id>| <IP-address>       |
|<host-id>| <IP-address>       |
+---------+--------------------+
...
+---------+--------------------+--------+-----------------------------------------------------+
|         | Host               | Status | Message                                             |
+---------+--------------------+--------+-----------------------------------------------------|
|<host-id>| <IP-address>       |Partial |Reissue completed, triggers failed                   |
|<host-id>| <IP-address>       |N/A     |[ Skipped… ]                                         |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
|<host-id>| <IP-address>       |N/A     |[ Skipped... ]                                       |
+---------+--------------------+--------+-----------------------------------------------------+
```

Display System and Service Logs

NetWitness provides views into system logs and service logs. When you view service logs, you can select messages for the service or host.

View System Logs

1. Go to **(Admin) > System.**
2. In the options panel, select **System Logging**.

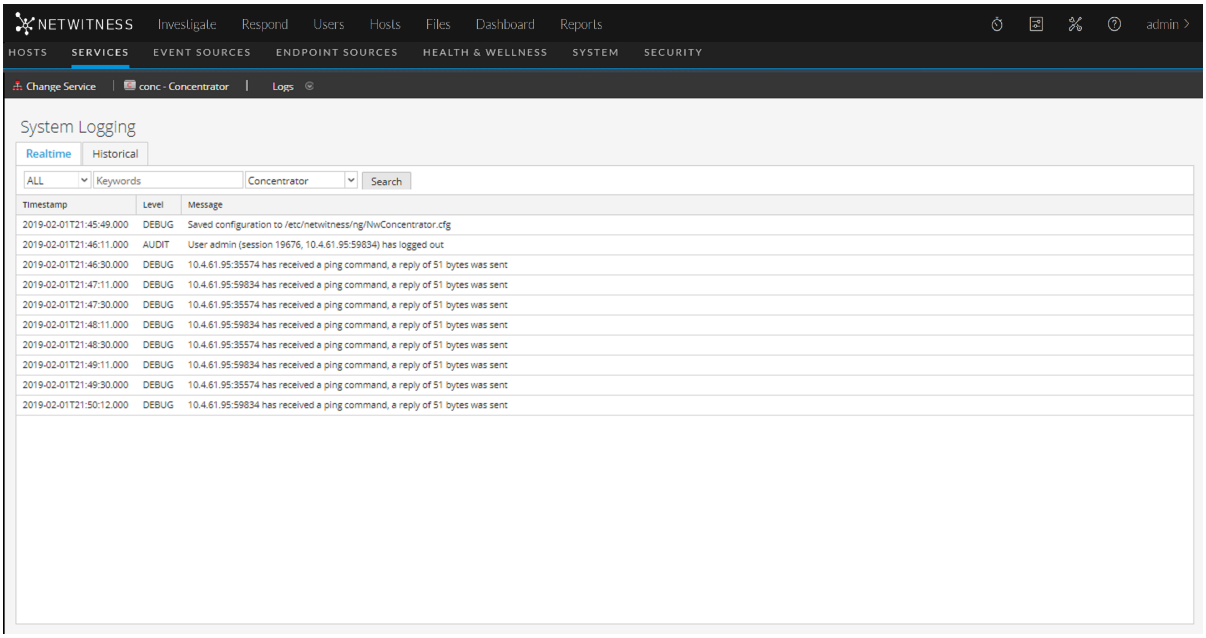


Display Service Logs

To display NetWitness service logs:

1. Go to **(Admin) > Services**.
2. In the **Services** list, select a service.

3. In the **Actions** column, select **View > Logs**.



Filter Log Entries

To filter the results shown in the Realtime tab:

1. (Optional) For system and service logs, select a **Log Level** and a **Keyword**, or both. System logs have seven log levels. Service logs have only six log levels because they do not include the **TRACE** level. The default is **ALL** log entries.
2. (Optional) For service logs, select the Service: host or service.
3. Click **Filter**.

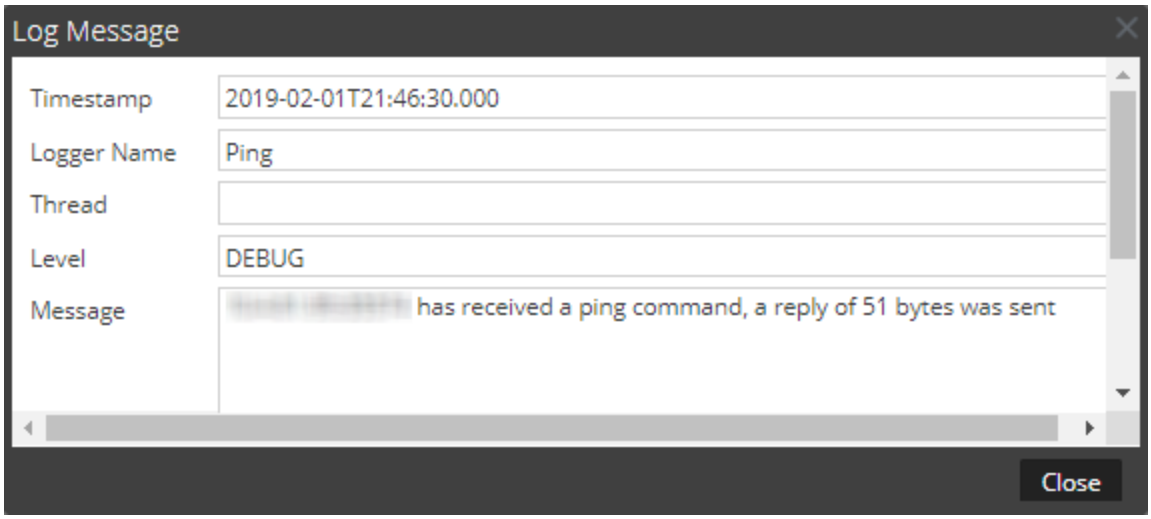
   The view is refreshed with the most recent 10 entries matching your filter. As new matching log entries become available, the view is updated to show those entries.

Show Details of a Log Entry

Each row of the Realtime tab Log list provides the summary information of a log entry. To view complete details:

1. Double-click a log entry.

   The **Log Message** dialog, which contains the Timestamp, Logger Name, Thread, Level and Message, is displayed.



2. After viewing, click **Close**.

Access Reporting Engine Log File

All Log Files

The Reporting Engine stores the following logs in the `rsasoc/rsa/soc/reporting-engine/log` directory:

- Current logs in the `reporting-engine.log` file.
- Backup copies of previous logs in the `reporting-engine.log.*` file.
- All UNIX script logs in the files that have the following syntax: `reporting-engine.sh_timestamp.log` (for example, `reporting-engine.sh_20120921.log`).

The Reporting Engine rarely writes command line error messages to the `rsasoc/nohup.out` file.

Upstart Logs

The Reporting Engine appends the log messages and output written by upstart daemon and the commands used to start the reporting-engine to the `/var/log/secure` directory.

An upstart log file is a system log file, which means that only the root user can read it. The Reporting Engine generates log files, retains backup copies of previous log files, stores UNIX script log files, and appends upstart log files to another directory.

Search and Export Historical Logs

NetWitness provides a searchable view of the NetWitness log or the service log in a paged format. When initially loaded, the grid shows the last page of the log entries for the system or the service. You can export logs from the current view.

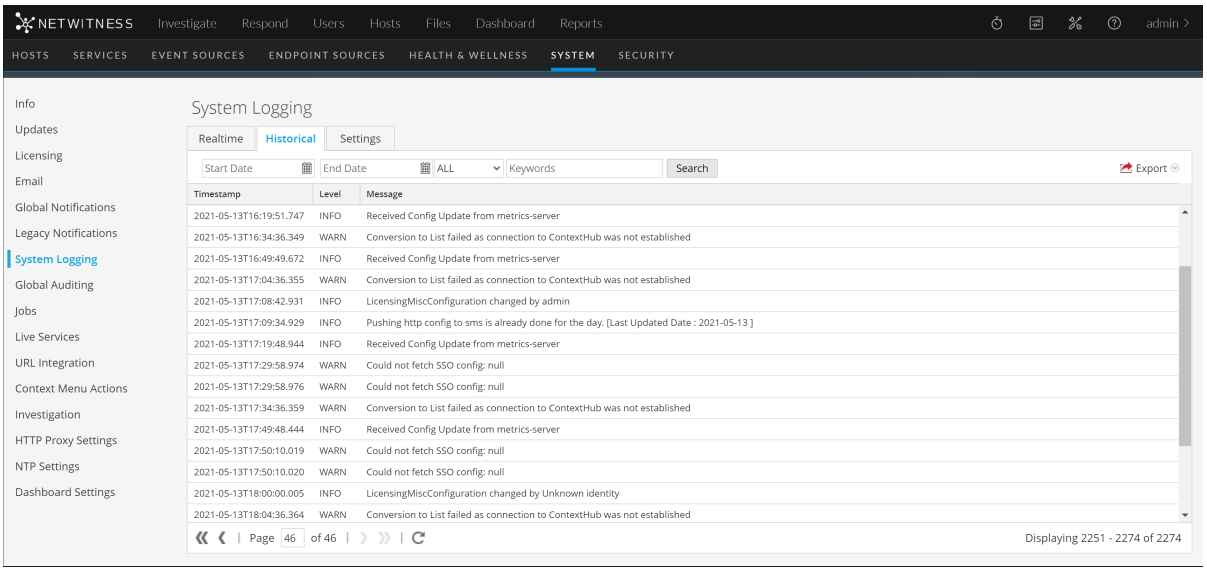
Display the Historical System Log

To display the historical log for the system:

1. Go to  **(Admin) > System.**
2. In the options panel, select **System Logging**.

   The System Logging panel is opened to the **Realtime** tab by default.
3. Click the **Historical** tab.

   A list of historical logs for the system is displayed.



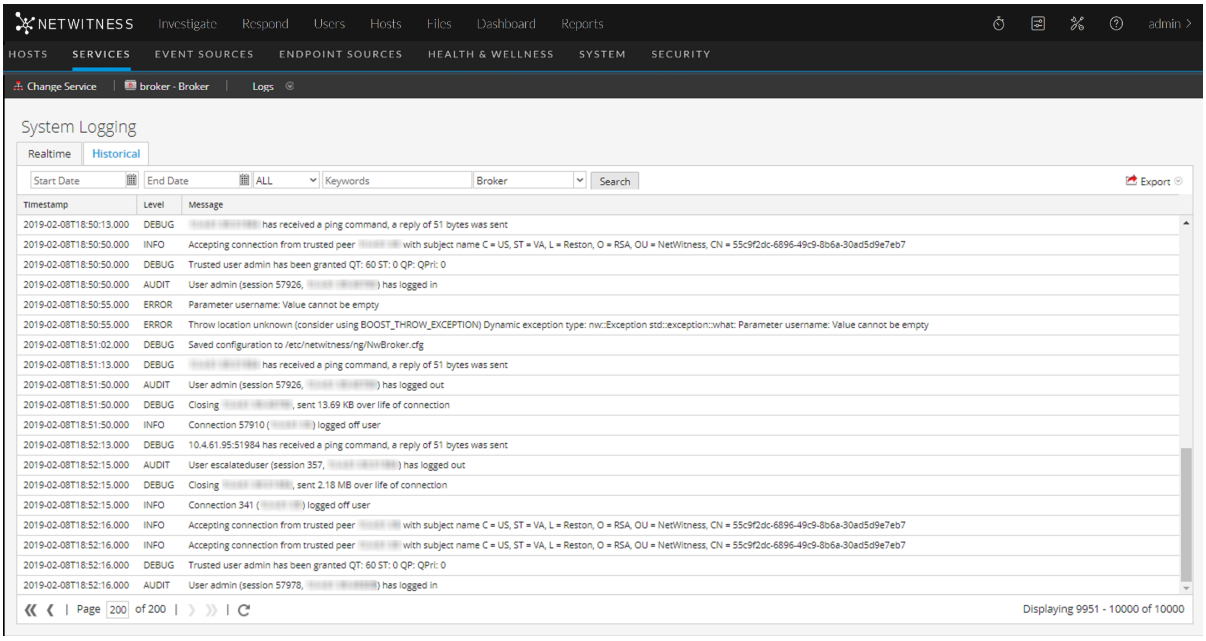
Display a Historical Service Log

To display the historical log for services:

1. Select  **(Admin) > Services**.
2. Select a service.
3. In the **Actions** column, select **View > Logs**.

   The service logs view is displayed with the Realtime tab open.
4. Click the **Historical** tab.

A list of historical logs for the selected service is displayed.



Search Log Entries

To search the results shown in the **Historical** tab:

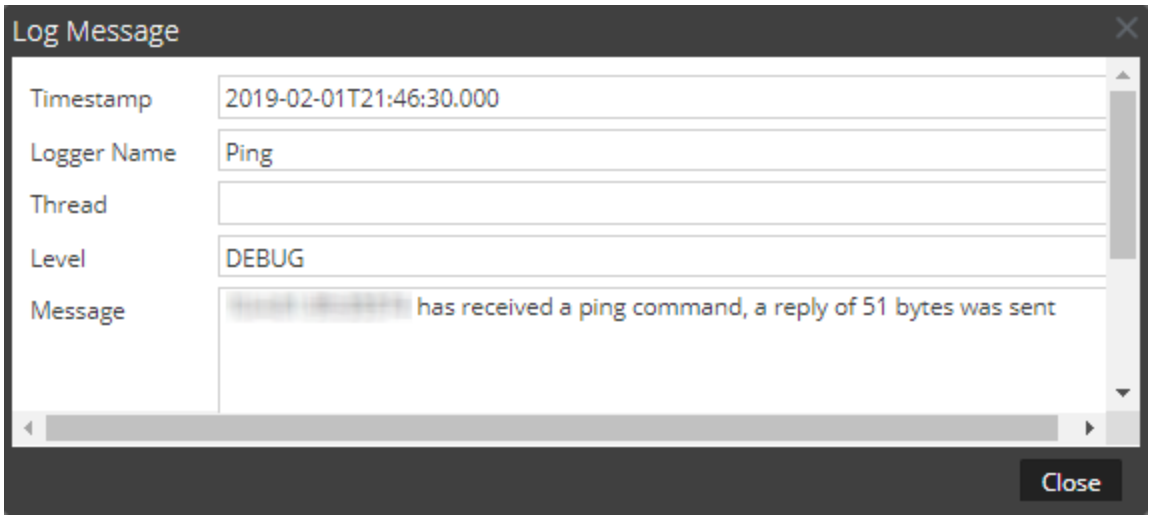
1. (Optional) Select a **Start Date** and **End Date**. Optionally, select a **Start Time** and **End Time**.
2. (Optional) For system and service logs, select a **Log Level** and a **Keyword**, or both. System logs have seven log levels. Service logs have only six log levels because they do not include the **TRACE** level. The default is **ALL** log entries.
3. (Optional) For service logs, select the Service: host or service.
4. Click **Search**.
   The view is refreshed with the most recent 10 entries matching your filter. As new matching log entries become available, the view is updated to show those entries.

Show Details of a Log Entry

Each row of the **Historical** tab Log grid provides the summary information of a log entry. To display all the details for a log message:

1. Double-click a log entry.

   The Log Message dialog, which contains the Timestamp, Logger Name, Thread, Level and Message, is displayed.



2. After viewing, click **Close**.

Page Through Log Entries

To peruse the different pages of the list, use the paging controls on the bottom of the grid as follows:

- Use the navigation buttons
- Manually type the page number you want to view, and press **ENTER**.

Export a Log File

To export the logs in the current view:

Click **Export**, and select one of the drop-down options: **CSV Format** or **Tab Delimited**.
The file is downloaded with a filename that identifies the log type and the field delimiter. For example, a NetWitness system log exported with comma-separated values is named `UAP_log_export_CSV.txt`, and a host log exported with tab-separated values is named `APPLIANCE_log_export_TAB.txt`.

Maintain Queries Using URL Integration

A URL integration provides a way to represent the bread crumbs, or query path, you take when actively investigating a service in the Navigate view. You do not need to display and edit these objects often.

A URL integration maps a unique ID that is automatically created each time you click on a navigation link in the Navigation view to drill into data. When the drill-down completes, the URL reflects the query IDs for the current drill point. The Display Name is displayed in the bread crumb in the Navigate view.

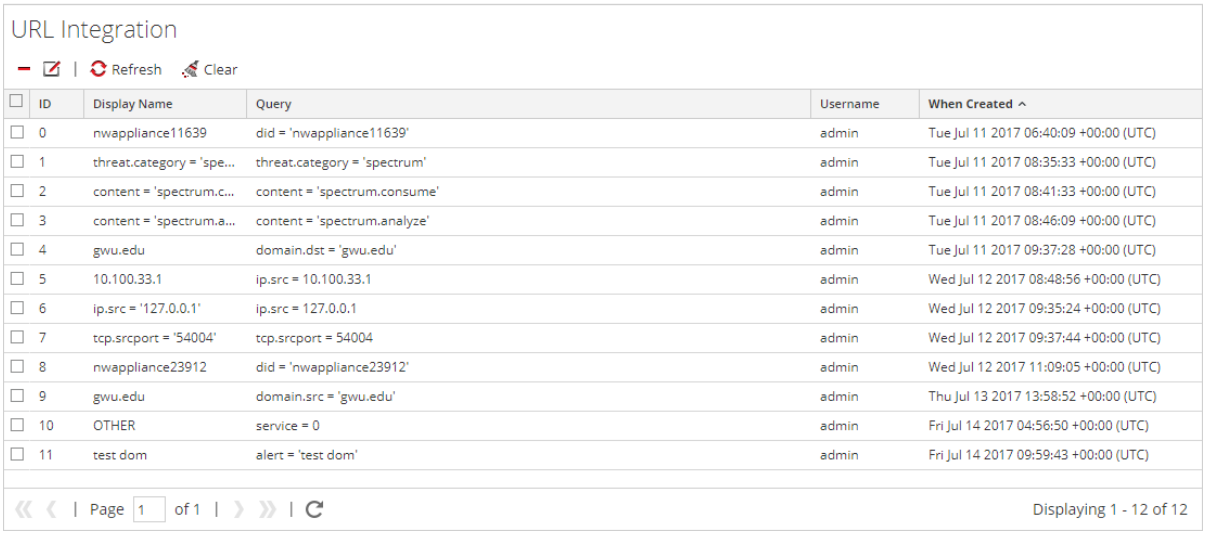
The URL Integration panel provides a list of queries and allows users who have the proper permissions to modify this underlying source of data and analyze the query patterns of other users of the NetWitness system. Within the panel, you can:

- Refresh the list.
- Edit a query.
- Delete a query.
- Clear all queries in the list.

Caution: After a query has been removed from the system, any Investigation URLs that included the ID of that query will no longer function.

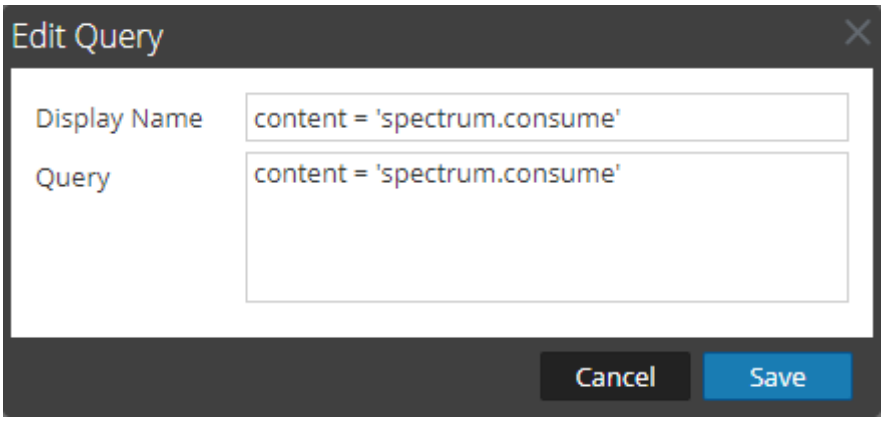
Edit a Query

1. Go to **(Admin) > System**.
2. In the options panel, select **URL Integration**.



3. Select the row in the grid and either double-click the row or click .

   The **Edit Query Dialog** is displayed.



4. Edit the **Display Name** and the **Query**, but do not leave either field blank.
5. To save the changes, click **Save**.

Delete a Query

Caution: After a query has been removed from the system, any Investigation URLs that included the ID of that query will no longer function.

To remove a query from NetWitness entirely:

1. Select the query.
2. Click −

   A dialog requests confirmation that you want to delete the query.
3. Click **Yes**.

Clear All Queries

To clear all queries from the list:

- Click Clear

  The entire list is cleared.

Use a Query in a URI

URL integration facilitates integrations with third-party products by allowing a search against the NetWitness architecture. By using a query in a URI, you can pivot directly from any product that allows custom links, into a specific drill point in the Investigation view in NetWitness.

The format for entering a URI using a URL-encoded query is:

```
http://<nw host:port>/investigation/<serviceId>/navigate/query/<encoded query>/date/<start date>/<enddate>
```

where

- <nw host: port> is the IP address or DNS, with or without a port, as appropriate (ssl or not). This designation is only needed if access is configured over a non-standard port through a proxy.
- <serviceId> is the internal Service ID in the NetWitness instance for the service to query against. The service ID can be represented only as an integer. You can see the relevant service ID from the url when accessing the investigation view within NetWitness. This value will change based on the service being connected to for analysis.
- <encoded query> is the URL-encoded NetWitness query. The length of query is limited by the HTML URL limitations.
- <start date> and <end date> define the date range for the query. The format is <yyyy-mm-dd>T<hh:mm>. The start and end dates are required. Relative ranges (for example, Last Hour) are not supported in this version. All times are run as UTC.

For example:
http://localhost:9191/investigation/12/navigate/query/alias%20exists/date/2018-09-01T00:00/2018-10-31T00:00

Examples

These are query examples where the NetWitness Server is 192.168.1.10 and the serviceID is identified as 2.

All activity on 03/12/2018 between 5:00 and 6:00 AM with a hostname registered

- Custom Pivot: alias.host exists
- https://192.168.1.10/investigation/2...13-03-12T06:00

All activity on 3/12/2018 between 5:00 and 5:10 PM with http traffic to and from IP address 10.10.10.3

- Custom Pivot: service=80 && (ip.src=10.10.10.3 || ip.dst=10.0.3.3)
- Encoded Pivot Dissected:
  - service=80 => service&3D80
  - ip.src=10.10.10.3 => ip%2Esrc%3D10%2E10%2E10%2E3
  - ip.dst=10.10.10.3 => ip%2Esrc%3D10%2E10%2E10%2E3
  - https://192.168.1.10/investigation/2...13-03-12T17:10

Additional Notes

Some values may not need to be encoded as part of the query. For example, commonly the IP source (src) and destination (dst) is used for this integration point. If leveraging a third-party application for integration of this feature, it is possible to reference those without encoding applied.

Manage the deploy_admin Account

The depploy_admin account is used on every NetWitness host, and must be kept in sync between all hosts. Prior to 11.4.1, the process to change the deploy_admin account required administrators to log into every NetWitness host and run the /opt/rsa/saTools/bin/set-deploy-admin-password script on each system. Starting with 11.4.1, the deploy_admin password is centrally managed with the nw-manage script on the NW Server. nw-manage script execution updates the password on all NetWitness component hosts that use the deploy_admin account. The nw-manage script output displays the password update results for each host. If a NetWitness component host is down or unreachable for any reason, the nw-manage script provides an additional option to synchronize the deploy_admin password on the previously unresponsive host with the NW Server when that host becomes available again.

The following procedures describe how to change the deploy_admin password for all hosts in your environment, for hosts in a mixed version environment, and for hosts that are unavailable during the first attempt to change the deploy_admin password.

Change the deploy_admin Account Password

1. Log in to the NW Server host using SSH or the NwConsole.
2. Run the following command:
   nw-manage --update-deploy-admin-pw
   A prompt for the new password is displayed.
3. Enter the new password.

Change the deploy_admin Account Password in a Mixed Version Environment

If you are operating in a mixed version environment (for example, NW Server is on a newer version (greater than or equal to 11.4.1) and the NW component hosts are still on an older version of NetWitness (less than 11.4.1), the nw-manage script prompts you to run the /opt/rsa/saTools/bin/set-deploy-admin-password script on those older component hosts **first**. After the hosts on the older versions are updated, you rerun the nw-manage script on the NW Server with the --skip-version-checks argument.

1. On each component host that is on an older version, reset the deploy_admin password by running the following command:
   /opt/rsa/saTools/bin/set-deploy-admin-password
   This resets the deploy_admin password on all the component hosts with the older versions.
2. Log in to the NW Server host using SSH or the NwConsole and run the following command:
   nw-manage --update-deploy-admin-pw --skip-version-checks
   A prompt for the new password is displayed.
3. Enter the new password.

Change the deploy_admin Account Password for a Component Host that is Unavailable

If a component host is down or otherwise unreachable the first time you run the `nw-manage` script, it is identified as skipped in the `nw-manage --update-deploy-admin-pw` output. When the host is back online, its `deploy_admin` password must be synchronized with the NW Server.

To synchronize the previously unreachable host with the NW Server:

1. Log in to the NW Server host using SSH or the NwConsole.
2. Run the following command:
   ```
   nw-manage --sync-deploy-admin-pw --host-key <ID, IP, hostname or display name of host>
   ```

NW Server Host Secondary IP Configuration Management

Starting with 11.5, the NW Server has a new attribute, `secondary IP` that facilitates automated failover and IP address change management.

Note: For clarification, this attribute is simply another metadata host attribute that is specific to the NW Server; it does NOT describe, nor is it related to, a secondary network interface on the NW Server.

In the failover use case, the active NW Server has a `secondary IP` value that matches the standby server IP address. This `secondary IP` value is known to the other NetWitness hosts, and when the primary NW Server fails to respond in a timely fashion due to an IP address change of the active NW Server, or when there is a scheduled failover to the NW standby server, the other NetWitness hosts and services are configured to automatically attempt to connect to the secondary IP address, which in this case is the NW standby server.

The `secondary IP` attribute is also used for NW Server IP address change and for the NW standby server. The NW Server IP address change procedures automatically populate the secondary IP address of the NW Server to the new IP address and propagate that information to the other NetWitness hosts and services. Then, similar to the failover use case when NW Server's original IP address is no longer valid and the new IP address is active, the NetWitness hosts and services automatically switch to the new secondary IP address, which becomes the new NW Server IP address. If the NW standby server's IP address is changed, the `secondary IP` attribute on the NW Server needs to be updated with the new NW standby server IP address.

The `secondary IP` attribute is managed primarily with automated processes, but there are use cases that require manual management of the NW Server host `secondary IP` attribute, such as an IP address change of the NW standby server or a one-time registration during an NW standby server upgrade.

The `secondary IP` is managed using the `nw-manage` script on the NW Server:

- To add a secondary IP address to the NW Server host, run the following command:
  `nw-manage --add-nws-secondary-ip --ipv4 <secondary ip address>`
- To remove a secondary IP address from the NW Server host, run the following command:
  `nw-manage --remove-nws-secondary-ip --ipv4 <secondary ip address>`
- To view secondary IP addresses assigned to the NW Server host, run the following command:
  `nw-manage --get-nws-secondary-ip`

For information about changing IP addresses, see Change Host Network Configuration.

For information about failover procedures, see "Warm Standby NW Server Host" in the *Deployment Guide for NetWitness Platform*.

Change Host Network Configuration

This topic describes how to change the network configuration for NW Server and component hosts in your environment. **The instructions in this section assume that all the hosts in your environment are on version 11.7 or later.**

Note: If your NW Server is referenced by other NW hosts that use a Network Address Translation (NAT) IP address, and you want to change the NAT IP address, you must remove the old NAT IP address and add the new NAT IP address using the instructions provided in NW Server Host Secondary IP Configuration Management.

This section contains the following procedures:

- Change Host Network Configuration
- Change Network Configuration for Warm Standby (Secondary) Server
- Reconnecting Component Hosts with NW Server Hosts

Note: Changing IPv6 addresses is not supported in version 11.7 and later.

Change Host Network Configuration

Use this procedure to update the network configuration for any host type in your environment for version 11.7 or later.

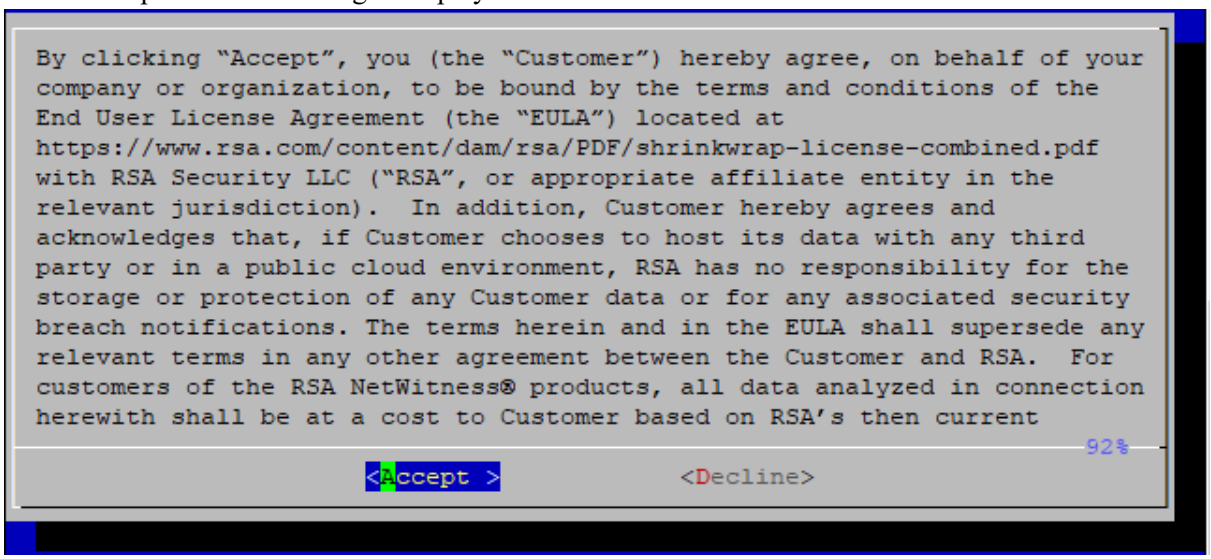
To change the network configuration of a host:

1. From the console, log in to the host for which you wish to change the network configuration.

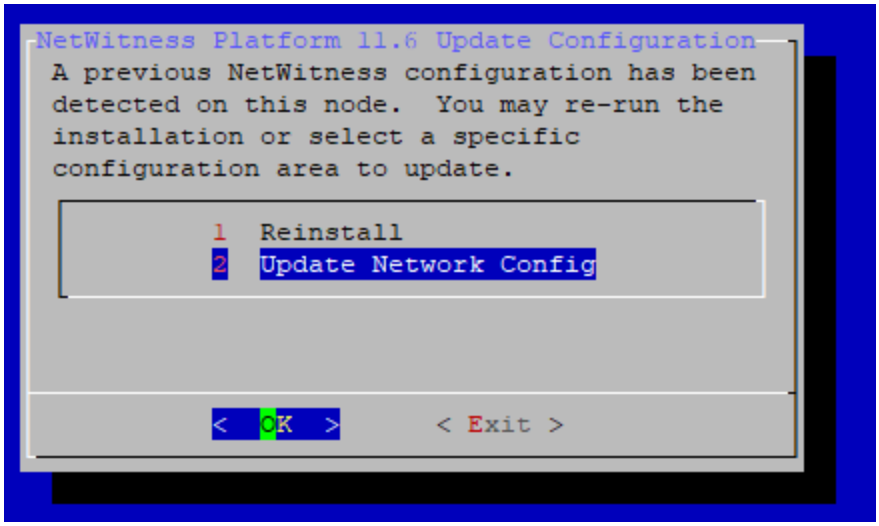
   Note: If you are updating the IP address of your NW Server and you are using DHCP, run the following command before you go to step 2:
   ```
   nw-manage --add-nws-secondary-ip --ipv4 <new DHCP allocated ip address of NW Server>
   ```

2. Run the following command:
   ```
   nwsetup-tui
   ```

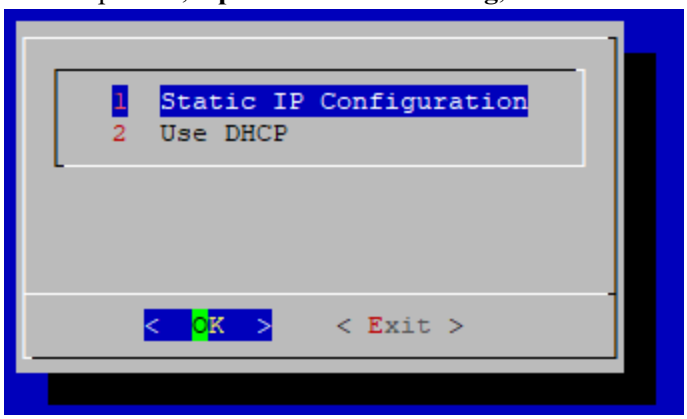
The nwsetup-tui license dialog is displayed.



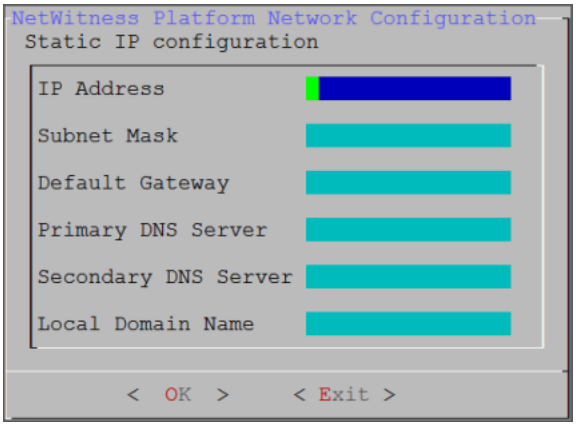
3. Click Accept. The NetWitness Platform Update Configuration dialog is displayed:



4. Select option 2, **Update Network Config**, and click **OK**.



5. Select option 1, **Static IP Configuration**, and click **OK**.
   The NetWitness Platform Network Configuration Static IP configuration dialog is displayed.



6. Enter the new network and DNS configuration and click **OK**.
   The new network and DNS configuration is applied to the host.

   Note: While changing the IP address, the user interface may become temporarily unavailable while the update is in process. The user interface will come back up shortly.

Note: After upgrading the NW Server host or a component host to 11.7 or later version, review the contents of the `/etc/hosts.user` file for any obsolete host entries. The `/etc/hosts.user` file contains system and user-generated entries that are not managed by NetWitness Platform. However, entries from `/etc/hosts.user` are merged with NetWitness Platform-generated host mappings to create and update `/etc/hosts`. To avoid conflicts with NetWitness Platform-generated mappings, and to avoid generating connectivity errors resulting from an IP address change, RSA recommends that you remove any entries in `/etc/hosts.user` that include a non-loopback IP address of a NetWitness Platform host. After updating `/etc/hosts.user`, you must refresh the system by running the following command:

```
nw-manage --refresh-host --host-key <ID, IP, hostname or display name of host>
```

Note: While changing a host's IP address or during failover, component hosts can become disconnected from NW Server hosts. Follow these steps to reconnect a host system to its NW Server system.
1. Log in to the component host using SSH or the console.
2. Run the command `nw-manage --override-nws-ip --ipv4 <current IP address of the NW Server.`
When this command completes, the component host is reconnected to the NW Server at the specified IP address.

Follow the steps in the sections that apply to your environment.

- SSO
- Reporting Engine
- UCF
- PAM

- ECAT
- RSA NetWitness Orchestrator (By Demisto)
- Audit Logging
- Health and Wellness
- Malware Analysis
- Windows Legacy Collection

SSO

Update Configuration for Single Sign-On

Note: You must disable SSO configurations ONLY when NW Server IP is changed.

When the host network is configured with a new IP address, the SSO configurations also must be updated.

To do this:

1. Disable the SSO configuration using `nw-shell` after failover from new IP.
   To resolve this issue you must disable SSO manually, using the following commands:
   a. SSH to `admin server node`.
   b. Connect to `nw-shell`.
   c. Connect to admin server service using the `connect --service admin-server` command.
   d. Log in to admin server using the `login` command.
   e. Enter the admin username and password.
   f. Execute the following commands:
      - `cd /rsa/security/authentication/web/saml/sso-enabled`
      - `set false`
      - `logout`
      - `exit`

- `systemctl restart rsa-nw-admin-server`



2. Change the host IP address to the new IP.
3. Generate the new metadata and reupload it in ADFS. For more information, see the "Configure SAML 2.0 provider settings for portals" topic in the Microsoft documentation.

For more information, see the "Troubleshooting" topic in the *System Security and User Management Guide.*

Reporting Engine

Update Configuration for Reporting Engine

Note: You must update the Reporting Engine configurations ONLY when NW Server IP is changed.

When the host network is configured with a new IP address, you must update and verify the Reporting Engin configurations. The hostname for NetWitness configurations under the Output Actions must be updated with the new IP.

To manually configure the new IP, perform the following steps:

1. Log in to NetWitness Platform.
2. Navigate to  (Admin) > **Services > Reporting Engine > View > Config**.
3. Click the **Output Actions** tab.

4. Add the new IP address in the **Hostname** field.
5. Click **Apply**.

UCF

To enable UCF to communicate with NetWitness Platform:

1. On the UCF server, execute the `runConnectionManager.bat` file (the same file that is used for adding connection details).
2. Select **Option #2, Edit endpoints**.
3. Select the NW Server connection from the options that are displayed.
4. When you are prompted for Host Address (the old IP address is shown in parentheses) enter the new IP address.

   Note: Do not change any other setting.

PAM

If you have PAM configured, after the failover, you must configure the system again using the instructions in the "Configure PAM Login Capability" topic in the *System Security and User Management Guide*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

ECAT

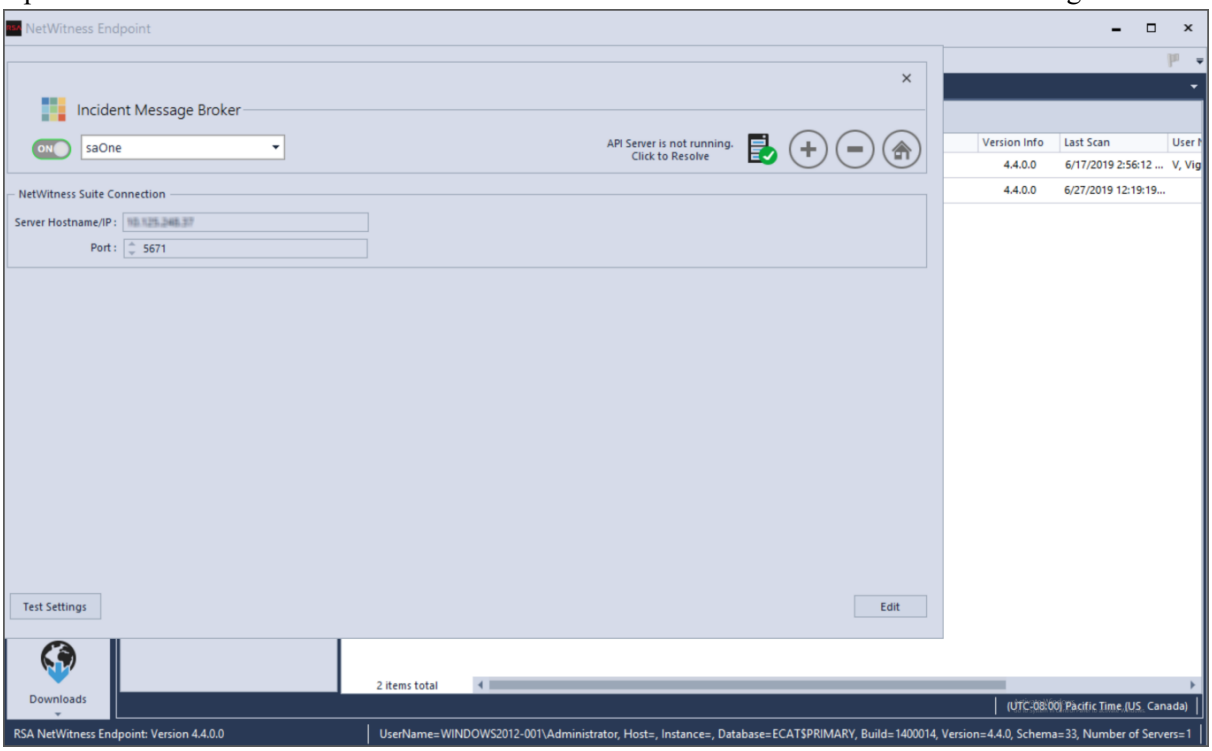
Update the following services:

- Incident Message Broker
- NetWitness Suite
- Syslog Server Settings

Incident Message Broker

1. Log in to the NetWitness Endpoint user interface and go to **Configure** > **Monitoring and External Components Configuration** > **Incident Message Broker**.

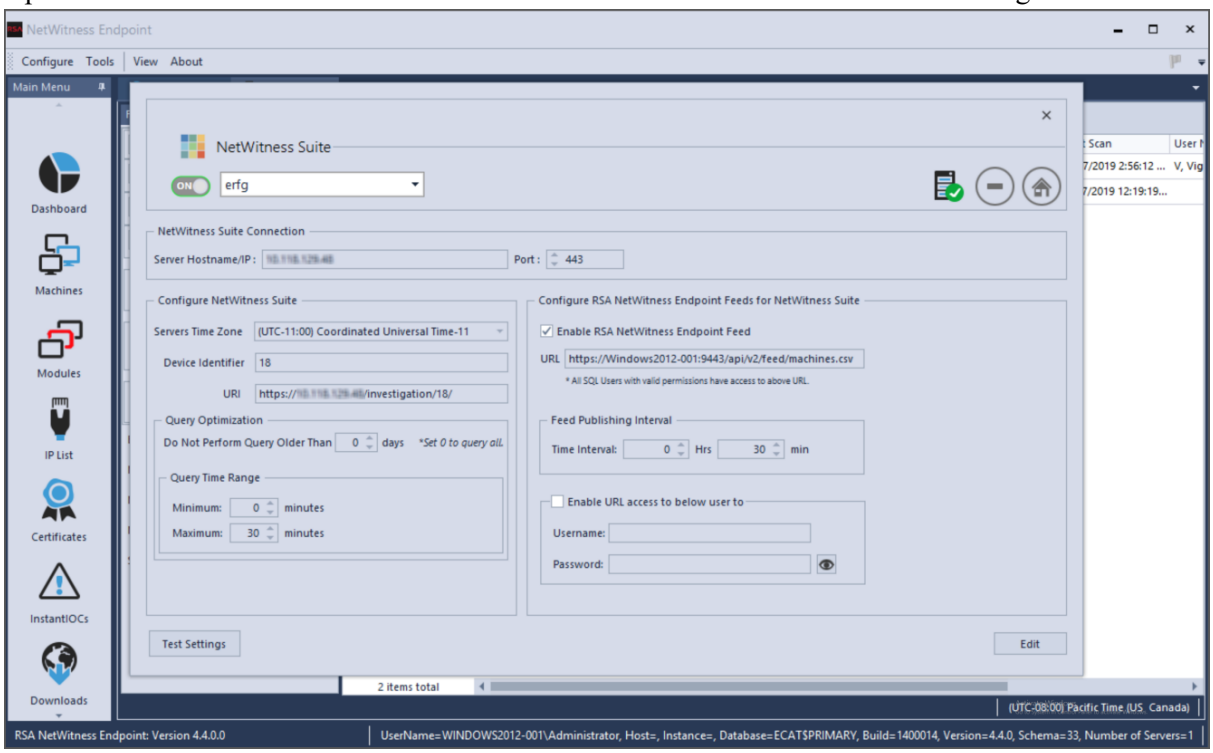
2. Update the server Hostname and IP Address to the current active server and test the settings.



NetWitness Suite

1. Log in to the NetWitness Endpoint user interface and go to **Configure** > **Monitoring and External Components Configuration** > **Netwitness Suite**.

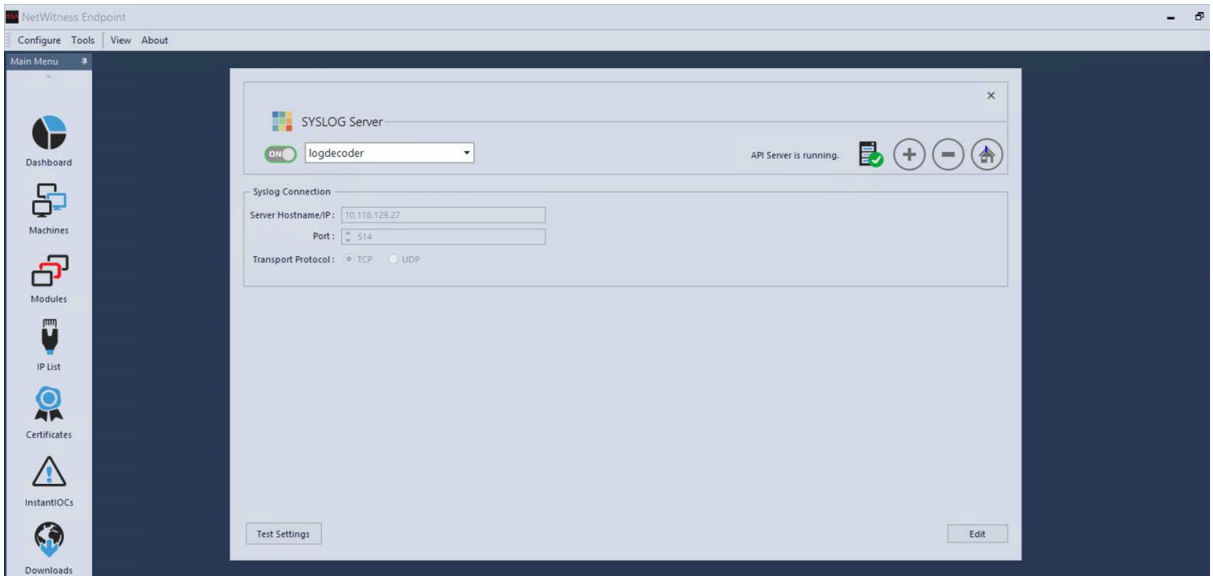
2. Update the server Hostname and IP address to the current active server and test settings.



Syslog Server Settings

If you are forwarding syslog messages to a NetWitness Platform Log Decoder, update the syslog server settings to point to the new IP address of the Log Decoder host.

1. Log in to the NetWitness Endpoint user interface and go to **Configure** > **Syslog Server**.

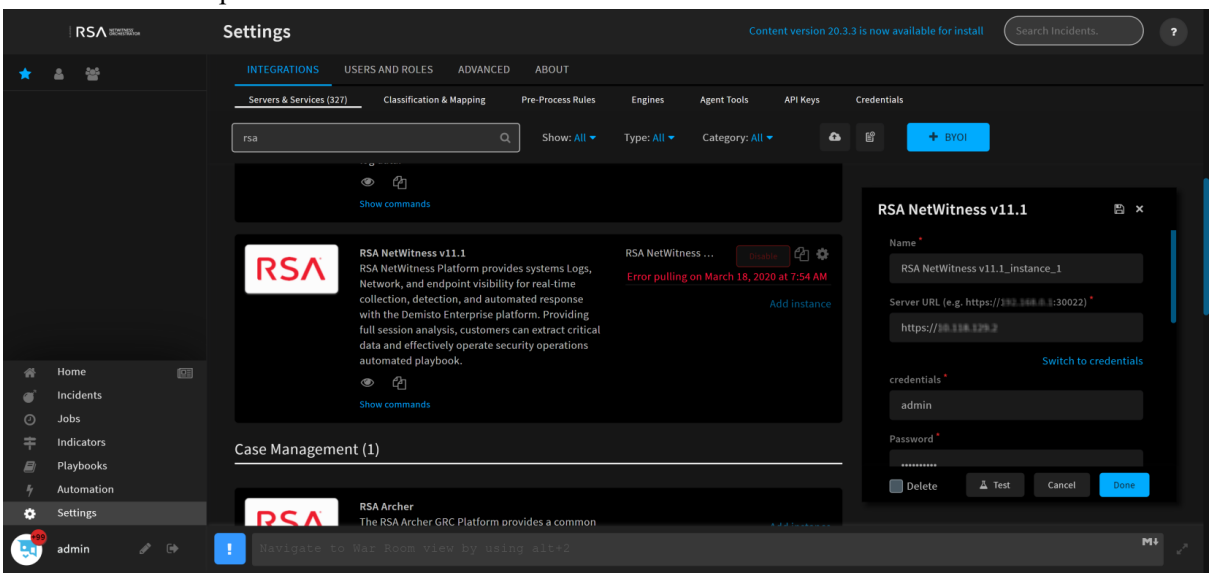


2. Select **logdecoder**, and in **Server Hostname/IP**, enter the new IP address of the Log Decoder host.

RSA NetWitness Orchestrator (By Demisto)

Update the Current Active NW Server to Fetch Respond Incidents and Alerts

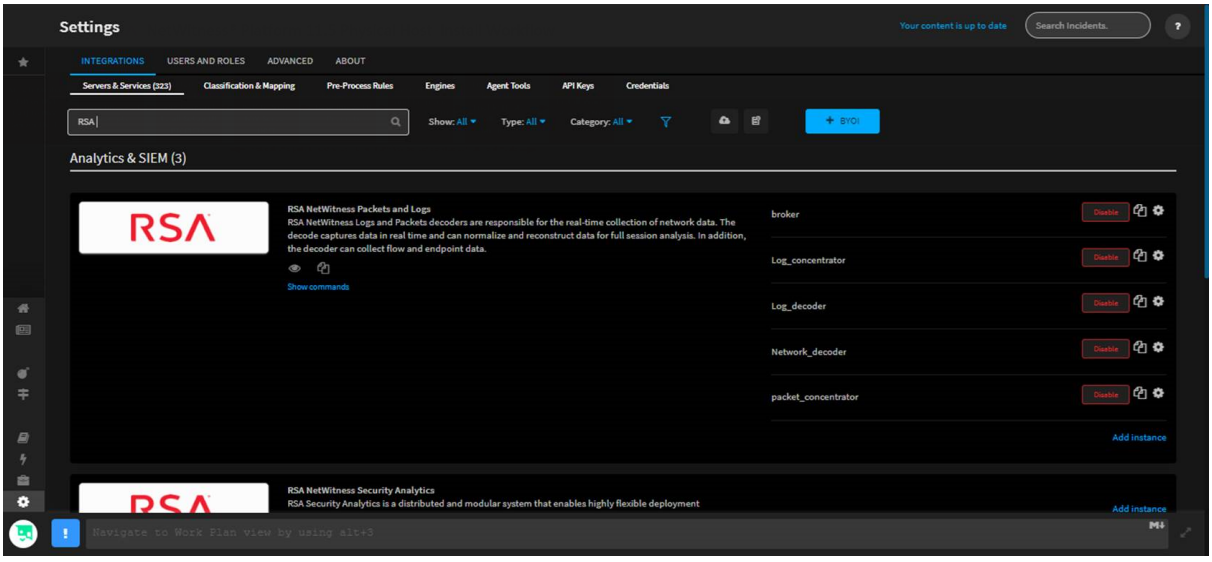
1. Log in to Orchestrator and go to **Settings** > **server&services**.
2. Edit the RSA NetWitness V11.1 instance by updating the server URL to the current active NW Server to fetch respond incidents and alerts.



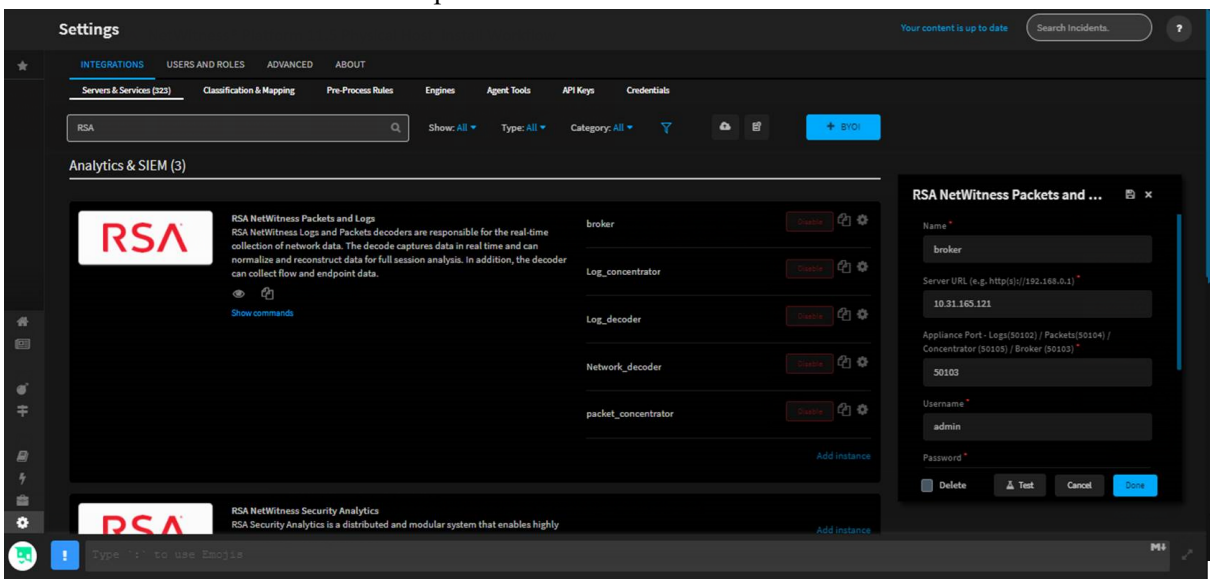
Update Component Hosts Acting as Data Sources

If you change the IP address of a component host, for example, a Concentrator, Network or Log Decoder, or Broker, that is acting as data source to the Orchestrator, update the following settings to point to the new IP address of the host.

1. Log in to Orchestrator and go to **Settings** > **server&services** and select the component host.



2. Enter the new IP address of the component host in **Server URL** and click **Done**.



Audit Logging

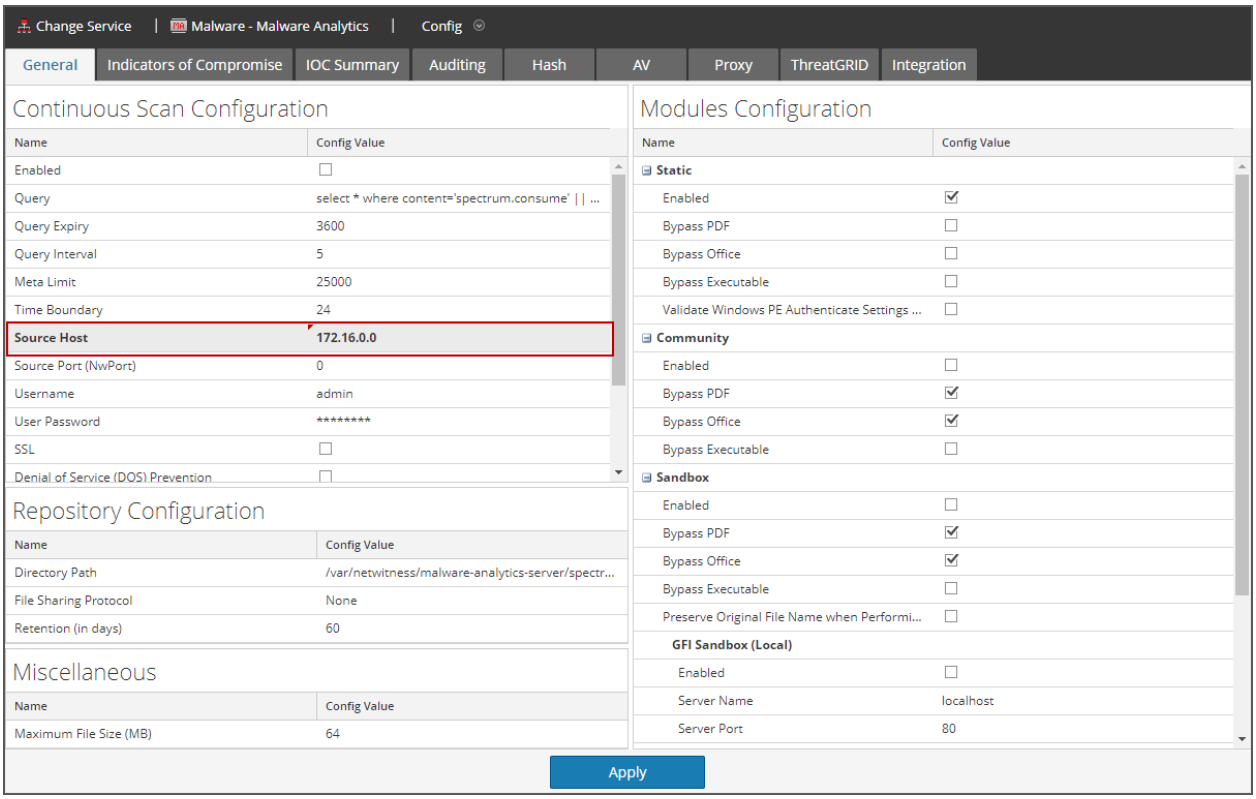
If you have changed the IP address of the NW Server, you must reconfigure audit logging. For instructions, see "Configure Global Audit Logging" in the *System Configuration Guide*.

Health and Wellness

If you have any Health and Wellness rules that contain IP addresses that have been changed, you must update those rules with the new IP addresses. For information about managing Health and Wellness rules, see "Monitor Health and Wellness using NetWitness Platform UI" in the *System Maintenance Guide*.

Malware Analysis

Source host IP address changes are not updated in the NetWitness user interface for Malware Analysis continuous scan configurations. You must manually update this configuration in the Malware Analysis Config view > **General** > **Continuous Scan Configuration** and update the source host IP address to the new host IP address.



Windows Legacy Collection

On occasion, you may need to change the IP address of your Windows Legacy Collector. You may also need to edit any Destination Groups that you have configured.

Change WLC IP Address

The following procedure describes how to change the IP address for your system.

1. Log onto the Windows Legacy Collector system and manually change the IP address on the system.
2. In the UI, confirm that the Log Collector service corresponding to the WLC system shows up in error (Red). It might take some time for it to reflect the changed status.
3. On the NetWitness Server, use the **nw-manage** utility to view the host information for the WLC using the following command:

   ```
   nw-manage --list-hosts
   ```

   Sample output from running the command is shown here:

   ```
   {
   "id" : "fdb8150c-e040-459e-8cc5-3c60ec2c65ae",
   "displayName" : "WLC-HOST-104",
   "hostname" : "10.101.216.102",
   "ipv4" : "10.101.216.102",
   "ipv4Public" : null
   } ]
   ```

   You use the value of **"id"** from your output in the following step.

4. Use the **nw-manage** utility to change the IP address of the WLC. For the **host-id** argument, use the value for the **"id"** that you noted from step 3. For the **ipv4** value, use the new IP Address to which you are changing.

   ```
   nw-manage --update-host --host-id "fdb8150c-e040-459e-8cc5-3c60ec2c65ae" --ipv4 10.101.216.105
   ```

5. After you see the message that the previous command ran successfully, go to the NetWitness Server UI and verify that the WLC service is running without any errors.

Edit Destination Groups For Log Collectors and VLCs

The Windows Legacy Collector is often configured with Destination Groups to forward events to Log Collectors or Virtual Log Collectors. If the IP address of any such Destination LC or VLC is changed, the Windows Legacy Collector can no longer forward events. To remediate this, you must edit the Destination groups for the WLC, making sure to select the new LC or VLC IP Address.

Change Network Configuration for Warm Standby (Secondary) Server

You can change the network configuration of a warm standby (secondary server) by following these steps:

1. Follow the steps described in Change Host Network Configuration to change the IP address on the secondary server.
2. Log in to the active NW Server and remove the previous secondary server IP address by running the following command:

   ```
   nw-manage --remove-nws-secondary-ip --ipv4 <previous standby server ip address>
   ```

3. On the active NW Server, add the new standby server secondary IP address value by running the following command:

   ```
   nw-manage --add-nws-secondary-ip --ipv4 <new standby server ip address>
   ```

4. Schedule the backup of the primary NW Server and the copying of backed-up data to the secondary NW Server. See step 18 in "Setup Secondary NW Server in Standby Role" in the *Deployment Guide for NetWitness Platform*.

For information about configuring warm standby servers, see "Warm Standby NW Server Host" in the *Deployment Guide for NetWitness Platform* and NW Server Host Secondary IP Configuration Management.

Reconnecting Component Hosts with NW Server Hosts

While changing a host's IP address or during failover, component hosts can become disconnected from NW Server hosts. Follow these steps to reconnect a host system to its NW Server system.

1. Log in to the component host using SSH or the console.
2. Run the following command:

   ```
   nw-manage --override-nws-ip --ipv4 <current IP address of the NW Server>
   ```

When this command completes, the component host is reconnected to the NW Server at the specified IP address.

Manage Custom Host Entries

If you have hosts in your environment that use custom entries, or if you want to use NAT IPv4 addresses for your hosts, this topic provides instructions for configuring these settings.

Manage Custom Host Entries in /etc/hosts

If you identify hosts in your environment that need custom entries, you can add custom entries for your hosts in the `/etc/hosts.user` file.

To add custom host entries:

1. From the console, log in to the host on which to define custom entries.
2. Add the custom entry to the `/etc/hosts.user` file.

To update or refresh a custom entry in `/etc/hosts`:

Run the following command:

```
nw-manage --refresh-host --host-key <ID, IP, hostname or display name of host>
```

Manage Public or NAT IPv4 Addresses for Hosts

To add or update public or Network Address Translation (NAT) IPv4 addresses for NW Hosts, run the following command:

```
nw-manage --update-host --host-id <NW Server Host UUID> --ipv4-public <IP address>
```

Manage Custom Jetty Configuration

If you want to have custom jetty configurations or additional JAVA_OPTIONS that persist during the upgrade, this topic provides the instructions for configuring custom jetty settings.

To add custom jetty configuration

1. From SSH prompt, create a copy of jetty configuration as `jetty.user` using the below command.
   `$ cp /etc/default/jetty /etc/default/jetty.user`
2. Edit the `/etc/default/jetty.user` and remove all lines EXCEPT one of the `JAVA_OPTIONS` additional settings lines.
   For example, `JAVA_OPTIONS="${JAVA_OPTIONS} -Drsa.primary.host=true "`
3. Edit the line with the option (in this example increasing the max memory option from 8G to 24G):
   `JAVA_OPTIONS="${JAVA_OPTIONS} -Xmx24G "`
4. Make sure to use the `JAVA_OPTIONS="${JAVA_OPTIONS} "` format to EXTEND the `JAVA OPTIONS`.

5. REMOVE the `[ -f /etc/default/jetty.user ] && source /etc/default/jetty.user` file.
6. Save the file.
7. Restart the jetty service.
   `systemctl restart jetty`

Note: When you overwrite the jetty.user custom config file, ensure that all the jetty configurations (JAVA_OPTIONS) present in /etc/default/jetty are available in this file. If any jetty configuration (JAVA_OPTIONS) is not available in the jetty.user file, you must copy it to the file and then restart the jetty service.

Configure FIPS Support

NetWitness 11.x ships with FIPS-validated 140-2 Cryptographic Modules that support all cryptographic operations within NetWitness. NetWitness leverages two modules that support a level three design assurance:

- RSA BSAFE Crypto-J
- RSA OwB

Both modules have been certified with an operational environment comparable to the standard NetWitness configuration.

By default, the cryptographic modules enforce the usage of FIPS-certified cipher suites wherever possible. For exceptions, refer to the information below and to the release notes. For additional information about the FIPS modules, see https://csrc.nist.gov/publications/detail/fips/140/2/final.

The RSA BSAFE Crypto-J FIPS Certificate number is 3172, and OwB uses the CCME FIPS Module in FIPS-approved mode.

In 11.x, FIPS is enabled on all services except Log Collector. This includes Log Decoder and Decoder if they were FIPS-enabled in 10.6.x or any previous version. FIPS cannot be disabled on any services except for Log Collector, Log Decoder and Decoder.

Note: For a fresh installation of 11.x, by default, all core services will be FIPS enforced except Log Collector and Log Decoder. FIPS cannot be disabled on any services except for Log Collector, Log Decoder and Network Decoder.

Note: For upgrades to 11.x from previous versions, the following conditions apply for the Log Collector, Log Decoder and Decoder services:
- Log Collector is not FIPS enabled after upgrading to 11.x, even if FIPS was enabled in a previous version. You must enable FIPS support after upgrading to 11.x. See the instructions in FIPS support for Log Collectors.
- If FIPS was enabled for the Log Decoder and Network Decoder services in a prevous version, FIPS will also be enabled in 11.x. However, if Log Decoder and Network Decoder were NOT FIPS enabled in a previous version, they will not be enabled in 11.x, and you can manually enable FIPS for these services if required. See the instructions in FIPS support for Log Decoders and Decoders.

FIPS support for Log Collectors

To enable FIPS for Log Collectors:

1. Stop the Log Collector service.
2. Open the `/etc/systemd/system/nwlogcollector.service.d/nwlogcollector-opts-managed.conf` file.
3. Change the value of the following variable to **off** as described here:

   `Environment="OWB_ALLOW_NON_FIPS=on"`

   to

   `Environment="OWB_ALLOW_NON_FIPS=off"`
4. Reload the system daemon by running the following command:

   `systemctl daemon-reload`
5. Restart the Log Collector service.
6. Set the FIPS mode for the Log Collector service in the UI:

   Note: This step is not required if you are upgrading from 10.6.x to 11.x and FIPS was enabled in 10.6.x.

   a. Go to  (Admin) > **Services**.

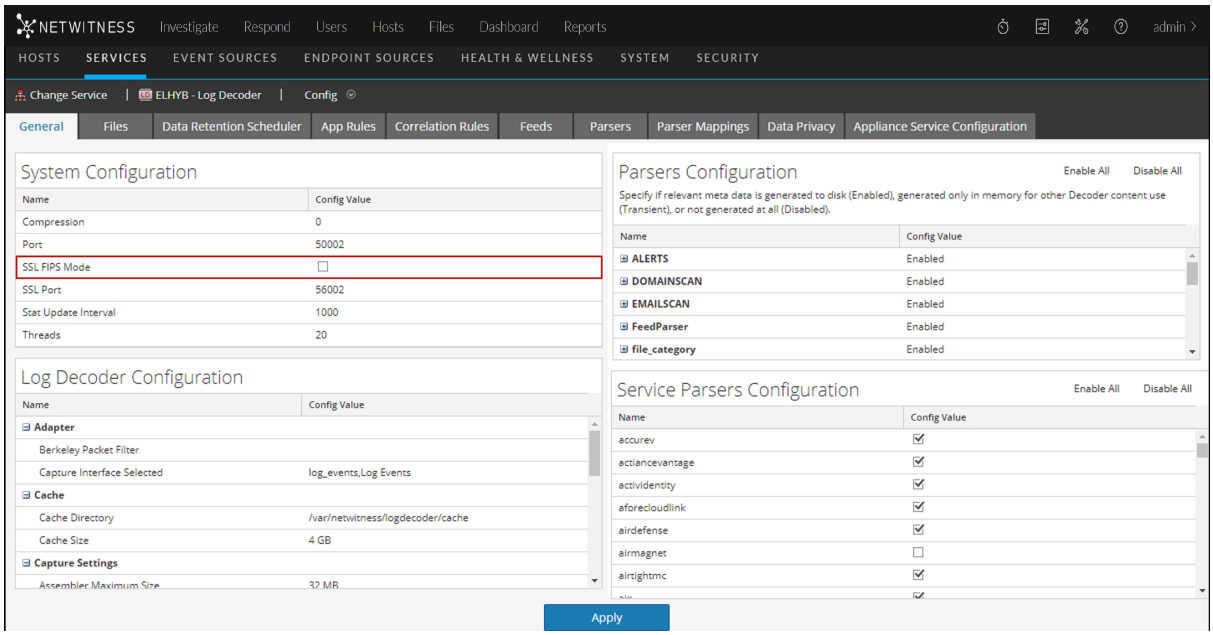
   b. Select the Log Collector service and go to **View** > **Config**.

   c. In SSL FIPS Mode, select the checkbox under Config Value and click **Apply**.

FIPS support for Log Decoders and Decoders

To enable FIPS for Log Decoders and Decoders that did not have FIPS enabled in 10.6.x:

1. Go to  (Admin) > **Services** and select a Log Decoder or Network Decoder service.
2. Select **View** > **Config**, and in **System Configuration**, enable **SSL FIPS Mode** by selecting the check box in the **Config Value** column.



3. Restart the service.
4. Click **Apply**.

DISA STIG

Note: 11.3.1 feature - DISA STIG (Defense Information Systems Agency Security Technical Implementation Guide) support was introduced in NetWitness Platform 11.3.1. Versions 11.0.0.0 to 11.3.0.0 do not support DISA STIG.

NetWitness Platform version 11.3.1 supports all Audit Rules in the DISA STIG Control Group. The supported version for DISA STIG is Red Hat Enterprise Linux V3R2. NetWitness will expand its support of STIG rules in future NetWitness Platform versions.

This section includes the following topics.

How STIG Limits Account Access

NetWitness Passwords

Generate the OpenSCAP Report

Manage STIG Controls Script (`manage-stig-controls`)

Rules List

Exceptions to STIG Compliance

IMPORTANT: All rules are enabled by default except for **control goup 1-ssh-prevent-root** an **control group 3-fips-kernel**. You can enable or disable rules by control group using the `manage-stig-controls` script.

How STIG Limits Account Access

The STIG hardening RPM helps to lock down information, systems, and software, which might otherwise be vulnerable to a malicious computer attack by limiting account access to a system. For example, the STIG script:

- Ensures that the account password has a length, complexity, expiration period, and lockout period that are in accordance with DISA best practices.
- Applies auditing and logging of user actions on the host.

NetWitness Passwords

NetWitness Platform requires passwords that are STIG compliant.

Generate the OpenSCAP Report

Security Content Automation Protocol (SCAP) is a line of standards or rules managed by the National Institute of Standards and Technology (NIST). It was created to provide a standardized approach to maintaining the security of enterprise systems, such as automatically verifying the presence of patches, checking system security configuration settings, and examining systems for signs of compromise.

The OpenSCAP report evaluates your environment against the SCAP rules. The results are sent to the `HOSTNAME-ssg-results.(XML|HTML)` depending on the output format you select.

Disable Rules in OpenSCAP Report that Hang the Report

There may be STIG rules that you do not want to include in the OpenSCAP report because they make the report hang. Use the following command to disable items on the SCAP report:

```
sed -i 's/select idref="rule-id" selected="true"/select idref="rule-id"
selected="false"/g' /usr/share/xml/scap/ssg/content/ssg-rhel7-xccdf.xml
```

where `rule-id` is the Rule ID that you can replace with the Rule ID that may hang during a test.

For example, the report has a rule ID called `partition_for_audit` (shown as `Rule ID: partition_for_audit`). If you disable a rule, OpenSCAP does not check against that rule. This means that you need to check for compliance to the `partition_for_audit` rule manually.

Install OpenSCAP

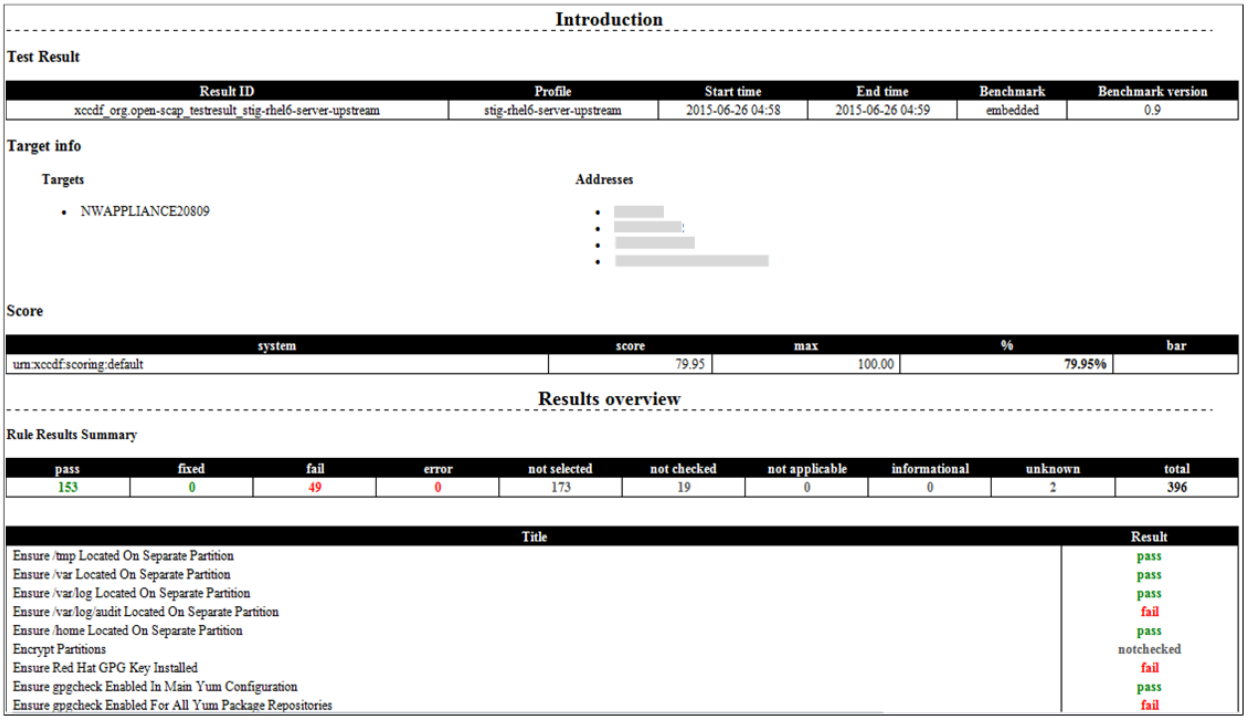
You must

1. SSH to the host
2. Execute the following commands.

```
yum install scap-security-guide
```

Sample Report

The following report is a sample section from an OpenSCAP report.



Introduction

Test Result

Result ID	Profile	Start time	End time	Benchmark	Benchmark version
xccdf_org.open-scap_testresult_stig-rhel6-server-upstream	stig-rhel6-server-upstream	2015-06-26 04:58	2015-06-26 04:59	embedded	0.9

Target info

Targets

- NWAPPLIANCE20809

Addresses

- [redacted]

Score

system	score	max	%	bar
urn:xccdf:scoring:default	79.95	100.00	79.95%	

Results overview

Rule Results Summary

pass	fixed	fail	error	not selected	not checked	not applicable	informational	unknown	total
153	0	49	0	173	19	0	0	2	396

Title	Result
Ensure /tmp Located On Separate Partition	pass
Ensure /var Located On Separate Partition	pass
Ensure /var/log Located On Separate Partition	pass
Ensure /var/log/audit Located On Separate Partition	fail
Ensure /home Located On Separate Partition	pass
Encrypt Partitions	notchecked
Ensure Red Hat GPG Key Installed	fail
Ensure gpgcheck Enabled In Main Yum Configuration	pass
Ensure gpgcheck Enabled For All Yum Package Repositories	fail

Report Fields

Section	Field	Description
Introduction - Test Result	Result ID	The Extensible Configuration Checklist Description Format (XCCDF) identifier of the report results.
	Profile	XCCDF profile under which the report results are categorized.
	Start time	When the report started.
	End time	When the report ended.
	Benchmark	XCCDF benchmark
	Benchmark version	Version number of the benchmark.
Introduction - Score	system	XCCDF scoring method.
	score	Score attained after running the report.
	max	Highest score attainable.
	%	Score attained after running the report as a percentage.
	bar	Not Applicable.
Results overview - Rule Results Summary	pass	Passed rule check.
	fixed	Rule check that failed previously is now fixed.
	fail	Failed rule check.
	error	Could not perform rule check.
	not selected	This check was not applicable to your NetWitness Platform deployment.
	not checked	Rule could not be checked. There are several reasons why a rule cannot be checked. For example, the rule check requires a check engine not supported by the OpenSCAP report.
	not applicable	Rule check does not apply to your NetWitness Platform deployment.
	informational	Rule checks for informational purposes only (no action required for **fail**).
	unknown	Report was able to check the rule. Run steps manually as described in the report to check the rule.
	total	Total number of rules checked.

Section	Field	Description
Exceptions	Title	Name of rule being checked.
	Result	Valid values are **pass**, **fixed**, **fail**, **error**, **not selected**, **not checked**, **not applicable**, **informational**, or unknown. **Note:** Results values are defined the **Results overview - Rule Results Summary**.

Create the OpenSCAP Report

The following tasks show you how to create the OpenSCAP Report in HTML, XML, or both HTML and XML.

Create Report in HTML Only

To create an OpenSCAP report in HTML only:

1. SSH to the host.
2. Submit the following command:

   ```
   mkdir -p /opt/rsa/openscap
   ```

3. Submit the following command for report upgrades only:

   ```
   sed -i -r -e "s/<platform.*//g" /usr/share/xml/scap/ssg/content/ssg-rhel7-xccdf.xml
   ```

4. Submit the following command:

   ```
   oscap xccdf eval --profile "stig" --report /opt/rsa/openscap/`hostname`.html --cpe /usr/share/xml/scap/ssg/content/ssg-rhel7-cpe-dictionary.xml /usr/share/xml/scap/ssg/content/ssg-rhel7-xccdf.xml
   ```

5. Report will be available under following location:

   ```
   /opt/rsa/openscap/
   ```

Create Report in XML Only

To create an OpenSCAP report in xml only:

1. SSH to the host.
2. Submit the following command:

   ```
   mkdir -p /opt/rsa/openscap
   ```

3. Submit the following command for report upgrades only:

   ```
   sed -i -r -e "s/<platform.*//g" /usr/share/xml/scap/ssg/content/ssg-rhel7-xccdf.xml
   ```

4. Submit the following command:

```
oscap xccdf eval --profile "stig" --results
/opt/rsa/openscap/`hostname`.xml --cpe /usr/share/xml/scap/ssg/content/ssg-
rhel7-cpe-dictionary.xml /usr/share/xml/scap/ssg/content/ssg-rhel7-
xccdf.xml
```

5. Report will be available under following location:
   `/opt/rsa/openscap/`

Create Report in Both XML and HTML

To create an OpenSCAP report in both xml and html:

1. SSH to the host.
2. Submit the following command:
   ```
   mkdir -p /opt/rsa/openscap
   ```
3. Submit the following command for report upgrades only:
   ```
   sed -i -r -e "s/<platform.*//g" /usr/share/xml/scap/ssg/content/ssg-rhel7-
   xccdf.xml
   ```
4. Submit the following command:
   ```
   oscap xccdf eval --profile "stig" --results
   /opt/rsa/openscap/`hostname`.xml --report /opt/rsa/openscap/`hostname`.html
   --cpe /usr/share/xml/scap/ssg/content/ssg-rhel7-cpe-dictionary.xml
   /usr/share/xml/scap/ssg/content/ssg-rhel7-xccdf.xml
   ```
5. Report will be available under following location:
   `/opt/rsa/openscap/`

Manage STIG Controls Script (`manage-stig-controls`)

You can use the `manage-stig-controls` script and its arguments to enable or disable STIG Control groups for which you want to apply STIG configuration. You can specify all hosts or individual hosts as arguments and you can enable or disable all control groups or individual control groups. This script is available in /usr/bin/ directory.

To manage STIG controls for a host:

1. SSH to the NW Server host or use the Console from the NetWitness Platform User Interface.
2. Submit the `manage-stig-controls` script with the commands, control groups, and other arguments you want to apply.
3. Reboot the host.

Commands

Command	Description
`--enable-all-controls`	Enables all STIG controls. For example: `manage-stig-controls --enable-all-controls`

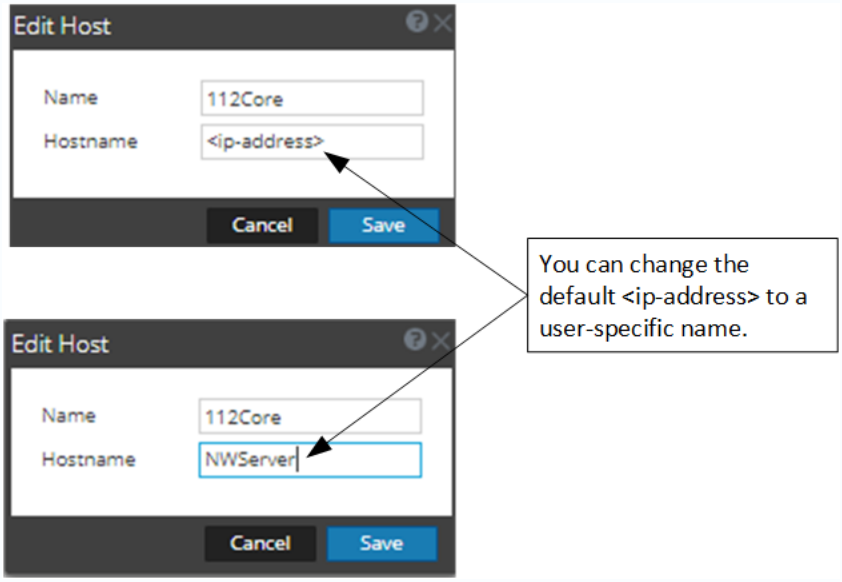
Command	Description
`--disable-all-controls`	Disables all STIG controls. For example: `manage-stig-controls --disable-all-controls`
`--enable-default-controls`	Enables all STIG Controls except ssh-prevent-root and fips-kernel. For example: `manage-stig-controls --enable-default-controls`
`--enable-control-groups <IDs>`	Enables (comma delimited) list of STIG Control GroupIDs. For example: `manage-stig-controls --enable-control-groups '1, 2, 3'`
`--disable-control-groups <IDs>`	Disables (comma delimited) list of STIG Control Group IDs For example: `manage-stig-controls --disable-control-groups '1, 2, 3'`

Control Groups

You use the ID as an argument for the control group or groups.

ID	Group	Description	Specified by Default
1	`ssh-prevent-root`	Prevent root login through SSH.	no
2	`ssh`	SSH STIG configuration.	yes
3	`fips-kernel`	FIPS Kernel configuration	no
4	`auth`	Authentication STIG configuration	yes
5	`audit`	Audit STIG configuration	yes
6	`packages`	RPM Package STIG configuration	yes
7	`services`	Services STIG configuration	yes

Other Arguments

Argument	Description
`--host-all`	Apply STIG configuration to all hosts. For example: `manage-stig-controls --host-all`
`--skip-health-checks`	Disable health checks for all hosts (not recommended). For example: `manage-stig-controls --skip-health-checks`
`--host-id <id>`	Apply STIG configuration for the host identified by `<id>` (host identification code). For example: `manage-stig-controls --host-id <id>`
`--host-name <display-name>`	Apply STIG configuration for host identified by `<display-name>`. `display-name` is the value shown under **Name** in the **(Admin)** > **Hosts** View in the NetWitness Platform Interface. For example: `manage-stig-controls --host-name <display-name>`
`--host-addr <Hostname-in UI>` or `--host-addr <hostname>`	Apply STIG configuration for the host identified by the value shown under **Hostname** in the **(Admin)** > **Hosts** > **Edit** dialog in the NetWitness Platform Interface. This value can be an ip-addres (default) or a user-specified name. For example: `manage-stig-controls --host-addr <hostname>` 
`-v, --verbose`	Enable verbose output. For example: `manage-stig-controls -v`

Rules List

The following table lists all the STIG rules with their:

- Control Group - you can use the Control Group ID as an argument in the manage-stig-controls script to expand on reduce the scope of rules checked. (1= ssh-prevent-root, 2 = ssh, 3 = fips-kernel, 4 = auth, 5 = audit, 6 = packages, 7 = services)
- Default Status - tells you if the rule is enabled or disabled by default.
- Passed or Exception status - tells you if the rule passed (that is, complies with STIG) or is an exception.

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-26404-4	Ensure /var Located On Separate Partition	n/a	n/a	Exception
CCE-26631-2	Set Password Strength Minimum Different Characters	auth	enabled	Passed
CCE-26828-4	Disable DCCP Support	n/a	n/a	Passed
CCE-26884-7	Set Lockout Time For Failed Password Attempts	auth	enabled	Exception
CCE-26892-0	Set the GNOME3 Login Warning Banner Text	n/a	enabled	Passed
CCE-26923-3	Limit Password Reuse	n/a	enabled	Passed
CCE-26952-2	Configure Periodic Execution of AIDE	audit	enabled	Exception
CCE-26970-4	Enable GNOME3 Login Warning Banner	audit	enabled	Passed
CCE-26971-2	Ensure /var/log/audit Located On Separate Partition	audit	enabled	Exception
CCE-26989-4	Ensure gpgcheck Enabled In Main Yum Configuration	n/a	enabled	Passed
CCE-27002-5	Set Password Minimum Age	n/a	enabled	Passed
CCE-27051-2	Set Password Maximum Age	auth	enabled	Passed
CCE-27053-8	Set Password Hashing Algorithm in /etc/libuser.conf	n/a	enabled	Passed
CCE-27082-7	Set SSH Client Alive Count	ssh	disabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-27096-7	Install AIDE	n/a	n/a	Exception
CCE-27127-0	Enable Randomized Layout of Virtual Address Space	n/a	enabled	Exception
CCE-27157-7	Verify File Hashes with RPM	n/a	n/a	Exception
CCE-27160-1	Set Password Retry Prompts Permitted Per-Session	n/a	enabled	Passed
CCE-27200-5	Set Password Strength Minimum Uppercase Characters	auth	enabled	Passed
CCE-27209-6	Verify and Correct File Permissions with RPM	n/a	n/a	Exception
CCE-27213-8	Record Events that Modify the System's Discretionary Access Controls - setxattr	audit	enabled	Passed
CCE-27214-6	Set Password Strength Minimum Digit Characters	auth	enabled	Passed
CCE-27218-7	Remove the X Windows Package Group	n/a	enabled	Passed
CCE-27275-7	Set Last Logon/Access Notification	n/a	enabled	Passed
CCE-27277-3	Disable Modprobe Loading of USB Storage Driver	services	enabled	Exception
CCE-27279-9	Configure SELinux Policy	n/a	enabled	Passed
CCE-27280-7	Record Events that Modify the System's Discretionary Access Controls - lsetxattr	audit	enabled	Passed
CCE-27286-4	Prevent Log In to Accounts With Empty Password	n/a	enabled	Passed
CCE-27287-2	Require Authentication for Single User Mode	n/a	enabled	Passed
CCE-27293-0	Set Password Minimum Length	auth	enabled	Passed
CCE-27295-5	Use Only FIPS 140-2 Validated Ciphers	n/a	enabled	Exception

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-27297-1	Set Interval For Counting Failed Password Attempts	auth	enabled	Passed
CCE-27303-7	Modify the System Login Banner	ssh	enabled	Exception
CCE-27309-4	Set Boot Loader Password in grub2	n/a	enabled	Exception
CCE-27311-0	Verify Permissions on SSH Server Public *.pub Key Files	n/a	enabled	Passed
CCE-27314-4	Enable SSH Warning Banner	ssh	enabled	Passed
CCE-27320-1	Allow Only SSH Protocol 2	n/a	enabled	Passed
CCE-27326-8	Ensure No Device Files are Unlabeled by SELinux	n/a	enabled	Passed
CCE-27334-2	Ensure SELinux State is Enforcing	n/a	enabled	Exception
CCE-27339-1	Record Events that Modify the System's Discretionary Access Controls - chmod	audit	enabled	Passed
CCE-27342-5	Uninstall rsh-server Package	n/a	enabled	Passed
CCE-27343-3	Ensure Logs Sent To Remote Host	n/a	n/a	Passed
CCE-27345-8	Set Password Strength Minimum Lowercase Characters	auth	enabled	Passed
CCE-27349-0	Set Default firewalld Zone for Incoming Packets	n/a	n/a	Exception
CCE-27350-8	Set Deny For Failed Password Attempts	auth	enabled	Passed
CCE-27351-6	Install the screen Package	n/a	enabled	Passed
CCE-27353-2	Record Events that Modify the System's Discretionary Access Controls - fremovexattr	audit	enabled	Passed
CCE-27355-7	Set Account Expiration Following Inactivity	n/a	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-27356-5	Record Events that Modify the System's Discretionary Access Controls - fchown	audit	enabled	Passed
CCE-27358-1	Deactivate Wireless Network Interfaces	n/a	enabled	Passed
CCE-27360-7	Set Password Strength Minimum Special Characters	auth	enabled	Passed
CCE-27363-1	Do Not Allow SSH Environment Options	ssh	enabled	Passed
CCE-27364-9	Record Events that Modify the System's Discretionary Access Controls - chown	audit	enabled	Passed
CCE-27367-2	Record Events that Modify the System's Discretionary Access Controls - removexattr	audit	enabled	Passed
CCE-27375-5	Configure auditd space_left Action on Low Disk Space	audit	enabled	Passed
CCE-27377-1	Disable SSH Support for .rhosts Files	n/a	enabled	Passed
CCE-27386-2	Ensure Default SNMP Password Is Not Used	n/a	n/a	Exception
CCE-27387-0	Record Events that Modify the System's Discretionary Access Controls - fchownat	audit	enabled	Passed
CCE-27388-8	Record Events that Modify the System's Discretionary Access Controls - fchmodat	audit	enabled	Passed
CCE-27389-6	Record Events that Modify the System's Discretionary Access Controls - fsetxattr	audit	enabled	Passed
CCE-27393-8	Record Events that Modify the System's Discretionary Access Controls - fchmod	audit	enabled	Passed
CCE-27394-6	Configure auditd mail_acct Action on Low Disk Space	audit	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-27401-9	Uninstall telnet-server Package	n/a	enabled	Passed
CCE-27399-5	Uninstall ypserv Package	n/a	enabled	Passed
CCE-27407-6	Enable auditd Service	audit	enabled	Passed
CCE-27410-0	Record Events that Modify the System's Discretionary Access Controls - lremovexattr	audit	enabled	Passed
CCE-27413-4	Disable Host-Based Authentication	n/a	enabled	Passed
CCE-27433-2	Set SSH Idle Timeout Interval	ssh	enabled	Passed
CCE-27434-0	Configure Kernel Parameter for Accepting IPv4 Source-Routed Packets for All Interfaces	n/a	enabled	Passed
CCE-27437-3	Ensure auditd Collects Information on the Use of Privileged Commands	audit	enabled	Passed
CCE-27445-6	Disable SSH Root Login	n/a	n/a	Exception
CCE-27447-2	Ensure auditd Collects Information on Exporting to Media (successful)	audit	enabled	Passed
CCE-27455-5	Use Only FIPS 140-2 Validated MACs	n/a	enabled	Passed
CCE-27458-9	Mount Remote Filesystems with Kerberos Security	n/a	enabled	Passed
CCE-27461-3	Ensure auditd Collects System Administrator Actions	audit	enabled	Passed
CCE-27471-2	Disable SSH Access via Empty Passwords	n/a	enabled	Exception
CCE-27485-2	Verify Permissions on SSH Server Private *_key Key Files	n/a	n/a	Passed
CCE-27498-5	Disable the Automounter	n/a	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-27503-2	All GIDs referenced in /etc/passwd must be defined in /etc/group	n/a	enabled	Passed
CCE-27511-5	Disable Ctrl-Alt-Del Reboot Activation	services	enabled	Passed
CCE-27512-3	Set Password Maximum Consecutive Repeating Characters	n/a	enabled	Passed
CCE-27557-8	Set Interactive Session Timeout	auth	disabled	Passed
CCE-80104-3	Disable GDM Automatic Login	n/a	enabled	Passed
CCE-80105-0	Disable GDM Guest Login	n/a	enabled	Passed
CCE-80108-4	Enable the GNOME3 Login Smartcard Authentication	n/a	enabled	Passed
CCE-80110-0	Set GNOME3 Screensaver Inactivity Timeout	n/a	enabled	Passed
CCE-80111-8	Enable GNOME3 Screensaver Idle Activation	n/a	enabled	Passed
CCE-80112-6	Enable GNOME3 Screensaver Lock After Idle Period	n/a	enabled	Passed
CCE-80127-4	Install McAfee Virus Scanning Software	n/a	n/a	Exception
CCE-80129-0	Virus Scanning Software Definitions Are Updated	n/a	n/a	Exception
CCE-80134-0	Ensure All Files Are Owned by a User	n/a	enabled	Passed
CCE-80135-7	Ensure All Files Are Owned by a Group	n/a	enabled	Passed
CCE-80136-5	Ensure All World-Writable Directories Are Owned by a System Account	n/a	enabled	Passed
CCE-80144-9	Ensure /home Located On Separate Partition	n/a	enabled	Passed
CCE-80148-0	Add nosuid Option to Removable Media Partitions	n/a	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80156-3	Disable Kernel Parameter for Sending ICMP Redirects for All Interfaces	n/a	n/a	Exception
CCE-80157-1	Disable Kernel Parameter for IP Forwarding	n/a	n/a	Exception
CCE-80158-9	Configure Kernel Parameter for Accepting ICMP Redirects for All Interfaces	n/a	n/a	Exception
CCE-80162-1	Configure Kernel Parameter for Accepting Source-Routed Packets By Default	n/a	enabled	Passed
CCE-80163-9	Configure Kernel Parameter for Accepting ICMP Redirects By Default	n/a	n/a	Exception
CCE-80165-4	Configure Kernel Parameter to Ignore ICMP Broadcast Echo Requests	n/a	n/a	Exception
CCE-80174-6	Ensure System is Not Acting as a Network Sniffer	n/a	enabled	Passed
CCE-80179-5	Configure Kernel Parameter for Accepting IPv6 Source-Routed Packets for All Interfaces	n/a	n/a	Exception
CCE-80192-8	Ensure rsyslog Does Not Accept Remote Messages Unless Acting As Log Server	n/a	enabled	Passed
CCE-80205-8	Ensure the Default Umask is Set Correctly in login.defs	n/a	enabled	Passed
CCE-80207-4	Enable Smart Card Login	n/a	n/a	Exception
CCE-80213-2	Uninstall tftp-server Package	n/a	enabled	Passed
CCE-80214-0	Ensure tftp Daemon Uses Secure Mode	n/a	enabled	Passed
CCE-80215-7	Install the OpenSSH Server Package	n/a	enabled	Passed
CCE-80216-5	Enable the OpenSSH Service	n/a	enabled	Passed
CCE-80220-7	Disable GSSAPI Authentication	ssh	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80221-5	Disable Kerberos Authentication	n/a	enabled	Passed
CCE-80222-3	Enable Use of Strict Mode Checking	n/a	enabled	Passed
CCE-80223-1	Enable Use of Privilege Separation	n/a	enabled	Passed
CCE-80224-9	Disable Compression Or Set Compression to delayed	n/a	enabled	Passed
CCE-80225-6	Print Last Log	n/a	enabled	Exception
CCE-80226-4	Enable Encrypted X11 Forwarding	n/a	n/a	Exception
CCE-80240-5	Mount Remote Filesystems with nosuid	n/a	enabled	Passed
CCE-80245-4	Uninstall vsftpd Package	n/a	enabled	Passed
CCE-80258-7	Disable KDump Kernel Crash Analyzer (kdump)	services	enabled	Passed
CCE-80346-0	Ensure YUM Removes Previous Package Versions	packages	enabled	Passed
CCE-80347-8	Ensure gpgcheck Enabled for Local Packages	packages	enabled	Passed
CCE-80348-6	Ensure gpgcheck Enabled for Repository Metadata	n/a	n/a	Exception
CCE-80350-2	Ensure Users Re-Authenticate for Privilege Escalation - sudo !authenticate	n/a	enabled	Passed
CCE-80351-0	Ensure Users Re-Authenticate for Privilege Escalation - sudo NOPASSWD	n/a	enabled	Passed
CCE-80352-8	Ensure the Logon Failure Delay is Set Correctly in login.defs	auth	enabled	Passed
CCE-80353-6	Configure the root Account for Failed Password Attempts	auth	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80354-4	Set the UEFI Boot Loader Password	fips-kernel	disabled	Passed
CCE-80359-3	Enable FIPS Mode in GRUB2	fips-kernel	disabled	Exception
CCE-80370-0	Set GNOME3 Screensaver Lock Delay After Activation Period	n/a	enabled	Passed
CCE-80371-8	Ensure Users Cannot Change GNOME3 Screensaver Settings	n/a	enabled	Passed
CCE-80372-6	Disable SSH Support for User Known Hosts	ssh	enabled	Passed
CCE-80373-4	Disable SSH Support for Rhosts RSA Authentication	audit	enabled	Passed
CCE-80374-2	Configure Notification of Post-AIDE Scan Details	n/a	n/a	Exception
CCE-80375-9	Configure AIDE to Verify Access Control Lists (ACLs)	n/a	n/a	Exception
CCE-80376-7	Configure AIDE to Verify Extended Attributes	n/a	n/a	Exception
CCE-80377-5	Configure AIDE to Use FIPS 140-2 for Validating Hashes	n/a	n/a	Exception
CCE-80378-3	Verify User Who Owns /etc/cron.allow file	n/a	enabled	Passed
CCE-80379-1	Verify Group Who Owns /etc/cron.allow file	n/a	enabled	Passed
CCE-80380-9	Ensure cron Is Logging To Rsyslog	n/a	enabled	Passed
CCE-80381-7	Shutdown System When Auditing Failures Occur	audit	enabled	Passed
CCE-80382-5	Record Attempts to Alter Logon and Logout Events - tallylog	audit	enabled	Passed
CCE-80383-3	Record Attempts to Alter Logon and Logout Events - faillock	n/a	n/a	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80384-1	Record Attempts to Alter Logon and Logout Events - lastlog	audit	enabled	Passed
CCE-80385-8	Record Unauthorized Access Attempts to Files (unsuccessful) - creat	audit	enabled	Passed
CCE-80386-6	Record Unauthorized Access Attempts to Files (unsuccessful) - open	audit	enabled	Passed
CCE-80387-4	Record Unauthorized Access Attempts to Files (unsuccessful) - openat	audit	enabled	Passed
CCE-80388-2	Record Unauthorized Access Attempts to Files (unsuccessful) - open_by_ handle_at	audit	enabled	Passed
CCE-80389-0	Record Unauthorized Access Attempts to Files (unsuccessful) - truncate	audit	enabled	Passed
CCE-80390-8	Record Unauthorized Access Attempts to Files (unsuccessful) - ftruncate	audit	enabled	Passed
CCE-80391-6	Record Any Attempts to Run semanage	audit	enabled	Passed
CCE-80392-4	Record Any Attempts to Run setsebool	audit	enabled	Passed
CCE-80393-2	Record Any Attempts to Run chcon	audit	enabled	Passed
CCE-80395-7	Ensure auditd Collects Information on the Use of Privileged Commands - passwd	audit	enabled	Passed
CCE-80396-5	Ensure auditd Collects Information on the Use of Privileged Commands - unix_chkpwd	audit	enabled	Passed
CCE-80397-3	Ensure auditd Collects Information on the Use of Privileged Commands - gpasswd	audit	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80398-1	Ensure auditd Collects Information on the Use of Privileged Commands - chage	audit	enabled	Passed
CCE-80399-9	Ensure auditd Collects Information on the Use of Privileged Commands - userhelper	audit	enabled	Passed
CCE-80400-5	Ensure auditd Collects Information on the Use of Privileged Commands - su	audit	enabled	Passed
CCE-80401-3	Ensure auditd Collects Information on the Use of Privileged Commands - sudo	audit	enabled	Passed
CCE-80402-1	Ensure auditd Collects Information on the Use of Privileged Commands - sudoedit	audit	enabled	Passed
CCE-80403-9	Ensure auditd Collects Information on the Use of Privileged Commands - newgrp	audit	enabled	Passed
CCE-80404-7	Ensure auditd Collects Information on the Use of Privileged Commands - chsh	audit	enabled	Passed
CCE-80405-4	Ensure auditd Collects Information on the Use of Privileged Commands - umount	audit	enabled	Passed
CCE-80406-2	Ensure auditd Collects Information on the Use of Privileged Commands - postdrop	audit	enabled	Passed
CCE-80407-0	Ensure auditd Collects Information on the Use of Privileged Commands - postqueue	audit	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80408-8	Ensure auditd Collects Information on the Use of Privileged Commands - ssh-keysign	audit	enabled	Passed
CCE-80410-4	Ensure auditd Collects Information on the Use of Privileged Commands - crontab	audit	enabled	Passed
CCE-80411-2	Ensure auditd Collects Information on the Use of Privileged Commands - pam_timestamp_check	audit	enabled	Passed
CCE-80412-0	Ensure auditd Collects File Deletion Events by User - rmdir	audit	enabled	Passed
CCE-80413-8	Ensure auditd Collects File Deletion Events by User - renameat	audit	enabled	Passed
CCE-80414-6	Ensure auditd Collects Information on Kernel Module Loading - init_ module	audit	enabled	Passed
CCE-80415-3	Ensure auditd Collects Information on Kernel Module Unloading - delete_ module	audit	enabled	Passed
CCE-80430-2	Record Events that Modify User/Group Information - /etc/security/opasswd	audit	enabled	Passed
CCE-80431-0	Record Events that Modify User/Group Information - /etc/shadow	audit	enabled	Passed
CCE-80432-8	Record Events that Modify User/Group Information - /etc/gshadow	audit	enabled	Passed
CCE-80433-6	Record Events that Modify User/Group Information - /etc/group	audit	enabled	Passed
CCE-80434-4	Ensure Home Directories are Created for New Users	n/a	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80435-1	Record Events that Modify User/Group Information - /etc/passwd	audit	enabled	Passed
CCE-80436-9	Mount Remote Filesystems with noexec	n/a	enabled	Passed
CCE-80437-7	Configure PAM in SSSD Services	n/a	n/a	Exception
CCE-80438-5	Configure Multiple DNS Servers in /etc/resolv.conf	n/a	n/a	Exception
CCE-80439-3	Configure Time Service Maxpoll Interval	services	enabled	Passed
CCE-80447-6	Configure the Firewalld Ports	n/a	n/a	Exception
CCE-80513-5	Remove Host-Based Authentication Files	n/a	enabled	Passed
CCE-80514-3	Remove User Host-Based Authentication Files	n/a	enabled	Passed
CCE-80515-0	Configure SSSD LDAP Backend Client CA Certificate Location	n/a	n/a	Exception
CCE-80519-2	Install Smart Card Packages For Multifactor Authentication	n/a	n/a	Exception
CCE-80537-4	Configure auditd space_left on Low Disk Space	audit	enabled	Passed
CCE-80544-0	Ensure Users Cannot Change GNOME3 Session Idle Settings	n/a	enabled	Passed
CCE-80545-7	Verify and Correct Ownership with RPM	n/a	n/a	Exception
CCE-80546-5	Configure SSSD LDAP Backend to Use TLS For All Transactions	n/a	n/a	Exception
CCE-80547-3	Ensure auditd Collects Information on Kernel Module Loading and Unloading - finit_module	audit	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-80563-0	Ensure Users Cannot Change GNOME3 Screensaver Lock After Idle Period	n/a	enabled	Passed
CCE-80564-8	Ensure Users Cannot Change GNOME3 Screensaver Idle Activation	n/a	enabled	Passed
CCE-80660-4	Record Any Attempts to Run setfiles	audit	enabled	Exception
CCE-80661-2	Ensure auditd Collects Information on Kernel Module Loading - create_ module	audit	enabled	Exception
CCE-80995-4	Ensure auditd Collects File Deletion Events by User - rename	audit	enabled	Exception
CCE-80996-2	Ensure auditd Collects File Deletion Events by User - unlinkat	audit	enabled	Exception
CCE-80998-8	Verify firewalld Enabled	n/a	n/a	Exception
CCE-81106-7	Ensure auditd Collects File Deletion Events by User - unlink	audit	enabled	Passed
CCE-81153-9	Add nosuid Option to /home	n/a	enabled	Passed
CCE-82020-9	Set Password Strength Minimum Different Characters	auth	enabled	Passed
CCE-82030-8	Limit Password Reuse	n/a	enabled	Passed
CCE-82035-7	Ensure /var/log/audit Located On Separate Partition	audit	enabled	Exception
CCE-82036-5	Set Password Minimum Age	n/a	enabled	Passed
CCE-82038-1	Set Password Hashing Algorithm in /etc/libuser.conf	n/a	enabled	Passed
CCE-82041-5	Limit the Number of Concurrent Login Sessions Allowed Per User	auth	enabled	Passed
CCE-82043-1	Set PAM's Password Hashing Algorithm	n/a	enabled	Passed

CCE Number	Rule Name	Control Group	Default Status	Passed/ Exception
CCE-82045-6	Set Password Strength Minimum Different Categories	audit	enabled	Passed
CCE-82050-6	Set Password Hashing Algorithm in /etc/login.defs	n/a	enabled	Passed
CCE-82053-0	Ensure /tmp Located On Separate Partition	n/a	n/a	Exception
CCE-82054-8	Verify Only Root Has UID 0	n/a	enabled	Passed
CCE-82353-4	Ensure /var Located On Separate Partition	n/a	n/a	Exception

Exceptions to STIG Compliance

This topic contains:

- Rule exceptions that are the responsibility of the customer to resolve.
- Rule exceptions that are "Not a Finding" which means that they do not apply to NetWitness Platform. NetWitness has verified that the system meets these requirements.
- Rules to be supported in future release.

Key to Elements in Exception Descriptions

CCE Number

The Common Configuration Enumeration (CCE), assigns unique entries (also called CCE numbers) to configuration guidance statements and configuration controls to improve workflow by facilitating fast and accurate correlation of configuration issues present in disparate domains. In this way, it is similar to other comparable data standards such as the **Common Vulnerability and Exposure (CVE®) List** (http://cve.mitre.org/cve), which assigns identifiers to publicly known system vulnerabilities. The OpenSCAP report lists exceptions by CCE number.

This sections lists the exceptions you can receive when you run the OpenSCAP report. The ID or Common Configuration Enumeration (CCE) number in the table is the identification number for the exception from the OpenSCAP report.

Control Group ID

Number that identifies the control group you specify in the manage-stig-controls script to enable or disable the rule.

ID	Group	Description	Specified by Default
1	`ssh-prevent-root`	Prevent root login through SSH.	no
2	`ssh`	SSH STIG configuration.	yes
3	`fips-kernel`	FIPS Kernel configuration	no
4	`auth`	Authentication STIG configuration	yes
5	`audit`	Audit STIG configuration	yes
6	`packages`	RPM Package STIG configuration	yes
7	`services`	Services STIG configuration	yes

Check

Describes what the rule checks to identify exceptions to DISA STIG compliance.

Comments

Provides insight on why you would receive this exception. This section includes one of the following comments that describes the exception:

- **Customer Responsibility** - You are responsible to make sure the system meets this requirement.
- **Not a Finding** - Exception does not apply to NetWitness Platform. NetWitness has verified that the system meets this requirement.
- **Future Feature** - NetWitness Platform does not meet this requirement. NetWitness plans to fix this in a future release of NetWitness Platform.

Customer Responsibility Exceptions

CCE-26952-2 Configure Periodic Execution of AIDE (Control Group = audit)

Check	At a minimum, configure AIDE to run a weekly scan and at most, daily. To implement a daily execution of AIDE at 4:05am using cron, add the following line to the `/etc/crontab` file:
	`05 4 * * * root /usr/sbin/aide --check`
	To implement a weekly execution of AIDE at 4:05am using cron, add the following line to the `/etc/crontab` file:
	`05 4 * * 0 root /usr/sbin/aide --check`
	AIDE can be executed periodically through other means; this is merely one example. The usage of cron's special time codes, such as `@daily` and `@weekly` is acceptable.

Comments	**Customer Responsibility.** NetWitness Platform does not provide AIDE because it has a negative impact on performance. If you must install it, run as infrequently as possible to adhere to your security policy.

CCE-27096-7 Install AIDE (Control Group = n/a)

Check	Install the AIDE package with the following command: `$ sudo yum install aid`
Comments	**Customer Responsibility.** NetWitness Platform does not provide AIDE because it has a negative impact on performance. If you must install it, run as infrequently possible to adhere to your security policy.

CCE-27218-7 Remove the X Windows Package Group

Check	The Rule CCE-27218-7 "Remove the X Windows Package Group" is an exception for Log Collector and Log Decoder services.
Comments	**Customer Responsibility.**Log Collector plugin collection framework uses SELinux sandbox technology that has a direct dependency on the given rpm. Removing of the rpm will lead to loss of plugin collection functionality in Log Collector service.

CCE-27295-5 Use Only FIPS 140-2 Validated Ciphers (Control Group = n/a)

Check	Limit the ciphers to those algorithms which are FIPS-approved. Counter (CTR) mode is also preferred over cipher-block chaining (CBC) mode. The following line in `/etc/ssh/sshd_config` demonstrates use of FIPS 140-2 validated ciphers:
	`Ciphers aes128-ctr,aes192-ctr,aes256-ctr`
	The following ciphers are FIPS 140-2 certified on RHEL 7:
	`- aes128-ctr - aes192-ctr - aes256-ctr - aes128-cbc - aes192-cbc - aes256-cbc - 3des-cbc - rijndael-cbc@lysator.liu.se`
	Any combination of the above ciphers will pass this check. Official FIPS 140-2 paperwork for RHEL7 can be found at http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140sp/140sp2630.pdf.

Comments	Testing approach results in failed test, but the ciphers defined meet the STIG Rule definition. You can check cipher under this file: `/etc/ssh/sshd_config`

CCE-27334-2 Ensure SELinux State is Enforcing

Check	Ensure SELinux State is Enforcing.
Comments	SELinux state is default it is set to 'permissive' by default for all the NetWitness Platform hosts instead of 'Enforcing' due to performance impact.

CCE-27445-6 Disable SSH Root Login (Control Group = ssh-prevent-root)

Check	The root user should never be allowed to login to a system directly over a network.
Comments	**Customer Responsibility.** Disable root login through SSH by adding or editing the following line in the `/etc/ssh/sshd_config` file: `PermitRootLoginNetWitness`.

CCE-80127-4 Install McAfee Virus Scanning Software (Control Group = n/a)

Check	Install McAfee VirusScan Enterprise for Linux antivirus software which is provided for DoD systems and uses signatures to search for the presence of viruses on the filesystem.
Comments	**Customer Responsibility.** Install virus scanning software. NetWitness does not provide this software.

CCE-80129-0 Virus Scanning Software Definitions Are Updated (Control Group = n/a)

Check	Make sure that virus definition files are no older than 7 days or their last release.
Comments	**Customer Responsibility.** NetWitness does not provide this software.

CCE-80207-4 Enable Smart Card Login (Control Group = n/a)

Check	For guidance on enabling SSH to authenticate against a Common Access Card (CAC), consult documentation at: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/System-Level_Authentication_Guide/smartcards.html#authconfig-smartcards https://access.redhat.com/solutions/82273
Comments	**Customer Responsibility.** The NetWitness Platform supports username/certificate for authentication to shell. If you want to configure a smart card log in, you must do this outside of NetWitness NetWitness.

CCE-80359-3 Enable FIPS Mode in GRUB2 (Control Group = fips-kernel)

Check	To ensure FIPS mode is enabled, install the `dracut-fips` package and rebuild `initramfs` by running the following commands:
	`$ sudo yum install dracut-fips dracut -f`
	After the `dracut` command has been run, add the `fips=1` argument to the default GRUB 2 command line for the Linux operating system in the `/etc/default/grub` file as shown in the following example:
	`GRUB_CMDLINE_LINUX='crashkernel=auto rd.lvm.lv=VolGroup/LogVol06 rd.lvm.lv=VolGroup/lv_swap rhgb quiet rd.shell=0 fips=1'`
	Finally, rebuild the `grub.cfg` file by using the `grub2-mkconfig -o` command as follows (On BIOS-based machines, issue the following command as root):
	`~]# grub2-mkconfig -o /boot/grub2/grub.cfg`
	On UEFI-based machines, issue the following command as root:
	`~]# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg`

Comments	**Customer Responsibility.** NetWitness Platform does not enabled by default. You can enable FIPS by following the procedures in the Configure FIPS Support.

CCE-80374-2 Configure Notification of Post-AIDE Scan Details (Control Group = n/a)

Check	AIDE should notify appropriate personnel of the details of a scan after the scan has been run. If AIDE has already been configured for periodic execution in the `/etc/crontab` file, append the following line to the existing AIDE line:
	`| /bin/mail -s '$(hostname) - AIDE Integrity Check' root@localhost`
	Otherwise, add the following line to the `/etc/crontab` file:
	`05 4 * * * root /usr/sbin/aide --check | /bin/mail -s '$(hostname) - AIDE Integrity Check' root@localhost`
	AIDE can be executed periodically through other means. This is just one example.

Comments	**Customer Responsibility.** NetWitness Platform does not provide AIDE because it has a negative impact on performance. If you must install it, run as infrequently possible to adhere to your security policy.

CCE-80375-9 Configure AIDE to Verify Access Control Lists (Control Group = n/a)

Check	By default, the `acl` option is added to the FIPSR ruleset in AIDE. If using a custom ruleset or the `acl` option is missing, add `acl` to the appropriate ruleset. For example, add `acl` to the following line in the `/etc/aide.conf` file: `FIPSR = p+i+n+u+g+s+m+c+acl+selinux+xattrs+sha256` AIDE rules can be configured in multiple ways; this is merely one example that is already configured by default.
Comments	**Customer Responsibility.** NetWitness Platform does not provide AIDE because it has a negative impact on performance. If you must install it, run as infrequently possible to adhere to your security policy.

CCE-80376-7 Configure AIDE to Verify Extended Attributes (Control Group = n/a)

Check	By default, the `xattrs` option is added to the FIPSR ruleset in AIDE. If using a custom ruleset or the `xattrs` option is missing, add `xattrs` to the appropriate ruleset. For example, add xattrs to the following line in the `/etc/aide.conf` file: `FIPSR = p+i+n+u+g+s+m+c+acl+selinux+xattrs+sha256` AIDE rules can be configured in multiple ways. This is just one example that is already configured by default.
Comments	**Customer Responsibility.** NetWitness Platform does not provide AIDE because it has a negative impact on performance. If you must install it, run as infrequently possible to adhere to your security policy.

CCE-80377-5 Configure AIDE to Use FIPS 140-2 for Validating Hashes (Control Group = n/a)

Check	By default, the `sha512` option is added to the `ORMAL` ruleset in AIDE. If using a custom `ruleset` or the `sha512` option is missing, add `sha512` to the appropriate ruleset. For example, add `sha512` to the following line in the `/etc/aide.conf` file:`ORMAL = FIPSR+sha512` AIDE rules can be configured in multiple ways; this is merely one example that is already configured by default.
Comments	**Customer Responsibility.** NetWitness Platform does not provide AIDE because it has a negative impact on performance. If you must install it, run as infrequently possible to adhere to your security policy.

CCE-80519-2 Install Smart Card Packages For Multi-Factor Authentication (Control Group = n/a)

Check	Configure the operating system to implement multifactor authentication by installing the required packages with the following command: `$ sudo yum install esc pam_pkcs11 authconfig-gtk`

Comments	**Customer Responsibility.** The NetWitness Platform supports username/certificate for authentication to shell. If you want to configure a smart card log in, you must do this outside of NetWitness NetWitness.

Exceptions That Are Not a Finding

The following exceptions do not apply to NetWitness Platform. NetWitness has verified that the system meets these requirements.

CCE-26404-4 Ensure /var Located On Separate Partition (Control Group = n/a)

Check	The `/var` directory is used by daemons and other system services to store frequently-changing data. Ensure that `/var` has its own partition or logical volume at installation time, or migrate it using LVM.
Comments	**Not a Finding.**NetWitness software is installed in `/var/netwitness` by default and has a separate partition on `/var/netwitness`.

CCE-26828-4 Set GNOME Login Inactivity timeout (Control Group = n/a)

Check	Verify that the GNOME Login Inactivity Timeout is set on the host (The graphical desktop environment must set the idle timeout to no more than 15 minutes.).
Comments	**Not a Finding.** NetWitness Platform does not use Gnome Graphical User Interface (GUI) Desktop.

CCE-26884-7 Set Lockout Time For Failed Password Attempts (Control Group = auth)

Check	To configure the system to lock out accounts after a number of incorrect login attempts and require an administrator to unlock the account using `pam_faillock.so`, modify the content of both `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` by adding the following line immediately before the `pam_unix.so` statement in the `AUTH` section:
	`auth required pam_faillock.so preauth silent deny= unlock_time= fail_interval=`
	Add the following line immediately after the `pam_unix.so` statement in the `AUTH` section:
	`auth [default=die] pam_faillock.so authfail deny= unlock_time= fail_interval=`
	Add the following line immediately before the `pam_unix.so` statement in the `ACCOUNT` section:
	`account required pam_faillock.s`

Comments	**Not a Finding.** `root_unlock_time` is set to `600` seconds.

CCE-26971-2 Ensure `/var/log/audit` Located On Separate Partition (Control Group = audit)

Check	Audit logs are stored in the `/var/log/audit` directory. Ensure that it has its own partition or logical volume at installation time, or migrate it later using LVM. Make absolutely certain that it is large enough to store all audit logs that will be created by the auditing daemon.
Comments	**Not a Finding.**NetWitness Platform has the `/var/log` directory as a separate partition.

CCE-27127-0 Enable Randomized Layout of Virtual Address Space (Control Group = n/a)

Check	To set the runtime status of the `kernel.randomize_va_space` kernel parameter, run the following command:
	`$ sudo sysctl -w kernel.randomize_va_space=2`
	If this is not the system default value, add the following line to the `/etc/sysctl.conf` file:
	`kernel.randomize_va_space = 2`

Comments	**Not a Finding.** Value of `/proc/sys/kernel/randomize_va_space` is already `2`.

CCE-27157-7 Verify File Hashes with RPM (Control Group = n/a)

Check	Without cryptographic integrity protections, system executables and files can be altered by unauthorized users without detection. The RPM package management system can check the hashes of installed software packages, including many that are important to system security. To verify that the cryptographic hash of system files and commands match vendor values, run the following command to list which files on the system with hashes that differ from what is expected by the RPM database:
	`$ rpm -Va | grep '^..5' A 'c'`
	in the second column indicates that a file is a configuration file, which may appropriately be expected to change. If the file was not expected to change, investigate the cause of the change using audit logs or other means. The package can then be reinstalled to restore the file. Run the following command to determine which package owns the file:
	`$ rpm -qf`
	The package can be reinstalled from a yum repository using the command:
	`FILENAME $ sudo yum reinstall`
	Alternatively, the package can be reinstalled from trusted media using the command:
	`PACKAGENAME $ sudo rpm -Uvh PACKAGENAME`

Comments	**Not a Finding.** Only mismatched files not marked as config files in rpms are Commercial Off the Shelf (COTS) product based that cannot be updated. Most File Hash/RPM combinations are in sync. Any discrepancies are COTS products that cannot be updated.

CCE-27339-1 Record Events that Modify the System's Discretionary Access Controls - chmod

Check	Verify that the host records events that modify the system's discretionary access controls - `chown`.

Comments	**Not a Finding.** Make sure that you have the correct `chown` configuration on the host. The following settings are the correct configuration.
	`[root@localhost nwadmin]# grep chown /etc/audit/*`
	`/etc/audit/audit.rules:-a exit,always -F arch=b64 -S chown -S chmod -S fchmod -S fchmodat -S fchown -S fchownat -S fremovexattr -S fsetxattr -S lchown -S lremovexattr -S lremovexattr -S lsetxattr -S removexattr -S setxattr -F auid>=500 -F auid!=4294967295 -k perm_mod`
	`/etc/audit/audit.rules:-a exit,always -F arch=b32 -S chown -S fchmod -S fchmodat -S fchown -S fchownat -S fremovexattr -S fsetxattr -S lchown -S lremovexattr -S lremovexattr -S lsetxattr -S removexattr -S setxattr -F auid>=500 -F auid!=4294967295 -k perm_mod`

CCE-27209-6 Verify and Correct File Permissions with RPM (Control Group = n/a)

Rule Name	

Check	The RPM package management system can check file access permissions of installed software packages, including many that are important to system security. Verify that the file permissions of system files and commands match vendor values. Check the file permissions with the following command:
	`$ sudo rpm -Va | grep '^.M'`
	Output indicates files that do not match vendor defaults. After locating a file with incorrect permissions, run the following command to determine which package owns it:
	`$ rpm -qf FILENAME`
	Next, run the following command to reset its permissions to the correct values:
	`$ sudo rpm --quiet --setperms PACKAGENAME`

Comments	**Not a Finding.** The file permissions do not match the rpm, they are configured to be stricter during configuration management.

CCE-27303-7 (Control ID = 2) Modify the System Login Banner (Control Group = ssh)

Check	To configure the system login banner edit the `/etc/issue` file. Replace the default text with a message compliant with the local site policy or a legal disclaimer. The DoD required text is either: " You are accessing a U.S. Government (USG) Information System (IS) that is provided for USG-authorized use only. By using this IS (which includes any device attached to this IS), you consent to the following conditions: • The USG routinely intercepts and monitors communications on this IS for purposes including, but not limited to, penetration testing, COMSEC monitoring, network operations and defense, personnel misconduct (PM), law enforcement (LE), and counterintelligence (CI) investigations. • At any time, the USG may inspect and seize data stored on this IS. • Communications using, or data stored on, this IS are not private, are subject to routine monitoring, interception, and search, and may be disclosed or used for any USG-authorized purpose. • This IS includes security measures (e.g., authentication and access controls) to protect USG interests -- not for your personal benefit or privacy. -Notwithstanding the above, using this IS does not constitute consent to PM, LE or CI investigative searching or monitoring of the content of privileged communications, or work product, related to personal representation or services by attorneys, psychotherapists, or clergy, and their assistants. Such communications and work product are private and confidential. See User Agreement for details." or " I've read & consent to terms in IS user agreem't."
Comments	**Not a Finding.** The login banner is displayed but does not hyphenate "agreem't"

CCE-27311-0 Very Permissions on SHH Server *.pub Key Files (Control Group = na)

Check	
Comments	**Not a Finding.** All public keys are set to with permissions 640 in the `/etc/ssh/` directory.

CCE-27314-4 Enable SSH Warning Banner (Control Group = na)

Check	
Comments	**Not a Finding.** The required configuration exists in the `etc/ssh/sshd_conf` file.

CCE-27349-0 Set Default `firewalld` Zone for Incoming Packets (Control Group = n/a)

Check	To set the default zone to drop for the built-in default zone which processes incoming IPv4 and IPv6 packets, modify the following line in the `/etc/firewalld/firewalld.conf` file to be: `DefaultZone=drop`
Comments	**Not a Finding.** NetWitness Platform `firewalld`service is disabled because it uses IP Tables, not FirewallD.

CCE-27386-2 Ensure Default SNMP Password Is Not Used (Control Group = n/a)

Check	Edit `/etc/snmp/snmpd.conf` file by removing or changing the default community strings of public and private. After the default community strings have been changed, restart the SNMP service: `$ sudo service snmpd restart`
Comments	**Not a Finding.** NetWitness Platform does not use `snmp`, and the `snmpd` service not enabled.

CCE-27455-5 Use Only FIPS 140-2 Validated MACs (Control Group = na)

Check	
Comments	**Not a Finding.** The following configuration exists in `/etc/ssh/sshd_config` file: `MACs hmac-sha1,hmac-sha2-256,hmac-sha2-512`

CCE-27471-2 Disable SSH Access via Empty Passwords (Control Group = n/a)

Check	Explicitly disallow SSH login from accounts with empty passwords, add or correct the following line in the `/etc/ssh/sshd_config` file.
Comments	**Not a Finding.** NetWitness Platform sets the **permitemptypasswords** parameter to **no** by default. This should pass the DISA STIG rule check.

CCE-27485-2 Very Permissions on SHH Server Private *.key Key Files (Control Group = na)

Check	
Comments	**Not a Finding.** All private keys are set to with permissions 640 in the `/etc/ssh/` directory.

CCE-80156-3 Disable Kernel Parameter for Sending ICMP Redirects for All Interfaces (Control Group = n/a)

Check	To set the runtime status of the `t.ipv4.conf.all.send_redirects` kernel parameter, run the following command:
	`$ sudo sysctl -w net.ipv4.conf.all.send_redirects=0`
	If this is not the system default value, add the following line to the `/etc/sysctl.conf` file:
	`t.ipv4.conf.all.send_redirects = 0`

Comments	**Not a Finding.** NetWitness Platform does not accept incoming Internet Control Message Protocol (ICMP) traffic.

CCE-80157-1 Disable Kernel Parameter for IP Forwarding (Control Group = n/a)

Check	To set the runtime status of the `t.ipv4.ip_forward` kernel parameter, run the following command:
	`$ sudo sysctl -w net.ipv4.ip_forward=0`
	If this is not the system default value, add the following line to the `/etc/sysctl.conf` file:
	`t.ipv4.ip_forward = 0`

Comments	**Not a Finding.** NetWitness Platform only uses FIPS certified MACs (for example, `MACs hmac-sha1,hmac-sha2-256,hmac-sha2-512`).

CCE-80158-9 Configure Kernel Parameter for Accepting ICMP Redirects for All Interfaces (Control Group = n/a)

Check	To set the runtime status of the `t.ipv4.conf.all.accept_redirects` kernel parameter, run the following command:
	`$ sudo sysctl -w net.ipv4.conf.all.accept_redirects=0`
	If this is not the system default value, add the following line to the `/etc/sysctl.conf` file:
	`t.ipv4.conf.all.accept_redirects = 0`

Comments	**Not a Finding** NetWitness Platform does not accept incoming Internet Control Message Protocol (ICMP) traffic.

CCE-80163-9 Configure Kernel Parameter for Accepting ICMP Redirects By Default (Control Group = n/a)

Check	To set the runtime status of the `t.ipv4.conf.default.accept_redirects` kernel parameter, run the following command:
	`$ sudo sysctl -w net.ipv4.conf.default.accept_redirects=0`
	If this is not the system default value, add the following line to the `/etc/sysctl.conf` file:
	`t.ipv4.conf.default.accept_redirects = 0`

Comments	**Not a Finding** NetWitness Platform does not accept incoming Internet Control Message Protocol (ICMP) traffic.

CCE-80165-4 Configure Kernel Parameter to Ignore ICMP Broadcast Echo Requests (Control Group = n/a)

Rule Name	

Check	To set the runtime status of the `t.ipv4.icmp_echo_ignore_broadcasts` kernel parameter, run the following command:
	`$ sudo sysctl -w net.ipv4.icmp_echo_ignore_broadcasts=1`
	If this is not the system default value, add the following line to the `/etc/sysctl.conf` file:
	`t.ipv4.icmp_echo_ignore_broadcasts = 1`

Comments	**Not a Finding**NetWitness Platform does not accept incoming Internet Control Message Protocol (ICMP) traffic.

CCE-80225-6 Print Last Log (Control Group = n/a)

Check	When enabled, SSH will display the date and time of the last successful account log in. To enable `LastLog` in SSH, add or correct the following line in the `/etc/ssh/sshd_config` file: `PrintLastLog yes`
Comments	**Not a Finding.** NetWitness Platform sets `printlastlog` to `yes` by default.

CCE-80226-4 Enable Encrypted X11 Forwarding (Control Group = n/a)

Check	Enable Encrypted X11 Forwarding - By default, remote X11 connections are not encrypted when initiated by users. SSH has the capability to encrypt remote X11 connections when SSH's `X11Forwarding` option is enabled. To enable X11 Forwarding, add or correct the following line in the `/etc/ssh/sshd_config` file: `X11Forwarding yes`
Comments	**Not a Finding.**NetWitness Platform does not have X11 installed or running.

CCE-80348-6 Ensure gpgcheck Enabled for Repository Metadata (Control Group = n/a)

Check	Verify the operating system prevents the installation of patches, service packs, device drivers, or operating system components of local packages without verification of the repository metadata. Check that yum verifies the repository metadata prior to install with the following command. This should be configured by `setting repo_ gpgcheck` to `1` in `/etc/yum.conf`.
Comments	**Not a Finding.** .NetWitness Platform rpm signing procedures do not support signing the repo metadata

CCE-80383-3 Record Attempts to ALter Logon Events - failock (Control Group = na)

Check	
Comments	**Not a Finding.** The required rules are configured in the `/etc/audit/rules.d/nw-stig.rules` file.

CCE-80399-9 Ensure auditd Collects Information on the Use of Privileged Commands - userhelper (Control Group = na)

Check	
Comments	**Not a Finding.** The required rules are configured in the `/etc/audit/rules.d/nw-stig.rules` file.

CCE-80437-7 Configure PAM in SSSD Services (Control Group = n/a)

Check	SSSD should be configured to run SSSD pam services. To configure SSSD to known SSH hosts, add pam to services under the `[sssd]` section in `/etc/sssd/sssd.conf` file. For example: `[sssd] services = sudo, autofs, pam`
Comments	**Not a Finding.** NetWitness Platform does not currently support Multi-Factor authentication. As a result, SSSD service is not installed on a NetWitness Host.

CCE-80438-5 Configure Multiple DNS Servers in `/etc/resolv.conf` (Control Group = n/a)

Check	Multiple Domain Name System (DNS) Servers should be configured in the `/etc/resolv.conf` file. This provides redundant name resolution services in the event that a domain server crashes. To configure the system to contain as least 2 DNS servers, add a corresponding `nameserver` entry in `ip_address /etc/resolv.conf` file for each DNS server where `ip_address` is the IP address of a valid DNS server. For example: `search example.com nameserver 192.168.0.1 nameserver 192.168.0.2`
Comments	**Not a Finding.** NetWitness Platform orchestrates and configures an internal DNS server that all NetWitness hosts use for name resolution. You can configure external DNS servers, but it is dependent on your environment.

CCE-80439-3 Configure Time Service Maxpoll Interval (Control Group = na)

Check	
Comments	**Not a Finding.** The required `maxpoll 10` value is set in the `/etc/ntp.conf` file.

CCE-80447-6 Configure the Firewalld Ports (Control Group = n/a)

Check	Configure the `firewalld` ports to allow approved services to have access to the system. To configure `firewalld` to open ports, run the following command:
	`$ sudo firewall-cmd --permanent --add-port= or port_number/tcp $ sudo firewall-cmd --permanent --add-port=`
	Run the command list above for each of the ports listed below: <ports>
	To configure `service_nam firewalld` to allow access, run the following command(s):
	`firewall-cmd --permanent --add-service=ssh`

Comments	**Not a Finding.** NetWitness Platform `firewalld` service is disabled because it uses IP Tables, not FirewallD.

CCE-80515-0 Configure SSSD LDAP Backend Client CA Certificate Location (Control Group = n/a)

Check	Configure SSSD to implement cryptography to protect the integrity of LDAP remote access sessions. By setting the `ldap_tls_cacertdir` option in `/etc/sssd/sssd.conf` to point to the path for the X.509 certificates used for peer authentication. `ldap_tls_cacertdir /path/to/tls/cacert`
Comments	**Not a Finding.** NetWitness Platform does not currently support Multi-Factor authentication. As a result, SSSD service is not installed on a NetWitness Host.

CCE-80545-7 Verify and Correct Ownership with RPM (Control Group = n/a)

Check	The RPM package management system can check file ownership permissions of installed software packages, including many that are important to system security. After locating a file with incorrect permissions, which can be found with `rpm -Va | grep '^.....\(U|.G\)'`
	Run the following command to determine which package owns it:
	`$ rpm -qf`
	Next, run the following command to reset its permissions to the correct values:
	`FILENAME $ sudo rpm --setugids PACKAGENAME`

Comments	**Not a Finding.**Files/Directories with ownership differing from the rpm are generally COTS based and have been changed from root ownership to a specified COTS related account.

CCE-80546-5 Configure SSSD LDAP Backend to Use TLS For All Transactions (Control Group = n/a)

Check	This check verifies that RHEL7 implements cryptography to protect the integrity of remote LDAP authentication sessions. To determine if LDAP is being used for authentication, use the following command:
	`$ sudo grep -i useldapauth /etc/sysconfig/authconfig If USELDAPAUTH=yes`
	To check if LDAP is configured to use TLS, use the following command:
	`$ sudo grep -i ldap_id_use_start_tls /etc/sssd/sssd.conf`

Comments	**Not a Finding.**NetWitness Platform does not currently support Multi-Factor authentication. As a result, the SSSD service is not installed on a NetWitness Host.

CCE-80998-8 Verify firewall Enabled

Check	Verify the operating system enabled an application firewall. Check to see if "firewalld" is installed with the following command:
	`yum list installed firewalld`
	`firewalld-0.3.9-11.el7.noarch.rpm`
	If the "firewalld" package is not installed, ask the System Administrator if another firewall application (such as iptables) is installed.
	If an application firewall is not installed, this is a finding.
	Check to see if the firewall is loaded and active with the following command:
	`systemctl status firewalld`
	`firewalld.service - firewalld - dynamic firewall daemon`
	Loaded: loaded (/usr/lib/systemd/system/firewalld.service; enabled)
	Active: active (running) since Tue 2014-06-17 11:14:49 CEST; 5 days ago
	If "firewalld" does not show a status of "loaded" and "active", this is a finding.
	Check the state of the firewall:
	`firewall-cmd --state`
	`running`
	If "firewalld" does not show a state of "running", this is a finding.

Comments	**Not a Finding.** NetWitness Platform firewalldservice is disabled because it uses IP Tables, not FirewallD.

CCE-82035-7 Ensure /var/log/audit Located On Separate Partition

Check	Determine if the operating system is configured to have the "/var/log/audit" path is on a separate file system.
	`grep /var/log/audit /etc/fstab`
	If no result is returned, or the operating system is not configured to have "/var/log/audit" on a separate file system, this is a finding.
	Verify that "/var/log/audit" is mounted on a separate file system:
	`mount | grep "/var/log/audit"`
	If no result is returned, or "/var/log/audit" is not on a separate file system, this is a finding.

Comments	**Not a Finding**. NetWitness Platform has the /var/log directory as a separate partition.

CCE-82053-0 Ensure /tmp Located On Separate Partition

Check	Verify that a separate file system/partition has been created for "/tmp". Check that a file system/partition has been created for "/tmp" with the following command:
	`systemctl is-enabled tmp.mount`
	`enabled`
	If the "tmp.mount" service is not enabled, check to see if "/tmp" is defined in the fstab with a device and mount point:
	`grep -i /tmp /etc/fstab`
	UUID=a411dc99-f2a1-4c87-9e05-184977be8539 /tmp ext4 rw,relatime,discard,data=ordered,nosuid,noexec, 0 0
	If "tmp.mount" service is not enabled or the "/tmp" directory is not defined in the fstab with a device and mount point, this is a finding.

Comments	**Future Feature** - NetWitness Platform does not meet this requirement. NetWitness plans to fix this in a future release of NetWitness Platform.

CCE-82353-4 Ensure /var Located On Separate Partition

Check	Verify that a separate file system/partition has been created for "/var". Check that a file system/partition has been created for "/var" with the following command: `grep /var /etc/fstab` UUID=c274f65f /var ext4 noatime,nobarrier 1 2 If a separate entry for "/var" is not in use, this is a finding.
Comments	**Not a Finding.** Hardware is dedicated for NetWitness, and NetWitness software is installed in /var/netwitness by default and a separate partition is on /var/netwitness.

Rules Supported in a Future Release

The following checks for non-compliance to STIG rules are not supported in NetWitness Platform and will be added in a future release.

CCE-27277-3 Disable Modprobe Loading of USB Storage Driver (Control Group = services)

Check	To prevent USB storage devices from being used, configure the kernel module loading system to prevent automatic loading of the USB storage driver. To configure the system to prevent the usb-storage kernel module from being loaded, add the following line to a file in the `/etc/modprobe.d` directory : `install usb-storage /bin/true` This will prevent the `modprobe` program from loading the usb-storage module, but will not prevent an administrator (or another program) from using the `insmod` program to load the module manually.
Comments	**Future Feature.**

CCE-27309-4 Set Boot Loader Password in `grub2` (Control Group = fips-kernel)

Check	The `grub2` boot loader should have a superuser account and password protection enabled to protect boot-time settings. To do so, select a superuser account name and password and and modify the `/etc/grub.d/01_users` configuration file with the new account name. Because plain text passwords are a security risk, generate a hash for the password by running the following command:
	`$ grub2-setpassword`
	When prompted, enter the password that was selected.
	NOTE: It is recommended not to use common administrator account names like root, admin, or administrator for the `grub2` superuser account. Change the superuser to a different username (The default is 'root').
	`$ sed -i s/root/bootuser/g /etc/grub.d/01_users`
	To meet FISMA Moderate, the bootloader superuser account and password MUST differ from the root account and password. Once the superuser account and password have been added, update the `grub.cfg` file by running:
	`grub2-mkconfig -o /boot/grub2/grub.cfg`
	NOTE: Do NOT manually add the superuser account and password to the `grub.cfg` file as the `grub2-mkconfig` command overwrites this file.

Comments	**Future Feature.**

CCE-80179-5 Configure Kernel Parameter for Accepting IPv6 Source-Routed Packets for All Interfaces

Check	To set the runtime status of the `t.ipv6.conf.all.accept_source_route` kernel parameter, run the following command:
	`$ sudo sysctl -w net.ipv6.conf.all.accept_source_route=0`
	If this is not the system default value, add the following line to the `/etc/sysctl.conf` file:
	`t.ipv6.conf.all.accept_source_route = 0`

Comments	**Future Feature.**

CCE-80660-4 Record Any Attempts to Run setfiles (Control Group = audit)

Check	At a minimum, the audit system should collect any execution attempt of the setfiles command for all users and root. If the auditd daemon is configured to use the augenrules program to read audit rules during daemon startup (the default), add the following lines to a file with `.rules` in `/etc/audit/rules.d: -a always,exit -F path=/usr/sbin/setfiles -F perm=x -F auid>=1000 -F auid!=429496729as a suffix 5 -F key=privileged-priv_change`. If the `auditd` daemon is configured to use the `auditctl` utility to read audit rules during daemon startup, add the following lines to `/etc/audit/audit.rules` file: `-a always,exit -F path=/usr/sbin/setfiles -F perm=x -F auid>=1000 -F auid!=4294967295 -F key=privileged-priv_chang`
Comments	**Future Feature.**

CCE-80661-2 Ensure `auditd` Collects Information on Kernel Module Loading - `create_module` (Control Group = audit)

Check	To capture kernel module loading events, use following line, setting ARCH to either b32 for 32-bit system, or having two lines for both b32 and b64 in case your system is 64-bit: `-a always,exit -F arch=` The place where you add the line depends on the way `ARCH -S create_module -F key=modules auditd` daemon is configured. If it is configured to use the augenrules program (the default), add the line to a file with the `.rules` suffix in the `/etc/audit/rules.d` directory. If the `auditd` daemon is configured to use the `auditctl` utility, add the line to the `/etc/audit/audit.rules`file .
Comments	**Future Feature.**

Troubleshoot NetWitness Platform

For information about troubleshooting NetWitness, see the following topics:

- Debugging Information
- Error Notification
- Miscellaneous Tips
- NwLogPlayer: see the *Log Parser Customization Guide* for details.
- Troubleshoot Feeds

Debugging Information

NetWitness Log Files

The following files contain NetWitness log information.

Component	File
rabbitmq	/var/log/rabbitmq/nw@localhost.log /var/log/rabbitmq/nw@localhost-sasl.log
collectd	/var/log/messages
nwlogcollector	/var/log/messages
nwlogdecoder	/var/log/messages
sms	/opt/rsa/sms/wrapper.log
sms	/opt/rsa/sms/logs/sms.log
sms	/opt/rsa/sms/logs/audit/audit.log
NetWitness	/var/lib/netwitness/uax/logs/nw.log
NetWitness	/var/lib/netwitness/uax/logs/audit/audit.log
NetWitness	/opt/rsa/jetty9/logs

Files of Interest

The following files are used in key NetWitness components, and can be useful when trying to track down miscellaneous issues.

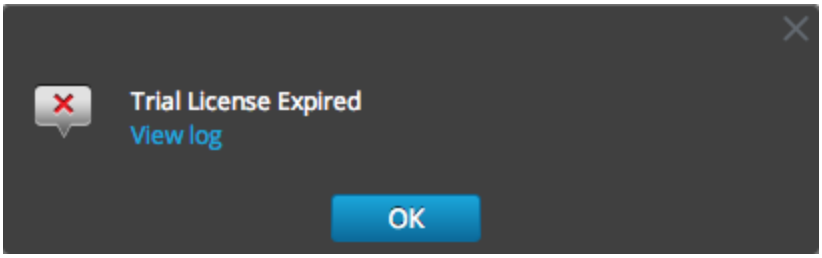
Component	File	Description
rabbit	/etc/rabbitmq/rabbitmq.config	RabbitMQ configuration file. This configuration file partially drives the behavior of RabbitMQ, particularly around network/SSL settings.
rabbit	/etc/rabbitmq/rabbitmq-env.conf	RabbitMQ environment configuration file. This file specifies the RabbitMQ node name and location of the enabled plugins file.

Component	File	Description
rabbit	/etc/rabbitmq/rsa_enabled_plugins	This file specifies the list of enabled plugins in RabbitMQ. This file is managed by the RabbitMQ server, with the `rabbitmq-plugins` command. This file overrides the `/etc/rabbitmq/enabled_plugins` path to work around issues with upgrading the Log Collector from early versions.
rabbit	/etc/rabbitmq/ssl/truststore.pem	The RabbitMQ trust store. This file contains a sequence of PEM-encoded X.509 certificates, represented trust CAs. Any clients that connect to RabbitMQ and present a certificate that is signed by a CA in this list is considered a trusted client.
rabbit	/var/log/rabbitmq/mnesia/nw@localhost	The RabbitMQ Mnesia directory. Mnesia is the Erlang/OTP database technology, for storing Erlang objects persistently. RabbitMQ uses this technology for storing information such as the current set of policies, persistent exchanges and queues, and so forth. Importantly, the `msg_store_persistent` and `msg_store_transient` directories are where RabbitMQ stores messages that have been spooled to disk, for example, if messages are published as persistent messages, or have paged off to disk due to memory limitations. Keep a close eye on this directory if disk or memory alarms have tripped in RabbitMQ. **Caution:** Do not delete these files manually. Use RabbitMQ tools to purge or delete queues. Modifying these files manually may render your RabbitMQ instance inoperable.

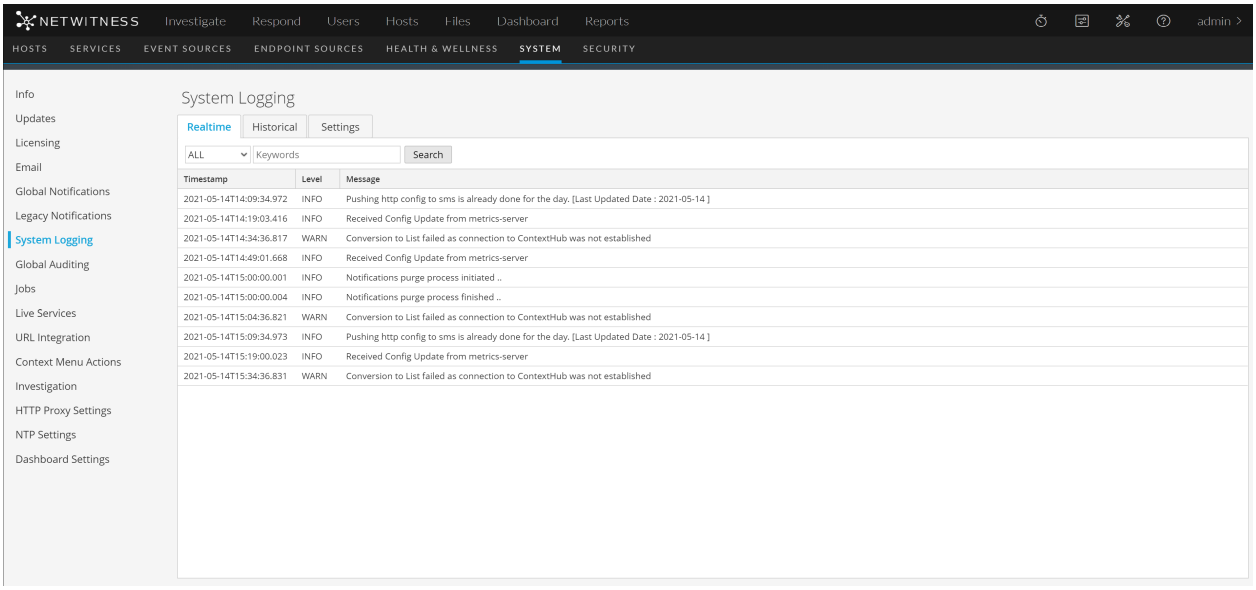
Error Notification

NetWitness has a set of error message types associated with different components and operations. NetWitness displays feedback in the form of a simple error notification and a log entry.

When an error notification dialog is displayed, you have two options: simply acknowledge the message or view the system log for more information.



If you want to view the system log for more information when an error notification is displayed, click **View log**. The log opens in the **(Admin)** > **System** view with a list of messages. Timestamp and message level are also listed.



Miscellaneous Tips

Audit Log Messages

It can be useful to see which user actions result in which log message types in the `/var/log/messages` file.

The event categories spreadsheet included in the log parser package in the NetWitness Parser v2.0.zip archive lists the event categories and the event parser lines to help with building reports, alerts, and queries.

NwConsole for Health & Wellness

RSA has added the command option `logParse` in NwConsole. This command option supports log parsing, a convenient way to check a log parser without setting up the full system to perform log parsing. For more information about the `logParse` command, at the command line, type `help logParse`.

Thick Client Error: remote content device entry not found

The remote content device entry was not found error can be generated for a correlation rule applied to a Concentrator. In NetWitness Investigate, if you click the `correlation-rule-name` meta value in the Alert meta key, you do not get session information.

Instead of using correlation rules on Decoders and Concentrators, use ESA rules. The ESA rules **do** record the correlation sessions that match the ESA rule.

View Example Parsers

Since Flex and Lua parsers are encrypted when they are delivered by Live, you cannot easily view their contents.

However, some plain text examples are available here: https://www.dell.com/support/home/en-us.

Configure WinRM Event Sources

The following Inside Dell article has a video that walks through the process of setting up Windows RM (Remote Management) collection: https://inside.dell.com/docs/DOC-122732.

Additionally, it contains two scripts that are shortcuts for procedures described in the "Windows Event Source Configuration Guide."

Troubleshoot Feeds

Overview

The purpose of the feed generator is to generate a mapping of an event source to the list of groups to which it belongs.

If you have an event source from which you are collecting messages, and it is not displayed in the correct event source groups, this topic provides background and information to help you track down the problem.

Details

The ESM Feed maps multiple keys to single value. It maps the DeviceAddress, Forwarder, and DeviceType attributes to groupName.

The purpose of the ESM feed is to enrich event source meta with the groupName collected on the Log Decoder.

How it Works

The feed generator is scheduled to update every minute. However, it is triggered only if there are any changes (create, update, or delete) in event sources or groups.

It generates a single feed file with event-source-to-group mapping, and pushes the same feed to all of the Log Decoders that are connected to NetWitness.

Once the feed file is uploaded on the Log Decoders, for any new events, it enriches events meta data with groupName, and appends this groupName to logstats.

Once the groupName is in logstats, the ESM Aggregator groups information and sends it to ESM. At this point, you should see the **Group Name** column under the **Event Source Monitoring** tab.

The entire process can take some time. Therefore, you may need to wait for several seconds after you add a new group or event source, before the Group name is displayed.

Note: If the event source type attribute changes when the feed is updated, NetWitness adds a new logstats entry rather than updating the existing one. Thus, there will be two different logstats entries in `logdecoder`. Previously existing messages would have been listed under the previous type, and all new messages are logged for the new event source type.

Feed File

The format of the feed file is as follows:

```
DeviceAddress, Forwarder, DeviceType, GroupName
```

The `DeviceAddress` is either `ipv4, ipv6`, or `hostname`, depending on which of these have been defined for the event source.

The following is a sample of the feed file:

```
"12.12.12.12","d6","NETFLOW","grp1"
"12.12.12.12","ld4","netflow","grp1"
"12.12.12.12","d6","netfow","grp1"
"0:E:507:E6:D4DB:E:59C:A","10.25.50.243","apache","Apachegrp"
"1.2.3.4","LCC","apache","Apachegrp"
"10.100.33.234","LC1","apache","Apachegrp"
"10.25.50.248","10.25.50.242","apache","Apachegrp"
"10.25.50.251","10.25.50.241","apache","Apachegrp"
"10.25.50.252","10.25.50.255","apache","Apachegrp"
"10.25.50.253","10.25.50.251","apache","Apachegrp"
"10.25.50.254","10.25.50.230","apache","Apachegrp"
"10.25.50.255","10.25.50.254","apache","Apachegrp"
"13.13.13.13","LC1","apache","Apachegrp"
"AB:F255:9:8:6C88:EEC:44CE:7",,"apache","Apachegrp"
"Appliance1234",,"apache","Apachegrp"
"CB:F255:9:8:6C88:EEC:44CE:7","10.25.50.253","apache","Apachegrp"
```

Troubleshooting

You can check the following items to narrow down where the problem is occurring.

Feed File Existence

Verify that the feeds ZIP archive exists in the following location:

`/opt/rsa/sms/esmfeed.zip`

Do not modify this file.

Group Meta Populated on LD

Verify that the group meta is populated on the Log Decoder. Navigate to the Log Decoder REST and check logstats:

`http://LogDecoderIP:50102/decoder?msg=logStats&force-content-type=text/plain`

This is a sample logstats file with group information:

```
device=apache  forwarder=NWAPPLIANCE10304  source=1.2.3.4  count=338
lastSeenTime=2015-Feb-04 22:30:19  lastUpdatedTime=2015-Feb-04 22:30:19
groups=IP1234Group,apacheGroup
device=apachetomcat  forwarder=NWAPPLIANCE10304  source=5.6.7.8  count=1301
```

```
lastSeenTime=2015-Feb-04 22:30:19  lastUpdatedTime=2015-Feb-04 22:30:19
groups=AllOtherGroup,ApacheTomcatGroup
```

In the above text, the group information is **bold**.

Device Group Meta on Concentrator

Verify that the **Device Group** meta data exists on the Concentrator, and that events have values for the `device.group` field.



SMS Log File

Check the SMS log file in the following location to view informational and error messages: `/opt/rsa/sms/logs/sms.log`

The following are examples informational messages:

```
Feed generator triggered...
Created CSV feed file.
Created zip feed file.
Pushed ESM Feed to LogDeocder : <logdecoder IP>
```

The following are examples of error messages:

```
Error creating CSV File : <reason>Unable to push the ESM Feed: Unable to
create feed zip archive.
Failed to add Group in CSV: GroupName: <groupName> : Error: <error>
Unable to push the ESM Feed: CSV file is empty, make sure you have al-least on
group with al-least one eventsource.
Unable to push the ESM Feed: No LogDecoders found.
```

```
Unable to push the ESM Feed: Unable to push feed file on LogDecoder-
<logdecoderIP>Unable to push the ESM Feed:
admin@<logdecoderIP>:50002/decoder/parsers received error: The zip archive
"/etc/netwitness/ng/upload/<esmfeedfileName>.zip" could not be opened
Unable to push the ESM Feed: <reason>
```

Verify Logstats Data is Getting Read and Published by ESMReader and ESMAggregator

These are the steps to verify that logstats are collected by `collectd` and published to Event Source Management.

ESMReader

1. On LogDecoders add the **debug "true"** flag in **/etc/collectd.d/NwLogDecoder_ESM.conf**:

```
#
# Copyright (c) 2014 RSA The Security Division of EMC
#
<Plugin generic_cpp> PluginModulePath "/usr/lib64/collectd"
    debug "true"
    <Module "NgEsmReader" "all"> port "56002"
        ssl       "yes"
        keypath   "/var/lib/puppet/ssl/private_keys/d4c6dcd4-6737-4838-
a2f7-ba7e9a165aae.pem"
        certpath  "/var/lib/puppet/ssl/certs/d4c6dcd4-6737-4838-a2f7-
ba7e9a165aae.pem"
        interval  "600"
        query     "all"
        <stats></stats></Module><Module "NgEsmReader" "update"> port
"56002" ssl "yes"
        keypath   "/var/lib/puppet/ssl/private_keys/d4c6dcd4-6737-4838-
a2f7-ba7e9a165aae.pem"
        certpath  "/var/lib/puppet/ssl/certs/d4c6dcd4-6737-4838-a2f7-
ba7e9a165aae.pem"
        interval  "60"
        query     "update"
        <stats></stats></Module></Plugin>
```

2. Run the command:
   `collectd service restart`
3. Run the following command:
   `tail –f /var/log/messages | grep collectd`
   Verify that ESMReader is reading logstats and there are no errors. If there are any read issues, you will see errors similar to the following:

```
Apr 29 18:47:45 NWAPPLIANCE15788 collectd[14569]: DEBUG: NgEsmReader_all:
error getting ESM data for field "groups" from logstat device=checkpointfw1
forwarder=PSRTEST source=1.11.51.212. Reason: <reason>Apr 29 18:58:36
NWAPPLIANCE15788 collectd[14569]: DEBUG: NgEsmReader_update: error getting ESM
data for field "forwarder" from logstat device=apachetomcat
source=10.31.204.240. Reason: <reason>
```

ESMAggregator

1. On NetWitness, uncomment the verbose flag in **/etc/collectd.d/ESMAggregator.conf**:

```
# ESMAggregator module collectd.conf configuration file
#
# Copyright (c) 2014 RSA The Security Divsion of EMC
#

<Plugin generic_cpp>    PluginModulePath "/usr/lib64/collectd"

<Module "ESMAggregator">
        verbose 1
        interval "60"
        cache_save_interval "600"
        persistence_dir "/var/lib/netwitness/collectd"
</Module>    </Plugin>
```

2. Run the following command:
   `collectd service restart.`
3. Run the following command:
   `run “tail –f /var/log/messages | grep ESMA`
   Look for for ESMAggregator data and make sure your logstat entry is available in logs.

Sample output:

```
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[0]
logdecoder[0] = d4c6dcd4-6737-4838-a2f7-ba7e9a165aae
```

```
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[1]
logdecoder_utcLastUpdate[0] = 1425174451
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[2]
groups = Cacheflowelff,Mixed
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[3]
logdecoders = d4c6dcd4-6737-4838-a2f7-ba7e9a165aae
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[4]
utcLastUpdate = 1425174451
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: Dispatching
ESM stat NWAPPLIANCE15788/esma_update-cacheflowelff/esm_counter-3.3.3.3 with a
value of 1752 for NWAPPLIANCE15788/cacheflowelff/esm_counter-3.3.3.3
aggregated from 1 log decoders
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[0]
logdecoder[0] = 767354a8-5e84-4317-bc6a-52e4f4d8bfff
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[1]
logdecoder_utcLastUpdate[0] = 1425174470
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[2]
groups = Cacheflowelff,Mixed
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[3]
logdecoders = 767354a8-5e84-4317-bc6a-52e4f4d8bfff
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: MetaData[4]
utcLastUpdate = 1425174470
Mar  1 02:32:08 NWAPPLIANCE15936 collectd[11203]: ESMAggregator: Dispatching
RRD stat NWAPPLIANCE15788/esma_rrd-cacheflowelff/esm_counter-3.3.3.3 with a
value of 1752 for NWAPPLIANCE15788/cacheflowelff/esm_counter-3.3.3.3
aggregated from 1 log
```

Configure JMX Feed Generator Job Interval

Although the feed generation job is scheduled to execute every minute by default, you can change this by using `jconsole`.

To change the feed generator job interval:

1. Open **jconsole** for the SMS service.
2. On the MBeans tab, navigate to **com.rsa.netwitness.sms** > **API** > **esmConfiguration** > **Attributes**.
3. Modify the value for the property **FeedGeneratorJobIntervalInMinutes**.

4. Go to **Operations** under the same navigation tree, an click **commit()**. This persists the new value in the corresponding json file under **/opt/rsa/sms/conf**, and uses the value if SMS is restarted.

Setting a new value reschedules the feed generator job for the new interval.

Troubleshooting Cert-Reissue Command

You must contact Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897) to troubleshoot problems. You know there is a problem if any `<host-id>` does not return a `Success`**Status**. `Success` indicates that certificates were reissued for a host.

Argument Options Used for Troubleshooting

You use the following argument options with `cert-reissue --host-all` to troubleshoot problems.

You can run `cert-reissue --host-all<arguments>` multiple times without an adverse effect.

Note: Use the following Argument Options with caution. They force the `cert-reissue` command to execute for all the hosts.

Argument Option	Description
`--skip-health-checks`	Reissues certificates for all hosts at one time without applying system health checks (force Reissue). This means that the command does not: • verify that all hosts are online line. • verify that all services are running. **Use case:** You have numerous hosts and you know that a small minority of them will fail. This updates all the hosts that conform to the checking rules and you can reissue certificates for the others subsequently with the help of Customer Support.
`--skip-version-checks`	Do not verify that hosts are running version 11.4.0.0 or later. **Use Case:** You have numerous hosts and your know that some of them are not updated to 11.4 or later. This reissues certificates for all the hosts that are at 11.4 or later and you can reissue certificates for the others subsequently with the help of Customer Support.
`--ignore-trigger-errors`	Ignore any errors that trigger failures. This option forces the cert reissue process to continue disregarding the errors instead of aborting or failing the cert reissue command quickly. When a cert reissue for a host succeeds, the reissued certificates on that host are not provisioned to other dependent hosts (referred to as trusts). In this case, the: • host with reissued certificates is reported as “Partial.” • the hosts with trusts that failed to update are listed separately in the summary table to tell you that these hosts may require a refresh using the new `--refresh-trusts-only` option.

Argument Option	Description
`--refresh-trusts-only`	Refreshes trusts exclusively for host identified by `<id>` (does not reissue certificates for that host).

Problems and How to Troubleshoot Them

This section describes solutions to problems that you may encounter when running the cert-reissue command to reissue certificates with suggested causes and solutions.

Status	Failed!
Error Message	... 2019-02-06 13:34:39.646 INFO 8540 --- [main] c.r.n.i.o.client.OrchestrationClient : Checking host connections... ... 2019-02-06 13:34:57.861 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Host '192.168.200.99' (nw-platform-esa-primary) verification failed! ... 2019-02-06 13:34:57.862 INFO 8540 --- [main] c.r.n.i.o.client.OrchestrationClient : Checking status of services... 2019-02-06 13:35:57.931 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Service 'nw-platform-node-zero - Investigate Server' not available! ... +---------+--------------------+--------+--+ | | Host | Status | Message | +---------+--------------------+--------+--+ |<host-id>| <IP-address> |Failed! |failed to connect, is host online? | |<host-id>| <IP-address> |Failed! |service(s) down | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | +---------+--------------------+--------+--+
Cause	`cert-reissue --host-all` failed because one or more hosts are offline or one or more run time services are unreachable. You can force this command to run in spite of this error by specifying the `--skip-health-checks` option, that is: `cert-reissue --host-all--skip-health-checks`
Solution	1. Bring appropriate hosts back online or make sure the NW Server hosts run time services are running. 2. Run `cert-reissue` for the hosts affected.

Status	Failed!
Error Message	... 2019-02-06 13:34:39.643 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Host '192.168.200.102' (nw-platform-decoder) version '11.2.0.0' not supported, minimum required version: 11.3.0.0 2019-02-06 13:34:39.644 ERROR 8540 --- [main] c.r.n.i.o.client.HostValidator : Host '192.168.200.101' (nw-platform-concentrator) version '11.2.0.0' not supported, minimum required version: 11.3.0.0 ... +---------+--------------------+--------+---+ | | Host | Status | Message | +---------+--------------------+--------+---+ |<host-id>| <IP-address> |Failed! |version <version-earlier-than-11.3.0.0> not supported| |<host-id>| <IP-address> |Failed! |version <version-earlier-than-11.3.0.0> not supported| |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | +---------+--------------------+--------+---+
Cause	`cert-reissue -host-all` command string failed because one or more hosts are running a version earlier than 11.4.0.0 **Note:** You can force the reissue of certificates for the remaining hosts using the `-skip-version-checks` argument.
Solution	Update the host to 11.4 or later and run `cert-reissue` for that host again.

Status	Partial
Error Message	... 2019-02-06 02:27:09.078 ERROR 20647 --- [main] c.r.n.i.o.client.OrchestrationClient : Trigger failed for host '<IP-address>' (nw-platform-decoder) 2019-02-06 02:27:09.079 ERROR 20647 --- [main] c.r.n.i.o.client.OrchestrationClient : Trigger failed for host '<IP-address>' (nw-platform-concentrator) ... 2019-02-06 02:27:09.118 WARN 20647 --- [main] c.r.n.i.o.client.OrchestrationClient : One or more host(s) may require manual refresh due to failed triggers: +---------+--------------------+ | | Host | +---------+--------------------+ |<host-id>| <IP-address> | |<host-id>| <IP-address> | +---------+--------------------+ ... +---------+--------------------+--------+---+ | | Host | Status | Message | +---------+--------------------+--------+---+ |<host-id>| <IP-address> |Partial |Reissue completed, triggers failed | |<host-id>| <IP-address> |N/A |[Skipped…] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | |<host-id>| <IP-address> |N/A |[Skipped...] | +---------+--------------------+--------+---+
Cause	`cert-reissue` command completed on NW Server host however one or more triggers failed. This aborted the `cert-reissue` command for other hosts.
Solution	Address all the errors and run the `cert-reissue --host-all<arguments>` command string again.

Status	Partial
Error Message	... 2019-02-06 14:18:03.208 ERROR 17800 --- [main] c.r.n.i.o.client.OrchestrationClient : Trigger failed for host '192.168.200.82' (nw-platform-node-x) ... ... 2019-02-06 14:29:05.200 WARN 17800 --- [main] c.r.n.i.o.client.OrchestrationClient : One or more host(s) may require manual refresh due to failed triggers: +---------+--------------------+ | | Host | +---------+--------------------+ |<host-id>| <IP-address> | +---------+--------------------+ ... +---------+--------------------+--------+---+ | | Host | Status | Message | +---------+--------------------+--------+---+ |<host-id>| <IP-address> |Failed! |Cert reissue failed! | |<host-id>| <IP-address> |Partial |Reissue completed, triggers failed | |<host-id>| <IP-address> |Success | Cert reissue successful | |<host-id>| <IP-address> |Success | Cert reissue successful | |<host-id>| <IP-address> |Success | Cert reissue successful | +---------+--------------------+--------+---+
Cause	One or more hosts did not pass system health checks. In addition, one or more of the unhealthy hosts are running core services, which will result in the NW Server host `cert-reissue` to fail (because of failed triggers explained above). By disabling health checks and trigger errors, you can continue the process and reissue certificates for the remaining hosts. The NW Server host **Status** is reported as `Partial` because the `cert-reissue` command completed for the NW Server but downstream triggers failed for other hosts.
Solution	Manually refresh the failed core hosts (to synchronize trust peers). Submit the following command string to reissue certificates for healthy hosts. `cert-reissue --host-all --skip-health-checks --ignore-trigger-errors`

References

This section describes the NetWitness user interface views in which you can perform system maintenance tasks. You use this interface to:

- Monitor and maintain services (settings, statistics, command and message syntax, REST API, NetWitness Console utility, and protocols supported in NetWitness).
- Display the current NetWitness version and license status.
- Manage your Local Update Repository from which you apply software version updates to hosts.

The following topics describe each interface in detail:

- Health and Wellness View
  - Health and Wellness View - Alarms View
  - Event Source Monitoring View
  - Health and Wellness Historical Graphs
    - Historical Graph View for Events Collected from an Event Source
    - Historical Graph for System Stats
  - Health and Wellness Settings View - Archiver
  - Health and Wellness Settings View - Event Sources
  - Health and Wellness Settings View - Warehouse Connector
  - Monitoring View
  - Policies View
  - System Stats Browser View
- New Health & Wellness Tab
- System View - System Info Panel
- System Updates Panel - Settings Tab
- System Logging - Settings View
- System Logging - Realtime Tab
- System Logging - Historical Tab

Health and Wellness View

The Health and Wellness settings allow you to set and view alarms, monitor events, and view policies and system statistics. For more details on each of these, see the following topics:

- Health and Wellness View - Alarms View
- Event Source Monitoring View
- Health and Wellness Historical Graphs
- Health and Wellness Settings View - Archiver
- Health and Wellness Settings View - Event Sources
- Health and Wellness Settings View - Warehouse Connector
- Monitoring View
- Policies View
- System Stats Browser View

Health and Wellness View - Alarms View

You can monitor hosts and services to determine when user-defined limitations have been reached by viewing all the active alarms. Policy rules, that you define or assign to hosts and services, in the **Policies tab** trigger these alarms. You can:

- View all the alarms that are currently active for all your systems and services
- Select an alarm and view its details

What do you want to do?

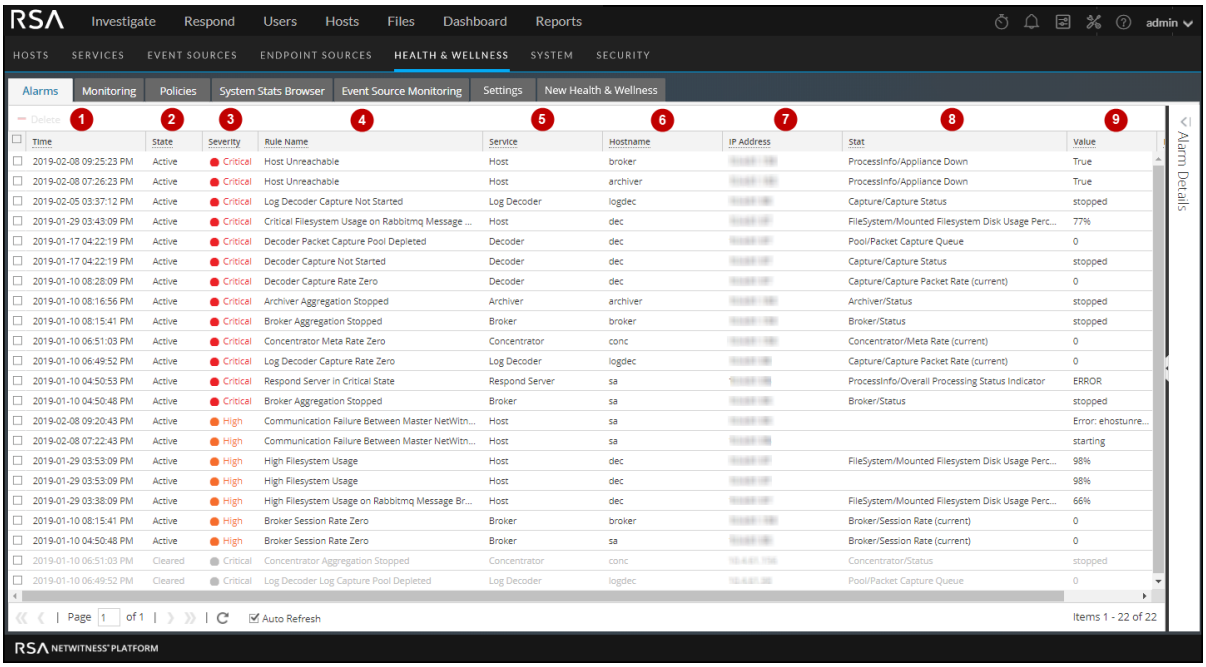
Role	I want to ...	Show me how
Administrator	View the alarm status of NetWitness Servers and services.	Monitor Alarms
Administrator	View detailed information about a specific alarm.	Monitor Alarms

Related Topics

Manage Policies

Quick Look

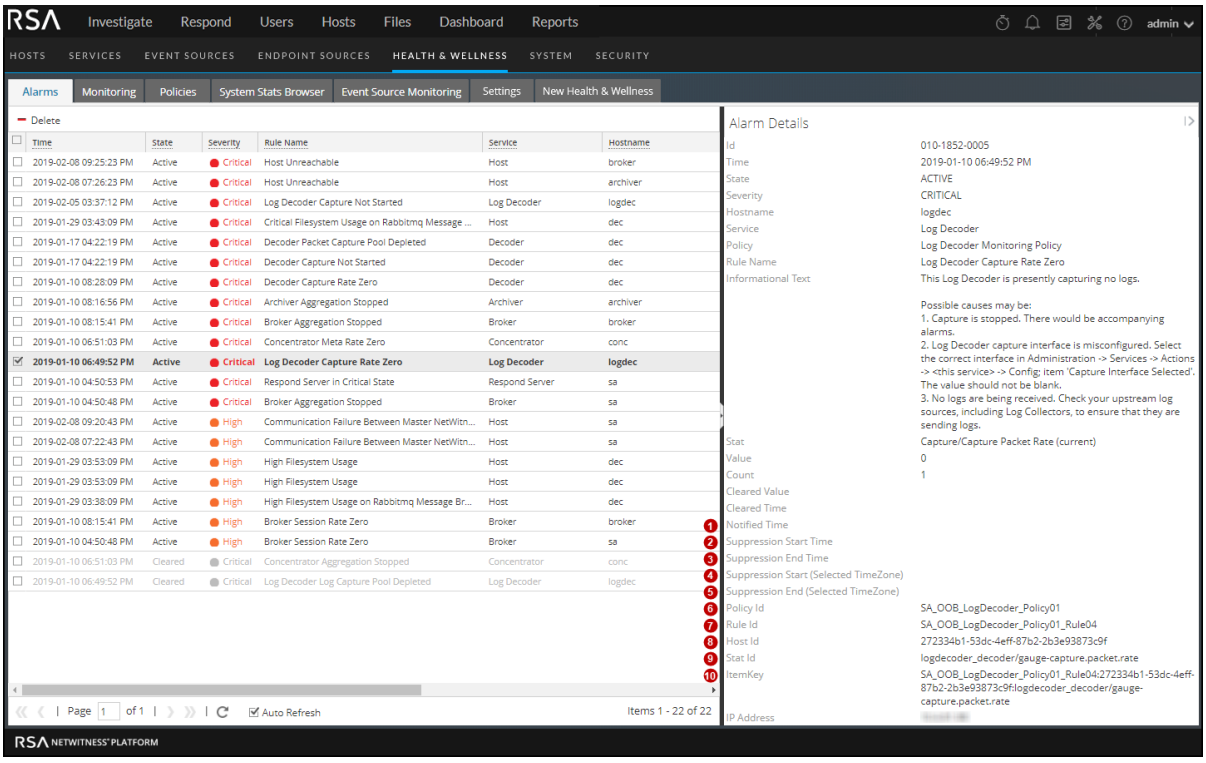
The required permission to access this view is **Manage services**. To access the Alarms view, go to **(Admin)** > **Health & Wellness**. The Health & Wellness view opens with the Alarms tab displayed. The Alarms tab contains an alarms list and an Alarm Details panel.



1	Time when the alarm was triggered.
2	Status of the alarm: • **Active** - the statistical threshold was crossed triggering the alarm. • **Cleared** - the clearing threshold was crossed and the alarm is no longer active.
3	Severity assigned to this alarm: • **Critical** • **High** • **Medium** • **Low**
4	Name of the rule that triggers the alarm.
5	Service defined in the rule.
6	Host on which the alarm is triggered.
7	Statistic selected in the rule that triggers the alarm.
8	Value of the statistic that triggered the alarm.
9	Identification number of the alarm.

Note: NetWitness sorts the alarms in time order. You can sort the relevant parameters in ascending or descending order.

This figure shows the Alarms tab with the Alarm Details panel expanded.



Alarm Details Panel

The Alarm Details panel displays information for the alarm selected in the Alarms list. It contains all the information in the Alarms list plus the following fields.

1. Alarm Notified time
2. Suppression start time
3. Suppression end time
4. Suppression start (selected time zone)
5. Suppression end (selected time zone)
6. The Policy ID
7. The Rule ID
8. The Host ID
9. The Stat ID
10. Item key

Event Source Monitoring View

Note: For NetWitness 11.3.1, this view has been deprecated: this functionality is now available in the **(Admin) > Event Sources > Discovery** view. To manage Event Sources, see "About Event Source Management" in the *NetWitness Event Source Management Guide*.

Health and Wellness Historical Graphs

Configuring Archiver monitoring enables you to automatically generate notifications when critical thresholds concerning Archiver aggregation and storage have been met. The Historical Graph view provides a visualization of historical data.

Note: Historical graphs are not available for non-numeric statistics, and is indicated by a greyed-out icon.

See the following topics for more details:

- Historical Graph View for Events Collected from an Event Source
- Historical Graph for System Stats

Historical Graph View for Events Collected from an Event Source

Note: For NetWitness 11.4.1, this view has been deprecated. To manage Event Sources, use the  (Admin) > Event Sources view. For details, see "About Event Source Management" in the*NetWitness Event Source Management Guide*.

Historical Graph for System Stats

To access the Historical Graph for the System Stats:

1. Go to **(Admin) > Health & Wellness**.

   The Health & Wellness view is displayed with the Alarms tab open.

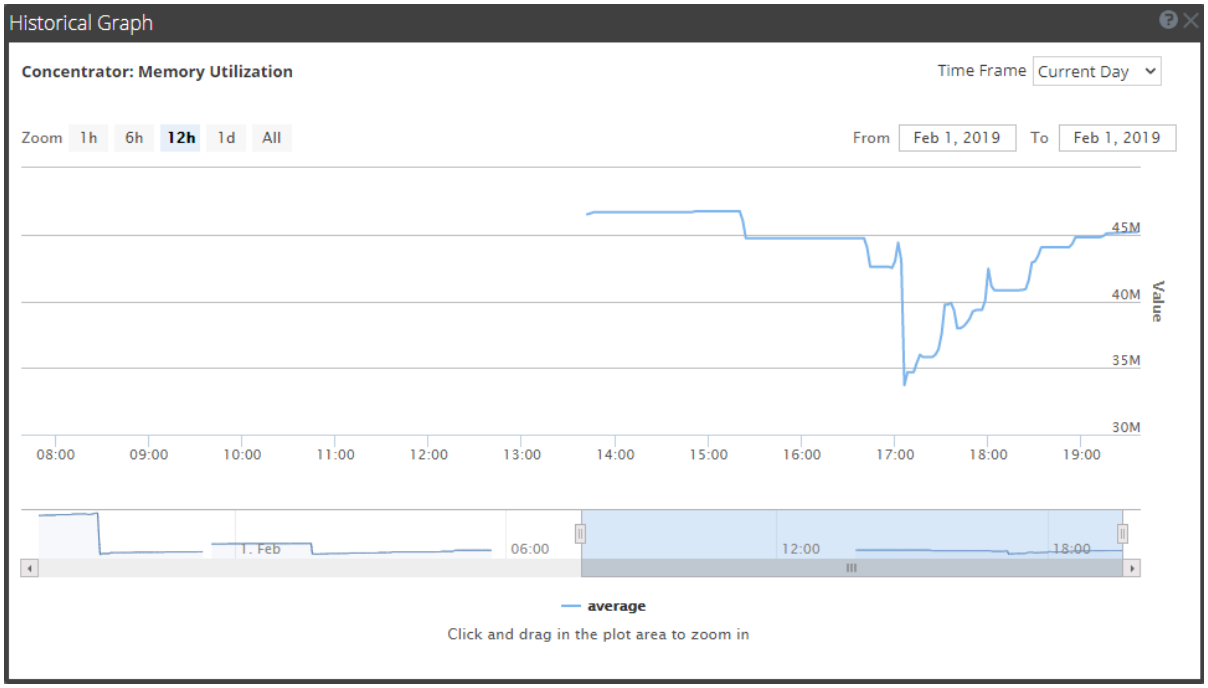
2. Click the **System Stats Browser** tab.

   The System Stats Browser tab is displayed.

3. In the **Historical Graph** column, select .

   The Historical graph for the selected statistic for a host is displayed.

   The figure displays the system stats view for the Memory Utilization statistics.



Parameters

You can customize the graph view as required. The table lists the various parameters used to customize the historical graph view.

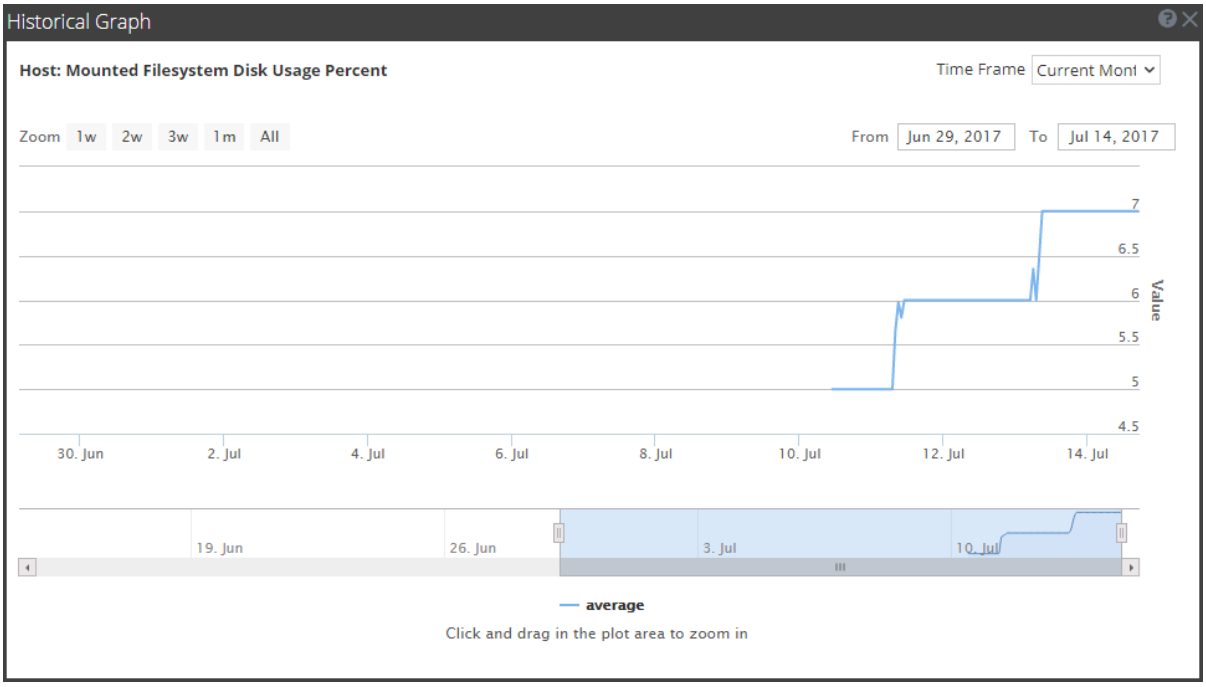
Parameter	Description
Time Frame	Select the time frame for which you want to view the historical data. The available options are: **Current Day**, **Current Week**, **Current Month**, and **Current Year**.
From <date> To <date>	Select the date range for which you want to view the historical data,

You can zoom in for a detailed view of the data in the Historical graph.

Zoom in function 1 and 2:

You can select one of the values to view the historical data for the selected value. The figure below displays an example for the 6h frame selected for zoom in. The slider bar at the right bottom corner is also changed to a 6h window.

Alternatively, you can slide the bar in the right hand corner to zoom in to a required frame.



Zoom in function 3:

You can click and drag in the plot area to zoom in for a required frame of time.

The figure below displays an example of how the graph appears while you click and drag.



Health and Wellness Settings View - Archiver

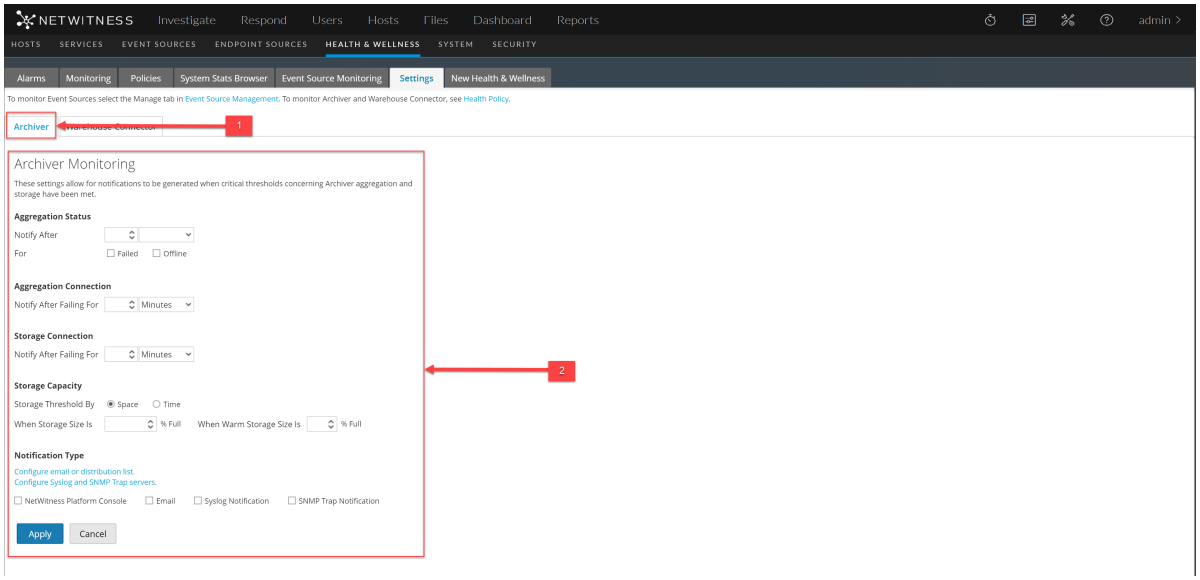
To access the Archiver Monitoring view:

1. Go to  **(Admin) > Health & Wellness**.
2. Select **Settings > Archiver**.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Monitor service details of Archiver	Monitor Service Details

Quick Look



1	Displays Archiver Monitoring Panel
2	Configures Archiver Monitoring Panel to automatically receive notification

Features

The following table lists the parameters required to configure Archiver to automatically generate notification when critical thresholds are reached.

Parameter	Value	Description
Aggregation Status	Notify After	Number of minutes or hours after which you are notified of aggregation status
	For	Failed - If enabled, you are notified when the Archiver aggregation status is failed for the defined number of minutes or hours. Offline - If enabled, you are notified when the Archiver aggregation status is offline for the defined number of minutes or hours.
Aggregation Connection	Notify After Failing for	Number of minutes or hours after which you receive notification if the Archiver aggregation connection fails.
Storage Connection	Notify After Failing for	Number of minutes or hours after which you receive notification if the Archiver storage connection fails.
Storage Capacity	Storage Threshold By	Select **Space** if you want to receive a notification when the Archiver storage capacity exceeds the percentage defined in the **When Storage Size Is** field. Select **Time** if you want to receive a notification when the files stored in the Archiver exceeds the defined number of days in the **When Oldest Storage File Is** field.
	When Storage Size Is	Enter the percentage of used storage to trigger a notification.
	When Warm Storage Size Is	Enter the percentage of used storage on the warm server to trigger a notification.
Notification Type	Configure email or distribution list	Click to configure email so that you can receive notifications in NetWitness.
	Configure Syslog and SNMP Trap servers	Click to configure audit logs.
	NW Console, Email, Syslog Notification, SNMP Trap Notification	Enable NW Console to get notifications on the NetWitness UI notification toolbar. Enable Email to get email notifications. Enable Syslog Notification to generate syslog events. Enable SNMP Trap Notification to get audit events as SNMP traps.

Health and Wellness Settings View - Event Sources

Note: For NetWitness 11.4.1, this view has been deprecated. To manage Event Sources, use the  (Admin) > Event Sources view. For details, see "About Event Source Management" in the*NetWitness Event Source Management Guide*.

Health and Wellness Settings View - Warehouse Connector

Configuring the Warehouse Connector monitoring enables you to automatically generate notification when critical thresholds concerning Warehouse Connector and storage have been met.

Access the Warehouse Connector Monitoring view

1. Go to  **(Admin) > Health & Wellness**.
2. Select **Settings > Warehouse Connector**.

What do you want to do?

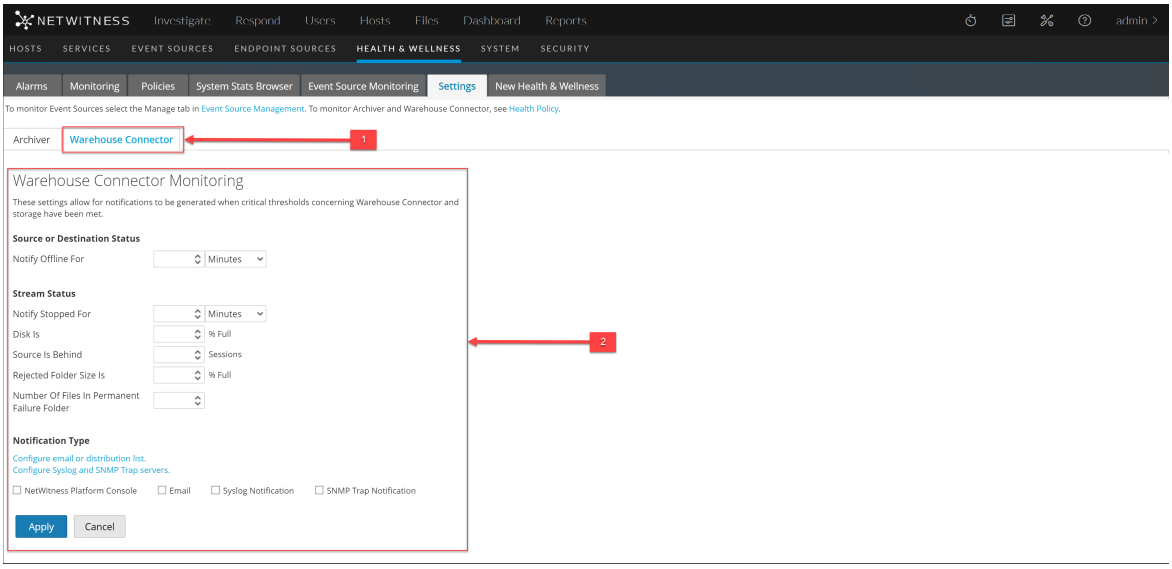
Role	I want to ...	Show me how
Administrator	View the details of Warehouse connector	Warehouse Connector Details View

Related topics

Monitor Service Details

Quick Look

The Warehouse Connector Monitoring view is displayed.



1 Displays the Warehouse Connector Monitoring view panel.
2 Allows you to configure Warehouse Connector Monitoring parameters.

Warehouse Connector Monitoring parameters

The following table lists the parameters required to configure the Warehouse Connector to automatically generate notification when critical thresholds are reached.

Parameter	Value	Description
Source or Destination Status	Notify Offline For	Number of minutes or hours after which you receive a notification if the source or destination connection fails.
Stream Status	Notify Stopped For	Number of minutes or hours after which you receive a notification when the Stream goes offline.
	Disk Is	The limit on the percentage of disk usage after which you would like to receive a notification.
	Source Is Behind	Number of sessions after which a notification is raised if the source goes behind the defined number of sessions.
	Rejected Folder Size Is	Limit on the percentage of folder usage after which you receive a notification.
	Number Of Files in Permanent Failure Folder	Limit on the number of files in the permanent failure folder after which you receive a notification.
Notification Type	Configure email or distribution list	Click to configure email so that you can receive notifications in NetWitness.
	Configure Syslog and SNMP Trap servers	Click to configure audit logs.
	NW Console, Email, Syslog Notification, SNMP Trap Notification	Enable NW Console to get notifications on the NetWitness UI notification toolbar. Enable Email to get email notifications. Enable Syslog Notification to generate syslog events. Enable SNMP Trap Notification to get audit events as SNMP traps.

Monitoring View

NetWitness provides detailed statistics and other information about the host and the individual NetWitness services in Details views. You can view the current health of all the hosts and the services running on the hosts in the Monitoring view.

What do you want to do?

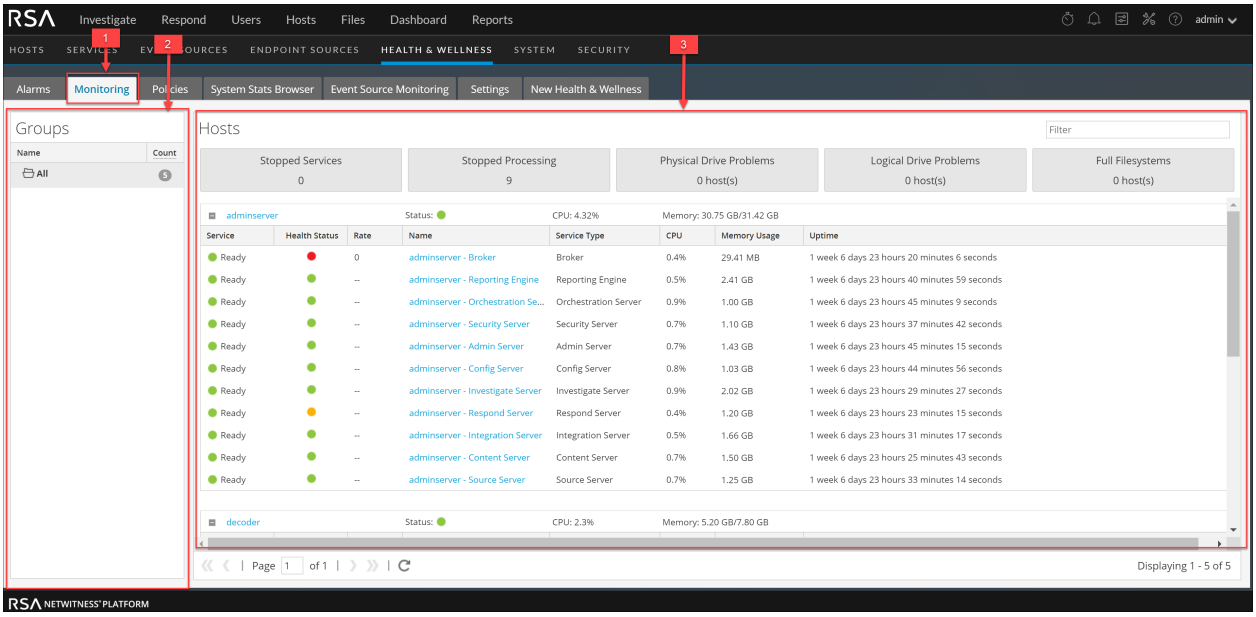
Role	I want to ...	Show me how
Administrator	View and Perform Procedures	Monitor Hosts and Services

Quick Look

To access this view:

1. Go to **(Admin) > Health & Wellness**.
2. Click the **Monitoring** tab.

The Monitoring view is displayed.



1 The Monitoring tab shows health statistics for the NetWitness Platform hosts and services.
2 The Group panel enables you to view statistics for a selected group.
3 The Hosts panel displays operational statistics.

Groups Panel

The Groups panel lists all of the groups of hosts available. When you select a group, the associated content is displayed in the Hosts panel.

Note: If the total host count in the Groups panel is lower than the actual number of hosts displayed in the Hosts panel, refer to the Troubleshooting Health & Wellness topic for possible causes of this issue and recommended solutions.

Hosts Panel

The Hosts panel displays operational statistics for hosts and the services running on each host.

Parameter	Description
Filter	Type a host name or a service name in the Filter field to display the corresponding hosts and services in the Host panel.
Stopped Services	Click **Stopped Services** to display a list of all stopped services. It also displays the host on which the service is installed.
Stopped Processing	Click **Stopped Processing** to display a list of all the hosts that have services installed on them that are in the stopped processing status.
Physical drive Problems <#> host(s)	Click to view the hosts that have physical drive problems.
Logical Drive Problems <#> host(s)	Click to view the hosts that have logical drive problems.
Full Filesystems <#> host(s)	Click to view the hosts that have full file systems.

Note: The summary information in the boxes at the top displays the System Statistics for all of the hosts configured in NetWitness and does not change with the application of filters on the groups.

Below the boxes at the top of the Hosts panel is a list of hosts, the services installed on them, and information regarding the hosts and services.

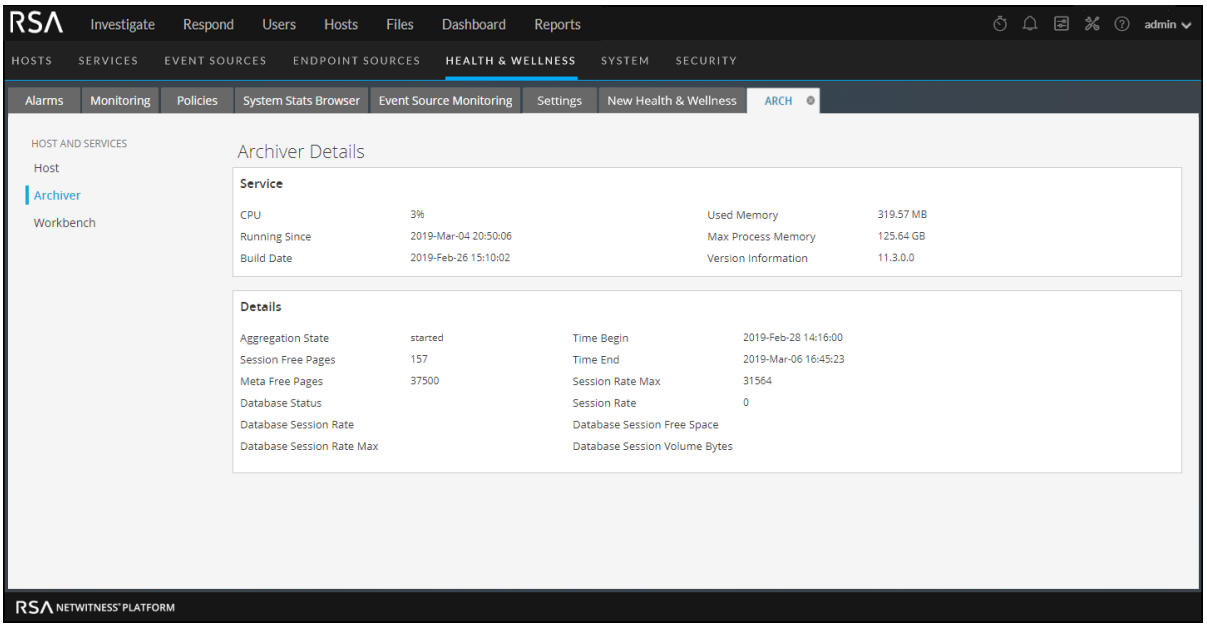
Parameter	Description
Host Name	Displays the host name. If a host has services installed that are not in view, you will see a ⊞ prefixed to the host name. Click ⊞ to view all the services installed on the host.
Status	Displays the status of the Host. ● - The host is active and running. ● - The host is stopped or yet to start processing.
CPU	Displays the current CPU usage of the host.
Memory	Displays the Memory used by the host.

When you click ⊞ prefixed to the host name, a list of all the services installed on the host is displayed. The table below describes parameters displayed for a service and their description.

Parameter	Description
Service	Displays the status of the service. ● Ready - The service is active and running. ● Stopped - The service is stopped or yet to start processing.
Health Status	Displays the processing status of the service.  - The process is running and the data is being processed at a rate greater than zero. ● - The processing is stopped. ● - The processing is turned on but the data is not being processed.
Rate	Shows the rate at which data is being processed.
Name	Name of the service in the format `<host> - <service>`. Click the link in the Name field to get additional service details.
Service Type	Name of the type of service.
CPU	Shows the current CPU usage of the service.
Memory Usage	Displays the Memory used by the service.
Uptime	Displays the time for which the service has been running.

Archiver Details View

The Archiver Details view provides information about the Archiver. The following figure shows the Archiver details.



For the related procedure, see Monitor Service Details

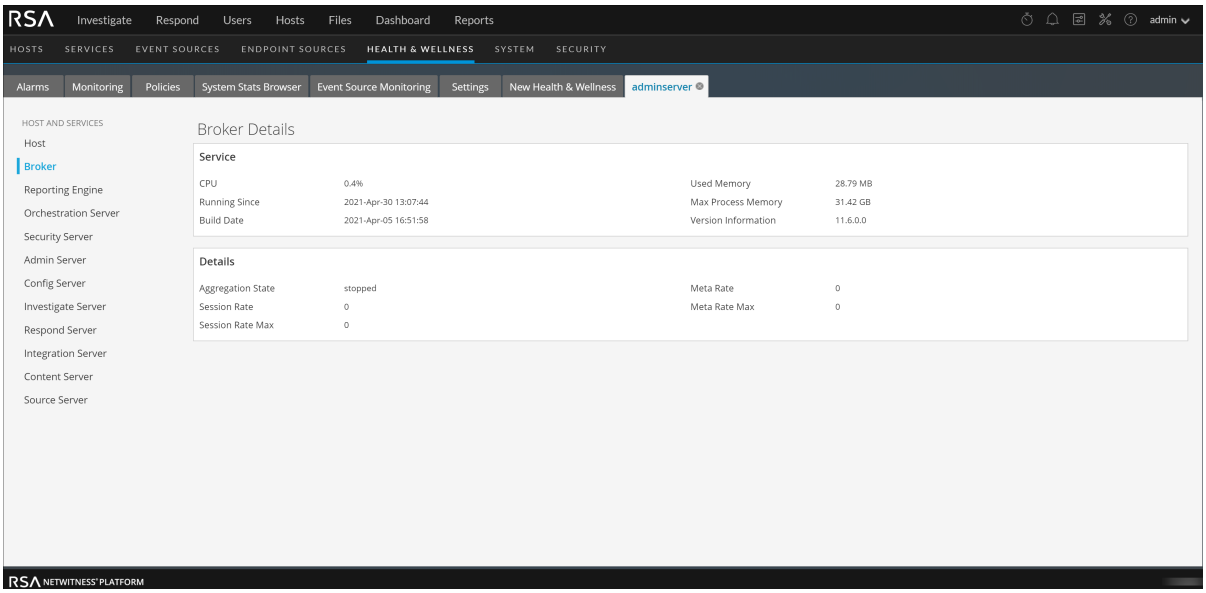
This section displays the current generic statistics for the service.

Statistic	Description
Aggregation State	State of data aggregation.
Time Begin	Time (UTC) when the first session was tracked by the index.
Session Free Pages	Session pages available for aggregation.
Time End	Time (UTC) when the last session was tracked by the index.
Meta Free Pages	Pages available for aggregation.
Session Rate Max	Maximum sessions per second rate.

Statistic	Description
Database Status	Status of databases. Valid values are: • `closed` - not available for QUERY and UPDATE (databases are being initialized). This value is seldom seen. • `opened` - available for QUERY and UPDATE. • `failure` - failed to open. This can happen for any number of reasons. You can check this if CAPTURE fails to start or if queries fail to return data. This is normally caused by database corruption.
Session Rate	Sessions per second rate.
Database Session Rate	Per second rate at which the service is writing sessions to the database.
Database Session Free Space	Amount of session free space available for aggregation.
Database Session Rate Max	Maximum per second rate at which the service is writing sessions to the database.
Database Session Volume Bytes	Number of session bytes in the database.

Broker Details View

The Broker Details view provides information about the Broker. The following figure shows the Broker details.



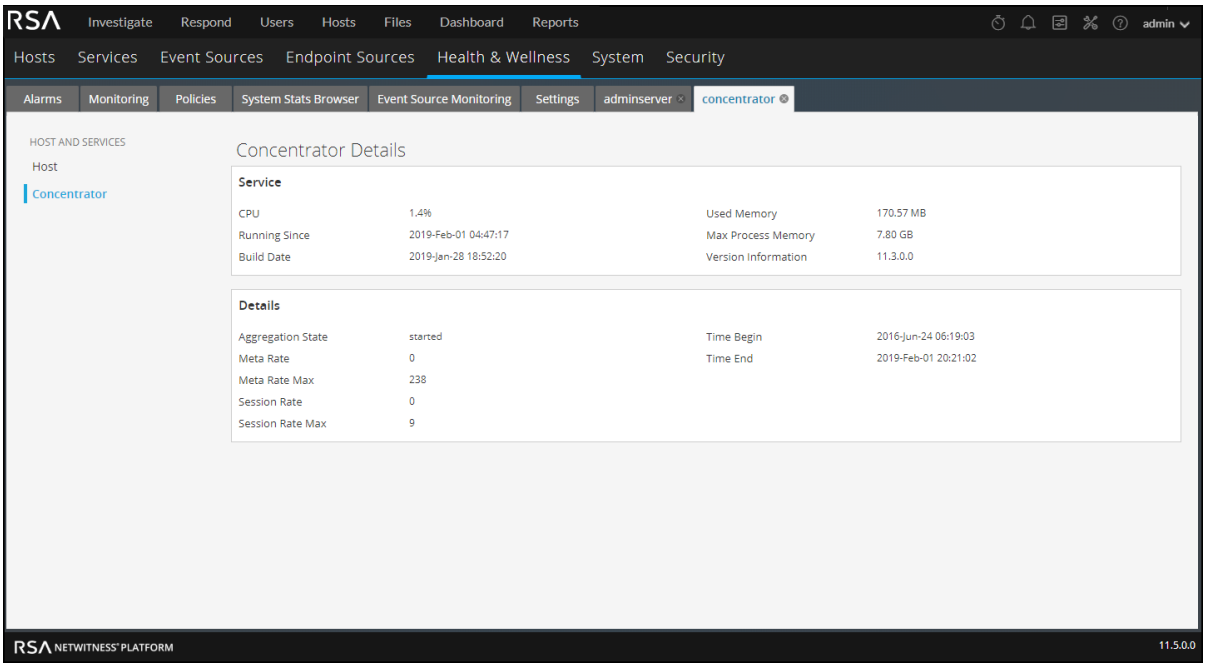
For the related procedure, see Monitor Service Details.

This section displays the current generic statistics for the service.

Statistic	Description
Aggregation State	State of data aggregation.
Meta Rate	Metadata objects per second rate.
Session Rate	Sessions per second rate.
Meta Rate Max	Maximum metadata objects per second rate.
Session Rate Max	Maximum sessions per second rate.

Concentrator Details View

The Concentrator Details view provides information about the Concentrator. The following figure shows the Concentrator details.



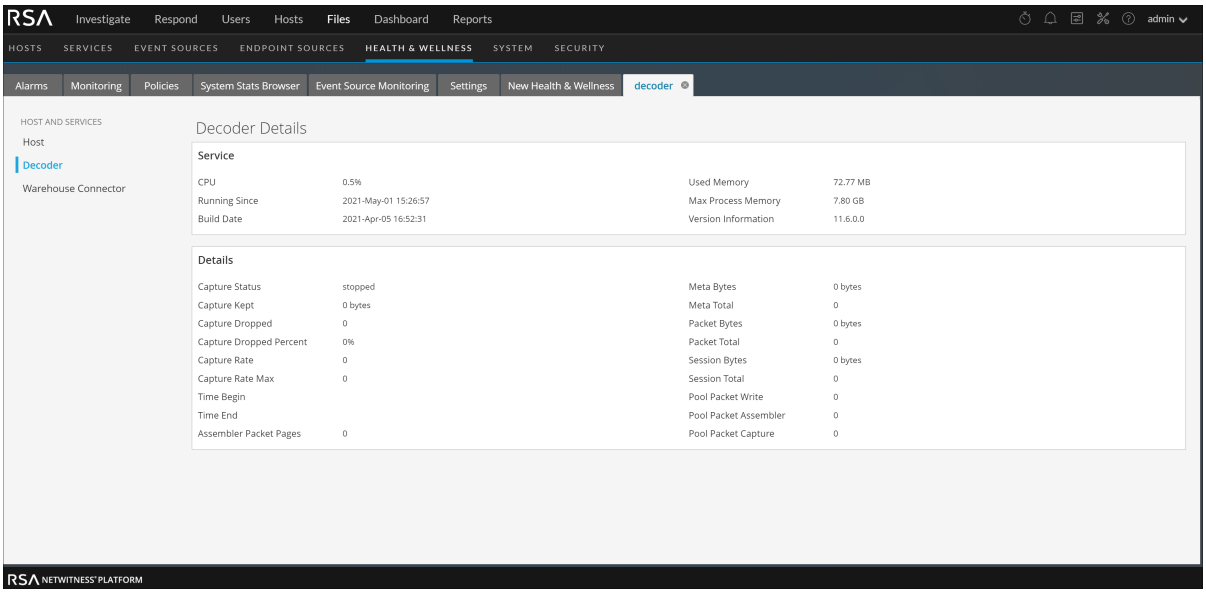
For the related procedure, see Monitor Service Details

The section displays the current generic statistics for the service.

Statistic	Description
Aggregation State	State of data aggregation.
Time Begin	Time (UTC) when the first session was tracked by the index.
Meta Rate	Metadata objects per second rate.
Time End	Time (UTC) when the last session was tracked by the index.
Meta Rate Max	Maximum metadata objects per second rate.
Session Rate	Sessions per second rate.
Session Rate Max	Maximum sessions per second rate.

Decoder Details View

The Decoder Details view provides information about the Decoder. The following figure shows the Decoder details.



For the related procedure, see Monitor Service Details.

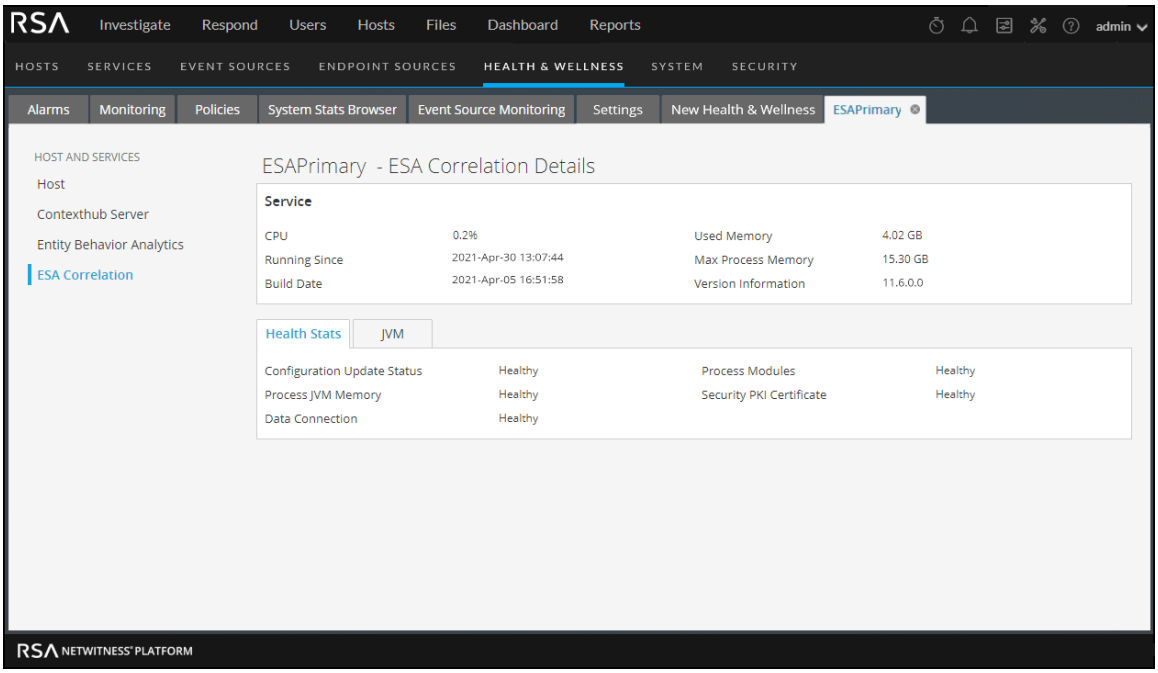
This section displays the current generic statistics for the service.

Statistic	Description
Capture Status	Status of data capture. Valid values are: • `starting` - Starting data capture (not capturing data yet). • `started` - Capturing data. • `stopping` - Stopping data capture (received request to stop data capture, but not have not stopped capturing data yet). • `stopped` - Not capturing data. • `disabled` - Not configured as a Decoder service.
Meta Bytes	Number of meta bytes in the database.
Capture Kept	Number of packets kept during capture.
Meta Total	Amount of metadata in the database.
Capture Dropped	Number of packets reported by the network card as dropped. After the service stops capturing data, the rate is reset to zero.
Packet Bytes	Number of packet bytes in the database.

Statistic	Description
Capture Dropped Percent	Packets reported by the network card as dropped as a percentage.
Packet Total	Number of packet objects held in the packet database. The total decreases when the database rolls files off due to size constraints. After the service stops capturing data, the number is not reset.
Capture Rate	Megabits per second rate at which the service is capturing data. Rate is a rolling average sample over a short time period (10 seconds). After the service stops capturing data, rate is reset to zero.
Session Bytes	Number of session bytes in the database.
Capture Rate Max	Maximum megabits per second rate at which the service is capturing data. Rate is a rolling average sample over a short time period (10 seconds). After the service stops capturing data, displays the maximum rate during data capture.
Session Total	Number of sessions held in the session database. This value shrinks when the database rolls files off due to size constraints. After the service stops capturing data, the number is not reset.
Time Begin	Time when first packet was captured (time when the first packet was stored in the packet database). This time increases as packets are rolled out of the packet database.
Pool Packet Write	Number of packet pages currently in the PCS pipeline that need to be written to the database.
Time End	Time when the last packet was captured (time when packet was written to the database). The time increases as new packets are captured.
Pool Packet Assembler	Number of pool packet pages waiting to be assembled.
Assembler Packet Pages	Number of packet pages waiting to be assembled.
Pool Packet Capture	Number of packet pages available for capture.

ESA Correlation Details View

The ESA Correlation Details view provides information for the ESA Correlation service. The following figure shows the ESA Correlation service details.



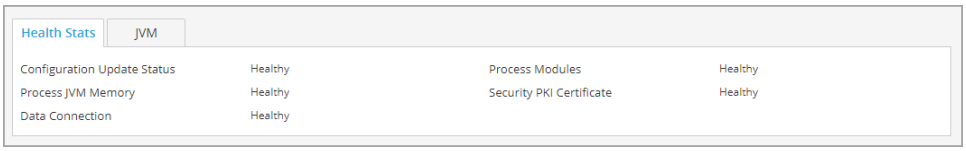
For the related procedure, see Monitor Service Details.

Many services, including the ESA Correlation service, have Health Stats and Java Virtual Machine (JVM) tabs. The Health Stats tab provides information about the health status of the service. The JVM tab shows the total memory used by the selected service and the total memory capacity of the host.

For more information on the ESA Correlation service and ESA Rule memory usage, see the *Alerting with ESA Correlation Rules User Guide*.

Health Stats Tab

The Health Stats tab provides information about the health status of the selected service.



The services on this tab can show one of three states:

- `Healthy`: The service is healthy.
- `Unhealthy`: The service is mostly functional, but it needs attention to mitigate potential down time.
- `Fatal`: Action needs to be taken to restore the service.

Statistic	Description
Configuration Update Status	Indicates whether the service requires a restart for configuration changes to take effect. If the Configuration Update Status shows as Unhealthy, restart the service.
Process JVM Memory	Indicates the memory usage status. An Unhealthy status occurs when heap memory usage is greater than or equal to 80%. A Fatal status occurs when heap memory usage is greater than or equal to 95%. If the service is using too much memory, you can add more memory or move services to other hosts. For more details on memory usage, go to the JVM tab.
Data Connection	Indicates the health of the database connection of the service to MongoDB.
Process Modules	Indicates the health of the service. If a service is starting up, it shows as Unhealthy since its health is not yet determined. A service is Healthy if it is up and running properly. A service shows as Fatal if it is running in upgrade mode or if the service is running in safe or degraded mode.
Security PKI Certificate	Indicates the service certificate health. It shows as Healthy if a given X509 certificate is self-signed.

JVM Tab

The JVM tab shows the total memory used by the selected service and the total memory capacity of the host.

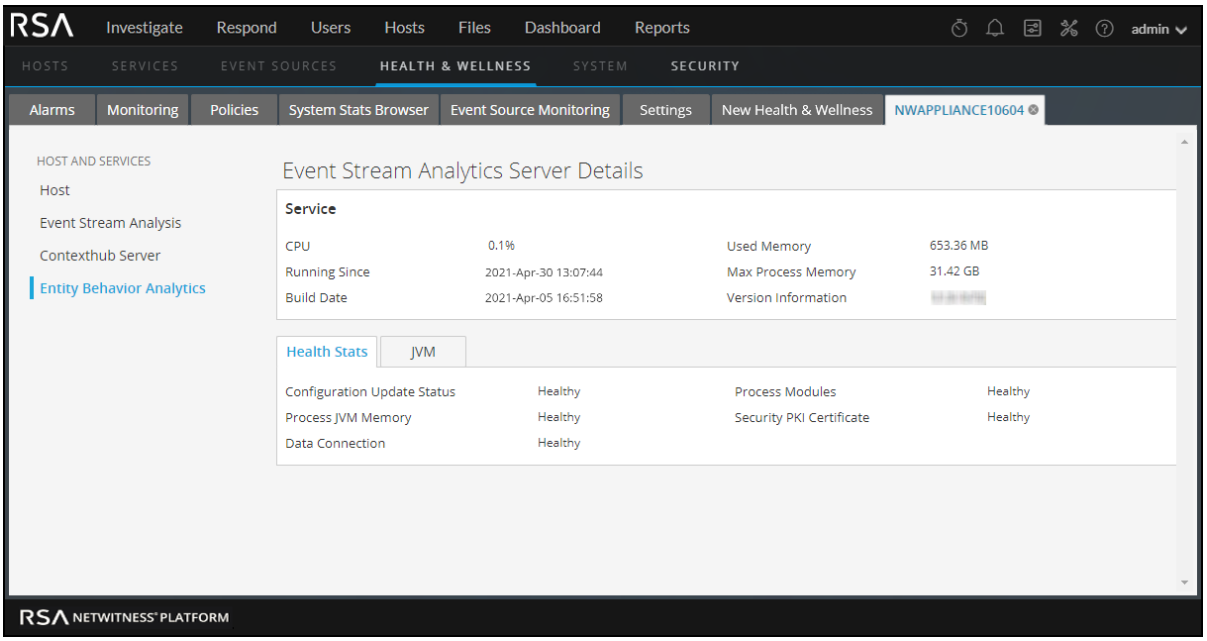


Statistic	Description
JVM Total Memory Max	Shows the total memory capacity for the entire host.
JVM Total Memory Used	Shows the total memory used by all services and processes on the host.

ESA Analytics Details View

The ESA Analytics Details view provides health status information about the selected ESA Analytics service. ESA Analytics services process the data for automated threat detection. It is important that you address any item that shows a status other than healthy, so that data processing is not interrupted and critical events are not missed.

The following figure shows the ESA Analytics Details view.



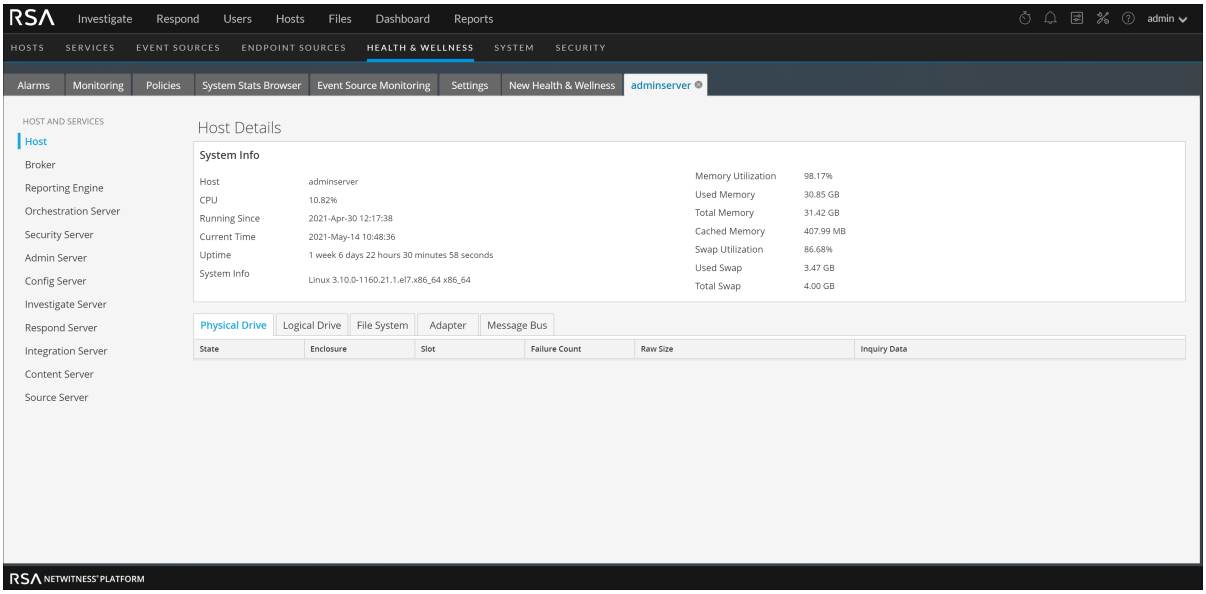
For the related procedure, see Monitor Service Details.

Many services, including the ESA Analytics service, have **Health Stats** and Java Virtual Machine (**JVM**) tabs. The **Health Stats** tab provides information about the health status of the service. The **JVM tab** shows the total memory used by the selected service and the total memory capacity of the host. For more information, see Health Stats Tab and JVM Tab.

For more information on ESA Analytics, see the *Automated Threat Detection Guide* and the *ESA Configuration Guide*.

Host Details View

The Host Details view provides information about a host, as shown in the following figure.



The options panel on the left displays the host and the services installed on the host. You can click on a host or service to view the statistics and other pertinent information for that host or service.

The Details panel displays information that is specific to the host and provides additional information regarding the hardware of the host.

For the related procedure, see Monitor Service Details

The top section displays the current performance, capacity, and historical statistics for the host.

Parameter	Description
Host	Hostname.
CPU	Current CPU usage of the host.
Running Since	Time when the host was started.
Current Time	Current time on the host
Uptime	Time for which the host has been active.
System Info	OS version installed on the host.
Memory Utilization	Percentage of memory utilized by the host.
Used Memory	Memory used in GB.
Total Memory	Capacity of the memory installed on the system.
Cached Memory	Memory that is cached to disk in GB.
Swap Utilization	Percentage of system swap in use.

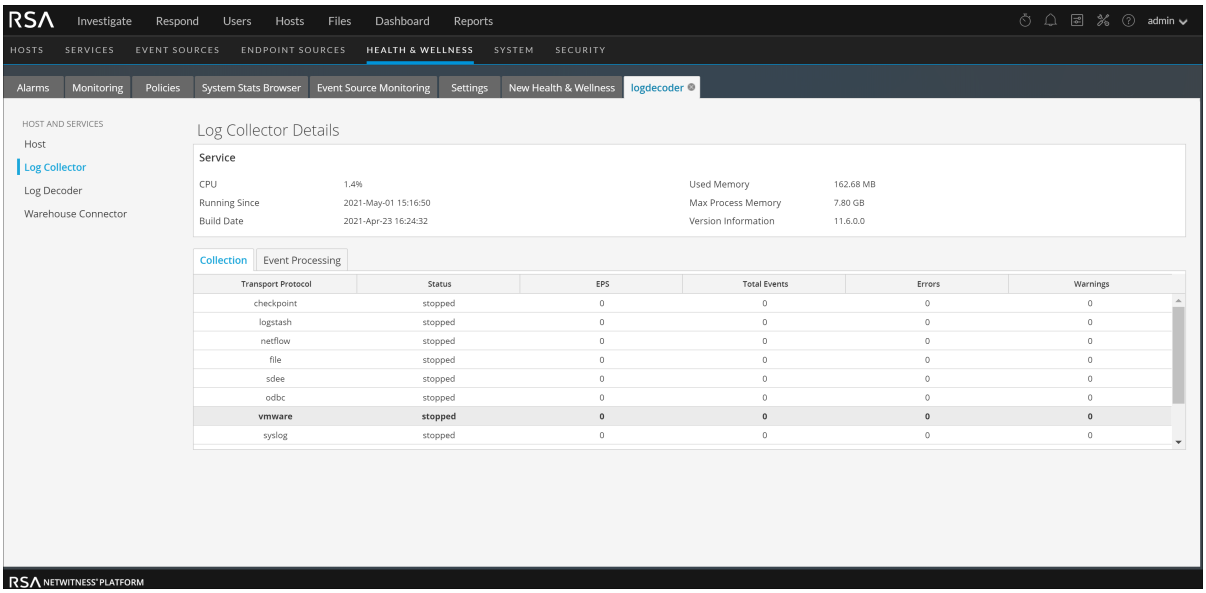
Parameter	Description
Used Swap	Swap used in GB.
Total Swap	Capacity of the swap installed on the system.

The lower section displays the current generic statistics for the host in the tabs described in the following table.

Tab	Description
Physical Drive	Type of physical drive, its usage and additional information of the physical drive on the host.
Logical Drive	Logical drive on the host.
File System	File system information, the size, current usage, and available capacity on the host.
Adapter	Adapter used on the host.
Message Bus	Publish In Rate - rate at which incoming messages are published to the message bus queue. Total Messages Queued - number of messages in the message queue. Memory Used - amount of memory used by the message bus (in bytes). Disk Free - free disk space available for the message bus (in bytes). Memory Limit - system memory limit. If the memory usage exceeds this value, this trips the Memory Alarm and NetWitness stops accepting messages. Disk Free Limit - limit of free disk space available for the message bus. If the available disk space falls below this value, this trips the Disk Free Alarm and NetWitness stops accepting messages. Memory Limit Available - Amount of memory available to this message broker (in bytes) before the Memory Used Alarm is tripped. Disk Limit Available - Amount of free disk space available to this message broker (in bytes) before the Disk Free Limit alarm is tripped. Disk Free Alarm - True or False. True indicates that the available disk space is below the value set in Disk Free Limit and NetWitness has stopped accepting messages. Memory Alarm - True or False. True indicates that the available memory is below the value set in Memory Limit and NetWitness has stopped accepting messages.

Log Collector Details View

The Log Collector Details view provides information for the Log Collector. The following figure depicts the Log Collector Details.



For the related procedure, see Monitor Service Details.

The lower section consists of the Collection and Event Processing tabs that display generic statistics for the service.

Collection Tab

Displays the event collection statistics for each Log Collection protocol you have implemented in NetWitness (see "Log Collection Getting Started Guide" in the Log Collection Guides).

Event Processing Tab

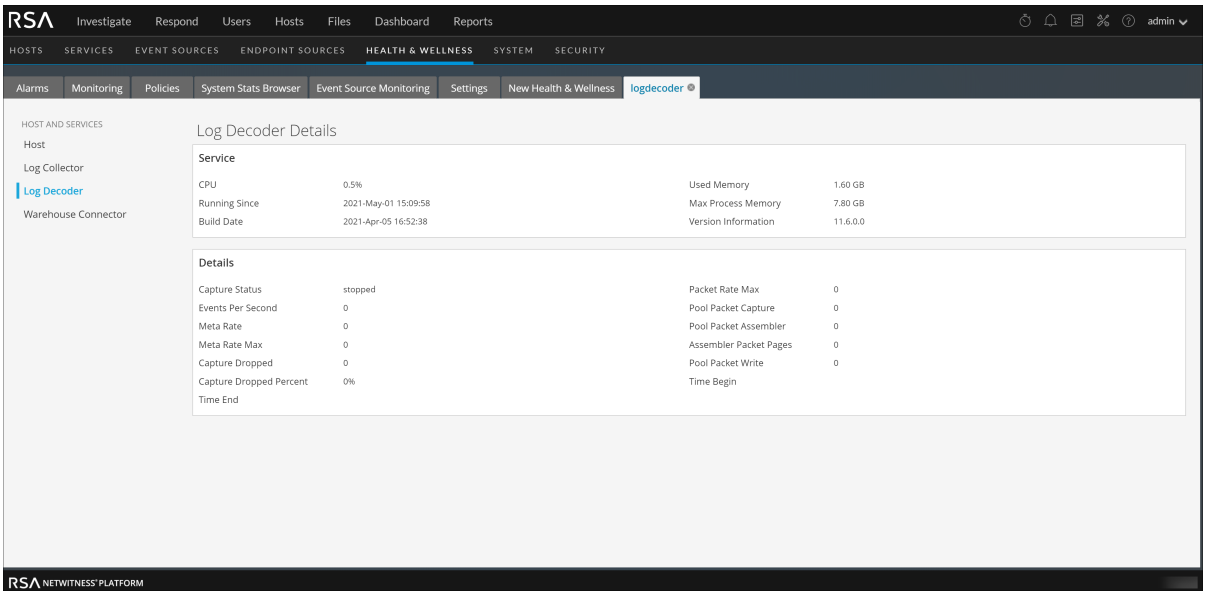
Displays statistics for the NetWitness internal event processing protocol (that is, the Log Decoder) for Log Collection.

Parameter	Description
Transport Protocol	NetWitness protocol use for Log Collections (that is, the Log Decoder).

Parameter	Description
Status	Status of the Log Decoder. Valid values are: • `starting` - Starting data capture (not capturing data yet). • `started` - Capturing data. • `stopping` - Stopping data capture (received request to stop data capture, but not have not stopped capturing data yet). • `stopped` - Not capturing data. • `disabled` - Not configured as a Decoder service.
EPS	Rate (events per second) at which this the Log Decoder is processing events from the Log Collector.
Total Events	Total events processed by the Log Decoder.
Errors	Number of errors encountered.
Warnings	Number of warnings encountered.
Byte Rate	Current throughput in bytes per second.

Log Decoder Details View

The Log Decoder Details view provides information for the Log Decoder. The following figure shows the Log Decoder Details.



For the related procedure, see Monitor Service Details.

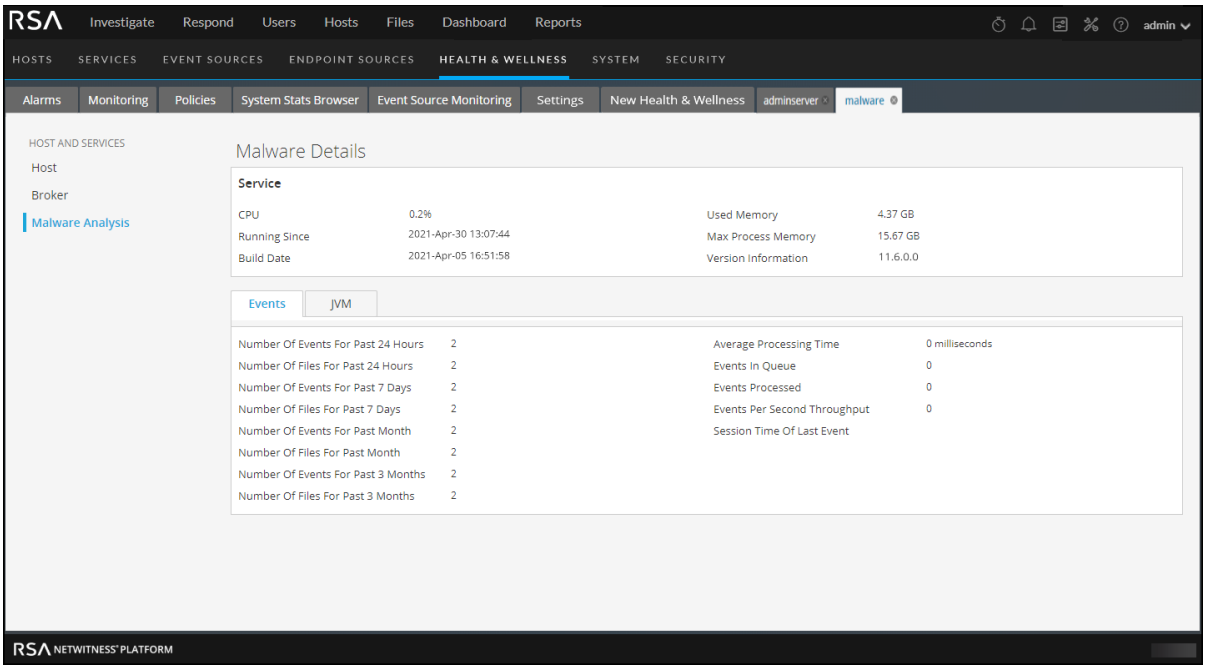
This section displays the current generic statistics for the service.

Statistic	Description
Capture Status	Status of data capture. Valid values are: • `starting` - Starting data capture (not capturing data yet). • `started` - Capturing data. • `stopping` - Stopping data capture (received request to stop data capture, but not have not stopped capturing data yet). • `stopped` - Not capturing data. • `disabled` - Not configured as a Log Decoder service.
Packet Rate Max	Maximum per second rate at which the service is writing packets to the database. Rate is a rolling average sample over a short time period (10 seconds). After the service stops capturing data, displays the maximum rate during data capture.
Events Per Second	Rate (events per second) at which the Log Decoder is processing events from the Log Collector.
Pool Packet Capture	Number of packet pages available for capture.

Statistic	Description
Meta Rate	Per second rate at which the service is writing metadata objects to the database. Rate is a rolling average sample over a short time period (10 seconds). After the service stops capturing data, rate is reset to zero.
Pool Packet Assembler	Number of packet pages waiting to be assembled.
Meta Rate Max	Maximum per second rate at which the service is writing metadata objects to the database. Rate is a rolling average sample over a short time period (10 seconds). After the service stops capturing data, displays the maximum rate reached during data capture.
Assembler Packet Pages	Number of packet pages waiting to be assembled.
Capture Dropped	Number of packets reported by the network card as dropped. After the service stops capturing data, rate is reset to zero.
Pool Packet Write	Number of packet pages in the PCS pipeline that need to be written to the database.
Capture Dropped Percent	Packets reported by the network card as dropped as a percentage.
Time Begin	Time when first packet was captured (time when the first packet was stored in the packet database). This time increases as packets are rolled out of the packet database.
Time End	Time when the last packet was captured (time when packet was written to the database). The time increases as new packets are captured.

Malware Details View

The Malware Details view provides information for Malware Analysis. The following figure shows the Malware Details.



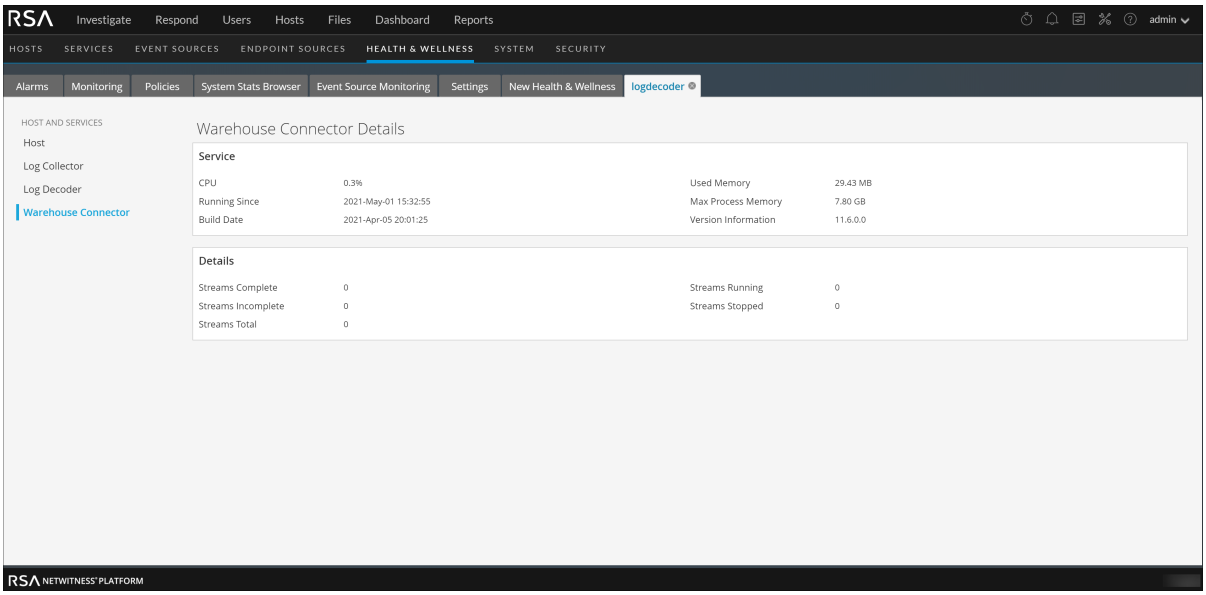
For the related procedure, see Monitor Service Details.

Displays the following event-related statistical information for the Malware Analysis service.

- Number of events for the past 24 hours
- Average processing time
- Number of files for the past 24 hours
- Events in queue
- Number of events for the past 7 days
- Events processed
- Number of events for the past 7 days
- Events per second throughput
- Number of events for the past month
- Session time of the last event
- Number of files for the past month
- Number of events for the past 3 months
- Number of files for the past 3 months

Warehouse Connector Details View

The Warehouse Connector Details tab provides information for the Warehouse Connector, such as the date it was built, CPU, and version information. The following figure shows the Warehouse Connector Details.



For the related procedure, see Monitor Service Details.

Policies View

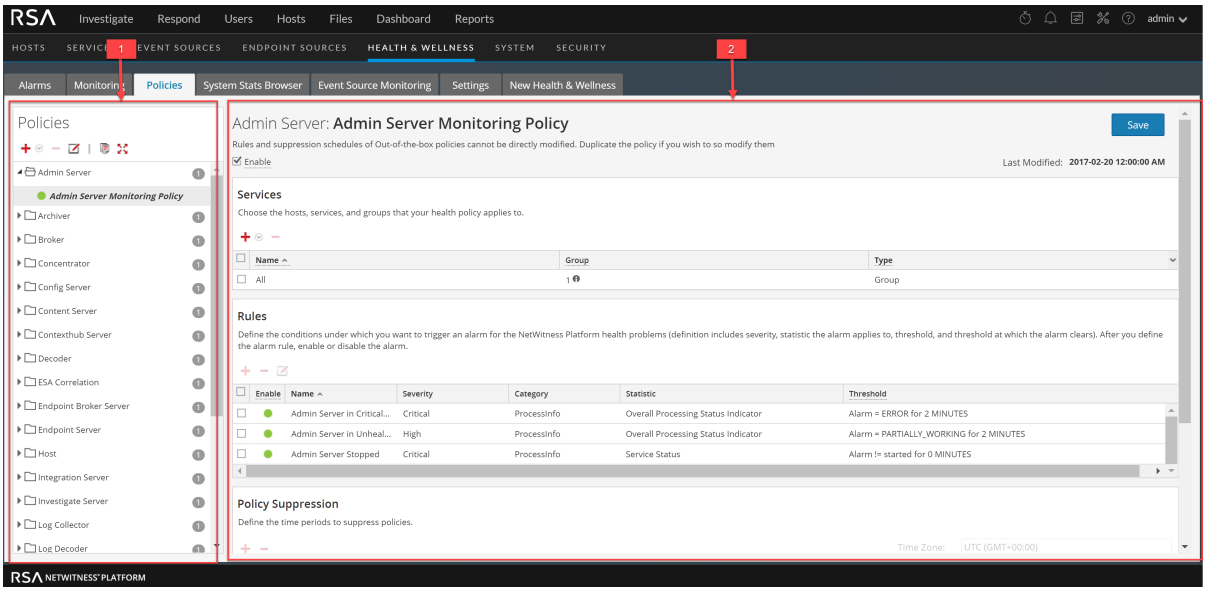
The required permission to access this view is **Manage services**.

What do you want to do?

Role	I want to ...	Show me how
Administrator	View the policies NetWitness Server and Services	Manage Policies
Administrator	Add, Edit, Duplicate, and Delete Policies	Manage Policies

Quick Look

The figure depicts the Policies view.



1 Policies Panel
2 Policy Detail Panel

1. Go to **(Admin) > Health & Wellness**.
2. Click the **Policies** tab.

Policies Panel

In the Policies panel, you can add or delete policies for hosts and services in this panel.

Feature	Description
	Displays available service types to create a new policy. Select one so that you can define a policy or policies for it.
	Deletes the selected policy from the Policies panel. You can only delete one policy at a time.
	Allows you to change the name of the policy.
	Creates a copy of the selected policy. For example, if you select **First Policy** and click , NetWitness creates a copy of this policy and names it First Policy (1).
	Expands the list of policies under the services and hosts in the Policies panel.
	Contracts the list of policies under the services and hosts in the Policies panel.
	List of: • Services and hosts for which you have defined policies. • NetWitness standard policies that you can apply to hosts and services.

Policy Detail Panel

The Policy Detail panel displays the policy selected from the Policies panel.

Feature	Description
Save	Saves any changes you made in this panel.
Policy Type	Displays the type of policy you selected.
Modified Date	Displays the last date this policy was modified.
 Enable	Enables or disables the policy.
Services	
	Displays menu in which you select: • **Groups** to display the Groups dialog from which you select service groups to this policy. • **Service/Host** to display the Services/Hosts dialog from which you select services to add to this policy. If the policy type is Host, the menu displays Host (and not Service). You can select services based on policy type.
	Deletes the selected service or group from this policy.
Rules	

Feature	Description
	Displays the Add Rule dialog in which you define a rule for this policy.
	Deletes the selected rule from this policy.
	Displays the Edit Rule dialog for the selected rule.
Policy Suppression	
	Adds a policy suppression timeframe row.
	Deletes the selected policy suppression timeframe row.
Time Zone	Selects the time zone for the Policy from the drop-down list. This time zone applies to both Policy Suppression and Rule Suppression.
	Selects the checkbox to select a policy suppression timeframe row.
Days	Days of the week that you want to suppress the policy according to the time range specified. Click on the day of the week that you want to suppress the policy. You can select any combination of days including all days.
Time Range	Time range during which the policy is suppressed for the days selected.
Notification	
	Adds an EMAIL notification row.
	Deletes the selected policy suppression timeframe row.
Notification Settings	Opens the Notification Servers view in which you can define the Email notification settings.
	Selects a policy suppression time frame row.
Output	The type of notification defined on the Global Notifications page. Can be email, SNMP, Syslog, or Script.
Recipient	The name of the person receiving the notification.
Notification Server	Selects the EMAIL notification server. See "Configure Notification Servers" in the *System Configuration Guide* for the source of the values in this drop-down list.
Template	Selects the Template for this EMAIL notification. NetWitness provides the Health & Wellness Default SMTP Template and the alarms template. See" Configure Notification Templates"in the *System Configuration Guide* for the source of the other values in this drop-down list. **Note:** Refer to Include the Default Email Subject Line if you want to include the default Email subject line from the Health & Wellness template in your Health & Wellness Email notifications for specified recipients.

Groups dialog

Feature	Description
Groups panel	
Name	Displays the service groups you have defined. You can select: • **All** to display all your services in the Services panel. • A group to display the services in comprise that group in the Services panel.
Services panel	
Name	Displays the name of the service.
Host	Displays the host on which the service is running.
Type	Displays the type of service.

Rules Dialog

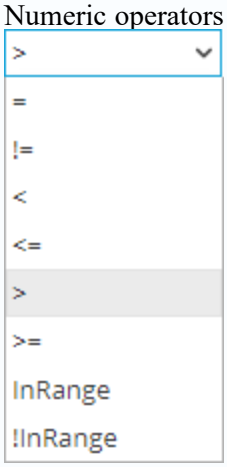
Feature	Description
☐ Enable	Enables or disables the rule for this policy.
Name	Describes the name of the rule.
Description	Describes the rule. Include the following information in this field. • Informational description - purpose of the rule and what problem it monitors. • Remediation - steps to take to resolve the condition that triggers the alarm for this rule.
Severity	Defines the severity of the rule. Valid values are: • Critical • High • Medium • Low

Feature	Description
Statistic	Defines the statistics you want to check with this rule. You can select: • Statistical category from the left drop-down list. • Statistic from the right drop-down list. **Note:** For Public Key Infrastructure (PKI) policy, select PKI in the category and statistics as any one of the following: - NetWitness Server PKI Certificate Expiration - Displays the time left before the certificate expires. - NetWitness Server PKI CRL Expiration - Displays the time left before the Certificate Revocation List (CRL) expires. - NetWitness Server PKI CRL Status - Displays the current status of the CRL. Refer to the System Stats Browser View for examples of the statistics you may want to check with a rule.
Alarm Threshold	Defines he threshold of the rule that triggers the policy alarm: • Amount **Note:** For CRL expiry the supported format is ddddhhmm, for example: - 10000 represents 1 day - 2359 represents 23 hours and 59 minutes - 10023 represents 1 day and 23 minutes - 3650100 represents 365 days and 1 hour • Time in minutes
Recovery	Defines when to clear the threshold of the rule: • Operator • Amount • Time in minutes
Rule Suppression	
	Adds a rule suppression timeframe row.
	Deletes the selected rule suppression time frame row.
	Selects a rule suppression time frame row.
Time Zone: *time-zone*	Displays the Policy time zone. You select the time zone for a policy in the Policy Suppression panel.
Days	Defines days of the week that you want to suppress the rule according to the time range specified. Click on the day of the week that you want to suppress the rule. You can select any combination of days including all days.
Time Range	Defines the time range during which the rule is suppressed for the days selected.

Threshold Operators

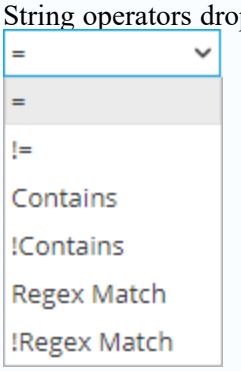
The **Alarm Threshold** and **Recovery Threshold** fields in the Rules dialog prompt you for either numeric or string operators based on the statistic criteria you specify.

Numeric operators drop-down menu:



=
!=
<
<=
>
>=
InRange
!InRange

String operators drop-down menu:

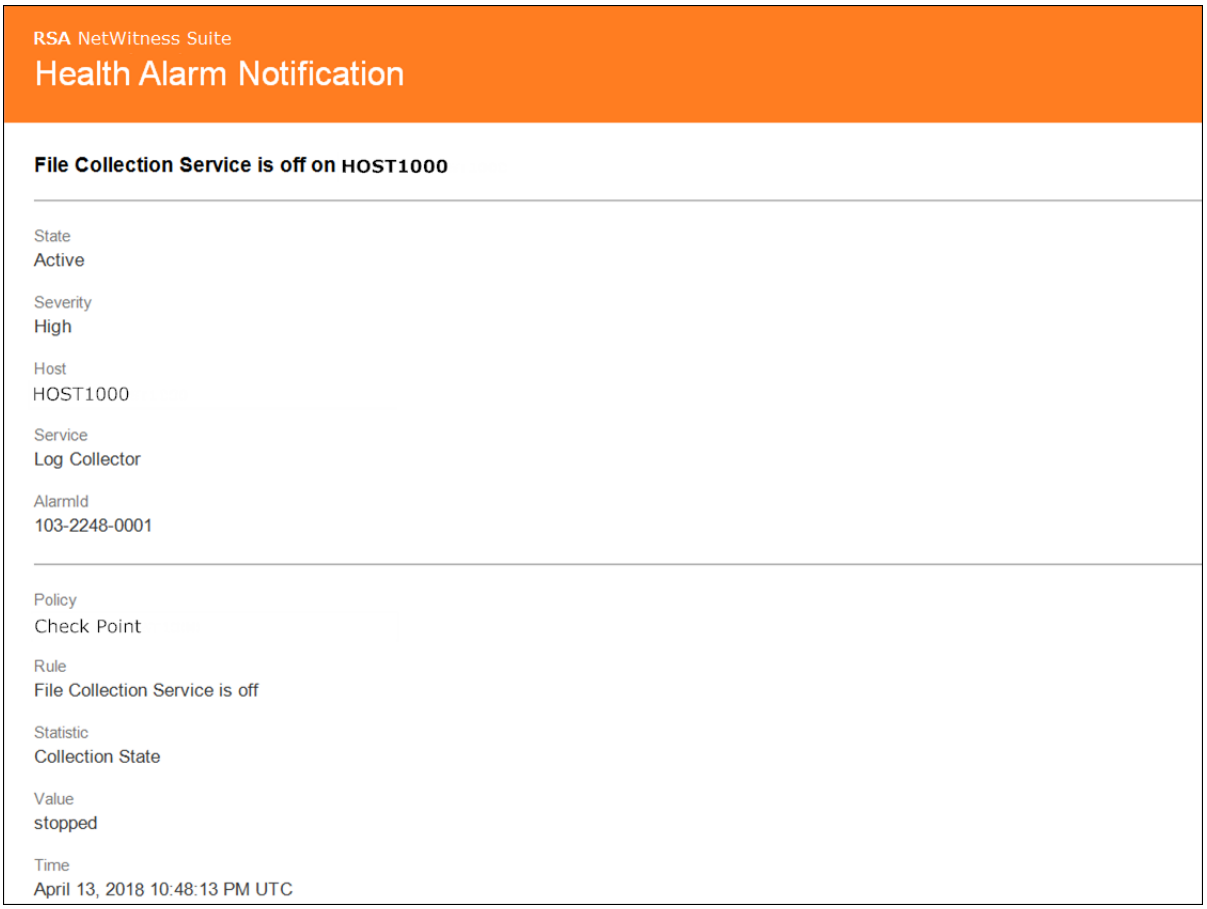


=
!=
Contains
!Contains
Regex Match
!Regex Match

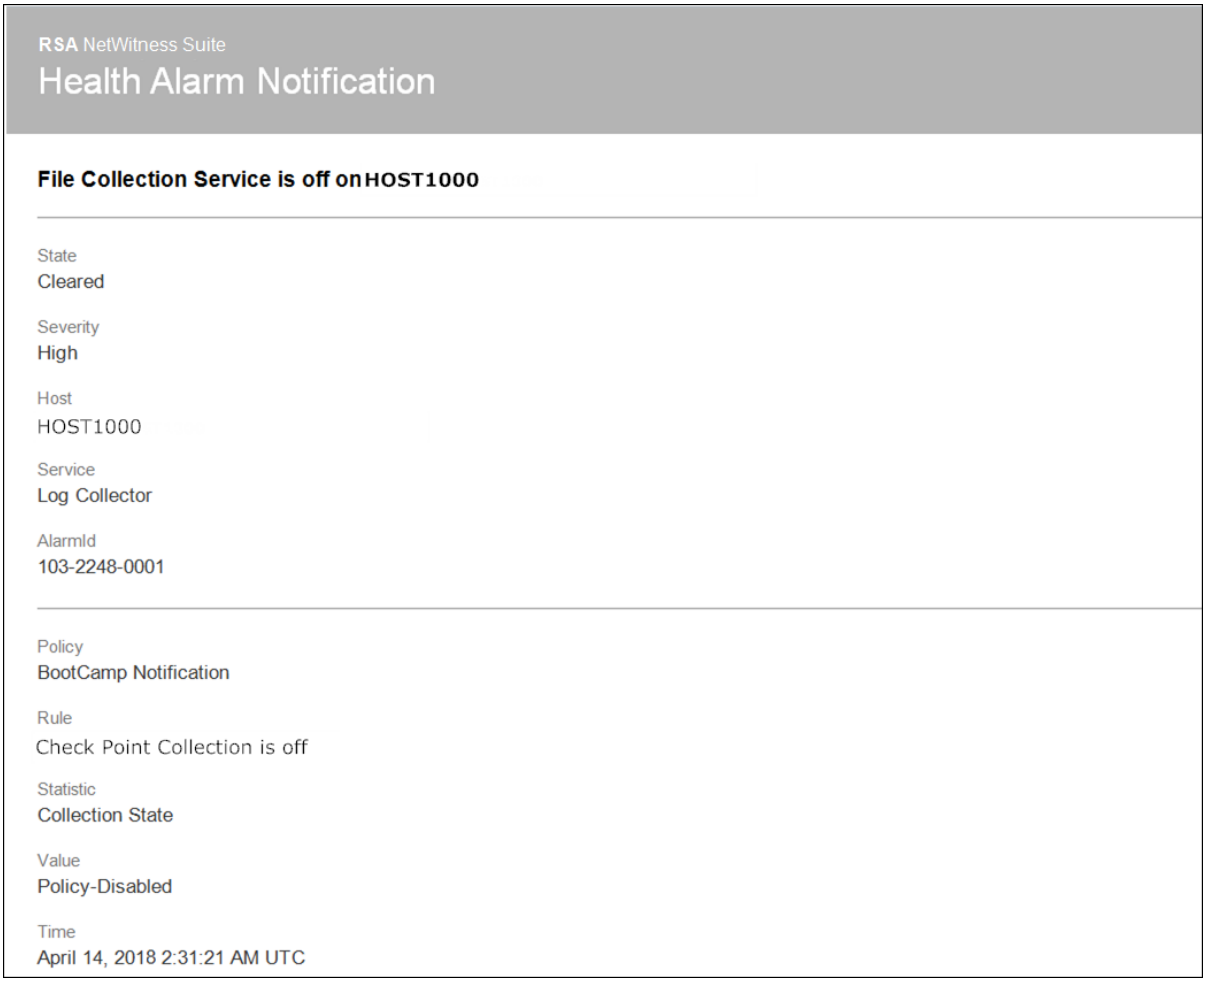
Health & Wellness Email Templates

Note: Please refer to Include the Default Email Subject Line if you want to include the default Email subject line from the Health & Wellness template in your Health & Wellness Email notifications for specified recipients.

Health & Wellness Default SMTP Template



Alarms Template



NetWitness Platform Out-of-the-Box Policies

The following table lists the NetWitness Out-of-the-Box Policies with the rules defined for each policy.

You can perform the following tasks on any of these policies:

- Change service and group assignments.
- Disable or enable policies.

You cannot perform the following tasks on any of these policies:

- Delete them.
- Edit Policy names.

Note: Additional information about the Out-of-the-Box Policies can be found in the User Interface under
Health & Wellness > Policies.

Policy Name	Rule Name	Alarm Triggered
	Communication Failure Between Master NetWitness Server Host and a Remote Host	Host is down, Network is down, Message Broker is Down, or Invalid or missing security certificates for 10 minutes or more.

Policy Name	Rule Name	Alarm Triggered
NetWitness Server Monitoring Policy	Critical Usage on Rabbitmq Message Broker Filesystem	For `var/lib/rabbitmq`, Mounted Filesystem Disk Usage goes over 75%.
	Filesystem is Full	Overall Mounted Filesystem Disk Usage reaches 100%.
	High Filesystem Usage	Overall Mounted Filesystem Disk Usage goes over 95%.
	High System Swap Utilization	Swap Utilization goes under 5 % for 5 minutes or more.
	High Usage on Rabbitmq Message Broker Filesystem	Mounted Filesystem Disk Usage for `var/lib/rabbitmq` goes over 60%.
	Host Unreachable	Host down.
	LogCollector Event Processor Exchange Bindings Status	Issue with Log Collection Message Broker Queues for 10 minutes or more.
	LogCollector Event Processor Queue with No Bindings	Issue with Log Collection Message Broker Queues for 10 minutes or more.
	LogCollector Event Processor Queue with No Consumers	Issue with Log Collection Message Broker Queues for 10 minutes or more.
	Power Supply Failure	Host not receiving power.
	RAID Logical Drive Degraded	For Raid Logical Drive, Drive State equals Degraded or Partially Degraded.
	RAID Logical Drive Failed	For Raid Logical Drive, Logical Drive State equals Offline, Failed, or Unknown.
	RAID Logical Drive Rebuilding	For Raid Logical Drive, Logical Drive State equals Rebuild.
	RAID Physical Drive Failed	For Raid Physical Drive, Physical Drive State does not equal Online, Online Spun Up, or Hotspare.
	RAID Physical Drive Failure Predicted	For Raid Physical Drive, Physical Drive Predictive Failure Count is greater than 1.
	RAID Physical Drive Rebuilding	For Raid Physical Drive, Physical Drive State equals Rebuild.
	RAID Physical Drive Unconfigured	For Raid Physical Drive, Physical Drive State contains Unconfigured (good).
	SD Card Failure	SD Card Status does not equal ok.

Policy Name	Rule Name	Alarm Triggered
NetWitness Archiver Monitoring Policy	Archiver Aggregation Stopped	Archiver Status does not equal started.
	Archiver Database(s) Not Open	Database Status does not equal opened.
	Archiver Not Consuming From Service	Devices Status does not equal consuming.
	Archiver Service in Bad State	Service State does not equal started or ready.
	Archiver Service Stopped	Server Status does not equal started.
NetWitness Broker Monitoring Policy	Broker >5 Pending Queries	Queries Pending greater than or equal to 5 for 10 minutes or more.
	Broker Aggregation Stopped	Broker Status does not equal started.
	Broker Not Consuming From Service	Devices Status does not equal consuming.
	Broker Service in Bad State	Service State does not equal started or ready.
	Broker Service Stopped	Server Status does not equal started.
	Broker Session Rate Zero	Session Rate (current) equals 0 for 2 minutes or more.
NetWitness Concentrator Monitoring Policy	Concentrator >5 Pending Queries	Queries Pending greater than or equal to 5 for 10 minutes or more.
	Concentrator Aggregation Behind >100K Sessions	Devices Sessions Behind is greater than or equal to 100000 for 1 minute or more.
	Concentrator Aggregation Behind >1M Sessions	Devices Sessions Behind is greater than or equal to 1000000 for 1 minute or more.
	Concentrator Aggregation Behind >50M Sessions	Devices Sessions Behind is greater than or equal to 50000000 for 1 minute or more.
	Concentrator Aggregation Stopped	Broker Status does not equal started.
	Concentrator Database(s) Not Open	Database Status does not equal opened.
	Concentrator Meta Rate Zero	Concentrator Meta Rate (current) equals 0 for 2 minutes or more.
	Concentrator Not Consuming From Service	Devices Status does not equal consuming.
	Concentrator Service in Bad State	Service State does not equal started or ready.
	Concentrator Service Stopped	Server Status does not equal started.

Policy Name	Rule Name	Alarm Triggered
NetWitness Decoder Monitoring Policy	Decoder Capture Not Started	Capture Status does not equal started.
	Decoder Capture Rate Zero	Capture Rate (current) equals 0 for 2 minutes or more.
	Decoder Database Not Open	Database Status does not equal opened.
	Decoder Dropping >1% of Packets	Capture Packets Percent Dropped (current) is greater than or equal to 1%.
	Decoder Dropping >10% of Packets	Capture Packets Percent Dropped (current) is greater than or equal to 10%.
	Decoder Dropping >5% of Packets	Capture Packets Percent Dropped (current) is greater than or equal to 5%.
	Decoder Packet Capture Pool Depleted	Packet Capture Queue equals 0 for 2 minutes or more.
	Decoder Service in Bad State	Service State does not equal started or ready.
	Decoder Service Stopped	Server Status does not equal started.
NetWitness Event Steam Analysis Monitoring Policy	ESA Overall Memory Utilization > 85%	Total ESA Memory Usage % is greater than or equal to 85 %.
	ESA Overall Memory Utilization > 95%	Total ESA Memory Usage % is greater than or equal to 95 %.
	ESA Service Stopped	Server Status does not equal started.
	ESA Trial Rules Disabled	Trial Rules Status does not equal enabled.
NetWitness IPDB Extractor Monitoring Policy	IPDB Extractor Service in Bad State	Service State does not equal started or ready.
	IPDB Extractor Service Stopped	Server Status does not equal started.
NetWitness Incident Management Monitoring Policy	Incident Management Service Stopped	Server Status does not equal started.

Policy Name	Rule Name	Alarm Triggered
NetWitness Log Collector Monitoring Policy	Log Collector Service Stopped	Server Status does not equal started.
	Log Decoder Event Queue > 50% Full	Number of events currently in the queue is using 50% or more of the queue.
	Log Decoder Event Queue > 80% Full	Number of events currently in the queue is using 80% or more of the queue.
	Log Collector Service in Bad State	Service State does not equal started or ready.
NetWitness Log Decoder Monitoring Policy	Decoder Dropping>10% of Packets	Capture Packets Percent Dropped (current) is greater than or equal to 10%
	Log Capture Not Started	Capture Status does not equal started.
	Log Decoder Capture Rate Zero	Capture Rate (current) equals 0 for 2 minutes or more.
	Log Decoder Database Not Open	Database Status does not equal opened.
	Log Decoder Dropping >1% of Logs	Capture Packets Percent Dropped (current) is greater than or equal to 1%.
	Log Decoder Dropping >5% of Logs	Capture Packets Percent Dropped (current) is greater than or equal to 5%.
	Log Decoder Packet Capture Pool Depleted	Packet Capture Queue equals 0 for 2 minutes or more.
	Log Decoder Service Stopped	Server Status does not equal started.
	Log Decoder Service in Bad State	Service State does not equal started or ready.
NetWitness Malware Analysis Monitoring Policy	Malware Analysis Service Stopped	Server Status does not equal started.

Policy Name	Rule Name	Alarm Triggered
NetWitness Reporting Engine Monitoring Policy	Reporting Engine Alerts Critical Utilization	Alerts Utilization is greater than or equal to 10 for 5 minutes or more.
	Reporting Engine Available Disk <10%	Available disk space is less than 10%.
	Reporting Engine Available Disk <5%	Available disk space is less than or equal to 5%.
	Reporting Engine Charts Critical Utilization	Charts Utilization is greater than or equal to 10 for 5 minutes or more.
	Reporting Engine Rules Critical Utilization	Rules Utilization is greater than or equal to 10 for 5 minutes or more.
	Reporting Engine Schedule Task Pool Critical Utilization	Schedule Task Pool Utilization is greater than or equal to 10 for 15 minutes or more.
	Reporting Engine Service Stopped	Server Status does not equal started.
	Reporting Engine Shared Task Critical Utilization	Shared Task Pool Utilization is greater than or equal to 10 for 5 minutes or more.
NetWitness Warehouse Connector Monitoring Policy	Warehouse Connector Service in Bad State	Service State does not equal started or ready.
	Warehouse Connector Service Stopped	Server Status does not equal started.
	Warehouse Connector Stream Behind	Stream Behind is greater than or equal to 2000000.
	Warehouse Connector Stream Disk Utilization > 75%	Stream Disk Usage (Pending Destination Load) is greater than or equal to 75.
	Warehouse Connector Stream in Bad State	Stream Status does not equal consuming or online for 10 minutes r more.
	Warehouse Connector Stream Permanently Rejected Files > 300	Number of files in the permanently rejected files is greater than or equal to 300.
	Warehouse Connector Stream Permanently Rejected Folder > 75% Full	Rejected folder usage is greater than or equal to 75%.
NetWitness Workbench Monitoring Policy	Workbench Service in Bad State	Service State does not equal started or ready.
	Workbench Service Stopped	Server Status does not equal started.

System Stats Browser View

NetWitness provides a way to monitor the status and operations of hosts and services. The System Stats Browser tab displays key statistics, service system information, and host system information for a host or service.

You can customize the stats view depending on the parameter you select to filter the data.

To access the System Stats Browser view:

1. Go to  **(Admin) > Health & Wellness**.

   The Health & Wellness view is displayed with the Alarms tab open.
2. Click the **System Stats Browser** tab.

What do you want to do?

Role	I want to ...	Show me how
Administrator	View the System Stat Historical Graph	Historical Graph for System Stats

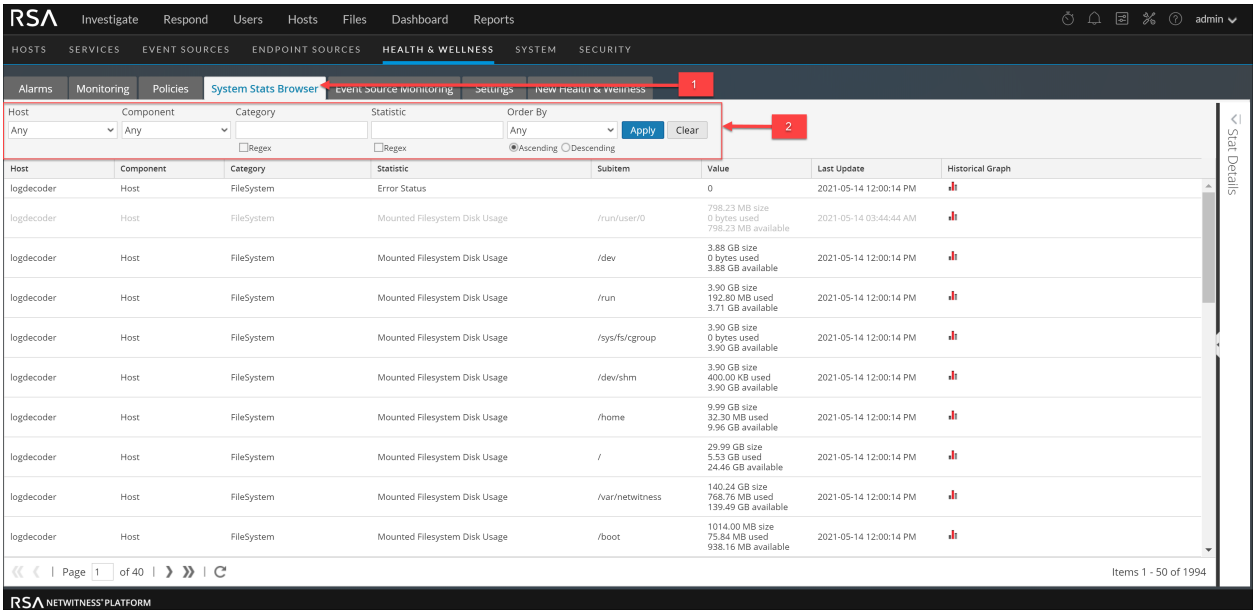
Related Topics

Monitor Service Statistics

Filter System Statistics

Quick Look

The System Stats Browser view is displayed.



1 Displays System Stats Browser View
2 Toolbar used to filter and customize the System Stats Browser View

Note: Historical graphs are enabled, and can be displayed, for statistics with numeric values. However, historical graphs are disabled for statistics with string values, for example, Health checks (Healthy), and are displayed as gray in the UI.

Filters

This table lists the various parameters you can use to filter and customize the System Stats view.

Parameter	Description
Host	Select a host from the drop-down menu to display the stats of the selected host. Select **Any** to list all the available hosts.
Component	Select a component from the drop-down menu to display the stats for the selected component. Select **Any** to list out all the components on a selected host.
Category	Type the category to display the stats for the required category. Select Regex to enable Regex filter. It performs a regular expression search against text and lists out the specified category. If Regex is not selected it supports globbing pattern matching.
Statistic	Type the statistic to display the required statistic on all the hosts or components. Select Regex to enable Regex filter. This performs a regular expression search against text and lists out the specified category. If Regex is not selected it supports globbing pattern matching.
Order By	Select the order in which the list needs to be filtered. Select Ascending to filter the list it in an ascending order.

Commands

Command	Action
Apply	Click to apply the filters chosen and display the list accordingly.
Clear	Click to clear the chosen filters.

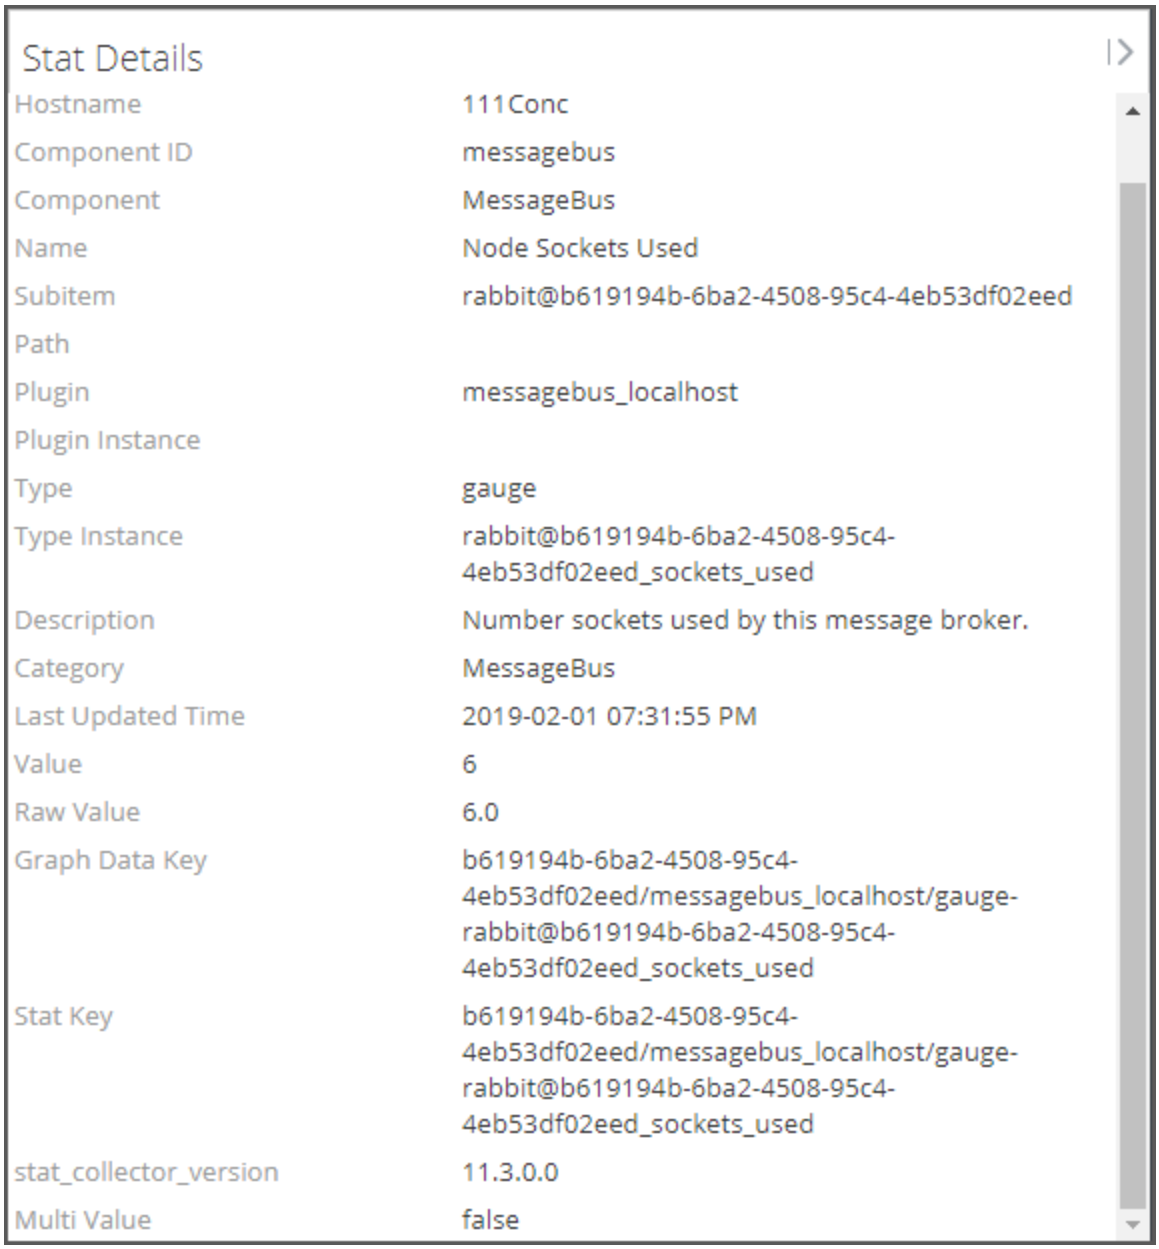
System Stats View Display

Displays statistics, service system information, and host system information for a host or service.

Access Stats Details

Select one of the stats and click **Stats Details** on the right hand side of the panel.

The Stats details panel opens with details of the selected stats.


Stat Details
Hostname
111Conc
Component ID
messagebus
Component
MessageBus
Name
Node Sockets Used
Subitem
rabbit@b619194b-6ba2-4508-95c4-4eb53df02eed
Path
Plugin
messagebus_localhost
Plugin Instance
Type
gauge
Type Instance
rabbit@b619194b-6ba2-4508-95c4-4eb53df02eed_sockets_used
Description
Number sockets used by this message broker.
Category
MessageBus
Last Updated Time
2019-02-01 07:31:55 PM
Value
6
Raw Value
6.0
Graph Data Key
b619194b-6ba2-4508-95c4-4eb53df02eed/messagebus_localhost/gauge-rabbit@b619194b-6ba2-4508-95c4-4eb53df02eed_sockets_used
Stat Key
b619194b-6ba2-4508-95c4-4eb53df02eed/messagebus_localhost/gauge-rabbit@b619194b-6ba2-4508-95c4-4eb53df02eed_sockets_used
stat_collector_version
11.3.0.0
Multi Value
false

New Health & Wellness Tab

In New Health & Wellness tab, you can filter alerts (monitors) and add the notifications which sends a notification message when an alert is triggered, suppress the notification for a time period and pivot to New Health and Wellness dashboards.

To access this view:

1. Go to  **(Admin) > Health & Wellness**.
2. Click the **New Health & Wellness** tab.

The figure depicts the New Health & Wellness view.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Manage Notification	Manage Notification and Suppression
Administrator	Manage Suppression	Manage Notification and Suppression

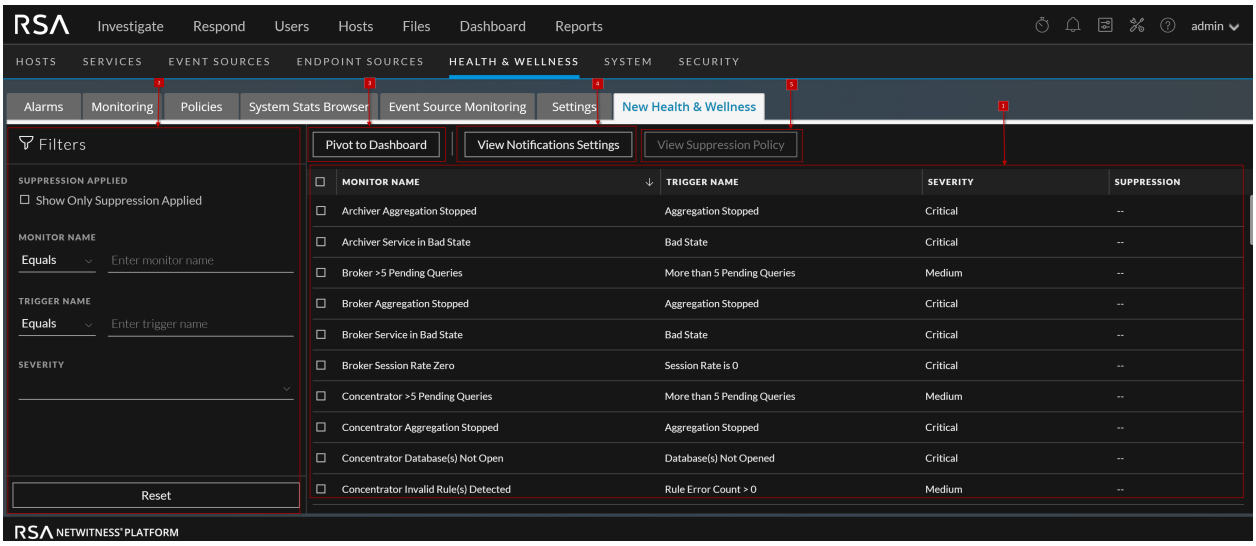
Related Topics

Monitoring through Dashboards

Monitoring through Alerts

Managing Notifications

Quick Look



1 **Notification Panel**: You can view the list of alerts (monitors) along with the trigger name, severity and suppression.

Monitor Name - Name of the monitor.

Trigger Name - Name of the trigger.

Severity - Severity assigned to the alert. The options are Critical, High, Medium, and Low.

Suppression - Indicates whether suppression policy is applied or not. Green tick indicates suppression policy is applied.

2 **Filters Alerts**. You can filter alerts on Suppression Applied, Monitor Name, Trigger Name, and Severity. To reset filters, click **Reset**.

3 **Pivot to Dashboard**. You can view New Health and Wellness Dashboards. For more information, see Accessing New Health and Wellness Dashboards.

4 **View Notification Settings**. You can define notification settings which sends a notification message when an alert is triggered. For more information, see Adding Alert Notifications.

5 **View Suppression Policy**. You can suppress notifications for a time period by specifying a suppression policy. For more information, see Suppressing Notifications.

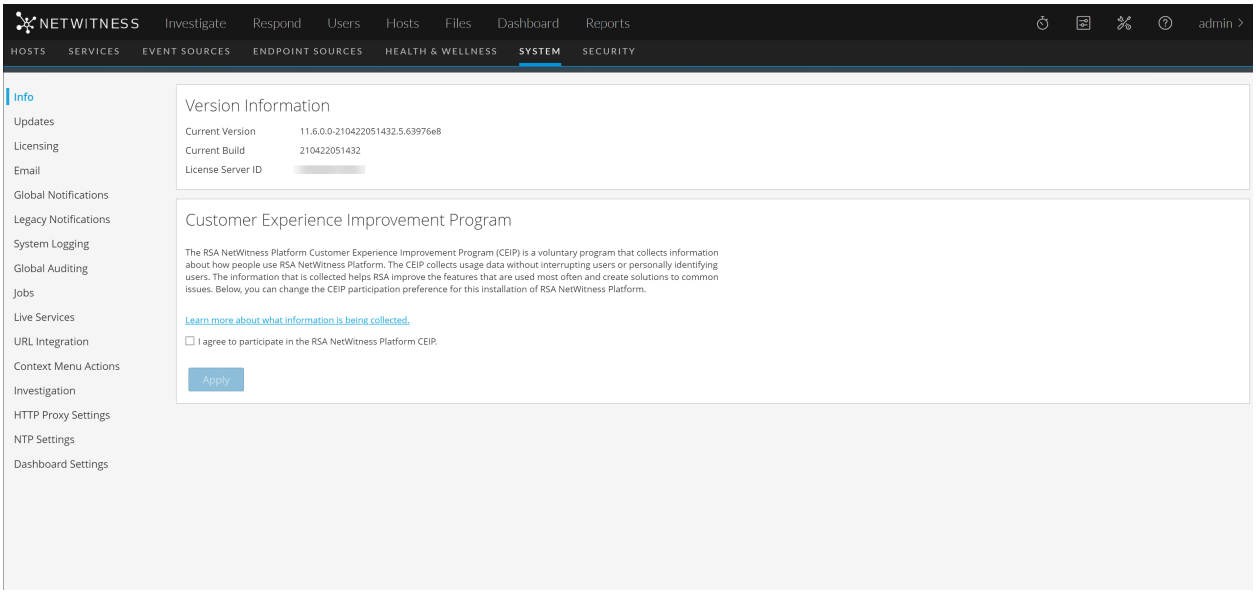
System View - System Info Panel

This topic describes the System Information panel, which displays information about the system version and license status.

The required role to access this view is **Manage System Settings**.

To access this view, do one of the following:

- Go to  **(Admin) > System**.
  The System Information panel is displayed by default.
- When you receive a notification that a new version of NetWitness is available in the Notifications tray, click **View**.



The Version Information section displays version information about the version of NetWitness that is currently installed. The following table describes the features of the Version Information section.

Name	Description
Current Version	Displays the version of NetWitness that is currently running. The format of the version is *major-relase.minor-release.stability-id.build-number*. Possible values for the *stability-id* are: • 1 - Development • 2 - Alpha • 3 - Beta • 4 - RC • 5 - Gold

Name	Description
Current Build	Identifies the current build revision for use in troubleshooting situations.
License Server ID	Each client host is shipped with the Local Licensing Server (LLS) installed to manage host licenses. This field indicates whether the LLS is installed for this instance of NetWitness. • When the LLS is installed, the Licensing Server ID is displayed. • **Unknown** indicates that the LLS is not installed.
License Status	Indicates whether or not the license is enabled. If the license is: • Enabled, **Enabled** is displayed in this field and there is a **Disable** button to the right so you can disable it. • Disabled, **Disabled** is displayed in this field and there is an **Enable** button to the right so you can enable it.

System Updates Panel - Settings Tab

System Updates Settings tab describes the interface you use to set up a connection to Live Update Repository. These settings ensure that the NetWitness can reach the Live Update Repository and synchronize it with your Local Update Repository.

The required permission to access this view is **Apply System Updates**.

To access this view:

1. Go to  **(Admin) > System**.
2. Select **Updates**.

What do you want to do?

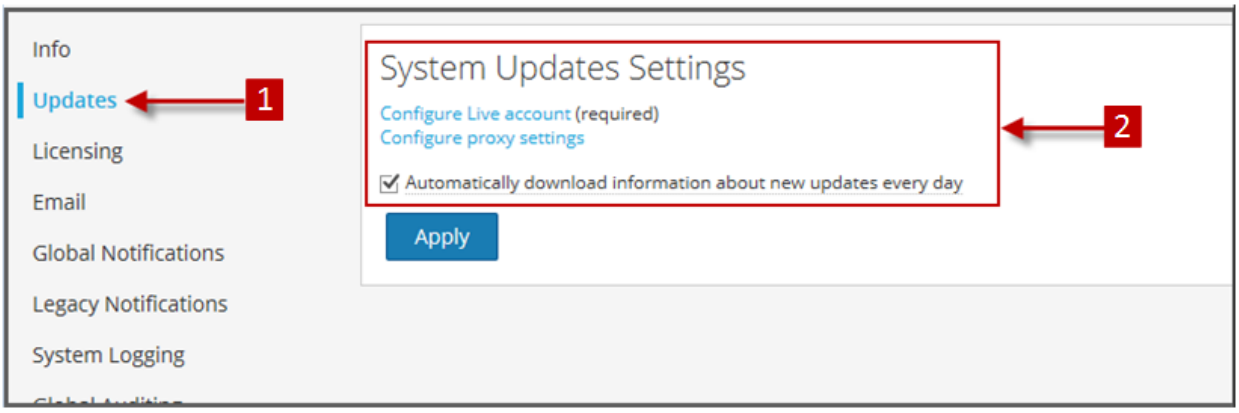
Role	I want to ...	Show me how
Administrator	Automatically download updates	Enable automatic synchronization with the NetWitness update repository.

Related Topics

Manage NetWitness Platform Updates

Quick Look

The System Updates Settings panel is displayed.



1 Displays System Update Setting Tab
2 Configure Account and Setting for Automatic Updates

Features

This table describes the features in the System Updates Settings panel.

Feature	Description
Configure Live account	Displays the **(Admin)** > **System** > **Live Services** panel in which you can configure your Live Account credentials if they are not configured.
Configure proxy settings	Displays the **(Admin)** > **System** > **HTTP Proxy Settings** panel in which you can configure an HTTP proxy if it is not configured.
Automatically download information about new updates every day	Select to enable automatic synchronization with the NetWitness update repository. If there are new updates available, information will automatically be displayed in the **(Admin)** > **HOSTS** panel.
Apply	Applies the settings in this tab.

System Logging - Settings View

The NetWitness Settings view in the System Logging panel configures the size of the log files, the number of backup log files maintained, as well as the default logging levels for the packages within NetWitness. The "Configure Log File Settings" topic in the *System Configuration Guide* provides detailed procedures.

To access the Settings tab:

1. Go to **(Admin) > System.**
2. In the options panel, select **System Logging**.

   The System Logging panel opens to the Realtime tab by default.
3. Click the **Settings** tab.

What do you want to do?

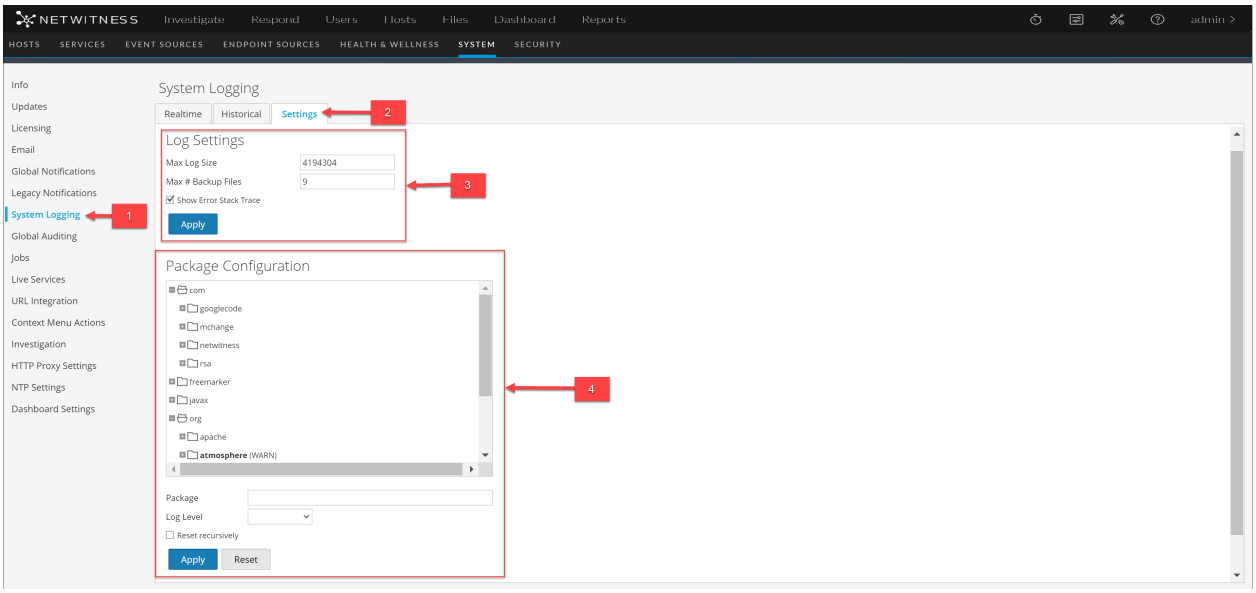
Role	I want to ...	Show me how
Administrator	Configure the size of the Log files	See the "Configure Log File Settings" topic in the *System Configuration Guide*

Related Topics

System Logging - Historical Tab

System Logging - Realtime Tab

Quick Look



1 Displays System Logging Panel
2 Displays Settings Tab
3 Configure Log Settings
4 Configure Packages

Features

The Settings tab has two sections: Log Settings and Package Configuration.

Log Settings

The Log Settings section configures the size of the NetWitness log files and the number of backup logs that NetWitness maintains.

Feature	Description
Max Log Size	Configures the maximum size in bytes of each log file. The minimum value for this setting is **4096**.
Max # Backup Files	Specifies how many backup log files are maintained. The minimum value for this setting is **0**. When the maximum number of log files is attained, and a new backup file is made, the oldest backup is discarded.
 Show Error Stack Trace	Displays ERROR, STACK, and TRACE log messages.
Apply	Puts the settings into effect immediately for all future logs.

Package Configuration

The Package Configuration section shows the NetWitness packages in a tree structure.

Feature	Description
Package tree	Contains all the packages used within NetWitness. You can drill down into the tree to view the log levels of each package. The **root** logging level represents the default log level for all packages that are not explicitly set. The root level is set to **INFO**
Package field	The name of the selected package when you select a package in the **Package** tree.
Log Level	If the selected package has a log level explicitly set, the value is displayed in the **Log Level** field.
 Reset recursively	Resets the log recursively.

Feature	Description
Apply	Puts settings into effect immediately for all future logs.
Reset	Resets the selected package to the log level of **root**.

System Logging - Realtime Tab

This topic describes the features of the System Logging > Realtime tab and the Services Logs view > Realtime tab.

The Realtime tab is a view of the NetWitness log or a service log. When it is initially loaded, the view contains the last 10 log entries. As new entries become available, the view is updated with those entries.

To access the Realtime tab:

1. Go to  **(Admin) > System.**
2. In the options panel, select **System Logging**.
   The System Logging panel opens to the **Realtime** tab by default.

What do you want to do?

Role	I want to ...	Show me how
Administrator	See details of Log entry	Display System and Service Logs

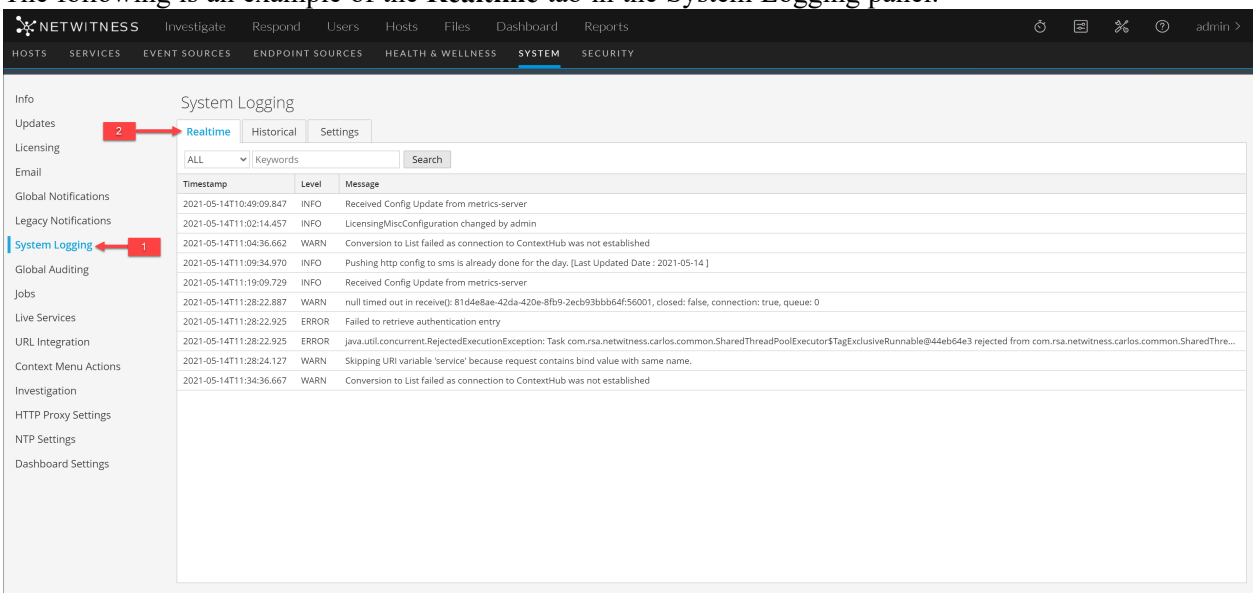
Related Topics

System Logging - Settings View

System Logging - Historical Tab

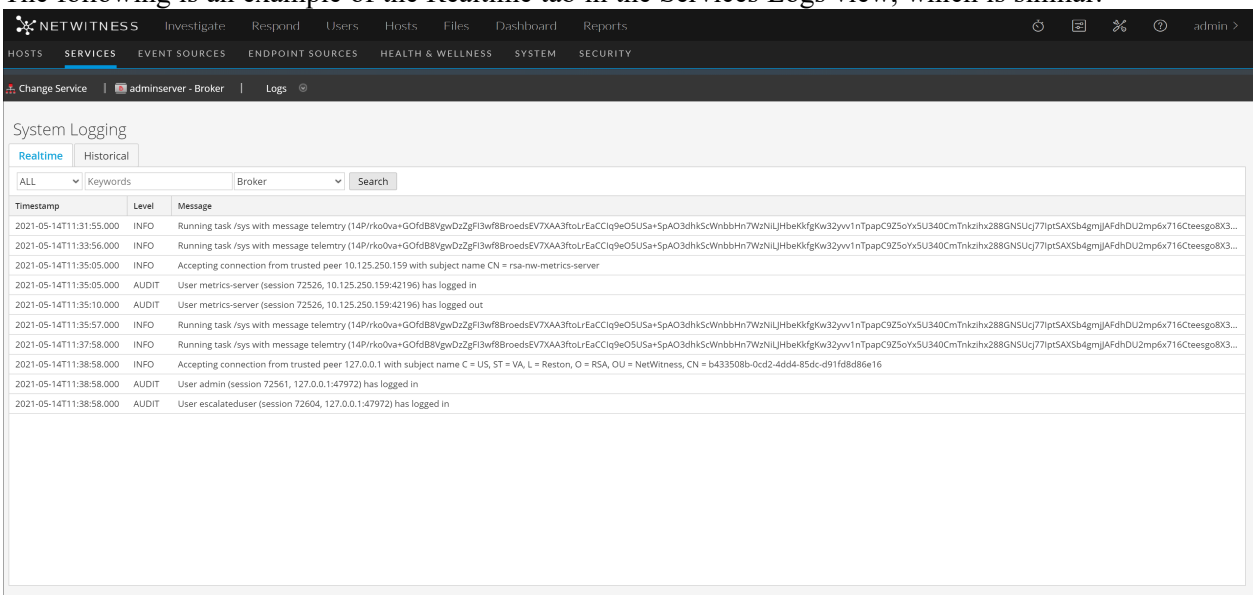
Quick Look

The following is an example of the **Realtime** tab in the System Logging panel.



1 Displays System Logging Panel
2 Displays Realtime Tab

The following is an example of the Realtime tab in the Services Logs view, which is similar.



Features

The Realtime tab has a toolbar with input fields to allow filtering of the entries, and below the toolbar is a grid containing the log entries.

Toolbar

Feature	Description
Log Level drop-down 	Selects the log level for entries to display in the grid. The **Log Level** drop-down shows the available log levels for the system or the service. • System logs have seven log levels. • Service logs have only six log levels because they do not include the **TRACE** level. • The default is **ALL** log entries.
Keywords field	Specifies a keyword to use when filtering entries. This field is the same for system and service log filtering.
Service field (Service Logs only)	Specifies the service type to use when filtering service log entries. Possible values are the host or the service.
Search button	Click to activate filtering based on the log level, keyword, and service selections.

Log Grid Columns

Column	Description
Timestamp	This is the timestamp for the entry.
Level	This is the log level for the message.
Message	This is the text of the log entry.

System Logging - Historical Tab

The Historical tab provides a searchable view of a NetWitness log or a service log in a paged format. When initially loaded, the grid shows the last page of the log entries for the system or the system.

To access the Historical tab:

1. Go to **(Admin) > System.**
2. In the options panel, select **System Logging**.

   The System Logging panel opens to the **Realtime** tab by default.
3. Click the **Historical** tab.

What do you want to do?

Role	I want to ...	Show me how
Administrator	View the Historical Graph	Historical Graph for System Stats

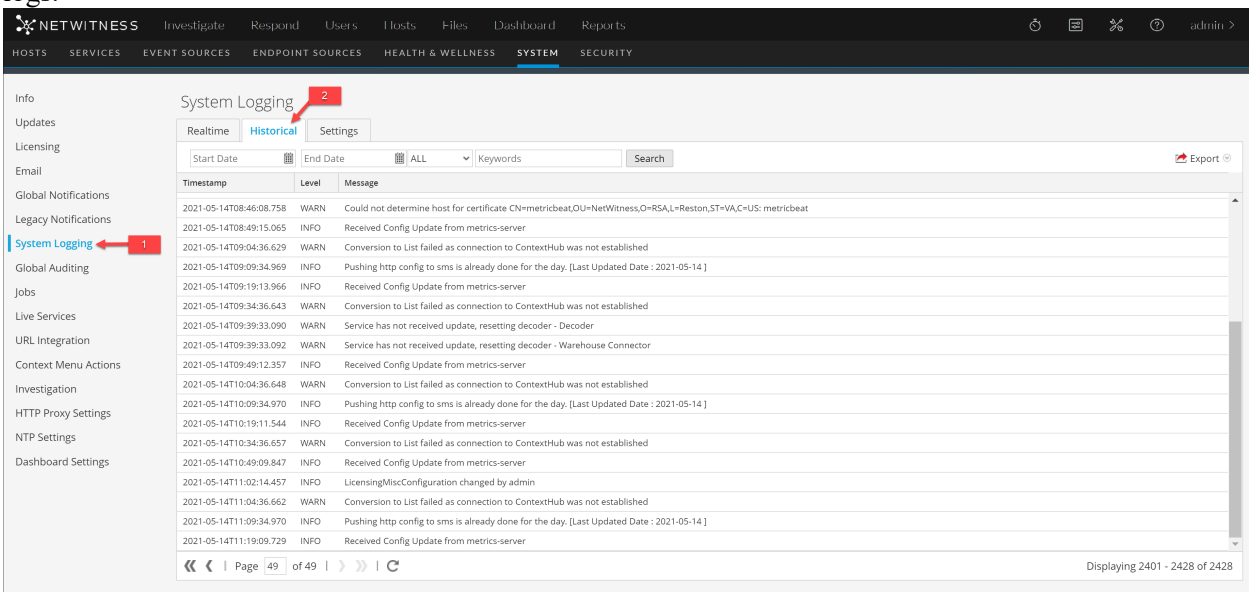
Related Topics

System Logging - Realtime Tab

System Logging - Settings View

Quick Look

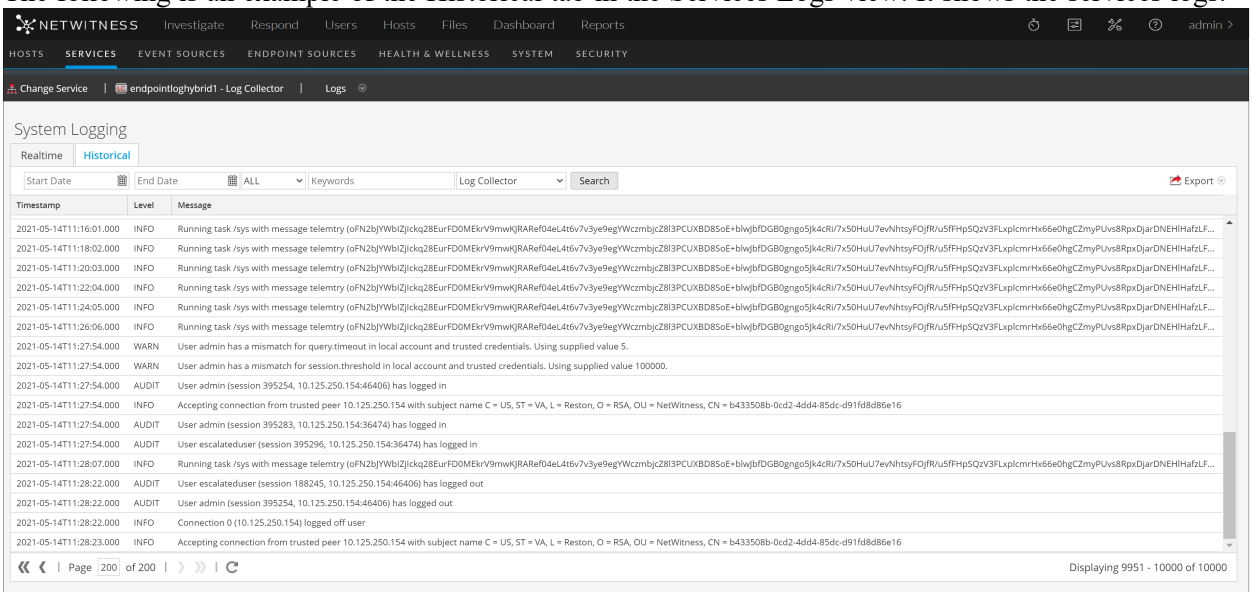
The following is an example of the **Historical** tab in the System Logging panel. It shows the NetWitness logs.



1 Displays System Logging Tab

2 Displays Historical Tab

The following is an example of the Historical tab in the Services Logs view. It shows the services logs.



Features

The Historical tab has a toolbar with input fields to allow filtering of the entries, a grid containing the log entries, and paging tools.

Feature	Description
Start Date and End Date	The Start Date and End Date range search options limit the log entries to a point in time. When used, you must provide both a start and end date. The times are optional. The date range is validated to assure that the end date is not before the start date.
Log Level drop-down 	Selects the log level for entries to display in the grid. The Log Level drop-down shows the available log levels for the system or the service. • System logs have seven log levels. • Service logs have only six log levels because they do not include the TRACE level. • The default is ALL log entries.
Keyword field	Specifies a keyword to use when filtering entries. This field is the same for system and service log filtering.
Service field (Service Logs only)	Specifies the service type to use when filtering service log entries. Possible values are the host or the service.
Search button	Click to activate a search based on the start and end date, log level, keyword, and service selections.

Feature	Description
Export	Click to export the currently viewed grid entries to a text file. You can select either comma-separated or tab-separated format for the entries in the file.

Column	Description
Timestamp	This is the timestamp for the entry.
Level	This is the log level for the message.
Message	This is the text of the log entry.

The paging tools below the grid provide a way to navigate through the pages of log entries.



Search Log Entries

To search the results shown in the Historical tab:

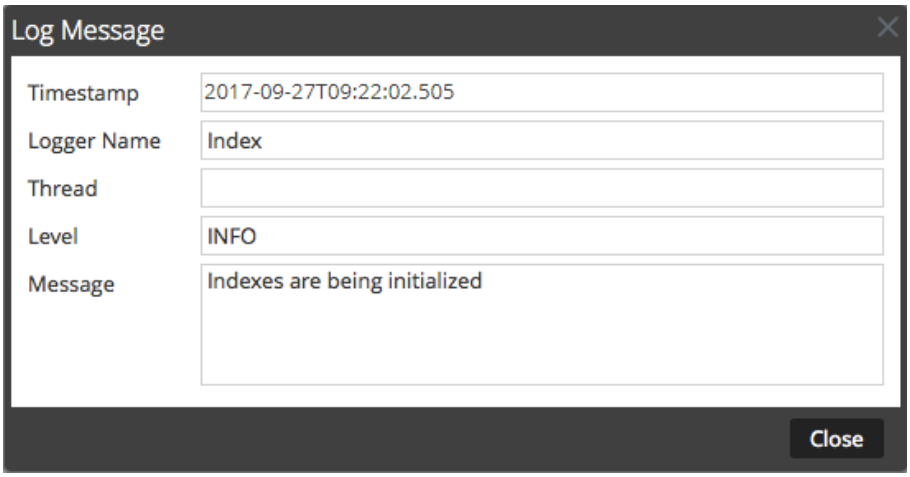
1. (Optional) Select a **Start Date** and **End Date**. Optionally, select a **Start Time** and **End Time**.
2. (Optional) For system and service logs, select a **Log Level** and a **Keyword**, or both.
3. (Optional) For service logs, select the **Service**: host or service.
4. Click **Search**.
   The view is refreshed with the most recent 10 entries matching your filter. As new matching log entries become available, the view is updated to show those entries.

Show Details of a Log Entry

Each row of the Historical tab Log grid provides the summary information of a log entry. To view complete details:

1. Double-click a log entry.

   The Log Message dialog, which contains the Timestamp, Logger Name, Thread, Level and Message, is displayed.



2. When finished viewing, click **Close**.

Page Through the Entries

To view the different pages of the grid, use the paging controls on the bottom of the grid as follows:

- Use the navigation buttons
- Manually enter the page you want to view, and press **ENTER**.

Export

To export the logs in the current view:

Click **Export**, and select one of the drop-down options, **CSV Format** or **Tab Delimited**.
The file is downloaded with a filename that identifies the log type and the field delimiter. For example, a NetWitness system log exported with comma-separated values is named `UAP_log_export_CSV.txt`, and an appliance log exported with tab-separated values is named `APPLIANCE_log_export_TAB.txt`.

NetWitness® Platform XDR

Version 12.0

UEBA Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Introduction **4**

UEBA Supported Sources by Schema 4

Authentication Schema 4

File Schema: 4

Active Directory Schema 4

Endpoint Process Schema 5

Endpoint Registry Schema 5

Packet Schema 5

UEBA Configuration **6**

ueba-server-config script 6

reset-presidio script 7

UEBA Indicator Forwarder 8

Update Data Source Details 8

Add Features for UEBA Packet Schema 9

Add the Hunting Pack: 9

Add JA3 and JA3s: 10

Assign User Access to UEBA 10

Create an Analysts Role 11

Enable User Entity Behavior Analytics Incident Rule 12

Learning Period Per Scale **14**

Learning Period Per Scale for 11.5 14

Physical Machine 14

Virtual Machine 15

Learning Period Per Scale for 11.5.1 and later versions 18

Physical Machine 18

Virtual Machine 19

Troubleshooting UEBA Configurations **21**

Task Failure Issues in Airflow 21

User Interface Inaccessible 22

Get UEBA Configuration Parameters 23

Check UEBA Progress Status using Airflow: 23

Check if data is received on the UEBA by Kibana: 24

Scaling Limitation Issue 25

UEBA Policy Issue 26

Troubleshoot Using Kibana 26

Troubleshoot Using Airflow 27

Introduction

NetWitness® UEBA configuration is designed for analysts to perform analytics for leveraged data collected from netwitness logs and networks to perform UEBA analytics.

Note: Mixed mode is not supported for UEBA in NetWitness Platform. The NetWitness server, and UEBA must all be installed and configured on the same NetWitness Platform version.

UEBA Supported Sources by Schema

Authentication Schema

- Windows Logon and Authentication Activity in Version 11.2 - Supported Event IDs: 4624, 4625, 4769, 4648 (device.type=winevent_snare|winevent_nic)
- RSASecurID Token in Version 11.3.1 - device.type = 'rsaacesrv' ec.activity = 'Logon'
- RedHat Linux in Version 11.3.1- device.type = 'rhlinux'
- Windows Remote Management in Version 11.3.2 - Supported Event IDs: 4624,4625,4769,4648 (device.type=windows)
- VPN Logs and in Version 11.5 - event.type = 'vpn' ec.activity = 'logon'
- Azure AD Logs in Version 11.5 - device.type = 'microsoft_azure_signin_events'

Note: Make sure you have configured the Azure Monitor plugin in your deployment. This enables UEBA to run a query for Azure AD log events for monitoring purposes in the correct format. For more information on how to configure the Azure Monitor plugin, see the *Azure Monitor Event Source Configuration Guide*.

File Schema:

- Windows File Servers in Version 11.2 - Supported Event IDs: 4663,4660,4670,5145 (device.type=winevent_snare|winevent_nic)
- device.type=windows in Version 11.3.1

Active Directory Schema

- Windows Active Directory in Version 11.2 - Supported Event IDs: 4741,4742,4733,4734,4740,4794,5376,5377,5136,4764,4743,4739,4727,4728,4754,4756,4757,4758,4720,4722,4723,4724,4725,4726,4738,4767,4717,4729,4730,4731,4732 (device.type=winevent_snare|winevent_nic)
- device.type=windows in Version 11.3.1

Endpoint Process Schema

- Endpoint Process in Version 11.3 - Category = 'Process Event'

Endpoint Registry Schema

- Endpoint Registry in Version 11.3 - Category = 'Registry Event'

Packet Schema

- TLS in Version 11.4 - Service 443 (direction='outbound')

Note: The TLS Packet requires adding the hunting package and enabling the JA3 features as described in Add required features for UEBA Packets Schema.

UEBA Configuration

This topic provides the high-level tasks required to configure UEBA.

IMPORTANT: Changing the UEBA start-date or the UEBA processed schemas requires a re-run of the UEBA system as well as cleanup of the UEBA databases. In order to avoid deleting the information in the UI, you can use the reset_presidio.py script as described in reset-presidio script, it will keep the data in the UI (e.g. Alerts, Indicators, Entities and Scores).

Note: Steps 1 to 4 must be executed as root on the UEBA machine.

ueba-server-config script

The `ueba-server-config` script is usually used to configure and run the UEBA component after the deployment. Also, it can be used to update the UEBA configuration during run time.

IMPORTANT: If you change the start-time or the processing schemas, you must re-run UEBA. All script arguments (except the boolean arguments) are mandatory and must be filled.

For more information on the script parameters, see the NetWitness Installation Guide for Version 11.5.

To run the script use the following command `/opt/rsa/saTools/bin/ueba-server-config --help`

Argument	Variable	Description
`-u`	`<user>`	User name of the credentials for the Broker or Concentrator instance that you are using as a data source.
`-p`	`<password>`	Password of the credentials for the Broker or Concentrator instance that you are using as a data source. The following special characters are supported in a password. `!"#$%&()*+,-:;<=>?@[\]^_`\{\|}` If you want to include a special character or special characters, you must delimit the password with an apostrophe sign, for example: `sh /opt/rsa/saTools/bin/ueba-server-config -u brokeruser -p '!"UHfz?@ExMn#$' -h 10.64.153.104 -t 2018-08-01T00:00:00Z -s 'AUTHENTICATION FILE ACTIVE_DIRECTORY TLS PROCESS REGISTRY' -o broker -v`
`-h`	`<host>`	IP address of the Broker or Concentrator used as the data source. Currently, only one data source is supported.
`-o`	`<type>`	Data source host type (`broker` or `concentrator`).

Argument	Variable	Description
-t	<startTime>	Historical start time as of which you start collecting data from the data source in YYYY-MM-DDTHH-MM-SSZ format (for example, 2018-08-15T00:00:00Z). **Note:** The script interprets the time you enter as UTC (Coordinated Universal Time) and it does not adjust the time to your local time zone.
-s	<schemas>	Array of data schemas. If you want to specify multiple schemas, use a space to separate each schema (for example, AUTHENTICATION FILE ACTIVE_DIRECTORY PROCESS REGISTRY TLS).
-v		verbose mode.
-e	<argument>	Boolean Argument. This enables the UEBA indicator forwarder to Respond. **Note:** If your NetWitness deployment includes an active Respond server, you can transfer NetWitness UEBA indicators to the Respond server and create incidents by enabling the indicator forwarder, from this data. For more information on how to enable the NetWitness UEBA incidents aggregation, see Enable User Entity Behavior Analytics Incident Rule.

Note: The TLS packet requires adding the hunting package and enabling the JA3 features. For more information, see Add Features for UEBA Packet Schema.

reset-presidio script

IMPORTANT: The reset_presidio.py script deletes the UEBA back-end databases and can also delete the front-end database that is present in the UI.

The reset_presidio.py script is used to re-run the UEBA system as well as to update the UEBA start-date and the processing schemas easily without having to provide all the other parameters required by the ueba-server-config script. This script re-runs the UEBA while it deletes the backed data (models, aggregations, etc.). To delete the front-end data (UI entities and alerts, etc.) use the clean option. If you don't specify a date, the script will set the default start date, a 28 days earlier than the current date. NetWitness recommends that the UEBA start date is set to 28 days earlier than the current date. For UEBA systems that intend to process TLS data, you must verify that the start date is set to no later than 14 days earlier than the current date.

Note: UEBA requires to process 28 days of data before the alerts can be created.
- If you choose a start date that is less than 28 days before the current date, for example 10 days earlier from the current date, you will have to wait for another 18 days from the current date to see alerts in your UEBA system (if created).
- If you choose a start date that is greater than 27 days, it's recommended to delete the front-end database as well (use the -c) to avoid duplicate alerts.

To run the script, load the Airflow virtual environment variables as follows:

source /etc/sysconfig/airflow

source $AIRFLOW_VENV/bin/activate

```
OWB_ALLOW_NON_FIPS=on python
/var/netwitness/presidio/airflow/venv/lib/python2.7/site-packages/presidio_
workflows-1.0-py2.7.egg/presidio/utils/airflow/reset_presidio.py --help
```

```
deactivate
```

Argument	Variable	Description
-h, --help		Script Help
-c, --clean	<argument>	If true, clean any existing data in Elasticsearch DB (as Alerts, Indicators, Entities, etc), all data will be deleted form the UEBA UI
-s	<schema>	Reconfigure the UEBA engine array of schemas (e.g. [AUTHENTICATION FILE ACTIVE_DIRECTORY PROCESS REGISTRY TLS])
-d	<date>	Reconfigure the UEBA engine to start from midnight UTC of this date. If not set, by default reset the start date to 27 days before the current system day, at midnight UTC, to avoid duplicate alerts in the UEBA UI, in case you didn't cleaned the elasticsearch data (-c) (e.g. 2010-12-31)

UEBA Indicator Forwarder

Note: The UEBA Indicator Forwarder is supported by the UEBA from version 11.3 and later. If your NetWitness environment includes an active respond server, you can transfer the UEBA indicators to the respond server and to the correlation server in order to create Incidents. For more information, see Enable User Entity Behavior Analytics Incident Rule.

Run the following command to activate the UEBA Indicator Forwarder:

```
curl -X PATCH http://localhost:8881/configuration -H 'content-type:
application/json' -d '{"operations":
[{"op":"replace","path":"/outputForwarding/enableForwarding","value":true}]}'
```

To deactivate the UEBA indicator forwarder, change the “value":true at the request body to be “value":false.

Update Data Source Details

In order to update the details of the data source you must use the `ueba-server-config` script. For more information, see ueba-server-config script.

The data sources details are:

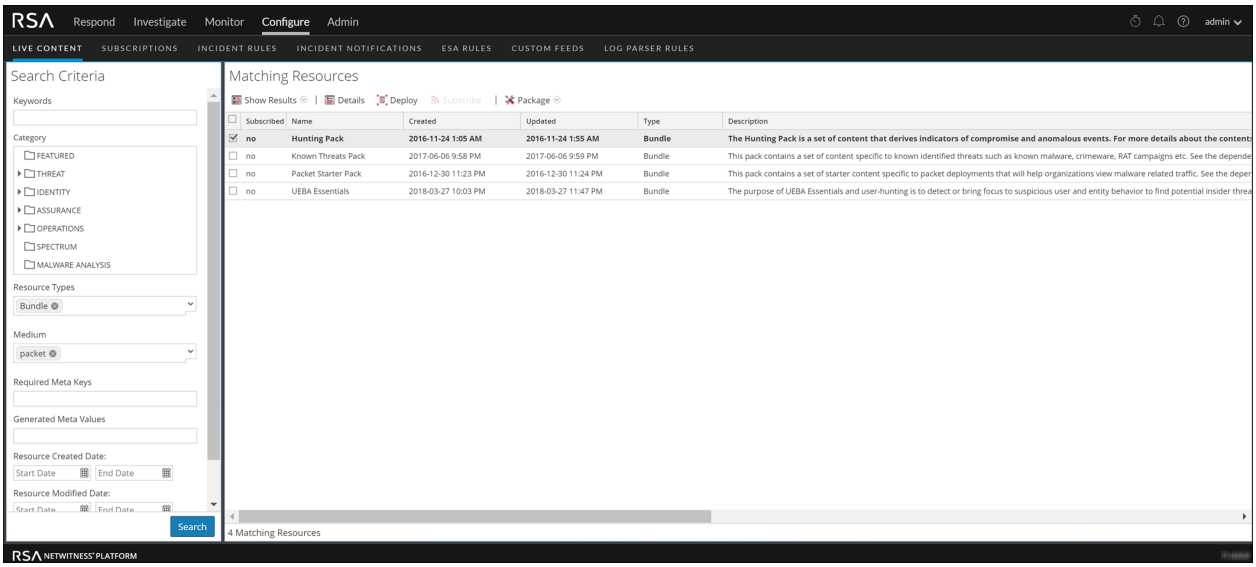
- Data Source type (Broker / Concentrator).
- Data Source username.
- Data Source password.
- Data Source host.

Add Features for UEBA Packet Schema

Add the Hunting Pack:

In NetWitness Platform, add the hunting pack or verify it it's available:

1. Login to NetWitness Platform
2. Navigate to **(Admin)** and select **Admin Server**
3. Click and select **Configure** > **Live Content**



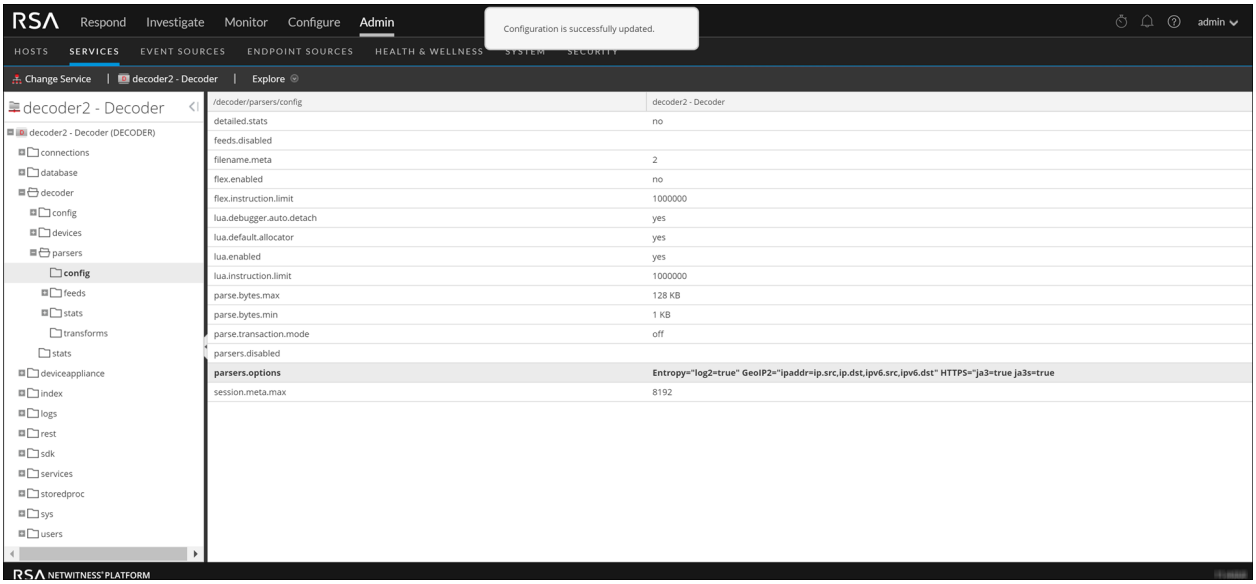
4. On the left menu, select the following:
   a. Bundle under Resources Type.
   b. Packet under Medium
5. Click **Search**.
   A list of matching resources is displayed.
6. Select **Hunting Pack** from the list and click **Deploy**.
   The hunting pack is added.

Add JA3 and JA3s:

The JA3 and JA3s fields are supported by the Network Decoder in 11.3.1 and later. Verify that your Network Decoder is upgraded to one of these versions.

To add JA3 and Ja3s:

1. Log in to NetWitness Platform.
2. Go to  (Admin) and select Decoder.
3. Navigate to `/decoder/parsers/config/parsers.options`.
4. Add `HTTPS="ja3=true ja3s=true`.

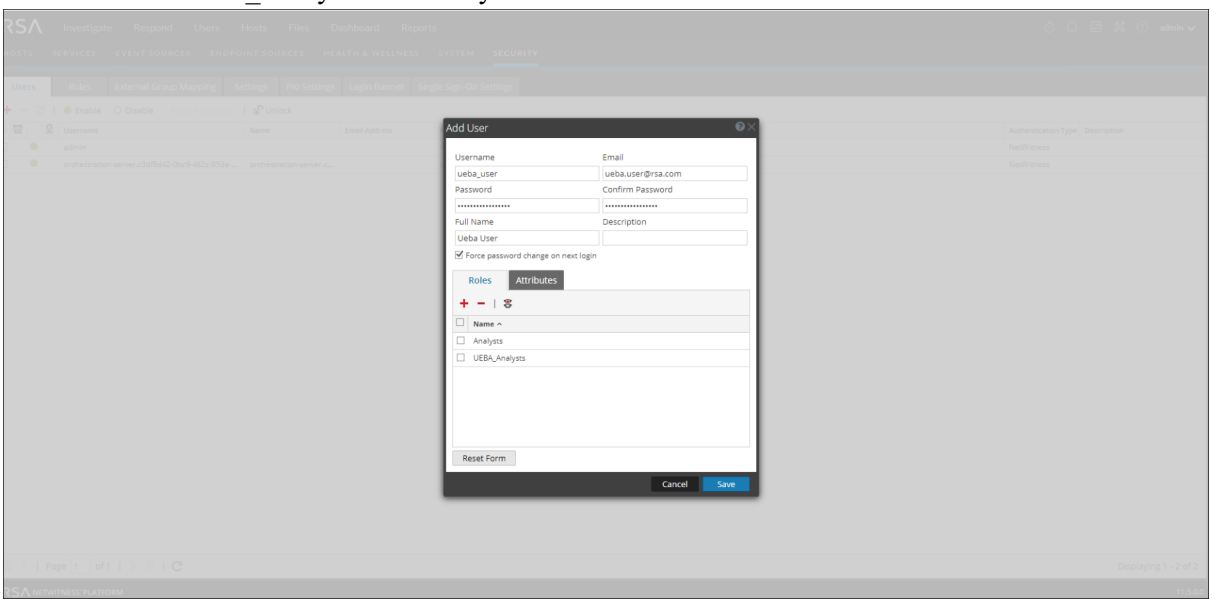


The JA3 and JA3s fields are configured.

Assign User Access to UEBA

To create a user with privileges to access the UEBA pages (Users tab) on the Netwitness UI do the following:

1. Navigate to (Admin) > **Security**.
2. Create a new UEBA_Analysts and Analysts user roles.



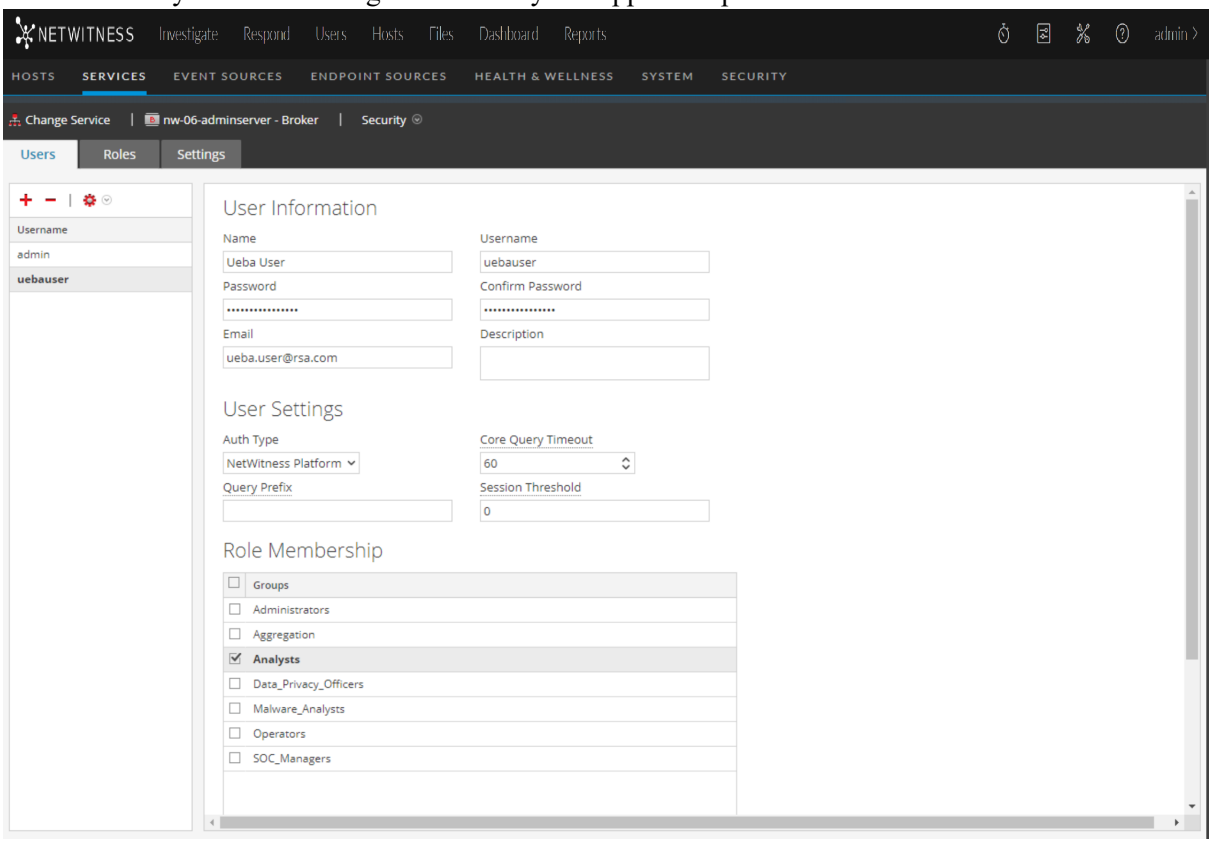
For more information, see the "Manage Users with Roles and Permissions" topic in the *System Security and User Management Guide*.

Create an Analysts Role

In order to fetch data from the data source (Broker / Concentrator), you need to create a user using the Docktor-UEBA: Validation Too" role in the data source service.

1. Navigate to the security tab at the data source service page.
2. (Admin) > **Services** > **Security**

3. Create an analyst user and assign it to the any of supported special characters.

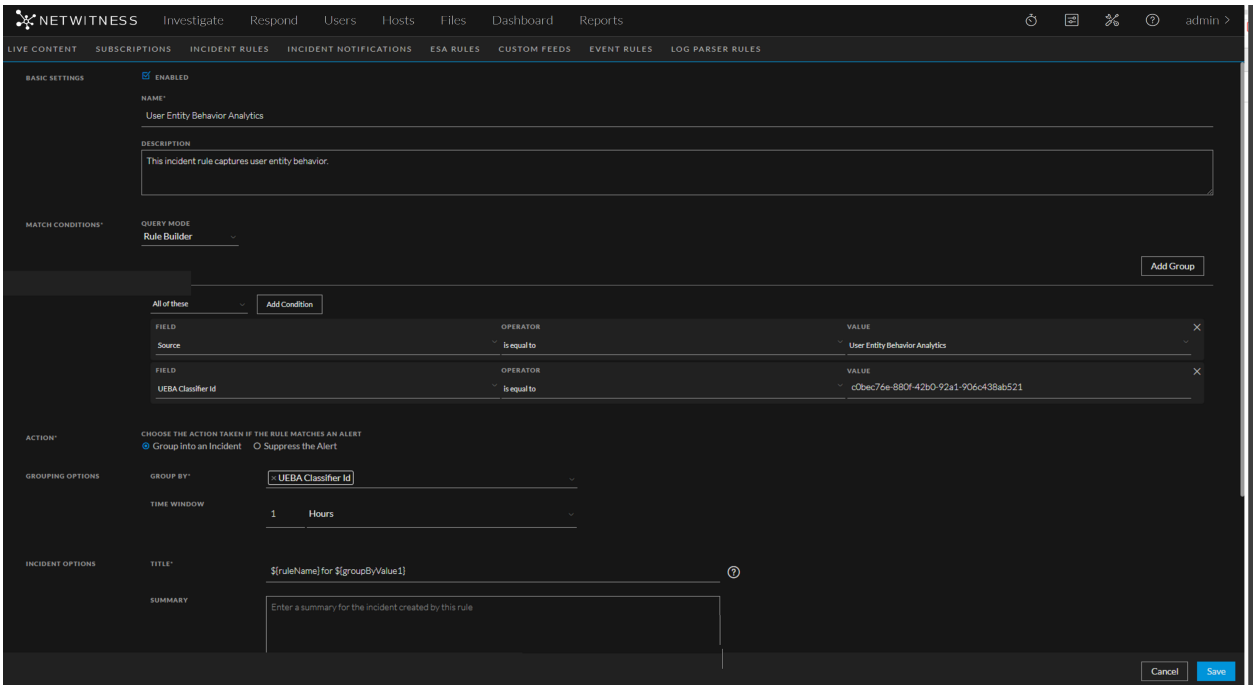


Enable User Entity Behavior Analytics Incident Rule

In order to aggregate the UEBA indicators under Incident rule, follow the instructions below:

Enable the UEBA Forwarding process as described in Enable UEBA Indicator Forwarder.

1. Go to (Configure) > **Incident Rules**.
2. Select the **User Entity Behavior Analytics** rule.
3. Select the enable check box and click **Save**.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
LIVE CONTENT
SUBSCRIPTIONS
INCIDENT RULES
INCIDENT NOTIFICATIONS
ESA RULES
CUSTOM FEEDS
EVENT RULES
LOG PARSER RULES
BASIC SETTINGS
ENABLED
NAME*
User Entity Behavior Analytics
DESCRIPTION
This incident rule captures user entity behavior.
MATCH CONDITIONS*
QUERY MODE
Rule Builder
Add Group
All of these
Add Condition
FIELD
OPERATOR
VALUE
Source
is equal to
User Entity Behavior Analytics
FIELD
OPERATOR
VALUE
UEBA Classifier Id
is equal to
c0bec76e-880f-42b0-92a1-906c438ab521
ACTION*
CHOOSE THE ACTION TAKEN IF THE RULE MATCHES AN ALERT
Group into an Incident
Suppress the Alert
GROUPING OPTIONS
GROUP BY*
UEBA Classifier Id
TIME WINDOW
1
Hours
INCIDENT OPTIONS
TITLE*
${ruleName} for ${groupByValue1}
SUMMARY
Enter a summary for the incident created by this rule
Cancel
Save

Learning Period Per Scale

Learning Period Per Scale for 11.5

Physical Machine

SERIES 5 (DELL R630) SPECIFICATIONS

Supported Scale	Existing NetWitness customer (historical data available)	Learning Period Alerts will be generated when the learning period is complete
Logs and Endpoint data for 100,000 users + 20 million network events per day.	Yes	**11.5 Installation** Up to 4 days with 28 days of historical data.
	Yes	**11.5 Upgrade from 11.4.x with no schema changes** No learning period. • UEBA rerun is not required.
	Yes	**11.5 Upgrade from 11.3.x or prior versions with no schema changes** Up to 4 days with 28 days of historical data. • UEBA rerun is required.
	Yes	**11.5 Upgrade with schema changes** Up to 4 days with 28 days of historical data. • UEBA rerun is required

Supported Scale	Existing NetWitness customer (historical data available)	Learning Period Alerts will be generated when the learning period is complete
Logs and Endpoint data for 100,000 users + 60 million network events per day.	Yes	**11.5 Installation** Up to 14 days with 14 days of historical data.
	Yes	**11.5 Upgrade from 11.4.x with no schema changes** No learning period. • UEBA rerun is not required.
	Yes	**11.5 Upgrade from 11.3.x or prior versions with no schema changes** Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5.*
	Yes	**11.5 Upgrade with schema changes** Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5.*
Logs and Endpoint data for up to 100,000 users + 60 million network events per day.	No	**11.5 Installation** 28 days

Virtual Machine

CPU	Memory	Read IOPS	Write IOPS
16 cores	64GB	500MB	500MB

Note: NetWitness recommends you to deploy UEBA on a virtual host, only if your log collection volume is low. If you have a moderate to high log collection volume, NetWitness recommends you to deploy UEBA on the physical host as described in the "NetWitness UEBA Host Hardware Specifications" topic of the *Physical Host Installation Guide*. Contact Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897) for advice on choosing which host, virtual or physical, to use for UEBA.

Supported Scale	Existing NetWitness customer (historical data available)	Learning Period Alerts will be generated when the learning period is complete
Logs and Endpoint data for up to 100,000 users with 30 million events per day (no network data).	Yes	**11.5 Installation** Up to 4 days with 28 days of historical data.
	Yes	**11.5 Upgrade from 11.4.x with no schema changes** No learning period. • UEBA rerun is not required.
	Yes	**11.5 Upgrade from 11.3.x or prior versions with no schema changes** Up to 4 days with 28 days of historical data. • UEBA rerun is required.
	Yes	**11.5 Upgrade with schema changes** Up to 4 days with 28 days of historical data. • UEBA rerun is required

Supported Scale	Existing NetWitness customer (historical data available)	Learning Period Alerts will be generated when the learning period is complete
Logs and Endpoint data for up to 100,000 users with 30 million events per day + 20 million network events per day.	Yes	**11.5 Installation** Up to 14 days with 14 days of historical data.
		11.5 Upgrade from 11.4.x with no schema changes No learning period. • UEBA rerun is not required.
		11.5 Upgrade from 11.3.x or prior versions with no schema changes Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5.*
		11.5 Upgrade with schema changes Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5.*
Logs and Endpoint data for up to 100,000 users with 30 million events per day + 20 million network events per day.	No	**11.5 Installation** 28 days

Note: Network events per day refers to number of events consumed by UEBA per day. To determine the scale of network events for existing customers, see Troubleshooting UEBA Configurations.

Learning Period Per Scale for 11.5.1 and later versions

Note: For all supported scales, when historical data is not available, the learning period is 28 days.

Physical Machine

SERIES 5 (DELL R630) SPECIFICATIONS

Supported Scale for existing NetWitness customers (historical data is available)	Learning Period Alerts will be generated when the learning period is complete
Logs and Endpoint data for 100,000 users + 20 million network events per day.	**11.5.1 Installation** Up to 4 days with 28 days of historical data.
	11.5.1 Upgrade from 11.4.x No learning period. • UEBA rerun is not required.
	11.5.1 Upgrade from 11.3.x or prior versions Up to 4 days with 28 days of historical data. • UEBA rerun is required.
	11.5.1 Upgrade with schema removal Up to 4 days with 28 days of historical data. • UEBA rerun is required

Supported Scale for existing NetWitness customers (historical data is available)	Learning Period Alerts will be generated when the learning period is complete
Logs and Endpoint data for 100,000 users + 60 million network events per day.	**11.5.1 Installation** Up to 14 days with 14 days of historical data.
	11.5.1 Upgrade from 11.4.x No learning period. • UEBA rerun is not required.
	11.5.1 Upgrade from 11.3.x or prior versions Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5*.
	11.5.1 Upgrade with schema removal Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5*.

Virtual Machine

If there is not historical data, then the learning period will be 28 days.

CPU	Memory	Read IOPS	Write IOPS
16 cores	64GB	500MB	500MB

Note: NetWitness recommends you to deploy UEBA on a virtual host, only if your log collection volume is low. If you have a moderate to high log collection volume, NetWitness recommends you to deploy UEBA on the physical host as described in the " NetWitness UEBA Host Hardware Specifications" topic of the *Physical Host Installation Guide*. Contact Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897) for advice on choosing which host, virtual or physical, to use for UEBA.

Supported Scale for existing NetWitness customers (historical data is available)	Learning Period Alerts will be generated when the learning period is complete
Logs and Endpoint data for up to 100,000 users with 30 million events per day (no network data).	**11.5.1 Installation** Up to 4 days with 28 days of historical data.
	11.5.1 Upgrade from 11.4.x No learning period. • UEBA rerun is not required.
	11.5.1 Upgrade from 11.3.x or prior versions Up to 4 days with 28 days of historical data. • UEBA rerun is required.
	11.5.1 Upgrade with schema removal Up to 4 days with 28 days of historical data. • UEBA rerun is required
Logs and Endpoint data for up to 100,000 users with 30 million events per day + 20 million network events per day.	**11.5.1 Installation** Up to 14 days with 14 days of historical data.
	11.5.1 Upgrade from 11.4.x No learning period. • UEBA rerun is not required.
	11.5.1 Upgrade from 11.3.x or prior versions Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5.*
	11.5.1 Upgrade with schema removal Up to 14 days with 14 days of historical data. • UEBA rerun is required. **Note:** This scenario is impacted by ASOC-101686 known issue. For more information, see *NetWitness Release Notes for 11.5.*

Note: Network events per day refers to number of events consumed by UEBA per day. To determine the scale of network events for existing customers, see Troubleshooting UEBA Configurations.

Troubleshooting UEBA Configurations

This section provides information about possible issues when using NetWitness UEBA.

Task Failure Issues in Airflow

Problem	The `userId_output_entities` task fails when the username contains a backslash.
Cause	When events with usernames containing a backslash character is passed through UEBA, then the `userId_output_entities` task fails.
Solution	To resolve these issue contact the customer success to obtain the relevant files and execute the following steps: • Stop airflow-scheduler service. • Remove all MongoDB documents in the "aggr", "accm" and "input" collections that contains `context.userId` with hashtag. These documents can be located using the FindCollecionsContainsBackslash.js script. • Replace the `/var/netwitness/presidio/flume/conf/adapter/transformers/authentication.json` file with the updated authentication .json. • Restart the airflow-scheduler service. • Validate that the next run of the `userId_output_entities` task is completed successfully.

Problem	The `AUTHENTICATION_userId_build_feature_historical_data` task fails when the username contains a hashtag.
Cause	When events with usernames containing a hashtag character is passed through UEBA, then the `AUTHENTICATION_userId_build_feature_historical_data` task fails.
Solution	To resolve these issue contact the customer success to obtain the relevant files and execute the following steps: • Stop airflow-scheduler service. • Remove all MongoDB documents in the "aggr", "accm" and "input" collections that contains `context.userId` with hashtag. These documents can be located using the FindCollecionsContainsHashtagContextUserId.js script. • Replace the `/var/netwitness/presidio/flume/conf/adapter/transformers/authentication.json` file with the updated authentication .json. • Restart the airflow-scheduler service.

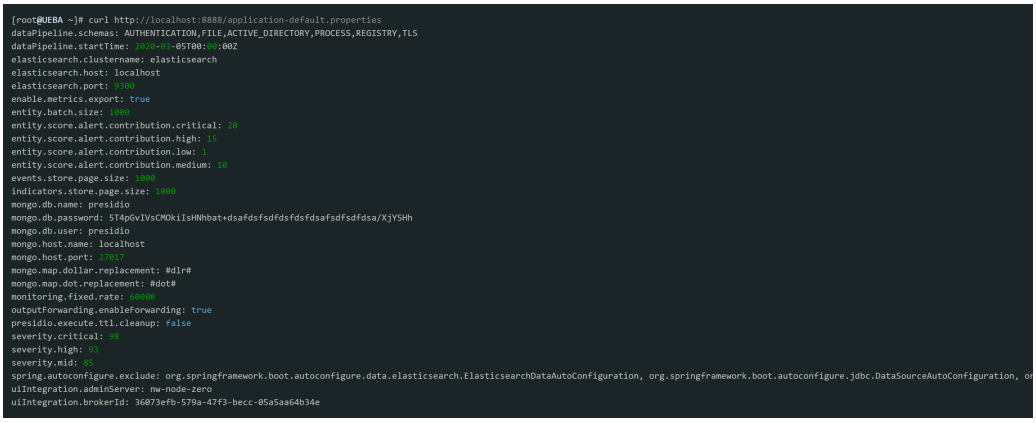
	• Validate that the next run of `AUTHENTICATION_userId_build_feature_historical_data` task is completed successfully.

User Interface Inaccessible

Problem	The User Interface is not accessible.
Cause	You have more than one NetWitness UEBA service existing in your NetWitness deployment and you can only have NetWitness UEBA service in your deployment.

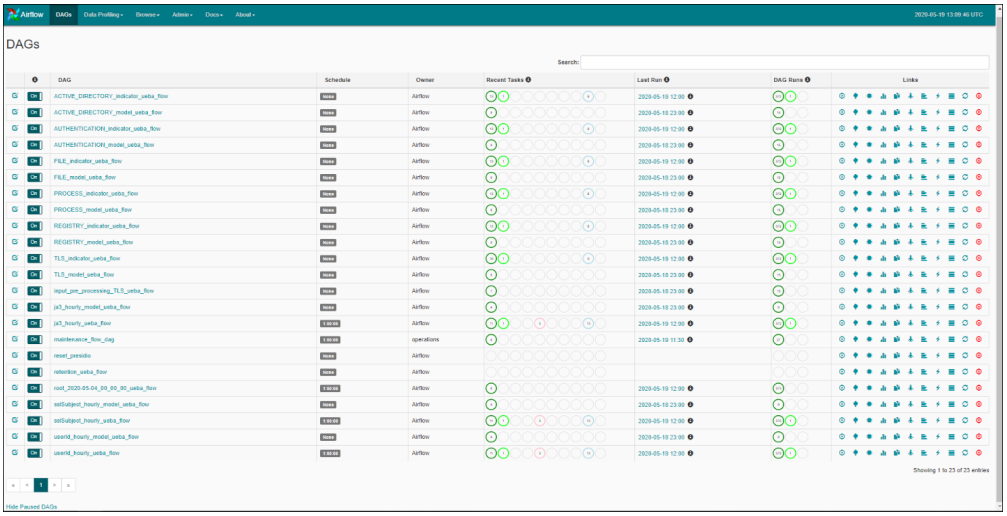
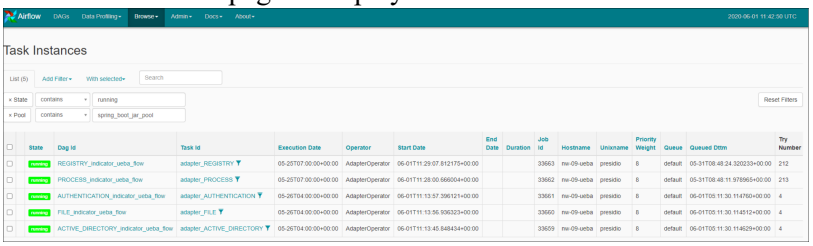
Solution	Complete the following steps to remove the extra NetWitness UEBA service.
	1. SSH to NW Server and run the following commands to query the list of installed NetWitness UEBA services.
	`# orchestration-cli-client --list-services|grep presidio-airflow`
	`... Service: ID=7e682892-b913-4dee-ac84-ca2438e522bf, NAME=presidio-airflow, HOST=xxx.xxx.xxx.xxx:null, TLS=true`
	`... Service: ID=3ba35fbe-7220-4e26-a2ad-9e14ab5e9e15, NAME=presidio-airflow, HOST=xxx.xxx.xxx.xxx:null, TLS=true`
	2. From the list of services, determine which instance of the `presidio-airflow` service should be removed (by looking at the host addresses).
	3. Run the following command to remove the extra service from Orchestration (use the matching service ID from the list of services):
	`# orchestration-cli-client --remove-service --id <ID-for-presidio-airflow-form-previous-output>`
	Note: Run the following command to update NW Server to restore NGINX:
	`# orchestration-cli-client --update-admin-node`
	4. Log in to NetWitness, go to  **(Admin) > Hosts**, and remove the extra NetWitness UEBA host.

Get UEBA Configuration Parameters

Issue	How to get UEBA configuration parameters?
Explanation	In order to get the UEBA configuration main parameters, run the `curl http://localhost:8888/application-default.properties` command from the UEBA machine.  The main parameters which will be returned are the following: • uiIntegration.brokerId: The Service ID of the NW data source (Broker / Concentrator) • dataPipeline.schemas: List of schemas processed by the UEBA • dataPipeline.startTime: The date the UEBA started consuming data from the NW data source • outputForwarding.enableForwarding: The UEBA Forwarder status
Resolution	See the resolution for these statistics in the Troubleshooting UEBA Configurations section.

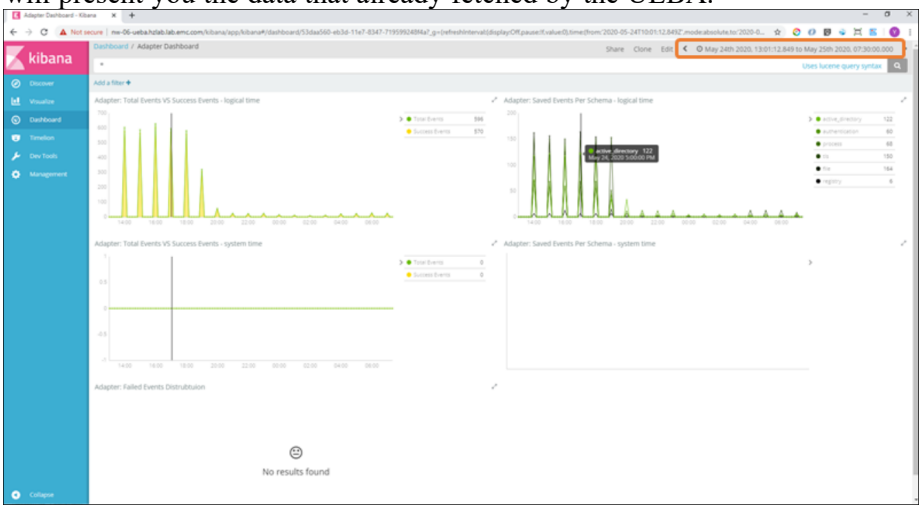
Check UEBA Progress Status using Airflow:

Issue	How to check UEBA progress status using Airflow?

Resolution	1. Navigate to- https://<UEBA-host-name>/admin. Enter the admin username and the deploy-admin password. The following image is of the Airflow home page that shows the system is working as expected.  2. Make sure that no red or yellow circles appear in the main page: • red circle indicates that a task has failed. • yellow circle indicates that a task has failed and is “awaiting” for a retry. If a “failed” or “up-for-retry” task appears, investigate what is the root cause of the problem. 3. Make sure the system continues to run. 4. Tap the **Browse** button and select **Task Instance**. 5. Add the following filters: **State** = `running` and **Pool** = `spring_boot_jar_pool`. The Task Instance page is displayed.  The **Execution Date** column shows the current time window for each running task. Make sure the execution date is greater than the UEBA start-date and that new tasks have an updated date are added to the table.

Check if data is received on the UEBA by Kibana:

Issue	How to check if data is received on the UEBA by Kibana:

Explanation	Navigate to- https://<UEBA-host-name>/kibana. Enter the admin username and password: To check that the data is flowing to the UEBA go to the Adapter Dashboard: Tap the Dashboard tab in the left menu Tap Adapter Dashboard at the right menu Select the relevant time range at the top bar The charts on this dashboard will present you the data that already fetched by the UEBA. 

Scaling Limitation Issue

When installed on a Virtual Machine, UEBA can process up to 20 million network events per day. Based on this limitation, you may encounter the following issues.

Issue	How to determine the scale of network events currently available, to know if it exceeds the UEBA limitation.
Solution	To know the network data limit, perform the following : • Run the query on the Broker or Concentrator that connects to UEBA using NetWitness UI: `service=443 && direction='outbound' && analysis.service!='quic' && ip.src exists && ip.dst exists && tcp.srcport!=443` Calculate the total number of events for the selected days (including weekdays with standard workload). If the average is above 20 million per day then it indicates that UEBA's supported scale is exceeded.

Issue	Can UEBA for Packets be used if UEBA's supported scale is exceeded?
Solution	You must create or choose a Broker that is connected to a subset of Concentrators that does not exceed the supported limit. To know the network data limit, perform the following : • Run the query on the Concentrator that connects to UEBA using NetWitness UI:

	`service=443 && direction='outbound' && analysis.service!='quic' && ip.src exists && ip.dst exists && tcp.srcport!=443` Calculate the total number of events for the selected days (including weekdays with standard workload). If the average is above 20 million per day then it indicates that UEBA's supported scale is exceeded.

Note: The Broker must query all the available and needed data needed such as logs, endpoint and network (packets). UEBA packets models are based on the whole environment. Hence, make sure that the data parsed from the subset of Concentrators is consistent.

UEBA Policy Issue

Issue	After you create a rule under UEBA policy, duplicate values are displayed in the Statistics drop-down.

Solution	To remove the duplicate values, perform the following:
	1. Log in to MongoDB using following command:`mongo admin -u deploy_admin -p {Enter the password}`
	2. Run the following command on MongoDB:
	`use sms;`
	`db.getCollection('sms_statdefinition').find({componentId :"presidioairflow"})`
	`db.getCollection('sms_statdefinition').deleteMany ({componentId :"presidioairflow"})`

Troubleshoot Using Kibana

Issue	After you deploy NetWitness UEBA, the connection between the NetWitness and NetWitness UEBA is successful but there are very few or no events in the **Users** > **OVERVIEW** tab. 1. Log in to **Kibana**. 2. Go to **Table of Content** > **Dashboards** > **Adapter Dashboard**. 3. Adjust the **Time Range** on the top-right corner of the page and review the following: • If the new events are flowing. • In the **Saved Events Per Schema** graph, see the number of successful events per schema per hour. • In the **Total Events vs. Success Events** graph, see the total number of events and number of successful events. The number of successful events should be more every hour. For example, in an environment with 1000 users or more, there should be thousands

	of authentication and file access events and more than 10 Active Directory events. If there are very few events, there is likely an issue with Windows auditing.
Solution	You must identify the missing events and reconfigure the Windows auditing. 1. Go to **INVESTIGATE** > **Navigate**. 2. Filter by **devide.type= device.type "winevent_snare"** or **"winevent_nic"**. 3. Review the events using **reference.id** meta key to identify the missing events. 4. Reconfigure the Windows auditing. For more information, see **NetWitness UEBA Windows Audit Policy** topic.

Issue	The historical load is complete and the events are coming from Adapter dashboard but no alerts are displayed in the **Users** > **OVERVIEW** tab.
Solution	1. Go to **Kibana** > **Table of content** > **Scoring and model cache**. 2. Adjust the **Time Range** from the top-right corner of the page, and see if the events are scored.

Issue	The historical load is complete but no alerts are displayed in the **Investigate** > **Users** tab.
Solution	1. Go to **Kibana** > **Dashboard** > **Overview**. 2. Adjust the **Time Range** from the top-right corner of the page, and see how many users are analyzed and if any anomalies are found.

Troubleshoot Using Airflow

Issue	After you start running the UEBA it is not possible to remove a data source during the run process else the process stops.
Solution	You must either continue the process till it completes or remove the required data source from UEBA and rerun the process.

Issue	After you deploy UEBA and if there are no events displayed in the **Kibana** > **Table of content** > **Adapter** dashboard and Airflow has already processed the hours but there are no events. This is due to some communication issue.
Solution	You must check the logs and resolve the issue. 1. Log in to **Airflow**. 2. Go to **Admin** > **REST API Plugin**. 3. In the **Failed Tasks Logs**, click **execute**.

A zip file is downloaded.

4. Unzip the file and open the log file to view and resolve the error.
5. In the **DAGs** > **reset_presidio**, click **Trigger Dag**.
   This deletes all the data and compute all the alert from the beginning.

Note: During initial installation, if the hours are processed successfully but there are no events, you must click reset_presidio after fixing the data in the Broker. Do not reset if there are alerts.

NetWitness® Platform XDR

Version 12.0

NetWitness UEBA Quick Start Guide

NETWITNESS

Platform XDR

NetWitness® Platform XDR

Version 12.0

NetWitness UEBA Quick Start Guide

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable

NETWITNESS
Platform XDR

NetWitness® Platform XDR

Version 12.0

NetWitness UEBA Quick Start Guide

to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

NETWITNESS
Platform XDR

What is NetWitness UEBA?

NetWitness UEBA (User and Entity Behavior Analytics) is an advanced analytics solution for discovering, investigating, and monitoring risky behaviors across all users and entities in your network environment. NetWitness UEBA is used for:

- Detecting malicious and rogue users
- Pinpointing high-risk behaviors
- Discovering attacks
- Investigating emerging security threats
- Identify potential attacker activity

About this Guide

This guide provides end-to-end instructions to configure NetWitness UEBA and to use UEBA features.

NetWitness Platform 11.7 Documentation in NETWITNESS community Link

NetWitness Platform product documentation is organized along functional lines. If you are looking for a specific guide or version, go to the Version 11.x Master Table of Contents.

Use these links to view the NetWitness Platform 11.5 documentation. Both links provide the same documentation, in these two formats:

- HTML Guides include the latest information for currently supported 11.x versions: NetWitness Platform 11.x Documentation.
- PDF Guides provide the information for a specific version: NetWitness Platform 11.7 PDFs.

Use these links to access documentation that is not related to a particular version of the software:

- Hardware setup guides: https://community.netwitness.com/t5/netwitness-platform-hardware/tkb-p/netwitness-hardware-documentation
- Documentation for Content such as feeds, parsers, application rules, and reports: https://community.netwitness.com/t5/threat-intelligence/ct-p/threat-intelligence.

Getting Started

The following tasks can be performed in any sequence.

Description	References
	
View information about product updates, improvements, and known issues.	Release Notes
Understand NetWitness UEBA	*NetWitness UEBA User Guide*

Setup and Installation

Standalone Installation

The following tasks must be performed in the following sequence.

Description	References
	
Review the supported hardware.	"System Requirement" topic in *UEBA Standalone Installation Guide*
Review the UEBA deployment.	"NetWitness UEBA Standalone Installation " topic in *UEBA Standalone Installation Guide*
Configure the ports on your firewall.	"NetWitness UEBA Standalone Installation " topic in *UEBA Standalone Installation Guide*
Install NetWitness Server host.	"Installation Tasks" topic in *UEBA Standalone Installation Guide*
Install 11.3 Log Hybrid Host.	"Installation Tasks" topic in *UEBA Standalone Installation Guide*
Install and Configure NetWitness UEBA.	"Installation Tasks" topic in *UEBA Standalone Installation Guide*
Assign the UEBA_Analysts and Analysts roles to the UEBA users.	"Role Permissions" in the System Security and User Management Guide

Fresh Installation

The following tasks must be performed in the following sequence.

Description	References
	
Review the supported hardware.	"Supported Hardware" in the *Physical Host Installation Guide*
Review the UEBA architecture.	"NetWitness Platform Network Architecture Diagram" topic in the *Deployment Guide*
Configure the ports on your firewall.	"Network Architecture and Ports" topic in the *Deployment Guide*
Install NetWitness Server host and other components.	"Task 1 - Install 11.7 on the NetWitness Server (NW Server) Host" and "Task 2 - Install 11.7 on Other Component Hosts" in *Physical Host Installation Guide* "Install NetWitness Platform Virtual Host in Virtual Environment" in the *Virtual Host Installation Guide*
Install UEBA.	"NetWitness® UEBA" in *Physical Host Installation Guide*
Assign the UEBA_Analysts and Analysts roles to the UEBA users.	"Role Permissions" in the *System Security and User Management Guide*

Update

The following tasks must be performed in the following sequence.

Description	References
	
Deploy the Endpoint Pack from Live, which contains File Category Lua Parser for the UEBA integration with Endpoint.	During deployment, you must specify Endpoint Log Hybrid Log Decoder service. In case of multiple Endpoint servers, select all the Endpoint Log Hybrid Log Decoder services
Enable Endpoint data sources such as Process and Registry to generate alerts in UEBA.	"Enable Endpoint Data Sources" in the *Update Instructions*
Enable UEBA indicator forwarder to transfer the UEBA indicators to the NetWitness Respond server and to the correlation server to create an incidents.	"Enable UEBA Indicator Forwarder" in the *Update Instructions*

Description	References
After you update to NetWitness Platform 11.7 the Broker or Concentrator UUID changes. You must update the NetWitness Platform core services, and update the Broker or Concentrator UUID.	"Update Broker or Concentrator UUID" in the *Update Instructions*
Update Airflow Configuration.	"Update Airflow Configuration" in the *Update Instructions*
Restart the Airflow scheduler service after the presidio_ upgrade DAG is successful.	"Restart Airflow scheduler service" in the *Update Instructions*

Investigation

The following tasks can be performed in any sequence.

Description	References
	
Investigate high risk users.	"Investigate High-Risk Users" topic in the *NetWitness UEBA User Guide*
Investigate top alerts.	"Investigate Top Alerts" topic in the *NetWitness UEBA User Guide*

Monitoring

The following tasks can be performed in any sequence.

Description	References
	
Review NetWitness UEBA metrics in Health and Wellness.	"View NetWitness UEBA Metrics in Health and Wellness" topic in the NetWitness UEBA User Guide
Monitor Health and Wellness of UEBA.	"Monitor Health and Wellness of UEBA" topic in the NetWitness UEBA User Guide

NetWitness® Platform XDR

Version 12.0

UEBA Standalone Installation Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Introduction **4**

NetWitness UEBA Standalone Installation **5**

System Requirements **6**

Physical Host Hardware Specifications 6

Virtual Host Specifications 7

Installation Media 7

Installation Tasks **8**

For Physical Hosts: 8

Task 1. Install 11.6 on the NetWitness Server Host 8

Task 2. Install 11.6 Log Hybrid Host 8

Task 3. Install and Configure NetWitness® UEBA 8

Configure NetWitness UEBA 10

Enable Access Permission for the NetWitness UEBA User Interface 13

For Virtual Hosts: 13

Task 1. Install 11.6 on the NetWitness Server Host 13

Task 2. Install 11.6 Log Hybrid Host 14

Task 3. Install and NetWitness® UEBA 14

Post Installation Task **15**

Set a user with privileges to access the UEBA pages 15

Introduction

NetWitness® UEBA standalone installation is designed to support other security tools such as Security Information and Event Management (SIEM). It allows users of a third-party SIEM solutions to leverage NetWitness for UEBA.

NetWitness® UEBA standalone installation includes:

1. NetWitness Admin Server
2. NetWitness Log Hybrid
3. NetWitness UEBA

Note: NetWitness® UEBA standalone installation supports up to 100,000 users on a single running instance. If you have more than 100,000 users, contact NetWitness Customer Support.

Windows Log Sources

NetWitness UEBA standalone installation natively supports the following Windows log sources:

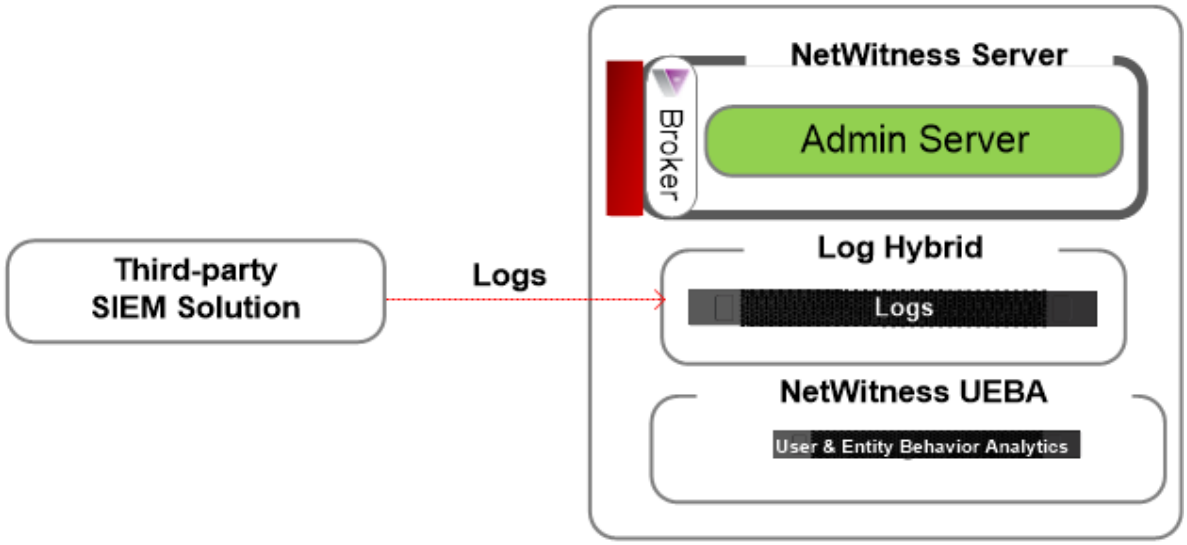
- Windows Active Directory
- Windows Logon and Authentication Activity
- Windows File Servers
- Windows Remote Management

For more information, see the "NetWitness UEBA Use Cases for Windows Logs" topic in the *NetWitness UEBA User Guide*.

NetWitness UEBA Standalone Installation

This section contains a high-level UEBA standalone installation diagram and a list of NetWitness UEBA ports.

The following diagram illustrates the NetWitness UEBA standalone installation.



Note:

- UEBA connects to the Broker or Log Hybrid host.
- To configure the Log Hybrid for logs from a third-party SIEM, contact your RSA Account Manager.

UEBA Host and Service Ports

Source Host	Destination Host	Destination Ports	Comments
UEBA Server	NW Server	TCP 443	Update Repository
UEBA Server	NW Server	TCP 56003 (SSL), 50003 (Non-SSL), 50103 (REST)	Broker Application Ports
UEBA Server	NW Server	TCP 56005 (SSL), 50005 (Non-SSL), 50105 (REST)	Log Hybrid, Concentrator Application Ports
Admin Workstation	UEBA Server	443	UEBA Monitoring
Admin Workstation	UEBA Server	22	SSH

System Requirements

You can run the NetWitness® UEBA standalone installation as follows:

- Physical Hosts (software running on hardware supplied)
  - Install physical hosts and connect to the network as described in the NetWitness® Platform Hardware Setup Guides and the *NetWitness® Platform Physical Host Installation Guide*.
  - Set up licensing for NetWitness Platform as described in the *NetWitness® Platform Licensing Guide*.
- On-Premises (On-Prem) Virtual Hosts (Software Only provided by NetWitness)

Physical Host Hardware Specifications

You must install the NetWitness UEBA host on the S5 (Dell R630 appliance) hardware.

SERIES 5 (DELL R630) SPECIFICATIONS

Specification	Capacity
Model	Dell PowerEdge R630xl
Processor Type	Intel Xeon E5 -2680v3
Processor Speed Cache Number of Cores Number of Processors Number of Threads	2.5 GHz 30MB 12 2 24
Total Memory	256GB
Internal Disk Controller	Dell PERC H730
External Disk Controller	Dell PERC H830
SAN Connectivity (HBA) - Optional	N/A
Remote Management Card	iDRAC8 Enterprise
Drives	Total - 6 Drives 2 x 1TB, 2.5" HDD 4 x 2TB, 2.5" HDD
Chassis	1U
Weight	18.4 kg (40.5 lbs)

Specification	Capacity
NIC Card*	On Board 2 x 10 Gb Copper 2 x 10 Gb & 2 x 1Gb Copper (Other options are available)
Dimensions	H: 4.28 cm (1.68 in.) x W: 48.23 cm (18.98 in.) x D: 75.51 cm (29.72 in.)
Power	1100W Redundant
BTU/hr	4100 BTU/hr (max)
Amps (Spec)	1100W / 220VAC = 5A
Actual Amp Draw (Post Startup)	2.1 Amps
Events Per Second (EPS)	100K EPS
Throughput	N/A

* NIC Card options are available for swap with on-board daughter card or add on.

Virtual Host Specifications

Following is the recommended system requirements for a UEBA virtual host.

CPU	Memory	Read IOPS	Write IOPS
16 or 2.4GHz	64 GB	500	500

Note: NetWitness recommends that you only deploy UEBA on a virtual host if your log collection volume is low. If you have a moderate to high log collection volume, NetWitness recommends that you deploy UEBA on the physical host.

Installation Media

Installation media are in the form of OVA packages, which are available for download and installation from Download Central (https://download.rsasecurity.com). As part of your order fulfillment, NetWitness gives you access to the OVA.

Installation Tasks

This topic contains the tasks you must complete to install NetWitness UEBA standalone installation.

Note: Download or make sure you have access to the *Physical Host Installation Guide* for Version 11.6 before beginning the tasks. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

For Physical Hosts:

You must complete the following tasks in the order shown below.

Task 1. Install 11.6 on the NetWitness Server Host

Task 2. Install 11.6 Log Hybrid Host

Task 3. Install and Configure NetWitness® UEBA

Task 1. Install 11.6 on the NetWitness Server Host

For the NetWitness Server (NW Server), this task:

- Creates a base image.
- Sets up the 11.6 NW Server host.

For more information on how to install the NetWitness Server host, see "Install 11.6 on the NetWitness Server (NW Server) Host" section in the *Physical Host Installation Guide for Version 11.6*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Task 2. Install 11.6 Log Hybrid Host

For a non-NW Server host, this task:

- Creates a base image.
- Sets up the 11.6 non-NW Server host or Log Hybrid.

For more information on how to install the Log Hybrid host, see "Task 2 - Install 11.6 on Other Component Hosts" section in the *Physical Host Installation Guide for Version 11.6*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Task 3. Install and Configure NetWitness® UEBA

To set up NetWitness UEBA, you must install and configure the NetWitness UEBA service.

The following procedure shows you how to install the NetWitness UEBA service on a NetWitness UEBA Host Type and configure the service.

1. Complete steps 1 - 14 under "Task 2 - Install 11.6 on Other Component Hosts" in "Installation Tasks" of the *NetWitness Platform Physical Host Installation Guide for Version 11.6*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

   Note: The Kibana and Airflow webserver User Interface password is the same as the deploy_admin password. Make sure that you record this password and store it in a safe location.

2. Log in to the NetWitness Platform and go to **(Admin) > Hosts**.

   The New Hosts dialog is displayed with the Hosts view grayed out in the background.

   Note: If the New Hosts dialog is not displayed, click **Discover** in the **Hosts** view toolbar.

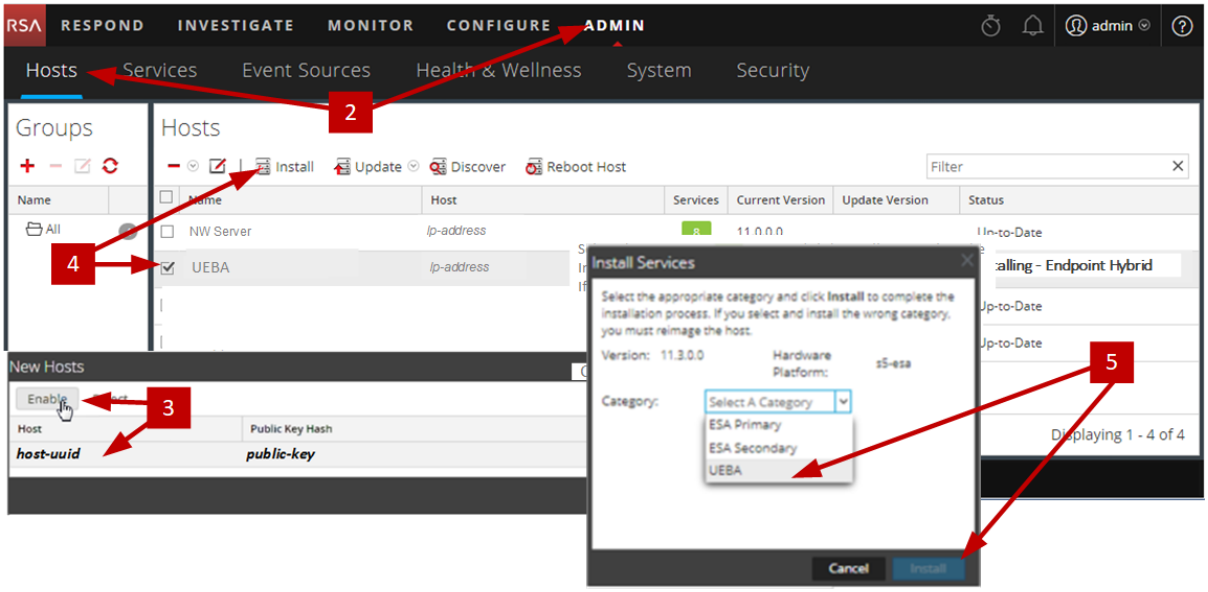
3. Select the host in the **New Hosts** dialog and click **Enable**.

   The New Hosts dialog closes and the host is displayed in the Hosts view.

4. Select that host in the **Hosts** view (for example, **UEBA**) and click Install.

   The Install Services dialog is displayed.

5. Select the **UEBA** Host Type and click **Install**.



6. Make sure that the UEBA service is running.
7. Complete licensing requirements for NetWitness UEBA.

   See the *NetWitness Platform 11.6 Licensing Management Guide* for more information. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

   Note: NetWitness Platform supports the User and Entity Behavior Analytics License (UEBA). This license is used based on the number of users. The Out-of-the-Box Trial License is a 90-day trial license. In case of UEBA licenses, the 90-day trial period begins from the time the UEBA service deployed on the NetWitness Platform product.

Configure NetWitness UEBA

To start running UEBA:

1. Define the following parameters: data schemas, data source (NetWitness Broker or Concentrator) and start date.

   a. Define UEBA schemas:
   Choose schemas from the following list:

   AUTHENTICATION, FILE, ACTIVE_DIRECTORY, PROCESS, REGISTRY and TLS.

   Note: The TLS packet requires adding the hunting package and enabling the JA3 feature. For more information regarding events that each schema contains, see the *NetWitness UEBA Configuration Guide*.

   b. Define the data source:
   If your deployment has multiple Concentrators, we recommend that you assign a Broker at the top of your deployment hierarchy for the NetWitness UEBA data source.

   c. Define the UEBA start-date:

   Note: The selected start date must contain events from all configured schemas.

   NetWitness recommends that the UEBA start date is set to 28 days earlier than the current date. For UEBA systems that intend to process TLS data, you must make sure that the start date is set to no later than 14 days earlier than the current date.

2. . Create a user account for the data source (Broker or Concentrator) to authenticate to the data.

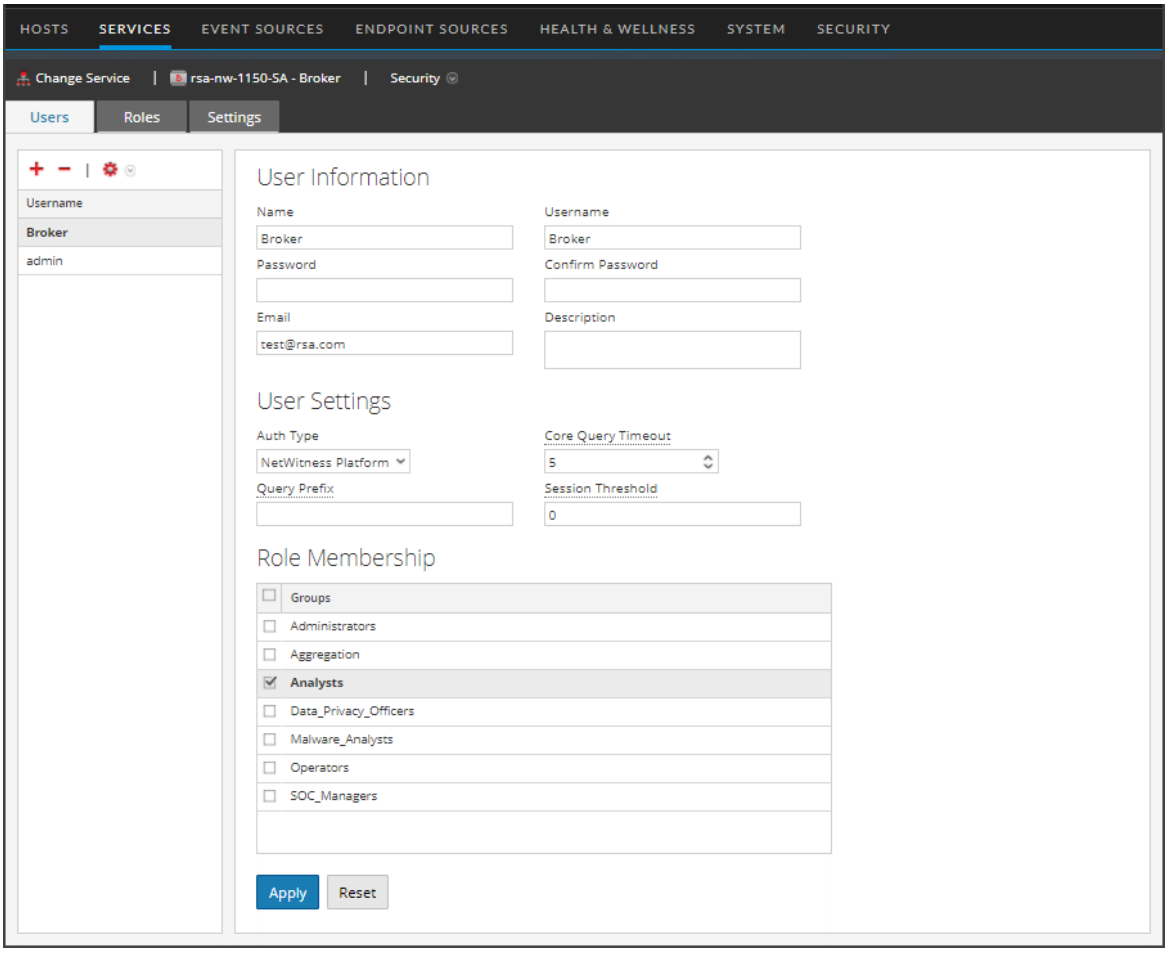
   a. Log into NetWitness Platform.

   b. Go to  (Admin) > **Services**.

   c. Locate the data source service (Broker or Concentrator).
   Select that service, and select (Actions) > **View** > **Security**.

   d. Create a new user and assign the “Analysts” role to that user.

The following example shows a user account created for a Broker.



3. SSH to the NetWitness UEBA server host.
4. (For Virtual Machines Only) Set the appropriate parallelism value:
   If the UEBA system runs on VM, update the airflow parallelism value to be 64 by running the following command:
   ```
   sed -i "s|parallelism = 256|parallelism = 64|g"
   /var/netwitness/presidio/airflow/airflow.cfg
   ```
5. Submit the following commands with the above parameters that you already defined.
   ```
   /opt/rsa/saTools/bin/ueba-server-config -u <user> -p <password> -h <host> -
   o <type> -t <startTime> -s <schemas> -v -e <argument>
   ```
   Where:

Argument	Variable	Description
`-u`	`<user>`	User name of the credentials for the Broker or Concentrator instance that you are using as a data source.

Argument	Variable	Description
-p	<password>	Password of the credentials for the Broker or Concentrator instance that you are using as a data source. The following special characters are supported in a password. !"#$%&()*+,-:;<=>?@[\]^_`\{\|} If you want to include a special character or special characters, you must delimit the password with an apostrophe sign, for example: sh /opt/rsa/saTools/bin/ueba-server-config -u brokeruser -p **'!"UHfz?@ExMn#$'** -h 10.64.153.104 -t 2018-08-01T00:00:00Z -s 'AUTHENTICATION FILE ACTIVE_DIRECTORY TLS PROCESS REGISTRY' -o broker -v
-h	<host>	IP address of the Broker or Concentrator used as the data source. Currently, only one data source is supported.
-o	<type>	Data source host type (broker or concentrator).
-t	<startTime>	Historical start time as of which you start collecting data from the data source in YYYY-MM-DDTHH-MM-SSZ format (for example, 2018-08-15T00:00:00Z). **Note:** The script interprets the time you enter as UTC (Coordinated Universal Time) and it does not adjust the time to your local time zone.
-s	<schemas>	Array of data schemas. If you want to specify multiple schemas, use a space to separate each schema (for example, AUTHENTICATION FILE ACTIVE_ DIRECTORY PROCESS REGISTRY TLS).
-v		**verbose mode**.
-e	<argument>	Boolean Argument. This enables the UEBA indicator forwarder to Respond. **Note:** If your NetWitness deployment includes an active Respond server, you can transfer NetWitness UEBA indicators to the Respond server and create incidents by enabling the indicator forwarder, from this data. For more information on how to enable the NetWitness UEBA incidents aggregation, see Installation Tasks.

6. Set the appropriate "Boot Jar Pools" slots:

   - **Physical Appliance**: Update the `spring_boot_jar_pool`" slots to be 18.
   - **Virtual Appliance**: If the UEBA system is running on VM, update the `spring_boot_jar_pool` and the `retention_spring_boot_jar_pool` slots values to 5.
     To update the “Spring Boot Jar Pools” slots, Go to the Airflow main page, tap the “Admin” tab at the top bar and tap “Pools”.

   a. To access the Airflow UI, go to `https://<UEBA_host>/admin` and enter the credentials.

   • User: admin

   • Password: The environment deploy admin password.

   b. Click on the pencil mark of the polls to update the slot values.



Enable Access Permission for the NetWitness UEBA User Interface

After you install NetWitness UEBA standalone 11.6, you need to assign the UEBA_Analysts and Analysts roles to the UEBA users. For more information, see 'Assign User Access to UEBA' topic in the *NetWitness UEBA Configuration Guide*. After this configuration, UEBA users can access the **Investigate** > **Users** view.

Note: To complete NetWitness UEBA configuration according to the needs of your organization, See*NetWitness UEBA Configuration Guide*.

For Virtual Hosts:

You must complete the following tasks in the order shown below.

Task 1. Install 11.6 on the NetWitness Server Host

Task 2. Install 11.6 Log Hybrid Host

Task 3. Install and NetWitness® UEBA

Task 1. Install 11.6 on the NetWitness Server Host

On the host you have deployed for the NetWitness Server (NW Server), this task installs:

- The 11.6.0.0 NW Server environmental platform.
- The NW Admin Server.
- A repository with the RPM files required to install the other functional components or services.

For more information on how to install the NetWitness Server host, see "Task 1- Install 11.6.0.0 on the NW Server Host" section in the *Virtual Host Installation Guide for Version 11.6*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Task 2. Install 11.6 Log Hybrid Host

Complete the following tasks on a non-NW Server host:

- Install the 11.6.0.0 environmental platform.
- Apply the 11.6.0.0 RPM files to the service from the NW Server Update Repository.

Note: You must install the Log Hybrid host.

For more information on how to install the non-NetWitness Server host, see "Task 3 - Install 11.6 for on Other Component Hosts" section in the *Virtual Host Installation Guide for Version 11.6*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues..

Task 3. Install and NetWitness® UEBA

Prerequisite: Increase Memory for Virtual Deployment

Virtual Machines are deployed with approximately 104 GB in the storage mount by default. To install NetWitness UEBA, you must increase the storage space in your virtual environment to at least 800 GB.

To set up NetWitness UEBA, you must install and configure the NetWitness UEBA service.

The following procedure shows you how to install the NetWitness UEBA service on a NetWitness UEBA Host Type and configure the service.

1. Complete steps 1 - 15 for Virtual Hosts under "Task 3 - Install 11.6 on Other Component Hosts" in "Installation Tasks" of the *NetWitness Platform Virtual Host Installation Guide for Version 11.6*. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

   Note: The Kibana and Airflow webserver User Interface password is the same as the deploy admin password. Make sure that you record this password and store it in a safe location.

2. Complete steps 2 - 9 under Task 3. Install and Configure NetWitness® UEBA.

Post Installation Task

Set a user with privileges to access the UEBA pages

After you install NetWitness UEBA standalone 11.6, you need to assign the UEBA_Analysts and Analysts roles to the UEBA users. For more information, see *NetWitness UEBA Configuration Guide*.

NetWitness® Platform XDR

Version 12.0

NetWitness UEBA User Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Introduction **6**

Users 6

Network 6

How NetWitness UEBA Works 7

Retrieve Data 8

Create Baselines 8

Create Baselines for Users 8

Create Baselines for Network 9

Detect Anomalies 9

Generate Indicators 9

User Indicators 9

Network Indicators 9

Generate Alerts 10

Users 10

Network 10

Prioritize User or Network Entity with Risky Behavior 10

Supported Sources 12

Log Sources 12

Network Sources 12

Recommended Workflows 12

Detection Workflow 12

Forensic Workflow 14

NetWitness UEBA Use Cases **16**

Use Case for Users 16

Use Case for Network Entities 16

Alert Types **17**

Alert Types for a User 17

Alert Types for Network Entities 21

NetWitness UEBA Indicators **23**

Indicators for Users 23

Windows File Servers 23

Active Directory 24

Logon Activity 24

Process 25

Registry 26

Indicators for Network Entities ... 26
Access NetWitness UEBA ... 29
UEBA Licensing ... 29
Investigate High-Risk Entities ... 30
Identify High-Risk Entities ... 31
View Top Ten Risky Entities ... 32
View All High-Risk Entities ... 32
View Entities of Specific Group ... 33
View Entity Based on Forensic Investigation ... 34
Begin an Investigation of High-Risk Entity ... 35
Take Action on High-Risk Users ... 37
Specify that an alert is not risky. ... 37
Save Behavioral Profile ... 37
Add All Users to the Watchlist ... 39
Watch Profile ... 39
Export a list of High-Risk Users ... 40
Investigate Top Alerts ... 42
Begin an Investigation of Critical Alerts ... 44
Filter Alerts ... 47
Investigate Events ... 48
Manage Top Alerts ... 50
Set Alert Feedback Status ... 52
View Modeled Behaviors ... 54
Reading an Indicator Chart ... 56
Types of Charts ... 57
Continuous Bar Chart ... 57
Dotted Chart ... 58
Time Chart ... 59
View NetWitness UEBA Metrics in Health and Wellness ... 60
Monitor Health and Wellness of UEBA ... 63
Access Kibana ... 63
Access Airflow ... 63
Kibana ... 64
Overview Dashboard ... 64
System Host overview ... 65
Adapter Dashboard ... 67
Support Dashboard Logical Time ... 68
Support Dashboard System Time ... 69
Scoring and Model Cache ... 70

Airflow 72
Reference 76
Overview Tab 76
Top Risky User or Network Entity Panel 77
Top Alerts Panel 78
Alerts Severity Panel 78
Entities Tab 79
Filters Panel 80
Risk Indicator panel 81
Entities List Panel 82
Alerts Tab 82
Filters Panel 83
Alerts Panel 84
User or Network Entity Profile View 85
Appendix: NetWitness UEBA Windows Audit Policy 93

Introduction

NetWitness User and Entity Behavior Analytics (UEBA) is an advanced analytics solution that empowers enterprise SOC managers and analysts to discover, investigate, and monitor risky behaviors across entities namely Users and Network (packets) in your environment.

NetWitness UEBA enables analyst to:

- Detect
  - malicious and rogue users
  - abnormal network traffic
- Identify high-risk behaviors
- Discover attacks
- Investigate emerging security threats
- Identify potential attacker's activity

This guide provides information and instructions for using the NetWitness UEBA functionalities and capabilities. It describes the key investigation methodologies, the main system capabilities, common use cases, and step-by-step instructions for the recommended workflow strategies.

Users

UEBA helps to analyze all users in your organization using logs and endpoint data for user activities, which is retrieved and parsed from the NetWitness Database (NWDB).

Network

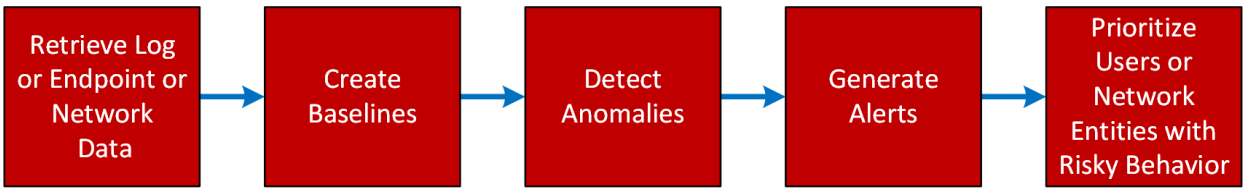
Note: Network entities are supported on NetWitness Platform 11.4 and later.

UEBA helps to analyze malicious outbound traffic masked within a legitimate HTTPS session. It can detect various network abnormalities, such as abnormal outbound traffic volume sent to a specific port, domain, organization or SSL Subject. The network (packet) data is retrieved and parsed from the NWDB into the new TLS data source, which supports two new entities: JA3 and SSL Subject. These entities validate the false negatives and true positives, and detect abnormal network traffic for JA3 and SSL Subject fingerprints.

- **JA3** - JA3 is a method of creating client-side SSL/TLS fingerprints to identify the client application initiating the session. The JA3 fingerprints perform JA3-signature-based analysis and detect abnormal network traffic, such as abnormal number of bytes sent over HTTPS.
- **SSL Subject** - The subject field of the certificate identifies the entity associated with the public key stored in the subject public key field, which is the entity for which the certificate was issued.

How NetWitness UEBA Works

NetWitness UEBA uses analytics to detect anomalies in the log and endpoint or network data to derive behavioral results from them. The following diagram displays the basic workflow:



The following table provides a brief description of each step.

Step	Description	More Information
1. Retrieve Log and Endpoint or Network Data	NetWitness UEBA retrieves logs or endpoint or network data from the NWDB and uses the data to create analytic results.	See Retrieve Data
2. Create Baselines	Baselines are derived from detailed analysis of normal user or network entity behavior, and are used as a basis for comparison to user or network entity behavior over time.	See Create Baselines
3. Detect Anomalies	An anomaly is a deviation of a user or network entity from the normal baseline behavior. NetWitness UEBA performs statistical analysis to compare each new activity to the baseline. User or network entity activities that deviate from expected baseline values are scored accordingly to reflect the severity of the deviation.	See Detect Anomalies
4. Generate Alerts	All the anomalies found in step 3 are grouped into hourly batches. Each batch is scored based on the uniqueness of its indicators. If the indicator composition is unique compared to a user or network entity's historic hourly batch compositions, it is likely that this batch is transformed into an alert.	See Generate Indicators and Generate Alerts
5. Prioritize User or Network Entities with Risky Behavior	NetWitness UEBA prioritizes the potential risk from a user or network entity by using a simplified additive scoring formula. Each alert is assigned a severity that increases a user or network entity's score by a predefined number of points. User or network entity with high scores either have multiple alerts, or alerts of high levels of severity associated with them.	See Prioritize User or Network Entity with Risky Behavior

Retrieve Data

NetWitness UEBA connects to a Concentrator service to retrieve log and endpoint data for the user entity or network data for the network entities. In case of multiple Concentrators, the NetWitness UEBA server connects to a Broker service. You can use the Broker service that is available on the NetWitness Platform Admin server if you do not have an exclusive Broker in your deployment. During NetWitness UEBA installation, the administrator specifies the IP address of the Broker service. For more information, see the "(Optional) Task 2 - Install NetWitness UEBA" topic in the *NetWitness Platform 11.5 Physical Host Installation Guide*

Note: In 11.4 and later, and when installed on a virtual machine, UEBA can process up to 20 million network events per day. For more information to resolve these issues, see the 'Troubleshooting UEBA Configurations' topic in the *UEBA Configuration Guide*.

Create Baselines

NetWitness UEBA uses machine learning to analyze multiple aspects of a user or network entity behavior within a stream of log and endpoint or network data and gradually builds a multi-dimensional baseline of typical behavior for each user or network entity.

Behavioral baselines are also created on a global level to describe common activities observed throughout the network. For example, if a working hour is abnormal for a user entity, but is not abnormal for the organization, the false-positive reduction algorithms decrease the impact on the alert score. Models are updated frequently and are constantly improving as time goes on.

Note: NetWitness UEBA requires 28 days of historical log and endpoint data for users and network data for network entities to create a proper baseline for all entities in your network. However, NetWitness recommends that you configure NetWitness UEBA to start baselining your data two months before the deployment date `<today-60days>`. The first 28 days are used for model training and are not scored. The remaining 32 days are leveraged to improve and update the model, and are also scored to provide initial value.

Note: For version 11.2 or later, there is limited support for environments with multiple domains. Distinct username values that are registered under different domains are normalized, and combined into one modeled entity. As a result, different users, who share the same username in different domains, will incorrectly be attributed to a single normalized entity.

Create Baselines for Users

NetWitness UEBA analyzes user actions to build a multi-dimensional baseline that reflects the typical behavior of the user. For example, the baseline can include information about the hours in which a user typically logs on.

Create Baselines for Network

NetWitness UEBA analyzes the network traffic pattern of JA3 or SSL Subject within a stream of network data to create a multi-dimensional baseline. For example, the baseline can be the allowed limit of data sent from an application or specific port that is connected to an application.

Detect Anomalies

The data is parsed hourly, to detect abnormal behavior. After establishing a behavioral baseline for all entities in your environment, each incoming event is compared to the baseline, to determine abnormalities. Based on the deviation the event is scored. The score is high if the deviation is strong and vice-versa. If anomalies are detected, they are turned into indicators that can be viewed on the user interface (UI).

For example, if a user's normal working hours are 9:00 AM to 5:00 PM, a new activity at 6:00 PM or 7:00 PM is not a strong deviation, and is probably not scored as an anomaly. However, an authentication at midnight is a strong deviation and is scored as an anomaly.

For example, in an organization, when a session is authenticated into a website for a SSL handshake, and communicates to five different ports or domains, it is not a strong deviation, and is probably not scored as an anomaly. But if the website communicates to an abnormal port or domain, it is a strong deviation. This indicates an abnormal behavior and is scored as an anomaly and triggers an alert.

Generate Indicators

If anomalies are detected, they are turned into indicators. NetWitness UEBA uses indicators to define validated anomalous activities. Indicators represent anomalies found in either a single event or multiple events batched over time.

User Indicators

User behavior or abnormal user activities, such as suspicious user logons, brute-force password attacks, unusual user changes, and abnormal file access are anomalous activities. Every anomalous activity is associated to an indicator. For more information, see Indicators for Users

Network Indicators

Network behavior or abnormal network traffic that contribute to data exfiltration or phishing, are examples of anomalous activities. Every anomalous activity is associated to an indicator. For more information, see Indicators for Network Entities.

Generate Alerts

All anomalies that are found are grouped into hourly batches by the user or network entity name. Each batch is scored based on the uniqueness of the composition of its indicators. If a composition is unique compared to the user or network entity's history, it is likely that this batch is transformed into an alert, and the anomalies into indicators. A high-scored batch of anomalies becomes an alert that contains valid indicators of compromise.

An abnormal activity by itself, even if it happens hundreds of times a day in a large corporate environment, does not necessarily reflect an account compromise. However, an abnormal behavior that occurs with a lot of other abnormal behaviors can indicate that the account is compromised and is an indication that additional analysis is required.

For example, if the following combination of one or more abnormal user or network behaviors occur, an alert is triggered.

Users

- Authentication from an abnormal computer.
- Multiple authentication attempts identified in a short time frame.
- Multiple files are deleted by this user from the corporate file share.
- Download or transfer files larger that the allowed limits.

Network

- Abnormal destination port for source netname.
- Abnormal organization for source netname.
- Abnormal traffic volume sent to organization.
- Abnormal traffic volume sent to port.

Note: The NetWitness UEBA User Interface can initially appear as empty because alerts are not generated until the baselines are established. If there is no historical audit data when NetWitness UEBA is enabled, the system starts generating the baselines from the time it is deployed, and requires 28 full days to elapse before generating new alerts. If historical audit data is processed when NetWitness UEBA is enabled, alerts appear after the historical data is processed, usually within two to four days.

Prioritize User or Network Entity with Risky Behavior

The entities scores are a primary tool for incident prioritization. The entities score is based on a simple additive calculation of an entity's alerts. Alerts and analyst feedback are the only factors in the entities score calculation, with the impact on the scores determined by their levels of severity. A unified color code is used for entities and alert scores:

Severity	Color	Score
Critical	Red	+20
High	Orange	+15
Medium	Yellow	+10
Low	Green	+1

Supported Sources

Log Sources

NetWitness UEBA natively supports the following data sources:

- Windows Active Directory in Version 11.2
- Windows Logon and Authentication Activity in Version 11.2
- Windows File Servers in Version 11.2
- Windows Remote Management in Version 11.3.2
- NetWitness Endpoint Process in Version 11.3
- NetWitness Endpoint Registry in Version 11.3
- SecurID Token in Version 11.3.1
- RedHat Linux in Version 11.3.1
- VPN Logs in Version 11.5
- Azure Active Directory Logs in Version 11.5

Network Sources

- TLS in Version 11.4

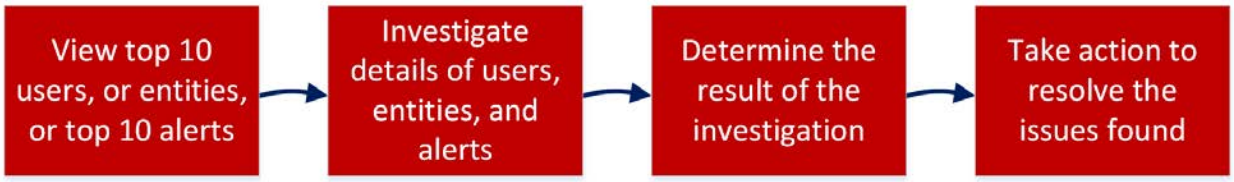
Recommended Workflows

To use NetWitness UEBA more effectively, there are two workflows - Detection and Forensic workflow.

Detection Workflow

The detection workflow gives you an overview of the health of your environment, and then focuses on investigating the top high-risk users, entities, and alerts that are displayed in the Overview tab.

The following flowchart illustrates the steps to follow for detecting suspicious behavior in your environment.



The following table describes each step in the workflow.

Step	Description	Instructions
View top ten users, or entities, or top 10 alerts,	In the Overview tab, note the users and network entity with the risky behaviors and the top most critical alerts.	Investigate High-Risk User or Network Entity and Investigate Top Alerts
Investigate details of users, entities, and alerts	Drill-down into detailed information about risky user or entity behaviors and critical alerts to determine the cause of these actions and how to resolve them.	Investigate High-Risk User or Network Entity and Investigate Events
Determine the result of the investigation	Analyze the summary information provided in the UI from the previous steps and identify focus areas on to resolve the issues.	Identify High-Risk User or Network Entity and Investigate Events
Take action to resolve the issues found	Target specific user or entity behaviors and events to address, and use results of this investigation to improve and sharpen future investigations.	Take Action on High-Risk User or Network Entity

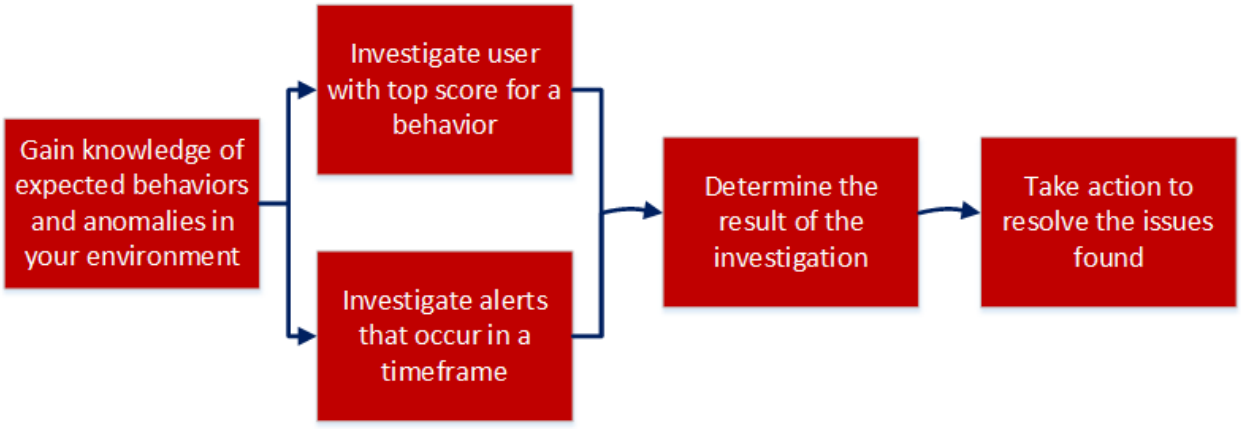
Forensic Workflow

The forensic workflow is recommended when you have an understanding of the typical user or entity behaviors and anomalies in your environment, and helps you focus on specific forensic information that is based on a user or entity behavior, or a specific time frame in which suspicious events occurred.

Using forensics information, analysts may determine actions and behaviors that the attacker is likely to attempt using the following questions:

- What fundamental techniques and behaviors are common across all intrusions?
- What evidence do these techniques leave behind?
- What do attackers do?
- What are normal behaviors of my accounts and entities?
- Which are my sensitive machines and where are they located?

The following flowchart illustrates how to perform investigation on forensic information that is based on a specific user or entity behavior, or a specific time frame in which suspicious events occurred.



The following table describes each step in the workflow.

Step	Description	Instructions
Gain knowledge of expected behaviors and anomalies in your environment	Establish a baseline of normal behaviors, expected anomalies, and unexpected anomalies, to focus on anomalies that are significant for your environment.	Retrieve Data, Detect Anomalies, and Generate Alerts.
Investigate a user or network entity with top score for a specific behavior	Select a user or network entity with a high score for a specific behavior and gather detailed information.	Investigate High-Risk User or Network Entity and Investigate Events.
Investigate alerts that occur in a specific time frame	Determine a time frame of interest, and in the Alerts tab, select that time frame to see detailed information about alerts that occurred during that period.	Investigate Events

Step	Description	Instructions
Determine the result of the investigation	Based on your knowledge of expected user or network entity behavior, focus on the indicators that are displayed during the specified time period and determine if the anomalies that were discovered need to be resolved.	Investigate Events and Identify High-Risk Entities
Take action to resolve the issues found	Target specific user or network entity behaviors and events to address, and use the results of this investigation to improve and sharpen future investigations.	Take Action on High-Risk User or Network Entity

NetWitness UEBA Use Cases

NetWitness UEBA focuses on providing advanced detection capabilities to guard enterprises from insider threats. These could either be compromised trusted users or network entity within a network, or alternatively, an external attacker malicious taking advantage of credentials acquired by using advanced account takeover techniques.

Identity theft typically begins with the theft of credentials, which are then used to obtain unauthorized access to resources and to gain control over the network. Attackers may also exploit compromised non-admin users to obtain access to resources for which they have administrative rights, and then escalate those privileges.

NetWitness UEBA helps you separate possibly malicious activity from the otherwise abnormal, but not risky, user or network entity actions.

Use Case for Users

An attacker who uses stolen credentials may trigger suspicious network events while accessing resources. Detecting illicit credential use is possible, but requires that you separate attacker activity from the high volume of legitimate events. The following use cases define certain risk types, and the corresponding system capabilities used for their detection. You can review the use cases, represented by their alert type and description, to gain an initial understanding of the related risky behavior of each use case. Using NetWitness UEBA, you can then drill down into the indicators that reflect the possibly risky user activities to learn more. For more information about NetWitness UEBA-supported indicators, see Indicators for Users. When anomalies are detected, they are compared to the baseline and compiled into hourly alerts. For more information on types of alerts for Users, see Alert Types for a User .

Use Case for Network Entities

UEBA can detect malicious traffic masked within a legitimate HTTPS session. Based on this alert analysis, the analyst can drill down to the indicators and determine if the activity was normal or not. For more information about NetWitness UEBA-supported entity indicators, see Indicators for Network Entities. For example, the analyst can detect if there was any abnormal number of bytes sent to a port or a domain. If this type of events or a combination of such events are detected an alert is triggered. For more information on types of alerts for network entity, see Alert Types for Network Entities.

Alert Types

Alert Types for a User

Alert Type	Description
Mass Changes to Groups	An abnormal number of changes are made to groups. Investigate which elements are changed, and decide if the changes were legitimate or possibly the result of risky or malicious behavior. This activity is associated with the **Multiple Group Membership Changes** indicator.
Multiple Failed Logons	In traditional password cracking attempts, the attacker tries to obtain a password through guesswork or by employing other low-tech methods to gain initial access. The attacker risks getting caught or being locked out by explicitly attempting to authenticate; but with some prior knowledge of the victim's password history, may be able to successfully authenticate. Look for additional abnormal indications that the account owner is not the one attempting to access this account. This activity is usually associated with the **Multiple Failed Authentications** indicator.
User Login to Abnormal Host	Attackers often need to reacquire credentials and perform other sensitive activities, like using remote access. Tracing the access chain backwards may lead to the discovery of other computers involved in possibly risky activity. If an attacker's presence is limited to a single compromised host or to many compromised hosts, that activity can be associated with the **Abnormal Host** indicator.
Data Exfiltration	Data exfiltration is the unauthorized copying, transfer, or retrieval of data from a computer or server. Data exfiltration is a malicious activity performed through various techniques, typically by cyber criminals over the Internet or other network. This activity can be associated with the **Excessive Number of File Rename Events**, **Excessive Number of Files Moved from File System**, and **Excessive Number of Files Moved to File System** indicators.
Mass File Rename	Ransomware is a type of malware that encrypts desktop and system files, making them inaccessible. Some ransom ware, for example, Locky, encrypts and renames files as part of their initial execution. Use this indication of mass-file-renaming to determine if your file system is infected with ransomware. This activity can be associated with the **Multiple File Rename Events** indicator.

Alert Type	Description
Snooping User	Snooping is unauthorized access to another person's or company's data. Snooping can be as simple as the casual observance of an e-mail on others computer, or watching what someone else is typing. More sophisticated snooping uses software programs to remotely monitor activity on a computer or network device. This activity can be associated with the **Multiple File Access Events**, **Multiple Failed File Access Events**, **Multiple File Open Events**, and **Multiple Folder Open Events** indicators.
Multiple Logons by User	All authentication activity, malicious or not, appears as normal logons. Therefore, administrators should monitor unexpected authorized activity. The key is that attackers use these stolen credentials for unauthorized access, which may provide an opportunity for detection. When an account is used for unusual activities, for example, authenticating an unusual amount of times, the account may have been compromised. This activity can be associated with the **Multiple Successful Authentications** indicator.
User Logged into Multiple Hosts	Attackers typically need to reacquire credentials periodically. This is because their key chain of stolen credentials naturally degrades over time, due to password changes and resets. Therefore, attackers frequently maintain a foothold in the compromised organization by installing backdoors and maintaining credentials from many computers in the environment. This activity can be associated with the **Logged onto Multiple Hosts** indicator.
Mass Permission Changes	Some credential theft techniques, for example, Pass-the-Hash, use an iterative, two-stage process. First, an attacker obtains elevated read-write permission to privileged areas of volatile memory and file systems, which are typically accessible only to system-level processes on at least one computer. Second, the attacker attempts to increase access to other computers on the network. Investigate if abnormal permission changes have taken place on the file systems to ensure that they were not compromised by an attacker. This activity can be associated with the **Multiple File Access Permission Changes**, **Multiple Failed File Access Permission Changes**, and **Abnormal File Access Permission Change** indicators.
Abnormal Active Directory (AD) Changes	If an attacker gains highly-privileged access to an Active Directory domain or domain controller, that access can be leveraged to access, control, or even destroy the entire forest. If a single domain controller is compromised and an attacker modifies the AD database, those modifications replicate to every other domain controller in the domain, and depending on the partition in which the modifications are made, the forest as well. Investigate abnormal changes conducted by admins and non-admins in AD to determine if they represent a possible true compromise to the domain. This activity can be associated with the **Abnormal Active Directory Change**, **Multiple Account Management Changes**, **Multiple User Account Management Changes**, and **Multiple Failed Account Management Changes** indicators.

Alert Type	Description
Sensitive User Status Changes	A domain or enterprise administrator account has the default ability to exercise control over all resources in a domain, regardless of whether it operates with malicious or benign intent. This control includes the ability to create and change accounts; read, write, or delete data; install or alter applications; and erase operating systems. Some of these activities trigger organically as part of the account's natural life cycle. Investigate these security sensitive user account changes, and determine if it is compromised. This activity can be associated with the **User Account Enabled**, **User Account Disabled**, **User Account Unlocked**, **User Account Type Changed**, **User Account Locked**, **User Password Never Expires Option Changed**, **User Password Changed by Non-Owner**, and **User Password Change** indicators.
Abnormal File Access	Monitor for abnormal file access to prevent improper access to confidential files and theft of sensitive data. By selectively monitoring file views, modifications and deletions, you can detect possibly unauthorized changes to sensitive files, whether caused by an attack or a change management error. This activity can be associated with the **Abnormal File Access Event** and **Multiple File Delete Events** indicators.
Non-Standard Hours	All authentication activity, malicious or not, appears as normal logons. Therefore, administrators should monitor unexpected authorized activity. The key is that attackers use these stolen credentials for unauthorized access, which may provide an opportunity for detection. When an account is being used for unusual activities, for example, authenticating an unusual number of times, the account may have been compromised. Use the indication of an abnormal activity time to determine if the account is taken over by an external actor. This activity can be associated with the **Abnormal File Access Time**, **Abnormal VPN Logon Time** **Abnormal Azure AD Logon Time**, **Abnormal Active Directory Change Time**, and **Abnormal Logon Time** indicators.
Credential Dumping	Credential dumping is the process of obtaining account login and password information, in the form of a hash or a clear text password, from the operating system and software. Credentials can then be used to perform lateral movement and access restricted information. This activity can be associated with the **Abnormal Process Created a Remote Thread in LSASS** indicator.

Alert Type	Description
Discovery & Reconnaissance	Discovery consists of techniques that allow the adversary to gain knowledge about the system and internal network. When attackers gain access to a new system, they must orient themselves to what they now have control of and what benefits operating from that system give to their current objective or overall goals during the intrusion. The operating system provides many native tools that aid in this post-compromise information-gathering phase. This activity can be associated with the **Abnormal Reconnaissance Tool Execute , Multiple Distinct Reconnaissance Tools Executed, Multiple Reconnaissance Tool Activities Executed** and **User Executed a Reconnaissance Tool Multiple Times** indicators.
PowerShell & Scripting	PowerShell is a powerful interactive command-line interface and scripting environment included in the Windows operating system. Attackers can use PowerShell to perform a number of actions, including discovery of information and execution of code. Examples include the Start-Process cmdlet which can be used to run an executable and the Invoke-Command cmdlet which runs a command locally or on a remote computer. This activity can be associated with the **User Ran an Abnormal Process to Execute a Scripting Tool, Abnormal Process Executed a Scripting Tool, Scripting Tool Triggered an Abnormal Application, User Ran a Scripting Tool that Triggered an Abnormal Application, User Ran a Scripting Tool to Open an Abnormal Process** and **Scripting Tool Opened an Abnormal Process** indicators.
Registry Run Keys & Start Folder	Adding an entry to the "run keys" in the Registry or startup folder will cause the program referenced to be executed when a user logs in. The program will be executed under the context of the user and will have the account's associated permissions level. Attackers can use these configuration locations to execute malware, such as remote access tools, to maintain persistence through system reboots. Attackers may also use Masquerading to make the Registry entries look as if they are associated with legitimate programs. Adding an entry to the "run keys" in the Registry or startup folder will cause the program referenced to be executed when a user logs in. These programs will be executed under the context of the user and will have the account's associated permissions level. This activity can be associated with the **Abnormal Process Modified a Registry Key Group** indicator.
Multiple Failed Authentications - External Access	As organizations increase their reliance on external authentication infrastructures, attackers may attempt to leverage these infrastructures to their advantage. Brute force techniques as well as more traditional password cracking methods like guesswork can be utilized to gain initial access. These activities can be associated with the **Multiple Failed Azure AD Authentications** and **Multiple Failed VPN Authentications** indicators.

Alert Type	Description
Abnormal Country	As organizations increase their reliance on external authentication infrastructures, attackers may attempt to leverage these infrastructures to their advantage. When devices or accounts are compromised as well as when credentials are wrongly shared, attackers may utilize them to gain initial access from an abnormal location. These activities can be associated with the **Abnormal Azure AD Logon Country** and **Abnormal VPN Logon Country** indicators.
Snooping User - Cloud Service Account	Snooping is unauthorized access to company data or data belonging to another person. Snooping can be as simple as the casual observance of an email on another person's computer. More sophisticated snooping uses software programs to remotely monitor activity on a computer or a cloud service account. This activity can be associated with the **Azure AD - Logon Attempts to Multiple Applications** indicator.
Abnormal Remote Application	Attackers may leverage compromised account details or devices to access remote applications that genuine end users do not frequently access to collect and even exfiltrate sensitive information. This activity can be associated with the **Azure AD - Abnormal Application** indicator.
Process Injection	Process injection is a method of executing arbitrary code in the address space of a separate live process. Running code in the context of another process may allow access to the process's memory, system/network resources, and possibly elevated privileges. Execution via process injection may also evade detection from security products since the execution is masked under a legitimate process. This activity can be associated with the **Abnormal Process Created a Remote Thread in a Windows Process** indicator.

Alert Types for Network Entities

Alert Type	Description
Phishing	Phishing is the fraudulent attempt to obtain sensitive information such as usernames, passwords and credit card details by disguising oneself as a trustworthy entity in an electronic communication. This activity can be associated with **Abnormal Country for SSL Subject**, and **Abnormal SSL Subject for JA3.** indicators.
Data Exfiltration	Data exfiltration is the unauthorized copying, transfer, or retrieval of data from a computer or server. Data Exfiltration is a malicious activity performed through various techniques, typically by cyber criminals over the Internet or other network. This activity can be associated with **Abnormal Traffic Volume Sent to Domain** and **Abnormal Traffic Volume Sent from JA3** indicators.

Alert Type	Description
Command & Control (C&C)	Command and control infrastructure can be leveraged by attackers as a communication channel between a compromised asset within the impacted network and an attacker-controlled server. Attackers may attempt to mask this malicious communication within regular network traffic; consequently, this activity can be associated with numerous network indicators such as **Abnormal Destination Port for JA3** and **High Number of IPs Contact a New SSL Subject**.

NetWitness UEBA Indicators

Indicators for Users

The following tables list indicators that display when a potentially malicious activity is detected for users.

Windows File Servers

Indicator	Alert Type	Description	
Abnormal File Access Time	Non-Standard Hours	A user has accessed a file at an abnormal time.	
Abnormal File Access Permission Change	Mass Permission Changes	A user changed multiple share permissions.	
Abnormal File Access Event	Abnormal File Access	A user has accessed a file abnormally.	
Multiple File Access Permission Changes	Mass Permission Changes	A user changed multiple file share permissions.	
Multiple File Access Events	Snooping User	A user accessed multiple file events.	
Multiple Failed File Access Events	Snooping User	A user failed multiple times to access a file.	
Multiple File Open Events	Snooping User	A user opened multiple files.	
Multiple Folder Open Events	Snooping User	A user opened multiple folders.	
Multiple File Delete Events	Abnormal File Access	A user deleted multiple files.	
Multiple Failed File Access Permission Changes	Mass Permission Changes	A user failed multiple attempts to change file access permissions	

Active Directory

Indicator	Alert Type	Description
Abnormal Active Directory Change Time	Non-Standard Hours	A user made Active Directory changes at an abnormal time.
Abnormal Active Directory Object Change	Abnormal AD Changes	A user made Active Directory attribute changes abnormally.
Multiple Group Membership Changes	Mass Changes to Groups	A user made multiple changes to groups successfully.
Multiple Active Directory Object Changes	Abnormal AD Changes	A user made multiple Active Directory changes successfully.
Multiple User Account Changes	Abnormal AD Changes	A user made multiple sensitive Active Directory changes successfully.
Multiple Failed Account Changes	Abnormal AD Changes	A user failed to make multiple Active Directory changes.
Admin Password Changed	Admin Password Change	The password of an admin was changed.
User Account Enabled	Sensitive User Status Changes	An account of a user was enabled.
User Account Disabled	Sensitive User Status Changes	An account of a user was disabled.
User Account Unlocked	Sensitive User Status Changes	An account of a user was unlocked.
User Account Type Changed	Sensitive User Status Changes	The type of user was changed.
User Account Locked	Sensitive User Status Changes	An account of a user was locked.
User Password Reset	Sensitive User Status Changes	The password of a user was reset.
User Password Never Expires Option Changed	Sensitive User Status Changes	The password policy of a user was changed.

Logon Activity

Indicator	Alert Type	Description
Abnormal Remote Computer	Abnormal Computer Access	A user accessed a remote computer abnormally.

Indicator	Alert Type	Description
Abnormal Logon Time	Non-Standard Hours	A user logged on at an abnormal time.
Abnormal Computer	User Login to Abnormal Host	A user attempted to access a computer abnormally.
Multiple Successful Authentications	Multiple Logons by User	A user logged on multiple times.
Multiple Failed Authentications	Multiple Failed Logons	A user failed multiple authentication attempts.
Logon Attempts to Multiple Source Computers	User Logged into Multiple Hosts	A user attempted to log on from multiple computers.

Process

Indicator	Alert Type	Description
Abnormal Process Created a Remote Thread in LSASS	Credential Dumping	An abnormal process was created into the LSASS process.
Abnormal Reconnaissance Tool Executed	Discovery and Reconnaissance	An abnormal process was executed.
Abnormal Process Executed a Scripting Tool	PowerShell and Scripting	An abnormal process executed a scripting tool.
Abnormal Process Executed a Scripting Tool	PowerShell and Scripting	An abnormal process was triggered by a scripting tool.
Scripting Tool Triggered an Abnormal Application	PowerShell and Scripting	An abnormal process was opened by a scripting tool.
Abnormal Process Created a Remote Thread in a Windows	PowerShell and Scripting	An abnormal process was injected into a known windows process .
Multiple Distinct Reconnaissance Tools Executed	Discovery and Reconnaissance	Multiple reconnaissance tools were executed in an hour.
Multiple Reconnaissance Tool Activities Executed	Discovery and Reconnaissance	Multiple reconnaissance tool activities were executed in an hour.
User Ran an Abnormal Process to Execute a Scripting Tool	PowerShell / Scripting	An abnormal process executed a scripting tool.
User Ran a Scripting Tool that Triggered an Abnormal Application	PowerShell / Scripting	A scripting tool was executed that triggered an abnormal application.
User Ran a Scripting Tool to Open an Abnormal Process	PowerShell / Scripting	A scripting tool was executed to open an abnormal process.

Registry

Indicator	Alert Type	Description
Abnormal Process Modified a Registry Key Group	Registry Run Keys	An abnormal process modified a service key registry.

Indicators for Network Entities

The following tables list indicators that display when a potentially malicious activity is detected for JA3 and SSL Subject entities.

Note: Indicators are for JA3, and in some instances the JA3 hash can be mapped to more than one client application.

Indicator	Entity Type	Alert Type	Description
Abnormal Traffic Volume Sent from IP to SSL Subject	SSL Subject	Data exfiltration	An IP address in the organization sent an unexpectedly high amount of data to an SSL Subject.
Abnormal Traffic Volume Sent from IP to Domain	SSL Subject	Data exfiltration	An IP address in the organization sent an unexpectedly high amount of data to a domain and SSL Subject.
Abnormal Traffic Volume Sent from IP to Organization	SSL Subject	Data exfiltration	An IP address in the organization sent an unexpectedly high amount of data to an organization and SSL Subject.
Abnormal Traffic Volume Sent from IP to Port	SSL Subject	Data exfiltration	An IP address in the organization sent an unexpectedly high amount of data to a port and SSL Subject.
Abnormal Traffic Volume Sent to SSL Subject	SSL Subject	Data exfiltration	An unexpectedly high amount of data was sent to an SSL Subject.
Abnormal Traffic Volume Sent to Domain	SSL Subject	Data exfiltration	An unexpectedly high amount of data was sent to a domain and SSL Subject.
Abnormal Traffic Volume Sent to Port	SSL Subject	Data exfiltration	An unexpectedly high amount of data was sent to a port and SSL Subject.
Abnormal Traffic Volume Sent to Organization	SSL Subject	Data exfiltration	An unexpectedly high amount of data was sent to an organization and SSL Subject.
Abnormal Traffic Volume Sent from JA3	JA3	Data exfiltration	Abnormal number of bytes sent from JA3 .

Indicator	Entity Type	Alert Type	Description
High Number of IPs Use JA3	JA3	C&C	An abnormally high number of IPs use JA3.
Abnormal SSL Subject for Source Netname	SSL Subject and JA3	Phishing	A source netname contacted an abnormal SSL Subject.
Abnormal Domain for Source Netname	SSL Subject and JA3	Phishing	A source netname contacted an abnormal domain
Abnormal Destination Port for Source Netname	SSL Subject and JA3	C&C	A source netname contacted an abnormal destination port.
Abnormal Organization for Source Netname	SSL Subject and JA3	Phishing	A source netname contacted an abnormal organization.
Abnormal Country for SSL Subject	SSL Subject and JA3	Phishing	An SSL Subject was contacted with an abnormal destination country.
Abnormal Destination Port for SSL Subject	SSL Subject and JA3	C&C	An SSL Subject was contacted through an abnormal destination port.
Abnormal Time for SSL Subject	SSL Subject and JA3	Non-Standard Hours	An SSL Subject was contacted at an abnormal time.
Abnormal Destination Port for Domain	SSL Subject and JA3	C&C	A domain was accessed through an abnormal destination port.
Abnormal Destination Port for Organization	SSL Subject and JA3	C&C	An organization was accessed through an abnormal destination port.
Abnormal Time for JA3	SSL Subject and JA3	Non-Standard Hours	JA3 was used at an abnormal time.
Abnormal JA3 for Source Netname	SSL Subject and JA3	C&C	A source netname utilized an abnormal client application.
Abnormal SSL Subject for JA3	SSL Subject and JA3	Phishing	JA3 contacted an abnormal SSL Subject.
Abnormal Domain for JA3	SSL Subject and JA3	Phishing	JA3 contacted an abnormal domain.
Abnormal Destination Port for JA3	SSL Subject and JA3	C&C	JA3 contacted an abnormal destination port.
High Number of IPs Contact a New SSL Subject	SSL Subject	C&C	High number of IPs contacted SSL Subject.
High Number of IPs Contact a New Domain	SSL Subject	C&C	High number of IPs contacted a new domain.

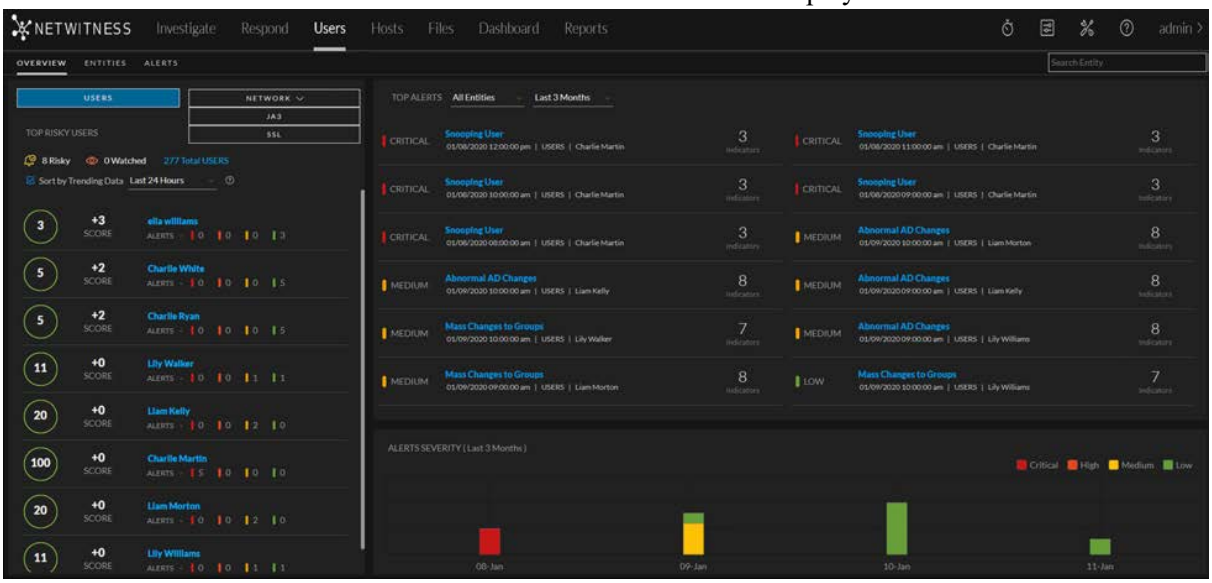
Indicator	Entity Type	Alert Type	Description
High Number of IPs Contact a New Organization	SSL Subject	C&C	High number of IPs contacted a new organization.
High Number of IPs Contact a New Port	SSL Subject	C&C	High number of IPs contacted a new port.
Abnormal Traffic Volume Sent from an IP to a New SSL Subject	SSL Subject	Data Exfiltration	Abnormal number of bytes sent from an IPs to an SSL Subject.
Abnormal Traffic Volume Sent from an IP to a New Domain	SSL Subject	Data Exfiltration	Abnormal number of bytes were sent an IP to a domain.
Abnormal Traffic Volume Sent from an IP to a New Port	SSL Subject	Data Exfiltration	Abnormal number of bytes were sent from an IP to a port.
Abnormal Traffic Volume Sent from an IP to a New Organization	SSL Subject	Data Exfiltration	Abnormal number of bytes were sent from an IP to an organization.
Abnormal Traffic Volume Sent to a New SSL Subject	SSL Subject	Data Exfiltration	Abnormal number of bytes were sent to a SSL Subject. n
Abnormal Traffic Volume Sent to a New Domain	SSL Subject	Data Exfiltration	Abnormal number of bytes were sent to a new domain.
Abnormal Traffic Volume Sent to a New Port	SSL Subject	Data Exfiltration	Abnormal number of bytes were sent to a new port.
Abnormal Traffic Volume Sent to a New Organization	SSL Subject	Data Exfiltration	Abnormal number of bytes were sent to an organization for an SSL Subject.
Abnormal Traffic Volume Sent from a New JA3	JA3	Data Exfiltration	Abnormal number for bytes were sent to JA3.

Access NetWitness UEBA

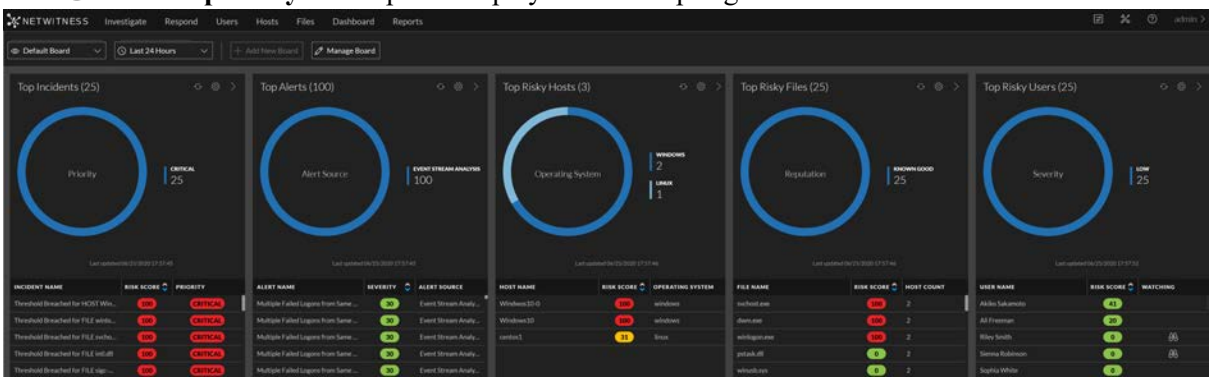
Note: To access the NetWitness UEBA service and Users tab, you must be assigned to either the UEBA_Analyst role or Administrators role. For information about how to assign these roles, see the "How Role-Based Access Control Works" topic in the *System Security and User Maintenance Guide*

To access NetWitness UEBA, log in to NetWitness and do one of the following:

1. Go to **Users** > **Overview** to view the NetWitness UEBA feature displayed.



2. Click ❯ in the **Top Risky Users** panel displayed on the Springboard to view the **Users** tab.



You can choose a dark or a light theme for the view. For more information, see the "Choose the Appearance of NetWitness" topic in the *NetWitness Getting Started Guide*.

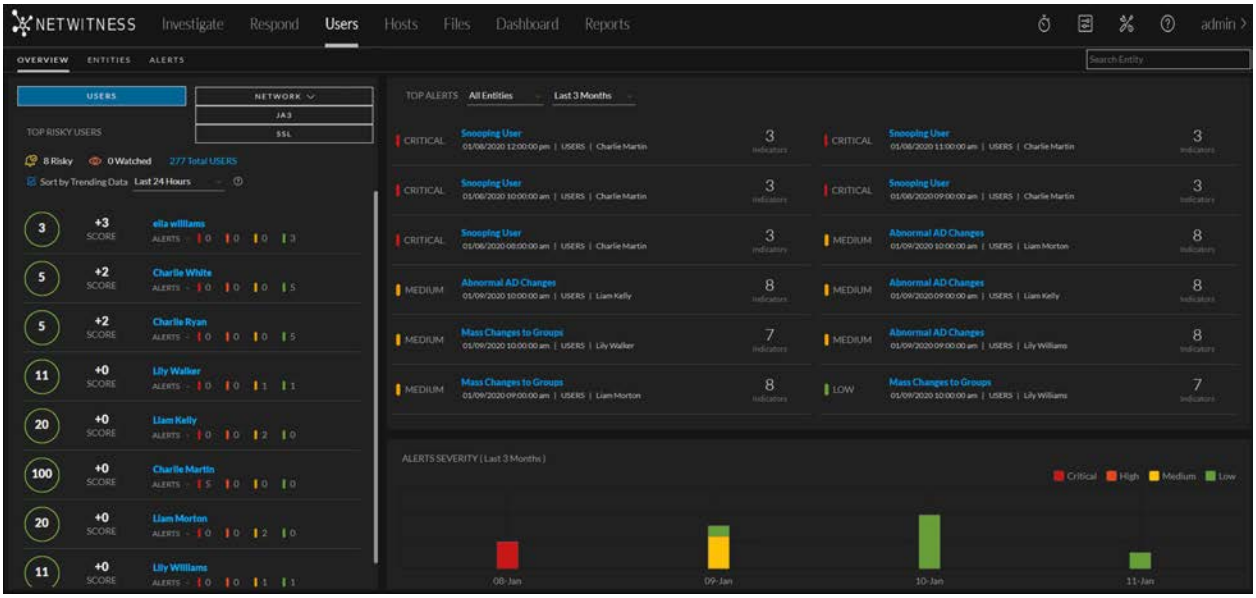
UEBA Licensing

You must alo ensure that you have NetWitness UEBA licensing configured. For information about NetWitness UEBA licensing, see the "User and Entity Behavior Analytics License" topic in the *Licensing Management Guide*.

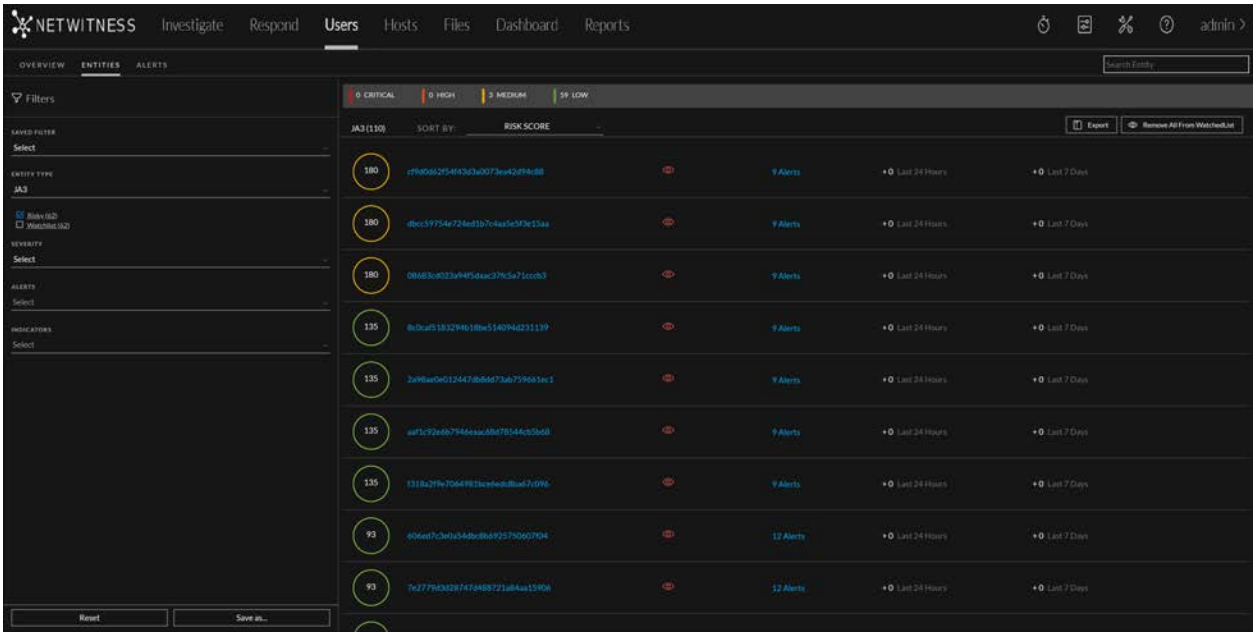
Investigate High-Risk Entities

An entity score is built based on the alert score and the alert severity. Using the entity score, you can identify entities that require immediate attention, perform deeper investigation, and take required action. You can identify high-risk entities from either the **Overview** tab or the **ENTITIES** tab.

The following figure is an example of top ten High-Risk entities in the **Overview** tab.



The following figure is an example of all the risky entities in your environment in the **ENTITIES** tab.



The following is a high-level process to investigate high-risk entities in your environment.

1. Identify high-risk entities. You can identify high-risk users using the following ways:
   - The **Overview** tab shows the top ten risky entities in your environment. From the listed entities identify the entities with a critical severity or entity score more than 100.
   - The **ENTITIES** tab shows all the risky entities in your environment, you can sort by Risk Score (default), Trending Data Last Day or Last Week (marked with +), Name, Alerts. Identify how many entities are marked Critical, High and Medium or based on the forensic investigation, identify malicious entity behavior and build use-case driven target entity lists using behavioral filters. Additionally, you can also use different types of filters (Risky or Watchlist) to identify targeted group of high-risk entities.

   Note: The investigation should mostly focus on Critical, High and Medium severities. Low scoring users are not typically worth much investigation.

   Hover over the number of alerts associated with the risky entities to quickly see what the alerts are and determine if there is a good mix.

   For more information, see the Identify High-Risk Entities topic.
2. In the **User Profile** view, investigate the alerts and indicators of the user.
   a. Review the list of alerts associated with the user and the alert score for each alert, sorted by severity.
   b. Expand the alert names to identify a threat narrative. The strongest contributing indicator determines the alert's name that suggests why this hour is flagged.
   c. Use the alert flow timeline to understand the abnormal activities.
   d. Review each indicator associated with the alert to see the details about the indicator, including the timeline in which the anomaly occurred. Also, you can further investigate the incident using external resources such as SIEM, network forensics, directly reaching out to the user or a managing director and so on.

   For more information, see the Begin an Investigation of High-Risk Entity topic.
3. On completion of the investigation, you can record your observation as follows:
   a. Specify if an alert is not a risk.
   b. Save the behavioral profile for the use case found in your environment.
   c. If you want to keep a track of user activity, you can add users to the watchlist, and watch user profile.

   For more information, see the Take Action on High-Risk Users topic.

Identify High-Risk Entities

You can identify high-risk user in your environment in the following ways:

- View top five high-risk entities
- View all the high-risk entities

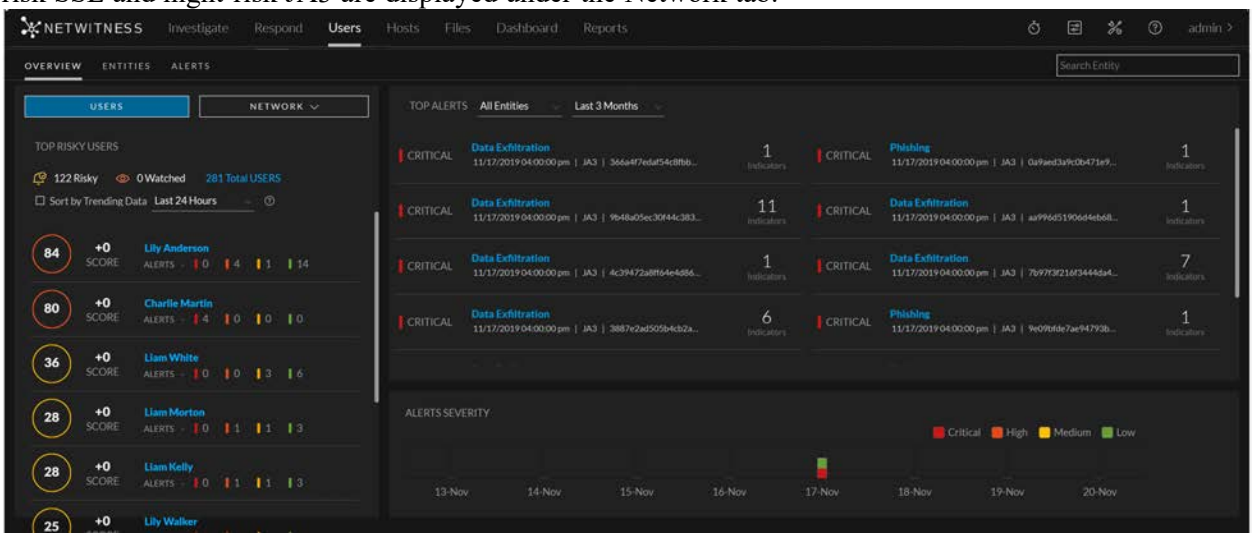
- View users of a specific group
- View users and other entities based on forensic investigation

View Top Ten Risky Entities

In the **OVERVIEW** tab, you can view the list of top five high-risk entities in your environment along with the risky score.

To view the top risky entities:

Log into **NetWitness** and go to **Investigate** > **ENTITIES**.
The Overview tab is displayed with the high-risk users displayed in the High Risk Users tab, and high-risk SSL and hight-risk JA3 are displayed under the Network tab.



View All High-Risk Entities

In the **Network** tab, you can view the list of all the high risk users in your environment along with the user score and total number of alerts associated with the users.

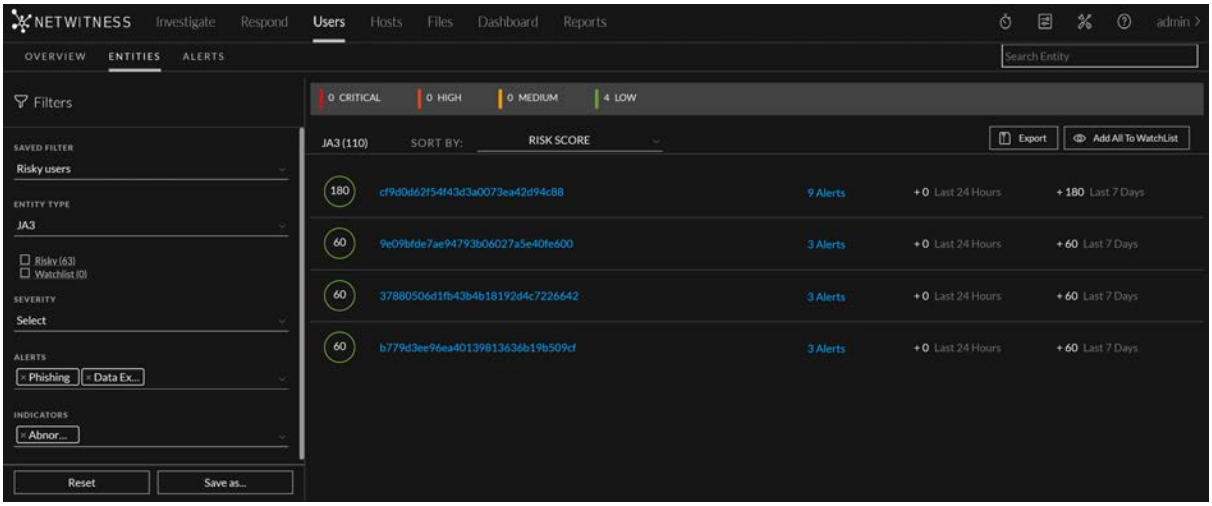
To view all high-risk users:

1. Log into **NetWitness** and go to **Investigate** > **ENTITIES**.

   The Overview tab is displayed.

2. Click **ENTITIES** tab.

The list of all high-risk entities are displayed.



View Entities of Specific Group

In the **Network** tab, you can use different types of filters to identify targeted group of high-risk entity.

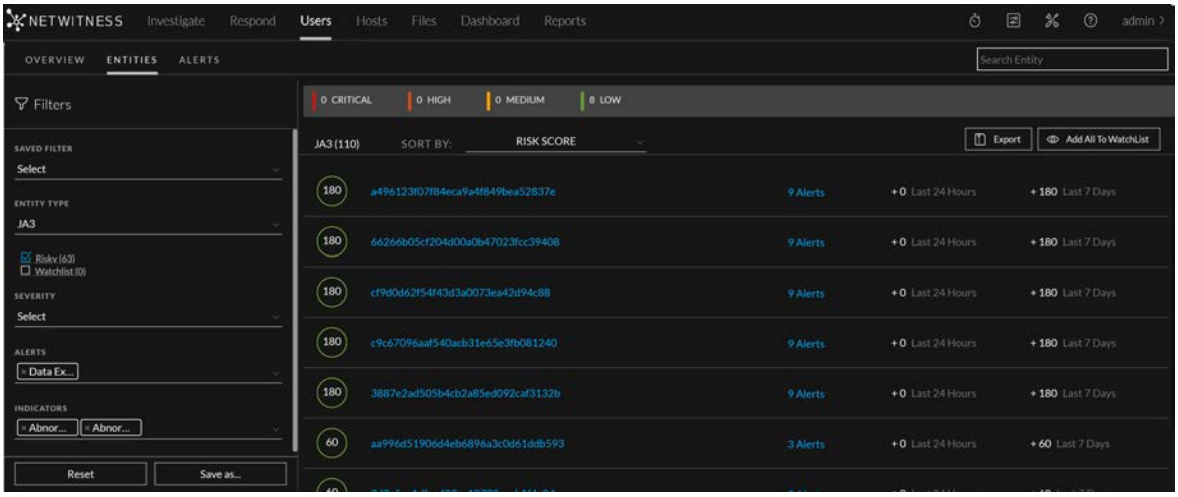
To view users of specific group:

1. Log into **NetWitness** and go to **Investigate** > **ENTITIES**.

   The Overview tab is displayed.

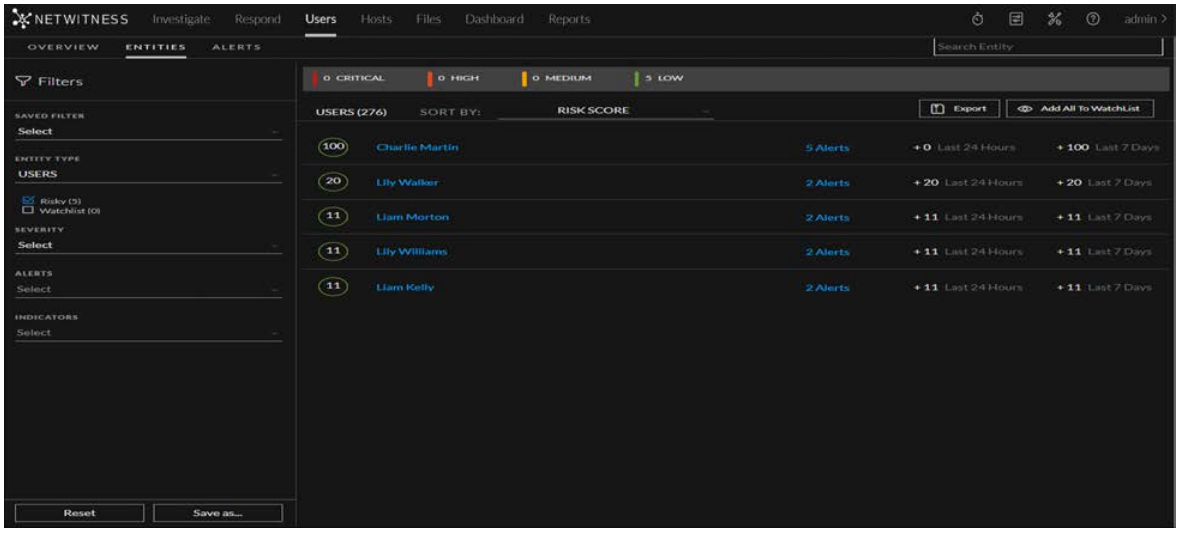
2. Click the **ENTITIES** tab.
3. In the **Filters** panel, do any of the following:

   - **Risky Entities**: To view all the risky entities in your environment, select **Risky** . By default, risky entities along with their risky score are displayed.



   - **Watchlist** : To view the list of entities that you added to the watchlist to monitor for specific

changes, select **Watchlist**.



Note: You can view users of one or more group by selecting one or more filters. For example, if you want to view the list of admin users who are risky users, select the **Admin Users** and **Risky Users** filters.

View Entity Based on Forensic Investigation

In the **ENTITIES** tab, you can use Alert Types and Indicators which are behavioral filters to view high-risk users based on forensic investigation. For more information on forensic investigation, see *Forensic Workflow* in the Introduction topic.

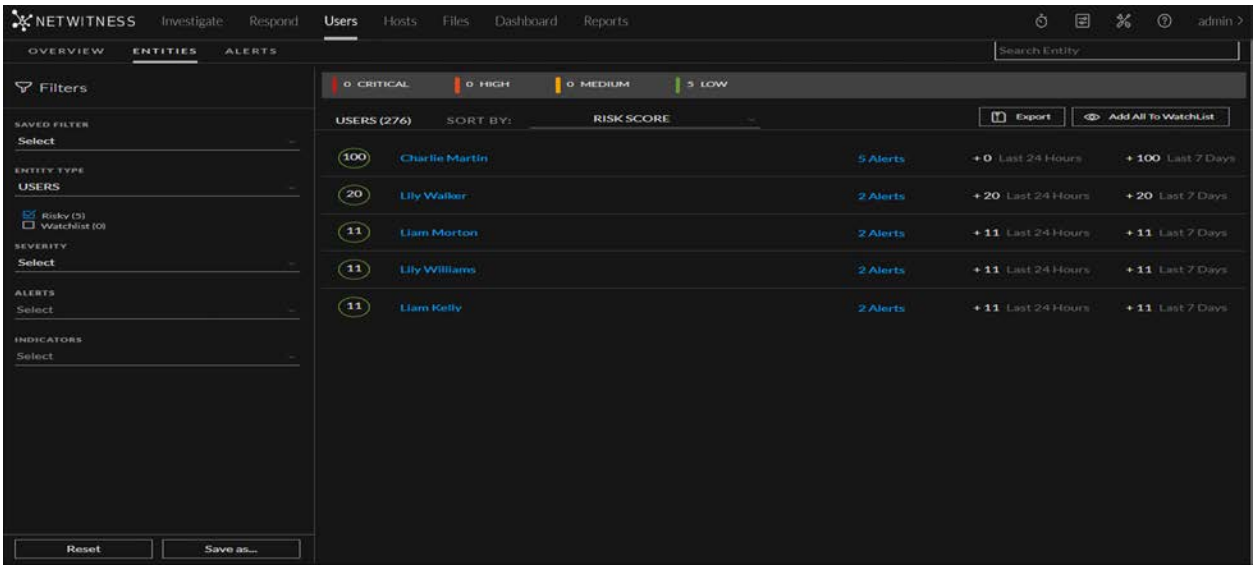
To view users based on specific forensic investigation:

1. Log into **NetWitness** and go to **Investigate** > **ENTITIES**.

   The Overview tab is displayed.

2. Click **ENTITIES** tab.
3. To create a behavioral filter using alert types, select one or more alerts in the **ALERTS** drop-down list.
4. To create a behavioral filter using indicators, select one or more indicators in the **INDICATORS** drop-down list.
5. To filter the result for JA3 entity, select JA3 from the **ENTITY TYPE** drop-down list.

Note: You can select combination of one or more alert types and indicators to create a behavioral filter based on your requirement. For example, to monitor abnormal access to confidential files and theft of sensitive data, you can create a behavioral filter with Alert Types = **Data Exfiltration** and Indicators = **Abnormal JA3 for Source Netname (3 JA3)**.



To save these behavioral filters as favorites for future investigation, click **Save as...**.

To delete the filters click **Reset Filters**.

Similarly, you can view the results for the SSL entity based on forensic investigation.

Begin an Investigation of High-Risk Entity

After identifying the high-risk entities, you can begin the investigation of high-risk entities.

To investigate high-risk entities:

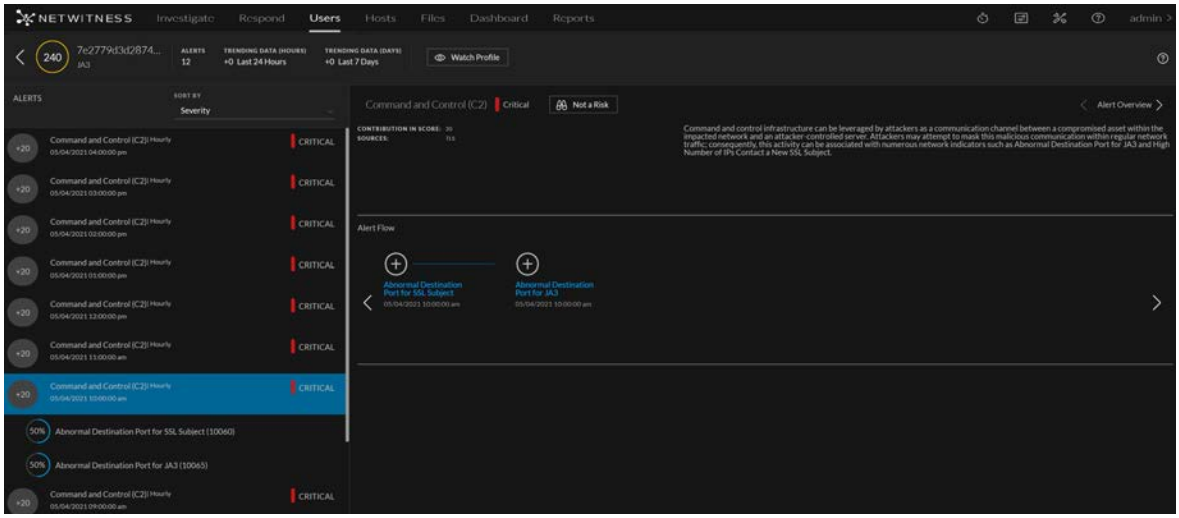
1. Log into **NetWitness** and go to **INVESTIGATE** > **ENTITIES**.

   In the **Overview** tab, select **ENTITIES** tab, select an entity from the **NETWORK** drop-down to investigate. For example, if you select **Network** > **JA3**, all the JA3 high-risk entities will be displayed. Also, can sort the results by selecting the **Sort by Trending Data** checkbox. If you select this option, the data will be sorted by the Entity score (marked with +) that changed in the past 24 hours or in the past one week. By default, the result is sorted by the Entity risk score.

2. To further investigate the alert of the entity, click an alert in the **High Risk JA3** panel. The following information is displayed:

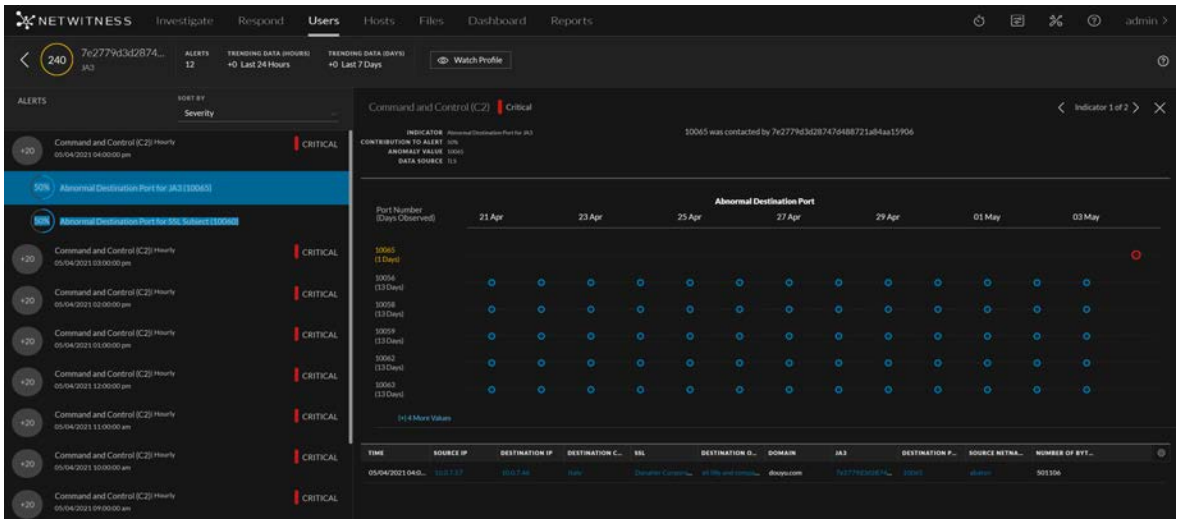
   - The alert name
   - The timeframe of the alert (Hourly or Daily)
   - The severity level icon
   - The contribution to the entity score value (for example, 20)
   - The data sources for the alert (for example, TLS)
     The middle panel is the Alert Flow panel. This panel provides a timeline of events that are related to the formation of the alert. The timeline of events can help to determine if the alert is an actual

risk.



3. To investigate the indicators associated with an alert for an entity, in the **High Risk** panel, select an alert and then select an indicator. The following information is displayed:

   - The indicator name and a description of the indicator type
   - Contribution to Alert
   - The anomaly values
   - The data source of the events found in the indicator
     The central panel display changes depending on which indicator is selected.



Note: You can investigate a high risk SSL entity using the above procedure.

For more information on how to read the indicator charts, see Reading an Indicator Chart topic.

Take Action on High-Risk Users

After investigation, you can take action on the risky users to reduce or prevent further damage caused by malicious attackers in your organization. You can take any of the following actions:

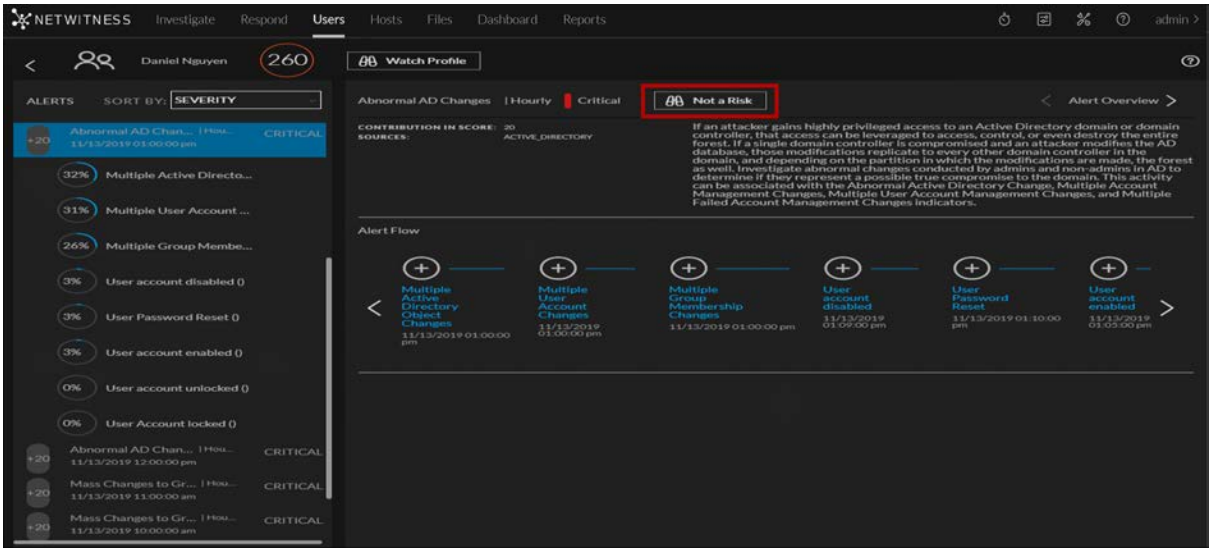
- Specify if the alert is not risky
- Save the behavioral profile for the use case found in your environment
- Add users to the watchlist, and the watch user profile, if you want to keep a track of the user activity

Specify that an alert is not risky.

If and alert is not a risk, you can mark it so that the user score for the user is automatically reduced.

To specify if the alert is not risky:

1. Log into **NetWitness** and go to **Investigate** > **ENTITIES**.
2. Take action on the users from any of the following tabs:
   a. In the **OVERVIEW** tab, in the **High Risk Users** panel, select a user and click either on the username or on the user score.
   b. In the **ENTITIES** tab, select a user and click on the username.
   The User Profile view is displayed.
3. If the alert is not a risk, you can specify by clicking **Not a Risk**.



When an alert is marked as **Not a Risk**, the user score is reduced automatically.

Save Behavioral Profile

The combination of the alert types and indicators you select during the forensics investigation is a behavioral profile. You can save the behavioral profile, so you can monitor this use case in future.

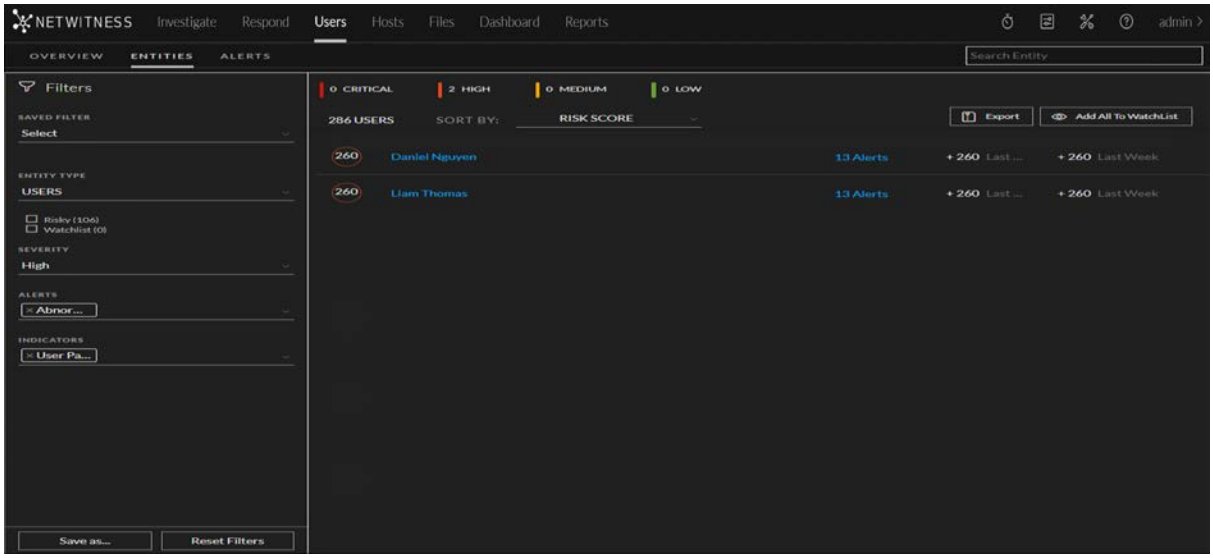
For example, if your organization is attacked and the attackers penetrated by brute forcing user accounts, you can select filters using the brute force alert type. This can be saved as favorite. You can proactively monitor for future brute force attempts. To do so, you can click the favorite to see if new users were subjected to this type of attack.

To save a behavioral profile:

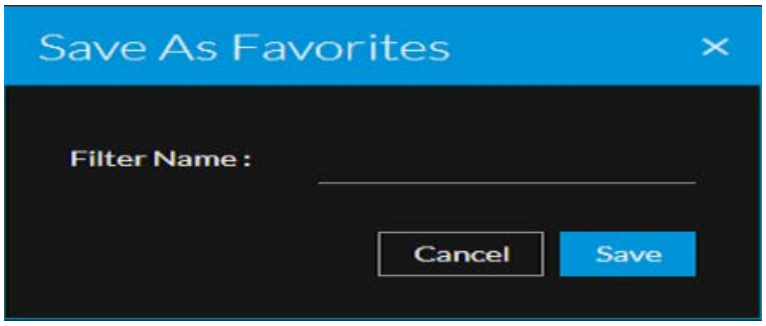
1. Log into **NetWitness** and go to **Investigate** > **ENTITIES**.

   The Overview tab is displayed.

2. Click the **Users** tab.
3. In the **Filters** panel, select the alert in the **Alert Type** drop-down and Indicators in the **Indicators** drop-down.
4. Click **Save to Favorites**.



5. In the **Save Filter** dialog, enter the name of the filter and click **OK**.



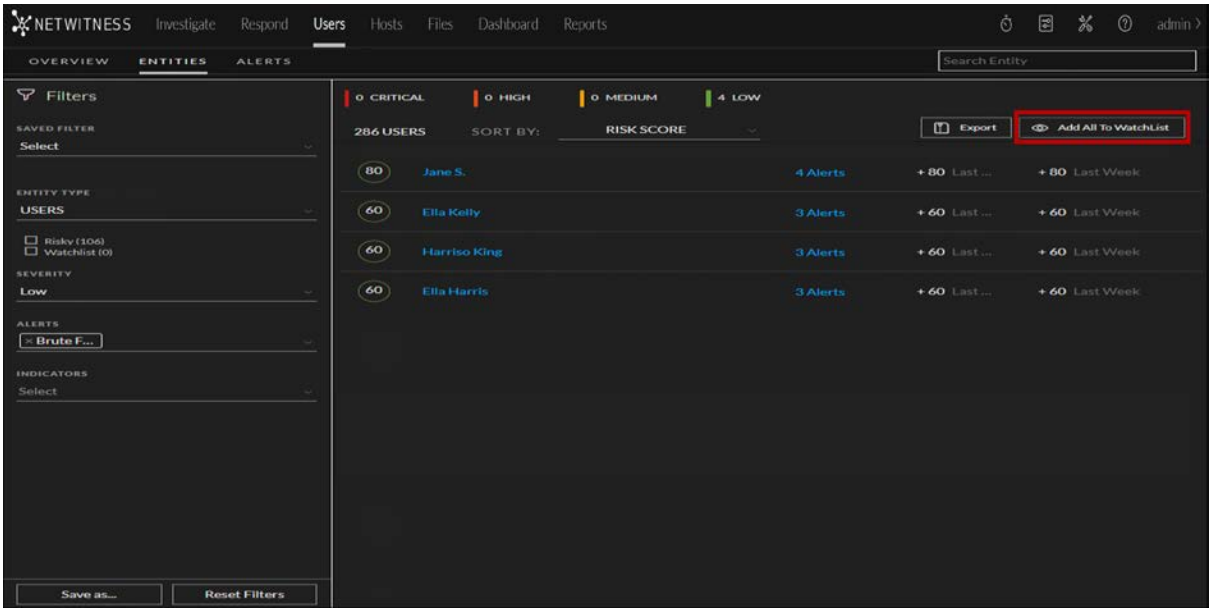
The behavioral profile is saved and displayed in the Favorites panel. You can click on the profile in the Favorites to monitor the users.

Add All Users to the Watchlist

If you want to keep track of users with recent activity but do not want to follow up with an immediate investigation, you can add the users to the watchlist and revisit over time to see if the risk score is elevated.

To add all users to the watchlist:

1. Log into **NetWitness** and go to **Investigate** > **ENTITIES**.
   The Overview tab is displayed.
2. Select the **ENTITIES** tab.
3. Select the users of specific categories using filters.
4. Click **Add All to Watchlist**.



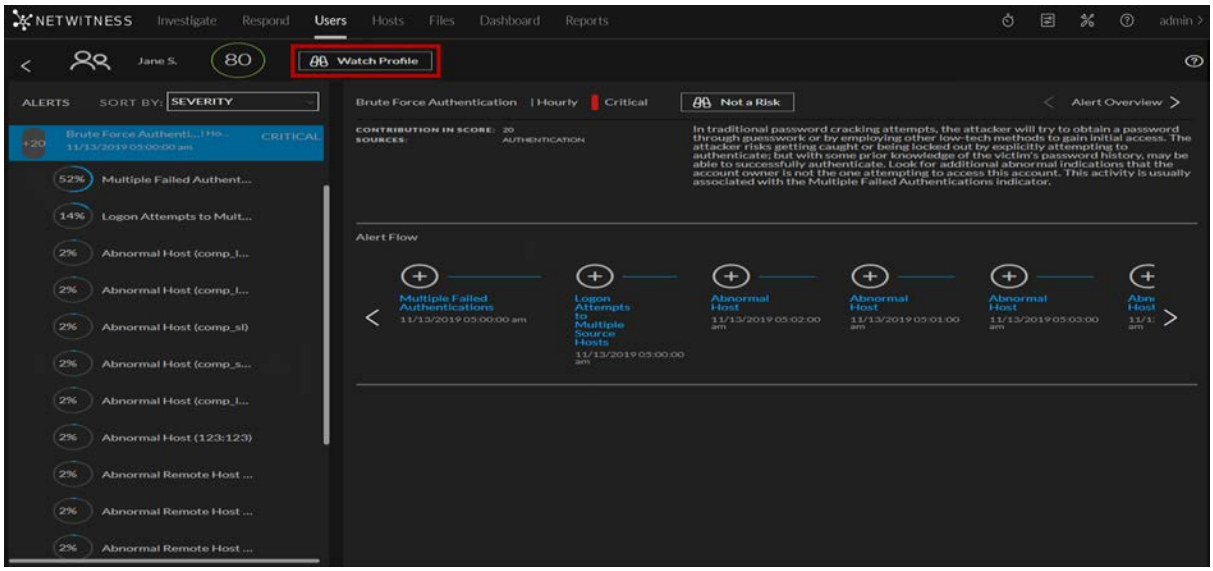
The list of users are added to the watchlist.

Watch Profile

The watch user profile is a list of users that you want to monitor for potential threats. The watch user profile marks a user so that the users can be quickly referenced on the dashboard. This is essentially a bookmark to monitor suspicious users.

To watch user profile:

1. Log into **NetWitness** and go to **Investigate** > **ENTITIES**. Do any of the following:
   a. In the **Overview** tab, under **High Risk Users** panel, select a user and click on either the username or the user score.
   b. In the **Users** tab, select a user and click the username.
   The User Profile view is displayed.
2. Click **Watch Profile** in the upper right corner of the User Profile.



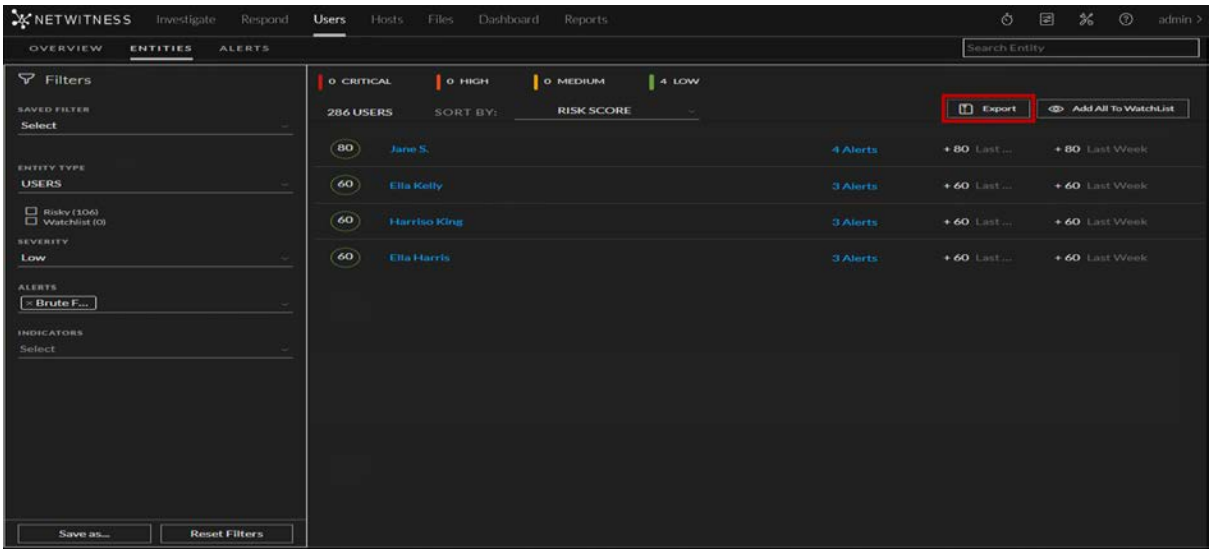
The user is added to the watchlist.

Export a list of High-Risk Users

You can export a list of all users and their scores in a .csv file format. You can use this information to compare with other data analysis tools like tableau, powerbi, and zeppelin.

To export a list of high-risk users:

1. Go to **INVESTIGATE** > **Users**.
   The Overview tab is displayed.
2. Select the **Users** tab.
3. Click **Export**.

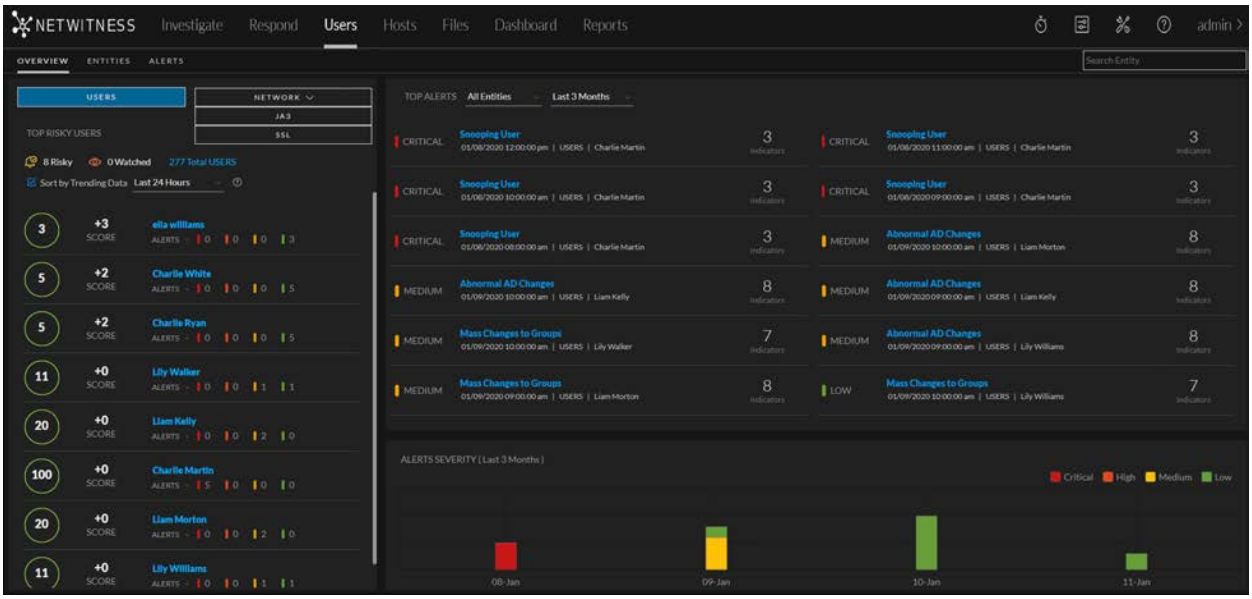


The list of all users and the associated user score is downloaded in the .csv file format.

Investigate Top Alerts

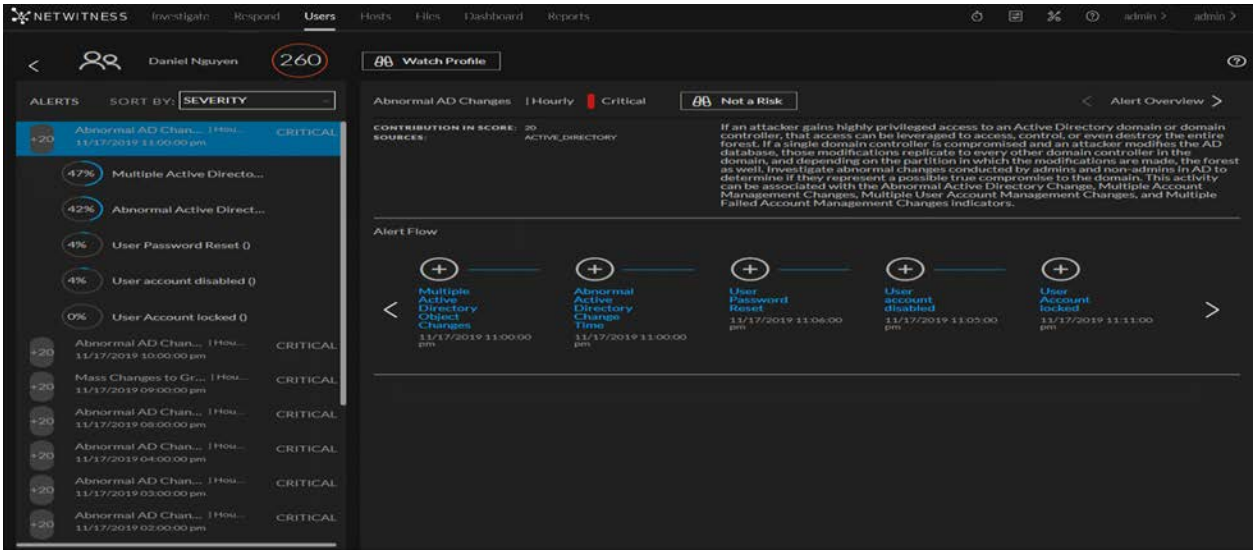
Anomalies that are found as incoming events are compared to the baseline and compiled into hourly alerts. Relatively strong deviations from the baseline, together with a unique composition of anomalies, are more likely to get a higher alert score.

You can quickly view the most critical alerts in your environment, and start investigating them from either the OVERVIEW tab or the ALERTS tab. The following figure is an example of top alerts in the OVERVIEW tab. The alerts are listed in order of severity and the number of users who generate the alerts.

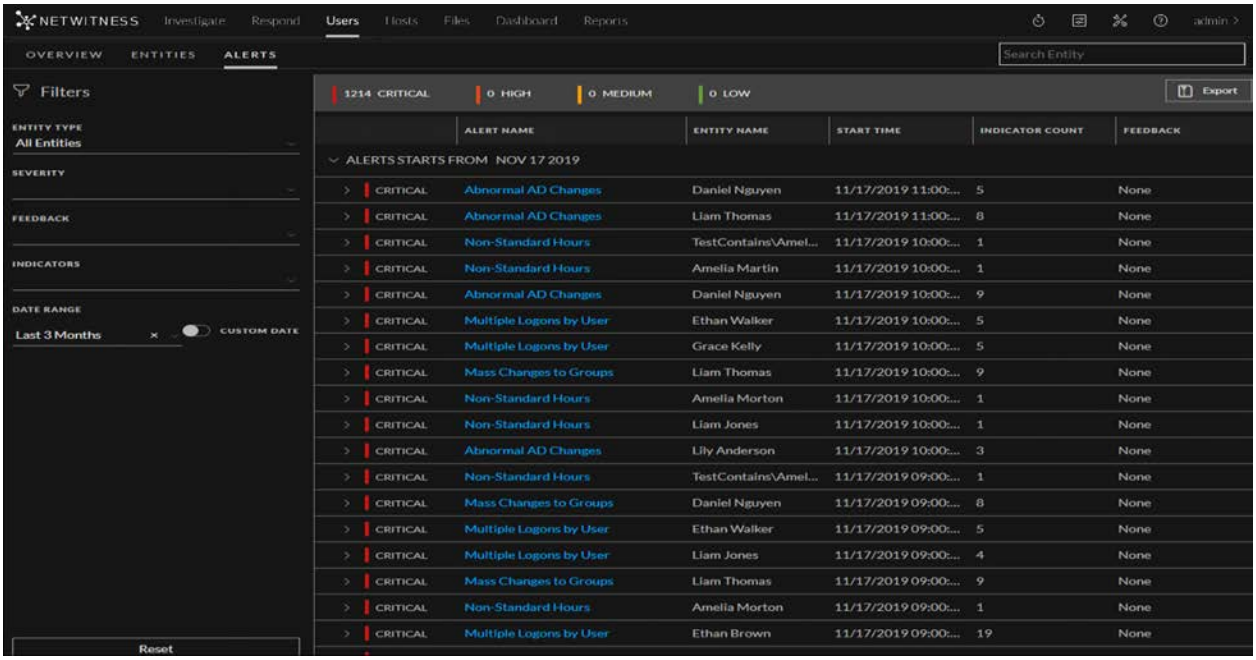


To investigate an alert on this page, click an alert in the **Top Alerts** section.

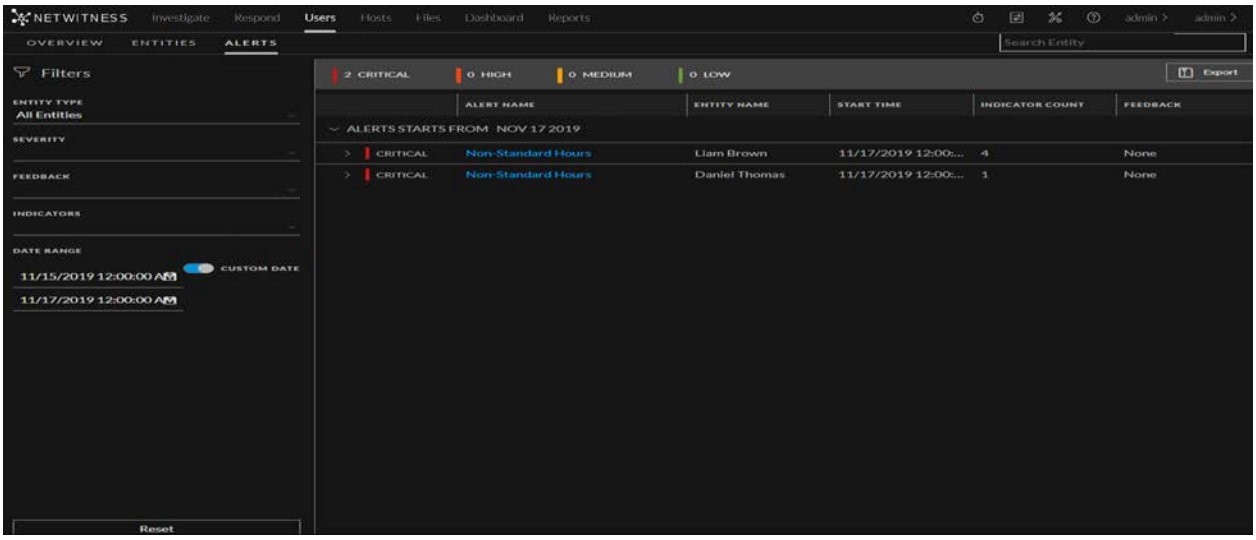
The following figure shows details about the event that caused the alert, and the time frame in which it occurred.



From the Alerts Severity panel at the bottom of the Overview tab, you can click on a bar in the graph to review top alerts in the ALERTS tab. The following figure shows the top alerts listed in the Alerts tab.



Investigating alerts is particularly useful when you want to focus on a timeframe in which you believe your systems were compromised. You can view forensic information based on a timeframe and gather detailed information about events that occurred during that time in the Alerts tab.

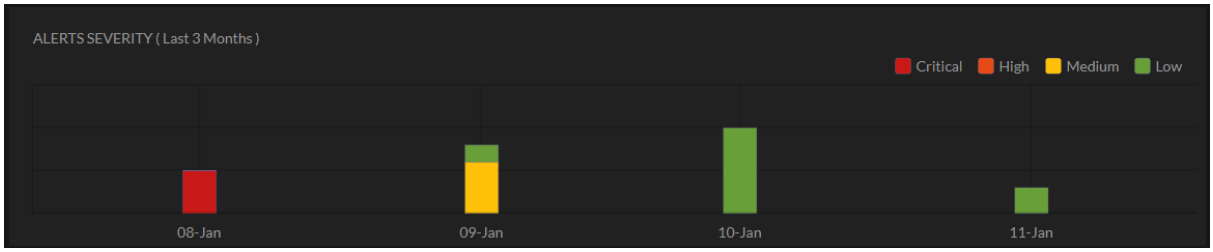


Begin an Investigation of Critical Alerts

You can begin your investigation of critical alerts in the following ways:

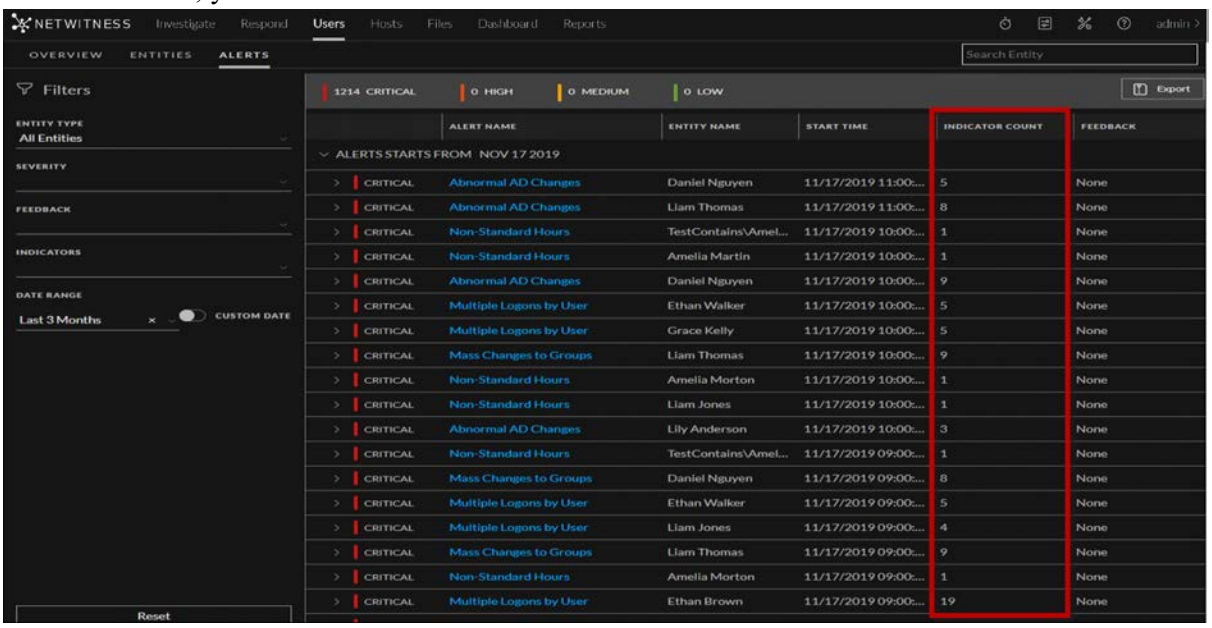


1. On the Overview tab, look at the ALL ALERTS.



Is there an even distribution of alerts or are there a few days when there was a noticeable spike? A spike could indicate something suspicious like malware. Make a note of those days so you can inspect the alerts (the bar from the chart links directly to the alerts for that specific day).

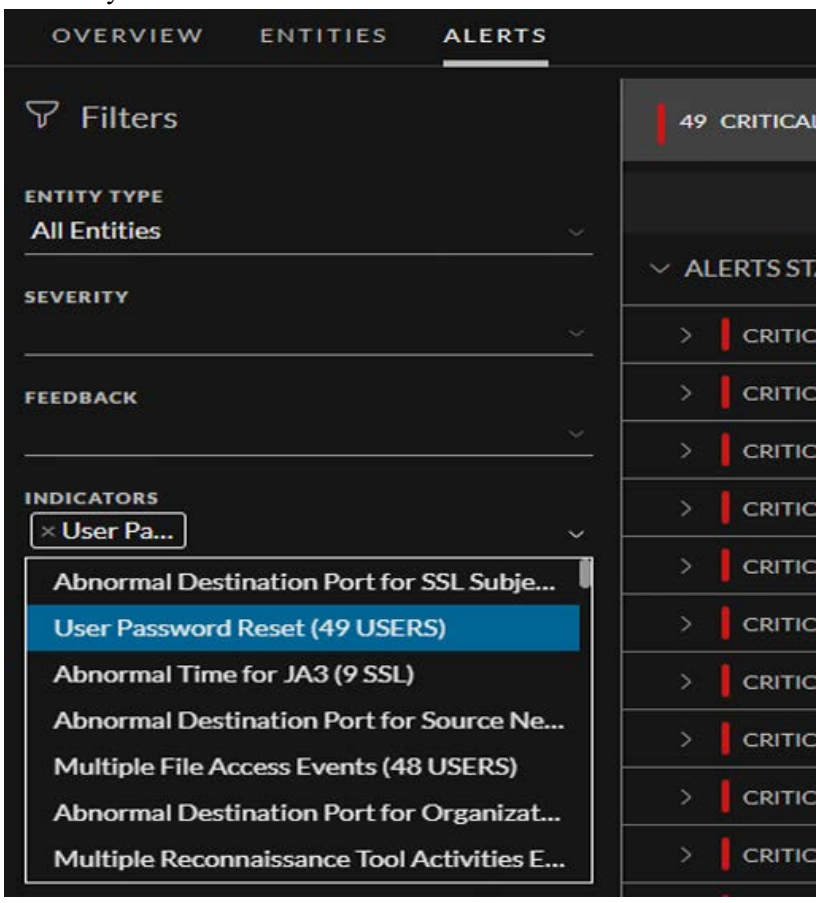
2. In the Alerts tab, you can view the indicator count:



To identify users with the highest number of alerts, more indicators help illustrate more insights and provide a more rigid timeline that you can follow.

3. Expand the top alerts in the list:
   - Look for alerts that have varied data sources. These show a broader pattern of behavior.
   - Look for a variety of different indicators.
   - Look for indicators with high numeric values, specifically for high values that are not indicative of a manual activity (for example, a user accessed 8,000 files).
4. Look for unique Windows event types that users do not typically change as these can indicate suspicious administrative activity.

5. Search by indicators.



The list shows the number of alerts raised that contain each indicator.

- Look for the top volume indicators; filter by an indicator and review by user to find users who experienced the highest number of these indicators.
- In general, you can ignore time-based alerts (for example, Abnormal Logon Time) as these are very common. However, they provide good context when combined with higher interest indicators.

6. Drill into more detail:
   - Leverage alert names to begin establishing a threat narrative. Use the strongest contributing indicator that usually determines the alert's name to begin explaining why this user is flagged.
   - Use the timeline to layout the activities found and try to understand the observed behaviors.
   - Follow up by reviewing each indicator and demonstrating the supporting information, in the form of graphs and events, that can help you verify an incident. Suggest possible next stages of investigation using external resources (for example, SIEM, network forensics, and directly reaching out to the user, or a managing director).
   - Conclude the investigation by prompting for feedback and leaving a comment.
7. Take action to address threats determined by the investigation of alerts. For more information, see Take Action on High-Risk User or Network Entity .

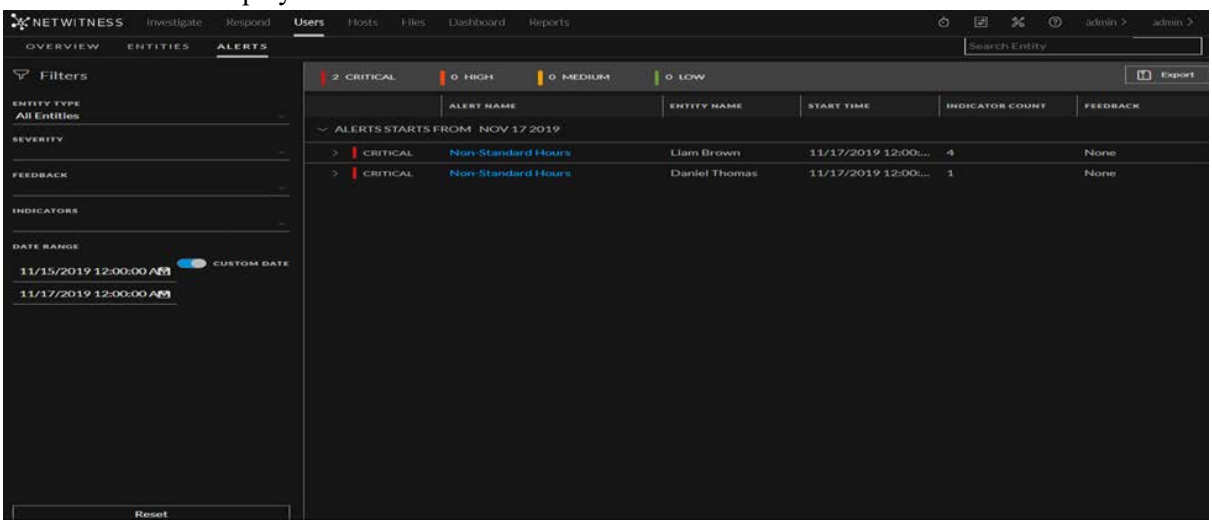
The following topics explain various ways to investigate alerts.

- Filter Alerts
- Investigate Events
- Manage Top Alerts
- View NetWitness UEBA Metrics in Health and Wellness

Filter Alerts

You can filter alerts displayed in the Alerts tab by severity, feedback, entity, indicators, and date range.

1. Go to **Users** > **Alerts**.
   The Alerts tab is displayed.



2. To filter by severity, click the down arrow under **Severity** in the **Alerts Filters** panel, and select any one option. The options are Critical, High, Medium, and Low.
3. To filter by feedback, click the down arrow under **Feedback**, and select any one option. The options are None, and Rejected.
4. To filter by entity, click the down arrow under **Entity Type**, and select any one option. The options are All Entities, Users, JA3, and SSL.
5. To filter by date range,
   - Click the down arrow under **Date Range** and select any one option. The Options are Last 7 Days, Last 2 Weeks, Last 1 Month, and Last 3 Months.
   - Select **Custom Date** under **Date Range**. In the **Start Date**, select the start range date range, and in the **End Date** select the end range date that you want the investigate.

The alerts are displayed in the right panel according to the filter you selected. To reset filters, in the bottom of left panel, click **Reset**.

Investigate Events

You can view all alerts and indicator associated with a user or network entity in the User Profile view.

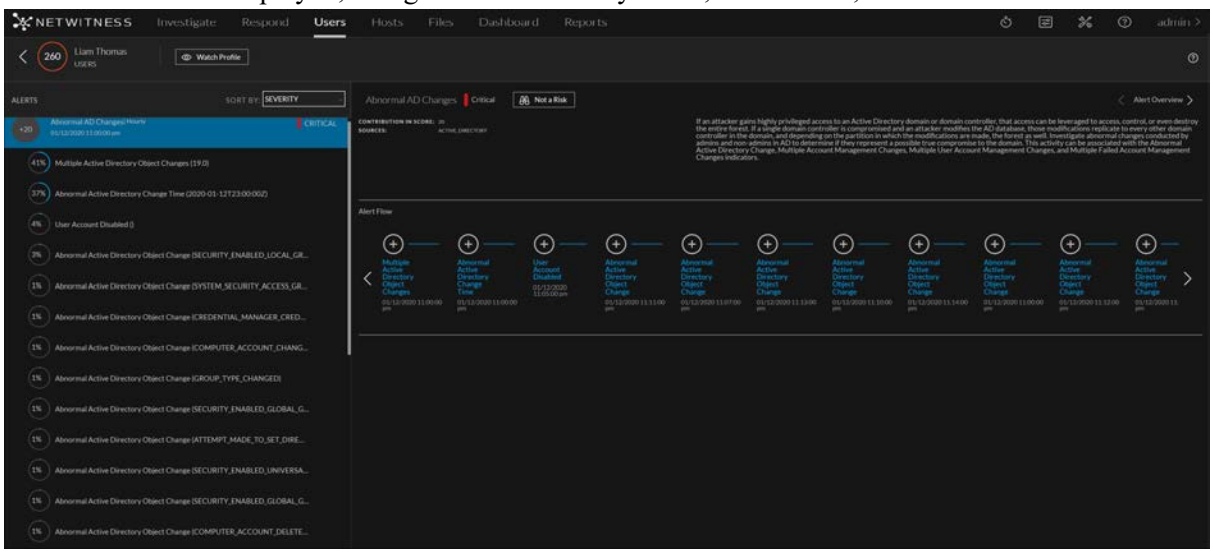
In the events table, you can find all events contributed to the specific indicator for the specific user or network entity.

For example, you can further investigate on events by clicking on a username or a network entity that pivots to **Investigate** > **Events**. In the Events view, you can see the list of events that occurred on that day for the specific user or network entity. By default, the time range is set to one hour. You can change the time range.

In case of Endpoint Indicators, you can pivot to **Host Details** view and can have deeper insight about that host. And, pivot to **Analyze Process** view for detailed investigation on the process for that event for that week as the time range is set to seven days. By default, the time range is set to seven days however, it can be customized.

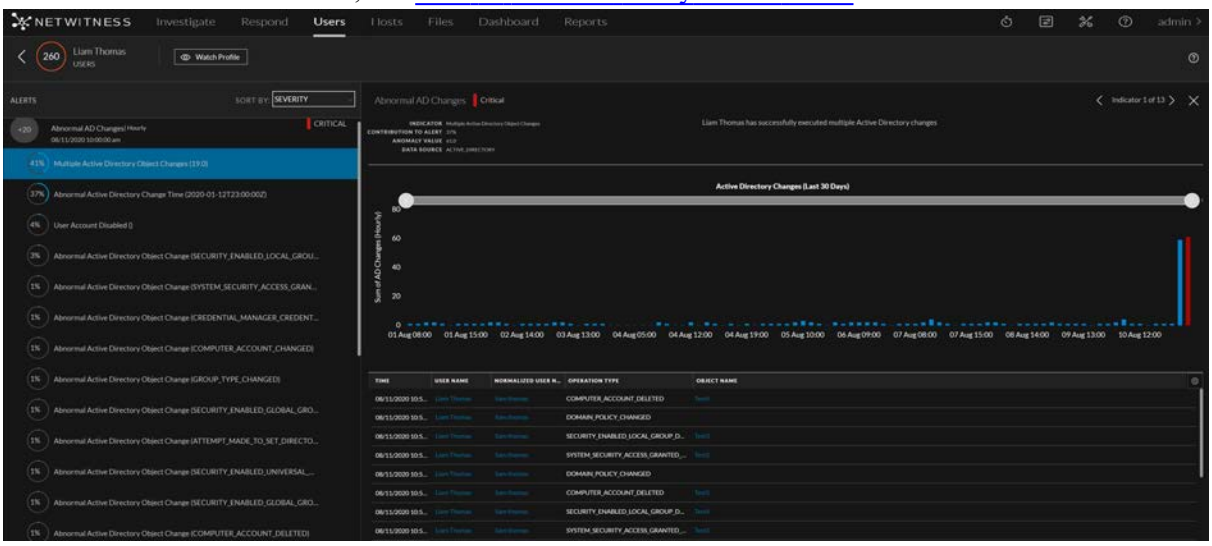
To view the events:

1. Go to **Users** > **Alerts**.
2. Under **Filters**, select the **Entity Type**.
   The indicators are displayed, along with the anomaly value, data source, and start time.



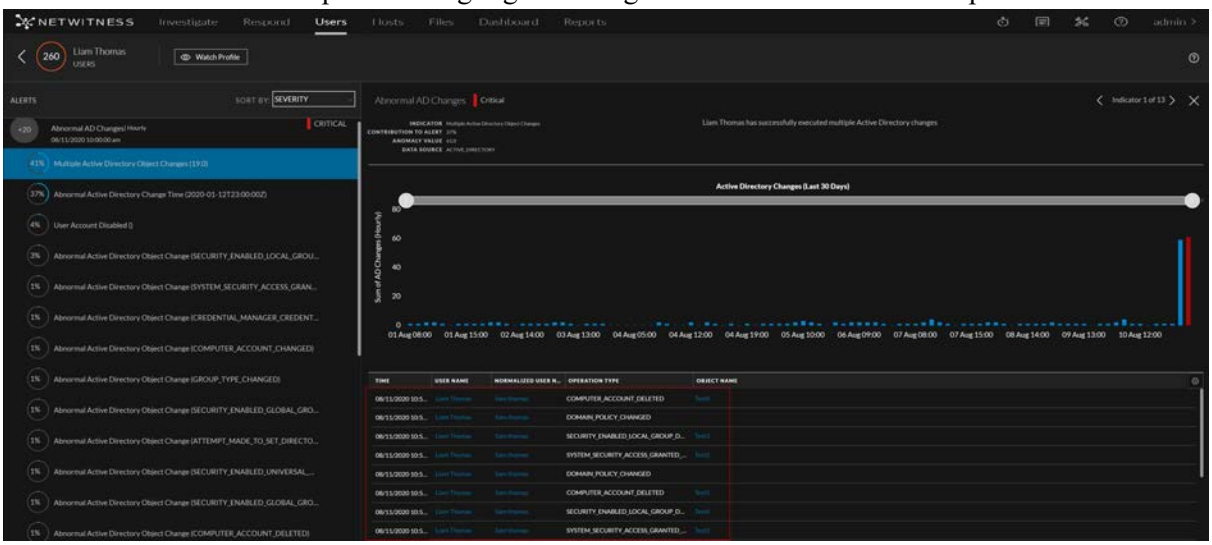
3. Click an alert name, and under **Alert Flow**, click the  icon.
   A graph is displayed that shows details about a specific indicator, including the timeline in which the anomaly occurred and the user associated with the indicator. The following figure shows an example of a graph. The type of graph can vary, depending on the type of analysis performed by NetWitness

UEBA. For more information, see User or Network Entity Profile View.



To pivot to the Events view:

1. Go to **Users** > **Alerts**, and select an alert or user or network entity.
   Indicators are displayed under the alert.
2. Select an indicator of interest.
   Values that can be used to pivot are highlighted in light blue at the bottom of the panel.



3. In the Events table, click the link highlighted in blue and pivot to the alert in the Events view.
   The Events view is displayed.
   For User events the username is a clickable pivot link. For JA3 and SSL Subject network entities events source IP, destination IP, destination country, destination organization, destination port, JA3 or SSL subject and source netname are the clickable pivot links.

For information about investigating items of interest in the Events view, see "Reconstructing and Analyzing Events" topic in the *NetWitness Investigate User Guide*.

To pivot to the Hosts Details view:

If you have NetWitness Endpoint installed, you can pivot to Hosts Details view for detailed information of the host.

1. Go to **Users** > **Alerts**, and select an alert or user or network entity.
   Indicators are displayed under the alert.
2. Select an indicator of interest.
   Details about the indicator are displayed in the right panel.
3. In the events table, click the events related to the host.
   The Host Details view is displayed.

For information about investigating items of interest in the Hosts view, see "Investigating Hosts" topic in the *NetWitness Endpoint User Guide*.

To pivot to the Analyze Process view:

If you have NetWitness Endpoint installed, you can pivot to Analyze Process view for detailed information about the process.

1. Go to **Users** > **Alerts**, and select an alert or user or network entity.
2. Select an alert name. Indicators are displayed under the alert.
3. Select an indicator of interest.
   Details about the indicator are displayed in the right panel.
4. In the Events table, click the events related to the process.
   The Analyze process view is displayed.

For more information, see "Investigating a Process" topic in the *NetWitness Endpoint User Guide*.

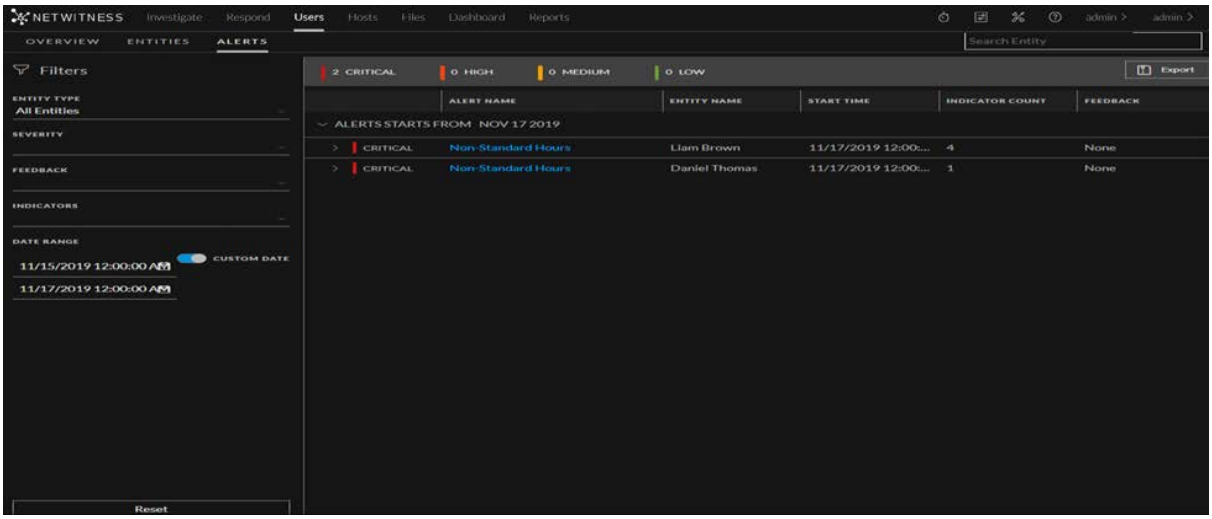
For more information on how to read the indicator charts, see Reading an Indicator Chart topic.

Manage Top Alerts

You can export a list of all alerts to a .csv file format. You can use this information to compare the data from other sources in other data analysis tools like tableau, powerbi, and zeppelin.

To export alert data to a .csv file:

1. Go to **Investigate** > **Entities** > **Alerts**.
   The Alerts tab is displayed.


NETWITNESS
Investigate
Respond
Users
Hosts
Files
Dashboard
Reports
admin
admin
OVERVIEW
ENTITIES
ALERTS
Search Entity
Filters
ENTITY TYPE
All Entities
SEVERITY
FEEDBACK
INDICATORS
DATE RANGE
11/15/2019 12:00:00 AM
CUSTOM DATE
11/17/2019 12:00:00 AM
Reset
2 CRITICAL
0 HIGH
0 MEDIUM
0 LOW
Export
ALERT NAME
ENTITY NAME
START TIME
INDICATOR COUNT
FEEDBACK
ALERTS STARTS FROM NOV 17 2019
CRITICAL
Non-Standard Hours
Liam Brown
11/17/2019 12:00:...
4
None
CRITICAL
Non-Standard Hours
Daniel Thomas
11/17/2019 12:00:...
1
None

2. On the top right, click **Export**.
   All the alert data is downloaded in a .csv file format. The following figure is an example of the exported alert data in .csv format:

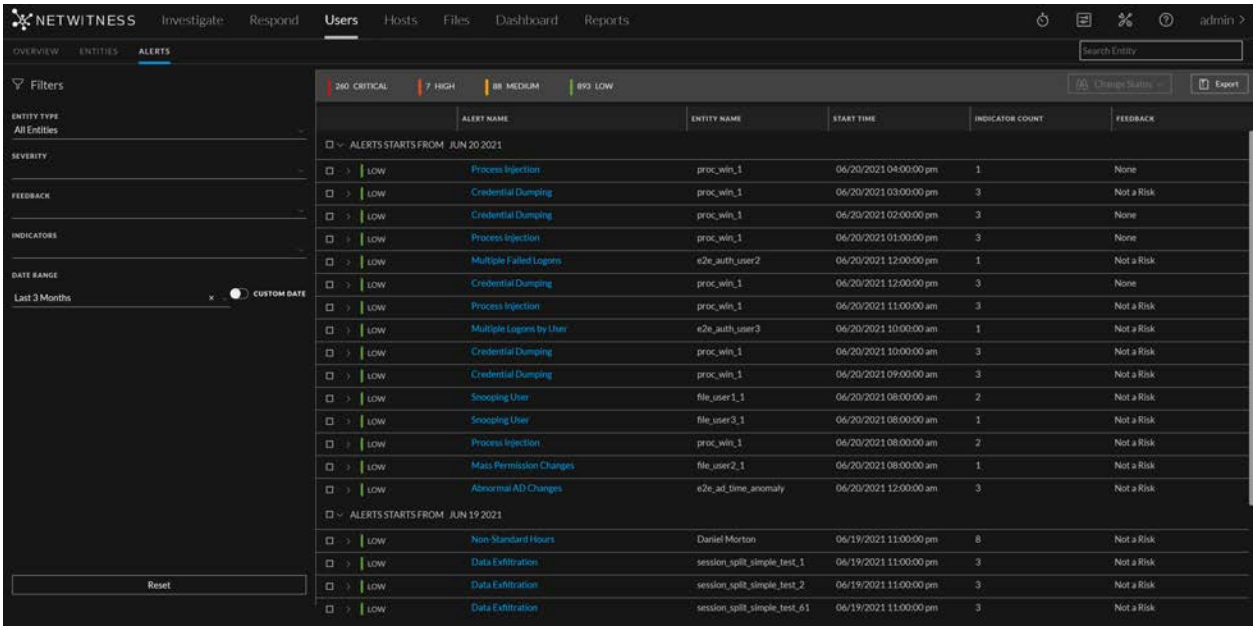
	A	B	C	D	E	F	G
1	Alert Name	Entity Name	Start Time	# of Indica	Status	Feedback	Severity
2	Brute Force Authenticati	e2e_auth_user2	Mar 06 20	1	Reviewed	No Feedback	Low
3	Multiple Logons by User	e2e_auth_user3	Mar 06 20	1	Reviewed	No Feedback	Low
4	Snooping User (Hourly)	file_user1_1	Mar 06 20	2	Reviewed	No Feedback	Low
5	Snooping User (Hourly)	file_user3_1	Mar 06 20	1	Reviewed	No Feedback	Low
6	Mass Permission Change	file_user2_1	Mar 06 20	1	Reviewed	No Feedback	Low
7	Abnormal AD Changes (H	e2e_ad_time_anor	Mar 06 20	4	Reviewed	No Feedback	Low
8	Abnormal AD Changes (H	Amelia Thompson	Mar 05 20	13	Reviewed	No Feedback	Medium
9	Abnormal AD Changes (H	Lily Walker	Mar 05 20	11	Reviewed	No Feedback	Low
10	Multiple Logons by User	Matilda Martin	Mar 05 20	6	Reviewed	No Feedback	Low
11	Multiple Logons by User	Matilda Robinson	Mar 05 20	6	Reviewed	No Feedback	Low

Set Alert Feedback Status

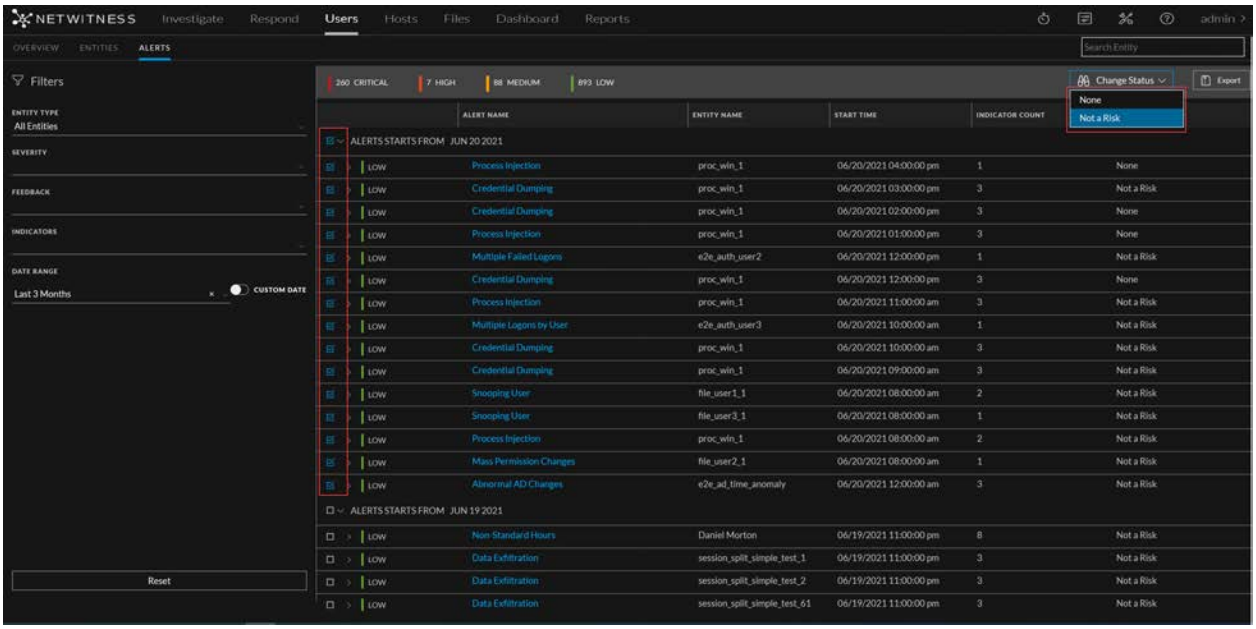
The analyst can mark alerts as **Not a Risk** or they can remove the **Not a Risk** feedback status by updating it to **None**. Also, multiple alerts grouped by date can be selected to perform this action.

To set an alert feedback status:

1. Go to **Users** > **Alerts**.
   The Alerts tab is displayed with a list of alerts.



2. Select the checkbox ☐ and click Not a Risk to mark all the alerts as not risky.



A confirmation message is displayed to confirm the change.

3. Click **OK**, to confirm the status or click **Cancel** if you don't want to continue with the update.

Note: All alerts that are marked as Not a Risk will not contribute to the entity's score.

View Modeled Behaviors

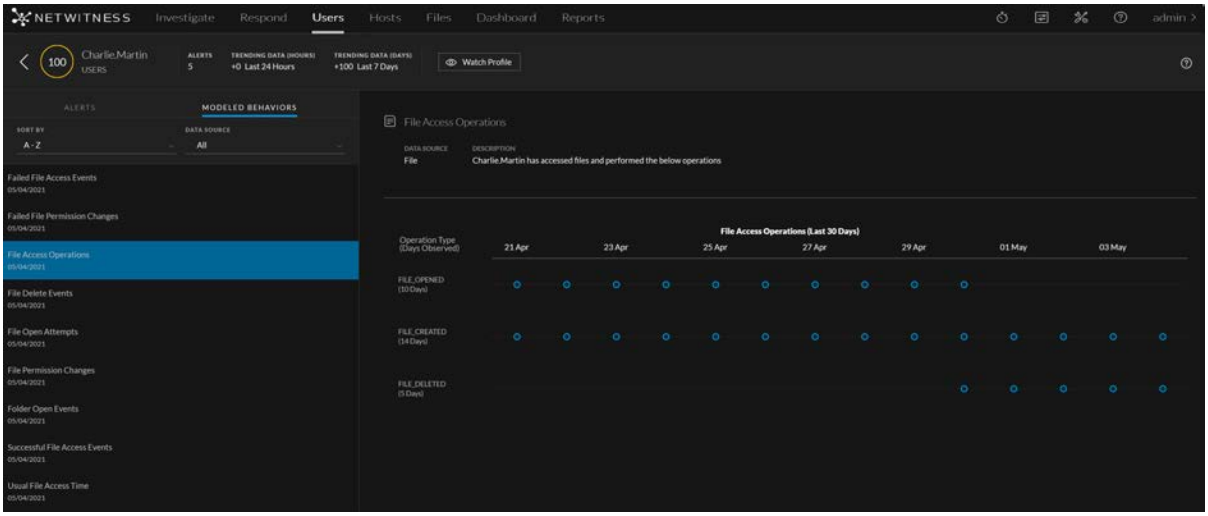
(From 11.5.1 and later) NetWitness UEBA Modeled Behavior provides analysts with visibility into all the activities of a user. These modeled behaviors are based on the log data leveraged by UEBA and are available a day after the UEBA service is configured and running. UEBA monitors abnormal user behaviors to identify risky users and this requires data to be processed for a certain period of time. However, modeled behaviors reflect all the activities of the user within a day of the service configuration.

For example, if a user fails multiple times by logging in with incorrect credentials within an hour, analysts can view these behaviors as **Failed Authentications** for the user.

Note:
- Modeled Behaviors are not applicable to network entities.
- By default, the sorting and data source options are disabled when there is no data available.

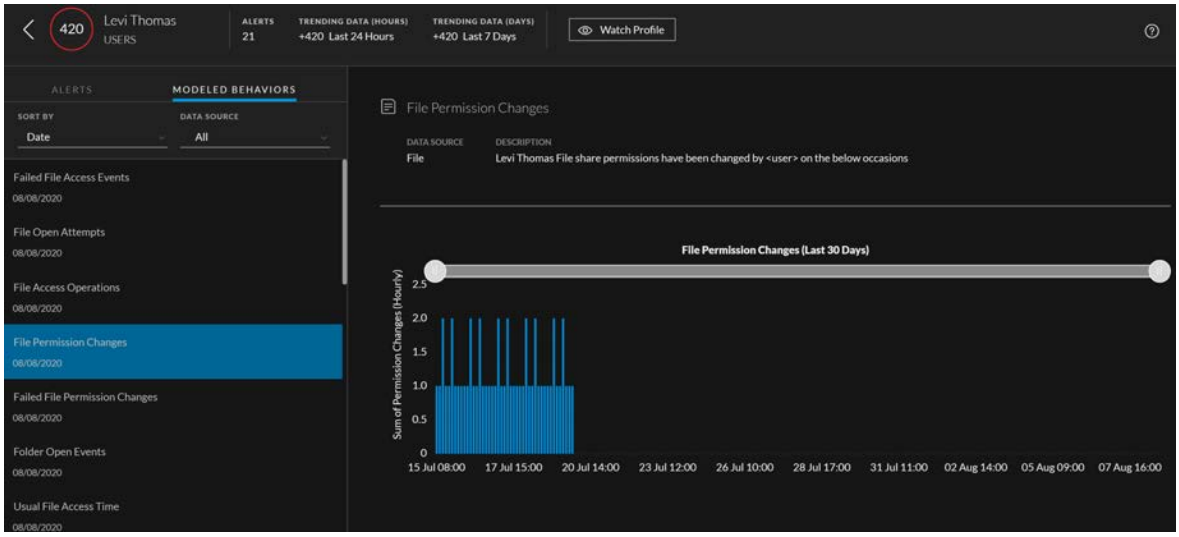
To view the Modeled Behaviors:

1. Log into NetWitness Platform and click **Users**.
2. Do one of the following:
   - In the **Overview** tab, under **Top Risky Users** panel, click on a username.
   - In the **Entities** tab, click on a username.
3. Click the **Modeled Behaviors** tab, to view the Modeled Behaviors highlighted with a blue line in the left panel. The results can be sorted by the date or in alphabetical order.



4. Select the data source from the drop-down according to your preference and filter the modeled behaviors:
   - Active Directory
   - File
   - Authentication

5. Based on the data source you select, the following information is displayed on the right panel:
   - Data source name
   - Modeled Behavior description
   - A graph is displayed that contain details of a specific Modeled Behavior. You can view the modeled behaviors of a user for the last 30 days. The type of graph can vary, depending on the type of analysis performed by UEBA. The following figure is an example of a Modeled Behavior.



Reading an Indicator Chart

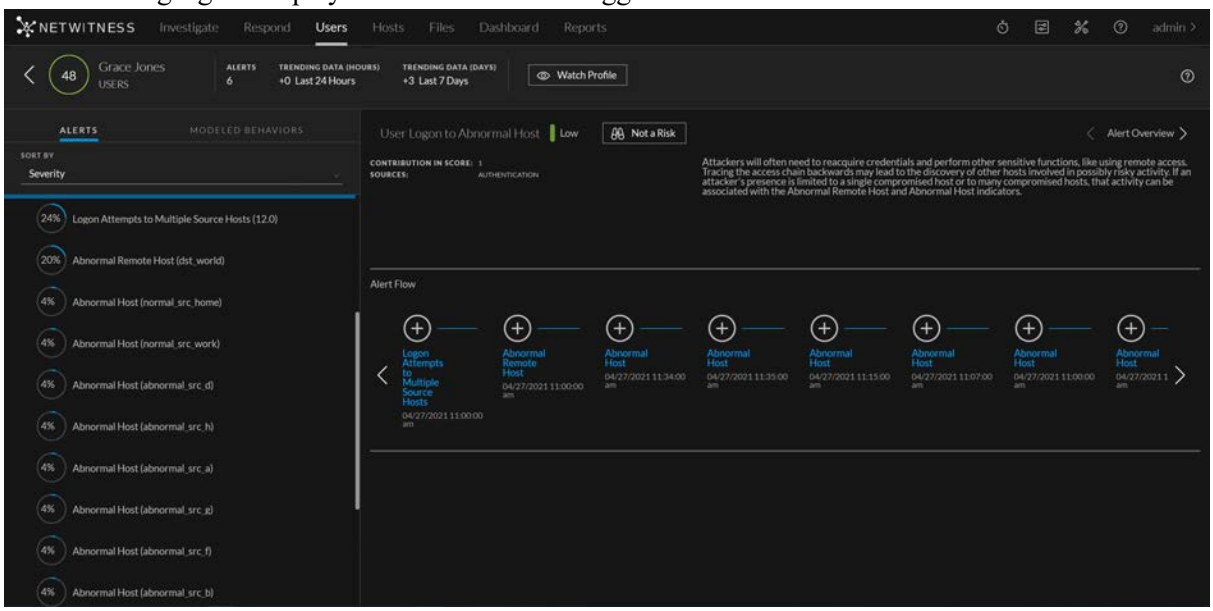
Note: In version 11.6 and later, the pie chart was replaced by the dotted chart.

An indicator chart is a pictorial illustration of the anomaly and baseline values of an entity that you want to further investigate.

The chart gives the Analyst a better insight of the indicator which in turn will help determine the next steps. The chart provides the analyst with the user's baseline values over time to better understand the context of the anomaly.

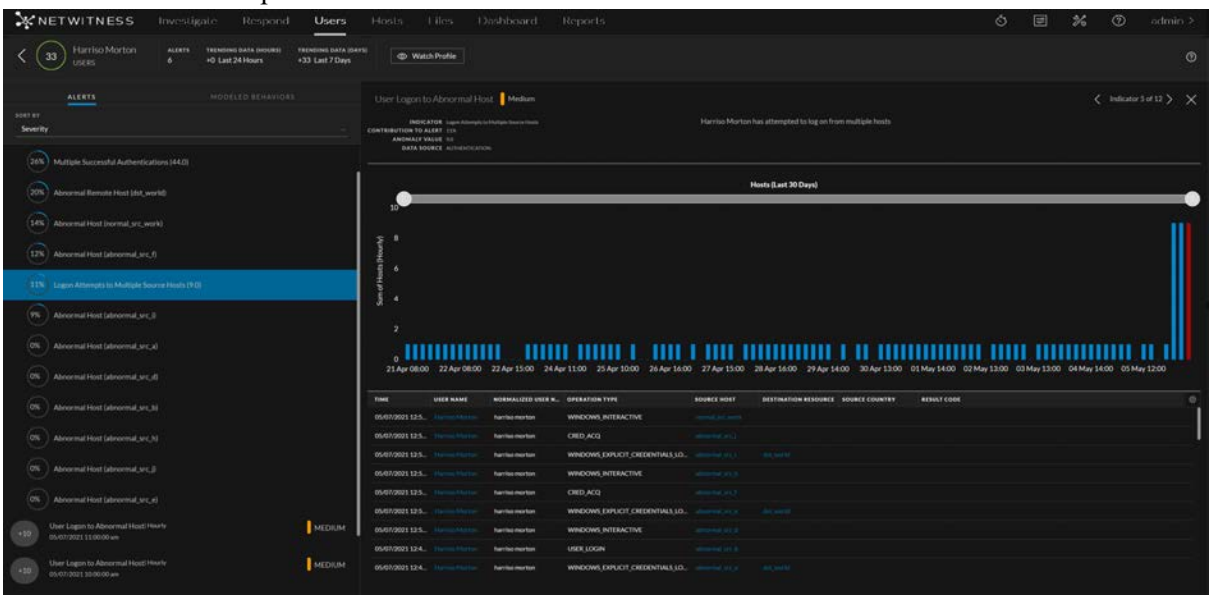
To view an indicator chart,

1. Log in to NetWitness Platform.
2. Navigate to **Users** > **Entities**.
3. Select the user you want to investigate.
   The following figure displays an alert for a user logged on to an abnormal host.



4. In the Alert Flow section, select the Logon Attempts to Multiple Source Hosts.

5. Click the + icon to expand and view the details.



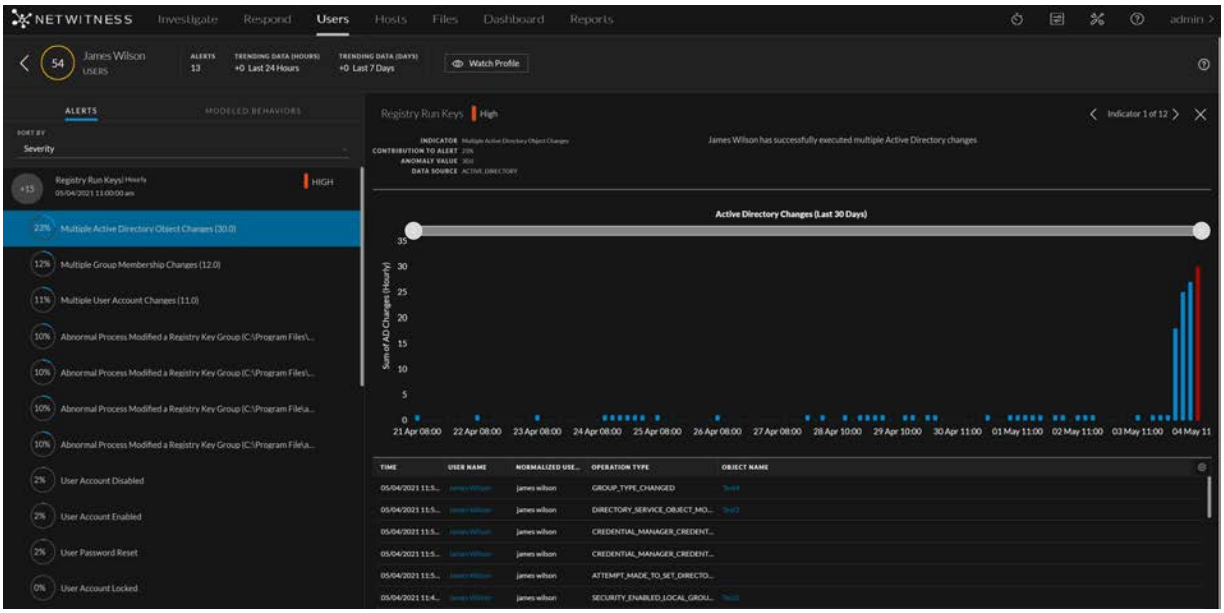
Types of Charts

There are three main types of charts currently available.

Continuous Bar Chart

In this type of an indicator chart, the bar color differentiates the behavior by displaying a blue bar and a red bar. For example, the following figure displays in a span of 30-days the number of Active Directory changes made on a daily basis which are displayed by blue bars and indicates the baseline behavior.

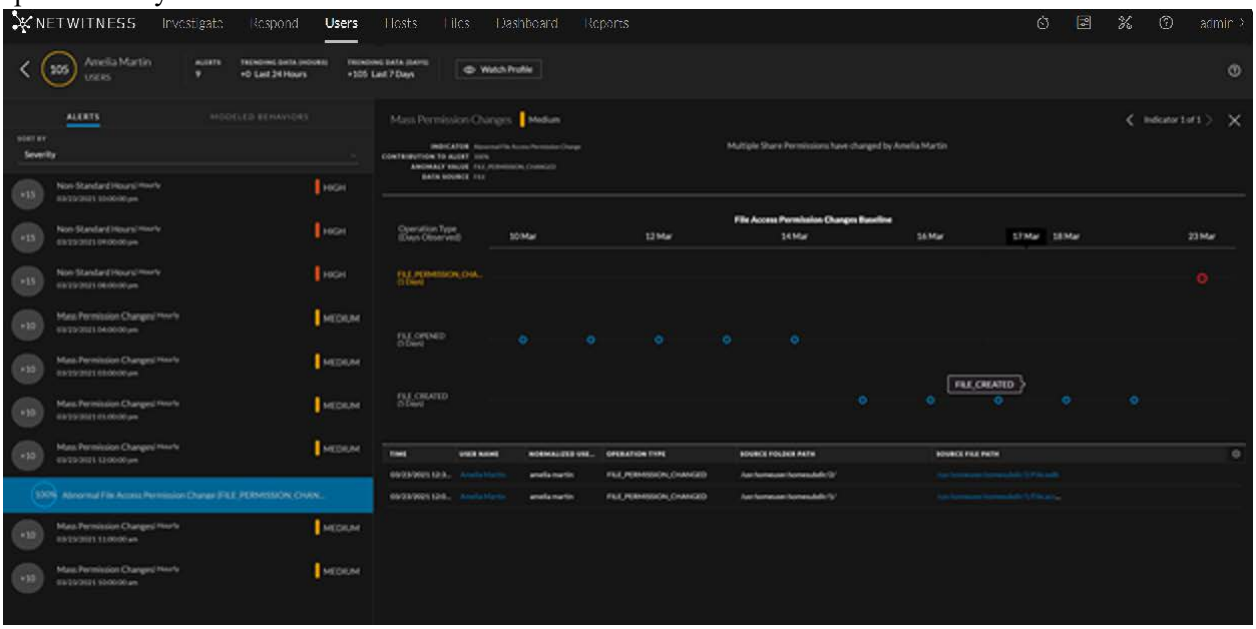
The red bar indicates that the user has accessed a high number of files in a specific hour.



Another variation in the visualization of the chart is where you see an additional series of grey bars that represents the baseline values of the model. In this case, if the blue bars series is displayed, it depicts the specific entity trend that the anomaly is also a part of.

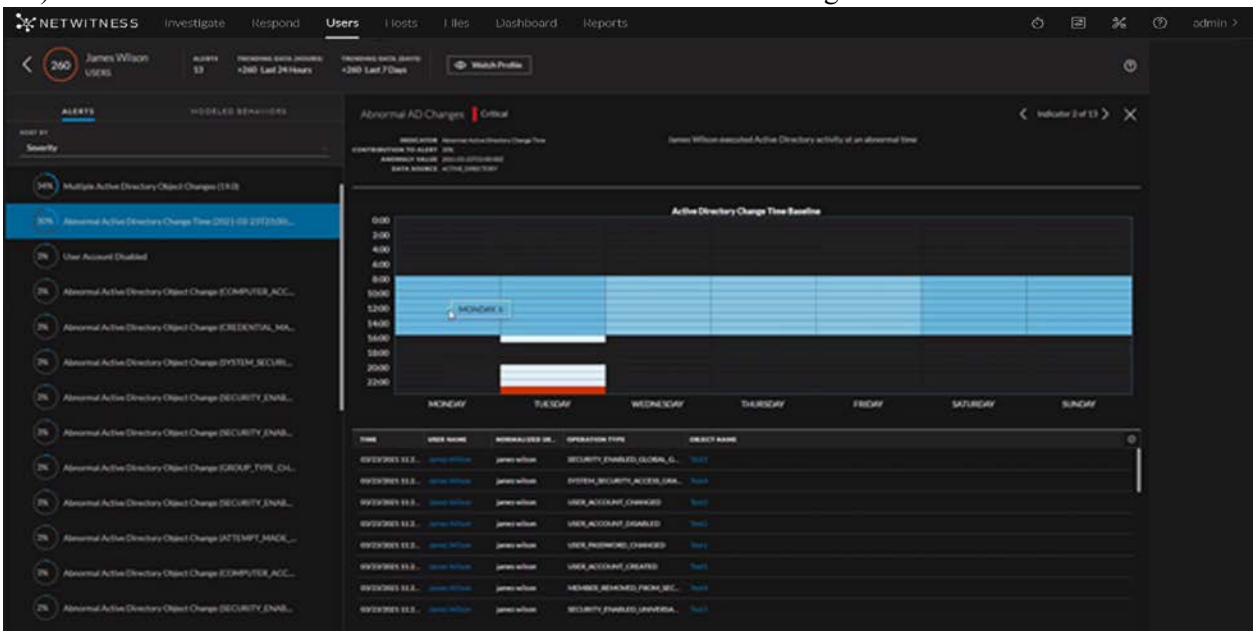
Dotted Chart

In the dotted indicator chart, the anomaly is displayed on top of the graph indicated by yellow color text and red color circle. The chart provides the analyst with the user's baseline values over time to better understand the context of the anomaly. The additional values (apart from the anomaly value) depicted in the Y-axis, represent the baseline values and the total number of days they were observed for this specific entity.



Time Chart

The time indicator chart displays the time the user has accessed a particular information. For example, in the following figure, the user has accessed Active Directory at an abnormal time over the past 30 days. The regular working hours of the user are the baseline values and the anomaly value (the hour marked in red) indicates that this is an abnormal time for this user to make changes in AD.



View NetWitness UEBA Metrics in Health and Wellness

NetWitness UEBA sends metrics to the System Stats Browser tab in  **(Admin)** > **Health and Wellness**. Along with basic system usage information, metrics that are specific to NetWitness UEBA users, alerts, and events are provided.

Analysts can use these metrics in the following ways:

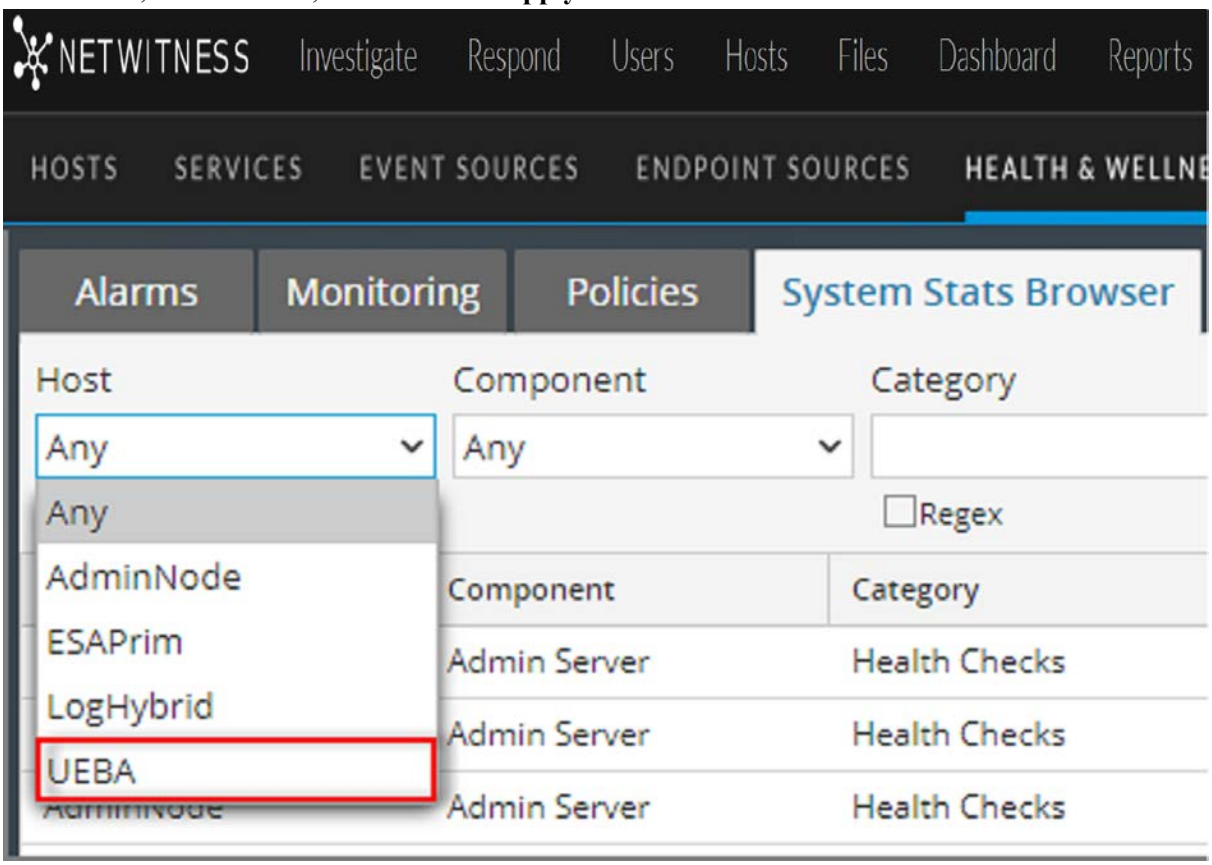
- Confirm that the currently procured license is in compliance with their license agreements, and by how much per day.
- Determine if the system is functioning as required.
- Actively monitor new events.
- Monitor the creation of new indicators and alerts.

If these critical metrics are reported as "0", it may indicate a system malfunction.

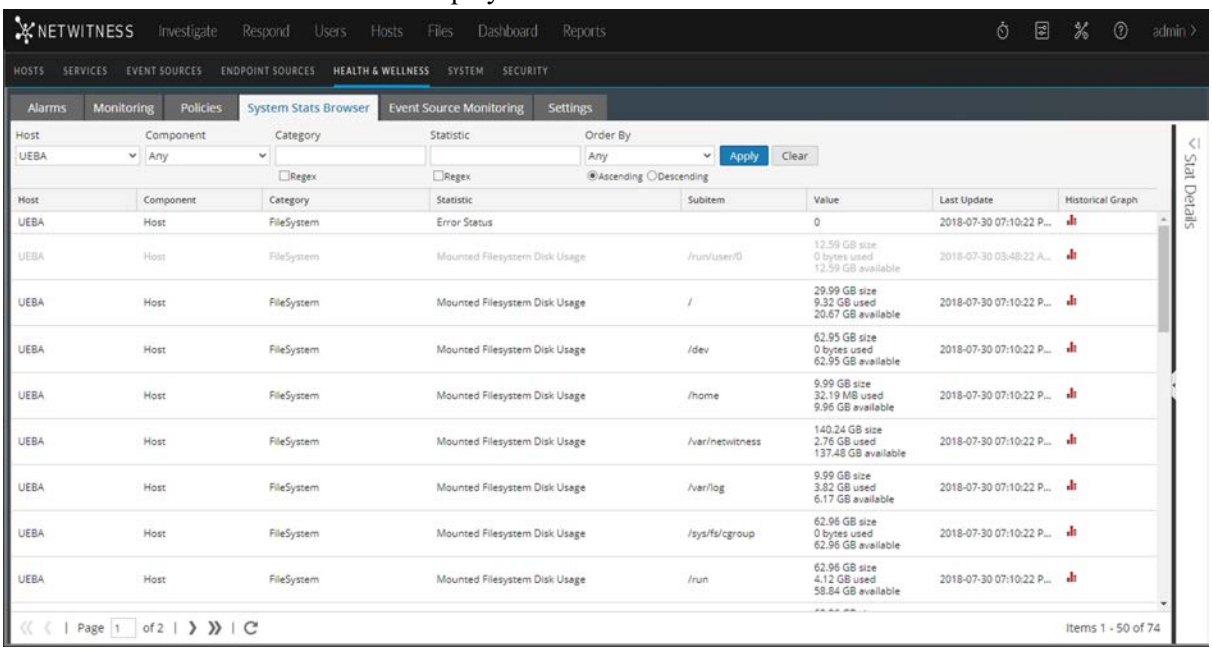
To view NetWitness UEBA metrics in the System Stats Browser in Health and Wellness:

1. Go to  **(Admin)** > **Health & Wellness**.
2. Click the **System Stats Browser** tab.
   The System Stats Browser is displayed.

3. Under Host, select **UEBA**, and then click **Apply**.

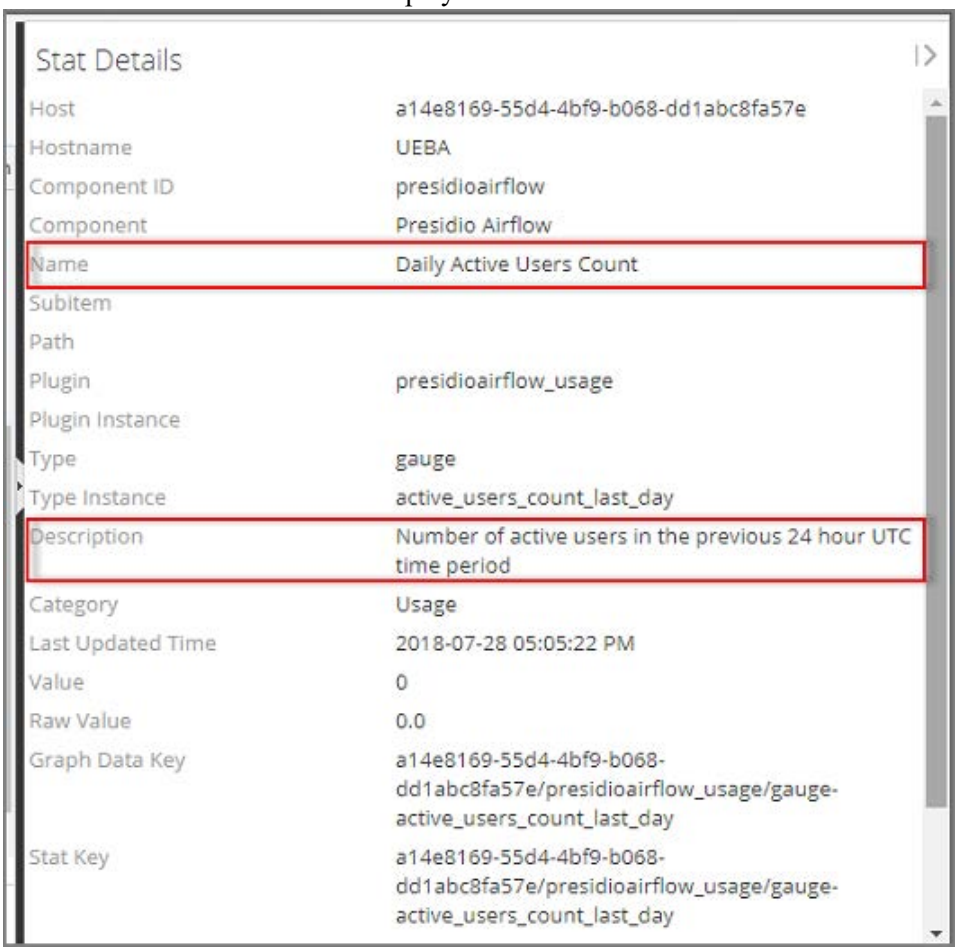


Results for NetWitness UEBA are displayed.



4. To view details for a statistic, click **Stat Details**.

The details of the statistics are displayed.



The **Name** and **Description** fields provide a summary of the metrics that are displayed.

For more information about Health and Wellness and System Stats Browser tab, see "Monitor System Statistics" topic in the *System Maintenance Guide*.

Monitor Health and Wellness of UEBA

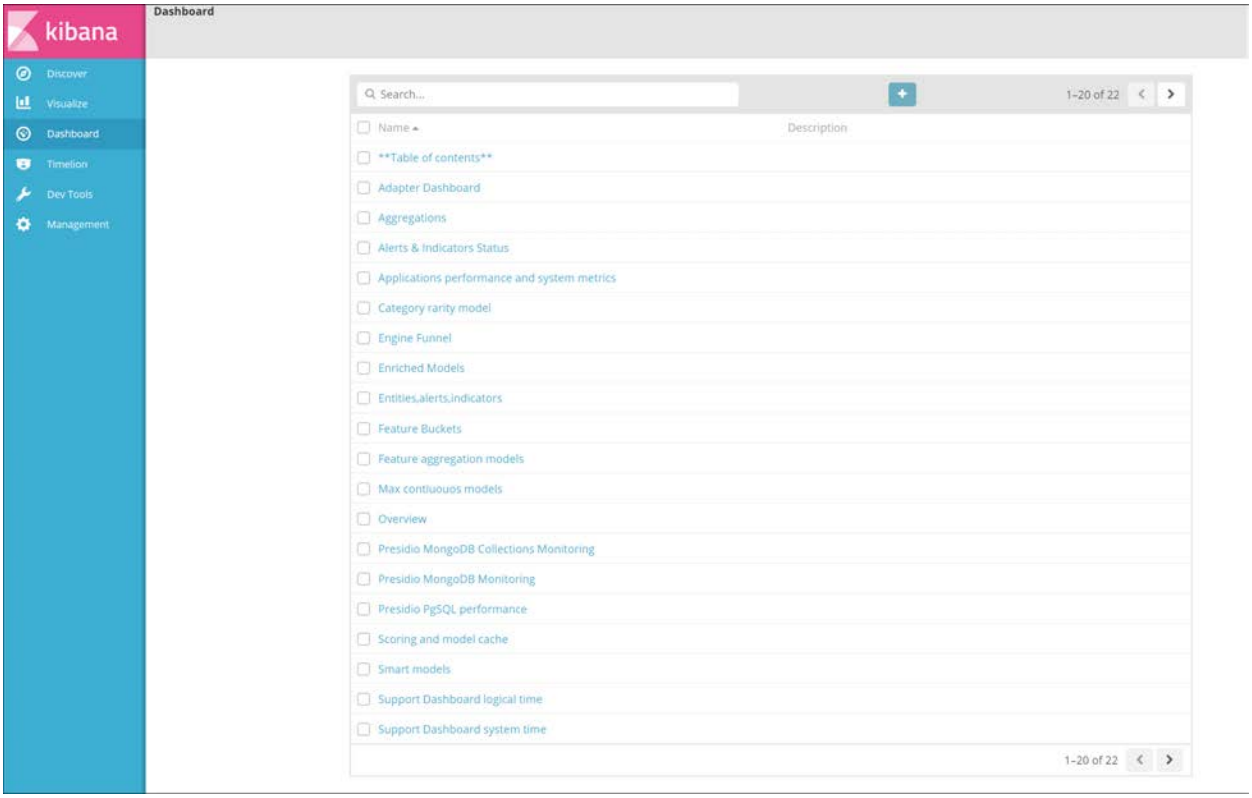
You can view the status of UEBA host in the**Users** > **OVERVIEW** tab.

The UEBA system should generate at least one alert weekly. If the system stops generating the alerts for a period of seven days or more, advanced monitoring is required to monitor statistics about the total number of events versus successful events, total number of alerts generated, and so on.

Advanced monitoring is enabled through a third-party tools prepackaged in NetWitness: Kibana and Airflow.

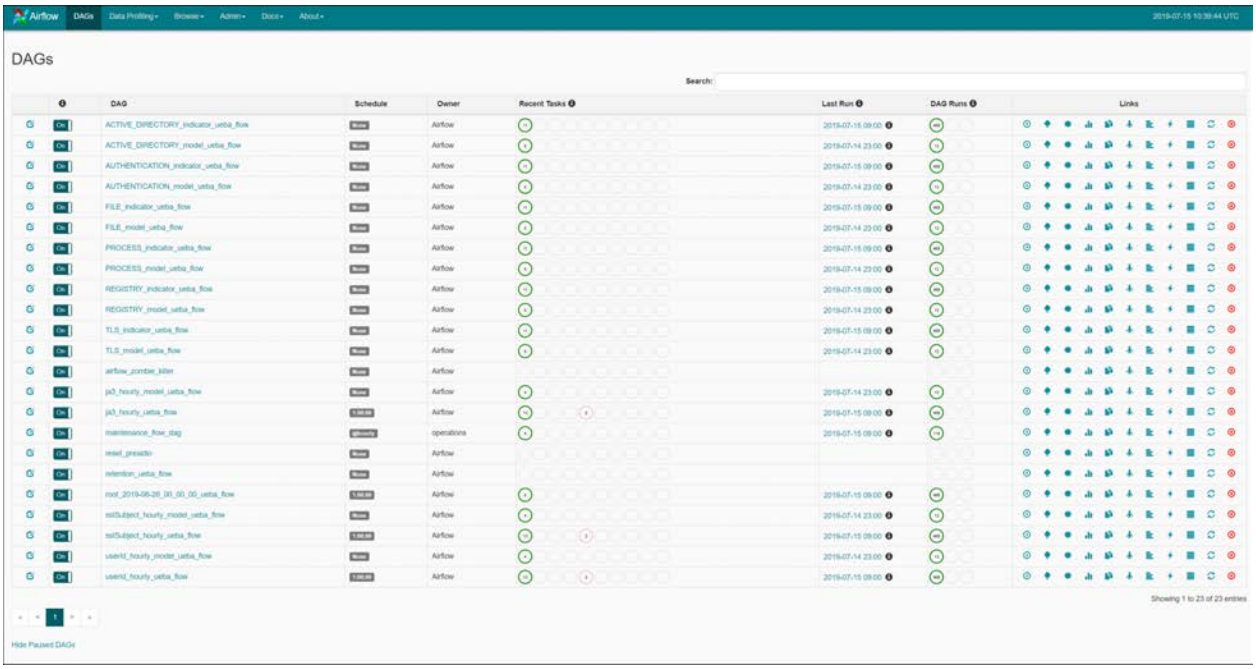
Access Kibana

To access kibana, go to `https://<UEBA_host>/kibana/app/kibana#/`, enter user name and password and the Dashboard is displayed.



Access Airflow

To access Airflow, go to `https://<UEBA_host>/admin/`, enter user name and password and the DAGs view is displayed.



Note: The Kibana and Airflow web server user interface password is the same as the deploy_admin password. Make sure that you record this password and store it in a safe location.

Kibana

Kibana is an open source analytics and visualization platform. You can monitor the health of UEBA through various dashboards:

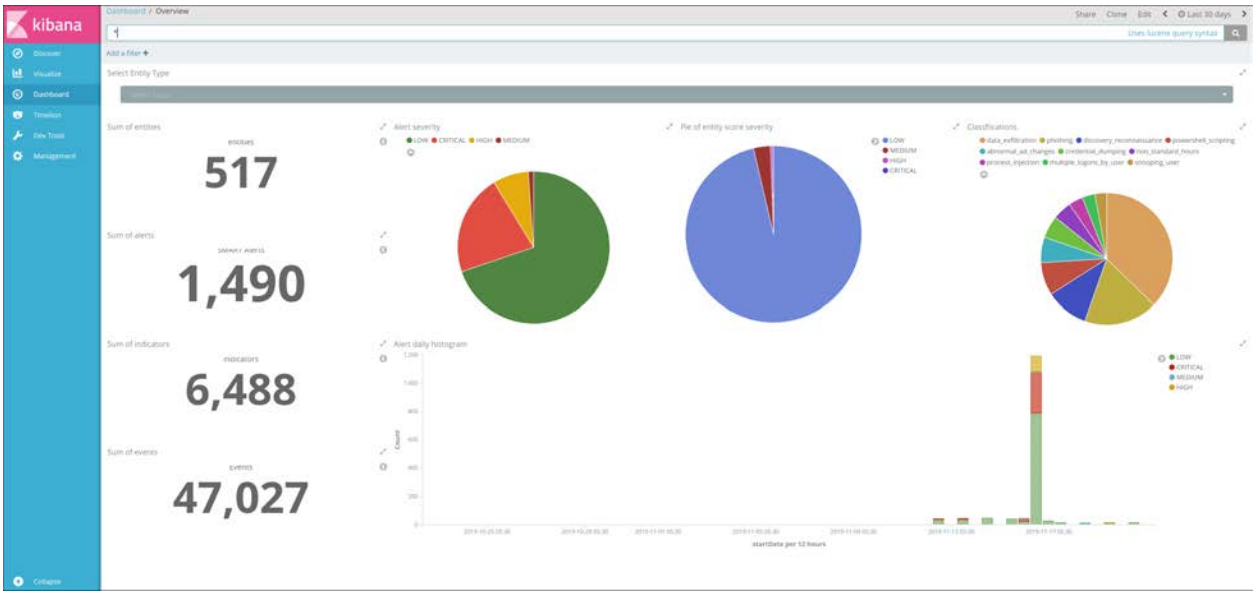
Overview Dashboard

The Overview dashboard provides the statistics over the analytics about the users, entities, alerts, and indicators, such as:

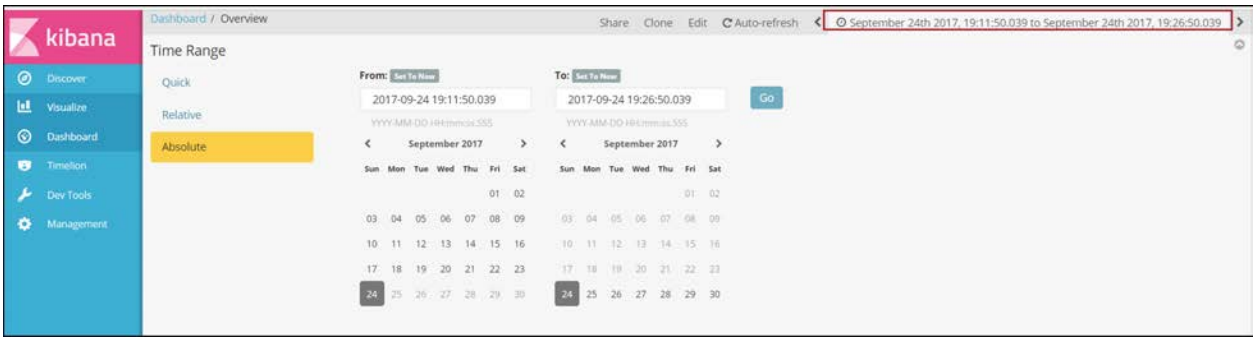
- The alerts type that are generated, and the alert severity distribution with the severity types (Low, Medium, High, Critical).
- Total number of active entities and how many alerts are generated for those entities.
- The number of indicators and events processed.
- The pie chart for entity score severity and distribution for the alerts classification.
- Alert daily histogram, which is the total number of alert per each severity triggered over time.

To access the overview dashboard:

1. Go to Kibana, click **Dashboards** > **Overview**.
   The Overview dashboard is displayed with the aggregate results for all entities.



2. To view the data for a specific entity, select a value from the **Select Entity Type** drop-down. For example, ja3, sslSubject, or userid.
3. Adjust the time range on the top-right corner of the page based on your requirement to view the statistics.



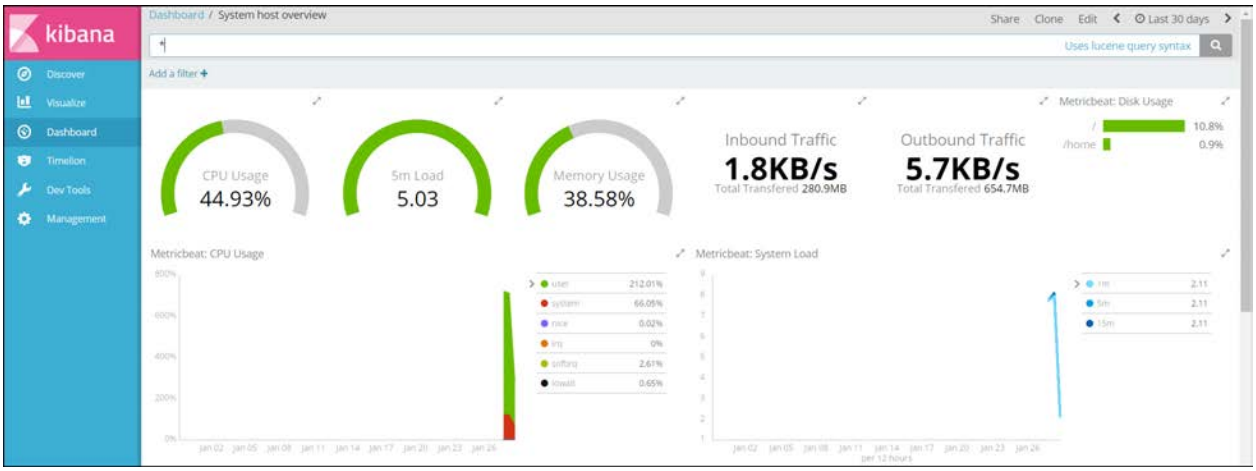
System Host overview

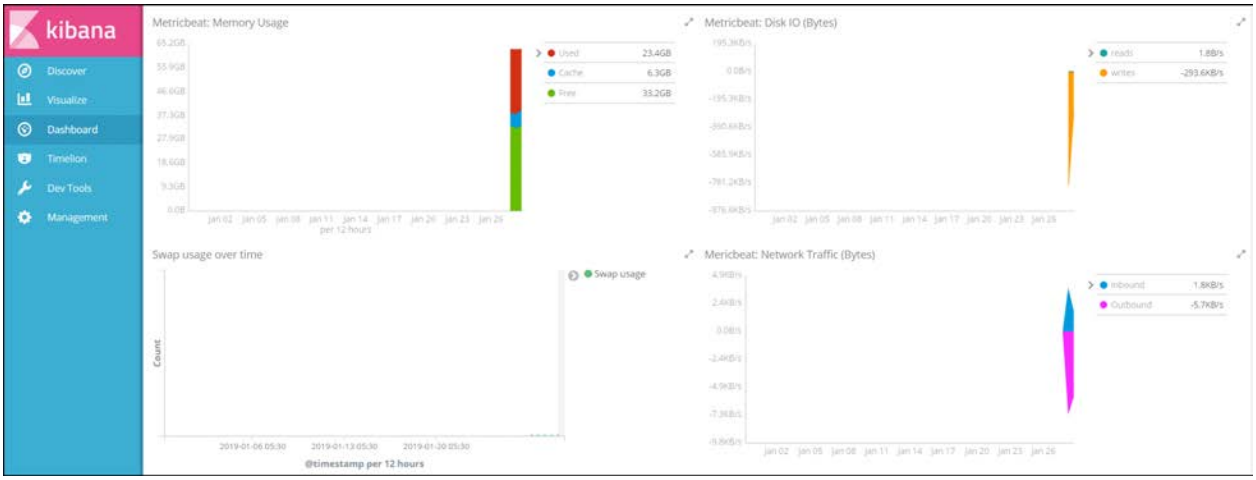
The System Host overview dashboard monitors the performance and health of UEBA hosts, such as:

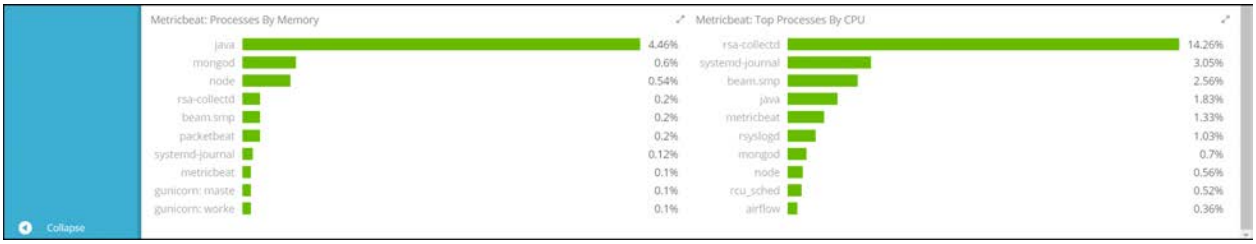
- CPU usage.
- Memory consumption, and network.
- Process consuming CPU and Memory, for example MongoDB.
- Statistics over the disk usage.
- Inbound data is the amount of data transferred by user to view the UEBA UI.
- Outbound data is the amount of data fetched by UEBA from Broker or Concentrator.

To access System Host overview dashboard

1. Go to Kibana, click **Dashboards** > **System host overview**.
   The System host overview dashboard is displayed.







2. Adjust the time range on the top-right corner of the page based on your requirement to view the statistics.



Note: During historical load, the system works in high parallelism. Due to that IO, CPU, and Memory is in high utilization. The pace would be 30 logical days four wall clock time. Once the UEBA server is online, the resource utilization reduces.

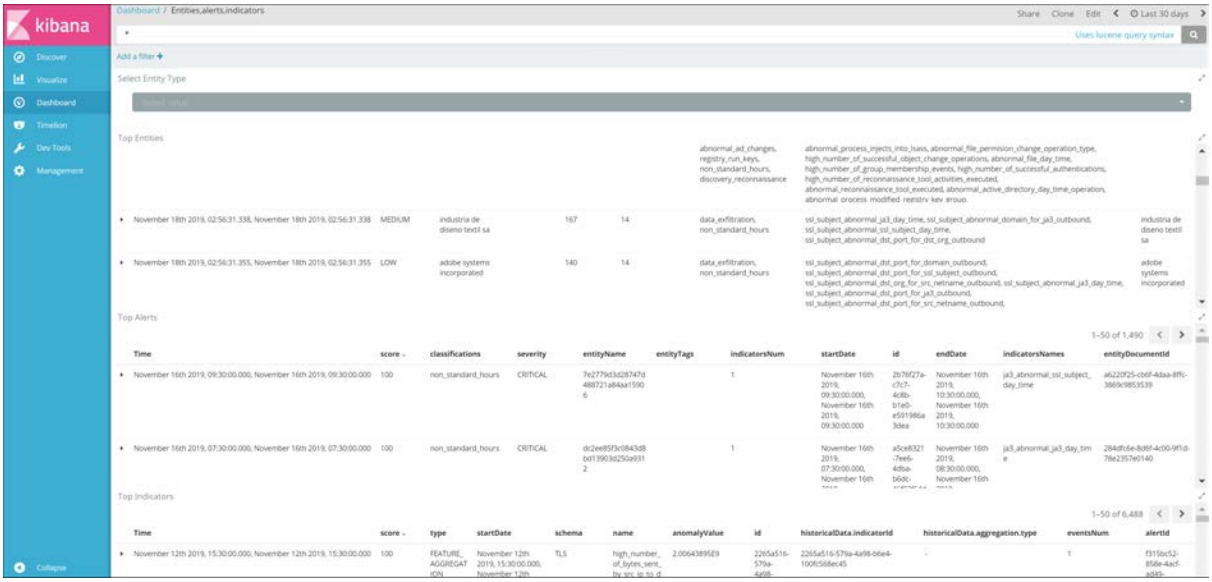
Adapter Dashboard

The **Adapter** dashboard is used to monitor the following:

- The failed events distribution.
- Total number of events versus successful events.
- Saved events per schema.

To access the entities, alerts and indicators

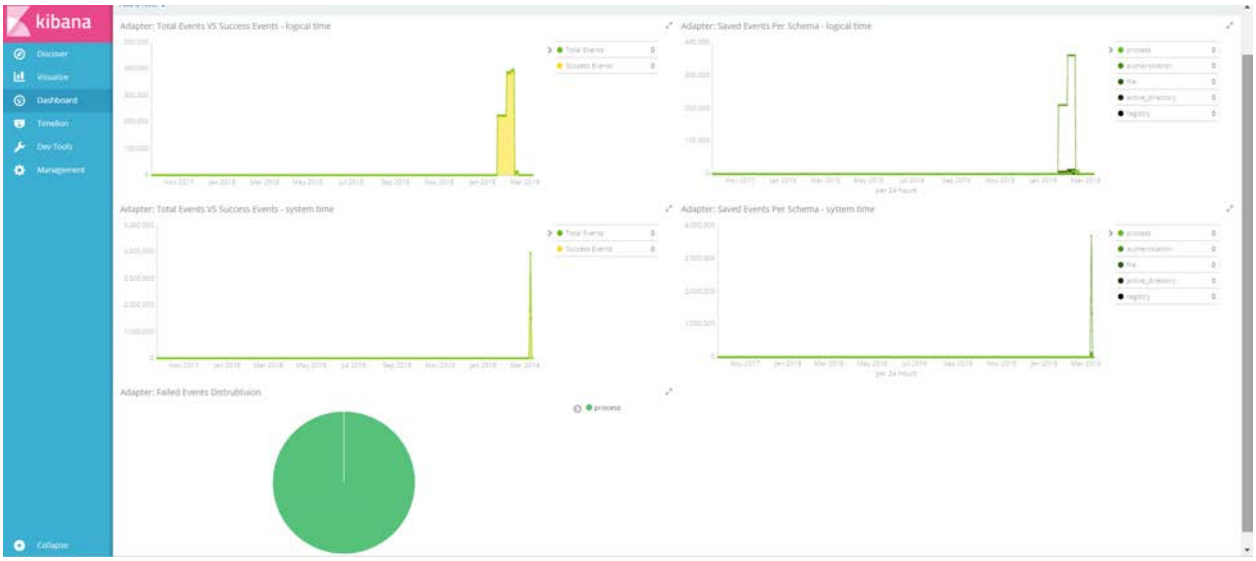
1. Go to Kibana, click **Dashboards** > **Entities, alerts, indicators**.
   The Entities, alerts, indicators Dashboard is displayed with an aggregate data for all entities.



2. To view the data for a specific entity, select a value from the **Select Entity Type** drop-down. For example, ja3, sslSubject, or userid.

To access the adapter dashboard system Time

1. Go to Kibana, click **Dashboards** > **Adapter**.
   The Adapter Dashboard is displayed.



2. Adjust the time range on the top-right corner of the page based on your requirement to view the statistics.



Support Dashboard Logical Time

The **Support Dashboard Logical Time** provides the capability to detect events processed time, which is different from the system time, such as:

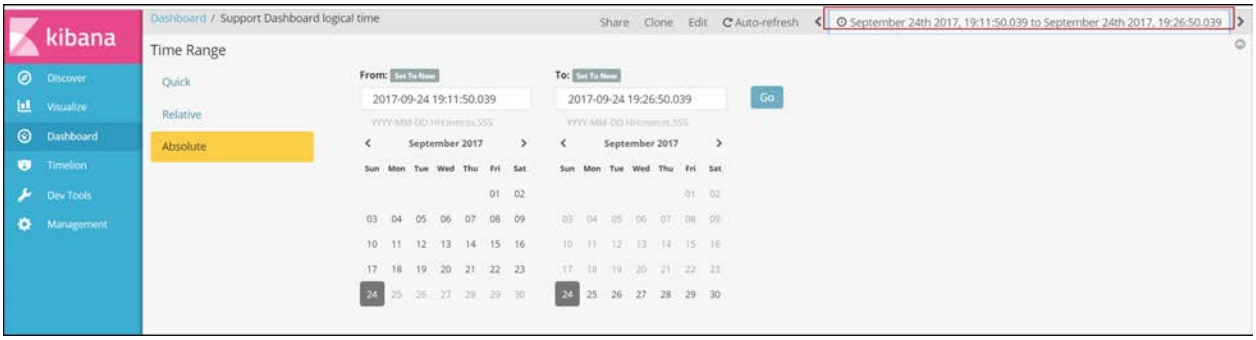
- The amount of filtered events over time per schema
- The total number of alerts generated
- The alert types distribution
- The events that are related to an alert

To access support dashboard logical time:

1. Go to Kibana, click **Dashboards** > **Support Dashboard Logical Time**.
   The Support Dashboard logical time is displayed.



2. Adjust the time range on the top-right corner of the page based on your requirement to view the statistics.



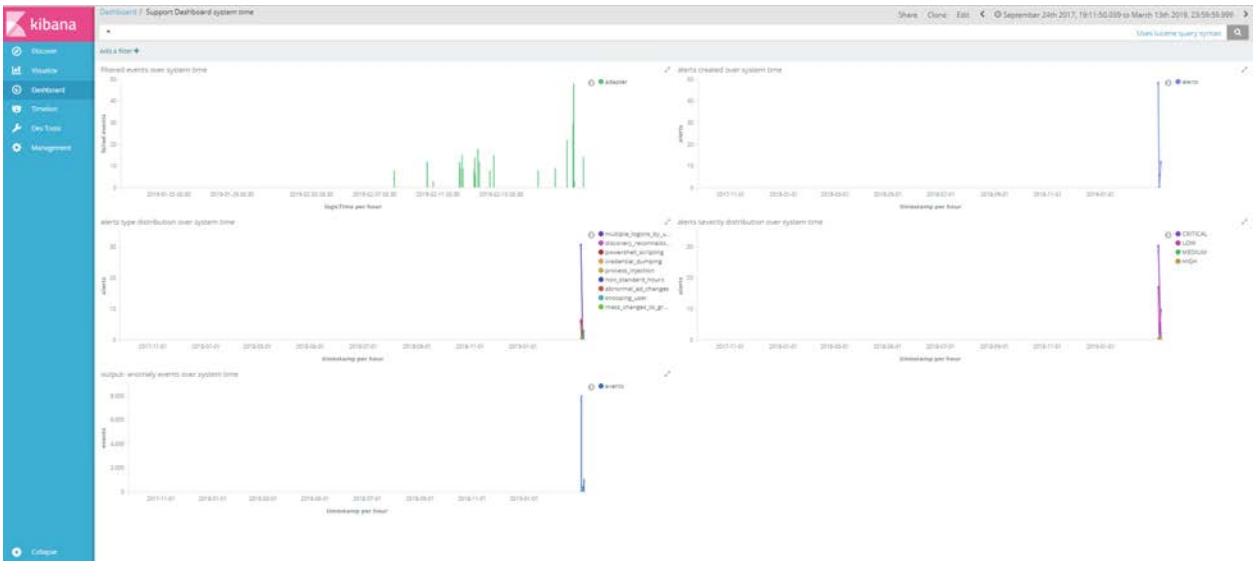
Support Dashboard System Time

The support dashboard system time allows you to monitor the system time when events are processed.

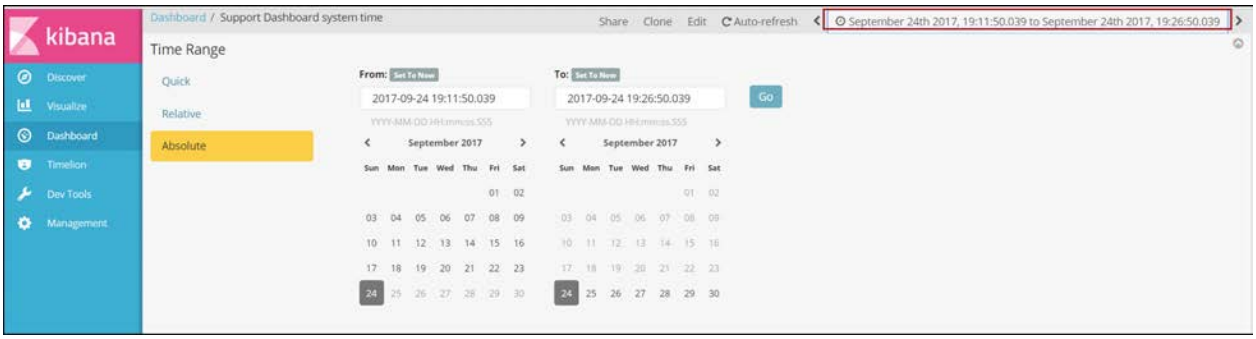
- The amount of filtered events over time per schema.
- The total number of alerts generated.
- The alert types distribution.
- The events that are related to an alert.

To access support dashboard system Time:

1. Go to Kibana, click **Dashboards** > **Support Dashboard system Time**.



2. Adjust the time range on the top-right corner of the page to view the statistics.

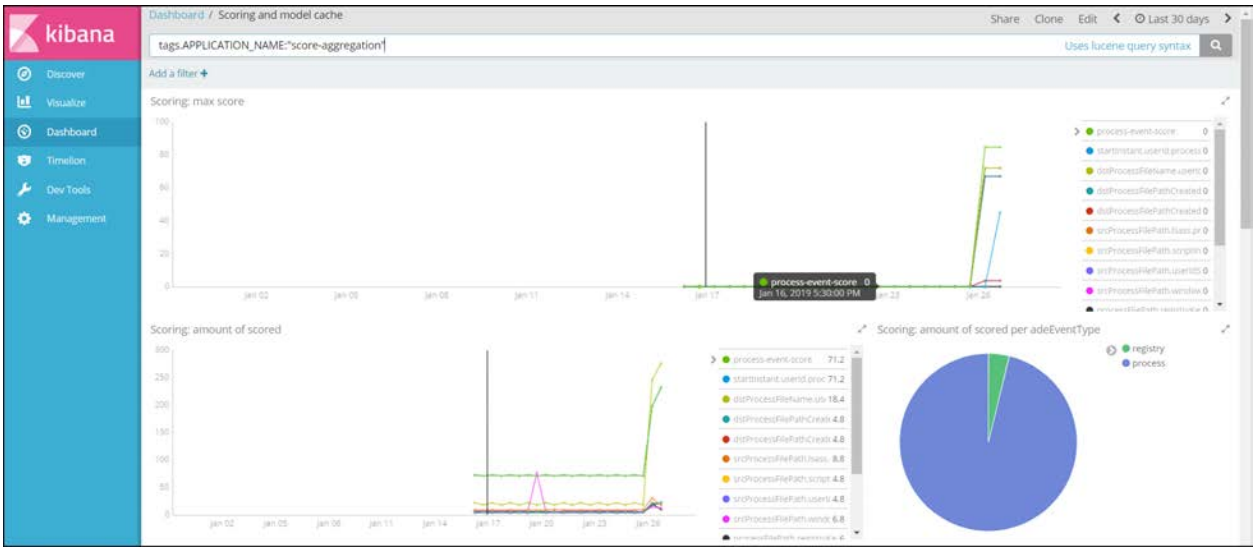


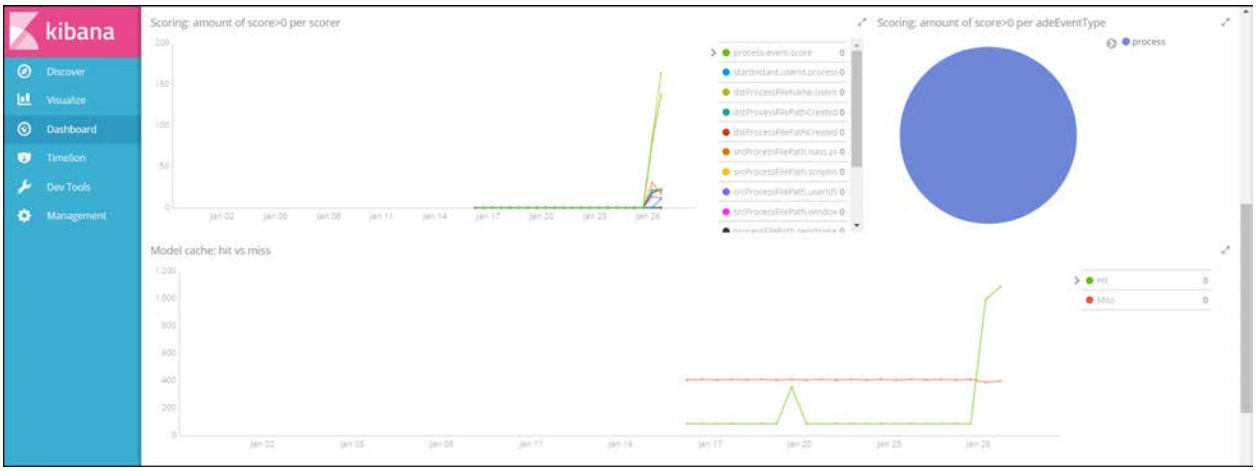
Scoring and Model Cache

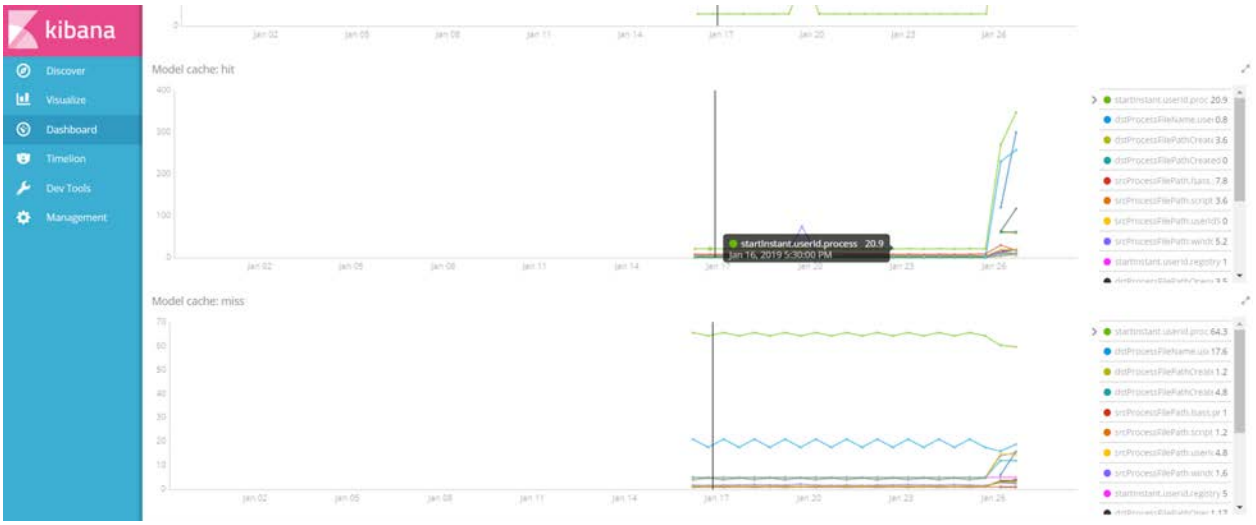
The Scoring and Model cache dashboard provides the capability to view events being scored.

To access scoring and model cache dashboard:

1. Go to Kibana, click **Dashboards** > **Scoring and Model Cache**.
   The Scoring and model cache dashboard is displayed.

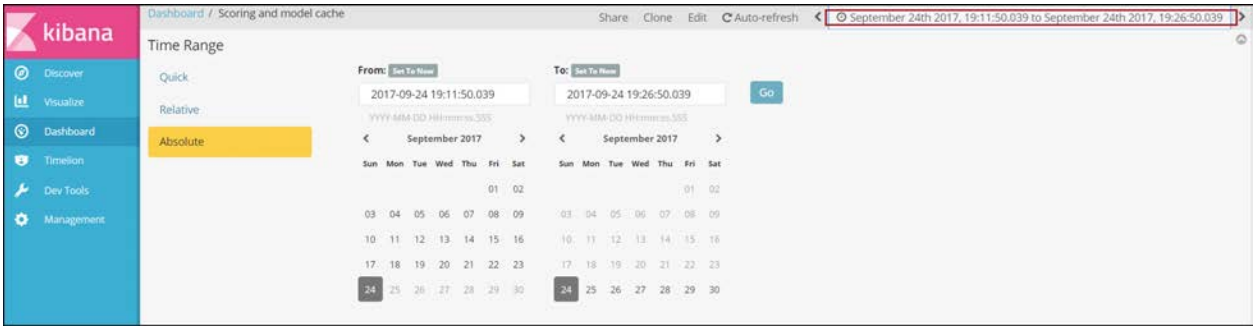

kibana
Dashboard / Scoring and model cache
Share
Clone
Edit
Last 30 days
tags.APPLICATION_NAME:"score-aggregation"
Uses lucene query syntax
Add a filter
Discover
Visualize
Dashboard
Timelion
Dev Tools
Management
Scoring: max score
process-event-score 0
Jan 16, 2019 5:30:00 PM
Scoring: amount of scored
Scoring: amount of scored per adeEventType
registry
process


kibana
Discover
Visualize
Dashboard
Timelion
Dev Tools
Management
Scoring: amount of score>0 per scorer
Scoring: amount of score>0 per adeEventType
process
Model cache: hit vs miss


kibana
Discover
Visualize
Dashboard
Timelion
Dev Tools
Management
Model cache: hit
startInstant.userId.process 20.9
Jan 16, 2019 5:30:00 PM
Model cache: miss



2. Adjust the time range on the top-right corner of the page to view the statistics.



Airflow

Airflow is a tool for describing, executing, and monitoring the UEBA tasks. In Airflow, a DAG is a collection of all tasks you want to run, organized based on the schemas that reflects their relationships and dependencies. For example, schemas such as Active Directory, Authentication, File, Process, TLS and Registry. Each schema is divided into two:

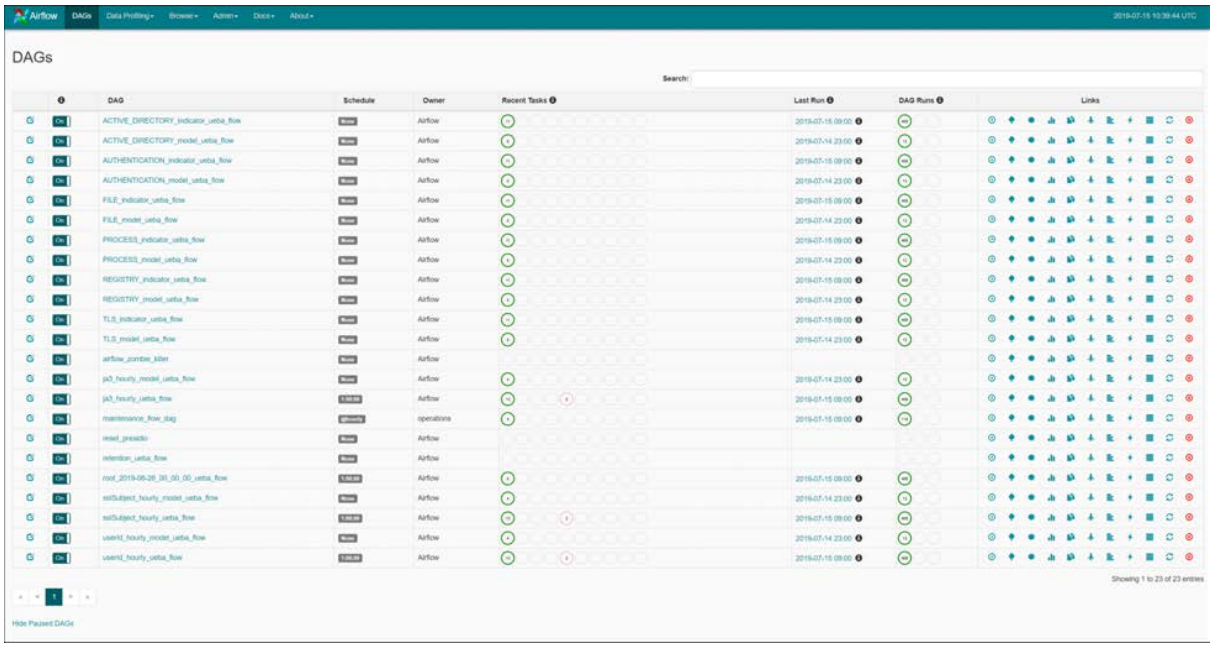
- Indicator DAG which is responsible to read events from broker and score the events based on the models.
- Model DAG which is responsible in building the models.

You can monitor the scheduled task by seeing how many tasks are successful, failed, or currently running.

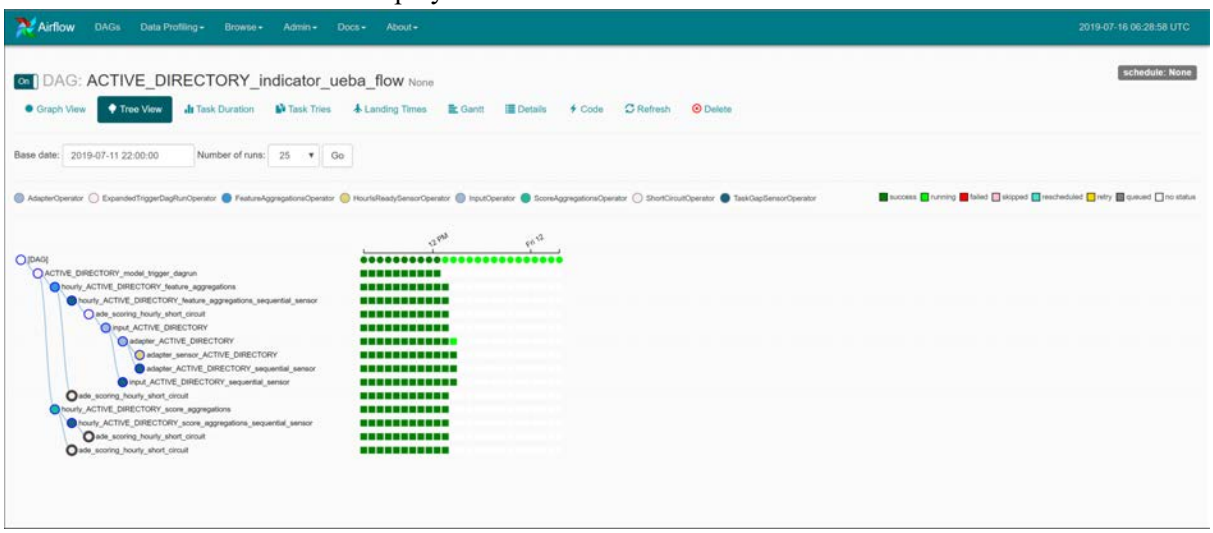
There are several DAGs and each DAG is a workflow.

To monitor UEBA service tasks, perform the following:

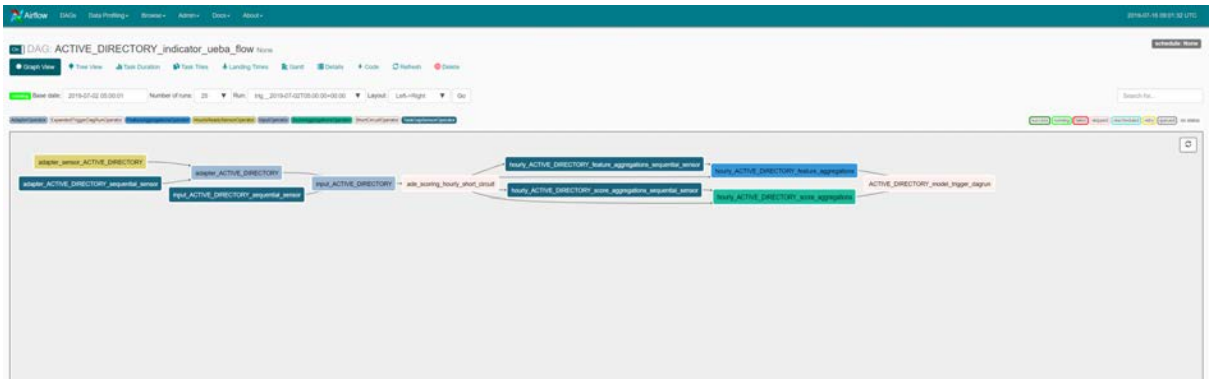
1. Go to **Airflow**.
   The DAGs view is displayed.



2. In the **DAG Runs** section, see the status of the tasks. For example, how many tasks are successful, failed or currently running.
3. To view the different tasks associated with the DAG, click **Tree view**.
   The Tree view of the DAG is displayed.



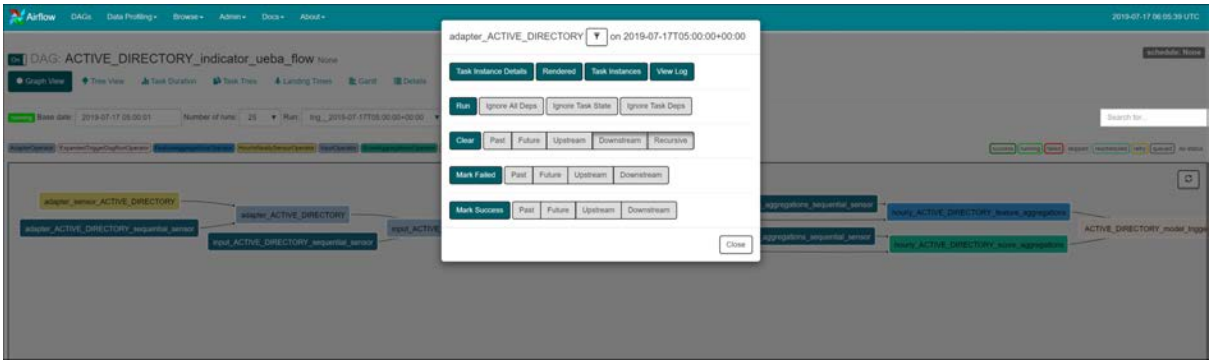
4. To view the DAG's dependencies and the current status of a specific task, in the DAG, click **Graph view**.



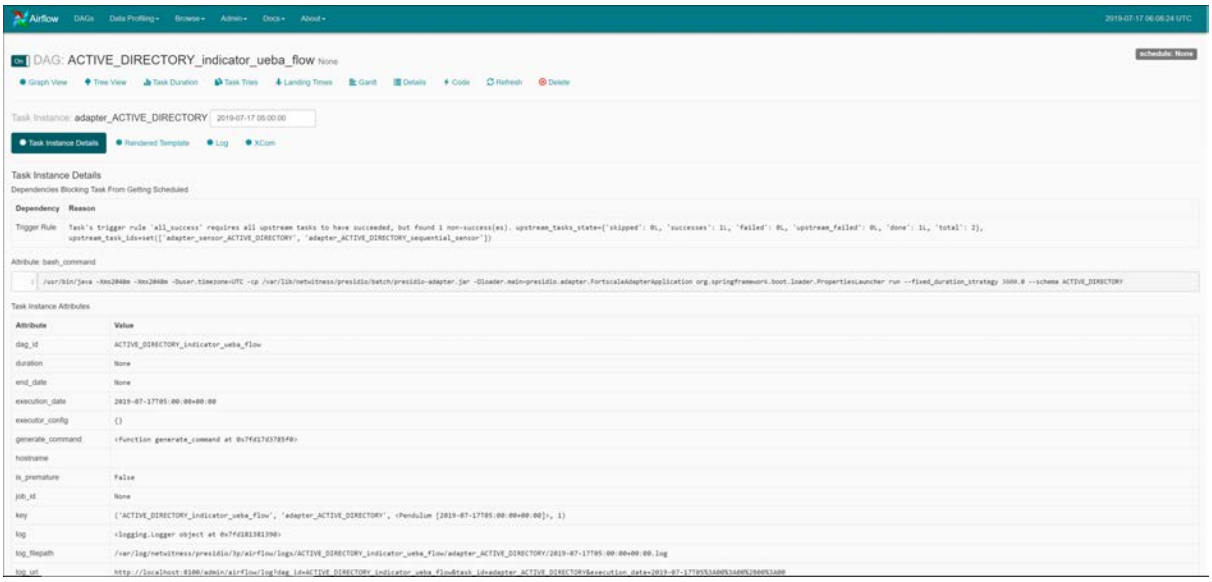
In the **Graph** view, hover over the task to see the status of the specific task.



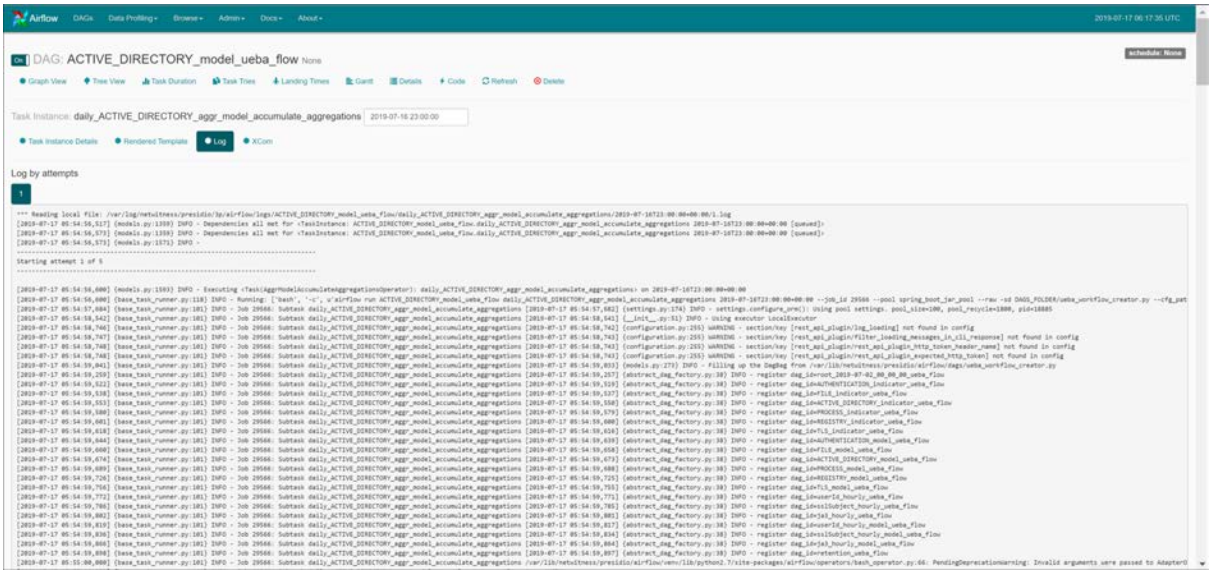
For detailed information about the specific task, click **Task** and click **Task Instance Details.**



The Task Instance Details view is displayed.



To view the logs of the specific task, click **Log**.



Note: After you begin to run a DAG, schemas cannot be removed from UEBA, otherwise the process will stop. For more information see, 'Troubleshooting UEBA Configurations' topic in the *UEBA Configuration Guide*..

Reference

This section provides information about the RSA NetWitness UEBA user interface.

Overview Tab

The Overview tab provides an initial view into the recent and most important user or network entity activities in the environment. Each panel shows either prioritized incidents for investigation or consolidated metrics reflecting potential risks to the enterprise.

Workflow



What do you want to do?

User Role	I want to ...	Documentation
UEBA Analyst	View top ten high-risk users or network entities.*	Identify High-Risk User or Network Entity
UEBA Analyst	View risky user or network entities, and watchlist or network entities.*	Identify High-Risk User or Network Entity
UEBA Analyst	View user based on alert type and indicator.	Identify High-Risk User or Network Entity
UEBA Analyst	Investigate alerts in my environment.	Investigate Top Alerts
UEBA Analyst	Begin an investigation of critical alerts.	Investigate Top Alerts
UEBA Analyst	Sort alerts to focus my investigation.	Filter Alerts
UEBA Analyst	Investigate threat indicators.	Investigate Events
UEBA Analyst	Export alert data.	Manage Top Alerts

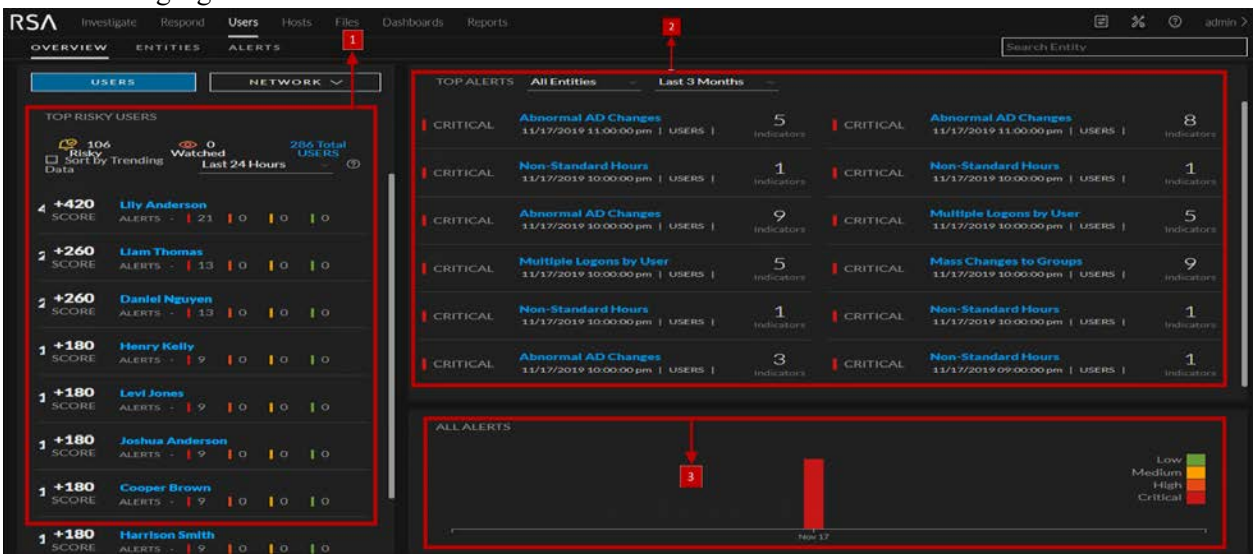
*You can complete the tasks here.

Related Topics

- Begin an Investigation of High-Risk User Or Network Entity
- Investigate Top Alerts
- Filter Alerts
- Manage Top Alerts

Quick Look

The following figure shows the Overview tab.



The Overview tab consists of the following panels:

1 Top Risky User or Network entities panel
2 Top Alerts panel
3 Alerts Severity panel

Top Risky User or Network Entity Panel

The High Risk User or Network entities panel lists the top ten high-risk users or network entities along with the user or network entity score.

In this example, the following table describes the high risk users panel elements.

Name	Description
Risky	All user or network entities with a risk score greater than 0.
Watched	All user or network entities who are currently flagged as Watched.
Total Users	All user or network entities in the network.

Name	Description
User or Network entity name	The name of the user or network entity.
User or Network Entity Score	The score of the user or network entity, with the color indicating the severity of the score. • red indicates critical • orange represents a high risk • yellow indicates a medium risk • green represents a low risk

Top Alerts Panel

The Top Alerts panel displays a list of alerts for the associated user or network entity, severity, alert creation date, and number of indicators. The list consists of the top ten alerts in the Last 24 Hours, Last 7 days, Last 1 Month and Last 3 Months.

The following table describes the top alerts panel elements.

Name	Description
Severity Icon	The alert severity icon. The options are Critical, High, Medium, or Low.
Alert Name	The name of the alert.
Alert Creation Date	The date when an alert is generated.
Number of Indicators	The number of indicators associated with the alert.

Alerts Severity Panel

The Alert Severity panel graphically displays the number of alerts.

The following table describes alert severity panel elements.

Name	Description
Severity level	The severity is color coded, where red indicates a Critical alert, orange represents a High risk alert, yellow indicates a Medium risk alert, and green represents a Low risk alert. For example: 

Entities Tab

The Entities tab is a proactive threat hunting console. You can use behavioral filters to build use-case driven target lists, and to continuously monitor the environment for specific risky behavior patterns.

Workflow



What do you want to do?

User Role	I want to ...	Documentation
UEBA Analyst	View high-risk users or network entities*.	Identify High-Risk User or Network Entity
UEBA Analyst	View user or network entity based on alert type and indicator*.	Identify High-Risk User or Network Entity
UEBA Analyst	Begin an investigation of high-risk user or network entities.	Begin an Investigation of High-Risk User Or Network Entity
UEBA Analyst	Take action on high-risk users or network entities*.	Take Action on High-Risk User or Network Entity
UEBA Analyst	Export high-risk users or network entities*.	Export a list of High-Risk User or Network Entity
UEBA Analyst	Begin an investigation of critical alerts.	Investigate Top Alerts
UEBA Analyst	Investigate threat indicators.	Investigate Events

*You can complete the tasks here.

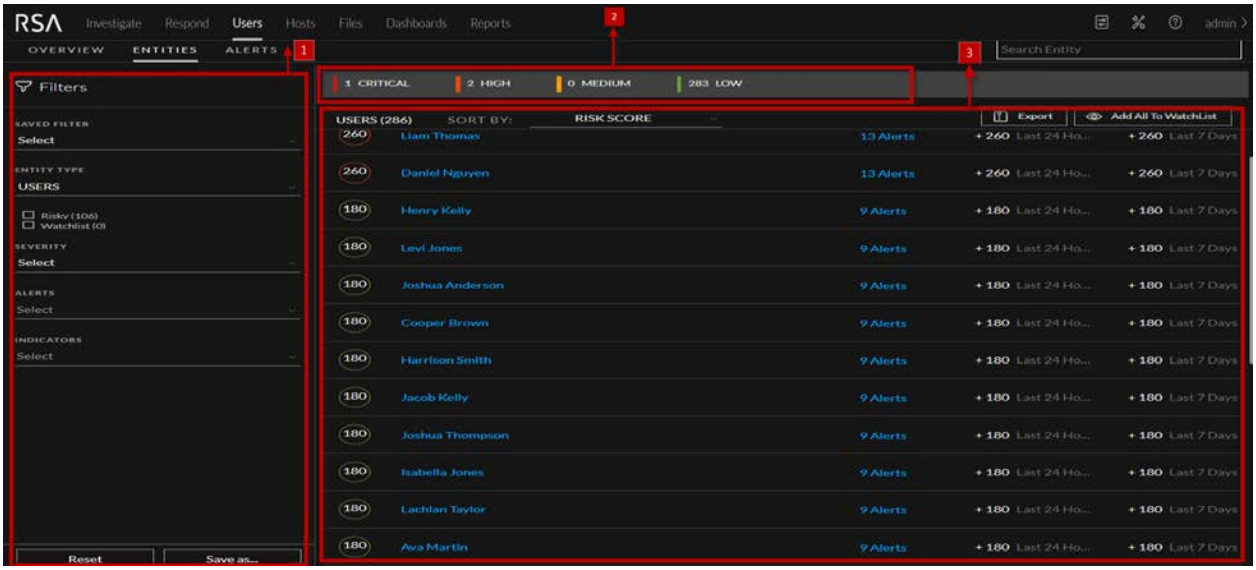
Related Topics

- Begin an Investigation of High-Risk User Or Network Entity
- Investigate Top Alerts

- Filter Alerts
- Investigate Events
- Export a list of High-Risk User or Network Entity

Quick Look

The following figure shows the Entities tab.



The Users tab consists of the following panels:

1 Filters panel
2 Risk Indicator Panel
3 User or Entity List panel

Filters Panel

The Filters panel lists two pre-defined filters, with the number of users associated with each in parentheses, and the list of behavioral profiles that are saved as favorites.

Filter Type	Description
Saved Filter	Previously saved behavioral filters.
Entity Type	Entity type such as Users, JA3, and SSL.
Risky User or Network Entities	All user or network entities with a risk score greater than 0.
Watchlist User or Network Entities	All user or network entities that are currently flagged as Watched.
Severity	Severity type, such as critical, high, medium and low.

Filter Type	Description
Alerts	Any of the existing alert types that describe the supported distinct use cases (Brute Force Attempt, Snooping User, Abnormal AD Change, Data Exfiltration).
Indicators	Any of the existing behavioral features modeled by NetWitness UEBA. This filter can also be used to target only alerts from a specific data source or application.
Reset	Reset the filter.
Save as	Save the filters as favorites.

Risk Indicator panel

The Risk indicator provides a severity-based breakdown of the target user or network entities.



The following table describes the risk indicator panel elements.

Color	Severity
Red	Critical
Orange	High
Yellow	Medium
Green	Low

Entities List Panel

The Entities List panel displays the list of all the user or network entities in your environment along with the user or network entity score and number of alerts associated with the user or network entity.

The following table describes the Entities List panel elements.

User Data	Description
Username or Network entity name	The name of the user or network entity.
Score	The user or the network entity.
Number of alerts	The total number of alerts generated for the user or network entity.
Sort by	The Sort by drop-down menu allows you to select the sorting method for the list. The options are: Risk Score, Name, Alerts, Trending last 24 hours, and Trending last 7 days.
Export	Export a list of all user or network entities and their scores in a .csv file format.
Add All to Watchlist	Adds all user or network entities in the filtered view to the watchlist.
Search Entity	Searches for a user name or a network entity that you typed, allows you to select it from the list that is displayed matching your entry.

Alerts Tab

The Alerts tab displays details about all alerts in your environment. You can view forensic information about suspicious activity in your environment that is based on a specific timeframe.

Workflow



What do you want to do?

User Role	I want to ...	Documentation
UEBA Analyst	Investigate alerts in my environment*.	Investigate Top Alerts

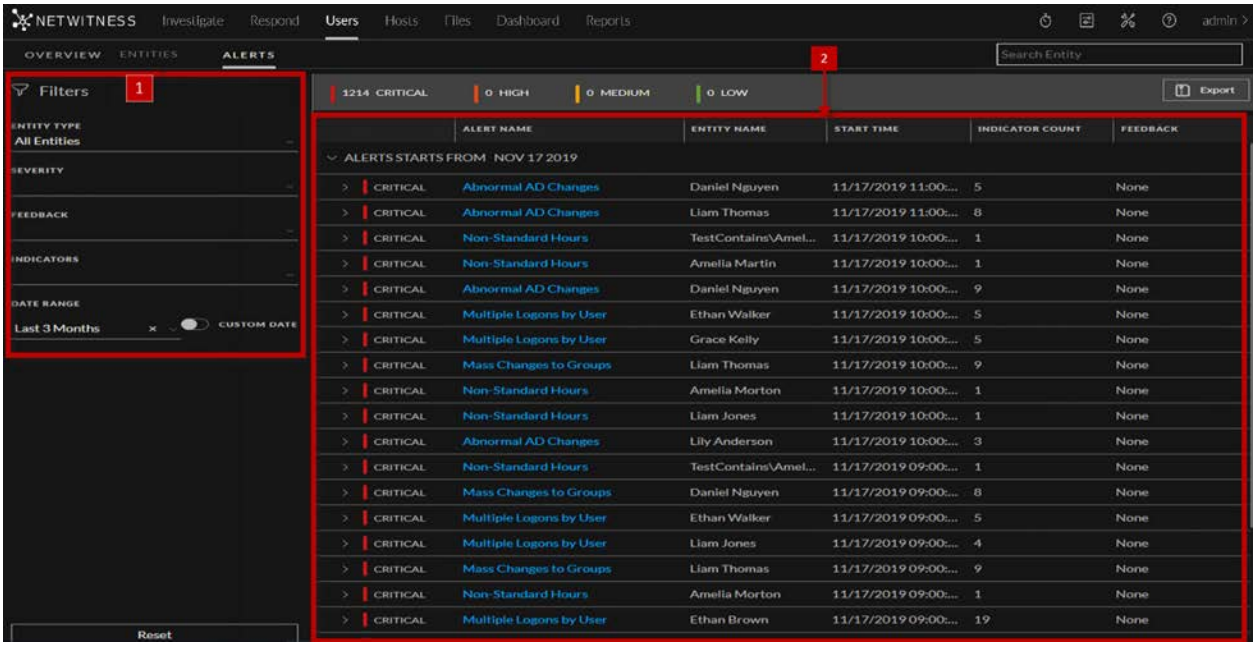
User Role	I want to ...	Documentation
UEBA Analyst	Sort alerts to focus my investigation*.	Filter Alerts
UEBA Analyst	Investigate incidents based on threat indicators*.	Investigate Events
UEBA Analyst	Share alert data in spreadsheet format.	Manage Top Alerts

*You can complete the tasks here.

Related Topics

- Investigate Top Alerts
- Filter Alerts
- Investigate Events
- Manage Top Alerts

Quick Look



The Alerts tab consists of the following panels:

1 Filters panel
2 Alerts panel

Filters Panel

Use the filters panel to refine your investigation of alerts. The filters are automatically applied as you make your selections. You can reset all currently set filters by clicking **Reset**.

The following table describes the filters types.

Filter Name	Description	Options
Entity Type	Filters the list of alerts to include only alerts for a specific user name.	All Entities, Users, JA3, and SSL
Severity	Filters the list of alerts to include alerts for one or more severity levels.	Critical, High, Medium, or Low.
Feedback	Filters the list of alerts to include alerts for one or more feedback types.	Select All, No Feedback, or Not a Risk.
Indicators	Filters the list of alerts to include alerts for one or more indicators.	Examples of indicators are: • Active Directory - Abnormal Logon Time • Authentication - Logged onto Multiple Computers • Multiple File Access Failures
Date Range	Filters the list of alerts to include alerts created during a specific time range.	Last 7 days, Last 2 weeks, Last 1 month, Last 3 months, Last 6 month or specified range.

Alerts Panel

The Alerts panel displays the following information for each alert:

- Severity Icon: An icon next to the alert name that indicates the severity level of the alert.
- Alert Name: The name of the alert and the alert timeframe.
- Entity Name: The name of the entity that generated the alert.
- Start Time: The date and time when this alert was first detected.
- Indicator Count: The number of unique behavior anomalies (indicators) associated with the alert.
- Feedback: Indicates if a feedback value assigned for the alert.

At the beginning of each alert line is an arrow that expands the alert to display additional details. When you expand, the following fields are displayed:

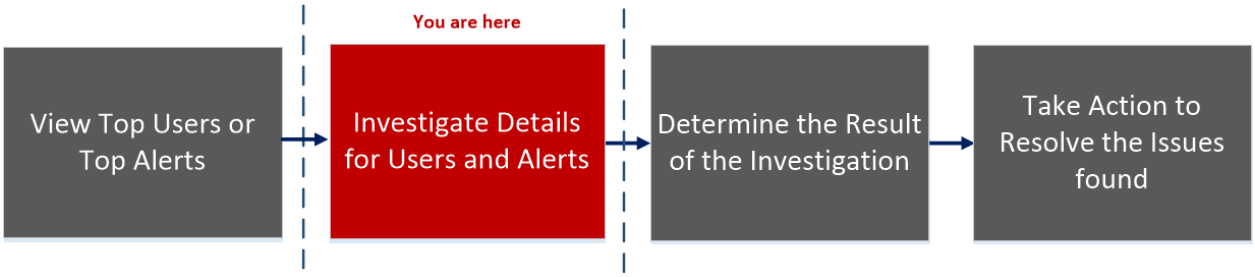
- Indicator Name – The name of each unique indicator that is associated with the alert.
- Anomaly Value – The indicator's value, representing the deviation amount or value as it differs from the user's normal behavior.
- Data Source – The type of data where the indicator was found.
- Start Time – The date and time when this indicator was first detected.

The data that is currently displayed in the central pane can be exported to a .csv file by clicking **Export** at the top right of the pane.

User or Network Entity Profile View

The **User Network Entity Profile** view provides detailed information about all alerts and related indicators of a user or network entity.

Workflow



What do you want to do?

User Role	I want to ...	Documentation
UEBA Analyst	View high-risk user or network entities*	Identify High-Risk User or Network Entity
UEBA Analyst	Begin an investigation of high-risk user or network entities*	Begin an Investigation of High-Risk User Or Network Entity
UEBA Analyst	Take action on high-risk user or network entities.	Take Action on High-Risk User or Network Entity
UEBA Analyst	Export high-risk user or network entities.	Export a list of High-Risk User or Network Entity
UEBA Analyst	Begin an investigation of critical alerts*	Investigate Top Alerts

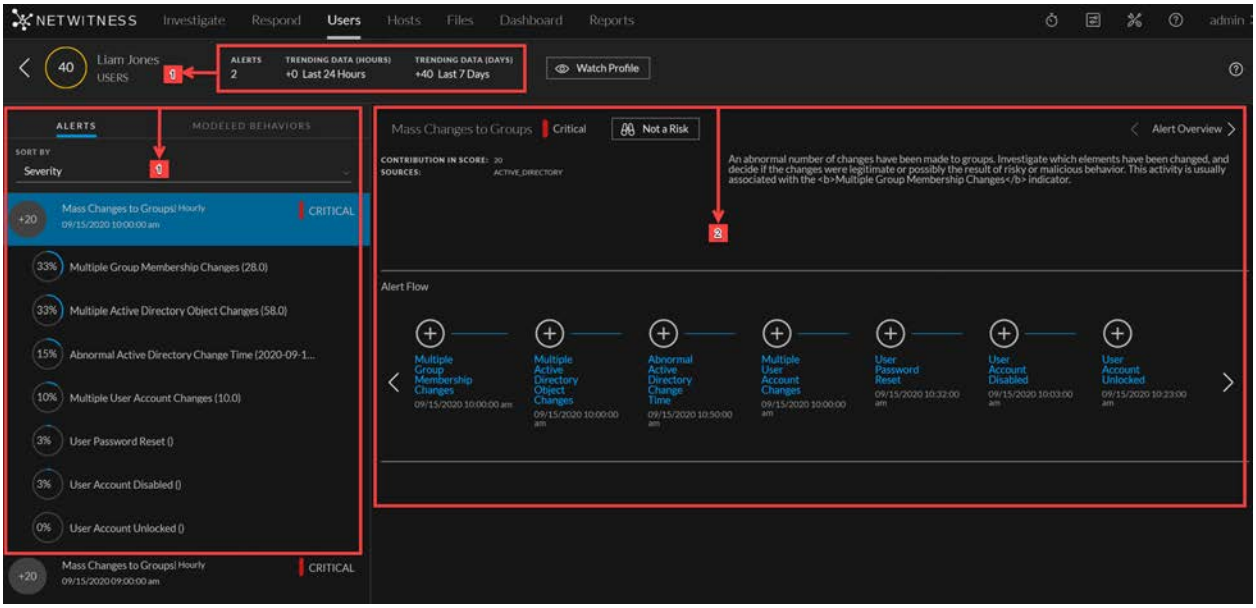
User Role	I want to ...	Documentation
UEBA Analyst	Investigate threat indicators.	Investigate Events
UEBA Analyst	View Modeled Behaviors for users **Note:** From 11.5.1 version and later, Modeled Behaviors for users can be viewed.	View Modeled Behaviors

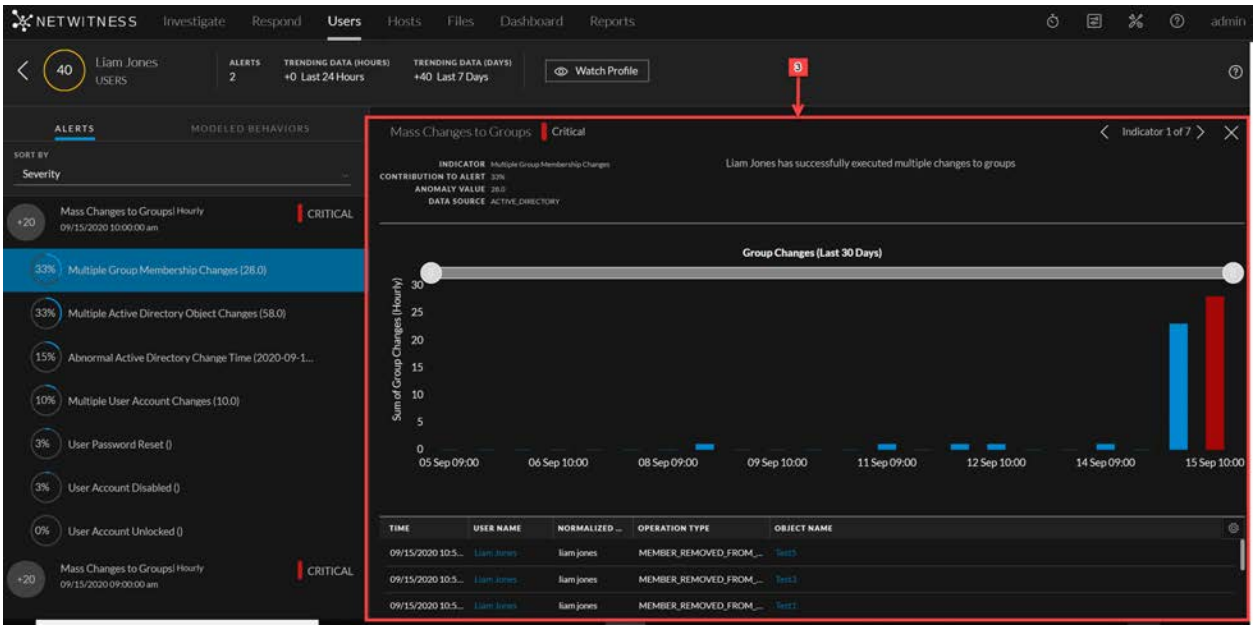
*You can complete the tasks here.

Related Topics

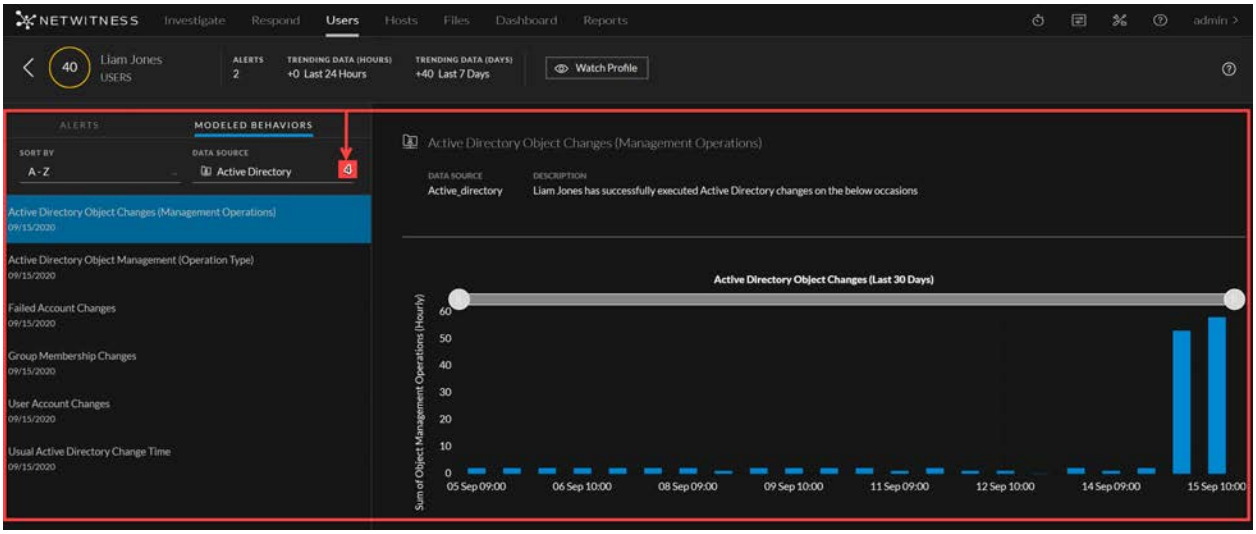
- Begin an Investigation of High-Risk User Or Network Entity
- Investigate Top Alerts
- Filter Alerts
- Investigate Events
- Export a list of High-Risk User or Network Entity
- View Modeled Behaviors

Quick Look





The following figure shows the User Modeled Behaviors view.



The Users Profile consist of the following panels:

1. User Risk Score panel
2. Alerts Flow panel
3. Indicator panel
4. Modeled Behaviors panel

User or Network Entity Risk Score Panel

The User or Network Entity Risk Score panel contains the following information:

Name	Description
User Score	The user score of the user highlighted based on the severity.
Alerts	The total number of alerts generated for the user in the last 90 days.
Trending Data (Hours)	The trending data for last 24 hours shows the increase in the user's score in the last 24 hours.
Trending Data (Days)	The trending data for last 7 days shows the increase in the user's score in the last 7 days.
Alerts	The following information is displayed: • alert names • severity level icon • start date and time for the alert • timeframe of the alert (Hourly) • risk score of the alert (+20) • list of alert indicator names and the number of times the indicator events occurred.
Sort by	The alerts are sorted based on Severity and Date. By default, it is sorted by severity.

Alert Flow Panel

The Alert Flow panel displays the following information:

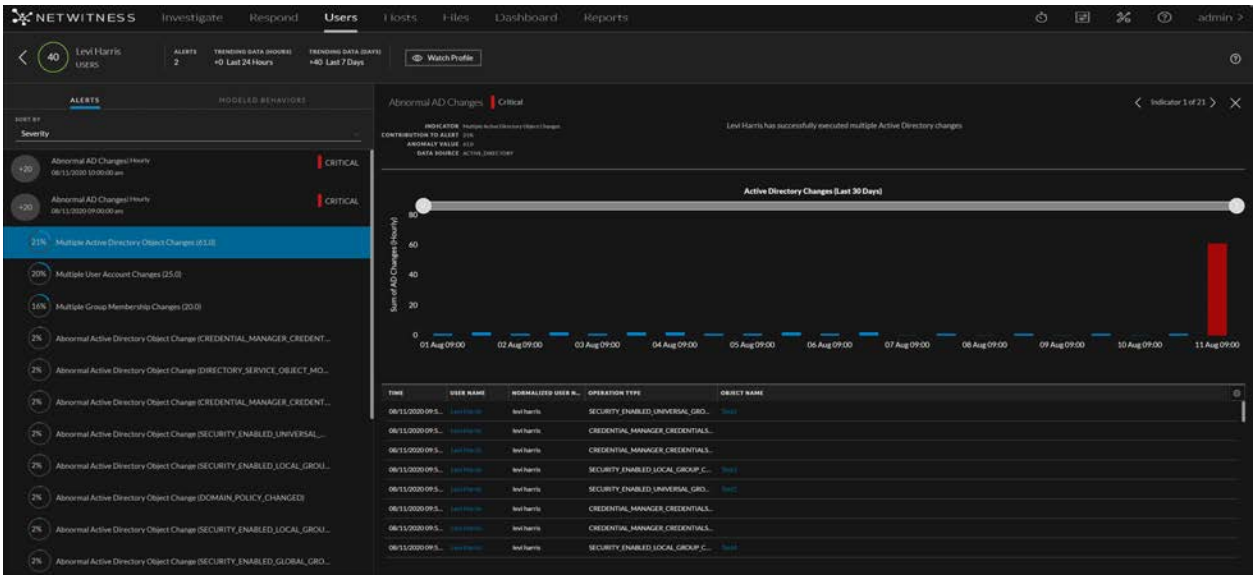
Name	Description
Alert name	The name of the alert.
Time frame	The timeframe of the alert (hourly).
Severity level	The severity of the alert.
Contribution in score	The contribution to the user score value (for example, +20).
Sources	The data sources for the alert (for example, Active Directory).
Tamerlane graph	The timeline of events that are related to the formation of the alert.

Indicator Panel

Click on a graph icon in the Alert Flow panel to open the Indicator panel. The following table describes the indicator panel elements:

Name	Description
Indicator	The name of the indicator with timeframe of the indicator in parentheses. For example, Multiple Group Membership Changes (Hourly).
Contribution to Alert	The alert contribution percentage.
Anomaly Value	The anomaly value.
Data source	The data source from where the alert is triggered.

In the Indicator panel the events table list events specific to the data sources.



- **Common events for User Entity**

The following tables list events specific to all the data sources.

Event Name	Description
Time	The date and time when an event is triggered.
Username	The name of user for whom an indicator is triggered.
Normalized user name	The name of user for whom an indicator is triggered.
Operation Type	The action performed by the user. For example, Member Added To Group.

Event Name	Description
Result	The status of the action performed by the user.

- **Windows File Servers**

The following tables list events specific to Windows file servers.

Event Name	Description
Source Folder Path	Absolute folder path of a file for which an event is triggered.
Source File Path	Absolute file path for which an event is triggered.

- **Active Directory**

The following tables list event specific to Active Directory.

Event Name	Description
Object Name	Object name defined in the Active Directory.

- **Logon Activity**

The following tables list events specific to Logon Activity.

Event Name	Description
Computer	Host name from where an event is triggered.

- **Process**

The following tables list events specific to Process.

Event Name	Description
Machine Name	Name of the host from where this event is triggered for the user.
Source Process	Process triggered by the event
Destination Process	Process triggered by source process.

- **Registry**

The following tables list events specific to Registry.

Event Name	Description
Machine Name	Name of the host from where this event is triggered for the user.
Process Directory	Absolute directory path of the process for which an event is triggered.
Process File Name	Process file name for which an event is triggered.
Registry Key Group	Type of registry key.
Registry Key	Registry key path.
Registry Value Name	Registry value name that is created or modified.
Operation Type	The action performed by the user. For example, Member Added To Group.

Network Entities

The following tables list events specific to JA3 and SSL Subject.

Event Name	Description
Source IP	The IP address from which network data is sent.
Destination IP	The IP address to which network data is sent.
Destination Country	The country name to which the network data is sent.
SSL	The SSL Subject.
Destination Organization	The organization name where the network data is sent.
Domain	The domain name to which the network data is sent.
JA3	The JA3 hash value.
Destination Port	The port number to which the network data is sent.
Source Netname	The name of the source netname.
Number of Bytes Sent	The number of bytes sent.
Destination ASN	
JA3S	The JA3S hash value.

Event Name	Description
Destination Netname	The name of the destination netname.
Number of Bytes Received	The number of bytes received.

Modeled Behaviors Panel

The Modeled Behaviors panel displays the following information:

Name	Description
Modeled Behaviors	The following information is displayed: • The data source names • The date of the user's last activity • Description of the Modeled Behaviors.
Data Source	The data source can be selected from the drop-down.
Sort by	The Modeled Behaviors are sorted based on date and alphabetical order. By default, it is sorted by alphabetical order.

Appendix: NetWitness UEBA Windows Audit Policy

To achieve maximum benefit from NetWitness UEBA, NetWitness recommends that you implement the Windows audit policies described here.

For a base set of policies to audit, see the "Windows Server 2016, Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008 Audit Settings Recommendations" section of this article from Microsoft: Audit Policy Recommendations.

The policies under "Stronger Recommendation" are required, and the following policies, to ensure that all of the required Authentication and Active Directory events are audited:

- Audit Detailed File Share
- Audit File Share
- Audit File System

NetWitness recommends that you enable auditing for both success and failures.

The following Windows events must be audited:

For the Authentication models:

4624	4625	4769	4628

For the AD models:

4670	4717	4720	4722	4723	4724	4725	4726
4727	4728	4729	4730	4731	4732	4733	4734
4735	4737	4738	4739	4740	4741	4742	4743
4754	4755	4756	4757	4758	4764	4767	4794
5136	5376	5377					

For File Access Models:

4660	4663	4670	5145

NetWitness® Platform XDR

Version 12.0

Virtual Host Installation Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Virtual Deployment Overview **6**

Process 6

NetWitness High-Level Deployment Diagram 7

Installation Media 7

Install NetWitness Platform Virtual Host in Virtual Environment **8**

Work Flow: 8

Prerequisites 8

Checklist 9

Step 1a. Create Virtual Machine - VMware 9

Prerequisites 9

Procedure 9

Step 1b. Create Virtual Machine - Microsoft Hyper-V 13

Prerequisites 13

Procedure 14

Step 1c. Create Virtual Machine - Nutanix AHV 21

Prerequisites 21

Procedure 22

Create Image on Nutanix AHV 22

Create VM on Nutanix AHV 24

Step 1d. Create Virtual Machine in Esxi 26

Prerequisites 26

Procedure 27

Create Image on Esxi 27

Create VM on Esxi 28

Step 2. Configure Block Storage to Accommodate NetWitness Platform 30

Task 1. Add New Disk 30

Add New Disk 30

Add New Disk in VMware ESXi 30

Add New Disk in Hyper-V 34

Add New Disk in Nutanix AHV 40

Task 2. Add New Volume and Extend Existing File Systems 43

Admin Server 43

ESAPrimary/ESASecondary/Malware 43

Log Collector 44

Log Decoder 44

Virtual Drive Space Ratios 44

Extending File Systems ... 45
Concentrator ... 47
Virtual Drive Space Ratios ... 47
Extending File Systems ... 47
Archiver ... 49
Decoder ... 50
Virtual Drive Space Ratios ... 50
Extending File Systems ... 50
Endpoint Log Hybrid ... 52
Virtual Drive Space Ratios ... 52
Extending File Systems ... 52
UEBA ... 53
Task 3. Storage Configurations ... 53
Step 3. Install NetWitness Platform ... 54
Prerequisites ... 54
Install NetWitness Platform ... 54
Set Up ESA Hosts ... 62
Install Component Services on Hosts ... 62
Complete Licensing Requirements ... 63
(Optional) Install Warm Standby NW Server ... 63
Step 4. Configure Host-Specific Parameters ... 63
Configure Log Ingest in the Virtual Environment ... 63
Configure Packet Capture in the Virtual Environment ... 63
Set a vSwitch to Promiscuous Mode ... 63
Use of a Third-Party Virtual Tap ... 64
Step 5. Post Installation Tasks ... 65
Event Stream Analysis (ESA) ... 65
Configure Meta Keys on New ESA Hosts to Match Upgraded ESA Hosts in the Same NetWitness Platform Network ... 65
NetWitness Endpoint ... 65
Install Endpoint Log Hybrid ... 66
Configuring Multiple Endpoint Log Hybrids ... 66
Step 1: Install additional Endpoint Log Hybrid ... 66
Step 2: Setup the Endpoint Log Hybrid ... 66
Step 3: Switch to the NetWitness UI and add Hosts ... 67
Add Hosts to the Endpoint Log Hybrid: ... 67
(Optional) Configure an Endpoint Service on an Existing Log Decoder Host ... 68
Do You Need to Install an Endpoint Service onto Separate Hardware ... 68
Install an Endpoint Service Category on an Existing Log Decoder ... 68
NetWitness UEBA ... 69
Install UEBA ... 69

Configure NetWitness UEBA 70
Enable Access Permission for the NetWitness UEBA User Interface 73
Deployment Options 73
Appendix A. Troubleshooting 74
Command Line Interface (CLI) 75
Event Stream Analysis 76
Appendix B. Silent Installation Using CLI 77
Appendix C. Virtual Host Recommended System Requirements 80
Scenario One 80
Scenario Two 83
Scenario Three 86
Scenario Four 88
Legacy Windows Collectors Sizing Guidelines 89
Appendix D. Update the Virtual ESA Host Memory 90

Virtual Deployment Overview

This document provides instructions on the installation and configuration of NetWitness 11.7.0.0 hosts running in a virtual environment.

Virtual hosts have the same functionality as the NetWitness Azure, AWS, and hardware hosts. NetWitness recommends that you perform the following tasks when you set up your virtual environment.

Before you can deploy NetWitness in virtual environment, you need to:

IMPORTANT: Get familiar with the NetWitness Storage Guide to understand the types of drives and volumes needed to support NetWitness instances. For more information, see Storage Guide for NetWitness Platform 11.x.

- Review the recommended compute and memory specifications needed for each NetWitness instance.
- Make sure that you have a NetWitness Throughput license.
- Review of the network architecture and port usage.

Process

The components and topology of a NetWitness network can vary greatly between installations, and should be carefully planned before the process begins. Initial planning includes:

- Consideration of site requirements and safety requirements.
- Support of group aggregation on Archivers and Concentrators, and virtual hosts.

When updating hosts and services, follow recommended guidelines under the "Running in Mixed Mode" topic in the *Host and Services Getting Started Guide*.

You should also become familiar with Hosts, Host Types, and Services as they are used in the context of NetWitness also described in the *Host and Services Getting Started Guide*.

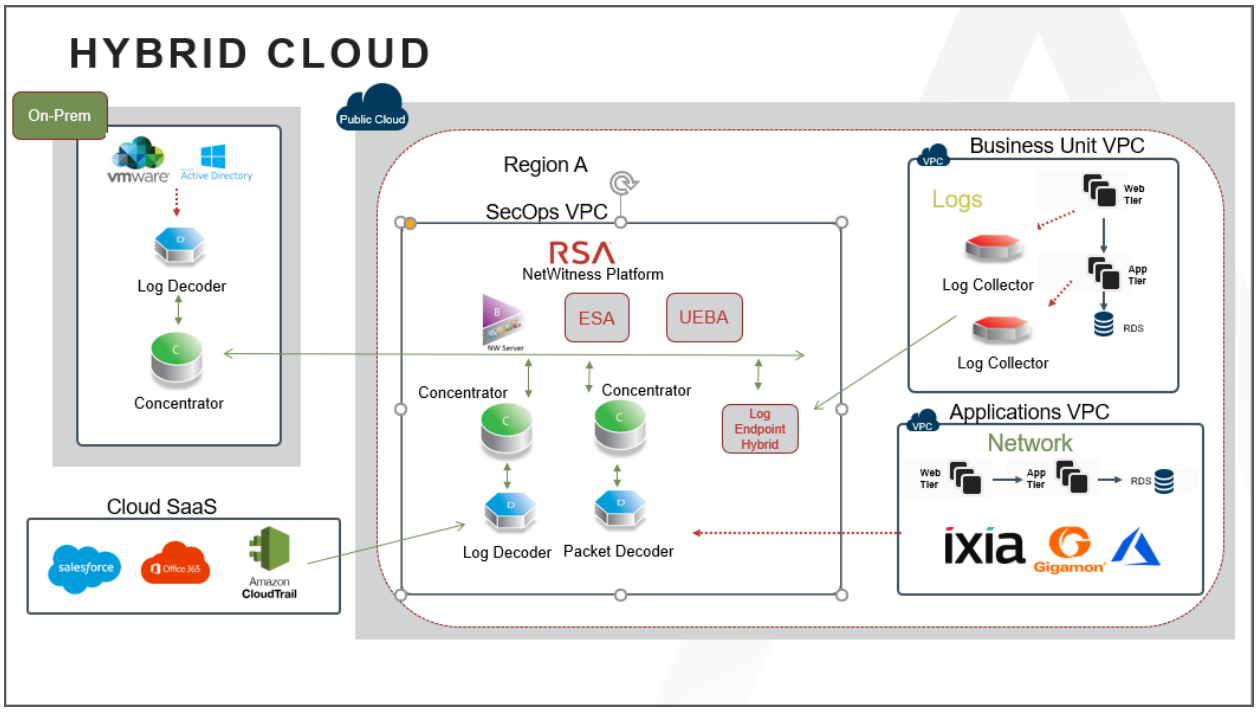
NetWitness High-Level Deployment Diagram

NetWitness is inherently modular. Whether organizations are looking to deploy on-premise or in the cloud, the NetWitness components are decoupled in a way which allows flexible deployment architectures to satisfy a variety of use cases.

The following figure is an example of a hybrid cloud deployment, where the base of the components are residing within the SecOps VPC. Centralizing these components make management easier while keeping network latency to a minimum.

Network, log and endpoint traffic could then be aggregated up to the SecOps VPC. The on-premise location would function just like a normal physical deployment and would be accessible for investigations and analytics.

Cloud SaaS visibility could be captured from a Log Decoder residing in either the cloud or on-premise locations.



For information on VMware concepts, refer to the VMware product documentation.

The virtual hosts are provided as an OVA. You need to deploy the OVA file as a virtual machine in your virtual infrastructure.

Installation Media

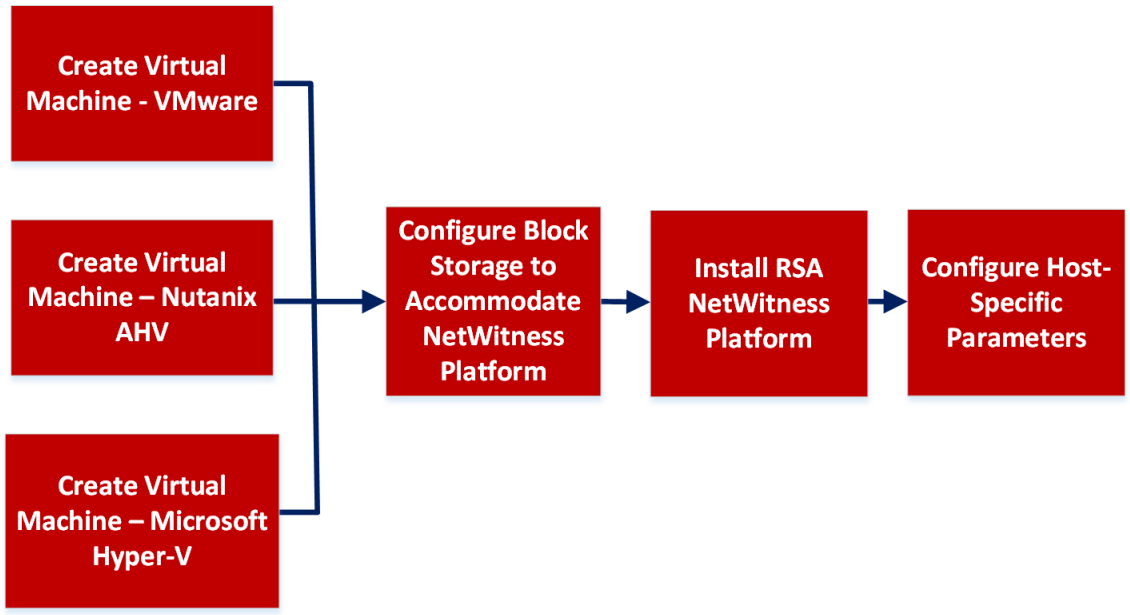
Installation media are in the form of OVA and VHDX packages, which are available for download and installation from Downloads (https://community.netwitness.com/t5/netwitness-platform-downloads/tkb-p/netwitness-downloads). As part of your order fulfillment, NetWitness gives you access to the OVA and VHDX.

Install NetWitness Platform Virtual Host in Virtual Environment

Complete the following procedures according to their numbered sequence to install NetWitness in a virtual environment.

Work Flow:

This figure shows the high-level workflow mandatory for installing NetWitness Platform virtual host.



Note: When you configure databases to accommodate NetWitness Platform, the default database space allocation after you deploy databases from OVA or VHDX will not be adequate to support the NetWitness deployment. You must expand the datastores after initial deployment to avoid any issues. For more information, see Step 2. Configure Block Storage to Accommodate NetWitness Platform.

IMPORTANT: Review the ***Network Architecture and Ports*** topic in the ***Deployment Guide*** in the NetWitness help so that you can configure NetWitness services and your firewalls. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Caution: Do not proceed with the installation until the ports on your firewall are configured.

Prerequisites

Make sure that you have:

- A VMware ESX Server that meets the requirements described in the above section. Supported versions are 7.0, 6.7, 6.5, 6.0, and 5.5.
- Administrator rights to create the virtual machines on the VMware ESX Server.

Checklist

Step	Description	√
1.	Step 1a. Create Virtual Machine - VMware	
2.	Step 1b. Create Virtual Machine - Microsoft Hyper-V	
4.	Step 1c. Create Virtual Machine - Nutanix AHV	
3.	Step 2. Configure Block Storage to Accommodate NetWitness Platform	
5.	Step 3. Install NetWitness Platform	
6.	Step 4. Configure Host-Specific Parameters	
7.	Step 5. Post Installation Tasks	

Step 1a. Create Virtual Machine - VMware

Complete the following steps to deploy the OVA file on the vCenter Server or ESX Server using the vSphere client.

Prerequisites

Make sure that you have:

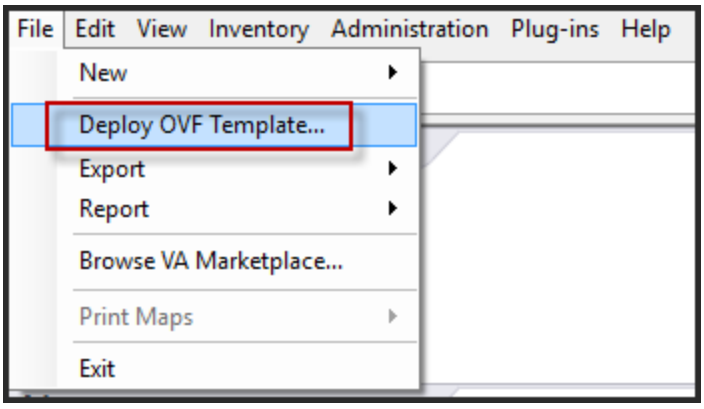
- Network IP addresses, netmask, and gateway IP addresses for the virtual host.
- Network names for all virtual hosts, if you are creating a cluster.
- DNS or host information.
- Password for virtual host access. The default username is `root` and the default password is `netwitness`.
- The NetWitness virtual host package file for example, rsanw-11.7.0.0.xxxx.el7-x86_64.ova. (You download this package from Download Central (https://community.netwitness.com/t5/netwitness-platform-downloads/tkb-p/netwitness-downloads).)

Procedure

Note: The following instructions illustrate an example of deploying an OVA host in the ESXi environment. The screens you see may be different from this example.

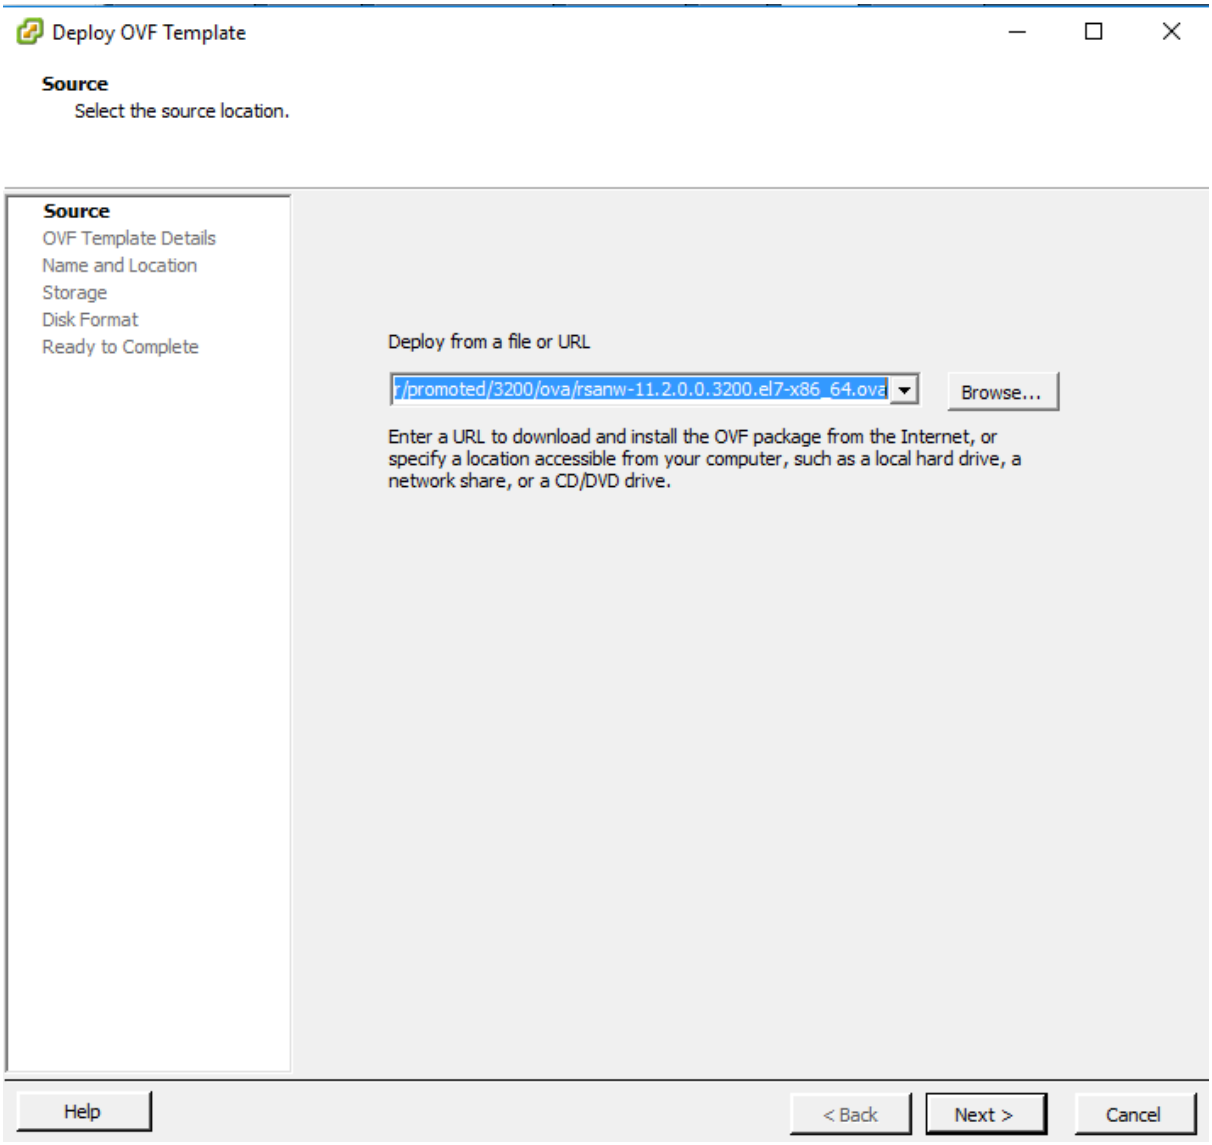
To deploy the OVA host:

1. Log on to the ESXi environment.
2. In the **File** drop-down, select **Deploy OVF Template**.



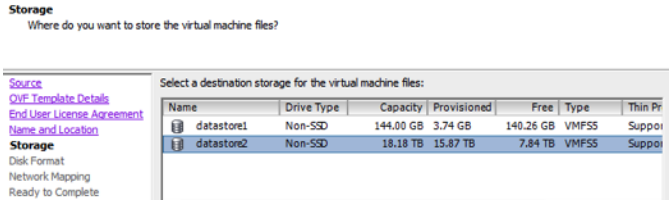
3. The Deploy OVF Template dialog is displayed. In the **Deploy OVF Template** dialog, select the OVF for the host that you want to deploy in the virtual environment (for example, **V11.7**

GOLD\\rsanw-11.7.0.0.xxxx.el7-x86_64.ova), and click **Next**.

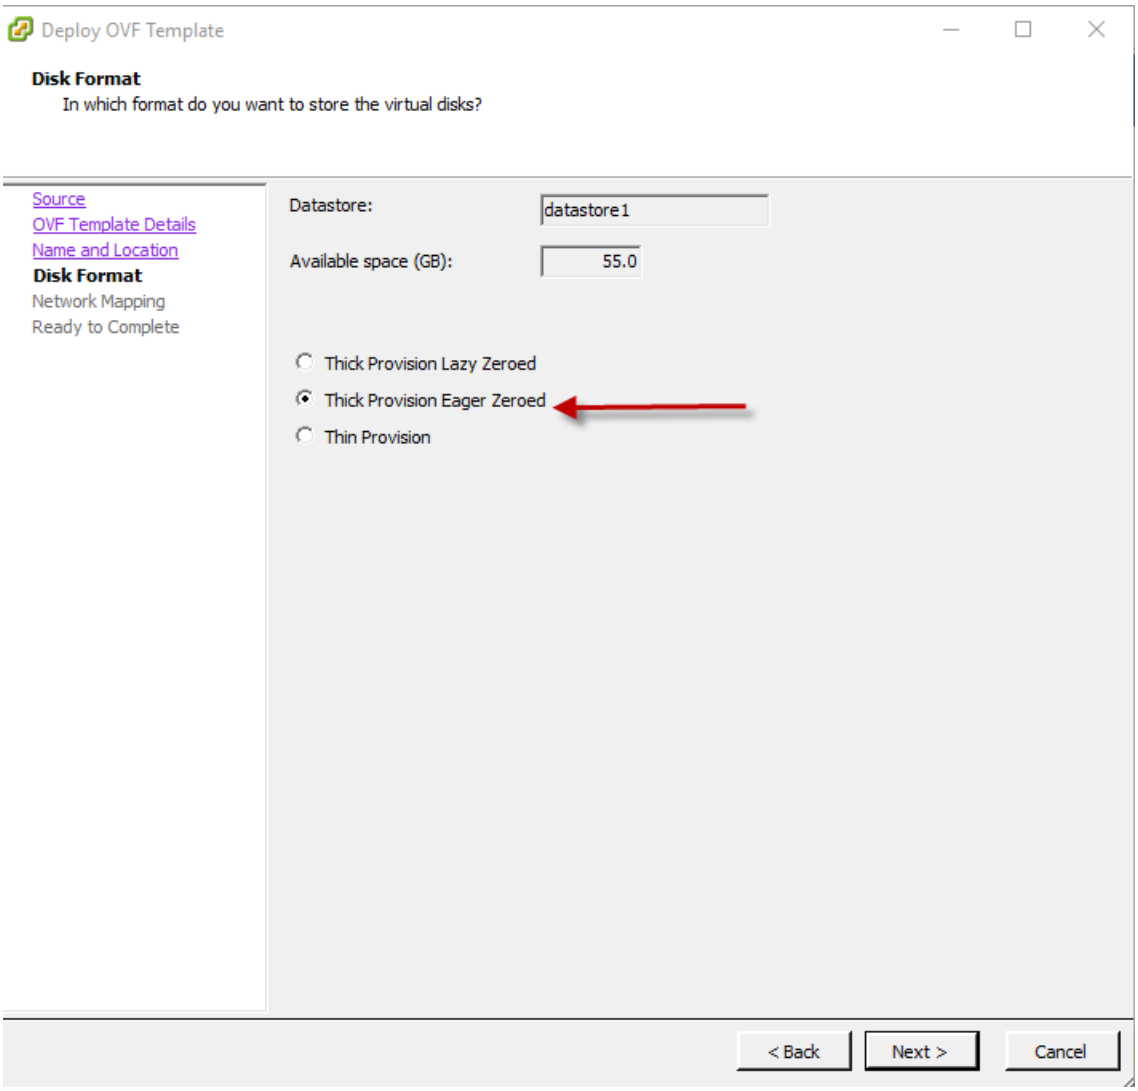


4. The Name and Location dialog is displayed. The designated name does not reflect the server hostname. The name displayed is useful for inventory reference from within ESXi.
5. Make a note of the name, and click **Next**.

   Storage Options are displayed.



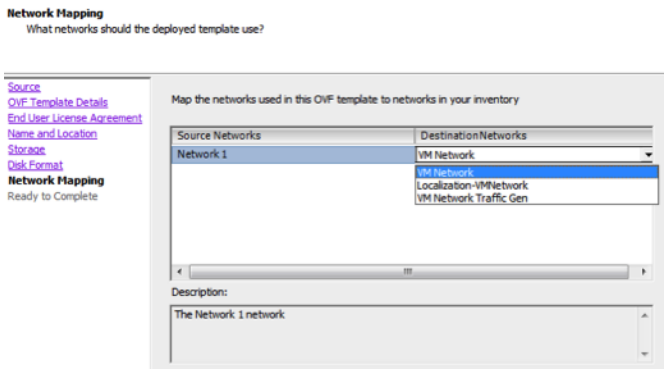
6. For Storage options, designate the datastore location for the virtual host and click **Next**.



Note: This location is for the host operating system (OS) exclusively. It does not have to be the same datastore needed to set up and configure additional volumes for the NetWitness databases on certain hosts (covered in the following sections).

7. Click **Next**.

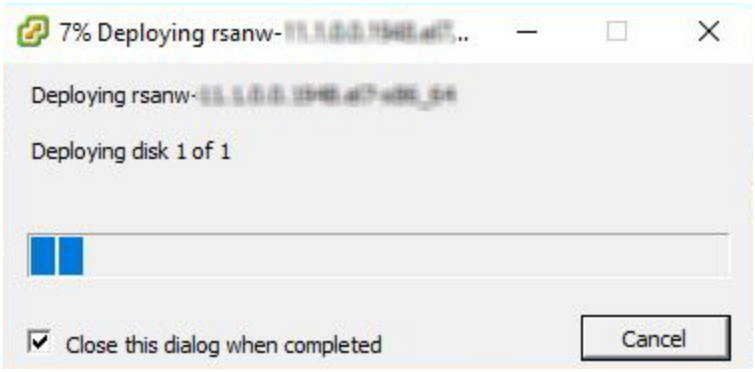
   The Network Mapping options are displayed.



8. Select the **Network label** based on your requirement (For example, VM Network), and click **Next**.

   Note: If you want to configure Network Mapping now, you can select options, but NetWitness recommends that you keep the default values and configure network mapping after you configure the OVA. You configure the OVA in Step 4: Configure Host-Specific Parameters.

   A status window showing deployment status is displayed.



   After the process is complete, the new OVA is presented in the designated resource pool visible on ESXi from within vSphere. At this point, the core virtual host is installed, but is still not configured.

Step 1b. Create Virtual Machine - Microsoft Hyper-V

Complete the following steps according to their numbered sequence to deploy virtual host in Hyper-V.

Prerequisites

Make sure that you have:

- Network IP addresses, netmask, and gateway IP addresses for the virtual host.
- Network names for all virtual hosts, if you are creating a cluster.
- DNS or host information.

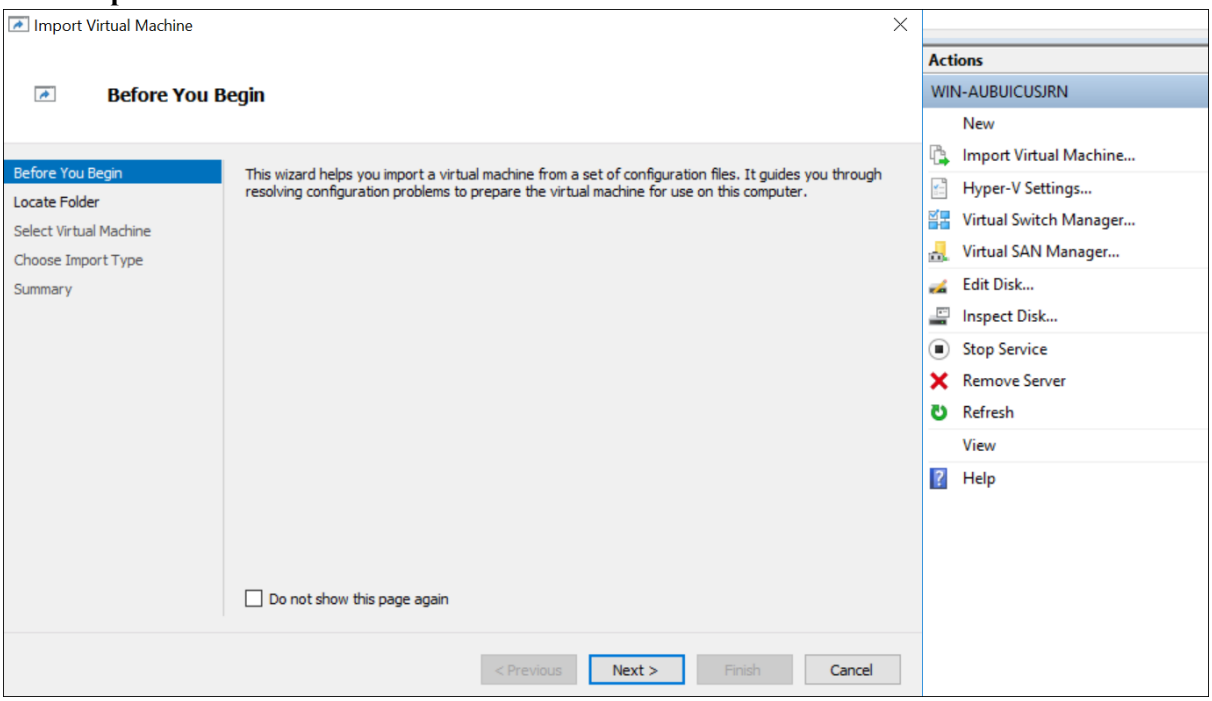
- Password for virtual host access. The default username is `root` and the default password is `netwitness`.
- The NetWitness virtual host package file for example, `rsa-nw-11.7.0.0.3274.zip`. (You download this package from Downloads https://community.netwitness.com/t5/netwitness-platform-downloads/tkb-p/netwitness-downloads)

Procedure

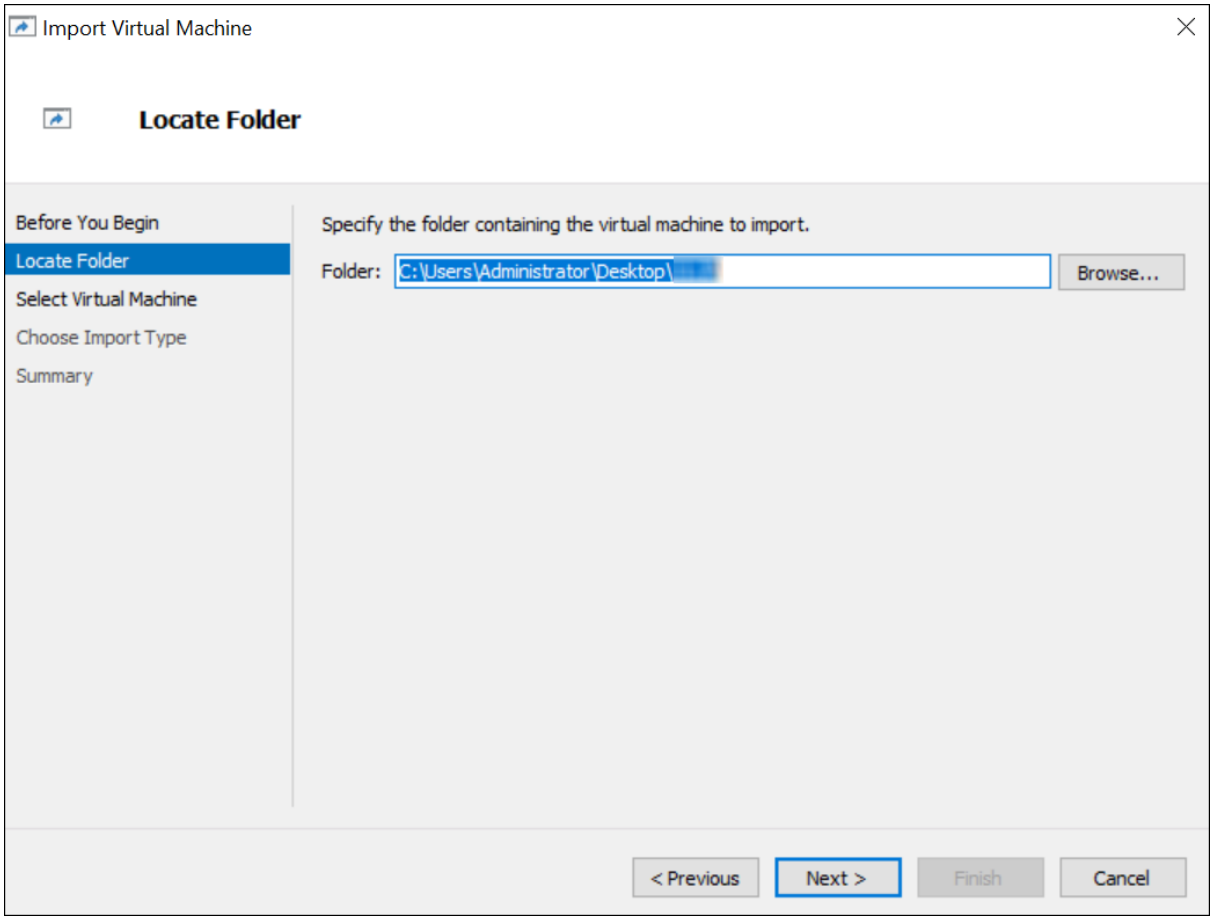
Note: The following instructions illustrate an example of deploying a VM in the Hyper-V environment. The screens you see may be different from this example.

To deploy virtual host in Hyper-V.

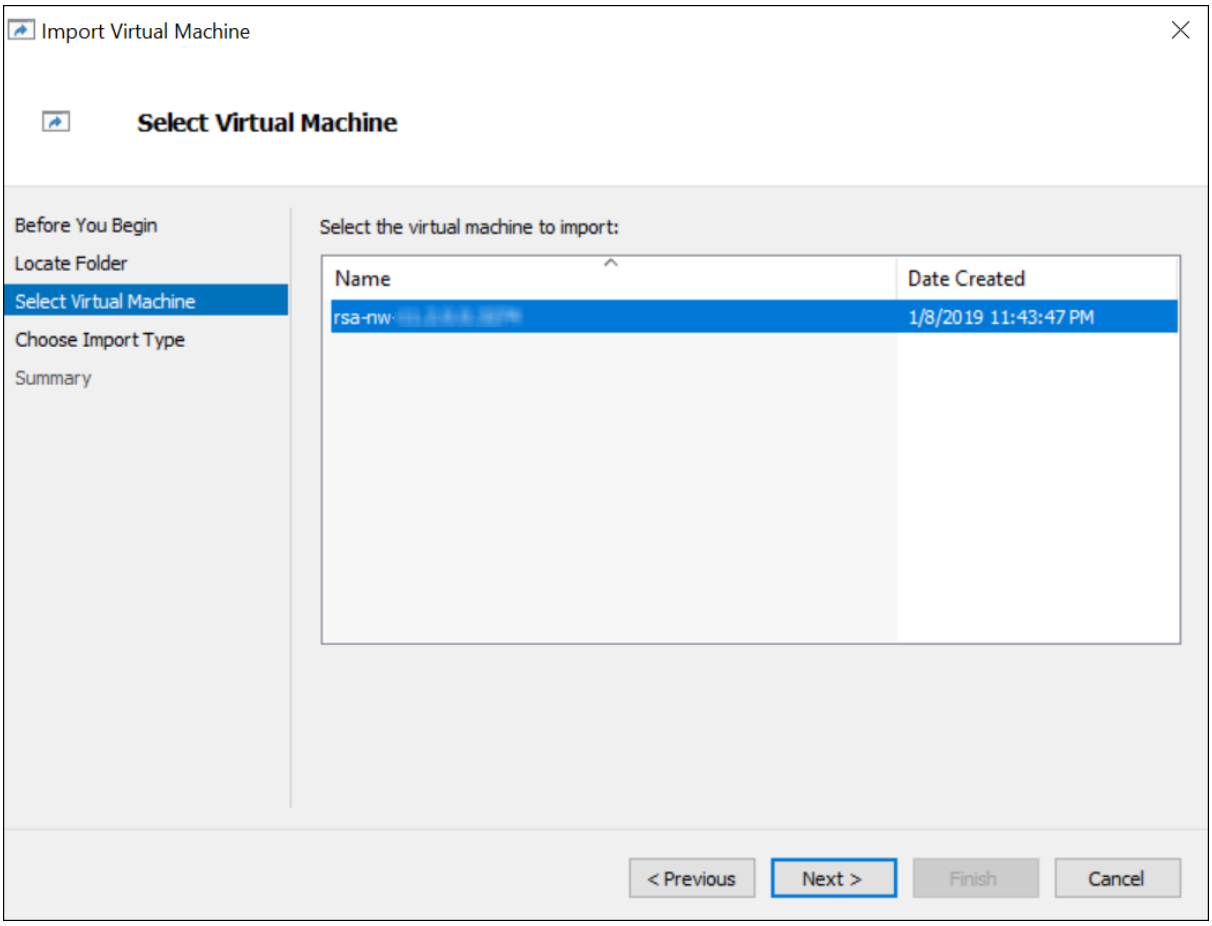
1. Log on to Hyper-V Manager.
2. Click **Import Virtual Machine** and Click **Next**.



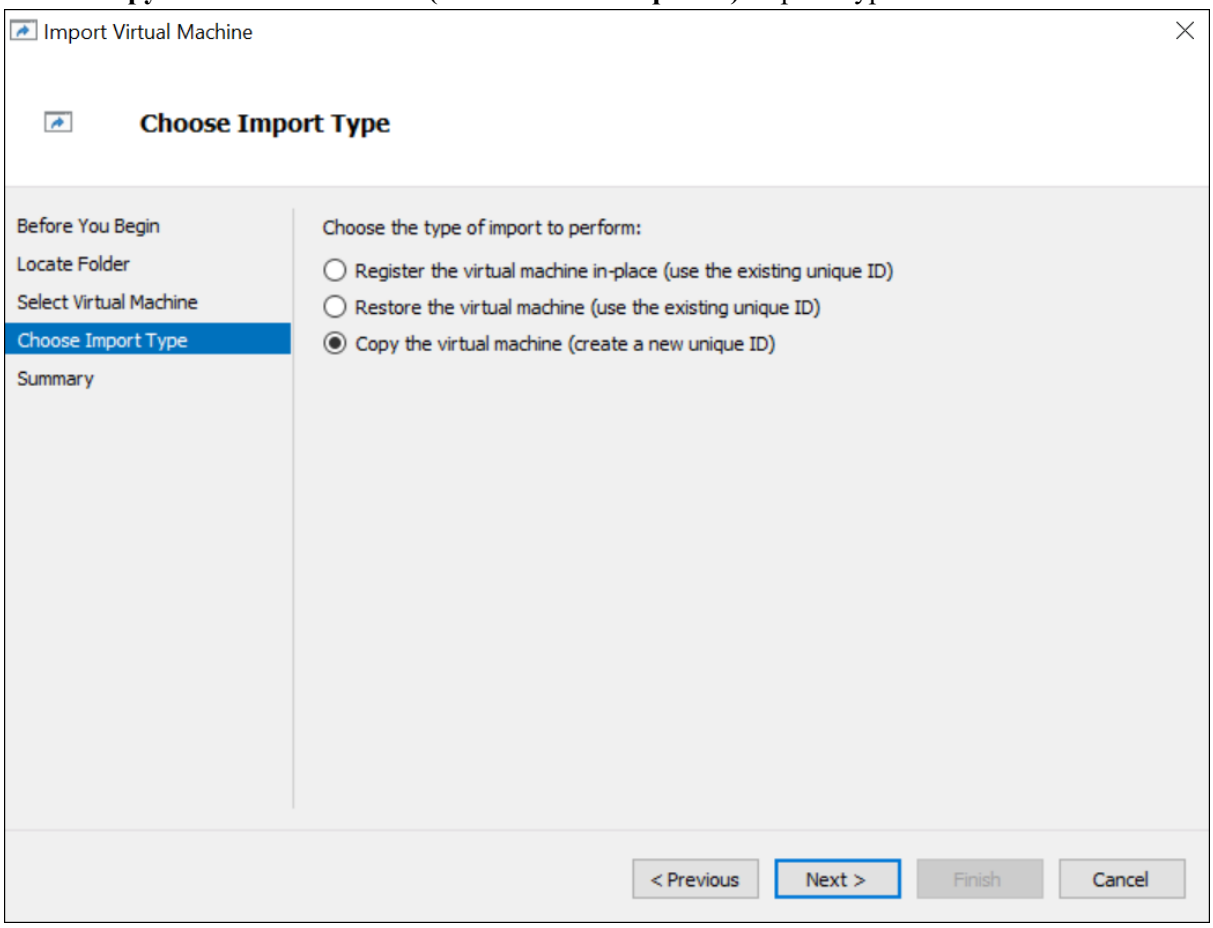
3. In the **Import Virtual Machine** dialog, specify the path where the zip file is extracted and Click **Next**.



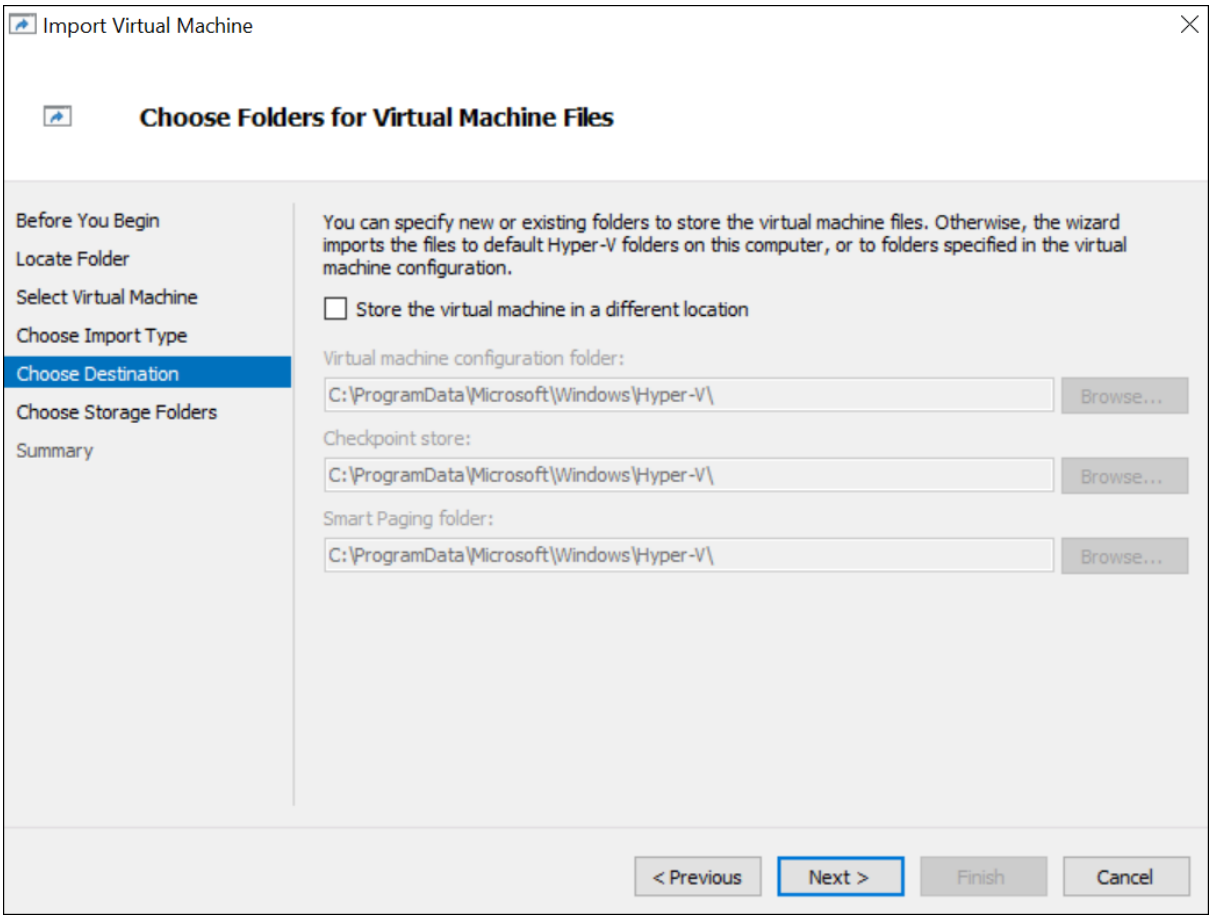
4. Select the Virtual Machine and Click **Next**.



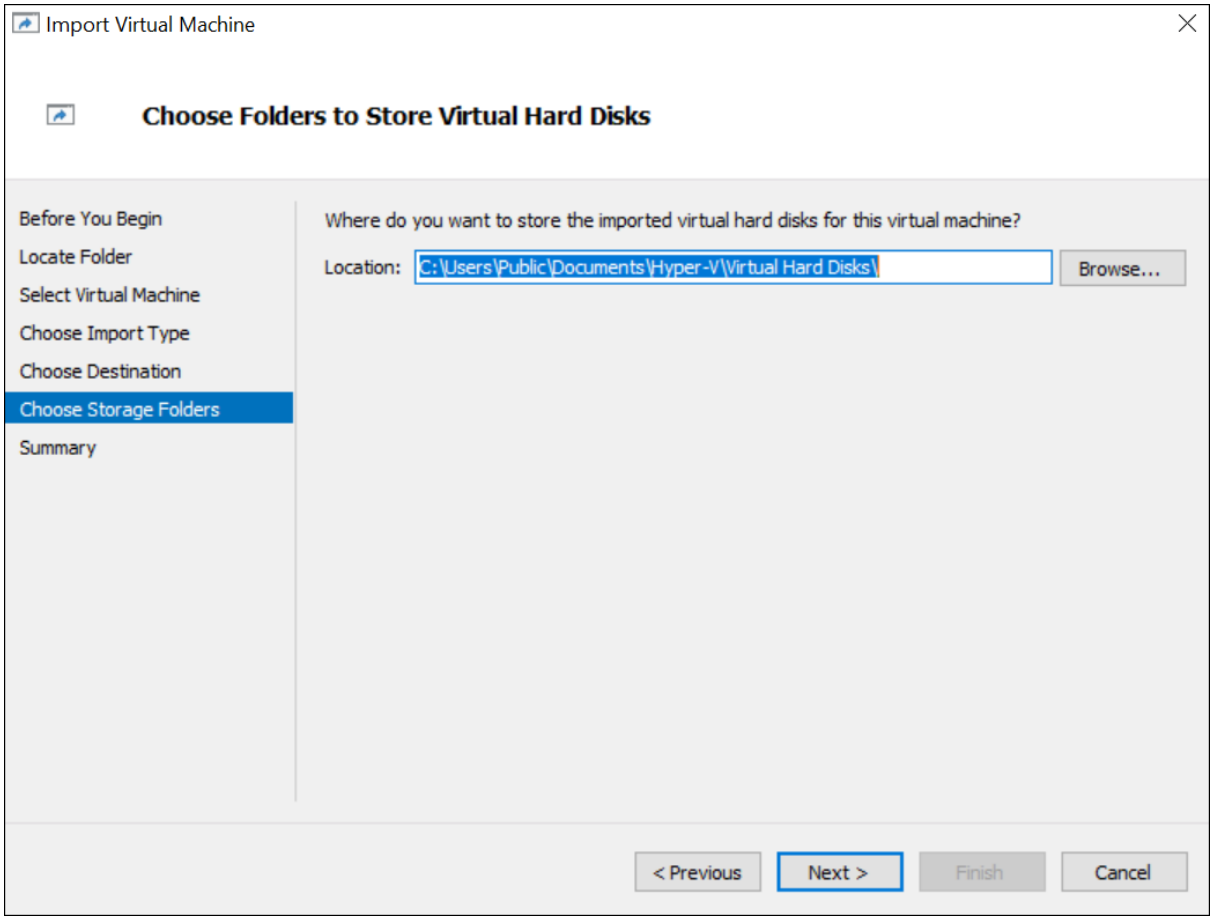
5. Choose **copy the Virtual machine (create a new unique ID)** Import Type.



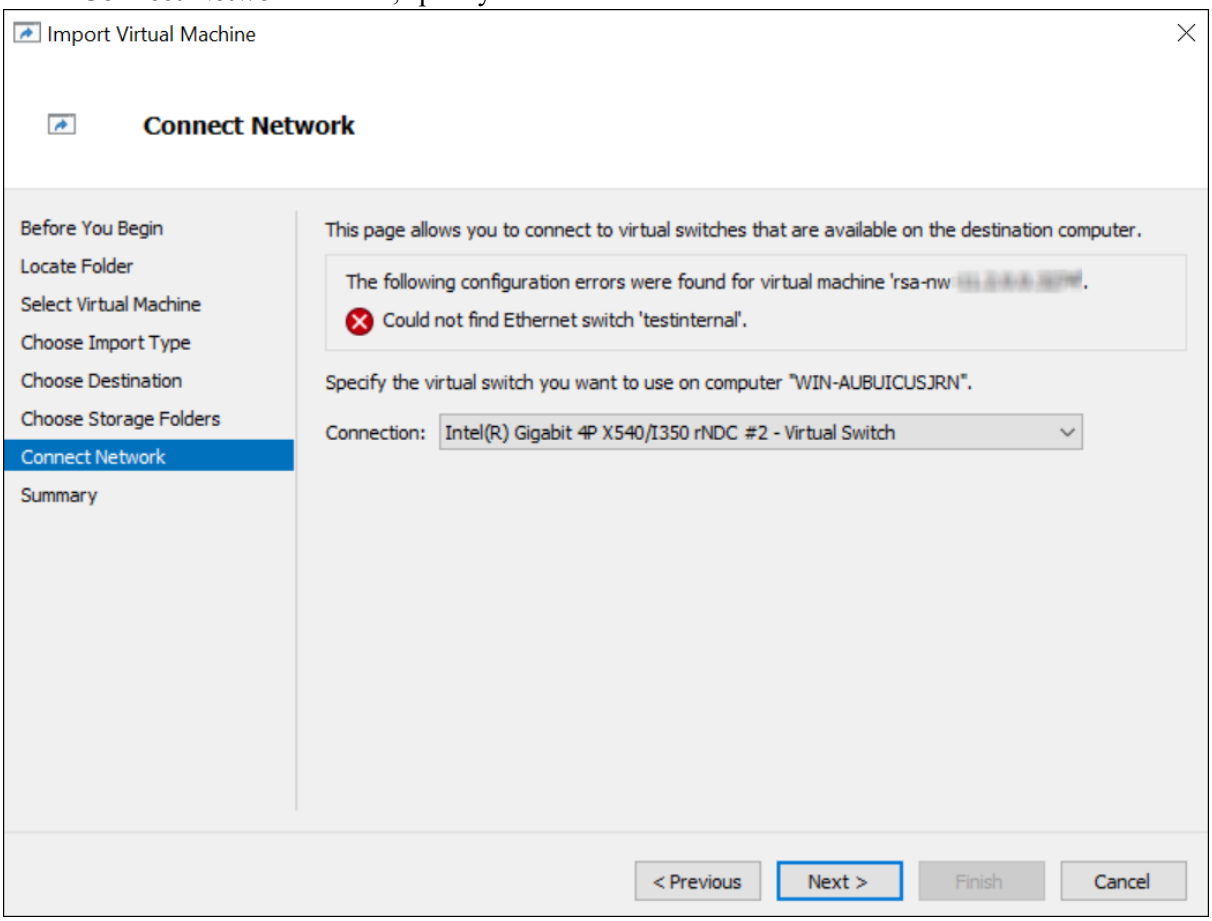
6. In the **Choose Destination** section, specify the new or existing folder to store the Virtual Machine files.



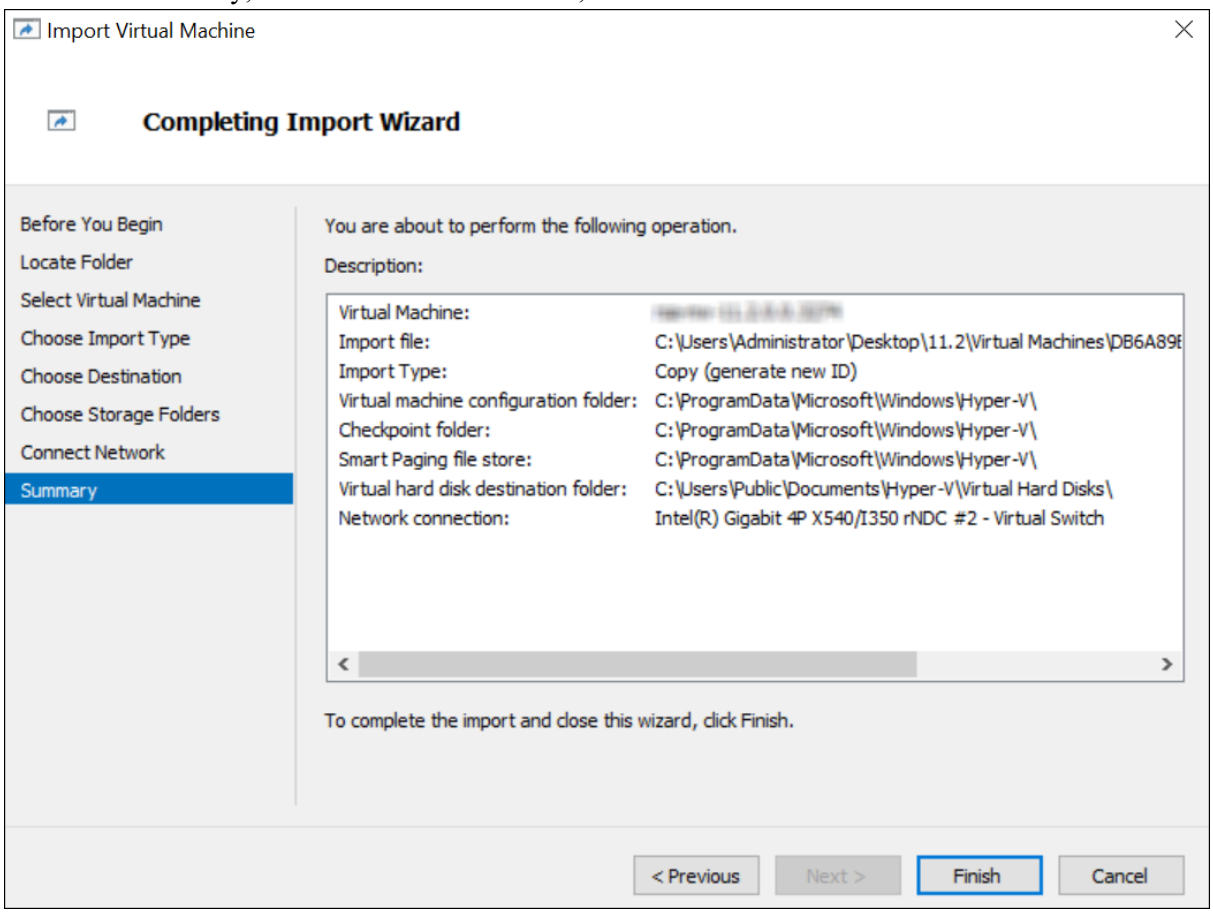
7. In the **Choose Storage Folder** section, specify the location where you want to store multiple Virtual Machine deployments.



8. In the **Connect Network** section, specify the Network name for the Virtual Machine to connect.



9. Check the Summary, if all the details are correct, click **Finish**.



Step 1c. Create Virtual Machine - Nutanix AHV

Nutanix AHV is an enterprise-ready hypervisor, offering integrated virtualization, app mobility, management, operational insights, and security. Complete the following steps according to their numbered sequence to deploy virtual host in Nutanix Acropolis Hypervisor (AHV).

Prerequisites

Make sure that you have:

- Network IP addresses, netmask, and gateway IP addresses for the virtual host.
- Network names for all virtual hosts, if you are creating a cluster.
- DNS or host information.
- Password for virtual host access. The default username is `root` and the default password is `netwitness`.

- The NetWitness 11.5 or later ISO file which is available for download from Downloads (https://community.netwitness.com/t5/netwitness-platform-downloads/tkb-p/netwitness-downloads).
- Administrative user account on Nutanix Prism.

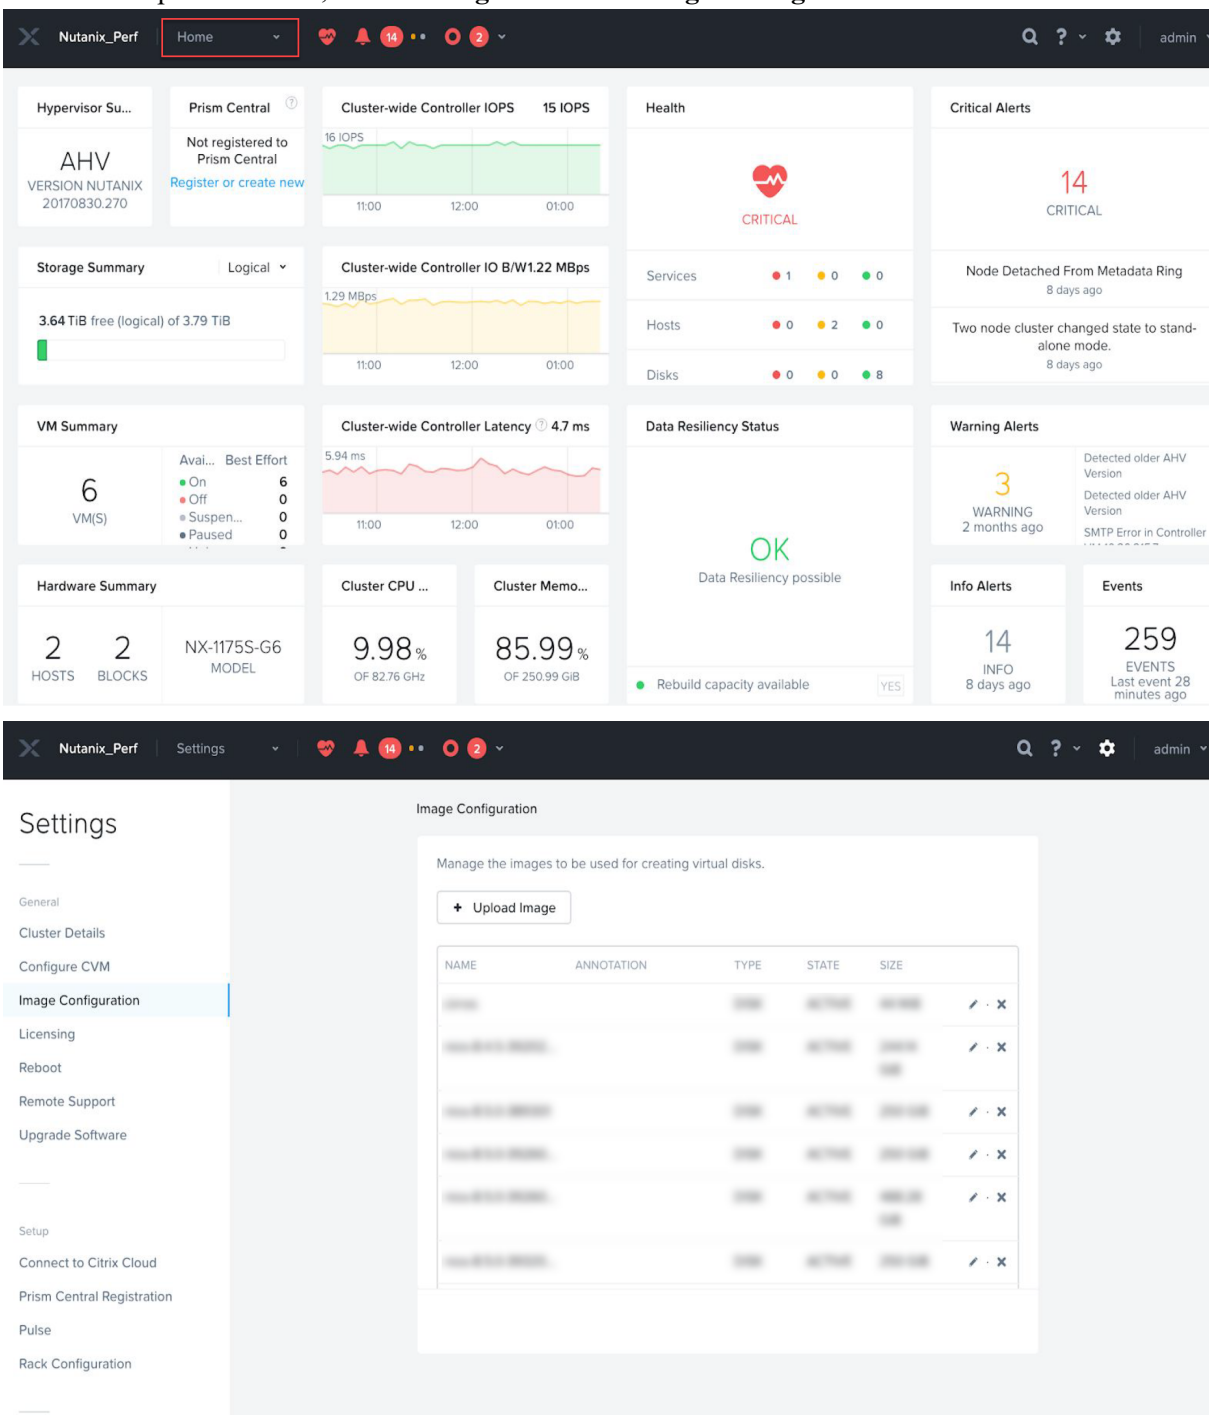
Procedure

Note: The following instructions illustrate an example of deploying a VM in the Nutanix AHV environment. The screens you see may be different from this example.

Create Image on Nutanix AHV

1. Log in to the Nutanix Prism GUI.

2. From the drop-down menu, click **Settings** and select **Image Configuration**.

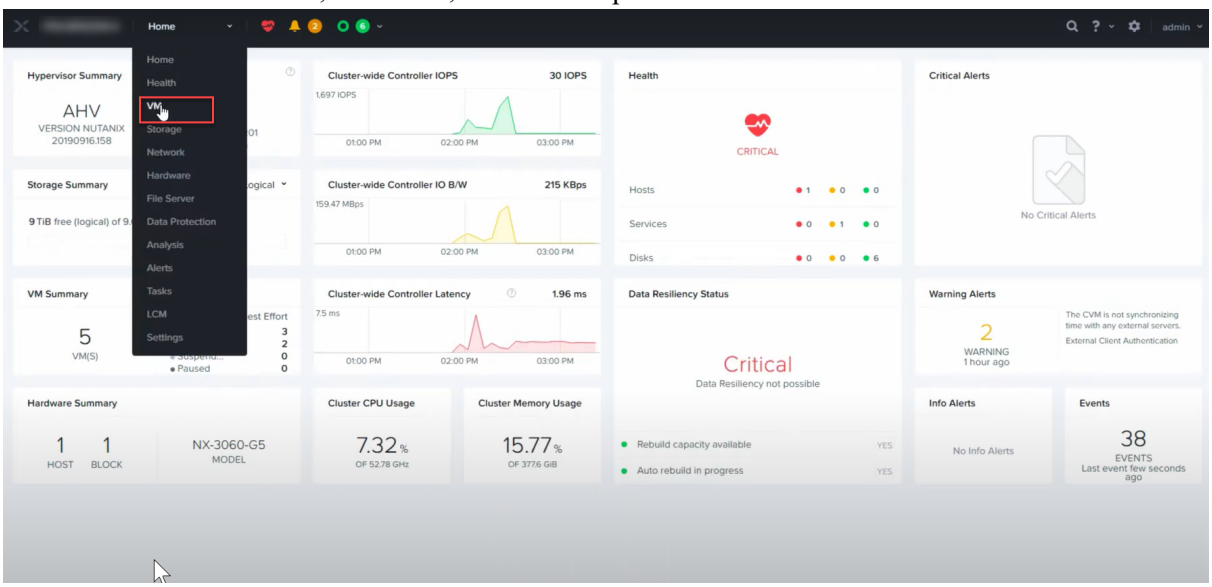


3. Click **Upload Image**.
4. In the **Create Image** dialog, select **Image Type** as **ISO**. Specify other details based on your requirements.
5. Select **Upload a file** and click **Choose File**.
6. Browse to the stored location of the NetWitness 11.5 or later ISO file and select it.

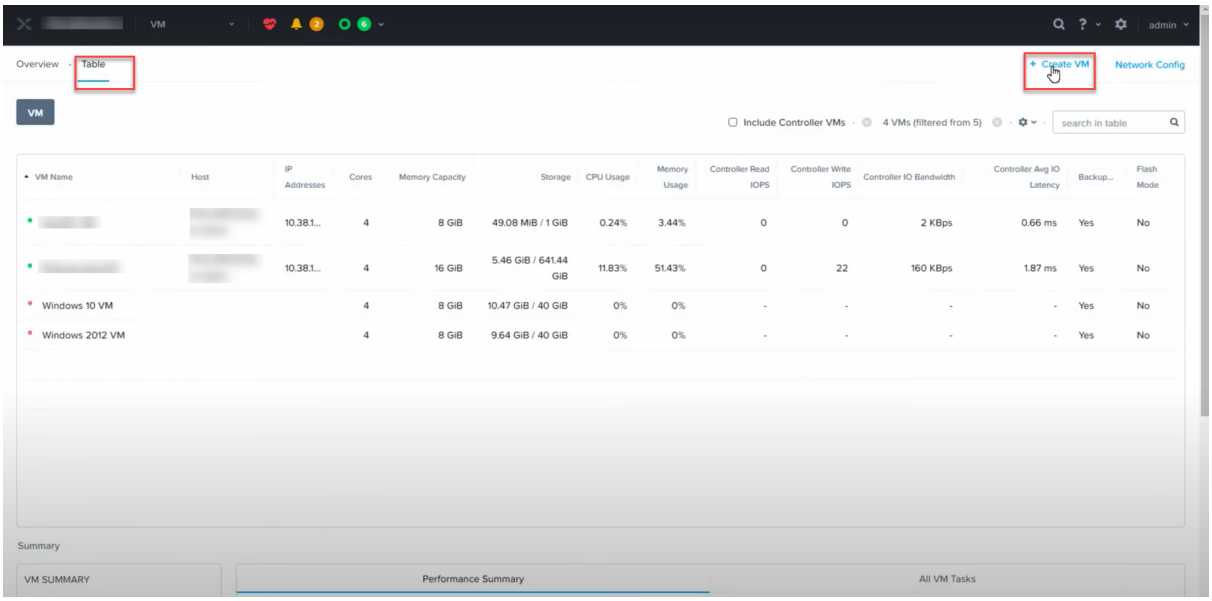
7. Click **Save**.

Create VM on Nutanix AHV

1. In the Nutanix Primis GUI, click **VM**, from the drop-down menu.

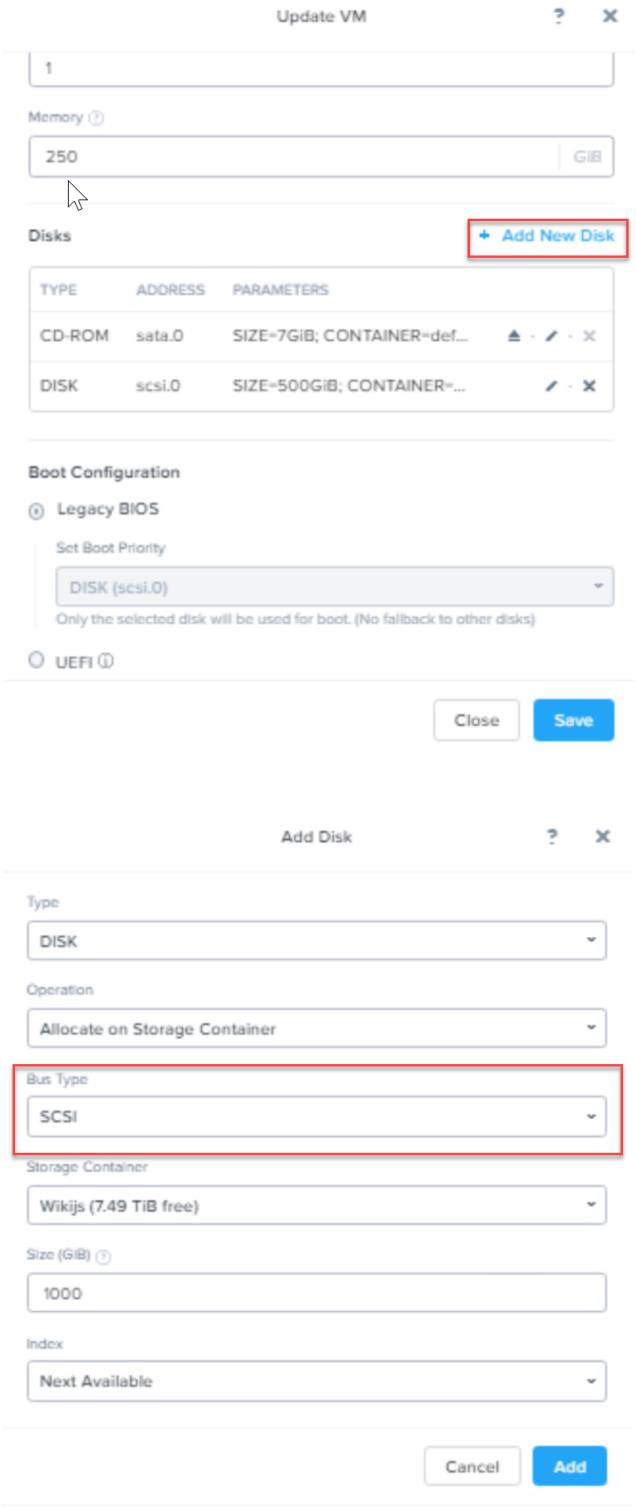


2. Go to the **Table** view and click **Create VM**.

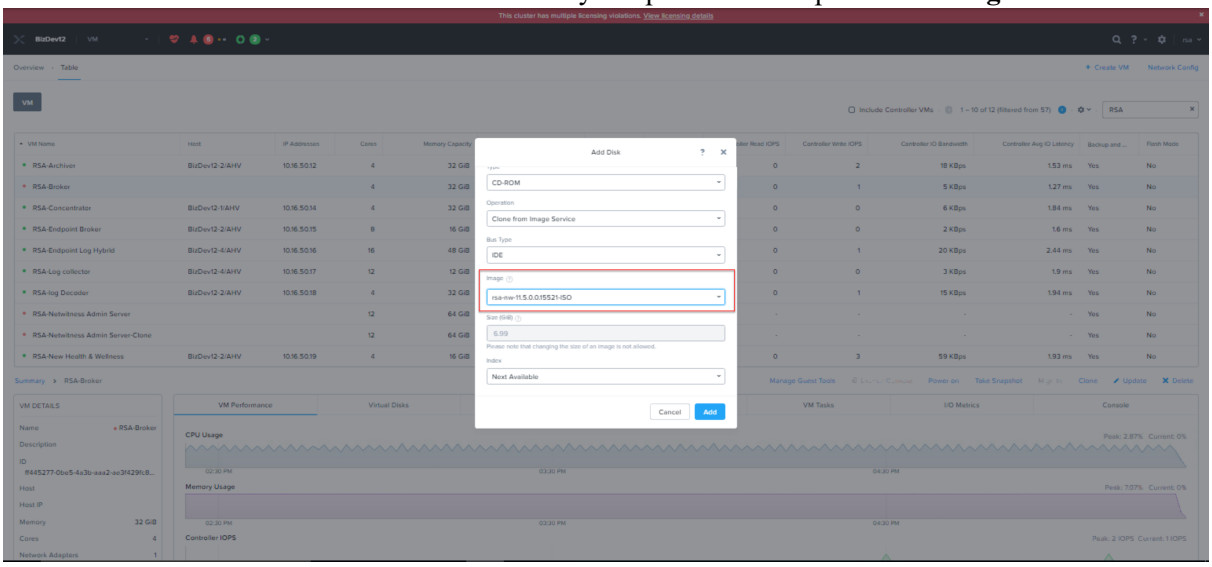


3. In the **Create VM** dialog, enter the required details such as name, description, time zone, vCPU, cores per vCPU, etc. For more information, see Creating a VM (AHV) in the Prism Web Console Guide.

4. In the **Create VM** dialog, scroll down to the **Disks** section and click **Add New Disk**.



5. In the **Add Disk** dialog, do the following:
   a. Select **DISK** in the **Type** field and specify other details based on your requirement.
   b. Select **Allocate on Storage Container** in the **Operations** field.
   c. Select **SCSI** as **Bus Type**.
6. In the **Add Disk** dialog, select **CD-ROM** as the **Type**.
7. Select **Clone from Image Service** in the **Operations** field and **IDE** as **Bus Type**.
8. Select the NetWitness 11.5 or later ISO file that you uploaded in Step 6 in the **Image** field.



9. Select **Clone from Image Service** in the **Operations** field and **IDE** as **Bus Type**.
10. Click **Add** and enter other details such as adding network adapters, GPU, etc.
11. Click **Save**. You can now see your VM in the VM list.

IMPORTANT: Ensure that you eject the CD-ROM after installing the NetWitness Platform. Otherwise, every time you reboot the VM it will boot from CD-ROM.

Step 1d. Create Virtual Machine in Esxi

Complete the following steps according to their numbered sequence to create the virtual machine in Esxi environment.

Prerequisites

Make sure that you have:

- Network IP addresses, netmask, and gateway IP addresses for the virtual host.
- Network names for all virtual hosts, if you are creating a cluster.

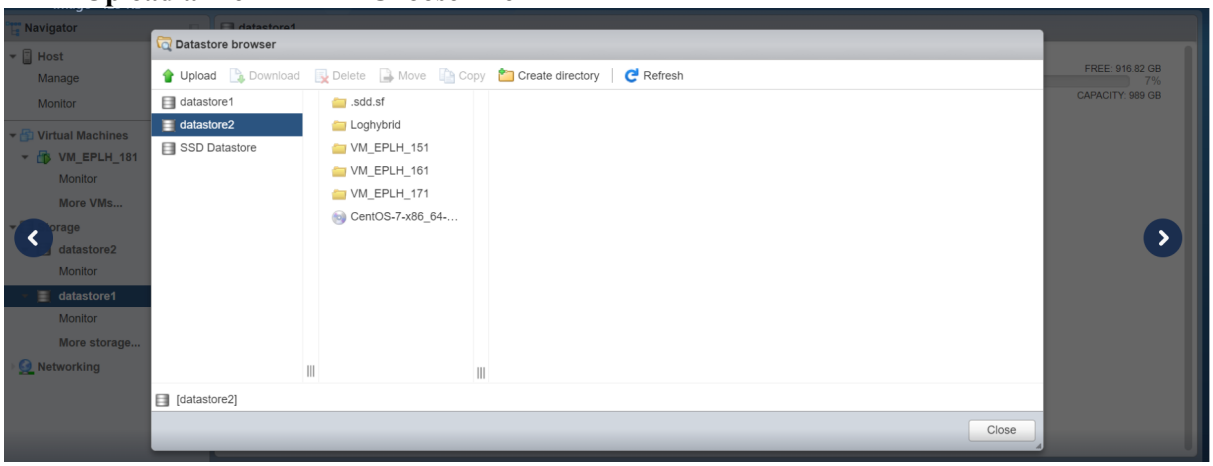
- DNS or host information.
- Password for virtual host access. The default username is `root` and the default password is `netwitness`.
- The NetWitness virtual host package file for example, `rsa-nw-11.7.0.0.3274.zip`. (You download this package from Downloads https://community.netwitness.com/t5/netwitness-platform-downloads/tkb-p/netwitness-downloads)

Procedure

Note: The following instructions illustrate an example of creating a VM in Esxi environment. The screens you see may be different from this example.

Create Image on Esxi

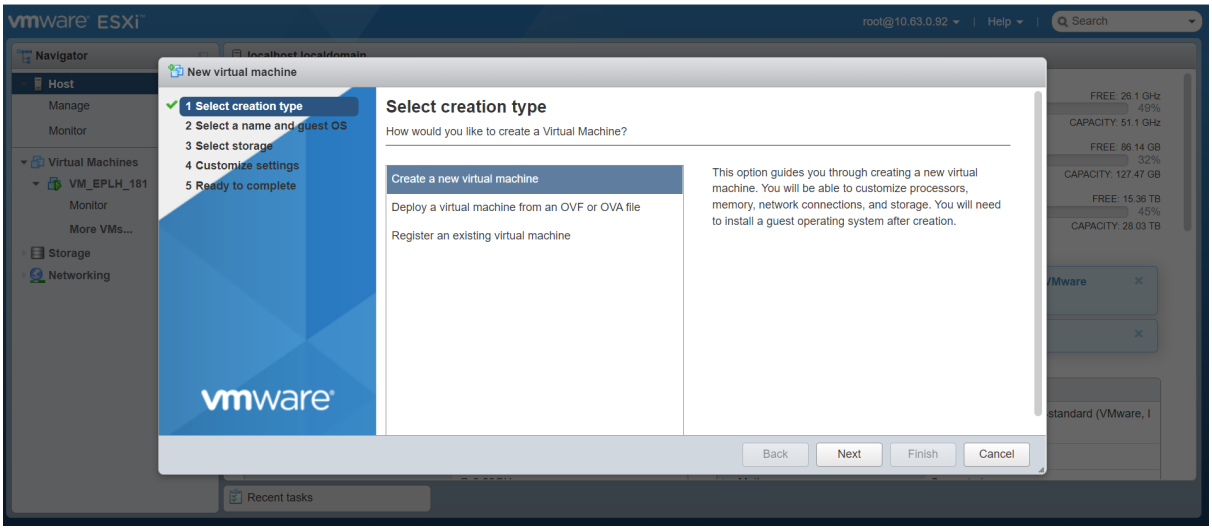
1. Log in to the Esxi GUI.
2. From the drop-down menu, click **Settings** and select **Image Configuration**.
3. Click **Upload Image**.
4. In the **Create Image** dialog, select **Image Type** as **ISO**. Specify other details based on your requirements.
5. Select **Upload a file** and click **Choose File**.



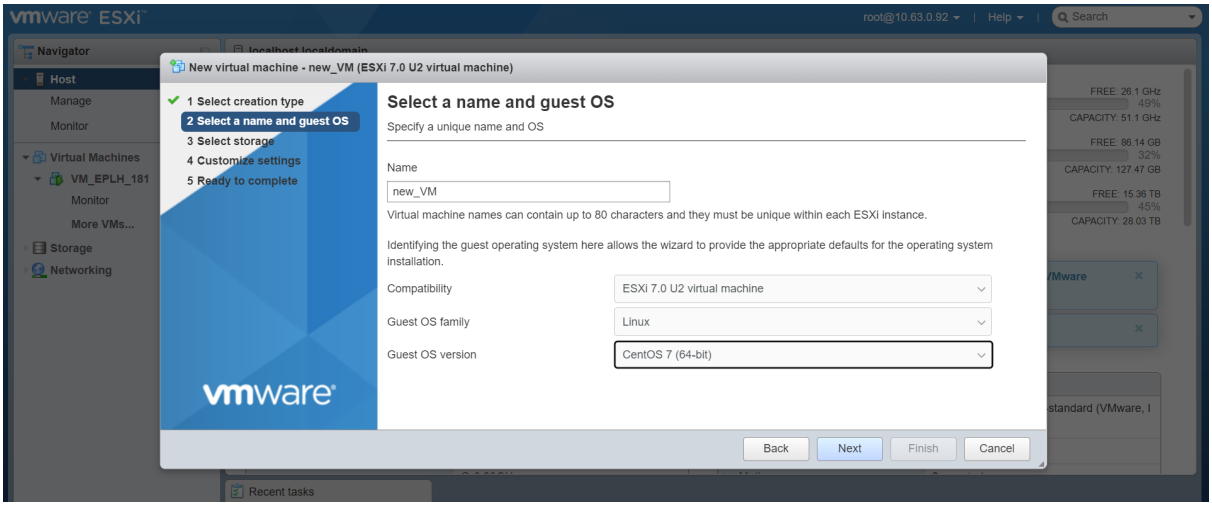
6. Browse to the stored location of the NetWitness 11.5 or later ISO file and select it.
7. Click **Save**.

Create VM on Esxi

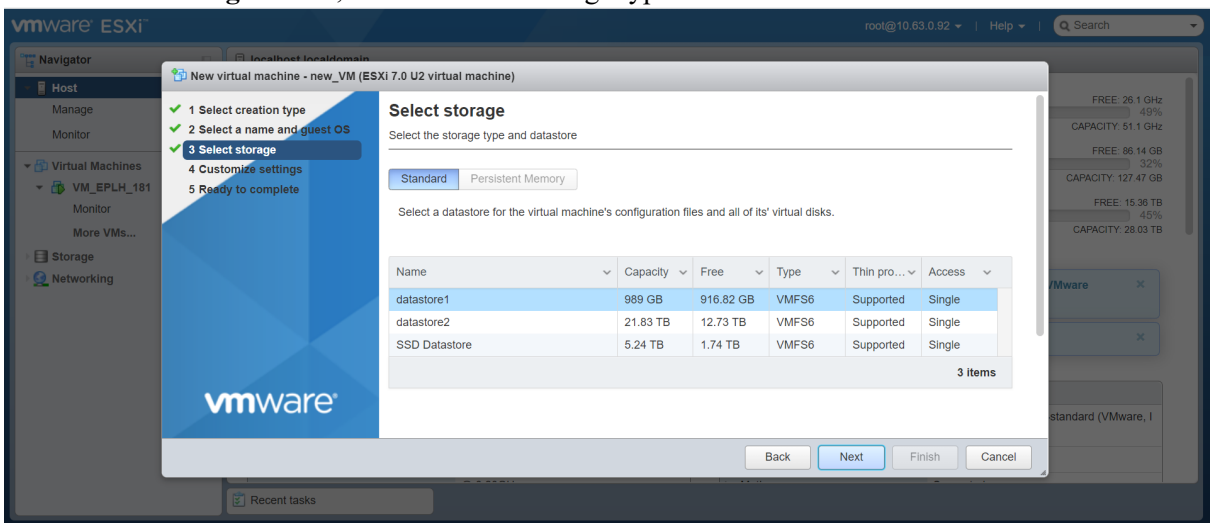
1. In the Esxi GUI, click **VM**, from the drop-down menu.
2. Go to the **Host** screen and click **Create a new virtual machine**



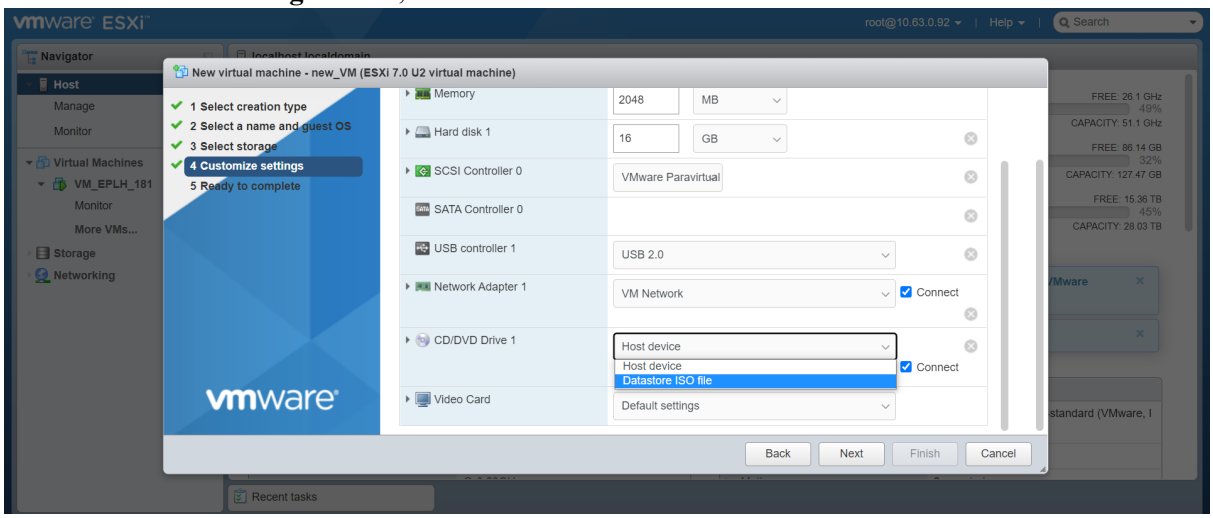
3. In the **Select a name and guest OS** screen, enter the required details such as name, compatibility, guest OS family, guest OS version and click **NEXT**.



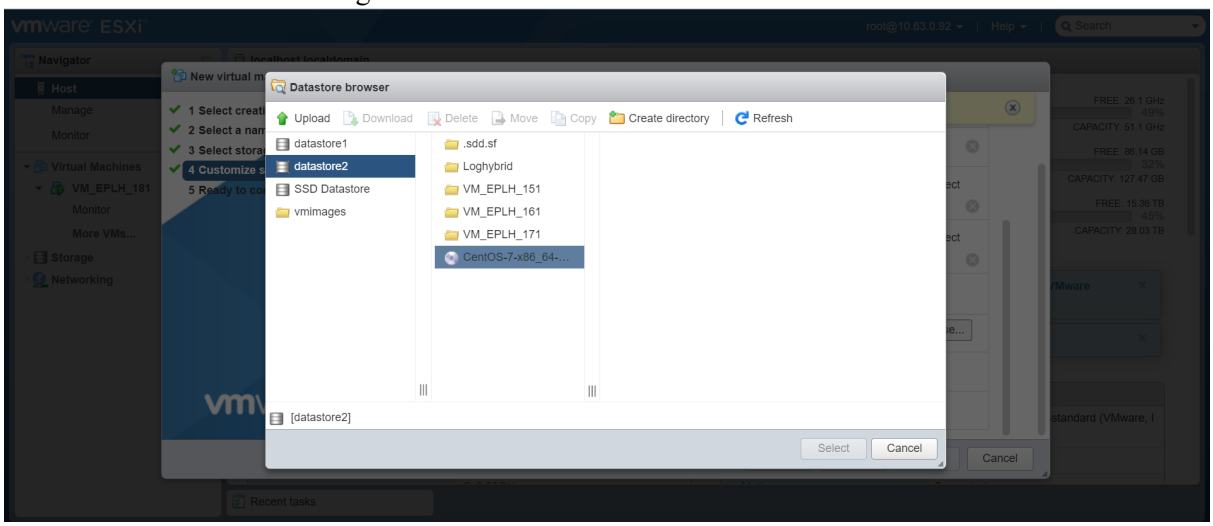
4. In the **Select storage** screen, select the data storage type and click **Next**.



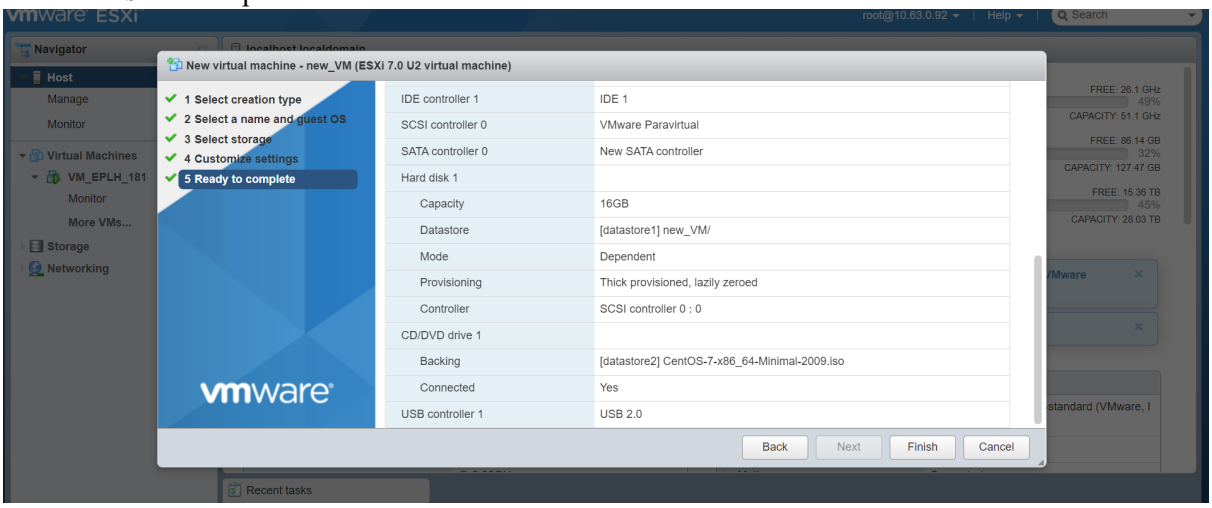
5. In the **Customize settings** screen, select **CD/DVD Drive 1** as **Datastore ISO file**.



6. Browse the mounted ISO image and click **Next**.



7. Click **Finish** to complete the installation.



Step 2. Configure Block Storage to Accommodate NetWitness Platform

When you deploy databases from OVA or VHDX, the initial database space allocation is not adequate for production environment. You need to add or expand the datastores volume after installation.

Task 1. Add New Disk

Get familiar with the NetWitness® Platform Storage Guide to understand the types of drives and volumes needed to support NetWitness instances. For more information, see Storage Guide for NetWitness Platform 11.x

Add New Disk

Add New Disk in VMware ESXi

Add New Disk in Hyper-V

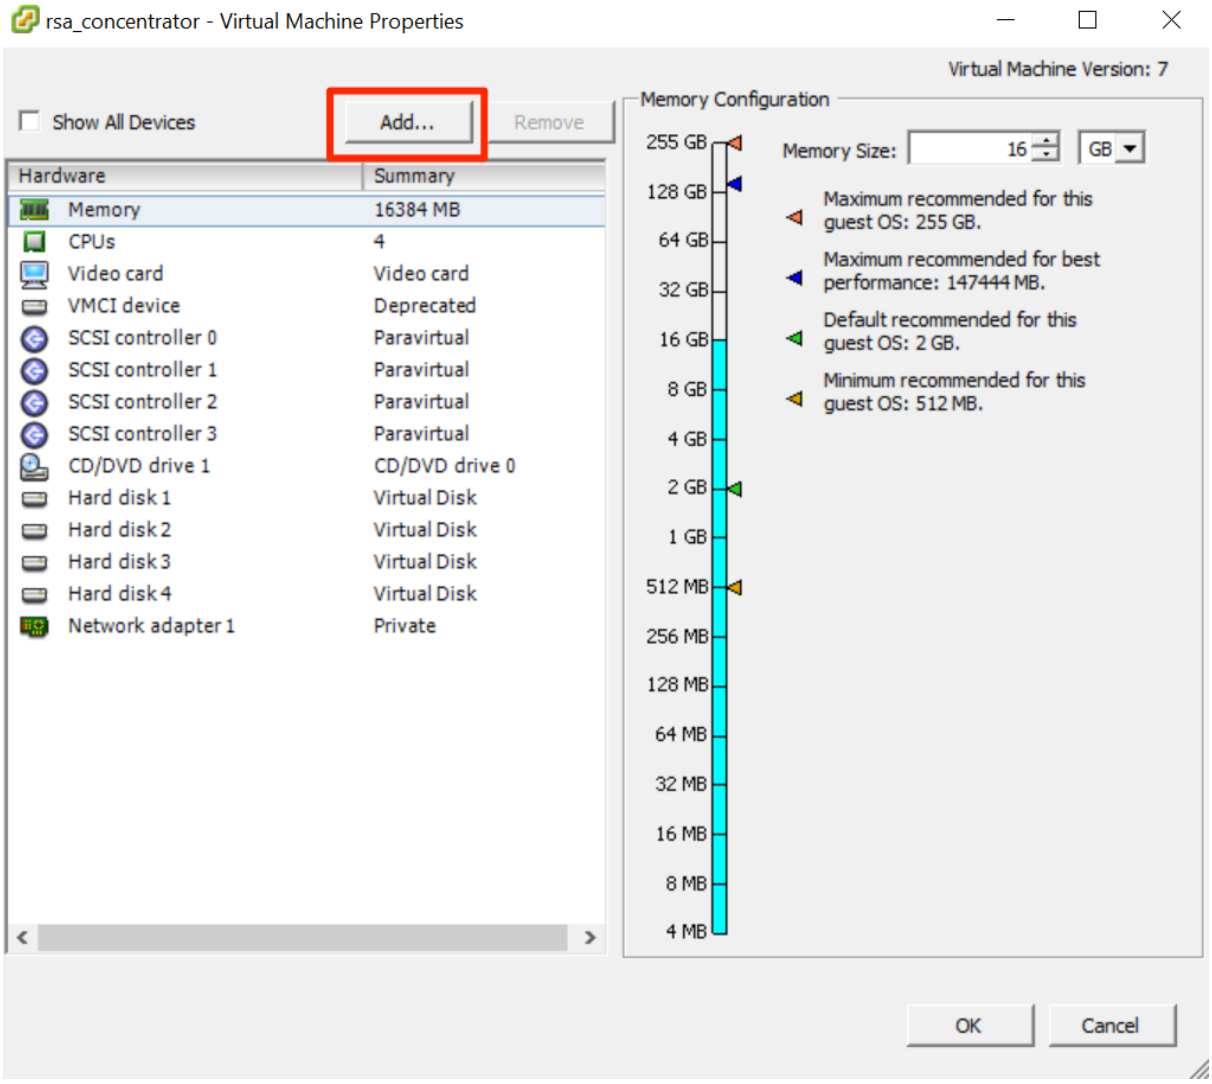
Add New Disk in Nutanix AHV

Add New Disk in VMware ESXi

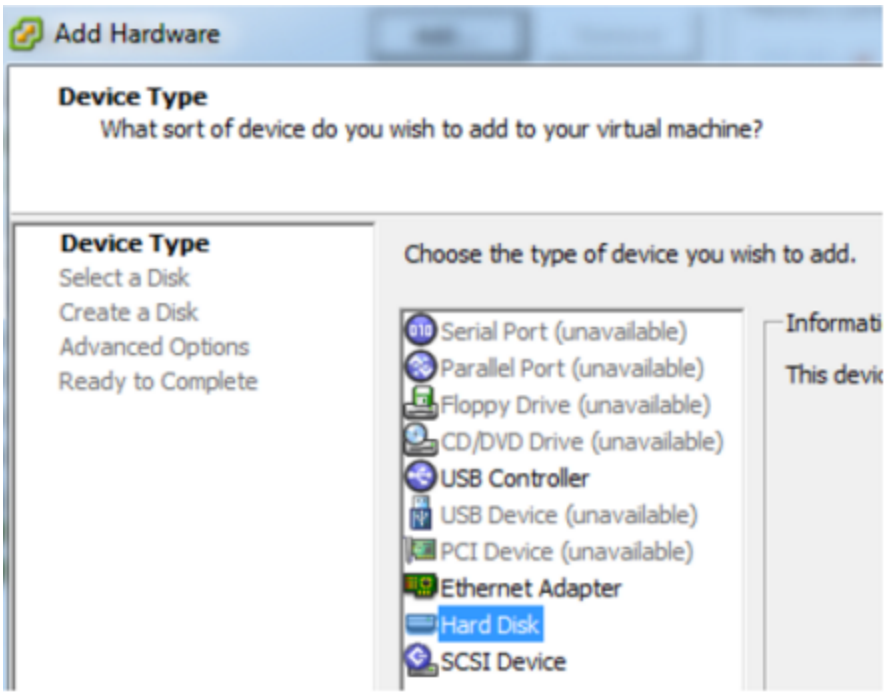
This procedure shows you how to add a new 100 GB disk on the same datastore.

Note: The procedure to add a disk on different datastore is similar to the procedure shown here.

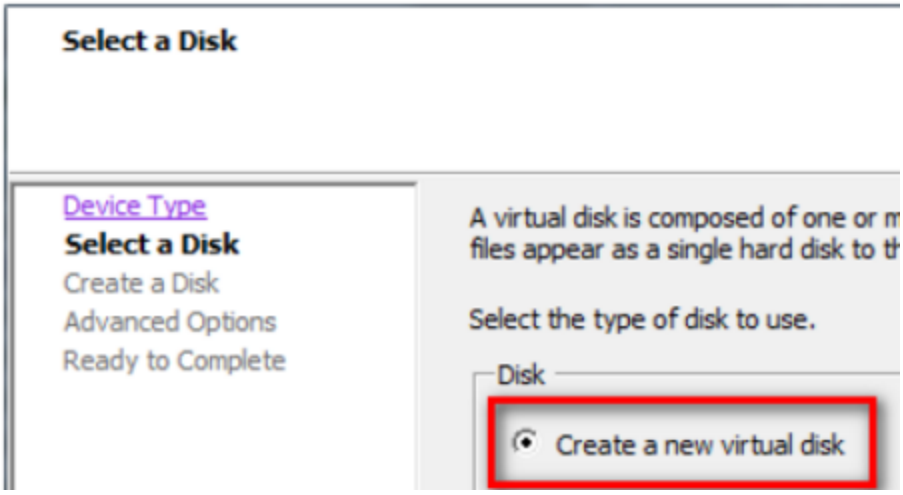
1. Shut down the machine, edit **Virtual Machine Properties**, click **Hardware** tab, and click **Add**.



2. Select **Hard Disk** as the device type.

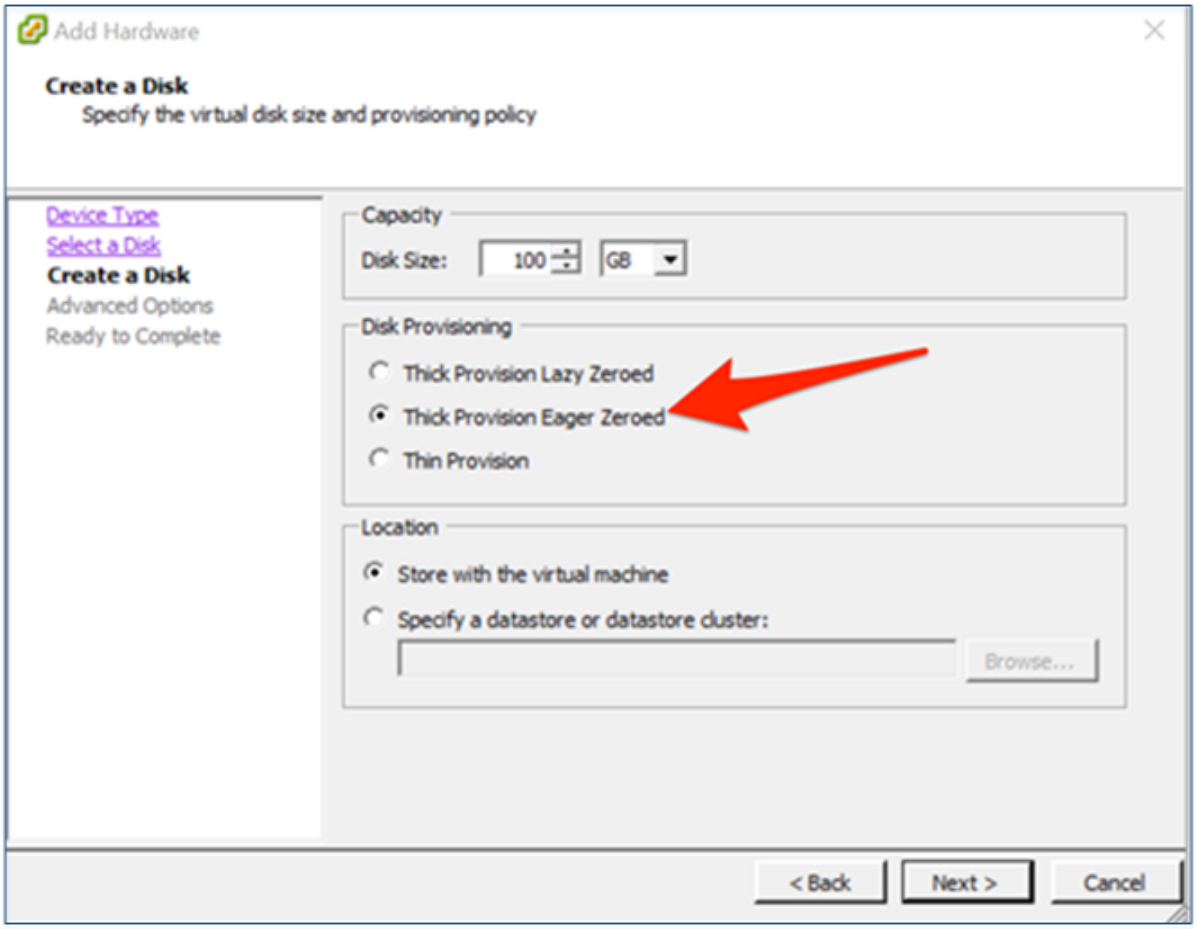


3. Select **Create a new virtual disk**.

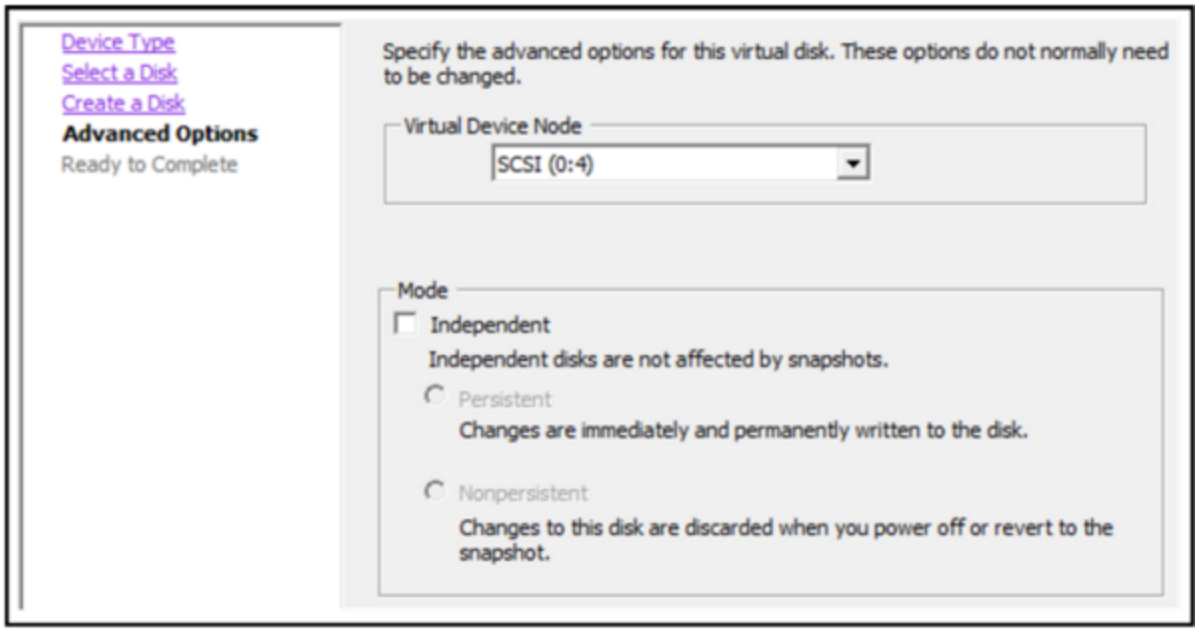


4. Choose the size of the new disk and where you want to create it (on the same datastore or a different datastore).

Note: Choose data provisioning based on your requirements

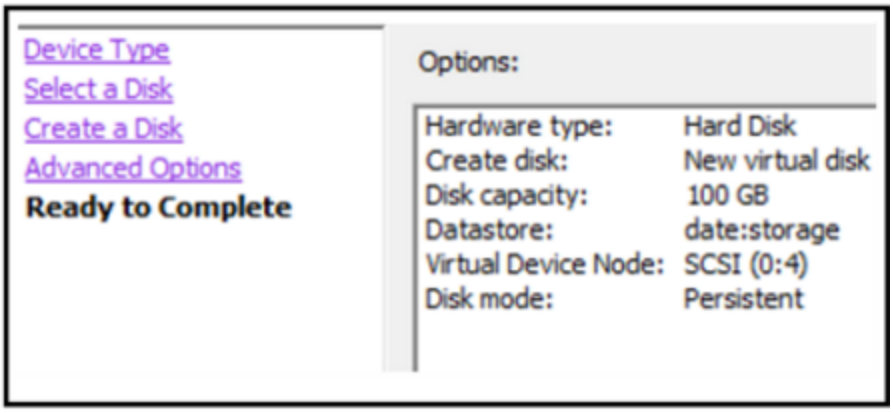


5. Approve the proposed Virtual Device Node.



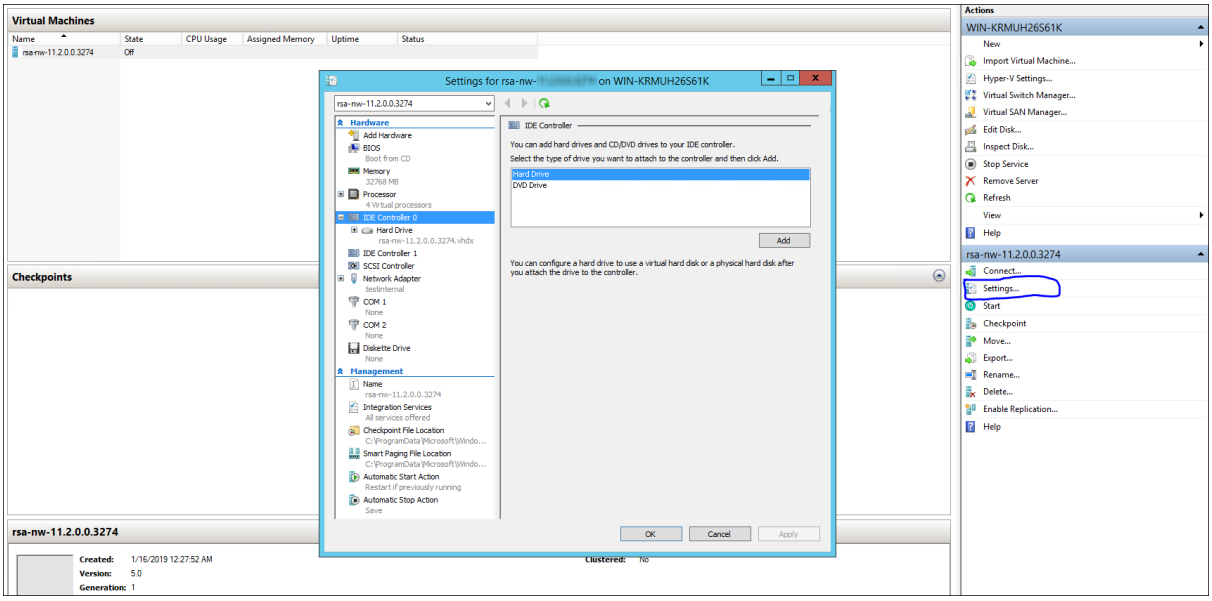
Note: The Virtual Device Node can vary, but it is pertinent to `/dev/sdX` mappings.

6. Confirm the settings.

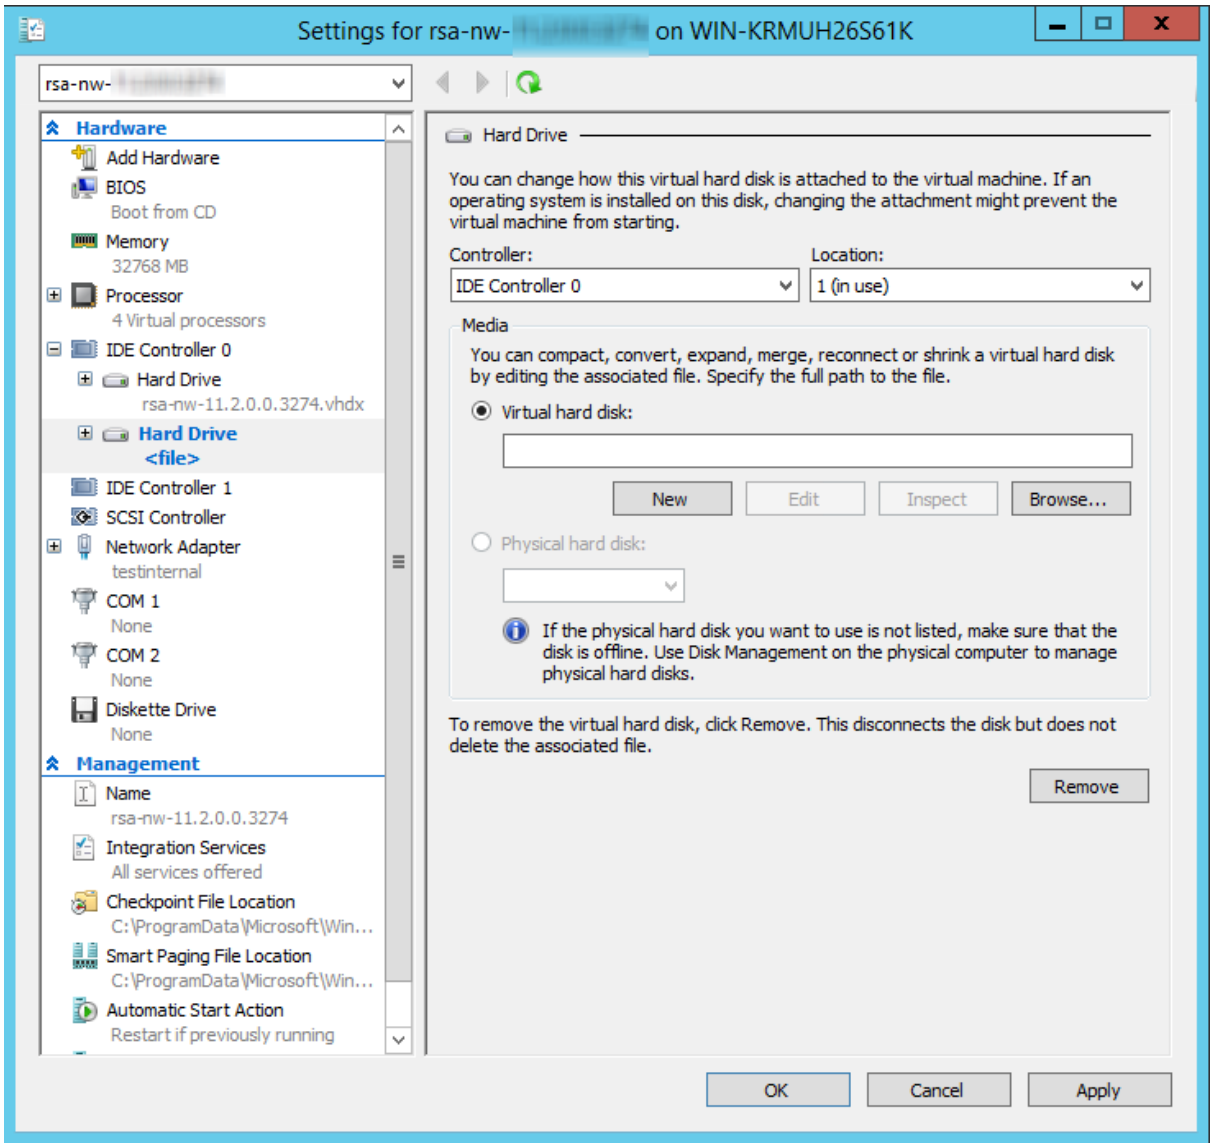


Add New Disk in Hyper-V

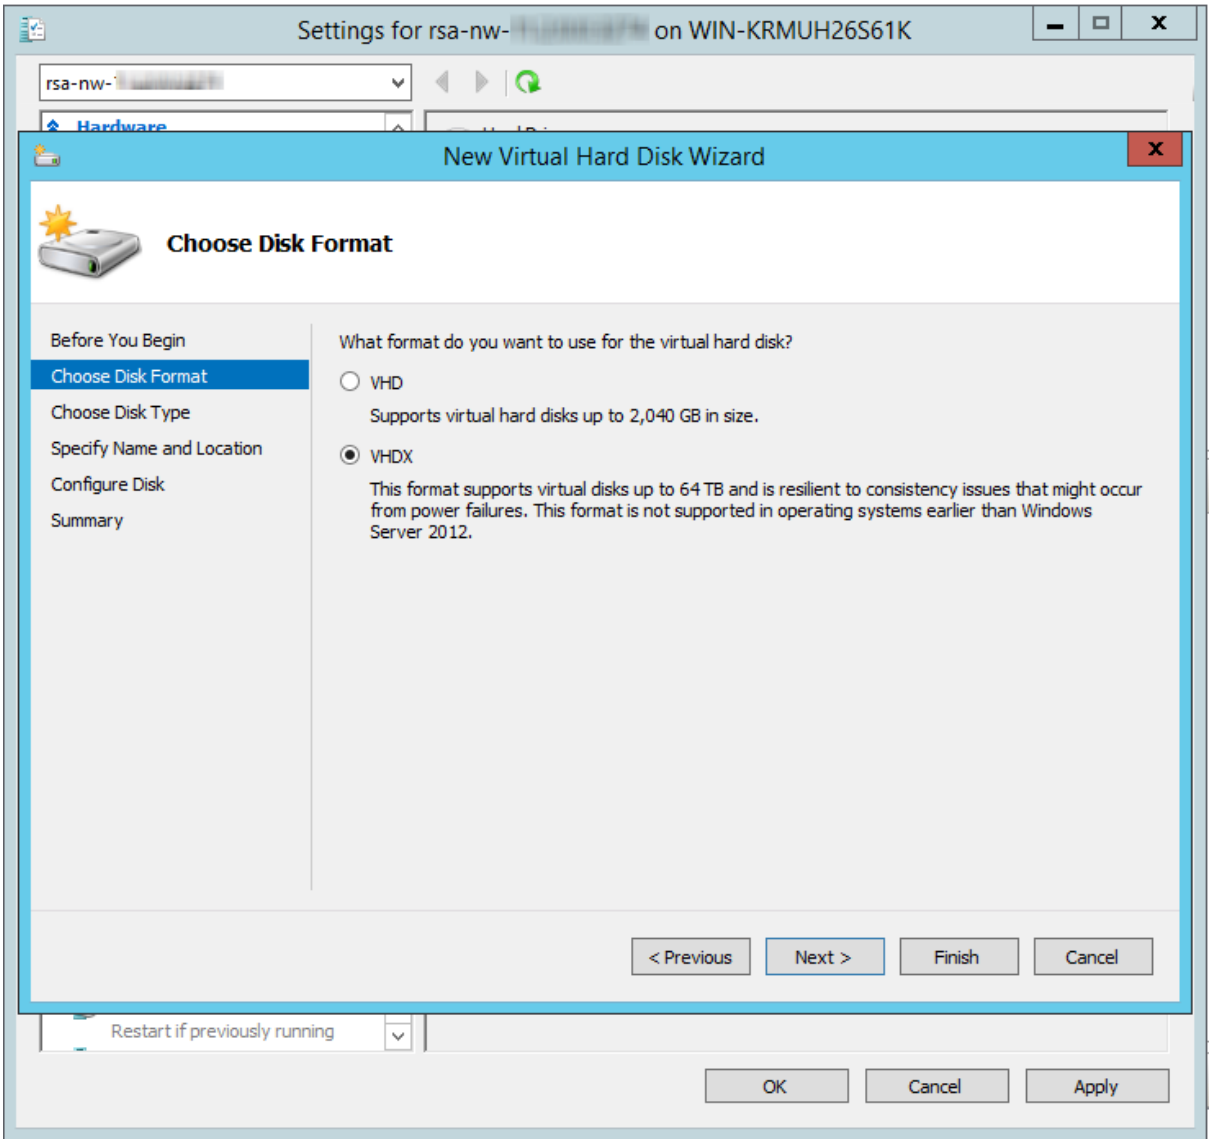
1. Shut down the VM and click **Settings and IDE Controller**, select the **Hard Drive** and click **Add**.



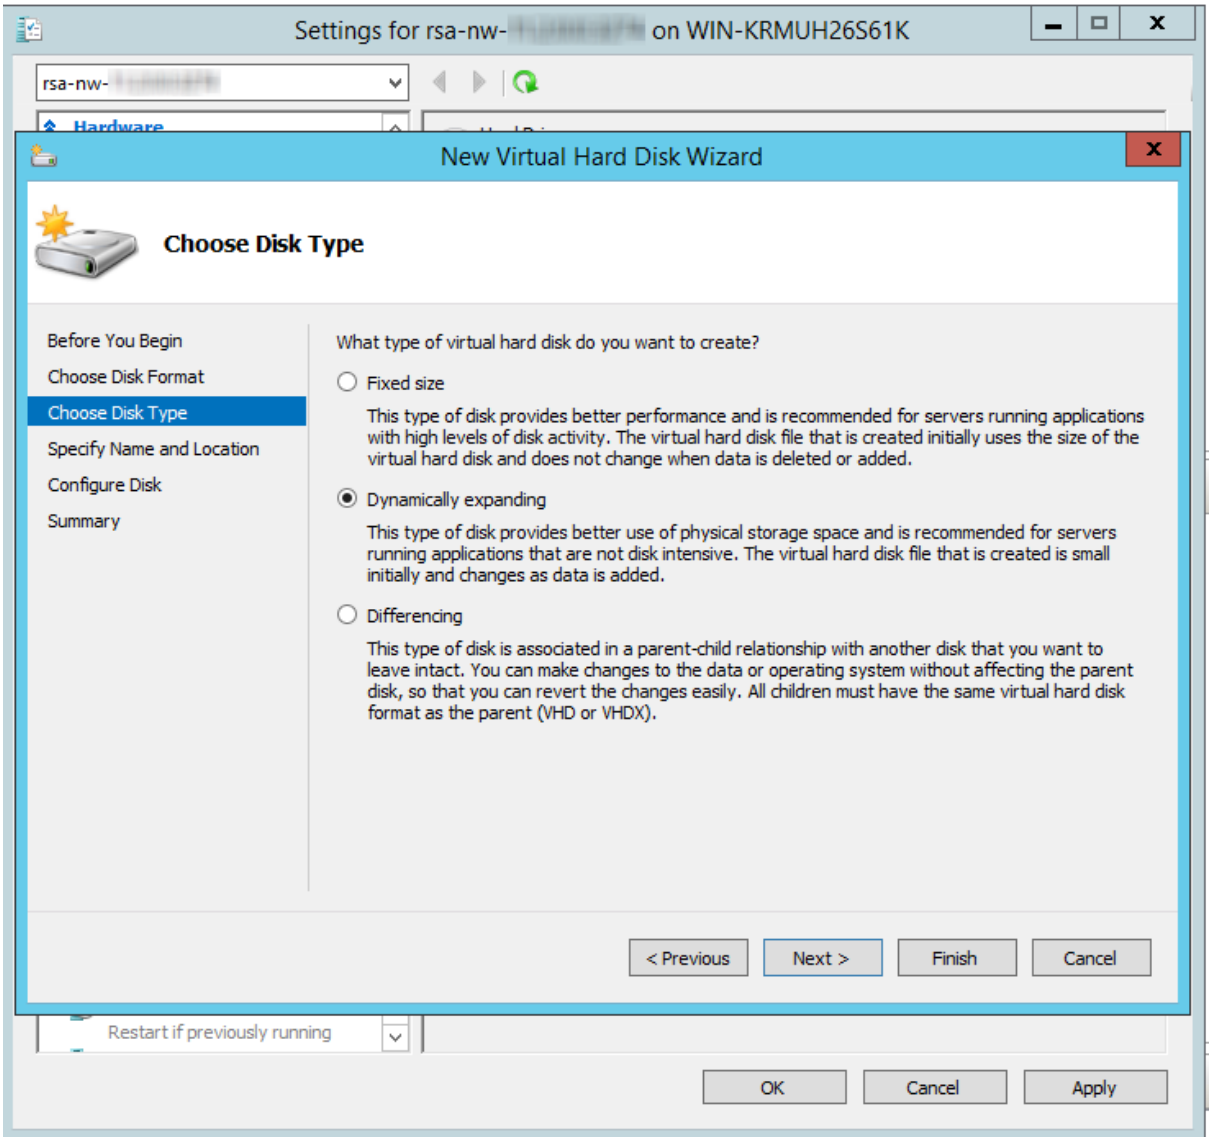
2. Select the New Virtual Hard disk.



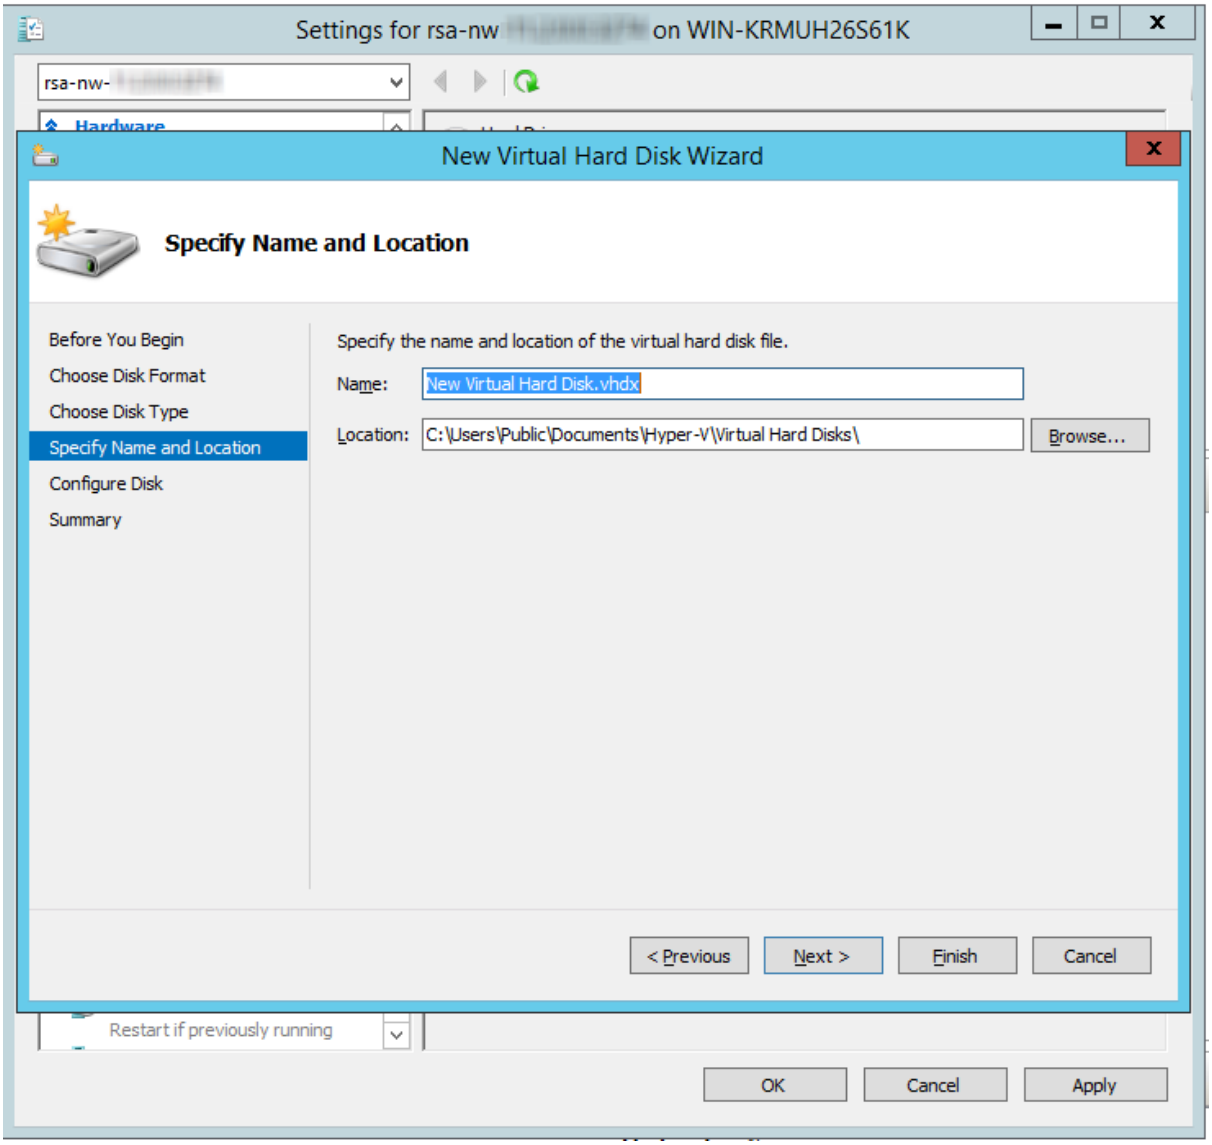
3. Select **VHDX** as a disk format.



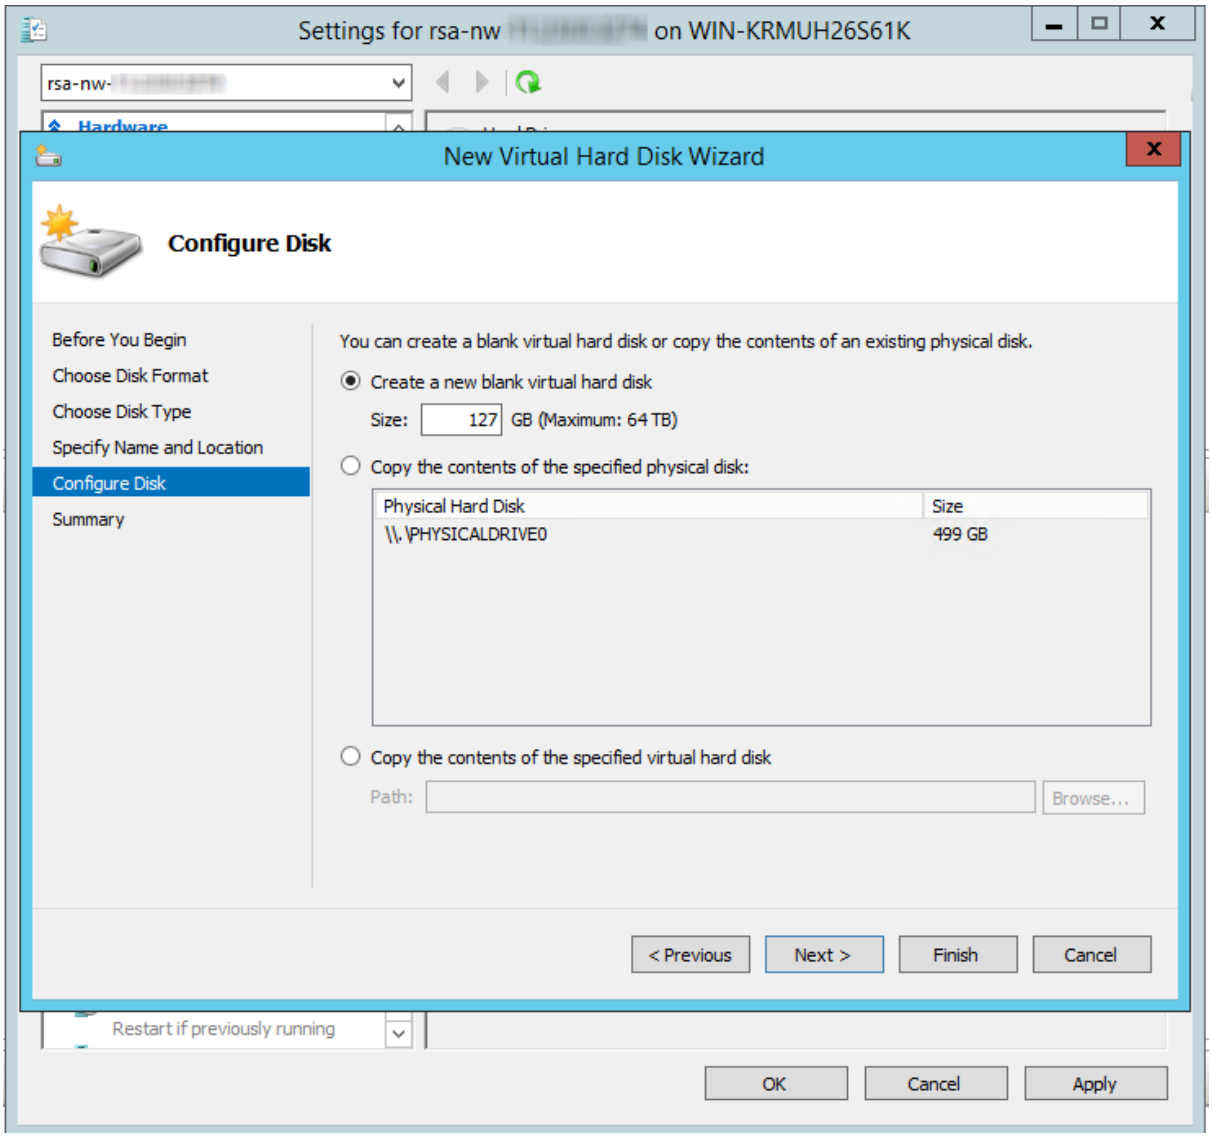
4. Select **Dynamically expanding** as a disk type.



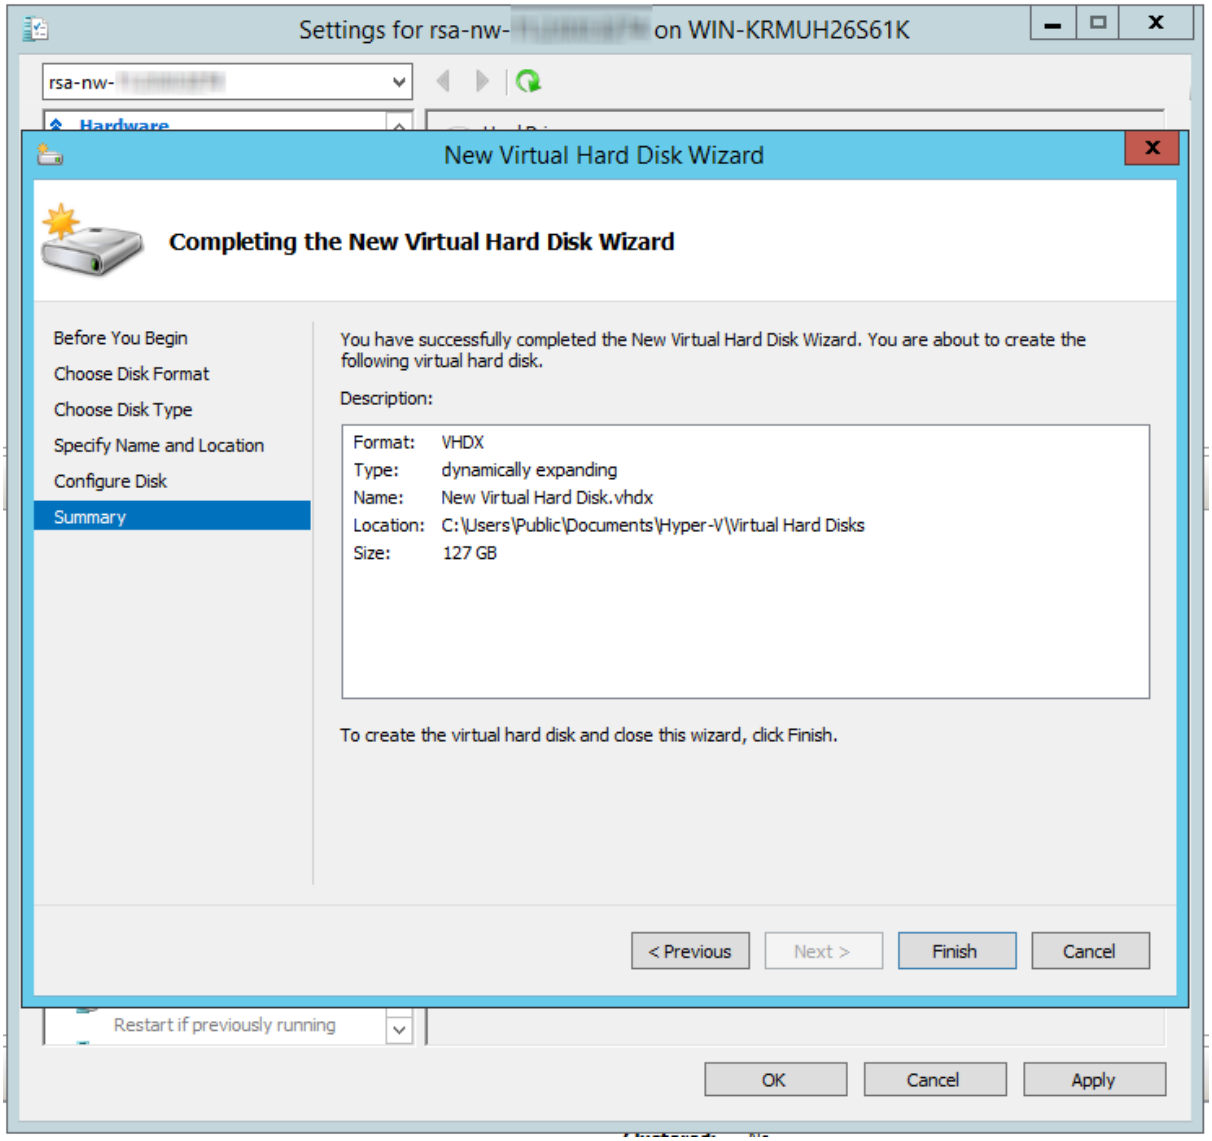
5. Specify the **Name** and **Location** of the virtual hard disk file.



6. Select **create a new blank virtual hard disk** and specify the size.



7. In the **Summary**, review the settings and click **Finish**.

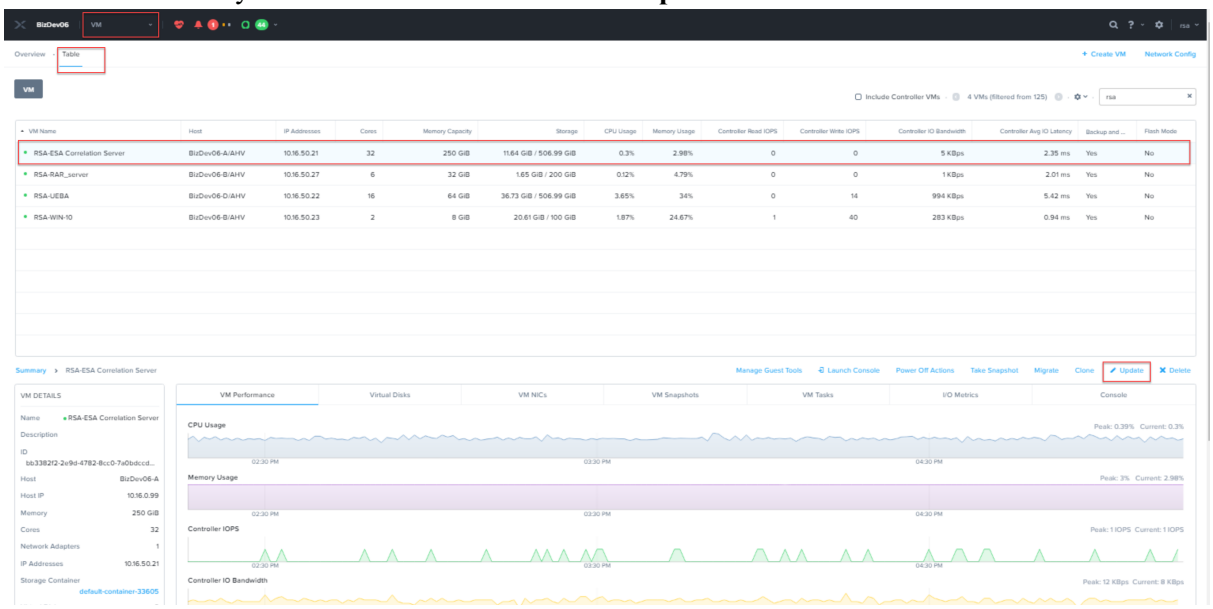


Add New Disk in Nutanix AHV

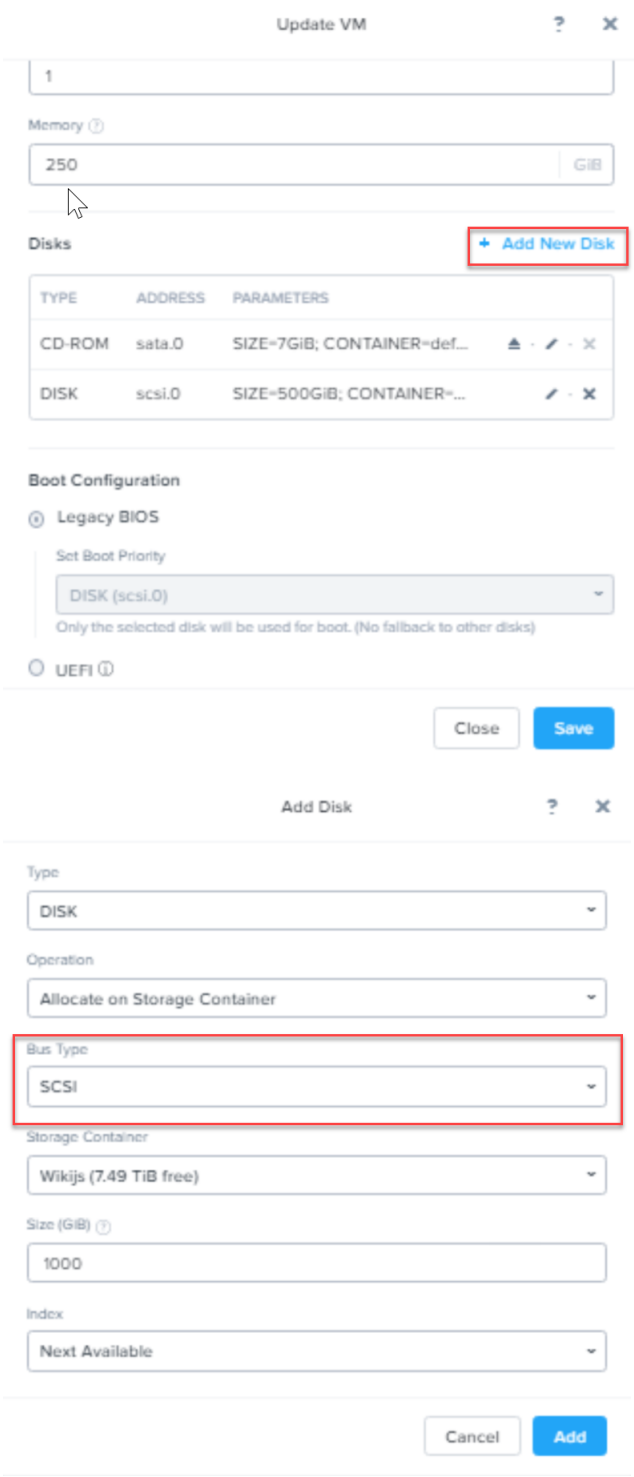
Perform the following steps to add a new disk to your Nutanix AHV VM.

1. Log in to the Nutanix Prism GUI.
2. From the drop-down menu, click **VM** and select the **Table** view.

3. Select the VM that you want to add disk to and click **Update**.



4. In the **Update VM** dialog, scroll down to the **Disks** section and click **Add New Disk**.



5. In the **Add Disk** dialog, do the following:
   a. Select **DISK** in the **Type** field and specify other details based on your requirement.
   b. Select **Allocate on Storage Container** in the **Operations** field.

c. Select **SCSI** as **Bus Type**.

6. Click **Add**.
7. In the **Update VM** dialog, click **Save**.

Task 2. Add New Volume and Extend Existing File Systems

Following commands are commonly used for the file extension.

- `/dev/sdc` for extending `nw-home` or `/var/netwitness`.
- `/dev/sdd` for creating `/var/netwitness/xxxxxx`.
- `/dev/<>` for creating `/var/netwitness/xxxxxx/metadb`.
- `/dev/<>` for creating `/var/netwitness/xxxxx/sessiondb`.
- `/dev/sde` for creating `/var/netwitness/xxxxx/index`.

Note: The number of `/dev/<>` varies based on the retention days or the number of disks attached.

Admin Server

RSA recommended partition for AdminServer (Can be changed based on the retention days).

LVM	Folder	Size	Disk Type
`/dev/netwitness_vg00/nwhome`	`/var/netwitness/`	2TB	SSD

Attach external disk for extension of `/var/netwitness/` (refer to the steps in attaching the disk) partition. Create an additional disk with suffix as `nwhome`.

Follow these steps:

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, for example if you attach one 2TB disk.
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

or,

```
lvextend -l +100%FREE /dev/netwitness_vg00/nwhome
```

6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

ESAPrimary/ESASecondary/Malware

RSA recommended partition for ESAPrimary/ESASecondary/Malware (Can be changed based on the retention days).

LVM	Folder	Size	Disk Type
`/dev/netwitness_vg00/nwhome`	`/var/netwitness/`	6TB	HDD

Attach external disk for extension of `/var/netwitness/` partition, create an external disk with suffix as nwhome.

Follow these steps:

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, for example, if you attach one 6TB disk
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvextend -L 5.9T /dev/netwitness_vg00/nwhome`
6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

Log Collector

RSA recommends the following partition for the LogCollector (Can be changed based on the retention days).

LVM	Folder	Size	Disk Type
`/dev/netwitness_vg00/nwhome`	`/var/netwitness/`	500GB	HDD

Attach an external disk for extension of /var/netwitness/ partition, create an external disk with suffix as nwhome.

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, for example if you attach one 500GB disk
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvextend -L 488G /dev/netwitness_vg00/nwhome`
6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

Log Decoder

Virtual Drive Space Ratios

The following table provides optimal configurations for packet and log hosts.

Log Decoder			
Persistent Datastores		Cache Datastore	
PacketDB	SessionDB	Meta DB	Index

100% as calculated by Sizing & Scoping Calculator	1 GB per 1000 EPS of traffic sustained provides 8 hours cache	20 GB per 1000 EPS of traffic sustained provides 8 hours cache	0.5 GB per 1000 EPS of traffic sustained provides 4 hours cache

Extending File Systems

Follow the below instructions to extend the file systems.

Attach an external disk for extension of `/var/netwitness/` partition, create an external disk with suffix as nwhome, attach other external disks for LogDecoder database partition. For extending /var/netwitness partition follow these steps:

Note: No other partition should reside on this volume, only to be used for `/var/netwitness/`

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, suppose if you had add attach one 2TB disk
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

or,

`lvextend -l +100%FREE /dev/netwitness_vg00/nwhome`

6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

Other partitions are also required. Create the following partitions on the logdecodersmall volume group.

Folder	LVM	Volume Group
`/var/netwitness/logdecoder`	`decoroot`	`logdecodersmall`
`/var/netwitness/logdecoder/index`	`index`	`logdecodersmall`
/var/netwitness/logdecoder/metadb	metadb	logdecodersmall
`/var/netwitness/logdecoder/sessiondb`	`sessiondb`	`logdecodersmall`

Follow these steps to create the partitions mentioned in the table above:

1. Execute `lsblk` and get the physical volume names from the output
2. `pvcreate /dev/sdd`
3. `vgcreate -s 32 logdecodersmall /dev/sdd`
4. `lvcreate -L <disk_size> -n <lvm_name> logdecodersmall`
5. `mkfs.xfs /dev/logdecodersmall/<lvm_name>`
6. Repeat steps 4 and 5 for all the LVM's mentioned

The following partition should be on volume group LogDecoder

Folder	LVM	Volume Group
/var/netwitness/logdecoder/packetdb	packetdb	logdecoder

Follow these steps:

1. Execute `lsblk` and get the physical volume names from the output
2. `pvcreate /dev/sde`
3. `vgcreate -s 32 logdecoder /dev/sde`
4. `lvcreate -L <disk_size> -n packetdb logdecoder`
5. `mkfs.xfs /dev/logdecoder/packetdb`

RSA recommends below sizing partition for LogDecoder (Can be changed based on the retention days)

LVM	Folder	Size	Disk Type
/dev/netwitness_vg00/nwhome	/var/netwitness/	1TB	HDD
/dev/logdecodersmall/decoroot	/var/netwitness/logdecoder	10GB	HDD
/dev/logdecodersmall/index	/var/netwitness/logdecoder/index	30GB	HDD
/dev/logdecodersmall/metadb	/var/netwitness/logdecoder/metadb	3TB	HDD
/dev/logdecodersmall/sessiondb	/var/netwitness/logdecoder/sessiondb	370GB	HDD
/dev/logdecoder/packetdb	/var/netwitness/logdecoder/packetdb	18TB	HDD

Create each directory and mount the LVM on it in a serial manner, except /var/netwitness which will be already created.

Note: Create the folder `/var/netwitness/logdecoder` and mount on `/dev/logdecodersmall/decoroot` then create the other folders and mount them.

After that add the below entries in `/etc/fstab` in the same order and mount them using `mount -a`.

```
/dev/logdecodersmall/decoroot /var/netwitness/logdecoder xfs noatime,nosuid 1 2
/dev/logdecodersmall/index /var/netwitness/logdecoder/index xfs noatime,nosuid 1 2
/dev/logdecodersmall/metadb /var/netwitness/logdecoder/metadb xfs noatime,nosuid 1 2
/dev/logdecodersmall/sessiondb /var/netwitness/logdecoder/sessiondb xfs noatime,nosuid 1 2
/dev/logdecoder/packetdb /var/netwitness/logdecoder/packetdb xfs noatime,nosuid 1 2
```

Concentrator

Virtual Drive Space Ratios

The following table provides optimal configurations for packet and log hosts.

Concentrator		
Persistent Datastores	Cache Datastores	
Meta DB	SessionDB Index	Index
Calculated as 10% of the PacketDB required for a 1:1 retention ratio	30 GB per 1TB of PacketDB for standard multi protocol network deployments as seen at typical internet gateways	5% of the calculated MetaDB on the Concentrator. Preferred High Speed Spindles or SSD for fast access

Log Concentrator		
Persistent Datastores	Cache Datastores	
Meta DB	SessionDB Index	Index
Calculated as 100% of the PacketDB required for a 1:1 retention ratio	3 GB per 1000 EPS of sustained traffic per day of retention	5% of the calculated MetaDB on the Concentrator. Preferred High Speed Spindles or SSD for fast access

Extending File Systems

Attach external disk for extension of `/var/netwitness/` partition, Create an external disk with suffix as `nwhome`, attach other external disks for Concentrator database partition.

For extending `/var/netwitness` partition follow below steps:

Note: No other partition should reside on this volume, only to be used for `/var/netwitness/`.

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, for example if you attach one 2TB disk
3. `pvcreate /dev/sdc` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

or,

`lvextend -l +100%FREE /dev/netwitness_vg00/nwhome`

6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

The following partitions are also required on volume group concentrator.

Folder	LVM	Volume Group
/var/netwitness/concentrator	root	concentrator
/var/netwitness/concentrator/sessiondb	sessiondb	concentrator
/var/netwitness/concentrator/metadb	metadb	concentrator

Follow these steps:

1. Execute `lsblk` and get the physical volume names from the output
2. `pvcreate /dev/sdd`
3. `vgcreate -s 32 concentrator /dev/sdd`
4. `lvcreate -L <disk_size> -n <lvm_name> concentrator`
5. `mkfs.xfs /dev/concentrator/<lvm_name>`
6. Repeat steps 4 and 5 for all the LVM's mentioned

Below partition should be on volume group index

Folder	LVM	Volume Group
/var/netwitness/concentrator/index	index	index

Follow these steps:

1. Execute `lsblk` and get the physical volume names from the output
2. `pvcreate /dev/sde`
3. `vgcreate -s 32 index /dev/sde`
4. `lvcreate -L <disk_size> -n index index`
5. `mkfs.xfs /dev/index/index`

RSA recommends below sizing partition for Concentrator (Can be changed based on the retention days)

LVM	Folder	Size	Disk Type
/dev/netwitness_vg00/nwhome	/var/netwitness/	1TB	HDD
/dev/concentrator/root	/var/netwitness/concentrator	10GB	HDD
/dev/concentrator/metadb	/var/netwitness/concentrator/metadb	3TB	HDD
/dev/concentrator/sessiondb	/var/netwitness/concentrator/sessiondb	370GB	HDD
/dev/index/index	/var/netwitness/concentrator/index	2TB	HDD

Create each directory and mount the LVM on it in a serial manner, except `/var/netwitness` which will be already created.

Note: Create the folder `/var/netwitness/concentrator` and mount on `/dev/concentrator/root` then create the other folders and mount them.

After that add the below entries in `/etc/fstab` in the same order

```
/dev/concentrator/root /var/netwitness/concentrator xfs noatime,nosuid 1 2
/dev/concentrator/sessiondb /var/netwitness/concentrator/sessiondb xfs noatime,nosuid 1 2
/dev/concentrator/metadb /var/netwitness/concentrator/metadb xfs noatime,nosuid 1 2 2
/dev/index/index /var/netwitness/concentrator/index xfs noatime,nosuid 1 2
```

Archiver

The following partition is required for the Archiver volume group.

Folder	LVM	Volume Group
`/var/netwitness/archiver`	`archiver`	`archiver`

Follow these steps:

1. Execute `lsblk` and get the physical volume names from the output
2. `pvcreate /dev/sde`
3. `vgcreate -s 32 archiver /dev/sde`
4. `lvcreate -L <disk_size> -n archiver archiver`
5. `mkfs.xfs /dev/archiver/archiver`

Attach an external disk for extension of `/var/netwitness/` partition, create an external disk with suffix as `nwhome`, attach other external disks for Archiver database partition.

For extending `/var/netwitness` partition follow these steps:

Note: No other partition should reside on this volume, only to be used for /var/netwitness.

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, suppose if you had add attach one 2TB disk
3. `pvcreate /dev/sdc` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

or,

```
lvextend -l +100%FREE /dev/netwitness_vg00/nwhome
```

RSA recommends the following sizing partition for the Archiver (Can be changed based on the retention days).

LVM	Folder	Size	Disk Type
`/dev/netwitness_vg00/nwhome`	`/var/netwitness/`	1TB	HDD
`/dev/archiver/archiver`	`/var/netwitness/archiver`	4TB	HDD

Create each directory and mount the LVM on it in a serial manner, except `/var/netwitness` which will be already created.

After that add the below entries in `/etc/fstab` in the same order

```
/dev/archiver/archiver /var/netwitness/archiver xfs noatime,nosuid 1 2
```

Decoder

Virtual Drive Space Ratios

The following table provides optimal configurations for packet and log hosts.

Decoder			
Persistent Datastores		Cache Datstore	
PacketDB	SessionDB	Meta DB	Index
100% as calculated by Sizing & Scoping Calculator	6 GB per 100Mb/s of traffic sustained provides 4 hours cache	60 GB per 100Mb/s of traffic sustained provides 4 hours cache	3 GB per 100Mb/s of traffic sustained provides 4 hours cache

Extending File Systems

Attach an external disk for extension of `/var/netwitness/` partition, create an external disk with suffix as nwhome, attach other external disks for decoder database partition. For extending `/var/netwitness` partition follow these steps:

Note: No other partition should reside on `/var/netwitness/`.

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, suppose if you had add attach one 2TB disk
3. `pvcreate /dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

or,

```
lvextend -l +100%FREE /dev/netwitness_vg00/nwhome
```

6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

The following four partitions should be on the decodersmall volume group.

Folder	LVM	Volume Group
`/var/netwitness/decoder`	`decoroot`	`decodersmall`

Folder	LVM	Volume Group
/var/netwitness/decoder/index	index	decodersmall
/var/netwitness/decoder/metadb	metadb	decodersmall
/var/netwitness/decoder/sessiondb	sessiondb	decodersmall

Follow these steps:

1. Execute `lsblk` and get the physical volume names from the output
2. `pvcreate /dev/sdd`
3. `vgcreate -s 32 decodersmall /dev/sdd`
4. `lvcreate -L <disk_size> -n <lvm_name> decodersmall`
5. `mkfs.xfs /dev/decodersmall/<lvm_name>`
6. Repeat steps 4 and 5 for all the LVM's mentioned

The following partition should be on the decoder volume group.

Folder	LVM	Volume Group
/var/netwitness/decoder/packetdb	packetdb	decoder

1. Execute `lsblk` and get the physical volume names from the output
2. `pvcreate /dev/sde`
3. `vgcreate -s 32 decoder /dev/sde`
4. `lvcreate -L <disk_size> -n packetdb decoder`
5. `mkfs.xfs /dev/decoder/packetdb`

RSA recommends the following sizing partition for the Decoder (Can be changed based on the retention days).

LVM	Folder	Size	Disk Type
/dev/netwitness_vg00/nwhome	/var/netwitness/	1TB	HDD
/dev/decodersmall/decoroot	/var/netwitness/decoder	10GB	HDD
/dev/decodersmall/index	/var/netwitness/decoder/index	30GB	HDD
/dev/decodersmall/metadb	/var/netwitness/decoder/metadb	3TB	HDD
/dev/decodersmall/sessiondb	/var/netwitness/decoder/sessiondb	370GB	HDD
/dev/decoder/packetdb	/var/netwitness/decoder/packetdb	18TB	HDD

Create each directory and mount the LVM on it in serial manner, except /var/netwitness which will be already created.

Note: Create the folder `/var/netwitness/decoder` and mount on `/dev/decodersmall/decoroot` then create the other folders and mount them.

After that add the below entries in `/etc/fstab` in the same order and mount them using `mount -a`.

```
/dev/decodersmall/decoroot /var/netwitness/decoder xfs noatime,nosuid 1
2
/dev/decodersmall/index /var/netwitness/decoder/index xfs
noatime,nosuid 1 2
/dev/decodersmall/metadb /var/netwitness/decoder/metadb xfs
noatime,nosuid 1 2
/dev/decodersmall/sessiondb /var/netwitness/decoder/sessiondb xfs
noatime,nosuid 1 2
/dev/decoder/packetdb /var/netwitness/decoder/packetdb xfs
noatime,nosuid 1 2
```

Endpoint Log Hybrid

Virtual Drive Space Ratios

The following table provides optimal configurations for packet and log hosts.

EndPoint Log Decoder					
	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	120 GB	26GB	6Gb	NA	152GB
Concentrator	206GB	NA	6GB	4GB	216GB
MongoDB	NA	NA	NA	NA	13GB (12 GB tracking data, 1 GB scan data)

Note: The above Endpoint Log Hybrid sizing guidelines are for 20 K agents and 20 K events per day per agent with an event size of 1500 bytes.
The same sizing guidelines are applicable for scan data with 20 K sessions per day per agent except MongoDB as mentioned above.

Extending File Systems

For Endpoint Server, attach external disk for extension of `/var/netwitness/` partition, create an external disk with suffix as nwhome.

Follow these steps:

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, for example, if you attach one 6TB disk
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`

5. `lvextend -L 5.9T /dev/netwitness_vg00/nwhome`
6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

RSA recommended partition for Endpoint Server (Can be changed based on the retention days).

LVM	Folder	Size	Disk Type
`/dev/netwitness_vg00/nwhome`	`/var/netwitness/`	6TB	HDD

For Mongo DB, attach external disk for extension of `/var/netwitness/mongo` partition, create an external disk with suffix as `nwhome`.

Follow these steps:

1. Ensure you have added a new disk. For more information, see Task 1. Add New Disk.
2. Execute `lsblk` and get the physical volume name, for example, if you attach one 6TB disk
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc1`
4. `vgextend hybrid /dev/sdc1`
5. `lvextend -L 5.9T /dev/hybrid-vlmng`
6. `xfs_growfs /dev/mapper/hybrid-vlmng`

RSA recommended partition for Mongo DB (Can be changed based on the retention days).

LVM	Folder	Size	Disk Type
`/dev/hybrid-vlmng`	`/var/netwitness/mongo`	6TB	HDD

For Log Decoder, Log Collector, and Concentrator see Log Decoder, Log Collector, and Concentrator.

UEBA

The following procedure attaches an external disk and extends the `/var/netwitness/` partition. You must use nwhome as the eternal disk suffix. This procedure illustrates how to add a 2TB disk.

Note: `/var/netwitness` is the only partition that can reside on this volume.

1. List the physical volume name.

`lsblk` (for example, `dev/mapper/sdc`)

2. Extend the `/var/netwitness/` partition.

```
pvcreate <pv_name>where pv_name is dev/mapper/sdc
vgextend netwitness_vg00 /dev/mapper/sdc
lvextend -L 1.9T /dev/mapper/netwitness_vg00/nwhome
xfs_growfs /dev/mapper/netwitness_vg00-nwhome
```

This partition is the RSA recommended partition for UEBA. You can change it based on retention days.

Task 3. Storage Configurations

For storage allocations of all host types, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness Platform 11.x*.

For more information on storage configurations using the REST API, see the "Configure Storage Using the REST API" topic in the *Storage Guide for NetWitness Platform 11.x*.

Step 3. Install NetWitness Platform

Complete the steps below to install NetWitness 11.7.

Note: Before installing the hosts make sure that the time on each host is synchronized with the time on the NetWitness Server.
To synchronize the time do one of the following:
- Configure the NTP Server. For more information, see"Configure NTP Servers" in the *System Configuration Guide*.
- Run the following commands on each hosts:
1. SSH to NW host.
2. Run the following commands.

```
systemctl stop ntpd
ntpdate nw-node-zero
systemctl start ndpd
```

This task installs:

- The 11.7.0.0 NW Server environmental platform.
- The NW Server components (that is, Admin Server, Config Server, Orchestration Server, Integration Server, Broker, Investigate Server, Reporting Engine, Respond Server and Security server).
- A repository with the RPM files required to install the other functional components or services.

Prerequisites

Deploy your 11.7.0.0 environment:

1. Add new VM.
2. Configure storage.
3. Set up firewalls.

Install NetWitness Platform

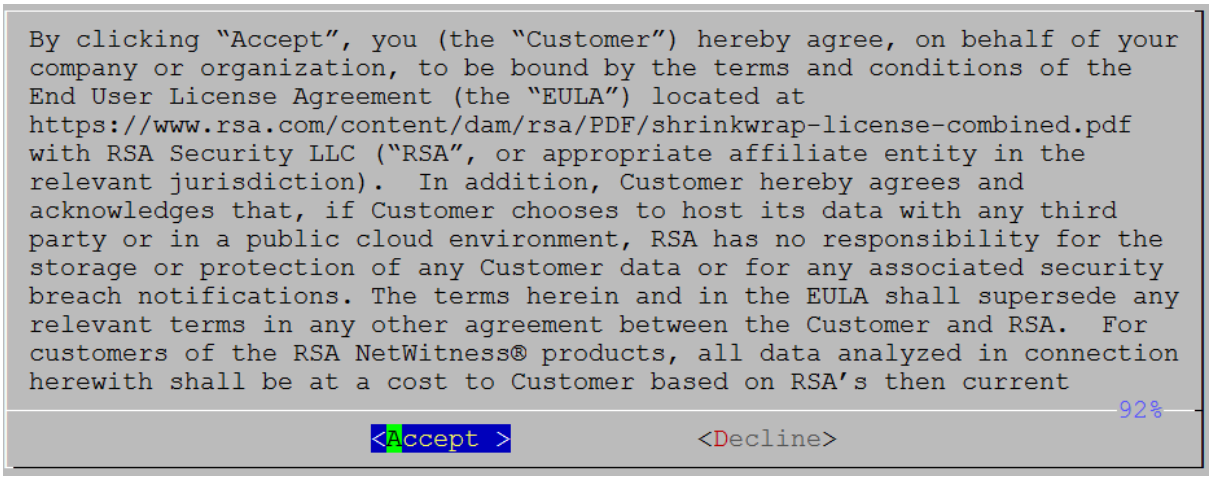
Caution: If you want to install the Endpoint Relay Server, do not run the `nwsetup-tui` script. Follow the instructions in "(Optional) Installing and Configuring Relay Server" in the *NetWitness Endpoint Configuration Guide*.

IMPORTANT: In NetWitness Platform version 11.6 or Later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`.
If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

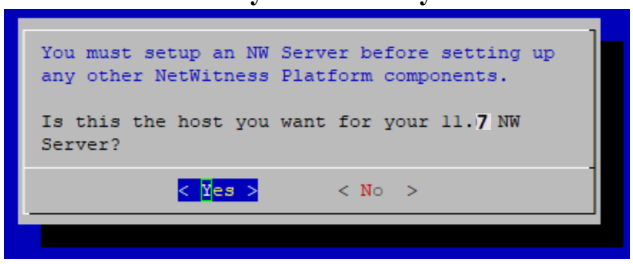
1. Log in to the host with the `root` credentials and run the `nwsetup-tui` command to set up the host.

   This initiates the `nwsetup-tui` (Setup program) and the EULA is displayed.

   Note: Use the following options to navigate the Setup prompts.
   1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, and use the Tab key to move to and from commands (such as **<Yes>**, **<No>**, **<OK>**, and **<Cancel>**). Press **Enter** to register your command response and move to the next prompt.
   2.) The Setup program adopts the color scheme of the desktop or console you use to access the host.
   3.) If you specify DNS servers during the Setup program (`nwsetup-tui`) execution, they MUST be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` script to proceed. Any misconfigured DNS servers cause the Setup program to fail. If you need to reach a DNS server after setup that is unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see "Change Host Network Configuration" topic in the System Maintenance Guide.
   If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Update Repository** prompt in step 12 (the DNS servers are not defined so the system cannot access the external repo).

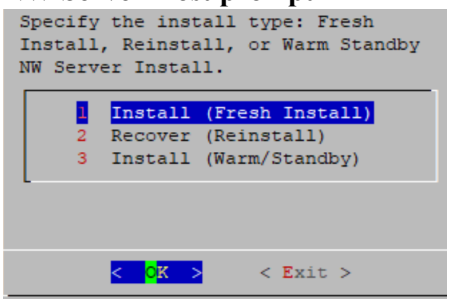
2. Tab to **Accept** and press **Enter**.
   The **Is this the host you want for your 12.0 NW Server** prompt is displayed.

3. Tab to **Yes** and press **Enter** to install 12.0 on the NW Server.
   Tab to **No** and press **Enter** to install 12.0 on other component hosts.

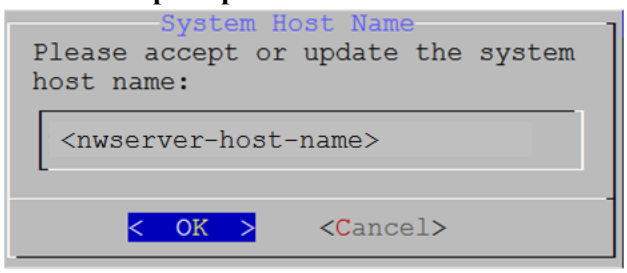
   Caution: If you choose the wrong host for the NW Server and complete the Setup, you must restart the Setup Program (step 2) and complete steps all the subsequent steps to correct this error.

4. The **Install** prompt is displayed (**Recover** does not apply to the installation.).
   NW Server Host prompt:



   Other Component Hosts, the prompt is the same, but does not include option 3 Install (Warm/Standby)

5. Press **Enter**. **Install (Fresh Install)** is selected by default.
   The **System Host Name** prompt is displayed.
   NW Server prompt:

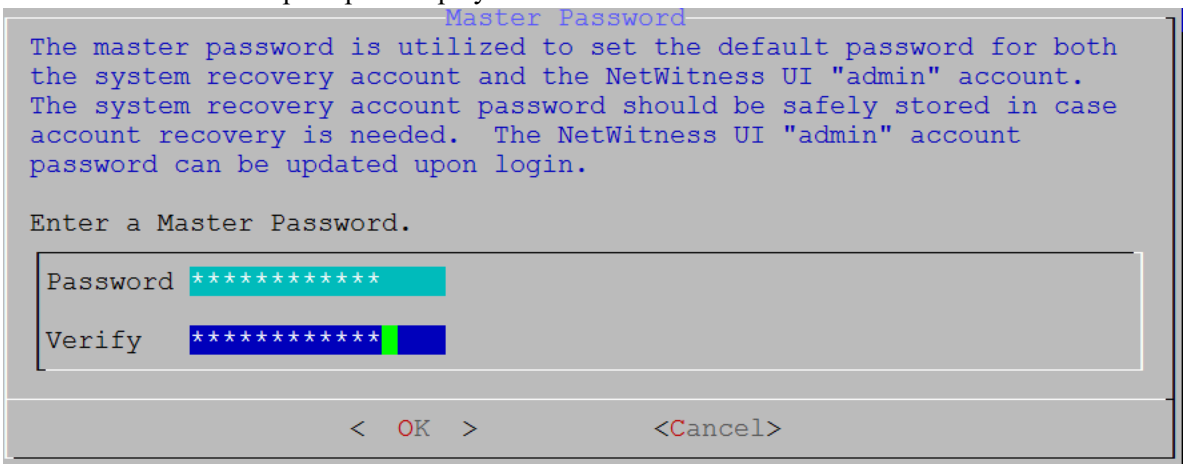


   Other Component Hosts prompt says <non-nwserver-host-name>

   Caution: If you include "." in a host name, the host name must also include a valid domain name.

   Press **Enter** if want to keep this name. If not, edit the host name, tab to **OK**, and press **Enter** to change it.

6. **This step applies only to NW Server hosts**.
   The **Master Password** prompt is displayed.



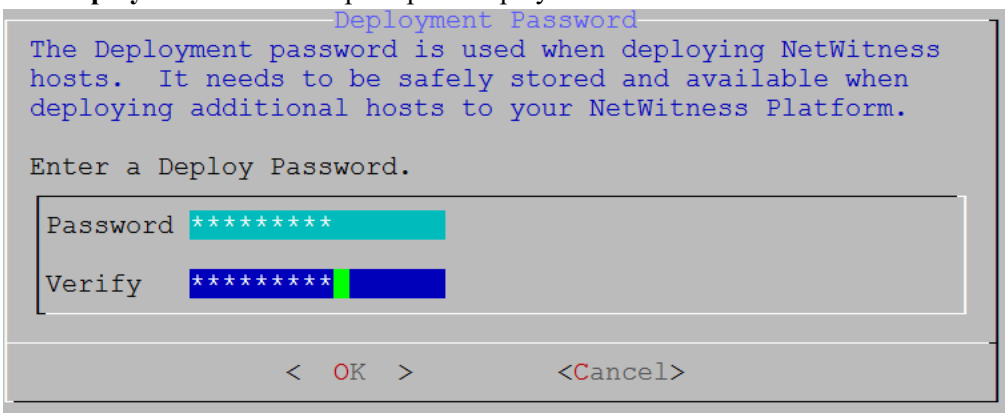
   The following list of characters are supported for Master Password and Deployment Password:

- Symbols: **! @ # % ^ +**
- Numbers: **0-9**
- Lowercase Characters: **a-z**
- Uppercase Characters: **A-Z**

No ambiguous characters are supported for Master Password and Deployment Password. For example:
space { } [ ] ( ) / \ ' " ` ~ ; : . < > -

Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

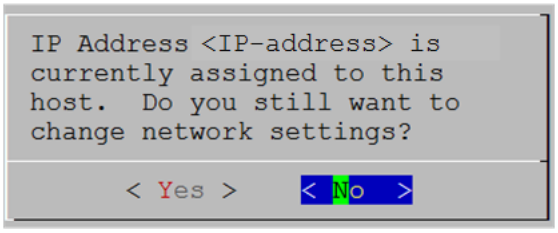
7. **This step applies to both NW Server hosts and component hosts.**
   The **Deployment Password** prompt is displayed.



   Type the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

8. One of the following conditional prompts is displayed.

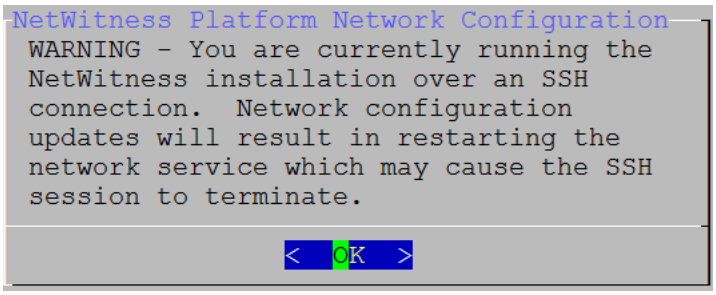
   - If the Setup program finds a valid IP address for this host, the following prompt is displayed.



     Press **Enter** if you want to use this IP and avoid changing your network settings. Tab to **Yes** and press **Enter** if you want to change the IP configuration on the host.

   - If you are using an SSH connection, the following warning is displayed.

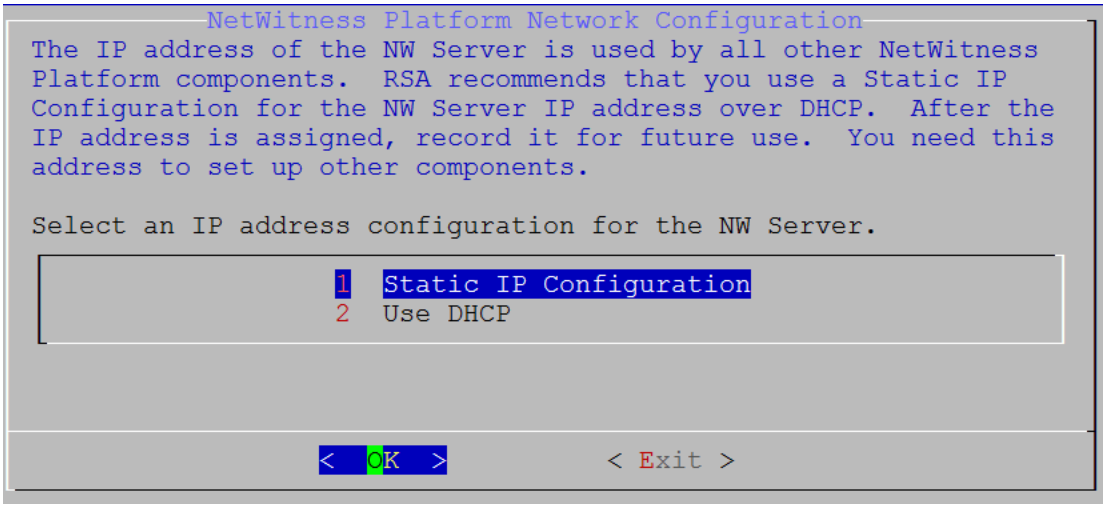
     Note: If you connect directly from the host console, the following warning is not displayed.



Press **Enter** to close warning prompt.

- If the Setup Program finds an IP configuration and you choose to use it, the **Update Repository** prompt is displayed. Go to step 12 and complete the installation.
- If the Setup Program did not find an IP configuration or if you choose to change the existing IP configuration, the **Network Configuration** prompt is displayed.

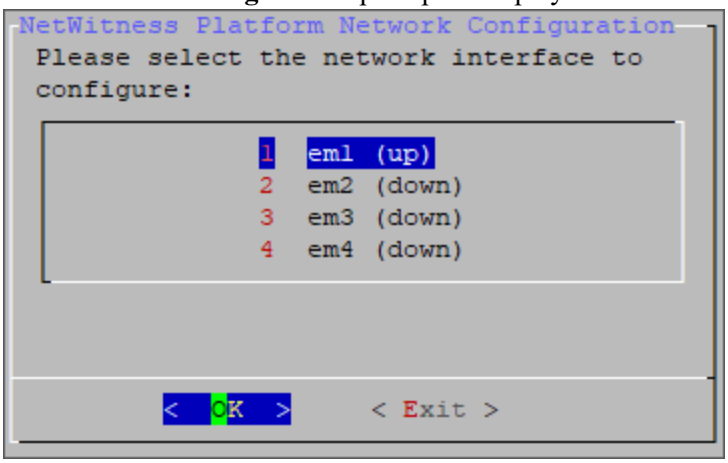
Caution: Only select "**Use DHCP**" as an IP address configuration for the NW Server if DHCP issues static IP addresses.



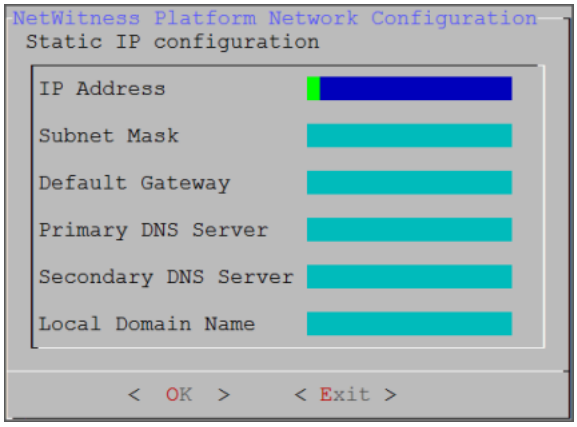
Tab to **OK** and press **Enter** to use **Static IP**.

If you want to use DHCP, down arrow to **2 Use DHCP** and press **Enter**.

The **Network Configuration** prompt is displayed.



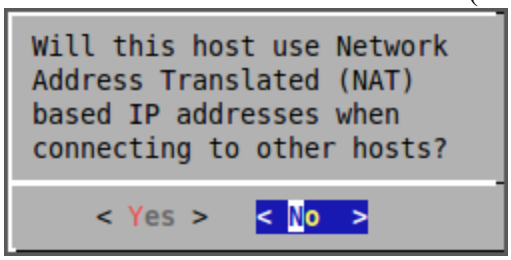
9. Down arrow to the network interface you want, tab to **OK**, and press **Enter**. If you do not want to continue, tab to **Exit.**
   The following **Static IP Configuration** prompt is displayed.



10. Type the configuration values, tab to **OK**, and press **Enter**. If you do not complete all the required fields, an `All fields are required` error message is displayed (**Secondary DNS Server** and **Local Domain Name** fields are not required). If you use the wrong syntax or character length for any of the fields, an `Invalid <field-name>` error message is displayed.

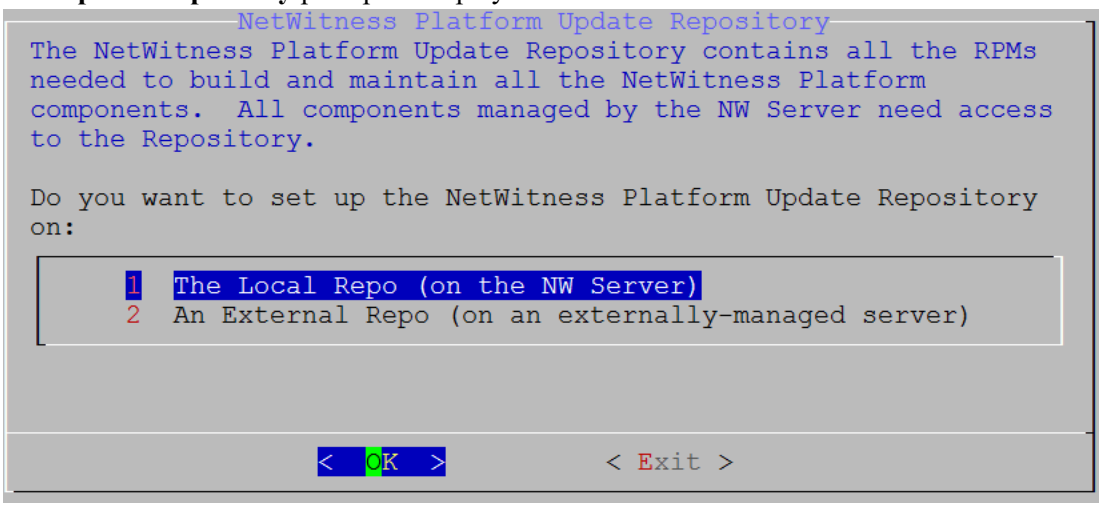
    Caution: If you select **DNS Server**, make sure that the DNS Server is correct and the host can access it before proceeding with the installation.

11. The Use Network Address Translation (NAT) prompt is displayed.



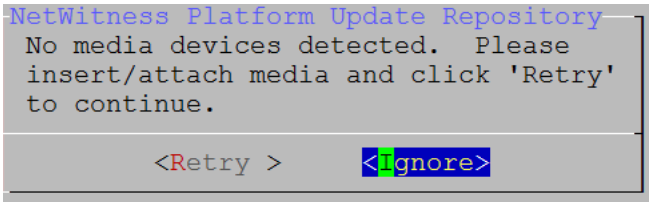
    For the NW Server, tab to **No** and press **Enter**.
    For component hosts, if this host requires the use of NAT-based addresses to communicate with the NW Server, tab to **Yes**. Otherwise, tab to **No** and press **Enter**.

12. The **Update Repository** prompt is displayed.

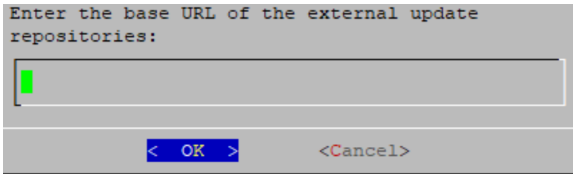


For the NW Server:

- Press **Enter** to choose the **Local Repo**.
- If you want to use an external repo, down arrow to **External Repo**, tab to **OK**, and press **Enter**. If you select **1 The Local Repo (on the NW Server)** in the Setup program, make sure that you have the appropriate media attached to the host (media that contains the ISO file, for example a build stick) from which it can install NetWitness 12.0. If the program cannot find the attached media, you receive the following prompt.



- If you select **2 An External Repo (on an externally-managed server)**, the UI prompts you for a URL. The repositories give you access to RSA updates and CentOS updates. Refer to "Appendix B. Create an External Repo" in this guide for instructions on how to create this repo and its external repo URL so you can enter it in the following prompt.

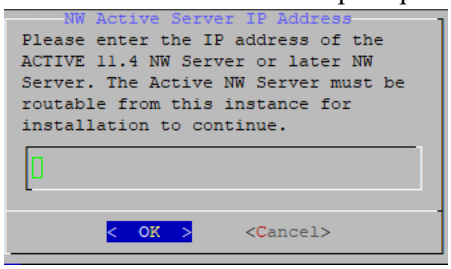


Enter the base URL of the NetWitness external repo and click **OK.** The **Start Install** prompt is displayed.

For component hosts:

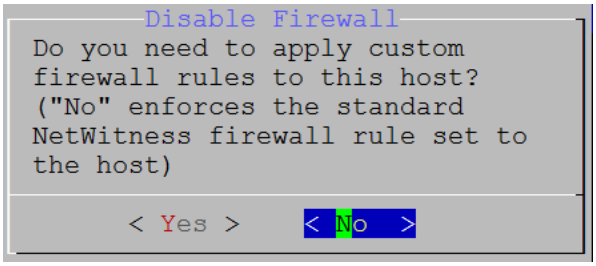
- Select the same repo that you selected when you installed the NW Server host and follow the steps above.

- The NW Server IP Address prompt is displayed.

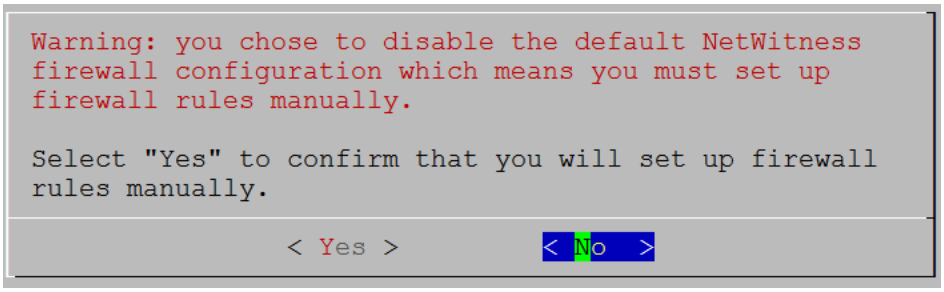


Type the NW Server IP address. Tab to **OK** and press **Enter**.

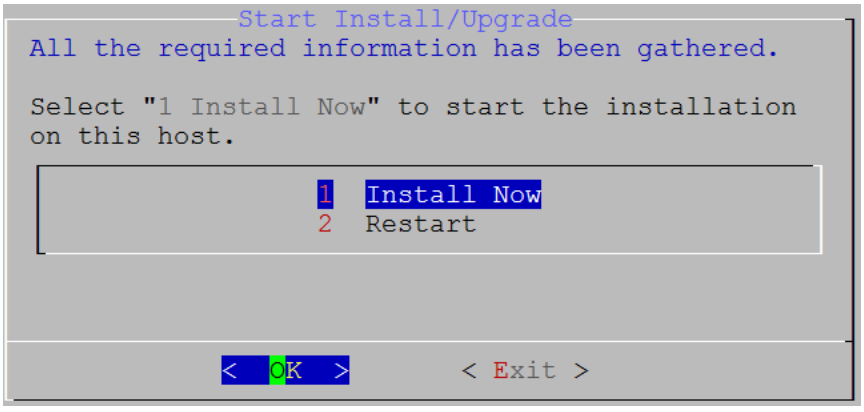
13. The Disable firewall prompt is displayed.



Tab to **No** (default), and press **Enter** to use the standard firewall configuration.
To disable the standard firewall configuration, tab to **Yes**, and press **Enter**.
If you select **Yes**, confirm your selection(select **Yes** again) or select **No** to use the standard firewall configuration.

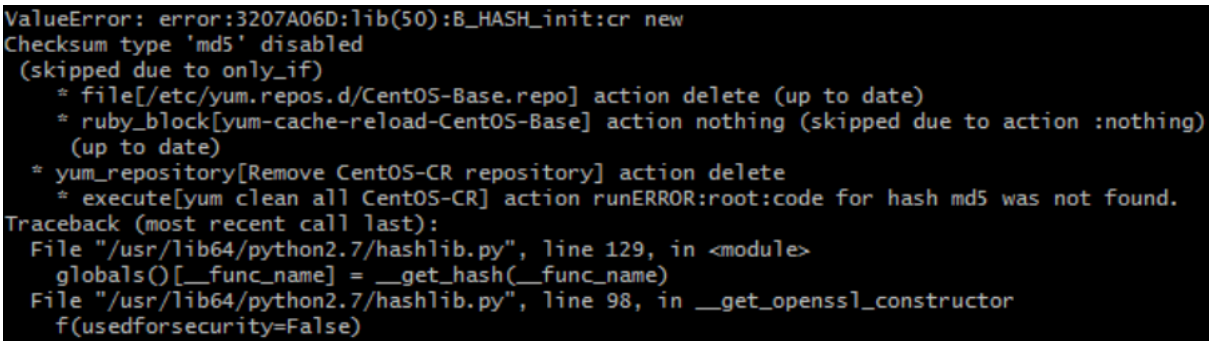


14. The **Start Install** prompt is displayed.



15. Press **Enter** to install 12.0.
When **Installation complete** is displayed, you have installed 12.0 on this host.

Note: Ignore the hash code errors similar to the errors shown in the following figure that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.



16. (Optional) If your system configuration requires that a component host must use a NAT IP address to reach the NW Server host, you must configure the NAT IP address of the NW Server by running the following command:
    `nw-manage --update-host --host-id <NW Server Host UUID> --ipv4-public <NAT IP address>`

Set Up ESA Hosts

After you install your NW Server and component hosts, follow these steps to set up your ESA hosts.

- Install your primary ESA host following the instructions in "Install 12.0 on the NetWitness Server (NW Server) Host and Other Component Hosts" in this guide, and install the **ESA Primary** service on it after you finish the Set Up program in the UI in (Admin) > **Hosts** > Install:
- (Conditional) If you have a secondary ESA host, install it and install the **ESA Secondary** service on it after you finish the Set Up program in the UI in (Admin) > **Hosts** > Install:

Install Component Services on Hosts

After you have installed NW Server and component hosts, and set up your ESA hosts, follow these steps to install component services, such as Decoders and Concentrators, on your host systems.

1. Install a component service on the host.
    a. Log into NetWitness and go to (missing or bad snippet)> **Hosts**.
    The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background.

    Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

    b. Select the host in the **New Hosts** dialog and click **Enable**.
    The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

    c. Select that host in the **Hosts** view and click .

The **Install Services** dialog is displayed.

d. Select the appropriate host type (for example, **Concentrator**) in **Category** and click **Install**.

Complete Licensing Requirements

Complete licensing requirements for installed services. See the *NetWitness Platform 12.0 Licensing Management Guide* for more information. Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

(Optional) Install Warm Standby NW Server

Refer to "Warm Standby NW Server Host" under "Deployment Option Setup Procedures" in the *Deployment Guide for NetWitness Platform 12.0* for instructions on how to set up a Warm Standby NW Server.

Step 4. Configure Host-Specific Parameters

Certain application-specific parameters are required to configure log ingest and packet capture in the Virtual Environment.

Configure Log Ingest in the Virtual Environment

Log ingest is easily accomplished by sending the logs to the IP address you have specified for the Decoder. The Decoder's management interface allows you to then select the proper interface to listen for traffic on if it has not already selected it by default.

Configure Packet Capture in the Virtual Environment

There are two options for capturing packets in a VMware environment. The first is setting your vSwitch in promiscuous mode and the second is to use a third-party Virtual Tap.

Set a vSwitch to Promiscuous Mode

The option of putting a switch whether virtual or physical into promiscuous mode, also described as a SPAN port (Cisco services) and port mirroring, is not without limitations. Whether virtual or physical, depending on the amount and type of traffic being copied, packet capture can easily lead to over subscription of the port, which equates to packet loss. Taps, being either physical or virtual, are designed and intended for loss less 100% capture of the intended traffic.

Promiscuous mode is disabled by default, and should not be turned on unless specifically required. Software running inside a virtual machine may be able to monitor any and all traffic moving across a vSwitch if it is allowed to enter promiscuous mode as well as causing packet loss due to over subscription of the port..

To configure a portgroup or virtual switch to allow promiscuous mode:

1. Log on to the ESXi/ESX host or vCenter Server using the vSphere Client.
2. Select the ESXi/ESX host in the inventory.

3. Select the **Configuration** tab.
4. In the **Hardware** section, click **Networking**.
5. Select **Properties** of the virtual switch for which you want to enable promiscuous mode.
6. Select the virtual switch or portgroup you want to modify, and click **Edit**.
7. Click the **Security** tab. In the **Promiscuous Mode** drop-down menu, select **Accept**.

Use of a Third-Party Virtual Tap

Installation methods of a virtual tap vary depending on the vendor. Please refer to the documentation from your vendor for installation instructions. Virtual taps are typically easy to integrate, and the user interface of the tap simplifies the selection and type of traffic to be copied.

Virtual taps encapsulate the captured traffic in a GRE tunnel. Depending on the type you choose, either of these scenarios may apply:

- An external host is required to terminate the tunnel, and the external host directs the traffic to the Decoder interface.
- The tunnel send traffic directly to the Decoder interface, where NetWitness handles the de-encapsulation of the traffic.

Step 5. Post Installation Tasks

This topic contains the tasks you complete after you install 11.7.

- Event Stream Analysis (ESA)
- NetWitness Endpoint
- NetWitness UEBA

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Event Stream Analysis (ESA)

Configure Meta Keys on New ESA Hosts to Match Upgraded ESA Hosts in the Same NetWitness Platform Network

If you have one or more ESA hosts in a NetWitness Platform network, which were upgraded from a version before 11.3.0.2 to 11.7, and you add a new ESA host, you must configure the meta keys on the new ESA host to match the other ESA hosts. All ESA Correlation services on the same NetWitness Platform network must have the same Meta Key configurations.

1. For each ESA Correlation service on an upgraded ESA host and for the ESA Correlation service on the newly installed ESA host:
   a. Open a new tab, go to  (Admin) > **Services**, and in the Services view, select the ESA Correlation service and then select  > **View** > **Explore**.
   b. In the Explore view node list for the ESA Correlation service, select **correlation** > **stream**.
2. Ensure that the **`multi-valued`** and **`single-valued`** meta key values are the same on each of the upgraded ESA Correlation services.
3. Ensure that the **`multi-valued`** and **`single-valued`** meta key values on the newly installed ESA host are the same as those on the upgraded services.
4. To apply any changes on the ESA Correlation services, go to  **(Configure)** > **ESA Rules** and click the **Settings** tab. In the Meta Key References, click the **Meta Re-Sync (Refresh)** icon ().
5. If you updated the ESA Correlation services, redeploy the ESA rule deployments.

For more information, see "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*.

NetWitness Endpoint

The tasks in this section only apply to customers that use the NetWitness Endpoint component of NetWitness Platform.

Install Endpoint Log Hybrid

Depending on the number of agents and the location of the agents, you can choose to deploy a single Endpoint Log Hybrid host or multiple Endpoint Log Hybrid hosts. To deploy a host, you provision it and install a category on it.

- **Single Endpoint Log Hybrid host** - Deploy NetWitness Server host, Endpoint Log Hybrid host, and ESA host or hosts.
- **Multiple Endpoint Log Hybrid hosts** - Deploy NetWitness Server host, ESA host or hosts, Endpoint Log Hybrid hosts. For a consolidated view of all endpoint data from multiple Endpoint Log Hybrid hosts, install the Endpoint Broker.

  Note: NetWitness recommends that you co-locate the Endpoint Broker on the NetWitness Broker host. However, you can deploy the Endpoint Broker on a separate host or co-locate it on the Endpoint Log Hybrid.

  Note: You must plan to scale your ESA deployment to support multiple Endpoint Log Hybrid hosts.

Follow these steps to deploy an Endpoint Log Hybrid host.

Complete the following steps first:

- For a physical host, complete steps 1 - 16 in "Install NetWitness Platform" under Installation Tasks in the *Physical Host Installation Guide for NetWitness Platform 11.7*
- For a virtual host, complete steps 1 - 16 in "Step 4. Install NetWitness Platform" under Install NetWitness Platform Virtual Host in Virtual Environment in the *Virtual Host Installation Guide for NetWitness Platform 11.7*

Configuring Multiple Endpoint Log Hybrids

Follow these steps to install another Endpoint Log Hybrid.

Step 1: Install additional Endpoint Log Hybrid

- To install a physical host, complete steps 1 - 16 in "Install NetWitness Platform" under Installation Tasks in the *Physical Host Installation Guide for NetWitness Platform 12.0*
- To install a virtual host, complete steps 1 - 16 in "Step 4. Install NetWitness Platform" under Install NetWitness Platform Virtual Host in Virtual Environment in the *Virtual Host Installation Guide for NetWitness Platform 12.0*

Step 2: Setup the Endpoint Log Hybrid

1. Create a directory

   ```
   mkdir -p /etc/pki/nw/nwe-ca.
   ```

2. Copy the following certificates from the first Endpoint Log Hybrid to the secondary Endpoint Log Hybrid:

```
/etc/pki/nw/nwe-ca/nwerootca-cert.pem
/etc/pki/nw/nwe-ca/nwerootca-key.pem
```

Note: NetWitness recommends that you copy certificates from Endpoint Log Hybrid to secondary Endpoint Log Hybrid CentOS to Windows using the `SCP` command to avoid any corruption caused by Antivirus or third-party tools.

Step 3: Switch to the NetWitness UI and add Hosts

- See Add Hosts to the Endpoint Log Hybrid: for more information.

Add Hosts to the Endpoint Log Hybrid:

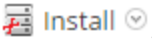
1. Log into NetWitness Platform and click  **(Admin) > Hosts**.

   The New Hosts dialog is displayed with the Hosts view grayed out in the background.

   Note: If the New Hosts dialog is not displayed, click **Discover** in the **Hosts** view toolbar.

2. Select the host in the **New Hosts** dialog and click **Enable**.

   The New Hosts dialog closes and the host is displayed in the Hosts view.

3. Select that host in the **Hosts** view (for example, **Endpoint**) and click 

.

   The Install Services dialog is displayed.

4. Select the **Endpoint Log Hybrid** category and click **Install**.
5. Make sure that the Endpoint Log Hybrid service is running.
6. Configure Endpoint Meta forwarding.

   See the *Endpoint Configuration Guide* for instructions on how to configure Endpoint Meta forwarding.

7. Deploy the ESA Rules from the Endpoint Rule Bundle. For more information, see "Deploy Endpoint Risk Scoring Rules on ESA" section in the ESA Configuration Guide.

   Note: The Endpoint IIOCs are available as OOTB Endpoint Application rules.

8. Review the default policies and create groups to manage your agents. See *Endpoint Configuration Guide*.

   Note: In 11.3 or later, agents can operate in Insights or Advanced mode depending on the policy configuration. The default policy enables the agent in an advanced mode. If you want to continue to use the Insights agent, before updating, review the policy, and make sure that the Agent mode is set to Insights.

9. Install the Endpoint Agent. You can install an Insights (free version) or an Advanced agent (licensed). See *Endpoint Agent Installation Guide* for detailed instructions on how to install the agent.

Note: You can migrate the Endpoint Agent from 4.4.0.x to 11.7. For more information, see the *NetWitness Endpoint 4.4.0.x to NetWitness Platform 11.7 Migration Guide*.

(Optional) Configure an Endpoint Service on an Existing Log Decoder Host

You can install an Endpoint service category on an existing Log Decoder host. For an overview of installing service categories on hosts, see "Hosts and Services Set Up Procedures" in the *Host and Services Getting Started Guide*. (missing or bad snippet)

- If you have an existing Endpoint Log Hybrid, you must copy certificates from that Endpoint Hybrid host to the Log Decoder before you install the Endpoint service category on the Log Decoder.
- If you do not have an Endpoint Log Hybrid host, you do not need to copy over the certificates before you install the Endpoint service category on the Log Decoder.

Do You Need to Install an Endpoint Service onto Separate Hardware

If you are only using NW Platform for collecting and analyzing logs, you can co-locate your Endpoint Server on the same physical hardware as your Log Decoder. For more information, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness Platform*.

If you exceed these guidelines, the amount of disk space usage and CPU might become so high as to create alarms for your Endpoint Server in Health and Wellness. If you notice this, and are running both log collection and EDR scans, you can use Throttling to control the amount of data coming into the Log Decoder.

If that doesn't help, NetWitness recommends that you move your Endpoint Server onto separate hardware from that used by your Log Decoder.

Install an Endpoint Service Category on an Existing Log Decoder

To install an Endpoint service category on an existing Log Decoder if you have an existing Endpoint Log Hybrid:

1. Create a directory `mkdir -p /etc/pki/nw/nwe-ca`.
2. Copy the following certificates from the first Endpoint Log Hybrid to the Log Decoder on which you are going to install the additional **Endpoint** service category.

   Note: NetWitness recommends that you copy certificates from Endpoint Log Hybrid to secondary Endpoint Log Hybrid using the `SCP` command to avoid any corruption caused by Antivirus or third-party tools.

   `/etc/pki/nw/nwe-ca/nwerootca-cert.pem`

   `/etc/pki/nw/nwe-ca/nwerootca-key.pem`

3. Log into NetWitness Platform and click  **(Admin) > Hosts**.
4. Select the Log Decoder host in the **Hosts** view and click .

   The Install Services dialog is displayed.
5. Select the **Endpoint** category and click **Install**.

To install an Endpoint service category on an existing Log Decoder if you do not have an existing Endpoint Log Hybrid:

1. Log into NetWitness Platform and click **(Admin) > Hosts**.
2. Select the Log Decoder host in the **Hosts** view and click Install.
   The Install Services dialog is displayed.
3. Select the **Endpoint** category and click **Install**.

NetWitness UEBA

The tasks in this section only apply to customers that use the UEBA component of NetWitness Platform.

Install UEBA

To set up NetWitness UEBA in NetWitness Platform 11.7, you must install and configure the NetWitness UEBA service.

The following procedure shows you how to install the NetWitness UEBA service on a NetWitness UEBA Host Type and configure the service.

1. For:
   - A physical host, complete steps 1 - 16 in "Install NetWitness Platform" under "Installation Tasks" in the *Physical Host Installation Guide for NetWitness Platform 11.7*.
   - A virtual host, complete steps 1 - 16 in "Step 4. Install NetWitness Platform" in the *Virtual Host Installation Guide for NetWitness Platform 11.7*.

   Note: The Kibana and Airflow webserver User Interface password is the same as the deploy admin password. Make sure that you record this password and store it in a safe location.
2. Log into NetWitness Platform and go to **(Admin) > Hosts**.
   The New Hosts dialog is displayed with the Hosts view grayed out in the background.

   Note: If the New Hosts dialog is not displayed, click **Discover** in the **Hosts** view toolbar.
3. Select the host in the **New Hosts** dialog and click **Enable**.
   The New Hosts dialog closes and the host is displayed in the Hosts view.
4. Select that host in the **Hosts** view (for example, **UEBA**) and click Install.
   The Install Services dialog is displayed.
5. Select the **UEBA** Host Type and click **Install**.
6. Make sure that the UEBA service is running.
7. Complete licensing requirements for NetWitness UEBA.
   See the *Licensing Management Guide* for more information.

Note: NetWitness Platform supports the User and Entity Behavior Analytics License (UEBA). This license is used based on the number of users. The Out-of-the-Box Trial License is a 90-day trial license. In case of UEBA licenses, the 90-day trial period begins from the time the UEBA service deployed on the NetWitness Platform product.

Configure NetWitness UEBA

To start running UEBA:

1. Define the following parameters: data schemas, data source (NetWitness Broker or Concentrator) and start date.
   a. Define UEBA schemas:
   Choose schemas from the following list:

   AUTHENTICATION, FILE, ACTIVE_DIRECTORY, PROCESS, REGISTRY and TLS.

   Note: The TLS packet requires adding the hunting package and enabling the JA3 feature. For more information regarding events that each schema contains, see the *NetWitness UEBA Configuration Guide*.

   b. Define the data source:
   If your deployment has multiple Concentrators, we recommend that you assign a Broker at the top of your deployment hierarchy for the NetWitness UEBA data source.

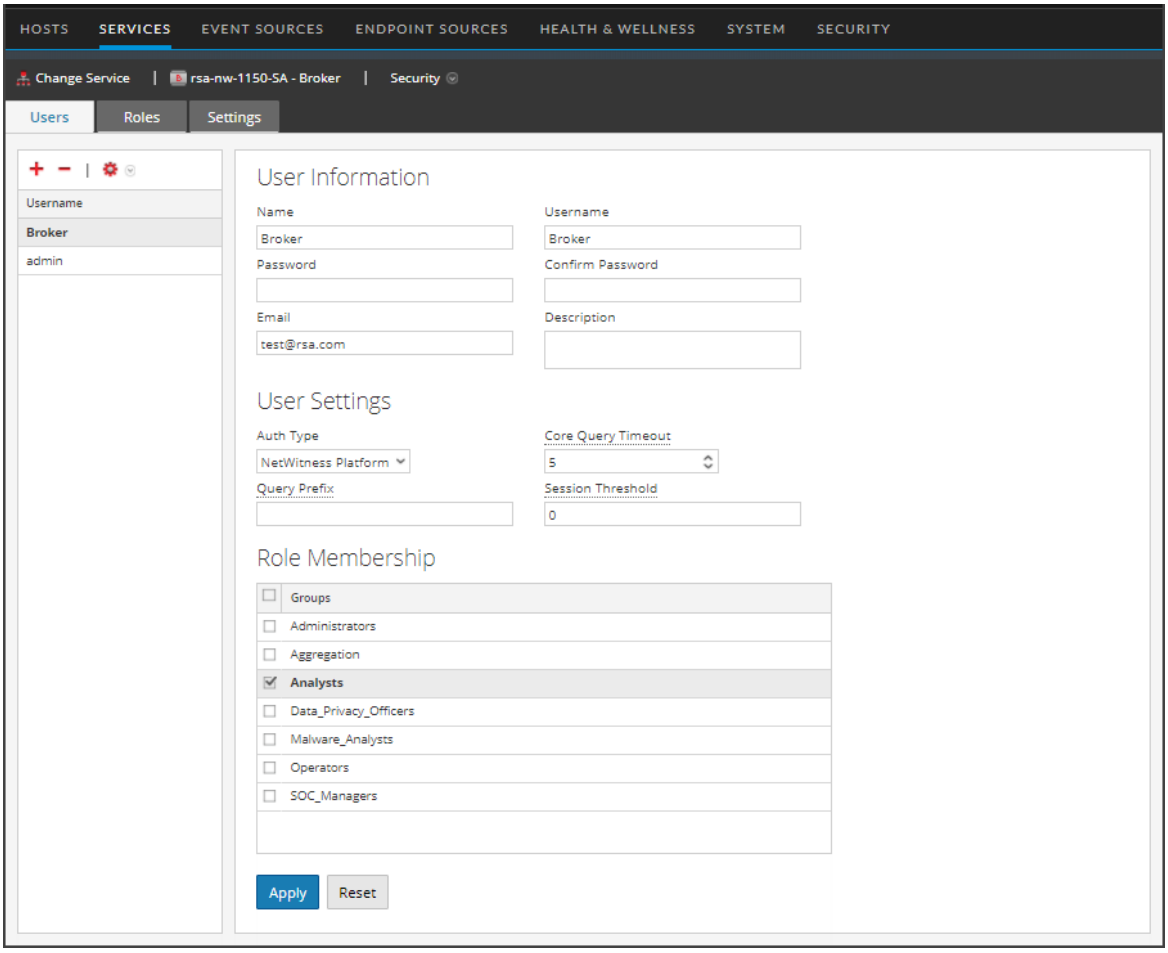
   c. Define the UEBA start-date:

   Note: The selected start date must contain events from all configured schemas.

   NetWitness recommends that the UEBA start date is set to 28 days earlier than the current date. For UEBA systems that intend to process TLS data, you must make sure that the start date is set to no later than 14 days earlier than the current date.

2. . Create a user account for the data source (Broker or Concentrator) to authenticate to the data.
   a. Log into NetWitness Platform.
   b. Go to (missing or bad snippet) **> Services**.
   c. Locate the data source service (Broker or Concentrator).
   Select that service, and select  (Actions) > **View** > **Security**.
   d. Create a new user and assign the “Analysts” role to that user.

The following example shows a user account created for a Broker.



3. SSH to the NetWitness UEBA server host.
4. Set the appropriate parallelism value:
   If the UEBA system runs on VM, update the airflow parallelism value to be 64 by running the following command:
   ```
   sed -i "s|parallelism = 256|parallelism = 64|g"
   /var/netwitness/presidio/airflow/airflow.cfg
   ```
5. Submit the following commands with the above parameters that you already defined.
   ```
   /opt/rsa/saTools/bin/ueba-server-config -u <user> -p <password> -h <host> -
   o <type> -t <startTime> -s <schemas> -v -e <argument>
   ```
   Where:

Argument	Variable	Description
`-u`	`<user>`	User name of the credentials for the Broker or Concentrator instance that you are using as a data source.

Argument	Variable	Description
-p	<password>	Password of the credentials for the Broker or Concentrator instance that you are using as a data source. The following special characters are supported in a password. !"#$%&()*+,-:;<=>?@[\]^_`\{\|} If you want to include a special character or special characters, you must delimit the password with an apostrophe sign, for example: sh /opt/rsa/saTools/bin/ueba-server-config -u brokeruser -p **'!"UHfz?@ExMn#$'** -h 10.64.153.104 -t 2018-08-01T00:00:00Z -s 'AUTHENTICATION FILE ACTIVE_DIRECTORY TLS PROCESS REGISTRY' -o broker -v
-h	<host>	IP address of the Broker or Concentrator used as the data source. Currently, only one data source is supported.
-o	<type>	Data source host type (broker or concentrator).
-t	<startTime>	Historical start time as of which you start collecting data from the data source in YYYY-MM-DDTHH-MM-SSZ format (for example, 2018-08-15T00:00:00Z). **Note:** The script interprets the time you enter as UTC (Coordinated Universal Time) and it does not adjust the time to your local time zone.
-s	<schemas>	Array of data schemas. If you want to specify multiple schemas, use a space to separate each schema (for example, AUTHENTICATION FILE ACTIVE_ DIRECTORY PROCESS REGISTRY TLS).
-v		**verbose mode**.
-e	<argument>	Boolean Argument. This enables the UEBA indicator forwarder to Respond. **Note:** If your NetWitness deployment includes an active Respond server, you can transfer NetWitness UEBA indicators to the Respond server and create incidents by enabling the indicator forwarder, from this data. For more information on how to enable the NetWitness UEBA incidents aggregation, see Step 5. Post Installation Tasks.

6. If you are deploying a hot fix on 11.x.x.x version, you must do the following:
   a. Run the presidio-upgrade DAG.

b. Press the play sign next to the DAG and then click the trigger button.

7. Set the appropriate "Boot Jar Pools" slots:

- **Virtual Appliance**: If the UEBA system is running on VM and update the `spring_boot_jar_pool` and the `retention_spring_boot_jar_pool` slots values to 22.
  To update the “Spring Boot Jar Pools” slots, Go to the Airflow main page, tap the “Admin” tab at the top bar and tap “Pools”.

a. To access the Airflow UI, go to `https://<UEBA_host>/admin` and enter the credentials.
   User: admin
   Password: The environment deploy admin password

b. Click on the pencil mark of the polls to update the slot values.



Enable Access Permission for the NetWitness UEBA User Interface

After you install NetWitness UEBA 11.7, you need to assign the UEBA_Analysts and Analysts roles to the UEBA users. For more information, see 'Assign User Access to UEBA' topic in the *NetWitness UEBA Configuration Guide*. After this configuration, UEBA users can access the **Investigate** > **Users** view.

Note: To complete NetWitness UEBA configuration according to the needs of your organization, See the *NetWitness UEBA Configuration Guide*.

Deployment Options

NetWitness Platform has the following deployment options. See the *NetWitness Deployment Guide* for detailed instructions on how to deploy these options.

- **Analyst User Interface** - gives you access to a subset of features in the NetWitness Platform UI that you can set up in individual locations when you deploy NetWitness Platform in multiple locations. It is designed to reduce latency and improve the performance that can occur when accessing all functionality from the Primary User Interface on the NW Server Host (Primary UI).
- **Group Aggregation** - configures multiple Archiver or Concentrator services as a group and share the aggregation tasks between them.
- **New Health and Wellness Search** - New Health and Wellness is an advanced monitoring and alerting system that provides insights on the operational state of the host and services in your deployment, and helps identify potential issues.
- **Second Endpoint Server** - deploys a second Endpoint Server.

Appendix A. Troubleshooting

This section describes solutions to problems that you may encounter during installations and upgrades. In most cases, NetWitness creates log messages when it encounters these problems.

Note: If you cannot resolve an upgrade issue using the following troubleshooting solutions, contact Customer Support (https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/ta-p/563897).

This section has troubleshooting documentation for the following services, features, and processes.

- Command Line Interface (CLI)
- Event Stream Analysis

Go to the NetWitness All Versions Documents page and find NetWitness Platform guides to troubleshoot issues.

Command Line Interface (CLI)

Error Message	Command Line Interface (CLI) displays: "Orchestration failed." `Mixlib::ShellOut::ShellCommandFailed: Command execution failed. STDOUT/STDERR suppressed for sensitive resource in/var/log/netwitness/config-management/chef-solo.log`
Cause	Entered the wrong `deploy_admin` password in `nwsetup-tui`.
Solution	Retrieve your `deploy_admin` password. 1. SSH to the NW Server host. `security-cli-client --get-config-prop --prop-hierarchy nw.security-client --prop-name deployment.password` SSH to the host that failed. 2. Run the `nwsetup-tui` again using correct `deploy_admin` password.

Error Message	`ERROR com.rsa.smc.sa.admin.web.controller.ajax.health.AlarmsController - Cannot connect to System Management Service`
Cause	NetWitness sees the Service Management Service (SMS) as down after successful upgrade even though the service is running.
Solution	Restart SMS service. `systemctl restart rsa-sms`

Error Message	You receive a message in the User Interface to reboot the host after you update and reboot the host offline. 
Cause	You cannot use CLI to reboot the host. You must use the User Interface.
Solution	Reboot the host in the Host View in the User Interface.

Event Stream Analysis

For ESA Correlation troubleshooting information, see the *Alerting with ESA Correlation Rules User Guide*.

Appendix B. Silent Installation Using CLI

You can use the following Command Line Interface commands to run the installation script (`nwsetup-tui`) without getting prompted for inputs. This enables you to automate the installation of a host by supplying response to the scripts prompts through the command line.

1. After you have created a base image on the host, log in to the host with the `root` credentials.
2. Submit the `nwsetup-tui` script with the `--silent` command and the arguments that you want to apply.

   The following command string is an example of how you would install a basic NW Server host.

   ```
   nwsetup-tui --silent --is-head=true --host-name=new-host --master-pass=netwitness --deploy-pass=netwitness --repo-type=1 --custom-firewall=false --ip-override=false --eula=true
   ```

Note: In NetWitness Platform version 11.6 or later, deployment account password must contain at least one number, one upper and lower case letter, and one special characters (!@#%^,+ .) along with the existing policy. The same password policy applies while updating deploy_admin password using `nw-manage script`.
If `deploy_admin` password is changed on Primary NW Server, It must be changed on the Warm Standby Server if it exists.

3. (Conditional - For Component Hosts Only) Install the appropriate service **Category** on the newly provisioned host in the NetWitness Platform Hosts view.

   a. Log into NetWitness and go to  **(Admin)** > **Hosts**.

   The **New Hosts** dialog is displayed with the **Hosts** view grayed out in the background

   Note: If the **New Hosts** dialog is not displayed, click **Discover** in the Hosts view toolbar.

   b. Select the host in the **New Hosts** dialog and click **Enable**.

   The **New Hosts** dialog closes and the host is displayed in the **Hosts** view.

   c. Select that host in the **Hosts** view (for example, **Event Stream Analysis**) and click 

.

   The **Install Services** dialog is displayed.

   d. Select the appropriate host type in **Category** and click **Install**.

Arguments

Argument	Description
`--help-install-opts`	Display all the arguments in this table.

Argument	Description
`--eula`	Accept or decline the End User License Agreement (EULA). Specify: • `true` (default) to accept the agreement • `false` to decline it and cancel the installation. For example: `--eula=true`
`--is-head`	Designate the host as the NW Server host or a component host. Specify: • `true` for NW Server host. • `false` for Component host. For example: `--is-head=true`
`--host-name`	Specify new hostname. If you do not specify this argument, NetWitness Platform retains the existing hostname. For example: `--host-name=<hostname>`
`--master-pass`	Enter master password. For example: `--master-pass=<password>`
`--deploy-pass`	Enter deployment password. For example: `--deploy-pass=<password>`
`--iface-name`	Specify network interface. For example: `--iface-name=eth0`
`--ip-override`	Accept or override IP address found for this host or change the IP configuration found on the host. Specify: • `true` provide IP address. • `false` use IP address found on the host. For example: `--ip-override=false`
`--ip-type`	Select ip address configuration type. Specify: • `1` Static IP Configuration) • `2` DCHP For example: `--ip-type=1`
`--ip-addr`	For Static IP configuration, enter IP Address for static address. For example: `--ip-addr=<ip-address>`
`--ip-netmask`	For Static IP configuration, enter Subnet Mask for static address. For example: `--ip-gateway=<subnet-mask>`

Argument	Description
`--ip-gateway`	For Static IP configuration, enter default gateway for static address. For example: `--ip-gateway=<default-gateway>`
`--ip-nameserver`	IP address assigned to DNS server. `--ip-nameserver=<ip-address>`
`--ip-nameserver-secondary`	Optional - IP address assigned to a secondary DNS server. For example: `--ip-nameserver-secondary=<ip-address>`
`--ip-domain`	For Static IP configuration, enter Local Domain Name for static address. For example: `--ip-domain=<default-gateway>`
`--repo-type`	Select type of update repository. Specify: • `1` Local repository • `2` External repository For example: `--repo-type=1`
`--repo-url`	For an external update repository, specify the url of the repository. For example: `--repo-url=<url>`
`--head-ip`	For a component host, specify IP Address of the NW Server. For example: `--head-ip=<ip-address>`
`--custom-firewall`	Disable default firewall configuration and use your custom configuration. Specify: • `true` use custom firewall configuration. • `false` use default firewall configuration. For example: `--custom-firewall=true`
`--use-nat`	Configure the host to use Network Address Translation (NAT) based IP addresses: • `true` use NAT IPs to connect to other hosts • `false` do not use NAT IPs to connect to other hosts (default) For example: `--use-nat=false`

Appendix C. Virtual Host Recommended System Requirements

The following tables list the vCPU, vRAM, and Read and Write IOPS recommended requirements for the virtual hosts based on the EPS or capture rate for each component.

- Storage allocation is covered in Step 3 “Configure Databases to Accommodate NetWitness”.
- vRAM and vCPU recommendations may vary depending on capture rates, configuration and content enabled.
- The recommendations were tested at ingest rates of up to 25,000 EPS for logs and two Gbps for packets, for non SSL.
- The vCPU specifications for all the components listed in the following tables are Intel Xeon CPU @2.59 Ghz.
- All ports are SSL tested at 15,000 EPS for logs and 1.5 Gbps for packets.

Note: The above recommended values might differ for 11.7.0.0 installation when you install and try the new features and enhancements.

Scenario One

The requirements in these tables were calculated under the following conditions.

- All the components were integrated.
- The Log stream included a Log Decoder, Concentrator, and Archiver.
- The Packet Stream included a Network Decoder and Concentrator.
- Additional Packet Stream included a Network Hybrid with query load.
- The background load included hourly and daily reports.
- Charts were configured.

Note: Intel x86 64-bit chip architecture is 2.599 GHz or greater speed per core.

Log Decoder

EPS	CPU	Memory	Read IOPS	Write IOPS
2,500	6 cores	32 GB	50	75
5,000	8 cores	32 GB	100	100
7,500	10 cores	32 GB	150	150

Network Decoder

Mbps	CPU	Memory	Read IOPS	Write IOPS
50	4 cores	32 GB	50	150
100	4 cores	32 GB	50	250
250	4 cores	32 GB	50	350

Concentrator - Log Stream

EPS	CPU	Memory	Read IOPS	Write IOPS
2,500	4 cores	32 GB	300	1,800
5,000	4 cores	32 GB	400	2,350
7,500	6 cores	32 GB	500	4,500

Concentrator - Packet Stream

Mbps	CPU	Memory	Read IOPS	Write IOPS
50	4 cores	32 GB	50	1,350
100	4 cores	32 GB	100	1,700
250	4 cores	32 GB	150	2,100

Archiver

EPS	CPU	Memory	Read IOPS	Write IOPS
2,500	4 cores	32 GB	150	250
5,000	4 cores	32 GB	150	250
7,500	6 cores	32 GB	150	350

Note: NetWitness recommends using Virtual Machine as a hybrid only for lower EPS rates. In case of high query load or high EPS, consider using Physical Appliance.

(For version 11.7.1 and Later) Log Hybrid

Rate (EPS)	vCPU	vRAM	Total IOPS	Read IOPS	Write IOPS
2500	10 Cores	48	2325	450 (Concentrator 400, Decoder 50)	1875 (Concentrator 1800, Decoder 75)
5000	12 Cores	64	3100	650 (Concentrator 500, Decoder 100)	2450 (Concentrator 2350, Decoder 100)

(For version 11.7.1 and Later) Network Hybrid

Mbps	CPU	Memory	Read IOPS	Write IOPS
50	8 cores	48 GB	350 (Concentrator 300, Decoder 50)	1650 (Concentrator 1500, Decoder 150)
100	8 cores	64 GB	550 (Concentrator 500, Decoder 50)	1950 (Concentrator 1700, Decoder 250)

Mbps	CPU	Memory	Read IOPS	Write IOPS
250	8 cores	64 GB	850 (Concentrator 800, Decoder 50)	2450 (Concentrator 2100, Decoder 350)

Scenario Two

The requirements in these tables were calculated under the following conditions.

- All the components were integrated.
- The Log stream included a Log Decoder, Concentrator, Warehouse Connector, and Archiver.
- The Packet Stream included a Network Decoder, Concentrator, and Warehouse Connector.
- Event Stream Analysis was aggregating at 90K EPS from three Hybrid Concentrators.
- Respond was receiving alerts from the Reporting Engine and Event Stream Analysis.
- The background load Included reports, charts, alerts, investigation, and Respond.
- Alerts were configured.

Log Decoder

EPS	CPU	Memory	Read IOPS	Write IOPS
10,000	16 cores	50 GB	300	50
15,000	20 cores	60 GB	550	100

Network Decoder

Mbps	CPU	Memory	Read IOPS	Write IOPS
500	8 cores	40 GB	150	200
1,000	12 cores	50 GB	200	400
1,500	16 cores	75 GB	200	500

Concentrator - Log Stream

EPS	CPU	Memory	Read IOPS	Write IOPS
10,000	10 cores	50 GB	1,550 + 50	6,500
15,000	12 cores	60 GB	1,200 + 400	7,600

Concentrator - Packet Stream

Mbps	CPU	Memory	Read IOPS	Write IOPS
500	12 cores	50 GB	250	4,600
1,000	16 cores	50 GB	550	5,500
1,500	24 cores	75 GB	1,050	6,500

Warehouse Connector - Log Stream

EPS	CPU	Memory	Read IOPS	Write IOPS
10,000	8 cores	30 GB	50	50
15,000	10 cores	35 GB	50	50

Warehouse Connector - Packet Stream

Mbps	CPU	Memory	Read IOPS	Write IOPS
500	6 cores	32 GB	50	50
1,000	6 cores	32 GB	50	50
1,500	8 cores	40 GB	50	50

Archiver - Log Stream

EPS	CPU	Memory	Read IOPS	Write IOPS
10,000	12 cores	40 GB	1,300	700
15,000	14 cores	45 GB	1,200	900

ESA Correlation service with Context Hub

EPS	CPU	Memory	Read IOPS	Write IOPS
90,000	32 cores	250 GB	50	50

New Health and Wellness

Minimum memory for a standalone virtual host is 16 GB.

Each NetWitness platform host writes 150 MB of Health and Wellness Metrics data into Elasticsearch data per day. For example, if you have 45 NetWitness Platform hosts then 6.6 GB of metrics data is written to Elasticsearch per day.

CPU	Memory
4 cores	16 GB

NetWitness Server and Co-Located Components

The NetWitness Server, Jetty, Broker, Respond, and Reporting Engine are in the same location.

CPU	Memory	Read IOPS	Write IOPS
12 cores	64 GB	100	350

Analyst UI

The NetWitness UI and the Broker, Investigate, Respond, and Reporting Engine services are in the same location.

CPU	Memory	Read IOPS	Write IOPS
8 cores	32 GB	100	350

Scenario Three

The requirements in these tables were calculated under the following conditions.

- All the components were integrated.
- The Log stream included a Log Decoder and Concentrator.
- The Packet stream included a Network Decoder and the Concentrator.
- Event Stream Analysis was aggregating at 90K EPS from three Hybrid Concentrators.
- Respond was receiving alerts from the Reporting Engine and Event Stream Analysis.
- The background load included hourly and daily reports.
- Charts were configured.

Log Decoder

EPS	CPU	Memory	Read IOPS	Write IOPS
25,000	32 cores	75 GB	1050	150

Network Decoder

Mbps	CPU	Memory	Read IOPS	Write IOPS
2,000	16 cores	75 GB	300	650

Concentrator - Log Stream

EPS	CPU	Memory	Read IOPS	Write IOPS
25,000	16 cores	75 GB	1,200 + 400	9,200

Concentrator - Packet Stream

Mbps	CPU	Memory	Read IOPS	Write IOPS
2,000	24 cores	75 GB	1250	7,050

Log Collector (Local and Remote)

The Remote Log Collector is a Log Collector service running on a remote host and the Remote Collector is deployed virtually.

EPS	CPU	Memory	Read IOPS	Write IOPS
15,000	8 cores	8 GB	50	50
30,000	8 cores	15 GB	100	100

Scenario Four

The requirements in these tables were calculated under the following conditions for Endpoint Log Hybrid.

- All the components were integrated.
- Endpoint Server is installed.
- The Log stream included a Log Decoder and Concentrator.

Endpoint Log Hybrid

The values provided below are qualified for NetWitness 11.7 for a dedicated Endpoint Log Hybrid with no other log sources configured.

Agents	CPU	Memory	IOPS Values		
<= 5K	16 core	32 GB		**Read IOPS**	**Write IOPS**
			Log Decoder	250	150
			Concentrator	150	7,050
			MongoDb	250	150

Agents	CPU	Memory	IOPS Values		
> 5K <= 15K	16 core	64 GB		**Read IOPS**	**Write IOPS**
			Log Decoder	250	150
			Concentrator	150	7,050
			MongoDb	250	150

Agents	CPU	Memory	IOPS Values		
> 15K <= 50K	24 core	128 GB		**Read IOPS**	**Write IOPS**
			Log Decoder	250	150
			Concentrator	150	7,050
			MongoDb	250	150

If you have more than 20K agents in your virtual deployment, NetWitness recommends you to do one of the following:

- Scale resources such as CPU, RAM, and storage
- Install a physical host (Series 5 Endpoint Log Hybrid)

For details on disk usage, see the Prepare Virtual or Cloud Storage topic in the *Storage Guide for NetWitness Platform 11.x*.

Endpoint Broker

Agents	CPU	RAM
50000	2%	4 GB

Log Collector (Local and Remote)

The Remote Log Collector is a Log Collector service running on a remote host and the Remote Collector is deployed virtually.

EPS	CPU	Memory	Read IOPS	Write IOPS
15,000	8 cores	8 GB	50	50
30,000	8 cores	15 GB	100	100

Legacy Windows Collectors Sizing Guidelines

Refer to the *Legacy Windows Collection Update & Installation* for sizing guidelines for the Legacy Windows Collector.

UEBA

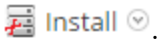
CPU	Memory	Read IOPS	Write IOPS
16 cores	64 GB	500MB	500MB

Note: NetWitness recommends that you only deploy UEBA on a virtual host if your log collection volume is low. If you have a moderate to high log collection volume, NetWitness recommends that you deploy UEBA on the physical host described under "NetWitness UEBA Host Hardware Specifications" in the Physical Host Installation Guide. Contact Customer Support for advice on choosing which host, virtual or physical, to use for UEBA.

Appendix D. Update the Virtual ESA Host Memory

ESA current memory is allocated to 65% of the available memory on the host. (For example, with 128 GB available memory, ESA memory will be 81 GB.)

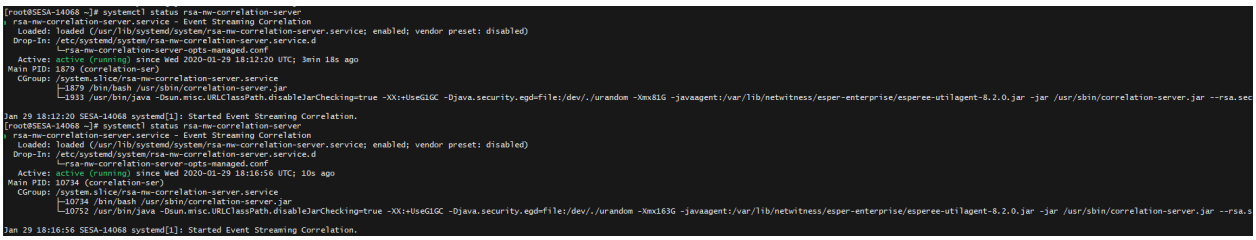
To Update the Memory of the Virtual ESA Host:

1. Power down the virtual machine host and update the virtual host memory from x GB to y GB. (Example: x = 128 GB and y = 256 GB).
2. Power on the virtual machine host.
3. Log in to NetWitness Platform and go to  **(Admin) > Hosts**.
4. Select the ESA host where the memory is updated and click .
   The Install Services dialog is displayed.
5. Select ESA Primary or ESA Secondary on the host, depending on the ESA host category, and click **Install**.
   After the installation completes, the memory settings update automatically.

To Check ESA Memory:

On your ESA host, run the following command:

```
systemctl status rsa-nw-correlation-server
```



NetWitness® Platform XDR

Version 12.0

Warehouse Connector Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

Warehouse Connector Configuration **4**
How Warehouse Connector Works 4
Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid **7**
Configure a Warehouse Connector Service **8**
Configure the Data Source for Warehouse Connector **9**
Configure the Data Source 9
Update the Port Number and SSL Settings of the Data Source 10
Configure the Destination **12**
Configure the Destination Using NFS 14
Configure the Destination Using SFTP 17
Configure the Destination Using WebHDFS 24
Configure a Stream **28**
Create a Stream 28
Finalize a Stream 30
Start a Stream 30
Monitor a Warehouse Connector **32**
Add Warehouse as a Data Source to Reporting Engine **34**
Analyze a Warehouse Report **35**
View the Warehouse Connector Service **36**
Troubleshoot the Warehouse Connector **37**
Manage a Stream **39**
Manage a Lockbox **45**
Warehouse Connector Configuration References **48**
General Tab Settings 49
Appliance Service Configuration Tab Settings 52
Sources and Destinations Configuration 55
Add Stream Dialog 58
Streams Configuration 61
Lockbox Settings 68

Warehouse Connector Configuration

How Warehouse Connector Works

Warehouse Connector collects meta and events from Decoder and Log Decoder and writes them in AVRO format into a Hadoop-based distributed computing system. You can set up the Warehouse Connector as a service on existing Log Decoders or Decoders.

The Warehouse Connector contains the following components:

- Data Source
- Destination
- Data Stream

Data Source

Warehouse Connector collects data from the data source to store it in the destination. The supported data sources are Log Decoder and Decoder.

Destination

Destination is the Hadoop-based distributed computing system that collects, manages, and enables reporting on security data. The following are the supported destinations:

- NetWitness Warehouse (MapR) deployments
- HortonWorks Data Platform
- Any Hadoop-based distributed computing system that supports WebHDFS or NFS mounting of HDFS file systems.
  Example: Commercial MapR M5 Enterprise Edition for Apache Hadoop

Data Streams

A data stream is a logical connection between the data source and destination. You can have multiple streams for different subsets of data collected. You can setup streams to segregate data from multiple Decoder and Log Decoder services. You can create a stream with a single data source and destination or with multiple data sources and a single destination.

The Warehouse Connector:

- Aggregates session and raw log data from Decoders and Log Decoders.
- Transfers the aggregated data to supported destinations like Hadoop based deployments.
- Serializes the aggregated data that includes both schema and data into AVRO format.

Meta Filters

Meta filters enables you to filter the meta keys that should be written into the Warehouse. For more information, see Specify Meta Filters for a Stream.

Multi-Valued Meta Keys

NetWitness Warehouse supports multi-valued meta keys. The multi-valued meta keys is the meta field with the array type. You can use the meta keys library to determine the meta fields of type array and write HIVE queries with the correct syntax for arrays. By default, the following meta keys are treated as multi-valued and are defined in the file, **multivalue-bootstrap.xml** located at **/etc/netwitness/ng** in the Warehouse Connector:

- alias.host
- action
- username
- alias.ip
- alias.ipv6
- email
- device.group
- event.class

Checksum Validation

You can validate the file integrity of the AVRO files that are transferred from the Warehouse Connector to the data destinations. You need to enable checksum validation option when you configure the Warehouse Connector.

Lockbox Support

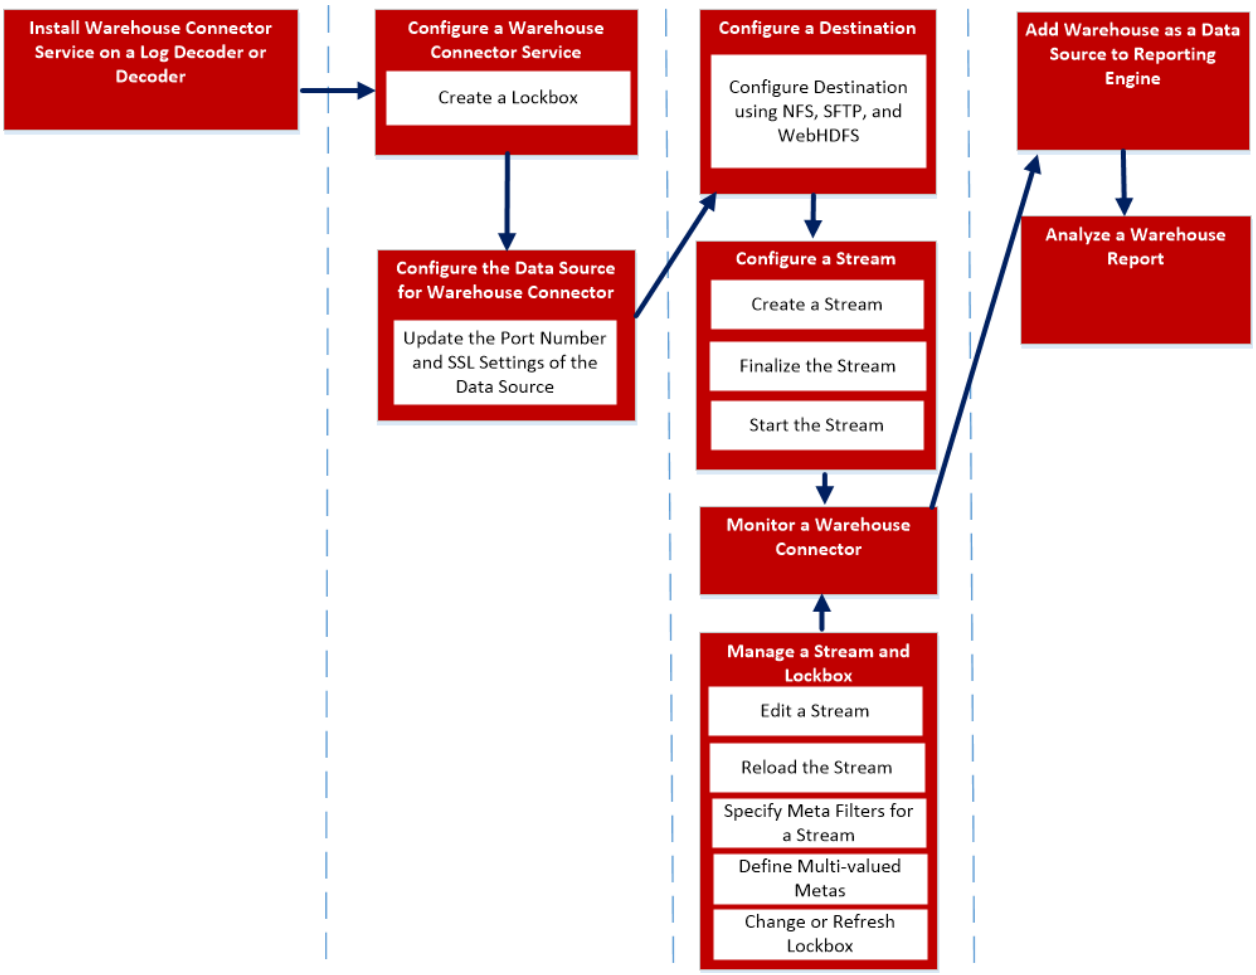
Lockbox provides an encrypted file that Warehouse Connector uses to store and protect sensitive data. You need to create the lockbox by providing a lockbox password while configuring the Warehouse Connector for the first time.

Note: You can orchestrate Warehouse Connector on a Log Decoder or a Decoder appliance.

The following is an overview on how to install and configure the Warehouse Connector service on Log Decoder or Decoder.

- Configuring the Warehouse Connector service on NetWitness,
- Configuring data sources,
- Destinations,
- Streams for Warehouse Connector,
- Configuring alert notifications on NetWitness.

Note: RSA NetWitness Platform has added a Health & Wellness stat for Warehouse Connector to indicate the status of its Lockbox. Also, an out-of-the-box rule has been added so that a Health & Wellness alarm is raised when the Lockbox does not exist or cannot be opened.



To install and configure the Warehouse Connector service, perform the following:

1. Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid
2. Configure a Warehouse Connector Service
3. Configure the Data Source for Warehouse Connector
4. Configure the Destination
5. Configure a Stream
6. Monitor a Warehouse Connector
7. Add Warehouse as a Data Source to Reporting Engine
8. Analyze a Warehouse Report
9. Manage a Stream

Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid

To install (fresh install) the Warehouse Connector service on a Log Decoder or Decoder or Hybrid:

1. Log in to NW Admin server
2. Enter the following command on NetWitness Server:
   `warehouse-installer --help`
   The command line interface (CLI) usage descriptions are displayed.
3. Install Warehouse Connector service by executing the following command:
   `warehouse-installer --host-key <ID, IP address, hostname or display name of Decoder, Log Decoder, or Hybrid host>`
4. Restart the Warehouse Connector using following command:
   `service nwwarehouseconnector restart`

The Warehouse Connector service is successfully installed on the Log Decoder or Decoder or Hybrid.

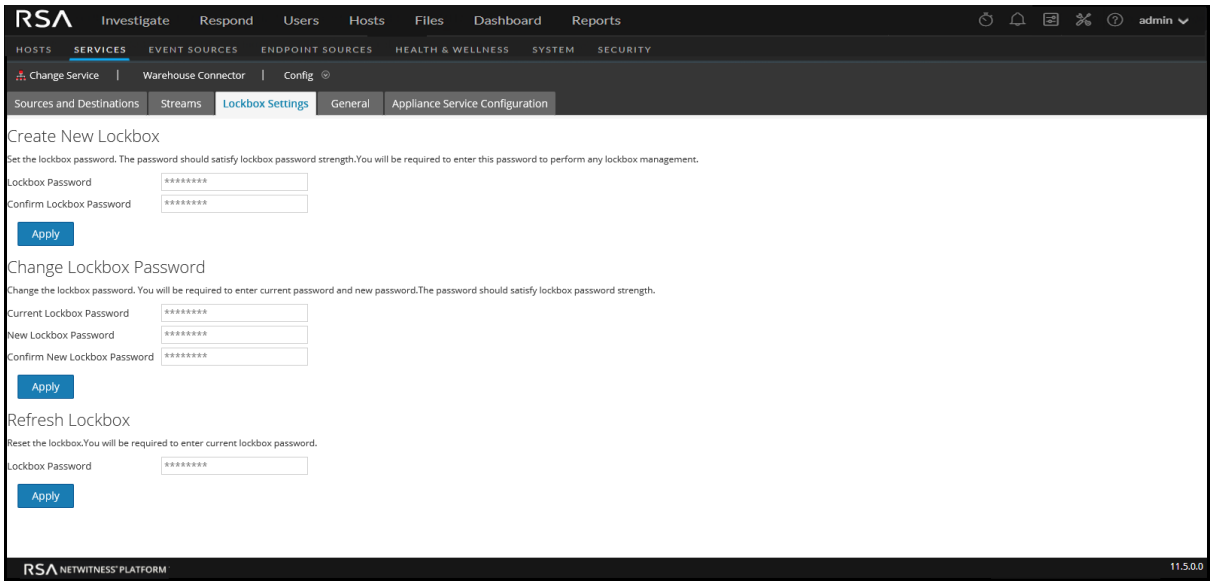
Note: If you have custom-defined host entries in `/etc/hosts`, make sure to have the same entires in `/etc/hosts.user`. Run the following command to ensure that the host names are updated in this installation:
`nw-manage --refresh-host --host-key <ID, IP, hostname or display name of host>`
For more information, see "Manage Custom Host Entries in /etc/hosts" in the *System Maintenance Guide*.

Configure a Warehouse Connector Service

You can configure the Warehouse Connector service using the following procedure.

To set the Lockbox password:

1. Log on to NetWitness.
2. Go to **(Admin) > Services**.
3. In the Services view, select the added Warehouse Connector service, and select **> View > Config**.
4. In the Services Config view of Warehouse Connector, click the **Lockbox Settings** tab.



5. In the **Create New Lockbox** section, perform the following:
   a. In the **Lockbox Password** field, enter the new lockbox password.

      Note: The lockbox password must be at least eight characters in length and it must contain at least three of the following groups: one uppercase character [A-Z], one lowercase character [a-z], one numeral [0-9], and one special character.

   b. In the **Confirm Lockbox Password** field, enter the added lockbox password to confirm.
   c. Click **Apply**.

   The Lockbox password is set.

Configure the Data Source for Warehouse Connector

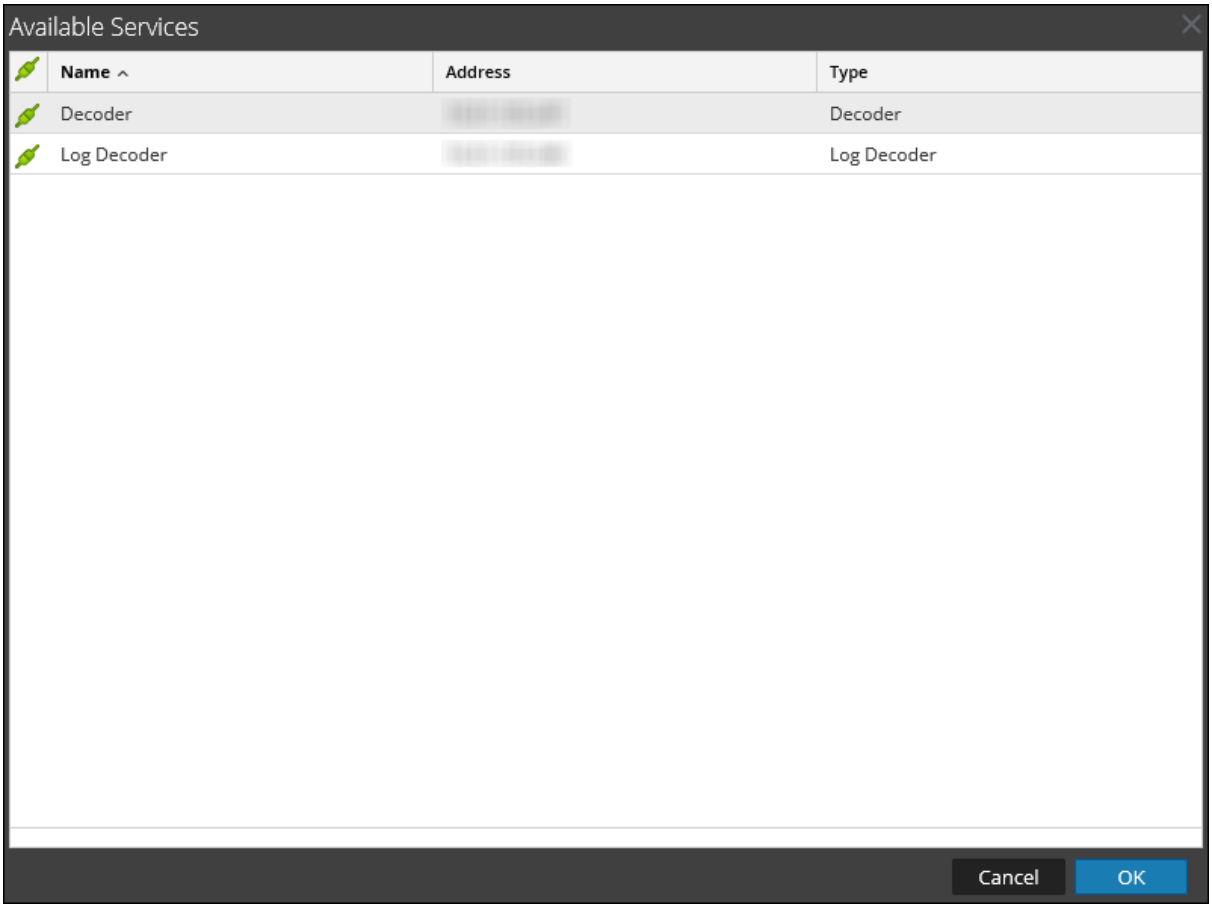
There are two procedures for configuring the data source for the Warehouse Connector:

1. First, you must configure the data source.
2. Then, you need to update the Port Number and SSL Settings for the Data Source.

Configure the Data Source

To configure the data source:

1. Log on to NetWitness.
2. Go to **(Admin) > Services**.
3. In the Services view, select the added Warehouse Connector service, and select **> View > Config**.
   The Services Config view of Warehouse Connector is displayed.
4. On the **Sources and Destinations** tab, in the **Source Configuration** section, click .



5. In the **Available Services** dialog, select the Log Decoder or Decoder services that you want to add as a source to the Warehouse Connector service and click **OK**.
   The selected Log Decoder and Decoder services are listed in the **Source Configuration** section.

Update the Port Number and SSL Settings of the Data Source

If there is change to the port number or SSL settings of the Warehouse Connect data sources, you can update the changes in Explore view of the Warehouse Connector.

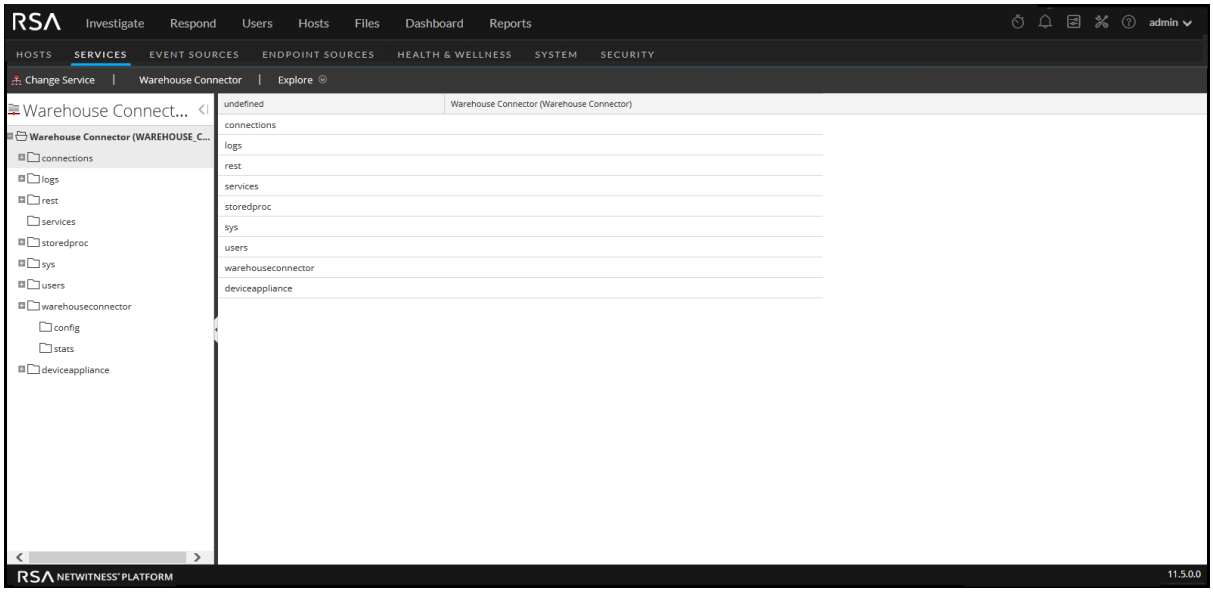
Make sure that:

- You have the changed port number or SSL settings of the data source.
- You have stopped the streams related to the data source that is updated with the port number or SSL settings.

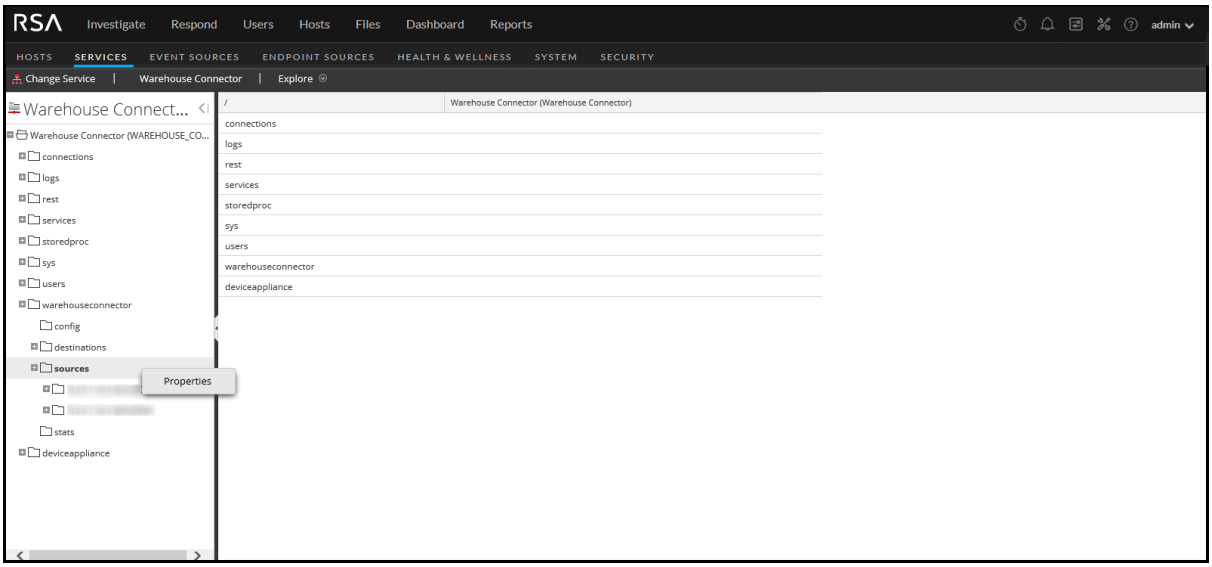
To update the port number or SSL settings:

1. Log on to NetWitness.
2. Go to  **(Admin) > Services**.
3. In the Services view, select the added Warehouse Connector service and select  **> View > Explore**.

   The Service Explore view of Warehouse Connector is displayed.



4. Navigate to **warehouseconnector/sources**, right-click the source, and click **Properties**.
   The Properties of the source is displayed.



5. In the drop-down menu, select **update**. In the Parameters field, perform the following:

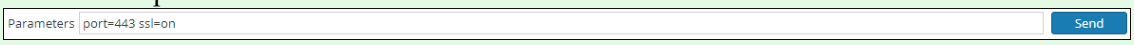
   - To update the port number of the source, enter `port=<new_source_portnumber>` and click **Send**.

   - To update the SSL settings of the source, enter `ssl=<new_ssl_settings>` and click **Send**.

     Note: You can also update the port number and ssl settings simultaneous by adding space between the parameters.

6. Restart the Warehouse Connector service.
7. Start the streams.

Configure the Destination

You can configure the destination using NFS, SFTP, and WebHDFS. Change the destination to which the Warehouse Connector service needs to write the collected data using NFS:

- NetWitness Warehouse (MapR) deployments
- Commercial MapR M5 Enterprise Edition for Apache Hadoop deployments

You can configure the Warehouse Connector to write to a remote destination using Secure File Transfer Protocol (SFTP). The remote destination can be a remote server that is NFS mounted to the MapR cluster or it can be a remote staging server.

By default, in the remote destination the Warehouse Connector writes data in the following directory structure:

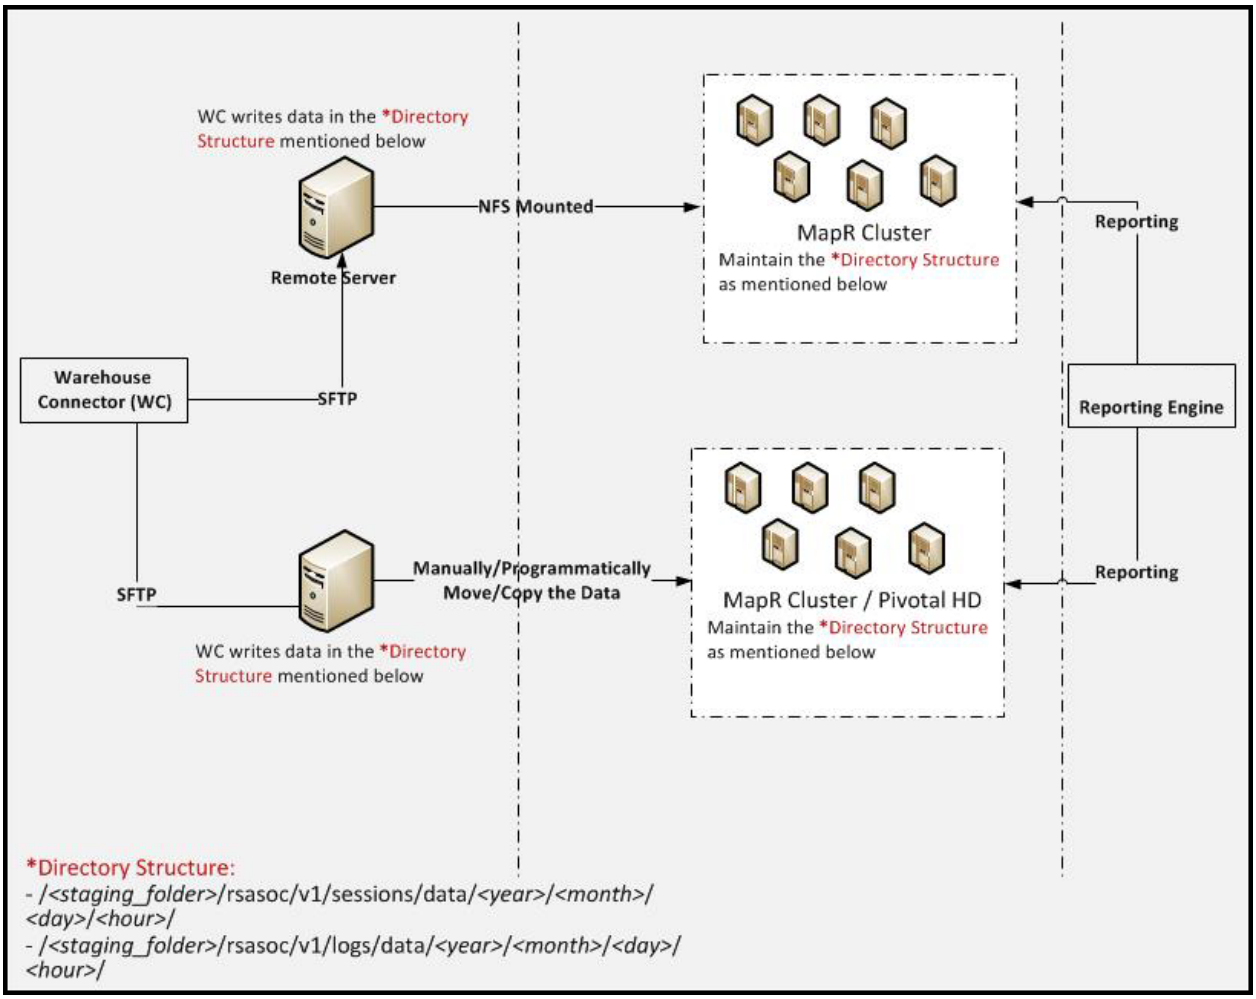
- `/<staging_folder>/rsasoc/v1/sessions/data/<year>/<month>/<day>/<hour>/`
- `/<staging_folder>/rsasoc/v1/logs/data/<year>/<month>/<day>/<hour>/`
  Where <staging_folder> is the folder on the remote server where the Warehouse Connector writes the data.

If you are using a remote staging server as the remote destination, you need to manually copy or move the directory structure to any of the following deployments:

- NetWitness Warehouse (MapR)
- Commercial MapR M5 Enterprise Edition for Apache Hadoop
- HortonWorks HD

To generate reports from the data written by Warehouse Connector, make sure that in your Hadoop deployment you maintain a similar directory structure that is created by Warehouse Connector in the remote destinations.

The following illustration describes how you can use SFTP to write data from Warehouse Connector to a remote destination.



You can configure the Warehouse Connector service to write the collected data to a Hadoop-based distributed computing system that supports WebHDFS.

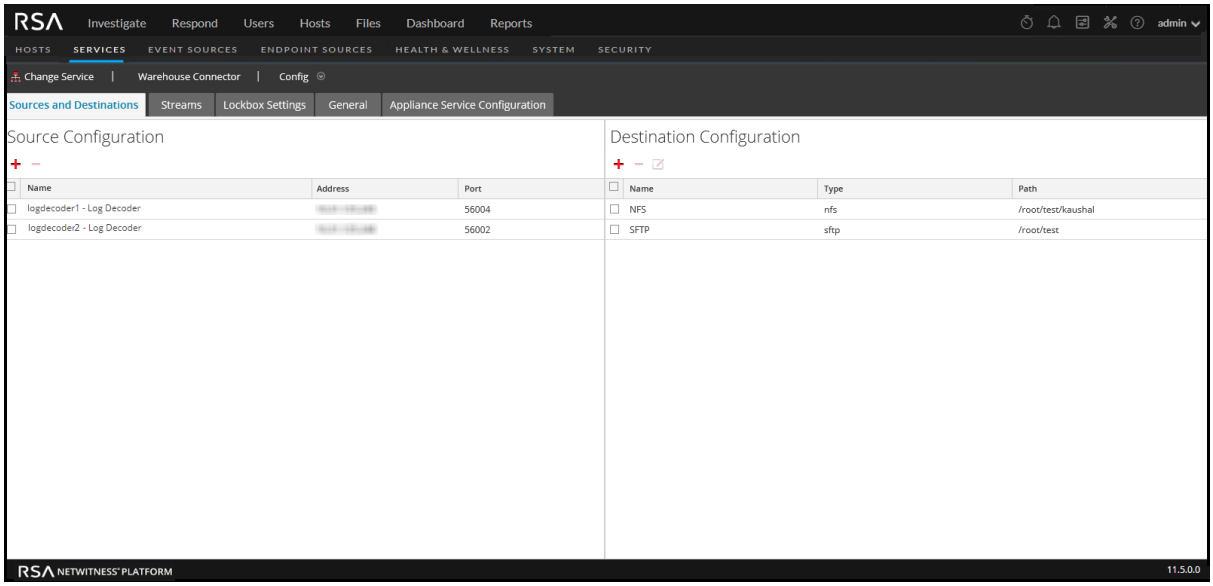
Configure the Destination Using NFS

Make sure that you have:

- Installed the Warehouse Connector service or virtual appliance in your network environment.
- Added the Warehouse Connector service to NetWitness. For more information, see "Add a Service to a Host" in the *Hosts and Services Getting Started Guide*.
- Set up NFS on Warehouse Connector. For more information on how to set up NFS on Warehouse Connector, see "Configure Warehouse Connector to Write to Warehouse" in the *Warehouse (MapR) Configuration Guide*.

To configure the destination using NFS:

1. Log on to NetWitness.
2. Go to **(Admin) > Services**.
3. In the Services view, select the Warehouse Connector service, and select **> View > Config**.
   The Services Config View of Warehouse Connector is displayed.

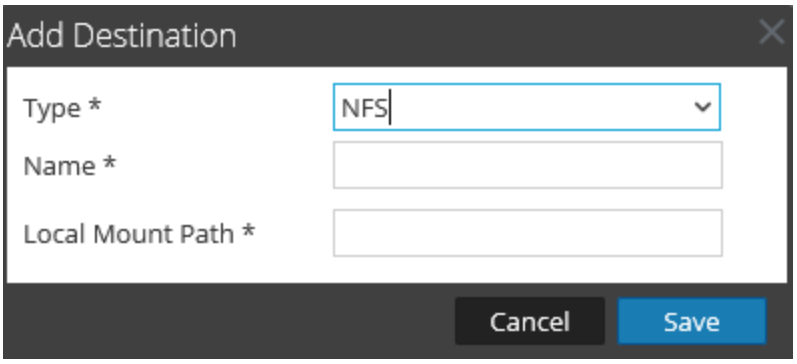


4. On the **Sources and Destinations** tab, in the **Destination Configuration** section, click .
5. In the **Add Destination** dialog, from the **Type** drop-down list, select **NFS**.
6. In the **Name** field, enter a unique symbolic name for the destination.

Note: The **Name** field does not support spaces or special characters except underscore (_).

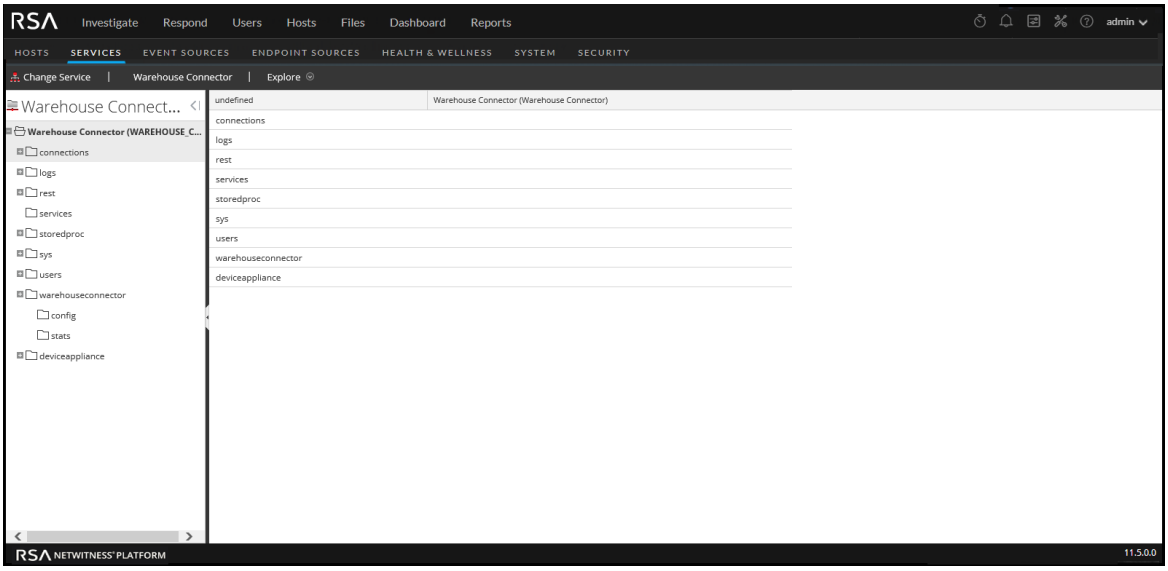
7. In the **Local Mount Path** field, enter the locally mounted directory for HDFS where you want the Warehouse Connector to write the data. For example:
   If **/saw** is the local mount point for HDFS that you have configured while mounting the mapr NFS cluster on the host where you have installed the Warehouse Connector service to write to NetWitness Warehouse (MapR), create a directory named **Ionsaw01** under **/saw** and the corresponding Local Mount Path for the destination would be **/saw/Ionsaw01**.

   For more information, see "Mount the Warehouse on the Warehouse Connector" topic in the *Warehouse (MapR) Configuration Guide*.



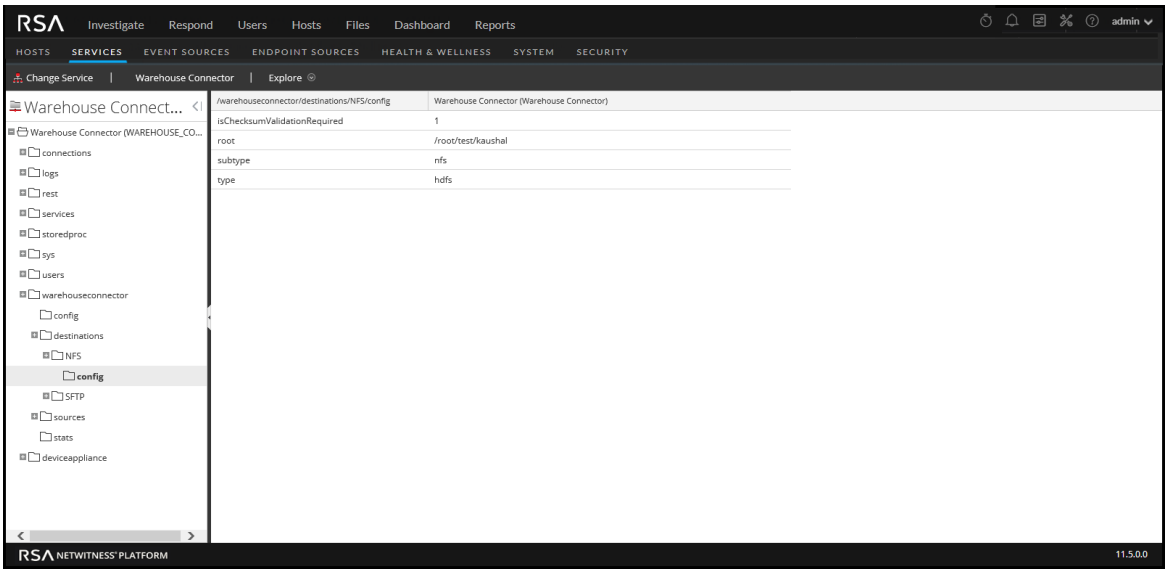
   The **/saw** mount point implies to / as the root path for HDFS. The Warehouse Connector writes the data to **/Ionsaw01** in HDFS.

8. Click **Save**.
9. (Optional) If you want to enable checksum validation, perform the following:
   a. Go to **(Admin) > Services**.
   b. In the Services view, select the added Warehouse Connector service, and select **> View > Explore**.
      The Explore view of Warehouse Connector is displayed.



c. In the options panel, navigate to **warehouseconnector/destinations/nfs/config**. This is the name of the destination and is dynamic.

d. Set the parameter `isChecksumValidationRequired` to **1**.



e. Restart the respective stream.

Configure the Destination Using SFTP

Make sure that you have:

- Installed the Warehouse Connector service or virtual appliance in your network environment.
- Added the Warehouse Connector service to NetWitness. For more information, see the "Add a Service to a Host" in the *Hosts and Services Getting Started Guide*.
- For the SFTP destination type, the destination host should be listed in the **/root/.ssh/known_hosts** file used by the ssh service (for example, sshd) running on the Warehouse Connector.

Add Destination from Warehouse Connector Host

To add the destination host to the **/root/.ssh/known_hosts** file, from the Warehouse Connector host, initiate a secure connection to the destination host:

1. Log in to the Warehouse Connector.
2. Enter `ssh root@<SAWIP>` or `ssh username@<SAWIP>`.
3. Select **Yes** and enter the password.
4. Add the host key in the **/root/.ssh/known_hosts** file

> **Note:** After you upgrade Warehouse Connector to 11.0, you must make sure that the destination host is listed in the `/root/.ssh/known_hosts` file used by the SSH service (i.e. sshd) running on the Warehouse Connector. If you do not perform this action, the streams configured with SFTP in Warehouse Connector will not start.

- If you want to use SFTP to write data into the destination using SSH key-based access, you need to configure SSH key-based access between the Warehouse Connector and the Warehouse host or Hadoop node. For more information, see **Configure SSH Keys** below.

> **Note:** If you want to enable checksum validation to validate the integrity of the AVRO files that are transferred from the Warehouse Connector to the destinations, make sure that you generate the keys without setting the passphrase and do a key exchange between Warehouse Connector and the warehouse nodes.

Configure SSH Keys

To configure SSH key-based access between the Warehouse Connector and the Warehouse host or Hadoop node:

1. Generate SSH keys on the Warehouse Connector at the default location. Perform the following:
   a. SSH to the Warehouse Connector.
   b. Type the following command and press ENTER:

      ```
      $ OWB_FORCE_FIPS_MODE_OFF=1 ssh-keygen -t ecdsa -b 521
      ```

   c. The command prompts you to enter the file in which to save the generated key.

```
Enter file in which to save the key (/root/.ssh/id_ecdsa):
```

d. Enter the file in which you want to save the key and press ENTER.

The command prompts you to enter and confirm the passphrase.

> **Note:** If you want to enable checksum validation to validate the integrity of the AVRO files that are transferred from the Warehouse Connector to the destinations, make sure that you do not set the **passphrase**. Then, the below steps e, f, g, and h are not applicable.

```
Enter passphrase (empty for no passphrase):
 Enter same passphrase again:
```

The public key is generated and is saved in the location that you provided.

> **Note:** If the SSH key is not generated in the default location (`/root/.ssh/id_ecdsa`), you need to configure the destination for warehouse connector through Explore view. For more information, see To configure the destination through Explore view:.

e. Change the directory by entering the following command:

```
cd /root/.ssh/
```

f. Move the generated key to the below location:

```
mv ~/.ssh/id_ecdsa ~/.ssh/id_ecdsa.old
```

g. Type the following command and press ENTER:

```
$ OWB_FORCE_FIPS_MODE_OFF=1 openssl pkcs8 -topk8 -v2 des3 -in id_
ecdsa.old -out id_ecdsa
```

The command prompts you to enter and confirm the passphrase.

h. Enter the encryption passphrase.

i. Run the following command to change the file permission:

```
chmod 600 ~/.ssh/id_ecdsa
```

j. Copy the generated public key to append to the remote Warehouse host or Hadoop node.

```
ssh-copy-id -i ~/.ssh/id_ecdsa.pub root@<destination host ipaddress>
```

2. SSH to remote Warehouse host or Hadoop node as "`ssh '<user>@<ip address>`", if identity key file is at default location.
or
SSH to remote Warehouse host or Hadoop node as "`ssh '<user>@<ip address> -i <identity file path>`", If identity key file is not at default location.

3. Append the generated public key to the remote Warehouse host or Hadoop node's authorized keys list located at `~/.ssh/authorized_keys`.

> **Note:** Make sure that you copy the public keys to the Hadoop node and while copying the public key ensure that you provide the login details of the user using which the WebHDFS destination would be added.

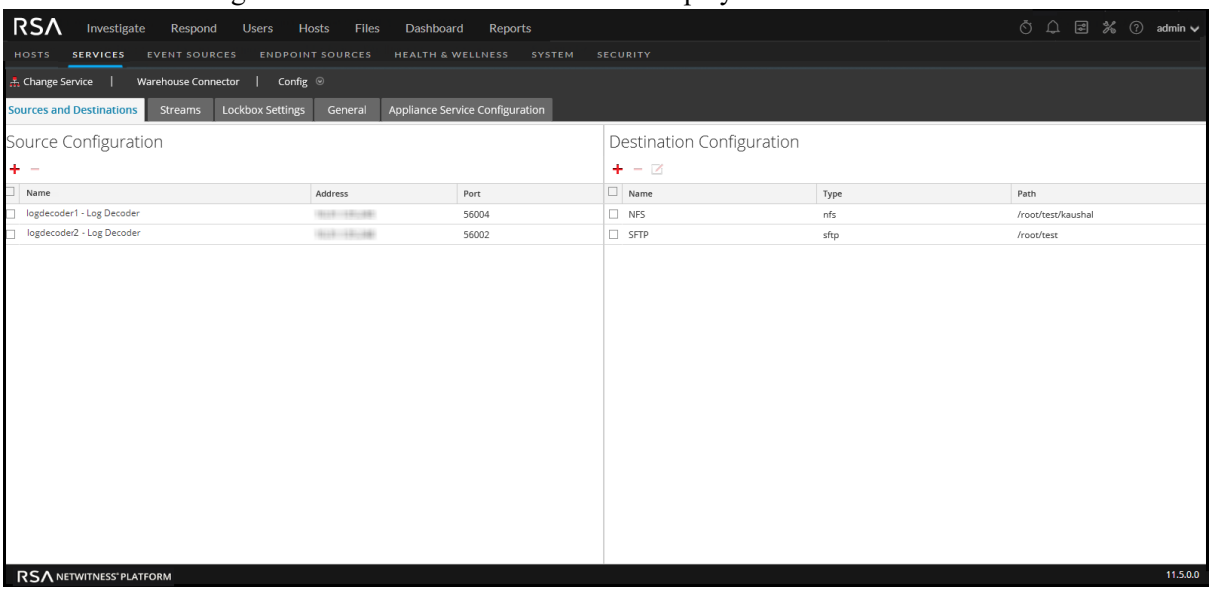
You can now securely communicate between Warehouse Connector and Warehouse nodes or Hadoop nodes.

Configure Warehouse Connector to use SFTP destination

Note: If the SSH key is not generated in the default location (`/root/.ssh/id_ecdsa`), you need to configure the destination through Explore view. For more information, see To configure the destination through Explore view:.

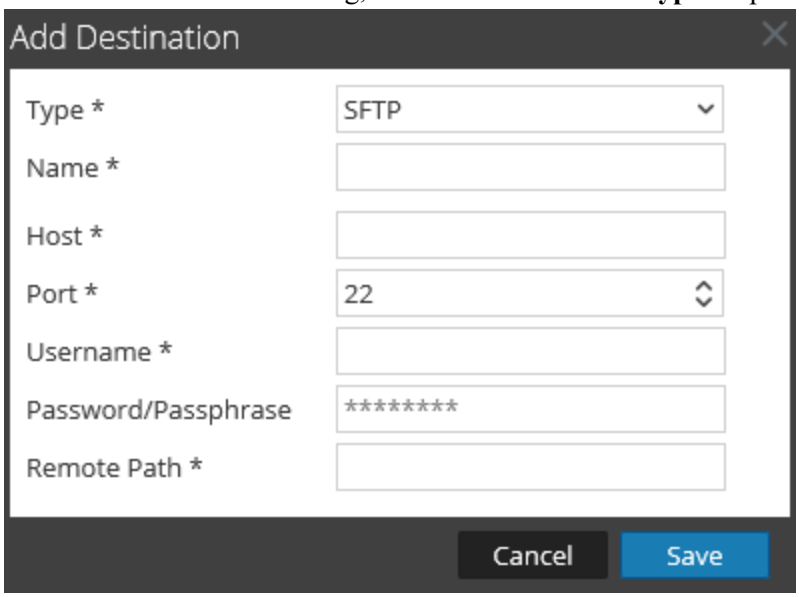
To configure the destination through User Interface:

1. Log on to NetWitness
2. Go to **(Admin) > Services**.
3. In the Services view, select the added Warehouse Connector service, and select **> View > Config**.
   The Services Config view of Warehouse Connector is displayed.



4. On the **Sources and Destinations** tab, in the **Destination Configuration** section, click .

5. In the **Add Destination** dialog, select **SFTP** from the **Type** drop-down list.



6. In the **Name** field, enter a unique symbolic name for the destination.

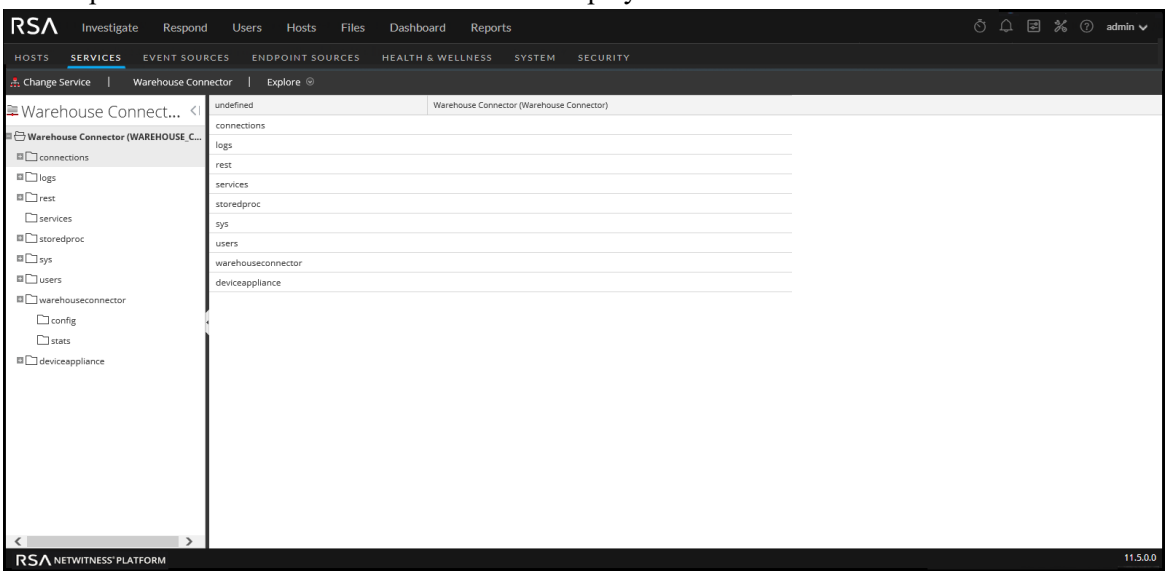
   Note: The **Name** field does not support spaces or special characters except underscore (_).

7. In the **Host** field, enter the remote server IP address.
8. In the **Port** field, retain the default port, **22**.
9. In the **Username** field, enter the SSH username.

   Note: In the case of HortonWorks HD, ensure that the username is gpadmin and for password based access the password for gpadmin should be used. For passphrase-based access, the passphrase used to generate the keys for gpadmin user should be used.

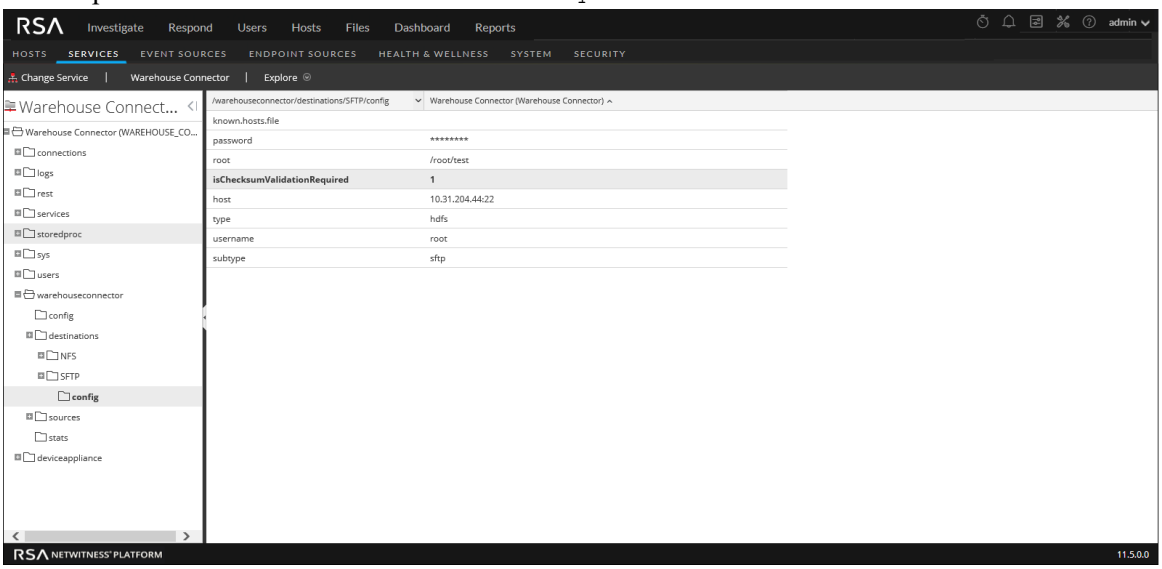
10. In the **Password/Passphrase** field, enter one of the following:
    - SSH password - If you are using SFTP to write data into the destination using password-based access.
    - SSH passphrase - If you are using SFTP to write data into the destination using SSH key-based access.
11. In the **Remote Path** field, enter the path of the directory present on the SFTP server.
12. Click **Save**.
13. (Optional) If you want to enable checksum validation, perform the following:
    a. Go to **(Admin) > Services**.
    b. In the Services view, select the added Warehouse Connector service, and select **> View > Explore**.

The Explore view of Warehouse Connector is displayed.



c. In the options panel, navigate to **warehouseconnector/destinations/sftp/config**.

d. Set the parameter `isChecksumValidationRequired` to **1**.



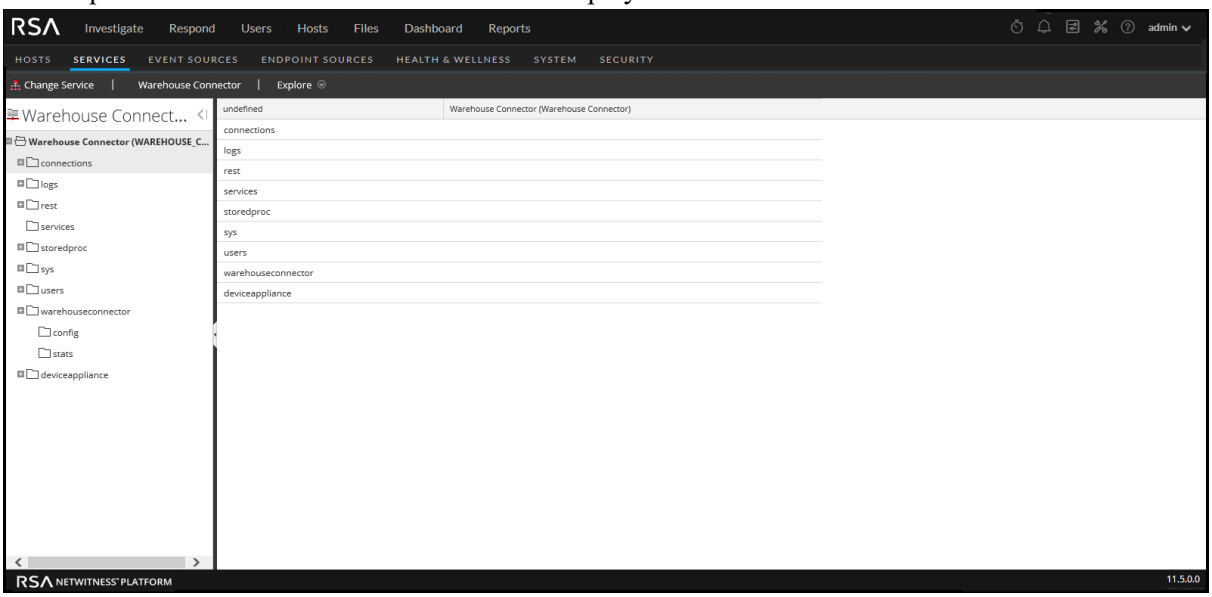
e. Restart the respective stream.

To configure the destination through Explore view:

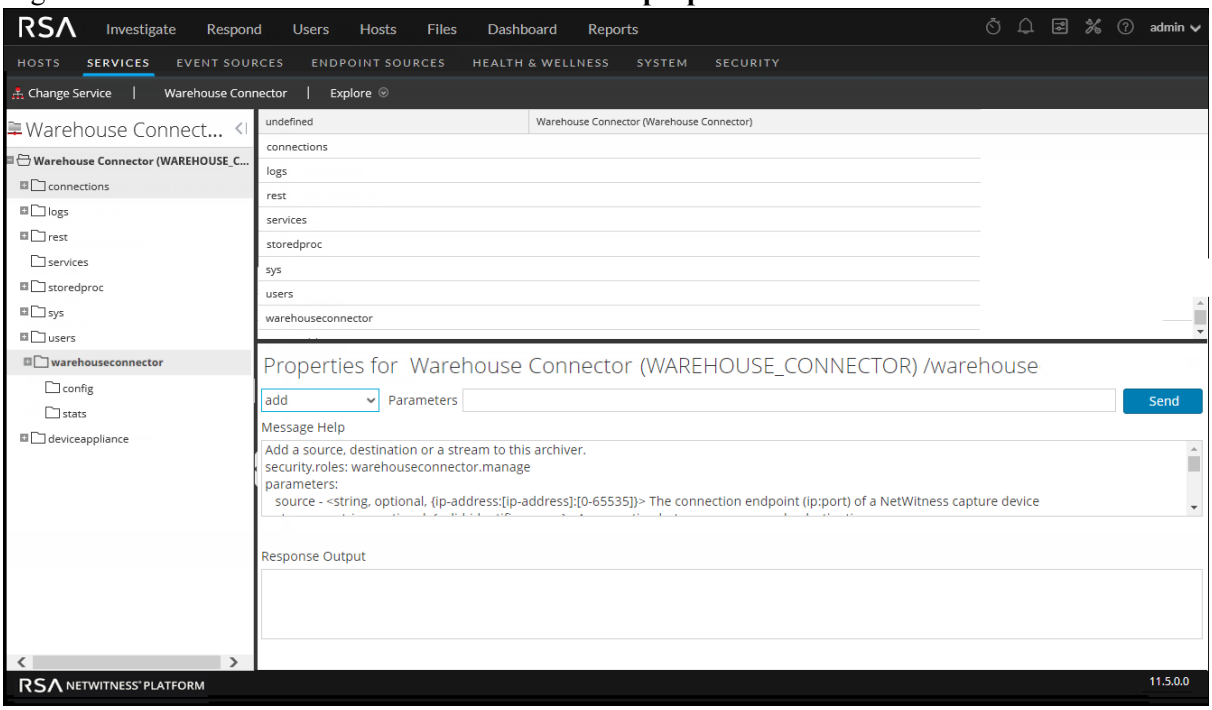
1. Go to  **(Admin) > Services**.

2. In the Services view, select the added Warehouse Connector service, and select  **> View > Explore**.

The Explore view of Warehouse Connector is displayed.



3. Right click on "warehouseconnector" node and select **properties**.



4. Select "add" property and manually enter the below config parameters.

```
name=<destination name> destination=sftp://<destination path>
host=<destination host ipaddress:port> type=hdfs port=22
username=<username> password=<password> privKeyFile=<private key file path>
```

Aggregate Metas and Raw Logs for a Log Session

To aggregate raw logs and metas from Log Decoder into a single AVRO file instead of two folders.

1. Go to **ADMIN** > **Services**.
2. Select a Warehouse Connector service and click  > **View** > **Explore**.
   The Explore view for the Warehouse Connector is displayed.
3. Open `warehouseconnector/streams/<stream name>/loader/config` and in the right pane, select the `export.logAndsession.avro.enabled` parameter.
4. Change the value to **yes**.
5. Restart the service.
6. Go to **ADMIN** > **Services**.
7. Select a Warehouse Connector service and click  > **View** > **Config**.
8. On the **Streams** tab, select the stream that you want to reload.
9. Click **Reload**.

Configure the Destination Using WebHDFS

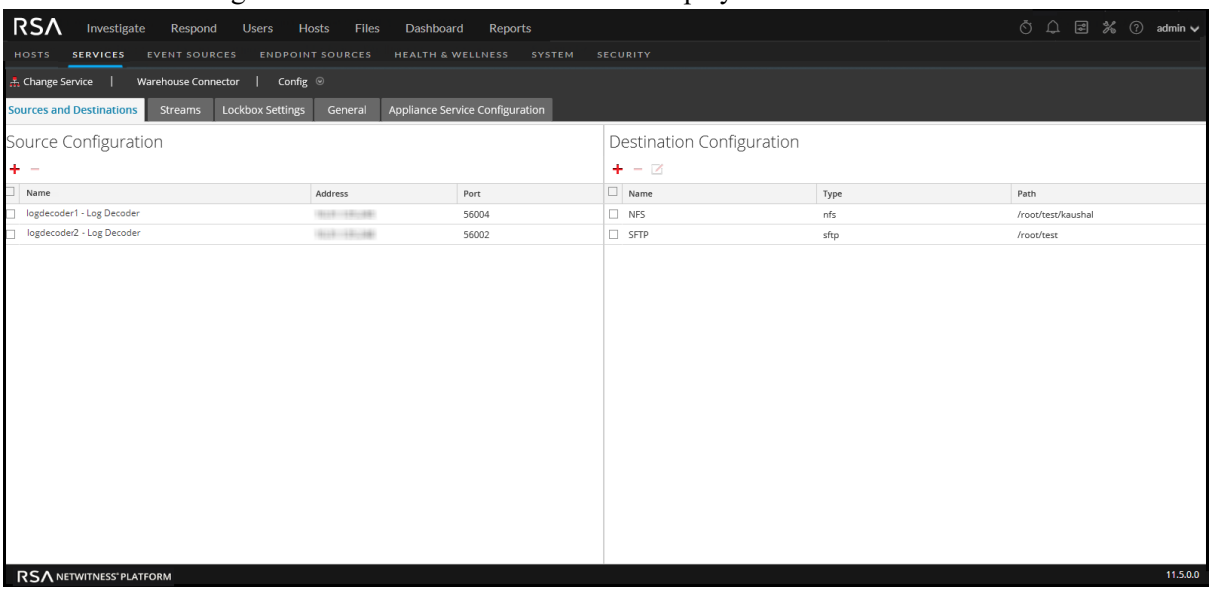
Make sure that you have:

- Installed the Warehouse Connector service or virtual appliance in your network environment.
- Added the hostname (or FQDN) and IP address of the warehouse nodes and Warehouse Connector to the DNS server. If the DNS server is not configured, add the hostname (or FQDN) and IP address of the warehouse nodes and Warehouse Connector to the file in the host on which the Warehouse Connector service is installed.
- If you want Kerberos authentication between the warehouse connector and the warehouse cluster, make sure that you perform the following:
  - Kerberos Key Distribution Center (KDC) Server is configured in your network environment and the Kerberos Keytab file is copied to the host on which you have installed Warehouse Connector.
  - Kerberos authentication is enabled in the warehouse cluster.
- If you want to enable checksum validation to validate the integrity of the AVRO files that are transferred from the Warehouse Connector to the destinations, make sure that you generate the keys without setting the passphrase and do a key exchange between the Warehouse Connector and the warehouse nodes. You need to configure SSH key-based access between the Warehouse Connector and the Warehouse host or hadoop node. For more information, see 'Configure SSH Keys' in Configure the Destination Using SFTP.

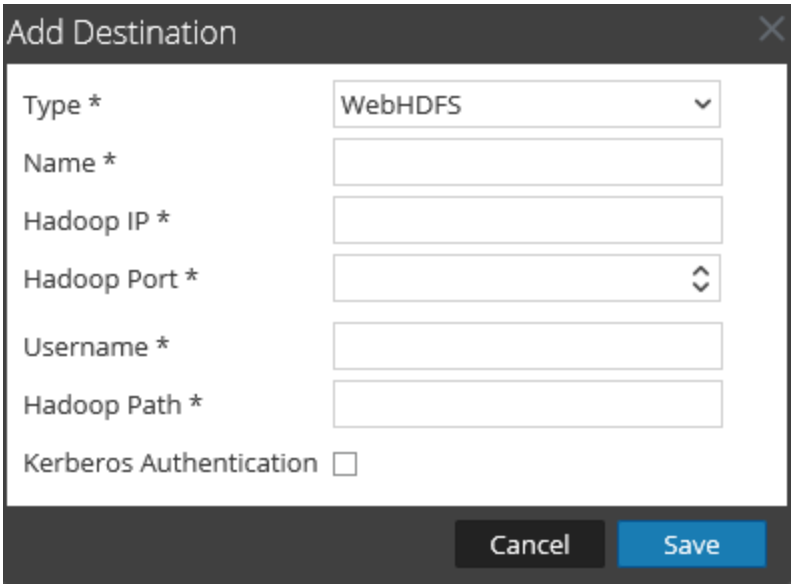
Configure Warehouse Connector to Write to SFTP destination

To configure the destination:

1. Log on to NetWitness.
2. Go to **(Admin) > Services**.
3. In the Services view, select the added Warehouse Connector service and select **> View > Config**.
   The Services Config view of Warehouse Connector is displayed.



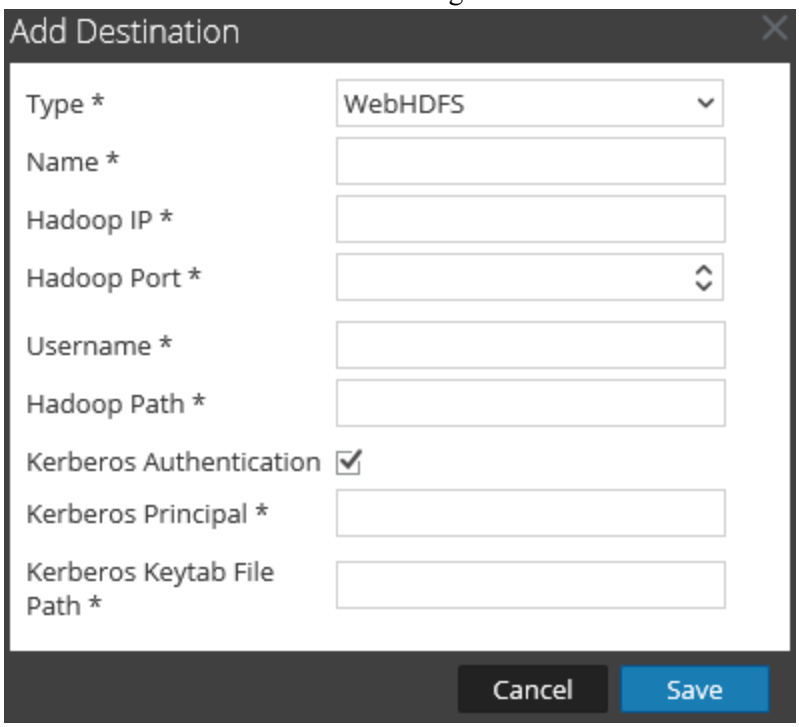
4. On the **Sources and Destinations** tab, in the **Destination Configuration** section, click .
5. In the **Add Destination** dialog, select **WebHDFS** from the drop-down list.



6. In the **Name** field, enter a unique symbolic name for the destination.

   Note: The **Name** field does not support spaces or special characters except underscore (_).

7. In the **Hadoop IP** field, enter the namenode IP address of the warehouse cluster.
8. In the **Hadoop Port** field, enter the base port that is used by the namenode web user interface.
9. In the **Username** field, enter the owner of the directory in the warehouse to which Warehouse Connector should write the data.
10. In the **Hadoop Path** field, enter the path of the directory in the warehouse to which Warehouse Connector should write the data.
11. Select the **Kerberos Authentication** checkbox, if you want the warehouse connector to securely communicate with the warehouse using Kerberos authentication.

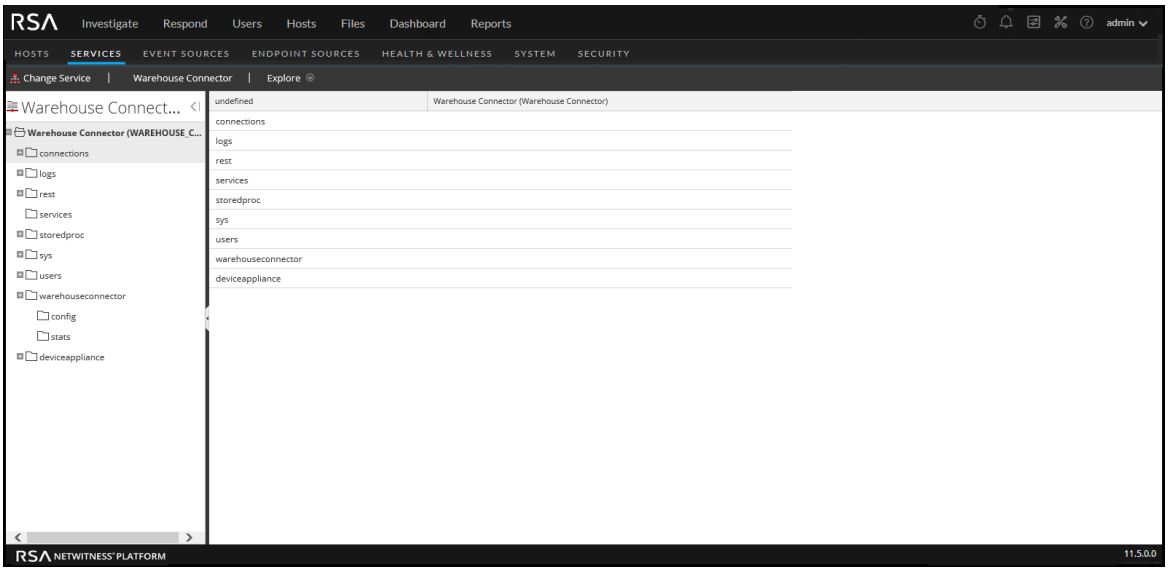


   Perform the following:

   a. In the **Kerberos Principal** field, enter the KDC Principal used for Kerberos authentication.
   b. In the **Kerberos Keytab File Path** field, enter the path of the Kerberos Keytab file in the Warehouse Connector.

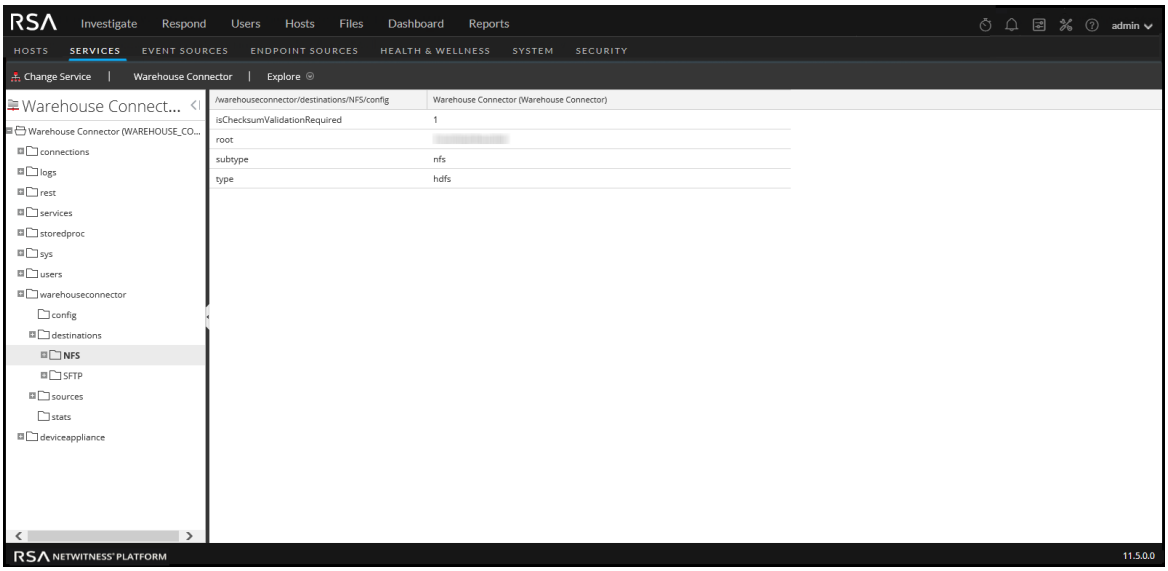
12. Click **Save**.
13. (Optional) If you want to enable checksum validation, perform the following:

   a. Go to **(Admin) > Services**.
   b. In the Services view, select the added Warehouse Connector service and select **> View > Explore**.
      The Explore view of Warehouse Connector is displayed.



c. In the options panel, navigate to **warehouseconnector/destinations/webhdfs/config**.

d. Set the parameter `isChecksumValidationRequired` to **1**.



e. Restart the respective stream.

Configure a Stream

You can configure the data stream to define the data source and destination combinations.

Make sure that you have:

- Installed the Warehouse Connector service or virtual appliance in your network environment.
- Added the Warehouse Connector service to NetWitness. For more information, see "Add a Service to a Host" in the *Hosts and Services Getting Started Guide*.
- Configured the data source from which the Warehouse Connector service needs to collect data. For more information, see Configure the Data Source for Warehouse Connector.
- Configured the destination to which the Warehouse Connector service needs to write the collected data. For more information, see Configure the Destination.

To configure the stream:

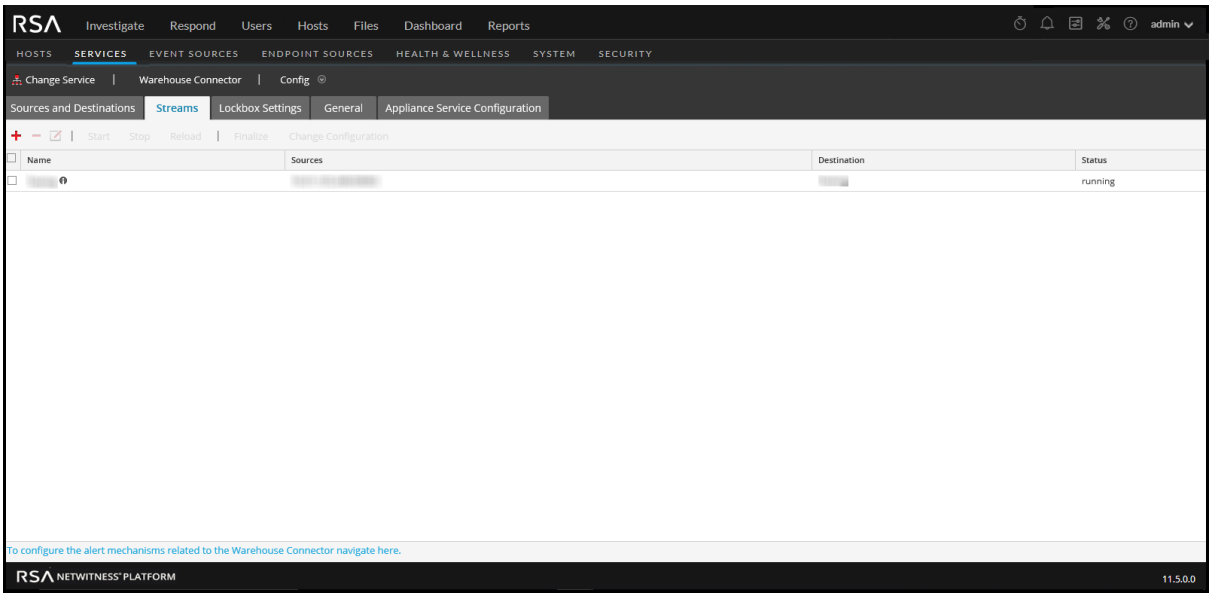
1. Create a stream
2. Finalize the stream
3. Start the stream

Create a Stream

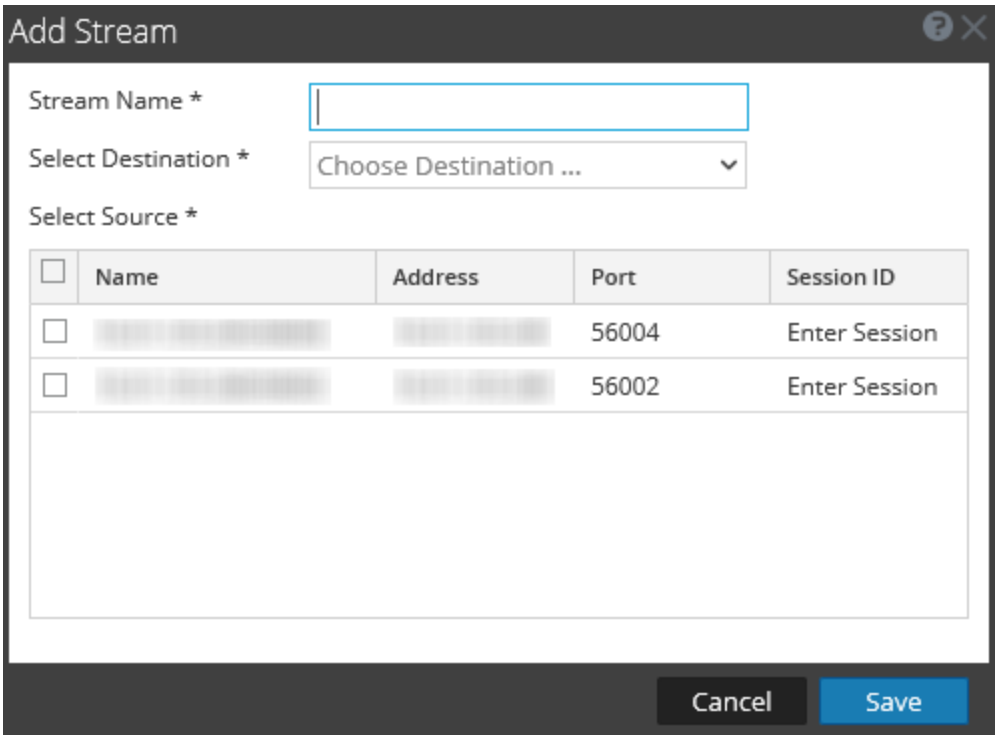
To create a stream:

1. Go to  **(Admin) > Services**.
2. In the Services view, select the added Warehouse Connector service and select  **> View > Config**.

   The Services Config view of Warehouse Connector is displayed.
3. Click the **Streams** tab.



4. On the **Streams** tab, click .



5. In the **Add Stream** dialog, perform the following steps:

   a. In the **Stream Name** field, enter a name for the stream.

   Note: The **Stream Name** field does not support spaces or special characters except underscore (_).

   b. In the **Select Destination** drop-down menu, select a destination from the list of destinations added to the Warehouse Connector.

c. In the **Select Source** field, select sources from the list of sources displayed.

d. In the **Session ID** column, enter the last session id.
If you provide any session id, the Warehouse Connector will start the aggregation from that session, whereas if this is left blank, the aggregation will start from the current session.

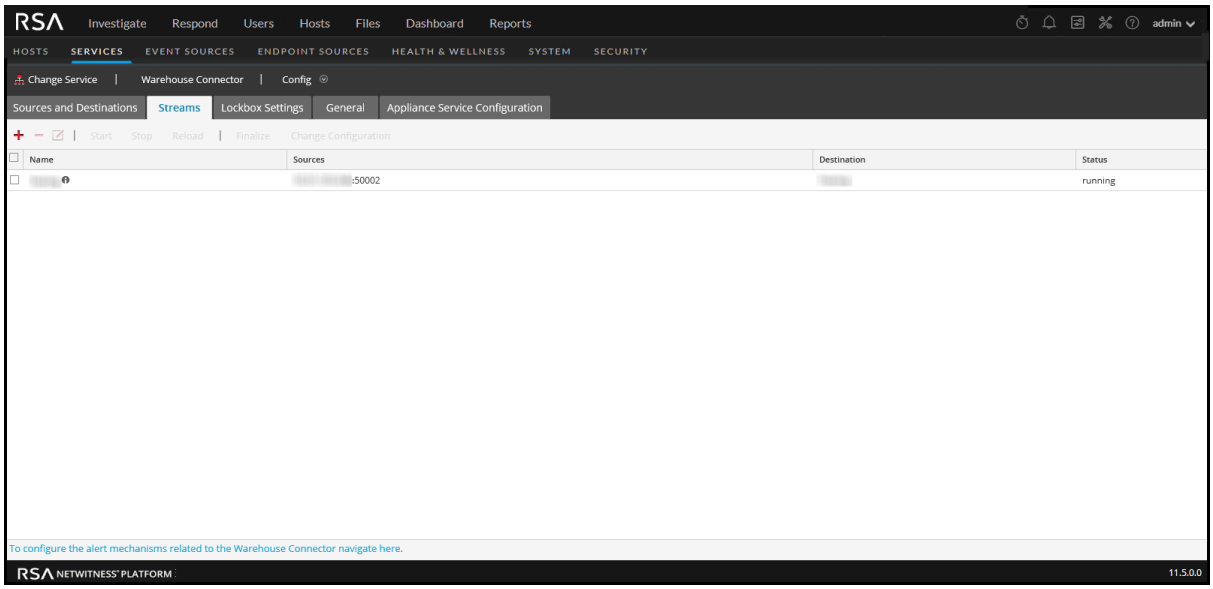
e. Click **Save**.

Finalize a Stream

To finalize the stream:

1. Go to **(Admin) > Services**.
2. In the Services view, select the added Warehouse Connector service and select **> View > Config**.

   The Services Config view of Warehouse Connector is displayed.
3. On the **Streams** tab, select the stream that you have created.



4. Click **Finalize**.

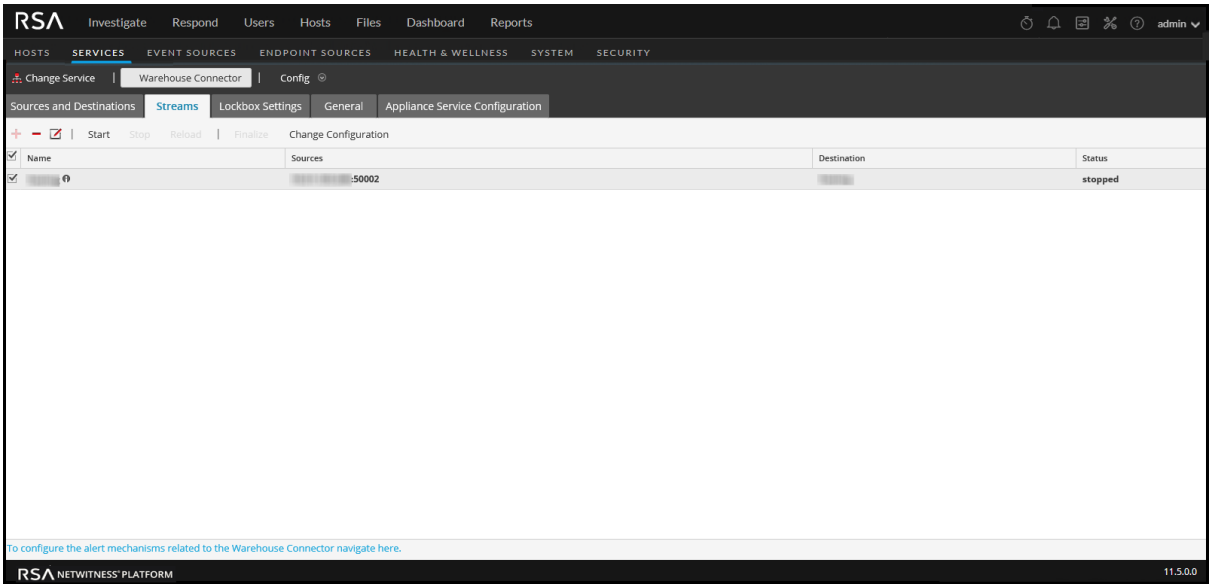
Start a Stream

Note: If you have deployed a Warehouse Connector Virtual Appliance, make sure that you change the default value of the `Maximum Message Hold Count` parameter to 800000. For more information, see General Tab Settings.

To start the stream:

1. Go to  **(Admin) > Services**.
2. In the Services view, select the added Warehouse Connector service and select  **> View > Config**.

   The Services Config view of Warehouse Connector is displayed.
3. On the **Streams** tab, select the stream that you have created.



4. Click **Start**.

Monitor a Warehouse Connector

By monitoring a Warehouse Connector, you can automatically generate notifications when critical thresholds concerning Warehouse Connector and its storage have been met.

To monitor a Warehouse Connector:

1. Go to  **(Admin) > Services**.
2. In the Services view, select the added Warehouse Connector service and select  **> View > Config**.
   The Services Config view of Warehouse Connector is displayed.
3. Click the **Streams** tab.
4. At the bottom of the **Streams** tab, click **To configure the alert mechanisms related to the Warehouse Connector navigate here**.
   The Warehouse Connector Monitoring view is displayed.

   This page is deprecated and will be removed in a future release.
5. In the **Source or Destination Status** section, select the number of minutes or hours in the **Notify Offline For** field.
   You will receive a notification if the source or destination connection fails for the defined number of minutes or hours.
6. In the **Stream Status** section, perform the following:
   a. In the **Notify Stopped For** field, define the number of minutes or hours after which you would like to receive a notification when the stream goes offline.
   b. In the **Disk Is** field, define the limit on the percentage of disk usage after which you would like to receive a notification.
   c. In the **Source is Behind** field, define the number of sessions. A notification is raised if the source goes behind the defined number of sessions.
   d. In the **Rejected Folder Size is** field, define the limit on the percentage of folder usage after which would like to receive a notification.
   e. In the **Number Of Files in Permanent Failure Folder** field, define the limit on the number of files in the permanent failure folder after which you would like to receive a notification.
7. In the **Notification Type** field, perform the following:
   a. Click **Configure email or distribution list** to configure email so that you can receive notifications in NetWitness. For more information, see the "Configure Email Server and Notification Account" topic in the *System Configuration* guide.
   b. Click **Configure Syslog and SNMP Trap servers** to configure audit logs. For more information, see the "Configure Syslog and SNMP Settings" topic in the *System Configuration* guide.
   c. Select the following notification mechanisms as per your requirement:

- **NetWitness Console** - To get notifications on the NetWitness UI notification toolbar.
- **Email** - To get email notifications.
- **Syslog Notification** - To generate syslog events.
- **SNMP Trap Notifications** - To get audit events as SNMP traps.

Add Warehouse as a Data Source to Reporting Engine

You must add Warehouse as a data source to Reporting Engine to make this data source available to report against this Reporting Engine. For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.

Analyze a Warehouse Report

The Warehouse modules provide analysts with reports of early indicators of compromise. The following Warehouse reports can be analyzed in NetWitness:

- Suspicious Domains report
- Suspicious DNS Activity report
- Host Profile report

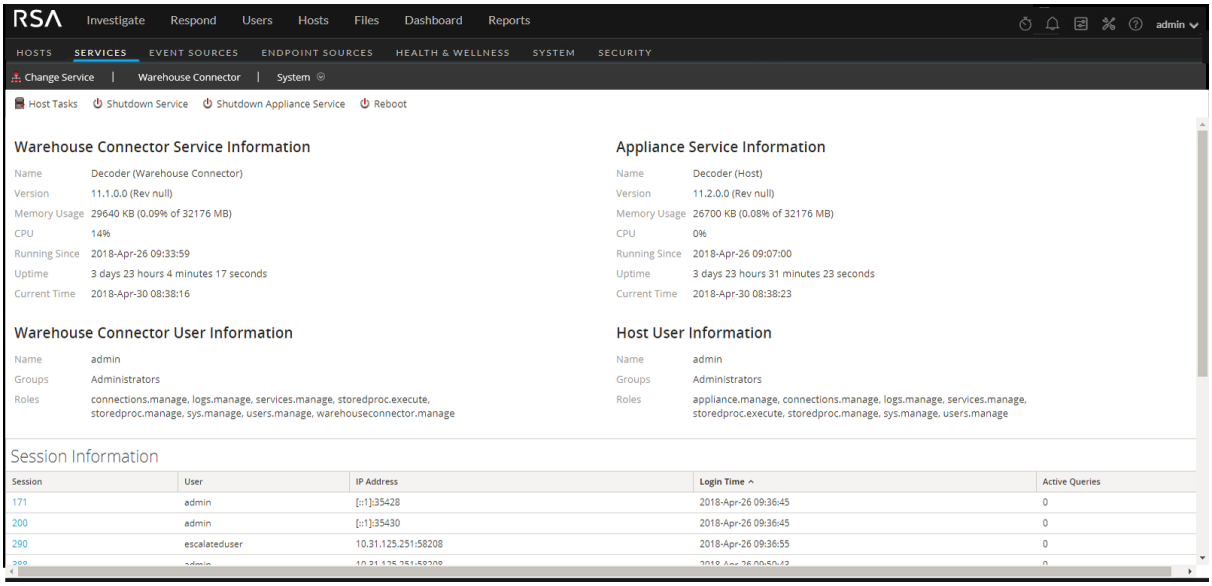
For more information, see "Step 4. Analyze a Warehouse Report" in the *Warehouse Guide*.

View the Warehouse Connector Service

While the information displayed in the Services System view is the same for all types of core services, several options in the toolbar are relevant only for Warehouse Connector.

To access this view:

1. Go to **(Admin) > Services**.
2. In the Services view, select a Warehouse Connector and select **> View > System**.
   The Systems view for the selected Warehouse Connector is displayed.



The following is an example of toolbar options for Warehouse Connectors.



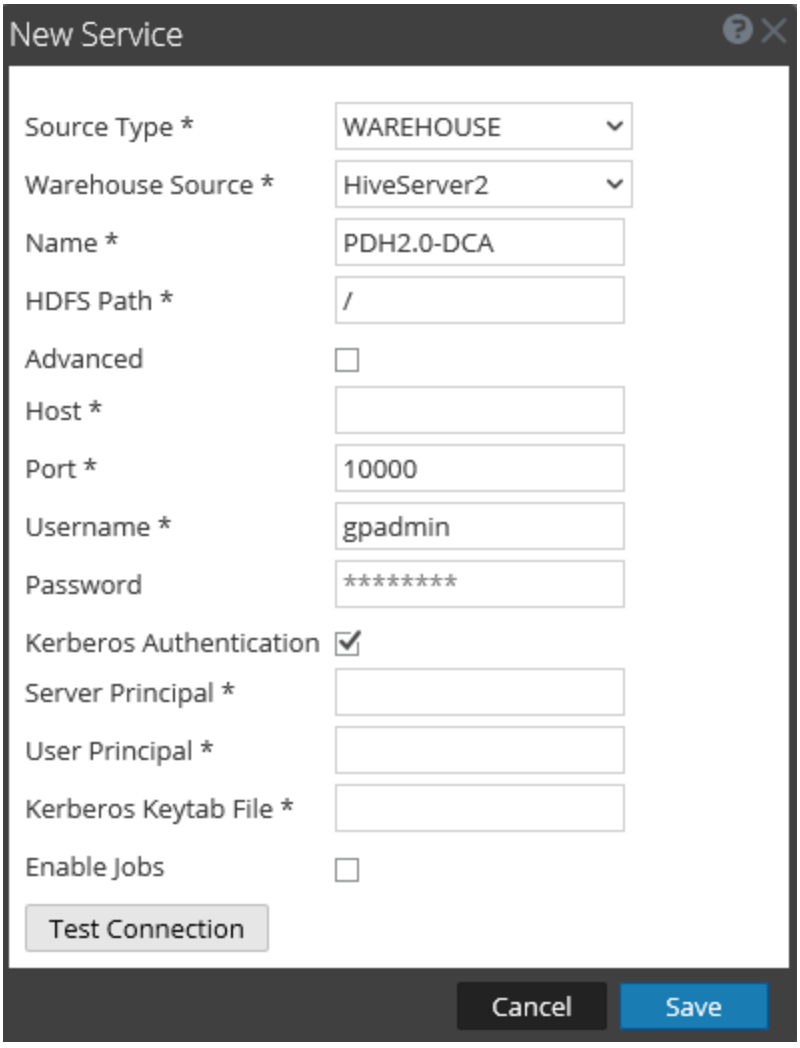
Host Tasks, Shutdown Service, Shutdown Appliance Service or (Shutdown Appliance), and Reboot are common to all services and are described in the *Hosts and Services Getting Started Guide*.

Troubleshoot the Warehouse Connector

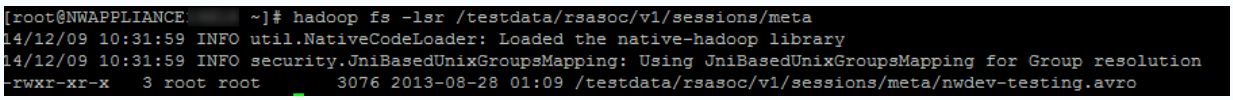
The following information suggests the possible issues that NetWitness users may encounter when adding a Warehouse service to the Reporting Engine as a data source for reporting in NetWitness. Look for explanations and solutions in this section.

While adding a Warehouse service to the Reporting Engine as a data source for reporting, you may observe some of the errors listed in this document. Information is provided on how to troubleshoot the errors and add the data source successfully.

The following figure shows the New Service dialog.



For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.

Error	Possible Solutions
Could not open connection to HiveServer	• Ensure that the HiveServer2 is running on the Host. • Check if the port provided can be accessible from the Reporting Engine server.
No Schema found in HDFS path	Ensure that meta avro data file(s) are available in the HDFS path (**<HDFS Path>/rsasoc/v1/sessions/meta**) mentioned. The following figure shows an example of the command to check the files in hdfs.
	
Could not open connection to HiveServer, GSS initiate failed	GSS initiate failed errors will be observed only in the case of Kerberos enabled Hive. Ensure that the proper keytab file is provided and it should have read options for the rsasoc user (user on which the Reporting Engine Server runs). Ensure that the system time is synchronized between KDC, Hadoop (HortonWorks) server, and the Reporting Engine system.

Add SFTP Destination

To add SFTP destination to Warehouse Connector, you need to follow the below steps:

1. SSH to NetWitness Host and run the following command to edit `ssd_config` file:
   `vi/etc/ssh/sshd_config file`
2. Do not comment `HostKey/etc/ssh/ssh_host_rsa_key`
3. You must append `KexAlgorithms +diffie-hellman-group14-sha1,diffie-hellman-group-exchange-sha1`
4. Run the following command to restart NetWitness Host:
   `systemctl restart sshd`
5. On Warehouse Connector Host SSH to NetWitness Host and run the following command:
   `ssh -o "HostKeyAlgorithms ssh-rsa" root@<NWhost_IP>`
   For example, `ssh -o "HostKeyAlgorithms ssh-rsa" root@10.10.10.10`
6. Run the following curl command on Warehouse Connector Host to authenticate:
   `curl -vvvv sftp://root:netwitness@<NWhost_ip>:22/root/`
   For example, `curl -vvvv sftp://root:netwitness@<10.10.10.10>:22/root/`
7. You can now add SFTP Destination from UI.

Manage a Stream

You can manage a stream using the following procedures:

- Edit a Stream
- Reload the Stream
- Specify meta filters for a Stream
- Define multi-valued metas

Edit a Stream

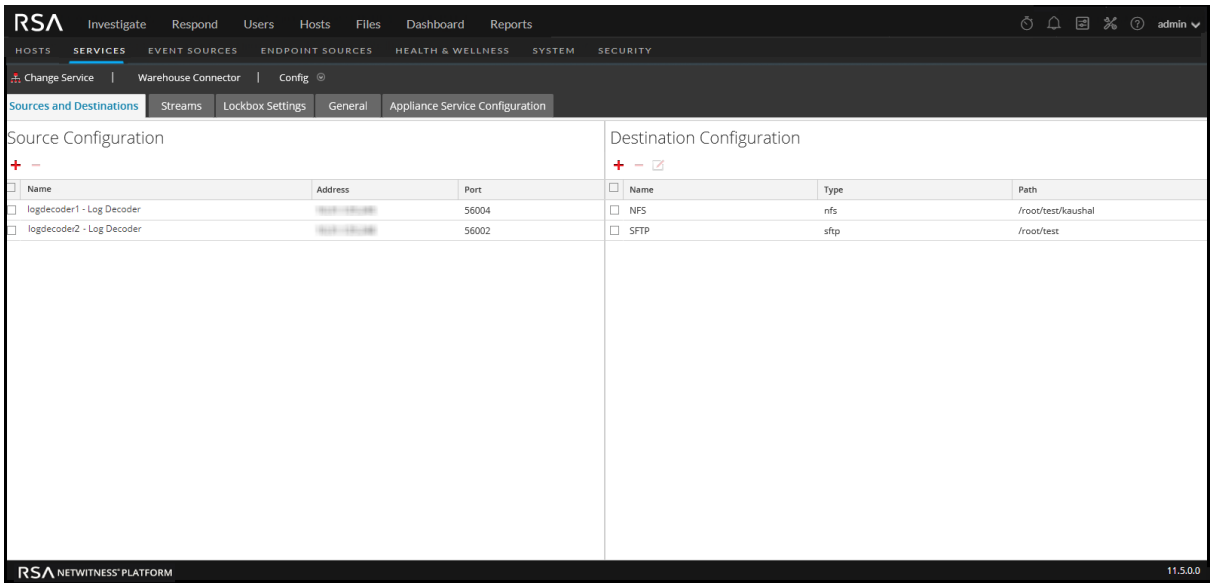
You can edit a stream to perform the following:

- Add data sources to the stream.
- Delete existing data sources from the stream.

To edit a stream:

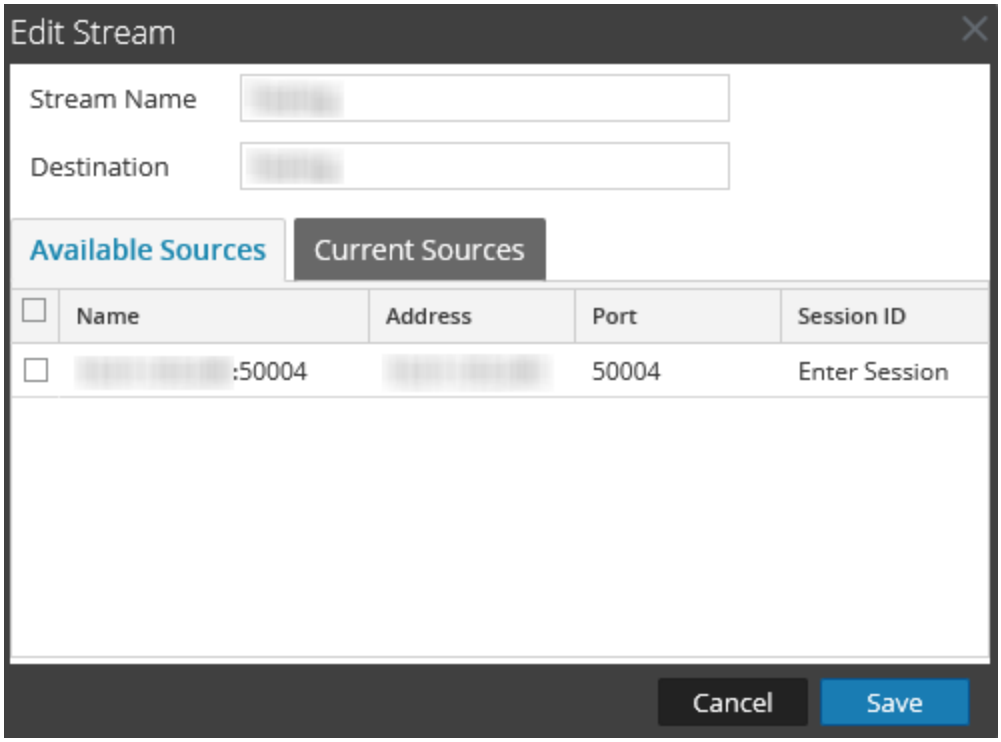
1. Go to **(Admin) > Services**.
2. In the Services view, select the added Warehouse Connector service and select **> View > Config**.

   The Services Config view of Warehouse Connector is displayed.

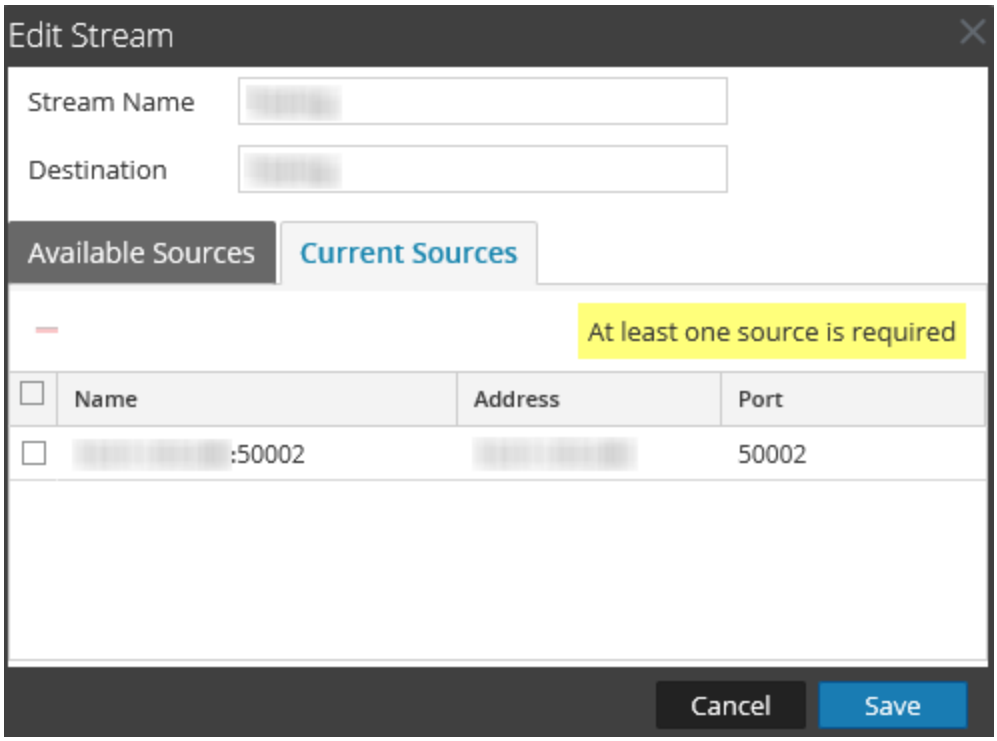


3. On the **Streams** tab, click .
4. In the **Edit Stream** dialog, you can perform the following:

- On the **Available Sources** tab, you can select the available data sources to add to the stream and click **Save**.



- On the **Current Sources** tab, you can delete an existing data source from the stream. Select the data source and click − .



Reload the Stream

When you reload the stream, the Warehouse Connector updates the schema file for the stream. You must reload the stream when you add a new custom meta to the Log Decoder or Decoder.

To reload the stream:

1. Go to **(Admin) > Services**.
2. In the Services view, select the added Warehouse Connector service and select **> View > Config**.

   The Services Config view of Warehouse Connector is displayed.
3. On the **Streams** tab, select the stream that you want to reload.
4. Click **Reload**.

Specify Meta Filters for a Stream

You need to specify the filter for each stream in the `export.session.meta.fields` parameter in the Explore view of the Warehouse Connector.

The following table lists the values that you can provide as a filter:

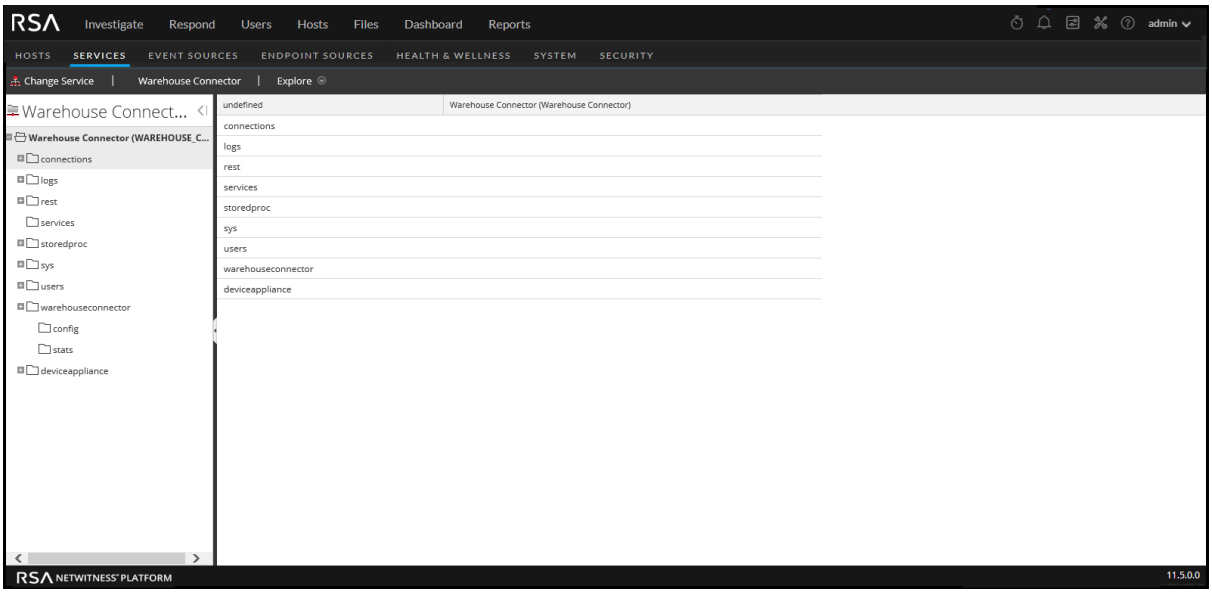
Values	Description
*	All the collected metas are written to SAW.
*, *meta1*, *meta2*	All the metas except the defined metas are written to SAW. For example, **Filter:**`*,ip.src` All the metas except `ip.src` is written to SAW.
meta1, meta2, meta3	Only the defined metas are written to SAW.

Note: By default, the following metas are written to Warehouse even if you specify them in the filter:
- ng_source
- unique_id
- time

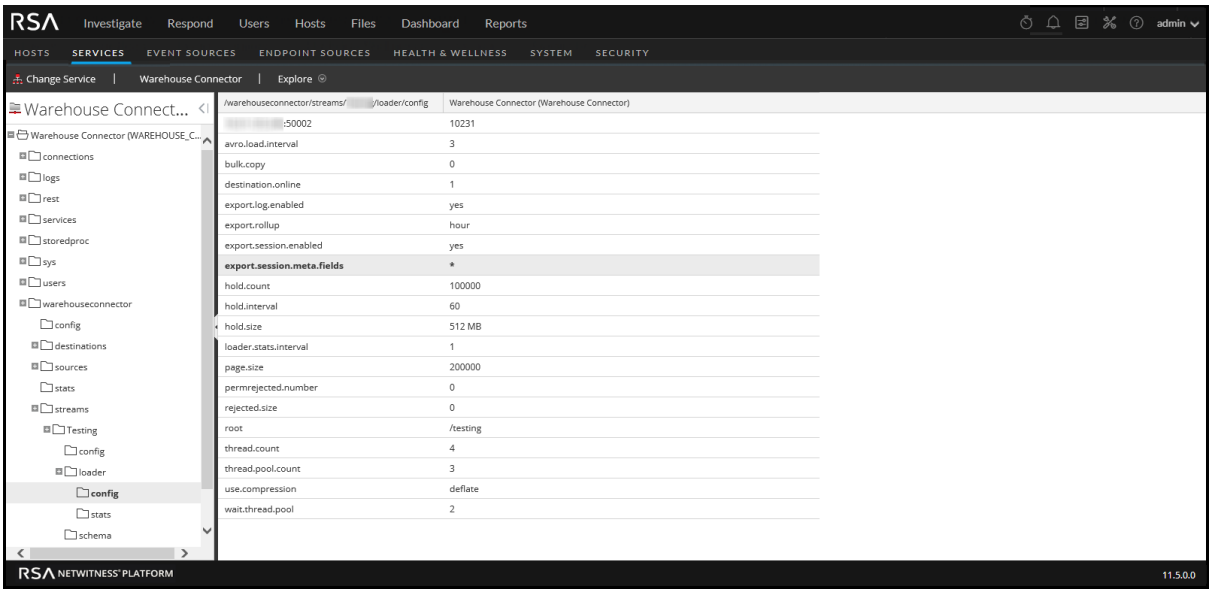
To specify meta filters for a Stream:

1. Go to **(Admin) > Services**.
2. In the Services view, select a Warehouse Connector services and select **> View > Config**.

   The Explore view of the Warehouse Connector service is displayed.



3. In the options panel, select **warehouseconnector > streams > <stream_name> > loader > config**.
4. In the `export.session.meta.fields` parameter, enter the filter.



5. Restart the stream.

Define Multi-valued Metas

You can also define an existing meta or a custom meta to be treated as multi-valued meta.

To define multi-valued metas:

Caution: Defining an existing meta to be treated as multi-valued may change the data type of the meta and cause the associated reports to fail.

1. Create a new file with the filename **multivalue-users.xml** in the **/etc/netwitness/ng** directory.
2. Add the following entries:

```
<?xml version="1.0" encoding="utf-8"?>
<NetWitness>
   <MultiValueMetas>
      <Meta>NEWMETANAME</Meta>
   </MultiValueMetas>
</NetWitness>
```

   Where *NEWMETANAME* is the existing meta or a custom meta to be treated as multi-valued meta.

   Caution: Make sure that you do not add metas that are by default treated as non multi-value.

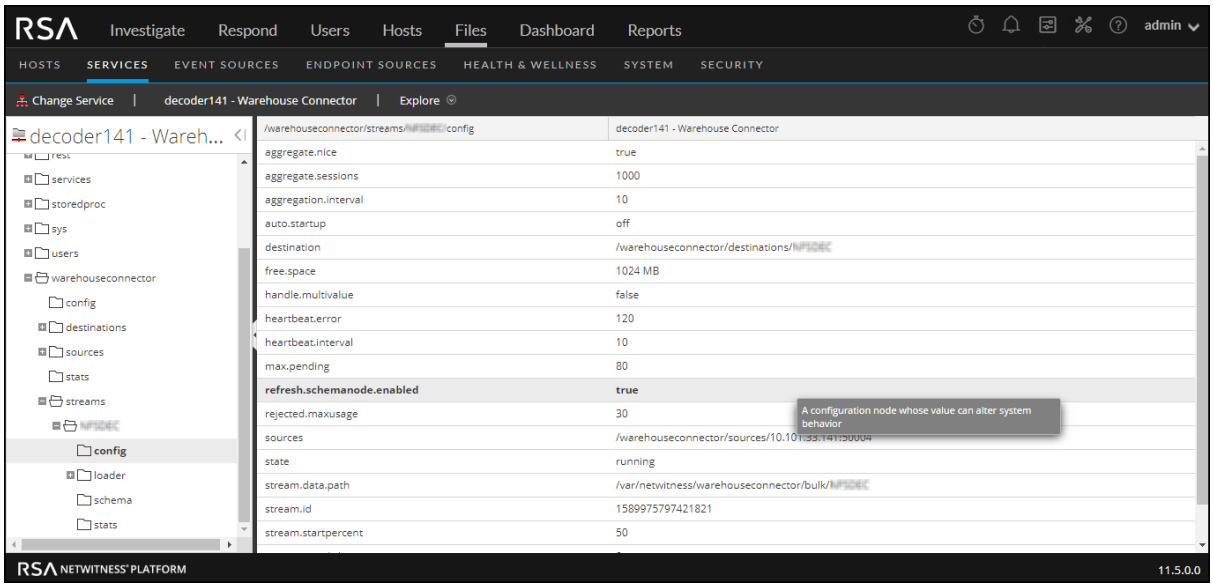
3. Restart the stream.

View the current schema

You can view the current schema that is used by warehouse connector for writing in AVRO files.

To view the current schema:

1. Go to **(Admin) > Services**.
2. In the Services view, select a Warehouse Connector services and select **> View > Explore**.
   The Explore view of the Warehouse Connector service is displayed.
3. In the options panel, select **warehouseconnector > streams > <stream_name> > config**.
4. Set the value for `refresh.schemanode.enabled` parameter to **true**. By default, this value is set to false.

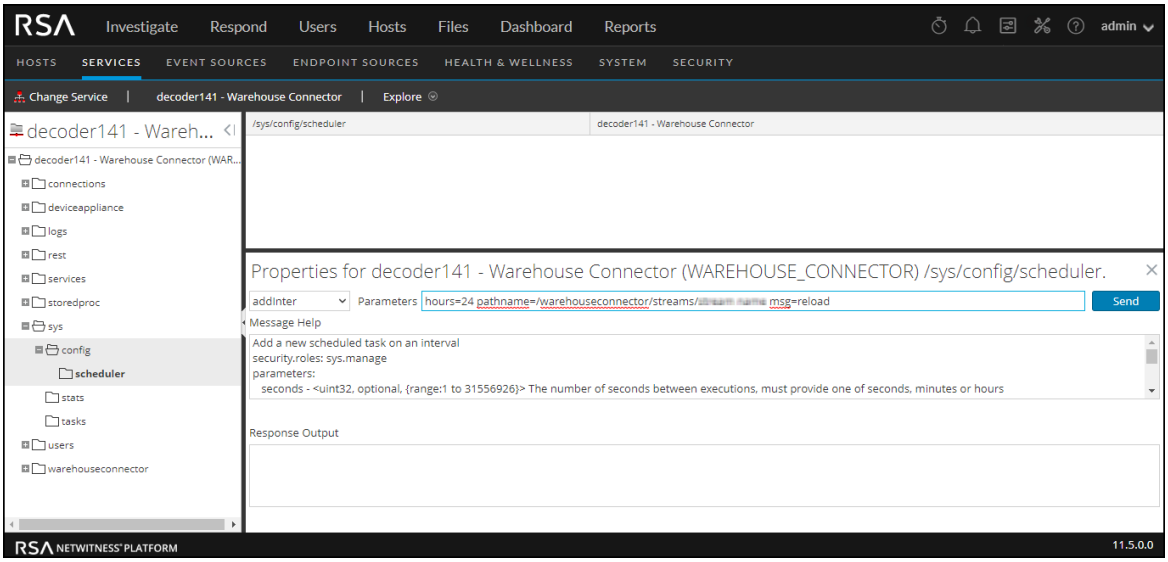


5. Reload the stream. For more information see, Reload the Stream.
6. Restart the Warehouse Connector service.

7. Go to **(Admin) > Services**.
8. In the Services view, select a Warehouse Connector services and select **> View > Explore**.
   The Explore view of the Warehouse Connector service is displayed.
9. In the options panel, select **warehouseconnector > streams > <stream_name> > schema**, to view the current schema.

Note: You must to reload the stream every time, before you want to view the current schema or you can add a scheduler to reload the stream automatically at regular intervals as mentioned in step 10.

10. (Optional) To reload the stream automatically at regular intervals, follow the below steps.
    a. Go to **(Admin) > Services**.
    b. In the Services view, select a Warehouse Connector services and select **> View > Explore**.
       The Explore view of the Warehouse Connector service is displayed.
    c. In options panel, select **warehouseconnector > sys > config > scheduler**. Right click and select **properties**.
    d. In the property drop down select `addInter` and in the **Parameters** text box, add "`hours=24 pathname=/warehouseconnector/streams/<stream name> msg=reload`".



Note: Standard time format used is hours. You can use seconds or minutes format in lieu of hours.

    e. Then click **Send**.

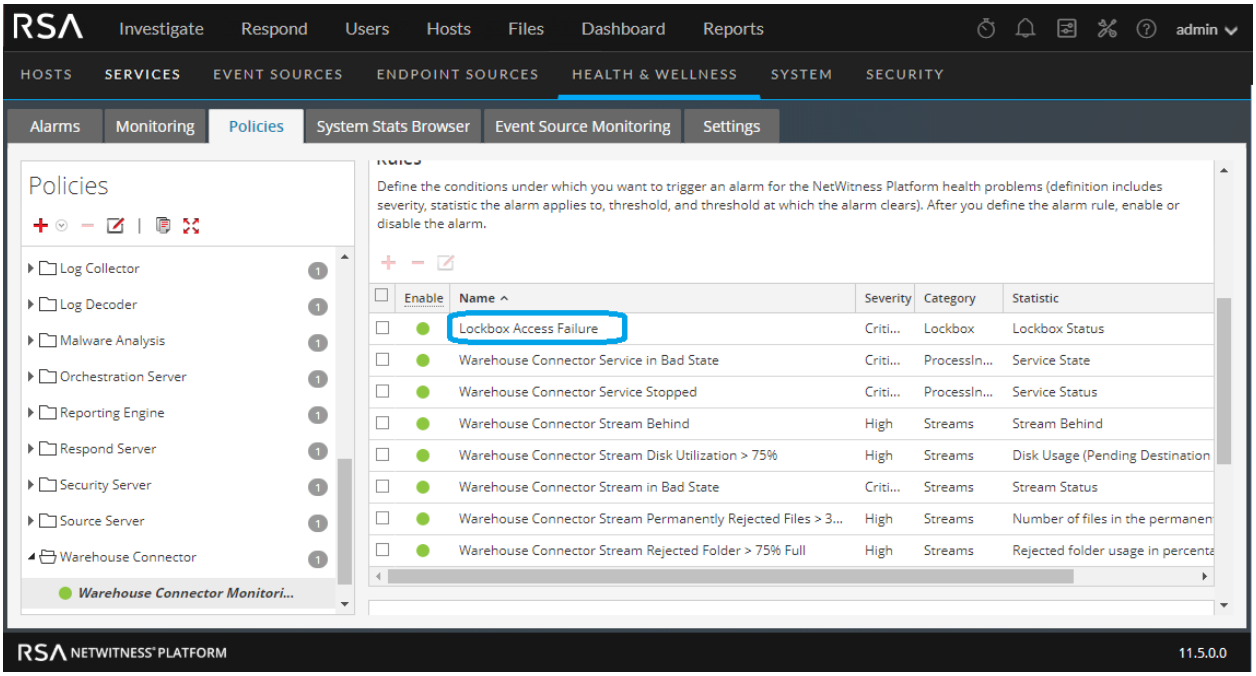
Manage a Lockbox

You can change the lockbox password, as well as refresh the lockbox.

Lockbox Status

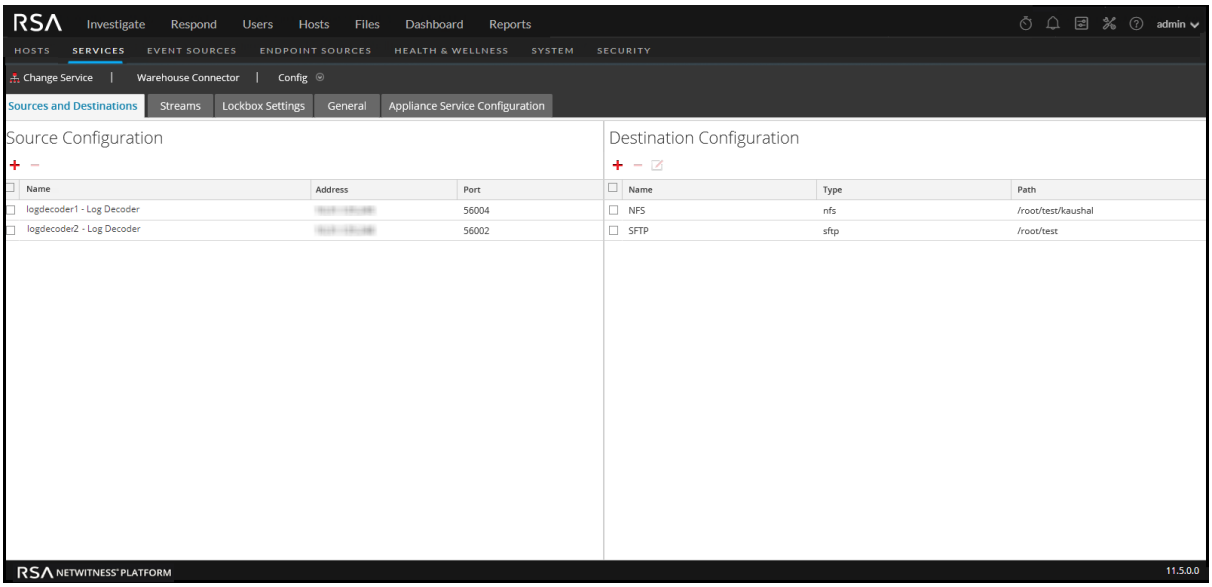
Note: RSA NetWitness Platform has added a Health & Wellness stat for Warehouse Connector to indicate the status of its Lockbox. Also, an out-of-the-box rule has been added so that a Health & Wellness alarm is raised when the Lockbox does not exist or cannot be opened.

To see the rule details, go to **(Admin) > Health & Wellness > Policies > Warehouse Connector** and select **Warehouse Connector Monitoring Policy**.

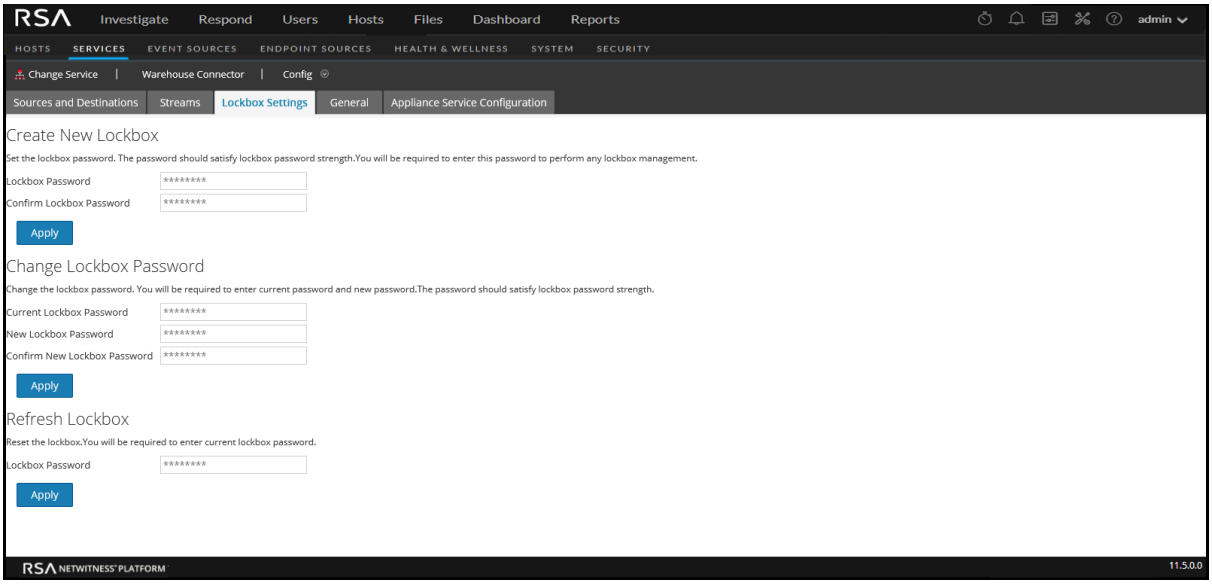


To change the Lockbox password:

1. Log on to NetWitness.
2. Go to **(Admin) > Services**.
3. In the Services view, select a Warehouse Connector service, and select **> View > Config**.
   The Services Config view of Warehouse Connector is displayed.



4. Click the **Lockbox Settings** tab.



5. In the **Change Lockbox Password** section, perform the following:
   a. In the **Current Lockbox Password** field, enter the current lockbox password.
   b. In the **New Lockbox Password** field, enter the new lockbox password.

      Note: The lockbox password must be at least eight characters in length and it must contain at least three of the following groups: one uppercase character [A-Z], one lowercase character [a-z], one numeral [0-9], and one special character.

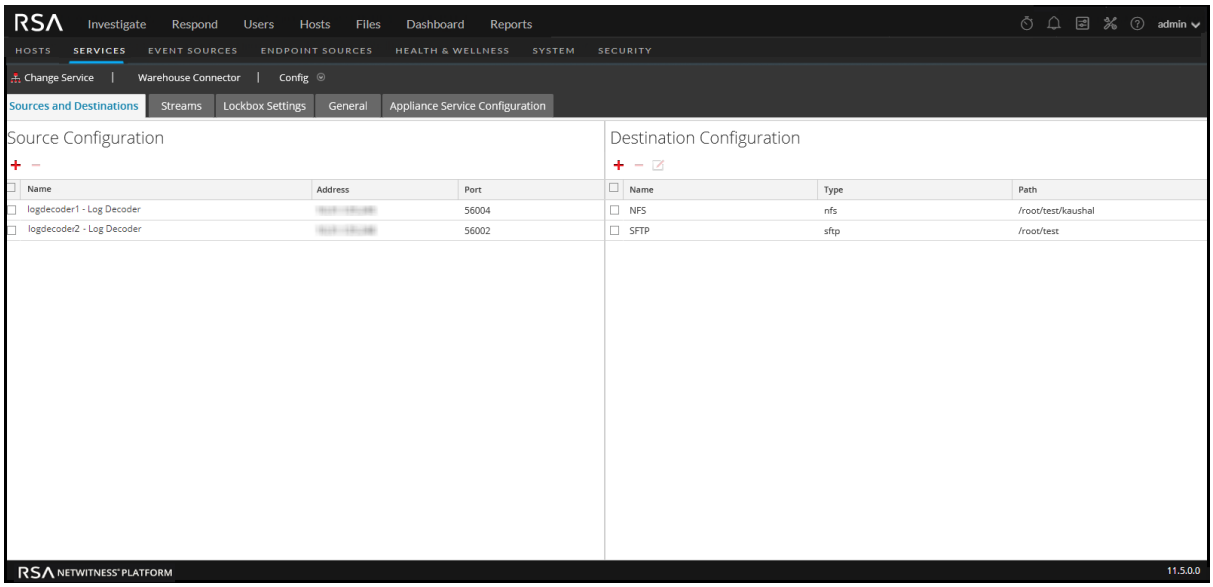
   c. In the **Confirm New Lockbox Password** field, enter the new lockbox password to confirm.
   d. Click **Apply**.

The Lockbox password is successfully changed.

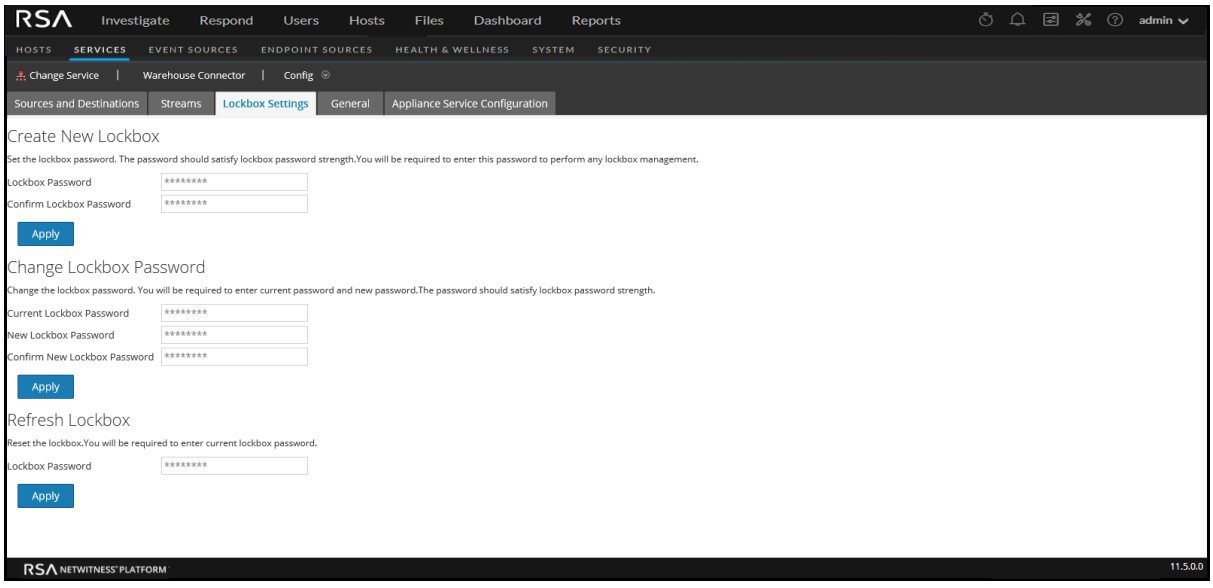
To refresh the Lockbox:

1. Log on to NetWitness.
2. Go to  **(Admin) > Services**.
3. In the Services view, select the added Warehouse Connector service, and select  **> View > Config**.

   The Services Config view of Warehouse Connector is displayed.



4. Click the **Lockbox Settings** tab.



5. In the **Refresh Lockbox** section, enter the current lockbox password in the **Lockbox Password** field.
6. Click **Apply**.
   The Lockbox is reset.

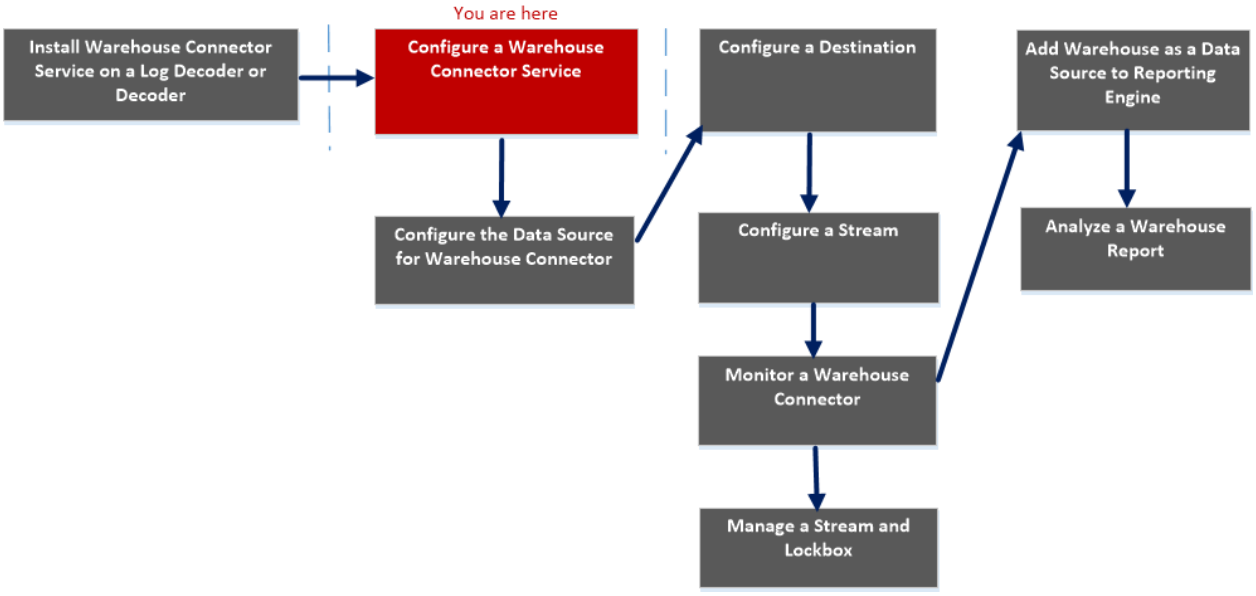
Warehouse Connector Configuration References

This section contains descriptions of the user interface as well as other reference information.

General Tab Settings

The General tab displays the general configuration settings for Warehouse Connector service.

Workflow



What do you want to do?

Role	I want to...	Refer to...
Administrator	Install Warehouse Connector Service on a Log Decoder or Decoder	Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid
Administrator	Configure a Warehouse Connector Service*	Configure a Warehouse Connector Service
Administrator	Configure the Data Source for Warehouse Connector	Configure the Data Source for Warehouse Connector
Administrator	Configure the Destination using NFS, SFTP, WebHDFS.	Configure the Destination Using NFS Configure the Destination Using SFTP Configure the Destination Using WebHDFS
Administrator	Configure a Stream	Configure a Stream
Administrator	Monitor a Warehouse Connector	Monitor a Warehouse Connector

Role	I want to...	Refer to...
Administrator	Add Warehouse as Data Source to Reporting Engine	For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator	Analyze a Warehouse Report	For more information, see "Step 4. Analyze a Warehouse Report" in the *Warehouse Guide*.
Administrator	Manage a Stream and Lockbox	Manage a Stream

*You can perform this task in the current view

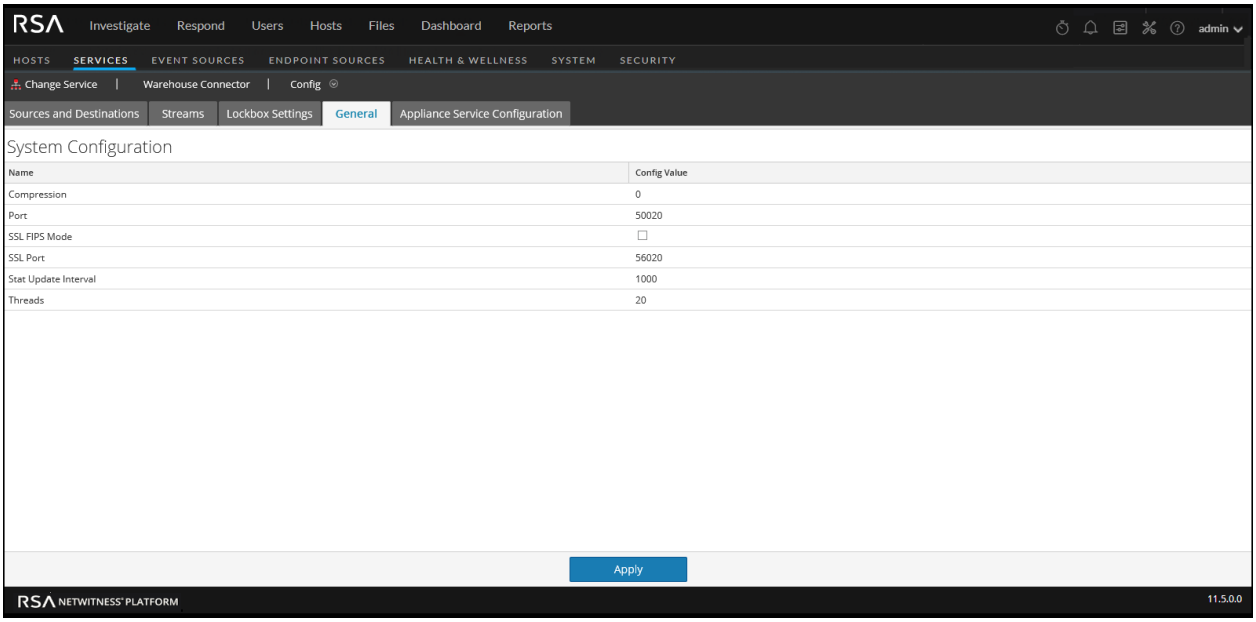
Related topics

- Configure a Warehouse Connector Service

Quick Look

The following figure shows the General tab on the Warehouse Connector Services Config view.

The General tab displays the system configuration parameters for the Warehouse Connector service.



When you add a Warehouse Connector service, default values are in effect. NetWitness designed the default values to accommodate most environments and recommends that you do not modify these values because it may adversely affect performance.

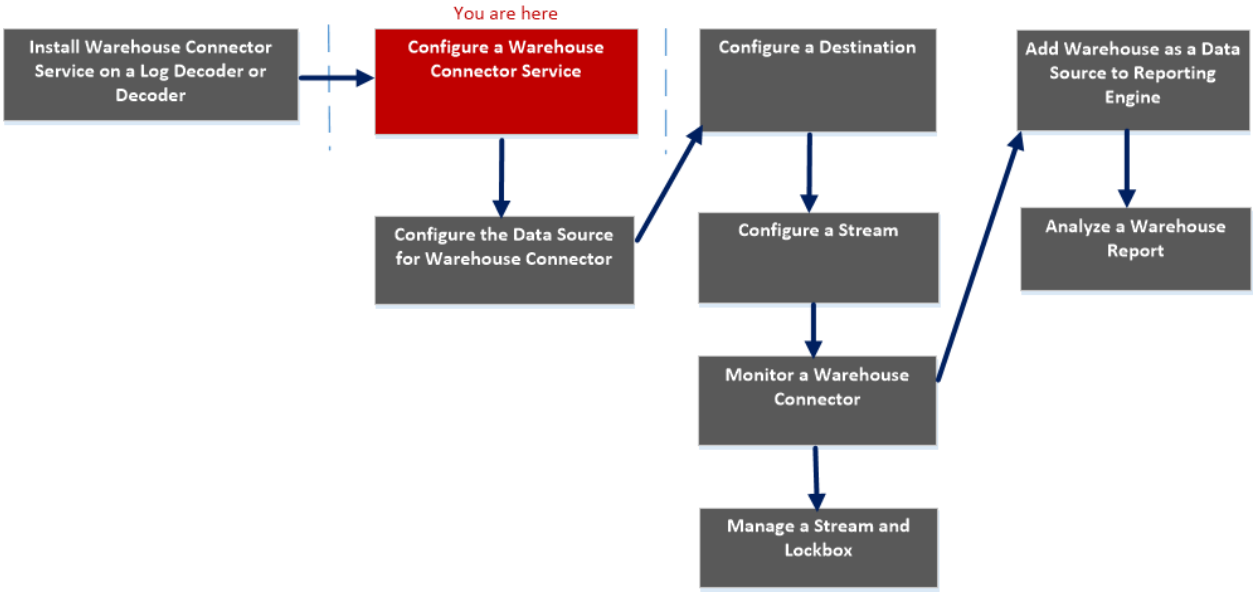
The following table describes the System Configuration parameters:

Name	Config Value
Compression	Determines the minimum amount of bytes before a message is compressed. If set to zero, messages are not compressed.
Port	Determines the port used by the service. **Note:** If you change the port number, ensure that you restart the service.
SSL	If enabled, all the data transferred in the network will be encrypted using SSL.
Stat Update Interval	Determines how often (in milliseconds) statistic nodes are updated in the system.
Threads	Determines the number of threads in the thread pool to handle incoming requests.

Appliance Service Configuration Tab Settings

The Appliance Service Configuration tab displays the appliance configuration settings for Warehouse Connector service. For more information, see "Appliance Service Configuration" in the *Hosts and Services Getting Started Guide*.

Workflow



What do you want to do?

Role	I want to...	Refer to...
Administrator	Install Warehouse Connector Service on a Log Decoder or Decoder	Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid
Administrator	Configure a Warehouse Connector Service*	Configure a Warehouse Connector Service
Administrator	Configure the Data Source for Warehouse Connector	Configure the Data Source for Warehouse Connector
Administrator	Configure the Destination using NFS, SFTP, WebHDFS.	Configure the Destination Using NFS Configure the Destination Using SFTP Configure the Destination Using WebHDFS
Administrator	Configure a Stream	Configure a Stream
Administrator	Monitor a Warehouse Connector	Monitor a Warehouse Connector

Role	I want to...	Refer to...
Administrator	Add Warehouse as Data Source to Reporting Engine	For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator	Analyze a Warehouse Report	For more information, see "Step 4. Analyze a Warehouse Report" in the *Warehouse Guide*.
Administrator	Manage a Stream and Lockbox*	Manage a Stream

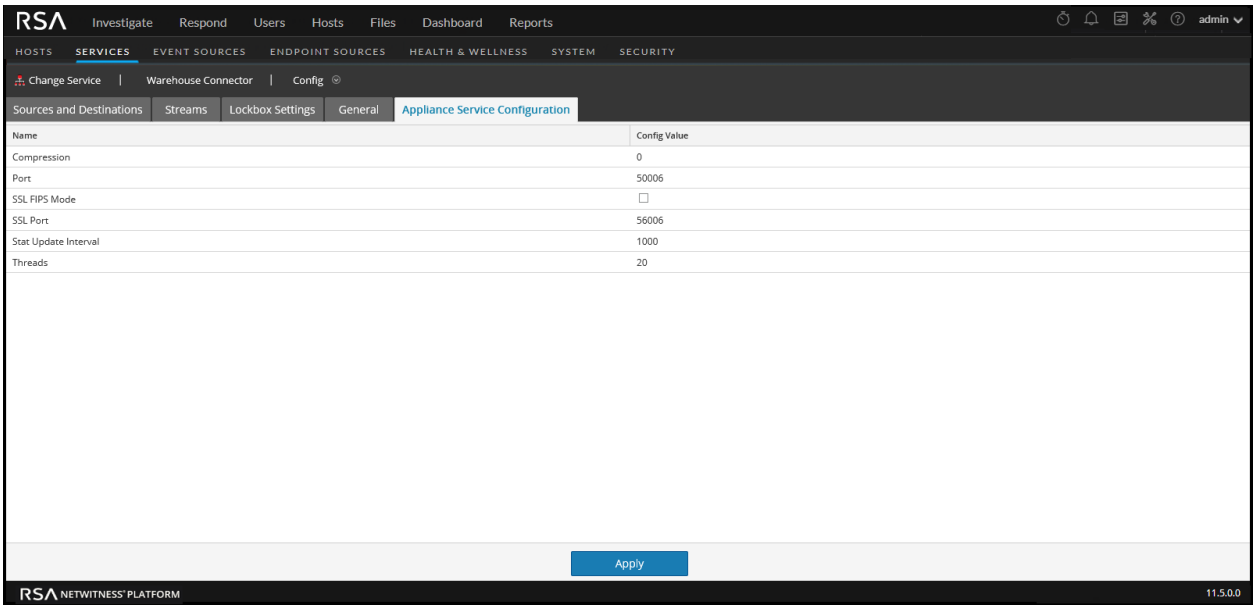
*You can perform this task in the current view

Related topics

- Configure a Warehouse Connector Service

Quick Look

The following figure shows the different settings on the Appliance Service Configuration tab.



When you add a Warehouse Connector service, default values are in effect. NetWitness designed the default values to accommodate most environments and recommends that you do not edit these values because it may adversely affect performance.

The following table describes the Appliance Service Configuration parameters:

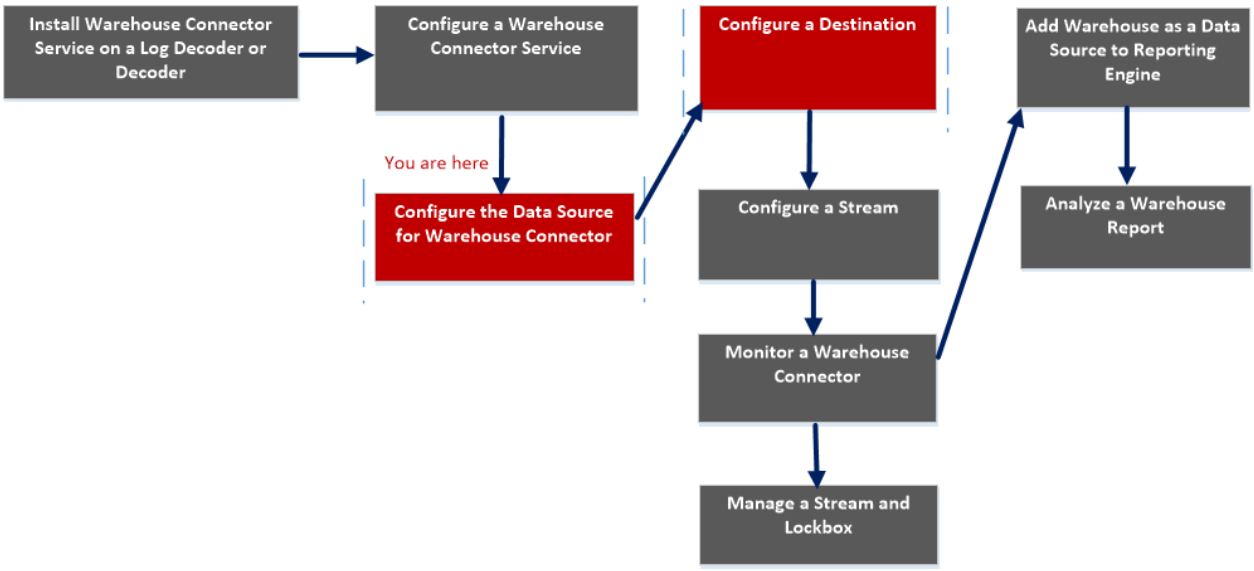
Name	Configuration Value
Compression	Determines the minimum amount of bytes before a message is compressed. If set to zero, messages are not compressed.

Name	Configuration Value
Port	Determines the port used by the service. **Note:** If you change the port number, ensure that you restart the service.
SSL FIPS Mode	If enabled, all the data transferred in the network will be encrypted using SSL FIPS.
SSL Port	Determines the SSL port used by the service.
Stat Update Interval	Determines how often (in milliseconds) statistic nodes are updated in the system.
Threads	Determines the number of threads in the thread pool to handle incoming requests.

Sources and Destinations Configuration

The Sources and Destinations tab for a Warehouse Connector in the Services Config view provides a way to manage basic service configuration and configure source and destination.

Workflow



What do you want to do?

Role	I want to...	Refer to...
Administrator	Install Warehouse Connector Service on a Log Decoder or Decoder	Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid
Administrator	Configure a Warehouse Connector Service	Configure a Warehouse Connector Service
Administrator	Configure the Data Source for Warehouse Connector*	Configure the Data Source for Warehouse Connector
Administrator	Configure the Destination using NFS, SFTP, WebHDFS*	Configure the Destination Using NFS Configure the Destination Using SFTP Configure the Destination Using WebHDFS
Administrator	Configure a Stream	Configure a Stream
Administrator	Monitor a Warehouse Connector	Monitor a Warehouse Connector

Role	I want to...	Refer to...
Administrator	Add Warehouse as Data Source to Reporting Engine	For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator	Analyze a Warehouse Report	For more information, see "Step 4. Analyze a Warehouse Report" in the *Warehouse Guide*.
Administrator	Manage a Stream and Lockbox	Manage a Stream

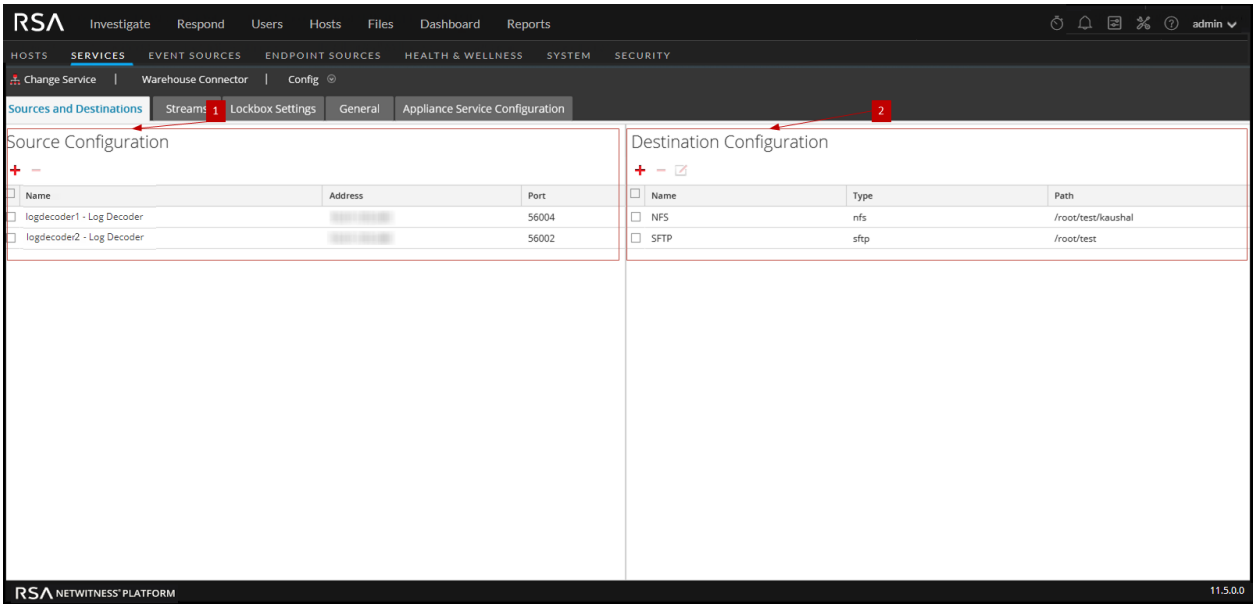
*You can perform this task in the current view

Related topics

- Configure the Data Source for Warehouse Connector
- Configure the Destination

Quick Look

The following figure shows the Sources and Destinations tab on the Warehouse Connector Services Config view.



The Sources and Destinations tab includes the following two sections:

1 Source Configuration
2 Destination Configuration

Source Configuration

The Source Configuration section allows you to configure the data sources from which the Warehouse Connector service needs to collect data.

The following is an example of the Source Configuration section.

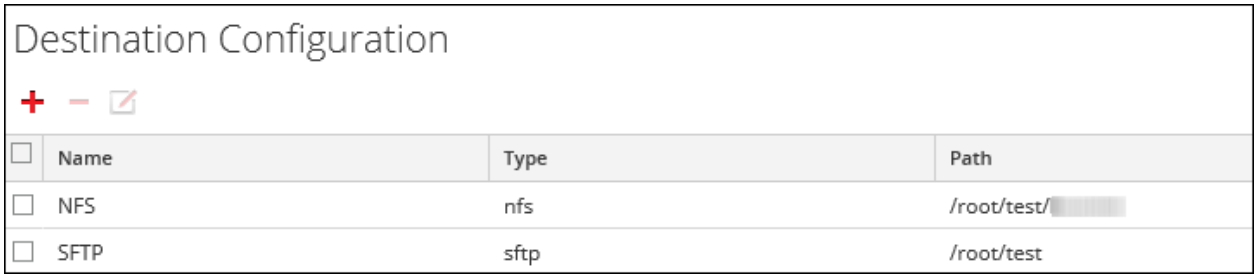


The Source Configuration section allows you to perform the following:

Features	Description
+	Add the data source.
−	Delete the data source.

Destination Configuration

The Destination Configuration section allows you to configure the destination to which the Warehouse Connector service needs to write the collected data.



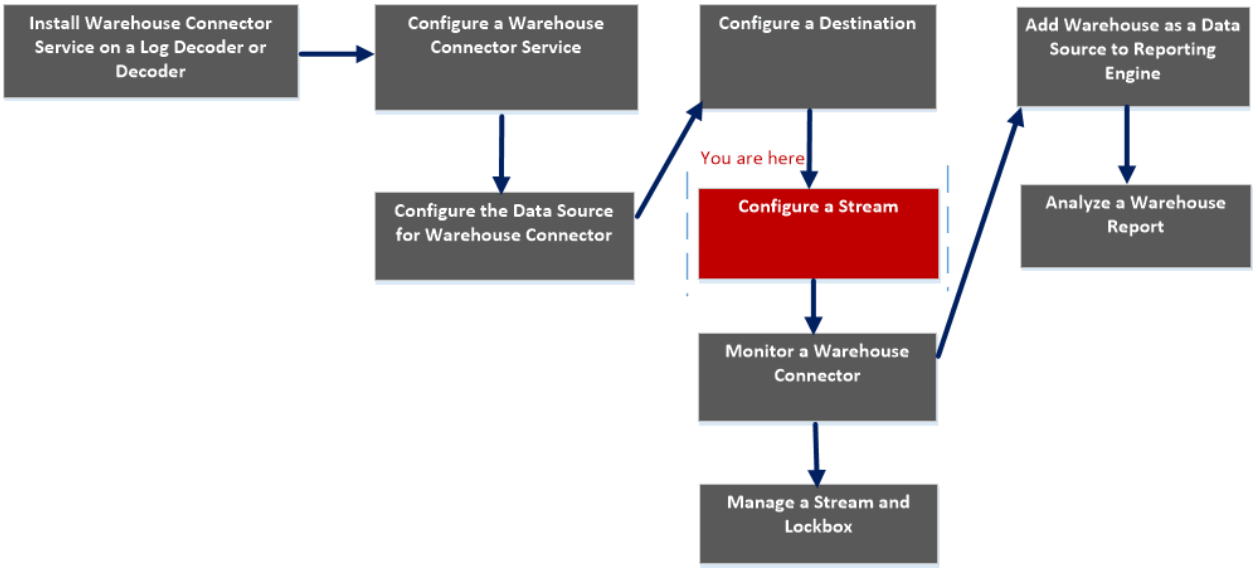
The Destination Configuration section allows you to perform the following:

Features	Description
+	Add the destination.
−	Delete the destination.
	Edit the destination. **Note:** You can only edit the SFTP destination type.

Add Stream Dialog

You can configure and add a stream to a Warehouse Connector in this dialog

Workflow



What do you want to do?

Role	I want to...	Refer to...
Administrator	Install Warehouse Connector Service on a Log Decoder or Decoder	Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid
Administrator	Configure a Warehouse Connector Service	Configure a Warehouse Connector Service
Administrator	Configure the Data Source for Warehouse Connector	Configure the Data Source for Warehouse Connector
Administrator	Configure the Destination using NFS, SFTP, WebHDFS.	Configure the Destination Using NFS Configure the Destination Using SFTP Configure the Destination Using WebHDFS
Administrator	Configure a Stream*	Configure a Stream
Administrator	Monitor a Warehouse Connector	Monitor a Warehouse Connector

Role	I want to...	Refer to...
Administrator	Add Warehouse as Data Source to Reporting Engine	For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator	Analyze a Warehouse Report	For more information, see "Step 4. Analyze a Warehouse Report" in the *Warehouse Guide*.
Administrator	Manage a Stream and Lockbox	Manage a Stream

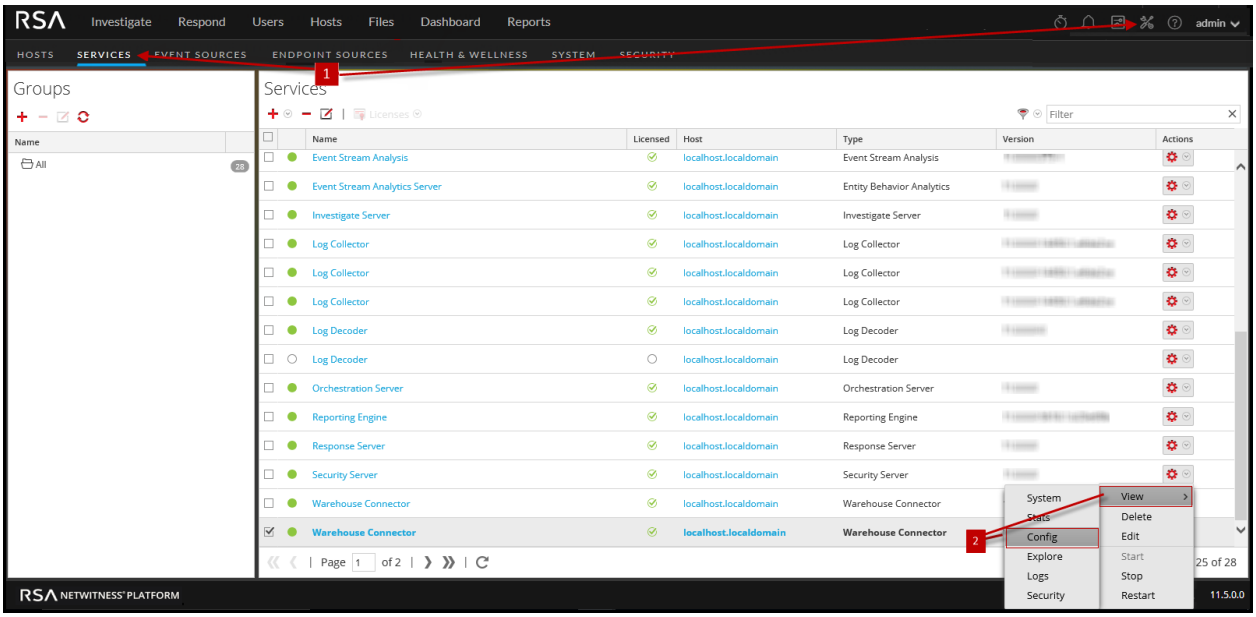
*You can perform this task in the current view

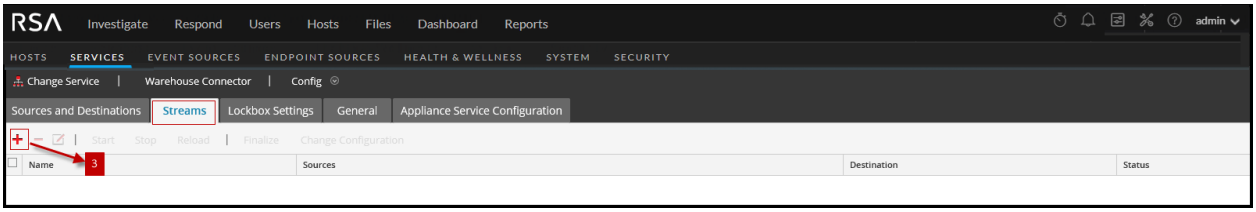
Related Topics

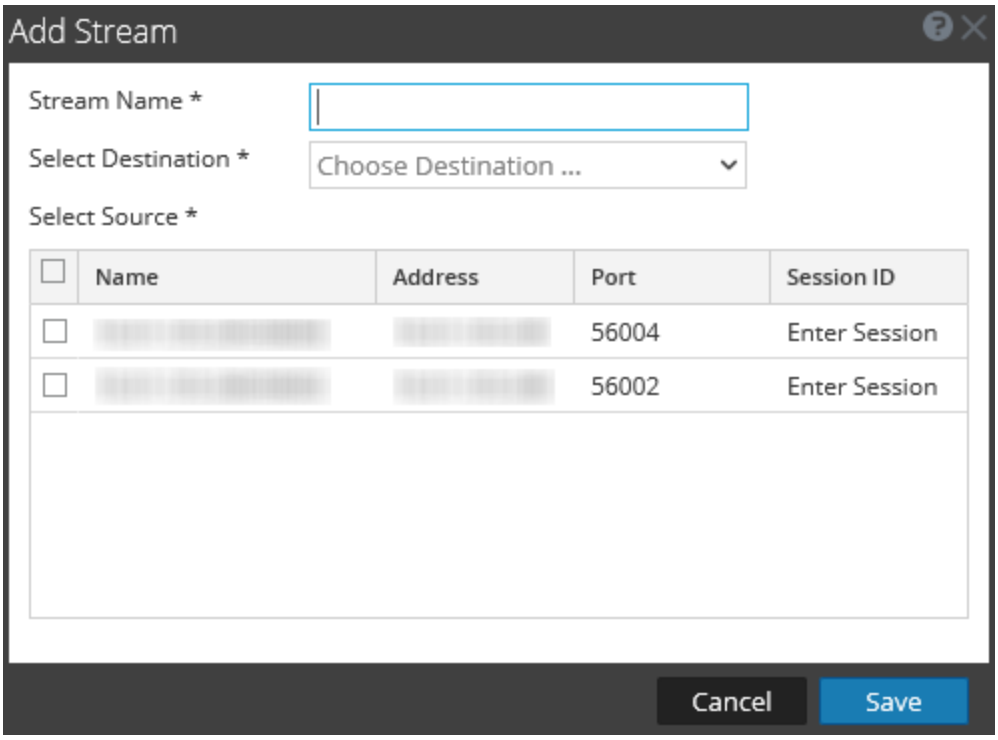
- Configure a Stream

Quick Look

The following figure is an example with the important features labeled.







1. Go to **(Admin) > Services**.
2. In the services view, select a Warehouse Connector service and select **>view>config**
3. In the **Streams** tab, click to view the add stream dialog.

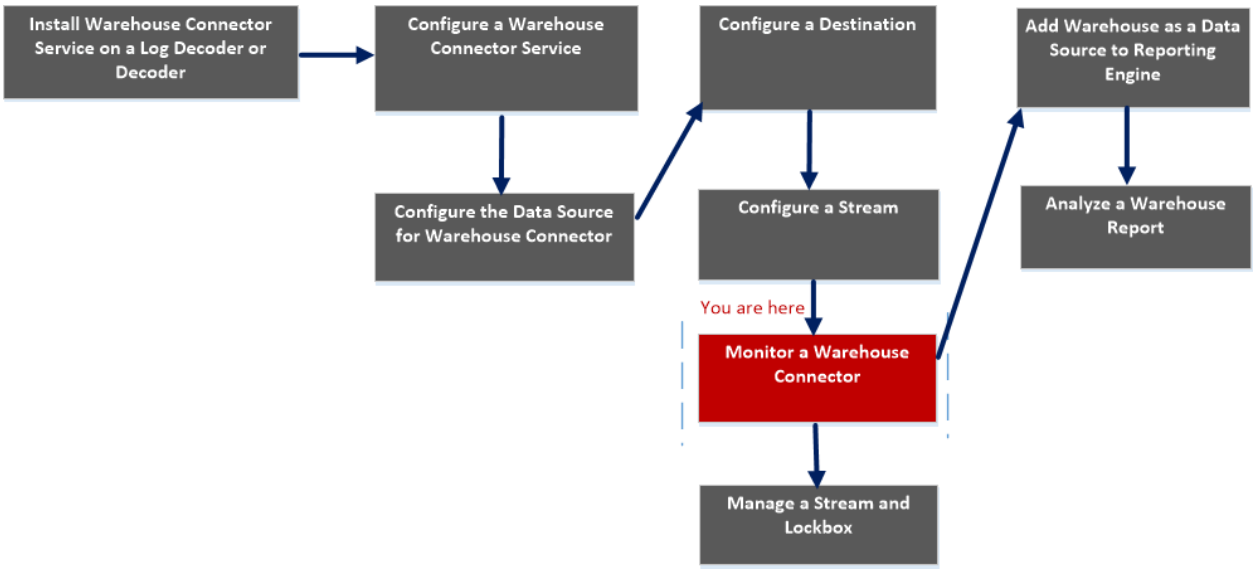
The following table describes the fields in the Add Stream dialog:

Parameter	Description
Stream Name	Type the name of the stream. The stream name may only contain alphanumeric characters and underscores. It cannot exceed 20 characters in length.
Select Destination	Select a destination from the drop-down list.
Select Source	Select a source from the grid at the bottom section of the dialog.
Name	The name of the source.
Address	The address of the source.
Port	The port of the source.
Session ID	The session ID of the source.

Streams Configuration

The Streams tab for a Warehouse Connector in the Services Config view provides a way to manage stream configuration.

Workflow



What do you want to do?

Role	I want to...	Refer to...
Administrator	Install Warehouse Connector Service on a Log Decoder or Decoder	Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid
Administrator	Configure a Warehouse Connector Service	Configure a Warehouse Connector Service
Administrator	Configure the Data Source for Warehouse Connector	Configure the Data Source for Warehouse Connector
Administrator	Configure the Destination using NFS, SFTP, WebHDFS.	Configure the Destination Using NFS Configure the Destination Using SFTP Configure the Destination Using WebHDFS
Administrator	Configure a Stream*	Configure a Stream
Administrator	Monitor a Warehouse Connector*	Monitor a Warehouse Connector

Role	I want to...	Refer to...
Administrator	Add Warehouse as Data Source to Reporting Engine	For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator	Analyze a Warehouse Report	For more information, see "Step 4. Analyze a Warehouse Report" in the *Warehouse Guide*.
Administrator	Manage a Stream and Lockbox	Manage a Stream

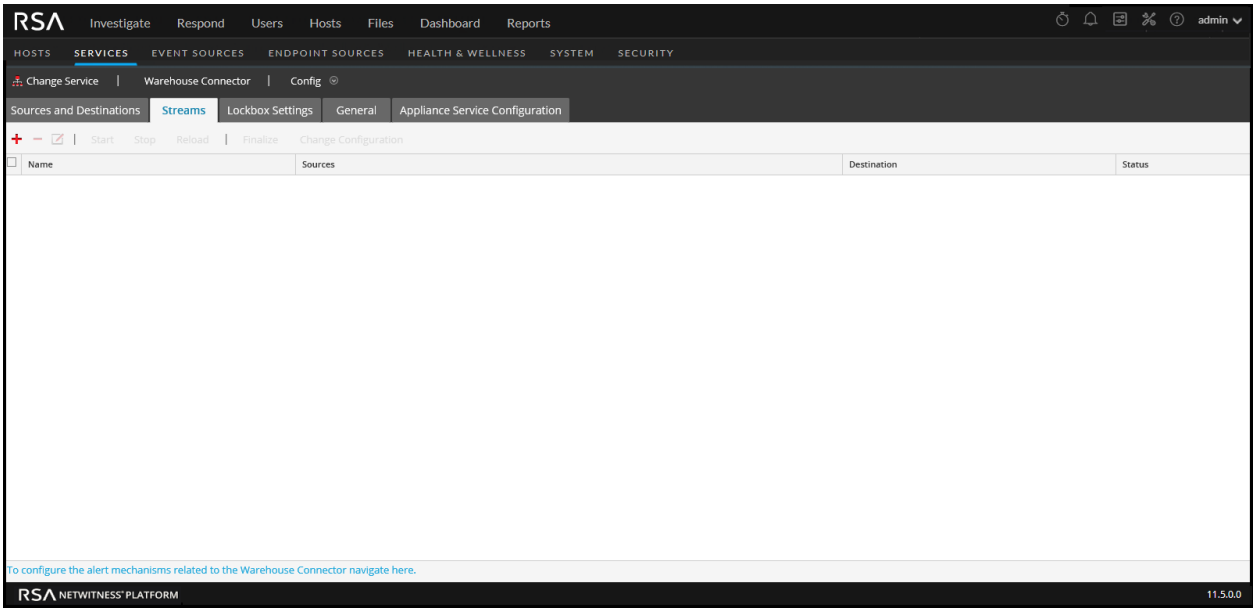
*You can perform this task in the current view

Related topics

Configure a Stream

Quick Look

The following figure shows the Streams tab on the Warehouse Connector Services Config view.



The Streams tab allows you to perform the following:

Features	Description
	Add a stream.
	Delete a stream.
	Edit the stream.

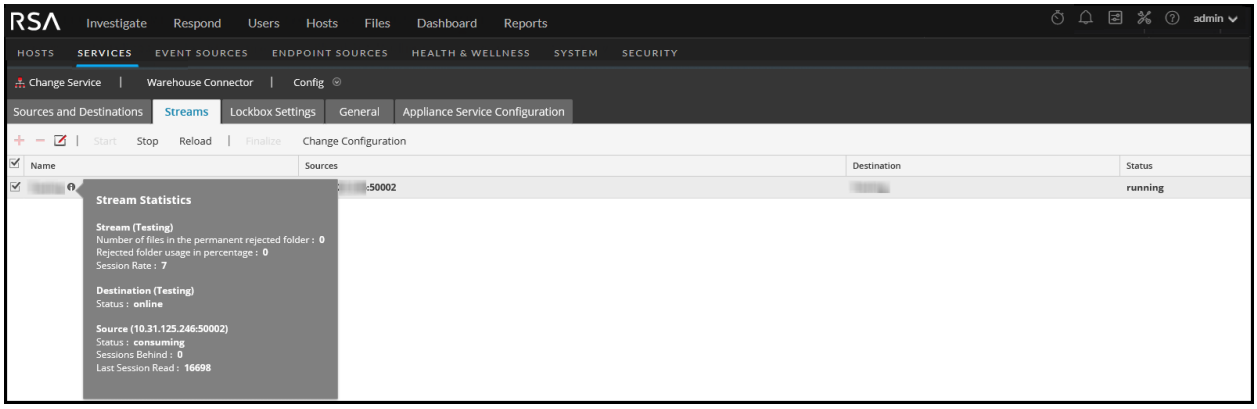
Features	Description
Start	Start the stream.
Stop	Stop the stream.
Finalize	Finalize the stream.
Reload	Reload the stream. If you have added a new meta or if a new meta is added as part of content update to any of the sources, Log Decoder or Decoder, you need to reload the stream for the meta to be visible in the schema for the Reporting Engine. Reloading a stream does not have any impact on the data, but only the new meta list is fetched from the sources.

The following table describes the fields in the Streams tab:

Parameter	Description
Name	Name of the stream.
Sources	The sources associated with the stream.
Destination	The destinations associated with the stream.
Status	Status of the stream.

Stream Statistics

You can view the statistics of a configured stream. Click the icon next to the name of the stream.



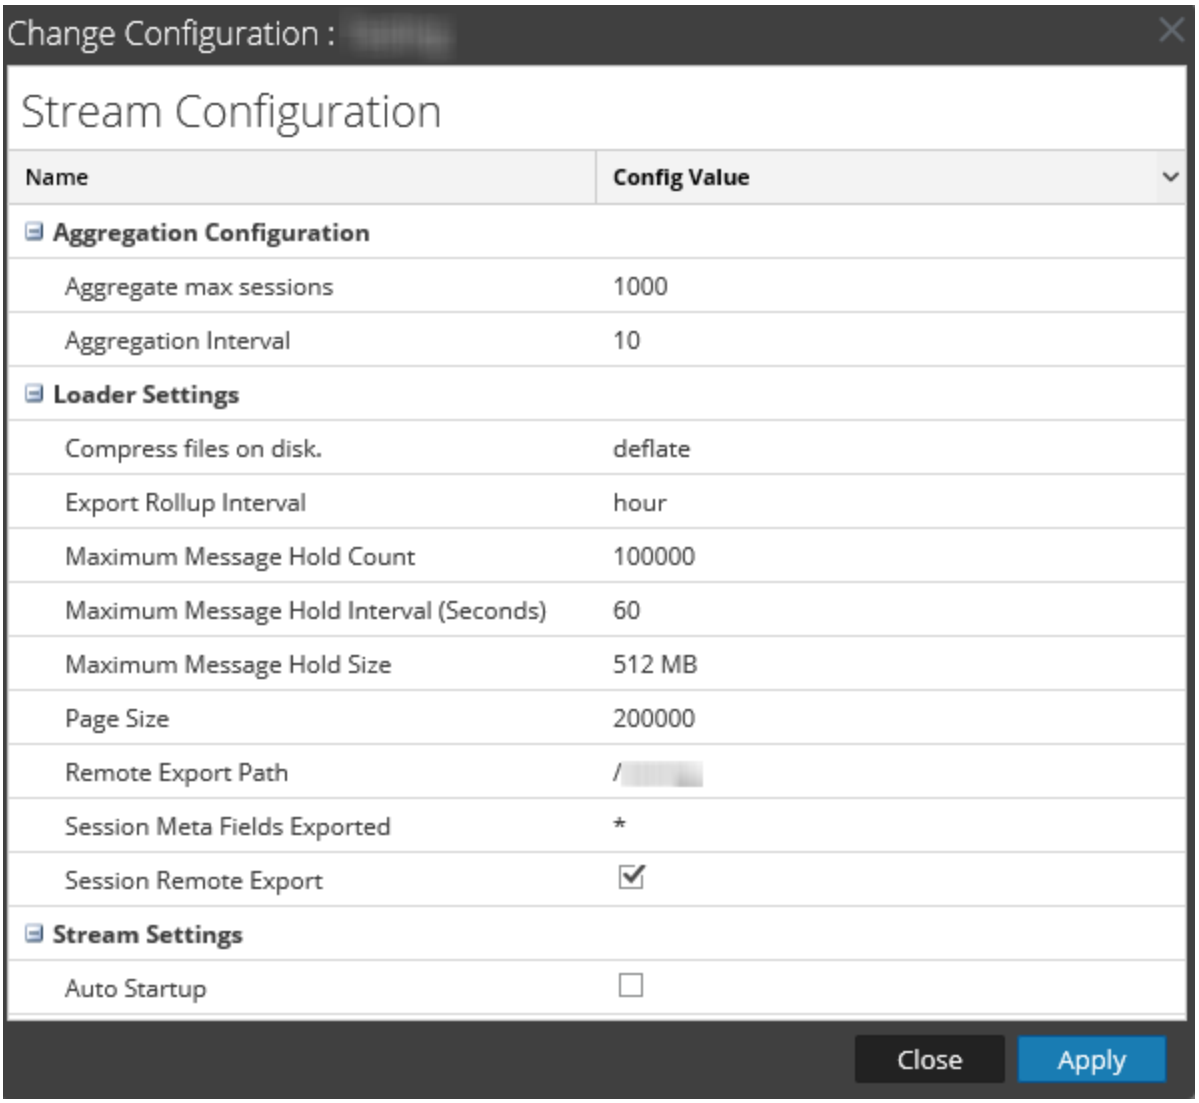
The following parameters are displayed in the Stream Statistics:

Section	Parameter	Description
Stream		

Section	Parameter	Description
	Number of files in the permanent rejected folder	Determines the number of files in the permanent rejected folder (named, **permfail**) in the Warehouse Connector. The permanent rejected folder contains the files that Warehouse Connector failed to write to the destination.
	Rejected folder usage in percentage	Determines the disk usage of the rejected folder.
	Session Rate	Determines the rate at which the session is processed by the Warehouse Connector for the source.
Destination		
	Status	Indicates the status of the destination.
Source		
	Status	Indicate the status of the source.
	Sessions Behind	Determines that number of sessions that needs to be processed by the Warehouse Connector.
	Last Session read	Determines the last session id processed by the Warehouse Connector.

Change Stream Configuration

You can change configuration of a stream in runtime. In the **Streams** tab, click **Change Configuration** to change the configuration of the selected stream.



You can change the following parameters of the Stream Configuration:

Note: If you change the value of any parameter in stream configuration, make sure that you restart the stream.

After upgrading, if the values of Maximum Message Hold Count, Maximum Message Hold Interval and Maximum Message Hold Size are 3000000, 60 and 128 respectively, ensure that you assign the following values to the streams:

- Maximum Message Hold Count - 2400000
- Maximum Message Hold Interval - 600
- Maximum Message Hold Size - 512

You can assign these values by modifying the existing Stream configuration.

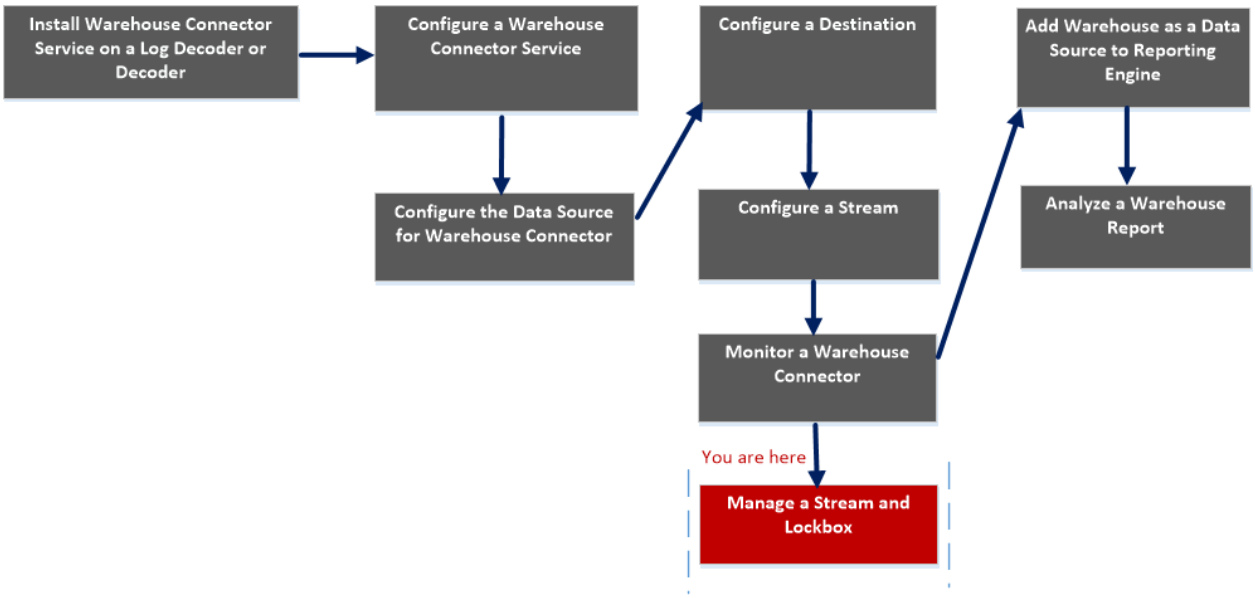
Parameter	Description
	Aggregation Configuration
Aggregate max sessions	Determines the maximum number of sessions in a response for an aggregation request from the Warehouse Connector to the source .
Aggregation Interval	Determines the time between the responses from the source.
	Loader Settings
Compress files on disk	Enable to compress files on disk. Supported values: • Deflate - Provides smaller compressed files and good performance while generating reports. • Off By default, the parameter is set to **deflate**.
Export Rollup Interval	Determines the roll-up interval for export files and also the directory structure the Warehouse Connector writes to the destination. For example: If the parameter is set to: **Value** — **Directory Structure** hour — /rsasoc/v1/[logs | sessions]/data/{year}/{month}/{day}/{hour} minute — /rsasoc/v1/[logs | sessions]/data/{year}/{month}/{day}/{hour}/{minute} day — /rsasoc/v1/[logs | sessions]/data/{year}/{month}/{day} If you change the value of the parameter, ensure that you restart the stream. Recommended value is **hour**.
Maximum Message Hold Count	Determines the maximum number of sessions to store in the memory before processing. **Note:** If you have deployed a Warehouse Connector Virtual Appliance, make sure that you change the default value of the parameter to 800000.
Maximum Message Hold Interval (Seconds)	Determines the maximum time (in seconds) to hold the sessions in memory before processing.
Maximum Message Hold Size	Determines the maximum size for the sessions to store in the memory before processing.
Remote Export Path	Determines the remote local mount point for HDFS (nfs://) and the location to export the data.

Parameter	Description
Page Size	Determines the maximum pages.
Stream Settings	
Auto Startup	Enable to automatically start the stream whenever the Warehouse connector process is restarted. By default, the parameter is set to **off**.

Lockbox Settings

The Lockbox Settings tab for a Warehouse Connector in the Services Config view provide a way to manage the lockbox settings.

Workflow



What do you want to do?

Role	I want to...	Refer to...
Administrator	Install Warehouse Connector Service on a Log Decoder or Decoder	Install Warehouse Connector Service on a Log Decoder or Decoder or Hybrid
Administrator	Configure a Warehouse Connector Service*	Configure a Warehouse Connector Service
Administrator	Configure the Data Source for Warehouse Connector	Configure the Data Source for Warehouse Connector
Administrator	Configure the Destination using NFS, SFTP, WebHDFS.	Configure the Destination Using NFS Configure the Destination Using SFTP Configure the Destination Using WebHDFS
Administrator	Configure a Stream	Configure a Stream
Administrator	Monitor a Warehouse Connector	Monitor a Warehouse Connector

Role	I want to...	Refer to...
Administrator	Add Warehouse as Data Source to Reporting Engine	For more information, see "Add Warehouse as a Data Source to Reporting Engine" in the *Reporting Engine Configuration Guide*.
Administrator	Analyze a Warehouse Report	For more information, see "Step 4. Analyze a Warehouse Report" in the *Warehouse Guide*.
Administrator	Manage a Stream and Lockbox*	Manage a Stream Manage a Lockbox

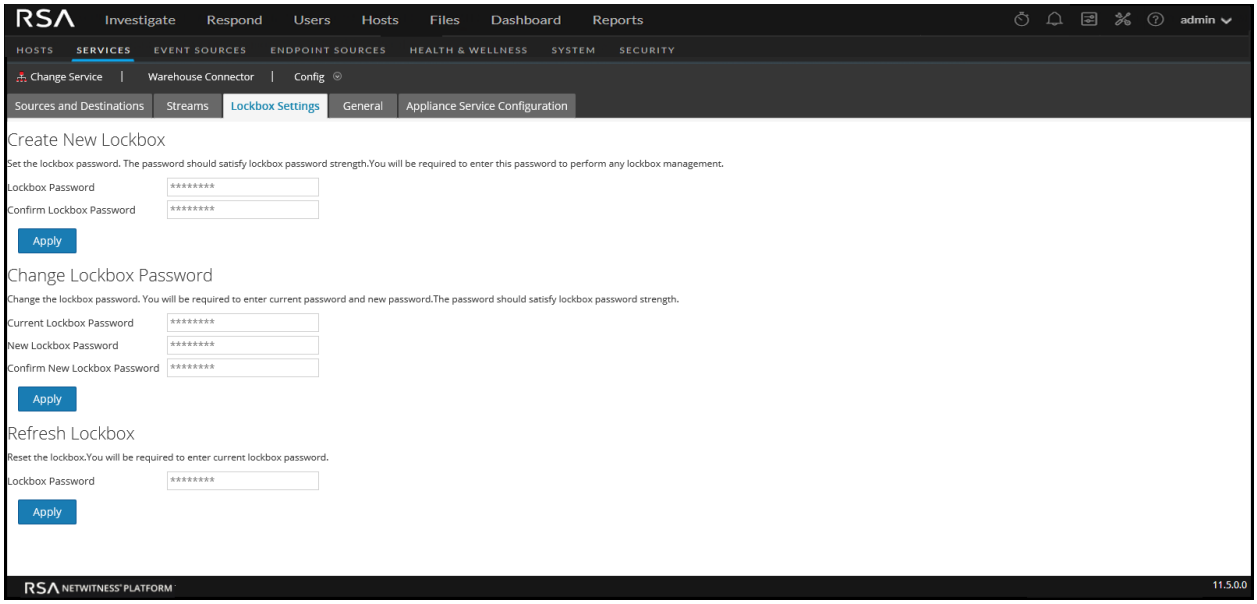
*You can perform this task in the current view

Related topics

- Configure a Warehouse Connector Service
- Manage a Stream

Quick Look

The following figure shows the Lockbox settings tab on the Warehouse Connector Services Config view.



The Lockbox Settings tab allows you to set, change, or refresh the lockbox password of the Warehouse Connector.

NetWitness® Platform XDR

Version 11.x

Windows Legacy Collection

NETWITNESS
Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

Contents

NetWitness Legacy Windows Collection Update & Installation Instructions **4**
Setup Requirements **5**
Update the NetWitness Legacy Windows Collector from 10.6.x to 12.0 ..... **7**
Fresh Install 12.0 Legacy Windows Collector **12**
Configure the Windows Server **16**
Change the Windows Legacy Collector IP Address **17**
Troubleshoot a Fresh or Upgrade Install **18**
Logs to Examine for Information 18
Issues with the Lockbox 18
(Optional) Backup and Restore Legacy Windows Collector **19**
Restore the Windows Legacy Collection Backup after Upgrade 19
Revert Windows Legacy Collection from 12.0 Back to 10.6.4 19
Add a Windows Legacy Collector Host and Service in NetWitness Platform **21**

NetWitness Legacy Windows Collection Update & Installation Instructions

NetWitness Legacy Windows collection collects event data from multiple Windows Event Source domains.

It supports collection from:

- Windows 2003 and earlier event sources
- NetApp ONTAP host evt files

This document contains the following sections:

- Setup Requirements
- Update the NetWitness Legacy Windows Collector from 10.6.x to 12.0
- Fresh Install 12.0 Legacy Windows Collector
- Configure the Windows Server
- Change the Windows Legacy Collector IP Address
- Troubleshoot a Fresh or Upgrade Install
- (Optional) Backup and Restore Legacy Windows Collector
- Add a Windows Legacy Collector Host and Service in NetWitness Platform

Setup Requirements

This section provides the NetWitness Legacy Windows Collector Setup requirements.

Caution: If you are installing or updating to version 11.x, in order to use the Security Analytics Legacy Windows Collector with NetWitness, you need to first install the following windows updates:

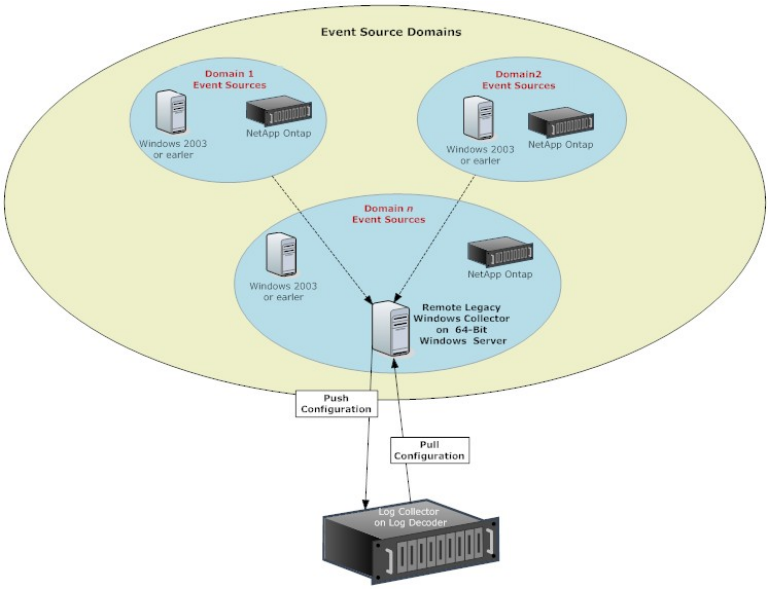
• KB2919355
• KB2919442
• KB2999226
• KB3173424

If these updates are not installed, you will get an error message, and the Legacy Windows Collector will not be installed.

To set up the NetWitness ® Platform Legacy Windows Collector, you need:

- Any of the following physical or virtual systems that can reach the Windows 2003 event source domains:
  - Windows 2008 R2 SP1 64-Bit Server,
  - Windows 2012 Server, or
  - Windows 2016 Server, or
  - Windows 2019 Server
- A minimum of 20% free disk space. For example, you need at least 20 GB of free space if your system drive is 100 GB in size.

IMPORTANT: Do not install the Legacy Windows Collector on a domain controller.



- For NetWitness Platform 11.6 and Later (Upgrade and Fresh Install), you must download and install Visual studio 2019 x64 runtime before running the Windows Legacy Collector installer. You can download the Visual Studio 2019 x64 at https://support.microsoft.com/en-us/help/2977003/the-latest-supported-visual-c-downloads.

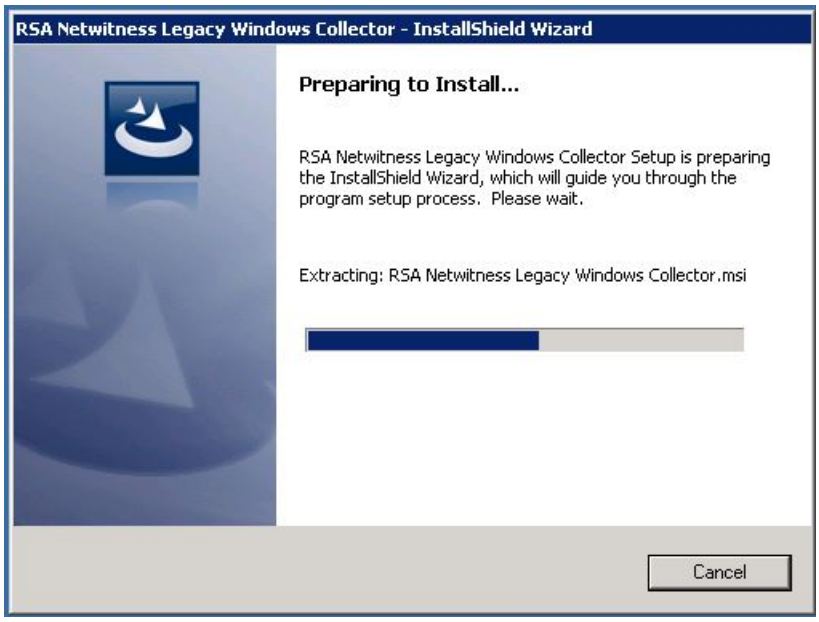
Update the NetWitness Legacy Windows Collector from 10.6.x to 12.0

This section tells you how to update the NetWitness 10.6.x Legacy Windows Collector to 12.0

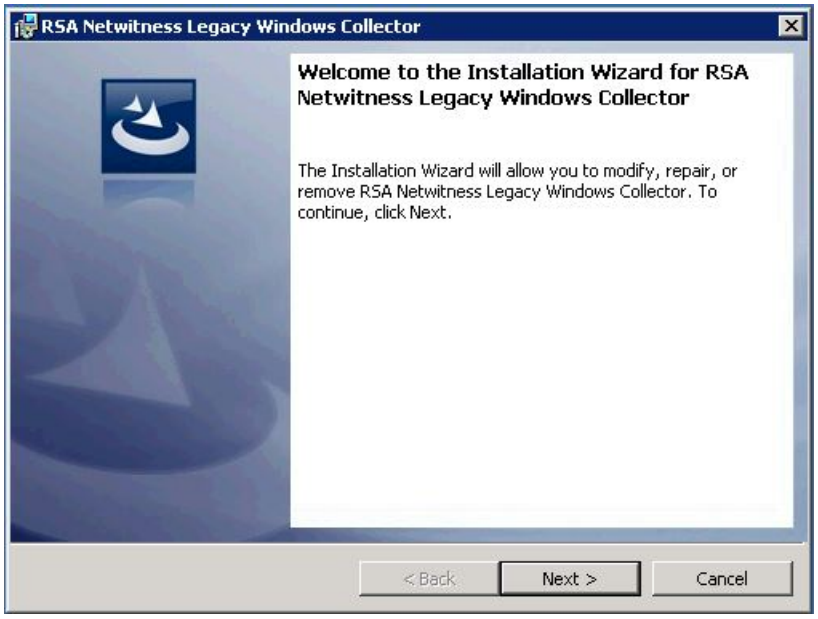
To update the NetWitness 10.6.x Legacy Windows Collector to 12.0 on a Windows 64-Bit server:

1. Depending on your version of NetWitness, navigate to one of the following URLs on NetWitness Community:
   - For NetWitness 11.6.1.1, go to https://community.netwitness.com//t5/netwitness-platform-downloads/rsa-netwitness-platform-11-6-1-1-patch-download/ta-p/643401 and click **NetWitness Platform 11.6.1.1 Upgrade Pack to download the ZIP archive**.
   - For NetWitness 11.6, go to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-6-legacy-windows-collector-download/ta-p/606860 and click **NetWitness Logs & Network 11.6 Legacy Windows Collector** to download the ZIP archive.
   - For NetWitness 11.5, go to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-5-legacy-windows-collector-download/ta-p/572571 and click **NetWitness Logs & Network 11.5 Legacy Windows Collector** to download the ZIP archive.
   - For NetWitness 11.4, go to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-4-legacy-windows-collector-download/ta-p/565269 and click **NetWitness Logs & Packets 11.4 Legacy Windows Collector** to download the ZIP archive.
   - For NetWitness 11.3, go to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-3-legacy-windows-collector-download/ta-p/556065 and click **NetWitness Logs & Packets 11.3 Legacy Windows Collector** to download the ZIP archive.
   - For NetWitness 11.2, go to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-2-legacy-windows-collector-download/ta-p/526984 and click **NetWitness Logs & Packets 11.2 Legacy Windows Collector** to download the ZIP archive.
   - For NetWitness 11.1, go to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-1-legacy-windows-collector-download/ta-p/551450 and click **NetWitness Logs & Packets 11.1 Legacy Windows Collector** to download the ZIP archive.
2. Unzip the downloaded file.
3. Log on to a Windows 2008, 201, 2016, or 2019 Server.
4. Copy **NWLegacyWindowsCollector-*version-number*.exe** to the Windows Server.
5. Right click on **NWLegacyWindowsCollector-*version-number*.exe** and select **Run As Administrator**.

   The Preparing to Install…. page of update installation wizard is displayed.

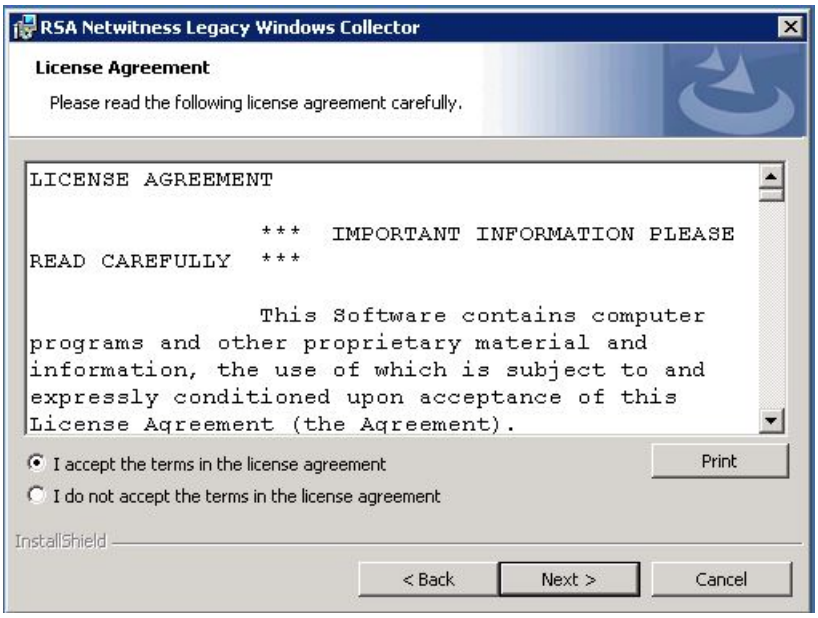


After the update installation program extracts NetWitness Legacy Windows Collector installation files, the **Welcome** page is displayed.



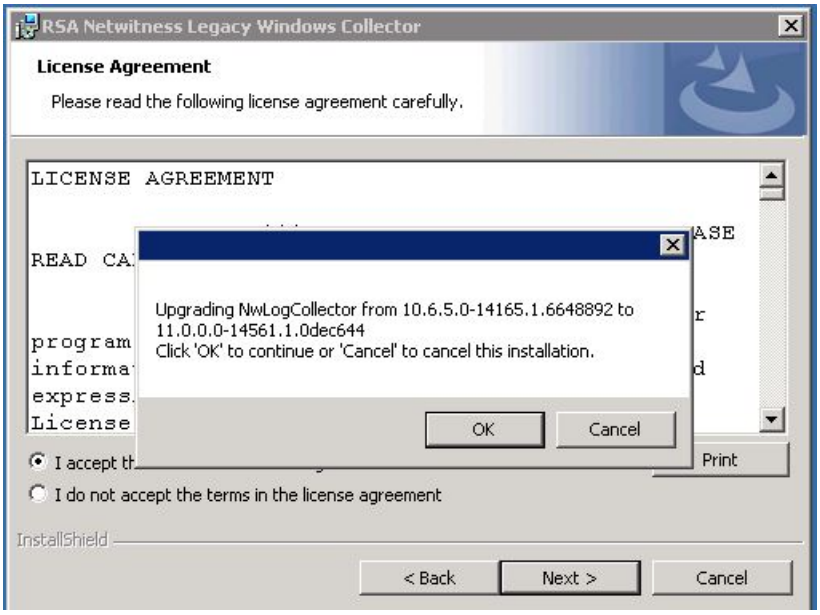
6. Click **Next**.

   The License Agreement page is displayed.



7. Read the License agreement carefully, select the **I accept the terms in the license agreement** radio button, and click **Next**.

   Before it starts the update, the wizard asks if you want to continue or cancel the installation of the update.



8. Click **OK** to continue installing the update.
9. Click Install.

   The Installation screens for the Legacy Windows Collector page is displayed.

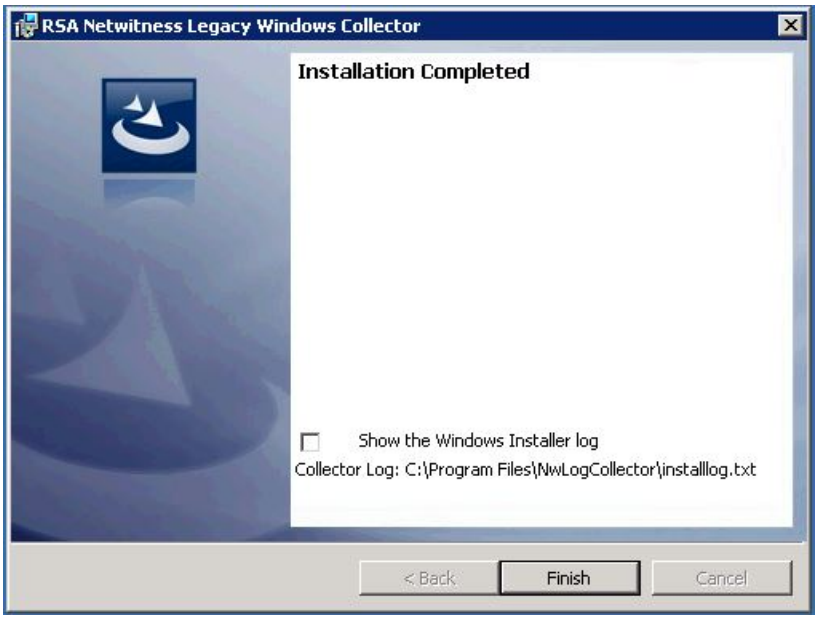




After the update installation completes, the **Next** button becomes active.

10. Click **Next**.

    The Installation Completed page is displayed.



11. (Optional) If you want to review a log of the update installation, select the **Show the Windows Installer** log checkbox.
12. Click **Finish**.
13. Reboot the machine.

This completes the update of the Legacy Windows Collector to NetWitness 12.0.

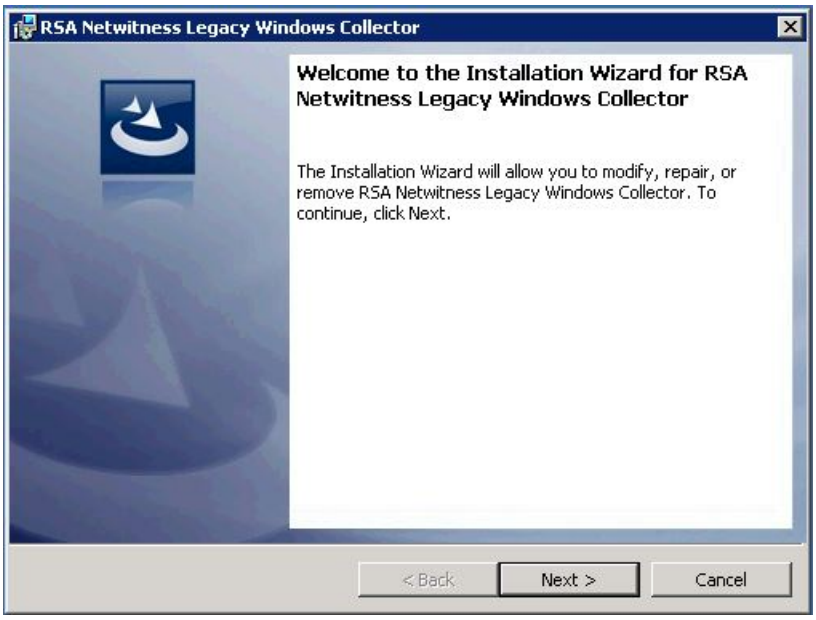
Fresh Install 12.0 Legacy Windows Collector

This section describes how to install the 12.0 Legacy Windows Collector on a Windows 2008, 2012, 2016, or 2019 64-Bit server

To install the NetWitness Legacy Windows Collector on a Windows 2008, 2012, 2016, or 2019 64-Bit server:

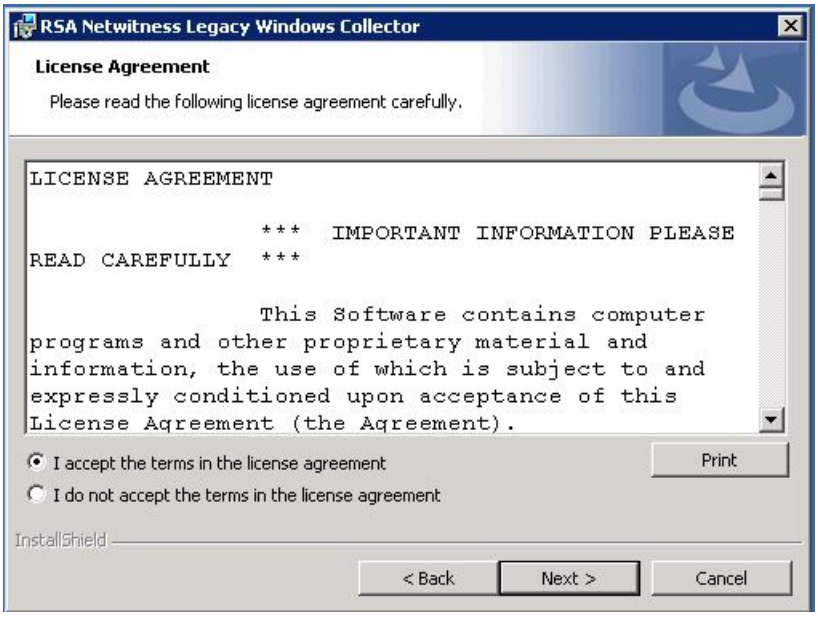
1. Depending on your version of NetWitness Platform, navigate to one of the following URLs on NetWitness Community:
   - For NetWitness 11.6.1.1 only, navigate to https://community.netwitness.com//t5/netwitness-platform-downloads/rsa-netwitness-platform-11-6-1-1-patch-download/ta-p/643401 on NetWitness Community. Click NetWitness Logs & Network 11.6.1.1 Legacy Windows Collector to download the ZIP archive.
   - For NetWitness 11.6, navigate to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-6-legacy-windows-collector-download/ta-p/606860 on NetWitness Community. Click NetWitness Logs & Packets 11.6 - Legacy Windows Collector to download the ZIP archive.
   - For NetWitness 11.7, navigate to https://community.netwitness.com/t5/netwitness-platform-downloads/rsa-netwitness-11-7-legacy-windows-collector-download/ta-p/644208 on RSA link. Click RSA NetWitness Platform 11.7 Legacy Windows Collector to download the ZIP archive.
   - For NetWitness 12.0, navigate to https://community.netwitness.com/t5/rsa-netwitness-platform-staged/netwitness-12-0-legacy-windows-collector-download/ta-p/684864 on RSA link. Click RSA NetWitness Platform 12.0 Legacy Windows Collector to download the ZIP archive.
2. Unzip the downloaded file.
3. Copy the **NWLegacyWindowsCollector-*version-number*.exe** to the Windows Server.
4. Right click on the **NWLegacyWindowsCollector-*version-number*.exe** and select **Run As Administrator**.

   The **Welcome** page of installation wizard is displayed.



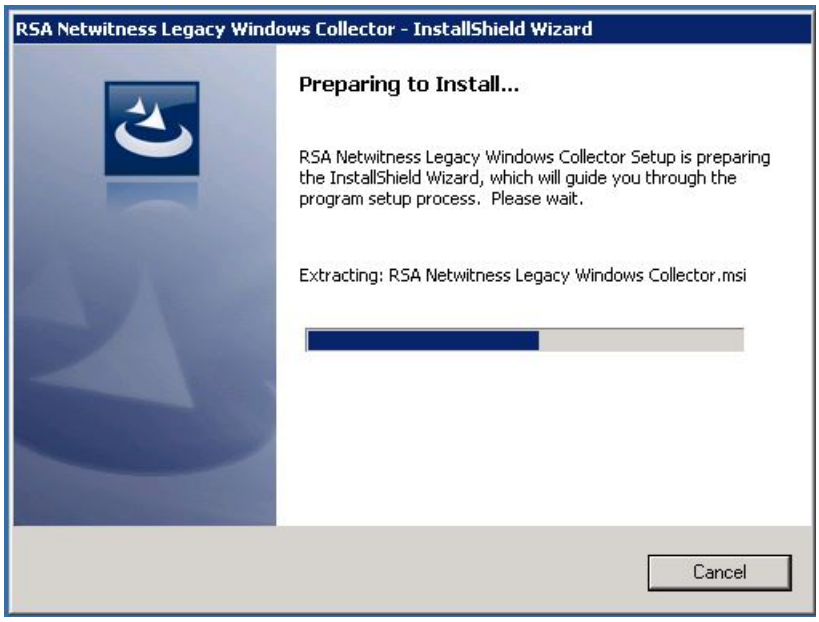
5. Click **Next**.

   The License Agreement page is displayed.



6. Read the License agreement carefully, select the **I accept the terms in the license agreement** radio button, and click **Next**.

   The Ready to Install the Program page is displayed.



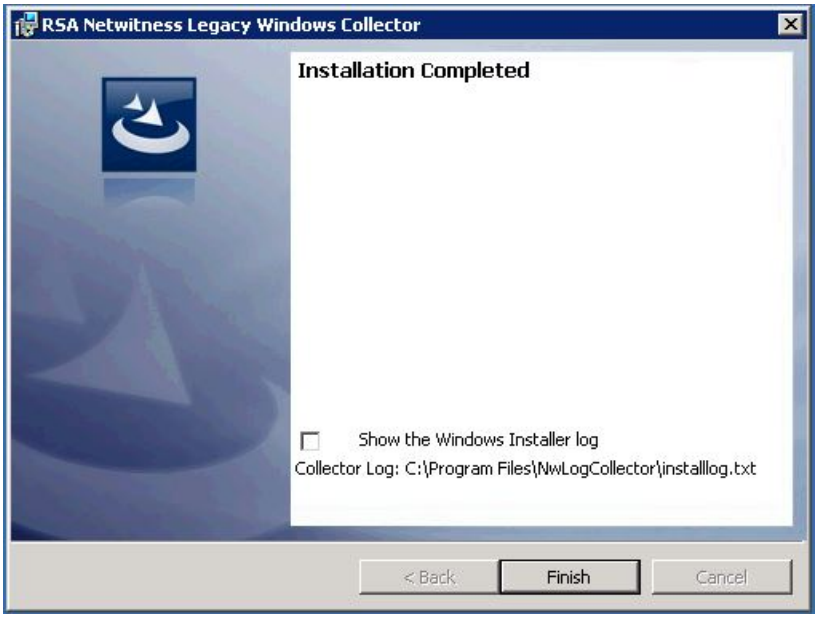
7. Click **Install**.

   The Installation screens for the Legacy Windows Collector page are displayed.





The Installation Completed page is displayed.



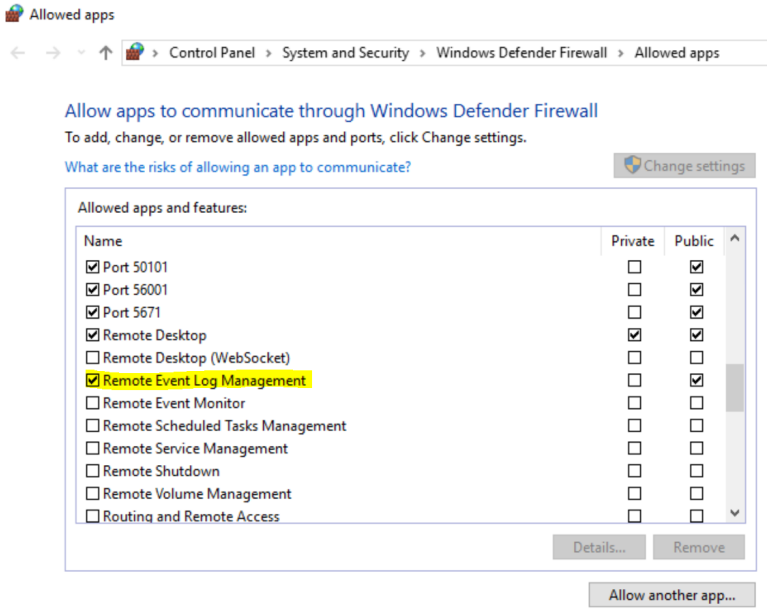
8. (Optional) If you want to review a log of the installation, select the **Show the Windows Installer** log checkbox.
9. Click **Finish**.
10. Reboot the machine.

This completes the installation of the 12.0 Legacy Windows Collector. Please refer to the **Windows Legacy and NetApp Collection Configuration Guide** on NetWitness Community for instructions on how to configure Legacy Windows collection in NetWitness.

Configure the Windows Server

For the NetWitness to communicate with the Windows Server, you need to allow Remote Event Log Management on the Windows Server.

1. On the Windows Server, in Services, start the Remote Registry Service.
2. In Firewall, enable Remote Event Log Management for your network, as shown below.



Change the Windows Legacy Collector IP Address

Note: The procedures in this section apply to NetWitness 11.5 and later only.

On occasion, you may need to change the IP address of your Windows Legacy Collector. You may also need to edit any Destination Groups that you have configured.

Change WLC IP Address

The following procedure describes how to change the IP address for your system.

1. Log onto the Windows Legacy Collector system and manually change the IP address on the system.
2. In the UI, confirm that the Log Collector service corresponding to the WLC system shows up in error (Red). It might take some time for it to reflect the changed status.
3. On the NetWitness Server, use the **nw-manage** utility to view the host information for the WLC using the following command:

   ```
   nw-manage --list-hosts
   ```

   Sample output from running the command is shown here:

   ```
   {
   "id" : "fdb8150c-e040-459e-8cc5-3c60ec2c65ae",
   "displayName" : "WLC-HOST-104",
   "hostname" : "10.101.216.102",
   "ipv4" : "10.101.216.102",
   "ipv4Public" : null
   } ]
   ```

   You use the value of **"id"** from your output in the following step.
4. Use the **nw-manage** utility to change the IP address of the WLC. For the **host-id** argument, use the value for the **"id"** that you noted from step 3. For the **ipv4** value, use the new IP Address to which you are changing.

   ```
   nw-manage --update-host --host-id "fdb8150c-e040-459e-8cc5-3c60ec2c65ae" --
   ipv4 10.101.216.105
   ```
5. After you see the message that the previous command ran successfully, go to the NetWitness Server UI and verify that the WLC service is running without any errors.

Edit Destination Groups For Log Collectors and VLCs

The Windows Legacy Collector is often configured with Destination Groups to forward events to Log Collectors or Virtual Log Collectors. If the IP address of any such Destination LC or VLC is changed, the Windows Legacy Collector can no longer forward events. To remediate this, you must edit the Destination groups for the WLC, making sure to select the new LC or VLC IP Address.

Troubleshoot a Fresh or Upgrade Install

Logs to Examine for Information

Refer to the following log files if you need to troubleshoot problems:

- `%systemDrive%\Netwitness\ng\logcollector\MessageBroker.log`
- `%systemDrive%\Program Files\NwLogCollector\installlog.txt`

Run `C:\Program Files\NwLogCollector\ziplogfiles.vbs` to generate the ***`hostname`***`_WLCversion_timestamp.zip` that contains all the log files and other information needed for troubleshooting.

Issues with the Lockbox

When you create a lockbox password on a new Windows Legacy Collector, you might see the following error:

```
failed to set secure storage password: failed to create lockbox: The Lockbox or cryptography library could not be found.
```

This can occur if you are running Windows Legacy Collector version 11.x.

If you encounter this issue, download and install both of the following redistributable packages:

- Visual C++ 2010: https://www.microsoft.com/en-us/download/details.aspx?id=14632
- Visual C++ 2012: https://www.microsoft.com/en-us/download/details.aspx?id=30679

(Optional) Backup and Restore Legacy Windows Collector

This section tells you how to upgrade from 10.6.4 to NetWitness 12.0 for the Legacy Windows Collector.

Note: You only need to do this if you are changing the Windows VM where you run the Windows Legacy Collector.

During upgrade to NetWitness 12.0, the backup script for the Windows Legacy Collector is invoked automatically, and creates the 10.6.4 configuration and run-time backups. After the 12.0 installation is completed, run the Restore script to restore the configuration and run-time files for the updated Windows Legacy Collection.

Restore the Windows Legacy Collection Backup after Upgrade

To restore the Windows Legacy Collection setup on a newly upgraded NetWitness 11 platform:

1. On the Windows Legacy Collector, open a command prompt window.
2. Navigate to **C:\Program Files\NwLogCollector**, where the scripts are stored.
3. Run the following commands for restoring a backup:
   - Backup configuration files: `WLC-Restore.bat "Config-bkup_`***`timestamp`***`.zip"`
   - Backup run-time files: `WLC-Restore.bat "Runtime-bkup_`***`timestamp`***`.zip"`
4. Once the restore is completed, set the lockbox SSV to use the password that you created during 10.6.4 setup.
   a. In the **Security Analytics** menu, select **Services**, then select your Windows Legacy Collector and choose **Explore**.
   b. From the left navigation pane, expand **logcollection > properties > crypto**.
   c. Run the following command: `op=setssv pw=`***`password_for_10.6.x_lockbox`***, and hit **Send**.

Revert Windows Legacy Collection from 12.0 Back to 10.6.4

To revert the Windows Legacy Collection setup from 12.0 back to 10.6.4:

1. Uninstall the 12.0 Setup. Note the location of the backup folder created by the system during the uninstall procedure.

2. Install the 10.6.4 version of the Windows Legacy Collector.
3. Navigate to **C:\Program Files\NwLogCollector**, where the scripts are stored.
4. Run the Restore script from backup folder present in **C:\Program Files\NwLogCollector** to restore the configuration and run-time setup on the 10.6.4 Windows Legacy Collector.
   - Backup configuration files: `WLC-Restore.bat "Config-bkup_`***`timestamp`***`.zip"`
   - Backup run-time files: `WLC-Restore.bat "Runtime-bkup_`***`timestamp`***`.zip"`
5. Once the restore is completed, set the lockbox SSV to use the password that you created during 10.6.4 setup.
   a. In the **Security Analytics** menu, select **Services**, then select your Windows Legacy Collector and choose **Explore**.
   b. From the left navigation pane, expand **logcollection > properties > crypto**.
   c. Run the following command: `op=setssv pw=`***`password_for_10.6.x_lockbox`***, and hit **Send**.

Add a Windows Legacy Collector Host and Service in NetWitness Platform

For this version of the Windows Legacy Collector, NetWitness has provided a script that replaces the manual steps of adding a Windows Legacy Collector host and service in the NetWitness UI.

To create a Windows Legacy Collector Host and Service in NetWitness:

1. SSH to your NetWitness server.
2. Run the following command:

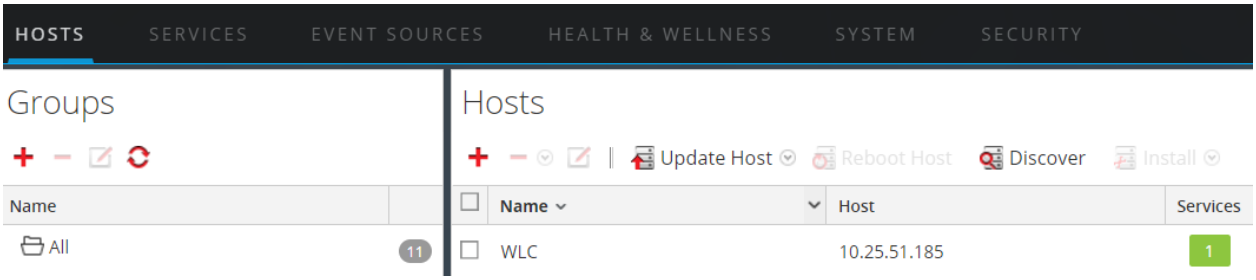
   ```
   wlc-cli-client --host-display-name hostDisplayName --service-display-name
   serviceDisplayName --host WLChostIPAddress --port 50101 --use-ssl false
   ```

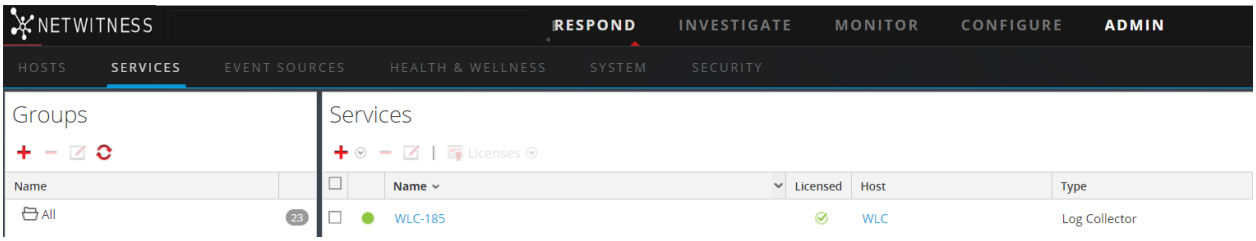
   The parameters are explained below:

   - **--host-display-name**: the name for the host as it is displayed in the NetWitness Hosts page
   - ***--service-display-name***: the name for the host as it is displayed in the NetWitness Services page
   - **--host**: the IP address for the Windows Legacy Collector
   - **--port**: the port NetWitness uses to communicate with the Windows Legacy Collector. The recommended value is 50101.
3. You will be prompted to supply the following information:
   - **Windows Log Collector REST Username** and **Windows Log Collector REST Password**: you must supply admin credentials for the Windows Legacy Collector.
   - **Security Server Username** and **Security Server Password**: you must supply admin credentials for NetWitness.

Note: If the **Security Server Password** contains any special character, you must use backslash (\) before the special character. For example, if the password is **netwitness@123**, enter the password as **netwitness\@123**.

After you complete this procedure, you should see the Windows Legacy Collector Host and Service as shown in the following screenshots.




NETWITNESS
RESPOND
INVESTIGATE
MONITOR
CONFIGURE
ADMIN
HOSTS
SERVICES
EVENT SOURCES
HEALTH & WELLNESS
SYSTEM
SECURITY
Groups
Name
All
23
Services
Licenses
Name
Licensed
Host
Type
WLC-185
WLC
Log Collector

NetWitness® Platform XDR

Version 12.0

Workbench Configuration Guide

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

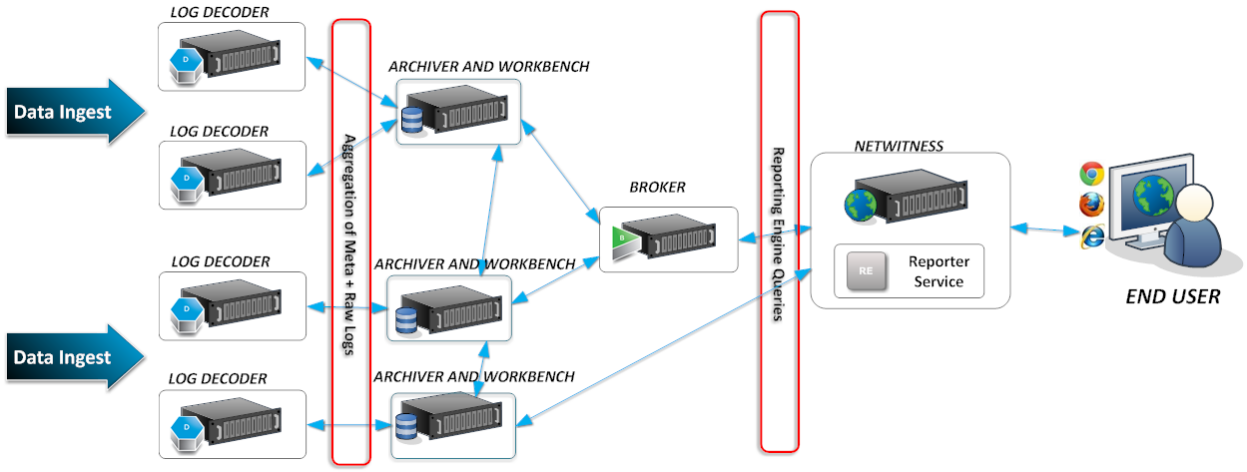
Contents

Workbench Overview **4**
Workbench Configuration Procedures **5**
Adding Workbench Service as a Data Source to Broker **7**
Adding Workbench as a Data Source to Reporting Engine **10**
Managing Collections **12**
Mount Archiver Directories 12
Create a Collection 12
Delete a Collection 14
Example Procedure: How to Restore a Collection for Reporting and Investigation 16
Investigate a Collection 17
View Workbench Collection Statistics 19
View Workbench Logs 20
Services Config View - Workbench **21**
Services Config View - Collections Tab **24**
Services Config View - General Tab **27**
System Configuration Panel 28
Workbench Configuration Panel 29
Troubleshooting **30**

Workbench Overview

NetWitness Workbench service allows collections to be created with restored data that was saved offline from an Archiver. Once the data is copied and saved into a collection, it can be analyzed from Investigation and Reporting.

The following diagram depicts the architecture of a NetWitness network that implements the Workbench.

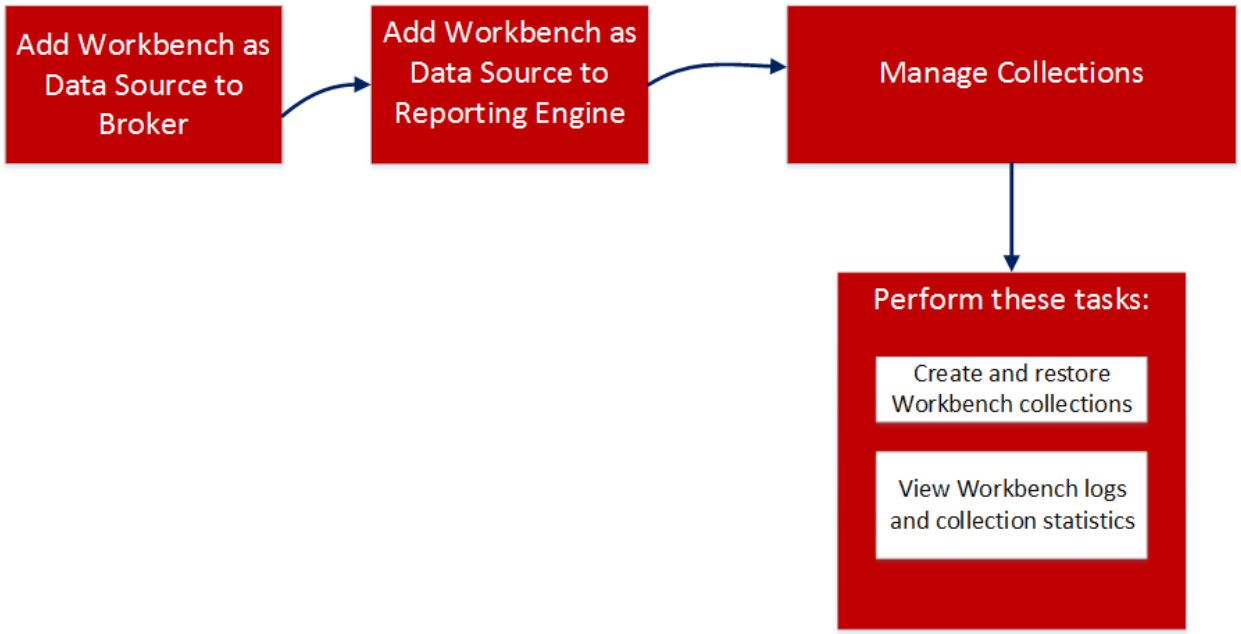


Workbench Configuration Procedures

Note: While NetWitness 11.0 and later continues to support the Workbench, and some customers may have configured Workbench to handle restoring of data. The best practice for restoring data is to use the Archiver to configure archival and restoring of data. For more information, see *Archiver Configuration Guide*.

Workflow

These are the basic steps for configuring and managing a Workbench service.



1. Add a Workbench service as a data source to Broker (see Adding Workbench Service as a Data Source to Broker).
2. Add a Workbench service as a data source to Reporting Engine (see Adding Workbench as a Data Source to Reporting Engine
3. Manage collections on a Workbench service (see Managing Collections).
4. Investigate a Workbench (see Managing Collections).

Prerequisites

Before configuring the Workbench service, you must:

- Add the NetWitness workbench service to the host in your network environment. For more information, see Workbench Overview.

- Install the NetWitness Workbench host in your network environment. For more information, see *Host and Services Getting Started Guide*.

The steps to configure the Workbench service are:

1. Adding Workbench Service as a Data Source to Broker
2. Adding Workbench as a Data Source to Reporting Engine

After configuration is complete, you can create and manage collections as described in Managing Collections.

Adding Workbench Service as a Data Source to Broker

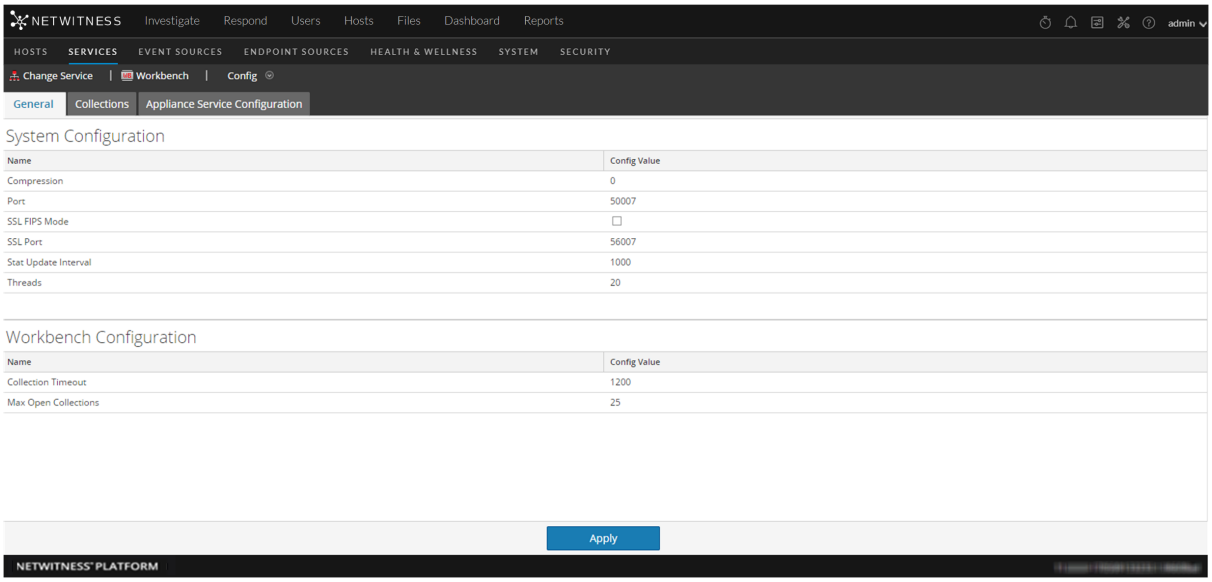
Prerequisites

Before adding the Workbench service, you must:

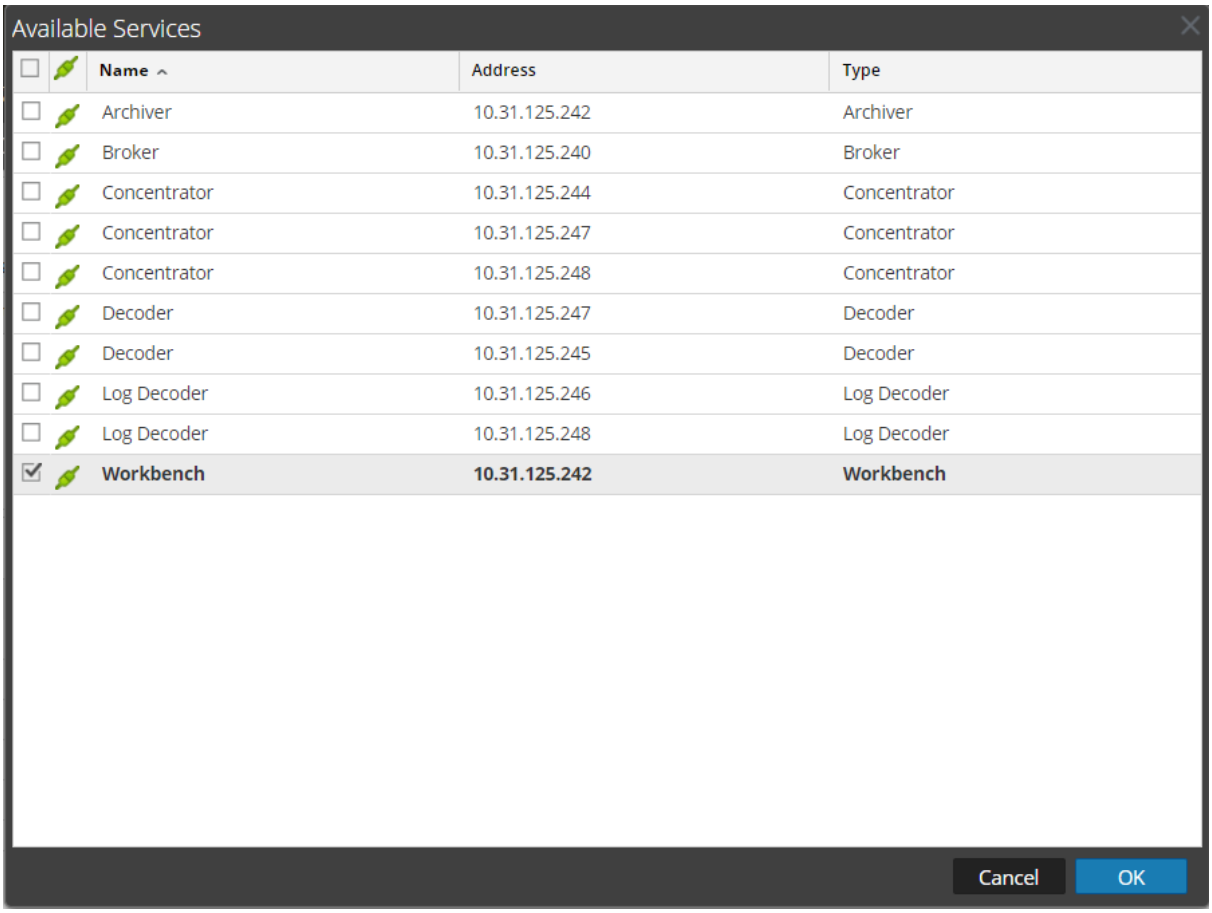
- Install the Workbench service on the Archiver appliance.
- Add a collection on the workbench service.

To add the Workbench service as a data source on the Broker:

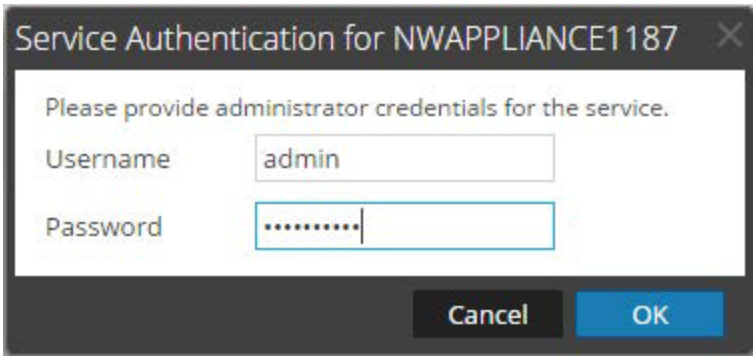
1. Go to **(Admin) > Services**.
2. Select a **Broker** service, and select **> View > Config.**
   The Service Config view is displayed with General tab open.



3. In the Aggregate Services section, Click .
   The Available Services dialog is displayed.

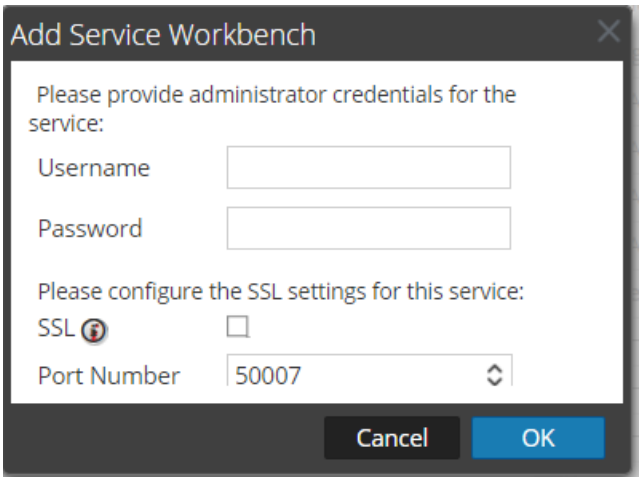


4. Select the Workbench service and click **OK**.
5. If the Workbench service is using a Trust Model, a Service Authentication dialog for the selected service is displayed.



6. Type the username and password for admin credentials for the service and click **OK**.

   The Add Service Workbench dialog is displayed.



7. Type the username and password for admin credentials for the service and click **OK**.

   The workbench service is now added as a data source to the Broker and is listed in the NWDATA Sources list.

Note: You must perform this procedure for each collection.

Adding Workbench as a Data Source to Reporting Engine

Prerequisites

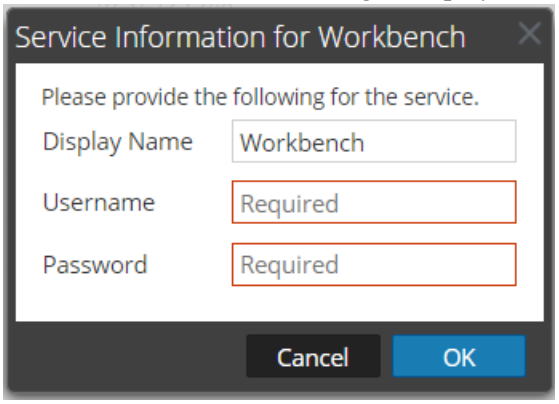
These are the tasks required before adding the Workbench as a data source to Reporting:

1. Add the Reporting Engine as a service to your NetWitness deployment.
2. Add the Workbench as a service to your NetWitness Archiver host (if not already installed).

Note: Adding Workbench collections as a data source to Reporting Engine depends on a trusted connection. If the Workbench is established with a trusted connection, you must manually add Workbenchcollections as a source to the Reporting Engine.

To associate the Workbench data source with the Reporting Engine:

1. Go to  **(Admin) > Services**.
2. Select a **Reporting Engine** in the Services panel. Select  **View > Config**.
   The Services Config view is displayed with General tab Open.
3. Click the **Sources** tab, click  in the toolbar and select **Available Services**.
   The Available Services dialog is displayed.
4. Select a Workbench service and Click **OK**.
   The Service Information dialog is displayed.



5. Enter User Name and Password.
   - Required if the Workbench service is Trusted.
   - Optional if the Workbench service is not trusted (added manually).
6. Click **OK**.

7. Select **Collection** in Add a Collection from Workbench dialog.
8. Click **OK**.

Result

You can now create reports on the data collected by the Workbench.

Managing Collections

An Administrator can create and delete Workbench collections and view Workbench statistics and logs. This topic provides all of these procedures and an example procedure for restoring a collection for Reporting and Investigation.

- Mount Archiver Directories
- Create a Collection
- Delete a Collection
- Investigate a Collection
- View Workbench Collection Statistics
- View Workbench Logs

Mount Archiver Directories

If data is in offline storage or cold-tier storage, you need to mount the Archiver directories in order to restore the data for reporting and investigation purposes:

1. Go to  **(Admin) > Services**.
2. Select an **Archiver** from the Services panel and select  **> View > Explore**.
   The Explorer view for the Archiver is displayed
3. Right-click on the Database node in left-hand tree and select **properties** to open them in the right-hand panel.
4. Run the **manifest** command for a time range, for example, 2017-April-01 to 2017-April-10.
   The search returns all files that need to be restored for the selected query.

Create a Collection

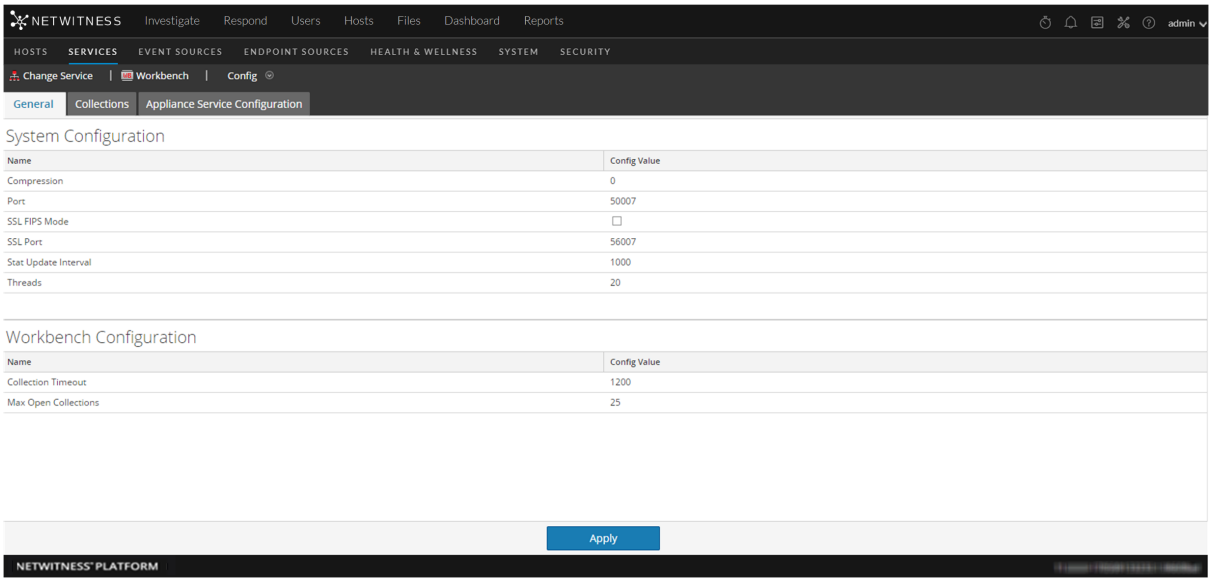
Administrators can create collections of restored data from a backup or from an existing set of data.

Note: You can point the source path to the location of the database files and the restore command copies them to the workbench. You need to mount those directories to the Archiver (where the Workbench is installed) before a restoration collection can be created.

To create a collection using data restored from the backed up data or existing subset of data:

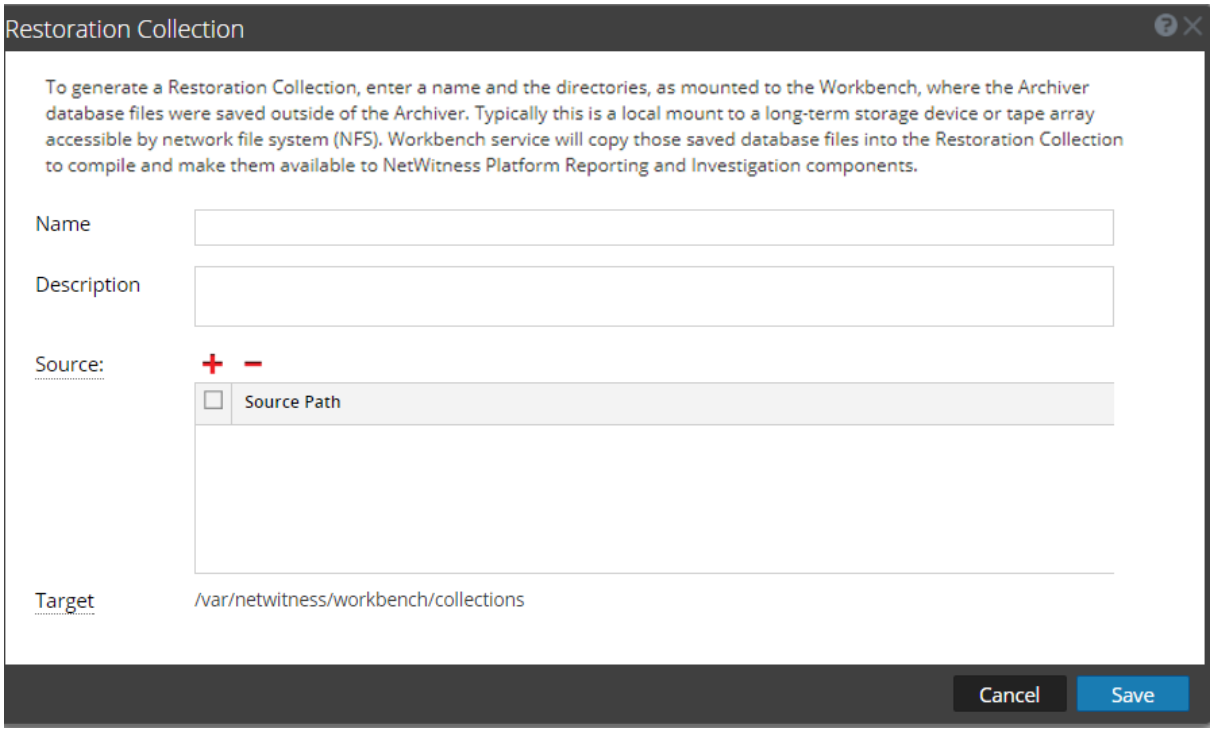
1. Go to  **(Admin) > Services.**

2. In the **Services** view, select a **Workbench**, then select  **> View > Config**.
   The Services Config view is displayed with the General tab open.



3. Click the **Collections** tab.
   The Collections panel is displayed.

4. Click  in the toolbar.
   The Restoration Collection dialog is displayed.



5. Provide the following information:

- **Name:** Name of the Workbench collection that you want to restore.
- **Source:** Location where the Archiver database files have been moved from cold storage.

Note: Target is the location where the collection is created.

6. Click **Save** to restore the collection.

Note: If the source path provided to create the restoration collection does not exist, the following error message is displayed:

```
The source path does not exist '/xxx/xxx/'.
```

If there is insufficient storage to restore your collection, the following error is displayed:

```
Error during disk space checking. Insufficient disk space in location
'/xxx/xxx'.
```

The Schedule Job dialog is displayed with the following message:

```
Restoring data into a new collection. Check the jobs page for progress.
```

7. Click the **Jobs** icon  in the NetWitness toolbar to expand the list of restoration collection jobs with their current status.

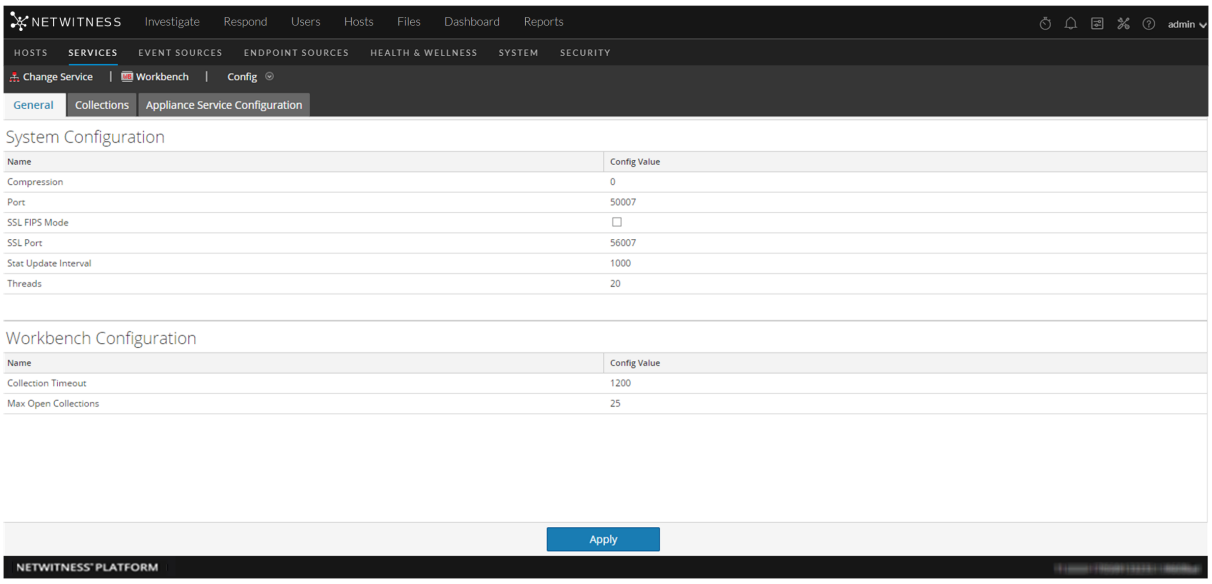
Note: Restoring a collection that is larger than 550 GB may take several hours to process.

Delete a Collection

Administrators can delete collections from the Workbench service.

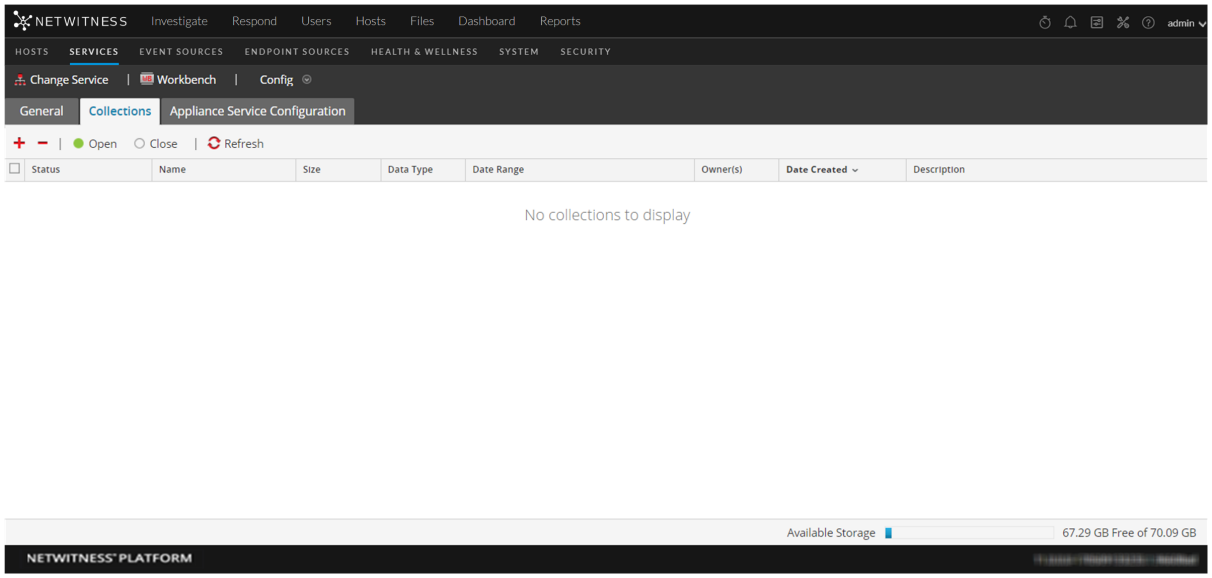
Perform the following steps to delete a collection:

1. Go to  **(Admin) > Services.**
2. From the **Services** view, select a **Workbench** and click  **> View > Config**.
   The Services Config view opens with the General tab displayed.



3. Select the **Collections** tab.

   The Collections panel is displayed.



4. In the Collections panel, select the collection that you want to delete.
5. Click − from the toolbar.

   A warning dialog requests confirmation.
6. If you want to delete the collection, click **Yes**.

   The collection is removed from the Workbench service.

Example Procedure: How to Restore a Collection for Reporting and Investigation

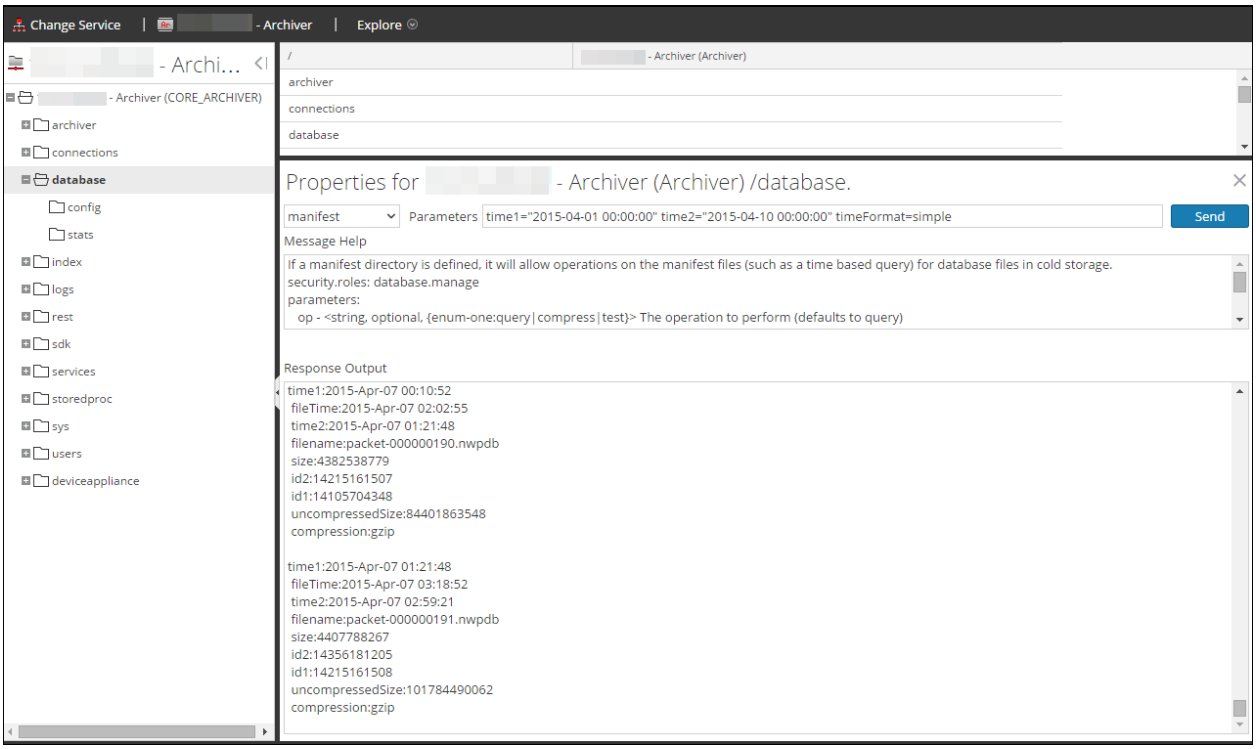
The following steps illustrate how to restore data for reporting and investigation purposes that is in offline storage or cold-tier storage. In the following example, data is restored for the time range beginning on 2015-April-01 through 2015-April-10.

To restore data for reporting and investigation purposes:

1. Go to **(Admin) > Services**.
2. Select the **Archiver** from the **Services** panel and Click **> View > Explore**.
   The Explorer view for Archiver is displayed
3. Right click on Database node in left-hand tree and select **properties** to open them in the right-hand panel.
4. Run the **manifest** command for the selected time range 2015-April-01 to 2015-April-10.
   The search returns all files that need to be restored for your selected query.

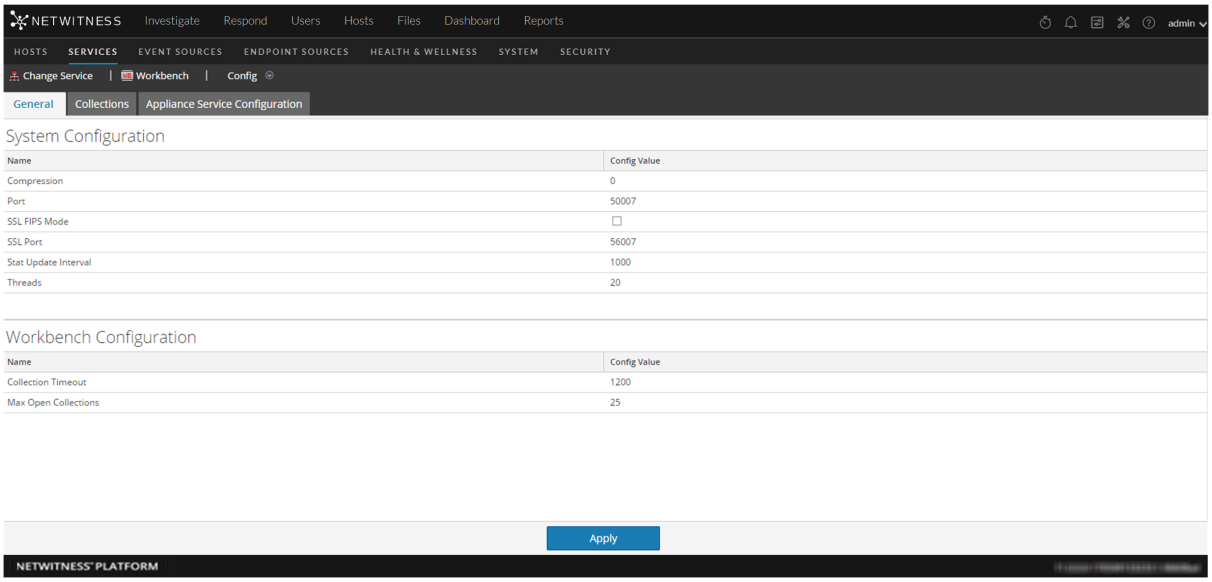
Example Search:

```
time1="2015-04-01 00:00:00" time2="2015-04-10 00:00:00" timeFormat=simple
```



6. Go to **(Admin) > Services**.
7. In the **Services** view, select a **Workbench**, then select **> View > Config**.

The Services Config view is displayed with the General tab open.



8. Select the **Collections tab.**
9. Create a restoration collection with the source path pointing to files listed in the manifest command output.
10. Save the collection.

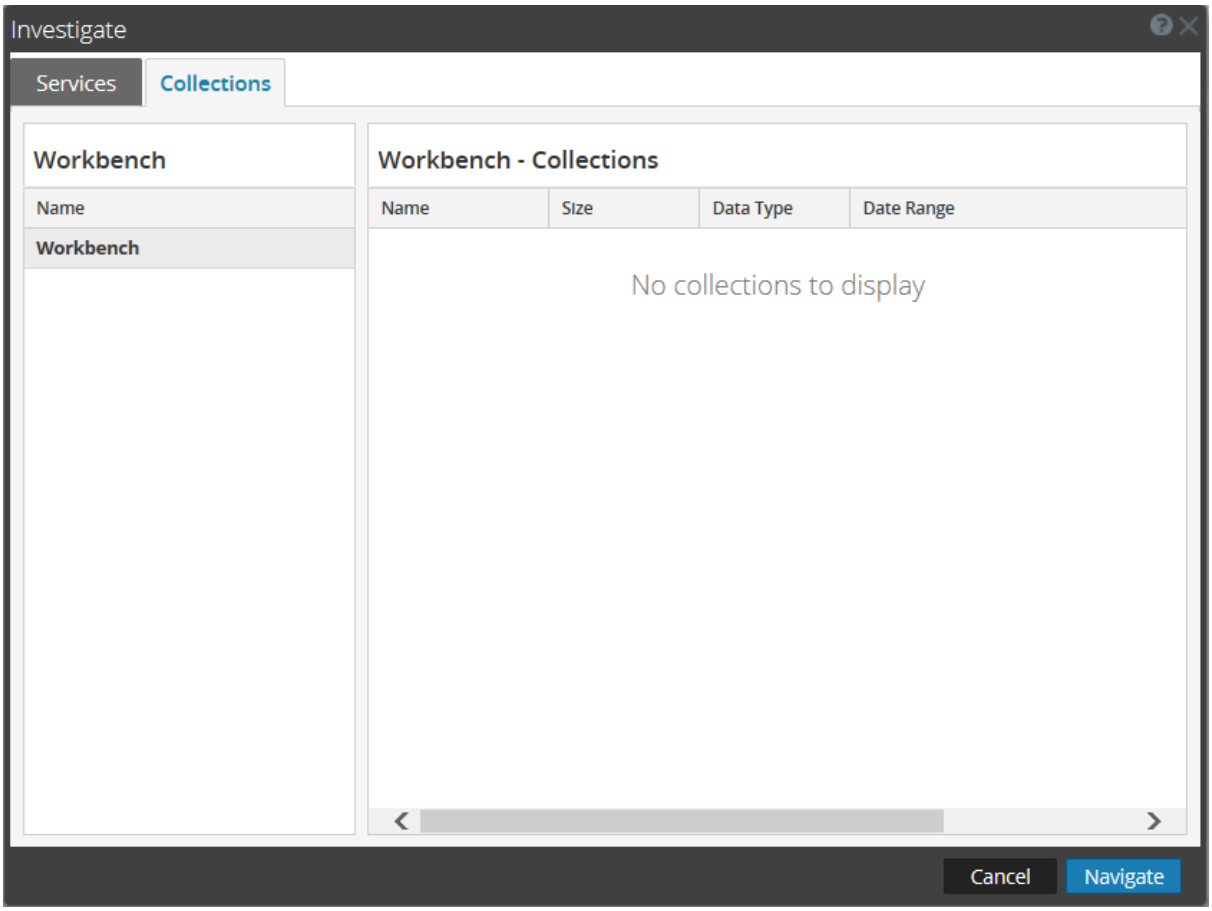
    After successfully creating a collection, you can use this collection for reporting and investigation purposes.

Investigate a Collection

To perform an investigation on a workbench collection:

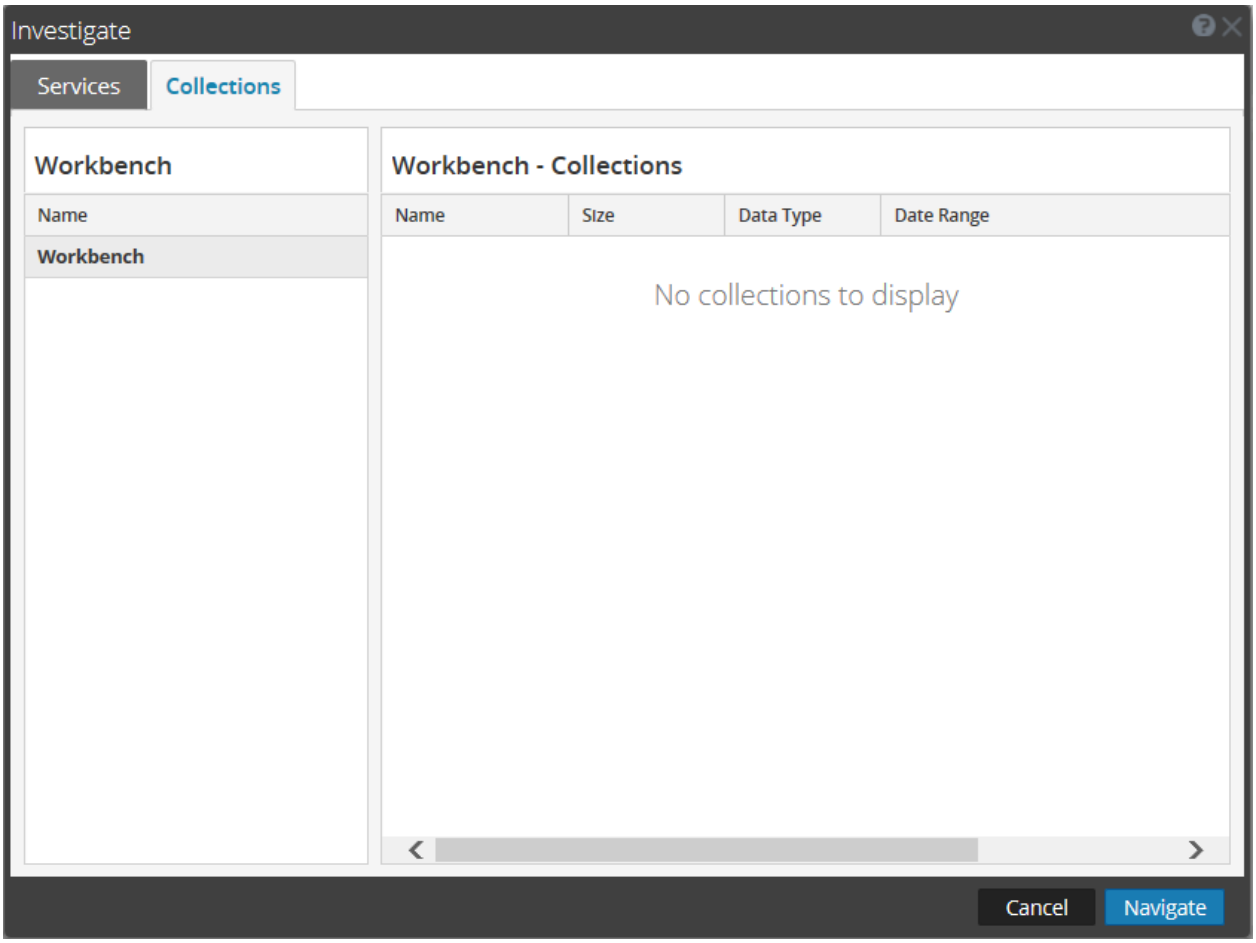
1. Select **Investigate**.

   The Investigate dialog is displayed.



2. Click the **Collections** tab in the Investigate dialog.
3. Select a Workbench service in the left panel.
4. Select the collection you want to investigate in the right panel.
5. Click **Navigate**.

The Navigate view is displayed showing data pertaining to the Workbench collection that you selected.



Note: For detailed information about using Investigation, see *Investigation and Malware Analysis Guide*.

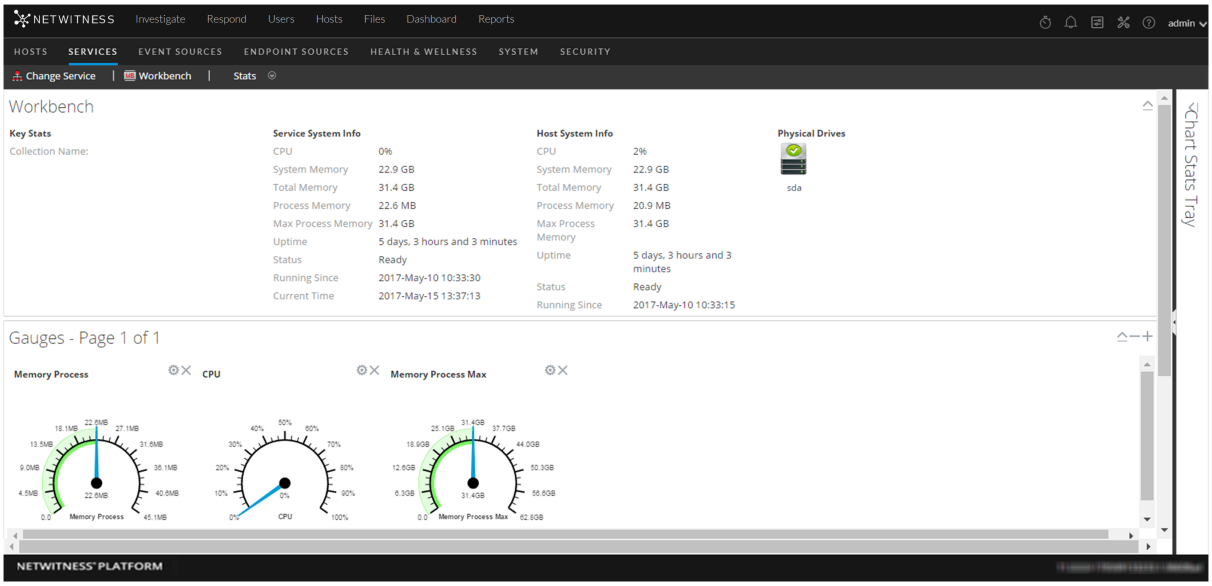
View Workbench Collection Statistics

The same statistics available for other services are provided for the Workbench service. The Services Stats view displays key statistics and system information that pertain to your selected Workbench service. The information is displayed in several different sections within the Stats view: Workbench, Gauges, Timeline Charts and Chart Stats Tray. The Chart Stats Tray lists all available statistics for the Workbench. Any statistic in the Chart Stats Tray can be displayed in a gouge or a timeline chart.

Perform the following steps to view workbench statistics:

1. Go to **(Admin) > Services**.
2. In the **Services** view, select a **Workbench** , then select **> View > Stats**.

The Services Stats view is displayed.



Note: For more information about Workbench statistics, see the *Host and Services Getting Started Guide*.

View Workbench Logs

Perform the following steps to view logs on a Workbench service:

1. Go to **(Admin) > Services.**
2. In the **Services** view, select a **Workbench**, then select **> View > Logs**.
   The Services Logs panel is displayed.

Note: For information about viewing and configuring audit logs, see the topic "Configure Global Audit Logging" in the *System Configuration Guide* .

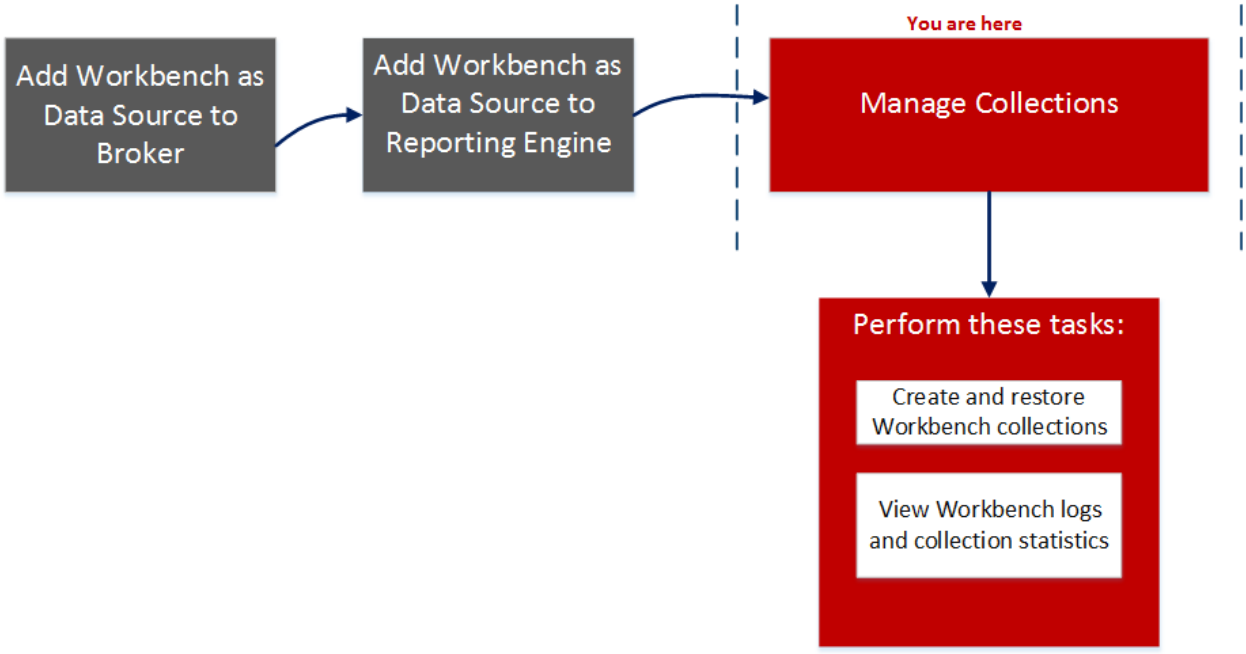
Services Config View - Workbench

In the Services Config view for workbench, some of the parameters are the same as other NetWitness services, while others are specific to the Workbench service.

The Services Config view - Workbench ((Admin) > Services >, select a Workbench service and select View > Config) provides a way to configure a Workbench service.

Workflow

These are the basic steps for configuring and managing a Workbench service.



What do you want to do?

Role	I want to...	Show me how...
Administrator	Add Workbench as data source to Broker	Adding Workbench Service as a Data Source to Broker
Administrator	Add Workbench as a Data Source to Reporting Engine	Adding Workbench as a Data Source to Reporting Engine
Administrator	*Create or delete a collection	Managing Collections
Administrator	*View Workbench statistics and logs	Managing Collections

Role	I want to...	Show me how...
Administrator	View configuration information about appliances that are connected to the Workbench service.	Select the **Appliance Service Configuration** tab. The Appliance Service Configuration tab is the same for all NetWitness services. It provides configuration information about appliances that are connected to the Workbench service. For information on the **Appliance Service Configuration** tab, see **Appliance Service Configuration Tab** in the *Host and Services Getting Started Guide*.

*You can perform this task here.

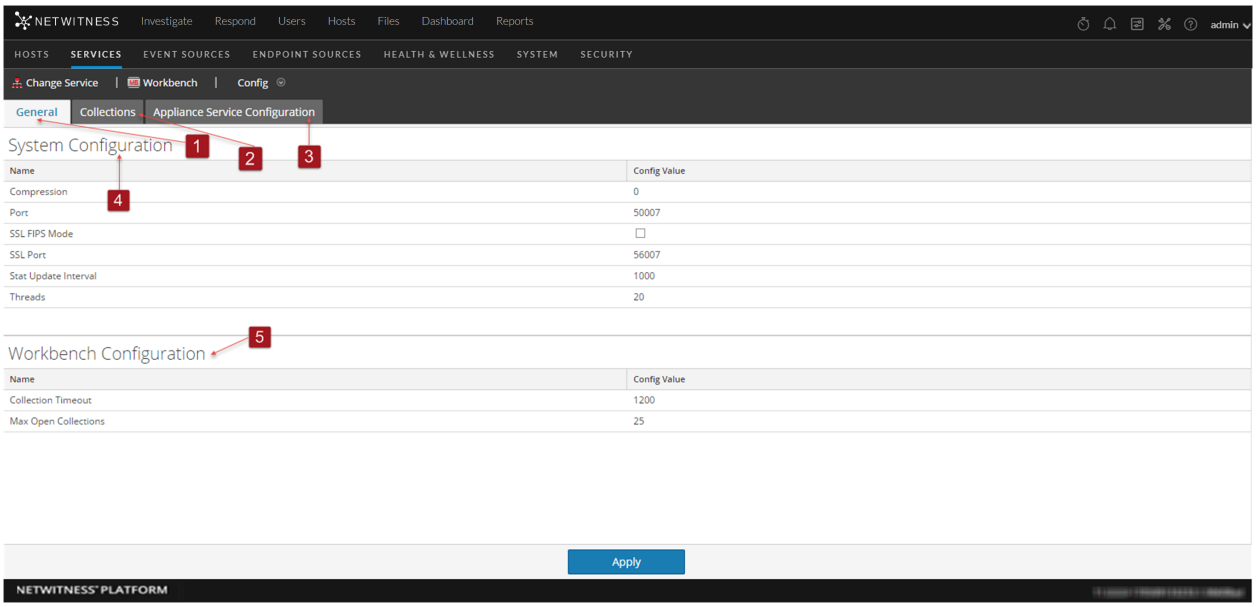
Related Topics

- Workbench Configuration Procedures

Quick Look

The Workbench service has three tabs and two panels in the Config view:

- General tab
- Collections tab
- Appliance Service Configuration tab
- System Configuration panel
- Workbench Configuration panel



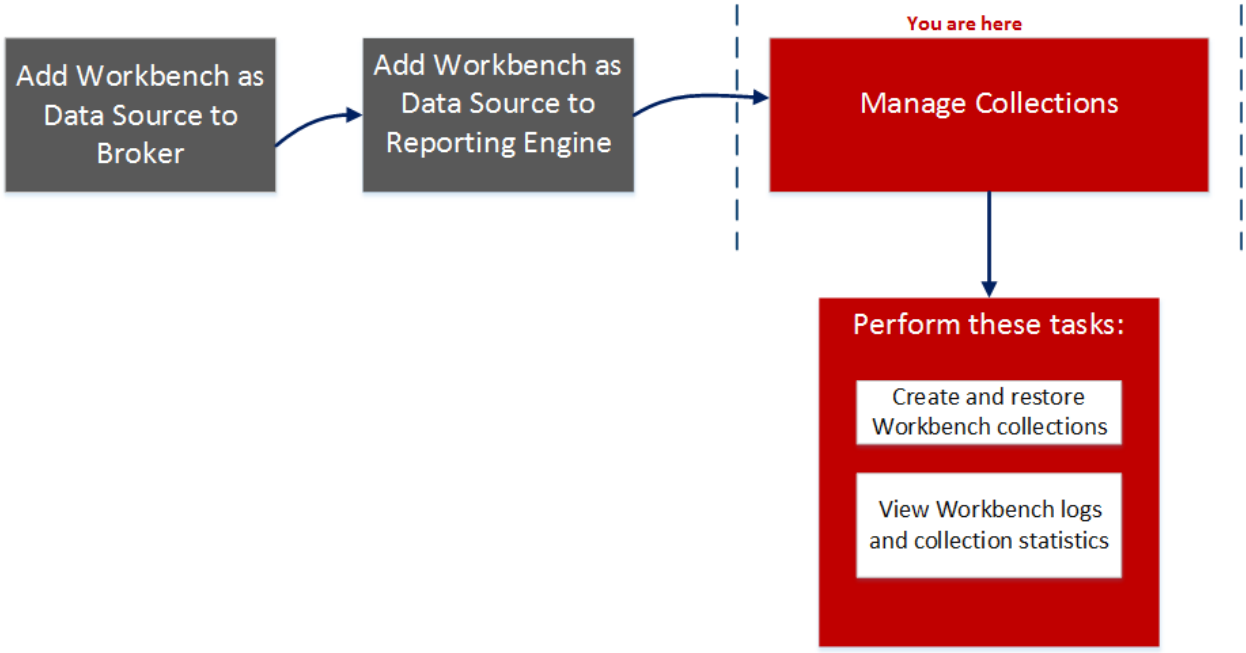
1	General tab provides a way to manage basic Workbench service configuration.
2	Collections tab provides a way to manage collections on a Workbench service.
3	Appliance Service Configuration tab provides a way to configure a Workbench service.
4	System Configuration panel provides a way to manage service configuration for a Workbench service.
5	Workbench Configuration panel provides a way to start and stop a Workbench service.

Services Config View - Collections Tab

The Collections tab for the Workbench service provides a way to manage workbench collections. To access the Collections tab, go to **(Admin) > Services** > select a **Workbench** service, select **View > Config** and select the **Collections** tab.

Workflow

These are the basic steps for configuring and managing a Workbench service.



What do you want to do?

Role	I want to...	Documentation
Administrator	*Create and restore Workbench collections.	Managing Collections
Administrator	*View Workbench logs and collection statistics.	Managing Collections

Role	I want to...	Documentation
Administrator	View configuration information about appliances that are connected to the Workbench service.	Select the **Appliance Service Configuration** tab. The Appliance Service Configuration tab is the same for all NetWitness services. It provides configuration information about appliances that are connected to the Workbench service. For information on the **Appliance Service Configuration** tab, see **Appliance Service Configuration Tab** in the *Host and Services Getting Started Guide*.

*You can perform this task here.

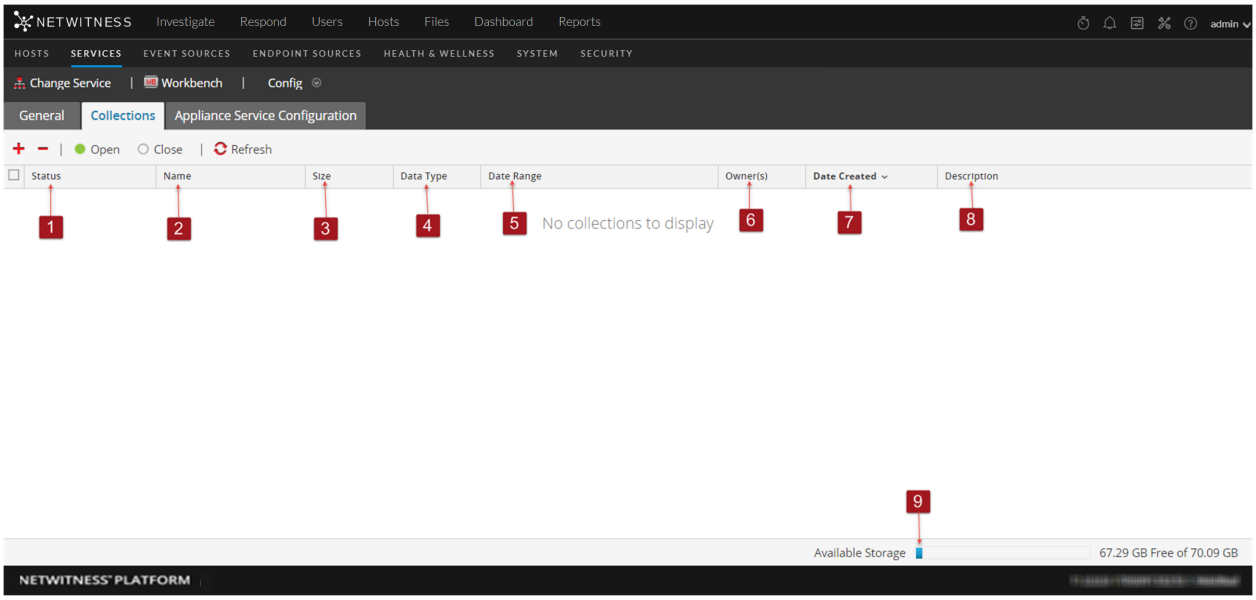
Related Topics

- Workbench Configuration Procedures

Quick Look

The Collections tab has a toolbar and a panel that lists relevant information about the Workbench collections.

The following figure is an example of the Collections panel.



1 Status of the Restoration Collection:

- **Resorting Data** - Data restoration is in progress.
- **Closed** - Data is restored.

	• **Opening** - Data is being indexed. • **Ready** - Indexing is complete. • **Closing** - Collection is closing.
2	**Name**: Name of the file being restored.
3	**Size**: Collection size.
4	**Data Type**: Logs.
5	**Date Range**: Lists the range of dates when the collection is being restored.
6	**Owner**: Lists the Collection creator.
7	**Date Created**: Shows the date when the collection was created.
8	**Description**: Description of the Restoration collection.
9	**Available Storage Indicator**: Shows the available disk space, given in gigabytes (GB). The Workbench validates to ensure there is enough available space when attempting to create a restoration collection.

Toolbar

These are the toolbar options.

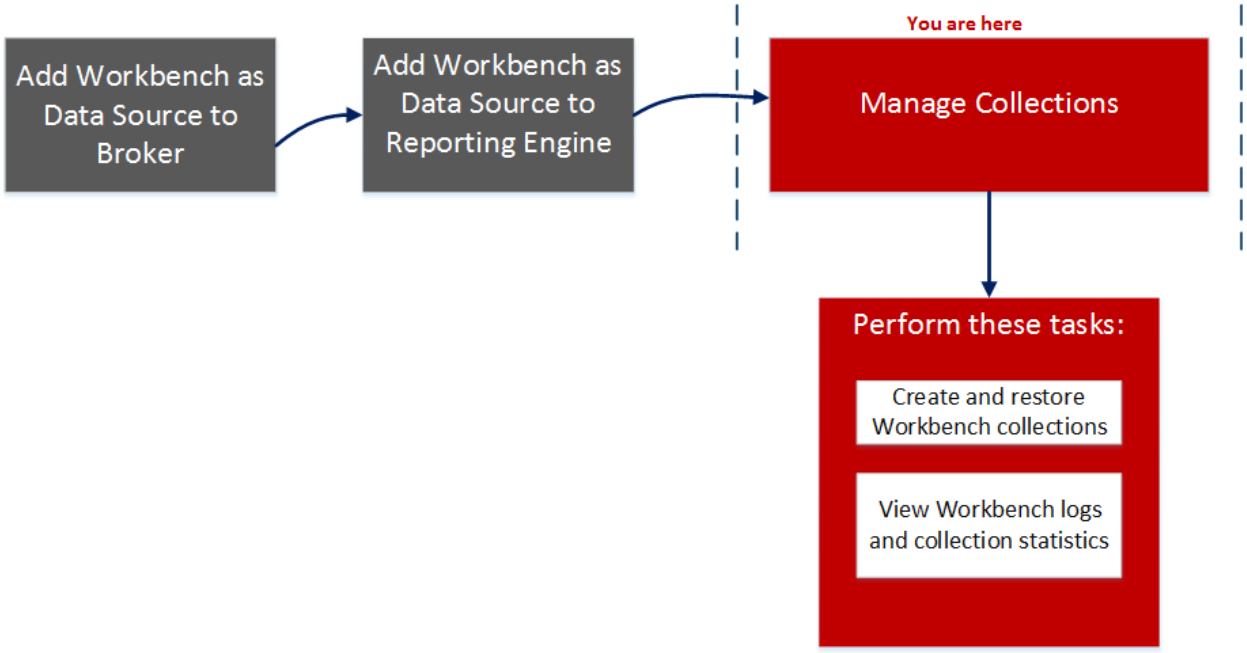
Parameter	Description
	Creates a new restoration collection.
	Deletes the selected Workbench collection.
Open and Close - Refers to the status of the restoration collection.	Open - Makes collection available for investigation and reporting. Close - Makes collection unavailable for investigation and reporting while preserving resources.
	Refreshes the list of Workbench collections.

Services Config View - General Tab

The General tab for the Workbench service provides a way to manage basic service configuration. To access the General tab, go to **(Admin) > Services >** select a service and select **> View > Config**.

Workflow

These are the basic steps for configuring and managing a Workbench service.



What do you want to do?

Role	I want to...	Show me how...
Administrator	*Create and restore Workbench service collections.	Managing Collections
Administrator	*View Workbench logs and collection statistics.	Managing Collections
Administrator	*Process Workbench collections.	Managing Collections

*You can perform this task here.

Related Topics

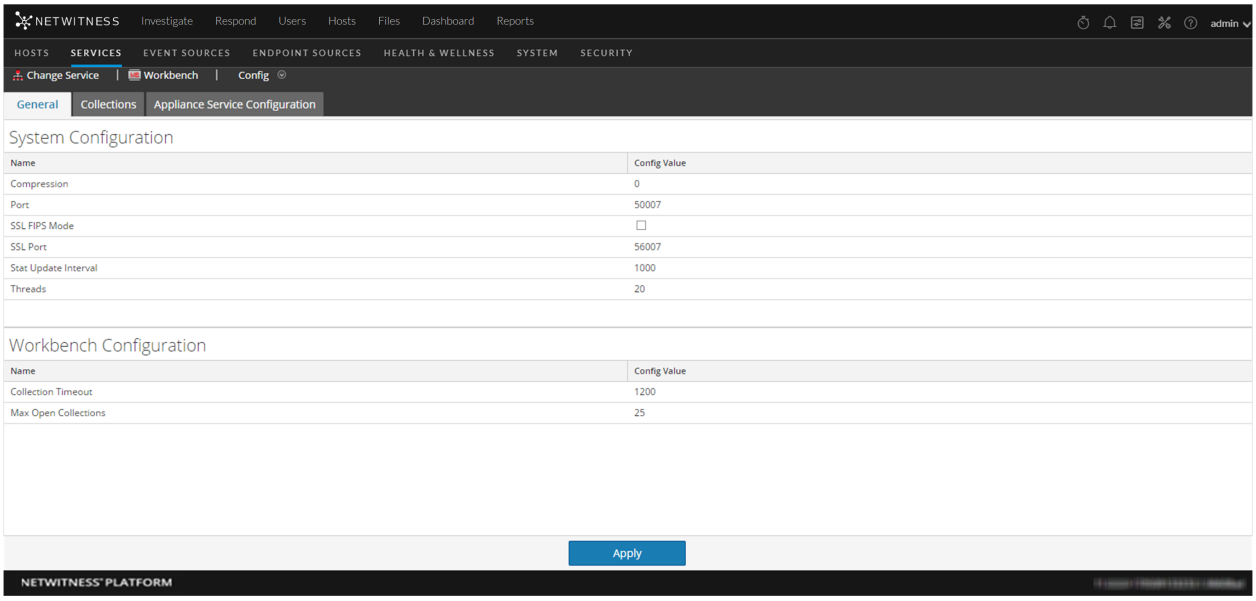
- Workbench Configuration Procedures

Quick Look

The General tab has two panels:

- System Configuration
- Workbench Configuration

The following figure is an example of the General tab.



System Configuration Panel

The System Configuration panel displays configuration parameters for the Workbench service. The following table describes the System Configuration panel features.

Parameter	Description
Compression	When set to a positive value, the minimum amount of bytes before a message is compressed. **0** means no compression for any message. Change takes effect on subsequent connections.
Port	The unencrypted port this service will listen on. **0** means disabled. Change takes effect on service restart.
SSL FIPS Mode	Determines whether the OpenSSL library will enter FIPS mode. Change takes effect on service restart.
SSL Port	The SSL port this service will listen on. **0** means disabled. Change takes effect on service restart.
Stat Update Interval	Determines how often (in milliseconds) statistic nodes are updated in the system. Change takes effect immediately.
Threads	The number of threads in the thread pool to handle incoming requests. Change takes effect immediately.

Workbench Configuration Panel

The Workbench Configuration panel displays configuration parameters for the Workbench collections. The following table describes the Workbench Configuration panel features.

Parameter	Description
Collection Timeout	The number of seconds before an idle collection is automatically closed.
Max Open Collections	The number of collections that can be open at once. A setting of 0 disables the limit.
Apply	Updates the modified configurations in the panel.

Troubleshooting

NetWitness notifies users of issues using pop-up notifications.

NetWitness Workbench returns the following types of error messages explained in the following table.

Problem	Possible Causes	Solutions

Unable to connect to workbench service from NetWitness **user interface Administration page.**	NetWitness service is not running.	Verify that your NetWitness service is running. Log in to your NetWitness Server and run the following command:
		`status nwworkbench`
		Firewall rules should allow connections from 50007, 50607 and 50107.
		Verify your connection by running the following command:
		`service iptables status`
		Verify that you are able to launch REST. Execute the following command for your appliance:
		`https://<IPAddress>:50107 service`
		If you are able to launch REST service for your appliance, you can confirm that there is no problem with the appliance. Navigate to the NetWitness side for further investigation as follows:
		• Enable debug mode and watch for sa.log errors located at:
		`/var/lib/netwitness/uax/logs`
		• Enable developer tools using the shortcut `Ctrl+Shift+I` for Chrome and verify the preview and response for the request.

Not able to view Appliance service configuration tab for workbench appliance running in SSL mode.		Enable SSL for appliance service and restart the appliance service.

Problem	Possible Causes	Solutions

The following error message is displayed when trying to load meta in order to create a report on a workbench collection: "Unable to fetch schema from data source when trying to load meta."		Load meta for the appliance from the NetWitness User Interface Rule library and watch for any errors in Reporting Engine log located at:
		`/home/rsasoc/rsa/soc/reporting-engine/logs`
		Launch REST for the device and watch for any error if you run the following query:
		`/sdk?msg=language&force-content-type=text/plain&expiry=600&size=10`

No results are displayed after running query from NetWitness User Interface via the Reporting Engine.		Run the query on the Reporting Engine and watch for `/var/log/messages` on the data source. Look for an exact query that matches the data source. **TIP**: Search for `[SDK-Query]` in log file. Copy the exact query and run from REST SDK to see if you get any results. `REST Query: /sdk?msg=query&force-contenttype=text/plain&expiry=600&query=select%20user.dst&size=10`
Workbench Available storage indicator in Workbench Collections Tab is not accurate.	Available storage indicator in the User Interface displays the default Collections directory shown below: `/VAR/NETWITNESS/WORKBENCH/COLLECTIONS`	None.
Unable to open new collections after opening existing collections.	There is a workbench configuration called “Max Open Collections” that is set to 25 by default. This configuration specifies the number of collections that can be open at the same time.	You can modify this number. A setting of zero disables the limit of maximum open collections.

Problem	Possible Causes	Solutions
Successfully opened a collection that got to Ready state. But after a while, the collection automatically changed to Closed state.	There is a workbench configuration called “collection.timeout” that is set to 1200 seconds by default. This configuration specifies the number of seconds before an idle collection is automatically closed. Maximum time allowed before timeout occurs is 86,400 seconds (24 hours).	A setting of zero disables the timeout.
Querying for a time range using `/database manifest` command returned blank output.	Blank output indicates that there are no **nwdb** files available for the time range.	None.
Created collection, but collection status is not available in Jobs, and collection is not displayed in workbench Collections tab.	You might be running in a mixed mode environment (for example, creating a collection on a 10.4.x version of workbench from a 10.5 NetWitness User Interface.	The collection is displayed in the workbench Collections tab after you reload the page.
Noticed blank Date Range and Date Created values for collections.	All collections display blank Date Range and blank Date Created values.	Date Range and Date Created values are displayed after upgrading to 10.5.
Discrepancy in behavior of adding workbench collections as a data source to Reporting Engine.	This behavior depends on whether you have a trusted connection or a non-trusted connection.	If your workbench service is established with a trusted connection, you should manually add workbench collections as a source to Reporting Engine. If your workbench service is not established with a trusted connection when the workbench restoration collection was created, it automatically sends a message to the Reporting Engine to add it as a source in the Reporting Engine.

Problem	Possible Causes	Solutions
Collection attributes (size, date range and date created) are not displayed.	Date range is not displayed for a collection if Jetty service is restarted while restoration is in process. Restoration collections created from an Explorer view display a blank Date Range. Any collections created on a 10.4 Workbench will display blank Date Range and blank Date Created values after upgrading to 10.5. In a mixed mode environment (10.5 NetWitness Server and 10.4.x workbench), size, date range, and date created are not displayed.	None.
Exception or blank page is displayed when drilling down on a workbench collection.	Collection closed because it exceeded the collection time out.	Investigate the collection from the beginning.
Empty collection is created.	Empty collection is displayed if restoration fails because Workbench service is restarted during collection creation.	None.
Service abruptly shuts down.		Run the service from command line and watch for any errors. For an example, run the command from the server console `/usr/sbin/NwWorkbench` for workbench.
REST request denied.		Verify user.agent.whitelist config located at `/rest/config/`. If non-blank, this should be a regex expression to match valid HTTP user agents. If the regex fails to match, all REST requests will be denied (see allow.missing.user.agent for the potential exception). If blank, all requests are allowed.
Queries with raw meta return blank values for Raw field.		Verify that you have a relevant **`packet db`**.

NetWitness® Platform XDR

Version 12.0

NetWitness Shell User Guide

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

NETWITNESS
Platform XDR

NetWitness® Platform XDR

Version 12.0

NetWitness Shell User Guide

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

July, 2022

NETWITNESS
Platform XDR

NetWitness® Platform XDR

Version 12.0

NetWitness Shell User Guide

Contents

Shell **6**

Features 6

Installation 6

Usage 7

System Commands 7

Available Commands 7

Help, History Command Usage 9

Authentication Commands 9

Context-Changing Commands 11

Connecting to a Service 11

Change Node 11

Node Display Commands 12

Node Value Commands 13

NETWITNESS
Platform XDR

NetWitness® Platform XDR

Version 12.0

NetWitness Shell User Guide

Node List Commands 14
Method Node Commands 16
Run State of a Service 17
Health 17
Respond Server APIs 17
Retrieve Incidents 18
Retrieve Incident Details 18
Retrieve Alerts 19
Add Alerts to an Incident 20
Delete Alerts from an Incident 21
Metrics 21
Config 22
Snapshot 23
Scripting 23
Troubleshooting Commands 24
fix-keystore 24
print-keystore 25

NETWITNESS
Platform XDR

NetWitness® Platform XDR

Version 12.0

NetWitness Shell User Guide

reconstruct-keystore 25
Advanced Customization 25
Tree View 27
Features 27
Implementation 27
Node Types 27
Node Structure 28

NETWITNESS
Platform XDR

Shell

This guide describes the shell utility `nw-shell` that can be used to troubleshoot the operations of NetWitness management services like `security-server`, `investigate-server`, and `correlation-server`. It is the equivalent of the `NwConsole` utility used to interact with NetWitness capture services like Decoders and Concentrators. The shell utility is independent of the business logic of the service, and works the same with most NetWitness services.

Features

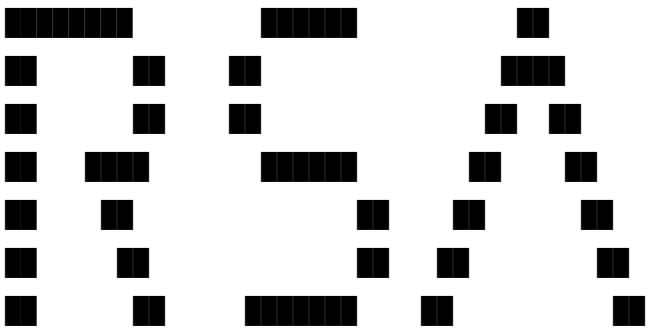
The `nw-shell` utility implements the following features:

- Supports secure connections to the local NetWitness service instances.
- Supports navigation of the service `tree` to explore its operational state (for more information, see Tree View).
- Provides an intuitive display of configuration, metrics, and health-check information to help with troubleshooting.
- Supports scripting to automate simple administration tasks in field deployments.
- Supports Linux, OSX and Windows terminals.

Installation

The `nw-shell` utility can be installed with the `rsa-nw-shell` RPM as shown here:

```
$ sudo yum install rsa-nw-shell
$ /usr/local/bin/nw-shell
```



```
RSA NetWitness Shell. Version: 4.0.0
See  "help" to list available commands, "help connect" to get started.
offline »
```

Usage

The primary goal of `nw-shell` is to help a human operator explore the runtime state of a NetWitness service. It is, essentially, an interactive program that invokes APIs on running NetWitness services. Each NetWitness service includes a system API that exposes its runtime state as a *logical* tree. The shell leverages the Tree API structure to present a hierarchical view of a service that is similar to a file system view. Users can navigate the tree by using the `cd` command to access directories, and can view or modify the corresponding configuration at the location, or invoke API methods and view the current state of components in the node.

The set of commands available to users at a given time depend on the current shell *context*, for example, their placement inside the *logical* tree. Certain commands, however, are always available, and we begin with a description of those commands.

System Commands

Shell *system* commands operate on the shell itself, instead of the nodes of a connected service. These commands are always available.

Command	Function	Example
`help [<command>]`	Provides help on available commands or a particular command.	`help` or `help connect`
`clear`	Clears the screen (shortcut `Ctrl-l`).	`clear`
`exit` or `quit`	Exits the shell.	`quit`
`history`	Displays the history of previously-run commands.	`history`

- Use the `help <command>` to explore the available commands. It is always available.
- `nw-shell` supports tab completion of a command, and the applicable parameter names wherever possible. For example, pressing the **Tab** key after typing **`e`** completes the command to `exit`.
- `nw-shell` also supports a non-interactive mode where it executes scripts from a provided file, used by specifying the absolute filepath prefixed by the `@` argument. See Shell for details.

Available Commands

```
Built-In Commands
        clear: Clear the shell screen.
        exit, quit: Exit the shell.
        help: Display help about available commands.
        history: Display or save the history of previously run commands
```

```
        script: Read and execute commands from a file.
        stacktrace: Display the full stacktrace of the last error.

Context Commands
      * cd: Change the current node. Usage: cd <path>
        connect: Connect to a service. One of --service or --port must be
specified.Usage: connect [--service <service>[.<id>]] [--broker
amqp://localhost/rsa/system] [--host localhost] [--port] [--insecure false]
      * where: Which service shell is connected to?

Token Commands
        login: Authenticate to a service. Usage: login [connect-parameters]
        login-insecure: Authenticate to a service providing user and password
on the command prompt. The password is recorded in the shell history so this
command must be used with care.Usage: login-insecure --user <user> --password
<password> [connect-parameters]
      * logout: Clear the authentication context: logout
      * whoami: Who am I?

Tree Node Commands
      * json: Print the current node as a JSON string
      * show: Pretty print the current node

Tree Node List Commands
      * config: Summarize configuration of the current subtree
      * health: Summarize health of the current subtree
      * ls: List the children of the current node. Usage: ls [<filter>] [--
values] [--types]
      * lsv: Shorthand for ls --values. Usage: lsv [<filter>] [--types]
      * method: Summarize methods of the current subtree
      * metrics: Summarize metrics of the current subtree
      * snapshot: Snapshot the current subtree

Tree Node Method Commands
      * invoke: Invokes the method that exists on the current method type
```

```
node. Usage: invoke [argument] [--file jsonfile]

Tree Node Value Commands
      * get: Get the value of the current node
      * set: Set the value of the current node. Usage: set <new-value>

Commands marked with (*) are currently unavailable.
Type `help <command>` to learn more.
```

Help, History Command Usage

```
offline » help connect
    connect - Connect to a service. One of --service or --port must be
specified. Usage: connect [--service <service>[.<id>]] [--broker
amqp://localhost/rsa/system] [--host localhost] [--port] [--insecure false]
    ...
    ...
offline » history
help
help connect
history
offline »
```

You can navigate previously-typed commands in `nw-shell` by using the Up and Down arrow keys, which can help minimize typing by recalling previously-executed commands.

Authentication Commands

As an administration and monitoring tool, it is important that `nw-shell` authenticates users before handing them control over a running service. The following commands manage the shell authentication context:

Command	Function
`login`	Authenticates and retains tokens for future interactions.
`logout`	Clears the security context and token.
`whoami`	Prints a summary of the current identity.

The `login` command can be used to authenticate the user and establish an identity. Service operators can log in once (against the NetWitness Security Server) and then use the token to connect to multiple services and perform administration based on the roles assigned. This single-sign-on workflow is simplified by separating the `login` and `connect` commands to allow the operator to authenticate once (using `login`) and then switch services seamlessly using `connect`.

In general, the `login` command takes the same parameters as `connect` (described below) to specify the service that performs the credential validation. For example:

```
offline » login
user: admin
password: ********
admin@offline » connect --service respond-server
INFO: Connected to respond-server (23e1dab7-0658-41a5-bb1e-d716a37d5ea5)
admin@respond-server:Folder:/rsa »
admin@respond-server:Folder:/rsa » connect --service investigate-server
INFO: Connected to investigate-server 2f21db20-4b50-48cf-8f7f-c0be0d1d1d12)
admin@investigate-server:Folder:/rsa »
```

The current logged-in identity can be confirmed at any point using the `whoami` command. For example:

```
admin@offline » whoami
```

Subject	admin
Issued By	security-server-ebeb90da-761b-4f14-b9ec-30410499d2b8
Issued At	2018-10-04T17:22:52.041Z
Expires After	2018-10-05T05:22:52.041Z
Roles	[ebeb90da-761b-4f14-b9ec-30410499d2b8]

The shell security context can be cleared by using the `logout` command.

```
security-server:Folder:/rsa » logout
security-server:Folder:/rsa » whoami
You are not logged in.
```

Some operations, such as setting certain configuration properties or invoking a method, require a certain RBAC. To perform such privileged operations, you need to authenticate with an identity that has the necessary permissions. For example:

```
security-server:Configuration » /rsa/logging/operations/max-file-count » get
10
security-server:Configuration:/rsa/logging/operations/max-file-count » set 15
ERROR: Failed to set the node value: Access is denied
security-server:Configuration:/rsa/logging/operations/max-file-count » login
user: admin
```

```
password: **********
admin@security-server:Configuration:/rsa/logging/operations/max-file-count »
set 15
security-server:Configuration:/rsa/logging/operations/max-file-count » get
15
```

Context-Changing Commands

The following commands change the current shell context.

Command	Function	Options	Examples
`connect`	Connect to a service	`servicehost`, `port`, `http`, `broker`	`connect --service security-server.d6a55b48...`
`cd`	Change context to a node		`cd /rsa/security`, `cd ../sys`, `cd ..`, `cd`

Connecting to a Service

The shell can connect to services over AMQP or HTTP(S), however, AMQP is the preferred and default option.

To connect to a specific instance of a service, `name` and `serviceId` needs to be supplied in the format:
`connect --service {service-name}.{serivceId}`
For example, `foo-server.d6a55b48-6103-46bd-9ead-3b4d589b302b`.

If the service identifier is skipped, it is assumed to be `any`, for example, `connect --service foo-server` will connect with any service named `foo-server` that is connected to the AMQP broker.

Change Node

The `cd` command can be used to change the current node. Just like `cd` on a file system shell, it takes relative or absolute paths as input and changes the current context to the node at that path.

```
offline » cd
Command 'cd' was found but is not currently available because you are not
connected to any service.


offline » connect
```

```
security-server:Folder:/rsa » cd log
security-server:Folder:/rsa/logging » cd ../security
security-server:Component:/rsa/security » cd /rsa/security/fips-mode
security-server:Configuration:/rsa/security/fips-mode » cd
security-server:Folder:/rsa »
```

The shell prompt summarizes and presents the user's current context. It starts off with `offline`, and once the shell is connected to a service, it displays the service name it is connected to, and the type and the path of the current node. Once the user is authenticated, the userId is included in the prompt.

Invoking `cd` before connecting to a service reminds the user that some commands work only in certain contexts.

Note: Some commands are available only in certain contexts. For example, `cd` works only when the shell is *online*. The Help command lists all the commands, but commands marked with (*) are unavailable in the given context.

Certain node-specific commands are enabled only when the current node is of a certain type. For example, method nodes support a command `invoke`, which is not enabled for any other node type.

Node Display Commands

The following commands are available with all node types. They do not take any arguments, and display the node details for the user to review.

Command	Function
`show`	Pretty-print the current node.
`json`	Print the current node as a JSON string.

For example:

```
security-server:Configuration:/rsa/security/authentication/token-
lifetime » show
```

Configuration	/rsa/security/authentication/token-lifetime
value	10 hours
valueType	com.rsa.asoc.launch.api.helpers.Seconds
defaultValue	10
description	The time-to-live on a token.

```
security-server:Configuration » /rsa/security/authentication/token-lifetime »
json
{
  "path"  :  "/rsa/security/authentication/token-lifetime",
  "type"  :  "Configuration",
```

```
  "value"  :  "10 HOURS",
  "parent"  : {
    "path"  :  "/rsa/security/authentication",
    "type"  :  "Component"
  },
  "attributes"  : {
    "defaultValue"  : 10,
    "valueType"  :  "com.rsa.asoc.launch.api.helpers.Seconds",
    "description"  :  "The time-to-live on a token."
  }
}
```

The output of `json` dumps the API payload and may contain more details than those shown by `show`.

Node Value Commands

`Configuration`, `Metric` and `Health` nodes have values. Their current values can be obtained using the `get` command.

Command	Function	Examples
`get`	Get the current value of the node.	
`set`	Update the value of the node.	`set false`, `set '{"enabled": false}'`

For example:

```
security-server:Gauge:/rsa/process/hostname » get
"hostxyz.corp.emc.com"
security-server:Gauge:/rsa/process/hostname » cd /rsa/transport/http/secure
security-server:Configuration:/rsa/transport/http/secure » get
true
```

The value of a `Configuration` node can be changed by invoking the `set` command. For example:

```
admin@security-server:Configuration:/rsa/security/pki/tls-protocols »
show
```

Configuration	/rsa/security/pki/tls-protocols
value valueType description	[TLSv1.2] java.lang.String[] This property controls the TLS protocol versions supported by the applications.

```
admin@security-server:Configuration:/rsa/security/pki/tls-protocols » set '
["SSLv3"]'
admin@security-server:Configuration:/rsa/security/pki/tls-protocols » value
[
  "SSLv3"
]
```

Node List Commands

`Folder`, `Component` and `Method` nodes can contain other nodes as their children. You can use the `ls` command to list the children nodes, their types, and where available, their current values.

Command	Function	Options	Examples
`ls <filter>`	List the child nodes of the current node.	`types`, `values`	`ls`, `ls cfg`, `ls --values`
`lsv <filter>`	Shorthand for `ls <filter> --values`	`types`	`lsv metrics ls name:fips`

For example:

```
security-server:Component:/rsa/process » ls
jvm                       Component
modules                   Component
current-time-utc          Gauge
current-time-utc-pretty   Gauge
fips140-mode              Gauge
hostname                  Gauge
mode                      Gauge
service-id                Gauge
service-name              Gauge
status                    Gauge
uptime                    Gauge
uptime-pretty             Gauge
version                   Gauge
version-full              Gauge
version-raw               Gauge
```

```
ready                    Method
shutdown                 Method

security-server:Component:/rsa/process » ls comp
jvm      Component
modules  Component

security-server:Component:/rsa/process » ls name:uptime
uptime         Gauge
uptime-pretty  Gauge

security-server:Component:/rsa/process » ls name:uptime --values
uptime         Gauge  3713854
uptime-pretty  Gauge  1 hour 1 minute 53 seconds

security-server:Component:/rsa/process » lsv
jvm
modules
current-time-utc         1483978142782
current-time-utc-pretty  2017-01-09T16:09:02.782Z
fips140-mode             true
hostname                 hostxyz.corp.emc.com
mode                     Normal
service-id               1fb7572a-4d87-497e-a4da-802819c10a72
service-name             no-op-server
status                   Running
uptime                   512064
uptime-pretty            8 minutes 32 seconds
version                  0.0
version-full             0.0.0.0
version-raw              0.0.0.0
ready
shutdown
```

Listing commands are not available on nodes that do not have children (for example, configuration or metric nodes):

```
security-server:Component:/rsa/process » cd hostname
security-server:Gauge:/rsa/process/hostname » ls
Command 'ls' exists but is not currently available because you are not logged
in.
```

Note: Values displayed in `ls` and `lsv` may be truncated for presentation purposes. Use `get` on the value node to get the complete (unaltered) value of a node.

Method Node Commands

The `method` command shows all the available methods and all the sub-folders under the current node.

```
security-server:Component:/rsa/process » method
/rsa/process/ready
/rsa/process/shutdown
```

On a node that is of `Method` type , the `invoke` command will invoke the method. For example:

```
security-server:Method:/rsa/process/ready » invoke
{
  "ready" : true,
  "serviceId" : "84f36740-5ae7-409f-b14f-b17e98703983",
  "marketingVersion" : "0.0"
}
```

If the `method` signature takes `input`, then `invoke <input>` will invoke the `method` with given parameters.

```
security-server:Method:/rsa/health/get » show
```

Method	/rsa/health/get
output	java.util.List<com.rsa.asoc.launch.api.health.api.HealthStatus>
input	java.lang.String
description	Get the details of a single composite health check @param id The health check id @return The composite health check

Metric	Value
failed	0
invoked	0
timer	0.0

```
Method:/rsa/health/get » invoke 'rsa.health.checks.security-pki'
[
  {
    "name" : "rsa.security.pki.pki-health",
    "status" : "Unhealthy",
    "details" : {
      "Reason" : "Using a self-signed certificate"    }
  }
]
```

For inline `method` execution with input payload, in cases where the input contains any special characters such as `\, \n, \r, \t , ",` the input needs to be properly escaped with a backslash (`\`).

In cases where `method` takes a complex payload, payload can be supplied with a file reference. In this case, payload does not need any special escaping.

```
Method:/rsa/configuration/collections/register » invoke --file
/tmp/jsonfile.txt
```

Run State of a Service

Once `nw-shell` is connected to a service, commands like `health`, `metrics`, `config` or `snapshot` can be used to retrieve the current state of the service at the current node, and all the sub-folders under the current node.

Health

The `health` command lists the health of the components, and the all the sub-folders, under the current node.

```
admin@security-server:Component:/rsa/process » health
/rsa/process/jvm/memory-health      Healthy
/rsa/process/modules/module-health  Healthy
```

Respond Server APIs

You can use `nw-shell` to retrieve and modify the incident details. In order for this to work, make sure that the relevant APIs are exposed. You can perform following actions using these APIs.

Retrieve Incidents

Retrieve Incident Details

Retrieve Incidents

Add Alerts to an Incident
Delete Alerts from an Incident

Retrieve Incidents

You can retrieve incidents using a single criteria. The criteria can include name, priority, risk score, status, alert count, etc. The `invoke` command returns all the incidents with the matching criteria.

Note: You must export the output to a file to avoid rollover in the shell output. You can query using single meta filters only, for example, filter by name or priority. Depending on the load, a large number of results might cause the request time-out.

Sample Query

```
admin@respond-server:Folder:/rsa » cd respond/incident
admin@respond-server:Component:/rsa/respond/incident » cd find-incident-by-criteria
admin@respond-server:Method:/rsa/respond/incident/find-incident-by-criteria »
invoke '["<meta name>", "<meta value>", "<number of records>"]'
```

Examples

The following example command reads the query from an input file and exports the results to an output file.

```
invoke --file /root/incident_req.txt outputFile /root/output_incidents.json
```

In the above example, the `incident_req.txt` file contains the following information.

```
[root@SA ~]# cat incident_req.txt
["status", "ASSIGNED", "500"]
```

The output is exported to the `output_incidents.json` file that contains 500 incidents with ASSIGNED status.

The following are some more examples:

- Retrieve incidents with `testinc for rule` as name and a maximum of 10 incidents.
  `["name", "testinc for rule", "10"]`
- Retrieve incidents with HIGH priority and a maximum of 50 incidents.
  `["priority", "HIGH", "50"]"]`
- Retrieve incidents with ASSIGNED status and a maximum of 500 incidents.
  `["status", "ASSIGNED", "500"]`

Retrieve Incident Details

The `invoke` command is used to list the details of specific incidents.

Sample Query

```
admin@respond-server:Folder:/rsa » cd respond/incident
admin@respond-server:Component:/rsa/respond/incident » cd find-by-ids
```

```
admin@respond-server:Method:/rsa/respond/incident/find-by-ids » invoke '["INC-1", "INC-2"]'
```

The output is incident details in JSON format displayed in `nw-shell`.

Examples

You can use an input file containing the request to list the incident details.

```
admin@respond-server:Method:/rsa/respond/incident/find-by-ids » invoke --file /tmp/findIncidents.txt
```

The output is incident details in JSON format displayed in `nw-shell`.

In the above example, the file `/tmp/findIncidents.txt` file contains the following information.

```
[root@NWAPPLIANCE27144 ~]# cat /tmp/findIncidents.txt

["INC-10", "INC-12"]
```

You can store the incident details to an output file.

```
admin@respond-server:Folder:/rsa » cd respond/incident
admin@respond-server:Component:/rsa/respond/incident » cd /rsa/respond/alert/add-alerts-to-incident
admin@respond-server:Method:/rsa/respond/incident/find-by-ids » invoke '["INC-1", "INC-2"]' outputFile /root/result_output.json
```

The output is incident details in JSON format stored in the `result_output.json` file.

Retrieve Alerts

You can retrieve alerts using a single criteria. The criteria can include any field that is part of the alert record (alert name, alert type, alert source, etc). The `invoke` command returns the alert with the matching criteria.

Note: You must export the output to a file to avoid rollover in the shell output. You can query using single meta filters only, for example, filter by name or type. Depending on the load, a large number of results might cause the request time-out.

Sample Query

```
admin@respond-server:Folder:/rsa » cd respond/alert
admin@respond-server:Component:/rsa/respond/alert » cd find-alerts-by-criteria
admin@respond-server:Method:/rsa/respond/alert/find-alerts-by-criteria »
invoke '["<meta name>", "<meta value>", "<number of alerts to retrieve>",
"<list of alert fields (include fields)>"]'
```

Note: If you provide `<null>` as the `<list of alert fields>` value, the output returns all the alert fields.

Examples

The following example command reads the query from an input file and exports the results to an output file.

```
invoke --file /root/alerts_req.txt outputFile /root/output_alerts.json
```

In the above example, the `alerts_req.txt` file contains the following information.

```
[root@SA ~]# cat alerts_req.txt
["alert.source","Event Stream Analysis", "20", "alert.severity"]
```

The output is exported to the `output_alerts.json` file that contains 20 alerts from Event Stream Analysis and the alert severity field included.

The following are some more examples:

- Retrieve 20 alerts from Event Stream Analysis with the alert severity field included.
  `["alert.source","Event Stream Analysis", "20", "alert.severity"]`
- Retrieve 20 alerts from Event Stream Analysis with all the alert fields included.
  `["alert.source","Event Stream Analysis", "20", "null"]`

Add Alerts to an Incident

The `invoke` command is used to add alerts to an existing incident. You will require alert IDs and Incident ID.

Sample Query:

```
admin@respond-server:Folder:/rsa » cd respond/alert

admin@respond-server:Component:/rsa/respond/incident » cd
/rsa/respond/alert/add-alerts-to-incident
admin@respond-server:Method:/rsa/respond/alert/add-alerts-to-incident » invoke
'["INC-1","603dcc5091e818391ee57e1d","603dcc5091e818391ee57e1c",
"603dcc5091e818391ee57e1e"]'
```

Alternatively, you can specify the add alert query in an input file.

```
invoke --file /tmp/add_alerts.txt
```

In the above example, the `/tmp/add_alerts.txt` file contains the following information.

```
[root@NWAPPLIANCE27144 ~]# cat /tmp/add_alerts.txt

["INC-1","603dcc5091e818391ee57e1d","603dcc5091e818391ee57e1c",
"603dcc5091e818391ee57e1e"]
```

The first parameter is the incident ID (INC-1). The rest of the parameters are alert IDs (for example, 603dcc5091e818391ee57e1e) to be added to the incident mentioned in the first parameter.

The output is a list of alert IDs added to the incident.

Delete Alerts from an Incident

The `invoke` command is used to remove alerts from an existing incident. You will require alert IDs and Incident ID.

Sample Query:

```
admin@respond-server:Folder:/rsa » cd respond/alert

admin@respond-server:Component:/rsa/respond/incident » cd
/rsa/respond/alert/delete-alerts-by-alert-ids
admin@respond-server:Method:/rsa/respond/alert/delete-alerts-by-alert-ids »
invoke '["INC-1","603dcc5091e818391ee57e1d","603dcc5091e818391ee57e1c",
"603dcc5091e818391ee57e1e"]'
```

Alternatively, you can specify the delete alert query in an input file.

```
invoke --file /tmp/delete_alerts.txt
```

In the above example, the file `/tmp/delete_alerts.txt` file contains the following information.

```
[root@NWAPPLIANCE27144 ~]# cat /tmp/delete_alerts.txt

["INC-1","603dcc5091e818391ee57e1d","603dcc5091e818391ee57e1c",
"603dcc5091e818391ee57e1e"]
```

The first parameter is the incident ID (INC-1). The rest of the parameters are alerts IDs (for example, 603dcc5091e818391ee57e1e) to be removed from the incident mentioned in the first parameter.

The output is a list of alert IDs removed from the incident.

Metrics

The `metrics` command lists all of the available metrics, and all the sub-folders, under the current node.

```
admin@respond-server:Component:/rsa/tree/node » metrics
/rsa/tree/node/get/invoked   14
/rsa/tree/node/get/timer     5666541.077918259
/rsa/tree/node/list/invoked  12
/rsa/tree/node/list/timer    6652400.379371729
```

In the above example, `/rsa/tree/node/get/invoked 14` shows that the method `get` is invoked `14` times, and `timer` shows that it took on average `5666541.077918259` nano seconds to process the `get` request.

Using `cd` into the `timer` node shows more details, such as rate of requests, for the current node for per-minute and five minutes. For example:

```
admin@respond-server:Timer:/rsa/tree/node/get/timer » show
```

Timer	/rsa/tree/node/get/timer
value	3267513.7521543643
valueType	java.lang.Double
stdDeviation	4381863.808014679
fiveMinuteRate	0.018725581010540414
max	181326026
count	27
sampleSize	27
oneMinuteRate	0.038516144812437314
min	810240
fifteenMinuteRate	0.009425563630894371
median	2146170.0
meanRate	5.189274295248098E-4
mean	3267513.7521543643

Config

The `config` command shows all the configuration properties in the current node and all the sub-folders under the current node.

```
admin@respond-server:Component:/rsa/logging » config
/rsa/logging/audit/max-file-count       10
/rsa/logging/audit/max-file-size        10 MB
/rsa/logging/forward/categories         [Audit]
/rsa/logging/forward/destination        SYSLOG_UDP
/rsa/logging/forward/enabled            true
/rsa/logging/forward/host               localhost
/rsa/logging/forward/port               50514
/rsa/logging/forward/secure             false
/rsa/logging/levels
/rsa/logging/operations/max-file-count  10
/rsa/logging/operations/max-file-size   10 MB
```

Snapshot

The `snapshot` command combines `metrics`, `configuration`, and `health` of the current state of the service node and all the sub-folders under the current node. This command is very useful for troubleshooting purposes. When opening a support case, take a `snapshot` dump at the root node `\rsa` of a service that might be having issues, and attach it to the case.

```
admin@security-server:Component:/rsa/security » snapshot
/rsa/security/account/external/get-all/invoked      16
/rsa/security/account/external/get-all/timer        NaN
/rsa/security/account/external/get/failed           6
/rsa/security/account/external/get/invoked          6
/rsa/security/account/external/get/timer            NaN
/rsa/security/account/force-password-change/invoke  1
/rsa/security/account/force-password-change/timer   NaN
/rsa/security/account/get/invoked                   123
/rsa/security/account/get/timer                     902163.0
/rsa/security/account/password-policy/cannot-inclu  false
/rsa/security/account/password-policy/min-chars     8
/rsa/security/account/password-policy/min-lower-ch  0
/rsa/security/account/password-policy/min-non-lati  0
/rsa/security/account/password-policy/min-numeric-  0
/rsa/security/account/password-policy/min-special-  0


----
```

Scripting

`nw-shell` supports non-interactive executions that can take commands from files. For example:

```
> cat /tmp/foo.script
connect --service security-server
cd /rsa/security/pki/ciphers
get
> time nw-shell @/tmp/foo.script
INFO: Connected to security-server (52c7b92c-23d5-4b3e-9973-31d8b1b27ec4)
[
```

```
  "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256",
  "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA",
  "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256",
  "TLS_DHE_RSA_WITH_AES_128_CBC_SHA"
]
nw-shell    5.89s user 0.33s system 305% cpu 2.032 total
```

`nw-shell` also has a built-in `script` which can be used to execute scripts in the shell once the shell is launched.

```
 offline » script --file /tmp/foo.script
 INFO: Connected to security-server (52c7b92c-23d5-4b3e-9973-31d8b1b27ec4)
 [
   "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256",
   "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256",
   "TLS_DHE_RSA_WITH_AES_128_GCM_SHA256",
   "TLS_RSA_WITH_AES_128_CBC_SHA256"
 ]
```

Note: `nw-shell` is primarily meant to help a human to explore the runtime state of a NetWitness Platform service. For most other purposes, it is usually best to use the service published APIs.

Troubleshooting Commands

The `nw-shell` utility provides troubleshooting commands to help recover from misconfigurations that may cause the service to work improperly.

fix-keystore

When a service's keystore is corrupt or its SSL trust is broken with the rest of NetWitness, you can use the `fix-keystore` command to fix this issue. This command reads the keystore file of the services running on the host and attempts to reestablish trust with the NetWitness Platform CA. For example, the following is the command to repair the respond-server 's keystore:

```
>> fix-keystore --service respond-server
```

This command repairs the keystore `keystore.p12` that is located at `/etc/netwitness/respond-server`. It does not modify the service's original keystore. The repaired keystore can be found with the suffix

`.good`, for example, `keystore.p12.good`. If no fixes can be made to the keystore, this command does not write the new `keystore.p12.good` file.

Once the keystore repair is successful, the system admin can restore the service's communication by renaming `keystore.p12.good` to `keystore.p12` and restart the service.

print-keystore

The `print-keystore` command prints the requested NetWitness Platform Service's keystore certificates to the given output file in JSON format. Unless verbose is enabled, only a few key attributes of the certificates are printed to the file. For example, the following command prints the respond-server's keystore certificates.

```
>> print-keystore respond-server outputFile /tmp/crts
```

With verbose enabled:

```
>> print-keystore respond-server outputFile /tmp/crts verbose
```

reconstruct-keystore

The `reconstruct-keystore` command can be used when a service's keystore needs to be restored to its original clean state.

The following command reconstructs the respond-server's keystore to the new file `/etc/netwitness/respond- server/keystore.p12.new` (it does not modify the original keystore).

```
>> reconstruct-keystore --service respond-server
```

To restore the service back to its healthy state, restore `keystore.p12.new` to `keystore.p12`, and restart the service.

Advanced Customization

The following JVM system properties can be used to customize certain aspects of the `nw-shell` presentation:

System Property	Controls	Default Value
`console.width`	The width of text.	100
`console.colors`	The use of colors.	true
`console.prompt`	The prompt.	`%s%s%s »`
`timeout`	The API timeout.	30s

These can be specified using the `JAVA_OPTS` environment variable. For example:

```
> export JAVA_OPTS="-Dtimeout=100s -Dconsole.width=40 -
Dconsole.prompt=%s%s%s>"
> export JAVA_OPTS="-Dconsole.colors=false"
> nw-shell
```

The API timeout is a time unit which can use typical time unit values such as 30 MINUTES, 30 mins, 30 m, 5 seconds, 5 sec, 5s, 5, 8000 milliseconds, or 8000 ms (ms = milliseconds).

The value of `console.prompt` must contain placeholders (for example, `%s`) for three strings which are replaced, in order, by:

1. The node type
2. A separator (`:`)
3. The node path

Setting `console.colors` to `false` will not turn off *all* use of colors, it only disables colors used by `nw-shell` in its output. The underlying libraries that the program uses may still output some text in color.

Tree View

A NetWitness service, such as `security-server` or `investigate-server`, is a collection of components working together to implement the goals of the service. This topic describes the support for exposing these components and their associated elements (for example, configuration parameters, metrics, API methods, health checks, and so on) in a `tree-view`. The `tree-view` is a hierarchically-nested arrangement of `tree-nodes`, each of which corresponds to a distinct functional area of the service. The arrangement enables the end-user to navigate the nodes at run-time and to explore the current state of the service. See Shell for an interactive tool that uses the tree-view.

Features

The `tree-view` implements the following features:

- Supports encapsulation of functionally related elements under the same `tree-node`.
- Supports hierarchical linkage of `tree-node` s to enable navigation from one node to another.

Implementation

The current implementation of the `tree-view` is described in the following sections.

Node Types

The `tree-view` is comprised of `tree-node` instances taken from the following types.

Type	Represents
`Component`	A service component with associated configurations, metrics, health, APIs, etc.
`Configuration`	A configuration parameter that controls some aspect of the service behavior (for example, `rsa.filesystem.prefix`).
`Metric`	A *quantitative* measure of some aspect of the component operations (for example, Requests Processed).
`Health`	A *qualitative* assessment of some aspect of the component (for example, "Using too much memory").
`Method`	An API that can be invoked remotely to interact with the component (for example, `/rsa/security/pki/setCertificate`)
`Parameter`	A payload parameter that is an input or an output from an API.
`Folder`	A logical grouping of nodes that are related to one another but are not a component (for example, the root `/rsa`).

Node Structure

All `tree-nodes` have the same structure. In particular, they include the following elements:

Item	Information
`path`	The node path is a backslash (/) separated path, for example, `/rsa/security/pki/api-secure`.
`type`	The type of the node, taken from the previous table.
`attributes`	A map of string keys to attribute values. The available keys depend on the type of the node.
`value`	The value of this node. This can be null if the node has no associated value (or if it has a null value).
`children`	A list of child nodes, if the node has any.
`parent`	The parent node. This is only null for the root node.

NetWitness® Platform XDR

Version 12.0

System Security and User Management

NETWITNESS

Platform XDR

Contact Information

NetWitness Community at https://community.netwitness.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

October, 2022

Contents

Set Up System Security **9**

Setting Up System Security **10**

Set Up System Security Workflow 10

Configure Password Complexity 10

Password Strength 11

Configure Password Strength 11

Change the Default Admin Passwords 14

Best Practices 14

Change the admin Password for the NetWitness 14

Change the admin Password for Core Services 14

Remove and Re-add a Data Source on the Reporting Engine 15

Change the admin Password for a Service Using the REST API 15

Configure System-Level Security Settings 17

Configure Security Settings 17

Restrict Access to Incidents 18

(Optional) Configure External Authentication 20

Configure Active Directory 20

Configure Active Directory Authentication 20

Add a New Active Directory Configuration 21

Edit an Active Directory Configuration 22

Test an Active Directory Configuration 24

Delete an Active Directory Configuration 24

Configure PAM Login Capability 24

Pluggable Authentication Modules 24

Name Service Switch 25

PAM and NSS Combination 25

Process Overview 25

Configure and Test the PAM Module 25

PAM Kerberos 26

PAM RADIUS 27

Add a RADIUS Client and Associated Agent 28

PAM Agent for SecurID 31

Configure and Test the NSS UNIX Service 35

Configuration 35

Test NSS Functionality 35

Enable PAM in NetWitness Server 35

Test External Authentication for PAM 36
Create Group Mappings in NetWitness Server 37
(Optional) Configure PKI Authentication 38
(Optional) Use a Custom Server Certificate 38
Supported Keystore Formats 38
(Optional) Create a Certificate Signing Request (CSR) and Certificate Store for a Server Certificate 38
Import an NW Server Certificate with its Private Key 40
(Optional) Create a Customized Login Banner 42
How Role-Based Access Control Works 44
Preconfigured Roles 44
Trusted Connections Between Server and Service 45
How Trusted Connections Are Established 46
Common Role Names on the Server and Services 46
End-to-End Workflow for User Setup and Service Access 47
Role Permissions 48
Service Permissions Format for New Services 48
Admin-server 49
Administration 49
Alerting 52
Config-server 53
Content-server 53
Contexthub-server 55
Correlation-server 57
Dashboard 59
Endpoint-broker-server 61
Endpoint-server 63
Incidents 65
Integration-server 66
Investigate 67
Investigate-server 68
License-server 71
Live 71
Malware 72
Metrics-server 73
Orchestration-server 74
Reports 75
Respond-server 78
Incident Email Notification Settings Permissions 81
Respond Event Analysis Permissions 82
Respond Saved Filter Permissions 82

Security-server ... 82
Source-server ... 83
Springboard ... 85
Manage Users with Roles and Permissions ... 86
Manage Users Workflow ... 86
Review the Preconfigured NetWitness Platform Roles ... 86
(Optional) Add a Role and Assign Permissions ... 88
Add a Role and Assign Permissions ... 89
Change Permissions Assigned to a Role ... 90
Duplicate a Role ... 90
Delete a Role ... 90
Verify Query and Session Attributes per Role ... 91
Query and Session Attributes ... 91
How Query-Handling Attribute Settings Apply to Individual Users ... 91
Set Query Handling Attributes for a User Role ... 92
Set Up Users ... 94
Add a User and Assign a Role ... 94
Add a Local User ... 94
Add a User for External Authentication ... 97
Change User Information or Roles ... 99
Delete a User ... 99
Reset a User Password ... 99
Enable, Unlock, and Delete User Accounts ... 100
Enable Disabled NetWitness User Accounts ... 100
Disable NetWitness User Accounts ... 101
Unlock Locked NetWitness User Accounts ... 101
Delete NetWitness User Accounts ... 101
(Optional) Map User Roles to External Groups ... 102
Add Role Mapping for an External Group ... 102
Edit Role Mapping for a Group ... 103
Search for External Groups ... 104
Set Up Multi-Factor Authentication ... 106
ADFS Log in to NetWitness with SecurID Passcode ... 106
Configure Authentication Manager ... 106
Configure NetWitness ... 106
Configure ADFS ... 107
PAM SecurID Log in to NetWitness for AD Users ... 107
Configure Authentication Manager ... 108
Configure NetWitness ... 108

Set Up Single Sign-On Authentication **109**
NetWitness Single Sign-On Authentication Workflow 109
Configure Single Sign-On 110
Enable Single Sign-On 110
(Optional) Set Up Public Key Infrastructure (PKI) Authentication **112**
NetWitness PKI Authentication Workflow 113
Configure PKI Authentication 113
Import Server Certificate and Trusted CA Certificate 114
Certificate Revocation List 114
User Principal Settings 115
Lookup Query 115
Import NW Server Certificate with its Private Key 115
Import Trusted CAs, Configuring CRL and User Principal Settings 115
(Optional) Configure the CRL Manually 125
Enable PKI Authentication 127
Disable PKI 127
Disable PKI Authentication 128
Disable PKI using command line 128
Delete Server Certificate and Trusted CA Certificate 129
Troubleshooting **131**
Users are able to create a password of 8-chracters or less despite the configured minimum password length of 9 characters in Version 11.3 131
Unable to log in to NetWitness Platform using SSO 132
Manual Steps to Disable SSO 132
References **134**
Admin Security View 135
What do you want to do? 135
Related Topics 135
Quick Look 135
Users Tab 137
Workflow 137
What do you want to do? 137
Related Topics 137
Quick Look 137
Add or Edit User Dialog 140
What do you want to do? 140
Related Topics 140
Quick Look 140
Add User Dialog 140
Edit User Dialog 141

User Information 142
Roles Tab 142
Attributes Tab 143
Roles Tab 144
Workflow 144
Related Topics 144
What do you want to do? 144
Related Topics 144
Quick Look 144
Add or Edit Role Dialog 146
What do you want to do? 146
Related Topics 146
Quick Look 146
Role Info 147
Attributes 147
Permissions 148
External Group Mapping Tab 149
Workflow 149
What do you want to do? 149
Related Topics 149
Quick Look 149
Add Role Mapping Dialog 151
What do you want to do? 151
Quick Look 151
Group Mapping 153
Mapped Roles 154
Search External Groups Dialog 154
What do you want to do? 154
Quick Look 155
Settings Tab 157
Workflow 157
What do you want to do? 157
Related Topics 157
Quick Look 157
Password Settings 159
Security Settings 160
PAM Authentication 161
Active Directory Configurations 161
Restrict Access to Incidents 162
PKI Settings Tab 163

Workflow 163
What do you want to do? 163
Quick Look 163
Login Banner Tab 166
What do you want to do? 166
Quick Look 166
Single Sign-On Settings Tab 168
What do you want to do? 168
Quick Look 168

Set Up System Security

This guide provides information about setting up system-wide security in NetWitness Platform and controlling user access. The system administrator needs to understand system-wide security settings, user accounts, system roles, permissions, and access to services.

This guide contains the following sections:

- Setting Up System Security
- How Role-Based Access Control Works
- Manage Users with Roles and Permissions
- (Optional) Set Up Public Key Infrastructure (PKI) Authentication
- References

Setting Up System Security

Setting up system security involves configuring password complexity and other system-level security settings for users to securely log in to NetWitness and to prevent unauthorized user access.

Set Up System Security Workflow

This figure shows the high-level workflow for setting up the system security.



First, you can configure password complexity by setting up password strength by specifying maximum password length, decimals, special characters, and so on. Next, you should change the default administrator passwords. Then you can configure system-level security by adjusting the security settings such as lockout period, maximum login failures, session timeout, idle period, and case-sensitive usernames. You can also configure optional settings for external authentication, PKI authentication, and a customized login banner.

These are the procedures for setting up system security:

- Configure Password Complexity
- Change the Default Admin Passwords
- Configure System-Level Security Settings
- (Optional) Configure External Authentication
- (Optional) Configure PKI Authentication
- (Optional) Create a Customized Login Banner

Configure Password Complexity

Passwords are an important part of your network security strategy. They provide critical front-line protection for your computer systems and help prevent attacks and unauthorized access to private information.

Password policies, designed to enhance the security of corporate networks, vary depending on the industry, corporate requirements, and regulations. Because of these password policy variations, NetWitness software allows you to configure the password complexity requirements for internal NetWitness users to conform to your corporate password policy guidelines.

Password complexity requirements apply only to internal users and are not enforced for external users. External users rely on their own methods and systems to enforce password complexity.

In addition, you can set a global default user expiration period and determine if and when internal users receive notification that their passwords are about to expire. The password expiration notification consists of a password expiration message when a user logs on to NetWitness.

Password Strength

Strong passwords make it more difficult for attackers to guess user passwords and help prevent unauthorized access to your organization's network. You can define the appropriate level of password strength for your NetWitness users. When you configure the password strength settings, they apply to internal NetWitness users, including the admin user.

You can choose to enforce any combination of the following password strength requirements when a NetWitness user creates or changes their password:

- Minimum password length
- Minimum number of uppercase characters
- Minimum number of lowercase characters
- Minimum number of decimals (0 through 9)
- Minimum number of special characters
- Minimum number of non-Latin alphabetic characters (includes Unicode characters from Asian languages)
- Whether or not the password can contain the username

For example, you can create a strong password requirement that has a minimum of 9 characters, cannot contain the username of the user, and contains a mix of uppercase and lowercase letters, numbers, and special characters.

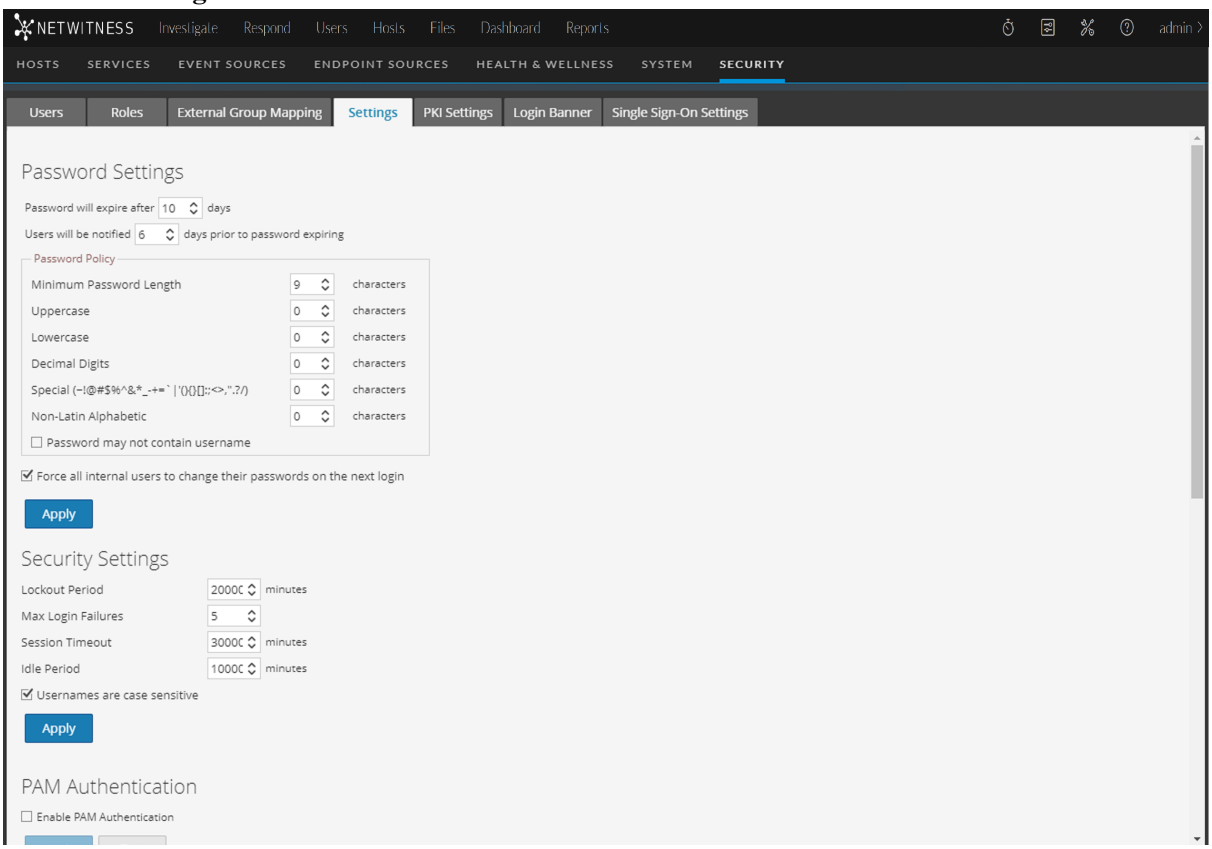
If you choose to enforce a minimum number of non-Latin alphabetic characters, ensure that your users have these characters available to them when setting their passwords.

For an example of a strong password policy, see the "STIG Compliant Passwords" in the *System Maintenance Guide*.

Configure Password Strength

1. In NetWitness, go to  **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.

2. Click the **Settings** tab.



3. In the **Password Settings** section, select the password complexity requirements to enforce when NetWitness users set their passwords and specify the minimum characters required, if applicable. Set the value to 0 for requirements you do not want to enforce, except for Minimum Password Length, which has a minimum value of 9 characters.

Note: In 11.2 and previous versions, the minimum password length is 8. Hence, on upgrade or update from previous versions to 11.3, you must set the minimum password length to 9 characters.

Requirement	Description
Password will expire after <n> days	The default number of days before a password expires for all internal NetWitness users. A value of zero (0) disables password expiration. For new installations, the default value is 30. For upgrades, the previous value will migrate automatically to the upgraded installation.
Users will be notified <n> days prior to password expiring	The number of days before the password expiration date, to notify a user that their password is about to expire. Users see a Password Expiration Message dialog when they log on to NetWitness. The minimum value is 1 day.

Requirement	Description
Minimum Password Length	Specifies a minimum password length. A minimum password length prevents users from using short passwords that are easy to guess. There is a minimum password length of 9 characters required by default. **Note:** In Version 11.2 and earlier versions, the minimum password length is 8. In Version 11.3, the minimum password length changed to 9. On upgrade or update from earlier versions to Version 11.3, users can still create a password with 8 characters until you explicitly set the minimum password length to 9 characters.
Uppercase	Specifies a minimum number of uppercase characters for the password. This includes European language characters A through Z, with diacritic marks, Greek characters, and Cyrillic characters. For example: • Cyrillic uppercase: Д Ц • Greek uppercase: Π Λ
Lowercase	Specifies a minimum number of lowercase characters for the password. This includes European language characters a through z, sharp-s, with diacritic marks, Greek characters, and Cyrillic characters. For example: • Cyrillic lowercase: д ц • Greek lowercase: π λ
Decimal Digits	Specifies a minimum number of decimal characters (0 through 9) for the password.
Special (~!@#$%^&*_-+=`|'(){}[];:<>,".?/)	Specifies a minimum number of special characters for the password: `~!@#$%^&*_-+=`|'(){}[];:<>,".?/`
Non-Latin Alphabetic	Specifies a minimum number of Unicode alphabetic characters that are not uppercase or lowercase. This includes Unicode characters from Asian languages. For example: • Kanji (Japanese): 頁 (leaf) 枒 (tree)
Password May Not Contain Username	Specifies that a password cannot contain the case-insensitive username of the user.

4. If you want your password policy changes to take effect at the next login instead of the next password change, select **Force all internal users to change their passwords on the next login**. Note that this setting is selected by default.
5. Click **Apply**.
   The password strength settings take effect when internal users create or change their passwords. If you selected **Force all internal users to change their passwords on the next login**, all internal users must change their password the next time they log on to NetWitness.

Change the Default Admin Passwords

The system administrator's user account is installed with NetWitness. The username is **admin** and the default password is the password that was entered in the Text-based User Interface (TUI) during the NetWitness installation process. The **Administrators** role is assigned to admin. This role has full system privileges to control what a user can do and which services a user can access. The only modification you can make to this account is to change the password. Unlike other NetWitness users, changes to the **admin** user password do not automatically propagate to downstream services. When you configure the password strength settings, they apply to all NetWitness users, including the admin user.

Passwords, an important aspect of computer security, are the front line of protection for your system. The **admin** user is pre-installed in NetWitness and on each Core service. For security, you create the users and roles for your organization in NetWitness, and on each Core service.

Best Practices

NetWitness recommends the following best practices:

- Change the **admin** password of each service from the default.
- Create a different password for the **admin** account on each service.

Change the admin Password for the NetWitness

Change the **admin** password for the NetWitness in the Profile view. See "Change Password" in the *NetWitness Getting Started Guide*. The password of the **admin** user does not propagate to Core services.

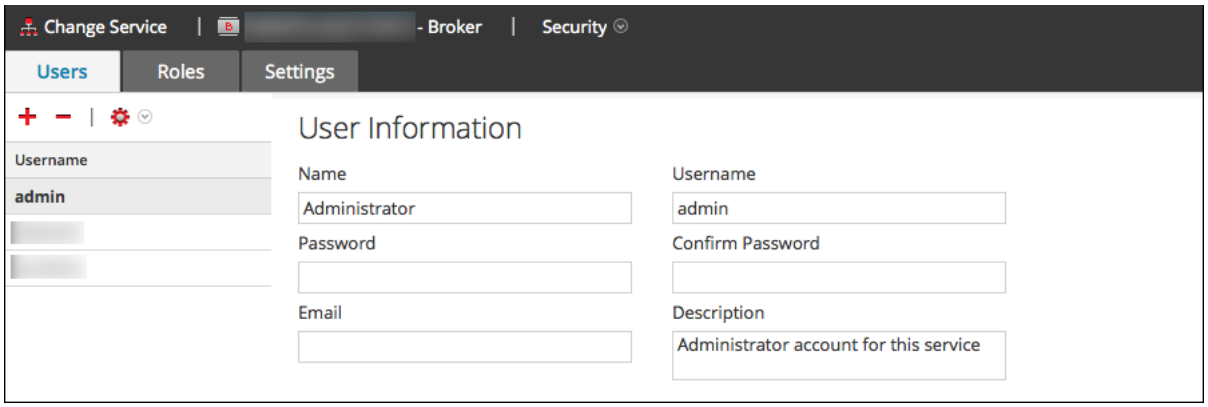
Note: After you change the admin password, you must remove and re-add a data source on the Reporting Engine. For more information, see the **Remove and re-add a Data Source on the Reporting Engine** section below.

Change the admin Password for Core Services

To change the admin password for a Core service:

1. Go to  **(Admin) > Services.**
2. Select a service, and then select  **> View > Security**.

3. On the **Users** tab, select the **admin** user.



4. In the **Password** field, type a new admin password for the selected service.
5. In the **Confirm Password** field, retype the new password.
6. Click **Apply**.

Note: After you change the admin password, you must remove and re-add a data source on the Reporting Engine. For more information, see **Remove and re-add a Data Source on the Reporting Engine** below.

Remove and Re-add a Data Source on the Reporting Engine

Reporting Engine validates a data source using the data source username and password. If you change the username or password of a data source, you must remove and re-add the data source.

To remove and re-add a data source on the Reporting Engine:

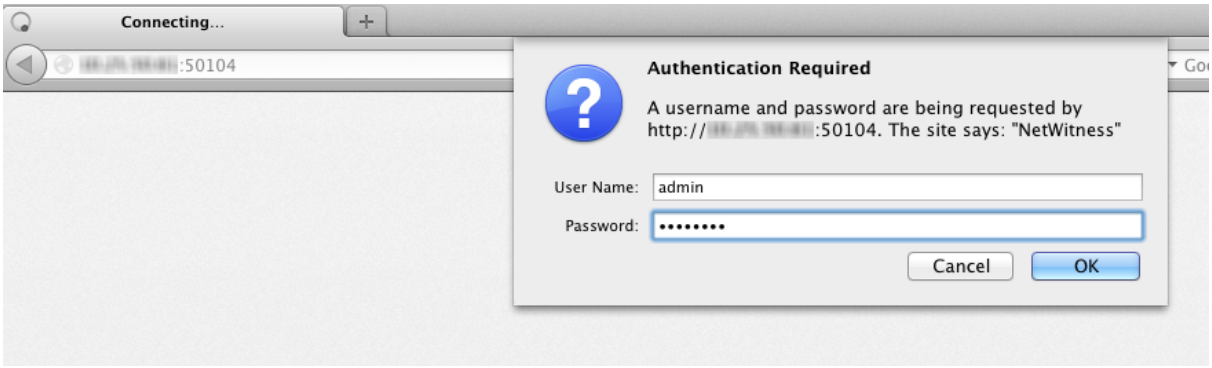
1. Go to **(Admin) > Services**.
2. In the Services view, select Reporting Engine and **View > Config**.
3. Click the **Sources** tab.
4. Select a service to remove and click .
5. Click and select **Available Services.**
6. Select the service you removed in step 4 and click **OK**.
7. When prompted, enter the new username and password for the service.

Change the admin Password for a Service Using the REST API

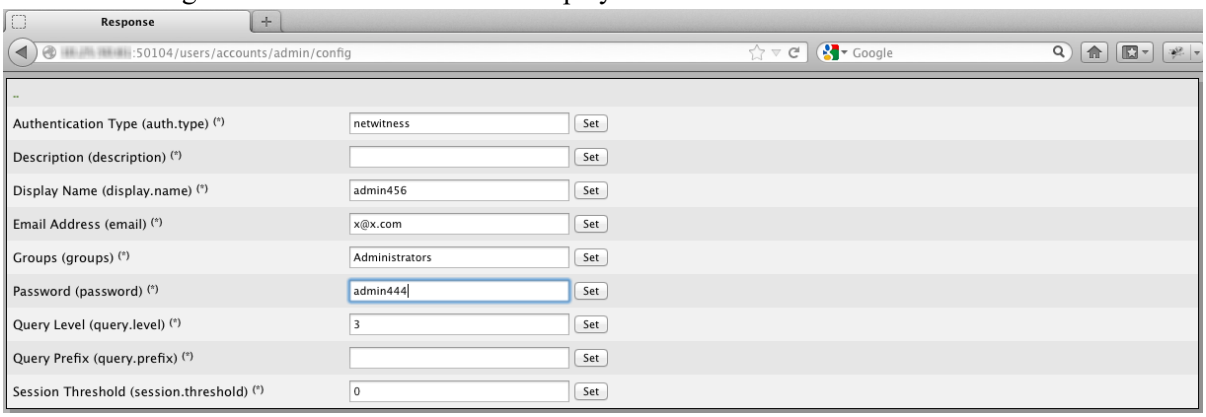
In rare circumstances, you may need to change the admin password for a Core service outside of the NetWitness user interface. This is simply another way to perform the Core service password change, and is not the preferred method.

To change the admin password for the service using the REST User Interface:

1. Open a web browser, and go to the following URL:
   <hostname>:<port>
   where the **hostname** is the name of a NetWitness Core service and **port** is the port used for REST communication. Here is an example for a Decoder: `http://10.20.30.40:50104`
   The authentication dialog is displayed.



2. In the dialog, enter the user name and password used for authentication as admin on the service, and click **OK**. The default user name is **admin** and the default password is **netwitness**.
   The REST window for the service is displayed.
3. Navigate through the node structure to **users/accounts/admin/config**.
   The user configuration fields for admin are displayed in the browser window.



4. In the Password field, type a new admin password and click **Set**.

Configure System-Level Security Settings

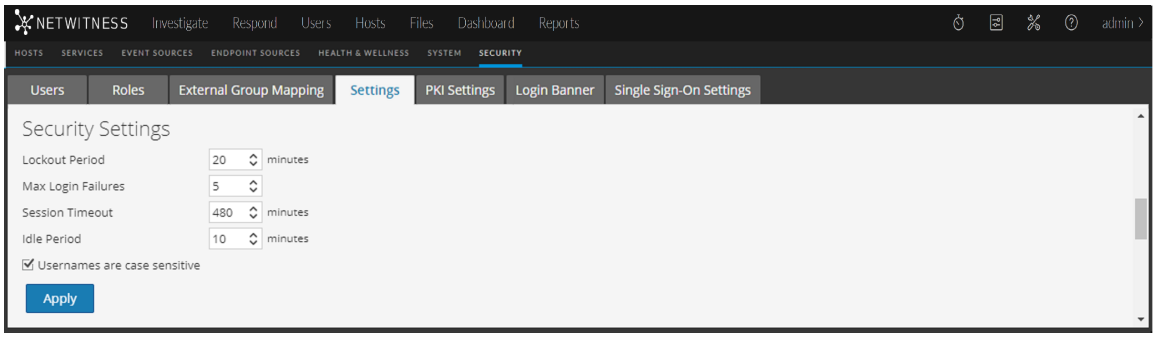
Most global security settings, such as the maximum number of failed login attempts to allow, apply to all NetWitness users and sessions. Settings related to passwords in the Password Strength section, such as password expiration period and the default number of days before user passwords expire, apply to internal NetWitness users, but not external users.

Configure Security Settings

1. Go to  **(Admin) > Security**.

   The Security view is displayed with the **Users** tab open.

2. Click the **Settings** tab.



3. In the **Security Settings** section, specify values for the fields as described in the following table.

Field	Description
Lockout Period	Number of minutes to lock a user out of NetWitness after the configured number of failed logins is exceeded. The default value is 20 minutes.
Max Login Failures	The maximum number of unsuccessful login attempts before a user is locked out. The default value is 5.
Session Timeout	The maximum duration of a user session before timing out in minutes. The default value is 480. The session times out when the configured time has elapsed, after which the user must log in again. The maximum allowed value is 30,000. **Note:** If you migrated to NetWitness 11.x from version 10.6.x and previously used a value of 0 for an unlimited session timeout, the value was reset automatically to 30,000 minutes, as a value of 0 is no longer supported.

Field	Description
Idle Period	Number of minutes of inactivity before a session times out. The default value is 10. The maximum allowed value is 30,000. **Note:** If you migrated to NetWitness 11.x from version 10.6.x and previously used a value of 0 for an unlimited idle period, the value was reset automatically to the default value of 10, as a value of 0 is no longer supported.
Usernames are case sensitive	Select this option if you want the Username field on the NetWitness login to be case sensitive. For example, if usernames are case sensitive, you could use admin to log on to NetWitness, but you could not use Admin.

4. Click **Apply**. The Security Settings take effect immediately. If a password expires, the user receives a prompt to change the password when they log on to NetWitness.

Restrict Access to Incidents

By default, analysts can view all of the incidents, alerts, and tasks in the Respond view. If you have sensitive or restricted information that should not be shared, you can restrict what analysts and other users can see in the Respond view.

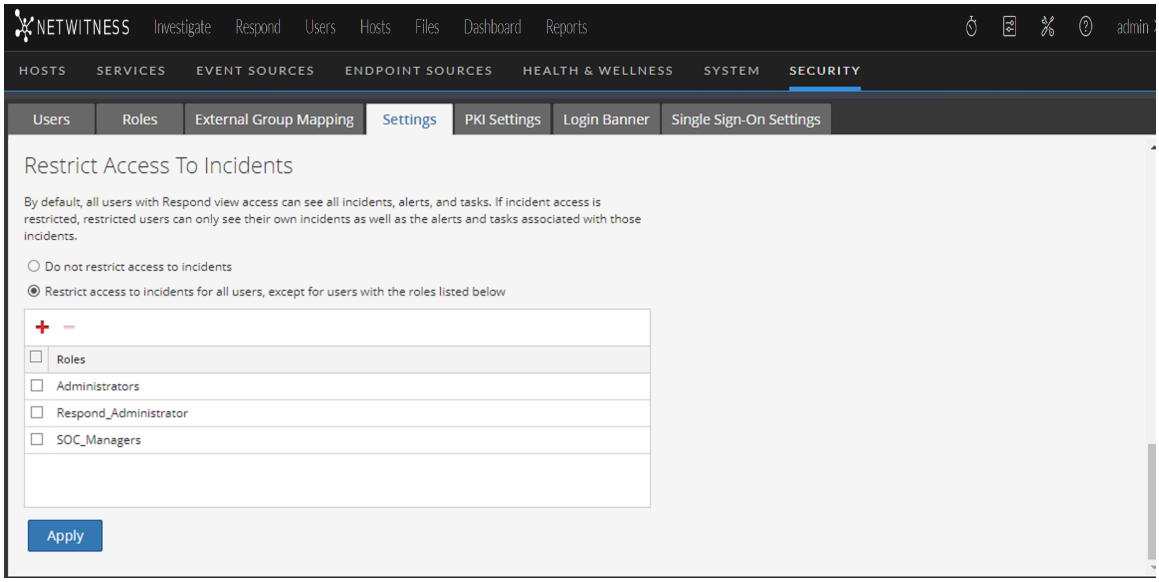
If you restrict access to incidents:

- Analysts can only see incidents assigned to them as well as the alerts and tasks associated with those incidents. Likewise, they can only change the status of and add journal entries (notes) to their own incidents.
- Analysts cannot see the Alerts and Tasks tabs in the Respond view (Respond > Tasks and Respond > Alerts are hidden), so they cannot view all alerts and tasks.
- Analysts cannot see the Assignee button or change the assignee of an incident.
- Analysts cannot see the Related Indicators (alerts) panel (Incident Details view > Find Related tab in the left-side panel).
- When adding events to incidents from the Investigate views, users can only add events to incidents to which they have access. The list of incidents to which users can add events only shows incidents that the user can access.
- When creating incidents from the Investigate views, users must have access to those incidents to view them in the Respond view. For example, when creating incidents from the Investigate view, Analysts must assign the incidents to themselves to view them in the Respond view.

Caution: These restrictions apply to all NetWitness users, except users with the **Administrators**, **Respond_Administrator**, and **SOC_Managers** roles. However, you can adjust the list of user roles whose access to incidents should not be restricted.

To restrict access to incidents:

1. Go to **(Admin) > Security** and click the **Settings** tab.
2. In the **Restrict Access to Incidents** section, select **Restrict access to incidents for all users, except for users with the roles listed below**.



3. In the list, add the user roles whose access to incidents should not be restricted.
4. Click **Apply**.
   Changes take effect on the next log in to NetWitness.

(Optional) Configure External Authentication

When a user logs in, NetWitness first attempts to authenticate locally. If no local user is found, and External Authentication configuration is enabled, an attempt is made to authenticate externally.

External authentication allows users who do not have an internal NetWitness user account to log on to NetWitness and receive role-based permissions.

NetWitness supports two methods of external authentication, Active Directory and Pluggable Authentication Modules (PAM). The below procedures describe how to configure and test each method.

Configure Active Directory

When a user logs in, NetWitness first attempts to authenticate locally. If no local user is found, and Active Directory configuration is enabled, an attempt is made to authenticate with Active Directory Service. You can configure Active Directory settings to enable authentication of external groups in the  (Admin) > Security view > Settings tab.

In an environment with multiple authentication servers, LDAP forwarding allows LDAP referral following for AD group lookups. LDAP forwarding can increase the time required to log on because AD group lookups are extended to connected authentication servers. When your AD instance attempts to contact domain controllers that are blocked by your firewall, users can experience a delay of several minutes in logging on to NetWitness. NetWitness has a configuration option that specifies whether LDAP forwarding occurs; by default, LDAP referrals are disabled. When disabled, your AD instance does not attempt to contact referred domain controllers.

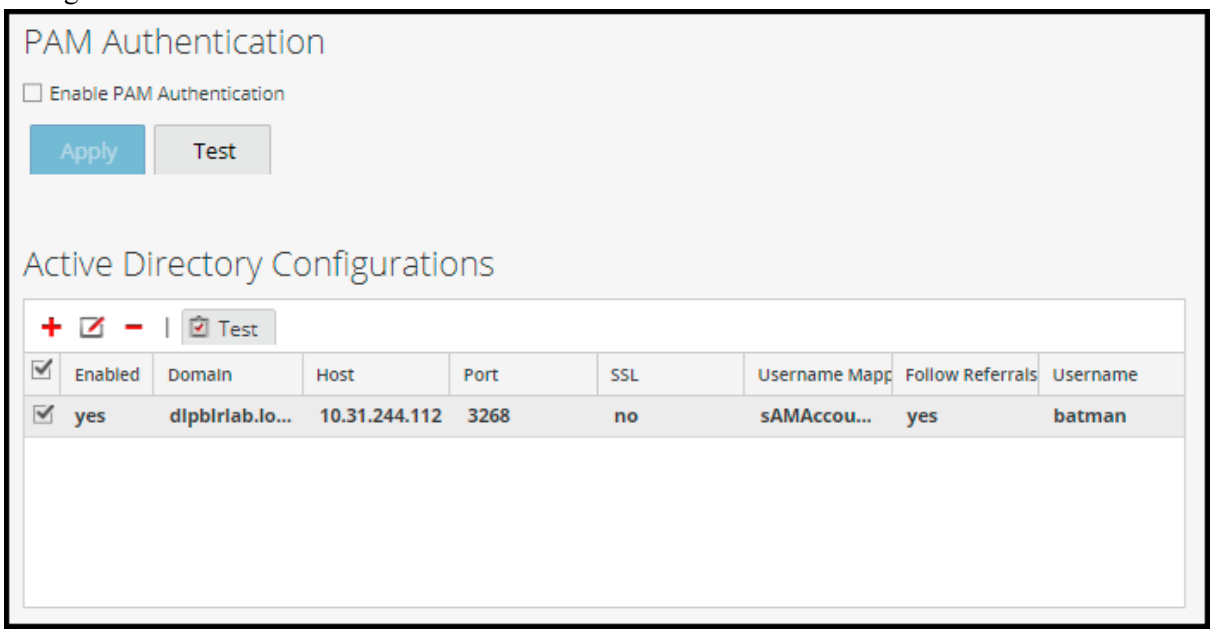
Note: The Settings tab also provides the option to enable PAM configuration, which can be used simultaneously with Active Directory configurations. For information on enabling and configuring PAM authentication, see Configure PAM Login Capability.

Configure Active Directory Authentication

To configure Active Directory authentication:

1. Go to  **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **Settings** tab.
   The Active Directory Configurations list is displayed in the panel so that you can add or edit a

configuration.



3. Add, edit, or delete domains as necessary, as described in the following sections.
   The domains added to this list are automatically populated in the External Group Mapping tab so that you can map security roles to each group.

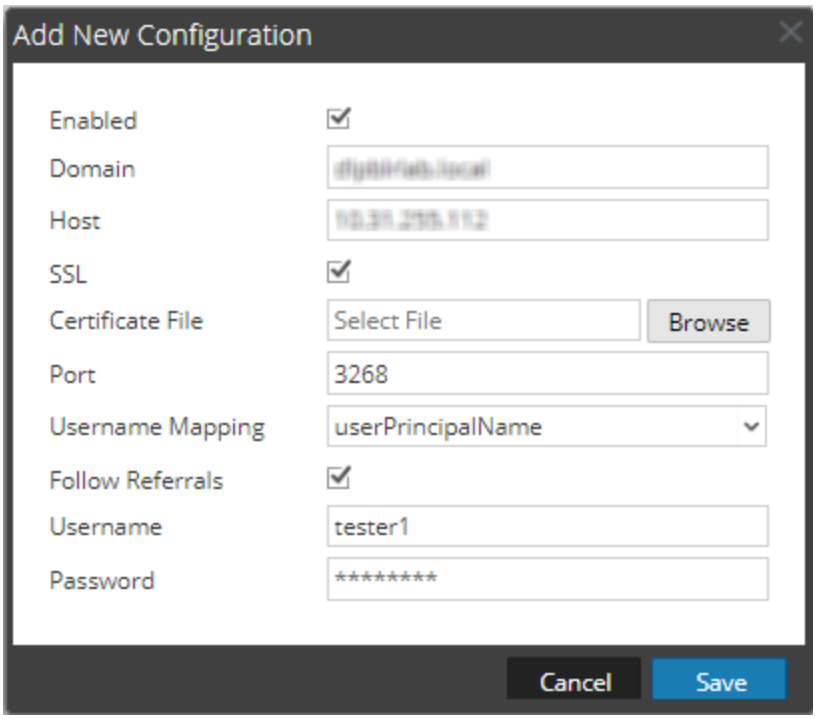
Note: To configure security roles used for Active Directory access, see (Optional) Map User Roles to External Groups.

Note: After configuring an active directory with a domain in **Admin** > **Security** > **Settings**, once you authenticate successfully using the respective active directory credentials, you are tied to that specific active directory and the domain thus configured.

Add a New Active Directory Configuration

To add a new active directory configuration in the Active Directory Configurations list:

1. Under Active Directory Configurations, click +.
   The Add New Configuration dialog is displayed.



2. Select the **Enabled** checkbox.
3. Enter **Domain**, **Host** and **Port** information for the Active Directory Service.
4. (Optional) To select SSL for this configuration, select the **SSL** checkbox. You must then enter the Active Directory server certificate file by clicking **Browse** and selecting the desired file to upload.
5. In the **Username Mapping** field, select the Active Directory search field to use for username mapping. You can select userPrincipalName (UPN) or sAMAccountName.
6. For sites that have multiple authentication servers, click **Follow Referrals** to enable or disable LDAP referral following for AD group lookups.
7. In the **Username** and **Password** fields, enter the username and password for a bind user to access Active Directory. This is usually a service account that has permissions to query the domain and validate user accounts and group membership.

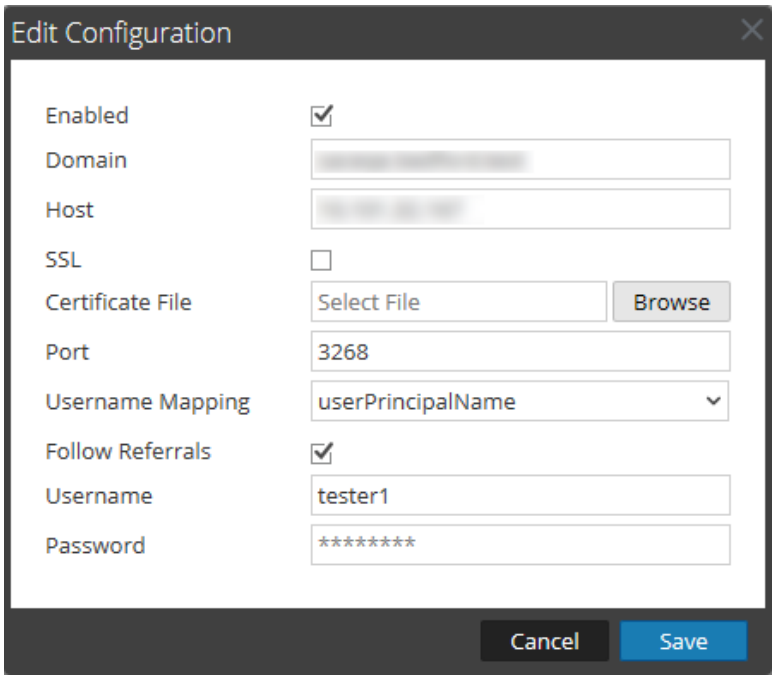
Note: If you selected sAMAccountName in the **Username Mapping** field, you must enter the username in the format "domain\user" to authenticate.

8. Click **Save**.
   The new configuration is listed in the Active Directory Configurations list.

Edit an Active Directory Configuration

To edit an active directory configuration in the Active Directory Configurations list:

1. Under **Active Directory Configurations**, select the configuration you wish to edit and click .
   The Edit Configuration dialog is displayed.



2. (Optional) Enter the **Domain**, **Host** and **Port** information for the Active Directory Service.
3. (Optional) To select SSL for this configuration, select the **SSL** checkbox. You must then enter the Active Directory server certificate file by clicking **Browse** and selecting the desired file to upload.
4. (Optional) In the **Username Mapping** field, select the the Active Directory search field to use for username mapping.
5. To specify the Follow LDAP referrals behavior in environments with multiple authentication servers, select the **Follow Referrals** checkbox.
   a. If you want to disable LDAP forwarding, clear the box.
   b. If you want to enable LDAP forwarding, select the box.
6. In the **Username** and **Password** fields, enter the username and password for a bind user to access Active Directory. This is usually a service account that has permissions to query the domain and validate user accounts and group membership.
7. Click **Save**.
   The configuration is listed in the Active Directory Configurations list.

Test an Active Directory Configuration

To test an Active Directory configuration:

1. Select the configuration to be tested from the Active Directory Configurations list.
2. In the toolbar, click Test.
   A message that the test is successful is displayed.
3. If the test does not succeed, review and edit the configuration.

Delete an Active Directory Configuration

To delete an Active Directory configuration:

1. Under Active Directory Configurations, select the configuration to be deleted from the Active Directory Configurations list.
2. In the toolbar, click −.
   A message is displayed warning you that all users in the selected Active Directory configuration will not be able to log in to NetWitness if it is deleted.
3. Do one of the following:
   a. To confirm the deletion, click **Yes**.
   b. To cancel the deletion, click **No**.

Configure PAM Login Capability

Pluggable Authentication Module (PAM) login capability involves two separate components:

- PAM for user authentication
- NSS for group authorization

Together they provide external users the capability to log on to NetWitness without having an internal NetWitness account, and to receive permissions or roles determined by mapping the external group to a NetWitness security role. Both components are required for a login to succeed.

External authentication is a system-level setting. Before configuring PAM, carefully review all of the information here.

Pluggable Authentication Modules

PAM is a Linux-provided library responsible for authenticating users against authentication providers such as RADIUS, Kerberos, and Agent for SecurID. For implementation, each authentication provider uses its own module, which is in the form of an operating system (OS) package such as pam_krb5. NetWitness uses the OS-provided PAM library, and the module that the PAM library is configured to use, to authenticate users.

Note: PAM provides only the ability to authenticate.

Name Service Switch

NSS is a Linux feature that provides databases that the OS and applications use to discover information like hostnames; user attributes like home directory, primary group, and login shell; and to list users that belong to a given group. Similar to PAM, NSS is configurable and uses modules to interact with different types of providers. NetWitness uses OS-provided NSS capabilities to authorize external PAM users by looking up whether a user is known to NSS and then requesting from NSS the groups of which that user is a member. NetWitness compares the results of the request to the NetWitness External Group Mapping and if a matching group is found, the user is granted access to log on to NetWitness with the level of security defined in the External Group Mapping.

Note: NSS does not provide authentication.

PAM and NSS Combination

Both PAM (authentication) and NSS (authorization) must succeed in order for an external user to be allowed to log on to NetWitness. The procedure for configuring and troubleshooting PAM is different than the procedure for configuring and troubleshooting NSS. The PAM examples in this guide include Kerberos, RADIUS, and SecurID. The NSS example includes UNIX. The PAM and NSS module combination used is determined by site needs.

Process Overview

To configure PAM login capability, follow the instructions in this document to complete each step:

1. Configure and test the PAM module.
2. Configure and test the NSS service.
3. Enable PAM in NetWitness Server.
4. Create group mappings in NetWitness Server.

Before beginning the setup of PAM, review the procedure and gather the external authentication server details depending on the PAM module you want to implement.

Before beginning the setup of NSS, review the procedure, identify the group names that you will use in the External Group mapping, and gather the external authentication server details, depending on the NSS service being used.

Before beginning setup of PAM in NetWitness, identify the group names that you will use in the External Group mapping. When mapping roles, the role in NetWitness must match a group name that exists in the external authentication server.

Configure and Test the PAM Module

Choose one of the following sections to set up and configure the PAM component:

- PAM Kerberos
- PAM RADIUS
- PAM Agent for SecurID

PAM Kerberos

Kerberos Communication Ports – TCP 88

To configure PAM authentication using Kerberos:

1. Execute the following command (but first verify that the `krb5-workstation` package is installed in your environment):
   `yum install krb5-workstation pam_krb5`
2. Edit the following lines in the Kerberos configuration file `/etc/krb5.conf`. Replace variables, which are delimited by <angle brackets>, with your values and omitting the angle brackets. Capitalization is required where shown.

```
# Configuration snippets may be placed in this directory as well
includedir /etc/krb5.conf.d/

[logging]
default = FILE:/var/log/krb5libs.log
kdc = FILE:/var/log/krb5kdc.log
admin_server = FILE:/var/log/kadmind.log

[libdefaults]
 dns_lookup_realm = false
 ticket_lifetime = 24h
 dns_lookup_kdc = true
 renew_lifetime = 7d
 forwardable = true
 rdns = false
 default_realm = <DOMAIN.COM>
 default_ccache_name = KEYRING:persistent:%{uid}

[realms]
 <DOMAIN.COM> = {
 kdc = <SERVER.DOMAIN.COM>
 admin_server = <SERVER.DOMAIN.COM>
 }

 [domain_realm]
 <domain.com> = <DOMAIN.COM>
 <.domain.com> = <DOMAIN.COM>
```

3. Test the Kerberos configuration with the command:
   `kinit <user>@<DOMAIN.COM>`
   No output after entering the password indicates success.
4. Edit the NetWitness Server PAM configuration file `/etc/pam.d/securityanalytics` to add the following line. If the file does not exist, create it and add the following line:
   `auth sufficient pam_krb5.so no_user_check`

This completes the configuration for PAM Kerberos. Now, go to the next section, Configure and Test the NSS UNIX Service.

PAM RADIUS

Radius Communication Ports - UDP 1812 or UDP 1813

To configure PAM authentication using Radius you must add the NetWitness Server to your Radius Server's Client list and configure a shared secret. Contact the Radius Server Administrator for this procedure.

To configure PAM authentication using RADIUS:

1. Execute the following command (but first verify that the `pam_radius_auth` package is installed in your environment):

   ```
   yum install pam_radius_auth
   ```

2. Edit the RADIUS configuration file, `/etc/raddb/server` as follows:

   ```
   # server[:port] shared_secret  timeout (s)
   server    secret     3
   ```

3. Edit the NetWitness Server PAM configuration file `/etc/pam.d/securityanalytics` to add the following line. If the file does not exist, create it and add the following line:

   ```
   auth sufficient pam_radius_auth.so
   ```

4. Execute the following command to copy the RADIUS library:

   ```
   cp /usr/lib/security/pam_radius_auth.so /usr/lib64/security/
   ```

Caution: For PAM RADIUS to work, the `/etc/raddb/server` files must have write permission. The command needed for this is: `chown netwitness:netwitness /etc/raddb/server`.

Caution: You must restart the Jetty server after making the above changes for PAM RADIUS. The command for this is:
`systemctl restart jetty`

The PAM Modules and associated services output information to `/var/log/messages` and `/var/log/secure`. These outputs can be used to assist in troubleshooting configuration problems.

The following procedure is an example of the steps to configure PAM authentication for RADIUS using SecurID:

Note: The examples in these tasks use Authentication Manager as the RADIUS server.

1. Execute the following command (but first verify that the `pam_radius_auth` package is installed in your environment):

   ```
   yum install pam_radius_auth
   ```

2. Edit the RADIUS configuration file, `/etc/raddb/server` and update it with the authentication manager instance hostname, shared secret and timeout value:

   ```
   # server[:port] shared_secret timeout (s)
    111.222.33.44         secret        1
    #other-server        other-secret 3
    192.168.12.200:6369 securid      10
   ```

Note: You must comment out `127.0.0.1` and `other-server` lines and add the IP address of the authentication manager primary instance with RADIUS port number (for example, `192.168.12.200:1812`), RADIUS shared secret, and a timeout value of 10.

3. Edit the NetWitness Server PAM configuration file `/etc/pam.d/securityanalytics` to add the following line. If the file does not exist, create it and add the following line:

   ```
   auth sufficient pam_radius_auth.so
   ```

   Note: You can add `debug` to the end of the above line in the `/etc/pam.d/securityanalytics` file to enables PAM debugging (for example, `auth sufficient pam_radius_auth.so debug`)

4. Run the following command to copy the RADIUS library:

   ```
   cp /usr/lib/security/pam_radius_auth.so /usr/lib64/security/
   ```

The PAM Modules and associated services output information to `/var/log/messages` and `/var/log/secure`.These outputs can be used to assist in troubleshooting configuration problems.

Add a RADIUS Client and Associated Agent

Note: The examples in these tasks use Authentication Manager as the RADIUS server.
You must use administrative account credentials to log on Authentication Manager Security Console.

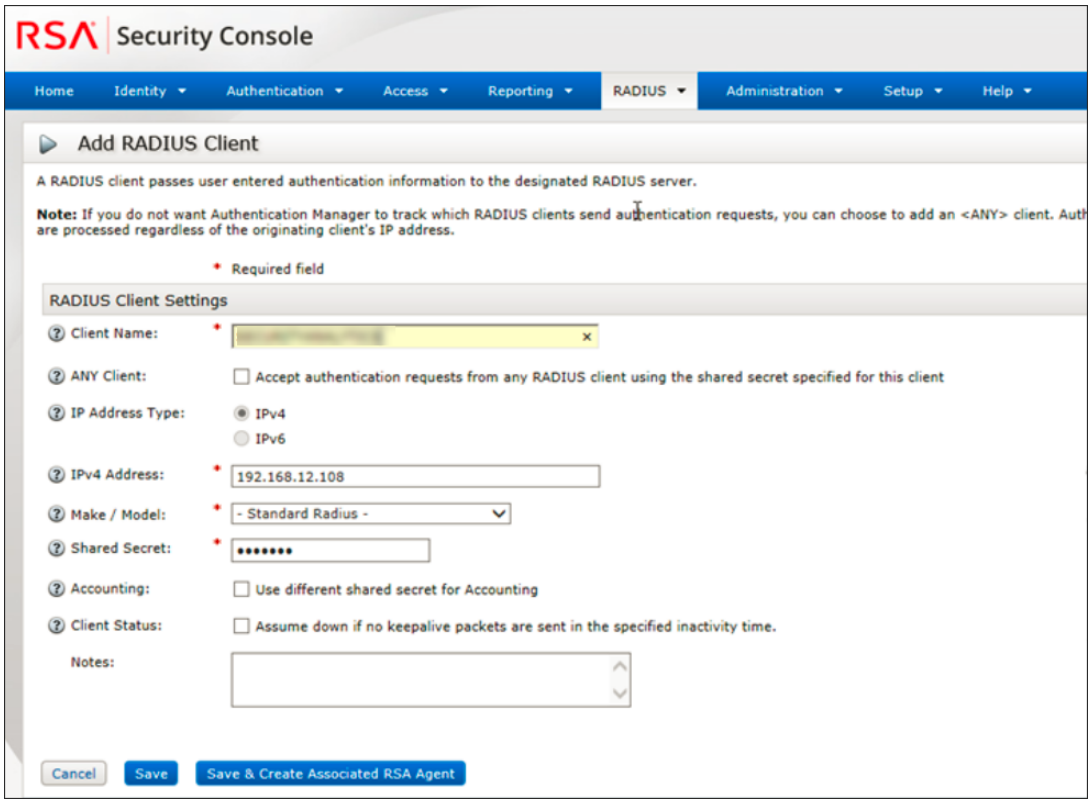
To add a RADIUS Client and Associated Agent:

1. Log on to Authentication Manager.

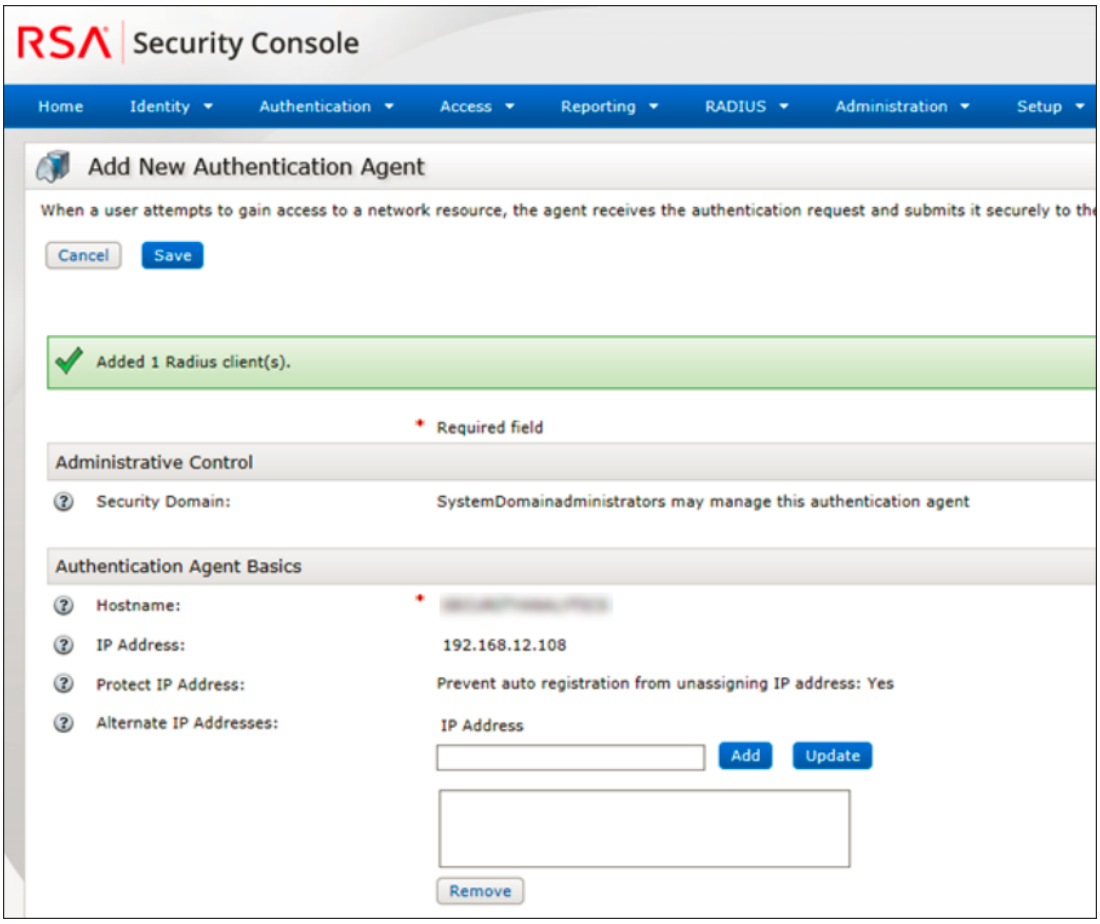
   The Security Console is displayed.

2. In the Security Console, click **RADIUS > RADIUS Client > Add New**.

   The Add RADIUS Client page is displayed.



3. In RADIUS Client Settings, provide the following information:
   a. In the **Client Name** field, enter the name of the client, for example, NetWitness.
   b. In the **IPv4 Address** field, enter the IPv4 address of the RADIUS client, for example, 192.168.12.108.
   c. In the **Make/Model** drop-down list, select the type of RADIUS client, for example, Fortinet.
   d. In the **Shared Secret** field, enter the authentication shared secret.
4. Click **Save & Create Associated Netwitness Agent**.



5. Click **Save**.

If the Authentication Manager instance is unable to find the authentication agent on the network, a warning page is displayed. Click **Yes, Save Agent**.

For more information, see the "Add a RADIUS Client" topic in *Authentication Manager 8.2 Administrator's Guide*.

This completes the configuration for PAM RADIUS. Now, go to the next section, Configure and Test the NSS UNIX Service.

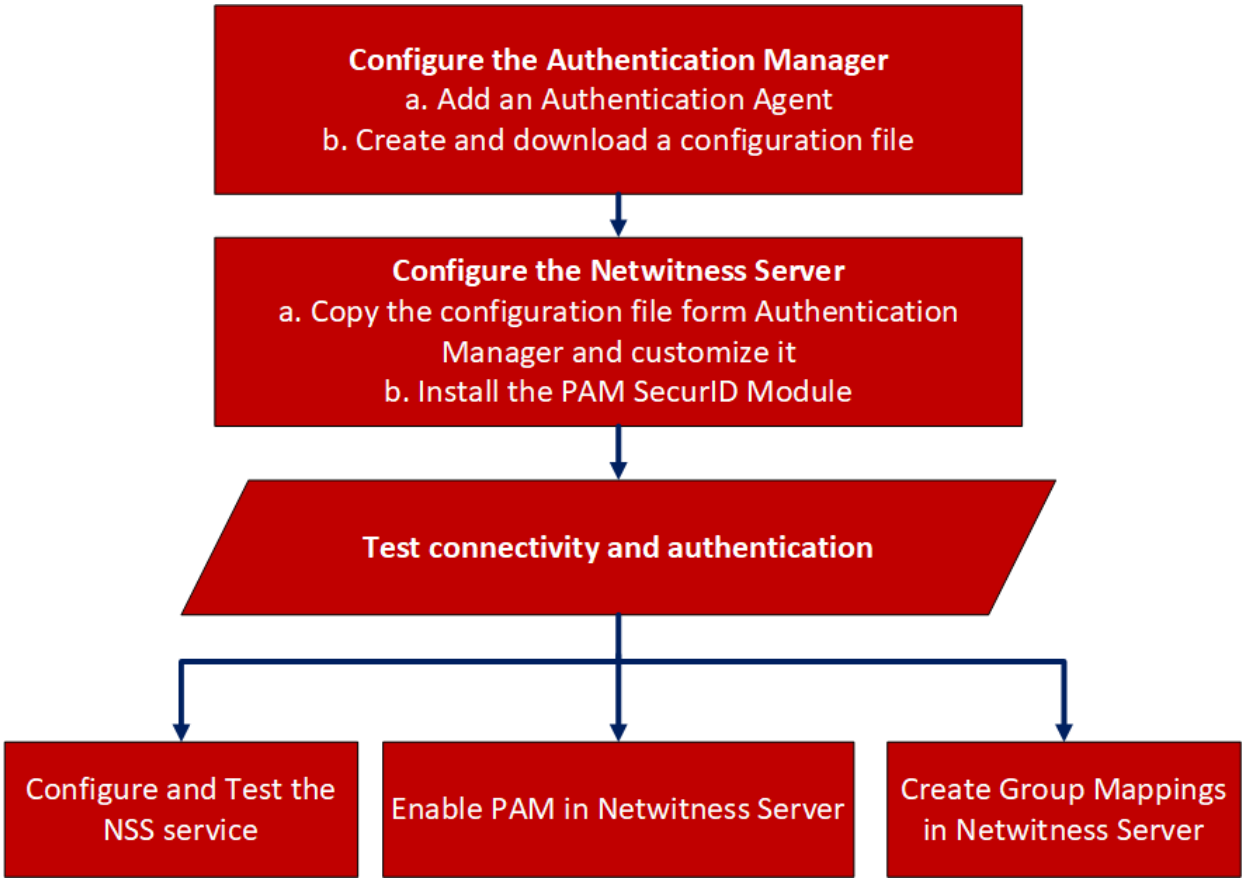
PAM Agent for SecurID

PAM Communication Port - UDP 5500

The SecurID PAM module is supported only under the following condition:

Trusted connections must be enabled and functioning between NetWitness and Core services.

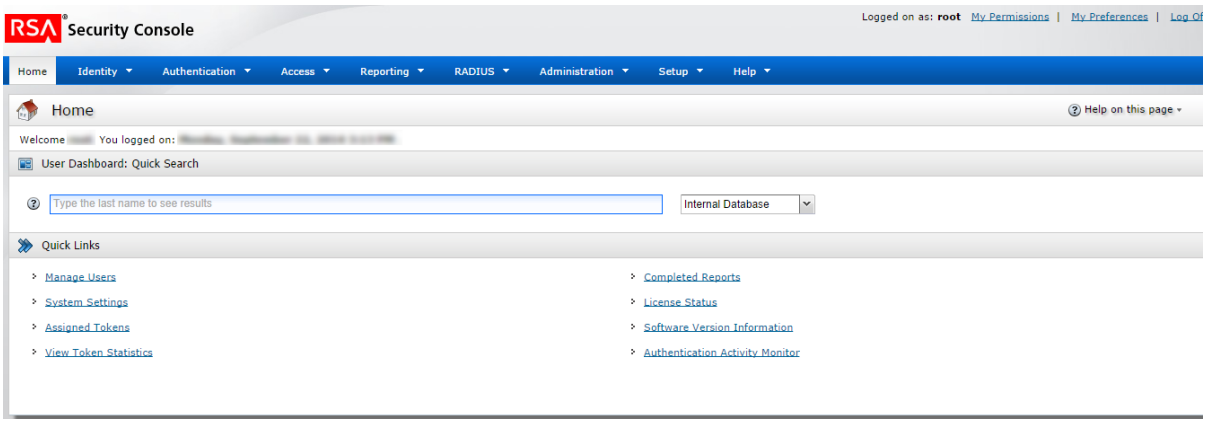
Configure the SecurID PAM Module



To configure Authentication Manager:

1. Log on to Authentication Manager.

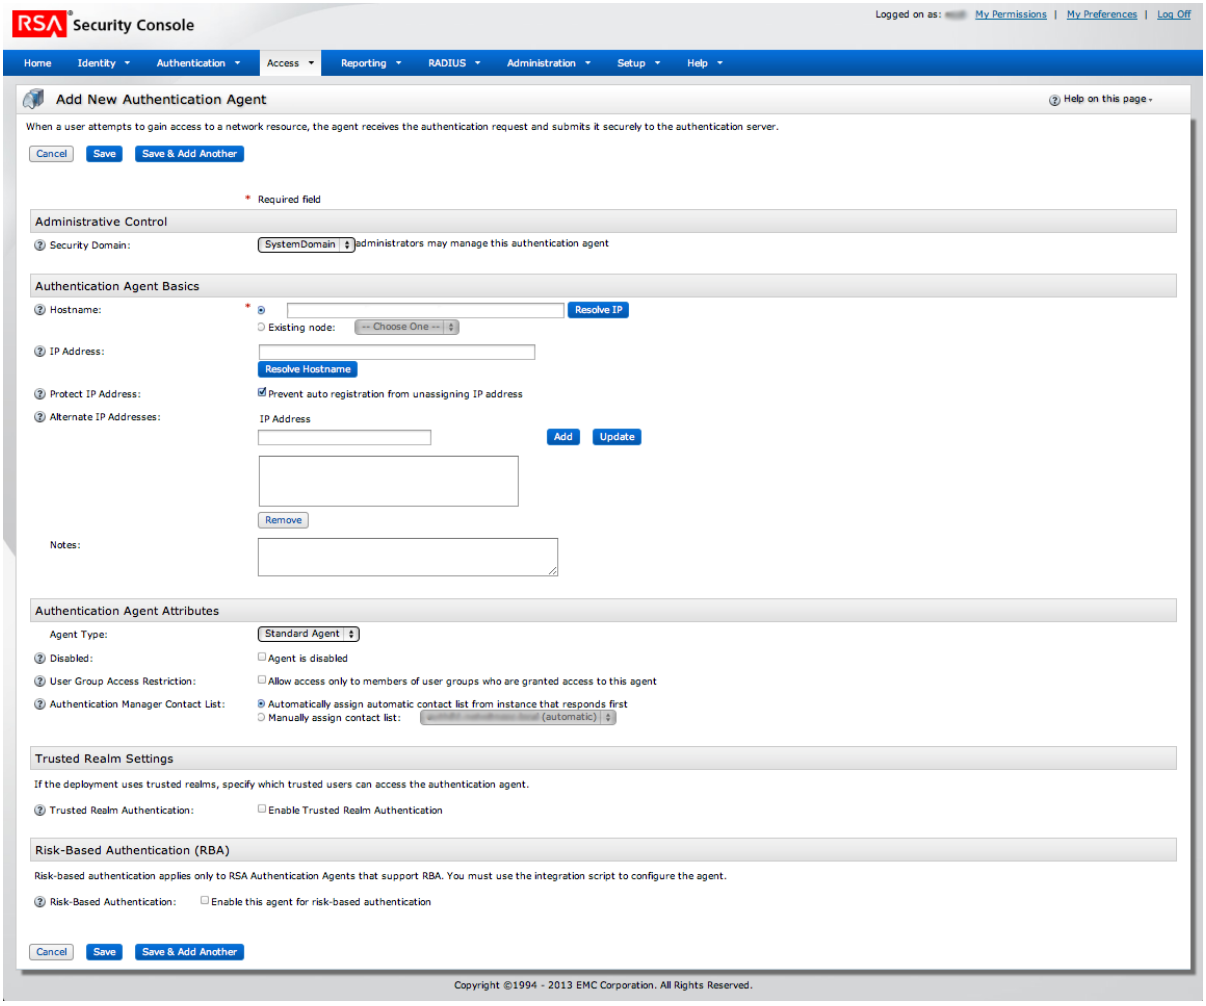
   The Security Console is displayed.



2. In the Security Console, add a new authentication agent.

   Click **Access > Authentication Agents > Add New**.

   The Add New Authentication Agent page is displayed.



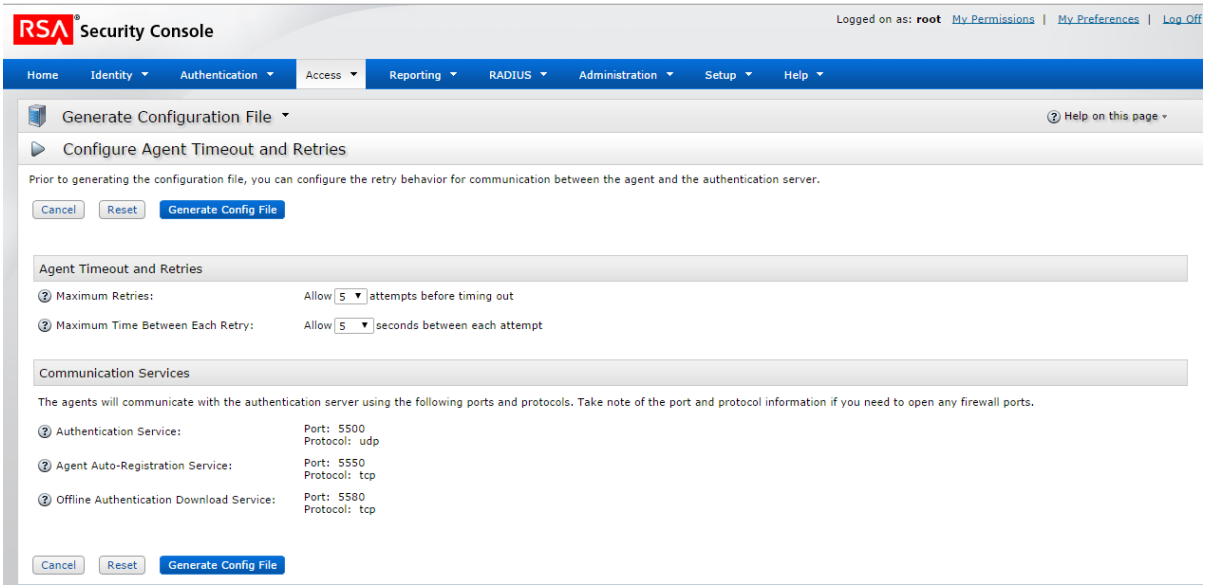
3. In the **Hostname** field, type the hostname of the NetWitness Server.
4. Click **Resolve IP**.

The IP address of the NetWitness Server is automatically displayed in the **IP Address** field.

5. Keep the default settings and click **Save**.
6. Generate a configuration file.

   Go to **Access > Authentication Agents > Generate Configuration File**.

   The Generate Configuration File page is displayed.



7. Keep the defaults and click **Generate Config File**.

   This creates **AM_Config.zip**, which contains two files.

8. Click **Download Now**.

To install and configure the PAM SecurID module:

1. On the NetWitness Server, make the following directory:

   `mkdir /var/ace`

2. On the NetWitness Server, copy `sdconf.rec` from the .zip file to `/var/ace`.
3. Create the text file `sdopts.rec` in the `/var/ace` directory.
4. Insert the following line:

   `CLIENT_IP=<IP address of NetWitness Server>`

5. Install the SecurID Authorization Agent for PAM, which is available in the yum repository:

   `yum install sid-pam-installer`

6. Run the install script:

   `/opt/rsa/pam-agent-installer/install_pam.sh`

7. Follow the prompts to accept or change the defaults.

8. Edit the NetWitness Server PAM configuration file, `/etc/pam.d/securityanalytics` to add the following line. If the file does not exist, create it and add the following line:

   ```
   auth sufficient pam_securid.so
   ```

This completes the installation of the SecurID PAM module. Next, test the connectivity and authentication. Then, follow the procedures in Configure and Test the NSS UNIX Service.

Note: If the PAM SecurID setup is not complete, it may crash the Jetty server and the NetWitness UI will not be displayed. You must wait until the PAM authentication configuration is complete and then restart the Jetty server.

To test connectivity and authentication:

1. Run `/opt/pam/bin/64bit/acetest`, and enter the **username** and **passcode**.
2. (Optional) If acetest fails, turn on debugging:

   ```
   vi/etc/sd_pam.conf
   RSATRACELEVEL=15
   ```

3. Run `/opt/pam/bin/64bit/acestatus`. The output is displayed as shown below.

   ```
   RSA ACE/Server Limits
    ---------------------
    Configuration Version : 15 Client Retries : 5
    Client Timeout : 5 DES Enabled : Yes

   RSA ACE/Static Information
    --------------------------
    Service : securid Protocol : udp Port Number : 5500

   RSA ACE/Dynamic Information
    ---------------------------
    Server Release : 8.1.0.0 Communication : 5

   RSA ACE/Server List
    -------------------
    Server Name :  auth81.netwitness.local
    Server Address :  192.168.100.10
    Server Active Address : 192.168.100.10
    Master : Yes Slave : No Primary : Yes
    Usage : Available for Authentications
   ```

4. (Optional) To troubleshoot the Authentication Manager server, go to **Reporting > Real-time Activity Monitors > Authentication Activity Monitor.** Then click **Start Monitor**.
5. If you changed the setting, reset RSATRACELEVEL to 0:

   ```
   vi/etc/sd_pam.conf
   RSATRACELEVEL=0
   ```

Caution: After installation, verify that VAR_ACE in the `/etc/sd_pam.conf` file points to the correct location of the sdconf.rec file. This is the path to the configuration files. The command needed for this is: `chown -R netwitness:netwitness /var/ace`.

This completes the configuration for PAM Agent for SecurID. Now, go to the next section, Configure and Test the NSS UNIX Service.

Configure and Test the NSS UNIX Service

Configuration

No configuration is necessary to enable the NSS UNIX module; it is enabled in the host operating system by default. To authorize a user for a specific group, simply add that user to the operating system and add them to a group:

1. Create an OS group to use add your external user to with this command:

   ```
   groupadd <groupname>
   ```

2. Add the external user to the OS with this command:

   ```
   adduser -G <groupname> -M -N <externalusername>
   ```

Note: This does NOT permit or allow access to the NetWitness Server console.

This completes the configuration for NSS UNIX. Next, go to Test NSS Functionality.

Test NSS Functionality

To test whether NSS is working with any of the previous NSS services, use the following commands:

```
getent passwd <pamUser>
 getent group <groupOfPamUser>
```

Output should be similar to:

```
[root@~]# getent passwd myuser
 myuser:*:10000:10000::/home/myuser:/bin/sh
[root@~]# getent group mygroup
 mygroup:*:10000:myuser3
```

- If neither command produces output, NSS is not working properly for external authorization. Refer to the troubleshooting guidance for your NSS module provided in this document.
- If `getent` commands succeed and authentication success is confirmed in `/var/log/secure` but NetWitness still fails to allow External users to login:
  - Was the correct group name specified for the NSS group in NW External Group Mapping? See Enable PAM and Create Group Mappings below.
  - It is possible that the NSS configuration has changed and NetWitness has not picked up the change. A reboot of the NetWitness host will cause NetWitness to pick up NSS configuration changes. A restart of the Jetty server is not sufficient.

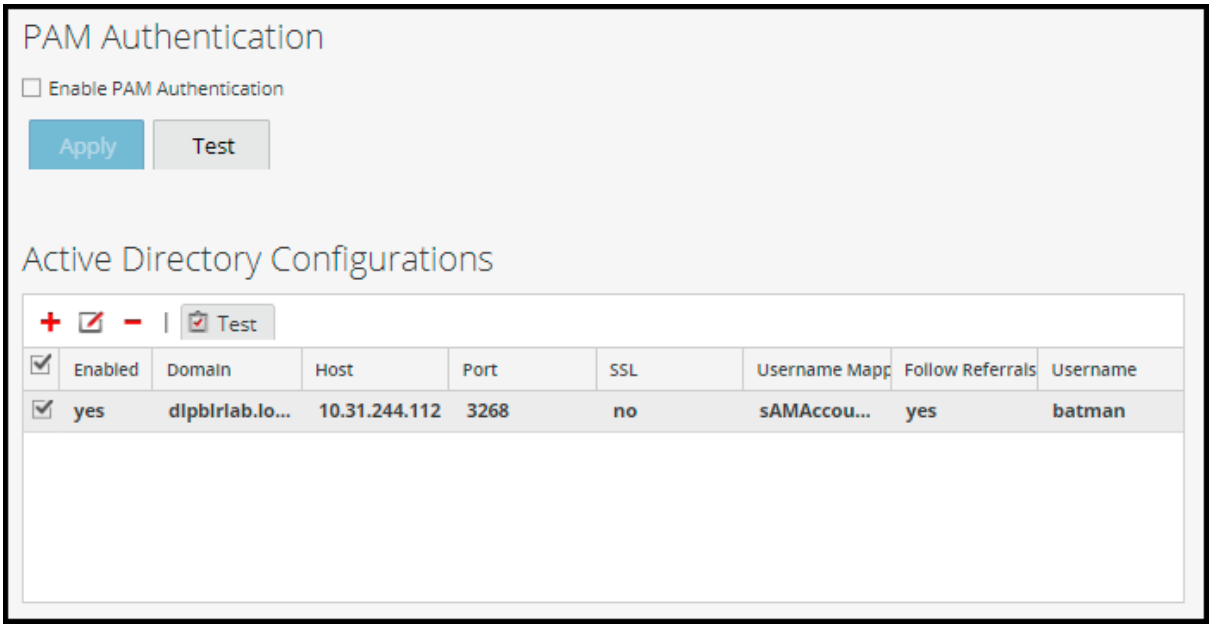
Enable PAM in NetWitness Server

1. Go to  **(Admin) > Security**.

   The Security view is displayed with the Users tab open.

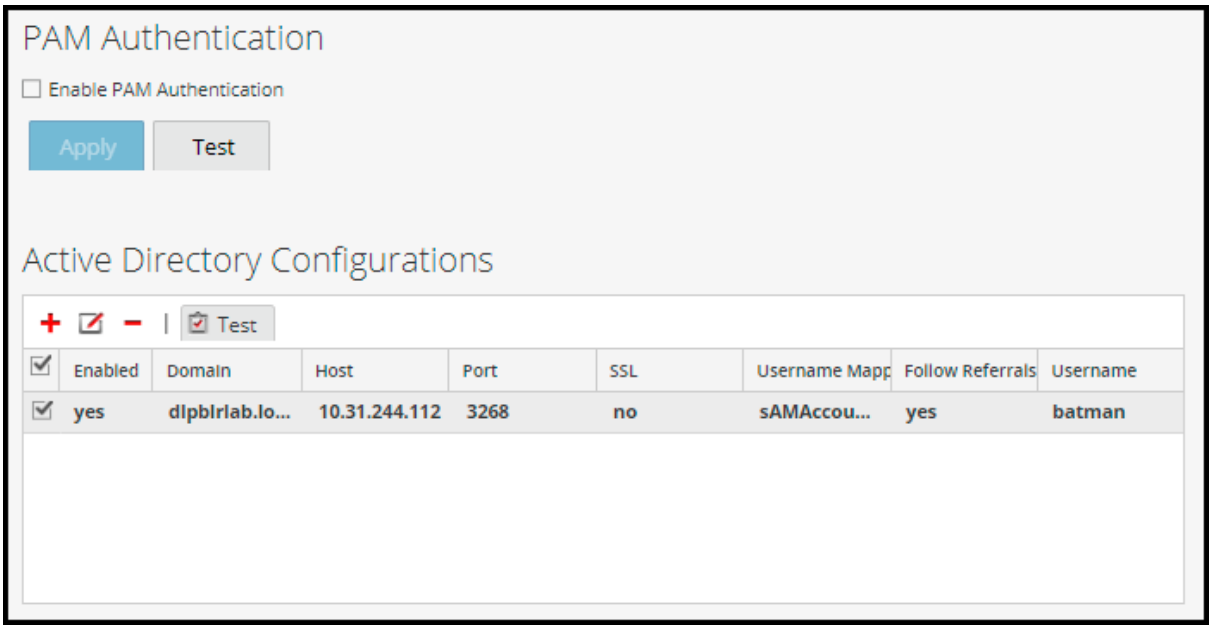
2. Click the **Settings** tab.

3. Under **PAM Authentication**, select **Enable PAM Authentication** and click **Apply**.



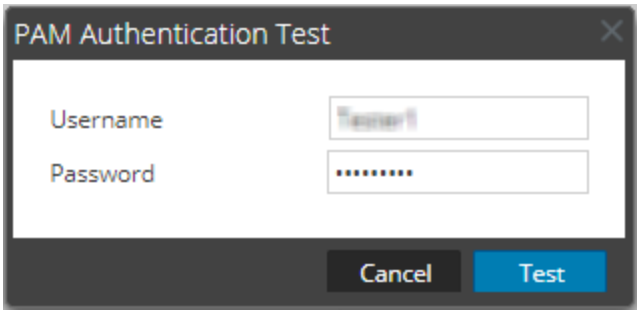
Test External Authentication for PAM

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **Settings** tab.
3. Under **PAM Authentication**, select **Enable PAM Authentication**.



4. Under **PAM Authentication** options, click **Test**.

The PAM Authentication Test dialog is displayed.



5. Type a user name and password that you want to test for authentication using the current PAM configuration.
6. Click **Test**.

   The external authentication method is tested to ensure connectivity.
7. If the test does not succeed, review and edit the configuration.

PAM is enabled, and Active Directory configurations will also remain enabled. PAM configurations are automatically populated in the External Group Mapping tab so that you can map security roles to each group.

Create Group Mappings in NetWitness Server

To configure security roles used for PAM access, see (Optional) Map User Roles to External Groups.

(Optional) Configure PKI Authentication

PKI (Public Key Infrastructure) authentication feature enables the user to securely login to NetWitness Platform using a smart card. The smart card is a hardware device and contains a Certificate and a password protected Private Key corresponding to the Certificate which enables PKI to validate the user. The certificate is provided by the CA and is unique for a user. For more information on how to configure PKI Authentication, see (Optional) Set Up Public Key Infrastructure (PKI) Authentication.

(Optional) Use a Custom Server Certificate

NetWitness also allows you to configure custom web server certificate to be used as NetWitness Server certificate. By default NetWitness Server uses a web server certificate generated by NetWitness for HTTPS connection. You can configure custom web server certificate even if PKI is not enabled.

Supported Keystore Formats

You must select the format that meets your requirement. The following keystore formats are supported:

- For server certificate with its private key:
  - pfx/pkcs/p12 (PKCS8/PKCS12 are the standards)
  - jks (JKS standard)

Note: The .pfx, .p12 and .jks are containers that can contain one or more private keys and its corresponding chains or certificates.

(Optional) Create a Certificate Signing Request (CSR) and Certificate Store for a Server Certificate

Note: The steps provided in this procedure allows you to create a CSR and Certificate Store for a Server Certificate.

If a server certificate is already created along with its private key, you can directly upload the certificate to the NetWitness Server. If the server certificate is not created, based on the CSR created, the CSR can be submitted to the Certificate Authority (CA) server to obtain a server certificate. Once the certificate is created, perform the following steps to package the private key and the signed certificate that must uploaded to the NetWitness Server to be used as a server certificate.

To create a CSR for a Server Certificate:

1. Change the directory to /root:

```
cd /root
```

2. Create a new directory:

```
mkdir nw_pki_server_cert
```

3. Change the directory to the newly created directory:

```
cd nw_pki_server_cert
```

4. Create a Private Key of 2048 Bits:

```
openssl genrsa -out nw_server_pki_private_key.key 2048
```

5. Create a CSR:

```
openssl req -new -sha256 -key nw_server_pki_private_key.key -out server_cert_
request.csr
```

For example, if country: US, location: RT, and unit:

CN: ABCD (Hostname or IP Address of the Machine)

For multiple names, use values such as : CN=ABCD, CN=10.XX.XXX.XX

email: example@rsa.com

6. Check the CSR and Private Key match.

```
openssl req -noout -modulus -in server_cert_request.csr | openssl sha256
openssl rsa -noout -modulus -in nw_server_pki_private_key.key | openssl sha256
```

For example:

```
[root@ABCD open_ssl_test]# openssl rsa -noout -modulus -in server_private.key
| openssl sha256
(stdin)= 88df3d1ea5b2f411712b96d2ed4a72f5
[root@ABCD open_ssl_test]# openssl req -noout -modulus -in server_cert_
request.csr | openssl sha256
(stdin)= 88df3d1ea5b2f411712b96d2ed4a72f5
```

Note: You make a note of both stdin's.

7. Submit the CSR to the CA and get a signed Server Certificate.

8. Copy the certificate in PEM format to the new directory:

```
/root/nw_pki_server_cert/signed_certificate.pem
```

9. Check the certificate for the correct public key.

```
openssl x509 -noout -modulus -in certificate.pem | openssl sha256
```

For example :

```
[root@ABCD open_ssl_test]# mv test.pem certificate.pem
[root@ABCD open_ssl_test]# openssl x509 -noout -modulus -in certificate.pem |
openssl sha256
(stdin)= 3e2f4bbd1f32ae097902afcc1893089e
[root@ABCD open_ssl_test]# openssl rsa -noout -modulus -in sa_server_pki_
private_key.key | openssl sha256
(stdin)= 3e2f4bbd1f32ae097902afcc1893089e
[root@ABCD open_ssl_test]# openssl req -noout -modulus -in server_cert_
request.csr | openssl sha256
(stdin)= 3e2f4bbd1f32ae097902afcc1893089e
```

10. Copy the Private Key and Certificate to a Key Store.

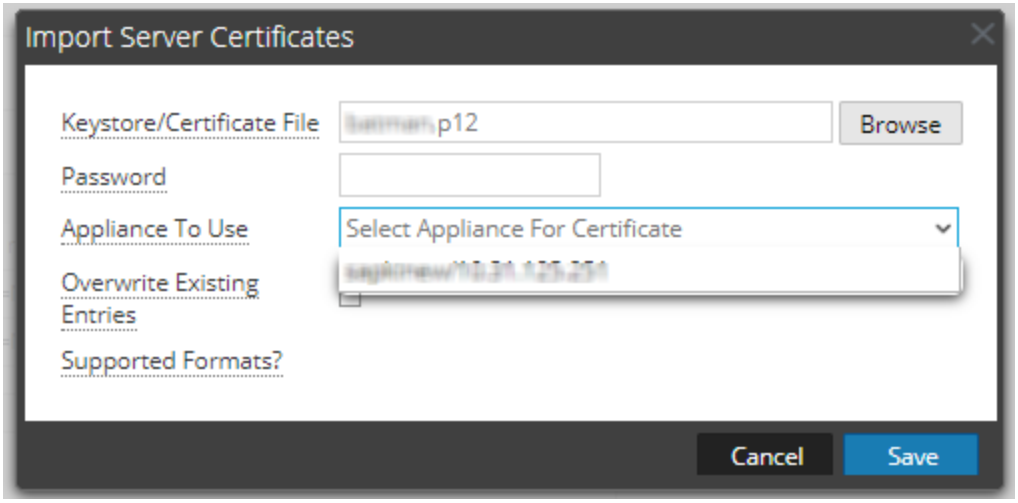
```
openssl pkcs12 -export -descert -name <myservercert> -in signed_
certificate.pem -inkey nw_server_pki_private_key.key -out keystore.p12
```

11. Enter the keystore password, for example **NetWitness@123**, to the Keystore.

Import an NW Server Certificate with its Private Key

Note: .p12, .jks, and .pfx are the supported server certificate formats. Execute the following OpenSSL command to convert certificates to the supported format:
`openssl pkcs12 -export -out cert.p12 -in cert.cer -inkey nw_server_pki_private_key.key`
Only certificates that are configured with an export password can be uploaded to the NetWitness UI.

1. Go to **(Admin) > Security.**
   The Security view is displayed with the **Users** tab open.
2. Click the **PKI Settings** tab.
3. In the **Server Certificates** section, click +.
   The Import Server Certificates dialog is displayed.



4. In the **Keystore/Certificate File** field, click **Browse** and select the keystore.
5. In the **Password** field, enter the keystore password.
6. In the **Appliance To Use** field, select the appliance for which you want to use this certificate.
7. (Optional) Select the **Overwrite Existing Entries** checkbox to overwrite the entries of the certificate that is already added.
8. Click **Save.**
   The NetWitness Server certificate with its private key is successfully added to NetWitness.

Note: When the certificate is being applied on the selected appliance, no other operation on PKI can be performed until the process is completed.
Double-click on the added entries to view the details of the certificate.

9. To apply the server certificate on a server, select a certificate and click .

Note: Uploading a keystore will add the server certificate and its private key locally. To apply a server certificate on a server, you need to select a server certificate and click the synchronization button . All server certificates are also synchronized on the appliances when PKI is enabled.

Note: To complete the certificate update process, log into the Admin Server (node-0) stop and start the nginx with the command `systemctl restart nginx`.

(Optional) Create a Customized Login Banner

You can create and enable customized login banner that is displayed before users log on to NetWitness asking users to agree the conditions. Users who do not agree are not able to log on.

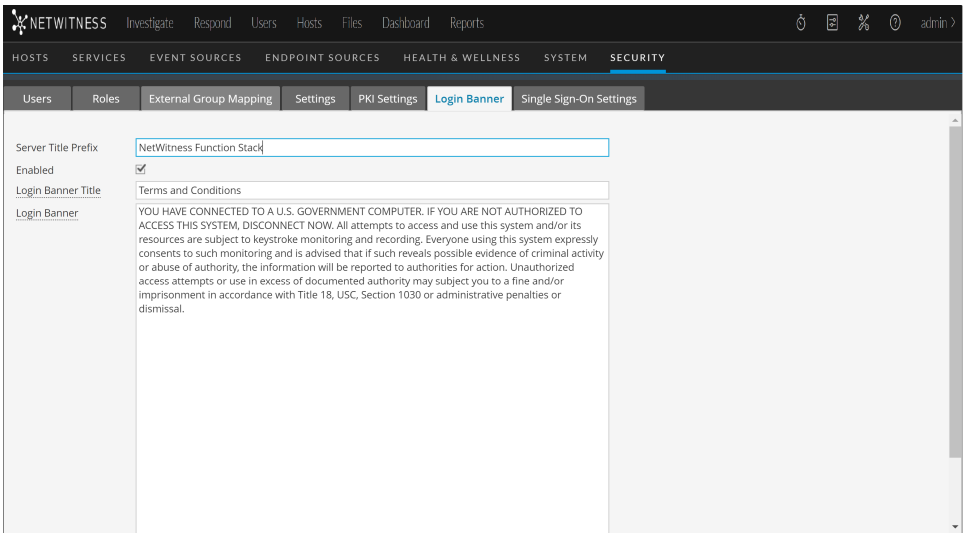
To create and enable a customized login banner:

1. Go to  **(Admin) > Security**.

   The Security view is displayed with the Users tab open.

2. Click the **Login Banner** tab and select the **Enabled** checkbox to toggle between enabling and disabling the banner.

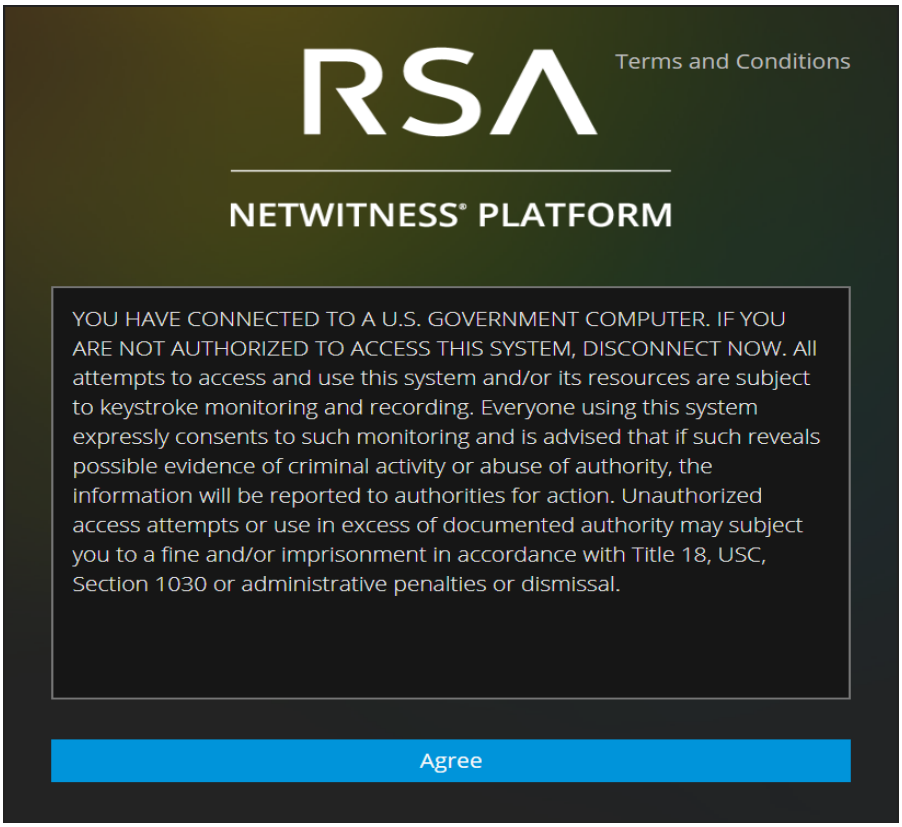
   When Enable is selected, the Login Banner Title and Login Banner fields become active with default content in place.



3. Use the default content or type the custom title and content for your banner and click **Apply**. The banner is enabled and becomes active immediately.

   Note: While both plain text and text with HTML tags are allowed, any suspicious tags will be removed. For example, all links must use HTTPS protocols.

4. To test the banner, log out. The banner is displayed in front of the fields for entering NetWitness credentials.



5. Click **Agree**.

   The banner closes and you can log on.

How Role-Based Access Control Works

In the NetWitness Platform, roles determine what users can do. A role has permissions assigned to it and you must assign a role to each user. The user then has permission to do what the role allows. Role-bsed access control (RBAC) is established when there is a trusted connection between NetWitness Server and a Core service.

Preconfigured Roles

To simplify the process of creating roles and assigning permissions, there are preconfigured roles in NetWitness. You can also add roles customized for your organization.

The following table lists each preconfigured role and the permissions assigned to it. All permissions are assigned to the Administrators role. A subset of permissions is assigned to each of the other roles.

Role	Permission
Administrators	Full system access. The System Administrators persona is granted all permissions by default.
Analysts	Access to meta and session content but not to configurations. The Security Operation Center (SOC) Analysts persona is centered around investigation, ESA Alerting, Reporting, and Respond, but not system configuration.
Reporting_Engine_ Content_Administrators	Access to manage the Live content. Users with the Reporting Engine Content Administrator role can deploy Reporting Engine content (rules, reports, charts, and lists) from Live Content, view and manage permissions to the deployed content in Reporting Engine.
Data_Privacy_Officers	The Data Privacy Officer (DPO) persona is similar to Administrators with additional focus on configuration options that manage obfuscation and viewing of sensitive data within the system (see the *Data Privacy Management Guide*). Users with the DPO role can see which meta keys are flagged for obfuscation, and they also see obfuscated meta keys and values created for the flagged meta keys.
Malware_Analysts	Access to investigations and malware events. The only access granted to the Malware Analysts persona is the Malware Analysis module.
Operators	Access to configurations but not to meta and session content. The System Operators persona is focused on system configuration, but not investigation, ESA, Alerting, Reporting, and Respond.
Respond_Administrator	Access to all Respond permissions. The Respond Administrator persona is focused on system configuration of Respond.
SOC_Managers	Same access as Analysts plus additional permission to handle incidents. The SOC Managers persona is identical to Analysts, but with permissions necessary to configure Respond.

Role	Permission
UEBA_Analysts	Access to the NetWitness UEBA service in the **Investigate** > **Users** view. NetWitness UEBA is an advanced analytics solution for discovering, investigating, and monitoring risky behaviors across all entities in your network environment. **Note:** You do not need to set up specific permissions for this role. You only need to assign this role to a user, and that user will have access to NetWitness UEBA.

Trusted Connections Between Server and Service

In a trusted connection, a service explicitly trusts NetWitness Server to manage and authenticate users. This reduces administration on each service because authenticated users do not have to be defined locally in each Core service.

As the following table shows, you perform all user management tasks on the server.

Task	Location
Add a user	Server
Maintain usernames	Server
Maintain passwords	Server
Authenticate internal NetWitness users	Server
(Optional) Authenticate external users with: - Active Directory - PAM	 Server Server
Install and configure PAM	Server

The benefits of a trusted connection and centralized user management are that:

- You perform all user administration tasks once, on NetWitness Server only.
- You control access to services but do not have to set up and authenticate users on the services.
- Users enter passwords once at NetWitness logon and are authenticated by the server.
- Users, already authenticated by the server, access every Core service in (Admin) > Services without entering a password.

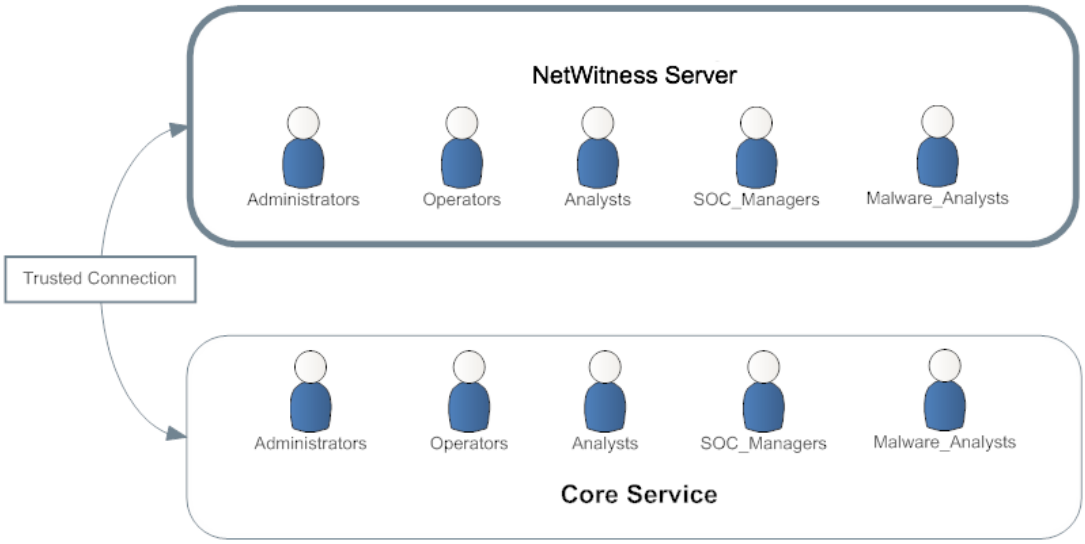
How Trusted Connections Are Established

When you install or upgrade to 11.x, trusted connections are established by default with two settings:

- SSL is enabled.
- The Core service is connected to an encrypted SSL port.

Common Role Names on the Server and Services

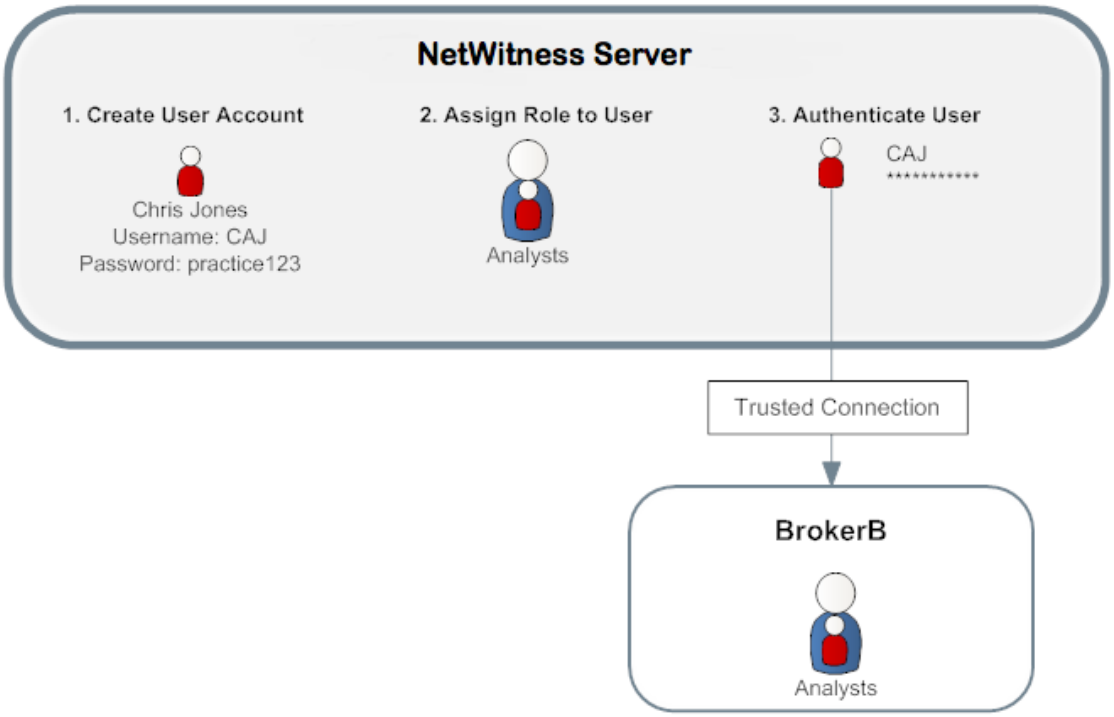
Trusted connections rely on common role names on the server and service. On a fresh installation, NetWitness installs the five preconfigured roles on the server and each Core service.



If you add a custom role, such as JuniorAnalysts, you must add the role to each service, such as ArchiverA and BrokerB. Role names are case-senstive, cannot contain spaces and must be identical. For example, JuniorAnalyst (singular) and JuniorAnalysts (plural) do not meet the requirements for common role names.

End-to-End Workflow for User Setup and Service Access

This workflow shows how role-based access control works when there is a trusted connection between NetWitness Server and the service BrokerB.



1. On NetWitness Server, create an account for a new user:
   Name: Chris Jones
   Username: CAJ
   Password: practice123
2. Determine if you want to assign a preconfigured or custom role to Chris Jones:
   - **Preconfigured role**
     a. Keep or modify the default permissions assigned to the **Analysts role,** which include permissions such as access to the Alerting, Investigation and Malware modules,
     b. Assign the Analysts role to Chris Jones.
   - **Custom role**
     a. Create the custom role, such as JuniorAnalysts.
     b. Assign permissions to the **JuniorAnalysts role**.
     c. Assign the JuniorAnalysts role to Chris Jones.
     d. Add the JuniorAnalysts role to the service, such as BrokerB.
3. The user, Chris Jones, logs on to NetWitness Server:

- Username: CAJ
- Password: practice123

4. The server authenticates Chris.
5. The trusted connection allows the authenticated user, Chris, to access BrokerB without entering another password.

For more detailed descriptions and procedures, see Manage Users with Roles and Permissions.

Role Permissions

InNetWitness, user can access each module, dashlet, and view is restricted based on the assigned permissions. You can locate these role permissions in the Add or Edit Roles dialogs accessible from the  (Admin) > Security > Roles tab.

In the Add or Edit Role dialogs, the tabs in the Permission section represent different areas of NetWitness and show the available permissions for those areas. For example, the Administration tab shows the permissions available in the Admin view.

Note: There is no Configure tab in the Add/Edit Role dialogs that corresponds to the Configure view. To assign permissions in the Configure view, assign permissions to the views contained within the Configure view: Live Content (Live), Incident Rules (Incidents), Respond Notifications (Incidents, Respond-server, Integration server), ESA Rules (Alerting), Subscriptions (Live), and Custom Feeds (Live).

Note: To the left of the Administration tab is a tab marked with an asterisk (*). This tab indicates access to management of backend services only.

The tables that follow show the default permissions assigned to each NetWitness user role:

- Administrators
- Respond Administrators (RAs)
- Reporting Engine Content Administrators (RE CAs)
- Data Privacy Officers (DPOs)
- SOC Managers (SOC Mgrs)
- Operators
- Malware Analysts (MAs)
- Analysts
- UEBA Analysts

Since the Administrators role has all of the permissions by default, it is not included in the tables.

Service Permissions Format for New Services

The service permissions for some new NetWitness services contain three parts in the following format:

<service name>.<resource>.<action>

For example, for the **investigate-server.metrics.read** permission:

- `service name` = **investigate-server**
- `resource` = **metrics**
- `action` = **read**

Users assigned this permission can read any metrics that the investigate-server service exposes.

Admin-server

The following table describes the permissions in the Admin-server tab.

Permission	Description
admin-server.configuration.manage	Permission to modify all service configuration parameters
admin-server.health.read	Permission to view any health notifications that the service exposes
admin-server.logs.manage	Permission to change log-related configuration
admin-server.metrics.read	Permission to view any metrics that the service exposes
admin-server.process.manage	Permission to start and stop the service
admin-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
admin-server.security.read	Permission to view security-related resources

Administration

The following table describes the list of permissions in Administration tab.

Permission	Description
Access Administration Module	Permission to access all the administration modules
Access Health & Wellness	Permission to access the health and wellness module
Apply System Updates	Permission to update the system
Can Opt In to Live Intelligence Sharing	Permission to opt for Live Intelligence sharing
Manage Advanced Settings	Permission to modify the advanced settings
Manage ATD Settings	Permission to modify the ATD settings
Manage Auditing	Permission to modify the auditing

Permission	Description
Manage Email	Permission to change the email settings
Manage Global Auditing	Permission to modify global auditing
Manage Health & Wellness Policy	Permission to update the health & wellness policy
Manage Jobs	Permission to change the job settings
Manage LLS	Permission to modify LLS
Manage Logs	Permission to modify log related configurations
Manage Notifications	Permission to change notification settings
Manage Plugins	Permission to modify the plugins
Manage Predicates	Permission to modify the predicates
Manage Reconstruction	Permission to change the reconstruction
Manage Security	Permission to update the security settings
Manage Services	Permission to start and stop the services
Manage SSL Security	Permission to manage PKI setting
Manage System Settings	Permission to the modify the system settings
Modify ESA Settings	Permission to modify the ESA settings
Modify Event Sources	Permission to modify the ESA sources
Modify Hosts	Permission to modify the hosts
Modify Services	Permission to modify the services
View Event Sources	Permission to view the event sources
View Health & Wellness Policy	Permission to view the health & wellness policy
View Health & Wellness Stats Browser	Permission to view the health and wellness status in the browser
View Hosts	Permission to view the hosts
View Services	Permission to view the services
View Unified Sources	Permission to view the unified sources

The following table lists the permissions in the Administration tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts	UEBA Analysts
Access Administration Module		Yes	Yes	Yes	Yes	Yes	
Access Health & Wellness		Yes	Yes	Yes	Yes	Yes	
Apply System Updates				Yes			
Can Opt In to Live Intelligence Sharing				Yes			
Manage Advanced Settings				Yes			
Manage ATD Settings	Yes	Yes	Yes	Yes			
Manage Auditing		Yes		Yes			
Manage Email				Yes			
Manage Global Auditing		Yes		Yes			
Manage Health & Wellness Policy				Yes			
Manage Jobs		Yes	Yes	Yes			
Manage LLS				Yes			
Manage Logs		Yes		Yes			
Manage Notifications				Yes			
Manage Plugins		Yes	Yes	Yes		Yes	
Manage Predicates				Yes			
Manage Reconstruction				Yes			

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts	UEBA Analysts
Manage Security		Yes		Yes			
Manage Services		Yes		Yes			
Manage SSL Security							
Manage System Settings		Yes	Yes	Yes		Yes	
Modify ESA Settings				Yes			
Modify Event Sources				Yes			
Modify Hosts				Yes			
Modify Services		Yes		Yes			
View Event Sources			Yes	Yes			
View Health & Wellness Policy			Yes	Yes		Yes	
View Health & Wellness Stats Browser		Yes	Yes	Yes		Yes	
View Hosts		Yes		Yes			
View Services		Yes		Yes			
View Unified Sources		Yes	Yes	Yes		Yes	

Alerting

The following table describes the permissions in the Alerting tab.

Permission	Description
Access Alerting Module	Permission to access the alerting module
Manage Rules	Permission to update the rules

Permission	Description
View Alerts	Permission to view the alerts
View Rules	Permission to view the rules

The following table lists the permissions in the Alerting tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Access Alerting Module	Yes	Yes	Yes	Yes		Yes
Manage Rules	Yes	Yes	Yes	Yes		
View Alerts	Yes	Yes	Yes			Yes
View Rules		Yes	Yes	Yes		

Config-server

The following table describes the permissions in the Config-server tab. The Administrators role has all of the permissions and is the only role granted permissions by default.

Permission	Description
config-server.*	All permissions (everything below)
config-server.configuration.manage	Permission to modify all service configuration parameters
config-server.health.read	Permission to view any health notifications that the service exposes
config-server.logs.manage	Permission to change log-related configuration
config-server.metrics.read	Permission to view any metrics that the service exposes
config-server.process.manage	Permission to start and stop the service
config-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
config-server.security.read	Permission to view security-related resources

Content-server

The following table describes the permissions in the Content-server tab.

Permission	Description
content-server.*	All permissions (everything below)
content-server.collection.read	Permission to read selective collection content
content-server.configuration.manage	Permission to modify all service configuration parameters
content-server.health.read	Permission to view any health notifications that the service exposes
content-server.logparser.manage	Permission to manage log parser configurations
content-server.logparser.read	Permission to view log parser configurations
content-server.logs.manage	Permission to change log-related configuration
content-server.metrics.read	Permission to view any metrics that the service exposes
content-server.policy.read	Permission to read policies
content-server.process.manage	Permission to start and stop the service
content-server.rule.manage	Permission to manage content rules
content-server.rule.read	Permission to view content rules
content-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
content-server.security.read	Permission to view security-related resources

The following table lists the permissions in the Content-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrator role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
content-server.*		Yes		Yes		
content-server.collection.read						
content-server.configuration.manage						
content-server.health.read						
content-server.logparser.manage						
content-server.logparser.read			Yes			Yes
content-server.logs.manage						

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
content-server.metrics.read						
content-server.policy.read						
content-server.process.manage						
content-server.rule.manage						
content-server.rule.read						
content-server.security.manage						
content-server.security.read						

Contexthub-server

The following table describes the permissions in the Contexthub-server tab.

Permission	Description
contexthub-server.*	All permissions (everything below)
contexthub-server.configuration.manage	Permission to modify all service configuration parameters
contexthub-server.connection.manage	Permission to modify all connection settings
contexthub-server.connection.read	Permission to view all connection settings
contexthub-server.connectiontypes.read	Permission to view all configured connection types
contexthub-server.datasource.manage	Permission to modify data source settings
contexthub-server.datasource.read	Permission to view data source settings
contexthub-server.health.read	Permission to view any health notifications that the service exposes
contexthub-server.listentries.manage	Permission to modify list entries
contexthub-server.logs.manage	Permission to change log-related configuration
contexthub-server.metrics.read	Permission to view any metrics that the service exposes
contexthub-server.process.manage	Permission to start and stop the service

Permission	Description
contexthub-server.query.read	Permission to view queries
contexthub-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
contexthub-server.security.read	Permission to view security-related resources
contexthub-server.stix.read	Permission to view stix settings
contexthub-server.taxiidatasource.manage	Permission to modify settings for the taxii data source
contexthub-server.taxiidatasource.read	Permission to view settings for the taxii data source

The following table lists the permissions in the Contexthub-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrator role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
contexthub-server.*		Yes				
contexthub-server.configuration.manage						
contexthub-server.connection.manage						
contexthub-server.connection.read	Yes		Yes		Yes	Yes
contexthub-server.connectiontypes.read			Yes			
contexthub-server.datasource.manage	Yes		Yes		Yes	Yes
contexthub-server.datasource.read	Yes		Yes		Yes	Yes
contexthub-server.health.read						
contexthub-server.listentries.manage	Yes		Yes		Yes	Yes
contexthub-server.logs.manage						
contexthub-server.metrics.read						

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
contexthub-server.process.manage						
contexthub-server.query.read	Yes		Yes		Yes	Yes
contexthub-server.security.manage						
contexthub-server.security.read						
contexthub-server.stix.read			Yes		Yes	Yes
contexthub-server.taxiidatasource.manage			Yes		Yes	Yes
contexthub-server.taxiidatasource.read			Yes		Yes	Yes

Correlation-server

The following table describes the permissions in the Correlation-server tab. These permissions pertain to ESA Correlation.

Permission	Description
correlation-server.*	All permissions (everything below)
correlation-server.configuration.manage	Permission to modify all service configuration parameters
correlation-server.endpoint.manage	Permission to modify all endpoint configuration parameters
correlation-server.endpoint.read	Permission to view all endpoint configuration parameters
correlation-server.engine.manage	Permission to modify all engine configuration parameters
correlation-server.engine.read	Permission to view all engine configuration parameters
correlation-server.esperrule.manage	Permission to modify all esperrule configuration parameters
correlation-server.esperrule.read	Permission to view all esperrule configuration parameters
correlation-server.health.read	Permission to view any health notifications that the service exposes
correlation-server.keyvaluerule.manage	Permission to modify all keyvaluerule configuration parameters
correlation-server.keyvaluerule.read	Permission to view all keyvaluerule configuration parameters

Permission	Description
correlation-server.logs.manage	Permission to change log-related configuration
correlation-server.metrics.read	Permission to view any metrics that the service exposes
correlation-server.module.manage	Permission to modify each module
correlation-server.module.read	Permission to view each module
correlation-server.process.manage	Permission to start and stop the service
correlation-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
correlation-server.security.read	Permission to view security-related resources
correlation-server.stream.manage	Permission to edit stream configuration settings
correlation-server.stream.read	Permission to view stream configuration settings
correlation-server.telemetry.read	Permission to view telemetry configuration settings

The following table lists the permissions in the Correlation-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrator role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
correlation-server.*		Yes				
correlation-server.configuration.manage						
correlation-server.endpoint.manage						
correlation-server.endpoint.read						
correlation-server.engine.manage	Yes		Yes	Yes		
correlation-server.engine.read	Yes		Yes	Yes		
correlation-server.esperrule.manage						
correlation-server.esperrule.read						
correlation-server.health.read						

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
correlation-server.keyvaluerule.manage						
correlation-server.keyvaluerule.read						
correlation-server.logs.manage						
correlation-server.metrics.read						
correlation-server.module.manage	Yes		Yes	Yes		
correlation-server.module.read	Yes		Yes	Yes		
correlation-server.process.manage						
correlation-server.security.manage						
correlation-server.security.read						
correlation-server.stream.manage	Yes		Yes	Yes		
correlation-server.stream.read	Yes		Yes	Yes		
correlation-server.telemetry.read						

Dashboard

The following table describes the permissions in the Dashboard tab.

Permission	Description
Dashlet Access - Admin Device List Dashlet	Permission to access Admin Device List Dashlet
Dashlet Access - Admin Device Monitor Dashlet	Permission to access Admin Device Monitor Dashlet
Dashlet Access - Admin News Dashlet	Permission to access Admin News Dashlet
Dashlet Access - Alert Variance Dashlet	Permission to access Alert Variance Dashlet

Permission	Description
Dashlet Access - Alerting Recent Alerts Dashlet	Permission to access Alerting Recent Alerts Dashlet
Dashlet Access - Investigation Jobs Dashlet	Permission to access Investigation Jobs Dashlet
Dashlet Access - Investigation Top Values Dashlet	Permission to access Investigation Top Values Dashlet
Dashlet Access - Live Featured Resources Dashlet	Permission to access Live Featured Resources Dashlet
Dashlet Access - Live New Resources Dashlet	Permission to access Live New Resources Dashlet
Dashlet Access - Live Subscriptions Dashlet	Permission to access Live Subscriptions Dashlet
Dashlet Access - Live Updated Resources Dashlet	Permission to access Live Updated Resources Dashlet
Dashlet Access - Malware Jobs Dashlet	Permission to access Malware Jobs Dashlet
Dashlet Access - Reporting Recent Report Dashlet	Permission to access Reporting Recent Report Dashlet
Dashlet Access - Reporting Charts Dashlet	Permission to access Reporting Charts Dashlet
Dashlet Access - Top Alerts Dashlet	Permission to access Top Alerts Dashlet
Dashlet Access - Unified First Watch Dashlet	Permission to access Unified First Watch Dashlet
Dashlet Access - Unified Shortcuts Dashlet	Permission to access Unified Shortcuts Dashlet

The following table lists the permissions in the Dashboard tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RA	DPOs	SOC Mgrs	Operators	MAs	Analysts
Dashlet Access - Admin Device List Dashlet	Yes	Yes	Yes	Yes		Yes
Dashlet Access - Admin Device Monitor Dashlet		Yes				
Dashlet Access - Admin News Dashlet	Yes	Yes	Yes	Yes		Yes
Dashlet Access - Alert Variance Dashlet	Yes	Yes	Yes			Yes
Dashlet Access - Alerting Recent Alerts Dashlet	Yes	Yes	Yes			Yes

Permission	RA	DPOs	SOC Mgrs	Operators	MAs	Analysts
Dashlet Access - Investigation Jobs Dashlet	Yes	Yes	Yes			Yes
Dashlet Access - Investigation Top Values Dashlet	Yes	Yes	Yes			Yes
Dashlet Access - Live Featured Resources Dashlet	Yes	Yes	Yes	Yes		Yes
Dashlet Access - Live New Resources Dashlet	Yes	Yes	Yes	Yes		Yes
Dashlet Access - Live Subscriptions Dashlet	Yes	Yes	Yes	Yes		Yes
Dashlet Access - Live Updated Resources Dashlet	Yes	Yes	Yes	Yes		Yes
Dashlet Access - Malware Jobs Dashlet	Yes	Yes	Yes			Yes
Dashlet Access - Reporting Recent Report Dashlet	Yes	Yes	Yes			Yes
Dashlet Access - Reporting Charts Dashlet	Yes	Yes	Yes			Yes
Dashlet Access - Top Alerts Dashlet	Yes	Yes	Yes			Yes
Dashlet Access - Unified First Watch Dashlet	Yes	Yes	Yes	Yes		Yes
Dashlet Access - Unified Shortcuts Dashlet	Yes	Yes	Yes	Yes		Yes

Endpoint-broker-server

The following table describes the permissions in the Endpoint Broker server tab.

Permission	Description
endpoint-broker-server*	All permissions (everything below)
endpoint-broker-server.agent.manage	Permission to manage the agent, that is start or stop scan, downloading file from host, delete agent data from the Endpoint Log Hybrid and so on.

Permission	Description
endpoint-broker-server.agent.read	Permission to view the endpoint data received from the agent such as host, file, certificate, events and so on.
endpoint-broker-server.configuration.manage	Permission to modify all endpoint broker configuration parameters
endpoint-broker-server.health.read	Permission to view any health notifications that the service exposes
endpoint-broker-server.logs.manage	Permission to change log-related configuration
endpoint-broker-server.metrics.read	Permission to view any metrics that the service exposes
endpoint-broker-server.policy.read	Permission to view existing policy details
endpoint-broker-server.process.manage	Permission to start and stop the service
endpoint-broker-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
endpoint-broker-server.security.read	Permission to view security-related resources

The following table lists the permissions in the Endpoint-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrator role has all of the permissions by default and is not listed.

Permission	RA	DPOs	SOC Mgrs	Operators	MAs	Analysts
endpoint-broker-server*						
endpoint-broker-server.agent.manage				Yes		Yes
endpoint-broker-server.agent.read				Yes		Yes
endpoint-broker-server.configuration.manage						
endpoint-broker-server.health.read						
endpoint-broker-server.logs.manage						
endpoint-broker-server.metrics.read						
endpoint-broker-server.policy.read						Yes

Permission	RA	DPOs	SOC Mgrs	Operators	MAs	Analysts
endpoint-broker-server.process.manage						
endpoint-broker-server.security.manage						
endpoint-broker-server.security.read						

Endpoint-server

The following table describes the permissions in the Endpoint-server tab.

Permission	Description
endpoint-server*	All permissions (everything below)
endpoint-server.agent.manage	Permission to generate and download the agent packager. Permission to manage the agent, that is start or stop scan, downloading files, master file table (MFT), memory dumps from host, isolate host from network, delete agent data from the Endpoint Log Hybrid and so on Note: Analyze file, Scan With OPSWAT, Save Local Copy have been moved to endpoint-server.file.analyze (from version 11.7)
endpoint-server.agent.read	Permission to view the agent packager configuration Permission to view the endpoint data received from the agent such as host, file, certificate, events, and so on
endpoint-server.agentupdate.manage	Permission to upgrade agent and uninstall agent
endpoint-server.ca.manage	Permission to generate and download the agent packager Permission to upgrade agent
endpoint-server.ca.read	Permission to generate and download the agent packager
endpoint-server.configuration.manage	Permission to modify all endpoint configuration parameters
endpoint-server.file.analyze	Permission to analyze file, save local copy, and initiate OPSWAT scans
endpoint-server.filter.manage	Permission to save, modify, and delete filters
endpoint-server.filter.read	Permission to view filters

Permission	Description
endpoint-server.health.read	Permission to view any health notifications that the service exposes
endpoint-server.logs.manage	Permission to change log-related configuration
endpoint-server.metrics.read	Permission to view any metrics that the service exposes
endpoint-server.policy.read	Permission to view existing policy details
endpoint-server.process.manage	Permission to start and stop the service
endpoint-server.relay.manage	Permission to modify Relay Server Configuration
endpoint-server.relay.read	Permissions to view Relay Server details
endpoint-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
endpoint-server.security.read	Permission to view security-related resources
endpoint-server.tag.manage	Permission to manage Tags (Create and Delete)

The following table lists the permissions in the Endpoint-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrator role has all of the permissions by default and is not listed.

Permission	RA	DPOs	SOC Mgrs	Operators	MAs	Analysts
endpoint-server*						
endpoint-server.agent.manage				Yes		Yes
endpoint-server.agent.read				Yes		Yes
endpoint-server.agentupdate.manage						
endpoint-server.ca.manage				Yes		
endpoint-server.ca.read				Yes		
endpoint-server.configuration.manage						
endpoint-server.filter.manage						Yes
endpoint-server.filter.read						Yes
endpoint-server.health.read						
endpoint-server.logs.manage						
endpoint-server.metrics.read						

Permission	RA	DPOs	SOC Mgrs	Operators	MAs	Analysts
endpoint-server.policy.read						Yes
endpoint-server.process.manage						
endpoint-server.rar.manage						
endpoint-server.rar.read						
endpoint-server.relay.manage				Yes		
endpoint-server.relay.read				Yes		
endpoint-server.security.manage						
endpoint-server.security.read						

Incidents

The following table describes the permissions in the Incidents tab.

Permission	Description
Access Incident Module	Permission to access the Incident module
Configure Incident Management Integration	Permission to configure incident management integration
Delete Alerts and incidents	Permission o delete alerts and incidents
Manage Alert Handling Rules	Permission to modify the alert handling rules
View and Manage Incidents	Permission to modify the incidents

The following table lists the permissions in the Incidents tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Access Incident Module	Yes	Yes	Yes		Yes	Yes
Configure Incident Management Integration	Yes	Yes	Yes			

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Delete Alerts and incidents	Yes	Yes				
Manage Alert Handling Rules	Yes	Yes	Yes			
View and Manage Incidents	Yes	Yes	Yes		Yes	Yes

Integration-server

(The Integration-server permissions are available in NetWitness version 11.1 and later.)

The following table describes the permissions in the Integration-server tab.

Permission	Description
integration-server.*	All permissions (everything below)
integration-server.api.access	Permission to authorize external requests from 3rd party applications
integration-server.configuration.manage	Permission to view and modify all service integration configuration parameters
integration-server.health.read	Permission to read any health notifications that the service exposes
integration-server.logs.manage	Permission to change log-related integration configurations
integration-server.metrics.read	Permission to read any metrics that the service exposes
integration-server.notification.manage	Permission to change global notification configurations (for example, SMTP server)
integration-server.notification.read	Permission to read global notification configurations (for example, SMTP server)
integration-server.notification.send	Permission to send notifications (for example, Email)
integration-server.process.manage	Permission to start and stop the service
integration-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
integration-server.security.read	Permission to read security-related resources
integration-server.template.manage	Permission to change notification template
integration-server.template.read	Permission to read notification template

The following table lists the permissions in the Integration-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrator role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
integration-server.*		Yes				
integration-server.api.access						
integration-server.configuration.manage						
integration-server.health.read						
integration-server.logs.manage						
integration-server.metrics.read						
integration-server.notification.manage	Yes		Yes	Yes		
integration-server.notification.read	Yes		Yes	Yes		
integration-server.notification.send	Yes		Yes	Yes		
integration-server.process.manage						
integration-server.security.manage						
integration-server.security.read						
integration-server.template.manage	Yes		Yes	Yes		
integration-server.template.read	Yes		Yes	Yes		

Investigate

The following table describes the permissions in the Investigate tab.

Permission	Description
Access Investigation Module	Permission to access investigation module
Context Lookup	Permission to access context lookup
Create Incidents from Investigation	Permission to create incidents from investigation
Manage List from Investigation	Permission to modify the list of investigation
Navigate Events	Permission to navigate the events
Navigate Values	Permission to navigate the values

The following table lists the permissions in the Investigate tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Access Investigation Module	Yes	Yes	Yes		Yes	Yes
Context Lookup	Yes		Yes		Yes	Yes
Create Incidents from Investigation	Yes		Yes		Yes	Yes
Manage List from Investigation	Yes		Yes		Yes	Yes
Navigate Events	Yes	Yes	Yes		Yes	Yes
Navigate Values	Yes	Yes	Yes		Yes	Yes

Investigate-server

The following table describes the permissions in the Investigate-server tab.

Permission	Description
investigate-server.*	All permissions (everything below) for 11.4 and above Events view, and 11.3 and earlier Event Analysis view.
investigate-server.column group.read	Permission to access column groups

Permission	Description
investigate-server.configuration.manage	Permission to change any configuration properties for the service
investigate-server.content.export	Permission to export content from the service
investigate-server-content.manage	Permission to clear all per service or per user reconstruction cache
investigate-server.content.reconstruct	Permission to view the summary view, the packet, packet map, text, log, and file reconstructions, as well as the packet count
investigate-server.event.filter	Permission to view the Filter Events panel in the Events view
investigate-server.event.read	Permission to view events that the service exposes
investigate-server.health.read	Permission to view any health notifications that the service exposes
investigate-server.incident.manage	Create or update an incident in Respond
investigate-server.logs.manage	Permission to change log-related configuration
investigate-server.metagroup.manage	Permission to manage meta groups
investigate-server.metagroup.read	Permission to view and use meta groups
investigate-server.metrics.read	Permission to view any metrics that the service exposes
investigate-server.predicate.manage	Permission to edit or remove one or more predicates
investigate-server.predicate.read	Permission to filter events in the Navigate view, Legacy EventsEvents view, and Events view. Note: This permission is required with investigate-server.event.read permission to provide access to the and Events view.
investigate-server.process.manage	Permission to start and stop the service
investigate-server.profile.read	Permission to access profiles.
investigate-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
investigate-server.security.read	Permission to view security-related resources

The following table lists the permissions in the Investigate-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed. The UEBA Analysts and Content Administrators roles have none of these permissions by default and are not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
investigate-server.*	Yes	Yes				
investigate-server.columngroup.read			Yes		Yes	Yes
investigate-server.configuration.manage						
investigate-server.content.export			Yes		Yes	Yes
investigate-server.content.manage						
investigate-server.content.reconstruct			Yes		Yes	Yes
investigate-server.event.filter	Yes	Yes	Yes		Yes	Yes
investigate-server.event.read			Yes		Yes	Yes
investigate-server.health.read						
investigate-server.incident.manage						Yes
investigate-server.logs.manage						
investigate-server.metagroup.manage						
investigate-server.metagroup.read			Yes		Yes	Yes
investigate-server.metrics.read						
investigate-server.predicate.manage						
investigate-server.predicate.read			Yes		Yes	Yes
investigate-server.process.manage						
investigate-server.profile.read			Yes		Yes	Yes
investigate-server.security.manage						

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
investigate-server.security.read						

License-server

The following table describes the permissions in the License-server tab. The Administrator and Operator have all of the permissions and are the only roles granted permissions by default.

Permission	Description
license-server.*	All permissions (everything below)
license-server.configuration.manage	Permission to modify all service configuration parameters
license-server.health.read	Permission to view any health notifications that the service exposes
license-server.license.manage	Permission to manage license related configurations
license-server.license.read	Permission to view license related configurations
license-server.logs.manage	Permission to change log-related configuration
license-server.metrics.read	Permission to view any metrics that the service exposes
license-server.process.manage	Permission to start and stop the service
license-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
license-server.security.read	Permission to view security-related resources

Live

The following table describes the permissions in the Live tab.

Permission	Permission
Live	
Access Live Module	Permission to access live module
Manage Live System Settings	Permission to modify the live system settings
Resources	

Permission	Permission
Deploy Live Resources	Permission to deploy live resources
Manage Live Feeds	Permission to modify live feeds
Manage Live Resources	Permission to modify live resources
Search Live Resources	Permission to search live resources
View Live Resource Details	Permission to view live resource details

The following table lists the permissions in the Live tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Live						
Access Live Module		Yes	Yes	Yes		Yes
Manage Live System Settings				Yes		
Resources						
Deploy Live Resources		Yes		Yes		
Manage Live Feeds		Yes		Yes		
Manage Live Resources		Yes		Yes		
Search Live Resources		Yes	Yes	Yes		Yes
View Live Resource Details		Yes	Yes	Yes		

Malware

The following table describes the permissions in the Malware tab.

Permission	Operators
Download Malware File(s)	Permission to donwnload the malware files for investigation
Initiate Malware Analysis Scan	Permission to start the malware analysis scan
View Malware Analysis Events	Permission to view the malware analysis events

The following table lists the permissions in the Malware tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Download Malware File (s)		Yes	Yes		Yes	Yes
Initiate Malware Analysis Scan		Yes	Yes		Yes	Yes
View Malware Analysis Events		Yes	Yes		Yes	Yes

Metrics-server

The following table describes the permissions in the Metrics-server tab. The Administrators role have all of the permissions and are the only roles granted permissions by default.

Permission	Description
metrics-server .*	All permissions (everything below)
metrics-server.configuration.manage	Permission to modify all service configuration parameters
metrics-server-content.manage	Permission to modify configuration parameters in the service
metrics-server-content.read	Permission to view configuration parameters of the service
metrics-server.health.read	Permission to view any health notifications that the service exposes
metrics-server.logs.manage	Permission to change log-related configuration
metrics-server.metric.manage	Permission to modify all the configuration parameters
metrics-server.metric.read	Permission to view configuration of New Health and Wellness
metrics-server.metrics.read	Permission to view any metrics that the service exposes
metrics-server.process.manage	Permission to start and stop the service
metrics-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
metrics-server.security.read	Permission to view security-related resources

The following table lists the permissions in the Metrics-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RA	DPOs	SOC Mgrs	Operator	MAs	Analysts	UEBA Analysts
metrics-server .*							
metrics-server.configuration.manage							
metrics-server-content.manage							
metrics-server-content.read							
metrics-server.health.read							
metrics-server.logs.manage							
metrics-server.metric.manage							
metrics-server.metric.read		Yes	Yes	Yes	Yes	Yes	
metrics-server.metrics.read							
metrics-server.process.manage							
metrics-server.security.manage							
metrics-server.security.read							

Orchestration-server

The following table describes the permissions in the Orchestration-server tab. The Administrators, Operators, and Data Privacy Officers roles have all of the permissions and are the only roles granted permissions by default.

Permission	Description
orchestration-server.*	All permissions (everything below)
orchestration-server.configuration.manage	Permission to modify all service configuration parameters
orchestration-server.file.read	Permission to view files

Permission	Description
orchestration-server.health.read	Permission to view any health notifications that the service exposes
orchestration-server.logs.manage	Permission to change log-related configuration
orchestration-server.metrics.read	Permission to view any metrics that the service exposes
orchestration-server.process.manage	Permission to start and stop the service
orchestration-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
orchestration-server.security.read	Permission to view security-related resources

Reports

The following table describes the permissions in the Reports tab.

Permission	Description
Alert	
Define RE Alert	Permission to define the RE alerts
Export RE Alert Definition	Permission to export the RE alert definistions
Manage RE Alerts	Permission to to modify the RE alerts
View RE Alerts	Permission to view the RE alerts
View Scheduled RE Alerts	Permission to view the scheduled RE alerts
Chart	
Define Chart	Permission to define the charts
Delete Chart	Permission to delete the charts
Export Chart Definition	Permission to export the chart definistions
Manage Charts	Permission to modify the charts
View Charts	Permission to view the charts
List	
Define Lists	Permission to define the lists
Delete List	Permission to delete the lists

Permission	Description
Export List	Permission to export the lists
Manage Lists	Permission to modify the lists
Report	
Define Report	Permission to define the reports
Delete Report	Permission to delete the reports
Export Report	Permission to export the reports
Manage Reports	Permission to modify the reports
View Reports	Permission to view the reports
Reports	
Access Configure	Permission to access Configure module
Access Reporter Module	Permission to access Reporter module
Access Reporter search	Permission to access Reporter search
Access View	Permission to access Reports view
Rule	
Add RE Alert Definition from Rule	Permission to add RE alert definition from the rules
Define Rule	Permission to define the rules
Delete Rule	Permission to delete the rules
Export Rule	Permission to export the rules
Manage Rules	Permission to modify the rules
View Rule Usage	Permission to view the rules usage
Schedules	
Define Schedule	Permission to define the schedules
Delete Schedule	Permission to delete the schedules
View Schedules	Permission to view the schedules
Warehouse Analytics	
Define Jobs	Permission to define the warehouse analytics jobs

Permission	Description
Delete Jobs	Permission to delete the warehouse analytics jobs
Manage Jobs	Permission to modify the warehouse analytics jobs
View Jobs	Permission to view the warehouse analytics jobs

The following table lists the permissions in the Reports tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all of the permissions by default and is not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Alert						
Define RE Alert		Yes	Yes			Yes
Export RE Alert Definition		Yes	Yes			Yes
Manage RE Alerts		Yes	Yes			Yes
View RE Alerts	Yes	Yes	Yes			Yes
View Scheduled RE Alerts		Yes	Yes			Yes
Chart						
Define Chart		Yes	Yes			Yes
Delete Chart		Yes	Yes			Yes
Export Chart Definition		Yes	Yes			Yes
Manage Charts		Yes	Yes			Yes
View Charts		Yes	Yes			Yes
List						
Define Lists		Yes	Yes			Yes
Delete List		Yes	Yes			Yes
Export List		Yes	Yes			Yes
Manage Lists		Yes	Yes			Yes
Report						
Define Report		Yes	Yes			Yes
Delete Report		Yes	Yes			Yes
Export Report		Yes	Yes			Yes

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
Manage Reports		Yes	Yes			Yes
View Reports		Yes	Yes			Yes
Reports						
Access Configure		Yes	Yes			Yes
Access Reporter Module		Yes	Yes			Yes
Access Reporter search		Yes	Yes			Yes
Access View		Yes	Yes			Yes
Rule						
Add RE Alert Definition from Rule		Yes	Yes			Yes
Define Rule		Yes	Yes			Yes
Delete Rule		Yes	Yes			Yes
Export Rule		Yes	Yes			Yes
Manage Rules		Yes	Yes			Yes
View Rule Usage		Yes	Yes			Yes
Schedules						
Define Schedule		Yes	Yes			Yes
Delete Schedule		Yes	Yes			Yes
View Schedules		Yes	Yes			Yes
Warehouse Analytics						
Define Jobs		Yes	Yes			Yes
Delete Jobs		Yes	Yes			Yes
Manage Jobs		Yes	Yes			Yes
View Jobs		Yes	Yes			Yes

Respond-server

The following table describes the permissions in the Respond-server tab.

Note: For viewing and managing the Risk Score feature, users who have installed NetWitness Platform 11.3 or upgraded from NetWitness 10.6.x to 11.3, risk score permissions will be already present for Analysts. For users updating from NetWitness 11.x to NetWitness Platform 11.3, the Administrator has to provide Analysts' permissions to manage and view risk score.

Permission	Description
respond-server.*	All permissions (everything below)
respond-server.alert.delete	Permission to delete alerts
respond-server.alert.manage	Permission to create, update, or delete alerts and alert filters
respond-server.alert.read	Permission to view alerts and alert filters
respond-server.alertrule.manage	Permission to create, update, or delete alert aggregation rules
respond-server.alertrule.read	Permission to view alert aggregation rules
respond-server.configuration.manage	Permission to change any configuration properties for the service
respond-server.health.read	Permission to view any health notifications that the service exposes
respond-server.incident.delete	Permission to delete incidents
respond-server.incident.manage	Permission to create, update, or delete incidents and incident filters including permission to view the Create Incident and Add to Incident options in the Investigate > Events view
respond-server.incident.read	Permission to view incidents and incident filters
respond-server.journal.manage	Permission to create, update, or delete journal entries for an incident
respond-server.journal.read	Permission to view journal entries for an incident
respond-server.logs.manage	Permission to change log-related configuration
respond-server.metrics.read	Permission to view any metrics that the service exposes
respond-server.notification.manage	(This permission is available in NetWitness version 11.1 and later.) Permission to configure incident email notification settings such as the selected email server, SOC Managers, and who will be sent the notifications (Assignee and SOC Managers)
respond-server.notification.read	(This permission is available in NetWitness version 11.1 and later.) Permission to view incident email notification settings
respond-server.process.manage	Permission to start and stop the service
respond-server.remediation.manage	Permission to create, update, or delete remediation tasks

Permission	Description
respond-server.remediation.read	Permission to view remediation tasks
respond-server.risk.manage	Permission to manage risk score
respond-server.risk.read	Permission to view risk score
respond-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
respond-server.security.read	Permission to view security-related resources

The following table lists the permissions in the Respond-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrator has all of the permissions by default and are not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
respond-server.*	Yes	Yes				
respond-server.alert.delete						
respond-server.alert.manage			Yes		Yes	Yes
respond-server.alert.read			Yes		Yes	Yes
respond-server.alertrule.manage			Yes			
respond-server.alertrule.read			Yes			
respond-server.configuration.manage						
respond-server.health.read						
respond-server.incident.delete						
respond-server.incident.manage			Yes		Yes	Yes
respond-server.incident.read			Yes		Yes	Yes
respond-server.journal.manage			Yes		Yes	Yes
respond-server.journal.read			Yes		Yes	Yes
respond-server.logs.manage						
respond-server.metrics.read						

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
respond-server.notification.manage			Yes			
respond-server.notification.read			Yes			
respond-server.process.manage						
respond-server.remediation.manage			Yes		Yes	Yes
respond-server.remediation.read			Yes		Yes	Yes
respond-server.risk.manage						Yes
respond-server.risk.read						Yes
respond-server.security.manage						
respond-server.security.read						

Incident Email Notification Settings Permissions

Note: Incident email notification setting permissions are available in NetWitness version 11.1 and later.
If you are updating from NetWitness version 11.0 to 11.1 or later, you will need to add additional permissions to your existing built-in NetWitness user roles. For all upgrades to 11.1 or later, you will need to add additional permissions to custom roles.

The following permissions are required for Respond Administrators, Data Privacy Officers, and SOC Managers to access Incident Email Notification Settings [(Configure) > Incident Notifications].

Incidents tab:

- Configure Incident Management Integration

Respond-server tab:

- respond-server.notification.manage
- respond-server.notification.read

Integration-server tab:

- integration-server.notification.read
- integration-server.notification.manage

Respond Event Analysis Permissions

Note: The Event Analysis panel in the Respond view is available in NetWitness version 11.2 and later.

The Events panel in the Respond view, formerly known as the Event Analysis panel, shows the Events view from Investigate for specific indicator events. The following permissions are required to view the Events panel in the Respond view. These permissions are provided by default for users with the Analysts role.

Investigate-server tab:

- investigate-server.event.read
- investigate-server.content.reconstruct
- investigate-server.content.export

Administration tab:

- Access Administration Module

Respond Saved Filter Permissions

Note: Saved filters for the incidents and alerts lists in Respond are available in NetWitness version 11.5 and later.

The following permissions are required for the incidents and alerts filters (Respond > Incidents and Respond > Alerts). The Analysts role has the required Respond filter permissions by default.

Respond-server tab:

- respond-server.incident.manage
- respond-server.incident.read
- respond-server.alert.manage
- respond-server.alert.read

Security-server

The following table describes the permissions in the Security-server tab. The Administrators, Operators, and Data Privacy Officers roles have all of the permissions and are the only roles granted permissions by default.

Permission	Description
security-server.*	All permissions (everything below)
security-server.account.manage	Permission to view, create, modify, or remove NetWitness local accounts
security-server.account.read	Permission to view NetWitness local accounts

Permission	Description
security-server.ca.manage	Permission to manage NetWitness deployment PKI parameters (for example, sign certificates, and so on)
security-server.ca.read	Permission to view NetWitness deployment PKI parameters
security-server.configuration.manage	Permission to modify all service configuration parameters
security-server.connection.manage	Permission to modify all connection configuration parameters
security-server.health.read	Permission to view any health notifications that the service exposes
security-server.logs.manage	Permission to change log-related configuration
security-server.metrics.read	Permission to view any metrics that the service exposes
security-server.permission.manage	Permission to create or remove NetWitness permissions
security-server.pki.manage	Permission to modify all pki configuration parameters
security-server.process.manage	Permission to start and stop the service
security-server.role.manage	Permission to create, modify, or remove NetWitness roles (for example, add role permissions)
security-server.role.read	Permission to view NetWitness role definitions
security-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
security-server.security.read	Permission to view security-related resources
security-server.test.manage	Permission to modify all test configuration parameters
security-server.user.manage	Permission to view, create, modify, or remove NetWitness user profiles
security-server.user.read	Permission to view NetWitness user profile details (for example, roles, login times, and so on)

Source-server

The following table describes the permissions in the Source-server tab.

Permission	Description
source-server.*	All permissions (everything below)
source-server.configuration.manage	Permission to change any configuration properties for the service

Permission	Description
source-server.group.manage	Permission to create and manage USM groups
source-server.group.manage.nopolicy	Permission to manage nopolicy
source-server.group.read	Permission to view USM groups
source-server.grouppolicy.read	Permission to view the canonical groups and policies
source-server.health.read	Permission to view any health notifications that the service exposes
source-server.logs.manage	Permission to change log-related configuration
source-server.metrics.read	Permission to view any metrics that the service exposes
source-server.policy.manage	Permission to create and manage USM policies
source-server.policy.read	Permission to view USM policies
source-server.process.manage	Permission to start and stop the service
source-server.security.manage	Permission to edit security-related resources (passwords, keys, and so on)
source-server.security.read	Permission to view security-related resources

The following table lists the permissions in the Source-server tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators and Operators roles have all of the permissions by default and are not listed.

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
source-server.*						
source-server.configuration.manage						
source-server.group.manage						
source-server.group.manage.nopolicy						
source-server.group.read		Yes	Yes			Yes
source-server.grouppolicy.read						
source-server.health.read						
source-server.logs.manage						
source-server.metrics.read						

Permission	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
source-server.policy.manage						
source-server.policy.read		Yes	Yes			Yes
source-server.process.manage						
source-server.security.manage						
source-server.security.read						

Springboard

The following table describes the permissions in Springboard tab.

Permission	Description
springboard.*	All Permissions (everything below)
springboard.manage	Permission to manage the Springboard, that is view, add, delete, and rearrange panels, and also restore system default settings.
springboard.read	Permission to view Springboard.

The following table lists the permissions in the Springboard tab assigned to each role. A blank field indicates that the role does not have the permission. The Administrators role has all the permissions by default and is not listed.

Permissions	RAs	DPOs	SOC Mgrs	Operators	MAs	Analysts
springboard.*						
springboard.manage						
springboard.read	Yes		Yes		Yes	Yes

Manage Users with Roles and Permissions

A set of end-to-end procedures for managing users in NetWitness is described in this section. These steps explain how to add a user in NetWitness and then how to control what the user can do.

Manage Users Workflow

This figure shows the high-level workflow for managing the users in NetWitness.



First, review the preconfigured NetWitness roles such as Administrators, Respond Administrators, Data Privacy Officers, SOC managers, Operators, Malware analysts, Analysts, and UEBA analysts. Determine if you need to make any adjustments to these roles for your environment. You can add new custom roles and assign permissions to each role. After you define the roles, you can verify the query and session attributes that are set for each role. Then you can set up the users by adding new users and assigning roles to them. You can also map external group users to the NetWitness roles.

These are the procedures for setting up and managing users:

- Review the Preconfigured NetWitness Platform Roles
- (Optional) Add a Role and Assign Permissions
- Verify Query and Session Attributes per Role
- Set Up Users
- (Optional) Map User Roles to External Groups

Review the Preconfigured NetWitness Platform Roles

To simplify the process of creating roles and assigning permissions, there are preconfigured roles in NetWitness.

Role	Permission
Administrators	Full system access. The System Administrators persona is granted all permissions by default.
Analysts	Access to meta and session content but not to configurations. The Security Operation Center (SOC) Analysts persona is centered around investigation, ESA Alerting, Reporting, and Respond, but not system configuration.

Role	Permission
Reporting_Engine_ Content_Administrators	Access to manage the Live content. Users with the Reporting Engine Content Administrator role can deploy Reporting Engine content (rules, reports, charts, and lists) from Live Content, view and manage permissions to the deployed content in Reporting Engine.
Data_Privacy_Officers	The Data Privacy Officer (DPO) persona is similar to Administrators with additional focus on configuration options that manage obfuscation and viewing of sensitive data within the system (see the *Data Privacy Management Guide*). Users with the DPO role can see which meta keys are flagged for obfuscation, and they also see obfuscated meta keys and values created for the flagged meta keys.
Malware_Analysts	Access to investigations and malware events. The only access granted to the Malware Analysts persona is the Malware Analysis module.
Operators	Access to configurations but not to meta and session content. The System Operators persona is focused on system configuration, but not investigation, ESA, Alerting, Reporting, and Respond.
Respond_Administrator	Access to all Respond permissions. The Respond Administrator persona is focused on system configuration of Respond.
SOC_Managers	Same access as Analysts plus additional permission to handle incidents. The SOC Managers persona is identical to Analysts, but with permissions necessary to configure Respond.
UEBA_Analysts	Access to the NetWitness UEBA service in the **Investigate** > **Users** view. NetWitness UEBA is an advanced analytics solution for discovering, investigating, and monitoring risky behaviors across all entities in your network environment. **Note:** You do not need to set up specific permissions for this role. You only need to assign this role to a user, and that user will have access to NetWitness UEBA.

The administrator can also add custom roles.

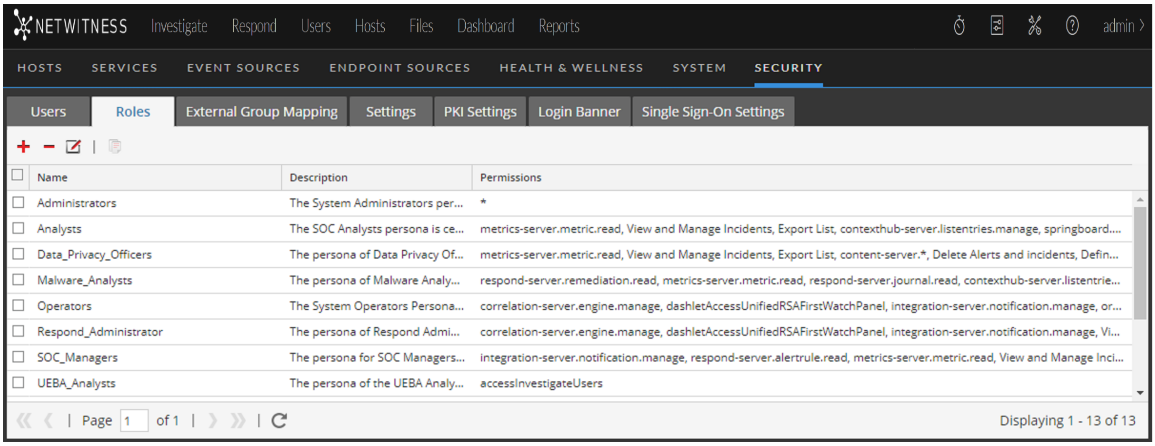
(Optional) Add a Role and Assign Permissions

Although NetWitness has preconfigured roles, you can add custom roles. For example, in addition to the preconfigured Analysts role you could add custom roles for AnalystsEurope and AnalystsAsia. For a detailed list of permissions, see Role Permissions.

Each of the following procedures starts on the **Roles** tab.

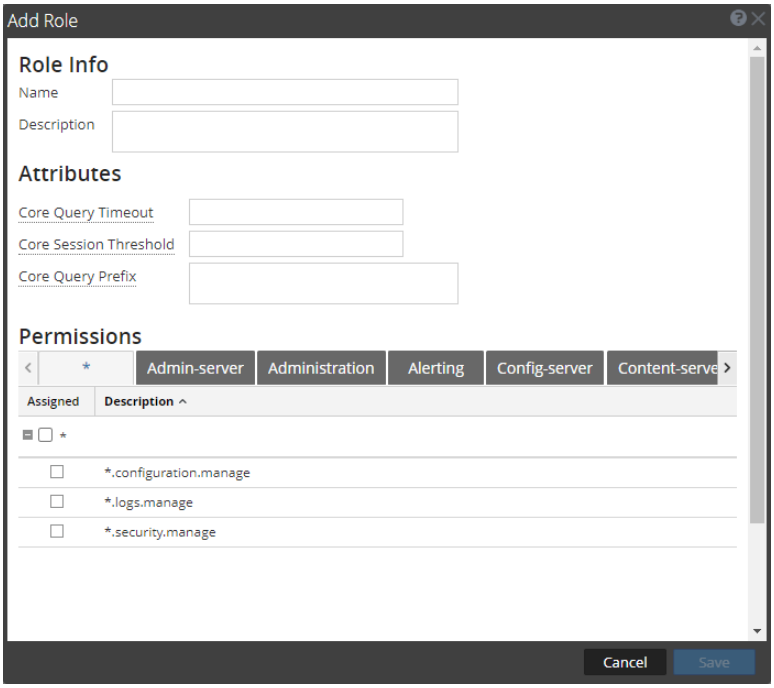
To navigate to the Roles tab:

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **Roles** tab.



Add a Role and Assign Permissions

1. In the **Roles** tab, click ✚ in the toolbar.
2. The **Add Role** dialog is displayed.



3. In the **Role Info** section, type the following information for the role:
   - **Name**
   - (Optional) **Description**
4. In the **Attributes** section, enter the desired values for each attribute. For more information on attributes, see Verify Query and Session Attributes per Role.
5. In the **Permissions** section:
   - Click ‹ and › to scroll through the modules.
   - Select a module the role will access.
   - Select each permission the role will have.
6. Repeat the previous step until you select all permissions to assign to the role.
7. Click **Save** to add the new role, which is effective immediately. You can now assign the new role to users.

Change Permissions Assigned to a Role

1. In the **Roles** tab, select the role and click .
   The **Edit Role** dialog is displayed.
2. In the **Permissions** section:

- Click  and  to scroll through the modules. Select a module to revise permissions for it. Select or deselect each permission.
- Select a module to revise permissions for it.
- Select or deselect each permission.

3. Repeat the previous step until the role has the required permissions.
4. Click **Save**. The revised permissions are effective immediately.

Duplicate a Role

An efficient way to add a new role is to duplicate a similar role, save it with a new name and revise the permissions that are already assigned.

1. In the **Roles** tab, select the role you want to duplicate and click .
2. Type a new role name and click **Save**.
3. To change the permissions, follow the steps in the next procedure.

Delete a Role

You can delete a role if it is not assigned to any users.

1. In the **Roles** tab, select the role and click .
2. A dialog requests confirmation that you want to delete the role. Click **Yes**.

Verify Query and Session Attributes per Role

After you define your user roles, it is important to verify the query and session attributes that are set for each role. You can adjust these settings according to your requirements.

You can understand how these role settings impact individual user settings and what happens if a user is a member of multiple roles.

Query and Session Attributes

Query and session attributes determine how to handle the queries that a user runs. These attributes enable you to lock down the information that users can retrieve, and the attributes apply to all sessions of users assigned to a role.

Depending on your requirements, you can specify the following query-handling attributes for a user role:

- **Core Query Timeout** is an optional setting that applies to NetWitness Core services. It specifies the maximum number of minutes that a user can run a query. If this value is set, it must be zero (0) or greater. A value of zero represents no timeout. The default value is 5 minutes.
- **Core Session Threshold** is a required setting. This value must be zero (0) or greater. The default is 100000. The limit you specify here overrides the **Max Session Export** value defined in the Investigate view settings. If the threshold is greater than zero, a query optimization will extrapolate the total session counts that exceed the threshold. When the meta value count returned by the query reaches the threshold, the system will:
  - Stop its determination of the session count.
  - Show the threshold and percentage of query time used to reach the threshold.
- **Core Query Prefix** is an optional filter applied to queries the user runs. The prefix restricts query results that the user sees. For example, the `'service' = 80` query prefix is prepended to any queries run by the user, and the user can only access metadata of HTTP sessions. If you set up data privacy using a whitelist, every meta key specified in the core query prefix must also be whitelisted as described in the *Data Privacy Management Guide*.

Note: In Version 11.1 and later, you can use configured meta entities in a Core Query Prefix. For additional information about configuring meta entities, refer to the *Core Database Tuning Guide*.

The query-handing attribute settings applied for a user depend on the role memberships of the user. It is important to verify the query-handling attribute settings for your roles.

How Query-Handling Attribute Settings Apply to Individual Users

If a user is a member of multiple roles, the following logic applies for the user:

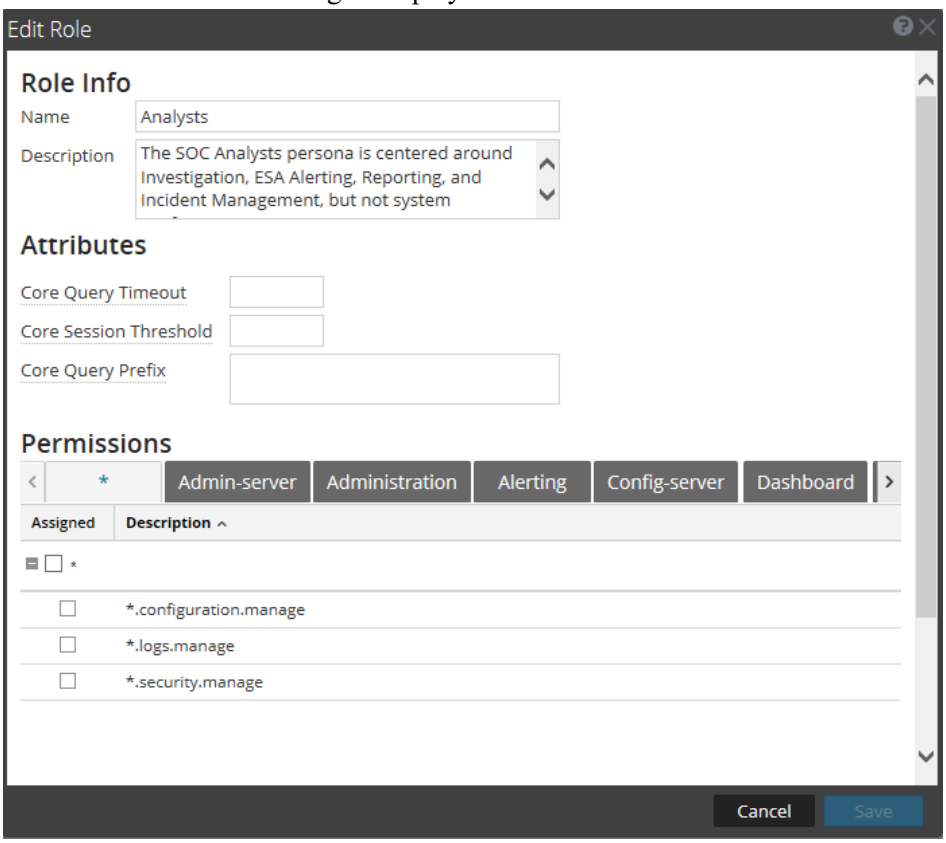
- **Query Timeout:** The most permissive (highest) value of all assigned roles is applied to the user.
- **Query Prefix:** The query prefixes of each of the user roles are AND'd together.

- **Session Threshold:** The highest value of all the assigned roles is applied to the user.

Set Query Handling Attributes for a User Role

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.

2. Click the **Roles** tab. If you are adding a role, click . If you are editing a role, select the role and click .
   The Add or Edit Role dialog is displayed.



3. To set the attributes for the role, in the **Attributes** section:
   - (Optional) In the **Core Query Timeout** field, type the maximum number of minutes that a user can run a query. This timeout applies to queries performed from Investigate.
   - Type a **Core Session Threshold** for the system to stop its determination of the session count.
   - (Optional) Type a **Core Query Prefix** to filter query results that role members see in the Investigate Navigate view, Events view, and Event Analysis view. You can specify a query that is prepended to all queries executed by users with a specific role. For example, if the `'service' = 80` query prefix is prepended to all queries by users in this role, the users can only access

metadata of HTTP sessions. If users attempt to navigate to non-HTTP event, the view is not displayed.

Caution: If you add or modify a query prefix, and a user who has access to everything falls into a user role that is now restricted by a query prefix, that user will still be able to access cached data and view restricted content. To remove visibility of restricted content and enforce the query prefix in the Navigate or Events view, go to  (Admin) > System > Investigation, and use Reconstruction Cache Settings to clear the cache for all services available to affected users. The user will still be able to access cached data and view restricted content in the Event Analysis view until the analyst restarts the NetWitness Investigate services or the cache is cleared automatically after 24 hours.

4. Click **Save.**

Set Up Users

The procedures to set up a new user are described below.

- Add a User and Assign a Role
- Enable, Unlock, and Delete User Accounts

Add a User and Assign a Role

All NetWitness users must have a local or external user account. You can add a new user to each type of user account, local and external. You can assign role to a local user.

The following considerations are important when managing local and external user accounts.

Local User Account	External User Account
Managed within NetWitness.	Managed externally and outside the scope of this document.
Roles assigned directly.	Roles assigned by external group mapping.
Derives permissions from each role assigned to the user	Derives permissions from each role mapped to the account's external user group, as explained in (Optional) Map User Roles to External Groups.
NetWitness manages all user information.	NetWitness manages user identification only. This includes Username, Full Name and Email.

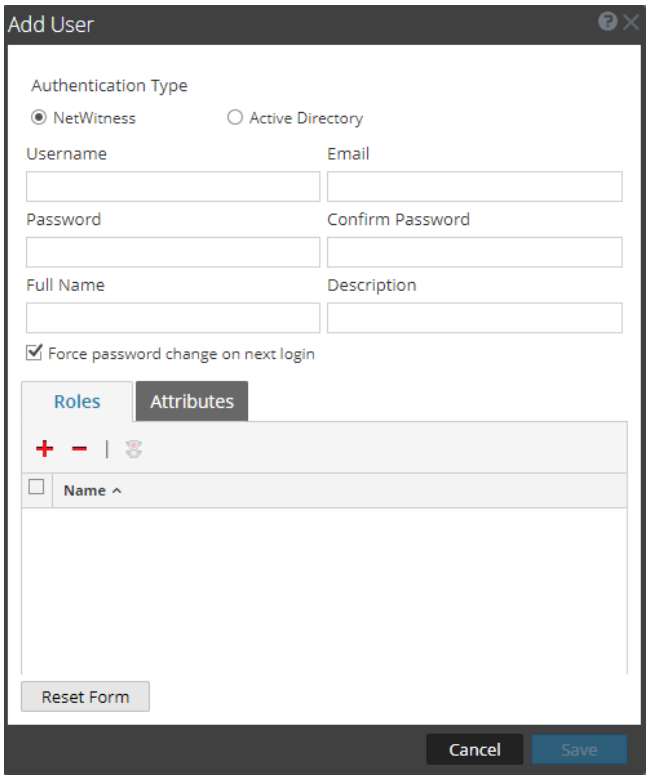
Each of the following procedures starts on the Users tab. To navigate to the Users tab, go to  (Admin) > Security. The Security view is displayed with the Users tab open.

Add a Local User

To add a local user account and assign a role to the user:

1. In the **Users** tab, click  in the toolbar.

   The **Add User** dialog is displayed.



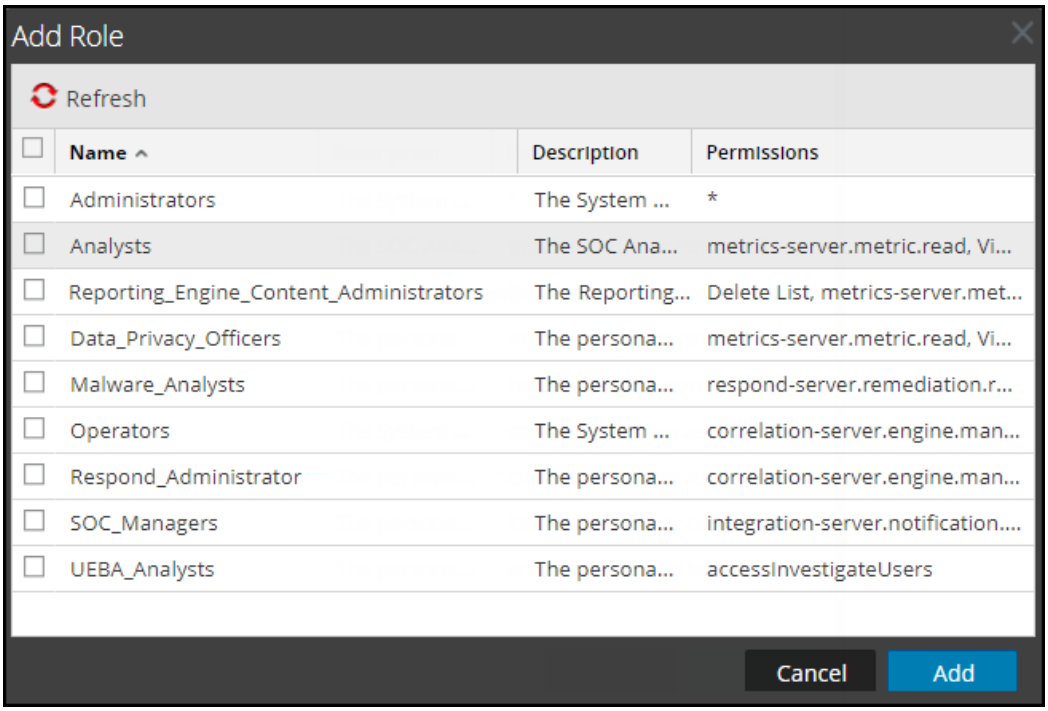
2. Type the following account information for the new user:
   - **Authentication Type**: **NetWitness** is selected by default and is the correct choice when adding a local user. This option is only displayed when there are AD or PAM configurations set up to allow for selecting that authentication type.

     Note: If there are no AD or PAM configurations, the authentication type is set to NetWitness automatically and there are no other options available.

   - **Username** for logging on to NetWitness
   - **Email** address
   - Password for logging on to NetWitness, in the **Password** and **Confirm Password** fields
   - **Full Name** of the new user
   - (Optional) **Description** of the user account
3. To expire the user password the next time the user logs on, select **Force password change on next login**.

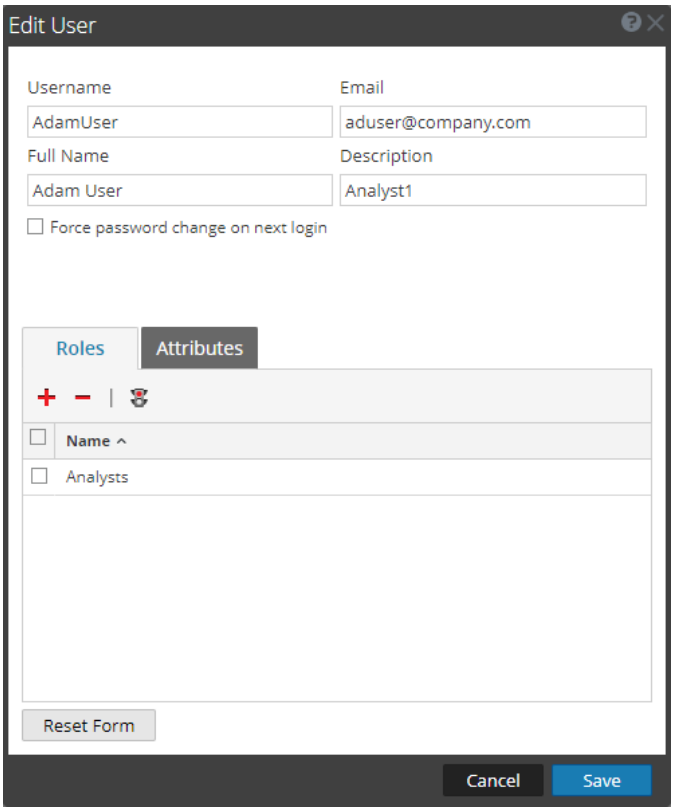
   This does not affect any active user sessions. The  appears in the user row to show that the user password expired. After a password is expired, you cannot undo it. This checkbox is cleared the next time you edit the user account.
4. To assign a role to the user, click + in the **Roles** tab.

   The **Add Role** selection dialog shows the list of available roles.

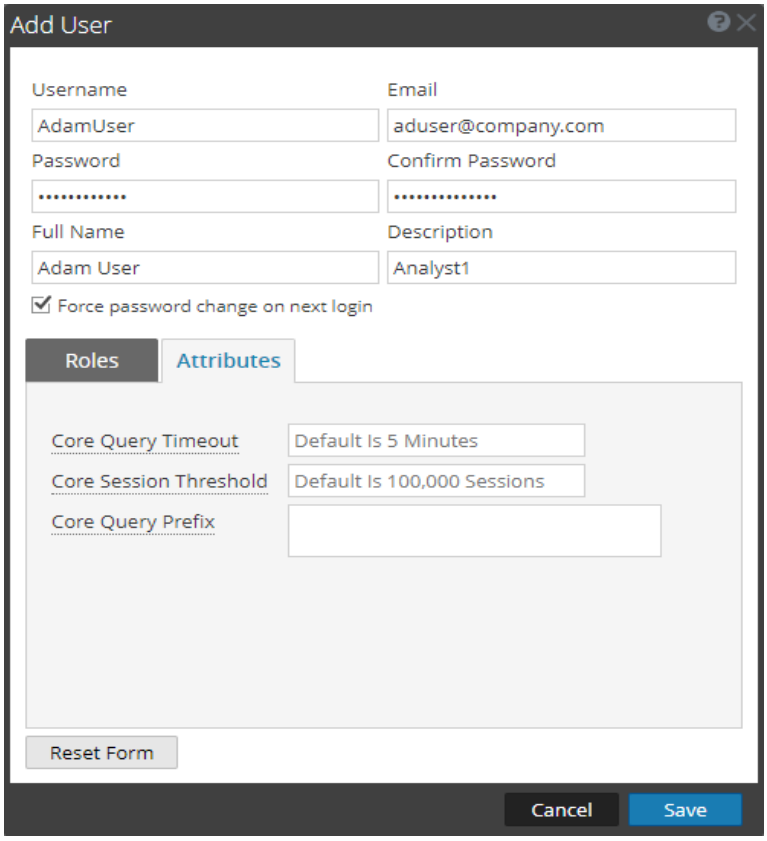


5. Select each role to assign and click **Add**.

   The **Add User** dialog shows each role assigned to the user.

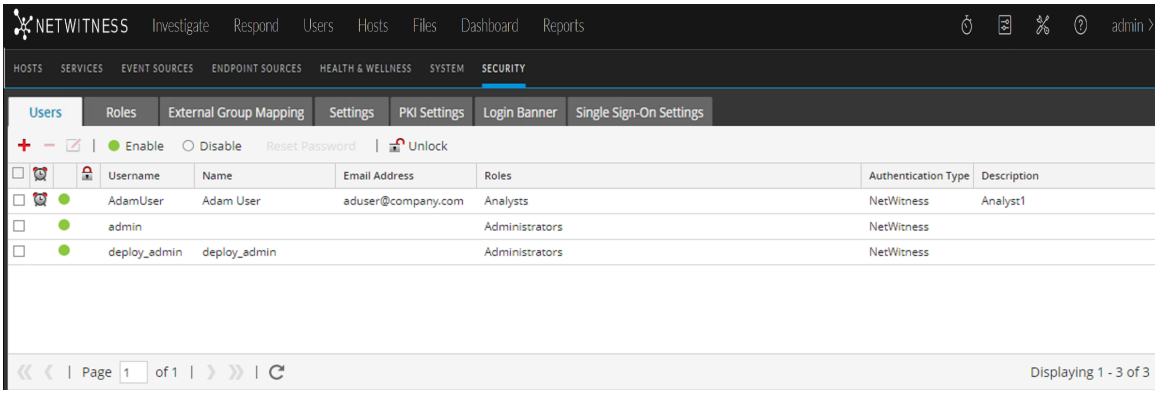


6. (Optional) To assign attributes to a user, go to **Attributes** and modify the appropriate values. These attributes are unique to the user and follow all the same rules for attributes within roles. For more information on attributes, see Query and Session Attributes.



7. (Optional) Select a role and click to **Show all permissions** for the role.
8. Click **Save**.

   The **Users** tab shows the new user and each role assigned to the user. The account is active immediately.



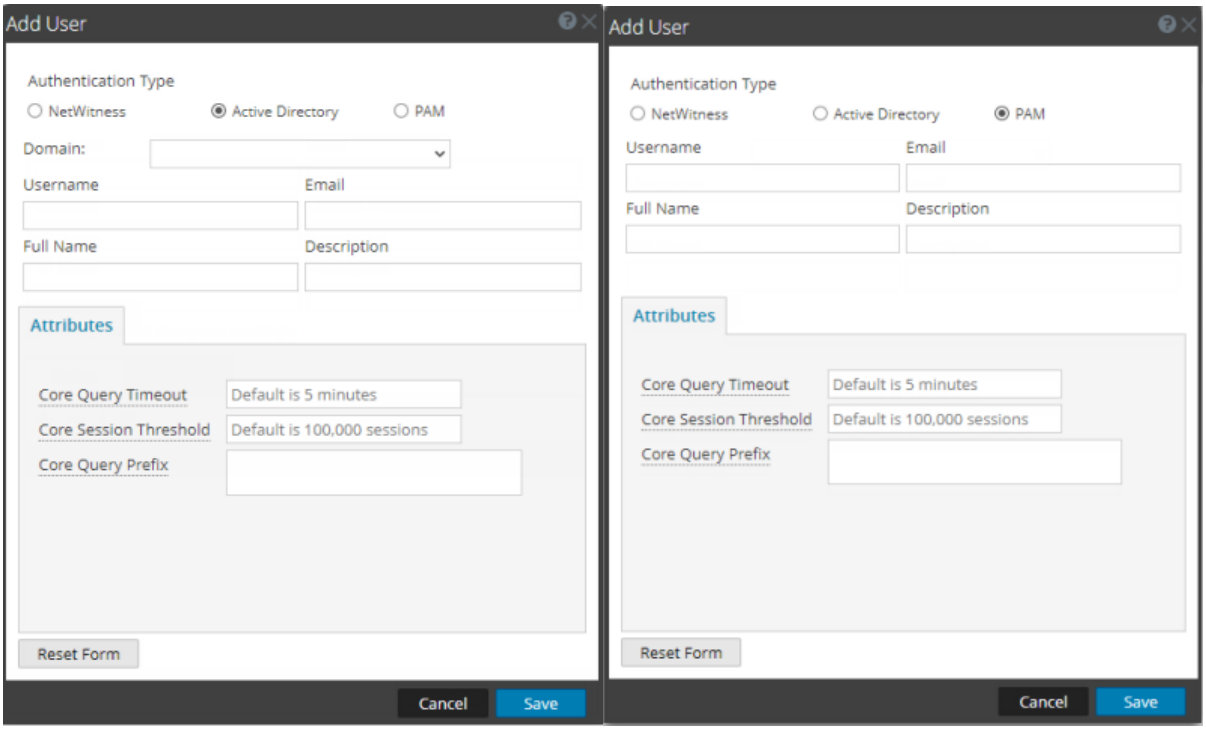
Add a User for External Authentication

To add a user for external authentication:

1. In the **Users** tab, click + in the toolbar.

   The **Add User** dialog is displayed.

2. For **Authentication Type**, select either **Active Directory** or **PAM**. The dialog will update to show the required fields for the selected external authentication type.



3. Type the following information:

   - **Domain** (if select Active Directory authentication only): Select the Active Directory domain for the user from the drop-down list of available domains.
   - **Username** for logging on to NetWitness
   - **Email** address
   - **Full Name** of the new user
   - (Optional) **Description** of the user account

4. In the Attributes section, type the following information.

   a. **Core Query Timeout**- most permissive (highest) value of all assigned roles is applied to the user.

   b. **Core Session Threshold** - query prefixes of each of the user roles are AND'd together.

   c. **Core Query Prefix** - highest value of all the assigned roles is applied to the user.

5. Click **Save**. The Users tab shows the new user account, which still needs a role and permissions.

6. To map a role to the new user, see (Optional) Map User Roles to External Groups.

Change User Information or Roles

To change a user's account information or assigned roles:

1. In the **Users** tab, select a user and click  in the toolbar.

   The **Edit User** dialog is displayed.
2. To edit user information, change any of the following fields:
   - **Email**
   - **Full Name**
   - **Description**
3. To expire the **internal** user password the next time the user logs on, select **Force password change on next login**.

   This does not affect any active user sessions. The  appears in the user row to show that the user password expired. After a password is expired, you cannot undo it. This checkbox is cleared the next time you edit the user account.
4. In the **Roles** section:
   - To assign another role, click  , select a role and click **Add**.
   - To remove an assigned role, select the role and click .
5. Click **Save**.

Delete a User

To delete a user:

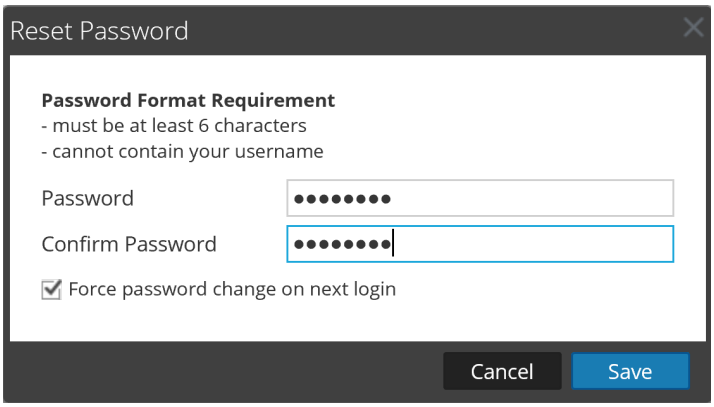
1. In the **Users** tab, select a user.
2. In the toolbar, click .
3. Click **Save**.

Note: To fully delete a user that is externally authenticated by Active Directory, you must also delete the user from the AD Group.

Reset a User Password

To reset a user password:

1. In the **Users** tab, select a user.
2. In the toolbar, click **Reset Password**.



The **Password Format Requirement** section lists the specific requirements for the password. Administrators can adjust these requirements for all internal users in the password policy. See Configure Password Complexity.

3. Choose whether to force a password change the next time the user logs in to NetWitness.
4. Click **Save**.

Enable, Unlock, and Delete User Accounts

All users of NetWitness must either have a local user account with username and password or have an external user account. Within NetWitness, you can enable, disable, and delete local user accounts.

The first time an external user logs into NetWitness, a new user entry is automatically created with NetWitness. NetWitness manages only user identification information; for example, Full Name and Email.

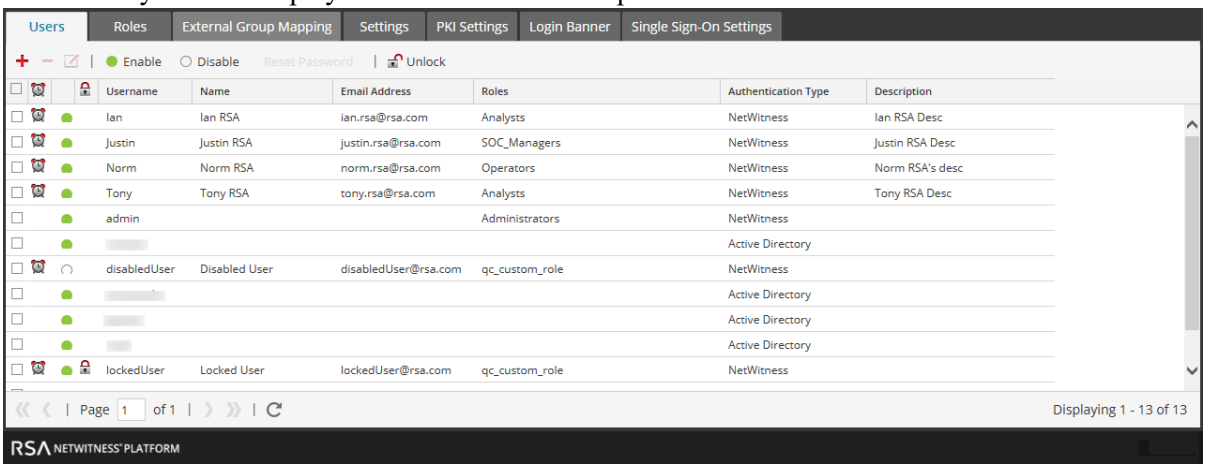
You can unlock locked accounts for both local and external users.

Enable Disabled NetWitness User Accounts

To enable NetWitness user accounts that have been disabled:

1. Go to **(Admin) > Security**.

   The Security view is displayed with the **Users** tab open.



2. In the **Users** grid, select one or more accounts.
3. Click Enable.
   A successful message displays for enabled accounts, and the users can log in to NetWitness.

Disable NetWitness User Accounts

You can block user access by disabling users. Disabling the user does not delete user preferences. This action blocks user access without deleting user preferences so that upon re-enabling users, user preferences are intact. You can re-enable users to restore user access. Disabling users applies only to Local users and not External Users.

To disable NetWitness user accounts:

1. In the **Users** grid, select one or more accounts.
2. Click Disable.
   A successful message displays for disabled accounts, and the users can no longer log in to NetWitness.

Unlock Locked NetWitness User Accounts

A user is locked out for a period of time after a number of failed consecutive login attempts. To unlock NetWitness user accounts that are locked due to excessive failed login attempts:

1. In the **Users** grid, select one or more accounts.
2. Click Unlock.
   A successful message displays for unlocked accounts, and the users can log in to NetWitness.

Delete NetWitness User Accounts

If not using External Authentication, a user can log on to NetWitness using a local account. These local accounts are directly managed using NetWitness. To revoke access to a local user, either disable the account or delete the account completely from the system.

Note: This deletes all user preferences for the account from NetWitness. If this is not the intention, disable the user instead of deleting the user.

To delete NetWitness user accounts:

1. Go to  **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. In the Users list, select one or more accounts.
3. Click .
   A warning dialog requests confirmation.
4. If you want to delete the accounts, click **Yes**.
   The accounts are removed from NetWitness, and the users can no longer log in to NetWitness.

(Optional) Map User Roles to External Groups

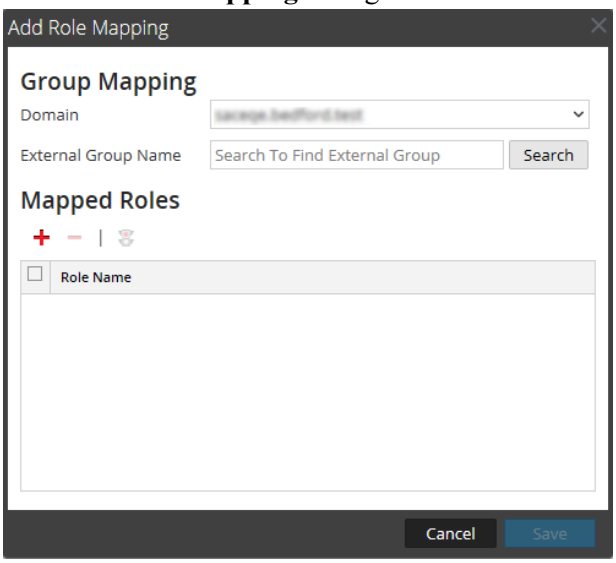
In NetWitness, external groups derive permissions for various modules and views from NetWitness user roles, which have permissions assigned to them. To provide access to an external group, map user roles to it. To modify an external group's access, edit the roles mapped to it. Add and delete roles until the external group has the necessary access. Changes take effect immediately.

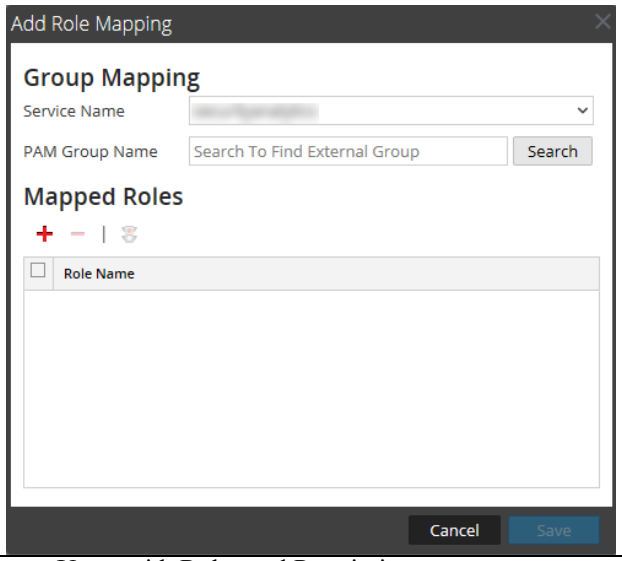
In the Settings tab, you must set up a method for external user authentication to make external groups visible to NetWitness.

Add Role Mapping for an External Group

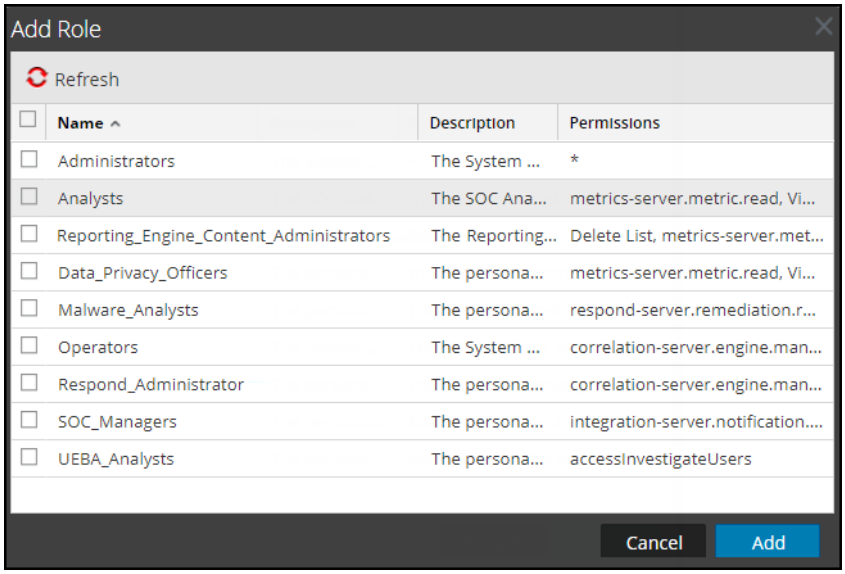
To add role mapping for an external group:

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **External Group Mapping** tab.
3. In the toolbar, click +.
   The **Add Role Mapping** dialog for the external authentication method you selected is displayed.





4. Click **Search** and search for an external group name in the Search for External Groups, then select an external group name.
5. To add roles to the group mapping, click + in the **Mapped Roles** section.
   The **Add Role** dialog is displayed.



6. Select the checkbox in the title bar to select all roles, or select roles individually.
7. To add the roles to the **Mapped Roles** section in the Add Role Mapping dialog, click **Add**.
   The dialog closes and the selected roles are displayed in the Mapped Roles section.
8. If you want to delete roles from the **Mapped Roles** section, select the roles and click − .
9. When the **Add Role Mapping** dialog reflects the role mapping that you want to define for the group, click **Save**.
   The Add Role Mapping dialog closes, and the new role mapping is listed in the External Group Mapping tab list.

Edit Role Mapping for a Group

To edit role mapping in a group:

1. In the **External Group Mapping** action bar, click **Edit**.
   The **Edit Role Mapping** dialog is displayed with the group name in the **External Group Name** field.
2. To add roles to the mapping, click + in the **Mapped Roles** section.
   The Add Role dialog is displayed.
3. Select the checkbox in the title bar to select all roles, or select roles individually.
4. To add the roles to the **Mapped Roles** section in the **Add Role Mapping** dialog, click **Add**.
   The dialog closes, and the selected roles are displayed in the Mapped Roles section.

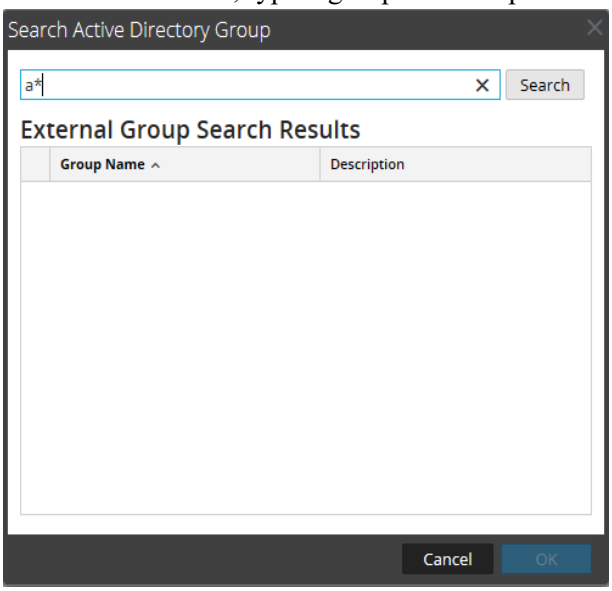
5. If you want to delete roles from the **Mapped Roles** section, select the roles and click .
6. When the **Edit Role Mapping** dialog reflects the role mapping that you want to define for the group, click **Save**.
   The dialog closes, and the edited role mapping is listed in the External Group Mapping tab.

Search for External Groups

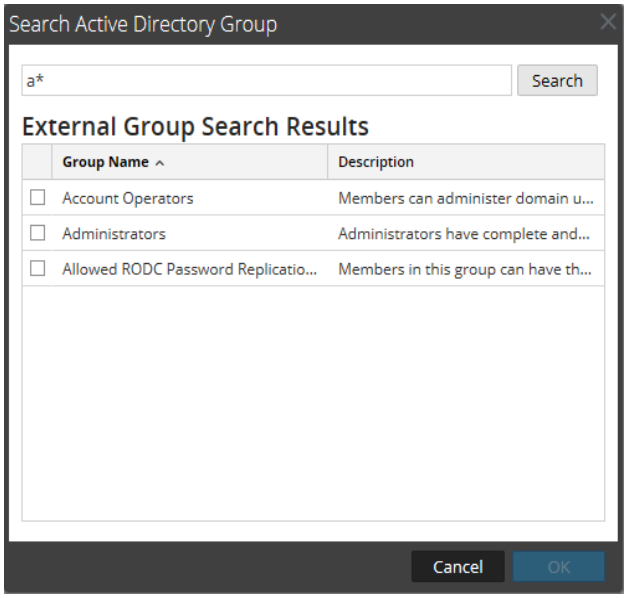
The instructions for searching for external groups that have NetWitness user roles mapped to them.

To search for an external group:

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **External Group Mapping** tab.
3. In the toolbar, click or .
   The **Add Role Mapping** dialog for the external authentication method you selected is displayed.
4. The **Group Mapping** section is dependent on the selected external authentication method.
   - For **Active Directory**, select a **Domain**. Then click **Search** next to **External Group Name**.
   - For **PAM**, click **Search** next to **PAM Group Name**.
     The **Search External Groups** dialog is displayed.
5. In **Common Name**, type a group name or part of a group name with the wild card character (*).



6. Click **Search**.
   The results are displayed in the **External Group Search Results** section.



7. Select the group to which you want to assign roles and click **OK**.

Set Up Multi-Factor Authentication

You can set up Multi-Factor Authentication (MFA) for NetWitness using one of the following methods:

- ADFS Log in to NetWitness with SecurID Passcode.
- PAM SecurID Log in to NetWitness for AD Users.

ADFS Log in to NetWitness with SecurID Passcode

Single Sign-On (SSO) functionality of NetWitness can be leveraged where the user authenticates using ADFS log in user interface with AD user credentials followed by the SecurID passcode.

Note: This method is suitable only for single AD users.

Prerequisites

- NetWitness Platform (NW) version 11.4 or later
- MS Active Directory Federation Services (ADFS) - MS Windows Server 2012 R2 or later
- MS Active Directory (AD) – MS Windows Server 2008 R2 or later
- Authentication Manager (AM) 8.4 or later
- Authentication Agent for ADFS 1.0 or later

Perform the following configurations:

1. Configure Authentication Manager.
2. Configure NetWitness.
3. Configure ADFS.

Configure Authentication Manager

Configure Active Directory as an Identity Source in Authentication Manager using the steps described in the section Add an Identity Source.

Configure NetWitness

1. Configure Active Directory for External Authentication to NetWitness using the steps described in the section Configure Active Directory.
2. NetWitness must be configured for SSO using the steps described in the section Configure Single Sign-On.

Configure ADFS

ADFS must be configured for SSO in NetWitness. You must copy the exported metadata (see step 9 in Configure Single Sign-On) to ADFS and perform the following steps:

1. Go to **Server Manager > Tools > ADFS management > Trust Relationships**.
2. On the right-side, click **Add Relying Part Trust > Start**.
3. Click **Import data about the relying party from a file** and select the metadata file.
4. Click **Next**, and enter a display name.
5. Click **Next** until the **Close** button is displayed.
6. Ensure the **Open the Edit Claim Rules** option is selected.
7. Click **Close**.
8. In the **Edit Claim Rules** dialog, click **Add Rule**.
9. In the **Add Transform Claim Wizard** dialog, click **Next**.
10. Enter a claim rule name.
11. In the **Attribute Store** drop-down menu, select **Active Directory**.
12. In the Mapping of LDAP attributes table, on the left-side select **SAM-Account-Name** and on right-side, select **Name ID**.

    Note: Only one mapping is required.

13. Click **Finish**.
14. Click **Apply**.
15. Click **OK**.

    Next you need to configure MFA using Authentication Agent in ADFS. The agent is freely available at (https://community.securid.com/t5/securid-authentication-agent-for/tkb-p/auth-agent-ad-fs-documentation) and for more information on configuration, see ® Authentication Agent 2.0.3 for Microsoft® AD FS Administrator's Guide.

PAM SecurID Log in to NetWitness for AD Users

In this method only SecurID Passcode is required for authenticating to NetWitness. Authentication Manager takes care of the authentication to AD without requiring the password from the user.

After the configuration, the user registered in the Active Directory can log in to NetWitness using the SecurID passcode.

Prerequisites

- NetWitness Platform (NW) version 11.0 or later
- MS Active Directory (AD) – MS Windows Server 2008 R2 or later
- Authentication Manager (AM) 8.2 or later

Perform the following configurations:

1. Configure Authentication Manager.
2. Configure NetWitness.

Configure Authentication Manager

Configure Active Directory as a Identity Source in Authentication Manager using the steps described in the section Add an Identity Source.

Configure NetWitness

Complete the additional configuration for Authentication Manager and NetWitness configuration as described in the section Configure PAM Login Capability.

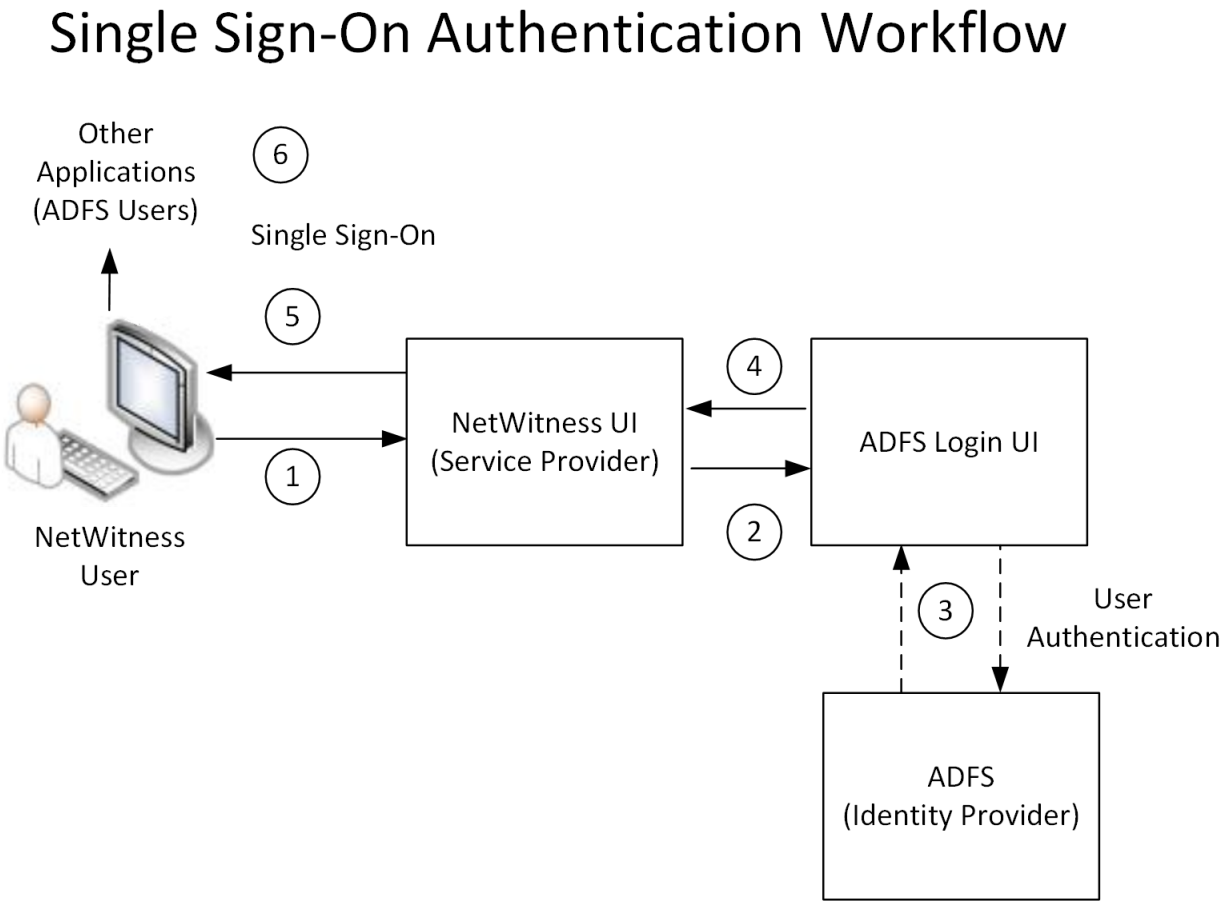
Set Up Single Sign-On Authentication

Note: In 11.4 or later, Single Sign-On (SSO) authentication can be used to access the UI however only one Active Directory is supported. SSO authentication is not supported on an Analyst UI Deployment.

Single Sign-On authentication enables the user to log in to NetWitness or any other application if the user is authenticated by the same Identity Provider (IDP). The Active Directory Federation Services (ADFS) is the only supported IDP and the protocol used for SSO is SAML 2.0.

NetWitness Single Sign-On Authentication Workflow

The following workflow shows how the user can access NetWitness using Single Sign-On authentication.



The workflow of SSO authentication shows the following:

1. User tries to access the NetWitness UI using the web browser. For example, https://nw-host/login.
2. The user is prompted to login into the IDP (ADFS) login page.

3. The user enters the credentials for authentication.
4. If the authentication is successful, NetWitness authorizes the user based on the user groups configured on the Active Directory Server and External Role Mapping in NetWitness.
5. If the authorization is successful, the user is logged into the NetWitness.

Note: If the single sign-on authentication fails, the user cannot access the NetWitness. For more information, see Troubleshooting.

Configure Single Sign-On

The following workflow describes the tasks to be performed in sequence to configure Single Sign-On authentication on NetWitness.



Configure ADFS as IDP for NetWitness

For instructions on how to configure ADFS as IDP for NetWitness, see the *Configure SAML 2.0 provider settings for portals* topic in Microsoft documentation.

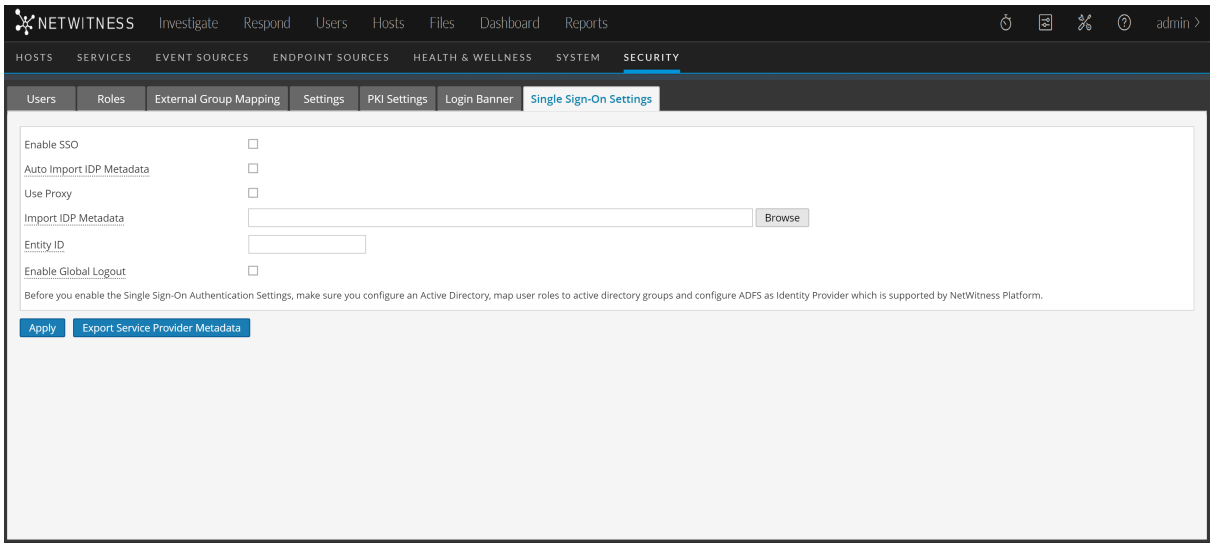
Map User Roles to External Groups

Atleast one Active Directory group should be mapped to an administrator role in NetWitness. For instructions on how to map user roles to Active Directory groups, see (Optional) Map User Roles to External Groups.

Enable Single Sign-On

1. Go to  **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **Single Sign-On Settings** tab.

3. Select the **Enable SSO** checkbox.



4. Select the **Auto Import IDP Metadata** if you want the latest IDP metadata to be automatically downloaded at regular intervals.
   When you select this check box, a **Metadata URL** field will be displayed where you must enter the IDP metadata URL.

5. Select **Use proxy** checkbox for the requests to IDP tobe routed through the proxy configured in **(Admin)** > **System** > **HTTP Proxy settings**.

6. Select **Import IDP Metadata** to manually import the meta data and enter the IDP metadata URL.
   Note: Make sure you update the link every time the IDP metadata is updated.

7. Enter a unique entity ID to identify the NetWitness instance in the Identity Provider.

8. (Optional) Select the **Enable Global Logout** checkbox if you want to be logged out of NetWitness along with all the other associated sessions authenticated by IDP.

9. Click **Apply**.
   This may take some time however we recommend you to restart the admin-server immediately. To export the metadata in an XML format either click the link in the notification tray and download the metadata or click **Export Service Provider Metadata** .

Note: The exported Service Provider metadata must be imported to IDP. For more information, see the *Configure SAML 2.0 provider settings for portals* topic in Microsoft documentation.

(Optional) Set Up Public Key Infrastructure (PKI) Authentication

Note: In 11.3 or later, PKI authentication can be used to access the NetWitness UI.

PKI is an authentication method which allows the users to access the NetWitness User Interface (UI) using digital certificates.

The certificates are issued by a Third-Party Certificate Authority (CA) which is external to NetWitness. The following categories of certificates are required for PKI authentication:

- Trusted CA Certificates
- User Certificate (issued by the CA)
- NetWitness Server Certificate (private key and its certificate chain) - Optional

Trusted CA Certificates

Trusted CA certificates are a collection of certificates used by NetWitness as trusted authorities to validate the certificate provided by the user. If the user does not have a certificate signed by one of these CA(s), the user cannot access the NetWitness Platform UI.

User Certificate

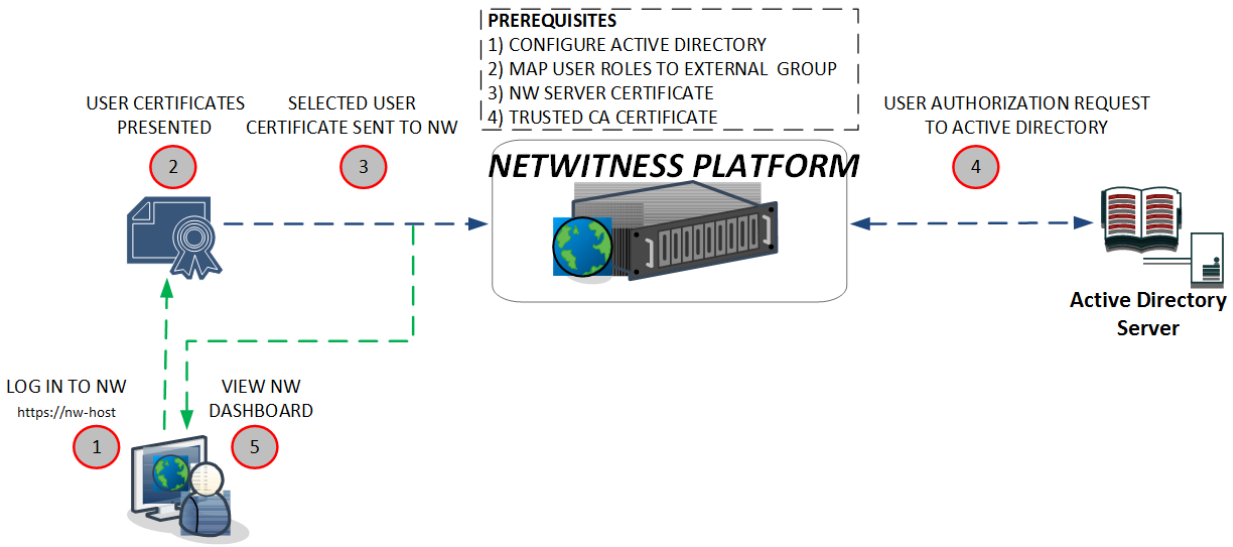
A user certificate is issued by a CA that is trusted by NetWitness and is used by the NetWitness Platform user to present the user identity. By default, user certificates are identified by most browsers. If the certificates are not displayed, you must import the certificates into the browser certificates store.

NetWitness Server Certificate (Optional)

A NetWitness Server certificate is issued by a trusted CA and is used by NetWitness Server to present its identity. If you access the NetWitness Platform UI using HTTPS, the certificate is displayed in the web browser.

NetWitness PKI Authentication Workflow

The following workflow shows how the user can access NetWitness using PKI authentication.



The workflow of PKI authentication shows the following:

1. Access the NetWitness UI using the web browser. For example, https://nw-host/login.
2. The user is prompted to select the user certificate.

Note: The certificate prompt appears differently depending on the browser.

3. The user selects the certificate. The browser sends the selected certificate to the NetWitness for authentication.
4. If the authentication is successful, the NetWitness authorizes the user based on the user groups configured on the Active Directory Server and External Role Mapping in NetWitness.
5. If the authorization is successful, the user is logged into the NetWitness.

Note: If the certificate validation fails, the user cannot access the NetWitness.

Configure PKI Authentication

Perform the following tasks in sequence to configure PKI authentication on NetWitness.

1. Configure Active Directory
2. Map User Roles to External Groups
3. Import Server Certificate and Trusted CA Certificate
4. Enable PKI Authentication

Note: Only Active Directory is supported for PKI Authentication.

Configure Active Directory

For instructions on how to use Active Directory to authenticate external user logins, see Configure Active Directory.

Map User Roles to External Groups

To enable PKI based authentication, a minimum of one external group should be mapped to administrator role in NetWitness. For instructions on how to map user roles to external groups to provide necessary access, see (Optional) Map User Roles to External Groups.

Import Server Certificate and Trusted CA Certificate

To enable PKI authentication, you must import the Trusted CA certificate into the NetWitness. You can import the server certificate and the trusted CA certificate in the  (Admin) > Security view > PKI Settings tab.

Certificate Revocation List

A Certificate Revocation List (CRL) is a file that contains a list of revoked certificates with details such as the serial number and revocation date of each certificate. When a certificate validity is expired, it must be revoked to avoid any compromise of the certificate by unauthorized users. For example, if a NetWitness user resigns from an organization, then the user's certificate must be revoked by the issuing CA.

You can import the CRL or specify the HTTP URL issued by your trusted CA, so that NetWitness can validate with the CRL to block unauthorized users from accessing NetWitness.

To import the CRL or specify a HTTP URL into NetWitness, choose one of the following methods:

- **HTTP server** - This is the most common CRL location where the CA publishes the CRL to external applications using an HTTP server. The NetWitness reads the CRL using the HTTP URL.
- **Local CRL** - This allows the users to manually download the CRL from the CA and upload it to the NetWitness.
- **OCSP (Online Certificate Status Protocol) Responder** - This allows NetWitness to verify the revocation status of a particular certificate instead of validating the complete CRL. To specify a OCSP Responder, you need to provide the HTTP URL and optionally the OCSP Responder's Signing certificate. Make sure the OCSP Responder is online while adding the entry. When the OCSP Responder Signing Certificate is updated, you need to manually update the certificate in NetWitness.

Note: You can configure the CRL when you import the CA certificate or after importing the CA certificate.

User Principal Settings

User Principal Settings allow NetWitness to uniquely identify the user from the user certificate for PKI authentication. To identify the user, you must specify an attribute in the user certificate to extract the user name or user id. NetWitness must be configured to read the value of this attribute. NetWitness uses the extracted value of this attribute as username or user id for authorization and retrieves the user groups from an Active Directory (AD) server. By default, NetWitness extracts the entire value of the selected attribute, without filtering any characters. You can apply regular expression (RegEx) to refine the value extracted.

Note: The conversion of Distinguished Names (attribute in Subject Name or Subject Alternative Name) to a human readable format is done based on RFC2253 (LDAPv3). Therefore, any Relative Distinguished Name apart from the one defined in RFC2253 (LDAPv3) may display in hex format. For example, email attribute value may display #1160bcghryy637bchs774. You can apply RegEX to extract the value. NetWitness tries its best to extract values for such attributes.

User Principal settings can be configured when you import the Trusted CA certificate.

Lookup Query

A Lookup Query sends a specific query to Active Directory to retrieve the user object. Here is an example of a sample query for retrieving a user object from the Active Directory.

```
Sample Query:(&(objectCategory=Person)(objectClass=User)(CN=${nw-pki-user}))
nw-pki-user is replaced with the value extracted from the user certificate.
```

Caution: Make sure that the AD user account is active. AD does not return the user account expiry (accountExpires) information to NetWitness along with user details. Therefore, the NetWitness cannot validate the AD user account is expired or not.

Note: NetWitness does not validate the syntax of the lookup query. You must ensure proper query syntax is used to retrieve the user object from Active Directory.

Import NW Server Certificate with its Private Key

For instructions on how to import the NetWitness Server Certificate with its private key, see (Optional) Use a Custom Server Certificate.

Import Trusted CAs, Configuring CRL and User Principal Settings

To import a trusted CA:

1. Go to  **(Admin) > Security.**
   The Security view is displayed with the **Users** tab open.
2. Click the **PKI Settings** tab.

3. In the **Trusted CAs** section, click  .
   The Import Certificate Authority dialog is displayed.

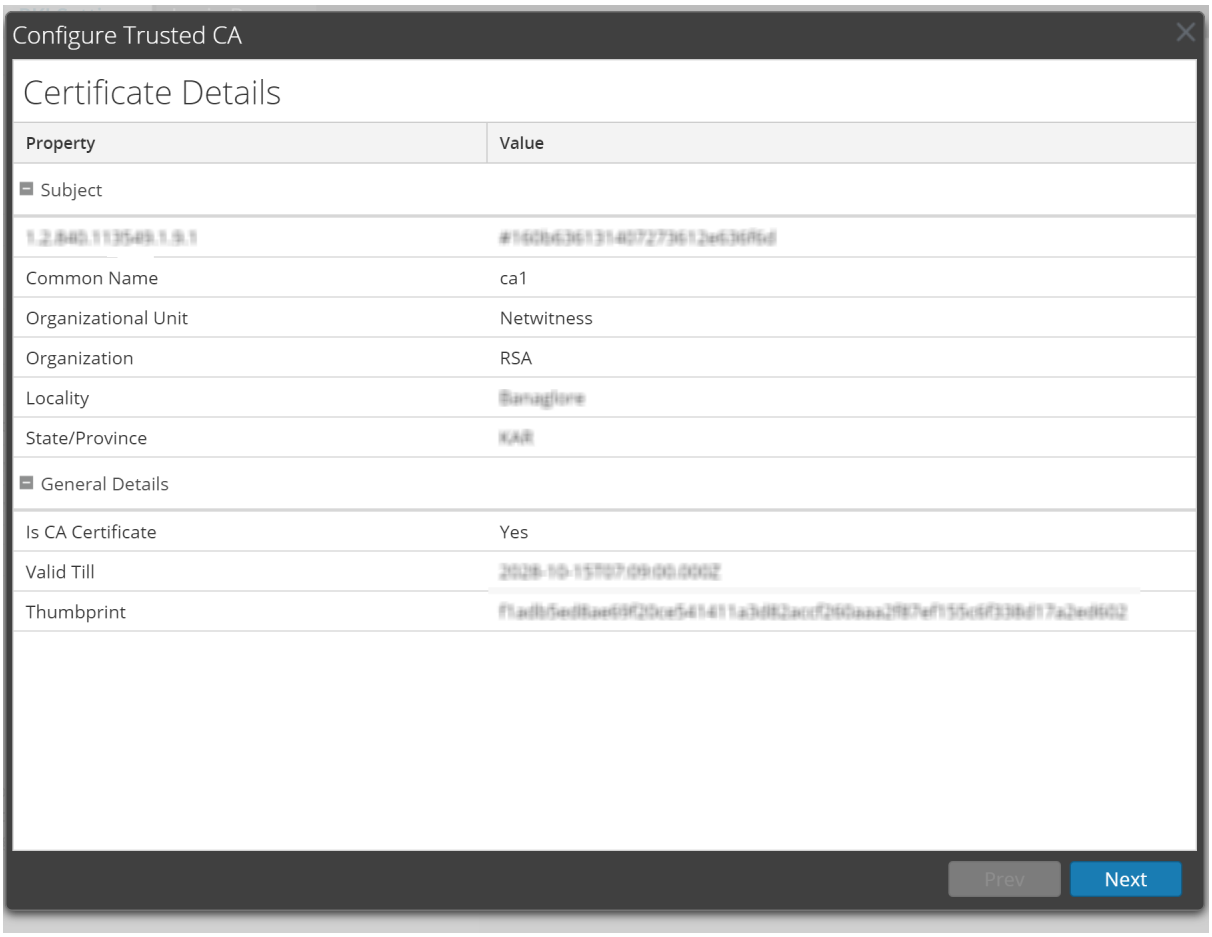


Note: The supported formats are .p12, .jks, .pfx, .pem, .crt, .der, and .cer

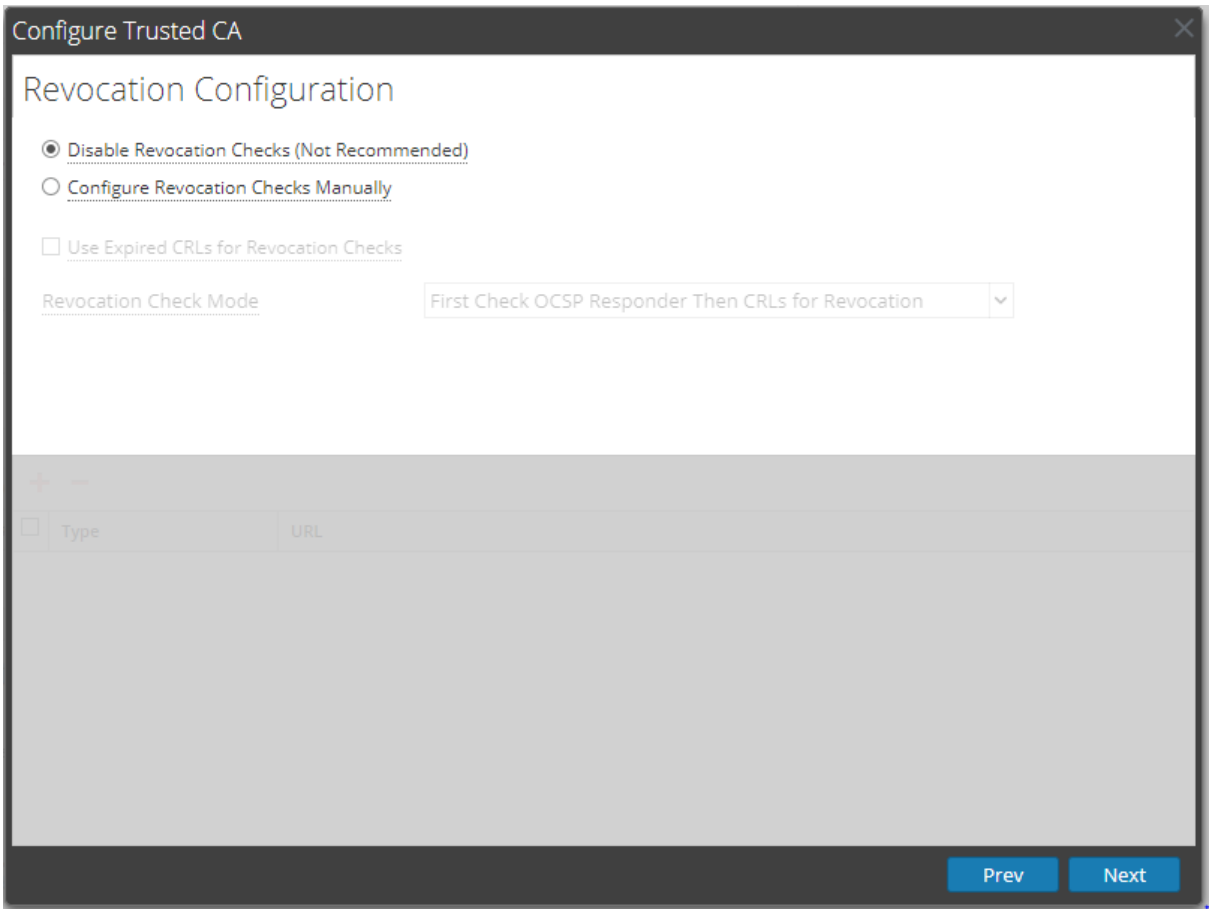
4. In the **CA Store File** field, click **Browse** and select a certificate.
5. In the **Password** field, enter the password of the certificate.

Note: The password is applicable only for .p12, .pfx and .jks certificate formats.

6. Click **Save**.
   A Configure Trusted CA dialog is displayed with Certificate Details.


Configure Trusted CA
Certificate Details
Property
Value
Subject
Common Name
ca1
Organizational Unit
Netwitness
Organization
RSA
Locality
State/Province
General Details
Is CA Certificate
Yes
Valid Till
Thumbprint
Prev
Next

7. Click **Next**.



8. In the Revocation Configuration section, do one of the following to configure the CRL revocation check.

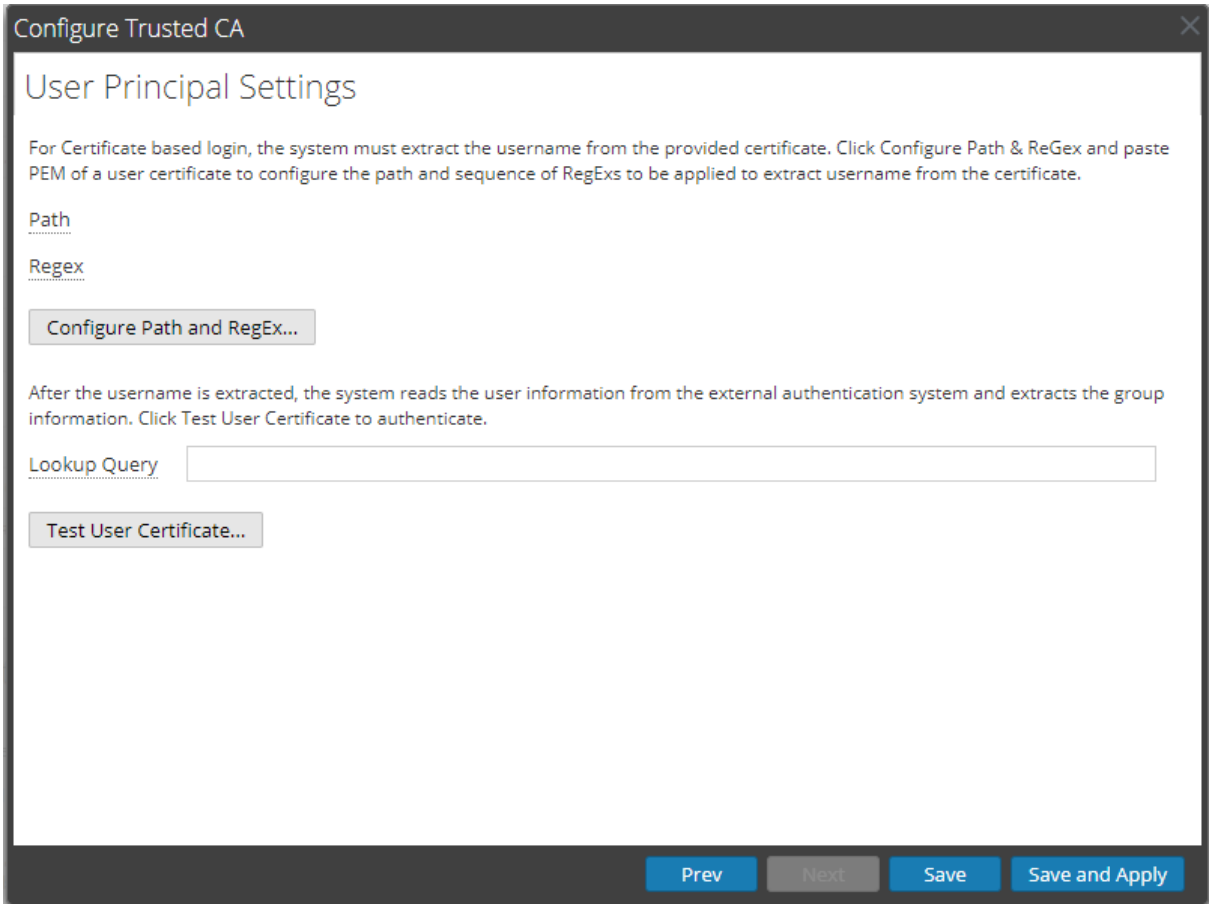
   - Select **Disable Revocation Checks** to disable the CRL revocation check.

   Warning: Disabling the revocation check may increase the risk of unauthorized users logging in to NetWitness.

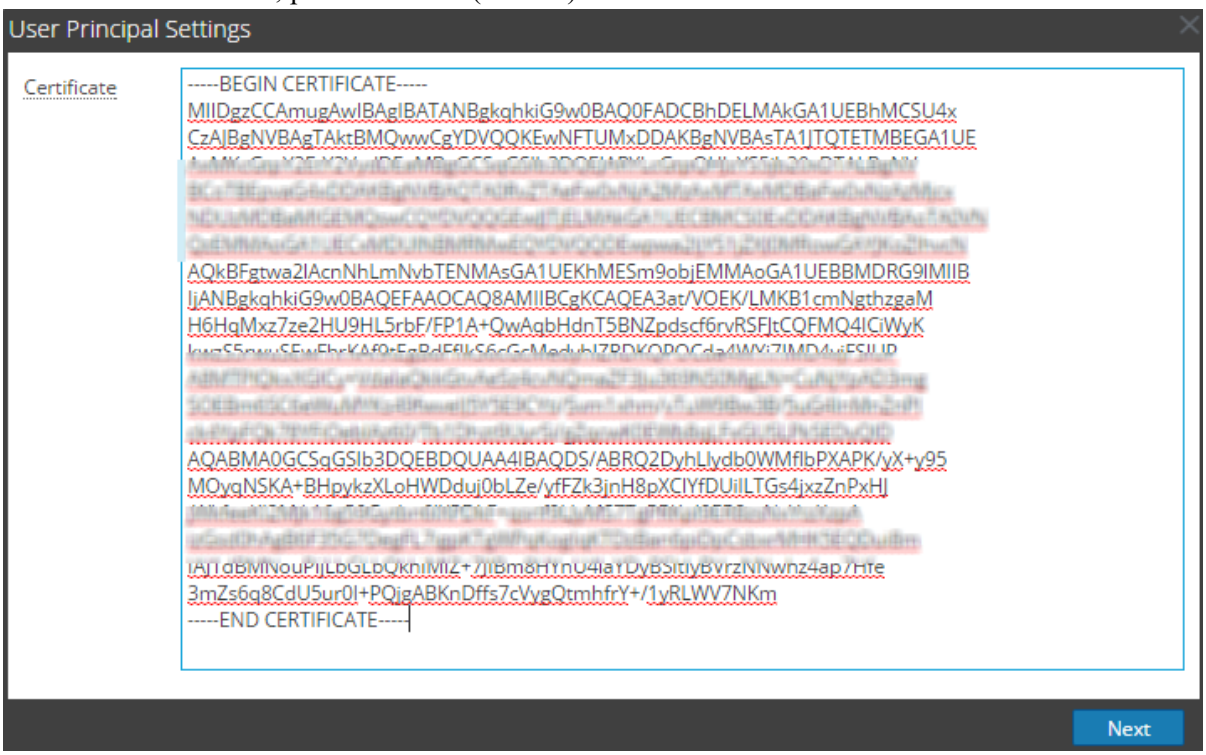
   - Select **Configure Revocation Checks Manually** to manually configure the CRL revocation check.
     For more information, see (Optional) Configure the CRL Manually.

9. Click **Next**.

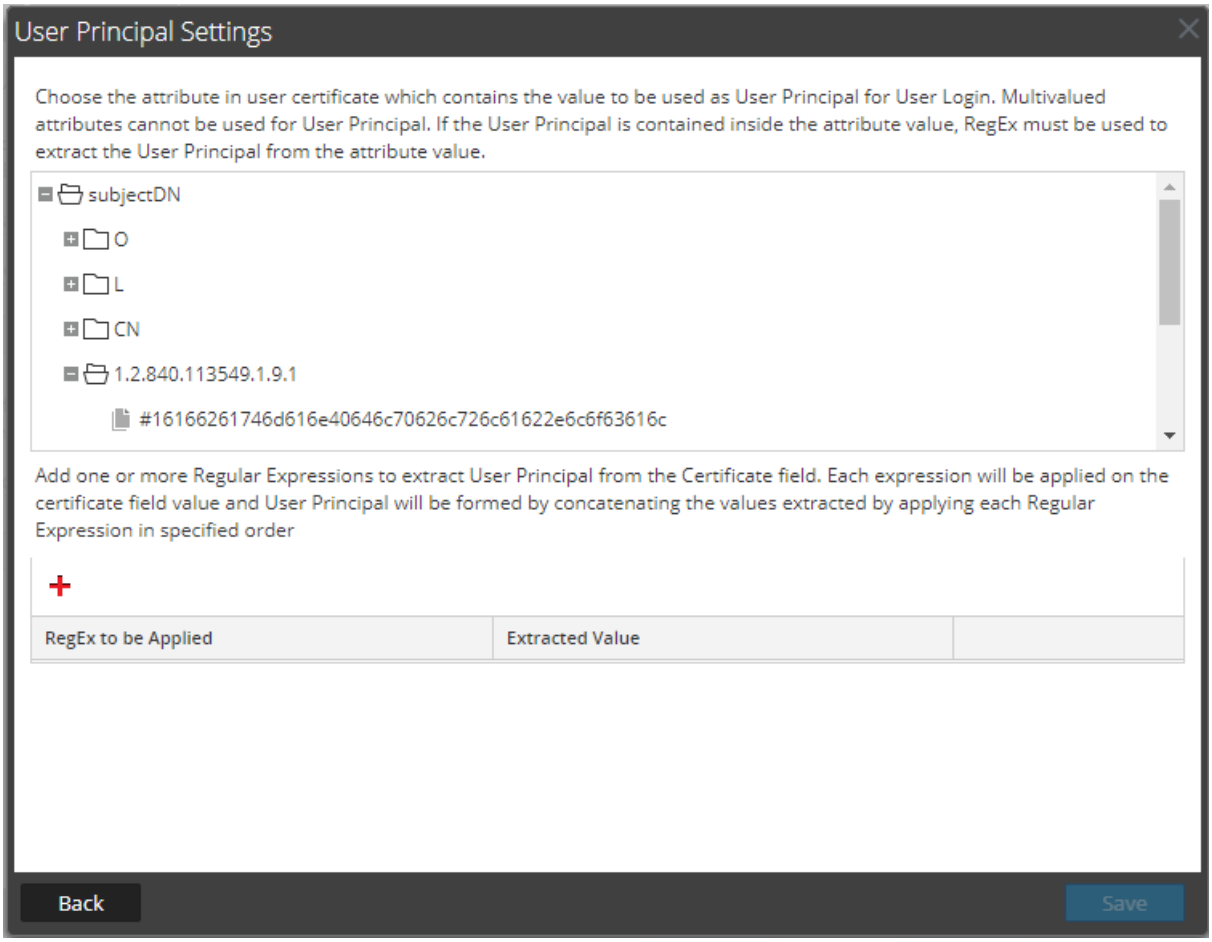
10. To configure user principal settings, in the User Principal Settings section, click **Configure Path and RegEx**.



11. In the **Certificate** field, paste the PEM (Base64) encoded user certificate.



12. Click **Next**.
    The **subjectDN**, **subjectAltnames**, and **extenstions** attributes are displayed.

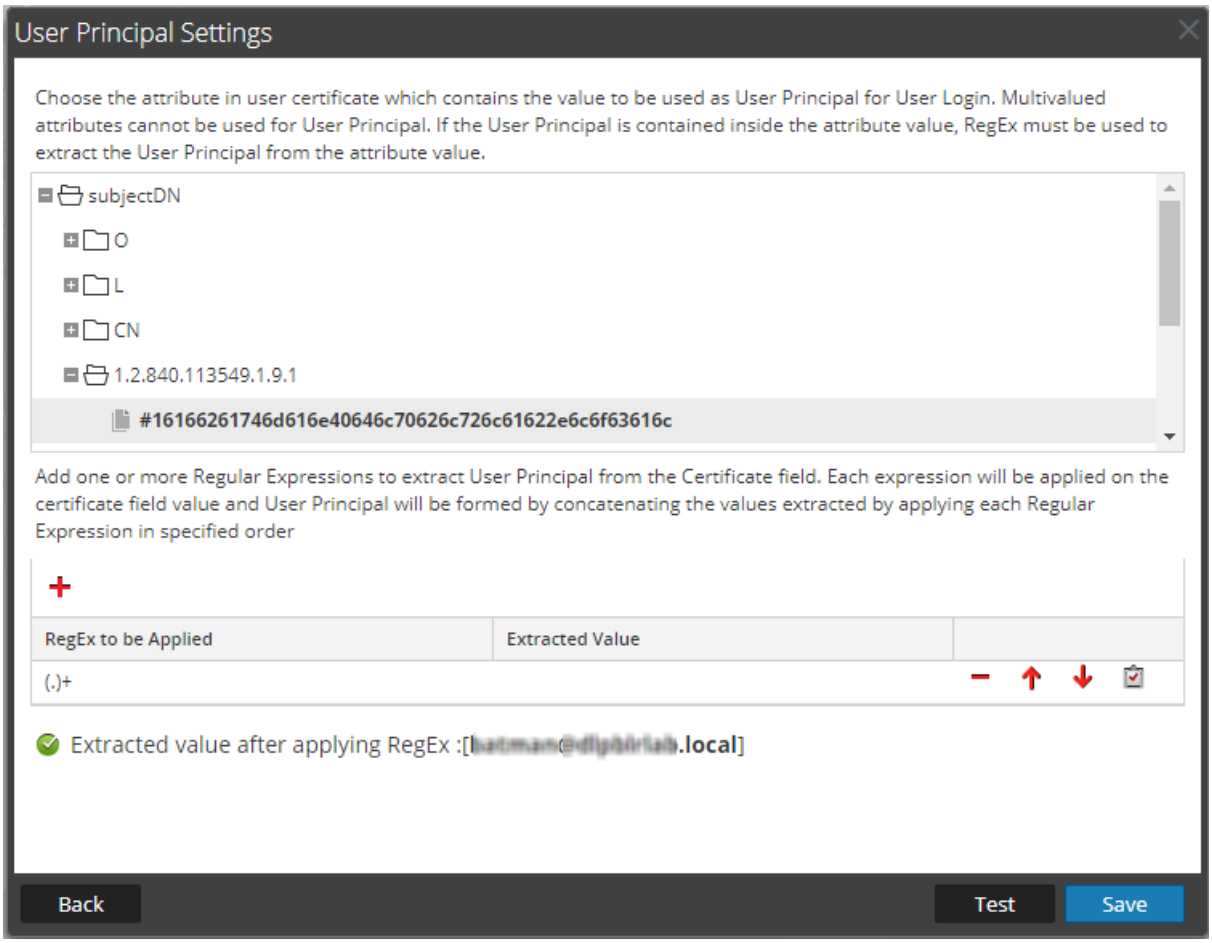


13. Select a unique attribute value to extract the username or user id.

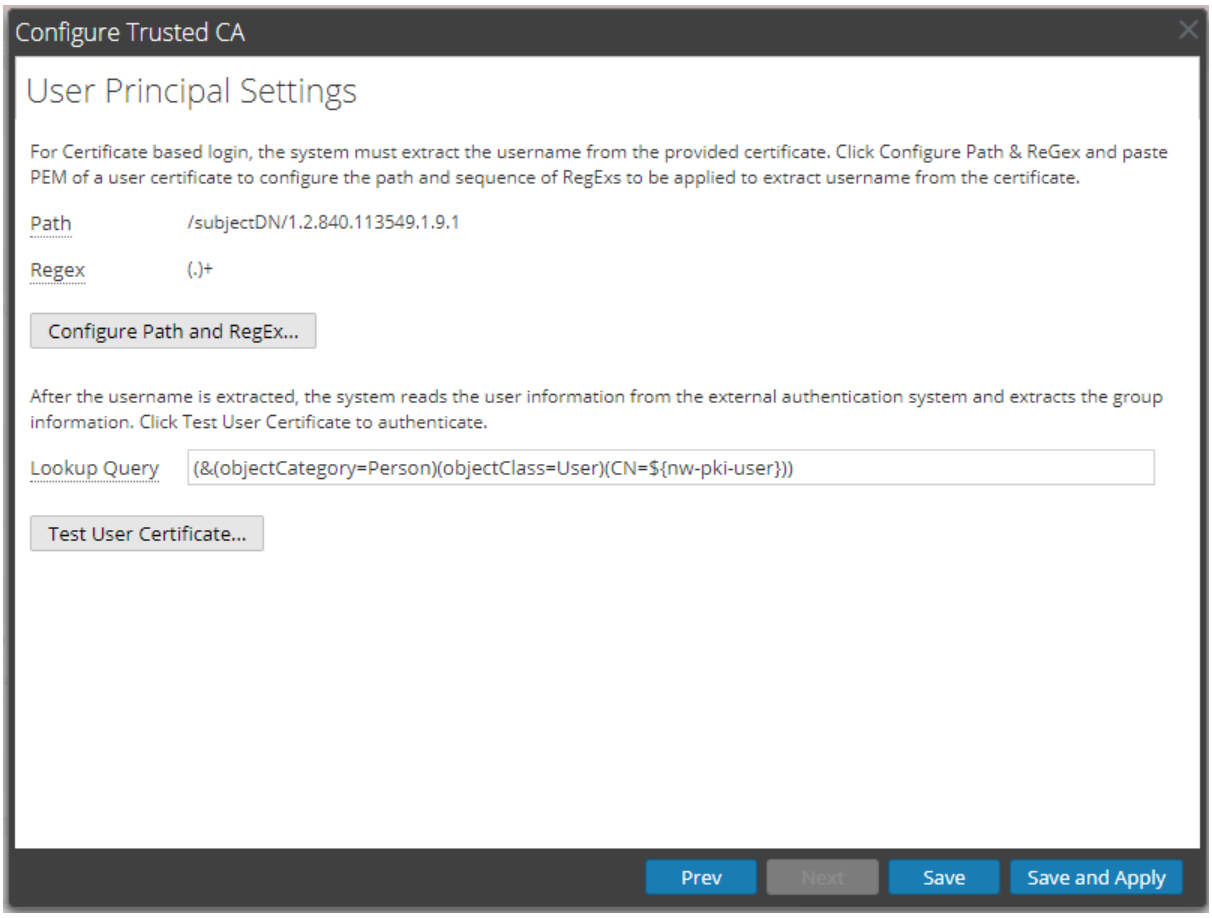
14. To apply RegEx, click .

> **Note:** You can apply multiple RegEx to extract a username or user id. All of the extracted values are concatenated to generate the final username or user id.

15. Click **Test**.
    The final user name or user id after applying RegEx is extracted and displayed.

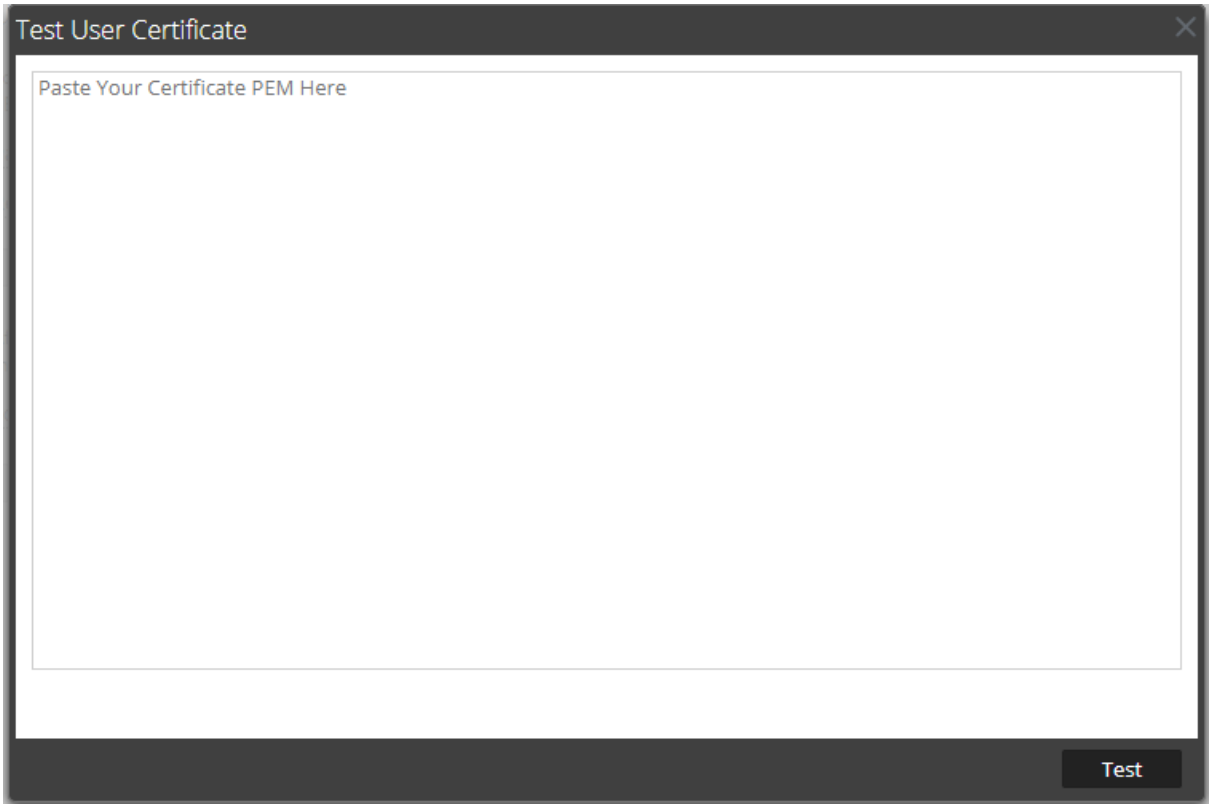

User Principal Settings
Choose the attribute in user certificate which contains the value to be used as User Principal for User Login. Multivalued attributes cannot be used for User Principal. If the User Principal is contained inside the attribute value, RegEx must be used to extract the User Principal from the attribute value.
subjectDN
O
L
CN
1.2.840.113549.1.9.1
#16166261746d616e40646c70626c726c61622e6c6f63616c
Add one or more Regular Expressions to extract User Principal from the Certificate field. Each expression will be applied on the certificate field value and User Principal will be formed by concatenating the values extracted by applying each Regular Expression in specified order
RegEx to be Applied
Extracted Value
(.)+
Extracted value after applying RegEx :[.local]
Back
Test
Save

16. Click **Save**.



17. Enter the query in the **Lookup Query** field to query the external authentication system for retrieving the user objects.

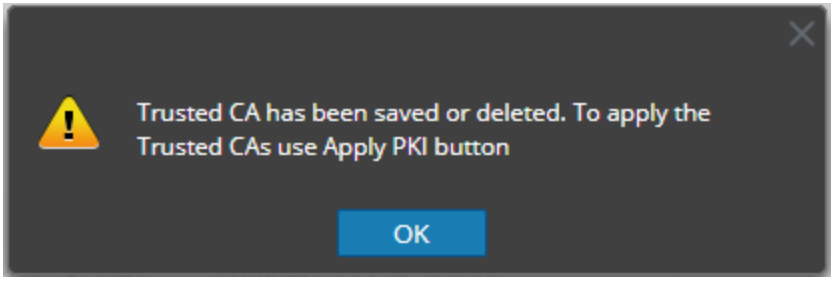
18. Click **Test User Certificate**.
    The Test User Certificate dialog is displayed.



19. Paste the PEM of the user certificate and click **Test**.

Note: The user Certificate is used by NetWitness for a dry-run of the trusted CA configuration. If the user certificate validation and the user id extraction from the certificate is successful, a confirmation message is displayed.

20. After the certificate is validated, click **Save**.
    The following message is displayed.



21. Click **OK**.
    The Trusted CA certificate is added to the NetWitness.

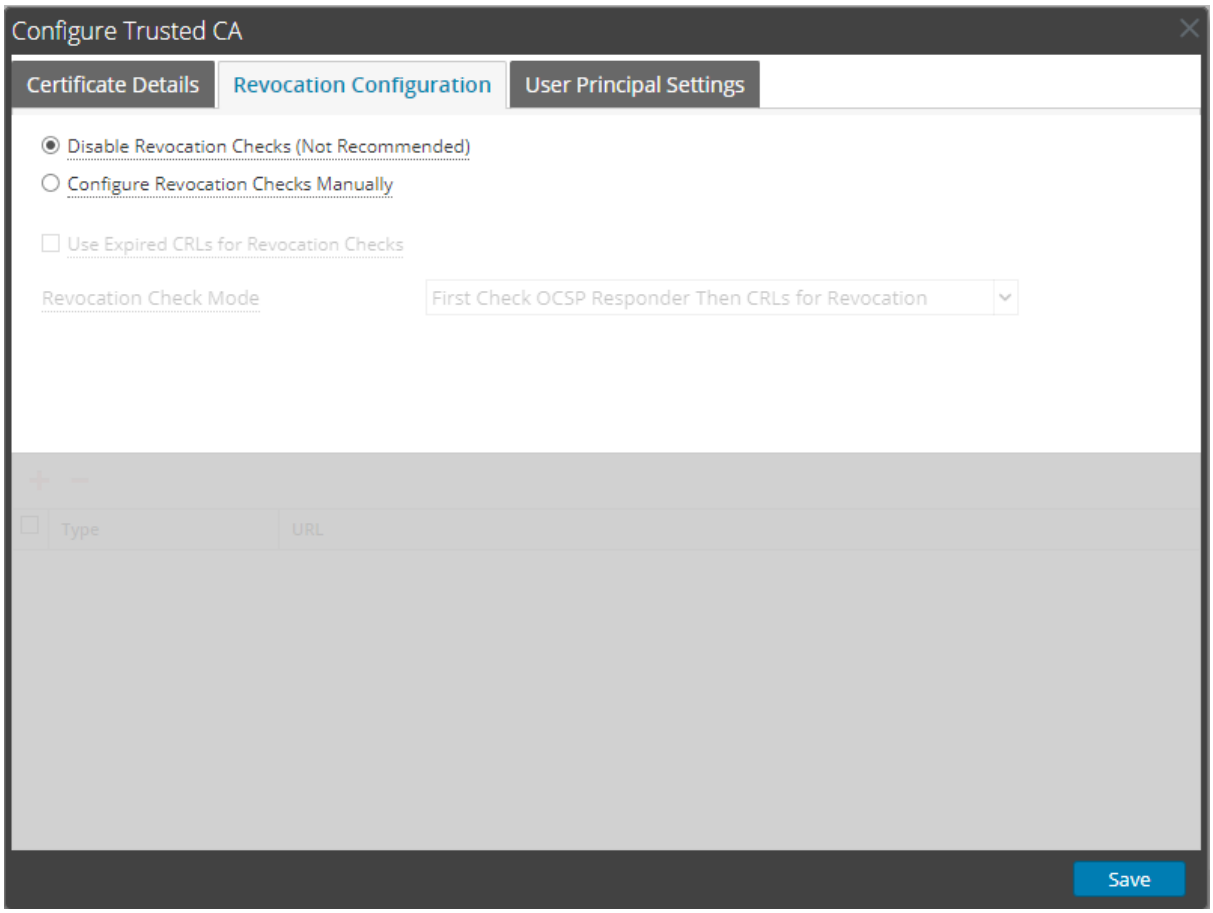
Note: You can add multiple Trusted CA certificates.

(Optional) Configure the CRL Manually

To configure the CRL manually:

1. Double click on the imported CA certificate.

   The Configure Trusted CA dialog is displayed.



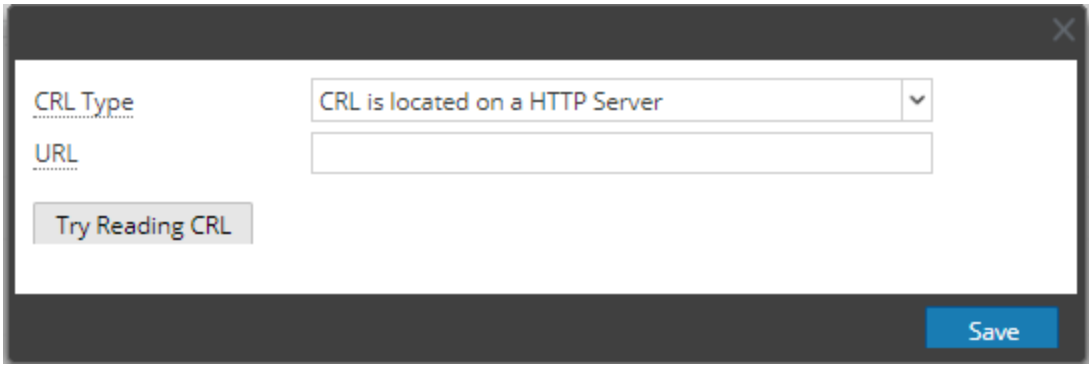
2. In the Revocation Configuration section, select **Configure Revocation Checks Manually**.
3. If the latest CRL is not available, select **Use Expired CRLs for Revocation Checks** to use the expired CRL for revocation.

   Caution: If the above option is enabled, the first CRL in the sequence will be used. This option is useful when you want the NetWitness to work even if PKI system is not available. Make note that first CRL in the sequence is always valid and it will not expire.

4. In the **Revocation Check Mode** field, do one of the following to validate the user certificate.
   - Select **Check only CRLs for Revocation** to use only the CRLs.
   - Select **Check only OCSP Responder for Revocation** to use only the OCSP Responders.
   - Select **First Check CRLs then OCSP Responder for Revocation** to use the CRL. If all the CRLs are expired, use the OCSP Responders.

- Select **First Check OCSP Responder Then CRLs for Revocation** to use the OCSP Responders. If all the Responders are offline or unavailable, use the CRLs.

5. Click  to add the CRL.



6. To add a CRL published on a HTTP server:
   a. In the **CRL Type** field, select **CRL is located on a HTTP server**.
   b. In the **URL** field, specify the HTTP URL to access the CRL

Note: Make sure that the CRL is available and HTTP server is accessible from NetWitness.

7. To upload a CRL file downloaded from the CA:
   a. In the **CRL Type** field, select **CRL is available as a File**.
   b. In the **CRL file** field, click **Browse** to upload the CRL file.

Note: Make sure that the CRL is downloaded from CDP location.

8. To add a OCSP Responder:
   a. In the **CRL Type** field, select **HTTP URL for OCSP Responder.**
   b. In the **URL** field, specify the HTTP URL.
   c. In the **Certificate** field, click **Browse** to upload the OCSP Responder Signing Certificate.

Note: Make sure that the OCSP Responder is accessible from NetWitness.

9. Click **Try Reading CRL**.

   The NetWitness UI displays the extracted information from the CRL.

Note: The CRL revocation check is done in the sequence that the CRL is added.
For example:
- If there are two CRLs configured and both are valid, only the first CRL is considered for revocation. The second CRL is considered for revocation only after the first CRL expires.
- If there are two CRLs configured, if the first CRL is expired and you select **Use Expired CRLs for Revocation Checks**, the first CRL is only considered for revocation check and second CRL is ignored.

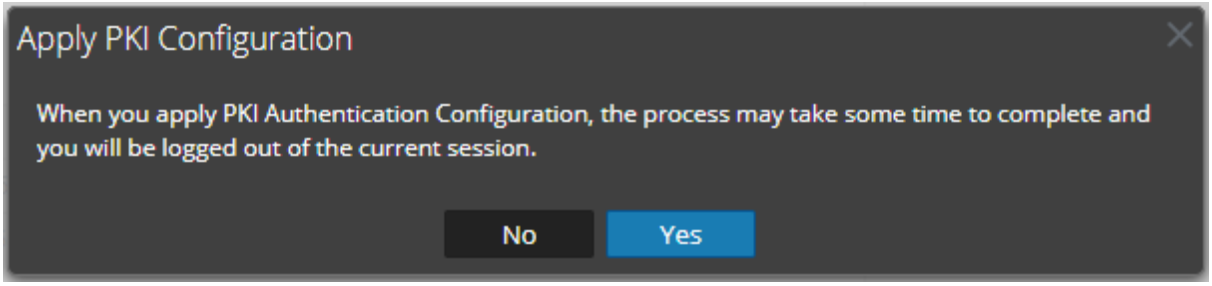
If the HTTP URL is located on the HTTPS location, the NetWitness does not validate the web server certificate of the HTTP server on which the CRL is located.

10. Click **Save**.

    The CRL file is added to the NetWitness.

Enable PKI Authentication

1. Go to  **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **PKI Settings** tab.
3. In the PKI Based Authentication Status section, select **Enabled**.
4. Click **Apply PKI Configuration**.
   The PKI Based Authentication Enabled dialog is displayed.



5. Click **Yes**.

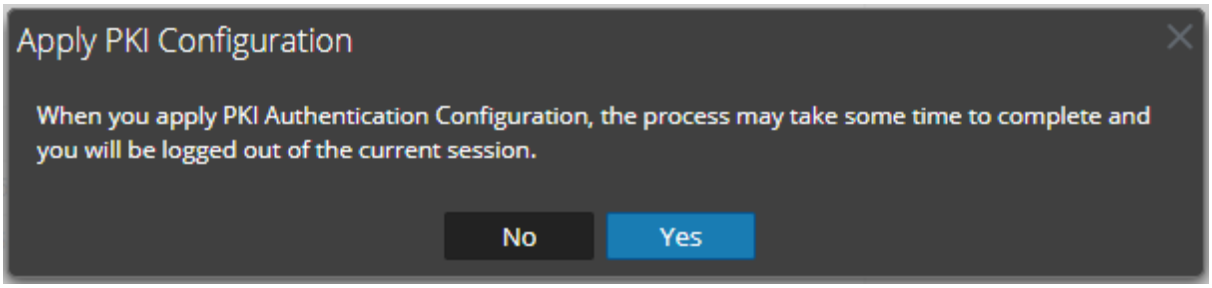
Note: After you enable PKI,
- Do not delete the AD configuration and external group mapping that corresponds to the user certificate's domain.
- Log out from a PKI based session by closing the browser used to access NetWitness.
- If audit log is enabled, the user login and activity is logged using the user DN.

Disable PKI

Note: If NetWitness users or Administrators are unable to access the NetWitness UI and want to use user name and password based authentication, you must disable PKI using the command line. See, Disable PKI using command line.

Disable PKI Authentication

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **PKI Settings** tab.
3. In the PKI Based Authentication Status section, select **Disabled**.
   The PKI Based Authentication Disabled dialog is displayed.



4. Click **Yes**.

Note: After disabling the PKI, wait for some time, close the browser, and open NetWitness in a new browser.

Disable PKI using command line

1. SSH to access root@node-0.
2. Run the following command:
   a. `nw-shell.`
   b. `login` {Enter the username and password to login}
   c. `connect --service orchestration-server`
   d. `cd /rsa/orchestration/userpki/disable-pki-on-hosts`
   e. `invoke`
   f. `connect --service admin-server`
   g. `cd /rsa/security/authentication/web/pki-enabled`
   h. `set false`
   i. `cd/rsa/security/pki/client-auth`
   j. `set WANT`
   k. `exit`
   l. `systemctl reset rsa-nw-admin-server`
   m. `systemctl reset-nw-security-server`

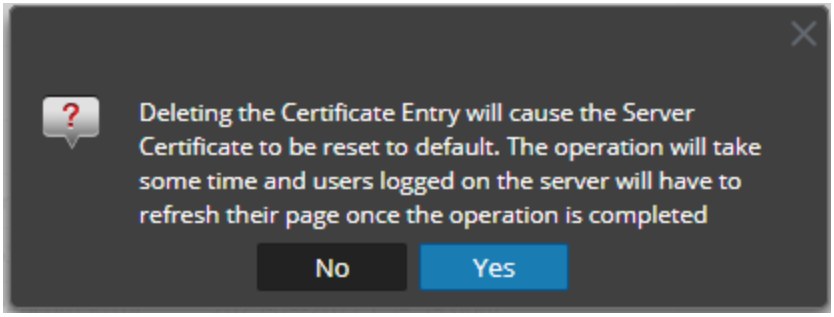
Note: The command `systemctl restart rsa-nw-admin-server` restarts the admin server and `systemctl restart rsa-nw-security-server` restarts the security server.

Delete Server Certificate and Trusted CA Certificate

Note: Deleting the server certificate may cause NetWitness to apply the default server certificate on the selected appliance.

To delete a NetWitness Server Certificate with its Private Key:

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **PKI Settings** tab.
3. In the **Server Certificates** section, select the certificate to delete.
4. Click .
   The following message is displayed.

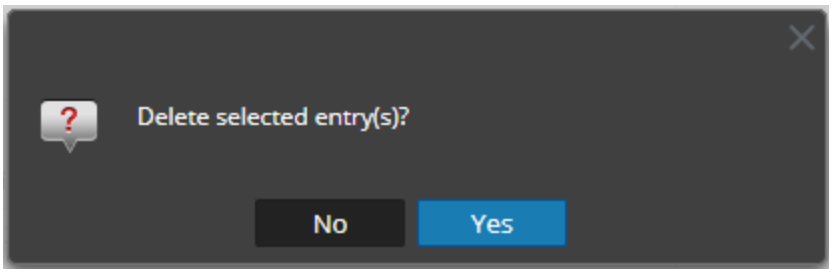


5. Click **Yes**.

Note: When the certificate is being applied on the selected appliance, no other operation on PKI can be performed until the process is completed.

To delete a Trusted CA Certificate:

1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **PKI Settings** tab.
3. In the **Trusted CAs** section, select the certificate to delete.
4. Click .
   The following message is displayed.



5. Click **Yes**.

Note: Make sure you apply the PKI configuration again after deleting the Trusted CA certificate.

Troubleshooting

This topic provides information about possible issues that NetWitness users may encounter when configuring the System Security and User Management settings in NetWitness. You can look up explanations of issues and their solutions.

Users are able to create a password of 8-chracters or less despite the configured minimum password length of 9 characters in Version 11.3

Problem	Solutions
When NetWitness was upgraded from 11.2 and previous versions to Version 11.3, the administrator did not set the minimum password length to 9 characters.	In 11.2 and earlier versions, the minimum password length is 8. The minimum password length changed to 9, in Versions 11.3. If you upgrade or update from earlier versions to 11.3, users can still create a password of 8 characters until you explicitly set the minimum password length to 9 characters as described in Configure Password Complexity.

Unable to log in to NetWitness Platform using SSO

Problem	Solutions
When the Administrator configures the SSO incorrectly and is unable to log in to NetWitness.	**Manual Steps to Disable SSO** To resolve this issue you must disable SSO manually, using the following commands: 1. SSH to `admin server node`. 2. Connect to `nw-shell`. 3. Connect to admin server service using the `connect --service admin-server` command. 4. Log in to admin server using the `login` command. 5. Enter the admin username and password. 6. Execute the following commands: • `cd /rsa/security/authentication/web/saml/sso-enabled` • `set false` • `logout` • `exit` • `systemctl restart rsa-nw-admin-server`



Problem	Solutions
Unable to connect to IDP and request session has timed out	• Check if the admin server is able to reach the specific IDP metadata URL. • Check if the IDP can be to accessed over the internet, if not configure the proxy and try again.
SSL handshake failed as the certificate is not verified	• Enable the `trust-all-certs-for-idp-metadata` flag in the explorer view of admin-server by navigating to RSA>Security>Authentication>Web>SAML. • Import the SSL certificate of the IDP metadata server to the JVM trust store, run the command `keytool -import -trustcacerts -file /root/selfsignedadfs.cer -alias selfsignedcert -keystore /usr/lib/jvm/java-1.8.0-openjdk-1.8.0.222.b10-0.el7_6.x86_64/jre/lib/security/cacerts` on the Admin node.
SSL handshake failed as the hostname is not verified	• Check the IDP metadata server's SSL certificate has a valid DN and matches the server hostname. • Enable the `trust-all-certs-for-idp-metadata` flag in the explorer view of admin-server by navigating to RSA>Security>Authentication>Web>SAML.
Fail over IP address changed	Perform the following manual steps to configure the new IP address. 1. Disable SSO using nw-shell after failover from new IP. For more information, see Manual Steps to Disable SSO 2. Generate the new metadata and reupload it in ADFS. For more information, see see the *Configure SAML 2.0 provider settings for portals* topic in Microsoft documentation.

References

This section is a collection of references for system security and user management in NetWitness.

- Admin Security View
- Users Tab
- Add or Edit User Dialog
- Roles Tab
- Add or Edit Role Dialog
- External Group Mapping Tab
- Add Role Mapping Dialog
- Search External Groups Dialog
- Settings Tab
- PKI Settings Tab
- Login Banner Tab

Admin Security View

The Admin Security view provides the capability to manage user accounts, manage user roles, map external groups to NetWitness roles, and modify other security-related system parameters. These apply to the NetWitness system and are used in conjunction with the security settings for individual services.

What do you want to do?

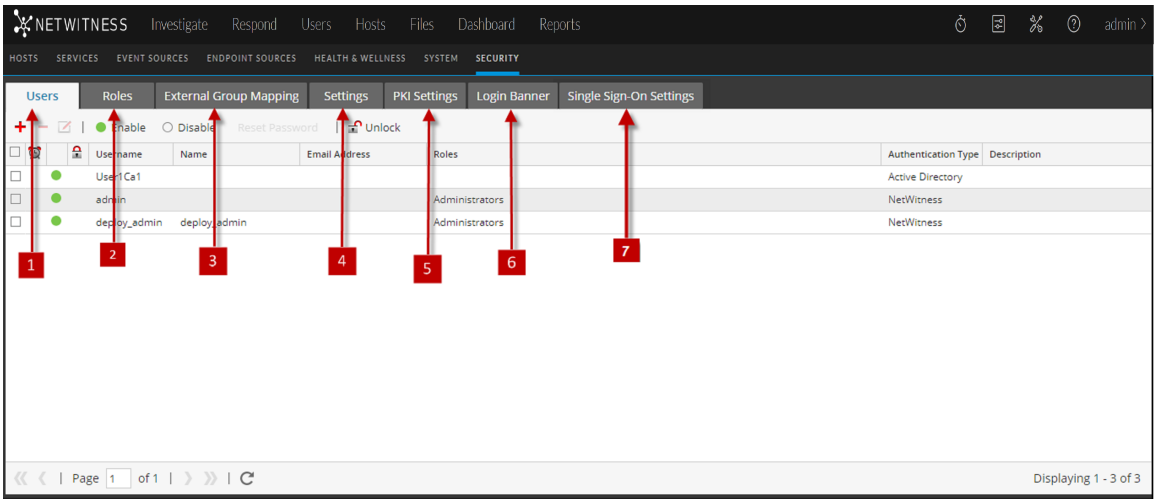
Role	I want to ...	Show me how
Admin	Manage users	Set Up Users
Admin	Manage roles	Review the Preconfigured NetWitness Platform Roles (Optional) Add a Role and Assign Permissions
Admin	Configure settings	Configure System-Level Security Settings
Admin	(Optional) Configure external group mappings	(Optional) Map User Roles to External Groups
Admin	(Optional) Set login conditions	(Optional) Create a Customized Login Banner
Admin	(Optional) Set up PKI	Configure PKI Authentication

Related Topics

- Users Tab
- Roles Tab
- External Group Mapping Tab
- Settings Tab
- PKI Settings Tab
- Login Banner Tab

Quick Look

To display the Admin Security view, go to  **(Admin) > Security**.

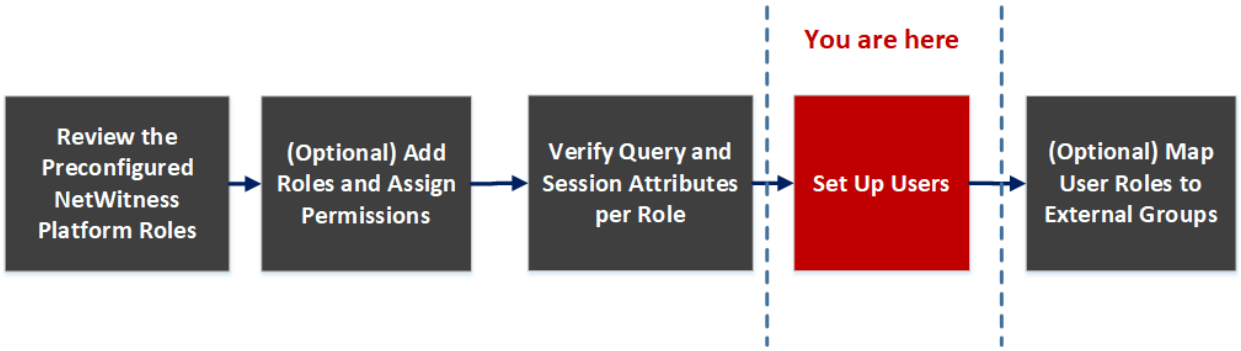


1. **Users** tab allows you to manage user accounts.
2. **Roles** tab allows you to define security roles and assign roles to user accounts.
3. **External Group Mapping** tab allows you to manage access parameters for LDAP groups.
4. **Settings** tab allows you to configure password complexity and expiration for internal NetWitness users and to configure system behavior due to failed logins and inactivity. It also allows you to configure external authentication.
5. **PKI Settings** allows you to enable and disable the PKI authentication.
6. **Login Banner** tab allows you to set conditions which must be agreed to before gaining access to the login screen.
7. **Single Sign-On Settings** tab allows you to enable and disable the SSO authentication.

Users Tab

The Users tab in the Admin Security view allows you to set up user accounts. Each NetWitness user must have a user account. In the Users tab, you can create, edit, delete, enable, disable, and unlock a user account.

Workflow



What do you want to do?

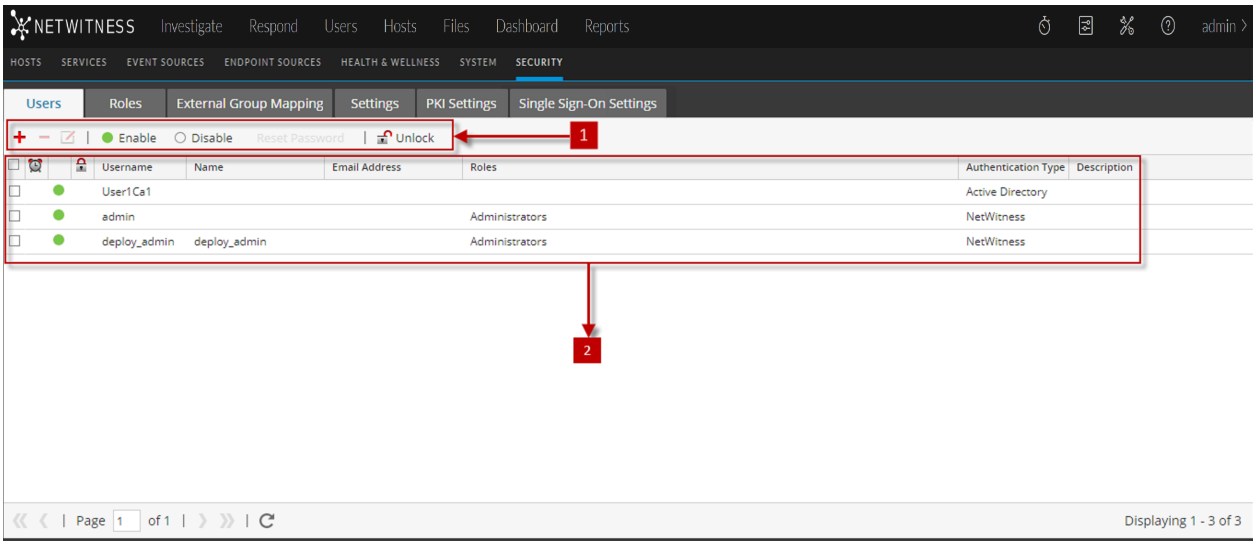
Role	I want to ...	Show me how
Admin	Set up a new user	Set Up Users
Admin	Manage user accounts	

Related Topics

- Add or Edit User Dialog

Quick Look

To access this view, go to  **(Admin) > Security**. The Security view displayed with **Users** tab open.



The Users tab includes the following panels.

1 **Users** toolbar
2 **Users** list

Users Toolbar

The following table describes the Users toolbar.

Feature	Description
+	Opens the Add User dialog.
−	Deletes the selected user.
	Opens the Edit User dialog for the selected user.
Enable	Enables a disabled user account with all user preferences intact.
Disable	Blocks user access without deleting user preferences so that upon re-enabling users, user preferences are intact.
Reset Password	Opens the Reset Password dialog, which enables you to change the password of the selected user. This dialog lists the password format requirements necessary to change the password and allows you to force the user to change their password on the next login.
Unlock	Unlocks a user account that has been locked due to too many failed login attempts.

Users List

The following table describes the columns in the Users list.

Column	Description
	If this icon appears in a user row, it indicates that the user password has expired.
Username	Username to log on to NetWitness.
Name	Name of the user to whom the account belongs.
Email Address	Email address of the user.
Roles	Role assigned to the user.
Authentication Type	Authentication method, which could be external by Active Directory or PAM or internal by NetWitness.
Description	Description of the user account.

Add or Edit User Dialog

All users must either have a local user account with a username and password or an external user account that is mapped to NetWitness.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Add a User and Assign a Role	
Administrator	Change User Information	
Administrator	Reset a User Password	
Administrator	Add a User for External Authentication	

Related Topics

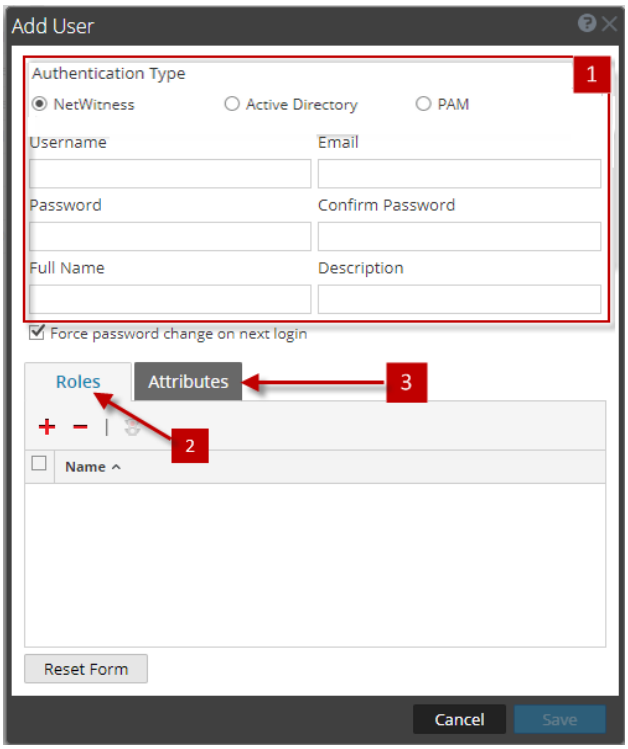
- Manage Users with Roles and Permissions

Quick Look

To access Add or Edit User dialog, go to  **(Admin) > Security > Users** tab and in the toolbar, click , or select a user and click .

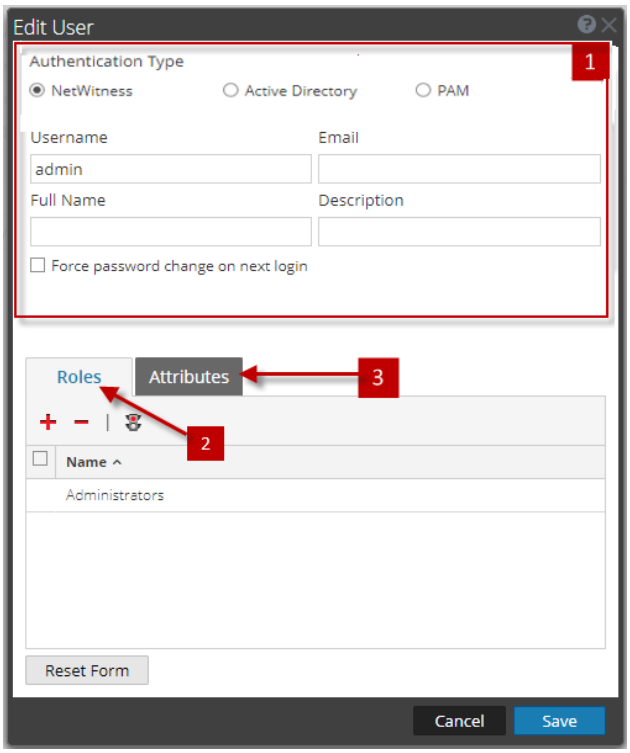
Add User Dialog

This is the Add User dialog for an internal user.



Edit User Dialog

This is the Edit User dialog for an internal user.



The Add User and Edit User dialogs are the same except that the Add User dialog contains additional **Password** and **Confirm Password** fields. You can add a password for a new user in the Add User dialog. Users can change their own passwords in the user preferences. You can reset a password for a user directly from the Users tab.

The Add User and Edit User dialogs have following sections:

1. User information
2. **Roles** tab
3. **Attributes** tab

User Information

The following table provides descriptions of the user information.

Field	Description
NetWitness	Authenticate with NetWitness.
Active Directory	Authenticate with Active Directory.
PAM	Authenticate with PAM.
Username	Username for the NetWitness user account.
Email	Email address of the user.
Password	(Add User dialog only) Password to log on to NetWitness.
Confirm Password	(Add User dialog only) Password confirmation for adding the user password.
Full Name	Name of the user.
Description	(Optional) Description of the user.
Force password change on next login	Expires the user password the next time the user logs on to NetWitness. This field applies only to internal users. This does not affect any active user sessions. The ⏰ appears in the user row to show that the user password expired. After a password is expired, you cannot undo it. This checkbox is cleared the next time you edit the user account.
Reset Form	Removes any changes in process.

Roles Tab

The following table provides descriptions of the **Roles** tab options. The Roles tab shows the roles that are assigned to the user.

Option	Description
+	Opens the Add Role dialog that lists roles you could assign to the user.

Option	Description
	Removes the selected role from being assigned to the user.
	Shows permissions for the selected role.
Name	Lists each role assigned to the user.

Attributes Tab

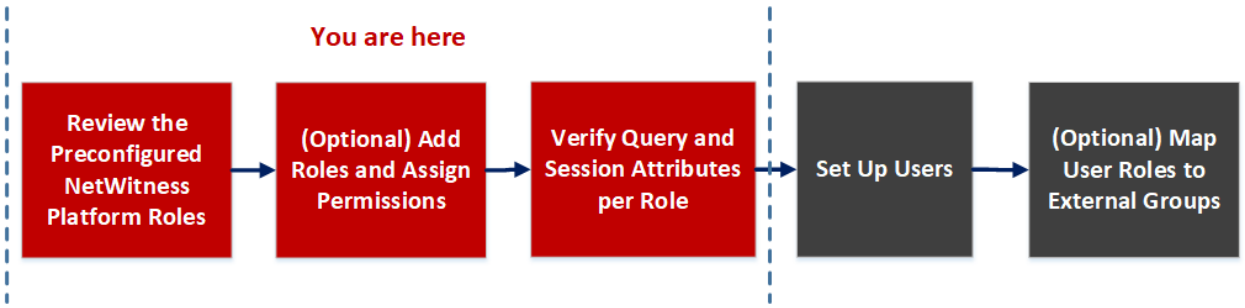
The following table lists the descriptions of the **Attributes** tab options. The **Attributes** tab shows the attributes that are assigned to the user.

Option	Description
Core Query Timeout	(Optional) Specifies the maximum number of minutes that a user can run a query. The default value is 5 minutes. This timeout only applies to queries performed from Investigation. If this value is set, it must be zero (0) or greater. A value of zero represents no timeout.
Core Session Threshold	Controls how the service scans meta values to determine session counts. This value must be zero (0) or greater. If this value is greater than zero, a query optimization will extrapolate the total session counts that exceed the threshold. When the meta value returned by the query reaches the threshold, the system will: • Stop its determination of the session count • Show the threshold and percentage of query time used to reach the threshold The default value is `100000`. The limit you specify here overrides the **Max Session Export** value defined in the Investigate view settings.
Core Query Prefix	(Optional) Filters query results to restrict what the role members see. By default, this is blank. For example, the `'service' = 80` query prefix prepends to any queries run by the user and the user can only access meta of HTTP sessions.

Roles Tab

Roles are assigned to all NetWitness users. Users receive the permissions the roles allow. In the Roles tab you can create, duplicate, edit and delete a role. You can also see a list of all roles and their respective permissions.

Workflow



Related Topics

- Add or Edit Role Dialog

What do you want to do?

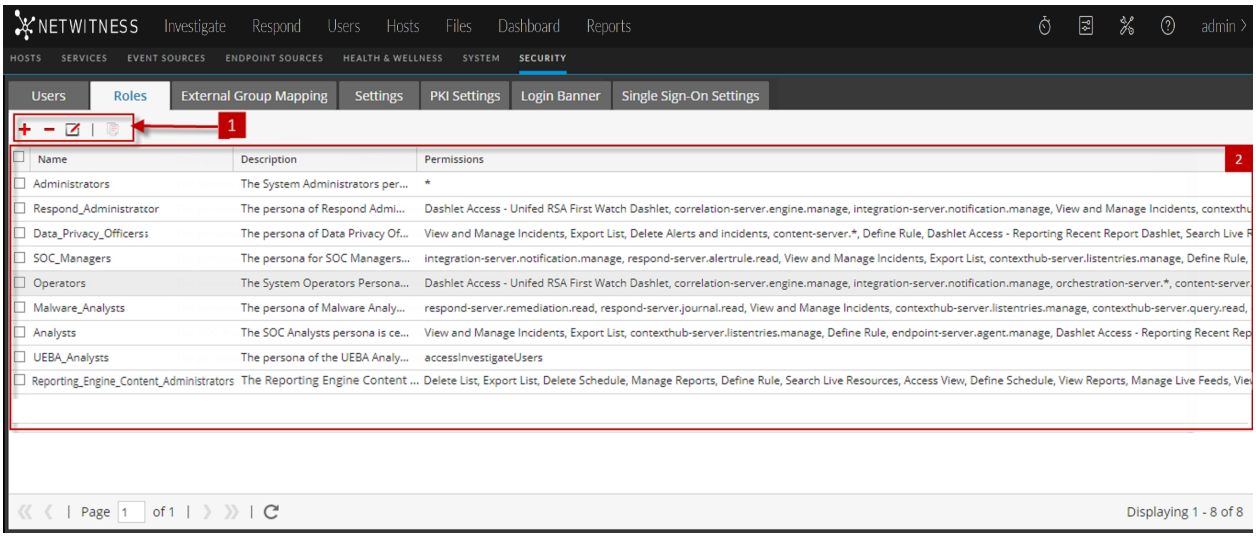
Role	I want to ...	Show me how
Admin	View preconfigured roles	Review the Preconfigured NetWitness Platform Roles
Admin	Create a new role	(Optional) Add a Role and Assign Permissions

Related Topics

- Add or Edit Role Dialog

Quick Look

To access **Roles** tab view, go to  **(Admin) > Security** and click the **Roles** tab.



The Roles tab has the following sections.

1 **Roles** toolbar
2 **Roles** list

Roles Toolbar

The following table describes the **Roles** toolbar.

Icon	Description
	Displays the Add Role dialog.
	Displays the Edit Role dialog.
	Displays a warning message, and asks for confirmation that you want to delete a role.
	Duplicates a role to save with a different name.

Roles list

The following table describes the columns in the **Roles** list.

Column	Description
Name	Displays the name of a role that can be given to a user.
Description	Displays a description of the role.
Permissions	Displays the permissions assigned to the role.

Add or Edit Role Dialog

In the Add Role and Edit Role dialogs, you can add or edit a role and the permissions assigned to it. You can also specify the query-handling attributes for role members to lock down the information that they can retrieve. The structure of these dialogs is the same. The only difference is that you either add a new role or modify an existing role.

When you change permissions for a role, the change is immediately applied to users who are assigned the particular role after the role is saved.

What do you want to do?

Role	I want to ...	Show me how
Admin	View preconfigured roles	Review the Preconfigured NetWitness Platform Roles
Admin	Create a new role	(Optional) Add a Role and Assign Permissions
Admin	Edit a role	(Optional) Add a Role and Assign Permissions
Admin	Delete a role	(Optional) Add a Role and Assign Permissions

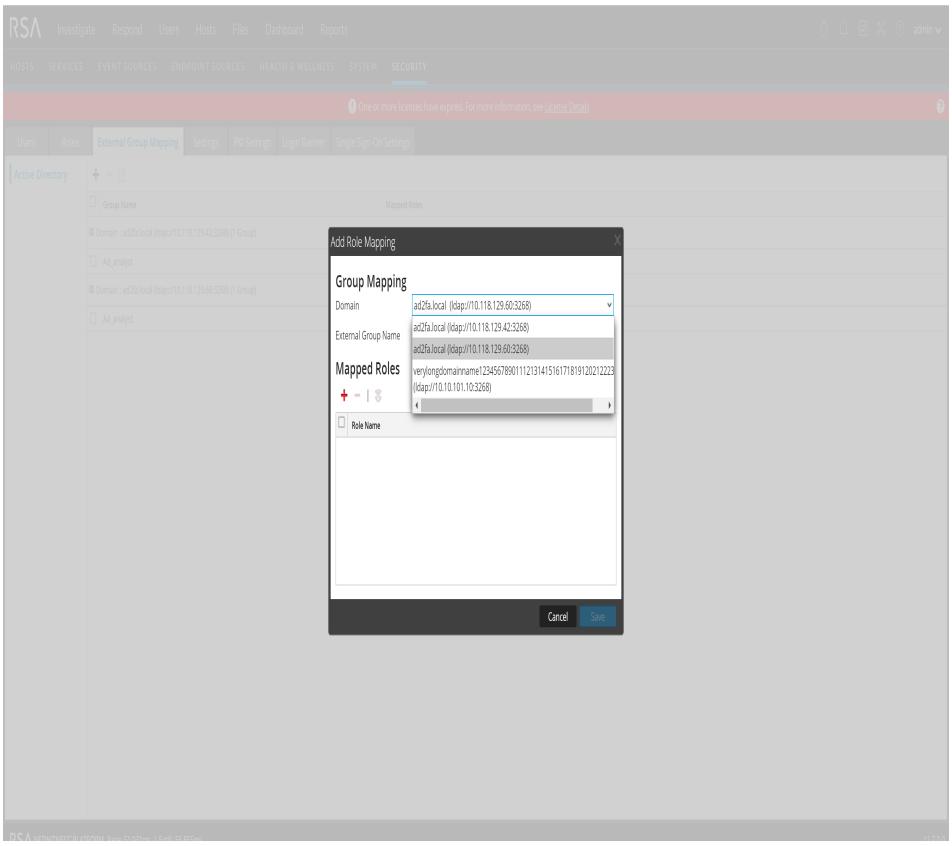
Related Topics

Verify Query and Session Attributes per Role

Role Permissions

Quick Look

To access this view, go to  **(Admin) > Security > Roles** tab and in the toolbar, click , or select a role and click .



The Add Role and Edit Role dialogs include three sections.

1. **Role info**
2. **Attributes**
3. **Permissions**

Role Info

This is the information in the **Role Info** section.

Feature	Description
Name	The name of the user role.
Description	An optional description of the user role.

Attributes

The following table describes the fields in the **Attributes** section..

Field	Description
Core Query Timeout	(Optional) Specifies the maximum number of minutes that a user can run a query. The default value is 5 minutes. This timeout only applies to queries performed from Investigation. If this value is set, it must be zero (0) or greater. A value of zero represents no timeout.
Core Session Threshold	Controls how the service scans meta values to determine session counts. This value must be zero (0) or greater. If this value is greater than zero, a query optimization will extrapolate the total session counts that exceed the threshold. When the meta value returned by the query reaches the threshold, the system will: • Stop its determination of the session count • Show the threshold and percentage of query time used to reach the threshold The default value is `100000`. The limit you specify here overrides the **Max Session Export** value defined in the Investigate view settings.
Core Query Prefix	(Optional) Filters query results to restrict what the role members see. By default, this is blank. For example, the `'service' = 80` query prefix prepends to any queries run by the user and the user can only access meta of HTTP sessions.

Permissions

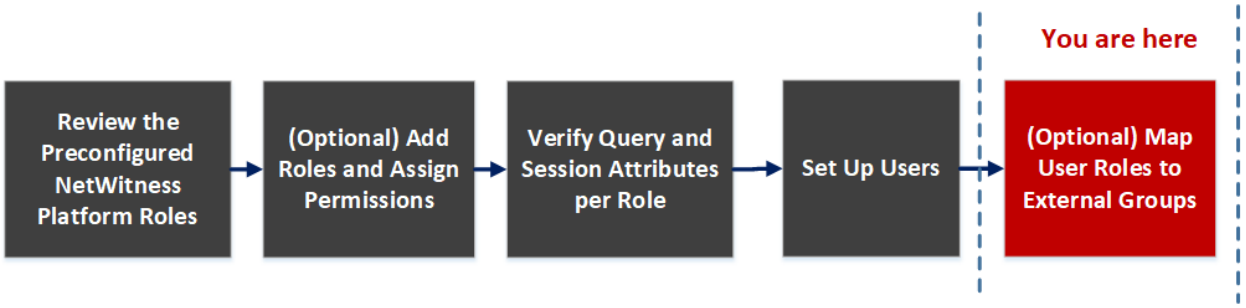
The following table describes the features in the **Permission** section.

Feature	Description
Module tabs	There are fifteen default tabs, one for each module: Administration, Admin-server, Alerting, Config-server, Incidents, Investigation, Investigation-server, Integration-server, Live, Malware, Orchestration-server, Reports, Response-server, Security-server and Dashboard. Additional tabs may be available based on the installation. Each tab lists the permissions for a module.
Assigned column	Checkbox that indicates if a module permission is assigned to the role.
Description column	List of all permissions for the module.
Save	Saves the role with the selected permissions assigned to it.
Cancel	Cancels any work and closes the dialog.

External Group Mapping Tab

If you set up external user authentication, you can map NetWitness user roles to an external group. The External Group Mapping tab provides information about each external group to which you have mapped roles.

Workflow



What do you want to do?

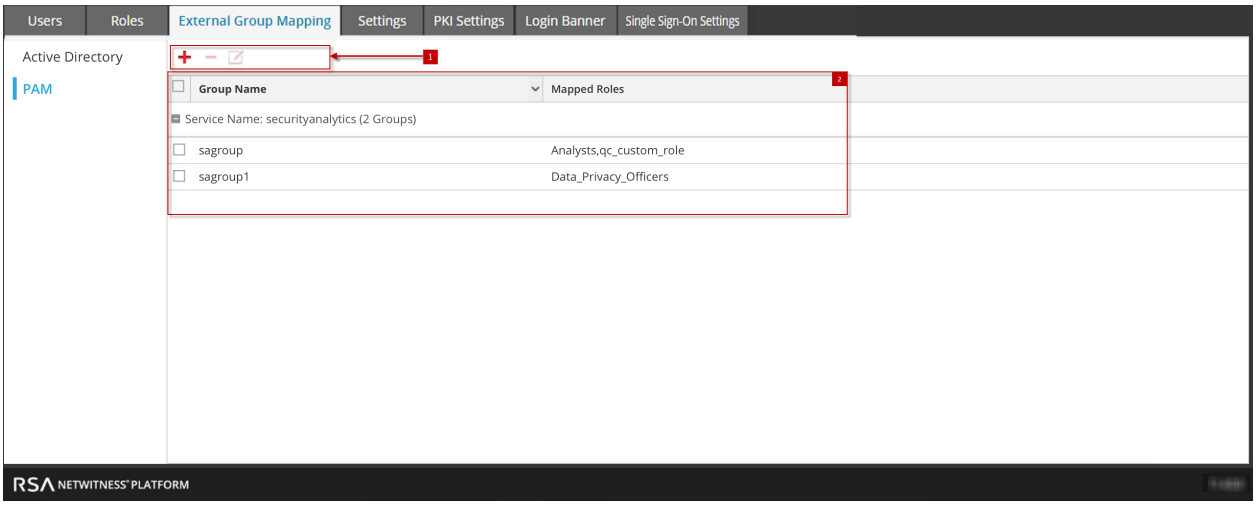
Role	I want to ...	Show me how
Admin	Map a role to an external group	(Optional) Map User Roles to External Groups
Admin	Search for an external group	Search for External Groups

Related Topics

- Add Role Mapping Dialog
- Search External Groups Dialog

Quick Look

To access this view, go to **(Admin) > Security** and click the **External Group Mapping** tab.



The External Group Mapping includes the following sections.

1 Toolbar
2 List

Toolbar

The following table describes the **External Group Mapping** toolbar.

Icon	Description
	Displays the Add Role Mapping dialog in which you can select an external group and map it to a NetWitness role.
	Displays a warning message and asks for confirmation to remove all NetWitness roles mapped to the external group.
	Displays the Edit Role Mapping dialog in which you can add or remove NetWitness roles from the external group.

List

The following table describes the **External Group Mapping** list.

Feature	Description
Group type	In the column on the left, click either **Active Directory** or **PAM** to show groups for the selected type.
Selection box	In a row, toggles selection of a group name. In the title bar, toggles selection of all group names.
Group Name	Displays the name of the external group that has access to NetWitness.
Mapped Roles	Displays the NetWitness roles mapped to the external group.

Add Role Mapping Dialog

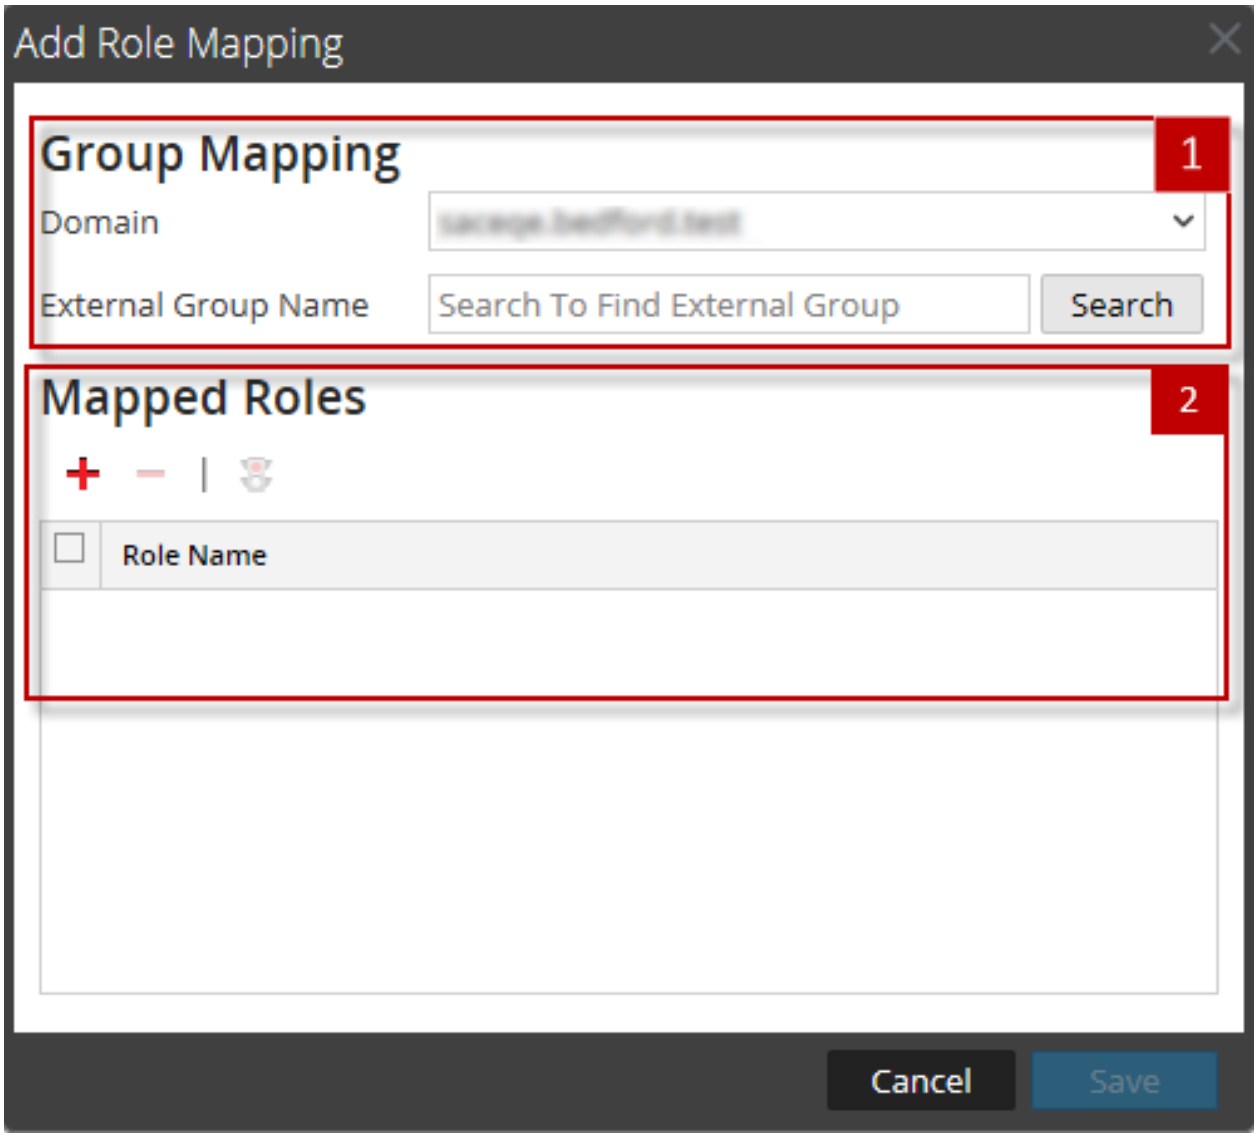
In NetWitness each user role has its own set of permissions. You can map one or more NetWitness roles to an external group, which grants the group the same set of permissions that each role has.

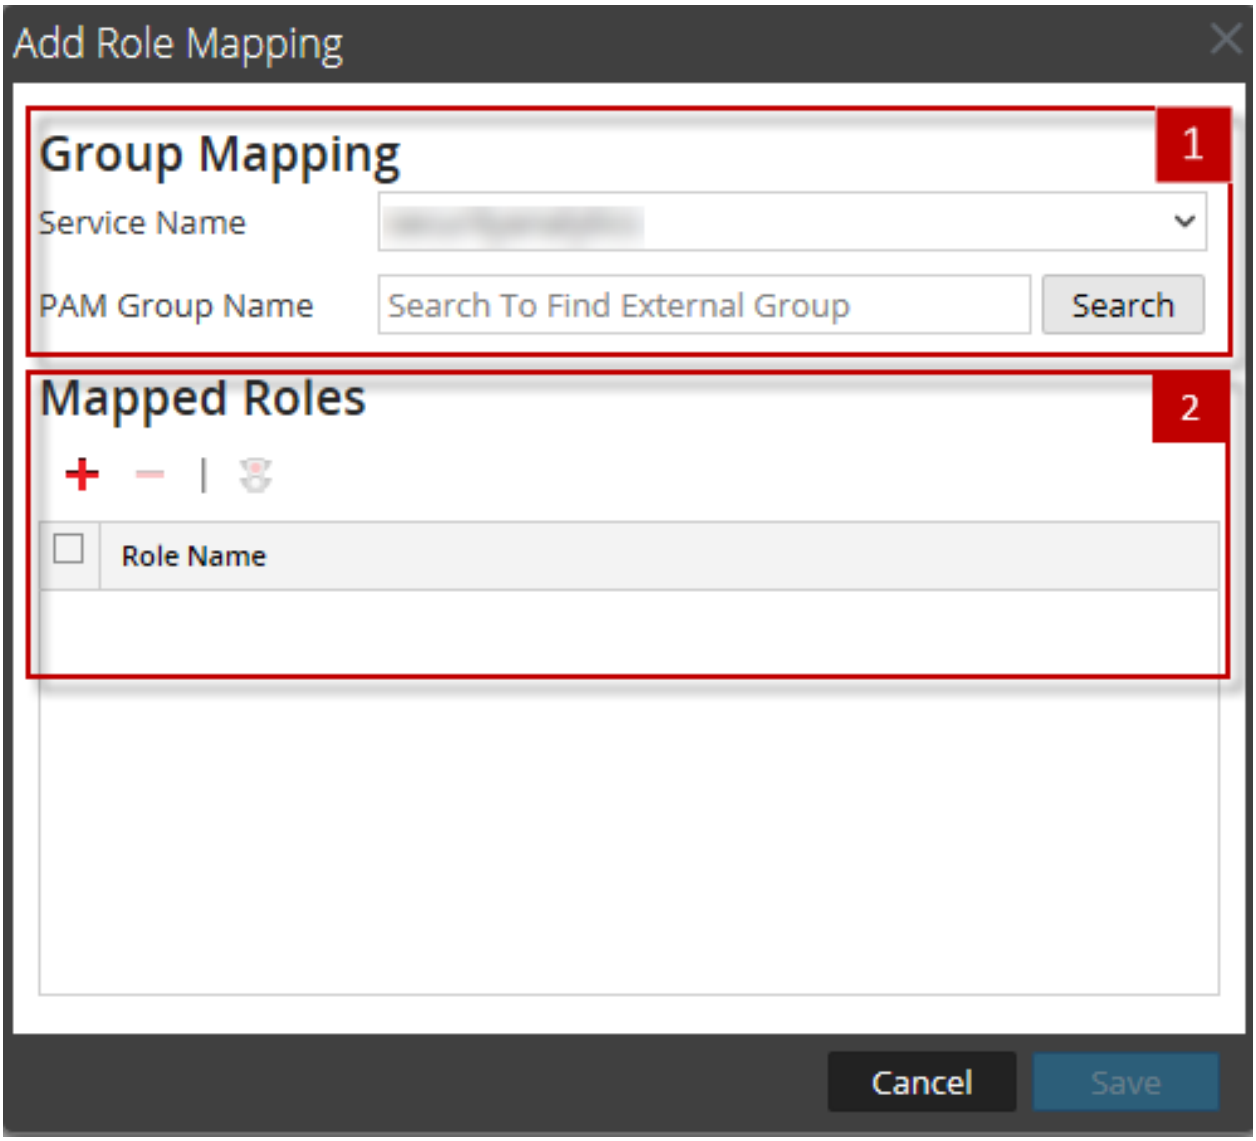
What do you want to do?

Role	I want to ...	Show me how
Admin	Map a role to an external group	(Optional) Map User Roles to External Groups
Admin	Search for an external group	Search for External Groups

Quick Look

To access this dialog, go to  **(Admin) > Security**, click the **External Group Mapping** tab, and in the toolbar, click +. The **Add Role Mapping** dialog for the external authentication method that you set up is displayed.


Add Role Mapping
Group Mapping
1
Domain
External Group Name
Search To Find External Group
Search
Mapped Roles
2
Role Name
Cancel
Save



Note: The Add Role Mapping and the Edit Role Mapping dialogs are nearly identical. The only difference is that you cannot search in the Edit Role Mapping dialog.

The **Add Role Mapping** tab includes the following sections.

1 **Group Mapping**
2 **Mapped Roles**

Group Mapping

The following table describes the **Group Mapping** section features.

Feature	Description
Domain	Displayed if you set up Active Directory for external user authentication. The domain name of the external AD group to which roles are mapped.
External Group Name	Displayed if you set up Active Directory for external user authentication. The external group to which roles are mapped.
PAM Group Name	Displayed if you configured PAM for external user authentication. The name of the external group to which roles are mapped.
Search	Displays a search dialog in which you can search for external groups using domain name and host. Search is not available in the Edit Role Mapping dialog.

Mapped Roles

The following table describes the **Mapped Roles** section features.

Feature	Description
	Opens the Add Role dialog, in which configured NetWitness user roles to add are listed.
	Removes selected roles from the Mapped Roles grid.
Name	Displays the name of the NetWitness user role.
Permissions	Displays the permissions associated with the NetWitness user role.
Cancel	Cancels the new group mapping or changed group mapping and closes the dialog.
Save	Saves the new group mapping or changed group mapping and closes the dialog.

Search External Groups Dialog

If you set up external user authentication, you can map NetWitness user roles to external groups. You search for external groups to select the groups to which you want to map NetWitness roles.

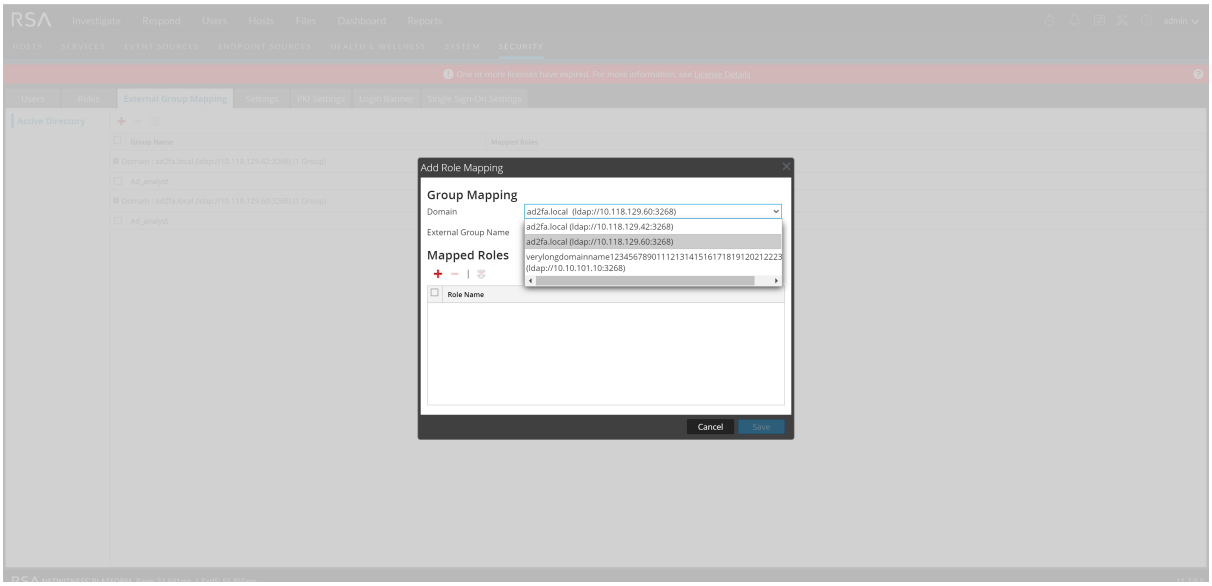
What do you want to do?

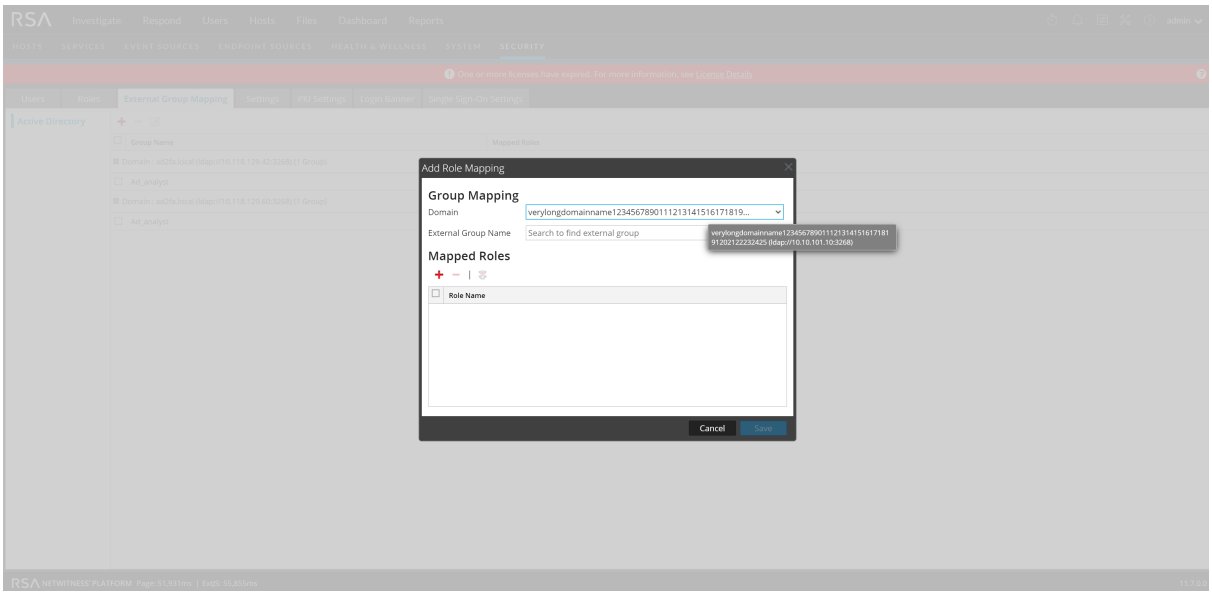
Role	I want to ...	Show me how
Admin	Map a role to an external group	(Optional) Map User Roles to External Groups
Admin	View external group mappings	External Group Mapping Tab
Admin	Search for external groups	Search for External Groups

Quick Look

To access this dialog:

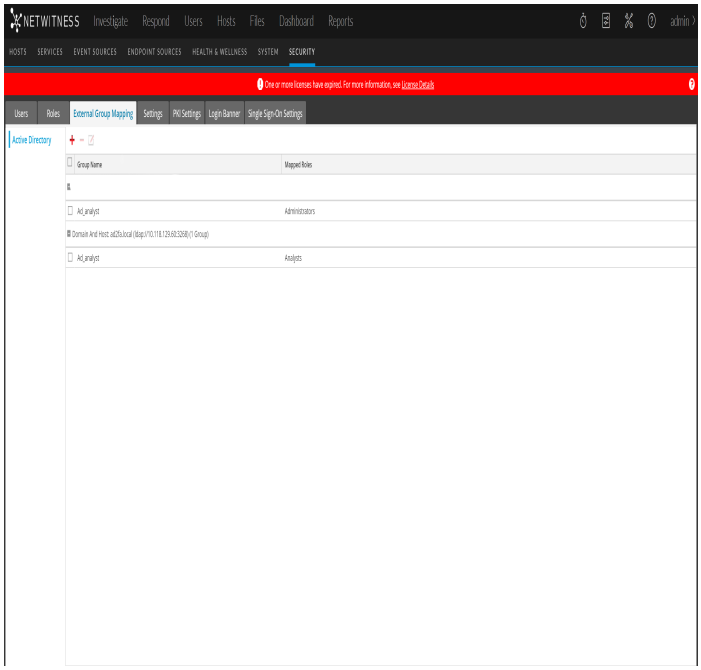
1. Go to **(Admin) > Security**.
   The Security view is displayed with the **Users** tab open.
2. Click the **External Group Mapping** tab.
3. In the toolbar, click +.
   The Add Role Mapping dialog for the external authentication method you set up is displayed.





4. In the Group Mapping section, select a **Domain**.

5. In the Group Mapping section, click **Search** next to **External Group Name** field.
   The **Search External Groups** dialog is displayed.



The following table describes the features of the **Search External Group** dialog.

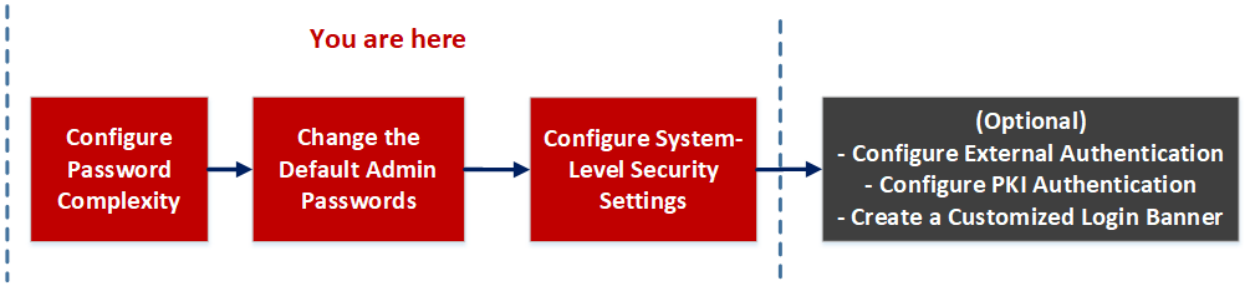
Feature	Description
Retrieve all groups or enter a search criteria	Group name for which you are searching. Can be the exact name or can contain the wild card character (*) to match any character.
Group Name	External group to which you could map roles.
Description	Optional text about the group.
OK	Displays the Add Role Mapping dialog, showing the external group you selected.
Cancel	Closes the dialog.

Settings Tab

In the Settings tab, you can configure password complexity for internal NetWitness users and set system-wide security parameters. You can also restrict access to incidents in the Respond view.

Password complexity requirements apply only to internal users and are not enforced for external users. External users rely on their own methods and systems to enforce password complexity.

Workflow



What do you want to do?

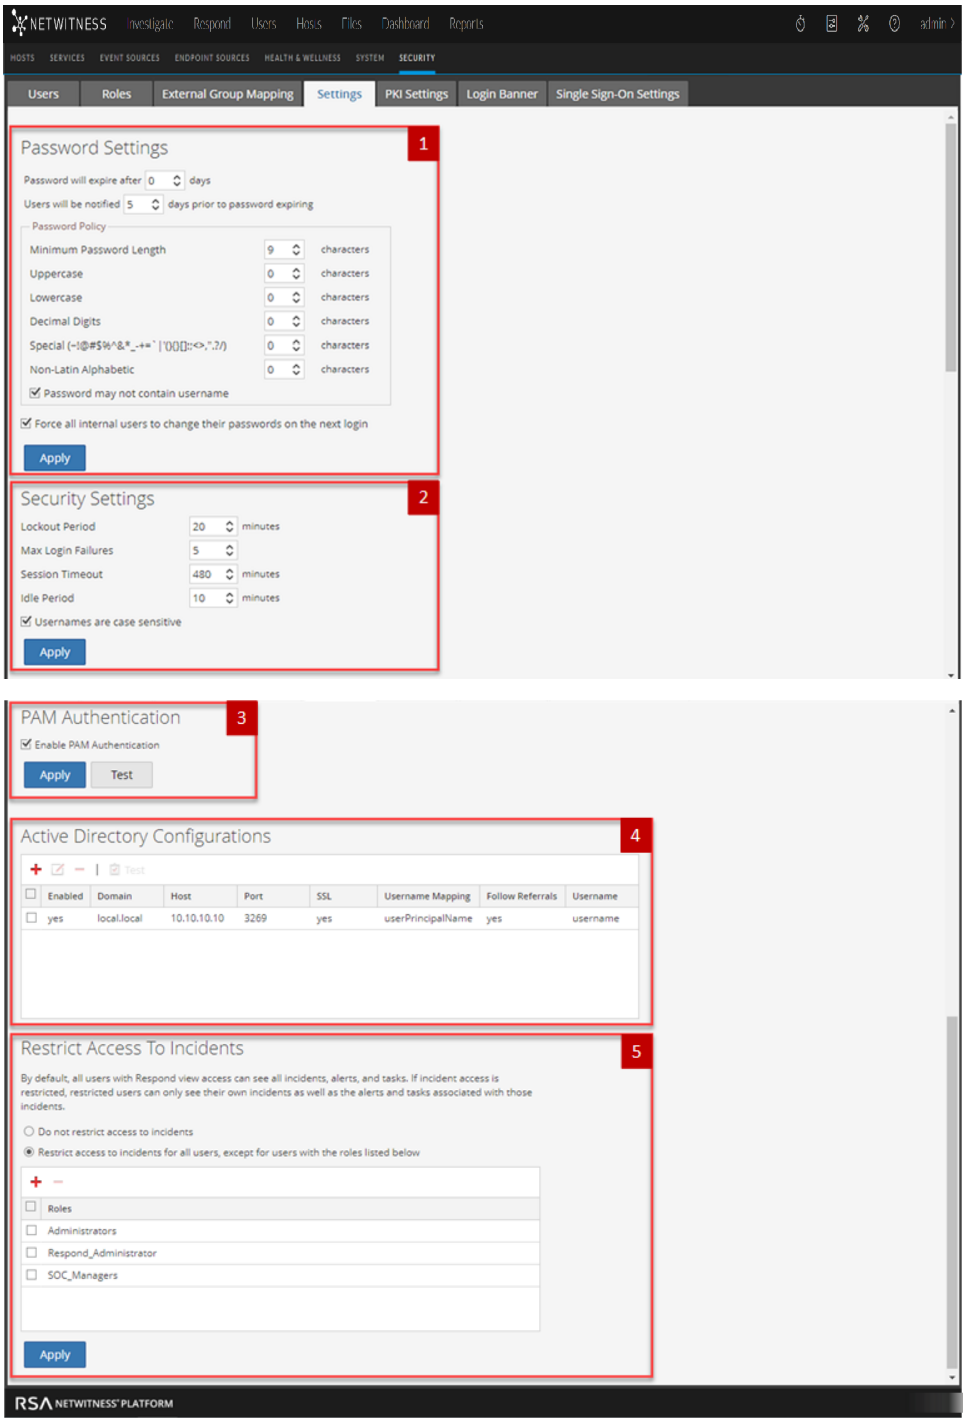
Role	I want to ...	Show me how
Admin	Configure password complexity	Configure Password Complexity
Admin	Configure system-level security settings	Configure System-Level Security Settings
Admin	Restrict access to incidents.	Configure System-Level Security Settings
Admin	(Optional) Configure external authentication	(Optional) Configure External Authentication

Related Topics

- Set Up System Security

Quick Look

To access this view, go to  **(Admin) > Security** and click the **Settings** tab.



The **Settings** tab includes the following sections.

1 **Password Settings**
2 **Security Settings**
3 **PAM Authentication**
4 **Active Directory Configurations**
5 **Restrict Access To Incidents**

Password Settings

The Password Policy section enables you to configure password complexity requirements for internal NetWitness users when they set their passwords.

Option	Description
Password will expire after <n> days	The default number of days before a password expires for all internal NetWitness users. A value of zero (0) disables password expiration. For new installations, the default value is 30. For upgrades, the previous value will migrate automatically to the upgraded installation.
Users will be notified <n> days prior to password expiring	The number of days before the password expiration date, to notify a user that their password is about to expire. Users receive a one-time email on the specified date before their passwords expire. They also see a Password Expiration Message dialog when they log on to NetWitness. The minimum value is 1 day.
Minimum Password Length	Specifies a minimum password length requirement for NetWitness user passwords. A minimum password length prevents users from using short passwords that are easy to guess.
Uppercase	Specifies a minimum number of uppercase characters for the password. This includes European language characters A through Z, with diacritic marks, Greek characters, and Cyrillic characters. For example: • Cyrillic uppercase: Д Ц • Greek uppercase: Π Λ
Lowercase	Specifies a minimum number of lowercase characters for the password. This includes European language characters a through z, sharp-s, with diacritic marks, Greek characters, and Cyrillic characters. For example: • Cyrillic lowercase: д ц • Greek lowercase: π λ
Decimal Digits	Specifies a minimum number of decimal characters (0 through 9) for the password.
Special (~!@#$%^&*_-+=`|'(){}[]:;<>,".?/)	Specifies a minimum number of special characters for the password: `~!@#$%^&*_-+=`|'(){}[]:;<>,".?/`
Non-Latin Alphabetic	Specifies a minimum number of Unicode alphabetic characters that are not uppercase or lowercase. This includes Unicode characters from Asian languages. For example: • Kanji (Japanese): 頁 (leaf) 枒 (tree)
Password May Not Contain Username	Specifies that a password cannot contain the case-insensitive username of the user.

Option	Description
Force all internal users to change their passwords on the next login	Forces all internal users to change their passwords the next time they log on to NetWitness instead of when they create or change their passwords. Note that this setting is checked by default.
Apply	Password strength settings take effect when NetWitness users create or change their passwords. If **Force all internal users to change their passwords on the next login** is selected, all internal users must change their password the next time they log on to NetWitness.

Security Settings

The Security Settings section enables you to configure global security settings for NetWitness users.

Option	Description
Lockout Period	Number of minutes to lock a user out of NetWitness after the configured number of failed logins is exceeded. The default value is 20 minutes.
Max Login Failures	The maximum number of unsuccessful login attempts before a user is locked out. The default value is 5
Session Timeout	The maximum duration of a user session before timing out in minutes. The default value is 600. If the value is 0, there is no maximum time for a session. If the value is a positive integer, the session times out when the configured time has elapsed. The user must log in again.
Idle Period	Number of minutes of inactivity before a session times out. The default value is 10. If the value is 0, the session will not timeout.
Usernames are case sensitive	Select this option if you want the Username field on the NetWitness login screen to be case sensitive. For example, if usernames are case sensitive, you could use admin to log on to NetWitness, but you could not use Admin. This is a mandatory field.
Password	Enter the password if you want to add or edit the Active Directory Security Settings. This is a mandatory field.
Apply	Makes the settings become effective immediately.

PAM Authentication

The PAM Authentication section enables you to configure NetWitness to use Active Directory or PAM to authenticate and test external user logins.

Option	Description
Enable PAM Authentication	Allows NetWitness to use Pluggable Authentication Modules (PAM) to authenticate external user logons.
Apply	Makes the PAM configuration settings become effective in the next logon.
Test	Prompts for a username and password, then tests the currently enabled PAM authentication method.

Active Directory Configurations

The Active Directory Configuration section enables you to configure NetWitness to use Active Directory to authenticate external user logins.

Option	Description
Enabled	Enables Active Directory authentication for NetWitness users.
Domain	Domain name where the Active Directory Service is located.
Host	Host name or IP address where the Active Directory Service is located.
Port	Port on the host that is used for Active Directory Service authentication.
SSL	Indicates whether the Active Directory Service uses Secure Sockets Layer (SSL). To enable SSL so that your Active Directory Service can communicate with NetWitness version 11.1 and later, you must upload an Active Directory server certificate.
Username Mapping	Indicates the Active Directory search field to use for username mapping. You can specify userPrincipalName (UPN) or sAMAccountName.
Follow Referrals	Indicates whether NetWitness will follow LDAP referrals made by Active Directory.
Username	Username of the user that binds to the Active Directory Service while searching Active Directory groups. This is usually a service account that has permissions to query the domain and validate user accounts and group membership. This credential is not used for any other purpose.
Password	Password of the user that binds to the Active Directory Service while searching Active Directory groups. This is usually a service account that has permissions to query the domain and validate user accounts and group membership. This credential is not used for any other purpose.

Restrict Access to Incidents

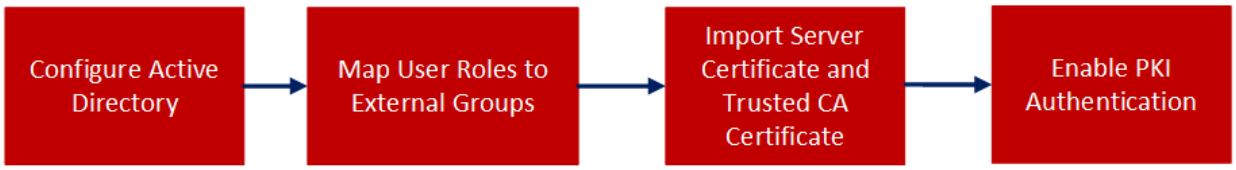
By default, analysts can view all of the incidents, alerts, and tasks in the Respond view. If you have sensitive or restricted information that should not be shared, you can restrict what analysts and other users can see in the Respond view.

Option	Description
Do not restrict access to incidents	Enables analysts to view all of the incidents, alerts, and tasks in the Respond view. This is the default setting.
Restrict access to incidents for all users, except for users with the roles listed below	Restricts what analysts and other users can see in the Respond view. Analysts can only see incidents assigned to them as well as the alerts and tasks associated with those incidents.
(List of user roles whose access to incidents is not restricted.)	When you select to restrict access to incidents, this list shows the user roles that do not have restricted access to incidents. You can adjust this list of user roles independently from your restriction selection. For example, you can add and adjust the list of non-restricted user roles before you make the selection to restrict access to incidents.
Apply	Changes take effect on the next log in to NetWitness Platform.

PKI Settings Tab

The Public Key Infrastructure (PKI) Settings tab enables you to configure PKI authentication for NetWitness.

Workflow

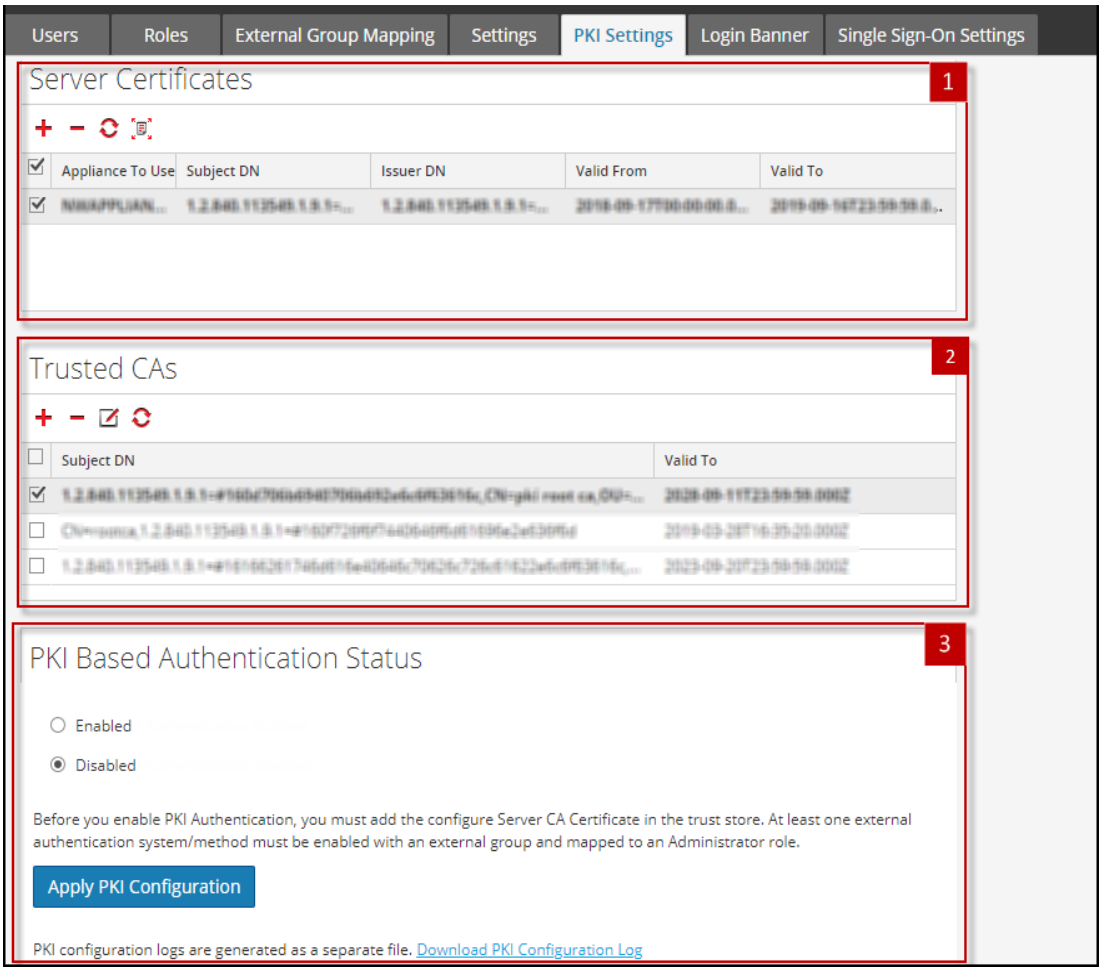


What do you want to do?

Role	I want to ...	Show me how
Admin	Configure External Authentication	Configure Active Directory
Admin	Map user Roles to External Groups	(Optional) Map User Roles to External Groups
Admin	Import Server Certificate	Import Server Certificate and Trusted CA Certificate
Admin	Enable PKI Authentication	Enable PKI Authentication

Quick Look

To access this view, go to  **(Admin) > Security** and click the **PKI Settings** tab.



The **PKI Settings** tab includes the following sections.

1 **Server Certificates**
2 **Trusted CAs**
3 **PKI Based Authentication Status**

Server Certificates

The **Server Certificates** section enables you to import a server certificate with its private key to NetWitness Server. The following table lists the **Server Certificates** features and their description.

Feature	Description
Appliance to Use	The device to use the certificate
Subject DN	The entity to which the certificate is issued.
Issuer DN	The entity which issued the certificate.
Valid From	The start date for the certificate validity.
Valid To	The end date till when a certificate is valid.

Trusted CAs

The **Trusted CAs** section enables you to import a Certificate Authority (CA) certificate to NetWitness. The following table lists the **Trusted CAs** features and their description.

Feature	Description
Subject DN	The entity to which the certificate is issued.
CA	Indicates whether the certificate is Certificate Authority (CA).
Valid To	The end date till when a certificate is valid.

PKI Based Authentication Status

The **PKI Based Authentication Status** section enables you to enable PKI authentication in NetWitness. The following table list the **PKI Based Authentication Status** features and their description.

Feature	Description
Enabled	Select the option to enable PKI.
Disabled	Select the option to disable PKI.
Apply PKI Configuration	Enables PKI authentication for NetWitness users.

Login Banner Tab

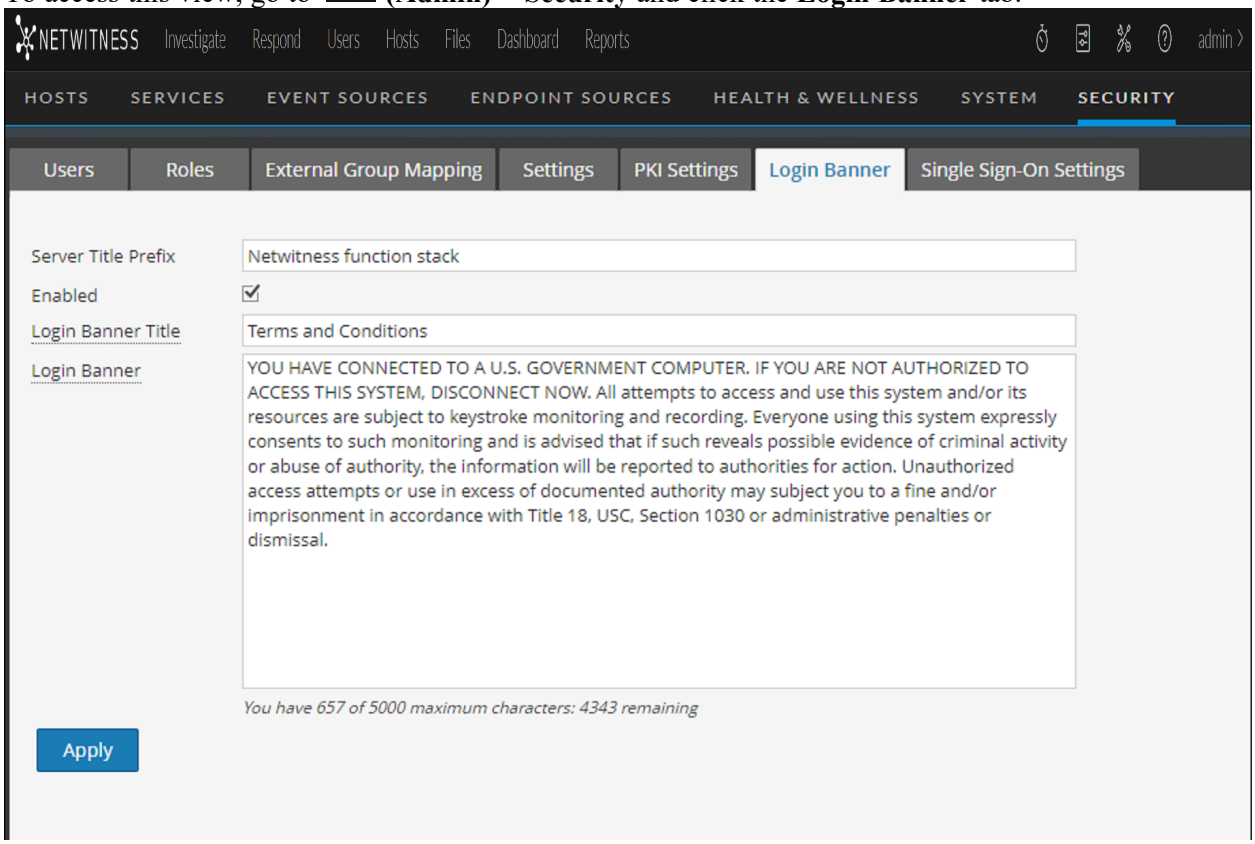
The Login Banner tab provides a way to add a banner to the NetWitness login screen, which will prevent a user from logging on until they agree to the conditions. Add the server title prefix to differentiate the NetWitness Server of the current tab, when you have multiple deployed in your system. You can customize the default title and text of the login banner. The banner is disabled by default.

What do you want to do?

Role	I want to ...	Show me how
Admin	Create or enable a login banner	(Optional) Create a Customized Login Banner

Quick Look

To access this view, go to **(Admin) > Security** and click the **Login Banner** tab.



The following table lists the features of the Login Banner tab.

Feature	Description
Server Title Prefix	Displays the prefix of the NetWitness Server on the title bar.
Enabled	Checkbox that indicates whether or not the login banner is enabled. This box is cleared by default.
Login Banner Title	Shows the title of the dialog box that contains the login conditions.
Login Banner	Shows the conditions the user must acknowledge.

Single Sign-On Settings Tab

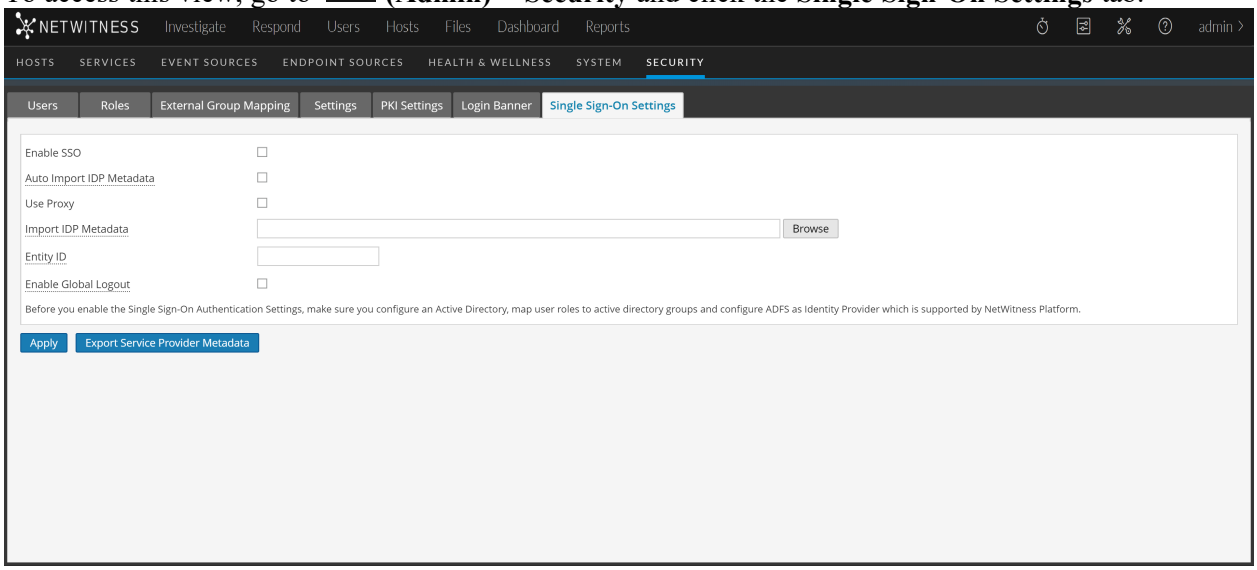
The SSO Settings tab provides a way to enable SSO.

What do you want to do?

Role	I want to ...	Show me how
Admin	Enable SSO Settings	Configure Single Sign-On
Admin	Configure Active Directory	Configure Active Directory
Admin	Configure ADFS as IDP for NetWitness	For instructions on how to configure ADFS as IDP for NetWitness, see the *Configure SAML 2.0 provider settings for portals* topic in Microsoft documentation.
Admin	Map User Roles to External Groups	Manage Users with Roles and Permissions

Quick Look

To access this view, go to **(Admin) > Security** and click the **Single Sign-On Settings** tab.



The following table lists the features of the Single Sign-On tab.

Feature	Description
Enable SSO	Checkbox that indicates whether or not the single sign-on is enabled. This box is cleared by default.

Feature	Description
Auto Import IDP Metadata	If selected the latest IDP metadata is downloaded at regular intervals.
Metadata URL	Enter the metadata URL generated when the connection was established with the ADFS.
Use Proxy	If enabled, the requests to IDP will be routed through the proxy configured in  (Admin) > System > HTTP Proxy settings.
Import IDP Metadata	Enter the metadata URL generated when the connection was established with the ADFS. **Note:** Make sure you update the link every time the IDP metadata is updated.
Entity ID	A unique identifier for NetWitness unique amongst all the applications managed by the same IDP.
Enable Global Logout	Checkbox that enables Global Logout setting for users. When Global Logout is enabled, the user is logged out of NetWitness and also from other applications authenticated by ADFS.
Apply	The admin-server is restarted after which you will get notified in the notification tray when the metadata is ready to be downloaded.
Export Service Provider Metadata	Exports an XML file which is uploaded to IDP to establish the connection between NetWitness (SP) and IDP (ADFS) for authentication.